SnapCenter software documentation

SnapCenter software

NetApp
September 29, 2025

This PDF was generated from https://docs.netapp.com/us-en/snapcenter-61/index.html on September 29, 2025. Always check docs.netapp.com for the latest.

Table of Contents

SnapCenter software documentation . . . 1
Release notes . . . 2
  Release notes . . . 2
  What's new in SnapCenter software 6.1 . . . 2
  Supported upgrade paths for SnapCenter . . . 2
Get started . . . 4
  Learn about SnapCenter software . . . 4
    SnapCenter overview . . . 4
    Security features in SnapCenter . . . 8
    Role-based access control in SnapCenter . . . 9
    Disaster recovery in SnapCenter . . . 14
    Licenses required by SnapCenter . . . 15
    SnapMirror active sync in SnapCenter . . . 18
    Key concepts of data protection . . . 18
    Storage systems and applications supported by SnapCenter . . . 20
    Authentication methods for SnapCenter credentials . . . 21
  Supported SnapCenter operations for ASA r2 systems . . . 22
  Quick start for SnapCenter software . . . 23
Install and configure SnapCenter Server . . . 25
  Prepare for installing the SnapCenter Server . . . 25
    Requirements to install SnapCenter Server . . . 25
    Register to access the SnapCenter software . . . 31
    Multi-factor authentication (MFA) . . . 31
  Install the SnapCenter Server . . . 41
    Install the SnapCenter Server on Windows host . . . 41
    Install the SnapCenter Server on Linux host . . . 45
    Register SnapCenter . . . 49
    Log in to SnapCenter using RBAC authorization . . . 49
  Configure the SnapCenter Server . . . 53
    Add and provision the storage system . . . 53
    Add SnapCenter Standard controller-based licenses . . . 73
    Configure High Availability . . . 78
    Configure role-based access control (RBAC) . . . 82
    Configure audit log settings . . . 110
    Configure secured MySQL connections with SnapCenter Server . . . 111
  Configure Certificate-based authentication . . . 117
    Enable Certificate-based authentication . . . 117
    Export Certificate Authority (CA) certificates from SnapCenter Server . . . 117
    Import CA certificate to the Windows plug-in hosts . . . 118
    Import CA Certificate to the UNIX plug-in hosts . . . 118
    Export SnapCenter certificates . . . 120
  Configure CA Certificate for Windows host . . . 121
    Generate CA Certificate CSR file . . . 121

Import CA certificates . . . 121
Get the CA certificate thumbprint . . . 122
Configure CA certificate with Windows host plug-in services . . . 122
Configure CA certificate with SnapCenter site . . . 123
Enable CA certificates for SnapCenter . . . 124
Configure CA Certificate for Linux host . . . 124
Configure nginx certificate . . . 125
Configure audit log certificate . . . 125
Configure SnapCenter services certificate . . . 125
Configure and enable two-way SSL communication on Windows host . . . 126
Configure two-way SSL communication on Windows host . . . 126
Enable two-way SSL communication on Windows host . . . 128
Configure and enable two-way SSL communication on Linux host . . . 129
Configure two-way SSL communication on Linux host . . . 130
Enable SSL communication on Linux host . . . 131
Configure Active Directory, LDAP, and LDAPS . . . 131
Register untrusted Active Directory domains . . . 132
Configure IIS Application Pools to enable Active Directory read permissions . . . 133
Configure CA client certificate for LDAPS . . . 134
Protect Microsoft SQL Server databases . . . 135
SnapCenter Plug-in for Microsoft SQL Server . . . 135
SnapCenter Plug-in for Microsoft SQL Server overview . . . 135
What you can do with the SnapCenter Plug-in for Microsoft SQL Server . . . 135
SnapCenter Plug-in for Microsoft SQL Server features . . . 136
Support for Asymmetric LUN Mapping in Windows clusters . . . 137
Storage types supported by SnapCenter Plug-in for Microsoft SQL Server . . . 138
Storage layout recommendations for SnapCenter Plug-in for Microsoft SQL Server . . . 140
Minimum ONTAP privileges required for SQL plug-in . . . 142
Prepare storage systems for SnapMirror and SnapVault replication for Plug-in for SQL server . . . 145
Backup strategy for SQL Server resources . . . 145
Restoration strategy for SQL Server . . . 150
Define a cloning strategy for SQL Server . . . 153
Prepare to install the SnapCenter Plug-in for Microsoft SQL Server . . . 153
Installation workflow for SnapCenter Plug-in for Microsoft SQL Server . . . 153
Prerequisites to add hosts and install SnapCenter Plug-in for Microsoft SQL Server . . . 154
Host requirements to install SnapCenter Plug-ins Package for Windows . . . 155
Set up credentials for the SnapCenter Plug-ins Package for Windows . . . 156
Configure credentials for an individual SQL Server resource . . . 157
Configure gMSA on Windows Server 2016 or later . . . 159
Install SnapCenter Plug-in for Microsoft SQL Server . . . 161
Configure CA Certificate . . . 167
Configure Disaster recovery . . . 170
Install SnapCenter Plug-in for VMware vSphere . . . 172
Deploy CA certificate . . . 172
Configure the CRL file . . . 172

Prepare for data protection . . . 172
Prerequisites for using SnapCenter Plug-in for Microsoft SQL Server . . . 172
How resources, resource groups, and policies are used for protecting SQL Server . . . 173
Back up SQL Server database, or instance, or availability group . . . 174
Backup workflow . . . 174
Determine whether resources are available for backup . . . 175
Migrate resources to NetApp storage system . . . 177
Create backup policies for SQL Server databases . . . 178
Create resource groups and attach policies for SQL Server . . . 184
Create resource groups and enable secondary protection for Microsoft SQL Server resources on ASA r2 systems . . . 187
Requirements for backing up SQL resources . . . 189
Create a storage system connection and a credential using PowerShell cmdlets . . . 190
Back up SQL resources . . . 191
Back up SQL Server resource groups . . . 196
Monitor SQL resources backup operations in the SnapCenter Jobs page . . . 196
Cancel the SnapCenter Plug-in for Microsoft SQL Server backup operations . . . 197
View SQL Server backups and clones in the Topology page . . . 198
Clean up the secondary backup count using PowerShell cmdlets . . . 200
Restore SQL Server resources . . . 201
Restore workflow . . . 201
Requirements for restoring a database . . . 201
Restore SQL Server database backups . . . 202
Restore an SQL Server database from secondary storage . . . 209
Restore resources using PowerShell cmdlets . . . 210
Reseed Availability Group databases . . . 212
Monitor SQL resources restore operations . . . 213
Cancel SQL resources restore operations . . . 214
Clone SQL Server database resources . . . 214
Clone workflow . . . 214
Clone from a SQL Server database backup . . . 215
Perform Clone Lifecycle . . . 222
Monitor SQL database clone operations . . . 225
Cancel SQL resource clone operations . . . 225
Split a clone . . . 226
Protect SAP HANA databases . . . 228
SnapCenter Plug-in for SAP HANA Databases . . . 228
SnapCenter Plug-in for SAP HANA Database overview . . . 228
What you can do using the SnapCenter Plug-in for SAP HANA Database . . . 228
SnapCenter Plug-in for SAP HANA Database features . . . 228
Storage types supported by SnapCenter Plug-in for SAP HANA Database . . . 229
Minimum ONTAP privileges required for SAP HANA plug-in . . . 230
Prepare storage systems for SnapMirror and SnapVault replication for SAP HANA databases . . . 233
Backup strategy for SAP HANA databases . . . 233
Restore and recovery strategy for SAP HANA databases . . . 236

Prepare to install the SnapCenter Plug-in for SAP HANA Database . . . 239
Installation workflow of SnapCenter Plug-in for SAP HANA Database . . . 239
Prerequisites for adding hosts and installing SnapCenter Plug-in for SAP HANA Database . . . 239
Host requirements to install SnapCenter Plug-ins Package for Windows . . . 242
Host requirements for installing the SnapCenter Plug-ins Package for Linux . . . 243
Set up credentials for the SnapCenter Plug-in for SAP HANA Database . . . 244
Configure gMSA on Windows Server 2016 or later . . . 247
Install the SnapCenter Plug-in for SAP HANA Databases . . . 248
Configure CA Certificate . . . 254
Install SnapCenter Plug-in for VMware vSphere . . . 261
Deploy CA certificate . . . 261
Configure the CRL file . . . 261
Prepare for data protection . . . 261
Prerequisites for using the SnapCenter Plug-in for SAP HANA Database . . . 262
How resources, resource groups, and policies are used for protecting SAP HANA databases . . . 262
Back up SAP HANA resources . . . 263
Back up SAP HANA resources . . . 263
Configure HDB User Store Key and HDBSQL OS User for the SAP HANA database . . . 264
Discover resources and prepare multitenant database containers for data protection . . . 265
Add resources manually to the plug-in host . . . 267
Create backup policies for SAP HANA databases . . . 269
Create resource groups and attach policies . . . 274
Create resource groups and enable secondary protection for SAP HANA resources on ASA r2 systems . . . 278
Create a storage system connection and a credential using PowerShell cmdlets for SAP HANA database . . . 280
Back up SAP HANA databases . . . 282
Back up resource groups . . . 289
Monitor SAP HANA databases backup operations . . . 290
Cancel backup operations for SAP HANA . . . 291
View SAP HANA database backups and clones in the Topology page . . . 292
Restore SAP HANA Databases . . . 294
Restore workflow . . . 294
Restore and recover a manually added resource backup . . . 294
Restore and recover an auto discovered database backup . . . 299
Restore resources using PowerShell cmdlets . . . 303
Monitor SAP HANA databases restore operations . . . 305
Clone SAP HANA resource backups . . . 306
Clone workflow . . . 306
Clone a SAP HANA database backup . . . 307
Monitor SAP HANA database clone operations . . . 310
Split a clone . . . 311
Delete or split SAP HANA database clones after upgrading SnapCenter . . . 312
Protect Oracle databases . . . 314
Overview of SnapCenter Plug-in for Oracle Database . . . 314

- What can you do with the Plug-in for Oracle Database . . . 314
- Features of Plug-in for Oracle Database . . . 314
- Storage types supported by Plug-in for Oracle Database . . . 316
- Prepare storage systems for SnapMirror and SnapVault replication for Plug-in for Oracle . . . 318
- Minimum ONTAP privileges required for Plug-in for Oracle . . . 318

Install SnapCenter Plug-in for Oracle Database . . . 320
- Installation workflow of SnapCenter Plug-in for Oracle Database . . . 320
- Prerequisites for adding hosts and installing Plug-ins Package for Linux or AIX . . . 321
- Add hosts and install Plug-ins Package for Linux or AIX using GUI . . . 330
- Alternate ways to install Plug-ins Package for Linux or AIX . . . 333
- Configure the SnapCenter Plug-in Loader service . . . 336
- Configure CA certificate with SnapCenter Plug-in Loader (SPL) service on Linux host . . . 340
- Enable CA Certificates for plug-ins . . . 342
- Import data from SnapManager for Oracle and SnapManager for SAP to SnapCenter . . . 343

Install SnapCenter Plug-in for VMware vSphere . . . 348
- Deploy CA certificate . . . 348
- Configure the CRL file . . . 348

Prepare for protecting Oracle databases . . . 348

Back up Oracle databases . . . 350
- Overview of backup procedure . . . 350
- Backup configuration information . . . 351
- Requirements for backing up an Oracle database . . . 363
- Discover Oracle databases available for backup . . . 363
- Create backup policies for Oracle databases . . . 365
- Create resource groups and attach policies for Oracle databases . . . 371
- Create resource groups and enable secondary protection for Oracle resources on ASA r2 systems . . . 373
- Back up Oracle resources . . . 376
- Back up Oracle database resource groups . . . 378
- Monitor Oracle database backup . . . 379
- Other back up operations . . . 380

Mount and unmount database backups . . . 384
- Mount a database backup . . . 384
- Unmount a database backup . . . 386

Restore and recover Oracle databases . . . 386
- Restore workflow . . . 386
- Define a restore and recovery strategy for Oracle databases . . . 386
- Predefined environment variables for restore specific prescript and postscript . . . 391
- Requirements for restoring an Oracle database . . . 392
- Restore and recover Oracle database . . . 393
- Restore and recover tablespaces using point-in-time recovery . . . 398
- Restore and recover pluggable database using point-in-time recovery . . . 399
- Restore and recover Oracle databases using UNIX commands . . . 402
- Monitor Oracle database restore operations . . . 402
- Cancel Oracle database restore operations . . . 403

Clone Oracle database . . . 404

Clone workflow . . . 404
Define a clone strategy for Oracle databases . . . 404
Predefined environment variables for clone specific prescript and postscript . . . 405
Requirements for cloning an Oracle database . . . 407
Clone an Oracle database backup . . . 409
Clone a pluggable database . . . 417
Clone Oracle database backups using UNIX commands . . . 420
Split an Oracle Database Clone . . . 421
Split clone of a pluggable database . . . 422
Monitor Oracle database clone operations . . . 423
Refresh a clone . . . 423
Delete clone of a pluggable database . . . 424
Manage application volumes . . . 425
What are application volumes . . . 425
Add application volumes . . . 425
Backup application volumes . . . 426
Clone application volume backup . . . 429
Protect Windows file systems . . . 431
SnapCenter Plug-in for Microsoft Windows concepts . . . 431
SnapCenter Plug-in for Microsoft Windows overview . . . 431
What you can do with the SnapCenter Plug-in for Microsoft Windows . . . 431
SnapCenter Plug-in for Windows features . . . 431
How SnapCenter backs up Windows file systems . . . 433
Storage types supported by SnapCenter Plug-in for Microsoft Windows . . . 433
Minimum ONTAP privileges required for Windows plug-in . . . 435
Prepare storage systems for SnapMirror and SnapVault replication . . . 437
Define a backup strategy for Windows file systems . . . 438
Sources and destinations of clones for Windows file systems . . . 439
Install SnapCenter Plug-in for Microsoft Windows . . . 440
Installation workflow of SnapCenter Plug-in for Microsoft Windows . . . 440
Installation requirements for SnapCenter Plug-in for Microsoft Windows . . . 440
Add hosts and install SnapCenter Plug-in for Microsoft Windows . . . 444
Install SnapCenter Plug-in for Microsoft Windows on multiple remote hosts using PowerShell cmdlets . . . 448
Install the SnapCenter Plug-in for Microsoft Windows silently from the command line . . . 448
Monitor SnapCenter plug-in package installation status . . . 450
Configure CA certificate . . . 450
Install SnapCenter Plug-in for VMware vSphere . . . 454
Deploy CA certificate . . . 454
Configure the CRL file . . . 454
Back up Windows file systems . . . 454
Back up Windows file systems . . . 454
Determine resource availability for Windows file systems . . . 456
Create backup policies for Windows file systems . . . 456
Create resource groups for Windows file systems . . . 459
Create resource groups and enable secondary protection for Windows file systems on ASA r2

systems . . . 462
Create a storage system connection and a credential using PowerShell cmdlets . . . 464
Back up a single resource on demand for Windows file systems . . . 465
Back up resource groups for Windows file systems . . . 469
Monitor backup operations . . . 470
Cancel backup operations . . . 471
View related backups and clones in the Topology page . . . 472
Clean up the secondary backup count using PowerShell cmdlets . . . 474
Restore Windows file systems . . . 475
Restore Windows file system backups . . . 475
Restore resources using PowerShell cmdlets . . . 479
Monitor restore operations . . . 482
Cancel restore operations . . . 483
Clone Windows file systems . . . 484
Clone from a Windows file system backup . . . 484
Monitor clone operations . . . 490
Cancel clone operations . . . 491
Split a clone . . . 491
Protect Microsoft Exchange Server databases . . . 494
SnapCenter Plug-in for Microsoft Exchange Server concepts . . . 494
SnapCenter Plug-in for Microsoft Exchange Server overview . . . 494
What you can do with SnapCenter Plug-in for Microsoft Exchange Server . . . 494
Storage types supported by SnapCenter Plug-in for Microsoft Windows and for Microsoft Exchange Server . . . 495
Minimum ONTAP privileges required for Exchange plug-in . . . 496
Prepare storage systems for SnapMirror and SnapVault replication . . . 499
Define a backup strategy for Exchange Server resources . . . 499
Define a restore strategy for Exchange databases . . . 502
Install SnapCenter Plug-in for Microsoft Exchange Server . . . 503
Installation workflow of SnapCenter Plug-in for Microsoft Exchange Server . . . 503
Prerequisites to add hosts and install SnapCenter Plug-in for Microsoft Exchange Server . . . 503
Set up credentials for SnapCenter Plug-in for Windows . . . 507
Configure gMSA on Windows Server 2016 or later . . . 508
Add hosts and install Plug-in for Exchange . . . 510
Install Plug-in for Exchange from the SnapCenter Server host using PowerShell cmdlets . . . 515
Install the SnapCenter Plug-in for Exchange silently from the command line . . . 515
Monitor SnapCenter plug-in package installation status . . . 517
Configure CA Certificate . . . 517
Configure SnapManager 7.x for Exchange and SnapCenter to coexist . . . 521
Install SnapCenter Plug-in for VMware vSphere . . . 522
Deploy CA certificate . . . 522
Configure the CRL file . . . 523
Prepare for data protection . . . 523
Prerequisites for using the SnapCenter Plug-in for Microsoft Exchange Server . . . 523
How resources, resource groups, and policies are used for protecting Exchange Server . . . 524

Back up Exchange resources . . . 525
Backup workflow . . . 525
Exchange database and backup verification . . . 526
Determine whether Exchange resources are available for backup . . . 526
Create backup policies for Exchange Server databases . . . 528
Create resource groups and attach policies for Exchange Servers . . . 535
Create a storage system connection and a credential using PowerShell cmdlets for Exchange Server . . . 537
Back up Exchange databases . . . 538
Back up Exchange resources groups . . . 543
Monitor backup operations . . . 544
Cancel backup operations for Exchange database . . . 545
View Exchange backups in the Topology page . . . 546
Restore Exchange resources . . . 548
Restore workflow . . . 548
Requirements for restoring an Exchange database . . . 548
Restore Exchange databases . . . 549
Granular recovery of mails and mailbox . . . 553
Restore an Exchange Server database from secondary storage . . . 553
Reseed a passive Exchange node replica . . . 553
Reseed a replica using PowerShell cmdlets for Exchange database . . . 554
Monitor restore operations . . . 555
Cancel restore operations for Exchange database . . . 556
Protect IBM Db2 . . . 557
SnapCenter Plug-in for IBM Db2 . . . 557
SnapCenter Plug-in for IBM Db2 overview . . . 557
What you can do using the SnapCenter Plug-in for IBM Db2 . . . 557
SnapCenter Plug-in for IBM Db2 features . . . 557
Storage types supported by SnapCenter Plug-in for IBM Db2 . . . 558
Minimum ONTAP privileges required for IBM Db2 plug-in . . . 559
Prepare storage systems for SnapMirror and SnapVault replication for IBM Db2 . . . 562
Backup strategy for IBM Db2 . . . 562
Restore and recovery strategy for IBM Db2 . . . 564
Prepare to install the SnapCenter Plug-in for IBM Db2 . . . 566
Installation workflow of SnapCenter Plug-in for IBM Db2 . . . 566
Prerequisites for adding hosts and installing Plug-ins package for Windows, Linux, or AIX . . . 566
Host requirements to install SnapCenter Plug-ins Package for Windows . . . 571
Host requirements for installing the SnapCenter Plug-ins Package for Linux . . . 572
Set up credentials for the SnapCenter Plug-in for IBM Db2 . . . 573
Configure gMSA on Windows Server 2016 or later . . . 575
Install the SnapCenter Plug-in for IBM Db2 . . . 576
Configure CA Certificate . . . 582
Prepare for data protection . . . 589
Prerequisites for using the SnapCenter Plug-in for IBM Db2 . . . 589
How resources, resource groups, and policies are used for protecting IBM Db2 . . . 590
Back up IBM Db2 resources . . . 590

Back up IBM Db2 resources . . . 590
Discover the databases automatically . . . 592
Add resources manually to the plug-in host . . . 592
Create backup policies for IBM Db2 . . . 593
Create resource groups and attach policies . . . 596
Create resource groups and enable secondary protection for IBM Db2 resources on ASA r2 systems . . . 599
Create a storage system connection and a credential using PowerShell cmdlets for IBM Db2 . . . 602
Back up Db2 databases . . . 603
Back up resource groups . . . 609
Monitor IBM Db2 backup operations . . . 610
Cancel backup operations for IBM Db2 . . . 611
View IBM Db2 backups and clones in the Topology page . . . 612
Restore IBM Db2 . . . 613
Restore workflow . . . 613
Restore a manually added resource backup . . . 614
Restore and recover an auto discovered database backup . . . 618
Monitor IBM Db2 restore operations . . . 619
Clone IBM Db2 resource backups . . . 620
Clone workflow . . . 620
Clone a IBM Db2 backup . . . 621
Monitor IBM Db2 clone operations . . . 628
Split a clone . . . 629
Delete or split IBM Db2 database clones after upgrading SnapCenter . . . 630
Protect PostgreSQL . . . 631
SnapCenter Plug-in for PostgreSQL . . . 631
SnapCenter Plug-in for PostgreSQL overview . . . 631
What you can do using the SnapCenter Plug-in for PostgreSQL . . . 631
SnapCenter Plug-in for PostgreSQL features . . . 631
Storage types supported by SnapCenter Plug-in for PostgreSQL . . . 632
Minimum ONTAP privileges required for PostgreSQL plug-in . . . 633
Prepare storage systems for SnapMirror and SnapVault replication for PostgreSQL . . . 635
Backup strategy for PostgreSQL . . . 636
Restore and recovery strategy for PostgreSQL . . . 638
Prepare to install the SnapCenter Plug-in for PostgreSQL . . . 640
Installation workflow of SnapCenter Plug-in for PostgreSQL . . . 640
Prerequisites to add hosts and install SnapCenter Plug-in for PostgreSQL . . . 640
Host requirements to install SnapCenter Plug-ins Package for Windows . . . 643
Host requirements for installing the SnapCenter Plug-ins Package for Linux . . . 644
Set up credentials for the SnapCenter Plug-in for PostgreSQL . . . 645
Configure gMSA on Windows Server 2016 or later . . . 647
Install the SnapCenter Plug-in for PostgreSQL . . . 648
Configure CA Certificate . . . 654
Prepare for data protection . . . 661
Prerequisites for using the SnapCenter Plug-in for PostgreSQL . . . 661
How resources, resource groups, and policies are used for protecting PostgreSQL . . . 661

Back up PostgreSQL resources . . . 662
Back up PostgreSQL resources . . . 662
Discover the clusters automatically . . . 664
Add resources manually to the plug-in host . . . 664
Create backup policies for PostgreSQL . . . 666
Create resource groups and attach policies . . . 669
Create resource groups and enable secondary protection for PostgreSQL resources on ASA r2 systems . . . 673
Create a storage system connection and a credential using PowerShell cmdlets for PostgreSQL . . . 675
Back up PostgreSQL . . . 677
Back up resource groups . . . 682
Monitor PostgreSQL backup operations . . . 683
Cancel backup operations for PostgreSQL . . . 684
View PostgreSQL backups and clones in the Topology page . . . 685
Restore PostgreSQL . . . 686
Restore workflow . . . 686
Restore and recover a manually added resource backup . . . 687
Restore and recover an auto discovered cluster backup . . . 691
Restore resources using PowerShell cmdlets . . . 693
Monitor PostgreSQL restore operations . . . 696
Clone PostgreSQL resource backups . . . 697
Clone workflow . . . 697
Clone a PostgreSQL backup . . . 698
Monitor PostgreSQL clone operations . . . 701
Split a clone . . . 702
Delete or split PostgreSQL cluster clones after upgrading SnapCenter . . . 703
Protect MySQL . . . 705
SnapCenter Plug-in for MySQL . . . 705
SnapCenter Plug-in for MySQL overview . . . 705
What you can do using the SnapCenter Plug-in for MySQL . . . 705
SnapCenter Plug-in for MySQL features . . . 705
Storage types supported by SnapCenter Plug-in for MySQL . . . 706
Minimum ONTAP privileges required for MySQL plug-in . . . 707
Prepare storage systems for SnapMirror and SnapVault replication for MySQL . . . 709
Backup strategy for MySQL . . . 710
Restore and recovery strategy for MySQL . . . 712
Prepare to install the SnapCenter Plug-in for MySQL . . . 713
Installation workflow of SnapCenter Plug-in for MySQL . . . 713
Prerequisites to add hosts and install SnapCenter Plug-in for MySQL . . . 714
Host requirements to install SnapCenter Plug-ins Package for Windows . . . 717
Host requirements for installing the SnapCenter Plug-ins Package for Linux . . . 718
Set up credentials for the SnapCenter Plug-in for MySQL . . . 719
Install the SnapCenter Plug-in for MySQL . . . 721
Configure CA Certificate . . . 726
Prepare for data protection . . . 733

Prerequisites for using the SnapCenter Plug-in for MySQL 733
How resources, resource groups, and policies are used for protecting MySQL 734
Back up MySQL resources 734
Back up MySQL resources 734
Discover the databases automatically 736
Add resources manually to the plug-in host 736
Create backup policies for MySQL 737
Create resource groups and attach policies 741
Create resource groups and enable secondary protection for MySQL resources on ASA r2 systems 744
Create a storage system connection and a credential using PowerShell cmdlets for MySQL 747
Back up MySQL 748
Back up resource groups 753
Monitor MySQL backup operations 754
Cancel backup operations for MySQL 755
View MySQL backups and clones in the Topology page 756
Restore MySQL 757
Restore workflow 757
Restore and recover a manually added resource backup 758
Restore and recover an auto discovered database backup 762
Restore resources using PowerShell cmdlets 764
Monitor MySQL restore operations 766
Clone MySQL resource backups 767
Clone workflow 767
Clone a MySQL backup 768
Monitor MySQL clone operations 771
Split a clone 772
Delete or split MySQL database clones after upgrading SnapCenter 773
Protect applications using NetApp supported plug-ins 774
NetApp supported plug-ins 774
NetApp supported plug-ins overview 774
What you can do with the NetApp supported plug-ins 774
NetApp supported plug-ins features 775
Storage types supported by NetApp supported plug-ins 776
Minimum ONTAP privileges required for NetApp supported plug-in 776
Prepare storage systems for SnapMirror and SnapVault replication for NetApp supported plug-ins 778
Define a backup strategy 779
Backup strategy for NetApp supported plug-ins 780
Types of restore strategies supported for manually added NetApp supported plug-in resources 780
Prepare to install NetApp supported plug-ins 781
Installation workflow of SnapCenter NetApp supported plug-ins 781
Prerequisites for adding hosts and installing Plug-ins package for Windows, Linux, or AIX 781
Host requirements to install SnapCenter Plug-ins Package for Windows 786
Host requirements for installing the SnapCenter Plug-ins Package for Linux and AIX 787
Set up credentials for NetApp supported plug-ins 788
Configure gMSA on Windows Server 2016 or later 789

Install the NetApp supported plug-ins . . . 791
Configure CA Certificate . . . 797
Prepare for data protection . . . 804
Prerequisites for using the NetApp supported plug-ins . . . 804
How resources, resource groups, and policies are used for protecting NetApp supported plug-in resources . . . 805
Back up NetApp supported plug-ins resources . . . 806
Back up NetApp supported plug-ins resources . . . 806
Add resources to NetApp supported plug-ins . . . 806
Create policies for NetApp supported plug-in resources . . . 811
Create resource groups and attach policies . . . 814
Create resource groups and enable secondary protection for resources on ASA r2 systems . . . 817
Create a storage system connection and a credential using PowerShell cmdlets . . . 820
Back up individual NetApp supported plug-ins resources . . . 821
Back up resource groups of NetApp supported plug-in resources . . . 826
Monitor NetApp supported plug-in resources backup operations . . . 827
Cancel backup operations for NetApp supported plug-ins . . . 828
View NetApp supported plug-ins resource related backups and clones in the Topology page . . . 829
Restore NetApp supported plug-ins resources . . . 830
Restore NetApp supported plug-in resources . . . 830
Restore a resource backup . . . 831
Monitor NetApp supported plug-in resources restore operations . . . 835
Clone NetApp supported plug-ins resource backups . . . 836
Clone NetApp supported plug-ins resource backups . . . 836
Clone from a backup . . . 837
Monitor NetApp supported plug-in resource clone operations . . . 843
Protect Unix file systems . . . 844
What you can do with the SnapCenter Plug-in for Unix file systems . . . 844
Supported configurations . . . 844
Limitations . . . 845
Features . . . 845
Install SnapCenter Plug-in for Unix file systems . . . 845
Prerequisites for adding hosts and installing Plug-ins Package for Linux . . . 845
Add hosts and install Plug-ins Package for Linux using GUI . . . 846
Configure the SnapCenter Plug-in Loader service . . . 849
Configure CA certificate with SnapCenter Plug-in Loader (SPL) service on Linux host . . . 852
Enable CA Certificates for plug-ins . . . 855
Install SnapCenter Plug-in for VMware vSphere . . . 855
Deploy CA certificate . . . 855
Configure the CRL file . . . 856
Prepare for protecting Unix file systems . . . 856
Back up Unix file systems . . . 856
Discover the UNIX file systems available for backup . . . 856
Create backup policies for Unix file systems . . . 857
Create resource groups and attach policies for Unix file systems . . . 859

Create resource groups and enable secondary protection for Unix file systems on ASA r2 systems . . . 861
Back up Unix file systems . . . 864
Back up Unix file systems resource groups . . . 865
Monitor Unix file systems backup . . . 865
View protected Unix file systems in the Topology page . . . 867
Restore and recover Unix file systems . . . 869
Restore Unix file systems . . . 869
Monitor Unix file systems restore operations . . . 870
Clone Unix file systems . . . 871
Clone Unix file system backup . . . 871
Split a clone . . . 872
Monitor Unix file systems clone operations . . . 874
Protect applications running on Azure NetApp Files . . . 875
Protect applications running on Azure NetApp Files . . . 875
Limitations . . . 875
Install SnapCenter and create credentials . . . 875
Install SnapCenter on Azure Virtual Machine . . . 875
Create the Azure credential in SnapCenter . . . 877
Configure the Azure storage account . . . 877
Create the credential to add the plug-in host . . . 878
Protect SAP HANA databases . . . 879
Add hosts and install SnapCenter plug-in for SAP HANA database . . . 879
Add SAP HANA database . . . 879
Create backup policies for SAP HANA databases . . . 880
Create resource groups and attach SAP HANA backup policies . . . 880
Back up SAP HANA databases running on Azure NetApp Files . . . 881
Back up SAP HANA resource groups . . . 882
Restore and recover SAP HANA databases . . . 883
Clone SAP HANA database backup . . . 884
Protect Microsoft SQL Server databases . . . 885
Add hosts and install SnapCenter plug-in for SQL Server database . . . 885
Create backup policies for SQL Server databases . . . 885
Create resource groups and attach SQL backup policies . . . 887
Back up SQL Server databases running on Azure NetApp Files . . . 888
Back up SQL Server resource groups . . . 889
Restore and recover SQL Server databases . . . 889
Clone SQL Server database backup . . . 890
Protect Oracle databases . . . 892
Add hosts and install SnapCenter plug-in for Oracle database . . . 892
Create backup policies for Oracle databases . . . 892
Create resource groups and attach Oracle backup policies . . . 893
Back up Oracle databases running on Azure NetApp Files . . . 894
Back up Oracle resource groups . . . 895
Restore and recover Oracle databases . . . 895
Clone Oracle database backup . . . 898

Manage SnapCenter Server and plug-ins . . . 901
View dashboard . . . 901
Overview of dashboard . . . 901
How to view information on the dashboard . . . 904
Request status reports of the jobs from the dashboard . . . 905
Request reports of the protection status from the dashboard . . . 905
Manage RBAC . . . 906
Modify a role . . . 906
Modify users and groups . . . 907
Manage hosts . . . 907
Refresh virtual machine information . . . 908
Modify plug-in hosts . . . 909
Start or restart plug-in services . . . 910
Suspend schedules for host maintenance . . . 910
Operations supported from the Resources page . . . 911
Manage policies . . . 912
Modify policies . . . 912
Detach policies . . . 912
Delete policies . . . 913
Manage resource groups . . . 913
Stop and resume operations on resource groups . . . 914
Delete resource groups . . . 914
Manage backups . . . 914
Rename backups . . . 914
Delete backups . . . 915
Remove protection . . . 916
Delete clones . . . 916
Monitor jobs, schedules, events, and logs . . . 917
Monitor jobs . . . 917
Monitor schedules . . . 918
Monitor events . . . 918
Monitor logs . . . 919
Remove jobs and logs from SnapCenter . . . 920
Overview of SnapCenter reporting capabilities . . . 920
Access reports . . . 921
Filter your report . . . 922
Export or print reports . . . 922
Set the SMTP server for email notifications . . . 922
Configure the option to email reports . . . 923
Manage the SnapCenter Server repository . . . 923
Prerequisites for protecting the SnapCenter repository . . . 923
Back up the SnapCenter repository . . . 924
View backups of the SnapCenter repository . . . 924
Restore the SnapCenter database repository . . . 924
Migrate the SnapCenter repository . . . 925

Reset the SnapCenter repository password . . . 926
Manage resources of untrusted domains . . . 926
Modify untrusted domains . . . 926
Unregister untrusted Active Directory domains . . . 927
Manage the storage system . . . 928
Modify storage system configuration . . . 928
Delete the storage system . . . 930
REST API support . . . 931
Manage EMS data collection . . . 931
Stop EMS data collection . . . 931
Start EMS data collection . . . 932
Change EMS data collection schedule and target SVM . . . 932
Monitor EMS data collection status . . . 932
Upgrade SnapCenter Server and plug-ins . . . 933
Configure SnapCenter to check for available updates . . . 933
Upgrade workflow . . . 933
Supported upgrade paths . . . 933
Upgrade the SnapCenter Server on Windows host . . . 934
Upgrade the SnapCenter Server on Linux host . . . 936
Upgrade your plug-in packages . . . 937
Tech refresh . . . 939
Tech refresh of SnapCenter Server host . . . 939
Tech refresh of a node in F5 cluster . . . 940
Decommissioning the old SnapCenter Server host . . . 940
Rollback to the old SnapCenter Server host . . . 940
Disaster recovery . . . 940
Tech refresh of SnapCenter plug-in hosts . . . 942
Tech refresh of storage system . . . 944
Update the backups of the primary storage . . . 945
Update the backups of the secondary storage . . . 946
Uninstall SnapCenter Server and plug-ins . . . 948
Uninstall SnapCenter plug-in packages . . . 948
Prerequisites for removing a host . . . 948
Remove a host . . . 949
Uninstall plug-ins using the SnapCenter GUI . . . 949
Uninstall Windows plug-ins using the PowerShell cmdlet . . . 950
Uninstall plug-ins locally on a host . . . 951
Uninstall plug-ins package for Linux or AIX using CLI . . . 951
Uninstall the SnapCenter Server on Windows host . . . 952
Uninstall the SnapCenter Server on Linux host . . . 952
Automate using REST APIs . . . 954
SnapCenter automation using REST APIs . . . 954
How to access SnapCenter REST API natively . . . 954
REST web services foundation . . . 954
Resources and state representation . . . 954

URI endpoints . . . 955
HTTP messages . . . 955
JSON formatting . . . 955
Basic operational characteristics . . . 955
Request and response API transaction . . . 955
Support for CRUD operations . . . 955
Object identifiers . . . 956
Object instances and collections . . . 956
Synchronous and asynchronous operations . . . 956
Security . . . 956
Input variables controlling an API request . . . 957
HTTP methods . . . 957
Request headers . . . 957
Request body . . . 957
Filtering objects . . . 957
Requesting specific object fields . . . 958
Sorting objects in the output set . . . 959
Pagination when retrieving objects in a collection . . . 959
Size properties . . . 960
Interpretation of an API response . . . 960
HTTP status code . . . 960
Response headers . . . 961
Response body . . . 961
Errors . . . 962
REST APIs supported for SnapCenter Server and plug-ins . . . 963
Auth . . . 963
Domains . . . 963
Jobs . . . 963
Settings . . . 963
Hosts . . . 963
Resources . . . 964
Backups . . . 965
Clones . . . 966
Clone split . . . 966
Resource Groups . . . 966
Policies . . . 967
Storage . . . 967
Share . . . 967
Plugins . . . 967
Reports . . . 968
Alerts . . . 969
Rbac . . . 969
Configuration . . . 969
CertificateSettings . . . 969
Repository . . . 969

Version . . . 970
How to access REST APIs using the Swagger API web page . . . 970
Get started with the REST API . . . 970
Hello World . . . 971
Legal notices . . . 972
Copyright . . . 972
Trademarks . . . 972
Patents . . . 972
Privacy policy . . . 972
Open source . . . 972

SnapCenter software documentation

Release notes

Release notes

Learn about the new and enhanced features available in SnapCenter 6.1.

For a complete list of new features and enhancements, see What's new in SnapCenter 6.1.

For details about known issues, limitations, fixed issues, and new and changed commands, see the SnapCenter Software 6.1 Release Notes. You must sign in with your NetApp account or create an account to access the Release Notes.

What's new in SnapCenter software 6.1

The 6.1 release of SnapCenter software includes several new features and enhancements.

Update	Description
Support for ASA r2 systems	SnapCenter now supports ASA r2 systems that create a simplified experience specific to the needs of SAN-only customers.
Supports secondary tamperproof snapshot	SnapCenter allows you to specify the locking period for secondary snapshots.

Supported upgrade paths for SnapCenter

The upgrade path helps you to understand from which all earlier versions of SnapCenter you can upgrade to the latest versions of SnapCenter and which versions of the plug-ins are supported.

If you are on SnapCenter Server version…	You can directly upgrade SnapCenter Server to…	Supported plug-in versions
5.0	6.0	• 5.0 • 6.0
	6.0.1	• 6.0.1
	6.1	• 6.1
6.0	6.0.1	• 6.0 • 6.0.1
	6.1	• 6.1

If you are on SnapCenter Server version…	You can directly upgrade SnapCenter Server to…	Supported plug-in versions
6.0.1	6.1	• 6.0.1 • 6.1

For information about upgrading the SnapCenter Plug-in for VMware vSphere, see Upgrade SnapCenter Plug-in for VMware vSphere.

Get started

Learn about SnapCenter software

SnapCenter overview

SnapCenter software is a simple, centralized, and scalable platform for application-consistent data protection. It protects applications, databases, host file systems, and VMs on ONTAP systems in the Hybrid Cloud.

SnapCenter uses NetApp Snapshot, SnapRestore, FlexClone, SnapMirror, and SnapVault technologies to provide:

- Fast, space-efficient, application-consistent, disk-based backups
- Fast, detailed restore, and application-consistent recovery
- Quick, space-efficient cloning

SnapCenter includes SnapCenter Server and lightweight plug-ins. You can automate plug-in deployment to remote application hosts, schedule backup, verification, and clone operations, and monitor data protection operations.

You can install SnapCenter either on on-premises or on a public cloud to protect data.

- On-premises to protect the following:
  - Data that is on ONTAP FAS, AFF, or ASA primary systems and replicated to ONTAP FAS, AFF, or ASA secondary systems
  - Data that is on ONTAP Select primary systems
  - Data that is on ONTAP FAS, AFF, or ASA primary and secondary systems and protected to local StorageGRID object storage
  - Data that is on ONTAP ASA r2 primary and secondary systems
- On-premises in a Hybrid Cloud to protect the following:
  - Data that is on ONTAP FAS, AFF, or ASA primary systems and replicated to Cloud Volumes ONTAP
  - Data that is on ONTAP FAS, AFF, or ASA primary and secondary systems and protected to object and archive storage in cloud (using BlueXP backup and recovery integration)
- In a public cloud to protect the following:
  - Data that is on Cloud Volumes ONTAP (formerly ONTAP Cloud) primary systems
  - Data that is on Amazon FSX for ONTAP
  - Data that is on primary Azure NetApp Files (Oracle, Microsoft SQL, and SAP HANA)

Key features

SnapCenter provides the following key features:

- Centralized, application-consistent data protection of different applications

  Data protection is supported for Microsoft Exchange Server, Microsoft SQL Server, Oracle Databases on

Linux or AIX, SAP HANA database, IBM Db2, PostgreSQL, MySQL, and Windows Host Filesystems running on ONTAP systems. SnapCenter also supports protection of applications such as MongoDB, Storage, MaxDB, Sybase ASE, ORASCPM.

- Policy-based backups

  Policy-based backups leverage NetApp Snapshot technology to create fast, space-efficient, application-consistent, disk-based backups. You can also set up automatic protection of these backups to secondary storage by updating existing protection relationships.

- Backups for multiple resources

  You can back up multiple resources (applications, databases, or host file systems) of the same type at once using SnapCenter resource groups.

- Restore and recovery

  SnapCenter provides rapid, granular restores of backups and application-consistent, time-based recovery. You can restore from any destination in the Hybrid Cloud.

- Cloning

  SnapCenter provides quick, space-efficient, and application-consistent cloning. You can clone on any destination in the Hybrid Cloud.

- Single user management graphical user interface

  SnapCenter provides a single interface to manage backups and clones in any Hybrid Cloud destination.

- REST APIs, Windows cmdlets, UNIX commands

  SnapCenter provides REST APIs for most functionality for integration with any orchestration software, and use of Windows PowerShell cmdlets and command-line interface.

- Centralized data protection dashboard and reporting
- Role-Based Access Control (RBAC) for security and delegation
- A built-in repository database with high-availability to store all backup metadata
- Automated push install of plug-ins
- High Availability
- Disaster Recovery (DR)
- SnapLock Learn More
- SnapMirror active sync (initially released as SnapMirror Business Continuity [SM-BC])
- Synchronous mirroring Learn More

SnapCenter architecture and components

SnapCenter uses a layered design with a central management server and plug-in hosts. The server and plug-in hosts can be in different locations.

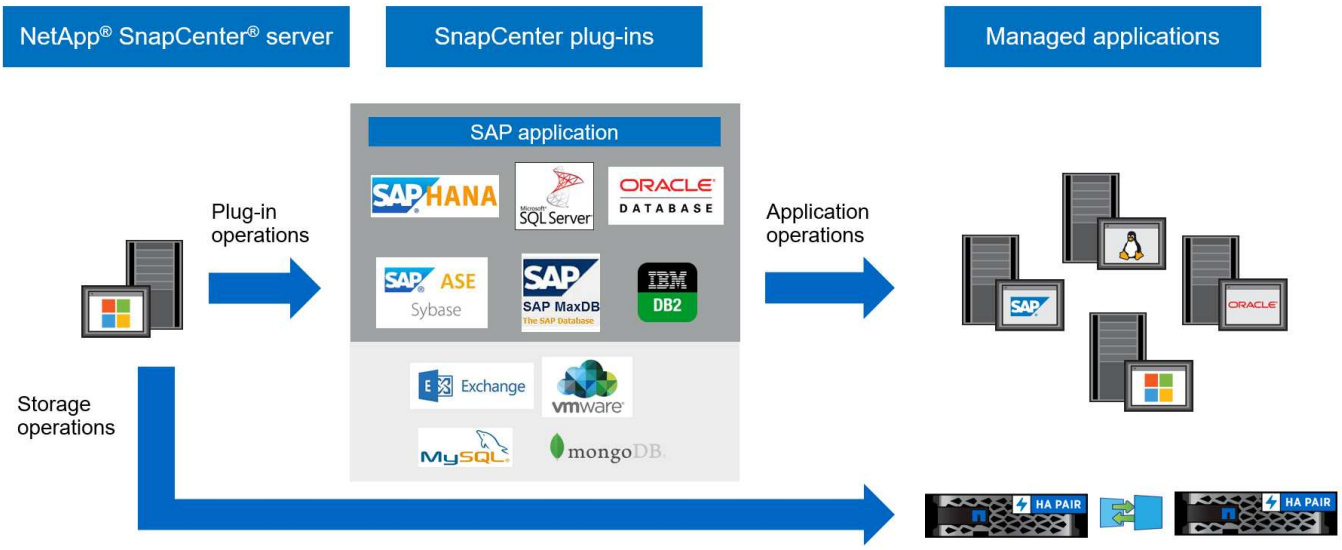

SnapCenter includes the SnapCenter Server, the SnapCenter Plug-in package for Windows, and the SnapCenter Plug-In Package for Linux. Each package contains plug-ins for various applications and infrastructure components.

SnapCenter Server

The SnapCenter Server supports Microsoft Windows and Linux (RHEL 8.x, RHEL 9.x, SLES 15 SP5) operating systems. SnapCenter server includes a web server, a centralized HTML5-based user interface, PowerShell cmdlets, REST APIs, and the SnapCenter repository.

SnapCenter stores information about its operations in the SnapCenter repository.

SnapCenter plug-ins

Each SnapCenter plug-in supports specific environments, databases, and applications.

Plug-in name	Included in install package	Requires other plug-ins	Installed on host	Platform supported
SnapCenter plug-in for Microsoft SQL Server	Plug-ins package for Windows	Plug-in for Windows	SQL Server host	Windows
SnapCenter plug-in for Windows	Plug-ins package for Windows		Windows host	Windows
SnapCenter plug-in for Microsoft Exchange Server	Plug-ins package for Windows	Plug-in for Windows	Exchange Server host	Windows
SnapCentre plug-in for Oracle Database	Plug-ins package for Linux and plug-ins Package for AIX	Plug-in for UNIX	Oracle host	Linux or AIX

Plug-in name	Included in install package	Requires other plug-ins	Installed on host	Platform supported
SnapCenter plug-in for SAP HANA Database	Plug-ins package for Linux and plug-ins package for Windows	Plug-in for UNIX or plug-in for Windows	HDBSQL client host	Linux or Windows
SnapCenter plug-in for IBM Db2	Plug-ins package for Linux and plug-ins Package for Windows	Plug-in for UNIX or plug-in for Windows	Db2 host	Linux, AIX, or Windows
SnapCenter plug-in for PostgreSQL	Plug-ins package for Linux and plug-ins package for Windows	Plug-in for UNIX or plug-in for Windows	PostgreSQL host	Linux or Windows
SnaoCenter plug-in for MySQL	Plug-ins package for Linux and plug-ins package for Windows	Plug-in for UNIX or Plug-in for Windows	MySQL host	Linux or Windows
SnapCenter plug-in for MongoDB	Plug-ins package for Linux and plug-ins package for Windows	Plug-in for UNIX or plug-in for Windows	MongoDB host	Linux or Windows
SnapCenter plug-in for ORASCPM (Oracle Applications)	Plug-ins package for Linux and plug-ins package for Windows	Plug-in for UNIX or plug-in for Windows	Oracle host	Linux or Windows
SnapCenter plug-in for SAP ASE	Plug-ins package for Linux and plug-ins package for Windows	Plug-in for UNIX or plug-in for Windows	SAP host	Linux or Windows
SnapCenter plug-in for SAP MaxDB	Plug-ins package for Linux and plug-ins package for Windows	Plug-in for UNIX or plug-in for Windows	SAP MaxDB host	Linux or Windows
SnapCenter plug-in for Storage plug-in	Plug-ins package for Linux and Plug-ins package for Windows	Plug-in for UNIX or plug-in for Windows	Storage host	Linux or Windows

The SnapCenter Plug-in for VMware vSphere supports crash-consistent and VM-consistent backup and restore operations for virtual machines (VMs), datastores, and Virtual Machine Disks (VMDKs). It also supports

application-consistent backup and restore operations for virtualized databases and file systems.

To protect databases, filesystems, VMs, or datastores on VMs, deploy the SnapCenter Plug-in for VMware vSphere appliance. For information, refer SnapCenter Plug-in for VMware vSphere documentation.

SnapCenter repository

The SnapCenter repository, sometimes referred to as the NSM database, stores information and metadata for every SnapCenter operation.

The SnapCenter Server installation installs the MySQL Server repository database by default. If you have already installed MySQL Server and want to perform a fresh installation of SnapCenter Server, you must uninstall MySQL Server.

SnapCenter supports MySQL Server 8.0.37 or later as the SnapCenter repository database. If you use an earlier version of MySQL Server with an earlier release of SnapCenter, the SnapCenter upgrade process upgrades MySQL Server to version 8.0.37 or later.

The SnapCenter repository stores the following information and metadata:

- Backup, clone, restore, and verification metadata
- Reporting, job, and event information
- Host and plug-in information
- Role, user, and permission details
- Storage system connection information

Security features in SnapCenter

SnapCenter employs strict security and authentication features to enable you to keep your data secure.

SnapCenter includes the following security features:

- All communication to SnapCenter uses HTTP over SSL (HTTPS).
- All credentials in SnapCenter are protected using Advanced Encryption Standard (AES) encryption.
- Supports security algorithms that are compliant with the Federal Information Processing Standard (FIPS).
- Supports using the authorized CA certificates provided by the customer.
- Supports Transport Layer Security (TLS) 1.3 for communication with ONTAP. You can also use TLS 1.2 for communication between clients and servers.
- Supports a certain set of SSL Cipher suites to provide security across network communication. Learn more.
- SnapCenter is installed inside your company's firewall to enable access to the SnapCenter Server and to enable communication between the SnapCenter Server and the plug-ins.
- SnapCenter API and operation access uses tokens encrypted with AES encryption, which expire after 24 hours.
- SnapCenter integrates with Windows Active Directory for login and role-based access control (RBAC) that govern access permissions.
- IPsec is supported with SnapCenter on ONTAP for Windows and Linux host machines. Learn more.

- SnapCenter PowerShell cmdlets are session secured.
- After a default period of 15 minutes of inactivity, SnapCenter warns you that you will be logged out in 5 minutes.

  After 20 minutes of inactivity, SnapCenter logs you out, and you must log in again. You can modify the log out period.

- Login is temporarily disabled after 5 incorrect login attempts.
- Supports CA certificate authentication between SnapCenter Server and ONTAP. Learn more.
- Integrity Verifier is added to the SnapCenter Server and the plug-ins and it validates all the shipped binaries during fresh installation and upgrade operations.

CA Certificate Overview

The SnapCenter Server installer enables the Centralized SSL Certificate Support during installation. To enhance the secured communication between the server and the plug-in, SnapCenter supports using the authorized CA certificates provided by the customer.

You should deploy CA certificates after installing the SnapCenter Server and the respective plug-ins. For more information, see Generate CA Certificate CSR file.

You can also deploy CA certificate for SnapCenter plug-in for VMware vSphere. For more information, see Create and import certificates.

Two-way SSL communication

Two-way SSL communication secures the mutual communication between SnapCenter Server and the plug-ins.

Certificate based authentication Overview

Certificate based authentication verifies the authenticity of respective users who try to access the SnapCenter plug-in host. User should export the SnapCenter Server certificate without private key and import it in the plug-in host trusted store. Certificate based authentication works only if the two-way SSL feature is enabled.

Multi-factor authentication (MFA)

MFA uses a third-party Identity Provider (IdP) via the Security Assertion Markup Language (SAML) to manage user sessions. This functionality enhances the authentication security by having an option to use multiple factors such as TOTP, biometrics, push notifications etc. along with the existing username & password. Also, it enables the customer to use their own user identity providers to get unified user login (SSO) across their portfolio.

MFA is applicable only for SnapCenter Server UI login. The logins are authenticated through the IdP Active Directory Federation Services (AD FS). You can configure various authentication factors at AD FS. SnapCenter is the service provider and you should configure SnapCenter as a relying party in AD FS. To enable MFA in SnapCenter, you will require the AD FS metadata.

For information to enable MFA, see Enable Multi-factor authentication.

Role-based access control in SnapCenter

SnapCenter role-based access control (RBAC) and ONTAP permissions enable

SnapCenter administrators to delegate control of SnapCenter resources to different users or groups of users. This centrally managed access empowers application administrators to work securely within delegated environments.

You can create and modify roles, and add resource access to users any time. However, when you are setting up SnapCenter for the first time, you should at least add Active Directory users or group to roles, and then add resource access to those users or groups.

> You cannot use SnapCenter to create user or group accounts. You should create user or group accounts in Active Directory of the operating system or database.

Types of RBAC in SnapCenter

SnapCenter uses the following types of role-based access control:

- SnapCenter RBAC
- Application-level RBAC
- SnapCenter plug-in for VMware vSphere RBAC
- ONTAP permissions

SnapCenter RBAC

SnapCenter has predefined roles and you can assign users or groups of users to these roles. The predefined roles are:

- SnapCenter Admin role
- App Backup and Clone Admin role
- Backup and Clone Viewer role
- Infrastructure Admin role

When you assign a role to a user, only jobs that are relevant to that user are visible in the Jobs page unless you assigned the SnapCenterAdmin role.

You can also create new roles and manage permissions and users. You can assign permissions to users or groups to access SnapCenter objects such as hosts, storage connections, and resource groups.

You can assign RBAC permissions to users and groups within the same forest and to users belonging to different forests. You cannot assign RBAC permissions to users belonging to nested groups across forests.

> If you create a custom role, it must contain all of the permissions of the SnapCenterAdmin role. If you only copy some of the permissions, for example, Host add or Host remove, you cannot perform those operations.

Users are required to provide authentication during login, through the graphical user interface (GUI) or using PowerShell cmdlets. If users are members of more than one role, after entering login credentials, they are prompted to specify the role they want to use. Users are also required to provide authentication to run the APIs.

Application-level RBAC

SnapCenter uses credentials to verify that authorized SnapCenter users also have application-level permissions.

For example, if you want to perform data protection operations in a SQL Server environment, you must set credentials with the proper Windows or SQL credentials. The SnapCenter Server authenticates the credentials set using either method. If you want to perform data protection operations in a Windows file system environment on ONTAP storage, the SnapCenter admin role must have admin privileges on the Windows host.

Similarly, if you want to perform data protection operations on an Oracle database and if the operating system (OS) authentication is disabled in the database host, you must set credentials with the Oracle database or Oracle ASM credentials. The SnapCenter Server authenticates the credentials set using one of these methods depending on the operation.

SnapCenter Plug-in for VMware vSphere RBAC

If you are using the SnapCenter VMware plug-in for VM-consistent data protection, the vCenter Server provides an additional level of RBAC. The SnapCenter VMware plug-in supports both vCenter Server RBAC and ONTAP RBAC. Learn More

Best Practice: NetApp recommends that you create one ONTAP role for SnapCenter Plug-in for VMware vSphere operations and assign it all the required privileges.

ONTAP permissions

You should create vsadmin account with required permissions to access the storage system. Learn More

Permissions assigned to the pre-defined SnapCenter roles

When you add a user to a role, you must assign either the StorageConnection permission to enable storage virtual machine (SVM) communication, or assign an SVM to the user to enable permission to use the SVM. The Storage Connection permission enables users to create SVM connections.

For example, a user with the SnapCenter Admin role can create SVM connections and assign them to a user with the App Backup and Clone Admin role, which by default does not have permission to create or edit SVM connections. Without an SVM connection, users cannot complete any backup, clone, or restore operations.

SnapCenter Admin role

The SnapCenter Admin role has all permissions enabled. You cannot modify the permissions for this role. You can add users and groups to the role or remove them.

App Backup and Clone Admin role

The App Backup and Clone Admin role has the permissions required to perform administrative actions for application backups and clone-related tasks. This role does not have permissions for host management, provisioning, storage connection management, or remote installation.

Permissions	Enabled	Create	Read	Update	Delete
Resource Group	Not applicable	Yes	Yes	Yes	Yes
Policy	Not applicable	Yes	Yes	Yes	Yes

Permissions	Enabled	Create	Read	Update	Delete
Backup	Not applicable	Yes	Yes	Yes	Yes
Host	Not applicable	Yes	Yes	Yes	Yes
Storage Connection	Not applicable	No	Yes	No	No
Clone	Not applicable	Yes	Yes	Yes	Yes
Provision	Not applicable	No	Yes	No	No
Dashboard	Yes	Not applicable	Not applicable	Not applicable	Not applicable
Reports	Yes	Not applicable	Not applicable	Not applicable	Not applicable
Restore	Yes	Not applicable	Not applicable	Not applicable	Not applicable
Resource	Yes	Yes	Yes	Yes	Yes
Plug-in Install/Uninstall	No	Not applicable		Not applicable	Not applicable
Migration	No	Not applicable	Not applicable	Not applicable	Not applicable
Mount	Yes	Yes	Not applicable	Not applicable	Not applicable
Unmount	Yes	Yes	Not applicable	Not applicable	Not applicable
Full Volume Restore	No	No	Not applicable	Not applicable	Not applicable
SecondaryProtection	No	No	Not applicable	Not applicable	Not applicable
Job Monitor	Yes	Not applicable	Not applicable	Not applicable	Not applicable

Backup and Clone Viewer role

The Backup and Clone Viewer role has read-only view of all permissions. This role also has permissions enabled for discovery, reporting, and access to the Dashboard.

Permissions	Enabled	Create	Read	Update	Delete
Resource Group	Not applicable	No	Yes	No	No

Permissions	Enabled	Create	Read	Update	Delete
Policy	Not applicable	No	Yes	No	No
Backup	Not applicable	No	Yes	No	No
Host	Not applicable	No	Yes	No	No
Storage Connection	Not applicable	No	Yes	No	No
Clone	Not applicable	No	Yes	No	No
Provision	Not applicable	No	Yes	No	No
Dashboard	Yes	Not applicable	Not applicable	Not applicable	Not applicable
Reports	Yes	Not applicable	Not applicable	Not applicable	Not applicable
Restore	No	No	Not applicable	Not applicable	Not applicable
Resource	No	No	Yes	Yes	No
Plug-in Install/Uninstall	No	Not applicable	Not applicable	Not applicable	Not applicable
Migration	No	Not applicable	Not applicable	Not applicable	Not applicable
Mount	Yes	Not applicable	Not applicable	Not applicable	Not applicable
Unmount	Yes	Not applicable	Not applicable	Not applicable	Not applicable
Full Volume Restore	No	Not applicable	Not applicable	Not applicable	Not applicable
SecondaryProtection	No	Not applicable	Not applicable	Not applicable	Not applicable
Job Monitor	Yes	Not applicable	Not applicable	Not applicable	Not applicable

Infrastructure Admin role

The Infrastructure Admin role has permissions enabled for host management, storage management, provisioning, resource groups, remote installation reports, and access to the Dashboard.

Permissions	Enabled	Create	Read	Update	Delete
Resource Group	Not applicable	Yes	Yes	Yes	Yes
Policy	Not applicable	No	Yes	Yes	Yes
Backup	Not applicable	Yes	Yes	Yes	Yes
Host	Not applicable	Yes	Yes	Yes	Yes
Storage Connection	Not applicable	Yes	Yes	Yes	Yes
Clone	Not applicable	No	Yes	No	No
Provision	Not applicable	Yes	Yes	Yes	Yes
Dashboard	Yes	Not applicable	Not applicable	Not applicable	Not applicable
Reports	Yes	Not applicable	Not applicable	Not applicable	Not applicable
Restore	Yes	Not applicable	Not applicable	Not applicable	Not applicable
Resource	Yes	Yes	Yes	Yes	Yes
Plug-in Install/Uninstall	Yes	Not applicable	Not applicable	Not applicable	Not applicable
Migration	No	Not applicable	Not applicable	Not applicable	Not applicable
Mount	No	Not applicable	Not applicable	Not applicable	Not applicable
Unmount	No	Not applicable	Not applicable	Not applicable	Not applicable
Full Volume Restore	No	No	Not applicable	Not applicable	Not applicable
SecondaryProtection	No	No	Not applicable	Not applicable	Not applicable
Job Monitor	Yes	Not applicable	Not applicable	Not applicable	Not applicable

Disaster recovery in SnapCenter

The SnapCenter disaster recovery (DR) feature lets you recover from disasters like resource corruption or server crashes. It helps to restore the SnapCenter repository,

server schedules, configuration components, and the SnapCenter Plug-in for SQL Server and its storage.

This section explains the two types of DR in SnapCenter:

SnapCenter Server DR

- SnapCenter Server data is backed up and can be recovered without any plug-in added to or managed by the SnapCenter Server.
- Secondary SnapCenter Server should be installed on the same installation directory and on the same port as the primary SnapCenter Server.
- For Multi-factor authentication (MFA), during SnapCenter Server DR, close all the browser tabs and reopen a browser to log in again. This will clear the existing or active session cookies and update that the correct configuration data.
- SnapCenter disaster recovery functionality uses REST APIs to backup the SnapCenter Server. See REST API workflows for disaster recovery of SnapCenter Server.
- Audit settings-related configuration file is not backed up in DR backup and neither on the DR server after the restore operation. You should manually repeat the Audit log settings.

SnapCenter Plug-in and Storage DR

DR is available only for SnapCenter Plug-in for SQL Server. If the plug-in is down, switch to another SQL host and recover the data by following a few steps. See Disaster recovery of SnapCenter Plug-in for SQL Server.

SnapCenter uses ONTAP SnapMirror to replicate data, which can be used for DR by keeping data synchronized at to a secondary site. To initiate failover, break the SnapMirror replication. During fallback, reverse the synchronization to replicate data from the DR site back to the primary location.

Licenses required by SnapCenter

SnapCenter requires several licenses to enable data protection of applications, databases, file systems, and virtual machines. The type of SnapCenter licenses you install depends on your storage environment and the features that you want to use.

License	Where required
SnapCenter Standard controller-based	Required for FAS, AFF, ASA SnapCenter Standard license is a controller-based license and is included as part of NetApp ONTAP One. If you have the SnapManager Suite license, you also get the SnapCenter Standard license entitlement. If you want to install SnapCenter on a trial basis with FAS, AFF, or ASA storage, you can obtain a NetApp ONTAP One evaluation license by contacting the sales representative. For information on licenses included with NetApp ONTAP One, refer to Licenses included with NetApp ONTAP One. SnapCenter is also offered as part of data protection bundle. If you have purchased A400 or later, you should purchase the data protection bundle.
SnapMirror or SnapVault	ONTAP Either SnapMirror or SnapVault license is required if replication is enabled in SnapCenter.
SnapRestore	Required to restore and verify backups. On primary storage systems • Required on SnapVault destination systems to perform remote verification and to restore from a backup. • Required on SnapMirror destination systems to perform remote verification.
FlexClone	Required to clone databases and verification operations. On primary and secondary storage systems • Required on SnapVault destination systems to create clones from secondary vault backup. • Required on SnapMirror destination systems to create clones from secondary SnapMirror backup.

License	Where required
Protocol licenses	• iSCSI or FC license for LUNs • CIFS license for SMB shares • NFS license for NFS type VMDKs • iSCSI or FC license for VMFS type VMDKs Required on SnapMirror destination systems to serve data if a source volume is unavailable.
SnapCenter Standard licenses (optional)	Secondary destinations It is recommended, but not required, that you add SnapCenter Standard licenses to secondary destinations. If SnapCenter Standard licenses are not enabled on secondary destinations, you cannot use SnapCenter to backup resources on the secondary destination after performing a failover operation. However, a FlexClone license is required on secondary destinations to perform clone and verification operations.
Single Mailbox Recovery (SMBR) licenses	If you are using SnapCenter Plug-in for Exchange to manage Microsoft Exchange Server databases and Single Mailbox Recovery (SMBR), you would need additional license for SMBR which needs to be purchased separately based on user mailbox. NetApp® Single Mailbox Recovery has come to the end of availability (EOA) on May 12, 2023. For more information, refer CPC-00507. NetApp will continue to support customers that have purchased mailbox capacity, maintenance, and support through marketing part numbers introduced on June 24, 2020, for the duration of the support entitlement. NetApp Single Mailbox Recovery is a partner product provided by Ontrack. Ontrack PowerControls offers capabilities that are similar to those of NetApp Single Mailbox Recovery. Customers can procure new Ontrack PowerControls software licenses and Ontrack PowerControls maintenance and support renewals from Ontrack (through licensingteam@ontrack.com) for granular mailbox recovery after the May 12, 2023, EOA date.

SnapCenter Advanced and SnapCenter NAS File Services licenses are deprecated, and are no longer available. The standard license and capacity-based license are no longer required for Amazon FSx for NetApp ONTAP, ONTAP Select, Cloud Volumes ONTAP, and Azure NetApp Files.

You should install one or more SnapCenter licenses. For information on how to add licenses, see Add SnapCenter Standard controller-based licenses.

SnapMirror active sync in SnapCenter

SnapMirror active sync enables business services to continue operating even through a complete site failure, supporting applications to fail over transparently using a secondary copy. Neither manual intervention nor additional scripting is required to trigger a failover with SnapMirror active sync.

For more information on SnapMirror active sync refer SnapMirror active sync overview.

For SnapMirror active sync, ensure that you have met the various hardware, software, and system configuration requirements. For information refer Prerequisites

The plug-ins supported for this feature are SnapCenter Plug-in for SQL Server, SnapCenter Plug-in for Windows, SnapCenter Plug-in for Oracle database, SnapCenter Plug-in for SAP HANA database, SnapCenter Plug-in for Microsoft Exchange Server, and SnapCenter Plug-in for Unix.

To support host initiator proximity in SnapCenter, it's value, either source or destination should be set in ONTAP.

The use cases not supported in SnapCenter:

- If you convert the existing asymmetric SnapMirror active sync workloads to symmetric by changing the policy on the SnapMirror active sync relationships from *automatedfailover* to *automatedfailoverduplex* in ONTAP, the same is not supported in SnapCenter.
- If there are backups of a resource group (already protected in SnapCenter) and then storage policy is changed on the SnapMirror active sync relationships from *automatedfailover* to *automatedfailoverduplex* in ONTAP, the same is not supported in SnapCenter.

Key concepts of data protection

Before using SnapCenter, understand key concepts for backup, clone, and restore.

Resources

Resources include databases, Windows file systems, or file shares backed up or cloned with SnapCenter. Depending on your environment, resources might also be database instances, SQL Server availability groups, Oracle databases, RAC databases, or custom application groups.

Resource group

A resource group is a collection of resources on a host or cluster, potentially from multiple hosts and clusters. Operations performed on a resource group apply to all its resources based on the specified schedule. You can perform on-demand or scheduled backups for individual resources or groups.

If one host in a shared resource group enters maintenance mode, all scheduled operations for that group will be suspended across all hosts.

Use relevant plug-ins to back up specific resources: database plug-ins for databases, file system plug-ins for file systems, and SnapCenter Plug-in for VMware vSphere for VMs and datastores.

Policies

Policies specify the backup frequency, copy retention, replication, scripts, and other characteristics of data protection operations.

One or more policies can be selected when creating a resource group or when performing an on-demand backup.

A resource group defines what needs to be protected and when it should be protected in terms of day and time. A policy describes how the protection will be carried out. For example, if backing up all databases or file systems of a host is necessary, a resource group including all databases or file systems in the host might be created. Two policies could then be attached to the resource group: a daily policy and an hourly policy.

When creating the resource group and attaching the policies, it is possible to configure it to perform a full backup daily and another schedule for log backups hourly.

Custom prescripts and postscripts can be used in data protection operations. These scripts allow automation either before or after the data protection job. For instance, a script could automatically notify of data protection job failures or warnings. Understanding the requirements for creating these scripts is crucial before setting up prescripts and postscripts.

Usage of prescripts and postscripts

Custom prescripts and postscripts can automate your data protection tasks before or after the job. For instance, you can add a script to notify you of job failures or warnings. Before setting them up, ensure you understand the requirements for these scripts.

Supported script types

The following types of scripts are supported for Windows:

- Batch files
- PowerShell scripts
- Perl scripts

The following types of scripts are supported for UNIX:

- Perl scripts
- Python scripts
- Shell scripts

Along with default bash shell other shells like sh-shell, k-shell, and c-shell are also supported.

Script path

All prescripts and postscripts that are run as part of SnapCenter operations on both nonvirtualized and virtualized storage systems, are executed on the plug-in host.

- The Windows scripts should be located on the plug-in host.

  The prescripts or postscripts path should not include drives or shares. The path should be relative to the SCRIPTS_PATH.

- The UNIX scripts should be located on the plug-in host.

  The script path is validated at the time of execution.

Where to specify scripts

Scripts are specified in backup policies. When a backup job starts, the policy automatically associates the script with the resources being backed up. When you create a backup policy, you can specify the prescript and postscript arguments.

You cannot specify multiple scripts.

Script timeouts

The timeout is set to 60 seconds, by default. You can modify the timeout value.

Script output

The default directory for the Windows prescripts and postscripts output files is Windows\System32.

There is no default location for the UNIX prescripts and postscripts. You can redirect the output file to any preferred location.

Storage systems and applications supported by SnapCenter

You should know the storage systems, applications, and databases supported by SnapCenter.

Supported storage systems

- NetApp ONTAP 9.12.1 and later
- Azure NetApp Files
- Amazon FSx for NetApp ONTAP

  Supports non-volatile memory express (NVMe) over Transport Control Protocol (TCP).

  For information about Amazon FSx for NetApp ONTAP, see Amazon FSx for NetApp ONTAP documentation.

- NetApp ASA r2 systems that are running NetApp ONTAP 9.16.1.

Supported applications and databases

SnapCenter supports protection of different applications and databases. For detailed information about the supported applications and databases, see NetApp Interoperability Matrix Tool.

SnapCenter supports protection of Oracle and Microsoft SQL workloads in VMware Cloud on Amazon Web Services (AWS) Software-Defined Data Center (SDDC) environments.
Learn More.

Authentication methods for SnapCenter credentials

Credentials use different authentication methods depending on the application or environment. Credentials authenticate users so they can perform SnapCenter operations. You should create one set of credentials for installing plug-ins and another for data protection operations.

Windows authentication

The Windows authentication method authenticates against Active Directory. For Windows authentication, Active Directory is set up outside of SnapCenter. SnapCenter authenticates with no additional configuration. You need a Windows credential to add hosts, install plug-in packages, and schedule jobs.

Untrusted domain authentication

SnapCenter allows users and groups belonging to untrusted domains to create Windows credentials. For the authentication to succeed, you should register the untrusted domains with SnapCenter.

Local workgroup authentication

SnapCenter allows the creation of Windows credentials with local workgroup users and groups. The Windows authentication for local workgroup users and groups does not happen during Windows credential creation but is deferred until the host registration and other host operations are performed.

SQL Server authentication

The SQL authentication method authenticates against a SQL Server instance. This means that a SQL Server instance must be discovered in SnapCenter. Therefore, before adding a SQL credential, you must add a host, install plug-in packages, and refresh resources. You need SQL Server authentication to perform operations such as scheduling on SQL Server or discovering resources.

Linux authentication

The Linux authentication method authenticates against a Linux host. You need Linux authentication during the initial step of adding the Linux host and installing the SnapCenter Plug-ins Package for Linux remotely from the SnapCenter GUI.

AIX authentication

The AIX authentication method authenticates against an AIX host. You need AIX authentication during the initial step of adding the AIX host and installing the SnapCenter Plug-ins Package for AIX remotely from the SnapCenter GUI.

Oracle database authentication

The Oracle database authentication method authenticates against an Oracle database. You need an Oracle database authentication to perform operations on the Oracle database if the operating system (OS) authentication is disabled on the database host. Therefore, before adding an Oracle database credential, you should create an Oracle user in the Oracle database with sysdba privileges.

Oracle ASM authentication

The Oracle ASM authentication method authenticates against an Oracle Automatic Storage Management (ASM) instance. Oracle ASM authentication is required if you need to access an Oracle ASM instance and OS authentication is disabled on the database host. Before adding an Oracle ASM credential, create an Oracle user with system privileges in the ASM instance.

RMAN catalog authentication

The RMAN catalog authentication method authenticates against the Oracle Recovery Manager (RMAN) catalog database. If you have configured an external catalog mechanism and registered your database to catalog database, you need to add RMAN catalog authentication.

Supported SnapCenter operations for ASA r2 systems

ASA r2 storage systems are supported beginning with SnapCenter 6.1. Learn more about ASA r2 systems

SnapCenter leverages REST APIs to perform all operations on ASA r2 systems, which do not support ZAPIs.

Operations supported by SnapCenter for ASA r2 systems

- Creating primary backups of applications over VMDK
- Transferring consistency group snapshots to secondary storage system
- Restoring the backups from primary and secondary storage systems to the original host or to the alternate host
  - In-place restore from both primary and secondary storage systems using VMware vMotion
  - Connect and copy restore from both primary and secondary storage systems
- Cloning the backups to the original host or to the alternate host

SnapCenter can discover or create ONTAP consistency groups. It can also provision and initialize SnapMirror relationships to the destination cluster for secondary protection.

For information on enabling secondary protection on ASA r2 systems for your application, refer to:

- Enable secondary protection for Microsoft SQL Server resources
- Enable secondary protection for SAP HANA resources
- Enable secondary protection for Oracle resources
- Enable secondary protection for Windows file systems
- Enable secondary protection for IBM Db2 resources
- Enable secondary protection for PostgreSQL resources
- Enable secondary protection for MySQL resources

- Enable secondary protection for Unix file systems

Operations not supported by SnapCenter for ASA r2 systems

- Raw Device Mapping (RDM)
- Application volumes for Oracle
- SAP HANA NDV
- LockVault
- Tamperproof snapshots
- FlexGroup volumes
- Hierarchical consistency group
- Migration from ASA, AFF, or FAS storage systems to ASA r2 storage systems
- Protection of databases that have mix of ASA, AFF, or FAS resources and ASA r2 resources
- Renaming of snapshots
- Secondary provisioning of the log directory of SQL plug-in host

Quick start for SnapCenter software

The quick start guide outlines the basic steps to install and configure the SnapCenter software.

Prepare for installing SnapCenter Server

You should ensure that all the requirements to install the SnapCenter Server are met.

- Requirements
- Register to access SnapCenter software
- Enable multifactor authentication

Install SnapCenter Server

The SnapCenter Server can be installed on either Windows or Linux hosts. Download the SnapCenter Server installation package from the NetApp Support Site and execute the installer.

- Install SnapCenter server on Windows
- Install SnapCenter Server on Linux

Configure SnapCenter Server

After installing the SnapCenter Server, you should configure it based on your environment.

Install the plug-in for your application

Ensure that all requirements for installing the application-specific plug-in are fulfilled based on the application in

use, then proceed with installing the respective plug-in.

Protect your application

After successfully installing the SnapCenter Server and the necessary plug-ins, you can initiate the creation of application backups. These backups can later be utilized for restoration and cloning purposes when required.

Install and configure SnapCenter Server

Prepare for installing the SnapCenter Server

Requirements to install SnapCenter Server

Before you install SnapCenter Server either on Windows or Linux host, you should review and ensure that all the requirements are met for your environment.

Domain and workgroup requirements for Windows host

The SnapCenter Server can be installed on a Windows host that is either in a domain or in a workgroup.

The user having admin privileges is allowed to install the SnapCenter server.

- Active Directory domain: You must use a Domain user with local administrator rights. The Domain user must be a member of the local Administrator group on the Windows host.
- Workgroups: You must use a local account that has local administrator rights.

While domain trusts, multi-domain forests, and cross-domain trusts are supported, cross-forest domains are not supported. The Microsoft documentation about Active Directory Domains and Trusts contains more information.

> After installing the SnapCenter Server, you should not change the domain in which the SnapCenter host is located. If you remove the SnapCenter Server host from the domain it was in when the SnapCenter Server was installed and then try to uninstall SnapCenter Server, the uninstall operation fails.

Space and sizing requirements

You should be familiar with the space and sizing requirements.

Item	Windows host requirements	Linux host requirements
Operating Systems	Microsoft Windows Only English, German, Japanese, and simplified Chinese version of the operating systems are supported. For the latest information about supported versions, see NetApp Interoperability Matrix Tool.	• Red Hat Enterprise Linux (RHEL) 8 and 9 • SUSE Linux Enterprise Server (SLES) 15 For the latest information about supported versions, see NetApp Interoperability Matrix Tool.
Minimum CPU count	4 cores	4 cores

Item	Windows host requirements	Linux host requirements
Minimum RAM	8 GB The MySQL Server buffer pool uses 20 percent of the total RAM.	8 GB
Minimum hard drive space for the SnapCenter Server software and logs	7 GB If you have the SnapCenter repository in the same drive where SnapCenter Server is installed, then it is recommended to have 15 GB.	15 GB
Minimum hard drive space for the SnapCenter repository	8 GB NOTE: If you have the SnapCenter Server in the same drive where SnapCenter repository is installed, then it is recommended to have 15 GB.	Not applicable
Required software packages	• ASP.NET Core Runtime 8.0.12 (and all subsequent 8.0.x patches) Hosting Bundle • PowerShell 7.4.2 or later For .NET specific troubleshooting information, see SnapCenter upgrade or install fails for legacy systems that do not have internet connectivity.	• .NET Framework 8.0.12 (and all subsequent 8.0.x patches) • PowerShell 7.4.2 or later • Nginx is a web server that can be used as a reverse proxy • Pam-devel PAM (Pluggable Authentication Modules) is a system security tool which allows system administrators to set authentication policy without having to recompile programs which do authentication.

ASP.NET core needs IIS_IUSRS to access the temp file system in SnapCenter Server on Windows.

SAN host requirements

SnapCenter does not include host utilities or a DSM. If the SnapCenter host is part of a SAN (FC/iSCSI) environment, you might need to install and configure additional software on the SnapCenter Server host.

- Host Utilities: The Host Utilities support FC and iSCSI, and it enables you to use MPIO on your Windows Servers. Learn More.
- Microsoft DSM for Windows MPIO: This software works with Windows MPIO drivers to manage multiple paths between NetApp and Windows host computers. A DSM is required for high availability configurations.

If you were using ONTAP DSM, you should migrate to Microsoft DSM. For more information, see How to migrate from ONTAP DSM to Microsoft DSM.

Browser requirements

SnapCenter software supports Chrome 125 and later and Microsoft Edge 110.0.1587.17 and later.

Port requirements

The SnapCenter software requires different ports for communication between different components.

- Applications cannot share a port.
- For customizable ports, you can select a custom port during installation if you do not want to use the default port.
- For fixed ports, you should accept the default port number.
- Firewalls
  - Firewalls, proxies, or other network devices should not interfere with connections.
  - If you specify a custom port when you install SnapCenter, you should add a firewall rule on the plug-in host for that port for the SnapCenter Plug-in Loader.

The following table lists the different ports and their default values.

Port Name	Port Numbers	Protocol	Direction	Description
SnapCenter port	8146	HTTPS	Bidirectional	This port is used for communication between the SnapCenter client (the SnapCenter user) and the SnapCenter Server and is also used for communication from the plug-in hosts to the SnapCenter Server. You can customize the port number.
SnapCenter SMCore communication port	8145	HTTPS	Bidirectional	This port is used for communication between the SnapCenter Server and the hosts where the SnapCenter plug-ins are installed. You can customize the port number.
Scheduler Service Port	8154	HTTPS		This is port is used to orchestrate the SnapCenter scheduler workflows for all the managed plug-ins within the SnapCenter server host in centralized manner. You can customize the port number.
RabbitMQ Port	5672	TCP		This is the default port that RabbitMQ listens on and it is used for publisher-subscriber model communication between Scheduler service and SnapCenter.

Port Name	Port Numbers	Protocol	Direction	Description
MySQL port	3306	HTTPS		The port is used for communicating with SnapCenter repository database. You can create secured connections from the SnapCenter Server to the MySQL server. Learn more
Windows plug-in hosts	135, 445	TCP		This port is used for communication between the SnapCenter Server and the host on which the plug-in is being installed. Additional dynamic port range specified by Microsoft should also be open.
Linux or AIX plug-in hosts	22	SSH	Unidirectional	This port is used for communication between the SnapCenter Server and the host, initiated from the server to client host.
SnapCenter Plug-ins Package for Windows, Linux or AIX	8145	HTTPS	Bidirectional	This port is used for communication between SMCore and hosts where the plug-ins package is installed. Customizable. You can customize the port number.
SnapCenter Plug-in for Oracle Database	27216			The default JDBC port is used by the plug-in for Oracle for connecting to the Oracle database.

Port Name	Port Numbers	Protocol	Direction	Description
SnapCenter Plug-in for Exchange Database	909			The default NET.TCP port is used by the plug-in for Windows for connecting to the Exchange VSS call-backs.
NetApp supported plug-ins for SnapCenter	9090	HTTPS		This is an internal port that is used only on the plug-in host; no firewall exception is required. Communication between the SnapCenter Server and plug-ins is routed through port 8145.
ONTAP cluster or SVM communication port	• 443 (HTTPS) • 80 (HTTP)	• HTTPS • HTTP	Bidirectional	The port is used by the SAL (Storage Abstraction Layer) for communication between the host running SnapCenter Server and SVM. The port is currently also used by the SAL on SnapCenter for Windows Plug-in hosts for communication between the SnapCenter plug-in host and SVM.
SnapCenter Plug-in for SAP HANA Database	• 3instance_number13 • 3instance_number15	• HTTPS • HTTP	Bidirectional	For a multitenant database container (MDC) single tenant, the port number ends with 13; for non MDC, the port number ends with 15. You can customize the port number.

Port Name	Port Numbers	Protocol	Direction	Description
SnapCenter Plug-in for PostgreSQL	5432			This port is the default PostgreSQL port used for communication by the plug-in for PostgreSQL to the PostgreSQL cluster. You can customize the port number.

Register to access the SnapCenter software

You should register to access the SnapCenter software if you are new to Amazon FSx for NetApp ONTAP or Azure NetApp Files and do not have an existing NetApp account.

Before you begin

- You should have access to the corporate email ID.
- If you using Azure NetApp Files, you should have the Azure subscription ID.
- If you are using Amazon FSx for NetApp ONTAP, you should have the File System ID of your FSx for ONTAP file system.

About this task

Your registration is subject to information validations and may take up to a day to confirm and upgrade new NetApp Support Site (NSS) account to **full** access from **guest** access.

Steps

1. Click https://mysupport.netapp.com/site/user/registration for registration.
2. Enter your corporate email ID, complete the captcha, accept NetApp's privacy policy, and click **Submit**.
3. Authenticate the registration by entering the OTP sent to your email ID and click **Continue**.
4. On the registration completion page, enter the following details to complete the registration.
   a. Select **NetApp Customer / End User**.
   b. In the SERIAL NUMBER field, either enter the Azure subscription ID if you are using Azure NetApp Files or the File System ID if you are using Amazon FSx for NetApp ONTAP.

You can raise a ticket at https://mysupport.netapp.com/site/help if you face any issue during registration or to know the status.

Multi-factor authentication (MFA)

Manage multi-factor authentication (MFA)

You can manage Multi-factor authentication (MFA) functionality in the Active Directory Federation Service (AD FS) Server and SnapCenter Server.

Enable multi-factor authentication (MFA)

You can enable MFA functionality for SnapCenter Server using PowerShell commands.

About this task

- SnapCenter supports SSO based logins when other applications are configured in the same AD FS. In certain AD FS configurations, SnapCenter might require user authentication for security reasons depending on the AD FS session persistence.
- The information regarding the parameters that can be used with the cmdlet and their descriptions can be obtained by running `Get-Help command_name`. Alternatively, you can also see SnapCenter Software Cmdlet Reference Guide.

Before you begin

- Windows Active Directory Federation Service (AD FS) should be up and running in the respective domain.
- You should have an AD FS supported Multi-factor authentication service such as Azure MFA, Cisco Duo, and so on.
- SnapCenter and AD FS server timestamp should be the same regardless of the timezone.
- Procure and configure the authorized CA certificate for SnapCenter Server.

  CA Certificate is mandatory for the following reasons:

  - Ensures that the ADFS-F5 communications do not break because the self-signed certificates are unique at the node level.
  - Ensures that during upgrade, repair, or disaster recovery (DR) in a standalone or high availability configuration, the self-signed certificate does not get recreated thus avoiding MFA reconfiguration.
  - Ensures IP-FQDN resolutions.

    For information on CA certificate, see Generate CA Certificate CSR file.

Steps

1. Connect to the Active Directory Federation Services (AD FS) host.
2. Download AD FS federation metadata file from "https://<host FQDN>/FederationMetadata/2007-06/FederationMetadata.xml".
3. Copy the downloaded file to SnapCenter Server to enable MFA feature.
4. Log in to SnapCenter Server as the SnapCenter Administrator user through PowerShell.
5. Using the PowerShell session, generate the SnapCenter MFA metadata file by using the *New-SmMultifactorAuthenticationMetadata -path* cmdlet.

   The path parameter specifies the path to save the MFA metadata file in the SnapCenter Server host.

6. Copy the generated file to the AD FS host to configure SnapCenter as the client entity.
7. Enable MFA for SnapCenter Server using the `Set-SmMultiFactorAuthentication` cmdlet.
8. (Optional) Check the MFA configuration status and settings by using `Get-SmMultiFactorAuthentication` cmdlet.
9. Go to the Microsoft management console (MMC) and perform the following steps:
   a. Click **File** > **Add/Remove Snapin**.

b. In the Add or Remove Snap-ins window, select **Certificates** and then click **Add**.
c. In the Certificates snap-in window, select the **Computer account** option, and then click **Finish**.
d. Click **Console Root** > **Certificates – Local Computer** > **Personal** > **Certificates**.
e. Right-click on the CA certificate bound to SnapCenter and then select **All Tasks** > **Manage Private Keys**.
f. On the permissions wizard perform the following steps:
   i. Click **Add**.
   ii. Click **Locations** and select the concerned host (top of hierarchy).
   iii. Click **OK** in the **Locations** pop-up window.
   iv. In the object name field, enter 'IIS_IUSRS' and click **Check Names** and click **OK**.

   If the check is successful, click **OK**.

10. In the AD FS host, open AD FS management wizard and perform the following steps:
    a. Right click on **Relying Party Trusts** > **Add Relying Party Trust** > **Start**.
    b. Select the second option and browse the SnapCenter MFA Metadata file and click **Next**.
    c. Specify a display name and click **Next**.
    d. Choose an access control policy as required and click **Next**.
    e. Select the settings in the next tab to default.
    f. Click **Finish**.

    SnapCenter is now reflected as a relying party with the provided display name.

11. Select the name and perform the following steps:
    a. Click **Edit Claim Issuance Policy**.
    b. Click **Add Rule** and click **Next**.
    c. Specify a name for the claim rule.
    d. Select **Active Directory** as the attribute store.
    e. Select the attribute as **User-Principal-Name** and the outgoing claim type as **Name-ID**.
    f. Click **Finish**.
12. Run the following PowerShell commands on the ADFS server.

```
Set-AdfsRelyingPartyTrust -TargetName '<Display name of relying party >'
-SigningCertificateRevocationCheck None

Set-AdfsRelyingPartyTrust -TargetName '<Display name of relying party >'
-EncryptionCertificateRevocationCheck None
```

13. Perform the following steps to confirm that the metadata was imported successfully.
    a. Right-click the relying party trust and select **Properties**.
    b. Ensure that the Endpoints, Identifiers, and Signature fields are populated.
14. Close all the browser tabs and reopen a browser to clear the existing or active session cookies, and login again.

SnapCenter MFA functionality can also be enabled using REST APIs.

For troubleshooting information, see Simultaneous login attempts in multiple tabs shows MFA error.

Update AD FS MFA Metadata

You should update the AD FS MFA metadata in SnapCenter whenever there is any modification in the AD FS Server, such as upgrade, CA certificate renewal, DR, and so on.

Steps

1. Download AD FS federation metadata file from "https://<host FQDN>/FederationMetadata/2007-06/FederationMetadata.xml"
2. Copy the downloaded file to SnapCenter Server to update the MFA configuration.
3. Update the AD FS metadata in SnapCenter by running the following cmdlet:

   ```
   Set-SmMultiFactorAuthentication -Path <location of ADFS MFA metadata xml file>
   ```

4. Close all the browser tabs and reopen a browser to clear the existing or active session cookies, and login again.

Update SnapCenter MFA metadata

You should update the SnapCenter MFA metadata in AD FS whenever there is any modification in ADFS server such as repair, CA certificate renewal, DR, and so on.

Steps

1. In the AD FS host, open AD FS management wizard and perform the following steps:
   a. Select **Relying Party Trusts**.
   b. Right click on the relying party trust that was created for SnapCenter and select **Delete**.

      The user defined name of the relying party trust will be displayed.

   c. Enable Multi-factor authentication (MFA).

      See Enable Multi-factor authentication.
2. Close all the browser tabs and reopen a browser to clear the existing or active session cookies, and login again.

Disable Multi-factor authentication (MFA)

Steps

1. Disable MFA and clean up the configuration files that were created when MFA was enabled by using the `Set-SmMultiFactorAuthentication` cmdlet.
2. Close all the browser tabs and reopen a browser to clear the existing or active session cookies, and login again.

Manage multi-factor authentication (MFA) using Rest API, PowerShell, and SCCLI

MFA login is supported from browser, REST API, PowerShell, and SCCLI. MFA is supported through an AD FS identity manager. You can enable MFA, disable MFA, and

configure MFA from GUI, REST API, PowerShell, and SCCLI.

Setup AD FS as OAuth/OIDC

Configure AD FS using Windows GUI wizard

1. Navigate to **Server Manager Dashboard** > **Tools** > **ADFS Management**.
2. Navigate to **ADFS** > **Application Groups**.
   a. Right-click on **Application Groups**.
   b. Select **Add Application group** and enter **Application Name**.
   c. Select **Server Application**.
   d. Click **Next**.
3. Copy **Client Identifier**.

   This is the Client ID.
   .. Add Callback URL (SnapCenter Server URL) in Redirect URL.
   .. Click **Next**.

4. Select **Generate shared secret**.

   Copy the secret value. This is the client's secret.
   .. Click **Next**.

5. On the **Summary** page, click **Next**.
   a. On the **Complete** page, click **Close**.
6. Right-click on the newly added **Application Group** and select **Properties**.
7. Select **Add application** from App Properties.
8. Click **Add application**.

   Select Web API and click **Next**.

9. On the Configure Web API page, enter the SnapCenter Server URL and Client Identifier created in the previous step into the Identifier section.
   a. Click **Add**.
   b. Click **Next**.
10. On the **Choose Access Control Policy** page, select control policy based on your requirement (For example, Permit everyone and require MFA) and click **Next**.
11. On the **Configure Application Permission** page, by default openid is selected as a scope, click **Next**.
12. On the **Summary** page, click **Next**.

    On the **Complete** page, click **Close**.

13. On the **Sample Application Properties** page, click **OK**.
14. JWT token issued by an authorization server (AD FS) and intended to be consumed by the resource.

    The 'aud' or audience claim of this token must match the identifier of the resource or Web API.

15. Edit the selected WebAPI and check that Callback URL (SnapCenter Server URL) and the client identifier were added correctly.

    Configure OpenID Connect to provide a username as claims.

16. Open the **AD FS Management** tool located under the **Tools** menu at the top right of the Server Manager.
    a. Select the **Application Groups** folder from the left sidebar.
    b. Select the Web API and click **EDIT**.
    c. Go-to Issuance Transform Rules Tab
17. Click **Add Rule**.
    a. Select the **Send LDAP Attributes as Claims** in the Claim rule template dropdown.
    b. Click **Next**.
18. Enter the **Claim rule** name.
    a. Select **Active Directory** in the Attribute store dropdown.
    b. Select **User-Principal-Name** in the **LDAP Attribute** dropdown and **UPN** in the O*utgoing Claim Type* dropdown.
    c. Click **Finish**.

Create Application Group using PowerShell commands

You can create the application group, web API, and add the scope and claims using PowerShell commands. These commands are available in automated script format. For more information see <link to KB article>.

1. Create the new Application Group in AD FS by using the following comamnd.

   ```
   New-AdfsApplicationGroup -Name $ClientRoleIdentifier
   -ApplicationGroupIdentifier $ClientRoleIdentifier
   ```

   `ClientRoleIdentifier` name of your application group

   `redirectURL` valid URL for redirection after authorization

2. Create the AD FS Server Application and generate the client secret.

   ```
   Add-AdfsServerApplication -Name "$ClientRoleIdentifier - Server app"
   -ApplicationGroupIdentifier $ClientRoleIdentifier -RedirectUri $redirectURL
   -Identifier $identifier -GenerateClientSecret
   ```

3. Create the ADFS Web API application and configure the policy name it should use.

   ```
   $identifier = (New-Guid).Guid

   Add-AdfsWebApiApplication -ApplicationGroupIdentifier $ClientRoleIdentifier
   -Name "App Web API"

   -Identifier $identifier -AccessControlPolicyName "Permit everyone"
   ```

4. Get the client ID and client secret from the output of the following commands because, it is shown only one time.

```
"client_id = $identifier"

"client_secret: "$($ADFSApp.ClientSecret)
```

5. Grant the AD FS Application the allatclaims and openid permissions.

```
Grant-AdfsApplicationPermission -ClientRoleIdentifier $identifier
-ServerRoleIdentifier $identifier -ScopeNames @('openid')

$transformrule = @"

@RuleTemplate = "LdapClaims"

@RuleName = "AD User properties and Groups"

c:[Type ==
"http://schemas.microsoft.com/ws/2008/06/identity/claims/windowsaccountname",
Issuer ==

"AD AUTHORITY"]

⇒ issue(store = "Active Directory", types =
("http://schemas.xmlsoap.org/ws/2005/05/identity/claims/upn"), query =
";userPrincipalName;{0}", param = c.Value);

"@
```

6. Write out the transform rules file.

```
$transformrule |Out-File -FilePath .\issueancetransformrules.tmp -force
-Encoding ascii
$relativePath = Get-Item .\issueancetransformrules.tmp
```

7. Name the Web API Application and define its Issuance Transform Rules using an external file.

```
Set-AdfsWebApiApplication -Name "$ClientRoleIdentifier - Web API"
-TargetIdentifier

$identifier -Identifier $identifier,$redirectURL -IssuanceTransformRulesFile

$relativePath
```

Update access token expiry time

You can update the access token expiry time using the PowerShell command.

About this task

- An access token can be used only for a specific combination of user, client, and resource. Access tokens cannot be revoked and are valid until their expiry.
- By default, the expiry time of an access token is 60 minutes. This minimal expiry time is sufficient and scaled. You must provide sufficient value to avoid any ongoing business-critical jobs.

Step

To update the access token expiry time for an application group WebApi, use the following command in AD FS server.

+
```
Set-AdfsWebApiApplication -TokenLifetime 3600 -TargetName "<Web API>"
```

Get the bearer token from AD FS

You should fill the below-mentioned parameters in any REST client (like Postman) and it prompts you to fill in the user credentials. Additionally, you should enter the second-factor authentication (something you have & something you are) to get the bearer token.

+
The validity of the bearer token is configurable from the AD FS server per application and the default validity period is 60 minutes.

Field	Value
Grant type	Authorization Code
Callback URL	Enter your application's base URL if you do not have a callback URL.
Auth URL	[adfs-domain-name]/adfs/oauth2/authorize
Access token URL	[adfs-domain-name]/adfs/oauth2/token
Client ID	Enter the AD FS client ID
Client secret	Enter the AD FS client secret
Scope	OpenID
Client Authentication	Send as Basic AUTH Header
Resource	In the **Advance Options** tab, add the Resource field with the same value as the Callback URL, which comes as an "aud" value in the JWT token.

Configure MFA in SnapCenter Server using PowerShell, SCCLI, and REST API

You can configure MFA in SnapCenter Server using PowerShell, SCCLI, and REST API.

SnapCenter MFA CLI authentication

In PowerShell and SCCLI, the existing cmdlet (Open-SmConnection) is extended with one more field called "AccessToken" to use the bearer token to authenticate the user.

```
Open-SmConnection -Credential <PSCredential> [-SMSbaseUrl <String>] [-Port
<String>] [-RoleName <String>] [ -AccessToken <string>]
```

After the above cmdlet is executed, a session is created for the respective user to execute further SnapCenter cmdlets.

SnapCenter MFA Rest API Authentication

Use bearer token in the format *Authorization=Bearer <access token>* in REST API client (like Postman or swagger) and mention the user RoleName in the header to get a successful response from SnapCenter.

MFA Rest API Workflow

When MFA is configured with AD FS, you should authenticate using an access (bearer) token to access the SnapCenter application by any Rest API.

About this task

- You can use any REST client like Postman, Swagger UI or FireCamp.
- Get an access token and use it to authenticate subsequent requests (SnapCenter Rest API) to perform any operation.

Steps

To authenticate through AD FS MFA

1. Configure the REST client to call AD FS endpoint to get the access token.

   When you hit the button to get an access token for an application, you will be redirected to the AD FS SSO page where you must provide your AD credentials and authenticate with MFA.
   1. In the AD FS SSO page, type your username or email in the Username text box.

   +
   Usernames must be formatted as user@domain or domain\user.

2. In the Password text box, type your password.
3. Click **Log in**.
4. From the **Sign-in Options** section, select an authentication option and authenticate (depending on your configuration).
   - Push: Approve the push notification that is sent to your phone.
   - QR Code: Use the AUTH Point mobile app to scan the QR code, then type the verification code shown in the app
   - One-Time Password: Type the one-time password for your token.
5. After successful authentication, a popup will open that contains the Access, ID, and Refresh Token.

   Copy the access token and use it in the SnapCenter Rest API to perform the operation.

6. In the Rest API, you should pass the access token and role name in the header section.
7. SnapCenter validates this access token from AD FS.

   If it is a valid token, SnapCenter decodes it and gets the username.

8. Using the Username and Role Name, SnapCenter authenticates the user for an API execution.

   If the authentication succeeds, SnapCenter returns the result else an error message is displayed.

Enable or disable SnapCenter MFA functionality for Rest API, CLI, and GUI

GUI

Steps

1. Log into the SnapCenter Server as the SnapCenter Administrator.
2. Click **Settings** > **Global Settings** > **MultiFactorAuthentication(MFA) Settings**
3. Select the interface (GUI/RST API/CLI) to enable or disable the MFA login.

PowerShell interface

Steps

1. Run the PowerShell or CLI commands for enabling MFA for GUI, Rest API, PowerShell, and SCCLI.

   ```
   Set-SmMultiFactorAuthentication -IsGuiMFAEnabled -IsRestApiMFAEnabled
   -IsCliMFAEnabled -Path
   ```

   The path parameter specifies the location of the AD FS MFA metadata xml file.

   Enables MFA for SnapCenter GUI, Rest API, PowerShell, and SCCLI configured with specified AD FS metadata file path.

2. Check the MFA configuration status and settings by using the `Get-SmMultiFactorAuthentication` cmdlet.

SCCLI Interface

Steps

1. ```
   # sccli Set-SmMultiFactorAuthentication -IsGuiMFAEnabled true
   -IsRESTAPIMFAEnabled true -IsCliMFAEnabled true -Path
   "C:\ADFS_metadata\abc.xml"
   ```
2. ```
   # sccli Get-SmMultiFactorAuthentication
   ```

REST APIs

1. Run the following post API for enabling MFA for GUI, Rest API, PowerShell, and SCCLI.

Parameter	Value
Requested URL	/api/4.9/settings/multifactorauthentication
HTTP method	Post

Request Body	{
	"IsGuiMFAEnabled": false,
	"IsRestApiMFAEnabled": true,
	"IsCliMFAEnabled": false,
	"ADFSConfigFilePath":
	"C:\\ADFS_metadata\\abc.xml"
	}

Response Body	{
	"MFAConfiguration": {
	"IsGuiMFAEnabled": false,
	"ADFSConfigFilePath":
	"C:\\ADFS_metadata\\abc.xml",
	"SCConfigFilePath": null,
	"IsRestApiMFAEnabled": true,
	"IsCliMFAEnabled": false,
	"ADFSHostName": "win-adfs-
	sc49.winscedom2.com"
	}
	}

2. Check the MFA configuration status and settings by using the following API.

Parameter	Value
Requested URL	/api/4.9/settings/multifactorauthentication
HTTP method	Get

Response Body	{
	"MFAConfiguration": {
	"IsGuiMFAEnabled": false,
	"ADFSConfigFilePath":
	"C:\\ADFS_metadata\\abc.xml",
	"SCConfigFilePath": null,
	"IsRestApiMFAEnabled": true,
	"IsCliMFAEnabled": false,
	"ADFSHostName": "win-adfs-
	sc49.winscedom2.com"
	}
	}

Install the SnapCenter Server

Install the SnapCenter Server on Windows host

You can run the SnapCenter Server installer executable to install the SnapCenter Server.

You can optionally perform several installation and configuration procedures by using PowerShell cmdlets. You

should be using PowerShell 7.4.2 or later.

Silent installation of the SnapCenter Server from the command-line is not supported.

Before you begin

- The SnapCenter Server host must be up to date with Windows updates with no pending system restarts.
- You should have ensured that MySQL Server is not installed on the host where you plan to install the SnapCenter Server.
- You should have enabled Windows installer debugging.

  See the Microsoft web site for information about enabling Windows installer logging.

  You should not install the SnapCenter Server on a host that has Microsoft Exchange Server, Active Directory, or Domain Name Servers.

Steps

1. Download the SnapCenter Server installation package from NetApp Support Site.
2. Initiate the SnapCenter Server installation by double-clicking the downloaded .exe file.

   After you initiate the installation, all the prechecks are performed and if the minimum requirements are not met appropriate error or warning messages are displayed.

   You can ignore the warning messages and proceed with installation; however, errors should be fixed.

3. Review the pre-populated values required for the SnapCenter Server installation and modify if required.

   You do not have to specify the password for MySQL Server repository database. During SnapCenter Server installation the password is auto generated.

   The special character “%” is not supported in the custom path for the repository database. If you include “%” in the path, installation fails.

4. Click **Install Now**.

   If you have specified any values that are invalid, appropriate error messages will be displayed. You should reenter the values, and then initiate the installation.

   If you click the **Cancel** button, the step that is being executed will be completed, and then start the rollback operation. The SnapCenter Server will be completely removed from the host.

   However, if you click **Cancel** when "SnapCenter Server site restart" or "Waiting for SnapCenter Server to start" operations are being performed, installation will proceed without cancelling the operation.

   Log files are always listed (oldest first) in the %temp% folder of the admin user. If you want to redirect the log locations, initiate the SnapCenter Server installation from the command prompt by running:`C:\installer_location\installer_name.exe /log"C:\"`

Features enabled on Windows host during installation

The SnapCenter Server installer enables the Windows features and roles on your Windows host during installation. These might be of interest for troubleshooting and maintaining the host system.

Category	Feature
Web Server	• Internet Information Services • World Wide Web Services • Common HTTP Features ◦ Default Document ◦ Directory Browsing ◦ HTTP Errors ◦ HTTP Redirection ◦ Static Content ◦ WebDAV Publishing • Health and Diagnostics ◦ Custom Logging ◦ HTTP Logging ◦ Logging Tools ◦ Request Monitor ◦ Tracing • Performance Features ◦ Static Content Compression • Security ◦ IP Security ◦ Basic Authentication ◦ Centralized SSL Certificate Support ◦ Client Certificate Mapping Authentication ◦ IIS Client Certificate Mapping Authentication ◦ IP and Domain Restrictions ◦ Request Filtering ◦ URL Authorization ◦ Windows Authentication • Application Development Features ◦ .NET Extensibility 4.5 ◦ Application Initialization ◦ ASP.NET Core Runtime 8.0.12 (and all subsequent 8.0.x patches) Hosting Bundle ◦ Server-Side Includes ◦ WebSocket Protocol • Management Tools ◦ IIS Management Console

Category	Feature
IIS Management Scripts and Tools	• IIS Management Service • Web Management Tools
.NET Framework 8.0.12 Features	• ASP.NET Core Runtime 8.0.12 (and all subsequent 8.0.x patches) Hosting Bundle • Windows Communication Foundation (WCF) HTTP Activation45 ◦ TCP Activation ◦ HTTP Activation For .NET specific troubleshooting information, see SnapCenter upgrade or install fails for legacy systems that do not have internet connectivity.
Message Queuing	• Message Queuing Services Ensure that no other applications uses the MSMQ service that SnapCenter creates and manages. • RabbitMQ • Erlang
Windows Process Activation Service	• Process Model
Configuration APIs	All

Install the SnapCenter Server on Linux host

You can run the SnapCenter Server installer executable to install the SnapCenter Server.

Before you begin

- If you want to install the SnapCenter Server using non-root user who does not have enough privileges to install SnapCenter, get the sudoers checksum file from the NetApp Support site. You should use appropriate checksum file based on the Linux Version.
- If the sudo package is not available in SUSE Linux, then install the sudo package to avoid authentication failure.
- For SUSE Linux, configure the hostname to avoid the installation failure.
- Check the secure Linux status by running the command `sestatus`. If the *SELinux status* is "enabled" and the *Current mode* is "enforcing", perform the following:
  - Run the command: `sudo semanage port -a -t http_port_t -p tcp <WEBAPP_EXTERNAL_PORT_>`

    The default value of *WEBAPP_EXTERNAL_PORT* is 8146

- If the firewall blocks the port, run `sudo firewall-cmd --add-port <WEBAPP_EXTERNAL_PORT_>/tcp`

  The default value of *WEBAPP_EXTERNAL_PORT* is 8146

- Run the following commands from the directory where you have read and write permission:
  - `sudo ausearch -c 'nginx' --raw | audit2allow -M my-nginx`

    If the command return "nothing to do", rerun the command after installing SnapCenter Server.

  - If the command creates *my-nginx.pp*, run the command to make the policy package active: `sudo semodule -i my-nginx.pp`
- The path used for MySQL PID directory is */var/opt/mysqld*. Run the following commands to set the permissions for MySQL installation.
  - `mkdir /var/opt/mysqld`
  - `sudo semanage fcontext -a -t mysqld_var_run_t "/var/opt/mysqld(/.*)?"`
  - `sudo restorecon -Rv /var/opt/mysqld`
- The path used for MySQL Data directory is */INSTALL_DIR/NetApp/snapcenter/SnapManagerWeb/Repository/MySQL/*. Run the following commands to set the permissions for MySQL data directory.
  - `mkdir -p /INSTALL_DIR/NetApp/snapcenter/SnapManagerWeb/Repository/MySQL`
  - `sudo semanage fcontext -a -t mysqld_db_t "/INSTALL_DIR/NetApp/snapcenter/SnapManagerWeb/Repository/MySQL(/.*)?"`
  - `sudo restorecon -Rv /INSTALL_DIR/NetApp/snapcenter/SnapManagerWeb/Repository/MySQL`

About this task

- When SnapCenter Server is installed on the Linux host, third-party services such as MySQL, RabbitMq, Errlang gets installed. You should not uninstall them.
- The SnapCenter Server installed on the Linux host does not support:
  - High availability
  - Windows plug-ins
  - Active Directory (Supports only the local users, both root and non-root user with creds)
  - Key based authentication to log into SnapCenter
- During the installation of .NET runtime, if the installation fails to resolve the dependencies of *libicu* library, then install *libicu* by running the command: `yum install -y libicu`
- If the installation of SnapCenter Server fails due to the non-availability of *Perl*, then install *Perl* by running the command: `yum install -y perl`

Steps

1. Download the following from NetApp Support Site to */home directory*.
   - SnapCenter Server installation package - **snapcenter-linux-server-(el8/el9/sles15).bin**
   - Public key file - **snapcenter_public_key.pub**

◦ Respective signature file - **snapcenter-linux-server-(el8/el9/sles15).bin.sig**

2. Validate the signature file.
   `$openssl dgst -sha256 -verify snapcenter_public_key.pub -signature <path to signature file> <path to bin file>`
3. For non-root user installation, add the visudo content specified in **snapcenter_server_checksum_(el8/el9/sles15).txt** available along with the .bin installer.
4. Assign the execute permission for the .bin installer.
   `chmod +x snapcenter-linux-server-(el8/el9/sles15).bin`
5. Perform one of the actions to install SnapCenter Server.

If you want to perform…	Do this…
Interactive installation	`./snapcenter-linux-server-(el8/el9/sles15).bin` You will be prompted to enter the following details: • The webapp external port that is used to access SnapCenter Server outside the Linux host. The default value is 8146. • The SnapCenter Server user who will install SnapCenter Server. • The installation directory where packages will be installed.

If you want to perform…	Do this…
Non interactive installation	`sudo ./snapcenter-linux-server-(el8/el9/sles15).bin -i silent -DWEBAPP_EXTERNAL_PORT=<port> -DWEBAPP_INTERNAL_PORT=<port> -DSMCORE_PORT=<port> -DSCHEDULER_PORT=<port> -DSNAPCENTER_SERVER_USER=<user> -DUSER_INSTALL_DIR=<dir> -DINSTALL_LOG_NAME=<filename>` Example: sudo ./snapcenter_linux_server.bin -i silent -DWEBAPP_EXTERNAL_PORT=8146 -DSNAPCENTER_SERVER_USER=root -DUSER_INSTALL_DIR=/opt -DINSTALL_LOG_NAME=InstallerLog.log Logs will be stored at */var/opt/snapcenter/logs*. Parameters to be passed for installing SnapCenter Server: • DWEBAPP_EXTERNAL_PORT: Webapp external port that is used to access SnapCenter Server outside the Linux host. The default value is 8146. • DWEBAPP_INTERNAL_PORT: Webapp internal port that is used to access SnapCenter Server within the Linux host. The default value is 8147. • DSMCORE_PORT: SMCore port on which the smcore services are running. The default value is 8145. • DSCHEDULER_PORT: Scheduler port on which the scheduler services are running. The default value is 8154. • DSNAPCENTER_SERVER_USER: SnapCenter Server user who will install SnapCenter Server. For *DSNAPCENTER_SERVER_USER*, the default is the user running the installer. • DUSER_INSTALL_DIR: Installation directory where packages will be installed. For *DUSER_INSTALL_DIR*, the default installation directory is */opt*. • DINSTALL_LOG_NAME: Log file name where installation logs will be stored. This is an optional parameter and if specified no logs will be displayed on the console.If you do not specify this parameter, logs will be displayed on the console and also stored in the default log file. DSELINUX: If the *SELinux status* is "enabled", the *Current mode* is "enforcing", and you have

What's next?

- If the *SELinux status* is "enabled" and the *Current mode* is "enforcing", the **nginx** service fails to start. You should run the the following commands:
  1. Go to home directory.
  2. Run the command: `journalctl -x|grep nginx.`
  3. If the Webapp internal port (8147) is not allowed to listen, run the following commands:
     - `ausearch -c 'nginx' --raw | audit2allow -M my-nginx`
     - `semodule -i my-nginx.pp`
  4. Run `setsebool -P httpd_can_network_connect on`

Features enabled on Linux host during installation

The SnapCenter Server installs below software packages which can help in troubleshooting and maintaining the host system.

- Rabbitmq
- Erlang

Register SnapCenter

If you are new to NetApp products and do not have an existing NetApp account, you should register SnapCenter to enable support.

Steps

1. After installing SnapCenter, navigate to **Help > About**.
2. In the *About SnapCenter* dialog box, make a note of the SnapCenter Instance, a 20 digit number that starts with 971.
3. Click https://register.netapp.com.
4. Click **I am not a registered NetApp Customer**.
5. Specify your details to register yourself.
6. Leave the NetApp Reference SN field blank.
7. Select **SnapCenter** from the Product Line drop-down.
8. Select the billing provider.
9. Enter the 20-digit SnapCenter instance ID.
10. Click **Submit**.

Log in to SnapCenter using RBAC authorization

SnapCenter supports role-based access control (RBAC). SnapCenter admin assigns roles and resources through SnapCenter RBAC to either a user in workgroup or active directory, or to groups in active directory. The RBAC user can now log in to SnapCenter with the assigned roles.

Before you begin

- executed the commands mentioned in Before you begin section, you should specify this parameter and assign the value as 1. The default value is 0.
- DUPGRADE: The default value is 0. Specify this parameter and it's value as any integer other than 0 to upgrade the SnapCenter Server.

- You should enable Windows Process Activation Service (WAS) in Windows Server Manager.
- If you want to use Internet Explorer as the browser to log in to the SnapCenter Server, you should ensure that the Protected Mode in Internet Explorer is disabled.
- If SnapCenter Server is installed on Linux host, you should log in using the user account which was used to install the SnapCenter Server.

About this task

During installation, the SnapCenter Server Install wizard creates a shortcut and places it on the desktop and in the Start menu of the host where SnapCenter is installed. Additionally, at the end of the installation, the Install wizard displays the SnapCenter URL based on the information that you provided during installation, which you can copy if you want to log in from a remote system.

If you have multiple tabs open in your web browser, closing just the SnapCenter browser tab does not log you out of SnapCenter. To end your connection with SnapCenter, you must log out of SnapCenter either by clicking the **Sign out** button, or by closing the entire web browser.

Best Practice: For security reasons, it is recommended that you do not enable your browser to save your SnapCenter password.

The default GUI URL is a secure connection to the default port 8146 on the server where the SnapCenter Server is installed (*https://server:8146*). If you provided a different server port during the SnapCenter installation, that port is used instead.

For High Availability (HA) deployment, you must access SnapCenter using the virtual cluster IP *https://Virtual_Cluster_IP_or_FQDN:8146*. If you do not see the SnapCenter UI when you navigate to *https://Virtual_Cluster_IP_or_FQDN:8146* in Internet Explorer (IE), you must add the Virtual Cluster IP address or FQDN as a trusted site in IE on each plug-in host, or you must disable IE Enhanced Security on each plug-in host.
For more information, see Unable to access cluster IP address from outside network.

In addition to using the SnapCenter GUI, you can use PowerShell cmdlets to create scripts to perform configuration, backup, and restore operations. Some cmdlets might have changed with each SnapCenter release. The SnapCenter Software Cmdlet Reference Guide has the details.

If you are logging in to SnapCenter for the first time, you must log in using the credentials that you provided during the install process.

Steps

1. Launch SnapCenter from the shortcut located on your local host desktop, or from the URL provided at the end of the installation, or from the URL provided by your SnapCenter administrator.
2. Enter user credentials.

To specify the following…	Use one of these formats…
Domain administrator	• NetBIOS\UserName • UserName@UPN suffix For example, username@netapp.com • Domain FQDN\UserName
Local administrator	UserName

3. If you are assigned more than one role, from the Role box, select the role that you want to use for this login session.

   Your current user and associated role are shown in the upper right of SnapCenter after you are logged in.

Result

The Dashboard page is displayed.

If the logging fails with the error that site cannot be reached, you should map the SSL certificate to SnapCenter. Learn more

After you finish

After logging to SnapCenter Server as an RBAC user for the first time, refresh the resources list.

If you have untrusted Active Directory domains that you want SnapCenter to support, you must register those domains with SnapCenter before configuring the roles for the users on untrusted domains. Learn more.

If you want to add the plug-in host in SnapCenter running on Linux host, you should get the checksum file from the location: */opt/NetApp/snapcenter/SnapManagerWeb/Repository*.

From 6.0 release, a shortcut for SnapCenter PowerShell is created on the desktop. You can directly access the SnapCenter PowerShell cmdlets by using the shortcut.

Log in to SnapCenter using Multi-Factor Authentication (MFA)

SnapCenter Server supports MFA for domain account, which is part of the active directory.

Before you begin

You should have enabled MFA. For information on how to enable MFA, see Enable Multi-factor authentication

About this task

- Only FQDN is supported
- Workgroup and cross domain users cannot login using MFA

Steps

1. Launch SnapCenter from the shortcut located on your local host desktop, or from the URL provided at the end of the installation, or from the URL provided by your SnapCenter administrator.

2. In the AD FS login page, enter Username and Password.

   When the username or password invalid error message is displayed on the AD FS page, you should check for the following:

   - Whether the username or password is valid

     The user account should exist in the Active Directory (AD)

   - Whether you exceeded the maximum allowed attempts that was set in AD
   - Whether AD and AD FS is up and running

Modify the SnapCenter default GUI session timeout

You can modify the SnapCenter GUI session timeout period to make it less than or greater than the default timeout period of 20 minutes.

As a security feature, after a default period of 15 minutes of inactivity, SnapCenter warns you that you will be logged out of the GUI session in 5 minutes. By default, SnapCenter logs you out of the GUI session after 20 minutes of inactivity, and you must log in again.

Steps

1. In the left navigation pane, click **Settings** > **Global Settings**.
2. In the Global Settings page, click **Configuration Settings**.
3. In the Session Timeout field, enter the new session timeout in minutes, and then click **Save**.

Secure the SnapCenter web server by disabling SSL 3.0

For security purposes, you should disable Secure Socket Layer (SSL) 3.0 protocol in Microsoft IIS if it is enabled on your SnapCenter web server.

There are flaws in the SSL 3.0 protocol that an attacker can use to cause connection failures, or to perform man-in-the-middle attacks and observe the encryption traffic between your website and its visitors.

Steps

1. To launch Registry Editor on the SnapCenter web server host, click **Start** > **Run**, and then enter regedit.
2. In Registry Editor, navigate to HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\SecurityProviders\SCHANNEL\Protocols\SSL 3.0\.
   - If the Server key already exists:
     i. Select the Enabled DWORD, and then click **Edit** > **Modify**.
     ii. Change the value to 0, and then click **OK**.
   - If the Server key does not exist:
     i. Click **Edit** > **New** > **Key**, and then name the key Server.
     ii. With the new Server key selected, click **Edit** > **New** > **DWORD**.
     iii. Name the new DWORD Enabled, and then enter 0 as the value.
3. Close Registry Editor.

Configure the SnapCenter Server

Add and provision the storage system

Add storage systems

You should set up the storage system that gives SnapCenter access to ONTAP storage, ASA r2 systems, or Amazon FSx for NetApp ONTAP to perform data protection and provisioning operations.

You can either add a stand-alone SVM or a cluster comprising of multiple SVMs. If you are using Amazon FSx for NetApp ONTAP, you can either add FSx admin LIF comprising of multiple SVMs using fsxadmin account or add FSx SVM in SnapCenter.

Before you begin

- You should have the required permissions in the Infrastructure Admin role to create storage connections.
- You should ensure that the plug-in installations are not in progress.

  Host plug-in installations must not be in progress while adding a storage system connection because the host cache might not be updated and databases status might be displayed in the SnapCenter GUI as “Not available for backup” or “Not on NetApp storage”.

- Storage system names should be unique.

  SnapCenter does not support multiple storage systems with the same name on different clusters. Each storage system that is supported by SnapCenter should have a unique name and a unique data LIF IP address.

About this task

- When you configure storage systems, you can also enable Event Management System (EMS) & AutoSupport features. The AutoSupport tool collects data about the health of your system and automatically sends the data to NetApp technical support, enabling them to troubleshoot your system.

  If you enable these features, SnapCenter sends AutoSupport information to the storage system and EMS messages to the storage system syslog when a resource is protected, a restore or clone operation finishes successfully, or an operation fails.

- If you are planning to replicate Snapshots to a SnapMirror destination or SnapVault destination, you must set up storage system connections for the destination SVM or Cluster as well as the source SVM or Cluster.

If you change the storage system password, scheduled jobs, on demand backup, and restore operations might fail. After you change the storage system password, you can update the password by clicking **Modify** in the Storage tab.

Steps

1. In the left navigation pane, click **Storage Systems**.
2. In the Storage Systems page, click **New**.
3. In the Add Storage System page, provide the following information:

For this field…	Do this…
Storage System	Enter the storage system name or IP address. Storage system names, not including the domain name, must have 15 or fewer characters, and the names must be resolvable. To create storage system connections with names that have more than 15 characters, you can use the Add-SmStorageConnectionPowerShell cmdlet. For storage systems with MetroCluster configuration (MCC), it is recommended to register both local and peer clusters for non-disruptive operations. SnapCenter does not support multiple SVMs with the same name on different clusters. Each SVM that is supported by SnapCenter must have a unique name. After adding the storage connection to SnapCenter, you should not rename the SVM or the Cluster using ONTAP. If SVM is added with a short name or FQDN then it has to be resolvable from both the SnapCenter and the plug-in host.
User name/Password	Enter the credentials of the storage user that has the required privileges to access the storage system.

For this field…	Do this…
Event Management System (EMS) & AutoSupport Settings	If you want to send EMS messages to the storage system syslog or if you want to have AutoSupport messages sent to the storage system for applied protection, completed restore operations, or failed operations, select the appropriate checkbox. When you select the **Send AutoSupport Notification for failed operations to storage system** checkbox, the **Log SnapCenter Server events to syslog** checkbox is also selected because EMS messaging is required to enable AutoSupport notifications.

4. Click **More Options** if you want to modify the default values assigned to platform, protocol, port, and timeout.
   a. In Platform, select one of the options from the drop-down list.

      If the SVM is the secondary storage system in a backup relationship, select the **Secondary** checkbox. When the **Secondary** option is selected, SnapCenter does not perform a license check immediately.

      If you have added SVM in SnapCenter then, user need to select the platform type from the dropdown manually.

   b. In Protocol, select the protocol that was configured during SVM or Cluster setup, typically HTTPS.
   c. Enter the port that the storage system accepts.

      The default port 443 typically works.

   d. Enter the time in seconds that should elapse before communication attempts are halted.

      The default value is 60 seconds.

   e. If the SVM has multiple management interfaces, select the **Preferred IP** checkbox, and then enter the preferred IP address for SVM connections.
   f. Click **Save**.
5. Click **Submit**.

Result

In the Storage Systems page, from the **Type** drop-down perform one of the following actions:

- Select **ONTAP SVMs** if you want to view all the SVMs that were added.

  If you have added FSx SVMs, the FSx SVMs are listed here.

- Select **ONTAP Clusters** if you want to view all the clusters that were added.

  If you have added FSx clusters using fsxadmin, the FSx clusters are listed here.

  When you click on the cluster name, all the SVMs that are part of the cluster are displayed in the Storage Virtual Machines section.

If a new SVM is added to the ONTAP cluster using ONTAP GUI, click **Rediscover** to view the newly added SVM.

After you finish

A cluster administrator must enable AutoSupport on each storage system node to send email notifications from all storage systems to which SnapCenter has access, by running the following command from the storage system command line:

```
autosupport trigger modify -node nodename -autosupport-message client.app.info
-to enable -noteto enable
```

The Storage Virtual Machine (SVM) administrator has no access to AutoSupport.

Storage connections and credentials

Before performing data protection operations, you should set up the storage connections and add the credentials that the SnapCenter Server and the SnapCenter plug-ins will use.

Storage connections

The storage connections give the SnapCenter Server and SnapCenter plug-ins access to the ONTAP storage. Setting up these connections also involves configuring AutoSupport and Event Management System (EMS) features.

Credentials

- Domain administrator or any member of the administrator group

  Specify the domain administrator or any member of the administrator group on the system on which you are installing the SnapCenter plug-in. Valid formats for the Username field are:

  - *NetBIOS\UserName*
  - *Domain FQDN\UserName*
  - *UserName@upn*
- Local administrator (for workgroups only)

  For systems that belong to a workgroup, specify the built-in local administrator on the system on which you are installing the SnapCenter plug-in. You can specify a local user account that belongs to the local administrators group if the user account has elevated privileges or the User Access control feature is disabled on the host system.

  The valid format for the Username field is: *UserName*
- Credentials for individual resource groups

  If you set up credentials for individual resource groups and the username does not have full admin privileges, you must assign at least the resource group and backup privileges to the username.

Provision storage on Windows hosts

Create and manage igroups

You create initiator groups (igroups) to specify which hosts can access a given LUN on the storage system. You can use SnapCenter to create, rename, modify, or delete an igroup on a Windows host.

Create an igroup

You can use SnapCenter to create an igroup on a Windows host. The igroup will be available in the Create Disk or Connect Disk wizard when you map the igroup to a LUN.

Steps

1. In the left navigation pane, click **Hosts**.
2. In the Hosts page, click **Igroup**.
3. In the Initiator Groups page, click **New**.
4. In the Create Igroup dialog box, define the igroup:

In this field…	Do this…
Storage System	Select the SVM for the LUN you will map to the igroup.
Host	Select the host on which you want to create the igroup.
Igroup Name	Enter the name of the igroup.
Initiators	Select the initiator.
Type	Select the initiator type, iSCSI, FCP, or mixed (FCP and iSCSI).

5. When you are satisfied with your entries, click **OK**.

   SnapCenter creates the igroup on the storage system.

Rename an igroup

You can use SnapCenter to rename an existing igroup.

Steps

1. In the left navigation pane, click **Hosts**.
2. In the Hosts page, click **Igroup**.
3. In the Initiator Groups page, click in the **Storage Virtual Machine** field to display a list of available SVMs, and then select the SVM for the igroup you want to rename.

4. In the list of igroups for the SVM, select the igroup you want to rename and click **Rename**.
5. In the Rename igroup dialog box, enter the new name for the igroup and click **Rename**.

Modify an igroup

You can use SnapCenter to add igroup initiators to an existing igroup. While creating an igroup you can add only one host. If you want to create an igroup for a cluster, you can modify the igroup to add other nodes to that igroup.

Steps

1. In the left navigation pane, click **Hosts**.
2. In the Hosts page, click **Igroup**.
3. In the Initiator Groups page, click in the **Storage Virtual Machine** field to display a drop-down list of available SVMs, then select the SVM for the igroup you want to modify.
4. In the list of igroups, select an igroup and click **Add Initiator to igroup**.
5. Select a host.
6. Select the initiators and click **OK**.

Delete an igroup

You can use SnapCenter to delete an igroup when you no longer need it.

Steps

1. In the left navigation pane, click **Hosts**.
2. In the Hosts page, click **Igroup**.
3. In the Initiator Groups page, click in the **Storage Virtual Machine** field to display a drop-down list of available SVMs, then select the SVM for the igroup you want to delete.
4. In the list of igroups for the SVM, select the igroup you want to delete and click **Delete**.
5. In the Delete igroup dialog box, click **OK**.

   SnapCenter deletes the igroup.

Create and manage disks

The Windows host sees LUNs on your storage system as virtual disks. You can use SnapCenter to create and configure an FC-connected or iSCSI-connected LUN.

- SnapCenter supports only basic disks. The dynamic disks are not supported.
- For GPT only one data partition and for MBR one primary partition is allowed that has one volume formatted with NTFS or CSVFS and has one mount path.
- Supported partition styles: GPT, MBR; in a VMware UEFI VM, only iSCSI disks are supported

SnapCenter does not support renaming a disk. If a disk that is managed by SnapCenter is renamed, SnapCenter operations will not succeed.

View the disks on a host

You can view the disks on each Windows host you manage with SnapCenter.

Steps

1. In the left navigation pane, click **Hosts**.
2. In the Hosts page, click **Disks**.
3. Select the host from the **Host** drop-down list.

   The disks are listed.

View clustered disks

You can view clustered disks on the cluster that you manage with SnapCenter. The clustered disks are displayed only when you select the cluster from the Hosts drop-down.

Steps

1. In the left navigation pane, click **Hosts**.
2. In the Hosts page, click **Disks**.
3. Select the cluster from the **Host** drop-down list.

   The disks are listed.

Establish an iSCSI session

If you are using iSCSI to connect to a LUN, you must establish an iSCSI session before you create the LUN to enable communication.

Before you begin

- You must have defined the storage system node as an iSCSI target.
- You must have started the iSCSI service on the storage system. Learn more

About this task

You can establish an iSCSI session only between the same IP versions, either from IPv6 to IPv6, or from IPv4 to IPv4.

You can use a link-local IPv6 address for iSCSI session management and for communication between a host and a target only when both are in the same subnet.

If you change the name of an iSCSI initiator, access to iSCSI targets is affected. After changing the name, you might require to reconfigure the targets accessed by the initiator so that they can recognize the new name. You must ensure to restart the host after changing the name of an iSCSI initiator.

If your host has more than one iSCSI interface, once you have established an iSCSI session to SnapCenter using an IP address on the first interface, you cannot establish an iSCSI session from another interface with a different IP address.

Steps

1. In the left navigation pane, click **Hosts**.
2. In the Hosts page, click **iSCSI Session**.
3. From the **Storage Virtual Machine** drop-down list, select the storage virtual machine (SVM) for the iSCSI target.
4. From the **Host** drop-down list, select the host for the session.
5. Click **Establish Session**.

   The Establish Session wizard is displayed.

6. In the Establish Session wizard, identify the target:

In this field…	Enter…
Target node name	The node name of the iSCSI target If there is an existing target node name, the name is displayed in read-only format.
Target portal address	The IP address of the target network portal
Target portal port	The TCP port of the target network portal
Initiator portal address	The IP address of the initiator network portal

7. When you are satisfied with your entries, click **Connect**.

   SnapCenter establishes the iSCSI session.

8. Repeat this procedure to establish a session for each target.

Create FC-connected or iSCSI-connected LUNs or disks

The Windows host sees the LUNs on your storage system as virtual disks. You can use SnapCenter to create and configure an FC-connected or iSCSI-connected LUN.

If you want to create and format disks outside of SnapCenter, only NTFS and CSVFS file systems are supported.

Before you begin

- You must have created a volume for the LUN on your storage system.

  The volume should hold LUNs only, and only LUNs created with SnapCenter.

  > ⓘ You cannot create a LUN on a SnapCenter-created clone volume unless the clone has already been split.

- You must have started the FC or iSCSI service on the storage system.
- If you are using iSCSI, you must have established an iSCSI session with the storage system.
- The SnapCenter Plug-ins Package for Windows must be installed only on the host on which you are

creating the disk.

About this task

- You cannot connect a LUN to more than one host unless the LUN is shared by hosts in a Windows Server failover cluster.
- If a LUN is shared by hosts in a Windows Server failover cluster that uses CSV (Cluster Shared Volumes), you must create the disk on the host that owns the cluster group.

Steps

1. In the left navigation pane, click **Hosts**.
2. In the Hosts page, click **Disks**.
3. Select the host from the **Host** drop-down list.
4. Click **New**.

   The Create Disk wizard opens.

5. In the LUN Name page, identify the LUN:

In this field…	Do this…
Storage System	Select the SVM for the LUN.
LUN path	Click **Browse** to select the full path of the folder containing the LUN.
LUN name	Enter the name of the LUN.
Cluster size	Select the LUN block allocation size for the cluster. Cluster size depends upon the operating system and applications.
LUN label	Optionally, enter descriptive text for the LUN.

6. In the Disk Type page, select the disk type:

Select…	If…
Dedicated disk	The LUN can be accessed by only one host. Ignore the **Resource Group** field.
Shared disk	The LUN is shared by hosts in a Windows Server failover cluster. Enter the name of the cluster resource group in the **Resource Group** field. You need to create the disk on only one host in the failover cluster.

Select…	If…
Cluster Shared Volume (CSV)	The LUN is shared by hosts in a Windows Server failover cluster that uses CSV. Enter the name of the cluster resource group in the **Resource Group** field. Make sure that the host on which you are creating the disk is the owner of the cluster group.

7. In the Drive Properties page, specify the drive properties:

Property	Description
Auto assign mount point	SnapCenter automatically assigns a volume mount point based on the system drive. For example, if your system drive is C:, auto assign creates a volume mount point under your C: drive (C:\scmnpt\). Auto assign is not supported for shared disks.
Assign drive letter	Mount the disk to the drive you select in the adjacent drop-down list.
Use volume mount point	Mount the disk to the drive path you specify in the adjacent field. The root of the volume mount point must be owned by the host on which you are creating the disk.
Do not assign drive letter or volume mount point	Choose this option if you prefer to mount the disk manually in Windows.
LUN size	Specify the LUN size; 150 MB minimum. Select MB, GB, or TB in the adjoining drop-down list.
Use thin provisioning for the volume hosting this LUN	Thin provision the LUN. Thin provisioning allocates only as much storage space as is needed at one time, allowing the LUN to grow efficiently to the maximum available capacity. Make sure there is enough space available on the volume to accommodate all the LUN storage you think you will need.

Property	Description
Choose partition type	Select GPT partition for a GUID Partition Table, or MBR partition for a Master Boot Record. MBR partitions might cause misalignment issues in Windows Server failover clusters. Unified extensible firmware interface (UEFI) partition disks are not supported.

8. In the Map LUN page, select the iSCSI or FC initiator on the host:

In this field…	Do this…
Host	Double-click the cluster group name to display a drop-down list that shows the hosts that belong to the cluster, and then select the host for the initiator. This field is displayed only if the LUN is shared by hosts in a Windows Server failover cluster.
Choose host initiator	Select **Fibre Channel** or **iSCSI**, and then select the initiator on the host. You can select multiple FC initiators if you are using FC with multipath I/O (MPIO).

9. In the Group Type page, specify whether you want to map an existing igroup to the LUN, or create a new igroup:

Select…	If…
Create new igroup for selected initiators	You want to create a new igroup for the selected initiators.
Choose an existing igroup or specify a new igroup for selected initiators	You want to specify an existing igroup for the selected initiators, or create a new igroup with the name you specify. Type the igroup name in the **igroup name** field. Type the first few letters of the existing igroup name to autocomplete the field.

10. In the Summary page, review your selections and then click **Finish**.

    SnapCenter creates the LUN and connects it to the specified drive or drive path on the host.

Resize a disk

You can increase or decrease the size of a disk as your storage system needs change.

About this task

- For thin provisioned LUN, the ONTAP lun geometry size is shown as the maximum size.
- For thick provisioned LUN, the expandable size (available size in the volume) is shown as the maximum size.
- LUNs with MBR-style partitions have a size limit of 2 TB.
- LUNs with GPT-style partitions have a storage system size limit of 16 TB.
- It is a good idea to make a Snapshot before resizing a LUN.
- If you need to restore a LUN from a Snapshot made before the LUN was resized, SnapCenter automatically resizes the LUN to the size of the Snapshot.

  After the restore operation, data added to the LUN after it was resized must be restored from a Snapshot made after it was resized.

Steps

1. In the left navigation pane, click **Hosts**.
2. In the Hosts page, click **Disks**.
3. Select the host from the Host drop-down list.

   The disks are listed.

4. Select the disk you want to resize and then click **Resize**.
5. In the Resize Disk dialog box, use the slider tool to specify the new size of the disk, or enter the new size in the Size field.

   If you enter the size manually, you need to click outside the Size field before the Shrink or Expand button is enabled appropriately. Also, you must click MB, GB, or TB to specify the unit of measurement.

6. When you are satisfied with your entries, click **Shrink** or **Expand**, as appropriate.

   SnapCenter resizes the disk.

Connect a disk

You can use the Connect Disk wizard to connect an existing LUN to a host, or to reconnect a LUN that has been disconnected.

Before you begin

- You must have started the FC or iSCSI service on the storage system.
- If you are using iSCSI, you must have established an iSCSI session with the storage system.
- You cannot connect a LUN to more than one host unless the LUN is shared by hosts in a Windows Server failover cluster.

- If the LUN is shared by hosts in a Windows Server failover cluster that uses CSV (Cluster Shared Volumes), then you must connect the disk on the host that owns the cluster group.
- The Plug-in for Windows needs to be installed only on the host on which you are connecting the disk.

Steps

1. In the left navigation pane, click **Hosts**.
2. In the Hosts page, click **Disks**.
3. Select the host from the **Host** drop-down list.
4. Click **Connect**.

   The Connect Disk wizard opens.

5. In the LUN Name page, identify the LUN to connect to:

In this field…	Do this…
Storage System	Select the SVM for the LUN.
LUN path	Click **Browse** to select the full path of the volume containing the LUN.
LUN name	Enter the name of the LUN.
Cluster size	Select the LUN block allocation size for the cluster. Cluster size depends upon the operating system and applications.
LUN label	Optionally, enter descriptive text for the LUN.

6. In the Disk Type page, select the disk type:

Select…	If…
Dedicated disk	The LUN can be accessed by only one host.
Shared disk	The LUN is shared by hosts in a Windows Server failover cluster. You need only connect the disk to one host in the failover cluster.
Cluster Shared Volume (CSV)	The LUN is shared by hosts in a Windows Server failover cluster that uses CSV. Make sure that the host on which you are connecting to the disk is the owner of the cluster group.

7. In the Drive Properties page, specify the drive properties:

Property	Description
Auto assign	Let SnapCenter automatically assign a volume mount point based on the system drive. For example, if your system drive is C:, the auto assign property creates a volume mount point under your C: drive (C:\scmnpt\). The auto assign property is not supported for shared disks.
Assign drive letter	Mount the disk to the drive you select in the adjoining drop-down list.
Use volume mount point	Mount the disk to the drive path you specify in the adjoining field. The root of the volume mount point must be owned by the host on which you are creating the disk.
Do not assign drive letter or volume mount point	Choose this option if you prefer to mount the disk manually in Windows.

8. In the Map LUN page, select the iSCSI or FC initiator on the host:

In this field…	Do this…
Host	Double-click the cluster group name to display a drop-down list that shows the hosts that belong to the cluster, then select the host for the initiator. This field is displayed only if the LUN is shared by hosts in a Windows Server failover cluster.
Choose host initiator	Select **Fibre Channel** or **iSCSI**, and then select the initiator on the host. You can select multiple FC initiators if you are using FC with MPIO.

9. In the Group Type page, specify whether you want to map an existing igroup to the LUN or create a new igroup:

Select…	If…
Create new igroup for selected initiators	You want to create a new igroup for the selected initiators.

Select…	If…
Choose an existing igroup or specify a new igroup for selected initiators	You want to specify an existing igroup for the selected initiators, or create a new igroup with the name you specify. Type the igroup name in the **igroup name** field. Type the first few letters of the existing igroup name to automatically complete the field.

10. In the Summary page, review your selections and click **Finish**.

    SnapCenter connects the LUN to the specified drive or drive path on the host.

Disconnect a disk

You can disconnect a LUN from a host without affecting the contents of the LUN, with one exception: If you disconnect a clone before it has been split off, you lose the contents of the clone.

Before you begin

- Make sure that the LUN is not in use by any application.
- Make sure that the LUN is not being monitored with monitoring software.
- If the LUN is shared, make sure to remove the cluster resource dependencies from the LUN and verify that all nodes in the cluster are powered on, functioning properly, and available to SnapCenter.

About this task

If you disconnect a LUN in a FlexClone volume that SnapCenter has created and no other LUNs on the volume are connected, SnapCenter deletes the volume. Before disconnecting the LUN, SnapCenter displays a message warning you that the FlexClone volume might be deleted.

To avoid automatic deletion of the FlexClone volume, you should rename the volume before disconnecting the last LUN. When you rename the volume, make sure that you change multiple characters than just the last character in the name.

Steps

1. In the left navigation pane, click **Hosts**.
2. In the Hosts page, click **Disks**.
3. Select the host from the **Host** drop-down list.

   The disks are listed.

4. Select the disk you want to disconnect, and then click **Disconnect**.
5. In the Disconnect Disk dialog box, click **OK**.

   SnapCenter disconnects the disk.

Delete a disk

You can delete a disk when you no longer need it. After you delete a disk, you cannot undelete it.

Steps

1. In the left navigation pane, click **Hosts**.
2. In the Hosts page, click **Disks**.
3. Select the host from the **Host** drop-down list.

   The disks are listed.

4. Select the disk you want to delete, and then click **Delete**.
5. In the Delete Disk dialog box, click **OK**.

   SnapCenter deletes the disk.

Create and manage SMB shares

To configure an SMB3 share on a storage virtual machine (SVM), you can use either the SnapCenter user interface or PowerShell cmdlets.

Best Practice: Using the cmdlets is recommended because it enables you to take advantage of templates provided with SnapCenter to automate share configuration.

The templates encapsulate best practices for volume and share configuration. You can find the templates in the Templates folder in the installation folder for the SnapCenter Plug-ins Package for Windows.

If you feel comfortable doing so, you can create your own templates following the models provided. You should review the parameters in the cmdlet documentation before creating a custom template.

Create an SMB share

You can use the SnapCenter Shares page to create an SMB3 share on a storage virtual machine (SVM).

You cannot use SnapCenter to back up databases on SMB shares. SMB support is limited to provisioning only.

Steps

1. In the left navigation pane, click **Hosts**.
2. In the Hosts page, click **Shares**.
3. Select the SVM from the **Storage Virtual Machine** drop-down list.
4. Click **New**.

   The New Share dialog opens.

5. In the New Share dialog, define the share:

In this field…	Do this…
Description	Enter descriptive text for the share.
Share name	Enter the share name, for example, test_share. The name you enter for the share will also be used as the volume name. The share name: • Must be a UTF-8 string. • Must not include the following characters: control characters from 0x00 to 0x1F (both inclusive), 0x22 (double quotes), and the special characters `\ / [ ] : (vertical bar) < > + = ; , ?`
Share path	• Click in the field to enter a new file system path, for example, /. • Double-click in the field to select from a list of existing file system paths.

6. When you are satisfied with your entries, click **OK**.

   SnapCenter creates the SMB share on the SVM.

Delete an SMB share

You can delete an SMB share when you no longer need it.

Steps

1. In the left navigation pane, click **Hosts**.
2. In the Hosts page, click **Shares**.
3. In the Shares page, click in the **Storage Virtual Machine** field to display a drop-down with a list of available storage virtual machines (SVMs), then select the SVM for the share you want to delete.
4. From the list of shares on the SVM, select the share you want to delete and click **Delete**.
5. In the Delete Share dialog box, click **OK**.

   SnapCenter deletes the SMB share from the SVM.

Reclaim space on the storage system

Although NTFS tracks the available space on a LUN when files are deleted or modified, it does not report the new information to the storage system. You can run the space reclamation PowerShell cmdlet on the Plug-in for Windows host to ensure that newly freed blocks are marked as available in storage.

If you are running the cmdlet on a remote plug-in host, you must have run the SnapCenterOpen-SMConnection cmdlet to open a connection to the SnapCenter Server.

Before you begin

- You must ensure that the space reclamation process has completed before performing a restore operation.
- If the LUN is shared by hosts in a Windows Server failover cluster, you must perform space reclamation on the host that owns the cluster group.
- For optimum storage performance, you should perform space reclamation as often as possible.

  You should ensure that the entire NTFS file system has been scanned.

About this task

- Space reclamation is time-consuming and CPU-intensive, so it is usually best to run the operation when storage system and Windows host usage is low.
- Space reclamation reclaims nearly all available space, but not 100 percent.
- You should not run disk defragmentation at the same time as you are performing space reclamation.

  Doing so can slow the reclamation process.

Step

From the application server PowerShell command prompt, enter the following command:

```
Invoke-SdHostVolumeSpaceReclaim -Path drive_path
```

drive_path is the drive path mapped to the LUN.

Provision storage using PowerShell cmdlets

If you do not want to use the SnapCenter GUI to perform host provisioning and space reclamation jobs, you can use the PowerShell cmdlets. You can use cmdlets directly or add them to scripts.

If you are running the cmdlets on a remote plug-in host, you must run the SnapCenter Open-SMConnection cmdlet to open a connection to the SnapCenter Server.

The information regarding the parameters that can be used with the cmdlet and their descriptions can be obtained by running *Get-Help command_name*. Alternatively, you can also refer to the SnapCenter Software Cmdlet Reference Guide.

If SnapCenter PowerShell cmdlets are broken due to removal of SnapDrive for Windows from the server, refer to SnapCenter cmdlets broken when SnapDrive for Windows is uninstalled.

Provision storage in VMware environments

You can use the SnapCenter Plug-in for Microsoft Windows in VMware environments to create and manage LUNs and manage Snapshots.

Supported VMware guest OS platforms

- Supported versions of Windows Server
- Microsoft cluster configurations

  Support for up to a maximum of 16 nodes supported on VMware when using the Microsoft iSCSI Software Initiator, or up to two nodes using FC

- RDM LUNs

  Support for a maximum of 56 RDM LUNs with four LSI Logic SCSI controllers for normal RDMS, or 42 RDM LUNs with three LSI Logic SCSI controllers on a VMware VM MSCS box-to-box Plug-in for Windows configuration

  Supports VMware ParaVirtual SCSI Controller. 256 disks can be supported on RDM disks.

For the latest information about supported versions, see NetApp Interoperability Matrix Tool.

VMware ESXi server-related limitations

- Installing the Plug-in for Windows on a Microsoft cluster on virtual machines using ESXi credentials is not supported.

  You should use your vCenter credentials when installing the Plug-in for Windows on clustered virtual machines.

- All clustered nodes must use the same target ID (on the virtual SCSI adapter) for the same clustered disk.
- When you create an RDM LUN outside of the Plug-in for Windows, you must restart the plug-in service to enable it to recognize the newly created disk.
- You cannot use iSCSI and FC initiators at the same time on a VMware guest OS.

Minimum vCenter privileges required for SnapCenter RDM operations

You should have the following vCenter privileges on the host to perform RDM operations in a guest OS:

- Datastore: Remove File
- Host: Configuration > Storage Partition Configuration
- Virtual Machine: Configuration

You must assign these privileges to a role at the Virtual Center Server level. The role to which you assign these privileges cannot be assigned to any user without root privileges.

After you assign these privileges, you can install the Plug-in for Windows on the guest OS.

Manage FC RDM LUNs in a Microsoft cluster

You can use the Plug-in for Windows to manage a Microsoft cluster using FC RDM LUNs, but you must first create the shared RDM quorum and shared storage outside the plug-in, and then add the disks to the virtual machines in the cluster.

Starting with ESXi 5.5, you can also use ESX iSCSI and FCoE hardware to manage a Microsoft cluster. The Plug-in for Windows includes out-of-box support for Microsoft clusters.

Requirements

The Plug-in for Windows provides support for Microsoft clusters using FC RDM LUNs on two different virtual machines that belong to two different ESX or ESXi servers, also known as cluster across boxes, when you meet specific configuration requirements.

- The virtual machines (VMs) must be running the same Windows Server version.
- ESX or ESXi server versions must be the same for each VMware parent host.
- Each parent host must have at least two network adapters.
- There must be at least one VMware Virtual Machine File System (VMFS) datastore shared between the two ESX or ESXi servers.
- VMware recommends that the shared datastore be created on an FC SAN.

  If necessary, the shared datastore can also be created over iSCSI.

- The shared RDM LUN must be in physical compatibility mode.
- The shared RDM LUN must be created manually outside of the Plug-in for Windows.

  You cannot use virtual disks for shared storage.

- A SCSI controller must be configured on each virtual machine in the cluster in physical compatibility mode:

  Windows Server 2008 R2 requires you to configure the LSI Logic SAS SCSI controller on each virtual machine. Shared LUNs cannot use the existing LSI Logic SAS controller if only one of its type exists and it is already attached to the C: drive.

  SCSI controllers of type paravirtual are not supported on VMware Microsoft clusters.

  When you add a SCSI controller to a shared LUN on a virtual machine in physical compatibility mode, you must select the **Raw Device Mappings** (RDM) option and not the **Create a new disk** option in the VMware Infrastructure Client.

- Microsoft virtual machine clusters cannot be part of a VMware cluster.
- You must use vCenter credentials and not ESX or ESXi credentials when you install the Plug-in for Windows on virtual machines that belongs to a Microsoft cluster.
- The Plug-in for Windows cannot create a single igroup with initiators from multiple hosts.

  The igroup containing the initiators from all ESXi hosts must be created on the storage controller prior to creating the RDM LUNs that will be used as shared cluster disks.

- Ensure that you create an RDM LUN on ESXi 5.0 using an FC initiator.

  When you create an RDM LUN, an initiator group is created with ALUA.

Limitations

The Plug-in for Windows supports Microsoft clusters using FC/iSCSI RDM LUNs on different virtual machines belonging to different ESX or ESXi servers.

This feature is not supported in releases before ESX 5.5i.

- The Plug-in for Windows does not support clusters on ESX iSCSI and NFS datastores.
- The Plug-in for Windows does not support mixed initiators in a cluster environment.

  Initiators must be either FC or Microsoft iSCSI, but not both.

- ESX iSCSI initiators and HBAs are not supported on shared disks in a Microsoft cluster.
- The Plug-in for Windows does not support virtual machine migration with vMotion if the virtual machine is part of a Microsoft cluster.
- The Plug-in for Windows does not support MPIO on virtual machines in a Microsoft cluster.

Create a shared FC RDM LUN

Before you can use FC RDM LUNs to share storage between nodes in a Microsoft cluster, you must first create the shared quorum disk and shared storage disk, and then add them to both virtual machines in the cluster.

The shared disk is not created using the Plug-in for Windows. You should create and then add the shared LUN to each virtual machine in the cluster.
For information, see Cluster Virtual Machines Across Physical Hosts.

Add SnapCenter Standard controller-based licenses

A SnapCenter Standard controller-based license is required if you are using FAS, AFF, or ASA storage controllers.

The controller-based license has the following characteristics:

- SnapCenter Standard entitlement included with purchase of Premium or Flash Bundle (not with the base pack)
- Unlimited storage usage
- Added directly to the FAS, AFF, or ASA storage controller by using either ONTAP System Manager or the ONTAP CLI.

  You do not enter any license information in the SnapCenter user interface for the SnapCenter controller-based licenses.

- Locked to the controller's serial number

For information on the licenses required, see SnapCenter licenses.

Step 1: Verify if the SnapManager Suite license is installed

You can use the SnapCenter user interface to check if a SnapManager Suite license is installed on FAS, AFF, or ASA primary storage systems and identify which systems that need licenses. SnapManager Suite licenses apply only to FAS, AFF, and ASA SVMs or clusters on primary storage systems.

If you already have a SnapManager Suite license on your controller, SnapCenter automatically provides the Standard controller-based license entitlement. The names SnapManagerSuite license and SnapCenter Standard controller-based license are used interchangeably, but they refer to the same license.

Steps

1. In the left navigation pane, select **Storage Systems**.
2. In the Storage Systems page, from the **Type** drop-down, select whether to view all the SVMs or clusters that were added:
   - To view all of the SVMs that were added, select **ONTAP SVMs**.
   - To view all of the clusters that were added, select **ONTAP Clusters**.

     When you select the cluster name, all of the SVMs that are part of the cluster are displayed in the Storage Virtual Machines section.
3. In the Storage Connections list, locate the Controller License column.

   The Controller License column displays the following status:

   - indicates that a SnapManager Suite license is installed on a FAS, AFF, or ASA primary storage system.
   - indicates that a SnapManager Suite license is not installed on a FAS, AFF, or ASA primary storage system.
   - Not applicable indicates that a SnapManager Suite license is not applicable because the storage controller is on Amazon FSx for NetApp ONTAP, Cloud Volumes ONTAP, ONTAP Select, or Secondary storage platforms.

Step 2: Identify the licenses installed on the controller

You can use the ONTAP command line to view all the licenses installed on your controller. You should be a cluster administrator on the FAS, AFF, or ASA system.

The controller displays the SnapCenter Standard controller-based license as the SnapManagerSuite license.

Steps

1. Log in to the NetApp controller using the ONTAP command line.
2. Enter the license show command, and then view the output to see if the SnapManagerSuite license is installed.

Example output

```
cluster1::> license show
(system license show)

Serial Number: 1-80-0000xx
Owner: cluster1
Package           Type     Description              Expiration
----------------- -------- ---------------------    ---------------
Base              site     Cluster Base License     -

Serial Number: 1-81-0000000000000000000000000xx
Owner: cluster1-01
Package           Type     Description              Expiration
----------------- -------- ---------------------    ---------------
NFS               license  NFS License              -
CIFS              license  CIFS License             -
iSCSI             license  iSCSI License            -
FCP               license  FCP License              -
SnapRestore       license  SnapRestore License      -
SnapMirror        license  SnapMirror License       -
FlexClone         license  FlexClone License        -
SnapVault         license  SnapVault License        -
SnapManagerSuite  license  SnapManagerSuite License -
```

In the example, the SnapManagerSuite license is installed, therefore, no additional SnapCenter licensing action is required.

Step 3: Retrieve the controller serial number

Get the controller serial number using the ONTAP command line. You must be a cluster administrator on the FAS, AFF, or ASA system to get your controller-based license serial number.

Steps

1. Log in to the controller using the ONTAP command line.
2. Enter the system show -instance command, and then review the output to locate the controller serial number.

Example output

```
cluster1::> system show -instance

Node: fasxxxx-xx-xx-xx
Owner:
Location: RTP 1.5
Model: FAS8080
Serial Number: 123451234511
Asset Tag: -
Uptime: 143 days 23:46
NVRAM System ID: xxxxxxxxx
System ID: xxxxxxxxxx
Vendor: NetApp
Health: true
Eligibility: true
Differentiated Services: false
All-Flash Optimized: false

Node: fas8080-41-42-02
Owner:
Location: RTP 1.5
Model: FAS8080
Serial Number: 123451234512
Asset Tag: -
Uptime: 144 days 00:08
NVRAM System ID: xxxxxxxxx
System ID: xxxxxxxxxx
Vendor: NetApp
Health: true
Eligibility: true
Differentiated Services: false
All-Flash Optimized: false
2 entries were displayed.
```

3. Record the serial numbers.

Step 4: Retrieve the serial number of the controller-based license

If you are using FAS, ASA, or AFF storage, you can retrieve the SnapCenter controller-based license from the NetApp Support Site before you install it using the ONTAP command line.

Before you begin

- You should have a valid NetApp Support Site login credentials.

  If you do not enter valid credentials, the system does not return any information for your search.

- You should have the controller serial number.

Steps

1. Log in to the NetApp Support Site.
2. Navigate to **Systems** > **Software Licenses**.
3. In the Selection Criteria area, ensure Serial Number (located on back of unit) is selected, enter the controller serial number, and then select **Go!**.

   A list of licenses for the specified controller is displayed.

4. Locate and record the SnapCenter Standard or SnapManagerSuite license.

Step 5: Add controller-based license

You can use the ONTAP command line to add a SnapCenter controller-based license when you are using FAS, AFF, or ASA systems, and you have a SnapCenter Standard or SnapManagerSuite license.

Before you begin

- You should be a cluster administrator on the FAS, AFF, or ASA system.
- You should have the SnapCenter Standard or SnapManagerSuite license.

About this task

If you want to install SnapCenter on a trial basis with FAS, AFF, or ASA storage, you can obtain a Premium Bundle evaluation license to install on your controller.

If you want to install SnapCenter on a trial basis, you should contact your sales representative to obtain a Premium Bundle evaluation license to install on your controller.

Steps

1. Log in to the NetApp cluster using the ONTAP command line.
2. Add the SnapManagerSuite license key:

   ```
   system license add -license-code license_key
   ```

   This command is available at the admin privilege level.

3. Verify that the SnapManagerSuite license is installed:

   ```
   license show
   ```

Step 6: Remove the trial license

If you are using a controller-based SnapCenter Standard license and need to remove the capacity-based trial license (serial number ending with “50”), you should use MySQL commands to remove the trial license manually. The trial license cannot be deleted using the SnapCenter user interface.

Removing a trial license manually is only required if you are using a SnapCenter Standard controller-based license.

Steps

1. On the SnapCenter Server, open a PowerShell window to reset the MySQL password.
   a. Run the Open-SmConnection cmdlet to establish connection with the SnapCenter Server for a SnapCenterAdmin account.
   b. Run the Set-SmRepositoryPassword to reset the MySQL password.

      For information about the cmdlets, see SnapCenter Software Cmdlet Reference Guide.

2. Open the command prompt and run mysql -u root -p to log into MySQL.

   MySQL prompts you for the password. Enter the credentials you provided while resetting the password.

3. Remove the trial license from the database:

```
use nsm;DELETE FROM nsm_License WHERE nsm_License_Serial_Number='510000050';
```

Configure High Availability

Configure SnapCenter Servers for High Availability

To support High Availability (HA) in SnapCenter running either on Windows or on Linux, you can install the F5 load balancer. F5 enables the SnapCenter Server to support active-passive configurations in up to two hosts that are in the same location. To use F5 Load Balancer in SnapCenter, you should configure the SnapCenter Servers and configure F5 load balancer.

You can also configure Network Load Balancing (NLB) to set up SnapCenter High Availability. You should manually configure NLB outside of SnapCenter installation for high availability.

For cloud environment, you can configure high availability using either Amazon Web Services (AWS) Elastic Load Balancing (ELB) and Azure load balancer.

Configure high availability using F5

For instruction to configure SnapCenter Servers for high availability using F5 load balancer, refer to How to configure SnapCenter Servers for high availability using F5 Load Balancer.

You must be a member of the Local Administrators group on the SnapCenter Servers (in addition to being assigned to the SnapCenterAdmin role) to use the following cmdlets for adding and removing F5 clusters:

- Add-SmServerCluster
- Add-SmServer
- Remove-SmServerCluster

  For more information, refer to SnapCenter Software Cmdlet Reference Guide.

Additional information

- After you install and configure SnapCenter for high availability, edit the SnapCenter desktop shortcut to point to the F5 cluster IP.
- If a failover occurs between SnapCenter Servers and if there is also an existing SnapCenter session, you must close the browser and log on to SnapCenter again.
- In load balancer setup (NLB or F5), if you add a host that is partially resolved by the NLB or F5 host and if the SnapCenter host is not able to reach out to this host, then the SnapCenter host page switches between hosts down and running state frequently. To resolve this issue, you should ensure that both the SnapCenter hosts are able to resolve the host in NLB or F5 host.
- SnapCenter commands for MFA settings should be executed on all the hosts. Relying party configuration should be done in the Active Directory Federation Services (AD FS) server using F5 cluster details. The host level SnapCenter UI access will be blocked after MFA is enabled.
- During failover, the audit log settings will not reflect on the second host. Hence, you should manually repeat the audit log settings on F5 passive host when it becomes active.

Configure high availability using Network Load Balancing (NLB)

You can configure Network Load Balancing (NLB) to set up SnapCenter High Availability. You should manually configure NLB outside of SnapCenter installation for high availability.

For information about how to configure Network Load Balancing (NLB) with SnapCenter refer to How to configure NLB with SnapCenter.

Configure high availability using AWS Elastic Load Balancing (ELB)

You can configure high availability SnapCenter environment in Amazon Web Services (AWS) by setting up two SnapCenter servers in separate availability zones (AZs) and configuring them for automatic failover. The architecture includes virtual private IP addresses, routing tables, and synchronization between active and standby MySQL databases.

Steps

1. Configure virtual private overlay IP in AWS. For information, refer to Configure virtual private overlay IP.
2. Prepare your Windows host
   a. Force IPv4 being prioritized above IPv6:
      - Location: HKLM\SYSTEM\CurrentControlSet\Services\Tcpip6\Parameters

- Key: DisabledComponents
- Type: REG_DWORD
- Value: 0x20

b. Ensure that the fully qualified domain names can be resolved via DNS or via local host configuration to the IPv4 addresses.

c. Ensure that you do not have a system proxy configured.

d. Ensure that the administrator password is same on both the Windows Server when using a setup without an Active Directory and the servers are not in one domain.

e. Add virtual IP on both Windows Servers.

3. Create the SnapCenter cluster.

a. Start Powershell and connect to SnapCenter.
```
Open-SmConnection
```

b. Create the cluster.
```
Add-SmServerCluster -ClusterName <cluster_name> -ClusterIP <cluster_ip>
-PrimarySCServerIP <primary_ip> -Verbose -Credential administrator
```

c. Add the secondary server.
```
Add-SmServer -ServerName <server_name> -ServerIP <server_ip>
-CleanUpSecondaryServer -Verbose -Credential administrator
```

d. Get the high availability details.
```
Get-SmServerConfig
```

4. Create the Lamda function to adjust the routing table in case the virtual private IP endpoint becomes unavailable, monitored by AWS CloudWatch. For information, refer to Create a Lambda function.

5. Create a monitor in CloudWatch to monitor the availability of the SnapCenter endpoint. An alarm is configured to trigger a Lambda function if the endpoint is unreachable. The Lambda function adjusts the routing table to redirect traffic to the active SnapCenter server. For information, refer to Create synthetic canaries.

6. Implement workflow using a step function as an alternative to CloudWatch monitoring, providing smaller failover times. The workflow includes a Lambda probe function to test the SnapCenter URL, a DynamoDB table for storing failure counts, and the Step Function itself.

a. Use a lambda function for probing the SnapCenter URL. For information, refer to Create Lambda function.

b. Create a DynamoDB table for storing the failure count between two Step Function iterations. For information, refer to Get started with DynamoDB table.

c. Create the Step Function. For information, refer to Step Function documentation.

d. Test a single step.

e. Test the complete function.

f. Create IAM Role and adjust permissions to be allowed to execute Lambda function.

g. Create schedule to trigger Step Function. For information, refer to Using Amazon EventBridge Scheduler to start a Step Functions.

Configure high availability using Azure load balancer

You can configure high availability SnapCenter environment using Azure load balancer.

Steps

1. Create virtual machines in a scale set using Azure portal. The Azure virtual machine scale set allows you to create and manage a group of load balanced virtual machines. The number of virtual machine instances can automatically increase or decrease in response to demand or a defined schedule. For information, refer to Create virtual machines in a scale set using Azure portal.
2. After configuring the virtual machines, log into each virtual machine in VM set and install SnapCenter Server in both the nodes.
3. Create the cluster in host 1.
   ```
   Add-SmServerCluster -ClusterName <cluster_name> -ClusterIP <specify the load balancer front end virtual ip> -PrimarySCServerIP <ip address> -Verbose -Credential <credentials>
   ```
4. Add the secondary server.
   ```
   Add-SmServer -ServerName <name of node2> -ServerIP <ip address of node2> -Verbose -Credential <credentials>
   ```
5. Obtain the high availability details.
   ```
   Get-SmServerConfig
   ```
6. If required, rebuild the secondary host.
   ```
   Set-SmRepositoryConfig -RebuildSlave -Verbose
   ```
7. Failover to the second host.
   ```
   Set-SmRepositoryConfig ActiveMaster <name of node2> -Verbose
   ```

== Switch from NLB to F5 for high availability

You can change your SnapCenter HA configuration from Network Load Balancing (NLB) to use F5 Load Balancer.

Steps

1. Configure SnapCenter Servers for high availability using F5. Learn more.
2. On the SnapCenter Server host, launch PowerShell.
3. Start a session by using the Open-SmConnection cmdlet, and then enter your credentials.
4. Update the SnapCenter Server to point to the F5 cluster IP address using the Update-SmServerCluster cmdlet.

   The information regarding the parameters that can be used with the cmdlet and their descriptions can be obtained by running *Get-Help command_name*. Alternatively, you can also refer to the SnapCenter Software Cmdlet Reference Guide.

High availability for the SnapCenter MySQL repository

MySQL replication is a feature of MySQL Server that enables you to replicate data from one MySQL database server (master) to another MySQL database server (slave). SnapCenter supports MySQL replication for high availability only on two Network Load Balancing-enabled (NLB-enabled) nodes.

SnapCenter performs read or write operations on the master repository and routes its connection to the slave repository when there is a failure on the master repository. The slave repository then becomes the master repository. SnapCenter also supports reverse replication, which is enabled only during failover.

If you want to use the MySQL high availability (HA) feature, you must configure Network Load Balancer (NLB) on the first node. The MySQL repository is installed on this node as part of the installation. While installing SnapCenter on the second node, you must join to the F5 of the first node and create a copy of the MySQL repository on the second node.

SnapCenter provides the *Get-SmRepositoryConfig* and *Set-SmRepositoryConfig* PowerShell cmdlets to manage MySQL replication.

The information regarding the parameters that can be used with the cmdlet and their descriptions can be obtained by running *Get-Help command_name*. Alternatively, you can also refer to the SnapCenter Software Cmdlet Reference Guide.

You must be aware of the limitations related to the MySQL HA feature:

- NLB and MySQL HA are not supported beyond two nodes.
- Switching from a SnapCenter standalone installation to an NLB installation or vice versa and switching from a MySQL standalone setup to MySQL HA are not supported.
- Automatic failover is not supported if the slave repository data is not synchronized with the master repository data.

  You can initiate a forced failover by using the *Set-SmRepositoryConfig* cmdlet.

- When failover is initiated, jobs that are running might fail.

  If failover happens because MySQL Server or SnapCenter Server is down, then any jobs that are running might fail. After failing over to the second node, all subsequent jobs run successfully.

For information about configuring high availability, see How to configure NLB and ARR with SnapCenter.

Configure role-based access control (RBAC)

Create a role

In addition to using the existing SnapCenter roles, you can create your own roles and customize the permissions.

To create your own roles, it is necessary to log in as the "SnapCenterAdmin" role.

Steps

1. In the left navigation pane, click **Settings**.
2. In the Settings page, click **Roles**.
3. Click .
4. Specify a name and description for the new role.

   Only the following special characters can be used in user names and group names: space ( ), hyphen (-), underscore (_), and colon (:).

5. Select **All members of this role can see other members' objects** to enable other members of the role to see resources such as volumes and hosts after they refresh the resources list.

   You should deselect this option if you do not want members of this role to see objects to which other

members are assigned.

> When this option is enabled, assigning users access to objects or resources is not required if users belong to the same role as the user who created the objects or resources.

6. In the Permissions page, select the permissions that you want to assign to the role, or click **Select All** to grant all permissions to the role.
7. Click **Submit**.

Add an NetApp ONTAP RBAC role using security login commands

You can use the security login commands to add a NetApp ONTAP RBAC role when your storage systems are running clustered ONTAP.

Before you begin

- Identify the task (or tasks) that you want to perform and the privileges required to perform these tasks.
- Grant privileges to commands and/or command directories.

  There are two levels of access for each command/command directory: all-access and read-only.

  You must always assign the all-access privileges first.

- Assign roles to users.
- Identify your configuration depending on whether your SnapCenter plug-ins are connected to the Cluster Administrator IP for the entire cluster or directly connected to a SVM within the cluster.

About this task

To simplify the configuration of these roles on storage systems, you can use the RBAC User Creator for NetApp ONTAP tool, which is posted on the NetApp Communities Forum.

This tool automatically handles setting up the ONTAP privileges correctly. For example, RBAC User Creator for NetApp ONTAP tool automatically adds the privileges in the correct order so that the all-access privileges appear first. If you add the read-only privileges first and then add the all-access privileges, ONTAP marks the all-access privileges as duplicates and ignores them.

> If you later upgrade SnapCenter or ONTAP, you should re-run the RBAC User Creator for NetApp ONTAP tool to update the user roles you created previously. User roles created for an earlier version of SnapCenter or ONTAP do not work properly with upgraded versions. When you re-run the tool, it automatically handles the upgrade. You do not need to recreate the roles.

More information about setting up ONTAP RBAC roles, see the ONTAP 9 Administrator Authentication and RBAC Power Guide.

Steps

1. On the storage system, create a new role by entering the following command:

```
security login role create <role_name\> -cmddirname "command" -access all
–vserver <svm_name\>
```

   - svm_name is the name of the SVM. If you leave this blank, it defaults to cluster administrator.
   - role_name is the name you specify for the role.

- command is the ONTAP capability.

  > You must repeat this command for each permission. Remember that all-access commands must be listed before read-only commands.

  For information about the list of permissions, see ONTAP CLI commands for creating roles and assigning permissions.

2. Create a user name by entering the following command:

   ```
   security login create -username <user_name\> -application ontapi -authmethod
   <password\> -role <name_of_role_in_step_1\> -vserver <svm_name\> -comment
   "user_description"
   ```

   - user_name is the name of the user you are creating.
   - <password> is your password. If you do not specify a password, the system will prompt you for one.
   - svm_name is the name of the SVM.

3. Assign the role to the user by entering the following command:

   ```
   security login modify username <user_name\> -vserver <svm_name\> -role
   <role_name\> -application ontapi -application console -authmethod <password\>
   ```

   - <user_name> is the name of the user you created in Step 2. This command lets you modify the user to associate it with the role.
   - <svm_name> is the name of the SVM.
   - <role_name> is the name of the role you created in Step 1.
   - <password> is your password. If you do not specify a password, the system will prompt you for one.

4. Verify that the user was created correctly by entering the following command:

   ```
   security login show -vserver <svm_name\> -user-or-group-name <user_name\>
   ```

   user_name is the name of the user you created in Step 3.

Create SVM roles with minimum privileges

There are several ONTAP CLI commands you must run when you create a role for a new SVM user in ONTAP. This role is required if you configure SVMs in ONTAP to use with SnapCenter and you do not want to use the vsadmin role.

Steps

1. On the storage system, create a role and assign all the permissions to the role.

   ```
   security login role create -vserver <svm_name\>- role <SVM_Role_Name\>
   -cmddirname <permission\>
   ```

   > You should repeat this command for each permission.

2. Create a user and assign the role to that user.

```
security login create -user <user_name\> -vserver <svm_name\> -application
ontapi -authmethod password -role <SVM_Role_Name\>
```

3. Unlock the user.

```
security login unlock -user <user_name\> -vserver <svm_name\>
```

ONTAP CLI commands for creating SVM roles and assigning permissions

There are several ONTAP CLI commands you should run to create SVM roles and assign permissions.

- `security login role create -vserver SVM_Name -role SVM_Role_Name -cmddirname "snapmirror list-destinations" -access all`
- `security login role create -vserver SVM_Name -role SVM_Role_Name -cmddirname "event generate-autosupport-log" -access all`
- `security login role create -vserver SVM_Name -role SVM_Role_Name -cmddirname "job history show" -access all`
- `security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "job show" -access all`
- `security login role create -vserver SVM_Name -role SVM_Role_Name -cmddirname "job stop" -access all`
- `security login role create -vserver SVM_Name -role SVM_Role_Name -cmddirname "lun" -access all`
- `security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "lun create" -access all`
- `security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "lun delete" -access all`
- `security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "lun igroup add" -access all`
- `security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "lun igroup create" -access all`
- `security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "lun igroup delete" -access all`
- `security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "lun igroup rename" -access all`
- `security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "lun igroup show" -access all`
- `security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "lun mapping add-reporting-nodes" -access all`
- `security login role create -vserver SVM_Name -role SVM_Role_Name -cmddirname "lun mapping create" -access all`
- `security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "lun mapping delete" -access all`

- `security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "lun mapping remove-reporting-nodes" -access all`
- `security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "lun mapping show" -access all`
- `security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "lun modify" -access all`
- `security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "lun move-in-volume" -access all`
- `security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "lun offline" -access all`
- `security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "lun online" -access all`
- `security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "lun resize" -access all`
- `security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "lun serial" -access all`
- `security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "lun show" -access all`
- `security login role create -vserver SVM_Name -role SVM_Role_Name -cmddirname "network interface" -access readonly`
- `security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "snapmirror policy add-rule" -access all`
- `security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "snapmirror policy modify-rule" -access all`
- `security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "snapmirror policy remove-rule" -access all`
- `security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "snapmirror policy show" -access all`
- `security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "snapmirror restore" -access all`
- `security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "snapmirror show" -access all`
- `security login role create -vserver SVM_Name -role SVM_Role_Name -cmddirname "snapmirror show-history" -access all`
- `security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "snapmirror update" -access all`
- `security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "snapmirror update-ls-set" -access all`
- `security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "version" -access all`
- `security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname`

"volume clone create" -access all

- security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "volume clone show" -access all
- security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "volume clone split start" -access all
- security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "volume clone split stop" -access all
- security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "volume create" -access all
- security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "volume destroy" -access all
- security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "volume file clone create" -access all
- security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "volume file show-disk-usage" -access all
- security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "volume modify" -access all
- security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "volume offline" -access all
- security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "volume online" -access all
- security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "volume qtree create" -access all
- security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "volume qtree delete" -access all
- security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "volume qtree modify" -access all
- security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "volume qtree show" -access all
- security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "volume restrict" -access all
- security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "volume show" -access all
- security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "volume snapshot create" -access all
- security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "volume snapshot delete" -access all
- security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "volume snapshot modify" -access all
- security login role create -vserver SVM_Name -role SVM_Role_Name -cmddirname "volume snapshot modify-snaplock-expiry-time" -access all

- security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "volume snapshot rename" -access all
- security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "volume snapshot restore" -access all
- security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "volume snapshot restore-file" -access all
- security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "volume snapshot show" -access all
- security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "volume snapshot show-delta" -access all
- security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "volume unmount" -access all
- security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "vserver cifs share create" -access all
- security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "vserver cifs share delete" -access all
- security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "vserver cifs share show" -access all
- security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "vserver cifs show" -access all
- security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "vserver export-policy create" -access all
- security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "vserver export-policy delete" -access all
- security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "vserver export-policy rule create" -access all
- security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "vserver export-policy rule show" -access all
- security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "vserver export-policy show" -access all
- security login role create -vserver SVM_Name -role SVM_Role_Name -cmddirname "vserver iscsi connection show" -access all
- security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "vserver" -access readonly
- security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "vserver export-policy" -access all
- security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "vserver iscsi" -access all
- security login role create -vserver SVM_Name -role SVM_Role_Name -cmddirname "volume clone split status" -access all
- security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname

```
"volume managed-feature" -access all
```

- `security login role create -vserver SVM_Name -role SVM_Role_Name -cmddirname "nvme subsystem map" -access all`
- `security login role create -vserver SVM_Name -role SVM_Role_Name -cmddirname "nvme subsystem create" -access all`
- `security login role create -vserver SVM_Name -role SVM_Role_Name -cmddirname "nvme subsystem delete" -access all`
- `security login role create -vserver SVM_Name -role SVM_Role_Name -cmddirname "nvme subsystem modify" -access all`
- `security login role create -vserver SVM_Name -role SVM_Role_Name -cmddirname "nvme subsystem host" -access all`
- `security login role create -vserver SVM_Name -role SVM_Role_Name -cmddirname "nvme subsystem controller" -access all`
- `security login role create -vserver SVM_Name -role SVM_Role_Name -cmddirname "nvme subsystem show" -access all`
- `security login role create -vserver SVM_Name -role SVM_Role_Name -cmddirname "nvme namespace create" -access all`
- `security login role create -vserver SVM_Name -role SVM_Role_Name -cmddirname "nvme namespace delete" -access all`
- `security login role create -vserver SVM_Name -role SVM_Role_Name -cmddirname "nvme namespace modify" -access all`
- `security login role create -vserver SVM_Name -role SVM_Role_Name -cmddirname "nvme namespace show" -access all`

Create SVM roles for ASA r2 systems

There are several ONTAP CLI commands you must run to you create a role for a new SVM user in ASA r2 systems. This role is required if you configure SVMs in ASA r2 systems to use with SnapCenter and you do not want to use the vsadmin role.

Steps

1. On the storage system, create a role and assign all the permissions to the role.

   ```
   security login role create –vserver <svm_name\>- role <SVM_Role_Name\>
   -cmddirname <permission\>
   ```

   You should repeat this command for each permission.

2. Create a user and assign the role to that user.

   ```
   security login create -user <user_name\> -vserver <svm_name\> -application
   http -authmethod password -role <SVM_Role_Name\>
   ```

3. Unlock the user.

```
security login unlock -user <user_name\> -vserver <svm_name\>
```

ONTAP CLI commands for creating SVM roles and assigning permissions

There are several ONTAP CLI commands you should run to create SVM roles and assign permissions.

- `security login role create -vserver SVM_Name -role SVM_Role_Name -cmddirname "snapmirror list-destinations" -access all`
- `security login role create -vserver SVM_Name -role SVM_Role_Name -cmddirname "event generate-autosupport-log" -access all`
- `security login role create -vserver SVM_Name -role SVM_Role_Name -cmddirname "job history show" -access all`
- `security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "job show" -access all`
- `security login role create -vserver SVM_Name -role SVM_Role_Name -cmddirname "job stop" -access all`
- `security login role create -vserver SVM_Name -role SVM_Role_Name -cmddirname "lun" -access all`
- `security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "lun create" -access all`
- `security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "lun delete" -access all`
- `security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "lun igroup add" -access all`
- `security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "lun igroup create" -access all`
- `security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "lun igroup delete" -access all`
- `security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "lun igroup rename" -access all`
- `security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "lun igroup show" -access all`
- `security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "lun mapping add-reporting-nodes" -access all`
- `security login role create -vserver SVM_Name -role SVM_Role_Name -cmddirname "lun mapping create" -access all`
- `security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "lun mapping delete" -access all`
- `security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "lun mapping remove-reporting-nodes" -access all`
- `security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "lun mapping show" -access all`

- security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "lun modify" -access all
- security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "lun move-in-volume" -access all
- security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "lun offline" -access all
- security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "lun online" -access all
- security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "lun resize" -access all
- security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "lun serial" -access all
- security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "lun show" -access all
- security login role create -vserver SVM_Name -role SVM_Role_Name -cmddirname "network interface" -access readonly
- security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "snapmirror policy add-rule" -access all
- security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "snapmirror policy modify-rule" -access all
- security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "snapmirror policy remove-rule" -access all
- security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "snapmirror policy show" -access all
- security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "snapmirror restore" -access all
- security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "snapmirror show" -access all
- security login role create -vserver SVM_Name -role SVM_Role_Name -cmddirname "snapmirror show-history" -access all
- security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "snapmirror update" -access all
- security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "snapmirror update-ls-set" -access all
- security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "version" -access all
- security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "volume clone create" -access all
- security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "volume clone show" -access all
- security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname

"volume clone split start" -access all

- security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "volume clone split stop" -access all
- security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "volume create" -access all
- security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "volume destroy" -access all
- security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "volume file clone create" -access all
- security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "volume file show-disk-usage" -access all
- security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "volume modify" -access all
- security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "volume offline" -access all
- security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "volume online" -access all
- security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "volume qtree create" -access all
- security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "volume qtree delete" -access all
- security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "volume qtree modify" -access all
- security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "volume qtree show" -access all
- security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "volume restrict" -access all
- security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "volume show" -access all
- security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "volume snapshot create" -access all
- security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "volume snapshot delete" -access all
- security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "volume snapshot modify" -access all
- security login role create -vserver SVM_Name -role SVM_Role_Name -cmddirname "volume snapshot modify-snaplock-expiry-time" -access all
- security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "volume snapshot rename" -access all
- security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "volume snapshot restore" -access all

- security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "volume snapshot restore-file" -access all
- security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "volume snapshot show" -access all
- security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "volume snapshot show-delta" -access all
- security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "volume unmount" -access all
- security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "vserver cifs share create" -access all
- security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "vserver cifs share delete" -access all
- security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "vserver cifs share show" -access all
- security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "vserver cifs show" -access all
- security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "vserver export-policy create" -access all
- security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "vserver export-policy delete" -access all
- security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "vserver export-policy rule create" -access all
- security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "vserver export-policy rule show" -access all
- security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "vserver export-policy show" -access all
- security login role create -vserver SVM_Name -role SVM_Role_Name -cmddirname "vserver iscsi connection show" -access all
- security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "vserver" -access readonly
- security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "vserver export-policy" -access all
- security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "vserver iscsi" -access all
- security login role create -vserver SVM_Name -role SVM_Role_Name -cmddirname "volume clone split status" -access all
- security login role create -vserver SVM_name -role SVM_Role_Name -cmddirname "volume managed-feature" -access all
- security login role create -vserver SVM_Name -role SVM_Role_Name -cmddirname "nvme subsystem map" -access all
- security login role create -vserver SVM_Name -role SVM_Role_Name -cmddirname

"nvme subsystem create" -access all

- security login role create -vserver SVM_Name -role SVM_Role_Name -cmddirname "nvme subsystem delete" -access all
- security login role create -vserver SVM_Name -role SVM_Role_Name -cmddirname "nvme subsystem modify" -access all
- security login role create -vserver SVM_Name -role SVM_Role_Name -cmddirname "nvme subsystem host" -access all
- security login role create -vserver SVM_Name -role SVM_Role_Name -cmddirname "nvme subsystem controller" -access all
- security login role create -vserver SVM_Name -role SVM_Role_Name -cmddirname "nvme subsystem show" -access all
- security login role create -vserver SVM_Name -role SVM_Role_Name -cmddirname "nvme namespace create" -access all
- security login role create -vserver SVM_Name -role SVM_Role_Name -cmddirname "nvme namespace delete" -access all
- security login role create -vserver SVM_Name -role SVM_Role_Name -cmddirname "nvme namespace modify" -access all
- security login role create -vserver SVM_Name -role SVM_Role_Name -cmddirname "nvme namespace show" -access all
- security login role create -vserver SVM_Name -role SVM_Role_Name -cmddirname "storage-unit show" -access all
- security login role create -vserver SVM_Name -role SVM_Role_Name -cmddirname "consistency-group" -access all
- security login role create -vserver SVM_Name -role SVM_Role_Name -cmddirname "snapmirror protect" -access all
- security login role create -vserver SVM_Name -role SVM_Role_Name -cmddirname "volume delete" -access all
- security login create -user-or-group-name user_name -application http -authentication-method password -role SVM_Role_Name -vserver SVM_Name
- security login create -user-or-group-name user_name -application ssh -authentication-method password -role SVM_Role_Name -vserver SVM_Name

Create ONTAP cluster roles with minimum privileges

You should create an ONTAP cluster role with minimum privileges so that you do not have to use the ONTAP admin role to perform operations in SnapCenter. You can run several ONTAP CLI commands to create the ONTAP cluster role and assign minimum privileges.

Steps

1. On the storage system, create a role and assign all the permissions to the role.

   security login role create –vserver <cluster_name\>- role <role_name\>

```
-cmddirname <permission\>
```

You should repeat this command for each permission.

2. Create a user and assign the role to that user.

```
security login create -user <user_name\> -vserver <cluster_name\> -application
ontapi http -authmethod password -role <role_name\>
```

3. Unlock the user.

```
security login unlock -user <user_name\> -vserver <cluster_name\>
```

ONTAP CLI commands for creating cluster roles and assigning permissions

There are several ONTAP CLI commands you should run to create cluster roles and assign permissions.

- `security login role create -vserver Cluster_name or cluster_name -role Role_Name -cmddirname "metrocluster show" -access readonly`
- `security login role create -vserver Cluster_name or cluster_name -role Role_Name -cmddirname "cluster identity modify" -access all`
- `security login role create -vserver Cluster_name -role Role_Name -cmddirname "cluster identity show" -access all`
- `security login role create -vserver Cluster_name -role Role_Name -cmddirname "cluster modify" -access all`
- `security login role create -vserver Cluster_name -role Role_Name -cmddirname "cluster peer show" -access all`
- `security login role create -vserver Cluster_name -role Role_Name -cmddirname "cluster show" -access all`
- `security login role create -vserver Cluster_name -role Role_Name -cmddirname "event generate-autosupport-log" -access all`
- `security login role create -vserver Cluster_name -role Role_Name -cmddirname "job history show" -access all`
- `security login role create -vserver Cluster_name -role Role_Name -cmddirname "job show" -access all`
- `security login role create -vserver Cluster_name -role Role_Name -cmddirname "job stop" -access all`
- `security login role create -vserver Cluster_name -role Role_Name -cmddirname "lun" -access all`
- `security login role create -vserver Cluster_name -role Role_Name -cmddirname "lun create" -access all`
- `security login role create -vserver Cluster_name -role Role_Name -cmddirname "lun delete" -access all`
- `security login role create -vserver Cluster_name -role Role_Name -cmddirname "lun igroup add" -access all`

- `security login role create -vserver Cluster_name -role Role_Name -cmddirname "lun igroup create" -access all`
- `security login role create -vserver Cluster_name -role Role_Name -cmddirname "lun igroup delete" -access all`
- `security login role create -vserver Cluster_name -role Role_Name -cmddirname "lun igroup modify" -access all`
- `security login role create -vserver Cluster_name -role Role_Name -cmddirname "lun igroup rename" -access all`
- `security login role create -vserver Cluster_name -role Role_Name -cmddirname "lun igroup show" -access all`
- `security login role create -vserver Cluster_name -role Role_Name -cmddirname "lun mapping add-reporting-nodes" -access all`
- `security login role create -vserver Cluster_name -role Role_Name -cmddirname "lun mapping create" -access all`
- `security login role create -vserver Cluster_name -role Role_Name -cmddirname "lun mapping delete" -access all`
- `security login role create -vserver Cluster_name -role Role_Name -cmddirname "lun mapping remove-reporting-nodes" -access all`
- `security login role create -vserver Cluster_name -role Role_Name -cmddirname "lun mapping show" -access all`
- `security login role create -vserver Cluster_name -role Role_Name -cmddirname "lun modify" -access all`
- `security login role create -vserver Cluster_name -role Role_Name -cmddirname "lun move-in-volume" -access all`
- `security login role create -vserver Cluster_name -role Role_Name -cmddirname "lun offline" -access all`
- `security login role create -vserver Cluster_name -role Role_Name -cmddirname "lun online" -access all`
- `security login role create -vserver Cluster_name -role Role_Name -cmddirname "lun persistent-reservation clear" -access all`
- `security login role create -vserver Cluster_name -role Role_Name -cmddirname "lun resize" -access all`
- `security login role create -vserver Cluster_name -role Role_Name -cmddirname "lun serial" -access all`
- `security login role create -vserver Cluster_name -role Role_Name -cmddirname "lun show" -access all`
- `security login role create -vserver Cluster_name -role Role_Name -cmddirname "network interface create" -access readonly`
- `security login role create -vserver Cluster_name -role Role_Name -cmddirname "network interface delete" -access readonly`
- `security login role create -vserver Cluster_name -role Role_Name -cmddirname`

"network interface modify" -access readonly

- security login role create -vserver Cluster_name -role Role_Name -cmddirname "network interface show" -access readonly
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "nvme subsystem map" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "nvme subsystem create" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "nvme subsystem delete" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "nvme subsystem modify" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "nvme subsystem host" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "nvme subsystem controller" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "nvme subsystem show" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "nvme namespace create" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "nvme namespace delete" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "nvme namespace modify" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "nvme namespace show" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "security login" -access readonly
- security login role create -role Role_Name -cmddirname "snapmirror create" -vserver Cluster_name -access all
- security login role create -role Role_Name -cmddirname "snapmirror list-destinations" -vserver Cluster_name -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "snapmirror policy add-rule" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "snapmirror policy create" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "snapmirror policy delete" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "snapmirror policy modify" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "snapmirror policy modify-rule" -access all

- security login role create -vserver Cluster_name -role Role_Name -cmddirname "snapmirror policy remove-rule" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "snapmirror policy show" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "snapmirror restore" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "snapmirror show" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "snapmirror show-history" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "snapmirror update" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "snapmirror update-ls-set" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "system license add" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "system license clean-up" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "system license delete" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "system license show" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "system license status show" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "system node modify" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "system node show" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "system status show" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "version" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "volume clone create" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "volume clone show" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "volume clone split start" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "volume clone split stop" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname

"volume create" -access all

- security login role create -vserver Cluster_name -role Role_Name -cmddirname "volume destroy" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "volume file clone create" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "volume file show-disk-usage" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "volume modify" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "volume snapshot modify-snaplock-expiry-time" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "volume offline" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "volume online" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "volume qtree create" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "volume qtree delete" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "volume qtree modify" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "volume qtree show" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "volume restrict" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "volume show" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "volume snapshot create" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "volume snapshot delete" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "volume snapshot modify" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "volume snapshot promote" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "volume snapshot rename" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "volume snapshot restore" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "volume snapshot restore-file" -access all

- security login role create -vserver Cluster_name -role Role_Name -cmddirname "volume snapshot show" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "volume snapshot show-delta" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "volume unmount" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "vserver" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "vserver cifs create" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "vserver cifs delete" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "vserver cifs modify" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "vserver cifs share modify" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "vserver cifs share create" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "vserver cifs share delete" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "vserver cifs share modify" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "vserver cifs share show" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "vserver cifs show" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "vserver create" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "vserver export-policy create" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "vserver export-policy delete" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "vserver export-policy rule create" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "vserver export-policy rule delete" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "vserver export-policy rule modify" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "vserver export-policy rule show" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname

"vserver export-policy show" -access all

- security login role create -vserver Cluster_name -role Role_Name -cmddirname "vserver iscsi connection show" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "vserver modify" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "vserver show" -access all

Create ONTAP cluster roles for ASA r2 systems

You should create an ONTAP cluster role with minimum privileges so that you do not have to use the ONTAP admin role to perform operations in SnapCenter. You can run several ONTAP CLI commands to create the ONTAP cluster role and assign minimum privileges.

Steps

1. On the storage system, create a role and assign all the permissions to the role.

   ```
   security login role create –vserver <cluster_name\>- role <role_name\>
   -cmddirname <permission\>
   ```

   You should repeat this command for each permission.

2. Create a user and assign the role to that user.

   ```
   security login create -user <user_name\> -vserver <cluster_name\> -application
   http -authmethod password -role <role_name\>
   ```

3. Unlock the user.

   ```
   security login unlock -user <user_name\> -vserver <cluster_name\>
   ```

ONTAP CLI commands for creating cluster roles and assigning permissions

There are several ONTAP CLI commands you should run to create cluster roles and assign permissions.

- security login role create -vserver Cluster_name or cluster_name -role Role_Name -cmddirname "metrocluster show" -access readonly
- security login role create -vserver Cluster_name or cluster_name -role Role_Name -cmddirname "cluster identity modify" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "cluster identity show" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "cluster modify" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "cluster peer show" -access all

- security login role create -vserver Cluster_name -role Role_Name -cmddirname "cluster show" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "event generate-autosupport-log" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "job history show" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "job show" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "job stop" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "lun" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "lun create" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "lun delete" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "lun igroup add" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "lun igroup create" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "lun igroup delete" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "lun igroup modify" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "lun igroup rename" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "lun igroup show" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "lun mapping add-reporting-nodes" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "lun mapping create" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "lun mapping delete" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "lun mapping remove-reporting-nodes" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "lun mapping show" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "lun modify" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname

"lun move-in-volume" -access all

- security login role create -vserver Cluster_name -role Role_Name -cmddirname "lun offline" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "lun online" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "lun persistent-reservation clear" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "lun resize" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "lun serial" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "lun show" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "network interface create" -access readonly
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "network interface delete" -access readonly
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "network interface modify" -access readonly
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "network interface show" -access readonly
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "nvme subsystem map" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "nvme subsystem create" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "nvme subsystem delete" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "nvme subsystem modify" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "nvme subsystem host" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "nvme subsystem controller" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "nvme subsystem show" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "nvme namespace create" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "nvme namespace delete" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "nvme namespace modify" -access all

- security login role create -vserver Cluster_name -role Role_Name -cmddirname "nvme namespace show" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "security login" -access readonly
- security login role create -role Role_Name -cmddirname "snapmirror create" -vserver Cluster_name -access all
- security login role create -role Role_Name -cmddirname "snapmirror list-destinations" -vserver Cluster_name -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "snapmirror policy add-rule" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "snapmirror policy create" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "snapmirror policy delete" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "snapmirror policy modify" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "snapmirror policy modify-rule" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "snapmirror policy remove-rule" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "snapmirror policy show" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "snapmirror restore" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "snapmirror show" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "snapmirror show-history" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "snapmirror update" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "snapmirror update-ls-set" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "system license add" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "system license clean-up" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "system license delete" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "system license show" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname

"system license status show" -access all

- security login role create -vserver Cluster_name -role Role_Name -cmddirname "system node modify" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "system node show" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "system status show" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "version" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "volume clone create" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "volume clone show" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "volume clone split start" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "volume clone split stop" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "volume create" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "volume destroy" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "volume file clone create" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "volume file show-disk-usage" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "volume modify" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "volume snapshot modify-snaplock-expiry-time" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "volume offline" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "volume online" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "volume qtree create" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "volume qtree delete" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "volume qtree modify" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "volume qtree show" -access all

- `security login role create -vserver Cluster_name -role Role_Name -cmddirname "volume restrict" -access all`
- `security login role create -vserver Cluster_name -role Role_Name -cmddirname "volume show" -access all`
- `security login role create -vserver Cluster_name -role Role_Name -cmddirname "volume snapshot create" -access all`
- `security login role create -vserver Cluster_name -role Role_Name -cmddirname "volume snapshot delete" -access all`
- `security login role create -vserver Cluster_name -role Role_Name -cmddirname "volume snapshot modify" -access all`
- `security login role create -vserver Cluster_name -role Role_Name -cmddirname "volume snapshot promote" -access all`
- `security login role create -vserver Cluster_name -role Role_Name -cmddirname "volume snapshot rename" -access all`
- `security login role create -vserver Cluster_name -role Role_Name -cmddirname "volume snapshot restore" -access all`
- `security login role create -vserver Cluster_name -role Role_Name -cmddirname "volume snapshot restore-file" -access all`
- `security login role create -vserver Cluster_name -role Role_Name -cmddirname "volume snapshot show" -access all`
- `security login role create -vserver Cluster_name -role Role_Name -cmddirname "volume snapshot show-delta" -access all`
- `security login role create -vserver Cluster_name -role Role_Name -cmddirname "volume unmount" -access all`
- `security login role create -vserver Cluster_name -role Role_Name -cmddirname "vserver" -access all`
- `security login role create -vserver Cluster_name -role Role_Name -cmddirname "vserver cifs create" -access all`
- `security login role create -vserver Cluster_name -role Role_Name -cmddirname "vserver cifs delete" -access all`
- `security login role create -vserver Cluster_name -role Role_Name -cmddirname "vserver cifs modify" -access all`
- `security login role create -vserver Cluster_name -role Role_Name -cmddirname "vserver cifs share modify" -access all`
- `security login role create -vserver Cluster_name -role Role_Name -cmddirname "vserver cifs share create" -access all`
- `security login role create -vserver Cluster_name -role Role_Name -cmddirname "vserver cifs share delete" -access all`
- `security login role create -vserver Cluster_name -role Role_Name -cmddirname "vserver cifs share modify" -access all`
- `security login role create -vserver Cluster_name -role Role_Name -cmddirname`

"vserver cifs share show" -access all

- security login role create -vserver Cluster_name -role Role_Name -cmddirname "vserver cifs show" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "vserver create" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "vserver export-policy create" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "vserver export-policy delete" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "vserver export-policy rule create" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "vserver export-policy rule delete" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "vserver export-policy rule modify" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "vserver export-policy rule show" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "vserver export-policy show" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "vserver iscsi connection show" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "vserver modify" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "vserver show" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "storage-unit show" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "consistency-group" show" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "snapmirror protect" show" -access all
- security login role create -vserver Cluster_name -role Role_Name -cmddirname "volume delete" show" -access all

Add a user or group and assign role and assets

To configure role-based access control for SnapCenter users, you can add users or groups and assign role. The role determines the options that SnapCenter users can access.

Before you begin

- You must have logged in as the "SnapCenterAdmin" role.

- You must have created the user or group accounts in Active Directory in the operating system or database. You cannot use SnapCenter to create these accounts.

  You can include only the following special characters in user names and group names: space ( ), hyphen (-), underscore (_), and colon (:).

- SnapCenter includes several predefined roles.

  You can either assign these roles to the user or create new roles.

- AD Users and AD Groups that are added to SnapCenter RBAC must have the READ permission on the Users Container and the Computers Container in the Active Directory.
- After you assign a role to a user or group that contains the appropriate permissions, you must assign the user access to SnapCenter assets, such as hosts and storage connections.

  This enables users to perform the actions for which they have permissions on the assets that are assigned to them.

- You should assign a role to the user or group at some point to take advantage of RBAC permissions and efficiencies.
- You can assign assets like host, resource groups, policy, storage connection, plug-in, and credential to the user while creating the user or group.
- The minimum assets that you should assign an user to perform certain operations are as follows:

Operation	Assets assignment
Protect resources	host, policy
Backup	host, resource group, policy
Restore	host, resource group
Clone	host, resource group, policy
Clone lifecycle	host
Create a Resource Group	host

- When a new node is added to a Windows cluster or a DAG (Exchange Server Database Availability Group) asset and if this new node is assigned to a user, you must reassign the asset to the user or group to include the new node to the user or group.

  You should reassign the RBAC user or group to the cluster or DAG to include the new node to the RBAC user or group. For example, you have a two-node cluster and you have assigned an RBAC user or group to the cluster. When you add another node to the cluster, you should reassign the RBAC user or group to the cluster to include the new node for the RBAC user or group.

- If you are planning to replicate Snapshots, you must assign the storage connection for both the source and destination volume to the user performing the operation.

You should add assets before assigning access to the users.

If you are using the SnapCenter Plug-in for VMware vSphere functions to protect VMs, VMDKs, or datastores, you should use the VMware vSphere GUI to add a vCenter user to a SnapCenter Plug-in for VMware vSphere role. For information about VMware vSphere roles, see Predefined roles packaged with SnapCenter Plug-in for VMware vSphere.

Steps

1. In the left navigation pane, click **Settings**.
2. In the Settings page, click **Users and Access** > .
3. In the Add Users/Groups from Active Directory or Workgroup page:

For this field…	Do this…
Access Type	Select either Domain or workgroup For Domain authentication type, you should specify the domain name of the user or group to which you want to add the user to a role. By default, it is pre-populated with the logged in domain name. You must register the untrusted domain in the **Settings** > **Global Settings** > **Domain Settings** page.
Type	Select either User or Group SnapCenter supports only security group and not the distribution group.
User Name	a. Type the partial user name, and then click **Add**. The user name is case-sensitive. b. Select the user name from the search list. When you add users from a different domain or an untrusted domain, you should type the user name fully because there is no search list for cross domain users. Repeat this step to add additional users or groups to the selected role.
Roles	Select the role to which you want to add the user.

4. Click **Assign**, and then in the Assign Assets page:
   a. Select the type of asset from the **Asset** drop-down list.
   b. In the Asset table, select the asset.

      The assets are listed only if the user has added the assets to SnapCenter.

   c. Repeat this procedure for all of the required assets.
   d. Click **Save**.
5. Click **Submit**.

After adding users or groups and assigning roles, refresh the resources list.

Configure audit log settings

Audit logs are generated for each and every activity of the SnapCenter Server. By default, audit logs are secured in the default installed location *C:\Program Files\NetApp\SnapCenter WebApp\audit*\.

Audit logs are secured by means of generating digitally signed digest for each and every audit events to protect it from the unauthorized modification. The generated digest's are maintained in the separate audit checksum file and it under goes periodic integrity checks to ensure the integrity of the content.

You should have logged in as the "SnapCenterAdmin" role.

About this task

- Alerts are sent in the following scenarios:
  - Audit log integrity check schedule or Syslog server is enabled or disabled
  - Audit log integrity check, audit log, or Syslog server log failure
  - Low disk space
- Email is sent only when integrity check fails.
- You should modify both audit log directory and audit checksum log directory paths together. You cannot modify only one of them.
- When audit log directory and audit checksum log directory paths are modified, the integrity check cannot be performed on audit logs present in the earlier location.
- Audit log directory and Audit checksum log directory paths should be on the local drive of SnapCenter Server.

  Shared or network mounted drives are not supported.

- If UDP protocol is used in the Syslog server settings, errors due to port is down or unavailable cannot be captured as either an error or an alert in SnapCenter.
- You can use Set-SmAuditSettings and Get-SmAuditSettings commands to configure the audit logs.

  The information regarding the parameters that can be used with the cmdlet and their descriptions can be obtained by running Get-Help command_name. Alternatively, you can also refer the SnapCenter Software Cmdlet Reference Guide.

Steps

1. In the **Settings** page, navigate to **Settings** > **Global Settings** > **Audit log Settings**.
2. In the Audit log section, enter the details.
3. Enter the **Audit log directory** and **Audit checksum log directory**
   a. Enter the Maximum file size
   b. Enter the Maximum log files
   c. Enter the percentage of disk space usage to send an alert
4. (Optional) Enable **Log UTC time**.
5. (Optional) Enable **Audit Log Integrity Check Schedule** and click **Start Integrity Check** for on demand integrity check.

   You can also run **Start-SmAuditIntegrityCheck** command to start on demand integrity check.

6. (Optional) Enable Forwarded audit logs to remote syslog server and enter the Syslog Server details.

   You should import the certificate from the Syslog server into the 'Trusted Root' for TLS 1.2 protocol.

   a. Enter Syslog Server Host
   b. Enter Syslog Server Port
   c. Enter Syslog Server Protocol
   d. Enter RFC Format
7. Click **Save**.
8. You can see audit integrity checks and disk space checks by clicking **Monitor** > **Jobs**.

Configure secured MySQL connections with SnapCenter Server

You can generate Secure Sockets Layer (SSL) certificates and key files if you want to secure the communication between SnapCenter Server and MySQL Server in standalone configurations or Network Load Balancing (NLB) configurations.

Configure secured MySQL connections for standalone SnapCenter Server configurations

You can generate Secure Sockets Layer (SSL) certificates and key files, if you want to secure the communication between SnapCenter Server and MySQL Server. You must configure the certificates and key files in the MySQL Server and SnapCenter Server.

The following certificates are generated:

- CA certificate
- Server public certificate and private key file
- Client public certificate and private key file

Steps

1. Set up the SSL certificates and key files for MySQL servers and clients on Windows by using the openssl command.

For information, see MySQL Version 5.7: Creating SSL Certificates and Keys Using openssl

> The common name value that is used for the server certificate, client certificate, and key files must each differ from the common name value that is used for the CA certificate. If the common name values are the same, the certificate and key files fail for servers that are compiled by using OpenSSL.

Best Practice: You should use the server fully qualified domain name (FQDN) as the common name for the server certificate.

2. Copy the SSL certificates and key files to the MySQL Data folder.

   The default MySQL Data folder path is `C:\ProgramData\NetApp\SnapCenter\MySQL Data\Data\`.

3. Update the CA certificate, server public certificate, client public certificate, server private key, and client private key paths in the MySQL server configuration file (my.ini).

   The default MySQL server configuration file (my.ini) path is `C:\ProgramData\NetApp\SnapCenter\MySQL Data\my.ini`.

   > You must specify the CA certificate, server public certificate, and server private key paths in the [mysqld] section of the MySQL server configuration file (my.ini).

   You must specify the CA certificate, client public certificate, and client private key paths in the [client] section of the MySQL server configuration file (my.ini).

   The following example shows the certificates and key files copied to the [mysqld] section of the my.ini file in the default folder `C:/ProgramData/NetApp/SnapCenter/MySQL Data/Data`.

```
ssl-ca="C:/ProgramData/NetApp/SnapCenter/MySQL Data/Data/ca.pem"
```

```
ssl-cert="C:/ProgramData/NetApp/SnapCenter/MySQL Data/Data/server-cert.pem"
```

```
ssl-key="C:/ProgramData/NetApp/SnapCenter/MySQL Data/Data/server-key.pem"
```

   The following example shows the paths updated in the [client] section of the my.ini file.

```
ssl-ca="C:/ProgramData/NetApp/SnapCenter/MySQL Data/Data/ca.pem"
```

```
ssl-cert="C:/ProgramData/NetApp/SnapCenter/MySQL Data/Data/client-cert.pem"
```

```
ssl-key="C:/ProgramData/NetApp/SnapCenter/MySQL Data/Data/client-
key.pem"
```

4. Stop the SnapCenter Server web application in the Internet Information Server (IIS).
5. Restart the MySQL service.
6. Update the value of the MySQLProtocol key in the SnapManager.Web.UI.dll.config file.

   The following example shows the value of the MySQLProtocol key updated in the SnapManager.Web.UI.dll.config file.

```
<add key="MySQLProtocol" value="SSL" />
```

7. Update the SnapManager.Web.UI.dll.config file with the paths that were provided in the [client] section of the my.ini file.

   The following example shows the paths updated in the [client] section of the my.ini file.

```
<add key="ssl-client-cert" value="C:/ProgramData/NetApp/SnapCenter/MySQL
Data/Data/client-cert.pem" />
```

```
<add key="ssl-client-key" value="C:/ProgramData/NetApp/SnapCenter/MySQL
Data/Data/client-key.pem" />
```

```
<add key="ssl-ca" value="C:/ProgramData/NetApp/SnapCenter/MySQL
Data/Data/ca.pem" />
```

8. Start the SnapCenter Server web application in the IIS.

Configure secured MySQL connections for HA configurations

You can generate Secure Sockets Layer (SSL) certificates and key files for both the High Availability (HA) nodes if you want to secure the communication between SnapCenter Server and MySQL servers. You must configure the certificates and key files in the MySQL servers and on the HA nodes.

The following certificates are generated:

- CA certificate

  A CA certificate is generated on one of the HA nodes, and this CA certificate is copied to the other HA node.

- Server public certificate and server private key files for both the HA nodes
- Client public certificate and client private key files for both the HA nodes

Steps

1. For the first HA node, set up the SSL certificates and key files for MySQL servers and clients on Windows by using the openssl command.

   For information, see MySQL Version 5.7: Creating SSL Certificates and Keys Using openssl

   > The common name value that is used for the server certificate, client certificate, and key files must each differ from the common name value that is used for the CA certificate. If the common name values are the same, the certificate and key files fail for servers that are compiled by using OpenSSL.

   Best Practice: You should use the server fully qualified domain name (FQDN) as the common name for the server certificate.

2. Copy the SSL certificates and key files to the MySQL Data folder.

   The default MySQL Data folder path is C:\ProgramData\NetApp\SnapCenter\MySQL Data\Data\.

3. Update the CA certificate, server public certificate, client public certificate, server private key, and client private key paths in the MySQL server configuration file (my.ini).

   The default MySQL server configuration file (my.ini) path is C:\ProgramData\NetApp\SnapCenter\MySQL Data\my.ini.

   > You must specify CA certificate, server public certificate, and server private key paths in the [mysqld] section of the MySQL server configuration file (my.ini).

   You must specify CA certificate, client public certificate, and client private key paths in the [client] section of the MySQL server configuration file (my.ini).

   The following example shows the certificates and key files copied to the [mysqld] section of the my.ini file in the default folder C:/ProgramData/NetApp/SnapCenter/MySQL Data/Data.

   ```
   ssl-ca="C:/ProgramData/NetApp/SnapCenter/MySQL Data/Data/ca.pem"
   ```

   ```
   ssl-cert="C:/ProgramData/NetApp/SnapCenter/MySQL Data/Data/server-cert.pem"
   ```

   ```
   ssl-key="C:/ProgramData/NetApp/SnapCenter/MySQL Data/Data/server-key.pem"
   ```

   The following example shows the paths updated in the [client] section of the my.ini file.

   ```
   ssl-ca="C:/ProgramData/NetApp/SnapCenter/MySQL Data/Data/ca.pem"
   ```

```
ssl-cert="C:/ProgramData/NetApp/SnapCenter/MySQL Data/Data/client-
cert.pem"
```

```
ssl-key="C:/ProgramData/NetApp/SnapCenter/MySQL Data/Data/client-
key.pem"
```

4. For the second HA node, copy the CA certificate and generate server public certificate, server private key files, client public certificate, and client private key files. perform the following steps:
   a. Copy the CA certificate generated on the first HA node to the MySQL Data folder of the second NLB node.

      The default MySQL Data folder path is C:\ProgramData\NetApp\SnapCenter\MySQL Data\Data\.

      > You must not create a CA certificate again. You should create only the server public certificate, client public certificate, server private key file, and client private key file.

   b. For the first HA node, set up the SSL certificates and key files for MySQL servers and clients on Windows by using the openssl command.

      MySQL Version 5.7: Creating SSL Certificates and Keys Using openssl

      > The common name value that is used for the server certificate, client certificate, and key files must each differ from the common name value that is used for the CA certificate. If the common name values are the same, the certificate and key files fail for servers that are compiled by using OpenSSL.

      It is recommended to use the server FQDN as the common name for the server certificate.

   c. Copy the SSL certificates and key files to the MySQL Data folder.
   d. Update the CA certificate, server public certificate, client public certificate, server private key, and client private key paths in the MySQL server configuration file (my.ini).

      > You must specify the CA certificate, server public certificate, and server private key paths in the [mysqld] section of the MySQL server configuration file (my.ini).

      You must specify the CA certificate, client public certificate, and client private key paths in the [client] section of the MySQL server configuration file (my.ini).

      The following example shows the certificates and key files copied to the [mysqld] section of the my.ini file in the default folder C:/ProgramData/NetApp/SnapCenter/MySQL Data/Data.

      ```
      ssl-ca="C:/ProgramData/NetApp/SnapCenter/MySQL Data/Data/ca.pem"
      ```

```
ssl-cert="C:/ProgramData/NetApp/SnapCenter/MySQL Data/Data/server-cert.pem"
```

```
ssl-key="C:/ProgramData/NetApp/SnapCenter/MySQL Data/Data/server-key.pem"
```

The following example shows the paths updated in the [client] section of the my.ini file.

```
ssl-ca="C:/ProgramData/NetApp/SnapCenter/MySQL Data/Data/ca.pem"
```

```
ssl-cert="C:/ProgramData/NetApp/SnapCenter/MySQL Data/Data/server-cert.pem"
```

```
ssl-key="C:/ProgramData/NetApp/SnapCenter/MySQL Data/Data/server-key.pem"
```

5. Stop the SnapCenter Server web application in the Internet Information Server (IIS) on both the HA nodes.
6. Restart the MySQL service on both the HA nodes.
7. Update the value of the MySQLProtocol key in the SnapManager.Web.UI.dll.config file for both the HA nodes.

   The following example shows the value of MySQLProtocol key updated in the SnapManager.Web.UI.dll.config file.

   ```
   <add key="MySQLProtocol" value="SSL" />
   ```

8. Update the SnapManager.Web.UI.dll.config file with the paths that you specified in the [client] section of the my.ini file for both the HA nodes.

   The following example shows the paths updated in the [client] section of the my.ini files.

   ```
   <add key="ssl-client-cert" value="C:/ProgramData/NetApp/SnapCenter/MySQL Data/Data/client-cert.pem" />
   ```

   ```
   <add key="ssl-client-key" value="C:/ProgramData/NetApp/SnapCenter/MySQL Data/Data/client-key.pem" />
   ```

```
<add key="ssl-ca" value="C:/ProgramData/NetApp/SnapCenter/MySQL
Data/Data/ca.pem" />
```

9. Start the SnapCenter Server web application in the IIS on both the HA nodes.
10. Use the Set-SmRepositoryConfig -RebuildSlave -Force PowerShell cmdlet with the -Force option on one of the HA nodes to establish secured MySQL replication on both the HA nodes.

    Even if the replication status is healthy, the -Force option allows you to rebuild the slave repository.

Configure Certificate-based authentication

Certificate-based authentication enhances security by verifying the identity of both the SnapCenter Server and plug-in hosts, ensuring secure and encrypted communication.

Enable Certificate-based authentication

To enable certificate-based authentication for SnapCenter Server and the Windows plug-in hosts, run the following PowerShell cmdlet. For the Linux plug-in hosts, the certificate-based authentication will be enabled when you enable the two-way SSL.

- To enable client certificate-based authentication:

  ```
  Set-SmConfigSettings -Agent –configSettings
  @{"EnableClientCertificateAuthentication"="true"} -HostName [hostname]
  ```

- To disable client certificate-based authentication:

  ```
  Set-SmConfigSettings -Agent –configSettings
  @{"EnableClientCertificateAuthentication"="false"} -HostName [hostname]`
  ```

Export Certificate Authority (CA) certificates from SnapCenter Server

You should export the CA certificates from the SnapCenter Server to the plug-in hosts using the Microsoft management console (MMC).

Before you begin

You should have configured the two-way SSL.

Steps

1. Go to the Microsoft management console (MMC), and then click **File** > **Add/Remove Snapin**.
2. In the Add or Remove Snap-ins window, select **Certificates** and then click **Add**.
3. In the Certificates Snap-in window, select the **Computer Account** option, and then click **Finish**.
4. Click **Console Root** > **Certificates - Local Computer** > **Personal** > **Certificates**.
5. Right-click on the procured CA certificate, which is used for SnapCenter Server and then select **All Tasks** > **Export** to start the export wizard.
6. Perform the following actions in the wizard.

For this option…	Do the following…
Export Private Key	Select **No, do not export the private key**, and then click **Next**.
Export File Format	Click **Next**.
File Name	Click **Browse** and specify the file path to save the certificate, and click **Next**.
Completing the Certificate Export Wizard	Review the summary, and then click **Finish** to start the export.

Certificate based authentication is not supported for SnapCenter HA configurations and SnapCenter Plug-in for VMware vSphere.

Import CA certificate to the Windows plug-in hosts

To use the exported SnapCenter Server CA certificate, you should import the related certificate to the SnapCenter Windows plug-in hosts using the Microsoft management console (MMC).

Steps

1. Go to the Microsoft management console (MMC), and then click **File** > **Add/Remove Snapin**.
2. In the Add or Remove Snap-ins window, select **Certificates** and then click **Add**.
3. In the Certificates Snap-in window, select the **Computer Account** option, and then click **Finish**.
4. Click **Console Root** > **Certificates - Local Computer** > **Personal** > **Certificates**.
5. Right-click on the folder “Personal”, and then select **All Tasks** > **Import** to start the import wizard.
6. Perform the following actions in the wizard.

For this option…	Do the following…
Store Location	Click **Next**.
File to Import	Select the SnapCenter Server certificate that ends with .cer extension.
Certificate Store	Click **Next**.
Completing the Certificate Export Wizard	Review the summary, and then click **Finish** to start the import.

Import CA Certificate to the UNIX plug-in hosts

You should import the CA certificate to the UNIX plug-in hosts.

About this task

- You can manage the password for SPL keystore, and the alias of the CA signed key pair in use.
- The password for SPL keystore and for all the associated alias password of the private key should be same.

Steps

1. You can retrieve SPL keystore default password from SPL property file. It is the value corresponding to the key `SPL_KEYSTORE_PASS`.
2. Change the keystore password:
   `$ keytool -storepasswd -keystore keystore.jks`
3. Change the password for all aliases of private key entries in the keystore to the same password used for the keystore:
   `$ keytool -keypasswd -alias "<alias_name>" -keystore keystore.jks`
4. Update the same for the key SPL_KEYSTORE_PASS in `spl.properties` file.
5. Restart the service after changing the password.

Configure root or intermediate certificates to SPL trust-store

You should configure the root or intermediate certificates to SPL trust-store. You should add the root CA certificate and then the intermediate CA certificates.

Steps

1. Navigate to the folder containing the SPL keystore: `/var/opt/snapcenter/spl/etc`.
2. Locate the file `keystore.jks`.
3. List the added certificates in the keystore:
   `$ keytool -list -v -keystore keystore.jks`
4. Add a root or intermediate certificate:
   `$ keytool -import -trustcacerts -alias <AliasNameForCerticateToBeImported> -file /<CertificatePath> -keystore keystore.jks`
5. Restart the service after configuring the root or intermediate certificates to SPL trust-store.

Configure CA signed key pair to SPL trust-store

You should configure the CA signed key pair to SPL trust-store.

Steps

1. Navigate to the folder containing the SPL's keystore `/var/opt/snapcenter/spl/etc`.
2. Locate the file `keystore.jks`.
3. List the added certificates in the keystore:
   `$ keytool -list -v -keystore keystore.jks`
4. Add the CA certificate having both private and public key.
   `$ keytool -importkeystore -srckeystore <CertificatePathToImport> -srcstoretype pkcs12 -destkeystore keystore.jks -deststoretype JKS`
5. List the added certificates in the keystore.

```
$ keytool -list -v -keystore keystore.jks
```

6. Verify that the keystore contains the alias corresponding to the new CA certificate, which was added to the keystore.
7. Change the added private key password for CA certificate to the keystore password.

   Default SPL keystore password is the value of the key SPL_KEYSTORE_PASS in `spl.properties` file.

   ```
   $ keytool -keypasswd -alias "<aliasNameOfAddedCertInKeystore>" -keystore
   keystore.jks`
   ```

8. If the alias name in the CA certificate is long and contains space or special characters ("*",","), change the alias name to a simple name:
   ```
   $ keytool -changealias -alias "<OrignalAliasName>" -destalias "<NewAliasName>"
   -keystore keystore.jks`
   ```
9. Configure the alias name from the keystore located in `spl.properties` file.
   Update this value against the key SPL_CERTIFICATE_ALIAS.
10. Restart the service after configuring the CA signed key pair to SPL trust-store.

Export SnapCenter certificates

You should export the SnapCenter certificates in .pfx format.

Steps

1. Go to the Microsoft management console (MMC), and then click **File** > **Add/Remove Snap-in**.
2. In the Add or Remove Snap-ins window, select **Certificates** and then click **Add**.
3. In the Certificates snap-in window, select the **My user account** option, and then click **Finish**.
4. Click **Console Root** > **Certificates - Current User** > **Trusted Root Certification Authorities** > **Certificates**.
5. Right-click the certificate that has the SnapCenter Friendly Name, and then select **All Tasks** > **Export** to start the export wizard.
6. Complete the wizard, as follows:

In this wizard window…	Do the following…
Export Private Key	Select the option **Yes, export the private key**, and then click **Next**.
Export File Format	Make no changes; click **Next**.
Security	Specify the new password to be used for the exported certificate, and then click **Next**.
File to Export	Specify a file name for the exported certificate (you must use .pfx), and then click **Next**.

In this wizard window…	Do the following…
Completing the Certificate Export Wizard	Review the summary, and then click **Finish** to start the export.

Configure CA Certificate for Windows host

Generate CA Certificate CSR file

You can generate a Certificate Signing Request (CSR) and import the certificate that can be obtained from a Certificate Authority (CA) using the generated CSR. The certificate will have a private key associated with it.

CSR is a block of encoded text that is given to an authorized certificate vendor to procure the signed CA certificate.

CA Certificate RSA key length must be minimum 3072 bits.

For information to generate a CSR, see How to generate CA Certificate CSR file.

If you own the CA certificate for your domain (*.domain.company.com) or your system (machine1.domain.company.com), you can skip generating the CA Certificate CSR file. You can deploy the existing CA certificate with SnapCenter.

For cluster configurations, the cluster name (virtual cluster FQDN), and the respective host names should be mentioned in the CA certificate. The certificate can be updated by filling the Subject Alternative Name (SAN) field before procuring the certificate. For a wild card certificate (*.domain.company.com), the certificate will contain all the hostnames of the domain implicitly.

Import CA certificates

You must import the CA certificates to the SnapCenter Server and the Windows host plug-ins using the Microsoft management console (MMC).

Steps

1. Go to the Microsoft management console (MMC), and then click **File** > **Add/Remove Snapin**.
2. In the Add or Remove Snap-ins window, select **Certificates** and then click **Add**.
3. In the Certificates snap-in window, select the **Computer account** option, and then click **Finish**.
4. Click **Console Root** > **Certificates – Local Computer** > **Trusted Root Certification Authorities** > **Certificates**.
5. Right-click on the folder “Trusted Root Certification Authorities”, and then select **All Tasks** > **Import** to start the import wizard.
6. Complete the wizard, as follows:

In this wizard window…	Do the following…
Import Private Key	Select the option **Yes**, import the private key, and then click **Next**.
Import File Format	Make no changes; click **Next**.
Security	Specify the new password to be used for the exported certificate, and then click **Next**.
Completing the Certificate Import Wizard	Review the summary, and then click **Finish** to start the import.

Importing certificate should be bundled with the private key (supported formats are: *.pfx, *.p12, and *.p7b).

7. Repeat Step 5 for the “Personal” folder.

Get the CA certificate thumbprint

A certificate thumbprint is a hexadecimal string that identifies a certificate. A thumbprint is calculated from the content of the certificate using a thumbprint algorithm.

Steps

1. Perform the following on the GUI:
   a. Double-click the certificate.
   b. In the Certificate dialog box, click the **Details** tab.
   c. Scroll through the list of fields and click **Thumbprint**.
   d. Copy the hexadecimal characters from the box.
   e. Remove the spaces between the hexadecimal numbers.

      For example, if the thumbprint is: "a9 09 50 2d d8 2a e4 14 33 e6 f8 38 86 b0 0d 42 77 a3 2a 7b", after removing the spaces, it will be: "a909502dd82ae41433e6f83886b00d4277a32a7b".

2. Perform the following from PowerShell:
   a. Run the following command to list the thumbprint of the installed certificate and identify the recently installed certificate by the subject name.

      Get-ChildItem -Path Cert:\LocalMachine\My

   b. Copy the thumbprint.

Configure CA certificate with Windows host plug-in services

You should configure the CA certificate with Windows host plug-in services to activate the installed digital certificate.

Perform the following steps on the SnapCenter Server and all the plug-in hosts where CA certificates are

already deployed.

Steps

1. Remove the existing certificate binding with SMCore default port 8145, by running the following command:

   ```
   > netsh http delete sslcert ipport=0.0.0.0:_<SMCore Port>
   ```

   For example:

   ```
   > netsh http delete sslcert ipport=0.0.0.0:8145
   ```

2. Bind the newly installed certificate with the Windows host plug-in services, by running the following commands:

   ```
   > $cert = "_<certificate thumbprint>_"
   > $guid = [guid]::NewGuid().ToString("B")
   > netsh http add sslcert ipport=0.0.0.0: _<SMCore Port>_ certhash=$cert
   appid="$guid"
   ```

   For example:

   ```
   > $cert = "a909502dd82ae41433e6f83886b00d4277a32a7b"
   > $guid = [guid]::NewGuid().ToString("B")
   > netsh http add sslcert ipport=0.0.0.0: _<SMCore Port>_ certhash=$cert
   appid="$guid"
   ```

Configure CA certificate with SnapCenter site

You should configure the CA certificate with SnapCenter site on Windows host.

Steps

1. Open IIS Manager on the Windows Server where SnapCenter is installed.
2. In the left navigation pane, click **Connections**.
3. Expand the name of the server and **Sites**.
4. Select the SnapCenter website on which you want to install the SSL Certificate.
5. Navigate to **Actions** > **Edit Site**, click **Bindings**.
6. In the Bindings page, select **binding for https**.
7. Click **Edit**.
8. From the SSL certificate drop-down list, select the recently imported SSL Certificate.
9. Click **OK**.

The SnapCenter Scheduler site (default port: 8154, HTTPS) is configured with self-signed certificate. This port is communicating within the SnapCenter Server host and it is not mandatory to configure with a CA certificate. However, if your environment mandates you to use a CA Certificate, repeat steps 5 to 9 using the SnapCenter Scheduler site.

If the recently deployed CA certificate is not listed in the drop-down menu, check if the CA certificate is associated with the private key.

Ensure that the certificate is added using the following path: **Console Root > Certificates – Local Computer > Trusted Root Certification Authorities > Certificates**.

Enable CA certificates for SnapCenter

You should configure the CA certificates and enable the CA certificate validation for the SnapCenter Server.

Before you begin

- You can enable or disable the CA certificates using the Set-SmCertificateSettings cmdlet.
- You can display the certificate status for the SnapCenter Server using the Get-SmCertificateSettings cmdlet.

The information regarding the parameters that can be used with the cmdlet and their descriptions can be obtained by running *Get-Help command_name*. Alternatively, you can refer to the SnapCenter Software Cmdlet Reference Guide.

Steps

1. In the Settings page, navigate to **Settings** > **Global Settings** > **CA Certificate Settings**.
2. Select **Enable Certificate Validation**.
3. Click **Apply**.

After you finish

The Managed Hosts tab host displays a padlock and the color of the padlock indicates the status of the connection between SnapCenter Server and the plug-in host.

- indicates that there is no CA certificate enabled or assigned to the plug-in host.
- indicates that the CA certificate is successfully validated.
- indicates that the CA certificate could not be validated.
- indicates that the connection information could not be retrieved.

When the status is yellow or green, the data protection operations completes successfully.

Configure CA Certificate for Linux host

After installing the SnapCenter Server on Linux, the installer creates the self-signed certificate. If you want to use the CA certificate, you should configure the certificates for

nginx reverse proxy, audit logging, and SnapCenter services.

Configure nginx certificate

Steps

1. Navigate to */etc/nginx/conf.d*: `cd /etc/nginx/conf.d`
2. Open **snapcenter.conf** using vi or any text editor.
3. Navigate to the server section in the configuration file.
4. Modify the paths of *ssl_certificate* and *ssl_certificate_key* to point to CA certificate.
5. Save and close the file.
6. Reload nginx: `$nginx -s reload`

Configure audit log certificate

Steps

1. Open *INSTALL_DIR/NetApp/snapcenter/SnapManagerWeb/SnapManager.Web.UI.dll.config* using vi or any text editor.

   The default value of *INSTALL_DIR* is */opt*.

2. Edit the **AUDILOG_CERTIFICATE_PATH** and **AUDILOG_CERTIFICATE_PASSWORD** keys to include the CA certificate path and password respectively.

   Only *.pfx* format is supported for audit log certificate.

3. Save and close the file.
4. Restart the **snapmanagerweb** service: `$ systemctl restart snapmanagerweb`

Configure SnapCenter services certificate

Steps

1. Open the following configuration files using using vi or any text editor.
   - *INSTALL_DIR/NetApp/snapcenter/SnapManagerWeb/SnapManager.Web.UI.dll.config*
   - *INSTALL_DIR/NetApp/snapcenter/SMCore/SMCoreServiceHost.dll.config*
   - *INSTALL_DIR/NetApp/snapcenter/Scheduler/Scheduler.Api.dll.config*

     The default value of *INSTALL_DIR* is */opt*.

2. Edit the **SERVICE_CERTIFICATE_PATH** and **SERVICE_CERTIFICATE_PASSWORD** keys to include the CA certificate path and password respectively.

   Only *.pfx* format is supported for SnapCenter services certificate.

3. Save and close the files.
4. Restart all the services.
   - `$ systemctl restart snapmanagerweb`
   - `$ systemctl restart smcore`

- $ systemctl restart scheduler

Configure and enable two-way SSL communication on Windows host

Configure two-way SSL communication on Windows host

You should configure the two-way SSL communication to secure the mutual communication between SnapCenter Server on Windows host and the plug-ins.

Before you begin

- You should have generated the CA Certificate CSR file with the minimum supported key length of 3072.
- The CA certificate should support server authentication and client authentication.
- You should have a CA certificate with private key and thumbprint details.
- You should have enabled the one-way SSL configuration.

  For more details, see Configure CA certificate section.

- You must have enabled two-way SSL communication on all the plug-in hosts and the SnapCenter Server.

  Environment with some hosts or server not enabled for two-way SSL communication is not supported.

Steps

1. To bind the port, perform the following steps on SnapCenter Server host for SnapCenter IIS web server port 8146 (default) and once again for SMCore port 8145 (default) using PowerShell commands.
   a. Remove the existing SnapCenter self-signed certificate port binding using the following PowerShell command.

      ```
      > netsh http delete sslcert ipport=0.0.0.0:<SMCore port/IIS port>
      ```

      For example,

      ```
      > netsh http delete sslcert ipport=0.0.0.0:8145

      > netsh http delete sslcert ipport=0.0.0.0:8146
      ```

   b. Bind the newly procured CA certificate with the SnapCenter server and SMCore port.

      ```
      > $cert = "<CA_certificate thumbprint>"

      > $guid = [guid]::NewGuid().ToString("B")

      > netsh http add sslcert ipport=0.0.0.0: <SMCore Port/IIS port>
      certhash=$cert appid="$guid" clientcertnegotiation=enable
      verifyclientcertrevocation=disable

      > netsh http show sslcert ipport=0.0.0.0:<SMCore Port/IIS port>
      ```

      For example,

```
> $cert = "abc123abc123abc123abc123"

> $guid = [guid]::NewGuid().ToString("B")

> netsh http add sslcert ipport=0.0.0.0:8146 certhash=$cert appid="$guid"
clientcertnegotiation=enable verifyclientcertrevocation=disable

> $guid = [guid]::NewGuid().ToString("B")

> netsh http add sslcert ipport=0.0.0.0:8145 certhash=$cert appid="$guid"
clientcertnegotiation=enable verifyclientcertrevocation=disable

> netsh http show sslcert ipport=0.0.0.0:8146

> netsh http show sslcert ipport=0.0.0.0:8145
```

2. To access permission to the CA certificate, add the SnapCenter's default IIS web server user "**IIS AppPool\SnapCenter**" in the certificate permission list by performing the following steps to access the newly procured CA certificate.
   a. Go to the Microsoft management console (MMC), and then click **File** > **Add/Remove SnapIn**.
   b. In the Add or Remove Snap-ins window, select **Certificates** and then click **Add**.
   c. In the Certificates snap-in window, select the **Computer account** option, and then click **Finish**.
   d. Click **Console Root** > **Certificates – Local Computer** > **Personal** > **Certificates**.
   e. Select the SnapCenter certificate.
   f. To start the add user\permission wizard, right-click on the CA certificate and select **All Tasks** > **Manage private keys**.
   g. Click on **Add**, on Select users and groups wizard change the location to local computer name (top most in the hierarchy)
   h. Add the IIS AppPool\SnapCenter user, give full control permissions.
3. For **CA certificate IIS permission**, add the new DWORD registry keys entry in SnapCenter Server from the following path:

   In the windows registry editor, traverse to the below mentioned path,

   ```
   HKey_Local_Machine\SYSTEM\CurrentControlSet\Control\SecurityProv
   ders\SCHANNEL
   ```

4. Create new DWORD registry key entry under the context of SCHANNEL registry configuration.

   ```
   SendTrustedIssuerList = 0

   ClientAuthTrustMode = 2
   ```

Configure SnapCenter Windows plug-in for Two-way SSL communication

You should configure SnapCenter Windows plug-in for two-way SSL communication using PowerShell commands.

Before you begin

Ensure that the CA certificate thumbprint is available.

Steps

1. To bind the port, perform the following actions on Windows plug-in host for SMCore port 8145 (default).
   a. Remove the existing SnapCenter self-signed certificate port binding using the following PowerShell command.

      ```
      > netsh http delete sslcert ipport=0.0.0.0:<SMCore port>
      ```

      For example,

      ```
      > netsh http delete sslcert ipport=0.0.0.0:8145
      ```

   b. Bind the newly procured CA certificate with the SMCore port.

      ```
      > $cert = "<CA_certificate thumbprint>"

      > $guid = [guid]::NewGuid().ToString("B")

      > netsh http add sslcert ipport=0.0.0.0: <SMCore Port> certhash=$cert
      appid="$guid"
      clientcertnegotiation=enable verifyclientcertrevocation=disable

      > netsh http show sslcert ipport=0.0.0.0:<SMCore Port>
      ```

      For example,

      ```
      > $cert = "abc123abc123abc123abc123"

      > $guid = [guid]::NewGuid().ToString("B")

      > netsh http add sslcert ipport=0.0.0.0:8145 certhash=$cert appid="$guid"
      clientcertnegotiation=enable verifyclientcertrevocation=disable

      > netsh http show sslcert ipport=0.0.0.0:8145
      ```

Enable two-way SSL communication on Windows host

You can enable two-way SSL communication to secure the mutual communication between SnapCenter Server on Windows host and the plug-ins using PowerShell commands.

Before you begin

Execute the commands for all the plug-ins and the SMCore agent first and then for server.

Steps

1. To enable the two-way SSL communication, run the following commands on the SnapCenter Server for the plug-ins, server, and for each of the agents for which the two-way SSL communication is required.

   ```
   > Set-SmConfigSettings -Agent -configSettings @{"EnableTwoWaySSL"="true"}
   ```

```
-HostName <Plugin_HostName>

> Set-SmConfigSettings -Agent -configSettings @{"EnableTwoWaySSL"="true"}
-HostName localhost

> Set-SmConfigSettings -Server -configSettings @{"EnableTwoWaySSL"="true"}
```

2. Perform the IIS SnapCenter Application pool recycle operation by using the following command.
   `> Restart-WebAppPool -Name "SnapCenter"`
3. For Windows plug-ins, restart the SMCore service by running the following PowerShell command:

   ```
   > Restart-Service -Name SnapManagerCoreService
   ```

Disable two-way SSL Communication

You can disable the two-way SSL communication using PowerShell commands.

About this task

- Execute the commands for all the plug-ins and the SMCore agent first and then for server.
- When you disable the two-way SSL communication, the CA certificate and its configuration are not removed.
- To add a new host to SnapCenter Server, you must disable the two-way SSL for all plug-in hosts.
- NLB and F5 are not supported.

Steps

1. To disable the two-way SSL communication, run the following commands on SnapCenter Server for all the plug-in hosts and the SnapCenter host.

   ```
   > Set-SmConfigSettings -Agent -configSettings @{"EnableTwoWaySSL"="false"}
   -HostName <Agent_HostName>

   > Set-SmConfigSettings -Agent -configSettings @{"EnableTwoWaySSL"="false"}
   -HostName localhost

   > Set-SmConfigSettings -Server -configSettings @{"EnableTwoWaySSL"="false"}
   ```

2. Perform the IIS SnapCenter Application pool recycle operation by using the following command.
   `> Restart-WebAppPool -Name "SnapCenter"`
3. For Windows plug-ins, restart the SMCore service by running the following PowerShell command:

   ```
   > Restart-Service -Name SnapManagerCoreService
   ```

Configure and enable two-way SSL communication on Linux host

Configure two-way SSL communication on Linux host

You should configure the two-way SSL communication to secure the mutual communication between SnapCenter Server on Linux host and the plug-ins.

Before you begin

- You should have configured the CA certificate for Linux host.
- You must have enabled two-way SSL communication on all the plug-in hosts and the SnapCenter Server.

Steps

1. Copy **certificate.pem** to */etc/pki/ca-trust/source/anchors/*.
2. Add the certificates in the trust list of your Linux host.
    - `cp root-ca.pem /etc/pki/ca-trust/source/anchors/`
    - `cp certificate.pem /etc/pki/ca-trust/source/anchors/`
    - `update-ca-trust extract`
3. Verify if the certificates were added to the trust list.
   `trust list | grep "<CN of your certificate>"`
4. Update **ssl_certificate** and **ssl_certificate_key** in the SnapCenter **nginx** file and restart.
    - `vim /etc/nginx/conf.d/snapcenter.conf`
    - `systemctl restart nginx`
5. Refresh the SnapCenter Server GUI link.
6. Update the values of the following keys in **SnapManager.Web.UI.dll.config** located at _ /<installation path>/NetApp/snapcenter/SnapManagerWeb_ and **SMCoreServiceHost.dll.config** located at */<installation path>/NetApp/snapcenter/SMCore*.
    - <add key="SERVICE_CERTIFICATE_PATH" value="<path of certificate.pfx>" />
    - <add key="SERVICE_CERTIFICATE_PASSWORD" value="<password>"/>
7. Restart the following services.
    - `systemctl restart smcore.service`
    - `systemctl restart snapmanagerweb.service`
8. Verify that the certificate is attached to the SnapManager web port.
   `openssl s_client -connect localhost:8146 -brief`
9. Verify that the certificate is attached to the smcore port.
   `openssl s_client -connect localhost:8145 -brief`
10. Manage password for SPL keystore and alias.
    a. Retrieve SPL keystore default password assigned to the **SPL_KEYSTORE_PASS** key in SPL property file.
    b. Change the keystore password.
       `keytool -storepasswd -keystore keystore.jks`
    c. Change the password for all the aliases of private key entries.
       `keytool -keypasswd -alias "<alias_name>" -keystore keystore.jks`

d. Update the same password for the key **SPL_KEYSTORE_PASS** in *spl.properties*.

e. Restart the service.

11. On plug-in Linux host, add the root and intermediate certificates in SPL plug-in's keystore.

    - `keytool -import -trustcacerts -alias <any preferred alias name> -file <path of root-ca.pem> -keystore <path of keystore.jks mentioned in spl.properties file>`
    - `keytool -importkeystore -srckeystore <path of certificate.pfx> -srcstoretype pkcs12 -destkeystore <path of keystore.jks mentioned in spl.properties file> -deststoretype JKS`

      a. Check the entries in keystore.jks.
      `keytool -list -v -keystore <path to keystore.jks>`

      b. Rename any alias if required.
      `keytool -changealias -alias "old-alias" -destalias "new-alias" -keypass keypass -keystore </path/to/keystore> -storepass storepas`

12. Update the value of **SPL_CERTIFICATE_ALIAS** in *spl.properties* file with the alias of **certificate.pfx** stored in *keystore.jks* and restart the SPL service: `systemctl restart spl`

13. Verify that the certificate is attached to the smcore port.
`openssl s_client -connect localhost:8145 -brief`

Enable SSL communication on Linux host

You can enable two-way SSL communication to secure the mutual communication between SnapCenter Server on Linux host and the plug-ins using PowerShell commands.

Step

1. Perform the following to enable one-way SSL communication.
   a. Log into SnapCenter GUI.
   b. Click **Settings** > **Global Settings** and select **Enable certificate validation on SnapCenter Server**.
   c. Click **Hosts** > **Managed Hosts** and select the plug-in host for which you want to enable one-way SSL.
   d. Click icon, and then click **Enable certificate validation**.
2. Enable two-way SSL communication from the SnapCenter Server Linux host.
   - `Open-SmConnection`
   - `Set-SmConfigSettings -Agent -configSettings @{"EnableTwoWaySSL"="true"} -HostName <Plugin Host Name>`
   - `Set-SmConfigSettings -Agent -configSettings @{"EnableTwoWaySSL"="true"} -HostName localhost`
   - `Set-SmConfigSettings -Server -configSettings @{"EnableTwoWaySSL"="true"}`

Configure Active Directory, LDAP, and LDAPS

Register untrusted Active Directory domains

You should register the Active Directory with SnapCenter Server to manage hosts, users, and groups from multiple untrusted Active Directory domains.

Before you begin

LDAP and LDAPS protocols

- You can register the untrusted active directory domains using either LDAP or LDAPS protocol.
- You should have enabled bidirectional communication between the plug-in hosts and the SnapCenter Server.
- DNS resolution should be set up from the SnapCenter Server to the plug-in hosts and vice-versa.

LDAP protocol

- The fully qualified domain name (FQDN) should be resolvable from SnapCenter Server.

  You can register an untrusted domain with the FQDN. If the FQDN is not resolvable from the SnapCenter Server, you can register with a domain controller IP address and this should be resolvable from SnapCenter Server.

LDAPS protocol

- CA certificates are required for LDAPS to provide end-to-end encryption during the active directory communication.

  Configure CA client certificate for LDAPS

- Domain controller host names (DCHostName) should be reachable from SnapCenter Server.

About this task

- You can use either the SnapCenter user interface, PowerShell cmdlets, or REST API to register an untrusted domain.

Steps

1. In the left navigation pane, click **Settings**.
2. In the Settings page, click **Global Settings**..
3. In the Global Settings page, click **Domain Settings**.
4. Click to register a new domain.
5. In the Register New Domain page, select either **LDAP** or **LDAPS**.
   a. If you select **LDAP**, specify the information that is required for registering the untrusted domain for LDAP:

For this field…	Do this…
Domain Name	Specify the NetBIOS name for the domain.

For this field…	Do this…
Domain FQDN	Specify the FQDN and click **Resolve**.
Domain controller IP addresses	If the domain FQDN is not resolvable from the SnapCenter Server, specify one or more domain controller IP addresses. For more information, see Add domain controller IP for untrusted domain from GUI.

b. If you select **LDAPS**, specify the information that is required for registering the untrusted domain for LDAPS:

For this field…	Do this…
Domain Name	Specify the NetBIOS name for the domain.
Domain FQDN	Specify the FQDN.
Domain controller Names	Specify one or more domain controller names and click **Resolve**.
Domain controller IP addresses	If the domain controller names is not resolvable from SnapCenter Server, you should rectify the DNS resolutions.

6. Click **OK**.

Configure IIS Application Pools to enable Active Directory read permissions

You can configure Internet Information Services (IIS) on your Windows Server to create a custom Application Pool account when you need to enable Active Directory read permissions for SnapCenter.

Steps

1. Open IIS Manager on the Windows Server where SnapCenter is installed.
2. In the left navigation pane, click **Application Pools**.
3. Select SnapCenter in the Application Pools list, and then click **Advanced Settings** in the Actions pane.
4. Select Identity, and then click **…** to edit the SnapCenter application pool identity.
5. In the Custom Account field, enter a domain user or domain admin account name with Active Directory read permission.
6. Click OK.

   The custom account replaces the built-in ApplicationPoolIdentity account for the SnapCenter application pool.

Configure CA client certificate for LDAPS

You should configure the CA client certificate for LDAPS on the SnapCenter Server when the Windows Active Directory LDAPS is configured with the CA certificates.

Steps

1. Go to the Microsoft management console (MMC), and then click **File** > **Add/Remove Snapin**.
2. In the Add or Remove Snap-ins window, select **Certificates** and then click **Add**.
3. In the Certificates snap-in window, select the **Computer account** option, and then click **Finish**.
4. Click **Console Root** > **Certificates – Local Computer** > **Trusted Root Certification Authorities** > **Certificates**.
5. Right-click on the folder “Trusted Root Certification Authorities”, and then select **All Tasks** > **Import** to start the import wizard.
6. Complete the wizard, as follows:

In this wizard window…	Do the following…
In the second page of the wizard	Click **Browse**, select the *Root Certificate* and click **Next**.
Completing the Certificate Import Wizard	Review the summary, and then click **Finish** to start the import.

7. Repeat Steps 5 and 6 for the intermediate certificates.

Protect Microsoft SQL Server databases

SnapCenter Plug-in for Microsoft SQL Server

SnapCenter Plug-in for Microsoft SQL Server overview

The SnapCenter Plug-in for Microsoft SQL Server is a host-side component of the NetApp SnapCenter Software that enables application-aware data protection management of Microsoft SQL Server databases. The Plug-in for SQL Server automates SQL Server database backup, verification, restore, and clone operations in your SnapCenter environment.

When the Plug-in for SQL Server is installed, you can use SnapCenter with NetApp SnapMirror technology to create mirror copies of backup sets on another volume and with NetApp SnapVault technology to perform disk-to-disk backup replication for standards compliance or archival purposes.

- Automates application-aware backup, restore, and clone operations for Microsoft SQL Server databases in your SnapCenter environment.
- Supports Microsoft SQL Server databases on VMDK and raw device mapping (RDM) LUNs when you deploy the SnapCenter Plug-in for VMware vSphere and register the plug-in with SnapCenter
- Supports provisioning SMB shares only. Support is not provided for backing up SQL Server databases on SMB shares.
- Supports importing backups from SnapManager for Microsoft SQL Server to SnapCenter.

What you can do with the SnapCenter Plug-in for Microsoft SQL Server

When the SnapCenter Plug-in for Microsoft SQL Server is installed in your environment, you can use SnapCenter to back up, restore, and clone SQL Server databases.

You can perform the following tasks that support backup operations, restore operations, and clone operations of SQL Server databases and database resources:

- Back up SQL Server databases and associated transaction logs

  You cannot create a log backup for master and msdb system databases. However, you can create log backups for model system database.

- Restore database resources
  - You can restore master system databases, msdb system databases, and model system databases.
  - You cannot restore multiple databases, instances, and availability groups.
  - You cannot restore the system database to an alternate path.
- Create point-in-time clones of production databases

  You cannot perform backup, restore, clone, and clone lifecycle operations on tempdb system databases.

- Verify backup operations immediately or defer verification until later

  Verification of SQL Server system database is not supported. SnapCenter clones the databases to perform

verification operation. SnapCenter cannot clone SQL Server system databases, and therefore verification of these databases is not supported.

- Schedule backup operations and clone operations
- Monitor backup operations, restore operations, and clone operations

The Plug-in for SQL Server does not support backup and recovery of SQL Server databases on SMB shares.

SnapCenter Plug-in for Microsoft SQL Server features

The Plug-in for SQL Server integrates with Microsoft SQL Server on the Windows host and with NetApp Snapshot technology on the storage system. To work with the Plug-in for SQL Server, you use the SnapCenter interface.

The Plug-in for SQL Server includes these major features:

- **Unified graphical user interface powered by SnapCenter**

  The SnapCenter interface provides you with standardization and consistency across plug-ins and environments. The SnapCenter interface enables you to complete consistent backup and restore processes across plug-ins, use centralized reporting, use at-a-glance dashboard views, set up role-based access control (RBAC), and monitor jobs across all plug-ins. SnapCenter also offers centralized scheduling and policy management to support backup and clone operations.

- **Automated central administration**

  You can schedule routine SQL Server backups, configure policy-based backup retention, and set up point-in-time and up-to-the-minute restore operations. You can also proactively monitor your SQL Server environment by configuring SnapCenter to send email alerts.

- **Nondisruptive NetApp Snapshot technology**

  The Plug-in for SQL Server uses NetApp Snapshot technology with the NetApp SnapCenter Plug-in for Microsoft Windows. This enables you to back up databases in seconds and restore them quickly without taking SQL Server offline. Snapshots consume minimal storage space.

In addition to these major features, the Plug-in for SQL Server offers the following benefits:

- Backup, restore, clone, and verification workflow support
- RBAC-supported security and centralized role delegation
- Creation of space-efficient, point-in-time copies of production databases for testing or data extraction by using NetApp FlexClone technology

  A FlexClone license is required on the storage system holding the clone.

- Nondisruptive and automated backup verification
- Ability to run multiple backups at the same time across multiple servers
- PowerShell cmdlets for scripting of backup, verify, restore, and clone operations
- Support for AlwaysOn Availability Groups (AGs) in SQL Server to accelerate AG setup, backup, and restore operations

- In-memory database and Buffer Pool Extension (BPE) as part of SQL Server 2014
- Support for backup of LUNs and virtual machine disks (VMDKs)
- Support for physical and virtualized infrastructures
- Support for iSCSI, Fibre Channel, FCoE, raw device mapping (RDM), and VMDK over NFS and VMFS

  NAS volumes should have a default export policy in storage virtual machine (SVM).

- Support for FileStream and file group in SQL Server standalone databases.
- Support for non-volatile memory express (NVMe) on Windows Server 2022
  - Backup, restore, clone, and verification workflows on VMDK layout created on NVMe over TCP/IP.
  - Supports NVMe firmware version 1.3 starting from ESX 8.0 update 2 and requires Virtual hardware version 21.
  - Windows Server Failover Clustering (WSFC) is not supported for applications over VMDK on NVMe over TCP/IP.
- Supports SnapMirror active sync (initially released as SnapMirror Business Continuity [SM-BC]) that enables business services to continue operating even through a complete site failure, supporting applications to fail over transparently using a secondary copy. Neither manual intervention nor additional scripting is required to trigger a failover with SnapMirror active sync.

Support for Asymmetric LUN Mapping in Windows clusters

SnapCenter Plug-in for Microsoft SQL Server supports discovery in SQL Server 2012 and later, Asymmetric LUN Mapping (ALM) configurations for high availability, and availability groups for disaster recovery. When discovering resources, SnapCenter discovers databases on local hosts and on remote hosts in ALM configurations.

An ALM configuration is a single Windows server failover cluster that contains one or more nodes in a primary data center and one or more nodes in a disaster recovery center.

Following is an example of an ALM configuration:

- Two failover cluster instances (FCI) in a multi-site datacenter
- FCI for local high availability (HA) and Availability Group (AG) for disaster recovery with a stand-alone instance at the disaster recovery site

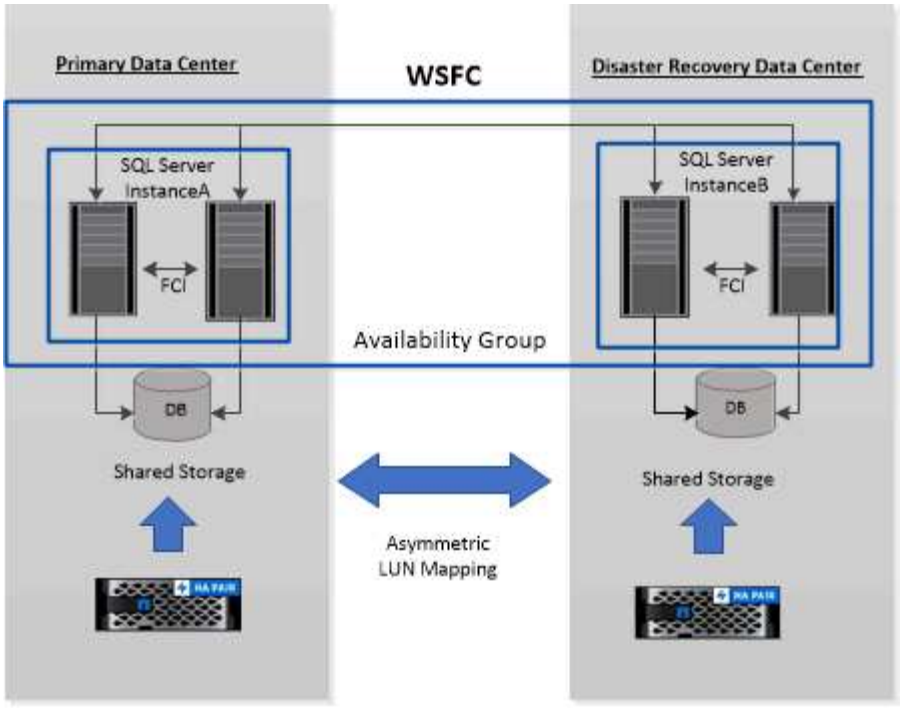

WSFC----Windows Server Failover Cluster

The storage in the primary datacenter is shared between the FCI nodes present in the primary datacenter. The storage in the disaster recovery datacenter is shared between the FCI nodes present in the disaster recovery datacenter.

The storage on the primary datacenter is not visible to the nodes on the disaster recovery datacenter, and vice versa.

ALM architecture combines two shared storage solution used by FCI, with non-shared or dedicated storage solution used by SQL AG. The AG solution uses identical drive letters for shared disk resources across data centers. This arrangement of storage, where a cluster disk is shared between a subset of nodes within a WSFC, is referred to as ALM.

Storage types supported by SnapCenter Plug-in for Microsoft SQL Server

SnapCenter supports a wide range of storage types on both physical machines and virtual machines. You must verify whether support is available for your storage type before installing the package for your host.

SnapCenter provisioning and data protection support is available on Windows Server. For the latest information about supported versions, see the NetApp Interoperability Matrix Tool.

Machine	Storage type	Provision using	Support notes
Physical server	FC-connected LUNs	SnapCenter graphical user interface (GUI) or PowerShell cmdlets	

Machine	Storage type	Provision using	Support notes
Physical server	iSCSI-connected LUNs	SnapCenter GUI or PowerShell cmdlets	
Physical server	SMB3 (CIFS) shares residing on a storage virtual machine (SVM)	SnapCenter GUI or PowerShell cmdlets	Support for provisioning only.
VMware VM	RDM LUNs connected by an FC or iSCSI HBA	PowerShell cmdlets	
VMware VM	iSCSI LUNs connected directly to the guest system by the iSCSI initiator	SnapCenter GUI or PowerShell cmdlets	
VMware VM	Virtual Machine File Systems (VMFS) or NFS datastores	VMware vSphere	
VMware VM	A guest system connected to SMB3 shares residing on an SVM	SnapCenter GUI or PowerShell cmdlets	Support for provisioning only.
VMware VM	vVol datastores on both NFS and SAN	ONTAP Tools for VMware vSphere	
Hyper-V VM	Virtual FC (vFC) LUNs connected by a virtual Fibre Channel Switch	SnapCenter GUI or PowerShell cmdlets	You must use Hyper-V Manager to provision Virtual FC (vFC) LUNs connected by a virtual Fibre Channel Switch. Hyper-V pass through disks and backing up databases on VHD(x) that are provisioned on NetApp storage are not supported.

Machine	Storage type	Provision using	Support notes
Hyper-V VM	iSCSI LUNs connected directly to the guest system by the iSCSI initiator	SnapCenter GUI or PowerShell cmdlets	Hyper-V pass through disks and backing up databases on VHD(x) that are provisioned on NetApp storage are not supported.
Hyper-V VM	A guest system connected to SMB3 shares residing on an SVM	SnapCenter GUI or PowerShell cmdlets	Support for provisioning only. Hyper-V pass through disks and backing up databases on VHD(x) that are provisioned on NetApp storage are not supported.

Storage layout recommendations for SnapCenter Plug-in for Microsoft SQL Server

A well-designed storage layout allows SnapCenter Server to back up your databases to meet your recovery objectives. You should consider several factors while defining your storage layout, including the size of the database, the rate of change of the database, and the frequency with which you perform backups.

The following sections define the storage layout recommendations and restrictions for LUNs and virtual machine disks (VMDKs) with SnapCenter Plug-in for Microsoft SQL Server installed in your environment.

In this case, LUNs can include VMware RDM disks and iSCSI direct-attached LUNs that are mapped to the guest.

LUN and VMDK requirements

You can optionally use dedicated LUNs or VMDKs for optimum performance and management for the following databases:

- Master and model system databases
- Tempdb
- User database files (.mdf and .ndf)
- User database transaction log files (.ldf)
- Log directory

To restore large databases, the best practice is to use dedicated LUNs or VMDKs. The time taken to restore a complete LUN or VMDK is less than the time taken to restore the individual files that are stored in the LUN or VMDK.

For the log directory, you should create a separate LUN or VMDK so that there is sufficient free space in the data or log file disks.

LUN and VMDK sample layouts

The following graphic shows how you can configure the storage layout for large databases on LUNs:

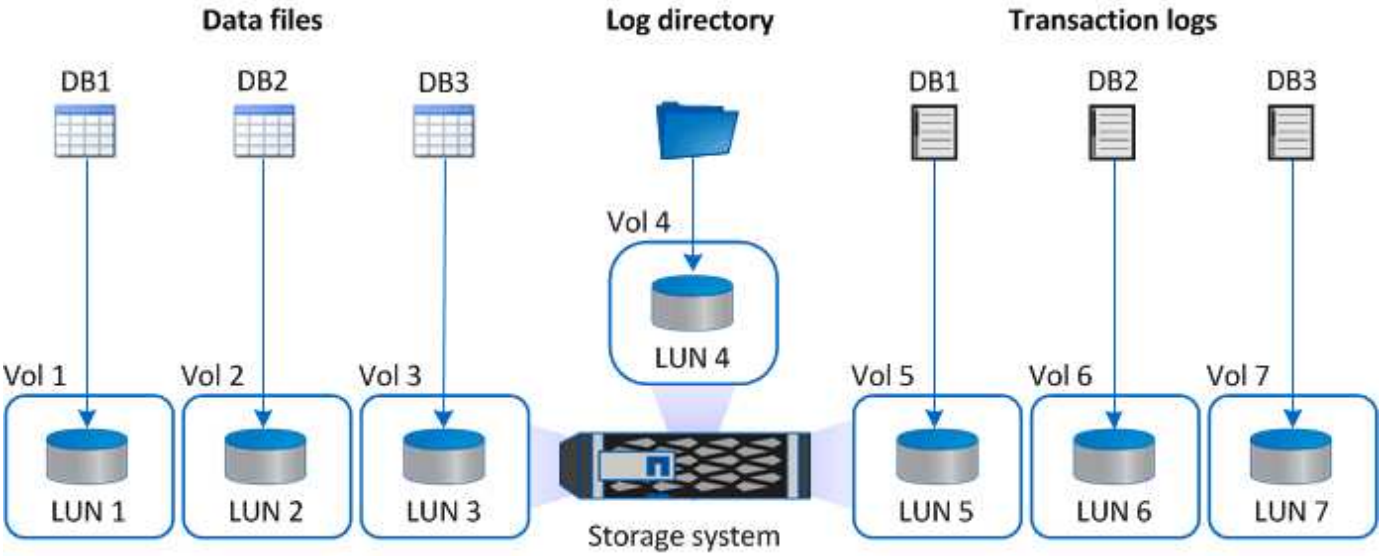

The following graphic shows how you can configure the storage layout for medium or small databases on LUNs:

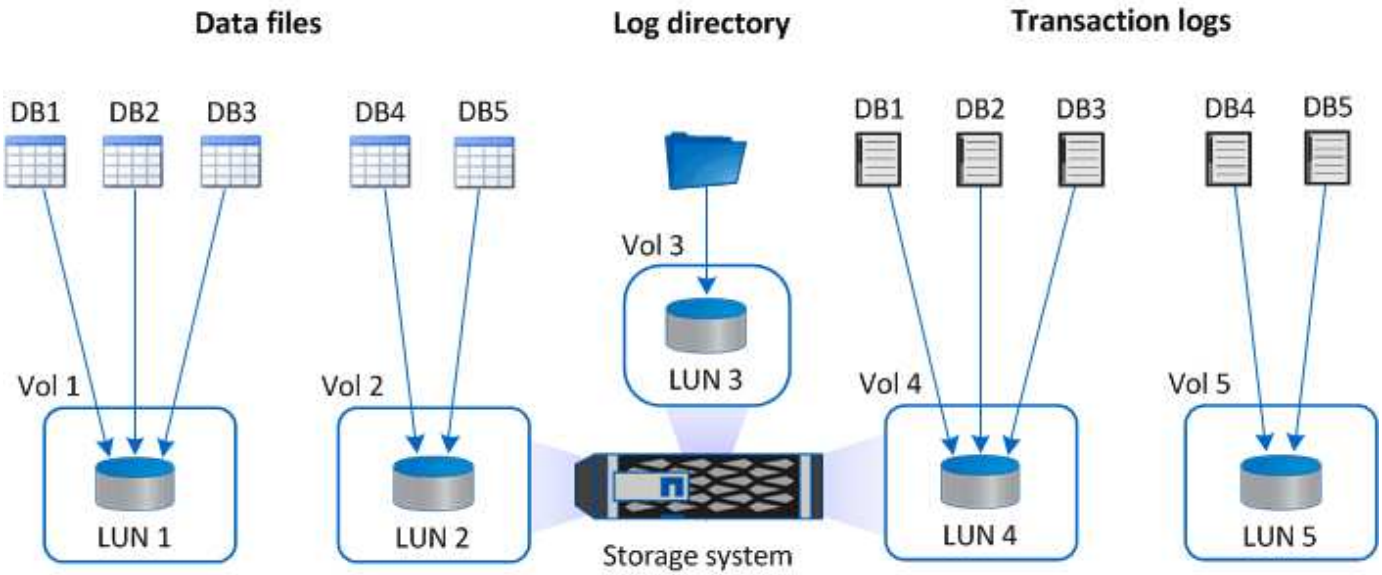

The following graphic shows how you can configure the storage layout for large databases on VMDKs:

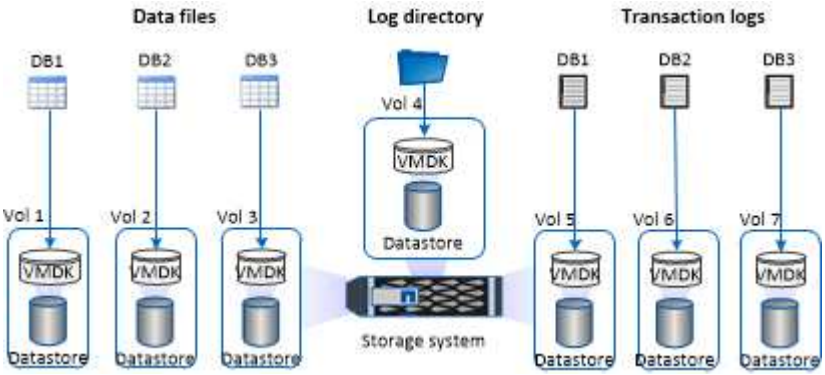

The following graphic shows how you can configure the storage layout for medium or small databases on VMDKs:

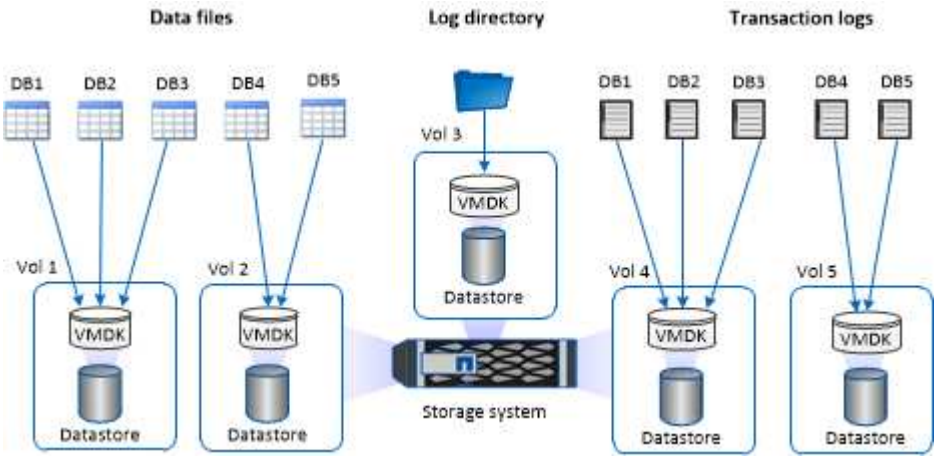

Minimum ONTAP privileges required for SQL plug-in

The minimum ONTAP privileges that are required vary according to the SnapCenter plug-ins you are using for data protection.

- All-access commands: Minimum privileges required for ONTAP 9.12.1 and later
  - event generate-autosupport-log
  - job history show
  - job stop
  - lun
  - lun create
  - lun delete
  - lun igroup add
  - lun igroup create
  - lun igroup delete
  - lun igroup rename
  - lun igroup show
  - lun mapping add-reporting-nodes
  - lun mapping create
  - lun mapping delete
  - lun mapping remove-reporting-nodes
  - lun mapping show
  - lun modify
  - lun move-in-volume
  - lun offline
  - lun online
  - lun resize
  - lun serial
  - lun show
  - snapmirror policy add-rule
  - snapmirror policy modify-rule
  - snapmirror policy remove-rule
  - snapmirror policy show
  - snapmirror restore
  - snapmirror show
  - snapmirror show-history
  - snapmirror update
  - snapmirror update-ls-set
  - snapmirror list-destinations
  - version
  - volume clone create
  - volume clone show

- volume clone split start
- volume clone split stop
- volume create
- volume destroy
- volume file clone create
- volume file show-disk-usage
- volume offline
- volume online
- volume modify
- volume qtree create
- volume qtree delete
- volume qtree modify
- volume qtree show
- volume restrict
- volume show
- volume snapshot create
- volume snapshot delete
- volume snapshot modify
- volume snapshot rename
- volume snapshot restore
- volume snapshot restore-file
- volume snapshot show
- volume unmount
- vserver cifs
- vserver cifs share create
- vserver cifs share delete
- vserver cifs shadowcopy show
- vserver cifs share show
- vserver cifs show
- vserver export-policy
- vserver export-policy create
- vserver export-policy delete
- vserver export-policy rule create
- vserver export-policy rule show
- vserver export-policy show
- vserver iscsi
- vserver iscsi connection show

- vserver show
- network interface
- network interface show
- vserver
- metrocluster show

Prepare storage systems for SnapMirror and SnapVault replication for Plug-in for SQL server

You can use a SnapCenter plug-in with ONTAP SnapMirror technology to create mirror copies of backup sets on another volume, and with ONTAP SnapVault technology to perform disk-to-disk backup replication for standards compliance and other governance-related purposes. Before you perform these tasks, you must configure a data-protection relationship between the source and destination volumes and initialize the relationship.

SnapCenter performs the updates to SnapMirror and SnapVault after it completes the Snapshot operation. SnapMirror and SnapVault updates are performed as part of the SnapCenter job; do not create a separate ONTAP schedule.

If you are coming to SnapCenter from a NetApp SnapManager product and are satisfied with the data protection relationships you have configured, you can skip this section.

A data protection relationship replicates data on primary storage (the source volume) to secondary storage (the destination volume). When you initialize the relationship, ONTAP transfers the data blocks referenced on the source volume to the destination volume.

SnapCenter does not support cascade relationships between SnapMirror and SnapVault volumes (**Primary** > **Mirror** > **Vault**). You should use fanout relationships.

SnapCenter supports the management of version-flexible SnapMirror relationships. For details about version-flexible SnapMirror relationships and how to set them up, see the ONTAP documentation.

Backup strategy for SQL Server resources

Define a backup strategy for SQL Server resources

Defining a backup strategy before you create your backup jobs helps ensure that you have the backups that you require to successfully restore or clone your databases. Your Service Level Agreement (SLA), Recovery Time Objective (RTO), and Recovery Point Objective (RPO) largely determine your backup strategy.

An SLA defines the level of service expected and addresses many service-related issues, including the availability and performance of service. The RTO is the time by when a business process must be restored after a disruption in service. An RPO defines the strategy for the age of the files that must be recovered from backup storage for regular operations to resume after a failure. The SLA, RTO, and RPO contribute to the backup strategy.

Type of backups supported

Backing up SQL Server system and user databases using SnapCenter requires that you choose the resource type, such as databases, SQL server instances, and Availability Groups (AG). Snapshot technology is leveraged to create online, read-only copies of the volumes on which the resources reside.

You can select the copy-only option to specify that the SQL Server does not truncate transaction logs. You should use this option when you are also managing the SQL Server with other backup applications. Keeping the transaction logs intact enables any backup application to restore the system databases. Copy-only backups are independent of the sequence of scheduled backups, and they do not affect the backup and restore procedures of the database.

Backup type	Description	Copy-only option with backup type
Full backup and log backup	Backs up the system database and truncates the transaction logs. The SQL Server truncates the transaction logs by removing the entries that are already committed to the database. After the full backup is complete, this option creates a transaction log that captures transaction information. Typically, you should choose this option. However, if your backup time is short, you can choose not to run a transaction log backup with full backup. You cannot create a log backup for master and msdb system databases. However, you can create log backups for model system database.	Backs up the system database files and the transaction logs without truncating the logs. A copy-only backup cannot serve as a differential base or differential backup, and does not affect the differential base. Restoring a copy-only full backup is the same as restoring any other full backup.
Full database backup	Backs up the system database files. You can create full database backup for master, model, and msdb system databases.	Backs up the system database files.

Backup type	Description	Copy-only option with backup type
Transaction log backup	Backs up the truncated transaction logs, copying only the transactions that were committed since the most recent transaction log was backed up. If you schedule frequent transaction log backups alongside full database backups, you can choose granular recovery points.	Backs up the transaction logs without truncating them. This backup type does not affect the sequencing of regular log backups. Copy-only log backups are useful for performing online restore operations.

Backup schedules for Plug-in for SQL server

Backup frequency (schedule type) is specified in policies; a backup schedule is specified in the resource group configuration. The most critical factor in determining a backup frequency or schedule is the rate of change for the resource and the importance of the data. You might back up a heavily used resource every hour, while you might back up a rarely used resource once a day. Other factors include the importance of the resource to your organization, your Service Level Agreement (SLA), and your Recover Point Objective (RPO).

An SLA defines the level of service expected and addresses many service-related issues, including the availability and performance of service. An RPO defines the strategy for the age of the files that must be recovered from backup storage for regular operations to resume after a failure. The SLA and RPO contribute to the data protection strategy.

Even for a heavily used resource, there is no requirement to run a full backup more than once or twice a day. For example, regular transaction log backups might be sufficient to ensure that you have the backups you need. The more often you back up your databases, the fewer transaction logs SnapCenter has to use at restore time, which can result in faster restore operations.

Backup schedules have two parts, as follows:

- Backup frequency

  Backup frequency (how often backups are to be performed), called *schedule type* for some plug-ins, is part of a policy configuration. You can select hourly, daily, weekly, or monthly as the backup frequency for the policy. If you do not select any of these frequencies, then the policy created is an on-demand-only policy. You can access policies by clicking **Settings** > **Policies**.

- Backup schedules

  Backup schedules (exactly when backups are to be performed) are part of a resource group configuration. For example, if you have a resource group that has a policy configured for weekly backups, you might configure the schedule to back up every Thursday at 10:00 PM. You can access resource group schedules by clicking **Resources** > **Resource Groups**.

Number of backup jobs needed for databases

Factors that determine the number of backup jobs that you need include the size of the database, the number of volumes used, the rate of change of the database, and your Service Level Agreement (SLA).

For database backups, the number of backup jobs that you choose typically depends on the number of volumes on which you placed your databases. For example, if you placed a group of small databases on one volume and a large database on another volume, you might create one backup job for the small databases and one backup job for the large database.

Backup naming conventions for Plug-in for SQL server

You can either use the default Snapshot naming convention or use a customized naming convention. The default backup naming convention adds a timestamp to Snapshot names that helps you identify when the copies were created.

The Snapshot uses the following default naming convention:

```
resourcegroupname_hostname_timestamp
```

You should name your backup resource groups logically, as in the following example:

```
dts1_mach1x88_03-12-2015_23.17.26
```

In this example, the syntax elements have the following meanings:

- *dts1* is the resource group name.
- *mach1x88* is the host name.
- *03-12-2015_23.17.26* is the date and timestamp.

Alternatively, you can specify the Snapshot name format while protecting resources or resource groups by selecting **Use custom name format for Snapshot copy**. For example, customtext_resourcegroup_policy_hostname or resourcegroup_hostname. By default, the time stamp suffix is added to the Snapshot name.

Backup retention options for Plug-in for SQL Server

You can choose either the number of days for which to retain backup copies or specify the number of backup copies you want to retain, up to a ONTAP maximum of 255 copies. For example, your organization might require that you retain 10 days of backup copies or 130 backup copies.

While creating a policy, you can specify the retention options for the backup type and the schedule type.

If you set up SnapMirror replication, the retention policy is mirrored on the destination volume.

SnapCenter deletes the retained backups that have retention labels that match the schedule type. If the schedule type was changed for the resource or resource group, backups with the old schedule type label might still remain on the system.

For long-term retention of backup copies, you should use SnapVault backup.

How long to retain transaction log backups on the source storage system

SnapCenter Plug-in for Microsoft SQL Server needs transaction log backups to perform up-to-the-minute restore operations, which restore your database to a time between two full backups.

For example, if Plug-in for SQL Server took a full backup at 8:00 a.m. and another full backup at 5:00 p.m., it could use the latest transaction log backup to restore the database to any time between 8:00 a.m. and 5:00 p.m. If transaction logs are not available, Plug-in for SQL Server can perform point-in-time restore operations only, which restore a database to the time that Plug-in for SQL Server completed a full backup.

Typically, you require up-to-the-minute restore operations for only a day or two. By default, SnapCenter retains a minimum of two days.

Multiple databases on the same volume

You can put all databases on the same volume, because the backup policy has an option to set the maximum databases per backup (default value is 100).

For example, if you have 200 databases in the same volume, two Snapshots are created with 100 databases in each of the two Snapshots.

Backup copy verification using the primary or secondary storage volume for Plug-in for SQL Server

You can verify backup copies on the primary storage volume or on either the SnapMirror or SnapVault secondary storage volume. Verification using a secondary storage volume reduces load on the primary storage volume.

When you verify a backup that is either on the primary or secondary storage volume, all the primary and the secondary Snapshots are marked as verified.

SnapRestore license is required to verify backup copies on SnapMirror and SnapVault secondary storage volume.

When to schedule verification jobs

Although SnapCenter can verify backups immediately after it creates them, doing so can significantly increase the time required to complete the backup job and is resource intensive. Hence, it is almost always best to schedule verification in a separate job for a later time. For example, if you back up a database at 5:00 p.m. every day, you might schedule verification to occur an hour later at 6:00 p.m.

For the same reason, it is usually not necessary to run backup verification every time you perform a backup. Performing verification at regular but less frequent intervals is usually sufficient to ensure the integrity of the backup. A single verification job can verify multiple backups at the same time.

Restoration strategy for SQL Server

Sources and destinations for a restore operation

You can restore a SQL Server database from a backup copy on either primary or secondary storage. You also can restore the database to different destinations in addition to its original location, enabling you to choose the destination that supports your requirements.

Sources for a restore operation

You can restore databases from primary or secondary storage.

Destinations for a restore operation

You can restore databases to various destinations:

Destination	Description
The original location	By default, SnapCenter restores the database to the same location on the same SQL Server instance.
A different location	You can restore the database to a different location on any SQL Server instance within the same host.
Original or different location using different database names	You can restore the database with a different name to any SQL Server instance on the same host where the backup was created.

Restore to alternate host across ESX servers for SQL databases on VMDKs (NFS and VMFS datastores) is not supported.

SQL Server recovery models supported by SnapCenter

Specific recovery models are assigned to each database type by default. The SQL Server database administrator can reassign each database to a different recovery model.

SnapCenter supports three types of SQL Server recovery models:

- Simple recovery model

  When you use the simple recovery model, you cannot back up the transaction logs.

- Full recovery model

  When you use the full recovery model, you can restore a database to its previous state from the point of failure.

- Bulk logged recovery model

  When you use the bulk logged recovery model, you must manually re-execute the bulk logged operation.

You must perform the bulk logged operation if the transaction log that contains the operation's commit record has not been backed up before restore. If the bulk logged operation inserts 10 million rows in a database, and the database fails before the transaction log is backed up, then the restored database will not contain the rows that were inserted by the bulk logged operation.

Types of restore operations

You can use SnapCenter to perform different types of restore operations on SQL Server resources.

- Restore up-to-the-minute
- Restore to a previous point in time

You can restore up to the minute or restore to a previous point in time in the following situations:

- Restore from SnapMirror or SnapVault secondary storage
- Restore to alternate path (location)

SnapCenter does not support volume-based SnapRestore.

Restore up to the minute

In an up-to-the-minute restore operation (selected by default), databases are recovered up to the point of failure. SnapCenter accomplishes this by performing the following sequence:

1. Backs up the last active transaction log before restoring the database.
2. Restores the databases from the full database backup that you select.
3. Applies all the transaction logs that were not committed to the databases (including transaction logs from the backups from the time the backup was created up to the most current time).

   Transaction logs are moved ahead and applied to any selected databases.

An up-to-the-minute restore operation requires a contiguous set of transaction logs.

Because the SnapCenter cannot restore SQL Server database transaction logs from log-shipping backup files (log-shipping enables you to automatically send transaction log backups from a primary database on a primary server instance to one or more secondary databases on separate secondary server instances), you are not able to perform an up-to-the-minute restore operation from the transaction log backups. For this reason, you should use the SnapCenter to back up your SQL Server database transaction log files.

If you do not need to retain up-to-the-minute restore capability for all backups, you can configure your system's transaction log backup retention through the backup policies.

Example of an up-to-the-minute restore operation

Assume that you run the SQL Server backup every day at noon, and on Wednesday at 4:00 p.m. you need to restore from a backup. For some reason, the backup from Wednesday noon failed verification, so you decide to restore from the Tuesday noon backup. After that, if the backup is restored, all the transaction logs are moved forward and applied to the restored databases, starting with those that were not committed when you created Tuesday's backup and continuing through the latest transaction log written on Wednesday at 4:00 p.m. (if the transaction logs were backed up).

Restore to a previous point in time

In a point-in-time restore operation, databases are restored only to a specific time from the past. A point-in-time restore operation occurs in the following restore situations:

- The database is restored to a given time in a backed-up transaction log.
- The database is restored, and only a subset of backed-up transaction logs are applied to it.

Restoring a database to a point in time results in a new recovery path.

The following image illustrates the issues when a point-in-time restore operation is performed:

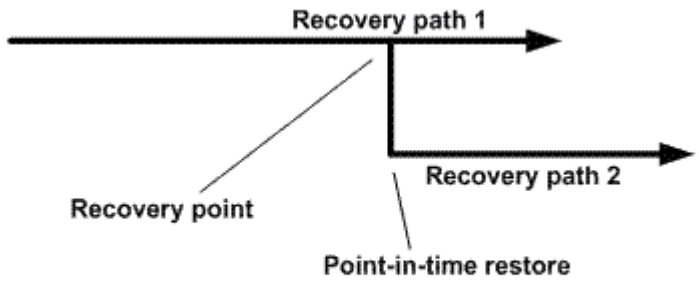

In the image, recovery path 1 consists of a full backup followed by several transaction log backups. You restore the database to a point in time. New transaction log backups are created after the point-in-time restore operation, which results in recovery path 2. The new transaction log backups are created without creating a new full backup. Due to data corruption or other problems, you cannot restore the current database until a new full backup is created. Also, it is not possible to apply the transaction logs created in recovery path 2 to the full backup belonging to recovery path 1.

If you apply transaction log backups, you can also specify a particular date and time at which you want to stop the application of backed up transactions. To do this, you specify a date and time within the available range and the SnapCenter removes any transactions that were not committed prior to that point in time. You can use this method to restore databases to a point in time before a corruption occurred, or to recover from an accidental database or table deletion.

Example of a point-in-time restore operation

Suppose you make full database backups once at midnight and a transaction log backup every hour. The database crashes at 9:45 a.m., but you still back up the transaction logs of the failed database. You can choose from among these point-in-time restore scenarios:

- Restore the full database backup made at midnight and accept the loss of the database changes made afterward. (Option: None)
- Restore the full database backup and apply all the transaction log backups until 9:45 a.m. (Option: Log until)
- Restore the full database backup and apply transaction log backups, specifying the time you want the transactions to restore from the last set of transaction log backups. (Option: By specific time)

In this case, you would calculate the date and time at which a certain error was reported. Any transactions that were not committed prior to the date and time specified are removed.

Define a cloning strategy for SQL Server

Defining a cloning strategy enables you to clone your database successfully.

1. Review the limitations related to clone operations.
2. Decide the type of clone you require.

Limitations of clone operations

You should be aware of the limitations of clone operations before you clone the databases.

- If you are using any version of Oracle from 11.2.0.4 to 12.1.0.1, the clone operation will be in hung state when you run the *renamedg* command. You can apply the Oracle patch 19544733 to fix this issue.
- Cloning of databases from a LUN that is directly attached to a host (for instance, by using Microsoft iSCSI Initiator on a Windows host) to a VMDK or an RDM LUN on the same Windows host, or another Windows host, or vice versa, is not supported.
- The root directory of the volume mount point cannot be a shared directory.
- If you move a LUN that contains a clone to a new volume, the clone cannot be deleted.

Types of clone operations

You can use SnapCenter to clone either a SQL Server database backup or a production database.

- Clone from a database backup

  The cloned database can serve as a baseline for developing new applications and help isolate application errors that occur in the production environment. The cloned database can also be used for recovery from soft database errors.

- Clone lifecycle

  You can use SnapCenter to schedule recurring clone jobs that will occur when the production database is not busy.

Prepare to install the SnapCenter Plug-in for Microsoft SQL Server

Installation workflow for SnapCenter Plug-in for Microsoft SQL Server

You should install and set up the SnapCenter Plug-in for Microsoft SQL Server if you want to protect SQL Server databases.

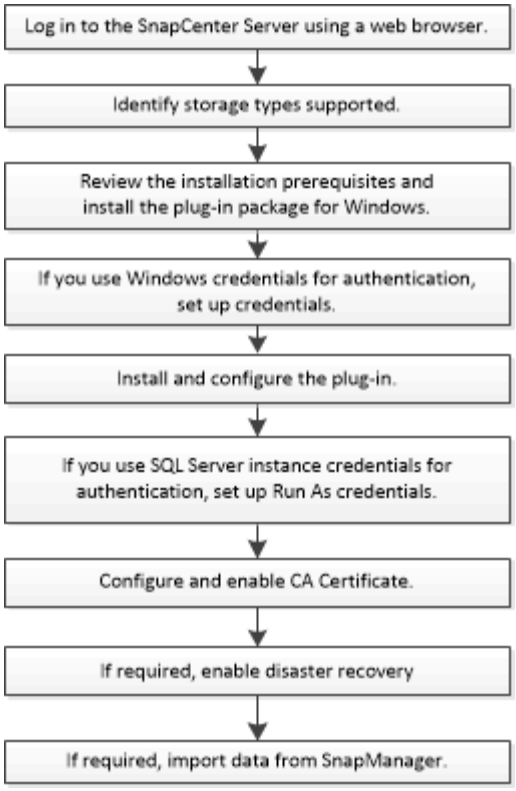

Prerequisites to add hosts and install SnapCenter Plug-in for Microsoft SQL Server

Before you add a host and install the plug-ins packages, you must complete all the requirements.

- If you are using iSCSI, the iSCSI service must be running.
- You must have a user with local administrator privileges with local login permissions on the remote host.
- If you manage cluster nodes in SnapCenter, you must have a user with administrative privileges to all the nodes in the cluster.
- You must have a user with sysadmin permissions on the SQL Server.

  SnapCenter Plug-in for Microsoft SQL Server uses Microsoft VDI Framework, which requires sysadmin access.

  Microsoft Support Article 2926557: SQL Server VDI backup and restore operations require Sysadmin privileges

- When installing a plug-in on a Windows host, if you specify a credential that is not built-in or if the user belongs to a local workgroup user, you must disable UAC on the host.
- If SnapManager for Microsoft SQL Server is installed, you must have stopped or disabled the service and schedules.

  If you plan to import backup or clone jobs into SnapCenter, do not uninstall SnapManager for Microsoft SQL Server.

- The host must be resolvable to the fully qualified domain name (FQDN) from the server.

  If the hosts file is modified to make it resolvable and if both the short name and the FQDN are specified in the hosts file, create an entry in the SnapCenter hosts file in the following format: <ip_address>

<host_fqdn> <host_name>

Host requirements to install SnapCenter Plug-ins Package for Windows

Before you install the SnapCenter Plug-ins Package for Windows, you should be familiar with some basic host system space requirements and sizing requirements.

Item	Requirements
Operating Systems	Microsoft Windows For the latest information about supported versions, see the NetApp Interoperability Matrix Tool.
Minimum RAM for the SnapCenter plug-in on host	1 GB
Minimum install and log space for the SnapCenter plug-in on host	5 GB You should allocate sufficient disk space and monitor the storage consumption by the logs folder. The log space required varies depending on the number of the entities to be protected and the frequency of data protection operations. If there is no sufficient disk space, the logs will not be created for the recently run operations.
Required software packages	• ASP.NET Core Runtime 8.0.12 (and all subsequent 8.0.x patches) Hosting Bundle • PowerShell Core 7.4.2 • Java 11 Oracle Java and OpenJDK Java 11 Oracle Java and OpenJDK is required only for SAP HANA, IBM Db2, PostgreSQL, MySQL, NetApp supported plug-ins, and other custom applications which can be installed on Windows host. For the latest information about supported versions, see the NetApp Interoperability Matrix Tool. For .NET specific troubleshooting information, see SnapCenter upgrade or install fails for legacy systems that do not have internet connectivity.

Set up credentials for the SnapCenter Plug-ins Package for Windows

SnapCenter uses credentials to authenticate users for SnapCenter operations. You should create credentials for installing SnapCenter plug-ins and additional credentials for performing data protection operations on databases or Windows file systems.

Before you begin

- You must set up Windows credentials before installing plug-ins.
- You must set up the credentials with administrator privileges, including administrator rights on the remote host.
- SQL authentication on Windows hosts

  You must set up SQL credentials after installing plug-ins.

  If you are deploying SnapCenter Plug-in for Microsoft SQL Server, you must set up SQL credentials after installing plug-ins. Set up a credential for a user with SQL Server sysadmin permissions.

  The SQL authentication method authenticates against a SQL Server instance. This means that a SQL Server instance must be discovered in SnapCenter. Therefore, before adding a SQL credential, you must add a host, install plug-in packages, and refresh resources. You need SQL Server authentication for performing operations such as scheduling or discovering resources.

Steps

1. In the left navigation pane, click **Settings**.
2. In the Settings page, click **Credential**.
3. Click **New**.
4. In the Credential page, specify the information required for configuring credentials:

For this field…	Do this…
Credential name	Enter a name for the credential.

For this field…	Do this…
User name/Password	Enter the user name and password that are to be used for authentication. • Domain administrator Specify the domain administrator on the system on which you are installing the SnapCenter plug-in. Valid formats for the Username field are: ◦ `NetBIOS\UserName` ◦ `Domain FQDN\UserName` • Local administrator (for workgroups only) For systems that belong to a workgroup, specify the built-in local administrator on the system on which you are installing the SnapCenter plug-in. You can specify a local user account that belongs to the local administrators group if the user account has elevated privileges or the User Access control feature is disabled on the host system. The valid format for the Username field is: `UserName` Do not use double quotes (") or backtick (`) in the passwords. You should not use the less than (<) and exclamation (!) symbols together in passwords. For example, lessthan<!10, lessthan10<!, backtick`12.
Authentication Mode	Select the authentication mode that you want to use. If you select the SQL authentication mode, you must also specify the SQL server instance and the host where the SQL instance is located.

5. Click **OK**.

After you finish setting up the credentials, you might want to assign credential maintenance to a user or group of users in the User and Access page.

Configure credentials for an individual SQL Server resource

You can configure credentials to perform data protection jobs on individual SQL Server resource for each user. While you can configure the credentials globally, you might want to do this only for a particular resource.

About this task

- If you are using Windows credentials for authentication, you must set up your credential before installing plug-ins.

However, if you are using an SQL Server instance for authentication, you must add the credential after installing plug-ins.

- If you have enabled SQL authentication while setting up the credentials, the discovered instance or database is shown with a red color padlock icon.

  If the padlock icon appears, you must specify the instance or database credentials to successfully add the instance or database to a resource group.

- You must assign the credential to a role-based access control (RBAC) user without sysadmin access when the following conditions are met:
  - The credential is assigned to an SQL instance.
  - The SQL instance or host is assigned to an RBAC user.

    The user must have both the resource group and backup privileges.

Step 1: Add and configure credentials

1. In the left navigation pane, select **Settings**.
2. In the Settings page, select **Credential**.
   a. To add a new credential, select **New**.
   b. In the Credential page, configure the credentials:

For this field…	Do this…
Credential name	Enter a name for the credentials.
Username	Enter the user name used for SQL Server authentication. • Domain administrator or any member of the administrator group Specify the domain administrator or any member of the administrator group on the system on which you are installing the SnapCenter plug-in. Valid formats for the **Username** field are: ◦ *NetBIOS\UserName* ◦ *Domain FQDN\UserName* • Local administrator (for workgroups only) For systems that belong to a workgroup, specify the built-in local administrator on the system on which you are installing the SnapCenter plug-in. You can specify a local user account that belongs to the local administrators group if the user account has elevated privileges or the User Access control feature is disabled on the host system. The valid format for the **Username** field is: *UserName*
Password	Enter the password used for authentication.

For this field…	Do this…
Authentication mode	Select the SQL Server authentication mode. You can also choose Windows authentication if the Windows user has sysadmin privileges on the SQL server.
Host	Select the host.
SQL Server Instance	Select the SQL Server instance.

c. Select **OK** to add the credential.

Step 2: Configure instances

1. In the left navigation pane, select **Resources**.
2. In the Resources page, select **Instance** from the **View** list.
   a. Select image::../media/filter_icon.png[filter icon], and then select the host name to filter the instances.
   b. Select image::../media/filter_icon.png[filter icon] to close the filter pane.
3. In the Instance Protect page, protect the instance, and if required, select **Configure Credentials**.

   If the user who is logged in to the SnapCenter Server does not have access to SnapCenter Plugin for Microsoft SQL Server, then the user has to configure the credentials.

   The credential option does not apply to databases and availability groups.

4. Select **Refresh Resources**.

Configure gMSA on Windows Server 2016 or later

Windows Server 2016 or later enables you to create a group Managed Service Account (gMSA) that provides automated service account password management from a managed domain account.

Before you begin

- You should have a Windows Server 2016 or later domain controller.
- You should have a Windows Server 2016 or later host, which is a member of the domain.

Steps

1. Create a KDS root key to generate unique passwords for each object in your gMSA.
2. For each domain, run the following command from the Windows domain controller: Add-KDSRootKey -EffectiveImmediately
3. Create and configure your gMSA:
   a. Create a user group account in the following format:

```
domainName\accountName$
```

b. Add computer objects to the group.

c. Use the user group you just created to create the gMSA.

For example,

```
New-ADServiceAccount -name <ServiceAccountName> -DNSHostName <fqdn>
-PrincipalsAllowedToRetrieveManagedPassword <group>
-ServicePrincipalNames <SPN1,SPN2,…>
```

d. Run `Get-ADServiceAccount` command to verify the service account.

4. Configure the gMSA on your hosts:

a. Enable the Active Directory module for Windows PowerShell on the host where you want to use the gMSA account.

To do this, run the following command from PowerShell:

```
PS C:\> Get-WindowsFeature AD-Domain-Services

Display Name                            Name                Install
State
------------                            ----
-------------
[ ] Active Directory Domain Services   AD-Domain-Services  Available


PS C:\> Install-WindowsFeature AD-DOMAIN-SERVICES

Success Restart Needed Exit Code      Feature Result
------- -------------- ---------      --------------
True    No             Success        {Active Directory Domain
Services, Active ...
WARNING: Windows automatic updating is not enabled. To ensure that
your newly-installed role or feature is
automatically updated, turn on Windows Update.
```

b. Restart your host.

c. Install the gMSA on your host by running the following command from the PowerShell command prompt: `Install-AdServiceAccount <gMSA>`

d. Verify your gMSA account by running the following command: `Test-AdServiceAccount <gMSA>`

5. Assign the administrative privileges to the configured gMSA on the host.

6. Add the Windows host by specifying the configured gMSA account in the SnapCenter Server.

SnapCenter Server will install the selected plug-ins on the host and the specified gMSA will be used as the service log on account during the plug-in installation.

Install SnapCenter Plug-in for Microsoft SQL Server

Add hosts and install the SnapCenter Plug-ins Package for Windows

You must use the SnapCenter **Add Host** page to add hosts and install the plug-ins package. The plug-ins are automatically installed on the remote hosts.

Before you begin

- If the operating system of the SnapCenter Server host is Windows 2019 and the operating system of the plug-in host is Windows 2022, you should perform the following:
  - Upgrade to Windows Server 2019 (OS Build 17763.5936) or later
  - Upgrade to Windows Server 2022 (OS Build 20348.2402) or later
- You must be a user that is assigned to a role that has the plug-in install and uninstall permissions, such as the SnapCenter Admin role.
- When installing a plug-in on a Windows host, if you specify a credential that is not built-in, you should disable UAC on the host.
- You should ensure that the message queueing service is in running state.
- If you are using group Managed Service Account (gMSA), you should configure gMSA with administrative privileges.

  Configure group Managed Service Account on Windows Server 2016 or later for SQL

About this task

You cannot add a SnapCenter Server as a plug-in host to another SnapCenter Server.

You can add a host and install the plug-in packages either for an individual host or for a cluster. If you are installing the plug-ins on a cluster or Windows Server Failover Clustering (WSFC), the plug-ins are installed on all of the nodes of the cluster.

For information on managing hosts, see Manage hosts.

Steps

1. In the left navigation pane, select **Hosts**.
2. Verify that the **Managed Hosts** tab is selected at the top.
3. Select **Add**.
4. In the Hosts page do the following:

For this field…	Do this…
Host Type	Select Windows as the host type. The SnapCenter Server adds the host, and then installs the Plug-in for Windows if the plug-in is not already installed on the host. If you select the Microsoft SQL Server option on the Plug-ins page, the SnapCenter Server installs the Plug-in for SQL Server.

For this field…	Do this…
Host name	Enter the fully qualified domain name (FQDN) or the IP address of the host. IP address is supported for untrusted domain hosts only if it resolves to the FQDN. SnapCenter depends on the proper configuration of the DNS. Therefore, the best practice is to enter the FQDN. You can enter the IP addresses or FQDN of one of the following: • Stand-alone host • WSFC If you are adding a host by using SnapCenter and the host is part of a subdomain, you must provide the FQDN.
Credentials	Select the credential name that you created or create new credentials. The credential must have administrative rights on the remote host. For details, see the information about creating credentials. You can view details about the credentials by positioning your cursor over the credential name that you specified. The credentials authentication mode is determined by the host type that you specify in the Add Host wizard.

5. In the **Select Plug-ins to Install** section, select the plug-ins to install.
6. Select **More Options**.

For this field…	Do this…
Port	Either retain the default port number or specify the port number. The default port number is 8145. If the SnapCenter Server was installed on a custom port, that port number will be displayed as the default port. If you manually installed the plug-ins and specified a custom port, you must specify the same port. Otherwise, the operation fails.

For this field…	Do this…
Installation Path	The default path is C:\Program Files\NetApp\SnapCenter. You can optionally customize the path.
Add all hosts in the cluster	Select this check box to add all of the cluster nodes in a WSFC or a SQL Availability Group. You should add all the cluster nodes by selecting the appropriate cluster check box in the GUI if you want to manage and identify multiple available SQL Availability Groups within a cluster.
Skip preinstall checks	Select this check box if you already installed the plug-ins manually and you do not want to validate whether the host meets the requirements for installing the plug-in.
Use group Managed Service Account (gMSA) to run the plug-in services	Select this check box if you want to use group Managed Service Account (gMSA) to run the plug-in services. Provide the gMSA name in the following format: domainName\accountName$. If the host is added with gMSA and if the gMSA has login and sys admin privileges, the gMSA will be used to connect to the SQL instance.

7. Select **Submit**.
8. For SQL Plug-in, select the host to configure the log directory.
   a. Select **Configure log directory** and in the Configure host log directory page, select **Browse** and complete the following steps:

      Only NetApp LUNs (drives) are listed for selection. SnapCenter backs up and replicates the host log directory as part of the backup operation.

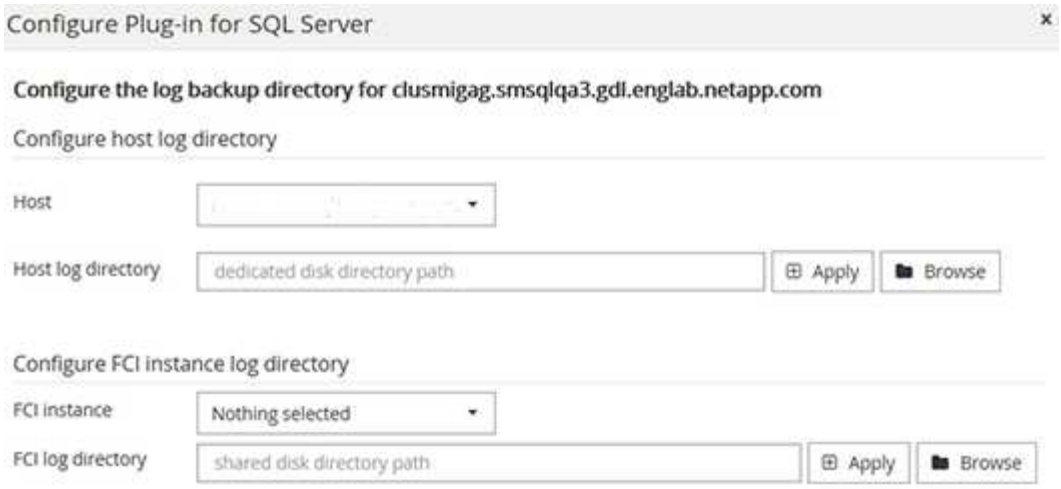

i. Select the drive letter or mount point on the host where the host log will be stored.

ii. Choose a subdirectory, if required.

iii. Select **Save**.

9. Select **Submit**.

   If you have not selected the **Skip prechecks** check box, the host is validated to verify whether it meets the requirements for installing the plug-in. The disk space, RAM, PowerShell version, .NET version, location (for Windows plug-ins), and Java version (for Linux plug-ins) are validated against the minimum requirements. If the minimum requirements are not met, appropriate error or warning messages are displayed.

   If the error is related to disk space or RAM, you can update the web.config file located at C:\Program Files\NetApp\SnapCenter WebApp to modify the default values. If the error is related to other parameters, you must fix the issue.

   In an HA setup, if you are updating web.config file, you must update the file on both nodes.

10. Monitor the installation progress.

Install SnapCenter Plug-in for Microsoft SQL Server on multiple remote hosts by using cmdlets

You can install the SnapCenter Plug-in for Microsoft SQL Server on multiple hosts simultaneously by using the Install-SmHostPackage PowerShell cmdlet.

Before you begin

You must have logged in to SnapCenter as a domain user with local administrator rights on each host on which you want to install the plug-in package.

Steps

1. Launch PowerShell.
2. On the SnapCenter Server host, establish a session using the Open-SmConnection cmdlet, and then enter your credentials.
3. Install the SnapCenter Plug-in for Microsoft SQL Server on multiple remote hosts using the Install-SmHostPackage cmdlet and the required parameters.

   The information regarding the parameters that can be used with the cmdlet and their descriptions can be

obtained by running *Get-Help command_name*. Alternatively, you can also refer to the SnapCenter Software Cmdlet Reference Guide.

You can use the -skipprecheck option when you have already installed the plug-ins manually and you do not want to validate whether the host meets the requirements for installing the plug-in.

4. Enter your credentials for remote installation.

Install the SnapCenter Plug-in for Microsoft SQL Server silently from the command line

You should install SnapCenter Plug-in for Microsoft SQL Server from within the SnapCenter user interface. However, if you cannot for some reason, you can run the Plug-in for SQL Server installation program unattended in silent mode from the Windows command line.

Before you begin

- You must delete the earlier version of SnapCenter Plug-in for Microsoft SQL Server before installing.

  For more information, see How to Install a SnapCenter Plug-In manually and directly from the Plug-In Host.

Steps

1. Validate whether C:\temp folder exists on the plug-in host and the logged in user has full access to it.
2. Download the Plug-in for SQL Server software from C:\ProgramData\NetApp\SnapCenter\Package Repository.

   This path is accessible from the host where the SnapCenter Server is installed.

3. Copy the installation file to the host on which you want to install the plug-in.
4. From a Windows command prompt on the local host, navigate to the directory to which you saved the plug-in installation files.
5. Install the Plug-in for SQL Server software:

   ```
   "snapcenter_windows_host_plugin.exe"/silent /debuglog"Debug_Log_Path"
   /log"Log_Path" BI_SNAPCENTER_PORT=Num
   SUITE_INSTALLDIR="Install_Directory_Path"
   BI_SERVICEACCOUNT=domain\\administrator BI_SERVICEPWD=password
   ISFeatureInstall=SCW,SCSQL
   ```

   Replace the placeholder values with your data

   - Debug_Log_Path is the name and location of the suite installer log file.
   - Log_Path is the location of the installation logs of the plug-in components (SCW, SCSQL, and SMCore).
   - Num is the port on which SnapCenter communicates with SMCore
   - Install_Directory_Path is the host plug-in package installation directory.
   - domain\administrator is the SnapCenter Plug-in for Microsoft Windows web service account.
   - password is the password for the SnapCenter Plug-in for Microsoft Windows web service account.

     ```
     "snapcenter_windows_host_plugin.exe"/silent
     ```

```
/debuglog"C:\HPPW_SCSQL_Install.log" /log"C:\" BI_SNAPCENTER_PORT=8145
SUITE_INSTALLDIR="C:\Program Files\NetApp\SnapCenter"
BI_SERVICEACCOUNT=domain\administrator BI_SERVICEPWD=password
ISFeatureInstall=SCW,SCSQL
```

All the parameters passed during the installation of Plug-in for SQL Server are case sensitive.

6. Monitor the Windows task scheduler, the main installation log file C:\Installdebug.log, and the additional installation files in C:\Temp.
7. Monitor the %temp% directory to verify that the msiexe.exe installers are installing the software without errors.

The installation of Plug-in for SQL Server registers the plug-in on the host and not on the SnapCenter Server. You can register the plug-in on the SnapCenter Server by adding the host using the SnapCenter GUI or PowerShell cmdlet. After the host is added, the plug-in is automatically discovered.

Monitor the status of installing Plug-in for SQL Server

You can monitor the progress of SnapCenter plug-in package installation by using the Jobs page. You might want to check the progress of installation to determine when it is complete or if there is an issue.

About this task

The following icons appear on the Jobs page and indicate the state of the operation:

- In progress
- Completed successfully
- Failed
- Completed with warnings or could not start due to warnings
- Queued

Steps

1. In the left navigation pane, click **Monitor**.
2. In the **Monitor** page, click **Jobs**.
3. In the **Jobs** page, to filter the list so that only plug-in installation operations are listed, do the following:
   a. Click **Filter**.
   b. Optional: Specify the start and end date.
   c. From the Type drop-down menu, select **Plug-in installation**.
   d. From the Status drop-down menu, select the installation status.
   e. Click **Apply**.
4. Select the installation job and click **Details** to view the job details.
5. In the **Job Details** page, click **View logs**.

Configure CA Certificate

Generate CA Certificate CSR file

You can generate a Certificate Signing Request (CSR) and import the certificate that can be obtained from a Certificate Authority (CA) using the generated CSR. The certificate will have a private key associated with it.

CSR is a block of encoded text that is given to an authorized certificate vendor to procure the signed CA certificate.

CA Certificate RSA key length must be minimum 3072 bits.

For information to generate a CSR, see How to generate CA Certificate CSR file.

If you own the CA certificate for your domain (*.domain.company.com) or your system (machine1.domain.company.com), you can skip generating the CA Certificate CSR file. You can deploy the existing CA certificate with SnapCenter.

For cluster configurations, the cluster name (virtual cluster FQDN), and the respective host names should be mentioned in the CA certificate. The certificate can be updated by filling the Subject Alternative Name (SAN) field before procuring the certificate. For a wild card certificate (*.domain.company.com), the certificate will contain all the hostnames of the domain implicitly.

Import CA certificates

You must import the CA certificates to the SnapCenter Server and the Windows host plug-ins using the Microsoft management console (MMC).

Steps

1. Go to the Microsoft management console (MMC), and then click **File** > **Add/Remove Snapin**.
2. In the Add or Remove Snap-ins window, select **Certificates** and then click **Add**.
3. In the Certificates snap-in window, select the **Computer account** option, and then click **Finish**.
4. Click **Console Root** > **Certificates – Local Computer** > **Trusted Root Certification Authorities** > **Certificates**.
5. Right-click on the folder “Trusted Root Certification Authorities”, and then select **All Tasks** > **Import** to start the import wizard.
6. Complete the wizard, as follows:

In this wizard window…	Do the following…
Import Private Key	Select the option **Yes**, import the private key, and then click **Next**.
Import File Format	Make no changes; click **Next**.
Security	Specify the new password to be used for the exported certificate, and then click **Next**.

In this wizard window…	Do the following…
Completing the Certificate Import Wizard	Review the summary, and then click **Finish** to start the import.

Importing certificate should be bundled with the private key (supported formats are: *.pfx, *.p12, and *.p7b).

7. Repeat Step 5 for the “Personal” folder.

Get the CA certificate thumbprint

A certificate thumbprint is a hexadecimal string that identifies a certificate. A thumbprint is calculated from the content of the certificate using a thumbprint algorithm.

Steps

1. Perform the following on the GUI:
   a. Double-click the certificate.
   b. In the Certificate dialog box, click the **Details** tab.
   c. Scroll through the list of fields and click **Thumbprint**.
   d. Copy the hexadecimal characters from the box.
   e. Remove the spaces between the hexadecimal numbers.

      For example, if the thumbprint is: "a9 09 50 2d d8 2a e4 14 33 e6 f8 38 86 b0 0d 42 77 a3 2a 7b", after removing the spaces, it will be: "a909502dd82ae41433e6f83886b00d4277a32a7b".

2. Perform the following from PowerShell:
   a. Run the following command to list the thumbprint of the installed certificate and identify the recently installed certificate by the subject name.

      Get-ChildItem -Path Cert:\LocalMachine\My

   b. Copy the thumbprint.

Configure CA certificate with Windows host plug-in services

You should configure the CA certificate with Windows host plug-in services to activate the installed digital certificate.

Perform the following steps on the SnapCenter Server and all the plug-in hosts where CA certificates are already deployed.

Steps

1. Remove the existing certificate binding with SMCore default port 8145, by running the following command:

   ```
   > netsh http delete sslcert ipport=0.0.0.0:_<SMCore Port>
   ```

   For example:

```
> netsh http delete sslcert ipport=0.0.0.0:8145
```

2. Bind the newly installed certificate with the Windows host plug-in services, by running the following commands:

```
> $cert = "_<certificate thumbprint>_"
> $guid = [guid]::NewGuid().ToString("B")
> netsh http add sslcert ipport=0.0.0.0: _<SMCore Port>_ certhash=$cert
appid="$guid"
```

For example:

```
> $cert = "a909502dd82ae41433e6f83886b00d4277a32a7b"
> $guid = [guid]::NewGuid().ToString("B")
> netsh http add sslcert ipport=0.0.0.0: _<SMCore Port>_ certhash=$cert
appid="$guid"
```

Enable CA Certificates for plug-ins

You should configure the CA certificates and deploy the CA certificates in the SnapCenter Server and the corresponding plug-in hosts. You should enable the CA certificate validation for the plug-ins.

Before you begin

- You can enable or disable the CA certificates using the run *Set-SmCertificateSettings* cmdlet.
- You can display the certificate status for the plug-ins using the *Get-SmCertificateSettings*.

The information regarding the parameters that can be used with the cmdlet and their descriptions can be obtained by running *Get-Help command_name*. Alternatively, you can also refer to the SnapCenter Software Cmdlet Reference Guide.

Steps

1. In the left navigation pane, click **Hosts**.
2. In the Hosts page, click **Managed Hosts**.
3. Select single or multiple plug-in hosts.
4. Click **More options**.
5. Select **Enable Certificate Validation**.

After you finish

The Managed Hosts tab host displays a padlock and the color of the padlock indicates the status of the connection between SnapCenter Server and the plug-in host.

- 🔒 indicates that the CA certificate is neither enabled nor assigned to the plug-in host.

- indicates that the CA certificate is successfully validated.
- indicates that the CA certificate could not be validated.
- indicates that the connection information could not be retrieved.

When the status is yellow or green, the data protection operations completes successfully.

Configure Disaster recovery

Disaster recovery of SnapCenter Plug-in for SQL Server

When the SnapCenter Plug-in for SQL Server is down, use the following steps to switch to a different SQL host and recover the data.

Before you begin

- The secondary host should have the same operating system, application, and host name as the primary host.
- Push the SnapCenter Plug-in for SQL Server to an alternative host using **Add host** or **Modify host** page. See Manage hosts for more information.

Steps

1. Select the host from the **Hosts** page to modify and install the SnapCenter Plug-in for SQL Server.
2. (Optional) Replace the SnapCenter Plug-in for SQL Server configuration files from disaster recovery (DR) backup to the new machine.
3. Import Windows and SQL schedules from the SnapCenter Plug-in for SQL Server folder from the DR backup.

Related information

See the Disaster Recovery APIs video.

Storage disaster recovery (DR) for SnapCenter Plug-in for SQL Server

You can recover the SnapCenter Plug-in for SQL Server storage by enabling the DR Mode for Storage in the Global Settings page.

Before you begin

- Ensure that the plug-ins are in maintenance mode.
- Break the SnapMirror/SnapVault relationship.
  Breaking SnapMirror relationships
- Attach the LUN from secondary to the host machine with same drive letter.
- Ensure that all the disks are connected using the same drive letters that were used prior to DR.
- Restart MSSQL server service.
- Ensure that the SQL resources are back online.

About this task

Disaster recovery (DR) is not supported on VMDK and RDM configurations.

Steps

1. In the Settings page, navigate to **Settings** > **Global Settings** > **Disaster Recovery**.
2. Select **Enable Disaster Recovery**.
3. Click **Apply**.
4. Verify whether the DR job is enabled or not by clicking **Monitor** > **Jobs**.

After you finish

- If new databases are created after the failover, the databases will be in non-DR mode.

  The new databases will continue to operate like they did before the failover.

- The new backups that were created in DR mode will be listed under SnapMirror or SnapVault (secondary) in the Topology page.

  An "i" icon is displayed next to the new backups to indicate that these backups were created during DR mode.

- You can delete the SnapCenter Plug-in for SQL Server backups that were created during failover either by using the UI or the following cmdlet: `Remove-SmBackup`
- After failover, if you want some of the resources to be in non-DR mode, use the following cmdlet: `Remove-SmResourceDRMode`

  For more information refer to the SnapCenter Software Cmdlet Reference Guide.

- SnapCenter Server will manage the individual storage resources(SQL databases) that are in DR or non-DR mode but not the resource group with storage resources that are in DR mode or non-DR mode.

Failback from SnapCenter Plug-in for SQL Server secondary storage to primary storage

After the SnapCenter Plug-in for SQL Server primary storage is back online, you should failback to the primary storage.

Before you begin

- Place the SnapCenter Plug-in for SQL Server in **Maintenance** mode from the Managed Hosts page.
- Disconnect the secondary storage from the host and connect from the primary storage.
- To failback to the primary storage, ensure that the relationship direction remains the same as it was before the failover by performing the reverse resync operation.

  To retain the roles of primary and secondary storage after the reverse resync operation, perform the reverse resync operation once again.

  For more information see Reverse resynchronizing mirror relationships

- Restart MSSQL server service.
- Ensure that the SQL resources are back online.

During failover or failback of the plug-in, the plug-in overall status is not refreshed immediately. The host and plug-in overall status is updated during the subsequent host refresh operation.

Steps

1. In the Settings page, navigate to **Settings** > **Global Settings** > **Disaster Recovery**.
2. Unselect **Enable Disaster Recovery**.
3. Click **Apply**.
4. Verify whether the DR job is enabled or not by clicking **Monitor** > **Jobs**.

After you finish

You can delete the SnapCenter Plug-in for SQL Server backups that were created during failover either by using the UI or the following cmdlet: `Remove-SmDRFailoverBackups`

Install SnapCenter Plug-in for VMware vSphere

If your database or filesystem is stored on virtual machines (VMs), or if you want to protect VMs and datastores, you must deploy the SnapCenter Plug-in for VMware vSphere virtual appliance.

For information to deploy, see Deployment Overview.

Deploy CA certificate

To configure the CA Certificate with SnapCenter Plug-in for VMware vSphere, see Create or import SSL certificate.

Configure the CRL file

SnapCenter Plug-in for VMware vSphere looks for the CRL files in a pre-configured directory. Default directory of the CRL files for SnapCenter Plug-in for VMware vSphere is */opt/netapp/config/crl*.

You can place more than one CRL file in this directory. The incoming certificates will be verified against each CRL.

Prepare for data protection

Prerequisites for using SnapCenter Plug-in for Microsoft SQL Server

Before you begin to use the Plug-in for SQL Server, the SnapCenter administrator must install and configure SnapCenter Server and perform prerequisite tasks.

- Install and configure SnapCenter Server.
- Log in to SnapCenter.
- Configure the SnapCenter environment by adding or assigning storage system connections and creating credentials.

  SnapCenter does not support multiple SVMs with the same name on different clusters. Each SVM supported by SnapCenter must have a unique name.

- Add hosts, install the plug-ins, discover (refresh) the resources, and configure the plug-ins.
- Move an existing Microsoft SQL Server database from a local disk to a NetApp LUN or vice versa by running Invoke-SmConfigureResources.

For information to run the cmdlet, see the SnapCenter Software Cmdlet Reference Guide

- If you are using SnapCenter Server to protect SQL databases that reside on VMware RDM LUNs or VMDKs, you must deploy the SnapCenter Plug-in for VMware vSphere and register the plug-in with SnapCenter. The SnapCenter Plug-in for VMware vSphere documentation has more information.

  SnapCenter Plug-in for VMware vSphere documentation

- Perform host-side storage provisioning using the SnapCenter Plug-in for Microsoft Windows.
- Set up SnapMirror and SnapVault relationships, if you want backup replication.

For details, see SnapCenter installation information.

For SnapCenter 4.1.1 users, the SnapCenter Plug-in for VMware vSphere 4.1.1 documentation has information on protecting virtualized databases and file systems. For SnapCenter 4.2.x users, the NetApp Data Broker 1.0 and 1.0.1, documentation has information on protecting virtualized databases and file systems using the SnapCenter Plug-in for VMware vSphere that is provided by the Linux-based NetApp Data Broker virtual appliance (Open Virtual Appliance format). For SnapCenter 4.3.x users, the SnapCenter Plug-in for VMware vSphere 4.3 documentation has information on protecting virtualized databases and file systems using the Linux-based SnapCenter Plug-in for VMware vSphere virtual appliance (Open Virtual Appliance format).

SnapCenter Plug-in for VMware vSphere documentation

How resources, resource groups, and policies are used for protecting SQL Server

Before you use SnapCenter, it is helpful to understand basic concepts related to the backup, clone, and restore operations you want to perform. You interact with resources, resource groups, and policies for different operations.

- Resources are typically databases, database instances, or Microsoft SQL Server availability groups that you back up or clone with SnapCenter.
- A SnapCenter resource group, is a collection of resources on a host or cluster.

  When you perform an operation on a resource group, you perform that operation on the resources defined in the resource group according to the schedule you specify for the resource group.

  You can back up on demand a single resource or a resource group. You also can perform scheduled backups for single resources and resource groups.

- The policies specify the backup frequency, copy retention, replication, scripts, and other characteristics of data protection operations.

  When you create a resource group, you select one or more policies for that group. You can also select a policy when you perform a backup on demand for a single resource.

Think of a resource group as defining *what* you want to protect and when you want to protect it in terms of day and time. Think of a policy as defining *how* you want to protect it. If you are backing up all databases or backing up all file systems of a host, for example, you might create a resource group that includes all the databases or all the file systems in the host. You could then attach two policies to the resource group: a daily policy and an hourly policy. When you create the resource group and attach the policies, you might configure the resource group to perform a full backup daily and another schedule that performs log backups hourly.

The following image illustrates the relationship between resources, resource groups, and policies for databases:

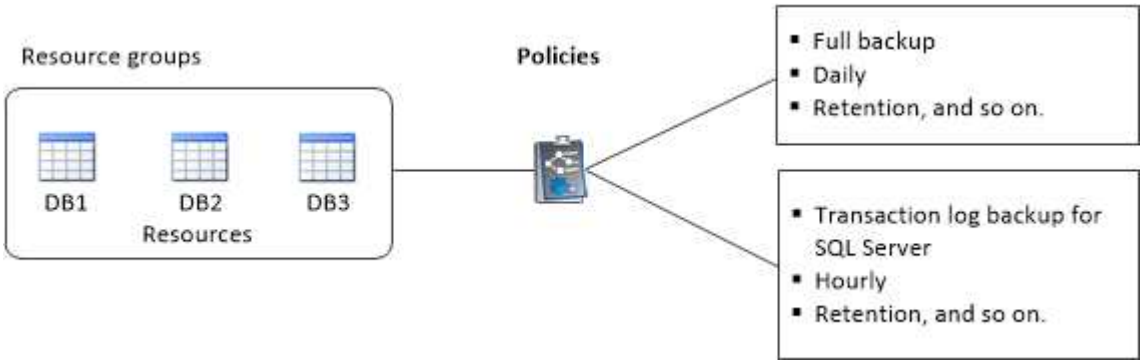

Back up SQL Server database, or instance, or availability group

Backup workflow

When you install the SnapCenter Plug-in for Microsoft SQL Server in your environment, you can use SnapCenter to back up the SQL Server resources.

You can schedule multiple backups to run across servers simultaneously.

Backup and restore operations cannot be performed simultaneously on the same resource.

The following workflow shows the sequence in which you must perform the backup operations:

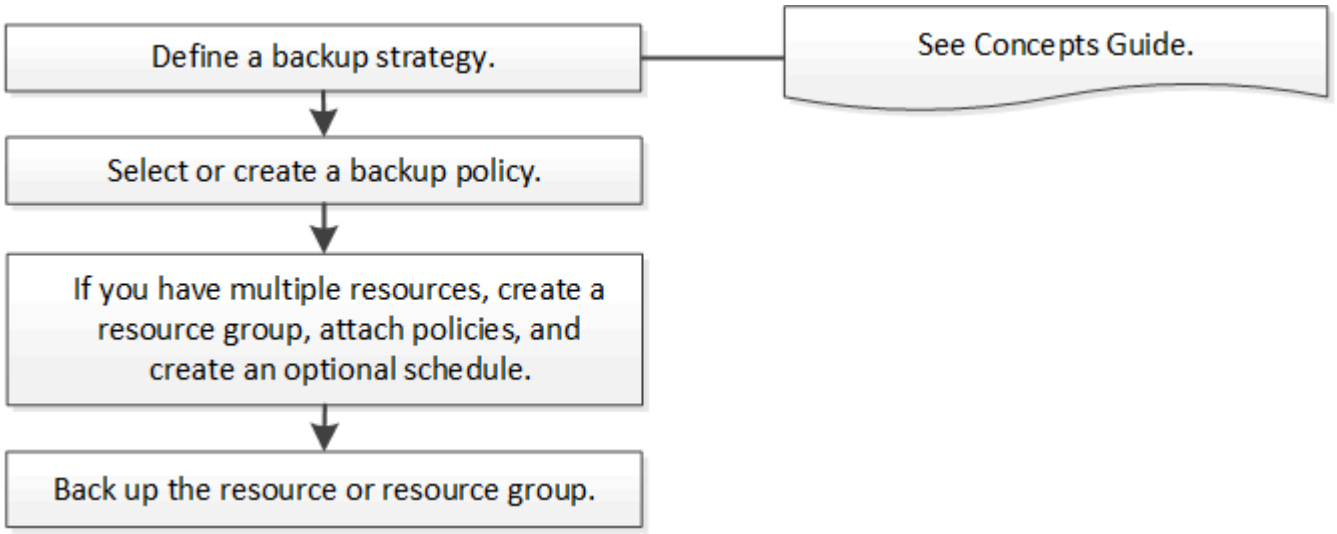

The Backup Now, Restore, Manage Backups, and Clone options on the Resources page are disabled if you select a non-NetApp LUN, a database that is corrupted, or a database that is being restored.

You can also use PowerShell cmdlets manually or in scripts to perform backup, restore, recovery, verify, and clone operations. For detailed information about PowerShell cmdlets, use the SnapCenter cmdlet help or see the SnapCenter Software Cmdlet Reference Guide

How SnapCenter backs up databases

SnapCenter uses Snapshot technology to back up the SQL Server databases that reside on LUNs or VMDKs. SnapCenter creates the backup by creating Snapshots of the databases.

When you select a database for a full database backup from the Resources page, SnapCenter automatically selects all the other databases that reside on the same storage volume. If the LUN or VMDK stores only a single database, you can clear or reselect the database individually. If the LUN or VMDK houses multiple databases, you must clear or reselect the databases as a group.

All the databases that reside on a single volume are backed up concurrently using Snapshots. If the maximum number of concurrent backup databases is 35, and if more than 35 databases reside in a storage volume, then the total number of Snapshots that are created equals the number of databases divided by 35.

You can configure the maximum number of databases for each Snapshot in the backup policy.

When SnapCenter creates a Snapshot, the entire storage system volume is captured in the Snapshot. However, the backup is valid only for the SQL host server for which the backup was created.

If data from other SQL host servers resides on the same volume, this data cannot be restored from the Snapshot.

Find more information

Quiesce or grouping resources operations fail

Determine whether resources are available for backup

Resources are the databases, application instances, Availability Groups, and similar components that are maintained by the plug-ins you have installed. You can add those resources to resource groups so that you can perform data protection jobs, but first you must identify which resources you have available. Determining available resources also verifies that the plug-in installation has completed successfully.

Before you begin

- You must have already completed tasks such as installing SnapCenter Server, adding hosts, creating storage system connections, and adding credentials.
- To discover the Microsoft SQL databases, one of the following conditions should be met.
  - The user that was used to add the plug-in host to SnapCenter Server should have the required permissions (sysadmin) on the Microsoft SQL Server.
  - If the above condition is not met, in the SnapCenter Server you should configure the user that has the required permissions (sysadmin) on the Microsoft SQL Server. The user should be configured at the Microsoft SQL Server instance level and the user can be a SQL or Windows user.
- To discover the Microsoft SQL databases in a Windows cluster, you must unblock the Failover Cluster Instance (FCI) TCP/IP port.
- If databases reside on VMware RDM LUNs or VMDKs, you must deploy the SnapCenter Plug-in for VMware vSphere and register the plug-in with SnapCenter.

  For more information, see Deploy SnapCenter Plug-in for VMware vSphere

- If the host is added with gMSA and if the gMSA has login and system admin privileges, the gMSA will be used to connect to the SQL instance.

About this task

You cannot back up databases when the **Overall Status** option in the Details page is set to Not available for backup. The **Overall Status** option is set to Not available for backup when any of the following is true:

- Databases are not on a NetApp LUN.
- Databases are not in normal state.

  Databases are not in normal state when they are offline, restoring, recovery pending, suspect, and so on.

- Databases have insufficient privileges.

  For example, if a user has only view access to the database, files and properties of the database cannot be identified and hence cannot be backed up.

SnapCenter can backup only the primary database if you have a availability group configuration on SQL Server Standard Edition.

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page select **Database**, or **Instance**, or **Availability Group**, from the **View** drop-down list.

   Click and select the host name and the SQL Server Instance to filter the resources. You can then click to close the filter pane.

3. Click **Refresh Resources**.

   The newly added, renamed, or deleted resources are updated to the SnapCenter Server inventory.

   You must refresh the resources if the databases are renamed outside of SnapCenter.

   The resources are displayed along with information such as resource type, host or cluster name, associated resource groups, backup type, policies and overall status.

   - If the database is on a non NetApp storage, `Not available for backup` is displayed in the **Overall Status** column.

     You cannot perform data protection operations on a database that is on a non NetApp storage.

   - If the database is on a NetApp storage and not protected, `Not protected` is displayed in the **Overall Status** column.
   - If the database is on a NetApp storage system and protected, the user interface displays `Backup not run` message in the **Overall Status** column.
   - If the database is on a NetApp storage system and protected and if the backup is triggered for the database, the user interface displays `Backup succeeded` message in the **Overall Status** column.

If you have enabled an SQL authentication while setting up the credentials, the discovered instance or database is shown with a red padlock icon. If the padlock icon appears, you must specify the instance or database credentials for successfully adding the instance or database to a resource group.

4. After the SnapCenter administrator assigns the resources to a RBAC user, the RBAC user must log in and click **Refresh Resources** to see the latest **Overall Status** of the resources.

Migrate resources to NetApp storage system

After you have provisioned your NetApp storage system using SnapCenter Plug-in for Microsoft Windows, you can migrate your resources to the NetApp storage system or from one NetApp LUN to another NetApp LUN using either the SnapCenter graphical user interface (GUI) or using the PowerShell cmdlets.

Before you begin

- You must have added storage systems to SnapCenter Server.
- You must have refreshed (discovered) the SQL Server resources.

Most of the fields on these wizard pages are self-explanatory. The following information describes some of the fields for which you might require guidance.

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page, select **Database** or **Instance** from the **View** drop-down list.
3. Select either the database or the instance from the list and click **Migrate**.
4. In the Resources page, perform the following actions:

For this field…	Do this…
Database Name (optional)	If you have selected an instance for migration, you must select the databases of that instance from the **Databases** drop-down list.
Choose Destinations	Select the target location for data and log files. The data and log files are moved to Data and Log folder respectively under the selected NetApp drive. If any folder in the folder structure is not present, then a folder is created, and the resource is migrated.
Show database file details (optional)	Select this option when you want to migrate multiple files of a single database. This option is not displayed when you select the **Instance** resource.

For this field…	Do this…
Options	Select **Delete copy of Migrated Database at Original Location** to delete copy of database from the source. Optional: **RUN UPDATE STATISTICS on tables before detaching the database**.

5. In the Verify page, perform the following actions:

For this field…	Do this…
Database Consistency Check Options	Select **Run before** to check the integrity of the database before migration. Select **Run after** to check the integrity of the database after migration.
DBCC CHECKDB options	• Select **PHYSICAL_ONLY** option to limit the integrity check to the physical structure of the database and to detect torn pages, checksum failures, and common hardware failures that impact the database. • Select **NO_INFOMSGS** option to suppress all of the informational messages. • Select **ALL_ERRORMSGS** option to display all of the reported errors per object. • Select **NOINDEX** option if you do not want to check nonclustered indexes. The SQL Server database uses Microsoft SQL Server Database Consistency Checker (DBCC) to check the logical and physical integrity of the objects in the database. ⓘ You might want to select this option to decrease the execution time. • Select **TABLOCK** option to limit the checks and obtain locks instead of using an internal database Snapshot.

6. Review the summary, and then click **Finish**.

Create backup policies for SQL Server databases

You can create a backup policy for the resource or the resource group before you use SnapCenter to back up SQL Server resources, or you can create a backup policy at the time you create a resource group or backup a single resource.

Before you begin

- You must have defined your data protection strategy.
- You must have prepared for data protection by completing tasks such as installing SnapCenter, adding hosts, identifying resources, and creating storage system connections.
- You must have configured the host log directory for log backup.
- You must have refreshed (discovered) the SQL Server resources.
- If you are replicating Snapshots to a mirror or vault, the SnapCenter administrator must have assigned the storage virtual machines (SVMs) for both the source volumes and destination volumes to you.

  For information about how administrators assign resources to users, see the SnapCenter installation information.

- If you want to run the PowerShell scripts in prescripts and postscripts, you should set the value of the usePowershellProcessforScripts parameter to true in the web.config file.

  The default value is false.

- Review the SnapMirror active sync specific prerequisites and limitations. For information, refer Object limits for SnapMirror active sync.

About this task

- A backup policy is a set of rules that governs how you manage and retain backups, and how frequently the resource or resource group is backed up. Additionally, you can specify replication and script settings. Specifying options in a policy saves time when you want to reuse the policy for another resource group.

  The SCRIPTS_PATH is defined using the PredefinedWindowsScriptsDirectory key located in the SMCoreServiceHost.exe.Config file of the plug-in host.

  If needed, you can change this path and restart SMcore service. It is recommended that you use the default path for security.

  The value of the key can be displayed from swagger through the API: API /4.7/configsettings

  You can use the GET API to display the value of the key. SET API is not supported.

- SnapLock
  - If 'Retain the backup copies for a specific number of days' option is selected, then the SnapLock retention period must be lesser than or equal to the mentioned retention days.

    Specifying a Snapshot locking period prevents deletion of the Snapshots until the retention period expires. This could lead to retaining a larger number of Snapshots than the count specified in the policy.

    For ONTAP 9.12.1 and below version, the clones created from the SnapLock Vault Snapshots as part of restore will inherit the SnapLock Vault expiry time. Storage admin should manually cleanup the clones post the SnapLock expiry time.

Step 1: Create Policy Name

1. In the left navigation pane, select **Settings**.
2. In the Settings page, select **Policies**.

3. Select **New**.
4. In the **Name** page, enter the policy name and details.

Step 2: Configure policy options

1. In the Policy type page, perform the following steps:
   a. Select your storage type.
   b. Select your policy scope.

Full Backup and Log Backup

Back up the database files and transaction logs and to truncate the transaction logs.

1. Select **Full backup and Log backup**.
2. Enter the maximum number of databases that should be backed up for each Snapshot.

> You must increase this value if you want to run multiple backup operations concurrently.

Full Backup

Back up the database files.

1. Select **Full backup**.
2. Enter the maximum number of databases that should be backed up for each Snapshot. Default value is 100

> You must increase this value if you want to run multiple backup operations concurrently.

Log Backup

1. Back up the transaction logs.
2. Select **Log backup**.

Copy Only Backup

1. If you are backing up your resources by using another backup application, select **Copy only backup**.

Keeping the transaction logs intact allows any backup application to restore the databases. You typically should not use the copy only option in any other circumstance.

> Microsoft SQL does not support the **Copy only backup** option together with the **Full backup and Log backup** option for secondary storage.

Step 3: Configure Availability Group settings

1. In the Availability Group Settings section, perform the following actions:

a. Backup on preferred backup replica only.

   Select this option to backup only on preferred backup replica. The preferred backup replica is decided by the backup preferences configured for the AG in the SQL Server.

b. Select replicas for backup.

   Choose the primary AG replica or the secondary AG replica for the backup.

c. Select Backup priority (Minimum and Maximum backup priority)

   Specify a minimum backup priority number, and a maximum backup priority number that decide the AG replica for backup. For example, you can have a minimum priority of 10 and a maximum priority of 50. In this case, all the AG replicas with a priority more than 10 and less than 50 are considered for backup.

   By default, the minimum priority is 1 and maximum priority is 100.

> In cluster configurations, the backups are retained at each node of the cluster according to the retention settings set in the policy. If the owner node of the AG changes, the backups are taken according to the retention settings and the backups of the previous owner node will be retained. The retention for AG is applicable only at the node level.

Step 4: Configure Snapshot and Replication settings

1. In the Snapshot and Replication page, perform the following steps:
   a. Specify the schedule type by selecting **On demand**, **Hourly**, **Daily**, **Weekly**, or **Monthly**.

      You can only select one schedule type for a policy.

> You can specify the schedule (start date, end date, and frequency) for backup operation while creating a resource group. This enables you to create resource groups that share the same policy and backup frequency, but lets you assign different backup schedules to each policy.

> If you have scheduled for 2:00 a.m., the schedule will not be triggered during daylight saving time (DST).

Step 5: Configure up-to-the minute retention settings

1. In the Up-to-the minute retention settings section, depending on the backup type selected in the backup type page, perform one or more of the following actions:

Specific number of copies

Retain only a specific number of Snapshots.

1. Select the **Keep log backups applicable to last <number> days** option, and specify the number of days to be retained. If you near this limit, you might want to delete older copies.

Specific number of days

Retain the backup copies for a specific number of days.

1. Select the **Keep log backups applicable to last <number> days of full backups** option, and specify the number of days to keep the log backup copies.

Step 6: Configure Snapshot settings

1. For the Full backup retention settings, perform the following actions:
   a. Specify total number of snapshots to keep
      i. To specify the number of snapshots to keep, select **Copies to keep**.
      ii. If the number of snapshots exceeds the specified number, the snapshots are deleted with the oldest copies deleted first.

      By default, the value of retention count is set to 2. If you set the retention count to 1, the retention operation might fail because the first snapshot is the reference snapshot for the SnapVault relationship until a newer snapshot is replicated to the target.

      The maximum retention value is 1018. Backups will fail if retention is set to a value higher than what the underlying NetApp ONTAP version supports.

   b. Length of time to keep snapshots
      i. If you want to specify the number of days for which you want to keep the snapshots before deleting them, select **Retain copies for**.
   c. Select **Snapshot copy locking period** and specify the duration in days, months, or years.

      Snaplock retention period should be less than 100 years.

   d. Select a Policy label.

      You can assign SnapMirror labels to primary snapshots for remote replication, allowing the primary snapshots to offload the snapshot replication operation from SnapCenter to ONTAP secondary systems. This can be done without enabling SnapMirror or SnapVault option in the policy page.

Step 7: Configure secondary replication options

1. In the Select secondary replication options section, select one or both of the following secondary replication options:

Update SnapMirror

Update SnapMirror after creating a local Snapshot copy.

1. Select this option to create mirror copies of backup sets on another volume (SnapMirror).

   This option should be enabled for SnapMirror active sync.

   During secondary replication, the SnapLock expiry time loads the primary SnapLock expiry time. Clicking the **Refresh** button in the Topology page refreshes the secondary and primary SnapLock expiry time that are retrieved from ONTAP.

   See View SQL Server backups and clones in the Topology page.

Update SnapVault

Update SnapVault after creating a Snapshot copy.

1. Select this option to perform disk-to-disk backup replication.

   During secondary replication, the SnapLock expiry time loads the primary SnapLock expiry time. Clicking the **Refresh** button in the Topology page refreshes the secondary and primary SnapLock expiry time that are retrieved from ONTAP.

   When SnapLock is configured only on the secondary from ONTAP known as SnapLock Vault, clicking the **Refresh** button in the Topology page refreshes the locking period on the secondary that is retrieved from ONTAP.

   For more information on SnapLock Vault see Commit Snapshot copies to WORM on a vault destination

   See View SQL Server backups and clones in the Topology page.

Error Retry Count

1. Enter the number of replication attempts that should occur before the process halts.

Step 8: Configure script settings

1. In the Script page, enter the path and the arguments of the prescript or postscript that should be run before or after the backup operation, respectively.

   For example, you can run a script to update SNMP traps, automate alerts, and send logs.

   The prescripts or postscripts path should not include drives or shares. The path should be relative to the SCRIPTS_PATH.

   You must configure the SnapMirror retention policy in ONTAP so that the secondary storage does not reach the maximum limit of Snapshots.

Step 9: Configure verification settings

In the Verification page, perform the following steps:

1. In the Run verification for following backup schedules section, select the schedule frequency.
2. In the Database consistency check options section, perform the following actions:
   a. Limit the integrity structure to physical structure of the database (PHYSICAL_ONLY)
      i. Select **Limit the integrity structure to physical structure of the database (PHYSICAL_ONLY)** to limit the integrity check to the physical structure of the database and to detect torn pages, checksum failures, and common hardware failures that impact the database.
   b. Suppress all information messages (NO INFOMSGS)
      i. Select **Supress all information messages (NO_INFOMSGS)** to suppress all informational messages. Selected by default.
   c. Display all reported error messages per object (ALL_ERRORMSGS)
      i. Select **Display all reported error messages per object (ALL_ERRORMSGS)** to display all the reported errors per object.
   d. Do not check nonclustered indexes (NOINDEX)
      i. Select **Do not check nonclustered indexes (NOINDEX)** if you do not want to check nonclustered indexes. The SQL Server database uses Microsoft SQL Server Database Consistency Checker (DBCC) to check the logical and physical integrity of the objects in the database.
   e. Limit the checks and obtain the locks instead of using an internal database Snapshot (TABLOCK)
      i. Select **Limit the checks and obtain the locks instead of using an internal database Snapshot copy (TABLOCK)** to limit the checks and obtain locks instead of using an internal database Snapshot.
3. In the **Log Backup** section, select **Verify log backup upon completion** to verify the log backup upon completion.
4. In the **Verification script settings** section, enter the path and the arguments of the prescript or postscript that should be run before or after the verification operation, respectively.

> The prescripts or postscripts path should not include drives or shares. The path should be relative to the SCRIPTS_PATH.

Step 10: Review summary

1. Review the summary, and then select **Finish**.

Create resource groups and attach policies for SQL Server

A resource group is a container to which you add resources that you want to back up and protect together. A resource group enables you to back up all of the data that is associated with a given application simultaneously. A resource group is required for any data protection job. You must also attach one or more policies to the resource group to define the type of data protection job that you want to perform.

You can protect resources individually without creating a new resource group. You can take backups on the protected resource.

About this task

- For ONTAP 9.12.1 and below version, the clones created from the SnapLock Vault snapshots as part of restore will inherit the SnapLock Vault expiry time. Storage admin should manually cleanup the clones post

the SnapLock expiry time.

- Adding new databases without SnapMirror active sync to an existing resource group which contains resources with SnapMirror active sync, is not supported.
- Adding new databases to an existing resource group in failover mode of SnapMirror active sync is not supported. You can add resources to the resource group only in regular or fail-back state.

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page, select **Database** from the **View** list.

   If you have recently added a resource to SnapCenter, click **Refresh Resources** to view the newly added resource.

3. Click **New Resource Group**.
4. In the Name page, perform the following actions:

For this field…	Do this…
Name	Enter the resource group name. The resource group name should not exceed 250 characters.
Tags	Enter one or more labels that will help you later search for the resource group. For example, if you add HR as a tag to multiple resource groups, you can later find all resource groups associated with the HR tag.
Use custom name format for Snapshot copy	Optional: Enter a custom Snapshot name and format. For example, customtext_resourcegroup_policy_hostname or resourcegroup_hostname. By default, a timestamp is appended to the Snapshot name.

5. In the Resources page, perform the following steps:
   a. Select the host name, resource type, and the SQL Server instance from drop-down lists to filter the list of resources.

      If you have recently added resources, they will appear on the list of Available Resources only after you refresh your resource list.

   b. To move resources from the **Available Resources** section to the Selected Resources section, perform one of the following steps:
      - Select **Autoselect all resources on same storage volume** to move all of the resources on the same volume to the Selected Resources section.
      - Select the resources from the **Available Resources** section and then click the right arrow to move them to the **Selected Resources** section.

6. In the Policies page, perform the following steps:

   a. Select one or more policies from the drop-down list.

      You can also create a policy by clicking + .

      In the Configure schedules for selected policies section, the selected policies are listed.

   b. In the Configure schedules for selected policies section, click + in the Configure Schedules column for the policy for which you want to configure the schedule.

   c. In the Add schedules for policy *policy_name* dialog box, configure the schedule by specifying the start date, expiration date, and frequency, and then click **OK**.

      You must do this for each frequency listed in the policy. The configured schedules are listed in the Applied Schedules column in the **Configure schedules for selected policies** section.

   d. Select the Microsoft SQL Server scheduler.

      You must also select a scheduler instance to associate with the scheduling policy.

      If you do not select Microsoft SQL Server scheduler, the default is Microsoft Windows scheduler.

   Third party backup schedules are not supported when they overlap with SnapCenter backup schedules. You should not modify the schedules and rename the backup job created in Windows scheduler or SQL Server agent.

7. In the Verification page, perform the following steps:

   a. Select the verification server from the **Verification server** drop-down list.

      The list includes all the SQL Servers added in SnapCenter. You can select multiple verification servers (local host or remote host).

      The verification server version should match the version and edition of the SQL server that is hosting the primary database.

   b. Click **Load locators** to load the SnapMirror and SnapVault volumes to perform verification on secondary storage.

   c. Select the policy for which you want to configure your verification schedule, and then click + .

   d. In the Add Verification Schedules policy_name dialog box, perform the following actions:

If you want to…	Do this…
Run verification after backup	Select **Run verification after backup**.
Schedule a verification	Select **Run scheduled verification**.

   e. Click **OK**.

The configured schedules are listed in the Applied Schedules column. You can review and then edit by clicking or delete by clicking ✖ .

8. In the Notification page, from the **Email preference** drop-down list, select the scenarios in which you want to send the emails.

   You must also specify the sender and receiver email addresses, and the subject of the email. If you want to attach the report of the operation performed on the resource group, select **Attach Job Report**.

   ⓘ For email notification, you must have specified the SMTP server details either using the GUI or PowerShell command Set-SmSmtpServer.

9. Review the summary, and then click **Finish**.

Related information

Create backup policies for SQL Server databases

Create resource groups and enable secondary protection for Microsoft SQL Server resources on ASA r2 systems

You should create the resource group to add the resources that are on ASA r2 systems. You can also provision the secondary protection while creating the resource group.

Before you begin

- You should ensure that you are not adding both ONTAP 9.x resources and ASA r2 resources to the same resource group.
- You should ensure that you do not have a database with both ONTAP 9.x resources and ASA r2 resources.

About this task

- The secondary protection is available only if the logged-in user is assigned to the role that has the **SecondaryProtection** capability enabled.
- If you enabled secondary protection, the resource group is put into maintenance mode while creating the primary and secondary consistency groups. After the primary and secondary consistency groups are created, the resource group is put out of maintenance mode.
- SnapCenter does not support secondary protection for a clone resource.

Steps

1. In the left navigation pane, select **Resources**, and the appropriate plug-in from the list.
2. In the Resources page, click **New Resource Group**.
3. In the Name page, perform the following actions:
   a. Enter a name for the resource group in the Name field.

      ⓘ The resource group name should not exceed 250 characters.

   b. Enter one or more labels in the Tag field to help you search for the resource group later.

      For example, if you add HR as a tag to multiple resource groups, you can later find all resource groups associated with the HR tag.

c. Select this check box, and enter a custom name format that you want to use for the Snapshot name.

   For example, customtext_resource group_policy_hostname or resource group_hostname. By default, a timestamp is appended to the Snapshot name.

d. Specify the destinations of the archive log files that you do not want to back up.

   You should use the exact same destination as it was set in the application including prefix, if needed.

4. In the Resources page, select the database host name from the **Host** drop-down list.

   The resources are listed in the Available Resources section only if the resource is discovered successfully. If you have recently added resources, they will appear on the list of available resources only after you refresh your resource list.

5. Select the ASA r2 resources from the Available Resources section and move them to the Selected Resources section.
6. In the Application Settings page, select the backup option.
7. In the Policies page, perform the following steps:

   a. Select one or more policies from the drop-down list.

      You can also create a policy by clicking + .

      In the Configure schedules for selected policies section, the selected policies are listed.

   b. Click + in the Configure Schedules column for the policy for which you want to configure a schedule.

   c. In the Add schedules for policy *policy_name* window, configure the schedule, and then click **OK**.

      Where, *policy_name* is the name of the policy that you have selected.

      The configured schedules are listed in the Applied Schedules column.

   Third party backup schedules are not supported when they overlap with SnapCenter backup schedules.

8. If the secondary protection is enabled for the policy that you have selected, then Secondary Protection page is displayed and you need to perform the following steps:

   a. Select the type of the replication policy.

      Synchronous replication policy is not supported.

   b. Specify the consistency group suffix that you want to use.

   c. From the Destination Cluster and Destination SVM drop-downs select the peered cluster and SVM that you want to use.

The cluster and SVM peering is not supported by SnapCenter. You should use System Manager or ONTAP CLIs to perform cluster and SVM peering.

If the resources are already protected outside of SnapCenter, those resources will be displayed in the Secondary Protected Resources section.

1. On the Verification page, perform the following steps:
   a. Click **Load locators** to load the SnapMirror or SnapVault volumes to perform verification on secondary storage.
   b. Click + in the Configure Schedules column to configure the verification schedule for all the schedule types of the policy.
   c. In the Add Verification Schedules policy_name dialog box, perform the following actions:

If you want to…	Do this…
Run verification after backup	Select **Run verification after backup**.
Schedule a verification	Select **Run scheduled verification** and then select the schedule type from the drop-down list.

   d. Select **Verify on secondary location** to verify your backups on secondary storage system.
   e. Click **OK**.

   The configured verification schedules are listed in the Applied Schedules column.

2. In the Notification page, from the **Email preference** drop-down list, select the scenarios in which you want to send the emails.

   You must also specify the sender and receiver email addresses, and the subject of the email. If you want to attach the report of the operation performed on the resource group, select **Attach Job Report**.

   For email notification, you must have specified the SMTP server details using the either the GUI or the PowerShell command Set-SmSmtpServer.

3. Review the summary, and then click **Finish**.

Requirements for backing up SQL resources

Before you backup a SQL resource, you must ensure that several requirements are met.

- You must have migrated a resource from a non-NetApp storage system to a NetApp storage system.
- You must have created a backup policy.
- If you want to back up a resource that has a SnapMirror relationship to a secondary storage, the ONTAP role assigned to the storage user should include the “snapmirror all” privilege. However, if you are using the “vsadmin” role, then the “snapmirror all” privilege is not required.
- The backup operation initiated by an active directory (AD) user fails if the SQL instance credential is not assigned to the AD user or group. You must assign the SQL instance credential to AD user or group from

the **Settings** > **User Access** page.

- You must have created a resource group with a policy attached.
- If a resource group has multiple databases from different hosts, the backup operation on some hosts might be triggered late because of network issues. You should configure the value of FMaxRetryForUninitializedHosts in web.config by using the Set-SmConfigSettings PS cmdlet.

Create a storage system connection and a credential using PowerShell cmdlets

You must create a storage virtual machine (SVM) connection and a credential before using PowerShell cmdlets to perform data protection operations.

Before you begin

- You should have prepared the PowerShell environment to execute the PowerShell cmdlets.
- You should have the required permissions in the Infrastructure Admin role to create storage connections.
- You should ensure that the plug-in installations are not in progress.

  Host plug-in installations must not be in progress while adding a storage system connection because the host cache might not be updated and databases status might be displayed in the SnapCenter GUI as “Not available for backup” or “Not on NetApp storage”.

- Storage system names should be unique.

  SnapCenter does not support multiple storage systems with the same name on different clusters. Each storage system that is supported by SnapCenter should have a unique name and a unique management LIF IP address.

Steps

1. Initiate a PowerShell Core connection session by using the Open-SmConnection cmdlet.

   This example opens a PowerShell session:

   ```
   PS C:\> Open-SmConnection
   ```

2. Create a new connection to the storage system by using the Add-SmStorageConnection cmdlet.

   This example creates a new storage system connection:

   ```
   PS C:\> Add-SmStorageConnection -Storage test_vs1 -Protocol Https
   -Timeout 60
   ```

3. Create a new credential by using the Add-SmCredential cmdlet.

   This example creates a new credential named FinanceAdmin with Windows credentials:

   ```
   PS C:> Add-SmCredential -Name FinanceAdmin -AuthMode Windows
   -Credential sddev\administrator
   ```

The information regarding the parameters that can be used with the cmdlet and their descriptions can be obtained by running *Get-Help command_name*. Alternatively, you can also refer to the SnapCenter Software Cmdlet Reference Guide.

Back up SQL resources

If a resource is not yet part of any resource group, you can back up the resource from the Resources page.

About this task

- To optimize the backup operation, you should create a reverse lookup record of Windows cluster names and IP addresses in the DNS server.
- For Windows credentials authentication, you must set up your credential before installing the plug-ins.
- For SQL Server instance authentication, you must add the credential after installing the plug-ins.
- For gMSA authentication, you must setup gMSA while registering the host with SnapCenter in the **Add Host** or **Modify Host** page to enable and use the gMSA.
- If the host is added with gMSA and if the gMSA has login and system admin privileges, then gMSA is allowed to connect to the SQL instance.
  - SnapCenter will verify whether authentication for SQL instances is configured. If the authentication is configured, the SQL instance will be accessed using this credential.
  - If the authentication is not configured, use gMSA to see if the SQL plug-in is currently operating. If the plug-in is operating, it will be used to establish a connection to the SQL instance.
  - The SQL instance will be accessed via Windows credentials authentication when both authentication for SQL instances is not configured and the plug-in is not operational.

SnapCenter UI

Steps

1. In the left navigation pane, select **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page, select **Database** ,or **Instance**, or **Availability Group** from the **View** drop-down list.
   a. Select the database, or instance, or availability group that you want to back up.

      When you take a backup of an instance, the information about the last backup status or the timestamp of that instance will not be available in the resources page.

      In the topology view, you cannot differentiate whether the backup status, timestamp, or backup is for an instance or a database.
3. In the Resources page, select the **custom name format for Snapshot copy** check box, and then enter a custom name format that you want to use for the Snapshot name.

   For example, customtext_policy_hostname or resource_hostname. By default, a timestamp is appended to the Snapshot name.
4. In the Policies page, perform the following tasks:
   a. In the Policies section, select one or more policies from the drop-down list.

      You can create a policy by selecting to start the policy wizard.

      In the **Configure schedules for selected policies** section, the selected policies are listed.

   b. Select in the Configure Schedules column for the policy for which you want to configure a schedule.
   c. In the **Add schedules for policy** `policy_name` dialog box, configure the schedule, and then select **OK**.

      Here `policy_name` is the name of the policy that you have selected.

      The configured schedules are listed in the **Applied Schedules** column.

   d. Select the **Use Microsoft SQL Server scheduler**, and then select the scheduler instance from the **Scheduler Instance** drop-down list that is associated with the scheduling policy.
5. In the Verification page, perform the following steps:
   a. Select the verification server from the **Verification server** drop-down list.

      You can select multiple verification servers (local host or remote host).

      The verification server version should be equal or above the version of the edition of the SQL server that is hosting the primary database.

   b. Select **Load secondary locators to verify backups on secondary** to verify your backups on secondary storage system.
   c.

Select the policy for which you want to configure your verification schedule, and then select + .

d. In the Add Verification Schedules *policy_name* dialog box, perform the following actions:

If you want to…	Do this…
Run verification after backup	Select **Run Verification after Backup**.
Schedule a verification	Select **Run scheduled verification**.

> If the verification server does not have a storage connection, the verification operation fails with error: Failed to mount disk.

e. Select **OK**.

The configured schedules are listed in the Applied Schedules column.

6. In the Notification page, from the **Email preference** drop-down list, select the scenarios in which you want to send the emails.

You must also specify the sender and receiver email addresses, and the subject of the email. If you want to attach the report of the operation performed on the resource group, select **Attach Job Report**.

> For email notification, you must have specified the SMTP server details either using the GUI or PowerShell command Set-SmSmtpServer.

7. Review the summary, and then select **Finish**.

The database topology page is displayed.

8. Select **Back up Now**.
9. In the Backup page, perform the following steps:

a. If you have applied multiple policies to the resource, from the **Policy** drop-down list, select the policy that you want to use for backup.

If the policy selected for the on-demand backup is associated with a backup schedule, the on-demand backups will be retained based on the retention settings specified for the schedule type.

b. Select **Verify after backup** to verify your backup.

c. Select **Backup**.

> You should not rename the backup job created in Windows scheduler or SQL Server agent.

If the policy selected for the on-demand backup is associated with a backup schedule, the on-demand backups will be retained based on the retention settings specified for the schedule type.

An implicit resource group is created. You can view this by selecting respective user or group from the User Access page. The implicit resource group type is “Resource”.

10. Monitor the operation progress by selecting **Monitor** > **Jobs**.

After you finish

- In MetroCluster configurations, SnapCenter might not be able to detect a protection relationship after a failover.

  Unable to detect SnapMirror or SnapVault relationship after MetroCluster failover

- If you are backing up application data on VMDKs and the Java heap size for the SnapCenter Plug-in for VMware vSphere is not large enough, the backup might fail. To increase the Java heap size, locate the script file /opt/netapp/init_scripts/scvservice. In that script, the `do_start method` command starts the SnapCenter VMware plug-in service. Update that command to the following: `Java -jar -Xmx8192M -Xms4096M`.

Related information

Create backup policies for SQL Server databases

Backup operations fails with MySQL connection error because of the delay in the TCP_TIMEOUT

Backup fails with Windows scheduler error

Quiesce or grouping resources operations fail

PowerShell cmdlets

Steps

1. Initiate a connection session with the SnapCenter Server for a specified user by using the Open-SmConnection cmdlet.

   ```
   Open-smconnection  -SMSbaseurl  https://snapctr.demo.netapp.com:8146
   ```

   The username and password prompt is displayed.

2. Create a backup policy by using the Add-SmPolicy cmdlet.

   This example creates a new backup policy with a SQL backup type of FullBackup:

   ```
   PS C:\> Add-SmPolicy -PolicyName TESTPolicy
   -PluginPolicyType SCSQL -PolicyType Backup
   -SqlBackupType FullBackup -Verbose
   ```

   This example creates a new backup policy with a Windows file system backup type of CrashConsistent:

   ```
   PS C:\> Add-SmPolicy -PolicyName FileSystemBackupPolicy
   -PluginPolicyType SCW -PolicyType Backup
   -ScwBackupType CrashConsistent -Verbose
   ```

3. Discover host resources by using the Get-SmResources cmdlet.

This example discovers the resources for the Microsoft SQL plug-in on the specified host:

```
C:\PS>PS C:\> Get-SmResources -HostName vise-f6.sddev.mycompany.com
-PluginCode SCSQL
```

This example discovers the resources for Windows file systems on the specified host:

```
C:\PS>PS C:\> Get-SmResources -HostName vise2-f6.sddev.mycompany.com
-PluginCode SCW
```

4. Add a new resource group to SnapCenter by using the Add-SmResourceGroup cmdlet.

   This example creates a new SQL database backup resource group with the specified policy and resources:

```
PS C:\> Add-SmResourceGroup -ResourceGroupName AccountingResource
-Resources @{"Host"="visef6.org.com";
"Type"="SQL Database";"Names"="vise-f6\PayrollDatabase"}
-Policies "BackupPolicy"
```

   This example creates a new Windows file system backup resource group with the specified policy and resources:

```
PS C:\> Add-SmResourceGroup -ResourceGroupName EngineeringResource
-PluginCode SCW -Resources @{"Host"="WIN-VOK20IKID5I";
"Type"="Windows Filesystem";"Names"="E:\"}
-Policies "EngineeringBackupPolicy"
```

5. Initiate a new backup job by using the New-SmBackup cmdlet.

```
PS C:> New-SmBackup -ResourceGroupName PayrollDataset -Policy
FinancePolicy
```

6. View the status of the backup job by using the Get-SmBackupReport cmdlet.

   This example displays a job summary report of all jobs that were run on the specified date:

```
PS C:\> Get-SmJobSummaryReport -Date '1/27/2016'
```

The information regarding the parameters that can be used with the cmdlet and their descriptions can be obtained by running *Get-Help command_name*. Alternatively, you can also refer to the SnapCenter Software Cmdlet Reference Guide.

Back up SQL Server resource groups

You can back up a resource group on demand from the Resources page. If a resource group has a policy attached and a schedule configured, then backups occur automatically according to the schedule.

Steps

1. In the left navigation pane, select **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page, select **Resource Group** from the **View** list.

   You can search the resource group either by entering the resource group name in the search box, or by selecting , and then selecting the tag. You can then select to close the filter pane.

3. In the Resource Groups page, select the resource group that you want to back up, and then select **Back up Now**.
4. In the Backup page, perform the following steps:
   a. If you have associated multiple policies with the resource group, from the **Policy** drop-down list, select the policy that you want to use for backup.

      If the policy selected for the on-demand backup is associated with a backup schedule, the on-demand backups will be retained based on the retention settings specified for the schedule type.

   b. After backup, select **Verify** to verify the on-demand backup.

      The **Verify** option in the policy applies only to scheduled jobs.

   c. Select **Backup**.
5. Monitor the operation progress by selecting **Monitor** > **Jobs**.

Related information

Create backup policies for SQL Server databases

Create resource groups and attach policies for SQL Server

Backup operations fails with MySQL connection error because of the delay in the TCP_TIMEOUT

Backup fails with Windows scheduler error

Monitor SQL resources backup operations in the SnapCenter Jobs page

You can monitor the progress of different backup operations by using the SnapCenterJobs page. You might want to check the progress to determine when it is complete or if there is an issue.

About this task

The following icons appear on the Jobs page and indicate the corresponding state of the operations:

- In progress
- Completed successfully

- Failed
- Completed with warnings or could not start due to warnings
- Queued
- Canceled

Steps

1. In the left navigation pane, click **Monitor**.
2. In the Monitor page, click **Jobs**.
3. In the Jobs page, perform the following steps:
   a. Click to filter the list so that only backup operations are listed.
   b. Specify the start and end dates.
   c. From the **Type** drop-down list, select **Backup**.
   d. From the **Status** drop-down, select the backup status.
   e. Click **Apply** to view the operations completed successfully.
4. Select a backup job, and then click **Details** to view the job details.

   Though the backup job status displays , when you click on job details you might see that some of the child tasks of the backup operation are still in progress or marked with warning signs.

5. In the Job Details page, click **View logs**.

   The **View logs** button displays the detailed logs for the selected operation.

Monitor data protection operations on SQL resources in the Activity pane

The Activity pane displays the five most recent operations performed. The Activity pane also displays when the operation was initiated and the status of the operation.

The Activity pane displays information regarding backup, restore, clone, and scheduled backup operations.

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. Click on the Activity pane to view the five most recent operations.

   When you click one of the operations, the operation details are listed in the **Job Details** page.

Cancel the SnapCenter Plug-in for Microsoft SQL Server backup operations

You can cancel backup operations that are running, queued, or non-responsive. When you cancel a backup operation, the SnapCenter Server stops the operation and removes all the Snapshots from the storage if the backup created is not registered with SnapCenter Server. If the backup is already registered with SnapCenter Server, it will not roll back the already created Snapshot even after the cancellation is triggered.

Before you begin

- You must be logged in as the SnapCenter Admin or job owner to cancel restore operations.
- You can cancel only the log or full backup operations that are queued or running.
- You cannot cancel the operation after the verification has started.

  If you cancel the operation before verification, the operation is canceled, and the verification operation will not be performed.

- You can cancel a backup operation from either the Monitor page or the Activity pane.
- In addition to using the SnapCenter GUI, you can use PowerShell cmdlets to cancel operations.
- The **Cancel Job** button is disabled for operations that cannot be canceled.
- If you selected **All members of this role can see and operate on other members objects** in Users\Groups page while creating a role, you can cancel the queued backup operations of other members while using that role.

Steps

Perform one of the following actions:

From the…	Action
Monitor page	a. In the left navigation pane, select **Monitor** > **Jobs**. b. Select the job and select **Cancel Job** .
Activity pane	a. After initiating the backup job, select on the Activity pane to view the five most recent operations. b. Select the operation. c. In the Job Details page, select **Cancel Job**.

Result

The operation is canceled, and the resource is reverted to the previous state. If the operation you canceled is non-responsive in the canceling or running state, you should run the `Cancel-SmJob -JobID <int> -Force` cmdlet to forcefully stop the backup operation.

View SQL Server backups and clones in the Topology page

When you are preparing to back up or clone a resource, you might find it helpful to view a graphical representation of all backups and clones on the primary and secondary storage.

About this task

In the Topology page, you can see all of the backups and clones that are available for the selected resource or resource group. You can view the details of those backups and clones, and then select them to perform data protection operations.

You can review the following icons in the **Manage Copies** view to determine whether the backups and clones are available on the primary or secondary storage (Mirror copies or Vault copies).

- displays the number of backups and clones that are available on the primary storage.
- displays the number of backups and clones that are mirrored on the secondary storage using SnapMirror technology.
- displays the number of backups and clones that are replicated on the secondary storage using SnapVault technology.
  - The number of backups displayed includes the backups deleted from the secondary storage.

    For example, if you created 6 backups using a policy to retain only 4 backups, the number of backups displayed is 6.

Clones of a backup of a version-flexible mirror on a mirror-vault type volume are displayed in the topology view, but the mirror backup count in the topology view does not include the version-flexible backup.

If you have secondary relationship as SnapMirror active sync (initially released as SnapMirror Business Continuity [SM-BC]), you can see following additional icons:

- The replica site is up.
- The replica site is down.
- The secondary mirror or vault relationship has not been re-established.

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page, either select the resource or resource group from the **View** drop-down list.
3. Select the resource either from the resource details view or from the resource group details view.

   If the resource selected is a cloned database, protect the cloned database, source of the clone is displayed in the Topology page. Click **Details** to view the backup used to clone.

   If the resource is protected, the Topology page of the selected resource is displayed.

4. Review the Summary card to see a summary of the number of backups and clones available on the primary and secondary storage.

   The **Summary Card** section displays the total number of backups and clones.

   Clicking the **Refresh** button starts a query of the storage to display an accurate count.

   If SnapLock enabled backup is taken, then clicking the **Refresh** button refreshes the primary and secondary SnapLock expiry time retrieved from ONTAP. A weekly schedule also refreshes the primary and secondary SnapLock expiry time retrieved from ONTAP.

When the application resource is spread across multiple volumes, the SnapLock expiry time for the backup will be the longest SnapLock expiry time that is set for a Snapshot in a volume. The longest SnapLock expiry time is retrieved from ONTAP.

For SnapMirror active sync, clicking the **Refresh** button refreshes the SnapCenter backup inventory by querying ONTAP for both primary and replica sites. A weekly schedule also performs this activity for all databases containing SnapMirror active sync relationship.

- For SnapMirror active sync and only for ONTAP 9.14.1, Async Mirror or Async MirrorVault relationships to the new primary destination should be manually configured after failover. From ONTAP 9.15.1 onwards Async Mirror or Async MirrorVault is auto configured to the new primary destination.
- After failover, a backup should be created for SnapCenter to be aware of the failover. You can click **Refresh** only after a backup has been created.

5. In the **Manage Copies** view, click **Backups** or **Clones** from the primary or secondary storage to see details of a backup or clone.

   The details of the backups and clones are displayed in a table format.

6. Select the backup from the table, and then click the data protection icons to perform restore, clone, rename, and delete operations.

   You cannot rename or delete backups that are on the secondary storage.

7. Select a clone from the table and click **Clone Split**.
8. If you want to delete a clone, select the clone from the table, and then click .

Clean up the secondary backup count using PowerShell cmdlets

You can use the Remove-SmBackup cmdlet to clean up the backup count for secondary backups that have no Snapshot. You might want to use this cmdlet when the total Snapshots displayed in the Manage Copies topology do not match the secondary storage Snapshot retention setting.

You must have prepared the PowerShell environment to execute the PowerShell cmdlets.

The information regarding the parameters that can be used with the cmdlet and their descriptions can be obtained by running *Get-Help command_name*. Alternatively, you can also refer to the SnapCenter Software Cmdlet Reference Guide.

Steps

1. Initiate a connection session with the SnapCenter Server for a specified user by using the Open-SmConnection cmdlet.

```
Open-SmConnection  -SMSbaseurl  https:\\snapctr.demo.netapp.com:8146/
```

2. Clean up secondary backups count using the -CleanupSecondaryBackups parameter.

   This example cleans up the backup count for secondary backups with no Snapshots:

```
Remove-SmBackup -CleanupSecondaryBackups
Remove-SmBackup
Are you sure want to remove the backup(s).
[Y] Yes  [A] Yes to All  [N] No  [L] No to All  [S] Suspend  [?] Help
(default is "Y"):
```

Restore SQL Server resources

Restore workflow

You can use SnapCenter to restore SQL Server databases by restoring the data from one or more backups to your active file system and then recovering the database. You can also restore databases that are in Availability Groups and then add the restored databases to the Availability Group. Before restoring an SQL Server database, you must perform several preparatory tasks.

The following workflow shows the sequence in which you must perform the database restoration operations:

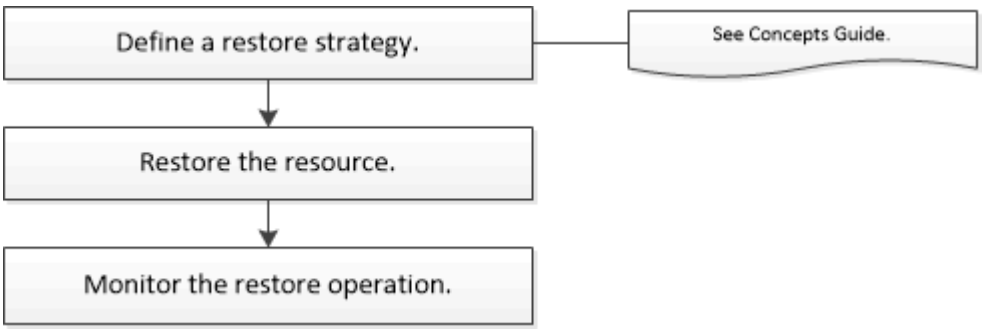

You can also use PowerShell cmdlets manually or in scripts to perform backup, restore, recovery, verify, and clone operations. For detailed information about PowerShell cmdlets, use the SnapCenter cmdlet help or see the SnapCenter Software Cmdlet Reference Guide

Find more information

Restore an SQL Server database from secondary storage

Restore and recover resources using PowerShell cmdlets

Restore operation might fail on Windows 2008 R2

Requirements for restoring a database

Before you restore a SQL Server database from a SnapCenter Plug-in for Microsoft SQL Server backup, you must ensure that several requirements are met.

- The target SQL Server instance must be online and running before you can restore a database.

  This applies to both user database restore operations and system database restore operations.

- SnapCenter operations that are scheduled to run against the SQL Server data you are restoring must be disabled, including any jobs scheduled on remote management or remote verification servers.
- If system databases are not functional, you must first rebuild the system databases using a SQL Server utility.
- If you are installing the plug-in, ensure that you grant permissions for other roles to restore the Availability Group (AG) backups.

  Restoring AG fails when one of the following conditions are met:

  - If the plug-in is installed by RBAC user and an admin tries to restore an AG backup
  - If the plug-in is installed by an admin and a RBAC user tries to restore an AG backup
- If you are restoring custom log directory backups to an alternate host, the SnapCenter Server and the plug-in host must have the same SnapCenter version installed.
- You must have installed Microsoft hotfix, KB2887595. The Microsoft Support Site contains more information about KB2887595.

  Microsoft Support Article 2887595: Windows RT 8.1, Windows 8.1, and Windows Server 2012 R2 update rollup: November 2013
- You must have backed up the resource groups or database.
- If you are replicating Snapshots to a mirror or vault, the SnapCenter administrator must have assigned you the storage virtual machines (SVMs) for both the source volumes and destination volumes.

  For information about how administrators assign resources to users, see the SnapCenter installation information.
- All backup and clone jobs must be stopped before restoring the database.
- The restore operation might timeout if the database size is in terabytes (TB).

  You must increase the value of the RESTTimeout parameter of SnapCenter Server to 20000000 ms by running the following command: Set-SmConfigSettings -Agent -configSettings @{"RESTTimeout" = "20000000"}. According to the size of the database, the timeout value can be changed and the maximum value that you can set is 86400000 ms.

If you want to restore while the databases are online, the online restore option should be enabled in the Restore page.

Restore SQL Server database backups

You can use SnapCenter to restore backed-up SQL Server databases. Database restoration is a multiphase process that copies all of the data and log pages from a specified SQL Server backup to a specified database.

About this task

- You can restore the backed-up SQL Server databases to a different SQL Server instance on the same host where the backup was created.

  You can use SnapCenter to restore the backed-up SQL Server databases to an alternate path so that you do not replace a production version.

- SnapCenter can restore databases in a Windows cluster without taking the SQL Server cluster group offline.
- If a cluster failure (a cluster group move operation) occurs during a restore operation (for example, if the node that owns the resources goes down), you must reconnect to the SQL Server instance, and then restart the restore operation.
- You cannot restore the database when the users or the SQL Server Agent jobs are accessing the database.
- You cannot restore system databases to an alternate path.
- The SCRIPTS_PATH is defined using the PredefinedWindowsScriptsDirectory key located in the SMCoreServiceHost.exe.Config file of the plug-in host.

  If needed, you can change this path and restart SMcore service. It is recommended that you use the default path for security.

  The value of the key can be displayed from swagger through the API: API /4.7/configsettings

  You can use the GET API to display the value of the key. SET API is not supported.
- Most of the fields on the Restore wizard pages are self-explanatory. The following information describes fields for which you might need guidance.
- For SnapMirror active sync restore operation, you must select the backup from the primary location.
- For SnapLock enabled policies, for ONTAP 9.12.1 and below version, if you specify a Snapshot locking period, the clones created from the tamper proof Snapshots as part of restore will inherit the SnapLock expiry time. Storage admin should manually cleanup the clones post the SnapLock expiry time.

SnapCenter UI

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page, select either **Database** or **Resource Group** from the **View** list.
3. Select the database or the resource group from the list.

   The topology page is displayed.

4. From the Manage Copies view, select **Backups** from the storage system.
5. Select the backup from the table, and then click the icon.

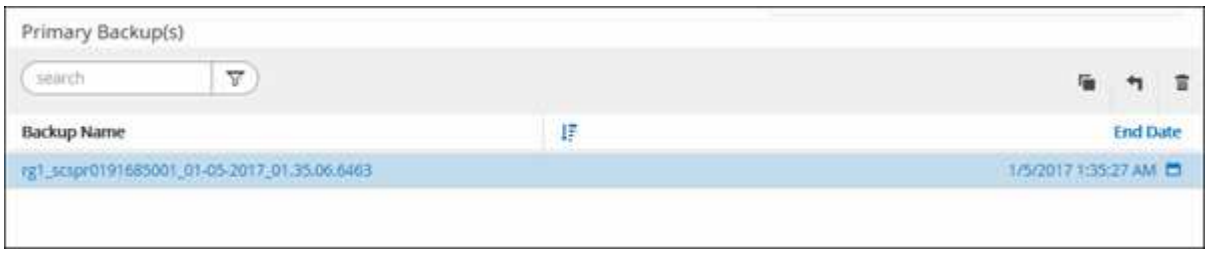

6. In the Restore Scope page, select one of the following options:

Option	Description
Restore the database to the same host where the backup was created	Select this option if you want to restore the database to the same SQL server where the backups are taken.
Restore the database to an alternate host	Select this option if you want the database to be restored to a different SQL server in the same or different host where backups are taken. Select a host name, provide a database name (optional), select an instance, and specify the restore paths. The file extension provided in the alternate path must be same as the file extension of the original database file. If the **Restore the database to an alternate host** option is not displayed in the Restore Scope page, clear the browser cache.

Option	Description
Restore the database using existing database files	Select this option if you want the database to be restored to an alternate SQL Server in the same or different host where backups are taken. Database files should be already present on the given existing file paths. Select a host name, provide a database name (optional), select an instance, and specify the restore paths.

7. In the Recovery Scope page, select one of the following options:

Option	Description
None	Select **None** when you need to restore only the full backup without any logs.
All log backups	Select **All log backups** up-to-the-minute backup restore operation to restore all of the available log backups after the full backup.
By log backups until	Select **By log backups** to perform a point-in-time restore operation, which restores the database based on backup logs until the backup log with the selected date.
By specific date until	Select **By specific date until** to specify the date and time after which transaction logs are not applied to the restored database. This point-in-time restore operation halts the restoration of transaction log entries that were recorded after the specified date and time.
Use custom log directory	If you have selected **All log backups**, **By log backups**, or **By specific date until** and the logs are located at a custom location, select **Use custom log directory**, and then specify the log location. The **Use Custom log directory** option is available only if you have selected **Restore the database to an alternate host** or **Restore the database using existing database files**. You can also use the shared path but ensure that the path is accessible by the SQL user. The custom log directory is not supported for availability group database.

8. In the Pre Ops page, perform the following steps:
   a. In the Pre Restore Options page, select one of the following options:
      - Select **Overwrite the database with same name during restore** to restore the database with the same name.
      - Select **Retain SQL database replication settings** to restore the database and retain the existing replication settings.
      - Select **Create transaction log backup before restore** to create a transaction log before the restore operation begins.
      - Select **Quit restore if transaction log backup before restore fails** to abort the restore operation if the transaction log backup fails.
   b. Specify optional scripts to run before performing a restore job.

      For example, you can run a script to update SNMP traps, automate alerts, send logs, and so on.

      The prescripts or postscripts path should not include drives or shares. The path should be relative to the SCRIPTS_PATH.

9. In the Post Ops page, perform the following steps:
   a. In the Choose database state after restore completes section, select one of the following options:
      - Select **Operational, but unavailable for restoring additional transaction logs** if you are restoring all of the necessary backups now.

        This is the default behavior, which leaves the database ready for use by rolling back the uncommitted transactions. You cannot restore additional transaction logs until you create a backup.

      - Select **Non-operational, but available for restoring additional transactional logs** to leave the database non-operational without rolling back the uncommitted transactions.

        Additional transaction logs can be restored. You cannot use the database until it is recovered.

      - Select **Read-only mode, available for restoring additional transactional logs** to leave the database in read-only mode.

        This option undoes uncommitted transactions, but saves the undone actions in a standby file so that recovery effects can be reverted.

        If the Undo directory option is enabled, more transaction logs are restored. If the restore operation for the transaction log is unsuccessful, the changes can be rolled back. The SQL Server documentation contains more information.

   b. Specify optional scripts to run after performing a restore job.

      For example, you can run a script to update SNMP traps, automate alerts, send logs, and so on.

      The prescripts or postscripts path should not include drives or shares. The path should be relative to the SCRIPTS_PATH.

10. In the Notification page, from the **Email preference** drop-down list, select the scenarios in which you want to send the emails.

You must also specify the sender and receiver email addresses, and the subject of the email.

11. Review the summary, and then click **Finish**.
12. Monitor the restore process by using the **Monitor** > **Jobs** page.

PowerShell cmdlets

Steps

1. Initiate a connection session with the SnapCenter Server for a specified user by using the Open-SmConnection cmdlet.

```
PS C:\> Open-Smconnection
```

2. Retrieve the information about the one or more backups that you want to restore by using the Get-SmBackup and Get-SmBackupReport cmdlets.

   This example displays information about all available backups:

```
PS C:\> Get-SmBackup

BackupId                       BackupName
BackupTime                     BackupType
--------                       ----------
----------                     ----------
  1               Payroll Dataset_vise-f6_08... 8/4/2015    11:02:32
AM                 Full Backup
  2               Payroll Dataset_vise-f6_08... 8/4/2015    11:23:17
AM
```

   This example displays detailed information about the backup from January 29th 2015 to February 3rd, 2015:

```
PS C:\> Get-SmBackupReport -FromDateTime "1/29/2015" -ToDateTime
"2/3/2015"

   SmBackupId           : 113
   SmJobId              : 2032
   StartDateTime        : 2/2/2015 6:57:03 AM
   EndDateTime          : 2/2/2015 6:57:11 AM
   Duration             : 00:00:07.3060000
   CreatedDateTime      : 2/2/2015 6:57:23 AM
   Status               : Completed
   ProtectionGroupName  : Clone
   SmProtectionGroupId  : 34
   PolicyName           : Vault
   SmPolicyId           : 18
   BackupName           : Clone_SCSPR0019366001_02-02-2015_06.57.08
   VerificationStatus   : NotVerified

   SmBackupId           : 114
   SmJobId              : 2183
   StartDateTime        : 2/2/2015 1:02:41 PM
   EndDateTime          : 2/2/2015 1:02:38 PM
   Duration             : -00:00:03.2300000
   CreatedDateTime      : 2/2/2015 1:02:53 PM
   Status               : Completed
   ProtectionGroupName  : Clone
   SmProtectionGroupId  : 34
   PolicyName           : Vault
   SmPolicyId           : 18
   BackupName           : Clone_SCSPR0019366001_02-02-2015_13.02.45
   VerificationStatus   : NotVerified
```

3. Restore data from the backup by using the Restore-SmBackup cmdlet.

```
Restore-SmBackup -PluginCode 'DummyPlugin' -AppObjectId
'scc54.sccore.test.com\DummyPlugin\NTP\DB1' -BackupId 269
-Confirm:$false
output:
Name                : Restore
'scc54.sccore.test.com\DummyPlugin\NTP\DB1'
Id                  : 2368
StartTime           : 10/4/2016 11:22:02 PM
EndTime             :
IsCancellable       : False
IsRestartable       : False
IsCompleted         : False
IsVisible           : True
IsScheduled         : False
PercentageCompleted : 0
Description         :
Status              : Queued
Owner               :
Error               :
Priority            : None
Tasks               : {}
ParentJobID         : 0
EventId             : 0
JobTypeId           :
ApisJobKey          :
ObjectId            : 0
PluginCode          : NONE
PluginName          :
```

The information regarding the parameters that can be used with the cmdlet and their descriptions can be obtained by running *Get-Help command_name*. Alternatively, you can also refer to the SnapCenter Software Cmdlet Reference Guide.

Restore an SQL Server database from secondary storage

You can restore the backed-up SQL Server databases from the physical LUNs (RDM, iSCSI, or FCP) on a secondary storage system. The Restore feature is a multiphase process that copies all of the data and the log pages from a specified SQL Server backup residing on the secondary storage system to a specified database.

Before you begin

- You must have replicated the Snapshots from primary to secondary storage system.
- You must ensure that the SnapCenter Server and the plug-in host are able to connect to the secondary storage system.

- Most of the fields on the Restore wizard pages are explained in the basic restore process. The following information describes some of the fields for which you might need guidance.

About this task

For SnapLock enabled policies, for ONTAP 9.12.1 and below version, if you specify a Snapshot locking period, the clones created from the tamper proof Snapshots as part of restore will inherit the SnapLock expiry time. Storage admin should manually cleanup the clones post the SnapLock expiry time.

Steps

1. In the left navigation pane, click **Resources**, and then select **SnapCenter Plug-in for SQL Server** from the list.
2. In the Resources page, select **Database** or **Resource Group** from the **View** drop-down list.
3. Select the database or resource group.

   The database or resource group topology page is displayed.
4. In the Manage Copies section, select **Backups** from the secondary storage system (mirrored or vault).
5. Select the backup from the list, and then click .
6. In the Location page, choose the destination volume for restoring selected resource.
7. Complete the Restore wizard, review the summary, and then click **Finish**.

If you restored a database to a different path that is shared by other databases, you should perform a full backup and backup verification to confirm that your restored database is free of physical-level corruption.

Restore resources using PowerShell cmdlets

Restoring a resource backup includes initiating a connection session with the SnapCenter Server, listing the backups and retrieving backup information, and restoring a backup.

You must have prepared the PowerShell environment to execute the PowerShell cmdlets.

Steps

1. Initiate a connection session with the SnapCenter Server for a specified user by using the Open-SmConnection cmdlet.

   ```
   PS C:\> Open-Smconnection
   ```

2. Retrieve the information about the one or more backups that you want to restore by using the Get-SmBackup and Get-SmBackupReport cmdlets.

   This example displays information about all available backups:

```
PS C:\> Get-SmBackup

BackupId                      BackupName                    BackupTime
BackupType
--------                      ----------                    ----------
----------
  1               Payroll Dataset_vise-f6_08... 8/4/2015    11:02:32 AM
Full Backup
  2               Payroll Dataset_vise-f6_08... 8/4/2015    11:23:17 AM
```

This example displays detailed information about the backup from January 29th 2015 to February 3rd, 2015:

```
PS C:\> Get-SmBackupReport -FromDateTime "1/29/2015" -ToDateTime
"2/3/2015"

   SmBackupId           : 113
   SmJobId              : 2032
   StartDateTime        : 2/2/2015 6:57:03 AM
   EndDateTime          : 2/2/2015 6:57:11 AM
   Duration             : 00:00:07.3060000
   CreatedDateTime      : 2/2/2015 6:57:23 AM
   Status               : Completed
   ProtectionGroupName  : Clone
   SmProtectionGroupId  : 34
   PolicyName           : Vault
   SmPolicyId           : 18
   BackupName           : Clone_SCSPR0019366001_02-02-2015_06.57.08
   VerificationStatus   : NotVerified

   SmBackupId           : 114
   SmJobId              : 2183
   StartDateTime        : 2/2/2015 1:02:41 PM
   EndDateTime          : 2/2/2015 1:02:38 PM
   Duration             : -00:00:03.2300000
   CreatedDateTime      : 2/2/2015 1:02:53 PM
   Status               : Completed
   ProtectionGroupName  : Clone
   SmProtectionGroupId  : 34
   PolicyName           : Vault
   SmPolicyId           : 18
   BackupName           : Clone_SCSPR0019366001_02-02-2015_13.02.45
   VerificationStatus   : NotVerified
```

3. Restore data from the backup by using the Restore-SmBackup cmdlet.

```
Restore-SmBackup -PluginCode 'DummyPlugin' -AppObjectId
'scc54.sccore.test.com\DummyPlugin\NTP\DB1' -BackupId 269
-Confirm:$false
output:
Name                : Restore
'scc54.sccore.test.com\DummyPlugin\NTP\DB1'
Id                  : 2368
StartTime           : 10/4/2016 11:22:02 PM
EndTime             :
IsCancellable       : False
IsRestartable       : False
IsCompleted         : False
IsVisible           : True
IsScheduled         : False
PercentageCompleted : 0
Description         :
Status              : Queued
Owner               :
Error               :
Priority            : None
Tasks               : {}
ParentJobID         : 0
EventId             : 0
JobTypeId           :
ApisJobKey          :
ObjectId            : 0
PluginCode          : NONE
PluginName          :
```

The information regarding the parameters that can be used with the cmdlet and their descriptions can be obtained by running *Get-Help command_name*. Alternatively, you can also refer to the SnapCenter Software Cmdlet Reference Guide.

Reseed Availability Group databases

Reseed is an option to restore Availability Group (AG) databases. If a secondary database gets out of synchronization with the primary database in an AG, you can reseed the secondary database.

Before you begin

- You must have created backup of secondary AG database that you want to restore.
- The SnapCenter Server and the plug-in host must have the same SnapCenter version installed.

About this task

- You cannot perform reseed operation on primary databases.
- You cannot perform a reseed operation if the replica database is removed from the availability group. When the replica is removed, the reseed operation fails.
- While running the reseed operation on SQL Availability Group database, you should not trigger log backups on the replica databases of that availability group database. If you trigger log backups during reseed operation, the reseed operation fails with The mirror database, "database_name" has insufficient transaction log data to preserve the log backup chain of the principal database error message.

Steps

1. In the left navigation pane, click **Resources**, and then select **SnapCenter Plug-in for SQL Server** from the list.
2. In the Resources page, select **Database** from the **View** list.
3. Select secondary AG database from the list.
4. Click **Reseed**.
5. Monitor the operation progress by clicking **Monitor** > **Jobs**.

Monitor SQL resources restore operations

You can monitor the progress of different SnapCenter restore operations by using the Jobs page. You might want to check the progress of an operation to determine when it is complete or if there is an issue.

About this task

Post-restore states describe the conditions of the resource after a restore operation and any further restore actions that you can take.

The following icons appear on the Jobs page, and indicate the state of the operation:

- In progress
- Completed successfully
- Failed
- Completed with warnings or could not start due to warnings
- Queued
- Canceled

Steps

1. In the left navigation pane, click **Monitor**.
2. In the **Monitor** page, click **Jobs**.
3. In the **Jobs** page, perform the following steps:
   a. Click to filter the list so that only restore operations are listed.
   b. Specify the start and end dates.
   c. From the **Type** drop-down list, select **Restore**.
   d. From the **Status** drop-down list, select the restore status.

e. Click **Apply** to view the operations that have been completed successfully.

4. Select the restore job, and then click **Details** to view the job details.
5. In the **Job Details** page, click **View logs**.

   The **View logs** button displays the detailed logs for the selected operation.

Cancel SQL resources restore operations

You can cancel restore jobs that are queued.

You should be logged in as the SnapCenter Admin or job owner to cancel restore operations.

About this task

- You can cancel a queued restore operation from either the **Monitor** page or the **Activity** pane.
- You cannot cancel a running restore operation.
- You can use the SnapCenter GUI, PowerShell cmdlets, or CLI commands to cancel the queued restore operations.
- The **Cancel Job** button is disabled for restore operations that cannot be canceled.
- If you selected **All members of this role can see and operate on other members objects** in Users\Groups page while creating a role, you can cancel the queued restore operations of other members while using that role.

Step

Perform one of the following actions:

From the…	Action
Monitor page	a. In the left navigation pane, click **Monitor** > **Jobs**. b. Select the job and click **Cancel Job**.
Activity pane	a. After initiating the restore operation, click on the Activity pane to view the five most recent operations. b. Select the operation. c. In the Job Details page, click **Cancel Job**.

Clone SQL Server database resources

Clone workflow

You must perform several tasks using SnapCenter Server before cloning database resources from a backup. Database cloning is the process of creating a point-in-time copy of a production database or its backup set. You can clone databases to test functionality that has to be implemented using the current database structure and content during application development cycles, to use the data extraction and manipulation tools

when populating data warehouses, or to recover data that was mistakenly deleted or changed.

A database cloning operation generates reports based on the job IDs.

The following workflow shows the sequence in which you must perform the cloning operations:

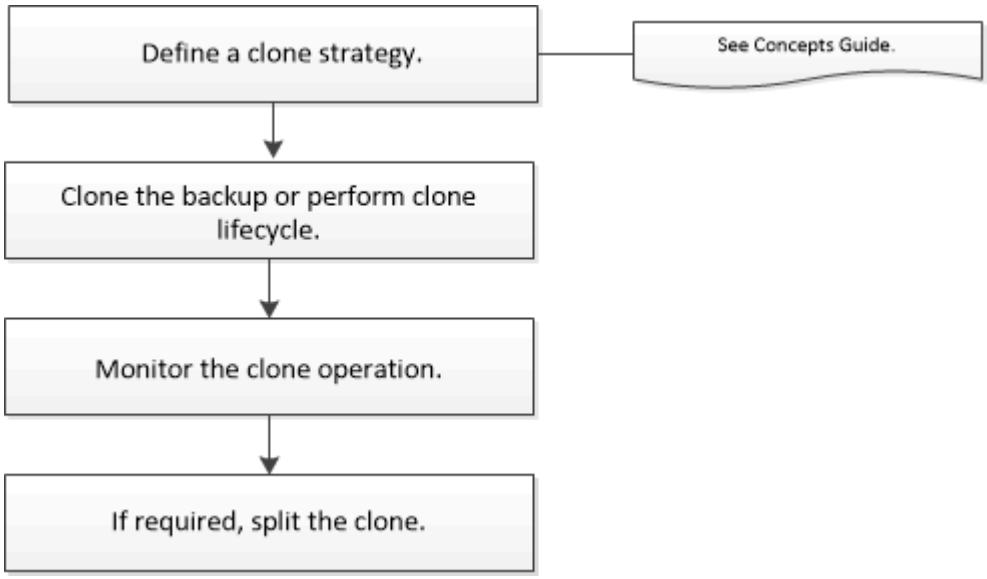

You can also use PowerShell cmdlets manually or in scripts to perform backup, restore, recovery, verify, and clone operations. For detailed information about PowerShell cmdlets, use the SnapCenter cmdlet help or see the SnapCenter Software Cmdlet Reference Guide

Find more information

Clone from a SQL Server database backup

Perform Clone Lifecycle

Clone operation might fail or take longer time to complete with default TCP_TIMEOUT value

Clone from a SQL Server database backup

You can use SnapCenter to clone a SQL Server database backup. If you want to access or restore an older version of the data, you can clone database backups on demand.

Before you begin

- You should have prepared for data protection by completing tasks such as adding hosts, identifying resources, and creating storage system connections.
- You should have backed up databases or resource groups.
- The protection type such as mirror, vault, or mirror-vault for data LUN and log LUN should be same to discover secondary locators during cloning to an alternate host using log backups.
- If the mounted clone drive cannot be found during a SnapCenter clone operation, you should change the CloneRetryTimeout parameter of SnapCenter Server to 300.
- You should ensure that the aggregates hosting the volumes should be in the assigned aggregates list of the storage virtual machine (SVM).

About this task

- While cloning to a standalone database instance, ensure that the mount point path exists and it is a dedicated disk.
- While cloning to a Failover Cluster Instance (FCI), ensure that the mount points exists, it is a shared disk, and the path and the FCI should belong to the same SQL resource group.
- Ensure that there is only one vFC or FC initiator attached to each host. This is because, SnapCenter supports only one initiator per host.
- If the source database or the target instance is on a cluster shared volume (csv), then the cloned database will be on the csv.
- The SCRIPTS_PATH is defined using the PredefinedWindowsScriptsDirectory key located in the SMCoreServiceHost.exe.Config file of the plug-in host.

  If needed, you can change this path and restart SMcore service. It is recommended that you use the default path for security.

  The value of the key can be displayed from swagger through the API: API /4.7/configsettings

  You can use the GET API to display the value of the key. SET API is not supported.

  For virtual environments (VMDK/RDM), ensure that the mount point is a dedicated disk.

- For ONTAP 9.12.1 and below version, the clones created from the SnapLock Vault Snapshots as part of restore will inherit the SnapLock Vault expiry time. Storage admin should manually cleanup the clones post the SnapLock expiry time.

SnapCenter UI

Steps

1. In the left navigation pane, select **Resources**, and then select **SnapCenter Plug-in for SQL Server** from the list.
2. In the Resources page, select either **Database** or **Resource Group** from the **View** list.

   Cloning of a backup of an instance is not supported.

3. Select the database or resource group.
4. From the **Manage Copies** view page, select the backup either from primary or secondary (mirrored or vaulted) storage system.
5. Select the backup, and then select .
6. In the **Clone Options** page, perform the following actions:

For this field…	Do this…
Clone server	Choose a host on which the clone should be created.
Clone instance	Choose a clone instance to which you want to clone the database backup. This SQL instance must be located in the specified clone server.
Clone suffix	Enter a suffix that will be appended to the clone file name to identify that the database is a clone. For example, *db1_clone*. If you are cloning to the same location as the original database, you must provide a suffix to differentiate the cloned database from the original database. Otherwise, the operation fails.
Auto assign mount point or Auto assign volume mount point under path	Choose whether to automatically assign a mount point or a volume mount point under a path. Auto assign volume mount point under path: The mount point under a path allows you to provide a specific directory. The mount points will be created within that directory. Before you choose this option, you must ensure that the directory is empty. If there is a database in the directory, the database will be in an invalid state after the mount operation.

7. In the Logs page, select one of the following options:

For this field…	Do this…
None	Choose this option when you want to clone only the full backup without any logs.
All log backups	Choose this option to clone all the available log backups dated after the full backup.
By log backups until	Choose this option to clone the database based on the backup logs that were created up to the backup log with the selected date.
By specific date until	Specify the date and time after which the transaction logs are not applied to the cloned database. This point-in-time clone halts the clone of the transaction log entries that were recorded after the specified date and time.

8. In the **Script** page, enter the script timeout, path, and the arguments of the prescript or postscript that should be run before or after the clone operation, respectively.

   For example, you can run a script to update SNMP traps, automate alerts, send logs, and so on.

   > The prescripts or postscripts path should not include drives or shares. The path should be relative to the SCRIPTS_PATH.

   The default script timeout is 60 seconds.

9. In the **Notification** page, from the **Email preference** drop-down list, select the scenarios in which you want to send the emails.

   You must also specify the sender and receiver email addresses, and the subject of the email. If you want to attach the report of the clone operation performed, select **Attach Job Report**.

   > For email notification, you must have specified the SMTP server details using the either the GUI or the PowerShell command Set-SmSmtpServer.

   For EMS, you can refer to Manage EMS data collection

10. Review the summary, and then select **Finish**.
11. Monitor the operation progress by selecting **Monitor** > **Jobs**.

After you finish

After the clone is created, you should never rename it.

Related information

Clone operation might fail or take longer time to complete with default TCP_TIMEOUT value

The failover cluster instance database clone fails

PowerShell cmdlets

Steps

1. Initiate a connection session with the SnapCenter Server for a specified user by using the Open-SmConnection cmdlet.

```
Open-SmConnection  -SMSbaseurl  https://snapctr.demo.netapp.com:8146
```

2. List the backups that can be cloned by using the Get-SmBackup or Get-SmResourceGroup cmdlet.

   This example displays information about all available backups:

```
C:\PS>PS C:\> Get-SmBackup

BackupId   BackupName                      BackupTime   BackupType
--------   ----------                      ----------   ----------
1          Payroll Dataset_vise-f6_08...  8/4/2015     Full Backup
                                          11:02:32 AM

2          Payroll Dataset_vise-f6_08...  8/4/2015
                                          11:23:17 AM
```

   This example displays information about a specified resource group, its resources, and associated policies:

```
PS C:\> Get-SmResourceGroup -ListResources –ListPolicies

Description :
CreationTime : 8/4/2015 3:44:05 PM
ModificationTime : 8/4/2015 3:44:05 PM
EnableEmail : False
EmailSMTPServer :
EmailFrom :
EmailTo :
EmailSubject :
EnableSysLog : False
ProtectionGroupType : Backup
EnableAsupOnFailure : False
Policies : {FinancePolicy}
HostResourceMaping : {}
Configuration : SMCoreContracts.SmCloneConfiguration
LastBackupStatus :
VerificationServer :
EmailBody :
EmailNotificationPreference : Never
VerificationServerInfo : SMCoreContracts.SmVerificationServerInfo
```

```
SchedulerSQLInstance :
CustomText :
CustomSnapshotFormat :
SearchResources : False
ByPassCredential : False
IsCustomSnapshot :
MaintenanceStatus : Production
PluginProtectionGroupTypes : {SMSQL}
Name : Payrolldataset
Type : Group
Id : 1
Host :
UserName :
Passphrase :
Deleted : False
Auth : SMCoreContracts.SmAuth
IsClone : False
CloneLevel : 0
ApplySnapvaultUpdate : False
ApplyRetention : False
RetentionCount : 0
RetentionDays : 0
ApplySnapMirrorUpdate : False
SnapVaultLabel :
MirrorVaultUpdateRetryCount : 7
AppPolicies : {}
Description : FinancePolicy
PreScriptPath :
PreScriptArguments :
PostScriptPath :
PostScriptArguments :
ScriptTimeOut : 60000
DateModified : 8/4/2015 3:43:30 PM
DateCreated : 8/4/2015 3:43:30 PM
Schedule : SMCoreContracts.SmSchedule
PolicyType : Backup
PluginPolicyType : SMSQL
Name : FinancePolicy
Type :
Id : 1
Host :
UserName :
Passphrase :
Deleted : False
Auth : SMCoreContracts.SmAuth
IsClone : False
```

```
CloneLevel : 0
clab-a13-13.sddev.lab.netapp.com
DatabaseGUID :
SQLInstance : clab-a13-13
DbStatus : AutoClosed
DbAccess : eUndefined
IsSystemDb : False
IsSimpleRecoveryMode : False
IsSelectable : True
SqlDbFileGroups : {}
SqlDbLogFiles : {}
AppFileStorageGroups : {}
LogDirectory :
AgName :
Version :
VolumeGroupIndex : -1
IsSecondary : False
Name : TEST
Type : SQL Database
Id : clab-a13-13\TEST
Host : clab-a13-13.sddev.mycompany.com
UserName :
Passphrase :
Deleted : False
Auth : SMCoreContracts.SmAuth
IsClone : False
```

3. Initiate a clone operation from an existing backup by using the New-SmClone cmdlet.

   This example creates a clone from a specified backup with all logs:

```
PS C:\> New-SmClone
-BackupName payroll_dataset_vise-f3_08-05-2015_15.28.28.9774
-Resources @{"Host"="vise-f3.sddev.mycompany.com";
"Type"="SQL Database";"Names"="vise-f3\SQLExpress\payroll"}
-CloneToInstance vise-f3\sqlexpress -AutoAssignMountPoint
-Suffix _clonefrombackup
-LogRestoreType All -Policy clonefromprimary_ondemand

PS C:> New-SmBackup -ResourceGroupName PayrollDataset -Policy
FinancePolicy
```

   This example creates a clone to a specified Microsoft SQL Server instance:

```
PS C:\> New-SmClone
-BackupName "BackupDS1_NY-VM-SC-SQL_12-08-2015_09.00.24.8367"
-Resources @{"host"="ny-vm-sc-sql";"Type"="SQL Database";
"Names"="ny-vm-sc-sql\AdventureWorks2012_data"}
-AppPluginCode SMSQL -CloneToInstance "ny-vm-sc-sql"
-Suffix _CLPOSH -AssignMountPointUnderPath "C:\SCMounts"
```

4. View the status of the clone job by using the Get-SmCloneReport cmdlet.

   This example displays a clone report for the specified job ID:

```
PS C:\> Get-SmCloneReport -JobId 186

SmCloneId : 1
SmJobId : 186
StartDateTime : 8/3/2015 2:43:02 PM
EndDateTime : 8/3/2015 2:44:08 PM
Duration : 00:01:06.6760000
Status : Completed
ProtectionGroupName : Draper
SmProtectionGroupId : 4
PolicyName : OnDemand_Clone
SmPolicyId : 4
BackupPolicyName : OnDemand_Full_Log
SmBackupPolicyId : 1
CloneHostName : SCSPR0054212005.mycompany.com
CloneHostId : 4
CloneName : Draper__clone__08-03-2015_14.43.53
SourceResources : {Don, Betty, Bobby, Sally}
ClonedResources : {Don_DRAPER, Betty_DRAPER, Bobby_DRAPER,
                   Sally_DRAPER}
```

The information regarding the parameters that can be used with the cmdlet and their descriptions can be obtained by running *Get-Help command_name*. Alternatively, you can also refer to the SnapCenter Software Cmdlet Reference Guide.

Perform Clone Lifecycle

Using SnapCenter, you can create clones from a resource group or database. You can either perform on-demand clone or you can schedule recurring clone operations of a resource group or database. If you clone a backup periodically, you can use the clone to develop applications, populate data, or recover data.

SnapCenter enables you to schedule multiple clone operations to run simultaneously across multiple servers.

Before you begin

- While cloning to a standalone database instance, ensure that the mount point path exists and it is a dedicated disk.
- While cloning to a Failover Cluster Instance (FCI), ensure that the mount points exists, it is a shared disk, and the path and the FCI should belong to the same SQL resource group.
- If the source database or the target instance is on a cluster shared volume (csv), then the cloned database will be on the csv.

For virtual environments (VMDK/RDM), ensure that the mount point is a dedicated disk.

About this task

- The SCRIPTS_PATH is defined using the PredefinedWindowsScriptsDirectory key located in the SMCoreServiceHost.exe.Config file of the plug-in host.

  If needed, you can change this path and restart SMcore service. It is recommended that you use the default path for security.

  The value of the key can be displayed from swagger through the API: API /4.7/configsettings

  You can use the GET API to display the value of the key. SET API is not supported.

- Most of the fields on the Clone lifecycle wizard pages are self-explanatory. The following information describes fields for which you might need guidance.
- For ONTAP 9.12.1 and below version, if you specify a Snapshot locking period, the clones created from the tamper proof Snapshots will inherit the SnapLock expiry time. Storage admin should manually cleanup the clones post the SnapLock expiry time.

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page, select either **Database** or **Resource Group** from the **View** list.
3. Select the resource group or database, and then click **Clone Lifecycle**.
4. In the Options page, perform the following actions:

For this field…	Do this…
Clone job name	Specify the clone life cycle job name that helps in monitoring and modifying the clone life cycle job.
Clone server	Choose the host on which the clone should be placed.
Clone instance	Choose the clone instance to which you want to clone the database. This SQL instance must be located in the specified clone server.

For this field…	Do this…
Clone suffix	Enter a suffix that will be appended to the clone database to identify that it is a clone. Each SQL instance that is used to create a clone resource group must have a unique database name. For example, if the clone resource group contains a source database “db1” from an SQL instance “inst1”, and if “db1” is cloned to “inst1”, then the clone database name should be “db1*clone”. “*clone” is a mandatory user-defined suffix because the database is cloned to the same instance. If “db1” is cloned to the SQL instance “inst2”, then the clone database name can remain “db1” (the suffix is optional) because the database is cloned to a different instance.
Auto assign mount point or Auto assign volume mount point under path	Choose whether to automatically assign a mount point or volume mount point under a path. Choosing to auto assign a volume mount point under a path enables you to provide a specific directory. The mount points will be created within that directory. Before you choose this option, you must ensure that the directory is empty. If there is a database in the directory, the database will be in an invalid state after the mount operation.

5. In the Location page, select a storage location to create a clone.
6. In the Script page, enter the path and the arguments of the prescript or postscript that should be run before or after the clone operation, respectively.

   For example, you can run a script to update SNMP traps, automate alerts, send logs, and so on.

   > The prescripts or postscripts path should not include drives or shares. The path should be relative to the SCRIPTS_PATH.

   The default script timeout is 60 seconds.

7. In the Schedule page, perform one of the following actions:
   - Select **Run now** if you want to execute the clone job immediately.
   - Select **Configure schedule** when you want to determine how frequently the clone operation should occur, when the clone schedule should start, on which day the clone operation should occur, when the schedule should expire, and whether the clones have to be deleted after the schedule expires.
8. In the Notification page, from the **Email preference** drop-down list, select the scenarios in which you want to send the emails.

   You must also specify the sender and receiver email addresses, and the subject of the email. If you want to attach the report of the clone operation performed, select **Attach Job Report**.

For email notification, you must have specified the SMTP server details using the either the GUI or the PowerShell command Set-SmSmtpServer.

For EMS, you can refer to Manage EMS data collection

9. Review the summary, and then click **Finish**.

You should monitor the cloning process using the **Monitor** > **Jobs** page.

Monitor SQL database clone operations

You can monitor the progress of SnapCenter clone operations by using the Jobs page. You might want to check the progress of an operation to determine when it is complete or if there is an issue.

About this task

The following icons appear on the Jobs page, and indicate the state of the operation:

- In progress
- Completed successfully
- Failed
- Completed with warnings or could not start due to warnings
- Queued
- Canceled

Steps

1. In the left navigation pane, click **Monitor**.
2. In the **Monitor** page, click **Jobs**.
3. In the **Jobs** page, perform the following steps:
   a. Click to filter the list so that only clone operations are listed.
   b. Specify the start and end dates.
   c. From the **Type** drop-down list, select **Clone**.
   d. From the **Status** drop-down list, select the clone status.
   e. Click **Apply** to view the operations that are completed successfully.
4. Select the clone job, and then click **Details** to view the job details.
5. In the Job Details page, click **View logs**.

Cancel SQL resource clone operations

You can cancel clone operations that are queued.

You should be logged in as the SnapCenter Admin or job owner to cancel clone operations.

About this task

- You can cancel a queued clone operation from either the **Monitor** page or the **Activity** pane.
- You cannot cancel a running clone operation.
- You can use the SnapCenter GUI, PowerShell cmdlets, or CLI commands to cancel the queued clone operations.
- If you selected **All members of this role can see and operate on other members objects** in Users\Groups page while creating a role, you can cancel the queued clone operations of other members while using that role.

Step

Perform one of the following actions:

From the…	Action
Monitor page	a. In the left navigation pane, click **Monitor** > **Jobs**. b. Select the operation, and click **Cancel Job**.
Activity pane	a. After initiating the clone operation, click on the Activity pane to view the five most recent operations. b. Select the operation. c. In the **Job Details** page, click **Cancel Job**.

Split a clone

You can use SnapCenter to split a cloned resource from the parent resource. The clone that is split becomes independent of the parent resource.

About this task

- You cannot perform the clone split operation on an intermediate clone.

  For example, after you create clone1 from a database backup, you can create a backup of clone1, and then clone this backup (clone2). After you create clone2, clone1 is an intermediate clone, and you cannot perform the clone split operation on clone1. However, you can perform the clone split operation on clone2.

  After splitting clone2, you can perform the clone split operation on clone1 because clone1 is no longer the intermediate clone.

- When you split a clone, the backup copies and clone jobs of the clone are deleted.
- For information about FlexClone volume split operations, see, Split a FlexClone volume from its parent volume.
- Ensure that the volume or aggregate on the storage system is online.

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. In the **Resources** page, select the appropriate option from the View list:

Option	Description
For database applications	Select **Database** from the View list.
For file systems	Select **Path** from the View list.

3. Select the appropriate resource from the list.

   The resource topology page is displayed.

4. From the **Manage Copies** view, select the cloned resource (for example, the database or LUN), and then click .
5. Review the estimated size of the clone that is to be split and the required space available on the aggregate, and then click **Start**.
6. Monitor the operation progress by clicking **Monitor** > **Jobs**.

   The clone split operation stops responding if the SMCore service restarts. You should run the Stop-SmJob cmdlet to stop the clone split operation, and then retry the clone split operation.

   If you want a longer poll time or shorter poll time to check whether the clone is split or not, you can change the value of *CloneSplitStatusCheckPollTime* parameter in *SMCoreServiceHost.exe.config* file to set the time interval for SMCore to poll for the status of the clone split operation. The value is in milliseconds and the default value is 5 minutes.

   For example:

```
<add key="CloneSplitStatusCheckPollTime" value="300000" />
```

   The clone split start operation fails if backup, restore, or another clone split is in progress. You should restart the clone split operation only after the running operations are complete.

Related information

SnapCenter clone or verification fails with aggregate does not exist

Protect SAP HANA databases

SnapCenter Plug-in for SAP HANA Databases

SnapCenter Plug-in for SAP HANA Database overview

The SnapCenter Plug-in for SAP HANA Database is a host-side component of the NetApp SnapCenter software that enables application-aware data protection management of SAP HANA databases. The Plug-in for SAP HANA Database automates the backup, restore, and cloning of SAP HANA databases in your SnapCenter environment.

SnapCenter supports single container and multitenant database containers (MDC). You can use the Plug-in for SAP HANA Database in both Windows and Linux environments. The plug-in that is not installed on the HANA database host is known as the centralized host plug-in. The centralized host plug-in can manage multiple HANA databases across different hosts.

When the Plug-in for SAP HANA Database is installed, you can use SnapCenter with NetApp SnapMirror technology to create mirror copies of backup sets on another volume. You can also use the plug-in with NetApp SnapVault technology to perform disk-to-disk backup replication for standards compliance.

The Plug-in for SAP HANA Database supports SnapMirror active sync (initially released as SnapMirror Business Continuity [SM-BC]) that enables business services to continue operating even through a complete site failure, supporting applications to fail over transparently using a secondary copy. Neither manual intervention nor additional scripting is required to trigger a failover with SnapMirror active sync.

What you can do using the SnapCenter Plug-in for SAP HANA Database

When you install the Plug-in for SAP HANA Database in your environment, you can use SnapCenter to back up, restore, and clone SAP HANA databases and their resources. You can also perform tasks supporting those operations.

- Add databases.
- Create backups.
- Restore from backups.
- Clone backups.
- Schedule backup operations.
- Monitor backup, restore, and clone operations.
- View reports for backup, restore, and clone operations.

SnapCenter Plug-in for SAP HANA Database features

SnapCenter integrates with the plug-in application and with NetApp technologies on the storage system. To work with the Plug-in for SAP HANA Database, you use the SnapCenter graphical user interface.

- **Unified graphical user interface**

The SnapCenter interface provides standardization and consistency across plug-ins and environments. The SnapCenter interface enables you to complete consistent backup, restore, and clone operations across plug-ins, use centralized reporting, use at-a-glance dashboard views, set up role-based access control (RBAC), and monitor jobs across all plug-ins.

- **Automated central administration**

  You can schedule backup operations, configure policy-based backup retention, and perform restore operations. You can also proactively monitor your environment by configuring SnapCenter to send email alerts.

- **Nondisruptive NetApp Snapshot copy technology**

  SnapCenter uses NetApp Snapshot technology with the Plug-in for SAP HANA Database to back up resources.

Using the Plug-in for SAP HANA Database also offers the following benefits:

- Support for backup, restore, and clone workflows
- RBAC-supported security and centralized role delegation

  You can also set the credentials so that the authorized SnapCenter users have application-level permissions.

- Creation of space-efficient and point-in-time copies of resources for testing or data extraction by using NetApp FlexClone technology

  A FlexClone license is required on the storage system where you want to create the clone.

- Support for the consistency group (CG) Snapshot feature of ONTAP as part of creating backups.
- Capability to run multiple backups simultaneously across multiple resource hosts

  In a single operation, Snapshots are consolidated when resources in a single host share the same volume.

- Capability to create Snapshots using external commands.
- Support for file-based backup.
- Support for Linux LVM on XFS file system.

Storage types supported by SnapCenter Plug-in for SAP HANA Database

SnapCenter supports a wide range of storage types on both physical machines and virtual machines (VMs). You must verify the support for your storage type before installing SnapCenter Plug-in for SAP HANA Database.

Machine	Storage type
Physical server	• FC-connected LUNs • iSCSI-connected LUNs • NFS-connected volumes

Machine	Storage type
VMware ESXi	• RDM LUNs connected by an FC or iSCSI ESXi HBAScanning of host bus adapters (HBAs) might take long time to complete because SnapCenter scans all the host bus adaptors present in the host. You can edit the **LinuxConfig.pm** file located at */opt/NetApp/snapcenter/spl/plugins/scu/scucore/modules/SCU/Config* to set the value of the **SCSI_HOSTS_OPTIMIZED_RESCAN** parameter to 1 to rescan only those HBAs that are listed in HBA_DRIVER_NAMES. • iSCSI LUNs connected directly to the guest system by the iSCSI initiator • VMDKs on NFS datastores • VMDKs on VMFS created • NFS volumes connected directly to the guest system • vVol datastores on both NFS and SAN vVol datastore can only be provisioned with ONTAP Tools for VMware vSphere.

Minimum ONTAP privileges required for SAP HANA plug-in

The minimum ONTAP privileges that are required vary according to the SnapCenter plug-ins you are using for data protection.

- All-access commands: Minimum privileges required for ONTAP 9.12.1 and later
  - event generate-autosupport-log
  - job history show
  - job stop
  - lun
  - lun create
  - lun create
  - lun create
  - lun delete
  - lun igroup add
  - lun igroup create
  - lun igroup delete
  - lun igroup rename

- lun igroup rename
- lun igroup show
- lun mapping add-reporting-nodes
- lun mapping create
- lun mapping delete
- lun mapping remove-reporting-nodes
- lun mapping show
- lun modify
- lun move-in-volume
- lun offline
- lun online
- lun persistent-reservation clear
- lun resize
- lun serial
- lun show
- snapmirror policy add-rule
- snapmirror policy modify-rule
- snapmirror policy remove-rule
- snapmirror policy show
- snapmirror restore
- snapmirror show
- snapmirror show-history
- snapmirror update
- snapmirror update-ls-set
- snapmirror list-destinations
- version
- volume clone create
- volume clone show
- volume clone split start
- volume clone split stop
- volume create
- volume destroy
- volume file clone create
- volume file show-disk-usage
- volume offline
- volume online
- volume modify

    - volume qtree create
    - volume qtree delete
    - volume qtree modify
    - volume qtree show
    - volume restrict
    - volume show
    - volume snapshot create
    - volume snapshot delete
    - volume snapshot modify
    - volume snapshot modify-snaplock-expiry-time
    - volume snapshot rename
    - volume snapshot restore
    - volume snapshot restore-file
    - volume snapshot show
    - volume unmount
    - vserver cifs
    - vserver cifs share create
    - vserver cifs share delete
    - vserver cifs shadowcopy show
    - vserver cifs share show
    - vserver cifs show
    - vserver export-policy
    - vserver export-policy create
    - vserver export-policy delete
    - vserver export-policy rule create
    - vserver export-policy rule show
    - vserver export-policy show
    - vserver iscsi
    - vserver iscsi connection show
    - vserver show
- Read-only commands: Minimum privileges required for ONTAP 8.3.0 and later
    - network interface
    - network interface show
    - vserver

Prepare storage systems for SnapMirror and SnapVault replication for SAP HANA databases

You can use a SnapCenter plug-in with ONTAP SnapMirror technology to create mirror copies of backup sets on another volume, and with ONTAP SnapVault technology to perform disk-to-disk backup replication for standards compliance and other governance-related purposes. Before you perform these tasks, you must configure a data-protection relationship between the source and destination volumes and initialize the relationship.

SnapCenter performs the updates to SnapMirror and SnapVault after it completes the Snapshot operation. SnapMirror and SnapVault updates are performed as part of the SnapCenter job; do not create a separate ONTAP schedule.

If you are coming to SnapCenter from a NetApp SnapManager product and are satisfied with the data protection relationships you have configured, you can skip this section.

A data protection relationship replicates data on primary storage (the source volume) to secondary storage (the destination volume). When you initialize the relationship, ONTAP transfers the data blocks referenced on the source volume to the destination volume.

SnapCenter does not support cascade relationships between SnapMirror and SnapVault volumes (**Primary** > **Mirror** > **Vault**). You should use fanout relationships.

SnapCenter supports the management of version-flexible SnapMirror relationships. For details about version-flexible SnapMirror relationships and how to set them up, see the ONTAP documentation.

Backup strategy for SAP HANA databases

Define a backup strategy for SAP HANA databases

Defining a backup strategy before you create your backup jobs helps you to have the backups that you require to successfully restore or clone your resources. Your service-level agreement (SLA), recovery time objective (RTO), and recovery point objective (RPO) largely determine your backup strategy.

About this task

An SLA defines the level of service that is expected and addresses many service-related issues, including the availability and performance of the service. RTO is the time by which a business process must be restored after a disruption in service. RPO defines the strategy for the age of the files that must be recovered from backup storage for regular operations to resume after a failure. SLA, RTO, and RPO contribute to the data protection strategy.

Steps

1. Determine when you should back up your resources.
2. Decide how many backup jobs you require.
3. Decide how to name your backups.
4. Decide whether you want to create a Snapshot copy-based policy to back up application-consistent Snapshots of the database.

5. Decide whether you want to verify the integrity of the database.
6. Decide whether you want to use NetApp SnapMirror technology for replication or NetApp SnapVault technology for long-term retention.
7. Determine the retention period for the Snapshots on the source storage system and the SnapMirror destination.
8. Determine whether you want to run any commands before or after the backup operation and provide a prescript or postscript.

Automatic discovery of resources on Linux host

Resources are SAP HANA databases and Non-data Volume on the Linux host that are managed by SnapCenter. After installing the SnapCenter Plug-in for SAP HANA Database plug-in, the SAP HANA databases on that Linux host are automatically discovered and displayed in the Resources page.

Automatic discovery is supported for the following SAP HANA resources:

- Single containers

  After installing or upgrading the plug-in, the single container resources located on a centralized host plug-in will continue as manually added resources.

  After installing or upgrading the plug-in, the SAP HANA databases are automatically discovered only on the SAP HANA Linux hosts, which are directly registered into SnapCenter.

- Multitenant database container (MDC)

  After installing or upgrading the plug-in, the MDC resources located on a centralized host plug-in will continue as manually added resource.

  You must continue to manually add the MDC resources on the centralized host plug-in after upgrading to SnapCenter 4.3.

  For SAP HANA Linux hosts directly registered in SnapCenter, installing or upgrading the plug-in will trigger an automatic discovery for resources on the host. After upgrading the plug-in, for every MDC resource that was located on the plug-in host, another MDC resource will be automatically discovered with a different GUID format and registered in SnapCenter. The new resource will be in locked state.

  For example, in SnapCenter 4.2, if E90 MDC resource was located on the plug-in host and registered manually, after upgrading to SnapCenter 4.3, another E90 MDC resource with a different GUID will be discovered and registered in SnapCenter.

Automatic discovery is not supported for the following configurations:

- RDM and VMDK layouts

  In case the above resources are discovered, the data protection operations are not supported on these resources.

- HANA multiple-host configuration
- Multiple instances on the same host

- Multitier scale out HANA System Replication
- Cascaded replication environment in System Replication mode

Type of backups supported

Backup type specifies the type of backup that you want to create. SnapCenter supports File-Based Backup and Snapshot copy-based backup types for SAP HANA databases.

File-Based Backup

File-Based Backups verify the integrity of the database. You can schedule the file-based backup operation to occur at specific intervals. Only active tenants are backed up. You cannot restore and clone File-Based backups from SnapCenter.

Snapshot copy based backup

Snapshot copy-based backups leverage NetApp Snapshot technology to create online, read-only copies of the volumes on which the SAP HANA databases reside.

How SnapCenter Plug-in for SAP HANA Database uses consistency group Snapshots

You can use the plug-in to create consistency group Snapshots for resource groups. A consistency group is a container that can house multiple volumes so that you can manage them as one entity. A consistency group is simultaneous Snapshots of multiple volumes, providing consistent copies of a group of volumes.

You can also specify the wait time for the storage controller to consistently group Snapshots. The available wait time options are **Urgent**, **Medium**, and **Relaxed**. You can also enable or disable Write Anywhere File Layout (WAFL) sync during consistent group Snapshot operation. WAFL sync improves the performance of a consistency group Snapshot.

How SnapCenter manages housekeeping of log and data backups

SnapCenter manages the housekeeping of log and data backups on the storage system and file system levels, and within the SAP HANA backup catalog.

The Snapshots on the primary or secondary storage and their corresponding entries in the SAP HANA catalog are deleted based on the retention settings. The SAP HANA catalog entries are also deleted during backup and resource group deletion.

Considerations for determining backup schedules for SAP HANA database

The most critical factor in determining a backup schedule is the rate of change for the resource. You might back up a heavily used resource every hour, while you might back up a rarely used resource once a day. Other factors include the importance of the resource to your organization, your service-level agreement (SLA) and your recovery point objective (RPO).

Backup schedules have two parts, as follows:

- Backup frequency (how often backups are to be performed)

Backup frequency, also called schedule type for some plug-ins, is part of a policy configuration. For example, you might configure the backup frequency as hourly, daily, weekly, or monthly.

- Backup schedules (exactly when backups are to be performed

  Backup schedules are part of a resource or resource group configuration. For example, if you have a resource group that has a policy configured for weekly backups, you might configure the schedule to back up every Thursday at 10:00 p.m.

Number of backup jobs needed for SAP HANA databases

Factors that determine the number of backup jobs that you need include the size of the resource, the number of volumes used, the rate of change of the resource, and your Service Level Agreement (SLA).

Backup naming conventions for Plug-in for SAP HANA databases

You can either use the default Snapshot naming convention or use a customized naming convention. The default backup naming convention adds a timestamp to Snapshot names that helps you identify when the copies were created.

The Snapshot uses the following default naming convention:

```
resourcegroupname_hostname_timestamp
```

You should name your backup resource groups logically, as in the following example:

```
dts1_mach1x88_03-12-2015_23.17.26
```

In this example, the syntax elements have the following meanings:

- *dts1* is the resource group name.
- *mach1x88* is the host name.
- *03-12-2015_23.17.26* is the date and timestamp.

Alternatively, you can specify the Snapshot name format while protecting resources or resource groups by selecting **Use custom name format for Snapshot copy**. For example, customtext_resourcegroup_policy_hostname or resourcegroup_hostname. By default, the time stamp suffix is added to the Snapshot name.

Restore and recovery strategy for SAP HANA databases

Define a restore and recovery strategy for SAP HANA resources

You must define a strategy before you restore and recover your database so that you can perform restore and recovery operations successfully.

Steps

1. Determine the restore strategies supported for manually added SAP HANA resources

2. Determine the restore strategies supported for auto discovered SAP HANA databases
3. Decide the type of recovery operations that you want to perform.

Types of restore strategies supported for manually added SAP HANA resources

You must define a strategy before you can successfully perform restore operations using SnapCenter. There are two types of restore strategies for manually added SAP HANA resources. You cannot recover manually added SAP HANA resources.

You cannot recover manually added SAP HANA resources.

Complete resource restore

- Restores all volumes, qtrees, and LUNs of a resource

If the resource contains volumes or qtrees, the Snapshots taken after the Snapshot selected for restore on such volumes or qtrees are deleted and cannot be recovered. Also, if any other resource is hosted on the same volumes or qtrees, then that resource is also deleted.

File level restore

- Restores files from volumes, qtrees, or directories
- Restores only the selected LUNs

Types of restore strategies supported for automatically discovered SAP HANA databases

You must define a strategy before you can successfully perform restore operations using SnapCenter. There are two types of restore strategies for automatically discovered SAP HANA databases.

Complete resource restore

- Restores all volumes, qtrees, and LUNs of a resource
  - The **Volume Revert** option should be selected to restore the entire volume.

If the resource contains volumes or qtrees, the Snapshots taken after the Snapshot selected for restore on such volumes or qtrees are deleted and cannot be recovered. Also, if any other resource is hosted on the same volumes or qtrees, then that resource is also deleted.

Tenant Database

- Restores the tenant database

If **Tenant Database** option is selected, then HANA studio or HANA recovery scripts external to SnapCenter must be used to perform the recovery operation.

Types of restore operations for auto discovered SAP HANA databases

SnapCenter supports volume-based SnapRestore (VBSR), Single File SnapRestore, and connect-and-copy restore types for automatically discovered SAP HANA databases.

Volume-based SnapRestore (VBSR) is performed in NFS environments for the following scenarios:

- When the backup selected for restore is taken on releases earlier than SnapCenter 4.3, and only if the **Complete Resource**option is selected
- When the backup selected for restore is taken in SnapCenter 4.3, and if the **Volume Revert** option is selected

Single File SnapRestore is performed in NFS environments for the following scenarios:

- When the backup selected for restore is taken in SnapCenter 4.3, and if only the **Complete Resource** option is selected
- For multitenant database containers (MDC), when the backup selected for restore is taken on SnapCenter 4.3, and the **Tenant Database** option is selected
- When the backup selected is from a SnapMirror or SnapVault secondary location, and the **Complete Resource** option is selected

Single File SnapRestore is performed in SAN environments for the following scenarios:

- When backups are taken on releases earlier than SnapCenter 4.3, and only if the **Complete Resource** option is selected
- When backups are taken in SnapCenter 4.3, and only if the **Complete Resource** option is selected
- When the backup is selected from a SnapMirror or SnapVault secondary location, and the **Complete Resource** option is selected

Connect-and-copy based restore is performed in SAN environments for the following scenario:

- For MDC, when the backup selected for restore is taken in SnapCenter 4.3, and the **Tenant Database** option is selected

Complete Resource, **Volume Revert**, and **Tenant Database** options are available on the Restore Scope page.

Types of recovery operations supported for SAP HANA databases

SnapCenter enables you to perform different types of recovery operations for SAP HANA databases.

- Recover the database up to the most recent state
- Recover the database up to a specific point in time

  You must specify the date and time for recovery.

- Recover the database up to a specific data backup

SnapCenter also provides the No recovery option for SAP HANA databases.

Prepare to install the SnapCenter Plug-in for SAP HANA Database

Installation workflow of SnapCenter Plug-in for SAP HANA Database

You should install and set up the SnapCenter Plug-in for SAP HANA Database if you want to protect SAP HANA databases.

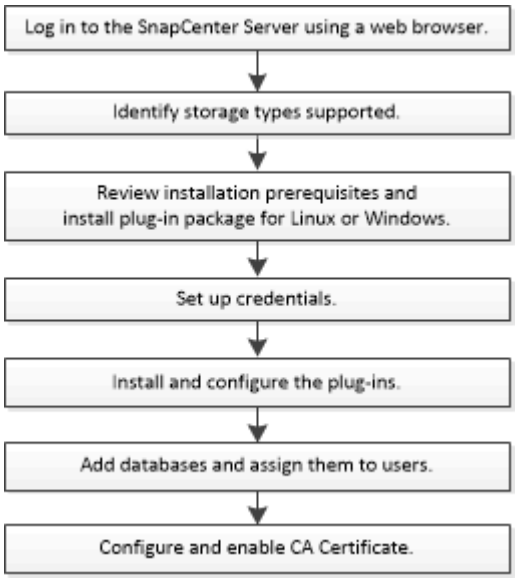

Prerequisites for adding hosts and installing SnapCenter Plug-in for SAP HANA Database

Before you add a host and install the plug-in packages, you must complete all the requirements. SnapCenter Plug-in for SAP HANA Database is available in both Windows and Linux environments.

- You must have installed Java 11 on your host.

  IBM Java is not supported on Windows and Linux host.

- You must have installed SAP HANA database interactive terminal (HDBSQL client) on the host.
- For Windows, plug-in Creator Service should be running using the “LocalSystem” windows user, which is the default behavior when Plug-in for SAP HANA Database is installed as domain administrator.
- For Windows, user store keys should be created as SYSTEM user.
- When installing a plug-in on a Windows host, if you specify a credential that is not built-in or if the user belongs to a local workgroup user, you must disable UAC on the host. SnapCenter Plug-in for Microsoft Windows will be deployed by default with the SAP HANA plug-in on Windows hosts.
- For Linux host, HDB Secure User Store keys are accessed as HDBSQL OS user.
- SnapCenter Server should have access to the 8145 or custom port of Plug-in for SAP HANA Database host.

Windows hosts

- You must have a domain user with local administrator privileges with local login permissions on the remote host.
- While installing Plug-in for SAP HANA Database on a Windows host, SnapCenter Plug-in for Microsoft Windows is installed automatically.
- You must have enabled the password-based SSH connection for the root or non-root user.
- You must have installed Java 11 on your Windows host.

  Download JAVA for All Operating Systems

  NetApp Interoperability Matrix Tool

Linux hosts

- You must have enabled the password-based SSH connection for the root or non-root user.
- You must have installed Java 11 on your Linux host.

  Download JAVA for All Operating Systems

  NetApp Interoperability Matrix Tool

- For SAP HANA databases that are running on a Linux host, while installing Plug-in for SAP HANA Database, SnapCenter Plug-in for UNIX is installed automatically.
- You should have **bash** as the default shell for plug-in installation.

Supplemental commands

To run a supplemental command on the SnapCenter Plug-in for SAP HANA, you must include it in the *allowed_commands.config* file.

- Default location on the Windows host: *C:\Program Files\NetApp\SnapCenter\Snapcenter Plug-in Creator\etc\allowed_commands.config*
- Default location on the Linux host: */opt/NetApp/snapcenter/scc/etc/allowed_commands.config*

To allow supplemental commands on the plug-in host, open *allowed_commands.config* file in an editor. Enter each command on a separate line and the commands are not case sensitive. Ensure that you specify the fully qualified pathname and enclose the pathname in quotation marks (") if it contains spaces.

For example:

command: mount

command: umount

command: "C:\Program Files\NetApp\SnapCreator commands\sdcli.exe"

command: myscript.bat

If the *allowed_commands.config* file is not present, the commands or script execution will be blocked and the workflow will fail with the following error:

"[/mnt/mount -a] execution not allowed. Authorize by adding the command in the file %s on the plugin host."

If the command or script is not present in the *allowed_commands.config*, the command or script execution will be blocked and the workflow will fail with the following error:

"[/mnt/mount -a] execution not allowed. Authorize by adding the command in the file %s on the plugin host."

You should not use a wildcard entry (*) to allow all commands.

Configure sudo privileges for non-root users for Linux host

SnapCenter 2.0 and later releases allow a non-root user to install the SnapCenter Plug-ins Package for Linux and to start the plug-in process. The plug-in processes will be running as an effective non-root user. You should configure sudo privileges for the non-root user to provide access to several paths.

What you will need

- Sudo version 1.8.7 or later.
- If the umask is 0027, ensure that the java folder and all the files inside should have permission of 555. Otherwise plug-in installation might fail.
- For the non-root user, ensure that the name of the non-root user and the user's group should be the same.
- Edit the */etc/ssh/sshd_config* file to configure the message authentication code algorithms: MACs hmac-sha2-256 and MACs hmac-sha2-512.

  Restart the sshd service after updating the configuration file.

  Example:

```
#Port 22
#AddressFamily any
#ListenAddress 0.0.0.0
#ListenAddress ::
#Legacy changes
#KexAlgorithms diffie-hellman-group1-sha1
#Ciphers aes128-cbc
#The default requires explicit activation of protocol
Protocol 2
HostKey/etc/ssh/ssh_host_rsa_key
MACs hmac-sha2-256
```

About this task

You should configure sudo privileges for the non-root user to provide access to the following paths:

- /home/*LINUX_USER*/.sc_netapp/snapcenter_linux_host_plugin.bin
- /custom_location/NetApp/snapcenter/spl/installation/plugins/uninstall
- /custom_location/NetApp/snapcenter/spl/bin/spl

Steps

1. Log in to the Linux host on which you want to install the SnapCenter Plug-ins Package for Linux.
2. Add the following lines to the /etc/sudoers file by using the visudo Linux utility.

```
Cmnd_Alias HPPLCMD = sha224:checksum_value== /home/
LINUX_USER/.sc_netapp/snapcenter_linux_host_plugin.bin,
/opt/NetApp/snapcenter/spl/installation/plugins/uninstall,
/opt/NetApp/snapcenter/spl/bin/spl, /opt/NetApp/snapcenter/scc/bin/scc
Cmnd_Alias PRECHECKCMD = sha224:checksum_value== /home/
LINUX_USER/.sc_netapp/Linux_Prechecks.sh
Cmnd_Alias CONFIGCHECKCMD = sha224:checksum_value==
/opt/NetApp/snapcenter/spl/plugins/scu/scucore/configurationcheck/Config
_Check.sh
Cmnd_Alias SCCMD = sha224:checksum_value==
/opt/NetApp/snapcenter/spl/bin/sc_command_executor
Cmnd_Alias SCCCMDEXECUTOR =checksum_value==
/opt/NetApp/snapcenter/scc/bin/sccCommandExecutor
LINUX_USER ALL=(ALL) NOPASSWD:SETENV: HPPLCMD, PRECHECKCMD,
CONFIGCHECKCMD, SCCCMDEXECUTOR, SCCMD
Defaults: LINUX_USER env_keep += "IATEMPDIR"
Defaults: LINUX_USER env_keep += "JAVA_HOME"
Defaults: LINUX_USER !visiblepw
Defaults: LINUX_USER !requiretty
```

If you are having a RAC setup, along with the other allowed commands, you should add the following to the /etc/sudoers file: '/<crs_home>/bin/olsnodes'

You can obtain the value of *crs_home* from the */etc/oracle/olr.loc* file.

LINUX_USER is the name of the non-root user that you created.

You can obtain the *checksum_value* from the **sc_unix_plugins_checksum.txt** file, which is located at:

- *C:\ProgramData\NetApp\SnapCenter\Package Repository\sc_unix_plugins_checksum.txt* if SnapCenter Server is installed on Windows host.
- */opt/NetApp/snapcenter/SnapManagerWeb/Repository/sc_unix_plugins_checksum.txt* if SnapCenter Server in installed on Linux host.

.

The example should be used only as a reference for creating your own data.

Host requirements to install SnapCenter Plug-ins Package for Windows

Before you install the SnapCenter Plug-ins Package for Windows, you should be familiar with some basic host system space requirements and sizing requirements.

Item	Requirements
Operating Systems	Microsoft Windows For the latest information about supported versions, see the NetApp Interoperability Matrix Tool.
Minimum RAM for the SnapCenter plug-in on host	1 GB
Minimum install and log space for the SnapCenter plug-in on host	5 GB You should allocate sufficient disk space and monitor the storage consumption by the logs folder. The log space required varies depending on the number of the entities to be protected and the frequency of data protection operations. If there is no sufficient disk space, the logs will not be created for the recently run operations.
Required software packages	• ASP.NET Core Runtime 8.0.12 (and all subsequent 8.0.x patches) Hosting Bundle • PowerShell Core 7.4.2 • Java 11 Oracle Java and OpenJDK Java 11 Oracle Java and OpenJDK is required only for SAP HANA, IBM Db2, PostgreSQL, MySQL, NetApp supported plug-ins, and other custom applications which can be installed on Windows host. For the latest information about supported versions, see the NetApp Interoperability Matrix Tool. For .NET specific troubleshooting information, see SnapCenter upgrade or install fails for legacy systems that do not have internet connectivity.

Host requirements for installing the SnapCenter Plug-ins Package for Linux

Before you install the SnapCenter Plug-ins Package for Linux, you should be familiar with some basic host system space and sizing requirements.

Item	Requirements
Operating systems	• Red Hat Enterprise Linux • SUSE Linux Enterprise Server (SLES) For the latest information about supported versions, see the NetApp Interoperability Matrix Tool.
Minimum RAM for the SnapCenter plug-in on host	1 GB
Minimum install and log space for the SnapCenter plug-in on host	2 GB You should allocate sufficient disk space and monitor the storage consumption by the logs folder. The log space required varies, depending on the number of the entities to be protected and the frequency of data protection operations. If there is no sufficient disk space, the logs will not be created for the recently run operations.
Required software packages	Java 11 Oracle Java and OpenJDK If you have upgraded JAVA to the latest version, you must ensure that the JAVA_HOME option located at /var/opt/snapcenter/spl/etc/spl.properties is set to the correct JAVA version and the correct path. For the latest information about supported versions, see the NetApp Interoperability Matrix Tool.

Set up credentials for the SnapCenter Plug-in for SAP HANA Database

SnapCenter uses credentials to authenticate users for SnapCenter operations. You should create credentials for installing SnapCenter plug-ins and additional credentials for performing data protection operations on databases or Windows file systems.

About this task

- Linux hosts

  You must set up credentials for installing plug-ins on Linux hosts.

  You must set up the credentials for the root user or for a non-root user who has sudo privileges to install and start the plug-in process.

  Best Practice: Although you are allowed to create credentials for Linux after deploying hosts and installing plug-ins, the best practice is to create credentials after you add SVMs, before you deploy hosts and install plug-ins.

- Windows hosts

  You must set up Windows credentials before installing plug-ins.

  You must set up the credentials with administrator privileges, including administrator rights on the remote host.

If you set up credentials for individual resource groups and the username does not have full admin privileges, you must assign at least the resource group and backup privileges to the username.

Steps

1. In the left navigation pane, click **Settings**.
2. In the Settings page, click **Credential**.
3. Click **New**.

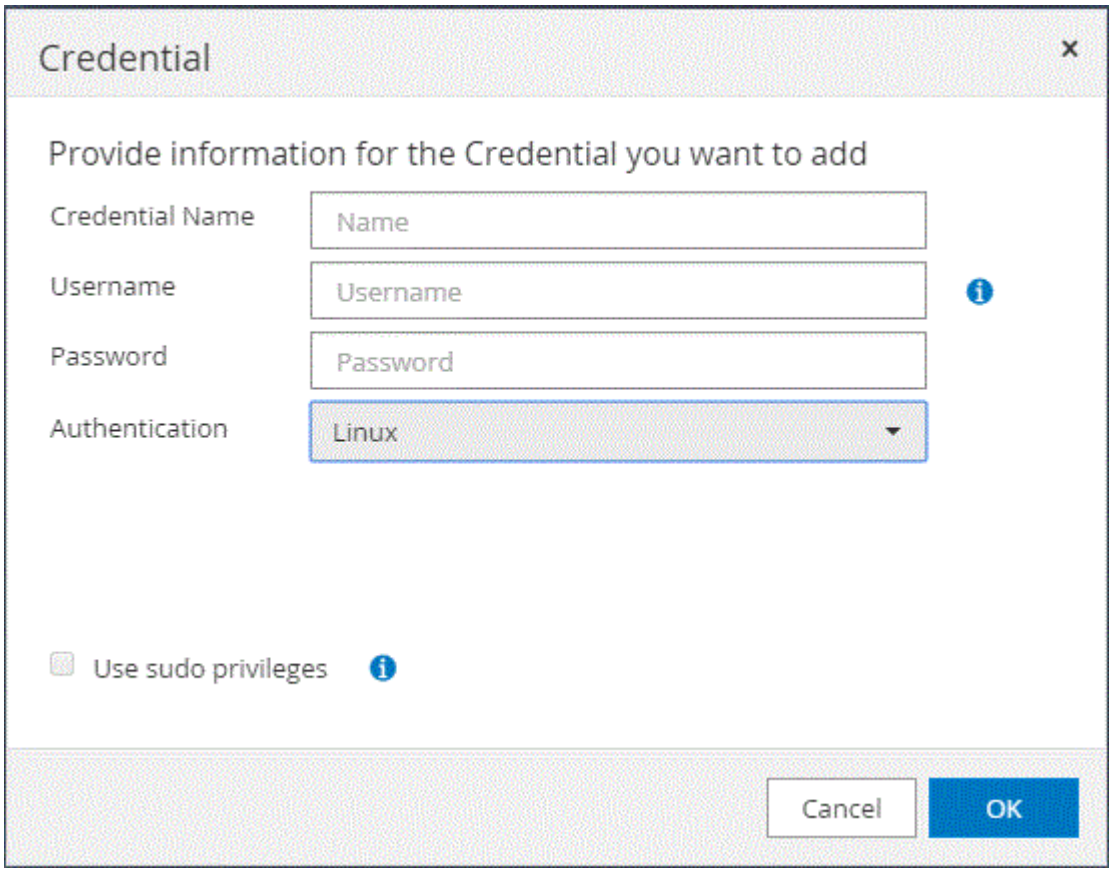

4. In the Credential page, specify the information required for configuring credentials:

For this field…	Do this…
Credential name	Enter a name for the credentials.

For this field…	Do this…
User name	Enter the user name and password that are to be used for authentication. • Domain administrator or any member of the administrator group Specify the domain administrator or any member of the administrator group on the system on which you are installing the SnapCenter plug-in. Valid formats for the Username field are: ◦ *NetBIOS\UserName* ◦ *Domain FQDN\UserName* • Local administrator (for workgroups only) For systems that belong to a workgroup, specify the built-in local administrator on the system on which you are installing the SnapCenter plug-in. You can specify a local user account that belongs to the local administrators group if the user account has elevated privileges or the User Access control feature is disabled on the host system. The valid format for the Username field is: *UserName* Do not use double quotes (") or backtick (`) in the passwords. You should not use the less than (<) and exclamation (!) symbols together in passwords. For example, lessthan<!10, lessthan10<!, backtick`12.
Password	Enter the password used for authentication.
Authentication Mode	Select the authentication mode that you want to use.
Use sudo privileges	Select the **Use sudo privileges** check box if you are creating credentials for a non-root user. Applicable to Linux users only.

5. Click **OK**.

After you finish setting up the credentials, you might want to assign credential maintenance to a user or group of users in the User and Access page.

Configure gMSA on Windows Server 2016 or later

Windows Server 2016 or later enables you to create a group Managed Service Account (gMSA) that provides automated service account password management from a managed domain account.

Before you begin

- You should have a Windows Server 2016 or later domain controller.
- You should have a Windows Server 2016 or later host, which is a member of the domain.

Steps

1. Create a KDS root key to generate unique passwords for each object in your gMSA.
2. For each domain, run the following command from the Windows domain controller: Add-KDSRootKey -EffectiveImmediately
3. Create and configure your gMSA:
   a. Create a user group account in the following format:

      ```
      domainName\accountName$
      ```

   b. Add computer objects to the group.
   c. Use the user group you just created to create the gMSA.

      For example,

      ```
      New-ADServiceAccount -name <ServiceAccountName> -DNSHostName <fqdn>
      -PrincipalsAllowedToRetrieveManagedPassword <group>
      -ServicePrincipalNames <SPN1,SPN2,…>
      ```

   d. Run `Get-ADServiceAccount` command to verify the service account.
4. Configure the gMSA on your hosts:
   a. Enable the Active Directory module for Windows PowerShell on the host where you want to use the gMSA account.

      To do this, run the following command from PowerShell:

```
PS C:\> Get-WindowsFeature AD-Domain-Services

Display Name                            Name                Install
State
------------                            ----
-------------
[ ] Active Directory Domain Services   AD-Domain-Services  Available


PS C:\> Install-WindowsFeature AD-DOMAIN-SERVICES

Success Restart Needed Exit Code      Feature Result
------- -------------- ---------      --------------
True    No             Success        {Active Directory Domain
Services, Active ...
WARNING: Windows automatic updating is not enabled. To ensure that
your newly-installed role or feature is
automatically updated, turn on Windows Update.
```

b. Restart your host.

c. Install the gMSA on your host by running the following command from the PowerShell command prompt: `Install-AdServiceAccount <gMSA>`

d. Verify your gMSA account by running the following command: `Test-AdServiceAccount <gMSA>`

5. Assign the administrative privileges to the configured gMSA on the host.
6. Add the Windows host by specifying the configured gMSA account in the SnapCenter Server.

   SnapCenter Server will install the selected plug-ins on the host and the specified gMSA will be used as the service log on account during the plug-in installation.

Install the SnapCenter Plug-in for SAP HANA Databases

Add hosts and install plug-in packages on remote hosts

You must use the SnapCenter Add Host page to add hosts, and then install the plug-ins packages. The plug-ins are automatically installed on the remote hosts. You can add a host and install plug-in packages either for an individual host or for a cluster.

Before you begin

- If the operating system of the SnapCenter Server host is Windows 2019 and the operating system of the plug-in host is Windows 2022, you should perform the following:
  - Upgrade to Windows Server 2019 (OS Build 17763.5936) or later
  - Upgrade to Windows Server 2022 (OS Build 20348.2402) or later
- You must be a user that is assigned to a role that has the plug-in install and uninstall permissions, such as the SnapCenter Admin role.

- When installing a plug-in on a Windows host, if you specify a credential that is not built-in, or if the user belongs to a local workgroup user, you must disable UAC on the host.
- You should ensure that the message queueing service is running.
- The administration documentation contains information about managing hosts.
- If you are using group Managed Service Account (gMSA), you should configure gMSA with administrative privileges.

  Configure group Managed Service Account on Windows Server 2016 or later for SAP HANA

About this task

- You cannot add a SnapCenter Server as a plug-in host to another SnapCenter Server.
- For SAP HANA System Replication to discover resources on both primary and secondary systems, it is recommended to add both the primary and the secondary systems using root or sudo user.

Steps

1. In the left navigation pane, click **Hosts**.
2. Verify that the **Managed Hosts** tab is selected at the top.
3. Click **Add**.
4. In the Hosts page, perform the following actions:

For this field…	Do this…
Host Type	Select the type of host: • Windows • Linux The Plug-in for SAP HANA is installed on the HDBSQL client host, and this host can be on either a Windows system or a Linux system.
Host name	Enter the communication host name. Enter the fully qualified domain name (FQDN) or the IP address of the host. SnapCenter depends on the proper configuration of the DNS. Therefore, the best practice is to enter the FQDN. You must configure the HDBSQL client and HDBUserStore on this host.

For this field…	Do this…
Credentials	Either select the credential name that you created or create new credentials. The credential must have administrative rights on the remote host. For details, see the information about creating credentials. You can view details about the credentials by positioning your cursor over the credential name that you provided. The credentials authentication mode is determined by the host type that you specify in the Add Host wizard.

5. In the Select Plug-ins to Install section, select the plug-ins to install.
6. (Optional) Click **More Options**.

For this field…	Do this…
Port	Either retain the default port number or specify the port number. The default port number is 8145. If the SnapCenter Server was installed on a custom port, that port number will be displayed as the default port. If you manually installed the plug-ins and specified a custom port, you must specify the same port. Otherwise, the operation fails.
Installation Path	The Plug-in for SAP HANA is installed on the HDBSQL client host, and this host can be on either a Windows system or a Linux system. • For the SnapCenter Plug-ins Package for Windows, the default path is C:\Program Files\NetApp\SnapCenter. Optionally, you can customize the path. • For the SnapCenter Plug-ins Package for Linux, the default path is /opt/NetApp/snapcenter. Optionally, you can customize the path.
Skip preinstall checks	Select this check box if you already installed the plug-ins manually and you do not want to validate whether the host meets the requirements for installing the plug-in.

For this field…	Do this…
Use group Managed Service Account (gMSA) to run the plug-in services	For Windows host, select this check box if you want to use group Managed Service Account (gMSA) to run the plug-in services. Provide the gMSA name in the following format: domainName\accountName$. gMSA will be used as a log on service account only for SnapCenter Plug-in for Windows service.

7. Click **Submit**.

   If you have not selected the Skip prechecks checkbox, the host is validated to verify whether the host meets the requirements for installing the plug-in. The disk space, RAM, PowerShell version, .NET version, location (for Windows plug-ins), and Java version (for Linux plug-ins) are validated against the minimum requirements. If the minimum requirements are not met, appropriate error or warning messages are displayed.

   If the error is related to disk space or RAM, you can update the web.config file located at C:\Program Files\NetApp\SnapCenter WebApp to modify the default values. If the error is related to other parameters, you must fix the issue.

   In an HA setup, if you are updating web.config file, you must update the file on both nodes.

8. If host type is Linux, verify the fingerprint, and then click **Confirm and Submit**.

   In a cluster setup, you should verify the fingerprint of each of the nodes in the cluster.

   Fingerprint verification is mandatory even if the same host was added earlier to SnapCenter and the fingerprint was confirmed.

9. Monitor the installation progress.

   The installation-specific log files are located at /custom_location/snapcenter/logs.

Install SnapCenter Plug-in Packages for Linux or Windows on multiple remote hosts by using cmdlets

You can install the SnapCenter Plug-in Packages for Linux or Windows on multiple hosts simultaneously by using the Install-SmHostPackage PowerShell cmdlet.

Before you begin

You must have logged in to SnapCenter as a domain user with local administrator rights on each host on which you want to install the plug-in package.

Steps

1. Launch PowerShell.

2. On the SnapCenter Server host, establish a session using the Open-SmConnection cmdlet, and then enter your credentials.
3. Install the plug-in on multiple hosts using the Install-SmHostPackage cmdlet and the required parameters.

   The information regarding the parameters that can be used with the cmdlet and their descriptions can be obtained by running *Get-Help command_name*. Alternatively, you can also refer to the SnapCenter Software Cmdlet Reference Guide.

   You can use the -skipprecheck option when you have installed the plug-ins manually and do not want to validate whether the host meets the requirements to install the plug-in.

4. Enter your credentials for remote installation.

Install the SnapCenter Plug-in for SAP HANA Database on Linux hosts by using the command-line interface

You should install the SnapCenter Plug-in for SAP HANA Database by using the SnapCenter user interface (UI). If your environment does not allow remote installation of the plug-in from the SnapCenter UI, you can install the Plug-in for SAP HANA Database either in console mode or in silent mode by using the command-line interface (CLI).

Before you begin

- You should install the Plug-in for SAP HANA Database on each of the Linux host where the HDBSQL client resides.
- The Linux host on which you are installing the SnapCenter Plug-in for SAP HANA Database must meet the dependent software, database, and operating system requirements.

  The Interoperability Matrix Tool (IMT) contains the latest information about the supported configurations.

  https://imt.netapp.com/matrix/imt.jsp?components=121029;&solution=1259&isHWU&src=IMT

- The SnapCenter Plug-in for SAP HANA Database is part of SnapCenter Plug-ins Package for Linux. Before you install SnapCenter Plug-ins Package for Linux, you should have already installed SnapCenter on a Windows host.

Steps

1. Copy the SnapCenter Plug-ins Package for Linux installation file (snapcenter_linux_host_plugin.bin) from C:\ProgramData\NetApp\SnapCenter\Package Repository to the host where you want to install the Plug-in for SAP HANA Database.

   You can access this path from the host where the SnapCenter Server is installed.

2. From the command prompt, navigate to the directory where you copied the installation file.
3. Install the plug-in: `path_to_installation_bin_file/snapcenter_linux_host_plugin.bin -i silent -DPORT=port_number_for_host -DSERVER_IP=server_name_or_ip_address -DSERVER_HTTPS_PORT=port_number_for_server`

   - -DPORT specifies the SMCore HTTPS communication port.
   - -DSERVER_IP specifies the SnapCenter Server IP address.
   - -DSERVER_HTTPS_PORT specifies the SnapCenter Server HTTPS port.

- -DUSER_INSTALL_DIR specifies the directory where you want to install the SnapCenter Plug-ins Package for Linux.
- DINSTALL_LOG_NAME specifies the name of the log file.

```
/tmp/sc-plugin-installer/snapcenter_linux_host_plugin.bin -i silent
-DPORT=8145 -DSERVER_IP=scserver.domain.com -DSERVER_HTTPS_PORT=8146
-DUSER_INSTALL_DIR=/opt
-DINSTALL_LOG_NAME=SnapCenter_Linux_Host_Plugin_Install_2.log
-DCHOSEN_FEATURE_LIST=CUSTOM
```

4. Edit the /<installation directory>/NetApp/snapcenter/scc/etc/SC_SMS_Services.properties file, and then add the PLUGINS_ENABLED = hana:3.0 parameter.
5. Add the host to the SnapCenter Server using the Add-Smhost cmdlet and the required parameters.

   The information regarding the parameters that can be used with the command and their descriptions can be obtained by running *Get-Help command_name*. Alternatively, you can also refer to the SnapCenter Software Cmdlet Reference Guide.

Monitor the status of installing Plug-in for SAP HANA

You can monitor the progress of SnapCenter plug-in package installation by using the Jobs page. You might want to check the progress of installation to determine when it is complete or if there is an issue.

About this task

The following icons appear on the Jobs page and indicate the state of the operation:

- In progress
- Completed successfully
- Failed
- Completed with warnings or could not start due to warnings
- Queued

Steps

1. In the left navigation pane, click **Monitor**.
2. In the **Monitor** page, click **Jobs**.
3. In the **Jobs** page, to filter the list so that only plug-in installation operations are listed, do the following:
   a. Click **Filter**.
   b. Optional: Specify the start and end date.
   c. From the Type drop-down menu, select **Plug-in installation**.
   d. From the Status drop-down menu, select the installation status.
   e. Click **Apply**.
4. Select the installation job and click **Details** to view the job details.

5. In the **Job Details** page, click **View logs**.

Configure CA Certificate

Generate CA Certificate CSR file

You can generate a Certificate Signing Request (CSR) and import the certificate that can be obtained from a Certificate Authority (CA) using the generated CSR. The certificate will have a private key associated with it.

CSR is a block of encoded text that is given to an authorized certificate vendor to procure the signed CA certificate.

CA Certificate RSA key length must be minimum 3072 bits.

For information to generate a CSR, see How to generate CA Certificate CSR file.

If you own the CA certificate for your domain (*.domain.company.com) or your system (machine1.domain.company.com), you can skip generating the CA Certificate CSR file. You can deploy the existing CA certificate with SnapCenter.

For cluster configurations, the cluster name (virtual cluster FQDN), and the respective host names should be mentioned in the CA certificate. The certificate can be updated by filling the Subject Alternative Name (SAN) field before procuring the certificate. For a wild card certificate (*.domain.company.com), the certificate will contain all the hostnames of the domain implicitly.

Import CA certificates

You must import the CA certificates to the SnapCenter Server and the Windows host plug-ins using the Microsoft management console (MMC).

Steps

1. Go to the Microsoft management console (MMC), and then click **File** > **Add/Remove Snapin**.
2. In the Add or Remove Snap-ins window, select **Certificates** and then click **Add**.
3. In the Certificates snap-in window, select the **Computer account** option, and then click **Finish**.
4. Click **Console Root** > **Certificates – Local Computer** > **Trusted Root Certification Authorities** > **Certificates**.
5. Right-click on the folder “Trusted Root Certification Authorities”, and then select **All Tasks** > **Import** to start the import wizard.
6. Complete the wizard, as follows:

In this wizard window…	Do the following…
Import Private Key	Select the option **Yes**, import the private key, and then click **Next**.
Import File Format	Make no changes; click **Next**.

In this wizard window…	Do the following…
Security	Specify the new password to be used for the exported certificate, and then click **Next**.
Completing the Certificate Import Wizard	Review the summary, and then click **Finish** to start the import.

Importing certificate should be bundled with the private key (supported formats are: *.pfx, *.p12, and *.p7b).

7. Repeat Step 5 for the “Personal” folder.

Get the CA certificate thumbprint

A certificate thumbprint is a hexadecimal string that identifies a certificate. A thumbprint is calculated from the content of the certificate using a thumbprint algorithm.

Steps

1. Perform the following on the GUI:
   a. Double-click the certificate.
   b. In the Certificate dialog box, click the **Details** tab.
   c. Scroll through the list of fields and click **Thumbprint**.
   d. Copy the hexadecimal characters from the box.
   e. Remove the spaces between the hexadecimal numbers.

      For example, if the thumbprint is: "a9 09 50 2d d8 2a e4 14 33 e6 f8 38 86 b0 0d 42 77 a3 2a 7b", after removing the spaces, it will be: "a909502dd82ae41433e6f83886b00d4277a32a7b".

2. Perform the following from PowerShell:
   a. Run the following command to list the thumbprint of the installed certificate and identify the recently installed certificate by the subject name.

      Get-ChildItem -Path Cert:\LocalMachine\My

   b. Copy the thumbprint.

Configure CA certificate with Windows host plug-in services

You should configure the CA certificate with Windows host plug-in services to activate the installed digital certificate.

Perform the following steps on the SnapCenter Server and all the plug-in hosts where CA certificates are already deployed.

Steps

1. Remove the existing certificate binding with SMCore default port 8145, by running the following command:

```
> netsh http delete sslcert ipport=0.0.0.0:_<SMCore Port>
```

For example:

```
> netsh http delete sslcert ipport=0.0.0.0:8145
```

2. Bind the newly installed certificate with the Windows host plug-in services, by running the following commands:

```
> $cert = "_<certificate thumbprint>_"
> $guid = [guid]::NewGuid().ToString("B")
> netsh http add sslcert ipport=0.0.0.0: _<SMCore Port>_ certhash=$cert
appid="$guid"
```

For example:

```
> $cert = "a909502dd82ae41433e6f83886b00d4277a32a7b"
> $guid = [guid]::NewGuid().ToString("B")
> netsh http add sslcert ipport=0.0.0.0: _<SMCore Port>_ certhash=$cert
appid="$guid"
```

Configure the CA Certificate for the SnapCenter SAP HANA Plug-ins service on Linux host

You should manage the password of the plug-ins keystore and its certificate, configure the CA certificate, configure root or intermediate certificates to the plug-ins trust-store, and configure CA signed key pair to plug-ins trust-store with SnapCenter plug-ins service to activate the installed digital certificate.

The plug-ins uses the file 'keystore.jks', which is located at */opt/NetApp/snapcenter/scc/etc* both as its trust-store and key-store.

Manage password for plug-in keystore and alias of the CA signed key pair in use

Steps

1. You can retrieve plug-in keystore default password from plug-in agent property file.

   It is the value corresponding to the key 'KEYSTORE_PASS'.

2. Change the keystore password:

```
keytool -storepasswd -keystore keystore.jks
```

3. Change the password for all aliases of private key entries in the keystore to the same password used for the keystore:

```
keytool -keypasswd -alias "alias_name_in_cert" -keystore keystore.jks
```

Update the same for the key KEYSTORE_PASS in *agent.properties* file.

4. Restart the service after changing the password.

Password for plug-in keystore and for all the associated alias password of the private key should be same.

Configure root or intermediate certificates to plug-in trust-store

You should configure the root or intermediate certificates without the private key to plug-in trust-store.

Steps

1. Navigate to the folder containing the plug-in keystore: /opt/NetApp/snapcenter/scc/etc.
2. Locate the file 'keystore.jks'.
3. List the added certificates in the keystore:

```
keytool -list -v -keystore keystore.jks
```

4. Add a root or intermediate certificate:

```
keytool -import -trustcacerts -alias myRootCA -file
/root/USERTrustRSA_Root.cer -keystore keystore.jks
```

5. Restart the service after configuring the root or intermediate certificates to plug-in trust-store.

You should add the root CA certificate and then the intermediate CA certificates.

Configure CA signed key pair to plug-in trust-store

You should configure the CA signed key pair to the plug-in trust-store.

Steps

1. Navigate to the folder containing the plug-in keystore /opt/NetApp/snapcenter/scc/etc.
2. Locate the file 'keystore.jks'.
3. List the added certificates in the keystore:

```
keytool -list -v -keystore keystore.jks
```

4. Add the CA certificate having both private and public key.

```
keytool -importkeystore -srckeystore /root/snapcenter.ssl.test.netapp.com.pfx
-srcstoretype pkcs12 -destkeystore keystore.jks -deststoretype JKS
```

5. List the added certificates in the keystore.

```
keytool -list -v -keystore keystore.jks
```

6. Verify that the keystore contains the alias corresponding to the new CA certificate, which was added to the keystore.
7. Change the added private key password for CA certificate to the keystore password.

   Default plug-in keystore password is the value of the key KEYSTORE_PASS in agent.properties file.

   ```
   keytool -keypasswd -alias "alias_name_in_CA_cert" -keystore keystore.jks
   ```

8. If the alias name in the CA certificate is long and contains space or special characters ("*",","), change the alias name to a simple name:

   ```
   keytool -changealias -alias "long_alias_name" -destalias "simple_alias"
   -keystore keystore.jks
   ```

9. Configure the alias name from CA certificate in agent.properties file.

   Update this value against the key SCC_CERTIFICATE_ALIAS.

10. Restart the service after configuring the CA signed key pair to plug-in trust-store.

Configure certificate revocation list (CRL) for plug-ins

About this task

- SnapCenter Plug-ins will search for the CRL files in a pre-configured directory.
- Default directory for the CRL files for SnapCenter Plug-ins is ' opt/NetApp/snapcenter/scc/etc/crl'.

Steps

1. You can modify and update the default directory in agent.properties file against the key CRL_PATH.

   You can place more than one CRL file in this directory. The incoming certificates will be verified against each CRL.

Configure the CA Certificate for the SnapCenter SAP HANA Plug-ins service on Windows host

You should manage the password of the plug-ins keystore and its certificate, configure the CA certificate, configure root or intermediate certificates to the plug-ins trust-store, and configure CA signed key pair to plug-ins trust-store with SnapCenter plug-ins service to activate the installed digital certificate.

The plug-ins uses the file *keystore.jks*, which is located at *C:\Program Files\NetApp\SnapCenter\Snapcenter Plug-in Creator\etc* both as its trust-store and key-store.

Manage password for plug-in keystore and alias of the CA signed key pair in use

Steps

1. You can retrieve plug-in keystore default password from plug-in agent property file.

It is the value corresponding to the key *KEYSTORE_PASS*.

2. Change the keystore password:

   keytool -storepasswd -keystore keystore.jks

   If the "keytool" command is not recognized on the Windows command prompt, replace the keytool command with its complete path.

   C:\Program Files\Java\<jdk_version>\bin\keytool.exe" -storepasswd -keystore keystore.jks

3. Change the password for all aliases of private key entries in the keystore to the same password used for the keystore:

   keytool -keypasswd -alias "alias_name_in_cert" -keystore keystore.jks

   Update the same for the key KEYSTORE_PASS in *agent.properties* file.

4. Restart the service after changing the password.

   Password for plug-in keystore and for all the associated alias password of the private key should be same.

Configure root or intermediate certificates to plug-in trust-store

You should configure the root or intermediate certificates without the private key to plug-in trust-store.

Steps

1. Navigate to the folder containing the plug-in keystore *C:\Program Files\NetApp\SnapCenter\Snapcenter Plug-in Creator\etc*
2. Locate the file 'keystore.jks'.
3. List the added certificates in the keystore:

   keytool -list -v -keystore keystore.jks

4. Add a root or intermediate certificate:

   keytool -import -trustcacerts -alias myRootCA -file /root/USERTrustRSA_Root.cer -keystore keystore.jks

5. Restart the service after configuring the root or intermediate certificates to plug-in trust-store.

You should add the root CA certificate and then the intermediate CA certificates.

Configure CA signed key pair to plug-in trust-store

You should configure the CA signed key pair to the plug-in trust-store.

Steps

1. Navigate to the folder containing the plug-in keystore *C:\Program Files\NetApp\SnapCenter\Snapcenter Plug-in Creator\etc*
2. Locate the file *keystore.jks*.

3. List the added certificates in the keystore:

   keytool -list -v -keystore keystore.jks

4. Add the CA certificate having both private and public key.

   keytool -importkeystore -srckeystore /root/snapcenter.ssl.test.netapp.com.pfx -srcstoretype pkcs12 -destkeystore keystore.jks -deststoretype JKS

5. List the added certificates in the keystore.

   keytool -list -v -keystore keystore.jks

6. Verify that the keystore contains the alias corresponding to the new CA certificate, which was added to the keystore.
7. Change the added private key password for CA certificate to the keystore password.

   Default plug-in keystore password is the value of the key KEYSTORE_PASS in agent.properties file.

   keytool -keypasswd -alias "alias_name_in_CA_cert" -keystore keystore.jks

8. Configure the alias name from CA certificate in *agent.properties* file.

   Update this value against the key SCC_CERTIFICATE_ALIAS.

9. Restart the service after configuring the CA signed key pair to plug-in trust-store.

Configure certificate revocation list (CRL) for SnapCenter plug-ins

About this task

- To download the latest CRL file for the related CA certificate see How to update certificate revocation list file in SnapCenter CA Certificate.
- SnapCenter plug-ins will search for the CRL files in a pre-configured directory.
- Default directory for the CRL files for SnapCenter plug-ins is *'C:\Program Files\NetApp\SnapCenter\Snapcenter Plug-in Creator\ etc\crl'*.

Steps

1. You can modify and update the default directory in *agent.properties* file against the key CRL_PATH.
2. You can place more than one CRL file in this directory.

   The incoming certificates will be verified against each CRL.

Enable CA Certificates for plug-ins

You should configure the CA certificates and deploy the CA certificates in the SnapCenter Server and the corresponding plug-in hosts. You should enable the CA certificate validation for the plug-ins.

Before you begin

- You can enable or disable the CA certificates using the run *Set-SmCertificateSettings* cmdlet.
- You can display the certificate status for the plug-ins using the *Get-SmCertificateSettings*.

The information regarding the parameters that can be used with the cmdlet and their descriptions can be obtained by running *Get-Help command_name*. Alternatively, you can also refer to the SnapCenter Software Cmdlet Reference Guide.

Steps

1. In the left navigation pane, click **Hosts**.
2. In the Hosts page, click **Managed Hosts**.
3. Select single or multiple plug-in hosts.
4. Click **More options**.
5. Select **Enable Certificate Validation**.

After you finish

The Managed Hosts tab host displays a padlock and the color of the padlock indicates the status of the connection between SnapCenter Server and the plug-in host.

- indicates that the CA certificate is neither enabled nor assigned to the plug-in host.
- indicates that the CA certificate is successfully validated.
- indicates that the CA certificate could not be validated.
- indicates that the connection information could not be retrieved.

When the status is yellow or green, the data protection operations completes successfully.

Install SnapCenter Plug-in for VMware vSphere

If your database or filesystem is stored on virtual machines (VMs), or if you want to protect VMs and datastores, you must deploy the SnapCenter Plug-in for VMware vSphere virtual appliance.

For information to deploy, see Deployment Overview.

Deploy CA certificate

To configure the CA Certificate with SnapCenter Plug-in for VMware vSphere, see Create or import SSL certificate.

Configure the CRL file

SnapCenter Plug-in for VMware vSphere looks for the CRL files in a pre-configured directory. Default directory of the CRL files for SnapCenter Plug-in for VMware vSphere is */opt/netapp/config/crl*.

You can place more than one CRL file in this directory. The incoming certificates will be verified against each CRL.

Prepare for data protection

Prerequisites for using the SnapCenter Plug-in for SAP HANA Database

Before you use SnapCenter Plug-in for SAP HANA Database, the SnapCenter administrator must install and configure the SnapCenter Server and perform the prerequisite tasks.

- Install and configure SnapCenter Server.
- Log in to SnapCenter Server.
- Configure the SnapCenter environment by adding storage system connections and creating credentials, if applicable.
- Install Java 11 on your Linux or Windows host.

  You must set the Java path in the environmental path variable of the host machine.

- Set up SnapMirror and SnapVault, if you want backup replication.
- Install the HDBSQL client on the host where you will install the Plug-in for SAP HANA Database.

  Configure the user store keys for the SAP HANA nodes that you will manage through this host.

- For SAP HANA database 2.0SPS05, if you are using a SAP HANA database user account, ensure that you have the following permissions to perform backup, restore, and clone operations in SnapCenter Server:
  - Backup admin
  - Catalog read
  - Database backup admin
  - Database recovery operator

How resources, resource groups, and policies are used for protecting SAP HANA databases

Before you use SnapCenter, it is helpful to understand basic concepts related to the backup, clone, and restore operations you want to perform. You interact with resources, resource groups, and policies for different operations.

- Resources are typically SAP HANA databases that you back up or clone with SnapCenter.
- A SnapCenter resource group, is a collection of resources on a host.

  When you perform an operation on a resource group, you perform that operation on the resources defined in the resource group according to the schedule you specify for the resource group.

  You can back up on demand a single resource or a resource group. You also can perform scheduled backups for single resources and resource groups.

- The policies specify the backup frequency, replication, scripts, and other characteristics of data protection operations.

  When you create a resource group, you select one or more policies for that group. You can also select a policy when you perform a backup on demand for a single resource.

Think of a resource group as defining what you want to protect and when you want to protect it in terms of day

and time. Think of a policy as defining how you want to protect it. If you are backing up all databases, for example, you might create a resource group that includes all of the databases in the host. You could then attach two policies to the resource group: a daily policy and an hourly policy. When you create the resource group and attach the policies, you might configure the resource group to perform a full backup daily.

Back up SAP HANA resources

Back up SAP HANA resources

You can either create a backup of a resource (database) or resource group. The backup workflow includes planning, identifying the databases for backup, managing backup policies, creating resource groups and attaching policies, creating backups, and monitoring the operations.

The following workflow shows the sequence in which you must perform the backup operation:

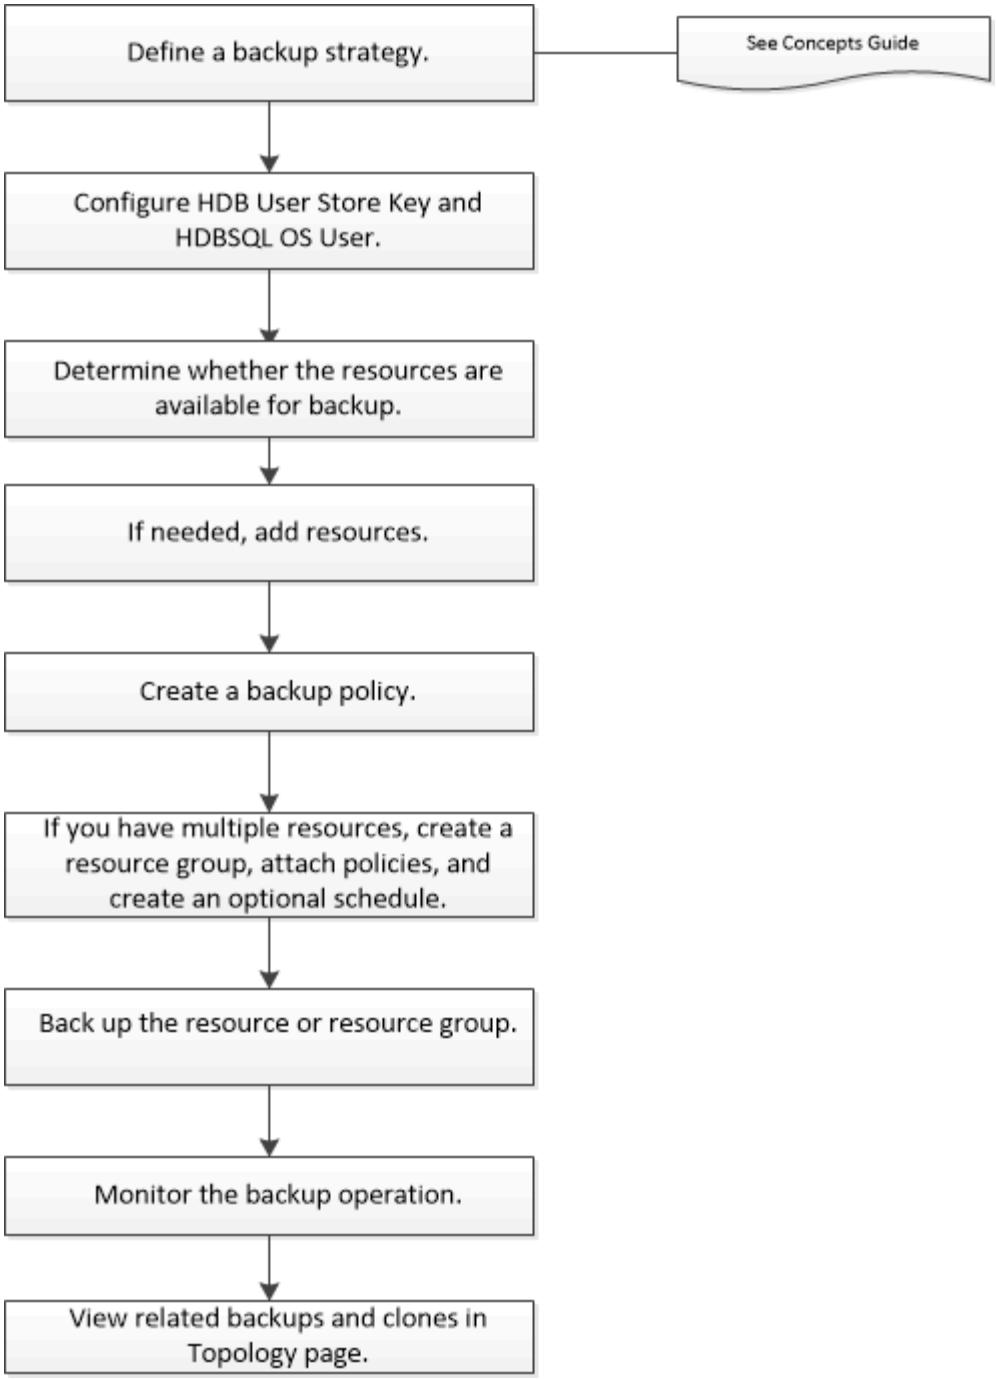

You can also use PowerShell cmdlets manually or in scripts to perform backup, restore, and clone operations. The SnapCenter cmdlet help and the cmdlet reference information contain more information about PowerShell cmdlets.
SnapCenter Software Cmdlet Reference Guide.

Configure HDB User Store Key and HDBSQL OS User for the SAP HANA database

You must configure HDB User Store Key and HDBSQL OS User to perform data protection operations on SAP HANA databases.

Before you begin

- If the SAP HANA database does not have the HDB Secure User Store Key and HDB SQL OS User configured, a red padlock icon appears only for the autodiscovered resources. If during a subsequent

discovery operation, the configured HDB Secure User Store Key was found to be incorrect or did not provide access to the database itself, then the red padlock icon will reappear.

- You must configure the HDB Secure User Store Key and the HDB SQL OS user to be able to protect the database or add it to a resource group to perform data protection operations.
- You must configure HDB SQL OS User to access the system database. If HDB SQL OS User is configured to access only tenant database, the discovery operation will fail.

Steps

1. In the left navigation pane, click **Resources** and then select SnapCenter Plug-in for SAP HANA Database from the list.
2. In the Resources page, select the resource type from the **View** list.
3. (Optional) Click and select the host name.

   You can then click to close the filter pane.

4. Select the database, and then click **Configure Database**.
5. In the Configure database settings section, enter HDB Secure User Store Key.

   The Plug-in host name is displayed and HDB SQL OS User is automatically populated to <sid>adm.

6. Click **OK**.

You can modify the database configuration from the Topology page.

Discover resources and prepare multitenant database containers for data protection

Discover the databases automatically

Resources are SAP HANA databases and Non-data Volume on the Linux host that are managed by SnapCenter. You can add these resources to resource groups to perform data protection operations after you discover the SAP HANA databases that are available.

Before you begin

- You must have already completed tasks such as installing the SnapCenter Server, adding HDB User Store Key, adding hosts, and setting up the storage system connections.
- You must have configured the HDB Secure User Store Key and HDB SQL OS user on the Linux host.
  - You must configure the HDB User Store Key with SID adm user. For example, for HANA system with A22 as the SID, the HDB User Store Key must be configured with a22adm.
- SnapCenter Plug-in for SAP HANA Database does not support automatic discovery of the resources residing on RDM/VMDK virtual environments. You must provide the storage information for virtual environments while adding the databases manually.

About this task

After installing the plug-in, all the resources on that Linux host are automatically discovered and displayed on the Resources page.

The automatically discovered resources cannot be modified or deleted.

Steps

1. In the left navigation pane, click **Resources**, and then select the Plug-in for SAP HANA Database from the list.
2. In the Resources page select the resource type from the View list.
3. (Optional) Click , and then select the host name.

   You can then click to close the filter pane.

4. Click **Refresh Resources** to discover the resources available on the host.

   The resources are displayed along with information such as resource type, host name, associated resource groups, backup type, policies and overall status.

   - If the database is on a NetApp storage and not protected, then Not protected is displayed in the Overall Status column.
   - If the database is on a NetApp storage system and protected, and if there is no backup operation performed, then Backup not run is displayed in the Overall Status column. The status will otherwise change to Backup failed or Backup succeeded based on the last backup status.

If the SAP HANA database does not have a HDB Secure User Store Key configured, a red padlock icon appears next to the resource. If during a subsequent discovery operation, the configured HDB Secure User Store Key was found to be incorrect or did not provide access to the database itself, then the red padlock icon will reappear.

You must refresh the resources if the databases are renamed outside of SnapCenter.

After you finish

You must configure the HDB Secure User Store Key and HDBSQL OS User to be able to protect the database or add it to the resource group to perform data protection operations.

Configure HDB User Store Key and HDBSQL OS User for the SAP HANA database

Prepare multitenant database containers for data protection

For SAP HANA hosts directly registered in SnapCenter, installing or upgrading the SnapCenter Plug-in for SAP HANA Database will trigger an automatic discovery for resources on the host. After installing or upgrading the plug-in, for every multitenant database containers (MDC) resource that was located on the plug-in host, another MDC resource will be automatically discovered with a different GUID format and registered in SnapCenter. The new resource will be in “locked” state.

About this task

For example, in SnapCenter 4.2, if E90 MDC resource was located on the plug-in host and registered manually, after upgrading to SnapCenter 4.3, another E90 MDC resource with a different GUID will be discovered and registered in SnapCenter.

The backups associated with the resource of SnapCenter 4.2 and earlier versions must be retained until the expiry of the retention period. After the retention period expires, you can delete the old MDC resource and continue to manage the new auto discovered MDC resource.

`Old MDC resource` is the MDC resource for a plug-in host that was manually added in SnapCenter 4.2 or earlier releases.

Perform the following steps to start using the new resource discovered in SnapCenter 4.3 for data protection operations:

Steps

1. In the Resources page, select the old MDC resource with backups added to the earlier SnapCenter release, and place it in “maintenance mode” from the Topology page.

   If the resource is part of a resource group, place the resource group in “maintenance mode”.

2. Configure the new MDC resource discovered after upgrading to SnapCenter 4.3 by selecting the new resource from the Resources page.

   “New MDC resource” is the newly discovered MDC resource that was discovered once the SnapCenter Server and the plug-in host was upgraded to 4.3. The new MDC resource can be identified as a resource with the same SID as the old MDC resource, for a given host, and with a red padlock icon next to it in the Resources page.

3. Protect the new MDC resource discovered after upgrading to SnapCenter 4.3 by selecting protection policies, schedules, and notification settings.
4. Delete the backups taken in SnapCenter 4.2 or earlier releases based on the retention settings.
5. Delete the resource group from the Topology page.
6. Delete the old MDC resource from the Resources page.

   For example, if the primary Snapshots retention period is 7 days and secondary Snapshots retention is 45 days, after 45 days are complete and after all the backups are deleted, you must delete the resource group and the old MDC resource.

Related information

Configure HDB User Store Key and HDBSQL OS User for the SAP HANA database

View SAP HANA database backups and clones in the Topology page

Add resources manually to the plug-in host

Automatic discovery is not supported for certain HANA instances. You must add these resources manually.

Before you begin

- You must have completed tasks such as installing the SnapCenter Server, adding hosts, setting up storage system connections, and adding HDB User Store Key.
- For SAP HANA system replication, it is recommended to add all the resources of that HANA system into one resource group and take a resource group backup. This ensures a seamless backup during takeover-failback mode.

Create resource groups and attach policies.

About this task

Automatic discovery is not supported for the following configurations:

- RDM and VMDK layouts

> In case the above resources are discovered, the data protection operations are not supported on these resources.

- HANA multiple-host configuration
- Multiple instances on the same host
- Multitier scale out HANA System Replication
- Cascaded replication environment in System Replication mode

Steps

1. In the left navigation pane, select the SnapCenter Plug-in for SAP HANA Database from the drop-down list, and then click **Resources**.
2. In the Resources page, click **Add SAP HANA Database**.
3. In the Provide Resource Details page, perform the following actions:

For this field…	Do this…
Resource Type	Enter the resource type. Resource types are Single Container, Multitenant Database Container (MDC), and Non-data Volume.
HANA System Name	Enter the descriptive SAP HANA system name. This option is available only if you selected Single Container or MDC resource types.
SID	Enter the system ID (SID). The installed SAP HANA system is identified by a single SID.
Plug-in Host	Select the plug-in host.
HDB Secure User Store Keys	Enter the key to connect to the SAP HANA system. The key contains the login information to connect to the database. For SAP HANA System Replication, secondary user key is not validated. This will be used during takeover.

For this field…	Do this…
HDBSQL OS User	Enter the user name for whom the HDB Secure User Store Key is configured. For Windows, it is mandatory for the HDBSQL OS User to be the SYSTEM user. Therefore, you must configure the HDB Secure User Store Key for the SYSTEM user.

4. In the Provide Storage Footprint page, select a storage system and choose one or more volumes, LUNs, and qtrees, and then click **Save**.

   Optional: You can click the + icon to add more volumes, LUNs, and qtrees from other storage systems.

5. Review the summary, and then click **Finish**.

   The databases are displayed along with information such as the SID, plug-in host, associated resource groups and policies, and overall status

If you want to provide users access to resources, you must assign the resources to the users. This enables users to perform the actions for which they have permissions on the assets that are assigned to them.

Add a user or group and assign role and assets

After adding the databases, you can modify the SAP HANA database details.

You cannot modify the following if there are backups associated with the SAP HANA resource:

- Multitenant database containers (MDC): SID, or HDBSQL Client (plug-in) Host
- Single Container: SID or HDBSQL Client (plug-in) Host
- Non-data Volume: Resource name, Associated SID, or Plug-in Host

Create backup policies for SAP HANA databases

Before you use SnapCenter to back up SAP HANA database resources, you must create a backup policy for the resource or resource group that you want to back up. A backup policy is a set of rules that governs how you manage, schedule, and retain backups.

Before you begin

- You must have defined your backup strategy.

  For details, see the information about defining a data protection strategy for SAP HANA databases.

- You must have prepared for data protection by completing tasks such as installing SnapCenter, adding hosts, setting up storage system connections, and adding resources.
- The SnapCenter administrator must have assigned the SVMs for both the source and destination volumes to you if you are replicating Snapshots to a mirror or vault.

  Additionally, you can specify replication, script, and application settings in the policy. These options saves time when you want to reuse the policy for another resource group.

- Review the SnapMirror active sync specific prerequisites and limitations. For information refer Object limits for SnapMirror active sync.

About this task

- SAP HANA System Replication
  - You can protect the primary SAP HANA system and all the data protection operations can be performed.
  - You can protect the secondary SAP HANA system, but the backups cannot be created.

    After the failover, all the data protection operation can be performed as the secondary SAP HANA system becomes the primary SAP HANA system.

    You cannot create a backup for SAP HANA data volume, but SnapCenter continues to protect the Non-data Volumes (NDV).
- SnapLock
  - If 'Retain the backup copies for a specific number of days' option is selected, then the SnapLock retention period must be lesser than or equal to the mentioned retention days.
  - Specifying a Snapshot locking period prevents deletion of the Snapshots until the retention period expires. This could lead to retaining a larger number of Snapshots than the count specified in the policy.
  - For ONTAP 9.12.1 and below version, the clones created from the SnapLock Vault Snapshots as part of restore will inherit the SnapLock Vault expiry time. Storage admin should manually cleanup the clones post the SnapLock expiry time.

Steps

1. In the left navigation pane, click **Settings**.
2. In the Settings page, click **Policies**.
3. Click **New**.
4. In the Name page, enter the policy name and details.
5. In the Policy type page, perform the following steps:
   - Choose storage type
   - Choose backup type:

If you want to…	Do this…
Create a backup using Snapshot technology	Select **Snapshot Based**.
Perform an integrity check of the database	Select **File-Based Backup**. Only active tenants are backed up.

6. In the Snapshot and Replication page, perform the following steps:
   - Specify the schedule type by selecting **On demand**, **Hourly**, **Daily**, **Weekly**, or **Monthly**.

You can specify the schedule (start date, end date, and frequency) for the backup operation while creating a resource group. This enables you to create resource groups that share the same policy and backup frequency, but also enables you to assign different backup schedules to each policy.

If you have scheduled for 2:00 a.m., the schedule will not be triggered during daylight saving time (DST).

7. In the Snapshot and Replication page, specify the retention settings for the backup type and the schedule type selected in the Backup Type page:

If you want to…	Then…
Keep a certain number of Snapshots	Select **Copies to keep**, and then specify the number of Snapshots that you want to keep. If the number of Snapshots exceeds the specified number, the oldest Snapshots are deleted first. The maximum retention value is 1018. Backups will fail if retention is set to a value higher than what the ONTAP version supports. For Snapshot copy-based backups, you must set the retention count to 2 or higher if you plan to enable SnapVault replication. If you set the retention count to 1, the retention operation might fail because the first Snapshot is the reference Snapshot for the SnapVault relationship until a newer Snapshot is replicated to the target. For SAP HANA system replication, it is recommended to add all the resources of the SAP HANA system into one resource group. This ensures that the right number of backups are retained. For SAP HANA System Replication, the total Snapshots taken will be equal to the retention set for the resource group. The removal of the oldest Snapshot is based on which node the oldest Snapshot is located. For example, the retention is set to 7 for a resource group with SAP HANA System Replication primary and SAP HANA System Replication secondary. You can take a maximum of 7 Snapshots at a time including both SAP HANA System Replication primary and SAP HANA System Replication secondary.
Keep the Snapshots for a certain number of days	Select **Retain copies for**, and then specify the number of days for which you want to keep the Snapshots before deleting them.

If you want to…	Then…
Snapshot copy locking period	Select **Snapshot copy locking period**, and specify days, months, or years. SnapLock retention period should be less than 100 years.

8. Select a Snapshot label.

> You can assign SnapMirror labels to primary snapshots for remote replication, allowing the primary snapshots to offload the snapshot replication operation from SnapCenter to ONTAP secondary systems. This can be done without enabling SnapMirror or SnapVault option in the policy page.

9. For Snapshot copy-based backups, in the Select secondary replication options section, select one or both of the following secondary replication options:

For this field…	Do this…
Update SnapMirror after creating a local Snapshot copy	Select this field to create mirror copies of the backup sets on another volume (SnapMirror replication). This option should be enabled for SnapMirror active sync. If the protection relationship in ONTAP is of type Mirror and Vault and if you select only this option, the Snapshot created on the primary will not be transferred to the destination, but will be listed in the destination. If this Snapshot is selected from the destination to perform a restore operation, then the Secondary Location is not available for the selected vaulted/mirrored backup error message is displayed. During secondary replication, the SnapLock expiry time loads the primary SnapLock expiry time. Clicking the **Refresh** button in the Topology page refreshes the secondary and primary SnapLock expiry time that are retrieved from ONTAP. See View SAP HANA database backups and clones in the Topology page.

For this field…	Do this…
Update SnapVault after creating a local Snapshot copy	Select this option to perform disk-to-disk backup replication (SnapVault backups). During secondary replication, the SnapLock expiry time loads the primary SnapLock expiry time. Clicking the **Refresh** button in the Topology page refreshes the secondary and primary SnapLock expiry time that are retrieved from ONTAP. When SnapLock is configured only on the secondary from ONTAP known as SnapLock Vault, clicking the **Refresh** button in the Topology page refreshes the locking period on the secondary that is retrieved from ONTAP. For more information on SnapLock Vault see Commit Snapshot copies to WORM on a vault destination See View SAP HANA database backups and clones in the Topology page.
Error retry count	Enter the maximum number of replication attempts that can be allowed before the operation stops.

You should configure SnapMirror retention policy in ONTAP for the secondary storage to avoid reaching the maximum limit of Snapshots on the secondary storage.

10. Review the summary, and then click **Finish**.

Create resource groups and attach policies

A resource group is the container to which you must add resources that you want to back up and protect. A resource group enables you to back up all the data that is associated with a given application simultaneously. A resource group is required for any data protection job. You must also attach one or more policies to the resource group to define the type of data protection job that you want to perform.

About this task

- To create SAP HANA system replication backups, it is recommended to add all the resources of the SAP HANA system into one resource group. This ensures a seamless backup during takeover-failback mode.
- For ONTAP 9.12.1 and below version, the clones created from the SnapLock Vault Snapshots as part of restore will inherit the SnapLock Vault expiry time. Storage admin should manually cleanup the clones post the SnapLock expiry time.
- Adding new databases without SnapMirror active sync to an existing resource group which contains resources with SnapMirror active sync, is not supported.
- Adding new databases to an existing resource group in failover mode of SnapMirror active sync is not supported. You can add resources to the resource group only in regular or fail-back state.

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page, click **New Resource Group**.
3. In the Name page, perform the following actions:

For this field…	Do this…
Name	Enter a name for the resource group. The resource group name should not exceed 250 characters.
Tags	Enter one or more labels that will help you later search for the resource group. For example, if you add HR as a tag to multiple resource groups, you can later find all resource groups associated with the HR tag.
Use custom name format for Snapshot copy	Select this check box, and enter a custom name format that you want to use for the Snapshot name. For example, customtext_resource group_policy_hostname or resource group_hostname. By default, a timestamp is appended to the Snapshot name.

4. In the Resources page, select a host name from the **Host** drop-down list and resource type from the **Resource Type** drop-down list.

   This helps to filter information on the screen.

5. Select the resources from the **Available Resources** section, and then click the right arrow to move them to the **Selected Resources** section.
6. In the Application Settings page, do the following:
   a. Click the **Backups** arrow to set additional backup options:

      Enable consistency group backup and perform the following tasks:

For this field…	Do this…
Afford time to wait for Consistency Group Snapshot operation to complete	Select **Urgent**, **Medium**, or **Relaxed** to specify the wait time for Snapshot operation to complete. Urgent = 5 seconds, Medium = 7 seconds, and Relaxed = 20 seconds.
Disable WAFL Sync	Select this to avoid forcing a WAFL consistency point.

b. Click the **Scripts** arrow and enter the pre and post commands for quiesce, Snapshot, and unquiesce operations. You can also enter the pre commands to be executed before exiting in the event of a failure.

c. Click the **Custom Configurations** arrow and enter the custom key-value pairs required for all data protection operations using this resource.

Parameter	Setting	Description
ARCHIVE_LOG_ENABLE	(Y/N)	Enables the archive log management to delete the archive logs.
ARCHIVE_LOG_RETENTION	number_of_days	Specifies the number of days the archive logs are retained. This setting must be equal to or greater than NTAP_SNAPSHOT_ RETENTIONS.
ARCHIVE_LOG_DIR	change_info_directory/logs	Specifies the path to the directory that contains the archive logs.

Parameter	Setting	Description
ARCHIVE_LOG_EXT	file_extension	Specifies the archive log file extension length. For example, if the archive log is log_backup_0_0_0_0.161518551942 9 and if the file_extension value is 5, then the extension of the log will retain 5 digits, which is 16151.
ARCHIVE_LOG_RECURSIVE_SE ARCH	(Y/N)	Enables the management of archive logs within subdirectories. You should use this parameter if the archive logs are located under subdirectories.

The custom key-value pairs are supported for SAP HANA Linux plug-in systems and not supported for SAP HANA database registered as a centralized windows plug-in.

d. Click the **Snapshot Copy Tool** arrow to select the tool to create Snapshots:

If you want…	Then…
SnapCenter to use the plug-in for Windows and put the file system into a consistent state before creating a Snapshot. For Linux resources, this option is not applicable.	Select **SnapCenter with File System Consistency**. This option is not applicable for SnapCenter Plug-in for SAP HANA Database.
SnapCenter to create a storage level Snapshot	Select **SnapCenter without File System Consistency**.
To enter the command to be executed on the host to create Snapshot copies.	Select **Other**, and then enter the command to be executed on the host to create a Snapshot.

7. In the Policies page, perform the following steps:
   a. Select one or more policies from the drop-down list.

You can also create a policy by clicking + .

The policies are listed in the Configure schedules for selected policies section.

b. In the Configure Schedules column, click + for the policy you want to configure.

c. In the Add schedules for policy *policy_name* dialog box, configure the schedule, and then click **OK**.

   Where, policy_name is the name of the policy that you have selected.

   The configured schedules are listed in the **Applied Schedules** column.

   Third party backup schedules are not supported when they overlap with SnapCenter backup schedules.

8. In the Notification page, from the **Email preference** drop-down list, select the scenarios in which you want to send the emails.

   You must also specify the sender and receiver email addresses, and the subject of the email. The SMTP server must be configured in **Settings** > **Global Settings**.

9. Review the summary, and then click **Finish**.

Create resource groups and enable secondary protection for SAP HANA resources on ASA r2 systems

You should create the resource group to add the resources that are on ASA r2 systems. You can also provision the secondary protection while creating the resource group.

Before you begin

- You should ensure that you are not adding both ONTAP 9.x resources and ASA r2 resources to the same resource group.
- You should ensure that you do not have a database with both ONTAP 9.x resources and ASA r2 resources.

About this task

- The secondary protection is available only if the logged-in user is assigned to the role that has the **SecondaryProtection** capability enabled.
- If you enabled secondary protection, the resource group is put into maintenance mode while creating the primary and secondary consistency groups. After the primary and secondary consistency groups are created, the resource group is put out of maintenance mode.
- SnapCenter does not support secondary protection for a clone resource.

Steps

1. In the left navigation pane, select **Resources**, and the appropriate plug-in from the list.
2. In the Resources page, click **New Resource Group**.
3. In the Name page, perform the following actions:
   a. Enter a name for the resource group in the Name field.

      The resource group name should not exceed 250 characters.

   b. Enter one or more labels in the Tag field to help you search for the resource group later.

      For example, if you add HR as a tag to multiple resource groups, you can later find all resource groups associated with the HR tag.

c. Select this check box, and enter a custom name format that you want to use for the Snapshot name.

   For example, customtext_resource group_policy_hostname or resource group_hostname. By default, a timestamp is appended to the Snapshot name.

d. Specify the destinations of the archive log files that you do not want to back up.

   You should use the exact same destination as it was set in the application including prefix, if needed.

4. In the Resources page, select the database host name from the **Host** drop-down list.

   The resources are listed in the Available Resources section only if the resource is discovered successfully. If you have recently added resources, they will appear on the list of available resources only after you refresh your resource list.

5. Select the ASA r2 resources from the Available Resources section and move them to the Selected Resources section.
6. In the Application Settings page, select the backup option.
7. In the Policies page, perform the following steps:

   a. Select one or more policies from the drop-down list.

      You can also create a policy by clicking + .

      In the Configure schedules for selected policies section, the selected policies are listed.

   b. Click + in the Configure Schedules column for the policy for which you want to configure a schedule.

   c. In the Add schedules for policy *policy_name* window, configure the schedule, and then click **OK**.

      Where, *policy_name* is the name of the policy that you have selected.

      The configured schedules are listed in the Applied Schedules column.

   Third party backup schedules are not supported when they overlap with SnapCenter backup schedules.

8. If the secondary protection is enabled for the policy that you have selected, then Secondary Protection page is displayed and you need to perform the following steps:

   a. Select the type of the replication policy.

      Synchronous replication policy is not supported.

   b. Specify the consistency group suffix that you want to use.

   c. From the Destination Cluster and Destination SVM drop-downs select the peered cluster and SVM that you want to use.

The cluster and SVM peering is not supported by SnapCenter. You should use System Manager or ONTAP CLIs to perform cluster and SVM peering.

If the resources are already protected outside of SnapCenter, those resources will be displayed in the Secondary Protected Resources section.

1. On the Verification page, perform the following steps:
   a. Click **Load locators** to load the SnapMirror or SnapVault volumes to perform verification on secondary storage.
   b. Click + in the Configure Schedules column to configure the verification schedule for all the schedule types of the policy.
   c. In the Add Verification Schedules policy_name dialog box, perform the following actions:

If you want to…	Do this…
Run verification after backup	Select **Run verification after backup**.
Schedule a verification	Select **Run scheduled verification** and then select the schedule type from the drop-down list.

   d. Select **Verify on secondary location** to verify your backups on secondary storage system.
   e. Click **OK**.

   The configured verification schedules are listed in the Applied Schedules column.

2. In the Notification page, from the **Email preference** drop-down list, select the scenarios in which you want to send the emails.

   You must also specify the sender and receiver email addresses, and the subject of the email. If you want to attach the report of the operation performed on the resource group, select **Attach Job Report**.

   For email notification, you must have specified the SMTP server details using the either the GUI or the PowerShell command Set-SmSmtpServer.

3. Review the summary, and then click **Finish**.

Create a storage system connection and a credential using PowerShell cmdlets for SAP HANA database

You must create a storage virtual machine (SVM) connection and a credential before using PowerShell cmdlets to back up, restore, or clone SAP HANA databases.

Before you begin

- You should have prepared the PowerShell environment to execute the PowerShell cmdlets.
- You should have the required permissions in the Infrastructure Admin role to create storage connections.
- You should ensure that the plug-in installations are not in progress.

Host plug-in installations must not be in progress while adding a storage system connection because the host cache might not be updated and databases status might be displayed in the SnapCenter GUI as “Not available for backup” or “Not on NetApp storage”.

- Storage system names should be unique.

  SnapCenter does not support multiple storage systems with the same name on different clusters. Each storage system that is supported by SnapCenter should have a unique name and a unique data LIF IP address.

Steps

1. Initiate a PowerShell connection session by using the Open-SmConnection cmdlet.

```
PS C:\> Open-SmStorageConnection
```

2. Create a new connection to the storage system by using the Add-SmStorageConnection cmdlet.

```
PS C:\> Add-SmStorageConnection -StorageType DataOntap -Type DataOntap
-OntapStorage 'scsnfssvm' -Protocol Https -Timeout 60
```

3. Create a new credential by using the Add-SmCredential cmdlet.

   This example shows how to create a new credential named FinanceAdmin with Windows credentials:

```
PS C:> Add-SmCredential -Name FinanceAdmin -AuthMode Windows
-Credential sddev\administrator
```

4. Add the SAP HANA communication host to SnapCenter Server.

```
PS C:> Add-SmHost -HostName 10.232.204.61 -OSType Windows -RunAsName
FinanceAdmin -PluginCode hana
```

5. Install the package and the SnapCenter Plug-in for SAP HANA Database on the host.

   For Linux:

```
PS C:> Install-SmHostPackage -HostNames 10.232.204.61 -ApplicationCode
hana
```

   For Windows:

```
Install-SmHostPackage -HostNames 10.232.204.61 -ApplicationCode hana
-FilesystemCode scw -RunAsName FinanceAdmin
```

6. Set the path to the HDBSQL client.

   For Windows:

```
PS C:> Set-SmConfigSettings -Plugin -HostName 10.232.204.61 -PluginCode
hana -configSettings @{"HANA_HDBSQL_CMD" = "C:\Program
Files\sap\hdbclient\hdbsql.exe"}
```

   For Linux:

```
Set-SmConfigSettings -Plugin -HostName scs-hana.gdl.englab.netapp.com
-PluginCode hana -configSettings
@{"HANA_HDBSQL_CMD"="/usr/sap/hdbclient/hdbsql"}
```

The information regarding the parameters that can be used with the cmdlet and their descriptions can be obtained by running *Get-Help command_name*. Alternatively, you can also refer to the SnapCenter Software Cmdlet Reference Guide.

Back up SAP HANA databases

If a resource is not yet part of any resource group, you can back up the resource from the Resources page.

Before you begin

- You must have created a backup policy.
- If you want to back up a resource that has a SnapMirror relationship with a secondary storage, the ONTAP role assigned to the storage user should include the “snapmirror all” privilege. However, if you are using the “vsadmin” role, then the “snapmirror all” privilege is not required.
- For Snapshot copy-based backup operation, ensure that all the tenant databases are valid and active.
- To create SAP HANA system replication backups, it is recommended to add all the resources of the SAP HANA system into one resource group. This ensures a seamless backup during takeover-failback mode.

  Create resource groups and attach policies.

  Back up resource groups

- If you want to create a file-based backup when one or more tenant databases are down, set the ALLOW_FILE_BASED_BACKUP_IFINACTIVE_TENANTS_PRESENT parameter to **YES** in the HANA properties file using `Set-SmConfigSettings` cmdlet.

  The information regarding the parameters that can be used with the cmdlet and their descriptions can be obtained by running Get-Help *command_name*. Alternatively, you can also refer to the SnapCenter Software Cmdlet Reference Guide

- For pre and post commands for quiesce, Snapshot, and unquiesce operations, you should check if the commands exist in the command list available on the plug-in host from the following paths:
  - Default location on the Windows host: *C:\Program Files\NetApp\SnapCenter\Snapcenter Plug-in*

Creator\etc\allowed_commands.config

- Default location on the Linux host: */opt/NetApp/snapcenter/scc/etc/allowed_commands.config*

If the commands do not exist in the command list, then the operation will fail.

SnapCenter UI

Steps

1. In the left navigation pane, select **Resources**, and then select the appropriate plug-in from the list.
2. In the Resource page, filter resources from the **View** drop-down list based on resource type.

   Select , and then select the host name and the resource type to filter the resources. You can then select to close the filter pane.

3. Select the resource that you want to back up.
4. In the Resource page, select **Use custom name format for Snapshot copy**, and then enter a custom name format that you want to use for the Snapshot name.

   For example, *customtext_policy_hostname* or *resource_hostname*. By default, a timestamp is appended to the Snapshot name.

5. In the Application Settings page, do the following:
   - Select the **Backups** arrow to set additional backup options:

     Enable consistency group backup, if needed, and perform the following tasks:

For this field…	Do this…
Afford time to wait for "Consistency Group Snapshot" operation to complete	Select **Urgent**, or **Medium**, or **Relaxed** to specify the wait time for Snapshot operation to finish. Urgent = 5 seconds, Medium = 7 seconds, and Relaxed = 20 seconds.
Disable WAFL Sync	Select this to avoid forcing a WAFL consistency point.

   - Select the **Scripts** arrow to run pre and post commands for quiesce, Snapshot, and unquiesce operations.

     You can also run pre commands before exiting the backup operation. Prescripts and postscripts are run in the SnapCenter Server.

   - Select the **Custom Configurations**arrow, and then enter the custom value pairs required for all jobs using this resource.
   - Select the **Snapshot Copy Tool** arrow to select the tool to create Snapshots:

If you want…	Then…
SnapCenter to create a storage-level Snapshot	Select **SnapCenter without File System Consistency**.
SnapCenter to use the plug-in for Windows to put the file system into a consistent state and then create a Snapshot	Select **SnapCenter with File System Consistency**.

If you want…	Then…
To enter the command to create a Snapshot	Select **Other**, and then enter the command to create a Snapshot.

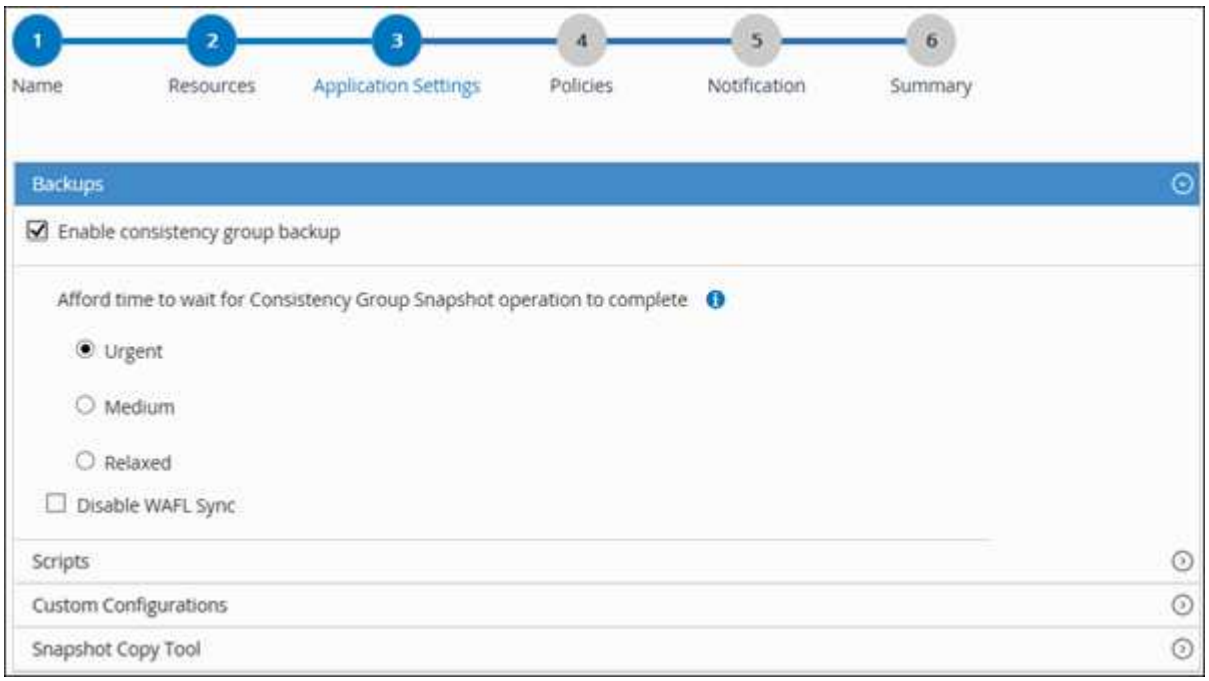

6. In the Policies page, perform the following steps:
   a. Select one or more policies from the drop-down list.

      You can also create a policy by clicking + .

      In the Configure schedules for selected policies section, the selected policies are listed.

   b. Select + in the Configure Schedules column for the policy for which you want to configure a schedule.
   c. In the Add schedules for policy *policy_name* dialog box, configure the schedule, and then select **OK**.

      policy_name is the name of the policy that you selected.

      The configured schedules are listed in the Applied Schedules column.

7. In the Notification page, from the **Email preference** drop-down list, select the scenarios in which you want to send the emails.

   You must also specify the sender and receiver email addresses, and the subject of the email. SMTP must also be configured in **Settings** > **Global Settings**.

8. Review the summary, and then select **Finish**.

   The resources topology page is displayed.

9. Select **Back up Now**.
10. In the Backup page, perform the following steps:
    a. If you applied multiple policies to the resource, from the **Policy** drop-down list, select the policy that you want to use for backup.

       If the policy selected for the on-demand backup is associated with a backup schedule, the on-demand backups will be retained based on the retention settings specified for the schedule type.

    b. Select **Backup**.
11. Monitor the operation progress by clicking **Monitor** > **Jobs**.
    - In MetroCluster configurations, SnapCenter might not be able to detect a protection relationship after a failover.

      For information, see: Unable to detect SnapMirror or SnapVault relationship after MetroCluster failover

    - If you are backing up application data on VMDKs and the Java heap size for the SnapCenter Plug-in for VMware vSphere is not large enough, the backup might fail.

      To increase the Java heap size, locate the script file */opt/netapp/init_scripts/scvservice*. In that script, the *do_start method* command starts the SnapCenter VMware plug-in service. Update that command to the following: *Java -jar -Xmx8192M -Xms4096M*

PowerShell cmdlets

Steps

1. Initiate a connection session with the SnapCenter Server for a specified user by using the Open-SmConnection cmdlet.

```
Open-smconnection  -SMSbaseurl
https:\\snapctr.demo.netapp.com:8146\
```

   The username and password prompt is displayed.

2. Add resources by using the Add-SmResources cmdlet.

   This example shows how to add a SAP HANA database of SingleContainer type:

```
C:\PS> Add-SmResource -HostName '10.232.204.42' -PluginCode 'HANA'
-DatabaseName H10 -ResourceType SingleContainer -StorageFootPrint
(@{"VolumeName"="HanaData10";"StorageSystem"="vserver_scauto_primary
"}) -SID 'H10' -filebackuppath '/tmp/HanaFileLog' -userstorekeys
'HS10' -osdbuser 'h10adm' -filebackupprefix 'H10_'
```

   This example shows how to add a SAP HANA database of MultipleContainers type:

```
C:\PS> Add-SmResource -HostName 'vp-hana2.gdl.englab.netapp.com'
-PluginCode 'HANA' -DatabaseName MDC_MT -ResourceType
MultipleContainers -StorageFootPrint
(@{"VolumeName"="VP_HANA2_data";"StorageSystem"="buck.gdl.englab.net
app.com"}) -sid 'A12' -userstorekeys 'A12KEY' -TenantType
'MultiTenant'
```

This example shows how to create a non-data volume resource:

```
C:\PS> Add-SmResource -HostName 'SNAPCENTERN42.sccore.test.com'
-PluginCode 'hana' -ResourceName NonDataVolume -ResourceType
NonDataVolume -StorageFootPrint
(@{"VolumeName"="ng_pvol";"StorageSystem"="vserver_scauto_primary"})
-sid 'S10'
```

3. Create a backup policy by using the Add-SmPolicy cmdlet.

   This example creates a backup policy for a Snapshot copy-based backup:

   ```
   C:\PS> Add-SmPolicy -PolicyName hana_snapshotbased -PolicyType
   Backup -PluginPolicyType hana -BackupType SnapShotBasedBackup
   ```

   This example creates a backup policy for a File-Based backup:

   ```
   C:\PS> Add-SmPolicy -PolicyName hana_Filebased -PolicyType Backup
   -PluginPolicyType hana -BackupType FileBasedBackup
   ```

4. Protect the resource or add a new resource group to SnapCenter by using the Add-SmResourceGroup cmdlet.

   This example protects a single container resource:

   ```
   C:\PS> Add-SmProtectResource -PluginCode HANA  -Policies
   hana_snapshotbased,hana_Filebased
    -Resources @{"Host"="host.example.com";"UID"="SID"} -Description
   test -usesnapcenterwithoutfilesystemconsistency
   ```

   This example protects a multiple containers resource:

```
C:\PS> Add-SmProtectResource -PluginCode HANA  -Policies
hana_snapshotbased,hana_Filebased
 -Resources @{"Host"="host.example.com";"UID"="MDC\SID"}
-Description test -usesnapcenterwithoutfilesystemconsistency
```

This example creates a new resource group with the specified policy and resources:

```
C:\PS> Add-SmResourceGroup -ResourceGroupName
'ResourceGroup_with_SingleContainer_MultipleContainers_Resources'
-Resources
@(@{"Host"="sccorelinux61.sccore.test.com";"Uid"="SID"},@{"Host"="sc
corelinux62.sccore.test.com";"Uid"="MDC\SID"})
  -Policies hana_snapshotbased,hana_Filebased
-usesnapcenterwithoutfilesystemconsistency  -plugincode 'HANA'
```

This example creates a non-data volume resource group:

```
C:\PS> Add-SmResourceGroup -ResourceGroupName
'Mixed_RG_backup_when_Remove_Backup_throguh_BackupName_windows'
-Resources
@(@{"Host"="SNAPCENTERN42.sccore.test.com";"Uid"="H11";"PluginName"=
"hana"},@{"Host"="SNAPCENTERN42.sccore.test.com";"Uid"="MDC\H31";"Pl
uginName"="hana"},@{"Host"="SNAPCENTERN42.sccore.test.com";"Uid"="No
nDataVolume\S10\NonDataVolume";"PluginName"="hana"}) -Policies
hanaprimary
```

5. Initiate a new backup job by using the New-SmBackup cmdlet.

   This example shows how to backup a resource group:

   ```
   C:\PS> New-SMBackup -ResourceGroupName
   'ResourceGroup_with_SingleContainer_MultipleContainers_Resources'
   -Policy hana_snapshotbased
   ```

   This example backs up a protected resource:

   ```
   C:\PS> New-SMBackup -Resources
   @{"Host"="10.232.204.42";"Uid"="MDC\SID";"PluginName"="hana"}
   -Policy hana_Filebased
   ```

6. Monitor the job status (running, completed, or failed) by using the Get-smJobSummaryReport cmdlet.

```
PS C:\> Get-smJobSummaryReport -JobID 123
```

7. Monitor the backup job details like backup ID, backup name to perform restore or clone operation by using the Get-SmBackupReport cmdlet.

```
PS C:\> Get-SmBackupReport -JobId 351
Output:
BackedUpObjects           : {DB1}
FailedObjects             : {}
IsScheduled               : False
HasMetadata               : False
SmBackupId                : 269
SmJobId                   : 2361
StartDateTime             : 10/4/2016 11:20:45 PM
EndDateTime               : 10/4/2016 11:21:32 PM
Duration                  : 00:00:46.2536470
CreatedDateTime           : 10/4/2016 11:21:09 PM
Status                    : Completed
ProtectionGroupName       : Verify_ASUP_Message_windows
SmProtectionGroupId       : 211
PolicyName                : test2
SmPolicyId                : 20
BackupName                : Verify_ASUP_Message_windows_scc54_10-04-
2016_23.20.46.2758
VerificationStatus        : NotVerified
VerificationStatuses      :
SmJobError                :
BackupType                : SCC_BACKUP
CatalogingStatus          : NotApplicable
CatalogingStatuses        :
ReportDataCreatedDateTime :
```

The information regarding the parameters that can be used with the cmdlet and their descriptions can be obtained by running *Get-Help command_name*. Alternatively, you can also refer to the SnapCenter Software Cmdlet Reference Guide.

Back up resource groups

A resource group is a collection of resources on a host. A backup operation on the resource group is performed on all resources defined in the resource group.

Before you begin

- You must have created a resource group with a policy attached.

- If you want to back up a resource that has a SnapMirror relationship with a secondary storage, the ONTAP role assigned to the storage user should include the “snapmirror all” privilege. However, if you are using the “vsadmin” role, then the “snapmirror all” privilege is not required.

About this task

You can back up a resource group on demand from the Resources page. If a resource group has a policy attached and a schedule configured, then backups occur automatically according to the schedule.

Steps

1. In the left navigation pane, select **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page, select **Resource Group** from the **View** list.

   You can search the resource group either by entering the resource group name in the search box or by selecting , and then selecting the tag. You can then select to close the filter pane.

3. In the Resource Groups page, select the resource group that you want to back up, and then select **Back up Now**.
4. In the Backup page, perform the following steps:
   a. If you associated multiple policies with the resource group, from the **Policy** drop-down list, select the policy that you want to use for backup.

      If the policy selected for the on-demand backup is associated with a backup schedule, the on-demand backups will be retained based on the retention settings specified for the schedule type.

   b. Select **Backup**.
5. Monitor the operation progress by selecting **Monitor** > **Jobs**.

Monitor SAP HANA databases backup operations

You can monitor the progress of different backup operations by using the SnapCenterJobs page. You might want to check the progress to determine when it is complete or if there is an issue.

About this task

The following icons appear on the Jobs page and indicate the corresponding state of the operations:

- In progress
- Completed successfully
- Failed
- Completed with warnings or could not start due to warnings
- Queued
- Canceled

Steps

1. In the left navigation pane, click **Monitor**.
2. In the Monitor page, click **Jobs**.

3. In the Jobs page, perform the following steps:
   a. Click to filter the list so that only backup operations are listed.
   b. Specify the start and end dates.
   c. From the **Type** drop-down list, select **Backup**.
   d. From the **Status** drop-down, select the backup status.
   e. Click **Apply** to view the operations completed successfully.
4. Select a backup job, and then click **Details** to view the job details.

   Though the backup job status displays , when you click on job details you might see that some of the child tasks of the backup operation are still in progress or marked with warning signs.

5. In the Job Details page, click **View logs**.

   The **View logs** button displays the detailed logs for the selected operation.

Monitor data protection operations on SAP HANA databases in the Activity pane

The Activity pane displays the five most recent operations performed. The Activity pane also displays when the operation was initiated and the status of the operation.

The Activity pane displays information regarding backup, restore, clone, and scheduled backup operations.

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. Click on the Activity pane to view the five most recent operations.

   When you click one of the operations, the operation details are listed in the **Job Details** page.

Cancel backup operations for SAP HANA

You can cancel backup operations that are queued.

What you will need

- You must be logged in as the SnapCenter Admin or job owner to cancel operations.
- You can cancel a backup operation from either the **Monitor** page or the **Activity** pane.
- You cannot cancel a running backup operation.
- You can use the SnapCenter GUI, PowerShell cmdlets, or CLI commands to cancel the backup operations.
- The **Cancel Job** button is disabled for operations that cannot be canceled.
- If you selected **All members of this role can see and operate on other members objects** in Users\Groups page while creating a role, you can cancel the queued backup operations of other members while using that role.

Steps

1. Perform one of the following actions:

From the…	Action
Monitor page	a. In the left navigation pane, click **Monitor** > **Jobs**. b. Select the operation, and then click **Cancel Job**.
Activity pane	a. After initiating the backup operation, click on the Activity pane to view the five most recent operations. b. Select the operation. c. In the Job Details page, click **Cancel Job**.

The operation is canceled, and the resource is reverted to the previous state.

View SAP HANA database backups and clones in the Topology page

When you are preparing to back up or clone a resource, you might find it helpful to view a graphical representation of all backups and clones on the primary and secondary storage.

About this task

You can review the following icons in the Manage Copies view to determine whether the backups and clones are available on the primary or secondary storage (Mirror copies or Vault copies).

- displays the number of backups and clones that are available on the primary storage.
- displays the number of backups and clones that are mirrored on the secondary storage using SnapMirror technology.
- displays the number of backups and clones that are replicated on the secondary storage using SnapVault technology.

The number of backups displayed includes the backups deleted from the secondary storage. For example, if you created 6 backups using a policy to retain only 4 backups, the number of backups displayed is 6.

Clones of a backup of a version-flexible mirror on a mirror-vault type volume are displayed in the topology view, but the mirror backup count in the topology view does not include the version-flexible backup.

In the Topology page, you can see all of the backups and clones that are available for the selected resource or resource group. You can view the details of those backups and clones, and then select them to perform data protection operations.

If you have secondary relationship as SnapMirror active sync (initially released as SnapMirror Business Continuity [SM-BC]), you can see following additional icons:

- The replica site is up.
- The replica site is down.
- The secondary mirror or vault relationship has not been re-established.

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page, either select the resource or resource group from the **View** drop-down list.
3. Select the resource either from the resource details view or from the resource group details view.

   If the resource is protected, the topology page of the selected resource is displayed.

4. Review the **Summary card** to see a summary of the number of backups and clones available on the primary and secondary storage.

   The **Summary Card** section displays the total number of File-Based backups, Snapshot copy-based backups, and clones.

   Clicking the **Refresh** button starts a query of the storage to display an accurate count.

   If SnapLock enabled backup is taken, then clicking the **Refresh** button refreshes the primary and secondary SnapLock expiry time retrieved from ONTAP. A weekly schedule also refreshes the primary and secondary SnapLock expiry time retrieved from ONTAP.

   When the application resource is spread across multiple volumes, the SnapLock expiry time for the backup will be the longest SnapLock expiry time that is set for a Snapshot in a volume. The longest SnapLock expiry time is retrieved from ONTAP.

   For SnapMirror active sync, clicking the **Refresh** button refreshes the SnapCenter backup inventory by querying ONTAP for both primary and replica sites. A weekly schedule also performs this activity for all databases containing SnapMirror active sync relationship.

   - For SnapMirror active sync and only for ONTAP 9.14.1, Async Mirror or Async MirrorVault relationships to the new primary destination should be manually configured after failover. From ONTAP 9.15.1 onwards Async Mirror or Async MirrorVault is auto configured to the new primary destination.
   - After failover, a backup should be created for SnapCenter to be aware of the failover. You can click **Refresh** only after a backup has been created.

5. In the Manage Copies view, click **Backups** or **Clones** from the primary or secondary storage to see details of a backup or clone.

   The details of the backups and clones are displayed in a table format.

6. Select the backup from the table, and then click the data protection icons to perform restore, clone, and delete operations.

   You cannot rename or delete backups that are on the secondary storage.

7. If you want to delete a clone, select the clone from the table, and then click .
8. If you want to split a clone, select the clone from the table, and then click .

Restore SAP HANA Databases

Restore workflow

The restore and recovery workflow includes planning, performing the restore operations, and monitoring the operations.

The following workflow shows the sequence in which you must perform the restore operation:

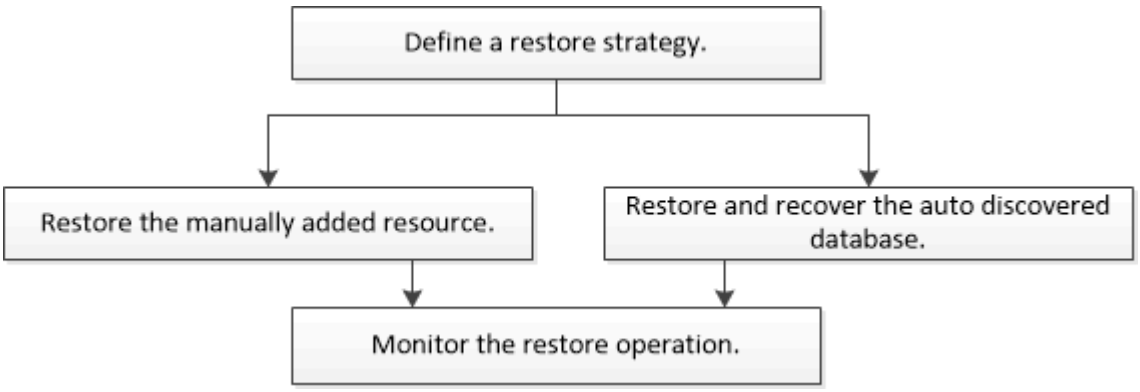

You can also use PowerShell cmdlets manually or in scripts to perform backup, restore, and clone operations. The SnapCenter cmdlet help and the cmdlet reference information contain detailed information about PowerShell cmdlets.

SnapCenter Software Cmdlet Reference Guide.

Restore and recover a manually added resource backup

You can use SnapCenter to restore and recover data from one or more backups.

Before you begin

- You must have backed up the resource or resource groups.
- You must have canceled any backup operation that is currently in progress for the resource or resource group that you want to restore.
- For pre restore, post restore, mount, and unmount commands, you should check if the commands exist in the command list available on the plug-in host from the following paths:
  - Default location on the Windows host: *C:\Program Files\NetApp\SnapCenter\Snapcenter Plug-in Creator\etc\allowed_commands.config*
  - Default location on the Linux host: */opt/NetApp/snapcenter/scc/etc/allowed_commands.config*

If the commands do not exist in the command list, then the operation will fail.

About this task

- File-based backup copies cannot be restored from SnapCenter.
- After upgrading to SnapCenter 4.3, the backups taken in SnapCenter 4.2 can be restored but cannot be

recovered. You must use HANA studio or HANA recovery scripts external to SnapCenter to recover the backups taken in SnapCenter 4.2.

- For ONTAP 9.12.1 and below version, the clones created from the SnapLock Vault Snapshots as part of restore will inherit the SnapLock Vault expiry time. Storage admin should manually cleanup the clones post the SnapLock expiry time.
- For SnapMirror active sync restore operation, you must select the backup from the primary location.

SnapCenter UI

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page, filter resources from the **View** drop-down list based on resource type.

   The resources are displayed along with type, host, associated resource groups and policies, and status.

   Although a backup might be for a resource group, when you restore, you must select the individual resources you want to restore.

   If the resource is not protected, “Not protected” is displayed in the Overall Status column. This can mean either that the resource is not protected, or that the resource was backed up by a different user.

3. Select the resource, or select a resource group and then select a resource in that group.

   The resource topology page is displayed.

4. From the Manage Copies view, select **Backups** either from the primary or secondary (mirrored or vaulted) storage systems.
5. In the Primary backup(s) table, select the backup that you want to restore from, and then click .

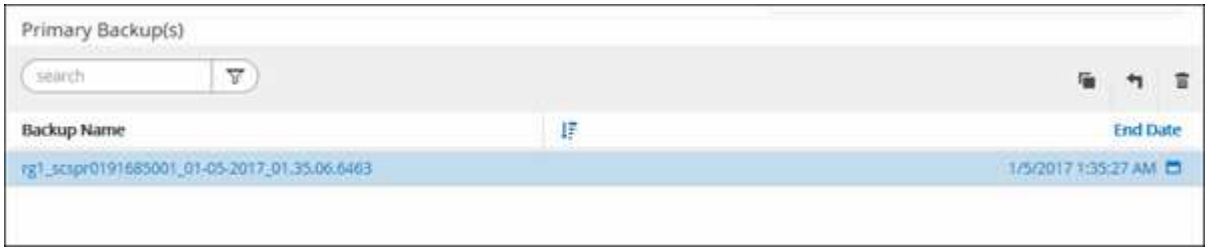

6. In the Restore Scope page, select either **Complete Resource** or **File Level**.

   a. If you select **Complete Resource**, all of the configured data volumes of the SAP HANA database are restored.

      If the resource contains volumes or qtrees, the Snapshots taken after the Snapshot selected for restore on such volumes or qtrees are deleted and cannot be recovered. Also, if any other resource is hosted on same volumes or qtrees, then that resource is also deleted.

   b. If you select **File Level**, then you can either select **All** or select the specific volumes or qtrees, and then enter the path related to those volumes or qtrees, separated by commas

      - You can select multiple volumes and qtrees.
      - If the resource type is LUN, the entire LUN is restored.

        You can select multiple LUNs.

   If you select **All**, all the files on the volumes, qtrees, or LUNs are restored.

7. In the Pre ops page, enter pre restore and unmount commands to run before performing a restore job.

Unmount commands are not available for auto discovered resources.

8. In the Post ops page, enter mount and post restore commands to run after performing a restore job.

   Mount commands are not available for auto discovered resources.

9. In the Notification page, from the **Email preference** drop-down list, select the scenarios in which you want to send the emails.

   You must also specify the sender and receiver email addresses and the subject of the email. SMTP must also be configured on the **Settings** > **Global Settings** page.

10. Review the summary, and then click **Finish**.
11. Monitor the operation progress by clicking **Monitor** > **Jobs**.

PowerShell cmdlets

Steps

1. Initiate a connection session with the SnapCenter Server for a specified user by using the Open-SmConnection cmdlet.

```
PS C:\> Open-Smconnection
```

2. Retrieve the information about the one or more backups that you want to restore by using the Get-SmBackup and Get-SmBackupReport cmdlets.

   This example displays information about all available backups:

```
PS C:\> Get-SmBackup

BackupId                      BackupName
BackupTime                    BackupType
--------                      ----------
----------                    ----------
  1               Payroll Dataset_vise-f6_08... 8/4/2015    11:02:32
AM                 Full Backup
  2               Payroll Dataset_vise-f6_08... 8/4/2015    11:23:17
AM
```

   This example displays detailed information about the backup from January 29th 2015 to February 3rd, 2015:

```
PS C:\> Get-SmBackupReport -FromDateTime "1/29/2015" -ToDateTime
"2/3/2015"

   SmBackupId           : 113
   SmJobId              : 2032
   StartDateTime        : 2/2/2015 6:57:03 AM
   EndDateTime          : 2/2/2015 6:57:11 AM
   Duration             : 00:00:07.3060000
   CreatedDateTime      : 2/2/2015 6:57:23 AM
   Status               : Completed
   ProtectionGroupName  : Clone
   SmProtectionGroupId  : 34
   PolicyName           : Vault
   SmPolicyId           : 18
   BackupName           : Clone_SCSPR0019366001_02-02-2015_06.57.08
   VerificationStatus   : NotVerified

   SmBackupId           : 114
   SmJobId              : 2183
   StartDateTime        : 2/2/2015 1:02:41 PM
   EndDateTime          : 2/2/2015 1:02:38 PM
   Duration             : -00:00:03.2300000
   CreatedDateTime      : 2/2/2015 1:02:53 PM
   Status               : Completed
   ProtectionGroupName  : Clone
   SmProtectionGroupId  : 34
   PolicyName           : Vault
   SmPolicyId           : 18
   BackupName           : Clone_SCSPR0019366001_02-02-2015_13.02.45
   VerificationStatus   : NotVerified
```

3. Restore data from the backup by using the Restore-SmBackup cmdlet.

```
Restore-SmBackup -PluginCode 'DummyPlugin' -AppObjectId
'scc54.sccore.test.com\DummyPlugin\NTP\DB1' -BackupId 269
-Confirm:$false
output:
Name                : Restore
'scc54.sccore.test.com\DummyPlugin\NTP\DB1'
Id                  : 2368
StartTime           : 10/4/2016 11:22:02 PM
EndTime             :
IsCancellable       : False
IsRestartable       : False
IsCompleted         : False
IsVisible           : True
IsScheduled         : False
PercentageCompleted : 0
Description         :
Status              : Queued
Owner               :
Error               :
Priority            : None
Tasks               : {}
ParentJobID         : 0
EventId             : 0
JobTypeId           :
ApisJobKey          :
ObjectId            : 0
PluginCode          : NONE
PluginName          :
```

The information regarding the parameters that can be used with the cmdlet and their descriptions can be obtained by running *Get-Help command_name*. Alternatively, you can also refer to the SnapCenter Software Cmdlet Reference Guide.

Restore and recover an auto discovered database backup

You can use SnapCenter to restore and recover data from one or more backups.

Before you begin

- You must have backed up the resource or resource groups.
- You must have canceled any backup operation that is currently in progress for the resource or resource group that you want to restore.
- For pre restore, post restore, mount, and unmount commands, you should check if the commands exist in the command list available on the plug-in host from the following paths:
  - Default location on the Windows host: *C:\Program Files\NetApp\SnapCenter\Snapcenter Plug-in*

Creator\etc\allowed_commands.config

- Default location on the Linux host: */opt/NetApp/snapcenter/scc/etc/allowed_commands.config*

If the commands do not exist in the command list, then the operation will fail.

About this task

- File-based backup copies cannot be restored from SnapCenter.
- After upgrading to SnapCenter 4.3, the backups taken in SnapCenter 4.2 can be restored but cannot be recovered. You must use HANA studio or HANA recovery scripts external to SnapCenter to recover the backups taken in SnapCenter 4.2.
- For ONTAP 9.12.1 and earlier versions, the clones created from the SnapLock Vault Snapshots as part of restore will inherit the SnapLock Vault expiry time. Storage admin should manually cleanup the clones post the SnapLock expiry time.
- For SnapMirror active sync restore operation, you must select the backup from the primary location.

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page, filter resources from the **View** drop-down list based on resource type.

   The resources are displayed along with type, host, associated resource groups and policies, and status.

   Although a backup might be for a resource group, when you restore, you must select the individual resources you want to restore.

   If the resource is not protected, “Not protected” is displayed in the Overall Status column. This can mean either that the resource is not protected, or that the resource was backed up by a different user.

3. Select the resource, or select a resource group and then select a resource in that group.

   The resource topology page is displayed.

4. From the Manage Copies view, select **Backups** either from the primary or secondary (mirrored or vaulted) storage systems.
5. In the Primary backup(s) table, select the backup that you want to restore from, and then click .

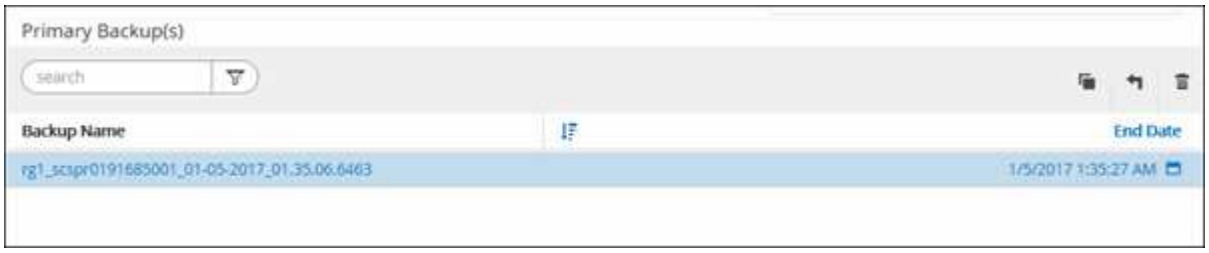

6. In the Restore Scope page, select **Complete Resource** to restore the configured data volumes of the SAP HANA database.

   You can select either **Complete Resource** (with or without **Volume Revert**), or **Tenant Database**.

Recovery operation is not supported by SnapCenter Server for multiple tenants when user selects either the **Tenant Database** or **Complete Restore** option. You must use HANA studio or HANA python script to perform the recovery operation.

a. Select **Volume Revert** if you want to restore the entire volume.

   This option is available for backups taken in SnapCenter 4.3 in NFS environments.

   If the resource contains volumes or qtrees, the Snapshots taken after the Snapshot selected for restore on such volumes or qtrees are deleted and cannot be recovered. Also, if any other resource is hosted on the same volumes or qtrees, then that resource is also deleted. This is applicable when **Complete Resource** with **Volume Revert** option is selected for restore.

b. Select **Tenant Database**.

   This option is available only for MDC resources.

   Ensure to stop the tenant database before performing the restore operation.

   If you select **Tenant Database** option, you must use HANA studio or use HANA recovery scripts external to SnapCenter to perform recovery operation.

7. In the Recovery scope page, select one of the following options:

If you…	Do this…
Want to recover as close as possible to the current time	Select **Recover to most recent state**. For single container resources specify one or more log and catalog backup locations. For multitenant database container (MDC) resources specify one or more log backup locations and the backup catalog location. For MDC resources, the path should contain both system database and tenant database logs.

If you…	Do this…
Want to recover to the specified point in time	Select **Recover to point in time**. a. Select the time zone. Browser timezone is populated by default. The selected time zone along with the input time is converted to absolute GMT. b. Enter date and time. For example, the HANA Linux host is located in Sunnyvale, CA and the user in Raleigh, NC is recovering the logs in to SnapCenter. The time difference between both these locations is 3 hours, and since the user has logged in from Raleigh, NC, the default browser time zone that will be selected in the GUI is GMT-04:00. If the user wants to perform a recovery to 5 a.m .Sunnyvale, CA, then the user has to set the browser time zone to the HANA Linux host time zone, which is GMT-07:00 and specify the date and time as 5:00 a.m. For single container resources specify one or more log and catalog backup locations. For MDC resources, specify one or more log backup locations and the backup catalog location. For MDC resources, the path should contain both system database and tenant database logs.
Want to recover to a specific data backup	Select **Recover to specified data backup**.
Do not want to recover	Select **No recovery**. You must perform the recovery operation manually from the HANA studio.

You can recover only those backups that are taken after upgrading to SnapCenter 4.3, provided both the host and the plug-in are upgraded to SnapCenter 4.3, and the backups selected for restore are taken after the resource is converted or discovered as auto discovered resource.

8. In the Pre ops page, enter pre restore and unmount commands to run before performing a restore job.

   Unmount commands are not available for auto discovered resources.

9. In the Post ops page, enter mount and post restore commands to run after performing a restore job.

Mount commands are not available for auto discovered resources.

10. In the Notification page, from the **Email preference** drop-down list, select the scenarios in which you want to send the emails.

    You must also specify the sender and receiver email addresses and the subject of the email. SMTP must also be configured on the **Settings** > **Global Settings** page.

11. Review the summary, and then click **Finish**.
12. Monitor the operation progress by clicking **Monitor** > **Jobs**.

Restore resources using PowerShell cmdlets

Restoring a resource backup includes initiating a connection session with the SnapCenter Server, listing the backups and retrieving backup information, and restoring a backup.

You must have prepared the PowerShell environment to execute the PowerShell cmdlets.

Steps

1. Initiate a connection session with the SnapCenter Server for a specified user by using the Open-SmConnection cmdlet.

```
PS C:\> Open-Smconnection
```

2. Retrieve the information about the one or more backups that you want to restore by using the Get-SmBackup and Get-SmBackupReport cmdlets.

   This example displays information about all available backups:

```
PS C:\> Get-SmBackup

BackupId                      BackupName                    BackupTime
BackupType
--------                      ----------                    ----------
----------
  1               Payroll Dataset_vise-f6_08... 8/4/2015    11:02:32 AM
Full Backup
  2               Payroll Dataset_vise-f6_08... 8/4/2015    11:23:17 AM
```

   This example displays detailed information about the backup from January 29th 2015 to February 3rd, 2015:

```
PS C:\> Get-SmBackupReport -FromDateTime "1/29/2015" -ToDateTime
"2/3/2015"

   SmBackupId           : 113
   SmJobId              : 2032
   StartDateTime        : 2/2/2015 6:57:03 AM
   EndDateTime          : 2/2/2015 6:57:11 AM
   Duration             : 00:00:07.3060000
   CreatedDateTime      : 2/2/2015 6:57:23 AM
   Status               : Completed
   ProtectionGroupName  : Clone
   SmProtectionGroupId  : 34
   PolicyName           : Vault
   SmPolicyId           : 18
   BackupName           : Clone_SCSPR0019366001_02-02-2015_06.57.08
   VerificationStatus   : NotVerified

   SmBackupId           : 114
   SmJobId              : 2183
   StartDateTime        : 2/2/2015 1:02:41 PM
   EndDateTime          : 2/2/2015 1:02:38 PM
   Duration             : -00:00:03.2300000
   CreatedDateTime      : 2/2/2015 1:02:53 PM
   Status               : Completed
   ProtectionGroupName  : Clone
   SmProtectionGroupId  : 34
   PolicyName           : Vault
   SmPolicyId           : 18
   BackupName           : Clone_SCSPR0019366001_02-02-2015_13.02.45
   VerificationStatus   : NotVerified
```

3. Restore data from the backup by using the Restore-SmBackup cmdlet.

```
Restore-SmBackup -PluginCode 'DummyPlugin' -AppObjectId
'scc54.sccore.test.com\DummyPlugin\NTP\DB1' -BackupId 269
-Confirm:$false
output:
Name                : Restore
'scc54.sccore.test.com\DummyPlugin\NTP\DB1'
Id                  : 2368
StartTime           : 10/4/2016 11:22:02 PM
EndTime             :
IsCancellable       : False
IsRestartable       : False
IsCompleted         : False
IsVisible           : True
IsScheduled         : False
PercentageCompleted : 0
Description         :
Status              : Queued
Owner               :
Error               :
Priority            : None
Tasks               : {}
ParentJobID         : 0
EventId             : 0
JobTypeId           :
ApisJobKey          :
ObjectId            : 0
PluginCode          : NONE
PluginName          :
```

The information regarding the parameters that can be used with the cmdlet and their descriptions can be obtained by running *Get-Help command_name*. Alternatively, you can also refer to the SnapCenter Software Cmdlet Reference Guide.

Monitor SAP HANA databases restore operations

You can monitor the progress of different SnapCenter restore operations by using the Jobs page. You might want to check the progress of an operation to determine when it is complete or if there is an issue.

About this task

Post-restore states describe the conditions of the resource after a restore operation and any further restore actions that you can take.

The following icons appear on the Jobs page, and indicate the state of the operation:

- In progress

- Completed successfully
- Failed
- Completed with warnings or could not start due to warnings
- Queued
- Canceled

Steps

1. In the left navigation pane, click **Monitor**.
2. In the **Monitor** page, click **Jobs**.
3. In the **Jobs** page, perform the following steps:
   a. Click to filter the list so that only restore operations are listed.
   b. Specify the start and end dates.
   c. From the **Type** drop-down list, select **Restore**.
   d. From the **Status** drop-down list, select the restore status.
   e. Click **Apply** to view the operations that have been completed successfully.
4. Select the restore job, and then click **Details** to view the job details.
5. In the **Job Details** page, click **View logs**.

   The **View logs** button displays the detailed logs for the selected operation.

Clone SAP HANA resource backups

Clone workflow

The clone workflow includes performing the clone operation and monitoring the operation.

About this task

- You can clone on the source SAP HANA server.
- You might clone resource backups for the following reasons:
  - To test functionality that has to be implemented using the current resource structure and content during application development cycles
  - For data extraction and manipulation tools when populating data warehouses
  - To recover data that was mistakenly deleted or changed

The following workflow shows the sequence in which you must perform the clone operation:

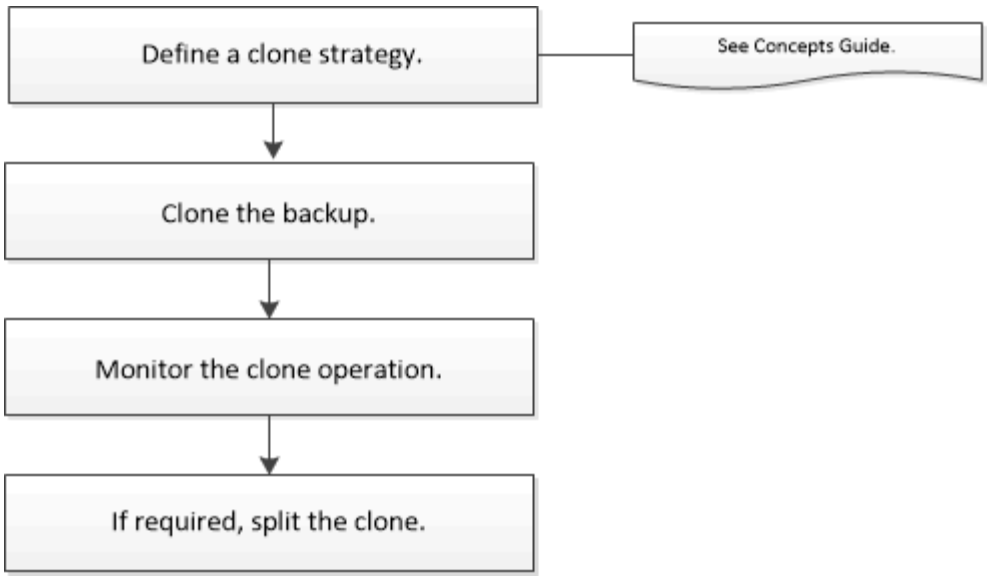

You can also use PowerShell cmdlets manually or in scripts to perform backup, restore, and clone operations. The SnapCenter cmdlet help and the cmdlet reference information contain detailed information about PowerShell cmdlets.

Clone a SAP HANA database backup

You can use SnapCenter to clone a backup. You can clone from primary or secondary backup.

Before you begin

- You should have backed up the resources or resource group.
- You should ensure that the aggregates hosting the volumes should be in the assigned aggregates list of the storage virtual machine (SVM).
- You cannot clone file-based backups.
- The target clone server should have the same SAP HANA instance SID that is provided in the Target Clone SID field.
- For pre clone or post clone commands, you should check if the commands exist in the command list available on the plug-in host from the following paths:
  - For Windows: *C:\Program Files\NetApp\SnapCenter\Snapcenter Plug-in Creator\etc\allowed_commands.config*
  - For Linux: */opt/NetApp/snapcenter/scc/etc/allowed_commands.config*

If the commands do not exist in the command list, then the operation will fail.

About this task

- For information about clone split operation limitations, see ONTAP 9 Logical Storage Management Guide.
- For ONTAP 9.12.1 and earlier version, the clones created from the SnapLock Vault Snapshots as part of restore will inherit the SnapLock Vault expiry time. Storage admin should manually cleanup the clones post the SnapLock expiry time.

SnapCenter UI

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page, filter resources from the **View** drop-down list based on resource type.

   The resources are displayed along with information such as type, host, associated resource groups and policies, and status.

3. Select the resource or resource group.

   You must select a resource if you select a resource group.

   The resource or resource group topology page is displayed.

4. From the Manage Copies view, select **Backups** either from the primary or secondary (mirrored or vaulted) storage systems.
5. Select the data backup from the table, and then click .
6. In the Location page, perform the following actions:

For this field…	Do this…
Plug-in host	Select the host on which the clone should be mounted and the plug-in is installed.
Target Clone SID	Enter the SAP HANA instance ID to clone from the existing backups.
NFS Export IP Address	Enter IP addresses or the host names on which the cloned volumes will be exported.
iSCSI Initiator	Enter iSCSI initiator name of the host to which the LUNs are exported. This option is available only if you selected LUN resource type.
Protocol	Enter the LUN protocol. This option is available only if you selected LUN resource type.

   If the resource selected is a LUN and you are cloning from a secondary backup, then the destination volumes are listed. A single source can have multiple destination volumes.

   Before cloning, you must ensure that the iSCSI initiator or the FCP is present and are configured and logged into alternate hosts.

7. In the Scripts page, perform the following steps:

   The scripts are run on the plug-in host.

   a. Enter the commands for pre clone or post clone that should be run before or after the clone

operation, respectively.

- Pre clone command: delete existing databases with the same name
- Post clone command: verify a database or start a database.

b. Enter the mount command to mount a file system to a host.

Mount command for a volume or qtree on a Linux machine:

Example for NFS: `mount VSERVER_DATA_IP:%VOLUME_NAME_Clone /mnt`

8. In the Notification page, from the **Email preference** drop-down list, select the scenarios in which you want to send the emails.

You must also specify the sender and receiver email addresses, and the subject of the email.

9. Review the summary, and then click **Finish**.
10. Monitor the operation progress by clicking **Monitor** > **Jobs**.

PowerShell cmdlets

Steps

1. Initiate a connection session with the SnapCenter Server for a specified user by using the Open-SmConnection cmdlet.

```
Open-SmConnection  -SMSbaseurl
https:\\snapctr.demo.netapp.com:8146/
```

2. Retrieve the backups to perform the clone operation by using the Get-SmBackup cmdlet.

This example shows that two backups are available for cloning:

```
 C:\PS> Get-SmBackup

        BackupId                    BackupName
BackupTime                    BackupType
        --------                    ----------
----------                    ----------
         1                     Payroll Dataset_vise-f6_08...
8/4/2015 11:02:32 AM          Full Backup
         2                     Payroll Dataset_vise-f6_08...
8/4/2015 11:23:17 AM
```

3. Initiate a clone operation from an existing backup and specify the NFS export IP addresses on which the cloned volumes are exported.

This example shows that the backup to be cloned has an NFSExportIPs address of 10.232.206.169:

```
New-SmClone -AppPluginCode hana -BackupName
scsccore1_sccore_test_com_hana_H73_scsccore1_06-07-
2017_02.54.29.3817 -Resources
@{"Host"="scsccore1.sccore.test.com";"Uid"="H73"}  -CloneToInstance
shivscc4.sccore.test.com -mountcommand 'mount
10.232.206.169:%hana73data_Clone /hana83data'
-preclonecreatecommands '/home/scripts/scpre_clone.sh'
-postclonecreatecommands '/home/scripts/scpost_clone.sh'
```

If NFSExportIPs is not specified, the default is exported to the clone target host.

4. Verify that the backups were cloned successfully by using the Get-SmCloneReport cmdlet to view the clone job details.

   You can view details such as clone ID, start date and time, end date and time.

```
PS C:\> Get-SmCloneReport -JobId 186

    SmCloneId           : 1
    SmJobId             : 186
    StartDateTime       : 8/3/2015 2:43:02 PM
    EndDateTime         : 8/3/2015 2:44:08 PM
    Duration            : 00:01:06.6760000
    Status              : Completed
    ProtectionGroupName : Draper
    SmProtectionGroupId : 4
    PolicyName          : OnDemand_Clone
    SmPolicyId          : 4
    BackupPolicyName    : OnDemand_Full_Log
    SmBackupPolicyId    : 1
    CloneHostName       : SCSPR0054212005.mycompany.com
    CloneHostId         : 4
    CloneName           : Draper__clone__08-03-2015_14.43.53
    SourceResources     : {Don, Betty, Bobby, Sally}
    ClonedResources     : {Don_DRAPER, Betty_DRAPER, Bobby_DRAPER,
Sally_DRAPER}
    SmJobError          :
```

Monitor SAP HANA database clone operations

You can monitor the progress of SnapCenter clone operations by using the Jobs page. You might want to check the progress of an operation to determine when it is complete or if there is an issue.

About this task

The following icons appear on the Jobs page, and indicate the state of the operation:

- In progress
- Completed successfully
- Failed
- Completed with warnings or could not start due to warnings
- Queued
- Canceled

Steps

1. In the left navigation pane, click **Monitor**.
2. In the **Monitor** page, click **Jobs**.
3. In the **Jobs** page, perform the following steps:
   a. Click to filter the list so that only clone operations are listed.
   b. Specify the start and end dates.
   c. From the **Type** drop-down list, select **Clone**.
   d. From the **Status** drop-down list, select the clone status.
   e. Click **Apply** to view the operations that are completed successfully.
4. Select the clone job, and then click **Details** to view the job details.
5. In the Job Details page, click **View logs**.

Split a clone

You can use SnapCenter to split a cloned resource from the parent resource. The clone that is split becomes independent of the parent resource.

About this task

- You cannot perform the clone split operation on an intermediate clone.

  For example, after you create clone1 from a database backup, you can create a backup of clone1, and then clone this backup (clone2). After you create clone2, clone1 is an intermediate clone, and you cannot perform the clone split operation on clone1. However, you can perform the clone split operation on clone2.

  After splitting clone2, you can perform the clone split operation on clone1 because clone1 is no longer the intermediate clone.

- When you split a clone, the backup copies and clone jobs of the clone are deleted.
- For information about FlexClone volume split operations, see, Split a FlexClone volume from its parent volume.
- Ensure that the volume or aggregate on the storage system is online.

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. In the **Resources** page, select the appropriate option from the View list:

Option	Description
For database applications	Select **Database** from the View list.
For file systems	Select **Path** from the View list.

3. Select the appropriate resource from the list.

   The resource topology page is displayed.

4. From the **Manage Copies** view, select the cloned resource (for example, the database or LUN), and then click .
5. Review the estimated size of the clone that is to be split and the required space available on the aggregate, and then click **Start**.
6. Monitor the operation progress by clicking **Monitor** > **Jobs**.

   The clone split operation stops responding if the SMCore service restarts. You should run the Stop-SmJob cmdlet to stop the clone split operation, and then retry the clone split operation.

   If you want a longer poll time or shorter poll time to check whether the clone is split or not, you can change the value of *CloneSplitStatusCheckPollTime* parameter in *SMCoreServiceHost.exe.config* file to set the time interval for SMCore to poll for the status of the clone split operation. The value is in milliseconds and the default value is 5 minutes.

   For example:

```
<add key="CloneSplitStatusCheckPollTime" value="300000" />
```

   The clone split start operation fails if backup, restore, or another clone split is in progress. You should restart the clone split operation only after the running operations are complete.

Related information

SnapCenter clone or verification fails with aggregate does not exist

Delete or split SAP HANA database clones after upgrading SnapCenter

After upgrading to SnapCenter 4.3, you will no longer see the clones. You can delete the clone or split the clones from the Topology page of the resource from which the clones were created.

About this task

If you want to locate the storage footprint of the hidden clones, run the following command: `Get-SmClone -ListStorageFootprint`

Steps

1. Delete the backups of the cloned resources by using the remove-smbackup cmdlet.
2. Delete the resource group of the cloned resources by using the remove-smresourcegroup cmdlet.
3. Remove the protection of the cloned resource by using the remove-smprotectresource cmdlet.
4. Select the parent resource from the Resources page.

   The resource topology page is displayed.

5. From the Manage Copies view, select the clones either from the primary or secondary (mirrored or replicated) storage systems.
6. Select the clones, and then click to delete clones or click to split the clones.
7. Click **OK**.

Protect Oracle databases

Overview of SnapCenter Plug-in for Oracle Database

What can you do with the Plug-in for Oracle Database

The SnapCenter Plug-in for Oracle Database is a host-side component of the NetApp SnapCenter Software that enables application-aware data protection management of Oracle databases.

The Plug-in for Oracle Database automates the backup, cataloging and uncataloging with Oracle Recovery Manager (RMAN), verification, mounting, unmounting, restore, recovery, and cloning of Oracle databases in your SnapCenter environment.
The Plug-in for Oracle Database installs SnapCenter Plug-in for UNIX to perform all the data protection operations.

You can use the Plug-in for Oracle Database to manage backups of Oracle databases running SAP applications. However, SAP BR*Tools integration is not supported.

- Back up datafiles, control files, and archive log files.

  Backup is supported only at container database (CDB) level.

- Restore and recovery of databases, CDBs, and pluggable databases (PDBs).

  Incomplete recovery of PDBs are not supported.

- Create clones of production databases up to a point-in-time.

  Cloning is supported only at CDB level.

- Verify backups immediately.
- Mount and unmount data and log backups for recovery operation.
- Schedule backup and verification operations.
- Monitor all operations.
- View reports for backup, restore, and clone operations.
- Automates application-aware backup, restore, recovery, verify, mount, unmount, and clone operations for Oracle databases in your SnapCenter environment
- Supports Oracle databases for SAP, however, SAP BR*Tools integration is not provided

Features of Plug-in for Oracle Database

The Plug-in for Oracle Database integrates with the Oracle database on the Linux or AIX host and with NetApp technologies on the storage system.

- Unified graphical user interface

  The SnapCenter interface provides standardization and consistency across plug-ins and environments. The SnapCenter interface enables you to complete consistent backup, restore, recovery, and clone

operations across plug-ins, use centralized reporting, use at-a-glance dashboard views, set up role-based access control (RBAC), and monitor jobs across all plug-ins.

- Automated central administration

  You can schedule backup and clone operations, configure policy-based backup retention, and perform restore operations. You can also proactively monitor your environment by configuring SnapCenter to send email alerts.

- Nondisruptive NetApp Snapshot technology

  SnapCenter uses NetApp Snapshot technology with the Plug-in for Oracle Database and Plug-in for UNIX to back up databases. Snapshots consume minimal storage space.

The Plug-in for Oracle Database also offers the following benefits:

- Support for backup, restore, clone, mount, unmount, and verification workflows
- Automatic discovery of Oracle databases configured on the host
- Support for cataloging and uncataloging using Oracle Recovery Manager (RMAN)
- RBAC-supported security and centralized role delegation

  You can also set the credentials so that the authorized SnapCenter users have application-level permissions.

- Support for Archive Log Management (ALM) for restore and clone operations
- Creation of space-efficient and point-in-time copies of production databases for testing or data extraction by using NetApp FlexClone technology

  A FlexClone license is required on the storage system where you want to create the clone.

- Support for consistency group (CG) feature of ONTAP as part of creating backups in SAN and ASM environments
- Nondisruptive and automated backup verification
- Capability to run multiple backups simultaneously across multiple database hosts

  In a single operation, Snapshots are consolidated when databases in a single host share the same volume.

- Support for physical and virtualized infrastructures
- Support for NFS, iSCSI, Fibre Channel (FC), RDM, VMDK over NFS and VMFS, and ASM over NFS, SAN, RDM, and VMDK
- Support for the Selective LUN Map (SLM) feature of ONTAP

  Enabled by default, the SLM feature periodically discovers the LUNs that do not have optimized paths and fixes them. You can configure SLM by modifying the parameters in the scu.properties file located at /var/opt/snapcenter/scu/etc.

  - You can disable this by setting the value of ENABLE_LUNPATH_MONITORING parameter to false.
  - You can specify the frequency in which the LUN paths will be fixed automatically by assigning the value (in hours) to LUNPATH_MONITORING_INTERVAL parameter.
    For information about SLM, see the ONTAP 9 SAN Administration section.

- Support for non-volatile memory express (NVMe) on Linux
  - NVMe util should be installed on the host.

    You must install NVMe util to clone or mount to alternate host.

  - Backup, restore, clone, mount, unmount, catalog, uncatalog, and verification operations are supported on the NVMe hardware except for the virtualized environments like RDM.

    The above operations are supported on devices without partitions or with single partition.

    > You can configure multipathing solution for NVMe devices by setting the native multipathing option in the kernel. Device Mapper (DM) multipathing is not supported.

  - Backup, restore, clone, mount, unmount, catalog, uncatalog, and verification workflows workflows are supported on NVMe over TCP/IP.
  - Backup, restore, clone, mount, unmount, catalog, uncatalog, and verification workflows are supported on VMDK layout created on NVMe over TCP/IP.
- Supports SnapMirror active sync (initially released as SnapMirror Business Continuity [SM-BC]) that enables business services to continue operating even through a complete site failure, supporting applications to fail over transparently using a secondary copy. Neither manual intervention nor additional scripting is required to trigger a failover with SnapMirror active sync.
- Supports any non-default user instead of oracle and grid.

  To support the non-default users, you should set the non-default users by modifying the values of the parameters in the **sco.properties** file located at *file /var/opt/snapcenter/sco/etc/*.

  The default values of the parameters are set as oracle and grid.

  - DB_USER=oracle
  - DB_GROUP=oinstall
  - GI_USER=grid
  - GI_GROUP=oinstall

Storage types supported by Plug-in for Oracle Database

SnapCenter supports a wide range of storage types on both physical and virtual machines. You must verify the support for your storage type before installing the SnapCenter Plug-ins Package for Linux or SnapCenter Plug-ins Package for AIX.

SnapCenter does not support storage provisioning for Linux and AIX.

Storage types supported on Linux

The following table lists the storage types supported on Linux.

Machine	Storage type
Physical server	• FC-connected LUNs • iSCSI-connected LUNs • NFS-connected volumes • NVMe-FC • NVMe/TCP
VMware ESXi	• RDM LUNs connected by an FC or iSCSI ESXi HBAScanning of host bus adapters (HBAs) might take long time to complete because SnapCenter scans all the host bus adaptors present in the host. You can edit the **LinuxConfig.pm** file located at */opt/NetApp/snapcenter/spl/plugins/scu/scucore/modules/SCU/Config* to set the value of the **SCSI_HOSTS_OPTIMIZED_RESCAN** parameter to 1 to rescan only those HBAs that are listed in HBA_DRIVER_NAMES. • iSCSI LUNs connected directly to the guest system by the iSCSI initiator • VMDKs on NFS datastores • VMDKs on VMFS created over NVMe/TCP RAC is supported on ESX 8.0U2 which has the support for shared VMDK • NFS volumes connected directly to the guest system • vVol datastores on both NFS and SAN vVol datastore can only be provisioned with ONTAP Tools for VMware vSphere.

Storage types supported on AIX

The following table lists the storage types supported on AIX.

Machine	Storage type
Physical server	• FC-connected and iSCSI-connected LUNs. In a SAN environment, ASM, LVM, and SAN file systems are supported. NFS on AIX and filesystem is not supported. • Enhanced Journaled File System (JFS2) Supports inline logging on SAN filesystems and LVM layout.

The NetApp Interoperability Matrix Tool contains the latest information about the supported versions.

Prepare storage systems for SnapMirror and SnapVault replication for Plug-in for Oracle

You can use a SnapCenter plug-in with ONTAP SnapMirror technology to create mirror copies of backup sets on another volume, and with ONTAP SnapVault technology to perform disk-to-disk backup replication for standards compliance and other governance-related purposes. Before you perform these tasks, you must configure a data-protection relationship between the source and destination volumes and initialize the relationship.

SnapCenter performs the updates to SnapMirror and SnapVault after it completes the Snapshot operation. SnapMirror and SnapVault updates are performed as part of the SnapCenter job; do not create a separate ONTAP schedule.

If you are coming to SnapCenter from a NetApp SnapManager product and are satisfied with the data protection relationships you have configured, you can skip this section.

A data protection relationship replicates data on primary storage (the source volume) to secondary storage (the destination volume). When you initialize the relationship, ONTAP transfers the data blocks referenced on the source volume to the destination volume.

SnapCenter does not support cascade relationships between SnapMirror and SnapVault volumes (**Primary** > **Mirror** > **Vault**). You should use fanout relationships.

SnapCenter supports the management of version-flexible SnapMirror relationships. For details about version-flexible SnapMirror relationships and how to set them up, see the ONTAP documentation.

Minimum ONTAP privileges required for Plug-in for Oracle

The minimum ONTAP privileges that are required vary according to the SnapCenter plug-ins you are using for data protection.

- All-access commands: Minimum privileges required for ONTAP 9.12.1 and later

- event generate-autosupport-log
- job history show
- job stop
- lun
- lun attribute show
- lun create
- lun delete
- lun geometry
- lun igroup add
- lun igroup create
- lun igroup delete
- lun igroup rename
- lun igroup show
- lun mapping add-reporting-nodes
- lun mapping create
- lun mapping delete
- lun mapping remove-reporting-nodes
- lun mapping show
- lun modify
- lun move-in-volume
- lun offline
- lun online
- lun persistent-reservation clear
- lun resize
- lun serial
- lun show
- snapmirror policy add-rule
- snapmirror policy modify-rule
- snapmirror policy remove-rule
- snapmirror policy show
- snapmirror restore
- snapmirror show
- snapmirror show-history
- snapmirror update
- snapmirror update-ls-set
- snapmirror list-destinations
- version

- volume clone create
- volume clone show
- volume clone split start
- volume clone split stop
- volume create
- volume destroy
- volume file clone create
- volume file show-disk-usage
- volume offline
- volume online
- volume modify
- volume qtree create
- volume qtree delete
- volume qtree modify
- volume qtree show
- volume restrict
- volume show
- volume snapshot create
- volume snapshot delete
- volume snapshot modify
- volume snapshot rename
- volume snapshot restore
- volume snapshot restore-file
- volume snapshot show
- volume unmount
- vserver
- vserver cifs
- vserver cifs shadowcopy show
- vserver show
- network interface
- network interface show
- metrocluster show

Install SnapCenter Plug-in for Oracle Database

Installation workflow of SnapCenter Plug-in for Oracle Database

You should install and set up the SnapCenter Plug-in for Oracle Database if you want to protect Oracle databases.

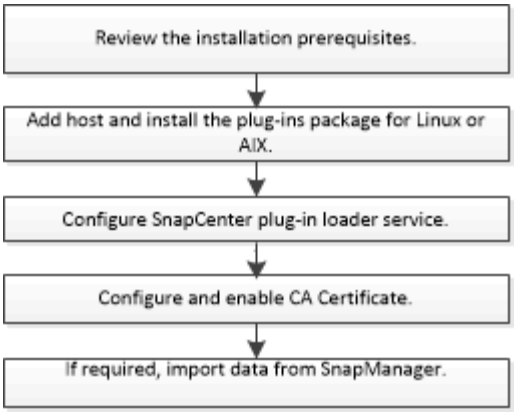

Prerequisites for adding hosts and installing Plug-ins Package for Linux or AIX

Before you add a host and install the plug-ins packages, you must complete all the requirements.

- If you are using iSCSI, the iSCSI service must be running.
- You must have enabled the password-based SSH connection for the root or non-root user.

  SnapCenter Plug-in for Oracle Database can be installed by a non-root user. However, you should configure the sudo privileges for the non-root user to install and start the plug-in process. After installing the plug-in, the processes will be running as an effective non-root user.

- If you are installing the SnapCenter Plug-ins Package for AIX on AIX host, you should have manually resolved the directory level symbolic links.

  The SnapCenter Plug-ins Package for AIX automatically resolves the file level symbolic link but not the directory level symbolic links to obtain the JAVA_HOME absolute path.

- Create credentials with authentication mode as Linux or AIX for the install user.
- You must have installed Java 11 on your Linux or AIX host.
  - Java from Oracle and OpenJDK is supported for Linux
  - IBM Java for AIX. You can download from IBM Semeru Runtimes Downloads

    Ensure that you have installed only the certified edition of JAVA 11 on the Linux host.

- For Oracle databases that are running on a Linux or AIX host, you should install both SnapCenter Plug-in for Oracle Database and SnapCenter Plug-in for UNIX.

  You can use the Plug-in for Oracle Database to manage Oracle databases for SAP as well. However, SAP BR*Tools integration is not supported.

- If you are using Oracle database 11.2.0.3 or later, you must install the 13366202 Oracle patch.

UUID mapping in the /etc/fstab file is not supported by SnapCenter.

- You should have **bash** as the default shell for plug-in installation.

Linux Host requirements

You should ensure that the host meets the requirements before installing the SnapCenter Plug-ins Package for Linux.

Item	Requirements
Operating systems	• Red Hat Enterprise Linux • Oracle Linux If you are using Oracle database on LVM in Oracle Linux or Red Hat Enterprise Linux 6.6 or 7.0 operating systems, you must install the latest version of Logical Volume Manager (LVM). • SUSE Linux Enterprise Server (SLES)
Minimum RAM for the SnapCenter plug-in on host	2 GB
Minimum install and log space for the SnapCenter plug-in on host	2 GB You should allocate sufficient disk space and monitor the storage consumption by the logs folder. The log space required varies depending on the number of the entities to be protected and the frequency of data protection operations. If there is no sufficient disk space, the logs will not be created for the recently run operations.
Required software packages	Java 11 Oracle and OpenJDK Ensure that you have installed only the certified edition of JAVA 11 on the Linux host. If you have upgraded JAVA to the latest version, you must ensure that the JAVA_HOME option located at /var/opt/snapcenter/spl/etc/spl.properties is set to the correct JAVA version and the correct path.

For the latest information about supported versions, see the NetApp Interoperability Matrix Tool.

Configure sudo privileges for non-root users for Linux host

SnapCenter 2.0 and later releases allow a non-root user to install the SnapCenter Plug-ins Package for Linux and to start the plug-in process. The plug-in processes will be running as an effective non-root user. You

should configure sudo privileges for the non-root user to provide access to several paths.

What you will need

- Sudo version 1.8.7 or later.
- If the umask is 0027, ensure that the java folder and all the files inside should have permission of 555. Otherwise plug-in installation might fail.
- For the non-root user, ensure that the name of the non-root user and the user's group should be the same.
- Edit the */etc/ssh/sshd_config* file to configure the message authentication code algorithms: MACs hmac-sha2-256 and MACs hmac-sha2-512.

  Restart the sshd service after updating the configuration file.

  Example:

```
#Port 22
#AddressFamily any
#ListenAddress 0.0.0.0
#ListenAddress ::
#Legacy changes
#KexAlgorithms diffie-hellman-group1-sha1
#Ciphers aes128-cbc
#The default requires explicit activation of protocol
Protocol 2
HostKey/etc/ssh/ssh_host_rsa_key
MACs hmac-sha2-256
```

About this task

You should configure sudo privileges for the non-root user to provide access to the following paths:

- /home/*LINUX_USER*/.sc_netapp/snapcenter_linux_host_plugin.bin
- /custom_location/NetApp/snapcenter/spl/installation/plugins/uninstall
- /custom_location/NetApp/snapcenter/spl/bin/spl

Steps

1. Log in to the Linux host on which you want to install the SnapCenter Plug-ins Package for Linux.
2. Add the following lines to the /etc/sudoers file by using the visudo Linux utility.

```
Cmnd_Alias HPPLCMD = sha224:checksum_value== /home/
LINUX_USER/.sc_netapp/snapcenter_linux_host_plugin.bin,
/opt/NetApp/snapcenter/spl/installation/plugins/uninstall,
/opt/NetApp/snapcenter/spl/bin/spl, /opt/NetApp/snapcenter/scc/bin/scc
Cmnd_Alias PRECHECKCMD = sha224:checksum_value== /home/
LINUX_USER/.sc_netapp/Linux_Prechecks.sh
Cmnd_Alias CONFIGCHECKCMD = sha224:checksum_value==
/opt/NetApp/snapcenter/spl/plugins/scu/scucore/configurationcheck/Config
_Check.sh
Cmnd_Alias SCCMD = sha224:checksum_value==
/opt/NetApp/snapcenter/spl/bin/sc_command_executor
Cmnd_Alias SCCCMDEXECUTOR =checksum_value==
/opt/NetApp/snapcenter/scc/bin/sccCommandExecutor
LINUX_USER ALL=(ALL) NOPASSWD:SETENV: HPPLCMD, PRECHECKCMD,
CONFIGCHECKCMD, SCCCMDEXECUTOR, SCCMD
Defaults: LINUX_USER env_keep += "IATEMPDIR"
Defaults: LINUX_USER env_keep += "JAVA_HOME"
Defaults: LINUX_USER !visiblepw
Defaults: LINUX_USER !requiretty
```

If you are having a RAC setup, along with the other allowed commands, you should add the following to the /etc/sudoers file: '/<crs_home>/bin/olsnodes'

You can obtain the value of *crs_home* from the */etc/oracle/olr.loc* file.

LINUX_USER is the name of the non-root user that you created.

You can obtain the *checksum_value* from the **sc_unix_plugins_checksum.txt** file, which is located at:

- *C:\ProgramData\NetApp\SnapCenter\Package Repository\sc_unix_plugins_checksum.txt* if SnapCenter Server is installed on Windows host.
- */opt/NetApp/snapcenter/SnapManagerWeb/Repository/sc_unix_plugins_checksum.txt* if SnapCenter Server in installed on Linux host.

The example should be used only as a reference for creating your own data.

AIX Host requirements

You should ensure that the host meets the requirements before installing the SnapCenter Plug-ins Package for AIX.

SnapCenter Plug-in for UNIX which is part of the SnapCenter Plug-ins Package for AIX, does not support concurrent volume groups.

Item	Requirements
Operating systems	AIX 7.1 or later
Minimum RAM for the SnapCenter plug-in on host	4 GB
Minimum install and log space for the SnapCenter plug-in on host	2 GB You should allocate sufficient disk space and monitor the storage consumption by the logs folder. The log space required varies depending on the number of the entities to be protected and the frequency of data protection operations. If there is no sufficient disk space, the logs will not be created for the recently run operations.
Required software packages	Java 11 IBM Java If you have upgraded JAVA to the latest version, you must ensure that the JAVA_HOME option located at /var/opt/snapcenter/spl/etc/spl.properties is set to the correct JAVA version and the correct path.

For the latest information about supported versions, see the NetApp Interoperability Matrix Tool.

Configure sudo privileges for non-root users for AIX host

SnapCenter 4.4 and later allows a non-root user to install the SnapCenter Plug-ins Package for AIX and to start the plug-in process. The plug-in processes will be running as an effective non-root user. You should configure sudo privileges for the non-root user to provide access to several paths.

What you will need

- Sudo version 1.8.7 or later.
- If the umask is 0027, ensure that the java folder and all the files inside should have permission of 555. Otherwise plug-in installation might fail.
- Edit the */etc/ssh/sshd_config* file to configure the message authentication code algorithms: MACs hmac-sha2-256 and MACs hmac-sha2-512.

  Restart the sshd service after updating the configuration file.

  Example:

```
#Port 22
#AddressFamily any
#ListenAddress 0.0.0.0
#ListenAddress ::
#Legacy changes
#KexAlgorithms diffie-hellman-group1-sha1
#Ciphers aes128-cbc
#The default requires explicit activation of protocol
Protocol 2
HostKey/etc/ssh/ssh_host_rsa_key
MACs hmac-sha2-256
```

About this task

You should configure sudo privileges for the non-root user to provide access to the following paths:

- /home/*AIX_USER*/.sc_netapp/snapcenter_aix_host_plugin.bsx
- /custom_location/NetApp/snapcenter/spl/installation/plugins/uninstall
- /custom_location/NetApp/snapcenter/spl/bin/spl

Steps

1. Log in to the AIX host on which you want to install the SnapCenter Plug-ins Package for AIX.
2. Add the following lines to the /etc/sudoers file by using the visudo Linux utility.

```
Cmnd_Alias HPPACMD = sha224:checksum_value== /home/
AIX_USER/.sc_netapp/snapcenter_aix_host_plugin.bsx,
/opt/NetApp/snapcenter/spl/installation/plugins/uninstall,
/opt/NetApp/snapcenter/spl/bin/spl
Cmnd_Alias PRECHECKCMD = sha224:checksum_value== /home/
AIX_USER/.sc_netapp/AIX_Prechecks.sh
Cmnd_Alias CONFIGCHECKCMD = sha224:checksum_value==
/opt/NetApp/snapcenter/spl/plugins/scu/scucore/configurationcheck/Config
_Check.sh
Cmnd_Alias SCCMD = sha224:checksum_value==
/opt/NetApp/snapcenter/spl/bin/sc_command_executor
AIX_USER ALL=(ALL) NOPASSWD:SETENV: HPPACMD, PRECHECKCMD,
CONFIGCHECKCMD, SCCMD
Defaults: LINUX_USER env_keep += "IATEMPDIR"
Defaults: LINUX_USER env_keep += "JAVA_HOME"
Defaults: AIX_USER !visiblepw
Defaults: AIX_USER !requiretty
```

If you are having a RAC setup, along with the other allowed commands, you should add the following to the /etc/sudoers file: '/<crs_home>/bin/olsnodes'

You can obtain the value of *crs_home* from the */etc/oracle/olr.loc* file.

AIX_USER is the name of the non-root user that you created.

You can obtain the *checksum_value* from the **sc_unix_plugins_checksum.txt** file, which is located at:

- *C:\ProgramData\NetApp\SnapCenter\Package Repository\sc_unix_plugins_checksum.txt* if SnapCenter Server is installed on Windows host.
- */opt/NetApp/snapcenter/SnapManagerWeb/Repository/sc_unix_plugins_checksum.txt* if SnapCenter Server in installed on Linux host.

The example should be used only as a reference for creating your own data.

Set up credentials

SnapCenter uses credentials to authenticate users for SnapCenter operations. You should create credentials for installing the plug-in package on Linux or AIX hosts.

About this task

The credentials are created either for the root user or for a non-root user who has sudo privileges to install and start the plug-in process.

For information, see: Configure sudo privileges for non-root users for Linux host or Configure sudo privileges for non-root users for AIX host

Best Practice: Although you are allowed to create credentials after deploying hosts and installing plug-ins, the best practice is to create credentials after you add SVMs, before you deploy hosts and install plug-ins.

Steps

1. In the left navigation pane, click **Settings**.
2. In the Settings page, click **Credential**.
3. Click **New**.
4. In the Credential page, enter the credential information:

For this field…	Do this…
Credential name	Enter a name for the credentials.

For this field…	Do this…
User name/Password	Enter the user name and password that are to be used for authentication. • Domain administrator Specify the domain administrator on the system on which you are installing the SnapCenter plug-in. Valid formats for the Username field are: ◦ *NetBIOS\UserName* ◦ *Domain FQDN\UserName* • Local administrator (for workgroups only) For systems that belong to a workgroup, specify the built-in local administrator on the system on which you are installing the SnapCenter plug-in. You can specify a local user account that belongs to the local administrators group if the user account has elevated privileges or the User Access control feature is disabled on the host system. The valid format for the Username field is: *UserName*
Authentication Mode	Select the authentication mode that you want to use. Depending on the operating system of the plug-in host, select either Linux or AIX.
Use sudo privileges	Select the **Use sudo privileges** check box if you are creating credentials for a non-root user.

5. Click **OK**.

After you finish setting up the credentials, you might want to assign credential maintenance to a user or group of users on the **User and Access** page.

Configure credentials for an Oracle database

You must configure credentials that are used to perform data protection operations on Oracle databases.

About this task

You should review the different authentication methods supported for Oracle database. For information, see Authentication methods for your credentials.

If you set up credentials for individual resource groups and the user name does not have full admin privileges, the user name must at least have resource group and backup privileges.

If you have enabled Oracle database authentication, a red padlock icon is shown in the resources view. You must configure database credentials to be able to protect the database or add it to the resource group to perform data protection operations.

If you specify incorrect details while creating a credential, an error message is displayed. You must click **Cancel**, and then retry.

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page, select **Database** from the **View** list.
3. Click , and then select the host name and the database type to filter the resources.

   You can then click to close the filter pane.
4. Select the database, and then click **Database Settings** > **Configure Database**.
5. In the Configure database settings section, from the **Use existing Credential** drop-down list, select the credential that should be used to perform data protection jobs on the Oracle database.

   The Oracle user should have sysdba privileges.

   You can also create a credential by clicking .
6. In the Configure ASM settings section, from the **Use existing Credential** drop-down list, select the credential that should be used to perform data protection jobs on the ASM instance.

   The ASM user should have sysasm privilege.

   You can also create a credential by clicking .
7. In the Configure RMAN catalog settings section, from the **Use existing credential** drop-down list, select the credential that should be used to perform data protection jobs on the Oracle Recovery Manager (RMAN) catalog database.

   You can also create a credential by clicking .

   In the **TNSName** field, enter the Transparent Network Substrate (TNS) file name that will be used by the SnapCenter Server to communicate with the database.
8. In the **Preferred RAC Nodes** field, specify the Real Application Cluster (RAC) nodes preferred for backup.

   The preferred nodes might be one or all cluster nodes where the RAC database instances are present. The backup operation is triggered only on these preferred nodes in the order of preference.

   In RAC One Node, only one node is listed in the preferred nodes, and this preferred node is the node where the database is currently hosted.

   After failover or relocation of RAC One Node database, refreshing of resources in the SnapCenter Resources page will remove the host from the **Preferred RAC Nodes** list where the database was earlier hosted. The RAC node where the database is relocated will be listed in **RAC Nodes** and will need to be manually configured as the preferred RAC node.

   For more information, see Preferred nodes in RAC setup.

9. Click **OK**.

Add hosts and install Plug-ins Package for Linux or AIX using GUI

You can use the Add Host page to add hosts, and then install the SnapCenter Plug-ins Package for Linux or SnapCenter Plug-ins Package for AIX. The plug-ins are automatically installed on the remote hosts.

About this task

You can add a host and install plug-in packages either for an individual host or for a cluster. If you are installing the plug-in on a cluster (Oracle RAC), the plug-in is installed on all of the nodes of the cluster. For Oracle RAC One Node, you should install the plug-in on both active and passive nodes.

Only password-based authentication is supported when you are installing the plug-in on an Oracle RAC. The SSH key based authentication is not supported.

You should be assigned to a role that has the plug-in install and uninstall permissions, such as the SnapCenter Admin role.

You cannot add a SnapCenter Server as a plug-in host to another SnapCenter Server.

Steps

1. In the left navigation pane, click **Hosts**.
2. Verify that the **Managed Hosts** tab is selected at the top.
3. Click **Add**.
4. In the Hosts page, perform the following actions:

For this field…	Do this…
Host Type	Select **Linux** or **AIX** as the host type. The SnapCenter Server adds the host, and then installs the Plug-in for Oracle Database and the Plug-in for UNIX if the plug-ins are not already installed on the host.

For this field…	Do this…
Host name	Enter the fully qualified domain name (FQDN) or the IP address of the host. SnapCenter depends on the proper configuration of the DNS. Therefore, the best practice is to enter the FQDN. You can enter the IP addresses or FQDN of one of the following: • Stand-alone host • Any node in the Oracle Real Application Clusters (RAC) environment Node VIP or scan IP is not supported If you are adding a host by using SnapCenter and the host is part of a subdomain, you must provide the FQDN.
Credentials	Either select the credential name that you created or create new credentials. The credential must have administrative rights on the remote host. For details, see the information about creating credentials. You can view details about the credentials by positioning the cursor over the credential name that you specified. The credentials authentication mode is determined by the host type that you specify in the Add Host wizard.

5. In the Select Plug-ins to Install section, select the plug-ins to install.
6. (Optional) Click **More Options**.

For this field…	Do this…
Port	Either retain the default port number or specify the port number. The default port number is 8145. If the SnapCenter Server was installed on a custom port, that port number will be displayed as the default port. If you manually installed the plug-ins and specified a custom port, you must specify the same port. Otherwise, the operation fails.
Installation Path	The default path is */opt/NetApp/snapcenter*. You can optionally customize the path.
Add all hosts in the Oracle RAC	Select this check box to add all the cluster nodes in an Oracle RAC. In a Flex ASM setup, all the nodes irrespective of whether it is a Hub or Leaf node, will be added.
Skip optional preinstall checks	Select this check box if you have already installed the plug-ins manually and you do not want to validate whether the host meets the requirements for installing the plug-in.

7. Click **Submit**.

   If you have not selected the Skip prechecks checkbox, the host is validated to verify whether the host meets the requirements for installing the plug-in.

   The precheck script does not validate the plug-in port firewall status if it is specified in the firewall reject rules.

   Appropriate error or warning messages are displayed if the minimum requirements are not met. If the error is related to disk space or RAM, you can update the web.config file located at *C:\Program Files\NetApp\SnapCenter WebApp* to modify the default values. If the error is related to other parameters, you should fix the issue.

   In an HA setup, if you are updating web.config file, you must update the file on both nodes.

8. Verify the fingerprint, and then click **Confirm and Submit**.

   In a cluster setup, you should verify the fingerprint of each of the nodes in the cluster.

   SnapCenter does not support ECDSA algorithm.

Fingerprint verification is mandatory even if the same host was added earlier to SnapCenter and the fingerprint was confirmed.

9. Monitor the installation progress.

   The installation-specific log files are located at */custom_location/snapcenter/logs*.

Result

All the databases on the host are automatically discovered and displayed in the Resources page. If nothing is displayed, click **Refresh Resources**.

Monitor installation status

You can monitor the progress of SnapCenter plug-in package installation by using the Jobs page. You might want to check the progress of installation to determine when it is complete or if there is an issue.

About this task

The following icons appear on the Jobs page and indicate the state of the operation:

- In progress
- Completed successfully
- Failed
- Completed with warnings or could not start due to warnings
- Queued

Steps

1. In the left navigation pane, click **Monitor**.
2. In the **Monitor** page, click **Jobs**.
3. In the **Jobs** page, to filter the list so that only plug-in installation operations are listed, do the following:
   a. Click **Filter**.
   b. Optional: Specify the start and end date.
   c. From the Type drop-down menu, select **Plug-in installation**.
   d. From the Status drop-down menu, select the installation status.
   e. Click **Apply**.
4. Select the installation job and click **Details** to view the job details.
5. In the **Job Details** page, click **View logs**.

Alternate ways to install Plug-ins Package for Linux or AIX

You can also install the Plug-ins Package for Linux or AIX manually by either using the cmdlets or CLIs.

Before installing the plug-in manually, you should validate the signature of the binary package by using the key **snapcenter_public_key.pub** and **snapcenter_linux_host_plugin.bin.sig** located at

C:\ProgramData\NetApp\SnapCenter\Package Repository.

Ensure that **OpenSSL 1.0.2g** is installed on the host where you want to install the plug-in.

Validate the signature of the binary package by running the command:

- For Linux host: `openssl dgst -sha256 -verify snapcenter_public_key.pub -signature snapcenter_linux_host_plugin.bin.sig snapcenter_linux_host_plugin.bin`
- For AIX host: `openssl dgst -sha256 -verify snapcenter_public_key.pub -signature snapcenter_aix_host_plugin.bsx.sig snapcenter_aix_host_plugin.bsx`

Install on multiple remote hosts using cmdlets

You should use the *Install-SmHostPackage* PowerShell cmdlet to install the SnapCenter Plug-ins Package for Linux or SnapCenter Plug-ins Package for AIX on multiple hosts.

What you will need

You should be logged in to SnapCenter as a domain user with local administrator rights on each host on which you want to install the plug-in package.

Steps

1. Launch PowerShell.
2. On the SnapCenter Server host, establish a session using the *Open-SmConnection* cmdlet, and then enter your credentials.
3. Install the SnapCenter Plug-ins Package for Linux or SnapCenter Plug-ins Package for AIX using the *Install-SmHostPackage* cmdlet and the required parameters.

   You can use the *-skipprecheck* option when you have already installed the plug-ins manually and do not want to validate whether the host meets the requirements to install the plug-in.

   The precheck script does not validate the plug-in port firewall status if it is specified in the firewall reject rules.

4. Enter your credentials for remote installation.

The information regarding the parameters that can be used with the cmdlet and their descriptions can be obtained by running *Get-Help command_name*. Alternatively, you can also refer to the SnapCenter Software Cmdlet Reference Guide.

Install on cluster host

You should install SnapCenter Plug-ins Package for Linux or SnapCenter Plug-ins Package for AIX on both the nodes of the cluster host.

Each of the nodes of the cluster host has two IPs. One of the IPs will be the public IP of the respective nodes and the second IP will be the cluster IP that is shared between both the nodes.

Steps

1. Install SnapCenter Plug-ins Package for Linux or SnapCenter Plug-ins Package for AIX on both the nodes

of the cluster host.

2. Validate that the correct values for SNAPCENTER_SERVER_HOST, SPL_PORT, SNAPCENTER_SERVER_PORT, and SPL_ENABLED_PLUGINS parameters are specified in the spl.properties file located at */var/opt/snapcenter/spl/etc/*.

   If SPL_ENABLED_PLUGINS is not specified in spl.properties, you can add it and assign the value SCO,SCU.

3. On the SnapCenter Server host, establish a session using the *Open-SmConnection* cmdlet, and then enter your credentials.
4. In each of the nodes, set the preferred IPs of the node using the *Set-PreferredHostIPsInStorageExportPolicy* sccli command and the required parameters.
5. In the SnapCenter Server host, add an entry for the cluster IP and corresponding DNS name in *C:\Windows\System32\drivers\etc\hosts*.
6. Add the node to the SnapCenter Server using the *Add-SmHost* cmdlet by specifying the cluster IP for the host name.

Discover the Oracle database on node 1 (assuming the cluster IP is hosted on node 1) and create a backup of the database. If a failover happens, you can use the backup created on node 1 to restore the database on node 2. You can also use the backup created on node 1 to create a clone on node 2.

There will be stale volumes, directories, and lock file if the failover happens while any other SnapCenter operations are running.

Install Plug-ins Package for Linux in silent mode

You can install the SnapCenter Plug-ins Package for Linux in silent mode by using the command-line interface (CLI).

What you will need

- You should review the prerequisites for installing the plug-ins package.
- You should ensure that the DISPLAY environment variable is not set.

  If the DISPLAY environment variable is set, you should run unset DISPLAY, and then try to manually install the plug-in.

About this task

You are required to provide the necessary installation information while installing in console mode, whereas in silent mode installation you do not have to provide any installation information.

Steps

1. Download the SnapCenter Plug-ins Package for Linux from the SnapCenter Server installation location.

   The default installation path is *C:\ProgramData\NetApp\SnapCenter\PackageRepository*. This path is accessible from the host where the SnapCenter Server is installed.

2. From the command prompt, navigate to the directory where you downloaded the installation file.
3. Run

```
./SnapCenter_linux_host_plugin.bin-i silent-DPORT=8145-
DSERVER_IP=SnapCenter_Server_FQDN-DSERVER_HTTPS_PORT=SnapCenter_Server_Port-
DUSER_INSTALL_DIR==/opt/custom_path
```

4. Edit the spl.properties file located at */var/opt/snapcenter/spl/etc/* to add SPL_ENABLED_PLUGINS=SCO,SCU, and then restart the SnapCenter Plug-in Loader service.

> The installation of the plug-ins package registers the plug-ins on the host and not on the SnapCenter Server. You should register the plug-ins on the SnapCenter Server by adding the host using the SnapCenter GUI or PowerShell cmdlet. While adding the host, select “None” as the credential. After the host is added, the installed plug-ins are automatically discovered.

Install Plug-ins Package for AIX in silent mode

You can install the SnapCenter Plug-ins Package for AIX in silent mode by using the command-line interface (CLI).

What you will need

- You should review the prerequisites for installing the plug-ins package.
- You should ensure that the DISPLAY environment variable is not set.

  If the DISPLAY environment variable is set, you should run unset DISPLAY, and then try to manually install the plug-in.

Steps

1. Download the SnapCenter Plug-ins Package for AIX from the SnapCenter Server installation location.

   The default installation path is *C:\ProgramData\NetApp\SnapCenter\PackageRepository*. This path is accessible from the host where the SnapCenter Server is installed.

2. From the command prompt, navigate to the directory where you downloaded the installation file.
3. Run

```
./snapcenter_aix_host_plugin.bsx-i silent-DPORT=8145-
DSERVER_IP=SnapCenter_Server_FQDN-DSERVER_HTTPS_PORT=SnapCenter_Server_Port-
DUSER_INSTALL_DIR==/opt/custom_path-
DINSTALL_LOG_NAME=SnapCenter_AIX_Host_Plug-in_Install_MANUAL.log-
DCHOSEN_FEATURE_LIST=CUSTOMDSPL_USER=install_user
```

4. Edit the spl.properties file located at */var/opt/snapcenter/spl/etc/* to add SPL_ENABLED_PLUGINS=SCO,SCU, and then restart the SnapCenter Plug-in Loader service.

> The installation of the plug-ins package registers the plug-ins on the host and not on the SnapCenter Server. You should register the plug-ins on the SnapCenter Server by adding the host using the SnapCenter GUI or PowerShell cmdlet. While adding the host, select “None” as the credential. After the host is added, the installed plug-ins are automatically discovered.

Configure the SnapCenter Plug-in Loader service

The SnapCenter Plug-in Loader service loads the plug-in package for Linux or AIX to

interact with the SnapCenter Server. The SnapCenter Plug-in Loader service is installed when you install the SnapCenter Plug-ins Package for Linux or SnapCenter Plug-ins Package for AIX.

About this task

After installing the SnapCenter Plug-ins Package for Linux or SnapCenter Plug-ins Package for AIX, the SnapCenter Plug-in Loader service starts automatically. If the SnapCenter Plug-in Loader service fails to start automatically, you should:

- Ensure that the directory where the plug-in is operating is not deleted
- Increase the memory space allotted to the Java Virtual Machine

The spl.properties file, which is located at */custom_location/NetApp/snapcenter/spl/etc/*, contains the following parameters. Default values are assigned to these parameters.

Parameter name	Description
LOG_LEVEL	Displays the log levels that are supported. The possible values are TRACE, DEBUG, INFO, WARN, ERROR, and FATAL.
SPL_PROTOCOL	Displays the protocol that is supported by SnapCenter Plug-in Loader. Only the HTTPS protocol is supported. You can add the value if the default value is missing.
SNAPCENTER_SERVER_PROTOCOL	Displays the protocol that is supported by SnapCenter Server. Only the HTTPS protocol is supported. You can add the value if the default value is missing.
SKIP_JAVAHOME_UPDATE	By default, the SPL service detects the java path and update JAVA_HOME parameter. Therefore the default value is set to FALSE. You can set to TRUE if you want to disable the default behavior and manually fix the java path.
SPL_KEYSTORE_PASS	Displays the password of the keystore file. You can change this value only if you change the password or create a new keystore file.

Parameter name	Description
SPL_PORT	Displays the port number on which the SnapCenter Plug-in Loader service is running. You can add the value if the default value is missing. You should not change the value after installing the plug-ins.
SNAPCENTER_SERVER_HOST	Displays the IP address or host name of the SnapCenter Server.
SPL_KEYSTORE_PATH	Displays the absolute path of the keystore file.
SNAPCENTER_SERVER_PORT	Displays the port number on which the SnapCenter Server is running.
LOGS_MAX_COUNT	Displays the number of SnapCenter Plug-in Loader log files that are retained in the */custom_location/snapcenter/spl/logs* folder. The default value is set to 5000. If the count is more than the specified value, then the last 5000 modified files are retained. The check for the number of files is done automatically every 24 hours from when SnapCenter Plug-in Loader service is started. If you manually delete the spl.properties file, then the number of files to be retained is set to 9999.
JAVA_HOME	Displays the absolute directory path of the JAVA_HOME which is used to start SPL service. This path is determined during installation and as part of starting SPL.
LOG_MAX_SIZE	Displays the maximum size of the job log file. Once the maximum size is reached, the log file is zipped, and the logs are written into the new file of that job.
RETAIN_LOGS_OF_LAST_DAYS	Displays the number of days up to which the logs are retained.

Parameter name	Description
ENABLE_CERTIFICATE_VALIDATION	Displays true when CA certificate validation is enabled for the host. You can enable or disable this parameter either by editing the spl.properties or by using the SnapCenter GUI or cmdlet.

If any of these parameters are not assigned to the default value or if you want to assign or change the value, then you can modify the spl.properties file. You can also verify the spl.properties file and edit the file to troubleshoot any issues related to the values that are assigned to the parameters. After you modify the spl.properties file, you should restart the SnapCenter Plug-in Loader service.

Steps

1. Perform one of the following actions, as required:
   - Start the SnapCenter Plug-in Loader service:
     - As a root user, run: `/custom_location/NetApp/snapcenter/spl/bin/spl start`
     - As a non-root user, run: `sudo /custom_location/NetApp/snapcenter/spl/bin/spl start`
   - Stop the SnapCenter Plug-in Loader service:
     - As a root user, run: `/custom_location/NetApp/snapcenter/spl/bin/spl stop`
     - As a non-root user, run: `sudo /custom_location/NetApp/snapcenter/spl/bin/spl stop`

     > You can use the -force option with the stop command to stop the SnapCenter Plug-in Loader service forcefully. However, you should use caution before doing so because it also terminates the existing operations.

   - Restart the SnapCenter Plug-in Loader service:
     - As a root user, run: `/custom_location/NetApp/snapcenter/spl/bin/spl restart`
     - As a non-root user, run: `sudo /custom_location/NetApp/snapcenter/spl/bin/spl restart`
   - Find the status of the SnapCenter Plug-in Loader service:
     - As a root user, run: `/custom_location/NetApp/snapcenter/spl/bin/spl status`
     - As a non root user, run: `sudo /custom_location/NetApp/snapcenter/spl/bin/spl status`
   - Find the change in the SnapCenter Plug-in Loader service:
     - As a root user, run: `/custom_location/NetApp/snapcenter/spl/bin/spl change`
     - As a non-root user, run: `sudo /custom_location/NetApp/snapcenter/spl/bin/spl change`

Configure CA certificate with SnapCenter Plug-in Loader (SPL) service on Linux host

You should manage the password of SPL keystore and its certificate, configure the CA certificate, configure root or intermediate certificates to SPL trust-store, and configure CA signed key pair to SPL trust-store with SnapCenter Plug-in Loader service to activate the installed digital certificate.

SPL uses the file 'keystore.jks', which is located at '/var/opt/snapcenter/spl/etc' both as its trust-store and key-store.

Manage password for SPL keystore and alias of the CA signed key pair in use

Steps

1. You can retrieve SPL keystore default password from SPL property file.

   It is the value corresponding to the key 'SPL_KEYSTORE_PASS'.

2. Change the keystore password:

   ```
   keytool -storepasswd -keystore keystore.jks
   ```

3. Change the password for all aliases of private key entries in the keystore to the same password used for the keystore:

   ```
   keytool -keypasswd -alias "<alias_name>" -keystore keystore.jks
   ```

   Update the same for the key SPL_KEYSTORE_PASS in spl.properties file.

4. Restart the service after changing the password.

Password for SPL keystore and for all the associated alias password of the private key should be same.

Configure root or intermediate certificates to SPL trust-store

You should configure the root or intermediate certificates without the private key to SPL trust-store.

Steps

1. Navigate to the folder containing the SPL keystore: */var/opt/snapcenter/spl/etc*.
2. Locate the file 'keystore.jks'.
3. List the added certificates in the keystore:

   ```
   keytool -list -v -keystore keystore.jks
   ```

4. Add a root or intermediate certificate:

```
keytool -import -trustcacerts -alias <AliasNameForCerticateToBeImported>
-file /<CertificatePath> -keystore keystore.jks
```

5. Restart the service after configuring the root or intermediate certificates to SPL trust-store.

You should add the root CA certificate and then the intermediate CA certificates.

Configure CA signed key pair to SPL trust-store

You should configure the CA signed key pair to the SPL trust-store.

Steps

1. Navigate to the folder containing the SPL's keystore /var/opt/snapcenter/spl/etc.
2. Locate the file 'keystore.jks'.
3. List the added certificates in the keystore:

```
keytool -list -v -keystore keystore.jks
```

4. Add the CA certificate having both private and public key.

```
keytool -importkeystore -srckeystore <CertificatePathToImport>
-srcstoretype pkcs12 -destkeystore keystore.jks -deststoretype JKS
```

5. List the added certificates in the keystore.

```
keytool -list -v -keystore keystore.jks
```

6. Verify that the keystore contains the alias corresponding to the new CA certificate, which was added to the keystore.
7. Change the added private key password for CA certificate to the keystore password.

   Default SPL keystore password is the value of the key SPL_KEYSTORE_PASS in spl.properties file.

```
keytool -keypasswd -alias "<aliasNameOfAddedCertInKeystore>" -keystore
keystore.jks
```

8. If the alias name in the CA certificate is long and contains space or special characters ("*",","), change the alias name to a simple name:

```
keytool -changealias -alias "<OrignalAliasName>" -destalias
"<NewAliasName>" -keystore keystore.jks
```

9. Configure the alias name from the keystore located in spl.properties file.

   Update this value against the key SPL_CERTIFICATE_ALIAS.

10. Restart the service after configuring the CA signed key pair to SPL trust-store.

Configure certificate revocation list (CRL) for SPL

You should configure the CRL for SPL

About this task

- SPL will look for the CRL files in a pre-configured directory.
- Default directory for the CRL files for SPL is */var/opt/snapcenter/spl/etc/crl*.

Steps

1. You can modify and update the default directory in spl.properties file against the key SPL_CRL_PATH.
2. You can place more than one CRL file in this directory.

   The incoming certificates will be verified against each CRL.

Enable CA Certificates for plug-ins

You should configure the CA certificates and deploy the CA certificates in the SnapCenter Server and the corresponding plug-in hosts. You should enable the CA certificate validation for the plug-ins.

Before you begin

- You can enable or disable the CA certificates using the run *Set-SmCertificateSettings* cmdlet.
- You can display the certificate status for the plug-ins using the *Get-SmCertificateSettings*.

The information regarding the parameters that can be used with the cmdlet and their descriptions can be obtained by running *Get-Help command_name*. Alternatively, you can also refer to the SnapCenter Software Cmdlet Reference Guide.

Steps

1. In the left navigation pane, click **Hosts**.
2. In the Hosts page, click **Managed Hosts**.
3. Select single or multiple plug-in hosts.
4. Click **More options**.
5. Select **Enable Certificate Validation**.

After you finish

The Managed Hosts tab host displays a padlock and the color of the padlock indicates the status of the

connection between SnapCenter Server and the plug-in host.

- indicates that the CA certificate is neither enabled nor assigned to the plug-in host.
- indicates that the CA certificate is successfully validated.
- indicates that the CA certificate could not be validated.
- indicates that the connection information could not be retrieved.

When the status is yellow or green, the data protection operations completes successfully.

Import data from SnapManager for Oracle and SnapManager for SAP to SnapCenter

Importing data from SnapManager for Oracle and SnapManager for SAP to SnapCenter enables you to continue to use your data from previous versions.

You can import data from SnapManager for Oracle and SnapManager for SAP to SnapCenter by running the import tool from the command-line interface (Linux host CLI).

The import tool creates policies and resource groups in SnapCenter. The policies and resource groups created in SnapCenter correspond to the profiles and operations performed using those profiles in SnapManager for Oracle and SnapManager for SAP. The SnapCenter import tool interacts with the SnapManager for Oracle and SnapManager for SAP repository databases and the database that you want to import.

- Retrieves all the profiles, schedules, and operations performed using the profiles.
- Creates a SnapCenter backup policy for each unique operation and each schedule attached to a profile.
- Creates a resource group for each target database.

You can run the import tool by executing the sc-migrate script located at */opt/NetApp/snapcenter/spl/bin*. When you install the SnapCenter Plug-ins Package for Linux on the database host that you want to import, the sc-migrate script is copied to */opt/NetApp/snapcenter/spl/bin*.

Importing data is not supported from SnapCenter graphical user interface (GUI).

SnapCenter does not support Data ONTAP operating in 7-Mode. You can use the 7-Mode Transition Tool to migrate data and configurations that are stored on a system running Data ONTAP operating in 7-Mode to an ONTAP system.

Configurations supported for importing data

Before you import data from SnapManager 3.4.x for Oracle and SnapManager 3.4.x for SAP to SnapCenter, you should be aware of the configurations that are supported with the SnapCenter Plug-in for Oracle Database.

The configurations that are supported with the SnapCenter Plug-in for Oracle Database are listed in the NetApp Interoperability Matrix Tool.

What gets imported to SnapCenter

You can import profiles, schedules, and operations performed using the profiles.

From SnapManager for Oracle and SnapManager for SAP	To SnapCenter
Profiles without any operations and schedules	A policy is created with default backup type as Online and backup scope as Full.
Profiles with one or more operations	Multiple policies are created based on a unique combination of a profile and operations performed using that profile. The policies created in SnapCenter contain the archive log pruning and retention details retrieved from the profile and corresponding operations.
Profiles with Oracle Recovery Manager (RMAN) configuration	Policies are created with the **Catalog backup with Oracle Recovery Manager** option enabled. If external RMAN cataloging was used in SnapManager, you must configure the RMAN catalog settings in SnapCenter. You can either select the existing credential or create a new credential. If RMAN was configured through control file in SnapManager, then you do not have to configure RMAN in SnapCenter.
Schedule attached to a profile	A policy is created just for the schedule.
Database	A resource group is created for each database that is imported. In a Real Application Clusters (RAC) setup, the node on which you run the import tool becomes the preferred node after importing and the resource group is created for that node.

When a profile is imported, a verification policy is created along with the backup policy.

When SnapManager for Oracle and SnapManager for SAP profiles, schedules, and any operations performed using the profiles are imported to SnapCenter, the different parameters values are also imported.

SnapManager for Oracle and SnapManager for SAP parameter and values	SnapCenter parameter and values	Notes
Backup Scope • Full • Data • Log	Backup Scope • Full • Data • Log	

SnapManager for Oracle and SnapManager for SAP parameter and values	SnapCenter parameter and values	Notes
Backup Mode • Auto • Online • Offline	Backup Type • Online • Offline Shutdown	If the backup mode is Auto, then the import tool checks the database state when the operation was performed, and appropriately sets the backup type as either Online or Offline Shutdown.
Retention • Days • Counts	Retention • Days • Counts	SnapManager for Oracle and SnapManager for SAP uses both Days and Counts to set the retention. In SnapCenter, there is either Days *OR* Counts. So, the retention is set with respect to days as the days get preference over counts in SnapManager for Oracle and SnapManager for SAP.
Pruning for Schedules • All • system change number (SCN) • Date • Logs created before specified hours, days, weeks, and months	Pruning for Schedules • All • Logs created before specified hours and days	SnapCenter does not support pruning based on SCN, Date, weeks, and months.
Notification • Emails sent only for successful operations • Emails sent only for failed operations • Emails sent for both success and failed operations	Notification • Always • On failure • Warning • Error	The email notifications are imported. However, you must manually update the SMTP server using the SnapCenter GUI. The subject of the email is left blank for you to configure.

What does not get imported to SnapCenter

The import tool does not import everything to SnapCenter.

You cannot import the following to SnapCenter:

- Backup metadata
- Partial backups

- Raw device mapping (RDM) and Virtual Storage Console (VSC) related backups
- Roles or any credentials available in the SnapManager for Oracle and SnapManager for SAP repository
- Data related to verification, restore, and clone operations
- Pruning for operations
- Replication details specified in the SnapManager for Oracle and SnapManager for SAP profile

  After importing, you must manually edit the corresponding policy created in SnapCenter to include the replication details.

- Cataloged backup information

Prepare to import data

Before you import data to SnapCenter, you must perform certain tasks to run the import operation successfully.

Steps

1. Identify the database that you want to import.
2. Using SnapCenter, add the database host and install SnapCenter Plug-ins Package for Linux.
3. Using SnapCenter, set up the connections for the storage virtual machines (SVMs) used by the databases on the host.
4. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
5. In the Resources page, ensure that the database to be imported is discovered and displayed.

   When you want to run the import tool, the database must be accessible or else the resource group creation fails.

   If the database has credentials configured, you must create a corresponding credential in SnapCenter, assign the credential to the database, and then re-run discovery of the database. If the database is residing on Automatic Storage Management (ASM), you must create credentials for the ASM instance, and assign the credential to the database.

6. Ensure that the user running the import tool has sufficient privileges to run SnapManager for Oracle or SnapManager for SAP CLI commands (such as the command to suspend schedules) from SnapManager for Oracle or SnapManager for SAP host.
7. Run the following commands on the SnapManager for Oracle or SnapManager for SAP host to suspend the schedules:
   a. If you want to suspend the schedules on the SnapManager for Oracle host, run:
      - `smo credential set -repository -dbname repository_database_name -host host_name -port port_number -login -username user_name_for_repository_database`
      - `smo profile sync -repository -dbname repository_database_name -host host_name -port port_number -login -username host_user_name_for_repository_database`
      - `smo credential set -profile -name profile_name`

        You must run the smo credential set command for each profile on the host.

b. If you want to suspend the schedules on the SnapManager for SAP host, run:

- `smsap credential set -repository -dbname repository_database_name -host host_name -port port_number -login -username user_name_for_repository_database`
- `smsap profile sync -repository -dbname repository_database_name -host host_name -port port_number -login -username host_user_name_for_repository_database`
- `smsap credential set -profile -name profile_name`

> You must run the smsap credential set command for each profile on the host.

8. Ensure that fully qualified domain name (FQDN) of the database host is displayed when you run hostname -f.

   If FQDN is not displayed, you must modify /etc/hosts to specify the FQDN of the host.

Import data

You can import data by running the import tool from the database host.

About this task

The SnapCenter backup policies that are created after importing have different naming formats:

- Policies created for the profiles without any operations and schedules have the SM_PROFILENAME_ONLINE_FULL_DEFAULT_MIGRATED format.

  When no operation is performed using a profile, the corresponding policy is created with default backup type as online and backup scope as full.

- Policies created for the profiles with one or more operations have the SM_PROFILENAME_BACKUPMODE_BACKUPSCOPE_MIGRATED format.
- Policies created for the schedules attached to the profiles have the SM_PROFILENAME_SMOSCHEDULENAME_BACKUPMODE_BACKUPSCOPE_MIGRATED format.

Steps

1. Log in to the database host that you want to import.
2. Run the import tool by executing the sc-migrate script located at */opt/NetApp/snapcenter/spl/bin*.
3. Enter the SnapCenter Server user name and password.

   After validating the credentials, a connection is established with SnapCenter.

4. Enter the SnapManager for Oracle or SnapManager for SAP repository database details.

   The repository database lists the databases that are available on the host.

5. Enter the target database details.

   If you want to import all the databases on the host, enter all.

6. If you want to generate a system log or send ASUP messages for failed operations, you must enable them either by running the *Add-SmStorageConnection* or *Set-SmStorageConnection* command.

> If you want to cancel an import operation, either while running the import tool or after importing, you must manually delete the SnapCenter policies, credentials, and resource groups that were created as part of import operation.

Results

The SnapCenter backup policies are created for profiles, schedules, and operations performed using the profiles. Resource groups are also created for each target database.

After importing the data successfully, the schedules associated with the imported database are suspended in SnapManager for Oracle and SnapManager for SAP.

> After importing, you must manage the imported database or file system using SnapCenter.

The logs for every execution of the import tool are stored in the */var/opt/snapcenter/spl/logs* directory with the name spl_migration_timestamp.log. You can refer to this log to review import errors and troubleshoot them.

Install SnapCenter Plug-in for VMware vSphere

If your database or filesystem is stored on virtual machines (VMs), or if you want to protect VMs and datastores, you must deploy the SnapCenter Plug-in for VMware vSphere virtual appliance.

For information to deploy, see Deployment Overview.

Deploy CA certificate

To configure the CA Certificate with SnapCenter Plug-in for VMware vSphere, see Create or import SSL certificate.

Configure the CRL file

SnapCenter Plug-in for VMware vSphere looks for the CRL files in a pre-configured directory. Default directory of the CRL files for SnapCenter Plug-in for VMware vSphere is */opt/netapp/config/crl*.

You can place more than one CRL file in this directory. The incoming certificates will be verified against each CRL.

Prepare for protecting Oracle databases

Before performing any data protection operation such as backup, clone, or restore operations, you must define your strategy and set up the environment. You can also set up the SnapCenter Server to use SnapMirror and SnapVault technology.

To take advantage of SnapVault and SnapMirror technology, you must configure and initialize a data protection relationship between the source and destination volumes on the storage device. You can use NetAppSystem Manager or you can use the storage console command line to perform these tasks.

Before you use the Plug-in for Oracle Database, the SnapCenter administrator should install and configure the SnapCenter Server and perform the prerequisite tasks.

- Install and configure SnapCenter Server. Learn more
- Configure the SnapCenter environment by adding storage system connections. Learn more

> SnapCenter does not support multiple SVMs with the same name on different clusters. Each SVM registered with SnapCenter using either SVM registration or cluster registration must be unique.

- Create credentials with authentication mode as Linux or AIX for the install user. Learn more
- Add hosts, install the plug-ins, and discover the resources.
- If you are using SnapCenter Server to protect Oracle databases that reside on VMware RDM LUNs or VMDKs, you must deploy the SnapCenter Plug-in for VMware vSphere and register the plug-in with SnapCenter.
- Install Java on your Linux or AIX host.

  See Linux host requirements or AIX host requirements for more information.

- You should set the time out value of the application firewall to 3 hours or more.
- If you have Oracle databases on NFS environments, you must have configured at least one NFS data LIF for primary or secondary storage to perform mount, clone, verification, and restore operations.
- If you have multiple data paths (LIFs) or a dNFS configuration, you can perform the following using the SnapCenter CLI on the database host:
  - By default, all the IP addresses of the database host are added to the NFS storage export policy in storage virtual machine (SVM) for the cloned volumes. If you want to have a specific IP address or restrict to a subset of the IP addresses, run the Set-PreferredHostIPsInStorageExportPolicy CLI.
  - If you have multiple data paths (LIFs) in SVM, SnapCenter chooses the appropriate data path (LIF) for mounting the NFS cloned volume. However, if you want to specify a specific data path (LIF), you must run the Set-SvmPreferredDataPath CLI.
    The command reference guide has more information.
- If you have Oracle databases on SAN environments, ensure that the SAN environment is configured as per the recommendation mentioned in the following guides:
  - Using Linux Hosts with ONTAP storage
  - Host Settings Affected by AIX Host Utilities
- If you have Oracle databases on LVM in Oracle Linux or RHEL operating systems, install the latest version of Logical Volume Management (LVM).
- If you are using SnapManager for Oracle and want to migrate to SnapCenter Plug-in for Oracle Database, you can migrate the profiles to policies and resource groups of SnapCenter by using the sccli command sc-migrate.
- Configure SnapMirror and SnapVault on ONTAP, if you want backup replication

For SnapCenter 4.1.1 users, the SnapCenter Plug-in for VMware vSphere 4.1.1 documentation has information on protecting virtualized databases and file systems. For SnapCenter 4.2.x users, the NetApp Data Broker 1.0 and 1.0.1, documentation has information on protecting virtualized databases and file systems using the SnapCenter Plug-in for VMware vSphere that is provided by the Linux-based NetApp Data Broker virtual appliance (Open Virtual Appliance format). For SnapCenter 4.3.x users, the SnapCenter Plug-in for VMware vSphere 4.3 documentation has information on protecting virtualized databases and file systems

using the Linux-based SnapCenter Plug-in for VMware vSphere virtual appliance (Open Virtual Appliance format).

Find more information

- Interoperability Matrix Tool
- SnapCenter Plug-in for VMware vSphere documentation
- Data protection operation fails in a non-multipath environment in RHEL 7 and later

Back up Oracle databases

Overview of backup procedure

You can either create a backup of a resource (database) or resource group. The backup procedure includes planning, identifying the resources for backup, creating backup policies, creating resource groups and attaching policies, creating backups, and monitoring the operations.

The following workflow shows the sequence in which you must perform the backup operation:

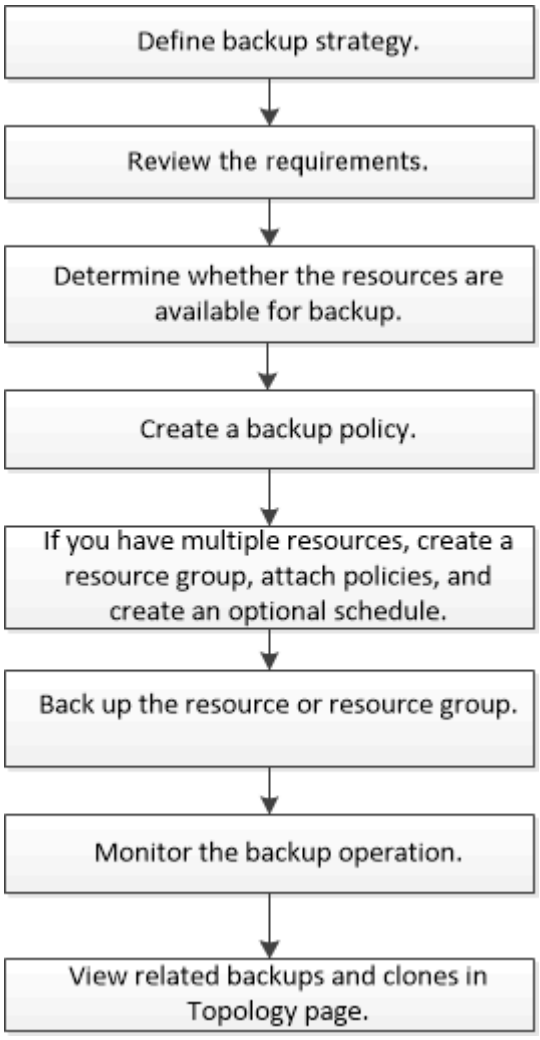

While creating a backup for Oracle databases, an operational lock file (*.sm_lock_dbsid*) is created on the

Oracle database host in the */var/opt/snapcenter/sco/lock* directory to avoid multiple operations being executed on the database. After the database has been backed up, the operational lock file is automatically removed.

However, if the previous backup was completed with a warning, the operational lock file might not get deleted, and the next backup operation gets into the wait queue. It might eventually get canceled if the **.sm_lock_dbsid** file is not deleted. In such scenario, you must manually delete the operational lock file by performing the following steps:

1. From the command prompt, navigate to */var/opt/snapcenter/sco/lock*.
2. Delete the operational lock:`rm -rf .sm_lock_dbsid.`

Backup configuration information

Supported Oracle database configurations for backups

SnapCenter supports backup of different Oracle database configurations.

- Oracle Standalone
- Oracle Real Application Clusters (RAC)
- Oracle Standalone Legacy
- Oracle Standalone Container Database (CDB)
- Oracle Data Guard standby

  You can only create offline-mount backups of Data Guard standby databases. Offline-shutdown backup, archive log only backup, and full backup are not supported.

- Oracle Active Data Guard standby

  You can only create online backups of Active Data Guard standby databases. Archive log only backup and full backup are not supported.

  Before creating a backup of Data Guard standby or Active Data Guard standby database, the managed recovery process (MRP) is stopped and once the backup is created, MRP is started.

- Automatic Storage Management (ASM)
  - ASM standalone and ASM RAC on Virtual Machine Disk (VMDK)

    Among all the restore methods supported for Oracle databases, you can perform only connect-and-copy restore of ASM RAC databases on VMDK.

  - ASM standalone and ASM RAC on Raw device mapping (RDM)

    You can perform backup, restore, and clone operations on Oracle databases on ASM, with or without ASMLib.

  - Oracle ASM Filter Driver (ASMFD)

    PDB migration and PDB cloning operations are not supported.

  - Oracle Flex ASM

For the latest information about supported Oracle versions, see the NetApp Interoperability Matrix Tool.

Types of backup supported for Oracle databases

Backup type specifies the type of backup that you want to create. SnapCenter supports online and offline backup types for Oracle databases.

Online backup

A backup that is created when the database is in the online state is called an online backup. Also called a hot backup, an online backup enables you to create a backup of the database without shutting it down.

As part of online backup, you can create a backup of the following files:

- Data files and control files only
- Archive log files only (the database is not brought to backup mode in this scenario)
- Full database that includes data files, control files, and archive log files

Offline backup

A backup created when the database is either in a mounted or shutdown state is called an offline backup. An offline backup is also called a cold backup. You can include only data files and control files in offline backups. You can create either an offline mount or offline shutdown backup.

- When creating an offline mount backup, you must ensure that the database is in a mounted state.

  If the database is in any other state, the backup operation fails.

- When creating an offline shutdown backup, the database can be in any state.

  The database state is changed to the required state to create a backup. After creating the backup, the database state is reverted to the original state.

How SnapCenter discovers Oracle databases

Resources are Oracle databases on the host that are maintained by SnapCenter. You can add these databases to resource groups to perform data protection operations after you discover the databases that are available.

The following sections describe the process that SnapCenter uses to discover different types and versions of Oracle databases.

For Oracle versions 11*g* to 12*c*R1

RAC database

The RAC databases are discovered only on the basis of /etc/oratab`entries. You should have the database entries in the /etc/oratab file.

Standalone

The standalone databases are discovered only on the basis of /etc/oratab entries.

ASM

The ASM instance entry should be available in the /etc/oratab file.

RAC One Node

The RAC One Node databases are discovered only on the basis of /etc/oratab entries.
The databases should be either in nomount, mount, or open state. You should have the database entries in the /etc/oratab file.

The RAC One Node database status will be marked as renamed or deleted if the database is already discovered and backups are associated with the database.

You should perform the following steps if the database is relocated:

1. Manually add the relocated database entry in the /etc/oratab file on the failed-over RAC node.
2. Manually refresh the resources.
3. Select the RAC One Node database from the resource page, and then click Database Settings.
4. Configure the database to set the preferred cluster nodes to the RAC node currently hosting the database.
5. Perform the SnapCenter operations.
6. If you have relocated a database from one node to another node and if the oratab entry in the earlier node is not deleted, manually delete the oratab entry to avoid the same database being displayed twice.

For Oracle versions 12cR2 to 18c, 19c, or 21c

RAC database

The RAC databases are discovered using the srvctl config command.
You should have the database entries in the /etc/oratab file.

Standalone

The standalone databases are discovered based on the entries in the /etc/oratab file and the output of the srvctl config command.

ASM

The ASM instance entry need not be in the /etc/oratab file.

RAC One Node

The RAC One Node databases are discovered using the srvctl config command only.
The databases should be either in nomount, mount, or open state. The RAC One Node database status will be marked as renamed or deleted if the database is already discovered and backups are associated with the database.

You should perform the following steps if the database is relocated:
. Manually refresh the resources.
. Select the RAC One Node database from the resource page, and then click Database Settings.
. Configure the database to set the preferred cluster nodes to the RAC node currently hosting the database.
. Perform the SnapCenter operations.

> If there are any Oracle 12*c*R2 and 18*c* database entries in the /etc/oratab file and the same database is registered with the srvctl config command, SnapCenter will eliminate the duplicate database entries.
> If there are stale database entries, the database will be discovered but the database will be unreachable and the status will be offline.

Preferred nodes in RAC setup

In Oracle Real Application Clusters (RAC) setup, you can specify the preferred nodes that SnapCenter uses to perform the backup operation. If you do not specify the preferred node, SnapCenter automatically assigns a node as the preferred node and backup is created on that node.

The preferred nodes might be one or all of the cluster nodes where the RAC database instances are present. The backup operation is triggered only on these preferred nodes in the order of the preference.

Example

The RAC database cdbrac has three instances: cdbrac1 on node1, cdbrac2 on node2, and cdbrac3 on node3.

The node1 and node2 instances are configured to be the preferred nodes, with node2 as the first preference and node1 as the second preference. When you perform a backup operation, the operation is first attempted on node2 because it is the first preferred node.

If node2 is not in the state to back up, which could be due to multiple reasons such as the plug-in agent is not running on the host, the database instance on the host is not in the required state for the specified backup type, or the database instance on node2 in a FlexASM configuration is not being served by the local ASM instance; then the operation will be attempted on node1.

The node3 will not be used for backup because it is not on the list of preferred nodes.

Flex ASM setup

In a Flex ASM setup, Leaf nodes will not be listed as preferred nodes if the cardinality is less than the number nodes in the RAC cluster. If there is any change in the Flex ASM cluster node roles, you should manually discover so that the preferred nodes are refreshed.

Required database state

The RAC database instances on the preferred nodes must be in the required state for the backup to finish successfully:

- One of the RAC database instances in the configured preferred nodes must be in the open state to create an online backup.
- One of the RAC database instances in the configured preferred nodes must be in the mount state, and all other instances, including other preferred nodes, must be in the mount state or lower to create an offline mount backup.
- RAC database instances can be in any state, but you must specify the preferred nodes to create an offline shutdown backup.

How to catalog backups with Oracle Recovery Manager

You can catalog the backups of Oracle databases using Oracle Recovery Manager (RMAN) to store the backup information in the Oracle RMAN repository.

The cataloged backups can be used later for block-level restore or tablespace point-in-time recovery operations. When you do not need these cataloged backups, you can remove the catalog information.

The database must be in mounted or higher state for cataloging. You can perform cataloging on data backups, archive log backups, and full backups. If cataloging is enabled for a backup of a resource group that has

multiple databases, cataloging is performed for each database. For Oracle RAC databases, cataloging will be performed on the preferred node where the database is at least in mounted state.

If you want to catalog backups of a RAC database, ensure that no other job is running for that database. If another job is running, the cataloging operation fails instead of getting queued.

External catalog database

By default, the target database control file is used for cataloging. If you want to add external catalog database, you can configure it by specifying the credential and Transparent Network Substrate (TNS) name of the external catalog using the Database Settings wizard from the SnapCenter graphical user interface (GUI). You can also configure the external catalog database from the CLI by running the Configure-SmOracleDatabase command with the -OracleRmanCatalogCredentialName and -OracleRmanCatalogTnsName options.

RMAN command

If you enabled the cataloging option while creating an Oracle backup policy from the SnapCenter GUI, the backups are cataloged using Oracle RMAN as a part of the backup operation. You can also perform deferred cataloging of backups by running the `Catalog-SmBackupWithOracleRMAN` command.

After cataloging the backups, you can run the `Get-SmBackupDetails` command to obtain the cataloged backup information such as the tag for cataloged datafiles, the control file catalog path, and the cataloged archive log locations.

Naming format

If the ASM disk group name is greater than or equal to 16 characters, from SnapCenter 3.0, the naming format used for the backup is SC_HASHCODEofDISKGROUP_DBSID_BACKUPID. However, If the disk group name is less than 16 characters, the naming format used for the backup is DISKGROUPNAME_DBSID_BACKUPID, which is the same format used in SnapCenter 2.0.

The HASHCODEofDISKGROUP is an automatically generated number (2 to 10 digit) unique for each ASM disk group.

Crosscheck operations

You can perform crosschecks to update outdated RMAN repository information about backups whose repository records do not match their physical status. For example, if a user removes archived logs from disk with an operating system command, the control file still indicates that the logs are on disk, when in fact they are not.

The crosscheck operation enables you to update the control file with the information. You can enable crosscheck by running the Set-SmConfigSettings command and assigning the value TRUE to the ENABLE_CROSSCHECK parameter. The default value is set to FALSE.

```
sccli Set-SmConfigSettings-ConfigSettingsTypePlugin-PluginCodeSCO-ConfigSettings
"KEY=ENABLE_CROSSCHECK, VALUE=TRUE"
```

Remove catalog information

You can remove the catalog information by running the Uncatalog-SmBackupWithOracleRMAN command. You cannot remove the catalog information using the SnapCenter GUI. However, information of a cataloged backup is removed while deleting the backup or while deleting the retention and resource group associated with that cataloged backup.

When you force a deletion of the SnapCenter host, the information of the cataloged backups associated with that host are not removed. You must remove information of all the cataloged backups for that host before forcing the deletion of the host.

If the cataloging and uncataloging fails because the operation time exceeded the time out value specified for the ORACLE_PLUGIN_RMAN_CATALOG_TIMEOUT parameter, you should modify the value of the parameter by running the following command:

```
/opt/Netapp/snapcenter/spl/bin/sccli Set-SmConfigSettings-ConfigSettingsType
Plugin -PluginCode SCO-ConfigSettings
"KEY=ORACLE_PLUGIN_RMAN_CATALOG_TIMEOUT,VALUE=user_defined_value"
```

After modifying the value of the parameter, restart the SnapCenter Plug-in Loader (SPL) service by running the following command:

```
/opt/NetApp/snapcenter/spl/bin/spl restart
```

The information regarding the parameters that can be used with the command and their descriptions can be obtained by running Get-Help command_name. Alternatively, you can refer to the SnapCenter Software Command Reference Guide.

Predefined environment variables for backup specific prescript and postscript

SnapCenter allows you to use the predefined environment variables when you execute the prescript and postscript while creating backup policies. This functionality is supported for all Oracle configurations except for VMDK.

SnapCenter predefines the values of the parameters that will be directly accessible in the environment where the shell scripts are executed. You do not have to manually specify the values of these parameters when executing the scripts.

Supported predefined environment variables for creating backup policy

- **SC_JOB_ID** specifies the job ID of the operation.

  Example: 256

- **SC_ORACLE_SID** specifies the system identifier of the database.

  If the operation involves multiple databases, the parameter will contain database names separated by pipe.

  This parameter will be populated for application volumes.

  Example: NFSB32|NFSB31

- **SC_HOST** specifies the host name of the database.

  For RAC, host name will be the name of the host on which backup is performed.

  This parameter will be populated for application volumes.

  Example: scsmohost2.gdl.englabe.netapp.com

- **SC_OS_USER** specifies the operating system owner of the database.

The data will be formatted as <db1>@<osuser1>|<db2>@<osuser2>.

Example: NFSB31@oracle|NFSB32@oracle

- **SC_OS_GROUP** specifies the operating system group of the database.

  The data will be formatted as <db1>@<osgroup1>|<db2>@<osgroup2>.

  Example: NFSB31@install|NFSB32@oinstall

- **SC_BACKUP_TYPE**" specifies the backup type (online full, online data, online log, offline shutdown, offline mount)

  Examples:

  - For full backup: ONLINEFULL
  - data only backup: ONLINEDATA
  - For log only backup: ONLINELOG

- **SC_BACKUP_NAME** specifies the name of the backup.

  This parameter will be populated for application volumes.

  Example: DATA@RG2_scspr2417819002_07-20-2021_12.16.48.9267_0|LOG@RG2_scspr2417819002_07-20-2021_12.16.48.9267_1|AV@RG2_scspr2417819002_07-20-2021_12.16.48.9267

- **SC_BACKUP_ID** specifies the backup ID.

  This parameter will be populated for application volumes.

  Example: DATA@203|LOG@205|AV@207

- **SC_ORACLE_HOME** specifies the path of the Oracle home directory.

  Example: NFSB32@/ora01/app/oracle/product/18.1.0/db_1|NFSB31@/ora01/app/oracle/product/18.1.0/db_1

- **SC_BACKUP_RETENTION** specifies the retention period defined in the policy.

  Examples:

  - For full backup: Hourly|DATA@DAYS:3|LOG@COUNT:4
  - For on-demand data only backup: Ondemand|DATA@COUNT:2
  - For on-demand log only backup: Ondemand|LOG@COUNT:2

- **SC_RESOURCE_GROUP_NAME** specifies the name of the resource group.

  Example: RG1

- **SC_BACKUP_POLICY_NAME** specifies the name of the backup policy.

  Example: backup_policy

- **SC_AV_NAME** specifies the names of the application volumes.

Example: AV1|AV2

- **SC_PRIMARY_DATA_VOLUME_FULL_PATH** specifies the storage mapping of SVM to volume for data files directory. It will be the name of the parent volume for luns and qtrees.

  The data will be formatted as <db1>@<SVM1:volume1>|<db2>@<SVM2:volume2>.

  Examples:

  - For 2 databases in the same resource group: NFSB32@buck:/vol/scspr2417819002_NFS_CDB_NFSB32_DATA|NFSB31@buck:/vol/scspr2417819002_NFS_CDB_NFSB31_DATA
  - For single database with data files spread across multiple volumes: buck:/vol/scspr2417819002_NFS_CDB_NFSB31_DATA,herculus:/vol/scspr2417819002_NFS
- **SC_PRIMARY_ARCHIVELOGS_VOLUME_FULL_PATH** specifies the storage mapping of SVM to volume for logs file directory. It will be the name of the parent volume for luns and qtrees.

  Examples:

  - For single database instance: buck:/vol/scspr2417819002_NFS_CDB_NFSB31_REDO
  - For multiple database instances: NFSB31@buck:/vol/scspr2417819002_NFS_CDB_NFSB31_REDO|NFSB32@buck:/vol/scspr2417819002_NFS_CDB_NFSB32_REDO
- **SC_PRIMARY_FULL_SNAPSHOT_NAME_FOR_TAG** specifies the list of Snapshots containing storage system name and volume name.

  Examples:

  - For single database instance: buck:/vol/scspr2417819002_NFS_CDB_NFSB32_DATA/RG2_scspr2417819002_07-21-2021_02.28.26.3973_0,buck:/vol/scspr2417819002_NFS_CDB_NFSB32_REDO/RG2_scspr2417819002_07-21-2021_02.28.26.3973_1
  - For multiple database instances: NFSB32@buck:/vol/scspr2417819002_NFS_CDB_NFSB32_DATA/RG2_scspr2417819002_07-21-2021_02.28.26.3973_0,buck:/vol/scspr2417819002_NFS_CDB_NFSB32_REDO/RG2_scspr2417819002_07-21-2021_02.28.26.3973_1|NFSB31@buck:/vol/scspr2417819002_NFS_CDB_NFSB31_DATA/RG2_scspr2417819002_07-21-2021_02.28.26.3973_0,buck:/vol/scspr2417819002_NFS_CDB_NFSB31_REDO/RG2_scspr2417819002_07-21-2021_02.28.26.3973_1
- **SC_PRIMARY_SNAPSHOT_NAMES** specifies the names of the primary Snapshots created during the backup.

  Examples:

  - For single database instance: RG2_scspr2417819002_07-21-2021_02.28.26.3973_0,RG2_scspr2417819002_07-21-2021_02.28.26.3973_1
  - For multiple database instances: NFSB32@RG2_scspr2417819002_07-21-2021_02.28.26.3973_0,RG2_scspr2417819002_07-21-2021_02.28.26.3973_1|NFSB31@RG2_scspr2417819002_07-21-2021_02.28.26.3973_0,RG2_scspr2417819002_07-21-2021_02.28.26.3973_1

    - For consistency group Snapshots that involves 2 volumes: cg3_R80404CBEF5V1_04-05-2021_03.08.03.4945_0_bfc279cc-28ad-465c-9d60-5487ac17b25d_2021_4_5_3_8_58_350
- **SC_PRIMARY_MOUNT_POINTS** specifies the mount point details which are part of the backup.

  The details include the directory on which volumes are mounted and not the immediate parent of the file under backup. For an ASM configuration, it is the name of the disk group.

  The data will be formatted as <db1>@<mountpoint1,mountpoint2>|<db2>@<mountpoint1,mountpoint2>.

  Examples:

    - For single database instance: /mnt/nfsdb3_data,/mnt/nfsdb3_log,/mnt/nfsdb3_data1
    - For multiple database instances: NFSB31@/mnt/nfsdb31_data,/mnt/nfsdb31_log,/mnt/nfsdb31_data1|NFSB32@/mnt/nfsdb32_data,/mnt/nfsdb32_log,/mnt/nfsdb32_data1
    - For ASM: +DATA2DG,+LOG2DG
- **SC_PRIMARY_SNAPSHOTS_AND_MOUNT_POINTS** specifies the names of the snapshots created during the backup of each of the mount points.

  Examples:

    - For single database instance: RG2_scspr2417819002_07-21-2021_02.28.26.3973_0:/mnt/nfsb32_data,RG2_scspr2417819002_07-21-2021_02.28.26.3973_1:/mnt/nfsb31_log
    - For multiple database instances: NFSB32@RG2_scspr2417819002_07-21-2021_02.28.26.3973_0:/mnt/nfsb32_data,RG2_scspr2417819002_07-21-2021_02.28.26.3973_1:/mnt/nfsb31_log|NFSB31@RG2_scspr2417819002_07-21-2021_02.28.26.3973_0:/mnt/nfsb31_data,RG2_scspr2417819002_07-21-2021_02.28.26.3973_1:/mnt/nfsb32_log
- **SC_ARCHIVELOGS_LOCATIONS** specifies the location of the archive logs directory.

  The directory names will be the immediate parent of the archive log files. If the archive logs are placed in more than one location then all the locations will be captured. This also includes the FRA scenarios. If softlinks are used for directory then the same will be populated.

  Examples:

    - For single database on NFS: /mnt/nfsdb2_log
    - For multiple databases on NFS and for the NFSB31 database archive logs that are placed in two different locations: NFSB31@/mnt/nfsdb31_log1,/mnt/nfsdb31_log2|NFSB32@/mnt/nfsdb32_log
    - For ASM: +LOG2DG/ASMDB2/ARCHIVELOG/2021_07_15
- **SC_REDO_LOGS_LOCATIONS** specifies the location of the redo logs directory.

  The directory names will be the immediate parent of the redo log files. If softlinks are used for directory then the same will be populated.

  Examples:

    - For single database on NFS: /mnt/nfsdb2_data/newdb1
    - For multiple databases on NFS: NFSB31@/mnt/nfsdb31_data/newdb31|NFSB32@/mnt/nfsdb32_data/newdb32

- For ASM: +LOG2DG/ASMDB2/ONLINELOG

- **SC_CONTROL_FILES_LOCATIONS** specifies the location of the control files directory.

  The directory names will be the immediate parent of the control files. If softlinks are used for directory then the same will be populated.

  Examples:

  - For single database on NFS: /mnt/nfsdb2_data/fra/newdb1,/mnt/nfsdb2_data/newdb1
  - For multiple databases on NFS: NFSB31@/mnt/nfsdb31_data/fra/newdb31,/mnt/nfsdb31_data/newdb31|NFSB32@/mnt/nfsdb32_data/fra/newdb32,/mnt/nfsdb32_data/newdb32
  - For ASM: +LOG2DG/ASMDB2/CONTROLFILE

- **SC_DATA_FILES_LOCATIONS**" specifies the location of the data files directory.

  The directory names will be the immediate parent of the data files. If softlinks are used for directory then the same will be populated.

  Examples:

  - For single database on NFS: /mnt/nfsdb3_data1,/mnt/nfsdb3_data/NEWDB3/datafile
  - For multiple databases on NFS: NFSB31@/mnt/nfsdb31_data1,/mnt/nfsdb31_data/NEWDB31/datafile|NFSB32@/mnt/nfsdb32_data1,/mnt/nfsdb32_data/NEWDB32/datafile
  - For ASM: +DATA2DG/ASMDB2/DATAFILE,+DATA2DG/ASMDB2/TEMPFILE

- **SC_SNAPSHOT_LABEL** specifies the name of the secondary labels.

  Examples: Hourly, Daily, Weekly, Monthly, or custom label.

Supported delimiters

- **:** is used to separate SVM name and volume name

  Example: buck:/vol/scspr2417819002_NFS_CDB_NFSB32_DATA/RG2_scspr2417819002_07-21-2021_02.28.26.3973_0,buck:/vol/scspr2417819002_NFS_CDB_NFSB32_REDO/RG2_scspr2417819002_07-21-2021_02.28.26.3973_1

- **@** is used to separate data from its database name and to separate the value from its key.

  Examples:

  - NFSB32@buck:/vol/scspr2417819002_NFS_CDB_NFSB32_DATA/RG2_scspr2417819002_07-21-2021_02.28.26.3973_0,buck:/vol/scspr2417819002_NFS_CDB_NFSB32_REDO/RG2_scspr2417819002_07-21-2021_02.28.26.3973_1|NFSB31@buck:/vol/scspr2417819002_NFS_CDB_NFSB31_DATA/RG2_scspr2417819002_07-21-2021_02.28.26.3973_0,buck:/vol/scspr2417819002_NFS_CDB_NFSB31_REDO/RG2_scspr2417819002_07-21-2021_02.28.26.3973_1
  - NFSB31@oracle|NFSB32@oracle

- **|** is used to separate the data between two different databases and to separate the data between two different entities for SC_BACKUP_ID, SC_BACKUP_RETENTION, and SC_BACKUP_NAME parameters.

Examples:

- DATA@203|LOG@205
- Hourly|DATA@DAYS:3|LOG@COUNT:4
- DATA@RG2_scspr2417819002_07-20-2021_12.16.48.9267_0|LOG@RG2_scspr2417819002_07-20-2021_12.16.48.9267_1

- **/** is used to separate the volume name from it's Snapshot for SC_PRIMARY_SNAPSHOT_NAMES and SC_PRIMARY_FULL_SNAPSHOT_NAME_FOR_TAG parameters.

  Example: NFSB32@buck:/vol/scspr2417819002_NFS_CDB_NFSB32_DATA/RG2_scspr2417819002_07-21-2021_02.28.26.3973_0,buck:/vol/scspr2417819002_NFS_CDB_NFSB32_REDO/RG2_scspr2417819002_07-21-2021_02.28.26.3973_1

- **,** is used to separate set of variables for the same DB.

  Example: NFSB32@buck:/vol/scspr2417819002_NFS_CDB_NFSB32_DATA/RG2_scspr2417819002_07-21-2021_02.28.26.3973_0,buck:/vol/scspr2417819002_NFS_CDB_NFSB32_REDO/RG2_scspr2417819002_07-21-2021_02.28.26.3973_1|NFSB31@buck:/vol/scspr2417819002_NFS_CDB_NFSB31_DATA/RG2_scspr2417819002_07-21-2021_02.28.26.3973_0,buck:/vol/scspr2417819002_NFS_CDB_NFSB31_REDO/RG2_scspr2417819002_07-21-2021_02.28.26.3973_1

Backup retention options

You can choose either the number of days for which to retain backup copies or specify the number of backup copies you want to retain, up to a ONTAP maximum of 255 copies. For example, your organization might require that you retain 10 days of backup copies or 130 backup copies.

While creating a policy, you can specify the retention options for the backup type and the schedule type.

If you set up SnapMirror replication, the retention policy is mirrored on the destination volume.

SnapCenter deletes the retained backups that have retention labels that match the schedule type. If the schedule type was changed for the resource or resource group, backups with the old schedule type label might still remain on the system.

For long-term retention of backup copies, you should use SnapVault backup.

Backup schedules

Backup frequency (schedule type) is specified in policies; a backup schedule is specified in the resource group configuration. The most critical factor in determining a backup frequency or schedule is the rate of change for the resource and the importance of the data. You might back up a heavily used resource every hour, while you might back up a rarely used resource once a day. Other factors include the importance of the resource to your organization, your Service Level Agreement (SLA), and your Recover Point

Objective (RPO).

An SLA defines the level of service expected and addresses many service-related issues, including the availability and performance of service. An RPO defines the strategy for the age of the files that must be recovered from backup storage for regular operations to resume after a failure. The SLA and RPO contribute to the data protection strategy.

Even for a heavily used resource, there is no requirement to run a full backup more than once or twice a day. For example, regular transaction log backups might be sufficient to ensure that you have the backups you need. The more often you back up your databases, the fewer transaction logs SnapCenter has to use at restore time, which can result in faster restore operations.

Backup schedules have two parts, as follows:

- Backup frequency

  Backup frequency (how often backups are to be performed), called *schedule type* for some plug-ins, is part of a policy configuration. You can select hourly, daily, weekly, or monthly as the backup frequency for the policy. If you do not select any of these frequencies, then the policy created is an on-demand-only policy. You can access policies by clicking **Settings** > **Policies**.

- Backup schedules

  Backup schedules (exactly when backups are to be performed) are part of a resource group configuration. For example, if you have a resource group that has a policy configured for weekly backups, you might configure the schedule to back up every Thursday at 10:00 PM. You can access resource group schedules by clicking **Resources** > **Resource Groups**.

Backup naming conventions

You can either use the default Snapshot naming convention or use a customized naming convention. The default backup naming convention adds a timestamp to Snapshot names that helps you identify when the copies were created.

The Snapshot uses the following default naming convention:

```
resourcegroupname_hostname_timestamp
```

You should name your backup resource groups logically, as in the following example:

```
dts1_mach1x88_03-12-2015_23.17.26
```

In this example, the syntax elements have the following meanings:

- *dts1* is the resource group name.
- *mach1x88* is the host name.
- *03-12-2015_23.17.26* is the date and timestamp.

Alternatively, you can specify the Snapshot name format while protecting resources or resource groups by selecting **Use custom name format for Snapshot copy**. For example, customtext_resourcegroup_policy_hostname or resourcegroup_hostname. By default, the time stamp suffix is

added to the Snapshot name.

Requirements for backing up an Oracle database

Before backing up an Oracle database, you should ensure that prerequisites are completed.

- You must have created a resource group with a policy attached.
- If you want to back up a resource that has a SnapMirror relationship with a secondary storage, the ONTAP role assigned to the storage user should include the “snapmirror all” privilege. However, if you are using the “vsadmin” role, then the “snapmirror all” privilege is not required.
- You must have assigned the aggregate that is being used by the backup operation to the storage virtual machine (SVM) used by the database.
- You should have verified that all data volumes and archive log volumes belonging to the database are protected if secondary protection is enabled for that database.
- You should have verified that the database that has files on the ASM disk groups should be in either “MOUNT” or “OPEN” state to verify its backups using the Oracle DBVERIFY utility.
- You should have verified that the volume mount point length does not exceed 240 characters.
- You should increase value of RESTTimeout to 86400000 ms in *C:\Program Files\NetApp \SMCore\SMCoreServiceHost.exe.config* file in the SnapCenter Server host, if the database being backed up is large (size in TBs).

  While modifying the values ensure that there are no running jobs and restart the SnapCenter SMCore service after increasing the value.

Discover Oracle databases available for backup

Resources are Oracle databases on the host that are managed by SnapCenter. You can add these databases to resource groups to perform data protection operations after you discover the databases that are available.

Before you begin

- You must have completed tasks such as installing the SnapCenter Server, adding hosts, creating storage system connections, and adding credentials.
- If the databases reside on a Virtual Machine Disk (VMDK) or raw device mapping (RDM), you must deploy the SnapCenter Plug-in for VMware vSphere and register the plug-in with SnapCenter.

  For more information, see Deploy SnapCenter Plug-in for VMware vSphere.
- If databases reside on a VMDK file system, you must have logged in to vCenter and navigated to **VM options** > **Advanced** > **Edit configuration** to set the value of *disk.enableUUID* to true for the VM.
- You must have reviewed the process that SnapCenter follows to discover different types and versions of Oracle databases.

Step 1: Prevent SnapCenter from discovering non-database entries

You can prevent SnapCenter from discovering non-database entries added in the oratab file.

Steps

1. After installing the plug-in for Oracle, the root user should create the **sc_oratab.config** file under the directory */var/opt/snapcenter/sco/etc/*.

   Grant the write permission to the Oracle binary owner and group so that the file could be maintained in future.

2. Database administrator should add the non-database entries in the **sc_oratab.config** file.

   It is recommended to maintain same format defined for the non-database entries in the */etc/oratab* file or the user can just add the non-database entity string.

   The string is case sensitive. Any text with # in the beginning is treated as a comment. The comment can be appended after the non-database name.

```
For example:
-----------------------------------------
# Sample entries
# Each line can have only one non-database name
# These are non-database name
oratar  # Added by the admin group -1
#Added by the script team
NEWSPT
DBAGNT:/ora01/app/oracle/product/agent:N
-----------------------------------------
```

3. Discover the resources.

   The non-database entries added in the **sc_oratab.config** will not be listed in the Resources page.

It is always recommended to take a backup of the sc_oratab.config file before upgrading the SnapCenter plug-in.

Step 2: Discover resources

After installing the plug-in, all of the databases on that host are automatically discovered and displayed in the Resources page.

The databases should be at least in the mounted state or above for the discovery of the databases to be successful. In an Oracle Real Application Clusters (RAC) environment, the RAC database instance in the host where the discovery is performed, should be at least in the mounted state or above for the discovery of the database instance to be successful. Only the databases that are discovered successfully can be added to the resource groups.

If you have deleted an Oracle database on the host, SnapCenter Server will not be aware and will list the deleted database. You should manually refresh the resources to update the SnapCenter resources list.

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.

2. In the Resources page, select **Database** from the **View** list.

   Click ilter icon], and then select the host name and the database type to filter the resources. You can then click the imagfilter_icon.gif[filter icon]_icon.png[] icon to close the filter pane.

3. Click **Refresh Resources**.

   In a RAC One Node scenario, the database is discovered as the RAC database on the node where it is currently hosted.

Results

The databases are displayed along with information such as database type, host or cluster name, associated resource groups and policies, and status.

You must refresh the resources if the databases are renamed outside of SnapCenter.

- If the database is on a non-NetApp storage system, the user interface displays a Not available for backup message in the Overall Status column.

  You cannot perform data protection operations on the database that is on a non-NetApp storage system.

- If the database is on a NetApp storage system and not protected, the user interface displays a Not protected message in the Overall Status column.
- If the database is on a NetApp storage system and protected, the user interface displays an Available for backup message in the Overall Status column.

If you have enabled an Oracle database authentication, a red padlock icon is shown in the resources view. You must configure database credentials to be able to protect the database or add it to the resource group to perform data protection operations.

Create backup policies for Oracle databases

Before you use SnapCenter to back up Oracle database resources, you must create a backup policy for the resource or the resource group that you want to back up. A backup policy is a set of rules that governs how you manage, schedule, and retain backups. You can also specify the replication, script, and backup type settings. Creating a policy saves time when you want to reuse the policy on another resource or resource group.

Before you begin

- You must have defined your backup strategy.
- You must have prepared for data protection by completing tasks such as installing SnapCenter, adding hosts, discovering databases, and creating storage system connections.
- If you are replicating Snapshots to a mirror or vault secondary storage, the SnapCenter administrator must have assigned the SVMs to you for both the source and destination volumes.
- If you have installed the plug-in as a non-root user, you should manually assign the execute permissions to the prescript and postscript directories.
- Review the SnapMirror active sync specific prerequisites and limitations. For information, refer Object limits for SnapMirror active sync.

About this task

If 'Retain the backup copies for a specific number of days' option is selected, then the SnapLock retention period must be lesser than or equal to the mentioned retention days.

+
Specifying a Snapshot locking period prevents deletion of the Snapshots until the retention period expires. This could lead to retaining a larger number of Snapshots than the count specified in the policy.

+
For ONTAP 9.12.1 and below version, the clones created from the SnapLock Vault Snapshots as part of restore will inherit the SnapLock Vault expiry time. Storage admin should manually cleanup the clones post the SnapLock expiry time.

Steps

1. In the left navigation pane, click **Settings**.
2. In the Settings page, click **Policies**.
3. Select **Oracle Database** from the drop-down list.
4. Click **New**.
5. In the Name page, enter the policy name and details.
6. In the Policy type page, perform the following steps:
   a. Select your storage type.
   b. Select policy scope:
      - If you want to **create an online backup**, select **Online backup**.

        You must specify whether you want to back up all the datafiles, control files, and archive log files, only datafiles and control files, or only archive log files.

      - If you want to **create an offline backup**, select **Offline backup**, and then select one of the following options:
        - If you want to create an offline backup when the database is in mounted state, select **Mount**.
        - If you want to create an offline shutdown backup by changing the database to shutdown state, select **Shutdown**.

          If you are having pluggable databases (PDBs), and want to save the state of the PDBs before creating the backup, you must select **Save state of PDBs**. This enables you to bring the PDBs to their original state after the backup is created.

   c. If you want to catalog backup using Oracle Recovery Manager (RMAN), select **Catalog backup with Oracle Recovery Manager (RMAN)**.

      You can perform deferred cataloging for one backup at a time either using the GUI or using the SnapCenter CLI command Catalog-SmBackupWithOracleRMAN.

      If you want to catalog backups of a RAC database, ensure that no other job is running for that database. If another job is running, the cataloging operation fails instead of getting queued.

   d. If you want to prune archive logs after backup, select **Prune archive logs after backup**.

Pruning of archive logs from the archive log destination that is unconfigured in the database, will be skipped.

If you are using Oracle Standard Edition, you can use LOG_ARCHIVE_DEST and LOG_ARCHIVE_DUPLEX_DEST parameters while performing archive log backup.

- You can delete archive logs only if you have selected the archive log files as part of your backup.

You must ensure that all the nodes in an RAC environment can access all the archive log locations for the delete operation to be successful.

If you want to…	Then…
Delete all archive logs	Select **Delete all archive logs**.
Delete archive logs that are older	Select **Delete archive logs older than**, and then specify the age of the archive logs that are to be deleted in days and hours.
Delete archive logs from all destinations	Select **Delete archive logs from all the destinations**.
Delete the archive logs from the log destinations that are part of the backup	Select **Delete archive logs from the destinations which are part of backup**.

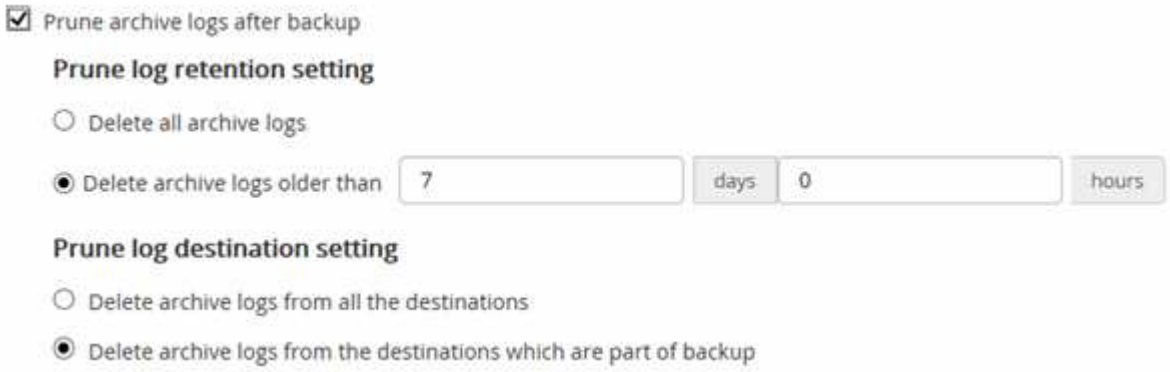

7. In the Snapshot and Replication page, perform the following steps:
   a. Specify the schedule frequency by selecting **On demand**, **Hourly**, **Daily**, **Weekly**, or **Monthly**.

You can specify the schedule (start date and end date) for the backup operation while creating a resource group. This enables you to create resource groups that share the same policy and backup frequency, but enables you to assign different backup schedules to each policy.

If you have scheduled for 2:00 a.m., the schedule will not be triggered during daylight saving time (DST).

   b. In the Data snapshot retention settings section, specify the retention settings for the backup type and

the schedule type selected in the Backup Type page:

If you want to…	Then…
Keep a certain number of Snapshots	Select **Copies to keep**, and then specify the number of Snapshots that you want to keep. If the number of Snapshots exceeds the specified number, the Snapshots are deleted with the oldest copies deleted first. The maximum retention value is 1018. Backups will fail if retention is set to a value higher than what the underlying ONTAP version supports. You must set the retention count to 2 or higher if you plan to enable SnapVault replication. If you set the retention count to 1, the retention operation might fail because the first Snapshot is the reference Snapshot for the SnapVault relationship until a newer Snapshot is replicated to the target.
Keep the Snapshots for a certain number of days	Select **Retain copies for**, and then specify the number of days for which you want to keep the Snapshots before deleting them.
Snapshot copy locking period	Select the **Snapshot copy locking period** and specify the duration in days, months, or years. SnapLock retention period should be less than 100 years.

c. In the Archive Log snapshot retention settings section, specify the retention settings for the backup type and the schedule type selected in the Backup Type page:

If you want to…	Then…

Keep a certain number of Snapshots	Select **Copies to keep**, and then specify the number of Snapshots that you want to keep. If the number of Snapshots exceeds the specified number, the Snapshots are deleted with the oldest copies deleted first. The maximum retention value is 1018. Backups will fail if retention is set to a value higher than what the underlying ONTAP version supports. You must set the retention count to 2 or higher if you plan to enable SnapVault replication. If you set the retention count to 1, the retention operation might fail because the first Snapshot is the reference Snapshot for the SnapVault relationship until a newer Snapshot is replicated to the target.
Keep the Snapshots for a certain number of days	Select **Retain copies for**, and then specify the number of days for which you want to keep the Snapshots before deleting them.
Snapshot copy locking period	Select the **Snapshot copy locking period** and specify the duration in days, months, or years. SnapLock retention period should be less than 100 years.

d. Select policy label.

You can assign SnapMirror labels to primary snapshots for remote replication, allowing the primary snapshots to offload the snapshot replication operation from SnapCenter to ONTAP secondary systems. This can be done without enabling SnapMirror or SnapVault option in the policy page.

8. In the Select secondary replication options section, select one or both of the following secondary replication options:

You must select the secondary replication options for **Secondary snapshot copy locking period** to be effective.

For this field…	Do this…
Update SnapMirror after creating a local Snapshot	Select this field to create mirror copies of the backup sets on another volume (SnapMirror replication). This option should be enabled for SnapMirror active sync. During secondary replication, the SnapLock expiry time loads the primary SnapLock expiry time. Clicking the **Refresh** button in the Topology page refreshes the secondary and primary SnapLock expiry time that are retrieved from ONTAP.
Update SnapVault after creating a local Snapshot	Select this option to perform disk-to-disk backup replication (SnapVault backups). When SnapLock is configured only on the secondary from ONTAP known as SnapLock Vault, clicking the **Refresh** button in the Topology page refreshes the locking period on the secondary that is retrieved from ONTAP. For more information on SnapLock Vault see Commit Snapshot copies to WORM on a vault destination See View Oracle database backups and clones in the Topology page.
Error retry count	Enter the maximum number of replication attempts that can be allowed before the operation stops.

You should configure SnapMirror retention policy in ONTAP for the secondary storage to avoid reaching the maximum limit of Snapshots on the secondary storage.

9. In the Script page, enter the path and the arguments of the prescript or postscript that you want to run before or after the backup operation, respectively.

   You must store the prescripts and postscripts either in */var/opt/snapcenter/spl/scripts* or in any folder inside this path. By default, the */var/opt/snapcenter/spl/scripts* path is populated. If you have created any folders inside this path to store the scripts, you must specify those folders in the path.

   You can also specify the script timeout value. The default value is 60 seconds.

   SnapCenter allows you to use the predefined environment variables when you execute the prescript and postscript. Learn more

10. In the Verification page, perform the following steps:

    a. Select the backup schedule for which you want to perform the verification operation.

b. In the Verification script commands section, enter the path and the arguments of the prescript or postscript that you want to run before or after the verification operation, respectively.

   You must store the prescripts and postscripts either in */var/opt/snapcenter/spl/scripts* or in any folder inside this path. By default, the */var/opt/snapcenter/spl/scripts* path is populated. If you have created any folders inside this path to store the scripts, you must specify those folders in the path.

   You can also specify the script timeout value. The default value is 60 seconds.

11. Review the summary, and then click **Finish**.

Create resource groups and attach policies for Oracle databases

A resource group is a container where you add resources that you want to back up and protect. A resource group allows you to back up all the data that is associated with a given application simultaneously.

About this task

- A database with files in ASM disk groups must be in "MOUNT" or "OPEN" state to verify its backups using the Oracle DBVERIFY utility.

  Attach one or more policies to the resource group to define the type of data protection job you want to perform.

  The following image illustrates the relationship between resources, resource groups, and policies for databases:

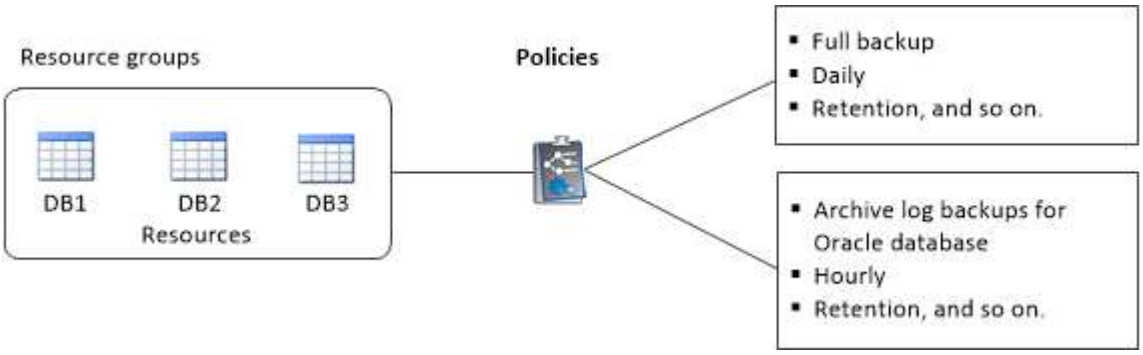

- For SnapLock enabled policies, for ONTAP 9.12.1 and below version, if you specify a Snapshot locking period, the clones created from the tamper proof Snapshots as part of restore will inherit the SnapLock expiry time. Storage admin should manually cleanup the clones post the SnapLock expiry time.
- Adding new databases without SnapMirror active sync to an existing resource group which contains resources with SnapMirror active sync, is not supported.
- Adding new databases to an existing resource group in failover mode of SnapMirror active sync is not supported. You can add resources to the resource group only in regular or fail-back state.

Steps

1. In the left navigation pane, select **Resources**, and the appropriate plug-in from the list.
2. In the Resources page, click **New Resource Group**.
3. In the Name page, perform the following actions:
   a. Enter a name for the resource group in the Name field.

The resource group name should not exceed 250 characters.

b. Enter one or more labels in the Tag field to help you search for the resource group later.

   For example, if you add HR as a tag to multiple resource groups, you can later find all resource groups associated with the HR tag.

c. Select this check box, and enter a custom name format that you want to use for the Snapshot name.

   For example, customtext_resource group_policy_hostname or resource group_hostname. By default, a timestamp is appended to the Snapshot name.

d. Specify the destinations of the archive log files that you do not want to back up.

   You should use the exact same destination as it was set in Oracle, including prefix, if needed.

4. In the Resources page, select an Oracle database host name from the **Host** drop-down list.

   The resources are listed in the Available Resources section only if the resource is discovered successfully. If you have recently added resources, they will appear on the list of available resources only after you refresh your resource list.

5. Select the resources from the Available Resources section and move them to the Selected Resources section.

   You can add databases from both Linux and AIX hosts in a single resource group.

6. In the Application Settings page, select the backup option.
7. In the Policies page, perform the following steps:

   a. Select one or more policies from the drop-down list.

      You can also create a policy by clicking + .

      In the Configure schedules for selected policies section, the selected policies are listed.

   b. Click + in the Configure Schedules column for the policy for which you want to configure a schedule.

   c. In the Add schedules for policy *policy_name* window, configure the schedule, and then click **OK**.

      Where, *policy_name* is the name of the policy that you have selected.

      The configured schedules are listed in the Applied Schedules column.

   Third party backup schedules are not supported when they overlap with SnapCenter backup schedules.

8. On the Verification page, perform the following steps:

   a. Click **Load locators** to load the SnapMirror or SnapVault volumes to perform verification on secondary

storage.

b. Click + in the Configure Schedules column to configure the verification schedule for all the schedule types of the policy.

c. In the Add Verification Schedules policy_name dialog box, perform the following actions:

If you want to…	Do this…
Run verification after backup	Select **Run verification after backup**.
Schedule a verification	Select **Run scheduled verification** and then select the schedule type from the drop-down list.

d. Select **Verify on secondary location** to verify your backups on secondary storage system.

e. Click **OK**.

The configured verification schedules are listed in the Applied Schedules column.

9. In the Notification page, from the **Email preference** drop-down list, select the scenarios in which you want to send the emails.

You must also specify the sender and receiver email addresses, and the subject of the email. If you want to attach the report of the operation performed on the resource group, select **Attach Job Report**.

ⓘ For email notification, you must have specified the SMTP server details using the either the GUI or the PowerShell command Set-SmSmtpServer.

10. Review the summary, and then click **Finish**.

Create resource groups and enable secondary protection for Oracle resources on ASA r2 systems

You should create the resource group to add the resources that are on ASA r2 systems. You can also provision the secondary protection while creating the resource group.

Before you begin

- You should ensure that you are not adding both ONTAP 9.x resources and ASA r2 resources to the same resource group.
- You should ensure that you do not have a database with both ONTAP 9.x resources and ASA r2 resources.

About this task

- The secondary protection is available only if the logged-in user is assigned to the role that has the **SecondaryProtection** capability enabled.
- If you enabled secondary protection, the resource group is put into maintenance mode while creating the primary and secondary consistency groups. After the primary and secondary consistency groups are created, the resource group is put out of maintenance mode.
- SnapCenter does not support secondary protection for a clone resource.

Steps

1. In the left navigation pane, select **Resources**, and the appropriate plug-in from the list.
2. In the Resources page, click **New Resource Group**.
3. In the Name page, perform the following actions:
   a. Enter a name for the resource group in the Name field.

      The resource group name should not exceed 250 characters.

   b. Enter one or more labels in the Tag field to help you search for the resource group later.

      For example, if you add HR as a tag to multiple resource groups, you can later find all resource groups associated with the HR tag.

   c. Select this check box, and enter a custom name format that you want to use for the Snapshot name.

      For example, customtext_resource group_policy_hostname or resource group_hostname. By default, a timestamp is appended to the Snapshot name.

   d. Specify the destinations of the archive log files that you do not want to back up.

      You should use the exact same destination as it was set in the application including prefix, if needed.

4. In the Resources page, select the database host name from the **Host** drop-down list.

   The resources are listed in the Available Resources section only if the resource is discovered successfully. If you have recently added resources, they will appear on the list of available resources only after you refresh your resource list.

5. Select the ASA r2 resources from the Available Resources section and move them to the Selected Resources section.
6. In the Application Settings page, select the backup option.
7. In the Policies page, perform the following steps:
   a. Select one or more policies from the drop-down list.

      You can also create a policy by clicking + .

      In the Configure schedules for selected policies section, the selected policies are listed.

   b. Click + in the Configure Schedules column for the policy for which you want to configure a schedule.
   c. In the Add schedules for policy *policy_name* window, configure the schedule, and then click **OK**.

      Where, *policy_name* is the name of the policy that you have selected.

      The configured schedules are listed in the Applied Schedules column.

Third party backup schedules are not supported when they overlap with SnapCenter backup schedules.

8. If the secondary protection is enabled for the policy that you have selected, then Secondary Protection page is displayed and you need to perform the following steps:
   a. Select the type of the replication policy.

      Synchronous replication policy is not supported.

   b. Specify the consistency group suffix that you want to use.
   c. From the Destination Cluster and Destination SVM drop-downs select the peered cluster and SVM that you want to use.

      The cluster and SVM peering is not supported by SnapCenter. You should use System Manager or ONTAP CLIs to perform cluster and SVM peering.

If the resources are already protected outside of SnapCenter, those resources will be displayed in the Secondary Protected Resources section.

1. On the Verification page, perform the following steps:
   a. Click **Load locators** to load the SnapMirror or SnapVault volumes to perform verification on secondary storage.
   b. Click in the Configure Schedules column to configure the verification schedule for all the schedule types of the policy.
   c. In the Add Verification Schedules policy_name dialog box, perform the following actions:

If you want to…	Do this…
Run verification after backup	Select **Run verification after backup**.
Schedule a verification	Select **Run scheduled verification** and then select the schedule type from the drop-down list.

   d. Select **Verify on secondary location** to verify your backups on secondary storage system.
   e. Click **OK**.

      The configured verification schedules are listed in the Applied Schedules column.

2. In the Notification page, from the **Email preference** drop-down list, select the scenarios in which you want to send the emails.

   You must also specify the sender and receiver email addresses, and the subject of the email. If you want to attach the report of the operation performed on the resource group, select **Attach Job Report**.

   For email notification, you must have specified the SMTP server details using the either the GUI or the PowerShell command Set-SmSmtpServer.

3. Review the summary, and then click **Finish**.

Back up Oracle resources

If a resource is not part of any resource group, you can back up the resource from the Resources page.

Steps

1. In the left navigation pane, select **Resources**, and the appropriate plug-in from the list.
2. In the Resources page, select **Database** from the View list.
3. Click , and then select the host name and the database type to filter the resources.

   You can then click to close the filter pane.

4. Select the database that you want to back up.

   The Database-Protect page is displayed.

5. In the Resources page, perform the following steps:
   a. Select the check box, and enter a custom name format that you want to use for the Snapshot name.

      For example, `customtext_policy_hostname` or `resource_hostname`. A timestamp is appended to the Snapshot name by default.

   b. Specify the destinations of the archive log files that you do not want to back up.
6. In the Policies page, perform the following steps:
   a. Select one or more policies from the drop-down list.

      You can create a policy by clicking .

      In the Configure schedules for selected policies section, the selected policies are listed.

   b. Click in the Configure Schedules column to configure a schedule for the policy you want.
   c. In the Add schedules for policy *policy_name* window, configure the schedule, and then select `OK`.

      policy_name is the name of the policy that you have selected.

      The configured schedules are listed in the Applied Schedules column.

7. In the Verification page, perform the following steps:
   a. Click **Load locators** to load the SnapMirror or SnapVault volumes to verify secondary storage.
   b. Click in the Configure Schedules column to configure the verification schedule for all of the schedule types of the policy.

      In the Add Verification Schedules *policy_name* dialog box, you can perform the following steps:
   c. Select **Run verification after backup**.
   d. Select **Run scheduled verification**, and select the schedule type from the drop-down list.

In a Flex ASM setup, you cannot perform verification operation on Leaf nodes if the cardinality is less than the number nodes in the RAC cluster.

e. Select **Verify on secondary location** to verify your backups on secondary storage.

f. Click **OK**.

The configured verification schedules are listed in the Applied Schedules column.

8. In the Notification page, select the scenarios in which you want to send the emails from the **Email preference** drop-down list.

You must specify the sender and receiver email addresses, and the subject of the email. If you want to attach the report of the backup operation performed on the resource, select **Attach Job Report**.

For email notification, you must have specified the SMTP server details using the either the GUI or the PowerShell command `Set-SmSmtpServer`.

9. Review the summary, and then click **Finish**.

The database topology page is displayed.

10. Click **Back up Now**.
11. In the Backup page, perform the following steps:

a. If you have applied multiple policies to the resource, from the Policy drop-down list, select the policy that you want to use for backup.

If the policy selected for the on-demand backup is associated with a backup schedule, the on-demand backups will be retained based on the retention settings specified for the schedule type.

b. Click **Backup**.

12. Monitor the operation progress by clicking **Monitor** > **Jobs**.

After you finish

- In AIX setup, you can use the `lkdev` command to lock and the `rendev` command to rename the disks on which the database that was backed up was residing.

Locking or renaming of devices will not affect the restore operation when you restore using that backup.

- If the backup operation fails because database query execution time exceeded the timeout value, you should change the value of the ORACLE_SQL_QUERY_TIMEOUT and ORACLE_PLUGIN_SQL_QUERY_TIMEOUT parameters by running the `Set-SmConfigSettings` cmdlet:

After modifying the value of the parameters, restart the SnapCenter Plug-in Loader (SPL) service by running the following command `/opt/NetApp/snapcenter/spl/bin/spl restart`

- If the file is not accessible and the mount point is unavailable during the verification process, the operation might fail with error code DBV-00100 specified file. You should modify the values of the VERIFICATION_DELAY and VERIFICATION_RETRY_COUNT parameters in sco.properties.

After modifying the value of the parameters, restart the SnapCenter Plug-in Loader (SPL) service by running the following command `/opt/NetApp/snapcenter/spl/bin/spl restart`

- In MetroCluster configurations, SnapCenter might not be able to detect a protection relationship after a failover.
- If you are backing up application data on VMDKs and the Java heap size for the SnapCenter Plug-in for VMware vSphere is not large enough, the backup might fail.

  To increase the Java heap size, locate the script file */opt/netapp/init_scripts/scvservice*. In that script, the `do_start method` command starts the SnapCenter VMware plug-in service. Update that command to the following: `Java -jar -Xmx8192M -Xms4096M`.

Find more information

- Unable to detect SnapMirror or SnapVault relationship after MetroCluster failover
- Oracle RAC One Node database is skipped for performing SnapCenter operations
- Failed to change the state of an Oracle 12c ASM database
- Customizable parameters for backup, restore and clone operations on AIX systems (Requires login)

Back up Oracle database resource groups

A resource group is a collection of resources on a host or cluster. The backup operation is performed on all resources defined in the resource group.

You can back up a resource group on demand from the Resources page. If a resource group has a policy attached and a schedule configured, then backups are created according to the schedule.

Steps

1. In the left navigation pane, select **Resources**, and the appropriate plug-in from the list.
2. In the Resources page, select **Resource Group** from the **View** list.
3. Enter the resource group name in the search box, or click , and select the tag.

   Click to close the filter pane.

4. In the Resource Group page, select the resource group to back up.

   If you have a federated resource group with two databases and one has data on non-NetApp storage, the backup operation is aborted even though the other database is on NetApp storage.

5. In the Backup page, perform the following steps:
   a. If you have multiple policies associated with the resource group, select the backup policy you want to use from the **Policy** drop-down list.

      If the policy selected for the on-demand backup is associated with a backup schedule, the on-demand backups will be retained based on the retention settings specified for the schedule type.

   b. Select **Backup**.
6. Monitor the progress by selecting **Monitor > Jobs**.

After you finish

- In AIX setup, you can use the `lkdev` command to lock and the `rendev` command to rename the disks on which the database that was backed up was residing.

Locking or renaming of devices will not affect the restore operation when you restore using that backup.

- If the backup operation fails because database query execution time exceeded the timeout value, you should change the value of the ORACLE_SQL_QUERY_TIMEOUT and ORACLE_PLUGIN_SQL_QUERY_TIMEOUT parameters by running the `Set-SmConfigSettings` cmdlet:

  After modifying the value of the parameters, restart the SnapCenter Plug-in Loader (SPL) service by running the following command `/opt/NetApp/snapcenter/spl/bin/spl restart`

- If the file is not accessible and the mount point is unavailable during the verification process, the operation might fail with error code DBV-00100 specified file. You should modify the values of the VERIFICATION_DELAY_ and VERIFICATION_RETRY_COUNT parameters in sco.properties.

  After modifying the value of the parameters, restart the SnapCenter Plug-in Loader (SPL) service by running the following command `/opt/NetApp/snapcenter/spl/bin/spl restart`

Monitor Oracle database backup

Learn how to monitor the progress of backup operations and data protection operations.

Monitor Oracle database backup operations

You can monitor the progress of different backup operations by using the SnapCenterJobs page. You might want to check the progress to determine when it is complete or if there is an issue.

About this task

The following icons appear on the Jobs page and indicate the corresponding state of the operations:

- In progress
- Completed successfully
- Failed
- Completed with warnings or could not start due to warnings
- Queued
- Canceled

Steps

1. In the left navigation pane, click **Monitor**.
2. In the Monitor page, click **Jobs**.
3. In the Jobs page, perform the following steps:
   a. Click to filter the list so that only backup operations are listed.
   b. Specify the start and end dates.
   c. From the **Type** drop-down list, select **Backup**.
   d. From the **Status** drop-down, select the backup status.
   e. Click **Apply** to view the operations completed successfully.

4. Select a backup job, and then click **Details** to view the job details.

   Though the backup job status displays ✔, when you click on job details you might see that some of the child tasks of the backup operation are still in progress or marked with warning signs.

5. In the Job Details page, click **View logs**.

   The **View logs** button displays the detailed logs for the selected operation.

Monitor data protection operations in the Activity pane

The Activity pane displays the five most recent operations performed. The Activity pane also displays when the operation was initiated and the status of the operation.

The Activity pane displays information regarding backup, restore, clone, and scheduled backup operations.

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. Click on the Activity pane to view the five most recent operations.

   When you click one of the operations, the operation details are listed in the **Job Details** page.

Other back up operations

Back up Oracle databases using UNIX commands

The backup workflow includes planning, identifying the resources for backup, creating backup policies, creating resource groups and attaching policies, creating backups, and monitoring the operations.

What you will need

- You should have added the storage system connections and created the credential using the commands *Add-SmStorageConnection* and *Add-SmCredential*.
- You should have established the connection session with the SnapCenter Server using the command *Open-SmConnection*.

  You can have only one SnapCenter account login session and the token is stored in the user home directory.

  The connection session is valid only for 24 hours. However, you can create a token with the TokenNeverExpires option to create a token that never expires and the session will always be valid.

About this task

You should execute the following commands to establish the connection with the SnapCenter Server, discover the Oracle database instances, add policy and resource group, backup and verify the backup.

The information regarding the parameters that can be used with the command and their descriptions can be obtained by running Get-Help *command_name*. Alternatively, you can also refer to the SnapCenter Software Command Reference Guide.

Steps

1. Initiate a connection session with the SnapCenter Server for a specified user: *Open-SmConnection*
2. Perform host resources discovery operation: *Get-SmResources*
3. Configure Oracle database credentials and preferred nodes for backup operation of a Real Application Cluster (RAC) database: *Configure-SmOracleDatabase*
4. Create a backup policy: *Add-SmPolicy*
5. Retrieve the information about the secondary (SnapVault or SnapMirror) storage location : *Get-SmSecondaryDetails*

   This command retrieves the primary to secondary storage mapping details of a specified resource. You can use the mapping details to configure the secondary verification settings while creating a backup resource group.

6. Add a resource group to SnapCenter: *Add-SmResourceGroup*
7. Create a backup: *New-SmBackup*

   You can poll the job using the WaitForCompletion option. If this option is specified, then the command continues to poll the server until the completion of the backup job.

8. Retrieve the logs from SnapCenter: *Get-SmLogs*

Cancel backup operations of Oracle databases

You can cancel backup operations that are either running, queued, or non-responsive.

You must be logged in as the SnapCenter Admin or job owner to cancel backup operations.

About this task

When you cancel a backup operation, the SnapCenter Server stops the operation and removes all the Snapshots from the storage if the backup created is not registered with SnapCenter Server. If the backup is already registered with SnapCenter Server, it will not roll back the already created Snapshot even after the cancellation is triggered.

- You can cancel only the log or full backup operation that are queued or running.
- You cannot cancel the operation after the verification has started.

  If you cancel the operation before verification, the operation is canceled, and the verification operation will not be performed.

- You cannot cancel the backup operation after the catalog operations has started.
- You can cancel a backup operation from either the Monitor page or the Activity pane.
- In addition to using the SnapCenter GUI, you can use CLI commands to cancel operations.
- The **Cancel Job** button is disabled for operations that cannot be canceled.
- If you selected **All members of this role can see and operate on other members objects** in

Users\Groups page while creating a role, you can cancel the queued backup operations of other members while using that role.

Step

Perform one of the following actions:

From the…	Action
Monitor page	a. In the left navigation pane, click **Monitor** > **Jobs**. b. Select the operation and click **Cancel Job**.
Activity pane	a. After initiating the backup job, click on the Activity pane to view the five most recent operations. b. Select the operation. c. In the Job Details page, click **Cancel Job**.

Results

The operation is canceled, and the resource is reverted to the original state.

If the operation you canceled is non-responsive in the canceling or running state, you should run the Cancel-SmJob -JobID <int> -Force to forcefully stop the backup operation.

View Oracle database backups and clones in the Topology page

When you are preparing to back up or clone a resource, you might find it helpful to view a graphical representation of all backups and clones on the primary and secondary storage.

About this task

In the Topology page, you can see all of the backups and clones that are available for the selected resource or resource group. You can view the details of those backups and clones, and then select them to perform data protection operations.

You can review the following icons in the Manage Copies view to determine whether the backups and clones are available on the primary or secondary storage (Mirror copies or Vault copies).

- displays the number of backups and clones that are available on the primary storage.
- displays the number of backups and clones that are mirrored on the secondary storage using SnapMirror technology.
- displays the number of backups and clones that are replicated on the secondary storage using SnapVault technology.

The number of backups displayed includes the backups deleted from the secondary storage. For example, if you created 6 backups using a policy to retain only 4 backups, the number of backups displayed is 6.

Clones of a backup of a version-flexible mirror on a mirror-vault type volume are displayed in the topology view, but the mirror backup count in the topology view does not include the version-flexible backup.

If you have secondary relationship as SnapMirror active sync (initially released as SnapMirror Business Continuity [SM-BC]), you can see following additional icons:

- The replica site is up.
- The replica site is down.
- The secondary mirror or vault relationship has not been re-established.

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page, either select the resource or resource group from the **View** drop-down list.
3. Select the resource either from the resource details view or from the resource group details view.

   If the resource is protected, the Topology page of the selected resource is displayed.

4. Review the Summary card to see a summary of the number of backups and clones available on the primary and secondary storage.

   The Summary Card section displays the total number of backups and clones and total number of log backups.

   Clicking the **Refresh** button starts a query of the storage to display an accurate count.

   If SnapLock enabled backup is taken, then clicking the **Refresh** button refreshes the primary and secondary SnapLock expiry time retrieved from ONTAP. A weekly schedule also refreshes the primary and secondary SnapLock expiry time retrieved from ONTAP.

   When the application resource is spread across multiple volumes, the SnapLock expiry time for the backup will be the longest SnapLock expiry time that is set for a Snapshot in a volume. The longest SnapLock expiry time is retrieved from ONTAP.

   For SnapMirror active sync, clicking the **Refresh** button refreshes the SnapCenter backup inventory by querying ONTAP for both primary and replica sites. A weekly schedule also performs this activity for all databases containing SnapMirror active sync relationship.

   - For SnapMirror active sync and only for ONTAP 9.14.1, Async Mirror or Async MirrorVault relationships to the new primary destination should be manually configured after failover. From ONTAP 9.15.1 onwards Async Mirror or Async MirrorVault is auto configured to the new primary destination.
   - After failover, a backup should be created for SnapCenter to be aware of the failover. You can click **Refresh** only after a backup has been created.

5. In the Manage Copies view, click **Backups** or **Clones** from the primary or secondary storage to see details

of a backup or clone.

The details of the backups and clones are displayed in a table format.

6. Select the backup from the table, and then click the data protection icons to perform restore, clone, mount, unmount, rename, catalog, uncatalog, and delete operations.

   You cannot rename or delete backups that are on the secondary storage.

   - If you have selected a log backup, you can only perform rename, mount, unmount, catalog, uncatalog, and delete operations.
   - If you have cataloged the backup using Oracle Recovery Manager (RMAN), you cannot rename those cataloged backups.

7. If you want to delete a clone, select the clone from the table, and then click .

If the value assigned to SnapmirrorStatusUpdateWaitTime is less, the Mirror and Vault backup copies are not listed on the topology page even if data and log volumes are successfully protected. You should increase the value assigned to SnapmirrorStatusUpdateWaitTime using *Set-SmConfigSettings* PowerShell cmdlet.

The information regarding the parameters that can be used with the command and their descriptions can be obtained by running Get-Help *command_name*.

Alternatively, you can also refer to the SnapCenter Software Command Reference Guide or SnapCenter Software Cmdlet Reference Guide.

Mount and unmount database backups

You can mount a single or multiple data and log only backups if you want to access the files in the backup. You can either mount the backup to the same host where the backup was created or to a remote host having same type of Oracle and host configurations. If you have manually mounted the backups, you should manually unmount the backups after completing the operation. At any given instance, a backup of a database can be mounted to any one of the host. While performing an operation, you can mount only a single backup.

In a Flex ASM setup, you cannot perform mount operation on Leaf nodes if the cardinality is less than the number nodes in the RAC cluster.

Mount a database backup

You should manually mount a database backup if you want to access the files in the backup.

What you will need

- If you have an Automatic Storage Management (ASM) database instance in an NFS environment and want to mount the ASM backups, you should have added the ASM disk path */var/opt/snapcenter/sco/backup*/*/*/*/** to the existing path defined in the asm_diskstring parameter.
- If you have an ASM database instance in an NFS environment and want to mount the ASM log backups as part of a recovery operation, you should have added the ASM disk path */var/opt/snapcenter/scu/clones/*/*_* to the existing path defined in the asm_diskstring parameter.

- In the asm_diskstring parameter, you should configure *AFD:** if you are using ASMFD or configure *ORCL:** if you are using ASMLIB.

  For information on how to edit the asm_diskstring parameter, see How to add disk paths to asm_diskstring.

- You should configure the ASM credentials and the ASM port if it differs from that of the source database host while mounting the backup.
- If you want to mount to an alternate host, you must verify that the alternate host meets the following requirements:
  - Same UID and GID as that of the original host
  - Same Oracle version as that of the original host
  - Same OS distribution and version as that of the original host
  - For NVMe, NVMe util should be installed
- You should ensure that the LUN is not mapped to the AIX host using iGroup consisting of mixed protocols iSCSI and FC. For more information, see Operation fails with error Unable to discover the device for LUN.

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page, select either **Database** or **Resource Group** from the **View** list.
3. Select the database either from the database details view or from the resource group details view.

   The database topology page is displayed.

4. From the Manage Copies view, select **Backups** either from the primary or secondary (mirrored or replicated) storage system.
5. Select the backup from the table, and then click .
6. In the Mount backups page, select the host on which you want to mount the backup from the **Choose the host to mount the backup** drop-down list.

   The mount path */var/opt/snapcenter/sco/backup_mount/backup_name/database_name* is displayed.

   If you are mounting the backup of an ASM database, the mount path +diskgroupname_SID_backupid is displayed.

7. Click **Mount**.

After you finish

- You can run the following command to retrieve the information related to the mounted backup:

  ```
  ./sccli Get-SmBackup -BackupName backup_name -ListMountInfo
  ```

- If you have mounted an ASM database, you can run the following command to retrieve the information related to the mounted backup:

  ```
  ./sccli Get-Smbackup -BackupNamediskgroupname_SID_backupid-listmountinfo
  ```

- To retrieve the backup ID, run the following command:

  ```
  ./sccli Get-Smbackup-BackupNamebackup_name
  ```

  The information regarding the parameters that can be used with the command and their descriptions can be obtained by running Get-Help *command_name*.
  Alternatively, you can also refer to the SnapCenter Software Command Reference Guide.

Unmount a database backup

You can manually unmount a mounted database backup when you no longer want to access files on the backup.

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page, select either **Database** or **Resource Group** from the **View** list.
3. Select the database either from the database details view or from the resource group details view.

   The database topology page is displayed.

4. Select the backup that is mounted, and then click .
5. Click **OK**.

Restore and recover Oracle databases

Restore workflow

The restore workflow includes planning, performing the restore operations, and monitoring the operations.

The following workflow shows the sequence in which you must perform the restore operation:

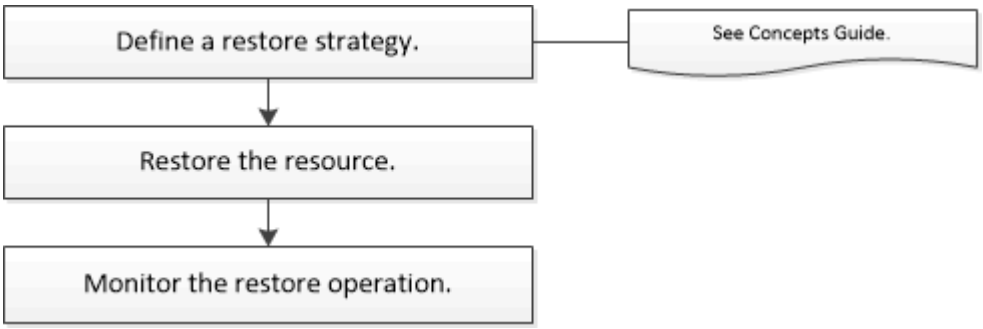

Define a restore and recovery strategy for Oracle databases

You must define a strategy before you restore and recover your database so that you can perform restore and recover operations successfully.

Types of backups supported for restore and recovery operations

SnapCenter supports restore and recovery of different types of Oracle database backups.

- Online data backup
- Offline shutdown data backup
- Offline mount data backup

> If you are restoring an offline shutdown or offline mount data backup, SnapCenter leaves the database in offline state. You should manually recover the database and reset the logs.

- Full backup
- Offline-mount backups of Data Guard standby databases
- Data-only online backups of Active Data Guard standby databases

> You cannot perform recovery of Active Data Guard standby databases.

- Online data backups, online full backups, offline mount backups, and offline shutdown backups in a Real Application Clusters (RAC) configuration
- Online data backups, online full backups, offline mount backups, and offline shutdown backups in an Automatic Storage Management (ASM) configuration

Types of restore methods supported for Oracle databases

SnapCenter supports connect-and-copy or in-place restore for Oracle databases. During a restore operation, SnapCenter determines the restore method that is appropriate for the file system to be used for restore without any data loss.

> SnapCenter does not support volume-based SnapRestore.

Connect-and-copy restore

If the database layout differs from the backup or if there are any new files after the backup was created, connect-and-copy restore is performed. In the connect-and-copy restore method, the following tasks are performed:

Steps

1. The volume is cloned from the Snapshot and the file system stack is built on the host using the cloned LUNs or volumes.
2. The files are copied from the cloned file systems to the original file systems.
3. The cloned file systems are then unmounted from the host and the cloned volumes are deleted from ONTAP.

> For a Flex ASM setup (where the cardinality is less than the number nodes in the RAC cluster) or ASM RAC databases on VMDK or RDM, only connect-and-copy restore method is supported.

Even if you have forcefully enabled in-place restore, SnapCenter performs connect-and-copy restore in the following scenarios:

- Restore from secondary storage system
- Restore of ASM disk groups present on nodes of an Oracle RAC setup on which database instance is not configured
- In Oracle RAC setup, on any of the peer nodes if the ASM instance or the cluster instance is not running or if the peer node is down
- Restore of control files only
- Restore a subset of tablespaces residing on a ASM disk group
- Disk group is shared between data files, sp file, and password file
- SnapCenter Plug-in Loader (SPL) service is not installed or not running on the remote node in a RAC environment
- New nodes are added to the Oracle RAC and the SnapCenter Server is not aware of the newly added nodes

In-place restore

If the database layout is similar to the backup and has not undergone any configuration change on the storage and database stack, in-place restore is performed, wherein the restore of file or LUN is performed on ONTAP. SnapCenter supports only Single File SnapRestore (SFSR) as part of the in-place restore method.

NetApp ONTAP supports in-place restore from secondary location.

If you want to perform in-place restore on the database, ensure that you have only datafiles on the ASM disk group. You must create a backup after any changes are made to the ASM disk group or in the physical structure of the database. After performing in-place restore, the disk group will contain the same number datafiles as at the time of backup.

The in-place restore will be applied automatically when disk group or mount point matches the following criteria:

- No new datafiles are added after backup (foreign file check)
- No addition, deletion, or recreation of ASM disk or LUN after backup (ASM disk group structural change check)
- No addition, deletion, or recreation of LUNs to LVM disk group (LVM disk group structural change check)

You can also forcefully enable in-place restore either using GUI, SnapCenter CLI, or PowerShell cmdlet to override the foreign file check and LVM disk group structural change check.

Performing In-place restore on ASM RAC

In SnapCenter, the node on which you perform restore is termed as primary node and all other nodes of the RAC on which ASM disk group resides are called peer nodes. SnapCenter changes the state of ASM disk group to dismount on all the nodes where the ASM disk group is in mount state before performing the storage restore operation. After the storage restore is complete, SnapCenter changes the state of ASM disk group as it was before the restore operation.

In SAN environments, SnapCenter removes devices from all the peer nodes and performs LUN unmap operation before storage restore operation. After storage restore operation, SnapCenter performs LUN map operation and constructs devices on all the peer nodes. In a SAN environment if the Oracle RAC ASM layout is residing on LUNs, then while restoring SnapCenter performs LUN unmap, LUN restore, and LUN map

operations on all the nodes of the RAC cluster where the ASM disk group resides. Before restoring even if all the initiators of the RAC nodes were not used for the LUNs, after restoring SnapCenter creates a new iGroup with all the initiators of all the RAC nodes.

- If there is any failure during prerestore activity on peer nodes, SnapCenter automatically rolls back the ASM disk group state as it was before performing restore on peer nodes on which prerestore operation was successful. Rollback is not supported for the primary and the peer node on which the operation failed. Before attempting another restore you must manually fix the issue on the peer node and bring the ASM disk group on the primary node back to mount state.
- If there is any failure during restore activity, then the restore operation fails and no roll back is performed. Before attempting another restore, you must manually fix the storage restore issue and bring the ASM disk group on the primary node back to mount state.
- If there is any failure during postrestore activity on any of the peer nodes, SnapCenter continues with the restore operation on the other peer nodes. You must manually fix the post restore issue on the peer node.

Types of restore operations supported for Oracle databases

SnapCenter enables you to perform different types of restore operations for Oracle databases.

Before restoring the database, backups are validated to identify whether any files are missing when compared to the actual database files.

Full restore

- Restores only the datafiles
- Restores only the control files
- Restores the datafiles and control files
- Restores datafiles, control files, and redo log files in Data Guard standby and Active Data Guard standby databases

Partial restore

- Restores only the selected tablespaces
- Restores only the selected pluggable databases (PDBs)
- Restores only the selected tablespaces of a PDB

Types of recovery operations supported for Oracle databases

SnapCenter enables you to perform different types of recovery operations for Oracle databases.

- The database up to the last transaction (all logs)
- The database up to a specific system change number (SCN)
- The database up to a specific date and time

  You must specify the date and time for recovery based on the database host's time zone.

  SnapCenter also provides the No recovery option for Oracle databases.

The plug-in for Oracle database does not support recovery if you have restored using a backup that was created with the database role as standby. You must always perform manual recovery for physical standby databases.

Limitations related to restore and recovery of Oracle databases

Before you perform restore and recovery operations, you must be aware of the limitations.

If you are using any version of Oracle from 11.2.0.4 to 12.1.0.1, the restore operation will be in hung state when you run the *renamedg* command. You can apply the Oracle patch 19544733 to fix this issue.

The following restore and recovery operations are not supported:

- Restore and recovery of tablespaces of the root container database (CDB)
- Restore of temporary tablespaces and temporary tablespaces associated with PDBs
- Restore and recovery of tablespaces from multiple PDBs simultaneously
- Restore of log backups
- Restore of backups to a different location
- Restore of redo log files in any configuration other than Data Guard standby or Active Data Guard standby databases
- Restore of SPFILE and Password file
- When you perform a restore operation on a database that was re-created using the preexisting database name on the same host, was managed by SnapCenter, and had valid backups, the restore operation overwrites the newly created database files even though the DBIDs are different.

  This can be avoided by performing either of following actions:

  - Discover the SnapCenter resources after the database is re-created
  - Create a backup of the re-created database

Limitations related to point-in-time recovery of tablespaces

- Point-in-time recovery (PITR) of SYSTEM, SYSAUX, and UNDO tablespaces is not supported
- PITR of tablespaces cannot be performed along with other types of restore
- If a tablespace is renamed and you want to recover it to a point before it was renamed, you should specify the earlier name of the tablespace
- If constraints for the tables in one tablespace are contained in another tablespace, you should recover both the tablespaces
- If a table and its indexes are stored in different tablespaces, then the indexes should be dropped before performing PITR
- PITR cannot be used to recover the current default tablespace
- PITR cannot be used to recover tablespaces containing any of the following objects:
  - Objects with underlying objects (such as materialized views) or contained objects (such as partitioned tables) unless all the underlying or contained objects are in the recovery set

    Additionally, if the partitions of a partitioned table are stored in different tablespaces, then you should either drop the table before performing PITR or move all the partitions to the same tablespace before

performing PITR.

- Undo or rollback segments
- Oracle 8 compatible advanced queues with multiple recipients
- Objects owned by the SYS user

  Examples of these types of objects are PL/SQL, Java classes, call out programs, views, synonyms, users, privileges, dimensions, directories, and sequences.

Sources and destinations for restoring Oracle databases

You can restore an Oracle database from a backup copy on either primary storage or secondary storage. You can only restore databases to the same location on the same database instance. However, in Real Application Cluster (RAC) setup, you can restore databases to other nodes.

Sources for restore operations

You can restore databases from a backup on primary storage or secondary storage. If you want to restore from a backup on the secondary storage in a multiple mirror configuration, you can select the secondary storage mirror as the source.

Destinations for restore operations

You can only restore databases to the same location on the same database instance.

In a RAC setup, you can restore RAC databases from any nodes in the cluster.

Predefined environment variables for restore specific prescript and postscript

SnapCenter allows you to use the predefined environment variables when you execute the prescript and postscript while restoring a database.

Supported predefined environment variables for restoring a database

- **SC_JOB_ID** specifies the job ID of the operation.

  Example: 257

- **SC_ORACLE_SID** specifies the system identifier of the database.

  If the operation involves multiple databases, this will contain database names separated by pipe.

  Example: NFSB31

- **SC_HOST** specifies the host name of the database.

  This parameter will be populated for application volumes.

  Example: scsmohost2.gdl.englabe.netapp.com

- **SC_OS_USER** specifies the operating system owner of the database.

  Example: oracle

- **SC_OS_GROUP** specifies the operating system group of the database.

  Example: oinstall

- **SC_BACKUP_NAME** specifies the name of the backup.

  This parameter will be populated for application volumes.

  Examples:

  - If the database is not running in ARCHIVELOG mode: DATA@RG2_scspr2417819002_07-20-2021_12.16.48.9267_0|LOG@RG2_scspr2417819002_07-20-2021_12.16.48.9267_1
  - If the database is running in ARCHIVELOG mode: DATA@RG2_scspr2417819002_07-20-2021_12.16.48.9267_0|LOG@RG2_scspr2417819002_07-20-2021_12.16.48.9267_1,RG2_scspr2417819002_07-21-2021_12.16.48.9267_1,RG2_scspr2417819002_07-22-2021_12.16.48.9267_1

- **SC_BACKUP_ID** specifies the ID of the backup.

  This parameter will be populated for application volumes.

  Examples:

  - If the database is not running in ARCHIVELOG mode: DATA@203|LOG@205
  - If the database is running in ARCHIVELOG mode: DATA@203|LOG@205,206,207

- **SC_RESOURCE_GROUP_NAME** specifies the name of the resource group.

  Example: RG1

- **SC_ORACLE_HOME** specifies the path of the Oracle home directory.

  Example: /ora01/app/oracle/product/18.1.0/db_1

- **SC_RECOVERY_TYPE** specifies the files that are recovered and also the recovery scope.

  Example: RESTORESCOPE:usingBackupControlfile=false|RECOVERYSCOPE:allLogs=true,noLogs=false,untiltime=false,untilscn=false.

For information about delimiters, see Supported delimiters.

Requirements for restoring an Oracle database

Before restoring an Oracle database, you should ensure that prerequisites are completed.

- You should have defined your restore and recovery strategy.
- The SnapCenter administrator should have assigned you the storage virtual machines (SVMs) for both the source volumes and destination volumes if you are replicating Snapshots to a mirror or vault.
- If archive logs are pruned as part of backup, you should have manually mounted the required archive log backups.
- If you want to restore Oracle databases that are residing on a Virtual Machine Disk (VMDK), you should

ensure that the guest machine has the required number of free slots for allocating the cloned VMDKs.

- You should ensure that all data volumes and archive log volumes belonging to the database are protected if secondary protection is enabled for that database.
- You should ensure that the RAC One Node database is in "nomount" state to perform control file or full database restore.
- If you have an ASM database instance in NFS environment, you should add the ASM disk path */var/opt/snapcenter/scu/clones/*/** to the existing path defined in the asm_diskstring parameter to successfully mount the ASM log backups as part of recovery operation.
- In the asm_diskstring parameter, you should configure *AFD:** if you are using ASMFD or configure *ORCL:** if you are using ASMLIB.

> For information on how to edit the asm_diskstring parameter, see How to add disk paths to asm_diskstring

- You should configure the static listener in the **listener.ora** file available at *$ORACLE_HOME/network/admin* for non ASM databases and *$GRID_HOME/network/admin* for ASM databases if you have disabled OS authentication and enabled Oracle database authentication for an Oracle database, and want to restore the datafiles and control files of that database.
- You should increase value of SCORestoreTimeout parameter by running the Set- SmConfigSettings command if the database size is in terabytes (TB).
- You should ensure that all the licenses required for vCenter are installed and up to date.

  If the licenses are not installed or up to date, a warning message is displayed. If you ignore the warning and proceed, restore from RDM fails.

- You should ensure that the LUN is not mapped to the AIX host using iGroup consisting of mixed protocols iSCSI and FC. For more information, see Operation fails with error Unable to discover the device for LUN.

Restore and recover Oracle database

In the event of data loss, you can use SnapCenter to restore data from one or more backups to your active file system and then recover the database.

Before you begin

If you have installed the plug-in as a non-root user, you should manually assign the execute permissions to the prescript and postscript directories.

About this task

- Recovery is performed using the archive logs available at the configured archive log location. If the database is running in ARCHIVELOG mode, Oracle database saves the filled groups of redo log files to one or more offline destinations, known collectively as the archived redo log. SnapCenter identifies and mounts optimal number of log backups based on the specified SCN, selected date and time, or all logs option.
  If the archive logs required for recovery are not available at the configured location, you should mount the Snapshot containing the logs and specify the path as external archive logs.

  If you migrate ASM database from ASMLIB to ASMFD, then the backups created with ASMLIB cannot be used to restore the database. You should create backups in the ASMFD configuration and use those backups to restore. Similarly, if ASM database is migrated from ASMFD to ASMLIB, you should create

backups in the ASMLIB configuration to restore.

When you restore a database, an operational lock file (.sm_lock_dbsid) is created on the Oracle database host in the */var/opt/snapcenter/sco/lock* directory to avoid multiple operations being executed on the database. After the database has been restored, the operational lock file is automatically removed.

Restore of SPFILE and Password file is not supported.

- For SnapLock enabled policies, for ONTAP 9.12.1 and below version, if you specify a Snapshot locking period, the clones created from the tamper proof Snapshots as part of restore will inherit the SnapLock expiry time. Storage admin should manually cleanup the clones post the SnapLock expiry time.

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page, select either **Database** or **Resource Group** from the **View** list.
3. Select the database from either the database details view or the resource group details view.

   The database topology page is displayed.

4. From the Manage Copies view, select **Backups** from either the primary or the secondary (mirrored or replicated) storage systems.
5. Select the backup from the table, and then click .
6. In the Restore Scope page, perform the following tasks:
   a. If you have selected a backup of a database in a Real Application Clusters (RAC) environment, select the RAC node.
   b. When you select a mirrored or vault data:
      - if there are no log backup at mirror or vault, nothing is selected and the locators are empty.
      - if log backups exist in mirror or vault, the latest log backup is selected and corresponding locator is displayed.

        If the selected log backup exists in both mirror and vault location, both the locators are displayed.

   c. Perform the following actions:

If you want to restore…	Do this…
All the datafiles of the database	Select **All Datafiles**. Only the datafiles of the database are restored. The control files, archive logs, or redo log files are not restored.
Tablespaces	Select **Tablespaces**. You can specify the tablespaces that you want to restore.

If you want to restore…	Do this…
Control files	Select **Control files**. While restoring control files, ensure that the directory structure either exists or should be created with the correct user and group ownerships, if any, to allow the files to be copied to the target location by the restore process. If the directory does not exist, the restore job will fail.
Redo log files	Select **Redo log files**. This option is available only for Data Guard standby or Active Data Guard standby databases. Redo log files are not backed up for non Data Guard databases. For non Data Guard databases the recovery is performed using archive logs.
Pluggable databases (PDBs)	Select **Pluggable databases**, and then specify the PDBs that you want to restore.
Pluggable database (PDB) tablespaces	Select **Pluggable database (PDB) tablespaces**, and then specify the PDB and the tablespaces of that PDB that you want to restore. This option is available only if you have selected a PDB for restore.

d. Select **Change database state if needed for restore and recovery** to change the state of the database to the state required to perform restore and recovery operations.

   The various states of a database from higher to lower are open, mounted, started, and shutdown. You must select this check box if the database is in a higher state but the state must be changed to a lower state to perform a restore operation. If the database is in a lower state but the state must be changed to a higher state to perform the restore operation, the database state is changed automatically even if you do not select the check box.

   If a database is in the open state, and for restore the database needs to be in the mounted state, then the database state is changed only if you select this check box.

e. Select **Force in place restore** if you want to perform in-place restore in the scenarios where new datafiles are added after backup or when LUNs are added, deleted, or re-created to an LVM disk group.

7. In the Recovery Scope page, perform the following actions:

If you…	Do this…
Want to recover to the last transaction	Select **All Logs**.
Want to recover to a specific System Change Number (SCN)	Select **Until SCN (System Change Number)**.
Want to recover to a specific data and time	Select **Date and Time**. You must specify the date and time of the database host's time zone.
Do not want to recover	Select **No recovery**.
Want to specify any external archive log locations	If the database is running in ARCHIVELOG mode, SnapCenter identifies and mounts optimal number of log backups based on the specified SCN, selected date and time, or all logs option. If you still want to specify the location of the external archive log files, select **Specify external archive log locations**. If archive logs are pruned as part of backup, and you have manually mounted the required archive log backups, you must specify the mounted backup path as the external archive log location for recovery. ⓘ You should verify the path and contents of the mount path before listing it as an external log location. • Oracle data protection with ONTAP • Operation fails with ORA-00308 error

You cannot perform restore with recovery from secondary backups if archive log volumes are not protected but data volumes are protected. You can restore only by selecting **No recovery**.

If you are recovering a RAC database with the open database option selected, only the RAC instance where the recovery operation was initiated is brought back to the open state.

ⓘ Recovery is not supported for Data Guard standby and Active Data Guard standby databases.

8. In the PreOps page, enter the path and the arguments of the prescript that you want to run before the restore operation.

   You must store the prescripts either in the */var/opt/snapcenter/spl/scripts* path or in any folder inside this path. By default, the */var/opt/snapcenter/spl/scripts* path is populated. If you have created any folders inside this path to store the scripts, you must specify those folders in the path.

You can also specify the script timeout value. The default value is 60 seconds.

SnapCenter allows you to use the predefined environment variables when you execute the prescript and postscript. Learn more

9. In the PostOps page, perform the following steps:
   a. Enter the path and the arguments of the postscript that you want to run after the restore operation.

      You must store the postscripts either in */var/opt/snapcenter/spl/scripts* or in any folder inside this path. By default, the */var/opt/snapcenter/spl/scripts* path is populated. If you have created any folders inside this path to store the scripts, you must specify those folders in the path.

      If the restore operation fails, postscripts will not be executed and cleanup activities will be triggered directly.

   b. Select the check box if you want to open the database after recovery.

      After restoring a container database (CDB) with or without control files, or after restoring only CDB control files, if you specify to open the database after recovery, then only the CDB is opened and not the pluggable databases (PDB) in that CDB.

      In a RAC setup, only the RAC instance that is used for recovery is opened after recovery.

      After restoring a user tablespace with control files, a system tablespace with or without control files, or a PDB with or without control files, only the state of the PDB related to the restore operation is changed to the original state. The state of the other PDBs that were not used for restore are not changed to the original state because the state of those PDBs were not saved. You must manually change the state of the PDBs that were not used for restore.

10. In the Notification page, from the **Email preference** drop-down list, select the scenarios in which you want to send the email notifications.

    You must also specify the sender and receiver email addresses, and the subject of the email. If you want to attach the report of the restore operation performed, you must select **Attach Job Report**.

    For email notification, you must have specified the SMTP server details by using the either the GUI or the PowerShell command Set-SmSmtpServer.

11. Review the summary, and then click **Finish**.
12. Monitor the operation progress by clicking **Monitor** > **Jobs**.

For more information

- Oracle RAC One Node database is skipped for performing SnapCenter operations
- Failed to restore from a secondary SnapMirror or SnapVault location
- Failed to restore from a backup of an orphan incarnation
- Customizable parameters for backup, restore and clone operations on AIX systems

Restore and recover tablespaces using point-in-time recovery

You can restore a subset of tablespaces that has been corrupted or dropped without impacting the other tablespaces in the database. SnapCenter uses RMAN to perform point-in-time recovery (PITR) of the tablespaces.

Before you begin

- The backups that are required to perform PITR of tablespaces should be cataloged and mounted.
- If you have installed the plug-in as a non-root user, you should manually assign the execute permissions to the prescript and postscript directories.

About this task

During PITR operation, RMAN creates an auxiliary instance at the specified auxiliary destination. The auxiliary destination could be a mount point or ASM disk group. If there is sufficient space in the mounted location, you can reuse one of the mounted locations instead of a dedicated mount point.

You should specify the date and time or SCN and the tablespace is restored on the source database.

You can select and restore multiple tablespaces residing on ASM, NFS, and SAN environments. For example, if tablespaces TS2 and TS3 reside on NFS and TS4 reside on SAN, you can perform on single PITR operation to restore all the tablespaces.

In a RAC setup, you can perform PITR of tablespaces from any node of the RAC.

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page, select either **Database** or **Resource Group** from the **View** list.
3. Select the database of type single instance (multitenant) either from the database details view or the resource group details view.

   The database topology page is displayed.

4. From the Manage Copies view, select **Backups** from either the primary or the secondary (mirrored or replicated) storage systems.

   If the backup is not cataloged, you should select the backup and click **Catalog**.

5. Select the cataloged backup, and then click .
6. In the Restore Scope page, perform the following tasks:
   a. If you have selected a backup of a database in a Real Application Clusters (RAC) environment, select the RAC node.
   b. Select **Tablespaces**, and then specify the tablespaces you want to restore.

      You cannot perform PITR on SYSAUX, SYSTEM, and UNDO tablespaces.

   c. Select **Change database state if needed for restore and recovery** to change the state of the database to the state required to perform restore and recovery operations.

7. In the Recovery Scope page, perform one of the following actions:
   - If you want to recover to a specific System Change Number (SCN), select **Until SCN** and specify the SCN and auxiliary destination.
   - If you want to recover to a specific date and time, select **Date and Time** and specify the date and time and the auxiliary destination.

     SnapCenter identifies and then mounts and catalogs the optimal number of data and log backups required to perform PITR based on specified SCN or the selected date and time.
8. In the PreOps page, enter the path and the arguments of the prescript that you want to run before the restore operation.

   You should store the prescripts either in the /var/opt/snapcenter/spl/scripts path or in any folder inside this path. By default, the /var/opt/snapcenter/spl/scripts path is populated. If you have created any folders inside this path to store the scripts, you must specify those folders in the path.

   You can also specify the script timeout value. The default value is 60 seconds.

   SnapCenter allows you to use the predefined environment variables when you execute the prescript and postscript. Learn more
9. In the PostOps page, perform the following steps:
   a. Enter the path and the arguments of the postscript that you want to run after the restore operation.

      If the restore operation fails, postscripts will not be executed and cleanup activities will be triggered directly.

   b. Select the check box if you want to open the database after recovery.
10. In the Notification page, from the **Email preference** drop-down list, select the scenarios in which you want to send the email notifications.
11. Review the summary, and then click **Finish**.
12. Monitor the operation progress by clicking **Monitor** > **Jobs**.

Restore and recover pluggable database using point-in-time recovery

You can restore and recover a pluggable database (PDB) that has been corrupted or dropped without impacting the other PDBs in the container database (CDB). SnapCenter uses RMAN to perform point-in-time recovery (PITR) of the PDB.

Before you begin

- The backups that are required to perform PITR of a PDB should be cataloged and mounted.

  In a RAC setup, you should manually close the PDB (changing the state to MOUNTED) on all the nodes of the RAC setup.
- If you have installed the plug-in as a non-root user, you should manually assign the execute permissions to the prescript and postscript directories.

About this task

During PITR operation, RMAN creates an auxiliary instance at the specified auxiliary destination. The auxiliary destination could be a mount point or ASM disk group. If there is sufficient space in the mounted location, you can reuse one of the mounted locations instead of a dedicated mount point.

You should specify the date and time or SCN to perform PITR of the PDB. RMAN can recover READ WRITE, READ ONLY, or dropped PDBs including datafiles.

You can restore and recover only:

- one PDB at a time
- one tablespace in a PDB
- multiple tablespaces of the same PDB

In a RAC setup, you can perform PITR of tablespaces from any node of the RAC.

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page, select either **Database** or **Resource Group** from the **View** list.
3. Select the database of type single instance (multitenant) either from the database details view or the resource group details view.

   The database topology page is displayed.

4. From the Manage Copies view, select **Backups** from either the primary or the secondary (mirrored or replicated) storage systems.

   If the backup is not cataloged, you should select the backup and click **Catalog**.

5. Select the cataloged backup, and then click .
6. In the Restore Scope page, perform the following tasks:
   a. If you have selected a backup of a database in a Real Application Clusters (RAC) environment, select the RAC node.
   b. Depending on whether you want to restore the PDB or tablespaces in a PDB, perform one of the actions:

If you want to…	Steps…
Restore a PDB	i. Select **Pluggable databases (PDBs)**. ii. Specify the PDB you want to restore. You cannot perform PITR on PDB$SEED database.

Restore tablespaces in a PDB	i. Select **Pluggable database (PDB) tablespaces**. ii. Specify the PDB. iii. Specify either a single tablespace or multiple tablespaces you want to restore. You cannot perform PITR on SYSAUX, SYSTEM, and UNDO tablespaces.

c. Select **Change database state if needed for restore and recovery** to change the state of the database to the state required to perform restore and recovery operations.

7. In the Recovery Scope page, perform one of the following actions:
   - If you want to recover to a specific System Change Number (SCN), select **Until SCN** and specify the SCN and auxiliary destination.
   - If you want to recover to a specific date and time, select **Date and Time** and specify the date and time and the auxiliary destination.

     SnapCenter identifies and then mounts and catalogs the optimal number of data and log backups required to perform PITR based on specified SCN or the selected date and time.

8. In the PreOps page, enter the path and the arguments of the prescript that you want to run before the restore operation.

   You should store the prescripts either in the /var/opt/snapcenter/spl/scripts path or in any folder inside this path. By default, the /var/opt/snapcenter/spl/scripts path is populated. If you have created any folders inside this path to store the scripts, you must specify those folders in the path.

   You can also specify the script timeout value. The default value is 60 seconds.

   SnapCenter allows you to use the predefined environment variables when you execute the prescript and postscript. Learn more

9. In the PostOps page, perform the following steps:
   a. Enter the path and the arguments of the postscript that you want to run after the restore operation.

      If the restore operation fails, postscripts will not be executed and cleanup activities will be triggered directly.

   b. Select the check box if you want to open the database after recovery.

      In a RAC setup, the PDB will be opened only on the node where the database was recovered. You should manually open the recovered PDB on all the other nodes of the RAC setup.

10. On the Notification page, from the **Email preference** drop-down list, select the scenarios in which you want to send the email notifications.
11. Review the summary, and then click **Finish**.
12. Monitor the operation progress by clicking **Monitor** > **Jobs**.

Restore and recover Oracle databases using UNIX commands

The restore and recovery workflow includes planning, performing the restore and recovery operations, and monitoring the operations.

About this task

- You should execute the following commands to establish the connection with the SnapCenter Server, list the backups and retrieve its information, and restore the backup.

  The information regarding the parameters that can be used with the command and their descriptions can be obtained by running Get-Help *command_name*. Alternatively, you can also refer to the SnapCenter Software Command Reference Guide.

- For SnapMirror active sync restore operation, you must select the backup from the primary location.

Steps

1. Initiate a connection session with the SnapCenter Server for a specified user: *Open-SmConnection*
2. Retrieve the information about the backups that you want to restore: *Get-SmBackup*
3. Retrieve the detailed information about the specified backup: *Get-SmBackupDetails*

   This command retrieves the detailed information about the backup of a specified resource with a given backup ID. The information includes database name, version, home, start and end SCN, tablespaces, pluggable databases, and its tablespaces.

4. Restore data from the backup: *Restore-SmBackup*

Monitor Oracle database restore operations

You can monitor the progress of different SnapCenter restore operations by using the Jobs page. You might want to check the progress of an operation to determine when it is complete or if there is an issue.

About this task

Post-restore states describe the conditions of the resource after a restore operation and any further restore actions that you can take.

The following icons appear on the Jobs page, and indicate the state of the operation:

- In progress
- Completed successfully
- Failed
- Completed with warnings or could not start due to warnings
- Queued
- Canceled

Steps

1. In the left navigation pane, click **Monitor**.
2. In the **Monitor** page, click **Jobs**.
3. In the **Jobs** page, perform the following steps:
   a. Click to filter the list so that only restore operations are listed.
   b. Specify the start and end dates.
   c. From the **Type** drop-down list, select **Restore**.
   d. From the **Status** drop-down list, select the restore status.
   e. Click **Apply** to view the operations that have been completed successfully.
4. Select the restore job, and then click **Details** to view the job details.
5. In the **Job Details** page, click **View logs**.

   The **View logs** button displays the detailed logs for the selected operation.

Cancel Oracle database restore operations

You can cancel restore jobs that are queued.

You should be logged in as the SnapCenter Admin or job owner to cancel restore operations.

About this task

- You can cancel a queued restore operation from either the **Monitor** page or the **Activity** pane.
- You cannot cancel a running restore operation.
- You can use the SnapCenter GUI, PowerShell cmdlets, or CLI commands to cancel the queued restore operations.
- The **Cancel Job** button is disabled for restore operations that cannot be canceled.
- If you selected **All members of this role can see and operate on other members objects** in Users\Groups page while creating a role, you can cancel the queued restore operations of other members while using that role.

Step

Perform one of the following actions:

From the…	Action
Monitor page	a. In the left navigation pane, click **Monitor** > **Jobs**. b. Select the job and click **Cancel Job**.
Activity pane	a. After initiating the restore operation, click on the Activity pane to view the five most recent operations. b. Select the operation. c. In the Job Details page, click **Cancel Job**.

Clone Oracle database

Clone workflow

The clone workflow includes planning, performing the clone operation, and monitoring the operation.

You might clone databases for the following reasons:

- To test functionality that has to be implemented using the current database structure and content during application development cycles.
- To populate data warehouses using data extraction and manipulation tools.
- To recover data that was mistakenly deleted or changed.

The following workflow shows the sequence in which you must perform the clone operation:

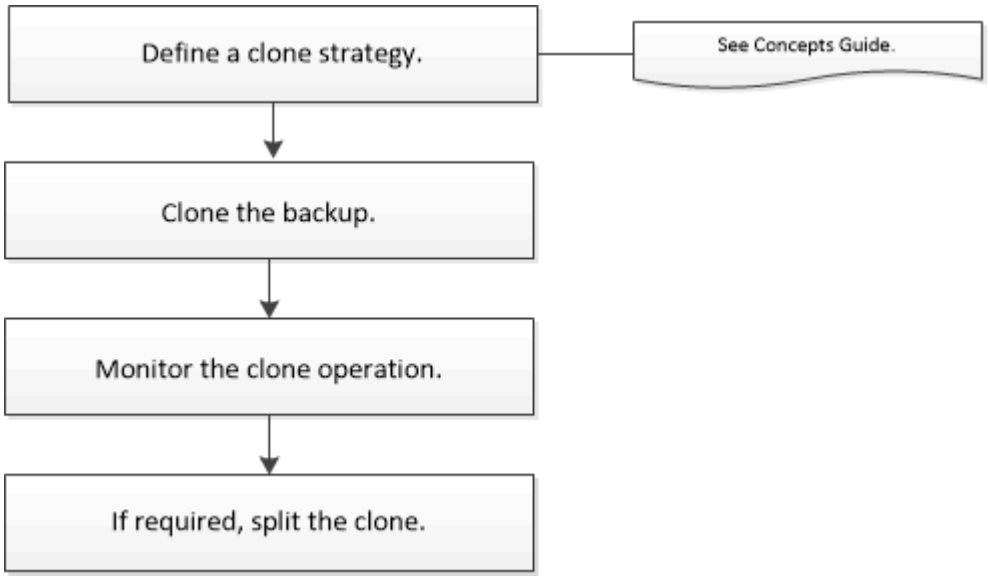

Define a clone strategy for Oracle databases

Defining a strategy before cloning your database ensures that the cloning operation is successful.

Types of backups supported for cloning

SnapCenter supports cloning of different types of backups of Oracle databases.

- Online data backup
- Online full backup
- Offline mount backup
- Offline shutdown backup
- Backups of Data Guard standby databases and Active Data Guard standby databases
- Online data backups, online full backups, offline mount backups, and offline shutdown backups in a Real Application Clusters (RAC) configuration

- Online data backups, online full backups, offline mount backups, and offline shutdown backups in an Automatic Storage Management (ASM) configuration

  SAN configurations are not supported if user_friendly_names option in the multipath configuration file is set to yes.

Cloning of archive log backups is not supported.

Types of cloning supported for Oracle databases

In an Oracle database environment, SnapCenter supports cloning of a database backup. You can clone the backup from primary and secondary storage systems.

The SnapCenter Server uses NetApp FlexClone technology to clone backups.

You can refresh a clone by running the "Refresh-SmClone" command. This command creates a backup of the database, deletes the existing clone, and creates a clone with the same name.

The clone refresh operation can only be performed using the UNIX commands.

Clone naming conventions for Oracle databases

From SnapCenter 3.0, the naming convention used for clones of file systems is different from the clones of ASM disk groups.

- The naming convention for SAN or NFS file systems is FileSystemNameofsourcedatabase_CLONESID.
- The naming convention for ASM disk groups is SC_HASHCODEofDISKGROUP_CLONESID.

  HASHCODEofDISKGROUP is an automatically generated number (2 to 10 digits) that is unique for each ASM disk group.

Limitations of cloning Oracle databases

You should be aware of the limitations of clone operations before you clone the databases.

- If you are using any version of Oracle from 11.2.0.4 to 12.1.0.1, the clone operation will be in hung state when you run the *renamedg* command. You can apply the Oracle patch 19544733 to fix this issue.
- Cloning of databases from a LUN that is directly attached to a host (for instance, by using Microsoft iSCSI Initiator on a Windows host) to a VMDK or an RDM LUN on the same Windows host, or another Windows host, or vice versa, is not supported.
- The root directory of the volume mount point cannot be a shared directory.
- If you move a LUN that contains a clone to a new volume, the clone cannot be deleted.

Predefined environment variables for clone specific prescript and postscript

SnapCenter allows you to use the predefined environment variables when you execute the prescript and postscript while cloning a database.

Supported predefined environment variables for cloning a database

- **SC_ORIGINAL_SID** specifies the SID of the source database.

  This parameter will be populated for application volumes.

  Example: NFSB32

- **SC_ORIGINAL_HOST** specifies the name of the source host.

  This parameter will be populated for application volumes.

  Example: asmrac1.gdl.englab.netapp.com

- **SC_ORACLE_HOME** specifies the path of the target database's Oracle home directory.

  Example: /ora01/app/oracle/product/18.1.0/db_1

- **SC_BACKUP_NAME**" specifies the name of the backup.

  This parameter will be populated for application volumes.

  Examples:

  - If the database is not running in ARCHIVELOG mode: DATA@RG2_scspr2417819002_07-20-2021_12.16.48.9267_0|LOG@RG2_scspr2417819002_07-20-2021_12.16.48.9267_1
  - If the database is running in ARCHIVELOG mode: DATA@RG2_scspr2417819002_07-20-2021_12.16.48.9267_0|LOG:RG2_scspr2417819002_07-20-2021_12.16.48.9267_1,RG2_scspr2417819002_07-21-2021_12.16.48.9267_1,RG2_scspr2417819002_07-22-2021_12.16.48.9267_1

- **SC_AV_NAME** specifies the names of the application volumes.

  Example: AV1|AV2

- **SC_ORIGINAL_OS_USER** specifies the operating system owner of the source database.

  Example: oracle

- **SC_ORIGINAL_OS_GROUP** specifies the operating system group of the source database.

  Example: oinstall

- **SC_TARGET_SID**" specifies the SID of the cloned database.

  For PDB clone workflow, the value of this parameter will not be predefined.

  This parameter will be populated for application volumes.

  Example: clonedb

- **SC_TARGET_HOST** specifies the name of the host where the database will be cloned.

  This parameter will be populated for application volumes.

  Example: asmrac1.gdl.englab.netapp.com

- **SC_TARGET_OS_USER** specifies the operating system owner of the cloned database.

For PDB clone workflow, the value of this parameter will not be predefined.

Example: oracle

- **SC_TARGET_OS_GROUP** specifies the operating system group of the cloned database.

  For PDB clone workflow, the value of this parameter will not be predefined.

  Example: oinstall

- **SC_TARGET_DB_PORT** specifies the database port of the cloned database.

  For PDB clone workflow, the value of this parameter will not be predefined.

  Example: 1521

For information about delimiters, see Supported delimiters.

Requirements for cloning an Oracle database

Before cloning an Oracle database, you should ensure that prerequisites are completed.

- You should have created a backup of the database using SnapCenter.

  You should have successfully created online data and log backups or offline (mount or shutdown) backups for the cloning operation to succeed.

- If you want to customize the control file or redo log file paths, you should have preprovisioned the required file system or Automatic Storage Management (ASM) disk group.

  By default, redo log and control files of the cloned database are created on the ASM disk group or the file system provisioned by SnapCenter for the data files of the clone database.

- If you are using ASM over NFS, you should add */var/opt/snapcenter/scu/clones/*/** to the existing path defined in the asm_diskstring parameter.
- In the asm_diskstring parameter, you should configure *AFD:** if you are using ASMFD or configure *ORCL:** if you are using ASMLIB.

  For information on how to edit the asm_diskstring parameter, see How to add disk paths to asm_diskstring.

- If you are creating the clone on an alternate host, the alternate host should meet the following requirements:
  - SnapCenter Plug-in for Oracle Database should be installed on the alternate host.
  - The clone host should be able to discover LUNs from primary or secondary storage.
    - If you are cloning from primary storage or secondary (Vault or Mirror) storage to an alternate host, then make sure that an iSCSI session is either established between the secondary storage and the alternate host, or zoned properly for FC.
    - If you are cloning from Vault or Mirror storage to the same host, then make sure that an iSCSI session is either established between the Vault or Mirror storage and the host, or zoned properly for FC.
    - If you are cloning in a virtualized environment, ensure that an iSCSI session is either established between the primary or secondary storage and the ESX server hosting the alternate host, or zoned

properly for FC.

For information, refer to host utilities documentation.

- If the source database is an ASM database:
  - The ASM instance should be up and running on the host where the clone will be performed.
  - The ASM disk group should be provisioned prior to the clone operation if you want to place archive log files of the cloned database in a dedicated ASM disk group.
  - The name of the data disk group can be configured, but ensure that the name is not used by any other ASM disk group on the host where the clone will be performed.

    Data files residing on the ASM disk group are provisioned as part of the SnapCenter clone workflow.
- For NVMe, NVMe util should be installed

- The protection type for the data LUN and the log LUN, such as mirror, vault, or mirror-vault, should be the same to discover secondary locators during cloning to an alternate host using log backups.
- You should set the value of exclude_seed_cdb_view to FALSE in the source database parameter file to retrieve seed PDB related information for cloning a backup of 12_c_ database.

  The seed PDB is a system-supplied template that the CDB can use to create PDBs. The seed PDB is named PDB$SEED. For information about PDB$SEED, see the Oracle Doc ID 1940806.1.

  You should set the value before backing up 12_c_ database.

- SnapCenter supports backup of file systems that are managed by the autofs subsystem. If you are cloning the database, ensure that data mount points are not under the root of the autofs mount point because the root user of the plug-in host does not have permission to create directories under the root of the autofs mount point.

  If control and redo log files are under data mount point, you should modify the control file path, and then redo log file path accordingly.

  You can manually register the new cloned mount points with the autofs subsystem. The new cloned mount points will not be registered automatically.

- If you have a TDE (auto login) and want to clone the database on the same or alternate host, you should copy wallet (key files) under */etc/ORACLE/WALLET/$ORACLE_SID* from the source database to the cloned database.
- You should set the value of use_lvmetad = 0 in */etc/lvm/lvm.conf* and stop the lvm2-lvmetad service to successfully perform cloning in storage area network (SAN) environments on Oracle Linux 7 or later or Red Hat Enterprise Linux (RHEL) 7 or later.
- You should install the 13366202 Oracle patch if you are using Oracle database 11.2.0.3 or later and the database ID for the auxiliary instance is changed using an NID script.
- You should ensure that the aggregates hosting the volumes should be in the assigned aggregates list of the storage virtual machine (SVM).
- For NVMe, if any target port has to be excluded from connecting, you should add the target node name and port name in the /var/opt/snapcenter/scu/etc/nvme.conf file.

  If the file does not exist, you should create the file as shown in the example below:

```
blacklist {
nn-0x<target_node_name_1>:pn-0x<target_port_name_1>
nn-0x<target_node_name_2>:pn-0x<target_port_name_2>
}
```

- You should ensure that the LUN is not mapped to the AIX host using iGroup consisting of mixed protocols iSCSI and FC. For more information, see Operation fails with error Unable to discover the device for LUN.

Clone an Oracle database backup

You can use SnapCenter to clone an Oracle database using the backup of the database.

Before you begin

If you have installed the plug-in as a non-root user, you should manually assign the execute permissions to the prescript and postscript directories.

About this task

- The cloning operation creates a copy of the database data files, and creates new online redo log files and control files. The database can be optionally recovered to a specified time, based on the specified recovery options.

  Cloning fails if you try to clone a backup that was created on a Linux host to an AIX host or vice-versa.

  SnapCenter creates a stand-alone database when cloned from an Oracle RAC database backup. SnapCenter supports creating clone from the backup of a Data Guard standby and Active Data Guard standby databases.

  During cloning, SnapCenter mounts the optimal number of log backups based on SCN or dat and time for recovery operations. After recovery, the log backup is unmounted. All such clones are mounted under */var/opt/snapcenter/scu/clones/*. If you are using ASM over NFS, you should add */var/opt/snapcenter/scu/clones/*/** to the existing path defined in the asm_diskstring parameter.

  While cloning a backup of an ASM database in a SAN environment, udev rules for the cloned host devices are created at */etc/udev/rules.d/999-scu-netapp.rules*. These udev rules associated with the cloned host devices are deleted when you delete the clone.

  In a Flex ASM setup, you cannot perform clone operation on Leaf nodes if the cardinality is less than the number nodes in the RAC cluster.

- For SnapLock enabled policies, for ONTAP 9.12.1 and below version, if you specify a Snapshot locking period, the clones created from the tamper proof Snapshots as part of restore will inherit the SnapLock expiry time. Storage admin should manually cleanup the clones post the SnapLock expiry time.

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page, select either **Database** or **Resource Group** from the **View** list.
3. Select the database either from the database details view or from the resource group details view.

The database topology page is displayed.

4. From the Manage Copies view, select the backups either from Local copies (primary), Mirror copies (secondary), or Vault copies (secondary).
5. Select the Data backup from the table, and then click .
6. In the Name page, perform one of the following actions:

If you want to…	Steps…
Clone a database (CDB or non CDB)	a. Specify the SID of the clone. The clone SID is not available by default, and the maximum length of the SID is 8 characters. You should ensure that no database with the same SID exists on the host where the clone will be created.
Clone a pluggable database (PDB)	a. Select **PDB Clone**. b. Specify the PDB that you want to clone. c. Specify the name of cloned PDB. For the detailed steps to clone a PDB, see Clone a pluggable database.

When you select a mirrored or vault data:

- if there are no log backup at mirror or vault, nothing is selected and the locators are empty.
- if log backups exist in mirror or vault, the latest log backup is selected and corresponding locator is displayed.

If the selected log backup exists in both mirror and vault location, both the locators are displayed.

7. In the Locations page, perform the following actions:

For this field…	Do this…
Clone host	By default, the source database host is populated. If you want to create the clone on an alternate host, select the host having the same version of Oracle and OS as that of the source database host.

For this field…	Do this…
Datafile locations	By default, the datafile location is populated. The SnapCenter default naming convention for SAN or NFS file systems is FileSystemNameofsourcedatabase_CLONESID. The SnapCenter default naming convention for ASM disk groups is SC_HASHCODEofDISKGROUP_CLONESID. The HASHCODEofDISKGROUP is an automatically generated number (2 to 10 digits) that is unique for each ASM disk group. If you are customizing the ASM disk group name, ensure that the name length adheres to the maximum length supported by Oracle. If you want to specify a different path, you must enter the datafile mount points or ASM disk group names for clone database. When you customize the datafile path, you must also change the control file and redo log file ASM disk group names or file system either to the same name used for data files or to an existing ASM disk groups or file system.
Control files	By default, the control file path is populated. The control files are placed in the same ASM disk group or file system as that of the data files. If you want to override the control file path, you can provide a different control file path. The file system or the ASM disk group should exist on the host. By default, the number of control files will be same as that of the source database. You can modify the number of control files but a minimum of one control file is required to clone the database. You can customize the control file path to a different file system (existing) than that of the source database.

For this field…	Do this…
Redo logs	By default, the redo log file group, path, and their sizes are populated. The redo logs are placed in the same ASM disk group or file system as that of the data files of the cloned database. If you want to override the redo log file path, you can customize the redo log file path to a different file system than that of the source database.. The new file system or the ASM disk group should exist on the host. By default, the number of redo log groups, redo log files, and their sizes will be same as that of the source database. You can modify the following parameters: • Number of redo log groups A minimum of two redo log groups are required to clone the database. • Redo log files in each group and their path You can customize the redo log file path to a different file system (existing) than that of the source database. A minimum of one redo log file is required in the redo log group to clone the database. • Sizes of the redo log file

8. On the Credentials page, perform the following actions:

For this field…	Do this…
Credential name for sys user	Select the Credential to be used for defining the sys user password of the clone database. If SQLNET.AUTHENTICATION_SERVICES is set to NONE in sqlnet.ora file on the target host, you should not select **None** as the Credential in the SnapCenter GUI.

For this field…	Do this…
ASM Instance Credential name	Select **None** if OS authentication is enabled for connecting to the ASM instance on the clone host. Otherwise, select the Oracle ASM credential configured with either “sys” user or an user having “sysasm” privilege applicable to the clone host.

The Oracle home, user name, and group details are automatically populated from the source database. You can change the values based on the Oracle environment of the host where the clone will be created.

9. In the PreOps page, perform the following steps:
   a. Enter the path and the arguments of the prescript that you want to run before the clone operation.

      You must store the prescript either in */var/opt/snapcenter/spl/scripts* or in any folder inside this path. By default, the */var/opt/snapcenter/spl/scripts* path is populated. If you have placed the script in any folder inside this path, you need to provide the complete path up to the folder where the script is placed.

      SnapCenter allows you to use the predefined environment variables when you execute the prescript and postscript. Learn more

   b. In the Database Parameter settings section, modify the values of prepopulated database parameters that are used to initialize the database.

      You can add additional parameters by clicking + .

      If you are using Oracle Standard Edition and the database is running in Archive log mode or you want restore a database from archive redo log, add the parameters and specify the path.

      - LOG_ARCHIVE_DEST
      - LOG_ARCHIVE_DUPLEX_DEST

      Fast recovery area (FRA) is not defined in the prepopulated database parameters. You can configure FRA by adding the related parameters.

      The default value of log_archive_dest_1 is $ORACLE_HOME/clone_sid and the archive logs of the cloned database will be created in this location. If you have deleted the log_archive_dest_1 parameter, the archive log location is determined by Oracle. You can define a new location for archive log by editing log_archive_dest_1 but ensure that the file system or disk group should be existing and made available on the host.

   c. Click **Reset** to get the default database parameter settings.
10. In the PostOps page, **Recover database** and **Until Cancel** are selected by default to perform recovery of the cloned database.

    SnapCenter performs recovery by mounting the latest log backup that have the unbroken sequence of archive logs after the data backup that was selected for cloning. The log and data backup should be on primary storage to perform the clone on primary storage and log and data backup should be on secondary storage to perform the clone on secondary storage.

The **Recover database** and **Until Cancel** options are not selected if SnapCenter fails to find the appropriate log backups. You can provide the external archive log location if log backup is not available in **Specify external archive log locations**. You can specify multiple log locations.

If you want to clone a source database that is configured to support flash recovery area (FRA) and Oracle Managed Files (OMF), the log destination for recovery must also adhere to OMF directory structure.

The PostOps page is not displayed if the source database is a Data Guard standby or an Active Data Guard standby database. For Data Guard standby or an Active Data Guard standby database, SnapCenter does not provide an option to select the type of recovery in the SnapCenter GUI but the database is recovered using Until Cancel recovery type without applying any logs.

Field name	Description
Until Cancel	SnapCenter performs recovery by mounting the latest log backup having the unbroken sequence of archive logs after that data backup that was selected for cloning. The cloned database is recovered till the missing or corrupt log file.
Date and time	SnapCenter recovers the database up to a specified date and time. The accepted format is mm/dd/yyyy hh:mm:ss. The time can be specified in 24 hour format.
Until SCN (System Change Number)	SnapCenter recovers the database up to a specified system change number (SCN).
Specify external archive log locations	If the database is running in ARCHIVELOG mode, SnapCenter identifies and mounts optimal number of log backups based on the specified SCN or the selected date and time. You can also specify the external archive log location. SnapCenter will not automatically identify and mount the log backups if you have selected Until Cancel.

Field name	Description
Create new DBID	By default **Create new DBID** check box is selected to generate a unique number (DBID) for the cloned database differentiating it from the source database. Clear the check box if you want to assign the DBID of the source database to the cloned database. In this scenario, if you want to register the cloned database with the external RMAN catalog where the source database is already registered, the operation fails.
Create tempfile for temporary tablespace	Select the check box if you want to create a tempfile for the default temporary tablespace of the cloned database. If the check box is not selected, the database clone will be created without the tempfile.
Enter sql entries to apply when clone is created	Add the sql entries that you want to apply when the clone is created.
Enter scripts to run after clone operation	Specify the path and the arguments of the postscript that you want to run after the clone operation. You should store the postscript either in */var/opt/snapcenter/spl/scripts* or in any folder inside this path. By default, the */var/opt/snapcenter/spl/scripts* path is populated. If you have placed the script in any folder inside this path, you need to provide the complete path up to the folder where the script is placed. If the clone operation fails, postscripts will not be executed and cleanup activities will be triggered directly.

11. In the Notification page, from the **Email preference** drop-down list, select the scenarios in which you want to send the emails.

    You must also specify the sender and receiver email addresses, and the subject of the email. If you want to attach the report of the clone operation performed, select **Attach Job Report**.

    For email notification, you must have specified the SMTP server details using the either the GUI or the PowerShell command Set-SmSmtpServer.

12. Review the summary, and then click **Finish**.

While performing recovery as part of clone create operation, even if recovery fails, the clone is created with a warning. You can perform manual recovery on this clone to bring the clone database to consistent state.

13. Monitor the operation progress by clicking **Monitor** > **Jobs**.

Result

After cloning the database you can refresh the resources page to list the cloned database as one of the resource available for backup. The cloned database can be protected like any other database using the standard backup workflow or can be included in a resource group (either newly created or existing). The cloned database can be further cloned (clone of clones).

After cloning, you should never rename the cloned database.

If you have not performed recovery while cloning, the backing up of the cloned database might fail due to improper recovery and you might have to perform manual recovery. The log backup can also fail if default location which was populated for archive logs is on a non-NetApp storage or if the storage system is not configured with SnapCenter.

In AIX setup, you can use the lkdev command to lock and the rendev command to rename the disks on which the cloned database resided.

Locking or renaming of devices will not affect the clone deletion operation. For AIX LVM layouts built on SAN devices, renaming of devices will not be supported for the cloned SAN devices.

Find more information

- Restore or cloning fails with ORA-00308 error message
- Failed to recover a cloned database
- Customizable parameters for backup, restore and clone operations on AIX systems

Update the preferred IP on the host

After the clone operation completes, the path provided by the Storage Access Layer (SAL) to the clone will be in *<nfs_lif_IP>:<JunctionPath>* format. To provide the preferred IP, you must configure it on the host using the SCCLI commands.

Steps

1. Login to the database host.
2. Initiate a PowerShell connection session with SnapCenter, for a specified user.

   Open-SmConnection

3. Create an empty file.

   touch /var/opt/snapcenter/scu/etc/storagepreference.properties

4. Configure the preferred data LIF for the SVM.

   Add-SvmPreferredDataPath -SVM <SVM Name> -DataPath <IP address or FQDN>

5. Verify the preferred path.

   Get-SvmPreferredDataPath

Clone a pluggable database

You can clone a pluggable database (PDB) to a different or same target CDB on the same host or alternate host. You can also recover the cloned PDB to a desired SCN or date and time.

Before you begin

If you have installed the plug-in as a non-root user, you should manually assign the execute permissions to the prescript and postscript directories.

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page, select either **Database** or **Resource Group** from the **View** list.
3. Select the database of type single instance (multitenant) either from the database details view or from the resource group details view.

   The database topology page is displayed.

4. From the Manage Copies view, select the backups either from Local copies (primary), Mirror copies (secondary), or Vault copies (secondary).
5. Select the backup from the table, and then click .
6. In the Name page, perform the following actions:
   a. Select **PDB Clone**.
   b. Specify the PDB that you want to clone.

      You can clone only one PDB at a time.

   c. Specify the name of the clone PDB.
7. In the Locations page, perform the following actions:

For this field…	Do this…
Clone host	By default, the source database host is populated. If you want to create the clone on an alternate host, select the host having the same version of Oracle and OS as that of the source database host.
Target CDB	Select the CDB where you want to include the cloned PDB. You should ensure that the target CDB is running.

Database State	Select the **Open the cloned PDB in READ-WRITE mode** checkbox if you want to open the PDB in READ-WRITE mode.
Datafile locations	By default, the datafile location is populated. The SnapCenter default naming convention for SAN or NFS file systems is FileSystemNameofsourcedatabase_SCJOBID. The SnapCenter default naming convention for ASM disk groups is SC_HASHCODEofDISKGROUP_SCJOBID. The HASHCODEofDISKGROUP is an automatically generated number (2 to 10 digits) that is unique for each ASM disk group. If you are customizing the ASM disk group name, ensure that the name length adheres to the maximum length supported by Oracle. If you want to specify a different path, you must enter the datafile mount points or ASM disk group names for clone database.

The Oracle home, user name, and group details are automatically populated from the source database. You can change the values based on the Oracle environment of the host where the clone will be created.

8. In the PreOps page, perform the following steps:
   a. Enter the path and the arguments of the prescript that you want to run before the clone operation.

      You should store the prescript either in /var/opt/snapcenter/spl/scripts or in any folder inside this path. By default, the /var/opt/snapcenter/spl/scripts path is populated. If you have placed the script in any folder inside this path, you need to provide the complete path up to the folder where the script is placed.

      SnapCenter allows you to use the predefined environment variables when you execute the prescript and postscript. Learn more

   b. In the Auxiliary CDB clone database parameter settings section, modify the values of prepopulated database parameters that are used to initialize the database.
9. Click **Reset** to get the default database parameter settings.
10. In the PostOps page, **Until Cancel** is selected by default to perform recovery of the cloned database.

    The **Until Cancel** option is not selected if SnapCenter fails to find the appropriate log backups. You can provide the external archive log location if log backup is not available in **Specify external archive log locations**. You can specify multiple log locations.

If you want to clone a source database that is configured to support flash recovery area (FRA) and Oracle Managed Files (OMF), the log destination for recovery must also adhere to OMF directory structure.

Field name	Description
Until Cancel	SnapCenter performs recovery by mounting the latest log backup having the unbroken sequence of archive logs after that data backup that was selected for cloning. The log and data backup should be on primary storage to perform the clone on primary storage and log and data backup should be on secondary storage to perform the clone on secondary storage. The cloned database is recovered till the missing or corrupt log file.
Date and time	SnapCenter recovers the database up to a specified date and time. The time can be specified in 24 hour format.
Until SCN (System Change Number)	SnapCenter recovers the database up to a specified system change number (SCN).
Specify external archive log locations	Specify the external archive log location.
Create new DBID	By default **Create new DBID** check box is not selected for the auxiliary clone database. Select the check box if you want to generate a unique number (DBID) for the auxiliary cloned database differentiating it from the source database.
Create tempfile for temporary tablespace	Select the check box if you want to create a tempfile for the default temporary tablespace of the cloned database. If the check box is not selected, the database clone will be created without the tempfile.
Enter sql entries to apply when clone is created	Add the sql entries that you want to apply when the clone is created.

Field name	Description
Enter scripts to run after clone operation	Specify the path and the arguments of the postscript that you want to run after the clone operation. You should store the postscript either in */var/opt/snapcenter/spl/scripts* or in any folder inside this path. By default, the */var/opt/snapcenter/spl/scripts* path is populated. If you have placed the script in any folder inside this path, you need to provide the complete path up to the folder where the script is placed. If the clone operation fails, postscripts will not be executed and cleanup activities will be triggered directly.

11. In the Notification page, from the **Email preference** drop-down list, select the scenarios in which you want to send the emails.

    You must also specify the sender and receiver email addresses, and the subject of the email. If you want to attach the report of the clone operation performed, select **Attach Job Report**.

    For email notification, you must have specified the SMTP server details using the either the GUI or the PowerShell command Set-SmSmtpServer.

12. Review the summary, and then click **Finish**.
13. Monitor the operation progress by clicking **Monitor** > **Jobs**.

After you finish

If you want to create a backup of the cloned PDB, you should backup the target CDB where the PDB is cloned because backing up only the cloned PDB is not possible. You should create a secondary relationship for the target CDB if you want to create the backup with secondary relationship.

In a RAC setup the storage for cloned PDB is attached only to the node where the PDB clone was performed. The PDBs on the other nodes of the RAC are in MOUNT state. If you want the cloned PDB to be accessible from the other nodes, you should manually attach the storage to the other nodes.

Find more information

- Restore or cloning fails with ORA-00308 error message
- Customizable parameters for backup, restore and clone operations on AIX systems

Clone Oracle database backups using UNIX commands

The clone workflow includes planning, performing the clone operation, and monitoring the operation.

About this task

You should execute the following commands to create the Oracle database clone specification file and initiate the clone operation.

The information regarding the parameters that can be used with the command and their descriptions can be obtained by running Get-Help *command_name*. Alternatively, you can also refer to the SnapCenter Software Command Reference Guide.

Steps

1. Create an Oracle database clone specification from a specified backup: *New-SmOracleCloneSpecification*

   If secondary data protection policy is unified mirror-vault, then specify only -IncludeSecondaryDetails. You do not have to specify -SecondaryStorageType.

   This command automatically creates an Oracle database clone specification file for the specified source database and its backup. You must also provide a clone database SID so that the specification file created has the automatically generated values for the clone database which you will be creating.

   The clone specification file is created at */var/opt/snapcenter/sco/clone_specs*.

2. Initiate a clone operation from a clone resource group or an existing backup: *New-SmClone*

   This command initiates a clone operation. You must also provide an Oracle clone specification file path for the clone operation. You can also specify the recovery options, host where the clone operation to be performed, prescripts, postscripts, and other details.

   By default, the archive log destination file for the clone database is automatically populated at *$ORACLE_HOME/CLONE_SIDs*.

Split an Oracle Database Clone

You can use SnapCenter to split a cloned resource from the parent resource. The clone that is split becomes independent of the parent resource.

About this task

- You cannot perform the clone split operation on an intermediate clone.

  For example, after you create clone1 from a database backup, you can create a backup of clone1, and then clone this backup (clone2). After you create clone2, clone1 is an intermediate clone, and you cannot perform the clone split operation on clone1. However, you can perform the clone split operation on clone2.

  After splitting clone2, you can perform the clone split operation on clone1 because clone1 is no longer the intermediate clone.

- When you split a clone, the backup copies of the clone are deleted.
- For information about FlexClone volume split operations, see, Split a FlexClone volume from its parent volume
- Ensure that the volume or aggregate on the storage system is online.

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page, select **Database** from the **View** list.
3. Select the cloned resource, (for example, the database or LUN) and then click .
4. Review the estimated size of the clone that is to be split and the required space available on the aggregate, and then click **Start**.
5. Monitor the operation progress by clicking **Monitor** > **Jobs**.

   The clone split operation stops responding if the SMCore service restarts and the databases on which the clone split operation was performed are listed as clones in the Resources page. You should run the *Stop-SmJob* cmdlet to stop the clone split operation, and then retry the clone split operation.

   If you want a longer poll time or shorter poll time to check whether the clone is split or not, you can change the value of CloneSplitStatusCheckPollTime parameter in SMCoreServiceHost.exe.config file to set the time interval for SMCore to poll for the status of the clone split operation. The value is in milliseconds and the default value is 5 minutes.

   For example,

```
<add key="CloneSplitStatusCheckPollTime" value="300000" />
```

The clone split start operation fails if backup, restore, or anther clone split is in progress. You should restart the clone split operation only after the running operations are complete.

Split clone of a pluggable database

You can use SnapCenter to split a cloned pluggable database (PDB).

About this task

If you created a backup of the target CDB where the PDB is cloned, when you split the PDB clone, the cloned PDB is also removed from all the backups of the target CDB containing the cloned PDB.

The PDB clones are not displayed in the inventory or resources view.

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. Select the source container database (CDB) from the resource or resource group view.
3. From the Manage Copies view, select **Clones** either from the primary or secondary (mirrored or replicated) storage systems.
4. Select the PDB clone (targetCDB:PDBClone) and then click .
5. Review the estimated size of the clone that is to be split and the required space available on the aggregate, and then click **Start**.
6. Monitor the operation progress by clicking **Monitor** > **Jobs**.

Monitor Oracle database clone operations

You can monitor the progress of SnapCenter clone operations by using the Jobs page. You might want to check the progress of an operation to determine when it is complete or if there is an issue.

About this task

The following icons appear on the Jobs page, and indicate the state of the operation:

- In progress
- Completed successfully
- Failed
- Completed with warnings or could not start due to warnings
- Queued
- Canceled

Steps

1. In the left navigation pane, click **Monitor**.
2. In the **Monitor** page, click **Jobs**.
3. In the **Jobs** page, perform the following steps:
   a. Click to filter the list so that only clone operations are listed.
   b. Specify the start and end dates.
   c. From the **Type** drop-down list, select **Clone**.
   d. From the **Status** drop-down list, select the clone status.
   e. Click **Apply** to view the operations that are completed successfully.
4. Select the clone job, and then click **Details** to view the job details.
5. In the Job Details page, click **View logs**.

Refresh a clone

You can refresh the clone by running the *Refresh-SmClone* command. This command creates a backup of the database, deletes the existing clone, and creates a clone with the same name.

You cannot refresh a PDB clone.

What you will need

- Create an online full backup or an offline data backup policy with no scheduled backups enabled.
- Configure the email notification in the policy for backup failures only.
- Define the retention count for the on-demand backups appropriately to ensure that there are no unwanted backups.

- Ensure that only an online full backup or an offline data backup policy is associated with resource group which is identified for refresh clone operation.
- Create a resource group with only one database.
- If a cron job is created for the clone refresh command, ensure that the SnapCenter schedules and the cron schedules are not overlapping for the database resource group.

  For a cron job created for the clone refresh command, ensure that you run Open-SmConnection after every 24hrs.

- Ensure that the clone SID is unique for a host.

  If multiple refresh clone operations use the same clone specification file or use the clone specification file with same clone SID, existing clone with the SID on the host will be deleted and then the clone will be created.

- Ensure that the backup policy is enabled with secondary protection and the clone specification file is created with “-IncludeSecondaryDetails” to create the clones using secondary backups.
  - If the primary clone specification file is specified but the policy has secondary update option selected, the backup will be created, and update will get transferred to secondary. However, the clone will be created from the primary backup.
  - If the primary clone specification file is specified and the policy does not have secondary update option selected, the backup will be created on primary and clone will be created from primary.

Steps

1. Initiate a connection session with the SnapCenter Server for a specified user: *Open-SmConnection*
2. Create an Oracle database clone specification from a specified backup: *New-SmOracleCloneSpecification*

   If secondary data protection policy is unified mirror-vault, then specify only -IncludeSecondaryDetails. You do not have to specify -SecondaryStorageType.

   This command automatically creates an Oracle database clone specification file for the specified source database and its backup. You must also provide a clone database SID so that the specification file created has the automatically generated values for the clone database which you will be creating.

   The clone specification file is created at */var/opt/snapcenter/sco/clone_specs*.

3. Run *Refresh-SmClone*.

   If the operation fails with the "PL-SCO-20032: canExecute operation failed with error: PL-SCO-30031: Redo log file +SC_2959770772_clmdb/clmdb/redolog/redo01_01.log exists" error messages, specify a higher value for *-WaitToTriggerClone*.

   For detailed information on UNIX commands, see the SnapCenter Software Command Reference Guide.

Delete clone of a pluggable database

You can delete the clone of a pluggable database (PDB) if you no longer require.

If you created a backup of the target CDB where the PDB is cloned, when you delete the PDB clone, the cloned PDB is also removed from the backup of the target CDB.

The PDB clones are not displayed in the inventory or resources view.

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. Select the source container database (CDB) from the resource or resource group view.
3. From the Manage Copies view, select **Clones** either from the primary or secondary (mirrored or replicated) storage systems.
4. Select the PDB clone (targetCDB:PDBClone) and then click .
5. Click **OK**.

Manage application volumes

What are application volumes

Application Volumes are the storage where information such as configuration, installer, and other non-data files related to Oracle database are stored.

SnapCenter Plug-in for Oracle Database allows you to create consistent backup of application volumes (non-data volumes) along with the Oracle databases.

The plug-in automates the backup and cloning of application volumes.

- Protect application volumes along with Oracle database volumes in a single resource group.
- Create backups of application volumes.
- Create backups of Oracle databases along with application volumes.
- Create clones of databases along with application volumes up to a point-in-time.
- Schedule backup operations.
- Monitor all operations.
- View reports of backup and clone operations.

Add application volumes

SnapCenter supports backing up and cloning of application volumes of Oracle database. You should manually add the application volumes. Auto discovery of application volumes is not supported.

Application volumes support only direct NFS and direct iSCSI connections.

Steps

1. In the left navigation pane, click **Resources**, and then select the Oracle Database plug-in from the list.
2. Click **Add Application Volume**.
3. In the Name page, perform the following actions:
    - In the Name field, enter the name of the application volume.

- In the Host Name field, enter the name of the host.

4. In the Storage Footprint page, enter the storage system name, select one or volumes, and specify the associated LUNs or Qtrees.

   You can add multiple storage systems.

5. Review the summary, and then click **Finish**.
6. In the Resources page, select **Application Volume** from the **View** list to view all the application volumes that you have added.

Modify application volume

You can modify all the values that you specified while adding the application volume, if no backups are created. If the backup is created, you can only modify the storage system details.

Steps

1. In the left navigation pane, click **Resources**, and then select the Oracle Database plug-in from the list.
2. In the Resources page, select **Application Volume** from the **View** list.
3. Click to modify the values.
4. Click to modify the values.

Delete application volume

When you delete an application volume, if there any backups associated with the application volume, the application volume will be put into maintenance mode and no new backups will be created and no earlier backups will be retained. If there are no backups associated, all the metadata will be deleted.

If required, SnapCenter allows you to undo the delete operation.

Steps

1. In the left navigation pane, click **Resources**, and then select the Oracle Database plug-in from the list.
2. In the Resources page, select **Application Volume** from the **View** list.
3. Click elete icon] to modify the values.

Backup application volumes

Back up application volume

If the application volume is not part of any resource group, you can back up the application volume from the Resources page.

About this task

By default, consistency group (CG) backups are created. If you want to create volume based backups, you should set the value of **EnableOracleNdvVolumeBasedBackup** to true in the *web.config* file.

Steps

1. In the left navigation pane, click **Resources**, and then select the Oracle Database plug-in from the list.
2. In the Resources page, select **Application Volume** from the **View** list.
3. Click , and then select the host name and the database type to filter the resources.

   You can then click to close the filter pane.

4. Select the application volume that you want to back up.

   The Application volume-Protect page is displayed.

5. In the Resource page, perform the following actions:

For this field…	Do this…
Use custom name format for Snapshot copy	Select this check box, and then enter a custom name format that you want to use for the Snapshot name. For example, customtext__policy_hostname or resource_hostname. By default, a timestamp is appended to the Snapshot name.
Exclude archive log destinations from backup	Specify the destinations of the archive log files that you do not want to back up.

6. In the Policies page, perform the following steps:

   a. Select one or more policies from the drop-down list.

   You can also create a policy by clicking .

   In the Configure schedules for selected policies section, the selected policies are listed.

   a. Click in the Configure Schedules column for the policy for which you want to configure a schedule.
   b. In the Add schedules for policy *policy_name* window, configure the schedule, and then click **OK**.

      policy_name is the name of the policy that you have selected.

      The configured schedules are listed in the Applied Schedules column.

7. In the Notification page, from the **Email preference** drop-down list, select the scenarios in which you want to send the emails.

   You must also specify the sender and receiver email addresses, and the subject of the email. If you want to attach the report of the backup operation performed on the resource, and then select **Attach Job Report**.

For email notification, you must have specified the SMTP server details using the either the GUI or the PowerShell command Set-SmSmtpServer.

8. Review the summary, and then click **Finish**.

   The application volume topology page is displayed.

9. Click **Back up Now**.
10. In the Backup page, perform the following steps:
    a. If you have applied multiple policies to the resource, from the **Policy** drop-down list, select the policy that you want to use for backup.
    b. Click **Backup**.
11. Monitor the operation progress by clicking **Monitor** > **Jobs**.

Back up the application volumes resource group

You can back up the resource group containing only application volumes or a mix of application volumes and database. A backup operation on the resource group is performed on all resources defined in the resource group.

If the resource group has multiple application volumes, all the application volumes should either have SnapMirror or SnapVault replication policy.

About this task

By default, consistency group (CG) backups are created. If you want to create volume based backups, you should set the value of **EnableOracleNdvVolumeBasedBackup** to true in the *web.config* file.

Steps

1. In the left navigation pane, click **Resources**, and then select the Oracle Database plug-in from the list.
2. In the Resources page, select **Resource Group** from the **View** list.

   You can search the resource group either by entering the resource group name in the search box, or by clicking , and then selecting the tag. You can then click to close the filter pane.

3. In the Resource Groups page, select the resource group that you want to back up, and then click **Back up Now**.
4. In the Backup page, perform the following steps:
    a. If you have associated multiple policies with the resource group, from the **Policy** drop-down list, select the policy that you want to use for backup.

       If the policy selected for the on-demand backup is associated with a backup schedule, the on-demand backups will be retained based on the retention settings specified for the schedule type.

    b. Click **Backup**.
5. Monitor the operation progress by clicking **Monitor** > **Jobs**.

Verification operation will be performed only for the databases and not for the application volumes.

Clone application volume backup

You can use SnapCenter to clone the application volume backups.

Before you begin

If you have installed the plug-in as a non-root user, you should manually assign the execute permissions to the prescript and postscript directories.

Steps

1. In the left navigation pane, click **Resources**, and then select the Oracle Database plug-in from the list.
2. In the Resources page, select **Application Volume** from the **View** list.
3. Select the application volume either from the application volume details view or from the resource group details view.

   The application volume topology page is displayed.

4. From the Manage Copies view, select the backups either from Local copies (primary), Mirror copies (secondary), or Vault copies (secondary).
5. Select the backup from the table, and then click .
6. In the Location page, perform the following actions:

For this field…	Do this…
Plug-in host	Select the host where you want to create the clone.
Target Resource Name	Specify the resource name.

7. In the Scripts page, specify the names of the scripts to be executed before cloning, commands to mount a file system, and names of the scripts to be executed after cloning.
8. In the Notification page, from the **Email preference** drop-down list, select the scenarios in which you want to send the emails.

   You must also specify the sender and receiver email addresses, and the subject of the email. If you want to attach the report of the clone operation performed, select **Attach Job Report**.

   For email notification, you must have specified the SMTP server details using the either the GUI or the PowerShell command Set-SmSmtpServer.

9. Review the summary, and then click **Finish**.

Split an application volume clone

You can use SnapCenter to split a cloned resource from the parent resource. The clone that is split becomes independent of the parent resource.

Steps

1. In the left navigation pane, click **Resources**, and then select the Oracle Database plug-in from the list.

2. In the Resources page, select **Application Volume** from the **View** list.
3. Select the cloned resource and click .
4. Review the estimated size of the clone that is to be split and the required space available on the aggregate, and then click **Start**.
5. Monitor the operation progress by clicking **Monitor** > **Jobs**.

Delete an application volume clone

You can delete clones if you find them no longer necessary. You cannot delete clones that acts like source for other clones.

Steps

1. In the left navigation pane, click **Resources**, and then select the Oracle Database plug-in from the list.
2. In the Resources page, select **Application Volume** from the **View** list.
3. Select the resource or resource group from the list.

   The resource or the resource group topology page is displayed.
4. From the Manage Copies view, select **Clones** either from the primary or secondary (mirrored or replicated) storage systems.
5. Select the clone, and then click .
6. In the Delete Clone page, perform the following actions:
   a. In the **Pre clone delete** field, enter the names of the scripts to be executed before deleting the clone.
   b. In the **Unmount** field, enter the commands to unmount the clone before deleting the clone.
7. Click **OK**.

Protect Windows file systems

SnapCenter Plug-in for Microsoft Windows concepts

SnapCenter Plug-in for Microsoft Windows overview

The SnapCenter Plug-in for Microsoft Windows is a host-side component of the NetApp SnapCenter Software that enables application-aware data protection management of Microsoft file system resources. In addition, it provides storage provisioning, Snapshot consistency, and space reclamation for Windows file systems. The Plug-in for Windows automates file system backup, restore, and cloning operations in your SnapCenter environment.

When the Plug-in for Windows is installed, you can use SnapCenter with NetApp SnapMirror technology to create mirror copies of backup sets on another volume and with NetApp SnapVault technology to perform disk-to-disk backup replication for archival or standards compliance.

- Enables application-aware data protection for other plug-ins that are running in Windows hosts in your SnapCenter environment
- Automates application-aware backup, restore, and clone operations for Microsoft file systems in your SnapCenter environment
- Supports storage provisioning, Snapshot consistency, and space reclamation for Windows hosts

The Plug-in for Windows provisions SMB shares and Windows file systems on physical and RDM LUNs but does not support backup operations for Windows file systems on SMB shares.

What you can do with the SnapCenter Plug-in for Microsoft Windows

When the Plug-in for Windows is installed in your environment, you can use SnapCenter to back up, restore, and clone Windows file systems. You can also perform tasks supporting those operations.

- Discover resources
- Back up Windows file systems
- Schedule backup operations
- Restore file system backups
- Clone file system backups
- Monitor backup, restore, and clone operations

The Plug-in for Windows does not support backup and restore of file systems on SMB shares.

SnapCenter Plug-in for Windows features

The Plug-in for Windows integrates with NetApp Snapshot technology on the storage system. To work with the Plug-in for Windows, you use the SnapCenter interface.

The Plug-in for Windows includes these major features:

- **Unified graphical user interface powered by SnapCenter**

  The SnapCenter interface provides you with standardization and consistency across plug-ins and environments. The SnapCenter interface enables you to complete consistent backup and restore processes across plug-ins, use centralized reporting, use at-a-glance dashboard views, set up role-based access control (RBAC), and monitor jobs across all plug-ins. SnapCenter also offers centralized scheduling and policy management to support backup and clone operations.

- **Automated central administration**

  You can schedule routine file system backups, configure policy-based backup retention, and set up restore operations. You can also proactively monitor your file system environment by configuring SnapCenter to send email alerts.

- **Nondisruptive NetApp Snapshot technology**

  The Plug-in for Windows uses NetApp Snapshot technology. This enables you to back up file systems in seconds and restore them quickly without taking host offline. Snapshots consume minimal storage space.

In addition to these major features, the Plug-in for Windows offers the following benefits:

- Backup, restore, and clone workflow support
- RBAC-supported security and centralized role delegation
- Creation of space-efficient copies of production file systems for testing or data extraction by using NetApp FlexClone technology

  For FlexClone licensing information, see SnapCenter licenses.

- Ability to run multiple backups at the same time across multiple servers
- PowerShell cmdlets for scripting of backup, restore, and clone operations
- Support for backup of file systems and virtual machine disks (VMDKs)
- Support for physical and virtualized infrastructures
- Support for iSCSI, Fibre Channel, FCoE, raw device mapping (RDM), Asymmetric LUN Mapping (ALM), VMDK over NFS and VMFS, and virtual FC
- Support for non-volatile memory express (NVMe) on Windows Server 2022
  - Backup, restore, clone, and verification workflows on VMDK layout created on NVMe over TCP/IP.
  - Supports NVMe firmware version 1.3 starting from ESX 8.0 update 2 and requires Virtual hardware version 21.
  - Windows Server Failover Clustering (WSFC) is not supported for applications over VMDK on NVMe over TCP/IP.
- Supports SnapMirror active sync (initially released as SnapMirror Business Continuity [SM-BC]) that enables business services to continue operating even through a complete site failure, supporting applications to fail over transparently using a secondary copy. Neither manual intervention nor additional scripting is required to trigger a failover with SnapMirror active sync.

How SnapCenter backs up Windows file systems

SnapCenter uses Snapshot technology to back up Windows file system resources that reside on LUNs, CSVs (cluster shared volumes), RDM (raw device mapping) volumes, ALM (asymmetric LUN mapping) in Windows clusters, and VMDKs based on VMFS/NFS (VMware Virtual Machine File System using NFS).

SnapCenter creates backups by creating Snapshots of the file systems. Federated backups, in which a volume contains LUNs from multiple hosts, are faster and more efficient than backups of each individual LUN because only one Snapshot of the volume is created compared to individual Snapshots of each file system.

When SnapCenter creates a Snapshot, the entire storage system volume is captured in the Snapshot. However, the backup is valid only for the host server for which the backup was created.

If data from other host servers resides on the same volume, this data cannot be restored from the Snapshot.

If a Windows file system contains a database, then backing up the file system is not the same as backing up the database. To back up a database, you must use one of the database plug-ins.

Storage types supported by SnapCenter Plug-in for Microsoft Windows

SnapCenter supports a wide range of storage types on both physical machines and virtual machines. You must verify whether support is available for your storage type before installing the package for your host.

SnapCenter provisioning and data protection support is available on Windows Server. For the latest information about supported versions, see the NetApp Interoperability Matrix Tool.

Machine	Storage type	Provision using	Support notes
Physical server	FC-connected LUNs	SnapCenter graphical user interface (GUI) or PowerShell cmdlets	
Physical server	iSCSI-connected LUNs	SnapCenter GUI or PowerShell cmdlets	
Physical server	SMB3 (CIFS) shares residing on a storage virtual machine (SVM)	SnapCenter GUI or PowerShell cmdlets	Support for provisioning only.
VMware VM	RDM LUNs connected by an FC or iSCSI HBA	PowerShell cmdlets	
VMware VM	iSCSI LUNs connected directly to the guest system by the iSCSI initiator	SnapCenter GUI or PowerShell cmdlets	

Machine	Storage type	Provision using	Support notes
VMware VM	Virtual Machine File Systems (VMFS) or NFS datastores	VMware vSphere	
VMware VM	A guest system connected to SMB3 shares residing on an SVM	SnapCenter GUI or PowerShell cmdlets	Support for provisioning only.
VMware VM	vVol datastores on both NFS and SAN	ONTAP Tools for VMware vSphere	
Hyper-V VM	Virtual FC (vFC) LUNs connected by a virtual Fibre Channel Switch	SnapCenter GUI or PowerShell cmdlets	You must use Hyper-V Manager to provision Virtual FC (vFC) LUNs connected by a virtual Fibre Channel Switch. Hyper-V pass through disks and backing up databases on VHD(x) that are provisioned on NetApp storage are not supported.
Hyper-V VM	iSCSI LUNs connected directly to the guest system by the iSCSI initiator	SnapCenter GUI or PowerShell cmdlets	Hyper-V pass through disks and backing up databases on VHD(x) that are provisioned on NetApp storage are not supported.

Machine	Storage type	Provision using	Support notes
Hyper-V VM	A guest system connected to SMB3 shares residing on an SVM	SnapCenter GUI or PowerShell cmdlets	Support for provisioning only. Hyper-V pass through disks and backing up databases on VHD(x) that are provisioned on NetApp storage are not supported.

Minimum ONTAP privileges required for Windows plug-in

The minimum ONTAP privileges that are required vary according to the SnapCenter plug-ins you are using for data protection.

- All-access commands: Minimum privileges required for ONTAP 9.12.1 and later
  - event generate-autosupport-log
  - job history show
  - job stop
  - lun
  - lun create
  - lun delete
  - lun igroup add
  - lun igroup create
  - lun igroup delete
  - lun igroup rename
  - lun igroup show
  - lun mapping add-reporting-nodes
  - lun mapping create
  - lun mapping delete
  - lun mapping remove-reporting-nodes
  - lun mapping show
  - lun modify
  - lun move-in-volume

- lun offline
- lun online
- lun resize
- lun serial
- lun show
- snapmirror policy add-rule
- snapmirror policy modify-rule
- snapmirror policy remove-rule
- snapmirror policy show
- snapmirror restore
- snapmirror show
- snapmirror show-history
- snapmirror update
- snapmirror update-ls-set
- snapmirror list-destinations
- version
- volume clone create
- volume clone show
- volume clone split start
- volume clone split stop
- volume create
- volume destroy
- volume file clone create
- volume file show-disk-usage
- volume offline
- volume online
- volume modify
- volume qtree create
- volume qtree delete
- volume qtree modify
- volume qtree show
- volume restrict
- volume show
- volume snapshot create
- volume snapshot delete
- volume snapshot modify
- volume snapshot rename

    - volume snapshot restore
    - volume snapshot restore-file
    - volume snapshot show
    - volume unmount
    - vserver cifs
    - vserver cifs share create
    - vserver cifs share delete
    - vserver cifs shadowcopy show
    - vserver cifs share show
    - vserver cifs show
    - vserver export-policy
    - vserver export-policy create
    - vserver export-policy delete
    - vserver export-policy rule create
    - vserver export-policy rule show
    - vserver export-policy show
    - vserver iscsi
    - vserver iscsi connection show
    - vserver show
- Read-only commands: Minimum privileges required for ONTAP 8.3.0 and later
    - network interface
    - network interface show
    - vserver

Prepare storage systems for SnapMirror and SnapVault replication

You can use a SnapCenter plug-in with ONTAP SnapMirror technology to create mirror copies of backup sets on another volume, and with ONTAP SnapVault technology to perform disk-to-disk backup replication for standards compliance and other governance-related purposes. Before you perform these tasks, you must configure a data-protection relationship between the source and destination volumes and initialize the relationship.

SnapCenter performs the updates to SnapMirror and SnapVault after it completes the Snapshot operation. SnapMirror and SnapVault updates are performed as part of the SnapCenter job; do not create a separate ONTAP schedule.

If you are coming to SnapCenter from a NetApp SnapManager product and are satisfied with the data protection relationships you have configured, you can skip this section.

A data protection relationship replicates data on primary storage (the source volume) to secondary storage (the destination volume). When you initialize the relationship, ONTAP transfers the data blocks referenced on the source volume to the destination volume.

SnapCenter does not support cascade relationships between SnapMirror and SnapVault volumes (**Primary** > **Mirror** > **Vault**). You should use fanout relationships.

SnapCenter supports the management of version-flexible SnapMirror relationships. For details about version-flexible SnapMirror relationships and how to set them up, see the ONTAP documentation.

Define a backup strategy for Windows file systems

Defining a backup strategy before you create your backups provides you with the backups that you require to successfully restore or clone your file systems. Your service-level agreement (SLA), recovery time objective (RTO), and recovery point objective (RPO) largely determine your backup strategy.

An SLA defines the level of service that is expected and addresses many service-related issues, including the availability and performance of the service. RTO is the time by which a business process must be restored after a disruption in service. RPO defines the strategy for the age of the files that must be recovered from backup storage for regular operations to resume after a failure. SLA, RTO, and RPO contribute to the data protection strategy.

Backup schedules for Windows file systems

Backup frequency is specified in policies; a backup schedule is specified in the resource group configuration. The most critical factor in determining a backup frequency or schedule is the rate of change for the resource and the importance of the data. You might back up a heavily used resource every hour, while you might back up a rarely used resource once a day. Other factors include the importance of the resource to your organization, your Service Level Agreement (SLA), and your Recover Point Objective (RPO).

An SLA defines the level of service expected and addresses many service-related issues, including the availability and performance of service. An RPO defines the strategy for the age of the files that must be recovered from backup storage for regular operations to resume after a failure. The SLA and RPO contribute to the data protection strategy.

Even for a heavily used resource, there is no requirement to run a full backup more than once or twice a day.

Backup schedules have two parts, as follows:

- Backup frequency

  Backup frequency (how often backups are to be performed), called *schedule type* for some plug-ins, is part of a policy configuration. For example, you might configure the backup frequency as hourly, daily, weekly, or monthly, or you can specify **None** which makes the policy an on-demand-only policy. You can access policies by clicking **Settings** > **Policies**.

- Backup schedules

  Backup schedules (exactly when backups are to be performed) are part of a resource group configuration. For example, if you have a resource group that has a policy configured for weekly backups, you might configure the schedule to back up every Thursday at 10:00 PM. You can access resource group schedules by clicking **Resources** > **Resource Groups**.

Number of backups needed for Windows file systems

Factors that determine the number of backups that you need include the size of the Windows file system, the number of volumes used, the rate of change of the file system, and your Service Level Agreement (SLA).

Backup naming convention for Windows file systems

Windows file system backups use the default Snapshot naming convention. The default backup naming convention adds a timestamp to Snapshot names that helps you identify when the copies were created.

The Snapshot uses the following default naming convention: resourcegroupname_hostname_timestamp

You should name your backup resource groups logically, as in the following example:

```
dts1_mach1x88_03-12-2015_23.17.26
```

In this example, the syntax elements have the following meanings:

- `dts1` is the resource group name.
- `mach1x88` is the host name.
- `03-12-2016_23.17.26` is the date and timestamp.

When creating a backup, you can also add a descriptive tag to help identify the backup. In contrast, if you want to use a customized backup naming convention, you need to rename the backup after the backup operation is complete.

Backup retention options

You can choose either the number of days for which to retain backup copies or specify the number of backup copies you want to retain, up to a ONTAP maximum of 255 copies. For example, your organization might require that you retain 10 days of backup copies or 130 backup copies.

While creating a policy, you can specify the retention options for the backup type and the schedule type.

If you set up SnapMirror replication, the retention policy is mirrored on the destination volume.

SnapCenter deletes the retained backups that have retention labels that match the schedule type. If the schedule type was changed for the resource or resource group, backups with the old schedule type label might still remain on the system.

For long-term retention of backup copies, you should use SnapVault backup.

Sources and destinations of clones for Windows file systems

You can clone a file system backup from primary storage or secondary storage. You also can choose the destination that supports your requirements; either the original backup location or a different destination on the same host or on a different host. The destination must be on the same volume as the clone source backup.

Clone destination	Description
Original, source, location	By default, SnapCenter stores the clone on the same location and the same host as the backup being cloned.
Different location	You can store the clone on a different location on the same host or on a different host. The host must have a configured connection to the storage virtual machine (SVM).

You can rename the clone after the clone operation is complete.

Install SnapCenter Plug-in for Microsoft Windows

Installation workflow of SnapCenter Plug-in for Microsoft Windows

You must install and set up SnapCenter Plug-in for Microsoft Windows if you want to protect Windows files that are not database files.

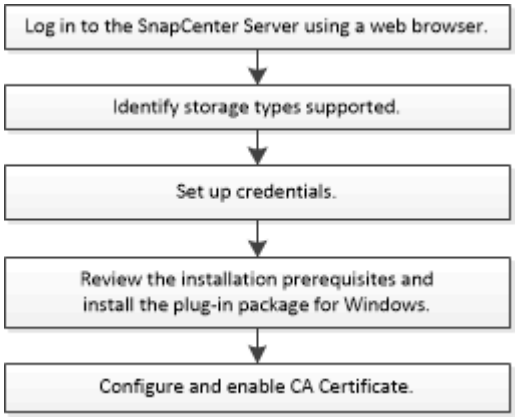

Installation requirements for SnapCenter Plug-in for Microsoft Windows

You should be aware of certain installation requirements before you install the Plug-in for Windows.

Before you begin to use the Plug-in for Windows, the SnapCenter administrator must install and configure SnapCenter Server and perform prerequisite tasks.

- You must have SnapCenter admin privileges to install the Plug-in for Windows.

  The SnapCenter admin role must have admin privileges.

- You must have installed and configured the SnapCenter Server.
- When installing a plug-in on a Windows host, if you specify a credential that is not built-in or if the user belongs to a local workgroup user, you must disable UAC on the host.
- You must set up SnapMirror and SnapVault if you want backup replication.

Host requirements to install SnapCenter Plug-ins Package for Windows

Before you install the SnapCenter Plug-ins Package for Windows, you should be familiar with some basic host system space requirements and sizing requirements.

Item	Requirements
Operating Systems	Microsoft Windows For the latest information about supported versions, see the NetApp Interoperability Matrix Tool. If you are on a Windows cluster setup, you should also install and configure the Windows Remote Management (WinRM).
Minimum RAM for the SnapCenter plug-in on host	1 GB
Minimum install and log space for the SnapCenter plug-in on host	5 GB You should allocate sufficient disk space and monitor the storage consumption by the logs folder. The log space required varies depending on the number of the entities to be protected and the frequency of data protection operations. If there is no sufficient disk space, the logs will not be created for the recently run operations.
Required software packages	• ASP.NET Core Runtime 8.0.12 (and all subsequent 8.0.x patches) Hosting Bundle • PowerShell Core 7.4.2 For the latest information about supported versions, see the NetApp Interoperability Matrix Tool. For .NET specific troubleshooting information, see SnapCenter upgrade or install fails for legacy systems that do not have internet connectivity.

Set up your credentials for the Plug-in for Windows

SnapCenter uses credentials to authenticate users for SnapCenter operations. You should create credentials for installing SnapCenter plug-ins, and additional credentials for performing data protection operations on Windows file systems.

What you will need

- You must set up Windows credentials before installing plug-ins.

- You must set up the credentials with administrator privileges, including administrator rights, on the remote host.
- If you set up credentials for individual resource groups, and the user does not have full admin privileges, you must assign at least the resource group and backup privileges to the user.

Steps

1. In the left navigation pane, click **Settings**.
2. In the Settings page, click **Credential**.
3. Click **New**.
4. In the Credential page, do the following:

For this field…	Do this…
Credential name	Enter a name for the credentials.
User name/Password	Enter the user name and password used for authentication. • Domain administrator or any member of the administrator group Specify the domain administrator or any member of the administrator group on the system on which you are installing the SnapCenter plug-in. Valid formats for the Username field are as follows: ◦ `NetBIOS\UserName` ◦ `Domain FQDN\UserName` ◦ `UserName@upn` • Local administrator (for workgroups only) For systems that belong to a workgroup, specify the built-in local administrator on the system on which you are installing the SnapCenter plug-in. You can specify a local user account that belongs to the local administrators group if the user account has elevated privileges or the User Access control feature is disabled on the host system. The valid format for the Username field is as follows: `UserName` Do not use double quotes (") or backtick (`) in the passwords. You should not use the less than (<) and exclamation (!) symbols together in passwords. For example, lessthan<!10, lessthan10<!, backtick`12.

For this field…	Do this…
Password	Enter the password used for authentication.

5. Click **OK**.

   After you finish setting up credentials, you might want to assign credential maintenance to a user or group of users on the User and Access page.

Configure gMSA on Windows Server 2016 or later

Windows Server 2016 or later enables you to create a group Managed Service Account (gMSA) that provides automated service account password management from a managed domain account.

Before you begin

- You should have a Windows Server 2016 or later domain controller.
- You should have a Windows Server 2016 or later host, which is a member of the domain.

Steps

1. Create a KDS root key to generate unique passwords for each object in your gMSA.
2. For each domain, run the following command from the Windows domain controller: Add-KDSRootKey -EffectiveImmediately
3. Create and configure your gMSA:
   a. Create a user group account in the following format:

      ```
      domainName\accountName$
      ```

   b. Add computer objects to the group.
   c. Use the user group you just created to create the gMSA.

      For example,

      ```
      New-ADServiceAccount -name <ServiceAccountName> -DNSHostName <fqdn>
      -PrincipalsAllowedToRetrieveManagedPassword <group>
      -ServicePrincipalNames <SPN1,SPN2,…>
      ```

   d. Run `Get-ADServiceAccount` command to verify the service account.
4. Configure the gMSA on your hosts:
   a. Enable the Active Directory module for Windows PowerShell on the host where you want to use the gMSA account.

      To do this, run the following command from PowerShell:

```
PS C:\> Get-WindowsFeature AD-Domain-Services

Display Name                            Name                Install
State
------------                            ----
-------------
[ ] Active Directory Domain Services   AD-Domain-Services  Available


PS C:\> Install-WindowsFeature AD-DOMAIN-SERVICES

Success Restart Needed Exit Code      Feature Result
------- -------------- ---------      --------------
True    No             Success        {Active Directory Domain
Services, Active ...
WARNING: Windows automatic updating is not enabled. To ensure that
your newly-installed role or feature is
automatically updated, turn on Windows Update.
```

b. Restart your host.

c. Install the gMSA on your host by running the following command from the PowerShell command prompt: `Install-AdServiceAccount <gMSA>`

d. Verify your gMSA account by running the following command: `Test-AdServiceAccount <gMSA>`

5. Assign the administrative privileges to the configured gMSA on the host.
6. Add the Windows host by specifying the configured gMSA account in the SnapCenter Server.

   SnapCenter Server will install the selected plug-ins on the host and the specified gMSA will be used as the service log on account during the plug-in installation.

Add hosts and install SnapCenter Plug-in for Microsoft Windows

You can use the SnapCenter Add Host page to add Windows hosts. The SnapCenter Plug-in for Microsoft Windows is automatically installed on the specified host. This is the recommended method for installing plug-ins. You can add a host and install a plug-in either for an individual host or a cluster.

Before you begin

- If the operating system of the SnapCenter Server host is Windows 2019 and the operating system of the plug-in host is Windows 2022, you should perform the following:
  - Upgrade to Windows Server 2019 (OS Build 17763.5936) or later
  - Upgrade to Windows Server 2022 (OS Build 20348.2402) or later
- You must be a user that is assigned to a role that has the plug-in install and uninstall permissions, such as the SnapCenter Admin role.
- When installing a plug-in on a Windows host, if you specify a credential that is not built-in or if the user

belongs to a local workgroup user, you must disable UAC on the host.

- The SnapCenter user should be added to the “Log on as a service” role of the Windows Server.
- You should ensure that the message queueing service is in running state.
- If you are using group Managed Service Account (gMSA), you should configure gMSA with administrative privileges.

  Configure group Managed Service Account on Windows Server 2016 or later for Windows File System

About this task

- You cannot add a SnapCenter Server as a plug-in host to another SnapCenter Server.
- Windows plug-ins
  - Microsoft Windows
  - Microsoft Exchange Server
  - Microsoft SQL Server
  - SAP HANA
- Installing plug-ins on a cluster

  If you install plug-ins on a cluster (WSFC, Oracle RAC, or Exchange DAG), they are installed on all of the nodes of the cluster.

- E-series storage

  You cannot install the Plug-in for Windows on a Windows host connected to E-series storage.

SnapCenter does not support adding of the same host (plug-in host) to SnapCenter if the host is already part of a workgroup and changed to another domain or vice versa.
If you want to add the same host, you should remove the host from SnapCenter and add it again.

Steps

1. In the left navigation pane, click **Hosts**.
2. Ensure that **Managed Hosts** is selected at the top.
3. Click **Add**.
4. In the Hosts page, do the following:

For this field…	Do this…
Host Type	Select the **Windows** type of host. SnapCenter Server adds the host and then installs the Plug-in for Windows if it is not already installed on the host.

For this field…	Do this…
Host name	Enter the fully qualified domain name (FQDN) or the IP address of the host. SnapCenter depends on the proper configuration of the DNS. Therefore, the best practice is to enter the fully qualified domain name (FQDN). You can enter the IP addresses or FQDN of one of the following: • Stand-alone host • Windows Server Failover Clustering (WSFC) If you are adding a host using SnapCenter and it is part of a subdomain, you must provide the FQDN.
Credentials	Select the credential name that you created or create the new credentials. The credential must have administrative rights on the remote host. For details, see information about creating a credential. Details about credentials, including the user name, domain, and host type, are displayed by placing your cursor over the credential name you provided. The authentication mode is determined by the host type that you specify in the Add Host wizard.

5. In the Select Plug-ins to Install section, select the plug-ins to install.

   For new deployments, no plug-in packages are listed.

6. (Optional) Click **More Options**.

For this field…	Do this…
Port	Either retain the default port number or specify the port number. The default port number is 8145. If the SnapCenter Server was installed on a custom port, that port number will be displayed as the default port. If you manually installed the plug-ins and specified a custom port, you must specify the same port. Otherwise, the operation fails.
Installation Path	The default path is C:\Program Files\NetApp\SnapCenter. You can optionally customize the path. For SnapCenter Plug-ins Package for Windows, the default path is C:\Program Files\NetApp\SnapCenter. However, if you want, you can customize the default path.
Add all hosts in the cluster	Select this check box to add all of the cluster nodes in a WSFC.
Skip preinstall checks	Select this check box if you already installed the plug-ins manually and you do not want to validate whether the host meets the requirements for installing the plug-in.
Use group Managed Service Account (gMSA) to run the plug-in services	Select this check box if you want to use group Managed Service Account (gMSA) to run the plug-in services. Provide the gMSA name in the following format: *domainName\accountName$*. gMSA will be used as a log on service account only for SnapCenter Plug-in for Windows service.

7. Click **Submit**.

   If you have not selected the **Skip prechecks** checkbox, the host is validated to see whether it meets the requirements to install the plug-in. The disk space, RAM, PowerShell version, .NET version, and location are validated against the minimum requirements. If the minimum requirements are not met, appropriate error or warning messages are displayed.

   If the error is related to disk space or RAM, you can update the web.config file located at `C:\Program Files\NetApp\SnapCenter` WebApp to modify the default values. If the error is related to other

parameters, you must fix the issue.

> In an HA setup, if you are updating web.config file, you must update the file on both nodes.

8. Monitor the installation progress.

Install SnapCenter Plug-in for Microsoft Windows on multiple remote hosts using PowerShell cmdlets

If you want to install SnapCenter Plug-in for Microsoft Windows on multiple hosts at one time, you can do so by using the `Install-SmHostPackage` PowerShell cmdlet.

You must have logged in to SnapCenter as a domain user with local administrator rights on each host on which you want to install plug-ins.

Steps

1. Launch PowerShell.
2. On the SnapCenter Server host, establish a session using the `Open-SmConnection` cmdlet, and then enter your credentials.
3. Add the standalone host or the cluster to SnapCenter using the `Add-SmHost` cmdlet and the required parameters.

   The information regarding the parameters that can be used with the cmdlet and their descriptions can be obtained by running *Get-Help command_name*. Alternatively, you can also refer to the SnapCenter Software Cmdlet Reference Guide.

4. Install the plug-in on multiple hosts using the `Install-SmHostPackage` cmdlet and the required parameters.

   You can use the `-skipprecheck` option when you have installed the plug-ins manually and do not want to validate whether the host meets the requirements to install the plug-in.

Install the SnapCenter Plug-in for Microsoft Windows silently from the command line

You can install the SnapCenter Plug-in for Microsoft Windows locally on a Windows host if you are unable to install the plug-in remotely from the SnapCenter GUI. You can run the SnapCenter Plug-in for Microsoft Windows installation program unattended, in silent mode, from the Windows command line.

Before you begin

- You must have installed ASP.NET Core Runtime 8.0.12 (and all subsequent 8.0.x patches) Hosting Bundle.
- You must have installed PowerShell 7.4.2 or later.
- You must be a local administrator on the host.

Steps

1. Download the SnapCenter Plug-in for Microsoft Windows from your install location.

   For example, the default installation path is C:\ProgramData\NetApp\SnapCenter\Package Repository.

This path is accessible from the host where the SnapCenter Server is installed.

2. Copy the installation file to the host on which you want to install the plug-in.
3. From the command prompt, navigate to the directory where you downloaded the installation file.
4. Enter the following command, replacing variables with your data:

```
"snapcenter_windows_host_plugin.exe"/silent / debuglog"" /log""
BI_SNAPCENTER_PORT= SUITE_INSTALLDIR="" BI_SERVICEACCOUNT= BI_SERVICEPWD=
ISFeatureInstall=SCW
```

For example:

```
`"C:\ProgramData\NetApp\SnapCenter\Package Repository
\snapcenter_windows_host_plugin.exe"/silent /debuglog"C:
\HPPW_SCW_Install.log" /log"C:\" BI_SNAPCENTER_PORT=8145
SUITE_INSTALLDIR="C: \Program Files\NetApp\SnapCenter"
BI_SERVICEACCOUNT=domain\administrator BI_SERVICEPWD=password
ISFeatureInstall=SCW`
```

All the parameters passed during the installation of Plug-in for Windows are case sensitive.

Enter the values for the following variables:

Variable	Value
/debuglog"<Debug_Log_Path>	Specify the name and location of the suite installer log file, as in the following example: Setup.exe /debuglog"C:\PathToLog\setupexe.log".
BI_SNAPCENTER_PORT	Specify the port on which SnapCenter communicates with SMCore.
SUITE_INSTALLDIR	Specify host plug-in package installation directory.
BI_SERVICEACCOUNT	Specify SnapCenter Plug-in for Microsoft Windows web service account.
BI_SERVICEPWD	Specify the password for SnapCenter Plug-in for Microsoft Windows web service account.
ISFeatureInstall	Specify the solution to be deployed by SnapCenter on remote host.

The *debuglog* parameter includes the path of the log file for SnapCenter. Writing to this log file is the preferred method of obtaining troubleshooting information, because the file contains the results of checks that the installation performs for plug-in prerequisites.

If necessary, you can find additional troubleshooting information in the log file for the SnapCenter for

Windows package. Log files for the package are listed (oldest first) in the *%Temp%* folder, for example, *C:\temp*\.

The installation of Plug-in for Windows registers the plug-in on the host and not on the SnapCenter Server. You can register the plug-in on the SnapCenter Server by adding the host using the SnapCenter GUI or PowerShell cmdlet. After the host is added, the plug-in is automatically discovered.

Monitor SnapCenter plug-in package installation status

You can monitor the progress of SnapCenter plug-in package installation by using the Jobs page. You might want to check the progress of installation to determine when it is complete or if there is an issue.

About this task

The following icons appear on the Jobs page and indicate the state of the operation:

- In progress
- Completed successfully
- Failed
- Completed with warnings or could not start due to warnings
- Queued

Steps

1. In the left navigation pane, click **Monitor**.
2. In the **Monitor** page, click **Jobs**.
3. In the **Jobs** page, to filter the list so that only plug-in installation operations are listed, do the following:
   a. Click **Filter**.
   b. Optional: Specify the start and end date.
   c. From the Type drop-down menu, select **Plug-in installation**.
   d. From the Status drop-down menu, select the installation status.
   e. Click **Apply**.
4. Select the installation job and click **Details** to view the job details.
5. In the **Job Details** page, click **View logs**.

Configure CA certificate

Generate CA Certificate CSR file

You can generate a Certificate Signing Request (CSR) and import the certificate that can be obtained from a Certificate Authority (CA) using the generated CSR. The certificate will have a private key associated with it.

CSR is a block of encoded text that is given to an authorized certificate vendor to procure the signed CA

certificate.

CA Certificate RSA key length must be minimum 3072 bits.

For information to generate a CSR, see How to generate CA Certificate CSR file.

If you own the CA certificate for your domain (*.domain.company.com) or your system (machine1.domain.company.com), you can skip generating the CA Certificate CSR file. You can deploy the existing CA certificate with SnapCenter.

For cluster configurations, the cluster name (virtual cluster FQDN), and the respective host names should be mentioned in the CA certificate. The certificate can be updated by filling the Subject Alternative Name (SAN) field before procuring the certificate. For a wild card certificate (*.domain.company.com), the certificate will contain all the hostnames of the domain implicitly.

Import CA certificates

You must import the CA certificates to the SnapCenter Server and the Windows host plug-ins using the Microsoft management console (MMC).

Steps

1. Go to the Microsoft management console (MMC), and then click **File** > **Add/Remove Snapin**.
2. In the Add or Remove Snap-ins window, select **Certificates** and then click **Add**.
3. In the Certificates snap-in window, select the **Computer account** option, and then click **Finish**.
4. Click **Console Root** > **Certificates – Local Computer** > **Trusted Root Certification Authorities** > **Certificates**.
5. Right-click on the folder “Trusted Root Certification Authorities”, and then select **All Tasks** > **Import** to start the import wizard.
6. Complete the wizard, as follows:

In this wizard window…	Do the following…
Import Private Key	Select the option **Yes**, import the private key, and then click **Next**.
Import File Format	Make no changes; click **Next**.
Security	Specify the new password to be used for the exported certificate, and then click **Next**.
Completing the Certificate Import Wizard	Review the summary, and then click **Finish** to start the import.

Importing certificate should be bundled with the private key (supported formats are: *.pfx, *.p12, and *.p7b).

7. Repeat Step 5 for the “Personal” folder.

Get the CA certificate thumbprint

A certificate thumbprint is a hexadecimal string that identifies a certificate. A thumbprint is calculated from the content of the certificate using a thumbprint algorithm.

Steps

1. Perform the following on the GUI:
   a. Double-click the certificate.
   b. In the Certificate dialog box, click the **Details** tab.
   c. Scroll through the list of fields and click **Thumbprint**.
   d. Copy the hexadecimal characters from the box.
   e. Remove the spaces between the hexadecimal numbers.

      For example, if the thumbprint is: "a9 09 50 2d d8 2a e4 14 33 e6 f8 38 86 b0 0d 42 77 a3 2a 7b", after removing the spaces, it will be: "a909502dd82ae41433e6f83886b00d4277a32a7b".

2. Perform the following from PowerShell:
   a. Run the following command to list the thumbprint of the installed certificate and identify the recently installed certificate by the subject name.

      Get-ChildItem -Path Cert:\LocalMachine\My

   b. Copy the thumbprint.

Configure CA certificate with Windows host plug-in services

You should configure the CA certificate with Windows host plug-in services to activate the installed digital certificate.

Perform the following steps on the SnapCenter Server and all the plug-in hosts where CA certificates are already deployed.

Steps

1. Remove the existing certificate binding with SMCore default port 8145, by running the following command:

   ```
   > netsh http delete sslcert ipport=0.0.0.0:_<SMCore Port>
   ```

   For example:

   ```
   > netsh http delete sslcert ipport=0.0.0.0:8145
   ```

2. Bind the newly installed certificate with the Windows host plug-in services, by running the following commands:

```
> $cert = "_<certificate thumbprint>_"
> $guid = [guid]::NewGuid().ToString("B")
> netsh http add sslcert ipport=0.0.0.0: _<SMCore Port>_ certhash=$cert
appid="$guid"
```

For example:

```
> $cert = "a909502dd82ae41433e6f83886b00d4277a32a7b"
> $guid = [guid]::NewGuid().ToString("B")
> netsh http add sslcert ipport=0.0.0.0: _<SMCore Port>_ certhash=$cert
appid="$guid"
```

Enable CA Certificates for plug-ins

You should configure the CA certificates and deploy the CA certificates in the SnapCenter Server and the corresponding plug-in hosts. You should enable the CA certificate validation for the plug-ins.

Before you begin

- You can enable or disable the CA certificates using the run *Set-SmCertificateSettings* cmdlet.
- You can display the certificate status for the plug-ins using the *Get-SmCertificateSettings*.

The information regarding the parameters that can be used with the cmdlet and their descriptions can be obtained by running *Get-Help command_name*. Alternatively, you can also refer to the SnapCenter Software Cmdlet Reference Guide.

Steps

1. In the left navigation pane, click **Hosts**.
2. In the Hosts page, click **Managed Hosts**.
3. Select single or multiple plug-in hosts.
4. Click **More options**.
5. Select **Enable Certificate Validation**.

After you finish

The Managed Hosts tab host displays a padlock and the color of the padlock indicates the status of the connection between SnapCenter Server and the plug-in host.

- indicates that the CA certificate is neither enabled nor assigned to the plug-in host.
- indicates that the CA certificate is successfully validated.
- indicates that the CA certificate could not be validated.
- indicates that the connection information could not be retrieved.

When the status is yellow or green, the data protection operations completes successfully.

Install SnapCenter Plug-in for VMware vSphere

If your database or filesystem is stored on virtual machines (VMs), or if you want to protect VMs and datastores, you must deploy the SnapCenter Plug-in for VMware vSphere virtual appliance.

For information to deploy, see Deployment Overview.

Deploy CA certificate

To configure the CA Certificate with SnapCenter Plug-in for VMware vSphere, see Create or import SSL certificate.

Configure the CRL file

SnapCenter Plug-in for VMware vSphere looks for the CRL files in a pre-configured directory. Default directory of the CRL files for SnapCenter Plug-in for VMware vSphere is */opt/netapp/config/crl*.

You can place more than one CRL file in this directory. The incoming certificates will be verified against each CRL.

Back up Windows file systems

Back up Windows file systems

When you install the SnapCenter Plug-in for Microsoft Windows in your environment, you can use SnapCenter to back up Windows file systems. You can back up a single file system or a resource group that contains multiple file systems. You can back up on demand or according to a defined protection schedule.

You can schedule multiple backups to run across servers simultaneously. Backup and restore operations cannot be performed simultaneously on the same resource.

The following workflow shows the sequence in which you must perform the backup operations:

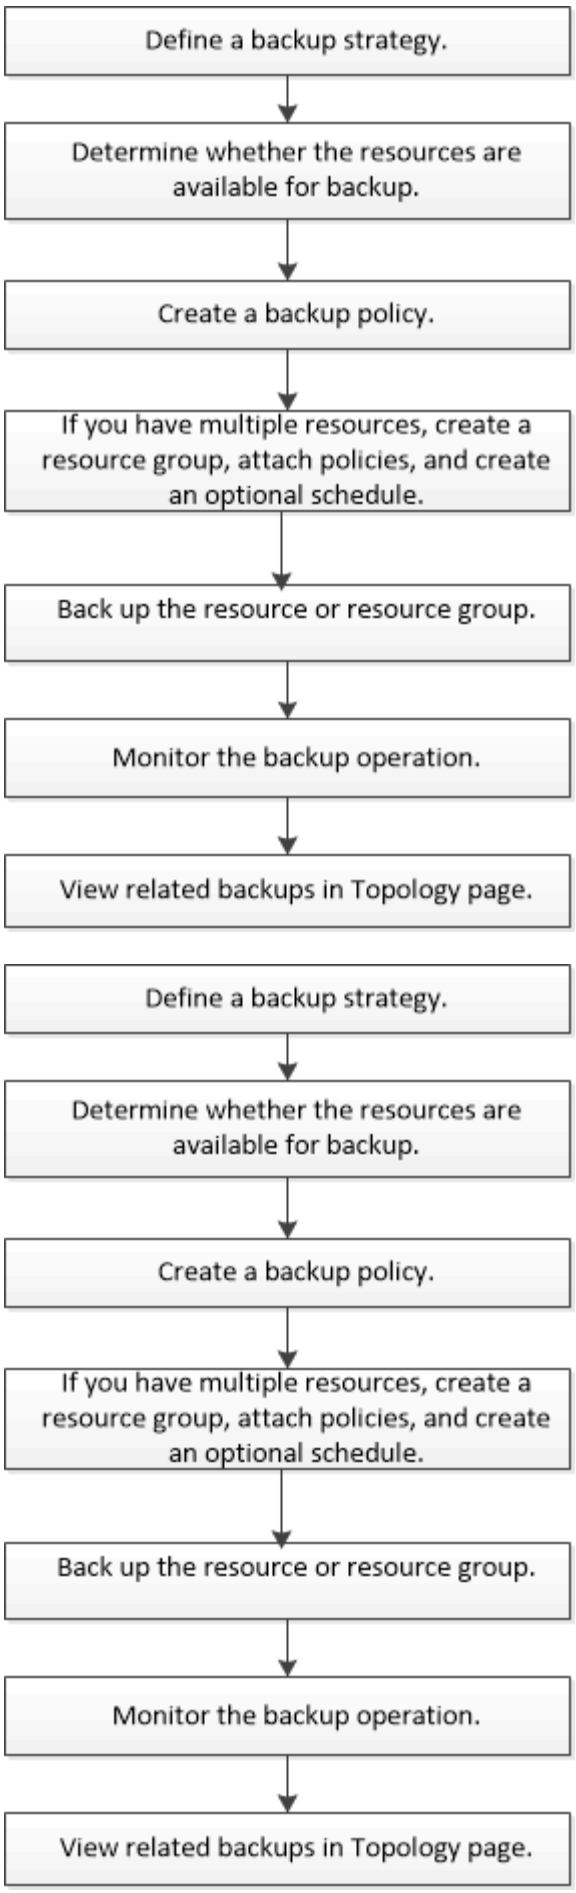
Define a backup strategy.
Determine whether the resources are available for backup.
Create a backup policy.
If you have multiple resources, create a resource group, attach policies, and create an optional schedule.
Back up the resource or resource group.
Monitor the backup operation.
View related backups in Topology page.
Define a backup strategy.
Determine whether the resources are available for backup.
Create a backup policy.
If you have multiple resources, create a resource group, attach policies, and create an optional schedule.
Back up the resource or resource group.
Monitor the backup operation.
View related backups in Topology page.

You can also use PowerShell cmdlets manually or in scripts to perform backup, restore, and clone operations. The SnapCenter cmdlet help or the SnapCenter Software Cmdlet Reference Guide contains detailed information about PowerShell cmdlets.

Determine resource availability for Windows file systems

Resources are the LUNs and similar components in your file system that are maintained by the plug-ins you have installed. You can add those resources to resource groups so that you can perform data protection jobs on multiple resources, but first you must identify which resources you have available. Discovering available resources also verifies that the plug-in installation was completed successfully.

Before you begin

- You must have already completed tasks such as installing SnapCenter Server, adding hosts, creating storage virtual machine (SVM) connections, and adding credentials.
- If files reside on VMware RDM LUNs or VMDKs, you must deploy the SnapCenter Plug-in for VMware vSphere and register the plug-in with SnapCenter. For more information, see SnapCenter Plug-in for VMware vSphere documentation.

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page, select **File Systems** from the list.
3. Select the host to filter the list of resources, and then click **Refresh Resources**.

   The newly added, renamed, or deleted file systems are updated to the SnapCenter Server inventory.

   You must refresh the resources if the databases are renamed outside of SnapCenter.

Create backup policies for Windows file systems

You can create a new backup policy for resources before you use SnapCenter to back up Windows file systems, or you can create a new backup policy at the time you create a resource group or when you back up a resource.

Before you begin

- You must have defined your backup strategy. Learn more
- You must have prepared for data protection.

  To prepare for data protection, you must complete tasks such as installing SnapCenter, adding hosts, discovering resources, and creating storage virtual machine (SVM) connections.

- If you are replicating Snapshots to a mirror or vault secondary storage, the SnapCenter administrator must have assigned the SVMs to you for both the source and destination volumes.
- If you want to run the PowerShell scripts in prescripts and postscripts, you should set the value of the usePowershellProcessforScripts parameter to true in the web.config file.

  The default value is false

- Review the SnapMirror active sync specific prerequisites and limitations. For information, refer Object limits for SnapMirror active sync.

About this task

- The SCRIPTS_PATH is defined using the PredefinedWindowsScriptsDirectory key located in the SMCoreServiceHost.exe.Config file of the plug-in host.

  If needed, you can change this path and restart SMcore service. It is recommended that you use the default path for security.

  The value of the key can be displayed from swagger through the API: API /4.7/configsettings

  You can use the GET API to display the value of the key. SET API is not supported.

- SnapLock
  - If 'Retain the backup copies for a specific number of days' option is selected, then the SnapLock retention period must be lesser than or equal to the mentioned retention days.
  - Specifying a Snapshot locking period prevents deletion of the Snapshots until the retention period expires. This could lead to retaining a larger number of Snapshots than the count specified in the policy.
  - For ONTAP 9.12.1 and below version, the clones created from the SnapLock Vault Snapshots as part of restore will inherit the SnapLock Vault expiry time. Storage admin should manually cleanup the clones post the SnapLock expiry time.

Steps

1. In the left navigation pane, select **Settings**.
2. In the Settings page, select **Policies**.
3. Select **New**.
4. In the Name page, enter the policy name and details.
5. In the Backup and Replication page, perform the following tasks:
   a. Select a backup setting.

Option	Description
File System Consistent Backup	Choose this option if you want SnapCenter to quiesce the disk drive on which the file system resides before the backup operation begins and then resume the disk drive after the backup operation ends.
File System Crash-consistent Backup	Choose this option if you do not want SnapCenter to quiesce the disk drive on which the file system resides.

   b. Select a schedule frequency (also called a policy type).

      The policy specifies the backup frequency only. The specific protection schedule for backing up is defined in the resource group. Therefore, two or more resource groups can share the same policy and backup frequency but have different backup schedules.

If you have scheduled for 2:00 a.m., the schedule will not be triggered during daylight saving time (DST).

c. Select a policy label.

Depending on the Snapshot label that you select, ONTAP applies the secondary Snapshot retention policy that matches the label.

If you have selected **Update SnapMirror after creating a local Snapshot copy**, you can optionally specify the secondary policy label. However, if you have selected **Update SnapVault after creating a local Snapshot copy**, you should specify the secondary policy label.

6. In the Select secondary replication options section, select one or both of the following secondary replication options:

For this field…	Do this…
Update SnapMirror after creating a local Snapshot copy	Select this option to create mirror copies of backup sets on another volume (SnapMirror). This option should be enabled for SnapSnapMirror active sync. During secondary replication, the SnapLock expiry time loads the primary SnapLock expiry time. Clicking the **Refresh** button in the Topology page refreshes the secondary and primary SnapLock expiry time that are retrieved from ONTAP. See View related backups and clones in the Topology page.
Update SnapVault after creating a Snapshot copy	Select this option to perform disk-to-disk backup replication. During secondary replication, the SnapLock expiry time loads the primary SnapLock expiry time. Clicking the Refresh button in the Topology page refreshes the secondary and primary SnapLock expiry time that are retrieved from ONTAP. When SnapLock is configured only on the secondary from ONTAP known as SnapLock Vault, clicking the Refresh button in the Topology page refreshes the locking period on the secondary that is retrieved from ONTAP. For more information on SnapLock Vault see Commit Snapshot copies to WORM on a vault destination

For this field…	Do this…
Error retry count	Enter the number of replication attempts that should occur before the process halts.

You should configure SnapMirror retention policy in ONTAP for the secondary storage to avoid reaching the maximum limit of Snapshots on the secondary storage.

7. In the Retention settings page, specify the retention settings for on-demand backups and for each schedule frequency you selected.

Option	Description
Total Snapshot copies to keep	Choose this option if you want to specify the number of Snapshots SnapCenter stores before automatically deleting them.
Keep Snapshot copies for	Choose this option if you want to specify the number of days SnapCenter retains a backup copy before deleting it.
Snapshot copy locking period	Select Snapshot locking period, and specify the duration days, months, or years. SnapLock retention period should be less than 100 years.

You should set the retention count to 2 or higher. The minimum value for the retention count is 2.

The maximum retention value is 1018. Backups will fail if retention is set to a value higher than what the ONTAP version supports.

8. In the Script page, enter the path of the prescript or postscript that you want the SnapCenter Server to run before or after the backup operation, respectively and a time limit that SnapCenter waits for the script to execute before timing out.

   For example, you can run a script to update SNMP traps, automate alerts, and send logs.

   The prescripts or postscripts path should not include drives or shares. The path should be relative to the SCRIPTS_PATH.

9. Review the summary, and then click **Finish**.

Create resource groups for Windows file systems

A resource group is the container to which you can add multiple file systems that you want to protect. You must also attach one or more policies to the resource group to define the type of data protection job that you want to perform, and then specify the backup

schedule.

About this task

- For ONTAP 9.12.1 and below version, the clones created from the SnapLock Vault Snapshots as part of restore will inherit the SnapLock Vault expiry time. Storage admin should manually cleanup the clones post the SnapLock expiry time.
- Adding new filesystems without SnapMirror active sync to an existing resource group which contains resources with SnapMirror active sync, is not supported.
- Adding new filesystems to an existing resource group in failover mode of SnapMirror active sync is not supported. You can add resources to the resource group only in regular or fail-back state.

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page, select **File Systems** from the list.

   If you have recently added a file system to SnapCenter, click **Refresh Resources** to view the newly added resource.

3. Click **New Resource Group**.
4. In the Name page in the wizard, do the following:

For this field…	Do this…
Name	Enter the resource group name. The resource group name should not exceed 250 characters.
Use custom name format for Snapshot copy	Optional: Enter a custom Snapshot name and format. For example, customtext_resourcegroup_policy_hostname or resourcegroup_hostname. By default, a timestamp is appended to the Snapshot name.
Tag	Enter a descriptive tag to help when finding a resource group.

5. In the Resources page, perform the following tasks:
   a. Select the host to filter the list of resources.

      If you have recently added resources, they will appear on the list of available resources only after you refresh your resource list.

   b. In the Available Resources section, click the file systems that you want to back up, and then click the right arrow to move them to the Added section.

      If you select the **Autoselect all resources on same storage volume** option, all of the resources on

the same volume are selected. When you move them to the Added section, all of the resources on that volume move together.

To add a single file system, clear the **Autoselect all resources on same storage volume** option and then select the file systems you want to move to the Added section.

6. In the Policies page, perform the following tasks:

   a. Select one or more policies from the drop-down list.

      You can select any existing policy and click **Details** to determine whether you can use that policy.

      If no existing policy meets your requirements, you can create a new policy by clicking to start the policy wizard.

      The selected policies are listed in the Policy column in the Configure schedules for selected policies section.

   b. In the Configure schedules for selected policies section, click in the Configure Schedules column for the policy for which you want to configure the schedule.

   c. If the policy is associated with multiple schedule types (frequencies), select the frequency that you want to configure.

   d. In the Add schedules for policy *policy_name* dialog box, configure the schedule by specifying the start date, expiration date, and frequency, and then click **Finish**.

      The configured schedules are listed in the Applied Schedules column in the Configure schedules for selected policies section.

   Third party backup schedules are not supported when they overlap with SnapCenter backup schedules. You should not modify the schedules from the Windows task scheduler and SQL Server Agent.

7. In the Notification page, provide notification information, as follows:

For this field…	Do this…
Email preference	Select **Always**, **On Failure**, or **On failure or warning**, to send emails to recipients after creating backup resource groups, attaching policies, and configuring schedules. Enter the SMTP server, default email subject line, and the To and From email addresses.
From	Email address
To	Email to address
Subject	Default email subject line

8. Review the summary, and then click **Finish**.

   You can perform a backup on demand or wait for the scheduled backup to occur.

Create resource groups and enable secondary protection for Windows file systems on ASA r2 systems

You should create the resource group to add the resources that are on ASA r2 systems. You can also provision the secondary protection while creating the resource group.

Before you begin

- You should ensure that you are not adding both ONTAP 9.x resources and ASA r2 resources to the same resource group.
- You should ensure that you do not have a database with both ONTAP 9.x resources and ASA r2 resources.

About this task

- The secondary protection is available only if the logged-in user is assigned to the role that has the **SecondaryProtection** capability enabled.
- If you enabled secondary protection, the resource group is put into maintenance mode while creating the primary and secondary consistency groups. After the primary and secondary consistency groups are created, the resource group is put out of maintenance mode.
- SnapCenter does not support secondary protection for a clone resource.

Steps

1. In the left navigation pane, select **Resources**, and the appropriate plug-in from the list.
2. In the Resources page, click **New Resource Group**.
3. In the Name page, perform the following actions:
   a. Enter a name for the resource group in the Name field.

      The resource group name should not exceed 250 characters.

   b. Enter one or more labels in the Tag field to help you search for the resource group later.

      For example, if you add HR as a tag to multiple resource groups, you can later find all resource groups associated with the HR tag.

   c. Select this check box, and enter a custom name format that you want to use for the Snapshot name.

      For example, customtext_resource group_policy_hostname or resource group_hostname. By default, a timestamp is appended to the Snapshot name.

   d. Specify the destinations of the archive log files that you do not want to back up.

      You should use the exact same destination as it was set in the application including prefix, if needed.

4. In the Resources page, select the database host name from the **Host** drop-down list.

   The resources are listed in the Available Resources section only if the resource is discovered successfully. If you have recently added resources, they will appear on the list of available resources only after you refresh your resource list.

5. Select the ASA r2 resources from the Available Resources section and move them to the Selected

Resources section.

6. In the Application Settings page, select the backup option.
7. In the Policies page, perform the following steps:
   a. Select one or more policies from the drop-down list.

      You can also create a policy by clicking + .

      In the Configure schedules for selected policies section, the selected policies are listed.

   b. Click + in the Configure Schedules column for the policy for which you want to configure a schedule.
   c. In the Add schedules for policy *policy_name* window, configure the schedule, and then click **OK**.

      Where, *policy_name* is the name of the policy that you have selected.

      The configured schedules are listed in the Applied Schedules column.

   Third party backup schedules are not supported when they overlap with SnapCenter backup schedules.

8. If the secondary protection is enabled for the policy that you have selected, then Secondary Protection page is displayed and you need to perform the following steps:
   a. Select the type of the replication policy.

      Synchronous replication policy is not supported.

   b. Specify the consistency group suffix that you want to use.
   c. From the Destination Cluster and Destination SVM drop-downs select the peered cluster and SVM that you want to use.

      The cluster and SVM peering is not supported by SnapCenter. You should use System Manager or ONTAP CLIs to perform cluster and SVM peering.

If the resources are already protected outside of SnapCenter, those resources will be displayed in the Secondary Protected Resources section.

1. On the Verification page, perform the following steps:
   a. Click **Load locators** to load the SnapMirror or SnapVault volumes to perform verification on secondary storage.
   b. Click + in the Configure Schedules column to configure the verification schedule for all the schedule types of the policy.
   c. In the Add Verification Schedules policy_name dialog box, perform the following actions:

If you want to…	Do this…
Run verification after backup	Select **Run verification after backup**.
Schedule a verification	Select **Run scheduled verification** and then select the schedule type from the drop-down list.

d. Select **Verify on secondary location** to verify your backups on secondary storage system.

e. Click **OK**.

The configured verification schedules are listed in the Applied Schedules column.

2. In the Notification page, from the **Email preference** drop-down list, select the scenarios in which you want to send the emails.

You must also specify the sender and receiver email addresses, and the subject of the email. If you want to attach the report of the operation performed on the resource group, select **Attach Job Report**.

For email notification, you must have specified the SMTP server details using the either the GUI or the PowerShell command Set-SmSmtpServer.

3. Review the summary, and then click **Finish**.

Create a storage system connection and a credential using PowerShell cmdlets

You must create a storage virtual machine (SVM) connection and a credential before using PowerShell cmdlets to perform data protection operations.

Before you begin

- You should have prepared the PowerShell environment to execute the PowerShell cmdlets.
- You should have the required permissions in the Infrastructure Admin role to create storage connections.
- You should ensure that the plug-in installations are not in progress.

  Host plug-in installations must not be in progress while adding a storage system connection because the host cache might not be updated and databases status might be displayed in the SnapCenter GUI as “Not available for backup” or “Not on NetApp storage”.

- Storage system names should be unique.

  SnapCenter does not support multiple storage systems with the same name on different clusters. Each storage system that is supported by SnapCenter should have a unique name and a unique management LIF IP address.

Steps

1. Initiate a PowerShell Core connection session by using the Open-SmConnection cmdlet.

This example opens a PowerShell session:

```
PS C:\> Open-SmConnection
```

2. Create a new connection to the storage system by using the Add-SmStorageConnection cmdlet.

   This example creates a new storage system connection:

```
PS C:\> Add-SmStorageConnection -Storage test_vs1 -Protocol Https
-Timeout 60
```

3. Create a new credential by using the Add-SmCredential cmdlet.

   This example creates a new credential named FinanceAdmin with Windows credentials:

```
PS C:> Add-SmCredential -Name FinanceAdmin -AuthMode Windows
-Credential sddev\administrator
```

The information regarding the parameters that can be used with the cmdlet and their descriptions can be obtained by running *Get-Help command_name*. Alternatively, you can also refer to the SnapCenter Software Cmdlet Reference Guide.

Back up a single resource on demand for Windows file systems

If a resource is not in a resource group, you can back up the resource on demand from the Resources page.

About this task

If you want to back up a resource that has a SnapMirror relationship with secondary storage, the role assigned to the storage user should include the “snapmirror all” privilege. However, if you are using the “vsadmin” role, then the “snapmirror all” privilege is not required.

When backing up a file system, SnapCenter does not back up LUNs that are mounted on a volume mount point (VMP) in the file system that is being backed up.

If you are working in a Windows file system context, do not back up database files. Doing so creates an inconsistent backup and a possible loss of data when restoring. To protect database files, you must use the appropriate SnapCenter plug-in for the database (for example, SnapCenter Plug-in for Microsoft SQL Server or SnapCenter Plug-in for Microsoft Exchange Server).

SnapCenter UI

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page, select the File System resource type, and then select the resource that you want to back up.
3. If the File System - Protect wizard does not start automatically, click **Protect** to start the wizard.

   Specify the protection settings, as described in the Creating resource groups tasks.

4. Optional: In the Resource page of the wizard, enter a custom name format for the Snapshot.

   For example, customtext_resourcegroup_policy_hostname or resourcegroup_hostname. By default, a timestamp is appended to the Snapshot name.

5. In the Policies page, perform the following tasks:
   a. Select one or more policies from the drop-down list.

      You can select any existing policy, and then click **Details** to determine whether you can use that policy.

      If no existing policy meets your requirements, you can copy an existing policy and modify it or you can create a new policy by clicking to start the policy wizard.
      If no existing policy meets your requirements, you can copy an existing policy and modify it or you can create a new policy by clicking to start the policy wizard.

      The selected policies are listed in the Policy column in the Configure schedules for selected policies section.

   b. In the Configure schedules for selected policies section, click in the Configure Schedules column for the policy for which you want to configure the schedule.
   c. In the Add schedules for policy *policy_name* dialog box, configure the schedule by specifying the start date, expiration date, and frequency, and then click **Finish**.

      The configured schedules are listed in the Applied Schedules column in the Configure schedules for selected policies section.

      Scheduled operations might fail

6. In the Notification page, perform the following tasks:

For this field…	Do this…
Email preference	Select **Always**, or **On Failure**, or **On failure or warning**, to send emails to recipients after creating backup resource groups, attaching policies, and configuring schedules. Enter the SMTP server information, default email subject line, and the “To” and “From” email addresses.
From	Email address
To	Email to address
Subject	Default email subject line

7. Review the summary, and then click **Finish**.

   The database topology page is displayed.

8. Click **Back up Now**.
9. In the Backup page, perform the following steps:
   a. If you have applied multiple policies to the resource, from the Policy drop-down list, select the policy that you want to use for backup.

      If the policy selected for the on-demand backup is associated with a backup schedule, the on-demand backups will be retained based on the retention settings specified for the schedule type.

   b. Click **Backup**.
10. Monitor the operation progress by clicking **Monitor** > **Jobs**.

PowerShell cmdlets

Steps

1. Initiate a connection session with the SnapCenter Server for a specified user by using the Open-SmConnection cmdlet.

   ```
   Open-smconnection  -SMSbaseurl  https://snapctr.demo.netapp.com:8146
   ```

   The username and password prompt is displayed.

2. Create a backup policy by using the Add-SmPolicy cmdlet.

   This example creates a new backup policy with a SQL backup type of FullBackup:

```
PS C:\> Add-SmPolicy -PolicyName TESTPolicy
-PluginPolicyType SCSQL -PolicyType Backup
-SqlBackupType FullBackup -Verbose
```

This example creates a new backup policy with a Windows file system backup type of CrashConsistent:

```
PS C:\> Add-SmPolicy -PolicyName FileSystemBackupPolicy
-PluginPolicyType SCW -PolicyType Backup
-ScwBackupType CrashConsistent -Verbose
```

3. Discover host resources by using the Get-SmResources cmdlet.

   This example discovers the resources for the Microsoft SQL plug-in on the specified host:

```
C:\PS>PS C:\> Get-SmResources -HostName vise-f6.sddev.mycompany.com
-PluginCode SCSQL
```

   This example discovers the resources for Windows file systems on the specified host:

```
C:\PS>PS C:\> Get-SmResources -HostName vise2-f6.sddev.mycompany.com
-PluginCode SCW
```

4. Add a new resource group to SnapCenter by using the Add-SmResourceGroup cmdlet.

   This example creates a new SQL database backup resource group with the specified policy and resources:

```
PS C:\> Add-SmResourceGroup -ResourceGroupName AccountingResource
-Resources @{"Host"="visef6.org.com";
"Type"="SQL Database";"Names"="vise-f6\PayrollDatabase"}
-Policies "BackupPolicy"
```

   This example creates a new Windows file system backup resource group with the specified policy and resources:

```
PS C:\> Add-SmResourceGroup -ResourceGroupName EngineeringResource
-PluginCode SCW -Resources @{"Host"="WIN-VOK20IKID5I";
"Type"="Windows Filesystem";"Names"="E:\"}
-Policies "EngineeringBackupPolicy"
```

5. Initiate a new backup job by using the New-SmBackup cmdlet.

```
PS C:> New-SmBackup -ResourceGroupName PayrollDataset -Policy
FinancePolicy
```

6. View the status of the backup job by using the Get-SmBackupReport cmdlet.

   This example displays a job summary report of all jobs that were run on the specified date:

```
PS C:\> Get-SmJobSummaryReport -Date '1/27/2016'
```

The information regarding the parameters that can be used with the cmdlet and their descriptions can be obtained by running *Get-Help command_name*. Alternatively, you can also refer to the SnapCenter Software Cmdlet Reference Guide.

Back up resource groups for Windows file systems

A resource group is a collection of resources on a host or cluster. A backup operation on the resource group is performed on all resources defined in the resource group. You can back up a resource group on demand from the Resources page. If a resource group has a policy attached and a schedule configured, then backups occur automatically according to the schedule.

Before you begin

- You must have created a resource group with a policy attached.
- If you want to back up a resource that has a SnapMirror relationship to secondary storage, the role assigned to the storage user should include the “snapmirror all” privilege. However, if you are using the “vsadmin” role, then the “snapmirror all” privilege is not required.
- If a resource group has multiple databases from different hosts, the backup operation on some of the hosts might trigger late because of network issues. You should configure the value of MaxRetryForUninitializedHosts in web.config by using the Set-SmConfigSettings PowerShell cmdlet

When backing up a file system, SnapCenter does not back up LUNs that are mounted on a volume mount point (VMP) in the file system that is being backed up.

If you are working in a Windows file system context, do not back up database files. Doing so creates an inconsistent backup and a possible loss of data when restoring. To protect database files, you must use the appropriate SnapCenter plug-in for the database (for example, SnapCenter Plug-in for Microsoft SQL Server or SnapCenter Plug-in for Microsoft Exchange Server).

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page, select **Resource Group** from the **View** list.

You can search the resource group either by entering the resource group name in the search box or by clicking and selecting the tag. You can then click to close the filter pane.

3. In the Resource Groups page, select the resource group that you want to back up, and then click **Back up Now**.

   For SnapCenter Plug-in for Oracle Database, if you have a federated resource group with two databases and one of the database has datafile on a non-NetApp storage, the backup operation is aborted even though the other database is on a NetApp storage.

4. In the Backup page, perform the following steps:
   a. If you have associated multiple policies with the resource group, from the **Policy** drop-down list, select the policy that you want to use for backup.

      If the policy selected for the on-demand backup is associated with a backup schedule, the on-demand backups will be retained based on the retention settings specified for the schedule type.

   b. Click **Backup**.
5. Monitor the operation progress by clicking **Monitor** > **Jobs**.
   - In MetroCluster configurations, SnapCenter might not be able to detect a protection relationship after a failover.

     Unable to detect SnapMirror or SnapVault relationship after MetroCluster failover

   - If you are backing up application data on VMDKs and the Java heap size for the SnapCenter Plug-in for VMware vSphere is not large enough, the backup might fail. To increase the Java heap size, locate the script file `/opt/netapp/init_scripts/scvservice`. In that script, the `do_start method` command starts the SnapCenter VMware plug-in service. Update that command to the following: `Java -jar -Xmx8192M -Xms4096M`.

Monitor backup operations

You can monitor the progress of different backup operations by using the SnapCenterJobs page. You might want to check the progress to determine when it is complete or if there is an issue.

About this task

The following icons appear on the Jobs page and indicate the corresponding state of the operations:

- In progress
- Completed successfully
- Failed
- Completed with warnings or could not start due to warnings
- Queued
- Canceled

Steps

1. In the left navigation pane, click **Monitor**.
2. In the Monitor page, click **Jobs**.
3. In the Jobs page, perform the following steps:
   a. Click to filter the list so that only backup operations are listed.
   b. Specify the start and end dates.
   c. From the **Type** drop-down list, select **Backup**.
   d. From the **Status** drop-down, select the backup status.
   e. Click **Apply** to view the operations completed successfully.
4. Select a backup job, and then click **Details** to view the job details.

Though the backup job status displays , when you click on job details you might see that some of the child tasks of the backup operation are still in progress or marked with warning signs.

5. In the Job Details page, click **View logs**.

   The **View logs** button displays the detailed logs for the selected operation.

Monitor operations in the Activity pane

The Activity pane displays the five most recent operations performed. The Activity pane also displays when the operation was initiated and the status of the operation.

The Activity pane displays information regarding backup, restore, clone, and scheduled backup operations.

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. Click on the Activity pane to view the five most recent operations.

   When you click one of the operations, the operation details are listed in the **Job Details** page.

Cancel backup operations

You can cancel backup operations that are queued.

What you will need

- You must be logged in as the SnapCenter Admin or job owner to cancel operations.
- You can cancel a backup operation from either the **Monitor** page or the **Activity** pane.
- You cannot cancel a running backup operation.
- You can use the SnapCenter GUI, PowerShell cmdlets, or CLI commands to cancel the backup operations.
- The **Cancel Job** button is disabled for operations that cannot be canceled.
- If you selected **All members of this role can see and operate on other members objects** in Users\Groups page while creating a role, you can cancel the queued backup operations of other members while using that role.

Steps

1. Perform one of the following actions:

From the…	Action
Monitor page	a. In the left navigation pane, click **Monitor** > **Jobs**. b. Select the operation, and then click **Cancel Job**.
Activity pane	a. After initiating the backup operation, click on the Activity pane to view the five most recent operations. b. Select the operation. c. In the Job Details page, click **Cancel Job**.

The operation is canceled, and the resource is reverted to the previous state.

View related backups and clones in the Topology page

When you are preparing to back up or clone a resource, you can view a graphical representation of all backups and clones on the primary and secondary storage. In the Topology page, you can see all of the backups and clones that are available for the selected resource or resource group. You can view the details of those backups and clones, and then select them to perform data protection operations.

About this task

You can review the following icons in the Manage Copies view to determine whether the backups and clones are available on the primary or secondary storage (Mirror copies or Vault copies).

- displays the number of backups and clones that are available on the primary storage.
- displays the number of backups and clones that are mirrored on the secondary storage using SnapMirror technology.

  Clones of a backup of a version-flexible mirror on a mirror-vault type volume are displayed in the topology view but the mirror backup count in the topology view does not include the version-flexible backup.

- displays the number of backups and clones that are replicated on the secondary storage using SnapVault technology.
  - The number of backups displayed includes the backups deleted from the secondary storage. For

example, if you have created 6 backups using a policy to retain only 4 backups, the number of backups displayed are 6.

Clones of a backup of a version-flexible mirror on a mirror-vault type volume are displayed in the topology view but the mirror backup count in the topology view does not include the version-flexible backup.

If you have secondary relationship as SnapMirror active sync (initially released as SnapMirror Business Continuity [SM-BC]), you can see following additional icons:

- The replica site is up.
- The replica site is down.
- The secondary mirror or vault relationship has not been re-established.

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page, either select the resource or resource group from the **View** drop-down list.
3. Select the resource either from the resource details view or from the resource group details view.

   If the resource is protected, the topology page of the selected resource is displayed.

4. Review the Summary card to see a summary of the number of backups and clones available on the primary and secondary storage.

   The Summary Card section displays the total number of backups and clones. For Oracle database only, the Summary Card section also displays the total number of log backups.

   Clicking the **Refresh** button starts a query of the storage to display an accurate count.

   If SnapLock enabled backup is taken, then clicking the **Refresh** button refreshes the primary and secondary SnapLock expiry time retrieved from ONTAP. A weekly schedule also refreshes the primary and secondary SnapLock expiry time retrieved from ONTAP.

   When the application resource is spread across multiple volumes, the SnapLock expiry time for the backup will be the longest SnapLock expiry time that is set for a Snapshot in a volume. The longest SnapLock expiry time is retrieved from ONTAP.

   For SnapMirror active sync, clicking the **Refresh** button refreshes the SnapCenter backup inventory by querying ONTAP for both primary and replica sites. A weekly schedule also performs this activity for all databases containing SnapMirror active sync relationship.

   - For SnapMirror active sync and only for ONTAP 9.14.1, Async Mirror or Async MirrorVault relationships to the new primary destination should be manually configured after failover. From ONTAP 9.15.1 onwards Async Mirror or Async MirrorVault is auto configured to the new primary destination.
   - After failover, a backup should be created for SnapCenter to be aware of the failover. You can click **Refresh** only after a backup has been created.

5. In the Manage Copies view, click **Backups** or **Clones** from the primary or secondary storage to see details of a backup or clone.

The details of the backups and clones are displayed in a table format.

6. Select the backup from the table, and then click the data protection icons to perform restore, clone, rename, and delete operations.

   You cannot rename or delete backups that are on the secondary storage system.

7. If you want to delete a clone, then select the clone from the table and click to delete the clone.

Example showing backups and clones on the primary storage

Clean up the secondary backup count using PowerShell cmdlets

You can use the Remove-SmBackup cmdlet to clean up the backup count for secondary backups that have no Snapshot. You might want to use this cmdlet when the total Snapshots displayed in the Manage Copies topology do not match the secondary storage Snapshot retention setting.

You must have prepared the PowerShell environment to execute the PowerShell cmdlets.

The information regarding the parameters that can be used with the cmdlet and their descriptions can be obtained by running *Get-Help command_name*. Alternatively, you can also refer to the SnapCenter Software Cmdlet Reference Guide.

Steps

1. Initiate a connection session with the SnapCenter Server for a specified user by using the Open-SmConnection cmdlet.

   ```
   Open-SmConnection  -SMSbaseurl  https:\\snapctr.demo.netapp.com:8146/
   ```

2. Clean up secondary backups count using the -CleanupSecondaryBackups parameter.

   This example cleans up the backup count for secondary backups with no Snapshots:

```
Remove-SmBackup -CleanupSecondaryBackups
Remove-SmBackup
Are you sure want to remove the backup(s).
[Y] Yes  [A] Yes to All  [N] No  [L] No to All  [S] Suspend  [?] Help
(default is "Y"):
```

Restore Windows file systems

Restore Windows file system backups

You can use SnapCenter to restore file system backups. File system restoration is a multiphase process that copies all the data from a specified backup to the original location of the file system.

Before you begin

- You must have backed up the file system.
- If a scheduled operation, such as a backup operation, is currently in progress for a file system, then that operation must be cancelled before you can start a restore operation.
- You can only restore a file system backup to the original location, not to an alternate path.

  You cannot restore a single file from a backup because the restored file system overwrites any data on the original location of the file system. To restore a single file from a file system backup, you must clone the backup and access the file in the clone.

- You cannot restore a system or boot volume.
- SnapCenter can restore file systems in a Windows cluster without taking the cluster group offline.

About this task

- The SCRIPTS_PATH is defined using the PredefinedWindowsScriptsDirectory key located in the SMCoreServiceHost.exe.Config file of the plug-in host.

  If needed, you can change this path and restart SMcore service. It is recommended that you use the default path for security.

  The value of the key can be displayed from swagger through the API: API /4.7/configsettings

  You can use the GET API to display the value of the key. SET API is not supported.

- For SnapMirror active sync restore operation, you must select the backup from the primary location.
- For ONTAP 9.12.1 and below version, the clones created from the SnapLock Vault Snapshots as part of restore will inherit the SnapLock Vault expiry time. Storage admin should manually cleanup the clones post the SnapLock expiry time.

SnapCenter UI

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. To filter the list of resources, select the File System and Resource Group options.
3. Select a resource group from the list, and then click **Restore**.
4. In the Backups page, select whether you want to restore from primary or secondary storage systems, and then select a backup to restore.
5. Select your options in the Restore wizard.
6. You can enter the path and the arguments of the prescript or postscript that you want SnapCenter to run before or after the restore operation, respectively.

   For example, you can run a script to update SNMP traps, automate alerts, send logs, and so on.

   > The prescripts or postscripts path should not include drives or shares. The path should be relative to the SCRIPTS_PATH.

7. In the Notification page, select one of the following options:

For this field…	Do this…
Log SnapCenter server events to storage system syslog	Select this option to log SnapCenter Server events to the syslog of the storage system.
Send AutoSupport notification for failed operations to storage system	Select this option to send information about any failed operations to NetApp using AutoSupport.
Email preference	Select **Always**, **On Failure**, or **On failure or warning** to send email messages to recipients after restoring backups. Enter the SMTP server, default email subject line, and To and From email addresses.

8. Review the summary, and then click **Finish**.
9. Monitor the operation progress by clicking **Monitor** > **Jobs**.

> If the restored file system contains a database, then you must also restore the database. If you do not restore the database, then your database might be in an invalid state. For information on restoring databases, see the Data Protection Guide for that database.

PowerShell cmdlets

Steps

1. Initiate a connection session with the SnapCenter Server for a specified user by using the Open-SmConnection cmdlet.

```
PS C:\> Open-Smconnection
```

2. Retrieve the information about the one or more backups that you want to restore by using the Get-SmBackup and Get-SmBackupReport cmdlets.

   This example displays information about all available backups:

```
PS C:\> Get-SmBackup

BackupId                      BackupName
BackupTime                    BackupType
--------                      ----------
----------                    ----------
  1               Payroll Dataset_vise-f6_08... 8/4/2015    11:02:32
AM                 Full Backup
  2               Payroll Dataset_vise-f6_08... 8/4/2015    11:23:17
AM
```

   This example displays detailed information about the backup from January 29th 2015 to February 3rd, 2015:

```
PS C:\> Get-SmBackupReport -FromDateTime "1/29/2015" -ToDateTime
"2/3/2015"

   SmBackupId           : 113
   SmJobId              : 2032
   StartDateTime        : 2/2/2015 6:57:03 AM
   EndDateTime          : 2/2/2015 6:57:11 AM
   Duration             : 00:00:07.3060000
   CreatedDateTime      : 2/2/2015 6:57:23 AM
   Status               : Completed
   ProtectionGroupName  : Clone
   SmProtectionGroupId  : 34
   PolicyName           : Vault
   SmPolicyId           : 18
   BackupName           : Clone_SCSPR0019366001_02-02-2015_06.57.08
   VerificationStatus   : NotVerified

   SmBackupId           : 114
   SmJobId              : 2183
   StartDateTime        : 2/2/2015 1:02:41 PM
   EndDateTime          : 2/2/2015 1:02:38 PM
   Duration             : -00:00:03.2300000
   CreatedDateTime      : 2/2/2015 1:02:53 PM
   Status               : Completed
   ProtectionGroupName  : Clone
   SmProtectionGroupId  : 34
   PolicyName           : Vault
   SmPolicyId           : 18
   BackupName           : Clone_SCSPR0019366001_02-02-2015_13.02.45
   VerificationStatus   : NotVerified
```

3. Restore data from the backup by using the Restore-SmBackup cmdlet.

```
Restore-SmBackup -PluginCode 'DummyPlugin' -AppObjectId
'scc54.sccore.test.com\DummyPlugin\NTP\DB1' -BackupId 269
-Confirm:$false
output:
Name                : Restore
'scc54.sccore.test.com\DummyPlugin\NTP\DB1'
Id                  : 2368
StartTime           : 10/4/2016 11:22:02 PM
EndTime             :
IsCancellable       : False
IsRestartable       : False
IsCompleted         : False
IsVisible           : True
IsScheduled         : False
PercentageCompleted : 0
Description         :
Status              : Queued
Owner               :
Error               :
Priority            : None
Tasks               : {}
ParentJobID         : 0
EventId             : 0
JobTypeId           :
ApisJobKey          :
ObjectId            : 0
PluginCode          : NONE
PluginName          :
```

The information regarding the parameters that can be used with the cmdlet and their descriptions can be obtained by running *Get-Help command_name*. Alternatively, you can also refer to the SnapCenter Software Cmdlet Reference Guide.

Restore resources using PowerShell cmdlets

Restoring a resource backup includes initiating a connection session with the SnapCenter Server, listing the backups and retrieving backup information, and restoring a backup.

You must have prepared the PowerShell environment to execute the PowerShell cmdlets.

Steps

1. Initiate a connection session with the SnapCenter Server for a specified user by using the Open-SmConnection cmdlet.

```
PS C:\> Open-Smconnection
```

2. Retrieve the information about the one or more backups that you want to restore by using the Get-SmBackup and Get-SmBackupReport cmdlets.

   This example displays information about all available backups:

```
PS C:\> Get-SmBackup

BackupId                      BackupName                    BackupTime
BackupType
--------                      ----------                    ----------
----------
  1               Payroll Dataset_vise-f6_08... 8/4/2015    11:02:32 AM
Full Backup
  2               Payroll Dataset_vise-f6_08... 8/4/2015    11:23:17 AM
```

   This example displays detailed information about the backup from January 29th 2015 to February 3rd, 2015:

```
PS C:\> Get-SmBackupReport -FromDateTime "1/29/2015" -ToDateTime
"2/3/2015"

   SmBackupId           : 113
   SmJobId              : 2032
   StartDateTime        : 2/2/2015 6:57:03 AM
   EndDateTime          : 2/2/2015 6:57:11 AM
   Duration             : 00:00:07.3060000
   CreatedDateTime      : 2/2/2015 6:57:23 AM
   Status               : Completed
   ProtectionGroupName  : Clone
   SmProtectionGroupId  : 34
   PolicyName           : Vault
   SmPolicyId           : 18
   BackupName           : Clone_SCSPR0019366001_02-02-2015_06.57.08
   VerificationStatus   : NotVerified

   SmBackupId           : 114
   SmJobId              : 2183
   StartDateTime        : 2/2/2015 1:02:41 PM
   EndDateTime          : 2/2/2015 1:02:38 PM
   Duration             : -00:00:03.2300000
   CreatedDateTime      : 2/2/2015 1:02:53 PM
   Status               : Completed
   ProtectionGroupName  : Clone
   SmProtectionGroupId  : 34
   PolicyName           : Vault
   SmPolicyId           : 18
   BackupName           : Clone_SCSPR0019366001_02-02-2015_13.02.45
   VerificationStatus   : NotVerified
```

3. Restore data from the backup by using the Restore-SmBackup cmdlet.

```
Restore-SmBackup -PluginCode 'DummyPlugin' -AppObjectId
'scc54.sccore.test.com\DummyPlugin\NTP\DB1' -BackupId 269
-Confirm:$false
output:
Name                : Restore
'scc54.sccore.test.com\DummyPlugin\NTP\DB1'
Id                  : 2368
StartTime           : 10/4/2016 11:22:02 PM
EndTime             :
IsCancellable       : False
IsRestartable       : False
IsCompleted         : False
IsVisible           : True
IsScheduled         : False
PercentageCompleted : 0
Description         :
Status              : Queued
Owner               :
Error               :
Priority            : None
Tasks               : {}
ParentJobID         : 0
EventId             : 0
JobTypeId           :
ApisJobKey          :
ObjectId            : 0
PluginCode          : NONE
PluginName          :
```

The information regarding the parameters that can be used with the cmdlet and their descriptions can be obtained by running *Get-Help command_name*. Alternatively, you can also refer to the SnapCenter Software Cmdlet Reference Guide.

Monitor restore operations

You can monitor the progress of different SnapCenter restore operations by using the Jobs page. You might want to check the progress of an operation to determine when it is complete or if there is an issue.

About this task

Post-restore states describe the conditions of the resource after a restore operation and any further restore actions that you can take.

The following icons appear on the Jobs page, and indicate the state of the operation:

- In progress

- Completed successfully
- Failed
- Completed with warnings or could not start due to warnings
- Queued
- Canceled

Steps

1. In the left navigation pane, click **Monitor**.
2. In the **Monitor** page, click **Jobs**.
3. In the **Jobs** page, perform the following steps:
   a. Click to filter the list so that only restore operations are listed.
   b. Specify the start and end dates.
   c. From the **Type** drop-down list, select **Restore**.
   d. From the **Status** drop-down list, select the restore status.
   e. Click **Apply** to view the operations that have been completed successfully.
4. Select the restore job, and then click **Details** to view the job details.
5. In the **Job Details** page, click **View logs**.

   The **View logs** button displays the detailed logs for the selected operation.

Cancel restore operations

You can cancel restore jobs that are queued.

You should be logged in as the SnapCenter Admin or job owner to cancel restore operations.

About this task

- You can cancel a queued restore operation from either the **Monitor** page or the **Activity** pane.
- You cannot cancel a running restore operation.
- You can use the SnapCenter GUI, PowerShell cmdlets, or CLI commands to cancel the queued restore operations.
- The **Cancel Job** button is disabled for restore operations that cannot be canceled.
- If you selected **All members of this role can see and operate on other members objects** in Users\Groups page while creating a role, you can cancel the queued restore operations of other members while using that role.

Step

Perform one of the following actions:

From the…	Action
Monitor page	a. In the left navigation pane, click **Monitor** > **Jobs**. b. Select the job and click **Cancel Job**.

From the…	Action
Activity pane	a. After initiating the restore operation, click on the Activity pane to view the five most recent operations. b. Select the operation. c. In the Job Details page, click **Cancel Job**.

Clone Windows file systems

Clone from a Windows file system backup

You can use SnapCenter to clone a Windows file system backup. If you want a copy of a single file that was mistakenly deleted or changed, then you can clone a backup and access that file in the clone.

Before you begin

- You should have prepared for data protection by completing tasks such as adding hosts, identifying resources, and creating storage virtual machine (SVM) connections.
- You should have a backup of the file system.
- You should ensure that the aggregates hosting the volumes should be in the assigned aggregates list of the storage virtual machine (SVM).
- You cannot clone a resource group. You can only clone individual file system backups.
- If a backup resides on a virtual machine with a VMDK disk, SnapCenter cannot clone the backup to a physical server.
- If you clone a Windows cluster (for example, a shared LUN or a cluster shared volume (CSV) LUN), the clone is stored as a dedicated LUN on the host that you specify.
- For a cloning operation, the root directory of the volume mount point cannot be a shared directory.
- You cannot create a clone on a node that is not the home node for the aggregate.
- You cannot schedule recurring clone (clone lifecycle) operations for Windows file systems; you can only clone a backup on demand.
- If you move a LUN that contains a clone to a new volume, SnapCenter can no longer support the clone. For example, you cannot use SnapCenter to delete that clone.
- You cannot clone across environments. For example, cloning from a physical disk to a virtual disk or vice versa.

About this task

- The SCRIPTS_PATH is defined using the PredefinedWindowsScriptsDirectory key located in the SMCoreServiceHost.exe.Config file of the plug-in host.

  If needed, you can change this path and restart SMcore service. It is recommended that you use the default path for security.

  The value of the key can be displayed from swagger through the API: API /4.7/configsettings

  You can use the GET API to display the value of the key. SET API is not supported.

- For ONTAP 9.12.1 and below version, the clones created from the SnapLock Vault Snapshots as part of restore will inherit the SnapLock Vault expiry time. Storage admin should manually cleanup the clones post the SnapLock expiry time.

SnapCenter UI

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page, select **File Systems** from the list.
3. Select the host.

   The topology view is automatically displayed if the resource is protected.

4. From the resources list, select the backup that you want to clone, and then click the clone icon.
5. In the Options page, do the following:

For this field…	Do this…
Clone server	Choose the host on which the clone should be created.
“Auto assign mount point” or “Auto assign volume mount point under path”	Choose whether to automatically assign a mount point or a volume mount point under a path. Auto assign volume mount point under path: The mount point under a path enables you to provide a specific directory in which the mount points will be created. Before you choose this option, you must verify that the directory is empty. If there is a backup in the directory, the backup will be in an invalid state after the mount operation.
Archive location	Choose an archive location if you are cloning a secondary backup.

6. In the Script page, specify any prescripts or postscripts you want to execute.

   The prescripts or postscripts path should not include drives or shares. The path should be relative to the SCRIPTS_PATH.

7. Review the summary, and then click **Finish**.
8. Monitor the operation progress by clicking **Monitor** > **Jobs**.

PowerShell cmdlets

Steps

1. Initiate a connection session with the SnapCenter Server for a specified user by using the Open-SmConnection cmdlet.

```
Open-SmConnection  -SMSbaseurl  https://snapctr.demo.netapp.com:8146
```

2. List the backups that can be cloned by using the Get-SmBackup or Get-SmResourceGroup cmdlet.

This example displays information about all available backups:

```
C:\PS>PS C:\> Get-SmBackup

BackupId   BackupName                      BackupTime   BackupType
--------   ----------                      ----------   ----------
1          Payroll Dataset_vise-f6_08...  8/4/2015     Full Backup
                                          11:02:32 AM

2          Payroll Dataset_vise-f6_08...  8/4/2015
                                          11:23:17 AM
```

This example displays information about a specified resource group, its resources, and associated policies:

```
PS C:\> Get-SmResourceGroup -ListResources –ListPolicies

Description :
CreationTime : 8/4/2015 3:44:05 PM
ModificationTime : 8/4/2015 3:44:05 PM
EnableEmail : False
EmailSMTPServer :
EmailFrom :
EmailTo :
EmailSubject :
EnableSysLog : False
ProtectionGroupType : Backup
EnableAsupOnFailure : False
Policies : {FinancePolicy}
HostResourceMaping : {}
Configuration : SMCoreContracts.SmCloneConfiguration
LastBackupStatus :
VerificationServer :
EmailBody :
EmailNotificationPreference : Never
VerificationServerInfo : SMCoreContracts.SmVerificationServerInfo
SchedulerSQLInstance :
CustomText :
CustomSnapshotFormat :
SearchResources : False
ByPassCredential : False
IsCustomSnapshot :
MaintenanceStatus : Production
PluginProtectionGroupTypes : {SMSQL}
Name : Payrolldataset
```

```
Type : Group
Id : 1
Host :
UserName :
Passphrase :
Deleted : False
Auth : SMCoreContracts.SmAuth
IsClone : False
CloneLevel : 0
ApplySnapvaultUpdate : False
ApplyRetention : False
RetentionCount : 0
RetentionDays : 0
ApplySnapMirrorUpdate : False
SnapVaultLabel :
MirrorVaultUpdateRetryCount : 7
AppPolicies : {}
Description : FinancePolicy
PreScriptPath :
PreScriptArguments :
PostScriptPath :
PostScriptArguments :
ScriptTimeOut : 60000
DateModified : 8/4/2015 3:43:30 PM
DateCreated : 8/4/2015 3:43:30 PM
Schedule : SMCoreContracts.SmSchedule
PolicyType : Backup
PluginPolicyType : SMSQL
Name : FinancePolicy
Type :
Id : 1
Host :
UserName :
Passphrase :
Deleted : False
Auth : SMCoreContracts.SmAuth
IsClone : False
CloneLevel : 0
clab-a13-13.sddev.lab.netapp.com
DatabaseGUID :
SQLInstance : clab-a13-13
DbStatus : AutoClosed
DbAccess : eUndefined
IsSystemDb : False
IsSimpleRecoveryMode : False
IsSelectable : True
```

```
SqlDbFileGroups : {}
SqlDbLogFiles : {}
AppFileStorageGroups : {}
LogDirectory :
AgName :
Version :
VolumeGroupIndex : -1
IsSecondary : False
Name : TEST
Type : SQL Database
Id : clab-a13-13\TEST
Host : clab-a13-13.sddev.mycompany.com
UserName :
Passphrase :
Deleted : False
Auth : SMCoreContracts.SmAuth
IsClone : False
```

3. Initiate a clone operation from an existing backup by using the New-SmClone cmdlet.

   This example creates a clone from a specified backup with all logs:

```
PS C:\> New-SmClone
-BackupName payroll_dataset_vise-f3_08-05-2015_15.28.28.9774
-Resources @{"Host"="vise-f3.sddev.mycompany.com";
"Type"="SQL Database";"Names"="vise-f3\SQLExpress\payroll"}
-CloneToInstance vise-f3\sqlexpress -AutoAssignMountPoint
-Suffix _clonefrombackup
-LogRestoreType All -Policy clonefromprimary_ondemand

PS C:> New-SmBackup -ResourceGroupName PayrollDataset -Policy
FinancePolicy
```

   This example creates a clone to a specified Microsoft SQL Server instance:

```
PS C:\> New-SmClone
-BackupName "BackupDS1_NY-VM-SC-SQL_12-08-2015_09.00.24.8367"
-Resources @{"host"="ny-vm-sc-sql";"Type"="SQL Database";
"Names"="ny-vm-sc-sql\AdventureWorks2012_data"}
-AppPluginCode SMSQL -CloneToInstance "ny-vm-sc-sql"
-Suffix _CLPOSH -AssignMountPointUnderPath "C:\SCMounts"
```

4. View the status of the clone job by using the Get-SmCloneReport cmdlet.

   This example displays a clone report for the specified job ID:

```
PS C:\> Get-SmCloneReport -JobId 186

SmCloneId : 1
SmJobId : 186
StartDateTime : 8/3/2015 2:43:02 PM
EndDateTime : 8/3/2015 2:44:08 PM
Duration : 00:01:06.6760000
Status : Completed
ProtectionGroupName : Draper
SmProtectionGroupId : 4
PolicyName : OnDemand_Clone
SmPolicyId : 4
BackupPolicyName : OnDemand_Full_Log
SmBackupPolicyId : 1
CloneHostName : SCSPR0054212005.mycompany.com
CloneHostId : 4
CloneName : Draper__clone__08-03-2015_14.43.53
SourceResources : {Don, Betty, Bobby, Sally}
ClonedResources : {Don_DRAPER, Betty_DRAPER, Bobby_DRAPER,
                   Sally_DRAPER}
```

The information regarding the parameters that can be used with the cmdlet and their descriptions can be obtained by running *Get-Help command_name*. Alternatively, you can also refer to the SnapCenter Software Cmdlet Reference Guide.

Monitor clone operations

You can monitor the progress of SnapCenter clone operations by using the Jobs page. You might want to check the progress of an operation to determine when it is complete or if there is an issue.

About this task

The following icons appear on the Jobs page, and indicate the state of the operation:

- In progress
- Completed successfully
- Failed
- Completed with warnings or could not start due to warnings
- Queued
- Canceled

Steps

1. In the left navigation pane, click **Monitor**.
2. In the **Monitor** page, click **Jobs**.
3. In the **Jobs** page, perform the following steps:
   a. Click to filter the list so that only clone operations are listed.
   b. Specify the start and end dates.
   c. From the **Type** drop-down list, select **Clone**.
   d. From the **Status** drop-down list, select the clone status.
   e. Click **Apply** to view the operations that are completed successfully.
4. Select the clone job, and then click **Details** to view the job details.
5. In the Job Details page, click **View logs**.

Cancel clone operations

You can cancel clone operations that are queued.

You should be logged in as the SnapCenter Admin or job owner to cancel clone operations.

About this task

- You can cancel a queued clone operation from either the **Monitor** page or the **Activity** pane.
- You cannot cancel a running clone operation.
- You can use the SnapCenter GUI, PowerShell cmdlets, or CLI commands to cancel the queued clone operations.
- If you selected **All members of this role can see and operate on other members objects** in Users\Groups page while creating a role, you can cancel the queued clone operations of other members while using that role.

Step

Perform one of the following actions:

From the…	Action
Monitor page	a. In the left navigation pane, click **Monitor** > **Jobs**. b. Select the operation, and click **Cancel Job**.
Activity pane	a. After initiating the clone operation, click on the Activity pane to view the five most recent operations. b. Select the operation. c. In the **Job Details** page, click **Cancel Job**.

Split a clone

You can use SnapCenter to split a cloned resource from the parent resource. The clone that is split becomes independent of the parent resource.

About this task

- You cannot perform the clone split operation on an intermediate clone.

  For example, after you create clone1 from a database backup, you can create a backup of clone1, and then clone this backup (clone2). After you create clone2, clone1 is an intermediate clone, and you cannot perform the clone split operation on clone1. However, you can perform the clone split operation on clone2.

  After splitting clone2, you can perform the clone split operation on clone1 because clone1 is no longer the intermediate clone.

- When you split a clone, the backup copies and clone jobs of the clone are deleted.
- For information about FlexClone volume split operations, see, Split a FlexClone volume from its parent volume.
- Ensure that the volume or aggregate on the storage system is online.

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. In the **Resources** page, select the appropriate option from the View list:

Option	Description
For database applications	Select **Database** from the View list.
For file systems	Select **Path** from the View list.

3. Select the appropriate resource from the list.

   The resource topology page is displayed.

4. From the **Manage Copies** view, select the cloned resource (for example, the database or LUN), and then click .
5. Review the estimated size of the clone that is to be split and the required space available on the aggregate, and then click **Start**.
6. Monitor the operation progress by clicking **Monitor** > **Jobs**.

   The clone split operation stops responding if the SMCore service restarts. You should run the Stop-SmJob cmdlet to stop the clone split operation, and then retry the clone split operation.

   If you want a longer poll time or shorter poll time to check whether the clone is split or not, you can change the value of *CloneSplitStatusCheckPollTime* parameter in *SMCoreServiceHost.exe.config* file to set the time interval for SMCore to poll for the status of the clone split operation. The value is in milliseconds and the default value is 5 minutes.

   For example:

   ```
   <add key="CloneSplitStatusCheckPollTime" value="300000" />
   ```

   The clone split start operation fails if backup, restore, or another clone split is in progress. You should restart the clone split operation only after the running operations are complete.

Related information

SnapCenter clone or verification fails with aggregate does not exist

Protect Microsoft Exchange Server databases

SnapCenter Plug-in for Microsoft Exchange Server concepts

SnapCenter Plug-in for Microsoft Exchange Server overview

The SnapCenter Plug-in for Microsoft Exchange Server is a host-side component of the NetApp SnapCenter Software that enables application-aware data protection management of Exchange databases. The Plug-in for Exchange automates the backup and restore of Exchange databases in your SnapCenter environment.

When the Plug-in for Exchange is installed, you can use SnapCenter with NetApp SnapMirror technology to create mirror copies of backup sets on another volume and with NetApp SnapVault technology to perform disk-to-disk backup replication for standards compliance or archival purposes.

If you want to restore and recover mails or mailbox instead of the complete Exchange Database, you can use the Single Mailbox Recovery (SMBR) software.
NetApp® Single Mailbox Recovery has come to the end of availability (EOA) on May 12, 2023. NetApp will continue to support customers that have purchased mailbox capacity, maintenance, and support through marketing part numbers introduced on June 24, 2020, for the duration of the support entitlement.

NetApp Single Mailbox Recovery is a partner product provided by Ontrack. Ontrack PowerControls offers capabilities that are similar to those of NetApp Single Mailbox Recovery. Customers can procure new Ontrack PowerControls software licenses and Ontrack PowerControls maintenance and support renewals from Ontrack (through licensingteam@ontrack.com) for granular mailbox recovery.

The Plug-in for Exchange supports SnapMirror active sync (initially released as SnapMirror Business Continuity [SM-BC]) that enables business services to continue operating even through a complete site failure, supporting applications to fail over transparently using a secondary copy. Neither manual intervention nor additional scripting is required to trigger a failover with SnapMirror active sync.

It supports the asymmetric, failover, or non-duplex mode of SnapMirror Active Sync. This refers to the solution where the optimized path is only from the primary side LUN owning node. Any I/O coming on the secondary cluster paths is served by proxying it over to the primary cluster. Synchronous replication is uni-directional, in the direction of primary to secondary.

- Automates application-aware backup and restore operations for Microsoft Exchange Server databases and Database Availability Groups (DAGs) in your SnapCenter environment
- Supports virtualized Exchange Servers on RDM LUNs when you deploy the SnapCenter Plug-in for VMware vSphere and register the plug-in with SnapCenter.

What you can do with SnapCenter Plug-in for Microsoft Exchange Server

You can use the Plug-in for Exchange to back up and restore Exchange Server databases.

- View and manage an active inventory of Exchange Database Availability Groups (DAGs), databases, and replica sets
- Define policies that provide the protection settings for backup automation

- Assign policies to resource groups
- Protect individual DAGs and databases
- Back up primary and secondary Exchange mailbox databases
- Restore databases from primary and secondary backups

Storage types supported by SnapCenter Plug-in for Microsoft Windows and for Microsoft Exchange Server

SnapCenter supports a wide range of storage types on both physical machines and virtual machines. You must verify whether support is available for your storage type before installing the package for your host.

SnapCenter provisioning and data protection support is available on Windows Server. For the latest information about supported versions, see the https://imt.netapp.com/matrix/imt.jsp?components=121031;&solution=1259&isHWU&src=IMT [NetApp Interoperability Matrix Tool^].

Machine	Storage type	Provision using	Support notes
Physical server	FC-connected LUNs	SnapCenter graphical user interface (GUI) or PowerShell cmdlets	
Physical server	iSCSI-connected LUNs	SnapCenter GUI or PowerShell cmdlets	
VMware VM	RDM LUNs connected by an FC or iSCSI HBA	PowerShell cmdlets	Physical compatibility only VMDKs are not supported.
VMware VM	iSCSI LUNs connected directly to the guest system by the iSCSI initiator	SnapCenter GUI or PowerShell cmdlets	VMDKs are not supported.

Machine	Storage type	Provision using	Support notes
Hyper-V VM	Virtual FC (vFC) LUNs connected by a virtual Fibre Channel Switch	SnapCenter GUI or PowerShell cmdlets	You must use Hyper-V Manager to provision Virtual FC (vFC) LUNs connected by a virtual Fibre Channel Switch. Hyper-V pass through disks and backing up databases on VHD(x) that are provisioned on NetApp storage are not supported.
Hyper-V VM	iSCSI LUNs connected directly to the guest system by the iSCSI initiator	SnapCenter GUI or PowerShell cmdlets	Hyper-V pass through disks and backing up databases on VHD(x) that are provisioned on NetApp storage are not supported.

Minimum ONTAP privileges required for Exchange plug-in

The minimum ONTAP privileges that are required vary according to the SnapCenter plug-ins you are using for data protection.

- All-access commands: Minimum privileges required for ONTAP 9.12.1 and later
  - event generate-autosupport-log
  - job history show
  - job stop
  - lun
  - lun create
  - lun create

- lun create
- lun delete
- lun igroup add
- lun igroup create
- lun igroup delete
- lun igroup rename
- lun igroup rename
- lun igroup show
- lun mapping add-reporting-nodes
- lun mapping create
- lun mapping delete
- lun mapping remove-reporting-nodes
- lun mapping show
- lun modify
- lun move-in-volume
- lun offline
- lun online
- lun persistent-reservation clear
- lun resize
- lun serial
- lun show
- snapmirror policy add-rule
- snapmirror policy modify-rule
- snapmirror policy remove-rule
- snapmirror policy show
- snapmirror restore
- snapmirror show
- snapmirror show-history
- snapmirror update
- snapmirror update-ls-set
- snapmirror list-destinations
- version
- volume clone create
- volume clone show
- volume clone split start
- volume clone split stop
- volume create

    - volume destroy
    - volume file clone create
    - volume file show-disk-usage
    - volume offline
    - volume online
    - volume modify
    - volume qtree create
    - volume qtree delete
    - volume qtree modify
    - volume qtree show
    - volume restrict
    - volume show
    - volume snapshot create
    - volume snapshot delete
    - volume snapshot modify
    - volume snapshot modify-snaplock-expiry-time
    - volume snapshot rename
    - volume snapshot restore
    - volume snapshot restore-file
    - volume snapshot show
    - volume unmount
    - vserver cifs
    - vserver cifs share create
    - vserver cifs share delete
    - vserver cifs shadowcopy show
    - vserver cifs share show
    - vserver cifs show
    - vserver export-policy
    - vserver export-policy create
    - vserver export-policy delete
    - vserver export-policy rule create
    - vserver export-policy rule show
    - vserver export-policy show
    - vserver iscsi
    - vserver iscsi connection show
    - vserver show
- Read-only commands: Minimum privileges required for ONTAP 8.3.0 and later

- network interface
- network interface show
- vserver

Prepare storage systems for SnapMirror and SnapVault replication

You can use a SnapCenter plug-in with ONTAP SnapMirror technology to create mirror copies of backup sets on another volume, and with ONTAP SnapVault technology to perform disk-to-disk backup replication for standards compliance and other governance-related purposes. Before you perform these tasks, you must configure a data-protection relationship between the source and destination volumes and initialize the relationship.

SnapCenter performs the updates to SnapMirror and SnapVault after it completes the Snapshot operation. SnapMirror and SnapVault updates are performed as part of the SnapCenter job; do not create a separate ONTAP schedule.

If you are coming to SnapCenter from a NetApp SnapManager product and are satisfied with the data protection relationships you have configured, you can skip this section.

A data protection relationship replicates data on primary storage (the source volume) to secondary storage (the destination volume). When you initialize the relationship, ONTAP transfers the data blocks referenced on the source volume to the destination volume.

SnapCenter does not support cascade relationships between SnapMirror and SnapVault volumes (**Primary** > **Mirror** > **Vault**). You should use fanout relationships.

SnapCenter supports the management of version-flexible SnapMirror relationships. For details about version-flexible SnapMirror relationships and how to set them up, see the ONTAP documentation.

Define a backup strategy for Exchange Server resources

Defining a backup strategy before you create your backup jobs helps ensure that you have the backups that you require to successfully restore your databases. Your Service Level Agreement (SLA), Recovery Time Objective (RTO), and Recovery Point Objective (RPO) largely determine your backup strategy.

An SLA defines the level of service expected and addresses many service-related issues, including the availability and performance of service. The RTO is the time by when a business process must be restored after a disruption in service. An RPO defines the strategy for the age of the files that must be recovered from backup storage for regular operations to resume after a failure. The SLA, RTO, and RPO contribute to the backup strategy.

Types of backups supported for Exchange database

Backing up Exchange mailboxes using SnapCenter requires that you choose the resource type, such as databases and Database Availability Groups (DAG). Snapshot technology is leveraged to create online, read-only copies of the volumes on which the resources reside.

Backup type	Description
Full and log backup	Backs up the databases and all transaction logs, including the truncated logs. After a full backup is complete, the Exchange Server truncates the transaction logs that are already committed to the database. Typically, you should choose this option. However, if your backup time is short, you can choose not to run a transaction log backup with full backup.
Full backup	Backs up databases and transaction logs. The truncated transaction logs are not backed up.
Log backup	Backs up all the transaction logs. The truncated logs that are already committed to the database are not backed up. If you schedule frequent transaction log backups between full database backups, you can choose granular recovery points.

Backup schedules for database plug-ins

Backup frequency (schedule type) is specified in policies; a backup schedule is specified in the resource group configuration. The most critical factor in determining a backup frequency or schedule is the rate of change for the resource and the importance of the data. You might back up a heavily used resource every hour, while you might back up a rarely used resource once a day. Other factors include the importance of the resource to your organization, your Service Level Agreement (SLA), and your Recover Point Objective (RPO).

An SLA defines the level of service expected and addresses many service-related issues, including the availability and performance of service. An RPO defines the strategy for the age of the files that must be recovered from backup storage for regular operations to resume after a failure. The SLA and RPO contribute to the data protection strategy.

Even for a heavily used resource, there is no requirement to run a full backup more than once or twice a day. For example, regular transaction log backups might be sufficient to ensure that you have the backups you need. The more often you back up your databases, the fewer transaction logs SnapCenter has to use at restore time, which can result in faster restore operations.

Backup schedules have two parts, as follows:

- Backup frequency

  Backup frequency (how often backups are to be performed), called *schedule type* for some plug-ins, is part of a policy configuration. You can select hourly, daily, weekly, or monthly as the backup frequency for the policy. If you do not select any of these frequencies, then the policy created is an on-demand-only policy. You can access policies by clicking **Settings** > **Policies**.

- Backup schedules

Backup schedules (exactly when backups are to be performed) are part of a resource group configuration. For example, if you have a resource group that has a policy configured for weekly backups, you might configure the schedule to back up every Thursday at 10:00 PM. You can access resource group schedules by clicking **Resources** > **Resource Groups**.

Number of backup jobs needed for databases

Factors that determine the number of backup jobs that you need include the size of the resource, the number of volumes used, the rate of change of the resource, and your Service Level Agreement (SLA).

Backup naming conventions

You can either use the default Snapshot naming convention or use a customized naming convention. The default backup naming convention adds a timestamp to Snapshot names that helps you identify when the copies were created.

The Snapshot uses the following default naming convention:

```
resourcegroupname_hostname_timestamp
```

You should name your backup resource groups logically, as in the following example:

```
dts1_mach1x88_03-12-2015_23.17.26
```

In this example, the syntax elements have the following meanings:

- *dts1* is the resource group name.
- *mach1x88* is the host name.
- *03-12-2015_23.17.26* is the date and timestamp.

Alternatively, you can specify the Snapshot name format while protecting resources or resource groups by selecting **Use custom name format for Snapshot copy**. For example, customtext_resourcegroup_policy_hostname or resourcegroup_hostname. By default, the time stamp suffix is added to the Snapshot name.

Backup retention options

You can choose either the number of days for which to retain backup copies or specify the number of backup copies you want to retain, up to a ONTAP maximum of 255 copies. For example, your organization might require that you retain 10 days of backup copies or 130 backup copies.

While creating a policy, you can specify the retention options for the backup type and the schedule type.

If you set up SnapMirror replication, the retention policy is mirrored on the destination volume.

SnapCenter deletes the retained backups that have retention labels that match the schedule type. If the schedule type was changed for the resource or resource group, backups with the old schedule type label might still remain on the system.

For long-term retention of backup copies, you should use SnapVault backup.

How long to retain transaction log backups on the source storage volume for Exchange Server

SnapCenter Plug-in for Microsoft Exchange Server needs transaction log backups to perform up-to-the-minute restore operations, which restore your database to a time between two full backups.

For example, if Plug-in for Exchange took a full plus transaction log backup at 8:00 a.m. and another full plus transaction log backup at 5:00 p.m., it could use the latest transaction log backup to restore the database to any time between 8:00 a.m. and 5:00 p.m. If transaction logs are not available, Plug-in for Exchange can perform point-in-time restore operations only, which restore a database to the time that Plug-in for Exchange completed a full backup.

Typically, you require up-to-the-minute restore operations for only a day or two. By default, SnapCenter retains a minimum of two days.

Define a restore strategy for Exchange databases

Defining a restoration strategy for Exchange Server enables you to restore your database successfully.

Sources for a restore operation in Exchange Server

You can restore an Exchange Server database from a backup copy on primary storage.

You can restore databases from primary storage only.

Types of restore operations supported for Exchange Server

You can use SnapCenter to perform different types of restore operations on Exchange resources.

- Restore up-to-the-minute
- Restore to a previous point in time

Restore up to the minute

In an up-to-the-minute restore operation, databases are recovered up to the point of failure. SnapCenter accomplishes this by performing the following sequence:

1. Restores the databases from the full database backup that you select.
2. Applies all the transaction logs that were backed up, as well as any new logs that were created since the most recent backup.

   Transaction logs are moved ahead and applied to any selected databases.

   Exchange creates a new log chain after a restore completes.

   Best Practice: It is recommended that you perform a new full and log backup after a restore completes.

   An up-to-the-minute restore operation requires a contiguous set of transaction logs.

   After you perform an up-to-the-minute restore, the backup you used for the restore is available only for point-in-time restore operations.

   If you do not need to retain up-to-the-minute restore capability for all backups, you can configure your

system's transaction log backup retention through the backup policies.

Restore to a previous point in time

In a point-in-time restore operation, databases are restored only to a specific time from the past. A point-in-time restore operation occurs in the following restore situations:

- The database is restored to a given time in a backed-up transaction log.
- The database is restored, and only a subset of backed-up transaction logs are applied to it.

Install SnapCenter Plug-in for Microsoft Exchange Server

Installation workflow of SnapCenter Plug-in for Microsoft Exchange Server

You should install and set up SnapCenter Plug-in for Microsoft Exchange Server if you want to protect Exchange databases.

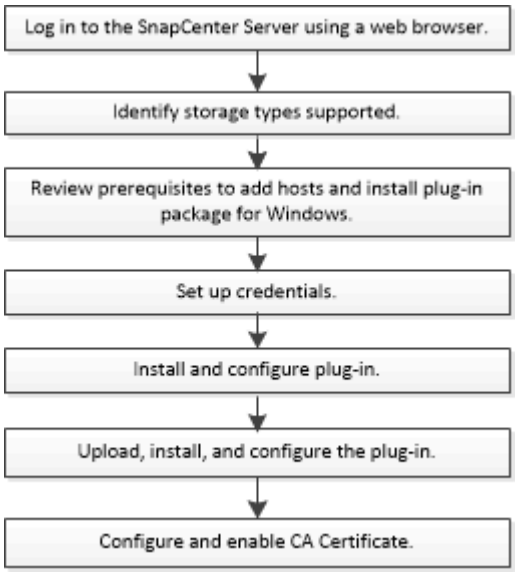

Prerequisites to add hosts and install SnapCenter Plug-in for Microsoft Exchange Server

Before you add a host and install the plug-in packages, you must complete all the requirements.

- If you are using iSCSI, the iSCSI service must be running.
- You must have a domain user with local administrator privileges with local login permissions on the remote host.
- You must be using Microsoft Exchange Server 2013, 2016, or 2019 for standalone and Database Availability Group configurations.
- When installing a plug-in on a Windows host, if you specify a credential that is not built-in or if the user belongs to a local workgroup user, you must disable UAC on the host.
- If you manage cluster nodes in SnapCenter, you must have a user with administrative privileges to all the nodes in the cluster.

- You must have a user with administrative permissions on the Exchange Server.
- If SnapManager for Microsoft Exchange Server and SnapDrive for Windows are already installed, you must unregister the VSS Hardware Provider used by SnapDrive for Windows before you install Plug-in for Exchange on the same Exchange Server to ensure successful data protection using SnapCenter.
- If SnapManager for Microsoft Exchange Server and Plug-in for Exchange are installed on the same server, you must suspend or delete from the Windows scheduler all schedules created by SnapManager for Microsoft Exchange Server.
- The host must be resolvable to the fully qualified domain name (FQDN) from the server. If the hosts file is modified to make it resolvable and if both the short name and the FQDN are specified in the hosts file, create an entry in the SnapCenter hosts file in the following format: *<ip_address> <host_fqdn> <host_name>*.
- Ensure the following ports are not blocked in the firewall, otherwise the add host operation fails. To resolve this issue, you must configure the dynamic port range. For more information, see Microsoft documentation.
  - Port range 50000 - 51000 for Windows 2016 and Exchange 2016
  - Port range 6000 - 6500 for Windows 2012 R2 and Exchange 2013
  - Port range 49152 - 65536 for Windows 2019

To identify the port range, execute the following commands:

- netsh int ipv4 show dynamicport tcp
- netsh int ipv4 show dynamicport udp
- netsh int ipv6 show dynamicport tcp
- netsh int ipv6 show dynamicport udp

Host requirements to install SnapCenter Plug-ins Package for Windows

Before you install the SnapCenter Plug-ins Package for Windows, you should be familiar with some basic host system space requirements and sizing requirements.

Item	Requirements
Operating Systems	Microsoft Windows For the latest information about supported versions, see the NetApp Interoperability Matrix Tool.
Minimum RAM for the SnapCenter plug-in on host	1 GB

Item	Requirements
Minimum install and log space for the SnapCenter plug-in on host	5 GB You should allocate sufficient disk space and monitor the storage consumption by the logs folder. The log space required varies depending on the number of the entities to be protected and the frequency of data protection operations. If there is no sufficient disk space, the logs will not be created for the recently run operations.
Required software packages	• ASP.NET Core Runtime 8.0.12 (and all subsequent 8.0.x patches) Hosting Bundle • PowerShell Core 7.4.2 • Java 11 Oracle Java and OpenJDK Java 11 Oracle Java and OpenJDK is required only for SAP HANA, IBM Db2, PostgreSQL, MySQL, NetApp supported plug-ins, and other custom applications which can be installed on Windows host. For the latest information about supported versions, see the NetApp Interoperability Matrix Tool. For .NET specific troubleshooting information, see SnapCenter upgrade or install fails for legacy systems that do not have internet connectivity.

Exchange Server privileges required

To enable SnapCenter to add Exchange Server or DAG, and to install SnapCenter Plug-in for Microsoft Exchange Server on a host or DAG, you must configure SnapCenter with credentials for a user with a minimum set of privileges and permissions.

You must have a domain user with local administrator privileges, and with local login permissions on the remote Exchange host, as well as administrative permissions on all the nodes in the DAG. The domain user requires the following minimum permissions:

- Add-MailboxDatabaseCopy
- Dismount-Database
- Get-AdServerSettings
- Get-DatabaseAvailabilityGroup
- Get-ExchangeServer
- Get-MailboxDatabase

- Get-MailboxDatabaseCopyStatus
- Get-MailboxServer
- Get-MailboxStatistics
- Get-PublicFolderDatabase
- Move-ActiveMailboxDatabase
- Move-DatabasePath -ConfigurationOnly:$true
- Mount-Database
- New-MailboxDatabase
- New-PublicFolderDatabase
- Remove-MailboxDatabase
- Remove-MailboxDatabaseCopy
- Remove-PublicFolderDatabase
- Resume-MailboxDatabaseCopy
- Set-AdServerSettings
- Set-MailboxDatabase -allowfilerestore:$true
- Set-MailboxDatabaseCopy
- Set-PublicFolderDatabase
- Suspend-MailboxDatabaseCopy
- Update-MailboxDatabaseCopy

Host requirements to install SnapCenter Plug-ins Package for Windows

Before you install the SnapCenter Plug-ins Package for Windows, you should be familiar with some basic host system space requirements and sizing requirements.

Item	Requirements
Operating Systems	Microsoft Windows For the latest information about supported versions, see the NetApp Interoperability Matrix Tool.
Minimum RAM for the SnapCenter plug-in on host	1 GB

Item	Requirements
Minimum install and log space for the SnapCenter plug-in on host	5 GB You should allocate sufficient disk space and monitor the storage consumption by the logs folder. The log space required varies depending on the number of the entities to be protected and the frequency of data protection operations. If there is no sufficient disk space, the logs will not be created for the recently run operations.
Required software packages	• ASP.NET Core Runtime 8.0.12 (and all subsequent 8.0.x patches) Hosting Bundle • PowerShell Core 7.4.2 • Java 11 Oracle Java and OpenJDK Java 11 Oracle Java and OpenJDK is required only for SAP HANA, IBM Db2, PostgreSQL, MySQL, NetApp supported plug-ins, and other custom applications which can be installed on Windows host. For the latest information about supported versions, see the NetApp Interoperability Matrix Tool. For .NET specific troubleshooting information, see SnapCenter upgrade or install fails for legacy systems that do not have internet connectivity.

Set up credentials for SnapCenter Plug-in for Windows

SnapCenter uses credentials to authenticate users for SnapCenter operations. You should create credentials for installing the plug-in package and additional credentials for performing data protection operations on databases.

About this task

You must set up credentials for installing plug-ins on Windows hosts. Although you can create credentials for Windows after deploying hosts and installing plug-ins, the best practice is to create credentials after you add SVMs, before you deploy hosts and install plug-ins.

Set up the credentials with administrator privileges, including administrator rights on the remote host.

If you set up credentials for individual resource groups and the username does not have full admin privileges, you must assign at least the resource group and backup privileges to the username.

Steps

1. In the left navigation pane, click **Settings**.
2. In the Settings page, click **Credential**.
3. Click **New**.

   The Credential window is displayed.

4. In the Credential page, do the following:

For this field…	Do this…
Credential name	Enter a name for the credential.
Username	Enter the user name used for authentication. • Domain administrator or any member of the administrator group Specify the domain administrator or any member of the administrator group on the system on which you are installing the SnapCenter plug-in. Valid formats for the Username field are: ◦ `NetBIOS\UserName` ◦ `Domain FQDN\UserName` • Local administrator (for workgroups only) For systems that belong to a workgroup, specify the built-in local administrator on the system on which you are installing the SnapCenter plug-in. You can specify a local user account that belongs to the local administrators group if the user account has elevated privileges or the User Access control feature is disabled on the host system. The valid format for the Username field is: `UserName`
Password	Enter the password used for authentication.
Authentication	Select Windows as the authentication mode.

5. Click **OK**.

Configure gMSA on Windows Server 2016 or later

Windows Server 2016 or later enables you to create a group Managed Service Account (gMSA) that provides automated service account password management from a managed domain account.

Before you begin

- You should have a Windows Server 2016 or later domain controller.
- You should have a Windows Server 2016 or later host, which is a member of the domain.

Steps

1. Create a KDS root key to generate unique passwords for each object in your gMSA.
2. For each domain, run the following command from the Windows domain controller: Add-KDSRootKey -EffectiveImmediately
3. Create and configure your gMSA:
   a. Create a user group account in the following format:

      ```
      domainName\accountName$
      ```

   b. Add computer objects to the group.
   c. Use the user group you just created to create the gMSA.

      For example,

      ```
      New-ADServiceAccount -name <ServiceAccountName> -DNSHostName <fqdn>
      -PrincipalsAllowedToRetrieveManagedPassword <group>
      -ServicePrincipalNames <SPN1,SPN2,…>
      ```

   d. Run `Get-ADServiceAccount` command to verify the service account.
4. Configure the gMSA on your hosts:
   a. Enable the Active Directory module for Windows PowerShell on the host where you want to use the gMSA account.

      To do this, run the following command from PowerShell:

```
PS C:\> Get-WindowsFeature AD-Domain-Services

Display Name                            Name                Install
State
------------                            ----
-------------
[ ] Active Directory Domain Services   AD-Domain-Services  Available


PS C:\> Install-WindowsFeature AD-DOMAIN-SERVICES

Success Restart Needed Exit Code      Feature Result
------- -------------- ---------      --------------
True    No             Success        {Active Directory Domain
Services, Active ...
WARNING: Windows automatic updating is not enabled. To ensure that
your newly-installed role or feature is
automatically updated, turn on Windows Update.
```

b. Restart your host.

c. Install the gMSA on your host by running the following command from the PowerShell command prompt: `Install-AdServiceAccount <gMSA>`

d. Verify your gMSA account by running the following command: `Test-AdServiceAccount <gMSA>`

5. Assign the administrative privileges to the configured gMSA on the host.
6. Add the Windows host by specifying the configured gMSA account in the SnapCenter Server.

   SnapCenter Server will install the selected plug-ins on the host and the specified gMSA will be used as the service log on account during the plug-in installation.

Add hosts and install Plug-in for Exchange

You can use the SnapCenter Add Host page to add Windows hosts. The Plug-in for Exchange is automatically installed on the specified host. This is the recommended method for installing plug-ins. You can add a host and install a plug-in either for an individual host or a cluster.

Before you begin

- If the operating system of the SnapCenter Server host is Windows 2019 and the operating system of the plug-in host is Windows 2022, you should perform the following:
  - Upgrade to Windows Server 2019 (OS Build 17763.5936) or later
  - Upgrade to Windows Server 2022 (OS Build 20348.2402) or later
- You must be a user that is assigned to a role that has the plug-in install and uninstall permissions, such as the SnapCenter Admin.
- When installing a plug-in on a Windows host, if you specify a credential that is not built-in, or if the user

belongs to a local workgroup user, you must disable UAC on the host.

- The message queueing service must be running.
- If you are using group Managed Service Account (gMSA), you should configure gMSA with administrative privileges. For information, see Configure group Managed Service Account on Windows Server 2016 or later for Microsoft Exchange Server.

About this task

- You cannot add a SnapCenter Server as a plug-in host to another SnapCenter Server.
- You can add a host and install plug-in packages either for an individual host or a cluster.
- If an exchange node is part of a DAG, you cannot add only one node into the SnapCenter Server.
- If you are installing plug-ins on a cluster (Exchange DAG), they are installed on all of the nodes of the cluster even if some of nodes do not have databases on NetApp LUNs.

Beginning with SnapCenter 4.6, SCE supports multitenancy and you can add a host using the following methods:

Add host operation	4.5 and earlier	4.6 and later
Add IP-less DAG in cross or different domain	Not supported	Supported
Add multiple IP DAGs with unique names, residing in the same or cross domain	Supported	Supported
Add multiple IP or IP-less DAGs which have same host names and/or DB name in cross domain	Not supported	Supported
Add multiple IP/IP-less DAGs with the same name and cross domain	Not supported	Supported
Add multiple standalone hosts with the same name and cross domain	Not supported	Supported

Plug-in for Exchange depends on SnapCenter Plug-ins Package for Windows, and the versions must be the same. During the Plug-in for Exchange installation, SnapCenter Plug-ins Package for Windows is selected by default and is installed along with the VSS Hardware Provider.

If SnapManager for Microsoft Exchange Server and SnapDrive for Windows are already installed, and you want to install Plug-in for Exchange on the same Exchange Server, you must unregister the VSS Hardware Provider used by SnapDrive for Windows because it is incompatible with the VSS Hardware Provider installed with Plug-in for Exchange and SnapCenter Plug-ins Package for Windows. For more information, see How to manually register the Data ONTAP VSS Hardware Provider.

Steps

1. In the left navigation pane, click **Hosts**.
2. Verify that **Managed Hosts** is selected at the top.
3. Click **Add**.
4. In the Hosts page, do the following:

For this field…	Do this…
Host Type	Select **Windows** as the host type. SnapCenter Server adds the host and then installs on the host the Plug-in for Windows and the Plug-in for Exchange if they are not already installed. Plug-in for Windows and Plug-in for Exchange must be the same version. If a different version of Plug-in for Windows was previously installed, SnapCenter updates the version as part of the installation.
Host name	Enter the fully qualified domain name (FQDN) or the IP address of the host. SnapCenter depends on the proper configuration of the DNS. Therefore, the best practice is to enter the fully qualified domain name (FQDN). An IP address is supported for untrusted domain hosts only if it resolves to the FQDN. If you are adding a host using SnapCenter and it is part of a subdomain, you must provide the FQDN. You can enter IP addresses or the FQDN of one of the following: • Stand-alone host • Exchange DAG For an Exchange DAG, you can: ◦ Add a DAG by providing the DAG name, DAG IP address, node name, or node IP address. ◦ Add the IP less DAG cluster by providing the IP address or the FQDN of one of the DAG cluster nodes. ◦ Add IP less DAG that resides in the same domain or different domain. You can also add multiple IP/IP less DAGs with the same name but different domains. For a stand-alone host or an Exchange DAG (cross-domain or same domain), it is recommended to provide FQDN or the IP address of the host or DAG.

For this field…	Do this…
Credentials	Select the credential name that you created, or create the new credentials. The credential must have administrative rights on the remote host. For details, see information about creating a credential. You can view details about the credentials by positioning your cursor over the credential name that you specified. Credentials authentication mode is determined by the host type that you specify in the Add Host wizard.

5. In the Select Plug-ins to Install section, select the plug-ins to install.

   When you select Plug-in for Exchange, SnapCenter Plug-in for Microsoft SQL Server is deselected automatically. Microsoft recommends that SQL Server and Exchange server not be installed on the same system due to the amount of memory used and other resource usage required by Exchange.

6. (Optional) Click **More Options**.

For this field…	Do this…
Port	Either retain the default port number or specify the port number. The default port number is 8145. If the SnapCenter Server was installed on a custom port, that port number will be displayed as the default port. If you manually installed the plug-ins and specified a custom port, you must specify the same port. Otherwise, the operation fails.
Installation Path	The default path is `C:\Program Files\NetApp\SnapCenter`. You can optionally customize the path.
Add all hosts in the DAG	Select this check box when you add a DAG.
Skip preinstall checks	Select this check box if you already installed the plug-ins manually and you do not want to validate whether the host meets the requirements for installing the plug-in.

For this field…	Do this…
Use group Managed Service Account (gMSA) to run the plug-in services	Select this check box if you want to use group Managed Service Account (gMSA) to run the plug-in services. Provide the gMSA name in the following format: *domainName\accountName$*. gMSA will be used as a log on service account only for SnapCenter Plug-in for Windows service.

7. Click **Submit**.

   If you have not selected the Skip prechecks check box, the host is validated to determine whether it meets the requirements to install the plug-in. If the minimum requirements are not met, the appropriate error or warning messages are displayed.

   If the error is related to disk space or RAM, you can update the web.config file located at `C:\Program Files\NetApp\SnapCenter` WebApp to modify the default values. If the error is related to other parameters, you must fix the issue.

   In an HA setup, if you are updating web.config file, you must update the file on both nodes.

8. Monitor the installation progress.

Configure custom port for NET TCP communication

By default, beginning with SnapCenter 6.0 release, SnapCenter plug-in for Windows uses the port 909 for NET TCP communication. If the port 909 is in use, you can configure another port for NET TCP communication.

Steps

1. Modify the value of the *NetTCPPort* key located at *C:\Program Files\NetApp\SnapCenter\SnapCenter Plug-in for Microsoft Windows\vssproviders\navssprv.exe.config* to the required port number.
   `<add key="NetTCPPort" value="new_port_number" />`
2. Modify the value of the *NetTCPPort* key located at *C:\Program Files\NetApp\SnapCenter\SnapCenter Plug-in for Microsoft Windows\SnapDriveService.dll.config* to the required port number.
   `<add key="NetTCPPort" value="new_port_number" />`
3. Unregister the *Data ONTAP VSS Hardware Provider* service by running below command:
   `"C:\Program Files\NetApp\SnapCenter\SnapCenter Plug-in for Microsoft Windows\navssprv.exe" -r service -u`

   Verify that the service is not displayed in the list of services in *services.msc*.

4. Register the *Data ONTAP VSS Hardware Provider* service by running below command:
   `"C:\Program Files\NetApp\SnapCenter\SnapCenter Plug-in for Microsoft Windows\vssproviders\navssprv.exe" -r service -a ".\LocalSystem"`

   Verify if the service is now displayed in the list of services in *services.msc*.

5. Restart the *Plug-in for Windows* service.

Install Plug-in for Exchange from the SnapCenter Server host using PowerShell cmdlets

You should install the Plug-in for Exchange from the SnapCenter GUI. If you do not want to use the GUI, you can use PowerShell cmdlets on the SnapCenter Server host or on a remote host.

Before you begin

- SnapCenter Server must have been installed and configured.
- You must be a local administrator on the host or a user with administrative privileges.
- You must be a user that is assigned to a role that has the plug-in, install, and uninstall permissions, such as the SnapCenter Admin.
- You must have reviewed the installation requirements and types of supported configurations before installing the Plug-in for Exchange.
- The host on which you want the Plug-in for Exchange installed must be a Windows host.

Steps

1. On the SnapCenter Server host, establish a session using the *Open-SmConnection* cmdlet, and then enter your credentials.
2. Add the host on which you want to install the Plug-in for Exchange using the *Add-SmHost* cmdlet with the required parameters.

   The information regarding the parameters that can be used with the cmdlet and their descriptions can be obtained by running *Get-Help command_name*. Alternatively, you can also refer to the SnapCenter Software Cmdlet Reference Guide.

   The host can be a standalone host or a DAG. If you specify a DAG, the *-IsDAG* parameter is required.
3. Install the Plug-in for Exchange using the *Install-SmHostPackage* cmdlet with the required parameters.

   This command installs the Plug-in for Exchange on the specified host, and then registers the plug-in with SnapCenter.

Install the SnapCenter Plug-in for Exchange silently from the command line

You should install Plug-in for Exchange from within the SnapCenter user interface. However, if you cannot for some reason, you can run the Plug-in for Exchange installation program unattended in silent mode from the Windows command line.

Before you begin

- You must have backed up your Microsoft Exchange Server resources.
- You must have installed the SnapCenter plug-in packages.
- You must delete the earlier release of SnapCenter Plug-in for Microsoft SQL Server before installing.

  For more information, see How to Install a SnapCenter Plug-In manually and directly from the Plug-In Host.

Steps

1. Validate whether *C:\temp* folder exists on the plug-in host and the logged in user has full access to it.
2. Download the SnapCenter Plug-in for Microsoft Windows from *C:\ProgramData\NetApp\SnapCenter\Package* Repository.

   This path is accessible from the host where the SnapCenter Server is installed.
3. Copy the installation file to the host on which you want to install the plug-in.
4. From a Windows command prompt on the local host, navigate to the directory to which you saved the plug-in installation files.
5. Enter the following command to install the plug-in.

   snapcenter_windows_host_plugin.exe"/silent /debuglog"<Debug_Log_Path>" /log"<Log_Path>" BI_SNAPCENTER_PORT=<Num> SUITE_INSTALLDIR="<Install_Directory_Path>" BI_SERVICEACCOUNT=<domain\administrator> BI_SERVICEPWD=<password> ISFeatureInstall=HPPW,SCW,SCE

   For example:

   C:\ProgramData\NetApp\SnapCenter\Package Repository\snapcenter_windows_host_plugin.exe"/silent /debuglog"C:\HPPW_SCSQL_Install.log" /log"C:\temp" BI_SNAPCENTER_PORT=8145 SUITE_INSTALLDIR="C:\Program Files\NetApp\SnapCenter" BI_SERVICEACCOUNT=domain\administrator BI_SERVICEPWD=password ISFeatureInstall=HPPW,SCW,SCE

   All the parameters passed during the installation of Plug-in for Exchange are case sensitive.

   Enter the following values for the variables:

Variable	Value
/debuglog"<Debug_Log_Path>	Specify the name and location of the suite installer log file, as in the following example: *Setup.exe /debuglog"C:\PathToLog\setupexe.log*
BI_SNAPCENTER_PORT	Specify the port on which SnapCenter communicates with SMCore.
SUITE_INSTALLDIR	Specify host plug-in package installation directory.
BI_SERVICEACCOUNT	Specify SnapCenter Plug-in for Microsoft Windows web service account.
BI_SERVICEPWD	Specify the password for SnapCenter Plug-in for Microsoft Windows web service account.
ISFeatureInstall	Specify the solution to be deployed by SnapCenter on remote host.

6. Monitor the Windows task scheduler, the main installation log file *C:\Installdebug.log*, and the additional installation files in *C:\Temp*.
7. Monitor the *%temp%* directory to check if the *msiexe.exe* installers are installing the software without errors.

The installation of Plug-in for Exchange registers the plug-in on the host and not on the SnapCenter Server. You can register the plug-in on the SnapCenter Server by adding the host using the SnapCenter GUI or PowerShell cmdlet. After the host is added, the plug-in is automatically discovered.

Monitor SnapCenter plug-in package installation status

You can monitor the progress of SnapCenter plug-in package installation by using the Jobs page. You might want to check the progress of installation to determine when it is complete or if there is an issue.

About this task

The following icons appear on the Jobs page and indicate the state of the operation:

- In progress
- Completed successfully
- Failed
- Completed with warnings or could not start due to warnings
- Queued

Steps

1. In the left navigation pane, click **Monitor**.
2. In the **Monitor** page, click **Jobs**.
3. In the **Jobs** page, to filter the list so that only plug-in installation operations are listed, do the following:
   a. Click **Filter**.
   b. Optional: Specify the start and end date.
   c. From the Type drop-down menu, select **Plug-in installation**.
   d. From the Status drop-down menu, select the installation status.
   e. Click **Apply**.
4. Select the installation job and click **Details** to view the job details.
5. In the **Job Details** page, click **View logs**.

Configure CA Certificate

Generate CA Certificate CSR file

You can generate a Certificate Signing Request (CSR) and import the certificate that can be obtained from a Certificate Authority (CA) using the generated CSR. The certificate will

have a private key associated with it.

CSR is a block of encoded text that is given to an authorized certificate vendor to procure the signed CA certificate.

CA Certificate RSA key length must be minimum 3072 bits.

For information to generate a CSR, see How to generate CA Certificate CSR file.

If you own the CA certificate for your domain (*.domain.company.com) or your system (machine1.domain.company.com), you can skip generating the CA Certificate CSR file. You can deploy the existing CA certificate with SnapCenter.

For cluster configurations, the cluster name (virtual cluster FQDN), and the respective host names should be mentioned in the CA certificate. The certificate can be updated by filling the Subject Alternative Name (SAN) field before procuring the certificate. For a wild card certificate (*.domain.company.com), the certificate will contain all the hostnames of the domain implicitly.

Import CA certificates

You must import the CA certificates to the SnapCenter Server and the Windows host plug-ins using the Microsoft management console (MMC).

Steps

1. Go to the Microsoft management console (MMC), and then click **File** > **Add/Remove Snapin**.
2. In the Add or Remove Snap-ins window, select **Certificates** and then click **Add**.
3. In the Certificates snap-in window, select the **Computer account** option, and then click **Finish**.
4. Click **Console Root** > **Certificates – Local Computer** > **Trusted Root Certification Authorities** > **Certificates**.
5. Right-click on the folder “Trusted Root Certification Authorities”, and then select **All Tasks** > **Import** to start the import wizard.
6. Complete the wizard, as follows:

In this wizard window…	Do the following…
Import Private Key	Select the option **Yes**, import the private key, and then click **Next**.
Import File Format	Make no changes; click **Next**.
Security	Specify the new password to be used for the exported certificate, and then click **Next**.
Completing the Certificate Import Wizard	Review the summary, and then click **Finish** to start the import.

Importing certificate should be bundled with the private key (supported formats are: *.pfx, *.p12, and *.p7b).

7. Repeat Step 5 for the “Personal” folder.

Get the CA certificate thumbprint

A certificate thumbprint is a hexadecimal string that identifies a certificate. A thumbprint is calculated from the content of the certificate using a thumbprint algorithm.

Steps

1. Perform the following on the GUI:
   a. Double-click the certificate.
   b. In the Certificate dialog box, click the **Details** tab.
   c. Scroll through the list of fields and click **Thumbprint**.
   d. Copy the hexadecimal characters from the box.
   e. Remove the spaces between the hexadecimal numbers.

      For example, if the thumbprint is: "a9 09 50 2d d8 2a e4 14 33 e6 f8 38 86 b0 0d 42 77 a3 2a 7b", after removing the spaces, it will be: "a909502dd82ae41433e6f83886b00d4277a32a7b".

2. Perform the following from PowerShell:
   a. Run the following command to list the thumbprint of the installed certificate and identify the recently installed certificate by the subject name.

      Get-ChildItem -Path Cert:\LocalMachine\My

   b. Copy the thumbprint.

Configure CA certificate with Windows host plug-in services

You should configure the CA certificate with Windows host plug-in services to activate the installed digital certificate.

Perform the following steps on the SnapCenter Server and all the plug-in hosts where CA certificates are already deployed.

Steps

1. Remove the existing certificate binding with SMCore default port 8145, by running the following command:

   ```
   > netsh http delete sslcert ipport=0.0.0.0:_<SMCore Port>
   ```

   For example:

   ```
   > netsh http delete sslcert ipport=0.0.0.0:8145
   ```

2. Bind the newly installed certificate with the Windows host plug-in services, by running the following commands:

```
> $cert = "_<certificate thumbprint>_"
> $guid = [guid]::NewGuid().ToString("B")
> netsh http add sslcert ipport=0.0.0.0: _<SMCore Port>_ certhash=$cert
appid="$guid"
```

For example:

```
> $cert = "a909502dd82ae41433e6f83886b00d4277a32a7b"
> $guid = [guid]::NewGuid().ToString("B")
> netsh http add sslcert ipport=0.0.0.0: _<SMCore Port>_ certhash=$cert
appid="$guid"
```

Enable CA Certificates for plug-ins

You should configure the CA certificates and deploy the CA certificates in the SnapCenter Server and the corresponding plug-in hosts. You should enable the CA certificate validation for the plug-ins.

Before you begin

- You can enable or disable the CA certificates using the run *Set-SmCertificateSettings* cmdlet.
- You can display the certificate status for the plug-ins using the *Get-SmCertificateSettings*.

The information regarding the parameters that can be used with the cmdlet and their descriptions can be obtained by running *Get-Help command_name*. Alternatively, you can also refer to the SnapCenter Software Cmdlet Reference Guide.

Steps

1. In the left navigation pane, click **Hosts**.
2. In the Hosts page, click **Managed Hosts**.
3. Select single or multiple plug-in hosts.
4. Click **More options**.
5. Select **Enable Certificate Validation**.

After you finish

The Managed Hosts tab host displays a padlock and the color of the padlock indicates the status of the connection between SnapCenter Server and the plug-in host.

- indicates that the CA certificate is neither enabled nor assigned to the plug-in host.
- indicates that the CA certificate is successfully validated.
- indicates that the CA certificate could not be validated.
- indicates that the connection information could not be retrieved.

> When the status is yellow or green, the data protection operations completes successfully.

Configure SnapManager 7.x for Exchange and SnapCenter to coexist

To enable SnapCenter Plug-in for Microsoft Exchange Server to coexist with SnapManager for Microsoft Exchange Server, you need to install SnapCenter Plug-in for Microsoft Exchange Server on the same Exchange Server on which SnapManager for Microsoft Exchange Server is installed, disable SnapManager for Exchange schedules, and configure new schedules and backups using SnapCenter Plug-in for Microsoft Exchange Server.

Before you begin

- SnapManager for Microsoft Exchange Server and SnapDrive for Windows are already installed, and SnapManager for Microsoft Exchange Server backups exist on the system and in the SnapInfo directory.
- You should have deleted or reclaimed backups taken by SnapManager for Microsoft Exchange Server that you no longer require.
- You should have suspended or deleted all schedules created by SnapManager for Microsoft Exchange Server from the Windows scheduler.
- SnapCenter Plug-in for Microsoft Exchange Server and SnapManager for Microsoft Exchange Server can coexist on the same Exchange Server, but you cannot upgrade existing SnapManager for Microsoft Exchange Server installations to SnapCenter.

  SnapCenter does not provide an option for the upgrade.

- SnapCenter does not support restoring Exchange databases from SnapManager for Microsoft Exchange Server backup.

  If you do not uninstall SnapManager for Microsoft Exchange Server after the SnapCenter Plug-in for Microsoft Exchange Server installation and later want to restore a SnapManager for Microsoft Exchange Server backup, you must perform additional steps.

Steps

1. Using PowerShell on all DAG nodes, determine whether the SnapDrive for Windows VSS Hardware Provider is registered: *vssadmin list providers*

   ```
   C:\Program Files\NetApp\SnapDrive>vssadmin list providers
   vssadmin 1.1 - Volume Shadow Copy Service administrative command-line
   tool
   (C) Copyright 2001-2013 Microsoft Corp.

   Provider name: 'Data ONTAP VSS Hardware Provider'
      Provider type: Hardware
      Provider Id: {ddd3d232-a96f-4ac5-8f7b-250fd91fd102}
      Version: 7. 1. 4. 6845
   ```

2. From the SnapDrive directory, unregister the VSS Hardware Provider from SnapDrive for Windows: *navssprv.exe -r service –u*
3. Verify that the VSS Hardware Provider was removed: *vssadmin list providers*
4. Add the Exchange host to SnapCenter, and then install the SnapCenter Plug-in for Microsoft Windows and

the SnapCenter Plug-in for Microsoft Exchange Server.

5. From the SnapCenter Plug-in for Microsoft Windows directory on all DAG nodes, verify that the VSS Hardware Provider is registered: *vssadmin list providers*

```
[PS] C:\Windows\system32>vssadmin list providers
vssadmin 1.1 - Volume Shadow Copy Service administrative command-line
tool
(C) Copyright 2001-2013 Microsoft Corp.

Provider name: 'Data ONTAP VSS Hardware Provider'
   Provider type: Hardware
   Provider Id: {31fca584-72be-45b6-9419-53a3277301d1}
   Version: 7. 0. 0. 5561
```

6. Stop the SnapManager for Microsoft Exchange Server backup schedules.
7. Using the SnapCenter GUI, create on-demand backups, configure scheduled backups, and configure retention settings.
8. Uninstall SnapManager for Microsoft Exchange Server.

   If you do not uninstall SnapManager for Microsoft Exchange Server now and later want to restore a SnapManager for Microsoft Exchange Server backup:

   a. Unregister SnapCenter Plug-in for Microsoft Exchange Server from all DAG nodes: *navssprv.exe -r service -u*

```
C:\Program Files\NetApp\SnapCenter\SnapCenter Plug-in for Microsoft
Windows>navssprv.exe -r service -u
```

   b. From the *C:\Program Files\NetApp\SnapDrive* directory, register SnapDrive for Windows on all DAG nodes: *navssprv.exe -r service -a hostname\\username -p password*

Install SnapCenter Plug-in for VMware vSphere

If your database or filesystem is stored on virtual machines (VMs), or if you want to protect VMs and datastores, you must deploy the SnapCenter Plug-in for VMware vSphere virtual appliance.

For information to deploy, see Deployment Overview.

Deploy CA certificate

To configure the CA Certificate with SnapCenter Plug-in for VMware vSphere, see Create or import SSL certificate.

Configure the CRL file

SnapCenter Plug-in for VMware vSphere looks for the CRL files in a pre-configured directory. Default directory of the CRL files for SnapCenter Plug-in for VMware vSphere is */opt/netapp/config/crl*.

You can place more than one CRL file in this directory. The incoming certificates will be verified against each CRL.

Prepare for data protection

Before performing any data protection operation such as backup, clone, or restore operations, you must define your strategy and set up the environment. You can also set up the SnapCenter Server to use SnapMirror and SnapVault technology.

To take advantage of SnapVault and SnapMirror technology, you must configure and initialize a data protection relationship between the source and destination volumes on the storage device. You can use NetAppSystem Manager or you can use the storage console command line to perform these tasks.

Find more information

Getting started with the REST API

Prerequisites for using the SnapCenter Plug-in for Microsoft Exchange Server

Before you use the Plug-in for Exchange, the SnapCenter administrator must install and configure the SnapCenter Server and perform the prerequisite tasks.

- Install and configure SnapCenter Server.
- Log in to SnapCenter.
- Configure the SnapCenter environment by adding or assigning storage system connections and creating a credential.

  SnapCenter does not support multiple SVMs with the same name on different clusters. Each SVM supported by SnapCenter must have a unique name.

- Add hosts, install the SnapCenter Plug-in for Microsoft Windows and the SnapCenter Plug-in for Microsoft Exchange Server, and discover (refresh) the resources.
- Perform host-side storage provisioning using the SnapCenter Plug-in for Microsoft Windows.
- If you are using SnapCenter Server to protect Exchange databases that reside on VMware RDM LUNs, you must deploy the SnapCenter Plug-in for VMware vSphere and register the plug-in with SnapCenter. The SnapCenter Plug-in for VMware vSphere documentation has more information.

  VMDKs are not supported.

- Move an existing Microsoft Exchange Server database from a local disk to supported storage using Microsoft Exchange tools.
- Set up SnapMirror and SnapVault relationships, if you want backup replication.

For SnapCenter 4.1.1 users, the SnapCenter Plug-in for VMware vSphere 4.1.1 documentation has information on protecting virtualized databases and file systems. For SnapCenter 4.2.x users, the NetApp Data

Broker 1.0 and 1.0.1, documentation has information on protecting virtualized databases and file systems using the SnapCenter Plug-in for VMware vSphere that is provided by the Linux-based NetApp Data Broker virtual appliance (Open Virtual Appliance format). For SnapCenter 4.3.x users, the SnapCenter Plug-in for VMware vSphere 4.3 documentation has information on protecting virtualized databases and file systems using the Linux-based SnapCenter Plug-in for VMware vSphere virtual appliance (Open Virtual Appliance format).

SnapCenter Plug-in for VMware vSphere documentation

How resources, resource groups, and policies are used for protecting Exchange Server

Before you use SnapCenter, it is helpful to understand basic concepts related to the backup, restore, and reseed operations you want to perform. You interact with resources, resource groups, and policies for different operations.

- Resources are typically mailbox databases or Microsoft Exchange Database Availability Group (DAG) that you back up with SnapCenter.
- A SnapCenter resource group, is a collection of resources on a host or Exchange DAG, and the resource group can include either a whole DAG or individual databases.

  When you perform an operation on a resource group, you perform that operation on the resources defined in the resource group according to the schedule you specify for the resource group.

  You can back up on demand a single resource or a resource group. You also can perform scheduled backups for single resources and resource groups.

  The resource groups were formerly known as datasets.

- The policies specify the backup frequency, copy retention, scripts, and other characteristics of data protection operations.

  When you create a resource group, you select one or more policies for that group. You can also select one or more policies when you perform a backup on demand for a single resource.

Think of a resource group as defining *what* you want to protect and when you want to protect it in terms of day and time. Think of a policy as defining *how* you want to protect it. If you are backing up all databases of a host, for example, you might create a resource group that includes all the databases in the host. You could then attach two policies to the resource group: a daily policy and an hourly policy. When you create the resource group and attach the policies, you might configure the resource group to perform a full backup daily and another schedule that performs log backups hourly. The following image illustrates the relationship between resources, resource groups, and policies for databases:

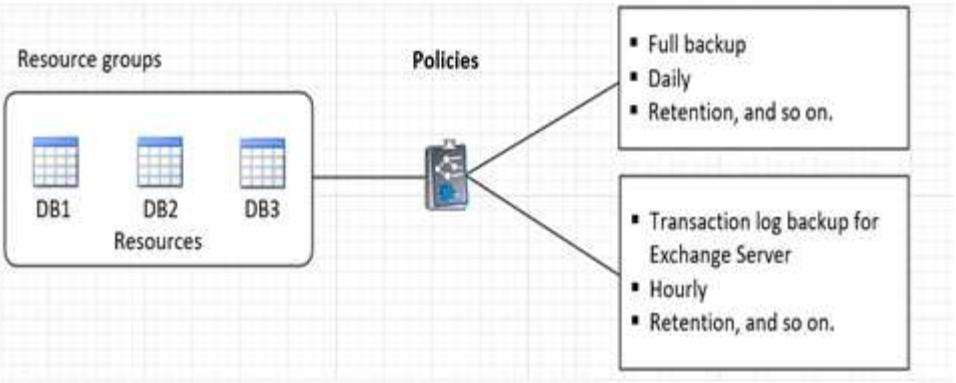

Back up Exchange resources

Backup workflow

When you install the SnapCenter Plug-in for Microsoft Exchange Server in your environment, you can use SnapCenter to back up Exchange resources.

You can schedule multiple backups to run across servers simultaneously. Backup and restore operations cannot be performed simultaneously on the same resource. Active and passive backup copies on the same volume are not supported.

The following workflow shows the sequence in which you must perform the backup operation:

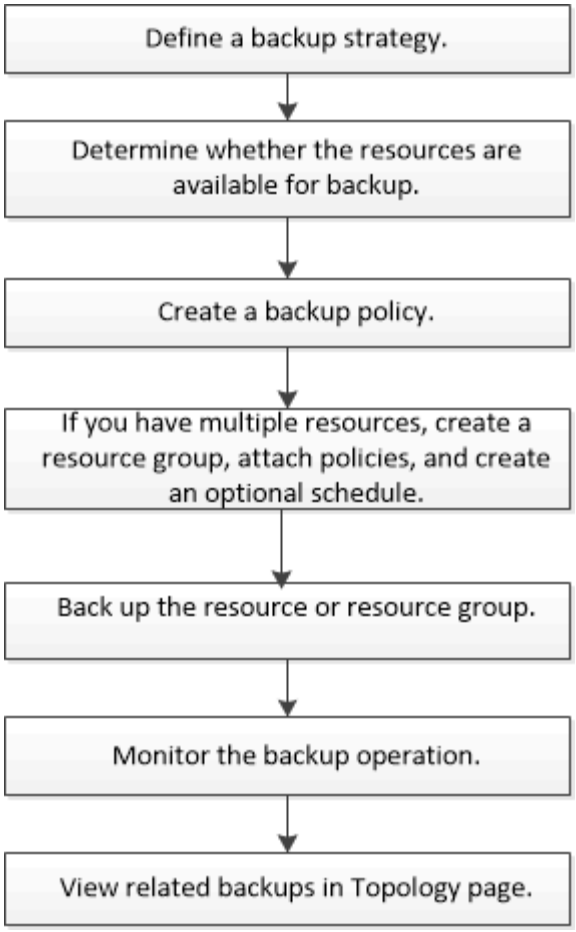

Exchange database and backup verification

SnapCenter Plug-in for Microsoft Exchange Server does not provide backup verification; however, you can use the Eseutil tool provided with Exchange to verify Exchange databases and backups.

The Microsoft Exchange Eseutil tool is a command line utility that is included with your Exchange server. The utility enables you to perform consistency checks to verify the integrity of Exchange databases and backups.

Best Practice: It is not necessary to perform consistency checks on databases that are part of a Database Availability Group (DAG) configuration with at least two replicas.

For additional information, see Microsoft Exchange Server documentation.

Determine whether Exchange resources are available for backup

Resources are the databases, Exchange Database Availability Groups that are maintained by the plug-ins you have installed. You can add those resources to resource groups so that you can perform data protection jobs, but first you must identify which resources you have available. Determining available resources also verifies that the plug-in installation has completed successfully.

Before you begin

- You must have already completed tasks such as installing SnapCenter Server, adding hosts, creating storage system connections, adding credentials, and installing Plug-in for Exchange.
- To take advantage of Single Mailbox Recovery software features, you must have located your active database on the Exchange Server where Single Mailbox Recovery software is installed.
- If databases reside on VMware RDM LUNs, you must deploy the SnapCenter Plug-in for VMware vSphere and register the plug-in with SnapCenter. The SnapCenter Plug-in for VMware vSphere documentation has more information.

About this task

- You cannot back up databases when the **Overall Status** option in the Details page is set to Not available for backup. The **Overall Status** option is set to Not available for backup when any of the following is true:
  - Databases are not on a NetApp LUN.
  - Databases are not in normal state.

    Databases are not in normal state when they are in mount, unmount, reseed, or recovery pending state.
- If you have a Database Availability Group (DAG), you can back up all databases in the group by running the backup job from the DAG.

Steps

1. In the left navigation pane, click **Resources**, and then select **Microsoft Exchange Server** from the plug-ins drop-down list located in the upper left corner of the Resources page.
2. In the Resources page select **Database**, or **Database Availability Group**, or **Resource Group**, from the **View** drop-down list.

   All the databases and DAGs are displayed with their DAG or hostnames in FQDN format, so you can distinguish between multiple databases.

   Click and select the host name and the Exchange Server to filter the resources. You can then click to close the filter pane.
3. Click **Refresh Resources**.

   The newly added, renamed, or deleted resources are updated to the SnapCenter Server inventory.

   You must refresh the resources if the databases are renamed outside of SnapCenter.

   The resources are displayed along with information such as resource name, Database Availability Group name, server in which the database is currently active, server with copies, time of last backup, and overall status.
   - If the database is on a non-NetApp storage, Not available for backup is displayed in the Overall Status column.

     In a DAG, if the active database copy is on non-NetApp storage and if at least one passive database copy is on NetApp storage, Not protected is displayed in the **Overall Status** column.

     You cannot perform data protection operations on a database that is on a non-NetApp storage type.
   - If the database is on NetApp storage and is not protected, Not protected is displayed in the **Overall Status** column.

- If the database is on a NetApp storage system and protected, the user interface displays the Backup not run message in the **Overall Status** column.
- If the database is on a NetApp storage system and is protected and if the backup is triggered for the database, the user interface displays the Backup succeeded message in the **Overall Status** column.

Create backup policies for Exchange Server databases

You can create a backup policy for the Exchange resources or for the resource groups before you use SnapCenter to back up Microsoft Exchange Server resources, or you can create a backup policy at the time you create a resource group or back up a single resource.

Before you begin

- You must have defined your data protection strategy.

  For details, see the information about defining a data protection strategy for Exchange databases.

- You must have prepared for data protection by completing tasks such as installing SnapCenter, adding hosts, identifying resources, and creating storage system connections.
- You must have refreshed (discovered) the Exchange Server resources.
- If you are replicating Snapshots to a mirror or vault, the SnapCenter administrator must have assigned the storage virtual machines (SVMs) for both the source volumes and destination volumes to you.
- If you want to run the PowerShell scripts in prescripts and postscripts, you should set the value of the `usePowershellProcessforScripts` parameter to true in the `web.config` file.

  The default value is false.

- Review the SnapMirror active sync specific prerequisites and limitations. For information, refer Object limits for SnapMirror active sync.

About this task

- A backup policy is a set of rules that governs how you manage and retain backups, and how frequently the resource or resource group is backed up. Additionally, you can specify script settings. Specifying options in a policy saves time when you want to reuse the policy for another resource group.
- Full backup retention is specific to a given policy. A database or resource using policy A with a full backup retention of 4 retains 4 full backups and has no effect on policy B for the same database or resource, which might have a retention of 3 to retain 3 full backups.
- Log backup retention is effective across policies, and applies to all log backups for a database or resource. Therefore, when a full backup is performed using policy B, the log retention setting affects log backups created by policy A on the same database or resource. Similarly, the log retention setting for policy A affects log backups created by policy B on the same database.
- The SCRIPTS_PATH is defined using the PredefinedWindowsScriptsDirectory key located in the SMCoreServiceHost.exe.Config file of the plug-in host.

  If needed, you can change this path and restart SMcore service. It is recommended that you use the default path for security.

  The value of the key can be displayed from swagger through the API: API /4.7/configsettings

  You can use the GET API to display the value of the key. SET API is not supported.

Best Practice: It's best that you configure the secondary retention policy based on the number of full and log backups, overall, that you want to retain. When you configure secondary retention policies, keep in mind that when databases and logs that are in different volumes, each backup can have three Snapshots, and when databases and logs are in the same volume, each backup can have two Snapshots.

- SnapLock
  - If 'Retain the backup copies for a specific number of days' option is selected, then the SnapLock retention period must be lesser than or equal to the mentioned retention days.

    Specifying a Snapshot locking period prevents deletion of the Snapshots until the retention period expires. This could lead to retaining a larger number of Snapshots than the count specified in the policy.

    For ONTAP 9.12.1 and below versions, the clones created from the SnapLock Vault Snapshots will inherit the SnapLock Vault expiry time. Storage admin should manually cleanup the clones post the SnapLock expiry time.

Steps

1. In the left navigation pane, click **Settings**.
2. In the Settings page, click **Policies**.
3. Click **New**.
4. In the Name page, enter the policy name and details.
5. In the Backup Type and Replication page, perform the following steps:
   a. Choose backup type:

If you want to…	Do this…
Back up the database files and the required transaction logs	Select **Full backup and Log backup**. Databases are backed up with log truncation, and all logs are backed up, including the truncated logs. This is the recommended backup type.
Back up the database files and the uncommitted transaction logs	Select **Full backup**. Databases are backed up with log truncation, and truncated logs are not backed up.

If you want to…	Do this…
Back up all the transaction logs	Select **Log backup**. All transaction logs on the active file system are backed up, and there is no log truncation. A *scebackupinfo* directory is created on the same disk as the live log. This directory contains the pointer to the incremental changes for the Exchange database and it is not equivalent to the complete log files.
Back up all database files and transaction logs without truncating the transaction log files	Select **Copy Backup**. All databases and all logs are backed up, and there is no log truncation. You typically use this backup type for reseeding a replica or for testing or diagnosing a problem.

ⓘ You should define the space required for log backups based on the full backup retention and not based on Up-to-the-minute (UTM) retention.

ⓘ Create separate vault policies for logs and databases when dealing with Exchange volumes (LUNs), and set the keep (retention) for the log policy to twice the number for each label as the database policy, using the same labels. For more information see, SnapCenter for Exchange Backups only keep half the Snapshots on the Vault destination log volume

b. In the Database Availability Group Settings section, select an action:

For this field…	Do this…
Back up active copies	Select this option to back up only the active copies of the selected database. For database availability groups (DAGs), this option backs up only active copies of all databases in the DAG. Passive copies are not backed up.
Back up copies on servers to be selected at backup job creation time	Select this option to back up any copies of the databases on the selected servers, both active and passive. For DAGs, this option backs up both active and passive copies of all databases on the selected servers.

In cluster configurations, the backups are retained at each node of the cluster according to the retention settings set in the policy. If the owner node of the cluster changes, the backups of the previous owner node will be retained. The retention is applicable only at the node level.

c. In the Schedule frequency section, select one or more of the frequency types: **On demand**, **Hourly**, **Daily**, **Weekly**, and **Monthly**.

You can specify the schedule (start date, end date) for backup operations while creating a resource group. This enables you to create resource groups that share the same policy and backup frequency, but lets you assign different backup schedules to each policy.

If you have scheduled for 2:00 a.m., the schedule will not be triggered during daylight saving time (DST).

d. Select the Policy label.

You can assign SnapMirror labels to primary snapshots for remote replication, allowing the primary snapshots to offload the snapshot replication operation from SnapCenter to ONTAP secondary systems. This can be done without enabling SnapMirror or SnapVault option in the policy page.

e. In the Select secondary replication options section, select one or both of the following secondary replication options:

For this field…	Do this…
Update SnapMirror after creating a local Snapshot	Select this option to keep mirror copies of backup sets on another volume (SnapMirror). During secondary replication, the SnapLock expiry time loads the primary SnapLock expiry time. This option should be enabled for SnapMirror active sync. The primary-only policy cannot be used if SnapMirror active sync is set up for Exchange ONTAP volumes. SnapCenter does not allow this. You should enable the "Mirror" option. Clicking the **Refresh** button in the Topology page refreshes the secondary and primary SnapLock expiry time that are retrieved from ONTAP. See View Exchange backups in the Topology page.

For this field…	Do this…
Update SnapVault after creating a local Snapshot	Select this option to perform disk-to-disk backup replication.
Error retry count	Enter the number of replication attempts that should occur before the process halts.

You should configure SnapMirror retention policy in ONTAP for the secondary storage to avoid reaching the maximum limit of Snapshots on the secondary storage.

6. In the Retention page, configure the retention settings.

   The options displayed depend upon the backup type and frequency type you previously selected.

   The maximum retention value is 1018. Backups will fail if retention is set to a value higher than what the underlying ONTAP version supports.

   You must set the retention count to 2 or higher if you plan to enable SnapVault replication. If you set the retention count to 1, the retention operation might fail because the first Snapshot is the reference Snapshot for the SnapVault relationship until a newer Snapshot is replicated to the target.

   a. In the Log backups retention settings section, select one of the following:

If you want to…	Do this…
Retain only a specific number of log backups	Select **Number of full backups for which logs are retained**, and specify the number of full backups for which you want up-to-the-minute restorability. Up-to-the-minute (UTM) retention applies to log backup created via full or log backup. For example, if UTM retention settings is configured to retain log backups of the last 5 full backups, then the log backups of the last 5 full backups are retained. The log folders created as part of full and log backups are automatically deleted as part of UTM. You cannot delete the log folders manually. For example, if the retention setting of full or full and log backup is set for 1 month and UTM retention is set to 10 Days, then the log folder created as part of these backups will be deleted as per UTM. As a result, only 10 days log folders will be there and all other backups are marked for point-in-time restore. You can set UTM retention value as 0, if you do not want to perform up-to-the-minute restore. This will enable point-in-time restore operation. **Best Practice:** It's best that the setting must be equal to the setting for Total Snapshots (full backups) in the Full backup retention settings section. This ensures that log files are retained for each full backup.
Retain the backup copies for a specific number of days	Select the **Keep log backups for last** option, and specify the number of days to keep the log backup copies. The log backups up to the number of days of full backups are retained.
Snapshot locking period	Select **Snapshot copy locking period**, and select days, months, or years. SnapLock retention period should be less than 100 years.

If you selected **Log backup** as the backup type, log backups are retained as part of the up-to-the-minute retention settings for full backups.

b. In the Full backup retention settings section, select one of the following for on-demand backups, and then select one for full backups:

For this field…	Do this…
Retain only a specific number of Snapshots	If you want to specify the number of full backups to keep, select the **Total Snapshot copies to keep** option, and specify the number of Snapshots (full backups) to retain. If the number of full backups exceeds the specified number, the full backups that exceed the specified number are deleted, with the oldest copies deleted first.
Retain full backups for a specific number of days	Select the **Keep Snapshot copies for** option, and specify the number of days to keep Snapshots (full backups).
Primary snapshot locking period	Select **Primary snapshot copy locking period**, and select days, months, or years. SnapLock retention period should be less than 100 years.
Secondary snapshot locking period	Select **Secondary snapshot copy locking period**, and select days, months, or years.

If you have a database with only log backups and no full backups on a host in a DAG configuration, the log backups are retained in the following ways:

- By default, SnapCenter finds the oldest full backup for this database in all the other hosts in the DAG, and deletes all log backups on this host that were taken before the full backup.
- You can override the above default retention behavior for a database on a host in a DAG with only log backups by adding the key **MaxLogBackupOnlyCountWithoutFullBackup** in the *C:\Program Files\NetApp\SnapCenter WebApp\web.config* file.

  ```
  <add key="MaxLogBackupOnlyCountWithoutFullBackup" value="10">
  ```

  In the example, the value 10 means you keep up to 10 log backups on the host.

7. In the Script page, enter the path and the arguments of the prescript or postscript that should be run before or after the backup operation, respectively.
   - Prescript backup arguments include “$Database” and “$ServerInstance”.
   - Postscript backup arguments include “$Database”, “$ServerInstance”, “$BackupName”, “$LogDirectory”, and “$LogSnapshot”.

   You can run a script to update SNMP traps, automate alerts, send logs, and so on.

   > The prescripts or postscripts path should not include drives or shares. The path should be relative to the SCRIPTS_PATH.

8. Review the summary, and then click **Finish**.

Create resource groups and attach policies for Exchange Servers

A resource group is required for any data protection job. You must also attach one or more policies to the resource group to define the type of data protection job that you want to perform and the protection schedule.

About this task

- The SCRIPTS_PATH is defined using the PredefinedWindowsScriptsDirectory key located in the SMCoreServiceHost.exe.Config file of the plug-in host.

  If needed, you can change this path and restart SMcore service. It is recommended that you use the default path for security.

  The value of the key can be displayed from swagger through the API: API /4.7/configsettings

  You can use the GET API to display the value of the key. SET API is not supported.

- For ONTAP 9.12.1 and below version, the clones created from the SnapLock Vault Snapshots as part of restore will inherit the SnapLock Vault expiry time. Storage admin should manually cleanup the clones post the SnapLock expiry time.
- Adding new databases without SnapMirror active sync to an existing resource group which contains resources with SnapMirror active sync, is not supported.
- Adding new databases to an existing resource group in failover mode of SnapMirror active sync is not supported. You can add resources to the resource group only in regular or fail-back state.

Steps

1. In the left navigation pane, click **Resources**, and then select the Microsoft Exchange Server plug-in from the list.
2. In the Resources page, select **Database** from the **View** list.

   If you have recently added a resource to SnapCenter, click **Refresh Resources** to view the newly added resource.

3. Click **New Resource Group**.
4. In the Name page, perform the following actions:

For this field…	Do this…
Name	Enter the resource group name. The resource group name should not exceed 250 characters.

For this field…	Do this…
Tags	Enter one or more labels that will help you later search for the resource group. For example, if you add HR as a tag to multiple resource groups, you can later find all resource groups associated with the HR tag.
Use custom name format for Snapshot copy	Optional: Enter a custom Snapshot name and format. For example, *customtext_resourcegroup_policy_hostname* or *resourcegroup_hostname*. By default, a timestamp is appended to the Snapshot name.

5. In the Resources page, perform the following steps:

   a. Select the resource type and the Database Availability Group from drop-down lists to filter the list of available resources.

   If you have recently added resources, they will appear in the list of Available Resources only after you refresh your resource list.

   In the Available Resources and Selected Resources sections, the database name is displayed with the FQDN of the host. This FQDN only indicates that the database is active on that specific host and might not take backup on this host. You should select one or more backup servers from the Server selection option, where you want to take backup in case you have selected the **Back up copies on servers to be selected at backup job creation time** option in the policy.

   a. Type the name of the resource in the search text box, or scroll to locate a resource.

   b. To move resources from the Available Resources section to the Selected Resources section, perform one of the following steps:

      - Select **Autoselect all resources on same storage volume** to move all of the resources on the same volume to the Selected Resources section.
      - Select the resources from the Available Resources section and then click the right arrow to move them to the Selected Resources section.

        Resource groups of SnapCenter for Microsoft Exchange Server cannot have more than 30 databases per Snapshot. If there are more than 30 databases in one resource group, a second Snapshot is created for the additional databases. Therefore, 2 sub jobs are created under the main backup job. For backups having secondary replication, while SnapMirror or SnapVault update is in progress, there could be scenarios where the update for both the sub-jobs overlap. The main backup job keeps on running forever even if the logs indicate that the job is completed.

6. In the Policies page, perform the following steps:

   a. Select one or more policies from the drop-down list.

   You can also create a policy by clicking + .

If a policy contains the **Back up copies on servers to be selected at backup job creation time** option, a server selection option is displayed to select one or more servers. The server selection option will list only the server where the selected database is on NetApp storage.

In the Configure schedules for selected policies section, the selected policies are listed.

b. In the Configure schedules for selected policies section, click + in the **Configure Schedules** column for the policy for which you want to configure the schedule.

c. In the Add schedules for policy *policy_name* dialog box, configure the schedule by specifying the start date, expiration date, and frequency, and then click **OK**.

You must do this for each frequency listed in the policy. The configured schedules are listed in the **Applied Schedules** column in the Configure schedules for selected policies section.

Third party backup schedules are not supported when they overlap with SnapCenter backup schedules.

7. In the Notification page, from the **Email preference** drop-down list, select the scenarios in which you want to send the emails.

You must also specify the sender and receiver email addresses, and the subject of the email. If you want to attach the report of the operation performed on the resource group, select **Attach Job Report**.

For email notification, you must have specified the SMTP server details either using the GUI or PowerShell command `Set-SmSmtpServer`.

The information regarding the parameters that can be used with the cmdlet and their descriptions can be obtained by running *Get-Help command_name*. Alternatively, you can also refer to the SnapCenter Software Cmdlet Reference Guide.

8. Review the summary, and then click **Finish**.

Create a storage system connection and a credential using PowerShell cmdlets for Exchange Server

You must create a storage virtual machine (SVM) connection and a credential before using PowerShell cmdlets to back up and restore.

Before you begin

- You should have prepared the PowerShell environment to execute the PowerShell cmdlets.
- You should have the required permissions in the Infrastructure Admin role to create storage connections.
- You should ensure that the plug-in installations are not in progress.

Host plug-in installations must not be in progress while adding a storage system connection because the host cache might not be updated and databases status might be displayed in the SnapCenter GUI as “Not available for backup” or “Not on NetApp storage”.

- Storage system names should be unique.

SnapCenter does not support multiple storage systems with the same name on different clusters. Each

storage system that is supported by SnapCenter should have a unique name and a unique data LIF IP address.

Steps

1. Initiate a PowerShell connection session by using the `Open-SmConnection` cmdlet.

   This example opens a PowerShell session:

   ```
   PS C:\> Open-SmConnection
   ```

2. Create a new connection to the storage system by using the `Add-SmStorageConnection` cmdlet.

   This example creates a new storage system connection:

   ```
   PS C:\> Add-SmStorageConnection -SVM test_vs1 -Protocol Https
   -Timeout 60
   ```

3. Create a new Run As account by using the `Add-Credential` cmdlet.

   This example creates a new Run As account named ExchangeAdmin with Windows credentials:

   ```
   PS C:> Add-SmCredential -Name ExchangeAdmin -AuthMode Windows
   -Credential sddev\administrator
   ```

The information regarding the parameters that can be used with the cmdlet and their descriptions can be obtained by running *Get-Help command_name*. Alternatively, you can also refer to the SnapCenter Software Cmdlet Reference Guide.

Back up Exchange databases

If a database is not part of any resource group, you can back up the database or Database Availability Group from the Resources page.

Before you begin

- You must have created a backup policy.
- You must have assigned the aggregate that is being used by the backup operation to the SVM used by the database.
- If you want to back up a resource that has a SnapMirror relationship with a secondary storage, the role assigned to the storage user should include the “snapmirror all” privilege. However, if you are using the “vsadmin” role, then the “snapmirror all” privilege is not required.
- If you want to perform backup of a Database or a Database Availability Group which has active/passive database copy on a NetApp and non-NetApp storage, and you have selected **Back up active copies** or **Back up copies on servers to be selected during backup job creation time** option in the policy, then the backup jobs will go in to warning state. The backup will succeed for active/passive database copy on NetApp storage and backup will fail for active/passive database copy on non-NetApp storage.

Best Practice: Do not run backups of active and passive databases at the same time. A race condition can occur and one of the backups might fail.

SnapCenter UI

Steps

1. In the left navigation pane, click **Resources**, and then select the **Microsoft Exchange Server plug-in** from the list.
2. In the Resources page, select either **Database**, or **Database Availability Group** from the **View** list.

   In the Resources page, the icon indicates that the database is on non-NetApp storage.

   > In a DAG, If an active database copy is on a non-NetApp storage and at least one passive database copy resides on a NetApp storage, then you can protect the database.

   Click , and then select the host name and the database type to filter the resources. You can then click to close the filter pane.

   - If you want to back up a database, click on the database name.
     i. If the Topology view is displayed, click **Protect**.
     ii. If the Database - Protect Resource wizard is displayed, continue to Step 3.
   - If you want to back up a Database Availability Group, click on the Database Availability Group name.
3. If you want to specify a custom Snapshot name, in the Resources page, select the **Use custom name format for Snapshot copy** check box, and then enter a custom name format that you want to use for the Snapshot name.

   For example, *customtext_policy_hostname* or *resource_hostname*. By default, a timestamp is appended to the Snapshot name.
4. In the Policies page, perform the following steps:

   a. Select one or more policies from the drop-down list.

   > You can also create a policy by clicking .

   > If a policy contains the **Back up copies on servers to be selected at backup job creation time** option, a server selection option is displayed to select one or more servers. The server selection option will list only the server where the selected database is on a NetApp storage.

   In the Configure schedules for selected policies section, the selected policies are listed.

   a. Click in the Configure Schedules column for the policy for which you want to configure a schedule.
   b. In the Add schedules for policy *policy_name* window, configure the schedule, and then click **OK**.

      Where, *policy_name* is the name of the policy that you have selected.

      The configured schedules are listed in the Applied Schedules column.

5. In the Notification page, from the **Email preference** drop-down list, select the scenarios in which you want to send the emails.

   You must also specify the sender and receiver email addresses, and the subject of the email. If you want to attach the report of the backup operation performed on the resource, select **Attach Job Report**.

   For email notification, you must have specified the SMTP server details using the either the GUI or the PowerShell command Set-SmSmtpServer.

6. Review the summary, and then click **Finish**.

   The database topology page is displayed.

7. Click **Back up Now**.
8. In the Backup page, perform the following steps:
   a. If you have applied multiple policies to the resource, from the **Policy** drop-down list, select the policy that you want to use for backup.

      If the policy selected for the on-demand backup is associated with a backup schedule, the on-demand backups will be retained based on the retention settings specified for the schedule type.

   b. Click **Backup**.
9. Monitor the backup's progress by double-clicking the job in the Activity pane at the bottom of the page to display the Job Details page.
   - In MetroCluster configurations, SnapCenter might not be able to detect a protection relationship after a failover.

     For information, see: Unable to detect SnapMirror or SnapVault relationship after MetroCluster failover

   - If you are backing up application data on VMDKs and the Java heap size for the SnapCenter Plug-in for VMware vSphere is not large enough, the backup might fail.

     To increase the Java heap size, locate the script file */opt/netapp/init_scripts/scvservice*. In that script, the *do_start method* command starts the SnapCenter VMware plug-in service. Update that command to the following: *Java -jar -Xmx8192M -Xms4096M*

PowerShell cmdlets

Steps

1. Initiate a connection session with the SnapCenter Server for a specified user by using the Open-SmConnection cmdlet.

```
Open-smconnection  -SMSbaseurl
https://snapctr.demo.netapp.com:8146/
```

   The username and password prompt is displayed.

2. Create a backup policy by using the Add-SmPolicy cmdlet.

This example creates a new backup policy with a full backup and log backup Exchange backup type:

```
C:\PS> Add-SmPolicy -PolicyName SCE_w2k12_Full_Log_bkp_Policy
-PolicyType Backup -PluginPolicytype SCE -SceBackupType
FullBackupAndLogBackup -BackupActiveCopies
```

This example creates a new backup policy with an hourly full backup and log backup Exchange backup type:

```
C:\PS> Add-SmPolicy -PolicyName SCE_w2k12_Hourly_Full_Log_bkp_Policy
-PolicyType Backup -PluginPolicytype SCE -SceBackupType
FullBackupAndLogBackup -BackupActiveCopies -ScheduleType Hourly
-RetentionSettings
@{'BackupType'='DATA';'ScheduleType'='Hourly';'RetentionCount'='10'}
```

This example creates a new backup policy to back up only Exchange logs:

```
Add-SmPolicy -PolicyName SCE_w2k12_Log_bkp_Policy -PolicyType Backup
-PluginPolicytype SCE -SceBackupType LogBackup -BackupActiveCopies
```

3. Discover host resources by using the Get-SmResources cmdlet.

   This example discovers the resources for the Microsoft Exchange Server plug-in on the specified host:

```
C:\PS> Get-SmResources -HostName vise-f6.sddev.mycompany.com
-PluginCode SCE
```

4. Add a new resource group to SnapCenter by using the Add-SmResourceGroup cmdlet.

   This example creates a new Exchange Server database backup resource group with the specified policy and resources:

```
C:\PS> Add-SmResourceGroup -ResourceGroupName SCE_w2k12_bkp_RG
-Description 'Backup ResourceGroup with Full and Log backup policy'
-PluginCode SCE -Policies
SCE_w2k12_Full_bkp_Policy,SCE_w2k12_Full_Log_bkp_Policy,SCE_w2k12_Lo
g_bkp_Policy -Resources @{'Host'='sce-w2k12-exch';'Type'='Exchange
Database';'Names'='sce-w2k12-exch.sceqa.com\sce-w2k12-exch_DB_1,sce-
w2k12-exch.sceqa.com\sce-w2k12-exch_DB_2'}
```

   This example creates a new Exchange Database Availability Group (DAG) backup resource group with the specified policy and resources:

```
Add-SmResourceGroup -ResourceGroupName SCE_w2k12_bkp_RG -Description
'Backup ResourceGroup with Full and Log backup policy' -PluginCode
SCE -Policies
SCE_w2k12_Full_bkp_Policy,SCE_w2k12_Full_Log_bkp_Policy,SCE_w2k12_Lo
g_bkp_Policy -Resources @{"Host"="DAGSCE0102";"Type"="Database
Availability Group";"Names"="DAGSCE0102"}
```

5. Initiate a new backup job by using the New-SmBackup cmdlet.

```
C:\PS> New-SmBackup -ResourceGroupName SCE_w2k12_bkp_RG -Policy
SCE_w2k12_Full_Log_bkp_Policy
```

This example creates a new backup to secondary storage:

```
New-SMBackup -DatasetName ResourceGroup1 -Policy
Secondary_Backup_Policy4
```

6. View the status of the backup job by using the Get-SmBackupReport cmdlet.

This example displays a job summary report of all jobs that were run on the specified date:

```
C:\PS> Get-SmJobSummaryReport -Date ?1/27/2018?
```

This example displays a job summary report for a specific job ID:

```
C:\PS> Get-SmJobSummaryReport -JobId 168
```

The information regarding the parameters that can be used with the cmdlet and their descriptions can be obtained by running *Get-Help command_name*. Alternatively, see SnapCenter Software Cmdlet Reference Guide.

Back up Exchange resources groups

A resource group is a collection of resources on a host or Exchange DAG, and the resource group can include either a whole DAG or individual databases. You can backup the resources groups from the Resources page.

Before you begin

- You must have created a resource group with a policy attached.
- You must have assigned the aggregate that is being used by the backup operation to the storage virtual machine (SVM) used by the database.

- If you want to back up a resource that has a SnapMirror relationship with a secondary storage, the role assigned to the storage user should include the “snapmirror all” privilege. However, if you are using the “vsadmin” role, then the “snapmirror all” privilege is not required.
- If a resource group has multiple databases from different hosts, the backup operation on some of the hosts might start late because of network issues. You should configure the value of `MaxRetryForUninitializedHosts` in `web.config` by using the `Set-SmConfigSettings` PowerShell cmdlet.
- In a resource group, if you include a Database or Database Availability Group which has active/passive database copy on a NetApp and non-NetApp storage, and you have selected **Back up active copies** or **Back up copies on servers to be selected during backup job creation time** option in the policy, then the backup jobs will go into warning state.

  The backup will succeed for active/passive database copy on NetApp storage and backup will fail for active/passive database copy on non-NetApp storage.

About this task

You can back up a resource group on demand from the Resources page. If a resource group has a policy attached and a schedule configured, then backups occur automatically according to the schedule.

Steps

1. In the left navigation pane, click **Resources**, and then select the **Microsoft Exchange Server plug-in** from the list.
2. In the Resources page, select **Resource Group** from the **View** list.

   You can search the resource group either by entering the resource group name in the search box or by clicking , and then selecting the tag. You can then click to close the filter pane.
3. In the Resource Groups page, select the resource group that you want to back up, and then click **Back up Now**.
4. In the Backup page, perform the following steps:
   a. If you have associated multiple policies with the resource group, from the **Policy** drop-down list, select the policy that you want to use for backup.

      If the policy selected for the on-demand backup is associated with a backup schedule, the on-demand backups will be retained based on the retention settings specified for the schedule type.

   b. Click **Backup**.
5. Monitor the backup’s progress by double-clicking the job in the Activity pane at the bottom of the page to display the Job Details page.

Monitor backup operations

You can monitor the progress of different backup operations by using the SnapCenterJobs page. You might want to check the progress to determine when it is complete or if there is an issue.

About this task

The following icons appear on the Jobs page and indicate the corresponding state of the operations:

-

In progress

- Completed successfully
- Failed
- Completed with warnings or could not start due to warnings
- Queued
- Canceled

Steps

1. In the left navigation pane, click **Monitor**.
2. In the Monitor page, click **Jobs**.
3. In the Jobs page, perform the following steps:
   a. Click to filter the list so that only backup operations are listed.
   b. Specify the start and end dates.
   c. From the **Type** drop-down list, select **Backup**.
   d. From the **Status** drop-down, select the backup status.
   e. Click **Apply** to view the operations completed successfully.
4. Select a backup job, and then click **Details** to view the job details.

   Though the backup job status displays , when you click on job details you might see that some of the child tasks of the backup operation are still in progress or marked with warning signs.

5. In the Job Details page, click **View logs**.

   The **View logs** button displays the detailed logs for the selected operation.

Monitor operations in the Activity pane

The Activity pane displays the five most recent operations performed. The Activity pane also displays when the operation was initiated and the status of the operation.

The Activity pane displays information regarding backup, restore, clone, and scheduled backup operations.

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. Click on the Activity pane to view the five most recent operations.

   When you click one of the operations, the operation details are listed in the **Job Details** page.

Cancel backup operations for Exchange database

You can cancel backup operations that are queued.

What you will need

- You must be logged in as the SnapCenter Admin or job owner to cancel operations.
- You can cancel a backup operation from either the **Monitor** page or the **Activity** pane.
- You cannot cancel a running backup operation.
- You can use the SnapCenter GUI, PowerShell cmdlets, or CLI commands to cancel the backup operations.
- The **Cancel Job** button is disabled for operations that cannot be canceled.
- If you selected **All members of this role can see and operate on other members objects** in Users\Groups page while creating a role, you can cancel the queued backup operations of other members while using that role.

Steps

1. Perform one of the following actions:

From the…	Action
Monitor page	a. In the left navigation pane, click **Monitor** > **Jobs**. b. Select the operation, and then click **Cancel Job**.
Activity pane	a. After initiating the backup operation, click on the Activity pane to view the five most recent operations. b. Select the operation. c. In the Job Details page, click **Cancel Job**.

The operation is canceled, and the resource is reverted to the previous state.

View Exchange backups in the Topology page

When you are preparing to back up a resource, you might find it helpful to view a graphical representation of all backups on the primary and secondary storages.

About this task

In the Topology page, you can see all of the backups that are available for the selected resource or resource group. You can view the details of those backups, and then select them to perform data protection operations.

You can review the following icon in the Manage Copies view to determine whether the backups are available on the primary or secondary storage (Mirror copies or Vault copies).

- displays the number of backups that are available on the primary storage.
- displays the number of backups that are mirrored on the secondary storage using SnapMirror technology.

- displays the number of backups that are replicated on the secondary storage using SnapVault technology.
  - The number of backups displayed includes the backups deleted from the secondary storage.

    For example, if you created 6 backups using a policy to retain only 4 backups, the number of backups displayed is 6.

    Best Practice: To ensure the correct number of replicated backups is displayed, we recommend that you refresh the topology.

If you have secondary relationship as SnapMirror active sync (initially released as SnapMirror Business Continuity [SM-BC]), you can see following additional icons:

- The replica site is up.
- The replica site is down.
- The secondary mirror or vault relationship has not been re-established.

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page, select the database, or the resource, or the resource group from the **View** drop-down list.
3. Select the resource either from the database details view or from the resource group details view.

   If the resource is protected, the Topology page of the selected resource is displayed.

4. Review the Summary card section to see a summary of the number of backups available on the primary and secondary storage.

   The Summary Card section displays the total number of backups and total number of log backups.

   Clicking the **Refresh** button starts a query of the storage to display an accurate count.

   If SnapLock enabled backup is taken, then clicking the **Refresh** button refreshes the primary and secondary SnapLock expiry time retrieved from ONTAP. A weekly schedule also refreshes the primary and secondary SnapLock expiry time retrieved from ONTAP.

   When the application resource is spread across multiple volumes, the SnapLock expiry time for the backup will be the longest SnapLock expiry time that is set for a Snapshot in a volume. The longest SnapLock expiry time is retrieved from ONTAP.

   For SnapMirror active sync, clicking the **Refresh** button refreshes the SnapCenter backup inventory by querying ONTAP for both primary and replica sites. A weekly schedule also performs this activity for all databases containing SnapMirror active sync relationship.

   - For SnapMirror active sync and only for ONTAP 9.14.1, Async Mirror or Async MirrorVault relationships to the new primary destination should be manually configured after failover. From ONTAP 9.15.1 onwards Async Mirror or Async MirrorVault is auto configured to the new primary destination.after

failover.

- After failover, a backup should be created for SnapCenter to be aware of the failover. You can click **Refresh** only after a backup has been created.

5. In the Manage Copies view, click **Backups** from the primary or secondary storage to see details of a backup.

   The details of the backups are displayed in a table format.

6. Select the backup from the table, and then click the data protection icons to perform restore, rename, and delete operations.

   You cannot rename or delete backups that are on the secondary storage. Deleting Snapshots is handled by ONTAP retention settings.

Restore Exchange resources

Restore workflow

You can use SnapCenter to restore Exchange databases by restoring one or more backups to your active file system.

The following workflow shows the sequence in which you must perform the Exchange database restore operations:

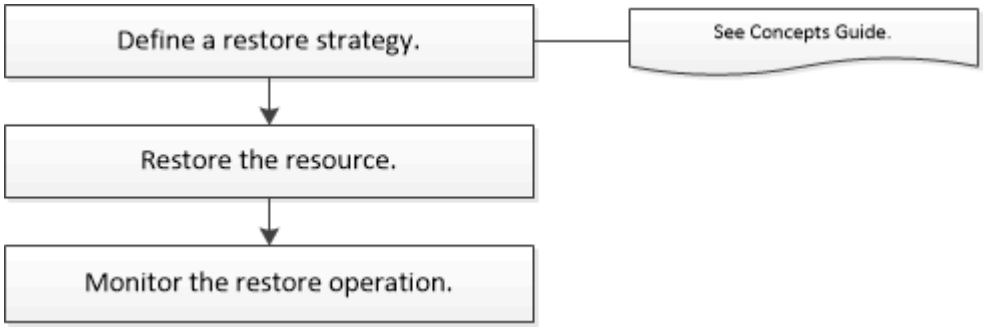

You can also use PowerShell cmdlets manually or in scripts to perform backup and restore operations. For detailed information about PowerShell cmdlets, use the SnapCenter cmdlet help or see SnapCenter Software Cmdlet Reference Guide.

Requirements for restoring an Exchange database

Before you restore an Exchange Server database from a SnapCenter Plug-in for Microsoft Exchange Server backup, you must ensure that several requirements are met.

To use the restore functionality completely, you must upgrade both SnapCenter Server and SnapCenter Plug-in for Exchange database to 4.6.

- The Exchange Server must be online and running before you can restore a database.
- The databases must exist on the Exchange Server.

Restoring deleted databases is not supported.

- SnapCenter schedules for the database must be suspended.
- The SnapCenter Server and the SnapCenter Plug-in for Microsoft Exchange Server host must be connected to the primary and secondary storage that contains the backups you want to restore.

Restore Exchange databases

You can use SnapCenter to restore backed-up Exchange databases.

Before you begin

- You must have backed up the resource groups, database, or Database Availability Groups (DAGs).
- When Exchange database is migrated to another location, restore operation does not work for old backups.
- If you are replicating Snapshots to a mirror or vault, the SnapCenter administrator must have assigned you the SVMs for both the source volumes and destination volumes.
- In a DAG, if an active database copy is on a non-NetApp storage and you want to restore from the passive database copy backup that is on a NetApp storage, make the passive copy (NetApp storage) as active copy, refresh the resources and perform the restore operation.

  Run the `Move-ActiveMailboxDatabase` command to make the passive database copy as active database copy.

  The Microsoft documentation contains information about this command.

About this task

- When restore operation is performed on a database, the database is mounted back on the same host and no new volume is created.
- DAG-level backups must be restored from individual databases.
- Full disk restore is not supported when files other than Exchange database (.edb) file exist.

  Plug-in for Exchange does not perform a full restore on a disk if the disk contains Exchange files such as those used for replication. When a full restore might impact Exchange functionality, Plug-in for Exchange performs a single file restore operation.

- Plug-in for Exchange cannot restore BitLocker encrypted drives.
- The SCRIPTS_PATH is defined using the PredefinedWindowsScriptsDirectory key located in the SMCoreServiceHost.exe.Config file of the plug-in host.

  If needed, you can change this path and restart SMcore service. It is recommended that you use the default path for security.

  The value of the key can be displayed from swagger through the API: API /4.7/configsettings

  You can use the GET API to display the value of the key. SET API is not supported.

- For ONTAP 9.12.1 and below version, the clones created from the SnapLock Vault Snapshots as part of restore will inherit the SnapLock Vault expiry time. Storage admin should manually cleanup the clones post the SnapLock expiry time.
- For SnapMirror active sync restore operation, you must select the backup from the primary location.

SnapCenter UI

Steps

1. On the left navigation pane, click **Resources** in the upper left corner of the Resource page.
2. Select the Exchange Server plug-in from the drop-down list.
3. In the Resources page, select **Database** from the View list.
4. Select the database from the list.
5. From the Manage Copies view, select **Backups**, from the Primary Backups table, and then click .
6. In the Options page, select one of the following log backup options:

Option	Description
All log backups	Choose **All log backups** to perform up-to-the-minute backup restore operation to restore all of the available log backups after the full backup.
By log backups until	Choose **By log backups until** to perform a point-in-time restore operation, which restores the database based on log backups until the selected log. The number of logs displayed in the drop-down list are based on UTM. For example, if full backup retention is 5 and UTM retention is 3, the number of log backups available are 5 but in the drop-down only 3 logs will be listed to perform restore operation.
By specific date until	Choose **By specific date until** to specify the date and time up to which transaction logs are applied to the restored database. This point-in-time restore operation restores transaction log entries that were recorded until the last backup on the specified date and time.
None	Choose **None** when you need to restore only the full backup without any log backups.

You can perform one of the following actions:

- **Recover and mount database after restore** - This option is selected by default.
- **Do not verify the integrity of transaction logs in the backup before restore** - By default, SnapCenter verifies the integrity of transaction logs in a backup before performing a restore operation.

Best Practice: You should not select this option.

7. In the Script page, enter the path and the arguments of the prescript or postscript that should be run before or after the restore operation, respectively.

   Restore prescript arguments include $Database and $ServerInstance.

   Restore postscript arguments include $Database, $ServerInstance, $BackupName, $LogDirectory, and $TargetServerInstance.

   You can run a script to update SNMP traps, automate alerts, send logs, and so on.

   The prescripts or postscripts path should not include drives or shares. The path should be relative to the SCRIPTS_PATH.

8. In the Notification page, from the **Email preference** drop-down list, select the scenarios in which you want to send the emails.

   You must also specify the sender and receiver email addresses, and the subject of the email.

9. Review the summary, and then click **Finish**.
10. You can view the status of the restore job by expanding the Activity panel at the bottom of the page.

    You should monitor the restore process by using the **Monitor** > **Jobs** page.

When you restore an active database from a backup, the passive database might go into suspended or failed state if there is a lag between the replica and the active database.

The state change can occur when the active database's log chain forks and begins a new branch which breaks replication. Exchange Server attempts to fix the replica, but if it is unable to do so, after restore, you should create a fresh backup, and then reseed the replica.

PowerShell cmdlets

Steps

1. Initiate a connection session with the SnapCenter Server for a specified user by using the `Open-SmConnection` cmdlet.

```
Open-smconnection  -SMSbaseurl
https://snapctr.demo.netapp.com:8146/
```

2. Retrieve the information about the one or more backups that you want to restore by using the `Get-SmBackup` cmdlet.

   This example displays information about all available backups:

```
PS C:\> Get-SmBackup

BackupId                        BackupName
BackupTime                      BackupType
--------                        ----------
----------                      ----------
341                             ResourceGroup_36304978_UTM...
12/8/2017 4:13:24 PM            Full Backup
342                             ResourceGroup_36304978_UTM...
12/8/2017 4:16:23 PM            Full Backup
355                             ResourceGroup_06140588_UTM...
12/8/2017 6:32:36 PM            Log Backup
356                             ResourceGroup_06140588_UTM...
12/8/2017 6:36:20 PM            Full Backup
```

3. Restore data from the backup by using the `Restore-SmBackup` cmdlet.

   This example restores an up-to-the-minute backup:

```
C:\PS> Restore-SmBackup -PluginCode SCE -AppObjectId 'sce-w2k12-
exch.sceqa.com\sce-w2k12-exch_DB_2' -BackupId 341
-IsRecoverMount:$true
```

   This example restores a point-in-time backup:

```
C:\ PS> Restore-SmBackup -PluginCode SCE -AppObjectId 'sce-w2k12-
exch.sceqa.com\sce-w2k12-exch_DB_2' -BackupId 341
-IsRecoverMount:$true -LogRestoreType ByTransactionLogs -LogCount 2
```

   This example restores a backup on secondary storage to primary story:

```
C:\ PS> Restore-SmBackup -PluginCode 'SCE' -AppObjectId 'DB2'
-BackupId 81 -IsRecoverMount:$true -Confirm:$false
-archive @{Primary="paw_vs:vol1";Secondary="paw_vs:vol1_mirror"}
-logrestoretype All
```

   The `-archive` parameter enables you to specify the primary and secondary volumes you want to use for the restore.

   The `-IsRecoverMount:$true` parameter enables you to mount the database after the restore.

The information regarding the parameters that can be used with the cmdlet and their descriptions can be obtained by running *Get-Help command_name*. Alternatively, you can also refer to the SnapCenter

Software Cmdlet Reference Guide.

Granular recovery of mails and mailbox

Single Mailbox Recovery (SMBR) software allows you to restore and recover mails or mailbox instead of the complete Exchange Database.

Restoring complete database just to recover a single mail will consume lot of time and resource. SMBR helps in quickly recovering the mails by creating clone copy of the Snapshot and then using Microsoft API's to mount the mailbox in SMBR.
For information on how to use SMBR, see SMBR Administration Guide.

For additional information on SMBR, refer the following:

- How to manually restore a single item with SMBR (also applicable for Ontrack Power Control restores)
- How to restore from secondary storage in SMBR with SnapCenter
- Recovering Microsoft Exchange Mail From SnapVault Using SMBR

Restore an Exchange Server database from secondary storage

You can restore a backed up Exchange Server database from secondary storage (mirror or vault).

You must have replicated the Snapshots from primary storage to a secondary storage.

About this task

- For ONTAP 9.12.1 and below version, the clones created from the SnapLock Vault Snapshots as part of restore will inherit the SnapLock Vault expiry time. Storage admin should manually cleanup the clones post the SnapLock expiry time.
- For SnapMirror active sync restore operation, you must select the backup from the primary location.

Steps

1. In the left navigation pane, click **Resources**, and then select **Microsoft Exchange Server plug-in** from the list.
2. In the Resources page, select **Database** or **Resource Group** from the **View** drop-down list.
3. Select the database or the resource group.

   The database or resource group topology page is displayed.
4. In the Manage Copies section, select **Backups** from the secondary storage system (mirror or vault).
5. Select the backup from the list, and then click .
6. In the Location page, choose the destination volume for restoring the selected resource.
7. Complete the Restore wizard, review the summary, and then click **Finish**.

Reseed a passive Exchange node replica

If you need to reseed a replica copy, for instance when a copy is corrupt, you can reseed

to the latest backup using the reseed feature in SnapCenter.

Before you begin

You must have created a backup of the database you want to reseed.

+
To avoid lagging between nodes, you can either create a new backup before you perform a reseed operation, or choose the host with the latest backup.

Steps

1. In the left navigation pane, click **Resources**, and then select **Microsoft Exchange Server plug-in** from the list.
2. In the Resources page, select the appropriate option from the View list:

Option	Description
To reseed a single database	Select **Database** from the View list.
To reseed databases in a DAG	Select **Database Availability Group** from the View list.

3. Select the resource you want to reseed.
4. In the Manage Copies page, click **Reseed**.
5. From the list of unhealthy databases copies in the Reseed wizard, select the one you want to reseed, and then click **Next**.
6. In the Host window, select the host with the backup from which you want to reseed, and then click **Next**.
7. In the Notification page, from the **Email preference** drop-down list, select the scenarios in which you want to send the emails.

   You must also specify the sender and receiver email addresses, and the subject of the email.

8. Review the summary, and then click **Finish**.
9. You can view the status of the job by expanding the Activity panel at the bottom of the page.

Reseed operation is not supported if the passive database copy resides on non-NetApp storage.

Reseed a replica using PowerShell cmdlets for Exchange database

You can use PowerShell cmdlets to restore an unhealthy replica by using either the most recent copy on the same host or the most recent copy from an alternate host.

The information regarding the parameters that can be used with the cmdlet and their descriptions can be obtained by running *Get-Help command_name*. Alternatively, you can also refer to the SnapCenter Software Cmdlet Reference Guide.

Steps

1. Initiate a connection session with the SnapCenter Server for a specified user by using the `Open-`

`SmConnection` cmdlet.

```
Open-smconnection  -SMSbaseurl  https:\\snapctr.demo.netapp.com:8146/
```

2. Reseed the database by using the `reseed-SmDagReplicaCopy` cmdlet.

   This example reseeds the failed copy of the database called execdb on the host "mva-rx200.netapp.com" using the latest backup on that host.

```
reseed-SmDagReplicaCopy -ReplicaHost "mva-rx200.netapp.com" -Database
execdb
```

   This example reseeds the failed copy of the database called execdb using the latest backup of the database (production/copy) on an alternate host "mva-rx201.netapp.com."

```
reseed-SmDagReplicaCopy -ReplicaHost "mva-rx200.netapp.com" -Database
execdb -BackupHost "mva-rx201.netapp.com"
```

Monitor restore operations

You can monitor the progress of different SnapCenter restore operations by using the Jobs page. You might want to check the progress of an operation to determine when it is complete or if there is an issue.

About this task

Post-restore states describe the conditions of the resource after a restore operation and any further restore actions that you can take.

The following icons appear on the Jobs page, and indicate the state of the operation:

- In progress
- Completed successfully
- Failed
- Completed with warnings or could not start due to warnings
- Queued
- Canceled

Steps

1. In the left navigation pane, click **Monitor**.
2. In the **Monitor** page, click **Jobs**.
3. In the **Jobs** page, perform the following steps:
   a. Click to filter the list so that only restore operations are listed.

b. Specify the start and end dates.
c. From the **Type** drop-down list, select **Restore**.
d. From the **Status** drop-down list, select the restore status.
e. Click **Apply** to view the operations that have been completed successfully.

4. Select the restore job, and then click **Details** to view the job details.
5. In the **Job Details** page, click **View logs**.

   The **View logs** button displays the detailed logs for the selected operation.

Cancel restore operations for Exchange database

You can cancel restore jobs that are queued.

You should be logged in as the SnapCenter Admin or job owner to cancel restore operations.

About this task

- You can cancel a queued restore operation from either the **Monitor** page or the **Activity** pane.
- You cannot cancel a running restore operation.
- You can use the SnapCenter GUI, PowerShell cmdlets, or CLI commands to cancel the queued restore operations.
- The **Cancel Job** button is disabled for restore operations that cannot be canceled.
- If you selected **All members of this role can see and operate on other members objects** in Users\Groups page while creating a role, you can cancel the queued restore operations of other members while using that role.

Step

Perform one of the following actions:

From the…	Action
Monitor page	a. In the left navigation pane, click **Monitor** > **Jobs**. b. Select the job and click **Cancel Job**.
Activity pane	a. After initiating the restore operation, click on the Activity pane to view the five most recent operations. b. Select the operation. c. In the Job Details page, click **Cancel Job**.

Protect IBM Db2

SnapCenter Plug-in for IBM Db2

SnapCenter Plug-in for IBM Db2 overview

The SnapCenter Plug-in for IBM Db2 Database is a host-side component of the NetApp SnapCenter software that enables application-aware data protection management of IBM Db2 databases. The Plug-in for IBM Db2 Database automates the backup, restore, and cloning of IBM Db2 databases in your SnapCenter environment.

- SnapCenter 6.0 supports IBM Db2 10.5 and later.
- SnapCenter 6.0.1 supports IBM Db2 9.7.x and later. Also, beginning with SnapCenter 6.0.1, IBM Db2 on AIX is supported.

SnapCenter supports single instance and multi instance Db2 setups. You can use the Plug-in for IBM Db2 Database in both Linux and Windows environments. In Windows environments, Db2 will be supported as manual resource.

Db2 pureScale environment and Db2 multi node (DPF) systems are not supported.

When the Plug-in for IBM Db2 Database is installed, you can use SnapCenter with NetApp SnapMirror technology to create mirror copies of backup sets on another volume. You can also use the plug-in with NetApp SnapVault technology to perform disk-to-disk backup replication for standards compliance.

SnapCenter Plug-in for Db2 supports NFS and SAN on ONTAP and Azure NetApp File storage layouts.

VMDK, vVol, and RDM virtual storage layout is supported.

What you can do using the SnapCenter Plug-in for IBM Db2

When you install the Plug-in for IBM Db2 Database in your environment, you can use SnapCenter to back up, restore, and clone IBM Db2 databases and their resources. You can also perform tasks supporting those operations.

- Add databases.
- Create backups.
- Restore from backups.
- Clone backups.
- Schedule backup operations.
- Monitor backup, restore, and clone operations.
- View reports for backup, restore, and clone operations.

SnapCenter Plug-in for IBM Db2 features

SnapCenter integrates with the plug-in application and with NetApp technologies on the storage system. To work with the Plug-in for IBM Db2 Database, you use the SnapCenter

graphical user interface.

- **Unified graphical user interface**

  The SnapCenter interface provides standardization and consistency across plug-ins and environments. The SnapCenter interface enables you to complete consistent backup, restore, and clone operations across plug-ins, use centralized reporting, use at-a-glance dashboard views, set up role-based access control (RBAC), and monitor jobs across all plug-ins.

- **Automated central administration**

  You can schedule backup operations, configure policy-based backup retention, and perform restore operations. You can also proactively monitor your environment by configuring SnapCenter to send email alerts.

- **Nondisruptive NetApp snapshot copy technology**

  SnapCenter uses NetApp snapshot technology with the Plug-in for IBM Db2 Database to back up resources.

Using the Plug-in for IBM Db2 also offers the following benefits:

- Support for backup, restore, and clone workflows
- RBAC-supported security and centralized role delegation

  You can also set the credentials so that the authorized SnapCenter users have application-level permissions.

- Creation of space-efficient and point-in-time copies of resources for testing or data extraction by using NetApp FlexClone technology

  A FlexClone license is required on the storage system where you want to create the clone.

- Support for the consistency group (CG) snapshot feature of ONTAP as part of creating backups.
- Capability to run multiple backups simultaneously across multiple resource hosts

  In a single operation, snapshots are consolidated when resources in a single host share the same volume.

- Capability to create snapshots using external commands.
- Support for Linux LVM on XFS file system.

Storage types supported by SnapCenter Plug-in for IBM Db2

SnapCenter supports a wide range of storage types on both physical machines and virtual machines (VMs). You must verify the support for your storage type before installing SnapCenter Plug-in for IBM Db2.

Machine	Storage type
Physical server	• FC-connected LUNs • iSCSI-connected LUNs • NFS-connected volumes
VMware ESXi	• RDM LUNs connected by an FC or iSCSI ESXi HBAScanning of host bus adapters (HBAs) might take long time to complete because SnapCenter scans all the host bus adaptors present in the host. You can edit the **LinuxConfig.pm** file located at */opt/NetApp/snapcenter/spl/plugins/scu/scucore/modules/SCU/Config* to set the value of the **SCSI_HOSTS_OPTIMIZED_RESCAN** parameter to 1 to rescan only those HBAs that are listed in HBA_DRIVER_NAMES. • iSCSI LUNs connected directly to the guest system by the iSCSI initiator • VMDKs on NFS datastores • VMDKs on VMFS created • NFS volumes connected directly to the guest system • vVol datastores on both NFS and SAN vVol datastore can only be provisioned with ONTAP Tools for VMware vSphere.

Minimum ONTAP privileges required for IBM Db2 plug-in

The minimum ONTAP privileges that are required vary according to the SnapCenter plug-ins you are using for data protection.

- All-access commands: Minimum privileges required for ONTAP 9.12.1 and later
  - event generate-autosupport-log
  - job history show
  - job stop
  - lun
  - lun create
  - lun create
  - lun create
  - lun delete
  - lun igroup add

- lun igroup create
- lun igroup delete
- lun igroup rename
- lun igroup rename
- lun igroup show
- lun mapping add-reporting-nodes
- lun mapping create
- lun mapping delete
- lun mapping remove-reporting-nodes
- lun mapping show
- lun modify
- lun move-in-volume
- lun offline
- lun online
- lun persistent-reservation clear
- lun resize
- lun serial
- lun show
- snapmirror policy add-rule
- snapmirror policy modify-rule
- snapmirror policy remove-rule
- snapmirror policy show
- snapmirror restore
- snapmirror show
- snapmirror show-history
- snapmirror update
- snapmirror update-ls-set
- snapmirror list-destinations
- version
- volume clone create
- volume clone show
- volume clone split start
- volume clone split stop
- volume create
- volume destroy
- volume file clone create
- volume file show-disk-usage

    - volume offline
    - volume online
    - volume modify
    - volume qtree create
    - volume qtree delete
    - volume qtree modify
    - volume qtree show
    - volume restrict
    - volume show
    - volume snapshot create
    - volume snapshot delete
    - volume snapshot modify
    - volume snapshot modify-snaplock-expiry-time
    - volume snapshot rename
    - volume snapshot restore
    - volume snapshot restore-file
    - volume snapshot show
    - volume unmount
    - vserver cifs
    - vserver cifs share create
    - vserver cifs share delete
    - vserver cifs shadowcopy show
    - vserver cifs share show
    - vserver cifs show
    - vserver export-policy
    - vserver export-policy create
    - vserver export-policy delete
    - vserver export-policy rule create
    - vserver export-policy rule show
    - vserver export-policy show
    - vserver iscsi
    - vserver iscsi connection show
    - vserver show
- Read-only commands: Minimum privileges required for ONTAP 8.3.0 and later
    - network interface
    - network interface show
    - vserver

Prepare storage systems for SnapMirror and SnapVault replication for IBM Db2

You can use a SnapCenter plug-in with ONTAP SnapMirror technology to create mirror copies of backup sets on another volume, and with ONTAP SnapVault technology to perform disk-to-disk backup replication for standards compliance and other governance-related purposes. Before you perform these tasks, you must configure a data-protection relationship between the source and destination volumes and initialize the relationship.

SnapCenter performs the updates to SnapMirror and SnapVault after it completes the Snapshot operation. SnapMirror and SnapVault updates are performed as part of the SnapCenter job; do not create a separate ONTAP schedule.

If you are coming to SnapCenter from a NetApp SnapManager product and are satisfied with the data protection relationships you have configured, you can skip this section.

A data protection relationship replicates data on primary storage (the source volume) to secondary storage (the destination volume). When you initialize the relationship, ONTAP transfers the data blocks referenced on the source volume to the destination volume.

SnapCenter does not support cascade relationships between SnapMirror and SnapVault volumes (**Primary** > **Mirror** > **Vault**). You should use fanout relationships.

SnapCenter supports the management of version-flexible SnapMirror relationships. For details about version-flexible SnapMirror relationships and how to set them up, see the ONTAP documentation.

Backup strategy for IBM Db2

Define a backup strategy for IBM Db2

Defining a backup strategy before you create your backup jobs helps you to have the backups that you require to successfully restore or clone your resources. Your service-level agreement (SLA), recovery time objective (RTO), and recovery point objective (RPO) largely determine your backup strategy.

About this task

An SLA defines the level of service that is expected and addresses many service-related issues, including the availability and performance of the service. RTO is the time by which a business process must be restored after a disruption in service. RPO defines the strategy for the age of the files that must be recovered from backup storage for regular operations to resume after a failure. SLA, RTO, and RPO contribute to the data protection strategy.

Steps

1. Determine when you should back up your resources.
2. Decide how many backup jobs you require.
3. Decide how to name your backups.
4. Decide whether you want to create a Snapshot copy-based policy to back up application-consistent snapshots of the database.
5. Decide whether you want to use NetApp SnapMirror technology for replication or NetApp SnapVault technology for long-term retention.

6. Determine the retention period for the snapshots on the source storage system and the SnapMirror destination.
7. Determine whether you want to run any commands before or after the backup operation and provide a prescript or postscript.

Automatic discovery of resources on Linux host

Resources are IBM Db2 databases and instances on the Linux host that are managed by SnapCenter. After installing the SnapCenter Plug-in for IBM Db2 plug-in, the IBM Db2 databases from all the instances on that Linux host are automatically discovered and displayed in the Resources page.

Type of backups supported

Backup type specifies the type of backup that you want to create. SnapCenter supports snapshot copy-based backup type for IBM Db2 databases.

Snapshot copy based backup

Snapshot copy-based backups leverage NetApp snapshot technology to create online, read-only copies of the volumes on which the IBM Db2 databases reside.

How SnapCenter Plug-in for IBM Db2 uses consistency group snapshots

You can use the plug-in to create consistency group snapshots for resource groups. A consistency group is a container that can house multiple volumes so that you can manage them as one entity. A consistency group is simultaneous snapshots of multiple volumes, providing consistent copies of a group of volumes.

You can also specify the wait time for the storage controller to consistently group snapshots. The available wait time options are **Urgent**, **Medium**, and **Relaxed**. You can also enable or disable Write Anywhere File Layout (WAFL) sync during consistent group snapshot operation. WAFL sync improves the performance of a consistency group snapshot.

How SnapCenter manages housekeeping of data backups

SnapCenter manages the housekeeping of data backups on the storage system and file system levels.

The snapshots on the primary or secondary storage and their corresponding entries in the IBM Db2 catalog are deleted based on the retention settings.

Considerations for determining backup schedules for IBM Db2

The most critical factor in determining a backup schedule is the rate of change for the resource. You might back up a heavily used resource every hour, while you might back up a rarely used resource once a day. Other factors include the importance of the resource to your organization, your service-level agreement (SLA) and your recovery point objective (RPO).

Backup schedules have two parts, as follows:

- Backup frequency (how often backups are to be performed)

  Backup frequency, also called schedule type for some plug-ins, is part of a policy configuration. For example, you might configure the backup frequency as hourly, daily, weekly, or monthly.

- Backup schedules (exactly when backups are to be performed)

  Backup schedules are part of a resource or resource group configuration. For example, if you have a resource group that has a policy configured for weekly backups, you might configure the schedule to back up every Thursday at 10:00 p.m.

Number of backup jobs needed for IBM Db2

Factors that determine the number of backup jobs that you need include the size of the resource, the number of volumes used, the rate of change of the resource, and your Service Level Agreement (SLA).

Backup naming conventions for Plug-in for IBM Db2 databases

You can either use the default Snapshot naming convention or use a customized naming convention. The default backup naming convention adds a timestamp to Snapshot names that helps you identify when the copies were created.

The Snapshot uses the following default naming convention:

```
resourcegroupname_hostname_timestamp
```

You should name your backup resource groups logically, as in the following example:

```
dts1_mach1x88_03-12-2015_23.17.26
```

In this example, the syntax elements have the following meanings:

- *dts1* is the resource group name.
- *mach1x88* is the host name.
- *03-12-2015_23.17.26* is the date and timestamp.

Alternatively, you can specify the Snapshot name format while protecting resources or resource groups by selecting **Use custom name format for Snapshot copy**. For example, customtext_resourcegroup_policy_hostname or resourcegroup_hostname. By default, the time stamp suffix is added to the Snapshot name.

Restore and recovery strategy for IBM Db2

Define a restore and recovery strategy for IBM Db2 resources

You must define a strategy before you restore and recover your database so that you can perform restore and recovery operations successfully.

Only manual recovery of database is supported.

Steps

1. Determine the restore strategies supported for manually added IBM Db2 resources
2. Determine the restore strategies supported for auto discovered IBM Db2 databases

Types of restore strategies supported for manually added IBM Db2 resources

You must define a strategy before you can successfully perform restore operations using SnapCenter. There are two types of restore strategies for manually added IBM Db2 resources.

You cannot recover manually added IBM Db2 resources.

Complete resource restore

- Restores all volumes, qtrees, and LUNs of a resource

If the resource contains volumes or qtrees, the snapshots taken after the snapshot selected for restore on such volumes or qtrees are deleted and cannot be recovered. Also, if any other resource is hosted on the same volumes or qtrees, then that resource is also deleted.

Type of restore strategy supported for automatically discovered IBM Db2

You must define a strategy before you can successfully perform restore operations using SnapCenter.

Complete resource restore is the restore strategy supported for automatically discovered IBM Db2 databases. This restores all the volumes, qtrees, and LUNs of a resource.

Types of restore operations for auto discovered IBM Db2

SnapCenter Plug-in for IBM Db2 supports Single File SnapRestore, and connect-and-copy restore types for automatically discovered IBM Db2 databases.

Single File SnapRestore is performed in NFS environments for the following scenarios:

- If only the **Complete Resource** option is selected
- When the backup selected is from a SnapMirror or SnapVault secondary location, and the **Complete Resource** option is selected

Single File SnapRestore is performed in SAN environments for the following scenarios:

- If only the **Complete Resource** option is selected
- When the backup is selected from a SnapMirror or SnapVault secondary location, and the **Complete Resource** option is selected

Prepare to install the SnapCenter Plug-in for IBM Db2

Installation workflow of SnapCenter Plug-in for IBM Db2

You should install and set up the SnapCenter Plug-in for IBM Db2 if you want to protect IBM Db2 databases.

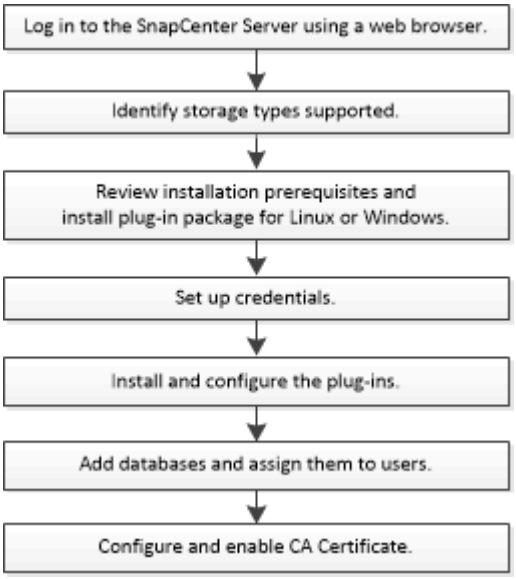

Prerequisites for adding hosts and installing Plug-ins package for Windows, Linux, or AIX

Before you add a host and install the plug-in packages, you must complete all the requirements. SnapCenter Plug-in for IBM Db2 supported on Windows, Linux, and AIX environments.

- You must have installed Java 11 on your host.

  IBM Java is not supported on Windows and Linux host.

- For Windows, plug-in Creator Service should be running using the “LocalSystem” windows user, which is the default behavior when Plug-in for IBM Db2 is installed as domain administrator.
- When installing a plug-in on a Windows host, if you specify a credential that is not built-in or if the user belongs to a local workgroup user, you must disable UAC on the host. SnapCenter Plug-in for Microsoft Windows will be deployed by default with the IBM Db2 plug-in on Windows hosts.
- SnapCenter Server should have access to the 8145 or custom port of Plug-in for IBM Db2 host.

Windows hosts

- You must have a domain user with local administrator privileges with local login permissions on the remote host.
- While installing Plug-in for IBM Db2 on a Windows host, SnapCenter Plug-in for Microsoft Windows is installed automatically.
- You must have enabled the password-based SSH connection for the root or non-root user.

- You must have installed Java 11 on your Windows host.

  Download JAVA for Windows

  NetApp Interoperability Matrix Tool

Linux and AIX hosts

- You must have enabled the password-based SSH connection for the root or non-root user.
- You must have installed Java 11 on your Linux host.

  Download JAVA for Linux

  Download JAVA for AIX

  NetApp Interoperability Matrix Tool

- For IBM Db2 databases that are running on a Linux host, while installing Plug-in for IBM Db2, SnapCenter Plug-in for UNIX is installed automatically.
- You should have **bash** as the default shell for plug-in installation.

Supplemental commands

To run a supplemental command on the SnapCenter Plug-in for IBM Db2, you must include it in the *allowed_commands.config* file.

- Default location on the Windows host: *C:\Program Files\NetApp\SnapCenter\Snapcenter Plug-in Creator\etc\allowed_commands.config*
- Default location on the Linux host: */opt/NetApp/snapcenter/scc/etc/allowed_commands.config*

To allow supplemental commands on the plug-in host, open *allowed_commands.config* file in an editor. Enter each command on a separate line and the commands are not case sensitive. Ensure that you specify the fully qualified pathname and enclose the pathname in quotation marks (") if it contains spaces.

For example:

command: mount

command: umount

command: "C:\Program Files\NetApp\SnapCreator commands\sdcli.exe"

command: myscript.bat

If the *allowed_commands.config* file is not present, the commands or script execution will be blocked and the workflow will fail with the following error:

"[/mnt/mount -a] execution not allowed. Authorize by adding the command in the file %s on the plugin host."

If the command or script is not present in the *allowed_commands.config*, the command or script execution will be blocked and the workflow will fail with the following error:

"[/mnt/mount -a] execution not allowed. Authorize by adding the command in the file %s on the plugin host."

You should not use a wildcard entry (*) to allow all commands.

Configure sudo privileges for non-root users for Linux host

SnapCenter allows a non-root user to install the SnapCenter Plug-ins Package for Linux and to start the plug-in process. The plug-in processes will be running as an effective non-root user. You should configure sudo privileges for the non-root user to provide access to several paths.

What you will need

- Sudo version 1.8.7 or later.
- If the umask is 0027, ensure that the java folder and all the files inside should have permission of 555. Otherwise plug-in installation might fail.
- For the non-root user, ensure that the name of the non-root user and the user's group should be the same.
- Edit the */etc/ssh/sshd_config* file to configure the message authentication code algorithms: MACs hmac-sha2-256 and MACs hmac-sha2-512.

  Restart the sshd service after updating the configuration file.

  Example:

```
#Port 22
#AddressFamily any
#ListenAddress 0.0.0.0
#ListenAddress ::
#Legacy changes
#KexAlgorithms diffie-hellman-group1-sha1
#Ciphers aes128-cbc
#The default requires explicit activation of protocol
Protocol 2
HostKey/etc/ssh/ssh_host_rsa_key
MACs hmac-sha2-256
```

About this task

You should configure sudo privileges for the non-root user to provide access to the following paths:

- /home/*LINUX_USER*/.sc_netapp/snapcenter_linux_host_plugin.bin
- /custom_location/NetApp/snapcenter/spl/installation/plugins/uninstall
- /custom_location/NetApp/snapcenter/spl/bin/spl

Steps

1. Log in to the Linux host on which you want to install the SnapCenter Plug-ins Package for Linux.
2. Add the following lines to the /etc/sudoers file by using the visudo Linux utility.

```
Cmnd_Alias HPPLCMD = sha224:checksum_value== /home/
LINUX_USER/.sc_netapp/snapcenter_linux_host_plugin.bin,
/opt/NetApp/snapcenter/spl/installation/plugins/uninstall,
/opt/NetApp/snapcenter/spl/bin/spl, /opt/NetApp/snapcenter/scc/bin/scc
Cmnd_Alias PRECHECKCMD = sha224:checksum_value== /home/
LINUX_USER/.sc_netapp/Linux_Prechecks.sh
Cmnd_Alias CONFIGCHECKCMD = sha224:checksum_value==
/opt/NetApp/snapcenter/spl/plugins/scu/scucore/configurationcheck/Config
_Check.sh
Cmnd_Alias SCCMD = sha224:checksum_value==
/opt/NetApp/snapcenter/spl/bin/sc_command_executor
Cmnd_Alias SCCCMDEXECUTOR =checksum_value==
/opt/NetApp/snapcenter/scc/bin/sccCommandExecutor
LINUX_USER ALL=(ALL) NOPASSWD:SETENV: HPPLCMD, PRECHECKCMD,
CONFIGCHECKCMD, SCCCMDEXECUTOR, SCCMD
Defaults: LINUX_USER env_keep += "IATEMPDIR"
Defaults: LINUX_USER env_keep += "JAVA_HOME"
Defaults: LINUX_USER !visiblepw
Defaults: LINUX_USER !requiretty
```

If you are having a RAC setup, along with the other allowed commands, you should add the following to the /etc/sudoers file: '/<crs_home>/bin/olsnodes'

You can obtain the value of *crs_home* from the */etc/oracle/olr.loc* file.

LINUX_USER is the name of the non-root user that you created.

You can obtain the *checksum_value* from the **sc_unix_plugins_checksum.txt** file, which is located at:

- *C:\ProgramData\NetApp\SnapCenter\Package Repository\sc_unix_plugins_checksum.txt* if SnapCenter Server is installed on Windows host.
- */opt/NetApp/snapcenter/SnapManagerWeb/Repository/sc_unix_plugins_checksum.txt* if SnapCenter Server in installed on Linux host.

The example should be used only as a reference for creating your own data.

Configure sudo privileges for non-root users for AIX host

SnapCenter 4.4 and later allows a non-root user to install the SnapCenter Plug-ins Package for AIX and to start the plug-in process. The plug-in processes will be running as an effective non-root user. You should configure sudo privileges for the non-root user to provide access to several paths.

What you will need

- Sudo version 1.8.7 or later.
- If the umask is 0027, ensure that the java folder and all the files inside should have permission of 555. Otherwise plug-in installation might fail.

- Edit the */etc/ssh/sshd_config* file to configure the message authentication code algorithms: MACs hmac-sha2-256 and MACs hmac-sha2-512.

  Restart the sshd service after updating the configuration file.

  Example:

```
#Port 22
#AddressFamily any
#ListenAddress 0.0.0.0
#ListenAddress ::
#Legacy changes
#KexAlgorithms diffie-hellman-group1-sha1
#Ciphers aes128-cbc
#The default requires explicit activation of protocol
Protocol 2
HostKey/etc/ssh/ssh_host_rsa_key
MACs hmac-sha2-256
```

About this task

You should configure sudo privileges for the non-root user to provide access to the following paths:

- /home/*AIX_USER*/.sc_netapp/snapcenter_aix_host_plugin.bsx
- /custom_location/NetApp/snapcenter/spl/installation/plugins/uninstall
- /custom_location/NetApp/snapcenter/spl/bin/spl

Steps

1. Log in to the AIX host on which you want to install the SnapCenter Plug-ins Package for AIX.
2. Add the following lines to the /etc/sudoers file by using the visudo Linux utility.

```
Cmnd_Alias HPPACMD = sha224:checksum_value== /home/
AIX_USER/.sc_netapp/snapcenter_aix_host_plugin.bsx,
/opt/NetApp/snapcenter/spl/installation/plugins/uninstall,
/opt/NetApp/snapcenter/spl/bin/spl
Cmnd_Alias PRECHECKCMD = sha224:checksum_value== /home/
AIX_USER/.sc_netapp/AIX_Prechecks.sh
Cmnd_Alias CONFIGCHECKCMD = sha224:checksum_value==
/opt/NetApp/snapcenter/spl/plugins/scu/scucore/configurationcheck/Config
_Check.sh
Cmnd_Alias SCCMD = sha224:checksum_value==
/opt/NetApp/snapcenter/spl/bin/sc_command_executor
AIX_USER ALL=(ALL) NOPASSWD:SETENV: HPPACMD, PRECHECKCMD,
CONFIGCHECKCMD, SCCMD
Defaults: LINUX_USER env_keep += "IATEMPDIR"
Defaults: LINUX_USER env_keep += "JAVA_HOME"
Defaults: AIX_USER !visiblepw
Defaults: AIX_USER !requiretty
```

If you are having a RAC setup, along with the other allowed commands, you should add the following to the /etc/sudoers file: '/<crs_home>/bin/olsnodes'

You can obtain the value of *crs_home* from the */etc/oracle/olr.loc* file.

AIX_USER is the name of the non-root user that you created.

You can obtain the *checksum_value* from the **sc_unix_plugins_checksum.txt** file, which is located at:

- *C:\ProgramData\NetApp\SnapCenter\Package Repository\sc_unix_plugins_checksum.txt* if SnapCenter Server is installed on Windows host.
- */opt/NetApp/snapcenter/SnapManagerWeb/Repository/sc_unix_plugins_checksum.txt* if SnapCenter Server in installed on Linux host.

The example should be used only as a reference for creating your own data.

Host requirements to install SnapCenter Plug-ins Package for Windows

Before you install the SnapCenter Plug-ins Package for Windows, you should be familiar with some basic host system space requirements and sizing requirements.

Item	Requirements
Operating Systems	Microsoft Windows For the latest information about supported versions, see the NetApp Interoperability Matrix Tool.

Item	Requirements
Minimum RAM for the SnapCenter plug-in on host	1 GB
Minimum install and log space for the SnapCenter plug-in on host	5 GB You should allocate sufficient disk space and monitor the storage consumption by the logs folder. The log space required varies depending on the number of the entities to be protected and the frequency of data protection operations. If there is no sufficient disk space, the logs will not be created for the recently run operations.
Required software packages	• ASP.NET Core Runtime 8.0.12 (and all subsequent 8.0.x patches) Hosting Bundle • PowerShell Core 7.4.2 • Java 11 Oracle Java and OpenJDK For the latest information about supported versions, see the NetApp Interoperability Matrix Tool. For .NET specific troubleshooting information, see SnapCenter upgrade or install fails for legacy systems that do not have internet connectivity.

Host requirements for installing the SnapCenter Plug-ins Package for Linux

Before you install the SnapCenter Plug-ins Package for Linux, you should be familiar with some basic host system space and sizing requirements.

Item	Requirements
Operating systems	• Red Hat Enterprise Linux • SUSE Linux Enterprise Server (SLES) For the latest information about supported versions, see the NetApp Interoperability Matrix Tool.
Minimum RAM for the SnapCenter plug-in on host	1 GB

Item	Requirements
Minimum install and log space for the SnapCenter plug-in on host	2 GB You should allocate sufficient disk space and monitor the storage consumption by the logs folder. The log space required varies, depending on the number of the entities to be protected and the frequency of data protection operations. If there is no sufficient disk space, the logs will not be created for the recently run operations.
Required software packages	Java 11 Oracle Java and OpenJDK If you have upgraded JAVA to the latest version, you must ensure that the JAVA_HOME option located at /var/opt/snapcenter/spl/etc/spl.properties is set to the correct JAVA version and the correct path. For the latest information about supported versions, see the NetApp Interoperability Matrix Tool.

Set up credentials for the SnapCenter Plug-in for IBM Db2

SnapCenter uses credentials to authenticate users for SnapCenter operations. You should create credentials for installing SnapCenter plug-ins and additional credentials for performing data protection operations on databases or Windows file systems.

About this task

- Linux hosts

  You must set up credentials for installing plug-ins on Linux hosts.

  You must set up the credentials for the root user or for a non-root user who has sudo privileges to install and start the plug-in process.

  Best Practice: Although you are allowed to create credentials for Linux after deploying hosts and installing plug-ins, the best practice is to create credentials after you add SVMs, before you deploy hosts and install plug-ins.

- Windows hosts

  You must set up Windows credentials before installing plug-ins.

  You must set up the credentials with administrator privileges, including administrator rights on the remote host.

If you set up credentials for individual resource groups and the username does not have full admin privileges,

you must assign at least the resource group and backup privileges to the username.

Steps

1. In the left navigation pane, click **Settings**.
2. In the Settings page, click **Credential**.
3. Click **New**.
4. In the Credential page, specify the information required for configuring credentials:

For this field…	Do this…
Credential name	Enter a name for the credentials.
User name	Enter the user name and password that are to be used for authentication. • Domain administrator or any member of the administrator group Specify the domain administrator or any member of the administrator group on the system on which you are installing the SnapCenter plug-in. Valid formats for the Username field are: ◦ *NetBIOS\UserName* ◦ *Domain FQDN\UserName* • Local administrator (for workgroups only) For systems that belong to a workgroup, specify the built-in local administrator on the system on which you are installing the SnapCenter plug-in. You can specify a local user account that belongs to the local administrators group if the user account has elevated privileges or the User Access control feature is disabled on the host system. The valid format for the Username field is: *UserName* Do not use double quotes (") or backtick (`) in the passwords. You should not use the less than (<) and exclamation (!) symbols together in passwords. For example, lessthan<!10, lessthan10<!, backtick`12.
Password	Enter the password used for authentication.
Authentication Mode	Select the authentication mode that you want to use.

For this field…	Do this…
Use sudo privileges	Select the **Use sudo privileges** check box if you are creating credentials for a non-root user. Applicable to Linux users only.

5. Click **OK**.

After you finish setting up the credentials, you might want to assign credential maintenance to a user or group of users in the User and Access page.

Configure gMSA on Windows Server 2016 or later

Windows Server 2016 or later enables you to create a group Managed Service Account (gMSA) that provides automated service account password management from a managed domain account.

Before you begin

- You should have a Windows Server 2016 or later domain controller.
- You should have a Windows Server 2016 or later host, which is a member of the domain.

Steps

1. Create a KDS root key to generate unique passwords for each object in your gMSA.
2. For each domain, run the following command from the Windows domain controller: Add-KDSRootKey -EffectiveImmediately
3. Create and configure your gMSA:
   a. Create a user group account in the following format:

      ```
      domainName\accountName$
      ```

   b. Add computer objects to the group.
   c. Use the user group you just created to create the gMSA.

      For example,

      ```
      New-ADServiceAccount -name <ServiceAccountName> -DNSHostName <fqdn>
      -PrincipalsAllowedToRetrieveManagedPassword <group>
      -ServicePrincipalNames <SPN1,SPN2,…>
      ```

   d. Run `Get-ADServiceAccount` command to verify the service account.
4. Configure the gMSA on your hosts:
   a. Enable the Active Directory module for Windows PowerShell on the host where you want to use the gMSA account.

To do this, run the following command from PowerShell:

```
PS C:\> Get-WindowsFeature AD-Domain-Services

Display Name                            Name                Install
State
------------                            ----
-------------
[ ] Active Directory Domain Services   AD-Domain-Services  Available


PS C:\> Install-WindowsFeature AD-DOMAIN-SERVICES

Success Restart Needed Exit Code      Feature Result
------- -------------- ---------      --------------
True    No             Success        {Active Directory Domain
Services, Active ...
WARNING: Windows automatic updating is not enabled. To ensure that
your newly-installed role or feature is
automatically updated, turn on Windows Update.
```

b. Restart your host.

c. Install the gMSA on your host by running the following command from the PowerShell command prompt: `Install-AdServiceAccount <gMSA>`

d. Verify your gMSA account by running the following command: `Test-AdServiceAccount <gMSA>`

5. Assign the administrative privileges to the configured gMSA on the host.
6. Add the Windows host by specifying the configured gMSA account in the SnapCenter Server.

SnapCenter Server will install the selected plug-ins on the host and the specified gMSA will be used as the service log on account during the plug-in installation.

Install the SnapCenter Plug-in for IBM Db2

Add hosts and install plug-in packages on remote hosts

You must use the SnapCenter Add Host page to add hosts, and then install the plug-ins packages. The plug-ins are automatically installed on the remote hosts. You can add a host and install plug-in packages either for an individual host or for a cluster.

Before you begin

- If the operating system of the SnapCenter Server host is Windows 2019 and the operating system of the plug-in host is Windows 2022, you should perform the following:
  - Upgrade to Windows Server 2019 (OS Build 17763.5936) or later
  - Upgrade to Windows Server 2022 (OS Build 20348.2402) or later

- You must be a user that is assigned to a role that has the plug-in install and uninstall permissions, such as the SnapCenter Admin role.
- When installing a plug-in on a Windows host, if you specify a credential that is not built-in, or if the user belongs to a local workgroup user, you must disable UAC on the host.
- You should ensure that the message queueing service is running.
- The administration documentation contains information about managing hosts.
- If you are using group Managed Service Account (gMSA), you should configure gMSA with administrative privileges.

  Configure group Managed Service Account on Windows Server 2016 or later for IBM Db2

About this task

- You cannot add a SnapCenter Server as a plug-in host to another SnapCenter Server.

Steps

1. In the left navigation pane, click **Hosts**.
2. Verify that the **Managed Hosts** tab is selected at the top.
3. Click **Add**.
4. In the Hosts page, perform the following actions:

For this field…	Do this…
Host Type	Select the type of host: • Windows • Linux The Plug-in for IBM Db2 is installed on the IBM Db2 client host, and this host can be on either a Windows system or a Linux system.
Host name	Enter the communication host name. Enter the fully qualified domain name (FQDN) or the IP address of the host. SnapCenter depends on the proper configuration of the DNS. Therefore, the best practice is to enter the FQDN.

For this field…	Do this…
Credentials	Either select the credential name that you created or create new credentials. The credential must have administrative rights on the remote host. For details, see the information about creating credentials. You can view details about the credentials by positioning your cursor over the credential name that you provided. The credentials authentication mode is determined by the host type that you specify in the Add Host wizard.

5. In the Select Plug-ins to Install section, select the plug-ins to install.

   While using the REST API to install Plug-in for Db2, you must pass the version as 3.0. For example, Db2:3.0

6. (Optional) Click **More Options**.

For this field…	Do this…
Port	Either retain the default port number or specify the port number. The default port number is 8145. If the SnapCenter Server was installed on a custom port, that port number will be displayed as the default port. If you manually installed the plug-ins and specified a custom port, you must specify the same port. Otherwise, the operation fails.
Installation Path	The Plug-in for IBM Db2 is installed on the IBM Db2 client host, and this host can be on either a Windows system or a Linux system. • For the SnapCenter Plug-ins Package for Windows, the default path is C:\Program Files\NetApp\SnapCenter. Optionally, you can customize the path. • For the SnapCenter Plug-ins Package for Linux, the default path is /opt/NetApp/snapcenter. Optionally, you can customize the path.

For this field…	Do this…
Skip preinstall checks	Select this check box if you already installed the plug-ins manually and you do not want to validate whether the host meets the requirements for installing the plug-in.
Use group Managed Service Account (gMSA) to run the plug-in services	For Windows host, select this check box if you want to use group Managed Service Account (gMSA) to run the plug-in services. Provide the gMSA name in the following format: domainName\accountName$. gMSA will be used as a log on service account only for SnapCenter Plug-in for Windows service.

7. Click **Submit**.

   If you have not selected the Skip prechecks checkbox, the host is validated to verify whether the host meets the requirements for installing the plug-in. The disk space, RAM, PowerShell version, .NET version, location (for Windows plug-ins), and Java 11 (for both Windows and Linux plug-ins) are validated against the minimum requirements. If the minimum requirements are not met, appropriate error or warning messages are displayed.

   If the error is related to disk space or RAM, you can update the web.config file located at C:\Program Files\NetApp\SnapCenter WebApp to modify the default values. If the error is related to other parameters, you must fix the issue.

   In an HA setup, if you are updating web.config file, you must update the file on both nodes.

8. If host type is Linux, verify the fingerprint, and then click **Confirm and Submit**.

   In a cluster setup, you should verify the fingerprint of each of the nodes in the cluster.

   Fingerprint verification is mandatory even if the same host was added earlier to SnapCenter and the fingerprint was confirmed.

9. Monitor the installation progress.
   - For Windows plug-in, the install and upgrade logs are located at: *C:\Windows\SnapCenter plugin\Install*<JOBID>\
   - For Linux plug-in, the install logs are located at: */var/opt/snapcenter/logs/SnapCenter_Linux_Host_Plug-in_Install*<JOBID>.log and the upgrade logs are located at: */var/opt/snapcenter/logs/SnapCenter_Linux_Host_Plug-in_Upgrade*<JOBID>.log

After you finish

If you want to upgrade to SnapCenter 6.0 or later, the existing PERL-based Plug-in for Db2 will be uninstalled from the remote plug-in server.

Install SnapCenter Plug-in Packages for Linux or Windows on multiple remote hosts by using cmdlets

You can install the SnapCenter Plug-in Packages for Linux or Windows on multiple hosts simultaneously by using the Install-SmHostPackage PowerShell cmdlet.

Before you begin

You must have logged in to SnapCenter as a domain user with local administrator rights on each host on which you want to install the plug-in package.

Steps

1. Launch PowerShell.
2. On the SnapCenter Server host, establish a session using the Open-SmConnection cmdlet, and then enter your credentials.
3. Install the plug-in on multiple hosts using the Install-SmHostPackage cmdlet and the required parameters.

   The information regarding the parameters that can be used with the cmdlet and their descriptions can be obtained by running *Get-Help command_name*. Alternatively, you can also refer to the SnapCenter Software Cmdlet Reference Guide.

   You can use the -skipprecheck option when you have installed the plug-ins manually and do not want to validate whether the host meets the requirements to install the plug-in.

4. Enter your credentials for remote installation.

Install the SnapCenter Plug-in for IBM Db2 on Linux hosts by using the command-line interface

You should install the SnapCenter Plug-in for IBM Db2 Database by using the SnapCenter user interface (UI). If your environment does not allow remote installation of the plug-in from the SnapCenter UI, you can install the Plug-in for IBM Db2 Database either in console mode or in silent mode by using the command-line interface (CLI).

Before you begin

- You should install the Plug-in for IBM Db2 Database on each of the Linux host where the IBM Db2 client resides.
- The Linux host on which you are installing the SnapCenter Plug-in for IBM Db2 Database must meet the dependent software, database, and operating system requirements.

  The Interoperability Matrix Tool (IMT) contains the latest information about the supported configurations.

  NetApp Interoperability Matrix Tool

- The SnapCenter Plug-in for IBM Db2 Database is part of SnapCenter Plug-ins Package for Linux. Before you install SnapCenter Plug-ins Package for Linux, you should have already installed SnapCenter on a Windows host.

About this task

If parameters are not mentioned, SnapCenter is installed with default values.

Steps

1. Copy the SnapCenter Plug-ins Package for Linux installation file (snapcenter_linux_host_plugin.bin) from C:\ProgramData\NetApp\SnapCenter\Package Repository to the host where you want to install the Plug-in

for IBM Db2.

You can access this path from the host where the SnapCenter Server is installed.

2. From the command prompt, navigate to the directory where you copied the installation file.
3. Install the plug-in: `path_to_installation_bin_file/snapcenter_linux_host_plugin.bin -i silent -DPORT=port_number_for_host -DSERVER_IP=server_name_or_ip_address -DSERVER_HTTPS_PORT=port_number_for_server`
   - -DPORT specifies the SMCore HTTPS communication port.
   - -DSERVER_IP specifies the SnapCenter Server IP address.
   - -DSERVER_HTTPS_PORT specifies the SnapCenter Server HTTPS port.
   - -DUSER_INSTALL_DIR specifies the directory where you want to install the SnapCenter Plug-ins Package for Linux.
   - DINSTALL_LOG_NAME specifies the name of the log file.

```
/tmp/sc-plugin-installer/snapcenter_linux_host_plugin.bin -i silent
-DPORT=8145 -DSERVER_IP=scserver.domain.com -DSERVER_HTTPS_PORT=8146
-DUSER_INSTALL_DIR=/opt
-DINSTALL_LOG_NAME=SnapCenter_Linux_Host_Plugin_Install_2.log
-DCHOSEN_FEATURE_LIST=CUSTOM
```

4. Edit the /<installation directory>/NetApp/snapcenter/scc/etc/SC_SMS_Services.properties file, and then add the PLUGINS_ENABLED = DB2:3.0 parameter.
5. Add the host to the SnapCenter Server using the Add-Smhost cmdlet and the required parameters.

   The information regarding the parameters that can be used with the command and their descriptions can be obtained by running *Get-Help command_name*. Alternatively, you can also refer to the SnapCenter Software Cmdlet Reference Guide.

Monitor the status of installing Plug-in for IBM Db2

You can monitor the progress of SnapCenter plug-in package installation by using the Jobs page. You might want to check the progress of installation to determine when it is complete or if there is an issue.

About this task

The following icons appear on the Jobs page and indicate the state of the operation:

- In progress
- Completed successfully
- Failed
- Completed with warnings or could not start due to warnings
- Queued

Steps

1. In the left navigation pane, click **Monitor**.
2. In the **Monitor** page, click **Jobs**.
3. In the **Jobs** page, to filter the list so that only plug-in installation operations are listed, do the following:
   a. Click **Filter**.
   b. Optional: Specify the start and end date.
   c. From the Type drop-down menu, select **Plug-in installation**.
   d. From the Status drop-down menu, select the installation status.
   e. Click **Apply**.
4. Select the installation job and click **Details** to view the job details.
5. In the **Job Details** page, click **View logs**.

Configure CA Certificate

Generate CA Certificate CSR file

You can generate a Certificate Signing Request (CSR) and import the certificate that can be obtained from a Certificate Authority (CA) using the generated CSR. The certificate will have a private key associated with it.

CSR is a block of encoded text that is given to an authorized certificate vendor to procure the signed CA certificate.

CA Certificate RSA key length must be minimum 3072 bits.

For information to generate a CSR, see How to generate CA Certificate CSR file.

If you own the CA certificate for your domain (*.domain.company.com) or your system (machine1.domain.company.com), you can skip generating the CA Certificate CSR file. You can deploy the existing CA certificate with SnapCenter.

For cluster configurations, the cluster name (virtual cluster FQDN), and the respective host names should be mentioned in the CA certificate. The certificate can be updated by filling the Subject Alternative Name (SAN) field before procuring the certificate. For a wild card certificate (*.domain.company.com), the certificate will contain all the hostnames of the domain implicitly.

Import CA certificates

You must import the CA certificates to the SnapCenter Server and the Windows host plug-ins using the Microsoft management console (MMC).

Steps

1. Go to the Microsoft management console (MMC), and then click **File** > **Add/Remove Snapin**.
2. In the Add or Remove Snap-ins window, select **Certificates** and then click **Add**.
3. In the Certificates snap-in window, select the **Computer account** option, and then click **Finish**.
4. Click **Console Root** > **Certificates – Local Computer** > **Trusted Root Certification Authorities** >

Certificates.

5. Right-click on the folder “Trusted Root Certification Authorities”, and then select **All Tasks** > **Import** to start the import wizard.
6. Complete the wizard, as follows:

In this wizard window…	Do the following…
Import Private Key	Select the option **Yes**, import the private key, and then click **Next**.
Import File Format	Make no changes; click **Next**.
Security	Specify the new password to be used for the exported certificate, and then click **Next**.
Completing the Certificate Import Wizard	Review the summary, and then click **Finish** to start the import.

Importing certificate should be bundled with the private key (supported formats are: *.pfx, *.p12, and *.p7b).

7. Repeat Step 5 for the “Personal” folder.

Get the CA certificate thumbprint

A certificate thumbprint is a hexadecimal string that identifies a certificate. A thumbprint is calculated from the content of the certificate using a thumbprint algorithm.

Steps

1. Perform the following on the GUI:
   a. Double-click the certificate.
   b. In the Certificate dialog box, click the **Details** tab.
   c. Scroll through the list of fields and click **Thumbprint**.
   d. Copy the hexadecimal characters from the box.
   e. Remove the spaces between the hexadecimal numbers.

      For example, if the thumbprint is: "a9 09 50 2d d8 2a e4 14 33 e6 f8 38 86 b0 0d 42 77 a3 2a 7b", after removing the spaces, it will be: "a909502dd82ae41433e6f83886b00d4277a32a7b".

2. Perform the following from PowerShell:
   a. Run the following command to list the thumbprint of the installed certificate and identify the recently installed certificate by the subject name.

      Get-ChildItem -Path Cert:\LocalMachine\My

   b. Copy the thumbprint.

Configure CA certificate with Windows host plug-in services

You should configure the CA certificate with Windows host plug-in services to activate the installed digital certificate.

Perform the following steps on the SnapCenter Server and all the plug-in hosts where CA certificates are already deployed.

Steps

1. Remove the existing certificate binding with SMCore default port 8145, by running the following command:

   ```
   > netsh http delete sslcert ipport=0.0.0.0:_<SMCore Port>
   ```

   For example:

   ```
   > netsh http delete sslcert ipport=0.0.0.0:8145
   ```

2. Bind the newly installed certificate with the Windows host plug-in services, by running the following commands:

   ```
   > $cert = "_<certificate thumbprint>_"
   > $guid = [guid]::NewGuid().ToString("B")
   > netsh http add sslcert ipport=0.0.0.0: _<SMCore Port>_ certhash=$cert
   appid="$guid"
   ```

   For example:

   ```
   > $cert = "a909502dd82ae41433e6f83886b00d4277a32a7b"
   > $guid = [guid]::NewGuid().ToString("B")
   > netsh http add sslcert ipport=0.0.0.0: _<SMCore Port>_ certhash=$cert
   appid="$guid"
   ```

Configure the CA Certificate for the SnapCenter IBM Db2 Plug-ins service on Linux host

You should manage the password of the plug-ins keystore and its certificate, configure the CA certificate, configure root or intermediate certificates to the plug-ins trust-store, and configure CA signed key pair to plug-ins trust-store with SnapCenter plug-ins service to activate the installed digital certificate.

The plug-ins uses the file 'keystore.jks', which is located at */opt/NetApp/snapcenter/scc/etc* both as its trust-store and key-store.

Manage password for plug-in keystore and alias of the CA signed key pair in use

Steps

1. You can retrieve plug-in keystore default password from plug-in agent property file.

It is the value corresponding to the key 'KEYSTORE_PASS'.

2. Change the keystore password:

```
keytool -storepasswd -keystore keystore.jks
```

3. Change the password for all aliases of private key entries in the keystore to the same password used for the keystore:

```
keytool -keypasswd -alias "alias_name_in_cert" -keystore keystore.jks
```

Update the same for the key KEYSTORE_PASS in *agent.properties* file.

4. Restart the service after changing the password.

Password for plug-in keystore and for all the associated alias password of the private key should be same.

Configure root or intermediate certificates to plug-in trust-store

You should configure the root or intermediate certificates without the private key to plug-in trust-store.

Steps

1. Navigate to the folder containing the plug-in keystore: /opt/NetApp/snapcenter/scc/etc.
2. Locate the file 'keystore.jks'.
3. List the added certificates in the keystore:

```
keytool -list -v -keystore keystore.jks
```

4. Add a root or intermediate certificate:

```
keytool -import -trustcacerts -alias myRootCA -file
/root/USERTrustRSA_Root.cer -keystore keystore.jks
```

5. Restart the service after configuring the root or intermediate certificates to plug-in trust-store.

You should add the root CA certificate and then the intermediate CA certificates.

Configure CA signed key pair to plug-in trust-store

You should configure the CA signed key pair to the plug-in trust-store.

Steps

1. Navigate to the folder containing the plug-in keystore /opt/NetApp/snapcenter/scc/etc.
2. Locate the file 'keystore.jks'.
3. List the added certificates in the keystore:

```
keytool -list -v -keystore keystore.jks
```

4. Add the CA certificate having both private and public key.

```
keytool -importkeystore -srckeystore /root/snapcenter.ssl.test.netapp.com.pfx
-srcstoretype pkcs12 -destkeystore keystore.jks -deststoretype JKS
```

5. List the added certificates in the keystore.

```
keytool -list -v -keystore keystore.jks
```

6. Verify that the keystore contains the alias corresponding to the new CA certificate, which was added to the keystore.
7. Change the added private key password for CA certificate to the keystore password.

   Default plug-in keystore password is the value of the key KEYSTORE_PASS in agent.properties file.

```
keytool -keypasswd -alias "alias_name_in_CA_cert" -keystore keystore.jks
```

8. If the alias name in the CA certificate is long and contains space or special characters ("*",","), change the alias name to a simple name:

```
keytool -changealias -alias "long_alias_name" -destalias "simple_alias"
-keystore keystore.jks
```

9. Configure the alias name from CA certificate in agent.properties file.

   Update this value against the key SCC_CERTIFICATE_ALIAS.

10. Restart the service after configuring the CA signed key pair to plug-in trust-store.

Configure certificate revocation list (CRL) for plug-ins

About this task

- SnapCenter Plug-ins will search for the CRL files in a pre-configured directory.
- Default directory for the CRL files for SnapCenter Plug-ins is ' opt/NetApp/snapcenter/scc/etc/crl'.

Steps

1. You can modify and update the default directory in agent.properties file against the key CRL_PATH.

   You can place more than one CRL file in this directory. The incoming certificates will be verified against each CRL.

Configure the CA Certificate for the SnapCenter IBM Db2 Plug-ins service on Windows host

You should manage the password of the plug-ins keystore and its certificate, configure the CA certificate, configure root or intermediate certificates to the plug-ins trust-store, and configure CA signed key pair to plug-ins trust-store with SnapCenter plug-ins service

to activate the installed digital certificate.

The plug-ins uses the file *keystore.jks*, which is located at *C:\Program Files\NetApp\SnapCenter\Snapcenter Plug-in Creator\etc* both as its trust-store and key-store.

Manage password for plug-in keystore and alias of the CA signed key pair in use

Steps

1. You can retrieve plug-in keystore default password from plug-in agent property file.

   It is the value corresponding to the key *KEYSTORE_PASS*.

2. Change the keystore password:

   keytool -storepasswd -keystore keystore.jks

   > If the "keytool" command is not recognized on the Windows command prompt, replace the keytool command with its complete path.

   C:\Program Files\Java\<jdk_version>\bin\keytool.exe" -storepasswd -keystore keystore.jks

3. Change the password for all aliases of private key entries in the keystore to the same password used for the keystore:

   keytool -keypasswd -alias "alias_name_in_cert" -keystore keystore.jks

   Update the same for the key KEYSTORE_PASS in *agent.properties* file.

4. Restart the service after changing the password.

   > Password for plug-in keystore and for all the associated alias password of the private key should be same.

Configure root or intermediate certificates to plug-in trust-store

You should configure the root or intermediate certificates without the private key to plug-in trust-store.

Steps

1. Navigate to the folder containing the plug-in keystore *C:\Program Files\NetApp\SnapCenter\Snapcenter Plug-in Creator\etc*
2. Locate the file 'keystore.jks'.
3. List the added certificates in the keystore:

   keytool -list -v -keystore keystore.jks

4. Add a root or intermediate certificate:

   keytool -import -trustcacerts -alias myRootCA -file /root/USERTrustRSA_Root.cer -keystore keystore.jks

5. Restart the service after configuring the root or intermediate certificates to plug-in trust-store.

> You should add the root CA certificate and then the intermediate CA certificates.

Configure CA signed key pair to plug-in trust-store

You should configure the CA signed key pair to the plug-in trust-store.

Steps

1. Navigate to the folder containing the plug-in keystore *C:\Program Files\NetApp\SnapCenter\Snapcenter Plug-in Creator\etc*
2. Locate the file *keystore.jks*.
3. List the added certificates in the keystore:

   keytool -list -v -keystore keystore.jks

4. Add the CA certificate having both private and public key.

   keytool -importkeystore -srckeystore /root/snapcenter.ssl.test.netapp.com.pfx -srcstoretype pkcs12 -destkeystore keystore.jks -deststoretype JKS

5. List the added certificates in the keystore.

   keytool -list -v -keystore keystore.jks

6. Verify that the keystore contains the alias corresponding to the new CA certificate, which was added to the keystore.
7. Change the added private key password for CA certificate to the keystore password.

   Default plug-in keystore password is the value of the key KEYSTORE_PASS in agent.properties file.

   keytool -keypasswd -alias "alias_name_in_CA_cert" -keystore keystore.jks

8. Configure the alias name from CA certificate in *agent.properties* file.

   Update this value against the key SCC_CERTIFICATE_ALIAS.

9. Restart the service after configuring the CA signed key pair to plug-in trust-store.

Configure certificate revocation list (CRL) for SnapCenter plug-ins

About this task

- To download the latest CRL file for the related CA certificate see How to update certificate revocation list file in SnapCenter CA Certificate.
- SnapCenter plug-ins will search for the CRL files in a pre-configured directory.
- Default directory for the CRL files for SnapCenter plug-ins is *'C:\Program Files\NetApp\SnapCenter\Snapcenter Plug-in Creator\ etc\crl'*.

Steps

1. You can modify and update the default directory in *agent.properties* file against the key CRL_PATH.
2. You can place more than one CRL file in this directory.

   The incoming certificates will be verified against each CRL.

Enable CA Certificates for plug-ins

You should configure the CA certificates and deploy the CA certificates in the SnapCenter Server and the corresponding plug-in hosts. You should enable the CA certificate validation for the plug-ins.

Before you begin

- You can enable or disable the CA certificates using the run *Set-SmCertificateSettings* cmdlet.
- You can display the certificate status for the plug-ins using the *Get-SmCertificateSettings*.

The information regarding the parameters that can be used with the cmdlet and their descriptions can be obtained by running *Get-Help command_name*. Alternatively, you can also refer to the SnapCenter Software Cmdlet Reference Guide.

Steps

1. In the left navigation pane, click **Hosts**.
2. In the Hosts page, click **Managed Hosts**.
3. Select single or multiple plug-in hosts.
4. Click **More options**.
5. Select **Enable Certificate Validation**.

After you finish

The Managed Hosts tab host displays a padlock and the color of the padlock indicates the status of the connection between SnapCenter Server and the plug-in host.

- indicates that the CA certificate is neither enabled nor assigned to the plug-in host.
- indicates that the CA certificate is successfully validated.
- indicates that the CA certificate could not be validated.
- indicates that the connection information could not be retrieved.

When the status is yellow or green, the data protection operations completes successfully.

Prepare for data protection

Prerequisites for using the SnapCenter Plug-in for IBM Db2

Before you use SnapCenter Plug-in for IBM Db2, the SnapCenter administrator must install and configure the SnapCenter Server and perform the prerequisite tasks.

- Install and configure SnapCenter Server.
- Log in to SnapCenter Server.
- Configure the SnapCenter environment by adding storage system connections and creating credentials, if applicable.
- Install Java 11 on your Linux or Windows host.

  You must set the Java path in the environmental path variable of the host machine.

- Set up SnapMirror and SnapVault, if you want backup replication.

How resources, resource groups, and policies are used for protecting IBM Db2

Before you use SnapCenter, it is helpful to understand basic concepts related to the backup, clone, and restore operations you want to perform. You interact with resources, resource groups, and policies for different operations.

- Resources are typically IBM Db2 databases that you back up or clone with SnapCenter.
- A SnapCenter resource group, is a collection of resources on a host.

  When you perform an operation on a resource group, you perform that operation on the resources defined in the resource group according to the schedule you specify for the resource group.

  You can back up on demand a single resource or a resource group. You also can perform scheduled backups for single resources and resource groups.

- The policies specify the backup frequency, replication, scripts, and other characteristics of data protection operations.

  When you create a resource group, you select one or more policies for that group. You can also select a policy when you perform a backup on demand for a single resource.

Think of a resource group as defining what you want to protect and when you want to protect it in terms of day and time. Think of a policy as defining how you want to protect it. If you are backing up all databases, for example, you might create a resource group that includes all of the databases in the host. You could then attach two policies to the resource group: a daily policy and an hourly policy. When you create the resource group and attach the policies, you might configure the resource group to perform a full backup daily.

Back up IBM Db2 resources

Back up IBM Db2 resources

You can either create a backup of a resource (database) or resource group. The backup workflow includes planning, identifying the databases for backup, managing backup policies, creating resource groups and attaching policies, creating backups, and monitoring the operations.

The following workflow shows the sequence in which you must perform the backup operation:

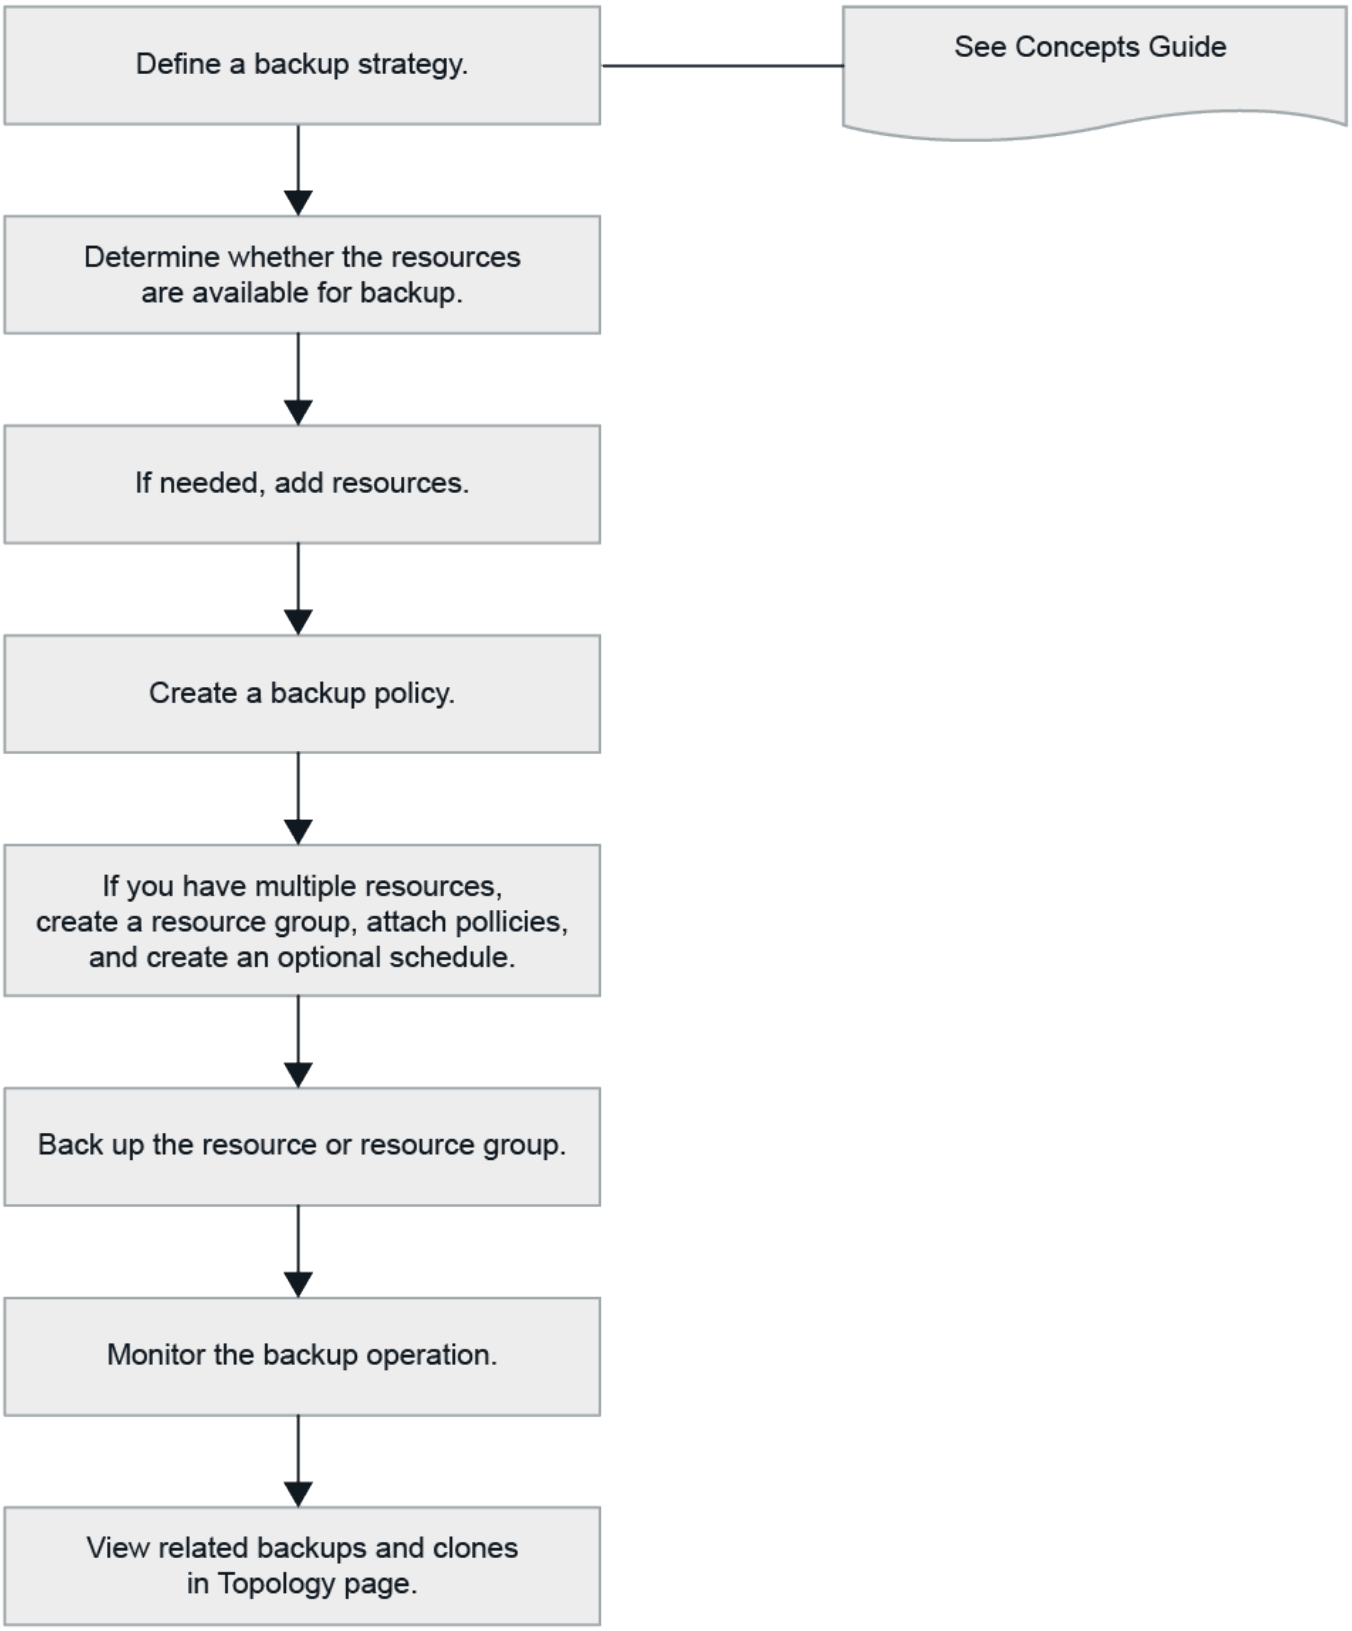

You can also use PowerShell cmdlets manually or in scripts to perform backup, restore, and clone operations. The SnapCenter cmdlet help and the cmdlet reference information contain more information about PowerShell cmdlets.
SnapCenter Software Cmdlet Reference Guide.

Discover the databases automatically

Resources are IBM Db2 databases on the Linux host that are managed by SnapCenter. You can add the resources to resource groups to perform data protection operations after you discover the IBM Db2 databases that are available.

Before you begin

- You must have already completed tasks such as installing the SnapCenter Server, adding hosts, and setting up the storage system connections.
- SnapCenter Plug-in for IBM Db2 does not support automatic discovery of the resources residing on RDM/VMDK virtual environments. You must provide the storage information for virtual environments while adding the databases manually.

About this task

- After installing the plug-in, all the databases on that Linux host are automatically discovered and displayed on the Resources page.
- Only databases are auto-discovered.

The automatically discovered resources cannot be modified or deleted.

Steps

1. In the left navigation pane, click **Resources**, and then select the Plug-in for IBM Db2 from the list.
2. In the Resources page select the resource type from the View list.
3. (Optional) Click , and then select the host name.

   You can then click to close the filter pane.
4. Click **Refresh Resources** to discover the resources available on the host.

   The resources are displayed along with information such as resource type, host name, associated resource groups, backup type, policies and overall status.

   - If the database is on a NetApp storage and not protected, then Not protected is displayed in the Overall Status column.
   - If the database is on a NetApp storage system and protected, and if there is no backup operation performed, then Backup not run is displayed in the Overall Status column. The status will otherwise change to Backup failed or Backup succeeded based on the last backup status.

You must refresh the resources if the databases are renamed outside of SnapCenter.

Add resources manually to the plug-in host

Automatic discovery is not supported on Windows host. You must add Db2 instances and database resources manually.

Before you begin

- You must have completed tasks such as installing the SnapCenter Server, adding hosts, and setting up storage system connections.

About this task

Manual discovery is not supported for the following configurations:

- RDM and VMDK layouts

Steps

1. In the left navigation pane, select **Resources**, and then select the SnapCenter Plug-in for IBM Db2 from the drop-down list.
2. In the Resources page, click **Add IBM DB2 Resource**.
3. In the Provide Resource Details page, perform the following actions:

For this field…	Do this…
Name	Specify the database name.
Host Name	Enter the host name.
Type	Select database or instance.
Instance	Specify the name of the instance, which is the parent of the database.
Credentials	Select the credentials or add information for the credential. This is optional.

4. In the Provide Storage Footprint page, select a storage type and choose one or more volumes, LUNs, and qtrees, and then click **Save**.

   Optional: You can click the + icon to add more volumes, LUNs, and qtrees from other storage systems.

5. Optional: In the Resource Settings page, for resources on the Windows host, enter custom key-value pairs for IBM Db2 plug-in
6. Review the summary, and then click **Finish**.

   The databases are displayed along with information such as the host name, associated resource groups and policies, and overall status

If you want to provide users access to resources, you must assign the resources to the users. This enables users to perform the actions for which they have permissions on the assets that are assigned to them.

Add a user or group and assign role and assets

After adding the databases, you can modify the IBM Db2 database details.

Create backup policies for IBM Db2

Before you use SnapCenter to back up IBM Db2 resources, you must create a backup policy for the resource or resource group that you want to back up. A backup policy is a

set of rules that governs how you manage, schedule, and retain backups.

Before you begin

- You must have defined your backup strategy.

  For details, see the information about defining a data protection strategy for IBM Db2 databases.

- You must have prepared for data protection by completing tasks such as installing SnapCenter, adding hosts, setting up storage system connections, and adding resources.
- The SnapCenter administrator must have assigned the SVMs for both the source and destination volumes to you if you are replicating snapshots to a mirror or vault.

  Additionally, you can specify replication, script, and application settings in the policy. These options saves time when you want to reuse the policy for another resource group.

About this task

- SnapLock
  - If 'Retain the backup copies for a specific number of days' option is selected, then the SnapLock retention period must be lesser than or equal to the mentioned retention days.
  - Specifying a snapshot locking period prevents deletion of the snapshots until the retention period expires. This could lead to retaining a larger number of snapshots than the count specified in the policy.
  - For ONTAP 9.12.1 and below version, the clones created from the SnapLock Vault snapshots as part of restore will inherit the SnapLock Vault expiry time. Storage admin should manually cleanup the clones post the SnapLock expiry time.

Steps

1. In the left navigation pane, click **Settings**.
2. In the Settings page, click **Policies**.
3. Click **New**.
4. In the Name page, enter the policy name and details.
5. In the Policy type page, perform the following:
   a. Select storage type.
   b. In the **Custom backup settings** section, provide any specific backup settings that have to be passed to the plug-in in key-value format.

      You can provide multiple key-values to be passed to the plug-in.

6. In the Snapshot and Replication page, perform the following actions.
   a. Specify the schedule frequency by selecting **On demand**, **Hourly**, **Daily**, **Weekly**, or **Monthly**.

      You can specify the schedule (start date, end date, and frequency) for the backup operation while creating a resource group. This enables you to create resource groups that share the same policy and backup frequency, but also enables you to assign different backup schedules to each policy.

      If you have scheduled for 2:00 a.m., the schedule will not be triggered during daylight saving time (DST).

b. In the Snapshot settings section, perform the following actions:

If you want to…	Then…
Keep a certain number of snapshots	Select **Copies to keep**, and then specify the number of snapshots that you want to keep. If the number of snapshots exceeds the specified number, the snapshots are deleted with the oldest copies deleted first.
Keep the Snapshots for a certain number of days	Select **Retain copies for**, and then specify the number of days for which you want to keep the Snapshots before deleting them.
Snapshot copy locking period	Select **Snapshot copy locking period** and specify days, months, or years. Snaplock retention period should be less than 100 years.

For Snapshot copy-based backups, you must set the retention count to 2 or higher if you plan to enable SnapVault replication. If you set the retention count to 1, the retention operation might fail because the first snapshot is the reference snapshot for the SnapVault relationship until a newer snapshot is replicated to the target.

c. Specify the Policy label.

You can assign SnapMirror labels to primary snapshots for remote replication, allowing the primary snapshots to offload the snapshot replication operation from SnapCenter to ONTAP secondary systems. This can be done without enabling SnapMirror or SnapVault option in the policy page.

7. In the Select secondary replication options section, select one or both of the following secondary replication options:

For this field…	Do this…
Update SnapMirror after creating a local Snapshot copy	Select this field to create mirror copies of the backup sets on another volume (SnapMirror replication). This option should be enabled for SnapMirror active sync.
Update SnapVault after creating a local Snapshot copy	Select this option to perform disk-to-disk backup replication (SnapVault backups).
Error retry count	Enter the maximum number of replication attempts that can be allowed before the operation stops.

You should configure SnapMirror retention policy in ONTAP for the secondary storage to avoid reaching the maximum limit of Snapshots on the secondary storage.

8. Review the summary, and then click **Finish**.

Create resource groups and attach policies

A resource group is the container to which you must add resources that you want to back up and protect. A resource group enables you to back up all the data that is associated with a given application simultaneously. A resource group is required for any data protection job. You must also attach one or more policies to the resource group to define the type of data protection job that you want to perform.

About this task

- For ONTAP 9.12.1 and below version, the clones created from the SnapLock Vault snapshots as part of restore will inherit the SnapLock Vault expiry time. Storage admin should manually cleanup the clones post the SnapLock expiry time.

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page, click **New Resource Group**.
3. In the Name page, perform the following actions:

For this field…	Do this…
Name	Enter a name for the resource group. The resource group name should not exceed 250 characters.
Tags	Enter one or more labels that will help you later search for the resource group. For example, if you add HR as a tag to multiple resource groups, you can later find all resource groups associated with the HR tag.
Use custom name format for snapshot copy	Select this check box, and enter a custom name format that you want to use for the snapshot name. For example, customtext_resource group_policy_hostname or resource group_hostname. By default, a timestamp is appended to the snapshot name.

4. In the Resources page, select a host name from the **Host** drop-down list and resource type from the **Resource Type** drop-down list.

   This helps to filter information on the screen.

5. Select the resources from the **Available Resources** section, and then click the right arrow to move them to the **Selected Resources** section.
6. In the Application Settings page, do the following:
   a. Click the **Backups** arrow to set additional backup options:

      Enable consistency group backup and perform the following tasks:

For this field…	Do this…
Afford time to wait for Consistency Group snapshot operation to complete	Select **Urgent**, **Medium**, or **Relaxed** to specify the wait time for snapshot operation to complete. Urgent = 5 seconds, Medium = 7 seconds, and Relaxed = 20 seconds.
Disable WAFL Sync	Select this to avoid forcing a WAFL consistency point.

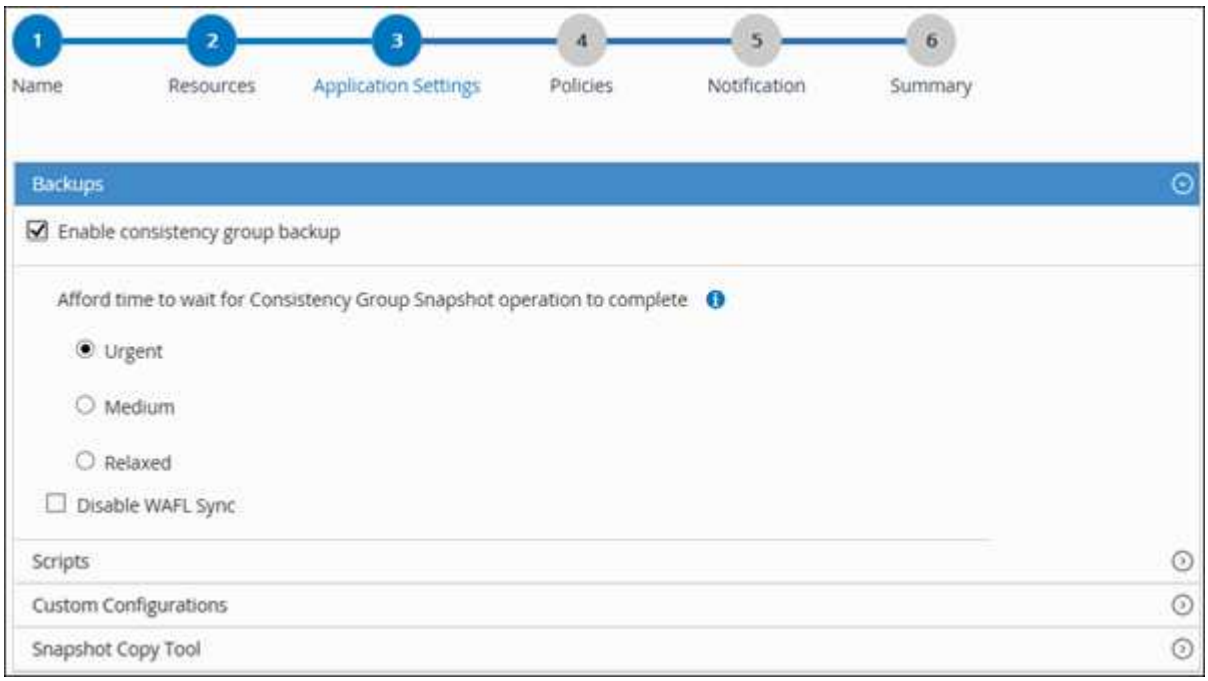

   b. Click the **Scripts** arrow and enter the pre and post commands for quiesce, snapshot, and unquiesce operations. You can also enter the pre commands to be executed before exiting in the event of a failure.
   c. Click the **Custom Configurations** arrow and enter the custom key-value pairs required for all data protection operations using this resource.

Parameter	Setting	Description
ARCHIVE_LOG_ENABLE	(Y/N)	Enables the archive log management to delete the archive logs.

Parameter	Setting	Description
ARCHIVE_LOG_RETENTION	number_of_days	Specifies the number of days the archive logs are retained. This setting must be equal to or greater than NTAP_SNAPSHOT_ RETENTIONS.
ARCHIVE_LOG_DIR	change_info_directory/logs	Specifies the path to the directory that contains the archive logs.
ARCHIVE_LOG_EXT	file_extension	Specifies the archive log file extension length. For example, if the archive log is log_backup_0_0_0_0.16151855 1942 9 and if the file_extension value is 5, then the extension of the log will retain 5 digits, which is 16151.
ARCHIVE_LOG_RECURSIVE_ SE ARCH	(Y/N)	Enables the management of archive logs within subdirectories. You should use this parameter if the archive logs are located under subdirectories.

The custom key-value pairs are supported for IBM Db2 Linux plug-in systems and not supported for IBM Db2 database registered as a centralized windows plug-in.

d. Click the **Snapshot Copy Tool** arrow to select the tool to create snapshots:

If you want…	Then…
SnapCenter to use the plug-in for Windows and put the file system into a consistent state before creating a snapshot. For Linux resources, this option is not applicable.	Select **SnapCenter with File System Consistency**.
SnapCenter to create a storage level snapshot	Select **SnapCenter without File System Consistency**.

If you want…	Then…
To enter the command to be executed on the host to create snapshot copies.	Select **Other**, and then enter the command to be executed on the host to create a snapshot.

7. In the Policies page, perform the following steps:
   a. Select one or more policies from the drop-down list.

      You can also create a policy by clicking + .

      The policies are listed in the Configure schedules for selected policies section.

   b. In the Configure Schedules column, click + for the policy you want to configure.
   c. In the Add schedules for policy *policy_name* dialog box, configure the schedule, and then click **OK**.

      Where, policy_name is the name of the policy that you have selected.

      The configured schedules are listed in the **Applied Schedules** column.

      Third party backup schedules are not supported when they overlap with SnapCenter backup schedules.

8. In the Notification page, from the **Email preference** drop-down list, select the scenarios in which you want to send the emails.

   You must also specify the sender and receiver email addresses, and the subject of the email. The SMTP server must be configured in **Settings** > **Global Settings**.

9. Review the summary, and then click **Finish**.

Create resource groups and enable secondary protection for IBM Db2 resources on ASA r2 systems

You should create the resource group to add the resources that are on ASA r2 systems. You can also provision the secondary protection while creating the resource group.

Before you begin

- You should ensure that you are not adding both ONTAP 9.x resources and ASA r2 resources to the same resource group.
- You should ensure that you do not have a database with both ONTAP 9.x resources and ASA r2 resources.

About this task

- The secondary protection is available only if the logged-in user is assigned to the role that has the **SecondaryProtection** capability enabled.
- If you enabled secondary protection, the resource group is put into maintenance mode while creating the primary and secondary consistency groups. After the primary and secondary consistency groups are created, the resource group is put out of maintenance mode.
- SnapCenter does not support secondary protection for a clone resource.

Steps

1. In the left navigation pane, select **Resources**, and the appropriate plug-in from the list.
2. In the Resources page, click **New Resource Group**.
3. In the Name page, perform the following actions:
   a. Enter a name for the resource group in the Name field.

      The resource group name should not exceed 250 characters.

   b. Enter one or more labels in the Tag field to help you search for the resource group later.

      For example, if you add HR as a tag to multiple resource groups, you can later find all resource groups associated with the HR tag.

   c. Select this check box, and enter a custom name format that you want to use for the Snapshot name.

      For example, customtext_resource group_policy_hostname or resource group_hostname. By default, a timestamp is appended to the Snapshot name.

   d. Specify the destinations of the archive log files that you do not want to back up.

      You should use the exact same destination as it was set in the application including prefix, if needed.

4. In the Resources page, select the database host name from the **Host** drop-down list.

   The resources are listed in the Available Resources section only if the resource is discovered successfully. If you have recently added resources, they will appear on the list of available resources only after you refresh your resource list.

5. Select the ASA r2 resources from the Available Resources section and move them to the Selected Resources section.
6. In the Application Settings page, select the backup option.
7. In the Policies page, perform the following steps:
   a. Select one or more policies from the drop-down list.

      You can also create a policy by clicking + .

      In the Configure schedules for selected policies section, the selected policies are listed.

   b. Click + in the Configure Schedules column for the policy for which you want to configure a schedule.
   c. In the Add schedules for policy *policy_name* window, configure the schedule, and then click **OK**.

      Where, *policy_name* is the name of the policy that you have selected.

      The configured schedules are listed in the Applied Schedules column.

Third party backup schedules are not supported when they overlap with SnapCenter backup schedules.

8. If the secondary protection is enabled for the policy that you have selected, then Secondary Protection page is displayed and you need to perform the following steps:
   a. Select the type of the replication policy.

      Synchronous replication policy is not supported.

   b. Specify the consistency group suffix that you want to use.
   c. From the Destination Cluster and Destination SVM drop-downs select the peered cluster and SVM that you want to use.

      The cluster and SVM peering is not supported by SnapCenter. You should use System Manager or ONTAP CLIs to perform cluster and SVM peering.

If the resources are already protected outside of SnapCenter, those resources will be displayed in the Secondary Protected Resources section.

1. On the Verification page, perform the following steps:
   a. Click **Load locators** to load the SnapMirror or SnapVault volumes to perform verification on secondary storage.
   b. Click in the Configure Schedules column to configure the verification schedule for all the schedule types of the policy.
   c. In the Add Verification Schedules policy_name dialog box, perform the following actions:

If you want to…	Do this…
Run verification after backup	Select **Run verification after backup**.
Schedule a verification	Select **Run scheduled verification** and then select the schedule type from the drop-down list.

   d. Select **Verify on secondary location** to verify your backups on secondary storage system.
   e. Click **OK**.

      The configured verification schedules are listed in the Applied Schedules column.

2. In the Notification page, from the **Email preference** drop-down list, select the scenarios in which you want to send the emails.

   You must also specify the sender and receiver email addresses, and the subject of the email. If you want to attach the report of the operation performed on the resource group, select **Attach Job Report**.

   For email notification, you must have specified the SMTP server details using the either the GUI or the PowerShell command Set-SmSmtpServer.

3. Review the summary, and then click **Finish**.

Create a storage system connection and a credential using PowerShell cmdlets for IBM Db2

You must create a storage virtual machine (SVM) connection and a credential before using PowerShell cmdlets to back up, restore, or clone IBM Db2 databases.

Before you begin

- You should have prepared the PowerShell environment to execute the PowerShell cmdlets.
- You should have the required permissions in the Infrastructure Admin role to create storage connections.
- You should ensure that the plug-in installations are not in progress.

  Host plug-in installations must not be in progress while adding a storage system connection because the host cache might not be updated and databases status might be displayed in the SnapCenter GUI as “Not available for backup” or “Not on NetApp storage”.

- Storage system names should be unique.

  SnapCenter does not support multiple storage systems with the same name on different clusters. Each storage system that is supported by SnapCenter should have a unique name and a unique data LIF IP address.

Steps

1. Click **SnapCenterPS** to launch the PowerShell Core.
2. Create a new connection to the storage system by using the Add-SmStorageConnection cmdlet.

   ```
   PS C:\> Add-SmStorageConnection -StorageType DataOntap -Type DataOntap
   -OntapStorage ‘scsnfssvm’ -Protocol Https -Timeout 60
   ```

3. Create a new credential by using the Add-SmCredential cmdlet.

   This example shows how to create a new credential named FinanceAdmin with Windows credentials:

   ```
   PS C:\> Add-SmCredential -Name 'FinanceAdmin' -Type Linux
   -AuthenticationType PasswordBased -Credential db2hostuser
   -EnableSudoPrevileges:$true
   ```

4. Add the IBM Db2 communication host to SnapCenter Server.

   For Linux:

   ```
   PS C:\> Add-SmHost -HostType Linux -HostName '10.232.204.61'
   -CredentialName 'defaultcreds'
   ```

   For Windows:

```
PS C:\> Add-SmHost -HostType Windows -HostName '10.232.204.61'
-CredentialName 'defaultcreds'
```

5. Install the package and the SnapCenter Plug-in for IBM Db2 on the host.

   For Linux:

```
PS C:\> Install-SmHostPackage -HostNames '10.232.204.61' -PluginCodes
DB2
```

   For Windows:

```
PS C:\> Install-SmHostPackage -HostNames '10.232.204.61' -PluginCodes
DB2,SCW
```

6. Set the path to the SQLLIB.

   For Windows, Db2 plug-in will use the default path for SQLLIB folder: “C:\Program Files\IBM\SQLLIB\BIN”

   If you want to override the default path, use the following command.

```
PS C:\> Set-SmConfigSettings -Plugin -HostName '10.232.204.61'
-PluginCode DB2 -configSettings
@{"DB2_SQLLIB_CMD"="<custom_path>\IBM\SQLLIB\BIN"}
```

The information regarding the parameters that can be used with the cmdlet and their descriptions can be obtained by running *Get-Help command_name*. Alternatively, you can also refer to the SnapCenter Software Cmdlet Reference Guide.

Back up Db2 databases

Backing up a database includes establishing a connection with the SnapCenter Server, adding resources, adding a policy, creating a backup resource group, and backing up.

Before you begin

- You must have created a backup policy.
- If you want to back up a resource that has a SnapMirror relationship with a secondary storage, the ONTAP role assigned to the storage user should include the “snapmirror all” privilege. However, if you are using the “vsadmin” role, then the “snapmirror all” privilege is not required.
- For Snapshot copy-based backup operation, ensure that all the tenant databases are valid and active.
- For pre and post commands for quiesce, Snapshot, and unquiesce operations, you should check if the commands exist in the command list available on the plug-in host from the following paths:
  - Default location on the Windows host: *C:\Program Files\NetApp\SnapCenter\Snapcenter Plug-in*

Creator\etc\allowed_commands.config

- Default location on the Linux host: */opt/NetApp/snapcenter/scc/etc/allowed_commands.config*

If the commands do not exist in the command list, then the operation will fail.

SnapCenter UI

Steps

1. In the left navigation pane, select **Resources**, and then select the appropriate plug-in from the list.
2. In the Resource page, filter resources from the **View** drop-down list based on resource type.

   Select , and then select the host name and the resource type to filter the resources. You can then select to close the filter pane.

3. Select the resource that you want to back up.
4. In the Resource page, select **Use custom name format for Snapshot copy**, and then enter a custom name format that you want to use for the Snapshot name.

   For example, *customtext_policy_hostname* or *resource_hostname*. By default, a timestamp is appended to the Snapshot name.

5. In the Application Settings page, do the following:
   - Select the **Backups** arrow to set additional backup options:

     Enable consistency group backup, if needed, and perform the following tasks:

For this field…	Do this…
Afford time to wait for "Consistency Group Snapshot" operation to complete	Select **Urgent**, or **Medium**, or **Relaxed** to specify the wait time for Snapshot operation to finish. Urgent = 5 seconds, Medium = 7 seconds, and Relaxed = 20 seconds.
Disable WAFL Sync	Select this to avoid forcing a WAFL consistency point.

   - Select the **Scripts** arrow to run pre and post commands for quiesce, Snapshot, and unquiesce operations.

     You can also run pre commands before exiting the backup operation. Prescripts and postscripts are run in the SnapCenter Server.

   - Select the **Custom Configurations**arrow, and then enter the custom value pairs required for all jobs using this resource.
   - Select the **Snapshot Copy Tool** arrow to select the tool to create Snapshots:

If you want…	Then…
SnapCenter to create a storage-level Snapshot	Select **SnapCenter without File System Consistency**.
SnapCenter to use the plug-in for Windows to put the file system into a consistent state and then create a Snapshot	Select **SnapCenter with File System Consistency**.

If you want…	Then…
To enter the command to create a Snapshot	Select **Other**, and then enter the command to create a Snapshot.

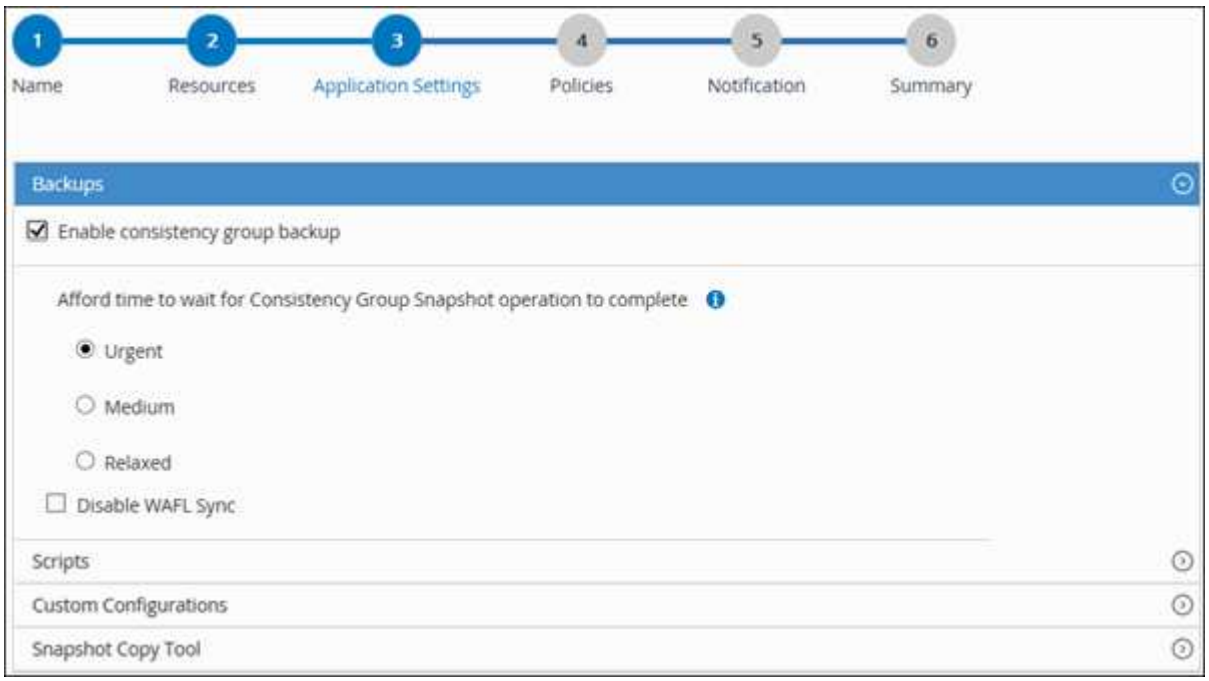

6. In the Policies page, perform the following steps:
   a. Select one or more policies from the drop-down list.

      You can also create a policy by clicking + .

      In the Configure schedules for selected policies section, the selected policies are listed.

   b. Select + in the Configure Schedules column for the policy for which you want to configure a schedule.
   c. In the Add schedules for policy *policy_name* dialog box, configure the schedule, and then select **OK**.

      policy_name is the name of the policy that you selected.

      The configured schedules are listed in the Applied Schedules column.

7. In the Notification page, from the **Email preference** drop-down list, select the scenarios in which you want to send the emails.

   You must also specify the sender and receiver email addresses, and the subject of the email. SMTP must also be configured in **Settings** > **Global Settings**.

8. Review the summary, and then select **Finish**.

   The resources topology page is displayed.

9. Select **Back up Now**.
10. In the Backup page, perform the following steps:
    a. If you applied multiple policies to the resource, from the **Policy** drop-down list, select the policy that you want to use for backup.

       If the policy selected for the on-demand backup is associated with a backup schedule, the on-demand backups will be retained based on the retention settings specified for the schedule type.

    b. Select **Backup**.
11. Monitor the operation progress by clicking **Monitor** > **Jobs**.
    - In MetroCluster configurations, SnapCenter might not be able to detect a protection relationship after a failover.

      For information, see: Unable to detect SnapMirror or SnapVault relationship after MetroCluster failover

PowerShell cmdlets

Steps

1. Initiate a connection session with the SnapCenter Server for a specified user by using the Open-SmConnection cmdlet.

   ```
   PS C:\> Open-SmConnection
   ```

   The username and password prompt is displayed.

2. Add manual resources by using the Add-SmResources cmdlet.

   This example shows how to add a IBM Db2 instance:

   ```
   PS C:\> Add-SmResource -HostName 10.32.212.13 -PluginCode DB2
   -ResourceType Instance -ResourceName db2inst1 -StorageFootPrint
   (@{"VolumeName"="windb201_data01";"LUNName"="windb201_data01";"Stora
   geSystem"="scsnfssvm"}) -MountPoints "D:\"
   ```

   For Db2 Database:

   ```
   PS C:\> Add-SmResource -HostName 10.32.212.13 -PluginCode DB2
   -ResourceType Database -ResourceName SALESDB -StorageFootPrint
   (@{"VolumeName"="windb201_data01";"LUNName"="windb201_data01";"Stora
   geSystem"="scsnfssvm"}) -MountPoints "D:\" -Instance DB2
   ```

3. Create a backup policy by using the Add-SmPolicy cmdlet.
4. Protect the resource or add a new resource group to SnapCenter by using the Add-SmResourceGroup cmdlet.

5. Initiate a new backup job by using the New-SmBackup cmdlet.

   This example shows how to backup a resource group:

   ```
   C:\PS> New-SMBackup -ResourceGroupName
   'ResourceGroup_with_Db2_Resources' -Policy db2_policy1
   ```

   This example backs up a Db2 instance:

   ```
   C:\PS> New-SMBackup -Resources
   @{"Host"="10.32.212.13";"Uid"="DB2INST1";"PluginName"="DB2"} -Policy
   db2_policy
   ```

   This example backs up a Db2 database:

   ```
   C:\PS> New-SMBackup -Resources
   @{"Host"="10.32.212.13";"Uid"="DB2INST1\WINARCDB";"PluginName"="DB2"
   } -Policy db2_policy
   ```

6. Monitor the job status (running, completed, or failed) by using the Get-smJobSummaryReport cmdlet.

   ```
   PS C:\> Get-SmJobSummaryReport -JobId 467

   SmJobId            : 467
   JobCreatedDateTime :
   JobStartDateTime   : 27-Jun-24 01:40:09
   JobEndDateTime     : 27-Jun-24 01:41:15
   JobDuration        : 00:01:06.7013330
   JobName            : Backup of Resource Group
   'SCDB201WIN_RAVIR1_OPENLAB_NETAPP_LOCAL_DB2_DB2_WINCIR' with policy
   'snapshot-based-db2'
   JobDescription     :
   Status             : Completed
   IsScheduled        : False
   JobError           :
   JobType            : Backup
   PolicyName         : db2_policy
   JobResultData      :
   ```

7. Monitor the backup job details like backup ID, backup name to perform restore or clone operation by using the Get-SmBackupReport cmdlet.

```
PS C:\> Get-SmBackupReport -JobId 467

BackedUpObjects           : {WINCIR}
FailedObjects             : {}
IsScheduled               : False
HasMetadata               : False
SmBackupId                : 84
SmJobId                   : 467
StartDateTime             : 27-Jun-24 01:40:09
EndDateTime               : 27-Jun-24 01:41:15
Duration                  : 00:01:06.7013330
CreatedDateTime           : 27-Jun-24 18:39:45
Status                    : Completed
ProtectionGroupName       : HOSTFQDN_DB2_DB2_WINCIR
SmProtectionGroupId       : 23
PolicyName                : db2_policy
SmPolicyId                : 13
BackupName                : HOSTFQDN _DB2_DB2_WINCIR_HOST_06-27-
2024_01.40.09.7397
VerificationStatus        : NotApplicable
VerificationStatuses      :
SmJobError                :
BackupType                : SCC_BACKUP
CatalogingStatus          : NotApplicable
CatalogingStatuses        :
ReportDataCreatedDateTime :
PluginCode                : SCC
PluginName                : DB2
PluginDisplayName         : IBM DB2
JobTypeId                 :
JobHost                   : HOSTFQDN
```

The information regarding the parameters that can be used with the cmdlet and their descriptions can be obtained by running *Get-Help command_name*. Alternatively, you can also refer to the SnapCenter Software Cmdlet Reference Guide.

Back up resource groups

A resource group is a collection of resources on a host. A backup operation on the resource group is performed on all resources defined in the resource group.

Before you begin

- You must have created a resource group with a policy attached.
- If you want to back up a resource that has a SnapMirror relationship with a secondary storage, the ONTAP

role assigned to the storage user should include the “snapmirror all” privilege. However, if you are using the “vsadmin” role, then the “snapmirror all” privilege is not required.

About this task

You can back up a resource group on demand from the Resources page. If a resource group has a policy attached and a schedule configured, then backups occur automatically according to the schedule.

Steps

1. In the left navigation pane, select **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page, select **Resource Group** from the **View** list.

   You can search the resource group either by entering the resource group name in the search box or by selecting , and then selecting the tag. You can then select to close the filter pane.
3. In the Resource Groups page, select the resource group that you want to back up, and then select **Back up Now**.
4. In the Backup page, perform the following steps:
   a. If you associated multiple policies with the resource group, from the **Policy** drop-down list, select the policy that you want to use for backup.

      If the policy selected for the on-demand backup is associated with a backup schedule, the on-demand backups will be retained based on the retention settings specified for the schedule type.

   b. Select **Backup**.
5. Monitor the operation progress by selecting **Monitor** > **Jobs**.

Monitor IBM Db2 backup operations

You can monitor the progress of different backup operations by using the SnapCenterJobs page. You might want to check the progress to determine when it is complete or if there is an issue.

About this task

The following icons appear on the Jobs page and indicate the corresponding state of the operations:

- In progress
- Completed successfully
- Failed
- Completed with warnings or could not start due to warnings
- Queued
- Canceled

Steps

1. In the left navigation pane, click **Monitor**.
2. In the Monitor page, click **Jobs**.
3. In the Jobs page, perform the following steps:

a. Click to filter the list so that only backup operations are listed.
b. Specify the start and end dates.
c. From the **Type** drop-down list, select **Backup**.
d. From the **Status** drop-down, select the backup status.
e. Click **Apply** to view the operations completed successfully.

4. Select a backup job, and then click **Details** to view the job details.

> Though the backup job status displays , when you click on job details you might see that some of the child tasks of the backup operation are still in progress or marked with warning signs.

5. In the Job Details page, click **View logs**.

   The **View logs** button displays the detailed logs for the selected operation.

Monitor data protection operations on IBM Db2 databases in the Activity pane

The Activity pane displays the five most recent operations performed. The Activity pane also displays when the operation was initiated and the status of the operation.

The Activity pane displays information regarding backup, restore, clone, and scheduled backup operations.

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. Click on the Activity pane to view the five most recent operations.

   When you click one of the operations, the operation details are listed in the **Job Details** page.

Cancel backup operations for IBM Db2

You can cancel backup operations that are queued.

What you will need

- You must be logged in as the SnapCenter Admin or job owner to cancel operations.
- You can cancel a backup operation from either the **Monitor** page or the **Activity** pane.
- You cannot cancel a running backup operation.
- You can use the SnapCenter GUI, PowerShell cmdlets, or CLI commands to cancel the backup operations.
- The **Cancel Job** button is disabled for operations that cannot be canceled.
- If you selected **All members of this role can see and operate on other members objects** in Users\Groups page while creating a role, you can cancel the queued backup operations of other members while using that role.

Steps

1. Perform one of the following actions:

From the…	Action
Monitor page	a. In the left navigation pane, click **Monitor** > **Jobs**. b. Select the operation, and then click **Cancel Job**.
Activity pane	a. After initiating the backup operation, click on the Activity pane to view the five most recent operations. b. Select the operation. c. In the Job Details page, click **Cancel Job**.

The operation is canceled, and the resource is reverted to the previous state.

View IBM Db2 backups and clones in the Topology page

When you are preparing to back up or clone a resource, you might find it helpful to view a graphical representation of all backups and clones on the primary and secondary storage.

About this task

You can review the following icons in the Manage Copies view to determine whether the backups and clones are available on the primary or secondary storage (Mirror copies or Vault copies).

- displays the number of backups and clones that are available on the primary storage.
- displays the number of backups and clones that are mirrored on the secondary storage using SnapMirror technology.
- displays the number of backups and clones that are replicated on the secondary storage using SnapVault technology.

The number of backups displayed includes the backups deleted from the secondary storage. For example, if you created 6 backups using a policy to retain only 4 backups, the number of backups displayed is 6.

Clones of a backup of a version-flexible mirror on a mirror-vault type volume are displayed in the topology view, but the mirror backup count in the topology view does not include the version-flexible backup.

In the Topology page, you can see all of the backups and clones that are available for the selected resource or resource group. You can view the details of those backups and clones, and then select them to perform data protection operations.

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page, either select the resource or resource group from the **View** drop-down list.
3. Select the resource either from the resource details view or from the resource group details view.

   If the resource is protected, the topology page of the selected resource is displayed.

4. Review the **Summary card** to see a summary of the number of backups and clones available on the primary and secondary storage.

   The **Summary Card** section displays the total number of Snapshot copy-based backups, and clones.

   Clicking the **Refresh** button starts a query of the storage to display an accurate count.

   If SnapLock enabled backup is taken, then clicking the **Refresh** button refreshes the primary and secondary SnapLock expiry time retrieved from ONTAP. A weekly schedule also refreshes the primary and secondary SnapLock expiry time retrieved from ONTAP.

   When the application resource is spread across multiple volumes, the SnapLock expiry time for the backup will be the longest SnapLock expiry time that is set for a Snapshot in a volume. The longest SnapLock expiry time is retrieved from ONTAP.

   After on demand backup, by clicking the **Refresh** button refreshes the details of backup or clone.

5. In the Manage Copies view, click **Backups** or **Clones** from the primary or secondary storage to see details of a backup or clone.

   The details of the backups and clones are displayed in a table format.

6. Select the backup from the table, and then click the data protection icons to perform restore, clone, and delete operations.

   You cannot rename or delete backups that are on the secondary storage.

7. If you want to delete a clone, select the clone from the table, and then click .
8. If you want to split a clone, select the clone from the table, and then click .

Restore IBM Db2

Restore workflow

The restore and recovery workflow includes planning, performing the restore operations, and monitoring the operations.

The following workflow shows the sequence in which you must perform the restore operation:

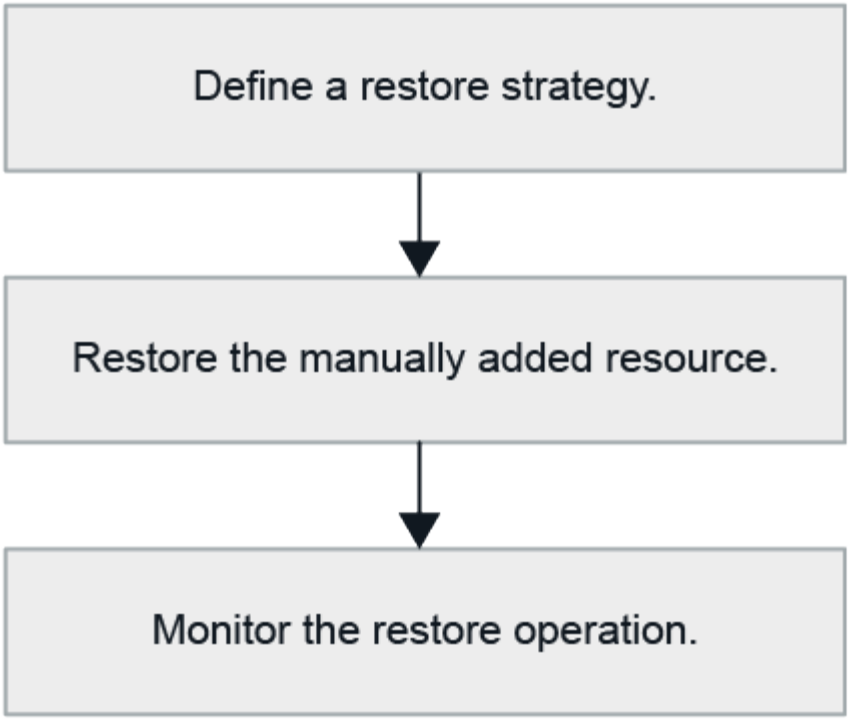

You can also use PowerShell cmdlets manually or in scripts to perform backup, restore, and clone operations. The SnapCenter cmdlet help and the cmdlet reference information contain detailed information about PowerShell cmdlets.

SnapCenter Software Cmdlet Reference Guide.

Restore a manually added resource backup

You can use SnapCenter to restore and recover data from one or more backups.

Before you begin

- You must have backed up the resource or resource groups.
- You must have canceled any backup operation that is currently in progress for the resource or resource group that you want to restore.
- For pre restore, post restore, mount, and unmount commands, you should check if the commands exist in the command list available on the plug-in host from the following paths:
  - Default location on the Windows host: *C:\Program Files\NetApp\SnapCenter\Snapcenter Plug-in Creator\etc\allowed_commands.config*
  - Default location on the Linux host: */opt/NetApp/snapcenter/scc/etc/allowed_commands.config*

If the commands do not exist in the command list, then the operation will fail.

About this task

- For ONTAP 9.12.1 and below version, the clones created from the SnapLock Vault Snapshots as part of restore will inherit the SnapLock Vault expiry time. Storage admin should manually cleanup the clones post the SnapLock expiry time.

SnapCenter UI

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page, filter resources from the **View** drop-down list based on resource type.

   The resources are displayed along with type, host, associated resource groups and policies, and status.

   > Although a backup might be for a resource group, when you restore, you must select the individual resources you want to restore.

   If the resource is not protected, “Not protected” is displayed in the Overall Status column. This can mean either that the resource is not protected, or that the resource was backed up by a different user.

3. Select the resource, or select a resource group and then select a resource in that group.

   The resource topology page is displayed.

4. From the Manage Copies view, select **Backups** either from the primary or secondary (mirrored or vaulted) storage systems.
5. In the Primary backup(s) table, select the backup that you want to restore from, and then click .

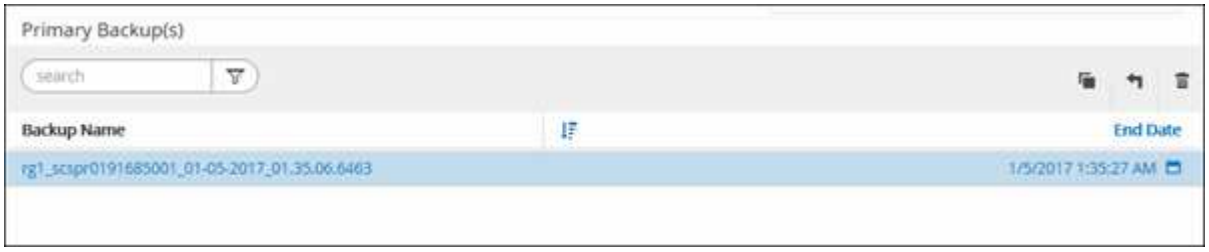

6. In the Restore Scope page, select **Complete Resource**.

   a. If you select **Complete Resource**, all of the configured data volumes of the IBM Db2 database are restored.

      If the resource contains volumes or qtrees, the Snapshots taken after the Snapshot selected for restore on such volumes or qtrees are deleted and cannot be recovered. Also, if any other resource is hosted on same volumes or qtrees, then that resource is also deleted.

      You can select multiple LUNs.

      > If you select **All**, all the files on the volumes, qtrees, or LUNs are restored.

7. In the Pre ops page, enter pre restore and unmount commands to run before performing a restore job.
8. In the Post ops page, enter mount and post restore commands to run after performing a restore job.
9. In the Notification page, from the **Email preference** drop-down list, select the scenarios in which you want to send the emails.

   You must also specify the sender and receiver email addresses and the subject of the email. SMTP must also be configured on the **Settings** > **Global Settings** page.

10. Review the summary, and then click **Finish**.
11. Monitor the operation progress by clicking **Monitor** > **Jobs**.

After you finish

Recovery is possible only if the Rollforward status is in "DB pending" status. This status is applicable to Db2 databases with archive logging enabled.

PowerShell cmdlets

Steps

1. Initiate a connection session with the SnapCenter Server for a specified user by using the Open-SmConnection cmdlet.

```
PS C:\> Open-SmConnection
```

2. Identify the backup that you want to restore by using the Get-SmBackup and Get-SmBackupReport cmdlets.

   This example shows that there are two backups available for the restore:

```
 PS C:\> Get-SmBackup -AppObjectId
cn24.sccore.test.com\DB2\db2inst1\Library

        BackupId    BackupName                    BackupTime
BackupType
        --------    ----------                    ----------
----------
         1          Payroll Dataset_vise-f6_08... 8/4/2015 11:02:32
AM  Full Backup
         2          Payroll Dataset_vise-f6_08... 8/4/2015 11:23:17
AM
```

   This example displays detailed information about the backup from January 29th 2015 to February 3rd, 2015:

```
PS C:\> Get-SmBackupReport -FromDateTime "1/29/2015" -ToDateTime
"2/3/2015"

SmBackupId           : 113
   SmJobId              : 2032
   StartDateTime        : 2/2/2015 6:57:03 AM
   EndDateTime          : 2/2/2015 6:57:11 AM
   Duration             : 00:00:07.3060000
   CreatedDateTime      : 2/2/2015 6:57:23 AM
   Status               : Completed
   ProtectionGroupName  : Clone
   SmProtectionGroupId  : 34
   PolicyName           : Vault
   SmPolicyId           : 18
   BackupName           : Clone_SCSPR0019366001_02-02-2015_06.57.08
   VerificationStatus   : NotVerified

   SmBackupId           : 114
   SmJobId              : 2183
   StartDateTime        : 2/2/2015 1:02:41 PM
   EndDateTime          : 2/2/2015 1:02:38 PM
   Duration             : -00:00:03.2300000
   CreatedDateTime      : 2/2/2015 1:02:53 PM
   Status               : Completed
   ProtectionGroupName  : Clone
   SmProtectionGroupId  : 34
   PolicyName           : Vault
   SmPolicyId           : 18
   BackupName           : Clone_SCSPR0019366001_02-02-2015_13.02.45
   VerificationStatus   : NotVerified
```

3. Restore data from the backup by using the Restore-SmBackup cmdlet.

> AppObjectId is "Host\Plugin\UID", where UID = <instance_name> is for manually discovered DB2 instance resource and UID = <instance_name>\<databse_name> is for IBM Db2 database resource. You can get the ResourceID from the Get-smResources cmdlet.

```
Get-smResources  -HostName cn24.sccore.test.com  -PluginCode DB2
```

This example shows how to restore the database from the primary storage:

```
Restore-SmBackup -PluginCode DB2 -AppObjectId
cn24.sccore.test.com\DB2\db2inst1\DB01 -BackupId 3
```

This example shows how to restore the database from the secondary storage:

```
Restore-SmBackup -PluginCode 'DB2' -AppObjectId
cn24.sccore.test.com\DB2\db2inst1\DB01 -BackupId 399 -Confirm:$false
-Archive @( @{"Primary"="<Primary
Vserver>:<PrimaryVolume>";"Secondary"="<Secondary
Vserver>:<SecondaryVolume>"})
```

The information regarding the parameters that can be used with the cmdlet and their descriptions can be obtained by running *Get-Help command_name*. Alternatively, you can also refer to the SnapCenter Software Cmdlet Reference Guide.

Restore and recover an auto discovered database backup

You can use SnapCenter to restore and recover data from one or more backups.

Before you begin

- You must have backed up the resource or resource groups.
- You must have canceled any backup operation that is currently in progress for the resource or resource group that you want to restore.
- For pre restore, post restore, mount, and unmount commands, you should check if the commands exist in the command list available on the plug-in host from the following paths:
  - Default location on the Windows host: *C:\Program Files\NetApp\SnapCenter\Snapcenter Plug-in Creator\etc\allowed_commands.config*
  - Default location on the Linux host: */opt/NetApp/snapcenter/scc/etc/allowed_commands.config*

If the commands do not exist in the command list, then the operation will fail.

About this task

- For Auto-discovered resources, restore is supported with SFSR.
- Auto-recovery is not supported.
- For ONTAP 9.12.1 and below version, the clones created from the SnapLock Vault snapshots as part of restore will inherit the SnapLock Vault expiry time. Storage admin should manually cleanup the clones post the SnapLock expiry time.

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page, filter resources from the **View** drop-down list based on resource type.

   The resources are displayed along with type, host, associated resource groups and policies, and status.

Although a backup might be for a resource group, when you restore, you must select the individual resources you want to restore.

If the resource is not protected, “Not protected” is displayed in the Overall Status column. This can mean either that the resource is not protected, or that the resource was backed up by a different user.

3. Select the resource, or select a resource group and then select a resource in that group.

   The resource topology page is displayed.

4. From the Manage Copies view, select **Backups** either from the primary or secondary (mirrored or vaulted) storage systems.
5. In the Primary backup(s) table, select the backup that you want to restore from, and then click .

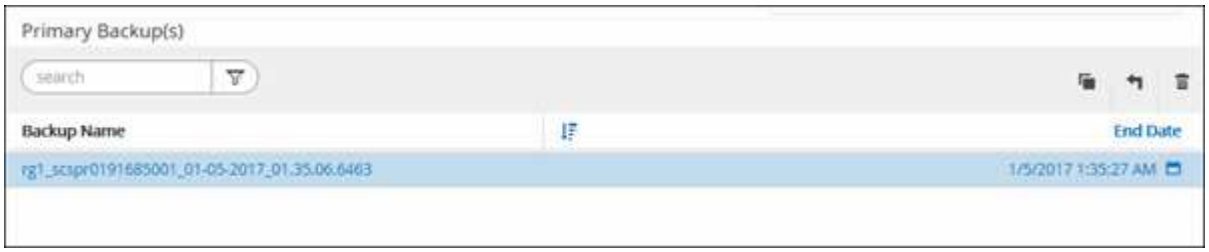

6. In the Restore Scope page, select **Complete Resource** to restore the configured data volumes of the IBM Db2 database.
7. In the Pre ops page, enter pre restore and unmount commands to run before performing a restore job.

   Unmount commands are not required for auto discovered resources.

8. In the Post ops page, enter mount and post restore commands to run after performing a restore job.

   Mount commands are not required for auto discovered resources.

9. In the Notification page, from the **Email preference** drop-down list, select the scenarios in which you want to send the emails.

   You must also specify the sender and receiver email addresses and the subject of the email. SMTP must also be configured on the **Settings** > **Global Settings** page.

10. Review the summary, and then click **Finish**.
11. Monitor the operation progress by clicking **Monitor** > **Jobs**.

After you finish

Recovery is possible only if the Rollforward status is in "DB pending" status. This status is applicable to Db2 databases with archive logging enabled.

Monitor IBM Db2 restore operations

You can monitor the progress of different SnapCenter restore operations by using the Jobs page. You might want to check the progress of an operation to determine when it is complete or if there is an issue.

About this task

Post-restore states describe the conditions of the resource after a restore operation and any further restore actions that you can take.

The following icons appear on the Jobs page, and indicate the state of the operation:

- In progress
- Completed successfully
- Failed
- Completed with warnings or could not start due to warnings
- Queued
- Canceled

Steps

1. In the left navigation pane, click **Monitor**.
2. In the **Monitor** page, click **Jobs**.
3. In the **Jobs** page, perform the following steps:
   a. Click to filter the list so that only restore operations are listed.
   b. Specify the start and end dates.
   c. From the **Type** drop-down list, select **Restore**.
   d. From the **Status** drop-down list, select the restore status.
   e. Click **Apply** to view the operations that have been completed successfully.
4. Select the restore job, and then click **Details** to view the job details.
5. In the **Job Details** page, click **View logs**.

   The **View logs** button displays the detailed logs for the selected operation.

Clone IBM Db2 resource backups

Clone workflow

The clone workflow includes performing the clone operation and monitoring the operation.

About this task

- You can clone on the source IBM Db2 server.
- You might clone resource backups for the following reasons:
  - To test functionality that has to be implemented using the current resource structure and content during application development cycles
  - For data extraction and manipulation tools when populating data warehouses
  - To recover data that was mistakenly deleted or changed

The following workflow shows the sequence in which you must perform the clone operation:

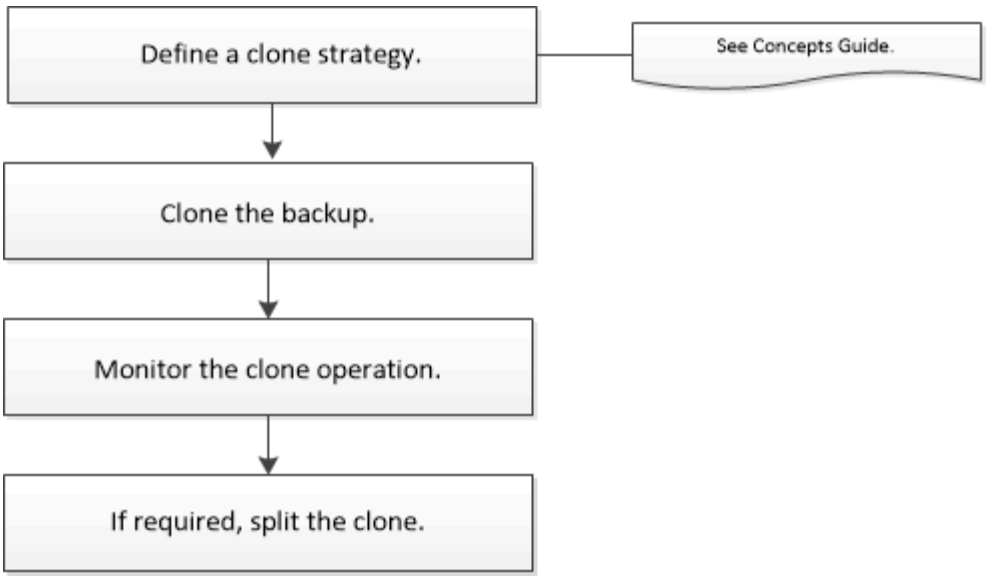

You can also use PowerShell cmdlets manually or in scripts to perform backup, restore, and clone operations. The SnapCenter cmdlet help and the cmdlet reference information contain detailed information about PowerShell cmdlets.

After you finish

After cloning the automatically discovered Db2 resources, the cloned resource is marked as manual resource. Click **Refresh resources** to recover the cloned Db2 resource. When you delete the clone, the storage and the host also gets cleaned up.

If you do not refresh the resources after the clone operation and attempt to delete the clone, the storage and the host will not be cleaned up. You must delete the entries manually in fstab.

Clone a IBM Db2 backup

You can use SnapCenter to clone a backup. You can clone from primary or secondary backup.

Before you begin

- You should have backed up the resources or resource group.
- You should ensure that the aggregates hosting the volumes should be in the assigned aggregates list of the storage virtual machine (SVM).
- While creating a clone for Db2 on an alternate host, you must create a n-1 directory structure for the clone mount path same as the original mount path on the other host. The mount path should have *755* execution permission.
- For pre clone or post clone commands, you should check if the commands exist in the command list available on the plug-in host from the following paths:
  - Default location on the Windows host: *C:\Program Files\NetApp\SnapCenter\Snapcenter Plug-in Creator\etc\allowed_commands.config*
  - Default location on the Linux host: */opt/NetApp/snapcenter/scc/etc/allowed_commands.config*

If the commands do not exist in the command list, then the operation will fail.

About this task

- For information about FlexClone volume split operations, see, Split a FlexClone volume from its parent volume.
- For ONTAP 9.12.1 and below version, the clones created from the SnapLock Vault Snapshots as part of restore will inherit the SnapLock Vault expiry time. Storage admin should manually cleanup the clones post the SnapLock expiry time.

SnapCenter UI

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page, filter resources from the **View** drop-down list based on resource type.

   The resources are displayed along with information such as type, host, associated resource groups and policies, and status.

3. Select the resource or resource group.

   You must select a resource if you select a resource group.

   The resource or resource group topology page is displayed.

4. From the Manage Copies view, select **Backups** either from the primary or secondary (mirrored or vaulted) storage systems.
5. Select the data backup from the table, and then click .
6. In the Location page, perform the following actions:

For this field…	Do this…
Clone server	Choose a host on which the clone should be created.
Target Clone Instance	Enter the target Db2 clone instance ID to clone from the existing backups. This is applicable only for ANF storage type resource.
Target Clone Name	Enter the name of the clone. This is applicable only for Db2 database resource.
NFS Export IP Address	Enter IP addresses or the host names on which the cloned volumes will be exported. This is applicable only to NFS storage type resource.
Capacity Pool Max. Throughput (MiB/s)	Enter the maximum throughput of a capacity pool.

7. In the Scripts page, perform the following steps:

   The scripts are run on the plug-in host.

a. Enter the commands for pre clone or post clone that should be run before or after the clone operation, respectively.
   - Pre clone command: delete existing databases with the same name
   - Post clone command: verify a database or start a database.

b. Enter the mount command to mount a file system to a host.

   Mount command for a volume or qtree on a Linux machine:

   Example for NFS: `mount VSERVER_DATA_IP:%VOLUME_NAME_Clone /mnt`

8. In the Notification page, from the **Email preference** drop-down list, select the scenarios in which you want to send the emails.

   You must also specify the sender and receiver email addresses, and the subject of the email.

9. Review the summary, and then click **Finish**.
10. Monitor the operation progress by clicking **Monitor** > **Jobs**.

After you finish

After cloning the automatically discovered Db2 resources, the cloned resource is marked as manual resource. Click **Refresh resources** to recover the cloned Db2 resource. When you delete the clone, the storage and the host also gets cleaned up.

If you do not refresh the resources after the clone operation and attempt to delete the clone, the storage and the host will not be cleaned up. You must delete the entries manually in fstab.

PowerShell cmdlets

Steps

1. Initiate a connection session with the SnapCenter Server for a specified user by using the Open-SmConnection cmdlet.

```
Open-SmConnection  -SMSbaseurl  https://snapctr.demo.netapp.com:8146
```

2. List the backups that can be cloned by using the Get-SmBackup or Get-SmResourceGroup cmdlet.

   This example displays information about all available backups:

```
C:\PS>PS C:\> Get-SmBackup

BackupId   BackupName                    BackupTime   BackupType
--------   ----------                    ----------   ----------
1          Payroll Dataset_vise-f6_08...  8/4/2015     Full Backup
                                          11:02:32 AM

2          Payroll Dataset_vise-f6_08...  8/4/2015
                                          11:23:17 AM
```

This example displays information about a specified resource group, its resources, and associated policies:

```
PS C:\> Get-SmResourceGroup -ListResources –ListPolicies

Description :
CreationTime : 8/4/2015 3:44:05 PM
ModificationTime : 8/4/2015 3:44:05 PM
EnableEmail : False
EmailSMTPServer :
EmailFrom :
EmailTo :
EmailSubject :
EnableSysLog : False
ProtectionGroupType : Backup
EnableAsupOnFailure : False
Policies : {FinancePolicy}
HostResourceMaping : {}
Configuration : SMCoreContracts.SmCloneConfiguration
LastBackupStatus :
VerificationServer :
EmailBody :
EmailNotificationPreference : Never
VerificationServerInfo : SMCoreContracts.SmVerificationServerInfo
SchedulerSQLInstance :
CustomText :
CustomSnapshotFormat :
SearchResources : False
ByPassCredential : False
IsCustomSnapshot :
MaintenanceStatus : Production
PluginProtectionGroupTypes : {SMSQL}
Name : Payrolldataset
Type : Group
Id : 1
Host :
UserName :
Passphrase :
Deleted : False
Auth : SMCoreContracts.SmAuth
IsClone : False
CloneLevel : 0
ApplySnapvaultUpdate : False
ApplyRetention : False
RetentionCount : 0
RetentionDays : 0
```

```
ApplySnapMirrorUpdate : False
SnapVaultLabel :
MirrorVaultUpdateRetryCount : 7
AppPolicies : {}
Description : FinancePolicy
PreScriptPath :
PreScriptArguments :
PostScriptPath :
PostScriptArguments :
ScriptTimeOut : 60000
DateModified : 8/4/2015 3:43:30 PM
DateCreated : 8/4/2015 3:43:30 PM
Schedule : SMCoreContracts.SmSchedule
PolicyType : Backup
PluginPolicyType : SMSQL
Name : FinancePolicy
Type :
Id : 1
Host :
UserName :
Passphrase :
Deleted : False
Auth : SMCoreContracts.SmAuth
IsClone : False
CloneLevel : 0
clab-a13-13.sddev.lab.netapp.com
DatabaseGUID :
SQLInstance : clab-a13-13
DbStatus : AutoClosed
DbAccess : eUndefined
IsSystemDb : False
IsSimpleRecoveryMode : False
IsSelectable : True
SqlDbFileGroups : {}
SqlDbLogFiles : {}
AppFileStorageGroups : {}
LogDirectory :
AgName :
Version :
VolumeGroupIndex : -1
IsSecondary : False
Name : TEST
Type : SQL Database
Id : clab-a13-13\TEST
Host : clab-a13-13.sddev.mycompany.com
UserName :
```

```
Passphrase :
Deleted : False
Auth : SMCoreContracts.SmAuth
IsClone : False
```

3. Initiate a clone operation from an existing backup by using the New-SmClone cmdlet.

   This example creates a clone from a specified backup with all logs:

```
PS C:\> New-SmClone
-BackupName payroll_dataset_vise-f3_08-05-2015_15.28.28.9774
-Resources @{"Host"="vise-f3.sddev.mycompany.com";
"Type"="SQL Database";"Names"="vise-f3\SQLExpress\payroll"}
-CloneToInstance vise-f3\sqlexpress -AutoAssignMountPoint
-Suffix _clonefrombackup
-LogRestoreType All -Policy clonefromprimary_ondemand

PS C:> New-SmBackup -ResourceGroupName PayrollDataset -Policy
FinancePolicy
```

   This example creates a clone to a specified Microsoft SQL Server instance:

```
PS C:\> New-SmClone
-BackupName "BackupDS1_NY-VM-SC-SQL_12-08-2015_09.00.24.8367"
-Resources @{"host"="ny-vm-sc-sql";"Type"="SQL Database";
"Names"="ny-vm-sc-sql\AdventureWorks2012_data"}
-AppPluginCode SMSQL -CloneToInstance "ny-vm-sc-sql"
-Suffix _CLPOSH -AssignMountPointUnderPath "C:\SCMounts"
```

4. View the status of the clone job by using the Get-SmCloneReport cmdlet.

   This example displays a clone report for the specified job ID:

```
PS C:\> Get-SmCloneReport -JobId 186

SmCloneId : 1
SmJobId : 186
StartDateTime : 8/3/2015 2:43:02 PM
EndDateTime : 8/3/2015 2:44:08 PM
Duration : 00:01:06.6760000
Status : Completed
ProtectionGroupName : Draper
SmProtectionGroupId : 4
PolicyName : OnDemand_Clone
SmPolicyId : 4
BackupPolicyName : OnDemand_Full_Log
SmBackupPolicyId : 1
CloneHostName : SCSPR0054212005.mycompany.com
CloneHostId : 4
CloneName : Draper__clone__08-03-2015_14.43.53
SourceResources : {Don, Betty, Bobby, Sally}
ClonedResources : {Don_DRAPER, Betty_DRAPER, Bobby_DRAPER,
                   Sally_DRAPER}
```

The information regarding the parameters that can be used with the cmdlet and their descriptions can be obtained by running *Get-Help command_name*. Alternatively, you can also refer to the SnapCenter Software Cmdlet Reference Guide.

Monitor IBM Db2 clone operations

You can monitor the progress of SnapCenter clone operations by using the Jobs page. You might want to check the progress of an operation to determine when it is complete or if there is an issue.

About this task

The following icons appear on the Jobs page, and indicate the state of the operation:

- In progress
- Completed successfully
- Failed
- Completed with warnings or could not start due to warnings
- Queued
- Canceled

Steps

1. In the left navigation pane, click **Monitor**.
2. In the **Monitor** page, click **Jobs**.
3. In the **Jobs** page, perform the following steps:
   a. Click to filter the list so that only clone operations are listed.
   b. Specify the start and end dates.
   c. From the **Type** drop-down list, select **Clone**.
   d. From the **Status** drop-down list, select the clone status.
   e. Click **Apply** to view the operations that are completed successfully.
4. Select the clone job, and then click **Details** to view the job details.
5. In the Job Details page, click **View logs**.

Split a clone

You can use SnapCenter to split a cloned resource from the parent resource. The clone that is split becomes independent of the parent resource.

About this task

- You cannot perform the clone split operation on an intermediate clone.

  For example, after you create clone1 from a database backup, you can create a backup of clone1, and then clone this backup (clone2). After you create clone2, clone1 is an intermediate clone, and you cannot perform the clone split operation on clone1. However, you can perform the clone split operation on clone2.

  After splitting clone2, you can perform the clone split operation on clone1 because clone1 is no longer the intermediate clone.

- When you split a clone, the backup copies and clone jobs of the clone are deleted.
- For information about FlexClone volume split operations, see, Split a FlexClone volume from its parent volume.
- Ensure that the volume or aggregate on the storage system is online.

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. In the **Resources** page, select the appropriate option from the View list:

Option	Description
For database applications	Select **Database** from the View list.
For file systems	Select **Path** from the View list.

3. Select the appropriate resource from the list.

   The resource topology page is displayed.

4. From the **Manage Copies** view, select the cloned resource (for example, the database or LUN), and then click .

5. Review the estimated size of the clone that is to be split and the required space available on the aggregate, and then click **Start**.
6. Monitor the operation progress by clicking **Monitor** > **Jobs**.

   The clone split operation stops responding if the SMCore service restarts. You should run the Stop-SmJob cmdlet to stop the clone split operation, and then retry the clone split operation.

   If you want a longer poll time or shorter poll time to check whether the clone is split or not, you can change the value of *CloneSplitStatusCheckPollTime* parameter in *SMCoreServiceHost.exe.config* file to set the time interval for SMCore to poll for the status of the clone split operation. The value is in milliseconds and the default value is 5 minutes.

   For example:

   ```
   <add key="CloneSplitStatusCheckPollTime" value="300000" />
   ```

   The clone split start operation fails if backup, restore, or another clone split is in progress. You should restart the clone split operation only after the running operations are complete.

Related information

SnapCenter clone or verification fails with aggregate does not exist

Delete or split IBM Db2 database clones after upgrading SnapCenter

After upgrading to SnapCenter 4.3, you will no longer see the clones. You can delete the clone or split the clones from the Topology page of the resource from which the clones were created.

About this task

If you want to locate the storage footprint of the hidden clones, run the following command: `Get-SmClone -ListStorageFootprint`

Steps

1. Delete the backups of the cloned resources by using the remove-smbackup cmdlet.
2. Delete the resource group of the cloned resources by using the remove-smresourcegroup cmdlet.
3. Remove the protection of the cloned resource by using the remove-smprotectresource cmdlet.
4. Select the parent resource from the Resources page.

   The resource topology page is displayed.

5. From the Manage Copies view, select the clones either from the primary or secondary (mirrored or replicated) storage systems.
6. Select the clones, and then click to delete clones or click to split the clones.
7. Click **OK**.

Protect PostgreSQL

SnapCenter Plug-in for PostgreSQL

SnapCenter Plug-in for PostgreSQL overview

The SnapCenter Plug-in for PostgreSQL cluster is a host-side component of the NetApp SnapCenter software that enables application-aware data protection management of PostgreSQL clusters. The Plug-in for PostgreSQL cluster automates the backup, restore, and cloning of PostgreSQL clusters in your SnapCenter environment.

SnapCenter supports single cluster and multi cluster PostgreSQL setups. You can use the Plug-in for PostgreSQL Clusters in both Linux and Windows environments. In Windows environments, PostgreSQL will be supported as manual resource.

When the Plug-in for PostgreSQL cluster is installed, you can use SnapCenter with NetApp SnapMirror technology to create mirror copies of backup sets on another volume. You can also use the plug-in with NetApp SnapVault technology to perform disk-to-disk backup replication for standards compliance.

SnapCenter Plug-in for PostgreSQL supports NFS and SAN on ONTAP and Azure NetApp File storage layouts.

VMDK, vVol, and RDM virtual storage layout is supported.

What you can do using the SnapCenter Plug-in for PostgreSQL

When you install the Plug-in for PostgreSQL cluster in your environment, you can use SnapCenter to back up, restore, and clone PostgreSQL clusters and their resources. You can also perform tasks supporting those operations.

- Add clusters.
- Create backups.
- Restore from backups.
- Clone backups.
- Schedule backup operations.
- Monitor backup, restore, and clone operations.
- View reports for backup, restore, and clone operations.

SnapCenter Plug-in for PostgreSQL features

SnapCenter integrates with the plug-in application and with NetApp technologies on the storage system. To work with the Plug-in for PostgreSQL Cluster, you use the SnapCenter graphical user interface.

- **Unified graphical user interface**

  The SnapCenter interface provides standardization and consistency across plug-ins and environments. The SnapCenter interface enables you to complete consistent backup, restore, and clone operations

across plug-ins, use centralized reporting, use at-a-glance dashboard views, set up role-based access control (RBAC), and monitor jobs across all plug-ins.

- **Automated central administration**

  You can schedule backup operations, configure policy-based backup retention, and perform restore operations. You can also proactively monitor your environment by configuring SnapCenter to send email alerts.

- **Nondisruptive NetApp snapshot copy technology**

  SnapCenter uses NetApp snapshot technology with the Plug-in for PostgreSQL cluster to back up resources.

Using the Plug-in for PostgreSQL also offers the following benefits:

- Support for backup, restore, and clone workflows
- RBAC-supported security and centralized role delegation

  You can also set the credentials so that the authorized SnapCenter users have application-level permissions.

- Creation of space-efficient and point-in-time copies of resources for testing or data extraction by using NetApp FlexClone technology

  A FlexClone license is required on the storage system where you want to create the clone.

- Support for the consistency group (CG) snapshot feature of ONTAP as part of creating backups.
- Capability to run multiple backups simultaneously across multiple resource hosts

  In a single operation, snapshots are consolidated when resources in a single host share the same volume.

- Capability to create snapshots using external commands.
- Support for Linux LVM on XFS file system.

Storage types supported by SnapCenter Plug-in for PostgreSQL

SnapCenter supports a wide range of storage types on both physical machines and virtual machines (VMs). You must verify the support for your storage type before installing SnapCenter Plug-in for PostgreSQL.

Machine	Storage type
Physical server	• FC-connected LUNs • iSCSI-connected LUNs • NFS-connected volumes

Machine	Storage type
VMware ESXi	• RDM LUNs connected by an FC or iSCSI ESXi HBAScanning of host bus adapters (HBAs) might take long time to complete because SnapCenter scans all the host bus adaptors present in the host. You can edit the **LinuxConfig.pm** file located at */opt/NetApp/snapcenter/spl/plugins/scu/scucore/modules/SCU/Config* to set the value of the **SCSI_HOSTS_OPTIMIZED_RESCAN** parameter to 1 to rescan only those HBAs that are listed in HBA_DRIVER_NAMES. • iSCSI LUNs connected directly to the guest system by the iSCSI initiator • VMDKs on NFS datastores • VMDKs on VMFS • NFS volumes connected directly to the guest system • vVol datastores on both NFS and SAN vVol datastore can only be provisioned with ONTAP Tools for VMware vSphere.

Minimum ONTAP privileges required for PostgreSQL plug-in

The minimum ONTAP privileges that are required vary according to the SnapCenter plug-ins you are using for data protection.

- All-access commands: Minimum privileges required for ONTAP 9.12.1 and later
  - event generate-autosupport-log
  - job history show
  - job stop
  - lun
  - lun create
  - lun create
  - lun create
  - lun delete
  - lun igroup add
  - lun igroup create
  - lun igroup delete
  - lun igroup rename

- lun igroup rename
- lun igroup show
- lun mapping add-reporting-nodes
- lun mapping create
- lun mapping delete
- lun mapping remove-reporting-nodes
- lun mapping show
- lun modify
- lun move-in-volume
- lun offline
- lun online
- lun persistent-reservation clear
- lun resize
- lun serial
- lun show
- snapmirror policy add-rule
- snapmirror policy modify-rule
- snapmirror policy remove-rule
- snapmirror policy show
- snapmirror restore
- snapmirror show
- snapmirror show-history
- snapmirror update
- snapmirror update-ls-set
- snapmirror list-destinations
- version
- volume clone create
- volume clone show
- volume clone split start
- volume clone split stop
- volume create
- volume destroy
- volume file clone create
- volume file show-disk-usage
- volume offline
- volume online
- volume modify

    - volume qtree create
    - volume qtree delete
    - volume qtree modify
    - volume qtree show
    - volume restrict
    - volume show
    - volume snapshot create
    - volume snapshot delete
    - volume snapshot modify
    - volume snapshot modify-snaplock-expiry-time
    - volume snapshot rename
    - volume snapshot restore
    - volume snapshot restore-file
    - volume snapshot show
    - volume unmount
    - vserver cifs
    - vserver cifs share create
    - vserver cifs share delete
    - vserver cifs shadowcopy show
    - vserver cifs share show
    - vserver cifs show
    - vserver export-policy
    - vserver export-policy create
    - vserver export-policy delete
    - vserver export-policy rule create
    - vserver export-policy rule show
    - vserver export-policy show
    - vserver iscsi
    - vserver iscsi connection show
    - vserver show
- Read-only commands: Minimum privileges required for ONTAP 8.3.0 and later
    - network interface
    - network interface show
    - vserver

Prepare storage systems for SnapMirror and SnapVault replication for PostgreSQL

You can use a SnapCenter plug-in with ONTAP SnapMirror technology to create mirror

copies of backup sets on another volume, and with ONTAP SnapVault technology to perform disk-to-disk backup replication for standards compliance and other governance-related purposes. Before you perform these tasks, you must configure a data-protection relationship between the source and destination volumes and initialize the relationship.

SnapCenter performs the updates to SnapMirror and SnapVault after it completes the Snapshot operation. SnapMirror and SnapVault updates are performed as part of the SnapCenter job; do not create a separate ONTAP schedule.

If you are coming to SnapCenter from a NetApp SnapManager product and are satisfied with the data protection relationships you have configured, you can skip this section.

A data protection relationship replicates data on primary storage (the source volume) to secondary storage (the destination volume). When you initialize the relationship, ONTAP transfers the data blocks referenced on the source volume to the destination volume.

SnapCenter does not support cascade relationships between SnapMirror and SnapVault volumes (**Primary** > **Mirror** > **Vault**). You should use fanout relationships.

SnapCenter supports the management of version-flexible SnapMirror relationships. For details about version-flexible SnapMirror relationships and how to set them up, see the ONTAP documentation.

Backup strategy for PostgreSQL

Define a backup strategy for PostgreSQL

Defining a backup strategy before you create your backup jobs helps you to have the backups that you require to successfully restore or clone your resources. Your service-level agreement (SLA), recovery time objective (RTO), and recovery point objective (RPO) largely determine your backup strategy.

About this task

An SLA defines the level of service that is expected and addresses many service-related issues, including the availability and performance of the service. RTO is the time by which a business process must be restored after a disruption in service. RPO defines the strategy for the age of the files that must be recovered from backup storage for regular operations to resume after a failure. SLA, RTO, and RPO contribute to the data protection strategy.

Steps

1. Determine when you should back up your resources.
2. Decide how many backup jobs you require.
3. Decide how to name your backups.
4. Decide whether you want to create a Snapshot copy-based policy to back up application-consistent snapshots of the cluster.
5. Decide whether you want to use NetApp SnapMirror technology for replication or NetApp SnapVault technology for long-term retention.
6. Determine the retention period for the snapshots on the source storage system and the SnapMirror destination.

7. Determine whether you want to run any commands before or after the backup operation and provide a prescript or postscript.

Automatic discovery of resources on Linux host

Resources are PostgreSQL clusters and instances on the Linux host that are managed by SnapCenter. After installing the SnapCenter Plug-in for PostgreSQL plug-in, the PostgreSQL clusters from all the instances on that Linux host are automatically discovered and displayed in the Resources page.

Type of backups supported

Backup type specifies the type of backup that you want to create. SnapCenter supports snapshot copy-based backup type for PostgreSQL clusters.

Snapshot copy based backup

Snapshot copy-based backups leverage NetApp snapshot technology to create online, read-only copies of the volumes on which the PostgreSQL clusters reside.

How SnapCenter Plug-in for PostgreSQL uses consistency group snapshots

You can use the plug-in to create consistency group snapshots for resource groups. A consistency group is a container that can house multiple volumes so that you can manage them as one entity. A consistency group is simultaneous snapshots of multiple volumes, providing consistent copies of a group of volumes.

You can also specify the wait time for the storage controller to consistently group snapshots. The available wait time options are **Urgent**, **Medium**, and **Relaxed**. You can also enable or disable Write Anywhere File Layout (WAFL) sync during consistent group snapshot operation. WAFL sync improves the performance of a consistency group snapshot.

How SnapCenter manages housekeeping of data backups

SnapCenter manages the housekeeping of data backups on the storage system and file system levels.

The snapshots on the primary or secondary storage and their corresponding entries in the PostgreSQL catalog are deleted based on the retention settings.

Considerations for determining backup schedules for PostgreSQL

The most critical factor in determining a backup schedule is the rate of change for the resource. You might back up a heavily used resource every hour, while you might back up a rarely used resource once a day. Other factors include the importance of the resource to your organization, your service-level agreement (SLA) and your recovery point objective (RPO).

Backup schedules have two parts, as follows:

- Backup frequency (how often backups are to be performed)

  Backup frequency, also called schedule type for some plug-ins, is part of a policy configuration. For example, you might configure the backup frequency as hourly, daily, weekly, or monthly.

- Backup schedules (exactly when backups are to be performed)

  Backup schedules are part of a resource or resource group configuration. For example, if you have a resource group that has a policy configured for weekly backups, you might configure the schedule to back up every Thursday at 10:00 p.m.

Number of backup jobs needed for PostgreSQL

Factors that determine the number of backup jobs that you need include the size of the resource, the number of volumes used, the rate of change of the resource, and your Service Level Agreement (SLA).

Backup naming conventions for Plug-in for PostgreSQL clusters

You can either use the default Snapshot naming convention or use a customized naming convention. The default backup naming convention adds a timestamp to Snapshot names that helps you identify when the copies were created.

The Snapshot uses the following default naming convention:

```
resourcegroupname_hostname_timestamp
```

You should name your backup resource groups logically, as in the following example:

```
dts1_mach1x88_03-12-2015_23.17.26
```

In this example, the syntax elements have the following meanings:

- *dts1* is the resource group name.
- *mach1x88* is the host name.
- *03-12-2015_23.17.26* is the date and timestamp.

Alternatively, you can specify the Snapshot name format while protecting resources or resource groups by selecting **Use custom name format for Snapshot copy**. For example, customtext_resourcegroup_policy_hostname or resourcegroup_hostname. By default, the time stamp suffix is added to the Snapshot name.

Restore and recovery strategy for PostgreSQL

Define a restore and recovery strategy for PostgreSQL resources

You must define a strategy before you restore and recover your cluster so that you can perform restore and recovery operations successfully.

Only manual recovery of cluster is supported.

Steps

1. Determine the restore strategies supported for manually added PostgreSQL resources
2. Determine the restore strategies supported for auto discovered PostgreSQL clusters
3. Decide the type of recovery operations that you want to perform.

Types of restore strategies supported for manually added PostgreSQL resources

You must define a strategy before you can successfully perform restore operations using SnapCenter.

You cannot recover manually added PostgreSQL resources.

Complete resource restore

- Restores all volumes, qtrees, and LUNs of a resource

If the resource contains volumes or qtrees, the snapshots taken after the snapshot selected for restore on such volumes or qtrees are deleted and cannot be recovered. Also, if any other resource is hosted on the same volumes or qtrees, then that resource is also deleted.

NOTE:
Plug-in for PostgreSQL creates a backup_label and tablespace_map in */<OS_temp_folder>/postgresql_sc_recovery*<Restore_JobId>/_ folder to help recover manually .

Type of restore strategy supported for automatically discovered PostgreSQL

You must define a strategy before you can successfully perform restore operations using SnapCenter.

Complete resource restore is the restore strategy supported for automatically discovered PostgreSQL clusters. This restores all the volumes, qtrees, and LUNs of a resource.

Types of restore operations for auto discovered PostgreSQL

SnapCenter Plug-in for PostgreSQL supports Single File SnapRestore, and connect-and-copy restore types for automatically discovered PostgreSQL clusters.

Single File SnapRestore is performed in NFS environments for the following scenarios:

- If only the **Complete Resource** option is selected
- When the backup selected is from a SnapMirror or SnapVault secondary location, and the **Complete Resource** option is selected

Single File SnapRestore is performed in SAN environments for the following scenarios:

- If only the **Complete Resource** option is selected

- When the backup is selected from a SnapMirror or SnapVault secondary location, and the **Complete Resource** option is selected

Types of recovery operations supported for PostgreSQL clusters

SnapCenter enables you to perform different types of recovery operations for PostgreSQL clusters.

- Recover the cluster up to the most recent state
- Recover the cluster up to a specific point in time

  You must specify the date and time for recovery.

SnapCenter also provides the No recovery option for PostgreSQL clusters.

Prepare to install the SnapCenter Plug-in for PostgreSQL

Installation workflow of SnapCenter Plug-in for PostgreSQL

You should install and set up the SnapCenter Plug-in for PostgreSQL if you want to protect PostgreSQL clusters.

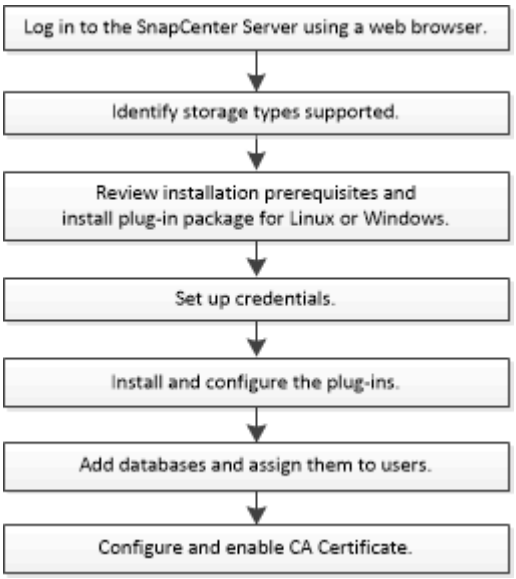

Prerequisites to add hosts and install SnapCenter Plug-in for PostgreSQL

Before you add a host and install the plug-in packages, you must complete all the requirements. SnapCenter Plug-in for PostgreSQL is available in both Windows and Linux environments.

- You must have installed Java 11 on your host.

  IBM Java is not supported on Windows and Linux host.

- For Windows, plug-in Creator Service should be running using the “LocalSystem” windows user, which is

the default behavior when Plug-in for PostgreSQL is installed as domain administrator.

- When installing a plug-in on a Windows host, if you specify a credential that is not built-in or if the user belongs to a local workgroup user, you must disable UAC on the host. SnapCenter Plug-in for Microsoft Windows will be deployed by default with the PostgreSQL plug-in on Windows hosts.
- SnapCenter Server should have access to the 8145 or custom port of Plug-in for PostgreSQL host.

Windows hosts

- You must have a domain user with local administrator privileges with local login permissions on the remote host.
- While installing Plug-in for PostgreSQL on a Windows host, SnapCenter Plug-in for Microsoft Windows is installed automatically.
- You must have enabled the password-based SSH connection for the root or non-root user.
- You must have installed Java 11 on your Windows host.

  Download JAVA for All Operating Systems

  NetApp Interoperability Matrix Tool

Linux hosts

- You must have enabled the password-based SSH connection for the root or non-root user.
- You must have installed Java 11 on your Linux host.

  Download JAVA for All Operating Systems

  NetApp Interoperability Matrix Tool

- For PostgreSQL clusters that are running on a Linux host, while installing Plug-in for PostgreSQL, SnapCenter Plug-in for UNIX is installed automatically.
- You should have **bash** as the default shell for plug-in installation.

Supplemental commands

To run a supplemental command on the SnapCenter Plug-in for PostgreSQL, you must include it in the *allowed_commands.config* file.

- Default location on the Windows host: *C:\Program Files\NetApp\SnapCenter\Snapcenter Plug-in Creator\etc\allowed_commands.config*
- Default location on the Linux host: */opt/NetApp/snapcenter/scc/etc/allowed_commands.config*

To allow supplemental commands on the plug-in host, open *allowed_commands.config* file in an editor. Enter each command on a separate line and the commands are not case sensitive. Ensure that you specify the fully qualified pathname and enclose the pathname in quotation marks (") if it contains spaces.

For example:

```
command: mount
command: umount
command: "C:\Program Files\NetApp\SnapCreator commands\sdcli.exe"
command: myscript.bat
```

If the *allowed_commands.config* file is not present, the commands or script execution will be blocked and the workflow will fail with the following error:

"[/mnt/mount -a] execution not allowed. Authorize by adding the command in the file %s on the plugin host."

If the command or script is not present in the *allowed_commands.config*, the command or script execution will be blocked and the workflow will fail with the following error:

"[/mnt/mount -a] execution not allowed. Authorize by adding the command in the file %s on the plugin host."

You should not use a wildcard entry (*) to allow all commands.

Configure sudo privileges for non-root users for Linux host

SnapCenter allows a non-root user to install the SnapCenter Plug-ins Package for Linux and to start the plug-in process. The plug-in processes will be running as an effective non-root user. You should configure sudo privileges for the non-root user to provide access to several paths.

What you will need

- Sudo version 1.8.7 or later.
- If the umask is 0027, ensure that the java folder and all the files inside should have permission of 555. Otherwise plug-in installation might fail.
- For the non-root user, ensure that the name of the non-root user and the user's group should be the same.
- Edit the */etc/ssh/sshd_config* file to configure the message authentication code algorithms: MACs hmac-sha2-256 and MACs hmac-sha2-512.

  Restart the sshd service after updating the configuration file.

  Example:

```
#Port 22
#AddressFamily any
#ListenAddress 0.0.0.0
#ListenAddress ::
#Legacy changes
#KexAlgorithms diffie-hellman-group1-sha1
#Ciphers aes128-cbc
#The default requires explicit activation of protocol
Protocol 2
HostKey/etc/ssh/ssh_host_rsa_key
MACs hmac-sha2-256
```

About this task

You should configure sudo privileges for the non-root user to provide access to the following paths:

- /home/*LINUX_USER*/.sc_netapp/snapcenter_linux_host_plugin.bin

- /custom_location/NetApp/snapcenter/spl/installation/plugins/uninstall
- /custom_location/NetApp/snapcenter/spl/bin/spl

Steps

1. Log in to the Linux host on which you want to install the SnapCenter Plug-ins Package for Linux.
2. Add the following lines to the /etc/sudoers file by using the visudo Linux utility.

```
Cmnd_Alias HPPLCMD = sha224:checksum_value== /home/
LINUX_USER/.sc_netapp/snapcenter_linux_host_plugin.bin,
/opt/NetApp/snapcenter/spl/installation/plugins/uninstall,
/opt/NetApp/snapcenter/spl/bin/spl, /opt/NetApp/snapcenter/scc/bin/scc
Cmnd_Alias PRECHECKCMD = sha224:checksum_value== /home/
LINUX_USER/.sc_netapp/Linux_Prechecks.sh
Cmnd_Alias CONFIGCHECKCMD = sha224:checksum_value==
/opt/NetApp/snapcenter/spl/plugins/scu/scucore/configurationcheck/Config
_Check.sh
Cmnd_Alias SCCMD = sha224:checksum_value==
/opt/NetApp/snapcenter/spl/bin/sc_command_executor
Cmnd_Alias SCCCMDEXECUTOR =checksum_value==
/opt/NetApp/snapcenter/scc/bin/sccCommandExecutor
LINUX_USER ALL=(ALL) NOPASSWD:SETENV: HPPLCMD, PRECHECKCMD,
CONFIGCHECKCMD, SCCCMDEXECUTOR, SCCMD
Defaults: LINUX_USER env_keep += "IATEMPDIR"
Defaults: LINUX_USER env_keep += "JAVA_HOME"
Defaults: LINUX_USER !visiblepw
Defaults: LINUX_USER !requiretty
```

LINUX_USER is the name of the non-root user that you created.

You can obtain the *checksum_value* from the **sc_unix_plugins_checksum.txt** file, which is located at:

- _C:\ProgramData\NetApp\SnapCenter\Package Repository\sc_unix_plugins_checksum.txt
_ if SnapCenter Server is installed on Windows host.
- _/opt/NetApp/snapcenter/SnapManagerWeb/Repository/sc_unix_plugins_checksum.txt
_ if SnapCenter Server in installed on Linux host.

The example should be used only as a reference for creating your own data.

Host requirements to install SnapCenter Plug-ins Package for Windows

Before you install the SnapCenter Plug-ins Package for Windows, you should be familiar with some basic host system space requirements and sizing requirements.

Item	Requirements
Operating Systems	Microsoft Windows For the latest information about supported versions, see the NetApp Interoperability Matrix Tool.
Minimum RAM for the SnapCenter plug-in on host	1 GB
Minimum install and log space for the SnapCenter plug-in on host	5 GB You should allocate sufficient disk space and monitor the storage consumption by the logs folder. The log space required varies depending on the number of the entities to be protected and the frequency of data protection operations. If there is no sufficient disk space, the logs will not be created for the recently run operations.
Required software packages	• ASP.NET Core Runtime 8.0.12 (and all subsequent 8.0.x patches) Hosting Bundle • PowerShell Core 7.4.2 For the latest information about supported versions, see the NetApp Interoperability Matrix Tool. For .NET specific troubleshooting information, see SnapCenter upgrade or install fails for legacy systems that do not have internet connectivity.

Host requirements for installing the SnapCenter Plug-ins Package for Linux

Before you install the SnapCenter Plug-ins Package for Linux, you should be familiar with some basic host system space and sizing requirements.

Item	Requirements
Operating systems	• Red Hat Enterprise Linux • SUSE Linux Enterprise Server (SLES) For the latest information about supported versions, see the NetApp Interoperability Matrix Tool.
Minimum RAM for the SnapCenter plug-in on host	1 GB

Item	Requirements
Minimum install and log space for the SnapCenter plug-in on host	2 GB You should allocate sufficient disk space and monitor the storage consumption by the logs folder. The log space required varies, depending on the number of the entities to be protected and the frequency of data protection operations. If there is no sufficient disk space, the logs will not be created for the recently run operations.
Required software packages	Java 11 Oracle Java and OpenJDK If you have upgraded JAVA to the latest version, you must ensure that the JAVA_HOME option located at /var/opt/snapcenter/spl/etc/spl.properties is set to the correct JAVA version and the correct path. For the latest information about supported versions, see the NetApp Interoperability Matrix Tool.

Set up credentials for the SnapCenter Plug-in for PostgreSQL

SnapCenter uses credentials to authenticate users for SnapCenter operations. You should create credentials for installing SnapCenter plug-ins and additional credentials for performing data protection operations on clusters or Windows file systems.

About this task

- Linux hosts

  You must set up credentials for installing plug-ins on Linux hosts.

  You must set up the credentials for the root user or for a non-root user who has sudo privileges to install and start the plug-in process.

  Best Practice: Although you are allowed to create credentials for Linux after deploying hosts and installing plug-ins, the best practice is to create credentials after you add SVMs, before you deploy hosts and install plug-ins.

- Windows hosts

  You must set up Windows credentials before installing plug-ins.

  You must set up the credentials with administrator privileges, including administrator rights on the remote host.

If you set up credentials for individual resource groups and the username does not have full admin privileges,

you must assign at least the resource group and backup privileges to the username.

Steps

1. In the left navigation pane, click **Settings**.
2. In the Settings page, click **Credential**.
3. Click **New**.
4. In the Credential page, specify the information required for configuring credentials:

For this field…	Do this…
Credential name	Enter a name for the credentials.
User name	Enter the user name and password that are to be used for authentication. • Domain administrator or any member of the administrator group Specify the domain administrator or any member of the administrator group on the system on which you are installing the SnapCenter plug-in. Valid formats for the Username field are: ◦ *NetBIOS\UserName* ◦ *Domain FQDN\UserName* • Local administrator (for workgroups only) For systems that belong to a workgroup, specify the built-in local administrator on the system on which you are installing the SnapCenter plug-in. You can specify a local user account that belongs to the local administrators group if the user account has elevated privileges or the User Access control feature is disabled on the host system. The valid format for the Username field is: *UserName* Do not use double quotes (") or backtick (`) in the passwords. You should not use the less than (<) and exclamation (!) symbols together in passwords. For example, lessthan<!10, lessthan10<!, backtick`12.
Password	Enter the password used for authentication.
Authentication Mode	Select the authentication mode that you want to use.

For this field…	Do this…
Use sudo privileges	Select the **Use sudo privileges** check box if you are creating credentials for a non-root user. Applicable to Linux users only.

5. Click **OK**.

After you finish setting up the credentials, you might want to assign credential maintenance to a user or group of users in the User and Access page.

Configure gMSA on Windows Server 2016 or later

Windows Server 2016 or later enables you to create a group Managed Service Account (gMSA) that provides automated service account password management from a managed domain account.

Before you begin

- You should have a Windows Server 2016 or later domain controller.
- You should have a Windows Server 2016 or later host, which is a member of the domain.

Steps

1. Create a KDS root key to generate unique passwords for each object in your gMSA.
2. For each domain, run the following command from the Windows domain controller: Add-KDSRootKey -EffectiveImmediately
3. Create and configure your gMSA:
   a. Create a user group account in the following format:

      ```
      domainName\accountName$
      ```

   b. Add computer objects to the group.
   c. Use the user group you just created to create the gMSA.

      For example,

      ```
      New-ADServiceAccount -name <ServiceAccountName> -DNSHostName <fqdn>
      -PrincipalsAllowedToRetrieveManagedPassword <group>
      -ServicePrincipalNames <SPN1,SPN2,…>
      ```

   d. Run `Get-ADServiceAccount` command to verify the service account.
4. Configure the gMSA on your hosts:
   a. Enable the Active Directory module for Windows PowerShell on the host where you want to use the gMSA account.

To do this, run the following command from PowerShell:

```
PS C:\> Get-WindowsFeature AD-Domain-Services

Display Name                                Name                Install
State
------------                                ----
-------------
[ ] Active Directory Domain Services   AD-Domain-Services  Available


PS C:\> Install-WindowsFeature AD-DOMAIN-SERVICES

Success Restart Needed Exit Code      Feature Result
------- -------------- ---------      --------------
True    No             Success        {Active Directory Domain
Services, Active ...
WARNING: Windows automatic updating is not enabled. To ensure that
your newly-installed role or feature is
automatically updated, turn on Windows Update.
```

b. Restart your host.

c. Install the gMSA on your host by running the following command from the PowerShell command prompt: `Install-AdServiceAccount <gMSA>`

d. Verify your gMSA account by running the following command: `Test-AdServiceAccount <gMSA>`

5. Assign the administrative privileges to the configured gMSA on the host.
6. Add the Windows host by specifying the configured gMSA account in the SnapCenter Server.

   SnapCenter Server will install the selected plug-ins on the host and the specified gMSA will be used as the service log on account during the plug-in installation.

Install the SnapCenter Plug-in for PostgreSQL

Add hosts and install plug-in packages on remote hosts

You must use the SnapCenter Add Host page to add hosts, and then install the plug-ins packages. The plug-ins are automatically installed on the remote hosts. You can add the host and install plug-in packages for an individual host.

Before you begin

- If the operating system of the SnapCenter Server host is Windows 2019 and the operating system of the plug-in host is Windows 2022, you should perform the following:
  - Upgrade to Windows Server 2019 (OS Build 17763.5936) or later
  - Upgrade to Windows Server 2022 (OS Build 20348.2402) or later

- You must be a user that is assigned to a role that has the plug-in install and uninstall permissions, such as the SnapCenter Admin role.
- When installing a plug-in on a Windows host, if you specify a credential that is not built-in, or if the user belongs to a local workgroup user, you must disable UAC on the host.
- You should ensure that the message queueing service is running.
- The administration documentation contains information about managing hosts.
- If you are using group Managed Service Account (gMSA), you should configure gMSA with administrative privileges.

  Configure group Managed Service Account on Windows Server 2016 or later for PostgreSQL

About this task

- You cannot add a SnapCenter Server as a plug-in host to another SnapCenter Server.

Steps

1. In the left navigation pane, click **Hosts**.
2. Verify that the **Managed Hosts** tab is selected at the top.
3. Click **Add**.
4. In the Hosts page, perform the following actions:

For this field…	Do this…
Host Type	Select the type of host: • Windows • Linux The Plug-in for PostgreSQL is installed on the PostgreSQL client host, and this host can be on either a Windows system or a Linux system.
Host name	Enter the communication host name. Enter the fully qualified domain name (FQDN) or the IP address of the host. SnapCenter depends on the proper configuration of the DNS. Therefore, the best practice is to enter the FQDN.

For this field…	Do this…
Credentials	Either select the credential name that you created or create new credentials. The credential must have administrative rights on the remote host. For details, see the information about creating credentials. You can view details about the credentials by positioning your cursor over the credential name that you provided. The credentials authentication mode is determined by the host type that you specify in the Add Host wizard.

5. In the Select Plug-ins to Install section, select the plug-ins to install.

   While using the REST API to install Plug-in for PostgreSQL, you must pass the version as 3.0. For example, PostgreSQL:3.0

6. (Optional) Click **More Options**.

For this field…	Do this…
Port	Either retain the default port number or specify the port number. The default port number is 8145. If the SnapCenter Server was installed on a custom port, that port number will be displayed as the default port. If you manually installed the plug-ins and specified a custom port, you must specify the same port. Otherwise, the operation fails.
Installation Path	The Plug-in for PostgreSQL is installed on the PostgreSQL client host, and this host can be on either a Windows system or a Linux system. • For the SnapCenter Plug-ins Package for Windows, the default path is C:\Program Files\NetApp\SnapCenter. Optionally, you can customize the path. • For the SnapCenter Plug-ins Package for Linux, the default path is /opt/NetApp/snapcenter. Optionally, you can customize the path.

For this field…	Do this…
Skip preinstall checks	Select this check box if you already installed the plug-ins manually and you do not want to validate whether the host meets the requirements for installing the plug-in.
Add all hosts in the cluster	Select this check box to add all the cluster nodes.
Use group Managed Service Account (gMSA) to run the plug-in services	For Windows host, select this check box if you want to use group Managed Service Account (gMSA) to run the plug-in services. Provide the gMSA name in the following format: domainName\accountName$. gMSA will be used as a log on service account only for SnapCenter Plug-in for Windows service.

7. Click **Submit**.

   If you have not selected the Skip prechecks checkbox, the host is validated to verify whether the host meets the requirements for installing the plug-in. The disk space, RAM, PowerShell version, .NET version, location (for Windows plug-ins), and Java version (for Linux plug-ins) are validated against the minimum requirements. If the minimum requirements are not met, appropriate error or warning messages are displayed.

   If the error is related to disk space or RAM, you can update the web.config file located at C:\Program Files\NetApp\SnapCenter WebApp to modify the default values. If the error is related to other parameters, you must fix the issue.

   In an HA setup, if you are updating web.config file, you must update the file on both nodes.

8. If host type is Linux, verify the fingerprint, and then click **Confirm and Submit**.

   In a cluster setup, you should verify the fingerprint of each of the nodes in the cluster.

   Fingerprint verification is mandatory even if the same host was added earlier to SnapCenter and the fingerprint was confirmed.

9. Monitor the installation progress.
   - For Windows plug-in, the install and upgrade logs are located at: *C:\Windows\SnapCenter plugin\Install*<JOBID>_
   - For Linux plug-in, the install logs are located at: */var/opt/snapcenter/logs/SnapCenter_Linux_Host_Plug-in_Install*<JOBID>.log_ and the upgrade logs are located at: */var/opt/snapcenter/logs/SnapCenter_Linux_Host_Plug-in_Upgrade*<JOBID>.log_

Install SnapCenter Plug-in Packages for Linux or Windows on multiple remote hosts by using cmdlets

You can install the SnapCenter Plug-in Packages for Linux or Windows on multiple hosts simultaneously by using the Install-SmHostPackage PowerShell cmdlet.

Before you begin

You must have logged in to SnapCenter as a domain user with local administrator rights on each host on which you want to install the plug-in package.

Steps

1. Launch PowerShell.
2. On the SnapCenter Server host, establish a session using the Open-SmConnection cmdlet, and then enter your credentials.
3. Install the plug-in on multiple hosts using the Install-SmHostPackage cmdlet and the required parameters.

   The information regarding the parameters that can be used with the cmdlet and their descriptions can be obtained by running *Get-Help command_name*. Alternatively, you can also refer to the SnapCenter Software Cmdlet Reference Guide.

   You can use the -skipprecheck option when you have installed the plug-ins manually and do not want to validate whether the host meets the requirements to install the plug-in.

4. Enter your credentials for remote installation.

Install the SnapCenter Plug-in for PostgreSQL on Linux hosts by using the command-line interface

You should install the SnapCenter Plug-in for PostgreSQL cluster by using the SnapCenter user interface (UI). If your environment does not allow remote installation of the plug-in from the SnapCenter UI, you can install the Plug-in for PostgreSQL cluster either in console mode or in silent mode by using the command-line interface (CLI).

Before you begin

- You should install the Plug-in for PostgreSQL cluster on each of the Linux host where the PostgreSQL client resides.
- The Linux host on which you are installing the SnapCenter Plug-in for PostgreSQL cluster must meet the dependent software, cluster, and operating system requirements.

  The Interoperability Matrix Tool (IMT) contains the latest information about the supported configurations.

  NetApp Interoperability Matrix Tool

- The SnapCenter Plug-in for PostgreSQL cluster is part of SnapCenter Plug-ins Package for Linux. Before you install SnapCenter Plug-ins Package for Linux, you should have already installed SnapCenter on a Windows host.

Steps

1. Copy the SnapCenter Plug-ins Package for Linux installation file (snapcenter_linux_host_plugin.bin) from C:\ProgramData\NetApp\SnapCenter\Package Repository to the host where you want to install the Plug-in for PostgreSQL.

   You can access this path from the host where the SnapCenter Server is installed.

2. From the command prompt, navigate to the directory where you copied the installation file.
3. Install the plug-in: `path_to_installation_bin_file/snapcenter_linux_host_plugin.bin -i silent -DPORT=port_number_for_host -DSERVER_IP=server_name_or_ip_address -DSERVER_HTTPS_PORT=port_number_for_server`
   - -DPORT specifies the SMCore HTTPS communication port.
   - -DSERVER_IP specifies the SnapCenter Server IP address.
   - -DSERVER_HTTPS_PORT specifies the SnapCenter Server HTTPS port.
   - -DUSER_INSTALL_DIR specifies the directory where you want to install the SnapCenter Plug-ins Package for Linux.
   - DINSTALL_LOG_NAME specifies the name of the log file.

```
/tmp/sc-plugin-installer/snapcenter_linux_host_plugin.bin -i silent
-DPORT=8145 -DSERVER_IP=scserver.domain.com -DSERVER_HTTPS_PORT=8146
-DUSER_INSTALL_DIR=/opt
-DINSTALL_LOG_NAME=SnapCenter_Linux_Host_Plugin_Install_2.log
-DCHOSEN_FEATURE_LIST=CUSTOM
```

4. Edit the /<installation directory>/NetApp/snapcenter/scc/etc/SC_SMS_Services.properties file, and then add the PLUGINS_ENABLED = PostgreSQL:3.0 parameter.
5. Add the host to the SnapCenter Server using the Add-Smhost cmdlet and the required parameters.

   The information regarding the parameters that can be used with the command and their descriptions can be obtained by running *Get-Help command_name*. Alternatively, you can also refer to the SnapCenter Software Cmdlet Reference Guide.

Monitor the status of installing Plug-in for PostgreSQL

You can monitor the progress of SnapCenter plug-in package installation by using the Jobs page. You might want to check the progress of installation to determine when it is complete or if there is an issue.

About this task

The following icons appear on the Jobs page and indicate the state of the operation:

- In progress
- Completed successfully
- Failed
- Completed with warnings or could not start due to warnings
- Queued

Steps

1. In the left navigation pane, click **Monitor**.
2. In the **Monitor** page, click **Jobs**.

3. In the **Jobs** page, to filter the list so that only plug-in installation operations are listed, do the following:
   a. Click **Filter**.
   b. Optional: Specify the start and end date.
   c. From the Type drop-down menu, select **Plug-in installation**.
   d. From the Status drop-down menu, select the installation status.
   e. Click **Apply**.
4. Select the installation job and click **Details** to view the job details.
5. In the **Job Details** page, click **View logs**.

Configure CA Certificate

Generate CA Certificate CSR file

You can generate a Certificate Signing Request (CSR) and import the certificate that can be obtained from a Certificate Authority (CA) using the generated CSR. The certificate will have a private key associated with it.

CSR is a block of encoded text that is given to an authorized certificate vendor to procure the signed CA certificate.

CA Certificate RSA key length must be minimum 3072 bits.

For information to generate a CSR, see How to generate CA Certificate CSR file.

If you own the CA certificate for your domain (*.domain.company.com) or your system (machine1.domain.company.com), you can skip generating the CA Certificate CSR file. You can deploy the existing CA certificate with SnapCenter.

For cluster configurations, the cluster name (virtual cluster FQDN), and the respective host names should be mentioned in the CA certificate. The certificate can be updated by filling the Subject Alternative Name (SAN) field before procuring the certificate. For a wild card certificate (*.domain.company.com), the certificate will contain all the hostnames of the domain implicitly.

Import CA certificates

You must import the CA certificates to the SnapCenter Server and the Windows host plug-ins using the Microsoft management console (MMC).

Steps

1. Go to the Microsoft management console (MMC), and then click **File** > **Add/Remove Snapin**.
2. In the Add or Remove Snap-ins window, select **Certificates** and then click **Add**.
3. In the Certificates snap-in window, select the **Computer account** option, and then click **Finish**.
4. Click **Console Root** > **Certificates – Local Computer** > **Trusted Root Certification Authorities** > **Certificates**.
5. Right-click on the folder “Trusted Root Certification Authorities”, and then select **All Tasks** > **Import** to start the import wizard.

6. Complete the wizard, as follows:

In this wizard window…	Do the following…
Import Private Key	Select the option **Yes**, import the private key, and then click **Next**.
Import File Format	Make no changes; click **Next**.
Security	Specify the new password to be used for the exported certificate, and then click **Next**.
Completing the Certificate Import Wizard	Review the summary, and then click **Finish** to start the import.

Importing certificate should be bundled with the private key (supported formats are: *.pfx, *.p12, and *.p7b).

7. Repeat Step 5 for the “Personal” folder.

Get the CA certificate thumbprint

A certificate thumbprint is a hexadecimal string that identifies a certificate. A thumbprint is calculated from the content of the certificate using a thumbprint algorithm.

Steps

1. Perform the following on the GUI:
   a. Double-click the certificate.
   b. In the Certificate dialog box, click the **Details** tab.
   c. Scroll through the list of fields and click **Thumbprint**.
   d. Copy the hexadecimal characters from the box.
   e. Remove the spaces between the hexadecimal numbers.

      For example, if the thumbprint is: "a9 09 50 2d d8 2a e4 14 33 e6 f8 38 86 b0 0d 42 77 a3 2a 7b", after removing the spaces, it will be: "a909502dd82ae41433e6f83886b00d4277a32a7b".

2. Perform the following from PowerShell:
   a. Run the following command to list the thumbprint of the installed certificate and identify the recently installed certificate by the subject name.

      Get-ChildItem -Path Cert:\LocalMachine\My

   b. Copy the thumbprint.

Configure CA certificate with Windows host plug-in services

You should configure the CA certificate with Windows host plug-in services to activate the installed digital certificate.

Perform the following steps on the SnapCenter Server and all the plug-in hosts where CA certificates are already deployed.

Steps

1. Remove the existing certificate binding with SMCore default port 8145, by running the following command:

   ```
   > netsh http delete sslcert ipport=0.0.0.0:_<SMCore Port>
   ```

   For example:

   ```
   > netsh http delete sslcert ipport=0.0.0.0:8145
   ```

2. Bind the newly installed certificate with the Windows host plug-in services, by running the following commands:

   ```
   > $cert = "_<certificate thumbprint>_"
   > $guid = [guid]::NewGuid().ToString("B")
   > netsh http add sslcert ipport=0.0.0.0: _<SMCore Port>_ certhash=$cert
   appid="$guid"
   ```

   For example:

   ```
   > $cert = "a909502dd82ae41433e6f83886b00d4277a32a7b"
   > $guid = [guid]::NewGuid().ToString("B")
   > netsh http add sslcert ipport=0.0.0.0: _<SMCore Port>_ certhash=$cert
   appid="$guid"
   ```

Configure the CA Certificate for the SnapCenter PostgreSQL Plug-ins service on Linux host

You should manage the password of the plug-ins keystore and its certificate, configure the CA certificate, configure root or intermediate certificates to the plug-ins trust-store, and configure CA signed key pair to plug-ins trust-store with SnapCenter plug-ins service to activate the installed digital certificate.

The plug-ins uses the file 'keystore.jks', which is located at */opt/NetApp/snapcenter/scc/etc* both as its trust-store and key-store.

Manage password for plug-in keystore and alias of the CA signed key pair in use

Steps

1. You can retrieve plug-in keystore default password from plug-in agent property file.

   It is the value corresponding to the key 'KEYSTORE_PASS'.

2. Change the keystore password:

```
keytool -storepasswd -keystore keystore.jks
```

3. Change the password for all aliases of private key entries in the keystore to the same password used for the keystore:

```
keytool -keypasswd -alias "alias_name_in_cert" -keystore keystore.jks
```

Update the same for the key KEYSTORE_PASS in *agent.properties* file.

4. Restart the service after changing the password.

Password for plug-in keystore and for all the associated alias password of the private key should be same.

Configure root or intermediate certificates to plug-in trust-store

You should configure the root or intermediate certificates without the private key to plug-in trust-store.

Steps

1. Navigate to the folder containing the plug-in keystore: /opt/NetApp/snapcenter/scc/etc.
2. Locate the file 'keystore.jks'.
3. List the added certificates in the keystore:

```
keytool -list -v -keystore keystore.jks
```

4. Add a root or intermediate certificate:

```
keytool -import -trustcacerts -alias myRootCA -file
/root/USERTrustRSA_Root.cer -keystore keystore.jks
```

5. Restart the service after configuring the root or intermediate certificates to plug-in trust-store.

You should add the root CA certificate and then the intermediate CA certificates.

Configure CA signed key pair to plug-in trust-store

You should configure the CA signed key pair to the plug-in trust-store.

Steps

1. Navigate to the folder containing the plug-in keystore /opt/NetApp/snapcenter/scc/etc.
2. Locate the file 'keystore.jks'.
3. List the added certificates in the keystore:

```
keytool -list -v -keystore keystore.jks
```

4. Add the CA certificate having both private and public key.

```
keytool -importkeystore -srckeystore /root/snapcenter.ssl.test.netapp.com.pfx
-srcstoretype pkcs12 -destkeystore keystore.jks -deststoretype JKS
```

5. List the added certificates in the keystore.

   ```
   keytool -list -v -keystore keystore.jks
   ```

6. Verify that the keystore contains the alias corresponding to the new CA certificate, which was added to the keystore.
7. Change the added private key password for CA certificate to the keystore password.

   Default plug-in keystore password is the value of the key KEYSTORE_PASS in agent.properties file.

   ```
   keytool -keypasswd -alias "alias_name_in_CA_cert" -keystore keystore.jks
   ```

8. If the alias name in the CA certificate is long and contains space or special characters ("*",","), change the alias name to a simple name:

   ```
   keytool -changealias -alias "long_alias_name" -destalias "simple_alias"
   -keystore keystore.jks
   ```

9. Configure the alias name from CA certificate in agent.properties file.

   Update this value against the key SCC_CERTIFICATE_ALIAS.

10. Restart the service after configuring the CA signed key pair to plug-in trust-store.

Configure certificate revocation list (CRL) for plug-ins

About this task

- SnapCenter Plug-ins will search for the CRL files in a pre-configured directory.
- Default directory for the CRL files for SnapCenter Plug-ins is ' opt/NetApp/snapcenter/scc/etc/crl'.

Steps

1. You can modify and update the default directory in agent.properties file against the key CRL_PATH.

   You can place more than one CRL file in this directory. The incoming certificates will be verified against each CRL.

Configure the CA Certificate for the SnapCenter PostgreSQL Plug-ins service on Windows host

You should manage the password of the plug-ins keystore and its certificate, configure the CA certificate, configure root or intermediate certificates to the plug-ins trust-store, and configure CA signed key pair to plug-ins trust-store with SnapCenter plug-ins service to activate the installed digital certificate.

The plug-ins uses the file *keystore.jks*, which is located at *C:\Program Files\NetApp\SnapCenter\Snapcenter Plug-in Creator\etc* both as its trust-store and key-store.

Manage password for plug-in keystore and alias of the CA signed key pair in use

Steps

1. You can retrieve plug-in keystore default password from plug-in agent property file.

   It is the value corresponding to the key *KEYSTORE_PASS*.

2. Change the keystore password:

   keytool -storepasswd -keystore keystore.jks

   If the "keytool" command is not recognized on the Windows command prompt, replace the keytool command with its complete path.

   C:\Program Files\Java\<jdk_version>\bin\keytool.exe" -storepasswd -keystore keystore.jks

3. Change the password for all aliases of private key entries in the keystore to the same password used for the keystore:

   keytool -keypasswd -alias "alias_name_in_cert" -keystore keystore.jks

   Update the same for the key KEYSTORE_PASS in *agent.properties* file.

4. Restart the service after changing the password.

   Password for plug-in keystore and for all the associated alias password of the private key should be same.

Configure root or intermediate certificates to plug-in trust-store

You should configure the root or intermediate certificates without the private key to plug-in trust-store.

Steps

1. Navigate to the folder containing the plug-in keystore *C:\Program Files\NetApp\SnapCenter\Snapcenter Plug-in Creator\etc*
2. Locate the file 'keystore.jks'.
3. List the added certificates in the keystore:

   keytool -list -v -keystore keystore.jks

4. Add a root or intermediate certificate:

   keytool -import -trustcacerts -alias myRootCA -file /root/USERTrustRSA_Root.cer -keystore keystore.jks

5. Restart the service after configuring the root or intermediate certificates to plug-in trust-store.

   You should add the root CA certificate and then the intermediate CA certificates.

Configure CA signed key pair to plug-in trust-store

You should configure the CA signed key pair to the plug-in trust-store.

Steps

1. Navigate to the folder containing the plug-in keystore *C:\Program Files\NetApp\SnapCenter\Snapcenter Plug-in Creator\etc*
2. Locate the file *keystore.jks*.
3. List the added certificates in the keystore:

   keytool -list -v -keystore keystore.jks

4. Add the CA certificate having both private and public key.

   keytool -importkeystore -srckeystore /root/snapcenter.ssl.test.netapp.com.pfx -srcstoretype pkcs12 -destkeystore keystore.jks -deststoretype JKS

5. List the added certificates in the keystore.

   keytool -list -v -keystore keystore.jks

6. Verify that the keystore contains the alias corresponding to the new CA certificate, which was added to the keystore.
7. Change the added private key password for CA certificate to the keystore password.

   Default plug-in keystore password is the value of the key KEYSTORE_PASS in agent.properties file.

   keytool -keypasswd -alias "alias_name_in_CA_cert" -keystore keystore.jks

8. Configure the alias name from CA certificate in *agent.properties* file.

   Update this value against the key SCC_CERTIFICATE_ALIAS.

9. Restart the service after configuring the CA signed key pair to plug-in trust-store.

Configure certificate revocation list (CRL) for SnapCenter plug-ins

About this task

- To download the latest CRL file for the related CA certificate see How to update certificate revocation list file in SnapCenter CA Certificate.
- SnapCenter plug-ins will search for the CRL files in a pre-configured directory.
- Default directory for the CRL files for SnapCenter plug-ins is *'C:\Program Files\NetApp\SnapCenter\Snapcenter Plug-in Creator\ etc\crl'*.

Steps

1. You can modify and update the default directory in *agent.properties* file against the key CRL_PATH.
2. You can place more than one CRL file in this directory.

   The incoming certificates will be verified against each CRL.

Enable CA Certificates for plug-ins

You should configure the CA certificates and deploy the CA certificates in the SnapCenter Server and the corresponding plug-in hosts. You should enable the CA certificate validation for the plug-ins.

Before you begin

- You can enable or disable the CA certificates using the run *Set-SmCertificateSettings* cmdlet.
- You can display the certificate status for the plug-ins using the *Get-SmCertificateSettings*.

The information regarding the parameters that can be used with the cmdlet and their descriptions can be obtained by running *Get-Help command_name*. Alternatively, you can also refer to the SnapCenter Software Cmdlet Reference Guide.

Steps

1. In the left navigation pane, click **Hosts**.
2. In the Hosts page, click **Managed Hosts**.
3. Select single or multiple plug-in hosts.
4. Click **More options**.
5. Select **Enable Certificate Validation**.

After you finish

The Managed Hosts tab host displays a padlock and the color of the padlock indicates the status of the connection between SnapCenter Server and the plug-in host.

- indicates that the CA certificate is neither enabled nor assigned to the plug-in host.
- indicates that the CA certificate is successfully validated.
- indicates that the CA certificate could not be validated.
- indicates that the connection information could not be retrieved.

When the status is yellow or green, the data protection operations completes successfully.

Prepare for data protection

Prerequisites for using the SnapCenter Plug-in for PostgreSQL

Before you use SnapCenter Plug-in for PostgreSQL, the SnapCenter administrator must install and configure the SnapCenter Server and perform the prerequisite tasks.

- Install and configure SnapCenter Server.
- Log in to SnapCenter Server.
- Configure the SnapCenter environment by adding storage system connections and creating credentials, if applicable.
- Install Java 11 on your Linux or Windows host.

  You must set the Java path in the environmental path variable of the host machine.

- Set up SnapMirror and SnapVault, if you want backup replication.

How resources, resource groups, and policies are used for protecting PostgreSQL

Before you use SnapCenter, it is helpful to understand basic concepts related to the

backup, clone, and restore operations you want to perform. You interact with resources, resource groups, and policies for different operations.

- Resources are typically PostgreSQL clusters that you back up or clone with SnapCenter.
- A SnapCenter resource group, is a collection of resources on a host.

  When you perform an operation on a resource group, you perform that operation on the resources defined in the resource group according to the schedule you specify for the resource group.

  You can back up on demand a single resource or a resource group. You also can perform scheduled backups for single resources and resource groups.

- The policies specify the backup frequency, replication, scripts, and other characteristics of data protection operations.

  When you create a resource group, you select one or more policies for that group. You can also select a policy when you perform a backup on demand for a single resource.

Think of a resource group as defining what you want to protect and when you want to protect it in terms of day and time. Think of a policy as defining how you want to protect it. If you are backing up all clusters, for example, you might create a resource group that includes all of the clusters in the host. You could then attach two policies to the resource group: a daily policy and an hourly policy. When you create the resource group and attach the policies, you might configure the resource group to perform a full backup daily.

Back up PostgreSQL resources

Back up PostgreSQL resources

You can either create a backup of a resource (cluster) or resource group. The backup workflow includes planning, identifying the clusters for backup, managing backup policies, creating resource groups and attaching policies, creating backups, and monitoring the operations.

The following workflow shows the sequence in which you must perform the backup operation:

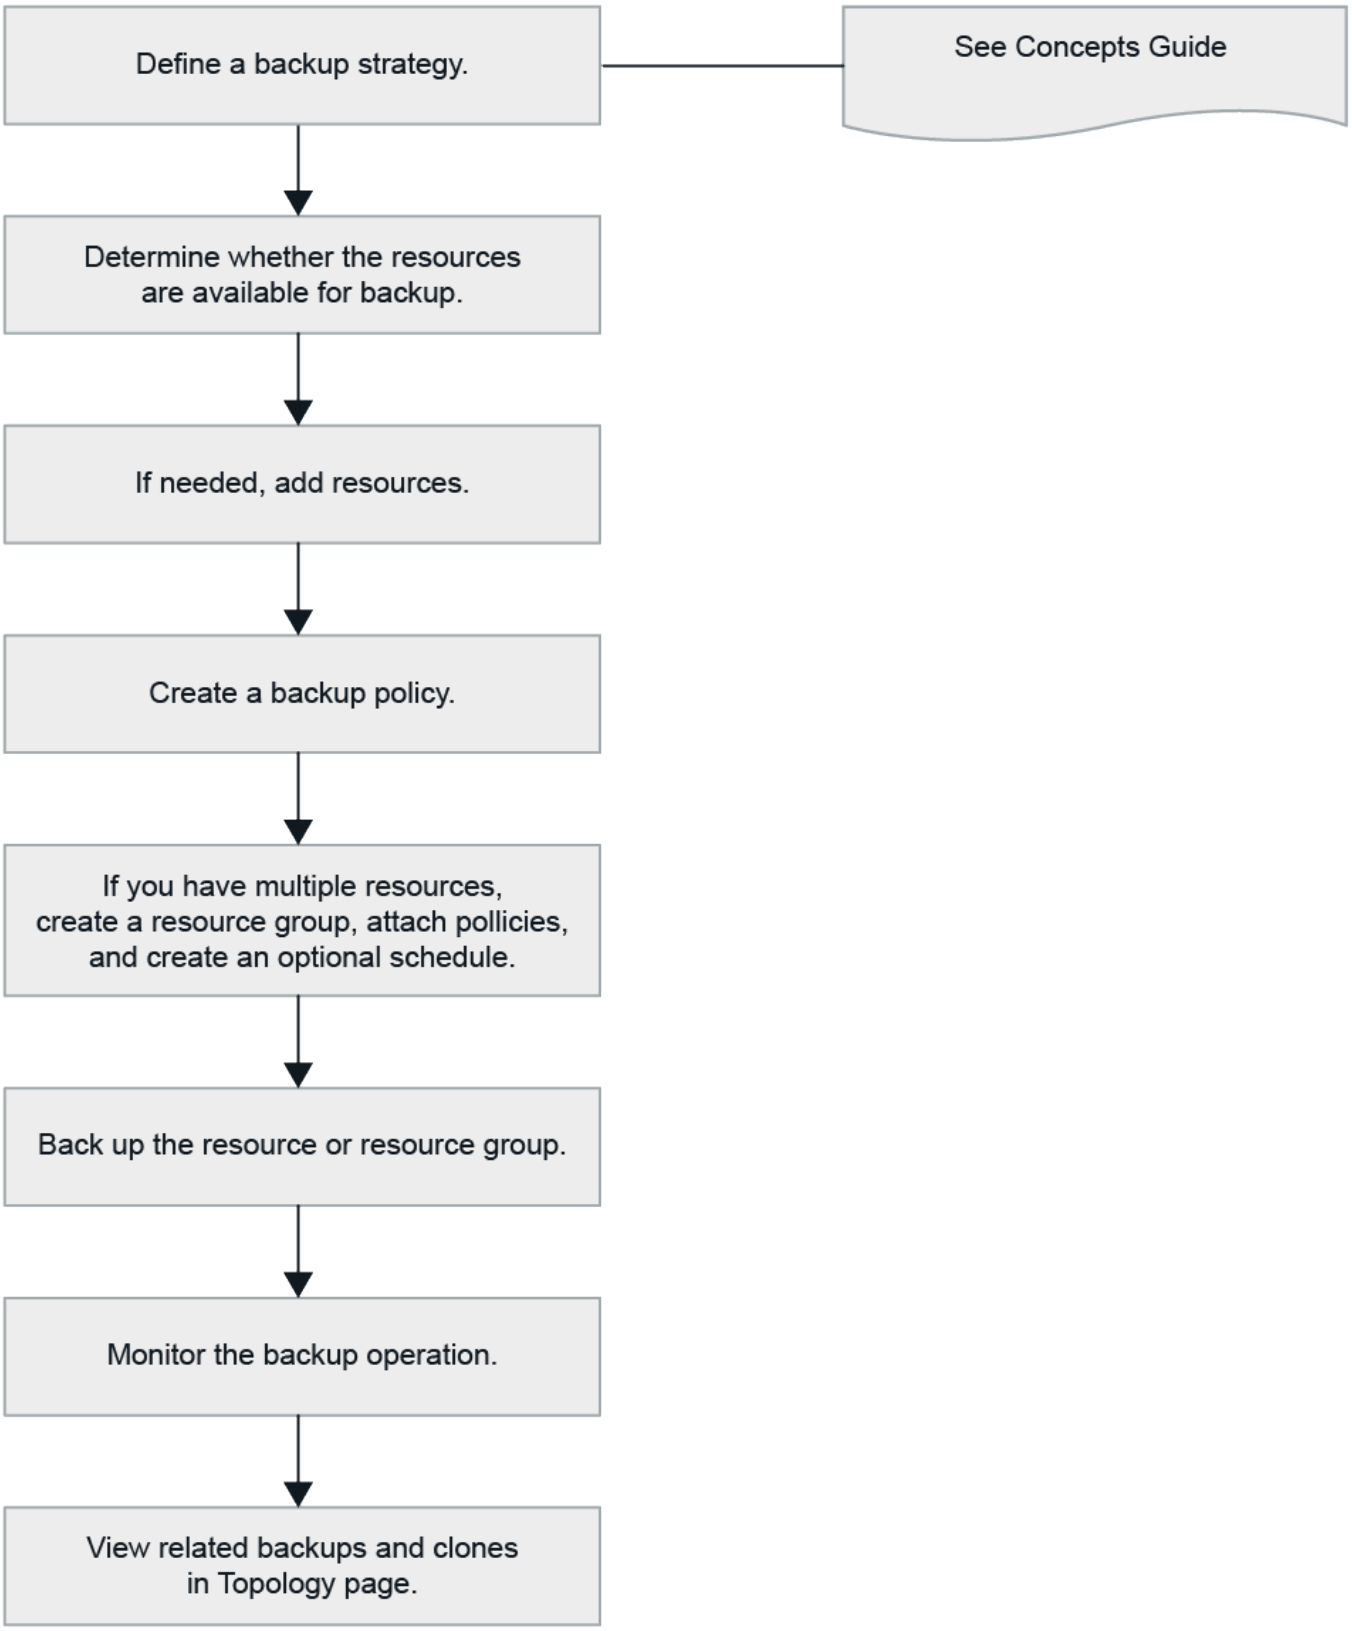

You can also use PowerShell cmdlets manually or in scripts to perform backup, restore, and clone operations. The SnapCenter cmdlet help and the cmdlet reference information contain more information about PowerShell cmdlets.
SnapCenter Software Cmdlet Reference Guide.

Discover the clusters automatically

Resources are PostgreSQL clusters on the Linux host that are managed by SnapCenter. You can add the resources to resource groups to perform data protection operations after you discover the PostgreSQL clusters that are available.

Before you begin

- You must have already completed tasks such as installing the SnapCenter Server, adding hosts, and setting up the storage system connections.
- SnapCenter Plug-in for PostgreSQL does not support automatic discovery of the resources residing on RDM/VMDK virtual environments.

About this task

- After installing the plug-in, all the clusters on that Linux host are automatically discovered and displayed on the Resources page.
- Only clusters are auto-discovered.

The automatically discovered resources cannot be modified or deleted.

Steps

1. In the left navigation pane, click **Resources**, and then select the Plug-in for PostgreSQL from the list.
2. In the Resources page select the resource type from the View list.
3. (Optional) Click , and then select the host name.

   You can then click to close the filter pane.

4. Click **Refresh Resources** to discover the resources available on the host.

   The resources are displayed along with information such as resource type, host name, associated resource groups, backup type, policies and overall status.

   - If the cluster is on a NetApp storage and not protected, then Not protected is displayed in the Overall Status column.
   - If the cluster is on a NetApp storage system and protected, and if there is no backup operation performed, then Backup not run is displayed in the Overall Status column. The status will otherwise change to Backup failed or Backup succeeded based on the last backup status.

You must refresh the resources if the clusters are renamed outside of SnapCenter.

Add resources manually to the plug-in host

Automatic discovery is not supported on Windows host. You must add Postgresql cluster resources manually.

Before you begin

- You must have completed tasks such as installing the SnapCenter Server, adding hosts, and setting up storage system connections.

About this task

Automatic discovery is not supported for the following configurations:

- RDM and VMDK layouts

Steps

1. In the left navigation pane, select the SnapCenter Plug-in for Postgresql from the drop-down list, and then click **Resources**.
2. In the Resources page, click **Add Postgresql resources**.
3. In the Provide Resource Details page, perform the following actions:

For this field…	Do this…
Name	Specify the cluster name.
Host Name	Enter the host name.
Type	Select cluster.
Instance	Specify the name of the instance, which is the parent of the cluster.
Credentials	Select the credentials or add information for the credential. This is optional.

4. In the Provide Storage Footprint page, select a storage type and choose one or more volumes, LUNs, and qtrees, and then click **Save**.

   Optional: You can click the + icon to add more volumes, LUNs, and qtrees from other storage systems.

5. Optional: In the Resource Settings page, for resources on the Windows host, enter custom key-value pairs for PostgreSQL plug-in
6. Review the summary, and then click **Finish**.

   The clusters are displayed along with information such as the host name, associated resource groups and policies, and overall status

If you want to provide users access to resources, you must assign the resources to the users. This enables users to perform the actions for which they have permissions on the assets that are assigned to them.

Add a user or group and assign role and assets

After you finish

- After adding the clusters, you can modify the PostgreSQL cluster details.
- The migrated resources (tablespace and clusters) from SnapCenter 5.0 will be tagged as PostgreSQL cluster type in SnapCenter 6.0.
- When you modify the manually added resources that are migrated from SnapCenter 5.0 or earlier, do the following in the **Resource Settings** page for custom key value pairs:
  - Specify the term "PORT" in the **Name** field.

◦ Specify the port number in the **Value** field.

Create backup policies for PostgreSQL

Before you use SnapCenter to back up PostgreSQL resources, you must create a backup policy for the resource or resource group that you want to back up. A backup policy is a set of rules that governs how you manage, schedule, and retain backups.

Before you begin

- You must have defined your backup strategy.

  For details, see the information about defining a data protection strategy for PostgreSQL clusters.

- You must have prepared for data protection by completing tasks such as installing SnapCenter, adding hosts, setting up storage system connections, and adding resources.
- The SnapCenter administrator must have assigned the SVMs for both the source and destination volumes to you if you are replicating snapshots to a mirror or vault.

  Additionally, you can specify replication, script, and application settings in the policy. These options saves time when you want to reuse the policy for another resource group.

About this task

- SnapLock
  - If 'Retain the backup copies for a specific number of days' option is selected, then the SnapLock retention period must be lesser than or equal to the mentioned retention days.
  - Specifying a snapshot locking period prevents deletion of the snapshots until the retention period expires. This could lead to retaining a larger number of snapshots than the count specified in the policy.
  - For ONTAP 9.12.1 and below version, the clones created from the SnapLock Vault snapshots as part of restore will inherit the SnapLock Vault expiry time. Storage admin should manually cleanup the clones post the SnapLock expiry time.

Steps

1. In the left navigation pane, click **Settings**.
2. In the Settings page, click **Policies**.
3. Click **New**.
4. In the Name page, enter the policy name and details.
5. In the Policy type page, perform the following:
   a. Select storage type.
   b. In the **Custom backup settings** section, provide any specific backup settings that have to be passed to the plug-in in key-value format.

      You can provide multiple key-values to be passed to the plug-in.

6. In the Backup and Replication page, perform the following actions:
   a. Specify the schedule frequency by selecting **On demand**, **Hourly**, **Daily**, **Weekly**, or **Monthly**.

You can specify the schedule (start date, end date, and frequency) for the backup operation while creating a resource group. This enables you to create resource groups that share the same policy and backup frequency, but also enables you to assign different backup schedules to each policy.

If you have scheduled for 2:00 a.m., the schedule will not be triggered during daylight saving time (DST).

b. In the Snapshot settings section, specify the retention settings for the backup type and the schedule type selected in the **Backup Type** page:

If you want to…	Then…
Keep a certain number of Snapshots	Select **Copies to keep**, and then specify the number of Snapshots that you want to keep. If the number of Snapshots exceeds the specified number, the Snapshots are deleted with the oldest copies deleted first. You must set the retention count to 2 or higher if you plan to enable SnapVault replication. If you set the retention count to 1, the retention operation might fail because the first Snapshot is the reference Snapshot for the SnapVault relationship until a newer Snapshot is replicated to the target. The maximum retention value is 1018. Backups will fail if retention is set to a value higher than what ONTAP version supports.
Keep the Snapshots for a certain number of days	Select **Retain copies for**, and then specify the number of days for which you want to keep the Snapshots before deleting them.
Snapshot copy locking period	Select **Snapshot copy locking period**, and specify days, months, or years. SnapLock retention period should be less than 100 years.

7. Select a Policy label.

You can assign SnapMirror labels to primary snapshots for remote replication, allowing the primary snapshots to offload the snapshot replication operation from SnapCenter to ONTAP secondary systems. This can be done without enabling SnapMirror or SnapVault option in the policy page.

8. In the Select secondary replication options section, select one or both of the following secondary replication options:

For this field…	Do this…
Update SnapMirror after creating a local Snapshot copy	Select this field to create mirror copies of the backup sets on another volume (SnapMirror replication). If the protection relationship in ONTAP is of type Mirror and Vault and if you select only this option, Snapshot created on the primary will not be transferred to the destination, but will be listed in the destination. If this Snapshot is selected from the destination to perform a restore operation, then the following error message is displayed: Secondary Location is not available for the selected vaulted/mirrored backup. During secondary replication, the SnapLock expiry time loads the primary SnapLock expiry time. Clicking the **Refresh** button in the Topology page refreshes the secondary and primary SnapLock expiry time that are retrieved from ONTAP. See View PostgreSQL resource related backups and clones in the Topology page.

For this field…	Do this…
Update SnapVault after creating a local Snapshot copy	Select this option to perform disk-to-disk backup replication (SnapVault backups). During secondary replication, the SnapLock expiry time loads the primary SnapLock expiry time. Clicking the **Refresh** button in the Topology page refreshes the secondary and primary SnapLock expiry time that are retrieved from ONTAP. When SnapLock is configured only on the secondary from ONTAP known as SnapLock Vault, clicking the **Refresh** button in the Topology page refreshes the locking period on the secondary that is retrieved from ONTAP. For more information on SnapLock Vault see Commit Snapshots to WORM on a vault destination See View PostgreSQL resource related backups and clones in the Topology page.
Error retry count	Enter the maximum number of replication attempts that can be allowed before the operation stops.

You should configure SnapMirror retention policy in ONTAP for the secondary storage to avoid reaching the maximum limit of Snapshots on the secondary storage.

9. Review the summary, and then click **Finish**.

Create resource groups and attach policies

A resource group is the container to which you must add resources that you want to back up and protect. A resource group enables you to back up all the data that is associated with a given application simultaneously. A resource group is required for any data protection job. You must also attach one or more policies to the resource group to define the type of data protection job that you want to perform.

About this task

- For ONTAP 9.12.1 and below version, the clones created from the SnapLock Vault snapshots as part of restore will inherit the SnapLock Vault expiry time. Storage admin should manually cleanup the clones post the SnapLock expiry time.

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page, click **New Resource Group**.
3. In the Name page, perform the following actions:

For this field…	Do this…
Name	Enter a name for the resource group. The resource group name should not exceed 250 characters.
Tags	Enter one or more labels that will help you later search for the resource group. For example, if you add HR as a tag to multiple resource groups, you can later find all resource groups associated with the HR tag.
Use custom name format for snapshot copy	Select this check box, and enter a custom name format that you want to use for the snapshot name. For example, customtext_resource group_policy_hostname or resource group_hostname. By default, a timestamp is appended to the snapshot name.

4. In the Resources page, select a host name from the **Host** drop-down list and resource type from the **Resource Type** drop-down list.

   This helps to filter information on the screen.

5. Select the resources from the **Available Resources** section, and then click the right arrow to move them to the **Selected Resources** section.
6. In the Application Settings page, do the following:
   a. Click the **Backups** arrow to set additional backup options:

      Enable consistency group backup and perform the following tasks:

For this field…	Do this…
Afford time to wait for Consistency Group snapshot operation to complete	Select **Urgent**, **Medium**, or **Relaxed** to specify the wait time for snapshot operation to complete. Urgent = 5 seconds, Medium = 7 seconds, and Relaxed = 20 seconds.
Disable WAFL Sync	Select this to avoid forcing a WAFL consistency point.

b. Click the **Scripts** arrow and enter the pre and post commands for quiesce, snapshot, and unquiesce operations. You can also enter the pre commands to be executed before exiting in the event of a failure.

c. Click the **Custom Configurations** arrow and enter the custom key-value pairs required for all data protection operations using this resource.

Parameter	Setting	Description
ARCHIVE_LOG_ENABLE	(Y/N)	Enables the archive log management to delete the archive logs.
ARCHIVE_LOG_RETENTION	number_of_days	Specifies the number of days the archive logs are retained. This setting must be equal to or greater than NTAP_SNAPSHOT_ RETENTIONS.
ARCHIVE_LOG_DIR	change_info_directory/logs	Specifies the path to the directory that contains the archive logs.

Parameter	Setting	Description
ARCHIVE_LOG_EXT	file_extension	Specifies the archive log file extension length. For example, if the archive log is log_backup_0_0_0_0.161518551942 9 and if the file_extension value is 5, then the extension of the log will retain 5 digits, which is 16151.
ARCHIVE_LOG_RECURSIVE_SE ARCH	(Y/N)	Enables the management of archive logs within subdirectories. You should use this parameter if the archive logs are located under subdirectories.

The custom key-value pairs are supported for PostgreSQL Linux plug-in systems and not supported for PostgreSQL cluster registered as a centralized windows plug-in.

d. Click the **Snapshot Copy Tool** arrow to select the tool to create snapshots:

If you want…	Then…
SnapCenter to use the plug-in for Windows and put the file system into a consistent state before creating a snapshot. For Linux resources, this option is not applicable.	Select **SnapCenter with File System Consistency**.
SnapCenter to create a storage level snapshot	Select **SnapCenter without File System Consistency**.
To enter the command to be executed on the host to create snapshot copies.	Select **Other**, and then enter the command to be executed on the host to create a snapshot.

7. In the Policies page, perform the following steps:
   a. Select one or more policies from the drop-down list.

You can also create a policy by clicking + .

The policies are listed in the Configure schedules for selected policies section.

b.

In the Configure Schedules column, click + for the policy you want to configure.

c. In the Add schedules for policy *policy_name* dialog box, configure the schedule, and then click **OK**.

Where, policy_name is the name of the policy that you have selected.

The configured schedules are listed in the **Applied Schedules** column.

Third party backup schedules are not supported when they overlap with SnapCenter backup schedules.

8. In the Notification page, from the **Email preference** drop-down list, select the scenarios in which you want to send the emails.

You must also specify the sender and receiver email addresses, and the subject of the email. The SMTP server must be configured in **Settings** > **Global Settings**.

9. Review the summary, and then click **Finish**.

Create resource groups and enable secondary protection for PostgreSQL resources on ASA r2 systems

You should create the resource group to add the resources that are on ASA r2 systems. You can also provision the secondary protection while creating the resource group.

Before you begin

- You should ensure that you are not adding both ONTAP 9.x resources and ASA r2 resources to the same resource group.
- You should ensure that you do not have a database with both ONTAP 9.x resources and ASA r2 resources.

About this task

- The secondary protection is available only if the logged-in user is assigned to the role that has the **SecondaryProtection** capability enabled.
- If you enabled secondary protection, the resource group is put into maintenance mode while creating the primary and secondary consistency groups. After the primary and secondary consistency groups are created, the resource group is put out of maintenance mode.
- SnapCenter does not support secondary protection for a clone resource.

Steps

1. In the left navigation pane, select **Resources**, and the appropriate plug-in from the list.
2. In the Resources page, click **New Resource Group**.
3. In the Name page, perform the following actions:
   a. Enter a name for the resource group in the Name field.

      The resource group name should not exceed 250 characters.

   b. Enter one or more labels in the Tag field to help you search for the resource group later.

      For example, if you add HR as a tag to multiple resource groups, you can later find all resource groups associated with the HR tag.

c. Select this check box, and enter a custom name format that you want to use for the Snapshot name.

   For example, customtext_resource group_policy_hostname or resource group_hostname. By default, a timestamp is appended to the Snapshot name.

d. Specify the destinations of the archive log files that you do not want to back up.

   You should use the exact same destination as it was set in the application including prefix, if needed.

4. In the Resources page, select the database host name from the **Host** drop-down list.

   The resources are listed in the Available Resources section only if the resource is discovered successfully. If you have recently added resources, they will appear on the list of available resources only after you refresh your resource list.

5. Select the ASA r2 resources from the Available Resources section and move them to the Selected Resources section.
6. In the Application Settings page, select the backup option.
7. In the Policies page, perform the following steps:

   a. Select one or more policies from the drop-down list.

      You can also create a policy by clicking + .

      In the Configure schedules for selected policies section, the selected policies are listed.

   b. Click + in the Configure Schedules column for the policy for which you want to configure a schedule.

   c. In the Add schedules for policy *policy_name* window, configure the schedule, and then click **OK**.

      Where, *policy_name* is the name of the policy that you have selected.

      The configured schedules are listed in the Applied Schedules column.

   Third party backup schedules are not supported when they overlap with SnapCenter backup schedules.

8. If the secondary protection is enabled for the policy that you have selected, then Secondary Protection page is displayed and you need to perform the following steps:

   a. Select the type of the replication policy.

      Synchronous replication policy is not supported.

   b. Specify the consistency group suffix that you want to use.

   c. From the Destination Cluster and Destination SVM drop-downs select the peered cluster and SVM that you want to use.

The cluster and SVM peering is not supported by SnapCenter. You should use System Manager or ONTAP CLIs to perform cluster and SVM peering.

If the resources are already protected outside of SnapCenter, those resources will be displayed in the Secondary Protected Resources section.

1. On the Verification page, perform the following steps:
   a. Click **Load locators** to load the SnapMirror or SnapVault volumes to perform verification on secondary storage.
   b. Click + in the Configure Schedules column to configure the verification schedule for all the schedule types of the policy.
   c. In the Add Verification Schedules policy_name dialog box, perform the following actions:

If you want to…	Do this…
Run verification after backup	Select **Run verification after backup**.
Schedule a verification	Select **Run scheduled verification** and then select the schedule type from the drop-down list.

   d. Select **Verify on secondary location** to verify your backups on secondary storage system.
   e. Click **OK**.

   The configured verification schedules are listed in the Applied Schedules column.

2. In the Notification page, from the **Email preference** drop-down list, select the scenarios in which you want to send the emails.

   You must also specify the sender and receiver email addresses, and the subject of the email. If you want to attach the report of the operation performed on the resource group, select **Attach Job Report**.

   For email notification, you must have specified the SMTP server details using the either the GUI or the PowerShell command Set-SmSmtpServer.

3. Review the summary, and then click **Finish**.

Create a storage system connection and a credential using PowerShell cmdlets for PostgreSQL

You must create a storage virtual machine (SVM) connection and a credential before using PowerShell cmdlets to back up, restore, or clone PostgreSQL clusters.

Before you begin

- You should have prepared the PowerShell environment to execute the PowerShell cmdlets.
- You should have the required permissions in the Infrastructure Admin role to create storage connections.
- You should ensure that the plug-in installations are not in progress.

Host plug-in installations must not be in progress while adding a storage system connection because the host cache might not be updated and clusters status might be displayed in the SnapCenter GUI as “Not available for backup” or “Not on NetApp storage”.

- Storage system names should be unique.

  SnapCenter does not support multiple storage systems with the same name on different clusters. Each storage system that is supported by SnapCenter should have a unique name and a unique data LIF IP address.

Steps

1. Initiate a PowerShell Core connection session by using the Open-SmConnection cmdlet.

```
PS C:\> Open-SmConnection
```

2. Create a new connection to the storage system by using the Add-SmStorageConnection cmdlet.

```
PS C:\> Add-SmStorageConnection -Storage test_vs1 -Protocol Https
-Timeout 60
```

3. Create a new credential by using the Add-SmCredential cmdlet.

   This example shows how to create a new credential named FinanceAdmin with Windows credentials:

```
PS C:> Add-SmCredential -Name FinanceAdmin -AuthMode Windows
-Credential sddev\administrator
```

4. Add the PostgreSQL communication host to SnapCenter Server.

```
PS C:> Add-SmHost -HostName 10.232.204.61 -OSType Windows -RunAsName
FinanceAdmin -PluginCode PostgreSQL
```

5. Install the package and the SnapCenter Plug-in for PostgreSQL on the host.

   For Linux:

```
PS C:> Install-SmHostPackage -HostNames 10.232.204.61 -ApplicationCode
PostgreSQL
```

   For Windows:

```
Install-SmHostPackage -HostNames 10.232.204.61 -ApplicationCode
PostgreSQL -FilesystemCode scw -RunAsName FinanceAdmin
```

6. Set the path to the SQLLIB.

   For Windows, PostgreSQL plug-in will use the default path for SQLLIB folder: “C:\Program Files\IBM\SQLLIB\BIN”

   If you want to override the default path, use the following command.

```
PS C:> Set-SmConfigSettings -Plugin -HostName 10.232.204.61 -PluginCode
PostgreSQL -configSettings @{“PostgreSQL_SQLLIB_CMD” =
“<custom_path>\IBM\SQLLIB\BIN”}
```

The information regarding the parameters that can be used with the cmdlet and their descriptions can be obtained by running *Get-Help command_name*. Alternatively, you can also refer to the SnapCenter Software Cmdlet Reference Guide.

Back up PostgreSQL

If a resource is not yet part of any resource group, you can back up the resource from the Resources page.

Before you begin

- You must have created a backup policy.
- If you want to back up a resource that has a SnapMirror relationship with a secondary storage, the ONTAP role assigned to the storage user should include the “snapmirror all” privilege. However, if you are using the “vsadmin” role, then the “snapmirror all” privilege is not required.
- For Snapshot copy-based backup operation, ensure that all the tenant clusters are valid and active.
- For pre and post commands for quiesce, Snapshot, and unquiesce operations, you should check if the commands exist in the command list available on the plug-in host from the following paths:
  - Default location on the Windows host: *C:\Program Files\NetApp\SnapCenter\Snapcenter Plug-in Creator\etc\allowed_commands.config*
  - Default location on the Linux host: */opt/NetApp/snapcenter/scc/etc/allowed_commands.config*

If the commands do not exist in the command list, then the operation will fail.

SnapCenter UI

Steps

1. In the left navigation pane, select **Resources**, and then select the appropriate plug-in from the list.
2. In the Resource page, filter resources from the **View** drop-down list based on resource type.

   Select , and then select the host name and the resource type to filter the resources. You can then select to close the filter pane.

3. Select the resource that you want to back up.
4. In the Resource page, select **Use custom name format for Snapshot copy**, and then enter a custom name format that you want to use for the Snapshot name.

   For example, *customtext_policy_hostname* or *resource_hostname*. By default, a timestamp is appended to the Snapshot name.

5. In the Application Settings page, do the following:
   - Select the **Backups** arrow to set additional backup options:

     Enable consistency group backup, if needed, and perform the following tasks:

For this field…	Do this…
Afford time to wait for "Consistency Group Snapshot" operation to complete	Select **Urgent**, or **Medium**, or **Relaxed** to specify the wait time for Snapshot operation to finish. Urgent = 5 seconds, Medium = 7 seconds, and Relaxed = 20 seconds.
Disable WAFL Sync	Select this to avoid forcing a WAFL consistency point.

   - Select the **Scripts** arrow to run pre and post commands for quiesce, Snapshot, and unquiesce operations.

     You can also run pre commands before exiting the backup operation. Prescripts and postscripts are run in the SnapCenter Server.

   - Select the **Custom Configurations**arrow, and then enter the custom value pairs required for all jobs using this resource.
   - Select the **Snapshot Copy Tool** arrow to select the tool to create Snapshots:

If you want…	Then…
SnapCenter to create a storage-level Snapshot	Select **SnapCenter without File System Consistency**.
SnapCenter to use the plug-in for Windows to put the file system into a consistent state and then create a Snapshot	Select **SnapCenter with File System Consistency**.

If you want…	Then…
To enter the command to create a Snapshot	Select **Other**, and then enter the command to create a Snapshot.

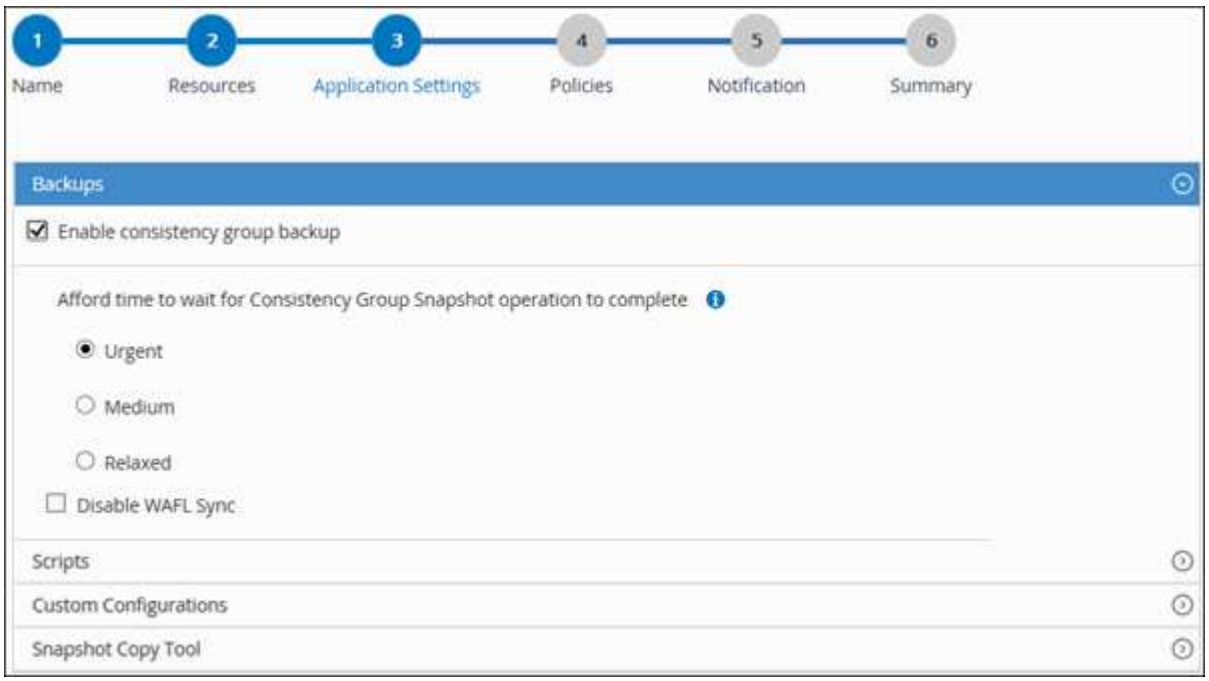

6. In the Policies page, perform the following steps:
   a. Select one or more policies from the drop-down list.

      You can also create a policy by clicking + .

      In the Configure schedules for selected policies section, the selected policies are listed.

   b. Select + in the Configure Schedules column for the policy for which you want to configure a schedule.
   c. In the Add schedules for policy *policy_name* dialog box, configure the schedule, and then select **OK**.

      policy_name is the name of the policy that you selected.

      The configured schedules are listed in the Applied Schedules column.
7. In the Notification page, from the **Email preference** drop-down list, select the scenarios in which you want to send the emails.

   You must also specify the sender and receiver email addresses, and the subject of the email. SMTP must also be configured in **Settings** > **Global Settings**.
8. Review the summary, and then select **Finish**.

   The resources topology page is displayed.

9. Select **Back up Now**.
10. In the Backup page, perform the following steps:
    a. If you applied multiple policies to the resource, from the **Policy** drop-down list, select the policy that you want to use for backup.

       If the policy selected for the on-demand backup is associated with a backup schedule, the on-demand backups will be retained based on the retention settings specified for the schedule type.

    b. Select **Backup**.
11. Monitor the operation progress by clicking **Monitor** > **Jobs**.
    - In MetroCluster configurations, SnapCenter might not be able to detect a protection relationship after a failover.

      For information, see: Unable to detect SnapMirror or SnapVault relationship after MetroCluster failover

    - If you are backing up application data on VMDKs and the Java heap size for the SnapCenter Plug-in for VMware vSphere is not large enough, the backup might fail.

      To increase the Java heap size, locate the script file */opt/netapp/init_scripts/scvservice*. In that script, the *do_start method* command starts the SnapCenter VMware plug-in service. Update that command to the following: *Java -jar -Xmx8192M -Xms4096M*

PowerShell cmdlets

Steps

1. Initiate a connection session with the SnapCenter Server for a specified user by using the Open-SmConnection cmdlet.

   ```
   PS C:\> Open-SmConnection
   ```

   The username and password prompt is displayed.

2. Add manual resources by using the Add-SmResources cmdlet.

   This example shows how to add a PostgreSQL instance:

   ```
   PS C:\> Add-SmResource -HostName 10.32.212.13 -PluginCode PostgreSQL
   -ResourceType Instance -ResourceName postgresqlinst1
   -StorageFootPrint
   (@{"VolumeName"="winpostgresql01_data01";"LUNName"="winpostgresql01_
   data01";"StorageSystem"="scsnfssvm"}) -MountPoints "D:\"
   ```

3. Create a backup policy by using the Add-SmPolicy cmdlet.
4. Protect the resource or add a new resource group to SnapCenter by using the Add-SmResourceGroup cmdlet.
5. Initiate a new backup job by using the New-SmBackup cmdlet.

This example shows how to backup a resource group:

```
C:\PS> New-SMBackup -ResourceGroupName 'ResourceGroup_wback-up-
clusters-using-powershell-cmdlets-postgresql.adocith_Resources'
-Policy postgresql_policy1
```

This example backs up a protected resource:

```
C:\PS> New-SMBackup -Resources
@{"Host"="10.232.204.42";"Uid"="MDC\SID";"PluginName"="postgresql"}
-Policy postgresql_policy2
```

6. Monitor the job status (running, completed, or failed) by using the Get-smJobSummaryReport cmdlet.

```
PS C:\> Get-smJobSummaryReport -JobID 123
```

7. Monitor the backup job details like backup ID, backup name to perform restore or clone operation by using the Get-SmBackupReport cmdlet.

```
PS C:\> Get-SmBackupReport -JobId 351
Output:
BackedUpObjects           : {DB1}
FailedObjects             : {}
IsScheduled               : False
HasMetadata               : False
SmBackupId                : 269
SmJobId                   : 2361
StartDateTime             : 10/4/2016 11:20:45 PM
EndDateTime               : 10/4/2016 11:21:32 PM
Duration                  : 00:00:46.2536470
CreatedDateTime           : 10/4/2016 11:21:09 PM
Status                    : Completed
ProtectionGroupName       : Verify_ASUP_Message_windows
SmProtectionGroupId       : 211
PolicyName                : test2
SmPolicyId                : 20
BackupName                : Verify_ASUP_Message_windows_scc54_10-04-
2016_23.20.46.2758
VerificationStatus        : NotVerified
VerificationStatuses      :
SmJobError                :
BackupType                : SCC_BACKUP
CatalogingStatus          : NotApplicable
CatalogingStatuses        :
ReportDataCreatedDateTime :
```

The information regarding the parameters that can be used with the cmdlet and their descriptions can be obtained by running *Get-Help command_name*. Alternatively, you can also refer to the SnapCenter Software Cmdlet Reference Guide.

Back up resource groups

A resource group is a collection of resources on a host. A backup operation on the resource group is performed on all resources defined in the resource group.

Before you begin

- You must have created a resource group with a policy attached.
- If you want to back up a resource that has a SnapMirror relationship with a secondary storage, the ONTAP role assigned to the storage user should include the “snapmirror all” privilege. However, if you are using the “vsadmin” role, then the “snapmirror all” privilege is not required.

About this task

You can back up a resource group on demand from the Resources page. If a resource group has a policy attached and a schedule configured, then backups occur automatically according to the schedule.

Steps

1. In the left navigation pane, select **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page, select **Resource Group** from the **View** list.

   You can search the resource group either by entering the resource group name in the search box or by selecting , and then selecting the tag. You can then select to close the filter pane.

3. In the Resource Groups page, select the resource group that you want to back up, and then select **Back up Now**.
4. In the Backup page, perform the following steps:
   a. If you associated multiple policies with the resource group, from the **Policy** drop-down list, select the policy that you want to use for backup.

      If the policy selected for the on-demand backup is associated with a backup schedule, the on-demand backups will be retained based on the retention settings specified for the schedule type.

   b. Select **Backup**.
5. Monitor the operation progress by selecting **Monitor** > **Jobs**.

Monitor PostgreSQL backup operations

You can monitor the progress of different backup operations by using the SnapCenterJobs page. You might want to check the progress to determine when it is complete or if there is an issue.

About this task

The following icons appear on the Jobs page and indicate the corresponding state of the operations:

- In progress
- Completed successfully
- Failed
- Completed with warnings or could not start due to warnings
- Queued
- Canceled

Steps

1. In the left navigation pane, click **Monitor**.
2. In the Monitor page, click **Jobs**.
3. In the Jobs page, perform the following steps:
   a. Click to filter the list so that only backup operations are listed.
   b. Specify the start and end dates.
   c. From the **Type** drop-down list, select **Backup**.
   d. From the **Status** drop-down, select the backup status.
   e. Click **Apply** to view the operations completed successfully.

4. Select a backup job, and then click **Details** to view the job details.

   Though the backup job status displays ✔, when you click on job details you might see that some of the child tasks of the backup operation are still in progress or marked with warning signs.

5. In the Job Details page, click **View logs**.

   The **View logs** button displays the detailed logs for the selected operation.

Monitor data protection operations on PostgreSQL clusters in the Activity pane

The Activity pane displays the five most recent operations performed. The Activity pane also displays when the operation was initiated and the status of the operation.

The Activity pane displays information regarding backup, restore, clone, and scheduled backup operations.

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. Click on the Activity pane to view the five most recent operations.

   When you click one of the operations, the operation details are listed in the **Job Details** page.

Cancel backup operations for PostgreSQL

You can cancel backup operations that are queued.

What you will need

- You must be logged in as the SnapCenter Admin or job owner to cancel operations.
- You can cancel a backup operation from either the **Monitor** page or the **Activity** pane.
- You cannot cancel a running backup operation.
- You can use the SnapCenter GUI, PowerShell cmdlets, or CLI commands to cancel the backup operations.
- The **Cancel Job** button is disabled for operations that cannot be canceled.
- If you selected **All members of this role can see and operate on other members objects** in Users\Groups page while creating a role, you can cancel the queued backup operations of other members while using that role.

Steps

1. Perform one of the following actions:

From the…	Action
Monitor page	a. In the left navigation pane, click **Monitor** > **Jobs**. b. Select the operation, and then click **Cancel Job**.

From the…	Action
Activity pane	a. After initiating the backup operation, click on the Activity pane to view the five most recent operations. b. Select the operation. c. In the Job Details page, click **Cancel Job**.

The operation is canceled, and the resource is reverted to the previous state.

View PostgreSQL backups and clones in the Topology page

When you are preparing to back up or clone a resource, you might find it helpful to view a graphical representation of all backups and clones on the primary and secondary storage.

About this task

You can review the following icons in the Manage Copies view to determine whether the backups and clones are available on the primary or secondary storage (Mirror copies or Vault copies).

- displays the number of backups and clones that are available on the primary storage.
- displays the number of backups and clones that are mirrored on the secondary storage using SnapMirror technology.
- displays the number of backups and clones that are replicated on the secondary storage using SnapVault technology.

The number of backups displayed includes the backups deleted from the secondary storage. For example, if you created 6 backups using a policy to retain only 4 backups, the number of backups displayed is 6.

Clones of a backup of a version-flexible mirror on a mirror-vault type volume are displayed in the topology view, but the mirror backup count in the topology view does not include the version-flexible backup.

In the Topology page, you can see all of the backups and clones that are available for the selected resource or resource group. You can view the details of those backups and clones, and then select them to perform data protection operations.

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page, either select the resource or resource group from the **View** drop-down list.
3. Select the resource either from the resource details view or from the resource group details view.

If the resource is protected, the topology page of the selected resource is displayed.

4. Review the **Summary card** to see a summary of the number of backups and clones available on the primary and secondary storage.

   The **Summary Card** section displays the total number of Snapshot copy-based backups, and clones.

   Clicking the **Refresh** button starts a query of the storage to display an accurate count.

   If SnapLock enabled backup is taken, then clicking the **Refresh** button refreshes the primary and secondary SnapLock expiry time retrieved from ONTAP. A weekly schedule also refreshes the primary and secondary SnapLock expiry time retrieved from ONTAP.

   When the application resource is spread across multiple volumes, the SnapLock expiry time for the backup will be the longest SnapLock expiry time that is set for a Snapshot in a volume. The longest SnapLock expiry time is retrieved from ONTAP.

   After on demand backup, by clicking the **Refresh** button refreshes the details of backup or clone.

5. In the Manage Copies view, click **Backups** or **Clones** from the primary or secondary storage to see details of a backup or clone.

   The details of the backups and clones are displayed in a table format.

6. Select the backup from the table, and then click the data protection icons to perform restore, clone, and delete operations.

   You cannot rename or delete backups that are on the secondary storage.

7. If you want to delete a clone, select the clone from the table, and then click .
8. If you want to split a clone, select the clone from the table, and then click .

Restore PostgreSQL

Restore workflow

The restore and recovery workflow includes planning, performing the restore operations, and monitoring the operations.

The following workflow shows the sequence in which you must perform the restore operation:

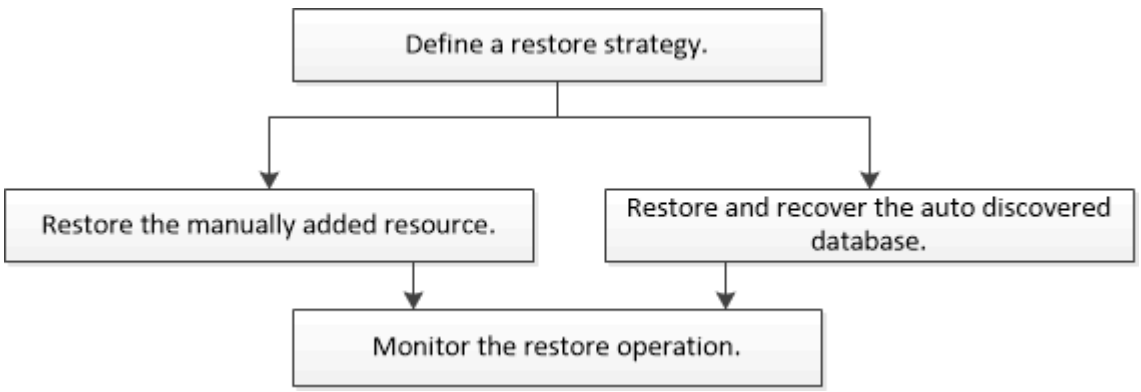

You can also use PowerShell cmdlets manually or in scripts to perform backup, restore, and clone operations. The SnapCenter cmdlet help and the cmdlet reference information contain detailed information about PowerShell cmdlets.

SnapCenter Software Cmdlet Reference Guide.

Restore and recover a manually added resource backup

You can use SnapCenter to restore and recover data from one or more backups.

Before you begin

- You must have backed up the resource or resource groups.
- You must have canceled any backup operation that is currently in progress for the resource or resource group that you want to restore.
- For pre restore, post restore, mount, and unmount commands, you should check if the commands exist in the command list available on the plug-in host from the following paths:
  - Default location on the Windows host: *C:\Program Files\NetApp\SnapCenter\Snapcenter Plug-in Creator\etc\allowed_commands.config*
  - Default location on the Linux host: */opt/NetApp/snapcenter/scc/etc/allowed_commands.config*

If the commands do not exist in the command list, then the operation will fail.

About this task

- For ONTAP 9.12.1 and below version, the clones created from the SnapLock Vault Snapshots as part of restore will inherit the SnapLock Vault expiry time. Storage admin should manually cleanup the clones post the SnapLock expiry time.

SnapCenter UI

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page, filter resources from the **View** drop-down list based on resource type.

   The resources are displayed along with type, host, associated resource groups and policies, and status.

   > Although a backup might be for a resource group, when you restore, you must select the individual resources you want to restore.

   If the resource is not protected, “Not protected” is displayed in the Overall Status column. This can mean either that the resource is not protected, or that the resource was backed up by a different user.
3. Select the resource, or select a resource group and then select a resource in that group.

   The resource topology page is displayed.
4. From the Manage Copies view, select **Backups** either from the primary or secondary (mirrored or vaulted) storage systems.
5. In the Primary backup(s) table, select the backup that you want to restore from, and then click .

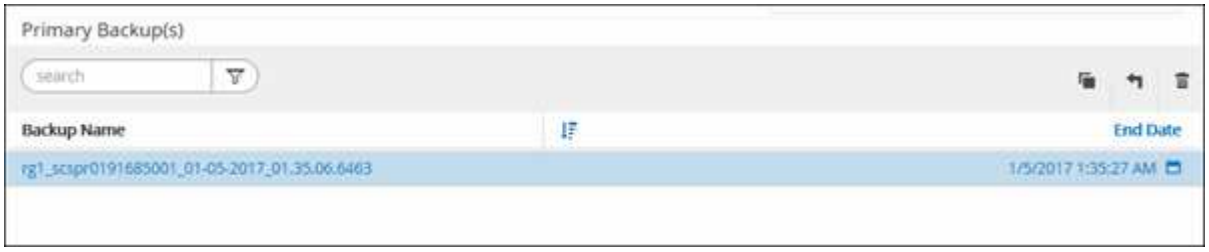

6. In the Restore Scope page, select **Complete Resource**.

   a. If you select **Complete Resource**, all of the configured data volumes of the PostgreSQL cluster are restored.

      If the resource contains volumes or qtrees, the Snapshots taken after the Snapshot selected for restore on such volumes or qtrees are deleted and cannot be recovered. Also, if any other resource is hosted on same volumes or qtrees, then that resource is also deleted.

      You can select multiple LUNs.

      > If you select **All**, all the files on the volumes, qtrees, or LUNs are restored.
7. In the Pre ops page, enter pre restore and unmount commands to run before performing a restore job.

   Unmount commands are not available for auto discovered resources.
8. In the Post ops page, enter mount and post restore commands to run after performing a restore job.

   Mount commands are not available for auto discovered resources.

9. In the Notification page, from the **Email preference** drop-down list, select the scenarios in which you want to send the emails.

   You must also specify the sender and receiver email addresses and the subject of the email. SMTP must also be configured on the **Settings** > **Global Settings** page.

10. Review the summary, and then click **Finish**.
11. Monitor the operation progress by clicking **Monitor** > **Jobs**.

PowerShell cmdlets

Steps

1. Initiate a connection session with the SnapCenter Server for a specified user by using the Open-SmConnection cmdlet.

```
PS C:\> Open-Smconnection
```

2. Retrieve the information about the one or more backups that you want to restore by using the Get-SmBackup and Get-SmBackupReport cmdlets.

   This example displays information about all available backups:

```
PS C:\> Get-SmBackup

BackupId                       BackupName
BackupTime                     BackupType
--------                       ----------
----------                     ----------
  1               Payroll Dataset_vise-f6_08... 8/4/2015    11:02:32
AM                 Full Backup
  2               Payroll Dataset_vise-f6_08... 8/4/2015    11:23:17
AM
```

   This example displays detailed information about the backup from January 29th 2015 to February 3rd, 2015:

```
PS C:\> Get-SmBackupReport -FromDateTime "1/29/2015" -ToDateTime
"2/3/2015"

   SmBackupId           : 113
   SmJobId              : 2032
   StartDateTime        : 2/2/2015 6:57:03 AM
   EndDateTime          : 2/2/2015 6:57:11 AM
   Duration             : 00:00:07.3060000
   CreatedDateTime      : 2/2/2015 6:57:23 AM
   Status               : Completed
   ProtectionGroupName  : Clone
   SmProtectionGroupId  : 34
   PolicyName           : Vault
   SmPolicyId           : 18
   BackupName           : Clone_SCSPR0019366001_02-02-2015_06.57.08
   VerificationStatus   : NotVerified

   SmBackupId           : 114
   SmJobId              : 2183
   StartDateTime        : 2/2/2015 1:02:41 PM
   EndDateTime          : 2/2/2015 1:02:38 PM
   Duration             : -00:00:03.2300000
   CreatedDateTime      : 2/2/2015 1:02:53 PM
   Status               : Completed
   ProtectionGroupName  : Clone
   SmProtectionGroupId  : 34
   PolicyName           : Vault
   SmPolicyId           : 18
   BackupName           : Clone_SCSPR0019366001_02-02-2015_13.02.45
   VerificationStatus   : NotVerified
```

3. Restore data from the backup by using the Restore-SmBackup cmdlet.

```
Restore-SmBackup -PluginCode 'DummyPlugin' -AppObjectId
'scc54.sccore.test.com\DummyPlugin\NTP\DB1' -BackupId 269
-Confirm:$false
output:
Name                : Restore
'scc54.sccore.test.com\DummyPlugin\NTP\DB1'
Id                  : 2368
StartTime           : 10/4/2016 11:22:02 PM
EndTime             :
IsCancellable       : False
IsRestartable       : False
IsCompleted         : False
IsVisible           : True
IsScheduled         : False
PercentageCompleted : 0
Description         :
Status              : Queued
Owner               :
Error               :
Priority            : None
Tasks               : {}
ParentJobID         : 0
EventId             : 0
JobTypeId           :
ApisJobKey          :
ObjectId            : 0
PluginCode          : NONE
PluginName          :
```

The information regarding the parameters that can be used with the cmdlet and their descriptions can be obtained by running *Get-Help command_name*. Alternatively, you can also refer to the SnapCenter Software Cmdlet Reference Guide.

Restore and recover an auto discovered cluster backup

You can use SnapCenter to restore and recover data from one or more backups.

Before you begin

- You must have backed up the resource or resource groups.
- You must have canceled any backup operation that is currently in progress for the resource or resource group that you want to restore.
- For pre restore, post restore, mount, and unmount commands, you should check if the commands exist in the command list available on the plug-in host from the following paths:
  - Default location on the Windows host: *C:\Program Files\NetApp\SnapCenter\Snapcenter Plug-in*

Creator\etc\allowed_commands.config

- Default location on the Linux host: */opt/NetApp/snapcenter/scc/etc/allowed_commands.config*

If the commands do not exist in the command list, then the operation will fail.

About this task

- File-based backup copies cannot be restored from SnapCenter.
- For Auto-discovered resources, restore is supported with SFSR.
- Auto-recovery is not supported.
- For ONTAP 9.12.1 and below version, the clones created from the SnapLock Vault snapshots as part of restore will inherit the SnapLock Vault expiry time. Storage admin should manually cleanup the clones post the SnapLock expiry time.

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page, filter resources from the **View** drop-down list based on resource type.

   The resources are displayed along with type, host, associated resource groups and policies, and status.

   Although a backup might be for a resource group, when you restore, you must select the individual resources you want to restore.

   If the resource is not protected, “Not protected” is displayed in the Overall Status column. This can mean either that the resource is not protected, or that the resource was backed up by a different user.

3. Select the resource, or select a resource group and then select a resource in that group.

   The resource topology page is displayed.

4. From the Manage Copies view, select **Backups** either from the primary or secondary (mirrored or vaulted) storage systems.
5. In the Primary backup(s) table, select the backup that you want to restore from, and then click .

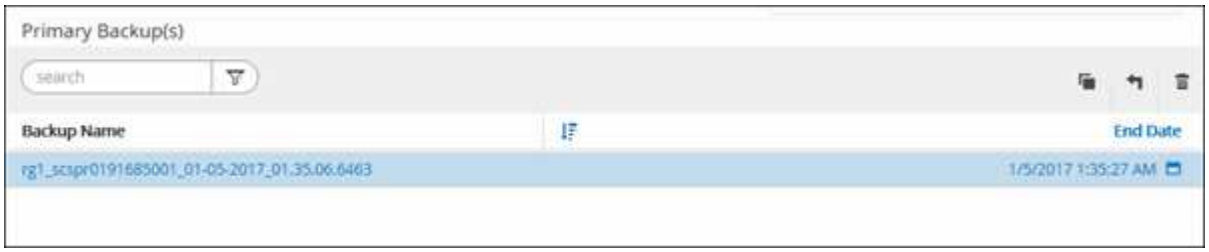

6. In the Restore Scope page, select **Complete Resource** to restore the configured data volumes of the PostgreSQL cluster.
7. In the Recovery scope page, select one of the following options:

If you…	Do this…

Want to recover as close as possible to the current time	Select **Recover to most recent state**. For single container resources specify one or more log and catalog backup locations.
Want to recover to the specified point in time	Select **Recover to point in time**. a. Enter date and time. Enter date and time. For example, the PostgreSQL Linux host is located in Sunnyvale, CA and the user in Raleigh, NC is recovering the logs in to SnapCenter. If the user wants to perform a recovery to 5 a.m .Sunnyvale, CA, then the user has to set the browser time zone to the PostgreSQL Linux host time zone, which is GMT-07:00 and specify the date and time as 5:00 a.m.
Do not want to recover	Select **No recovery**.

You cannot recover manually added PostgreSQL resources.

SnapCenter Plug-in for PostgreSQL creates a backup_label and tablespace_map in */<OS_temp_folder>/postgresql_sc_recovery*<Restore_JobId>/_ folder to help recover manually.

1. In the Pre ops page, enter pre restore and unmount commands to run before performing a restore job.

   Unmount commands are not available for auto discovered resources.

2. In the Post ops page, enter mount and post restore commands to run after performing a restore job.

   Mount commands are not available for auto discovered resources.

3. In the Notification page, from the **Email preference** drop-down list, select the scenarios in which you want to send the emails.

   You must also specify the sender and receiver email addresses and the subject of the email. SMTP must also be configured on the **Settings** > **Global Settings** page.

4. Review the summary, and then click **Finish**.
5. Monitor the operation progress by clicking **Monitor** > **Jobs**.

Restore resources using PowerShell cmdlets

Restoring a resource backup includes initiating a connection session with the SnapCenter Server, listing the backups and retrieving backup information, and restoring a backup.

You must have prepared the PowerShell environment to execute the PowerShell cmdlets.

Steps

1. Initiate a connection session with the SnapCenter Server for a specified user by using the Open-SmConnection cmdlet.

```
PS C:\> Open-Smconnection
```

2. Retrieve the information about the one or more backups that you want to restore by using the Get-SmBackup and Get-SmBackupReport cmdlets.

   This example displays information about all available backups:

```
PS C:\> Get-SmBackup

BackupId                      BackupName                    BackupTime
BackupType
--------                      ----------                    ----------
----------
  1               Payroll Dataset_vise-f6_08... 8/4/2015    11:02:32 AM
Full Backup
  2               Payroll Dataset_vise-f6_08... 8/4/2015    11:23:17 AM
```

   This example displays detailed information about the backup from January 29th 2015 to February 3rd, 2015:

```
PS C:\> Get-SmBackupReport -FromDateTime "1/29/2015" -ToDateTime
"2/3/2015"

   SmBackupId           : 113
   SmJobId              : 2032
   StartDateTime        : 2/2/2015 6:57:03 AM
   EndDateTime          : 2/2/2015 6:57:11 AM
   Duration             : 00:00:07.3060000
   CreatedDateTime      : 2/2/2015 6:57:23 AM
   Status               : Completed
   ProtectionGroupName  : Clone
   SmProtectionGroupId  : 34
   PolicyName           : Vault
   SmPolicyId           : 18
   BackupName           : Clone_SCSPR0019366001_02-02-2015_06.57.08
   VerificationStatus   : NotVerified

   SmBackupId           : 114
   SmJobId              : 2183
   StartDateTime        : 2/2/2015 1:02:41 PM
   EndDateTime          : 2/2/2015 1:02:38 PM
   Duration             : -00:00:03.2300000
   CreatedDateTime      : 2/2/2015 1:02:53 PM
   Status               : Completed
   ProtectionGroupName  : Clone
   SmProtectionGroupId  : 34
   PolicyName           : Vault
   SmPolicyId           : 18
   BackupName           : Clone_SCSPR0019366001_02-02-2015_13.02.45
   VerificationStatus   : NotVerified
```

3. Restore data from the backup by using the Restore-SmBackup cmdlet.

```
Restore-SmBackup -PluginCode 'DummyPlugin' -AppObjectId
'scc54.sccore.test.com\DummyPlugin\NTP\DB1' -BackupId 269
-Confirm:$false
output:
Name                : Restore
'scc54.sccore.test.com\DummyPlugin\NTP\DB1'
Id                  : 2368
StartTime           : 10/4/2016 11:22:02 PM
EndTime             :
IsCancellable       : False
IsRestartable       : False
IsCompleted         : False
IsVisible           : True
IsScheduled         : False
PercentageCompleted : 0
Description         :
Status              : Queued
Owner               :
Error               :
Priority            : None
Tasks               : {}
ParentJobID         : 0
EventId             : 0
JobTypeId           :
ApisJobKey          :
ObjectId            : 0
PluginCode          : NONE
PluginName          :
```

The information regarding the parameters that can be used with the cmdlet and their descriptions can be obtained by running *Get-Help command_name*. Alternatively, you can also refer to the SnapCenter Software Cmdlet Reference Guide.

Monitor PostgreSQL restore operations

You can monitor the progress of different SnapCenter restore operations by using the Jobs page. You might want to check the progress of an operation to determine when it is complete or if there is an issue.

About this task

Post-restore states describe the conditions of the resource after a restore operation and any further restore actions that you can take.

The following icons appear on the Jobs page, and indicate the state of the operation:

- In progress

- Completed successfully
- Failed
- Completed with warnings or could not start due to warnings
- Queued
- Canceled

Steps

1. In the left navigation pane, click **Monitor**.
2. In the **Monitor** page, click **Jobs**.
3. In the **Jobs** page, perform the following steps:
   a. Click to filter the list so that only restore operations are listed.
   b. Specify the start and end dates.
   c. From the **Type** drop-down list, select **Restore**.
   d. From the **Status** drop-down list, select the restore status.
   e. Click **Apply** to view the operations that have been completed successfully.
4. Select the restore job, and then click **Details** to view the job details.
5. In the **Job Details** page, click **View logs**.

   The **View logs** button displays the detailed logs for the selected operation.

Clone PostgreSQL resource backups

Clone workflow

The clone workflow includes performing the clone operation and monitoring the operation.

About this task

- You can clone on the source PostgreSQL server.
- You might clone resource backups for the following reasons:
  - To test functionality that has to be implemented using the current resource structure and content during application development cycles
  - For data extraction and manipulation tools when populating data warehouses
  - To recover data that was mistakenly deleted or changed

The following workflow shows the sequence in which you must perform the clone operation:

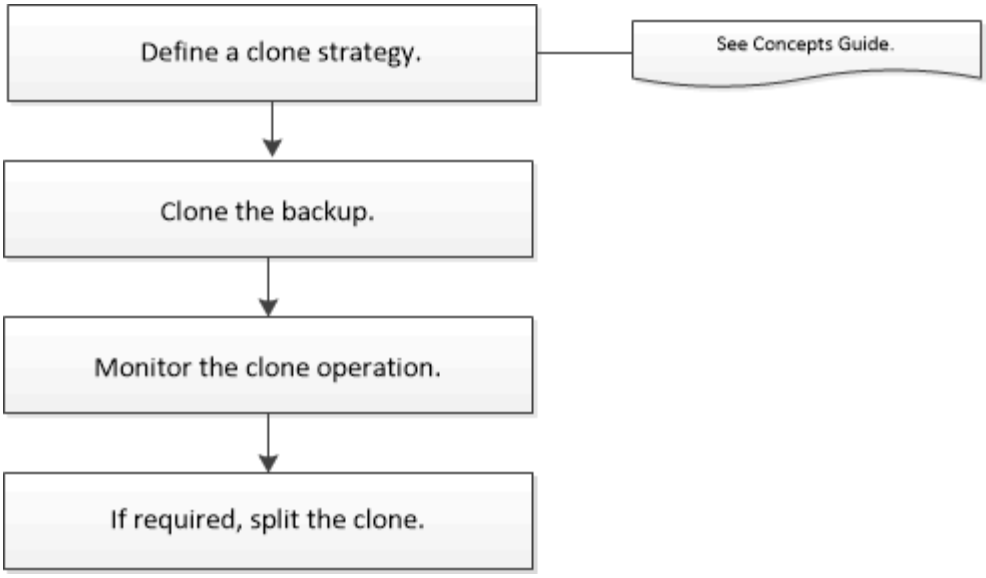

You can also use PowerShell cmdlets manually or in scripts to perform backup, restore, and clone operations. The SnapCenter cmdlet help and the cmdlet reference information contain detailed information about PowerShell cmdlets.

Clone a PostgreSQL backup

You can use SnapCenter to clone a backup. You can clone from primary or secondary backup.

Before you begin

- You should have backed up the resources or resource group.
- You should ensure that the aggregates hosting the volumes should be in the assigned aggregates list of the storage virtual machine (SVM).
- For pre clone or post clone commands, you should check if the commands exist in the command list available on the plug-in host from the following paths:
  - Default location on the Windows host: *C:\Program Files\NetApp\SnapCenter\Snapcenter Plug-in Creator\etc\allowed_commands.config*
  - Default location on the Linux host: */opt/NetApp/snapcenter/scc/etc/allowed_commands.config*

If the commands do not exist in the command list, then the operation will fail.

About this task

- For information about FlexClone volume split operations, see, Split a FlexClone volume from its parent volume.
- For ONTAP 9.12.1 and below version, the clones created from the SnapLock Vault Snapshots as part of restore will inherit the SnapLock Vault expiry time. Storage admin should manually cleanup the clones post the SnapLock expiry time.

SnapCenter UI

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page, filter resources from the **View** drop-down list based on resource type.

   The resources are displayed along with information such as type, host, associated resource groups and policies, and status.

3. Select the resource or resource group.

   You must select a resource if you select a resource group.

   The resource or resource group topology page is displayed.

4. From the Manage Copies view, select **Backups** either from the primary or secondary (mirrored or vaulted) storage systems.
5. Select the data backup from the table, and then click .
6. In the Location page, perform the following actions:

For this field…	Do this…
Clone server	Choose a host on which the clone should be created.
Target Port	Enter the target PostgreSQL target port to clone from the existing backups.
NFS Export IP Address	Enter IP addresses or the host names on which the cloned volumes will be exported. This is applicable only to NFS storage type resource.
Capacity Pool Max. Throughput (MiB/s)	Enter the maximum throughput of a capacity pool. This is applicable only for ANF storage type resource.

7. In the Scripts page, perform the following steps:

   The scripts are run on the plug-in host.

   a. Enter the commands for pre clone or post clone that should be run before or after the clone operation, respectively.
      - Pre clone command: delete existing clusters with the same name
      - Post clone command: verify a cluster or start a cluster.

b. Enter the mount command to mount a file system to a host.

   Mount command for a volume or qtree on a Linux machine:

   Example for NFS: `mount VSERVER_DATA_IP:%VOLUME_NAME_Clone /mnt`

8. In the Notification page, from the **Email preference** drop-down list, select the scenarios in which you want to send the emails.

   You must also specify the sender and receiver email addresses, and the subject of the email.

9. Review the summary, and then click **Finish**.
10. Monitor the operation progress by clicking **Monitor** > **Jobs**.

PowerShell cmdlets

Steps

1. Initiate a connection session with the SnapCenter Server for a specified user by using the Open-SmConnection cmdlet.

```
PS C:\> Open-SmConnection
```

2. Retrieve the backups to perform the clone operation by using the Get-SmBackup cmdlet.

   This example shows that two backups are available for cloning:

```
 C:\PS> Get-SmBackup

        BackupId                    BackupName
BackupTime                  BackupType
        --------                    ----------
----------                  ----------
         1                   Payroll Dataset_vise-f6_08...
8/4/2015 11:02:32 AM        Full Backup
         2                   Payroll Dataset_vise-f6_08...
8/4/2015 11:23:17 AM
```

3. Initiate a clone operation from an existing backup and specify the NFS export IP addresses on which the cloned volumes are exported.

   This example shows that the backup to be cloned has an NFSExportIPs address of 10.32.212.14:

   For PostgreSQL cluster:

```
PS C:\> New-SmClone -AppPluginCode PostgreSQL -BackupName "
scpostgresql01_ openenglab_netapp_com_PostgreSQL_postgres_5432_06-
26-2024_00_33_41_1570" -Resources @{"Host"="
10.32.212.13";"Uid"="postgres_5432"} -port 2345 -CloneToHost
10.32.212.14
```

If NFSExportIPs is not specified, the default is exported to the clone target host.

4. Verify that the backups were cloned successfully by using the Get-SmCloneReport cmdlet to view the clone job details.

   You can view details such as clone ID, start date and time, end date and time.

```
PS C:\> Get-SmCloneReport -JobId 186

    SmCloneId           : 1
    SmJobId             : 186
    StartDateTime       : 8/3/2015 2:43:02 PM
    EndDateTime         : 8/3/2015 2:44:08 PM
    Duration            : 00:01:06.6760000
    Status              : Completed
    ProtectionGroupName : Draper
    SmProtectionGroupId : 4
    PolicyName          : OnDemand_Clone
    SmPolicyId          : 4
    BackupPolicyName    : OnDemand_Full_Log
    SmBackupPolicyId    : 1
    CloneHostName       : SCSPR0054212005.mycompany.com
    CloneHostId         : 4
    CloneName           : Draper__clone__08-03-2015_14.43.53
    SourceResources     : {Don, Betty, Bobby, Sally}
    ClonedResources     : {Don_DRAPER, Betty_DRAPER, Bobby_DRAPER,
Sally_DRAPER}
    SmJobError          :
```

Monitor PostgreSQL clone operations

You can monitor the progress of SnapCenter clone operations by using the Jobs page. You might want to check the progress of an operation to determine when it is complete or if there is an issue.

About this task

The following icons appear on the Jobs page, and indicate the state of the operation:

- In progress
- Completed successfully
- Failed
- Completed with warnings or could not start due to warnings
- Queued
- Canceled

Steps

1. In the left navigation pane, click **Monitor**.
2. In the **Monitor** page, click **Jobs**.
3. In the **Jobs** page, perform the following steps:
   a. Click to filter the list so that only clone operations are listed.
   b. Specify the start and end dates.
   c. From the **Type** drop-down list, select **Clone**.
   d. From the **Status** drop-down list, select the clone status.
   e. Click **Apply** to view the operations that are completed successfully.
4. Select the clone job, and then click **Details** to view the job details.
5. In the Job Details page, click **View logs**.

Split a clone

You can use SnapCenter to split a cloned resource from the parent resource. The clone that is split becomes independent of the parent resource.

About this task

- You cannot perform the clone split operation on an intermediate clone.

  For example, after you create clone1 from a database backup, you can create a backup of clone1, and then clone this backup (clone2). After you create clone2, clone1 is an intermediate clone, and you cannot perform the clone split operation on clone1. However, you can perform the clone split operation on clone2.

  After splitting clone2, you can perform the clone split operation on clone1 because clone1 is no longer the intermediate clone.

- When you split a clone, the backup copies and clone jobs of the clone are deleted.
- For information about FlexClone volume split operations, see, Split a FlexClone volume from its parent volume.
- Ensure that the volume or aggregate on the storage system is online.

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. In the **Resources** page, select the appropriate option from the View list:

Option	Description
For database applications	Select **Database** from the View list.
For file systems	Select **Path** from the View list.

3. Select the appropriate resource from the list.

   The resource topology page is displayed.

4. From the **Manage Copies** view, select the cloned resource (for example, the database or LUN), and then click .
5. Review the estimated size of the clone that is to be split and the required space available on the aggregate, and then click **Start**.
6. Monitor the operation progress by clicking **Monitor** > **Jobs**.

   The clone split operation stops responding if the SMCore service restarts. You should run the Stop-SmJob cmdlet to stop the clone split operation, and then retry the clone split operation.

   If you want a longer poll time or shorter poll time to check whether the clone is split or not, you can change the value of *CloneSplitStatusCheckPollTime* parameter in *SMCoreServiceHost.exe.config* file to set the time interval for SMCore to poll for the status of the clone split operation. The value is in milliseconds and the default value is 5 minutes.

   For example:

   ```
   <add key="CloneSplitStatusCheckPollTime" value="300000" />
   ```

   The clone split start operation fails if backup, restore, or another clone split is in progress. You should restart the clone split operation only after the running operations are complete.

Related information

SnapCenter clone or verification fails with aggregate does not exist

Delete or split PostgreSQL cluster clones after upgrading SnapCenter

After upgrading to SnapCenter 4.3, you will no longer see the clones. You can delete the clone or split the clones from the Topology page of the resource from which the clones were created.

About this task

If you want to locate the storage footprint of the hidden clones, run the following command: `Get-SmClone -ListStorageFootprint`

Steps

1. Delete the backups of the cloned resources by using the remove-smbackup cmdlet.
2. Delete the resource group of the cloned resources by using the remove-smresourcegroup cmdlet.

3. Remove the protection of the cloned resource by using the remove-smprotectresource cmdlet.
4. Select the parent resource from the Resources page.

   The resource topology page is displayed.

5. From the Manage Copies view, select the clones either from the primary or secondary (mirrored or replicated) storage systems.
6. Select the clones, and then click to delete clones or click to split the clones.
7. Click **OK**.

Protect MySQL

SnapCenter Plug-in for MySQL

SnapCenter Plug-in for MySQL overview

The SnapCenter Plug-in for MySQL Database is a host-side component of the NetApp SnapCenter software that enables application-aware data protection management of MySQL databases. The Plug-in for MySQL Database automates the backup, restore, and cloning of MySQL databases in your SnapCenter environment.

SnapCenter supports single instance MySQL setups. You can use the Plug-in for MySQL Database in both Linux and Windows environments. In Windows environments, MySQL will be supported as manual resource.

When the Plug-in for MySQL Database is installed, you can use SnapCenter with NetApp SnapMirror technology to create mirror copies of backup sets on another volume. You can also use the plug-in with NetApp SnapVault technology to perform disk-to-disk backup replication for standards compliance.

SnapCenter Plug-in for MySQL supports NFS and SAN on ONTAP and Azure NetApp File storage layouts.

VMDK, vVol, and RDM virtual storage layout is supported.

Symbolic links are not supported.

What you can do using the SnapCenter Plug-in for MySQL

When you install the Plug-in for MySQL Database in your environment, you can use SnapCenter to back up, restore, and clone MySQL instances. You can also perform tasks supporting those operations.

- Add instances.
- Create backups.
- Restore from backups.
- Clone backups.
- Schedule backup operations.
- Monitor backup, restore, and clone operations.
- View reports for backup, restore, and clone operations.

SnapCenter Plug-in for MySQL features

SnapCenter integrates with the plug-in application and with NetApp technologies on the storage system. To work with the Plug-in for MySQL Database, you use the SnapCenter graphical user interface.

- **Unified graphical user interface**

  The SnapCenter interface provides standardization and consistency across plug-ins and environments. The SnapCenter interface enables you to complete consistent backup, restore, and clone operations

across plug-ins, use centralized reporting, use at-a-glance dashboard views, set up role-based access control (RBAC), and monitor jobs across all plug-ins.

- **Automated central administration**

  You can schedule backup operations, configure policy-based backup retention, and perform restore operations. You can also proactively monitor your environment by configuring SnapCenter to send email alerts.

- **Nondisruptive NetApp snapshot copy technology**

  SnapCenter uses NetApp snapshot technology with the Plug-in for MySQL Database to back up resources.

Using the Plug-in for MySQL also offers the following benefits:

- Support for backup, restore, and clone workflows
- RBAC-supported security and centralized role delegation

  You can also set the credentials so that the authorized SnapCenter users have application-level permissions.

- Creation of space-efficient and point-in-time copies of resources for testing or data extraction by using NetApp FlexClone technology

  A FlexClone license is required on the storage system where you want to create the clone.

- Support for the consistency group (CG) snapshot feature of ONTAP as part of creating backups.
- Capability to run multiple backups simultaneously across multiple resource hosts

  In a single operation, snapshots are consolidated when resources in a single host share the same volume.

- Capability to create snapshots using external commands.
- Support for Linux LVM on XFS file system.

Storage types supported by SnapCenter Plug-in for MySQL

SnapCenter supports a wide range of storage types on both physical machines and virtual machines (VMs). You must verify the support for your storage type before installing SnapCenter Plug-in for MySQL.

Machine	Storage type
Physical server	• FC-connected LUNs • iSCSI-connected LUNs • NFS-connected volumes

Machine	Storage type
VMware ESXi	• RDM LUNs connected by an FC or iSCSI ESXi HBAScanning of host bus adapters (HBAs) might take long time to complete because SnapCenter scans all the host bus adaptors present in the host. You can edit the **LinuxConfig.pm** file located at */opt/NetApp/snapcenter/spl/plugins/scu/scucore/modules/SCU/Config* to set the value of the **SCSI_HOSTS_OPTIMIZED_RESCAN** parameter to 1 to rescan only those HBAs that are listed in HBA_DRIVER_NAMES. • iSCSI LUNs connected directly to the guest system by the iSCSI initiator • VMDKs on NFS datastores • VMDKs on VMFS created • NFS volumes connected directly to the guest system • vVol datastores on both NFS and SAN vVol datastore can only be provisioned with ONTAP Tools for VMware vSphere.

Minimum ONTAP privileges required for MySQL plug-in

The minimum ONTAP privileges that are required vary according to the SnapCenter plug-ins you are using for data protection.

- All-access commands: Minimum privileges required for ONTAP 9.12.1 and later
  - event generate-autosupport-log
  - job history show
  - job stop
  - lun
  - lun create
  - lun create
  - lun create
  - lun delete
  - lun igroup add
  - lun igroup create
  - lun igroup delete
  - lun igroup rename

- lun igroup rename
- lun igroup show
- lun mapping add-reporting-nodes
- lun mapping create
- lun mapping delete
- lun mapping remove-reporting-nodes
- lun mapping show
- lun modify
- lun move-in-volume
- lun offline
- lun online
- lun persistent-reservation clear
- lun resize
- lun serial
- lun show
- snapmirror policy add-rule
- snapmirror policy modify-rule
- snapmirror policy remove-rule
- snapmirror policy show
- snapmirror restore
- snapmirror show
- snapmirror show-history
- snapmirror update
- snapmirror update-ls-set
- snapmirror list-destinations
- version
- volume clone create
- volume clone show
- volume clone split start
- volume clone split stop
- volume create
- volume destroy
- volume file clone create
- volume file show-disk-usage
- volume offline
- volume online
- volume modify

    - volume qtree create
    - volume qtree delete
    - volume qtree modify
    - volume qtree show
    - volume restrict
    - volume show
    - volume snapshot create
    - volume snapshot delete
    - volume snapshot modify
    - volume snapshot modify-snaplock-expiry-time
    - volume snapshot rename
    - volume snapshot restore
    - volume snapshot restore-file
    - volume snapshot show
    - volume unmount
    - vserver cifs
    - vserver cifs share create
    - vserver cifs share delete
    - vserver cifs shadowcopy show
    - vserver cifs share show
    - vserver cifs show
    - vserver export-policy
    - vserver export-policy create
    - vserver export-policy delete
    - vserver export-policy rule create
    - vserver export-policy rule show
    - vserver export-policy show
    - vserver iscsi
    - vserver iscsi connection show
    - vserver show
- Read-only commands: Minimum privileges required for ONTAP 8.3.0 and later
    - network interface
    - network interface show
    - vserver

Prepare storage systems for SnapMirror and SnapVault replication for MySQL

You can use a SnapCenter plug-in with ONTAP SnapMirror technology to create mirror

copies of backup sets on another volume, and with ONTAP SnapVault technology to perform disk-to-disk backup replication for standards compliance and other governance-related purposes. Before you perform these tasks, you must configure a data-protection relationship between the source and destination volumes and initialize the relationship.

SnapCenter performs the updates to SnapMirror and SnapVault after it completes the Snapshot operation. SnapMirror and SnapVault updates are performed as part of the SnapCenter job; do not create a separate ONTAP schedule.

If you are coming to SnapCenter from a NetApp SnapManager product and are satisfied with the data protection relationships you have configured, you can skip this section.

A data protection relationship replicates data on primary storage (the source volume) to secondary storage (the destination volume). When you initialize the relationship, ONTAP transfers the data blocks referenced on the source volume to the destination volume.

SnapCenter does not support cascade relationships between SnapMirror and SnapVault volumes (**Primary** > **Mirror** > **Vault**). You should use fanout relationships.

SnapCenter supports the management of version-flexible SnapMirror relationships. For details about version-flexible SnapMirror relationships and how to set them up, see the ONTAP documentation.

Backup strategy for MySQL

Define a backup strategy for MySQL

Defining a backup strategy before you create your backup jobs helps you to have the backups that you require to successfully restore or clone your resources. Your service-level agreement (SLA), recovery time objective (RTO), and recovery point objective (RPO) largely determine your backup strategy.

About this task

An SLA defines the level of service that is expected and addresses many service-related issues, including the availability and performance of the service. RTO is the time by which a business process must be restored after a disruption in service. RPO defines the strategy for the age of the files that must be recovered from backup storage for regular operations to resume after a failure. SLA, RTO, and RPO contribute to the data protection strategy.

Steps

1. Determine when you should back up your resources.
2. Decide how many backup jobs you require.
3. Decide how to name your backups.
4. Decide whether you want to create a Snapshot copy-based policy to back up application-consistent snapshots of the database.
5. Decide whether you want to use NetApp SnapMirror technology for replication or NetApp SnapVault technology for long-term retention.
6. Determine the retention period for the snapshots on the source storage system and the SnapMirror destination.

7. Determine whether you want to run any commands before or after the backup operation and provide a prescript or postscript.

Automatic discovery of resources on Linux host

Resources are MySQL instances on the Linux host that are managed by SnapCenter. After installing the SnapCenter Plug-in for MySQL plug-in, the MySQL instances on that Linux host are automatically discovered and displayed in the Resources page.

Type of backups supported

Backup type specifies the type of backup that you want to create. SnapCenter supports snapshot copy-based backup type for MySQL databases.

Snapshot copy based backup

Snapshot copy-based backups leverage NetApp snapshot technology to create online, read-only copies of the volumes on which the MySQL databases reside.

How SnapCenter Plug-in for MySQL uses consistency group snapshots

You can use the plug-in to create consistency group snapshots for resource groups. A consistency group is a container that can house multiple volumes so that you can manage them as one entity. A consistency group is simultaneous snapshots of multiple volumes, providing consistent copies of a group of volumes.

You can also specify the wait time for the storage controller to consistently group snapshots. The available wait time options are **Urgent**, **Medium**, and **Relaxed**. You can also enable or disable Write Anywhere File Layout (WAFL) sync during consistent group snapshot operation. WAFL sync improves the performance of a consistency group snapshot.

How SnapCenter manages housekeeping of log backups

SnapCenter manages the housekeeping of data backups on the storage system and file system levels.

Considerations for determining backup schedules for MySQL

The most critical factor in determining a backup schedule is the rate of change for the resource. You might back up a heavily used resource every hour, while you might back up a rarely used resource once a day. Other factors include the importance of the resource to your organization, your service-level agreement (SLA) and your recovery point objective (RPO).

Backup schedules have two parts, as follows:

- Backup frequency (how often backups are to be performed)

  Backup frequency, also called schedule type for some plug-ins, is part of a policy configuration. For example, you might configure the backup frequency as hourly, daily, weekly, or monthly.

- Backup schedules (exactly when backups are to be performed)

  Backup schedules are part of a resource or resource group configuration. For example, if you have a resource group that has a policy configured for weekly backups, you might configure the schedule to back up every Thursday at 10:00 p.m.

Number of backup jobs needed for MySQL

Factors that determine the number of backup jobs that you need include the size of the resource, the number of volumes used, the rate of change of the resource, and your Service Level Agreement (SLA).

Backup naming conventions for Plug-in for MySQL databases

You can either use the default Snapshot naming convention or use a customized naming convention. The default backup naming convention adds a timestamp to Snapshot names that helps you identify when the copies were created.

The Snapshot uses the following default naming convention:

```
resourcegroupname_hostname_timestamp
```

You should name your backup resource groups logically, as in the following example:

```
dts1_mach1x88_03-12-2015_23.17.26
```

In this example, the syntax elements have the following meanings:

- *dts1* is the resource group name.
- *mach1x88* is the host name.
- *03-12-2015_23.17.26* is the date and timestamp.

Alternatively, you can specify the Snapshot name format while protecting resources or resource groups by selecting **Use custom name format for Snapshot copy**. For example, customtext_resourcegroup_policy_hostname or resourcegroup_hostname. By default, the time stamp suffix is added to the Snapshot name.

Restore and recovery strategy for MySQL

Define a restore and recovery strategy for MySQL resources

You must define a strategy before you restore and recover your database so that you can perform restore and recovery operations successfully.

Only manual recovery of database is supported.

Steps

1. Determine the restore strategies supported for manually added MySQL resources

2. Determine the restore strategies supported for auto discovered MySQL databases
3. Decide the type of recovery operations that you want to perform.

Types of restore strategies supported for manually added MySQL resources

You must define a strategy before you can successfully perform restore operations using SnapCenter. There are two types of restore strategies for manually added MySQL resources.

You cannot recover manually added MySQL resources.

Complete resource restore

- Restores all volumes, qtrees, and LUNs of a resource

If the resource contains volumes or qtrees, the snapshots taken after the snapshot selected for restore on such volumes or qtrees are deleted and cannot be recovered. Also, if any other resource is hosted on the same volumes or qtrees, then that resource is also deleted.

Type of restore strategy supported for automatically discovered MySQL

You must define a strategy before you can successfully perform restore operations using SnapCenter.

Complete resource restore is the restore strategy supported for automatically discovered MySQL databases. This restores all the volumes, qtrees, and LUNs of a resource.

Types of restore operations for auto discovered MySQL

SnapCenter Plug-in for MySQL supports Single File SnapRestore, and connect-and-copy restore types for automatically discovered MySQL databases.

Single File SnapRestore is performed in NFS environments for the following scenarios:

- If only the **Complete Resource** option is selected
- When the backup selected is from a SnapMirror or SnapVault secondary location, and the **Complete Resource** option is selected

Single File SnapRestore is performed in SAN environments for the following scenarios:

- If only the **Complete Resource** option is selected
- When the backup is selected from a SnapMirror or SnapVault secondary location, and the **Complete Resource** option is selected

Prepare to install the SnapCenter Plug-in for MySQL

Installation workflow of SnapCenter Plug-in for MySQL

You should install and set up the SnapCenter Plug-in for MySQL if you want to protect

MySQL databases.

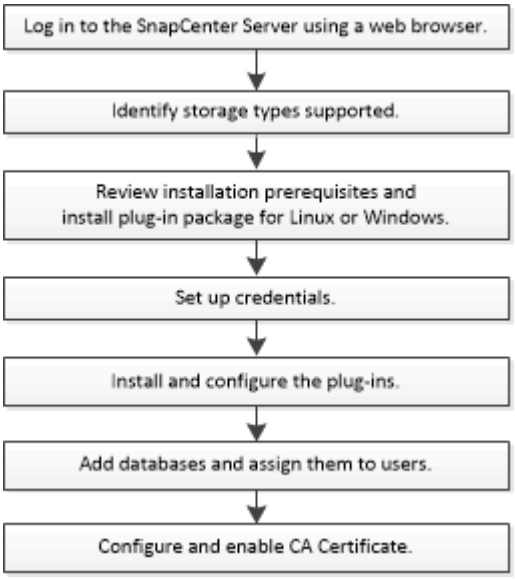

Prerequisites to add hosts and install SnapCenter Plug-in for MySQL

Before you add a host and install the plug-in packages, you must complete all the requirements. SnapCenter Plug-in for MySQL is available in both Windows and Linux environments.

- You must have installed Java 11 on your host.

  IBM Java is not supported on Windows and Linux host.

- For Windows, plug-in Creator Service should be running using the “LocalSystem” windows user, which is the default behavior when Plug-in for MySQL is installed as domain administrator.
- When installing a plug-in on a Windows host, if you specify a credential that is not built-in or if the user belongs to a local workgroup user, you must disable UAC on the host. SnapCenter Plug-in for Microsoft Windows will be deployed by default with the MySQL plug-in on Windows hosts.
- SnapCenter Server should have access to the 8145 or custom port of Plug-in for MySQL host.
- For MySQL 5.7, binlog should be specified in mysql config (my.cnf or mysql-server.cnf) file.

Windows hosts

- You must have a domain user with local administrator privileges with local login permissions on the remote host.
- While installing Plug-in for MySQL on a Windows host, SnapCenter Plug-in for Microsoft Windows is installed automatically.
- You must have enabled the password-based SSH connection for the root or non-root user.
- You must have installed Java 11 on your Windows host.

  Download JAVA for All Operating Systems

  NetApp Interoperability Matrix Tool

Linux hosts

- You must have enabled the password-based SSH connection for the root or non-root user.
- You must have installed Java 11 on your Linux host.

  Download JAVA for All Operating Systems

  NetApp Interoperability Matrix Tool

- For MySQL databases that are running on a Linux host, while installing Plug-in for MySQL, SnapCenter Plug-in for UNIX is installed automatically.
- You should have **bash** as the default shell for plug-in installation.

Supplemental commands

To run a supplemental command on the SnapCenter Plug-in for MySQL, you must include it in the *allowed_commands.config* file.

- Default location on the Windows host: *C:\Program Files\NetApp\SnapCenter\Snapcenter Plug-in Creator\etc\allowed_commands.config*
- Default location on the Linux host: */opt/NetApp/snapcenter/scc/etc/allowed_commands.config*

To allow supplemental commands on the plug-in host, open *allowed_commands.config* file in an editor. Enter each command on a separate line and the commands are not case sensitive. Ensure that you specify the fully qualified pathname and enclose the pathname in quotation marks (") if it contains spaces.

For example:

```
command: mount
command: umount
command: "C:\Program Files\NetApp\SnapCreator commands\sdcli.exe"
command: myscript.bat
```

If the *allowed_commands.config* file is not present, the commands or script execution will be blocked and the workflow will fail with the following error:

"[/mnt/mount -a] execution not allowed. Authorize by adding the command in the file %s on the plugin host."

If the command or script is not present in the *allowed_commands.config*, the command or script execution will be blocked and the workflow will fail with the following error:

"[/mnt/mount -a] execution not allowed. Authorize by adding the command in the file %s on the plugin host."

You should not use a wildcard entry (*) to allow all commands.

Configure sudo privileges for non-root users for Linux host

SnapCenter allows a non-root user to install the SnapCenter Plug-ins Package for Linux and to start the plug-in process. The plug-in processes will be running as an effective non-root user. You should configure sudo privileges for the non-root user to provide access to several paths.

What you will need

- Sudo version 1.8.7 or later.
- If the umask is 0027, ensure that the java folder and all the files inside should have permission of 555. Otherwise plug-in installation might fail.
- For the non-root user, ensure that the name of the non-root user and the user's group should be the same.
- Edit the */etc/ssh/sshd_config* file to configure the message authentication code algorithms: MACs hmac-sha2-256 and MACs hmac-sha2-512.

  Restart the sshd service after updating the configuration file.

  Example:

```
#Port 22
#AddressFamily any
#ListenAddress 0.0.0.0
#ListenAddress ::
#Legacy changes
#KexAlgorithms diffie-hellman-group1-sha1
#Ciphers aes128-cbc
#The default requires explicit activation of protocol
Protocol 2
HostKey/etc/ssh/ssh_host_rsa_key
MACs hmac-sha2-256
```

About this task

You should configure sudo privileges for the non-root user to provide access to the following paths:

- /home/*LINUX_USER*/.sc_netapp/snapcenter_linux_host_plugin.bin
- /custom_location/NetApp/snapcenter/spl/installation/plugins/uninstall
- /custom_location/NetApp/snapcenter/spl/bin/spl

Steps

1. Log in to the Linux host on which you want to install the SnapCenter Plug-ins Package for Linux.
2. Add the following lines to the /etc/sudoers file by using the visudo Linux utility.

```
Cmnd_Alias HPPLCMD = sha224:checksum_value== /home/
LINUX_USER/.sc_netapp/snapcenter_linux_host_plugin.bin,
/opt/NetApp/snapcenter/spl/installation/plugins/uninstall,
/opt/NetApp/snapcenter/spl/bin/spl, /opt/NetApp/snapcenter/scc/bin/scc
Cmnd_Alias PRECHECKCMD = sha224:checksum_value== /home/
LINUX_USER/.sc_netapp/Linux_Prechecks.sh
Cmnd_Alias CONFIGCHECKCMD = sha224:checksum_value==
/opt/NetApp/snapcenter/spl/plugins/scu/scucore/configurationcheck/Config
_Check.sh
Cmnd_Alias SCCMD = sha224:checksum_value==
/opt/NetApp/snapcenter/spl/bin/sc_command_executor
Cmnd_Alias SCCCMDEXECUTOR =checksum_value==
/opt/NetApp/snapcenter/scc/bin/sccCommandExecutor
LINUX_USER ALL=(ALL) NOPASSWD:SETENV: HPPLCMD, PRECHECKCMD,
CONFIGCHECKCMD, SCCCMDEXECUTOR, SCCMD
Defaults: LINUX_USER env_keep += "IATEMPDIR"
Defaults: LINUX_USER env_keep += "JAVA_HOME"
Defaults: LINUX_USER !visiblepw
Defaults: LINUX_USER !requiretty
```

If you are having a RAC setup, along with the other allowed commands, you should add the following to the /etc/sudoers file: '/<crs_home>/bin/olsnodes'

You can obtain the value of *crs_home* from the */etc/oracle/olr.loc* file.

LINUX_USER is the name of the non-root user that you created.

You can obtain the *checksum_value* from the **sc_unix_plugins_checksum.txt** file, which is located at:

- _C:\ProgramData\NetApp\SnapCenter\Package Repository\sc_unix_plugins_checksum.txt
  _ if SnapCenter Server is installed on Windows host.
- _/opt/NetApp/snapcenter/SnapManagerWeb/Repository/sc_unix_plugins_checksum.txt
  _ if SnapCenter Server in installed on Linux host.

The example should be used only as a reference for creating your own data.

Host requirements to install SnapCenter Plug-ins Package for Windows

Before you install the SnapCenter Plug-ins Package for Windows, you should be familiar with some basic host system space requirements and sizing requirements.

Item	Requirements
Operating Systems	Microsoft Windows For the latest information about supported versions, see the NetApp Interoperability Matrix Tool.
Minimum RAM for the SnapCenter plug-in on host	1 GB
Minimum install and log space for the SnapCenter plug-in on host	5 GB You should allocate sufficient disk space and monitor the storage consumption by the logs folder. The log space required varies depending on the number of the entities to be protected and the frequency of data protection operations. If there is no sufficient disk space, the logs will not be created for the recently run operations.
Required software packages	• ASP.NET Core Runtime 8.0.12 (and all subsequent 8.0.x patches) Hosting Bundle • PowerShell Core 7.4.2 For the latest information about supported versions, see the NetApp Interoperability Matrix Tool. For .NET specific troubleshooting information, see SnapCenter upgrade or install fails for legacy systems that do not have internet connectivity.

Host requirements for installing the SnapCenter Plug-ins Package for Linux

Before you install the SnapCenter Plug-ins Package for Linux, you should be familiar with some basic host system space and sizing requirements.

Item	Requirements
Operating systems	• Red Hat Enterprise Linux • SUSE Linux Enterprise Server (SLES) For the latest information about supported versions, see the NetApp Interoperability Matrix Tool.
Minimum RAM for the SnapCenter plug-in on host	1 GB

Item	Requirements
Minimum install and log space for the SnapCenter plug-in on host	2 GB You should allocate sufficient disk space and monitor the storage consumption by the logs folder. The log space required varies, depending on the number of the entities to be protected and the frequency of data protection operations. If there is no sufficient disk space, the logs will not be created for the recently run operations.
Required software packages	Java 11 Oracle Java and OpenJDK If you have upgraded JAVA to the latest version, you must ensure that the JAVA_HOME option located at /var/opt/snapcenter/spl/etc/spl.properties is set to the correct JAVA version and the correct path. For the latest information about supported versions, see the NetApp Interoperability Matrix Tool.

Set up credentials for the SnapCenter Plug-in for MySQL

SnapCenter uses credentials to authenticate users for SnapCenter operations. You should create credentials for installing SnapCenter plug-ins and additional credentials for performing data protection operations on databases or Windows file systems.

About this task

- Linux hosts

  You must set up credentials for installing plug-ins on Linux hosts.

  You must set up the credentials for the root user or for a non-root user who has sudo privileges to install and start the plug-in process.

  Best Practice: Although you are allowed to create credentials for Linux after deploying hosts and installing plug-ins, the best practice is to create credentials after you add SVMs, before you deploy hosts and install plug-ins.

- Windows hosts

  You must set up Windows credentials before installing plug-ins.

  You must set up the credentials with administrator privileges, including administrator rights on the remote host.

If you set up credentials for individual resource groups and the username does not have full admin privileges,

you must assign at least the resource group and backup privileges to the username.

Steps

1. In the left navigation pane, click **Settings**.
2. In the Settings page, click **Credential**.
3. Click **New**.
4. In the Credential page, specify the information required for configuring credentials:

For this field…	Do this…
Credential name	Enter a name for the credentials.
User name	Enter the user name and password that are to be used for authentication. • Domain administrator or any member of the administrator group Specify the domain administrator or any member of the administrator group on the system on which you are installing the SnapCenter plug-in. Valid formats for the Username field are: ◦ *NetBIOS\UserName* ◦ *Domain FQDN\UserName* • Local administrator (for workgroups only) For systems that belong to a workgroup, specify the built-in local administrator on the system on which you are installing the SnapCenter plug-in. You can specify a local user account that belongs to the local administrators group if the user account has elevated privileges or the User Access control feature is disabled on the host system. The valid format for the Username field is: *UserName* Do not use double quotes (") or backtick (`) in the passwords. You should not use the less than (<) and exclamation (!) symbols together in passwords. For example, lessthan<!10, lessthan10<!, backtick`12.
Password	Enter the password used for authentication.
Authentication Mode	Select the authentication mode that you want to use.

For this field…	Do this…
Use sudo privileges	Select the **Use sudo privileges** check box if you are creating credentials for a non-root user. Applicable to Linux users only.

5. Click **OK**.

After you finish setting up the credentials, you might want to assign credential maintenance to a user or group of users in the User and Access page.

Install the SnapCenter Plug-in for MySQL

Add hosts and install plug-in packages on remote hosts

You must use the SnapCenter Add Host page to add hosts, and then install the plug-ins packages. The plug-ins are automatically installed on the remote hosts. You can add the host and install plug-in packages for an individual host.

Before you begin

- If the operating system of the SnapCenter Server host is Windows 2019 and the operating system of the plug-in host is Windows 2022, you should perform the following:
  - Upgrade to Windows Server 2019 (OS Build 17763.5936) or later
  - Upgrade to Windows Server 2022 (OS Build 20348.2402) or later
- You must be a user that is assigned to a role that has the plug-in install and uninstall permissions, such as the SnapCenter Admin role.
- When installing a plug-in on a Windows host, if you specify a credential that is not built-in, or if the user belongs to a local workgroup user, you must disable UAC on the host.
- You should ensure that the message queueing service is running.
- The administration documentation contains information about managing hosts.

About this task

- You cannot add a SnapCenter Server as a plug-in host to another SnapCenter Server.

Steps

1. In the left navigation pane, click **Hosts**.
2. Verify that the **Managed Hosts** tab is selected at the top.
3. Click **Add**.
4. In the Hosts page, perform the following actions:

For this field…	Do this…
Host Type	Select the type of host: • Windows • Linux The Plug-in for MySQL has to be installed on the MySQL database server.
Host name	Enter the communication host name. Enter the fully qualified domain name (FQDN) or the IP address of the host. SnapCenter depends on the proper configuration of the DNS. Therefore, the best practice is to enter the FQDN.
Credentials	Either select the credential name that you created or create new credentials. The credential must have administrative rights on the remote host. For details, see the information about creating credentials. You can view details about the credentials by positioning your cursor over the credential name that you provided. The credentials authentication mode is determined by the host type that you specify in the Add Host wizard.

5. In the Select Plug-ins to Install section, select the plug-ins to install.

   While using the REST API to install Plug-in for MySQL, you must pass the version as 3.0. For example, MySQL:3.0

6. (Optional) Click **More Options**.

For this field…	Do this…
Port	Either retain the default port number or specify the port number. The default port number is 8145. If the SnapCenter Server was installed on a custom port, that port number will be displayed as the default port. If you manually installed the plug-ins and specified a custom port, you must specify the same port. Otherwise, the operation fails.

For this field…	Do this…
Installation Path	The Plug-in for MySQL is installed on the MySQL client host, and this host can be on either a Windows system or a Linux system. • For the SnapCenter Plug-ins Package for Windows, the default path is C:\Program Files\NetApp\SnapCenter. Optionally, you can customize the path. • For the SnapCenter Plug-ins Package for Linux, the default path is /opt/NetApp/snapcenter. Optionally, you can customize the path.
Skip preinstall checks	Select this check box if you already installed the plug-ins manually and you do not want to validate whether the host meets the requirements for installing the plug-in.
Add all hosts in the cluster	Not Applicable.
Use group Managed Service Account (gMSA) to run the plug-in services	Not Applicable.

7. Click **Submit**.

   If you have not selected the Skip prechecks checkbox, the host is validated to verify whether the host meets the requirements for installing the plug-in. The disk space, RAM, PowerShell version, .NET version, location (for Windows plug-ins), and Java version (for Linux plug-ins) are validated against the minimum requirements. If the minimum requirements are not met, appropriate error or warning messages are displayed.

   If the error is related to disk space or RAM, you can update the web.config file located at C:\Program Files\NetApp\SnapCenter WebApp to modify the default values. If the error is related to other parameters, you must fix the issue.

   In an HA setup, if you are updating web.config file, you must update the file on both nodes.

8. If host type is Linux, verify the fingerprint, and then click **Confirm and Submit**.

   In a cluster setup, you should verify the fingerprint of each of the nodes in the cluster.

   Fingerprint verification is mandatory even if the same host was added earlier to SnapCenter and the fingerprint was confirmed.

9. Monitor the installation progress.
   - For Windows plug-in, the install and upgrade logs are located at: *C:\Windows\SnapCenter plugin\Install*<JOBID>_
   - For Linux plug-in, the install logs are located at: */var/opt/snapcenter/logs/SnapCenter_Linux_Host_Plug-in_Install*<JOBID>.log_ and the upgrade logs

are located at: */var/opt/snapcenter/logs/SnapCenter_Linux_Host_Plug-in_Upgrade*<JOBID>.log_

After you finish

If you want to upgrade to SnapCenter 6.0 version, the existing PERL-based Plug-in for MySQL will be uninstalled from the remote plug-in server.

Install SnapCenter Plug-in Packages for Linux or Windows on multiple remote hosts by using cmdlets

You can install the SnapCenter Plug-in Packages for Linux or Windows on multiple hosts simultaneously by using the Install-SmHostPackage PowerShell cmdlet.

Before you begin

You must have logged in to SnapCenter as a domain user with local administrator rights on each host on which you want to install the plug-in package.

Steps

1. Launch PowerShell.
2. On the SnapCenter Server host, establish a session using the Open-SmConnection cmdlet, and then enter your credentials.
3. Install the plug-in on multiple hosts using the Install-SmHostPackage cmdlet and the required parameters.

   The information regarding the parameters that can be used with the cmdlet and their descriptions can be obtained by running *Get-Help command_name*. Alternatively, you can also refer to the SnapCenter Software Cmdlet Reference Guide.

   You can use the -skipprecheck option when you have installed the plug-ins manually and do not want to validate whether the host meets the requirements to install the plug-in.

4. Enter your credentials for remote installation.

Install the SnapCenter Plug-in for MySQL on Linux hosts by using the command-line interface

You should install the SnapCenter Plug-in for MySQL Database by using the SnapCenter user interface (UI). If your environment does not allow remote installation of the plug-in from the SnapCenter UI, you can install the Plug-in for MySQL Database either in console mode or in silent mode by using the command-line interface (CLI).

Before you begin

- You should install the Plug-in for MySQL Database on each of the Linux host where the MySQL instance must be protected.
- The Linux host on which you are installing the SnapCenter Plug-in for MySQL Database must meet the dependent software, database, and operating system requirements.

  The Interoperability Matrix Tool (IMT) contains the latest information about the supported configurations.

  NetApp Interoperability Matrix Tool

- The SnapCenter Plug-in for MySQL Database is part of SnapCenter Plug-ins Package for Linux. Before you install SnapCenter Plug-ins Package for Linux, you should have already installed SnapCenter on a Windows host.

Steps

1. Copy the SnapCenter Plug-ins Package for Linux installation file (snapcenter_linux_host_plugin.bin) from C:\ProgramData\NetApp\SnapCenter\Package Repository to the host where you want to install the Plug-in for MySQL.

   You can access this path from the host where the SnapCenter Server is installed.

2. From the command prompt, navigate to the directory where you copied the installation file.
3. Install the plug-in: `path_to_installation_bin_file/snapcenter_linux_host_plugin.bin -i silent -DPORT=port_number_for_host -DSERVER_IP=server_name_or_ip_address -DSERVER_HTTPS_PORT=port_number_for_server`

   - -DPORT specifies the SMCore HTTPS communication port.
   - -DSERVER_IP specifies the SnapCenter Server IP address.
   - -DSERVER_HTTPS_PORT specifies the SnapCenter Server HTTPS port.
   - -DUSER_INSTALL_DIR specifies the directory where you want to install the SnapCenter Plug-ins Package for Linux.
   - DINSTALL_LOG_NAME specifies the name of the log file.

   ```
   /tmp/sc-plugin-installer/snapcenter_linux_host_plugin.bin -i silent
   -DPORT=8145 -DSERVER_IP=scserver.domain.com -DSERVER_HTTPS_PORT=8146
   -DUSER_INSTALL_DIR=/opt
   -DINSTALL_LOG_NAME=SnapCenter_Linux_Host_Plugin_Install_2.log
   -DCHOSEN_FEATURE_LIST=CUSTOM
   ```

4. Edit the /<installation directory>/NetApp/snapcenter/scc/etc/SC_SMS_Services.properties file, and then add the PLUGINS_ENABLED = MySQL:3.0 parameter.
5. Add the host to the SnapCenter Server using the Add-Smhost cmdlet and the required parameters.

   The information regarding the parameters that can be used with the command and their descriptions can be obtained by running *Get-Help command_name*. Alternatively, you can also refer to the SnapCenter Software Cmdlet Reference Guide.

Monitor the status of installing Plug-in for MySQL

You can monitor the progress of SnapCenter plug-in package installation by using the Jobs page. You might want to check the progress of installation to determine when it is complete or if there is an issue.

About this task

The following icons appear on the Jobs page and indicate the state of the operation:

- In progress
- Completed successfully
- Failed
-

Completed with warnings or could not start due to warnings

- Queued

Steps

1. In the left navigation pane, click **Monitor**.
2. In the **Monitor** page, click **Jobs**.
3. In the **Jobs** page, to filter the list so that only plug-in installation operations are listed, do the following:
   a. Click **Filter**.
   b. Optional: Specify the start and end date.
   c. From the Type drop-down menu, select **Plug-in installation**.
   d. From the Status drop-down menu, select the installation status.
   e. Click **Apply**.
4. Select the installation job and click **Details** to view the job details.
5. In the **Job Details** page, click **View logs**.

Configure CA Certificate

Generate CA Certificate CSR file

You can generate a Certificate Signing Request (CSR) and import the certificate that can be obtained from a Certificate Authority (CA) using the generated CSR. The certificate will have a private key associated with it.

CSR is a block of encoded text that is given to an authorized certificate vendor to procure the signed CA certificate.

CA Certificate RSA key length must be minimum 3072 bits.

For information to generate a CSR, see How to generate CA Certificate CSR file.

If you own the CA certificate for your domain (*.domain.company.com) or your system (machine1.domain.company.com), you can skip generating the CA Certificate CSR file. You can deploy the existing CA certificate with SnapCenter.

For cluster configurations, the cluster name (virtual cluster FQDN), and the respective host names should be mentioned in the CA certificate. The certificate can be updated by filling the Subject Alternative Name (SAN) field before procuring the certificate. For a wild card certificate (*.domain.company.com), the certificate will contain all the hostnames of the domain implicitly.

Import CA certificates

You must import the CA certificates to the SnapCenter Server and the Windows host plug-ins using the Microsoft management console (MMC).

Steps

1. Go to the Microsoft management console (MMC), and then click **File** > **Add/Remove Snapin**.

2. In the Add or Remove Snap-ins window, select **Certificates** and then click **Add**.
3. In the Certificates snap-in window, select the **Computer account** option, and then click **Finish**.
4. Click **Console Root** > **Certificates – Local Computer** > **Trusted Root Certification Authorities** > **Certificates**.
5. Right-click on the folder “Trusted Root Certification Authorities”, and then select **All Tasks** > **Import** to start the import wizard.
6. Complete the wizard, as follows:

In this wizard window…	Do the following…
Import Private Key	Select the option **Yes**, import the private key, and then click **Next**.
Import File Format	Make no changes; click **Next**.
Security	Specify the new password to be used for the exported certificate, and then click **Next**.
Completing the Certificate Import Wizard	Review the summary, and then click **Finish** to start the import.

Importing certificate should be bundled with the private key (supported formats are: *.pfx, *.p12, and *.p7b).

7. Repeat Step 5 for the “Personal” folder.

Get the CA certificate thumbprint

A certificate thumbprint is a hexadecimal string that identifies a certificate. A thumbprint is calculated from the content of the certificate using a thumbprint algorithm.

Steps

1. Perform the following on the GUI:
   a. Double-click the certificate.
   b. In the Certificate dialog box, click the **Details** tab.
   c. Scroll through the list of fields and click **Thumbprint**.
   d. Copy the hexadecimal characters from the box.
   e. Remove the spaces between the hexadecimal numbers.

      For example, if the thumbprint is: "a9 09 50 2d d8 2a e4 14 33 e6 f8 38 86 b0 0d 42 77 a3 2a 7b", after removing the spaces, it will be: "a909502dd82ae41433e6f83886b00d4277a32a7b".

2. Perform the following from PowerShell:
   a. Run the following command to list the thumbprint of the installed certificate and identify the recently installed certificate by the subject name.

      Get-ChildItem -Path Cert:\LocalMachine\My

b. Copy the thumbprint.

Configure CA certificate with Windows host plug-in services

You should configure the CA certificate with Windows host plug-in services to activate the installed digital certificate.

Perform the following steps on the SnapCenter Server and all the plug-in hosts where CA certificates are already deployed.

Steps

1. Remove the existing certificate binding with SMCore default port 8145, by running the following command:

   ```
   > netsh http delete sslcert ipport=0.0.0.0:_<SMCore Port>
   ```

   For example:

   ```
   > netsh http delete sslcert ipport=0.0.0.0:8145
   ```

2. Bind the newly installed certificate with the Windows host plug-in services, by running the following commands:

   ```
   > $cert = "_<certificate thumbprint>_"
   > $guid = [guid]::NewGuid().ToString("B")
   > netsh http add sslcert ipport=0.0.0.0: _<SMCore Port>_ certhash=$cert
   appid="$guid"
   ```

   For example:

   ```
   > $cert = "a909502dd82ae41433e6f83886b00d4277a32a7b"
   > $guid = [guid]::NewGuid().ToString("B")
   > netsh http add sslcert ipport=0.0.0.0: _<SMCore Port>_ certhash=$cert
   appid="$guid"
   ```

Configure the CA Certificate for the SnapCenter MySQL Plug-ins service on Linux host

You should manage the password of the plug-ins keystore and its certificate, configure the CA certificate, configure root or intermediate certificates to the plug-ins trust-store, and configure CA signed key pair to plug-ins trust-store with SnapCenter plug-ins service to activate the installed digital certificate.

The plug-ins uses the file 'keystore.jks', which is located at */opt/NetApp/snapcenter/scc/etc* both as its trust-store and key-store.

Manage password for plug-in keystore and alias of the CA signed key pair in use

Steps

1. You can retrieve plug-in keystore default password from plug-in agent property file.

   It is the value corresponding to the key 'KEYSTORE_PASS'.

2. Change the keystore password:

   ```
   keytool -storepasswd -keystore keystore.jks
   ```

3. Change the password for all aliases of private key entries in the keystore to the same password used for the keystore:

   ```
   keytool -keypasswd -alias "alias_name_in_cert" -keystore keystore.jks
   ```

   Update the same for the key KEYSTORE_PASS in *agent.properties* file.

4. Restart the service after changing the password.

> Password for plug-in keystore and for all the associated alias password of the private key should be same.

Configure root or intermediate certificates to plug-in trust-store

You should configure the root or intermediate certificates without the private key to plug-in trust-store.

Steps

1. Navigate to the folder containing the plug-in keystore: /opt/NetApp/snapcenter/scc/etc.
2. Locate the file 'keystore.jks'.
3. List the added certificates in the keystore:

   ```
   keytool -list -v -keystore keystore.jks
   ```

4. Add a root or intermediate certificate:

   ```
   keytool -import -trustcacerts -alias myRootCA -file
   /root/USERTrustRSA_Root.cer -keystore keystore.jks
   ```

5. Restart the service after configuring the root or intermediate certificates to plug-in trust-store.

> You should add the root CA certificate and then the intermediate CA certificates.

Configure CA signed key pair to plug-in trust-store

You should configure the CA signed key pair to the plug-in trust-store.

Steps

1. Navigate to the folder containing the plug-in keystore /opt/NetApp/snapcenter/scc/etc.
2. Locate the file 'keystore.jks'.
3. List the added certificates in the keystore:

   ```
   keytool -list -v -keystore keystore.jks
   ```

4. Add the CA certificate having both private and public key.

   ```
   keytool -importkeystore -srckeystore /root/snapcenter.ssl.test.netapp.com.pfx
   -srcstoretype pkcs12 -destkeystore keystore.jks -deststoretype JKS
   ```

5. List the added certificates in the keystore.

   ```
   keytool -list -v -keystore keystore.jks
   ```

6. Verify that the keystore contains the alias corresponding to the new CA certificate, which was added to the keystore.
7. Change the added private key password for CA certificate to the keystore password.

   Default plug-in keystore password is the value of the key KEYSTORE_PASS in agent.properties file.

   ```
   keytool -keypasswd -alias "alias_name_in_CA_cert" -keystore keystore.jks
   ```

8. If the alias name in the CA certificate is long and contains space or special characters ("*",","), change the alias name to a simple name:

   ```
   keytool -changealias -alias "long_alias_name" -destalias "simple_alias"
   -keystore keystore.jks
   ```

9. Configure the alias name from CA certificate in agent.properties file.

   Update this value against the key SCC_CERTIFICATE_ALIAS.

10. Restart the service after configuring the CA signed key pair to plug-in trust-store.

Configure certificate revocation list (CRL) for plug-ins

About this task

- SnapCenter Plug-ins will search for the CRL files in a pre-configured directory.
- Default directory for the CRL files for SnapCenter Plug-ins is ' opt/NetApp/snapcenter/scc/etc/crl'.

Steps

1. You can modify and update the default directory in agent.properties file against the key CRL_PATH.

   You can place more than one CRL file in this directory. The incoming certificates will be verified against each CRL.

Configure the CA Certificate for the SnapCenter MySQL Plug-ins service on Windows host

You should manage the password of the plug-ins keystore and its certificate, configure the CA certificate, configure root or intermediate certificates to the plug-ins trust-store, and configure CA signed key pair to plug-ins trust-store with SnapCenter plug-ins service to activate the installed digital certificate.

The plug-ins uses the file *keystore.jks*, which is located at *C:\Program Files\NetApp\SnapCenter\Snapcenter Plug-in Creator\etc* both as its trust-store and key-store.

Manage password for plug-in keystore and alias of the CA signed key pair in use

Steps

1. You can retrieve plug-in keystore default password from plug-in agent property file.

   It is the value corresponding to the key *KEYSTORE_PASS*.

2. Change the keystore password:

   keytool -storepasswd -keystore keystore.jks

   > If the "keytool" command is not recognized on the Windows command prompt, replace the keytool command with its complete path.

   C:\Program Files\Java\<jdk_version>\bin\keytool.exe" -storepasswd -keystore keystore.jks

3. Change the password for all aliases of private key entries in the keystore to the same password used for the keystore:

   keytool -keypasswd -alias "alias_name_in_cert" -keystore keystore.jks

   Update the same for the key KEYSTORE_PASS in *agent.properties* file.

4. Restart the service after changing the password.

   > Password for plug-in keystore and for all the associated alias password of the private key should be same.

Configure root or intermediate certificates to plug-in trust-store

You should configure the root or intermediate certificates without the private key to plug-in trust-store.

Steps

1. Navigate to the folder containing the plug-in keystore *C:\Program Files\NetApp\SnapCenter\Snapcenter Plug-in Creator\etc*
2. Locate the file 'keystore.jks'.
3. List the added certificates in the keystore:

   keytool -list -v -keystore keystore.jks

4. Add a root or intermediate certificate:

keytool -import -trustcacerts -alias myRootCA -file /root/USERTrustRSA_Root.cer -keystore keystore.jks

5. Restart the service after configuring the root or intermediate certificates to plug-in trust-store.

> You should add the root CA certificate and then the intermediate CA certificates.

Configure CA signed key pair to plug-in trust-store

You should configure the CA signed key pair to the plug-in trust-store.

Steps

1. Navigate to the folder containing the plug-in keystore *C:\Program Files\NetApp\SnapCenter\Snapcenter Plug-in Creator\etc*
2. Locate the file *keystore.jks*.
3. List the added certificates in the keystore:

   keytool -list -v -keystore keystore.jks

4. Add the CA certificate having both private and public key.

   keytool -importkeystore -srckeystore /root/snapcenter.ssl.test.netapp.com.pfx -srcstoretype pkcs12 -destkeystore keystore.jks -deststoretype JKS

5. List the added certificates in the keystore.

   keytool -list -v -keystore keystore.jks

6. Verify that the keystore contains the alias corresponding to the new CA certificate, which was added to the keystore.
7. Change the added private key password for CA certificate to the keystore password.

   Default plug-in keystore password is the value of the key KEYSTORE_PASS in agent.properties file.

   keytool -keypasswd -alias "alias_name_in_CA_cert" -keystore keystore.jks

8. Configure the alias name from CA certificate in *agent.properties* file.

   Update this value against the key SCC_CERTIFICATE_ALIAS.

9. Restart the service after configuring the CA signed key pair to plug-in trust-store.

Configure certificate revocation list (CRL) for SnapCenter plug-ins

About this task

- To download the latest CRL file for the related CA certificate see How to update certificate revocation list file in SnapCenter CA Certificate.
- SnapCenter plug-ins will search for the CRL files in a pre-configured directory.
- Default directory for the CRL files for SnapCenter plug-ins is *'C:\Program Files\NetApp\SnapCenter\Snapcenter Plug-in Creator\ etc\crl'*.

Steps

1. You can modify and update the default directory in *agent.properties* file against the key CRL_PATH.

2. You can place more than one CRL file in this directory.

   The incoming certificates will be verified against each CRL.

Enable CA Certificates for plug-ins

You should configure the CA certificates and deploy the CA certificates in the SnapCenter Server and the corresponding plug-in hosts. You should enable the CA certificate validation for the plug-ins.

Before you begin

- You can enable or disable the CA certificates using the run *Set-SmCertificateSettings* cmdlet.
- You can display the certificate status for the plug-ins using the *Get-SmCertificateSettings*.

The information regarding the parameters that can be used with the cmdlet and their descriptions can be obtained by running *Get-Help command_name*. Alternatively, you can also refer to the SnapCenter Software Cmdlet Reference Guide.

Steps

1. In the left navigation pane, click **Hosts**.
2. In the Hosts page, click **Managed Hosts**.
3. Select single or multiple plug-in hosts.
4. Click **More options**.
5. Select **Enable Certificate Validation**.

After you finish

The Managed Hosts tab host displays a padlock and the color of the padlock indicates the status of the connection between SnapCenter Server and the plug-in host.

- indicates that the CA certificate is neither enabled nor assigned to the plug-in host.
- indicates that the CA certificate is successfully validated.
- indicates that the CA certificate could not be validated.
- indicates that the connection information could not be retrieved.

> When the status is yellow or green, the data protection operations completes successfully.

Prepare for data protection

Prerequisites for using the SnapCenter Plug-in for MySQL

Before you use SnapCenter Plug-in for MySQL, the SnapCenter administrator must install and configure the SnapCenter Server and perform the prerequisite tasks.

- Install and configure SnapCenter Server.
- Log in to SnapCenter Server.
- Configure the SnapCenter environment by adding storage system connections and creating credentials, if

applicable.

- Install Java 11 on your Linux or Windows host.

  You must set the Java path in the environmental path variable of the host machine.

- Set up SnapMirror and SnapVault, if you want backup replication.

How resources, resource groups, and policies are used for protecting MySQL

Before you use SnapCenter, it is helpful to understand basic concepts related to the backup, clone, and restore operations you want to perform. You interact with resources, resource groups, and policies for different operations.

- Resources are typically MySQL instances that you back up or clone with SnapCenter.
- A SnapCenter resource group, is a collection of resources on a host.

  When you perform an operation on a resource group, you perform that operation on the resources defined in the resource group according to the schedule you specify for the resource group.

  You can back up on demand a single resource or a resource group. You also can perform scheduled backups for single resources and resource groups.

- The policies specify the backup frequency, replication, scripts, and other characteristics of data protection operations.

  When you create a resource group, you select one or more policies for that group. You can also select a policy when you perform a backup on demand for a single resource.

Think of a resource group as defining what you want to protect and when you want to protect it in terms of day and time. Think of a policy as defining how you want to protect it. If you are backing up all databases, for example, you might create a resource group that includes all of the databases in the host. You could then attach two policies to the resource group: a daily policy and an hourly policy. When you create the resource group and attach the policies, you might configure the resource group to perform a full backup daily.

Back up MySQL resources

Back up MySQL resources

You can either create a backup of a resource (database) or resource group. The backup workflow includes planning, identifying the databases for backup, managing backup policies, creating resource groups and attaching policies, creating backups, and monitoring the operations.

The following workflow shows the sequence in which you must perform the backup operation:

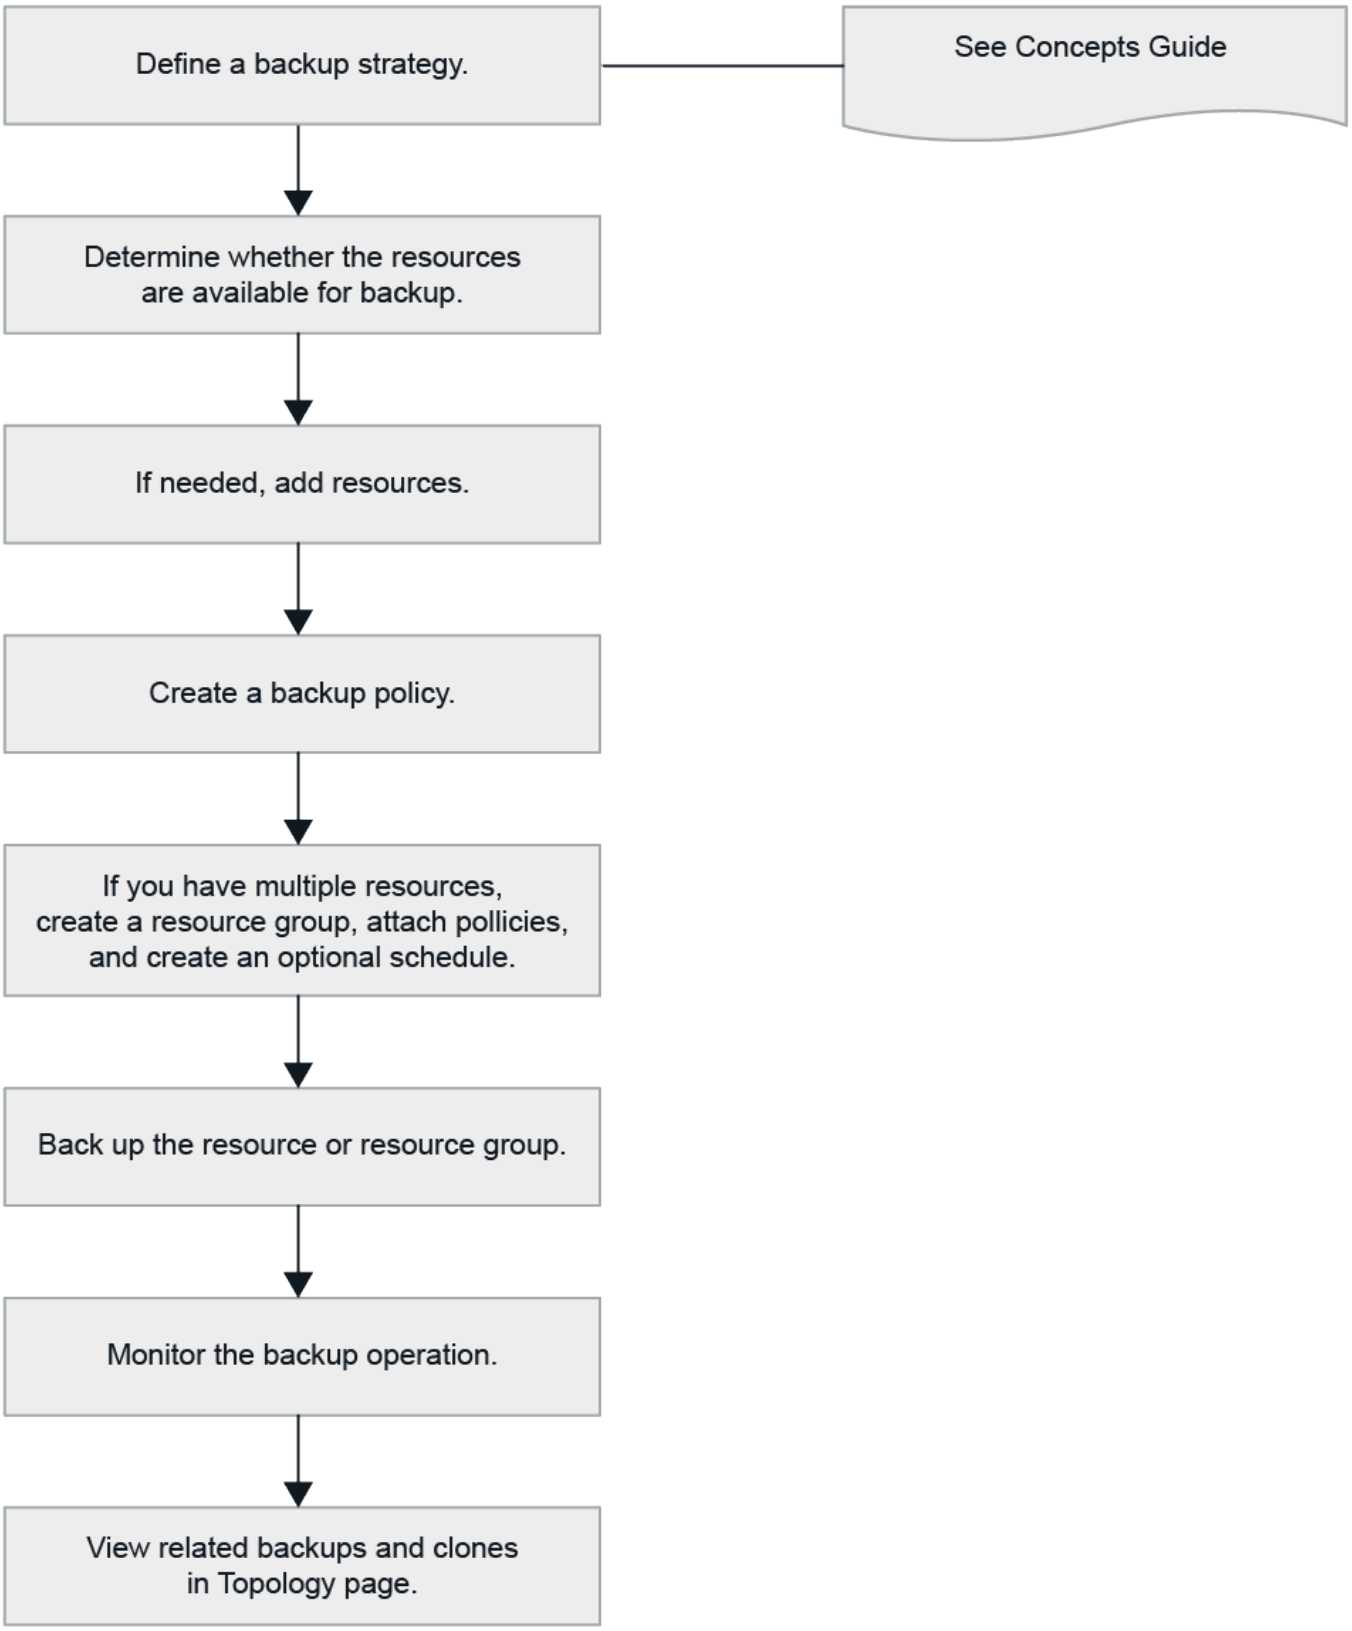

You can also use PowerShell cmdlets manually or in scripts to perform backup, restore, and clone operations. The SnapCenter cmdlet help and the cmdlet reference information contain more information about PowerShell cmdlets.
SnapCenter Software Cmdlet Reference Guide.

Discover the databases automatically

Resources are MySQL databases on the Linux host that are managed by SnapCenter. You can add the resources to resource groups to perform data protection operations after you discover the MySQL databases that are available.

Before you begin

- You must have already completed tasks such as installing the SnapCenter Server, adding hosts, and setting up the storage system connections.
- SnapCenter Plug-in for MySQL does not support automatic discovery of the resources residing on RDM/VMDK virtual environments. You must provide the storage information for virtual environments while adding the databases manually.

About this task

- After installing the plug-in, all the databases on that Linux host are automatically discovered and displayed on the Resources page.
- Only databases are auto-discovered.

The automatically discovered resources cannot be modified or deleted.

Steps

1. In the left navigation pane, click **Resources**, and then select the Plug-in for MySQL from the list.
2. In the Resources page select the resource type from the View list.
3. (Optional) Click , and then select the host name.

   You can then click to close the filter pane.

4. Click **Refresh Resources** to discover the resources available on the host.

   The resources are displayed along with information such as resource type, host name, associated resource groups, backup type, policies and overall status.

   - If the database is on a NetApp storage and not protected, then Not protected is displayed in the Overall Status column.
   - If the database is on a NetApp storage system and protected, and if there is no backup operation performed, then Backup not run is displayed in the Overall Status column. The status will otherwise change to Backup failed or Backup succeeded based on the last backup status.

You must refresh the resources if the instances are renamed outside of SnapCenter.

Add resources manually to the plug-in host

Automatic discovery is not supported on Windows host. You must add MySQL instances and database resources manually.

Before you begin

- You must have completed tasks such as installing the SnapCenter Server, adding hosts, and setting up storage system connections.

Steps

1. In the left navigation pane, select the SnapCenter Plug-in for MySQL from the drop-down list, and then click **Resources**.
2. In the Resources page, click **Add MySQL resources**.
3. In the Provide Resource Details page, perform the following actions:

For this field…	Do this…
Name	Specify the database name.
Host Name	Enter the host name.
Type	Select instance.
Instance	Not applicable.
Credentials	Select the credentials or add information for the credential. This is optional.

4. In the Provide Storage Footprint page, select a storage type and choose one or more volumes, LUNs, and qtrees, and then click **Save**.

   Optional: You can click the + icon to add more volumes, LUNs, and qtrees from other storage systems.

5. Optional: In the Resource Settings page, enter custom key-value pairs for MySQL plug-in.
6. Review the summary, and then click **Finish**.

   The databases are displayed along with information such as the host name, associated resource groups and policies, and overall status

If you want to provide users access to resources, you must assign the resources to the users. This enables users to perform the actions for which they have permissions on the assets that are assigned to them.

Add a user or group and assign role and assets

After adding the databases, you can modify the MySQL database details.

Create backup policies for MySQL

Before you use SnapCenter to back up MySQL resources, you must create a backup policy for the resource or resource group that you want to back up. A backup policy is a set of rules that governs how you manage, schedule, and retain backups.

Before you begin

- You must have defined your backup strategy.

  For details, see the information about defining a data protection strategy for MySQL databases.

- You must have prepared for data protection by completing tasks such as installing SnapCenter, adding hosts, setting up storage system connections, and adding resources.
- The SnapCenter administrator must have assigned the SVMs for both the source and destination volumes to you if you are replicating snapshots to a mirror or vault.

  Additionally, you can specify replication, script, and application settings in the policy. These options saves time when you want to reuse the policy for another resource group.

About this task

- SnapLock
  - If 'Retain the backup copies for a specific number of days' option is selected, then the SnapLock retention period must be lesser than or equal to the mentioned retention days.
  - Specifying a snapshot locking period prevents deletion of the snapshots until the retention period expires. This could lead to retaining a larger number of snapshots than the count specified in the policy.
  - For ONTAP 9.12.1 and below version, the clones created from the SnapLock Vault snapshots as part of restore will inherit the SnapLock Vault expiry time. Storage admin should manually cleanup the clones post the SnapLock expiry time.

Steps

1. In the left navigation pane, click **Settings**.
2. In the Settings page, click **Policies**.
3. Click **New**.
4. In the Name page, enter the policy name and details.
5. In the Policy type page, perform the following:
   a. Select storage type.
   b. In the Custom backup settings section, provide any specific backup settings that have to be passed to the plug-in in key-value format.

      You can provide multiple key-values to be passed to the plug-in.
6. In the Snapshot and Replication page, perform the following steps:
   a. Specify the schedule type by selecting **On demand**, **Hourly**, **Daily**, **Weekly**, or **Monthly**.

      You can specify the schedule (start date, end date, and frequency) for the backup operation while creating a resource group. This enables you to create resource groups that share the same policy and backup frequency, but enables you to assign different backup schedules to each policy.

      If you have scheduled for 2:00 a.m., the schedule will not be triggered during daylight saving time (DST).

   b. In the Snapshot settings section, specify the retention settings for the backup type and the schedule type selected in the **Backup Type** page:

If you want to…	Then…
Keep a certain number of Snapshots	Select **Copies to keep**, and then specify the number of Snapshots that you want to keep. If the number of Snapshots exceeds the specified number, the Snapshots are deleted with the oldest copies deleted first. ⓘ You must set the retention count to 2 or higher if you plan to enable SnapVault replication. If you set the retention count to 1, the retention operation might fail because the first Snapshot is the reference Snapshot for the SnapVault relationship until a newer Snapshot is replicated to the target. ⓘ The maximum retention value is 1018. Backups will fail if retention is set to a value higher than what the underlying ONTAP version supports.
Keep the Snapshots for a certain number of days	Select **Retain copies for**, and then specify the number of days for which you want to keep the Snapshots before deleting them.
Snapshot copy locking period	Select **Snapshot copy locking period**, and specify days, months, or years. SnapLock retention period should be less than 100 years.

c. Select a Policy label.

ⓘ You can assign SnapMirror labels to primary snapshots for remote replication, allowing the primary snapshots to offload the snapshot replication operation from SnapCenter to ONTAP secondary systems. This can be done without enabling SnapMirror or SnapVault option in the policy page.

7. In the Select secondary replication options section, select one or both of the following secondary replication options:

For this field…	Do this…
Update SnapMirror after creating a local Snapshot copy	Select this field to create mirror copies of the backup sets on another volume (SnapMirror replication). If the protection relationship in ONTAP is of type Mirror and Vault and if you select only this option, Snapshot created on the primary will not be transferred to the destination, but will be listed in the destination. If this Snapshot is selected from the destination to perform a restore operation, then the following error message is displayed: Secondary Location is not available for the selected vaulted/mirrored backup. During secondary replication, the SnapLock expiry time loads the primary SnapLock expiry time. Clicking the **Refresh** button in the Topology page refreshes the secondary and primary SnapLock expiry time that are retrieved from ONTAP. See View MySQL resource related backups and clones in the Topology page.
Update SnapVault after creating a local Snapshot copy	Select this option to perform disk-to-disk backup replication (SnapVault backups). During secondary replication, the SnapLock expiry time loads the primary SnapLock expiry time. Clicking the **Refresh** button in the Topology page refreshes the secondary and primary SnapLock expiry time that are retrieved from ONTAP. When SnapLock is configured only on the secondary from ONTAP known as SnapLock Vault, clicking the **Refresh** button in the Topology page refreshes the locking period on the secondary that is retrieved from ONTAP. For more information on SnapLock Vault see Commit Snapshots to WORM on a vault destination See View MySQL resource related backups and clones in the Topology page.
Error retry count	Enter the maximum number of replication attempts that can be allowed before the operation stops.

You should configure SnapMirror retention policy in ONTAP for the secondary storage to avoid reaching the maximum limit of Snapshots on the secondary storage.

8. Review the summary, and then click **Finish**.

Create resource groups and attach policies

A resource group is the container to which you must add resources that you want to back up and protect. A resource group enables you to back up all the data that is associated with a given application simultaneously. A resource group is required for any data protection job. You must also attach one or more policies to the resource group to define the type of data protection job that you want to perform.

About this task

- For ONTAP 9.12.1 and below version, the clones created from the SnapLock Vault snapshots as part of restore will inherit the SnapLock Vault expiry time. Storage admin should manually cleanup the clones post the SnapLock expiry time.

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page, click **New Resource Group**.
3. In the Name page, perform the following actions:

For this field…	Do this…
Name	Enter a name for the resource group. The resource group name should not exceed 250 characters.
Tags	Enter one or more labels that will help you later search for the resource group. For example, if you add HR as a tag to multiple resource groups, you can later find all resource groups associated with the HR tag.
Use custom name format for snapshot copy	Select this check box, and enter a custom name format that you want to use for the snapshot name. For example, customtext_resource group_policy_hostname or resource group_hostname. By default, a timestamp is appended to the snapshot name.

4. In the Resources page, select a host name from the **Host** drop-down list and resource type from the **Resource Type** drop-down list.

   This helps to filter information on the screen.

5. Select the resources from the **Available Resources** section, and then click the right arrow to move them to the **Selected Resources** section.

6. In the Application Settings page, do the following:
   a. Click the **Backups** arrow to set additional backup options:

      Enable consistency group backup and perform the following tasks:

For this field…	Do this…
Afford time to wait for Consistency Group snapshot operation to complete	Select **Urgent**, **Medium**, or **Relaxed** to specify the wait time for snapshot operation to complete. Urgent = 5 seconds, Medium = 7 seconds, and Relaxed = 20 seconds.
Disable WAFL Sync	Select this to avoid forcing a WAFL consistency point.

   b. Click the **Scripts** arrow and enter the pre and post commands for quiesce, snapshot, and unquiesce operations. You can also enter the pre commands to be executed before exiting in the event of a failure.
   c. Click the **Custom Configurations** arrow and enter the custom key-value pairs required for all data protection operations using this resource.

Parameter	Setting	Description
ARCHIVE_LOG_ENABLE	(Y/N)	Enables the archive log management to delete the archive logs.

Parameter	Setting	Description
ARCHIVE_LOG_RETENTION	number_of_days	Specifies the number of days the archive logs are retained. This setting must be equal to or greater than NTAP_SNAPSHOT_ RETENTIONS.
ARCHIVE_LOG_DIR	change_info_directory/logs	Specifies the path to the directory that contains the archive logs.
ARCHIVE_LOG_EXT	file_extension	Specifies the archive log file extension length. For example, if the archive log is log_backup_0_0_0_0.16151855 1942 9 and if the file_extension value is 5, then the extension of the log will retain 5 digits, which is 16151.
ARCHIVE_LOG_RECURSIVE_ SE ARCH	(Y/N)	Enables the management of archive logs within subdirectories. You should use this parameter if the archive logs are located under subdirectories.

The custom key-value pairs are supported for MySQL Linux plug-in systems and not supported for MySQL database registered as a centralized windows plug-in.

d. Click the **Snapshot Copy Tool** arrow to select the tool to create snapshots:

If you want…	Then…
SnapCenter to use the plug-in for Windows and put the file system into a consistent state before creating a snapshot. For Linux resources, this option is not applicable.	Select **SnapCenter with File System Consistency**.
SnapCenter to create a storage level snapshot	Select **SnapCenter without File System Consistency**.

If you want…	Then…
To enter the command to be executed on the host to create snapshot copies.	Select **Other**, and then enter the command to be executed on the host to create a snapshot.

7. In the Policies page, perform the following steps:
   a. Select one or more policies from the drop-down list.

      You can also create a policy by clicking + .

      The policies are listed in the Configure schedules for selected policies section.

   b. In the Configure Schedules column, click + for the policy you want to configure.
   c. In the Add schedules for policy *policy_name* dialog box, configure the schedule, and then click **OK**.

      Where, policy_name is the name of the policy that you have selected.

      The configured schedules are listed in the **Applied Schedules** column.

      Third party backup schedules are not supported when they overlap with SnapCenter backup schedules.
8. In the Notification page, from the **Email preference** drop-down list, select the scenarios in which you want to send the emails.

   You must also specify the sender and receiver email addresses, and the subject of the email. The SMTP server must be configured in **Settings** > **Global Settings**.
9. Review the summary, and then click **Finish**.

Create resource groups and enable secondary protection for MySQL resources on ASA r2 systems

You should create the resource group to add the resources that are on ASA r2 systems. You can also provision the secondary protection while creating the resource group.

Before you begin

- You should ensure that you are not adding both ONTAP 9.x resources and ASA r2 resources to the same resource group.
- You should ensure that you do not have a database with both ONTAP 9.x resources and ASA r2 resources.

About this task

- The secondary protection is available only if the logged-in user is assigned to the role that has the **SecondaryProtection** capability enabled.
- If you enabled secondary protection, the resource group is put into maintenance mode while creating the primary and secondary consistency groups. After the primary and secondary consistency groups are created, the resource group is put out of maintenance mode.
- SnapCenter does not support secondary protection for a clone resource.

Steps

1. In the left navigation pane, select **Resources**, and the appropriate plug-in from the list.
2. In the Resources page, click **New Resource Group**.
3. In the Name page, perform the following actions:
   a. Enter a name for the resource group in the Name field.

      The resource group name should not exceed 250 characters.

   b. Enter one or more labels in the Tag field to help you search for the resource group later.

      For example, if you add HR as a tag to multiple resource groups, you can later find all resource groups associated with the HR tag.

   c. Select this check box, and enter a custom name format that you want to use for the Snapshot name.

      For example, customtext_resource group_policy_hostname or resource group_hostname. By default, a timestamp is appended to the Snapshot name.

   d. Specify the destinations of the archive log files that you do not want to back up.

      You should use the exact same destination as it was set in the application including prefix, if needed.

4. In the Resources page, select the database host name from the **Host** drop-down list.

   The resources are listed in the Available Resources section only if the resource is discovered successfully. If you have recently added resources, they will appear on the list of available resources only after you refresh your resource list.

5. Select the ASA r2 resources from the Available Resources section and move them to the Selected Resources section.
6. In the Application Settings page, select the backup option.
7. In the Policies page, perform the following steps:
   a. Select one or more policies from the drop-down list.

      You can also create a policy by clicking + .

      In the Configure schedules for selected policies section, the selected policies are listed.

   b. Click + in the Configure Schedules column for the policy for which you want to configure a schedule.
   c. In the Add schedules for policy *policy_name* window, configure the schedule, and then click **OK**.

      Where, *policy_name* is the name of the policy that you have selected.

      The configured schedules are listed in the Applied Schedules column.

Third party backup schedules are not supported when they overlap with SnapCenter backup schedules.

8. If the secondary protection is enabled for the policy that you have selected, then Secondary Protection page is displayed and you need to perform the following steps:
   a. Select the type of the replication policy.

      Synchronous replication policy is not supported.

   b. Specify the consistency group suffix that you want to use.
   c. From the Destination Cluster and Destination SVM drop-downs select the peered cluster and SVM that you want to use.

      The cluster and SVM peering is not supported by SnapCenter. You should use System Manager or ONTAP CLIs to perform cluster and SVM peering.

If the resources are already protected outside of SnapCenter, those resources will be displayed in the Secondary Protected Resources section.

1. On the Verification page, perform the following steps:
   a. Click **Load locators** to load the SnapMirror or SnapVault volumes to perform verification on secondary storage.
   b. Click in the Configure Schedules column to configure the verification schedule for all the schedule types of the policy.
   c. In the Add Verification Schedules policy_name dialog box, perform the following actions:

If you want to…	Do this…
Run verification after backup	Select **Run verification after backup**.
Schedule a verification	Select **Run scheduled verification** and then select the schedule type from the drop-down list.

   d. Select **Verify on secondary location** to verify your backups on secondary storage system.
   e. Click **OK**.

      The configured verification schedules are listed in the Applied Schedules column.

2. In the Notification page, from the **Email preference** drop-down list, select the scenarios in which you want to send the emails.

   You must also specify the sender and receiver email addresses, and the subject of the email. If you want to attach the report of the operation performed on the resource group, select **Attach Job Report**.

   For email notification, you must have specified the SMTP server details using the either the GUI or the PowerShell command Set-SmSmtpServer.

3. Review the summary, and then click **Finish**.

Create a storage system connection and a credential using PowerShell cmdlets for MySQL

You must create a storage virtual machine (SVM) connection and a credential before using PowerShell cmdlets to back up, restore, or clone MySQL databases.

Before you begin

- You should have prepared the PowerShell environment to execute the PowerShell cmdlets.
- You should have the required permissions in the Infrastructure Admin role to create storage connections.
- You should ensure that the plug-in installations are not in progress.

  Host plug-in installations must not be in progress while adding a storage system connection because the host cache might not be updated and databases status might be displayed in the SnapCenter GUI as “Not available for backup” or “Not on NetApp storage”.

- Storage system names should be unique.

  SnapCenter does not support multiple storage systems with the same name on different clusters. Each storage system that is supported by SnapCenter should have a unique name and a unique data LIF IP address.

Steps

1. Initiate a PowerShell Core connection session by using the Open-SmConnection cmdlet.

```
PS C:\> Open-SmConnection
```

2. Create a new connection to the storage system by using the Add-SmStorageConnection cmdlet.

```
PS C:\> Add-SmStorageConnection -Storage test_vs1 -Protocol Https
-Timeout 60
```

3. Create a new credential by using the Add-SmCredential cmdlet.

   This example shows how to create a new credential named FinanceAdmin with Windows credentials:

```
PS C:> Add-SmCredential -Name FinanceAdmin -AuthMode Windows
-Credential sddev\administrator
```

4. Add the MySQL communication host to SnapCenter Server.

```
PS C:> Add-SmHost -HostName 10.232.204.61 -OSType Windows -RunAsName
FinanceAdmin -PluginCode mysql
```

5. Install the package and the SnapCenter Plug-in for MySQL on the host.

For Linux:

```
PS C:> Install-SmHostPackage -HostNames 10.232.204.61 -ApplicationCode
mysql
```

For Windows:

```
Install-SmHostPackage -HostNames 10.232.204.61 -ApplicationCode mysql
-FilesystemCode scw -RunAsName FinanceAdmin
```

6. Set the path to the SQLLIB.

   For Windows, MySQL plug-in will use the default path for SQLLIB folder: “C:\Program Files\IBM\SQLLIB\BIN”

   If you want to override the default path, use the following command.

```
PS C:> Set-SmConfigSettings -Plugin -HostName 10.232.204.61 -PluginCode
MySQL -configSettings @{“MySQL_SQLLIB_CMD” =
“<custom_path>\IBM\SQLLIB\BIN”}
```

The information regarding the parameters that can be used with the cmdlet and their descriptions can be obtained by running *Get-Help command_name*. Alternatively, you can also refer to the SnapCenter Software Cmdlet Reference Guide.

Back up MySQL

If a resource is not yet part of any resource group, you can back up the resource from the Resources page.

Before you begin

- You must have created a backup policy.
- If you want to back up a resource that has a SnapMirror relationship with a secondary storage, the ONTAP role assigned to the storage user should include the “snapmirror all” privilege. However, if you are using the “vsadmin” role, then the “snapmirror all” privilege is not required.
- For Snapshot copy-based backup operation, ensure that all the tenant databases are valid and active.
- For pre and post commands for quiesce, Snapshot, and unquiesce operations, you should check if the commands exist in the command list available on the plug-in host from the following paths:
  - Default location on the Windows host: *C:\Program Files\NetApp\SnapCenter\Snapcenter Plug-in Creator\etc\allowed_commands.config*
  - Default location on the Linux host: */opt/NetApp/snapcenter/scc/etc/allowed_commands.config*

If the commands do not exist in the command list, then the operation will fail.

SnapCenter UI

Steps

1. In the left navigation pane, select **Resources**, and then select the appropriate plug-in from the list.
2. In the Resource page, filter resources from the **View** drop-down list based on resource type.

   Select , and then select the host name and the resource type to filter the resources. You can then select to close the filter pane.

3. Select the resource that you want to back up.
4. In the Resource page, select **Use custom name format for Snapshot copy**, and then enter a custom name format that you want to use for the Snapshot name.

   For example, *customtext_policy_hostname* or *resource_hostname*. By default, a timestamp is appended to the Snapshot name.

5. In the Application Settings page, do the following:
   - Select the **Backups** arrow to set additional backup options:

     Enable consistency group backup, if needed, and perform the following tasks:

For this field…	Do this…
Afford time to wait for "Consistency Group Snapshot" operation to complete	Select **Urgent**, or **Medium**, or **Relaxed** to specify the wait time for Snapshot operation to finish. Urgent = 5 seconds, Medium = 7 seconds, and Relaxed = 20 seconds.
Disable WAFL Sync	Select this to avoid forcing a WAFL consistency point.

   - Select the **Scripts** arrow to run pre and post commands for quiesce, Snapshot, and unquiesce operations.

     You can also run pre commands before exiting the backup operation. Prescripts and postscripts are run in the SnapCenter Server.

   - Select the **Custom Configurations**arrow, and then enter the custom value pairs required for all jobs using this resource.
   - Select the **Snapshot Copy Tool** arrow to select the tool to create Snapshots:

If you want…	Then…
SnapCenter to create a storage-level Snapshot	Select **SnapCenter without File System Consistency**.
SnapCenter to use the plug-in for Windows to put the file system into a consistent state and then create a Snapshot	Select **SnapCenter with File System Consistency**.

If you want…	Then…
To enter the command to create a Snapshot	Select **Other**, and then enter the command to create a Snapshot.

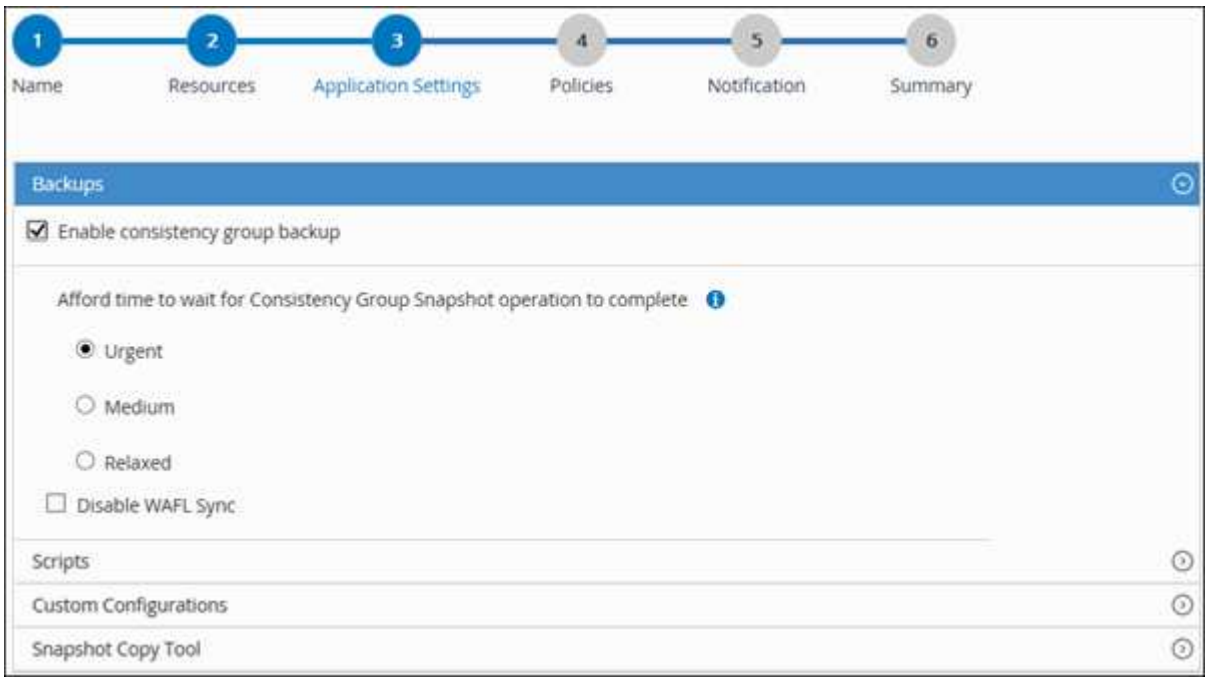

6. In the Policies page, perform the following steps:
   a. Select one or more policies from the drop-down list.

      You can also create a policy by clicking + .

      In the Configure schedules for selected policies section, the selected policies are listed.

   b. Select + in the Configure Schedules column for the policy for which you want to configure a schedule.
   c. In the Add schedules for policy *policy_name* dialog box, configure the schedule, and then select **OK**.

      policy_name is the name of the policy that you selected.

      The configured schedules are listed in the Applied Schedules column.

7. In the Notification page, from the **Email preference** drop-down list, select the scenarios in which you want to send the emails.

   You must also specify the sender and receiver email addresses, and the subject of the email. SMTP must also be configured in **Settings** > **Global Settings**.

8. Review the summary, and then select **Finish**.

   The resources topology page is displayed.

9. Select **Back up Now**.
10. In the Backup page, perform the following steps:
    a. If you applied multiple policies to the resource, from the **Policy** drop-down list, select the policy that you want to use for backup.

       If the policy selected for the on-demand backup is associated with a backup schedule, the on-demand backups will be retained based on the retention settings specified for the schedule type.

    b. Select **Backup**.
11. Monitor the operation progress by clicking **Monitor** > **Jobs**.
    - In MetroCluster configurations, SnapCenter might not be able to detect a protection relationship after a failover.

      For information, see: Unable to detect SnapMirror or SnapVault relationship after MetroCluster failover

    - If you are backing up application data on VMDKs and the Java heap size for the SnapCenter Plug-in for VMware vSphere is not large enough, the backup might fail.

      To increase the Java heap size, locate the script file */opt/netapp/init_scripts/scvservice*. In that script, the *do_start method* command starts the SnapCenter VMware plug-in service. Update that command to the following: *Java -jar -Xmx8192M -Xms4096M*

PowerShell cmdlets

Steps

1. Initiate a connection session with the SnapCenter Server for a specified user by using the Open-SmConnection cmdlet.

```
Open-SmConnection  -SMSbaseurl
https:\\snapctr.demo.netapp.com:8146\
```

   The username and password prompt is displayed.

2. Add manual resources by using the Add-SmResources cmdlet.

   This example shows how to add a MySQL instance:

```
PS C:\> Add-SmResource -HostName 10.32.212.13 -PluginCode MySQL
-ResourceType Instance -ResourceName mysqlinst1 -StorageFootPrint
(@{"VolumeName"="winmysql01_data01";"LUNName"="winmysql01_data01";"S
torageSystem"="scsnfssvm"}) -MountPoints "D:\"
```

3. Create a backup policy by using the Add-SmPolicy cmdlet.
4. Protect the resource or add a new resource group to SnapCenter by using the Add-SmResourceGroup cmdlet.
5. Initiate a new backup job by using the New-SmBackup cmdlet.

This example shows how to backup a resource group:

```
C:\PS> New-SmBackup -Resources
@{"Host"="scs000211748.gdl.englab.netapp.com";"Uid"="mysqld_3306";"P
luginName"="MySQL"} -Policy "MySQL_snapshotbased"
```

This example backs up a protected resource:

```
C:\PS> New-SMBackup -Resources
@{"Host"="10.232.204.42";"Uid"="MDC\SID";"PluginName"="hana"}
-Policy mysql_policy2
```

6. Monitor the job status (running, completed, or failed) by using the Get-smJobSummaryReport cmdlet.

```
PS C:\> Get-smJobSummaryReport -JobID 123
```

7. Monitor the backup job details like backup ID, backup name to perform restore or clone operation by using the Get-SmBackupReport cmdlet.

```
PS C:\> Get-SmBackupReport -JobId 351
Output:
BackedUpObjects           : {DB1}
FailedObjects             : {}
IsScheduled               : False
HasMetadata               : False
SmBackupId                : 269
SmJobId                   : 2361
StartDateTime             : 10/4/2016 11:20:45 PM
EndDateTime               : 10/4/2016 11:21:32 PM
Duration                  : 00:00:46.2536470
CreatedDateTime           : 10/4/2016 11:21:09 PM
Status                    : Completed
ProtectionGroupName       : Verify_ASUP_Message_windows
SmProtectionGroupId       : 211
PolicyName                : test2
SmPolicyId                : 20
BackupName                : Verify_ASUP_Message_windows_scc54_10-04-
2016_23.20.46.2758
VerificationStatus        : NotVerified
VerificationStatuses      :
SmJobError                :
BackupType                : SCC_BACKUP
CatalogingStatus          : NotApplicable
CatalogingStatuses        :
ReportDataCreatedDateTime :
```

The information regarding the parameters that can be used with the cmdlet and their descriptions can be obtained by running *Get-Help command_name*. Alternatively, you can also refer to the SnapCenter Software Cmdlet Reference Guide.

Back up resource groups

A resource group is a collection of resources on a host. A backup operation on the resource group is performed on all resources defined in the resource group.

Before you begin

- You must have created a resource group with a policy attached.
- If you want to back up a resource that has a SnapMirror relationship with a secondary storage, the ONTAP role assigned to the storage user should include the “snapmirror all” privilege. However, if you are using the “vsadmin” role, then the “snapmirror all” privilege is not required.

About this task

You can back up a resource group on demand from the Resources page. If a resource group has a policy attached and a schedule configured, then backups occur automatically according to the schedule.

Steps

1. In the left navigation pane, select **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page, select **Resource Group** from the **View** list.

   You can search the resource group either by entering the resource group name in the search box or by selecting , and then selecting the tag. You can then select to close the filter pane.

3. In the Resource Groups page, select the resource group that you want to back up, and then select **Back up Now**.
4. In the Backup page, perform the following steps:
   a. If you associated multiple policies with the resource group, from the **Policy** drop-down list, select the policy that you want to use for backup.

      If the policy selected for the on-demand backup is associated with a backup schedule, the on-demand backups will be retained based on the retention settings specified for the schedule type.

   b. Select **Backup**.
5. Monitor the operation progress by selecting **Monitor** > **Jobs**.

Monitor MySQL backup operations

You can monitor the progress of different backup operations by using the SnapCenterJobs page. You might want to check the progress to determine when it is complete or if there is an issue.

About this task

The following icons appear on the Jobs page and indicate the corresponding state of the operations:

- In progress
- Completed successfully
- Failed
- Completed with warnings or could not start due to warnings
- Queued
- Canceled

Steps

1. In the left navigation pane, click **Monitor**.
2. In the Monitor page, click **Jobs**.
3. In the Jobs page, perform the following steps:
   a. Click to filter the list so that only backup operations are listed.
   b. Specify the start and end dates.
   c. From the **Type** drop-down list, select **Backup**.
   d. From the **Status** drop-down, select the backup status.
   e. Click **Apply** to view the operations completed successfully.

4. Select a backup job, and then click **Details** to view the job details.

   Though the backup job status displays ✔, when you click on job details you might see that some of the child tasks of the backup operation are still in progress or marked with warning signs.

5. In the Job Details page, click **View logs**.

   The **View logs** button displays the detailed logs for the selected operation.

Monitor data protection operations on MySQL instances in the Activity pane

The Activity pane displays the five most recent operations performed. The Activity pane also displays when the operation was initiated and the status of the operation.

The Activity pane displays information regarding backup, restore, clone, and scheduled backup operations.

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. Click on the Activity pane to view the five most recent operations.

   When you click one of the operations, the operation details are listed in the **Job Details** page.

Cancel backup operations for MySQL

You can cancel backup operations that are queued.

What you will need

- You must be logged in as the SnapCenter Admin or job owner to cancel operations.
- You can cancel a backup operation from either the **Monitor** page or the **Activity** pane.
- You cannot cancel a running backup operation.
- You can use the SnapCenter GUI, PowerShell cmdlets, or CLI commands to cancel the backup operations.
- The **Cancel Job** button is disabled for operations that cannot be canceled.
- If you selected **All members of this role can see and operate on other members objects** in Users\Groups page while creating a role, you can cancel the queued backup operations of other members while using that role.

Steps

1. Perform one of the following actions:

From the…	Action
Monitor page	a. In the left navigation pane, click **Monitor** > **Jobs**. b. Select the operation, and then click **Cancel Job**.

From the…	Action
Activity pane	a. After initiating the backup operation, click on the Activity pane to view the five most recent operations. b. Select the operation. c. In the Job Details page, click **Cancel Job**.

The operation is canceled, and the resource is reverted to the previous state.

View MySQL backups and clones in the Topology page

When you are preparing to back up or clone a resource, you might find it helpful to view a graphical representation of all backups and clones on the primary and secondary storage.

About this task

You can review the following icons in the Manage Copies view to determine whether the backups and clones are available on the primary or secondary storage (Mirror copies or Vault copies).

- displays the number of backups and clones that are available on the primary storage.
- displays the number of backups and clones that are mirrored on the secondary storage using SnapMirror technology.
- displays the number of backups and clones that are replicated on the secondary storage using SnapVault technology.

> The number of backups displayed includes the backups deleted from the secondary storage. For example, if you created 6 backups using a policy to retain only 4 backups, the number of backups displayed is 6.

> Clones of a backup of a version-flexible mirror on a mirror-vault type volume are displayed in the topology view, but the mirror backup count in the topology view does not include the version-flexible backup.

In the Topology page, you can see all of the backups and clones that are available for the selected resource or resource group. You can view the details of those backups and clones, and then select them to perform data protection operations.

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page, either select the resource or resource group from the **View** drop-down list.
3. Select the resource either from the resource details view or from the resource group details view.

If the resource is protected, the topology page of the selected resource is displayed.

4. Review the **Summary card** to see a summary of the number of backups and clones available on the primary and secondary storage.

   The **Summary Card** section displays the total number of Snapshot copy-based backups, and clones.

   Clicking the **Refresh** button starts a query of the storage to display an accurate count.

   If SnapLock enabled backup is taken, then clicking the **Refresh** button refreshes the primary and secondary SnapLock expiry time retrieved from ONTAP. A weekly schedule also refreshes the primary and secondary SnapLock expiry time retrieved from ONTAP.

   When the application resource is spread across multiple volumes, the SnapLock expiry time for the backup will be the longest SnapLock expiry time that is set for a Snapshot in a volume. The longest SnapLock expiry time is retrieved from ONTAP.

   After on demand backup, by clicking the **Refresh** button refreshes the details of backup or clone.

5. In the Manage Copies view, click **Backups** or **Clones** from the primary or secondary storage to see details of a backup or clone.

   The details of the backups and clones are displayed in a table format.

6. Select the backup from the table, and then click the data protection icons to perform restore, clone, and delete operations.

   You cannot rename or delete backups that are on the secondary storage.

7. If you want to delete a clone, select the clone from the table, and then click .
8. If you want to split a clone, select the clone from the table, and then click .

Restore MySQL

Restore workflow

The restore and recovery workflow includes planning, performing the restore operations, and monitoring the operations.

The following workflow shows the sequence in which you must perform the restore operation:

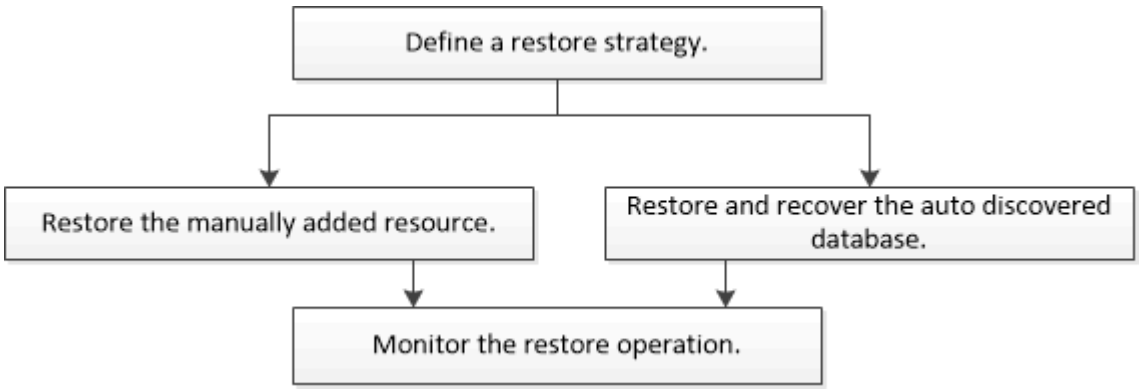

You can also use PowerShell cmdlets manually or in scripts to perform backup, restore, and clone operations. The SnapCenter cmdlet help and the cmdlet reference information contain detailed information about PowerShell cmdlets.

SnapCenter Software Cmdlet Reference Guide.

Restore and recover a manually added resource backup

You can use SnapCenter to restore and recover data from one or more backups.

Before you begin

- You must have backed up the resource or resource groups.
- You must have canceled any backup operation that is currently in progress for the resource or resource group that you want to restore.
- For pre restore, post restore, mount, and unmount commands, you should check if the commands exist in the command list available on the plug-in host from the following paths:
  - Default location on the Windows host: *C:\Program Files\NetApp\SnapCenter\Snapcenter Plug-in Creator\etc\allowed_commands.config*
  - Default location on the Linux host: */opt/NetApp/snapcenter/scc/etc/allowed_commands.config*

If the commands do not exist in the command list, then the operation will fail.

About this task

- For ONTAP 9.12.1 and earlier version, the clones created from the SnapLock Vault Snapshots as part of restore will inherit the SnapLock Vault expiry time. Storage admin should manually cleanup the clones post the SnapLock expiry time.

SnapCenter UI

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page, filter resources from the **View** drop-down list based on resource type.

   The resources are displayed along with type, host, associated resource groups and policies, and status.

   > Although a backup might be for a resource group, when you restore, you must select the individual resources you want to restore.

   If the resource is not protected, “Not protected” is displayed in the Overall Status column. This can mean either that the resource is not protected, or that the resource was backed up by a different user.
3. Select the resource, or select a resource group and then select a resource in that group.

   The resource topology page is displayed.
4. From the Manage Copies view, select **Backups** either from the primary or secondary (mirrored or vaulted) storage systems.
5. In the Primary backup(s) table, select the backup that you want to restore from, and then click .

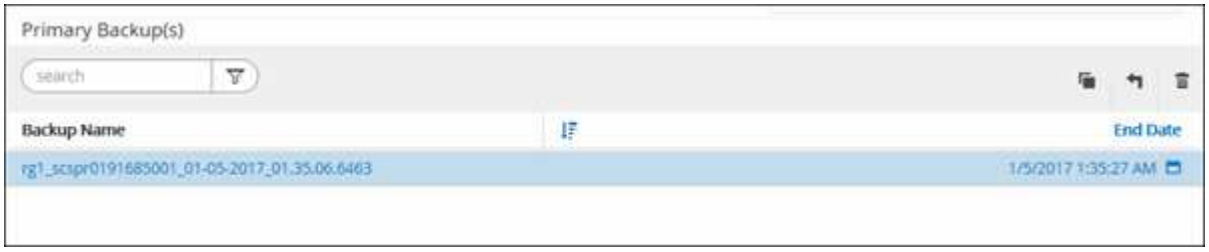

6. In the Restore Scope page, select **Complete Resource**.

   a. If you select **Complete Resource**, all of the configured data volumes of the MySQL database are restored.

      If the resource contains volumes or qtrees, the Snapshots taken after the Snapshot selected for restore on such volumes or qtrees are deleted and cannot be recovered. Also, if any other resource is hosted on same volumes or qtrees, then that resource is also deleted.

      You can select multiple LUNs.

      > If you select **All**, all the files on the volumes, qtrees, or LUNs are restored.
7. In the Pre ops page, enter pre restore and unmount commands to run before performing a restore job.

   Unmount commands are not available for auto discovered resources.
8. In the Post ops page, enter mount and post restore commands to run after performing a restore job.

   Mount commands are not available for auto discovered resources.

9. In the Notification page, from the **Email preference** drop-down list, select the scenarios in which you want to send the emails.

   You must also specify the sender and receiver email addresses and the subject of the email. SMTP must also be configured on the **Settings** > **Global Settings** page.

10. Review the summary, and then click **Finish**.
11. Monitor the operation progress by clicking **Monitor** > **Jobs**.

PowerShell cmdlets

Steps

1. Initiate a connection session with the SnapCenter Server for a specified user by using the Open-SmConnection cmdlet.

```
PS C:\> Open-Smconnection
```

2. Retrieve the information about the one or more backups that you want to restore by using the Get-SmBackup and Get-SmBackupReport cmdlets.

   This example displays information about all available backups:

```
PS C:\> Get-SmBackup

BackupId                      BackupName
BackupTime                    BackupType
--------                      ----------
----------                    ----------
  1               Payroll Dataset_vise-f6_08... 8/4/2015    11:02:32
AM                 Full Backup
  2               Payroll Dataset_vise-f6_08... 8/4/2015    11:23:17
AM
```

   This example displays detailed information about the backup from January 29th 2015 to February 3rd, 2015:

```
PS C:\> Get-SmBackupReport -FromDateTime "1/29/2015" -ToDateTime
"2/3/2015"

   SmBackupId           : 113
   SmJobId              : 2032
   StartDateTime        : 2/2/2015 6:57:03 AM
   EndDateTime          : 2/2/2015 6:57:11 AM
   Duration             : 00:00:07.3060000
   CreatedDateTime      : 2/2/2015 6:57:23 AM
   Status               : Completed
   ProtectionGroupName  : Clone
   SmProtectionGroupId  : 34
   PolicyName           : Vault
   SmPolicyId           : 18
   BackupName           : Clone_SCSPR0019366001_02-02-2015_06.57.08
   VerificationStatus   : NotVerified

   SmBackupId           : 114
   SmJobId              : 2183
   StartDateTime        : 2/2/2015 1:02:41 PM
   EndDateTime          : 2/2/2015 1:02:38 PM
   Duration             : -00:00:03.2300000
   CreatedDateTime      : 2/2/2015 1:02:53 PM
   Status               : Completed
   ProtectionGroupName  : Clone
   SmProtectionGroupId  : 34
   PolicyName           : Vault
   SmPolicyId           : 18
   BackupName           : Clone_SCSPR0019366001_02-02-2015_13.02.45
   VerificationStatus   : NotVerified
```

3. Restore data from the backup by using the Restore-SmBackup cmdlet.

```
Restore-SmBackup -PluginCode 'DummyPlugin' -AppObjectId
'scc54.sccore.test.com\DummyPlugin\NTP\DB1' -BackupId 269
-Confirm:$false
output:
Name                : Restore
'scc54.sccore.test.com\DummyPlugin\NTP\DB1'
Id                  : 2368
StartTime           : 10/4/2016 11:22:02 PM
EndTime             :
IsCancellable       : False
IsRestartable       : False
IsCompleted         : False
IsVisible           : True
IsScheduled         : False
PercentageCompleted : 0
Description         :
Status              : Queued
Owner               :
Error               :
Priority            : None
Tasks               : {}
ParentJobID         : 0
EventId             : 0
JobTypeId           :
ApisJobKey          :
ObjectId            : 0
PluginCode          : NONE
PluginName          :
```

The information regarding the parameters that can be used with the cmdlet and their descriptions can be obtained by running *Get-Help command_name*. Alternatively, you can also refer to the SnapCenter Software Cmdlet Reference Guide.

Restore and recover an auto discovered database backup

You can use SnapCenter to restore and recover data from one or more backups.

Before you begin

- You must have backed up the resource or resource groups.
- You must have canceled any backup operation that is currently in progress for the resource or resource group that you want to restore.
- For pre restore, post restore, mount, and unmount commands, you should check if the commands exist in the command list available on the plug-in host from the following paths:
  - Default location on the Windows host: *C:\Program Files\NetApp\SnapCenter\Snapcenter Plug-in*

Creator\etc\allowed_commands.config

- Default location on the Linux host: */opt/NetApp/snapcenter/scc/etc/allowed_commands.config*

If the commands do not exist in the command list, then the operation will fail.

About this task

- For Auto-discovered resources, restore is supported with SFSR.
- Point-in-time and Up-to-minute auto-recovery is not supported.
- For ONTAP 9.12.1 and below version, the clones created from the SnapLock Vault snapshots as part of restore will inherit the SnapLock Vault expiry time. Storage admin should manually cleanup the clones post the SnapLock expiry time.

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page, filter resources from the **View** drop-down list based on resource type.

   The resources are displayed along with type, host, associated resource groups and policies, and status.

   Although a backup might be for a resource group, when you restore, you must select the individual resources you want to restore.

   If the resource is not protected, “Not protected” is displayed in the Overall Status column. This can mean either that the resource is not protected, or that the resource was backed up by a different user.

3. Select the resource, or select a resource group and then select a resource in that group.

   The resource topology page is displayed.

4. From the Manage Copies view, select **Backups** either from the primary or secondary (mirrored or vaulted) storage systems.
5. In the Primary backup(s) table, select the backup that you want to restore from, and then click .

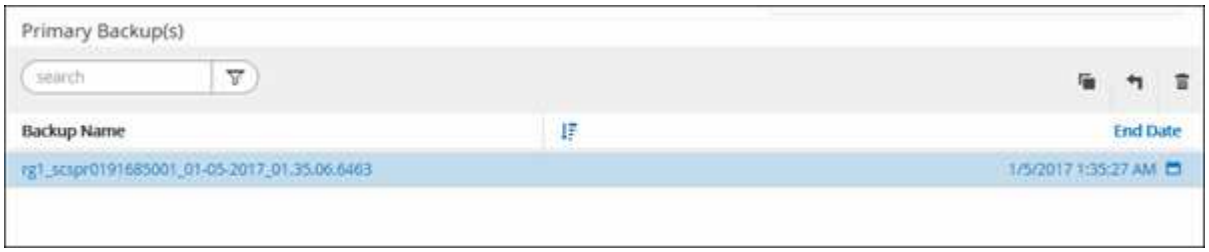

6. In the Restore Scope page, select **Complete Resource** to restore the configured data volumes of the MySQL database.
7. In the Pre ops page, enter pre restore and unmount commands to run before performing a restore job.

   Unmount commands are not available for auto discovered resources.

8. In the Post ops page, enter mount and post restore commands to run after performing a restore job.

   Mount commands are not available for auto discovered resources.

9. In the Notification page, from the **Email preference** drop-down list, select the scenarios in which you want to send the emails.

   You must also specify the sender and receiver email addresses and the subject of the email. SMTP must also be configured on the **Settings** > **Global Settings** page.

10. Review the summary, and then click **Finish**.
11. Monitor the operation progress by clicking **Monitor** > **Jobs**.

Restore resources using PowerShell cmdlets

Restoring a resource backup includes initiating a connection session with the SnapCenter Server, listing the backups and retrieving backup information, and restoring a backup.

You must have prepared the PowerShell environment to execute the PowerShell cmdlets.

Steps

1. Initiate a connection session with the SnapCenter Server for a specified user by using the Open-SmConnection cmdlet.

```
PS C:\> Open-Smconnection
```

2. Retrieve the information about the one or more backups that you want to restore by using the Get-SmBackup and Get-SmBackupReport cmdlets.

   This example displays information about all available backups:

```
PS C:\> Get-SmBackup

BackupId                      BackupName                    BackupTime
BackupType
--------                      ----------                    ----------
----------
  1               Payroll Dataset_vise-f6_08... 8/4/2015    11:02:32 AM
Full Backup
  2               Payroll Dataset_vise-f6_08... 8/4/2015    11:23:17 AM
```

   This example displays detailed information about the backup from January 29th 2015 to February 3rd, 2015:

```
PS C:\> Get-SmBackupReport -FromDateTime "1/29/2015" -ToDateTime
"2/3/2015"

   SmBackupId           : 113
   SmJobId              : 2032
   StartDateTime        : 2/2/2015 6:57:03 AM
   EndDateTime          : 2/2/2015 6:57:11 AM
   Duration             : 00:00:07.3060000
   CreatedDateTime      : 2/2/2015 6:57:23 AM
   Status               : Completed
   ProtectionGroupName  : Clone
   SmProtectionGroupId  : 34
   PolicyName           : Vault
   SmPolicyId           : 18
   BackupName           : Clone_SCSPR0019366001_02-02-2015_06.57.08
   VerificationStatus   : NotVerified

   SmBackupId           : 114
   SmJobId              : 2183
   StartDateTime        : 2/2/2015 1:02:41 PM
   EndDateTime          : 2/2/2015 1:02:38 PM
   Duration             : -00:00:03.2300000
   CreatedDateTime      : 2/2/2015 1:02:53 PM
   Status               : Completed
   ProtectionGroupName  : Clone
   SmProtectionGroupId  : 34
   PolicyName           : Vault
   SmPolicyId           : 18
   BackupName           : Clone_SCSPR0019366001_02-02-2015_13.02.45
   VerificationStatus   : NotVerified
```

3. Restore data from the backup by using the Restore-SmBackup cmdlet.

```
Restore-SmBackup -PluginCode 'DummyPlugin' -AppObjectId
'scc54.sccore.test.com\DummyPlugin\NTP\DB1' -BackupId 269
-Confirm:$false
output:
Name                : Restore
'scc54.sccore.test.com\DummyPlugin\NTP\DB1'
Id                  : 2368
StartTime           : 10/4/2016 11:22:02 PM
EndTime             :
IsCancellable       : False
IsRestartable       : False
IsCompleted         : False
IsVisible           : True
IsScheduled         : False
PercentageCompleted : 0
Description         :
Status              : Queued
Owner               :
Error               :
Priority            : None
Tasks               : {}
ParentJobID         : 0
EventId             : 0
JobTypeId           :
ApisJobKey          :
ObjectId            : 0
PluginCode          : NONE
PluginName          :
```

The information regarding the parameters that can be used with the cmdlet and their descriptions can be obtained by running *Get-Help command_name*. Alternatively, you can also refer to the SnapCenter Software Cmdlet Reference Guide.

Monitor MySQL restore operations

You can monitor the progress of different SnapCenter restore operations by using the Jobs page. You might want to check the progress of an operation to determine when it is complete or if there is an issue.

About this task

Post-restore states describe the conditions of the resource after a restore operation and any further restore actions that you can take.

The following icons appear on the Jobs page, and indicate the state of the operation:

- In progress

- Completed successfully
- Failed
- Completed with warnings or could not start due to warnings
- Queued
- Canceled

Steps

1. In the left navigation pane, click **Monitor**.
2. In the **Monitor** page, click **Jobs**.
3. In the **Jobs** page, perform the following steps:
   a. Click to filter the list so that only restore operations are listed.
   b. Specify the start and end dates.
   c. From the **Type** drop-down list, select **Restore**.
   d. From the **Status** drop-down list, select the restore status.
   e. Click **Apply** to view the operations that have been completed successfully.
4. Select the restore job, and then click **Details** to view the job details.
5. In the **Job Details** page, click **View logs**.

   The **View logs** button displays the detailed logs for the selected operation.

Clone MySQL resource backups

Clone workflow

The clone workflow includes performing the clone operation and monitoring the operation.

About this task

- You can clone on the source MySQL server.
- You might clone resource backups for the following reasons:
  - To test functionality that has to be implemented using the current resource structure and content during application development cycles
  - For data extraction and manipulation tools when populating data warehouses
  - To recover data that was mistakenly deleted or changed

The following workflow shows the sequence in which you must perform the clone operation:

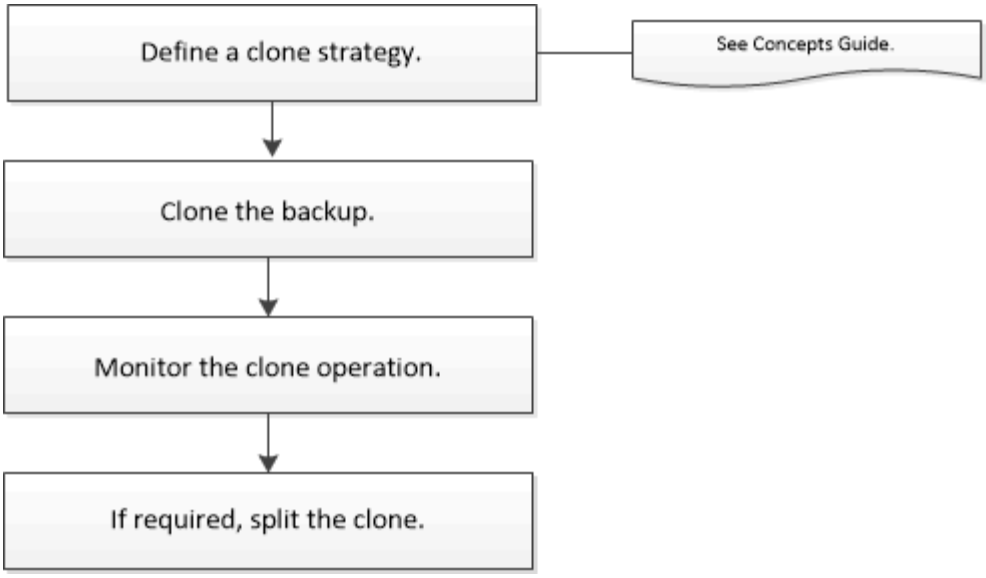

You can also use PowerShell cmdlets manually or in scripts to perform backup, restore, and clone operations. The SnapCenter cmdlet help and the cmdlet reference information contain detailed information about PowerShell cmdlets.

Clone a MySQL backup

You can use SnapCenter to clone a backup. You can clone from primary or secondary backup.

Before you begin

- You should have backed up the resources or resource group.
- You should ensure that the aggregates hosting the volumes should be in the assigned aggregates list of the storage virtual machine (SVM).
- For pre clone or post clone commands, you should check if the commands exist in the command list available on the plug-in host from the following paths:
  - Default location on the Windows host: *C:\Program Files\NetApp\SnapCenter\Snapcenter Plug-in Creator\etc\allowed_commands.config*
  - Default location on the Linux host: */opt/NetApp/snapcenter/scc/etc/allowed_commands.config*

If the commands do not exist in the command list, then the operation will fail.
* For MySQL 5.7 version, you should set IGNORE_MYSQLX_PORT = true (by default false) in the MySQL. properties file.

About this task

- You cannot protect the cloned MySQL instances.
- For information about FlexClone volume split operations, see, Split a FlexClone volume from its parent volume.
- For ONTAP 9.12.1 and earlier version, the clones created from the SnapLock Vault Snapshots as part of restore will inherit the SnapLock Vault expiry time. Storage admin should manually cleanup the clones post the SnapLock expiry time.

SnapCenter UI

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page, filter resources from the **View** drop-down list based on resource type.

   The resources are displayed along with information such as type, host, associated resource groups and policies, and status.

3. Select the resource or resource group.

   You must select a resource if you select a resource group.

   The resource or resource group topology page is displayed.

4. From the Manage Copies view, select **Backups** either from the primary or secondary (mirrored or vaulted) storage systems.
5. Select the data backup from the table, and then click .
6. In the Location page, perform the following actions:

For this field…	Do this…
Clone server	Choose a host on which the clone should be created.
Port	Provide the port on which the cloned MySQL instance to be started on.
NFS Export IP Address	Enter IP addresses or the host names on which the cloned volumes will be exported.

7. In the Scripts page, perform the following steps:

   The scripts are run on the plug-in host.

   a. Enter the commands for pre clone or post clone that should be run before or after the clone operation, respectively.
      - Pre clone command: delete existing databases with the same name
      - Post clone command: verify a database or start a database.
   b. Enter the mount command to mount a file system to a host.

      Mount command for a volume or qtree on a Linux machine:

      Example for NFS: `mount VSERVER_DATA_IP:%VOLUME_NAME_Clone /mnt`

8. In the Notification page, from the **Email preference** drop-down list, select the scenarios in which you want to send the emails.

   You must also specify the sender and receiver email addresses, and the subject of the email.

9. Review the summary, and then click **Finish**.
10. Monitor the operation progress by clicking **Monitor** > **Jobs**.

PowerShell cmdlets

Steps

1. Initiate a connection session with the SnapCenter Server for a specified user by using the Open-SmConnection cmdlet.

```
PS C:\> Open-SmConnection  -SMSbaseurl
https:\\snapctr.demo.netapp.com:8146/
```

2. Retrieve the backups to perform the clone operation by using the Get-SmBackup cmdlet.

   This example shows that two backups are available for cloning:

```
 C:\PS> Get-SmBackup

        BackupId                    BackupName
BackupTime                  BackupType
        --------                    ----------
----------                  ----------
         1                   Payroll Dataset_vise-f6_08...
8/4/2015 11:02:32 AM        Full Backup
         2                   Payroll Dataset_vise-f6_08...
8/4/2015 11:23:17 AM
```

3. Initiate a clone operation from an existing backup and specify the NFS export IP addresses on which the cloned volumes are exported.

   This example shows that the backup to be cloned has an NFSExportIPs address of 10.32.212.14:

```
PS C:\> New-SmClone -AppPluginCode MySQL -BackupName
"scs000211748_gdl_englab_netapp_com_MySQL_mysqld_3306_scs000211748_0
6-26-2024_06.08.35.4307" -Resources
@{"Host"="scs000211748.gdl.englab.netapp.com";"Uid"="mysqld_3306"}
-Port 3320 -CloneToHost shivarhel30.rtp.openenglab.netapp.com
```

   If NFSExportIPs is not specified, the default is exported to the clone target host.

4. Verify that the backups were cloned successfully by using the Get-SmCloneReport cmdlet to view the clone job details.

   You can view details such as clone ID, start date and time, end date and time.

```
PS C:\> Get-SmCloneReport -JobId 186

    SmCloneId           : 1
    SmJobId             : 186
    StartDateTime       : 8/3/2015 2:43:02 PM
    EndDateTime         : 8/3/2015 2:44:08 PM
    Duration            : 00:01:06.6760000
    Status              : Completed
    ProtectionGroupName : Draper
    SmProtectionGroupId : 4
    PolicyName          : OnDemand_Clone
    SmPolicyId          : 4
    BackupPolicyName    : OnDemand_Full_Log
    SmBackupPolicyId    : 1
    CloneHostName       : SCSPR0054212005.mycompany.com
    CloneHostId         : 4
    CloneName           : Draper__clone__08-03-2015_14.43.53
    SourceResources     : {Don, Betty, Bobby, Sally}
    ClonedResources     : {Don_DRAPER, Betty_DRAPER, Bobby_DRAPER,
Sally_DRAPER}
    SmJobError          :
```

Monitor MySQL clone operations

You can monitor the progress of SnapCenter clone operations by using the Jobs page. You might want to check the progress of an operation to determine when it is complete or if there is an issue.

About this task

The following icons appear on the Jobs page, and indicate the state of the operation:

- In progress
- Completed successfully
- Failed
- Completed with warnings or could not start due to warnings
- Queued
- Canceled

Steps

1. In the left navigation pane, click **Monitor**.
2. In the **Monitor** page, click **Jobs**.

3. In the **Jobs** page, perform the following steps:
   a. Click to filter the list so that only clone operations are listed.
   b. Specify the start and end dates.
   c. From the **Type** drop-down list, select **Clone**.
   d. From the **Status** drop-down list, select the clone status.
   e. Click **Apply** to view the operations that are completed successfully.
4. Select the clone job, and then click **Details** to view the job details.
5. In the Job Details page, click **View logs**.

Split a clone

You can use SnapCenter to split a cloned resource from the parent resource. The clone that is split becomes independent of the parent resource.

About this task

- You cannot perform the clone split operation on an intermediate clone.

  For example, after you create clone1 from a database backup, you can create a backup of clone1, and then clone this backup (clone2). After you create clone2, clone1 is an intermediate clone, and you cannot perform the clone split operation on clone1. However, you can perform the clone split operation on clone2.

  After splitting clone2, you can perform the clone split operation on clone1 because clone1 is no longer the intermediate clone.

- When you split a clone, the backup copies and clone jobs of the clone are deleted.
- For information about FlexClone volume split operations, see, Split a FlexClone volume from its parent volume.
- Ensure that the volume or aggregate on the storage system is online.

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. In the **Resources** page, select the appropriate option from the View list:

Option	Description
For database applications	Select **Database** from the View list.
For file systems	Select **Path** from the View list.

3. Select the appropriate resource from the list.

   The resource topology page is displayed.

4. From the **Manage Copies** view, select the cloned resource (for example, the database or LUN), and then click .
5. Review the estimated size of the clone that is to be split and the required space available on the aggregate, and then click **Start**.

6. Monitor the operation progress by clicking **Monitor** > **Jobs**.

   The clone split operation stops responding if the SMCore service restarts. You should run the Stop-SmJob cmdlet to stop the clone split operation, and then retry the clone split operation.

   If you want a longer poll time or shorter poll time to check whether the clone is split or not, you can change the value of *CloneSplitStatusCheckPollTime* parameter in *SMCoreServiceHost.exe.config* file to set the time interval for SMCore to poll for the status of the clone split operation. The value is in milliseconds and the default value is 5 minutes.

   For example:

   ```
   <add key="CloneSplitStatusCheckPollTime" value="300000" />
   ```

   The clone split start operation fails if backup, restore, or another clone split is in progress. You should restart the clone split operation only after the running operations are complete.

Related information

SnapCenter clone or verification fails with aggregate does not exist

Delete or split MySQL database clones after upgrading SnapCenter

After upgrading to SnapCenter 4.3, you will no longer see the clones. You can delete the clone or split the clones from the Topology page of the resource from which the clones were created.

About this task

If you want to locate the storage footprint of the hidden clones, run the following command: `Get-SmClone -ListStorageFootprint`

Steps

1. Delete the backups of the cloned resources by using the remove-smbackup cmdlet.
2. Delete the resource group of the cloned resources by using the remove-smresourcegroup cmdlet.
3. Remove the protection of the cloned resource by using the remove-smprotectresource cmdlet.
4. Select the parent resource from the Resources page.

   The resource topology page is displayed.

5. From the Manage Copies view, select the clones either from the primary or secondary (mirrored or replicated) storage systems.
6. Select the clones, and then click to delete clones or click to split the clones.
7. Click **OK**.

Protect applications using NetApp supported plug-ins

NetApp supported plug-ins

NetApp supported plug-ins overview

You can use the NetApp supported plug-ins like MongoDB, ORASCPM (Oracle Applications), SAP ASE, SAP MaxDB, and Storage plug-in for applications that you use and then use SnapCenter to backup, restore, or clone these applications. Your NetApp supported plug-ins act as host-side components of the NetApp SnapCenter Software, enabling application-aware data protection and management of resources.

When NetApp supported plug-ins are installed, you can use SnapCenter with NetApp SnapMirror technology to create mirror copies of backup sets on another volume and use NetApp SnapVault technology to perform disk-to-disk backup replication. The NetApp supported plug-ins can be used in both Windows and Linux environments.

SnapCenterCLI does not support NetApp supported plug-ins commands.

NetApp provides the Storage plug-in to perform data protection operations of the data volume on the ONTAP storage using the plug-in framework built into SnapCenter.

You can install the NetApp supported plug-ins from the Add Host page.

Add hosts and install plug-in packages on remote hosts.

SnapCenter support policy will cover support for the plug-in framework, core engine, and the associated APIs. Support will not cover the plug-in source code and the associated scripts built on the plug-in framework.

What you can do with the NetApp supported plug-ins

You can use the NetApp supported plug-ins like MongoDB, ORASCPM, Oracle Applications, SAP ASE, SAP MaxDB, and Storage plug-in for data protection operations.

- Add resources such as databases, instances, documents, or tablespaces.
- Create backups.
- Restore from backups.
- Clone backups.
- Schedule backup operations.
- Monitor backup, restore, and clone operations.
- View reports for backup, restore, and clone operations.

You can use the NetApp supported plug-ins for data protection operations.

- Take consistency group Snapshots of the storage volumes across ONTAP clusters.
- Backup custom applications using the built in pre and post scripting framework

  You can backup ONTAP volume, LUN, or a Qtree.

- Update Snapshots taken on the primary to an ONTAP secondary, leveraging the existing replication relationship (SnapVault/SnapMirror/unified replication) using SnapCenter policy

  ONTAP primary and secondary can be ONTAP FAS, AFF, ASA, ONTAP Select, or Cloud Volumes ONTAP.

- Recover complete ONTAP volume, LUN, or files.

  You should provide the respective file path manually as the browse or indexing features are not built into the product.

  Qtree or directory restore is not supported but you can clone and export only the Qtree if the backup scope is defined at a Qtree level.

NetApp supported plug-ins features

SnapCenter integrates with the plug-in application and with NetApp technologies on the storage system. To work with NetApp supported plug-ins like MongoDB, ORASCPM (Oracle Applications), SAP ASE, SAP MaxDB, and Storage plug-in you use the SnapCenter graphical user interface.

- **Unified graphical user interface**

  The SnapCenter interface provides standardization and consistency across plug-ins and environments. The SnapCenter interface enables you to complete consistent backup, restore, recovery, and clone operations across plug-ins, use centralized reporting, use at-a-glance dashboard views, set up role-based access control (RBAC), and monitor jobs across all plug-ins.

- **Automated central administration**

  You can schedule backup operations, configure policy-based backup retention, and perform restore operations. You can also proactively monitor your environment by configuring SnapCenter to send email alerts.

- **Nondisruptive NetApp Snapshot technology**

  SnapCenter uses NetApp Snapshot technology with the NetApp supported plug-ins to back up resources. Snapshots consume minimal storage space.

The NetApp supported plug-ins also offers the following benefits:

- Support for backup, restore, and clone workflows
- RBAC-supported security and centralized role delegation

  You can also set the credentials so that the authorized SnapCenter users have application-level permissions.

- Creation of space-efficient and point-in-time copies of resources for testing or data extraction by using

NetApp FlexClone technology

A FlexClone license is required on the storage system where you want to create the clone.

- Support for the consistency group (CG) Snapshot feature of ONTAP as part of creating backups.
- Capability to run multiple backups simultaneously across multiple resource hosts

  In a single operation, Snapshots are consolidated when resources in a single host share the same volume.

- Capability to create Snapshot using external commands.
- Capability to create file system consistent Snapshots in Windows environments.

Storage types supported by NetApp supported plug-ins

SnapCenter supports a wide range of storage types on both physical and virtual machines. You must verify the support for your storage type before installing NetApp supported plug-ins.

Machine	Storage type
Physical and NFS direct mounts on the VM hosts (VMDKs and RDM LUNs are not supported.)	FC-connected LUNs
Physical and NFS direct mounts on the VM hosts (VMDKs and RDM LUNs are not supported.)	iSCSI-connected LUNs
Physical and NFS direct mounts on the VM hosts (VMDKs and RDM LUNs are not supported.)	NFS-connected volumes
VMware ESXi	vVol datastores on both NFS and SAN vVol datastore can only be provisioned with ONTAP Tools for VMware vSphere.

Minimum ONTAP privileges required for NetApp supported plug-in

The minimum ONTAP privileges that are required vary according to the SnapCenter plug-ins you are using for data protection.

- All-access commands: Minimum privileges required for ONTAP 9.12.1 and later
  - event generate-autosupport-log
  - job history show
  - job stop
  - lun attribute show
  - lun create
  - lun delete
  - lun geometry

- lun igroup add
- lun igroup create
- lun igroup delete
- lun igroup rename
- lun igroup show
- lun mapping add-reporting-nodes
- lun mapping create
- lun mapping delete
- lun mapping remove-reporting-nodes
- lun mapping show
- lun modify
- lun move-in-volume
- lun offline
- lun online
- lun resize
- lun serial
- lun show
- network interface
- snapmirror policy add-rule
- snapmirror policy modify-rule
- snapmirror policy remove-rule
- snapmirror policy show
- snapmirror restore
- snapmirror show
- snapmirror show-history
- snapmirror update
- snapmirror update-ls-set
- snapmirror list-destinations
- version
- volume clone create
- volume clone show
- volume clone split start
- volume clone split stop
- volume create
- volume destroy
- volume file clone create
- volume file show-disk-usage

    - volume offline
    - volume online
    - volume modify
    - volume qtree create
    - volume qtree delete
    - volume qtree modify
    - volume qtree show
    - volume restrict
    - volume show
    - volume snapshot create
    - volume snapshot delete
    - volume snapshot modify
    - volume snapshot rename
    - volume snapshot restore
    - volume snapshot restore-file
    - volume snapshot show
    - volume unmount
    - vserver cifs
    - vserver cifs share create
    - vserver cifs share delete
    - vserver cifs shadowcopy show
    - vserver cifs share show
    - vserver cifs show
    - vserver export-policy create
    - vserver export-policy delete
    - vserver export-policy rule create
    - vserver export-policy rule show
    - vserver export-policy show
    - vserver iscsi connection show
    - vserver show
- Read-only commands: Minimum privileges required for ONTAP 8.3.0 and later
    - network interface

Prepare storage systems for SnapMirror and SnapVault replication for NetApp supported plug-ins

You can use a SnapCenter plug-in with ONTAP SnapMirror technology to create mirror copies of backup sets on another volume, and with ONTAP SnapVault technology to perform disk-to-disk backup replication for standards compliance and other governance-

related purposes. Before you perform these tasks, you must configure a data-protection relationship between the source and destination volumes and initialize the relationship.

SnapCenter performs the updates to SnapMirror and SnapVault after it completes the Snapshot operation. SnapMirror and SnapVault updates are performed as part of the SnapCenter job; do not create a separate ONTAP schedule.

If you are coming to SnapCenter from a NetApp SnapManager product and are satisfied with the data protection relationships you have configured, you can skip this section.

A data protection relationship replicates data on primary storage (the source volume) to secondary storage (the destination volume). When you initialize the relationship, ONTAP transfers the data blocks referenced on the source volume to the destination volume.

SnapCenter does not support cascade relationships between SnapMirror and SnapVault volumes (**Primary** > **Mirror** > **Vault**). You should use fanout relationships.

SnapCenter supports the management of version-flexible SnapMirror relationships. For details about version-flexible SnapMirror relationships and how to set them up, see the ONTAP documentation.

Define a backup strategy

Defining a backup strategy before you create your backup jobs ensures that you have the backups that you require to successfully restore or clone your resources. Your service-level agreement (SLA), recovery time objective (RTO), and recovery point objective (RPO) largely determine your backup strategy.

About this task

An SLA defines the level of service that is expected and addresses many service-related issues, including the availability and performance of the service. RTO is the time by which a business process must be restored after a disruption in service. RPO defines the strategy for the age of the files that must be recovered from backup storage for regular operations to resume after a failure. SLA, RTO, and RPO contribute to the data protection strategy.

Steps

1. Determine when you should back up your resources.
2. Decide how many backup jobs you require.
3. Decide how to name your backups.
4. Decide if you want Consistency Group Snapshots and decide on appropriate options for deleting Consistency Group Snapshots.
5. Decide whether you want to use NetApp SnapMirror technology for replication or NetApp SnapVault technology for long term retention.
6. Determine the retention period for the Snapshots on the source storage system and the SnapMirror destination.
7. Determine if you want to run any commands before or after the backup operation and provide a prescript or postscript.

Backup strategy for NetApp supported plug-ins

Backup schedules of NetApp supported plug-in resources

The most critical factor in determining a backup schedule is the rate of change for the resource. The more often you back up your resources, the fewer archive logs SnapCenter has to use for restoring, which can result in faster restore operations.

You might back up a heavily used resource every hour, while you might back up a rarely used resource once a day. Other factors include the importance of the resource to your organization, your service-level agreement (SLA) and your recovery point objective (RPO).

SLA defines the level of service expected and addresses many service-related issues, including the availability and performance of service. RPO defines the strategy for the age of the files that must be recovered from backup storage for regular operations to resume after a failure. SLA and RPO contribute to the data protection strategy.

Backup schedules have two parts, as follows:

- Backup frequency

  Backup frequency (how often backups are to be performed), also called schedule type for some plug-ins, is part of a policy configuration. For example, you might configure the backup frequency as hourly, daily, weekly or monthly. You can access policies in the SnapCenter GUI by clicking **Settings** > **Policies**.

- Backup schedules

  Backup schedules (exactly when backups are to be performed) are part of a resource or resource group configuration. For example, if you have a resource group that has a policy configured for weekly backups, you might configure the schedule to back up every Thursday at 10:00 p.m. You can access resource group schedules in the SnapCenter GUI by clicking **Resources**, then selecting the appropriate plug-in, and clicking **View** > **Resource Group**.

Number of backup jobs needed

Factors that determine the number of backup jobs that you need include the size of the resource, the number of volumes used, the rate of change of the resource, and your Service Level Agreement (SLA).

The number of backup jobs that you choose typically depends on the number of volumes on which you placed your resources. For example, if you placed a group of small resources on one volume and a large resource on another volume, you might create one backup job for the small resources and one backup job for the large resource.

Types of restore strategies supported for manually added NetApp supported plug-in resources

You must define a strategy before you can successfully perform restore operations using SnapCenter. There are two types of restore strategies for manually added NetApp supported plug-in resources.

You cannot recover manually added NetApp supported plug-in resources.

Complete resource restore

- Restores all volumes, qtrees, and LUNs of a resource

If the resource contains volumes or qtrees, the Snapshots taken after the Snapshot selected for restore on such volumes or qtrees are deleted and cannot be recovered. Also, if any other resource is hosted on the same volumes or qtrees, then that resource is also deleted.

File level restore

- Restores files from volumes, qtrees, or directories
- Restores only the selected LUNs

Prepare to install NetApp supported plug-ins

Installation workflow of SnapCenter NetApp supported plug-ins

You should install and set up SnapCenter NetApp supported plug-ins if you want to protect NetApp supported plug-in resources.

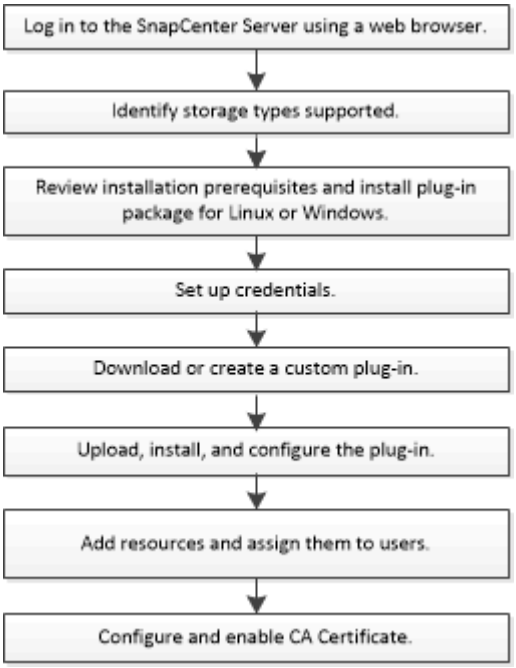

Prerequisites for adding hosts and installing Plug-ins package for Windows, Linux, or AIX

Before you add a host and install the plug-ins packages, you must complete all the requirements. The NetApp supported plug-ins are supported on Windows, Linux, and AIX environments.

Storage and Oracle applications are supported on AIX.

- You must have installed Java 11 on your Linux, Windows, or AIX host.

  IBM Java is not supported on Windows and Linux host.

- When installing a plug-in on a Windows host, if you specify a credential that is not built-in, or if the user belongs to a local workgroup user, you must disable UAC on the host.
- The NetApp supported plug-ins like MongoDB, ORASCPM, Oracle Applications, SAP ASE, SAP MaxDB, and Storage plug-in must be available on the client host from where the add host operation is performed.

General

If you are using iSCSI, the iSCSI service should be running.

Windows hosts

- You must have a domain user with local administrator privileges with local login permissions on the remote host.
- If you manage cluster nodes in SnapCenter, you must have a user with administrative privileges to all the nodes in the cluster.
- You must manually choose SnapCenter Plug-in for Microsoft Windows.

  Download JAVA for Windows

Linux and AIX hosts

Storage and Oracle applications are supported on AIX.

- You must have enabled the password-based SSH connection for the root or non-root user.
- You must have installed Java 11 on your Linux host.

  If you are using Windows Server 2019 or Windows Server 2016 for the SnapCenter Server host, you must install Java 11. The Interoperability Matrix Tool (IMT) contains the latest information about requirements.

  Download JAVA for Linux

  Download JAVA for AIX

  NetApp Interoperability Matrix Tool

- You must configure sudo privileges for the non-root user to provide access to several paths. Add the following lines to the /etc/sudoers file by using the visudo Linux utility.

  Ensure that you are using Sudo version 1.8.7 or later.

```
Cmnd_Alias HPPLCMD = sha224:checksum_value== /home/
LINUX_USER/.sc_netapp/snapcenter_linux_host_plugin.bin,
/opt/NetApp/snapcenter/spl/installation/plugins/uninstall,
/opt/NetApp/snapcenter/spl/bin/spl, /opt/NetApp/snapcenter/scc/bin/scc
Cmnd_Alias PRECHECKCMD = sha224:checksum_value== /home/
LINUX_USER/.sc_netapp/Linux_Prechecks.sh
Cmnd_Alias CONFIGCHECKCMD = sha224:checksum_value==
/opt/NetApp/snapcenter/spl/plugins/scu/scucore/configurationcheck/Config
_Check.sh
Cmnd_Alias SCCMD = sha224:checksum_value==
/opt/NetApp/snapcenter/spl/bin/sc_command_executor
Cmnd_Alias SCCCMDEXECUTOR =checksum_value==
/opt/NetApp/snapcenter/scc/bin/sccCommandExecutor
LINUX_USER ALL=(ALL) NOPASSWD:SETENV: HPPLCMD, PRECHECKCMD,
CONFIGCHECKCMD, SCCCMDEXECUTOR, SCCMD
Defaults: LINUX_USER env_keep += "IATEMPDIR"
Defaults: LINUX_USER env_keep += "JAVA_HOME"
Defaults: LINUX_USER !visiblepw
Defaults: LINUX_USER !requiretty
```

LINUX_USER is the name of the non-root user that you created.

You can obtain the *checksum_value* from the **sc_unix_plugins_checksum.txt** file, which is located at:

- *C:\ProgramData\NetApp\SnapCenter\Package Repository\sc_unix_plugins_checksum.txt* if SnapCenter Server is installed on Windows host.
- */opt/NetApp/snapcenter/SnapManagerWeb/Repository/sc_unix_plugins_checksum.txt* if SnapCenter Server in installed on Linux host.

> The example should be used only as a reference for creating your own data.

AIX Host requirements

You should ensure that the host meets the requirements before installing the SnapCenter Plug-ins Package for AIX.

> Storage and Oracle applications are supported on AIX.

> SnapCenter Plug-in for UNIX which is part of the SnapCenter Plug-ins Package for AIX, does not support concurrent volume groups.

Item	Requirements
Operating systems	AIX 7.1 or later
Minimum RAM for the SnapCenter plug-in on host	4 GB

Item	Requirements
Minimum install and log space for the SnapCenter plug-in on host	2 GB You should allocate sufficient disk space and monitor the storage consumption by the logs folder. The log space required varies depending on the number of the entities to be protected and the frequency of data protection operations. If there is no sufficient disk space, the logs will not be created for the recently run operations.
Required software packages	Java 11 IBM Java If you have upgraded JAVA to the latest version, you must ensure that the JAVA_HOME option located at /var/opt/snapcenter/spl/etc/spl.properties is set to the correct JAVA version and the correct path.

For the latest information about supported versions, see the NetApp Interoperability Matrix Tool.

Configure sudo privileges for non-root users for AIX host

SnapCenter 4.4 and later allows a non-root user to install the SnapCenter Plug-ins Package for AIX and to start the plug-in process. The plug-in processes will be running as an effective non-root user. You should configure sudo privileges for the non-root user to provide access to several paths.

What you will need

- Sudo version 1.8.7 or later.
- Edit the */etc/ssh/sshd_config* file to configure the message authentication code algorithms: MACs hmac-sha2-256 and MACs hmac-sha2-512.

  Restart the sshd service after updating the configuration file.

  Example:

```
#Port 22
#AddressFamily any
#ListenAddress 0.0.0.0
#ListenAddress ::
#Legacy changes
#KexAlgorithms diffie-hellman-group1-sha1
#Ciphers aes128-cbc
#The default requires explicit activation of protocol
Protocol 2
HostKey/etc/ssh/ssh_host_rsa_key
MACs hmac-sha2-256
```

About this task

You should configure sudo privileges for the non-root user to provide access to the following paths:

- /home/*AIX_USER*/.sc_netapp/snapcenter_aix_host_plugin.bsx
- /custom_location/NetApp/snapcenter/spl/installation/plugins/uninstall
- /custom_location/NetApp/snapcenter/spl/bin/spl

Steps

1. Log in to the AIX host on which you want to install the SnapCenter Plug-ins Package for AIX.
2. Add the following lines to the /etc/sudoers file by using the visudo Linux utility.

```
Cmnd_Alias HPPACMD = sha224:checksum_value== /home/
AIX_USER/.sc_netapp/snapcenter_aix_host_plugin.bsx,
/opt/NetApp/snapcenter/spl/installation/plugins/uninstall,
/opt/NetApp/snapcenter/spl/bin/spl
Cmnd_Alias PRECHECKCMD = sha224:checksum_value== /home/
AIX_USER/.sc_netapp/AIX_Prechecks.sh
Cmnd_Alias CONFIGCHECKCMD = sha224:checksum_value==
/opt/NetApp/snapcenter/spl/plugins/scu/scucore/configurationcheck/Config
_Check.sh
Cmnd_Alias SCCMD = sha224:checksum_value==
/opt/NetApp/snapcenter/spl/bin/sc_command_executor
AIX_USER ALL=(ALL) NOPASSWD:SETENV: HPPACMD, PRECHECKCMD,
CONFIGCHECKCMD, SCCMD
Defaults: LINUX_USER env_keep += "IATEMPDIR"
Defaults: LINUX_USER env_keep += "JAVA_HOME"
Defaults: AIX_USER !visiblepw
Defaults: AIX_USER !requiretty
```

If you are having a RAC setup, along with the other allowed commands, you should add the following to the /etc/sudoers file: '/<crs_home>/bin/olsnodes'

You can obtain the value of *crs_home* from the */etc/oracle/olr.loc* file.

AIX_USER is the name of the non-root user that you created.

You can obtain the *checksum_value* from the **sc_unix_plugins_checksum.txt** file, which is located at:

- *C:\ProgramData\NetApp\SnapCenter\Package Repository\sc_unix_plugins_checksum.txt* if SnapCenter Server is installed on Windows host.
- */opt/NetApp/snapcenter/SnapManagerWeb/Repository/sc_unix_plugins_checksum.txt* if SnapCenter Server in installed on Linux host.

The example should be used only as a reference for creating your own data.

Host requirements to install SnapCenter Plug-ins Package for Windows

Before you install the SnapCenter Plug-ins Package for Windows, you should be familiar with some basic host system space requirements and sizing requirements.

Item	Requirements
Operating Systems	Microsoft Windows For the latest information about supported versions, see the NetApp Interoperability Matrix Tool.
Minimum RAM for the SnapCenter plug-in on host	1 GB
Minimum install and log space for the SnapCenter plug-in on host	5 GB You should allocate sufficient disk space and monitor the storage consumption by the logs folder. The log space required varies depending on the number of the entities to be protected and the frequency of data protection operations. If there is no sufficient disk space, the logs will not be created for the recently run operations.

Item	Requirements
Required software packages	• ASP.NET Core Runtime 8.0.12 (and all subsequent 8.0.x patches) Hosting Bundle • PowerShell Core 7.4.2 • Java 11 Oracle Java and OpenJDK Java 11 Oracle Java and OpenJDK is required only for SAP HANA, IBM Db2, PostgreSQL, MySQL, NetApp supported plug-ins, and other custom applications which can be installed on Windows host. For the latest information about supported versions, see the NetApp Interoperability Matrix Tool. For .NET specific troubleshooting information, see SnapCenter upgrade or install fails for legacy systems that do not have internet connectivity.

Host requirements for installing the SnapCenter Plug-ins Package for Linux and AIX

You should ensure that the host meets the requirements before installing the SnapCenter Plug-ins Package for Linux or AIX.

Storage and Oracle applications are supported on AIX.

Item	Requirements
Operating systems	• Red Hat Enterprise Linux • Oracle Linux • SUSE Linux Enterprise Server (SLES)
Minimum RAM for the SnapCenter plug-in on host	1 GB
Minimum install and log space for the SnapCenter plug-in on host	2 GB You should allocate sufficient disk space and monitor the storage consumption by the logs folder. The log space required varies depending on the number of the entities to be protected and the frequency of data protection operations. If there is no sufficient disk space, the logs will not be created for the recently run operations.

Item	Requirements
Required software packages	Java 11 Oracle Java or OpenJDK If you have upgraded JAVA to the latest version, you must ensure that the JAVA_HOME option located at /var/opt/snapcenter/spl/etc/spl.properties is set to the correct JAVA version and the correct path.

For the latest information about supported versions, see the NetApp Interoperability Matrix Tool

Set up credentials for NetApp supported plug-ins

SnapCenter uses credentials to authenticate users for SnapCenter operations. You should create credentials for installing SnapCenter plug-ins and additional credentials for performing data protection operations on databases or Windows file systems.

Before you begin

- Linux or AIX hosts

  You must set up credentials for installing plug-ins on Linux or AIX hosts.

  You must set up the credentials for the root user or for a non-root user who has sudo privileges to install and start the plug-in process.

  Best Practice: Although you are allowed to create credentials for Linux after deploying hosts and installing plug-ins, the best practice is to create credentials after you add SVMs, before you deploy hosts and install plug-ins.

- Windows hosts

  You must set up Windows credentials before installing plug-ins.

  You must set up the credentials with administrator privileges, including administrator rights on the remote host.

- NetApp supported plug-ins applications

  The plug-in uses the credentials that are selected or created while adding a resource. If a resource does not require credentials during data protection operations, you can set the credentials as **None**.

About this task

If you set up credentials for individual resource groups and the username does not have full admin privileges, you must assign at least the resource group and backup privileges to the username.

Steps

1. In the left navigation pane, click **Settings**.
2. In the Settings page, click **Credential**.
3. Click **New**.
4. In the **Credential** page, specify the information required for configuring credentials:

For this field…	Do this…
Credential name	Enter a name for the credentials.
User name	Enter the user name and password that are to be used for authentication. • Domain administrator or any member of the administrator group Specify the domain administrator or any member of the administrator group on the system on which you are installing the SnapCenter plug-in. Valid formats for the Username field are: ◦ *NetBIOS\UserName* ◦ *Domain FQDN\UserName* • Local administrator (for workgroups only) For systems that belong to a workgroup, specify the built-in local administrator on the system on which you are installing the SnapCenter plug-in. You can specify a local user account that belongs to the local administrators group if the user account has elevated privileges or the User Access control feature is disabled on the host system. The valid format for the Username field is: *UserName*
Password	Enter the password used for authentication.
Authentication Type	Select the authentication type that you want to use.
Use sudo privileges	Select the **Use sudo privileges** check box if you are creating credentials for a non-root user. Applicable to Linux and AIX users only.

5. Click **OK**.

After you finish setting up the credentials, you might want to assign credential maintenance to a user or group of users on the User and Access page.

Configure gMSA on Windows Server 2016 or later

Windows Server 2016 or later enables you to create a group Managed Service Account (gMSA) that provides automated service account password management from a

managed domain account.

Before you begin

- You should have a Windows Server 2016 or later domain controller.
- You should have a Windows Server 2016 or later host, which is a member of the domain.

Steps

1. Create a KDS root key to generate unique passwords for each object in your gMSA.
2. For each domain, run the following command from the Windows domain controller: Add-KDSRootKey -EffectiveImmediately
3. Create and configure your gMSA:
   a. Create a user group account in the following format:

      ```
      domainName\accountName$
      ```

   b. Add computer objects to the group.
   c. Use the user group you just created to create the gMSA.

      For example,

      ```
      New-ADServiceAccount -name <ServiceAccountName> -DNSHostName <fqdn>
      -PrincipalsAllowedToRetrieveManagedPassword <group>
      -ServicePrincipalNames <SPN1,SPN2,…>
      ```

   d. Run `Get-ADServiceAccount` command to verify the service account.
4. Configure the gMSA on your hosts:
   a. Enable the Active Directory module for Windows PowerShell on the host where you want to use the gMSA account.

      To do this, run the following command from PowerShell:

```
PS C:\> Get-WindowsFeature AD-Domain-Services

Display Name                                Name                Install
State
------------                                ----
-------------
[ ] Active Directory Domain Services   AD-Domain-Services  Available


PS C:\> Install-WindowsFeature AD-DOMAIN-SERVICES

Success Restart Needed Exit Code      Feature Result
------- -------------- ---------      --------------
True    No             Success        {Active Directory Domain
Services, Active ...
WARNING: Windows automatic updating is not enabled. To ensure that
your newly-installed role or feature is
automatically updated, turn on Windows Update.
```

b. Restart your host.

c. Install the gMSA on your host by running the following command from the PowerShell command prompt: `Install-AdServiceAccount <gMSA>`

d. Verify your gMSA account by running the following command: `Test-AdServiceAccount <gMSA>`

5. Assign the administrative privileges to the configured gMSA on the host.
6. Add the Windows host by specifying the configured gMSA account in the SnapCenter Server.

   SnapCenter Server will install the selected plug-ins on the host and the specified gMSA will be used as the service log on account during the plug-in installation.

Install the NetApp supported plug-ins

Add hosts and install plug-in packages on remote hosts

You must use the SnapCenter Add Host page to add hosts, and then install the plug-in packages. The plug-ins are automatically installed on the remote hosts. You can add a host and install the plug-in packages either for an individual host or for a cluster.

Before you begin

- You should be a user that is assigned to a role that has the plug-in install and uninstall permissions, such as the SnapCenter Admin role.
- You should ensure that the message queueing service is running.
- If you are using group Managed Service Account (gMSA), you should configure gMSA with administrative privileges.

  Configure group Managed Service Account on Windows Server 2016 or later for custom applications

- For Windows host, you must ensure that you select SnapCenter Plug-in for Windows.

About this task

- You cannot add a SnapCenter Server as a plug-in host to another SnapCenter Server.
- If you install plug-ins on a cluster (WSFC), the plug-ins are installed on all of the nodes of the cluster.

Steps

1. In the left navigation pane, select **Hosts**.
2. Verify that the **Managed Hosts** tab is selected at the top.
3. Select **Add**.
4. In the Hosts page, perform the following actions:

For this field…	Do this…
Host Type	Select the host type: • Windows • Linux • AIX The NetApp supported plug-ins can be used in Windows, Linux, and AIX environments. Storage and Oracle applications are supported on AIX.
Host name	Enter the fully qualified domain name (FQDN) or the IP address of the host. SnapCenter depends on the proper configuration of the DNS. Therefore, the best practice is to enter the FQDN. For Windows environments, the IP address is supported for untrusted domain hosts only if it resolves to the FQDN. You can enter the IP addresses or FQDN of a stand-alone host. If you are adding a host by using SnapCenter and the host is part of a subdomain, you must provide the FQDN.

For this field…	Do this…
Credentials	Either select the credential name that you created, or create new credentials. The credentials must have administrative rights on the remote host. For details, see the information about creating credentials. You can view details about the credentials by positioning your cursor over the credential name that you specified. The credentials authentication mode is determined by the host type that you specify in the Add Host wizard.

5. In the **Select Plug-ins to Install** section, select the plug-ins to install.

   You can install the following plug-ins from the list:

   - MongoDB
   - ORASCPM (displayed as Oracle Applications)
   - SAP ASE
   - SAP MaxDB
   - Storage

6. (Optional) Select **More Options** to install the other plug-ins.

For this field…	Do this…
Port	Either retain the default port number, or specify the port number. The default port number is 8145. If the SnapCenter Server was installed on a custom port, that port number will be displayed as the default port. If you manually installed the plug-ins and specified a custom port, you must specify the same port. Otherwise, the operation fails.

For this field…	Do this…
Installation Path	The NetApp supported plug-ins can be installed on either a Windows system or Linux system. • For the SnapCenter Plug-ins Package for Windows, the default path is C:\Program Files\NetApp\SnapCenter. Optionally, you can customize the path. • For SnapCenter Plug-ins Package for Linux and SnapCenter Plug-ins Package for AIX, the default path is `/opt/NetApp/snapcenter`. Optionally, you can customize the path.
Skip preinstall checks	Select this check box if you already installed the plug-ins manually and you do not want to validate whether the host meets the requirements for installing the plug-in.
Use group Managed Service Account (gMSA) to run the plug-in services	For Windows host, select this check box if you want to use group Managed Service Account (gMSA) to run the plug-in services. Provide the gMSA name in the following format: domainName\accountName$. gMSA will be used as a log on service account only for SnapCenter Plug-in for Windows service.

7. Select **Submit**.

   If you have not selected the **Skip prechecks** checkbox, the host is validated to verify whether the host meets the requirements for installing the plug-in. The disk space, RAM, PowerShell version, .NET version, location (for Windows plug-ins), and Java version (for Linux plug-ins) are validated against the minimum requirements. If the minimum requirements are not met, appropriate error or warning messages are displayed.

   If the error is related to disk space or RAM, you can update the web.config file located at `C:\Program Files\NetApp\SnapCenter WebApp` to modify the default values. If the error is related to other parameters, you must fix the issue.

   In an HA setup, if you are updating SnapManager.Web.UI.dll.config, you must update the file on both nodes and restart the SnapCenter App Pool.

   Windows default path is `C:\Program Files\NetApp\SnapCenter WebApp\SnapManager.Web.UI.dll.config`

Linux default path is `/opt/NetApp/snapcenter/SnapManagerWeb/SnapManager.Web.UI.dll.config`

8. If host type is Linux, verify the fingerprint, and then select **Confirm and Submit**.

   Fingerprint verification is mandatory even if the same host was added earlier to SnapCenter and the fingerprint was confirmed.

9. Monitor the installation progress.

   The installation-specific log files are located at `/custom_location/snapcenter/` logs.

Install SnapCenter Plug-in Packages for Linux, Windows, or AIX on multiple remote hosts by using cmdlets

You can install the SnapCenter Plug-in Packages for Linux, Windows, or AIX on multiple hosts simultaneously by using the Install-SmHostPackage PowerShell cmdlet.

Before you begin

The user adding a host should have the administrative rights on the host.

Storage and Oracle applications are supported on AIX.

Steps

1. Launch PowerShell.
2. On the SnapCenter Server host, establish a session using the Open-SmConnection cmdlet, and then enter your credentials.
3. Install the plug-in on multiple hosts using the Install-SmHostPackage cmdlet and the required parameters.

   The information regarding the parameters that can be used with the cmdlet and their descriptions can be obtained by running *Get-Help command_name*. Alternatively, you can also refer to the SnapCenter Software Cmdlet Reference Guide.

   You can use the -skipprecheck option when you have installed the plug-ins manually and do not want to validate whether the host meets the requirements to install the plug-in.

4. Enter your credentials for remote installation.

Install the NetApp supported plug-ins on Linux hosts by using the command-line interface

You should install the NetApp supported plug-ins by using the SnapCenter user interface (UI). If your environment does not allow remote installation of the plug-in from the SnapCenter UI, you can install the NetApp supported plug-ins either in console mode or in silent mode by using the command-line interface (CLI).

Steps

1. Copy the SnapCenter Plug-ins Package for Linux installation file (snapcenter_linux_host_plugin.bin) from C:\ProgramData\NetApp\SnapCenter\Package Repository to the host where you want to install the NetApp supported plug-ins.

You can access this path from the host where the SnapCenter Server is installed.

2. From the command prompt, navigate to the directory where you copied the installation file.
3. Install the plug-in: `path_to_installation_bin_file/snapcenter_linux_host_plugin.bin -i silent -DPORT=port_number_for_host -DSERVER_IP=server_name_or_ip_address -DSERVER_HTTPS_PORT=port_number_for_server`
   - -DPORT specifies the SMCore HTTPS communication port.
   - -DSERVER_IP specifies the SnapCenter Server IP address.
   - -DSERVER_HTTPS_PORT specifies the SnapCenter Server HTTPS port.
   - -DUSER_INSTALL_DIR specifies the directory where you want to install the SnapCenter Plug-ins Package for Linux.
   - _DINSTALL_LOG_NAME specifies the name of the log file.

```
/tmp/sc-plugin-installer/snapcenter_linux_host_plugin.bin -i silent
-DPORT=8145 -DSERVER_IP=scserver.domain.com -DSERVER_HTTPS_PORT=8146
-DUSER_INSTALL_DIR=/opt
-DINSTALL_LOG_NAME=SnapCenter_Linux_Host_Plugin_Install_2.log
-DCHOSEN_FEATURE_LIST=CUSTOM
```

4. Add the host to the SnapCenter Server using the Add-Smhost cmdlet and the required parameters.

   The information regarding the parameters that can be used with the command and their descriptions can be obtained by running *Get-Help command_name*. Alternatively, you can also refer to the SnapCenter Software Cmdlet Reference Guide.

5. Log in to SnapCenter and upload the NetApp supported plug-in from the UI or by using PowerShell cmdlets.

   You can upload the NetApp supported plug-in from the UI by referring to Add hosts and install plug-in packages on remote hosts section.

   The SnapCenter cmdlet help and the cmdlet reference information contain more information about PowerShell cmdlets.

   SnapCenter Software Cmdlet Reference Guide.

Monitor the status of installing NetApp supported plug-ins

You can monitor the progress of SnapCenter plug-in package installation by using the Jobs page. You might want to check the progress of installation to determine when it is complete or if there is an issue.

About this task

The following icons appear on the Jobs page and indicate the state of the operation:

- In progress
- Completed successfully

- Failed
- Completed with warnings or could not start due to warnings
- Queued

Steps

1. In the left navigation pane, click **Monitor**.
2. In the **Monitor** page, click **Jobs**.
3. In the **Jobs** page, to filter the list so that only plug-in installation operations are listed, do the following:
   a. Click **Filter**.
   b. Optional: Specify the start and end date.
   c. From the Type drop-down menu, select **Plug-in installation**.
   d. From the Status drop-down menu, select the installation status.
   e. Click **Apply**.
4. Select the installation job and click **Details** to view the job details.
5. In the **Job Details** page, click **View logs**.

Configure CA Certificate

Generate CA Certificate CSR file

You can generate a Certificate Signing Request (CSR) and import the certificate that can be obtained from a Certificate Authority (CA) using the generated CSR. The certificate will have a private key associated with it.

CSR is a block of encoded text that is given to an authorized certificate vendor to procure the signed CA certificate.

CA Certificate RSA key length must be minimum 3072 bits.

For information to generate a CSR, see How to generate CA Certificate CSR file.

If you own the CA certificate for your domain (*.domain.company.com) or your system (machine1.domain.company.com), you can skip generating the CA Certificate CSR file. You can deploy the existing CA certificate with SnapCenter.

For cluster configurations, the cluster name (virtual cluster FQDN), and the respective host names should be mentioned in the CA certificate. The certificate can be updated by filling the Subject Alternative Name (SAN) field before procuring the certificate. For a wild card certificate (*.domain.company.com), the certificate will contain all the hostnames of the domain implicitly.

Import CA certificates

You must import the CA certificates to the SnapCenter Server and the Windows host plug-ins using the Microsoft management console (MMC).

Steps

1. Go to the Microsoft management console (MMC), and then click **File** > **Add/Remove Snapin**.
2. In the Add or Remove Snap-ins window, select **Certificates** and then click **Add**.
3. In the Certificates snap-in window, select the **Computer account** option, and then click **Finish**.
4. Click **Console Root** > **Certificates – Local Computer** > **Trusted Root Certification Authorities** > **Certificates**.
5. Right-click on the folder “Trusted Root Certification Authorities”, and then select **All Tasks** > **Import** to start the import wizard.
6. Complete the wizard, as follows:

In this wizard window…	Do the following…
Import Private Key	Select the option **Yes**, import the private key, and then click **Next**.
Import File Format	Make no changes; click **Next**.
Security	Specify the new password to be used for the exported certificate, and then click **Next**.
Completing the Certificate Import Wizard	Review the summary, and then click **Finish** to start the import.

Importing certificate should be bundled with the private key (supported formats are: *.pfx, *.p12, and *.p7b).

7. Repeat Step 5 for the “Personal” folder.

Get the CA certificate thumbprint

A certificate thumbprint is a hexadecimal string that identifies a certificate. A thumbprint is calculated from the content of the certificate using a thumbprint algorithm.

Steps

1. Perform the following on the GUI:
   a. Double-click the certificate.
   b. In the Certificate dialog box, click the **Details** tab.
   c. Scroll through the list of fields and click **Thumbprint**.
   d. Copy the hexadecimal characters from the box.
   e. Remove the spaces between the hexadecimal numbers.

      For example, if the thumbprint is: "a9 09 50 2d d8 2a e4 14 33 e6 f8 38 86 b0 0d 42 77 a3 2a 7b", after removing the spaces, it will be: "a909502dd82ae41433e6f83886b00d4277a32a7b".

2. Perform the following from PowerShell:
   a. Run the following command to list the thumbprint of the installed certificate and identify the recently installed certificate by the subject name.

Get-ChildItem -Path Cert:\LocalMachine\My

b. Copy the thumbprint.

Configure CA certificate with Windows host plug-in services

You should configure the CA certificate with Windows host plug-in services to activate the installed digital certificate.

Perform the following steps on the SnapCenter Server and all the plug-in hosts where CA certificates are already deployed.

Steps

1. Remove the existing certificate binding with SMCore default port 8145, by running the following command:

   ```
   > netsh http delete sslcert ipport=0.0.0.0:_<SMCore Port>
   ```

   For example:

   ```
   > netsh http delete sslcert ipport=0.0.0.0:8145
   ```

2. Bind the newly installed certificate with the Windows host plug-in services, by running the following commands:

   ```
   > $cert = "_<certificate thumbprint>_"
   > $guid = [guid]::NewGuid().ToString("B")
   > netsh http add sslcert ipport=0.0.0.0: _<SMCore Port>_ certhash=$cert
   appid="$guid"
   ```

   For example:

   ```
   > $cert = "a909502dd82ae41433e6f83886b00d4277a32a7b"
   > $guid = [guid]::NewGuid().ToString("B")
   > netsh http add sslcert ipport=0.0.0.0: _<SMCore Port>_ certhash=$cert
   appid="$guid"
   ```

Configure the CA Certificate for the NetApp supported plug-ins service on Linux host

You should manage the password of the plug-ins keystore and its certificate, configure the CA certificate, configure root or intermediate certificates to the plug-ins trust-store, and configure CA signed key pair to plug-ins trust-store with SnapCenter plug-ins service to activate the installed digital certificate.

The plug-ins uses the file 'keystore.jks', which is located at */opt/NetApp/snapcenter/scc/etc* both as its trust-store and key-store.

Manage password for plug-in keystore and alias of the CA signed key pair in use

Steps

1. You can retrieve plug-in keystore default password from plug-in agent property file.

   It is the value corresponding to the key 'KEYSTORE_PASS'.

2. Change the keystore password:

   ```
   keytool -storepasswd -keystore keystore.jks
   ```

3. Change the password for all aliases of private key entries in the keystore to the same password used for the keystore:

   ```
   keytool -keypasswd -alias "alias_name_in_cert" -keystore keystore.jks
   ```

   Update the same for the key KEYSTORE_PASS in *agent.properties* file.

4. Restart the service after changing the password.

> Password for plug-in keystore and for all the associated alias password of the private key should be same.

Configure root or intermediate certificates to plug-in trust-store

You should configure the root or intermediate certificates without the private key to plug-in trust-store.

Steps

1. Navigate to the folder containing the plug-in keystore: /opt/NetApp/snapcenter/scc/etc.
2. Locate the file 'keystore.jks'.
3. List the added certificates in the keystore:

   ```
   keytool -list -v -keystore keystore.jks
   ```

4. Add a root or intermediate certificate:

   ```
   keytool -import -trustcacerts -alias myRootCA -file
   /root/USERTrustRSA_Root.cer -keystore keystore.jks
   ```

5. Restart the service after configuring the root or intermediate certificates to plug-in trust-store.

> You should add the root CA certificate and then the intermediate CA certificates.

Configure CA signed key pair to plug-in trust-store

You should configure the CA signed key pair to the plug-in trust-store.

Steps

1. Navigate to the folder containing the plug-in keystore /opt/NetApp/snapcenter/scc/etc.
2. Locate the file 'keystore.jks'.
3. List the added certificates in the keystore:

   ```
   keytool -list -v -keystore keystore.jks
   ```

4. Add the CA certificate having both private and public key.

   ```
   keytool -importkeystore -srckeystore /root/snapcenter.ssl.test.netapp.com.pfx
   -srcstoretype pkcs12 -destkeystore keystore.jks -deststoretype JKS
   ```

5. List the added certificates in the keystore.

   ```
   keytool -list -v -keystore keystore.jks
   ```

6. Verify that the keystore contains the alias corresponding to the new CA certificate, which was added to the keystore.
7. Change the added private key password for CA certificate to the keystore password.

   Default plug-in keystore password is the value of the key KEYSTORE_PASS in agent.properties file.

   ```
   keytool -keypasswd -alias "alias_name_in_CA_cert" -keystore keystore.jks
   ```

8. If the alias name in the CA certificate is long and contains space or special characters ("*",","), change the alias name to a simple name:

   ```
   keytool -changealias -alias "long_alias_name" -destalias "simple_alias"
   -keystore keystore.jks
   ```

9. Configure the alias name from CA certificate in agent.properties file.

   Update this value against the key SCC_CERTIFICATE_ALIAS.

10. Restart the service after configuring the CA signed key pair to plug-in trust-store.

Configure certificate revocation list (CRL) for plug-ins

About this task

- SnapCenter Plug-ins will search for the CRL files in a pre-configured directory.
- Default directory for the CRL files for SnapCenter Plug-ins is ' opt/NetApp/snapcenter/scc/etc/crl'.

Steps

1. You can modify and update the default directory in agent.properties file against the key CRL_PATH.

   You can place more than one CRL file in this directory. The incoming certificates will be verified against each CRL.

Configure the CA Certificate for the NetApp supported plug-ins service on Windows host

You should manage the password of the plug-ins keystore and its certificate, configure the CA certificate, configure root or intermediate certificates to the plug-ins trust-store, and configure CA signed key pair to plug-ins trust-store with SnapCenter plug-ins service to activate the installed digital certificate.

The plug-ins uses the file *keystore.jks*, which is located at *C:\Program Files\NetApp\SnapCenter\Snapcenter Plug-in Creator\etc* both as its trust-store and key-store.

Manage password for plug-in keystore and alias of the CA signed key pair in use

Steps

1. You can retrieve plug-in keystore default password from plug-in agent property file.

   It is the value corresponding to the key *KEYSTORE_PASS*.

2. Change the keystore password:

   keytool -storepasswd -keystore keystore.jks

   > If the "keytool" command is not recognized on the Windows command prompt, replace the keytool command with its complete path.

   C:\Program Files\Java\<jdk_version>\bin\keytool.exe" -storepasswd -keystore keystore.jks

3. Change the password for all aliases of private key entries in the keystore to the same password used for the keystore:

   keytool -keypasswd -alias "alias_name_in_cert" -keystore keystore.jks

   Update the same for the key KEYSTORE_PASS in *agent.properties* file.

4. Restart the service after changing the password.

   > Password for plug-in keystore and for all the associated alias password of the private key should be same.

Configure root or intermediate certificates to plug-in trust-store

You should configure the root or intermediate certificates without the private key to plug-in trust-store.

Steps

1. Navigate to the folder containing the plug-in keystore *C:\Program Files\NetApp\SnapCenter\Snapcenter Plug-in Creator\etc*
2. Locate the file 'keystore.jks'.
3. List the added certificates in the keystore:

   keytool -list -v -keystore keystore.jks

4. Add a root or intermediate certificate:

keytool -import -trustcacerts -alias myRootCA -file /root/USERTrustRSA_Root.cer -keystore keystore.jks

5. Restart the service after configuring the root or intermediate certificates to plug-in trust-store.

> You should add the root CA certificate and then the intermediate CA certificates.

Configure CA signed key pair to plug-in trust-store

You should configure the CA signed key pair to the plug-in trust-store.

Steps

1. Navigate to the folder containing the plug-in keystore *C:\Program Files\NetApp\SnapCenter\Snapcenter Plug-in Creator\etc*
2. Locate the file *keystore.jks*.
3. List the added certificates in the keystore:

   keytool -list -v -keystore keystore.jks

4. Add the CA certificate having both private and public key.

   keytool -importkeystore -srckeystore /root/snapcenter.ssl.test.netapp.com.pfx -srcstoretype pkcs12 -destkeystore keystore.jks -deststoretype JKS

5. List the added certificates in the keystore.

   keytool -list -v -keystore keystore.jks

6. Verify that the keystore contains the alias corresponding to the new CA certificate, which was added to the keystore.
7. Change the added private key password for CA certificate to the keystore password.

   Default plug-in keystore password is the value of the key KEYSTORE_PASS in agent.properties file.

   keytool -keypasswd -alias "alias_name_in_CA_cert" -keystore keystore.jks

8. Configure the alias name from CA certificate in *agent.properties* file.

   Update this value against the key SCC_CERTIFICATE_ALIAS.

9. Restart the service after configuring the CA signed key pair to plug-in trust-store.

Configure certificate revocation list (CRL) for SnapCenter plug-ins

About this task

- To download the latest CRL file for the related CA certificate see How to update certificate revocation list file in SnapCenter CA Certificate.
- SnapCenter plug-ins will search for the CRL files in a pre-configured directory.
- Default directory for the CRL files for SnapCenter plug-ins is *'C:\Program Files\NetApp\SnapCenter\Snapcenter Plug-in Creator\ etc\crl'*.

Steps

1. You can modify and update the default directory in *agent.properties* file against the key CRL_PATH.

2. You can place more than one CRL file in this directory.

   The incoming certificates will be verified against each CRL.

Enable CA Certificates for plug-ins

You should configure the CA certificates and deploy the CA certificates in the SnapCenter Server and the corresponding plug-in hosts. You should enable the CA certificate validation for the plug-ins.

Before you begin

- You can enable or disable the CA certificates using the run *Set-SmCertificateSettings* cmdlet.
- You can display the certificate status for the plug-ins using the *Get-SmCertificateSettings*.

The information regarding the parameters that can be used with the cmdlet and their descriptions can be obtained by running *Get-Help command_name*. Alternatively, you can also refer to the SnapCenter Software Cmdlet Reference Guide.

Steps

1. In the left navigation pane, click **Hosts**.
2. In the Hosts page, click **Managed Hosts**.
3. Select single or multiple plug-in hosts.
4. Click **More options**.
5. Select **Enable Certificate Validation**.

After you finish

The Managed Hosts tab host displays a padlock and the color of the padlock indicates the status of the connection between SnapCenter Server and the plug-in host.

- indicates that the CA certificate is neither enabled nor assigned to the plug-in host.
- indicates that the CA certificate is successfully validated.
- indicates that the CA certificate could not be validated.
- indicates that the connection information could not be retrieved.

  When the status is yellow or green, the data protection operations completes successfully.

Prepare for data protection

Prerequisites for using the NetApp supported plug-ins

Before you use SnapCenter NetApp supported plug-ins, the SnapCenter administrator must install and configure the SnapCenter Server and perform the prerequisite tasks.

- Install and configure SnapCenter Server.
- Log in to SnapCenter Server.
- Configure the SnapCenter environment by adding storage system connections and creating credentials, if

applicable.

- Add hosts, and install and upload the plug-ins.
- If applicable, install Java 11 on the plug-in host.
- If you have multiple data paths (LIFs) or a dNFS configuration, you can perform the following using the SnapCenter CLI on the database host:
  - By default, all the IP addresses of the database host are added to the NFS storage export policy in storage virtual machine (SVM) for the cloned volumes. If you want to have a specific IP address or restrict to a subset of the IP addresses, run the Set-PreferredHostIPsInStorageExportPolicy CLI.
  - If you have multiple data paths (LIFs) in SVMs, SnapCenter chooses the appropriate data path (LIF) for mounting the NFS cloned volume. However, if you want to specify a specific data path (LIF), you must run the Set-SvmPreferredDataPath CLI.
    The information regarding the parameters that can be used with the command and their descriptions can be obtained by running *Get-Help command_name*. Alternatively, you can also refer to the SnapCenter Software Command Reference Guide.
- Set up SnapMirror and SnapVault, if you want backup replication.
- Ensure that port 9090 is not used by any other application on the host.

  Port 9090 must be reserved for use by NetApp supported plug-ins in addition to the other ports required by SnapCenter.

How resources, resource groups, and policies are used for protecting NetApp supported plug-in resources

Before you use SnapCenter, it is helpful to understand basic concepts related to the backup, clone, and restore operations you want to perform. You interact with resources, resource groups, and policies for different operations.

- Resources are typically databases, Windows file systems, or VMs that you back up or clone with SnapCenter.
- A SnapCenter resource group, is a collection of resources on a host or cluster.

  When you perform an operation on a resource group, you perform that operation on the resources defined in the resource group according to the schedule you specify for the resource group.

  You can back up on demand a single resource or a resource group. You also can perform scheduled backups for single resources and resource groups.
- The policies specify the backup frequency, copy retention, replication, scripts, and other characteristics of data protection operations.

  When you create a resource group, you select one or more policies for that group. You can also select a policy when you perform a backup on demand for a single resource.

Think of a resource group as defining *what* you want to protect and when you want to protect it in terms of day and time. Think of a policy as defining *how* you want to protect it. If you are backing up all databases or backing up all file systems of a host, for example, you might create a resource group that includes all the databases or all the file systems in the host. You could then attach two policies to the resource group: a daily policy and an hourly policy. When you create the resource group and attach the policies, you might configure the resource group to perform a File-Based backup daily and another schedule that performs Snapshot based

backup hourly.

Back up NetApp supported plug-ins resources

Back up NetApp supported plug-ins resources

The backup workflow includes planning, identifying the resources for backup, managing backup policies, creating resource groups and attaching policies, creating backups, and monitoring the operations.

The following workflow shows the sequence in which you must perform the backup operation:

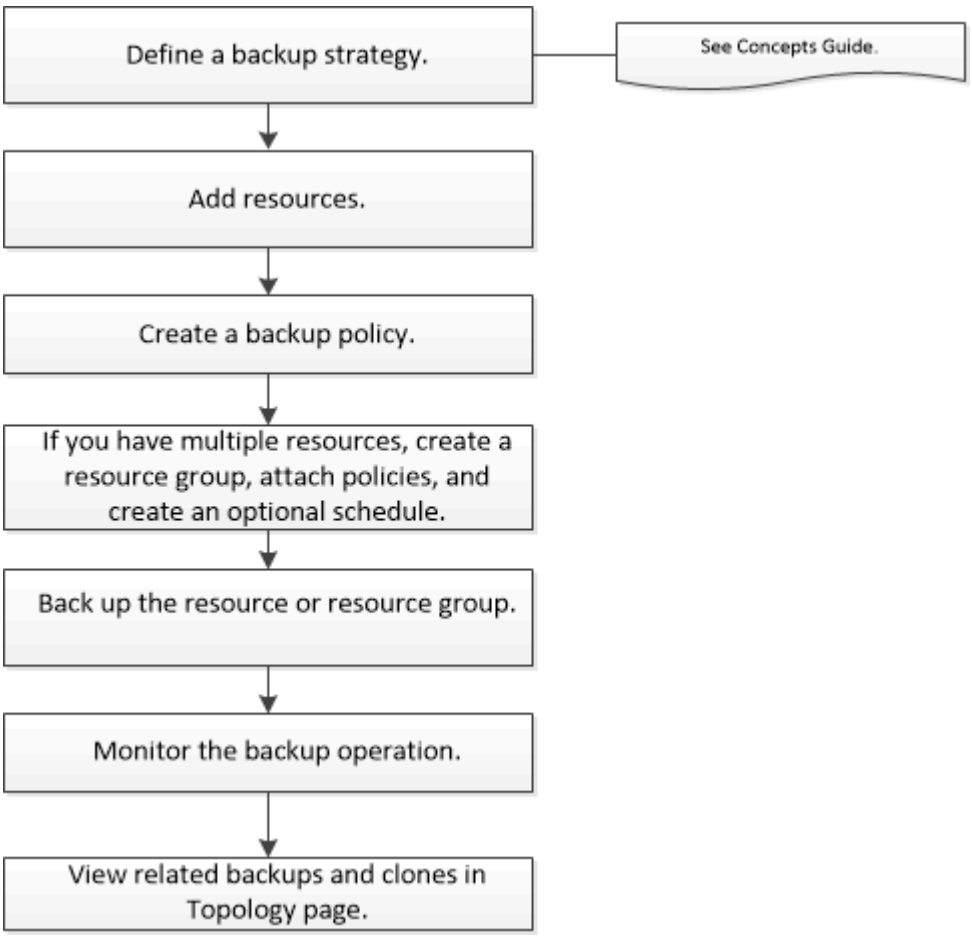

You can also use PowerShell cmdlets manually or in scripts to perform backup, restore, and clone operations. For detailed information about PowerShell cmdlets, use the SnapCenter cmdlet help or see the SnapCenter Software Cmdlet Reference Guide

Add resources to NetApp supported plug-ins

You must add the resources that you want to back up or clone. Depending on your environment, resources might be either database instances or collections that you want to back up or clone.

Before you begin

- You must have completed tasks such as installing the SnapCenter Server, adding hosts, creating storage

system connections, and adding credentials.

- You must have uploaded the plug-ins to SnapCenter Server.

Steps

1. In the left navigation pane, select **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page, select **Add Resource**.
3. In the Provide Resource Details page, perform the following actions:

For this field…	Do this…
Name	Enter the name of the resource.
Host name	Select the host.
Type	Select the type. Type is user defined as per the plug-in description file. For example, database and instance. In case the type selected has a parent, enter the details of the parent. For example, if the type is Database and the parent is Instance, enter the details of the Instance.
Credential name	Select Credential or create a new credential.
Mount Paths	Enter the mount paths where the resource is mounted. This is applicable only for a Windows host.

4. In the Provide Storage Footprint page, select a storage system and choose one or more volumes, LUNs, and qtrees, and then select **Save**.

   Optional: Select the + icon to add more volumes, LUNs, and qtrees from other storage systems.

   NetApp supported plug-ins does not support automatic discovery of the resources. The storage details of physical and virtual environments are also not discovered automatically. You must provide the storage information for physical and virtual environments while creating the resources.

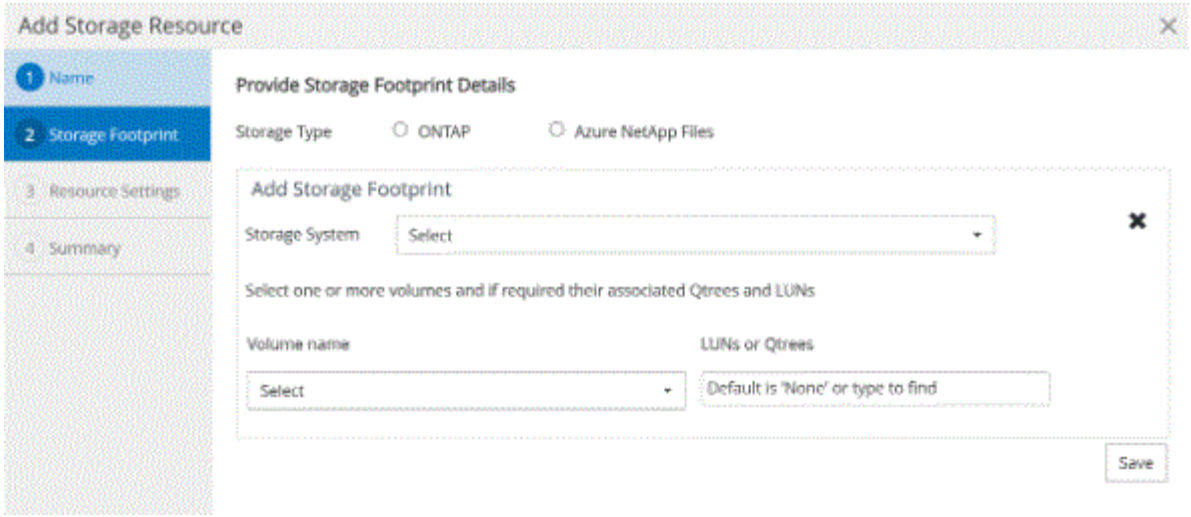

5. In the Resource Settings page, provide custom key-value pairs for the resource.

   Ensure that the custom keys name is in uppercase.

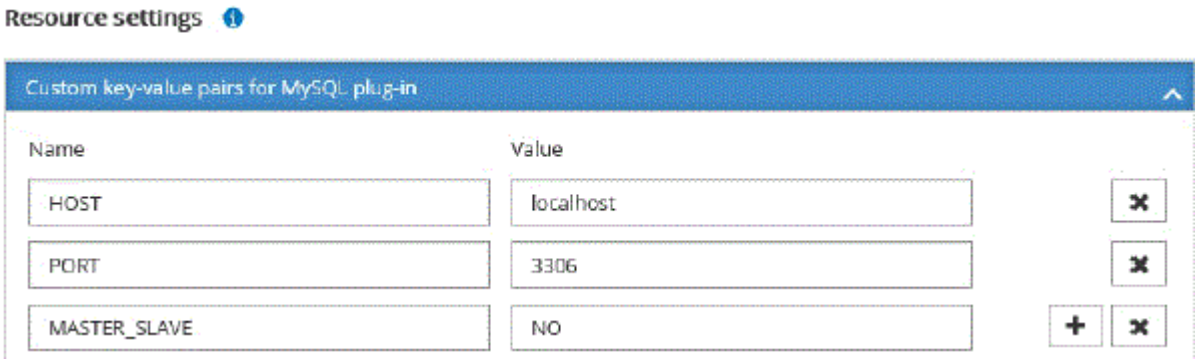

   For the respective plug-in parameters, refer Parameters to configure the resource

6. Review the summary, and then select **Finish**.

Result

The resources are displayed along with information such as type, host or cluster name, associated resource groups and policies, and overall status.

You must refresh the resources if the databases are renamed outside of SnapCenter.

After you finish

If you want to provide access to the assets to other users, the SnapCenter administrator must assign assets to those users. This enables users to perform the actions for which they have permissions on the assets that are assigned to them.

After adding the resources, you can modify the resource details. If a NetApp supported plug-ins resource has backups associated with it, the following fields cannot be modified: resource name, resource type, and host name.

Parameters to configure the resource

If you are adding the plug-ins manually, you can use the following parameters to configure the resource in the Resource Settings page.

Plug-in for MongoDB

Resource Settings:

- MONGODB_APP_SERVER=(for resource type as sharded cluster) or MONGODB_REPLICASET_SERVER=(for resource type as replicaset)
- OPLOG_PATH=(Optional parameter in case it is provided from MongoDB.propertiesfile)
- MONGODB_AUTHENTICATION_TYPE= (PLAIN for LDAP Authentication and None for others)

You must provide the following parameters needs to be provided n MongoDB.properties file:

- DISABLE_STARTING_STOPPING_SERVICES=
  - N if the start/stop services are performed by the plug-in.
  - Y if start/**stop services are performed by the user.
  - Optional parameter as default value is set to N.
- OPLOG_PATH_= (Optional parameter in case it is already provided as custom key-value pair in SnapCenter).

Plug-in for MaxDB

Resource Settings:

- XUSER_ENABLE (Y|N) enables or disables the use of an xuser for MaxDB so that a password is not required for the database user.
- HANDLE_LOGWRITER (Y|N) executes suspend logwriter (N) or resume logwriter (Y) operations.
- DBMCLICMD (path_to_dbmcli_cmd) specifies the path to the MaxDB dbmcli command. If not set, dbmcli on the search path is used.

  For Windows environment, the path must be within double-quotes ("…").

- SQLCLICMD (path_to_sqlcli_cmd) specifies the path to the MaxDB sqlcli command. If the path is not set, sqlcli is used on the search path.
- MAXDB_UPDATE_HIST_LOG (Y|N) instructs the MaxDB backup program whether it should update the MaxDB history log.
- MAXDB_CHECK_SNAPSHOT_DIR : Example, SID1:directory[,directory…]; [SID2:directoary[,directory…] checks that a Snap Creator Snapshot copy operation is successful and ensures that the snapshot is created.

  This applies to NFS only. The directory must point to the location that contains the .snapshot directory. Multiple directories can be included in a comma-separated list.

  In MaxDB 7.8 and later versions, the database backup request is marked Failed in the backup history.

- MAXDB_BACKUP_TEMPLATES: Specifies a backup template for each database.

The template must exist and be an external type of backup template. To enable snapshot integration for MaxDB 7.8 and later, you must have MaxDB background server functionality and already configured MaxDB backup template of the EXTERNAL type.

- MAXDB_BG_SERVER_PREFIX: Specifies the prefix for the background server name.

  If the MAXDB_BACKUP_TEMPLATES parameter is set, you must also set the MAXDB_BG_SERVER_PREFIX parameter. If you do not set the prefix, the default value na_bg_ is used.

Plug-in for SAP ASE

Resource Settings:

- SYBASE_SERVER (data_server_name) specifies the Sybase data server name (-S option on isql command). For example, p_test.
- SYBASE_DATABASES_EXCLUDE (db_name) allows databases to be excluded if the "ALL" construct is used.

  You can specify multiple databases by using a semicolon-separated list. For example: pubs2;test_db1.

- SYBASE_USER: user_name specifies the operating system user who can run the isql command.

  Required for UNIX. This parameter is required if the user running the Snap Creator Agent start and stop commands (usually the root user) and the user running the isql command are different.

- SYBASE_TRAN_DUMP db_name:directory_path enables you to perform a Sybase transaction dump after creating a snapshot. For example, pubs2:/sybasedumps/ pubs2

  You must specify each database requiring a transaction dump.

- SYBASE_TRAN_DUMP_COMPRESS (Y|N) enables or disables native Sybase transaction dump compression.
- SYBASE_ISQL_CMD (For example, /opt/sybase/OCS-15_0/bin/isql) defines the path to the isql command.
- SYBASE_EXCLUDE_TEMPDB (Y|N) allows you to auto exclude user created temporary databases.

Plug-in for Oracle applications (ORASCPM)

Resource Settings:

- SQLPLUS_CMD specifies the path to SQLplus.
- ORACLE_DATABASES lists the Oracle databases to be backed up and corresponding user (database:user).
- CNTL_FILE_BACKUP_DIR specifies the directory for control file back up.
- ORA_TEMP specifies the directory for temporary files.
- ORACLE_HOME specifies the directory where the Oracle software is installed.
- ARCHIVE_LOG_ONLY specifies whether to back up the archive logs or not.
- ORACLE_BACKUPMODE specifies whether to perform online or offline backup.
- ORACLE_EXPORT_PARAMETERS specifies whether the environment variables defined above should be re-exported while running */bin/su <user running sqlplus> -c sqlplus /nolog <cmd>*. This is typically the case,

when the user who is running sqlplus has not set all environment variables necessary to connect to the database using *connect / as sysdba*.

Create policies for NetApp supported plug-in resources

Before you use SnapCenter to back up NetApp supported plug-in specific resources, you must create a backup policy for the resource or resource group that you want to back up.

Before you begin

- You should have defined your backup strategy.

  For details, see the information about defining a data protection strategy for NetApp supported plug-ins.

- You should have prepared for data protection.

  Preparing for data protection includes tasks such as installing SnapCenter, adding hosts, creating storage system connections, and adding resources.

- The storage virtual machines (SVMs) should be assigned to you for mirror or vault operations.

  The SnapCenter administrator must have assigned the SVMs for both the source and destination volumes to you if you are replicating Snapshots to a mirror or vault.

- You should have manually added the resources that you want to protect.

About this task

- A backup policy is a set of rules that governs how you manage, schedule, and retain backups. Additionally, you can specify replication, script, and application settings.
- Specifying options in a policy saves time when you want to reuse the policy for another resource group.
- SnapLock
  - If 'Retain the backup copies for a specific number of days' option is selected, then the SnapLock retention period must be lesser than or equal to the mentioned retention days.
  - Specifying a Snapshot locking period prevents deletion of the Snapshots until the retention period expires. This could lead to retaining a larger number of Snapshots than the count specified in the policy.
  - For ONTAP 9.12.1 and below version, the clones created from the SnapLock Vault Snapshots as part of restore will inherit the SnapLock Vault expiry time. Storage admin should manually cleanup the clones post the SnapLock expiry time.

Primary SnapLock settings are managed in SnapCenter backup policy and the secondary SnapLock settings are managed by ONTAP.

Steps

1. In the left navigation pane, click **Settings**.
2. In the Settings page, click **Policies**.
3. Click **New**.
4. In the Name page, enter the policy name and details.
5. In the Policy type page, perform the following:

a. Select storage type.

b. In the Custom backup settings section, provide any specific backup settings that have to be passed to the plug-in in key-value format.

   You can provide multiple key-values to be passed to the plug-in.

6. In the Snapshot and Replication page, perform the following steps:

   a. Specify the schedule type by selecting **On demand**, **Hourly**, **Daily**, **Weekly**, or **Monthly**.

   You can specify the schedule (start date, end date, and frequency) for the backup operation while creating a resource group. This enables you to create resource groups that share the same policy and backup frequency, but enables you to assign different backup schedules to each policy.

   If you have scheduled for 2:00 a.m., the schedule will not be triggered during daylight saving time (DST).

   b. In the Snapshot settings section, specify the retention settings for the backup type and the schedule type selected in the **Backup Type** page:

If you want to…	Then…
Keep a certain number of Snapshots	Select **Copies to keep**, and then specify the number of Snapshots that you want to keep. If the number of Snapshots exceeds the specified number, the Snapshots are deleted with the oldest copies deleted first. You must set the retention count to 2 or higher if you plan to enable SnapVault replication. If you set the retention count to 1, the retention operation might fail because the first Snapshot is the reference Snapshot for the SnapVault relationship until a newer Snapshot is replicated to the target. The maximum retention value is 1018. Backups will fail if retention is set to a value higher than what the underlying ONTAP version supports.
Keep the Snapshots for a certain number of days	Select **Retain copies for**, and then specify the number of days for which you want to keep the Snapshots before deleting them.

If you want to…	Then…
Snapshot copy locking period	Select **Snapshot copy locking period**, and specify days, months, or years. SnapLock retention period should be less than 100 years.

c. Select a Policy label.

> You can assign SnapMirror labels to primary snapshots for remote replication, allowing the primary snapshots to offload the snapshot replication operation from SnapCenter to ONTAP secondary systems. This can be done without enabling SnapMirror or SnapVault option in the policy page.

7. In the Select secondary replication options section, select one or both of the following secondary replication options:

For this field…	Do this…
Update SnapMirror after creating a local Snapshot copy	Select this field to create mirror copies of the backup sets on another volume (SnapMirror replication). If the protection relationship in ONTAP is of type Mirror and Vault and if you select only this option, Snapshot created on the primary will not be transferred to the destination, but will be listed in the destination. If this Snapshot is selected from the destination to perform a restore operation, then the following error message is displayed: Secondary Location is not available for the selected vaulted/mirrored backup. During secondary replication, the SnapLock expiry time loads the primary SnapLock expiry time. Clicking the **Refresh** button in the Topology page refreshes the secondary and primary SnapLock expiry time that are retrieved from ONTAP. See View NetApp supported plug-ins resource related backups and clones in the Topology page.

For this field…	Do this…
Update SnapVault after creating a local Snapshot copy	Select this option to perform disk-to-disk backup replication (SnapVault backups). During secondary replication, the SnapLock expiry time loads the primary SnapLock expiry time. Clicking the **Refresh** button in the Topology page refreshes the secondary and primary SnapLock expiry time that are retrieved from ONTAP. When SnapLock is configured only on the secondary from ONTAP known as SnapLock Vault, clicking the **Refresh** button in the Topology page refreshes the locking period on the secondary that is retrieved from ONTAP. For more information on SnapLock Vault see Commit Snapshots to WORM on a vault destination. See View NetApp supported plug-ins resource related backups and clones in the Topology page.
Error retry count	Enter the maximum number of replication attempts that can be allowed before the operation stops.

ⓘ You should configure SnapMirror retention policy in ONTAP for the secondary storage to avoid reaching the maximum limit of Snapshots on the secondary storage.

8. Review the summary, and then click **Finish**.

Create resource groups and attach policies

A resource group is the container to which you must add resources that you want to back up and protect. It enables you to back up all the data that is associated with a given application simultaneously. You must also attach one or more policies to the resource group to define the type of data protection job that you want to perform.

Steps

1. In the left navigation pane, select **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page, select New Resource Group.
3. In the Name page, perform the following actions:

For this field…	Do this…
Name	Enter a name for the resource group. Note: The resource group name should not exceed 250 characters.

For this field…	Do this…
Tags	Enter one or more labels that will help you later search for the resource group. For example, if you add HR as a tag to multiple resource groups, you can later find all the resource groups associated with the HR tag.
Use custom name format for Snapshot copy	Select this check box, and enter a custom name format that you want to use for the Snapshot name. For example, *customtext_resource group_policy_hostname or resource group_hostname*. By default, a timestamp is appended to the Snapshot name.

4. Optional: In the Resources page, select a host name from the **Host** drop-down list and the resource type from the **Resource Type** drop-down list.

   This helps to filter information on the screen.

5. Select the resources from the **Available Resources** section, and then select the right arrow to move them to the **Selected Resources** section.
6. Optional: In the **Application Settings** page, do the following:
   a. Select the Backups arrow to set additional backup options:

      Enable consistency group backup and perform the following tasks:

For this field…	Do this…
Afford time to wait for Consistency Group Snapshot operation to complete	Select Urgent, Medium, or Relaxed to specify the wait time for Snapshot operation to complete. Urgent = 5 seconds, Medium = 7 seconds, and Relaxed = 20 seconds.
Disable WAFL Sync	Select this to avoid forcing a WAFL consistency point.

   b. Select the Scripts arrow and enter the pre and post commands for quiesce, Snapshot, and unquiesce operations. You can also enter the pre commands to be executed before exiting in the event of a failure.
   c. Select the Custom Configurations arrow and enter the custom key-value pairs required for all data protection operations using this resource.

Parameter	Setting	Description
ARCHIVE_LOG_ENABLE	(Y/N)	Enables the archive log management to delete the archive logs.
ARCHIVE_LOG_RETENTION	number_of_days	Specifies the number of days the archive logs are retained. This setting must be equal to or greater than NTAP_SNAPSHOT_ RETENTIONS.
ARCHIVE_LOG_DIR	change_info_directory/logs	Specifies the path to the directory that contains the archive logs.
ARCHIVE_LOG_EXT	file_extension	Specifies the archive log file extension length. For example, if the archive log is log_backup_0_0_0_0.161518551942 9 and if the file_extension value is 5, then the extension of the log will retain 5 digits, which is 16151.
ARCHIVE_LOG_RECURSIVE_SE ARCH	(Y/N)	Enables the management of archive logs within subdirectories. You should use this parameter if the archive logs are located under subdirectories.

d. Select the **Snapshot Copy Tool** arrow to select the tool to create Snapshots:

If you want…	Then…
SnapCenter to use the plug-in for Windows and put the file system into a consistent state before creating a Snapshot. For Linux resources, this option is not applicable.	Select **SnapCenter with File System Consistency**. This option is not applicable for SnapCenter Plug-in for SAP HANA Database.

If you want…	Then…
SnapCenter to create a storage level Snapshot	Select **SnapCenter without File System Consistency**.
To enter the command to be executed on the host to create Snapshots.	Select **Other**, and then enter the command to be executed on the host to create a Snapshot.

7. In the Policies page, perform the following steps:
   a. Select one or more policies from the drop-down list.

      You can also create a policy by selecting + .

      The policies are listed in the **Configure schedules for selected policies** section.

   b. In the **Configure Schedules** column, select + for the policy you want to configure.
   c. In the Add schedules for policy *policy_name* dialog box, configure the schedule and select OK.

      Where policy_name is the name of the policy that you have selected.

      The configured schedules are listed in the Applied Schedules column.
      Third party backup schedules are not supported when they overlap with SnapCenter backup schedules.

8. From the **Email preference** drop-down list on the **Notification** page, select the scenarios in which you want to send the emails.

   You must also specify the sender and receiver email addresses, and the subject of the email. The SMTP server must be configured in **Settings** > **Global Settings**.

9. Review the summary, and then select **Finish**.

Create resource groups and enable secondary protection for resources on ASA r2 systems

You should create the resource group to add the resources that are on ASA r2 systems. You can also provision the secondary protection while creating the resource group.

Before you begin

- You should ensure that you are not adding both ONTAP 9.x resources and ASA r2 resources to the same resource group.
- You should ensure that you do not have a database with both ONTAP 9.x resources and ASA r2 resources.

About this task

- The secondary protection is available only if the logged-in user is assigned to the role that has the **SecondaryProtection** capability enabled.
- If you enabled secondary protection, the resource group is put into maintenance mode while creating the primary and secondary consistency groups. After the primary and secondary consistency groups are

created, the resource group is put out of maintenance mode.

- SnapCenter does not support secondary protection for a clone resource.

Steps

1. In the left navigation pane, select **Resources**, and the appropriate plug-in from the list.
2. In the Resources page, click **New Resource Group**.
3. In the Name page, perform the following actions:
   a. Enter a name for the resource group in the Name field.

      The resource group name should not exceed 250 characters.

   b. Enter one or more labels in the Tag field to help you search for the resource group later.

      For example, if you add HR as a tag to multiple resource groups, you can later find all resource groups associated with the HR tag.

   c. Select this check box, and enter a custom name format that you want to use for the Snapshot name.

      For example, customtext_resource group_policy_hostname or resource group_hostname. By default, a timestamp is appended to the Snapshot name.

   d. Specify the destinations of the archive log files that you do not want to back up.

      You should use the exact same destination as it was set in the application including prefix, if needed.

4. In the Resources page, select the database host name from the **Host** drop-down list.

   The resources are listed in the Available Resources section only if the resource is discovered successfully. If you have recently added resources, they will appear on the list of available resources only after you refresh your resource list.

5. Select the ASA r2 resources from the Available Resources section and move them to the Selected Resources section.
6. In the Application Settings page, select the backup option.
7. In the Policies page, perform the following steps:
   a. Select one or more policies from the drop-down list.

      You can also create a policy by clicking + .

      In the Configure schedules for selected policies section, the selected policies are listed.

   b. Click + in the Configure Schedules column for the policy for which you want to configure a schedule.
   c. In the Add schedules for policy *policy_name* window, configure the schedule, and then click **OK**.

      Where, *policy_name* is the name of the policy that you have selected.

The configured schedules are listed in the Applied Schedules column.

Third party backup schedules are not supported when they overlap with SnapCenter backup schedules.

8. If the secondary protection is enabled for the policy that you have selected, then Secondary Protection page is displayed and you need to perform the following steps:
   a. Select the type of the replication policy.

      Synchronous replication policy is not supported.

   b. Specify the consistency group suffix that you want to use.
   c. From the Destination Cluster and Destination SVM drop-downs select the peered cluster and SVM that you want to use.

      The cluster and SVM peering is not supported by SnapCenter. You should use System Manager or ONTAP CLIs to perform cluster and SVM peering.

If the resources are already protected outside of SnapCenter, those resources will be displayed in the Secondary Protected Resources section.

1. On the Verification page, perform the following steps:
   a. Click **Load locators** to load the SnapMirror or SnapVault volumes to perform verification on secondary storage.
   b. Click + in the Configure Schedules column to configure the verification schedule for all the schedule types of the policy.
   c. In the Add Verification Schedules policy_name dialog box, perform the following actions:

If you want to…	Do this…
Run verification after backup	Select **Run verification after backup**.
Schedule a verification	Select **Run scheduled verification** and then select the schedule type from the drop-down list.

   d. Select **Verify on secondary location** to verify your backups on secondary storage system.
   e. Click **OK**.

      The configured verification schedules are listed in the Applied Schedules column.

2. In the Notification page, from the **Email preference** drop-down list, select the scenarios in which you want to send the emails.

   You must also specify the sender and receiver email addresses, and the subject of the email. If you want to attach the report of the operation performed on the resource group, select **Attach Job Report**.

   For email notification, you must have specified the SMTP server details using the either the GUI or the PowerShell command Set-SmSmtpServer.

3. Review the summary, and then click **Finish**.

Create a storage system connection and a credential using PowerShell cmdlets

You must create a storage virtual machine (SVM) connection and a credential before using PowerShell cmdlets to perform data protection operations.

Before you begin

- You should have prepared the PowerShell environment to execute the PowerShell cmdlets.
- You should have the required permissions in the Infrastructure Admin role to create storage connections.
- You should ensure that the plug-in installations are not in progress.

  Host plug-in installations must not be in progress while adding a storage system connection because the host cache might not be updated and databases status might be displayed in the SnapCenter GUI as “Not available for backup” or “Not on NetApp storage”.

- Storage system names should be unique.

  SnapCenter does not support multiple storage systems with the same name on different clusters. Each storage system that is supported by SnapCenter should have a unique name and a unique management LIF IP address.

Steps

1. Initiate a PowerShell Core connection session by using the Open-SmConnection cmdlet.

   This example opens a PowerShell session:

   ```
   PS C:\> Open-SmConnection
   ```

2. Create a new connection to the storage system by using the Add-SmStorageConnection cmdlet.

   This example creates a new storage system connection:

   ```
   PS C:\> Add-SmStorageConnection -Storage test_vs1 -Protocol Https
   -Timeout 60
   ```

3. Create a new credential by using the Add-SmCredential cmdlet.

   This example creates a new credential named FinanceAdmin with Windows credentials:

   ```
   PS C:> Add-SmCredential -Name FinanceAdmin -AuthMode Windows
   -Credential sddev\administrator
   ```

The information regarding the parameters that can be used with the cmdlet and their descriptions can be obtained by running *Get-Help command_name*. Alternatively, you can also refer to the SnapCenter Software Cmdlet Reference Guide.

Back up individual NetApp supported plug-ins resources

If an individual NetApp supported plug-ins resource is not part of any resource group, you can back up the resource from the Resources page. You can back up the resource on demand, or, if the resource has a policy attached and a schedule configured, then backups occur automatically according to the schedule.

Before you begin

- You must have created a backup policy.
- If you want to back up a resource that has a SnapMirror relationship with a secondary storage, the ONTAP role assigned to the storage user should include the “snapmirror all” privilege. However, if you are using the “vsadmin” role, then the “snapmirror all” privilege is not required.

SnapCenter UI

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page, filter resources from the **View** drop-down list based on resource type.

   Click , and then select the host name and the resource type to filter the resources. You can then click to close the filter pane.
3. Click the resource that you want to back up.
4. In the Resource page, if you want to use a custom name, select the **Use custom name format for Snapshot copy** check box , and then enter a custom name format for the Snapshot name.

   For example, *customtext_policy_hostname* or *resource_hostname*. By default, a timestamp is appended to the Snapshot name.
5. In the Application Settings page, do the following:
   a. Click the **Backups** arrow to set additional backup options:

      Enable consistency group backup, if needed, and perform the following tasks:

For this field…	Do this…
Afford time to wait for Consistency Group Snapshot operation to complete	Select Urgent, Medium, or Relaxed to specify the wait time for Snapshot operation to complete. Urgent = 5 seconds, Medium = 7 seconds, and Relaxed = 20 seconds.
Disable WAFL Sync	Select this to avoid forcing a WAFL consistency point.

b. Click the **Scripts** arrow to run pre and post commands for quiesce, Snapshot, and unquiesce operations. You can also run pre commands before exiting the backup operation.

   Prescripts and postscripts are run in the SnapCenter Server.

c. Click the **Custom Configurations** arrow, and then enter the custom value pairs required for all jobs using this resource.

d. Click the **Snapshot Copy Tool** arrow to select the tool to create Snapshots:

If you want…	Then…
SnapCenter to take a storage level Snapshot	Select **SnapCenter without File System Consistency**.
SnapCenter to use the plug-in for Windows to put the file system into a consistent state and then take a Snapshot	Select **SnapCenter with File System Consistency**.
To enter the command to create a Snapshot	Select **Other**, and then enter the command to create a Snapshot.

6. In the Policies page, perform the following steps:

   a. Select one or more policies from the drop-down list.

   ⓘ You can also create a policy by clicking + .

   In the Configure schedules for selected policies section, the selected policies are listed.

   b. Click + in the Configure Schedules column for the policy for which you want to configure a

schedule.

c. In the Add schedules for policy *policy_name* dialog box, configure the schedule, and then click **OK**.

Where, *policy_name* is the name of the policy that you have selected.

The configured schedules are listed in the Applied Schedules column.

7. In the Notification page, from the **Email preference** drop-down list, select the scenarios in which you want to send the emails.

You must also specify the sender and receiver email addresses, and the subject of the email. SMTP must also be configured in **Settings** > **Global Settings**.

8. Review the summary, and then click **Finish**.

The resources topology page is displayed.

9. Click **Back up Now**.
10. In the Backup page, perform the following steps:

a. If you have applied multiple policies to the resource, from the **Policy** drop-down list, select the policy that you want to use for backup.

If the policy selected for the on-demand backup is associated with a backup schedule, the on-demand backups will be retained based on the retention settings specified for the schedule type.

b. Click **Backup**.

11. Monitor the operation progress by clicking **Monitor** > **Jobs**.

PowerShell cmdlets

Steps

1. Initiate a connection session with the SnapCenter Server for a specified user by using the Open-SmConnection cmdlet.

```
Open-smconnection  -SMSbaseurl
https:\\snapctr.demo.netapp.com:8146\
```

The username and password prompt is displayed.

2. Add resources by using the Add-SmResources cmdlet.

This example adds resources:

```
Add-SmResource -HostName 'scc55.sccore.test.com' -PluginCode
'DummyPlugin' -ResourceName QDBVOL1 -ResourceType Database
-StorageFootPrint (
@{"VolumeName"="qtree_vol1_scc55_sccore_test_com";"QTREENAME"="qtree
Vol1";"StorageSystem"="vserver_scauto_primary"}) -Instance QTREE1
```

3. Create a backup policy by using the Add-SmPolicy cmdlet.

   This example creates a new backup policy:

```
Add-SMPolicy -PolicyName 'test2' -PolicyType 'Backup'
-PluginPolicyType DummyPlugin  -description 'testPolicy'
```

4. Add a new resource group to SnapCenter by using the Add-SmResourceGroup cmdlet.

   This example creates a new resource group with the specified policy and resources:

```
Add-SmResourceGroup -ResourceGroupName
'Verify_Backup_on_Multiple_Qtree_different_vserver_windows'
-Resources
@(@{"Host"="scc55.sccore.test.com";"Uid"="QTREE2";"PluginName"="Dumm
yPlugin"},@{"Host"="scc55.sccore.test.com";"Uid"="QTREE";"PluginName
"="DummyPlugin"}) -Policies test2 -plugincode 'DummyPlugin'
-usesnapcenterwithoutfilesystemconsistency
```

5. Initiate a new backup job by using the New-SmBackup cmdlet.

```
New-SMBackup -DatasetName
Verify_Backup_on_Multiple_Qtree_different_vserver_windows -Policy
test2
```

6. View the status of the backup job by using the Get-SmBackupReport cmdlet.

   This example displays a job summary report of all jobs that were run on the specified date:

```
Get-SmBackupReport -JobId 149

BackedUpObjects           : {QTREE2, QTREE}
FailedObjects             : {}
IsScheduled               : False
HasMetadata               : False
SmBackupId                : 1
SmJobId                   : 149
StartDateTime             : 1/15/2024 1:35:17 AM
EndDateTime               : 1/15/2024 1:36:19 AM
Duration                  : 00:01:02.4265750
CreatedDateTime           : 1/15/2024 1:35:51 AM
Status                    : Completed
ProtectionGroupName       :
Verify_Backup_on_Multiple_Qtree_different_vserver_windows
SmProtectionGroupId       : 1
PolicyName                : test2
SmPolicyId                : 4
BackupName                :
Verify_Backup_on_Multiple_Qtree_different_vserver_windows_scc55_01-
15-2024_01.35.17.4467
VerificationStatus        : NotApplicable
VerificationStatuses      :
SmJobError                :
BackupType                : SCC_BACKUP
CatalogingStatus          : NotApplicable
CatalogingStatuses        :
ReportDataCreatedDateTime :
PluginCode                : SCC
PluginName                : DummyPlugin
PluginDisplayName         : DummyPlugin
JobTypeId                 :
JobHost                   : scc55.sccore.test.com
```

Back up resource groups of NetApp supported plug-in resources

You can back up a resource group on demand from the Resources page. If a resource group has a policy attached and a schedule configured, then backups occur automatically according to the schedule.

Before you begin

- You must have created a resource group with a policy attached.
- If you want to back up a resource that has a SnapMirror relationship to secondary storage, the ONTAP role assigned to the storage user should include the “snapmirror all” privilege. However, if you are using the

"vsadmin" role, then the "snapmirror all" privilege is not required.

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page, select **Resource Group** from the **View** list.

   You can search the resource group either by entering the resource group name in the search box or by clicking and selecting the tag. You can then click to close the filter pane.
3. In the Resource Groups page, select the resource group that you want to back up, and then click **Back up Now**.
4. In the Backup page, perform the following steps:
   a. If you have associated multiple policies with the resource group, from the **Policy** drop-down list, select the policy that you want to use for backup.

      If the policy selected for the on-demand backup is associated with a backup schedule, the on-demand backups will be retained based on the retention settings specified for the schedule type.

   b. Click **Backup**.
5. Monitor the operation progress by clicking **Monitor** > **Jobs**.
   - In MetroCluster configurations, SnapCenter might not be able to detect a protection relationship after a failover.

     Unable to detect SnapMirror or SnapVault relationship after MetroCluster failover
   - If you are backing up application data on VMDKs and the Java heap size for the SnapCenter Plug-in for VMware vSphere is not large enough, the backup might fail. To increase the Java heap size, locate the script file /opt/netapp/init_scripts/scvservice. In that script, the `do_start method` command starts the SnapCenter VMware plug-in service. Update that command to the following: `Java -jar -Xmx8192M -Xms4096M`.

Monitor NetApp supported plug-in resources backup operations

You can monitor the progress of different backup operations by using the SnapCenterJobs page. You might want to check the progress to determine when it is complete or if there is an issue.

About this task

The following icons appear on the Jobs page and indicate the corresponding state of the operations:

- In progress
- Completed successfully
- Failed
- Completed with warnings or could not start due to warnings
- Queued
- Canceled

Steps

1. In the left navigation pane, click **Monitor**.
2. In the Monitor page, click **Jobs**.
3. In the Jobs page, perform the following steps:
   a. Click to filter the list so that only backup operations are listed.
   b. Specify the start and end dates.
   c. From the **Type** drop-down list, select **Backup**.
   d. From the **Status** drop-down, select the backup status.
   e. Click **Apply** to view the operations completed successfully.
4. Select a backup job, and then click **Details** to view the job details.

Though the backup job status displays , when you click on job details you might see that some of the child tasks of the backup operation are still in progress or marked with warning signs.

5. In the Job Details page, click **View logs**.

   The **View logs** button displays the detailed logs for the selected operation.

Cancel backup operations for NetApp supported plug-ins

You can cancel backup operations that are queued.

What you will need

- You must be logged in as the SnapCenter Admin or job owner to cancel operations.
- You can cancel a backup operation from either the **Monitor** page or the **Activity** pane.
- You cannot cancel a running backup operation.
- You can use the SnapCenter GUI, PowerShell cmdlets, or CLI commands to cancel the backup operations.
- The **Cancel Job** button is disabled for operations that cannot be canceled.
- If you selected **All members of this role can see and operate on other members objects** in Users\Groups page while creating a role, you can cancel the queued backup operations of other members while using that role.

Steps

1. Perform one of the following actions:

From the…	Action
Monitor page	a. In the left navigation pane, click **Monitor** > **Jobs**. b. Select the operation, and then click **Cancel Job**.

From the…	Action
Activity pane	a. After initiating the backup operation, click on the Activity pane to view the five most recent operations. b. Select the operation. c. In the Job Details page, click **Cancel Job**.

The operation is canceled, and the resource is reverted to the previous state.

View NetApp supported plug-ins resource related backups and clones in the Topology page

When you are preparing to back up or clone a resource, you might find it helpful to view a graphical representation of all backups and clones on the primary and secondary storage. In the Topology page, you can see all of the backups and clones that are available for the selected resource or resource group. You can view the details of those backups and clones, and then select them to perform data protection operations.

About this task

You can review the following icons in the Manage Copies view to determine whether the backups and clones are available on the primary or secondary storage (Mirror copies or Vault copies).

- displays the number of backups and clones that are available on the primary storage.
- displays the number of backups and clones that are mirrored on the secondary storage using SnapMirror technology.

  Clones of a backup of a version-flexible mirror on a mirror-vault type volume are displayed in the topology view but the mirror backup count in the topology view does not include the version-flexible backup.

- displays the number of backups and clones that are replicated on the secondary storage using SnapVault technology.

  The number of backups displayed includes the backups deleted from the secondary storage. For example, if you have created 6 backups using a policy to retain only 4 backups, the number of backups displayed are 6.

  Clones of a backup of a version-flexible mirror on a mirror-vault type volume are displayed in the topology view but the mirror backup count in the topology view does not include the version-flexible backup.

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page, either select the resource or resource group from the **View** drop-down list.
3. Select the resource either from the resource details view or from the resource group details view.

   If the resource is protected, the topology page of the selected resource is displayed.

4. Review the Summary card to see a summary of the number of backups and clones available on the primary and secondary storage.

   The Summary Card section displays the total number of backups and clones.

   Clicking the refresh button starts a query of the storage to display an accurate count.

   If SnapLock enabled backup is taken, then clicking the **Refresh** button refreshes the primary and secondary SnapLock expiry time retrieved from ONTAP. A weekly schedule also refreshes the primary and secondary SnapLock expiry time retrieved from ONTAP.

   When the application resource is spread across multiple volumes, the SnapLock expiry time for the backup will be the longest SnapLock expiry time that is set for a Snapshot in a volume. The longest SnapLock expiry time is retrieved from ONTAP.

   After on demand backup, by clicking the **Refresh** button refreshes the details of backup or clone.

5. In the Manage Copies view, click **Backups** or **Clones** from the primary or secondary storage to see details of a backup or clone.

   The details of the backups and clones are displayed in a table format.

6. Select the backup from the table, and then click the data protection icons to perform restore, clone, rename, and delete operations.

   You cannot rename or delete backups that are on the secondary storage system.

   You cannot rename the backups that are on the primary storage system.

7. If you want to delete a clone, then select the clone from the table and click to delete the clone.

Restore NetApp supported plug-ins resources

Restore NetApp supported plug-in resources

The restore and recovery workflow includes planning, performing the restore operations, and monitoring the operations.

About this task

The following workflow shows the sequence in which you must perform the restore operation:

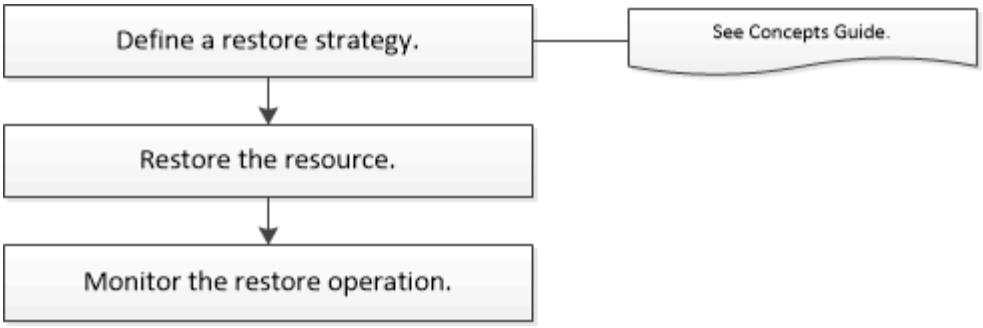

You can also use PowerShell cmdlets manually or in scripts to perform backup, restore, and clone operations. For information about PowerShell cmdlets, use the SnapCenter cmdlet help or see the SnapCenter Software Cmdlet Reference Guide.

Restore a resource backup

You can use SnapCenter to restore resources. The capabilities of the restore operations depends upon the plug-in that you use.

Before you begin

- You must have backed up the resource or resource groups.
- The SnapCenter administrator must have assigned you the storage virtual machines (SVMs) for both the source volumes and destination volumes if you are replicating Snapshots to a mirror or vault.
- You must have cancelled any backup operation that is currently in progress for the resource or resource group you want to restore.

About this task

- The default restore operation only restores storage objects. Restore operations at the application level can only be performed if the NetApp supported plug-in provides that capability.
- For ONTAP 9.12.1 and below version, the clones created from the SnapLock Vault Snapshots as part of restore will inherit the SnapLock Vault expiry time. Storage admin should manually cleanup the clones post the SnapLock expiry time.

SnapCenter UI

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page, filter resources from the **View** drop-down list based on resource type.

   The resources are displayed along with information such as type, host or cluster name, associated resource groups and policies, and status.

   Although a backup might be for a resource group, when you restore, you must select the individual resources you want to restore.

   If the resource is not protected, *Not protected* is displayed in the **Overall Status** column.

   The status *Not protected* in the **Overall Status** column can mean either that the resource is not protected, or that the resource was backed up by a different user.

3. Select the resource or select a resource group and then select a resource in that group.

   The resource topology page is displayed.

4. From the **Manage Copies** view, select **Backups** either from the primary or secondary (mirrored or vaulted) storage systems.
5. In the Primary backup(s) table, select the backup that you want to restore from, and then click .

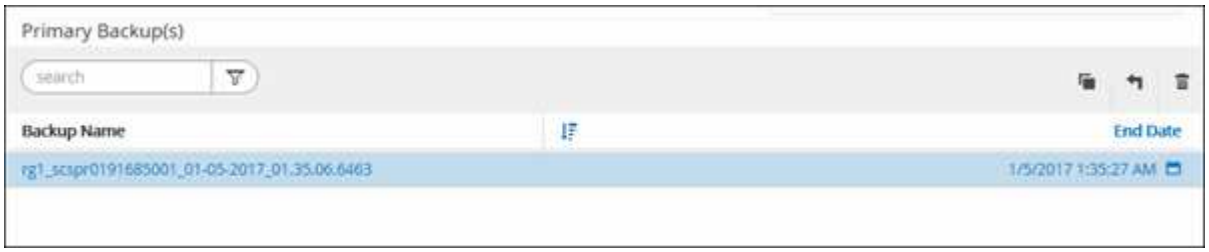

6. In the Restore Scope page, select either **Complete Resource** or **File Level**.
   a. If you selected **Complete Resource**, the resource backup is restored.

      If the resource contains volumes or qtrees as Storage Footprint, then newer Snapshots on such volumes or qtrees are deleted and cannot be recovered. Also, if any other resource is hosted on same volumes or qtrees, then that resource is also deleted.

   b. If you selected **File Level**, then you can either select **All**, or select volumes or qtrees, and then enter the path related to the volumes or qtrees that are selected separated by commas.
      - You can select multiple volumes and qtrees.
      - If resource type is LUN, entire LUN is restored. You can select multiple LUNs.

        NOTE: If you select **All**, all the files on the volumes, qtrees, or LUNs are restored.
7. In the **Pre ops** page, enter pre restore and unmount commands to run before performing a restore job.
8. In the **Post ops** page, enter mount and post restore commands to run after performing a restore job.

9. In the **Notification** page, from the **Email preference** drop-down list, select the scenarios in which you want to send the emails.

   You must also specify the sender and receiver email addresses, and the subject of the email. SMTP must also be configured in the **Settings** > **Global Settings** page.

10. Review the summary, and then click **Finish**.
11. Monitor the operation progress by clicking **Monitor** > **Jobs**.

PowerShell cmdlets

Steps

1. Initiate a connection session with the SnapCenter Server for a specified user by using the Open-SmConnection cmdlet.

```
PS C:\> Open-Smconnection
```

2. Retrieve the information about the one or more backups that you want to restore by using the Get-SmBackup and Get-SmBackupReport cmdlets.

   This example displays information about all available backups:

```
PS C:\> Get-SmBackup

BackupId                      BackupName
BackupTime                    BackupType
--------                      ----------
----------                    ----------
  1               Payroll Dataset_vise-f6_08... 8/4/2015    11:02:32
AM                 Full Backup
  2               Payroll Dataset_vise-f6_08... 8/4/2015    11:23:17
AM
```

   This example displays detailed information about the backup from January 29th 2015 to February 3rd, 2015:

```
PS C:\> Get-SmBackupReport -FromDateTime "1/29/2015" -ToDateTime
"2/3/2015"

   SmBackupId           : 113
   SmJobId              : 2032
   StartDateTime        : 2/2/2015 6:57:03 AM
   EndDateTime          : 2/2/2015 6:57:11 AM
   Duration             : 00:00:07.3060000
   CreatedDateTime      : 2/2/2015 6:57:23 AM
   Status               : Completed
   ProtectionGroupName  : Clone
   SmProtectionGroupId  : 34
   PolicyName           : Vault
   SmPolicyId           : 18
   BackupName           : Clone_SCSPR0019366001_02-02-2015_06.57.08
   VerificationStatus   : NotVerified

   SmBackupId           : 114
   SmJobId              : 2183
   StartDateTime        : 2/2/2015 1:02:41 PM
   EndDateTime          : 2/2/2015 1:02:38 PM
   Duration             : -00:00:03.2300000
   CreatedDateTime      : 2/2/2015 1:02:53 PM
   Status               : Completed
   ProtectionGroupName  : Clone
   SmProtectionGroupId  : 34
   PolicyName           : Vault
   SmPolicyId           : 18
   BackupName           : Clone_SCSPR0019366001_02-02-2015_13.02.45
   VerificationStatus   : NotVerified
```

3. Restore data from the backup by using the Restore-SmBackup cmdlet.

```
Restore-SmBackup -PluginCode 'DummyPlugin' -AppObjectId
'scc54.sccore.test.com\DummyPlugin\NTP\DB1' -BackupId 269
-Confirm:$false
output:
Name                : Restore
'scc54.sccore.test.com\DummyPlugin\NTP\DB1'
Id                  : 2368
StartTime           : 10/4/2016 11:22:02 PM
EndTime             :
IsCancellable       : False
IsRestartable       : False
IsCompleted         : False
IsVisible           : True
IsScheduled         : False
PercentageCompleted : 0
Description         :
Status              : Queued
Owner               :
Error               :
Priority            : None
Tasks               : {}
ParentJobID         : 0
EventId             : 0
JobTypeId           :
ApisJobKey          :
ObjectId            : 0
PluginCode          : NONE
PluginName          :
```

The information regarding the parameters that can be used with the cmdlet and their descriptions can be obtained by running *Get-Help command_name*. Alternatively, you can also refer to the SnapCenter Software Cmdlet Reference Guide.

Monitor NetApp supported plug-in resources restore operations

You can monitor the progress of different SnapCenter restore operations by using the Jobs page. You might want to check the progress of an operation to determine when it is complete or if there is an issue.

About this task

Post-restore states describe the conditions of the resource after a restore operation and any further restore actions that you can take.

The following icons appear on the Jobs page, and indicate the state of the operation:

- In progress
- Completed successfully
- Failed
- Completed with warnings or could not start due to warnings
- Queued
- Canceled

Steps

1. In the left navigation pane, click **Monitor**.
2. In the **Monitor** page, click **Jobs**.
3. In the **Jobs** page, perform the following steps:
   a. Click to filter the list so that only restore operations are listed.
   b. Specify the start and end dates.
   c. From the **Type** drop-down list, select **Restore**.
   d. From the **Status** drop-down list, select the restore status.
   e. Click **Apply** to view the operations that have been completed successfully.
4. Select the restore job, and then click **Details** to view the job details.
5. In the **Job Details** page, click **View logs**.

   The **View logs** button displays the detailed logs for the selected operation.

Clone NetApp supported plug-ins resource backups

Clone NetApp supported plug-ins resource backups

The clone workflow includes performing the clone operation and monitoring the operation.

About this task

You might clone resource backups for the following reasons:

- To test functionality that has to be implemented using the current resource structure and content during application development cycles
- For data extraction and manipulation tools when populating data warehouses
- To recover data that was mistakenly deleted or changed

The following workflow shows the sequence in which you must perform the clone operation:

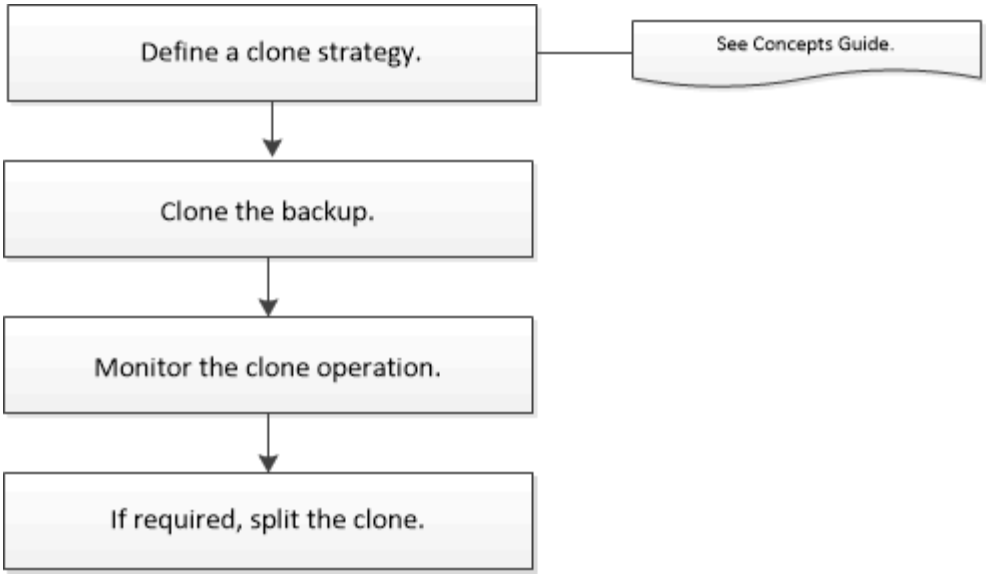

You can also use PowerShell cmdlets manually or in scripts to perform backup, restore, and clone operations. For detailed information about PowerShell cmdlets, use the SnapCenter cmdlet help or see the SnapCenter Software Cmdlet Reference Guide.

Clone from a backup

You can use SnapCenter to clone a backup. You can clone from primary or secondary backup. The capabilities of the clone operations depends upon the plug-in that you use.

Before you begin

- You must have backed up the resources or resource group.
- The default clone operation only clones storage objects. Clone operations at the application level can only be performed if the NetApp supported plug-in provides that capability.
- You should ensure that the aggregates hosting the volumes should be in the assigned aggregates list of the storage virtual machine (SVM).

About this task

For ONTAP 9.12.1 and below version, the clones created from the SnapLock Vault Snapshots as part of restore will inherit the SnapLock Vault expiry time. Storage admin should manually cleanup the clones post the SnapLock expiry time.

SnapCenter UI

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. In the **Resources** page, filter resources from the **View** drop-down list based on resource type.

   The resources are displayed along with information such as type, host or cluster name, associated resource groups and policies, and status.

3. Select the resource or resource group.

   You must select a resource if you select a resource group.

   The resource or resource group topology page is displayed.

4. From the Manage Copies view, select **Backups** either from the primary or secondary (mirrored or vaulted) storage systems.
5. Select the data backup from the table, and then click .
6. In the Locations page, perform the following:

For this field…	Do this…
Clone server	By default, the source host is populated. If you want to specify a different host, select the host on which the clone should be mounted and the plug-in is installed.
Clone suffix	This is mandatory when the clone destination is the same as the source. Enter a suffix that will be appended to the newly cloned resource name. The suffix ensures that the cloned resource is unique on the host. For example, rs1_clone. If you are cloning to the same host as the original resource, you must provide a suffix to differentiate the cloned resource from the original resource; otherwise, the operation fails.

   If the resource selected is a LUN and if you are cloning from a secondary backup, then the destination volumes are listed. Single source can have multiple destination volumes.

7. In the **Settings** page, perform the following:

For this field…	Do this…
Initiator name	Enter the host initiator name, which is either a IQDN or WWPN.

For this field…	Do this…
Igroup protocol	Select Igroup protocol.

Settings page is displayed only if the storage type is LUN.

8. In the Scripts page, enter the commands for pre clone or post clone that should be run before or after the clone operation, respectively. Enter the mount command to mount a file system to a host.

   For example:

   - Pre clone command: delete existing databases with the same name
   - Post clone command: verify a database or start a database.

     Mount command for a volume or qtree on a Linux machine:
     mount<VSERVER_NAME>:%<VOLUME_NAME_Clone /mnt>

9. In the **Notification** page, from the **Email preference** drop-down list, select the scenarios in which you want to send the emails.

   You must also specify the sender and receiver email addresses, and the subject of the email.

10. Review the summary and click **Finish**.
11. Monitor the operation progress by clicking **Monitor** > **Jobs**.

PowerShell cmdlets

Steps

1. Initiate a connection session with the SnapCenter Server for a specified user by using the Open-SmConnection cmdlet.

```
Open-SmConnection  -SMSbaseurl
https:\\snapctr.demo.netapp.com:8146/
```

2. List the backups that can be cloned using the Get-SmBackup or Get-SmResourceGroup cmdlet.

   This example displays information about all available backups:

```
 C:\PS>PS C:\> Get-SmBackup

BackupId                      BackupName
BackupTime                    BackupType
--------                      ----------
----------                    ----------
   1              Payroll Dataset_vise-f6_08... 8/4/2015    11:02:32
AM                    Full Backup
   2              Payroll Dataset_vise-f6_08... 8/4/2015    11:23:17
AM
```

This example displays information about a specified resource group:

```
PS C:\> Get-SmResourceGroup

    Description                 :
    CreationTime                : 10/10/2016 4:45:53 PM
    ModificationTime            : 10/10/2016 4:45:53 PM
    EnableEmail                 : False
    EmailSMTPServer             :
    EmailFrom                   :
    EmailTo                     :
    EmailSubject                :
    EnableSysLog                : False
    ProtectionGroupType         : Backup
    EnableAsupOnFailure         : False
    Policies                    : {}
    HostResourceMaping          : {}
    Configuration               :
SMCoreContracts.SmCloneConfiguration
    LastBackupStatus            : Completed
    VerificationServer          :
    EmailBody                   :
    EmailNotificationPreference : Never
    VerificationServerInfo      :
    SchedulerSQLInstance        :
    CustomText                  :
    CustomSnapshotFormat        :
    SearchResources             : False
    ByPassCredential            : False
    IsCustomSnapshot            :
    MaintenanceStatus           : Production
    PluginProtectionGroupTypes  : {SMSQL}
    Tag                         :
    IsInternal                  : False
    EnableEmailAttachment       : False
    VerificationSettings        : {}
    Name                        : NFS_DB
    Type                        : Group
    Id                          : 2
    Host                        :
    UserName                    :
    Passphrase                  :
    Deleted                     : False
    Auth                        : SMCoreContracts.SmAuth
    IsClone                     : False
    CloneLevel                  : 0
```

```
    Hosts                           :
    StorageName                     :
    ResourceGroupNames              :
    PolicyNames                     :

    Description                     :
    CreationTime                    : 10/10/2016 4:51:36 PM
    ModificationTime                : 10/10/2016 5:27:57 PM
    EnableEmail                     : False
    EmailSMTPServer                 :
    EmailFrom                       :
    EmailTo                         :
    EmailSubject                    :
    EnableSysLog                    : False
    ProtectionGroupType             : Backup
    EnableAsupOnFailure             : False
    Policies                        : {}
    HostResourceMaping              : {}
    Configuration                   :
SMCoreContracts.SmCloneConfiguration
    LastBackupStatus                : Failed
    VerificationServer              :
    EmailBody                       :
    EmailNotificationPreference     : Never
    VerificationServerInfo          :
    SchedulerSQLInstance            :
    CustomText                      :
    CustomSnapshotFormat            :
    SearchResources                 : False
    ByPassRunAs                     : False
    IsCustomSnapshot                :
    MaintenanceStatus               : Production
    PluginProtectionGroupTypes      : {SMSQL}
    Tag                             :
    IsInternal                      : False
    EnableEmailAttachment           : False
    VerificationSettings            : {}
    Name                            : Test
    Type                            : Group
    Id                              : 3
    Host                            :
    UserName                        :
    Passphrase                      :
    Deleted                         : False
    Auth                            : SMCoreContracts.SmAuth
    IsClone                         : False
```

```
    CloneLevel                  : 0
    Hosts                       :
    StorageName                 :
    ResourceGroupNames          :
    PolicyNames                 :
```

3. Initiate a clone operation from a clone resource group or an existing backup using the New-SmClone cmdlet.

   This example creates a clone from a specified backup with all logs:

```
New-SmClone -BackupName
Verify_delete_clone_on_qtree_windows_scc54_10-04-2016_19.05.48.0886
-Resources @{"Host"="scc54.sccore.test.com";"Uid"="QTREE1"}  -
CloneToInstance scc54.sccore.test.com -Suffix '_QtreeCloneWin9'
-AutoAssignMountPoint -AppPluginCode 'DummyPlugin' -initiatorname
'iqn.1991-
05.com.microsoft:scc54.sccore.test.com' -igroupprotocol 'mixed'
```

4. View the status of the clone job by using the Get-SmCloneReport cmdlet.

   This example displays a clone report for the specified job ID:

```
PS C:\> Get-SmCloneReport -JobId 186

    SmCloneId           : 1
    SmJobId             : 186
    StartDateTime       : 8/3/2015 2:43:02 PM
    EndDateTime         : 8/3/2015 2:44:08 PM
    Duration            : 00:01:06.6760000
    Status              : Completed
    ProtectionGroupName : Draper
    SmProtectionGroupId : 4
    PolicyName          : OnDemand_Clone
    SmPolicyId          : 4
    BackupPolicyName    : OnDemand_Full_Log
    SmBackupPolicyId    : 1
    CloneHostName       : SCSPR0054212005.mycompany.com
    CloneHostId         : 4
    CloneName           : Draper__clone__08-03-2015_14.43.53
    SourceResources     : {Don, Betty, Bobby, Sally}
    ClonedResources     : {Don_DRAPER, Betty_DRAPER, Bobby_DRAPER,
Sally_DRAPER}
    SmJobError          :
```

Monitor NetApp supported plug-in resource clone operations

You can monitor the progress of SnapCenter clone operations by using the Jobs page. You might want to check the progress of an operation to determine when it is complete or if there is an issue.

About this task

The following icons appear on the Jobs page, and indicate the state of the operation:

- In progress
- Completed successfully
- Failed
- Completed with warnings or could not start due to warnings
- Queued
- Canceled

Steps

1. In the left navigation pane, click **Monitor**.
2. In the **Monitor** page, click **Jobs**.
3. In the **Jobs** page, perform the following steps:
   a. Click to filter the list so that only clone operations are listed.
   b. Specify the start and end dates.
   c. From the **Type** drop-down list, select **Clone**.
   d. From the **Status** drop-down list, select the clone status.
   e. Click **Apply** to view the operations that are completed successfully.
4. Select the clone job, and then click **Details** to view the job details.
5. In the Job Details page, click **View logs**.

Protect Unix file systems

What you can do with the SnapCenter Plug-in for Unix file systems

When the Plug-in for Unix file systems is installed in your environment, you can use SnapCenter to back up, restore, and clone Unix file systems. You can also perform tasks supporting those operations.

- Discover resources
- Back up Unix file systems
- Schedule backup operations
- Restore file system backups
- Clone file system backups
- Monitor backup, restore, and clone operations

Supported configurations

Item	Supported configuration
Environments	• Physical server • Virtual server vVol datastores on both NFS and SAN. vVol datastore can only be provisioned with ONTAP Tools for VMware vSphere.
Operating systems	• Red Hat Enterprise Linux • Oracle Linux • SUSE Linux Enterprise Server (SLES)
File systems	• SAN: ◦ Both LVM and non LVM based file systems ◦ LVM over VMDK ext3, ext4, and xfs • NFS: NFS v3, NFS v4.x
Protocols	• FC • FCoE • iSCSI • NFS
Multipath	yes

Limitations

- Mix of RDMs and virtual disks in a volume group is not supported.
- File level restore is not supported.

  However, you can manually perform file level restore by cloning the backup and then copying the files manually.

- Mix of file systems spread across VMDKs coming from both NFS and VMFS datastore is not supported.
- NVMe is not supported.
- Provisioning is not supported.

Features

- Enables the Plug-in for Oracle Database to perform data protection operations on Oracle databases by handling the underlying host storage stack on Linux or AIX systems
- Supports Network File System (NFS) and storage area network (SAN) protocols on a storage system that is running ONTAP.
- For Linux systems, Oracle databases on VMDK and RDM LUNs is supported when you deploy the SnapCenter Plug-in for VMware vSphere and register the plug-in with SnapCenter.
- Supports Mount Guard for AIX on SAN filesystems and LVM layout.
- Supports Enhanced Journaled File System (JFS2) with inline logging on SAN filesystems and LVM layout for AIX systems only.

  SAN native devices, filesystems, and LVM layouts built on SAN devices are supported.

- Automates application-aware backup, restore, and clone operations for UNIX file systems in your SnapCenter environment

Install SnapCenter Plug-in for Unix file systems

Prerequisites for adding hosts and installing Plug-ins Package for Linux

Before you add a host and install the plug-ins package for Linux, you must complete all the requirements.

- If you are using iSCSI, the iSCSI service must be running.
- You can either use the password-based authentication for the root or non-root user or SSH key based authentication.

  SnapCenter Plug-in for Unix File Systems can be installed by a non-root user. However, you should configure the sudo privileges for the non-root user to install and start the plug-in process. After installing the plug-in, the processes will be running as an effective non-root user.

- Create credentials with authentication mode as Linux for the install user.
- You must have installed Java 11 on your Linux host.

  Ensure that you have installed only the certified edition of JAVA 11 on the Linux host.

For information to download JAVA, see: Java Downloads for All Operating Systems

- You should have **bash** as the default shell for plug-in installation.

Linux Host requirements

You should ensure that the host meets the requirements before installing the SnapCenter Plug-ins Package for Linux.

Item	Requirements
Operating systems	• Red Hat Enterprise Linux • Oracle Linux • SUSE Linux Enterprise Server (SLES)
Minimum RAM for the SnapCenter plug-in on host	2 GB
Minimum install and log space for the SnapCenter plug-in on host	2 GB You should allocate sufficient disk space and monitor the storage consumption by the logs folder. The log space required varies depending on the number of the entities to be protected and the frequency of data protection operations. If there is no sufficient disk space, the logs will not be created for the recently run operations.
Required software packages	Java 11 Oracle Java and OpenJDK Ensure that you have installed only the certified edition of JAVA 11 on the Linux host. If you have upgraded JAVA to the latest version, you must ensure that the JAVA_HOME option located at /var/opt/snapcenter/spl/etc/spl.properties is set to the correct JAVA version and the correct path.

For the latest information about supported versions, see the NetApp Interoperability Matrix Tool.

Add hosts and install Plug-ins Package for Linux using GUI

You can use the Add Host page to add hosts, and then install the SnapCenter Plug-ins Package for Linux. The plug-ins are automatically installed on the remote hosts.

Steps

1. In the left navigation pane, click **Hosts**.
2. Verify that the **Managed Hosts** tab is selected at the top.
3. Click **Add**.
4. In the Hosts page, perform the following actions:

For this field…	Do this…
Host Type	Select **Linux** as the host type.
Host name	Enter the fully qualified domain name (FQDN) or the IP address of the host. SnapCenter depends on the proper configuration of the DNS. Therefore, the best practice is to enter the FQDN. If you are adding a host by using SnapCenter and the host is part of a subdomain, you must provide the FQDN.
Credentials	Either select the credential name that you created or create new credentials. The credential must have administrative rights on the remote host. For details, see the information about creating credentials. You can view details about the credentials by positioning the cursor over the credential name that you specified. The credentials authentication mode is determined by the host type that you specify in the Add Host wizard.

5. In the Select Plug-ins to Install section, select **Unix File Systems**.
6. (Optional) Click **More Options**.

For this field…	Do this…
Port	Either retain the default port number or specify the port number. The default port number is 8145. If the SnapCenter Server was installed on a custom port, that port number will be displayed as the default port. If you manually installed the plug-ins and specified a custom port, you must specify the same port. Otherwise, the operation fails.
Installation Path	The default path is */opt/NetApp/snapcenter*. You can optionally customize the path. If you use the custom path, ensure that the default content of the sudoers is updated with the custom path.
Skip optional preinstall checks	Select this check box if you have already installed the plug-ins manually and you do not want to validate whether the host meets the requirements for installing the plug-in.

7. Click **Submit**.

   If you have not selected the Skip prechecks checkbox, the host is validated to verify whether the host meets the requirements for installing the plug-in.

   The precheck script does not validate the plug-in port firewall status if it is specified in the firewall reject rules.

   Appropriate error or warning messages are displayed if the minimum requirements are not met. If the error is related to disk space or RAM, you can update the web.config file located at *C:\Program Files\NetApp\SnapCenter WebApp* to modify the default values. If the error is related to other parameters, you should fix the issue.

   In an HA setup, if you are updating web.config file, you must update the file on both nodes.

8. Verify the fingerprint, and then click **Confirm and Submit**.

   SnapCenter does not support ECDSA algorithm.

   Fingerprint verification is mandatory even if the same host was added earlier to SnapCenter and the fingerprint was confirmed.

9. Monitor the installation progress.

   The installation-specific log files are located at */custom_location/snapcenter/logs*.

Result

All the file systems mounted on the host are automatically discovered and displayed under the Resources Page. If nothing is displayed, click **Refresh Resources**.

Monitor installation status

You can monitor the progress of SnapCenter plug-in package installation by using the Jobs page. You might want to check the progress of installation to determine when it is complete or if there is an issue.

About this task

The following icons appear on the Jobs page and indicate the state of the operation:

- In progress
- Completed successfully
- Failed
- Completed with warnings or could not start due to warnings
- Queued

Steps

1. In the left navigation pane, click **Monitor**.
2. In the **Monitor** page, click **Jobs**.
3. In the **Jobs** page, to filter the list so that only plug-in installation operations are listed, do the following:
   a. Click **Filter**.
   b. Optional: Specify the start and end date.
   c. From the Type drop-down menu, select **Plug-in installation**.
   d. From the Status drop-down menu, select the installation status.
   e. Click **Apply**.
4. Select the installation job and click **Details** to view the job details.
5. In the **Job Details** page, click **View logs**.

Configure the SnapCenter Plug-in Loader service

The SnapCenter Plug-in Loader service loads the plug-in package for Linux to interact with the SnapCenter Server. The SnapCenter Plug-in Loader service is installed when you install the SnapCenter Plug-ins Package for Linux.

About this task

After installing the SnapCenter Plug-ins Package for Linux, the SnapCenter Plug-in Loader service starts automatically. If the SnapCenter Plug-in Loader service fails to start automatically, you should:

- Ensure that the directory where the plug-in is operating is not deleted
- Increase the memory space allotted to the Java Virtual Machine

The spl.properties file, which is located at */custom_location/NetApp/snapcenter/spl/etc/*, contains the following parameters. Default values are assigned to these parameters.

Parameter name	Description
LOG_LEVEL	Displays the log levels that are supported. The possible values are TRACE, DEBUG, INFO, WARN, ERROR, and FATAL.
SPL_PROTOCOL	Displays the protocol that is supported by SnapCenter Plug-in Loader. Only the HTTPS protocol is supported. You can add the value if the default value is missing.
SNAPCENTER_SERVER_PROTOCOL	Displays the protocol that is supported by SnapCenter Server. Only the HTTPS protocol is supported. You can add the value if the default value is missing.
SKIP_JAVAHOME_UPDATE	By default, the SPL service detects the java path and update JAVA_HOME parameter. Therefore the default value is set to FALSE. You can set to TRUE if you want to disable the default behavior and manually fix the java path.
SPL_KEYSTORE_PASS	Displays the password of the keystore file. You can change this value only if you change the password or create a new keystore file.
SPL_PORT	Displays the port number on which the SnapCenter Plug-in Loader service is running. You can add the value if the default value is missing. You should not change the value after installing the plug-ins.
SNAPCENTER_SERVER_HOST	Displays the IP address or host name of the SnapCenter Server.
SPL_KEYSTORE_PATH	Displays the absolute path of the keystore file.
SNAPCENTER_SERVER_PORT	Displays the port number on which the SnapCenter Server is running.

Parameter name	Description
LOGS_MAX_COUNT	Displays the number of SnapCenter Plug-in Loader log files that are retained in the */custom_location/snapcenter/spl/logs* folder. The default value is set to 5000. If the count is more than the specified value, then the last 5000 modified files are retained. The check for the number of files is done automatically every 24 hours from when SnapCenter Plug-in Loader service is started. If you manually delete the spl.properties file, then the number of files to be retained is set to 9999.
JAVA_HOME	Displays the absolute directory path of the JAVA_HOME which is used to start SPL service. This path is determined during installation and as part of starting SPL.
LOG_MAX_SIZE	Displays the maximum size of the job log file. Once the maximum size is reached, the log file is zipped, and the logs are written into the new file of that job.
RETAIN_LOGS_OF_LAST_DAYS	Displays the number of days up to which the logs are retained.
ENABLE_CERTIFICATE_VALIDATION	Displays true when CA certificate validation is enabled for the host. You can enable or disable this parameter either by editing the spl.properties or by using the SnapCenter GUI or cmdlet.

If any of these parameters are not assigned to the default value or if you want to assign or change the value, then you can modify the spl.properties file. You can also verify the spl.properties file and edit the file to troubleshoot any issues related to the values that are assigned to the parameters. After you modify the spl.properties file, you should restart the SnapCenter Plug-in Loader service.

Steps

1. Perform one of the following actions, as required:
   - Start the SnapCenter Plug-in Loader service:
     - As a root user, run: `/custom_location/NetApp/snapcenter/spl/bin/spl start`
     - As a non-root user, run: `sudo /custom_location/NetApp/snapcenter/spl/bin/spl start`

- Stop the SnapCenter Plug-in Loader service:
  - As a root user, run: `/custom_location/NetApp/snapcenter/spl/bin/spl stop`
  - As a non-root user, run: `sudo /custom_location/NetApp/snapcenter/spl/bin/spl stop`

    You can use the -force option with the stop command to stop the SnapCenter Plug-in Loader service forcefully. However, you should use caution before doing so because it also terminates the existing operations.

- Restart the SnapCenter Plug-in Loader service:
  - As a root user, run: `/custom_location/NetApp/snapcenter/spl/bin/spl restart`
  - As a non-root user, run: `sudo /custom_location/NetApp/snapcenter/spl/bin/spl restart`
- Find the status of the SnapCenter Plug-in Loader service:
  - As a root user, run: `/custom_location/NetApp/snapcenter/spl/bin/spl status`
  - As a non root user, run: `sudo /custom_location/NetApp/snapcenter/spl/bin/spl status`
- Find the change in the SnapCenter Plug-in Loader service:
  - As a root user, run: `/custom_location/NetApp/snapcenter/spl/bin/spl change`
  - As a non-root user, run: `sudo /custom_location/NetApp/snapcenter/spl/bin/spl change`

Configure CA certificate with SnapCenter Plug-in Loader (SPL) service on Linux host

You should manage the password of SPL keystore and its certificate, configure the CA certificate, configure root or intermediate certificates to SPL trust-store, and configure CA signed key pair to SPL trust-store with SnapCenter Plug-in Loader service to activate the installed digital certificate.

SPL uses the file 'keystore.jks', which is located at '/var/opt/snapcenter/spl/etc' both as its trust-store and key-store.

Manage password for SPL keystore and alias of the CA signed key pair in use

Steps

1. You can retrieve SPL keystore default password from SPL property file.

   It is the value corresponding to the key 'SPL_KEYSTORE_PASS'.

2. Change the keystore password:

```
keytool -storepasswd -keystore keystore.jks
```

3. Change the password for all aliases of private key entries in the keystore to the same password used for the keystore:

```
keytool -keypasswd -alias "<alias_name>" -keystore keystore.jks
```

   Update the same for the key SPL_KEYSTORE_PASS in spl.properties file.

4. Restart the service after changing the password.

> Password for SPL keystore and for all the associated alias password of the private key should be same.

Configure root or intermediate certificates to SPL trust-store

You should configure the root or intermediate certificates without the private key to SPL trust-store.

Steps

1. Navigate to the folder containing the SPL keystore: */var/opt/snapcenter/spl/etc*.
2. Locate the file 'keystore.jks'.
3. List the added certificates in the keystore:

```
keytool -list -v -keystore keystore.jks
```

4. Add a root or intermediate certificate:

```
keytool -import -trustcacerts -alias <AliasNameForCerticateToBeImported>
-file /<CertificatePath> -keystore keystore.jks
```

5. Restart the service after configuring the root or intermediate certificates to SPL trust-store.

> You should add the root CA certificate and then the intermediate CA certificates.

Configure CA signed key pair to SPL trust-store

You should configure the CA signed key pair to the SPL trust-store.

Steps

1. Navigate to the folder containing the SPL's keystore /var/opt/snapcenter/spl/etc.
2. Locate the file 'keystore.jks'.
3. List the added certificates in the keystore:

```
keytool -list -v -keystore keystore.jks
```

4. Add the CA certificate having both private and public key.

```
keytool -importkeystore -srckeystore <CertificatePathToImport>
-srcstoretype pkcs12 -destkeystore keystore.jks -deststoretype JKS
```

5. List the added certificates in the keystore.

```
keytool -list -v -keystore keystore.jks
```

6. Verify that the keystore contains the alias corresponding to the new CA certificate, which was added to the keystore.
7. Change the added private key password for CA certificate to the keystore password.

   Default SPL keystore password is the value of the key SPL_KEYSTORE_PASS in spl.properties file.

```
keytool -keypasswd -alias "<aliasNameOfAddedCertInKeystore>" -keystore
keystore.jks
```

8. If the alias name in the CA certificate is long and contains space or special characters ("*",","), change the alias name to a simple name:

```
keytool -changealias -alias "<OrignalAliasName>" -destalias
"<NewAliasName>" -keystore keystore.jks
```

9. Configure the alias name from the keystore located in spl.properties file.

   Update this value against the key SPL_CERTIFICATE_ALIAS.

10. Restart the service after configuring the CA signed key pair to SPL trust-store.

Configure certificate revocation list (CRL) for SPL

You should configure the CRL for SPL

About this task

- SPL will look for the CRL files in a pre-configured directory.
- Default directory for the CRL files for SPL is */var/opt/snapcenter/spl/etc/crl*.

Steps

1. You can modify and update the default directory in spl.properties file against the key SPL_CRL_PATH.
2. You can place more than one CRL file in this directory.

   The incoming certificates will be verified against each CRL.

Enable CA Certificates for plug-ins

You should configure the CA certificates and deploy the CA certificates in the SnapCenter Server and the corresponding plug-in hosts. You should enable the CA certificate validation for the plug-ins.

Before you begin

- You can enable or disable the CA certificates using the run *Set-SmCertificateSettings* cmdlet.
- You can display the certificate status for the plug-ins using the *Get-SmCertificateSettings*.

The information regarding the parameters that can be used with the cmdlet and their descriptions can be obtained by running *Get-Help command_name*. Alternatively, you can also refer to the SnapCenter Software Cmdlet Reference Guide.

Steps

1. In the left navigation pane, click **Hosts**.
2. In the Hosts page, click **Managed Hosts**.
3. Select single or multiple plug-in hosts.
4. Click **More options**.
5. Select **Enable Certificate Validation**.

After you finish

The Managed Hosts tab host displays a padlock and the color of the padlock indicates the status of the connection between SnapCenter Server and the plug-in host.

- indicates that the CA certificate is neither enabled nor assigned to the plug-in host.
- indicates that the CA certificate is successfully validated.
- indicates that the CA certificate could not be validated.
- indicates that the connection information could not be retrieved.

When the status is yellow or green, the data protection operations completes successfully.

Install SnapCenter Plug-in for VMware vSphere

If your database or filesystem is stored on virtual machines (VMs), or if you want to protect VMs and datastores, you must deploy the SnapCenter Plug-in for VMware vSphere virtual appliance.

For information to deploy, see Deployment Overview.

Deploy CA certificate

To configure the CA Certificate with SnapCenter Plug-in for VMware vSphere, see Create or import SSL certificate.

Configure the CRL file

SnapCenter Plug-in for VMware vSphere looks for the CRL files in a pre-configured directory. Default directory of the CRL files for SnapCenter Plug-in for VMware vSphere is */opt/netapp/config/crl*.

You can place more than one CRL file in this directory. The incoming certificates will be verified against each CRL.

Prepare for protecting Unix file systems

Before performing any data protection operation such as backup, clone, or restore operations, you should set up your environment. You can also set up the SnapCenter Server to use SnapMirror and SnapVault technology.

To take advantage of SnapVault and SnapMirror technology, you must configure and initialize a data protection relationship between the source and destination volumes on the storage device. You can use NetAppSystem Manager or you can use the storage console command line to perform these tasks.

Before you use the Plug-in for Unix file systems, the SnapCenter administrator should install and configure the SnapCenter Server and perform the prerequisite tasks.

- Install and configure SnapCenter Server. Learn more
- Configure the SnapCenter environment by adding storage system connections. Learn more

> SnapCenter does not support multiple SVMs with the same name on different clusters. Each SVM registered with SnapCenter using either SVM registration or cluster registration must be unique.

- Add hosts, install the plug-ins, and discover the resources.
- If you are using SnapCenter Server to protect Unix file systems that reside on VMware RDM LUNs or VMDKs, you must deploy the SnapCenter Plug-in for VMware vSphere and register the plug-in with SnapCenter.
- Install Java on your Linux host.
- Configure SnapMirror and SnapVault on ONTAP, if you want backup replication.

Back up Unix file systems

Discover the UNIX file systems available for backup

After installing the plug-in, all the file systems on that host are automatically discovered and displayed in the Resources page. You can add these file systems to resource groups to perform data protection operations.

Before you begin

- You must have completed tasks such as installing the SnapCenter Server, adding hosts, and creating storage system connections.
- If the file systems reside on a Virtual Machine Disk (VMDK) or raw device mapping (RDM), you must deploy the SnapCenter Plug-in for VMware vSphere and register the plug-in with SnapCenter.

For more information, see Deploy SnapCenter Plug-in for VMware vSphere.

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page, select **Path** from the View list.
3. Click **Refresh Resources**.

   The file systems are displayed along with information such as type, host name, associated resource groups and policies, and status.

Create backup policies for Unix file systems

Before you use SnapCenter to back up Unix file systems, you must create a backup policy for the resource or the resource group that you want to back up. A backup policy is a set of rules that governs how you manage, schedule, and retain backups. You can also specify the replication, script, and backup type settings. Creating a policy saves time when you want to reuse the policy on another resource or resource group.

Before you begin

- You must have prepared for data protection by completing tasks such as installing SnapCenter, adding hosts, discovering the file systems, and creating storage system connections.
- If you are replicating Snapshots to a mirror or vault secondary storage, the SnapCenter administrator must have assigned the SVMs to you for both the source and destination volumes.
- Review the SnapMirror active sync specific prerequisites and limitations. For information, refer Object limits for SnapMirror active sync.

About this task

- SnapLock
  - If 'Retain the backup copies for a specific number of days' option is selected, then the SnapLock retention period must be lesser than or equal to the mentioned retention days.

    Specifying a Snapshot locking period prevents deletion of the Snapshots until the retention period expires. This could lead to retaining a larger number of Snapshots than the count specified in the policy.

    For ONTAP 9.12.1 and below version, the clones created from the SnapLock Vault Snapshots as part of restore will inherit the SnapLock Vault expiry time. Storage admin should manually cleanup the clones post the SnapLock expiry time.

Steps

1. In the left navigation pane, click **Settings**.
2. In the Settings page, click **Policies**.
3. Select **Unix File Systems** from the drop-down list.
4. Click **New**.
5. In the Name page, enter the policy name and details.
6. In the Backup and Replication page, perform the following actions:

a. Specify the backup settings.

b. Specify the schedule frequency by selecting **On demand**, **Hourly**, **Daily**, **Weekly**, or **Monthly**.

c. In the Select secondary replication options section, select one or both of the following secondary replication options:

For this field…	Do this…
Update SnapMirror after creating a local Snapshot copy	Select this field to create mirror copies of the backup sets on another volume (SnapMirror replication). This option should be enabled for SnapMirror active sync.
Update SnapVault after creating a local Snapshot copy	Select this option to perform disk-to-disk backup replication (SnapVault backups).
Error retry count	Enter the maximum number of replication attempts that can be allowed before the operation stops.

7. In the Retention page, specify the retention settings for the backup type and the schedule type selected in the Backup and Replication page:

If you want to…	Then…
Keep a certain number of Snapshots	Select **Copies to keep**, and then specify the number of Snapshots that you want to keep. If the number of Snapshots exceeds the specified number, the Snapshots are deleted with the oldest copies deleted first. The maximum retention value is 1018. Backups will fail if retention is set to a value higher than what the underlying ONTAP version supports. You must set the retention count to 2 or higher if you plan to enable SnapVault replication. If you set the retention count to 1, the retention operation might fail because the first Snapshot is the reference Snapshot for the SnapVault relationship until a newer Snapshot is replicated to the target.
Keep the Snapshots for a certain number of days	Select **Retain copies for**, and then specify the number of days for which you want to keep the Snapshots before deleting them.

Snapshot copy locking period	Select **Snapshot copy locking period** and specify the duration in days, months, or years. Snaplock retention period should be less than 100 years.

8. Select policy label.

   You can assign SnapMirror labels to primary snapshots for remote replication, allowing the primary snapshots to offload the snapshot replication operation from SnapCenter to ONTAP secondary systems. This can be done without enabling SnapMirror or SnapVault option in the policy page.

9. In the Script page, enter the path and the arguments of the prescript or postscript that you want to run before or after the backup operation, respectively.

   You should check if the commands exist in the command list available on the plug-in host from the _ /opt/NetApp/snapcenter/scc/etc/allowed_commands.config_ path.

   You can also specify the script timeout value. The default value is 60 seconds.

10. Review the summary, and then click **Finish**.

Create resource groups and attach policies for Unix file systems

A resource group is a container where you add resources that you want to back up and protect. A resource group allows you to back up all the data that is associated with the file systems.

About this task

- A database with files in ASM disk groups must be in "MOUNT" or "OPEN" state to verify its backups using the Oracle DBVERIFY utility.

  Attach one or more policies to the resource group to define the type of data protection job you want to perform.

  The following image illustrates the relationship between resources, resource groups, and policies for databases:

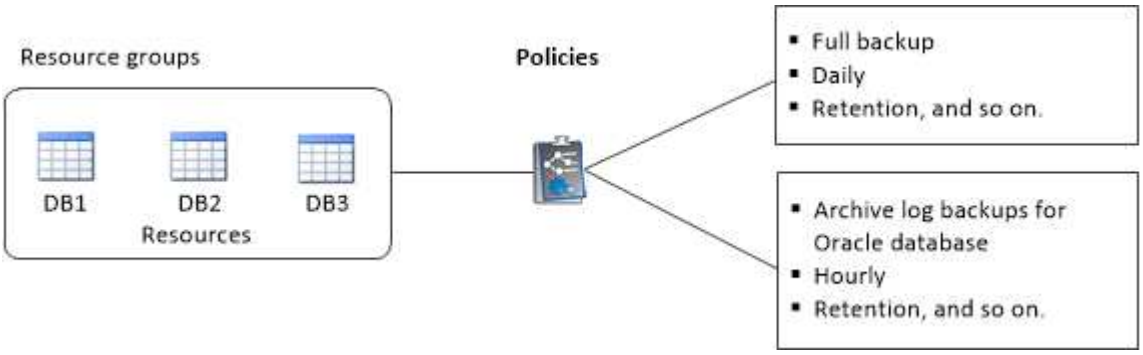

- For SnapLock enabled policies, for ONTAP 9.12.1 and below version, if you specify a Snapshot locking

period, the clones created from the tamper proof Snapshots as part of restore will inherit the SnapLock expiry time. Storage admin should manually cleanup the clones post the SnapLock expiry time.

- Adding new file systems without SnapMirror active sync to an existing resource group which contains resources with SnapMirror active sync, is not supported.
- Adding new file systems to an existing resource group in failover mode of SnapMirror active sync is not supported. You can add resources to the resource group only in regular or fail-back state.

Steps

1. In the left navigation pane, select **Resources**, and the appropriate plug-in from the list.
2. In the Resources page, click **New Resource Group**.
3. In the Name page, perform the following actions:
   a. Enter a name for the resource group in the Name field.

      The resource group name should not exceed 250 characters.

   b. Enter one or more labels in the Tag field to help you search for the resource group later.

      For example, if you add HR as a tag to multiple resource groups, you can later find all resource groups associated with the HR tag.

   c. Select the check box, and enter a custom name format that you want to use for the Snapshot name.

      For example, customtext_resource group_policy_hostname or resource group_hostname. By default, a timestamp is appended to the Snapshot name.
4. In the Resources page, select an Unix file systems host name from the **Host** drop-down list.

   The resources are listed in the Available Resources section only if the resource is discovered successfully. If you have recently added resources, they will appear on the list of available resources only after you refresh your resource list.
5. Select the resources from the Available Resources section and move them to the Selected Resources section.
6. In the Application Settings page, perform the following:
   - Select the Scripts arrow and enter the pre and post commands for quiesce, Snapshot, and unquiesce operations. You can also enter the pre commands to be executed before exiting in the event of a failure.
   - Select one of the backup consistency options:
     - Select **File System Consistent** if you want to ensure that file systems cached data is flushed before creating the backup and no input or output operations are allowed on filesystem while creating the backup.

       For File System Consistent, Consistency group snapshots will be taken for LUNs involved in Volume group.

     - Select **Crash Consistent** if you want to ensure that file systems cached data is flushed before creating the backup.

If you have added different file systems in the resource group, then all volumes from different file systems in the resource group will be put in a Consistency group.

7. In the Policies page, perform the following steps:

   a. Select one or more policies from the drop-down list.

      You can also create a policy by clicking + .

      In the Configure schedules for selected policies section, the selected policies are listed.

   b. Click + in the Configure Schedules column for the policy for which you want to configure a schedule.

   c. In the Add schedules for policy *policy_name* window, configure the schedule, and then click **OK**.

      Where, *policy_name* is the name of the policy that you have selected.

      The configured schedules are listed in the Applied Schedules column.

   Third party backup schedules are not supported when they overlap with SnapCenter backup schedules.

8. In the Notification page, from the **Email preference** drop-down list, select the scenarios in which you want to send the emails.

   You must also specify the sender and receiver email addresses, and the subject of the email. If you want to attach the report of the operation performed on the resource group, select **Attach Job Report**.

   For email notification, you must have specified the SMTP server details using the either the GUI or the PowerShell command Set-SmSmtpServer.

9. Review the summary, and then click **Finish**.

Create resource groups and enable secondary protection for Unix file systems on ASA r2 systems

You should create the resource group to add the resources that are on ASA r2 systems. You can also provision the secondary protection while creating the resource group.

Before you begin

- You should ensure that you are not adding both ONTAP 9.x resources and ASA r2 resources to the same resource group.
- You should ensure that you do not have a database with both ONTAP 9.x resources and ASA r2 resources.

About this task

- The secondary protection is available only if the logged-in user is assigned to the role that has the **SecondaryProtection** capability enabled.
- If you enabled secondary protection, the resource group is put into maintenance mode while creating the primary and secondary consistency groups. After the primary and secondary consistency groups are created, the resource group is put out of maintenance mode.

- SnapCenter does not support secondary protection for a clone resource.

Steps

1. In the left navigation pane, select **Resources**, and the appropriate plug-in from the list.
2. In the Resources page, click **New Resource Group**.
3. In the Name page, perform the following actions:
   a. Enter a name for the resource group in the Name field.

      The resource group name should not exceed 250 characters.

   b. Enter one or more labels in the Tag field to help you search for the resource group later.

      For example, if you add HR as a tag to multiple resource groups, you can later find all resource groups associated with the HR tag.

   c. Select this check box, and enter a custom name format that you want to use for the Snapshot name.

      For example, customtext_resource group_policy_hostname or resource group_hostname. By default, a timestamp is appended to the Snapshot name.

   d. Specify the destinations of the archive log files that you do not want to back up.

      You should use the exact same destination as it was set in the application including prefix, if needed.

4. In the Resources page, select the database host name from the **Host** drop-down list.

   The resources are listed in the Available Resources section only if the resource is discovered successfully. If you have recently added resources, they will appear on the list of available resources only after you refresh your resource list.

5. Select the ASA r2 resources from the Available Resources section and move them to the Selected Resources section.
6. In the Application Settings page, select the backup option.
7. In the Policies page, perform the following steps:
   a. Select one or more policies from the drop-down list.

      You can also create a policy by clicking + .

      In the Configure schedules for selected policies section, the selected policies are listed.

   b. Click + in the Configure Schedules column for the policy for which you want to configure a schedule.
   c. In the Add schedules for policy *policy_name* window, configure the schedule, and then click **OK**.

      Where, *policy_name* is the name of the policy that you have selected.

The configured schedules are listed in the Applied Schedules column.

Third party backup schedules are not supported when they overlap with SnapCenter backup schedules.

8. If the secondary protection is enabled for the policy that you have selected, then Secondary Protection page is displayed and you need to perform the following steps:
   a. Select the type of the replication policy.

      Synchronous replication policy is not supported.

   b. Specify the consistency group suffix that you want to use.
   c. From the Destination Cluster and Destination SVM drop-downs select the peered cluster and SVM that you want to use.

      The cluster and SVM peering is not supported by SnapCenter. You should use System Manager or ONTAP CLIs to perform cluster and SVM peering.

If the resources are already protected outside of SnapCenter, those resources will be displayed in the Secondary Protected Resources section.

1. On the Verification page, perform the following steps:
   a. Click **Load locators** to load the SnapMirror or SnapVault volumes to perform verification on secondary storage.
   b. Click + in the Configure Schedules column to configure the verification schedule for all the schedule types of the policy.
   c. In the Add Verification Schedules policy_name dialog box, perform the following actions:

If you want to…	Do this…
Run verification after backup	Select **Run verification after backup**.
Schedule a verification	Select **Run scheduled verification** and then select the schedule type from the drop-down list.

   d. Select **Verify on secondary location** to verify your backups on secondary storage system.
   e. Click **OK**.

      The configured verification schedules are listed in the Applied Schedules column.

2. In the Notification page, from the **Email preference** drop-down list, select the scenarios in which you want to send the emails.

   You must also specify the sender and receiver email addresses, and the subject of the email. If you want to attach the report of the operation performed on the resource group, select **Attach Job Report**.

   For email notification, you must have specified the SMTP server details using the either the GUI or the PowerShell command Set-SmSmtpServer.

3. Review the summary, and then click **Finish**.

Back up Unix file systems

If a resource is not part of any resource group, you can back up the resource from the Resources page.

Steps

1. In the left navigation pane, select **Resources**, and the appropriate plug-in from the list.
2. In the Resources page, select **Path** from the View list.
3. Click , and then select the host name and the Unix File Systems to filter the resources.
4. Select the file system that you want to back up.
5. In the Resources page, you can perform the following steps:
   a. Select the check box, and enter a custom name format that you want to use for the Snapshot name.

      For example, `customtext_policy_hostname` or `resource_hostname`. A timestamp is appended to the Snapshot name by default.
6. In the Application Settings page, perform the following:
   - Select the Scripts arrow and enter the pre and post commands for quiesce, Snapshot, and unquiesce operations. You can also enter the pre commands to be executed before exiting in the event of a failure.
   - Select one of the backup consistency options:
     - Select **File System Consistent** if you want to ensure that file systems cached data is flushed before creating the backup and no operations are performed on filesystem while creating the backup.
     - Select **Crash Consistent** if you want to ensure that file systems cached data is flushed before creating the backup.
7. In the Policies page, perform the following steps:
   a. Select one or more policies from the drop-down list.

      You can create a policy by clicking .

      In the Configure schedules for selected policies section, the selected policies are listed.

   b. Click in the Configure Schedules column to configure a schedule for the policy you want.
   c. In the Add schedules for policy *policy_name* window, configure the schedule, and then select `OK`.

      policy_name is the name of the policy that you have selected.

      The configured schedules are listed in the Applied Schedules column.
8. In the Notification page, select the scenarios in which you want to send the emails from the **Email preference** drop-down list.

   You must specify the sender and receiver email addresses, and the subject of the email. If you want to

attach the report of the backup operation performed on the resource, select **Attach Job Report**.

> For email notification, you must have specified the SMTP server details using the either the GUI or the PowerShell command `Set-SmSmtpServer`.

9. Review the summary, and then click **Finish**.

   The topology page is displayed.

10. Click **Back up Now**.
11. In the Backup page, perform the following steps:
    a. If you have applied multiple policies to the resource, from the Policy drop-down list, select the policy that you want to use for backup.

       If the policy selected for the on-demand backup is associated with a backup schedule, the on-demand backups will be retained based on the retention settings specified for the schedule type.

    b. Click **Backup**.
12. Monitor the operation progress by clicking **Monitor** > **Jobs**.

Back up Unix file systems resource groups

You can back up the Unix file systems defined in the resource group. You can back up a resource group on demand from the Resources page. If a resource group has a policy attached and a schedule configured, then backups are created according to the schedule.

Steps

1. In the left navigation pane, select **Resources**, and the appropriate plug-in from the list.
2. In the Resources page, select **Resource Group** from the **View** list.
3. Enter the resource group name in the search box, or click , and select the tag.

   Click to close the filter pane.

4. In the Resource Group page, select the resource group to back up.
5. In the Backup page, perform the following steps:
    a. If you have multiple policies associated with the resource group, select the backup policy you want to use from the **Policy** drop-down list.

       If the policy selected for the on-demand backup is associated with a backup schedule, the on-demand backups will be retained based on the retention settings specified for the schedule type.

    b. Select **Backup**.
6. Monitor the progress by selecting **Monitor > Jobs**.

Monitor Unix file systems backup

Learn how to monitor the progress of backup operations and data protection operations.

Monitor Unix file systems backup operations

You can monitor the progress of different backup operations by using the SnapCenterJobs page. You might want to check the progress to determine when it is complete or if there is an issue.

About this task

The following icons appear on the Jobs page and indicate the corresponding state of the operations:

- In progress
- Completed successfully
- Failed
- Completed with warnings or could not start due to warnings
- Queued
- Canceled

Steps

1. In the left navigation pane, click **Monitor**.
2. In the Monitor page, click **Jobs**.
3. In the Jobs page, perform the following steps:
   a. Click to filter the list so that only backup operations are listed.
   b. Specify the start and end dates.
   c. From the **Type** drop-down list, select **Backup**.
   d. From the **Status** drop-down, select the backup status.
   e. Click **Apply** to view the operations completed successfully.
4. Select a backup job, and then click **Details** to view the job details.

   Though the backup job status displays , when you click on job details you might see that some of the child tasks of the backup operation are still in progress or marked with warning signs.

5. In the Job Details page, click **View logs**.

   The **View logs** button displays the detailed logs for the selected operation.

Monitor data protection operations in the Activity pane

The Activity pane displays the five most recent operations performed. The Activity pane also displays when the operation was initiated and the status of the operation.

The Activity pane displays information regarding backup, restore, clone, and scheduled backup operations.

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. Click on the Activity pane to view the five most recent operations.

When you click one of the operations, the operation details are listed in the **Job Details** page.

View protected Unix file systems in the Topology page

When you are preparing to back up, restore, or clone a resource, you might find it helpful to view a graphical representation of all backups, restored file systems, and clones on the primary and secondary storage.

About this task

In the Topology page, you can see all the backups, restored file systems, and clones that are available for the selected resource or resource group. You can view the details of those backups, restored file systems, and clones, and then select them to perform data protection operations.

You can review the following icons in the Manage Copies view to determine whether the backups and clones are available on the primary or secondary storage (Mirror copies or Vault copies).

- displays the number of backups and clones that are available on the primary storage.
- displays the number of backups and clones that are mirrored on the secondary storage using SnapMirror technology.
- displays the number of backups and clones that are replicated on the secondary storage using SnapVault technology.

  The number of backups displayed includes the backups deleted from the secondary storage. For example, if you created 6 backups using a policy to retain only 4 backups, the number of backups displayed is 6.

Clones of a backup of a version-flexible mirror on a mirror-vault type volume are displayed in the topology view, but the mirror backup count in the topology view does not include the version-flexible backup.

If you have secondary relationship as SnapMirror active sync (initially released as SnapMirror Business Continuity [SM-BC]), you can see following additional icons:

- The replica site is up.
- The replica site is down.
- The secondary mirror or vault relationship has not been re-established.

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page, either select the resource or resource group from the **View** drop-down list.

3. Select the resource either from the resource details view or from the resource group details view.

   If the resource is protected, the Topology page of the selected resource is displayed.

4. Review the Summary card to see a summary of the number of backups and clones available on the primary and secondary storage.

   The Summary Card section displays the total number of backups and clones.

   Clicking the **Refresh** button starts a query of the storage to display an accurate count.

   If SnapLock enabled backup is taken, then clicking the **Refresh** button refreshes the primary and secondary SnapLock expiry time retrieved from ONTAP. A weekly schedule also refreshes the primary and secondary SnapLock expiry time retrieved from ONTAP.

   When the file system is spread across multiple volumes, the SnapLock expiry time for the backup will be the longest SnapLock expiry time that is set for a Snapshot in a volume. The longest SnapLock expiry time is retrieved from ONTAP.

   For SnapMirror active sync, clicking the **Refresh** button refreshes the SnapCenter backup inventory by querying ONTAP for both primary and replica sites. A weekly schedule also performs this activity for all databases containing SnapMirror active sync relationship.

   - For SnapMirror active sync and only for ONTAP 9.14.1, Async Mirror or Async MirrorVault relationships to the new primary destination should be manually configured after failover. From ONTAP 9.15.1 onwards Async Mirror or Async MirrorVault is auto configured to the new primary destination.
   - After failover, a backup should be created for SnapCenter to be aware of the failover. You can click **Refresh** only after a backup has been created.

5. In the Manage Copies view, click **Backups** or **Clones** from the primary or secondary storage to see details of a backup or clone.

   The details of the backups and clones are displayed in a table format.

6. Select the backup from the table, and then click the data protection icons to perform restore, clone, and delete operations.

   You cannot rename or delete backups that are on the secondary storage.

7. If you want to delete a clone, select the clone from the table, and then click .

Example showing backups and clones on the primary storage

Restore and recover Unix file systems

Restore Unix file systems

In the event of data loss, you can use SnapCenter to restore Unix file systems.

About this task

- You should execute the following commands to establish the connection with the SnapCenter Server, list the backups and retrieve its information, and restore the backup.

  The information regarding the parameters that can be used with the command and their descriptions can be obtained by running Get-Help *command_name*. Alternatively, you can also refer to the SnapCenter Software Command Reference Guide.

- For SnapMirror active sync restore operation, you must select the backup from the primary location.

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page, select either **Path** or **Resource Group** from the **View** list.
3. Select the file system from either the details view or the resource group details view.

   The topology page is displayed.

4. From the Manage Copies view, select **Backups** from either the primary or the secondary (mirrored or replicated) storage systems.
5. Select the backup from the table, and then click .
6. In the Restore Scope page:
   - For NFS file systems, by default **Connect and Copy** restore is selected. You can also select **Volume Revert** or **Fast Restore**.
   - For non NFS file systems, the restore scope is selected depending on the layout.

     The new files created after backup may not be available after restore depending on the file system type and layout.
7. In the PreOps page, enter pre restore commands to run before performing a restore job.
8. In the PostOps page, enter post restore commands to run after performing a restore job.

   You should check if the commands exist in the command list available on the plug-in host at location */opt/NetApp/snapcenter/scc/etc/allowed_commands.config* path.
9. In the Notification page, from the **Email preference** drop-down list, select the scenarios in which you want to send the email notifications.

   You must also specify the sender and receiver email addresses, and the subject of the email. If you want to attach the report of the restore operation performed, you must select **Attach Job Report**.

   For email notification, you must have specified the SMTP server details by using the either the GUI or the PowerShell command Set-SmSmtpServer.
10. Review the summary, and then click **Finish**.

    If restore operation fails, rollback is not supported.

    In case of restore of a filesystem residing on volume group, the old contents on the filesystem are not deleted. Only the content from the cloned filesystem will be copied to the source filesystem. This is applicable when there are multiple filesystems on the volume group and default NFS filesystem restores.
11. Monitor the operation progress by clicking **Monitor** > **Jobs**.

Monitor Unix file systems restore operations

You can monitor the progress of different SnapCenter restore operations by using the Jobs page. You might want to check the progress of an operation to determine when it is complete or if there is an issue.

About this task

Post-restore states describe the conditions of the resource after a restore operation and any further restore actions that you can take.

The following icons appear on the Jobs page, and indicate the state of the operation:

- In progress
- Completed successfully
- Failed
- Completed with warnings or could not start due to warnings
- Queued
- Canceled

Steps

1. In the left navigation pane, click **Monitor**.
2. In the **Monitor** page, click **Jobs**.
3. In the **Jobs** page, perform the following steps:
   a. Click to filter the list so that only restore operations are listed.
   b. Specify the start and end dates.
   c. From the **Type** drop-down list, select **Restore**.
   d. From the **Status** drop-down list, select the restore status.
   e. Click **Apply** to view the operations that have been completed successfully.
4. Select the restore job, and then click **Details** to view the job details.
5. In the **Job Details** page, click **View logs**.

   The **View logs** button displays the detailed logs for the selected operation.

Clone Unix file systems

Clone Unix file system backup

You can use SnapCenter to clone Unix file system using the backup of the filesystem.

Before you begin

- You can skip the fstab file update by setting the value of *SKIP_FSTAB_UPDATE* to **true** in the *agent.properties* file located at */opt/NetApp/snapcenter/scc/etc*.
- You can have a static clone volume name and junction path by setting the value of *USE_CUSTOM_CLONE_VOLUME_NAME_FORMAT* to **true** in the *agent.properties* file located at */opt/NetApp/snapcenter/scc/etc*. After updating the file, you should restart the SnapCenter Plug-in creator service by running the command: `/opt/NetApp/snapcenter/scc/bin/scc restart`.

  Example: Without this property the clone volume name and junction path will be like <Source_volume_name>_Clone_<Timestamp> but now it will be <Source_volume_name>_Clone_<Clone_Name>

  This keeps the name constant so that you can manually keep the fstab file updated if you do not prefer to update the fstab by SnapCenter.

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page, select either **Path** or **Resource Group** from the **View** list.
3. Select the file system from either the details view or the resource group details view.

   The topology page is displayed.

4. From the Manage Copies view, select the backups either from Local copies (primary), Mirror copies (secondary), or Vault copies (secondary).
5. Select the backup from the table, and then click .
6. In the Location page, perform the following actions:

For this field…	Do this…
Clone server	By default, the source host is populated.
Clone mount point	Specify the path where the file system will be mounted.

7. In the Scripts page, perform the following steps:
   a. Enter the commands for pre clone or post clone that should be run before or after the clone operation, respectively.

      You should check if the commands exist in the command list available on the plug-in host from the */opt/NetApp/snapcenter/scc/etc/allowed_commands.config* path.

8. In the Notification page, from the **Email preference** drop-down list, select the scenarios in which you want to send the emails.

   You must also specify the sender and receiver email addresses, and the subject of the email. If you want to attach the report of the clone operation performed, select **Attach Job Report**.

   For email notification, you must have specified the SMTP server details using the either the GUI or the PowerShell command Set-SmSmtpServer.

9. Review the summary, and then click **Finish**.
10. Monitor the operation progress by clicking **Monitor** > **Jobs**.

Split a clone

You can use SnapCenter to split a cloned resource from the parent resource. The clone that is split becomes independent of the parent resource.

About this task

- You cannot perform the clone split operation on an intermediate clone.

  For example, after you create clone1 from a database backup, you can create a backup of clone1, and then clone this backup (clone2). After you create clone2, clone1 is an intermediate clone, and you cannot perform the clone split operation on clone1. However, you can perform the clone split operation on clone2.

After splitting clone2, you can perform the clone split operation on clone1 because clone1 is no longer the intermediate clone.

- When you split a clone, the backup copies and clone jobs of the clone are deleted.
- For information about FlexClone volume split operations, see, Split a FlexClone volume from its parent volume.
- Ensure that the volume or aggregate on the storage system is online.

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. In the **Resources** page, select the appropriate option from the View list:

Option	Description
For database applications	Select **Database** from the View list.
For file systems	Select **Path** from the View list.

3. Select the appropriate resource from the list.

   The resource topology page is displayed.

4. From the **Manage Copies** view, select the cloned resource (for example, the database or LUN), and then click .
5. Review the estimated size of the clone that is to be split and the required space available on the aggregate, and then click **Start**.
6. Monitor the operation progress by clicking **Monitor** > **Jobs**.

   The clone split operation stops responding if the SMCore service restarts. You should run the Stop-SmJob cmdlet to stop the clone split operation, and then retry the clone split operation.

   If you want a longer poll time or shorter poll time to check whether the clone is split or not, you can change the value of *CloneSplitStatusCheckPollTime* parameter in *SMCoreServiceHost.exe.config* file to set the time interval for SMCore to poll for the status of the clone split operation. The value is in milliseconds and the default value is 5 minutes.

   For example:

```
<add key="CloneSplitStatusCheckPollTime" value="300000" />
```

   The clone split start operation fails if backup, restore, or another clone split is in progress. You should restart the clone split operation only after the running operations are complete.

Related information

SnapCenter clone or verification fails with aggregate does not exist

Monitor Unix file systems clone operations

You can monitor the progress of SnapCenter clone operations by using the Jobs page. You might want to check the progress of an operation to determine when it is complete or if there is an issue.

About this task

The following icons appear on the Jobs page, and indicate the state of the operation:

- In progress
- Completed successfully
- Failed
- Completed with warnings or could not start due to warnings
- Queued
- Canceled

Steps

1. In the left navigation pane, click **Monitor**.
2. In the **Monitor** page, click **Jobs**.
3. In the **Jobs** page, perform the following steps:
   a. Click to filter the list so that only clone operations are listed.
   b. Specify the start and end dates.
   c. From the **Type** drop-down list, select **Clone**.
   d. From the **Status** drop-down list, select the clone status.
   e. Click **Apply** to view the operations that are completed successfully.
4. Select the clone job, and then click **Details** to view the job details.
5. In the Job Details page, click **View logs**.

Protect applications running on Azure NetApp Files

Protect applications running on Azure NetApp Files

SnapCenter supports protecting your applications such as Oracle, SQL, and SAP HANA residing on Azure NetApp Files. Beginning with 6.0.1 release, SnapCenter supports the Azure NetApp Files backup feature that expands the data protection capabilities of Azure NetApp Files by providing fully managed backup solution for long-term recovery, archive, and compliance.

Azure NetApp Files is a premium storage solution that can be expensive for long-term backup retention. To optimize costs, you can move the backups from the Azure NetApp Files storage to an Azure object store. Beginning with SnapCenter 6.0.1, you can back up and clone applications residing on Azure NetApp Files to Azure Blob Storage (object store). You can maintain two copies of your data, volume snapshot copies on Azure NetApp Files storage for short term recovery and another copy on Azure Blob Storage for long term recovery.

When a policy with Azure NetApp Files backup is enabled and associated with an resource, SnapCenter handles the creation of volume snapshots and backing it to Azure Blob Storage. SnapCenter creates the Backup Vault and enable backup for the volume. If you have enabled backup for the volume, SnapCenter utilizes the existing vault.

Limitations

- Object storage functionalities for FAS, ASA, or AFF ONTAP and Amazon FSx for NetApp ONTAP storage systems is not supported.
- Oracle and SAP HANA mount and catalog workflows are not supported for object storage backups but are supported for snapshots.
- Oracle PDB clones are not supported for object storage backups but are supported for snapshots.
- Backup verification from object storage, REST API support, clone lifecycle management from object storage, and reporting features for object storage backups are not supported.
- Restoring from backups on Azure Blob Storage to Azure NetApp Files is not supported. You can use the clone option, alternatively.
- Clone split is not supported.

Install SnapCenter and create credentials

Install SnapCenter on Azure Virtual Machine

You can download the SnapCenter software from the NetApp Support site and install the software on the Azure virtual machine.

Before you begin

- Ensure that the Azure Windows virtual machine meets the requirements for SnapCenter Server installation. For information, see Requirements to install SnapCenter Server.
- If you are new to Azure NetApp Files and do not have an existing NetApp account, ensure that you have

registered so that you can access the SnapCenter Software. For inforamtion, see Register to access the SnapCenter software.

Steps

1. Download the SnapCenter Server installation package from NetApp Support Site.
2. Initiate the SnapCenter Server installation by double-clicking the downloaded .exe file.

   After you initiate the installation, all the pre-checks are performed and if the minimum requirements are not met appropriate error or warning messages are displayed. You can ignore the warning messages and proceed with installation; however, errors should be fixed.

3. Review the pre-populated values required for the SnapCenter Server installation and modify if required.

   You do not have to specify the password for MySQL Server repository database. During SnapCenter Server installation the password is auto generated.

   The special character “%” is not supported in the custom path for the repository database. If you include “%” in the path, installation fails.

4. Click **Install Now**.

   If you have specified any values that are invalid, appropriate error messages will be displayed. You should re-enter the values, and then initiate the installation.

   If you click the **Cancel** button, the step that is being executed will be completed, and then start the rollback operation. The SnapCenter Server will be completely removed from the host.

   However, if you click **Cancel** when "SnapCenter Server site restart" or "Waiting for SnapCenter Server to start" operations are being performed, installation will proceed without cancelling the operation.

Register the product to enable support

If you are new to NetApp and do not have an existing NetApp account, you should register the product to enable support.

Steps

1. After installing SnapCenter, navigate to **Help > About**.
2. In the *About SnapCenter* dialog box, make a note of the SnapCenter Instance, a 20 digit number that starts with 971.
3. Click https://register.netapp.com.
4. Click **I am not a registered NetApp Customer**.
5. Specify your details to register yourself.
6. Leave the NetApp Reference SN field blank.
7. Select **SnapCenter** from the Product Line drop-down.
8. Select the billing provider.
9. Enter the 20-digit SnapCenter instance ID.
10. Click **Submit**.

Create the Azure credential in SnapCenter

You should create the Azure credential in SnapCenter to access the Azure NetApp account.

Before you begin

- Ensure that you have created the service principal in Azure.
- Ensure that you have the tenant ID, client ID, and secret key associated with the service principal available with you.
  - The Tenant ID can be found in the Azure Portal on the Overview page for Entra ID.
  - The Client ID is also known as the Application ID for the Enterprise App/Service Principal. This can be found in the Azure Portal on the Overview page for the Enterprise Application you've created to act as the Service Principal for SnapCenter.
  - The Client Secret Key is also known as the Secret Value. You can create this client secret in the Azure Portal by navigating to **App registrations** under Entra ID. After selecting the Enterprise App that you created, navigate to **Certificates & secrets** and then **New client secret**.

    The secret value can only be accessed when created. You cannot access it later.

  - The Service Principal needs to be granted permissions. The role of Contributor will allow the service principal to perform required actions in Azure. This can be granted from the Access Control (IAM) page within the Subscription page.

Steps

1. In the left navigation pane, click **Settings**.
2. In the Settings page, click **Credential**.
3. Click **New**.
4. In the Credential page, specify the following information required to create the credential.

For this field…	Do this…
Credential Name	Enter a name for the credential.
Authentication Mode	Select **Azure Credential** from the drop-down list.
Tenant ID	Enter the tenant ID.
Client ID	Enter the client ID.
Client Secret Key	Enter the client secret key.

5. Click **OK**.

Configure the Azure storage account

You should configure the Azure storage account in SnapCenter.

The Azure storage account contains details about the subscription ID, Azure credential, and Azure NetApp account.

The standard licenses and capacity-based license is not required for Azure NetApp Files.

Steps

1. In the left navigation pane, click **Storage Systems**.
2. In the Storage Systems page, select **Azure NetApp Files** and click **New**.
3. Select the credential, subscription ID, and NetApp account from the respective drop-down lists.
4. Click **Submit**.

Create the credential to add the plug-in host

SnapCenter uses credentials to authenticate users for SnapCenter operations.

You should create credentials for installing SnapCenter plug-ins and additional credentials for performing data protection operations.

Steps

1. In the left navigation pane, click **Settings**.
2. In the Settings page, click **Credential**.
3. Click **New**.
4. In the Credential page, specify the following information required to create the credential.

For this field…	Do this…
Credential Name	Enter a name for the credential.
Authentication Mode	Select the authentication mode from the drop-down list.
Authentication Type	Select either **Password Based** or **SSH Key Based** (only for Linux host).
Username	Specify the username.
Password	If you have selected Password based authentication, specify the password.
SSH Private Key	If you have selected SSH Key Based authentication, specify the private key.
Use sudo privileges	Select the Use sudo privileges check box if you are creating credentials for a non-root user. This is applicable only for Linux users.

5. Click **OK**.

Protect SAP HANA databases

Add hosts and install SnapCenter plug-in for SAP HANA database

You must use the SnapCenter Add Host page to add hosts, and then install the plug-ins packages. The plug-ins are automatically installed on the remote hosts.

Before you begin

- You must be a user that is assigned to a role that has the plug-in install and uninstall permissions, such as the SnapCenter Admin role.
- When installing a plug-in on a Windows host, if you specify a credential that is not built-in, or if the user belongs to a local workgroup user, you must disable UAC on the host.
- If you are installing on the centralized host, ensure that the SAP HANA client software is installed on that host and open the required ports on the SAP HANA database host to run the HDB SQL queries remotely.

Steps

1. In the left navigation pane, click **Hosts**.
2. Verify that the **Managed Hosts** tab is selected.
3. Click **Add**.
4. In the Hosts page, perform the following actions:
   a. In the Host Type field, select the host type.
   b. In the Host name field, enter the fully qualified domain name (FQDN) or the IP address of the host.
   c. In the Credentials field, enter the credential that you created.
5. In the Select Plug-ins to Install section, select the plug-ins to install.
6. (Optional) Click **More Options** and specify the details.
7. Click **Submit**.
8. If host type is Linux, verify the fingerprint, and then click **Confirm and Submit**.

   In a cluster setup, you should verify the fingerprint of each of the nodes in the cluster.

9. Monitor the installation progress.

Add SAP HANA database

You should add the SAP HANA database manually.

About this task

Resources need to be added manually if the plug-in is installed on a centralized server. If the SAP HANA plug-in is installed on the HANA database host, then the HANA system is discovered automatically.

Automatic discovery is not supported for HANA multi-host configuration, they must be added through centralized plug-in only.

Steps

1. In the left navigation pane, select the SnapCenter Plug-in for SAP HANA Database from the drop-down list, and then click **Resources**.
2. In the Resources page, click **Add SAP HANA Database**.
3. In the Provide Resource Details page, perform the following actions:
   a. Enter the resource type either as Single Container, Multitenant Database Container, or Non-data Volume.
   b. Enter the SAP HANA system name.
   c. Enter the system ID (SID).
   d. Select the plug-in host.
   e. Enter the key to connect to the SAP HANA system.
   f. Enter the username for whom the HDB Secure User Store Key is configured.
4. In the Provide Storage Footprint page, select **Azure NetApp Files** as the storage type.
   a. Select the Azure NetApp account.
   b. Select the capacity pool and the associated volumes.
   c. Click **Save**.
5. Review the summary, and then click **Finish**.

Create backup policies for SAP HANA databases

Before you use SnapCenter to back up SAP HANA database resources, you must create a backup policy for the resource or resource group that you want to back up.

Steps

1. In the left navigation pane, click **Settings**.
2. In the Settings page, click **Policies**.
3. Click **New**.
4. In the Name page, enter the policy name and description.
5. In the Policy type page, perform the following steps:
   a. Select **Azure NetApp Files** as the storage type.
   b. Select **File-Based** if you want to perform an integrity check of the database.
   c. Select **Snapshot Based** if you want to create a backup using Snapshot technology.
6. In the Snapshot and backup page, perform the following steps:
   a. Select the frequency of the scheduled backup.
   b. Specify the retention settings.
   c. If you want to enable Azure NetApp Files backup, select **Enable backup** and specify the retention settings.
7. Review the summary and click **Finish**.

Create resource groups and attach SAP HANA backup policies

A resource group is the container to which you must add resources that you want to back

up and protect.

A resource group enables you to back up all the data that is associated with a given application simultaneously. A resource group is required for any data protection job. You must also attach one or more policies to the resource group to define the type of data protection job that you want to perform.

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page, click **New Resource Group**.
3. In the Name page, perform the following actions:

For this field…	Do this…
Name	Enter a name for the resource group.
Tags	Enter one or more labels that will help you later search for the resource group.
Use custom name format for Snapshot copy	Select this check box, and enter a custom name format that you want to use for the Snapshot name.

4. In the Resources page, select a host name from the **Host** drop-down list and resource type from the **Resource Type** drop-down list.
5. Select the resources from the **Available Resources** section, and then click the right arrow to move them to the **Selected Resources** section.
6. In the Policies page, perform the following steps:
   a. Select one or more policies from the drop-down list.
   b. In the Configure Schedules column, click for the policy you want to configure.
   c. In the Add schedules for policy *policy_name* dialog box, configure the schedule, and then click **OK**.
7. In the Notification page, from the **Email preference** drop-down list, select the scenarios in which you want to send the emails.
8. Review the summary, and then click **Finish**.

Back up SAP HANA databases running on Azure NetApp Files

If a resource is not yet part of any resource group, you can back up the resource from the Resources page.

Steps

1. In the left navigation pane, select **Resources**, and then select the appropriate plug-in from the list.
2. In the Resource page, filter resources from the **View** drop-down list based on resource type.
3. Select the resource that you want to back up.
4. In the Resource page, select **Use custom name format for Snapshot copy**, and then enter a custom name format that you want to use for the Snapshot name.
5. In the Application Settings page, do the following:

a. Select the **Backups** arrow to set additional backup options.

b. Select the **Scripts** arrow to run pre and post commands for quiesce, Snapshot, and unquiesce operations.

c. Select the **Custom Configurations** arrow, and then enter the custom value pairs required for all jobs using this resource.

d. Select the **Snapshot Copy Tool > SnapCenter without File System Consistency** to create Snapshots.

The **File System Consistency** option is applicable only for applications running on Windows hosts.

6. In the Policies page, perform the following steps:

a. Select one or more policies from the drop-down list.

b. Select + in the Configure Schedules column for the policy for which you want to configure a schedule.

c. In the Add schedules for policy *policy_name* dialog box, configure the schedule, and then select **OK**.

policy_name is the name of the policy that you selected.

7. In the Notification page, from the **Email preference** drop-down list, select the scenarios in which you want to send the emails.

You must also specify the sender and receiver email addresses, and the subject of the email. SMTP must also be configured in **Settings** > **Global Settings**.

8. Review the summary, and then select **Finish**.
9. Select **Back up Now**.
10. In the Backup page, perform the following steps:

a. If multiple policies are associated with the resource, from the **Policy** drop-down list, select the policy that you want to use for backup.

If the policy selected for the on-demand backup is associated with a backup schedule, the on-demand backups will be retained based on the retention settings specified for the schedule type.

11. Select **Backup**.
12. Monitor the operation progress by clicking **Monitor** > **Jobs**.

Back up SAP HANA resource groups

A resource group is a collection of resources on a host. A backup operation on the resource group is performed on all resources defined in the resource group.

Steps

1. In the left navigation pane, select **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page, select **Resource Group** from the **View** list.
3. In the Resource Groups page, select the resource group that you want to back up, and then select **Back up Now**.
4. In the Backup page, perform the following steps:

a. If multiple policies are associated with the resource group, from the **Policy** drop-down list, select the policy that you want to use for backup.

   If the policy selected for the on-demand backup is associated with a backup schedule, the on-demand backups will be retained based on the retention settings specified for the schedule type.

b. Select **Backup**.

5. Monitor the operation progress by selecting **Monitor** > **Jobs**.

Restore and recover SAP HANA databases

You can restore and recover data from the backups.

About this task

For Auto discovered HANA systems, if the **Complete Resource** option is selected, then restore is performed using Single File snapshot restore technology. If the **Fast Restore** check box is selected, then Volume Revert technology is used.

For manually added resources, Volume Revert technology is always used.

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page, filter resources from the **View** drop-down list based on resource type.
3. Select the resource or select a resource group and then select a resource in that group.
4. From the Manage Copies view, select **Backups** either from the primary or secondary (mirrored or vaulted) storage systems.
5. In the Primary backup(s) table, select the backup that you want to restore from, and then click .
6. In the Restore Scope page, select **Complete Resource**.

   All the configured data volumes of the SAP HANA database are restored.

7. For Auto discovered HANA systems, in the Recovery scope page, perform the following actions:
   a. Select **Recover to most recent state** if you want to recover as close as possible to the current time.
   b. Select **Recover to point in time** if you want to recover to the specified point in time.
   c. Select **Recover to specified data backup** if you want to recover to a specific data backup.
   d. Select **No recovery** if you do not want to recover now.
   e. Specify the log backup locations.
   f. Specify the backup catalog location.
8. In the Pre ops page, enter pre restore and unmount commands to run before performing a restore job.
9. In the Post ops page, enter mount and post restore commands to run after performing a restore job.
10. In the Notification page, from the **Email preference** drop-down list, select the scenarios in which you want to send the emails.

    You must also specify the sender and receiver email addresses and the subject of the email. SMTP must also be configured on the **Settings** > **Global Settings** page.

11. Review the summary, and then click **Finish**.
12. Monitor the operation progress by clicking **Monitor** > **Jobs**.

Clone SAP HANA database backup

You can use SnapCenter to clone a SAP HANA database using the backup of the database. The clones created are thick clones and are created on the parent capacity pool.

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page, filter resources from the **View** drop-down list based on resource type.
3. Select the resource or resource group.
4. From the Manage Copies view, select **Backups** from the primary storage system.
5. Select the data backup from the table, and then click .
6. In the Location page, perform the following actions:
   a. Select the host that has the SAP HANA plug-in installed for managing the cloned HANA system.

      It can be a centralized plug-in host or HANA system host.

      If the HANA plug-in is installed on a centralized host that manages HANA databases on other hosts, while creating or deleting clones, SnapCenter intentionally skips host side operations (mount or unmount filesystem) as the target server is a centralized host. You should use custom pre or post clone scripts to perform mount and unmount operations.

   b. Enter the SAP HANA SID to clone from the existing backups.
   c. Enter IP addresses or the host names on which the cloned volumes will be exported.
   d. If the SAP HANA database ANF volumes are configured in a manual QOS capacity pool, specify the QOS for the cloned volumes.

      If QOS for the cloned volumes is not specified, the QOS of the source volume will be used. If the automatic QOS capacity pool is used, the QOS value specified will be ignored.

7. In the Scripts page, perform the following steps:
   a. Enter the commands for pre clone or post clone that should be run before or after the clone operation, respectively.
   b. Enter the mount command to mount a file system to a host.

      If the source HANA system is auto discovered and the clone target host plug-in is installed on the SAP HANA host, then SnapCenter automatically unmounts the existing HANA data volumes on the clone target host and mounts the newly cloned HANA data volumes.

8. In the Notification page, from the **Email preference** drop-down list, select the scenarios in which you want to send the emails.
9. Review the summary, and then click **Finish**.
10. Monitor the operation progress by clicking **Monitor** > **Jobs**.

Clone Split is disabled for ANF clones because ANF clone is already an independent volume created from the selected Snapshot.

Protect Microsoft SQL Server databases

Add hosts and install SnapCenter plug-in for SQL Server database

SnapCenter supports data protection of SQL instances on SMB shares on Azure NetApp Files. The standalone and availability group (AG) configurations are supported.

You must use the SnapCenter Add Host page to add hosts, and then install the plug-ins package. The plug-ins are automatically installed on the remote hosts.

Before you begin

- You must be a user that is assigned to a role that has the plug-in install and uninstall permissions, such as the SnapCenter Admin role.
- When installing a plug-in on a Windows host, if you specify a credential that is not built-in, or if the user belongs to a local workgroup user, you must disable UAC on the host.

Steps

1. In the left navigation pane, select **Hosts**.
2. Verify that the **Managed Hosts** tab is selected at the top.
3. Select **Add**.
4. In the Hosts page do the following:
   a. In the Host Type field, select the host type.
   b. In the Host name field, enter the fully qualified domain name (FQDN) or the IP address of the host.
   c. In the Credentials field, enter the credential that you created.
5. In the **Select Plug-ins to Install** section, select the plug-ins to install.
6. (Optional) Click **More Options** and specify the details.
7. Select **Submit**.
8. Select **Configure log directory** and in the Configure host log directory page, enter the SMB path of the host log directory, and click **Save**.
9. Click **Submit** and monitor the installation progress.

Create backup policies for SQL Server databases

You can create a backup policy for the resource or the resource group before you use SnapCenter to back up SQL Server resources, or you can create a backup policy at the time you create a resource group or backup a single resource.

Steps

1. In the left navigation pane, click **Settings**.
2. In the Settings page, click **Policies**.
3. Click **New**.

4. In the Name page, enter the policy name and description.
5. In the Policy type page, perform the following steps:
    a. Select **Azure NetApp Files** as the storage type.
    b. Select the backup type.
        i. Select **Full Backup and Log Backup** if you want to back up database files and transaction logs.
        ii. Select **Full Backup** if you want to back up only the database files.
        iii. Select **Log Backup** if you want to back up only the transaction logs.
        iv. Select **Copy Only Backup** if you want to back up your resources by using another application.
    c. In the Availability Group Settings section, perform the following actions:
        i. Select Backup on preferred backup replica if you want to back up only on the replica.
        ii. Select primary AG replica or the secondary AG replica for the backup.
        iii. Select the backup priority.
6. In the Snapshot and backup page, perform the following steps:
    a. Select the frequency of the scheduled backup.
    b. Specify the retention settings depending on the backup type selected.
    c. If you want to enable Azure NetApp Files backup, select **Enable backup** and specify the retention settings.
7. In the Verification page, perform the following steps:
    a. In the Run verification for following backup schedules section, select the schedule frequency.
    b. In the Database consistency check options section, perform the following actions:
        i. Select **Limit the integrity structure to physical structure of the database (PHYSICAL_ONLY)** to limit the integrity check to the physical structure of the database and to detect torn pages, checksum failures, and common hardware failures that impact the database.
        ii. Select **Suppress all information messages (NO_INFOMSGS)** to suppress all informational messages.

            Selected by default.

        iii. Select **Display all reported error messages per object (ALL_ERRORMSGS)** to display all the reported errors per object.
        iv. Select **Do not check nonclustered indexes (NOINDEX)** if you do not want to check nonclustered indexes.

            The SQL Server database uses Microsoft SQL Server Database Consistency Checker (DBCC) to check the logical and physical integrity of the objects in the database.

        v. Select **Limit the checks and obtain the locks instead of using an internal database Snapshot copy (TABLOCK)** to limit the checks and obtain locks instead of using an internal database Snapshot.
    c. In the **Log Backup** section, select **Verify log backup upon completion** to verify the log backup upon completion.
    d. In the **Verification script settings** section, enter the path and the arguments of the prescript or postscript that should be run before or after the verification operation, respectively.

8. Review the summary and click **Finish**.

Create resource groups and attach SQL backup policies

A resource group is the container to which you must add resources that you want to back up and protect.

A resource group enables you to back up all the data that is associated with a given application simultaneously. A resource group is required for any data protection job. You must also attach one or more policies to the resource group to define the type of data protection job that you want to perform.

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page, click **New Resource Group**.
3. In the Name page, perform the following actions:

For this field…	Do this…
Name	Enter a name for the resource group.
Tags	Enter one or more labels that will help you later search for the resource group.
Use custom name format for Snapshot copy	Select this check box, and enter a custom name format that you want to use for the Snapshot name.

4. In the Resources page, select a host name from the **Host** drop-down list and resource type from the **Resource Type** drop-down list.
5. Select the resources from the **Available Resources** section, and then click the right arrow to move them to the **Selected Resources** section.
6. In the Policies page, perform the following steps:
   a. Select one or more policies from the drop-down list.
   b. In the Configure Schedules column, click + for the policy you want to configure.
   c. In the Add schedules for policy *policy_name* dialog box, configure the schedule, and then click **OK**.
   d. Select the Microsoft SQL Server scheduler.
7. In the Verification page, perform the following steps:
   a. Select the verification server.
   b. Select the policy for which you want to configure your verification schedule, and then click + .
   c. Either select **Run verification after backup** or **Run scheduled verification**.
   d. Click **OK**.
8. In the Notification page, from the **Email preference** drop-down list, select the scenarios in which you want to send the emails.
9. Review the summary, and then click **Finish**.

Back up SQL Server databases running on Azure NetApp Files

If a resource is not yet part of any resource group, you can back up the resource from the Resources page.

Before you begin

You should create a load balancer, if the Azure Windows Failover Cluster does not have a cluster IP assigned or if it is not reachable from SnapCenter. The load balancer's IP should be configured and reachable from the SnapCenter Server.

Steps

1. In the left navigation pane, select **Resources**, and then select the appropriate plug-in from the list.
2. In the Resource page, select **Database**, **Instance**, or **Availability Group** from the View drop-down list.
3. In the Resource page, select **Use custom name format for Snapshot copy**, and then enter a custom name format that you want to use for the Snapshot name.
4. In the Policies page, perform the following steps:
   a. Select one or more policies from the drop-down list.
   b. Select in the Configure Schedules column for the policy for which you want to configure a schedule.
   c. In the Add schedules for policy *policy_name* dialog box, configure the schedule, and then select **OK**.

      policy_name is the name of the policy that you selected.

   d. Select **Use Microsoft SQL Server scheduler**, and then select the scheduler instance from the **Scheduler Instance** drop-down list that is associated with the scheduling policy.
5. In the Verification page, perform the following steps:
   a. Select the verification server.
   b. Select the policy for which you want to configure your verification schedule, and then click .
   c. Either select **Run verification after backup** or **Run scheduled verification**.
   d. Click OK.
6. In the Notification page, from the **Email preference** drop-down list, select the scenarios in which you want to send the emails.
7. Review the summary, and then click **Finish**.
8. Select **Back up Now**.
9. In the Backup page, perform the following steps:
   a. If multiple policies are associated with the resource, from the **Policy** drop-down list, select the policy that you want to use for backup.
   b. Select **Verify after backup**.
   c. Select **Backup**.
10. Monitor the operation progress by clicking **Monitor** > **Jobs**.

Back up SQL Server resource groups

You can back up the resource groups that consist of multiple resources. A backup operation on the resource group is performed on all resources defined in the resource group.

Steps

1. In the left navigation pane, select **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page, select **Resource Group** from the **View** list.
3. In the Resource Groups page, select the resource group that you want to back up, and then select **Back up Now**.
4. In the Backup page, perform the following steps:
   a. If multiple policies are associated with the resource group, from the **Policy** drop-down list, select the policy that you want to use for backup.
   b. After backup, select **Verify** to verify the on-demand backup.
   c. Select **Backup**.
5. Monitor the operation progress by selecting **Monitor** > **Jobs**.

Restore and recover SQL Server databases

You can use SnapCenter to restore backed-up SQL Server databases. Database restoration is a multiphase process that copies all the data and log pages from a specified SQL Server backup to a specified database.

About this task

You should ensure that the target instance for restore is configured with an active directory user who belongs to the SMB ADactive directory domain and has permissions to set the file permissions appropriately. You should configure the credentials in SnapCenter at instance level.

The SQL authentication for target instance will not be supported for SMB configurations. The target instance should be configured in SnapCenter with the active directory user having the required permissions.

If the SnapCenter Plug-in services service account is not an active directory user then while performing restore to alternate host, the user who have the full control over the source volumes is required so that it can be impersonated and perform the required operation.

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page, select either **Database** or **Resource Group** from the View list.
3. Select the database or the resource group from the list.
4. From the Manage Copies view, select **Backups** from storage system.
5. Select the backup from the table, and then click the icon.
6. In the Restore Scope page, select one of the following options:
   a. Select **Restore the database to the same host where the backup was created** if you want to restore the database to the same SQL server where the backups are taken.

b. Select **Restore the database to an alternate host** if you want the database to be restored to a different SQL server in the same or different host where backups are taken.

7. In the Recovery Scope page, select one of the following options:
   a. Select **None** when you need to restore only the full backup without any logs.
   b. Select **All log backups** up-to-the-minute backup restore operation to restore all the available log backups after the full backup.
   c. Select **By log backups** to perform a point-in-time restore operation, which restores the database based on backup logs until the backup log with the selected date.
   d. Select **By specific date until** to specify the date and time after which transaction logs are not applied to the restored database.
   e. If you have selected **All log backups**, **By log backups**, or **By specific date until** and the logs are located at a custom location, select **Use custom log directory**, and then specify the log location.
8. In the Pre-Ops and Post Ops page, specify the required details.
9. In the Notification page, from the **Email preference** drop-down list, select the scenarios in which you want to send the emails.
10. Review the summary, and then click **Finish**.
11. Monitor the restore process by using the **Monitor** > **Jobs** page.

Clone SQL Server database backup

You can use SnapCenter to clone a SQL database using the backup of the database. The clones created are thick clones and are created on the parent capacity pool.

About this task

You should ensure that the target instance for clone is configured with an active directory user who belongs to the SMB ADactive directory domain and has permissions to set the file permissions appropriately. You should configure the credentials in SnapCenter at instance level.

The SQL authentication for target instance will not be supported for SMB configurations. The target instance should be configured in SnapCenter with the active directory user having the required permissions.

If the SnapCenter Plug-in services service account is not an active directory user then while performing clone, the user who have the full control over the source volumes is required so that it can be impersonated and perform the required operation.

Steps

1. In the left navigation pane, select **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page, select either **Database** or **Resource Group** from the **View** list.
3. Select the database or resource group.
4. From the **Manage Copies** view page, select the backup from primary storage system.
5. Select the backup, and then select .
6. In the **Clone Options** page, provide all the required details.
7. In the Location page, select a storage location to create a clone.

   If the SQL Server database ANF volumes are configured in a manual QOS capacity pool, specify the QOS

for the cloned volumes.

If QOS for the cloned volumes is not specified, the QOS of the source volume will be used. If the automatic QOS capacity pool is used, the QOS value specified will be ignored.

8. In the Logs page, select one of the following options:
   a. Select **None** if you want to clone only the full back up without any logs.
   b. Select **All log backups** if you want to clone all the available log backups dated after the full backup.
   c. Select **By log backups until** if you want to clone the database based on the backup logs that were created up to the backup log with the selected date.
   d. Select **By specific date until** if you do not want to apply the transaction logs after the specified date and time.
9. In the **Script** page, enter the script timeout, path, and the arguments of the prescript or postscript that should be run before or after the clone operation, respectively.
10. In the **Notification** page, from the **Email preference** drop-down list, select the scenarios in which you want to send the emails.
11. Review the summary, and then select **Finish**.
12. Monitor the operation progress by selecting **Monitor** > **Jobs**.

Perform Clone Lifecycle

Using SnapCenter, you can create clones from a resource group or database. You can either perform on-demand clone or you can schedule recurring clone operations of a resource group or database. If you clone a backup periodically, you can use the clone to develop applications, populate data, or recover data.

Steps

1. In the left navigation pane, select **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page, select either **Database** or **Resource Group** from the **View** list.
3. Select the database or resource group.
4. From the **Manage Copies** view page, select the backup from primary storage system.
5. Select the backup, and then select .
6. In the **Clone Options** page, provide all the required details.
7. In the Location page, select a storage location to create a clone.

   If the SQL Server database ANF volumes are configured in a manual QOS capacity pool, specify the QOS for the cloned volumes.

   If QOS for the cloned volumes is not specified, the QOS of the source volume will be used. If the automatic QOS capacity pool is used, the QOS value specified will be ignored.

8. In the **Script** page, enter the script timeout, path, and the arguments of the prescript or postscript that should be run before or after the clone operation, respectively.
9. In the Schedule page, perform one of the following actions:
   - Select **Run now** if you want to execute the clone job immediately.
   - Select **Configure schedule** when you want to determine how frequently the clone operation should occur, when the clone schedule should start, on which day the clone operation should occur, when the

schedule should expire, and whether the clones must be deleted after the schedule expires.

10. In the **Notification** page, from the **Email preference** drop-down list, select the scenarios in which you want to send the emails.
11. Review the summary, and then select **Finish**.
12. Monitor the operation progress by selecting **Monitor** > **Jobs**.

Protect Oracle databases

Add hosts and install SnapCenter plug-in for Oracle database

You can use the Add Host page to add hosts, and then install the SnapCenter Plug-ins Package for Linux or SnapCenter Plug-ins Package for AIX. The plug-ins are automatically installed on the remote hosts.

You can add a host and install plug-in packages either for an individual host or for a cluster. If you are installing the plug-in on a cluster (Oracle RAC), the plug-in is installed on all the nodes of the cluster. For Oracle RAC One Node, you should install the plug-in on both active and passive nodes.

Steps

1. In the left navigation pane, click **Hosts**.
2. Verify that the **Managed Hosts** tab is selected.
3. Click **Add**.
4. In the Hosts page, perform the following actions:
   a. In the Host Type field, select the host type.
   b. In the Host name field, enter the fully qualified domain name (FQDN) or the IP address of the host.
   c. In the Credentials field, enter the credential that you created.
5. In the Select Plug-ins to Install section, select the plug-ins to install.
6. (Optional) Click **More Options** and specify the details.
7. Click **Submit**.
8. Verify the fingerprint, and then click **Confirm and Submit**.

   In a cluster setup, you should verify the fingerprint of each of the nodes in the cluster.

9. Monitor the installation progress.

Create backup policies for Oracle databases

Before you use SnapCenter to back up Oracle database resources, you must create a backup policy for the resource or the resource group that you want to back up.

Steps

1. In the left navigation pane, click **Settings**.
2. In the Settings page, click **Policies**.
3. Select Oracle Database from the drop-down list.

4. Click **New**.
5. In the Name page, enter the policy name and description.
6. In the Policy type page, perform the following steps:
   a. Select **Azure NetApp Files** as the storage type.
   b. Select the backup type as either online or offline backup.
   c. If you want to catalog backup using Oracle Recovery Manager (RMAN), select **Catalog backup with Oracle Recovery Manager (RMAN)**.
   d. If you want to prune archive logs after backup, select **Prune archive logs after backup**.
   e. Specify the delete archive log settings.
7. In the Snapshot and backup page, perform the following steps:
   a. Select the frequency of the scheduled backup.
   b. Specify the retention settings.
   c. If you want to enable Azure NetApp Files backup, select **Enable backup** and specify the retention settings.
8. In the Script page, enter the path and the arguments of the prescript or postscript that you want to run before or after the backup operation, respectively.
9. In the Verification page, select the backup schedule for which you want to perform the verification operation and enter the path and the arguments of the prescript or postscript that you want to run before or after the verification operation, respectively.
10. Review the summary and click **Finish**.

Create resource groups and attach Oracle backup policies

A resource group is the container to which you must add resources that you want to back up and protect.

A resource group enables you to back up all the data that is associated with a given application simultaneously. A resource group is required for any data protection job. You must also attach one or more policies to the resource group to define the type of data protection job that you want to perform.

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page, click **New Resource Group**.
3. In the Name page, perform the following actions:

For this field…	Do this…
Name	Enter a name for the resource group.
Tags	Enter one or more labels that will help you later search for the resource group.
Use custom name format for Snapshot copy	Select this check box, and enter a custom name format that you want to use for the Snapshot name.

For this field…	Do this…
Archive log file destination	Specify the destinations of the archive log files.

4. In the Resources page, select a host name from the **Host** drop-down list and resource type from the **Resource Type** drop-down list.
5. Select the resources from the **Available Resources** section, and then click the right arrow to move them to the **Selected Resources** section.
6. In the Policies page, perform the following steps:
   a. Select one or more policies from the drop-down list.
   b. In the Configure Schedules column, click + for the policy you want to configure.
   c. In the Add schedules for policy *policy_name* dialog box, configure the schedule, and then click **OK**.
7. In the Verification page, perform the following steps:
   a. Select the verification server.
   b. Select the policy for which you want to configure your verification schedule, and then click * + .
   c. Either select **Run verification after backup** or **Run scheduled verification**.
   d. Click **OK**.
8. In the Notification page, from the **Email preference** drop-down list, select the scenarios in which you want to send the emails.
9. Review the summary, and then click **Finish**.

Back up Oracle databases running on Azure NetApp Files

If a resource is not yet part of any resource group, you can back up the resource from the Resources page.

Steps

1. In the left navigation pane, select **Resources**, and then select the appropriate plug-in from the list.
2. In the Resource page, select **Database** from the View drop-down list.
3. In the Resource page, select **Use custom name format for Snapshot copy**, and then enter a custom name format that you want to use for the Snapshot name.
4. In the Policies page, perform the following steps:
   a. Select one or more policies from the drop-down list.
   b. Select + in the Configure Schedules column for the policy for which you want to configure a schedule.
   c. In the Add schedules for policy *policy_name* dialog box, configure the schedule, and then select **OK**.
5. In the Verification page, perform the following steps:
   a. Select the verification server.
   b. Select the policy for which you want to configure your verification schedule, and then click + .

c. Either select **Run verification after backup** or **Run scheduled verification**.

d. Click OK.

6. In the Notification page, from the **Email preference** drop-down list, select the scenarios in which you want to send the emails.
7. Review the summary, and then click **Finish**.
8. Select **Back up Now**.
9. In the Backup page, perform the following steps:

   a. If multiple policies are associated with the resource, from the **Policy** drop-down list, select the policy that you want to use for backup.

   b. Click **Backup**.
10. Monitor the operation progress by clicking **Monitor** > **Jobs**.

Back up Oracle resource groups

You can back up the resource groups that consist of multiple resources. A backup operation on the resource group is performed on all resources defined in the resource group.

Steps

1. In the left navigation pane, select **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page, select **Resource Group** from the **View** list.
3. In the Resource Groups page, select the resource group that you want to back up, and then select **Back up Now**.
4. In the Backup page, perform the following steps:

   a. If multiple policies are associated with the resource group, from the **Policy** drop-down list, select the policy that you want to use for backup.

   b. Select **Backup**.
5. Monitor the operation progress by selecting **Monitor** > **Jobs**.

Restore and recover Oracle databases

In the event of data loss, you can use SnapCenter to restore data from one or more backups to your active file system and then recover the database.

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page, select either **Database** or **Resource Group** from the View list.
3. Select the database or the resource group from the list.
4. From the Manage Copies view, select **Backups** from the primary storage system.
5. Select the backup from the table, and then click .
6. In the Restore Scope page, perform the following tasks:

   a. Select RAC if you have selected a backup of a database in RAC environment.

b. Perform the following actions:

i. Select **All Datafiles** if you want to restore only the database files.

ii. Select **Tablespaces** if you want to restore only the tablespaces.

iii. Select **Redo log files** if you want to restore the redo log files of the Data Guard standby or Active Data Guard standby databases.

iv. Select **Pluggable databases** and specify the PDBs you want to restore.

v. Select **Pluggable database (PDB) tablespaces**, and then specify the PDB and the tablespaces of that PDB that you want to restore.

vi. Select **Restore the database to the same host where the backup was created** if you want to restore the database to the same SQL server where the backups are taken.

vii. Select **Restore the database to an alternate host** if you want the database to be restored to a different SQL server in the same or different host where backups are taken.

viii. Select **Change database state if needed for restore and recovery** to change the state of the database to the state required to perform restore and recovery operations.

ix. Select **Force in place restore** if you want to perform in-place restore in the scenarios where new datafiles are added after backup or when LUNs are added, deleted, or re-created to an LVM disk group.

7. In the Recovery Scope page, select one of the following options:

a. Select **All Logs** if you want to recover to the last transaction.

b. Select **Until SCN (System Change Number)** if you want to recover to a specific SCN.

c. Select **Date and Time** if you want to recover to a specific date and time.

d. Select **No recovery** if you do not want to recover.

e. Select **Specify external archive log locations** if you want to specify the location of the external archive log files.

8. In the Pre-Ops and Post Ops page, specify the required details.

9. In the Notification page, from the **Email preference** drop-down list, select the scenarios in which you want to send the emails.

10. Review the summary, and then click **Finish**.

11. Monitor the operation progress by clicking **Monitor** > **Jobs**.

Restore and recover tablespaces using point-in-time recovery

You can restore a subset of tablespaces that have been corrupted or dropped without impacting the other tablespaces in the database. SnapCenter uses RMAN to perform point-in-time recovery (PITR) of the tablespaces.

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page, select either **Database** or **Resource Group** from the View list.
3. Select the database of type single instance (multitenant).
4. From the Manage Copies view, select **Backups** from the storage system.

   If the backup is not cataloged, you should select the backup and click **Catalog**.

5. Select the catalogued backup, and then click .
6. In the Restore Scope page, perform the following tasks:
   a. Select **RAC** if you have selected a backup of a database in RAC environment.
   b. Select **Tablespaces** if you want to restore only the tablespaces.
   c. Select **Change database state if needed for restore and recovery** to change the state of the database to the state required to perform restore and recovery operations.
7. In the Recovery Scope page, select one of the following options:
   a. Select **Until SCN (System Change Number)** if you want to recover to a specific SCN.
   b. Select **Date and Time** if you want to recover to a specific date and time.
8. In the Pre-Ops and Post Ops page, specify the required details.
9. In the Notification page, from the **Email preference** drop-down list, select the scenarios in which you want to send the emails.
10. Review the summary, and then click **Finish**.
11. Monitor the restore process by using the **Monitor** > **Jobs** page.

Restore and recover pluggable database using point-in-time recovery

You can restore and recover a pluggable database (PDB) that has been corrupted or dropped without impacting the other PDBs in the container database (CDB). SnapCenter uses RMAN to perform point-in-time recovery (PITR) of the PDB.

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page, select either **Database** or **Resource Group** from the View list.
3. Select the database of type single instance (multitenant).
4. From the Manage Copies view, select **Backups** from the storage system.

   If the backup is not cataloged, you should select the backup and click **Catalog**.

5. Select the catalogued backup, and then click .
6. In the Restore Scope page, perform the following tasks:
   a. Select **RAC** if you have selected a backup of a database in RAC environment.
   b. Depending on whether you want to restore the PDB or tablespaces in a PDB, perform one of the actions:
      - Select **Pluggable databases (PDBs)** if you want to restore a PDB.
      - Select **Pluggable database (PDB) tablespaces** if you want to restore tablespaces in a PDB.
7. In the Recovery Scope page, select one of the following options:
   a. Select **Until SCN (System Change Number)** if you want to recover to a specific SCN.
   b. Select **Date and Time** if you want to recover to a specific date and time.
8. In the Pre-Ops and Post Ops page, specify the required details.
9. In the Notification page, from the **Email preference** drop-down list, select the scenarios in which you want

to send the emails.

10. Review the summary, and then click **Finish**.
11. Monitor the restore process by using the **Monitor** > **Jobs** page.

Clone Oracle database backup

You can use SnapCenter to clone an Oracle database using the backup of the database. The clones created are thick clones and are created on the parent capacity pool.

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page, select either **Database** or **Resource Group** from the View list.
3. Select the database.
4. From the Manage Copies view page, select the backup from primary storage system.
5. Select the Data backup, and then click .
6. In the Name page, select whether you want to clone a database (CDB or non CDB) or clone a pluggable database (PDB).
7. In the Locations page, specify the required details.

   If the Oracle database ANF volumes are configured in a manual QOS capacity pool, specify the QOS for the cloned volumes.

   If QOS for the cloned volumes is not specified, the QOS of the source volume will be used. If the automatic QOS capacity pool is used, the QOS value specified will be ignored.

8. In the Credentials page, perform one of the following:
   a. For Credential name for sys user, select the Credential to be used for defining the sys user password of the clone database.
   b. For ASM Instance Credential name, select **None** if OS authentication is enabled for connecting to the ASM instance on the clone host.

      Otherwise, select the Oracle ASM credential configured with either “sys” user or a user having “sysasm” privilege applicable to the clone host.

9. In the Pre-Ops page specify the path and arguments of the prescripts and in the Database Parameter settings section, modify the values of prepopulated database parameters that are used to initialize the database.
10. In the Post-Ops page, **Recover database** and **Until Cancel** are selected by default to perform recovery of the cloned database.
    a. If you select **Until Cancel**, SnapCenter performs recovery by mounting the latest log backup having the unbroken sequence of archive logs after that data backup that was selected for cloning.
    b. If you select **Date and time**, SnapCenter recovers the database up to a specified date and time.
    c. If you select **Until SCN**, SnapCenter recovers the database up to a specified SCN.
    d. If you select **Specify external archive log locations**, SnapCenter identifies and mounts optimal number of log backups based on the specified SCN or the selected date and time.
    e. By default, **Create new DBID** check box is selected to generate a unique number (DBID) for the cloned

database differentiating it from the source database.

Clear the check box if you want to assign the DBID of the source database to the cloned database. In this scenario, if you want to register the cloned database with the external RMAN catalog where the source database is already registered, the operation fails.

f. Select **Create tempfile for temporary tablespace** check box if you want to create a tempfile for the default temporary tablespace of the cloned database.

g. In **Enter sql entries to apply when clone is created**, add the sql entries that you want to apply when the clone is created.

h. In **Enter scripts to run after clone operation**, specify the path and the arguments of the postscript that you want to run after the clone operation.

11. In the Notification page, from the **Email preference** drop-down list, select the scenarios in which you want to send the emails.
12. Review the summary, and then select **Finish**.
13. Monitor the operation progress by selecting **Monitor** > **Jobs**.

Clone a pluggable database

You can clone a pluggable database (PDB) to a different or same target CDB on the same host or alternate host. You can also recover the cloned PDB to a desired SCN or date and time.

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page, select either **Database** or **Resource Group** from the View list.
3. Select the database of type single instance (multitenant).
4. From the Manage Copies view page, select the backup from primary storage system.
5. Select the backup, and then click .
6. In the Name page, select **PDB Clone** and specify the other details.
7. In the Locations page, specify the required details.
8. In the Pre-Ops page specify the path and arguments of the prescripts and in the Database Parameter settings section, modify the values of prepopulated database parameters that are used to initialize the database.
9. In the Post-Ops page, **Until Cancel** is selected by default to perform recovery of the cloned database.
   a. If you select **Until Cancel**, SnapCenter performs recovery by mounting the latest log backup having the unbroken sequence of archive logs after that data backup that was selected for cloning.
   b. If you select **Date and time**, SnapCenter recovers the database up to a specified date and time.
   c. If you select **Specify external archive log locations**, SnapCenter identifies and mounts optimal number of log backups based on the specified SCN or the selected date and time.
   d. By default, **Create new DBID** check box is selected to generate a unique number (DBID) for the cloned database differentiating it from the source database.

      Clear the check box if you want to assign the DBID of the source database to the cloned database. In this scenario, if you want to register the cloned database with the external RMAN catalog where the source database is already registered, the operation fails.

e. Select **Create tempfile for temporary tablespace** check box if you want to create a tempfile for the default temporary tablespace of the cloned database.

f. In **Enter sql entries to apply when clone is created**, add the sql entries that you want to apply when the clone is created.

g. In **Enter scripts to run after clone operation**, specify the path and the arguments of the postscript that you want to run after the clone operation.

10. In the Notification page, from the **Email preference** drop-down list, select the scenarios in which you want to send the emails.
11. Review the summary, and then select **Finish**.
12. Monitor the operation progress by selecting **Monitor** > **Jobs**.

Manage SnapCenter Server and plug-ins

View dashboard

Overview of dashboard

From the SnapCenter left navigation pane, the Dashboard gives you a first glance into the health of your system, including recent job activity, alerts, protection summary, storage efficiency and usage, status of SnapCenter jobs (Backup, Clone, Restore), configuration status for standalone and Windows cluster hosts, number of Storage Virtual Machines (SVMs) managed by SnapCenter, and license capacity.

Information displayed in the Dashboard view depends on the role assigned to the user that is currently logged in to SnapCenter. Some content might not be displayed if the user does not have permission to view that information.

In many cases, you can view more information about a display by hovering on **i**. In some cases, information in dashboard displays is linked to detailed source information in SnapCenter GUI pages such as Resources, Monitor, and Reports.

Recent Job Activities

The Recent Job Activities tile displays the latest job activity from any Backup, Restore, and Clone jobs that you have access to. Jobs in this display have one of the following states: Completed, Warning, Failed, Running, Queued, and Canceled.

Hovering over a job provides more information. You can view additional job information by clicking a specific job number, which redirects you to the Monitor page. From there, you can get job details or log information, and generate a report specific to that job.

Click **See All** to view a history of all SnapCenter jobs.

Alerts

The Alerts tile displays the latest unresolved Critical and Warning alerts for the hosts and SnapCenter Server.

The total count of Critical and Warning category alerts is shown at the top of the display. Clicking the Critical or Warning totals redirects you to the Alerts page with the specific filter applied in the Alerts page.

Clicking a specific alert redirects you to the Alerts page for details about that alert. Clicking **See All** at the bottom of the display redirects you to the Alerts page for a list of all alerts.

Latest Protection Summary

The Latest Protection Summary tile gives you the protection status for all entities that you have access to. By default, the display is set to provide the status for all plug-ins. Status information is provided for resources backed up to primary storage as Snapshots, and to secondary storage using SnapMirror and SnapVault technologies. The availability of protection status information for secondary storage is based on the selected plug-in type.

If you are using a mirror-vault protection policy, the counters for the protection summary are displayed in the SnapVault summary chart and not in the SnapMirror chart.

Protection status for individual plug-ins is available by selecting a plug-in from the drop-down menu. A donut chart shows the percentage of protected resources for the selected plug-in. Clicking a donut slice redirects you to the **Reports** > **Plug-in** page, which provides a detailed report of all primary and secondary storage activity for the specified plug-in.

Reports about secondary storage apply to SnapVault only; SnapMirror reports are not supported.

SAP HANA provides protection status information for primary and secondary storage for Snapshots. Only primary storage protection status is available for file-based backups.

Protection status	Primary storage	Secondary storage
Failed	Count of entities that are part of a Resource Group, where the Resource Group has run a backup, but the backup failed.	Count of entities with backups that have failed to transfer to a Secondary destination.
Successful	Count of entities in a resource group, where the Resource Group has been successfully backed up.	Count of entities with backups that have been successfully transferred to a Secondary destination.
Not configured	Count of entities that are not part of any Resource Group and have not been backed up.	Count of entities that are part of one or more Resource Groups that are not configured for backups to be transferred to a Secondary destination.
Not initiated	Count of entities that are part of a Resource Group, but no backup has been run.	Not applicable.

If you are using SnapCenter Server 4.2 and an earlier version of the plug-in (earlier than 4.2) to create backups, the **Latest Protection Summary** tile does not display the SnapMirror protection status of these backups.

Jobs

The Jobs tile provides you with a summary of backup, restore, and clone jobs that you have access to. You can customize the time frame for any report by using the drop-down menu. Time frame options are fixed at last 24 hours, last 7 days, and last 30 days. The default report shows data protection jobs run during the last 7 days.

Backup, restore, and clone job information is displayed in donut charts. Clicking a donut slice redirects you to the Monitor page with job filters pre-applied to the selection.

Job status	Description
Failed	Count of jobs that have failed.
Warning	Count of jobs that have experienced an error.
Successful	Count of jobs that have completed successfully.
Running	Count of jobs that are currently running.

Storage

The Storage tile displays the primary and secondary storage consumed by protection jobs over a 90-day period, graphically depicts consumption trends, and calculates primary storage savings. Storage information is updated once every 24 hours at 12 a.m.

The day's consumption total, which comprises the total number of backups that are available in SnapCenter and size occupied by these backups, will be displayed at the top of the display. A backup could have multiple Snapshots associated with it and the count will reflect the same. This is applicable to both primary and secondary Snapshots. For example, you have created 10 backups, out of which 2 are deleted due to policy-based backup retention and 1 backup is explicitly deleted by you. Thus, a count of 7 backups will be displayed along with the size occupied by these 7 backups.

The Storage Savings factor for primary storage is the ratio of logical capacity (clone and Snapshots savings plus storage consumed) to the physical capacity of primary storage. A bar chart illustrates the storage savings.

The line graph separately plots primary and secondary storage consumption on a day-by-day basis over a rolling 90-day period. Hovering over the charts provides detailed day-by-day results.

If you use SnapCenter Server 4.2 and an earlier version of the plug-in (earlier than 4.2) to create backups, the **Storage** tile does not display the number of backups, the storage consumed by these backups, the Snapshot savings, the clone savings, and the Snapshot size.

Configuration

The Configuration tile provides consolidated status information for all active stand-alone and Windows cluster hosts that SnapCenter is managing, and that you have access to. This includes the plug-in status information associated with those hosts.

Clicking the number adjacent to Hosts redirects you to the Managed Hosts section in the Hosts page. From there, you can obtain detailed information for a selected host.

Additionally, this display shows the sum of Standalone ONTAP SVMs and Cluster ONTAP SVMs that SnapCenter is managing and that you have access to. Clicking the number adjacent to SVM redirects you to the Storage Systems page. From there, you can obtain detailed information for a selected SVM.

The Host configuration state is presented as red (critical), yellow (warning), and green (active), along with the number of hosts in each state. Status messages are provided for each state.

Configuration status	Description
Upgrade mandatory	Count of hosts that are running unsupported plug-ins and need an upgrade. An unsupported plug-in is not compatible with this version of SnapCenter.
Migration mandatory	Count of hosts that are running unsupported plug-ins and need migration. An unsupported plug-in is not compatible with this version of SnapCenter.
No plug-ins installed	Count of hosts that are added successfully but the plug-ins need to be installed, or the plug-ins installation has failed.
Suspended	Count of hosts whose schedules are suspended and are under maintenance.
Stopped	Count of hosts that are up, but the plug-in services are not running.
Host down	Count of hosts that are down or not reachable.
Upgrade available (optional)	Count of hosts where a newer version of the plug-in package is available for upgrade.
Migration available (optional)	Count of hosts where a newer version of the plug-in is available for migration.
Configure log directory	Count of hosts where the log directory has to be configured for SCSQL to take transaction log backup.
Configure VMware plug-ins	Count of hosts where the SnapCenter Plug-in for VMware vSphere needs to be added.
Unknown	Count of hosts that have been registered but the installation is not yet triggered.
Running	Count of hosts that are up and plug-ins are running. And in the case of SCSQL plug-ins, log directory and hypervisor are configured.
Installing\Uninstalling plug-ins	Count of hosts where plug-in installation or uninstallation in progress.

How to view information on the dashboard

From the SnapCenter left navigation pane, you can view various Dashboard tiles, or displays, along with associated system details. The number of displays available in the

Dashboard is fixed and cannot be changed. The content provided within each display is dependent on role-based access control (RBAC).

Steps

1. In the left navigation pane, click **Dashboard**.
2. Click the active areas on each display to obtain additional information.

   For example, clicking a donut chart in **Jobs**, redirects you to the Monitor page for more information about your selection. Clicking a donut chart in **Protection Summary**, redirects you to the Reports page, which can provide more information about your selection.

Request status reports of the jobs from the dashboard

You can request reports about backup, restore, and clone jobs from the Dashboard page. This is useful if you want to identify the total number of successful or failed jobs in your SnapCenter environment.

Steps

1. In the left navigation pane, click **Dashboard**
2. Locate the Jobs tile in the Dashboard, and then select **Backup**, **Restore**, or **Clone**.
3. Using the pull-down menu, select the time frame for which you want Jobs information: 24 hours, 7 days, or 30 days.

   The systems display a donut chart covering the data.

4. Click the donut slice representing the job information for which you want a report.

   When you click the donut chart, you are redirected from the Dashboard page to the Monitor page. The Monitor page displays the jobs with the status you selected from the donut chart.

5. From the Monitor page list, click on a specific job to select it.
6. At the top of the Monitor page, click **Reports**.

Result

The report displays information only for the job you selected. You can review the report or download it to your local system.

Request reports of the protection status from the dashboard

You can request protection details for resources managed by specific plug-ins using the Dashboard. Only data backups are considered for data protection summary.

Steps

1. In the left navigation pane, click **Dashboard**.
2. Locate the Latest Protection Summary tile in the Dashboard and use the pull-down menu to select a plug-in.

The Dashboard displays a donut chart for resources backed up to Primary storage and, if applicable to the plug-in, a donut chart for resources backed up to secondary storage.

Data protection reports are available only for specific plug-ins types. Specifying **All Plug-ins** is not supported.

3. Click the donut slice representing the status for which you want a report.

   When you click the donut chart, you are redirected from Dashboard page to the Reports, and then to the Plug-in page. The report displays only status for the plug-in you selected. You can review the report or download it to your local system.

   Redirection to the Reports page for SnapMirror donut chart and File-based SAP HANA backup is not supported.

Manage RBAC

SnapCenter allows you to modify roles, users, and groups.

Modify a role

You can modify a SnapCenter role to remove users or groups and change the permissions associated with the role. It is especially useful to modify roles when you want to change or eliminate the permissions used by an entire role.

Before you begin

You must have logged in as the "SnapCenterAdmin" role.

You cannot modify or remove permissions for the SnapCenterAdmin role.

Steps

1. In the left navigation pane, click **Settings**.
2. In the Settings page, click **Roles**.
3. From the Role name field, click the role you want to modify.
4. Select **All members of this role can see other members' objects** to enable other members of the role to see resources such as volumes and hosts after they refresh the resources list.

   Deselect this option if you do not want members of this role to see objects to which other members are assigned.

   When this option is enabled, assigning users access to objects or resources is not required if users belong to the same role as the user who created the objects or resources.

5. In the Role Details page alter the permissions or unassign the members as needed.
6. Click **Submit**.

Modify users and groups

You can modify SnapCenter users or groups to alter their roles and assets.

Before you begin

You must be logged in as the SnapCenter administrator.

Steps

1. In the left navigation pane, click **Settings**.
2. In the Settings page, click **Users and Access**.
3. From the User or Group name list, click the user or group that you want to modify.
4. In the User or Group details page, alter roles and assets.
5. Click **Submit**.

Manage hosts

You can add hosts and install SnapCenter plug-in packages, add a verification server, remove hosts, migrate backup jobs, and update host to upgrade plug-in packages or add new plug-in packages. Depending on the plug-in you are using, you can also provision disks, manage SMB shares, manage initiator groups (igroups), manage iSCSI sessions, and migrate data.

You can perform these tasks…	For Microsoft Exchange Server	For Microsoft SQL Server	For Microsoft Windows	For Oracle Database	For SAP HANA Database	For NetApp supported plug-ins	For Db2	For PostgreSQL	For MySQL
Add hosts and install plug-in package	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Update ESXi information for a host	No	Yes	No	No	No	No	No	No	No
Suspend schedules and place hosts in maintenance mode	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

You can perform these tasks…	For Microsoft Exchange Server	For Microsoft SQL Server	For Microsoft Windows	For Oracle Database	For SAP HANA Database	For NetApp supported plug-ins	For Db2	For PostgreSQL	For MySQL
Modify hosts by adding, upgrading, or removing plug-ins	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Remove hosts from SnapCenter	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Start plug-in services (Only applicable for plug-ins running on Windows host)	Yes	Yes	Yes	No	Yes	Yes	Yes	Yes	Yes
Provision disks	No	No	Yes	No	No	No	No	No	No
Manage SMB shares	No	No	Yes	No	No	No	No	No	No
Manage iGroups	No	No	Yes	No	No	No	No	No	No
Manage iSCSI sessions	No	No	Yes	No	No	No	No	No	No

Refresh virtual machine information

You should refresh your virtual machine information when VMware vCenter credentials change or the database or file system host restarts. Refreshing your virtual machine information in SnapCenter initiates communication with the VMware vSphere vCenter and obtains vCenter credentials.

RDM-based disks are managed by the SnapCenter Plug-in for Microsoft Windows, which is installed on the database host. To manage RDMs, the SnapCenter Plug-in for Microsoft Windows communicates with the vCenter server that manages the database host.

Steps

1. In the SnapCenter left navigation pane, click **Hosts**.
2. In the Hosts page, click **Managed Hosts**.
3. In the Managed Hosts page, select the host you want to update.
4. Click **Refresh VM**.

Modify plug-in hosts

After installing a plug-in, you can modify the plug-in hosts details if required. You can modify credentials, installation path, plug-ins, log directory details for SnapCenter Plug-in for Microsoft SQL Server, group Managed Service Account (gMSA), and the plug-in port.

Ensure that the plug-in version is the same as that of the SnapCenter Server version.

About this task

- You can modify a plug-in port only after the plug-in is installed.

  You cannot modify the plug-in port while upgrade operations are in progress.

- While modifying a plug-in port, you should be aware of the following port rollback scenarios:
  - In a standalone setup, if SnapCenter fails to change the port of one of the components, the operation fails and the old port is retained for all of the components.

    If the port was changed for all of the components but one of the components fails to start with the new port, then the old port is retained for all of the components. For example, if you want to change the port for two plug-ins on the stand-alone host and SnapCenter fails to apply the new port to one of the plug-ins, the operation fails (with an appropriate error message) and the old port is retained for both the plug-ins.

  - In a clustered setup, if SnapCenter fails to change the port of the plug-in that is installed on one of the nodes, the operation fails and the old port is retained for all of the nodes.

    For example, if the plug-in is installed on four nodes in a clustered setup, and if the port is not changed for one of the nodes, the old port is retained for all of the nodes.

When plug-ins are installed with gMSA, you can modify in the **More Options** windows. When plug-ins are installed without gMSA, you can specify the gMSA account to use it as the plug-in service account.

Steps

1. In the left navigation pane, click **Hosts**.
2. Verify that **Managed Hosts** is selected at the top.
3. Select the host for which you want to modify and modify any one field.

   Only one field can be modified at a time.

4. Click **Submit**.

Result

The host is validated and added to SnapCenter Server.

Start or restart plug-in services

Starting the SnapCenter plug-in services enable you to start services if they are not running or restart them if they are running. You might want to restart services after maintenance has been performed.

You should ensure that no jobs are running when restarting the services.

Steps

1. In the left navigation pane, click **Hosts**.
2. In the Hosts page, click **Managed Hosts**.
3. In the Managed Hosts page, select the host you want to start.
4. Click icon and click **Start Service** or **Restart Service**.

   You can start or restart service of multiple hosts simultaneously.

Suspend schedules for host maintenance

When you want to prevent the host from running any SnapCenter scheduled jobs, you can place your host in maintenance mode. You should do this before you upgrade the plug-ins or if you are performing maintenance tasks on hosts.

> You cannot suspend the schedules on a host that is down because SnapCenter cannot communicate with that host.

Steps

1. In the left navigation pane, click **Hosts**.
2. In the Hosts page, click **Managed Hosts**.
3. In the Managed Hosts page, select the host that you want to suspend.
4. Click the icon, and then click **Suspend Schedule** to place the host for this plug-in in maintenance mode.

   You can suspend the schedule of multiple hosts simultaneously.

   > You do not have to stop the plug-in service first. The plug-in service can be in a running or stopped state.

Result

After you suspend the schedules on the host, the Managed Hosts page shows **Suspended** in the Overall status field for the host.

After you complete host maintenance, you can bring the host out of maintenance mode by clicking **Activate Schedule**.
You can activate the schedule of multiple hosts simultaneously.

Operations supported from the Resources page

You can discover resources and perform data protection operations from the Resources page. The operations you can perform differ based on the plug-in you are using to manage your resources.

From the Resources page, you can perform the following tasks:

You can perform these tasks…	For Microsoft Exchange Server	For Microsoft SQL Server	For Microsoft Windows	For Oracle Database	For SAP HANA Database
Determine whether resources are available for backup	Yes	Yes	Yes	Yes	Yes
Perform on-demand backup of a resource	Yes	Yes	Yes	Yes	Yes
Restore from backups	Yes	Yes	Yes	Yes	Yes
Clone backups	No	Yes	Yes	Yes	Yes
Manage backups	Yes	Yes	Yes	Yes	Yes
Manage clones	No	Yes	Yes	Yes	Yes
Manage policies	Yes	Yes	Yes	Yes	Yes
Manage storage connections	Yes	Yes	Yes	Yes	Yes
Mount backups	No	No	No	Yes	No
Unmount backups	No	No	No	Yes	No
View details	Yes	Yes	Yes	Yes	Yes

Manage policies

You can detach policies from a resource or resource group, modify, delete, view, and copy.

Modify policies

You can modify the replication options, Snapshot retention settings, error retry count, or scripts information while a policy is attached to a resource or resource group. You can modify the schedule type (frequency) only after you detach a policy.

About this task

Modifying the schedule type in a policy requires additional steps because the SnapCenter Server registers the schedule type only at the time the policy is attached to a resource or resource group.

If you want to…	Then…
Add an additional schedule type	Create a new policy and attach it to the necessary resources or resource groups. For example, if a resource group policy specifies only hourly backups and you want to add daily backups also, you can create a policy with a daily schedule type and add it to the resource group. The resource group would then have two policies: hourly and daily.
Remove or change a schedule type	Perform the following: 1. Detach the policy from every resource and resource group that uses that policy. 2. Modify the schedule type. 3. Attach the policy again to all the resources and resource groups. For example, if a policy specifies hourly backups and you want to change that to daily backups, you must detach the policy first.

Steps

1. In the left navigation pane, click **Settings**.
2. In the Settings page, click **Policies**.
3. Select the policy, and then click **Modify**.
4. Modify the information, and then click **Finish**.

Detach policies

You can detach policies from a resource or resource group any time that you no longer want those policies to govern data protection for the resources. You must detach a policy before you can delete it or before you

modify the schedule type.

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page, select **Resource Group** from the **View** list.
3. Select the resource group, and then click **Modify Resource Group**.
4. In the Policies page of the Modify Resource Group wizard, from the drop-down list, clear the check mark next to the policies you want to detach.
5. Make any additional modifications to the resource group in the rest of the wizard, and then click **Finish**.

Delete policies

If you no longer require policies, you might want to delete them.

Before you begin

You should detach the policy from resource or resource groups if the policy is associated with any resource or resource groups.

Steps

1. In the left navigation pane, click **Settings**.
2. In the Settings page, click **Policies**.
3. Select the policy, and then click **Delete**.
4. Click **Yes**.

Manage resource groups

You can perform various operations on resource groups.

You can perform the following tasks related to resource groups:

- Modify a resource group by selecting the resource group and clicking **Modify Resource Group** to edit the information you provided while creating the resource group.

  You can change the schedule while modifying the resource group. However, to change the schedule type you must modify the policy.

  If you remove resources from a resource group, the backup retention settings defined in the policies currently attached to the resource group will continue to be applied to the removed resources.

- Create a backup of a resource group.
- Create a clone of a backup.

  You can clone from the existing backups of SQL, Oracle, Windows file systems, custom applications, and SAP HANA database resources or resource groups.

- Create a clone of a resource group.

This operation is supported only for SQL resource groups (which contains only databases). You can configure a schedule for cloning a resource group (clone lifecycle).

- Prevent scheduled operations on resource groups from starting.
- Delete a resource group.

Stop and resume operations on resource groups

You can temporarily disable scheduled operations from starting on a resource group. Later when you want, you can enable those operations.

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page, select **Resource Group** from the **View** list.
3. Select the resource group and click **Maintenance**.
4. Click **OK**.

If you want to resume operations on the resource group that you had put on maintenance mode, select the resource group and click **Production**.

Delete resource groups

You can delete a resource group if you no longer need to protect the resources in the resource group. You must ensure that resource groups are deleted before you remove plug-ins from SnapCenter.

About this task

You should manually delete all clones created for any of the resources in the resource group. You can optionally force the deletion of all backups, metadata, policies, and Snapshots associated with the resource group.

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page, select **Resource Group** from the **View** list.
3. Select the resource group, and then click **Delete**.
4. Optional: Select the **Delete backups and detach policies associated with this Resource Group** check box to remove all backups, metadata, policies, and Snapshots associated with the resource group.
5. Click **OK**.

Manage backups

You can rename and delete backups. You can also delete multiple backups simultaneously.

Rename backups

You can rename backups if you want to provide a better name to improve searchability.

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page either select the resource or resource group from the **View** drop-down list.
3. Select the resource or resource group from the list.

   The resource or resource group topology page is displayed. If the resource or resource group is not configured for data protection, the Protect wizard is displayed instead of the topology page.

4. From the Manage Copies view, select **Backups** from the primary storage systems.

   You cannot rename the backups that are on the secondary storage system.

   If you have cataloged the backups of Oracle databases using Oracle Recovery Manager (RMAN), you cannot rename those cataloged backups.

5. Select the backup, and then click .
6. In the **Rename backup as** field, enter a new name and click **OK**.

Delete backups

You can delete backups if you no longer require the backup for other data protection operations.

Before you begin

You must have deleted the associated clones before deleting a backup.

If a backup is associated with a cloned resource, you cannot delete the backup.

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page either select the resource or resource group from the **View** drop-down list.
3. Select the resource or resource group from the list.

   The resource or resource group topology page is displayed.

4. From the Manage Copies view, select **Backups** from the primary storage systems.

   You cannot delete the backups that are on the secondary storage system.

5. Select the backup, and then click .

   If you are deleting a SAP HANA database backup, the associated SAP HANA catalogs of the backup are also deleted.

   If the last remaining backup is deleted, the associated HANA catalog entries cannot be deleted.

6. Click **OK**.

If you have some stale database backups in SnapCenter which do not have corresponding backups on the storage system, you must use remove-smbackup command to clean up these stale backup entries. If the stale backups were cataloged, they will be uncataloged from the recovery catalog database.

Remove protection

Remove protection deletes all the backups and detaches all the policies. Before removing protection, you should ensure that the backups are not mounted and no clones are associated with the backup.

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page either select the resource or resource group from the **View** drop-down list.
3. Select the resource or resource group from the list.

   The resource or resource group topology page is displayed.

4. Select the backup and click **Remove Protection**.

Delete clones

You can delete clones if you find them no longer necessary.

About this task

You cannot delete clones that acts like source for other clones.

For example, if the production database is db1, database clone1 is cloned from backup of db1 and subsequently clone1 is protected. The database clone2 is cloned from backup of clone1. If you decide to delete clone1, you must first delete clone2, and then delete clone1.

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page, either select the resource or resource group from the **View** drop-down list.
3. Select the resource or resource group from the list.

   The resource or the resource group topology page is displayed.

4. From the Manage Copies view, select **Clones** either from the primary or secondary (mirrored or replicated) storage systems.
5. Select the clone, and then click .

   If you are deleting SAP HANA database clones, in the Delete Clone page, perform the following actions:

   a. In the **Pre clone delete** field, enter the commands that should be run before deleting the clone.
   b. In the **Unmount** field, enter the command to unmount the clone before deleting the clone.
6. Click **OK**.

After you finish

Sometimes the file systems are not deleted. You must increase the value of the CLONE_DELETE_DELAY parameter by running the following command: `./sccli Set-SmConfigSettings`

> The CLONE_DELETE_DELAY parameter specifies the number of seconds to wait after completing the deletion of application clone and before starting the deletion of file system.

After modifying the value of the parameter, restart the SnapCenter Plug-in Loader (SPL) service.

Monitor jobs, schedules, events, and logs

You can monitor the progress of your jobs, get information about scheduled jobs, and review events and logs from the Monitor page.

Monitor jobs

You can view information about SnapCenter backup, clone, restore, and verification jobs. You can filter this view based on start and end date, type of job, resource group, policy, or SnapCenter plug-in. You can also get additional details and log files for specified jobs.

You can also monitor jobs related to SnapMirror and SnapVault operations.

> You can monitor only the jobs that you created and that are relevant to you unless you are assigned SnapCenter Admin or another super user role.

You can perform the following tasks related to monitoring jobs:

- Monitor backup, clone, restore, and verification operations.
- View job details and reports.
- Stop a scheduled job.

Manage scheduled backup jobs

Beginning with SnapCenter 6.0.1 release, a new parameter **JobConcurrencyThreshold** has been introduced that sets a threshold to the number of scheduled jobs that can run at any given time. This allows you to control the number of backups you want to run based on your systems hardware configuration.

The default value assigned to **JobConcurrencyThreshold** is 0 and is disabled. You can enable by assigning a value, if you observe performance degradation during the scheduled backup window.

> If you enable **JobConcurrencyThreshold** to manage concurrent jobs, SnapCenter will not allow you to control the order of the backups and the backups might not be triggered at the same time specified in the schedule.

Steps

1. Set the value of *JobConcurrencyThreshold* parameter located at *C:\Program Files\NetApp\SnapCenter WebApp\SnapManager.Web.UI.dll.config*.
2. Recycle the SnapCenter application pool by clicking IIS > Application Pools > SnapCenter > Restart.
3. Restart the SnapCenter Web service by clicking IIS > Sites > SnapCenter > Restart.

Manage stale jobs

Stale jobs are created from interruptions in SnapCenter or from improper job updates. Beginning with SnapCenter 6.0.1 release, a predefined schedule is introduced to clean up these stale jobs that are stuck for more than 72 hours. You can change the schedule frequency by editing the configurable parameter **CleanUpStaleJobsIntervalHours**.

You can trigger the clean up on demand by running the schedule from **Monitor** > **Schedules** > **SnapCenter_StaleJobCleanUp**.

Steps

1. Set the value of *CleanUpStaleJobsIntervalHours* parameter located at *C:\Program Files\NetApp\SnapCenter WebApp\SnapManager.Web.UI.dll.config*.
2. Recycle the SnapCenter application pool by clicking IIS > Application Pools > SnapCenter > Restart.
3. Restart the SnapCenter Web service by clicking IIS > Sites > SnapCenter > Restart.

Monitor schedules

You might want to view current schedules to determine when the operation starts, when it was last run, and when it runs next. You can also determine the host on which the operation runs, along with the operation's resource group and policy information.

Steps

1. In the left navigation pane, click **Monitor**.
2. In the Monitor page, click **Schedules**.
3. Select the resource group and the schedule type.
4. View the list of scheduled operations.

Monitor events

You can view a list of SnapCenter events in the system, such as when a user creates a resource group or when the system initiates activities, such as creating a scheduled backup. You might want to view events to determine if an operation such as a backup or a restore operation is currently in progress.

About this task

All job information appears in the Events page. For example, when a backup job starts, a "backup start" event appears. When the backup completes, a "backup complete" event appears.

Steps

1. In the left navigation pane, click **Monitor**.
2. In the Monitor page, click **Events**.
3. (Optional) In the Filter box, enter the start or end date, category of event (such as backup, resource group, or policy) and severity level, and click **Apply**. Alternatively, enter characters in the Search box.
4. View the list of events.

Monitor logs

You can view and download SnapCenter Server logs, SnapCenter host agent logs, and plug-in logs. You might want to view the logs to help with troubleshooting.

About this task

You can filter the logs to show only a specific log severity level:

- Debug
- Info
- Warn
- Error
- Fatal

You can also obtain job level logs, for example, logs that help you troubleshoot the reason for a backup job failure. For job level logs, use the **Monitor** > **Jobs** option.

Steps

1. In the left navigation pane, click **Monitor**.
2. In the Jobs page, select a job and click Download logs.

   The downloaded zipped folder contains the job logs and the common logs. The zipped folder name contains the job id and job type selected.

3. In the Monitor page, click **Logs**.
4. Select the log type, host, and instance.

   If you select log type as **plugin**, you can select a host or SnapCenter plug-in. You cannot do this if the log type is **server**.

5. To filter the logs by a specific source, message, or log level, click the filter icon at the top of the column heading.

   To show all logs, choose **Greater than or equal to** as the `Debug` level.

6. Click **Refresh**.
7. View the list of logs.
8. Click **Download** to download the logs.

   The downloaded zipped folder contains the job logs and the common logs. The zipped folder name contains the job id and job type selected.

In large configurations for optimum performance, you should set the log settings for SnapCenter to minimal level by using the PowerShell cmdlet.

```
Set-SmLogSettings -LogLevel All -MaxFileSize 10MB -MaxSizeRollBackups 10
-JobLogsMaxFileSize 10MB -Server
```

To access health or configuration information after a failover job finishes, run the cmdlet `Get-SmRepositoryConfig`.

Remove jobs and logs from SnapCenter

You can remove backup, restore, clone, and verification jobs and logs from SnapCenter. SnapCenter stores successful and failed job logs indefinitely unless you remove them. You might want to remove them to replenish storage.

About this task

There must be no jobs currently in operation.
You can remove a specific job by providing a Job ID or you can remove jobs within a specified period.

You do not need to place the host in maintenance mode to remove jobs.

Steps

1. Launch PowerShell.
2. From the command prompt, enter: `Open-SMConnection`
3. From the command prompt, enter: `Remove-SmJobs`
4. In the left navigation pane, click **Monitor**.
5. In the Monitor page, click **Jobs**.
6. In the Jobs page, review the status of the job.

Related information

The information regarding the parameters that can be used with the cmdlet and their descriptions can be obtained by running *Get-Help command_name*. Alternatively, you can also refer to the SnapCenter Software Cmdlet Reference Guide.

Overview of SnapCenter reporting capabilities

SnapCenter provides a variety of reporting options that enable you to monitor and manage your system health and operation success.

Report type	Description
Backup Report	The Backup Report provides overall data about backup trends for your SnapCenter environment, the backup success rate, and some information about each backup performed during the specified time. If a backup is deleted, the report does not display any status information for the deleted backup. The Backup Detail Report provides detailed information about a specified backup job and lists the resources successfully backed up and any that have failed.

Report type	Description
Clone Report	The Clone Report provides overall data about clone trends for your SnapCenter environment, the clone success rate, and some information about each clone job performed during the specified time. If a clone is deleted, the report does not display any status information for the deleted clone. The Clone Detail Report provides details about the specified clone, clone host, and clone job task status. If a task fails, the Clone Detail Report displays information about the failure.
Restore Report	The Restore Report provides overall information about restore jobs. The Restore Detail Report provides details about a specified restore job, including host name, backup name, job start and duration, and the status of individual job tasks. If a task fails, the Restore Detail Report displays information about the failure.
Protection Report	These reports provide protection details for resources managed by all SnapCenter plug-in instances. This report provides protection details for resources managed by all plug-in instances. You can see an overview, details of unprotected resources, resources that have not been backed up when the report was generated, resources of a resource group for which backup operations have failed, and SnapVault status.
Scheduled Report	These reports are scheduled to run periodically like daily, weekly or monthly. The reports are generated automatically on the specified date and time and the report is sent to the respective people through e-mail. You can enable, disable, modify, or delete the schedules. The enabled schedule can be run on demand by clicking on the **Run Now** button. The administrator can run any schedule, but the generated report will contain data based on the permission provided by the user who created the schedule. Any other user other than Administrator will be able to see or modify schedule based on their permission .If All members of this role can see other members' objects option is selected in the Add Role page, then other members of the role will be able to see and modify.

Access reports

You can use the SnapCenter Dashboard to get a quick overview of the health of your system. From the Dashboard you can drill into more details. Alternatively, you can access the detailed reports directly.

You can access reports by one of the following methods:

- In the left navigation pane, click **Dashboard**, and then click **Last Protection Summary** pie chart to see more details in the Reports page.
- In the left navigation pane, click **Reports**.

Filter your report

You might want to filter your report data according to a range of parameters, depending on the level of detail and time span of information you require.

Steps

1. In the left navigation pane, click **Reports**.
2. If the Parameter view is not displayed, click the **Toggle Parameters Area** icon from the report toolbar.
3. Specify the time range for which you want to run your report.

   If you omit the end date, you retrieve all available information.
4. Filter your report information based on any of the following criteria:
   - Resource group
   - Host
   - Policy
   - Resource
   - Status
   - Plug-in Name
5. Click **Apply**.

Export or print reports

Exporting SnapCenter reports enables you to view the report in a variety of alternative formats. You can also print reports.

Steps

1. In the left navigation pane, click **Reports**.
2. From the reports toolbar, perform one of the following:
   - Click the **Toggle Print Preview** icon to preview a printable report.
   - Select a format from the **Export** icon drop-down list to export a report to an alternate format.
3. To print a report, click the **Print** icon.
4. To view a specific report summary, scroll to the appropriate section of the report.

Set the SMTP server for email notifications

You can specify the SMTP server to use for sending data protection job reports to yourself or to others. You can also send a test email to verify the configuration. The settings are applied globally for any SnapCenter job for which you configure email notification.

This option configures the SMTP server for sending all data protection job reports. However, if you want to have regular SnapCenter data protection job updates for a particular resource sent to yourself or to others so that you can monitor the status of those updates, you can configure the option to email the SnapCenter reports when you are creating a resource group.

Steps

1. In the left navigation pane, click **Settings**.
2. In the Settings page, click **Global Settings**.
3. Enter the SMTP server and click **Save**.
4. To send a test email, enter the email address from and to which you will send the email, enter the subject, and click **Send**.

Configure the option to email reports

If you want to have regular SnapCenter data protection job updates sent to yourself or to others so that you can monitor the status of those updates, you can configure the option to email the SnapCenter reports when you are creating a resource group.

Before you begin

You must have configured your SMTP server in the Global Settings page under Settings.

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. Select the type of resource you want to view and click **New Resource Group**, or select an existing resource group and click **Modify** to configure email reports for an existing resource group.
3. In the Notification panel of the New Resource Group wizard, select from the pull-down menu whether you want to receive reports always, on failure, or on failure or warning.
4. Enter the address the email is sent from, the address the email is sent to, and the subject of the email.

Manage the SnapCenter Server repository

Information related to various operations performed from SnapCenter is stored in the SnapCenter Server database repository. You must create backups of the repository to protect the SnapCenter Server from data loss.

The SnapCenter Server repository is sometimes referred to as the NSM database.

Prerequisites for protecting the SnapCenter repository

Your environment should meet certain prerequisites to protect the SnapCenter repository.

- Managing storage virtual machine (SVM) connections

  You should configure the storage credentials.

- Provisioning hosts

  At least one NetApp storage disk should be present on the SnapCenter repository host. If a NetApp disk is

not present on the SnapCenter repository host, you must create one.

For details about adding hosts, setting up SVM connections, and provisioning hosts, see the installation instructions.

- Provisioning iSCSI LUN or VMDK

For high availability (HA) configuration, you can provision either a iSCSI LUN or a VMDK in one of the SnapCenter Servers.

Back up the SnapCenter repository

Backing up the SnapCenter Server repository helps protect it from data loss. You can back up the repository by running the *Protect-SmRepository* cmdlet.

About this task

The *Protect-SmRepository* cmdlet accomplishes the following tasks:

- Creates a resource group and a policy
- Creates a backup schedule for the SnapCenter repository

Steps

1. Launch PowerShell.
2. On the SnapCenter Server host, establish a session using the *Open-SmConnection* cmdlet, and then enter your credentials.
3. Back up the repository using the *Protect-SmRepository* cmdlet and the required parameters.

View backups of the SnapCenter repository

You can display a list of SnapCenter Server database repository backups by running the *Get-SmRepositoryBackups* cmdlet.

The repository backups are created according to the schedule specified in the *Protect-SmRepository* cmdlet.

Steps

1. Launch PowerShell.
2. From the command prompt, enter the following cmdlet, and then provide credentials to connect to the SnapCenter Server: *Open-SMConnection*
3. List all available SnapCenter database backups using the *Get-SmRepositoryBackups* cmdlet.

Restore the SnapCenter database repository

You can restore the SnapCenter repository by running the *Restore-SmRepositoryBackup* cmdlet.

When you are restoring the SnapCenter repository, other SnapCenter operations that are running will be impacted because during the restore operation the repository database is not accessible.

Steps

1. Launch PowerShell.
2. From the command prompt, enter the following cmdlet, and then provide credentials to connect to the SnapCenter Server: *Open-SMConnection*
3. Restore the repository backup using the *Restore-SmRepositoryBackup* cmdlet.

   The following cmdlet restores the SnapCenter MySQL database repository from the backups existing on either iSCSI LUN or VMDK:

   ```
   C:\PS>Restore-SmRepositoryBackup -BackupName MYSQL_DS_SC_Repository_mva-x3550-s09_09-15-2016_10.32.00.4445
   ```

   The following cmdlet restores the SnapCenter MySQL database when backup files are deleted accidentally in the iSCSI LUN. For VMDK manually restore the backup from ONTAP snapshots.

   ```
   C:\PS>Restore-SmRepositoryBackup -BackupName MYSQL_DS_SC_Repository_mva-x3550-s09_09-15-2016_10.32.00.4445 -RestoreFileSystem
   ```

> The backup that was used to perform the repository restore operation will not be listed when the repository backups are retrieved after performing the restore operation.

Migrate the SnapCenter repository

You can migrate the SnapCenter Server database repository from the default location to another disk. You might migrate the repository when you want to relocate it to a disk with more space.

Steps

1. Stop the MYSQL57 service in Windows.
2. Locate the MySQL data directory.

   You can usually find the data directory at C:\ProgramData\MySQL\MySQL Server 5.7\Data.

3. Copy the MySQL data directory to the new location, for example, E:\Data\nsm.
4. Right click on the new directory, and then select **Properties** > **Security** to add the Network Service local server account to the new directory, and then assign the account full control.
5. Rename the original database directory, for example, nsm_copy.
6. From a Windows command prompt, create a symbolic directory link by using the *mklink* command.

   ```
   "mklink /d "C:\ProgramData\MySQL\MySQL Server 5.7\Data\nsm" "E:\Data\nsm" "
   ```

7. Start the MYSQL57 service in Windows.
8. Verify that the database location change is successful by logging in to SnapCenter and checking repository entries, or by logging in to the MySQL utility and connecting to the new repository.
9. Delete the original, renamed, database repository directory (nsm_copy).

Reset the SnapCenter repository password

The MySQL Server repository database password is automatically generated during SnapCenter Server installation from SnapCenter 4.2. This automatically generated password is not known to SnapCenter user at any point. If you want to access the repository database, you should reset the password.

Before you begin

You should have the SnapCenter administrator privileges to reset the password.

Steps

1. Launch PowerShell.
2. From the command prompt, enter the following command, and then provide the credentials to connect to the SnapCenter Server: *Open-SMConnection*
3. Reset the repository password: *Set-SmRepositoryPassword*

   The following command resets the repository password:

```
Set-SmRepositoryPassword at command pipeline position 1
Supply values for the following parameters:
NewPassword: ********
ConfirmPassword: ********
Successfully updated the MySQL server password.
```

Related information

The information regarding the parameters that can be used with the cmdlet and their descriptions can be obtained by running *Get-Help command_name*. Alternatively, you can also refer to the SnapCenter Software Cmdlet Reference Guide.

Manage resources of untrusted domains

In addition to managing hosts in Active Directory (AD) trusted domains, SnapCenter also manages hosts in multiple AD untrusted domains. The untrusted AD domains must be registered with the SnapCenter Server. SnapCenter supports users and groups of multiple untrusted AD domains.

You can install the SnapCenter Server on a machine that is in either a domain or a workgroup. To install the SnapCenter Server, you should specify the domain credentials if the machine is in a domain or the local administrator credentials if the machine is in a workgroup.

Active Directory (AD) groups that belong to domains not registered with the SnapCenter Server are not supported. Although you can create SnapCenter roles with these AD groups, logging in to SnapCenter Server fails with the following error message: The user you are trying to login does not belong to any roles. Please contact your administrator.

Modify untrusted domains

You can modify an untrusted domain when you want to update the domain controller IP addresses or the fully qualified domain name (FQDN).

About this task

After you modify the FQDN, the associated assets (hosts, users, and groups) might not function as expected.

To modify an untrusted domain, you can use either the SnapCenter user interface or PowerShell cmdlets.

Steps

1. In the left navigation pane, click **Settings**.
2. In the Settings page, click **Global Settings**.
3. In the Global Settings page, click **Domain Settings**.
4. Click , and then provide the following details:

For this field…	Do this…
Domain FQDN	Specify the FQDN, and click **Resolve**.
Domain controller IP addresses	If the domain FQDN is not resolvable, specify one or more domain controller IP addresses.

5. Click **OK**.

Unregister untrusted Active Directory domains

You can unregister an untrusted Active Directory domain if you do not want to use the assets that are associated with that domain.

Before you begin

You should have removed the hosts, users, groups, and credentials that are associated with the untrusted domain.

About this task

- After the domain is unregistered from SnapCenter Server, users of that domain cannot access SnapCenter Server.
- If there are associated assets (hosts, users, and groups), after unregistering the domain, the assets will be non-operational.
- To unregister an untrusted domain, you can use either the SnapCenter user interface or PowerShell cmdlets.

Steps

1. In the left navigation pane, click **Settings**.
2. In the Settings page, click **Global Settings**.
3. In the Global Settings page, click **Domain Settings**.
4. From the list of domains, select the domain that you want to unregister.
5. Click , and then click **OK**.

Manage the storage system

After adding the storage system, you can modify the storage system configuration and connections, or delete the storage system.

Modify storage system configuration

You can use SnapCenter to modify your storage system configuration if you want to change the user name, password, platform, port, protocol, timeout period, preferred IP address, or messaging options.

About this task

You can modify storage connections for an individual user or for a group. If you belong to one or more groups with permission to the same storage system, the storage connection name is displayed multiple times in the storage connection list, once for each group with permission to the storage system.

Steps

1. In the left navigation pane, click **Storage Systems**.
2. In the Storage Systems page, from the **Type** drop-down perform one of the following actions:

Select…	Steps…
ONTAP SVMs	To view all the storage virtual machines (SVMs) that were added, and to modify the required SVM configuration. 1. In the Storage Connections page, click the appropriate SVM name. 2. Perform one of the following actions: ◦ If the SVM is not part of any cluster, in the Modify Storage System page, modify the configurations such as user name, password, EMS and AutoSupport settings, platform, protocol, port, timeout, and preferred IP. ◦ If the SVM is part of a cluster, then in the Modify Storage System page, select **Manage SVM Independently** and modify the configurations such as user name, password, EMS and AutoSupport settings, platform, protocol, port, timeout, and preferred IP. After modifying the SVM to be managed independently, if you decide to manage it through cluster, you should delete the SVM and then click **Rediscover**. The SVM will be added to the ONTAP cluster. When a storage system password is updated on SnapCenter GUI, you should restart the SMCore services of the respective plug-in or the server host because the updated password does not reflect in SMCore, and the backup jobs will fail with an incorrect credential error.

Select…	Steps…
ONTAP Clusters	To view all the clusters that were added and modify the required cluster configuration. 1. In the Storage Connections page, click the cluster name. 2. In the Modify Storage System page, click the edit icon next to Username and modify the user name and password. 3. Select or clear the EMS and AutoSupport settings. 4. Click **More Options** and modify other configurations such as platform, protocol, port, timeout, and preferred IP.

3. Click **Submit**.

Delete the storage system

You can use SnapCenter to delete any unused storage system.

About this task

You can delete storage connections for an individual user or for a group. If you belong to one or more groups with permission to the same storage system, the storage system name is displayed multiple times in the storage connection list, once for each group with permission to the storage system.

When you are deleting a storage system, all operations that are being performed on that storage system will fail.

Steps

1. In the left navigation pane, click **Storage Systems**.
2. In the Storage Systems page, from the **Type** drop-down, select either **ONTAP SVMs** or **ONTAP Clusters**.
3. In the Storage Connections page, either select the check box next to the SVM, or the cluster that you want to delete.

   You cannot select the SVM that is part of a cluster.

4. Click **Delete**.
5. In the Delete Storage System Connection Settings page, click **OK**.

   If an SVM is deleted from ONTAP cluster using ONTAP GUI, in the SnapCenter GUI click **Rediscover** to update the SVM list.

REST API support

All ASA, AFF, or FAS systems connections to ONTAP will go through ZAPI by default. REST API can be enabled for specific ONTAP versions.

SnapCenter leverages REST APIs to perform all operations on ASA r2 systems, which do not support ZAPIs.

You can modify the config keys in the following configuration files:

- IsRestEnabledForStorageConnection

  Default value is false.

- MinOntapVersionToUseREST

  Default value is 9.13.1.

Enable connection through REST API

1. Set IsRestEnabledForStorageConnection to true.
2. Add the key in SMCoreServiceHost.dll.config and SnapDriveService.dll.config on both the server and the Windows plug-in hosts.

   <add key="IsRestEnabledForStorageConnection" value="true" />

Limit connection through REST API to specific version of ONTAP

1. Set the config parameter MinOntapVersionToUseREST to true.
2. Add the key in SMCoreServiceHost.dll.config and SnapDriveService.dll.config on both the server and the Windows plug-in hosts.

   <add key="MinOntapVersionToUseREST" value="9.13.1" />

3. Restart the service for SmCore on the server and the plug-in and SnapDrive service on the plug-in machine.

Manage EMS data collection

You can schedule and manage Event Management System (EMS) data collection using PowerShell cmdlets. EMS data collection involves gathering details about the SnapCenter Server, the installed SnapCenter plug-in packages, the hosts, and similar information, and then sending it to a specified ONTAP storage virtual machine (SVM).

System CPU utilization is high when data-collection task is in progress. CPU utilization remains high as long as the operation is progress irrespective of the data size.

Stop EMS data collection

EMS data collection is enabled by default and runs every seven days after your installation date. You can disable data collection at any time by using the PowerShell cmdlet *Disable-SmDataCollectionEMS*.

Steps

1. From a PowerShell command line, establish a session with SnapCenter by entering *Open-SmConnection*.
2. Disable EMS data collection by entering *Disable-SmDataCollectionEms*.

Start EMS data collection

EMS data collection is enabled by default and is scheduled to run every seven days from the installation date. If you have disabled it, you can start EMS data collection again by using the *Enable-SmDataCollectionEMS* cmdlet.

The NetApp ONTAP event generate-autosupport-log permission has been granted to the storage virtual machine (SVM) user.

Steps

1. From a PowerShell command line, establish a session with SnapCenter by entering *Open-SmConnection*.
2. Enable EMS data collection by entering *Enable-SmDataCollectionEMS*.

Change EMS data collection schedule and target SVM

You can use PowerShell cmdlets to change the EMS data collection schedule or the target storage virtual machine (SVM).

Steps

1. From a PowerShell command line, to establish a session with SnapCenter, enter the *Open-SmConnection* cmdlet.
2. To change the EMS data collection target, enter the *Set-SmDataCollectionEmsTarget* cmdlet.
3. To change the EMS data collection schedule, enter the *Set-SmDataCollectionEmsSchedule* cmdlet.

Monitor EMS data collection status

You can monitor the status of your EMS data collection using several PowerShell cmdlets. You can get information about the schedule, storage virtual machine (SVM) target, and status.

Steps

1. From a PowerShell command line, establish a session with SnapCenter by entering *Open-SmConnection*.
2. Retrieve information about the EMS data collection schedule by entering *Get-SmDataCollectionEmsSchedule*.
3. Retrieve information about the EMS data collection status by entering *Get-SmDataCollectionEmsStatus*.
4. Retrieve information about the EMS data collection target by entering *Get-SmDataCollectionEmsTarget*.

Related information

The information regarding the parameters that can be used with the cmdlet and their descriptions can be obtained by running *Get-Help command_name*. Alternatively, you can also refer to the SnapCenter Software Cmdlet Reference Guide.

Upgrade SnapCenter Server and plug-ins

Configure SnapCenter to check for available updates

SnapCenter periodically communicates with the NetApp Support Site to notify you of available software updates. You can also create a schedule to specify the interval in which you want to receive information about available updates.

Steps

1. In the left navigation pane, click **Settings**.
2. In the **Settings** page, click **Software**.

   The Available Software page displays the available plug-in packages, versions available, and their installation status.

3. Click **Check for updates** to see if any newer versions of plug-in packages are available.
4. Click **Schedule Updates** to create a schedule to specify the interval in which you want to receive information about available updates:
   a. Select the interval in **Check for updates**.
   b. Select the SnapCenter Server Admin Windows credential and click **OK**.

Upgrade workflow

Each release of SnapCenter contains an updated SnapCenter Server and plug-in package. Plug-in package updates are distributed with the SnapCenter installer. You can configure SnapCenter to check for available updates.

The workflow shows the different tasks required to upgrade the SnapCenter Server and the Plug-in packages.

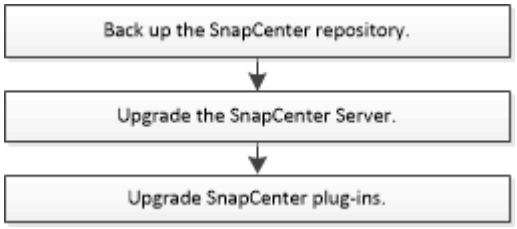

Supported upgrade paths

The upgrade path helps you to understand from which all earlier versions of SnapCenter you can upgrade to the latest versions of SnapCenter and which versions of the plug-ins are supported.

If you are on SnapCenter Server version…	You can directly upgrade SnapCenter Server to…	Supported plug-in versions
5.0	6.0	• 5.0 • 6.0
	6.0.1	• 6.0.1
	6.1	• 6.1
6.0	6.0.1	• 6.0 • 6.0.1
	6.1	• 6.1
6.0.1	6.1	• 6.0.1 • 6.1

For information about upgrading the SnapCenter Plug-in for VMware vSphere, see Upgrade SnapCenter Plug-in for VMware vSphere.

Upgrade the SnapCenter Server on Windows host

You should upgrade the SnapCenter Server to access the latest features and enhancements provided in the newest version.

Before you begin

- Update the SnapCenter Server host with the latest Windows updates and ensure there are no pending system restarts.
- Ensure no other operations are running before you start the upgrade.
- Install ASP.NET Core Runtime 8.0.12 (and all subsequent 8.0.x patches) Hosting Bundle and PowerShell 7.4.2 or later.
- Back up the SnapCenter repository (MySQL) database after ensuring that no jobs are running. This is recommended before upgrading SnapCenter Server and the Exchange plug-in.

  For information, see Back up the SnapCenter repository.

- Back up all the modified SnapCenter configuration files that you have modified either on the SnapCenter Server host or the plug-in host.

  Examples of SnapCenter configuration files: SnapDriveService.exe.config, SMCoreServiceHost.exe.config, and so on.

- If you have installed multiple versions of custom plug-in in SnapCenter 5.0, before upgrading to 6.0 or later, you should run the powershell cmdlets to remove all the earlier versions of the custom plug-in (except the latest one) from the SnapCenter repository (NSM database).
  - Run `Open-SmConnection` and login using the SnapCenterAdmin role credentials

- Run `Remove-SmPluginPackage -PluginName M<plug-in name> -PluginVersion <version number>`

  For more information, refer to Upgrade to SnapCenter 6.0 or later fails.

About this task

- During upgrade, SnapCenter executes a SQL script to update the Exchange data in the NSM database, converting the DAG and host shortname to FQDN. This applies only if you use SnapCenter Server with the Exchange plug-in.
- If you have manually placed the server host in maintenance mode, after the upgrade select **Hosts** > **Activate Schedule** to bring the server host out of maintenance mode.
- For the existing backup and verification schedules with prescripts and postscripts enabled in the policy, the backup operations will continue to work after upgrade.

  In the **Job details** page, a warning message recommends that the customer should copy the scripts to the SCRIPTS_PATH and edit the policy to provide a path that is relative to the SCRIPTS_PATH. For the clone lifecycle job, the warning message appears at the sub-job level.

Steps

1. Download the SnapCenter Server installation package from the NetApp Support Site.

   https://mysupport.netapp.com/site/products/all/details/snapcenter/downloads-tab

2. Create a copy of the web.config located at C:\Program Files\NetApp\SnapCenter WebApp.
3. Export the SnapCenter plug-in host schedules from the Windows Task Scheduler to restore them if the upgrade fails.

   ```
   md d:\\SCBackup``schtasks /query /xml /TN taskname >>
   "D:\SCBackup\taskname.xml"
   ```

4. Create the SnapCenter MySQL database dump if the repository backup is not configured.

   ```
   md d:\\SCBackup``mysqldump --all-databases --single-transaction --add-drop
   -database --triggers --routines --events -u root -p >
   D:\SCBackup\SCRepoBackup.dmp
   ```

   When prompted, enter the password.

5. Double-click the downloaded .exe file to initiate the SnapCenter Server upgrade.

   After you start the upgrade, SnapCenter runs prechecks. If the system does not meet the minimum requirements, SnapCenter shows error or warning messages. You can ignore warnings and continue with the installation, but you must fix any errors.

   SnapCenter continues to use the existing MySQL Server repository database password provided during installation of the earlier version of SnapCenter Server.

6. Select **Upgrade**.

   If you select **Cancel** at any stage, SnapCenter stops the upgrade. It will not rollback the SnapCenter Server to previous state.

Best Practice: Log out and log in again, or open a new browser to access the SnapCenter User Interface.

After you finish

- If the plug-in is installed using a sudo user, you should copy the sha224 keys available at *C:\ProgramData\NetApp\SnapCenter\Package Repository\sc_unix_plugins_checksum.txt* to update the */etc/sudoers* file.
- You should perform a fresh discovery of resources on the server host.

  If SnapCenter displays the server host status as stopped, wait for some time and perform a fresh discovery. You can also change the value of **HostRefreshInterval** parameter (default value is 3600 seconds) to any value more than 10 minutes.

- If the upgrade fails, clean up the failed installation, reinstall the earlier version of SnapCenter, and restore the NSM database to the state it was in before.
- After upgrading the server host, you must also upgrade the plug-ins before adding any storage system.

Upgrade the SnapCenter Server on Linux host

You can use the SnapCenter Server installer file to upgrade the SnapCenter Server.

Steps

1. Perform one of the actions to upgrade SnapCenter Server.

If you want to perform…	Do this…
Non interactive upgrade	`sudo ./snapcenter-linux-server-(el8/el9/sles15).bin -i silent -DUPGRADE=<value> -DINSTALL_LOG_NAME=<filename>` Example: sudo ./snapcenter_linux_server.bin -i silent -DUPGRADE=1 -DINSTALL_LOG_NAME=InstallerLog.log Logs will be stored at */var/opt/snapcenter/logs*. Parameters to be passed for upgrade: • DINSTALL_LOG_NAME: Log file name where installation logs will be stored. • DUPGRADE: The default value is 0. Specify this parameter and it's value as any integer other than 0 to upgrade the SnapCenter Server.

If you want to perform…	Do this…
Interactive installation	`./snapcenter-linux-server-(el8/el9/sles15).bin` You will be prompted for upgrade confirmation. Enter any value other than 0 to confirm SnapCenter Server upgrade.

You should either log out and then log into SnapCenter, or close and then open a new browser to access SnapCenter GUI.

Upgrade your plug-in packages

The plug-in packages are distributed as part of the SnapCenter upgrade.

You do not have to manually put each plug-in host that you want to upgrade into maintenance mode because the upgrade procedure places your Windows, Linux, or AIX plug-in hosts into maintenance mode. The maintenance mode prevents any scheduled jobs to be executed on the plug-in host during the upgrade.

Before you begin

- If you are a non-root user with access to the Linux machines, you should update the */etc/sudoers* file with the latest checksum values before performing the upgrade operation.
- By default SnapCenter detects JAVA_HOME from the environment. If you want to use a fixed JAVA_HOME and if you are upgrading the plug-ins on a Linux host, you should manually add the SKIP_JAVAHOME_UPDATE parameter in the *spl.properties* file located at */var/opt/snapcenter/spl/etc/* and set the value to TRUE.

  The value of JAVA_HOME gets updated when the plug-in is upgraded or when the SnapCenter plug-in loader (SPL) service restarts. Before upgrading or restarting the SPL, if you add the SKIP_JAVAHOME_UPDATE parameter and set the value to TRUE, the value of JAVA_HOME is not updated.

- You should have backed up all the SnapCenter configuration files that you have modified either on the SnapCenter Server host or the plug-in host.

  Examples of SnapCenter configuration files: SnapDriveService.exe.config, SMCoreServiceHost.exe.config, and so on.

About this task

- For SnapCenter Plug-in for Microsoft SQL Server, SnapCenter Plug-in for Microsoft Exchange Server, and SnapCenter Plug-in for Microsoft Windows, it is recommended to upgrade both the server and the plug-in hosts to the latest version for the SCRIPTS_PATH to run.

  For the existing backup and verification schedules with prescripts and postscripts enabled in the policy, the backup operations will continue to work after upgrade.

  In the **Job details** page, a warning message recommends that the customer should copy the scripts to the SCRIPTS_PATH and edit the policy to provide a path that is relative to the SCRIPTS_PATH. For the clone lifecycle job, the warning message appears at the sub-job level.

Steps

1. In the left navigation pane, click **Hosts** > **Managed Hosts**.
2. Upgrade the hosts by performing one of the following tasks:
   - If the Overall Status column displays "Upgrade available" for one of the plug-inhosts, click the plug-in host name and perform the following:
     a. Click **More Options**.
     b. Select **Skip prechecks** if you do not want to validate whether the plug-in host meets the requirements to upgrade the plug-in.
     c. Click **Upgrade**.
   - If you want to upgrade multiple hosts, select all the hosts, click , and then click **Upgrade** > **OK**.

     All the related services are restarted during the plug-in upgrade.

     > All the plug-ins in the package gets selected, but only the plug-ins that were installed with the earlier version of SnapCenter are upgraded, and the remaining plug-ins are not installed. You must use the **Add plug-ins** option to install any new plug-in.

     If you have not selected the **Skip prechecks** check box, the plug-in host is validated to see if it meets the requirements to install the plug-in. If the minimum requirements are not met, appropriate error or warning messages are displayed. After fixing the issue, click **Upgrade**.

     > If the error is related to disk space or RAM, you can update either the web.config located at C:\Program Files\NetApp\SnapCenter WebApp, or the PowerShell config files located at C:\Windows\System32\WindowsPowerShell\v1.0\Modules\SnapCenter\ to modify the default values. If the error is related to remaining parameters then you must fix the issue, and then validate the requirements again.

Tech refresh

Tech refresh of SnapCenter Server host

When the SnapCenter Server host requires refresh, you can install the same version of SnapCenter Server on the new host and then run the APIs to backup the SnapCenter from old server and restore it in on the new server.

Steps

1. Deploy the new host and perform the following tasks:
   a. Install the same version of the SnapCenter Server.
   b. (Optional) Configure CA certificates and enable two-way SSL. For more information, refer to Configure CA Certificate and Configure and enable two-way SSL.
   c. (Optional) Configure multi-factor authentication. For more information, refer to Enable multi-factor authentication.
2. Log in as the SnapCenter Admin user.
3. Create a backup of the SnapCenter Server on the old host using either the API: `/<snapcenter_version>/server/backup` or the cmdlet: *New-SmServerBackup*.

   Before taking the backup, suspend all the scheduled jobs and ensure that no jobs are running.

   If you want to restore the backup on the SnapCenter Server that is running on a new domain, before taking a backup you should add the new domain user in the old SnapCenter host and assign the SnapCenter admin role.

4. Copy the backup from the old host to new host.
5. Restore the backup of the SnapCenter Server on the new host using either the API: `/<snapcenter_version>/server/restore` or the cmdlet: *Restore-SmServerBackup*.

   Restore will update the new SnapCenter Server URL in all the hosts by default. If you want to skip the update, use the *-SkipSMSURLInHosts* attribute and separately update the server URL by running using either the API: `/<snapcenter_version>/server/configureurl` or the cmdlet: *Set-SmServerConfig*.

   If the plug-in host is not able to resolve the server hostname, log in to each of the plug-in host and add the *etc/host* entry for the new IP in the <New IP> SC_Server_Name format.

   The server *etc/host* entries will not be restored. You can restore it manually from the old server.

   If the backup is restored on the SnpCenter Server that is running on a new domain and if you want to continue to use the old domain users, you should register the old domain in the new SnapCenter Server.

   If you have manually updated the web.config file in old SnapCenter host, the updates will not be copied to the new host. You should manually make the same changes in the web.config file of the new host.

6. If you have skipped updating the SnapCenter Server URL or any of the host was down during the restore process, update the new server name in all the hosts or specified hosts that are managed by the SnapCenter using either the API: `/<snapcenter_version>/server/configureurl` or the cmdlet: *Set-SmServerConfig*.
7. Activate the scheduled jobs on all the hosts from the new SnapCenter Server.

Tech refresh of a node in F5 cluster

You can do tech refresh of any node in the F5 cluster by removing the node and adding the new node. If the node that needs to be refreshed is active, make another node of the cluster as active and then remove the node.

For information on how to add a node to F5 cluster, refer to Configure SnapCenter Servers for High Availability using F5.

If the url of the F5 cluster changes, the url can be updated in all the hosts using either the API: `/<snapcenter_version>/server/configureurl` or the cmdlet: *Set-SmServerConfig*.

Decommissioning the old SnapCenter Server host

You can remove the old SnapCenter Server host after verifying that the new SnapCenter Server is up and running and all the plug-in hosts are able to communicate with the the new SnapCenter Server host.

Rollback to the old SnapCenter Server host

In case of any issues, you can bring back the old SnapCenter Server host by updating the SnapCenter Server URL in all the hosts using either the API: `/<snapcenter_version>/server/configureurl` or the cmdlet: *Set-SmServerConfig*.

Disaster recovery

Disaster recovery of standalone SnapCenter host

You can perform disaster recovery by restoring the server backup to the new host.

Before you begin

Ensure that you have a backup of the old SnapCenter Server.

Steps

1. Deploy the new host and perform the following tasks:
   a. Install the same version of the SnapCenter Server.
   b. Configure CA certificates and enable two-way SSL. For more information, refer to Configure CA Certificate and Configure and enable two-way SSL.
2. Copy the old SnapCenter Server backup to the new host.
3. Log in as the SnapCenter Admin user.
4. Restore the backup of the SnapCenter Server on the new host using either the API: `/<snapcenter_version>/server/restore` or the cmdlet: *Restore-SmServerBackup*.

   Restore will update the new SnapCenter Server URL in all the hosts by default. If you want to skip the

update, use the *-SkipSMSURLInHosts* attribute and separately update the server URL by using either the API: `/<snapcenter_version>/server/configureurl` or the cmdlet: *Set-SmServerConfig*.

> If the plug-in host is not able to resolve the server hostname, log in to each of the plug-in host and add the *etc/host* entry for the new IP in the <New IP> SC_Server_Name format.

> The server *etc/host* entries will not be restored. You can restore it manually from the old server.

5. If you have skipped updating the URL or any of the host was down during the restore process, update the new server name in all the hosts or specified hosts that are managed by the SnapCenter using either the API: `/<snapcenter_version>/server/configureurl` or the cmdlet: *Set-SmServerConfig*.

Disaster recovery of SnapCenter F5 cluster

You can perform disaster recovery by restoring the server backup to the new host and then converting the standalone host to a cluster.

Before you begin

Ensure that you have a backup of the old SnapCenter Server.

Steps

1. Deploy the new host and perform the following tasks:
   a. Install the same version of the SnapCenter Server.
   b. Configure CA certificates and enable two-way SSL. For more information, refer to Configure CA Certificate and Configure and enable two-way SSL.
2. Copy the old SnapCenter Server backup to the new host.
3. Log in as the SnapCenter Admin user.
4. Restore the backup of the SnapCenter Server on the new host using either the API: `/<snapcenter_version>/server/restore` or the cmdlet: *Restore-SmServerBackup*.

   Restore will update the new SnapCenter Server URL in all the hosts by default. If you want to skip the update, use the *-SkipSMSURLInHosts* attribute and separately update the server URL by using either the API: `/<snapcenter_version>/server/configureurl` or the cmdlet: *Set-SmServerConfig*.

   > If the plug-in host is not able to resolve the server hostname, log in to each of the plug-in host and add the *etc/host* entry for the new IP in the <New IP> SC_Server_Name format.

   > The server *etc/host* entries will not be restored. You can restore it manually from the old server.

5. If you have skipped updating the URL or any of the host was down during the restore process, update the new server name in all the hosts or specified hosts that are managed by the SnapCenter using either the API: `/<snapcenter_version>/server/configureurl` or the cmdlet: *Set-SmServerConfig*.
6. Convert the standalone host to F5 cluster.

   For information on how to configure F5, refer to Configure SnapCenter Servers for High Availability using F5.

Related information

For information on the APIs, you need to access the Swagger page. see How to access REST APIs using the swagger API web page.

The information regarding the parameters that can be used with the cmdlet and their descriptions can be obtained by running *Get-Help command_name*. Alternatively, you can also refer the SnapCenter Software Cmdlet Reference Guide.

Tech refresh of SnapCenter plug-in hosts

When the SnapCenter plug-in hosts require refresh, you should move the resources from old host to new host. When the new host is added to SnapCenter, it will discover all the resources but will be treated as new resources.

About this task

You should run the API or cmdlet which will take old host name and new host name as input, compare the resources by name, and relink the objects of matching resources from old host to new host. The matching resources will be marked as protected.

- The *IsDryRun* parameter is set to True by default and this identifies the matching resources of the old and new host.

  After verifying the matching resources, you should set the *IsDryRun* parameter to False to relink the objects of the matching resources from the old host to new host.

- The *AutoMigrateManuallyAddedResources* parameter is set to True by default and this automatically copies the manually added resources from old host to the new host.

  The *AutoMigrateManuallyAddedResources* parameter is applicable only for Oracle and SAP HANA resources.

- The *SQLInstanceMapping* parameter should be used if the instance name is different between old host and new host. If it is a default instance then use *default_instance* as instance name.

Tech refresh is supported for the following SnapCenter Plug-ins:

- SnapCenter Plug-in for Microsoft SQL Server
  - If the SQL databases are protected at instance level and as part of host tech refresh only partial resources are moved to new host, then the existing instance level protection will be converted to resource group protection and instances from both the hosts will be added to the resource group.
  - If a SQL host (for example host1) is used as either scheduler or verification server for resources of another host (for example host2), then while performing tech refresh on host1, the schedule or the verification details will not be migrated and will continue to run on host1. If you have to modify, then you should manually change it in the respective hosts.
  - If you are using SQL Failover Cluster Instances (FCI) setup, you can perform the tech refresh by adding the new node to the FCI cluster and refreshing the plug-in host in SnapCenter.
  - If you are using SQL Availability Group (AG) setup, tech refresh is not required. You can add the new node to AG and refresh the host in SnapCenter.
- SnapCenter Plug-in for Windows
- SnapCenter Plug-in for Oracle Database

If you are using Oracle Real Application Cluster (RAC) setup, you can perform the tech refresh by adding the new node to the RAC cluster and refreshing the plug-in host in SnapCenter.

- SnapCenter Plug-in for SAP HANA Database

The supported use cases are:

- Migrating resources from one host to another host.
- Migrating resources from multiple hosts to one or fewer hosts.
- Migrating resources from one host to multiple hosts.

The supported scenarios are:

- New host has a different name from the old host
- Existing host has been renamed

Before you begin

As this workflow modifies the data in SnapCenter repository, it is recommended to backup the SnapCenter repository. Incase of any data issues, SnapCenter repository can be reverted to old state using the backup.

For more information, refer to Back up the SnapCenter repository.

Steps

1. Deploy the new host and install the application.
2. Suspend the schedules of the old host.
3. Move the required resources from the old host to the new host.
   a. Bring up the required databases in the new host from the same storage.
      - Ensure that the storage is mapped to the same drive or same mount path as that of old host. If the storage is not mapped correctly, backups created in old host cannot be used for restore.

        By default, Windows auto assigns the next available drive.

      - If storage DR is enabled, the respective storage should be mounted in the new host.
   b. Check for the compatibility if there is a change in application version.
   c. Only for Oracle plug-in host, ensure that the UIDs and GIDs of Oracle and it's group users are same as that of old host.

      For information, refer to:

      - How to migrate SQL database from old host to new host
      - How to migrate Oracle database from old host to new host
      - How to bring up SAP HANA database onto new host
4. Add the new host to SnapCenter.
5. Verify if all the resources are discovered.
6. Run the host refresh API: `/<snapcenter_version>/techrefresh/host` or the cmdlet: *Invoke-SmTechRefreshHost*.

The dry run is enabled by default and the matching resources to be relinked are identified. You can verify the resources by running either the API: '/jobs/{jobid}' or the cmdlet *Get-SmJobSummaryReport*.

If you have migrated the resources from multiple hosts, you should run the API or the cmdlet for all the hosts. If the drive or mount path in the new host is not same as the old host, following restore operations will fail:

- SQL in-place restore will fail. However, RTAL feature can be leveraged.
- Restore of Oracle and SAP HANA databases will fail.

If you want to migrate to multiple hosts, you should perform all the steps from step 1 for all the hosts.

You can run the API or cmdlet on the same host multiple times, it will relink only if there is a new resource identified.

7. (Optional) Remove the old host or hosts from SnapCenter.

Related information

For information on the APIs , you need to access the Swagger page. see How to access REST APIs using the swagger API web page.

The information regarding the parameters that can be used with the cmdlet and their descriptions can be obtained by running *Get-Help command_name*. Alternatively, you can also refer the SnapCenter Software Cmdlet Reference Guide.

Tech refresh of storage system

When the storage is tech refreshed, the data is migrated to new storage and the application hosts are mounted with new storage. The SnapCenter backup workflow identifies the new storage and creates the snapshot if the new storage is registered in SnapCenter.

You can perform restore, mount, and clone on the new backups created after storage refresh. However these operations will fail when performed on the backups that were created before storage refresh because the backups has the old storage details. You should run the storage tech refresh API or cmdlet to update the old backups in SnapCenter with the new storage details.

Tech refresh is supported for the following SnapCenter Plug-ins:

- SnapCenter Plug-in for Microsoft SQL Server
- SnapCenter Plug-in for Windows
- SnapCenter Plug-in for Oracle Database
- SnapCenter Plug-in for SAP HANA Database
- SnapCenter Plug-in for Microsoft Exchange Server

The supported use cases are:

- Primary storage refresh

The storage tech refresh is supported for replacing the primary storage with new storage. You cannot convert the existing secondary storage to a primary storage.

- Secondary storage refresh

Update the backups of the primary storage

When the storage is tech refreshed, you should run the storage tech refresh API or cmdlet to update the old backups in SnapCenter with the new storage details.

Before you begin

As this workflow modifies the data in SnapCenter repository, it is recommended to backup the SnapCenter repository. Incase of any data issues, SnapCenter repository can be reverted to old state using the backup.

For more information, refer to Back up the SnapCenter repository.

Steps

1. Migrate the data from old storage to new storage.

   For information on how to migrate, refer to:

   - How to migrate the data to new storage
   - How can I copy a volume and preserve all of the Snapshot copies?
2. Put the host to maintenance mode.
3. Mount the new storage in the respective hosts and bring up the databases.

   The new storage should be connected to host in the same way as before. For example, if it was connected as SAN, it needs to be connected as SAN.

   The new storage needs to be mounted on the same drive or path as that of the old storage.
4. Verify that all the resources are up and running.
5. Add the new storage in SnapCenter.

   Ensure that you have an unique SVM name across clusters in SnapCenter. If you are using the same SVM name in the new storage and if all the volumes of the SVM can be migrated before executing the storage refresh, then it is recommended to delete the SVM in old cluster and rediscover the old cluster in SnapCenter which will remove the SVM from cache.
6. Put the host in production mode.
7. In SnapCenter, create a backup of the resources whose storage is migrated. A new backup is necessary for SnapCenter to identify the latest storage footprint, and it will be used to update the metadata of existing old backups.

Whenever a new LUN is attached to host, it will have a new serial number. During discovery of Windows File System, SnapCenter will treat every unique serial number as new resource. During storage tech refresh when the LUN from new storage is attached to host with the same drive letter or path, the discovery of Windows File System in SnapCenter will mark the existing resource as deleted even if it is mounted with same drive letter or path and display the new LUN as new resource. As the resource is marked as deleted, it will not be considered for storage tech refresh in SnapCenter and all the backups of the old resource will be lost. When ever storage refresh happens, for Windows file system resources, resource discovery should not be performed before executing storage refresh API or cmdlet.

8. Run either the storage refresh API: `/<snapcenter_version>/techrefresh/primarystorage` or the cmdlet: *Invoke-SmTechRefreshPrimaryStorage*.

   If the resource is configured with a replication enabled policy, the latest backup after the storage refresh should have details of the secondary storage.

   a. If you are using SQL Failover Cluster Instances (FCI) setup, the backups are maintained at cluster level. You should provide the cluster name as input for storage tech refresh.
   b. If you are using SQL Availability Group (AG) setup, the backups are maintained at node level. You should provide the node name as input for storage tech refresh.
   c. If you are using Oracle Real Application Clusters (RAC) setup, you can perform storage tech refresh on any node.

      The *IsDryRun* attribute is set to True by default. It will identify the resources for which the storage is refreshed. You can view the resource and the changed storage details by running either the API: '<snapcenter_version>/jobs/{jobid}' or the cmdlet *Get-SmJobSummaryReport*.

9. After verifying the storage details, set the *IsDryRun* attribute to False and run the storage refresh API: `/<snapcenter_version>/techrefresh/primarystorage` or the cmdlet: *Invoke-SmTechRefreshPrimaryStorage*.

   This will update the storage details in the older backups.

   You can run the API or cmdlet on the same host multiple times, it will update the storage details in the older backups only if the storage is refreshed.

   The clone hierarchy cannot be migrated in ONTAP. If the storage being migrated has any clone metadata in SnapCenter, then the cloned resource will be marked as independent resource. Clones of clone metadata will be removed recursively.

10. (Optional) If all the snapshots are not moved from old primary storage to new primary storage, run the following API: `/<snapcenter_version>/hosts/primarybackupsexistencecheck` or the cmdlet *Invoke-SmPrimaryBackupsExistenceCheck*.

    This will perform the snapshot existence check on the new primary storage and mark the respective backups not available for any operation in SnapCenter.

Update the backups of the secondary storage

When the storage is tech refreshed, you should run the storage tech refresh API or cmdlet to update the old backups in SnapCenter with the new storage details.

Before you begin

As this workflow modifies the data in SnapCenter repository, it is recommended to backup the SnapCenter repository. Incase of any data issues, SnapCenter repository can be reverted to old state using the backup.

For more information, refer to Back up the SnapCenter repository.

Steps

1. Migrate the data from old storage to new storage.

   For information on how to migrate, refer to:

   - How to migrate the data to new storage
   - How can I copy a volume and preserve all of the Snapshot copies?

2. Establish the SnapMirror relationship between the primary storage and new secondary storage, and make sure relationship state is healthy.
3. In SnapCenter, create a backup of the resources whose storage is migrated.

   A new backup is necessary for SnapCenter to identify the latest storage footprint and it will be used to update the metadata of existing old backups.

   > You should wait until this operation is completed. If you proceed to the next step before completion, SnapCenter will loose old secondary snapshot metadata completely.

4. After successfully creating backup of all the resources in a host, run either the secondary storage refresh API: `/<snapcenter_version>/techrefresh/secondarystorage` or the cmdlet: *Invoke-SmTechRefreshSecondaryStorage*.

   This will update the secondary storage details of the older backups in the given host.

   If you want to run this at resource level, click **Refresh** for each resource to update the secondary storage metadata.

5. After successfully updating the older backups, you can break the old secondary storage relationship with primary.

Uninstall SnapCenter Server and plug-ins

Uninstall SnapCenter plug-in packages

Prerequisites for removing a host

You can remove hosts and uninstall individual plug-ins or plug-in packages using the SnapCenter GUI. You can also uninstall individual plug-ins or plug-in packages on remote hosts using the command-line interface (CLI) on your SnapCenter Server host or using the Windows **Uninstall a program** option locally on any host.

Before you remove a host from SnapCenter Server, you should complete the prerequisites.

- You should log in as an administrator.
- You should ensure that discovery jobs are not running on the host.
- You should be assigned a role with the required permissions to remove all of the objects associated with the host. Otherwise, the remove operation fails.
- You should confirm the fingerprint if the SSH key was modified after adding the host to SnapCenter.
- You should confirm the fingerprint if the SnapCenter host is upgraded to a later version of SnapCenter but the plug-in host is still running an earlier version of the plug-in.

Prerequisites to remove a host using role-based access control

- You should have logged in using an RBAC role that has read, delete host, installation, uninstallation of plug-in, and delete objects permissions.

  Objects can be clone, backup, resource group, storage system, and so on.

- You should have added the RBAC user to the RBAC role.
- You should assign the RBAC user to the host, plug-in, credential, resource groups, and storage system (for clones) that you want to delete.
- You should have logged in SnapCenter as an RBAC user.

Prerequisites to remove a host with clones created from clone lifecycle operation

- You should have created clone jobs using clone lifecycle management for SQL databases.
- You should have created an RBAC role with clone read and delete, resource read and delete, resource group read and delete, storage read and delete, provision read and delete, mount, unmount, plug-in installation and uninstallation, host read and delete permissions.
- You should have assigned the RBAC user to the RBAC role.
- You should have assigned the RBAC user to the host, SnapCenter Plug-in for Microsoft SQL Server, credential, clone lifecycle resource group, and storage system.
- You should have logged in SnapCenter as an RBAC user.

For information about uninstalling the SnapCenter Plug-in for VMware vSphere, see Remove SnapCenter Plug-in for VMware vSphere.

Remove a host

When the SnapCenter Server removes a host, it first removes the backup, clones, clone jobs, resource groups, and resources listed for that host on the SnapCenter Resources page, and then it uninstalls the plug-in packages on the host.

About this task

- If you delete a host, the backups, clones, and resource groups associated with the host are also deleted.
- When you remove the resource groups, all the associated schedules are also removed.
- If the host has a resource group that is shared with another host and you delete the host, then the resource group is also deleted.
- You should use the *Remove-SmHost* cmdlet to remove the decommissioned or unreachable plug-in hosts.

  The information regarding the parameters that can be used with the cmdlet and their descriptions can be obtained by running *Get-Help command_name*. Alternatively, you can also refer to the SnapCenter Software Cmdlet Reference Guide

- The time required to remove a host depends on the number of backups and the retention settings. This is because the Snapshots are deleted from each of the controllers and the metadata is cleaned.

Steps

1. In the left navigation pane, click **Hosts**.
2. In the **Hosts** page, click **Managed Hosts**.
3. Select the host you want to remove, and then click **Remove**.
4. For Oracle RAC clusters, to remove SnapCenter software from all the hosts in the cluster, select **Include all the hosts of cluster**.

   You can also remove one node of a cluster and in that way remove all the nodes one by one.

5. Click **OK**.

When you uninstall and reinstall host plug-ins on a cluster, the cluster resources are not automatically discovered. Select the cluster hostname, and then click **Refresh Resources** to automatically discover the cluster resources.

Uninstall plug-ins using the SnapCenter GUI

When you decide that you no longer require an individual plug-in or a plug-in package, you can uninstall it using the SnapCenter interface.

Before you begin

- You should have removed the resource groups for the plug-in package that you are uninstalling.
- You should have detached the policies associated with the resource groups for the plug-in package that you are uninstalling.

About this task

You can uninstall an individual plug-in. For example, you might need to uninstall the SnapCenter Plug-in for Microsoft SQL Server because a host is running out of resources and you want to move that plug-in to a more

powerful host. You can also uninstall an entire plug-in package. For example, you might need to uninstall the SnapCenter Plug-ins Package for Linux, which includes SnapCenter Plug-in for Oracle Database and SnapCenter Plug-in for UNIX.

- Removing a host includes uninstalling all plug-ins.

  When you remove a host from SnapCenter, SnapCenter uninstalls all the plug-in packages on the host before removing the host.

- SnapCenter GUI removes plug-ins from one host at a time.

  When you use the SnapCenter GUI, you can uninstall plug-ins on only one host at a time. However, you can have several uninstall operations running at the same time.

  You can also uninstall a plug-in from multiple hosts by using the *Uninstall-SmHostPackage* cmdlet and the required parameters. The information regarding the parameters that can be used with the cmdlet and their descriptions can be obtained by running *Get-Help command_name*. Alternatively, you can also refer to the SnapCenter Software Cmdlet Reference Guide.

> Uninstalling the SnapCenter Plug-ins Package for Windows from a host on which the SnapCenter Server is installed will damage the SnapCenter Server installation. Do not uninstall the SnapCenter Plug-ins Package for Windows unless you are certain that you no longer require the SnapCenter Server.

Steps

1. In the left navigation pane, click **Hosts**.
2. In the Hosts page, click **Managed Hosts**.
3. In the Managed Hosts page, select the host from which you want to uninstall the plug-in or plug-in package.
4. Adjacent to the plug-in that you want to remove, click **Remove** > **Submit**.

After you finish

You should wait for 5 minutes before you reinstall the plug-in on that host. This time period is sufficient for the SnapCenter GUI to refresh the status of the managed host. The installation fails if you immediately reinstall the plug-in.

If you are uninstalling SnapCenter Plug-ins Package for Linux, uninstallation-specific log files are available at: */custom_location/snapcenter/log*.

Uninstall Windows plug-ins using the PowerShell cmdlet

You can uninstall individual plug-ins or uninstall plug-ins packages from one or more hosts by using the *Uninstall-SmHostPackage* cmdlet on the SnapCenter Server host command-line interface.

You should have logged in to SnapCenter as a domain user with local administrator rights on each host on which you want to uninstall the plug-ins.

Steps

1. Launch PowerShell.

2. On the SnapCenter Server host, enter: *Open-SMConnection -SMSbaseUrl https://SNAPCENTER_SERVER_NAME/DOMAIN_NAME* command, and then enter your credentials.
3. Uninstall the Windows plug-ins using the *Uninstall-SmHostPackage* cmdlet and the required parameters.

   The information regarding the parameters that can be used with the cmdlet and their descriptions can be obtained by running *Get-Help command_name*. Alternatively, you can also refer the SnapCenter Software Cmdlet Reference Guide.

Uninstall plug-ins locally on a host

You can uninstall SnapCenter plug-ins locally on a host if you cannot reach the host from the SnapCenter Server.

About this task

The best practice for uninstalling individual plug-ins or plug-in packages is to either use the SnapCenter GUI or use the Uninstall-SmHostPackage cmdlet on the SnapCenter Server host command-line interface. These procedures help the SnapCenter Server to stay up to date with any changes.

However, you might have a rare need to uninstall plug-ins locally. For example, you might have run an uninstall job from the SnapCenter Server but the job failed, or you uninstalled your SnapCenter Server and orphan plug-ins remain on a host.

Uninstalling a plug-in package locally on a host does not delete data associated with the host; for example scheduled jobs and backup metadata.

Do not attempt to uninstall the SnapCenter Plug-ins Package for Windows locally from the Control Panel. You must use the SnapCenter GUI to ensure that SnapCenter Plug-in for Microsoft Windows is properly uninstalled.

Steps

1. On the host system, navigate to the Control Panel and click **Uninstall a program**.
2. In the list of programs, select the SnapCenter plug-in or plug-in package you want to uninstall and click **Uninstall**.

   Windows uninstalls all plug-ins in the selected package.

Uninstall plug-ins package for Linux or AIX using CLI

You can uninstall SnapCenter Plug-ins Package for Linux or SnapCenter Plug-ins Package for AIX by using the command-line interface.

Before you begin

- Ensure that you have deleted the scheduled jobs
- Ensure that all the running jobs are completed.

Step

Run */custom_location/NetApp/snapcenter/spl/installation/plugins/uninstall* to uninstall.

Uninstall the SnapCenter Server on Windows host

If you no longer wish to use the SnapCenter Server to manage data protection jobs, you can uninstall SnapCenter Server using the Programs and Features Control Panel on the SnapCenter Server host. Uninstalling the SnapCenter Server removes all its components.

Before you begin

- Ensure that you have at least 2 GB of free space on the drive where the SnapCenter Server is installed.
- Ensure that the domain in which the SnapCenter Server is installed is not removed.

  If you remove the domain where the SnapCenter Server was installed and then try to uninstall, the operation fails.

- You should have backed up the repository database because the repository database will be cleaned up and uninstalled.

Steps

1. On the SnapCenter Server host, navigate to the Control Panel.
2. Make sure you are in the **Category** view.
3. Under Programs, click **Uninstall a program**.

   Programs and Features window opens.

4. Select NetApp SnapCenter Server, and then click **Uninstall**.

   From SnapCenter 4.2, when you uninstall the SnapCenter Server, all it's components including the MySQL Server repository database is uninstalled.

   - Removing the NLB node from an NLB cluster requires that you restart the SnapCenter Server host. If you do not restart the host, you might experience a failure if you attempt to reinstall the SnapCenter Server.
   - You should manually uninstall .NET Framework which is not removed during uninstallation.

Uninstall the SnapCenter Server on Linux host

If you no longer wish to use the SnapCenter Server to manage data protection jobs, you can uninstall SnapCenter Server. Uninstalling the SnapCenter Server removes all its components.

Steps

1. Perform one of the actions to uninstall SnapCenter Server.

If you want to perform…	Do this…
Non interactive uninstallation	`$ sudo /opt/NetApp/snapcenter/SnapManagerWeb/installation/uninstall -i silent -DCONFIRM=1` Example: sudo /opt/NetApp/snapcenter/SnapManagerWeb/installation/uninstall
Interactive uninstallation	`$ sudo <USER_INSTALL_DIR>/NetApp/snapcenter/SnapManagerWeb/installation/uninstall` Enter any value other than 0 in the confirmation input to confirm uninstallation.

Automate using REST APIs

SnapCenter automation using REST APIs

You can use REST APIs to perform several SnapCenter management operations. REST APIs are exposed through the Swagger web page. You can access the Swagger web page available at *https://<SnapCenter_IP_address_or_name>:<SnapCenter_port>/swagger/* to display the REST API documentation, as well as to manually issue an API call.

The plug-ins that support REST APIs are:

- Plug-in for Microsoft SQL Server
- Plug-in for SAP HANA Database
- Plug-in for Oracle Database

For information on SnapCenter Plug-in for VMware vSphere REST APIs, see SnapCenter Plug-in for VMware vSphere REST APIs

How to access SnapCenter REST API natively

You can access the SnapCenter REST API directly using any programming language that supports a REST client. Popular language choices include Python, PowerShell, and Java.

REST web services foundation

Representational State Transfer (REST) is a style for creating distributed web applications. When applied to the design of a web services API, it establishes a set of technologies and best practices for exposing server-based resources and managing their states. It uses mainstream protocols and standards to provide a flexible foundation for managing SnapCenter.

Resources and state representation

Resources are the basic components of a web-based system. When creating a REST web services application, early design tasks include:

Identification of system or server-based resources

Every system uses and maintains resources. A resource can be a file, business transaction, process, or administrative entity. One of the first tasks in designing an application based on REST web services is to identify the resources.

Definition of resource states and associated state operations

Resources are always in one of a finite number of states. The states, as well as the associated operations used to affect the state changes, should be clearly defined.

URI endpoints

Every REST resource must be defined and made available using a well-defined addressing scheme. The endpoints where the resources are located and identified use a Uniform Resource Identifier (URI).

The URI provides a general framework for creating a unique name for each resource in the network. The Uniform Resource Locator (URL) is a type of URI used with web services to identify and access resources. Resources are typically exposed in a hierarchical structure similar to a file directory.

HTTP messages

Hypertext Transfer Protocol (HTTP) is the protocol used by the web services client and server to exchange request and response messages about the resources.

As part of designing a web services application, HTTP methods are mapped to the resources and corresponding state management actions. HTTP is stateless. Therefore, to associate a set of related requests and responses as part of one transaction, additional information must be included in the HTTP headers carried with the request and response data flows.

JSON formatting

While information can be structured and transferred between a web services client and server in several ways, the most popular option is JavaScript Object Notation (JSON).

JSON is an industry standard for representing simple data structures in plain text and is used to transfer state information describing the resources. The SnapCenter REST API uses JSON to format the data carried in the body of each HTTP request and response.

Basic operational characteristics

While REST establishes a common set of technologies and best practices, the details of each API can vary based on the design choices.

Request and response API transaction

Every REST API call is performed as an HTTP request to the SnapCenter Server system which generates an associated response to the client. This request and response pair is considered an API transaction.

Before using the API, you should be familiar with the input variables available to control a request and the contents of the response output.

Support for CRUD operations

Each of the resources available through the SnapCenter REST API is accessed based on the CRUD model:

- Create
- Read
- Update
- Delete

For some of the resources, only a subset of the operations is supported.

Object identifiers

Each resource instance or object is assigned a unique identifier when it is created. In most cases, the identifier is a 128-bit UUID. These identifiers are globally unique within a specific SnapCenter Server.

After issuing an API call that creates a new object instance, a URL with the associated ID is returned to the caller in the location header of the HTTP response. You can extract the identifier and use it on subsequent calls when referring to the resource instance.

The content and internal structure of the object identifiers can change at any time. You should only use the identifiers on the applicable API calls as needed when referring to the associated objects.

Object instances and collections

Depending on the resource path and HTTP method, an API call can apply to a specific object instance or a collection of objects.

Synchronous and asynchronous operations

SnapCenter performs an HTTP request received from a client either synchronously or asynchronously.

Synchronous processing

SnapCenter performs the request immediately and responds with an HTTP status code of 200 or 201 if it is successful.

Every request using the method GET is always performed synchronously. In addition, requests that use POST are designed to run synchronously if they are expected to complete in less than two seconds.

Asynchronous processing

If an asynchronous request is valid, SnapCenter creates a background task to process the request and a job object to anchor the task. The HTTP status code 202 is returned to the caller along with the job object. You should retrieve the state of the job to determine success or failure.

Requests that use the methods POST and DELETE are designed to run asynchronously if they are expected to take more than two seconds to complete.

Security

The security provided with the REST API is based primarily on the existing security features available with SnapCenter. The following security is used by the API:

Transport Layer Security

All traffic sent over the network between the SnapCenter Server and client is typically encrypted using TLS, based on the SnapCenter configuration settings.

HTTP authentication

At an HTTP level, basic authentication is used for the API transactions. An HTTP header with the user name and password in a base64 string is added to each request.

Input variables controlling an API request

You can control how an API call is processed through parameters and variables set in the HTTP request.

HTTP methods

The HTTP methods supported by the SnapCenter REST API are shown in the following table.

Not all the HTTP methods are available at each of the REST endpoints.

HTTP method	Description
GET	Retrieves object properties on a resource instance or collection.
POST	Creates a new resource instance based on the supplied input.
DELETE	Deletes an existing resource instance.
PUT	Modifies an existing resource instance.

Request headers

You should include several headers in the HTTP request.

Content-type

If the request body includes JSON, this header should be set to *application/json*.

Accept

This header should be set to *application/json*.

Authorization

Basic authentication should be set with the user name and password encoded as a base64 string.

Request body

The content of the request body varies depending on the specific call. The HTTP request body consists of one of the following:

- JSON object with input variables
- Empty

Filtering objects

When issuing an API call that uses GET, you can limit or filter the returned objects based on any attribute. For example, you can specify an exact value to match:

```
<field>=<query value>
```

In addition to an exact match, other operators are available to return a set of objects over a range of values. The SnapCenter REST API supports the filtering operators shown in the table below.

Operator	Description
=	Equal to
<	Less than
>	Greater than
⇐	Less than or equal to
>=	Greater than or equal to
UPDATE	Or
!	Not equal to
*	Greedy wildcard

You can also return a collection of objects based on whether a specific field is set or not set by using the **null** keyword or its negation **!null** as part of the query.

Any fields that are not set are generally excluded from matching queries.

Requesting specific object fields

By default, issuing an API call using GET returns only the attributes that uniquely identify the object or objects. This minimum set of fields acts as a key for each object and varies based on the object type. You can select additional object properties using the `fields` query parameter in the following ways:

Common or standard fields

Specify **fields=*** to retrieve the most commonly used object fields. These fields are typically maintained in local server memory or require little processing to access. These are the same properties returned for an object after using GET with a URL path key (UUID).

All fields

Specify **fields=**** to retrieve all the object fields, including those requiring additional server processing to access.

Custom field selection

Use **fields=<field_name>** to specify the exact field you want. When requesting multiple fields, the values must be separated using commas without spaces.

As a best practice, you should always identify the specific fields you want. You should only retrieve the set of common fields or all fields when needed. Which fields are classified as common, and returned using *fields=**, is determined by NetApp based on internal performance analysis. The classification of a field might change in future releases.

Sorting objects in the output set

The records in a resource collection are returned in the default order defined by the object. You can change the order using the `order_by` query parameter with the field name and sort direction as follows:

```
order_by=<field name> asc|desc
```

For example, you can sort the type field in descending order followed by id in ascending order:

```
order_by=type desc, id asc
```

- If you specify a sort field but do not provide a direction, the values are sorted in ascending order.
- When including multiple parameters, you must separate the fields with a comma.

Pagination when retrieving objects in a collection

When issuing an API call using GET to access a collection of objects of the same type, SnapCenter attempts to return as many objects as possible based on two constraints. You can control each of these constraints using additional query parameters on the request. The first constraint reached for a specific GET request terminates the request and therefore limits the number of records returned.

If a request ends before iterating over all the objects, the response contains the link needed to retrieve the next batch of records.

Limiting the number of objects

By default, SnapCenter returns a maximum of 10,000 objects for a GET request. You can change this limit using the *max_records* query parameter. For example:

```
max_records=20
```

The number of objects actually returned can be less than the maximum in effect, based on the related time constraint as well as the total number of objects in the system.

Limiting the time used to retrieve the objects

By default, SnapCenter returns as many objects as possible within the time allowed for the GET request. The default timeout is 15 seconds. You can change this limit using the *return_timeout* query parameter. For example:

```
return_timeout=5
```

The number of objects actually returned can be less than the maximum in effect, based on the related constraint on the number of objects as well as the total number of objects in the system.

Narrowing the result set

If needed, you can combine these two parameters with additional query parameters to narrow the result set. For example, the following returns up to 10 EMS events generated after the specified time:

```
time⇒ 2018-04-04T15:41:29.140265Z&max_records=10
```

You can issue multiple requests to page through the objects. Each subsequent API call should use a new time

value based on the latest event in the last result set.

Size properties

The input values used with some API calls as well as certain query parameters are numeric. Rather than provide an integer in bytes, you can optionally use a suffix as shown in the following table.

Suffix	Description
KB	KB Kilobytes (1024 bytes) or kibibytes
MB	MB Megabytes (KB x 1024 bytes) or mebibytes
GB	GB Gigabytes (MB x 1024 bytes) or gibibytes
TB	TB Terabytes (GB x 1024 byes) or tebibytes
PB	PB Petabytes (TB x 1024 byes) or pebibytes

Interpretation of an API response

Each API request generates a response back to the client. You should examine the response to determine whether it was successful and retrieve additional data as needed.

HTTP status code

The HTTP status codes used by the SnapCenter REST API are described below.

Code	Description
200	OK Indicates success for calls that do not create a new object.
201	Created An object is successfully created. The location header in the response includes the unique identifier for the object.
202	Accepted A background job has been started to perform the request, but has not completed yet.
400	Bad request The request input is not recognized or is inappropriate.
401	Unauthorized User authentication has failed.

Code	Description
403	Forbidden Access is denied due to an authorization (RBAC) error.
404	Not found The resource referred to in the request does not exist.
405	Method not allowed The HTTP method in the request is not supported for the resource.
409	Conflict An attempt to create an object failed because a different object must be created first or the requested object already exists.
500	Internal error A general internal error occurred at the server.

Response headers

Several headers are included in the HTTP response generated by the SnapCenter.

Location

When an object is created, the location header includes the complete URL to the new object including the unique identifier assigned to the object.

Content-type

This will normally be `application/json`.

Response body

The content of the response body resulting from an API request differs based on the object, processing type, and the success or failure of the request. The response is always rendered in JSON.

Single object

A single object can be returned with a set of fields based on the request. For example, you can use GET to retrieve selected properties of a cluster using the unique identifier.

Multiple objects

Multiple objects from a resource collection can be returned. In all cases, there is a consistent format used, with `num_records` indicating the number of records and records containing an array of the object instances. For example, you can retrieve the nodes defined in a specific cluster.

Job object

If an API call is processed asynchronously, a Job object is returned which anchors the background task. For example, the PATCH request used to update the cluster configuration is processed asynchronously and returns a Job object.

Error object

If an error occurs, an Error object is always returned. For example, you will receive an error when attempting to change a field not defined for a cluster.

Empty

In certain cases, no data is returned and the response body includes an empty JSON object.

Errors

If an error occurs, an error object is returned in the response body.

Format

An error object has the following format:

```
"error": {
"message": "<string>",
"code": <integer>[,
"target": "<string>"]
}
```

You can use the code value to determine the general error type or category, and the message to determine the specific error. When available, the target field includes the specific user input associated with the error.

Common error codes

The common error codes are described in the following table. Specific API calls can include additional error codes.

Code	Description
409	An object with the same identifier already exists.
400	The value for a field has an invalid value or is missing, or an extra field was provided.
400	The operation is not supported.
405	An object with the specified identifier cannot be not found.
403	Permission to perform the request is denied.
409	The resource is in use.

REST APIs supported for SnapCenter Server and plug-ins

The resources available through the SnapCenter REST API are organized in categories, as displayed on the SnapCenter API documentation page. A brief description of each of the resources with the base resource paths is presented below, along with additional usage considerations where appropriate.

Auth

You can use this API to log into the SnapCenter Server. This API returns a user authorization token that is used to authenticate subsequent requests.

Domains

You can use APIs to perform different operations.

- retrieve all the domains in SnapCenter
- retrieve details of a specific domain
- register or unregister a domain
- modify a domain

Jobs

You can use APIs to perform different operations.

- retrieve all the jobs in SnapCenter
- retrieve status of a job
- cancel or stop a job

Settings

You can use APIs to perform different operations.

- register, modify, or remove a credential
- displays the credential information registered in the SnapCenter Server
- configure notification settings
- retrieves information about the SMTP server currently configured to send email notifications and displays the name of the SMTP server, the name of the recipients, and the name of the sender
- displays multi-factor authentication (MFA) configuration of the SnapCenter Server login
- enable or disable and configure MFA for the SnapCenter Server login
- create the configuration file required to setup MFA

Hosts

You can use APIs to perform different operations.

- query all SnapCenter hosts

- remove one or more hosts from SnapCenter
- retrieve a host by name
- retrieve all resources on a host
- retrieve a resource using the resource ID
- retrieve the plug-in configuration details
- configure the plug-in host
- retrieve all resources of the plug-in for Microsoft SQL Server host
- retrieve all resources of the plug-in for Oracle database host
- retrieve all resources of the plug-in for custom application host
- retrieve all resources of the plug-in for SAP HANA host
- retrieve the plug-ins installed
- install plug-ins on an existing host
- upgrade host package
- remove plug-ins from an existing host
- add plug-in on a host
- add or modify host
- get the signature of the Linux host
- register the signature of the Linux host
- put the host to maintenance or production mode
- start or restart the plug-in services on the host
- rename a host

Resources

You can use APIs to perform different operations.

- retrieve all resources
- retrieve a resource using the resource ID
- retrieve all resources of the plug-in for Microsoft SQL Server host
- retrieve all resources of the plug-in for Oracle database host
- retrieve all resources of the plug-in for custom application host
- retrieve all resources of the plug-in for SAP HANA host
- retrieve a Microsoft SQL Server resource using a key
- retrieve a custom resource using a key
- modify a resource of the plug-in for custom application host
- remove a resource of the plug-in for custom application host using a key
- retrieve a SAP HANA resource using a key
- modify a resource of the plug-in for SAP HANA host
- remove a resource of the plug-in for SAP HANA host using a key

- retrieve an Oracle resource using a key
- create an Oracle application volume resource
- modify an Oracle application volume resource
- remove an Oracle application volume resource using a key
- retrieve the secondary details of the Oracle resource
- backup the Microsoft SQL Server resource using plug-in for Microsoft SQL Server
- backup the Oracle resource using plug-in for Oracle database
- backup the custom resource using plug-in for custom application
- configure the SAP HANA database
- configure the Oracle database
- restore a SQL database backup
- restore an Oracle database backup
- restore a custom application backup
- create a SAP HANA resource
- protect a custom resource using plug-in for custom application
- protect a Microsoft SQL Server resource using plug-in for Microsoft SQL Server
- modify a protected Microsoft SQL Server resource
- remove protection for Microsoft SQL Server resource
- protect an Oracle resource using plug-in for Oracle database
- modify a protected Oracle resource
- remove protection from Oracle resource
- clone a resource from the backup using plug-in for custom application
- clone an Oracle application volume from the backup using plug-in for Oracle database
- clone a Microsoft SQL Server resource from the backup using plug-in for Microsoft SQL Server
- create a clone life cycle of a Microsoft SQL Server resource
- modify clone life cycle of a Microsoft SQL Server resource
- delete clone life cycle of a Microsoft SQL Server resource
- move an existing Microsoft SQL Server database from a local disk to a NetApp LUN
- create a clone specification file for an Oracle database
- initiate an on-demand clone refresh job of an Oracle resource
- create an Oracle resource from the backup using the clone specification file
- restores the database to the secondary replica and joins the database back to the availability group
- create an Oracle application volume resource

Backups

You can use APIs to perform different operations.

- retrieve backup details by backup name, type, plug-in, resource, or date

- retrieve all backups
- retrieve backup details
- rename or delete backups
- mount an Oracle backup
- unmount an Oracle backup
- catalog an Oracle backup
- uncatalog an Oracle backup
- get all the backups required to be mounted to perform point-in-time recovery

Clones

You can use APIs to perform different operations.

- create, display, modify, and delete Oracle database clone specification file
- display Oracle database clone hierarchy
- retrieve clone details
- retrieve all clones
- delete clones
- retrieve clone details by ID
- initiate an on-demand clone refresh job of an Oracle resource
- clone an Oracle resource from the backup using the clone specification file

Clone split

You can use APIs to perform different operations.

- estimate the clone split operation of the cloned resource
- retrieve the status of a clone split operation
- start or stop a clone split operation

Resource Groups

You can use APIs to perform different operations.

- retrieve details of all resource groups
- retrieve the resource group by name
- create a resource group for plug-in for custom application
- create a resource group for plug-in for Microsoft SQL Server
- create a resource group for plug-in for Oracle database
- modify a resource group for plug-in for custom application
- modify a resource group for plug-in for Microsoft SQL Server
- modify a resource group for plug-in for Oracle database

- create, modify, or delete clone life cycle of a resource group for plug-in for Microsoft SQL Server
- back up a resource group
- put the resource group to maintenance or production mode
- remove a resource group

Policies

You can use APIs to perform different operations.

- retrieve policy details
- retrieve policy details by name
- delete a policy
- create a copy of an existing policy
- create or modify policy for plug-in for custom application
- create or modify policy for plug-in for Microsoft SQL Server
- create or modify policy for for plug-in for Oracle database
- create or modify policy for plug-in for SAP HANA database

Storage

You can use APIs to perform different operations.

- retrieve all the shares
- retrieve a share by name
- create or delete a share
- retrieve storage details
- retrieve storage details by name
- create, modify, or delete a storage
- discover resources on a storage cluster
- retrieve resources on a storage cluster

Share

You can use APIs to perform different operations.

- retrieve the details of a share
- retrieve details of all the shares
- create or delete a share on the storage
- retrieve a share by name

Plugins

You can use APIs to perform different operations.

- list all the plug-ins for a host
- retrieve a Microsoft SQL Server resource using a key
- modify a custom resource using a key
- remove a custom resource using a key
- retrieve a SAP HANA resource using a key
- modify a SAP HANA resource using a key
- remove a SAP HANA resource using a key
- retrieve an Oracle resource using a key
- modify an Oracle application volume resource using a key
- remove an Oracle application volume resource using a key
- backup the Microsoft SQL Server resource using plug-in for Microsoft SQL Server and a key
- backup the Oracle resource using plug-in for Oracle database and a key
- backup the custom application resource using plug-in for custom application and a key
- configure the SAP HANA database using a key
- configure the Oracle database using a key
- restore a custom application backup using a key
- create a SAP HANA resource
- create an Oracle application volume resource
- protect a custom resource using plug-in for custom application
- protect a Microsoft SQL Server resource using plug-in for Microsoft SQL Server
- modify a protected Microsoft SQL Server resource
- remove protection for Microsoft SQL Server resource
- protect an Oracle resource using plug-in for Oracle database
- modify a protected Oracle resource
- remove protection from Oracle resource
- clone a resource from the backup using plug-in for custom application
- clone an Oracle application volume from the backup using plug-in for Oracle database
- clone a Microsoft SQL Server resource from the backup using plug-in for Microsoft SQL Server
- create a clone life cycle of a Microsoft SQL Server resource
- modify clone life cycle of a Microsoft SQL Server resource
- delete clone life cycle of a Microsoft SQL Server resource
- create a clone specification file for an Oracle database
- initiate an on-demand clone life cycle of an Oracle resource
- clone an Oracle resource from the backup using the clone specification file

Reports

You can use APIs to perform different operations.

- retrieve reports of backup, restore, and clone operations for respective plug-ins
- add, run, delete, or modify schedules
- retrieve data for the scheduled reports

Alerts

You can use APIs to perform different operations.

- retrieve all the alerts
- retrieve alerts by IDs
- delete multiple alerts or delete an alert by ID

Rbac

You can use APIs to perform different operations.

- retrieve details of users, groups, and roles
- add or delete users
- assign user to role
- unassign user from role
- create, modify, or delete roles
- assign group to a role
- unassign group from a role
- add or delete groups
- create a copy of an existing role
- assign or unassign resources to user or group

Configuration

You can use APIs to perform different operations.

- view the configuration settings
- modify the configuration settings

CertificateSettings

You can use APIs to perform different operations.

- view the certificate status for the SnapCenter Server or plug-in host
- modify the certificate settings for the SnapCenter Server or plug-in host

Repository

You can use APIs to perform different operations.

- retrieve the repository backups

- view the configuration information about the repository
- protect and restore the SnapCenter repository
- unprotect the SnapCenter repository
- rebuild and failover the repository

Version

You can use this API to view the SnapCenter version.

How to access REST APIs using the Swagger API web page

REST APIs are exposed through the Swagger web page. You can access the Swagger web page to display the SnapCenter Server REST APIs, as well as to manually issue an API call. You can use REST APIs to help manage your SnapCenter Server or to perform data protection operations.

You should know the management IP address or domain name of the SnapCenter Server on which you want to execute the REST APIs.

You do not need special permissions to run the REST API client. Any user can access the Swagger web page. The respective permissions on the objects that are accessed via the REST API are based on the user who generates the token to login to the REST API.

Steps

1. From a browser, enter the URL to access the Swagger web page in the format *https://<SnapCenter_IP_address_or_name>:<SnapCenter_port>/swagger/*.

   Ensure that the REST API URL does not have the following characters: +, ., %, and &.

2. In the **Swagger Explore** field, if the Swagger API documentation does not display automatically, type: *https://<SnapCenter_IP_address_or_name>:<SnapCenter_port>/Content/swagger/SnapCenter.yaml*
3. Click **Explore**.

   A list of API resource types or categories are displayed.

4. Click an API resource type to display the APIs in that resource type.

If you encounter unexpected behavior when executing SnapCenter REST APIs, you can use the log files to identify the cause and resolve the problem.
You can download the log files from the SnapCenter user interface by clicking **Monitor** > **Logs** > **Download**.

Get started with the REST API

You can quickly get started using the SnapCenter REST API. Accessing the API provides some perspective before you begin using it with the more complex workflow processes on a live setup.

Hello World

You can run a simple command on your system to get started using the SnapCenter REST API and confirm its availability.

Before you begin

- Ensure that the Curl utility is available on your system.
- IP address or host name of the SnapCenter Server
- User name and password for an account with authority to access the SnapCenter REST API.

> If your credentials include special characters, you need to format them in a way that is acceptable to Curl based on the shell you are using. For example, you can insert a backslash before each special character or wrap the entire `username:password` string in single quotes.

Step

At the command line interface, run the following to retrieve the plug-in information:

```
curl -X GET -u username:password -k
"https://<ip_address>/api/hosts?fields=IncludePluginInfo"
```

Example:

```
curl -X GET -u admin:password -k
"'https://10.225.87.97/api/hosts?fields=IncludePluginInfo"
```

Legal notices

Legal notices provide access to copyright statements, trademarks, patents, and more.

Copyright

https://www.netapp.com/company/legal/copyright/

Trademarks

NETAPP, the NETAPP logo, and the marks listed on the NetApp Trademarks page are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.

https://www.netapp.com/company/legal/trademarks/

Patents

A current list of NetApp owned patents can be found at:

https://www.netapp.com/pdf.html?item=/media/11887-patentspage.pdf

Privacy policy

https://www.netapp.com/company/legal/privacy-policy/

Open source

Notice files provide information about third-party copyright and licenses used in NetApp software.

Notice for SnapCenter 6.1

Copyright information

Copyright © 2025 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP “AS IS” AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

Trademark information

NETAPP, the NETAPP logo, and the marks listed at http://www.netapp.com/TM are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.