Active IQ Unified Manager Documentation

Active IQ Unified Manager

NetApp
November 13, 2026

This PDF was generated from https://docs.netapp.com/us-en/active-iq-unified-manager-918/index.html on November 13, 2026. Always check docs.netapp.com for the latest.

Table of Contents

Active IQ Unified Manager Documentation . . . 1
Release notes for Active IQ Unified Manager . . . 2
Get started . . . 3
Get started with Active IQ Unified Manager on VMware vSphere . . . 3
System requirements . . . 3
Install Active IQ Unified Manager . . . 3
Get started with Active IQ Unified Manager on Linux . . . 4
System requirements . . . 4
Install Active IQ Unified Manager . . . 4
Get started with Active IQ Unified Manager on Windows . . . 5
System requirements . . . 5
Install Active IQ Unified Manager . . . 6
Install Unified Manager on VMware vSphere systems . . . 7
Learn about Active IQ Unified Manager . . . 7
Learn about Active IQ Unified Manager . . . 7
Learn what the Active IQ Unified Manager server does . . . 7
Learn about the Active IQ Unified Manager installation sequence . . . 7
Requirements for installing Unified Manager . . . 8
Ensure the server meets Active IQ Unified Manager installation requirements . . . 8
Learn about virtual infrastructure and hardware requirements for Active IQ Unified Manager . . . 8
Learn about VMware software and installation requirements for Active IQ Unified Manager . . . 9
Use supported browsers in Active IQ Unified Manager . . . 10
Learn about protocol and port requirements for Active IQ Unified Manager . . . 10
Complete the worksheet in Active IQ Unified Manager . . . 14
Installation, upgrade, and removal of Unified Manager on VMware vSphere . . . 15
Learn about the Active IQ Unified Manager deployment process . . . 15
Unified Manager virtual appliance deployment workflow . . . 16
Unified Manager virtual appliance upgrade workflow . . . 20
Restart the Active IQ Unified Manager virtual machine . . . 22
Remove Active IQ Unified Manager . . . 23
Install Unified Manager on Linux systems . . . 24
Introduction to Active IQ Unified Manager . . . 24
Learn about Active IQ Unified Manager . . . 24
Understand what the Active IQ Unified Manager server does . . . 24
View the installation sequence for Active IQ Unified Manager . . . 24
Requirements for installing Unified Manager . . . 25
Understand the requirements for installing Active IQ Unified Manager . . . 25
Understand virtual infrastructure and hardware requirements for Active IQ Unified Manager . . . 25
Learn about Linux software and installation requirements in Active IQ Unified Manager . . . 27
Use supported browsers for Active IQ Unified Manager . . . 29
Learn about protocol and port requirements in Active IQ Unified Manager . . . 30
Complete the worksheet in Active IQ Unified Manager . . . 32
Install, upgrade, and remove Unified Manager software . . . 34

Install, upgrade, and remove Active IQ Unified Manager software . . . 34
Understand the installation sequence for Active IQ Unified Manager on Red Hat Enterprise Linux . . . 34
Set up required software . . . 35
Configure SELinux for NFS and CIFS shares in Active IQ Unified Manager . . . 38
Install Unified Manager on Linux systems . . . 41
Upgrade Unified Manager on Red Hat Enterprise Linux . . . 47
Upgrade third-party products post Unified Manager installation . . . 52
Restart Active IQ Unified Manager . . . 53
Remove Active IQ Unified Manager . . . 53
Remove the custom umadmin user and maintenance group for Active IQ Unified Manager . . . 54
Install Unified Manager on Windows systems . . . 55
Introduction to Active IQ Unified Manager . . . 55
Learn about Active IQ Unified Manager . . . 55
Learn what the Active IQ Unified Manager server does . . . 55
Learn about the installation sequence for Active IQ Unified Manager . . . 55
Requirements for installing Unified Manager . . . 56
Ensure the server meets Active IQ Unified Manager installation requirements . . . 56
Understand virtual infrastructure and hardware requirements for Active IQ Unified Manager . . . 56
Understand Windows software and installation requirements for Active IQ Unified Manager . . . 58
Use supported browsers in Active IQ Unified Manager . . . 59
Understand protocol and port requirements for Active IQ Unified Manager . . . 59
Complete the worksheet in Active IQ Unified Manager . . . 62
Install, upgrade, and remove Unified Manager software . . . 64
Learn about the installation process for Active IQ Unified Manager on Windows . . . 64
Install Unified Manager on Windows . . . 64
Change the JBoss password on Windows in Active IQ Unified Manager . . . 67
Use the supported upgrade path for versions of Active IQ Unified Manager . . . 68
Upgrade Active IQ Unified Manager on Windows . . . 68
Upgrade third-party products . . . 70
Uninstall third-party products in Active IQ Unified Manager . . . 71
Restart Active IQ Unified Manager on Windows . . . 71
Uninstall Active IQ Unified Manager . . . 72
Perform configuration and administrative tasks . . . 74
Configure Active IQ Unified Manager . . . 74
Learn about configuring Active IQ Unified Manager . . . 74
Learn about the configuration sequence in Active IQ Unified Manager . . . 74
Access the web UI in Active IQ Unified Manager . . . 74
Perform the initial setup of the Active IQ Unified Manager web UI . . . 75
Add clusters to Active IQ Unified Manager . . . 77
Configure notifications . . . 79
Change the local user password in Active IQ Unified Manager . . . 87
Set the session inactivity timeout in Active IQ Unified Manager . . . 87
Set the session timeout using the CLI in Active IQ Unified Manager . . . 88
Change the Unified Manager host name . . . 88
Enable and disable policy-based storage management in Active IQ Unified Manager . . . 92

Configure Active IQ Unified Manager backup . . . 93
Manage feature settings . . . 93
Learn about feature settings in Active IQ Unified Manager . . . 93
Enable policy-based storage management in Active IQ Unified Manager . . . 93
Learn about the API Gateway feature in Active IQ Unified Manager . . . 94
Specify inactivity timeout in Active IQ Unified Manager . . . 94
Enable Active IQ portal events in Active IQ Unified Manager . . . 95
Enable and disable security settings for compliance in Active IQ Unified Manager . . . 95
Enable and disable script upload in Active IQ Unified Manager . . . 96
Add a login banner in Active IQ Unified Manager . . . 96
Use the maintenance console . . . 97
Learn about the maintenance console in Active IQ Unified Manager . . . 97
Learn about the maintenance user in Active IQ Unified Manager . . . 97
Diagnostic user capabilities in Active IQ Unified Manager . . . 97
Access the maintenance console in Active IQ Unified Manager . . . 97
Access the maintenance console using the vSphere VM console in Active IQ Unified Manager . . . 98
Maintenance console menus . . . 99
Change the maintenance user password on Windows in Active IQ Unified Manager . . . 104
Change the Active IQ Unified Manager umadmin password on Linux systems . . . 104
Change the ports Active IQ Unified Manager uses for HTTP and HTTPS protocols . . . 104
Add network interfaces to Active IQ Unified Manager . . . 105
Regenerate the client certificate in Active IQ Unified Manager . . . 106
Add disk space . . . 106
Manage user access . . . 110
Add users . . . 110
Edit user settings in Active IQ Unified Manager . . . 111
View users in Active IQ Unified Manager . . . 111
Delete users or groups in Active IQ Unified Manager . . . 112
Learn about role-based access control in Active IQ Unified Manager . . . 112
Learn about user types in Active IQ Unified Manager . . . 112
Learn about user roles in Active IQ Unified Manager . . . 113
Learn about user roles and capabilities in Active IQ Unified Manager . . . 114
Manage SAML authentication settings . . . 115
Enable SAML authentication in Active IQ Unified Manager . . . 115
Use an identity provider in Active IQ Unified Manager . . . 116
Enable SAML authentication in Active IQ Unified Manager . . . 117
Change the SAML authentication identity provider in Active IQ Unified Manager . . . 118
Update SAML authentication settings after an Active IQ Unified Manager security certificate change . . . 118
Disable SAML authentication in Active IQ Unified Manager . . . 119
Disable SAML authentication from the Active IQ Unified Manager maintenance console . . . 120
Use the SAML Authentication page in Active IQ Unified Manager . . . 121
Manage authentication . . . 121
Edit authentication servers in Active IQ Unified Manager . . . 122
Delete authentication servers in Active IQ Unified Manager . . . 122
Authenticate with Active Directory or OpenLDAP in Active IQ Unified Manager . . . 122

Audit Logging . . . 123
Use the Remote Authentication page in Active IQ Unified Manager . . . 125
Manage security certificates . . . 128
View the HTTPS security certificate in in Active IQ Unified Manager . . . 128
Download a HTTPS certificate signing request in Active IQ Unified Manager . . . 129
Install a CA-signed HTTPS certificate in Active IQ Unified Manager . . . 129
Learn about HTTPS certificates generated using external tools in Active IQ Unified Manager Unified Manager . . . 130
Page descriptions for certificate management . . . 132
Monitor and manage storage . . . 135
Introduction to Active IQ Unified Manager . . . 135
Learn about Active IQ Unified Manager . . . 135
Learn about health monitoring in Active IQ Unified Manager . . . 135
Learn about performance monitoring in Active IQ Unified Manager . . . 136
Use Active IQ Unified Manager REST APIs . . . 137
What the Unified Manager server does . . . 137
Understand the user interface . . . 138
Learn about the Active IQ Unified Manager user interface . . . 138
Typical window layouts in Active IQ Unified Manager . . . 138
Customizing window layout in Active IQ Unified Manager . . . 139
Use the Active IQ Unified Manager Help . . . 140
Bookmark your favorite Help topics in Active IQ Unified Manager . . . 141
Search for storage objects in Active IQ Unified Manager . . . 141
Export storage data as reports in Active IQ Unified Manager . . . 142
Filter inventory page content in Active IQ Unified Manager . . . 143
View active events from the notification bell in Active IQ Unified Manager . . . 144
Monitor and manage clusters from the dashboard . . . 145
Monitor and manage clusters from the dashboard in Active IQ Unified Manager . . . 145
Learn about the Active IQ Unified Manager Dashboard page . . . 146
Fix or manage ONTAP issues directly from Unified Manager . . . 149
Manage clusters . . . 155
Learn how the cluster discovery process works in Active IQ Unified Manager . . . 155
View the list of clusters monitored by Active IQ Unified Manager . . . 156
Add clusters to Active IQ Unified Manager . . . 156
Edit cluster settings in Active IQ Unified Manager . . . 158
Remove clusters from Active IQ Unified Manager . . . 159
Rediscover clusters in Active IQ Unified Manager . . . 160
Monitor VMware virtual infrastructure . . . 160
Understand how Active IQ Unified Manager monitors VMware virtual infrastructure . . . 160
View and add vCenter Servers in Active IQ Unified Manager . . . 164
Remove a vCenter Server from Active IQ Unified Manager . . . 166
Understand how Active IQ Unified Manager monitors virtual machines . . . 166
View virtual infrastructure visibility in a disaster recovery setup with Active IQ Unified Manager . . . 168
Provision and manage workloads . . . 170
Provision and manage workloads in Active IQ Unified Manager . . . 170

Workloads overview . . . 170
Manage Performance Service Levels . . . 177
Manage Storage Efficiency Policies . . . 183
Manage and monitor MetroCluster configurations . . . 184
Learn about volume behavior during switchover and switchback in Active IQ Unified Manager . . . 185
Understand cluster connectivity in a MetroCluster over FC configuration in Active IQ Unified Manager . . . 187
Learn about MetroCluster over FC data mirroring in Active IQ Unified Manager . . . 187
Monitor MetroCluster configurations in Active IQ Unified Manager . . . 188
Monitor MetroCluster replication in Active IQ Unified Manager . . . 191
Manage quotas . . . 192
Learn about quota limits in Active IQ Unified Manager . . . 192
View user and user group quotas in Active IQ Unified Manager . . . 192
Create rules to generate email addresses in Active IQ Unified Manager . . . 192
Create an email notification format for user and user group quotas in Active IQ Unified Manager . . . 193
Edit user and group quota email addresses in Active IQ Unified Manager . . . 193
Understand more about quotas . . . 194
Description of quotas dialog boxes . . . 195
Troubleshoot . . . 198
Add disk space to the Unified Manager database directory . . . 198
Change the performance statistics collection interval in Active IQ Unified Manager . . . 201
Change the length of time Active IQ Unified Manager retains event and performance data . . . 202
Resolve an unknown authentication error in Active IQ Unified Manager . . . 203
Resolve a User not found error in Active IQ Unified Manager . . . 203
Resolve the issue caused by adding LDAP using Other authentication service in Active IQ Unified Manager . . . 204
Resolve a NetApp Manageability SDK log rotation issue in Active IQ Unified Manager . . . 204
Manage events and alerts . . . 206
Manage events . . . 206
Learn about Active IQ platform events in Active IQ Unified Manager . . . 206
What Event Management System events are . . . 206
Understand how Active IQ Unified Manager handles a received event . . . 211
View event details in Active IQ Unified Manager . . . 212
View unassigned events in Active IQ Unified Manager . . . 213
Acknowledge and resolve events in Active IQ Unified Manager . . . 213
Assign events to specific users in Active IQ Unified Manager . . . 214
Disable unwanted events in Active IQ Unified Manager . . . 215
Fix issues using automatic remediation in Active IQ Unified Manager . . . 215
Enable and disable Active IQ event reporting in Active IQ Unified Manager . . . 216
Upload a new Active IQ rules file in Active IQ Unified Manager . . . 217
How Active IQ Unified Manager generates Active IQ platform events . . . 218
Resolve Active IQ platform events in Active IQ Unified Manager . . . 218
Configure event retention settings in Active IQ Unified Manager . . . 219
What a Unified Manager maintenance window is . . . 219
Manage host system resource events in Active IQ Unified Manager . . . 221
Understand more about events . . . 222

List of events and severity types . . . 227
Description of event windows and dialog boxes . . . 273
Manage alerts . . . 285
Learn about alerts in Active IQ Unified Manager . . . 285
Information contained in an Active IQ Unified Manager alert email . . . 286
Add alerts . . . 286
Add alerts for performance events in Active IQ Unified Manager . . . 289
Test alerts in Active IQ Unified Manager . . . 290
Disable alerts for Resolved and Obsolete events in Active IQ Unified Manager . . . 290
Exclude disaster recovery destination volumes from generating alerts in Active IQ Unified Manager . . . 291
View alerts in Active IQ Unified Manager . . . 292
Edit alerts in Active IQ Unified Manager . . . 292
Delete alerts in Active IQ Unified Manager . . . 292
Description of alert windows and dialog boxes . . . 293
Manage scripts . . . 298
Manage scripts in Active IQ Unified Manager . . . 298
How scripts work with alerts in Active IQ Unified Manager . . . 299
Add scripts in Active IQ Unified Manager . . . 300
Delete scripts in Active IQ Unified Manager . . . 301
Test script execution in Active IQ Unified Manager . . . 301
Learn about supported Active IQ Unified Manager CLI commands . . . 301
Description of script windows and dialog boxes . . . 307
Monitor and manage cluster performance . . . 309
Introduction to Active IQ Unified Manager performance monitoring . . . 309
Learn about performance monitoring in Active IQ Unified Manager . . . 309
Explore Active IQ Unified Manager performance monitoring features . . . 309
Use Active IQ Unified Manager interfaces to manage storage system performance . . . 310
View cluster configuration and performance data in Active IQ Unified Manager . . . 310
Understand data continuity collection cycles in Active IQ Unified Manager . . . 311
Learn about timestamps in collected data and events in Active IQ Unified Manager . . . 312
Navigating performance workflows in the Unified Manager GUI . . . 312
Navigate performance workflows in the Active IQ Unified Manager GUI . . . 312
Log in to the Active IQ Unified Manager UI . . . 313
Graphical interface and navigational paths . . . 314
Search for storage objects in Active IQ Unified Manager . . . 319
Filter inventory page content in Active IQ Unified Manager . . . 320
Monitor cluster performance from the Dashboard . . . 322
Monitor cluster performance from the Dashboard in Active IQ Unified Manager . . . 322
Review performance panels on the Active IQ Unified Manager Dashboard . . . 322
Interpret performance banner messages in Active IQ Unified Manager . . . 323
Change the performance statistics collection interval in Active IQ Unified Manager . . . 323
Troubleshoot workloads using the workload analyzer . . . 324
Troubleshoot workloads with the workload analyzer in Active IQ Unified Manager . . . 324
Understand workload analyzer data in Active IQ Unified Manager . . . 324
Understand when to use the workload analyzer in Active IQ Unified Manager . . . 326

Use the workload analyzer in Active IQ Unified Manager . . . 326
Monitor cluster performance from the Performance Cluster Landing page . . . 327
Monitor cluster performance from the Performance Cluster Landing page in Active IQ Unified Manager . . . 327
Understand the Performance Cluster Landing page in Active IQ Unified Manager . . . 327
Performance Cluster Landing page . . . 328
Monitor performance using the Performance Inventory pages . . . 332
Monitor object performance using Performance Inventory pages in Active IQ Unified Manager . . . 332
View Performance inventory pages for all storage objects in Active IQ Unified Manager . . . 333
Refine Performance inventory page contents . . . 338
Review Active IQ Unified Manager recommendations for tiering data to the cloud . . . 341
Monitor performance using the Performance Explorer pages . . . 342
Monitor performance using the Performance Explorer pages in Active IQ Unified Manager . . . 342
Use root objects for performance monitoring in Active IQ Unified Manager . . . 342
Filter correlated objects in the Active IQ Unified Manager grid . . . 343
Specify a time range for correlated objects . . . 343
Define correlated objects for comparison graphing in Active IQ Unified Manager . . . 344
Understand counter charts in Active IQ Unified Manager . . . 345
Understand the types of performance counter charts in Active IQ Unified Manager . . . 346
Select performance charts to display in Active IQ Unified Manager . . . 348
Expand the Counter Charts pane in Active IQ Unified Manager . . . 349
Change the Counter Charts focus to a shorter period in Active IQ Unified Manager . . . 349
View event details in the Active IQ Unified Manager Events Timeline . . . 350
Counter Charts Zoom View . . . 350
View volume latency by cluster component in Active IQ Unified Manager . . . 353
View SVM IOPS traffic by protocol in Active IQ Unified Manager . . . 353
View volume and LUN latency charts to verify performance guarantee in Active IQ Unified Manager . . . 354
View the performance for All SAN Array clusters in Active IQ Unified Manager . . . 354
View node IOPS based on local workloads in Active IQ Unified Manager . . . 355
Components of the Object Landing pages . . . 356
Manage performance using QoS policy group information . . . 361
Manage performance using QoS policy group information in Active IQ Unified Manager . . . 361
Control workload throughput with storage QoS in Active IQ Unified Manager . . . 361
View QoS policy groups available on all clusters in Active IQ Unified Manager . . . 362
View volumes or LUNs within the same QoS policy group in Active IQ Unified Manager . . . 363
View QoS policy group settings for volumes or LUNs in Active IQ Unified Manager . . . 363
Compare volumes or LUNs within the same QoS policy group in Active IQ Unified Manager . . . 364
View QoS policies in throughput charts with Active IQ Unified Manager . . . 365
View workload QoS minimum and maximum settings in Active IQ Unified Manager . . . 366
Manage performance using performance capacity and available IOPS information . . . 367
Manage performance capacity and available IOPS in Active IQ Unified Manager . . . 367
Understand the performance capacity used counter in Active IQ Unified Manager . . . 367
Use the performance capacity used value in Active IQ Unified Manager . . . 368
Understand available IOPS in Active IQ Unified Manager . . . 369
View node and aggregate performance capacity used values in Active IQ Unified Manager . . . 370

View available IOPS values for nodes and aggregates in Active IQ Unified Manager . . . 371
View performance capacity counter charts in Active IQ Unified Manager . . . 372
Configure performance capacity used thresholds in Active IQ Unified Manager . . . 373
Use the performance capacity used counter in Active IQ Unified Manager . . . 374
Understand and use the Node Failover Planning page . . . 374
Use the Node Failover Planning page in Active IQ Unified Manager . . . 375
Determine corrective actions from the Performance/Node Failover Planning page in Active IQ Unified Manager . . . 375
Review the components of the Performance/Node Failover Planning page in Active IQ Unified Manager . . . 375
Create threshold policies for node failover planning in Active IQ Unified Manager . . . 376
Use the Performance Capacity Used Breakdown chart for failover planning in Active IQ Unified Manager . . . 377
Collect data and monitor workload performance . . . 378
Collect data and monitor workload performance in Active IQ Unified Manager . . . 378
Monitor workload types with Active IQ Unified Manager . . . 379
Understand workload performance measurement values in Active IQ Unified Manager . . . 380
Understand how Active IQ Unified Manager calculates the expected performance range . . . 382
Use the latency forecast for performance analysis in Active IQ Unified Manager . . . 383
Use workload latency to identify performance issues in Active IQ Unified Manager . . . 384
Assess how cluster operations can affect workload latency in Active IQ Unified Manager . . . 385
Monitor MetroCluster performance with Active IQ Unified Manager . . . 385
Understand performance events and alerts . . . 386
Understand performance events and alerts in Active IQ Unified Manager . . . 386
Identify sources of performance events in Active IQ Unified Manager . . . 386
Prioritize performance events by severity in Active IQ Unified Manager . . . 387
Monitor configuration changes detected by Active IQ Unified Manager . . . 388
Understand system-defined performance threshold policies in Active IQ Unified Manager . . . 388
Analyze performance events and notifications in Active IQ Unified Manager . . . 391
Understand how Active IQ Unified Manager determines performance impact for an event . . . 393
Understand cluster components and contention in Active IQ Unified Manager . . . 393
Identify workload roles in performance events in Active IQ Unified Manager . . . 395
Manage performance thresholds . . . 396
Manage performance thresholds in Active IQ Unified Manager . . . 396
Use user-defined performance threshold policies in Active IQ Unified Manager . . . 397
Respond to breached performance threshold policies in Active IQ Unified Manager . . . 399
Understand performance counters tracked using thresholds in Active IQ Unified Manager . . . 399
Understand combination threshold policy objects and counters in Active IQ Unified Manager . . . 401
Create user-defined performance threshold policies in Active IQ Unified Manager . . . 402
Assign performance threshold policies to storage objects in Active IQ Unified Manager . . . 403
View performance threshold policies in Active IQ Unified Manager . . . 404
Edit user-defined performance threshold policies in Active IQ Unified Manager . . . 405
Remove performance threshold policies from storage objects in Active IQ Unified Manager . . . 405
Understand performance threshold policy changes in Active IQ Unified Manager . . . 406
Manage threshold policies when moving objects in Active IQ Unified Manager . . . 406

Analyze performance events . . . 407
Analyze performance events in Active IQ Unified Manager . . . 407
Display information about performance events in Active IQ Unified Manager . . . 407
Analyze events from user-defined performance thresholds . . . 408
Analyze events from system-defined performance thresholds . . . 409
Analyze events from dynamic performance thresholds . . . 415
Resolve performance events . . . 422
Resolve performance events in Active IQ Unified Manager . . . 422
Confirm that latency is within the expected range in Active IQ Unified Manager . . . 422
Assess how configuration changes affect workload performance in Active IQ Unified Manager . . . 422
Improve workload performance by adjusting client-side settings in Active IQ Unified Manager . . . 422
Understand how client or network issues affect performance in Active IQ Unified Manager . . . 423
Analyze high-activity volumes within a QoS policy group in Active IQ Unified Manager . . . 423
Move LIFs to improve performance in Active IQ Unified Manager . . . 424
Run storage efficiency operations at less busy times . . . 424
Understand how adding disks affects data reallocation in Active IQ Unified Manager . . . 425
Improve workload performance by enabling Flash Cache on a node in Active IQ Unified Manager . . . 426
Improve workload performance by enabling Flash Pool on a storage aggregate in Active IQ Unified Manager . . . 426
Check MetroCluster configuration health in Active IQ Unified Manager . . . 426
Verify MetroCluster configurations in Active IQ Unified Manager . . . 427
Move workloads to a different aggregate in Active IQ Unified Manager . . . 427
Move workloads to a different node in Active IQ Unified Manager . . . 428
Move workloads to an aggregate on a different node in Active IQ Unified Manager . . . 430
Move workloads to a node in a different HA pair in Active IQ Unified Manager . . . 431
Move workloads to another node in a different HA pair in Active IQ Unified Manager . . . 433
Use QoS policy settings to prioritize node workloads in Active IQ Unified Manager . . . 435
Remove inactive volumes and LUNs in Active IQ Unified Manager . . . 435
Understand how adding disks affects aggregate layout in Active IQ Unified Manager . . . 436
Set up a connection between a Unified Manager server and an external data provider . . . 436
Set up a connection between Active IQ Unified Manager and an external data provider . . . 436
Send Active IQ Unified Manager performance data to an external server . . . 436
Set up Graphite to receive performance data from Active IQ Unified Manager . . . 437
Configure a connection between Active IQ Unified Manager and an external data provider . . . 438
Monitor and manage cluster health . . . 441
Introduction to Active IQ Unified Manager health monitoring . . . 441
Learn about Active IQ Unified Manager health monitoring . . . 441
Understand Active IQ Unified Manager health monitoring features . . . 442
Use Active IQ Unified Manager interfaces to manage storage system health . . . 443
Manage and monitor clusters and cluster object health . . . 443
Understand cluster monitoring . . . 443
View the cluster list and details in Active IQ Unified Manager . . . 445
Check the health of clusters in a MetroCluster configuration in Active IQ Unified Manager . . . 446
View the health and capacity status of All SAN Array clusters in Active IQ Unified Manager . . . 448
View the node list and details in Active IQ Unified Manager . . . 448

Generate a hardware inventory report for contract renewal in Active IQ Unified Manager . . . 449
View the Storage VM list and details in Active IQ Unified Manager . . . 449
View the aggregate list and details in Active IQ Unified Manager . . . 450
View FabricPool capacity information in Active IQ Unified Manager . . . 451
View storage pool details in Active IQ Unified Manager . . . 452
View the volume list and details in Active IQ Unified Manager . . . 452
View details about NFS shares in Active IQ Unified Manager . . . 453
View details about SMB/CIFS shares in Active IQ Unified Manager . . . 454
View the list of Snapshot copies in Active IQ Unified Manager . . . 454
Delete Snapshot copies in Active IQ Unified Manager . . . 455
Calculate reclaimable space for Snapshot copies in Active IQ Unified Manager . . . 456
Understand cluster object windows and dialog boxes in Active IQ Unified Manager . . . 456
Common Unified Manager health workflows and tasks . . . 456
Understand common Active IQ Unified Manager health workflows and tasks . . . 457
Monitor and troubleshoot data availability . . . 458
Resolve capacity issues . . . 464
Manage health thresholds . . . 466
Manage cluster security objectives . . . 470
Manage backup and restore operations . . . 482
Manage scripts . . . 497
Manage and monitor groups . . . 500
Prioritize storage object events using annotations . . . 508
Send a support bundle through web UI and maintenance console . . . 516
Tasks and information related to several workflows . . . 523
Protect and restore data . . . 579
Create and troubleshoot protection relationships . . . 579
Types of SnapMirror protection in Active IQ Unified Manager . . . 579
Set up protection relationships in Unified Manager . . . 581
Perform a protection relationship failover and failback . . . 583
Resolve a protection job failure . . . 587
Resolve lag issues in Active IQ Unified Manager . . . 590
Manage and monitor protection relationships . . . 591
Learn about managing and monitoring protection relationships in Active IQ Unified Manager . . . 591
View volume protection status in Active IQ Unified Manager . . . 592
View volume protection relationships in Active IQ Unified Manager . . . 594
Monitor Consistency Group LUN relationships in Active IQ Unified Manager . . . 594
Create a SnapVault relationship from the Health: All Volumes view in Active IQ Unified Manager . . . 595
Create a SnapVault protection relationship from the Volume/Health details page in Active IQ Unified Manager . . . 596
Create a SnapMirror relationship from the Health: All Volumes view in Active IQ Unified Manager . . . 597
Create a SnapMirror relationship from the Volume Health details page in Active IQ Unified Manager . . . 598
Create a SnapMirror relationship with version-flexible replication in Active IQ Unified Manager . . . 599
Create SnapMirror relationships with version-flexible replication and back option capability in Active IQ Unified Manager . . . 600
Configure destination efficiency settings in Active IQ Unified Manager . . . 601

Create SnapMirror and SnapVault schedules in Active IQ Unified Manager . . . 602
Extend protection using cascade or fanout relationships in Active IQ Unified Manager . . . 603
Edit protection relationships from the Volume Relationships page in Active IQ Unified Manager . . . 603
Edit protection relationships from the Volume/ Health details page in Active IQ Unified Manager . . . 604
Create a SnapMirror policy to maximize transfer efficiency in Active IQ Unified Manager . . . 604
Create a SnapVault policy to maximize transfer efficiency in Active IQ Unified Manager . . . 605
Abort a transfer from the Volume Relationships page in Active IQ Unified Manager . . . 606
Abort a transfer from the Volume/Health details page in Active IQ Unified Manager . . . 606
Quiesce a protection relationship from the Volume Relationships page in Active IQ Unified Manager . . . 607
Quiesce a protection relationship from the Volume/ Health details page in Active IQ Unified Manager . . . 608
Break a SnapMirror relationship from the Active IQ Unified Manager Volume Relationships page . . . 609
Remove a protection relationship from the Volume Relationships page in Active IQ Unified Manager . . . 609
Resume scheduled transfers on a quiesced relationship from the Volume Relationships page in Active IQ Unified Manager . . . 610
Resume transfers on a quiesced relationship from the Volume / Health details page in Active IQ Unified Manager . . . 610
Initialize or update protection relationships from the Volume Relationships page in Active IQ Unified Manager . . . 611
Initialize or update protection relationships from the Volume/Health details page in Active IQ Unified Manager . . . 612
Resynchronize protection relationships from the Volume Relationships page in Active IQ Unified Manager . . . 613
Reverse protection relationships from the Volume Relationships page in Active IQ Unified Manager . . . 614
Restore data using the Volume and Volume/Health details pages in Active IQ Unified Manager . . . 615
Learn about resource pools in Active IQ Unified Manager . . . 616
Create resource pools in Active IQ Unified Manager . . . 616
Edit resource pools in Active IQ Unified Manager . . . 617
View resource pool inventory in Active IQ Unified Manager . . . 617
Add resource pool members in Active IQ Unified Manager . . . 617
Remove aggregates from resource pools in Active IQ Unified Manager . . . 618
Delete resource pools in Active IQ Unified Manager . . . 618
Monitor Storage VM Disaster Recovery protection relationships . . . 619
Manage storage VM associations in Active IQ Unified Manager . . . 621
SVM and resource pool requirements to support storage services in Active IQ Unified Manager . . . 623
Learn about jobs in Active IQ Unified Manager . . . 623
Monitor jobs in Active IQ Unified Manager . . . 624
View job details in Active IQ Unified Manager . . . 624
Abort jobs in Active IQ Unified Manager . . . 624
Retry a failed protection job in Active IQ Unified Manager . . . 625
Description of Protection relationships windows and dialog boxes . . . 625
Generate custom reports . . . 666
Unified Manager reporting . . . 666
Learn about Active IQ Unified Manager reporting . . . 666
Access points for generating reports . . . 666
Understand reports in Active IQ Unified Manager . . . 668

Understand the view and report relationship in Active IQ Unified Manager . . . 668
Learn about types of reports in Active IQ Unified Manager . . . 669
Understand reporting limitations in Active IQ Unified Manager . . . 670
Work with reports . . . 671
Understand report workflows in Active IQ Unified Manager . . . 671
Get started with reporting in Active IQ Unified Manager . . . 672
Search for a scheduled report in Active IQ Unified Manager . . . 675
Customize reports . . . 675
Download reports in Active IQ Unified Manager . . . 678
Schedule reports . . . 679
Schedule reports in Active IQ Unified Manager . . . 679
Schedule a report . . . 679
Manage report schedules . . . 681
Sample custom reports . . . 683
Understand sample custom reports in Active IQ Unified Manager . . . 683
Customize cluster storage reports . . . 684
Customize aggregate capacity reports . . . 686
Customize volume capacity reports . . . 688
Customize Qtree capacity reports . . . 691
Customize NFS share reports . . . 692
Customize storage VM reports . . . 694
Customize volume relationship reports . . . 695
Customize volume performance reports . . . 699
Sample Microsoft Excel reports . . . 699
Undestand sample Microsoft Excel reports in Active IQ Unified Manager . . . 700
Create a report to show an aggregate capacity table and chart in Active IQ Unified Manager . . . 700
Show aggregate total versus available capacity charts in Active IQ Unified Manager . . . 702
Create a report to view available volume capacity charts in Active IQ Unified Manager . . . 706
Create a report to view aggregates with the most available IOPS in Active IQ Unified Manager . . . 708
Manage storage using REST APIs . . . 711
Getting started with Active IQ Unified Manager . . . 711
Get started with Active IQ Unified Manager REST APIs . . . 711
Identify intended users of Active IQ Unified Manager REST APIs . . . 711
Manage Active IQ Unified Manager API access and categories . . . 711
Understand REST services offered in Active IQ Unified Manager . . . 713
Understand API version numbers in Active IQ Unified Manager . . . 713
Understand the ONTAP storage resource model in Active IQ Unified Manager . . . 714
REST API access and authentication in Active IQ Unified Manager . . . 715
Access and authenticate REST APIs in Active IQ Unified Manager . . . 715
Learn about Active IQ Unified Manager HTTP status codes . . . 718
Follow recommended practices for using Active IQ Unified Manager APIs . . . 719
Troubleshoot with system logs in Active IQ Unified Manager . . . 719
Understand how Active IQ Unified Manager uses Job objects for asynchronous processing . . . 720
Learn about the Hello API server sample program in Active IQ Unified Manager . . . 721
Unified Manager REST APIs . . . 725

Understand categories for Active IQ Unified Manager REST APIs . . . 725
Manage storage objects in a data center using APIs in Active IQ Unified Manager . . . 725
Access ONTAP APIs through proxy access in Active IQ Unified Manager . . . 730
Perform administrative tasks using APIs in Active IQ Unified Manager . . . 733
Manage users in Active IQ Unified Manager using APIs . . . 735
View performance metrics using APIs in Active IQ Unified Manager . . . 736
View jobs and system details using APIs in Active IQ Unified Manager . . . 744
Manage events and alerts using APIs in Active IQ Unified Manager . . . 745
Manage workloads . . . 748
Common workflows for storage management . . . 756
Understand common API workflows for storage management in Active IQ Unified Manager . . . 756
Determine space issues in aggregates using Active IQ Unified Manager APIs . . . 756
Determine storage object issues using Active IQ Unified Manager event APIs . . . 758
Troubleshoot ONTAP volumes using Active IQ Unified Manager gateway APIs . . . 759
Workflows for workload management . . . 762
Legal notices . . . 787
Copyright . . . 787
Trademarks . . . 787
Patents . . . 787
Privacy policy . . . 787
Open source . . . 787

Active IQ Unified Manager Documentation

Release notes for Active IQ Unified Manager

Provides a summary of new features, limitations, and known issues for Active IQ Unified Manager 9.18.

For more information, see the Active IQ Unified Manager Release Notes.

Get started

Get started with Active IQ Unified Manager on VMware vSphere

You can download the `.tar` file that contains a root certificate, a `README` file, and an `OVA` file, and deploy the Unified Manager as a virtual appliance.

System requirements

- Operating system: VMware ESXi 8.0
- RAM: 12 GB
- CPU: 9572 MHz total
- Free disk space: 5 GB (thin provisioning), 152 GB (thick provisioning)

For detailed system requirements, see the Requirements for installing Unified Manager and Interoperability Matrix.

Install Active IQ Unified Manager

Download the installer

1. Download the `.tar` file that contains a root certificate, a `README` file, and an `OVA` file.
2. Save the file to a local directory or network directory that is accessible to your vSphere Client.
3. In the directory where you downloaded the `.tar` file, enter the `tar -xvzf ActiveIQUnifiedManager-<version>.tar.gz` command.

   The required `OVA` file, a root certificate, and a `README` file are unzipped to the target directory.

Verify integrity

You can verify the integrity of the `OVA` file by following the steps provided in the `README` file.

Install Unified Manager

1. In the vSphere Client, click **File > Deploy OVF Template**.
2. Locate the OVA file and use the wizard to deploy the virtual appliance on the ESXi server.
3. On the Review Details page, in the Publisher section, the message `Entrust Code Signing - OVCS2 (Trusted certificate)` confirms the integrity of the downloaded `OVA` file. For the message `Entrust Code Signing - OVCS2 (Invalid certificate)`, upgrade the VMware vCenter Server to 7.0U3E or greater version.
4. On the Customize Template page, in the Properties tab, populate the fields as required for the type of installation you are performing:
   - For static configuration, enter the required information in all the fields. Adding information for the **Secondary DNS** field is not required.
   - For DHCP using IPv4, do not add any information in any field.

    - For DHCP using IPv6, check the “Enable Auto IPv6 addressing” box. Do not add any information in any other field.
5. Power on the VM.
6. Click the Console tab to view the initial boot process.
7. Configure the time zone.
8. Enter a Unified Manager maintenance user name and password.

At the end of the installation, the information to connect to the Unified Manager web UI is displayed.

Get started with Active IQ Unified Manager on Linux

You can download the installation package and install Unified Manager on a physical or virtual Red Hat Enterprise Linux platform.

System requirements

- Operating system: Red Hat Enterprise Linux versions 8.x and 9.4, based on x86_64 architecture, installed using the “Server with GUI” base environment from the **Software Selection** option of the OS installer
- RAM: 12 GB, CPU: 9572 MHz total
- Free disk space: 100 GB of disk space in the `/opt/netapp/data` directory, 50 GB in the root partition. For separately mounted `/opt` and `/var/log` directories, ensure that `/opt` has 15 GB, `/var/log` has 16 GB, and `/tmp` has 10 GB of free space.

For detailed system requirements and for information about installing the product in a secured site, see the Requirements for installing Unified Manager and the Interoperability Matrix.

Install Active IQ Unified Manager

Download the installer

1. Download the `ActiveIQUnifiedManager-<version>.zip` installation package along with code signing certificate (`.pem`) and digital signature (`.sig`) .
2. In the directory where you have downloaded the installation file, run:

```
# unzip ActiveIQUnifiedManager-<version>.zip
```

Verify integrity

Run the following commands to verify the integrity of the installer package:

- Run `openssl x509 -pubkey -noout -in AIQUM-RHEL-CLIENT-INTER-ROOT.pem > <public_key_file_name>` to create a file with the public key from the code signing certificate.
- Run `openssl dgst -sha256 -verify <public_key_file_name> -signature <signature_file_name> ActiveIQUnifiedManager-<version>.zip` to verify the signature on the installer package.

Verify repository configuration

The procedures for configuring a Red Hat Enterprise Linux repository is site specific. You can use the `pre_install_check.sh` script included in the installation package to verify the configuration of your operating system. If your system is connected to the internet, you automatically receive instructions for setting up the Red Hat Enterprise Linux repository.

```
# sudo ./pre_install_check.sh
```

Install Unified Manager

Unified Manager uses the `yum` utility to install the software and any dependent software. Because there are varying images of Red Hat Enterprise Linux, the packages installed depend on the software present in the images. The `yum` utility determines the dependent software packages for installation. If you need more information on the dependent software packages, see the Linux software and installation requirements.

To install Unified Manager, run the following command, either as the root user or using `sudo`, from the directory where the installation file was unzipped:

```
# yum install netapp-um<version>.x86_64.rpm
```

or

```
% sudo yum install netapp-um<version>.x86_64.rpm
```

At the end of the installation, the information to connect to the Unified Manager web UI is displayed. If you are unable to connect to the web UI, refer to the `README` file packaged with the software for more information about port 443 restrictions.

Get started with Active IQ Unified Manager on Windows

You can download the installation package and install Unified Manager to monitor and troubleshoot data storage capacity, availability, performance, and protection issues.

System requirements

- Operating Systems
  - Microsoft Windows Server 2019 Standard and Datacenter Edition
  - Microsoft Windows Server 2022 Standard and Datacenter Edition

    Unified Manager is supported on 64 bits Windows operating system for the following languages:

  - English
  - Japanese
  - Simplified Chinese
- RAM: 12 GB
- CPU: 9572 MHz total
- Free disk space: 100 GB of disk space for the installation directory, 50 GB of disk space for the MySQL data directory

For detailed system requirements, see the Requirements for installing Unified Manager and Interoperability Matrix.

Install Active IQ Unified Manager

Download the installer

1. Download the `ActiveIQUnifiedManager-<version>.exe` installation package.
2. Copy the installation file to a directory on the target system.

Install Unified Manager

For installing Unified Manager, ensure that you have Microsoft .NET 4.5.2, or a later version installed. As a part of the installation process, Unified Manager installs other third-party packages as required. For more information about the dependent software packages, refer to the Windows software and installation requirements.

1. Log in to Windows using the default local administrator account.
2. In the directory where you downloaded the installation file, right-click and run the Unified Manager executable (.exe) file as an administrator.
3. When prompted, enter the user name and password to create the Unified Manager maintenance user.
4. In the Database Connection wizard, enter the MySQL root password.
5. Follow the remaining prompts to complete the installation.
6. Click **Finish** at the end of the installation and the Unified Manager web UI is displayed.

Install Unified Manager on VMware vSphere systems

Learn about Active IQ Unified Manager

Learn about Active IQ Unified Manager

Active IQ Unified Manager (formerly OnCommand Unified Manager) enables you to monitor and manage the health and performance of your ONTAP storage systems from a single interface. You can deploy Unified Manager on a Linux server, on a Windows server, or as a virtual appliance on a VMware host.

After you have completed the installation and have added the clusters that you want to manage, Unified Manager provides a graphical interface that displays the capacity, availability, protection, and performance status of the monitored storage systems.

Related information

NetApp Interoperability Matrix Tool

Learn what the Active IQ Unified Manager server does

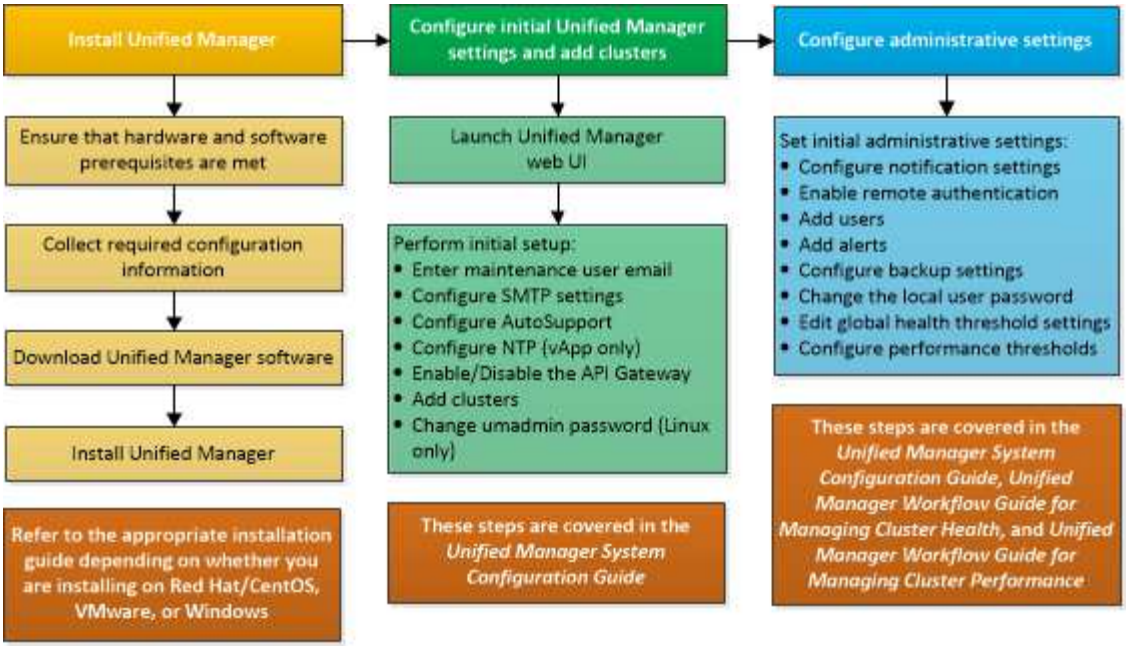
The Unified Manager server infrastructure consists of a data collection unit, a database, and an application server. It provides infrastructure services such as discovery, monitoring, role-based access control (RBAC), auditing, and logging.

Unified Manager collects cluster information, stores the data in the database, and analyzes the data to see if there are any cluster issues.

Learn about the Active IQ Unified Manager installation sequence

The installation workflow describes the tasks that you must perform before you can use Unified Manager.

These sections describe each of the items shown in the workflow below.



Requirements for installing Unified Manager

Ensure the server meets Active IQ Unified Manager installation requirements

Before you begin the installation process, ensure that the server on which you want to install Unified Manager meets the specific software, hardware, CPU, and memory requirements.

NetApp does not support any modification of the Unified Manager application code. If you need to apply any security measures to the Unified Manager server, you should make those changes to the operating system on which Unified Manager is installed.

For more details about applying security measures to the Unified Manager server, see the Knowledge Base article.

Supportability for Security Measures applied to Active IQ Unified Manager for Clustered Data ONTAP

Related information

For more information, see NetApp Interoperability Matrix Tool

Learn about virtual infrastructure and hardware requirements for Active IQ Unified Manager

Installing Unified Manager on virtual infrastructure or a physical system should meet the minimum requirements for memory, CPU, and disk space.

The following table displays the values that are recommended for memory, CPU, and disk space resources. These values have been qualified so that Unified Manager meets acceptable performance levels.

Hardware configuration	Recommended settings
RAM	12 GB

Hardware configuration	Recommended settings
Processors	4 CPUs
CPU cycle capacity	9572 MHz total (minimum requirement 9572 MHz)
Free disk space	• 5 GB (thin provisioned) • 152 GB (thick provisioned)

Unified Manager can be installed on systems with a small amount of memory, but the recommended 12 GB of RAM ensures that enough memory is available for optimal performance, and so that the system can accommodate additional clusters and storage objects as your configuration grows. You should not set any memory limits on the VM where Unified Manager is deployed, and should not enable any features (for example, ballooning) that hinder the software from utilizing the allocated memory on the system.

Additionally, there is a limit to the number of nodes that a single instance of Unified Manager can monitor before you install a second instance of Unified Manager. For node limits, CPU, and memory recommendations, see the Unified Manager Best Practices Guide.

Memory-page swapping negatively impacts the performance of the system and the management application. Competing for CPU resources that are unavailable because of overall host utilization can degrade performance.

Requirement for dedicated use

The physical or virtual system on which you install Unified Manager should be used exclusively for Unified Manager and should not be shared with other applications. Other applications might consume system resources and can drastically reduce the performance of Unified Manager.

Space requirements for backups

If you plan to use the Unified Manager backup and restore feature, allocate additional capacity so that the “data” directory or disk has 150 GB of space. A backup can be written to a local destination or to a remote destination. The best practice is to identify a remote location that is external to the Unified Manager host system that has a minimum of 150 GB of space.

Requirements for host connectivity

The physical system or virtual system on which you install Unified Manager should be configured in such a way that you can successfully `ping` the host name from the host itself. In case of IPv6 configuration, you should verify that `ping6` to the host name is successful to ensure that the Unified Manager installation succeeds.

You can use the host name (or the host IP address) to access the product web UI. If you configured a static IP address for your network during deployment, then you designated a name for the network host. If you configured the network using DHCP, you should obtain the host name from the DNS.

If you plan to allow users to access Unified Manager by using the short name instead of using the fully qualified domain name (FQDN) or IP address, then your network configuration has to resolve this short name to a valid FQDN.

Learn about VMware software and installation requirements for Active IQ Unified Manager

The VMware vSphere system on which you install Unified Manager requires a specific

version of the operating system and specific versions of the supporting software.

Operating system software

The supported operating system is VMware ESXi 8.0.

> Unified Manager OVA on VMware vSphere systems runs Debian OS 13 (Trixie) internally. For information about the versions of the virtual machine hardware that the supported versions of ESXi servers can support, refer to the VMware documentation.

The following versions of vSphere are supported:

- VMware vCenter Server 7.0 and 8.0

See the Interoperability Matrix for the complete and most current list of supported ESXi versions.

mysupport.netapp.com/matrix

The VMware ESXi server time should be the same as the NTP server time for the virtual appliance to function correctly. Synchronizing the VMware ESXi server time with the NTP server time prevents a time failure.

Installation requirements

VMware High Availability for the Unified Manager virtual appliance is supported.

If you deploy an NFS datastore on a storage system that is running ONTAP software, use the NetApp NFS Plug-in for VMware VAAI to use thick provisioning.

If deployment fails using your High Availability-enabled environment because of insufficient resources, you might need to modify the Cluster Features Virtual Machine Options by disabling the VM Restart Priority, and leaving the Host Isolation Response powered on.

> While installing or upgrading Unified Manager, the required third-party software and security patches are automatically installed or upgraded on a VMware vSphere system. Because the Unified Manager installation and upgrade processes control these components, you should not attempt a standalone install or upgrade of any third-party component.

Use supported browsers in Active IQ Unified Manager

To access the Unified Manager web UI, use a supported browser.

The Interoperability Matrix has the list of supported browser versions.

mysupport.netapp.com/matrix

For all browsers, disabling pop-up blockers ensures that software features are displayed properly.

If you plan to configure Unified Manager for SAML authentication, so that an identity provider (IdP) can authenticate users, you should check the list of browsers supported by the IdP as well.

Learn about protocol and port requirements for Active IQ Unified Manager

The required ports and protocols enable communication between the Unified Manager

server and the managed storage systems, servers, and other components.

Connections to the Unified Manager server

In typical installations you do not have to specify port numbers when connecting to the Unified Manager web UI, because default ports are always used. For example, because Unified Manager always attempts to run on its default port, you can enter `https://<host>` instead of `https://<host>:443`.

The Unified Manager server uses specific protocols to access the following interfaces:

Interface	Protocol	Port	Description
Unified Manager web UI	HTTP	80	Used to access the Unified Manager web UI; automatically redirects to the secure port 443.
Unified Manager web UI and programs using APIs	HTTPS	443	Used to securely access the Unified Manager web UI or to make API calls; API calls can only be made using HTTPS.
Maintenance console	SSH/SFTP	22	Used to access the maintenance console and retrieve support bundles.
Linux command line	SSH/SFTP	22	Used to access the Red Hat Enterprise Linux command line and retrieve support bundles.
Syslog	UDP	514	Used to access subscription-based EMS messages from ONTAP systems and to create events based on the messages.
REST	HTTPS	9443	Used to access realtime REST API-based EMS events from authenticated ONTAP systems.
MySQL database	MySQL	3306	Used to enable OnCommand Workflow Automation and OnCommand API Services access to Unified Manager.
AMQP QPID broker	TCP/IP	56072	Used for internal message communication.
AMQP QPID broker	WebSocket over TCP	56080	Used to listen on this port for messages received from ONTAP (cloud agent).

Interface	Protocol	Port	Description
AMQP QPID broker	WebSocket over TCP	56443	Used to listen on this port for messages received from ONTAP (cloud agent). The communication over this port supports encryption provided by TLS or SSL.

The default port for MySQL, 3306, is restricted only to localhost while installing Unified Manager on VMware vSphere systems. This does not impact any upgrade scenario, where the previous configuration is maintained. This configuration can be modified, and the connection can be made available to other hosts by using the `Control access to MySQL port 3306` option on the maintenance console. For information, see Additional menu options. The ports used for HTTP and HTTPS communication (ports 80 and 443) can be changed using the Unified Manager maintenance console. For more information, see Maintenance console menus.

Connections from the Unified Manager server

You should configure your firewall to open ports that enable communication between the Unified Manager server and managed storage systems, servers, and other components. If a port is not open, communication fails.

Depending on your environment, you can choose to modify the ports and protocols used by the Unified Manager server to connect to specific destinations.

The Unified Manager server connects using the following protocols and ports to the managed storage systems, servers, and other components:

Destination	Protocol	Port	Description
Storage system	HTTPS	443/TCP	Used to monitor and manage storage systems. If you are using this port, or any other port to connect to the VMware vCenter Server or ESXi server, ensure that the port is available and can be connected in a secured site.
Storage system	NDMP	10000/TCP	Used for certain Snapshot restore operations.
AutoSupport server	HTTPS	443	Used to send AutoSupport information. Requires the internet access to perform this function.
Authentication server	LDAP	389	Used to make authentication requests, and user and group lookup requests.
	LDAPS	636	Used for secure LDAP communication.
Mail server	SMTP	25	Used to send alert notification emails.
SNMP trap sender	SNMPv1 or SNMPv3	162/UDP	Used to send alert notification SNMP traps.
External data provider server	TCP	2003	Used to send performance data to an external data provider, such as Graphite.
NTP server	NTP	123/UDP	Used to synchronize the time on the Unified Manager server with an external NTP time server. (VMware systems only)

Destination	Protocol	Port	Description
AMQP QPID broker	TCP/IP	56072	Used for internal message communication.
AMQP QPID broker	WebSocket over TCP	56080	Used to listen on this port for messages received from ONTAP (cloud agent).
AMQP QPID broker	WebSocket over TCP	56443	Used to listen on this port for messages received from ONTAP (cloud agent). The communication over this port supports encryption provided by TLS or SSL.
Syslog	UDP	514	Used by Unified Manager to send audit logs to a remote syslog server.

Complete the worksheet in Active IQ Unified Manager

Before you install and configure Unified Manager, you should have specific information about your environment readily available. You can record the information in the worksheet.

Unified Manager installation information

The details required to install Unified Manager.

System on which software is deployed	Your value
ESXi server IP address	
Host fully qualified domain name	
Host IP address	
Network mask	
Gateway IP address	
Primary DNS address	
Secondary DNS address	
Search domains	
Maintenance user name	
Maintenance user password	

Unified Manager configuration information

The details to configure Unified Manager after installation. Some values are optional depending on your configuration.

Setting	Your value
Maintenance user email address	
NTP server	
SMTP server host name or IP address	
SMTP user name	
SMTP password	
SMTP port	25 (Default value)
Email from which alert notifications are sent	
Authentication server host name or IP address	
Active Directory administrator name or LDAP bind distinguished name	
Active Directory password or LDAP bind password	
Authentication server base distinguished name	
Identity provider (IdP) URL	
Identity provider (IdP) metadata	
SNMP trap destination host IP addresses	
SNMP port	

Cluster information

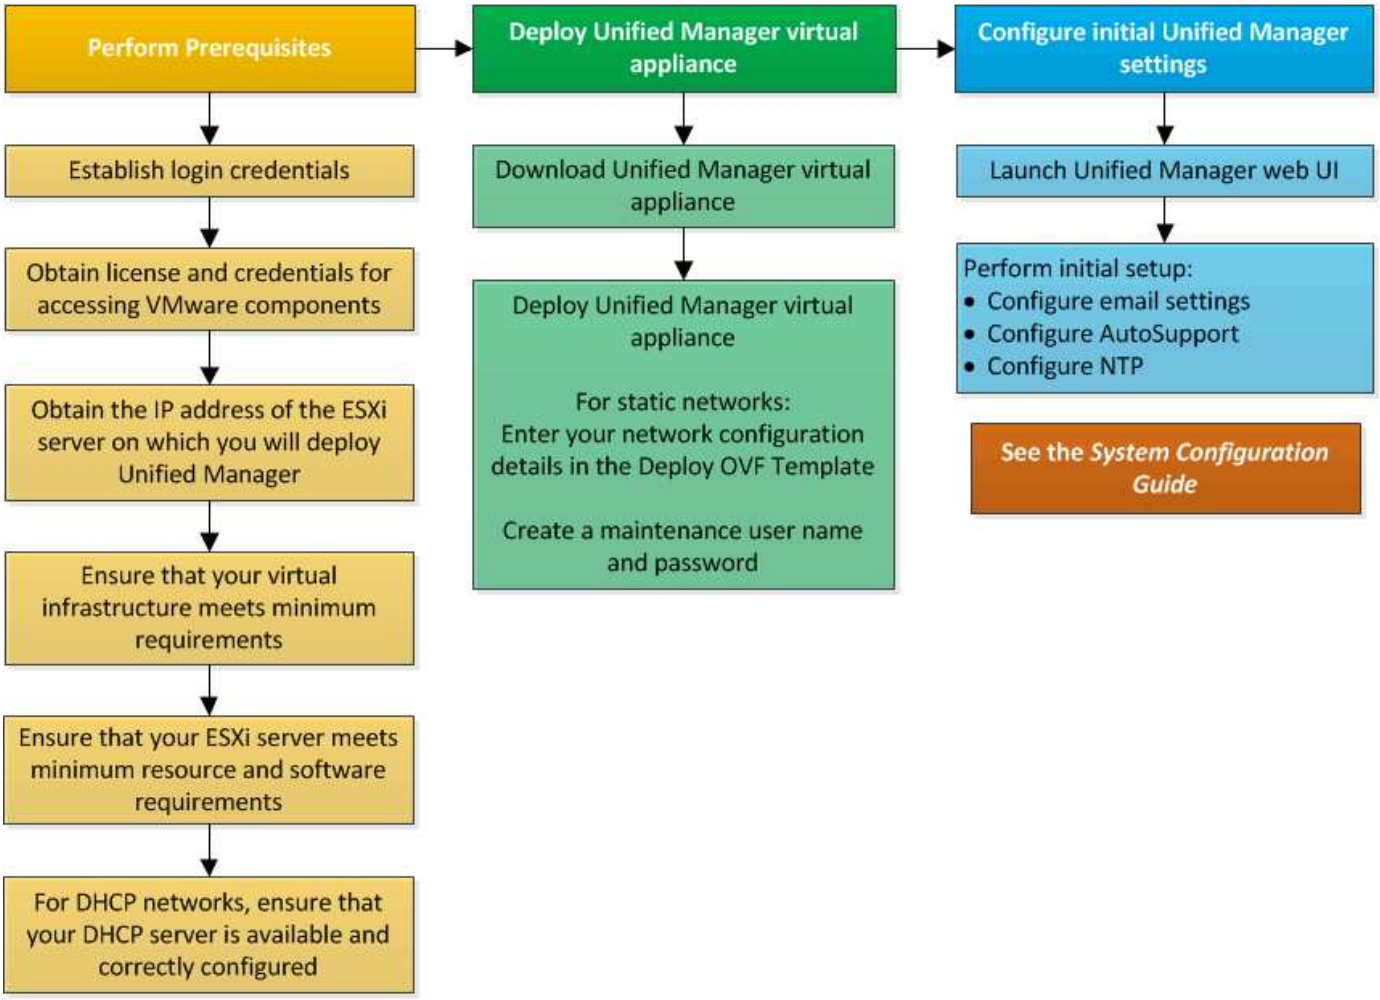
The details for the storage systems that you manage using Unified Manager.

Cluster 1 of N	Your value
Host name or cluster-management IP address	
ONTAP administrator user name The administrator must have been assigned the “admin” role.	
ONTAP administrator password	
Protocol	HTTPS

Installation, upgrade, and removal of Unified Manager on VMware vSphere

Learn about the Active IQ Unified Manager deployment process

The deployment workflow describes the tasks that you must perform before you can use Unified Manager.



Unified Manager virtual appliance deployment workflow

Deploy Active IQ Unified Manager

Deploying Unified Manager includes downloading software, deploying the virtual appliance, creating a maintenance user name and password, and performing the initial setup in the web UI.

Before you begin

- You should verify and complete the system requirements for deployment.

  See System requirements.

- Ensure that you have the following information:
  - Login credentials for the NetApp Support Site
  - Credentials for accessing the VMware vCenter Server and vSphere Web Client
  - IP address of the ESXi server on which you are deploying the Unified Manager virtual appliance
  - Details about the data center, such as storage space in the datastore and memory requirements
  - IPv6 should be enabled on the host if you are planning to use IPv6 addressing.

You can deploy Unified Manager as a virtual appliance on a VMware ESXi server.

You should access the maintenance console by using the VMware console, and not by using SSH.

> Starting with Unified Manager 9.8, VMware Tools has been replaced with Open VM Tools (`open-vm-tools`). You no longer have to install VMware Tools as part of the installation because `open-vm-tools` is included with the Unified Manager installation package.

After finishing the deployment and initial setup, you can either add clusters, or configure additional network settings in the maintenance console, and then access the web UI.

Steps

1. Follow the steps in Download Unified Manager.
2. Additionally, follow the steps in Deploy the Unified Manager virtual appliance.

Download the Active IQ Unified Manager installation file

Download the Unified Manager installation file from the NetApp Support Site to deploy Unified Manager as a virtual appliance.

Before you begin

You should have the login credentials for the NetApp Support Site.

The installation file is a `.tar` file that contains a root certificate, a `README` file, and an `OVA` file containing Unified Manager software configured for a virtual appliance.

Steps

1. Log in to the NetApp Support Site, and navigate to the Download page for Unified Manager:

   NetApp Support Site

2. Select the required version of Unified Manager and accept the end-users license agreement (EULA).
3. Download and save the `.tar` file for VMware vSphere installation to a local directory or network directory that is accessible to your vSphere Client.
4. Verify the checksum to ensure that the software downloaded correctly.
5. Navigate to the directory where you downloaded the `.tar` file and enter the following command in your terminal window to expand the Unified Manager bundle:

   ```
   tar -xvzf ActiveIQUnifiedManager-<version>.tar.gz
   ```

   The required `OVA` file, a root certificate, and a `README` file for Unified Manager are unzipped to the target directory.

6. Verify the integrity of the `OVA` file by following the steps provided in the `README` file.

Deploy the Active IQ Unified Manager virtual appliance

After downloading the installation file, you deploy Unified Manager as a virtual appliance. Use the vSphere Web Client to deploy the virtual appliance on an ESXi server. When you deploy the virtual appliance, a virtual machine is created.

Before you begin

You should review the system requirements. Make the necessary changes before deploying the Unified Manager virtual appliance.

See Virtual infrastructure requirements.

See VMware software and installation requirements.

If you use Dynamic Host Configuration Protocol (DHCP), ensure that the DHCP server is available, and that the DHCP and virtual machine (VM) network adapter configurations are correct. DHCP is configured by default.

If you use a static networking configuration, ensure that the IP address is not duplicated in the same subnet, and that the appropriate DNS server entries have been configured.

Obtain the following information before deploying the virtual appliance:

- Credentials for accessing the VMware vCenter Server and vSphere Web Client
- IP address of the ESXi server on which you are deploying the Unified Manager virtual appliance
- Details about the data center, such as availability of storage space
- If you are not using DHCP, obtain the IPv4 or IPv6 addresses for the networking devices to which you are planning to connect:
  - Fully qualified domain name (FQDN) of the host
  - IP address of the host
  - Network mask
  - IP address of the default gateway
  - Primary and secondary DNS addresses
  - Search domains

Starting with Unified Manager 9.8, VMware Tools has been replaced with Open VM Tools (`open-vm-tools`). You do not have to install VMware Tools as part of the installation process because `open-vm-tools` is included with the Unified Manager installation package.

When the virtual appliance is deployed, a unique self-signed certificate for HTTPS access is generated. When accessing the Unified Manager web UI, you might see a browser warning about untrusted certificates.

VMware High Availability for the Unified Manager virtual appliance is supported.

Steps

1. In the vSphere Client, click **File** > **Deploy OVF Template**.
2. Complete the Deploy OVF Template wizard to deploy the Unified Manager virtual appliance.

   On the Review Details page:

   - Verify the details of the Publisher section. The message **Entrust Code Signing - OVCS2 (Trusted certificate)** confirms the integrity of the downloaded `OVA` file.

     If you see the message **Entrust Code Signing - OVCS2 (Invalid certificate)**, then upgrade the VMware vCenter Server to 7.0U3E or greater version.

On the Customize Template page:

- Leave all the fields blank when using DHCP and IPv4 addressing.
- Check the “Enable Auto IPv6 addressing” box, and leave all the other fields blank when using DHCP and IPv6 addressing.
- If you want to use a static network configuration, you can complete the fields on this page and these settings are applied during deployment. Ensure that the IP address is unique to the host on which it is deployed, that it is not already in use, and that it has a valid DNS entry.

3. After the Unified Manager virtual appliance is deployed to the ESXi server, power on the VM by right-clicking the VM, and then selecting **Power On**.

   If the Power On operation fails because of insufficient resources, add resources and then retry the installation.

4. Click the **Console** tab.

   The initial boot process takes a few minutes to complete.

5. To configure your time zone, enter your geographic area and your city or region as prompted in the VM Console window.

   All the date information that is displayed uses the time zone that is configured for Unified Manager, regardless of the time zone setting on your managed devices. If your storage systems and the management server are configured with the same NTP server, they refer to the same instant in time, even if they appear differently. For example, if you create a Snapshot copy using a device that is configured using a different time zone than that of the management server, the time stamp is the management server time.

6. If no DHCP services are available, or if there is an error in the details for the static network configuration, select one of the following options:

If you use…	Then do this…
DHCP	Select **Retry DHCP**. If you plan to use DHCP, you should ensure that it is configured correctly. If you use a DHCP-enabled network, the FQDN and DNS server entries are given to the virtual appliance automatically. If DHCP is not properly configured with DNS, the host name “UnifiedManager” is automatically assigned and associated with the security certificate. If you have not set up a DHCP-enabled network, you should manually enter the networking configuration information.
A static network configuration	a. Select **Enter the details for static network configuration**. The configuration process takes a few minutes to complete. b. Confirm the values that you entered, and select **Y**.

7. At the prompt, enter a maintenance user name, and then click **Enter**.

   The maintenance user name should start with a letter from a-z, followed by any combination of -, a-z, or 0-9.

8. At the prompt, enter a password, and then click **Enter**.

   The VM console displays the URL for the Unified Manager web UI.

You can access the web UI to perform the initial setup of Unified Manager, as described in the Configuring Active IQ Unified Manager.

Unified Manager virtual appliance upgrade workflow

Upgrade Active IQ Unified Manager

You can upgrade Active IQ Unified Manager to release 9.18 from release 9.14 or 9.16 only.

During the upgrade process, Unified Manager is unavailable. You should complete any running operations before upgrading Unified Manager.

If Unified Manager is paired with an instance of OnCommand Workflow Automation, and there are new versions of software available for both products, you must disconnect the two products and then set up a new Workflow Automation connection after performing the upgrades. If you are performing an upgrade to only one of the products, then you should log into Workflow Automation after the upgrade and verify that it is still acquiring data from Unified Manager.

> If your system was upgraded from Unified Manager 9.7, you cannot upgrade directly from 9.14 to 9.18 because of file system changes in the Debian versions. Instead, follow this recommended upgrade path:
> 9.14 → 9.16P2 → 9.18.

Steps

1. Follow the steps in Download the Unified Manager ISO image.
2. Additionally, follow the steps described in Upgrade Unified Manager.

Use the supported upgrade path for versions of Active IQ Unified Manager

Active IQ Unified Manager supports a specific upgrade path for each version.

Not all versions of Unified Manager can perform an in-place upgrade to later versions. The Unified Manager upgrades are limited to an N-2 model. This means that you can only perform an upgrade to the next two releases on all platforms. For example, you can perform an upgrade to Unified Manager 9.18 from Unified Manager 9.14 and 9.16 only.

If you are running a version that is earlier than the supported versions, your Unified Manager instance will need to be upgraded to one of the supported versions first, and then upgraded to the current version.

For example, if your installed version is Unified Manager 9.9 and you want to upgrade to Unified Manager 9.14, you follow a sequence of upgrades.

Sample upgrade path:

1. Upgrade 9.11 → 9.13
2. Upgrade 9.13 → 9.14
3. Upgrade 9.13 → 9.16
4. Upgrade 9.14 or 9.16 → 9.18

For more information on the upgrade path matrix, see this knowledge base (KB) article.

Download the Active IQ Unified Manager upgrade file

Before upgrading Unified Manager, download the Unified Manager upgrade file from the NetApp Support Site.

Before you begin

You should have login credentials for the NetApp Support Site.

Steps

1. Log in to the NetApp Support Site:

   NetApp Support Site

2. Navigate to the Download page for upgrading Unified Manager on VMware vSphere.
3. Download the `.iso` image for upgrade, and save it to a local directory or network directory that is accessible to your vSphere Client.
4. Verify the checksum to ensure that the software has downloaded correctly.

Upgrade the Active IQ Unified Manager virtual appliance

You can upgrade the Active IQ Unified Manager virtual appliance from releases 9.14 or 9.16 to 9.18.

Before you begin

Ensure the following:

- You have downloaded the upgrade file, the ISO image from the NetApp Support Site.
- The system on which you are upgrading Unified Manager meet the system and software requirements.

  See Virtual infrastructure requirements.

  See VMware software and installation requirements.

- For vSphere 6.5 and later users, you have installed the VMware Remote Console (VMRC).
- During upgrade, you may be prompted to confirm on whether you want to keep the previous, default settings for retaining performance data for 13 months or to change it to 6 months. On confirming, the historical performance data is purged after 6 months.
- You have the following information:
  - Login credentials for the NetApp Support Site
  - Credentials for accessing the VMware vCenter Server and vSphere Web Client
  - Credentials for the Unified Manager maintenance user

During the upgrade process, Unified Manager is unavailable. You should complete any running operations before upgrading Unified Manager.

If you have paired Workflow Automation and Unified Manager, you should manually update the host name in Workflow Automation.

Steps

1. In the vSphere Client, click **Home** > **Inventory** > **VMs and Templates**.
2. Select the virtual machine (VM) on which the Unified Manager virtual appliance is installed.
3. If the Unified Manager VM is running, navigate to **Summary** > **Commands** > **Shut Down Guest**.
4. Create a backup copy, such as a snapshot or a clone, of the Unified Manager VM to create an application-consistent backup.
5. From the vSphere Client, power on the Unified Manager VM.
6. Launch the VMware Remote Console.
7. Click the **CDROM** icon, and select **Connect to Disk Image File (.iso)**.
8. Select the `ActiveIQUnifiedManager-<version>-virtual-update.iso` file, and click **Open**.
9. Click the **Console** tab.
10. Log in to the Unified Manager maintenance console.
11. In the Main Menu, select **Upgrade**.

    A message is displayed that Unified Manager is to be unavailable during the upgrade process, and is to resume after completion.

12. Type `y` to continue.

    A warning is displayed, reminding you to back up the virtual machine on which the virtual appliance resides.

13. Type `y` to continue.

    The upgrade process and the restart of Unified Manager services can take several minutes to complete.

14. Press any key to continue.

    You are automatically logged out of the maintenance console.

15. **Optional:** Log in to the maintenance console, and verify the version of Unified Manager.

You can launch the web UI on a new window on a supported web browser and log in to use the upgraded version of Unified Manager. Note that you need to wait for the discovery process to finish before performing any task in the UI.

Restart the Active IQ Unified Manager virtual machine

You can restart the Unified Manager virtual machine (VM) from the maintenance console. You should restart the VM after generating a new security certificate, or if there is a problem with the VM.

Before you begin

- The virtual appliance should be powered on.
- You should be logged in to the Unified Manager maintenance console as the maintenance user.

You can also restart the virtual machine from vSphere by using the VMware **Restart Guest** option. Refer to the Restart-VMGuest cmdlet in Broadcom's VMware vSphere PowerCLI CMDLET reference.

Steps

1. In the maintenance console, select **System Configuration** > **Reboot Virtual Machine**.
2. Start the Unified Manager web UI from your browser, and log in.

Remove Active IQ Unified Manager

You can uninstall Unified Manager by removing the virtual machine (VM) on which the Unified Manager software is installed.

Before you begin

- You should have credentials for accessing VMware vCenter Server and vSphere Web Client.
- Any active connection of the Unified Manager server to a Workflow Automation server should be closed.
- All clusters (data sources) should be removed from the Unified Manager server before you remove the virtual machine (VM).

Steps

1. Use the Unified Manager maintenance console to verify that the Unified Manager server does not have an active connection to an external data provider.
2. In the vSphere Client, click **Home** > **Inventory** > **VMs and Templates**.
3. Select the VM that you want to remove, and click the **Summary** tab.
4. If the VM is running, click **Power** > **Shut Down Guest**.
5. Right-click the VM that you want to remove, and click **Delete from Disk**.

Install Unified Manager on Linux systems

Introduction to Active IQ Unified Manager

Learn about Active IQ Unified Manager

Active IQ Unified Manager (formerly OnCommand Unified Manager) enables you to monitor and manage the health and performance of your ONTAP storage systems from a single interface. You can deploy Unified Manager on a Linux server, on a Windows server, or as a virtual appliance (vApp) on a VMware host.

After you have completed the installation and have added the clusters that you want to manage, Unified Manager provides a graphical interface that displays the capacity, availability, protection, and performance status of the monitored storage systems.

Related information

NetApp Interoperability Matrix Tool

Understand what the Active IQ Unified Manager server does

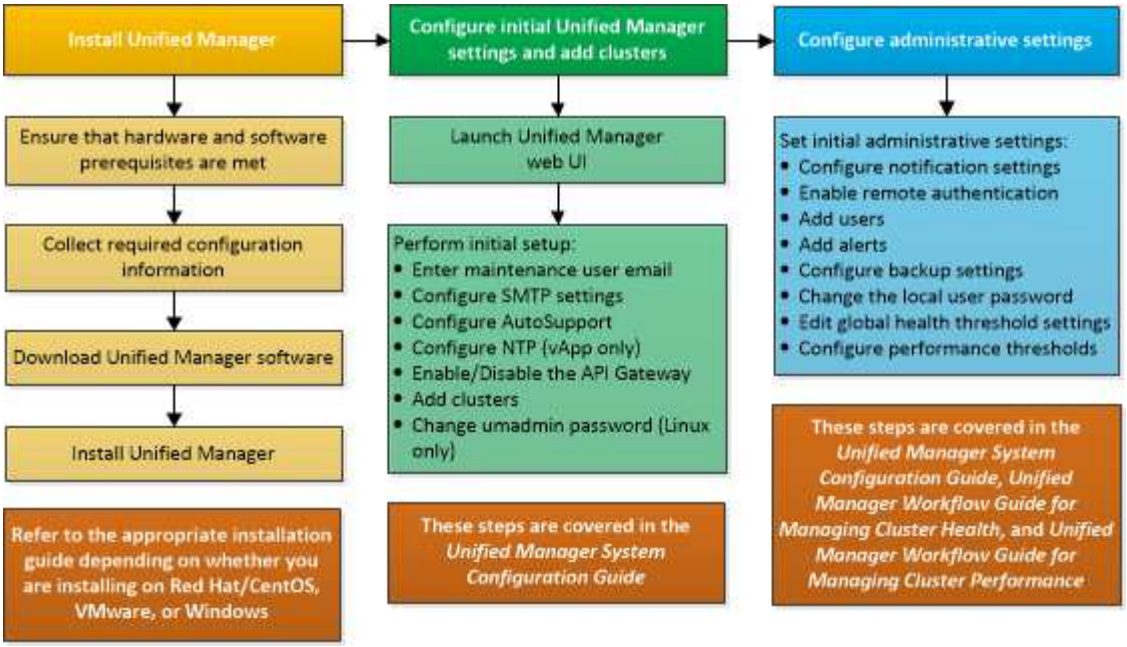
The Unified Manager server infrastructure consists of a data collection unit, a database, and an application server. It provides infrastructure services such as discovery, monitoring, role-based access control (RBAC), auditing, and logging.

Unified Manager collects cluster information, stores the data in the database, and analyzes the data to see if there are any cluster issues.

View the installation sequence for Active IQ Unified Manager

The installation workflow describes the tasks that you must perform before you can use Unified Manager.

These sections describe each of the items shown in the workflow below.



Requirements for installing Unified Manager

Understand the requirements for installing Active IQ Unified Manager

Before you begin the installation process, ensure that the server on which you want to install Unified Manager meets the specific software, hardware, CPU, and memory requirements.

NetApp does not support any modification of the Unified Manager application code. If you need to apply any security measures to the Unified Manager server, you should make those changes to the operating system on which Unified Manager is installed.

For more details about applying security measures to the Unified Manager server, see the Knowledge Base article.

Supportability for Security Measures applied to Active IQ Unified Manager for Clustered Data ONTAP

Related information

NetApp Interoperability Matrix Tool

Understand virtual infrastructure and hardware requirements for Active IQ Unified Manager

Installing Unified Manager on virtual infrastructure or a physical system should meet the minimum requirements for memory, CPU, and disk space.

The following table displays the values that are recommended for memory, CPU, and disk space resources. These values have been qualified so that Unified Manager meets acceptable performance levels.

Hardware configuration	Recommended settings
RAM	12 GB

Hardware configuration	Recommended settings
Processors	4 CPUs
CPU cycle capacity	9572 MHz total (minimum requirement 9572 MHz)
Free disk space	150 GB, where the capacity is allocated as follows: • 50 GB allotted to the root partition • 100 GB of free disk space allotted to the `/opt/netapp/data` directory, which is mounted on an LVM drive or on a separate local disk attached to the target system For separately mounted `/opt` and `/var/log` directories, ensure that `/opt` has 15 GB and `/var/log` has 16 GB of free space. The `/tmp` directory should have at least 10 GB of free space.

Unified Manager can be installed on systems with a small amount of memory, but the recommended 12 GB of RAM ensures that enough memory is available for optimal performance, and so that the system can accommodate additional clusters and storage objects as your configuration grows. You should not set any memory limits on the VM where Unified Manager is deployed, and should not enable any features (for example, ballooning) that hinder the software from utilizing the allocated memory on the system.

Additionally, there is a limit to the number of nodes that a single instance of Unified Manager can monitor before you install a second instance of Unified Manager. For node limits, CPU, and memory recommendations, see the Unified Manager Best Practices Guide.

Memory-page swapping negatively impacts the performance of the system and the management application. Competing for CPU resources that are unavailable because of overall host utilization can degrade performance.

Requirement for dedicated use

The physical or virtual system on which you install Unified Manager should be used exclusively for Unified Manager and should not be shared with other applications. Other applications might consume system resources and can drastically reduce the performance of Unified Manager.

Space requirements for backups

If you plan to use the Unified Manager backup and restore feature, allocate additional capacity so that the “data” directory or disk has 150 GB of space. A backup can be written to a local destination or to a remote destination. The best practice is to identify a remote location that is external to the Unified Manager host system that has a minimum of 150 GB of space.

Requirements for host connectivity

The physical system or virtual system on which you install Unified Manager should be configured in such a way that you can successfully `ping` the host name from the host itself. In case of IPv6 configuration, you should verify that `ping6` to the host name is successful to ensure that the Unified Manager installation succeeds.

You can use the host name (or the host IP address) to access the product web UI. If you configured a static IP address for your network during deployment, then you designated a name for the network host. If you configured the network using DHCP, you should obtain the host name from the DNS.

If you plan to allow users to access Unified Manager by using the short name instead of using the fully qualified domain name (FQDN) or IP address, then your network configuration has to resolve this short name to a valid FQDN.

Learn about Linux software and installation requirements in Active IQ Unified Manager

The Linux system on which you install Unified Manager requires specific versions of the operating system and supporting software. For Linux installations, Unified Manager can be installed only on a Red Hat Enterprise Linux operating system. It cannot be installed on any other operating system.

Operating system software

The Linux system must have the following versions of the operating system and supporting software installed:

- Red Hat Enterprise Linux versions 8.x and 9.6, based on x86_64 architecture

See the Interoperability Matrix for the complete and most current list of supported Red Hat Enterprise Linux versions.

NetApp Interoperability Matrix

The server should be dedicated to running Unified Manager. No other applications should be installed on the server. In some environments, a vulnerability scanner such as Qualys is installed on the Linux system because of company policy. You should disable the vulnerability scanner before installing Unified Manager to prevent installation failure.

> Vulnerability scanners (such as Qualys) might lead to high CPU usage when scanning the virtual machine (VM) either internally (where Unified Manager and the vulnerability scanner are installed on the same VM) or externally (where Unified Manager and the vulnerability scanner are installed on two different servers and the vulnerability scanner scans the VM where Unified Manager is installed). This issue often results in the VM being unresponsive and affects the Unified Manager services. Therefore, NetApp recommends disabling the vulnerability scanner in the VM where Unified Manager is installed. If disabling the scanner is not an option, then scan the VM during non-business hours and restart the services after the scan is complete.

Third-party software

Unified Manager is deployed on a WildFly web server. WildFly 27.0.1.Final is bundled and configured with Unified Manager.

The following third-party packages are required, but not included with Unified Manager. These packages are automatically installed by the `yum` installer during installation, provided you have configured the repositories as mentioned in the following sections.

- MySQL Community Edition version 8.4.10 (from the MySQL repository)
- OpenJDK version 17.0.14+6 (from the Red Hat Extra Enterprise Linux Server repository)

- Python 3.13.3 (from the Red Hat Extra Enterprise Linux Server repository)
- 7zip version 7z2301-linux-x64. For more information, view Configure the 7-zip software package.

> You must shut down a running instance of Unified Manager before upgrading any third-party software. After the third-party software installation is complete, you can restart Unified Manager.

User authorization requirements

Installation of Unified Manager on a Linux system can be performed by the root user or by non-root users by using the `sudo` command.

Additionally, the default Unified Manager installation includes the umadmin user (default maintenance user) and the jboss user in the sudoers files (`ocum_sudoers` and `ocie_sudoers`) in the `/etc/sudoers.d/` directory. If you remove this content from your environment because of security policies, or because of a security monitoring tool, you must add it back. You need to preserve the sudoers configuration because some Unified Manager operations require these sudo privileges.

Installation requirements

The best practices for installing Red Hat Enterprise Linux and the associated repositories on your system are listed below. Systems installed or configured differently, or deployed off premise (in the cloud), might require additional steps, and Unified Manager might not run properly in such deployments.

- You must install Red Hat Enterprise Linux according to Red Hat best practices, and you should select the following default options, which requires selecting the “Server with GUI” base environment.
- While installing Unified Manager on Red Hat Enterprise Linux, the system must have access to the appropriate repository so that the installation program can access and install all the required software dependencies.
- For the `yum` installer to find dependent software in the Red Hat Enterprise Linux repositories, you must have registered the system during the Red Hat Enterprise Linux installation or afterwards by using a valid Red Hat subscription.

  See the Red Hat documentation for information about the Red Hat Subscription Manager.

- You must enable the Extra Packages for Enterprise Linux (EPEL) repository to successfully install the required third-party utilities on your system.

  If the EPEL repository is not configured on your system, you must manually download and configure the repository.

  See Manually configuring the EPEL repository.

- If the correct version of MySQL is not installed, you must enable the MySQL repository to successfully install MySQL software on your system.

  If the MySQL repository is not configured on your system, you must manually download and configure the repository.

  See Manually configuring the MySQL repository.

- You must install only one version of Java on your Unified Manager host; otherwise, the system may not be able to determine which version to use. See After updating Java on the Active IQ Unified Manager host, services will not start.

If your system does not have internet access, and the repositories are not mirrored from an internet-connected system to the unconnected system, you should follow the installation instructions to determine the external software dependencies of your system. Then you can download the required software to the internet-connected system, and copy the `.rpm` files to the system on which you plan to install Unified Manager. To download the artifacts and packages, you must use the `yum install` command. You must ensure that the two systems are running the same operating system version and that the subscription license is for the appropriate Red Hat Enterprise Linux.

> You must not install the required third-party software from repositories other than the repositories that are listed here. Software installed from the Red Hat repositories is designed explicitly for Red Hat Enterprise Linux, and conforms to Red Hat best practices (directory layouts, permissions, and so on). Software from other locations might not follow these guidelines, which might cause the Unified Manager installation to fail, or might cause issues with future upgrades.

Port 443 requirement

Generic images of Red Hat Enterprise Linux might block external access to port 443. Due to this restriction, you might be unable to connect to the Administrator web UI after installing Unified Manager. Running the following command allows access to port 443 for all external users and applications on a generic Red Hat Enterprise Linux.

```
# firewall-cmd --zone=public --add-port=443/tcp --permanent; firewall-cmd
--reload
```

You must install Red Hat Enterprise Linux with the “Server with GUI” base environment. It provides the commands used by Unified Manager installation instructions. Other base environments might require you to install additional commands to validate or complete the installation. If the `firewall-cmd` is not available on your system, you must install it by running the following command:

```
# sudo yum install firewalld
```

Contact your IT department before running the commands to see if your security policies require a different procedure.

> THP (Transparent Huge Pages) should be disabled on Red Hat systems. When enabled, in some cases it can cause Unified Manager to be shut down when certain processes consume too much memory and are terminated.

Use supported browsers for Active IQ Unified Manager

To access the Unified Manager web UI, use a supported browser.

The Interoperability Matrix has the list of supported browser versions.

mysupport.netapp.com/matrix

For all browsers, disabling pop-up blockers ensures that software features are displayed properly.

If you plan to configure Unified Manager for SAML authentication, so that an identity provider (IdP) can authenticate users, you should check the list of browsers supported by the IdP as well.

Learn about protocol and port requirements in Active IQ Unified Manager

The required ports and protocols enable communication between the Unified Manager server and the managed storage systems, servers, and other components.

Connections to the Unified Manager server

In typical installations you do not have to specify port numbers when connecting to the Unified Manager web UI, because default ports are always used. For example, because Unified Manager always attempts to run on its default port, you can enter `https://<host>` instead of `https://<host>:443`.

The Unified Manager server uses specific protocols to access the following interfaces:

Interface	Protocol	Port	Description
Unified Manager web UI	HTTP	80	Used to access the Unified Manager web UI; automatically redirects to the secure port 443.
Unified Manager web UI and programs using APIs	HTTPS	443	Used to securely access the Unified Manager web UI or to make API calls; API calls can only be made using HTTPS.
Maintenance console	SSH/SFTP	22	Used to access the maintenance console and retrieve support bundles.
Linux command line	SSH/SFTP	22	Used to access the Red Hat Enterprise Linux command line and retrieve support bundles.
MySQL database	MySQL	3306	Used to enable OnCommand Workflow Automation and OnCommand API Services access to Unified Manager.
Syslog	UDP	514	Used to access subscription-based EMS messages from ONTAP systems and to create events based on the messages.
REST	HTTPS	9443	Used to access realtime REST API-based EMS events from authenticated ONTAP systems.

Interface	Protocol	Port	Description
MySQL database	MySQL	3306	Used to enable OnCommand Workflow Automation and OnCommand API Services access to Unified Manager.
AMQP QPID broker	TCP/IP	56072	Used for internal message communication.
AMQP QPID broker	WebSocket over TCP	56080	Used to listen on this port for messages received from ONTAP (cloud agent).
AMQP QPID broker	WebSocket over TCP	56443	Used to listen on this port for messages received from ONTAP (cloud agent). The communication over this port supports encryption provided by TLS or SSL.

The default port for MySQL, 3306, is restricted only to localhost while installing Unified Manager on Linux systems. This does not impact any upgrade scenario, where the previous configuration is maintained. This configuration can be modified, and the connection can be made available to other hosts by using the `Control access to MySQL port 3306` option on the maintenance console. For information, see Additional menu options. The ports used for HTTP and HTTPS communication (ports 80 and 443) can be changed using the Unified Manager maintenance console. For more information, see Maintenance console menus.

Connections from the Unified Manager server

You should configure your firewall to open ports that enable communication between the Unified Manager server and managed storage systems, servers, and other components. If a port is not open, communication fails.

Depending on your environment, you can choose to modify the ports and protocols used by the Unified Manager server to connect to specific destinations.

The Unified Manager server connects using the following protocols and ports to the managed storage systems, servers, and other components:

Destination	Protocol	Port	Description
Storage system	HTTPS	443/TCP	Used to monitor and manage storage systems.
Storage system	NDMP	10000/TCP	Used for certain Snapshot restore operations.

Destination	Protocol	Port	Description
AutoSupport server	HTTPS	443	Used to send AutoSupport information. Requires the internet access to perform this function.
Authentication server	LDAP	389	Used to make authentication requests, and user and group lookup requests.
	LDAPS	636	Used for secure LDAP communication.
Mail server	SMTP	25	Used to send alert notification emails.
SNMP trap sender	SNMPv1 or SNMPv3	162/UDP	Used to send alert notification SNMP traps.
External data provider server	TCP	2003	Used to send performance data to an external data provider, such as Graphite.
NTP server	NTP	123/UDP	Used to synchronize the time on the Unified Manager server with an external NTP time server. (VMware systems only)
Syslog	UDP	514	Used by Unified Manager to send audit logs to a remote syslog server.
AMQP QPID broker	TCP/IP	56072	Used for internal message communication.
AMQP QPID broker	WebSocket over TCP	56080	Used to listen on this port for messages received from ONTAP (cloud agent).
AMQP QPID broker	WebSocket over TCP	56443	Used to listen on this port for messages received from ONTAP (cloud agent). The communication over this port supports encryption provided by TLS or SSL.

Complete the worksheet in Active IQ Unified Manager

Before you install and configure Unified Manager, you should have specific information about your environment readily available. You can record the information in the worksheet.

Unified Manager installation information

The details required to install Unified Manager.

System on which software is deployed	Your value
Host fully qualified domain name	
Host IP address	
Network mask	
Gateway IP address	
Primary DNS address	
Secondary DNS address	
Search domains	
Maintenance user name	
Maintenance user password	

Unified Manager configuration information

The details to configure Unified Manager after installation. Some values are optional depending on your configuration.

Setting	Your value
Maintenance user email address	
SMTP server host name or IP address	
SMTP user name	
SMTP password	
SMTP port	25 (Default value)
Email from which alert notifications are sent	
Authentication server host name or IP address	
Active Directory administrator name or LDAP bind distinguished name	
Active Directory password or LDAP bind password	
Authentication server base distinguished name	
Identity provider (IdP) URL	
Identity provider (IdP) metadata	
SNMP trap destination host IP addresses	
SNMP port	

Cluster information

The details for the storage systems that you manage using Unified Manager.

Cluster 1 of N	Your value
Host name or cluster-management IP address	
ONTAP administrator user name The administrator must have been assigned the “admin” role.	
ONTAP administrator password	
Protocol	HTTPS

Install, upgrade, and remove Unified Manager software

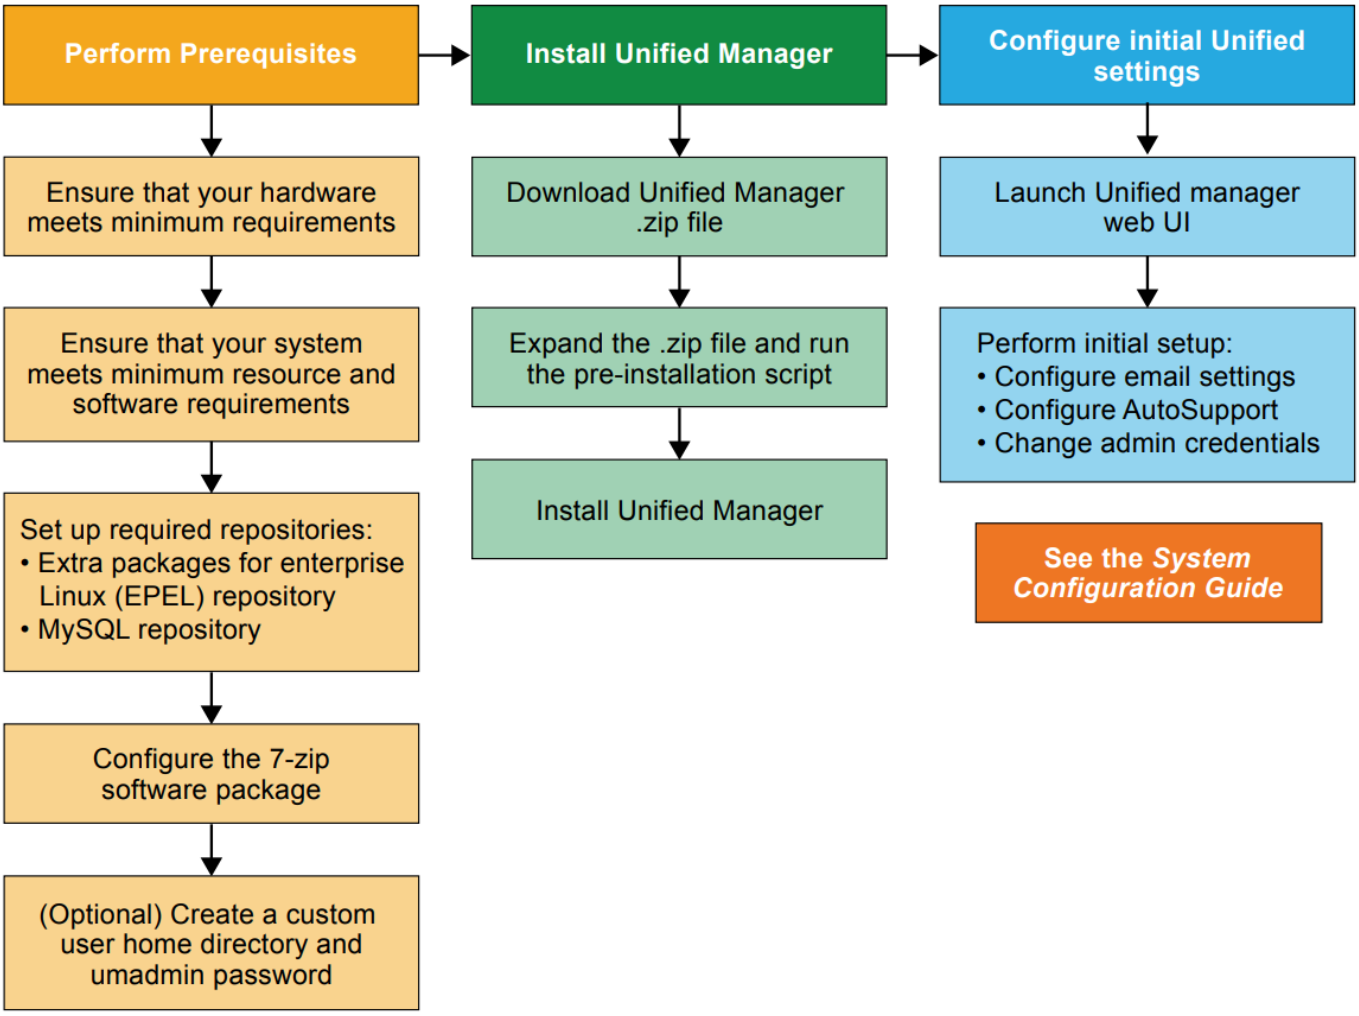
Install, upgrade, and remove Active IQ Unified Manager software

On Linux systems, you can install Unified Manager software, upgrade to a newer version of software, or remove Unified Manager.

Unified Manager can be installed on Red Hat Enterprise Linux servers. The Linux server on which you install Unified Manager can be running either on a physical machine or on a virtual machine running on VMware ESXi, Microsoft Hyper-V, or Citrix XenServer.

Understand the installation sequence for Active IQ Unified Manager on Red Hat Enterprise Linux

The installation workflow describes the tasks that you must perform before you can use Unified Manager.



Set up required software

Configure the EPEL repository for Active IQ Unified Manager

If the system on which you are installing Unified Manager does not have access to the Extra Packages for Enterprise Linux (EPEL) repository, then you must manually download and configure the repository for a successful installation.

The EPEL repository provides access to the required third-party utilities that must be installed on your system. You use the EPEL repository to install Unified Manager on a Red Hat Enterprise Linux system.

Steps

1. Download the EPEL repository for your installation. For version 8, download it from:

   ```
   wget https://dl.fedoraproject.org/pub/epel/epel-release-latest-8.noarch.rpm
   ```

   For version 9, download it from:

   ```
   wget https://dl.fedoraproject.org/pub/epel/epel-release-latest-9.noarch.rpm
   ```

2. Configure the EPEL repository:

```
yum install epel-release-latest-<version>.noarch.rpm
```

For Red Hat Enterprise Linux 8 systems, if you have internal repositories with modular RPM packages, for example, *javapackages-filesystem-<version>.module.rpm*, ensure that the metadata for the modular packages is also available in the same repository.

Configure the MySQL repository for Active IQ Unified Manager

If the system on which you are installing Unified Manager does not have access to the MySQL Community Edition repository, then you must manually download and configure the repository for a successful installation.

The MySQL repository provides access to the required MySQL software that must be installed on your system.

> This task might fail if the system does not have the internet connectivity. Refer to the MySQL documentation if the system on which you are installing Unified Manager does not have the internet access.

Steps

1. Download the appropriate MySQL repository for your installation. For version 8, download it from:

   ```
   wget http://repo.mysql.com/yum/mysql-8.4-community/el/8/x86_64/mysql84-community-release-el8-1.noarch.rpm
   ```

   For version 9, download it from:

   ```
   wget http://repo.mysql.com/yum/mysql-8.4-community/el/9/x86_64/mysql84-community-release-el9-1.noarch.rpm
   ```

2. Configure the MySQL repository:

   ```
   yum install mysql84-community-release-.noarch.rpm
   ```

For a Red Hat Enterprise Linux 8 system, if you have internal repositories with java-17-openjdk, 7-zip, and other software packages provided by the AppStream repository, you must disable your AppStream repository and install MySQL Community Server. Run the following command:

```
# sudo yum --disablerepo=rhel-8-for-x86_64-appstream-rpms install  mysql-community-server
```

If you receive an error on missing key or key mismatch, and your installation fails, try these steps:

- On a connected system, import the updated MySQL key by running the following command:

```
rpm --import https://repo.mysql.com/RPM-GPG-KEY-mysql-<xxxx>

for example:
```

```
rpm --import https://repo.mysql.com/RPM-GPG-KEY-mysql-2023
```

- On a system that does not have internet connectivity, update your MySQL repo file and disable `gpgcheck` by marking `gpgcheck=0`.

Configure the 7-zip software package in Active IQ Unified Manager

To successfully install Unified Manager on a Red Hat Enterprise Linux system, your system must be set up with access to the 7-zip software package. If network access is available and the 7-zip URL is unblocked by a firewall, 7-zip downloads automatically to your system during an installation or upgrade process. NetApp recommends the automatic download method.

Ensure the following URL is unblocked in your firewall to allow Unified Manager to automatically download the 7-zip software package:
`https://7-zip.org/a/7z2301-linux-x64.tar.xz`

Alternatively, you can pre-install 7-zip.

Pre-install the 7-zip package

You can pre-install 7-zip by downloading it manually or by using a bundled script. Choose from one of the following two methods:

Download using a bundled script

Step

1. Run the `install7zip.sh` script that is included in the Unified Manager zip file.

Download manually

Ensure you download the 7-zip package only from the official website.

Steps

1. Download the 7-zip package (version 23.01) from the official site:

   `https://7-zip.org/a/7z2301-linux-x64.tar.xz`

2. Copy the file to your RHEL virtual machine:

   ```
   7z2301-linux-x64.tar.xz
   ```

3. Extract the archive:

   ```
   sudo tar -xf 7z2301-linux-x64.tar.xz
   ```

4. Move the `7zz` file to the `/usr/bin/` directory:

```
sudo cp 7zz /usr/bin/
```

5. Set executable permissions:

```
sudo chmod 755 /usr/bin/7zz
```

Verify the installation

Verify that the 7-zip package has been successfully installed on your system:

```
7zz -version
```

Configure SELinux for NFS and CIFS shares in Active IQ Unified Manager

If you are planning to mount `/opt/netapp` or `/opt/netapp/data` on an NAS or SAN device, and you have SELinux enabled, you need to be aware of a few considerations.

If you are planning to mount `/opt/netapp` or `/opt/netapp/data` from anywhere other than the root file system, and you have SELinux enabled in your environment, you should set the correct context for the mounted directories.
For the applicable scenario in your environment, follow these steps for setting and confirming the correct SELinux context.

Configure the SELinux context when `/opt/netapp/data` is mounted

If you have mounted `/opt/netapp/data` in your system and SELinux is set to `Enforcing`, ensure that the SELinux context type for `/opt/netapp/data` is set to `mysqld_db_t`, which is the default context element for the location of the database files.

1. Run this command to check the context:

   ```
   ls -dZ /opt/netapp/data
   ```

   A sample output:

   ```
   drwxr-xr-x. mysql root unconfined_u:object_r:default_t:s0
   /opt/netapp/data
   ```

   In this output, the context is `default_t`. You should change this context to `mysqld_db_t`.

2. Perform these steps to set the context based on how you have mounted `/opt/netapp/data`.

   a. Run the following commands to set the context to `mysqld_db_t`:
   ```
   semanage fcontext -a -t mysqld_db_t "/opt/netapp/data"
   `restorecon -R -v /opt/netapp/data
   ```

b. If you have configured `/opt/netapp/data` in `/etc/fstab`, you should edit the `/etc/fstab` file. For the `/opt/netapp/data/` mount option, add the MySQL label as:

```
context=system_u:object_r:mysqld_db_t:s0
```

c. Unmount and remount `/opt/netapp/data/` for enabling the context.

d. If you have a direct NFS mount, run the following command to set the context to `mysqld_db_t`:

```
mount <nfsshare>:/<mountpoint> /opt/netapp/data -o
context=system_u:object_r:mysqld_db_t:s0
```

3. Verify whether the context is set correctly:

```
ls -dZ /opt/netapp/data/
```

A sample output:

```
drwxr-xr-x. mysql root unconfined_u:object_r:mysqld_db_t:s0
/opt/netapp/data/
```

Configure the SELinux context when `/opt/netapp` is mounted, and `/opt/netapp/data/` is also mounted separately

In this scenario, at first, you should set the context for `/opt/netapp/data/` as described in the previous section. After setting the correct context for `/opt/netapp/data/`, ensure that the parent directory `/opt/netapp` does not have the SELinux context set to `file_t`.

Steps

1. Run this command to check the context:

   ```
   ls -dZ /opt/netapp
   ```

   A sample output:

   ```
   drwxr-xr-x. mysql root unconfined_u:object_r:file_t:s0 /opt/netapp
   ```

   In this output, the context is `file_t` should be changed. The following commands set the context to `usr_t`. You can set the context to any value other than `file_t` based on your security requirements.

2. Perform these steps to set the context, based on how you have mounted `/opt/netapp`.

   a. Run the following commands to set the context:

```
semanage fcontext -a -t usr_t "/opt/netapp"
restorecon -v /opt/netapp
```

a. If you have configured `/opt/netapp` in `/etc/fstab`, you should edit the `/etc/fstab` file. For the

`/opt/netapp` mount option, add the MySQL label as:

```
context=system_u:object_r:usr_t:s0
```

b. Unmount, and then again mount `/opt/netapp` for enabling the context.

c. If you have a direct NFS mount, run the following command to set the context:

```
mount <nfsshare>:/<mountpoint> /opt/netapp -o
context=system_u:object_r:usr_t:s0
```

1. Verify whether the context is set correctly:

   ```
   ls -dZ /opt/netapp
   ```

   A sample output

```
drwxr-xr-x. mysql root unconfined_u:object_r:usr_t:s0 /opt/netapp
```

Configure the SELinux context when `/opt/netapp` is mounted, and `/opt/netapp/data/` is not mounted separately

If you have mounted `/opt/netapp` in your system and SELinux is set to `Enforcing`, ensure that the SELinux context type for `/opt/netapp` is set to `mysqld_db_t`, which is the default context element for the location of the database files.

Steps

1. Run this command to check the context:

   ```
   ls -dZ /opt/netapp
   ```

   A sample output:

   ```
   drwxr-xr-x. mysql root unconfined_u:object_r:default_t:s0 /opt/netapp
   ```

   In this output, the context is `default_t`. You should change this context to `mysqld_db_t`.

2. Perform the following steps to set the context based on how you have mounted `/opt/netapp`.

   a. Run the following commands to set the context to `mysqld_db_t`:
   `semanage fcontext -a -t mysqld_db_t "/opt/netapp"`
   `` `restorecon -R -v /opt/netapp ``

   b. If you have configured `/opt/netapp` in `/etc/fstab`, edit the `/etc/fstab` file. For the `/opt/netapp/` mount option, add the MySQL label as:
   `context=system_u:object_r:mysqld_db_t:s0`

   c. Unmount, and then again mount `/opt/netapp/` for enabling the context.

   d. If you have a direct NFS mount, run the following command to set the context to `mysqld_db_t`:
   `mount <nfsshare>:/<mountpoint> /opt/netapp -o`

```
context=system_u:object_r:mysqld_db_t:s0
```

3. Verify whether the context is set correctly:

```
ls -dZ /opt/netapp/
```

A sample output:

```
drwxr-xr-x. mysql root unconfined_u:object_r:mysqld_db_t:s0 /opt/netapp/
```

Install Unified Manager on Linux systems

Create a custom user home directory and umadmin password before you install Active IQ Unified Manager

You can create a custom home directory and define your own umadmin user password prior to installing Unified Manager. This task is optional, but some sites might need the flexibility to override Unified Manager installation default settings.

Before you begin

- The system must meet the requirements described in Hardware system requirements.
- You must be able to log in as the root user to the Red Hat Enterprise Linux system.

The default Unified Manager installation performs the following tasks:

- Creates the umadmin user with `/home/umadmin` as the home directory.
- Assigns the default password “admin” to the umadmin user.

Because some installation environments restrict access to `/home`, the installation fails. You must create the home directory in a different location. Additionally, some sites might have rules about password complexity or require that passwords be set by local administrators rather than being set by the installing program.

If your installation environment requires that you override these installation default settings, follow these steps to create a custom home directory and to define the umadmin user’s password.

When this information is defined prior to installation, the installation script discovers these settings and uses the defined values instead of using the installation default settings.

Additionally, the default Unified Manager installation includes the umadmin user in the sudoers files (`ocum_sudoers` and `ocie_sudoers`) in the `/etc/sudoers.d/` directory. If you remove this content from your environment because of security policies, or because of some security monitoring tool, you must add it back. You need to preserve the sudoers configuration because some Unified Manager operations require these sudo privileges.

The security policies in your environment must not restrict sudo privileges for the Unified Manager maintenance user. Some Unified Manager operations might fail if the privileges are restricted. Verify that you are able to run the following sudo command when logged in as the umadmin user after successful installation.

```
sudo systemctl status ocie
```

This command should return the appropriate status of the ocie service without any errors.

Steps

1. Log in as the root user to the server.
2. Create the umadmin group account called “maintenance”:

   ```
   groupadd maintenance
   ```

3. Create the user account “umadmin” in the maintenance group under a home directory of your choice:

   ```
   adduser --home <home_directory\> -g maintenance umadmin
   ```

4. Define the umadmin password:

   ```
   passwd umadmin
   ```

   The system prompts you to enter a new password string for the umadmin user.

After you have installed Unified Manager you must specify the umadmin user login shell.

Download Active IQ Unified Manager

You must download the Unified Manager `.zip` file from the NetApp Support Site to install Unified Manager.

Before you begin

You must have login credentials for the NetApp Support Site.

You download the Unified Manager installation package for the Red Hat Enterprise Linux system.

Steps

1. Log in to the NetApp Support Site, and navigate to the Download page for Unified Manager:

   NetApp Support Site

2. Select the required version of Unified Manager and accept the end-users license agreement (EULA).
3. Download the Unified Manager installer file for Linux, and save the `.zip` file to a directory on the target system.

   - Ensure that you download the correct version of the installer file for your Red Hat Enterprise Linux system. Based on whether you have Red Hat Enterprise Linux 8 or 9 installed, ensure that you download the appropriate version of the Unified Manager `.zip` file.
   - NetApp recommends that you download the code signing certificate (`.pem`) and digital signature (`.sig`) along with the `.zip` file.

4. Verify the checksum for integrity of the downloaded software.
5. If you have downloaded the code signing certificate and digital signature, you can verify the integrity of the installer file. You can use the following commands to verify the integrity of the installer file:
   - This command creates a file with the public key from the code signing certificate:

```
openssl x509 -pubkey -noout -in AIQUM-RHEL-CLIENT-INTER-ROOT.pem >
<public_key_file_name>
```

- Where **AIQUM-RHEL-CLIENT-INTER-ROOT.pem** is the file that contains the code signing certificate.

◦ This command verifies the signature on the installer file:

```
openssl dgst -sha256 -verify <public_key_file_name> -signature
<signature_file_name> ActiveIQUnifiedManager-<version>.zip
```

The message similar to `Verified Ok` confirms that the installer file is safe to use.

Install Active IQ Unified Manager on Red Hat Enterprise Linux

You can install Unified Manager on a physical or virtual Red Hat Enterprise Linux platform.

Before you begin

- The system on which you want to install Unified Manager must meet the system and software requirements.

  See Hardware system requirements.

  See Linux software and installation requirements.

- You must have downloaded the Unified Manager `.zip` file from the NetApp Support Site to the target system.
- You should have verified the integrity of the downloaded `.zip` file.
- You must have a supported web browser.
- Your terminal emulation software must have scrollback enabled.

The Red Hat Enterprise Linux system might have all the required versions of the required supporting software (Java, MySQL, additional utilities) installed, only some of the required software installed, or might be a newly installed system with none of the required software installed.

Steps

1. Log in to the server on which you are installing Unified Manager.
2. Enter the appropriate commands to assess what software might require installation or upgrade on the target system to support installation:

Required software and minimum version	Command to verify software and version
OpenJDK version 17.0.14+6	`java -version`
MySQL 8.4.10 Community Edition	`rpm -qa | grep -i mysql`
7-zip	`7zz -version`

3. If the installed version of MySQL is earlier than MySQL 8.4.10 Community Edition, enter the following command to uninstall it:

   ```
   rpm -e <mysql_package_name>
   ```

   If you receive dependency errors, you must add the `--nodeps` option to uninstall the component.

4. Navigate to the directory where you downloaded the installation `.zip` file and expand the Unified Manager bundle:

   ```
   unzip ActiveIQUnifiedManager-<version>.zip
   ```

   The required `.rpm` modules for Unified Manager are unzipped to the target directory.

5. Verify that the following module is available in the directory:

   ```
   ls *.rpm
   ```

   ```
   netapp-um<version>.x86_64.rpm
   ```

6. Run the pre-installation script to ensure that there are no system configuration settings or any installed software that might conflict with the installation of Unified Manager:

   ```
   sudo ./pre_install_check.sh
   ```

   The pre-installation script checks that the system has a valid Red Hat Enterprise Linux subscription, and that it has access to the required software repositories. If the script identifies any issues, you must fix the issues prior to installing Unified Manager.

   For a Red Hat Enterprise Linux 8 system, if you have internal repositories with JDK 17 - OpenJDK, 7-zip, and other software packages provided by the AppStream repository, you must disable your AppStream repository and install MySQL Community Server. Run the following command:

   ```
   # sudo yum --disablerepo=rhel-8-for-x86_64-appstream-rpms install
   mysql-community-server
   ```

7. **Optional:** You must perform step 7 only if your system is not connected to the internet and you have to manually download the packages that are required for your installation. If your system has the internet access and all the required packages are available, go to step 8. For systems that are not connected to the internet or that are not using the Red Hat Enterprise Linux repositories, perform the following steps to determine whether you are missing any required packages, and then download those packages:

   a. On the system on which you are installing Unified Manager, view the list of available and unavailable packages:

      ```
      yum install netapp-um<version>.x86_64.rpm --assumeno
      ```

      The items in the “Installing:” section are the packages that are available in the current directory, and the items in the “Installing for dependencies:” section are the packages that are missing on your system.

   b. On a system that has the internet access, download the missing packages:

      ```
      yum install <package_name> --downloadonly --downloaddir=.
      ```

Because the plug-in “yum-plugin-downloadonly” is not always enabled on Red Hat Enterprise Linux systems, you might need to enable the functionality to download a package without installing it:
`yum install yum-plugin-downloadonly`

c. Copy the missing packages from the internet-connected system to your installation system.

8. As the root user, or using `sudo`, run the following command to install the software:

   `yum install netapp-um<version>.x86_64.rpm`

   This command installs the .rpm packages, all other necessary supporting software, and the Unified Manager software.

   If the installation fails with the GPG NOKEY error , use `rpm --import` to import the keys from a URL:

   rpm --import https://repo.mysql.com/RPM-GPG-KEY-mysql-2023

   Do not attempt installation by using alternative commands (such as `rpm -ivh`). A successful installation of Unified Manager on a Red Hat Enterprise Linux system requires that all Unified Manager files and related files are installed in a specific order into a specific directory structure that is enforced automatically by the `yum install netapp-um<version>.x86_64.rpm` command.

9. Disregard the email notification that is displayed immediately after the installation messages.

   The email notifies the root user of an initial cron job failure, which has no adverse effect on the installation.

10. After the installation messages are complete, scroll back through the messages until you see the message in which the system displays an IP address or URL for the Unified Manager web UI, the maintenance user name (umadmin), and a default password.

    The message is similar to the following:

```
Active IQ Unified Manager installed successfully.
Use a web browser and one of the following URL(s) to configure and
access the Unified Manager GUI.
https://default_ip_address/    (if using IPv4)
https://[default_ip_address]/  (if using IPv6)
https://fully_qualified_domain_name/

Log in to Unified Manager in a web browser by using following details:
  username: umadmin
  password: admin
```

11. Record the IP address or URL, the assigned user name (umadmin), and the current password.
12. If you created a umadmin user account with a custom home directory prior to installing Unified Manager, then you must specify the umadmin user login shell:

    `usermod -s /bin/maintenance-user-shell.sh umadmin`

Access the web UI to change the default password of the umadmin user, and perform the initial setup of Unified Manager, as described in Configuring Active IQ Unified Manager. It is mandatory to change the default password of the umadmin user.

Learn about the users created during Active IQ Unified Manager installation

When you install Unified Manager on Red Hat Enterprise Linux, the following users are created by Unified Manager and third-party utilities: umadmin, jboss, and mysql.

- **umadmin**

  Used to log in to Unified Manager for the first time. This user is assigned an “Application Administrator” user role and is configured as the “Maintenance User” type. This user is created by Unified Manager.

- **jboss**

  Used to run Unified Manager services related to the JBoss utility. This user is created by Unified Manager.

- **mysql**

  Used to run MySQL database queries of Unified Manager. This user is created by the MySQL third-party utility.

In addition to these users, Unified Manager also creates corresponding groups: maintenance, jboss, and mysql. The maintenance and jboss groups are created by Unified Manager, while the mysql group is created by a third-party utility.

If you created a custom home directory and defined your own umadmin user password prior to installing Unified Manager, the installation program does not recreate the maintenance group or the umadmin user.

Change the JBoss password in Active IQ Unified Manager

You can reset the instance-specific JBoss password that is set during installation. You can reset the password optionally, in case your site requires this security capability to override the Unified Manager installation setting. This operation also changes the password JBoss uses to access MySQL.

- You must have root user access to the Red Hat Enterprise Linux system on which Unified Manager is installed.
- You must be able to access the NetApp-provided `password.sh` script in the directory `/opt/netapp/essentials/bin`.

Steps

1. Log in as root user on the system.
2. Stop the Unified Manager services by entering the following commands in the order shown:

```
systemctl stop ocieau

systemctl stop ocie
```

Do not stop the associated MySQL software.

3. Enter the following command to begin the password change process:

   ```
   /opt/netapp/essentials/bin/password.sh resetJBossPassword
   ```

4. When prompted, enter the new JBoss password, and then enter it a second time for confirmation.

   Note that the password should be between 8 and 16 characters, and must contain at least a digit, an upper case and lower case characters, and at least one of these special characters:

   ```
   !@%^*-_=[]:<>.?/~+
   ```

5. When the script completes, start the Unified Manager services by entering the following commands in the order shown:

   ```
   systemctl start ocie

   systemctl start ocieau
   ```

6. After all of the services are started, you can log in to the Unified Manager UI.

Upgrade Unified Manager on Red Hat Enterprise Linux

Upgrade Active IQ Unified Manager on Red Hat Enterprise Linux

You can upgrade Unified Manager when a new version is available.

Patch releases of Unified Manager software, when provided by NetApp, are installed using the same procedure as new releases.

If Unified Manager is paired with an instance of OnCommand Workflow Automation, and there are new versions of software available for both products, you must disconnect the two products and then set up a new Workflow Automation connection after performing the upgrades. If you are performing an upgrade to only one of the products, then you should log into Workflow Automation after the upgrade and verify that it is still acquiring data from Unified Manager.

Use the supported upgrade path for versions of Active IQ Unified Manager

Active IQ Unified Manager supports a specific upgrade path for each version.

Not all versions of Unified Manager can perform an in-place upgrade to later versions. The Unified Manager upgrades are limited to an N-2 model. This means that you can only perform an upgrade to the next two releases on all platforms. For example, you can perform an upgrade to Unified Manager 9.18 from Unified Manager 9.14 and 9.16 only.

If you are running a version that is earlier than the supported versions, your Unified Manager instance will need to be upgraded to one of the supported versions first, and then upgraded to the current version.

For example, if your installed version is Unified Manager 9.9 and you want to upgrade to Unified Manager 9.14, you follow a sequence of upgrades.

Sample upgrade path:

1. Upgrade 9.11 → 9.13
2. Upgrade 9.13 → 9.14
3. Upgrade 9.13 → 9.16
4. Upgrade 9.14 or 9.16 → 9.18

For more information on the upgrade path matrix, see this knowledge base (KB) article.

Upgrade Active IQ Unified Manager

You can upgrade from Unified Manager 9.14 or 9.16 to 9.18 by downloading and running the installation file on the Linux platform.

Before you begin

- The system on which you are upgrading Unified Manager must meet the system and software requirements.

  See Hardware system requirements.

  See Linux software and installation requirements.

- You must have a subscription to the Red Hat Enterprise Linux Subscription Manager.
- You must install, or upgrade to the correct version of OpenJDK prior to upgrading Unified Manager.

  See Upgrading JRE on Linux.

- You must install or upgrade to the correct version of 7-zip prior to upgrading Unified Manager. See Configure the 7-zip software package.
- To avoid data loss, you must have created a backup of the Unified Manager database in case there is an issue during the upgrade. NetApp recommends that you move the backup file from the `/opt/netapp/data` directory to an external location.
- During upgrade, you may be prompted to confirm whether you want to keep the previous, default settings for retaining performance data for 13 months or to change it to 6 months. On confirming, the historical performance data is purged after 6 months.
- You should have completed any running operations, because Unified Manager is unavailable during the upgrade process.
- MySQL Community Edition is automatically upgraded during Unified Manager upgrade. If the installed version of MySQL on your system is earlier than 8.4.10, the Unified Manager upgrade process automatically upgrades MySQL to 8.4.10.

Steps

1. Log in to the target Red Hat Enterprise Linux server.
2. Download the Unified Manager bundle to the server.

   See Downloading Unified Manager for Linux.

3. Navigate to the target directory and expand the Unified Manager bundle:

   ```
   unzip ActiveIQUnifiedManager-<version>.zip
   ```

   The required RPM modules for Unified Manager are unzipped to the target directory.

4. Verify that the following module is available in the directory:

   ```
   ls *.rpm
   ```

   ```
   netapp-um<version>.x86_64.rpm
   ```

5. Run the pre-installation script to ensure that there are no system configuration settings or any installed software that might conflict with the upgrade:

   ```
   sudo ./pre_install_check.sh
   ```

   The pre-installation script checks that the system has a valid Red Hat Enterprise Linux subscription, and that it has access to the required software repositories. If the script identifies any issues, you must fix the issues and continue with the upgrade.

   If there are any missing packages detected, perform the steps mentioned in Additional steps to perform for missing packages. If there are no missing packages, continue with the next steps.

6. Upgrade Unified Manager using the following script:

   ```
   upgrade.sh
   ```

   This script automatically executes the RPM modules, upgrading the necessary supporting software and the Unified Manager modules that run on them. Additionally, the upgrade script checks whether there are any system configuration settings or any installed software that might conflict with the upgrade. If the script identifies any issues, you must fix the issues before upgrading Unified Manager. If you have previously installed packages, such as *net-snmp* before you upgraded Unified Manager, the MySQL dependency might uninstall the package during upgrade. You need to install the package manually again for continuing to use it.

7. After the upgrade is complete, scroll back through the messages until you see the message displaying an IP address or URL for the Unified Manager web UI, the maintenance user name (umadmin), and the default password.

   The message is similar to the following:

   ```
   Active IQ Unified Manager upgraded successfully.
   Use a web browser and one of the following URLs to access the Unified
   Manager GUI:

   https://default_ip_address/    (if using IPv4)
   https://[default_ip_address]/  (if using IPv6)
   https://fully_qualified_domain_name/
   ```

Enter the specified IP address or URL in a new window of a supported web browser to start the Unified Manager web UI, and then log in by using the same maintenance user name (umadmin) and password that you set earlier.

Additional steps to perform for missing packages

If there are any missing packages detected at your site during upgrade, or if your system is not connected to the internet, or you are not using the Red Hat Enterprise Linux repositories, perform the following steps to

determine whether you are missing any required packages and download those packages.

> These steps need to be performed after step *5* of the main procedure. This procedure upgrades Unified Manager and you do not need to run any additional steps for upgrade.

1. View the list of available and unavailable packages:

   ```
   yum install netapp-um<version>.x86_64.rpm --assumeno
   ```

   The items in the “Installing:” section are the packages that are available in the current directory, and the items in the “Installing for dependencies:” section are the packages that are missing on your system.

2. On a different system that has the internet access, run the following command to download the missing packages.

   ```
   yum install package_name --downloadonly --downloaddir=.
   ```

   The packages are downloaded in the directory specified as `--downloaddir=`.

   Because the plug-in “yum-plugin-downloadonly” is not always enabled on Red Hat Enterprise Linux systems, you might need to enable the functionality to download a package without installing it:

   ```
   yum install yum-plugin-downloadonly
   ```

3. Copy the downloaded packages to the directory where you unzipped the Unified Manager bundle on the installation system.
4. Change directories to that directory and run the following command for installing the missing packages, along with their dependencies.

   ```
   yum install *.rpm
   ```

5. Start the Unified Manager server. Run these commands:

   ```
   systemctl start ocie

   systemctl start ocieau
   ```

This process completes the Unified Manager upgrade process. Enter the specified IP address or URL in a new window of a supported web browser to start the Unified Manager web UI, and then log in by using the same maintenance user name (umadmin) and password that you set earlier.

Upgrade the Red Hat Enterprise Linux host OS for Active IQ Unified Manager

If you previously installed Unified Manager on a Red Hat Enterprise Linux 7.x system and need to upgrade to Red Hat Enterprise Linux 8.x, you must follow one of the procedures listed in this topic. In both cases you must create a backup of Unified Manager on the Red Hat Enterprise Linux 7.x system, and then restore the backup onto a Red Hat Enterprise Linux 8.x system. Note that the supported versions of Red Hat Enterprise Linux are from 8.0 to 8.10.

The difference between the two options listed below is that in one case you are performing the Unified Manager restore onto a new 8.x server, and in the other case you are performing the restore operation onto the

same server.

Because this task requires that you create a backup of Unified Manager on the Red Hat Enterprise Linux 7.x system, you should create the backup only when you are prepared to complete the entire upgrade process so that Unified Manager is offline for the shortest period of time. Gaps in collected data appear in the Unified Manager UI for the period of time during which the Red Hat Enterprise Linux 7.x system is shut down and before the new Red Hat Enterprise Linux 8.x is started.

See Managing backup and restore operations if you need to review detailed instructions for the backup and restore processes.

Follow these steps if you have a spare system on which you can install Red Hat Enterprise Linux 8.x software so that you can perform the Unified Manager restore on that system while the Red Hat Enterprise Linux 7.x system is still available.

1. Install and configure a new server with Red Hat Enterprise Linux 8.x software.

   See Linux software and installation requirements.

2. On the Red Hat Enterprise Linux 8.x system, install the same version of Unified Manager software that you have on the existing Red Hat Enterprise Linux 7.x system.

   See Installing Unified Manager on Linux.

   Do not launch the UI or configure any clusters, users, or authentication settings when the installation is complete. The backup file populates this information during the restore process.

3. On the Red Hat Enterprise Linux 7.x system, from the Administration menu in the web UI, create a Unified Manager backup and then copy the backup file (`.7z` file) and the contents of the database repository directory (`/database-dumps-repo` subdirectory) to an external location.
4. On the Red Hat Enterprise Linux 7.x system, shut down Unified Manager.
5. On the Red Hat Enterprise Linux 8.x system, copy the backup file (`.7z` file) from the external location to `/opt/netapp/data/ocum-backup/` and the database repository files to the `/database-dumps-repo` subdirectory under the `/ocum-backup` directory.
6. Enter the following command to restore the Unified Manager database from the backup file:

   ```
   um backup restore -f /opt/netapp/data/ocum-backup/<backup_file_name>
   ```

7. Enter the IP address or URL into your web browser to start the Unified Manager web UI, and then log in to the system.

Once you have verified that the system is operating properly you can remove Unified Manager from the Red Hat Enterprise Linux 7.x system.

Upgrading the host OS on the same server

Follow these steps if you do not have a spare system on which you can install Red Hat Enterprise Linux 8.x software.

1. From the Administration menu in the web UI, create a Unified Manager backup and then copy the backup file (`.7z` file) and the contents of the database repository directory (`/database-dumps-repo` subdirectory) to an external location.
2. Remove the Red Hat Enterprise Linux 7.x image from the system and completely wipe the system.

3. Install and configure Red Hat Enterprise Linux 8.x software on the same system.

   See Linux software and installation requirements.

4. On the Red Hat Enterprise Linux 8.x system, install the same version of Unified Manager software that you had on the Red Hat Enterprise Linux 7.x system.

   See Installing Unified Manager on Linux.

   Do not launch the UI or configure any clusters, users, or authentication settings when the installation is complete. The backup file populates this information during the restore process.

5. Copy the backup file (`.7z` file) from the external location to `/opt/netapp/data/ocum-backup/` and the database repository files to the `/database-dumps-repo` subdirectory under the `/ocum-backup` directory.
6. Enter the following command to restore the Unified Manager database from the backup file:

   ```
   um backup restore -f /opt/netapp/data/ocum-backup/<backup_file_name>
   ```

7. Enter the IP address or URL into your web browser to start the Unified Manager web UI, and then log in to the system.

Upgrade third-party products post Unified Manager installation

Upgrade third-party products after installing Active IQ Unified Manager

You can upgrade third-party products, such as JRE, when Unified Manager is already installed on Linux systems.

The companies that develop these third-party products report security vulnerabilities on a regular basis. You can upgrade to newer versions of this software at your own schedule.

Upgrade OpenJDK on Linux for Active IQ Unified Manager

You can upgrade to a newer version of OpenJDK on the Linux server on which Unified Manager is installed to obtain fixes for security vulnerabilities.

Before you begin

You must have root privileges for the Linux system on which Unified Manager is installed.

You can update OpenJDK releases within release families. For example, you can upgrade from OpenJDK 11.0.14 to OpenJDK 11.0.17, but you cannot update directly from OpenJDK 11 to OpenJDK 12.

Steps

1. Log in as a root user on the Unified Manager host machine.
2. Download the appropriate version of OpenJDK (64-bit) to the target system.
3. Stop the Unified Manager services:

   ```
   systemctl stop ocieau

   systemctl stop ocie
   ```

4. Install the latest OpenJDK on the system.
5. Start the Unified Manager services:

```
systemctl start ocie

systemctl start ocieau
```

Restart Active IQ Unified Manager

You might have to restart Unified Manager after making configuration changes.

Before you begin

You must have root user access to the Red Hat Enterprise Linux server on which Unified Manager is installed.

Steps

1. Log in as root user to the server on which you want to restart the Unified Manager service.
2. Stop the Unified Manager service and the associated MySQL service in this order:

```
systemctl stop ocieau

systemctl stop ocie

systemctl stop mysqld
```

3. Start MySQL and Unified Manager services in this order:

```
systemctl start mysqld

systemctl start ocie

systemctl start ocieau
```

`mysqld` is a daemon program required to start and stop the MySQL server.

Remove Active IQ Unified Manager

You can stop and uninstall Unified Manager from the Red Hat Enterprise Linux host with a single command.

Before you begin

- You must have root user access to the server from which you want to remove Unified Manager.
- Security-Enhanced Linux (SELinux) must be disabled on the Linux system. Change the SELinux runtime mode to “Permissive” by using the `setenforce 0` command.
- All clusters (data sources) must be removed from the Unified Manager server before removing the software.
- You should manually delete the firewall rules that are created to allow or block MySQL port 3306. The firewall rules are not deleted automatically.

Steps

1. Log in as root user to the server on which you want to remove Unified Manager.
2. Stop and remove Unified Manager from the server:

   ```
   rpm -e netapp-um
   ```

   This step removes all the associated NetApp RPM packages. It does not remove the prerequisite software modules, such as Java, MySQL, and 7zip.

3. **Optional:** If appropriate, remove the supporting software modules, such as Java, MySQL, and 7zip:

   ```
   rpm -e mysql-community-client mysql-community-server mysql-community-common
   mysql-community-libs java-x.y
   ```

After this operation is complete, the software is removed. All the data from the `/opt/netapp/data` directory is moved to the `/opt/netapp/data/BACKUP` folder after uninstallation. Uninstalling Unified Manager removes the Java and MySQL packages also, unless the packages are required and used by any other application on the system. However, MySQL data is not deleted.

Remove the custom umadmin user and maintenance group for Active IQ Unified Manager

If you created a custom home directory to define your own umadmin user and maintenance account prior to installing Unified Manager, you should remove these items after you have uninstalled Unified Manager.

The standard Unified Manager uninstallation does not remove a custom-defined umadmin user and maintenance account. You must delete these items manually.

Steps

1. Log in as the root user to the Red Hat Enterprise Linux server.
2. Delete the umadmin user:

   ```
   userdel umadmin
   ```

3. Delete the maintenance group:

   ```
   groupdel maintenance
   ```

Install Unified Manager on Windows systems

Introduction to Active IQ Unified Manager

Learn about Active IQ Unified Manager

Active IQ Unified Manager (formerly OnCommand Unified Manager) enables you to monitor and manage the health and performance of your ONTAP storage systems from a single interface. You can deploy Unified Manager on a Linux server, on a Windows server, or as a virtual appliance on a VMware host.

After you have completed the installation and have added the clusters that you want to manage, Unified Manager provides a graphical interface that displays the capacity, availability, protection, and performance status of the monitored storage systems.

Related information

NetApp Interoperability Matrix Tool

Learn what the Active IQ Unified Manager server does

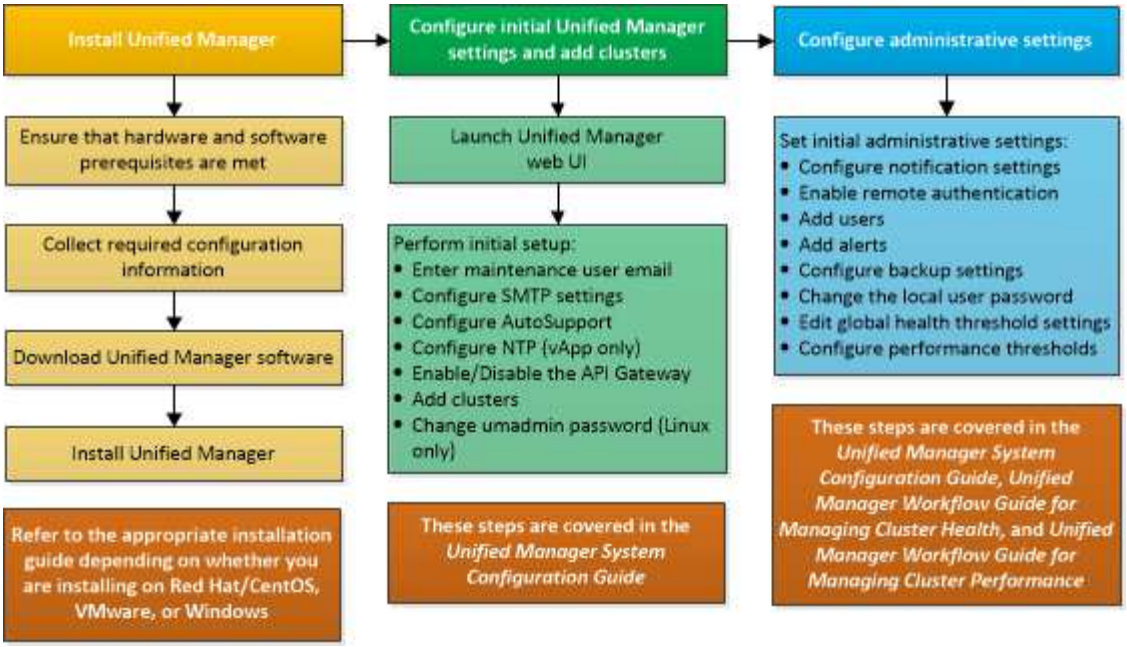
The Unified Manager server infrastructure consists of a data collection unit, a database, and an application server. It provides infrastructure services such as discovery, monitoring, role-based access control (RBAC), auditing, and logging.

Unified Manager collects cluster information, stores the data in the database, and analyzes the data to see if there are any cluster issues.

Learn about the installation sequence for Active IQ Unified Manager

The installation workflow describes the tasks that you must perform before you can use Unified Manager.

These sections describe each of the items shown in the workflow below.



Requirements for installing Unified Manager

Ensure the server meets Active IQ Unified Manager installation requirements

Before you begin the installation process, ensure that the server on which you want to install Unified Manager meets the specific software, hardware, CPU, and memory requirements.

NetApp does not support any modification of the Unified Manager application code. If you need to apply any security measures to the Unified Manager server, you should make those changes to the operating system on which Unified Manager is installed.

For more details about applying security measures to the Unified Manager server, see the Knowledge Base article.

Supportability for Security Measures applied to Active IQ Unified Manager for Clustered Data ONTAP

Related information

NetApp Interoperability Matrix Tool

Understand virtual infrastructure and hardware requirements for Active IQ Unified Manager

Installing Unified Manager on virtual infrastructure or a physical system should meet the minimum requirements for memory, CPU, and disk space.

The following table displays the values that are recommended for memory, CPU, and disk space resources. These values have been qualified so that Unified Manager meets acceptable performance levels.

Hardware configuration	Recommended settings
RAM	12 GB

Hardware configuration	Recommended settings
Processors	4 CPUs
CPU cycle capacity	9572 MHz total (minimum requirement 9572 MHz)
Free disk space	150 GB, where the capacity is allocated as follows: • 100 GB of disk space for the installation directory • 50 GB of disk space for the MySQL data directory

Unified Manager can be installed on systems with a small amount of memory, but the recommended 12 GB of RAM ensures that enough memory is available for optimal performance, and so that the system can accommodate additional clusters and storage objects as your configuration grows. You should not set any memory limits on the VM where Unified Manager is deployed, and should not enable any features (for example, ballooning) that hinder the software from utilizing the allocated memory on the system.

Additionally, there is a limit to the number of nodes that a single instance of Unified Manager can monitor before you install a second instance of Unified Manager. For node limits, CPU, and memory recommendations, see the Unified Manager Best Practices Guide.

Memory-page swapping negatively impacts the performance of the system and the management application. Competing for CPU resources that are unavailable because of overall host utilization can degrade performance.

Requirement for dedicated use

The physical or virtual system on which you install Unified Manager should be used exclusively for Unified Manager and should not be shared with other applications. Other applications might consume system resources and can drastically reduce the performance of Unified Manager.

Space requirements for backups

If you plan to use the Unified Manager backup and restore feature, allocate additional capacity so that the “data” directory or disk has 150 GB of space. A backup can be written to a local destination or to a remote destination. The best practice is to identify a remote location that is external to the Unified Manager host system that has a minimum of 150 GB of space.

Requirements for host connectivity

The physical system or virtual system on which you install Unified Manager should be configured in such a way that you can successfully `ping` the host name from the host itself. In case of IPv6 configuration, you should verify that `ping6` to the host name is successful to ensure that the Unified Manager installation succeeds.

You can use the host name (or the host IP address) to access the product web UI. If you configured a static IP address for your network during deployment, then you designated a name for the network host. If you configured the network using DHCP, you should obtain the host name from the DNS.

If you plan to allow users to access Unified Manager by using the short name instead of using the fully qualified domain name (FQDN) or IP address, then your network configuration has to resolve this short name to a valid FQDN.

Understand Windows software and installation requirements for Active IQ Unified Manager

For the successful installation of Unified Manager on Windows, you should ensure that the system on which Unified Manager is being installed meets the software requirements.

Operating system software

You can install Unified Manager on the following Windows editions:

- Microsoft Windows Server 2019 Standard and Datacenter Edition
- Microsoft Windows Server 2022 Standard and Datacenter Edition

Unified Manager is supported on 64 bits Windows operating system for the following languages:

- English
- Japanese
- Simplified Chinese

See the Interoperability Matrix for the complete and most current list of supported Windows versions.

mysupport.netapp.com/matrix

NetApp does not support installation of Unified Manager by using third-party tools, such as Microsoft System Center Configuration Manager (SCCM).

The server should be dedicated to running Unified Manager. No other applications should be installed on the server. It is possible that an active anti-virus software is installed on your Windows system because of company regulations. You should disable the anti-virus software before installing Unified Manager to prevent the installation from failing.

Third-party software

The following third-party packages are bundled with Unified Manager. If these third-party packages are not installed on your system, Unified Manager installs them as a part of the installation.

- Microsoft Visual C++ Redistributable Packages for Visual Studio 2015 version 14.29.30133.0
- MySQL Community Edition version 8.4.10
- Python 3.13.3
- OpenJDK version 17.0.15+6
- 7-zip version 25.01 or later

Starting with Unified Manager 9.5, OpenJDK is provided in the Unified Manager installation package and installed automatically. Oracle Java is not supported starting with Unified Manager 9.5.

If MySQL is pre-installed, you should ensure that:

- It is using the default port.

- The sample databases are not installed.
- The service name is “MYSQL84”.

Unified Manager is deployed on a WildFly web server. WildFly 27.0.1.Final is bundled and configured with Unified Manager.

> You should shut down a running instance of Unified Manager before upgrading any third-party software. After the third-party software installation is complete, you can restart Unified Manager.

Installation requirements

- Microsoft .NET 4.5.2, or greater, should be installed.
- The `temp` directory should be configured with 2 GB of disk space for extracting the installation files. To verify whether the directory is created, run the following command on the command-line interface: `echo %temp%`
- You should reserve 2 GB of disk space in the Windows drive for caching the Unified Manager MSI files.
- The Microsoft Windows Server on which you want to install Unified Manager should be configured with a fully qualified domain name (FQDN) such that `ping` responses to the host name and FQDN are successful.
- You should disable Microsoft IIS worldwide web publishing service and ensure that ports 80 and 443 are free.
- You should make sure that the Remote Desktop Session Host setting for “Windows Installer RDS Compatibility” is disabled during the installation.
- UDP port 514 should be free, and should not be used by any other service.
- You should disable all anti-virus software on your system before installing Unified Manager. After your installation is complete, make sure to manually exclude the following paths from anti-virus scan:
  - Unified Manager data directory, for example, `C:\ProgramData\NetApp\OnCommandAppData\`
  - Unified Manager installation directory, for example, `\C:\Program Files\NetApp\`
  - MySQL data directory, for example, `C:\ProgramData\MySQL\MySQLServerData`

Use supported browsers in Active IQ Unified Manager

To access the Unified Manager web UI, use a supported browser.

The Interoperability Matrix has the list of supported browser versions.

mysupport.netapp.com/matrix

For all browsers, disabling pop-up blockers ensures that software features are displayed properly.

If you plan to configure Unified Manager for SAML authentication, so that an identity provider (IdP) can authenticate users, you should check the list of browsers supported by the IdP as well.

Understand protocol and port requirements for Active IQ Unified Manager

The required ports and protocols enable communication between the Unified Manager server and the managed storage systems, servers, and other components.

Connections to the Unified Manager server

In typical installations you do not have to specify port numbers when connecting to the Unified Manager web UI, because default ports are always used. For example, because Unified Manager always attempts to run on its default port, you can enter `https://<host>` instead of `https://<host>:443`.

The Unified Manager server uses specific protocols to access the following interfaces:

Interface	Protocol	Port	Description
Unified Manager web UI	HTTP	80	Used to access the Unified Manager web UI; automatically redirects to the secure port 443.
Unified Manager web UI and programs using APIs	HTTPS	443	Used to securely access the Unified Manager web UI or to make API calls; API calls can only be made using HTTPS.
Maintenance console	SSH/SFTP	22	Used to access the maintenance console and retrieve support bundles.
Linux command line	SSH/SFTP	22	Used to access the Red Hat Enterprise Linux command line and retrieve support bundles.
Syslog	UDP	514	Used to access subscription-based EMS messages from ONTAP systems and to create events based on the messages.
REST	HTTPS	9443	Used to access realtime REST API-based EMS events from authenticated ONTAP systems.
MySQL database	MySQL	3306	Used to enable OnCommand Workflow Automation and OnCommand API Services access to Unified Manager.
AMQP QPID broker	TCP/IP	56072	Used for internal message communication.
AMQP QPID broker	WebSocket over TCP	56080	Used to listen on this port for messages received from ONTAP (cloud agent).

Interface	Protocol	Port	Description
AMQP QPID broker	WebSocket over TCP	56443	Used to listen on this port for messages received from ONTAP (cloud agent).The communication over this port supports encryption provided by TLS or SSL.

The default port for MySQL, 3306, is restricted only to localhost while installing Unified Manager on Windows systems. Enable the firewall to restrict access for the port MySQL, 3306, once the installation is complete. This does not impact any upgrade scenario, where the previous configuration is maintained. This configuration can be modified, and the connection can be made available to other hosts by using the `Control access to MySQL port 3306` option on the maintenance console. For information, see Additional menu options. The ports used for HTTP and HTTPS communication (ports 80 and 443) can be changed using the Unified Manager maintenance console. For more information, see Configuring Active IQ Unified Manager.

Connections from the Unified Manager server

You should configure your firewall to open ports that enable communication between the Unified Manager server and managed storage systems, servers, and other components. If a port is not open, communication fails.

Depending on your environment, you can choose to modify the ports and protocols used by the Unified Manager server to connect to specific destinations.

The Unified Manager server connects using the following protocols and ports to the managed storage systems, servers, and other components:

Destination	Protocol	Port	Description
Storage system	HTTPS	443/TCP	Used to monitor and manage storage systems.
Storage system	NDMP	10000/TCP	Used for certain Snapshot restore operations.
AutoSupport server	HTTPS	443	Used to send AutoSupport information. Requires the internet access to perform this function.
Authentication server	LDAP	389	Used to make authentication requests, and user and group lookup requests.
	LDAPS	636	Used for secure LDAP communication.
Mail server	SMTP	25	Used to send alert notification emails.

Destination	Protocol	Port	Description
SNMP trap sender	SNMPv1 or SNMPv3	162/UDP	Used to send alert notification SNMP traps.
External data provider server	TCP	2003	Used to send performance data to an external data provider, such as Graphite.
NTP server	NTP	123/UDP	Used to synchronize the time on the Unified Manager server with an external NTP time server. (VMware systems only)
AMQP QPID broker	TCP/IP	56072	Used for internal message communication.
AMQP QPID broker	WebSocket over TCP	56080	Used to listen on this port for messages received from ONTAP (cloud agent).
AMQP QPID broker	WebSocket over TCP	56443	Used to listen on this port for messages received from ONTAP (cloud agent). The communication through this port supports encryption provided by TLS or SSL.
Syslog	UDP	514	Used by Unified Manager to send audit logs to a remote syslog server.

Complete the worksheet in Active IQ Unified Manager

Before you install and configure Unified Manager, you should have specific information about your environment readily available. You can record the information in the worksheet.

Unified Manager installation information

The details required to install Unified Manager.

System on which software is deployed	Your value
Host fully qualified domain name	
Host IP address	
Network mask	
Gateway IP address	
Primary DNS address	

System on which software is deployed	Your value
Secondary DNS address	
Search domains	
Maintenance user name	
Maintenance user password	

Unified Manager configuration information

The details to configure Unified Manager after installation. Some values are optional depending on your configuration.

Setting	Your value
Maintenance user email address	
SMTP server host name or IP address	
SMTP user name	
SMTP password	
SMTP port	25 (Default value)
Email from which alert notifications are sent	
Authentication server host name or IP address	
Active Directory administrator name or LDAP bind distinguished name	
Active Directory password or LDAP bind password	
Authentication server base distinguished name	
Identity provider (IdP) URL	
Identity provider (IdP) metadata	
SNMP trap destination host IP addresses	
SNMP port	

Cluster information

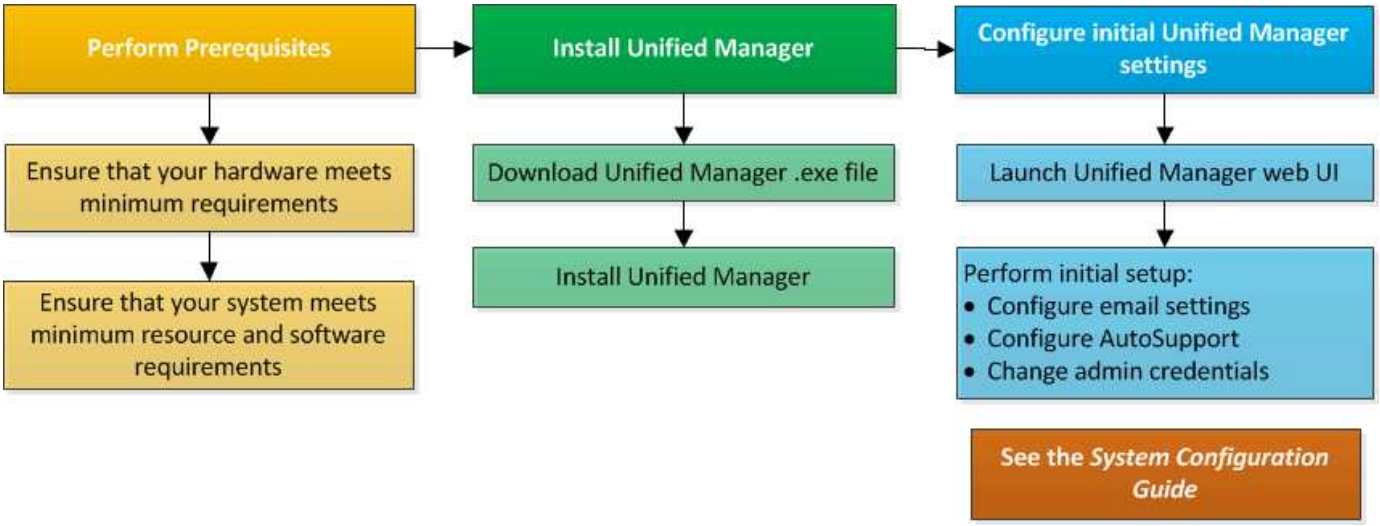
The details for the storage systems that you manage using Unified Manager.

Cluster 1 of N	Your value
Host name or cluster-management IP address	
ONTAP administrator user name The administrator must have been assigned the “admin” role.	
ONTAP administrator password	
Protocol	HTTPS

Install, upgrade, and remove Unified Manager software

Learn about the installation process for Active IQ Unified Manager on Windows

The installation workflow describes the tasks that you must perform before you can use Unified Manager.



Install Unified Manager on Windows

Install Active IQ Unified Manager on Windows

You can install Unified Manager to monitor and troubleshoot data storage capacity, availability, performance, and protection issues.

Before you begin

- The system on which you plan to install Unified Manager should meet the system and software requirements.

  See Hardware system requirements.

  See Windows software and installation requirements.

  Starting with Unified Manager 9.5, OpenJDK is provided in the installation package and installed automatically. Oracle Java is not supported starting with Unified Manager 9.5.

- You should have Windows administrator privileges. Ensure that your user name does not begin with an exclamation mark “!”. Installation of Unified Manager might fail if the user name of user running the installation begins with “!”.
- All clusters are supported, except for ASA r2 systems.
- You should have a supported web browser.
- The Unified Manager maintenance user password should be between 8 and 20 characters, should contain upper-case letters or lower-case letters, numerals, and special characters.
- The following special characters are not allowed in the password string for the maintenance user or for the MySQL root user: " ' ` % , = & < > | ^ \ / ( ) [ ] ;:

The following special characters are allowed: ~ ! @ # $ * - ? . + { }

Steps

1. Log in to Windows using the default local administrator account.
2. Log in to the NetApp Support Site, and navigate to the Download page for Unified Manager:

   NetApp Support Site

3. Select the required version of Unified Manager and accept the end-users license agreement (EULA).
4. Download the Unified Manager Windows installation file or zip file to a target directory on the Windows system.
5. Extract the zip file, if needed, and navigate to the directory where the installation file is located.
6. Right-click and run the Unified Manager installer executable (.exe) file as an administrator.

   Unified Manager detects missing or pre-installed third-party packages and lists them. If the required third-party packages are not installed in the system, Unified Manager installs them as part of the installation.

7. Click **Next**.
8. Enter the user name and password to create the maintenance user.
9. In the Database Connection wizard, enter the MySQL root password.
10. Click **Change** to specify a new location for the Unified Manager installation directory and MySQL data directory.

    If you do not change the installation directory, Unified Manager is installed in the default installation directory.

11. Click **Next**.
12. In the Ready to Install Shield wizard, click **Install**.
13. After the installation is complete, click **Finish**.
14. If you have an active anti-virus software installed on your Windows system, manually exclude the following paths from anti-virus scan after your installation is complete:
    - Unified Manager data directory
    - Unified Manager installation directory
    - MySQL data directory

The installation creates multiple directories:

- Installation directory

  This is the root directory for Unified Manager, which you specified during installation. Example: `C:\Program Files\NetApp\`

- MySQL data directory

  This is the directory where the MySQL databases are stored, which you specified during installation. Example: `C:\ProgramData\MySQL\MySQLServerData\`

- Java directory

This is the directory where OpenJDK is installed. Example: `C:\Program Files\NetApp\JDK\`

- Unified Manager application data directory (appDataDir)

  This is the directory where all the application-generated data is stored. This includes logs, support bundles, backup, and all other additional data. Example: `C:\ProgramData\NetApp\OnCommandAppData\`

You can access the web UI to perform the initial setup of Unified Manager, as described in the Configuring Active IQ Unified Manager.

Perform an unattended installation of Active IQ Unified Manager

You can install Unified Manager without user intervention by using the command-line interface. You can complete the unattended installation by passing the parameters in key-value pairs.

Steps

1. Log in to the Windows command-line interface by using the default local administrator account.
2. Navigate to the location where you want to install Unified Manager, and choose one of the following options:

Option	Instructions

If third-party packages are pre-installed	`ActiveIQUnifiedManager-x.y.exe /V"MYSQL_PASSWORD=mysql_password INSTALLDIR=\"Installation directory\" MYSQL_DATA_DIR=\"MySQL data directory\" MAINTENANCE_PASSWORD=maintenance_password MAINTENANCE_USERNAME=maintenance_username /qn /l*v CompletePathForLogFile"`
	
	Example:
	
	`ActiveIQUnifiedManager.exe /s /v"MYSQL_PASSWORD=netapp21! INSTALLDIR=\"C:\Program Files\NetApp\" MYSQL_DATA_DIR=\"C:\ProgramData\MYSQL\MySQLServer\" MAINTENANCE_PASSWORD=* MAINTENANCE_USERNAME=admin /qn /l*v C:\install.log"`

Option	Instructions

If third-party packages are not installed	`ActiveIQUnifiedManager-x.y.exe /V"MYSQL_PASSWORD=mysql_password INSTALLDIR=\"Installation directory\" MYSQL_DATA_DIR=\"MySQL data directory\" MAINTENANCE_PASSWORD=maintenance_password MAINTENANCE_USERNAME=maintenance_username /qr /l*v CompletePathForLogFile"`
	
	Example:
	
	`ActiveIQUnifiedManager.exe /s /v"MYSQL_PASSWORD=netapp21! INSTALLDIR=\"C:\Program Files\NetApp\" MYSQL_DATA_DIR=\"C:\ProgramData\MYSQL\MySQLServer\" MAINTENANCE_PASSWORD=* MAINTENANCE_USERNAME=admin /qr /l*v C:\install.log"`

The `/qr` option enables quiet mode with a reduced user interface. A basic user interface is displayed, which shows the installation progress. You are not prompted for inputs. If third-party packages such as JRE, MySQL, and 7zip are not pre-installed, you should use the `/qr` option. Installation fails if the `/qn` option is used on a server where third-party packages are not installed.

The `/qn` option enables quiet mode with no user interface. No user interface or details are displayed during installation. You should not use the `/qn` option when third-party packages are not installed.

3. Log in to the Unified Manager web user interface by using the following URL:

   `https://IP address`

Change the JBoss password on Windows in Active IQ Unified Manager

You can reset the instance-specific JBoss password that is set during installation. You can reset the password optionally, in case your site requires this security capability to override the Unified Manager installation setting. This operation also changes the password JBoss uses to access MySQL.

Before you begin

- You should have Windows admin privileges for the system on which Unified Manager is installed.
- You should have the password for the MySQL root user.
- You should be able to access the NetApp-provided `password.bat` script in the directory

  `C:\Program Files\NetApp\essentials\bin`.

Steps

1. Log in as the admin user on the Unified Manager host machine.

2. Use the Windows Services console to stop the following Unified Manager services:
   - NetApp Active IQ Acquisition Service (Ocie-au)
   - NetApp Active IQ Management Server Service (Oncommandsvc)
3. Launch the `password.bat` script to begin the password change process:

   ```
   C:\Program Files\NetApp\essentials\bin> password.bat resetJBossPassword
   ```

4. When prompted, enter the MySQL root user password.
5. When prompted, enter the new JBoss user password, and then enter it again for confirmation.

   Note that the password should be between 8 and 16 characters, and must contain at least a digit, an upper case and lower case characters, and at least one of these special characters:

   ```
   !@%^*-_=[]:<>.?/~+
   ```

6. When the script completes, start the Unified Manager services by using the Windows Services console:
   - NetApp Active IQ Management Server Service (Oncommandsvc)
   - NetApp Active IQ Acquisition Service (Ocie-au)
7. After all of the services are started, you can log in to the Unified Manager UI.

Use the supported upgrade path for versions of Active IQ Unified Manager

Active IQ Unified Manager supports a specific upgrade path for each version.

Not all versions of Unified Manager can perform an in-place upgrade to later versions. The Unified Manager upgrades are limited to an N-2 model. This means that you can only perform an upgrade to the next two releases on all platforms. For example, you can perform an upgrade to Unified Manager 9.18 from Unified Manager 9.14 and 9.16 only.

If you are running a version that is earlier than the supported versions, your Unified Manager instance will need to be upgraded to one of the supported versions first, and then upgraded to the current version.

For example, if your installed version is Unified Manager 9.9 and you want to upgrade to Unified Manager 9.14, you follow a sequence of upgrades.

Sample upgrade path:

1. Upgrade 9.11 → 9.13
2. Upgrade 9.13 → 9.14
3. Upgrade 9.13 → 9.16
4. Upgrade 9.14 or 9.16 → 9.18

For more information on the upgrade path matrix, see this knowledge base (KB) article.

Upgrade Active IQ Unified Manager on Windows

You can upgrade Unified Manager 9.14 or 9.16 to 9.18 by downloading and running the installation file on the Windows platform.

Before you begin

- The system on which you are upgrading Unified Manager should meet the system and software requirements.

  See Hardware system requirements.

  See Windows software and installation requirements.

  > Starting with Unified Manager 9.5, OpenJDK is provided in the installation package and installed automatically. Oracle Java is not supported starting with Unified Manager 9.5.

  > Ensure that Microsoft .NET 4.5.2 or greater is installed on your system before starting the upgrade.

- MySQL Community Edition is automatically upgraded during Unified Manager upgrade. If the installed version of MySQL on your system is earlier than 8.4.10, the Unified Manager upgrade process automatically upgrades MySQL to 8.4.10. You must not run a standalone upgrade of an earlier version of MySQL to 8.4.10.
- All clusters are supported, except for ASA r2 Systems.
- You should have Windows administrator privileges. Ensure that your user name does not begin with an exclamation mark “!”. Installation of Unified Manager might fail if the user name of user running the installation begins with “!”.
- You should have valid credentials to log in to the NetApp Support Site.
- To avoid data loss, you should have created a backup of the Unified Manager machine in case there is an issue during the upgrade.
- You should have adequate disk space available to perform the upgrade.

  The available space on the installation drive should be 2.5 GB larger than the size of the data directory. The upgrade stops and displays an error message indicating the amount of space to be added if there is not enough free space.

- During upgrade, you may be prompted to confirm whether you want to keep the previous, default settings for retaining performance data for 13 months or to change it to 6 months. On confirming, the historical performance data after 6 months is purged.
- Before upgrading, you should close any open files or folder in *<InstallDir>\JDK* and *MySQL Data Directory*.
- If you have an active anti-virus software installed on your Windows system, Unified Manager upgrade might fail. You should disable all virus scan software on your system before upgrading Unified Manager.

During the upgrade process, Unified Manager is unavailable. You should complete any running operations before upgrading Unified Manager.

If Unified Manager is paired with an instance of OnCommand Workflow Automation, and there are new versions of software available for both products, you must disconnect the two products and then set up a new Workflow Automation connection after performing the upgrades. If you are performing an upgrade to only one of the products, then you should log into Workflow Automation after the upgrade and verify that it is still acquiring data from Unified Manager.

Steps

1. Log in to the NetApp Support Site, and navigate to the Download page for Unified Manager:

   NetApp Support Site.

2. Select the required version of Unified Manager and accept the end-users license agreement (EULA).
3. Download the Unified Manager Windows installation file or zip file to a target directory on the Windows system. Extract the zip file, if needed, and navigate to the directory where the installation file is located. Right-click and run the Unified Manager installer executable (.exe) file as an administrator.

   Unified Manager prompts you with the following message:

   ```
   This setup will perform an upgrade of Unified Manager. Do you want to
   continue?
   ```

4. Click **Yes**, and then click **Next**.
5. Enter the MySQL84 root password that was set during installation, and then click **Next**.
6. Launch the web UI on a new window on a supported web browser and log in to use the upgraded version of Unified Manager.
7. If you have an active anti-virus software installed on your Windows system, make sure to manually exclude the following paths from anti-virus scan after your upgrade is complete:
   - Unified Manager data directory
   - Unified Manager installation directory
   - MySQL data directory

> To perform a silent upgrade of Unified Manager, run the following command:
> `ActiveIQUnifiedManager-<version\>.exe /s /v"MYSQL_PASSWORD=<password> /qn /l*v <system_drive>:\install.log"`

Upgrade third-party products

Upgrade third-party products in Active IQ Unified Manager on Windows

You can upgrade third-party products, such as JRE, on Unified Manager when installed on Windows systems.

The companies that develop these third-party products report security vulnerabilities on a regular basis. You can upgrade to newer versions of this software at your own schedule.

Upgrade OpenJDK on Windows for Active IQ Unified Manager

You can upgrade to a newer version of OpenJDK on the Windows server on which Unified Manager is installed to obtain fixes for security vulnerabilities.

Before you begin

You must have Windows admin privileges for the system on which Unified Manager is installed.

You can update OpenJDK releases within release families. For example, you can upgrade from OpenJDK 11.0.16 to OpenJDK 11.0.18, but you cannot update directly from OpenJDK 11 to OpenJDK 12.

Steps

1. Log in as the admin user on the Unified Manager host machine.

2. Download the appropriate version of OpenJDK (64-bit) from the OpenJDK site to the target system.

   For example, download the `jdk-17.0.15_windows-x64_bin.zip` file from the Oracle Java SE 17 Archive Downloads page.

   > An Oracle account is required to download the file. If you do not have an Oracle account, go to the Oracle account sign-in page to create one.

3. Use the Windows Services console to stop the following Unified Manager services:
   - NetApp Active IQ Acquisition Service (Ocie-au)
   - NetApp Active IQ Management Server Service (Oncommandsvc)
4. Expand the `zip` file.
5. Copy the directories and files from the resulting `jdk` directory (for example, `jdk-11.0.18` to the location where Java is installed. Example: `C:\Program Files\NetApp\JDK\`
6. Start the Unified Manager services by using the Windows Services console:
   - NetApp Active IQ Management Server Service (Oncommandsvc)
   - NetApp Active IQ Acquisition Service (Ocie-au)

Uninstall third-party products in Active IQ Unified Manager

After upgrading Unified Manager, you can uninstall third party products, such as Python or Microsoft Visual C++ 2013. This is optional.

Uninstall Python

Steps

1. From the **Control Panel** select **Programs** > **Programs and Features**.
2. Select Python from the list of programs.
3. Select **Uninstall this program**. Uninstall any version of Python except for 3.13.3. Ensure that you retain Python version 3.13.3.

Uninstall Microsoft Visual C++ 2013

Steps

1. From the **Control Panel** select **Programs** > **Programs and Features**.
2. Select **Microsoft Visual C++ 2013** from the list of programs.
3. Select **Uninstall this program**.

Restart Active IQ Unified Manager on Windows

You might have to restart Unified Manager after making configuration changes.

Before you begin

You must have Windows administrator privileges.

Steps

1. Log in to Windows using the default local administrator account.
2. Stop the Unified Manager services:

From the…	Stop the services in following order…
Command line	a. `sc stop ocie-au` b. `sc stop Oncommandsvc`
Microsoft Service Manager	a. NetApp Active IQ Acquisition Service (Ocie-au) b. NetApp Active IQ Management Server Service (Oncommandsvc)

3. Start the Unified Manager services:

From the…	Start the services in following order…
Command line	a. `sc start Oncommandsvc` b. `sc start ocie-au`
Microsoft Service Manager	a. NetApp Active IQ Management Server Service (Oncommandsvc) b. NetApp Active IQ Acquisition Service (Ocie-au)

Uninstall Active IQ Unified Manager

You can uninstall Unified Manager by using the Programs and Features wizard, or by performing an unattended uninstallation from the command-line interface.

Before you begin

- You must have Windows administrator privileges.
- All clusters (data sources) must be removed from the Unified Manager server before uninstalling the software.
- You should manually delete the firewall rules that are created to allow or block MySQL port 3306. The firewall rules are not deleted automatically.

Steps

1. Uninstall Unified Manager by choosing one of the following options:
   - If you uninstall Unified Manager from **Programs and Features** wizard, perform these steps:
     a. Navigate to **Control Panel** > **Program and Features**.
     b. Select Active IQ Unified Manager, and click **Uninstall**.
   - If you uninstall Unified Manager from Command line, then perform these steps:
     a. Log in to the Windows command line using administrator privileges.
     b. Navigate to the Active IQ Unified Manager directory, and run the following command:

     ```
     msiexec /x {A78760DB-7EC0-4305-97DB-E4A89CDFF4E1} /qn /l*v
     %systemdrive%\UmUnInstall.log
     ```

     If User Account Control (UAC) is enabled on the server, and you are logged in as a domain user,

you must use the command-line uninstallation method.

Unified Manager is uninstalled from your system.

2. Uninstall the following third-party packages and data that are not removed during the Unified Manager uninstallation:
   - Third-party packages: JRE, MySQL, Microsoft Visual C++ 2015 Redistributable, Python, and 7zip
   - MySQL application data generated by Unified Manager
   - Application logs and contents of application data directory

Perform configuration and administrative tasks

Configure Active IQ Unified Manager

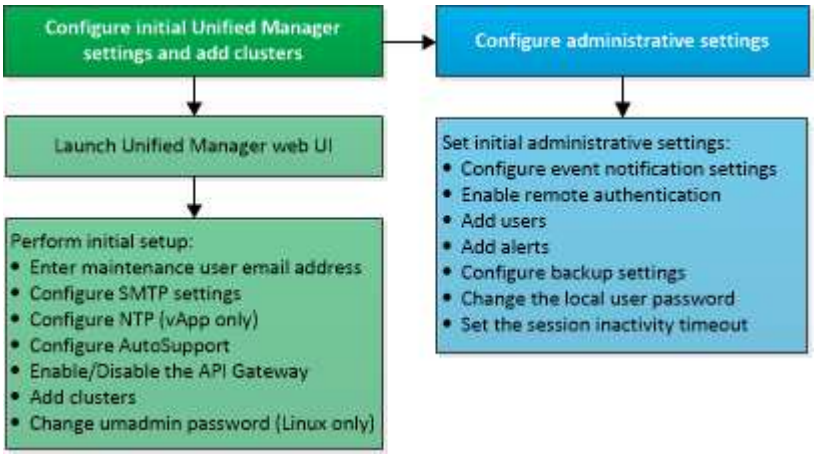
Learn about configuring Active IQ Unified Manager

After installing Active IQ Unified Manager (formerly OnCommand Unified Manager) you must complete the initial setup (also called the first experience wizard) to access the web UI. Then you can perform additional configuration tasks, such as adding clusters, configuring remote authentication, adding users, and adding alerts.

Some of the procedures described in this manual are required to complete the initial setup of your Unified Manager instance. Other procedures are recommended configuration settings that are helpful to set up on your new instance, or that are good to know about before you start the regular monitoring of your ONTAP systems.

Learn about the configuration sequence in Active IQ Unified Manager

The configuration workflow describes the tasks that you must perform before you can use Unified Manager.



Access the web UI in Active IQ Unified Manager

After you have installed Unified Manager, you can access the web UI to set up Unified Manager so that you can begin monitoring your ONTAP systems.

Before you begin

- If this is the first time you are accessing the web UI, you must log in as the maintenance user (or umadmin user for Linux installations).
- If you plan to allow users to access Unified Manager using the short name instead of using the fully qualified domain name (FQDN) or IP address, then your network configuration has to resolve this short name to a valid FQDN.
- If the server uses a self-signed digital certificate, the browser might display a warning indicating that the certificate is not trusted. You can either acknowledge the risk to continue the access or install a Certificate Authority (CA) signed digital certificate for server authentication.

Steps

1. Start the Unified Manager web UI from your browser by using the URL displayed at the end of the installation. The URL is the IP address or fully qualified domain name (FQDN) of the Unified Manager server.

   The link is in the following format: `https://URL`.

2. Log in to the Unified Manager web UI using your maintenance user credentials.

> If you make three consecutive unsuccessful attempts to log into the web UI within an hour, you will be locked out of the system, and will need to contact your system administrator. This is applicable for only local users.

Perform the initial setup of the Active IQ Unified Manager web UI

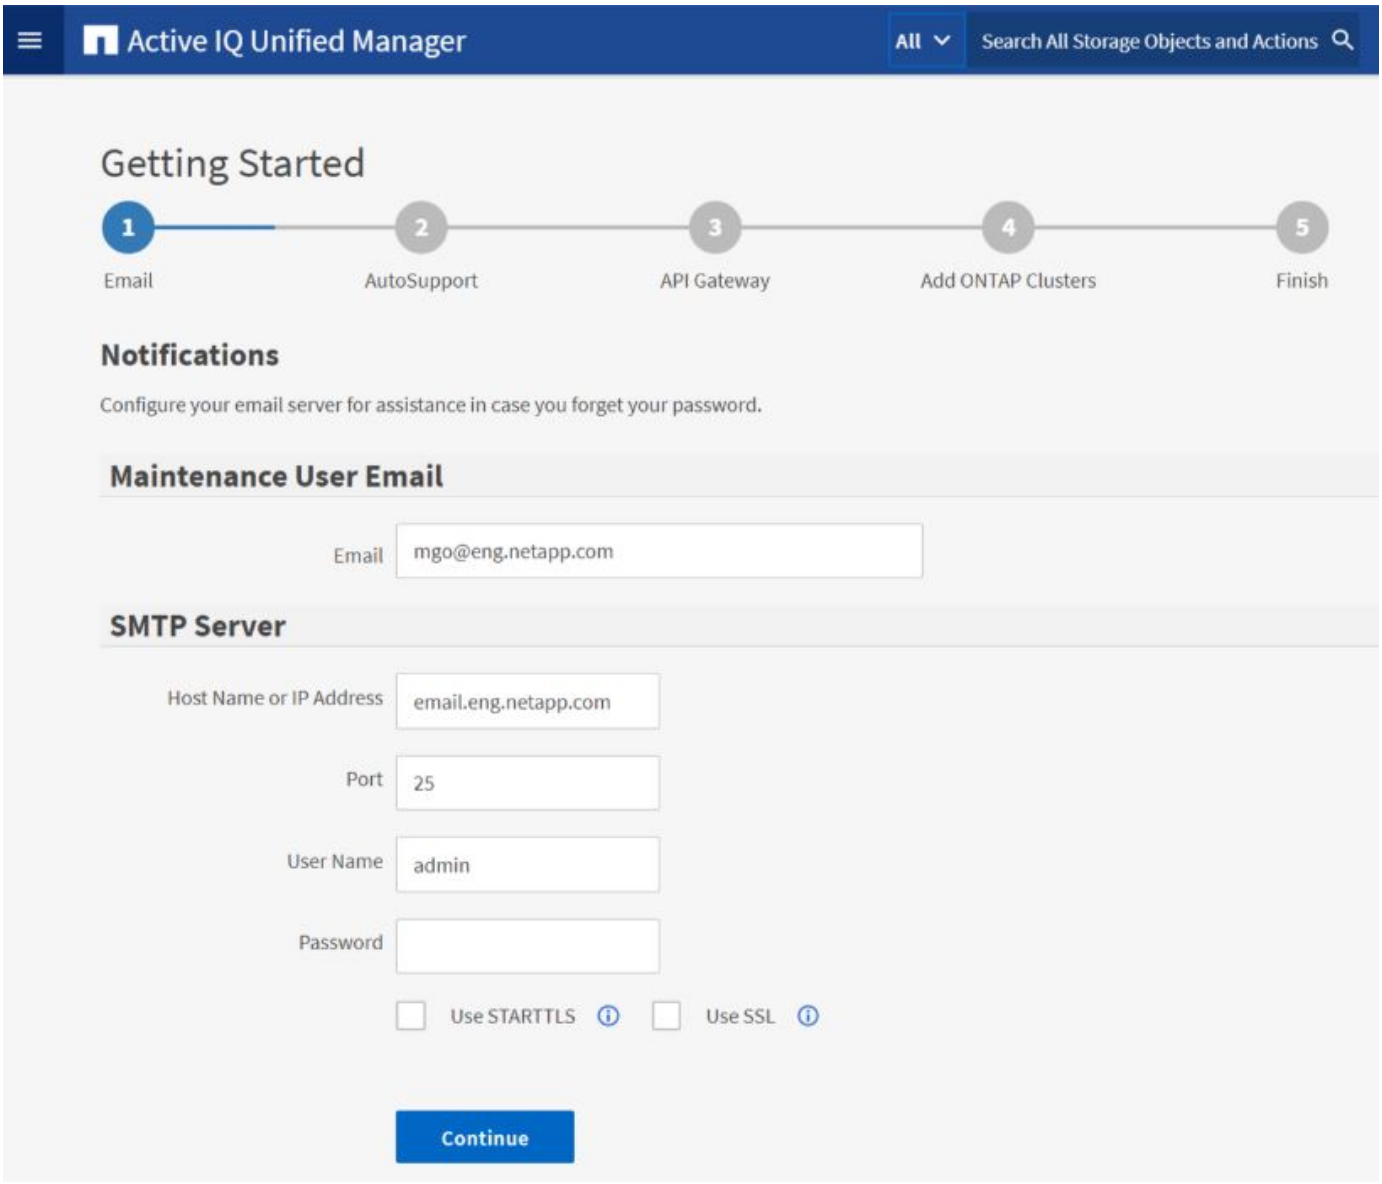
To use Unified Manager, you must first configure the initial setup options, including the NTP server, the maintenance user email address, the SMTP server host, and adding ONTAP clusters.

Before you begin

You must have performed the following operations:

- Launched the Unified Manager web UI using the URL provided after installation
- Logged in using the maintenance user name and password (umadmin user for Linux installations) created during installation

The Active IQ Unified ManagerGetting Started page appears only when you first access the web UI. The page below is from an installation on VMware.



If you want to change any of these options later, you can select your choice from the General options in the Unified Manager left-navigation pane. Note that the NTP setting is only for VMware installations, and it can be changed later using the Unified Manager maintenance console.

Steps

1. In the Active IQ Unified Manager Initial Setup page, enter the maintenance user email address, the SMTP server host name and any additional SMTP options, and the NTP server (VMware installations only). Then click **Continue**.

   If you have selected the **Use STARTTLS** or **Use SSL** option, a certificate page displays after you click the **Continue** button. Verify the certificate details and accept the certificate to continue with the initial setup settings of the web UI.

2. In the AutoSupport page click **Agree and Continue** to enable the sending of AutoSupport messages from Unified Manager to NetAppActive IQ.

   If you need to designate a proxy to provide internet access in order to send AutoSupport content, or if you want to disable AutoSupport, use the **General** > **AutoSupport** option from the web UI.

3. On Red Hat systems, change the umadmin user password from the default “admin” string to a personalized string.
4. In the Set up API Gateway page, select whether you want to use the API Gateway feature that allows Unified Manager to manage the ONTAP clusters you are planning to monitor using ONTAP REST APIs. Then click **Continue**.

   You can enable or disable this setting later in the web UI from **General** > **Feature Settings** > **API Gateway**. For more information about the APIs, see Getting started with Active IQ Unified Manager REST APIs.

5. Add the clusters that you want Unified Manager to manage, and then click **Next**. For each cluster you plan to manage, you must have the host name or cluster management IP address (IPv4 or IPv6) along with the user name and password credentials - the user must have the “admin” role.

   This step is optional. You can add clusters later in the web UI from **Storage Management** > **Cluster Setup**.

6. In the Summary page, verify that all the settings are correct and click **Finish**.

The Getting Started page closes and the Unified Manager Dashboard page is displayed.

Add clusters to Active IQ Unified Manager

You can add a cluster to Active IQ Unified Manager so that you can monitor the cluster. This includes the ability to obtain cluster information such as the health, capacity, performance, and configuration of the cluster so that you can find and resolve any issues that might occur.

Before you begin

- You must have the Application Administrator or Storage Administrator role.
- You must have the following information:
  - Unified Manager supports on-premises ONTAP clusters, ONTAP Select, Cloud Volumes ONTAP.
  - Host name or cluster-management IP address

    The host name is the FQDN or short name that Unified Manager uses to connect to the cluster. The host name must resolve to the cluster-management IP address.

    The cluster-management IP address must be the cluster-management LIF of the administrative storage virtual machine (SVM). If you use a node-management LIF, the operation fails.

  - The cluster must be running ONTAP version 9.1 software or greater.
  - ONTAP administrator user name and password

    This account must have the *admin* role with Application access set to *ontapi*, *amqp*, *console*, and *http*.

  - The port number to connect to the cluster using the HTTPS protocol (typically port 443)
  - You have the required certificates:

    SSL (HTTPS) certificate: This certificate is owned by Unified Manager. A default self-signed SSL (HTTPS) certificate is generated with a fresh installation of Unified Manager. NetApp recommends that you upgrade it to CA-signed certificate for better security. If the server certificate expires, you should

regenerate it and restart Unified Manager for the services to incorporate the new certificate. For more information about regenerating SSL certificate, see Generating an HTTPS security certificate.

EMS certificate: This certificate is owned by Unified Manager. It is used during authentication for EMS notifications received from ONTAP.

Certificates for Mutual TLS communication: Used during Mutual TLS communication between Unified Manager and ONTAP. The certificate-based authentication is enabled for a cluster, based on the ONTAP version. If the cluster running the ONTAP version is lower than the 9.5, certificate-based authentication is not enabled.

Certificate-based authentication is not enabled automatically for a cluster, if you are updating an older version of Unified Manager. However, you can enable it by modifying and saving the cluster details. If the certificate expires, you should regenerate it to incorporate the new certificate. For more information about viewing and regenerating the certificate, see Editing clusters.



- You can add a cluster from the web UI, and certificate-based authentication is automatically enabled.
- You can add a cluster through Unified Manager CLI, the certificate-based authentication is not enabled by default. If you add a cluster using Unified Manager CLI, it is required to edit the cluster using Unified Manager UI. You can see Supported Unified Manager CLI commands to add a cluster using Unified Manager CLI.
- If certificate-based authentication is enabled for a cluster, and you take the backup of Unified Manager from a server and restore to another Unified Manager server where hostname or IP address is changed, then monitoring of the cluster can fail. To avoid the failure, edit and save the cluster details. For more information about editing cluster details, see Editing clusters.

Cluster certificates: This certificate is owned by ONTAP. You cannot add a cluster to Unified Manager with an expired certificate and if the certificate has already expired, you should regenerate it before adding the cluster. For information about certificate generation, see the knowledge base (KB) article How to renew an ONTAP self-signed certificate in System Manager user interface.

- You must have adequate space on the Unified Manager server. You are prevented from adding a cluster to the server when greater than 90% of space in the database directory is already consumed.

For a MetroCluster configuration, you must add both the local and remote clusters, and the clusters must be configured correctly.

Steps

1. In the left navigation pane, click **Storage Management** > **Cluster Setup**.
2. On the Cluster Setup page, click **Add**.
3. In the Add Cluster dialog box, specify the required values, such as the host name or IP address of the cluster, user name, password, and port number.

   You can change the cluster-management IP address from IPv6 to IPv4 or from IPv4 to IPv6. The new IP address is reflected in the cluster grid and the cluster configuration page after the next monitoring cycle is complete.

4. Click **Submit**.
5. In the Authorize Host dialog box, click **View Certificate** to view the certificate information about the cluster.

6. Click **Yes**.

   After saving the cluster details, you can see the certificate for Mutual TLS communication for a cluster.

   If the certificate-based authentication is not enabled, Unified Manager checks the certificate only when the cluster is added initially. Unified Manager does not check the certificate for each API call to ONTAP.

After all the objects for a new cluster are discovered, Unified Manager starts to gather historical performance data for the previous 15 days. These statistics are collected using the data continuity collection functionality. This feature provides you with over two weeks of performance information for a cluster immediately after it is added. After the data continuity collection cycle is completed, real-time cluster performance data is collected, by default, every five minutes.

Because the collection of 15 days of performance data is CPU intensive, it is suggested that you stagger the addition of new clusters so that data continuity collection polls are not running on too many clusters at the same time. Additionally, if you restart Unified Manager during the data continuity collection period, the collection will be halted and you will see gaps in the performance charts for the missing timeframe.

If you receive an error message that you cannot add the cluster, check to see if the clocks on the two systems are not synchronized and the Unified Manager HTTPS certificate start date is later than the date on the cluster. You must ensure that the clocks are synchronized using NTP or a similar service.

Related information

Installing a CA signed and returned HTTPS certificate

Configure notifications

Configure Active IQ Unified Manager to send alert notifications

You can configure Unified Manager to send notifications that alert you about events in your environment. Before notifications can be sent, you must configure several other Unified Manager options.

Before you begin

You must have the Application Administrator role.

After deploying Unified Manager and completing the initial configuration, you should consider configuring your environment to trigger alerts and generate notification emails or SNMP traps based on the receipt of events.

Steps

1. Configure event notification settings.

   If you want alert notifications sent when certain events occur in your environment, you must configure an SMTP server and supply an email address from which the alert notification will be sent. If you want to use SNMP traps, you can select that option and provide the necessary information.

2. Enable remote authentication.

   If you want remote LDAP or Active Directory users to access the Unified Manager instance and receive

alert notifications, then you must enable remote authentication.

3. Add authentication servers.

   You can add authentication servers so that remote users within the authentication server can access Unified Manager.

4. Add users.

   You can add several different types of local or remote users and assign specific roles. When you create an alert, you assign a user to receive the alert notifications.

5. Add alerts.

   After you have added the email address for sending notifications, added users to receive the notifications, configured your network settings, and configured SMTP and SNMP options needed for your environment, then you can assign alerts.

Configure event notification settings in Active IQ Unified Manager

You can configure Unified Manager to send alert notifications when an event is generated or when an event is assigned to a user. You can configure the SMTP server that is used to send the alert, and you can set various notification mechanisms—for example, alert notifications can be sent as emails or SNMP traps.

Before you begin

You must have the following information:

- Email address from which the alert notification is sent

  The email address appears in the “From” field in sent alert notifications. If the email cannot be delivered for any reason, this email address is also used as the recipient for undeliverable mail.

- SMTP server host name, and the user name and password to access the server
- Host name or IP address for the trap destination host that will receive the SNMP trap, along with the SNMP version, outbound trap port, community, and other required SNMP configuration values

  To specify multiple trap destinations, separate each host with a comma. In this case, all other SNMP settings, such as version and outbound trap port, must be the same for all hosts in the list.

You must have the Application Administrator or Storage Administrator role.

Steps

1. In the left navigation pane, click **General** > **Notifications**.
2. In the Notifications page, configure the appropriate settings.

   Notes:

   - If the From Address is pre-filled with the address "ActiveIQUnifiedManager@localhost.com", you should change it to a real, working email address to make sure that all email notifications are delivered successfully.
   - If the host name of the SMTP server cannot be resolved, you can specify the IP address (IPv4 or IPv6) of the SMTP server instead of the host name.

3. Click **Save**.
4. If you have selected the **Use STARTTLS** or **Use SSL** option, a certificate page displays after you click the **Save** button. Verify the certificate details and accept the certificate to save the notification settings.

   You can click the **View Certificate Details** button to view the certificate details. If the existing certificate is expired, uncheck the **Use STARTTLS** or **Use SSL** box, save the notification settings, and again check the **Use STARTTLS** or **Use SSL** box to view a new certificate.

Enable remote authentication in Active IQ Unified Manager

You can enable remote authentication so that the Unified Manager server can communicate with your authentication servers. The users of the authentication server can access the Unified Manager graphical interface to manage storage objects and data.

Before you begin

You must have the Application Administrator role.

> The Unified Manager server must be connected directly with the authentication server. You must disable any local LDAP clients such as SSSD (System Security Services Daemon) or NSLCD (Name Service LDAP Caching Daemon).

You can enable remote authentication using either Open LDAP or Active Directory. If remote authentication is disabled, remote users cannot access Unified Manager.

Remote authentication is supported over LDAP and LDAPS (Secure LDAP). Unified Manager uses 389 as the default port for non-secure communication, and 636 as the default port for secure communication.

> The certificate that is used to authenticate users must conform to the X.509 format.

Steps

1. In the left navigation pane, click **General** > **Remote Authentication**.
2. Check the box for **Enable remote authentication…**.
3. In the Authentication Service field, select the type of service and configure the authentication service.

For Authentication type…	Enter the following information…
Active Directory	• Authentication server administrator name in one of following formats: ◦ `domainname\username` ◦ `username@domainname` ◦ `Bind Distinguished Name` (using the appropriate LDAP notation) • Administrator password • Base distinguished name (using the appropriate LDAP notation)

For Authentication type…	Enter the following information…
Open LDAP	• Bind distinguished name (in the appropriate LDAP notation) • Bind password • Base distinguished name

If the authentication of an Active Directory user takes a long time or times out, the authentication server is probably taking a long time to respond. Disabling support for nested groups in Unified Manager might reduce the authentication time.

If you select the Use Secure Connection option for the authentication server, then Unified Manager communicates with the authentication server using the Secure Sockets Layer (SSL) protocol.

4. **Optional:** Add authentication servers, and test the authentication.
5. Click **Save**.

Disable nested groups from remote authentication in Active IQ Unified Manager

If you have remote authentication enabled, you can disable nested group authentication so that only individual users, and not group members, can remotely authenticate to Unified Manager. You can disable nested groups when you want to improve Active Directory authentication response time.

Before you begin

- You must have the Application Administrator role.
- Disabling nested groups is only applicable when using Active Directory.

Disabling support for nested groups in Unified Manager might reduce the authentication time. If nested group support is disabled, and if a remote group is added to Unified Manager, individual users must be members of the remote group to authenticate to Unified Manager.

Steps

1. In the left navigation pane, click **General** > **Remote Authentication**.
2. Check the box for **Disable Nested Group Lookup**.
3. Click **Save**.

Set up authentication services in Active IQ Unified Manager

Authentication services enable the authentication of remote users or remote groups in an authentication server before providing them access to Unified Manager. You can authenticate users by using predefined authentication services (such as Active Directory or OpenLDAP), or by configuring your own authentication mechanism.

Before you begin

- You must have enabled remote authentication.
- You must have the Application Administrator role.

Steps

1. In the left navigation pane, click **General** > **Remote Authentication**.
2. Select one of the following authentication services:

If you select…	Then do this…
Active Directory	a. Enter the administrator name and password. b. Specify the base distinguished name of the authentication server. For example, if the domain name of the authentication server is ou@domain.com, then the base distinguished name is **cn=ou,dc=domain,dc=com**.
OpenLDAP	a. Enter the bind distinguished name and bind password. b. Specify the base distinguished name of the authentication server. For example, if the domain name of the authentication server is ou@domain.com, then the base distinguished name is **cn=ou,dc=domain,dc=com**.
Others	a. Enter the bind distinguished name and bind password. b. Specify the base distinguished name of the authentication server. For example, if the domain name of the authentication server is ou@domain.com, then the base distinguished name is **cn=ou,dc=domain,dc=com**. c. Specify the LDAP protocol version that is supported by the authentication server. d. Enter the user name, group membership, user group, and member attributes.

If you want to modify the authentication service, you must delete any existing authentication servers, and then add new authentication servers.

3. Click **Save**.

Add authentication servers to Active IQ Unified Manager

You can add authentication servers and enable remote authentication on the management server so that remote users within the authentication server can access Unified Manager.

Before you begin

- The following information must be available:

- Host name or IP address of the authentication server
- Port number of the authentication server

- You must have enabled remote authentication and configured your authentication service so that the management server can authenticate remote users or groups in the authentication server.
- You must have the Application Administrator role.

If the authentication server that you are adding is part of a high-availability (HA) pair (using the same database), you can also add the partner authentication server. This enables the management server to communicate with the partner when one of the authentication servers is unreachable.

Steps

1. In the left navigation pane, click **General** > **Remote Authentication**.
2. Enable or disable the **Use secure connection** option:

If you want to…	Then do this…
Enable it	a. Select the **Use Secure Connection** option. b. In the Authentication Servers area, click **Add**. c. In the Add Authentication Server dialog box, enter the authentication name or IP address (IPv4 or IPv6) of the server. d. In the Authorize Host dialog box, click View Certificate. e. In the View Certificate dialog box, verify the certificate information, and then click **Close**. f. In the Authorize Host dialog box, click **Yes**. When you enable the **Use Secure Connection authentication** option, Unified Manager communicates with the authentication server and displays the certificate. Unified Manager uses 636 as default port for secure communication and port number 389 for non-secure communication.
Disable it	a. Clear the **Use Secure Connection** option. b. In the Authentication Servers area, click **Add**. c. In the Add Authentication Server dialog box, specify either the host name or IP address (IPv4 or IPv6) of the server, and the port details. d. Click **Add**.

   The authentication server that you added is displayed in the Servers area.

3. Perform a test authentication to confirm that you can authenticate users in the authentication server that you added.

Test the configuration of authentication servers in Active IQ Unified Manager

You can validate the configuration of your authentication servers to ensure that the management server is able to communicate with them. You can validate the configuration by searching for a remote user or remote group from your authentication servers, and authenticating them using the configured settings.

Before you begin

- You must have enabled remote authentication, and configured your authentication service so that the Unified Manager server can authenticate the remote user or remote group.
- You must have added your authentication servers so that the management server can search for the remote user or remote group from these servers and authenticate them.
- You must have the Application Administrator role.

If the authentication service is set to Active Directory, and if you are validating the authentication of remote users who belong to the primary group of the authentication server, information about the primary group is not displayed in the authentication results.

Steps

1. In the left navigation pane, click **General** > **Remote Authentication**.
2. Click **Test Authentication**.
3. In the Test User dialog box, specify the user name and password of the remote user or the user name of the remote group, and then click **Test**.

   If you are authenticating a remote group, you must not enter the password.

Add alerts in Active IQ Unified Manager

You can configure alerts to notify you when a particular event is generated. You can configure alerts for a single resource, for a group of resources, or for events of a particular severity type. You can specify the frequency with which you want to be notified and associate a script to the alert.

Before you begin

- You must have configured notification settings such as the user email address, SMTP server, and SNMP trap host to enable the Active IQ Unified Manager server to use these settings to send notifications to users when an event is generated.
- You must know the resources and events for which you want to trigger the alert, and the user names or email addresses of the users that you want to notify.
- If you want to have a script execute based on the event, you must have added the script to Unified Manager by using the Scripts page.
- You must have the Application Administrator or Storage Administrator role.

You can create an alert directly from the Event details page after receiving an event in addition to creating an alert from the Alert Setup page, as described here.

Steps

1. In the left navigation pane, click **Storage Management** > **Alert Setup**.

2. In the Alert Setup page, click **Add**.
3. In the Add Alert dialog box, click **Name**, and enter a name and description for the alert.
4. Click **Resources**, and select the resources to be included in or excluded from the alert.

   You can set a filter by specifying a text string in the **Name contains** field to select a group of resources. Based on the text string that you specify, the list of available resources displays only those resources that match the filter rule. The text string that you specify is case-sensitive.

   If a resource conforms to both the include and exclude rules that you have specified, the exclude rule takes precedence over the include rule, and the alert is not generated for events related to the excluded resource.

5. Click **Events**, and select the events based on the event name or event severity type for which you want to trigger an alert.

   To select more than one event, press the Ctrl key while you make your selections.

6. Click **Actions**, and select the users that you want to notify, choose the notification frequency, choose whether an SNMP trap will be sent to the trap receiver, and assign a script to be executed when an alert is generated.

   If you modify the email address that is specified for the user and reopen the alert for editing, the Name field appears blank because the modified email address is no longer mapped to the user that was previously selected. Also, if you modified the email address of the selected user from the Users page, the modified email address is not updated for the selected user.

   You can also choose to notify users through SNMP traps.

7. Click **Save**.

Example of adding an alert

This example shows how to create an alert that meets the following requirements:

- Alert name: HealthTest
- Resources: includes all volumes whose name contains “abc” and excludes all volumes whose name contains “xyz”
- Events: includes all critical health events
- Actions: includes "sample@domain.com", a “Test” script, and the user has to be notified every 15 minutes

Perform the following steps in the Add Alert dialog box:

Steps

1. Click **Name**, and enter **HealthTest** in the **Alert Name** field.
2. Click **Resources**, and in the Include tab, select **Volumes** from the drop-down list.
   a. Enter **abc** in the **Name contains** field to display the volumes whose name contains “abc”.
   b. Select **<<All Volumes whose name contains 'abc'>>** from the Available Resources area, and move it to the Selected Resources area.
   c. Click **Exclude**, and enter **xyz** in the **Name contains** field, and then click **Add**.

3. Click **Events**, and select **Critical** from the Event Severity field.
4. Select **All Critical Events** from the Matching Events area, and move it to the Selected Events area.
5. Click **Actions**, and enter **sample@domain.com** in the Alert these users field.
6. Select **Remind every 15 minutes** to notify the user every 15 minutes.

   You can configure an alert to repeatedly send notifications to the recipients for a specified time. You should determine the time from which the event notification is active for the alert.

7. In the Select Script to Execute menu, select **Test** script.
8. Click **Save**.

Change the local user password in Active IQ Unified Manager

You can change your local user login password to prevent potential security risks.

Before you begin

You must be logged in as a local user.

The passwords for the maintenance user and for remote users cannot be changed using these steps. To change a remote user password, contact your password administrator. To change the maintenance user password, see Using the maintenance console.

Steps

1. Log in to Unified Manager.
2. From the top menu bar, click the user icon and then click **Change Password**.

   The **Change Password** option is not displayed if you are a remote user.

3. In the Change Password dialog box, enter the current password and the new password.
4. Click **Save**.

If Unified Manager is configured in a high-availability configuration, you must change the password on the second node of the setup. Both instances must have same password.

Set the session inactivity timeout in Active IQ Unified Manager

You can specify the inactivity timeout value for Unified Manager so that the session is terminated automatically after a certain period of inactivity. By default the timeout is set to 4,320 minutes (72 hours).

Before you begin

You must have the Application Administrator role.

This setting affects all logged in user sessions.



This option is not available if you have enabled Security Assertion Markup Language (SAML) authentication.

Steps

1. In the left navigation pane, click **General** > **Feature Settings**.
2. In the **Feature Settings** page, specify the inactivity timeout by choosing one of the following options:

If you want to…	Then do this…
Have no timeout set so that the session is never closed automatically	In the **Inactivity Timeout** panel, move the slider button to the left (off) and click **Apply**.
Set a specific number of minutes as the time out value	In the **Inactivity Timeout** panel, move the slider button to the right (on), specify the inactivity timeout value in minutes, and click **Apply**.

Set the session timeout using the CLI in Active IQ Unified Manager

You can set a maximum session timeout value for Unified Manager using the CLI so that the session is terminated automatically after a certain period of time.
By default, your session timeout is set to the maximum value, which is 4,320 minutes (72 hours). This means that your session ends automatically after 72 hours, even if you are logged in and actively using Unified Manager.

About this task

You must have the Application Administrator role.

The session timeout setting affects all logged in user sessions.

Steps

1. Log in to the Unified Manager CLI by entering the `um cli login` command. Use a valid username and password for authentication.
2. Enter the `um option set max.session.timeout.value=<in mins>` command to modify the session timeout value.

Change the Unified Manager host name

Change the Active IQ Unified Manager virtual appliance host name

The network host is assigned a name when the Unified Manager virtual appliance is first deployed. You can change the host name after deployment. If you change the host name, you must also regenerate the HTTPS certificate.

Before you begin

You must be logged in to Unified Manager as the maintenance user, or have the Application Administrator role assigned to you to perform these tasks.

You can use the host name (or the host IP address) to access the Unified Manager web UI. If you configured a static IP address for your network during deployment, then you would have designated a name for the network host. If you configured the network using DHCP, the host name should be taken from the DNS. If DHCP or DNS is not properly configured, the host name “Unified Manager” is automatically assigned and associated with the security certificate.

Regardless of how the host name was assigned, if you change the host name, and intend to use the new host name to access the Unified Manager web UI, you must generate a new security certificate.

If you access the web UI by using the server's IP address instead of the host name, you do not have to generate a new certificate if you change the host name. However, it is the best practice to update the certificate so that the host name in the certificate matches the actual host name.

If you change the host name in Unified Manager, you must manually update the host name in OnCommand Workflow Automation (WFA). The host name is not updated automatically in WFA.

The new certificate does not take effect until the Unified Manager virtual machine is restarted.

Steps

1. Generate an HTTPS security certificate

   If you want to use the new host name to access the Unified Manager web UI, you must regenerate the HTTPS certificate to associate it with the new host name.

2. Restart the Unified Manager virtual machine

   After you regenerate the HTTPS certificate, you must restart the Unified Manager virtual machine.

Generate a HTTPS security certificate in Active IQ Unified Manager

When Active IQ Unified Manager is installed for the first time, a default HTTPS certificate is installed. You might generate a new HTTPS security certificate that replaces the existing certificate.

Before you begin

You must have the Application Administrator role.

There can be multiple reasons to regenerate the certificate such as if you want to have better values for Distinguished Name (DN) or if you want a higher key size, or longer expiry period or if the current certificate has expired.

If you do not have access to the Unified Manager web UI, you can regenerate the HTTPS certificate with the same values using the maintenance console. While regenerating certificates, you can define the key size and the validity duration of the key. If you use the `Reset Server Certificate` option from the maintenance console, then a new HTTPS certificate is created which is valid for 397 days. This certificate will have an RSA key of size 2048 bits.

Steps

1. In the left navigation pane, click **General** > **HTTPS Certificate**.
2. Click **Regenerate HTTPS Certificate**.

   The Regenerate HTTPS Certificate dialog box is displayed.

3. Select one of the following options depending on how you want to generate the certificate:

If you want to…	Do this…
Regenerate the certificate with the current values	Click the **Regenerate Using Current Certificate Attributes** option.

If you want to…	Do this…
Generate the certificate using different values	Click the **Update the Current Certificate Attributes** option. The Common Name and Alternative Names fields will use the values from the existing certificate if you do not enter new values. The “Common Name” should be set to the FQDN of the host. The other fields do not require values, but you can enter values, for example, for the EMAIL, COMPANY, DEPARTMENT, City, State, and Country if you want those values to be populated in the certificate. You can also select from the available KEY SIZE (The key algorithm is “RSA”.) and VALIDITY PERIOD. • The permitted values for key size are `2048`, `3072` and `4096`. • The validity periods are minimum 1 day to maximum 36500 days. Even though a validity period of 36500 days is permitted, it is recommended you use a validity period of not more than 397 days or 13 months. Because if you select a validity period of more than 397 days and plan to export a CSR for this certificate and get it signed by a well known CA, the validity of the signed certificate returned to you by the CA will be reduced to 397 days. • You can select the “Exclude local identifying information (e.g. localhost)” checkbox if you want to remove the local identifying information from the Alternative Names field in the certificate. When this checkbox is selected, only what you enter in the field is used in the Alternative Names field. When left blank the resulting certificate will not have an Alternative Names field at all.

4. Click **Yes** to regenerate the certificate.
5. Restart the Unified Manager server so that the new certificate takes effect.
6. Verify the new certificate information by viewing the HTTPS certificate.

Restart the Active IQ Unified Manager virtual machine

You can restart the virtual machine from the maintenance console of Unified Manager. You must restart after generating a new security certificate or if there is a problem with the virtual machine.

Before you begin

The virtual appliance is powered on.

You are logged in to the maintenance console as the maintenance user.

You can also restart the virtual machine from vSphere by using the **Restart Guest** option. See the VMware documentation for more information.

Steps

1. Access the maintenance console.
2. Select **System Configuration** > **Reboot Virtual Machine**.

Change the Active IQ Unified Manager host name on Linux systems

At some point, you might want to change the host name of the Red Hat Enterprise Linux machine on which you have installed Unified Manager. For example, you might want to rename the host to more easily identify your Unified Manager servers by type, workgroup, or monitored cluster group when you list your Linux machines.

Before you begin

You must have root user access to the Linux system on which Unified Manager is installed.

You can use the host name (or the host IP address) to access the Unified Manager web UI. If you configured a static IP address for your network during deployment, then you would have designated a name for the network host. If you configured the network using DHCP, the host name should be taken from the DNS server.

Regardless of how the host name was assigned, if you change the host name and intend to use the new host name to access the Unified Manager web UI, you must generate a new security certificate.

If you access the web UI by using the server's IP address instead of the host name, you do not have to generate a new certificate if you change the host name. However, it is the best practice to update the certificate, so that the host name in the certificate matches the actual host name. The new certificate does not take effect until the Linux machine is restarted.

If you change the host name in Unified Manager, you must manually update the host name in OnCommand Workflow Automation (WFA). The host name is not updated automatically in WFA.

Steps

1. Log in as the root user to the Unified Manager system that you want to modify.
2. Stop the Unified Manager software and the associated MySQL software by entering the following command:

   ```
   systemctl stop ocieau ocie mysqld
   ```

3. Change the host name using the Linux `hostnamectl` command:

```
hostnamectl set-hostname new_FQDN

hostnamectl set-hostname nuhost.corp.widget.com
```

4. Regenerate the HTTPS certificate for the server:

```
/opt/netapp/essentials/bin/cert.sh create
```

5. Restart the network service:

```
systemctl restart NetworkManager.service
```

6. After the service is restarted, verify whether the new host name is able to ping itself:

```
ping new_hostname

ping nuhost
```

This command should return the same IP address that was set earlier for the original host name.

7. After you complete and verify your host name change, restart Unified Manager by entering the following command:

```
systemctl start mysqld ocie ocieau
```

Enable and disable policy-based storage management in Active IQ Unified Manager

Starting with Unified Manager 9.7, you can provision storage workloads (volumes and LUNs) on your ONTAP clusters, and manage those workloads based on assigned performance service levels. This functionality is similar to creating workloads in ONTAP System Manager and attaching QoS policies, but when applied using Unified Manager you can provision and manage workloads across all clusters that your Unified Manager instance is monitoring.

You must have the Application Administrator role.

This option is enabled by default, but you can disable it if you do not want to provision and manage workloads using Unified Manager.

When enabled, this option provides many new items in the user interface:

New Content	Location
A page to provision new workloads	Available from **Common Tasks** > **Provisioning**
A page to create performance service level policies	Available from **Settings** > **Policies** > **Performance Service Levels**
A page to create performance storage efficiency policies	Available from **Settings** > **Policies** > **Storage Efficiency**
Panels that describe your current Workload Performance and Workload IOPS	Available from the Dashboard

See the online help in the product for more information on these pages and on this functionality.

Steps

1. In the left navigation pane, click **General** > **Feature Settings**.
2. In the **Feature Settings** page, disable or enable policy-based storage management by choosing one of the following options:

If you want to…	Then do this…
Disable policy-based storage management	In the **Policy-based storage management** panel, move the slider button to the left.
Enable policy-based storage management	In the **Policy-based storage management** panel, move the slider button to the right.

Configure Active IQ Unified Manager backup

You can configure the backup capability on Unified Manager through a set of configuration steps to be performed on the host systems and the through maintenance console.

For information about the configuration steps, see Managing backup and restore operations.

Manage feature settings

Learn about feature settings in Active IQ Unified Manager

The Feature Settings page allows you to enable and disable specific features in Active IQ Unified Manager. This includes creating and managing storage objects based on policies, enabling the API Gateway and Login Banner, uploading scripts for managing alerts, timing out a web UI session based on inactivity time, and disabling receipt of Active IQ platform events.

The Feature Settings page is only available for users with Application Administrator role.

For information about Scripts Upload, see Enabling and disabling scripts upload.

Enable policy-based storage management in Active IQ Unified Manager

The **Policy-based storage management** option allows storage management based on service level objectives (SLOs). This option is enabled by default.

On activating this feature, you can provision storage workloads on the ONTAP clusters added to your Active IQ Unified Manager instance, and manage these workloads based on the assigned Performance Service Levels and Storage Efficiency Policies.

You can choose to activate or deactivate this feature from **General** > **Feature Settings** > **Policy-based storage management**. On activating this feature, the following pages are available for operation and monitoring:

- Provisioning (storage workload provisioning)
- **Policies** > **Performance Service Levels**
- **Policies** > **Storage Efficiency**
- Workloads Managed by Performance Service Level column on the Clusters Setup page
- Workload Performance panel on the **Dashboard**

You can use the screens to create Performance Service Levels and Storage Efficiency Policies, and provision storage workloads. You can also monitor the storage workloads that conform to the assigned Performance Service Levels, as well as the nonconforming ones. The Workload Performance and Workload IOPS panel also enables you to assess the total, available, and used capacity and performance (IOPS) of the clusters across your data center based on the storage workloads provisioned on them.

After activating this feature, you can run the Unified Manager REST APIs to perform some of these functions from **Menu Bar** > **Help button** > **API Documentation** > **storage-provider** category. Alternatively, you can enter the host name or IP address and the URL to access the REST API page in the format https://<hostname>/docs/api/

For more information about the APIs, see Getting started with Active IQ Unified Manager REST APIs.

Learn about the API Gateway feature in Active IQ Unified Manager

The API Gateway feature allows Active IQ Unified Manager to be a single control plane from which you can manage multiple ONTAP clusters, without logging in to them individually.

You can enable this feature from the configuration pages that appear when you first log in to Unified Manager. Alternatively, you can enable or disable this feature from **General** > **Feature Settings** > **API Gateway**.

Unified Manager REST APIs are different from the ONTAP REST APIs, and not all the functionalities of ONTAP REST APIs can be availed by using the Unified Manager REST APIs. However, if you have a specific business requirement of accessing the ONTAP APIs for managing specific features that are not exposed to Unified Manager, you can enable the API Gateway feature and execute the ONTAP APIs. The gateway acts as a proxy to tunnel the API requests by maintaining the header and body requests in the same format as in the ONTAP APIs. You can use your Unified Manager credentials and execute the specific APIs to access and manage the ONTAP clusters without passing individual cluster credentials. Unified Manager performs as a single point of management for running the APIs across the ONTAP clusters managed by your Unified Manager instance. The response returned by the APIs is the same as the response returned by the respective ONTAP REST APIs executed directly from ONTAP.

After enabling this feature, you can execute the Unified Manager REST APIs from **Menu Bar** > **Help button** > **API Documentation** > **gateway** category. Alternatively, you can enter the host name or IP address and the URL to access the REST API page in the format `https://<hostname>/docs/api/`

For more information about the APIs, see Getting started with Active IQ Unified Manager REST APIs.

Specify inactivity timeout in Active IQ Unified Manager

You can specify the inactivity timeout value for Active IQ Unified Manager. After an inactivity of the specified time, the application is automatically logged out. This option is enabled by default.

You can deactivate this feature or modify the time from **General** > **Feature Settings** > **Inactivity Timeout**. Once you activate this feature, you should specify the time limit of inactivity (in minutes) in the **LOGOUT AFTER** field, after which the system automatically logs out. The default value is 4320 minutes (72 hours).

This option is not available if you have enabled Security Assertion Markup Language (SAML) authentication.

Enable Active IQ portal events in Active IQ Unified Manager

You can specify whether you want to enable or disable Active IQ portal events. This setting allows the Active IQ portal to discover and display additional events about system configuration, cabling, and so forth. This option is enabled by default.

On enabling this feature, Active IQ Unified Manager displays events discovered by the Active IQ portal. These events are created by running a set of rules against AutoSupport messages generated from all monitored storage systems. These events are different from the other Unified Manager events, and they identify incidents or risks related to system configuration, cabling, best practice, and availability issues.

You can choose to activate or deactivate this feature from **General** > **Feature Settings** > **Active IQ Portal Events**. In sites with no external network access, you must upload the rules manually from **Storage Management** > **Event Setup** > **Upload Rules**.

This feature is enabled by default. Disabling this feature stops the Active IQ events from being discovered or displayed on Unified Manager. When disabled, enabling this feature allows Unified Manager to receive the Active IQ events on a cluster at a predefined time of 00:15 for that cluster timezone.

Enable and disable security settings for compliance in Active IQ Unified Manager

By using the **Customize** button on the **Security Dashboard** panel of the Features Settings page, you can enable or disable the security parameters for compliance monitoring on Unified Manager.

The settings that are enabled or disabled from this page govern the overall compliance status of the clusters and storage VMs on Unified Manager. Based on the selections, the corresponding columns are visible in the **Security: All Clusters** view of the Clusters inventory page and the **Security: All Storage VMs** view of the Storage VMs inventory page.

Only users with administrator role can edit these settings.

The security criteria for your ONTAP clusters, storage VMs, and volumes are evaluated against the recommendations defined in the Security Hardening Guide for NetApp ONTAP 9. The Security panel on the dashboard and the Security page display the default security compliance status of your clusters, storage VMs, and volumes. Security events are also generated and management actions enabled for the clusters and storage VMs that have security violations.

Customize security settings

To customize the settings for compliance monitoring as applicable to your ONTAP environment, follow these steps:

Steps

1. Click **General > Feature Settings > Security Dashboard > Customize**. The **Customize Security Dashboard Settings** pop-up appears.

> The security compliance parameters that you enable or disable can directly affect the default security views, reports, and scheduled reports on the Clusters and Storage VMs screens. If you had uploaded an excel report from these screens before modifying the security parameters, the downloaded excel reports might be faulty.

2. To enable or disable the custom settings for your ONTAP clusters, select the required general setting under **Cluster**. For information on the options for customizing cluster compliance, see Cluster compliance categories.
3. To enable or disable the custom settings for your storage VMs, select the required general setting under **Storage VM**. For information on the options for customizing storage VM compliance, see Storage VM compliance categories.

Customize AutoSupport and authentication settings

On the **AutoSupport Settings** section, you can specify whether HTTPS transport is to be used for sending AutoSupport messages from ONTAP.

From the **Authentication Settings** section, you can enable Unified Manager alerts to be raised for the default ONTAP administrator user.

Enable and disable script upload in Active IQ Unified Manager

The ability to upload scripts to Unified Manager and run them is enabled by default. If your organization does not want to allow this activity because of security reasons, you can disable this functionality.

Before you begin

You must have the Application Administrator role.

Steps

1. In the left navigation pane, click **General** > **Feature Settings**.
2. In the **Feature Settings** page, disable or enable scripting by choosing one of the following options:

If you want to…	Then do this…
Disable scripts	In the **Script Upload** panel, move the slider button to the left.
Enable scripts	In the **Script Upload** panel, move the slider button to the right.

Add a login banner in Active IQ Unified Manager

Adding a login banner enables your organization to display any information, such as, who is permitted access to the system and the terms and conditions of use during login and logout.

Any user, such as storage operators or administrators can view this login banner pop-up during login, logout, and session timeout.

Use the maintenance console

Learn about the maintenance console in Active IQ Unified Manager

The Unified Manager maintenance console enables you to maintain the settings on your Unified Manager system and to make any necessary changes to prevent issues from occurring.

Depending on the operating system on which you have installed Unified Manager, the maintenance console provides the following functions:

- Troubleshoot any issues with your virtual appliance, especially if the Unified Manager web interface is not available
- Upgrade to newer versions of Unified Manager
- Generate support bundles to send to technical support
- Configure network settings
- Change the maintenance user password
- Connect to an external data provider to send performance statistics
- Change the performance data collection internal
- Restore the Unified Manager database and configuration settings from a previously backed up version.

Learn about the maintenance user in Active IQ Unified Manager

The maintenance user is created during the installation of Unified Manager on a Red Hat Enterprise Linux system. The maintenance user name is the “umadmin” user. The maintenance user has the Application Administrator role in the web UI, and that user can create subsequent users and assign them roles.

The maintenance user, or umadmin user, can also access the Unified Manager maintenance console.

Diagnostic user capabilities in Active IQ Unified Manager

The purpose of diagnostic access is to enable technical support to assist you in troubleshooting, and you should only use it when directed by technical support.

The diagnostic user can execute OS-level commands when directed by technical support, for troubleshooting purposes.

Access the maintenance console in Active IQ Unified Manager

If the Unified Manager user interface is not in operation, or if you need to perform functions that are not available in the user interface, you can access the maintenance console to manage your Unified Manager system.

Before you begin

You must have installed and configured Unified Manager.

After 15 minutes of inactivity, the maintenance console logs you out.

When installed on VMware, if you have already logged in as the maintenance user through the VMware console, you cannot simultaneously log in using Secure Shell.

Step

1. Follow these steps to access the maintenance console:

On this operating system…	Follow these steps…
VMware	a. Using Secure Shell, connect to the IP address or fully qualified domain name of the Unified Manager virtual appliance. b. Log in to the maintenance console using your maintenance user name and password.
Linux	a. Using Secure Shell, connect to the IP address or fully qualified domain name of the Unified Manager system. b. Log in to the system with the maintenance user (umadmin) name and password. c. Enter the command `maintenance_console` and press Enter.
Windows	a. Log in to the Unified Manager system with administrator credentials. b. Launch PowerShell as a Windows administrator. c. Enter the command `maintenance_console` and press Enter.

The Unified Manager maintenance console menu is displayed.

Access the maintenance console using the vSphere VM console in Active IQ Unified Manager

If the Unified Manager user interface is not in operation, or if you need to perform functions that are not available in the user interface, you can access the maintenance console to reconfigure your virtual appliance.

Before you begin

- You must be the maintenance user.
- The virtual appliance must be powered on to access the maintenance console.

Steps

1. In vSphere Client, locate the Unified Manager virtual appliance.
2. Click the **Console** tab.

3. Click inside the console window to log in.
4. Log in to the maintenance console using your user name and password.

   After 15 minutes of inactivity, the maintenance console logs you out.

Maintenance console menus

Learn about maintenance console menus in Active IQ Unified Manager

The maintenance console consists of different menus that enable you to maintain and manage special features and configuration settings of the Unified Manager server.

Depending on the operating system on which you have installed Unified Manager, the maintenance console consists of the following menus:

- Upgrade Unified Manager (VMware only)
- Network Configuration (VMware only)
- System Configuration (VMware only)
  1. Support/Diagnostics
  2. Reset Server Certificate
  3. Regenerate Client Certificate
  4. Backup Restore
  5. External Data Provider
  6. Performance Polling Interval Configuration
  7. Disable SAML authentication
  8. View/Change Application Ports
  9. Debug Log Configuration
  10. Control access to MySQL port 3306
  11. Exit

You select the number from the list for accessing the specific menu option. For example, for backup and restore, you select *4*.

Network Configuration menu in Active IQ Unified Manager

The Network Configuration menu enables you to manage the network settings. You should use this menu when the Unified Manager user interface is not available.



This menu is not available if Unified Manager is installed on Red Hat Enterprise Linux or on Microsoft Windows.

The following menu choices are available.

- **Display IP Address Settings**

  Displays the current network settings for the virtual appliance, including the IP address, network, broadcast address, netmask, gateway, and DNS servers.

- **Change IP Address Settings**

  Enables you to change any of the network settings for the virtual appliance, including the IP address, netmask, gateway, or DNS servers. If you switch your network settings from DHCP to static networking using the maintenance console, you cannot edit the host name. You must select **Commit Changes** for the changes to take place.

- **Display Domain Name Search Settings**

  Displays the domain name search list used for resolving host names.

- **Change Domain Name Search Settings**

  Enables you to change the domain names for which you want to search when resolving host names. You must select **Commit Changes** for the changes to take place.

- **Display Static Routes**

  Displays the current static network routes.

- **Change Static Routes**

  Enables you to add or delete static network routes. You must select **Commit Changes** for the changes to take place.

  - **Add Route**

    Enables you to add a static route.

  - **Delete Route**

    Enables you to delete a static route.

  - **Back**

    Takes you back to the **Main Menu**.

  - **Exit**

    Exits the maintenance console.

- **Disable Network Interface**

  Disables any available network interfaces. If only one network interface is available, you cannot disable it. You must select **Commit Changes** for the changes to take place.

- **Enable Network Interface**

  Enables available network interfaces. You must select **Commit Changes** for the changes to take place.

- **Commit Changes**

  Applies any changes made to the network settings for the virtual appliance. You must select this option to enact any changes made, or the changes do not occur.

- **Ping a Host**

Pings a target host to confirm IP address changes or DNS configurations.

- **Restore to Default Settings**

  Resets all settings to the factory default. You must select **Commit Changes** for the changes to take place.

- **Back**

  Takes you back to the **Main Menu**.

- **Exit**

  Exits the maintenance console.

Use the System Configuration menu in Active IQ Unified Manager

The System Configuration menu enables you to manage your virtual appliance by providing various options, such as viewing the server status, and rebooting and shutting down the virtual machine.

When Unified Manager is installed on a Linux or Microsoft Windows system, only the “Restore from a Unified Manager Backup” option is available from this menu.

The following menu choices are available:

- **Display Server Status**

  Displays the current server status. Status options include Running and Not Running.

  If the server is not running, you might need to contact technical support.

- **Reboot Virtual Machine**

  Reboots the virtual machine, stopping all services. After rebooting, the virtual machine and services restart.

- **Shut Down Virtual Machine**

  Shuts down the virtual machine, stopping all services.

  You can select this option only from the virtual machine console.

- **Change <logged in user> User Password**

  Changes the password of the user that is currently logged in, which can only be the maintenance user.

- **Increase Data Disk Size**

  Increases the size of the data disk (disk 3) in the virtual machine.

- **Increase Swap Disk Size**

  Increases the size of the swap disk (disk 2) in the virtual machine.

- **Change Time Zone**

Changes the time zone to your location.

- **Change NTP Server**

  Changes the NTP Server settings, such as IP address or fully qualified domain name (FQDN).

- **Change NTP Service**

  Switches between the `ntp` and `systemd-timesyncd` services.

- **Restore from a Unified Manager Backup**

  Restores the Unified Manager database and configuration settings from a previously backed up version.

- **Reset Server Certificate**

  Resets the server security certificate.

- **Change hostname**

  Changes the name of the host on which the virtual appliance is installed.

- **Back**

  Exits the System Configuration menu and returns to the Main Menu.

- **Exit**

  Exits the maintenance console menu.

Use the Support and Diagnostics menu in Active IQ Unified Manager

The Support and Diagnostics menu enables you to generate a support bundle that you can send to technical support for troubleshooting assistance.

The following menu options are available:

- **Generate Light Support Bundle**

  Enables you to produce a lightweight support bundle that contains just 30 days of logs and configuration database records — it excludes performance data, acquisition recording files, and server heap dump.

- **Generate Support Bundle**

  Enables you to create a full support bundle (7-Zip file) containing diagnostic information in the diagnostic user's home directory. If your system is connected to the internet you can also upload the support bundle to NetApp.

  The file includes information generated by an AutoSupport message, the contents of the Unified Manager database, detailed data about the Unified Manager server internals, and verbose-level logs not normally included in AutoSupport messages or in the lightweight support bundle.

Learn about additional maintenance console menu options in Active IQ Unified Manager

The following menu options enable you to perform various administrative tasks on the

Unified Manager server.

The following menu choices are available:

- **Reset Server Certificate**

  Regenerates the HTTPS server certificate.

  You can regenerate the server certificate in the Unified Manager GUI by clicking **General** > **HTTPS Certificates** > **Regenerate HTTPS Certificate**.

- **Disable SAML authentication**

  Disables SAML authentication so that the identity provider (IdP) no longer provides sign-on authentication for users accessing the Unified Manager GUI. This console option is typically used when an issue with the IdP server or SAML configuration blocks users from accessing the Unified Manager GUI.

- **External Data Provider**

  Provides options for connecting Unified Manager to an external data provider. After you establish the connection, performance data is sent to an external server so that storage performance experts can chart the performance metrics using third-party software. The following options are displayed:

  - **Display Server Configuration**--Displays the current connection and configuration settings for an external data provider.
  - **Add / Modify Server Connection**--Enables you to enter new connection settings for an external data provider, or change existing settings.
  - **Modify Server Configuration**--Enables you to enter new configuration settings for an external data provider, or change existing settings.
  - **Delete Server Connection**--Deletes the connection to an external data provider.

    After the connection is deleted, Unified Manager loses its connection to the external server.

- **Backup Restore**

  For information, see the topics under Managing backup and restore operations.

- **Performance Polling Interval Configuration**

  Provides an option for configuring how frequently Unified Manager collects performance statistical data from clusters. The default collection interval is 5 minutes.

  You can change this interval to 10 or 15 minutes if you find that collections from large clusters are not completing on time.

- **View/Change Application Ports**

  Provides an option to change the default ports that Unified Manager uses for HTTP and HTTPS protocols, if required for security. The default ports are 80 for HTTP and 443 for HTTPS.

- **Control access to MySQL port 3306**

  Controls host access to the default MySQL port 3306. For reasons of security, the access through this port is restricted only to localhost during a fresh installation of Unified Manager on Linux, Windows, and

VMware vSphere systems. This option enables you to toggle the visibility of this port between the localhost and remote hosts, that is, if it is enabled for localhost only in your environment, you can make this port available to remote hosts as well. Alternately, when enabled for all hosts, you can restrict the access of this port to localhost only. If the access was enabled on remote hosts previously, the configuration is retained in an upgrade scenario. You should check the firewall settings on Windows systems after toggling the port visibility, and disable the firewall settings if the settings are configured to restrict access to MySQL port 3306.

- **Exit**

  Exits the maintenance console menu.

Change the maintenance user password on Windows in Active IQ Unified Manager

You can change the Unified Manager maintenance user password when required.

Steps

1. From the Unified Manager web UI login page, click **Forgot Password**.

   A page is displayed that prompts for the name of the user whose password you want to reset.

2. Enter the user name and click **Submit**.

   An email with a link to reset the password is sent to the email address that is defined for that user name.

3. Click the **reset password link** in the email and define the new password.
4. Return to the web UI and log in to Unified Manager using the new password.

Change the Active IQ Unified Manager umadmin password on Linux systems

For security reasons, you must change the default password for the Unified Manager umadmin user immediately after completing the installation process. If necessary, you can change the password again anytime later.

Before you begin

- Unified Manager must be installed on a Red Hat Enterprise Linux Linux system.
- You must have the root user credentials for the Linux system on which Unified Manager is installed.

Steps

1. Log in as the root user to the Linux system on which Unified Manager is running.
2. Change the umadmin password:

   ```
   passwd umadmin
   ```

   The system prompts you to enter a new password for the umadmin user.

Change the ports Active IQ Unified Manager uses for HTTP and HTTPS protocols

The default ports that Unified Manager uses for HTTP and HTTPS protocols can be changed after installation if required for security. The default ports are 80 for HTTP and

443 for HTTPS.

Before you begin

You must have a user ID and password authorized to log in to the maintenance console of the Unified Manager server.

> There are some ports that are considered unsafe when using the Mozilla Firefox or Google Chrome browsers. Check with your browser before assigning a new port number for HTTP and HTTPS traffic. Selecting an unsafe port could make the system inaccessible, which would require that you contact customer support for a resolution.

The instance of Unified Manager is restarted automatically after you change the port, so make sure this is a good time to take the system down for a short amount of time.

1. Log in using SSH as the maintenance user to the Unified Manager host.

   The Unified Managermaintenance console prompts are displayed.

2. Type the number of the menu option labeled **View/Change Application Ports**, and then press Enter.
3. If prompted, enter the maintenance user password again.
4. Type the new port numbers for the HTTP and HTTPS ports, and then press Enter.

   Leaving a port number blank assigns the default port for the protocol.

   You are prompted whether you want to change the ports and restart Unified Manager now.

5. Type **y** to change the ports and restart Unified Manager.
6. Exit out of the maintenance console.

After this change, users must include the new port number in the URL to access the Unified Manager web UI, for example https://host.company.com:1234, https://12.13.14.15:1122, or https://[2001:db8:0:1]:2123.

Add network interfaces to Active IQ Unified Manager

You can add new network interfaces if you need to separate network traffic.

Before you begin

You must have added the network interface to the virtual appliance using vSphere.

The virtual appliance must be powered on.

> You cannot perform this operation if Unified Manager is installed on Red Hat Enterprise Linux or on Microsoft Windows.

Steps

1. In the vSphere console Main Menu, select **System Configuration** > **Reboot Operating System**.

   After rebooting, the maintenance console can detect the newly added network interface.

2. Access the maintenance console.
3. Select **Network Configuration** > **Enable Network Interface**.

4. Select the new network interface and press **Enter**.

   Select **eth1** and press **Enter**.

5. Type **y** to enable the network interface.
6. Enter the network settings.

   You are prompted to enter the network settings if using a static interface, or if DHCP is not detected.

   After entering the network settings, you automatically return to the **Network Configuration** menu.

7. Select **Commit Changes**.

   You must commit the changes to add the network interface.

Regenerate the client certificate in Active IQ Unified Manager

You can use the maintenance console to regenerate your Unified Manager client certificate. Your certificate has an extended validity period of five years.

Steps

1. Log in as the maintenance user.
2. Select **Regenerate Client Certificate** and press **Enter**.
3. If the system prompts for authentication, enter the required credentials.
4. Select **Confirm Regeneration** and press **Enter**.
5. After the certificate is regenerated, restart the Unified Manager service to apply the changes.
6. To verify the updated certificate log in to ONTAP. View the client certificate status and the validity status.

Related information
Manage ONTAP certificates with System Manager

Add disk space

Add disk space to the Active IQ Unified Manager database directory

The Unified Manager database directory contains all of the health and performance data collected from ONTAP systems. Some circumstances may require that you increase the size of the database directory.

For example, the database directory may get full if Unified Manager is collecting data from a large number of clusters where each cluster has many nodes. You will receive a warning event when the database directory is 90% full, and a critical event when the directory is 95% full.



No additional data is collected from clusters after the directory reaches 95% full.

The steps required to add capacity to the data directory are different depending on whether Unified Manager is running on a VMware ESXi server, on a Red Hat server, or on a Microsoft Windows server.

Add space to the data directory on a Linux host in Active IQ Unified Manager

If you allotted insufficient disk space to the `/opt/netapp/data` directory to support Unified Manager when you originally set up the Linux host and then installed Unified Manager, you can add disk space after installation by increasing disk space on the `/opt/netapp/data` directory.

Before you begin

You must have root user access to the Red Hat Enterprise Linux machine on which Unified Manager is installed.

We recommend that you back up the Unified Manager database before increasing the size of the data directory.

Steps

1. Log in as root user to the Linux machine on which you want to add disk space.
2. Stop the Unified Manager service and the associated MySQL software in the order shown:

   ```
   systemctl stop ocieau ocie mysqld
   ```

3. Create a temporary backup folder (for example, `/backup-data`) with sufficient disk space to contain the data in the current `/opt/netapp/data` directory.
4. Copy the content and privilege configuration of the existing `/opt/netapp/data` directory to the backup data directory:

   ```
   cp -arp /opt/netapp/data/* /backup-data
   ```

5. If SE Linux is enabled:
   a. Get the SE Linux type for folders on existing `/opt/netapp/data` folder:

      ```
      se_type= ls -Z /opt/netapp/data | awk '{print $4}' | awk -F: '{print $3}' |
      head -1
      ```

      The system returns a confirmation similar to the following:

      ```
      echo $se_type
      mysqld_db_t
      ```

   b. Run the chcon command to set the SE Linux type for the backup directory:

      ```
      chcon -R --type=mysqld_db_t /backup-data
      ```

6. Remove the contents of the `/opt/netapp/data` directory:
   a. `cd /opt/netapp/data`
   b. `rm -rf *`
7. Expand the size of the `/opt/netapp/data` directory to a minimum of 150 GB through LVM commands or by adding extra disks.

If you have created `/opt/netapp/data` from a disk, then you should not try to mount `/opt/netapp/data` as an NFS or CIFS share. Because, in this case, if you try to expand the disk space, some LVM commands, such as `resize` and `extend` might not work as expected.

8. Confirm that the `/opt/netapp/data` directory owner (mysql) and group (root) are unchanged:

   ```
   ls -ltr /opt/netapp/ | grep data
   ```

   The system returns a confirmation similar to the following:

   ```
   drwxr-xr-x.  17 mysql root  4096 Aug 28 13:08 data
   ```

9. If SE Linux is enabled, confirm that the context for the `/opt/netapp/data` directory is still set to mysqld_db_t:

   a. `touch /opt/netapp/data/abc`

   b. `ls -Z /opt/netapp/data/abc`

      The system returns a confirmation similar to the following:

      ```
      -rw-r--r--. root root unconfined_u:object_r:mysqld_db_t:s0
      /opt/netapp/data/abc
      ```

10. Delete the file abc so that this extraneous file does not cause a database error in the future.
11. Copy the contents from backup-data back to the expanded `/opt/netapp/data` directory:

    ```
    cp -arp /backup-data/* /opt/netapp/data/
    ```

12. If SE Linux is enabled, run the following command:

    ```
    chcon -R --type=mysqld_db_t /opt/netapp/data
    ```

13. Start the MySQL service:

    ```
    systemctl start mysqld
    ```

14. After the MySQL service is started, start the ocie and ocieau services in the order shown:

    ```
    systemctl start ocie ocieau
    ```

15. After all of the services are started, delete the backup folder `/backup-data`:

    ```
    rm -rf /backup-data
    ```

Add space to the data disk on a VMware virtual machine in Active IQ Unified Manager

If you need to increase the amount of space on the data disk for the Unified Manager

database, you can add capacity after installation by increasing disk space using the Unified Manager maintenance console.

Before you begin

- You must have access to the vSphere Client.
- The virtual machine must have no snapshots stored locally.
- You must have the maintenance user credentials.

We recommend that you back up your virtual machine before increasing the size of virtual disks.

Steps

1. In the vSphere client, select the Unified Manager virtual machine, and then add more disk capacity to data `disk 3`. See the VMware documentation for details.

   In some rare cases the Unified Manager deployment uses “Hard Disk 2” for the data disk instead of “Hard Disk 3”. If this has occurred in your deployment, increase the space of whichever disk is larger. The data disk will always have more space than the other disk.

2. In the vSphere client, select the Unified Manager virtual machine, and then select the **Console** tab.
3. Click in the console window, and then log in to the maintenance console using your user name and password.
4. In the Main Menu, enter the number for the **System Configuration** option.
5. In the System Configuration Menu, enter the number for the **Increase Data Disk Size** option.

Add space to the logical drive on a Windows server in Active IQ Unified Manager

If you need to increase the amount of disk space for the Unified Manager database, you can add capacity to the logical drive on which Unified Manager is installed.

Before you begin

You must have Windows administrator privileges.

We recommend that you back up the Unified Manager database before adding disk space.

Steps

1. Log in as administrator to the Windows server on which you want to add disk space.
2. Follow the step that corresponds to method you want to use to add more space:

Option	Description
On a physical server, add capacity to the logical drive on which the Unified Manager server is installed.	Follow the steps in the Microsoft topic: Extend a Basic Volume
On a physical server, add a hard disk drive.	Follow the steps in the Microsoft topic: Adding Hard Disk Drives
On a virtual machine, increase the size of a disk partition.	Follow the steps in the VMware topic: Increasing the size of a disk partition

Manage user access

Add users

Add users in Active IQ Unified Manager

You can add local users or database users by using the Users page. You can also add remote users or groups that belong to an authentication server. You can assign roles to these users and, based on the privileges of the roles, users can manage the storage objects and data with Unified Manager, or view the data in a database.

Before you begin

- You must have the Application Administrator role.
- To add a remote user or group, you must have enabled remote authentication and configured your authentication server.
- If you plan to configure SAML authentication so that an identity provider (IdP) authenticates users accessing the graphical interface, make sure these users are defined as “remote” users.

  Access to the UI is not allowed for users of type “local” or “maintenance” when SAML authentication is enabled.

If you add a group from Windows Active Directory, then all direct members and nested subgroups can authenticate to Unified Manager, unless nested subgroups are disabled. If you add a group from OpenLDAP or other authentication services, then only the direct members of that group can authenticate to Unified Manager.

Steps

1. In the left navigation pane, click **General** > **Users**.
2. On the Users page, click **Add**.
3. In the Add User dialog box, select the type of user that you want to add, and enter the required information.

   When entering the required user information, you must specify an email address that is unique to that user. You must avoid specifying email addresses that are shared by multiple users.

4. Click **Add**.

Create a database user in Active IQ Unified Manager

To support a connection between Workflow Automation and Unified Manager, or to access database views, you must first create a database user with the Integration Schema or Report Schema role in the Unified Manager web UI.

Before you begin

You must have the Application Administrator role.

Database users provide integration with Workflow Automation and access to report-specific database views. Database users do not have access to the Unified Manager web UI or the maintenance console, and cannot execute API calls.

Steps

1. In the left navigation pane, click **General** > **Users**.
2. In the Users page, click **Add**.
3. In the Add User dialog box, select **Database User** in the **Type** drop-down list.
4. Type a name and password for the database user.
5. In the **Role** drop-down list, select the appropriate role.

If you are…	Choose this role
Connecting Unified Manager with Workflow Automation	Integration Schema
Accessing reporting and other database views	Report Schema

6. Click **Add**.

Edit user settings in Active IQ Unified Manager

You can edit user settings—such as the email address and role—that are specified each user. For example, you might want to change the role of a user who is a storage operator, and assign storage administrator privileges to the user.

Before you begin

You must have the Application Administrator role.

When you modify the role that is assigned to a user, the changes are applied when either of the following actions occur:

- The user logs out and logs back in to Unified Manager.
- Session timeout of 24 hours is reached.

Steps

1. In the left navigation pane, click **General** > **Users**.
2. In the Users page, select the user for which you want to edit settings, and click **Edit**.
3. In the Edit User dialog box, edit the appropriate settings that are specified for the user.
4. Click **Save**.

View users in Active IQ Unified Manager

You can use the Users page to view the list of users who manage storage objects and data using Unified Manager. You can view details about the users, such as the user name, type of user, email address, and the role that is assigned to the users.

Before you begin

You must have the Application Administrator role.

Step

1. In the left navigation pane, click **General** > **Users**.

Delete users or groups in Active IQ Unified Manager

You can delete one or more users from the management server database to prevent specific users from accessing Unified Manager. You can also delete groups so that all the users in the group can no longer access the management server.

Before you begin

- When you are deleting remote groups, you must have reassigned the events that are assigned to the users of the remote groups.

  If you are deleting local users or remote users, the events that are assigned to these users are automatically unassigned.

- You must have the Application Administrator role.

Steps

1. In the left navigation pane, click **General** > **Users**.
2. In the Users page, select the users or groups that you want to delete, and then click **Delete**.
3. Click **Yes** to confirm the deletion.

Learn about role-based access control in Active IQ Unified Manager

RBAC (role-based access control) provides the ability to control who has access to various features and resources in the Active IQ Unified Manager server. RBAC enables administrators to manage groups of users by defining roles. If you need to restrict access for specific functionality to selected administrators, you must set up administrator accounts for them. If you want to restrict the information that administrators can view and the operations they can perform, you must apply roles to the administrator accounts you create.

The management server uses RBAC for user login and role permissions. If you have not changed the management server's default settings for administrative user access, you do not need to log in to view them.

When you initiate an operation that requires specific privileges, the management server prompts you to log in. For example, to create administrator accounts, you must log in with Application Administrator account access.

Learn about user types in Active IQ Unified Manager

A user type specifies the kind of account the user holds and includes remote users, remote groups, local users, database users, and maintenance users. Each of these types has its own role, which is assigned by a user with the role of Administrator.

Unified Manager user types are as follows:

- **Maintenance user**

  Created during the initial configuration of Unified Manager. The maintenance user then creates additional users and assigns roles. The maintenance user is also the only user with access to the maintenance console. When Unified Manager is installed on a Red Hat Enterprise Linux system, the maintenance user is given the user name "umadmin."

- **Local user**

  Accesses the Unified Manager UI and performs functions based on the role given by the maintenance user or a user with the Application Administrator role.

- **Remote group**

  A group of users that access the Unified Manager UI using the credentials stored on the authentication server. The name of this account should match the name of a group stored on the authentication server. All users within the remote group are given access to the Unified Manager UI using their individual user credentials. Remote groups can perform functions according to their assigned roles.

- **Remote user**

  Accesses the Unified Manager UI using the credentials stored on the authentication server. A remote user performs functions based on the role given by the maintenance user or a user with the Application Administrator role.

- **Database user**

  Has read-only access to data in the Unified Manager database, has no access to the Unified Manager web interface or the maintenance console, and cannot execute API calls.

Learn about user roles in Active IQ Unified Manager

The maintenance user or Application Administrator assigns a role to every user. Each role contains certain privileges. The scope of activities that you can perform in Unified Manager depends on the role you are assigned and which privileges the role contains.

Unified Manager includes the following predefined user roles:

- **Operator**

  Views storage system information and other data collected by Unified Manager, including histories and capacity trends. This role enables the storage operator to view, assign, acknowledge, resolve, and add notes for the events.

- **Storage Administrator**

  Configures storage management operations within Unified Manager. This role enables the storage administrator to configure thresholds and to create alerts and other storage management-specific options and policies.

- **Application Administrator**

  Configures settings unrelated to storage management. This role enables the management of users, security certificates, database access, and administrative options, including authentication, SMTP, networking, and AutoSupport.

  When Unified Manager is installed on Linux systems, the initial user with the Application Administrator role is automatically named “umadmin”.

- **Integration Schema**

This role enables read-only access to Unified Manager database views for integrating Unified Manager with OnCommand Workflow Automation (WFA).

- **Report Schema**

  This role enables read-only access to reporting and other database views directly from the Unified Manager database. The databases that can be viewed include:

  - netapp_model_view
  - netapp_performance
  - ocum
  - ocum_report
  - ocum_report_birt
  - opm
  - scalemonitor

Learn about user roles and capabilities in Active IQ Unified Manager

Based on your assigned user role, you can determine which operations you can perform in Unified Manager.

The following table displays the functions that each user role can perform:

Function	Operator	Storage Administrator	Application Administrator	Integration Schema	Report Schema
View storage system information	•	•	•	•	•
View other data, such as histories and capacity trends	•	•	•	•	•
View, assign, and resolve events	•	•	•		
View storage service objects, such as SVM associations and resource pools	•	•	•		
View threshold policies	•	•	•		
Manage storage service objects, such as SVM associations and resource pools		•	•		

Function	Operator	Storage Administrator	Application Administrator	Integration Schema	Report Schema
Define alerts		•	•		
Manage storage management options		•	•		
Manage storage management policies		•	•		
Manage users			•		
Manage administrative options			•		
Define threshold policies			•		
Manage database access			•		
Manage integration with WFA and provide access to the database views				•	
Schedule and save reports		•	•		
Execute “Fix It” operations from Management Actions		•	•		
Provide read-only access to database views					•

Manage SAML authentication settings

Enable SAML authentication in Active IQ Unified Manager

After you have configured remote authentication settings, you can enable Security Assertion Markup Language (SAML) authentication so that remote users are authenticated by a secure identity provider (IdP) before they can access the Unified Manager web UI.

Note that only remote users will have access to the Unified Manager graphical user interface after SAML authentication has been enabled. Local users and Maintenance users will not be able to access the UI. This configuration does not impact users who access the maintenance console.

Use an identity provider in Active IQ Unified Manager

When configuring Unified Manager to use an identity provider (IdP) to perform SAML authentication for all remote users, you need to be aware of some required configuration settings so that the connection to Unified Manager is successful.

You must enter the Unified Manager URI and metadata into the IdP server. You can copy this information from the Unified ManagerSAML Authentication page. Unified Manager is considered the service provider (SP) in the Security Assertion Markup Language (SAML) standard.

Supported encryption standards

- Advanced Encryption Standard (AES): AES-128 and AES-256
- Secure Hash Algorithm (SHA): SHA-1 and SHA-256

Validated identity providers

- Shibboleth
- Active Directory Federation Services (ADFS)

ADFS configuration requirements

- You must define three claim rules in the following order that are required for Unified Manager to parse ADFS SAML responses for this relying party trust entry.

Claim rule	Value
SAM-account-name	Name ID
SAM-account-name	urn:oid:0.9.2342.19200300.100.1.1
Token groups — Unqualified Name	urn:oid:1.3.6.1.4.1.5923.1.5.1.1

- You must set the authentication method to “Forms Authentication” or users may receive an error when logging out of Unified Manager . Follow these steps:
  a. Open the ADFS Management Console.
  b. Click on the Authentication Policies folder on the left tree view.
  c. Under Actions on the right, click Edit Global Primary Authentication Policy.
  d. Set the Intranet Authentication Method to “Forms Authentication” instead of the default “Windows Authentication”.
- In some cases login through the IdP is rejected when the Unified Manager security certificate is CA-signed. There are two workarounds to resolve this issue:
  - Follow the instructions identified in the link to disable the revocation check on the ADFS server for chained CA cert associated relying party:

    Disable Revocation Check per Relying Party Trust

  - Have the CA server reside within the ADFS server to sign the Unified Manager server cert request.

Other configuration requirements

- The Unified Manager clock skew is set to 5 minutes, so the time difference between the IdP server and the Unified Manager server cannot be more than 5 minutes or authentication will fail.

Enable SAML authentication in Active IQ Unified Manager

You can enable Security Assertion Markup Language (SAML) authentication so that remote users are authenticated by a secure identity provider (IdP) before they can access the Unified Manager web UI.

Before you begin

- You must have configured remote authentication and verified that it is successful.
- You must have created at least one Remote User, or a Remote Group, with the Application Administrator role.
- The Identity provider (IdP) must be supported by Unified Manager and it must be configured.
- You must have the IdP URL and metadata.
- You must have access to the IdP server.

After you have enabled SAML authentication from Unified Manager, users cannot access the graphical user interface until the IdP has been configured with the Unified Manager server host information. So you must be prepared to complete both parts of the connection before starting the configuration process. The IdP can be configured before or after configuring Unified Manager.

Only remote users will have access to the Unified Manager graphical user interface after SAML authentication has been enabled. Local users and Maintenance users will not be able to access the UI. This configuration does not impact users who access the maintenance console, the Unified Manager commands, or ZAPIs.

Unified Manager is restarted automatically after you complete the SAML configuration on this page.

Steps

1. In the left navigation pane, click **General** > **SAML Authentication**.
2. Select the **Enable SAML authentication** checkbox.

   The fields required to configure the IdP connection are displayed.

3. Enter the IdP URI and the IdP metadata required to connect the Unified Manager server to the IdP server.

   If the IdP server is accessible directly from the Unified Manager server, you can click the **Fetch IdP Metadata** button after entering the IdP URI to populate the IdP Metadata field automatically.

4. Copy the Unified Manager host metadata URI, or save the host metadata to an XML text file.

   You can configure the IdP server with this information at this time.

5. Click **Save**.

   A message box displays to confirm that you want to complete the configuration and restart Unified Manager.

6. Click **Confirm and Logout** and Unified Manager is restarted.

The next time authorized remote users attempt to access the Unified Manager graphical interface they will enter their credentials in the IdP login page instead of the Unified Manager login page.

If not already completed, access your IdP and enter the Unified Manager server URI and metadata to complete the configuration.

When using ADFS as your identity provider, the Unified Manager GUI does not honor the ADFS timeout and will continue to work until the Unified Manager session timeout is reached. You can change the GUI session timeout by clicking **General** > **Feature Settings** > **Inactivity Timeout**.

Change the SAML authentication identity provider in Active IQ Unified Manager

You can change the identity provider (IdP) that Unified Manager uses to authenticate remote users.

Before you begin

- You must have the IdP URL and metadata.
- You must have access to the IdP.

The new IdP can be configured before or after configuring Unified Manager.

Steps

1. In the left navigation pane, click **General** > **SAML Authentication**.
2. Enter the new IdP URI and the IdP metadata required to connect the Unified Manager server to the IdP.

   If the IdP is accessible directly from the Unified Manager server, you can click the **Fetch IdP Metadata** button after entering the IdP URL to populate the IdP Metadata field automatically.

3. Copy the Unified Manager metadata URI, or save the metadata to an XML text file.
4. Click **Save Configuration**.

   A message box displays to confirm that you want to change the configuration.

5. Click **OK**.

Access the new IdP and enter the Unified Manager server URI and metadata to complete the configuration.

The next time authorized remote users attempt to access the Unified Manager graphical interface they will enter their credentials in the new IdP login page instead of the old IdP login page.

Update SAML authentication settings after an Active IQ Unified Manager security certificate change

Any change to the HTTPS security certificate installed on the Unified Manager server requires that you update the SAML authentication configuration settings. The certificate is updated if you rename the host system, assign a new IP address for the host system, or manually change the security certificate for the system.

After the security certificate is changed and the Unified Manager server is restarted, SAML authentication will

not function and users will not be able to access the Unified Manager graphical interface. You must update the SAML authentication settings on both the IdP server and on the Unified Manager server to re-enable access to the user interface.

Steps

1. Log into the maintenance console.
2. In the **Main Menu**, enter the number for the **Disable SAML authentication** option.

   A message displays to confirm that you want to disable SAML authentication and restart Unified Manager.

3. Launch the Unified Manager user interface using the updated FQDN or IP address, accept the updated server certificate into your browser, and log in using the maintenance user credentials.
4. In the **Setup/Authentication** page, select the **SAML Authentication** tab and configure the IdP connection.
5. Copy the Unified Manager host metadata URI, or save the host metadata to an XML text file.
6. Click **Save**.

   A message box displays to confirm that you want to complete the configuration and restart Unified Manager.

7. Click **Confirm and Logout** and Unified Manager is restarted.
8. Access your IdP server and enter the Unified Manager server URI and metadata to complete the configuration.

Identity provider	Configuration steps
ADFS	a. Delete the existing relying party trust entry in the ADFS management GUI. b. Add a new relying party trust entry using the `saml_sp_metadata.xml` from the updated Unified Manager server. c. Define the three claim rules that are required for Unified Manager to parse ADFS SAML responses for this relying party trust entry. d. Restart the ADFS Windows service.
Shibboleth	a. Update the new FQDN of Unified Manager server into the `attribute-filter.xml` and `relying-party.xml` files. b. Restart the Apache Tomcat web server and wait for port 8005 to come online.

9. Log in to Unified Manager and verify that SAML authentication works as expected through your IdP.

Disable SAML authentication in Active IQ Unified Manager

You can disable SAML authentication when you want to stop authenticating remote users through a secure identity provider (IdP) before they can log into the Unified Manager web UI. When SAML authentication is disabled, the configured directory service providers, such as Active Directory or LDAP, perform sign-on authentication.

After you disable SAML authentication, Local users and Maintenance users will be able to access the graphical user interface in addition to configured Remote users.

You can also disable SAML authentication using the Unified Manager maintenance console if you do not have access to the graphical user interface.

Unified Manager is restarted automatically after SAML authentication is disabled.

Steps

1. In the left navigation pane, click **General** > **SAML Authentication**.
2. Uncheck the **Enable SAML authentication** checkbox.
3. Click **Save**.

   A message box displays to confirm that you want to complete the configuration and restart Unified Manager.

4. Click **Confirm and Logout** and Unified Manager is restarted.

The next time remote users attempt to access the Unified Manager graphical interface they will enter their credentials in the Unified Manager login page instead of the IdP login page.

Access your IdP and delete the Unified Manager server URI and metadata.

Disable SAML authentication from the Active IQ Unified Manager maintenance console

You may need to disable SAML authentication from the maintenance console when there is no access to the Unified Manager GUI. This could happen in cases of misconfiguration or if the IdP is not accessible.

Before you begin

You must have access to the maintenance console as the maintenance user.

When SAML authentication is disabled, the configured directory service providers, such as Active Directory or LDAP, perform sign-on authentication. Local users and Maintenance users will be able to access the graphical user interface in addition to configured Remote users.

You can also disable SAML authentication from the Setup/Authentication page in the UI.

Unified Manager is restarted automatically after SAML authentication is disabled.

Steps

1. Log into the maintenance console.
2. In the **Main Menu**, enter the number for the **Disable SAML authentication** option.

   A message displays to confirm that you want to disable SAML authentication and restart Unified Manager.

3. Type **y**, and then press Enter and Unified Manager is restarted.

The next time remote users attempt to access the Unified Manager graphical interface they will enter their credentials in the Unified Manager login page instead of the IdP login page.

If required, access your IdP and delete the Unified Manager server URL and metadata.

Use the SAML Authentication page in Active IQ Unified Manager

You can use the SAML Authentication page to configure Unified Manager to authenticate remote users using SAML though a secure identity provider (IdP) before they can to log in to the Unified Manager web UI.

- You must have the Application Administrator role to create or modify the SAML configuration.
- You must have configured remote authentication.
- You must have configured at least one remote user or remote group.

After remote authentication and remote users have been configured, you can select the Enable SAML authentication checkbox to enable authentication using a secure identity provider.

- **IdP URI**

  The URI to access the IdP from the Unified Manager server. Example URIs are listed below.

  ADFS example URI:

  ```
  https://win2016-dc.ntap2016.local/federationmetadata/2007-
  06/federationmetadata.xml
  ```

  Shibboleth example URI:

  ```
  https://centos7.ntap2016.local/idp/shibboleth
  ```

- **IdP Metadata**

  The IdP metadata in XML format.

  If the IdP URL is accessible from the Unified Manager server, you can click the **Fetch IdP Metadata** button to populate this field.

- **Host System (FQDN)**

  The FQDN of the Unified Manager host system as defined during installation. You can change this value if necessary.

- **Host URI**

  The URI to access the Unified Manager host system from the IdP.

- **Host Metadata**

  The host system metadata in XML format.

Manage authentication

Edit authentication servers in Active IQ Unified Manager

You can change the port that the Unified Manager server uses to communicate with your authentication server.

Before you begin

You must have the Application Administrator role.

Steps

1. In the left navigation pane, click **General** > **Remote Authentication**.
2. Check the **Disable Nested Group Lookup** box.
3. In the **Authentication Servers** area, select the authentication server that you want to edit, and then click **Edit**.
4. In the **Edit Authentication Server** dialog box, edit the port details.
5. Click **Save**.

Delete authentication servers in Active IQ Unified Manager

You can delete an authentication server if you want to prevent the Unified Manager server from communicating with the authentication server. For example, if you want to change an authentication server that the management server is communicating with, you can delete the authentication server and add a new authentication server.

Before you begin

You must have the Application Administrator role.

When you delete an authentication server, remote users or groups of the authentication server will no longer be able to access Unified Manager.

Steps

1. In the left navigation pane, click **General** > **Remote Authentication**.
2. Select one or more authentication servers that you want to delete, and then click **Delete**.
3. Click **Yes** to confirm the delete request.

   If the **Use Secure Connection** option is enabled, then the certificates associated with the authentication server are deleted along with the authentication server.

Authenticate with Active Directory or OpenLDAP in Active IQ Unified Manager

You can enable remote authentication on the management server and configure the management server to communicate with your authentication servers so that users within the authentication servers can access Unified Manager.

You can use one of the following predefined authentication services or specify your own authentication service:

- Microsoft Active Directory

  You cannot use Microsoft Lightweight Directory Services.

- OpenLDAP

You can select the required authentication service and add the appropriate authentication servers to enable the remote users in the authentication server to access Unified Manager. The credentials for remote users or groups are maintained by the authentication server. The management server uses the Lightweight Directory Access Protocol (LDAP) to authenticate remote users within the configured authentication server.

For local users who are created in Unified Manager, the management server maintains its own database of user names and passwords. The management server performs the authentication and does not use Active Directory or OpenLDAP for authentication.

Audit Logging

Learn about audit logs in Active IQ Unified Manager

You can detect whether the audit logs have been compromised with using Audit Logs. All the activities performed by a user are monitored and logged in the Audit Logs. The audits are performed for all user interface and publicly exposed APIs' functionalities of Active IQ Unified Manager.

You can use the **Audit Log: File View** to view and access all the audit log files available in your Active IQ Unified Manager. The files in the Audit Log: File View are listed based on their creation date. This view displays information of all the audit log that are captured from the installation or upgrade to the present in the system. Whenever you perform an action in Unified Manager, the information is updated and is available in the logs. The status of each log file is captured using the "File Integrity Status" attribute which gets actively monitored to detect tampering or deletion of the log file. The audit logs can have one of the following states when the audit logs are available in the system:

State	Description
ACTIVE	File in which logs are being currently logged.
NORMAL	File which is inactive, compressed and stored in the system.
TAMPERED	File which has been compromised by a user who has manually edited the file.
MANUAL_DELETE	File which got deleted by an authorized user.
ROLLOVER_DELETE	File which got deleted due to Rolling off based on Rolling Configuration Policy.
UNEXPECTED_DELETE	File which got deleted due to unknown reasons.

The Audit Log page includes the following command buttons:

- Configure
- Delete
- Download

The **DELETE** button enables you to delete any of the audit logs listed in the Audit Logs view. You can delete an audit log and optionally provide a reason to delete the file which helps in future to determine a valid delete. The REASON column lists the reason along with the name of the user who performed the delete operation.

Deleting a log file will cause deletion of file from the system but the entry in the DB table will not be deleted.

You can download the audit logs from Active IQ Unified Manager using the **DOWNLOAD** button in the Audit Logs section and export the audit log files. The files that are marked “NORMAL” or “TAMPERED” are downloaded in a compressed `.gzip` format.

The audit log files are archived periodically and saved to the database for reference. Before archival, the audit logs are digitally signed for maintaining the security and integrity.

When a full AutoSupport bundle is generated, the support bundle includes both archived and active audit log files. But when a light support bundle is generated, it includes only the active audit logs. The archived audit logs are not included.

Configure Active IQ Unified Manager audit logs

You can use the **Configure** button in the Audit Logs section to configure rolling policy for Audit Log files and to also enable remote logging for the Audit Logs.

You can set the values in the **MAX FILE SIZE** and **AUDIT LOG RETENTION DAYS** as per the desired amount and frequency of data that you want to store in the system. The value in the field **TOTAL AUDIT LOG SIZE** is the size of the total audit log data present in the system. The roll over policy is determined by the values in the field **AUDIT LOG RETENTION DAYS**, **MAX FILE SIZE**, and **TOTAL AUDIT LOG SIZE**. When the size of the audit log backup reaches the value configured in **TOTAL AUDIT LOG SIZE**, then the file that was archived first is deleted. This means that the oldest file is deleted. But the file entry continues to be available in the database and is marked as “Rollover Delete”. The **AUDIT LOG RETENTION DAYS** value is for the number of the days the audit log files are preserved. Any file older than the value set in this field is rolled over.

Steps

1. Click **Audit Logs** > > **Configure**.
2. Enter values in the **MAX FILE SIZE**, **TOTAL AUDIT LOG SIZE**, and **AUDIT LOG RETENTION DAYS**.

   To enable remote logging, select **Enable Remote Logging**.

Enable remote logging of audit logs in Active IQ Unified Manager

You can select the **Enable Remote Logging** checkbox on the Configure Audit Logs dialog box to enable remote audit logging. You can use this feature to transfer audit logs to a remote Syslog server. This will enable you to manage your audit logs when there are space constraints.

The remote logging of audit logs provides a tamper-proof backup in case the audit log files on the Active IQ Unified Manager server are tampered.

Steps

1. In the **Configure Audit Logs** dialog box, select the **Enable Remote Logging** checkbox.

   Additional fields to configure remote logging are displayed.

2. Enter the **HOSTNAME** and **PORT** of the remote server you want to connect to.
3. In the **SERVER CA CERTIFICATE** field, click **BROWSE** to select a public certificate of the target server.

The certificate should be uploaded in `.pem` format. This certificate should be obtained from the target Syslog server and should not have expired. The certificate should contain the selected “hostname” as part of the `SubjectAltName` (SAN) attribute.

4. Enter the values for the following fields: **CHARSET**, **CONNECTION TIMEOUT**, **RECONNECTION DELAY**.

   The values should be in milliseconds for these fields.

5. Select the required Syslog format and TLS protocol version in the **FORMAT** and **PROTOCOL** fields.
6. Select the **Enable Client Authentication** checkbox if the target Syslog server requires certificate based authentication.

   You will need to download client authentication certificate and upload it to the Syslog server before saving the Audit Log configuration, otherwise the connection will fail. Depending on the type of Syslog server, you might need to create a hash of the client authentication certificate.

   Example: syslog-ng requires a <hash> of the certificate to be created using the command `openssl x509 -noout -hash -in cert.pem`, and then you should symbolically link the client authentication certificate to a file named after the <hash> .0.

7. Click **Save** to configure the connection with your server and enable remote logging.

   You will be redirected to the Audit Logs page.

> The **Connection Timeout** value can affect the configuration. If the configuration takes longer time to respond than the defined value, it can lead to configuration failure due to a connection error. To establish a successful connection, increase the **Connection Timeout** value, and try the configuration again.

Use the Remote Authentication page in Active IQ Unified Manager

You can use the Remote Authentication page to configure Unified Manager to communicate with your authentication server to authenticate remote users who attempt to log into the Unified Manager web UI.

You must have the Application Administrator or Storage Administrator role.

After you select the Enable remote authentication checkbox, you can enable remote authentication using an authentication server.

- **Authentication Service**

  Enables you to configure the management server to authenticate users in directory service providers, such as Active Directory, OpenLDAP, or specify your own authentication mechanism. You can specify an authentication service only if you have enabled remote authentication.

  - **Active Directory**
    - Administrator Name

      Specifies the administrator name of the authentication server.

    - Password

Specifies the password to access the authentication server.

- Base Distinguished Name

  Specifies the location of the remote users in the authentication server. For example, if the domain name of the authentication server is ou@domain.com, then the base distinguished name is **cn=ou,dc=domain,dc=com**.

- Disable Nested Group Lookup

  Specifies whether to enable or disable the nested group lookup option. By default this option is disabled. If you use Active Directory, you can speed up authentication by disabling support for nested groups.

- Use Secure Connection

  Specifies the authentication service used for communicating with authentication servers.

◦ **OpenLDAP**

- Bind Distinguished Name

  Specifies the bind distinguished name that is used along with the base distinguished name to find remote users in the authentication server.

- Bind Password

  Specifies the password to access the authentication server.

- Base Distinguished Name

  Specifies the location of the remote users in the authentication server. For example, if the domain name of the authentication server is ou@domain.com, then the base distinguished name is **cn=ou,dc=domain,dc=com**.

- Use Secure Connection

  Specifies that Secure LDAP is used for communicating with LDAP authentication servers.

◦ **Others**

- Bind Distinguished Name

  Specifies the bind distinguished name that is used along with the base distinguished name to find remote users in the authentication server that you have configured.

- Bind Password

  Specifies the password to access the authentication server.

- Base Distinguished Name

  Specifies the location of the remote users in the authentication server. For example, if the domain name of the authentication server is ou@domain.com, then the base distinguished name is **cn=ou,dc=domain,dc=com**.

- Protocol Version

Specifies the Lightweight Directory Access Protocol (LDAP) version that is supported by your authentication server. You can specify whether the protocol version must be automatically detected or set the version to 2 or 3.

- User Name Attribute

  Specifies the name of the attribute in the authentication server that contains user login names to be authenticated by the management server.

- Group Membership Attribute

  Specifies a value that assigns the management server group membership to remote users based on an attribute and value specified in the user's authentication server.

- UGID

  If the remote users are included as members of a GroupOfUniqueNames object in the authentication server, this option enables you to assign the management server group membership to the remote users based on a specified attribute in that GroupOfUniqueNames object.

- Disable Nested Group Lookup

  Specifies whether to enable or disable the nested group lookup option. By default this option is disabled. If you use Active Directory, you can speed up authentication by disabling support for nested groups.

- Member

  Specifies the attribute name that your authentication server uses to store information about the individual members of a group.

- User Object Class

  Specifies the object class of a user in the remote authentication server.

- Group Object Class

  Specifies the object class of all groups in the remote authentication server.

  The values that you enter for the *Member*, *User Object Class*, and *Group Object Class* attributes should be the same as those added in your Active Directory, OpenLDAP, and LDAP configurations. Otherwise, the authentication might fail.

- Use Secure Connection

  Specifies the authentication service used for communicating with authentication servers.

  If you want to modify the authentication service, ensure that you delete any existing authentication servers and add new authentication servers.

Authentication Servers area

The Authentication Servers area displays the authentication servers that the management server communicates with to find and authenticate remote users. The credentials for remote users or groups are maintained by the authentication server.

- **Command buttons**

  Enables you to add, edit, or delete authentication servers.

  - Add

    Enables you to add an authentication server.

    If the authentication server that you are adding is part of a high-availability pair (using the same database), you can also add the partner authentication server. This enables the management server to communicate with the partner when one of the authentication servers is unreachable.

  - Edit

    Enables you to edit the settings for a selected authentication server.

  - Delete

    Deletes the selected authentication servers.

- **Name or IP Address**

  Displays the host name or IP address of the authentication server that is used to authenticate the user on the management server.

- **Port**

  Displays the port number of the authentication server.

- **Test Authentication**

  This button validates the configuration of your authentication server by authenticating a remote user or group.

  While testing, if you specify only the user name, the management server searches for the remote user in the authentication server, but does not authenticate the user. If you specify both the user name and password, the management server searches and authenticates the remote user.

  You cannot test the authentication if remote authentication is disabled.

Manage security certificates

View the HTTPS security certificate in in Active IQ Unified Manager

You can compare the HTTPS certificate details to the retrieved certificate in your browser to ensure that your browser's encrypted connection to Unified Manager is not being intercepted.

Before you begin

You must have the Operator, Application Administrator, or Storage Administrator role.

Viewing the certificate enables you to verify the content of a regenerated certificate, or to view Subject Alt Names (SAN) from which you can access Unified Manager.

Step

1. In the left navigation pane, click **General** > **HTTPS Certificate**.

   The HTTPS certificate is displayed at the top of the page

If you need to view more detailed information about the security certificate than what is displayed on the HTTPS Certificate page, you can view the connection certificate in your browser.

Download a HTTPS certificate signing request in Active IQ Unified Manager

You can download a certification signing request for the current HTTPS security certificate so that you can provide the file to a Certificate Authority to sign. A CA-signed certificate helps prevent man-in-the-middle attacks and provides better security protection than a self-signed certificate.

Before you begin

You must have the Application Administrator role.

Steps

1. In the left navigation pane, click **General** > **HTTPS Certificate**.
2. Click **Download HTTPS Certificate Signing Request**.
3. Save the `<hostname>.csr` file.

You can provide the file to a Certificate Authority to sign, and then install the signed certificate.

Install a CA-signed HTTPS certificate in Active IQ Unified Manager

You can upload and install a security certificate after a Certificate Authority has signed and returned it. The file that you upload and install must be a signed version of the existing self-signed certificate. A CA-signed certificate helps prevent man-in-the middle attacks and provides better security protection than a self-signed certificate.

*What .Before you begin

You must have completed the following actions:

- Downloaded the Certificate Signing Request file and had it signed by a Certificate Authority
- Saved the certificate chain in PEM format
- Included all certificates in the chain, from the Unified Manager server certificate to the root signing certificate, including any intermediate certificates present

You must have the Application Administrator role.



If the validity of certificate for which a CSR was created is more than 397 days, then the validity will be reduced to 397 days by the CA before signing and returning the certificate

Steps

1. In the left navigation pane, click **General** > **HTTPS Certificate**.

2. Click **Install HTTPS Certificate**.
3. In the dialog box that is displayed, click **Choose file…** to locate the file to upload.
4. Select the file, and then click **Install** to install the file.

   For information, see Installing a HTTPS certificate generated using external tools.

Example certificate chain

The following example shows how the certificate chain file might appear:

```
-----BEGIN CERTIFICATE-----
<*Server certificate*>
-----END CERTIFICATE-----
-----BEGIN CERTIFICATE-----
<*Intermediate certificate \#1 (if present)*>
-----END CERTIFICATE-----
-----BEGIN CERTIFICATE-----
<*Intermediate certificate \#2 (if present)*>
-----END CERTIFICATE-----
-----BEGIN CERTIFICATE-----
<*Root signing certificate*>
-----END CERTIFICATE-----
```

Learn about HTTPS certificates generated using external tools in Active IQ Unified Manager Unified Manager

You can install certificates that are self signed or CA signed and are generated using an external tool like OpenSSL, BoringSSL, LetsEncrypt.

You should load the private key along with the certificate chain because these certificates are externally generated public-private key pair. The permitted key-pair algorithms are “RSA” and “EC”. The **Install HTTPS Certificate** option is available in the HTTPS Certificates page under the General section. The file you upload should be in the following input format.

1. Private Key of the server that belongs to the Active IQ Unified Manager host
2. Certificate of the server that matches with the private key
3. Certificate of the CAs in reverse till the root, which are used to sign the above certificate

Format for loading a certificate with an EC key pair

The permitted curves are “prime256v1” and “secp384r1”. Sample of certificate with an externally generated EC pair:

```
 -----BEGIN EC PRIVATE KEY-----
<EC private key of Server>
-----END EC PRIVATE KEY-----
```

```
-----BEGIN CERTIFICATE-----
<Server certificate>
-----END CERTIFICATE-----
-----BEGIN CERTIFICATE-----
<Intermediate certificate #1 (if present)>
-----END CERTIFICATE-----
-----BEGIN CERTIFICATE-----
<Intermediate certificate #2 (if present)>
-----END CERTIFICATE-----
-----BEGIN CERTIFICATE-----
<Root signing certificate>
-----END CERTIFICATE-----
```

Format for loading a certificate with an RSA key pair

The allowed key sizes for the RSA key-pair belonging to the host certificate are 2048, 3072, and 4096. certificate with an externally generated **RSA key pair**:

```
-----BEGIN RSA PRIVATE KEY-----
 <RSA private key of Server>
 -----END RSA PRIVATE KEY-----
 -----BEGIN CERTIFICATE-----
 <Server certificate>
 -----END CERTIFICATE-----
 -----BEGIN CERTIFICATE-----
 <Intermediate certificate #1 (if present)>
 -----END CERTIFICATE-----
 -----BEGIN CERTIFICATE-----
 <Intermediate certificate #2 (if present)>
 -----END CERTIFICATE-----
 -----BEGIN CERTIFICATE-----
 <Root signing certificate>
 -----END CERTIFICATE-----
```

After the certificate is uploaded, you should restart the Active IQ Unified Manager instance for the changes to take effect.

Checks while uploading externally generated certificates

The system performs checks while uploading a certificate generated using external tools. If any of the checks fail, then the certificate is rejected. There are also validation included for the certificates that are generated from the CSR within the product and for certificates that are generated using external tools.

- The private key in the input is validated against the host certificate in the input.
- The Common Name (CN) in the host certificate is checked against the FQDN of the host.

- The Common Name (CN) of the host certificate should not be empty or blank and should not be set to localhost.
- The validity start date should not be in future and the validity expiry date of the certificate should not be in the past.
- If Intermediate CA or CA exists the validity start date of certificate should not be in future and the validity expiry date should not be in the past.

The private key in the input should not be encrypted. If there are any private keys that are encrypted, then they are rejected by the system.

Example 1

```
----BEGIN ENCRYPTED PRIVATE KEY-----
<Encrypted private key>
-----END ENCRYPTED PRIVATE KEY-----
```

Example 2

```
-----BEGIN RSA PRIVATE KEY-----
Proc-Type: 4,ENCRYPTED
<content here>
-----END RSA PRIVATE KEY-----
```

Example 3

```
-----BEGIN EC PRIVATE KEY-----
Proc-Type: 4,ENCRYPTED
<content here>
-----END EC PRIVATE KEY-----
```

If the certificate installation fails, see the knowledge base (KB) article:
ActiveIQ Unified Manager fails to install an externally generated certificate

Page descriptions for certificate management

Learn about the HTTPS certificate page in Active IQ Unified Manager

The HTTPS Certificate page enables you to view the current security certificate, download a certificate signing request, generate a new self-signed HTTPS certificate, or install a new HTTPS certificate.

If you have not generated a new self-signed HTTPS certificate, the certificate that appears on this page is the certificate that was generated during installation.

Command buttons

The command buttons enable you to perform the following operations:

- **Download HTTPS Certificate Signing Request**

  Downloads a certification request for the currently installed HTTPS certificate. Your browser prompts you to save the <hostname>.csr file so that you can provide the file to a Certificate Authority to sign.

- **Install HTTPS Certificate**

  Enables you to upload and install a security certificate after a Certificate Authority has signed and returned it. The new certificate is in effect after you restart the management server.

- **Regenerate HTTPS Certificate**

  Enables you to generate a new self-signed HTTPS certificate, which replaces the current security certificate. The new certificate is in effect after you restart Unified Manager.

Regenerate a HTTPS Certificate dialog box in Active IQ Unified Manager

The Regenerate HTTPS Certificate dialog box enables you to customize the security information and then generate a new HTTPS certificate with that information.

The current certificate information appears on this page.

The “Regenerate Using Current Certificate Attributes” and “Update the Current Certificate Attributes” selection enables you to regenerate the certificate with the current information or generate a certificate with new information.

- **Common Name**

  Required. The fully qualified domain name (FQDN) that you wish to secure.

  In Unified Manager high availability configurations, use the virtual IP address.

- **Email**

  Optional. An email address to contact your organization; typically the email address of the certificate administrator or IT department.

- **Company**

  Optional. Typically the incorporated name of your company.

- **Department**

  Optional. The name of the department in your company.

- **City**

  Optional. The city location of your company.

- **State**

  Optional. The state or province location, not abbreviated, of your company.

- **Country**

  Optional. The country location of your company. This is typically a two-letter ISO code of the country.

- **Alternative Names**

  Required. Additional, non-primary domain names that can be used to access this server in addition to the existing localhost or other network addresses. Separate each alternate name with a comma.

  Select the “Exclude local identifying information (e.g. localhost)” checkbox if you want to remove the local identifying information from the Alternative Names field in the certificate. When this checkbox is selected only what you enter in the field is used in the Alternative Names field. When left blank the resulting certificate will not have an Alternative Names field at all.

- **KEY SIZE (KEY ALGORITHM: RSA)**

  The key algorithm is set to RSA. You can select from one of the key sizes: 2048, 3072 or 4096 bits. The default key size is set to 2048 bits.

- **VALIDITY PERIOD**

  The default validity period is 397 days. If you have upgraded from a previous version, you might see the previous certificate validity unchanged.

For more information, see Generating HTTPS certificates.

Monitor and manage storage

Introduction to Active IQ Unified Manager

Learn about Active IQ Unified Manager

Active IQ Unified Manager (formerly OnCommand Unified Manager) enables you to monitor and manage the health and performance of your ONTAP storage systems from a single interface.

Unified Manager provides the following features:

- Discovery, monitoring, and notifications for systems that are installed with ONTAP software.
- Dashboard to show capacity, security, and performance health of the environment.
- Enhanced alerts, events, and threshold infrastructure.
- Displays detailed graphs that plot workload activity over time; including IOPS (operations), MBps (throughput), latency (response time), utilization, performance capacity, and cache ratio.
- Identifies workloads that are overusing cluster components and the workloads whose performance is impacted by the increased activity.
- Provides suggested corrective actions that can be performed to address certain incidents and events, and a "Fix It" button for some events so you can resolve the issue immediately.
- Integrates with OnCommand Workflow Automation to execute automated protection workflows.
- Ability to create new workloads, such as a LUN or file share, directly from Unified Manager and assign a Performance Service Level to define the performance and storage objectives for the users accessing the application using that workload.

Learn about health monitoring in Active IQ Unified Manager

Active IQ Unified Manager (formerly OnCommand Unified Manager) helps you to monitor a large number of systems running ONTAP software through a centralized user interface. The Unified Manager server infrastructure delivers scalability, supportability, and enhanced monitoring and notification capabilities.

The key capabilities of Unified Manager include monitoring, alerting, managing availability and capacity of clusters, managing protection capabilities, and bundling of diagnostic data and sending it to technical support.

You can use Unified Manager to monitor your clusters. When issues occur in the cluster, Unified Manager notifies you about the details of such issues through events. Some events also provide you with a remedial action that you can take to rectify the issues. You can configure alerts for events so that when issues occur, you are notified through email, and SNMP traps.

You can use Unified Manager to manage storage objects in your environment by associating them with annotations. You can create custom annotations and dynamically associate clusters, storage virtual machines (SVMs), and volumes with the annotations through rules.

You can also plan the storage requirements of your cluster objects using the information provided in the capacity and health charts, for the respective cluster object.

Physical and logical capacity

Unified Manager makes use of the concepts of physical and logical space used for ONTAP storage objects.

- Physical capacity: Physical space refers to the physical blocks of storage used in the volume. “Physical used capacity” is typically smaller than logical used capacity due to the reduction of data from storage efficiency features (such as deduplication and compression).
- Logical capacity: Logical space refers to the usable space (the logical blocks) in a volume. Logical space refers to how theoretical space can be used, without accounting for results of deduplication or compression. “Logical space used” is physical space used plus the savings from storage efficiency features (such as deduplication and compression) that have been configured. This measurement often appears larger than the physical used capacity because this does not reflect the data compression and other reductions in the physical space. Thus, the total logical capacity could be higher than the provisioned space.

Capacity measurement units

Unified Manager calculates storage capacity based on binary units of 1024 (2^{10}) bytes. In ONTAP 9.10.0 and earlier, these units were displayed as KB, MB, GB, TB, and PB. Beginning with ONTAP 9.10.1, they are displayed in Unified Manager as KiB, MiB, GiB, TiB, and PiB.

The units used for throughput continue to be kilobytes per second (Kbps), Megabytes per second (Mbps), Gigabytes per second (Gbps), or Terabytes per second (Tbps) and so forth, for all releases of ONTAP.

Capacity unit displayed in Unified Manager for ONTAP 9.10.0 and earlier	Capacity unit displayed in Unified Manager for ONTAP 9.10.1	Calculation	Value in bytes
KB	KiB	1024	1024 bytes
MB	MiB	1024 * 1024	1,048,576 bytes
GB	GiB	1024 * 1024 * 1024	1,073,741,824 bytes
TB	TiB	1024 * 1024 * 1024 * 1024	1,099,511,627,776 bytes

Learn about performance monitoring in Active IQ Unified Manager

Active IQ Unified Manager (formerly OnCommand Unified Manager) provides performance monitoring capabilities and event root-cause analysis for systems that are running NetApp ONTAP software.

Unified Manager helps you to identify workloads that are overusing cluster components and decreasing the performance of other workloads on the cluster. By defining performance threshold policies you can also specify maximum values for certain performance counters so that events are generated when the threshold is breached. Unified Manager alerts you about these performance events so that you can take corrective action, and bring performance back to normal levels of operation. You can view and analyze events in the Unified Manager UI.

Unified Manager monitors the performance of two types of workloads:

- User-defined workloads

  These workloads consist of FlexVol volumes and FlexGroup volumes that you have created in your cluster.

- System-defined workloads

  These workloads consist of internal system activity.

Use Active IQ Unified Manager REST APIs

Active IQ Unified Manager provides you with REST APIs to view the information about monitoring and managing your storage environment. APIs also allow provisioning and managing storage objects based on policies.

You can also execute ONTAP APIs on all ONTAP-managed clusters by using the API gateway supported by Unified Manager.

For information about Unified Manager REST APIs, see Getting started with Active IQ Unified Manager REST APIs.

What the Unified Manager server does

Learn about the Active IQ Unified Manager server

The Unified Manager server infrastructure consists of a data collection unit, a database, and an application server. It provides infrastructure services such as discovery, monitoring, role-based access control (RBAC), auditing, and logging.

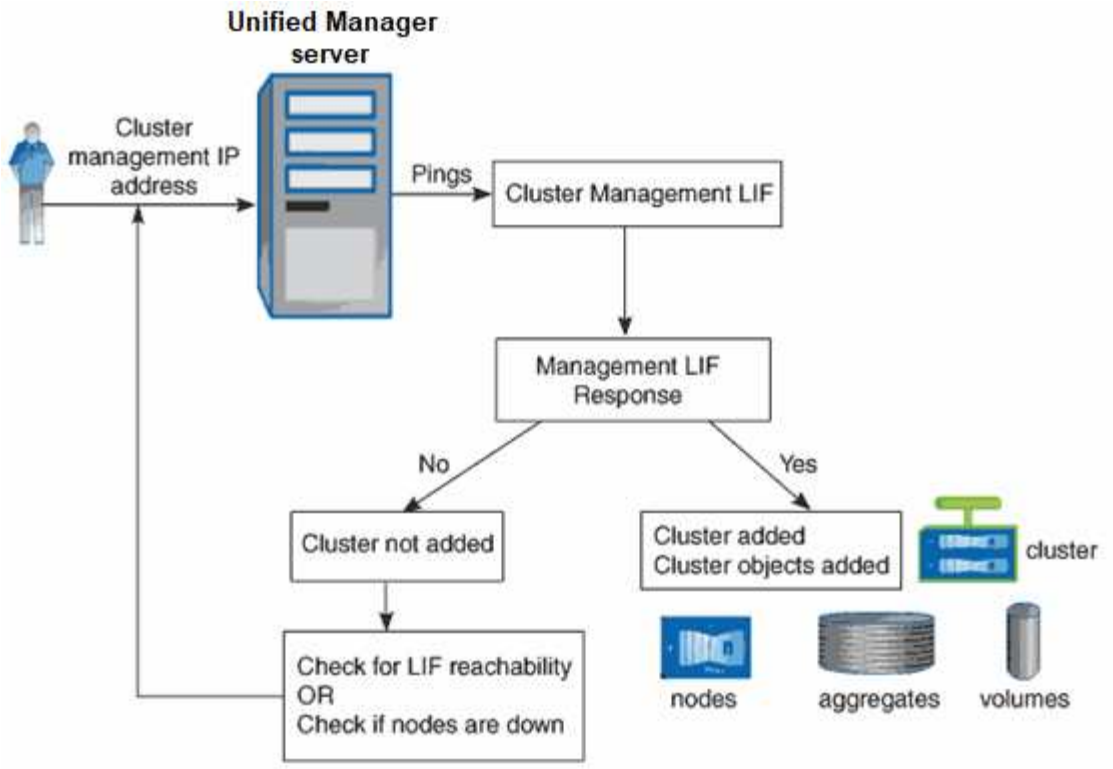
Unified Manager collects cluster information, stores the data in the database, and analyzes the data to see if there are any cluster issues.

Learn how the cluster discovery process works in Active IQ Unified Manager

After you have added the cluster to Unified Manager, the server discovers the cluster objects and adds them to its database. Understanding how the discovery process works helps you to manage your organization’s clusters and their objects.

The default monitoring interval is 15 minutes: if you have added a cluster to Unified Manager server, it takes 15 minutes to display the cluster details in the Unified Manager UI.

The following image illustrates the discovery process in Active IQ Unified Manager:



Understand the user interface

Learn about the Active IQ Unified Manager user interface

The Unified Manager user interface mainly consists of a dashboard that provides an at-a-glance view of the objects that are monitored. The user interface also provides access to viewing all the cluster objects.

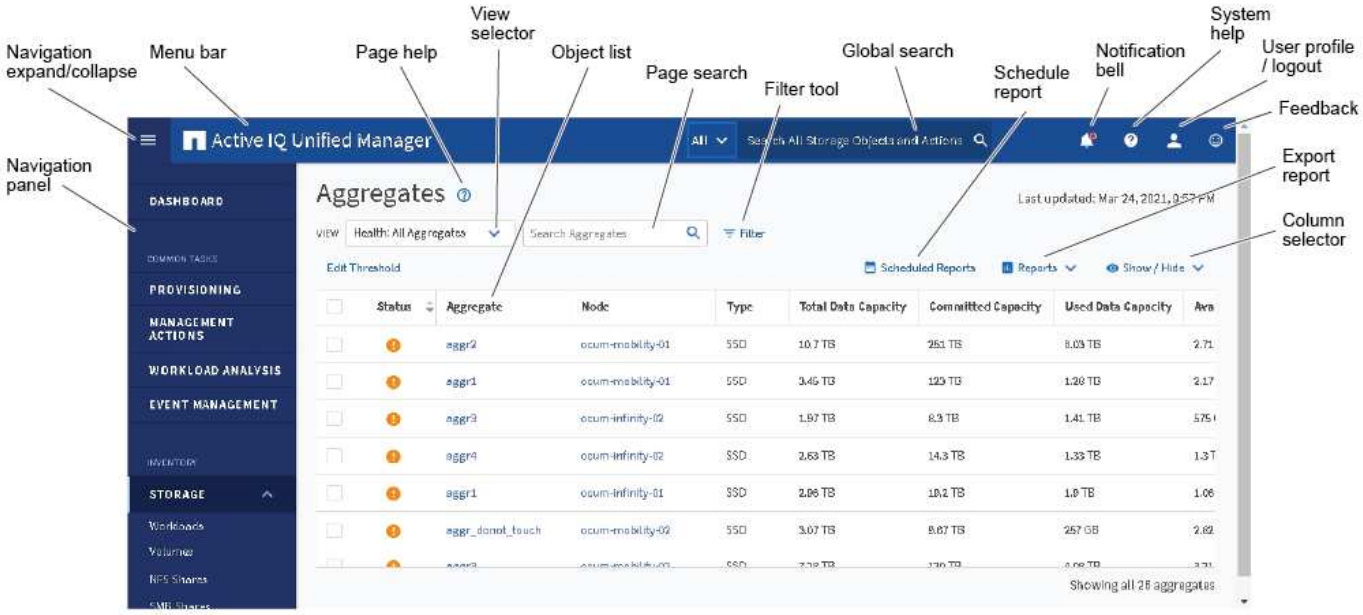
You can select a preferred view and use the action buttons as necessary. Your screen configuration is saved in a workspace so that all of the functionality you require is available when you start Unified Manager. However, when you navigate from one view to another, and then navigate back, the view might not be the same.

Typical window layouts in Active IQ Unified Manager

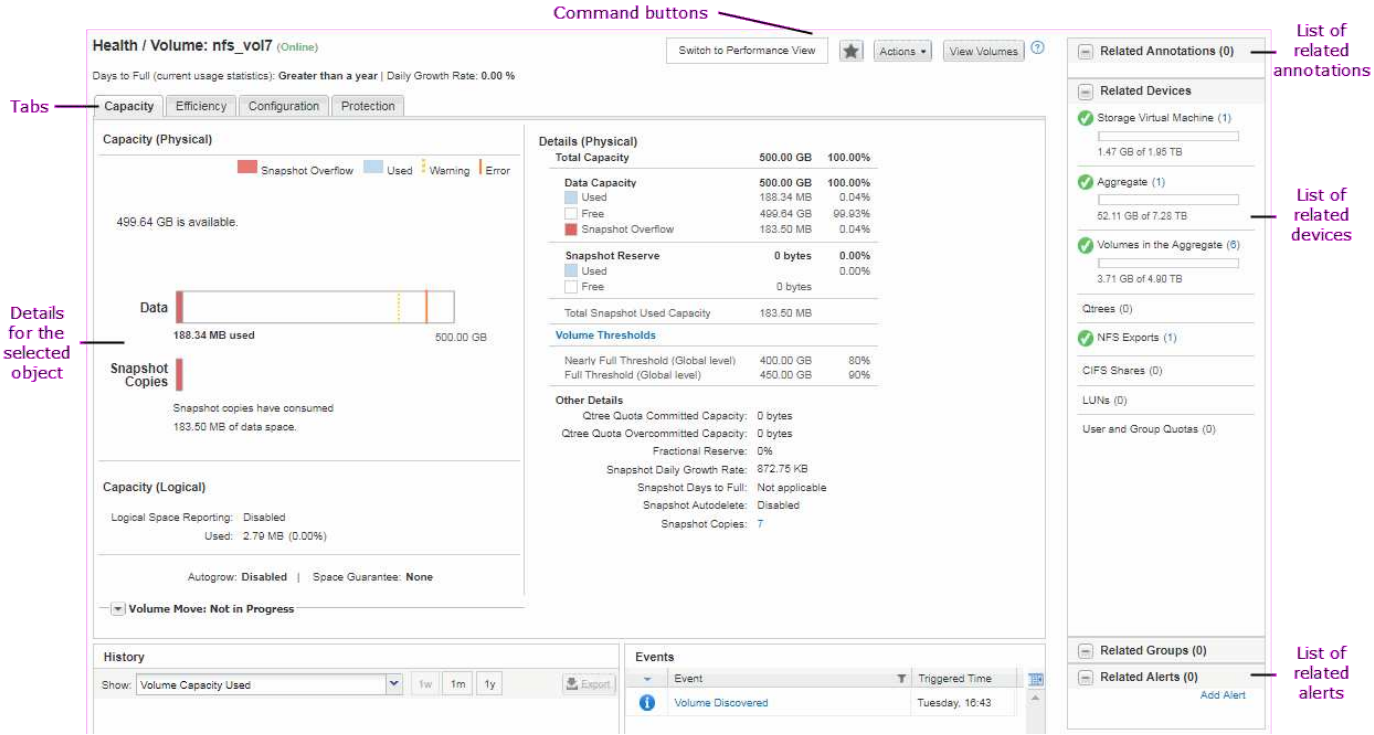
Understanding the typical window layouts helps you to navigate and use Active IQ Unified Manager effectively. Most Unified Manager windows are similar to one of two general layouts: object list or details. The recommended display setting is at least 1280 by 1024 pixels.

Not every window contains every element in the following diagrams.

Object list window layout



Object details window layout



Customizing window layout in Active IQ Unified Manager

Active IQ Unified Manager enables you to customize the layout of information on the storage and network object pages. By customizing the windows, you can control which data is viewable and how the data is displayed.

- **Sorting**

  You can click the column header to change the sort order of the column entries. When you click the column

header, the sort arrows (and) appear for that column.

- **Filtering**

  You can click the filter icon ( ) to apply filters to customize the display of information on the storage and network object pages so that only those entries that match the conditions that are provided are displayed. You apply filters from the Filters pane.

  The Filters pane enables you to filter most of the columns based on the options that are selected. For example, on the Health: All Volumes view, you can use the Filters pane to display all of the volumes that are offline by selecting the appropriate filter option under State.

  Capacity-related columns in any list always display capacity data in appropriate units rounded off to two decimal points. This also applies when filtering capacity columns. For example, if you use the filter in the Total Data Capacity column in the Health: All Aggregates view to filter data greater than 20.45 GB, the actual capacity of 20.454 GB is displayed as 20.45 GB. Similarly, if you filter data less than 20.45 GB, the actual capacity of 20.449 GB is displayed as 20.45 GB.

  If you use the filter in the Available Data % column in the Health: All Aggregates view to filter data greater than 20.45%, the actual capacity of 20.454% is displayed as 20.45%. Similarly, if you filter data less than 20.45%, the actual capacity of 20.449% is displayed as 20.45%.

- **Hiding or showing the columns**

  You can click the column display icon (**Show/Hide**) to select which columns you want to display. Once you have selected the appropriate columns you can re-order them by dragging them using your mouse.

- **Searching**

  You can use the search box to search for certain object attributes to help refine the list of items in the inventory page. For example, you can enter "cloud" to refine the list of volumes in the volumes inventory page to see all volumes that have the word "cloud" in them.

- **Exporting data**

  You can click the **Reports** button (or **Export** button to export data to a comma-separated values `(.csv)` file, `(.pdf)` document, or Microsoft Excel `(.xlsx)` file and use the exported data to build reports.

Use the Active IQ Unified Manager Help

The Help includes information about all features included in Active IQ Unified Manager. You can use the table of contents, the index, or the search tool to find information about the features and how to use them.

Help is available from each tab and from the menu bar of the Unified Manager user interface.

The search tool in the Help does not work for partial words.

- To learn about specific fields or parameters, click .
- To view all the Help contents, click > **Help/Documentation** in the menu bar.

  You can find more detailed information by expanding any portion of the Table of Contents in the navigation pane.

- To search the Help contents, click the **Search** tab in the navigation pane, type the word or series of words you want to find, and click **Go!**
- To print Help topics, click the printer icon.

Bookmark your favorite Help topics in Active IQ Unified Manager

In the Help Favorites tab, you can bookmark Help topics that you use frequently. Help bookmarks provide fast access to your favorite topics.

Steps

1. Navigate to the Help topic that you want to add as a favorite.
2. Click **Favorites**, and then click **Add**.

Search for storage objects in Active IQ Unified Manager

To quickly access a specific object, you can use the **Search all Storage Objects** field at the top of the menu bar. This method of global search across all objects enables you to quickly locate specific objects by type. Search results are sorted by storage object type and you can filter them further by object using the drop-down menu.

Before you begin

- You must have one of the following roles to perform this task: Operator, Application Administrator, or Storage Administrator.
- A valid search must contain at least three characters.

When using the drop-down menu value "All", the global search displays the total number of results found in all object categories; with a maximum of 25 search results for each object category. You can select a specific object type from the drop-down menu to refine the search within a specific object type. In this case the returned list is not restricted to the top 25 objects.

The object types you can search for include:

- Clusters
- Nodes
- Storage VMs
- Aggregates
- Volumes
- Qtrees
- SMB Shares
- NFS Shares
- User or Group Quotas
- LUNs
- NVMe Namespaces
- Initiator Groups
- Initiators

- Consistency Group

Entering a workload name returns the list of workloads under the appropriate Volumes or LUNs category.

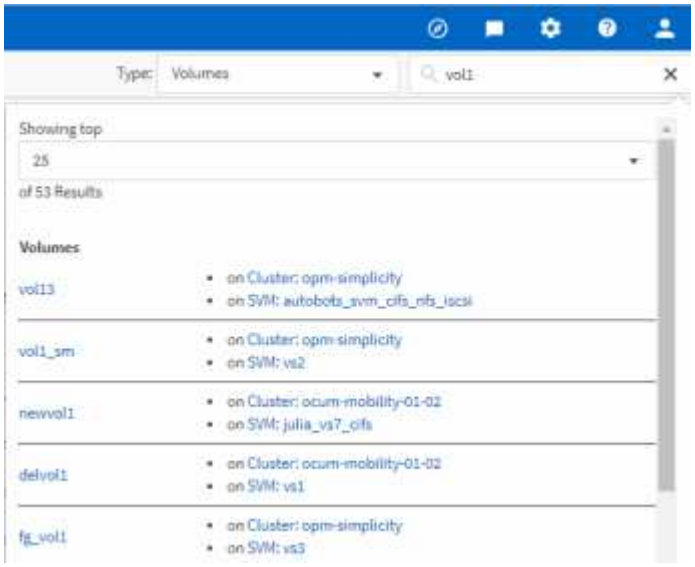
You can click any object in the search results to navigate to the Health details page for that object. If there is no direct health page for an object, then the Health page of the parent object is displayed. For example, when searching for a specific LUN, the SVM details page on which the LUN resides is displayed.

Ports and LIFs are not searchable in the global search bar.

Steps

1. Select an object type from the menu to refine the search results for only a single object type.
2. Type a minimum of three characters of the object name in the **Search all Storage Objects** field.

   In this example, the drop-down box has the Volumes object type selected. Typing "vol1" into the **Search all Storage Objects** field displays a list of all volumes whose names contain these characters.



Export storage data as reports in Active IQ Unified Manager

You can export storage data in a variety of output formats and then use the exported data to build reports. For example, if there are 10 critical events that have not been resolved, you can export the data from the Event Management inventory page to create a report, and then send the report to admins who can resolve the issues.

You can export data to a `.csv` file, `.xlsx` file, or `.pdf` document from the **Storage** and **Network** inventory pages and use the exported data to build reports. There are other locations in the product where only `.csv` or `.pdf` files can be generated.

Steps

1. Perform one of the following actions:

If you want to export…	Do this…
Storage object inventory details	Click **Storage** or **Network** from the left-navigation menu, and then select a storage object. Choose one of the system-provided views, or any custom view that you have created.
QoS Policy Group details	Click **Storage** > **QoS Policy Groups** from the left-navigation menu.
Storage capacity and protection history details	Click **Storage** > **Aggregates** or **Storage** > **Volumes**, then select a single aggregate or volume.
Event details	Click **Event Management** from the left-navigation menu.
Storage object top 10 performance details	Click **Storage** > **Clusters** > **Performance:All Clusters**, then select a cluster and choose the **Top Performers** tab. Then select a storage object and performance counter.

2. Click the **Reports** button (or **Export** button in some UI pages).
3. Click **Download CSV**, **Download PDF**, or **Download Excel** to confirm the export request.

   From the Top Performers tab you can choose to download a report of the statistics for the single cluster you are viewing or for all clusters in the data center.

   The file is downloaded.

4. Open the file in the appropriate application.

Related information

Health/Clusters inventory page

Scheduling a report

Filter inventory page content in Active IQ Unified Manager

You can filter inventory page data in Unified Manager to quickly locate data based on specific criteria. You can use filtering to narrow the contents of the Unified Manager pages to show only the results in which you are interested. This provides a very efficient method of displaying only the data in which you are interested.

Use **Filtering** to customize the grid view based on your preferences. Available filter options are based on the object type being viewed in the grid. If filters are currently applied, the number of applied filters displays at the right of the Filter button.

Three types of filter parameters are supported.

Parameter	Validation
String (text)	The operators are **contains**, **starts with**, **ends with**, and **does not contain**.

Parameter	Validation
Number	The operators are **greater than**, **less than**, **in the last**, and **between**.
Enum (text)	The operators are **is** and **is not**.

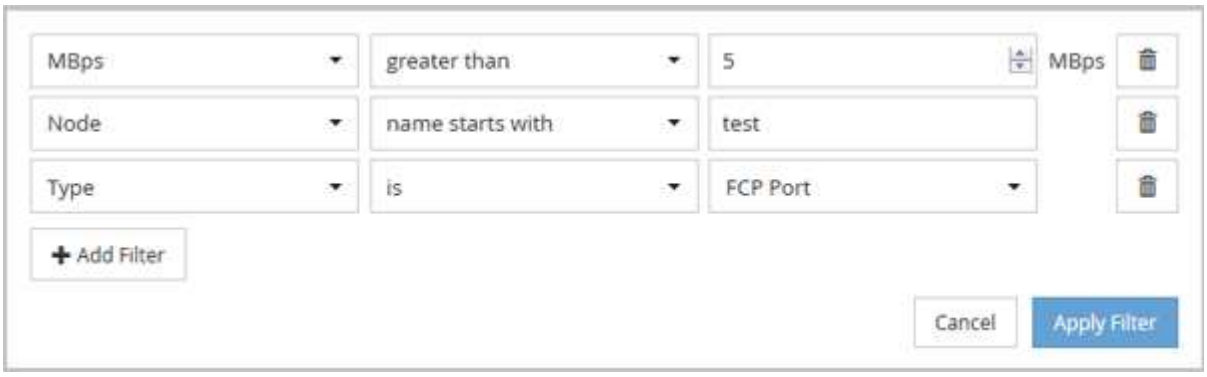
The Column, Operator, and Value fields are required for each filter; the available filters reflect the filterable columns on the current page. The maximum number of filters you can apply is four. Filtered results are based on combined filter parameters. Filtered results apply to all pages in your filtered search, not just the page currently displayed.

You can add filters using the Filtering panel.

1. At the top of the page, click the **Filter** button. The Filtering panel displays.
2. Click the left drop-down list and select an object; for example, *Cluster*, or a performance counter.
3. Click the center drop-down list, and select the operator you want to use.
4. In the last list, select or enter a value to complete the filter for that object.
5. To add another filter, click **+Add Filter**. An additional filter field displays. Complete this filter using the process described in the preceding steps. Note that upon adding your fourth filter, the **+Add Filter** button no longer displays.
6. Click **Apply Filter**. The filter options are applied to the grid and the number of filters is displayed to the right of the Filter button.
7. Use the Filtering panel to remove individual filters by clicking the trash icon at the right of the filter to be removed.
8. To remove all filters, click **Reset** at the bottom of the filtering panel.

Filtering example

The illustration shows the Filtering panel with three filters. The **+Add Filter** button displays when you have fewer than the maximum of four filters.



After clicking **Apply Filter**, the Filtering panel closes, applies your filters, and shows the number of filters applied (3).

View active events from the notification bell in Active IQ Unified Manager

The notification bell ( ) in the menu bar provides a fast way to view the most important active events that Unified Manager is tracking.

The list of active events provides a way to see the total number of critical, error, warning, and upgrade events on all clusters. This list includes events from the previous 7 days, and it does not include Information events. You can click a link to display the list of events that you are most interested in.

Note that when a cluster is not reachable, Unified Manager displays this information in this page. You can view detailed information about a cluster that is unreachable by clicking the **Details** button. This action opens the Event details page. Scale monitoring issues, such as low space or RAM on the management station, are also displayed on this page.

Steps

1. From the menu bar, click .
2. To view details for any of the active events, click the event text link, such as "2 Capacity" or "4 Performance".

Monitor and manage clusters from the dashboard

Monitor and manage clusters from the dashboard in Active IQ Unified Manager

The dashboard provides cumulative at-a-glance information about the current health of your monitored ONTAP systems. The dashboard provides “panels” that enable you to assess the overall capacity, performance, and security health of the clusters you are monitoring.

Additionally, there are certain ONTAP issues that you can fix directly from the Unified Manager user interface instead of having to use ONTAP System Manager or the ONTAP CLI.

At the top of the dashboard you can select whether the panels show information for all monitored clusters or for an individual cluster. You can start by viewing the status of all clusters and then drill down to individual clusters when you want to view detailed information.

Some of the panels listed below may not appear on the page based on your configuration.

Panels	Description
Management Actions	When Unified Manager can diagnose and determine a single resolution for an issue, those resolutions are displayed in this panel with a **Fix It** button.
Capacity	Displays the total and used capacity for the local tier and cloud tier, and the number of days until local capacity reaches the upper limit.
Performance Capacity	Displays the performance capacity value for each cluster and the number of days until performance capacity reaches the upper limit.
Workload IOPS	Displays the total number workloads that are currently running in a certain range of IOPS.
Workload Performance	Displays the total number of conforming and non-conforming workloads that are assigned to each defined Performance Service Level.

Panels	Description
Security	Displays the number of clusters that are compliant or not compliant, the number of SVMs that are compliant or not compliant, and the number of volumes that are encrypted or not encrypted.
Protection	Displays the number of Storage VMs that are protected by SVM-DR relationship, volumes protected by SnapMirror relationship, volumes protected by Snapshot, and clusters protected by MetroCluster.
Usage Overview	Displays clusters sorted by highest IOPS, highest throughput (MBps), or highest used physical capacity.

Learn about the Active IQ Unified Manager Dashboard page

The Dashboard page has "panels" that display the high level capacity, performance, and security health of the clusters you are monitoring. This page also provides a Management Actions panel that lists fixes that Unified Manager can make to resolve certain events.

Most of the panels also display the number of active events in that category, and the number of new events added over the previous 24 hours. This information helps you decide which clusters you may need to analyze further to resolve events. Clicking on the events displays the top events and provides a link to the Event Management inventory page filtered to show the active events in that category.

At the top of the dashboard you can select whether the panels show information for all monitored clusters ("All Clusters") or for an individual cluster. You can start by viewing the status of all clusters and then drill down to individual clusters when you want to view detailed information.

Some of the panels listed below appear on the dashboard based on your configuration.

Management Actions panel

There are certain issues that Unified Manager can diagnose thoroughly and provide a single resolution. When available, those resolutions are displayed in this panel with a **Fix It** or **Fix All** button. You can fix these issues immediately from Unified Manager instead of having to use ONTAP System Manager or the ONTAP CLI. For viewing all the issues, click on
See Fix ONTAP issues directly from Unified Manager for more information.

Capacity panel

When viewing all clusters, this panel displays the physical used capacity (after applying storage efficiency savings) and physical available capacity (not including potential storage efficiency savings) for each cluster, the number of days until the disks are projected to be full, and the data reduction ratio (without Snapshot copies) based on configured ONTAP storage efficiency settings. It also lists the used capacity for any configured cloud tiers. Clicking the bar chart takes you to the Aggregates inventory page for that cluster. Clicking the "Days To Full" text displays a message that identifies the aggregate with the least number of capacity days remaining; click the aggregate name to see more details.

When viewing a single cluster, this panel displays the physical used capacity and physical available capacity for the data aggregates sorted by each individual disk type on the local tier, and for the cloud tier. Clicking the bar chart for a disk type takes you to the Volumes inventory page for the volumes using that disk type.

Performance Capacity panel

When viewing all clusters, this panel displays the performance capacity value for each cluster (averaged over the previous 1 hour) and the number of days until performance capacity reaches the upper limit (based on daily growth rate). Clicking the bar chart takes you to the Nodes inventory page for that cluster. Note that the Nodes inventory page displays the performance capacity averaged over the previous 72 hours. Clicking the "Days To Full" text displays a message that identifies the node with the least number of performance capacity days remaining; click the node name to see more details.

When viewing a single cluster, this panel displays the cluster performance capacity used percentage, total IOPS, and total throughput (MB/s) values, and the number of days until each of these three metrics are anticipated to reach their upper limit.

Workload IOPS panel

When viewing a single cluster, this panel displays the total number workloads that are currently running in a certain range of IOPS, and indicates the number for each disk type when you hover your cursor over the chart.

Workload Performance panel

This panel displays the total number of conforming and non-conforming workloads that are assigned to each Performance Service Level (PSL) policy. It also displays the number of workloads that are not assigned a PSL. Clicking a bar chart takes you to the conforming workloads assigned to that policy in the Workloads page. Clicking the number that follows the bar chart takes you to the conforming and non-conforming workloads assigned to that policy.

Security panel

The Security panel presents the high-level security status for all clusters or a single cluster, depending on your current view. This panel displays:

- a list of the security events received in the past 24 hours. Click an event to view the details on the Event details page
- the cluster security status (count of compliant and non-compliant clusters)
- the storage VM security status (count of compliant and non compliant storage VMs)
- the volume encryption status (count of the volumes that are encrypted or not encrypted)
- the volume anti-ransomware status (count of volumes with anti-ransomware that are enabled or disabled)

You can click the bar charts of the compliant and non-compliant clusters, storage VMS, encrypted and unencrypted volumes, and volume anti-ransomware status to go to the respective pages and view the security details of the filtered clusters, storage VMs, and volumes.

Compliance is based on the NetApp Security Hardening Guide for ONTAP 9. Click the right-arrow at the top of the panel to view security details for all clusters on the Security page. For information, see Viewing detailed security status for clusters and Storage VMs.

Data Protection panel

This panel displays the data protection summary for a single or all the clusters in a data center. It displays the total number of data protection events, MetroCluster events, and number of active events raised in the last 24 hours in ONTAP. Clicking the link from each of these events takes you to the Event details page. You can click the **View All** link to view all active protection events in the Event Management inventory page. The panel displays:

- The number of volumes in a cluster or all the clusters in a data center protected by Snapshot copies.
- The number of volumes in a cluster or all the clusters in a data center protected by SnapMirror relationships. For SnapMirror relationships, the volume count at the source cluster is considered.
- The number of clusters or all the clusters in a data center protected by MetroCluster configuration over IP or FC.
- The number of volume relationships with SnapMirror recovery point objective (RPO) lag based on the lag status.

You can hover your mouse to view the respective counts and legends. You can click the right arrow at the top of the panel to view the details for a single or all the clusters on the Data Protection page. Also, you can click:

- The bar charts for unprotected volumes and volumes protected by Snapshot copies to go to the Volumes page and view the details.
- The bar charts for the clusters protected or not protected by MetroCluster configuration to go to the Clusters page and view the details.
- The bar charts for all relationships to go to the Relationships page, where the details are filtered based on the source cluster.

For more information, see Viewing volume protection status.

Usage Overview panel

When viewing all clusters, you can choose to view clusters sorted by highest IOPS, highest throughput (MB/s), or highest used physical capacity.

When viewing a single cluster, you can choose to view workloads sorted by highest IOPS, highest throughput (MB/s), or highest used logical capacity.

Related information

Fixing issues using Unified Manager automatic remediations

Displaying information about performance events

Managing performance using performance capacity and available IOPS information

Volume / Health details page

Performance event analysis and notification

Description of event severity types

Sources of performance events

Manage cluster security objectives

Monitor cluster performance from the Performance Cluster Landing page

Monitor performance using the Performance Inventory pages

Fix or manage ONTAP issues directly from Unified Manager

Manage ONTAP issues or features directly from Active IQ Unified Manager

You can fix certain ONTAP issues or manage certain ONTAP features directly from the Unified Manager user interface, instead of having to use ONTAP System Manager or the ONTAP CLI. The “Management Actions” option provides fixes to a number of ONTAP issues that have triggered Unified Manager events.

You can fix issues directly from the Management Actions page by selecting the **Management Actions** option on the left navigation pane. Management Actions are also available from the Management Actions panel on the Dashboard, Event details page, and Workload Analysis selection on the left-navigation menu.

There are certain issues that Unified Manager can diagnose thoroughly and provide a single resolution. For certain ONTAP features, such as anti-ransomware monitoring, Unified Manager performs internal checks and recommends specific actions. When available, those resolutions are displayed in Management Actions with a **Fix It** button. Click the **Fix It** button to fix the issue. You must have the Application Administrator or Storage Administrator role.

Unified Manager sends ONTAP commands to the cluster to make the requested fix. When the fix is complete the event is obsoleted.

Some management actions enable you to fix the same issue on multiple storage objects using the **Fix All** button. For example, there may be 5 volumes that have the "Volume Space Full" event that could be resolved by clicking the **Fix All** management action for "Enable volume autogrow". One click enables you to fix this issue on 5 volumes.

For information about the ONTAP issues and features that you can manage by using automatic remediation, see What issues can Unified Manager fix.

Learn about the Fix It and Fix All button options in Active IQ Unified Manager

The Management Actions page provides you with the **Fix It** or **Fix All** button to fix issues that Unified Manager has been notified about through an event.

We recommend that you click the buttons to fix an issue, as required. However, if you are not sure that you want to resolve the issue as recommended by Unified Manager, you can perform the following actions:

What do you want to do?	Action
Have Unified Manager fix the issue on all identified objects.	Click the **Fix All** button.
Do not fix the issue for any of the identified objects at this time and hide this management action until the event is raised again.	Click the down arrow and click **Dismiss All**.
Fix the issue on only some of the identified objects.	Click the name of the management action to expand the list and show all individual **Fix It** actions.Then follow the steps for fixing or dismissing individual management actions.

What do you want to do?	Action
Have Unified Manager fix the issue.	Click the **Fix It** button.
Do not fix the issue at this time and hide this management action until the event is raised again.	Click the down arrow and click **Dismiss**.
Display the details for this event so you can better understand the issue.	• Click the **Fix It** button and review the fix that will be applied in the resulting dialog box. • Click the down arrow and click **View Event Detail** to display the Event details page. Then click **Fix It** from either of these resulting pages if you want to fix the issue.
Display the details for this storage object so you can better understand the issue.	Click the name of the storage object to display details in either the Performance Explorer or Health Details page.

In some cases the fix is reflected in the next 15 minute configuration poll. In other cases it can take up to many hours for the configuration change to be verified and for the event to be obsoleted.

To see the list of completed or in progress management actions, click the filter icon and select **Completed** or **In Progress**.

Fix All operations run in a serial fashion, so when you view the **In Progress** panel some objects will have the Status **In Progress** whereas others will have the Status **Scheduled**; meaning they are still waiting to be implemented.

View the status of management actions to fix in Active IQ Unified Manager

You can view the status of all management actions that you have chosen to fix in the Management Actions page. Most actions are shown as **Completed** fairly quickly after Unified Manager sends the ONTAP command to the cluster. However, some actions, such as moving a volume, can take longer.

There are three filters available on the Management Actions page:

- **Completed** shows both management actions that completed successfully and those that have failed. **Failed** actions provide a reason for the failure so that you can address the issue manually.
- **In Progress** shows both those management actions that are being implemented, and those that are scheduled to be implemented.
- **Recommended** shows all the management actions that are currently active for all monitored clusters.

Steps

1. Click **Management Actions** on the left navigation pane. Alternately, click ⋮ at the top of the **Management Actions** panel on the **Dashboard** and select the View you want to see.

   The Management Actions page is displayed.

2. You can click the caret icon next to the management action in the **Description** field to see details about the issue and the command that is being used to fix the issue.
3. To view any actions that **failed**, sort on the **Status** column in the **Completed** View. You can use the **Filter**

tool for this same purpose.

4. If you want to view more information about a Failed management action, or if you decide that you want to fix a Recommended management action, you can click **View Event Detail** from the expanded area after you click the caret icon next to the management action. A **Fix It** button is available from that page.

Learn about ONTAP issues that Active IQ Unified Manager can fix

By using the automatic remediation feature of Active IQ Unified Manager, you can choose to resolve certain ONTAP issues or manage certain ONTAP features, such as anti-ransomware monitoring, effectively through Unified Manager.

This table describes these ONTAP issues or features that you can manage directly through the **Fix It** or **Fix All** button on the Unified Manager web UI.

Event Name and Description	Management Action	"Fix It" Operation
Volume Space Full The volume is almost out of space and it has breached the capacity full threshold. This threshold is set by default to 90% of the volume size.	Enable volume autogrow	Unified Manager determines that volume autogrow is not configured for this volume, so it enables this feature so the volume will grow in capacity when needed.
Inodes Full This volume has run out of inodes and cannot accept any new files.	Increase number of inodes on volume	Increases the number of inodes on the volume by 2 percent.
Storage Tier Policy Mismatch Detected The volume has lots of inactive data and the current tiering policy is set to "snapshot-only" or "none".	Enable automatic cloud tiering	Since the volume already resides on a FabricPool, it changes the tiering policy to "auto" so that inactive data is moved to the lower cost cloud tier.
Storage Tier Mismatch Detected The volume has lots of inactive data, but it does not reside on a cloud-enabled storage tier (FabricPool).	Change volumes' storage tier	Moves the volume to cloud-enabled storage tier and sets the tiering policy to "auto" to move inactive data to the cloud tier.
Audit Log Disabled The audit log is not enabled for the storage VM	Enable audit logging for the storage VM	Enables audit logging on the storage VM. Note that the storage VM must already have either a local or remote audit log location configured.

Event Name and Description	Management Action	"Fix It" Operation
Login Banner Disabled The login banner for the cluster should be enabled to increase security by making access restrictions clear.	Set login banner for the cluster	Sets the cluster login banner to “Access restricted to authorized users”.
Login Banner Disabled The login banner for the storage VM should be enabled to increase security by making access restrictions clear.	Set login banner for the storage VM	Sets the storage VM login banner to "Access restricted to authorized users".
SSH is Using Insecure Ciphers Ciphers with the suffix "-cbc" are considered insecure.	Remove insecure ciphers from the cluster	Removes the insecure ciphers — such as aes192-cbc and aes128-cbc — from the cluster.
SSH is Using Insecure Ciphers Ciphers with the suffix "-cbc" are considered insecure.	Remove insecure ciphers from the storage VM	Removes the insecure ciphers — such as aes192-cbc and aes128-cbc — from the storage VM.
AutoSupport HTTPS transport disabled The transport protocol used to send AutoSupport messages to technical support should be encrypted.	Set HTTPS as the transport protocol for AutoSupport messages	Sets HTTPS as the transport protocol for AutoSupport messages on the cluster.
Cluster Load Imbalance Threshold Breached Indicates that the load is imbalanced among the nodes in the cluster. This event is generated when the performance capacity used variance is more than 30% between nodes.	Balance cluster workloads	Unified Manager identifies the best volume to move from one node to the other to reduce the imbalance, and then moves the volume.
Cluster Capacity Imbalance Threshold Breached Indicates that the capacity is imbalanced among the aggregates in the cluster. This event is generated when the used capacity variance is more than 70% between aggregates.	Balance cluster capacity	Unified Manager identifies the best volume to move from one aggregate to another to reduce the imbalance, and then moves the volume.

Event Name and Description	Management Action	"Fix It" Operation
Performance Capacity Used Threshold Breached Indicates that the load on the node could become over utilized if you don’t reduce the utilization by one or more highly active workloads. This event is generated when the node performance capacity used value is more than 100% for more than 12 hours.	Limit high load on node	Unified Manager identifies the volume with the highest IOPS and it applies a QoS policy using the historical expected and peak IOPS levels to reduce the load on the node.
Dynamic Event Warning Threshold Breached Indicates that the node is already operating in an overloaded state due to the abnormally high load of some of the workloads.	Reduce overload in node	Unified Manager identifies the volume with the highest IOPS and it applies a QoS policy using the historical expected and peak IOPS levels to reduce the load on the node.
Takeover is not possible Failover is currently disabled, so access to the node’s resources during an outage or reboot would be lost until the node became available again.	Enable node failover	Unified Manager sends the appropriate command to enable failover on all nodes in the cluster.
Option Cf.takeover.on_panic is Configured OFF The nodeshell option "cf.takeover.on_panic" is set to **off**, which could cause an issue on HA-configured systems.	Enable takeover on panic	Unified Manager sends the appropriate command to the cluster to change this setting to **on**.
Disable nodeshell option snapmirror.enable The old nodeshell option "snapmirror.enable" is set to **on**, which could cause an issue during boot after upgrading to ONTAP 9.3 or greater.	Set snapmirror.enable option to off	Unified Manager sends the appropriate command to the cluster to change this setting to **off**.
Telnet enabled Indicates a potential security issue because Telnet is insecure and passes data in an unencrypted manner.	Disable Telnet	Unified Manager sends the appropriate command to the cluster to disable Telnet.

Event Name and Description	Management Action	"Fix It" Operation
Configure storage VM anti-ransomware learning Periodically checks for clusters with licenses for anti-ransomware monitoring. Validates whether a storage VM supports only NFS or SMB volumes in such a cluster.	Put storage VMs in a `learning` mode of anti-ransomware monitoring	Unified Manager sets anti-ransomware monitoring to `learning` state for the storage VMs through the cluster management console. Anti-ransomware monitoring on all the new volumes created on the storage VM are automatically moved to a learning mode. Through this enablement, ONTAP can learn the pattern of activity on the volumes and detect the anomalies due to potential malicious attacks.
Configure volume anti-ransomware learning Periodically checks for clusters with licenses for anti-ransomware monitoring. Validates whether a volume supports only NFS or SMB services in such a cluster.	Put volumes in `learning` mode of anti-ransomware monitoring	Unified Manager sets anti-ransomware monitoring to `learning` state for the volumes through the cluster management console. Through this enablement, ONTAP can learn the pattern of activity on the volumes and detect the anomalies due to potential malicious attacks.
Enable volume anti-ransomware Periodically checks for clusters with licenses for anti-ransomware monitoring. Detects whether the volumes are in the `learning` mode of anti-ransomware monitoring for more than 45 days, and determines the prospect of putting them in active mode.	Put volumes in `active` mode of anti-ransomware monitoring	Unified Manager sets anti-ransomware monitoring to `active` on the volumes through the cluster management console. Through this enablement, ONTAP can learn the pattern of activity on the volumes and detect the anomalies due to potential malicious attacks, and create alerts for data protection actions.
Disable volume anti-ransomware Periodically checks for clusters with licenses for anti-ransomware monitoring. Detects repetitious notifications during active anti-ransomware monitoring on the volumes (for example, multiple warnings of potential ransomware attacks are returned over 30 days).	Disable anti-ransomware monitoring on volumes	Unified Manager disables anti-ransomware monitoring on the volumes through the cluster management console.

Override management actions with scripts in Active IQ Unified Manager

You can create custom scripts and associate them to alerts to take specific actions for specific events, and not opt for the default management actions available for them on the Management Actions page or Unified Manager dashboard.

If you want to take specific actions for an event type and choose not to fix them as a part of the management action capability provided by Unified Manager, you can configure a custom script for the specific action. You can then associate the script with an alert for that event type and take care of such events individually. In this case, management actions are not generated for that specific event type on the Management Actions page or Unified Manager dashboard.

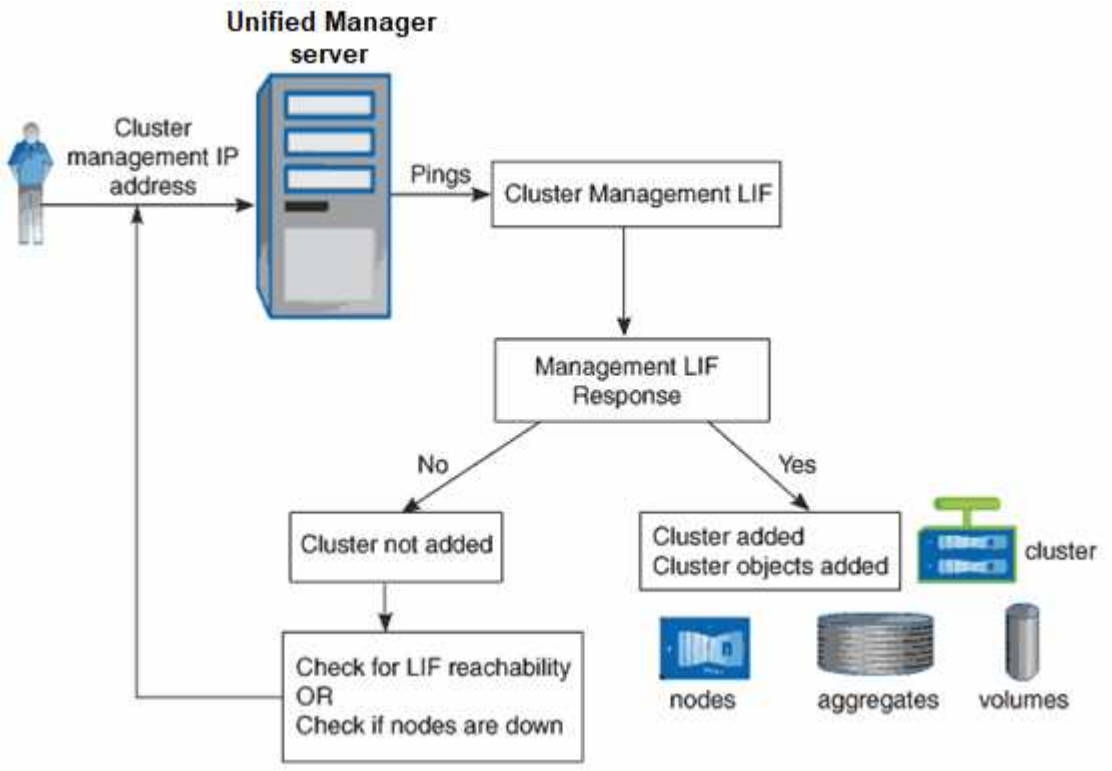
Manage clusters

Learn how the cluster discovery process works in Active IQ Unified Manager

After you have added a cluster to Unified Manager, the server discovers the cluster objects and adds them to its database. Understanding how the discovery process works helps you to manage your organization's clusters and their objects.

The monitoring interval for collecting cluster configuration information is 15 minutes. For example, after you have added a cluster, it takes 15 minutes to display the cluster objects in the Unified Manager UI. This time frame is also true when making changes to a cluster. For example, if you add two new volumes to an SVM in a cluster, you see those new objects in the UI after the next polling interval, which could be up to 15 minutes.

The following image illustrates the discovery process:



After all the objects for a new cluster are discovered, Unified Manager starts to gather historical performance data for the previous 15 days. These statistics are collected using the data continuity collection functionality. This feature provides you with over two weeks of performance information for a cluster immediately after it is added. After the data continuity collection cycle is completed, real-time cluster performance data is collected, by default, every five minutes.

Because the collection of 15 days of performance data is CPU intensive, it is suggested that you stagger the addition of new clusters so that data continuity collection polls are not running on too

many clusters at the same time.

View the list of clusters monitored by Active IQ Unified Manager

You can use the Cluster Setup page to view your inventory of clusters. You can view details about the clusters, such as their name or IP address and communication status.

Before you begin

You must have the Operator, Application Administrator, or Storage Administrator role.

Step

1. In the left navigation pane, click **Storage Management** > **Cluster Setup**.

   All the clusters in your storage environment managed by Unified Manager are displayed.The list of clusters is sorted by the collection state severity level column. You can click a column header to sort the clusters by different columns.

Add clusters to Active IQ Unified Manager

You can add a cluster to Active IQ Unified Manager so that you can monitor the cluster. This includes the ability to obtain cluster information such as the health, capacity, performance, and configuration of the cluster so that you can find and resolve any issues that might occur.

Before you begin

- You must have the Application Administrator role or the Storage Administrator role.
- You must have the following information:
  - Unified Manager supports on-premises ONTAP clusters, ONTAP Select, Cloud Volumes ONTAP.
  - You must have the host name or cluster management IP address (IPv4 or IPv6) for the cluster.

    When using the host name, it must resolve to the cluster management IP address for the cluster management LIF. If you use a node management LIF, the operation fails.

  - You must have the user name and password to access the cluster.

    This account must have the *admin* role with Application access set to *ontapi*, *amqp*, *console*, and *http*.

  - You must know the port number to connect to the cluster using the HTTPS protocol (typically port 443).
  - The cluster must be running ONTAP version 9.9 software or greater.
  - You must have adequate space on the Unified Manager server. You are prevented from adding a cluster to the server when greater than 90% of space is already consumed.
  - You have the required certificates:

    SSL (HTTPS) certificate: This certificate is owned by Unified Manager. A default self-signed SSL (HTTPS) certificate is generated with a fresh installation of Unified Manager. NetApp recommends that you upgrade it to CA-signed certificate for better security. If the server certificate expires, you should regenerate it and restart Unified Manager for the services to incorporate the new certificate. For more information about regenerating SSL certificate, see Generating an HTTPS security certificate.

EMS certificate: This certificate is owned by Unified Manager. It is used during authentication for EMS notifications received from ONTAP.

Certificates for Mutual TLS communication: Used during Mutual TLS communication between Unified Manager and ONTAP. The certificate-based authentication is enabled for a cluster, based on the ONTAP version. If the cluster running the ONTAP version is lower than the 9.5, certificate-based authentication is not enabled.

Certificate-based authentication is not enabled automatically for a cluster, if you are updating an older version of Unified Manager. However, you can enable it by modifying and saving the cluster details. If the certificate expires, you should regenerate it to incorporate the new certificate. For more information about viewing and regenerating the certificate, see Editing clusters.

- You can add a cluster from the web UI, and certificate-based authentication is automatically enabled.
- You can add a cluster through Unified Manager CLI, the certificate-based authentication is not enabled by default. If you add a cluster using Unified Manager CLI, it is required to edit the cluster using Unified Manager UI. You can see Supported Unified Manager CLI commands to add a cluster using Unified Manager CLI.
- If certificate-based authentication is enabled for a cluster, and you take the backup of Unified Manager from a server and restore to another Unified Manager server where hostname or IP address is changed, then monitoring of the cluster can fail. To avoid the failure, edit and save the cluster details. For more information about editing cluster details, see Editing clusters.
- On the cluster level, the Active IQ interface adds two new user group entries for the authentication method "cert".

Cluster certificates: This certificate is owned by ONTAP. You cannot add a cluster to Unified Manager with an expired certificate and if the certificate has already expired, you should regenerate it before adding the cluster. For information about certificate generation, see the knowledge base (KB) article How to renew an ONTAP self-signed certificate in System Manager user interface.

- A single instance of Unified Manager can support a specific number of nodes. If you need to monitor an environment that exceeds the supported node count, you must install an additional instance of Unified Manager to monitor some of the clusters. To view the list of supported node count, see the Unified Manager Best Practices Guide.

Steps

1. In the left navigation pane, click **Storage Management** > **Cluster Setup**.
2. On the Cluster Setup page, click **Add**.
3. In the Add Cluster dialog box, specify the values as required, and then click **Submit**.
4. In the Authorize Host dialog box, click **View Certificate** to view the certificate information about the cluster.
5. Click **Yes**.

   After saving the cluster details, you can see the certificate for Mutual TLS communication for a cluster.

   If the certificate-based authentication is not enabled, Unified Manager checks the certificate only when the cluster is added initially. Unified Manager does not check the certificate for each API call to ONTAP.

After all the objects for a new cluster are discovered, Unified Manager starts to gather historical performance

data for the previous 15 days. These statistics are collected using the data continuity collection functionality. This feature provides you with over two weeks of performance information for a cluster immediately after it is added. After the data continuity collection cycle is completed, real-time cluster performance data is collected, by default, every five minutes.

- The collection of 15 days of performance data is CPU intensive. NetApp recommends that you stagger the addition of new clusters so that data continuity collection polls are not running on too many clusters at the same time. Additionally, if you restart Unified Manager during the data continuity collection period, the collection is halted, and you will see gaps in the performance charts for the missing timeframe.
- When you add an ONTAP cluster with version 9.14.1 or later, communication occurs through the cloud agent feature. ONTAP automatically creates an internal service account user with read-only privileges (for example, clus-agent-xxxx). Unified Manager stops collecting data from those clusters if the `clus-agent` user is deleted.

If you receive an error message that you cannot add the cluster, check to see if the following issues exist:

- If the clocks on the two systems are not synchronized and the Unified Manager HTTPS certificate start date is later than the date on the cluster. You must ensure that the clocks are synchronized using NTP or a similar service.
- If the cluster has reached the maximum number of EMS notification destinations the Unified Manager address cannot be added. By default only 20 EMS notification destinations can be defined on the cluster.

Related information

Adding users

Viewing the cluster list and details

Installing a CA signed and returned HTTPS certificate

Edit cluster settings in Active IQ Unified Manager

You can modify the settings of an existing cluster, such as the host name or IP address, user name, password, and port, by using the Edit Cluster dialog box.

Before you begin

You must have the Application Administrator role or the Storage Administrator role.

Starting with Unified Manager 9.7, clusters can be added using HTTPS only.

Steps

1. In the left navigation pane, click **Storage Management** > **Cluster Setup**.
2. On the **Cluster Setup** page, select the cluster you want to edit, and then click **Edit**.
3. In the **Edit Cluster** dialog box, modify the values as required.

   If you have modified the details for a cluster added to Unified Manager, you can view the certificate details for Mutual TLS communication, based on the ONTAP version. For more information about ONTAP version,

see Certificates for Mutual TLS communication.

You can view the certificate details by clicking **Certificate Details**. If the certificate is expired, click the **Regenerate** button to incorporate the new certificate.

4. Click **Submit**.
5. In the Authorize Host dialog box, click **View Certificate** to view the certificate information about the cluster.
6. Click **Yes**.

Related information

Adding users

Viewing the cluster list and details

Remove clusters from Active IQ Unified Manager

You can remove a cluster from Unified Manager by using the Cluster Setup page. For example, you can remove a cluster if cluster discovery fails, or when you want to decommission a storage system.

Before you begin

You must have the Application Administrator role or the Storage Administrator role.

This task removes the selected cluster from Unified Manager. After a cluster is removed, it is no longer monitored. The instance of Unified Manager registered with the removed cluster is also unregistered from the cluster.

Removing a cluster also deletes all its storage objects, historical data, storage services, and all associated events from Unified Manager. These changes are reflected on the inventory pages and the details pages after the next data collection cycle.

> You must remove a cluster from Unified Manager before decommissioning it. If the cluster is already offline or unreachable when you attempt to remove it, Unified Manager cannot communicate with the cluster to remove the application record and configured Event Management System (EMS) subscriptions. This will cause the removal process to fail or become unresponsive before completion.

Steps

1. In the left navigation pane, click **Storage Management** > **Cluster Setup**.
2. On the Cluster Setup page, select the cluster that you want to remove and click **Remove**.
3. In the **Remove Data Source** message dialog, click **Remove** to confirm the remove request.

Related information

Adding users

Viewing the cluster list and details

Rediscover clusters in Active IQ Unified Manager

You can manually rediscover a cluster from the Cluster Setup page in order to obtain the latest information about the health, monitoring status, and performance status of the cluster.

You can manually rediscover a cluster when you want to update the cluster—such as by increasing the size of an aggregate when there is insufficient space—and you want Unified Manager to discover the changes that you make.

When Unified Manager is paired with OnCommand Workflow Automation (WFA), the pairing triggers the reacquisition of the data cached by WFA.

Steps

1. In the left navigation pane, click **Storage Management** > **Cluster Setup**.
2. On the **Cluster Setup** page, click **Rediscover**.

   Unified Manager rediscovers the selected cluster and displays the latest health and performance status.

Related information

Viewing the cluster list and details

Monitor VMware virtual infrastructure

Understand how Active IQ Unified Manager monitors VMware virtual infrastructure

Active IQ Unified Manager provides visibility into the virtual machines (VMs) in your virtual infrastructure, and enables monitoring and troubleshooting storage and performance issues in your virtual environment. You can use this feature to determine any latency issues in your storage environment or when there is a reported performance event on your vCenter Server.

A typical virtual infrastructure deployment on ONTAP has various components that are spread across compute, network, and storage layers. Any performance lag in a VM application might occur due to a combination of latencies faced by the various components at the respective layers. This feature is useful for storage and vCenter Server admins and IT generalists who need to analyze a performance issue in a virtual environment and understand in which component the issue has occurred.

You can now access the vCenter Server from the vCenter menu of the VMware section. The peek view of each virtual machine listed has the **VCENTER SERVER** link in the TOPOLOGY VIEW that launches the vCenter Server in a new browser. You can also use the **Expand Topology** button to launch the vCenter Server and click **View in vCenter** button to view the datastores in vCenter Server.

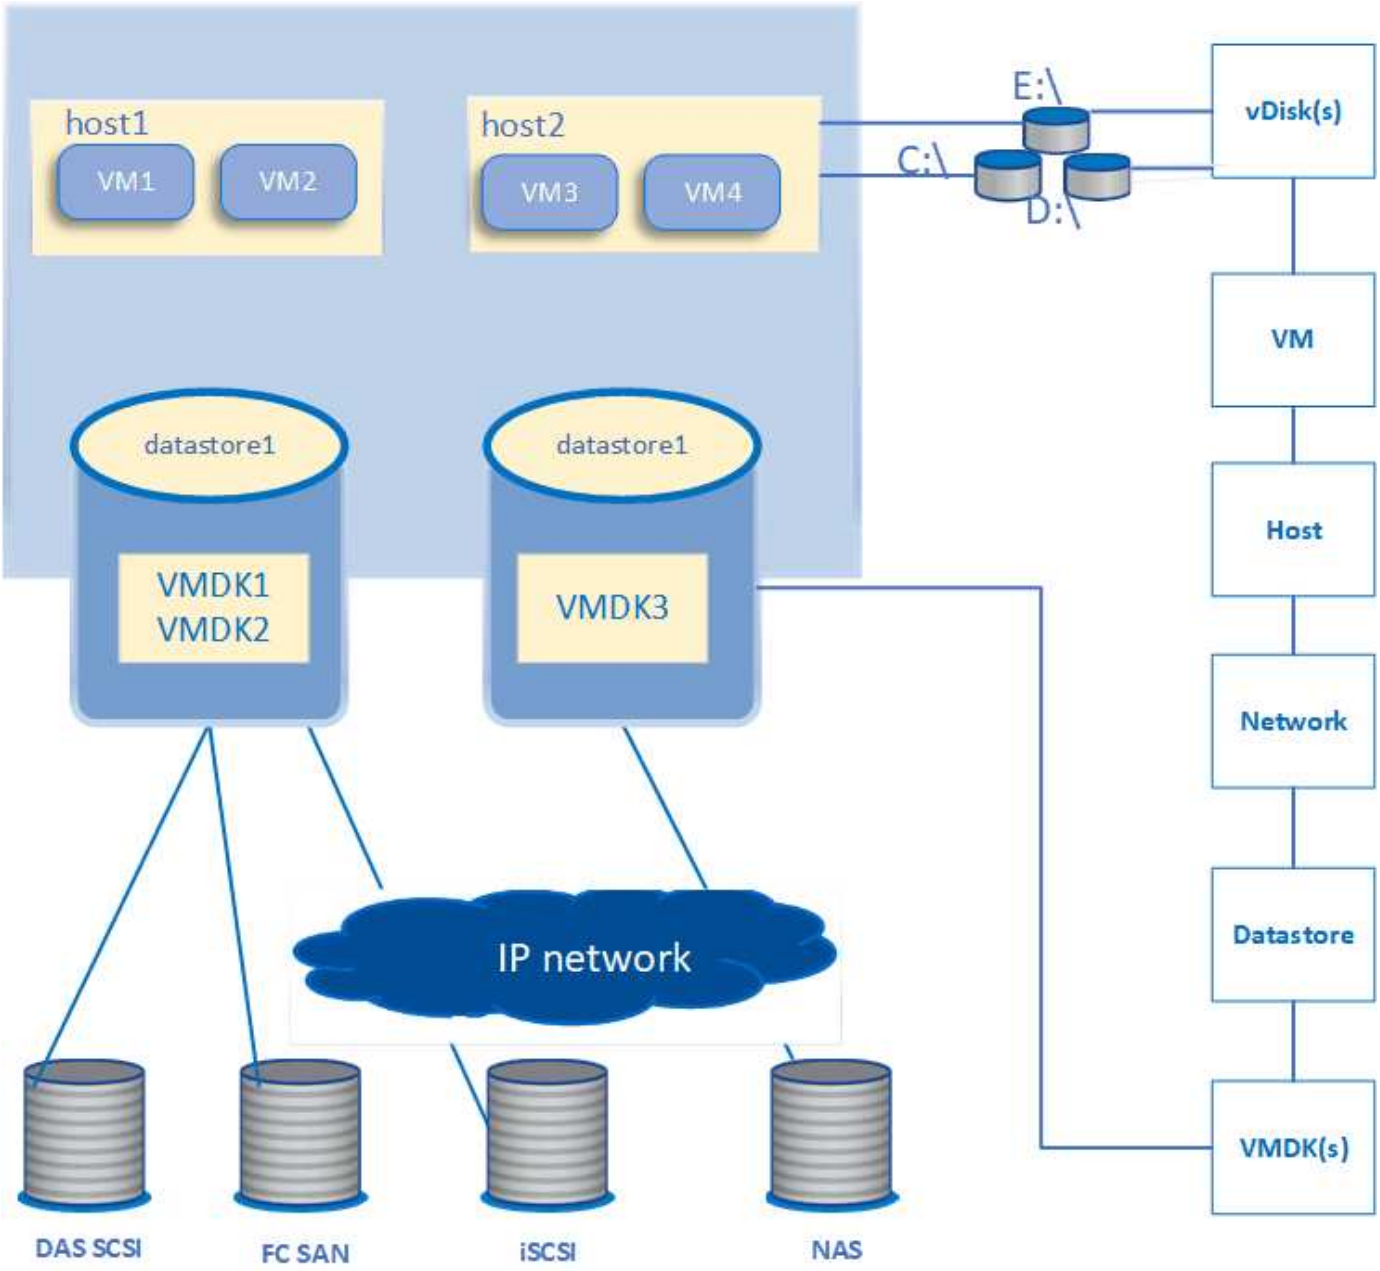
Unified Manager presents the underlying sub-system of a virtual environment in a topological view for determining whether a latency issue has occurred in the compute node, network, or storage. The view also highlights the specific object that causes the performance lag for taking remedial steps and addressing the underlying issue.

A virtual infrastructure deployed on ONTAP storage includes the following objects:

- vCenter Server: A centralized control plane for managing the VMware VMs, ESXi hosts, and all related components in a virtual environment. For more information about vCenter Server, see VMware documentation.
- Host: A physical or virtual system that runs ESXi, the virtualization software from VMware, and hosts the VM.
- Datastore: Datastores are virtual storage objects that are connected to the ESXi hosts. Datastores are manageable storage entities of ONTAP, such as LUNs or volumes, used as a repository for VM files, such as log files, scripts, configuration files, and virtual disks. They are connected to the hosts in the environment via a SAN or IP network connection. Datastores outside of ONTAP that are mapped to vCenter Server are not supported or displayed on Unified Manager.
- VM: A VMware virtual machine.
- Virtual disks: The virtual disks on datastores belonging to the VMs that have an extension as VMDK. Data from a virtual disk is stored on the corresponding VMDK.
- VMDK: A virtual machine disk on the datastore that provides storage space for virtual disks. For each virtual disk, there is a corresponding VMDK.

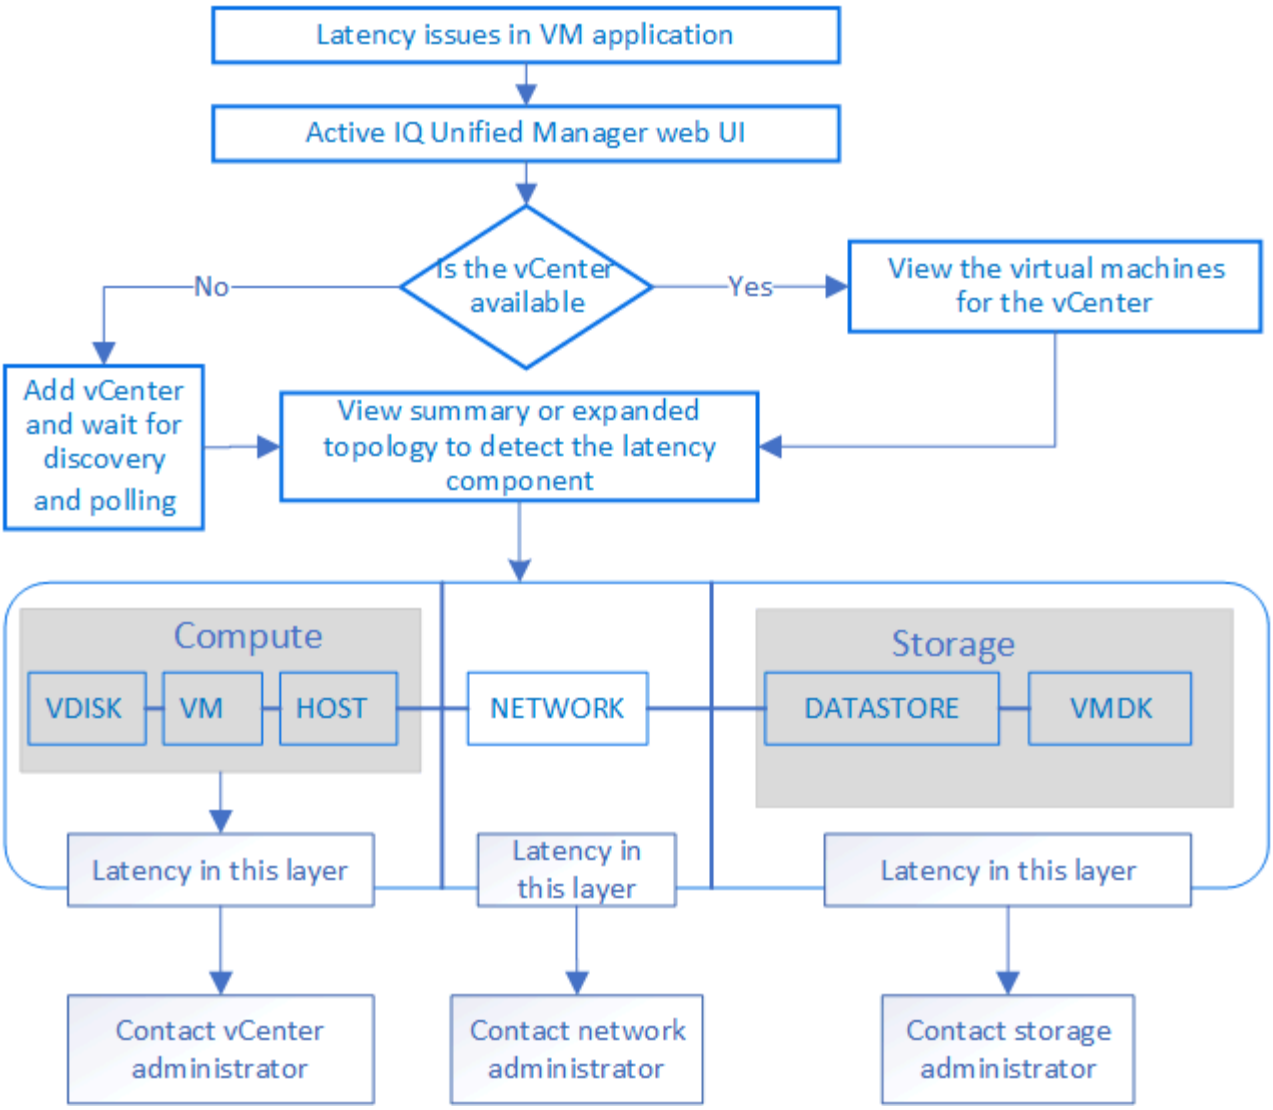
These objects are represented in a VM topology view.

VMware virtualization on ONTAP



User workflow

The following diagram displays a typical use case of using the VM topology view:



What is not supported

- Datastores that are outside of ONTAP and are mapped to the vCenter Server instances are not supported on Unified Manager. Any VMs with virtual disks on those datastores are also not supported.
- A datastore that spans across multiple LUNs is not supported.
- Datastores using Network address translation (NAT) for mapping data LIF (access endpoint) are not supported.
- Exporting volumes or LUNs as datastores on different clusters with same IP addresses in a multiple-LIFs configuration is not supported as Unified Manager is unable to identify which datastore belongs to which cluster.

  Example: Suppose cluster A has datastore A. Datastore A is exported via data LIF with same IP address x.x.x.x and VM A is created on this datastore. Similarly, cluster B has datastore B. The datastore B is exported via data LIF with same IP address x.x.x.x and VM B is created on the datastore B. UM will neither be able to map the datastore A for VM A's topology to corresponding ONTAP volume/LUN nor map VM B.

- Only NAS and SAN volumes (iSCSI and FCP for VMFS) are supported as datastores, virtual volumes (vVols) are not supported.
- Only iSCSI virtual disks are supported. Virtual disks of NVMe and SATA types are not supported.

- The views do not allow you to generate reports for analysing the performance of the various components.
- For the storage virtual machine (storage VM) disaster recovery (DR) setup that is supported for only virtual infrastructure on Unified Manager, the configuration has to be manually changed in vCenter Server to point to the active LUNs in switchover and switchback scenarios. Without a manual intervention, their datastores become inaccessible.

View and add vCenter Servers in Active IQ Unified Manager

For viewing and troubleshooting the performance of the virtual machines (VMs), the associated vCenter Servers must be added on your Active IQ Unified Manager instance.

Before you begin

Before adding or viewing vCenter Servers, ensure the following:

- You are aware of the vCenter Server names.
- You know the IP address of vCenter Server and have the required credentials. The credentials must be of a vCenter Server administrator or a root user with read-only access to vCenter Server.
- The vCenter Server that you want to add runs vSphere 6.5 or later.

  The support of Unified Manager for VMware ESXi and vCenter Server is available in English and Japanese languages.

- The data collection setting in vCenter Server is set to the statistics level of *`Level 3`*, ensuring the required level of metrics collection for all the monitored objects. The interval duration should be *`5 minutes`*, and the save period should be *`1 day`*.

  For more information, see “Data Collection Levels” section of *vSphere Monitoring and Performance Guide* in VMware documentation.

- The latency values in vCenter Server are configured in milliseconds, and not in microseconds, for successful calculations of the latency values.
- While adding the datastore to vCenter Server, you can use both the IP address of the host or the fully qualified domain name (FQDN). In case you are adding FQDN, ensure that the domain name can be resolved by the Unified Manager server. For example, for a Linux installation, ensure that the domain name is added in the `/etc/resolv.conf` file.
- The current time of vCenter Server is in sync with the vCenter Server time zone.
- vCenter Server is reachable for a successful discovery.
- You have the read access to VMware SDK when adding the vCenter Server to Unified Manager. This is required for configuration polling.

For every vCenter Server added and discovered, Unified Manager collects the configuration data, such as the vCenter Server and ESXi server details, ONTAP mapping, datastore details, and number of VMs hosted. It further collects the performance metrics of the components.

Steps

1. Go to **VMWARE** > **vCenter**, and check whether your vCenter Server is available on the list.

   If your vCenter Server is not available, you must add vCenter Server.

a. Click **Add**.

b. Add the correct IP address for vCenter Server and ensure that the device is reachable.

c. Add the user name and password of the administrator or root user with read-only access to vCenter Server.

d. Add the custom port number if you are using any port other than the default 443.

e. Click **Save**.

Upon successful discovery, a server certificate is displayed for you to accept.

When you accept the certificate, vCenter Server is added to the list of available vCenter Servers. Adding the device does not result into data collection for the associated VMs, and the collection occurs at the scheduled intervals.

2. If your vCenter Server is available on the **vCenters** page, check its status by hovering your mouse over the **Status** field to display whether your vCenter Server is performing as expected or whether there is a warning or error.

Adding vCenter Server allows you to view the following statuses. However, the performance and latency data of the corresponding VMs might take up to an hour after adding vCenter Server to be accurately reflected.

- Green: "Normal", indicating that vCenter Server has been discovered, and performance metrics have been successfully collected
- Yellow: "Warning" (for example, when the statistics level for vCenter Server has not been set to 3 or above to obtain statistics for each object)
- Orange: "Error" (indicates any internal errors, such as exception, failure in configuration data collection, or vCenter Server being unreachable)
  You can click the column display icon (**Show/Hide**) to view the status message for a vCenter Server status and troubleshoot the issue.

3. In case vCenter Server is unreachable or the credentials have changed, edit the vCenter Server details by selecting **vCenter** > **Edit**.
4. Make the necessary changes on the **Edit VMware vCenter Server** page.
5. Click **Save**.

vCenter Server data collection begins

vCenter Server collects real-time 20-second performance data samples and rolls them up to 5-minute samples. The schedule for performance data collection of Unified Manager is based on the default settings of vCenter Server. Unified Manager processes the 5-minute samples obtained from vCenter Server and calculates an hourly average of the IOPS and latency for the virtual disks, VMs, and hosts. For datastores, Unified Manager calculates an hourly average of the IOPS and latency from samples obtained from ONTAP. These values are available at the top of the hour. The performance metrics are not available immediately after vCenter Server is added, and is available only when the next hour starts. Performance data polling begins on completing a cycle of configuration data collection.

For polling vCenter Server configuration data, Unified Manager follows the same schedule as for collecting cluster configuration data. For information about vCenter Server configuration and performance data collection schedule, see "Cluster configuration and performance data collection activity".

Related information

Remove a vCenter Server from Active IQ Unified Manager

You can remove vCenter Servers from your Active IQ Unified Manager instance. For example, you can remove a vCenter Server if vCenter Server discovery fails or when it is no longer required.

Removing a vCenter Server also deletes all the virtual machines (VMs) hosted on that vCenter and its configuration data. After the vCenter Server is removed, it will no longer be monitored, along with its associated objects and historical data. These changes will be reflected on vCenter and virtual machine inventory pages.

Before you begin

Before removing vCenter Servers, ensure the following:

- You have the Application Administrator role or the Storage Administrator role.
- You should be aware of the vCenter Server names and respective IP addresses associated with them.

Steps

1. In the left navigation pane, click **VMWARE>vCenter**.
2. On the vCenters page, select the vCenter Server that you want to remove and click **Remove**.
3. In the **Remove vCenter** message dialog, click **OK** to confirm the remove request.

Understand how Active IQ Unified Manager monitors virtual machines

For any latency issue on the virtual machine (VM) applications, you might need to monitor the VMs to analyze and troubleshoot the cause. The VMs are available when their vCenter Server and the ONTAP clusters hosting the VM storage are added to Unified Manager.

You see the details of the VMs on the **VMWARE** > > **Virtual Machines** page. Information, such as the availability, status, used and allocated capacity, network latency, and the IOPS and latency of the VM, datastore, and host is displayed. For a VM supporting multiple datastores, the grid shows the metrics of the datastore with the worst latency, with an asterisk icon (*) indicating additional datastores. If you click on the icon, the metrics of the additional datastore is displayed. Some of these columns are not available for sorting and filtering.



For viewing a VM and its details, the discovery (polling or metrics collection) of the ONTAP cluster must be complete. If the cluster is removed from Unified Manager, the VM is no longer available, after the next cycle of discovery.

From this page, you can also view the detailed topology of a VM, displaying the components to which the VM is related, for example, the host, virtual disk, and datastore connected to it. The topology view displays the underlying components in their specific layers, in the following order: **Virtual Disk** > **VM** > **Host** > **Network** > **Datastore** > **VMDK**.

You can determine the I/O path and component-level latencies from a topological aspect and identify whether storage is the cause of the performance issue. The summary view of the topology displays the I/O path, and highlights the component that has IOPS and latency issues for you to decide on the troubleshooting steps. You

can also have an expanded view of the topology that depicts each component separately along with latency of that component. You can select a component to determine the I/O path highlighted through the layers.

View summary topology

To determine performance issues by viewing the VMs in a summary topology:

1. Go to **VMWARE** > **Virtual Machines**.
2. Search for your VM by typing its name in the search box. You can also filter your search results based on specific criteria by clicking on the **Filter** button. However, If you cannot find your VM, ensure that the corresponding vCenter Server has been added and discovered.

> vCenter Servers allow special characters (such as %, &, *, $, #, @, !, \, /, :, *, ?, "`, <, >, |, ;, ') in the names of vSphere entities, such as VM, cluster, datastore, folder, or file. The VMware vCenter Server and ESX/ESXi Server do not escape special characters used in the display names. However, when the name is processed in Unified Manager, it is displayed differently. For example, a VM named as `%$VC_AIQUM_clone_191124%` in vCenter Server is displayed as `%25$VC_AIQUM_clone_191124%25` in Unified Manager. You must keep a note of this issue when you query a VM with a name having a special characters in it.

3. Check the status of the VM. The VM statuses are retrieved from vCenter Server. The following statuses are available. For more information about these statuses, see VMware documentation.
   - Normal
   - Warning
   - Alert
   - Not monitored
   - Unknown
4. Click the down arrow beside the VM to see the summary view of the topology of the components across the compute, network, and storage layers. The node that has latency issues is highlighted. The summary view displays the worst latency of the components. For example, if a VM has more than one virtual disk, this view displays the virtual disk that has the worst latency among all the virtual disks.
5. To analyze the latency and throughput of the datastore over a period of time, click the **Workload Analyzer** button on top of the datastore object icon. You go to the Workload Analysis page, where you can select a time range and view the performance charts of the datastore. For more information about workload analyzer, see *Troubleshooting workloads using the workload analyzer*.

View expanded topology

You can drill down to each component separately by viewing the expanded topology of the VM.

Steps

1. From the summary topology view, click **Expand Topology**. You can see the detailed topology of each component separately with the latency numbers for each object. If there are multiple nodes in a category, for example multiple nodes in the datastore or VMDK, the node with worst latency is highlighted in red.
2. To check the IO path of a specific object, click on that object to see the IO path and the corresponding mapping. For example, to see the mapping of a virtual disk, click the virtual disk to view its highlighted mapping to the respective VMDK. In case of a performance lag of these components, you can collect more data from ONTAP and troubleshoot the issue.

Metrics are not reported for VMDKs. In the topology, only the VMDK names are displayed, and not metrics.

Related information

Troubleshooting workloads using the workload Analyzer

View virtual infrastructure visibility in a disaster recovery setup with Active IQ Unified Manager

You can view the configuration and performance metrics of the datastores hosted in a MetroCluster configuration or storage virtual machine (storage VM) disaster recovery (SVM DR) setup.

On Unified Manager, you can view the NAS volumes or LUNs in a MetroCluster configuration that are attached as datastores in vCenter Server. The datastores hosted in a MetroCluster configuration are represented in the same topological view as a datastore in a standard environment.

You can also view the NAS volumes or LUNs in a storage VM disaster recovery configuration that are mapped to the datastores in vCenter Server.

View datastores in MetroCluster configuration

Note the following prerequisites before viewing datastores in a MetroCluster configuration:

- In an event of switchover and switchback, the discovery of the primary and secondary clusters of the HA pair, and vCenter Servers should be complete.
- The primary and secondary clusters of the HA pair, and vCenter Servers must be managed by Unified Manager.
- The required setup must be completed on ONTAP and vCenter Server. For information, see ONTAP and vCenter documentation.

  ONTAP 9 Documentation Center

Follow these steps for viewing datastores:

1. On the **VMWARE** > **Virtual Machines** page, click the VM that hosts the datastore. Click the **Workload Analyzer** or the datastore object link. In the standard scenario when the primary site hosting the volume or LUN is functioning as expected, you can see the vServer cluster details of the primary site.
2. In case of a disaster, and a consecutive switchover to the secondary site, the datastore link points to the performance metrics of the volume or LUN in the secondary cluster. This is reflected after the next cycle of clusters and vServer discovery (acquisition) is complete.
3. After a successful switchback, the datastore link again reflects the performance metrics of the volume or LUN in the primary cluster. This is reflected after the next cycle of clusters and vServer discovery is complete.

View datastores in storage VM disaster recovery configuration

Note the following prerequisites before viewing datastores in a storage VM disaster recovery configuration:

- In an event of switchover and switchback, the discovery of the primary and secondary clusters of the HA

pair, and vCenter Servers should be complete.

- Both the source and destination cluster and storage VM peers should be managed by Unified Manager.
- The required setup must be completed on ONTAP and vCenter Server.
  - For NAS (NFS and VMFS) datastores, in case of a disaster, the steps include bringing up the secondary storage VM, verifying the data LIFs and routes, establishing lost connections on vCenter Server, and starting the VMs.

    For a switchback to the primary site, the data between the volumes should be synced before the primary site starts serving the data.

  - For SAN (iSCSI and FC for VMFS) datastores, vCenter Server formats the mounted LUN in a VMFS format. In case of a disaster, the steps include bringing up the secondary storage VM, verifying the data LIFs and routes. If the iSCSI target IPs are different from the primary LIFs, they need to be manually added. The new LUNs should be available as devices under the iSCSI adapter of the storage adapter of the host. Thereafter, new VMFS datastores with the new LUNs should be created and the old VMs registered with new names. The VMs must be up and running.

    In case of a recovery, the data between the volumes should be synced. New VMFS datastores should again be created using the LUNs and the old VMs registered with new names.

    For information about the setup, see ONTAP and vCenter Server documentation.

    ONTAP 9 Documentation Center

Follow these steps for viewing datastores:

1. On the **VMWARE** > **Virtual Machines** page, click the VM inventory that hosts the datastore. Click the datastore object link. In the standard scenario, you can see the performance data of the volumes and LUNs in the primary storage VM.
2. In case of a disaster, and a consecutive switchover to the secondary storage VM, the datastore link points to the performance metrics of the volume or LUN in the secondary storage VM. This is reflected after the next cycle of clusters and vServer discovery (acquisition) is complete.
3. After a successful switchback, the datastore link again reflects the performance metrics of the volume or LUN in the primary storage VM. This is reflected after the next cycle of clusters and vServer discovery is complete.

Unsupported scenarios

- For a MetroCluster configuration, note the following limitations:
  - Clusters in only the `NORMAL` and `SWITCHOVER` states are taken up. Other states, such as `PARTIAL_SWITCHOVER, PARTIAL_SWITCHBACK`, and `NOT_REACHABLE` are not supported.
  - Unless Automatic Switch Over (ASO) is enabled, if the primary cluster goes down, the secondary cluster cannot be discovered, and the topology continues to point to the volume or LUN in the primary cluster.
- For a storage VM disaster recovery configuration, note the following limitation:
  - A configuration with Site Recovery Manager (SRM) or Storage Replication Adapter (SRA) enabled for a SAN storage environment is not supported.

Provision and manage workloads

Provision and manage workloads in Active IQ Unified Manager

The active management feature of Active IQ Unified Manager provides Performance Service Levels, Storage Efficiency Policies, and storage provider APIs for provisioning, monitoring, and managing storage workloads in a data center.

Unified Manager provides this functionality by default. You can disable it from **Storage Management** > **Feature Settings** if you do not plan to use this functionality.

When enabled, you can provision workloads on the ONTAP clusters managed by your instance of Unified Manager. You can also assign policies, such as Performance Service Levels and Storage Efficiency Policies on the workloads and manage your storage environment based on those policies.

This feature enables the following functions:

- Automatic discovery of storage workloads on the added clusters enabling easy storage workload evaluation and deployment
- Provisioning NAS workloads supporting NFS and CIFS protocols
- Provisioning SAN workloads supporting iSCSI and FCP protocols
- Support for both NFS and CIFS protocols on the same file share
- Management of Performance Service Levels and Storage Efficiency Policies
- Assigning Performance Service Levels and Storage Efficiency Policies to storage workloads

The **Provisioning**, **Storage** > **Workloads**, and **Policies** options on the left pane of the UI enable you to modify various configurations.

You can perform the following functions by using these options:

- View storage workloads on the **Storage** > **Workloads** page
- Create storage workloads from the Provision Workload page
- Create and manage Performance Service Levels from Policies
- Create and manage Storage Efficiency Policies from Policies
- Assign policies to storage workloads from the Workloads page

Related information

Policy-based storage management

Workloads overview

Learn about the workloads dashboard in Active IQ Unified Manager

A workload represents the input/output (I/O) operations of a storage object, such as a volume or LUN. The way the storage is provisioned is based on the expected workload requirements. Workload statistics are tracked by Active IQ Unified Manager only after there is traffic to and from the storage object. For example, the workload IOPS and

latency values are available after users start using a database or email application.

The Workloads page displays a summary of the storage workloads of the ONTAP clusters managed by Unified Manager. It provides cumulative at-a-glance information about the storage workloads that conform to the Performance Service Level, as well as the non-conforming storage workloads. It also enables you to assess the total, available, and used capacity and performance (IOPS) of the clusters across your data center.



It is recommended that you assess the number of storage workloads that are non-conforming, unavailable, or not managed by any Performance Service Level, and take the necessary actions to ensure their conformance, capacity usage, and IOPS.

The Workloads page has the following two sections:

- Workloads overview: Provides an overview of the number of storage workloads on the ONTAP clusters managed by Unified Manager.
- Data center overview: Provides an overview of the capacity and IOPS of the storage workloads in the data center. The relevant data is displayed at a data center level and for individual .

Workloads overview section

The workloads overview section provides cumulative at-a-glance information of the storage workloads. The status of the storage workloads is displayed based on assigned and unassigned Performance Service Levels.

- **Assigned**: The following statuses are reported for storage workloads on which Performance Service Levels have been assigned:
  - **Conforming**: Performance of storage workloads is based on the Performance Service Levels assigned to them. If the storage workloads are within the threshold latency defined in the associated Performance Service Levels, they are marked as “conforming”. The conforming workloads are marked in blue.
  - **Non-conforming**: During performance monitoring, storage workloads are marked as “non-conforming” if the storage workloads latency exceeds the threshold latency defined in the associated Performance Service Level. The non-conforming workloads are marked in orange.
  - **Unavailable**: Storage workloads are marked as “unavailable” if they are offline, or if the corresponding cluster is unreachable. The unavailable workloads are marked in red.
- **Unassigned**: Storage workloads that do not have a Performance Service Level assigned to them, are reported as “unassigned”. The number is conveyed by the information icon.

The total workload count is the sum total of the assigned and unassigned workloads.

You can click the total number of workloads displayed in this section, and view them on the Workloads page.

The Conformance by Performance Service Levels subsection displays the total number of available storage workloads:

- Conforming to each type of Performance Service Level
- For which there is a mismatch between the assigned and the recommended Performance Service Levels

Data center overview section

The data center overview section graphically represents the available and used capacity, and IOPS for all of the clusters in the data center. By using this data, you should manage the capacity and IOPS of the storage workloads. The section also displays the following information for the storage workloads across all of the

clusters:

- The total, available, and used capacity for all of the clusters in your data center
- The total, available, and used IOPS for all of the clusters in your data center
- The available and used capacity based on each Performance Service Level
- The available and used IOPS based on each Performance Service Level
- The total space and IOPS used by the workloads that have no Performance Service Level assigned

How data center capacity and performance is calculated based on Performance Service Levels

The used capacity and IOPS is retrieved in terms of the total used capacity and performance of all of the storage workloads in the clusters.

The available IOPS is calculated based on the expected latency and recommended Performance Service Levels on the nodes. It includes the available IOPS for all of the Performance Service Levels whose expected latency is less than or equal to their own expected latency.

The available capacity is calculated based on the expected latency and recommended Performance Service Levels on aggregates. It includes the available capacity for all of the Performance Service Levels whose expected latency is less than or equal to their own expected latency.

View workloads in Active IQ Unified Manager

When you add clusters to Unified Manager, the storage workloads on each cluster are automatically discovered and displayed on the Workloads page.

Unified Manager begins to analyze the workloads for recommendation (recommended PSLs) only after I/O operations start on the storage workloads.

FlexGroup volumes and its constituents are excluded.

Workloads overview

The Workloads Overview page displays the overview of the workloads in the data center and the space and performance overview of the data center.

- **Workloads Overview** panel: Displays the total number of workloads and the number of workloads with or without PSLs assigned on them. The break-up of the workload count for each PSL is also displayed. Clicking the counts takes you to the **All Workloads** view with the filtered workloads. You can also view the number of workloads that do not conform with the system recommendation and assign the system-recommended PSLs to them by clicking the **Assign System-Recommended PSLs** button.
- **Data Center Overview** panel: Displays the available and used space (TiB) and performance (IOPS) of the data center. A break-up of the available and used space (TiB) and performance (IOPS) of all the workloads under each PSL is also displayed.

All workloads view

The **Storage > Workloads > All Workloads** page lists the storage workloads associated with the ONTAP clusters managed by Unified Manager.

For the newly-discovered storage workloads on which there have been no I/O operations, the status is “Waiting for I/O”. After I/O operations begin on the storage workloads, Unified Manager begins the analysis and the

workload status changes to “Learning…”. After the analysis is complete (within 24 hours from the beginning of the I/O operations), the recommended PSLs are displayed for the storage workloads.

The page also enables you to assign Storage Efficiency Policies (SEPs) and Performance Service Levels (PSLs) to storage workloads. You can perform multiple tasks:

- Add or provision storage workloads
- View and filter the list of workloads
- Assign PSLs to storage workloads
- Evaluate system-recommended PSLs and assign them to workloads
- Assign SEPs to storage workloads

Adding or provisioning storage workloads

You can add or provision the storage workloads to supported LUNs (supporting both iSCSI and FCP protocols), NFS file shares, and SMB shares.

Steps

1. Click **Storage > Workloads > All Workloads > Create**.
2. Create workloads. For information, see Provisioning and managing workloads.

Viewing and filtering workloads

On the All Workloads screen, you can view all the workloads in your data center or search for specific storage workloads based on either their PSLs or names. You can use the filter icon to enter specific conditions for your search. You can search by different filter conditions, such as by the host cluster or storage VM. The **Capacity Total** option allows filtering by the total capacity of the workloads (by MB). However, in this case, the number of workloads returned might vary, because the total capacity is compared at a byte level.

For each workload, information, such as the host cluster and storage VM is displayed, along with the assigned PSL and SEP.

The page also enables you to view the performance details of a workload. You can view detailed information about the IOPS, capacity, and latency of the workload by clicking the **Choose / Order Columns** button and selecting specific columns to view. The Performance View column displays the average and peak IOPS for a workload, and you can click the workload analyser icon to view the detailed IOPS analysis.

Analyzing performance and capacity criteria for a workload

The **Analyze Workload** button on the **IOPS Analysis** pop-up takes you to the Workload Analysis page, where you can select a time range and view the latency, throughput, and capacity trends for the selected workload. For more information about workload analyzer, see Troubleshooting workloads using the workload analyzer.

You can view performance information about a workload to help with troubleshooting by clicking the bar chart icon in the **Performance View** column. To view performance and capacity charts on the Workload Analysis page to analyze the object, click the **Analyze Workload** button.

For more information, see What data does the workload analyzer display.

Assign policies to workloads in Active IQ Unified Manager

You can assign Storage Efficiency Policies (SEPs) and Performance Service Levels

(PSLs) to storage workloads from the All Workloads page by using the different navigation options.

Assign policies to a single workload

You can assign either a PSL or an SEP or both, to a single workload. Follow these steps:

1. Select the workload.
2. Click the edit icon next to the row, and then click **Edit**.

   The **Assigned Performance Service Level** and **Storage Efficiency Policy** fields are enabled.

3. Select the required PSL or SEP, or both.
4. Click the check icon to apply the changes.

> You can also select a workload and click **More Actions** to assign the policies.

Assign policies to multiple storage workloads

You can assign a PSL or an SEP to multiple storage workloads together. Follow these steps:

1. Select the check boxes for the workloads to which you want to assign the policy, or select all the workloads in your data center.
2. Click **More Actions**.
3. For assigning a PSL, select **Assign Performance Service Level**. For assigning an SEP, select **Assign Storage Efficiency Policy**. A pop-up is displayed for you to select the policy.
4. Select the appropriate policy and click **Apply**. The number of workloads on which the policies are assigned are displayed. The workloads on which the policies are not assigned are also listed, with the cause.

> Applying policies on workloads in bulk might take some time depending on the number of workloads selected. You can click the **Run in background** button and continue with other tasks while the operation runs in the background. When the bulk assignment is complete, you can view the completion status. If you are applying a PSL on multiple workloads, you cannot trigger another request when the previous job of bulk assignment is running.

Assign system-recommended PSLs to workloads

You can assign system-recommended PSLs to those storage workloads in a data center that have no PSLs assigned, or the assigned PSLs do not match the system recommendation. To use this functionality, click the **Assign System Recommended PSLs** button. You do not have to select specific workloads.

The recommendation is internally determined by system analytics, and is skipped for those workloads whose IOPS and other parameters do not coincide with the definitions of any available PSL. Storage workloads with `Waiting for I/O` and Learning statuses are also excluded.

> There are special keywords that Unified Manager looks for in the workload name to override the system analytics and recommend a different PSL for the workload. When the workload has the letters “ora” in the name, the **Extreme Performance**PSL is recommended. And when the workload has the letters “vm” in the name, the **Performance**PSL is recommended.

Also see the knowledge base (KB) article ActiveIQ Unified Manager 'Assign System Recommended

Performance Service Level' is not adaptive to a highly variable workload

Provision file share volumes in Active IQ Unified Manager

You can create file share volumes that support CIFS/SMB and NFS protocols, on an existing cluster and Storage Virtual Machine (storage VM) from the Provision Workload page.

Before you begin

- The storage VM must have space for provisioning the file share volume.
- Either or both of the SMB and NFS services should be enabled on your storage VM.
- For selecting and assigning the Performance Service Level (PSL) and Storage Efficiency Policy (SEP) on the workload, the policies must have been created before you start creating the workload.

Steps

1. On the **Provision Workload** page, add the name of the workload that you want to create, and then select the cluster from the available list.
2. Based on the cluster that you have selected, the **STORAGE VM** field filters the available storage VMs for that cluster. Select the required storage VM from the list.

   Based on the SMB and NFS services supported on the storage VM, the NAS option is enabled in the Host Information section.

3. In the Storage and Optimization section, assign the storage capacity and PSL, and optionally, an SEP for the workload.

   The specifications for the SEP are assigned to the LUN and the definitions for the PSL are applied to the workload when it is created.

4. Select the **Enforce performance limits** check box if you want to enforce the PSL that you have assigned to the workload.

   Assigning a PSL to a workload ensures that the aggregate on which the workload is created can support the performance and capacity objectives defined in the respective policy. For example, if a workload is assigned "Extreme Performance "PSL, the aggregate on which the workload is to be provisioned should have the capability of supporting the performance and capacity objectives of the "Extreme Performance" policy, such as SSD storage.

   Unless you select this check box, the PSL is not applied to the workload, and the status of the workload on the dashboard appears as unassigned.

5. Select the **NAS** option.

   If you cannot see the **NAS** option enabled, verify whether the storage VM that you have selected supports either SMB or NFS, or both.

   If your storage VM is enabled for both SMB and NFS services, you can select the **Share by NFS** and **Share by SMB** check boxes and create a file share that supports both NFS and SMB protocols. If you want to create either an SMB or a CIFS share, select only the respective check box.

6. For NFS file share volumes, specify the IP address of the host or network to access the file share volume. You can enter comma-separated values for multiple hosts.

   On adding the host IP address, an internal check runs for matching the host details with the storage VM and the export policy for that host is created, or in case there is an existing policy, it is reused. If there are several NFS shares created for the same host, then an available export policy for the same host with matching rules is reused for all the files shares. The function of specifying rules of individual policies or reusing policies by providing specific policy keys is available when you provision the NFS share by using APIs.

7. For an SMB share, specify which users or user groups can access the SMB share and assign the required permissions. For each group of users, a new access control list (ACL) is generated during the file share creation.
8. Click **Save**.

   The workload is added to the list of storage workloads.

Provision LUNs in Active IQ Unified Manager

You can create LUNs that support CIFS/SMB and NFS protocols, on an existing cluster and Storage Virtual Machine (storage VM) from the Provision Workload page.

Before you begin

- The storage VM must have space for provisioning the LUN.
- Both iSCSI and FCP must be enabled on the storage VM on which you create the LUN.
- For selecting and assigning the Performance Service Level (PSL) and Storage Efficiency Policy (SEP) on the workload, the policies must have been created before you start creating the workload.

Steps

1. On the **Provision Workload** page, add the name of the workload that you want to create, and then select the cluster from the available list.

   Based on the cluster that you have selected, the **STORAGE VM** field filters the available storage VMs for that cluster.

2. Select the storage VM from the list that supports the iSCSI and FCP services.

   Based on your selection, the SAN option is enabled in the Host Information section.

3. In the **Storage and Optimization** section, assign the storage capacity and PSL, and optionally, the SEP for the workload.

   The specifications for the SEP are assigned to the LUN and the definitions for the PSL are applied to the workload when it is created.

4. Select the **Enforce performance limits** check box if you want to enforce the assigned PSL on the workload.

   Assigning a PSL to a workload ensures that the aggregate on which the workload is created can support the performance and capacity objectives defined in the respective policy. For example, if a workload is assigned the "Extreme Performance" PSL, the aggregate on which the workload is to be provisioned should have the capability of supporting the performance and capacity objectives of the "Extreme Performance" policy, such as SSD storage.

Unless you select this check box, the PSL is not applied to the workload, and the status of the workload on the dashboard appears as `unassigned`.

5. Select the **SAN** option. If you cannot see the **SAN** option enabled, verify whether the storage VM that you have selected supports iSCSI and FCP.
6. Select the host OS.
7. Specify the host mapping to control access of the initiators to the LUN. You can assign existing initiator groups (igroups), or define and map new igroups.

   If you create a new igroup while provisioning the LUN, you need to wait till the next discovery cycle (up to 15 minutes) for using it. It is therefore recommended that you use an existing igroup from the list of available igroups.

   If you want to create a new igroup, select the **Create a new initiator group** button, and enter the information for the igroup.

8. Click **Save**.

   The LUN is added to the list of storage workloads.

Manage Performance Service Levels

Manage Performance Service Levels in Active IQ Unified Manager

A Performance Service Level (PSL) enables you to define the performance and storage objectives for a workload. You can assign a PSL to a workload when initially creating the workload, or afterwards by editing the workload.

The management and monitoring of storage resources are based on Service Level Objectives (SLOs). SLOs are defined by service level agreements that are based on required performance and capacity. In Unified Manager, SLOs refer to the PSL definitions of the applications that are running on NetApp storage. Storage services are differentiated based on the performance and utilization of the underlying resources. A PSL is a description of the storage service objectives. A PSL enables the storage provider to specify the performance and capacity objectives for the workload. When you assign a PSL on a workload, the corresponding workload on ONTAP is managed by its performance and capacity objectives. Each PSL is governed by peak, expected, and absolute minimum IOPs, and expected latency.

Unified Manager has the following types of PSLs:

- **System-defined**: Unified Manager provides a few canned policies that cannot be changed. These predefined PSLs are:
  - Extreme Performance
  - Performance
  - Value

    The Extreme Performance, Performance, and Value PSLs are applicable for most of the common storage workloads in a data center.

    Unified Manager also offers three Performance Service Levels for database applications. These are extremely high-performance PSLs that support bursty IOPS and are appropriate for database

applications with the highest throughput demand.

  - Extreme for Database Logs
  - Extreme for Database Shared Data
  - Extreme for Database Data

- **User-defined**: If the predefined Performance Service Levels do not meet your requirements, then you can create new PSLs to meet your needs. For information, see Create and edit Performance Service Levels.
- **Beyond Extreme**: The Beyond Extreme PSLs are the system-recommended PSLs that are suggested for workloads that demand IOPs higher than Extreme. The workloads are internally analyzed based on their IOPS, capacity, and latency, and a Beyond Extreme PSL is recommended for each of these workloads on the **Storage > Workloads > All Workloads** screen. You can apply the PSLs to the workloads to ensure an optimum performance.

  The IOPs parameters for the workloads are dynamically generated, depending on the workload behavior, and appended to the name of the Beyond Extreme PSL in the format `Beyond Extreme <number-(peak IOPS/TB)> <number(expected IOPS/TB)>`. For example, if the system determines a workload to have the peak and expected IOPs as `106345` and `37929` respectively, the Beyond Extreme PSL that is generated for the workload is named as `Beyond Extreme 106345 37929`. Though these PSLs are recommended by the system, when you assign them to workloads, these PSLs are labeled as `User-defined` in type.

Manage workloads by assigning PSLs

You can access PSLs from the **Policies** > **Performance Service Levels** page and by using the storage provider APIs. Managing storage workloads by assigning PSLs to them is convenient as you do not have to individually manage the storage workloads. Any modifications can also be managed by reassigning another PSL rather than managing them individually. Unified Manager helps you to assign PSLs on your workloads based on internal assessment and recommendations.

For information on assigning system-recommended PSLs to workloads, see Assign system-recommended PSLs to workloads

The Performance Service Levels page lists the available PSL policies and enables you to add, edit, and delete them.

> You cannot modify a PSL that is system-defined or that is currently assigned to a workload. You cannot delete a PSL that is assigned to a workload, or if it is the only available PSL.

This page displays the following information:

Field	Description
Name	Name of the PSL.
Type	Whether the policy is system-defined or user-defined.
Expected IOPS/TB	Minimum number of IOPS that an application is expected to perform on a LUN or file share. Expected IOPS specifies the minimum expected IOPS allocated, based on the storage object allocated size.

Field	Description
Peak IOPS/TB	Maximum number of IOPS that an application can perform on a LUN or file share. Peak IOPS specifies the maximum possible IOPS allocated, based on the storage object allocated size or the storage object used size. Peak IOPS are based on an allocation policy. The allocation policy is either allocated-space or used-space. When the allocation policy is set to allocated-space, the peak IOPS is calculated based on the size of the storage object. When the allocation policy is set to used-space, the peak IOPS is calculated based on the amount of data stored in the storage object, taking into account storage efficiencies. By default, the allocation policy is set to used-space.
Absolute minimum IOPS	The absolute minimum IOPS is used as an override, when the expected IOPS is less than this value. The default values of the system-defined PSLs are the following: • Extreme Performance: If expected IOPS >= 6144/TB, then absolute minimum IOPS = 1000 • Performance: If expected IOPS >= 2048/TB and < 6144/TB, then absolute minimum IOPS = 500 • Value: If expected IOPS >= 128/TB and < 2048/TB, then absolute minimum IOPS = 75 The default values of the system-defined database PSLs are the following: • Extreme for Database Logs: If expected IOPS >= 22528, then absolute minimum IOPS = 4000 • Extreme for Database Shared Data: If expected IOPS >= 16384, then absolute minimum IOPS = 2000 • Extreme for Database Data: If expected IOPS >= 12288, then absolute minimum IOPS = 2000 The higher value of the absolute minimum IOPS for custom PSLs can be a maximum of 75000. The lower value is calculated as the following: 1000/expected latency
Expected latency	Expected latency for storage IOPS in milliseconds per operation (ms/op).
Capacity	Total available and used capacity in the clusters.
Workloads	Number of storage workloads that have been assigned the PSL.

For information about how the peak IOPS and expected IOPs help in achieving consistent differentiated performance on ONTAP clusters, see the following KB article:
What is Performance Budgeting?

Events generated for workloads breaching the threshold defined by PSLs

Note that if workloads exceed the expected latency value for 30% of the time during the previous hour, Unified Manager generates one of the following events to notify you of a potential performance issue:

- Workload Volume Latency Threshold Breached as defined by Performance Service Level Policy
- Workload LUN Latency Threshold Breached as defined by Performance Service Level Policy.

You may want to analyze the workload to see what may be causing the higher latency values.

For more information, see the following links:

- Volume events
- What happens when a performance threshold policy is breached
- How Unified Manager uses workload latency to identify performance issues
- What performance events are

System-defined PSLs

The following table provides information about the system-defined PSLs:

Performance Service Level	Description and use case	Expected latency (ms/op)	Peak IOPS	Expected IOPS	Absolute minimum IOPS
Extreme Performance	Provides extremely high throughput at a very low latency Ideal for latency-sensitive applications	1	12288	6144	1000
Performance	Provides high throughput at a low latency Ideal for database and virtualized applications	2	4096	2048	500

Performance Service Level	Description and use case	Expected latency (ms/op)	Peak IOPS	Expected IOPS	Absolute minimum IOPS
Value	Provides high storage capacity and moderate latency Ideal for high-capacity applications such as email, web content, file shares, and backup targets	17	512	128	75
Extreme for Database Logs	Provides maximum throughput at the lowest latency. Ideal for database applications supporting database logs. This PSL provides the highest throughput because database logs are extremely bursty and logging is constantly in demand.	1	45056	22528	4000
Extreme for Database Shared Data	Provides very high throughput at the lowest latency. Ideal for database applications data that is stored in a common data store, but is shared across databases.	1	32768	16384	2000

Performance Service Level	Description and use case	Expected latency (ms/op)	Peak IOPS	Expected IOPS	Absolute minimum IOPS
Extreme for Database Data	Provides high throughput at the lowest latency. Ideal for database applications data, such as database table information and metadata.	1	24576	12288	2000

Create and edit Performance Service Levels in Active IQ Unified Manager

When the system-defined Performance Service Levels do not match your workload requirements, you can create your own Performance Service Levels that are optimized for your workloads.

Before you begin

- You must have the Application Administrator role.
- The Performance Service Level name must be unique, and you cannot use the following reserved keywords:

  `Prime`, `Extreme`, `Performance`, `Value`, `Unassigned`, `Learning`, `Idle`, `Default`, and `None`.

You create and edit custom Performance Service Levels from the Performance Service Levels page by defining the service level objectives you require for the applications that will access storage.

You cannot modify a Performance Service Level if it is currently assigned to a workload.

Steps

1. In the left navigation pane under **Settings**, select **Policies** > **Performance Service Levels**.
2. In the **Performance Service Levels** page, click the appropriate button depending on whether you want to create a new Performance Service Level or if you want to edit an existing Performance Service Level.

To…	Follow these steps…
Create a new Performance Service Level	Click **Add**.
Edit an existing Performance Service Level	Select an existing Performance Service Level, and then click **Edit**.

   The page to add or edit a Performance Service Level is displayed.

3. Customize the Performance Service Level by specifying the performance objectives, and then click **Submit** to save the Performance Service Level.

You can apply the new or changed Performance Service Level to workloads (LUNs, NFS File Shares, CIFS Shares) from the Workloads page or when provisioning a new workload.

Manage Storage Efficiency Policies

Manage Storage Efficiency Policies in Active IQ Unified Manager

A Storage Efficiency Policy (SEP) enables you to define the storage efficiency characteristics of a workload. You can assign an SEP to a workload when initially creating the workload, or afterwards by editing the workload.

Storage efficiency includes using technologies, such as thin provisioning, deduplication, and data compression that increase storage utilization and decrease storage costs. While creating SEPs, you can use these space-saving technologies either individually or together to achieve maximum storage efficiency. When you associate the policies with your storage workloads, the specified policy settings are assigned to them. Unified Manager enables you to assign system-defined and user-defined SEPs to optimize storage resources in your data center.

Unified Manager provides two system-defined SEPs: High and Low. These SEPs are applicable to most of the storage workloads in a data center, however, you can create your own policies if the system-defined SEPs do not meet your requirements.

You cannot modify an SEP that is system-defined or that is currently assigned to a workload. You cannot delete an SEP that is assigned to a workload, or if it is the only available SEP.

The Storage Efficiency Policies page lists the available SEPs and enables you to add, edit, and delete customized SEPs. This page displays the following information:

Field	Description
Name	Name of the SEP.
Type	Whether the policy is system-defined or user-defined.
Space Reserve	Whether the volume is thin-provisioned or thick-provisioned.
Deduplication	Whether deduplication is enabled on the workload: • Inline: Deduplication occurs while being written on the workload • Background: Deduplication occurs in the workload • Disable: Deduplication is disabled on the workload
Compression	Whether data compression is enabled on the workload: • Inline: Data compression occurs while being written on the workload • Background: Data compression occurs in the workload • Disable: Data compression is disabled on the workload
Workloads	Number of storage workloads that have been assigned the SEP

Guidelines for creating a custom Storage Efficiency Policy

If the existing SEPs do not meet policy requirements for your storage workloads, you can create a custom SEP. However, it is recommended that you attempt to use the system-defined SEPs for your storage workloads, and only create custom SEPs if necessary.

You can view the SEP assigned to workloads in the All Workloads page and in the Volume / Health details page. You can view the cluster-level data reduction ratio (without Snapshot copies) based on these storage efficiencies in the Capacity panel on the Dashboard and in the Capacity: All Clusters view.

Create and edit Storage Efficiency Policies in Active IQ Unified Manager

When the system-defined Storage Efficiency Policies do not match your workload requirements, you can create your own Storage Efficiency Policies that are optimized for your workloads.

Before you begin

- You must have the Application Administrator role.
- The Storage Efficiency Policy name must be unique, and you cannot use the following reserved keywords:

  `High`, `Low`, `Unassigned`, `Learning`, `Idle`, `Default`, and `None`.

You create and edit custom Storage Efficiency Policies from the Storage Efficiency Policies page by defining the storage efficiency characteristics you require for the applications that will access storage.

> ⓘ You cannot modify a Storage Efficiency Policy if it is currently assigned to a workload.

Steps

1. In the left navigation pane under **Settings**, select **Policies** > **Storage Efficiency**.
2. In the **Storage Efficiency Policies** page, click the appropriate button depending on whether you want to create a new Storage Efficiency Policy or if you want to edit an existing Storage Efficiency Policy.

To…	Follow these steps…
Create a new Storage Efficiency Policy	Click **Add**
Edit an existing Storage Efficiency Policy	Select an existing Storage Efficiency Policy and click **Edit**

   The page to add or edit a Storage Efficiency Policy is displayed.

3. Customize the Storage Efficiency Policy by specifying the storage efficiency characteristics, and then click **Submit** to save the Storage Efficiency Policy.

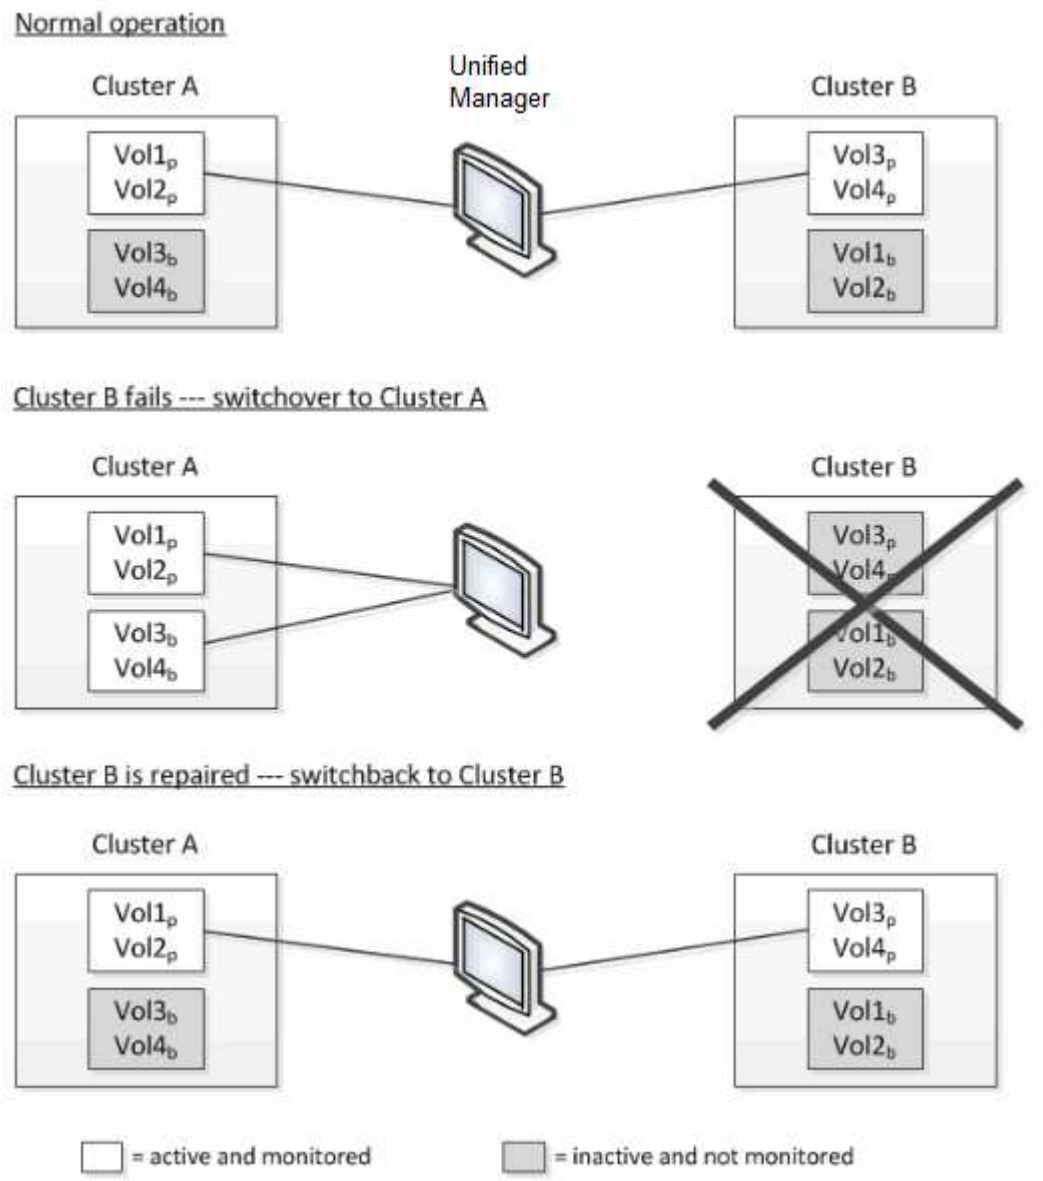
You can apply the new or changed Storage Efficiency Policy to workloads (LUNs, NFS File Shares, CIFS Shares) from the Workloads page or when provisioning a new workload.

Manage and monitor MetroCluster configurations

Learn about volume behavior during switchover and switchback in Active IQ Unified Manager

Events that trigger a switchover or switchback cause active volumes to be moved from one cluster to the other cluster in the disaster recovery group. The volumes on the cluster that were active and serving data to clients are stopped, and the volumes on the other cluster are activated and start serving data. Unified Manager monitors only those volumes that are active and running.

Because volumes are moved from one cluster to another, it is recommended that you monitor both clusters. A single instance of Unified Manager can monitor both clusters in a MetroCluster configuration, but sometimes the distance between the two locations necessitates using two Unified Manager instances to monitor both clusters. The following figure shows a single instance of Unified Manager:



The volumes with p in their names indicate the primary volumes, and the volumes with b in their names are mirrored backup volumes that are created by SnapMirror.

During normal operation:

- Cluster A has two active volumes: Vol1p and Vol2p.
- Cluster B has two active volumes: Vol3p and Vol4p.
- Cluster A has two inactive volumes: Vol3b and Vol4b.
- Cluster B has two inactive volumes: Vol1b and Vol2b.

Information pertaining to each of the active volumes (statistics, events, and so on) is collected by Unified Manager. Vol1p and Vol2p statistics are collected by Cluster A, and Vol3p and Vol4p statistics are collected by Cluster B.

After a catastrophic failure causes a switchover of active volumes from Cluster B to Cluster A:

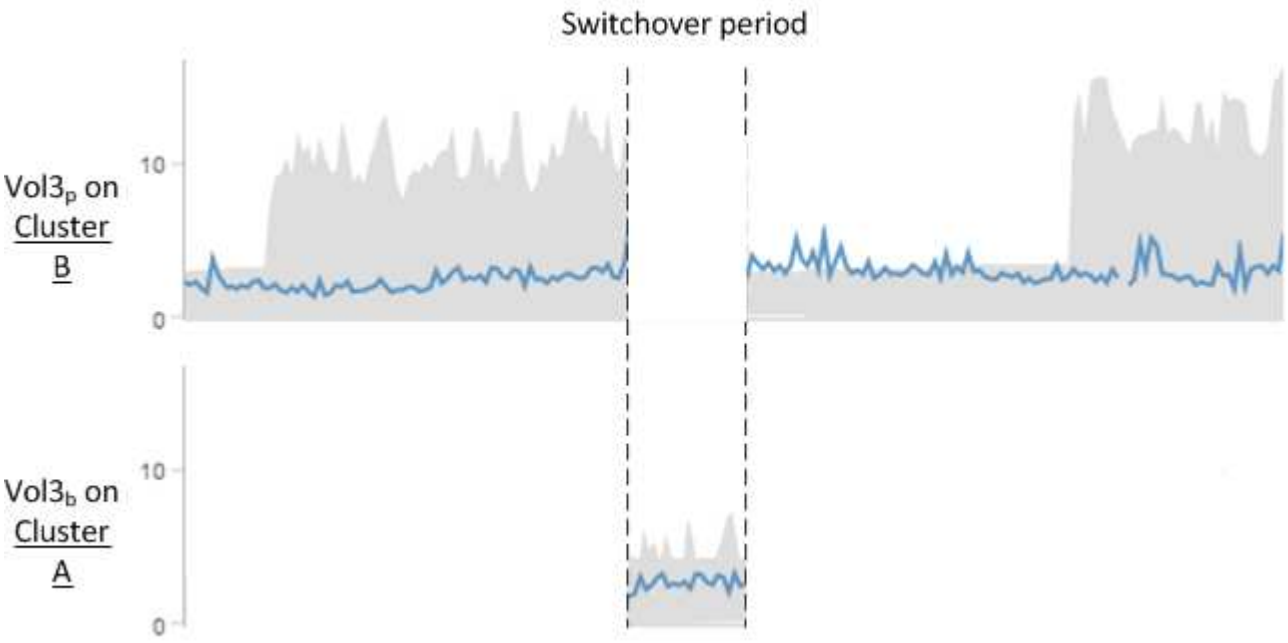
- Cluster A has four active volumes: Vol1p, Vol2p, Vol3b, and Vol4b.
- Cluster B has four inactive volumes: Vol3p, Vol4p, Vol1b, and Vol2b.

As during normal operation, information pertaining to each of the active volumes is collected by Unified Manager. But in this case, Vol1p and Vol2p statistics are collected by Cluster A, and Vol3b and Vol4b statistics are also collected by Cluster A.

Note that Vol3p and Vol3b are not the same volumes, because they are on different clusters. The information in Unified Manager for Vol3p is not the same as Vol3b:

- During switchover to Cluster A, Vol3p statistics and events are not visible.
- On the very first switchover, Vol3b looks like a new volume with no historical information.

When Cluster B is repaired and a switchback is performed, Vol3p is active again on Cluster B, with the historical statistics and a gap of statistics for the period during the switchover. Vol3b is not viewable from Cluster A until another switchover occurs:



- MetroCluster volumes that are inactive, for example, Vol3b on Cluster A after switchback, are identified with the message “This volume was deleted”. The volume is not actually deleted, but it is not currently being monitored by Unified Manager because it is not the active volume.

- If a single Unified Manager is monitoring both clusters in a MetroCluster configuration, volume search returns information for whichever volume is active at that time. For example, a search for "Vol3" would return statistics and events for Vol3b on Cluster A if a switchover has occurred and Vol3 has become active on Cluster A.

Understand cluster connectivity in a MetroCluster over FC configuration in Active IQ Unified Manager

Connectivity between the clusters in a MetroCluster over FC configuration can be one of the following statuses: Optimal, Impacted, or Down. Understanding the connectivity statuses enables you to manage your MetroCluster configurations effectively.

Connectivity status	Description	Icon displayed
Optimal	Connectivity between the clusters in the MetroCluster configuration is normal.	
Impacted	One or more errors compromise the status of failover availability; however, both of the clusters in the MetroCluster configuration are still up. For example, when the ISL link is down, when the intercluster IP link is down, or when the partner cluster is not reachable.	
Down	Connectivity between the clusters in the MetroCluster configuration is down because one or both of the clusters are down or the clusters are in failover mode. For example, when the partner cluster is down because of a disaster or when there is a planned switchover for testing purposes.	Switchover with errors:  Switchover successful: 

Learn about MetroCluster over FC data mirroring in Active IQ Unified Manager

MetroCluster over FC configurations provide data mirroring and the additional ability to initiate a failover if an entire site becomes unavailable. The status of data mirroring between the clusters in a MetroCluster over FC configuration can either be Normal or Mirroring Unavailable. Understanding the status enables you to manage your MetroCluster configurations effectively.

Data mirroring status	Description	Icon displayed
Normal	Data mirroring between the clusters in the MetroCluster configuration is normal.	
Mirroring Unavailable	Data mirroring between the clusters in the MetroCluster configuration is unavailable because of switchover. For example, when the partner cluster is down because of a disaster or when there is a planned switchover for testing purposes.	Switchover with errors:  Switchover successful: 

Monitor MetroCluster configurations in Active IQ Unified Manager

You can monitor connectivity issues in your MetroCluster configuration. The details include the status of the components and connectivity within a cluster and the connectivity status between the clusters in the MetroCluster configuration. Here, you will get to know how to monitor connectivity issues in clusters protected by MetroCluster over FC and MetroCluster over IP configurations.

You can monitor the MetroCluster configurations from the following views from the Active IQ Unified Manager left navigation pane:

- **Storage** > **Clusters** > **Protection: MetroCluster** view
- **Protection** > **Relationships** > **Relationship: MetroCluster** view

Unified Manager uses system health alerts to indicate the status of the components and connectivity in the MetroCluster configuration.

Before you begin

- Both the local and remote clusters in a MetroCluster configuration must be added to Active IQ Unified Manager.
- In a MetroCluster over IP configuration, if a Mediator is to be supported, the Mediator should be configured and added to the cluster by corresponding API.
- You must have the Operator, Application Administrator, or Storage Administrator role.

Monitor connectivity issues in MetroCluster over FC configuration

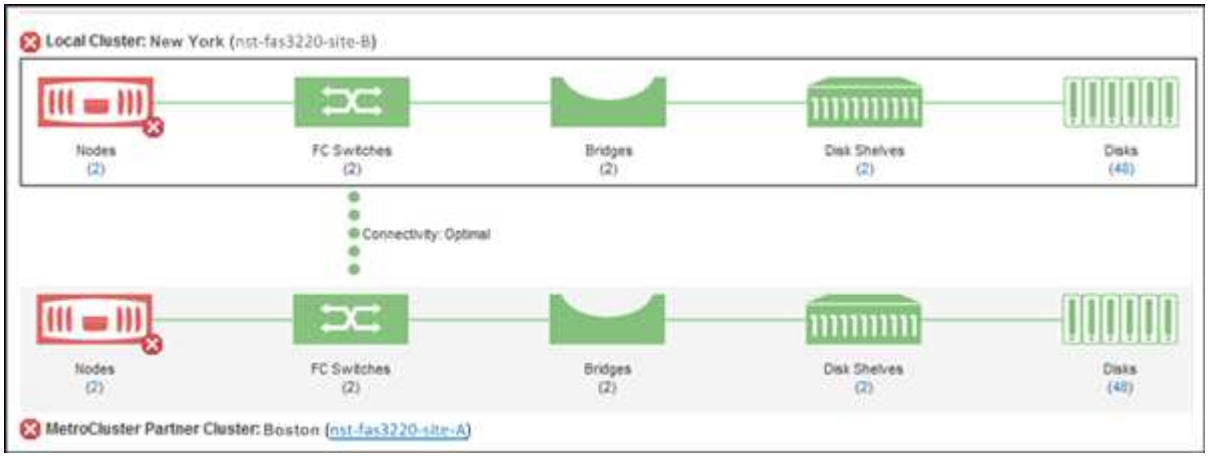
For clusters in a MetroCluster over FC configuration, the connectivity charts are displayed on the **Cluster / Health** details page. Follow these steps.

Steps

1. In the left navigation pane, click **Storage** > **Clusters**.

A list of all of the monitored clusters is displayed.

2. From the **Protection: MetroCluster** view, click the name of the cluster for which you want to view MetroCluster over FC configuration details. Alternately, you can filter by clusters in a MetroCluster configuration.
3. In the **Cluster / Health** details page, click the **MetroCluster Connectivity** tab. The **MetroCluster Connectivity** tab is available for only MetroCluster over FC configurations.



The topology of the MetroCluster configuration is displayed in the corresponding cluster object area. You can use the information displayed in the Cluster / Health details page to rectify any connectivity issues. For example, if the connectivity between the node and the switch in a cluster is down, the following icon is displayed:



If you move the pointer over the icon, you can view detailed information about the generated event.

If you discover connectivity issues in your MetroCluster configuration, you must log in to System Manager or access the ONTAP CLI to resolve the issues.

For more information about determining cluster health, see Determine cluster health in MetroCluster over FC configuration.

Monitor connectivity issues in MetroCluster over IP configuration

For clusters in a MetroCluster over IP configuration, the connectivity charts are displayed on the **Clusters** page. Follow these steps.

Steps

1. In the left navigation pane, click **Storage** > **Clusters**.

   A list of all of the monitored clusters is displayed.

2. From the **Protection: MetroClusters** view, click the name of the cluster for which you want to view MetroCluster over IP configuration details. Alternately, you can filter by clusters in a MetroCluster configuration.
3. Expand the row by clicking the caret ⌄ icon. The caret icon appears for only a cluster that this protected by MetroCluster over IP configuration.

You can view the topology of the source and mirror sites, as well as the Mediator, if any, used for the connection. You can view the following information:

- Connectivity across the sites
- Health and availability issues, if any, on both the sites
- Mediator-related issues
- Replication related issues.

> The following statuses are reported: Critical (), Error (), or Normal (). You can also view the aggregate data replication status of the primary and mirror data in the same topology.

In the following diagram, you can see that the intersite connectivity between the source and destination clusters is unavailable, and the Mediator between them is not configured.



4. Click the status icon. A message with the error definition is displayed. If an event has been raised for the issue in your MetroCluster over IP configuration, you can click the **View Event** button on the message and view the event details. When you have resolved the issue and the event, the status icon in this topology turns to normal ().
5. You can view further configuration details in the **MetroCluster Overview** and **Protection** sections on the **Configuration** tab of the **Cluster / Health** details page.

> Only for a MetroCluster over IP configuration, you can have view the cluster topology on the **Clusters** page. For clusters in a MetroCluster over FC configuration, the topology is displayed on the **MetroCluster Connectivity** tab on the **Cluster / Health** details page.

Related information

- Cluster / Health details page
- For information about **Relationship:MetroCluster** view, see Monitor MetroCluster configurations.
- For information about **Relationship: Last 1 month Transfer Status** view, see Relationship: Last 1 month Transfer Status view.
- For information about **Relationship: Last 1 month Transfer Rate** view, see Relationship: Last 1 month

Transfer Rate view.

- For information about **Relationship: All Relationships** view, see Relationship: All Relationships view.

Monitor MetroCluster replication in Active IQ Unified Manager

You can monitor and diagnose the overall health condition of the logical connections while mirroring the data. You can identify the issues or any risk that interrupts mirroring of cluster components such as aggregates, nodes, and storage virtual machines.

Unified Manager uses system health alerts to monitor the status of the components and connectivity in the MetroCluster configuration.

Before you begin

Both the local and remote cluster in the MetroCluster configuration must be added to Unified Manager

View replication for MetroCluster over IP configurations

For MetroCluster over IP configurations, the data replication status is displayed in the topology peek view for a cluster protected by MetroCluster over IP from the following views from the Unified Manager left navigation pane:

- **Storage** > **Clusters** > **Protection: MetroCluster** view
- **Protection** > **Relationships** > **Relationship: MetroCluster** view

For information, see Monitor connectivity issues in MetroCluster over IP.

View replication for MetroCluster over FC configurations

Follow these steps to determine any issues in the data replication for MetroCluster over FC configuration.

Steps

1. In the left navigation pane, click **Storage** > **Clusters**.

   A list of the monitored clusters is displayed.

2. From the **Health: All Clusters** view, click the name of the cluster for which you want to view MetroCluster replication details. On the **Cluster / Health details** page, click the **MetroCluster Replication** tab.

   The topology of the MetroCluster configuration to be replicated is displayed at the local site in the corresponding cluster object area with the information about the remote site where the data is being mirrored. If you move the pointer over the icon, you can view detailed information about the generated event.

You can use the information displayed in the Cluster / Health details page to rectify any replication issues. If you discover mirroring issues in your MetroCluster configuration, you must log in to System Manager or access the ONTAP CLI to resolve the issues.

Related information

Cluster / Health details page

Manage quotas

Learn about quota limits in Active IQ Unified Manager

User quota limits are values that the Unified Manager server uses to evaluate whether space consumption by a user is nearing the limit or has reached the limit that is set by the user's quota. If the soft limit is crossed or if the hard limit is reached, the Unified Manager server generates user quota events.

By default, the Unified Manager server sends a notification email to users who have crossed the quota soft limit or have reached the quota hard limit and for which user quota events are configured. Users with the Application Administrator role can configure alerts that notify the specified recipients of the user or user group quota events.

You can specify quota limits by using either ONTAP System Manager or the ONTAP CLI.

View user and user group quotas in Active IQ Unified Manager

The Storage VM / Health details page displays information about the user and user group quotas that are configured on the SVM. You can view the name of the user or user group, limits set on the disks and files, used disk and file space, and email address for notification.

Before you begin

You must have the Operator, Application Administrator, or Storage Administrator role.

Steps

1. In the left navigation pane, click **Storage** > **Storage VMs**.
2. In the **Health: All Storage VMs** view, select a Storage VM and then click the **User and Group Quotas** tab.

Related information

Adding users

Create rules to generate email addresses in Active IQ Unified Manager

You can create rules to specify the email address based on the user quota associated with clusters, storage virtual machines (SVMs), volumes, qtrees, users, or user groups. A notification is sent to the specified email address when there is a quota breach.

Before you begin

- You must have the Application Administrator or Storage Administrator role.
- You must have reviewed the guidelines on the Rules to Generate User and Group Quota Email Address page.

You must define the rules for quota email addresses and enter them in the order in which you want to execute them. For example, if you want to use the email address abc@xyz.com to receive notifications about quota breaches for abc and use the email address dl-$GROUP@$DOMAIN for all the other groups, the rules must be listed in the following order:

- if ($USER == 'abc') then abc@xyz.com
- if ($GROUP == *) then dl-$GROUP@$DOMAIN

If none of the criteria for the rules you specified are met, then the default rule is used:

if ($USER_OR_GROUP == *) then $USER_OR_GROUP@$DOMAIN

Steps

1. In the left navigation pane, click **General** > **Quota Email Rules**.
2. Enter the rule based on your criteria.
3. Click **Validate** to validate the syntax of the rule.

   An error message is displayed if the syntax of the rule is incorrect. You must correct the syntax and click **Validate** again.

4. Click **Save**.
5. Verify that the email address you created is displayed in the **User and Group Quotas** tab of the Storage **VM / Health** details page.

Create an email notification format for user and user group quotas in Active IQ Unified Manager

You can create a notification format for the emails that are sent to a user or a user group when there is a quota-related issue (soft limit breached or hard limit reached).

Before you begin

You must have the Application Administrator or Storage Administrator role.

Steps

1. In the left navigation pane, click **General** > **Quota Email Format**.
2. Enter or modify the details in the **From**, **Subject**, and **Email Details** fields.
3. Click **Preview** to preview the email notification.
4. Click **Close** to close the preview window.
5. Modify the content of the email notification, if required.
6. Click **Save**.

Edit user and group quota email addresses in Active IQ Unified Manager

You can modify the email addresses based on the user quota associated with clusters, storage virtual machines (SVMs), volumes, qtrees, users, or user groups. You can modify the email address when you want to override the email address generated by rules specified in the Rules to Generate User and Group Quota Email Address dialog box.

Before you begin

- You must have the Operator, Application Administrator, or Storage Administrator role.
- You must have reviewed the guidelines for creating rules.

If you edit an email address, the rules to generate the user and group quota email addresses are no longer applicable to the quota. For notifications to be sent to the email address generated by the rules specified, you must delete the email address and save the change.

Steps

1. In the left navigation pane, click **Storage** > **SVMs**.
2. In the **Health: All Storage VMs** view, select an SVM and then click the **User and Group Quotas** tab.
3. Click **Edit Email Address** below the row of tabs.
4. In the **Edit Email Address** dialog box, perform the appropriate action:

If…	Then…
You want notifications to be sent to the email address generated by the rules specified	a. Delete the email address in the **Email Address** field. b. Click **Save**. c. Refresh the browser by pressing F5 to reload the Edit Email Address dialog box. The email address generated by the specified rule is displayed in the **Email Address** field.
You want notifications to be sent to a specified email address	a. Modify the email address in the **Email Address** field. b. Click **Save**. The rules to generate the user and group quota email addresses are no longer applicable to the quota.

Understand more about quotas

Learn about the Active IQ Unified Manager quota process

Quotas can be soft or hard. Soft quotas cause ONTAP to send a notification when specified limits are exceeded, and hard quotas prevent a write operation from succeeding when specified limits are exceeded.

When ONTAP receives a request from a user or user group to write to a FlexVol volume, it checks to see whether quotas are activated on that volume for the user or user group and determines the following:

- Whether the hard limit will be reached

  If yes, the write operation fails when the hard limit is reached and the hard quota notification is sent.

- Whether the soft limit will be breached

  If yes, the write operation succeeds when the soft limit is breached and the soft quota notification is sent.

- Whether a write operation will not exceed the soft limit

  If yes, the write operation succeeds and no notification is sent.

Learn about quotas in Active IQ Unified Manager

Quotas provide a way to restrict or track the disk space and number of files used by a user, group, or qtree. You specify quotas using the `/etc/quotas` file. Quotas are applied to a specific volume or qtree.

Learn about quotas in Active IQ Unified Manager

You can use quotas to limit resource usage in FlexVol volumes, to provide notification when resource usage reaches specific levels, or to track resource usage.

You specify a quota for the following reasons:

- To limit the amount of disk space or the number of files that can be used by a user or group, or that can be contained by a qtree
- To track the amount of disk space or the number of files used by a user, group, or qtree, without imposing a limit
- To warn users when their disk usage or file usage is high

Description of quotas dialog boxes

Use the Quotas dialog box in Active IQ Unified Manager

You can use the appropriate option in the User and Group Quotas tab in the Health: All Storage VMs view to configure the format of the email notification that is sent when a quota-related issue occurs and to configure rules to specify email addresses based on the user quota.

Learn about the Email Notification Format page in Active IQ Unified Manager

The Email Notification Format page displays the rules of the email that is sent to a user or a user group when there is a quota-related issue (soft limit breached or hard limit reached).

The email notification is sent only when the following user or user group quota events are generated: User or Group Quota Disk Space Soft Limit Breached, User or Group Quota File Count Soft Limit Breached, User or Group Quota Disk Space Hard Limit Reached, or User or Group Quota File Count Hard Limit Reached.

- **From**

  Displays the email address from which the email is sent, which you can modify. By default, this is the email address that is specified Notifications page.

- **Subject**

  Displays the subject of the notification email.

- **Email Details**

  Displays the text of the notification email. You can modify the text based on your requirements. For example, you can provide information related to the quota attributes and reduce the number of keywords.

However, you should not modify the keywords.

Valid keywords are as follows:

- $EVENT_NAME

  Specifies the event name that caused the email notification.

- $QUOTA_TARGET

  Specifies the qtree or volume on which the quota is applicable.

- $QUOTA_USED_PERCENT

  Specifies the percentage of disk hard limit, disk soft limit, file hard limit, or file soft limit that is used by the user or user group.

- $QUOTA_LIMIT

  Specifies the disk hard limit or file hard limit that is reached by the user or user group and one of the following events is generated:

  - User or Group Quota Disk Space Hard Limit Reached
  - User or Group Quota Disk Space Soft Limit Reached
  - User or Group Quota File Count Hard Limit Reached
  - User or Group Quota File Count Soft Limit Reached

- $QUOTA_USED

  Specifies the disk space used or the number of files created by the user or user group.

- $QUOTA_USER

  Specifies the user or user group name.

Command buttons

The command buttons enable you to preview, save, or cancel the changes made to the email notification format:

- **Preview**

  Displays a preview of the notification email.

- **Restore to Factory Defaults**

  Enables you to restore the notification format to the factory default values.

- **Save**

  Saves the changes made to the notification format.

Learn about quota email address rules in Active IQ Unified Manager

The Rules to Generate User and Group Quota Email Address page enables you to create rules to specify email addresses based on the user quota associated with clusters, SVMs, volumes, qtrees, users, or user groups. A notification is sent to the specified email address when a quota is breached.

Rules area

You must define the rules for a quota email address. You can also add comments to explain the rules.

How you define rules

You must enter the rules in the order in which you want to execute them. If the first rule's criterion is met, then the email address is generated based on this rule. If the criterion is not met, then the criterion for the next rule is considered, and so on. Each line lists a separate rule. The default rule is the last rule in the list. You can change the priority order of rules. However, you cannot change the order of the default rule.

For example, if you want to use the email address qtree1@xyz.com to receive notifications about quota breaches for qtree1 and use the email address admin@xyz.com for all the other qtrees, the rules must be listed in the following order:

- if ($QTREE == 'qtree1') then qtree1@xyz.com
- if ($QTREE == *) then admin@xyz.com

If none of the criteria for the rules you specified are met, then the default rule is used:

if ($USER_OR_GROUP == *) then $USER_OR_GROUP@$DOMAIN

If more than one user has the same quota, the names of the users are displayed as comma-separated values and the rules are not applicable for the quota.

How you add comments

You can add comments to explain the rules. You should use # at the start of each comment and each line lists a separate comment.

Rules syntax

The syntax of the rule must be one of the following:

- if (`valid variableoperator *`) then `email ID@domain name`

  if is a keyword and is in lowercase. The operator is ==. The email ID can contain any character, the valid variables $USER_OR_GROUP, $USER, or $GROUP, or a combination of any character and the valid variables $USER_OR_GROUP, $USER, or $GROUP. The domain name can contain any character, the valid variable $DOMAIN, or a combination of any character and the valid variable $DOMAIN. Valid variables can be in uppercase or lowercase but must not be a combination of both. For example, $domain and $DOMAIN are valid, but $Domain is not a valid variable.

- if (`valid variableoperator` `` `string` ``) then `email ID@domain name`

  if is a keyword and is lowercase. The operator can be contains or ==. The email ID can contain any character, the valid variables $USER_OR_GROUP, $USER, or $GROUP, or a combination of any

character and the valid variables $USER_OR_GROUP, $USER, or $GROUP. The domain name can contain any character, the valid variable $DOMAIN, or a combination of any character and the valid variable $DOMAIN. Valid variables can be in uppercase or lowercase but must not be a combination of both. For example, $domain and $DOMAIN are valid, but $Domain is not a valid variable.

Command buttons

The command buttons enable you to save, validate, or cancel the created rules:

- **Validate**

  Validates the syntax of the created rule. If there are errors during validation, the rule that generates the error is displayed along with an error message.

- **Restore to Factory Defaults**

  Enables you to restore the address rules to the factory default values.

- **Save**

  Validates the syntax of the rule and saves the rule if there are no errors. If there are errors during validation, the rule that generates the error is displayed along with an error message.

Troubleshoot

Add disk space to the Unified Manager database directory

Add disk space to the database directory in Active IQ Unified Manager

The Unified Manager database directory contains all of the health and performance data collected from ONTAP systems. Some circumstances may require that you increase the size of the database directory.

For example, the database directory may get full if Unified Manager is collecting data from a large number of clusters where each cluster has many nodes. You will receive a warning event when the database directory is 90% full, and a critical event when the directory is 95% full.

No additional data is collected from clusters after the directory reaches 95% full.

The steps required to add capacity to the data directory are different depending on whether Unified Manager is running on a VMware ESXi server, on a Red Hat Linux server, or on a Microsoft Windows server.

Add space to the data disk of the VMware virtual machine in Active IQ Unified Manager

If you need to increase the amount of space on the data disk for the Unified Manager database, you can add capacity after installation by increasing disk space using the Unified Manager maintenance console.

Before you begin

- You must have access to the vSphere Client.
- The virtual machine must have no snapshots stored locally.

- You must have the maintenance user credentials.

We recommend that you back up your virtual machine before increasing the size of virtual disks.

Steps

1. In the vSphere client, select the Unified Manager virtual machine, and then add more disk capacity to data `disk 3`. See the VMware documentation for details.

   In some rare cases the Unified Manager deployment uses "Hard Disk 2" for the data disk instead of "Hard Disk 3". If this has occurred in your deployment, increase the space of whichever disk is larger. The data disk will always have more space than the other disk.

2. In the vSphere client, select the Unified Manager virtual machine, and then select the **Console** tab.
3. Click in the console window, and then log in to the maintenance console using your user name and password.
4. In the **Main Menu**, enter the number for the **System Configuration** option.
5. In the **System Configuration Menu**, enter the number for the **Increase Data Disk Size** option.

Add space to the the Linux host's data directory in Active IQ Unified Manager

If you allotted insufficient disk space to the `/opt/netapp/data` directory to support Unified Manager when you originally set up the Linux host and then installed Unified Manager, you can add disk space after installation by increasing disk space on the `/opt/netapp/data` directory.

Before you begin

You must have root user access to the Red Hat Enterprise Linux machine on which Unified Manager is installed.

We recommend that you back up the Unified Manager database before increasing the size of the data directory.

Steps

1. Log in as root user to the Linux machine on which you want to add disk space.
2. Stop the Unified Manager service and the associated MySQL software in the order shown: `systemctl stop ocieau ocie mysqld`
3. Create a temporary backup folder (for example, `/backup-data`) with sufficient disk space to contain the data in the current `/opt/netapp/data` directory.
4. Copy the content and privilege configuration of the existing `/opt/netapp/data` directory to the backup data directory:

   ```
   cp -arp /opt/netapp/data/* /backup-data
   ```

5. If SE Linux is enabled:
   a. Get the SE Linux type for folders on existing `/opt/netapp/data` folder:

      ```
      se_type=`ls -Z /opt/netapp/data| awk '{print $4}'| awk -F: '{print $3}'|
      head -1
      ```

The system returns a confirmation similar to the following:

```
echo $se_type
mysqld_db_t
```

b. Run the `chcon` command to set the SE Linux type for the backup directory:

```
chcon -R --type=mysqld_db_t /backup-data
```

6. Remove the contents of the `/opt/netapp/data` directory:
   a. `cd /opt/netapp/data`
   b. `rm -rf *`
7. Expand the size of the `/opt/netapp/data` directory to a minimum of 150 GB through LVM commands or by adding extra disks.

> If you have created `/opt/netapp/data` from a disk, then you should not try to mount `/opt/netapp/data` as an NFS or CIFS share. Because, in this case, if you try to expand the disk space, some LVM commands, such as `resize` and `extend` might not work as expected.

8. Confirm that the `/opt/netapp/data` directory owner (mysql) and group (root) are unchanged:

```
ls -ltr /opt/netapp/ | grep data
```

The system returns a confirmation similar to the following:

```
drwxr-xr-x.  17 mysql root  4096 Aug 28 13:08 data
```

9. If SE Linux is enabled, confirm that the context for the `/opt/netapp/data` directory is still set to mysqld_db_t:
   a. `touch /opt/netapp/data/abc`
   b. `ls -Z /opt/netapp/data/abc`

The system returns a confirmation similar to the following:

```
-rw-r--r--. root root unconfined_u:object_r:mysqld_db_t:s0
/opt/netapp/data/abc
```

10. Delete the file `abc` so that this extraneous file does not cause a database error in the future.
11. Copy the contents from `backup-data` back to the expanded `/opt/netapp/data` directory:

```
cp -arp /backup-data/* /opt/netapp/data/
```

12. If SE Linux is enabled, run the following command:

```
chcon -R --type=mysqld_db_t /opt/netapp/data
```

13. Start the MySQL service:

```
systemctl start mysqld
```

14. After the MySQL service is started, start the ocie and ocieau services in the order shown:

```
systemctl start ocie ocieau
```

15. After all of the services are started, delete the backup folder `/backup-data`:

```
rm -rf /backup-data
```

Add space to the logical drive of the Microsoft Windows server in Active IQ Unified Manager

If you need to increase the amount of disk space for the Unified Manager database, you can add capacity to the logical drive on which Unified Manager is installed.

Before you begin

You must have Windows administrator privileges.

We recommend that you back up the Unified Manager database before adding disk space.

Steps

1. Log in as administrator to the Windows server on which you want to add disk space.
2. Follow the step that corresponds to method you want to use to add more space:

Option	Description
On a physical server, add capacity to the logical drive on which the Unified Manager server is installed.	Follow the steps in the Microsoft topic: Extend a Basic Volume
On a physical server, add a hard disk drive.	Follow the steps in the Microsoft topic: Adding Hard Disk Drives
On a virtual machine, increase the size of a disk partition.	Follow the steps in the VMware topic: Increasing the size of a disk partition

Change the performance statistics collection interval in Active IQ Unified Manager

The default collection interval for performance statistics is 5 minutes. You can change this interval to 10 or 15 minutes if you find that collections from large clusters are not finishing within the default time. This setting affects the collection of statistics from all clusters that this instance of Unified Manager is monitoring.

Before you begin

You must have a user ID and password authorized to log in to the maintenance console of the Unified Manager

server.

The issue of performance statistics collections not finishing on time is indicated by the banner messages `Unable to consistently collect from cluster <cluster_name> or Data collection is taking too long on cluster <cluster_name>.`

You should change the collection interval only when required because of a statistics collections issue. Do not change this setting for any other reason.

> Changing this value from the default setting of 5 minutes can affect the number and frequency of performance events that Unified Manager reports. For example, system-defined performance thresholds trigger events when the policy is exceeded for 30 minutes. When using 5-minute collections, the policy must be exceeded for six consecutive collections. For 15-minute collections the policy must be exceeded for only two collection periods.

A message at the bottom of the Cluster Setup page indicates the current statistical data collection interval.

Steps

1. Log in using SSH as the maintenance user to the Unified Manager host.

   The Unified Managermaintenance console prompts are displayed.

2. Type the number of the menu option labeled **Performance Polling Interval Configuration**, and then press Enter.
3. If prompted, enter the maintenance user password again.
4. Type the number for the new polling interval that you want to set, and then press Enter.

If you changed the Unified Manager collection interval to 10 or 15 minutes, and you have a current connection to an external data provider (such as Graphite), you must change the data provider transmit interval so that it is equal to, or greater, than the Unified Manager collection interval.

Change the length of time Active IQ Unified Manager retains event and performance data

By default, Unified Manager stores event data and performance data for 6 months for all monitored clusters. After this time, older data is automatically deleted to make room for new data. This default timeframe works well for most configurations, but very large configurations with many clusters and nodes may need to reduce the retention period so that Unified Manager operates optimally.

Before you begin

You must have the Application Administrator role.

You can change the retention periods for these two types of data in the Data Retention page. These settings affect the retention of data from all clusters that this instance of Unified Manager is monitoring.

> Unified Manager collects performance statistics every 5 minutes. Each day the 5-minute statistics are summarized into hourly performance statistics. It retains 30 days of 5-minute historical performance data and 6 months of hourly summarized performance data (by default).

You should reduce the retention period only if you are running out of space or if backup and other operations

are taking a very long time to complete. Reducing the retention period has the following effects:

- Old performance data is deleted from the Unified Manager database after midnight.
- Old event data is deleted from the Unified Manager database immediately.
- Events prior to the retention period will no longer be available to view in the user interface.
- Locations in the UI where hourly performance statistics are displayed will be blank prior to the retention period.
- If the event retention period exceeds the performance data retention period, a message will be displayed under the performance slider warning that older performance events may not have backing data in their associated charts.

Steps

1. In the left navigation pane, click **Policies** > **Data Retention**.
2. In the **Data Retention** page, select the slider tool in the Event Retention or Performance Data Retention area and move it to the number of months that data should be retained, and click **Save**.

Resolve an unknown authentication error in Active IQ Unified Manager

When you are performing an authentication-related operation such as adding, editing, deleting, or testing remote users or groups, the following error message might be displayed: `Unknown authentication error.`

Cause

This problem can occur if you have set an incorrect value for the following options:

- Administrator Name of the Active Directory authentication service
- Bind Distinguished Name of the OpenLDAP authentication service

Corrective action

a. In the left navigation pane, click **General** > **Remote Authentication**.
b. Based on the authentication service that you have selected, enter the appropriate information for Administrator Name or Bind Distinguished Name.
c. Click **Test Authentication** to test the authentication with the details that you specified.
d. Click **Save**.

Resolve a User not found error in Active IQ Unified Manager

When you are performing an authentication-related operation such as adding, editing, deleting, or testing remote users or groups, the following error message is displayed: `User not found.`

Cause

This problem can occur if the user exists in the AD server or LDAP server, and if you have set the base distinguished name to an incorrect value.

Corrective action

a. In the left navigation pane, click **General** > **Remote Authentication**.
b. Enter the appropriate information for base distinguished name.
c. Click **Save**.

Resolve the issue caused by adding LDAP using Other authentication service in Active IQ Unified Manager

When you select Others as the Authentication service, the user and group Object Class retain the values from the previously selected template. If the LDAP server does not use the same values, the operation might fail.

Cause

The users are not configured correctly in OpenLDAP.

Corrective action

You can manually fix this issue by using one of the following workarounds.

If your LDAP user object class and group object class are user and group, respectively, perform the following steps:

a. In the left navigation pane, click**General** > **Remote Authentication**.
b. In the **Authentication Service** drop-down menu, select **Active Directory**, and then select **Others**.
c. Complete the text fields.

If your LDAP user object class and group object class are posixAccount and posixGroup, respectively, perform the following steps:

a. In the left navigation pane, click **General** > **Remote Authentication**.
b. In the **Authentication Service** drop-down menu, select **OpenLDAP**, and then select **Others**.
c. Complete the text fields.

If the first two workarounds do not apply, call the `option-set` API, and set the `auth.ldap.userObjectClass` and `auth.ldap.groupObjectClass` options to the correct values.

Resolve a NetApp Manageability SDK log rotation issue in Active IQ Unified Manager

After you add an ONTAPI API based cluster to Unified Manager on a Windows operating system, the `nmsdk.log` file size increases and exceeds the 10 MB size limit.

Cause

This problem can happen if log rotation does not occur.

Corrective action

a. Stop Unified Manager.

b. When you are installing Unified Manager on Windows, install Logrotate version 0.0.0.18. For more information, see the Security hardening guide for NetApp Manageability SDK technical report.

c. Start Unified Manager.

Manage events and alerts

Manage events

Learn about Active IQ platform events in Active IQ Unified Manager

Unified Manager can display events that have been discovered by the Active IQ platform. These events are created by running a set of rules against AutoSupport messages generated from all storage systems being monitored by Unified Manager.

For more information, see How Active IQ platform events are generated.

Unified Manager checks for a new rules file automatically and only downloads a new file when there are newer rules. In sites with no external network access, you need to upload the rules manually from **Storage Management** > **Event Setup** > **Upload Rules**.

These Active IQ events do not overlap with existing Unified Manager events, and they identify incidents or risks concerning system configuration, cabling, best practice, and availability issues.

For more information about enabling platform events, see Enabling Active IQ portal events.
For more information about uploading rules file, see Uploading a new Active IQ rules file.

NetApp Active IQ is a cloud based service that provides predictive analytics and proactive support to optimize storage system operations across the NetApp hybrid cloud. See NetApp Active IQ for more information.

What Event Management System events are

Learn about Event Management System events in Active IQ Unified Manager

The Event Management System (EMS) collects event data from different parts of the ONTAP kernel and provides event forwarding mechanisms. These ONTAP events can be reported as EMS events in Unified Manager. Centralized monitoring and management eases configuration of critical EMS events and alert notifications based on these EMS events.

The Unified Manager address is added as a notification destination to the cluster when you add the cluster to Unified Manager. An EMS event is reported as soon as the event occurs in the cluster.

There are two methods for receiving EMS events in Unified Manager:

- A certain number of important EMS events are reported automatically.
- You can subscribe to receive individual EMS events.

The EMS events that are generated by Unified Manager are reported differently depending on the method in which the event was generated:

Functionality	Automatic EMS messages	Subscribed EMS messages
Available EMS events	Subset of EMS events	All EMS events

Functionality	Automatic EMS messages	Subscribed EMS messages
EMS message name when triggered	Unified Manager event name (converted from EMS event name)	Non-specific in the format "Error EMS received". The detailed message provides the dot-notation format of the actual EMS event
Messages received	As soon as the cluster has been discovered	After adding each required EMS event to Unified Manager, and after the next 15 minute polling cycle
Event life cycle	Same as other Unified Manager events: New, Acknowledged, Resolved, and Obsolete states	The EMS event is made obsolete after the cluster is refreshed, after 15 minutes, from when the event was created
Captures events during Unified Manager downtime	Yes, when the system starts up it communicates with each cluster to acquire missing events	No
Event details	Suggested corrective actions are imported directly from ONTAP to provide consistent resolutions	Corrective actions not available in Event Details page

Some of the new automatic EMS events are Informational events that indicate that a previous event has been resolved. For example, the "FlexGroup Constituents Space Status All OK" Informational event indicates that the "FlexGroup Constituents Have Space Issues" Error event has been resolved. Informational events cannot be managed using the same event life cycle as other event severity types, however, the event is obsoleted automatically if the same volume receives another “Space Issues” Error event.

EMS events that are added automatically to Active IQ Unified Manager

The following ONTAP EMS events are added automatically to Unified Manager. These events will be generated when triggered on any cluster that Unified Manager is monitoring.

The following EMS events are available when monitoring clusters running ONTAP 9.5 or greater software:

Unified Manager Event name	EMS Event name	Affected resource	Unified Manager severity
Cloud Tier Access Denied for Aggregate Relocation	arl.netra.ca.check.failed	Aggregate	Error
Cloud Tier Access Denied for Aggregate Relocation During Storage Failover	gb.netra.ca.check.failed	Aggregate	Error
FabricPool Mirror Replication Resync Completed	wafl.ca.resync.complete	Cluster	Error
FabricPool Space Nearly Full	fabricpool.nearly.full	Cluster	Error

Unified Manager Event name	EMS Event name	Affected resource	Unified Manager severity
NVMe-oF Grace Period Started	nvmf.graceperiod.start	Cluster	Warning
NVMe-oF Grace Period Active	nvmf.graceperiod.active	Cluster	Warning
NVMe-oF Grace Period Expired	nvmf.graceperiod.expired	Cluster	Warning
LUN Destroyed	lun.destroy	LUN	Information
Cloud AWS MetaDataConnFail	cloud.aws.metadataConnFail	Node	Error
Cloud AWS IAMCredsExpired	cloud.aws.iamCredsExpired	Node	Error
Cloud AWS IAMCredsInvalid	cloud.aws.iamCredsInvalid	Node	Error
Cloud AWS IAMCredsNotFound	cloud.aws.iamCredsNotFound	Node	Error
Cloud AWS IAMCredsNotInitialized	cloud.aws.iamNotInitialized	Node	Information
Cloud AWS IAMRoleInvalid	cloud.aws.iamRoleInvalid	Node	Error
Cloud AWS IAMRoleNotFound	cloud.aws.iamRoleNotFound	Node	Error
Cloud Tier Host Unresolvable	objstore.host.unresolvable	Node	Error
Cloud Tier Intercluster Network Interface Down	objstore.interclusterlifDown	Node	Error
Request Mismatch Cloud Tier Signature	osc.signatureMismatch	Node	Error
One of NFSv4 Pools Exhausted	Nblade.nfsV4PoolExhaust	Node	Critical
QoS Monitor Memory Maxed	qos.monitor.memory.maxed	Node	Error
QoS Monitor Memory Abated	qos.monitor.memory.abated	Node	Information
NVMeNS Destroy	NVMeNS.destroy	Namespace	Information
NVMeNS Online	NVMeNS.offline	Namespace	Information
NVMeNS Offline	NVMeNS.online	Namespace	Information
NVMeNS Out of Space	NVMeNS.out.of.space	Namespace	Warning
Synchronous Replication Out Of Sync	sms.status.out.of.sync	SnapMirror relationship	Warning
Synchronous Replication Restored	sms.status.in.sync	SnapMirror relationship	Information

Unified Manager Event name	EMS Event name	Affected resource	Unified Manager severity
Synchronous Replication Auto Resync Failed	sms.resync.attempt.failed	SnapMirror relationship	Error
Many CIFS Connections	Nblade.cifsManyAuths	SVM	Error
Max CIFS Connection Exceeded	Nblade.cifsMaxOpenSameFile	SVM	Error
Max Number of CIFS Connection Per User Exceeded	Nblade.cifsMaxSessPerUsrConn	SVM	Error
CIFS NetBIOS Name Conflict	Nblade.cifsNbNameConflict	SVM	Error
Attempts to Connect Nonexistent CIFS Share	Nblade.cifsNoPrivShare	SVM	Critical
CIFS Shadow Copy Operation Failed	cifs.shadowcopy.failure	SVM	Error
Virus Found By AV Server	Nblade.vscanVirusDetected	SVM	Error
No AV Server Connection for Virus Scan	Nblade.vscanNoScannerConn	SVM	Critical
No AV Server Registered	Nblade.vscanNoRegdScanner	SVM	Error
No Responsive AV Server Connection	Nblade.vscanConnInactive	SVM	Information
AV Server too Busy to Accept New Scan Request	Nblade.vscanConnBackPressure	SVM	Error
Unauthorized User Attempt to AV Server	Nblade.vscanBadUserPrivAccess	SVM	Error
FlexGroup Constituents Have Space Issues	flexgroup.constituents.have.space.issues	Volume	Error
FlexGroup Constituents Space Status All OK	flexgroup.constituents.space.status.all.ok	Volume	Information
FlexGroup Constituents Have Inodes Issues	flexgroup.constituents.have.inodes.issues	Volume	Error
FlexGroup Constituents Inodes Status All OK	flexgroup.constituents.inodes.status.all.ok	Volume	Information
Volume Logical Space Nearly Full	monitor.vol.nearFull.inc.sav	Volume	Warning
Volume Logical Space Full	monitor.vol.full.inc.sav	Volume	Error
Volume Logical Space Normal	monitor.vol.one.ok.inc.sav	Volume	Information

Unified Manager Event name	EMS Event name	Affected resource	Unified Manager severity
WAFL Volume AutoSize Fail	wafl.vol.autoSize.fail	Volume	Error
WAFL Volume AutoSize Done	wafl.vol.autoSize.done	Volume	Information
WAFL READDIR File Operation Timeout	wafl.readdir.expired	Volume	Error

Subscribe to ONTAP EMS events in Active IQ Unified Manager

You can subscribe to receive Event Management System (EMS) events that are generated by systems that are installed with ONTAP software. A subset of EMS events are reported to Unified Manager automatically, but additional EMS events are reported only if you have subscribed to these events.

Before you begin

Do not subscribe to EMS events that are already added to Unified Manager automatically as this can cause confusion when receiving two events for the same issue.

You can subscribe to any number of EMS events. All the events to which you subscribe are validated, and only the validated events are applied to the clusters you are monitoring in Unified Manager. The *ONTAP 9 EMS Event Catalog* provides detailed information for all of the EMS messages for the specified version of ONTAP 9 software. Locate the appropriate version of the *EMS Event Catalog* from the ONTAP 9 Product Documentation page for a list of the applicable events.

ONTAP 9 Product Library

You can configure alerts for the ONTAP EMS events to which you subscribe, and you can create custom scripts to be executed for these events.

If you do not receive the ONTAP EMS events to which you have subscribed, there might be an issue with the DNS configuration of the cluster which is preventing the cluster from reaching the Unified Manager server. To resolve this issue, the cluster administrator must correct the DNS configuration of the cluster, and then restart Unified Manager. Doing so will flush the pending EMS events to the Unified Manager server.

Steps

1. In the left navigation pane, click **Storage Management** > **Event Setup**.
2. In the Event Setup page, click the **Subscribe to EMS events** button.
3. In the Subscribe to EMS events dialog box, enter the name of the ONTAP EMS event to which you want to subscribe.

   To view the names of the EMS events to which you can subscribe, from the ONTAP cluster shell, you can use the `event route show` command (prior to ONTAP 9) or the `event catalog show` command (ONTAP 9 or later).

   How to configure and receive alerts from ONTAP EMS Event Subscription in Active IQ Unified Manager

4. Click **Add**.

The EMS event is added to the Subscribed EMS events list, but the Applicable to Cluster column displays the status as “Unknown” for the EMS event that you added.

5. Click **Save and Close** to register the EMS event subscription with the cluster.
6. Click **Subscribe to EMS events** again.

   The status “Yes” appears in the Applicable to Cluster column for the EMS event that you added.

   If the status is not “Yes”, check the spelling of the ONTAP EMS event name. If the name is entered incorrectly, you must remove the incorrect event, and then add the event again.

When the ONTAP EMS event occurs, the event is displayed on the Events page. You can select the event to view details about the EMS event in the Event details page. You can also manage the disposition of the event or create alerts for the event.

Understand how Active IQ Unified Manager handles a received event

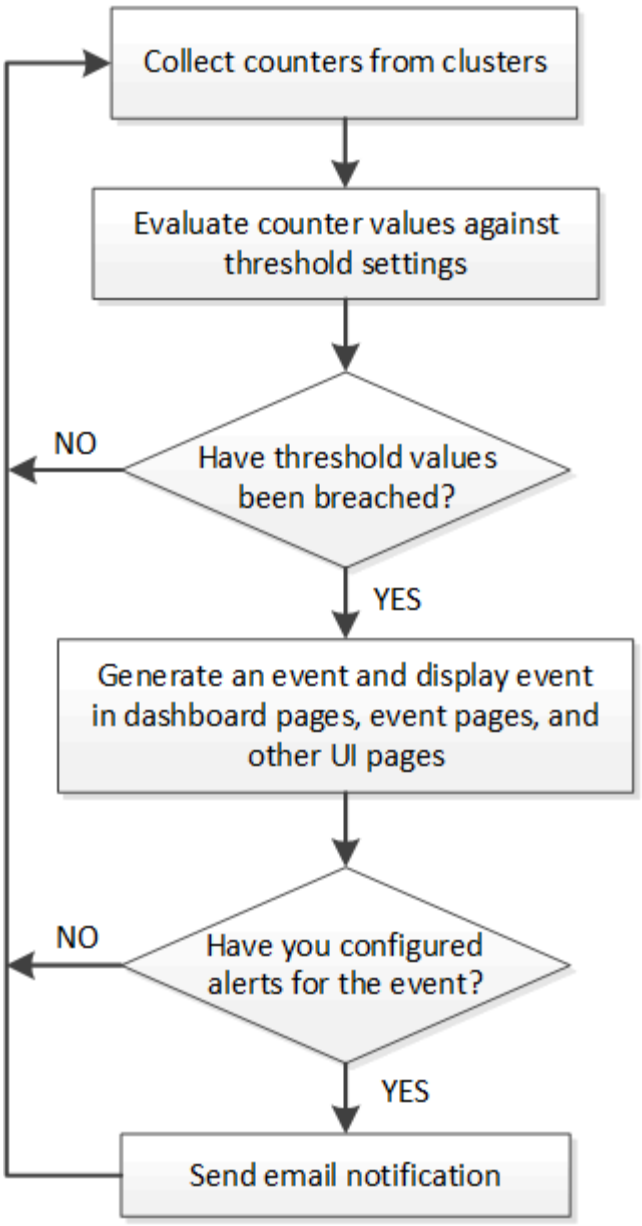
When Unified Manager receives an event, it is displayed in the Dashboard page, in the Event Management inventory page, in the Summary and Explorer tabs of the Cluster/Performance page, and in the object-specific inventory page (for example, the Volumes/Health inventory page).

When Unified Manager detects multiple continuous occurrences of the same event condition for the same cluster component, it treats all occurrences as a single event, not as separate events. The duration of the event is incremented to indicate that the event is still active.

Depending on how you configure settings in the Alert Setup page, you can notify other users about these events. The alert causes the following actions to be initiated:

- An email about the event can be sent to all Unified Manager Administrator users.
- The event can be sent to additional email recipients.
- An SNMP trap can be sent to the trap receiver.
- A custom script can be executed to perform an action.

This workflow is shown in the following diagram.



View event details in Active IQ Unified Manager

You can view details about an event that is triggered by Unified Manager to take corrective action. For example, if there is a health event Volume Offline, you can click that event to view the details and perform corrective actions.

Before you begin

You must have the Operator, Application Administrator, or Storage Administrator role.

The event details include information such as the source of the event, cause of the event, and any notes related to the event.

Steps

1. In the left navigation pane, click **Event Management**.

   By default, the All active events view displays the New and Acknowledged (active) events that have been generated over the previous 7 days that have an Impact Level of Incident or Risk.

2. If you want to view a particular category of events, for example, capacity events or performance events, click **View** and select from the menu of event types.
3. Click the event name for which you want to view the details.

   The event details are displayed in the Event details page.

View unassigned events in Active IQ Unified Manager

You can view unassigned events and then assign each of them to a user who can resolve them.

Before you begin

You must have the Operator, Application Administrator, or Storage Administrator role.

Steps

1. In the left navigation pane, click **Event Management**.

   By default, New and Acknowledged events are displayed on the Event Management inventory page.

2. From the **Filters** pane, select the **Unassigned** filter option in the **Assigned To** area.

Acknowledge and resolve events in Active IQ Unified Manager

You should acknowledge an event before you start working on the issue that generated the event so that you do not continue to receive repeat alert notifications. After you take corrective action for a particular event, you should mark the event as resolved.

Before you begin

You must have the Operator, Application Administrator, or Storage Administrator role.

You can acknowledge and resolve multiple events simultaneously.

You cannot acknowledge Information events.

Steps

1. In the left navigation pane, click **Event Management**.
2. From the events list, perform the following actions to acknowledge the events:

If you want to…	Do this…
Acknowledge and mark a single event as resolved	a. Click the event name. b. From the Event details page, determine the cause of the event. c. Click **Acknowledge**. d. Take appropriate corrective action. e. Click **Mark As Resolved**.

If you want to…	Do this…
Acknowledge and mark multiple events as resolved	a. Determine the cause of the events from the respective Event details page. b. Select the events. c. Click **Acknowledge**. d. Take appropriate corrective actions. e. Click **Mark As Resolved**.

After the event is marked resolved, the event is moved to the resolved events list.

3. **Optional**: In the **Notes and Updates** area, add a note about how you addressed the event, and then click **Post**.

Assign events to specific users in Active IQ Unified Manager

You can assign unassigned events to yourself or to other users, including remote users. You can reassign assigned events to another user, if required. For example, when frequent issues occur on a storage object, you can assign the events for these issues to the user who manages that object.

Before you begin

- The user's name and email ID must be configured correctly.
- You must have the Operator, Application Administrator, or Storage Administrator role.

Steps

1. In the left navigation pane, click **Event Management**.
2. In the **Event Management** inventory page, select one or more events that you want to assign.
3. Assign the event by choosing one of the following options:

If you want to assign the event to…	Then do this…
Yourself	Click **Assign To** > **Me**.
Another user	a. Click **Assign To** > **Another user**. b. In the Assign Owner dialog box, enter the user name, or select a user from the drop-down list. c. Click **Assign**. An email notification is sent to the user. If you do not enter a user name or select a user from the drop-down list, and click **Assign**, the event remains unassigned.

Disable unwanted events in Active IQ Unified Manager

All events are enabled by default. You can disable events globally to prevent the generation of notifications for events that are not important in your environment. You can enable events that are disabled when you want to resume receiving notifications for them.

Before you begin

You must have the Application Administrator or Storage Administrator role.

When you disable events, the previously generated events in the system are marked obsolete, and the alerts that are configured for these events are not triggered. When you enable events that are disabled, the notifications for these events are generated starting with the next monitoring cycle.

When you disable an event for an object (for example, the `vol offline` event), and then later you enable the event, Unified Manager does not generate new events for objects that went offline when the event was in the disabled state. Unified Manager generates a new event only when there is a change in the object state after the event was re-enabled.

Steps

1. In the left navigation pane, click **Storage Management** > **Event Setup**.
2. In the **Event Setup** page, disable or enable events by choosing one of the following options:

If you want to…	Then do this…
Disable events	a. Click **Disable**. b. In the Disable Events dialog box, select the event severity. c. In the Matching Events column, select the events that you want to disable based on the event severity, and then click the right arrow to move those events to the Disable Events column. d. Click **Save and Close**. e. Verify that the events that you disabled are displayed in the list view of the Event Setup page.
Enable events	a. Select the check box for the event, or events, that you want to enable. b. Click **Enable**.

Fix issues using automatic remediation in Active IQ Unified Manager

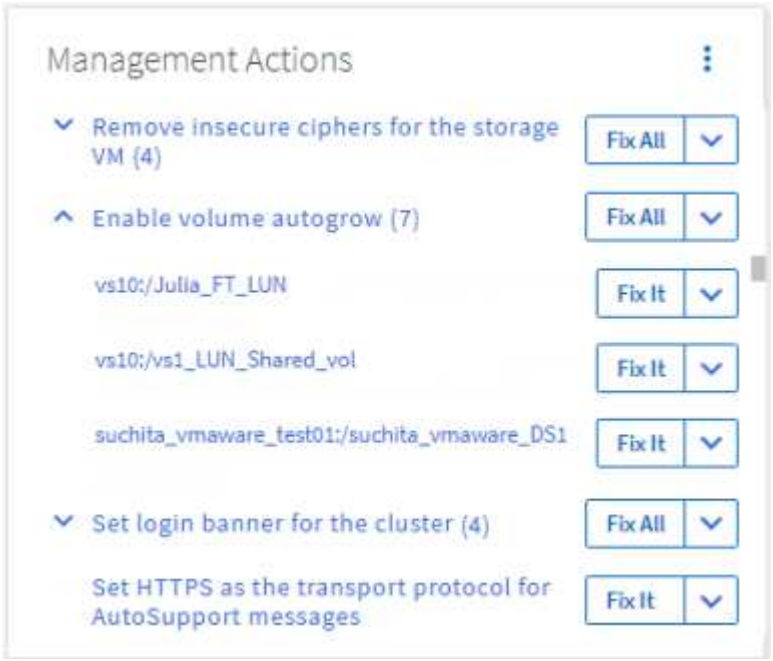
There are certain events that Unified Manager can diagnose thoroughly and provide a single resolution using the **Fix It** button. When available, those resolutions are displayed in the Dashboard, from the Event details page, and from the Workload Analysis selection on the left-navigation menu.

Most events have a variety of possible resolutions that are displayed in the Event details page so you can implement the best solution using ONTAP System Manager or the ONTAP CLI. A **Fix It** action is available

when Unified Manager has determined that there is a single resolution to fix the issue, and that it can be resolved using an ONTAP CLI command.

Steps

1. To view events that can be fixed from the **Dashboard**, click **Dashboard**.



2. To resolve any of the issues that Unified Manager can fix, click the **Fix It** button. To fix an issue that exists on multiple objects, click the **Fix All** button.

For information about the issues that can be fixed by automatic remediation, see What issues can Unified Manager fix.

Enable and disable Active IQ event reporting in Active IQ Unified Manager

Active IQ platform events are generated and displayed in the Unified Manager user interface by default. If you find that these events are too "noisy", or that you do not want to view these events in Unified Manager, then you can disable these events from being generated. You can enable them at a later time if you want to resume receiving these notifications.

Before you begin

You must have the Application Administrator role.

When you disable this feature, Unified Manager stops receiving Active IQ platform events immediately.

When you enable this feature, Unified Manager starts receiving Active IQ platform events shortly after midnight based on the timezone of the cluster. The start time is based on when Unified Manager receives AutoSupport messages from each cluster.

Steps

1. In the left navigation pane, click **General** > **Feature Settings**.
2. In the **Feature Settings** page, disable or enable Active IQ platform events by choosing one of the following

options:

If you want to…	Then do this…
Disable Active IQ platform events	In the **Active IQ Portal Events** panel, move the slider button to the left.
Enable Active IQ platform events	In the **Active IQ Portal Events** panel, move the slider button to the right.

Upload a new Active IQ rules file in Active IQ Unified Manager

Unified Manager checks for a new Active IQ events (rules) file automatically and downloads a new file when there are newer rules.
However, in sites with no external network access, you need to upload the rules file manually.

Active IQ rules are also referred to as Config Advisor (CA) secure rules.

When you are installing or upgrading Unified Manager to a specific release in a site without network connectivity, the bundled Active IQ rules are automatically available for upload. However, it is recommended that you download a new rules file approximately once a month from NetApp's Support Site to make sure that updated events are generated and your storage systems continue to perform optimally.

Before you begin

- Active IQ portal event reporting must be enabled. This feature is enabled by default. For information, see Enabling Active IQ portal events.
- You must download the rules file from the NetApp Support Site.

  The rules file is located at: https://mysupport.netapp.com/api/content-service/staticcontents/content/public/tools/unifiedmanager/ca/secure_rules.zip

Steps

1. On a computer that has network access, navigate to the NetApp Support Site and download the current rules `.zip` file.

   The bundled rules package includes the rules repository, data sources, and a NetApp KB article.

   On Windows systems, in a site without network connectivity, the NetApp KB article is not bundled by default with the installer. You can download the *secure_rules.zip* file from the support site and upload it to view the KB article for all the rules.

2. Transfer the rules file to some media that you can bring into the secure area and then copy it onto a system in the secure area.
3. In the left navigation pane, click **Storage Management** > **Event Setup**.
4. In the **Event Setup** page, click the **Upload Rules** button.
5. In the **Upload Rules** dialog box, navigate to and select the rules `.zip` file you downloaded and click **Upload**.

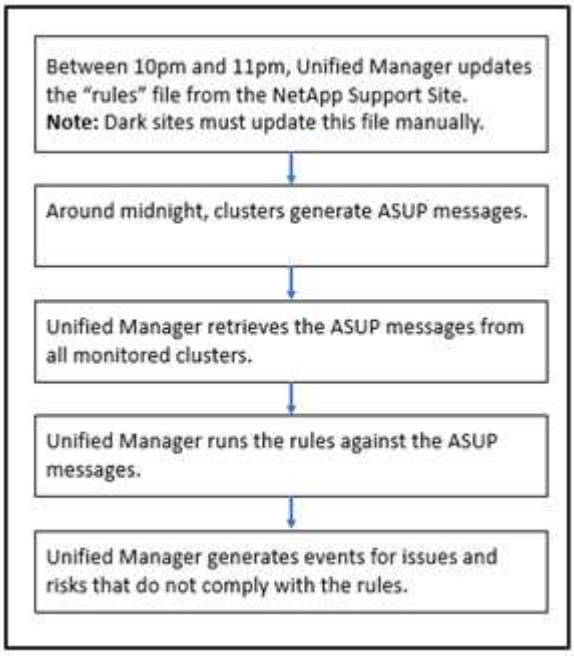
   This process can take a few minutes.

The rules file is unzipped on the Unified Manager server. After your managed clusters generate an AutoSupport file after midnight, Unified Manager checks the clusters against the rules file and generates new risk and incident events as required.

For more information, see this Knowledge Base (KB) article: How to update AIQCASecure rules manually in Active IQ Unified Manager.

How Active IQ Unified Manager generates Active IQ platform events

Active IQ platform incidents and risks are converted to Unified Manager events as shown in the following diagram.

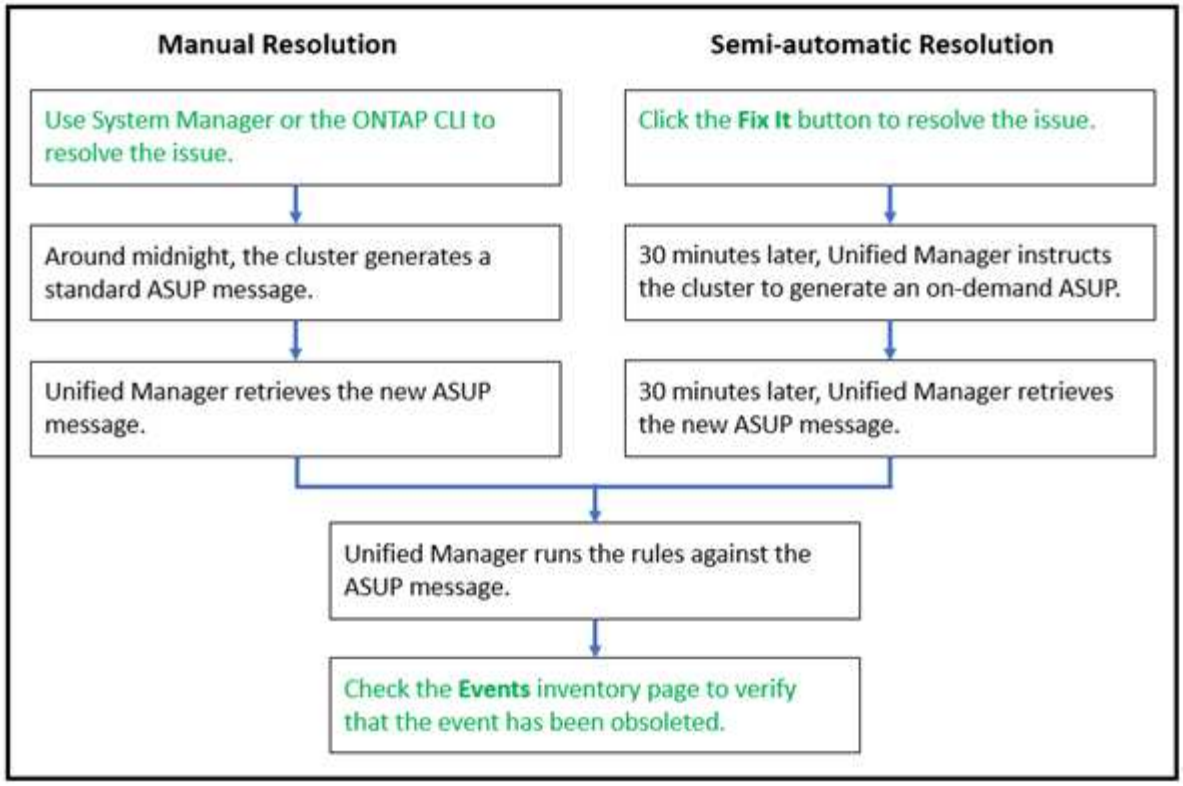


As you can see, the rules file that is compiled on the Active IQ platform is kept current, cluster AutoSupport messages are generated daily, and Unified Manager updates the list of events on a daily basis.

Resolve Active IQ platform events in Active IQ Unified Manager

Active IQ platform incidents and risks are similar to other Unified Manager events because they can be assigned to other users for resolution and they have the same available states. However, when you resolve these types of events using the **Fix It** button you can verify the resolution within hours.

The following diagram shows the actions you must take (in green) and the action that Unified Manager takes (in black) when resolving events that were generated from the Active IQ platform.



When performing a manual resolution you must log into System Manager or the ONTAP command-line interface to fix the issue. You will be able to verify the issue only after the cluster generates a new AutoSupport message at midnight.

When performing a semi-automatic resolution using the **Fix It** button you are able to verify that the fix was successful within hours.

Configure event retention settings in Active IQ Unified Manager

You can specify the number of months an event is retained in the Unified Manager server before it is automatically deleted.

Before you begin

You must have the Application Administrator role.

Retaining events for more than 6 months could affect the server performance and is not recommended.

Steps

1. In the left navigation pane, click **General** > **Data Retention**.
2. In the **Data Retention** page, select the slider in the Event Retention area and move it to the number of months that events should be retained, and click **Save**.

What a Unified Manager maintenance window is

Learn about Active IQ Unified Manager maintenance windows

You define a Unified Manager maintenance window to suppress events and alerts for a specific timeframe when you have scheduled cluster maintenance and you do not want to

receive a flood of unwanted notifications.

When the maintenance window starts, an "Object Maintenance Window Started" event is posted to the Event Management inventory page. This event is obsoleted automatically when the maintenance window ends.

During a maintenance window the events related to all objects on that cluster are still generated, but they do not appear in any of the UI pages, and no alerts or other types of notification are sent for these events. You can, however, view the events that were generated for all storage objects during a maintenance window by selecting one of the View options on the Event Management inventory page.

You can schedule a maintenance window to be initiated in the future, you can change the start and end times for a scheduled maintenance window, and you can cancel a scheduled maintenance window.

Schedule a maintenance window to disable cluster event notifications in Active IQ Unified Manager

If you have a planned downtime for a cluster, for example, to upgrade the cluster or to move one of the nodes, you can suppress the events and alerts that would normally be generated during that timeframe by scheduling a Unified Manager maintenance window.

Before you begin

You must have the Application Administrator or Storage Administrator role.

During a maintenance window, the events related to all objects on that cluster are still generated, but they do not appear in the event page, and no alerts or other types of notification are sent for these events.

The time you enter for the maintenance window is based on the time at the Unified Manager server.

Steps

1. In the left navigation pane, click **Storage Management** > **Cluster Setup**.
2. In the **Maintenance Mode** column for the cluster, select the slider button and move it to the right.

   The calendar window is displayed.

3. Select the start and end date and time for the maintenance window and click **Apply**.

   The message "Scheduled" appears next to the slider button.

When the start time is reached the cluster goes into maintenance mode and an "Object Maintenance Window Started" event is generated.

Change or cancel a scheduled maintenance window in Active IQ Unified Manager

If you have configured a Unified Manager maintenance window to occur in the future, you can change the start and end times or cancel the maintenance window from occurring.

Before you begin

You must have the Application Administrator or Storage Administrator role.

Canceling a currently running maintenance window is useful if you have completed cluster maintenance before the scheduled maintenance window end time and you want to start receiving events and alerts from the cluster again.

Steps

1. In the left navigation pane, click **Storage Management** > **Cluster Setup**.
2. In the **Maintenance Mode** column for the cluster:

If you want to…	Perform this step…
Change the timeframe for a scheduled maintenance window	a. Click the text "Scheduled" next to the slider button. b. Change the start and/or end date and time and click **Apply**.
Extend the length of an active maintenance window	a. Click the text "Active" next to the slider button. b. Change the end date and time and click **Apply**.
Cancel a scheduled maintenance window	Select the slider button and move it to the left.
Cancel an active maintenance window	Select the slider button and move it to the left.

View events generated during a maintenance window in Active IQ Unified Manager

If necessary, you can view the events that were generated for all storage objects during a Unified Manager maintenance window. Most events will appear in the Obsolete state once the maintenance window has completed and all system resources are back up and running.

Before you begin

At least one maintenance window must have completed before any events are available.

Events that occurred during a maintenance window do not appear on the Event Management inventory page by default.

Steps

1. In the left navigation pane, click **Events**.

   By default, all active (New and Acknowledged) events are displayed on the Event Management inventory page.

2. From the View pane, select the option **All events generated during maintenance**.

   The list of events trigged during the last 7 days from all maintenance window sessions and from all clusters are displayed.

3. If there have been multiple maintenance windows for a single cluster, you can click the **Triggered Time** calendar icon and select the period of time for the maintenance window events that you are interested in viewing.

Manage host system resource events in Active IQ Unified Manager

Unified Manager includes a service that monitors resource issues on the host system on which Unified Manager is installed. Issues such as lack of available disk space or lack of memory on the host system may trigger management station events that are displayed as banner messages across the top of the UI.

Management station events indicate an issue with the host system on which Unified Manager is installed. Examples of management station issues include disk space running low on the host system; Unified Manager missing a regular data collection cycle; and noncompletion, or late completion, of statistics analysis because the next collection poll was initiated.

Unlike all other Unified Manager event messages, these particular management station warning and critical events are displayed in banner messages.

Step

1. To view management station event information, perform these actions:

If you want to…	Do this…
View details of the event	Click the event banner to display the Event details page that includes suggested solutions for the issue.
View all management station events	a. In the left navigation pane, click **Event Management**. b. In the Filters pane on the Event Management inventory page, click the box for Management Station in the Source Type list.

Understand more about events

Event states in Active IQ Unified Manager

The state of an event helps you identify whether an appropriate corrective action is required. An event can be New, Acknowledged, Resolved, or Obsolete. Note that both New and Acknowledged events are considered to be active events.

The event states are as follows:

- **New**

  The state of a new event.

- **Acknowledged**

  The state of an event when you have acknowledged it.

- **Resolved**

  The state of an event when it is marked as resolved.

- **Obsolete**

  The state of an event when it is automatically corrected or when the cause of the event is no longer valid.

  You cannot acknowledge or resolve an obsolete event.

Example of different states of an event

The following examples illustrate the manual and automatic event state changes.

When the event Cluster Not Reachable is triggered, the event state is New. When you acknowledge the event, the event state changes to Acknowledged. When you have taken an appropriate corrective action, you must mark the event as resolved. The event state then changes to Resolved.

If the Cluster Not Reachable event is generated due to a power outage, then when the power is restored the cluster starts functioning without any administrator intervention. Therefore, the Cluster Not Reachable event is no longer valid, and the event state changes to Obsolete in the next monitoring cycle.

Unified Manager sends an alert when an event is in the Obsolete or Resolved state. The email subject line and email content of an alert provides information about the event state. An SNMP trap also includes information about the event state.

Learn about event severity types in Active IQ Unified Manager

Each event is associated with a severity type to help you prioritize the events that require immediate corrective action.

- **Critical**

  A problem occurred that might lead to service disruption if corrective action is not taken immediately.

  Performance critical events are sent from user-defined thresholds only.

- **Error**

  The event source is still performing; however, corrective action is required to avoid service disruption.

- **Warning**

  The event source experienced an occurrence that you should be aware of, or a performance counter for a cluster object is out of normal range and should be monitored to make sure it does not reach the critical severity. Events of this severity do not cause service disruption, and immediate corrective action might not be required.

  Performance warning events are sent from user-defined, system-defined, or dynamic thresholds.

- **Information**

  The event occurs when a new object is discovered, or when a user action is performed. For example, when any storage object is deleted or when there are any configuration changes, the event with severity type Information is generated.

  Information events are sent directly from ONTAP when it detects a configuration change.

Learn about event impact levels in Active IQ Unified Manager

Each event is associated with an impact level (Incident, Risk, Event, or Upgrade) to help you prioritize the events that require immediate corrective action.

- **Incident**

An incident is a set of events that can cause a cluster to stop serving data to the client and run out of space for storing data. Events with an impact level of Incident are the most severe. Immediate corrective action should be taken to avoid service disruption.

- **Risk**

  A risk is a set of events that can potentially cause a cluster to stop serving data to the client and run out of space for storing data. Events with an impact level of Risk can cause service disruption. Corrective action might be required.

- **Event**

  An event is a state or status change of storage objects and their attributes. Events with an impact level of Event are informational and do not require corrective action.

- **Upgrade**

  Upgrade events are a specific type of event reported from the Active IQ platform. These events identify issues where the resolution requires you to upgrade ONTAP software, node firmware, or operating system software (for security advisories). You may want to perform immediate corrective action for some of these issues, whereas other issues may be able to wait until your next scheduled maintenance.

Learn about event impact areas in Active IQ Unified Manager

Events are categorized into six impact areas (availability, capacity, configuration, performance, protection, and security) to enable you to concentrate on the types of events for which you are responsible.

- **Availability**

  Availability events notify you if a storage object goes offline, if a protocol service goes down, if an issue with storage failover occurs, or if an issue with hardware occurs.

- **Capacity**

  Capacity events notify you if your aggregates, volumes, LUNs, or namespaces are approaching or have reached a size threshold, or if the rate of growth is unusual for your environment.

- **Configuration**

  Configuration events inform you of the discovery, deletion, addition, removal, or renaming of your storage objects. Configuration events have an impact level of Event and a severity type of Information.

- **Performance**

  Performance events notify you of resource, configuration, or activity conditions on your cluster that might adversely affect the speed of data storage input or retrieval on your monitored storage objects.

- **Protection**

  Protection events notify you of incidents or risks involving SnapMirror relationships, issues with destination capacity, problems with SnapVault relationships, or issues with protection jobs. Any ONTAP object (especially aggregates, volumes, and SVMs) that host secondary volumes and protection relationships are categorized in the protection impact area.

- **Security**

  Security events notify you of how secure your ONTAP clusters, storage virtual machines (SVMs), and volumes are based on parameters defined in the NetApp Security Hardening Guide for ONTAP 9.

  Additionally, this area includes upgrade events that are reported from the Active IQ platform.

How Active IQ Unified Manager computes object status

Object status is determined by the most severe event that currently holds a New or Acknowledged state. For example, if an object status is Error, then one of the object's events has a severity type of Error. When corrective action has been taken, the event state moves to Resolved.

Learn about dynamic performance events in Active IQ Unified Manager

For dynamic performance events, the System Diagnosis section of the Event details page lists the top workloads with the highest latency or usage of the cluster component that is in contention.

The performance statistics are based on the time the performance event was detected up to the last time the event was analyzed. The charts also display historical performance statistics for the cluster component that is in contention.

For example, you can identify workloads with high utilization of a component to determine which workload to move to a less-utilized component. Moving the workload would reduce the amount of work on the current component, possibly bringing the component out of contention. At the top of this section is the time and date range when an event was detected and last analyzed. For active events (new or acknowledged), the last analyzed time is updated.

The latency and activity charts display the names of the top workloads when you hover your cursor over the chart. Clicking the Workload Type menu at the right of the chart enables you to sort the workloads based on their role in the event, including *sharks*, *bullies*, or *victims*, and displays details about their latency and their usage on the cluster component in contention. You can compare the actual value to the expected value to see when the workload was outside its expected range of latency or usage. For information, see Types of workloads monitored by Unified Manager.

> When you sort by peak deviation in latency, system-defined workloads are not displayed in the table, because latency applies only to user-defined workloads. Workloads with very low latency values are not displayed in the table.

For more information about the dynamic performance thresholds, see Analyzing events from dynamic performance thresholds.

For information about how Unified Manager ranks the workloads and determines the sort order, see How Unified Manager determines the performance impact for an event.

The data in the graphs shows 24 hours of performance statistics prior to the last time the event was analyzed. The actual values and expected values for each workload are based on the time the workload was involved in the event. For example, a workload might become involved in an event after the event was detected, so its performance statistics might not match the values at the time of event detection. By default, the workloads are sorted by peak (highest) deviation in latency.

Because Unified Manager retains a maximum of 30 days of 5-minute historical performance and event data, if the event is more than 30 days old, no performance data is displayed.

- **Workload Sort column**
  - **Latency chart**

    Displays the impact of the event to the latency of the workload during the last analysis.

  - **Component Usage column**

    Displays details about the workload usage of the cluster component in contention. In the graphs, the actual usage is a blue line. A red bar highlights the event duration, from the detection time to the last analyzed time. For more information, see Workload performance measurement values.

    For the network component, because network performance statistics come from activity off the cluster, this column is not displayed.

  - **Component Usage**

    Displays the history of utilization, in percent, for the network processing, data processing, and aggregate components or the history of activity, in percent, for the QoS policy group component. The chart is not displayed for the network or interconnect components. You can point to the statistics to view the usage statistics at a specific point in time.

  - **Total Write MB/s History**

    For the MetroCluster Resources component only, shows the total write throughput, in megabytes per second (MBps), for all volume workloads that are being mirrored to the partner cluster in a MetroCluster configuration.

  - **Event History**

    Displays red-shaded lines to indicate the historic events for the component in contention. For obsolete events, the chart displays events that occurred before the selected event was detected and after it was resolved.

Configuration changes detected by Active IQ Unified Manager

Unified Manager monitors your clusters for configuration changes to help you determine whether a change might have caused or contributed to a performance event. The Performance Explorer pages display a change event icon ( ) to indicate the date and time when the change was detected.

You can review the performance charts in the Performance Explorer pages and in the Workload Analysis page to see whether the change event impacted the performance of the selected cluster object. If the change was detected at or around the same time as a performance event, the change might have contributed to the issue, which caused the event alert to trigger.

Unified Manager can detect the following change events, which are categorized as Informational events:

- A volume moves between aggregates.

  Unified Manager can detect when the move is in progress, completed, or failed. If Unified Manager is down

during a volume move, when it is back up it detects the volume move and displays a change event for it.

- The throughput (MB/s or IOPS) limit of a QoS policy group that contains one or more monitored workloads changes.

  Changing a policy group limit can cause intermittent spikes in the latency (response time), which might also trigger events for the policy group. The latency gradually returns to normal and any events caused by the spikes become obsolete.

- A node in an HA pair takes over or gives back the storage of its partner node.

  Unified Manager can detect when the takeover, partial takeover, or giveback operation has been completed. If the takeover is caused by a panicked node, Unified Manager does not detect the event.

- An ONTAP upgrade or revert operation is completed successfully.

  The previous version and new version are displayed.

List of events and severity types

Learn about aggregate events in Active IQ Unified Manager

Aggregate events provide you with information about the status of aggregates so that you can monitor for potential problems. Events are grouped by impact area and include the event and trap name, impact level, source type, and severity.

Impact area: availability

An asterisk (*) identifies EMS events that have been converted to Unified Manager events.

Event name(Trap name)	Impact level	Source type	Severity
Aggregate Offline(ocumEvtAggregateStateOffline)	Incident	Aggregate	Critical
Aggregate Failed(ocumEvtAggregateStateFailed)	Incident	Aggregate	Critical
Aggregate Restricted(ocumEvtAggregateStateRestricted)	Risk	Aggregate	Warning
Aggregate Reconstructing(ocumEvtAggregateRaidStateReconstructing)	Risk	Aggregate	Warning
Aggregate Degraded(ocumEvtAggregateRaidStateDegraded)	Risk	Aggregate	Warning
Cloud Tier Partially Reachable(ocumEventCloudTierPartiallyReachable)	Risk	Aggregate	Warning

Event name(Trap name)	Impact level	Source type	Severity
Cloud Tier Unreachable(ocumEventCloudTierUnreachable)	Risk	Aggregate	Error
Cloud Tier Access Denied for Aggregate Relocation *(arlNetraCaCheckFailed)	Risk	Aggregate	Error
Cloud Tier Access Denied for Aggregate Relocation During Storage Failover *(gbNetraCaCheckFailed)	Risk	Aggregate	Error
MetroCluster Aggregate Left Behind(ocumEvtMetroClusterAggregateLeftBehind)	Risk	Aggregate	Error
MetroCluster Aggregate Mirroring Degraded(ocumEvtMetroClusterAggregateMirrorDegraded)	Risk	Aggregate	Error

Impact area: capacity

Event name(Trap name)	Impact level	Source type	Severity
Aggregate Space Nearly Full(ocumEvtAggregateNearlyFull)	Risk	Aggregate	Warning
Aggregate Space Full(ocumEvtAggregateFull)	Risk	Aggregate	Error
Aggregate Days Until Full(ocumEvtAggregateDaysUntilFullSoon)	Risk	Aggregate	Error
Aggregate Overcommitted(ocumEvtAggregateOvercommitted)	Risk	Aggregate	Error
Aggregate Nearly Overcommitted(ocumEvtAggregateAlmostOvercommitted)	Risk	Aggregate	Warning
Aggregate Snapshot Reserve Full(ocumEvtAggregateSnapReserveFull)	Risk	Aggregate	Warning
Aggregate Growth Rate Abnormal(ocumEvtAggregateGrowthRateAbnormal)	Risk	Aggregate	Warning

Impact area: configuration

Event name(Trap name)	Impact level	Source type	Severity
Aggregate Discovered(Not applicable)	Event	Aggregate	Information
Aggregate Renamed(Not applicable)	Event	Aggregate	Information
Aggregate Deleted(Not applicable)	Event	Node	Information

Impact area: performance

Event name(Trap name)	Impact level	Source type	Severity
Aggregate IOPS Critical Threshold Breached(ocumAggregateIopsIncident)	Incident	Aggregate	Critical
Aggregate IOPS Warning Threshold Breached(ocumAggregateIopsWarning)	Risk	Aggregate	Warning
Aggregate MB/s Critical Threshold Breached(ocumAggregateMbpsIncident)	Incident	Aggregate	Critical
Aggregate MB/s Warning Threshold Breached(ocumAggregateMbpsWarning)	Risk	Aggregate	Warning
Aggregate Latency Critical Threshold Breached(ocumAggregateLatencyIncident)	Incident	Aggregate	Critical
Aggregate Latency Warning Threshold Breached(ocumAggregateLatencyWarning)	Risk	Aggregate	Warning
Aggregate Performance Capacity Used Critical Threshold Breached(ocumAggregatePerfCapacityUsedIncident)	Incident	Aggregate	Critical

Event name(Trap name)	Impact level	Source type	Severity
Aggregate Performance Capacity Used Warning Threshold Breached(ocumAggregatePerfCapacityUsedWarning)	Risk	Aggregate	Warning
Aggregate Utilization Critical Threshold Breached (ocumAggregateUtilizationIncident)	Incident	Aggregate	Critical
Aggregate Utilization Warning Threshold Breached (ocumAggregateUtilizationWarning)	Risk	Aggregate	Warning
Aggregate Disks Over-utilized Threshold Breached (ocumAggregateDisksOverUtilizedWarning)	Risk	Aggregate	Warning
Aggregate Dynamic Threshold Breached (ocumAggregateDynamicEventWarning)	Risk	Aggregate	Warning

Learn about cluster events in Active IQ Unified Manager

Cluster events provide information about the status of clusters, which enables you to monitor the clusters for potential problems. The events are grouped by impact area, and include the event name, trap name, impact level, source type, and severity.

Impact area: availability

An asterisk (*) identifies EMS events that have been converted to Unified Manager events.

Event name (Trap name)	Impact level	Source type	Severity
Cluster Lacks Spare Disks(ocumEvtDisksNoSpares)	Risk	Cluster	Warning
Cluster Not Reachable(ocumEvtClusterUnreachable)	Risk	Cluster	Error
Cluster Monitoring Failed(ocumEvtClusterMonitoringFailed)	Risk	Cluster	Warning

Event name (Trap name)	Impact level	Source type	Severity
Cluster FabricPool License Capacity Limits Breached (ocumEvtExternalCapacityTierSpaceFull)	Risk	Cluster	Warning
NVMe-oF Grace Period Started *(nvmfGracePeriodStart)	Risk	Cluster	Warning
NVMe-oF Grace Period Active *(nvmfGracePeriodActive)	Risk	Cluster	Warning
NVMe-oF Grace Period Expired *(nvmfGracePeriodExpired)	Risk	Cluster	Warning
Object Maintenance Window Started(objectMaintenanceWindowStarted)	Event	Cluster	Critical
Object Maintenance Window Ended(objectMaintenanceWindowEnded)	Event	Cluster	Information
MetroCluster Spare Disks Left Behind(ocumEvtSpareDiskLeftBehind)	Risk	Cluster	Error
MetroCluster Automatic Unplanned Switchover Disabled(ocumEvtMccAutomaticUnplannedSwitchOverDisabled)	Risk	Cluster	Warning
Cluster user password changed *(cluster.passwd.changed)	Event	Cluster	Information

Impact area: capacity

Event name(Trap name)	Impact level	Source type	Severity
Cluster Capacity Imbalance Threshold Breached(ocumConformanceNodeImbalanceWarning)	Risk	Cluster	Warning

Event name(Trap name)	Impact level	Source type	Severity
Cluster Cloud Tier Planning (clusterCloudTierPlanningWarning)	Risk	Cluster	Warning
FabricPool Mirror Replication Resync Completed *(waflCaResyncComplete)	Event	Cluster	Warning
FabricPool Space Nearly Full *(fabricpoolNearlyFull)	Risk	Cluster	Error

Impact area: configuration

Event name(Trap name)	Impact level	Source type	Severity
Node Added(Not applicable)	Event	Cluster	Information
Node Removed(Not applicable)	Event	Cluster	Information
Cluster Removed(Not applicable)	Event	Cluster	Information
Cluster Add Failed(Not applicable)	Event	Cluster	Error
Cluster Name Changed(Not applicable)	Event	Cluster	Information
Emergency EMS received (Not applicable)	Event	Cluster	Critical
Critical EMS received (Not applicable)	Event	Cluster	Critical
Alert EMS received (Not applicable)	Event	Cluster	Error
Error EMS received (Not applicable)	Event	Cluster	Warning
Warning EMS received (Not applicable)	Event	Cluster	Warning
Debug EMS received (Not applicable)	Event	Cluster	Warning
Notice EMS received (Not applicable)	Event	Cluster	Warning
Informational EMS received (Not applicable)	Event	Cluster	Warning

ONTAP EMS events are categorized into three Unified Manager event severity levels.

Unified Manager event severity level	ONTAP EMS event severity level
Critical	Emergency Critical
Error	Alert
Warning	Error Warning Debug Notice Informational

Impact area: performance

Event name(Trap name)	Impact level	Source type	Severity
Cluster Load Imbalance Threshold Breached()	Risk	Cluster	Warning
Cluster IOPS Critical Threshold Breached(ocumClusterIopsIncident)	Incident	Cluster	Critical
Cluster IOPS Warning Threshold Breached(ocumClusterIopsWarning)	Risk	Cluster	Warning
Cluster MB/s Critical Threshold Breached(ocumClusterMbpsIncident)	Incident	Cluster	Critical
Cluster MB/s Warning Threshold Breached(ocumClusterMbpsWarning)	Risk	Cluster	Warning
Cluster Dynamic Threshold Breached(ocumClusterDynamicEventWarning)	Risk	Cluster	Warning

Impact area: security

Event name(Trap name)	Impact level	Source type	Severity
AutoSupport HTTPS Transport Disabled(ocumClusterASUPHttpsConfiguredDisabled)	Risk	Cluster	Warning
Log Forwarding Not Encrypted(ocumClusterAuditLogUnencrypted)	Risk	Cluster	Warning
Default Local Admin User Enabled(ocumClusterDefaultAdminEnabled)	Risk	Cluster	Warning
FIPS Mode Disabled(ocumClusterFipsDisabled)	Risk	Cluster	Warning
Login Banner Disabled(ocumClusterLoginBannerDisabled)	Risk	Cluster	Warning
Login Banner Changed(ocumClusterLoginBannerChanged)	Risk	Cluster	Warning
Log Forwarding Destinations Changed(ocumLogForwardDestinationsChanged)	Risk	Cluster	Warning
NTP Server Names Changed(ocumNtpServerNamesChanged)	Risk	Cluster	Warning
NTP Server Count is Low(securityConfigNTPServerCountLowRisk)	Risk	Cluster	Warning
Cluster Peer Communication Not Encrypted(ocumClusterPeerEncryptionDisabled)	Risk	Cluster	Warning
SSH is Using Insecure Ciphers(ocumClusterSSHInsecure)	Risk	Cluster	Warning
Telnet Protocol Enabled(ocumClusterTelnetEnabled)	Risk	Cluster	Warning
Passwords of some ONTAP user accounts use the less secure MD5 hash function(ocumClusterMD5PasswordHashUsed)	Risk	Cluster	Warning

Event name(Trap name)	Impact level	Source type	Severity
Cluster uses self-signed Certificate(ocumClusterSelfSignedCertificate)	Risk	Cluster	Warning
Cluster Remote Shell is Enabled(ocumClusterRshDisabled)	Risk	Cluster	Warning
Cluster Certificate About to Expire(ocumEvtClusterCertificateAboutToExpire)	Risk	Cluster	Warning
Cluster Certificate Expired(ocumEvtClusterCertificateExpired)	Risk	Cluster	Error

Learn about disk events in Active IQ Unified Manager

Disks events provide you with information about the status of disks so that you can monitor for potential problems. Events are grouped by impact area and include the event and trap name, impact level, source type, and severity.

Impact area: availability

Event name (Trap name)	Impact level	Source type	Severity
Flash Disks - Spare Blocks Almost Consumed(ocumEvtClusterFlashDiskFewerSpareBlockError)	Risk	Cluster	Error
Flash Disks - No Spare Blocks(ocumEvtClusterFlashDiskNoSpareBlockCritical)	Incident	Cluster	Critical
Some Unassigned Disks(ocumEvtClusterUnassignedDisksSome)	Risk	Cluster	Warning
Some Failed Disks(ocumEvtDisksSomeFailed)	Incident	Cluster	Critical

Learn about enclosures events in Active IQ Unified Manager

Enclosures events provide you with information about the status of disk shelf enclosures in your data center so that you can monitor for potential problems. Events are grouped by impact area and include the event and trap name, impact level, source type, and severity.

Impact area: availability

Event name (Trap name)	Impact level	Source type	Severity
Disk Shelf Fans Failed(ocumEvtShelfFanFailed)	Incident	Storage shelf	Critical
Disk Shelf Power Supplies Failed(ocumEvtShelfPowerSupplyFailed)	Incident	Storage shelf	Critical
Disk Shelf Multipath Not Configured(ocumDiskShelfConnectivityNotInMultiPath) This event does not apply to: • Clusters that are in a MetroCluster configuration • The following platforms: FAS2554, FAS2552, FAS2520, and FAS2240	Risk	Node	Warning
Disk Shelf Path Failure(ocumDiskShelfConnectivityPathFailure)	Risk	Storage Shelf	Warning

Impact area: configuration

Event name (Trap name)	Impact level	Source type	Severity
Disk Shelf Discovered(Not applicable)	Event	Node	Information
Disk Shelves Removed(Not applicable)	Event	Node	Information

Learn about fans events in Active IQ Unified Manager

Fans events provide you with information about the status fans on nodes in your data center so that you can monitor for potential problems. Events are grouped by impact area and include the event and trap name, impact level, source type, and severity.

Impact area: availability

Event name (Trap name)	Impact level	Source type	Severity
One or More Failed Fans(ocumEvtFansOneOrMoreFailed)	Incident	Node	Critical

Learn about flash card events in Active IQ Unified Manager

Flash card events provide you with information about the status of the flash cards installed on nodes in your data center so that you can monitor for potential problems. Events are grouped by impact area and include the event and trap name, impact level, source type, and severity.

Impact area: availability

Event name (Trap name)	Impact level	Source type	Severity
Flash Cards Offline(ocumEvtFlashCardOffline)	Incident	Node	Critical

Learn about inodes events in Active IQ Unified Manager

Inode events provide information when the inode is full or nearly full so that you can monitor for potential problems. Events are grouped by impact area and include the event and trap name, impact level, source type, and severity.

Impact area: capacity

Event name (Trap name)	Impact level	Source type	Severity
Inodes Nearly Full(ocumEvtInodesAlmostFull)	Risk	Volume	Warning
Inodes Full(ocumEvtInodesFull)	Risk	Volume	Error

Network interface events in Active IQ Unified Manager

Network interface events provide information about the status of your network interface (LIFs), so that you can monitor for potential problems. Events are grouped by impact area and include the event and trap name, impact level, source type, and severity.

Impact area: availability

Event name (Trap name)	Impact level	Source type	Severity
Network Interface Status Down(ocumEvtLifStatusDown)	Risk	Interface	Error
FC/FCoE Network Interface Status Down(ocumEvtFCLifStatusDown)	Risk	Interface	Error

Event name (Trap name)	Impact level	Source type	Severity
Network Interface Failover Not Possible(ocumEvtLifFailoverNotPossible)	Risk	Interface	Warning
Network Interface Not At Home Port(ocumEvtLifNotAtHomePort)	Risk	Interface	Warning

Impact area: configuration

Event name (Trap name)	Impact level	Source type	Severity
Network Interface Route Not Configured(Not applicable)	Event	Interface	Information

Impact area: performance

Event name (Trap name)	Impact level	Source type	Severity
Network Interface MB/s Critical Threshold Breached(ocumNetworkLifMbpsIncident)	Incident	Interface	Critical
Network Interface MB/s Warning Threshold Breached(ocumNetworkLifMbpsWarning)	Risk	Interface	Warning
FC Network Interface MB/s Critical Threshold Breached(ocumFcpLifMbpsIncident)	Incident	Interface	Critical
FC Network Interface MB/s Warning Threshold Breached(ocumFcpLifMbpsWarning)	Risk	Interface	Warning
NVMf FC Network Interface MB/s Critical Threshold Breached(ocumNvmfFcLifMbpsIncident)	Incident	Interface	Critical
NVMf FC Network Interface MB/s Warning Threshold Breached(ocumNvmfFcLifMbpsWarning)	Risk	Interface	Warning

Learn about LUN events in Active IQ Unified Manager

LUN events provide you with information about the status of your LUNs, so that you can monitor for potential problems. Events are grouped by impact area and include the event and trap name, impact level, source type, and severity.

Impact area: availability

An asterisk (*) identifies EMS events that have been converted to Unified Manager events.

Event name (Trap name)	Impact level	Source type	Severity
LUN Offline(ocumEvtLunOffline)	Incident	LUN	Critical
LUN Destroyed *(lunDestroy)	Event	LUN	Information
LUN mapped with unsupported operating system in igroup(igroupUnsupportedOsType)	Incident	LUN	Warning
Single Active Path To Access LUN(ocumEvtLunSingleActivePath)	Risk	LUN	Warning
No Active Paths To Access LUN(ocumEvtLunNotReachable)	Incident	LUN	Critical
No Optimized Paths To Access LUN(ocumEvtLunOptimizedPathInactive)	Risk	LUN	Warning
No Paths To Access LUN From HA Partner(ocumEvtLunHaPathInactive)	Risk	LUN	Warning
No Path to Access LUN from one Node in HA-pair(ocumEvtLunNodePathStatusDown)	Risk	LUN	Error

Impact area: capacity

Event name (Trap name)	Impact level	Source type	Severity
Insufficient Space For LUN Snapshot Copy(ocumEvtLunSnapshotNotPossible)	Risk	Volume	Warning

Impact area: configuration

Event name (Trap name)	Impact level	Source type	Severity
LUN mapped with unsupported operating system in igroup(igroupUnsupportedOsType)	Risk	LUN	Warning

Impact area: performance

Event name (Trap name)	Impact level	Source type	Severity
LUN IOPS Critical Threshold Breached(ocumLunIopsIncident)	Incident	LUN	Critical
LUN IOPS Warning Threshold Breached(ocumLunIopsWarning)	Risk	LUN	Warning
LUN MB/s Critical Threshold Breached(ocumLunMbpsIncident)	Incident	LUN	Critical
LUN MB/s Warning Threshold Breached(ocumLunMbpsWarning)	Risk	LUN	Warning
LUN Latency ms/op Critical Threshold Breached(ocumLunLatencyIncident)	Incident	LUN	Critical
LUN Latency ms/op Warning Threshold Breached(ocumLunLatencyWarning)	Risk	LUN	Warning
LUN Latency and IOPS Critical Threshold Breached(ocumLunLatencyIopsIncident)	Incident	LUN	Critical
LUN Latency and IOPS Warning Threshold Breached(ocumLunLatencyIopsWarning)	Risk	LUN	Warning
LUN Latency and MB/s Critical Threshold Breached(ocumLunLatencyMbpsIncident)	Incident	LUN	Critical

Event name (Trap name)	Impact level	Source type	Severity
LUN Latency and MB/s Warning Threshold Breached(ocumLunLatencyMbpsWarning)	Risk	LUN	Warning
LUN Latency and Aggregate Performance Capacity Used Critical Threshold Breached(ocumLunLatencyAggregatePerfCapacityUsedIncident)	Incident	LUN	Critical
LUN Latency and Aggregate Performance Capacity Used Warning Threshold Breached(ocumLunLatencyAggregatePerfCapacityUsedWarning)	Risk	LUN	Warning
LUN Latency and Aggregate Utilization Critical Threshold Breached(ocumLunLatencyAggregateUtilizationIncident)	Incident	LUN	Critical
LUN Latency and Aggregate Utilization Warning Threshold Breached(ocumLunLatencyAggregateUtilizationWarning)	Risk	LUN	Warning
LUN Latency and Node Performance Capacity Used Critical Threshold Breached(ocumLunLatencyNodePerfCapacityUsedIncident)	Incident	LUN	Critical
LUN Latency and Node Performance Capacity Used Warning Threshold Breached(ocumLunLatencyNodePerfCapacityUsedWarning)	Risk	LUN	Warning
LUN Latency and Node Performance Capacity Used - Takeover Critical Threshold Breached(ocumLunLatencyAggregatePerfCapacityUsedTakeoverIncident)	Incident	LUN	Critical

Event name (Trap name)	Impact level	Source type	Severity
LUN Latency and Node Performance Capacity Used - Takeover Warning Threshold Breached(ocumLunLatencyAggregatePerfCapacityUsedTakeoverWarning)	Risk	LUN	Warning
LUN Latency and Node Utilization Critical Threshold Breached(ocumLunLatencyNodeUtilizationIncident)	Incident	LUN	Critical
LUN Latency and Node Utilization Warning Threshold Breached(ocumLunLatencyNodeUtilizationWarning)	Risk	LUN	Warning
QoS LUN Max IOPS Warning Threshold Breached(ocumQosLunMaxIopsWarning)	Risk	LUN	Warning
QoS LUN Max MB/s Warning Threshold Breached(ocumQosLunMaxMbpsWarning)	Risk	LUN	Warning
Workload LUN Latency Threshold Breached as defined by Performance Service Level Policy(ocumConformanceLatencyWarning)	Risk	LUN	Warning

Learn about management station events in Active IQ Unified Manager

Management station events provide you with information about the status of server on which Unified Manager is installed so that you can monitor for potential problems. Events are grouped by impact area and include the event and trap name, impact level, source type, and severity.

Impact area: configuration

Event name (Trap name)	Impact level	Source type	Severity
Management Server Disk Space Nearly Full(ocumEvtUnifiedManagerDiskSpaceNearlyFull)	Risk	Management station	Warning

Event name (Trap name)	Impact level	Source type	Severity
Management Server Disk Space Full(ocumEvtUnifiedManagerDiskSpaceFull)	Incident	Management station	Critical
Management Server Low On Memory(ocumEvtUnifiedManagerMemoryLow)	Risk	Management station	Warning
Management Server Almost Out Of Memory(ocumEvtUnifiedManagerMemoryAlmostOut)	Incident	Management station	Critical
MySQL Log File Size Increased; Restart Required(ocumEvtMysqlLogFileSizeWarning)	Incident	Management station	Warning
Total Audit Log Size Allocation is About to Get Full	Risk	Management station	Warning
Syslog Server Certificate About to Expire	Risk	Management station	Warning
Syslog Server Certificate Expired	Risk	Management station	Error
Audit Log File Tampered	Risk	Management station	Warning
Audit Log File Deleted	Risk	Management station	Warning
Syslog Server Connection Error	Risk	Management station	Error
Syslog Server Configuration Changed	Event	Management station	Warning

Impact area: performance

Event name (Trap name)	Impact level	Source type	Severity
Performance Data Analysis Is Impacted(ocumEvtUnifiedManagerDataMissingAnalyze)	Risk	Management station	Warning
Performance Data Collection Is Impacted(ocumEvtUnifiedManagerDataMissingCollection)	Incident	Management station	Critical

These last two performance events were available for Unified Manager 7.2 only. If either of

these events exist in the New state, and then you upgrade to a newer version of Unified Manager software, the events will not be purged automatically. You will need to move the events to the Resolved state manually.

Learn about MetroCluster Bridge events in Active IQ Unified Manager

MetroCluster Bridge events provide you with information about the status of the bridges so that you can monitor for potential problems in a MetroCluster over FC configuration. Events are grouped by impact area and include the event and trap name, impact level, source type, and severity.

Impact area: availability

Event name (Trap name)	Impact level	Source type	Severity
Bridge Unreachable(ocumEvtBridgeUnreachable)	Incident	MetroCluster Bridge	Critical
Bridge Temperature Abnormal(ocumEvtBridgeTemperatureAbnormal)	Incident	MetroCluster Bridge	Critical

Learn about MetroCluster connectivity events in Active IQ Unified Manager

Connectivity events provide you with information about the connectivity between the components of a cluster and between clusters in MetroCluster over FC and MetroCluster over IP configurations, so that you can monitor for potential problems. Events are grouped by impact area and include the event and trap name, impact level, source type, and severity.

Events common in both configurations

These connectivity events are common for both MetroCluster over FC and MetroCluster over IP configurations.

Impact area: availability

Event name (Trap name)	Impact level	Source type	Severity
All Links Between MetroCluster Partners Down(ocumEvtMetroClusterAllLinksBetweenPartnersDown)	Incident	MetroCluster relationship	Critical
MetroCluster Partners Not Reachable Over Peering Network(ocumEvtMetroClusterPartnersNotReachableOverPeeringNetwork)	Incident	MetroCluster relationship	Critical

Event name (Trap name)	Impact level	Source type	Severity
MetroCluster Disaster Recovery Capability Impacted(ocumEvtMetroClusterDRStatusImpacted)	Risk	MetroCluster relationship	Critical
MetroCluster Configuration Switched Over(ocumEvtMetroClusterDRStatusImpacted)	Risk	MetroCluster relationship	Warning

MetroCluster over FC configuration

These events pertain to MetroCluster over FC configurations.

Impact area: availability

Event name (Trap name)	Impact level	Source type	Severity
All Inter-Switch Links Down(ocumEvtMetroClusterAllISLBetweenSwitchesDown)	Incident	MetroCluster inter-switch connection	Critical
FC-SAS Bridge To Storage Stack Link Down(ocumEvtBridgeSasPortDown)	Incident	MetroCluster bridge stack connection	Critical
MetroCluster Configuration Partially Switched Over(ocumEvtMetroClusterDRStatusPartiallyImpacted)	Risk	MetroCluster relationship	Error
Node To FC Switch All FC-VI Interconnect Links Down(ocumEvtMccNodeSwitchFcviLinksDown)	Incident	MetroCluster node switch connection	Critical
Node To FC Switch One Or More FC-Initiator Links Down(ocumEvtMccNodeSwitchFcLinksOneOrMoreDown)	Risk	MetroCluster node switch connection	Warning
Node To FC Switch All FC-Initiator Links Down(ocumEvtMccNodeSwitchFcLinksDown)	Incident	MetroCluster node switch connection	Critical
Switch To FC-SAS Bridge FC Link Down (ocumEvtMccSwitchBridgeFcLinksDown)	Incident	MetroCluster switch bridge connection	Critical

Event name (Trap name)	Impact level	Source type	Severity
Inter Node All FC VI InterConnect Links Down (ocumEvtMccInterNodeLinksDown)	Incident	Inter-node connection	Critical
Inter Node One Or More FC VI InterConnect Links Down (ocumEvtMccInterNodeLinksOneOrMoreDown)	Risk	Inter-node connection	Warning
Node To Bridge Link Down (ocumEvtMccNodeBridgeLinksDown)	Incident	Node bridge connection	Critical
Node to Storage Stack All SAS Links Down (ocumEvtMccNodeStackLinksDown)	Incident	Node stack connection	Critical
Node to Storage Stack One Or More SAS Links Down (ocumEvtMccNodeStackLinksOneOrMoreDown)	Risk	Node stack connection	Warning

MetroCluster over IP configuration

These events pertain to MetroCluster over IP configurations.

Impact area: availability

Event name (Trap name)	Impact level	Source type	Severity
MetroCluster IP intersite connectivity status is down (mccIntersiteconnectivityStatusDown)	Risk	MetroCluster Relationship	Critical
MetroCluster-IP Node to Switch Connection Offline (mccIpPortStatusOffline)	Risk	Node	Error

Learn about MetroCluster switch events in Active IQ Unified Manager

MetroCluster switch events for MetroCluster over FC configurations provide you with information about the status of the MetroCluster switches so that you can monitor for potential problems. Events are grouped by impact area and include the event and trap name, impact level, source type, and severity.

Impact area: availability

Event name (Trap name)	Impact level	Source type	Severity
Switch Temperature Abnormal(ocumEvtSwitchTemperatureAbnormal)	Incident	MetroCluster Switch	Critical
Switch Unreachable(ocumEvtSwitchUnreachable)	Incident	MetroCluster Switch	Critical
Switch Fans Failed(ocumEvtSwitchFansOneOrMoreFailed)	Incident	MetroCluster Switch	Critical
Switch Power Supplies Failed(ocumEvtSwitchPowerSuppliesOneOrMoreFailed)	Incident	MetroCluster Switch	Critical
Switch Temperature Sensors Failed(ocumEvtSwitchTemperatureSensorFailed) This event is applicable only for Cisco switches.	Incident	MetroCluster Switch	Critical

Learn about NVMe namespace events in Active IQ Unified Manager

NVMe Namespace events provide you with information about the status of your namespaces, so that you can monitor for potential problems. Events are grouped by impact area and include the event and trap name, impact level, source type, and severity.

An asterisk (*) identifies EMS events that have been converted to Unified Manager events.

Impact area: availability

Event name (Trap name)	Impact level	Source type	Severity
NVMeNS Offline *(nvmeNamespaceStatusOffline)	Event	Namespace	Information
NVMeNS Online *(nvmeNamespaceStatusOnline)	Event	Namespace	Information
NVMeNS Out of Space *(nvmeNamespaceSpaceOutOfSpace)	Risk	Namespace	Warning

Event name (Trap name)	Impact level	Source type	Severity
NVMeNS Destroy *(nvmeNamespaceDestroy)	Event	Namespace	Information

Impact area: performance

Event name (Trap name)	Impact level	Source type	Severity
NVMe Namespace IOPS Critical Threshold Breached(ocumNvmeNamespaceIopsIncident)	Incident	Namespace	Critical
NVMe Namespace IOPS Warning Threshold Breached(ocumNvmeNamespaceIopsWarning)	Risk	Namespace	Warning
NVMe Namespace MB/s Critical Threshold Breached(ocumNvmeNamespaceMbpsIncident)	Incident	Namespace	Critical
NVMe Namespace MB/s Warning Threshold Breached(ocumNvmeNamespaceMbpsWarning)	Risk	Namespace	Warning
NVMe Namespace Latency ms/op Critical Threshold Breached(ocumNvmeNamespaceLatencyIncident)	Incident	Namespace	Critical
NVMe Namespace Latency ms/op Warning Threshold Breached(ocumNvmeNamespaceLatencyWarning)	Risk	Namespace	Warning
NVMe Namespace Latency and IOPS Critical Threshold Breached(ocumNvmeNamespaceLatencyIopsIncident)	Incident	Namespace	Critical
NVMe Namespace Latency and IOPS Warning Threshold Breached(ocumNvmeNamespaceLatencyIopsWarning)	Risk	Namespace	Warning

Event name (Trap name)	Impact level	Source type	Severity
NVMe Namespace Latency and MB/s Critical Threshold Breached(ocumNvmeNamespaceLatencyMbpsIncident)	Incident	Namespace	Critical
NVMe Namespace Latency and MB/s Warning Threshold Breached(ocumNvmeNamespaceLatencyMbpsWarning)	Risk	Namespace	Warning

Learn about Node events in Active IQ Unified Manager

Node events provide you with information about node status so that you can monitor for potential problems. Events are grouped by impact area and include the event and trap name, impact level, source type, and severity.

An asterisk (*) identifies EMS events that have been converted to Unified Manager events.

Impact area: availability

Event name (Trap name)	Impact level	Source type	Severity
Node Root Volume Space Nearly Full(ocumEvtClusterNodeRootVolumeSpaceNearlyFull)	Risk	Node	Warning
Cloud AWS MetaDataConnFail *(ocumCloudAwsMetadataConnFail)	Risk	Node	Error
Cloud AWS IAMCredsExpired *(ocumCloudAwsIamCredsExpired)	Risk	Node	Error
Cloud AWS IAMCredsInvalid *(ocumCloudAwsIamCredsInvalid)	Risk	Node	Error
Cloud AWS IAMCredsNotFound *(ocumCloudAwsIamCredsNotFound)	Risk	Node	Error

Event name (Trap name)	Impact level	Source type	Severity
Cloud AWS IAMCredsNotInitialized *(ocumCloudAwsIamCredsNotInitialized)	Event	Node	Information
Cloud AWS IAMRoleInvalid *(ocumCloudAwsIamRoleInvalid)	Risk	Node	Error
Cloud AWS IAMRoleNotFound *(ocumCloudAwsIamRoleNotFound)	Risk	Node	Error
Cloud Tier Host Unresolvable *(ocumObjstoreHostUnresolvable)	Risk	Node	Error
Cloud Tier Intercluster Network Interface Down *(ocumObjstoreInterClusterLifDown)	Risk	Node	Error
One of NFSv4 Pools Exhausted *(nbladeNfsv4PoolEXhaust)	Incident	Node	Critical
Request Mismatch Cloud Tier Signature *(oscSignatureMismatch)	Risk	Node	Error

Impact area: capacity

Event name (Trap name)	Impact level	Source type	Severity
QoS Monitor Memory Maxed *(ocumQosMonitorMemoryMaxed)	Risk	Node	Error
QoS Monitor Memory Abated *(ocumQosMonitorMemoryAbated)	Event	Node	Information

Impact area: configuration

Event name (Trap name)	Impact level	Source type	Severity
Node Renamed(Not applicable)	Event	Node	Information

Impact area: performance

Event name (Trap name)	Impact level	Source type	Severity
Node IOPS Critical Threshold Breached(ocumNodeIopsIncident)	Incident	Node	Critical
Node IOPS Warning Threshold Breached(ocumNodeIopsWarning)	Risk	Node	Warning
Node MB/s Critical Threshold Breached(ocumNodeMbpsIncident)	Incident	Node	Critical
Node MB/s Warning Threshold Breached(ocumNodeMbpsWarning)	Risk	Node	Warning
Node Latency ms/op Critical Threshold Breached(ocumNodeLatencyIncident)	Incident	Node	Critical
Node Latency ms/op Warning Threshold Breached(ocumNodeLatencyWarning)	Risk	Node	Warning
Node Performance Capacity Used Critical Threshold Breached(ocumNodePerfCapacityUsedIncident)	Incident	Node	Critical
Node Performance Capacity Used Warning Threshold Breached(ocumNodePerfCapacityUsedWarning)	Risk	Node	Warning
Node Performance Capacity Used - Takeover Critical Threshold Breached(ocumNodePerfCapacityUsedTakeoverIncident)	Incident	Node	Critical
Node Performance Capacity Used - Takeover Warning Threshold Breached(ocumNodePerfCapacityUsedTakeoverWarning)	Risk	Node	Warning

Event name (Trap name)	Impact level	Source type	Severity
Node Utilization Critical Threshold Breached (ocumNodeUtilizationIncident)	Incident	Node	Critical
Node Utilization Warning Threshold Breached (ocumNodeUtilizationWarning)	Risk	Node	Warning
Node HA Pair Over-utilized Threshold Breached (ocumNodeHaPairOverUtilizedInformation)	Event	Node	Information
Node Disk Fragmentation Threshold Breached (ocumNodeDiskFragmentationWarning)	Risk	Node	Warning
Performance Capacity Used Threshold Breached (ocumNodeOverUtilizedWarning)	Risk	Node	Warning
Node Dynamic Threshold Breached (ocumNodeDynamicEventWarning)	Risk	Node	Warning

Impact area: security

Event name (Trap name)	Impact level	Source type	Severity
Advisory ID: NTAP-<*advisory ID*>(ocumx)	Risk	Node	Critical

Learn about NVRAM battery events in Active IQ Unified Manager

NVRAM battery events provide you with information about the status of your batteries so that you can monitor for potential problems. Events are grouped by impact area and include the event and trap name, impact level, source type, and severity.

Impact area: availability

Event name (Trap name)	Impact level	Source type	Severity
NVRAM Battery Low(ocumEvtNvramBatteryLow)	Risk	Node	Warning
NVRAM Battery Discharged(ocumEvtNvramBatteryDischarged)	Risk	Node	Error

Event name (Trap name)	Impact level	Source type	Severity
NVRAM Battery Overly Charged(ocumEvtNvramBatteryOverCharged)	Incident	Node	Critical

Learn about port events in Active IQ Unified Manager

Port events provide you with status about cluster ports so that you can monitor changes or problems on the port, like whether the port is down.

Impact area: availability

Event name (Trap name)	Impact level	Source type	Severity
Port Status Down(ocumEvtPortStatusDown)	Incident	Node	Critical

Impact area: performance

Event name (Trap name)	Impact level	Source type	Severity
Network Port MB/s Critical Threshold Breached(ocumNetworkPortMbpsIncident)	Incident	Port	Critical
Network Port MB/s Warning Threshold Breached(ocumNetworkPortMbpsWarning)	Risk	Port	Warning
FCP Port MB/s Critical Threshold Breached(ocumFcpPortMbpsIncident)	Incident	Port	Critical
FCP Port MB/s Warning Threshold Breached(ocumFcpPortMbpsWarning)	Risk	Port	Warning
Network Port Utilization Critical Threshold Breached(ocumNetworkPortUtilizationIncident)	Incident	Port	Critical
Network Port Utilization Warning Threshold Breached(ocumNetworkPortUtilizationWarning)	Risk	Port	Warning
FCP Port Utilization Critical Threshold Breached(ocumFcpPortUtilizationIncident)	Incident	Port	Critical

Event name (Trap name)	Impact level	Source type	Severity
FCP Port Utilization Warning Threshold Breached(ocumFcpPortUtilizationWarning)	Risk	Port	Warning

Learn about power supplies events in Active IQ Unified Manager

Power supplies events provide you with information about the status of your hardware so that you can monitor for potential problems. Events are grouped by impact area and include the event and trap name, impact level, source type, and severity.

Impact area: availability

Event name (Trap name)	Impact level	Source type	Severity
One or More Failed Power Supplies(ocumEvtPowerSupplyOneOrMoreFailed)	Incident	Node	Critical

Learn about protection events in Active IQ Unified Manager

Protection events tell you if a job has failed or been aborted so that you can monitor for problems. Events are grouped by impact area and include the event and trap name, impact level, source type, and severity.

Impact area: protection

Event name (Trap name)	Impact level	Source type	Severity
Protection Job Failed(ocumEvtProtectionJobTaskFailed)	Incident	Volume or storage service	Critical
Protection Job Aborted(ocumEvtProtectionJobAborted)	Risk	Volume or storage service	Warning

Learn about Qtree events in Active IQ Unified Manager

Qtree events provide you with information about the qtree capacity and the file and disk limits so that you can monitor for potential problems. Events are grouped by impact area and include the event and trap name, impact level, source type, and severity.

Impact area: capacity

Event name (Trap name)	Impact level	Source type	Severity
Qtree Space Nearly Full(ocumEvtQtreeSpaceNearlyFull)	Risk	Qtree	Warning

Event name (Trap name)	Impact level	Source type	Severity
Qtree Space Full(ocumEvtQtreeSpaceFull)	Risk	Qtree	Error
Qtree Space Normal(ocumEvtQtreeSpaceThresholdOk)	Event	Qtree	Information
Qtree Files Hard Limit Reached(ocumEvtQtreeFilesHardLimitReached)	Incident	Qtree	Critical
Qtree Files Soft Limit Breached(ocumEvtQtreeFilesSoftLimitBreached)	Risk	Qtree	Warning
Qtree Space Hard Limit Reached(ocumEvtQtreeSpaceHardLimitReached)	Incident	Qtree	Critical
Qtree Space Soft Limit Breached(ocumEvtQtreeSpaceSoftLimitBreached)	Risk	Qtree	Warning

Learn about Service Processor events in Active IQ Unified Manager

Service processor events provide you with information about the status of your processor so that you can monitor for potential problems. Events are grouped by impact area and include the event and trap name, impact level, source type, and severity.

Impact area: availability

Event name (Trap name)	Impact level	Source type	Severity
Service Processor Not Configured(ocumEvtServiceProcessorNotConfigured)	Risk	Node	Warning
Service Processor Offline(ocumEvtServiceProcessorOffline)	Risk	Node	Error

Learn about SnapMirror relationship events in Active IQ Unified Manager

SnapMirror relationship events provide you with information about the status of your Asynchronous and Synchronous SnapMirror relationships so that you can monitor for potential problems. Asynchronous SnapMirror relationship events are generated for both Storage VMs and volumes but Synchronous SnapMirror relationship events are generated only for volume relationships. There are no events generated for constituent volumes that are part of Storage VM disaster recovery relationships. Events are grouped by impact area and include the event and trap name, impact level, source type, and severity.

Impact area: protection

An asterisk (*) identifies EMS events that have been converted to Unified Manager events.

The SnapMirror relationships events are generated for Storage VMs that are protected by Storage VM disaster recovery but not for any constituent object relationships.

Event name (Trap name)	Impact level	Source type	Severity
Mirror Replication Unhealthy(ocumEvtSnapmirrorRelationshipUnhealthy)	Risk	SnapMirror relationship	Warning
Mirror Replication Broken-off(ocumEvtSnapmirrorRelationshipStateBrokenoff)	Risk	SnapMirror relationship	Error
Mirror Replication Initialize Failed(ocumEvtSnapmirrorRelationshipInitializeFailed)	Risk	SnapMirror relationship	Error
Mirror Replication Update Failed(ocumEvtSnapmirrorRelationshipUpdateFailed)	Risk	SnapMirror relationship	Error
Mirror Replication Lag Error(ocumEvtSnapMirrorRelationshipLagError)	Risk	SnapMirror relationship	Error
Mirror Replication Lag Warning(ocumEvtSnapMirrorRelationshipLagWarning)	Risk	SnapMirror relationship	Warning
Mirror Replication Resync Failed(ocumEvtSnapmirrorRelationshipResyncFailed)	Risk	SnapMirror relationship	Error
Synchronous Replication Out Of Sync *(syncSnapmirrorRelationshipOutofsync)	Risk	SnapMirror relationship	Warning
Synchronous Replication Restored *(syncSnapmirrorRelationshipInSync)	Event	SnapMirror relationship	Information
Synchronous Replication Auto Resync Failed *(syncSnapmirrorRelationshipAutoSyncRetryFailed)	Risk	SnapMirror relationship	Error

Event name (Trap name)	Impact level	Source type	Severity
Ontap Mediator is added on the Cluster (snapmirrorMediatorAdded)	Event	Cluster	Information
Ontap Mediator is removed from the Cluster (snapmirrorMediatorRemoved)	Event	Cluster	Information
Ontap Mediator is unreachable from the cluster (snapmirrorMediatorUnreachable)	Risk	Mediator	Warning
Ontap Mediator is not accessible from the cluster (snapmirrorMediatorMisconfigured)	Risk	Mediator	Error
Ontap Mediator Connectivity has been reestablished, and is resynced and ready for SnapMirror active sync (snapmirrorMediatorInQuorum)	Event	Mediator	Information

Learn about Asynchronous Mirror and Vault relationship events in Active IQ Unified Manager

Asynchronous Mirror and Vault relationship events provide you with information about the status of your Asynchronous SnapMirror and Vault relationships so that you can monitor for potential problems. Asynchronous Mirror and Vault relationship events are supported for both volume and Storage VM protection relationships. But only Vault relationships are not supported for Storage VM disaster recovery. Events are grouped by impact area and include the event and trap name, impact level, source type, and severity.

Impact area: protection

The SnapMirror and Vault relationships events are also generated for Storage VMs that are protected by Storage VM disaster recovery but not for any constituent object relationships.

Event name (Trap name)	Impact level	Source type	Severity
Asynchronous Mirror and Vault Unhealthy(ocumEvtMirrorVaultRelationshipUnhealthy)	Risk	SnapMirror relationship	Warning

Event name (Trap name)	Impact level	Source type	Severity
Asynchronous Mirror and Vault Broken-off(ocumEvtMirrorVaultRelationshipStateBrokenoff)	Risk	SnapMirror relationship	Error
Asynchronous Mirror and Vault Initialize Failed(ocumEvtMirrorVaultRelationshipInitializeFailed)	Risk	SnapMirror relationship	Error
Asynchronous Mirror and Vault Update Failed(ocumEvtMirrorVaultRelationshipUpdateFailed)	Risk	SnapMirror relationship	Error
Asynchronous Mirror and Vault Lag Error(ocumEvtMirrorVaultRelationshipLagError)	Risk	SnapMirror relationship	Error
Asynchronous Mirror and Vault Lag Warning(ocumEvtMirrorVaultRelationshipLagWarning)	Risk	SnapMirror relationship	Warning
Asynchronous Mirror and Vault Resync Failed(ocumEvtMirrorVaultRelationshipResyncFailed)	Risk	SnapMirror relationship	Error

"SnapMirror update failure" event is raised by Active IQ portal (Config Advisor).

Learn about Snapshot events in Active IQ Unified Manager

Snapshot events provide information about the status of snapshots which enables you to monitor the snapshots for potential problems. The events are grouped by impact area, and include the event name, trap name, impact level, source type, and severity.

Impact area: availability

Event name (Trap name)	Impact level	Source type	Severity
Snapshot Auto-delete Disabled(Not applicable)	Event	Volume	Information
Snapshot Auto-delete Enabled(Not applicable)	Event	Volume	Information
Snapshot Auto-delete Configuration Modified(Not applicable)	Event	Volume	Information

Learn about SnapVault relationship events in Active IQ Unified Manager

SnapVault relationship events provide you with information about the status of your SnapVault relationships so that you can monitor for potential problems. Events are grouped by impact area and include the event and trap name, impact level, source type, and severity.

Impact area: protection

Event name (Trap name)	Impact level	Source type	Severity
Asynchronous Vault Unhealthy(ocumEvtSnapVaultRelationshipUnhealthy)	Risk	SnapMirror relationship	Warning
Asynchronous Vault Broken-off(ocumEvtSnapVaultRelationshipStateBrokenoff)	Risk	SnapMirror relationship	Error
Asynchronous Vault Initialize Failed(ocumEvtSnapVaultRelationshipInitializeFailed)	Risk	SnapMirror relationship	Error
Asynchronous Vault Update Failed(ocumEvtSnapVaultRelationshipUpdateFailed)	Risk	SnapMirror relationship	Error
Asynchronous Vault Lag Error(ocumEvtSnapVaultRelationshipLagError)	Risk	SnapMirror relationship	Error
Asynchronous Vault Lag Warning(ocumEvtSnapVaultRelationshipLagWarning)	Risk	SnapMirror relationship	Warning
Asynchronous Vault Resync Failed(ocumEvtSnapvaultRelationshipResyncFailed)	Risk	SnapMirror relationship	Error

Learn about storage failover settings events in Active IQ Unified Manager

Storage failover (SFO) settings events provide you with information about whether your storage failover is disabled or not configured so that you can monitor for potential problems. Events are grouped by impact area and include the event and trap name, impact level, source type, and severity.

Impact area: availability

Event name (Trap name)	Impact level	Source type	Severity
Storage Failover Interconnect One Or More Links Down(ocumEvtSfoInterconnectOneOrMoreLinksDown)	Risk	Node	Warning
Storage Failover Disabled(ocumEvtSfoSettingsDisabled)	Risk	Node	Error
Storage Failover Not Configured(ocumEvtSfoSettingsNotConfigured)	Risk	Node	Error
Storage Failover State - Takeover(ocumEvtSfoStateTakeover)	Risk	Node	Warning
Storage Failover State - Partial Giveback(ocumEvtSfoStatePartialGiveback)	Risk	Node	Error
Storage Failover Node Status Down(ocumEvtSfoNodeStatusDown)	Risk	Node	Error
Storage Failover Takeover Not Possible(ocumEvtSfoTakeoverNotPossible)	Risk	Node	Error

Learn about storage services events in Active IQ Unified Manager

Storage services events provide you with information about the creation and subscription of storage services so that you can monitor for potential problems. Events are grouped by impact area and include the event and trap name, impact level, source type, and severity.

Impact area: configuration

Event name (Trap name)	Impact level	Source type	Severity
Storage Service Created(Not applicable)	Event	Storage service	Information
Storage Service Subscribed(Not applicable)	Event	Storage service	Information
Storage Service Unsubscribed(Not applicable)	Event	Storage service	Information

Impact area: protection

Event name (Trap name)	Impact level	Source type	Severity
Unexpected Deletion of Managed SnapMirror RelationshipocumEvtStorageServiceUnsupportedRelationshipDeletion	Risk	Storage service	Warning
Unexpected Deletion of Storage Service Member Volume(ocumEvtStorageServiceUnexpectedVolumeDeletion)	Incident	Storage service	Critical

Learn about storage shelf events in Active IQ Unified Manager

Storage shelf events tell you if your storage shelf has abnormal so that you can monitor for potential problems. Events are grouped by impact area and include the event and trap name, impact level, source type, and severity.

Impact area: availability

Event name (Trap name)	Impact level	Source type	Severity
Abnormal Voltage Range(ocumEvtShelfVoltageAbnormal)	Risk	Storage shelf	Warning
Abnormal Current Range(ocumEvtShelfCurrentAbnormal)	Risk	Storage shelf	Warning
Abnormal Temperature(ocumEvtShelfTemperatureAbnormal)	Risk	Storage shelf	Warning

Learn about Storage VM events in Active IQ Unified Manager

Storage VM (storage virtual machine, also known as SVM) events provide you with information about the status of your storage VMs (SVMs) so that you can monitor for potential problems. Events are grouped by impact area and include the event and trap name, impact level, source type, and severity.

An asterisk (*) identifies EMS events that have been converted to Unified Manager events.

Impact area: availability

Event name (Trap name)	Impact level	Source type	Severity
Storage VM CIFS Service Down(ocumEvtVserverCifsServiceStatusDown)	Incident	SVM	Critical

Event name (Trap name)	Impact level	Source type	Severity
SVM CIFS Service Not Configured(Not applicable)	Event	SVM	Information
Attempts to Connect Nonexistent CIFS Share *(nbladeCifsNoPrivShare)	Incident	SVM	Critical
CIFS NetBIOS Name Conflict *(nbladeCifsNbNameConflict)	Risk	SVM	Error
CIFS Shadow Copy Operation Failed *(cifsShadowCopyFailure)	Risk	SVM	Error
Many CIFS Connections *(nbladeCifsManyAuths)	Risk	SVM	Error
Max CIFS Connection Exceeded *(nbladeCifsMaxOpenSameFile)	Risk	SVM	Error
Max Number of CIFS Connection Per User Exceeded *(nbladeCifsMaxSessPerUsrConn)	Risk	SVM	Error
SVM FC/FCoE Service Down(ocumEvtVserverFcServiceStatusDown)	Incident	SVM	Critical
SVM iSCSI Service Down(ocumEvtVserverIscsiServiceStatusDown)	Incident	SVM	Critical
Storage VM NFS Service Down(ocumEvtVserverNfsServiceStatusDown)	Incident	SVM	Critical
SVM FC/FCoE Service Not Configured(Not applicable)	Event	SVM	Information
SVM iSCSI Service Not Configured(Not applicable)	Event	SVM	Information
SVM NFS Service Not Configured(Not applicable)	Event	SVM	Information
Storage VM Stopped(ocumEvtVserverDown)	Risk	SVM	Warning

Event name (Trap name)	Impact level	Source type	Severity
AV Server too Busy to Accept New Scan Request *(nbladeVscanConnBackPressure)	Risk	SVM	Error
No AV Server Connection for Virus Scan *(nbladeVscanNoScannerConn)	Incident	SVM	Critical
No AV Server Registered *(nbladeVscanNoRegdScanner)	Risk	SVM	Error
No Responsive AV Server Connection *(nbladeVscanConnInactive)	Event	SVM	Information
Unauthorized User Attempt to AV Server *(nbladeVscanBadUserPrivAccess)	Risk	SVM	Error
Virus Found By AV Server *(nbladeVscanVirusDetected)	Risk	SVM	Error

Impact area: configuration

Event name (Trap name)	Impact level	Source type	Severity
SVM Discovered(Not applicable)	Event	SVM	Information
SVM Deleted(Not applicable)	Event	Cluster	Information
SVM Renamed(Not applicable)	Event	SVM	Information

Impact area: performance

Event name (Trap name)	Impact level	Source type	Severity
SVM IOPS Critical Threshold Breached(ocumSvmIopsIncident)	Incident	SVM	Critical
SVM IOPS Warning Threshold Breached(ocumSvmIopsWarning)	Risk	SVM	Warning

Event name (Trap name)	Impact level	Source type	Severity
SVM MB/s Critical Threshold Breached(ocumSvmMbpsIncident)	Incident	SVM	Critical
SVM MB/s Warning Threshold Breached(ocumSvmMbpsWarning)	Risk	SVM	Warning
SVM Latency Critical Threshold Breached(ocumSvmLatencyIncident)	Incident	SVM	Critical
SVM Latency Warning Threshold Breached(ocumSvmLatencyWarning)	Risk	SVM	Warning

Impact area: security

Event name (Trap name)	Impact level	Source type	Severity
Audit Log Disabled(ocumVserverAuditLogDisabled)	Risk	SVM	Warning
Login Banner Disabled(ocumVserverLoginBannerDisabled)	Risk	SVM	Warning
SSH is Using Insecure Ciphers(ocumVserverSSHInsecure)	Risk	SVM	Warning
Login Banner Changed(ocumVserverLoginBannerChanged)	Risk	SVM	Warning
Storage VM anti-ransomware monitoring is Disabled (antiRansomwareSvmStateDisabled)	Risk	SVM	Warning
Storage VM anti-ransomware monitoring is Enabled (Learning Mode) (antiRansomwareSvmStateDryrun)	Event	SVM	Information

Event name (Trap name)	Impact level	Source type	Severity
Storage VM suitable for anti-ransomware monitoring (Learning Mode) (ocumEvtSvmArwCandidate)	Event	SVM	Information

Learn about user and group quota events in Active IQ Unified Manager

User and group quota events provide you with information about the capacity of the user and user group quota as well as the file and disk limits so that you can monitor for potential problems. Events are grouped by impact area and include the event and trap name, impact level, source type, and severity.

Impact area: capacity

Event name (Trap name)	Impact level	Source type	Severity
User or Group Quota Disk Space Soft Limit Breached(ocumEvtUserOrGroupQuotaDiskSpaceSoftLimitBreached)	Risk	User or group quota	Warning
User or Group Quota Disk Space Hard Limit Reached(ocumEvtUserOrGroupQuotaDiskSpaceHardLimitReached)	Incident	User or group quota	Critical
User or Group Quota File Count Soft Limit Breached(ocumEvtUserOrGroupQuotaFileCountSoftLimitBreached)	Risk	User or group quota	Warning
User or Group Quota File Count Hard Limit Reached(ocumEvtUserOrGroupQuotaFileCountHardLimitReached)	Incident	User or group quota	Critical

Learn about volume events in Active IQ Unified Manager

Volume events provide information about the status of volumes which enables you to monitor for potential problems. The events are grouped by impact area, and include the event name, trap name, impact level, source type, and severity.

An asterisk (*) identifies EMS events that have been converted to Unified Manager events.

Impact area: availability

Event name (Trap name)	Impact level	Source type	Severity
Volume Restricted(ocumEvtVolumeRestricted)	Risk	Volume	Warning
Volume Offline(ocumEvtVolumeOffline)	Incident	Volume	Critical
Volume Partially Available(ocumEvtVolumePartiallyAvailable)	Risk	Volume	Error
Volume Unmounted(Not applicable)	Event	Volume	Information
Volume Mounted(Not applicable)	Event	Volume	Information
Volume Remounted(Not applicable)	Event	Volume	Information
Volume Junction Path Inactive(ocumEvtVolumeJunctionPathInactive)	Risk	Volume	Warning
Volume Autosize Enabled(Not applicable)	Event	Volume	Information
Volume Autosize-Disabled(Not applicable)	Event	Volume	Information
Volume Autosize Maximum Capacity Modified(Not applicable)	Event	Volume	Information
Volume Autosize Increment Size Modified(Not applicable)	Event	Volume	Information

Impact area: capacity

Event name (Trap name)	Impact level	Source type	Severity
Thin-Provisioned Volume Space At Risk(ocumThinProvisionVolumeSpaceAtRisk)	Risk	Volume	Warning
Volume Efficiency Operation Error(ocumEvtVolumeEfficiencyOperationError)	Risk	Volume	Error
Volume Space Full(ocumEvtVolumeFull)	Risk	Volume	Error

Event name (Trap name)	Impact level	Source type	Severity
Volume Space Nearly Full(ocumEvtVolumeNearlyFull)	Risk	Volume	Warning
Volume Logical Space Full *(volumeLogicalSpaceFull)	Risk	Volume	Error
Volume Logical Space Nearly Full *(volumeLogicalSpaceNearlyFull)	Risk	Volume	Warning
Volume Logical Space Normal *(volumeLogicalSpaceAllOK)	Event	Volume	Information
Volume Snapshot Reserve Space Full(ocumEvtSnapshotFull)	Risk	Volume	Warning
Too Many Snapshot Copies(ocumEvtSnapshotTooMany)	Risk	Volume	Error
Volume Qtree Quota Overcommitted(ocumEvtVolumeQtreeQuotaOvercommitted)	Risk	Volume	Error
Volume Qtree Quota Nearly Overcommitted(ocumEvtVolumeQtreeQuotaAlmostOvercommitted)	Risk	Volume	Warning
Volume Growth Rate Abnormal(ocumEvtVolumeGrowthRateAbnormal)	Risk	Volume	Warning
Volume Days Until Full(ocumEvtVolumeDaysUntilFullSoon)	Risk	Volume	Error
Volume Space Guarantee Disabled(Not applicable)	Event	Volume	Information
Volume Space Guarantee Enabled(Not Applicable)	Event	Volume	Information
Volume Space Guarantee Modified(Not applicable)	Event	Volume	Information

Event name (Trap name)	Impact level	Source type	Severity
Volume Snapshot Reserve Days Until Full(ocumEvtVolumeSnapshotReserveDaysUntilFullSoon)	Risk	Volume	Error
FlexGroup Constituents Have Space Issues *(flexGroupConstituentsHaveSpaceIssues)	Risk	Volume	Error
FlexGroup Constituents Space Status All OK *(flexGroupConstituentsSpaceStatusAllOK)	Event	Volume	Information
FlexGroup Constituents Have Inodes Issues *(flexGroupConstituentsHaveInodesIssues)	Risk	Volume	Error
FlexGroup Constituents Inodes Status All OK *(flexGroupConstituentsInodesStatusAllOK)	Event	Volume	Information
WAFL Volume AutoSize Fail *(waflVolAutoSizeFail)	Risk	Volume	Error
WAFL Volume AutoSize Done *(waflVolAutoSizeDone)	Event	Volume	Information
FlexGroup Volume Is Over 80% Utilized*	Incident	Volume	Error
FlexGroup Volume Is Over 90% Utilized*	Incident	Volume	Critical
Volume storage efficiency anomaly (ocumVolumeAbnormalStorageEfficiencyWarning)	Risk	Volume	Warning
Volume Snapshot Reserve Under-utilized (volumeSnaphotReserveUnderutilizedWarning)	Event	Volume	Warning
Volume Snapshot Reserve Under-utilized (volumeSnaphotReserveUnderutilizedCleared)	Event	Volume	Warning

Impact area: configuration

Event name (Trap name)	Impact level	Source type	Severity
Volume Renamed(Not applicable)	Event	Volume	Information
Volume Discovered(Not applicable)	Event	Volume	Information
Volume Deleted(Not applicable)	Event	Volume	Information

Impact area: performance

Event name (Trap name)	Impact level	Source type	Severity
QoS Volume Max IOPS Warning Threshold Breached(ocumQosVolumeMaxIopsWarning)	Risk	Volume	Warning
QoS Volume Max MB/s Warning Threshold Breached(ocumQosVolumeMaxMbpsWarning)	Risk	Volume	Warning
QoS Volume Max IOPS/TB Warning Threshold Breached(ocumQosVolumeMaxIopsPerTbWarning)	Risk	Volume	Warning
Workload Volume Latency Threshold Breached as defined by Performance Service Level Policy(ocumConformanceLatencyWarning)	Risk	Volume	Warning
Volume IOPS Critical Threshold Breached(ocumVolumeIopsIncident)	Incident	Volume	Critical
Volume IOPS Warning Threshold Breached(ocumVolumeIopsWarning)	Risk	Volume	Warning
Volume MB/s Critical Threshold Breached(ocumVolumeMbpsIncident)	Incident	Volume	Critical
Volume MB/s Warning Threshold Breached(ocumVolumeMbpsWarning)	Risk	Volume	Warning

Event name (Trap name)	Impact level	Source type	Severity
Volume Latency Critical Threshold Breached(ocumVolumeLatencyIncident)	Incident	Volume	Critical
Volume Latency Warning Threshold Breached(ocumVolumeLatencyWarning)	Risk	Volume	Warning
Volume Cache Miss Ratio Critical Threshold Breached(ocumVolumeCacheMissRatioIncident)	Incident	Volume	Critical
Volume Cache Miss Ratio Warning Threshold Breached(ocumVolumeCacheMissRatioWarning)	Risk	Volume	Warning
Volume Latency and IOPS Critical Threshold Breached(ocumVolumeLatencyIopsIncident)	Incident	Volume	Critical
Volume Latency and IOPS Warning Threshold Breached(ocumVolumeLatencyIopsWarning)	Risk	Volume	Warning
Volume Latency and MB/s Critical Threshold Breached(ocumVolumeLatencyMbpsIncident)	Incident	Volume	Critical
Volume Latency and MB/s Warning Threshold Breached(ocumVolumeLatencyMbpsWarning)	Risk	Volume	Warning
Volume Latency and Aggregate Performance Capacity Used Critical Threshold Breached(ocumVolumeLatencyAggregatePerfCapacityUsedIncident)	Incident	Volume	Critical
Volume Latency and Aggregate Performance Capacity Used Warning Threshold Breached(ocumVolumeLatencyAggregatePerfCapacityUsedWarning)	Risk	Volume	Warning

Event name (Trap name)	Impact level	Source type	Severity
Volume Latency and Aggregate Utilization Critical Threshold Breached(ocumVolumeLatencyAggregateUtilizationIncident)	Incident	Volume	Critical
Volume Latency and Aggregate Utilization Warning Threshold Breached(ocumVolumeLatencyAggregateUtilizationWarning)	Risk	Volume	Warning
Volume Latency and Node Performance Capacity Used Critical Threshold Breached(ocumVolumeLatencyNodePerfCapacityUsedIncident)	Incident	Volume	Critical
Volume Latency and Node Performance Capacity Used Warning Threshold Breached(ocumVolumeLatencyNodePerfCapacityUsedWarning)	Risk	Volume	Warning
Volume Latency and Node Performance Capacity Used - Takeover Critical Threshold Breached(ocumVolumeLatencyAggregatePerfCapacityUsedTakeoverIncident)	Incident	Volume	Critical
Volume Latency and Node Performance Capacity Used - Takeover Warning Threshold Breached(ocumVolumeLatencyAggregatePerfCapacityUsedTakeoverWarning)	Risk	Volume	Warning
Volume Latency and Node Utilization Critical Threshold Breached(ocumVolumeLatencyNodeUtilizationIncident)	Incident	Volume	Critical

Event name (Trap name)	Impact level	Source type	Severity
Volume Latency and Node Utilization Warning Threshold Breached(ocumVolumeLatencyNodeUtilizationWarning)	Risk	Volume	Warning

Impact area: security

Event name (Trap name)	Impact level	Source type	Severity
Volume anti-ransomware monitoring is Enabled (Active Mode) (antiRansomwareVolumeStateEnabled)	Event	Volume	Information
Volume anti-ransomware monitoring is Disabled (antiRansomwareVolumeStateDisabled)	Risk	Volume	Warning
Volume anti-ransomware monitoring is Enabled (Learning Mode) (antiRansomwareVolumeStateDryrun)	Event	Volume	Information
Volume anti-ransomware monitoring is Paused (Learning Mode) (antiRansomwareVolumeStateDryrunPaused)	Risk	Volume	Warning
Volume anti-ransomware monitoring is Paused (Active Mode) (antiRansomwareVolumeStateEnablePaused)	Risk	Volume	Warning
Volume anti-ransomware monitoring is Disabling (antiRansomwareVolumeStateDisableInProgress)	Risk	Volume	Warning
Ransomware Activity Seen (callHomeRansomwareActivitySeen)	Incident	Volume	Critical
Volume suitable for anti-ransomware monitoring (Learning Mode) (ocumEvtVolumeArwCandidate)	Event	Volume	Information

Event name (Trap name)	Impact level	Source type	Severity
Volume suitable for anti-ransomware monitoring (Active Mode) (ocumVolumeSuitedForActiveAntiRansomwareDetection)	Risk	Volume	Warning
Volume exhibits noisy anti-ransomware alerting (antiRansomwareFeatureNoisyVolume)	Risk	Volume	Warning

Impact area: data protection

Event name (Trap name)	Impact level	Source type	Severity
Volume Has Insufficient Local Snapshot Protection (volumeLacksLocalProtectionWarning)	Risk	Volume	Warning
Volume Has Insufficient Local Snapshot Protection (volumeLacksLocalProtectionCleared)	Risk	Volume	Warning

Learn about volume move status events in Active IQ Unified Manager

Volume move status events tell you about the status of your volume move so that you can monitor for potential problems. Events are grouped by impact area and include the event and trap name, impact level, source type, and severity.

Impact area: capacity

Event name (Trap name)	Impact level	Source type	Severity
Volume Move Status: In Progress(Not applicable)	Event	Volume	Information
Volume Move Status - Failed(ocumEvtVolumeMoveFailed)	Risk	Volume	Error
Volume Move Status: Completed(Not applicable)	Event	Volume	Information
Volume Move - Cutover Deferred(ocumEvtVolumeMoveCutoverDeferred)	Risk	Volume	Warning

Description of event windows and dialog boxes

Learn about the Notifications page in Active IQ Unified Manager

You can configure the Unified Manager server to send notifications when an event is generated or when it is assigned to a user. You can also configure the notification mechanisms. For example, notifications can be sent as emails or SNMP traps.

You must have the Application Administrator or Storage Administrator role.

Email

This area enables you to configure the following email settings for alert notification:

- **From Address**

  Specifies the email address from which the alert notification is sent. This value is also used as the from address for a report when shared. If the From Address is pre-filled with the address "ActiveIQUnifiedManager@localhost.com", you should change it to a real, working email address to make sure that all email notifications are delivered successfully.

SMTP Server

This area enables you to configure the following SMTP server settings:

- **Host Name or IP Address**

  Specifies the host name of your SMTP host server, which is used to send the alert notification to the specified recipients.

- **User Name**

  Specifies the SMTP user name. SMTP user name is required only when the SMTPAUTH is enabled in the SMTP server.

- **Password**

  Specifies the SMTP password. SMTP user name is required only when the SMTPAUTH is enabled in the SMTP server.

- **Port**

  Specifies the port that is used by the SMTP host server to send alert notification.

  The default value is 25.

- **Use START/TLS**

  Checking this box provides secure communication between the SMTP server and the management server by using the TLS/SSL protocols (also known as start_tls and StartTLS).

- **Use SSL**

  Checking this box provides secure communication between the SMTP server and the management server by using the SSL protocol.

SNMP

This area enables you to configure the following SNMP trap settings:

- **Version**

  Specifies the SNMP version you want to use depending on the type of security you require. Options include Version 1, Version 3, Version 3 with Authentication, and Version 3 with Authentication and Encryption. The default value is Version 1.

- **Trap Destination Host**

  Specifies the host name or IP address (IPv4 or IPv6) that receives the SNMP traps that are sent by the management server. To specify multiple trap destinations, separate each host with a comma.

  All other SNMP settings, such as "Version" and "Outbound Port", must be the same for all hosts in the list.

- **Outbound Trap Port**

  Specifies the port through which the SNMP server receives the traps that are sent by the management server.

  The default value is 162.

- **Community**

  The community string to access the host.

- **Engine ID**

  Specifies the unique identifier of the SNMP agent and is automatically generated by the management server. Engine ID is available with SNMP Version 3, SNMP Version 3 with Authentication, and SNMP Version 3 with Authentication and Encryption.

- **Username**

  Specifies the SNMP user name. User name is available with SNMP Version 3, SNMP Version 3 with Authentication, and SNMP Version 3 with Authentication and Encryption.

- **Authentication Protocol**

  Specifies the protocol used to authenticate a user. Protocol options include MD5 and SHA. MD5 is the default value. Authentication protocol is available with SNMP Version 3 with Authentication and SNMP Version 3 with Authentication and Encryption.

- **Authentication Password**

  Specifies the password used when authenticating a user. Authentication password is available with SNMP Version 3 with Authentication and SNMP Version 3 with Authentication and Encryption.

- **Privacy Protocol**

  Specifies the privacy protocol used to encrypt SNMP messages. Protocol options include AES 128 and DES. The default value is AES 128. Privacy protocol is available with SNMP Version 3 with Authentication and Encryption.

- **Privacy Password**

  Specifies the password when using privacy protocol. Privacy password is available with SNMP Version 3 with Authentication and Encryption.

For more information about SNMP objects and traps, you can download the Active IQ Unified Manager MIB from the NetApp Support Site.

Learn about the Event Management inventory page in Active IQ Unified Manager

The Event Management inventory page enables you to view a list of current events and their properties. You can perform tasks such as acknowledging, resolving, and assigning events. You can also add an alert for specific events.

The information on this page is refreshed automatically every 5 minutes to ensure that the most current new events are displayed.

Filter components

Enable you to customize the information that is displayed in the events list. You can refine the list of events that are displayed using the following components:

- View menu to select from a pre-defined list of filter selections.

  This includes items such as all active (new and acknowledged) events, active performance events, events assigned to me (the logged in user), and all events generated during all maintenance windows.

- Search pane to refine the list of events by entering full or partial terms.
- Filter button that launches the Filters pane so you can select from every available field and field attribute to refine the list of events.

Command buttons

The command buttons enable you to perform the following tasks:

- **Assign To**

  Enables you to select the user to whom the event is assigned. When you assign an event to a user, the user name and the time when you assigned the event is added in the events list for the selected events.

  - Me

    Assigns the event to the currently logged in user.

  - Another user

    Displays the Assign Owner dialog box, which enables you to assign or reassign the event to other users. You can also unassign events by leaving the ownership field blank.

- **Acknowledge**

  Acknowledges the selected events.

  When you acknowledge an event, your user name and the time when you acknowledged the event are

added in the events list for the selected events. When you acknowledge an event, you are responsible for managing that event.

You cannot acknowledge Information events.

- **Mark As Resolved**

  Enables you to change the event state to resolved.

  When you resolve an event, your user name and the time when you resolved the event are added in the events list for the selected events. After you have taken corrective action for the event, you must mark the event as resolved.

- **Add Alert**

  Displays the Add Alert dialog box, which enables you to add alerts for the selected events.

- **Reports**

  Enables you to export details of the current event view to a comma-separated values (.csv) file or PDF document.

- **Show/Hide Column Selector**

  Enables you to choose the columns that display on the page and select the order in which they are displayed.

Events list

Displays details of all the events ordered by triggered time.

By default the All active events view is displayed to show the New and Acknowledged events for the previous seven days that have an Impact Level of Incident or Risk.

- **Triggered Time**

  The time at which the event was generated.

- **Severity**

  The event severity: Critical (), Error (), Warning (), and Information ().

- **State**

  The event state: New, Acknowledged, Resolved, or Obsolete.

- **Impact Level**

  The event impact level: Incident, Risk, Event, or Upgrade.

- **Impact Area**

  The event impact area: Availability, Capacity, Performance, Protection, Configuration, or Security.

- **Name**

The event name. You can select the name to display the Event details page for that event.

- **Source**

  The name of the object on which the event has occurred. You can select the name to display the health or performance details page for that object.

  When a shared QoS policy breach occurs, only the workload object that is consuming the most IOPS or MB/s is shown in this field. Additional workloads that are using this policy are displayed in the Event details page.

- **Source Type**

  The object type (for example, Storage VM, Volume, or Qtree) with which the event is associated.

- **Assigned To**

  The name of the user to whom the event is assigned.

- **Event Origin**

  Whether the event originated from the "Active IQ Portal" or directly from "Active IQ Unified Manager".

- **Annotation Name**

  The name of the annotation that is assigned to the storage object.

- **Notes**

  The number of notes that are added for an event.

- **Days Outstanding**

  The number of days since the event was initially generated.

- **Assigned Time**

  The time that has elapsed since the event was assigned to a user. If the time elapsed exceeds a week, the timestamp when the event was assigned to a user is displayed.

- **Acknowledged By**

  The name of the user who acknowledged the event. The field is blank if the event is not acknowledged.

- **Acknowledged Time**

  The time that has elapsed since the event was acknowledged. If the time elapsed exceeds a week, the timestamp when the event was acknowledged is displayed.

- **Resolved By**

  The name of the user who resolved the event. The field is blank if the event is not resolved.

- **Resolved Time**

  The time that has elapsed since the event was resolved. If the time elapsed exceeds a week, the

timestamp when the event was resolved is displayed.

- **Obsoleted Time**

  The time when the state of the event became Obsolete.

Event details page

Learn about the Event details page in Active IQ Unified Manager

From the Event details page, you can view the details of a selected event, such as the event severity, impact level, impact area, and event source. You can also view additional information about possible remediations to resolve the issue.

- **Event Name**

  The name of the event and the time the event was last seen.

  For non-performance events, while the event is in the New or Acknowledged state the last seen information is not known and is therefore hidden.

- **Event Description**

  A brief description of the event.

  In some cases a reason for the event being triggered is provided in the event description.

- **Component in Contention**

  For dynamic performance events, this section displays icons that represent the logical and physical components of the cluster. If a component is in contention, its icon is circled and highlighted red.

  See *Cluster components and why they can be in contention* for a description of the components that are displayed here.

The Event Information, System Diagnosis, and Suggested Actions sections are described in other topics.

Command buttons

The command buttons enable you to perform the following tasks:

- **Notes icon**

  Enables you to add or update a note about the event, and review all notes left by other users.

Actions menu

- **Assign to Me**

  Assigns the event to you.

- **Assign to Others**

  Opens the Assign Owner dialog box, which enables you to assign or reassign the event to other users.

When you assign an event to a user, the user's name and the time when the event was assigned are added in the events list for the selected events.

You can also unassign events by leaving the ownership field blank.

- **Acknowledge**

  Acknowledges the selected events so that you do not continue to receive repeat alert notifications.

  When you acknowledge an event, your user name and the time that you acknowledged the event are added in the events list (Acknowledged By) for the selected events. When you acknowledge an event, you take responsibility for managing that event.

- **Mark As Resolved**

  Enables you to change the event state to Resolved.

  When you resolve an event, your user name and the time that you resolved the event are added in the events list (Resolved By) for the selected events. After you have taken corrective action for the event, you must mark the event as resolved.

- **Add Alert**

  Displays the Add Alert dialog box, which enables you to add an alert for the selected event.

Learn about the Event Information section in Active IQ Unified Manager

You use the Event Information section on the Event details page to view the details about a selected event, such as the event severity, impact level, impact area, and event source.

Fields that are not applicable to the event type are hidden. You can view the following event details:

- **Event Trigger Time**

  The time at which the event was generated.

- **State**

  The event state: New, Acknowledged, Resolved, or Obsolete.

- **Obsoleted Cause**

  The actions that caused the event to be obsoleted, for example, the issue was fixed.

- **Event Duration**

  For active (new and acknowledged) events, this is the time between detection and the time when the event was last analyzed. For obsolete events, this is the time between detection and when the event was resolved.

  This field is displayed for all performance events, and for other event types only after they have been resolved or obsoleted.

- **Last Seen**

The date and time at which the event was last seen as active.

For performance events this value may be more recent than the Event Trigger Time as this field is updated after each new collection of performance data as long as the event is active. For other types of events, when in the New or Acknowledged state, this content is not updated and the field is therefore hidden.

- **Severity**

  The event severity: Critical (), Error (), Warning (), and Information ().

- **Impact Level**

  The event impact level: Incident, Risk, Event, or Upgrade.

- **Impact Area**

  The event impact area: Availability, Capacity, Performance, Protection, Configuration, or Security.

- **Source**

  The name of the object on which the event has occurred.

  When viewing the details for a shared QoS policy event, up to three of the workload objects that are consuming the most IOPS or MBps are listed in this field.

  You can click the source name link to display the health or performance details page for that object.

- **Source Annotations**

  Displays the annotation name and value for the object to which the event is associated.

  This field is displayed only for health events on clusters, SVMs, and volumes.

- **Source Groups**

  Displays the names of all the groups of which the impacted object is a member.

  This field is displayed only for health events on clusters, SVMs, and volumes.

- **Source Type**

  The object type (for example, SVM, Volume, or Qtree) with which the event is associated.

- **On Cluster**

  The name of the cluster on which the event occurred.

  You can click the cluster name link to display the health or performance details page for that cluster.

- **Affected Objects Count**

  The number of objects affected by the event.

  You can click the object link to display the inventory page populated with the objects that are currently affected by this event.

This field is displayed only for performance events.

- **Affected Volumes**

  The number of volumes that are being affected by this event.

  This field is displayed only for performance events on nodes or aggregates.

- **Triggered Policy**

  The name of the threshold policy that issued the event.

  You can hover your cursor over the policy name to see the details of the threshold policy. For adaptive QoS policies the defined policy, block size, and allocation type (allocated space or used space) is also displayed.

  This field is displayed only for performance events.

- **Rule Id**

  For Active IQ platform events, this is the number of the rule that was triggered to generate the event.

- **Acknowledged by**

  The name of the person who acknowledged the event and the time that the event was acknowledged.

- **Resolved by**

  The name of the person who resolved the event and the time that the event was resolved.

- **Assigned to**

  The name of the person who is assigned to work on the event.

- **Alert Settings**

  The following information about alerts is displayed:

  - If there are no alerts associated with the selected event, an **Add alert** link is displayed.

    You can open the Add Alert dialog box by clicking the link.

  - If there is one alert associated with the selected event, the alert name is displayed.

    You can open the Edit Alert dialog box by clicking the link.

  - If there is more than one alert associated with the selected event, the number of alerts is displayed.

    You can open the Alert Setup page by clicking the link to view more details about these alerts.

  Alerts that are disabled are not displayed.

- **Last Notification Sent**

  The date and time at which the most recent alert notification was sent.

- **Send by**

  The mechanism that was used to send the alert notification: email or SNMP trap.

- **Previous Script Run**

  The name of the script that was executed when the alert was generated.

Learn about Suggested Actions in Active IQ Unified Manager

The Suggested Actions section of the Event details page provides possible reasons for the event and suggests a few actions so that you can try to resolve the event on your own. The suggested actions are customized based on the type of event or type of threshold that has been breached.

This area is displayed only for some types of events.

In some cases there are **Help** links provided on the page that reference additional information for many suggested actions, including instructions for performing a specific action. Some of the actions may involve using Unified Manager, ONTAP System Manager, OnCommand Workflow Automation, ONTAP CLI commands, or a combination of these tools.

You should consider the actions suggested here as only a guidance in resolving this event. The action you take to resolve this event should be based on the context of your environment.

If you want to analyze the object and event in more detail, click the **Analyze Workload** button to display the Workload Analysis page.

There are certain events that Unified Manager can diagnose thoroughly and provide a single resolution. When available, those resolutions are displayed with a **Fix It** button. Click this button to have Unified Manager fix the issue causing the event.

For Active IQ platform events, this section may contain a link to a NetApp Knowledgebase article, when available, that describes the issue and possible resolutions. In sites with no external network access, a PDF of the Knowledgebase article is opened locally; the PDF is part of the rules file that you manually download to the Unified Manager instance.

Learn about System Diagnosis in Active IQ Unified Manager

The System Diagnosis section of the Event details page provides information that can help you diagnose issues that may have been responsible for the event.

This area is displayed only for some events.

Some performance events provide charts that are relevant to the particular event that has been triggered. Typically this includes and IOPS or MBps chart and a latency chart for the previous ten days. When arranged this way you can see which storage components are most affecting latency, or being affected by latency, when the event is active.

For dynamic performance events, the following charts are displayed:

- Workload Latency - Displays the history of latency for the top victim, bully, or shark workloads at the component in contention.

- Workload Activity - Displays details about the workload usage of the cluster component in contention.
- Resource Activity - Display historical performance statistics for the cluster component in contention.

Other charts are displayed when some cluster components are in contention.

Other events provide a brief description of the type of analysis the system is performing on the storage object. In some cases there will be one or more lines; one for each component that has been analyzed, for system-defined performance policies that analyze multiple performance counters. In this scenario, a green or red icon displays next to the diagnosis to indicate whether an issue was found, or not, in that particular diagnosis.

Learn about the Event Setup page in Active IQ Unified Manager

The Event Setup page displays the list of events that are disabled, and provides information such as the associated object type and severity of the event. You can also perform tasks such as disabling or enabling events globally.

You can access this page only if you have the Application Administrator or Storage Administrator role.

Command buttons

The command buttons enable you to perform the following tasks for selected events:

- **Disable**

  Launches the Disable Events dialog box, which you can use to disable events.

- **Enable**

  Enables selected events that you had chosen to disable previously.

- **Upload Rules**

  Launches the Upload Rules dialog box, which enables sites with no external network access to manually upload the Active IQ rules file to Unified Manager. The rules are run against cluster AutoSupport messages to generate events for system configuration, cabling, best practice, and availability as defined by the Active IQ platform.

- **Subscribe to EMS Events**

  Launches the Subscribe to EMS Events dialog box, which enables you to subscribe to receive specific Event Management System (EMS) events from the clusters that you are monitoring. The EMS collects information about events that occur on the cluster. When a notification is received for a subscribed EMS event, a Unified Manager event is generated with the appropriate severity.

List view

The List view displays (in tabular format) information about events that are disabled. You can use the column filters to customize the data that is displayed.

- **Event**

  Displays the name of the event that is disabled.

- **Severity**

Displays the severity of the event. The severity can be Critical, Error, Warning, or Information.

- **Source Type**

  Displays the source type for which the event is generated.

Disable the Events dialog box in Active IQ Unified Manager

The Disable Events dialog box displays the list of event types for which you can disable events. You can disable events for an event type based on a particular severity or for a set of events.

You must have the Application Administrator or Storage Administrator role.

Event Properties area

The Event Properties area specifies the following event properties:

- **Event Severity**

  Enables you to select events based on the severity type, which can be Critical, Error, Warning, or Information.

- **Event Name Contains**

  Enables you to filter events whose name contains the specified characters.

- **Matching events**

  Displays the list of events matching the event severity type and the text string you specify.

- **Disable events**

  Displays the list of events that you have selected for disabling.

  The severity of the event is also displayed along with the event name.

Command buttons

The command buttons enable you to perform the following tasks for the selected events:

- **Save and close**

  Disables the event type and closes the dialog box.

- **Cancel**

  Discards the changes and closes the dialog box.

Manage alerts

Learn about alerts in Active IQ Unified Manager

While events occur continuously, the Unified Manager generates an alert only when an

event meets specified filter criteria. You can choose the events for which alerts should be generated—for example, when a space threshold is exceeded or an object goes offline. You can also associate an alert with a script that is executed when an alert is triggered.

Filter criteria include object class, name, or event severity.

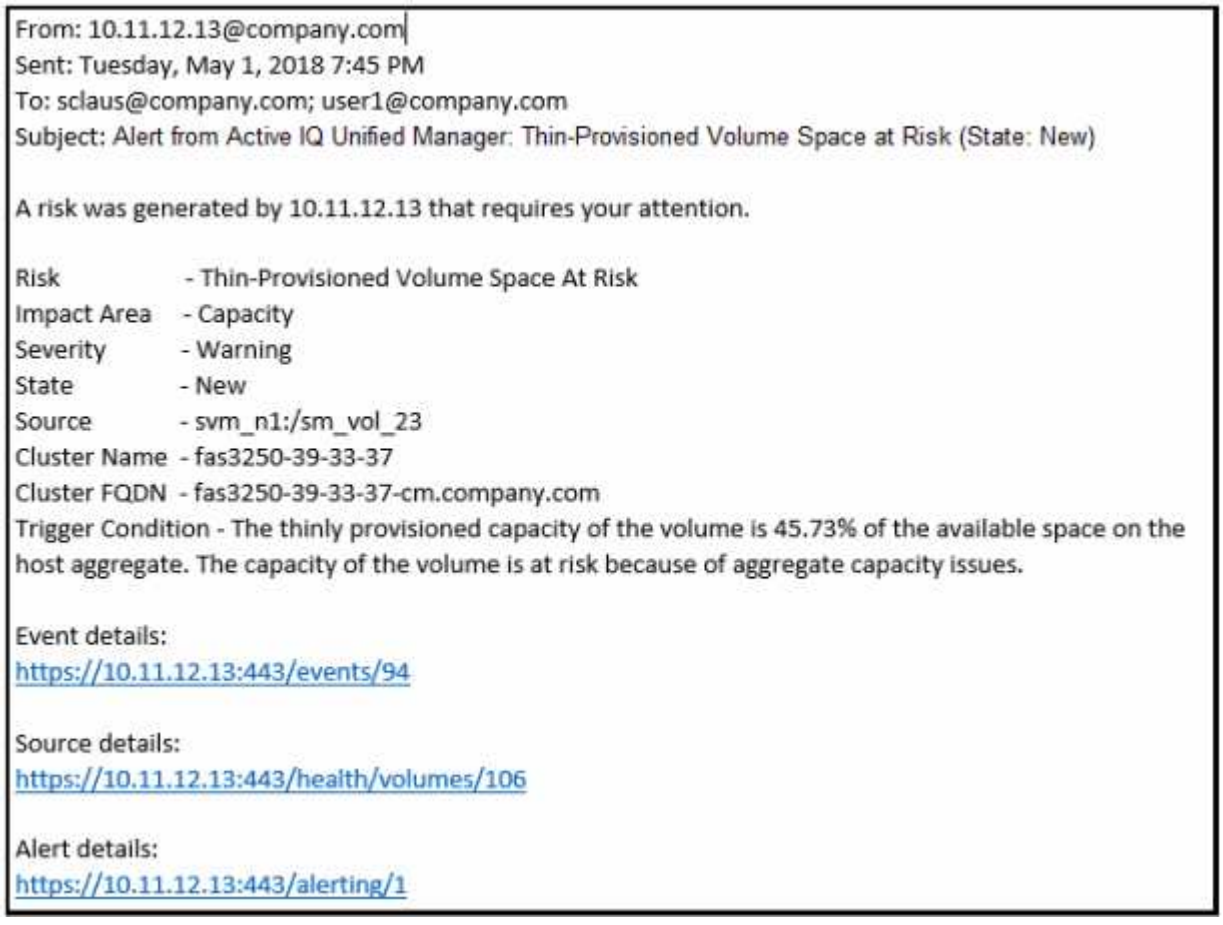
Information contained in an Active IQ Unified Manager alert email

Unified Manager alert emails provide the type of event, the severity of the event, the name of the policy or threshold that was breached to cause the event, and a description of the event. The email message also provides a hyperlink for each event that enables you to view the details page for the event in the UI.

Alert emails are sent to all users who have subscribed to receive alerts.

If a performance counter or capacity value has a large change during a collection period, it could cause both a critical and a warning event to be triggered at the same time for the same threshold policy. In this case, you may receive one email for the warning event and one for the critical event. This is because Unified Manager enables you to subscribe separately to receive alerts for warning and critical threshold breaches.

A sample alert email is shown below:


From: 10.11.12.13@company.com
Sent: Tuesday, May 1, 2018 7:45 PM
To: sclaus@company.com; user1@company.com
Subject: Alert from Active IQ Unified Manager: Thin-Provisioned Volume Space at Risk (State: New)

A risk was generated by 10.11.12.13 that requires your attention.

Risk - Thin-Provisioned Volume Space At Risk
Impact Area - Capacity
Severity - Warning
State - New
Source - svm_n1:/sm_vol_23
Cluster Name - fas3250-39-33-37
Cluster FQDN - fas3250-39-33-37-cm.company.com
Trigger Condition - The thinly provisioned capacity of the volume is 45.73% of the available space on the host aggregate. The capacity of the volume is at risk because of aggregate capacity issues.

Event details:
https://10.11.12.13:443/events/94

Source details:
https://10.11.12.13:443/health/volumes/106

Alert details:
https://10.11.12.13:443/alerting/1

Add alerts

Add alerts in Active IQ Unified Manager

You can configure alerts to notify you when a particular event is generated. You can configure alerts for a single resource, for a group of resources, or for events of a particular severity type. You can specify the frequency with which you want to be notified and associate a script to the alert.

Before you begin

- You must have configured notification settings such as the user email address, SMTP server, and SNMP trap host to enable the Active IQ Unified Manager server to use these settings to send notifications to users when an event is generated.
- You must know the resources and events for which you want to trigger the alert, and the user names or email addresses of the users that you want to notify.
- If you want to have a script execute based on the event, you must have added the script to Unified Manager by using the Scripts page.
- You must have the Application Administrator or Storage Administrator role.

You can create an alert directly from the Event details page after receiving an event in addition to creating an alert from the Alert Setup page, as described here.

Steps

1. In the left navigation pane, click **Storage Management** > **Alert Setup**.
2. In the **Alert Setup** page, click **Add**.
3. In the **Add Alert** dialog box, click **Name**, and enter a name and description for the alert.
4. Click **Resources**, and select the resources to be included in or excluded from the alert.

   You can set a filter by specifying a text string in the **Name contains** field to select a group of resources. Based on the text string that you specify, the list of available resources displays only those resources that match the filter rule. The text string that you specify is case-sensitive.

   If a resource conforms to both the include and exclude rules that you have specified, the exclude rule takes precedence over the include rule, and the alert is not generated for events related to the excluded resource.

5. Click **Events**, and select the events based on the event name or event severity type for which you want to trigger an alert.

   To select more than one event, press the Ctrl key while you make your selections.

6. Click **Actions**, and select the users that you want to notify, choose the notification frequency, choose whether an SNMP trap will be sent to the trap receiver, and assign a script to be executed when an alert is generated.

   If you modify the email address that is specified for the user and reopen the alert for editing, the Name field appears blank because the modified email address is no longer mapped to the user that was previously selected. Also, if you modified the email address of the selected user from the Users page, the modified email address is not updated for the selected user.

   You can also choose to notify users through SNMP traps.

7. Click **Save**.

Example of adding an alert

This example shows how to create an alert that meets the following requirements:

- Alert name: HealthTest
- Resources: includes all volumes whose name contains "abc" and excludes all volumes whose name contains "xyz"
- Events: includes all critical health events
- Actions: includes "sample@domain.com", a "Test" script, and the user has to be notified every 15 minutes

Perform the following steps in the Add Alert dialog box:

1. Click **Name**, and enter **`HealthTest`** in the **Alert Name** field.
2. Click **Resources**, and in the Include tab, select **Volumes** from the drop-down list.
   a. Enter **`abc`** in the **Name contains** field to display the volumes whose name contains "abc".
   b. Select **<<All Volumes whose name contains 'abc'>>** from the Available Resources area, and move it to the Selected Resources area.
   c. Click **Exclude**, and enter **`xyz`** in the **Name contains** field, and then click **Add**.
3. Click **Events**, and select **Critical** from the Event Severity field.
4. Select **All Critical Events** from the Matching Events area, and move it to the Selected Events area.
5. Click **Actions**, and enter **`sample@domain.com`** in the Alert these users field.
6. Select **Remind every 15 minutes** to notify the user every 15 minutes.

   You can configure an alert to repeatedly send notifications to the recipients for a specified time. You should determine the time from which the event notification is active for the alert.

7. In the Select Script to Execute menu, select **Test** script.
8. Click **Save**.

Guidelines for adding alerts in Active IQ Unified Manager

You can add alerts based on a resource, such as a cluster, node, aggregate, or volume, and events of a particular severity type. As a best practice, you can add an alert for any of your critical objects after you have added the cluster to which the object belongs.

You can use the following guidelines and considerations to create alerts to manage your systems effectively:

- Alert description

  You should provide a description for the alert so that it helps you track your alerts effectively.

- Resources

  You should decide which physical or logical resource requires an alert. You can include and exclude resources, as required. For example, if you want to closely monitor your aggregates by configuring an alert, you must select the required aggregates from the list of resources.

If you select a category of resources, for example, **<<All User or Group Quotas>>**, then you will receive alerts for all objects in that category.

> Selecting a cluster as the resource does not automatically select the storage objects within that cluster. For example, if you create an alert for all critical events for all clusters you will receive alerts only for cluster critical events. You will not receive alerts for critical events on nodes, aggregates, and so forth.

- Event severity

  You should decide if an event of a specified severity type (Critical, Error, Warning) should trigger the alert and, if so, which severity type.

- Selected Events

  If you add an alert based on the type of event generated, you should decide which events require an alert.

  If you select an event severity, but do not select any individual events (if you leave the "Selected Events" column empty) then you will receive alerts for all events in the category.

- Actions

  You must provide the user names and email addresses of the users who receive the notification. You can also specify an SNMP trap as a mode of notification. You can associate your scripts to an alert so that they are executed when an alert is generated.

- Notification frequency

  You can configure an alert to repeatedly send notification to the recipients for a specified time. You should determine the time from which the event notification is active for the alert. If you want the event notification to be repeated until the event is acknowledged, you should determine how often you want the notification to be repeated.

- Execute Script

  You can associate your script with an alert. Your script is executed when the alert is generated.

Add alerts for performance events in Active IQ Unified Manager

You can configure alerts for individual performance events just like any other events received by Unified Manager. Additionally, if you want to treat all performance events alike and have email sent to the same person, you can create a single alert to notify you when any critical or warning performance events are triggered.

Before you begin

You must have the Application Administrator or Storage Administrator role.

The example below shows how to create an event for all critical latency, IOPS, and MBps events. You can use this same methodology to select events from all performance counters, and for all warning events.

Steps

1. In the left navigation pane, click **Storage Management** > **Alert Setup**.

2. In the **Alert Setup** page, click **Add**.
3. In the **Add Alert** dialog box, click **Name**, and enter a name and description for the alert.
4. Do not select any resources on the **Resources** page.

   Because no resources are selected, the alert is applied to all clusters, aggregates, volumes, and so on, for which these events are received.

5. Click **Events** and perform the following actions:
   a. In the Event Severity list, select **Critical**.
   b. In the Event Name Contains field, enter `latency` and then click the arrow to select all the matching events.
   c. In the Event Name Contains field, enter `iops` and then click the arrow to select all the matching events.
   d. In the Event Name Contains field, enter `mbps` and then click the arrow to select all the matching events.
6. Click **Actions** and then select the name of the user who will receive the alert email in the **Alert these users** field.
7. Configure any other options on this page for issuing SNMP traps and executing a script.
8. Click **Save**.

Test alerts in Active IQ Unified Manager

You can test an alert to verify that you have configured it correctly. When an event is triggered, an alert is generated, and an alert email is sent to the configured recipients. You can verify whether the notification is sent and whether your script is executed by using the test alert.

Before you begin

- You must have configured notification settings such as the email address of the recipients, SMTP server, and SNMP trap.

  The Unified Manager server can use these settings to send notifications to users when an event is generated.

- You must have assigned a script and configured the script to run when the alert is generated.
- You must have the Application Administrator role.

Steps

1. In the left navigation pane, click **Storage Management** > **Alert Setup**.
2. In the **Alert Setup** page, select the alert that you want to test, and then click **Test**.

   A test alert email is sent to the email addresses that you specified while creating the alert.

Disable alerts for Resolved and Obsolete events in Active IQ Unified Manager

For all events that you have configured to send alerts, an alert message is sent when those events transition through all available states: New, Resolved, and Obsolete. If you

do not want to receive alerts for events as they move into the Resolved and Obsolete states, you can configure a global setting to suppress those alerts.

Before you begin

You must have the Application Administrator or Storage Administrator role.

By default, alerts are not sent for events as they move into the Resolved and Obsolete states.

Steps

1. In the left navigation pane, click **Storage Management** > **Alert Setup**.
2. In the **Alert Setup** page, perform one of the following actions using the slider control next to the item **Alerts for Resolved and Obsolete events**:

To…	Do this…
Stop sending alerts as events are resolved or obsoleted	Move the slider control to the left
Start sending alerts as events are resolved or obsoleted	Move the slider control to the right

Exclude disaster recovery destination volumes from generating alerts in Active IQ Unified Manager

When configuring volume alerts you can specify a string in the Alert dialog box that identifies a volume or group of volumes. If you have configured disaster recovery for SVMs, however, the source and destination volumes have the same name, so you will receive alerts for both volumes.

Before you begin

You must have the Application Administrator or Storage Administrator role.

You can disable alerts for disaster recovery destination volumes by excluding volumes that have the name of the destination SVM. This is possible because the identifier for volume events contains both the SVM name and volume name in the format "<svm_name>:/<volume_name>".

The example below shows how to create alerts for volume "vol1" on the primary SVM "vs1", but exclude the alert from being generated on a volume with the same name on SVM "vs1-dr".

Perform the following steps in the Add Alert dialog box:

Steps

1. Click **Name** and enter a name and description for the alert.
2. Click **Resources**, and then select the **Include** tab.
   a. Select **Volume** from the drop-down list, and then enter **`vol1`** in the **Name contains** field to display the volumes whose name contains "vol1".
   b. Select **<<All Volumes whose name contains 'vol1'>>** from the **Available Resources** area, and move it to the **Selected Resources** area.
3. Select the **Exclude** tab, select **Volume**, enter **`vs1-dr`** in the **Name contains** field, and then click **Add**.

This excludes the alert from being generated for volume "vol1" on SVM "vs1-dr".

4. Click **Events** and select the event or events that you want to apply to the volume or volumes.
5. Click **Actions** and then select the name of the user who will receive the alert email in the **Alert these users** field.
6. Configure any other options on this page for issuing SNMP traps and executing a script, and then click **Save**.

View alerts in Active IQ Unified Manager

You can view the list of alerts that is created for various events from the Alert Setup page. You can also view alert properties such as the alert description, notification method and frequency, events that trigger the alert, email recipients of the alerts, and affected resources such as clusters, aggregates, and volumes.

Before you begin

You must have the Operator, Application Administrator, or Storage Administrator role.

Step

1. In the left navigation pane, click **Storage Management** > **Alert Setup**.

   The list of alerts is displayed in the Alert Setup page.

Edit alerts in Active IQ Unified Manager

You can edit alert properties such as the resource with which the alert is associated, events, recipients, notification options, notification frequency, and associated scripts.

Before you begin

You must have the Application Administrator role.

Steps

1. In the left navigation pane, click **Storage Management** > **Alert Setup**.
2. In the **Alert Setup** page, select the alert that you want to edit, and click **Edit**.
3. In the **Edit Alert** dialog box, edit the name, resources, events, and actions sections, as required.

   You can change or remove the script that is associated with the alert.

4. Click **Save**.

Delete alerts in Active IQ Unified Manager

You can delete an alert when it is no longer required. For example, you can delete an alert that was created for a particular resource when that resource is no longer monitored by Unified Manager.

Before you begin

You must have the Application Administrator role.

Steps

1. In the left navigation pane, click **Storage Management** > **Alert Setup**.
2. On the **Alert Setup** page, select the alerts that you want to delete, and click **Delete**.
3. Click **Yes** to confirm the delete request.

Description of alert windows and dialog boxes

Learn about the Alert Setup page in Active IQ Unified Manager

The Alert Setup page displays a list of alerts and provides information about the alert name, status, notification method, and notification frequency. You can also add, edit, remove, enable, or disable alerts from this page.

You must have the Application Administrator or Storage Administrator role.

Command buttons

- **Add**

  Displays the Add Alert dialog box, which enables you to add new alerts.

- **Edit**

  Displays the Edit Alert dialog box, which enables you to edit selected alerts.

- **Delete**

  Deletes the selected alerts.

- **Enable**

  Enables the selected alerts to send notifications.

- **Disable**

  Disables the selected alerts when you want to temporarily stop sending notifications.

- **Test**

  Tests the selected alerts to verify their configuration after being added or edited.

- **Alerts for Resolved and Obsolete Events**

  Allows you to enable or disable the sending of alerts when events are moved to the Resolved or Obsolete states. This can help users from receiving unnecessary notifications.

List view

The list view displays, in tabular format, information about the alerts that are created. You can use the column filters to customize the data that is displayed. You can also select an alert to view more information about it in the details area.

- **Status**

Specifies whether an alert is enabled () or disabled ().

- **Alert**

  Displays the name of the alert.

- **Description**

  Displays a description for the alert.

- **Notification Method**

  Displays the notification method that is selected for the alert. You can notify users through email or SNMP traps.

- **Notification Frequency**

  Specifies the frequency (in minutes) with which the management server continues to send notifications until the event is acknowledged, resolved, or moved to the Obsolete state.

Details area

The details area provides more information about the selected alert.

- **Alert Name**

  Displays the name of the alert.

- **Alert Description**

  Displays a description for the alert.

- **Events**

  Displays the events for which you want to trigger the alert.

- **Resources**

  Displays the resources for which you want to trigger the alert.

- **Includes**

  Displays the group of resources for which you want to trigger the alert.

- **Excludes**

  Displays the group of resources for which you do not want to trigger the alert.

- **Notification Method**

  Displays the notification method for the alert.

- **Notification Frequency**

  Displays the frequency with which the management server continues to send alert notifications until the event is acknowledged, resolved, or moved to the Obsolete state.

- **Script Name**

  Displays the name of the script associated with the selected alert. This script is executed when an alert is generated.

- **Email Recipients**

  Displays the email addresses of users who receive the alert notification.

Learn about the Add Alert dialog box in Active IQ Unified Manager

You can create alerts to notify you when a particular event is generated, so that you can address the issue quickly and thereby minimize impact to your environment. You can create alerts for a single resource or a set of resources, and for events of a particular severity type. You can also specify the notification method and frequency of the alerts.

You must have the Application Administrator or Storage Administrator role.

Name

This area enables you to specify a name and description for the alert:

- **Alert Name**

  Enables you to specify an alert name.

- **Alert Description**

  Enables you to specify a description for the alert.

Resources

This area enables you to select an individual resource or group the resources based on a dynamic rule for which you want to trigger the alert. A *dynamic rule* is the set of resources filtered based on the text string you specify. You can search for resources by selecting a resource type from the drop-down list or you can specify the exact resource name to display a specific resource.

If you are creating an alert from any of the storage object details pages, the storage object is automatically included in the alert.

- **Include**

  Enables you to include the resources for which you want to trigger alerts. You can specify a text string to group resources that match the string and select this group to be included in the alert. For example, you can group all volumes whose name contains the "abc" string.

- **Exclude**

  Enables you to exclude resources for which you do not want to trigger alerts. For example, you can exclude all volumes whose name contains the "xyz" string.

  The Exclude tab is displayed only when you select all resources of a particular resource type: for example, <<All Volumes>> or <<All Volumes whose name contains 'xyz'>>.

  If a resource conforms to both the include and exclude rules that you have specified, the exclude rule takes

precedence over the include rule and the alert is not generated for the event.

Events

This area enables you to select the events for which you want to create the alerts. You can create alerts for events based on a particular severity or for a set of events.

To select more than one event, you should hold down the Ctrl key while you make your selections.

- **Event Severity**

  Enables you to select events based on the severity type, which can be Critical, Error, or Warning.

- **Event Name Contains**

  Enables you to select events whose name contains specified characters.

Actions

This area enables you to specify the users that you want to notify when an alert is triggered. You can also specify the notification method and the frequency of notification.

- **Alert these users**

  Enables you to specify the email address or user name of the user to receive notifications.

  If you modify the email address that is specified for the user and reopen the alert for editing, the Name field appears blank because the modified email address is no longer mapped to the user that was previously selected. Also, if you have modified the email address of the selected user from the Users page, the modified email address is not updated for the selected user.

- **Notification Frequency**

  Enables you to specify the frequency with which the management server sends notifications until the event is acknowledged, resolved, or moved to the obsolete state.

  You can choose the following notification methods:

  - Notify only once
  - Notify at a specified frequency
  - Notify at a specified frequency within the specified time range

- **Issue SNMP trap**

  Selecting this box enables you to specify whether SNMP traps should be sent to the globally configured SNMP host.

- **Execute Script**

  Enables you to add your custom script to the alert. This script is executed when an alert is generated.



If you do not see this capability available in the user interface it is because the functionality has been disabled by your administrator. If required, you can enable this functionality from **Storage Management** > **Feature Settings**.

Command buttons

- **Save**

  Creates an alert and closes the dialog box.

- **Cancel**

  Discards the changes and closes the dialog box.

Edit alert properties in Active IQ Unified Manager

You can edit alert properties such as the resource with which the alert is associated, events, script, and notification options.

Name

This area enables you to edit the name and description for the alert.

- **Alert Name**

  Enables you to edit the alert name.

- **Alert Description**

  Enables you to specify a description for the alert.

- **Alert State**

  Enables you to enable or disable the alert.

Resources

This area enables you to select an individual resource or group the resources based on a dynamic rule for which you want to trigger the alert. You can search for resources by selecting a resource type from the drop-down list or you can specify the exact resource name to display a specific resource.

- **Include**

  Enables you to include the resources for which you want to trigger alerts. You can specify a text string to group resources that match the string and select this group to be included in the alert. For example, you can group all volumes whose name contains the “vol0” string.

- **Exclude**

  Enables you to exclude resources for which you do not want to trigger alerts. For example, you can exclude all volumes whose name contains the “xyz” string.

  The Exclude tab is displayed only when you select all resources of a particular resource type—for example, <<All Volumes>> or <<All Volumes whose name contains 'xyz'>>.

Events

This area enables you to select the events for which you want to trigger the alerts. You can trigger an alert for events based on a particular severity or for a set of events.

- **Event Severity**

  Enables you to select events based on the severity type, which can be Critical, Error, or Warning.

- **Event Name Contains**

  Enables you to select events whose name contains the specified characters.

Actions

This area enables you to specify the notification method and the frequency of notification.

- **Alert these users**

  Enables you to edit the email address or user name, or specify a new email address or user name to receive notifications.

- **Notification Frequency**

  Enables you to edit the frequency with which the management server sends notifications until the event is acknowledged, resolved, or moved to the obsolete state.

  You can choose the following notification methods:

  - Notify only once
  - Notify at a specified frequency
  - Notify at a specified frequency within the specified time range

- **Issue SNMP trap**

  Enables you to specify whether SNMP traps should be sent to the globally configured SNMP host.

- **Execute Script**

  Enables you to associate a script with the alert. This script is executed when an alert is generated.

Command buttons

- **Save**

  Saves the changes and closes the dialog box.

- **Cancel**

  Discards the changes and closes the dialog box.

Manage scripts

Manage scripts in Active IQ Unified Manager

You can use scripts to automatically modify or update multiple storage objects in Unified Manager. The script is associated with an alert. When an event triggers an alert, the script is executed. You can upload custom scripts and test their execution when an alert is

generated.

The ability to upload scripts to Unified Manager and run them is enabled by default. If your organization does not want to allow this functionality because of security reasons, you can disable this functionality from **Storage Management** > **Feature Settings**.

Related information

Enabling and disabling the ability to upload scripts

How scripts work with alerts in Active IQ Unified Manager

You can associate an alert with your script so that the script is executed when an alert is raised for an event in Unified Manager. You can use the scripts to resolve issues with storage objects or identify which storage objects are generating the events.

When an alert is generated for an event in Unified Manager, an alert email is sent to the specified recipients. If you have associated an alert with a script, the script is executed. You can get the details of the arguments passed to the script from the alert email.

If you have created a custom script and associated it with an alert for a specific event type, actions are taken based on your custom script for that event type, and the **Fix it** actions are not available by default on the Management Actions page or Unified Manager dashboard.

The script uses the following arguments for execution:

- `-eventID`
- `-eventName`
- `-eventSeverity`
- `-eventSourceID`
- `-eventSourceName`
- `-eventSourceType`
- `-eventState`
- `-eventArgs`

You can use the arguments in your scripts and gather related event information or modify storage objects.

Example for obtaining arguments from scripts

```
`print "$ARGV[0] : $ARGV[1]\n"`
`print "$ARGV[7] : $ARGV[8]\n"`
```

When an alert is generated, this script is executed and the following output is displayed:

```
-`eventID : 290`
```

```
-`eventSourceID : 4138`
```

Add scripts in Active IQ Unified Manager

You can add scripts in Unified Manager, and associate the scripts with alerts. These scripts are executed automatically when an alert is generated, and enable you to obtain information about storage objects for which the event is generated.

Before you begin

- You must have created and saved the scripts that you want to add to the Unified Manager server.
- The supported file formats for scripts are Perl, Shell, PowerShell, Python, and `.bat` files.

Platform on which Unified Manager is installed	Supported languages
VMware	Perl and Shell scripts
Linux	Perl, Python, and Shell scripts
Windows	PowerShell, Perl, Python, and .bat scripts

  - For Perl scripts, Perl must be installed on the Unified Manager server. For VMware installations, Perl 5 is installed by default and scripts will support only what Perl 5 supports. If Perl was installed after Unified Manager, you must restart the Unified Manager server.
  - For PowerShell scripts, the appropriate PowerShell execution policy must be set on the Windows server so that the scripts can be executed.

  If your script creates log files to track the alert script progress, you must make sure that the log files are not created anywhere within the Unified Manager installation folder.

- You must have the Application Administrator or Storage Administrator role.

You can upload custom scripts and gather event details about the alert.

If you do not see this capability available in the user interface it is because the functionality has been disabled by your administrator. If required, you can enable this functionality from **Storage Management** > **Feature Settings**.

Steps

1. In the left navigation pane, click **Storage Management** > **Scripts**.
2. In the **Scripts** page, click **Add**.
3. In the **Add Script** dialog box, click **Browse** to select your script file.
4. Enter a description for the script that you select.
5. Click **Add**.

Related information

Enabling and disabling the ability to upload scripts

Delete scripts in Active IQ Unified Manager

You can delete a script from Unified Manager when the script is no longer required or valid.

Before you begin

- You must have the Application Administrator or Storage Administrator role.
- The script must not be associated with an alert.

Steps

1. In the left navigation pane, click **Storage Management** > **Scripts**.
2. In the **Scripts** page, select the script that you want to delete, and then click **Delete**.
3. In the **Warning** dialog box, confirm the deletion by clicking **Yes**.

Test script execution in Active IQ Unified Manager

You can verify that your script is executed correctly when an alert is generated for a storage object.

Before you begin

- You must have the Application Administrator or Storage Administrator role.
- You must have uploaded a script in the supported file format to Unified Manager.

Steps

1. In the left navigation pane, click **Storage Management** > **Scripts**.
2. In the **Scripts** page, add your test script.
3. In the left navigation pane, click **Storage Management** > **Alert Setup**.
4. In the **Alert Setup** page, perform one of the following actions:

To…	Do this…
Add an alert	a. Click **Add**. b. In the Actions section, associate the alert with your test script.
Edit an alert	a. Select an alert, and then click **Edit**. b. In the Actions section, associate the alert with your test script.

5. Click **Save**.
6. In the **Alert Setup** page, select the alert that you added or modified, and then click **Test**.

   The script is executed with the "-test" argument, and a notification alert is sent to the email addresses that were specified when the alert was created.

Learn about supported Active IQ Unified Manager CLI commands

As a storage administrator you can use the CLI commands to perform queries on the

storage objects; for example, on clusters, aggregates, volumes, qtrees, and LUNs. You can use the CLI commands to query the Unified Manager internal database and the ONTAP database. You can also use CLI commands in scripts that are executed at the beginning or end of an operation or are executed when an alert is triggered.

All commands must be preceded with the command `um cli login` and a valid user name and password for authentication.

For running the *um run* command, ensure that your account has the *console* application access.

CLI command	Description	Output
`um cli login -u <username> [-p <password>]`	Logs in to the CLI. Because of security implications, you should enter only the user name following the "-u" option. When used in this manner you will be prompted for the password, and the password will not be captured in the history or process table. The session expires after three hours from the time of login, after which the user must login again.	Displays the corresponding message.
`um cli logout`	Logs out of the CLI.	Displays the corresponding message.
`um help`	Displays all first level subcommands.	Displays all first level subcommands.
`um run cmd [ -t <timeout> ] <cluster> <command>`	The simplest way to run a command on one or more hosts. Mainly used for alert scripting to get or perform an operation on ONTAP. The optional timeout argument sets a maximum time limit (in seconds) for the command to complete on the client. The default is 0 (wait forever).	As received from ONTAP.
`um run query <sql command>`	Executes an SQL query. Only queries that read from the database are allowed. Any update, insert, or delete operations are not supported.	Results are displayed in a tabular form. If an empty set is returned, or if there is any syntax error or bad request, it displays the appropriate error message.

CLI command	Description	Output
`um datasource add -u <username> -P <password> [ -t <protocol> ] [ -p <port> ] <hostname-or-ip>`	Adds a datasource to the list of managed storage systems. A datasource describes how connections to storage systems are made. The options -u (username) and -P (password) must be specified when adding a datasource. The option -t (protocol) specifies the protocol used to communicate with the cluster (http or https). If the protocol is not specified, then both protocols will be attempted The option -p (port) specifies the port used to communicate with the cluster. If the port is not specified, then the default value of the appropriate protocol will be attempted. This command can be executed only by the storage admin.	Prompts for the user accept the certificate and prints the corresponding message.
`um datasource list [ <datasource-id>]`	Displays the datasources for managed storage systems.	Displays the following values in tabular format: `ID Address Port, Protocol Acquisition Status, Analysis Status, Communication status, Acquisition Message, and Analysis Message.`
`um datasource modify [ -h <hostname-or-ip> ] [ -u <username> ] [ -P <password> ] [ -t <protocol> ] [ -p <port> ] <datasource-id>`	Modifies one or more datasource options. Can be executed only by the storage admin.	Displays the corresponding message.
`um datasource remove <datasource-id>`	Removes the datasource (cluster) from Unified Manager.	Displays the corresponding message.
`um option list [ <option> .. ]`	Lists all the options that you can configure using the set command.	Displays the following values in tabular format: `Name, Value, Default Value, and Requires Restart.`
`um option set <option-name>=<option-value> [ <option-name>=<option-value> … ]`	Sets one or more options. The command can be executed only by the storage admin.	Displays the corresponding message.
`um version`	Displays the Unified Manager software version.	`Version ("9.6")`

CLI command	Description	Output
`um lun list [-q] [ -ObjectType <object-id>]`	Lists the LUNs after filtering on the specified object. -q is applicable for all commands to show no header. ObjectType can be lun, qtree, cluster, volume, quota, or svm. For example: **`um lun list -cluster 1`** In this example, "-cluster" is the objectType and "1" is the objectId. The command lists all the LUNs within the cluster with ID 1.	Displays the following values in tabular format: `ID and LUN path.`
`um svm list [-q] [ -ObjectType <object-id>]`	Lists the storage VMs after filtering on the specified object. ObjectType can be lun, qtree, cluster, volume, quota, or svm. For example: **`um svm list -cluster 1`** In this example, "-cluster" is the objectType and "1" is the objectId. The command lists all the storage VMs within the cluster with ID 1.	Displays the following values in tabular format: `Name and Cluster ID.`
`um qtree list [-q] [ -ObjectType <object-id>]`	Lists the qtrees after filtering on the specified object. -q is applicable for all commands to show no header. ObjectType can be lun, qtree, cluster, volume, quota, or svm. For example: **`um qtree list -cluster 1`** In this example, "-cluster" is the objectType and "1" is the objectId. The command lists all the qtrees within the cluster with ID 1.	Displays the following values in tabular format: `Qtree ID and Qtree Name.`

CLI command	Description	Output
`um disk list [-q] [-ObjectType <object-id>]`	Lists the disks after filtering on the specified object. ObjectType can be disk, aggr, node, or cluster. For example: **`um disk list -cluster 1`** In this example, "-cluster" is the objectType and "1" is the objectId. The command lists all the disks within the cluster with ID 1.	Displays the following values in tabular format `ObjectType and object-id`.
`um cluster list [-q] [-ObjectType <object-id>]`	Lists the clusters after filtering on the specified object. ObjectType can be disk, aggr, node, cluster, lun, qtree, volume, quota, or svm. For example: **`um cluster list -aggr 1`** In this example, "-aggr" is the objectType and "1" is the objectId. The command lists the cluster to which the aggregate with ID 1 belongs.	Displays the following values in tabular format: `Name, Full Name, Serial Number, Datasource Id, Last Refresh Time, and Resource Key`.
`um cluster node list [-q] [-ObjectType <object-id>]`	Lists the cluster nodes after filtering on the specified object. ObjectType can be disk, aggr, node, or cluster. For example: **`um cluster node list -cluster 1`** In this example, "-cluster" is the objectType and "1" is the objectId. The command lists all the nodes within the cluster with ID 1.	Displays the following values in tabular format `Name and Cluster ID`.

CLI command	Description	Output
`um volume list [-q] [-ObjectType <object-id>]`	Lists the volumes after filtering on the specified object. ObjectType can be lun, qtree, cluster, volume, quota, svm, or aggregate. For example: **`um volume list -cluster 1`** In this example, "-cluster" is the objectType and "1" is the objectId. The command lists all the volumes within the cluster with ID 1.	Displays the following values in tabular format `Volume ID and Volume Name`.
`um quota user list [-q] [-ObjectType <object-id>]`	Lists the quota users after filtering on the specified object. ObjectType can be qtree, cluster, volume, quota, or svm. For example: **`um quota user list -cluster 1`** In this example, "-cluster" is the objectType and "1" is the objectId. The command lists all the quota users within the cluster with ID 1.	Displays the following values in tabular format `ID, Name, SID and Email`.
`um aggr list [-q] [-ObjectType <object-id>]`	Lists the aggregates after filtering on the specified object. ObjectType can be disk, aggr, node, cluster, or volume. For example: **`um aggr list -cluster 1`** In this example, "-cluster" is the objectType and "1" is the objectId. The command lists all the aggregates within the cluster with ID 1.	Displays the following values in tabular format `Aggr ID, and Aggr Name`.
`um event ack <event-ids>`	Acknowledges one or more events.	Displays the corresponding message.
`um event resolve <event-ids>`	Resolves one or more events.	Displays the corresponding message.
`um event assign -u <username> <event-id>`	Assigns an event to a user.	Displays the corresponding message.

CLI command	Description	Output
`um event list [ -s <source> ] [ -S <event-state-filter-list>.. ] [ <event-id> .. ]`	Lists the events generated by the system or user. Filters events based on source, state, and IDs.	Displays the following values in tabular format `Source, Source type, Name, Severity, State, User and Timestamp`.
`um backup restore -f <backup_file_path_and_name>`	Restores a MySQL database backup using .7z files.	Displays the corresponding message.

Description of script windows and dialog boxes

Learn about the Scripts page in Active IQ Unified Manager

The Scripts page enables you to add your custom scripts to Unified Manager. You can associate these scripts with alerts to enable automatic reconfiguration of storage objects.

The Scripts page enables you to add or delete scripts from Unified Manager.

Command buttons

- **Add**

  Displays the Add Script dialog box, which enables you to add scripts.

- **Delete**

  Deletes the selected script.

List view

The list view displays, in tabular format, the scripts that you added to Unified Manager.

- **Name**

  Displays the name of the script.

- **Description**

  Displays the description of the script.

Add Script dialog box in Active IQ Unified Manager

The Add Script dialog box enables you to add scripts to Unified Manager. You can configure alerts with your scripts to automatically resolve events that are generated for storage objects.

You must have the Application Administrator or Storage Administrator role.

- **Select Script File**

  Enables you to select a script for the alert.

- **Description**

  Enables you to specify a description for the script.

Monitor and manage cluster performance

Introduction to Active IQ Unified Manager performance monitoring

Learn about performance monitoring in Active IQ Unified Manager

Active IQ Unified Manager (formerly OnCommand Unified Manager) provides performance monitoring capabilities and event root-cause analysis for systems that are running NetApp ONTAP software.

Unified Manager helps you to identify workloads that are overusing cluster components and decreasing the performance of other workloads on the cluster. By defining performance threshold policies you can also specify maximum values for certain performance counters so that events are generated when the threshold is breached. Unified Manager alerts you about these performance events so that you can take corrective action, and bring performance back to normal levels of operation. You can view and analyze events in the Unified Manager UI.

Unified Manager monitors the performance of two types of workloads:

- User-defined workloads

  These workloads consist of FlexVol volumes and FlexGroup volumes that you have created in your cluster.

- System-defined workloads

  These workloads consist of internal system activity.

Explore Active IQ Unified Manager performance monitoring features

Unified Manager collects and analyzes performance statistics from systems running ONTAP software. It uses dynamic performance thresholds and user-defined performance thresholds to monitor a variety of performance counters over many cluster components.

A high response time (latency) indicates that the storage object, for example, a volume, is performing slower than normal. This issue also indicates that the performance has decreased for client applications that are using the volume. Unified Manager identifies the storage component where the performance issue lies and provides a list of suggested actions you can take to address the performance issue.

Unified Manager includes the following features:

- Monitors and analyzes workload performance statistics from a system running ONTAP software.
- Tracks performance counters for clusters, nodes, aggregates, ports, SVMs, volumes, LUNs, NVMe namespaces, and network interfaces (LIFs).
- Displays detailed graphs that plot workload activity over time; including IOPS (operations), MB/s (throughput), latency (response time), utilization, performance capacity, and cache ratio.
- Enables you to create user-defined performance threshold policies that trigger events and send email alerts when the thresholds are breached.
- Uses system-defined thresholds and dynamic performance thresholds that learn about your workload

activity to identify and alert you to performance issues.

- Identifies the quality of service (QoS) policies and Performance Service Level policies (PSLs) that are applied to your volumes and LUNs.
- Clearly identifies the cluster component that is in contention.
- Identifies workloads that are overusing cluster components and the workloads whose performance is impacted by the increased activity.

Use Active IQ Unified Manager interfaces to manage storage system performance

These sections contain information about the two user interfaces that Active IQ Unified Manager provides for troubleshooting data storage capacity, availability, and protection issues. The two UIs are the Unified Manager web UI and the maintenance console.

If you want to use the protection features in Unified Manager, you must also install and configure OnCommand Workflow Automation (WFA).

Unified Manager web UI

The Unified Manager web UI enables an administrator to monitor and troubleshoot cluster issues relating to data storage capacity, availability, and protection.

These sections describe some common workflows that an administrator can follow to troubleshoot storage capacity, data availability, or protection issues displayed in the Unified Manager web UI.

Maintenance console

The Unified Manager maintenance console enables an administrator to monitor, diagnose, and address operating system issues, version upgrade issues, user access issues, and network issues related to the Unified Manager server itself. If the Unified Manager web UI is unavailable, the maintenance console is the only form of access to Unified Manager.

You can use this information for accessing the maintenance console and using it to resolve issues related to the functioning of the Unified Manager server.

View cluster configuration and performance data in Active IQ Unified Manager

The collection interval for *cluster configuration data* is 15 minutes. For example, after you have added a cluster, it takes 15 minutes to display the cluster details in the Unified Manager UI. This interval applies when making changes to a cluster too.

For example, if you add two new volumes to an SVM in a cluster, you see those new objects in the UI after the next polling interval, which could be up to 15 minutes.

Unified Manager collects current *performance statistics* from all monitored clusters every five minutes. It analyzes this data to identify performance events and potential issues. It retains 30 days of five-minute historical performance data and 180 days of one-hour historical performance data. This enables you to view very granular performance details for the current month, and general performance trends for up to a year.

The collection polls are offset by a few minutes so that data from every cluster is not sent at the same time, which could affect performance.

The following table describes the collection activities that Unified Manager performs:

Activity	Time interval	Description
Performance statistics poll	Every 5 minutes	Collects real-time performance data from each cluster.
Statistical analysis	Every 5 minutes	After every statistics poll, Unified Manager compares the collected data against user-defined, system-defined, and dynamic thresholds. If any performance thresholds have been breached, Unified Manager generates events and sends email to specified users, if configured to do so.
Configuration poll	Every 15 minutes	Collects detailed inventory information from each cluster to identify all the storage objects (nodes, SVMs, volumes, and so on).
Summarization	Every hour	Summarizes the latest 12 five-minute performance data collections into hourly averages. The hourly average values are used in some of the UI pages, and they are retained for 180 days.
Forecast analysis and data pruning	Every day after midnight	Analyzes cluster data to establish dynamic thresholds for volume latency and IOPS for the next 24 hours. Deletes from the database any five-minute performance data older than 30 days.
Data pruning	Every day after 2 a.m.	Deletes from the database any events older than 180 days and dynamic thresholds older than 180 days.
Data pruning	Every day after 3:30 a.m.	Deletes from the database any one-hour performance data older than 180 days.

Understand data continuity collection cycles in Active IQ Unified Manager

A data continuity collection cycle retrieves performance data outside of the real-time cluster performance collection cycle that runs, by default, every five minutes. Data continuity collections enable Unified Manager to fill in gaps of statistical data that occur when it was unable to collect real-time data.

Unified Manager performs data continuity collection polls of historical performance data when the following events occur:

- A cluster is initially added to Unified Manager.

  Unified Manager gathers historical performance data for the previous 15 days. This enables you to view two weeks of historical performance information for a cluster a few hours after it is added.

  Additionally, system-defined threshold events are reported for the previous period, if any exist.

- The current performance data collection cycle does not finish on time.

  If the real-time performance poll goes beyond the five-minute collection period, a data continuity collection cycle is initiated to gather that missing information. Without the data continuity collection, the next collection period is skipped.

- Unified Manager has been inaccessible for a period of time and then it comes back online, as in the following situations:
  - It was restarted.
  - It was shut down during a software upgrade or when creating a backup file.
  - A network outage is repaired.
- A cluster has been inaccessible for a period of time and then it comes back online, as in the following situations:
  - A network outage is repaired.
  - A slow wide area network connection delayed the normal collection of performance data.

A data continuity collection cycle can collect a maximum of 24 hours of historical data. If Unified Manager is down for longer than 24 hours, a gap in performance data appears in the UI pages.

A data continuity collection cycle and a real-time data collection cycle cannot run at the same time. The data continuity collection cycle must finish before the real-time performance data collection is initiated. When the data continuity collection is required to collect more than one hour of historical data, then you see a banner message for that cluster at the top of the Notifications pane.

Learn about timestamps in collected data and events in Active IQ Unified Manager

The timestamp that appears in collected health and performance data, or that appears as the detection time for an event, is based on the ONTAP cluster time, adjusted to the time zone set on the web browser.

It is highly recommended that you use a Network Time Protocol (NTP) server to synchronize the time on your Unified Manager servers, ONTAP clusters, and web browsers.

If you see timestamps that look incorrect for a particular cluster, you might want to check that the cluster time has been set correctly.

Navigating performance workflows in the Unified Manager GUI

Navigate performance workflows in the Active IQ Unified Manager GUI

The Unified Manager interface provides many pages for the collection and display of

performance information. You use the left navigation panel to navigate to pages in the GUI, and you use tabs and links on the pages to view and configure information.

You use all of the following pages to monitor and troubleshoot cluster performance information:

- dashboard page
- storage and network object inventory pages
- storage object details pages (including the performance explorer)
- configuration and setup pages
- events pages

Log in to the Active IQ Unified Manager UI

You can log in to the Unified Manager UI using a supported web browser.

Before you begin

- The web browser must meet minimum requirements.

  See the Interoperability Matrix at mysupport.netapp.com/matrix for the complete list of supported browser versions.

- You must have the IP address or URL of the Unified Manager server.

You are automatically logged out of the session after 1 hour of inactivity. This timeframe can be configured under **General** > **Feature Settings**.

Steps

1. Enter the URL in your web browser, where URL is the IP address or fully qualified domain name (FQDN) of the Unified Manager server:
   - For IPv4: `https://URL/`
   - For IPv6: `https://[URL]/`

If the server uses a self-signed digital certificate, the browser might display a warning indicating that the certificate is not trusted. You can either acknowledge the risk to continue the access or install a Certificate Authority (CA) signed digital certificate for server authentication.
. At the login screen, enter your user name and password.

If login to the Unified Manager user interface is protected using SAML authentication you will enter your credentials in the identity provider (IdP) login page instead of the Unified Manager login page.

The Dashboard page is displayed.

If the Unified Manager server is not initialized, a new browser window displays the first experience wizard. You must enter an initial email recipient to which email alerts will be sent, the SMTP server that will handle email communications, and whether AutoSupport is enabled to send information about your Unified Manager installation to technical support. The Unified Manager UI appears after you complete this information.

Graphical interface and navigational paths

Explore graphical interface and navigational paths in Active IQ Unified Manager

Unified Manager has great flexibility and enables you to accomplish multiple tasks in various ways. There are many navigation paths you will discover as you work in Unified Manager. While not all of the possible combinations of navigations can be shown, you should be familiar with a few of the more common scenarios.

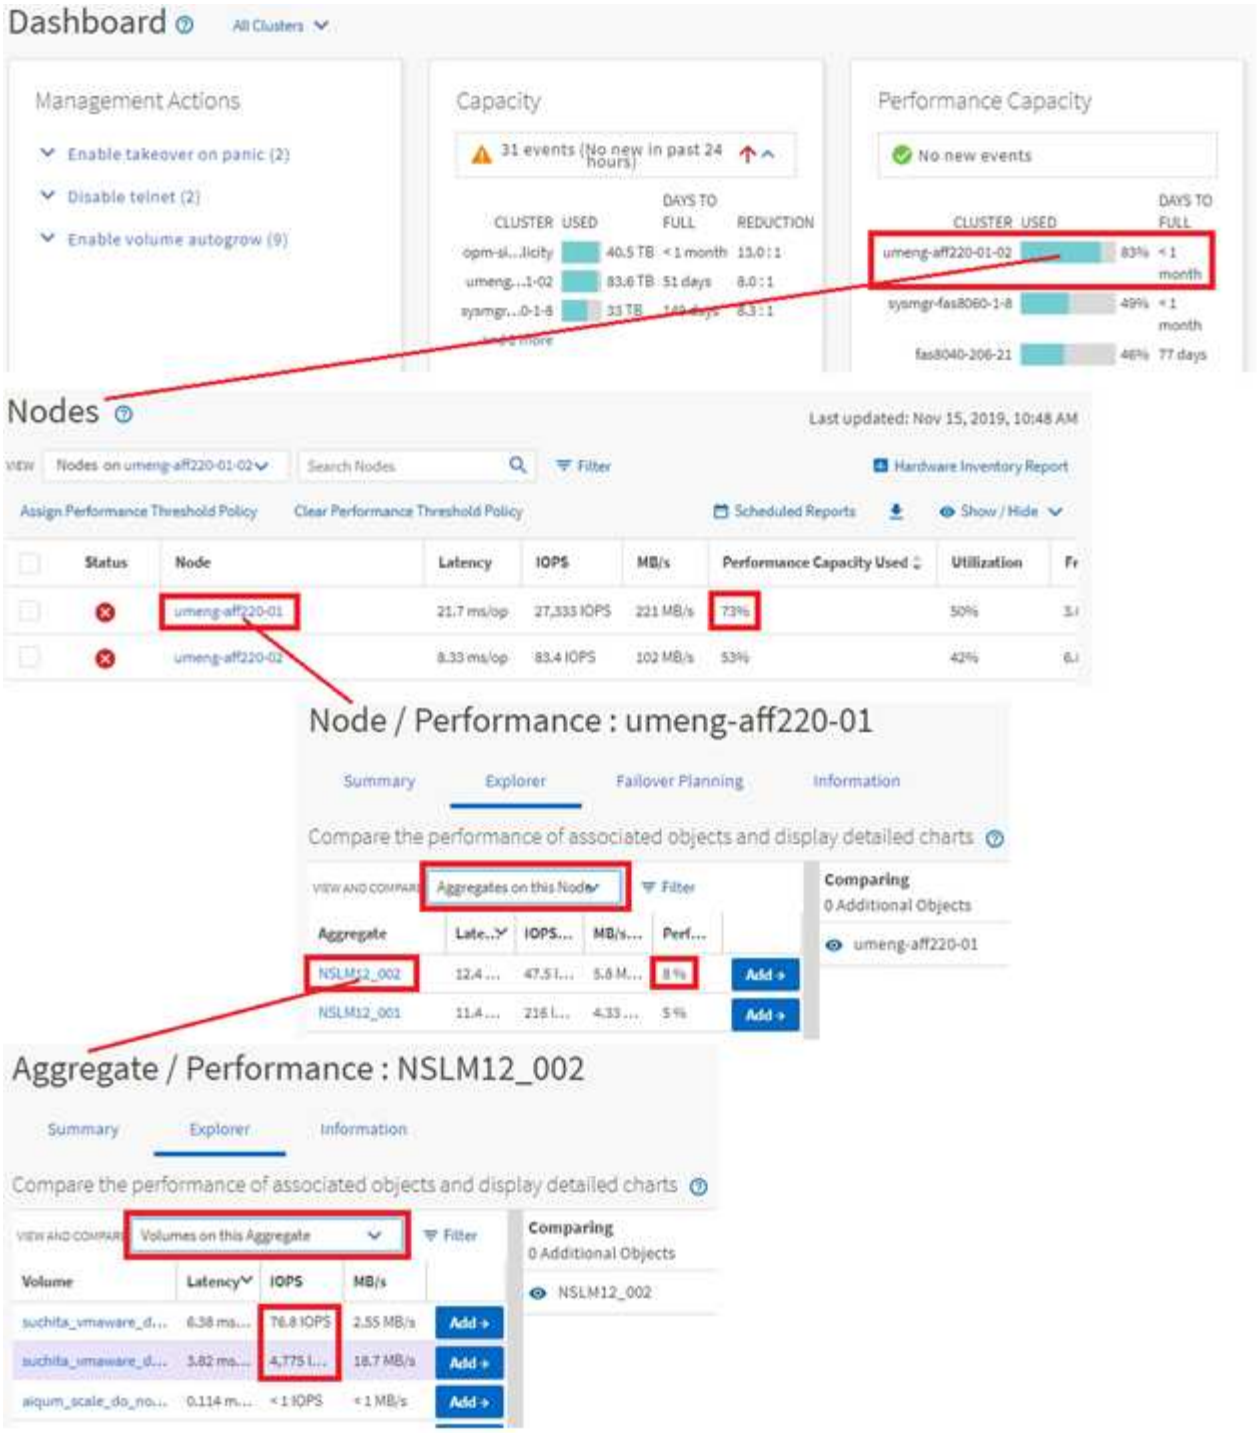
Monitor cluster object navigation in Active IQ Unified Manager

You can monitor the performance of all objects in any cluster managed by Unified Manager. Monitoring your storage objects provides you with an overview of cluster and object performance, and includes performance event monitoring. You can view performance and events at a high level, or you can further investigate any details of object performance and performance events.

This is one example of many possible cluster object navigations:

1. From the Dashboard page, review the details in the Performance Capacity panel to identify the cluster that is using the most performance capacity and click the bar chart to navigate to the list of nodes for that cluster.
2. Identify the node with the highest performance capacity used value and click that node.
3. From the Node / Performance Explorer page, click **Aggregates on this Node** from the View and Compare menu.
4. Identify the aggregate that is using the most performance capacity and click that aggregate.
5. From the Aggregate / Performance Explorer page, click **Volumes on this Aggregate** from the View and Compare menu.
6. Identify the volumes that are using the most IOPS.

You should investigate these volumes to see if you should apply a QoS policy or Performance Service Level policy, or change the policy settings, so that those volumes do not use such a large percentage of IOPS on the cluster.



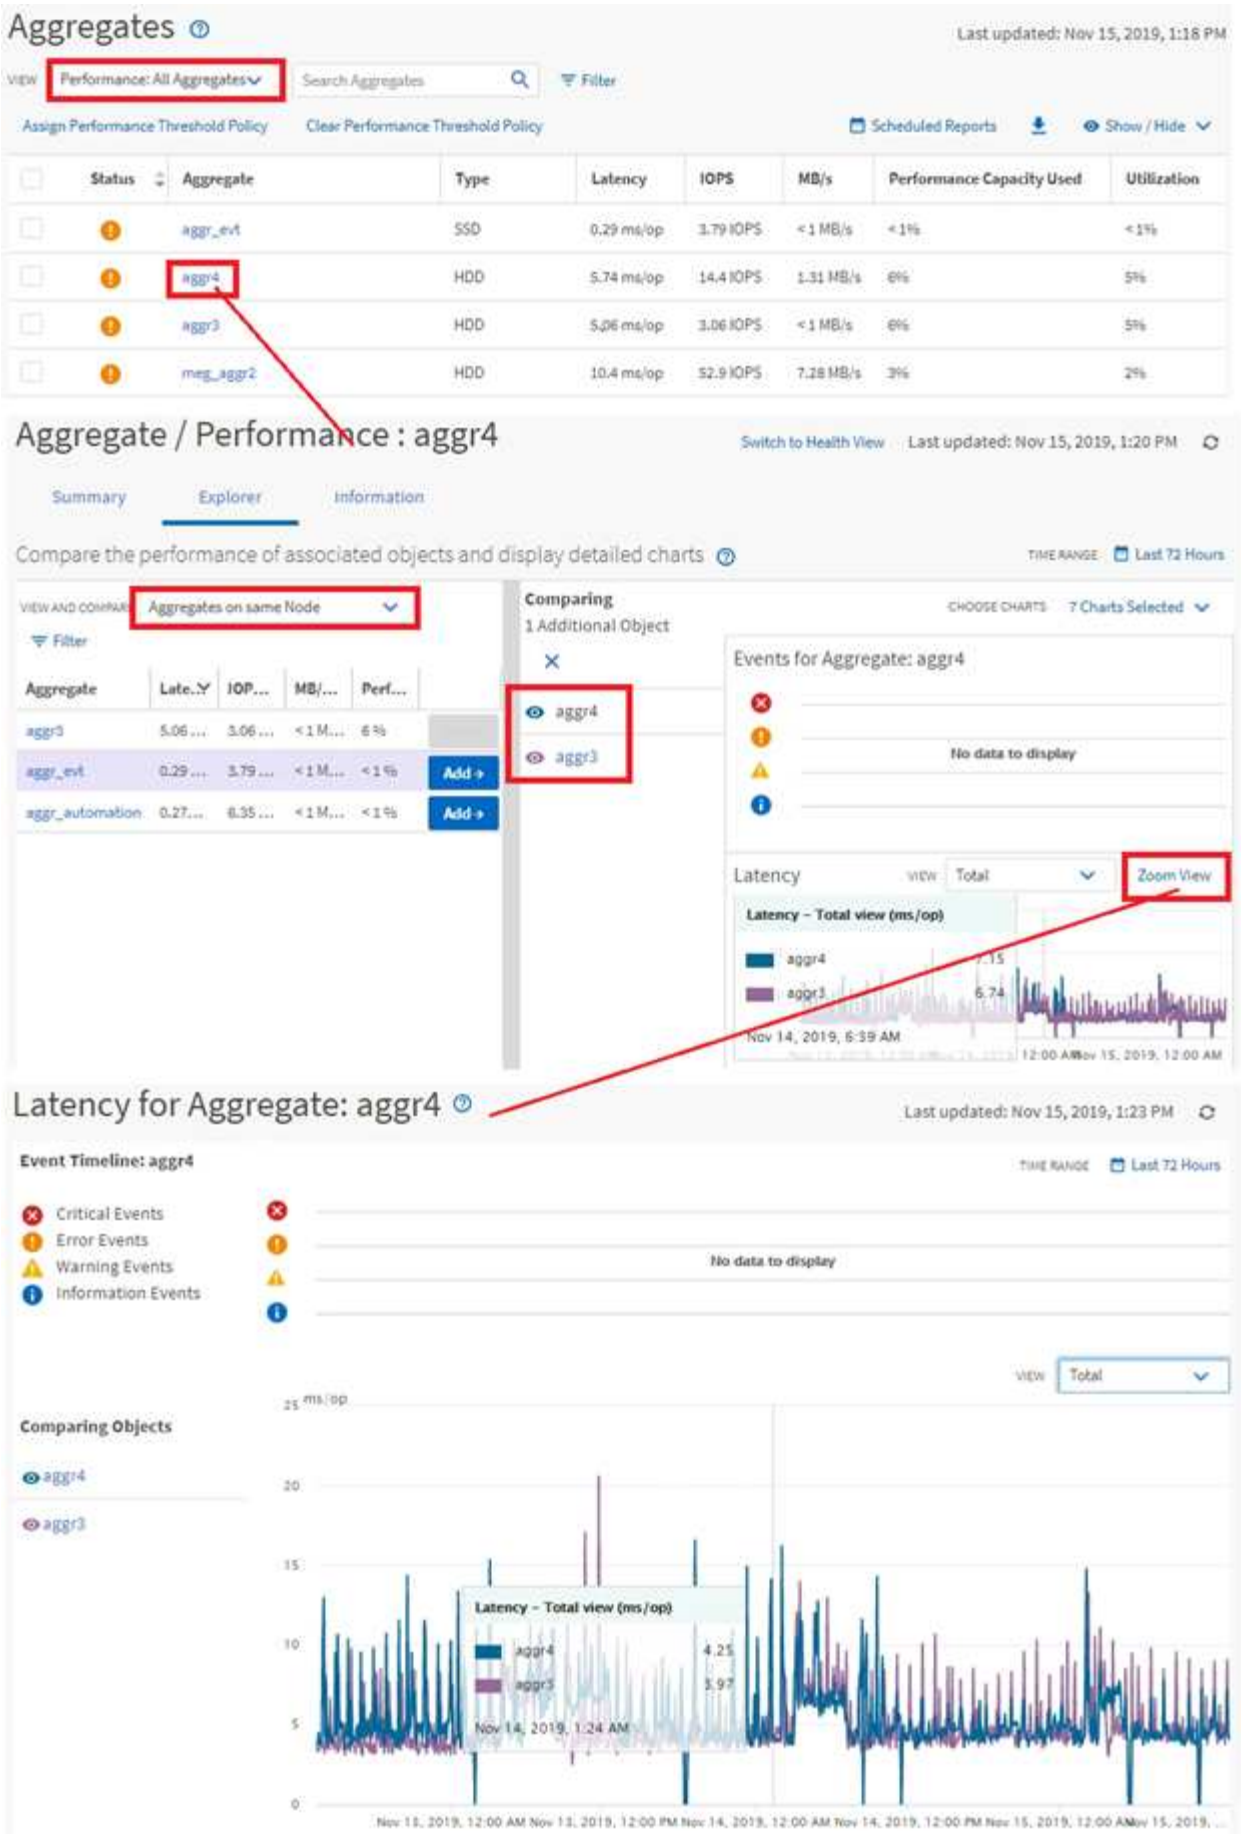
Monitor cluster performance navigation in Active IQ Unified Manager

You can monitor the performance of all clusters managed by Unified Manager. Monitoring your clusters provides you with an overview of cluster and object performance and includes performance event monitoring. You can view performance and events at a high level, or you can further investigate any details of cluster and object performance and performance events.

This is one example of many possible cluster performance navigational paths:

1. In the left navigation pane, click **Storage** > **Aggregates**.
2. To view information about the performance in those aggregates, select the Performance: All Aggregates view.
3. Identify the aggregate you want to investigate and click that aggregate name to navigate to the Aggregate / Performance Explorer page.
4. Optionally, select other objects to compare with this aggregate in the View and Compare menu, and then add one of the objects to the comparing pane.

   Statistics for both objects will appear in the counter charts for comparison.
5. In the Comparing pane at the right on the Explorer page, click **Zoom View** in one of the counter charts to view details about the performance history for that aggregate.


Aggregates
Last updated: Nov 15, 2019, 1:18 PM
VIEW
Performance: All Aggregates
Search Aggregates
Filter
Assign Performance Threshold Policy
Clear Performance Threshold Policy
Scheduled Reports
Show / Hide
Status
Aggregate
Type
Latency
IOPS
MB/s
Performance Capacity Used
Utilization
aggr_evt
SSD
0.29 ms/op
3.79 IOPS
< 1 MB/s
< 1%
aggr4
HDD
5.74 ms/op
14.4 IOPS
1.31 MB/s
6%
5%
aggr3
HDD
5.06 ms/op
3.06 IOPS
< 1 MB/s
6%
5%
meg_aggr2
HDD
10.4 ms/op
52.9 IOPS
7.28 MB/s
3%
2%
Aggregate / Performance : aggr4
Switch to Health View
Last updated: Nov 15, 2019, 1:20 PM
Summary
Explorer
Information
Compare the performance of associated objects and display detailed charts
TIME RANGE
Last 72 Hours
VIEW AND COMPARE
Aggregates on same Node
Filter
Aggregate
Late...
IOP...
MB/...
Perf...
aggr3
5.06 ...
3.06 ...
< 1 M...
6 %
aggr_evt
0.29 ...
3.79 ...
< 1 M...
< 1 %
Add
aggr_automation
0.27...
6.35 ...
< 1 M...
< 1 %
Add
Comparing
1 Additional Object
CHOOSE CHARTS
7 Charts Selected
aggr4
aggr3
Events for Aggregate: aggr4
No data to display
Latency
VIEW
Total
Zoom View
Latency – Total view (ms/op)
aggr4
7.15
aggr3
6.74
Nov 14, 2019, 6:39 AM
12:00 AM
Nov 15, 2019, 12:00 AM
Latency for Aggregate: aggr4
Last updated: Nov 15, 2019, 1:23 PM
Event Timeline: aggr4
TIME RANGE
Last 72 Hours
Critical Events
Error Events
Warning Events
Information Events
No data to display
VIEW
Total
Comparing Objects
aggr4
aggr3
25
ms/op
20
15
10
5
0
Latency – Total view (ms/op)
aggr4
4.25
aggr3
3.97
Nov 14, 2019, 1:24 AM
Nov 13, 2019, 12:00 AM
Nov 13, 2019, 12:00 PM
Nov 14, 2019, 12:00 AM
Nov 14, 2019, 12:00 PM
Nov 15, 2019, 12:00 AM
Nov 15, 2019,

Investigate performance events in Active IQ Unified Manager

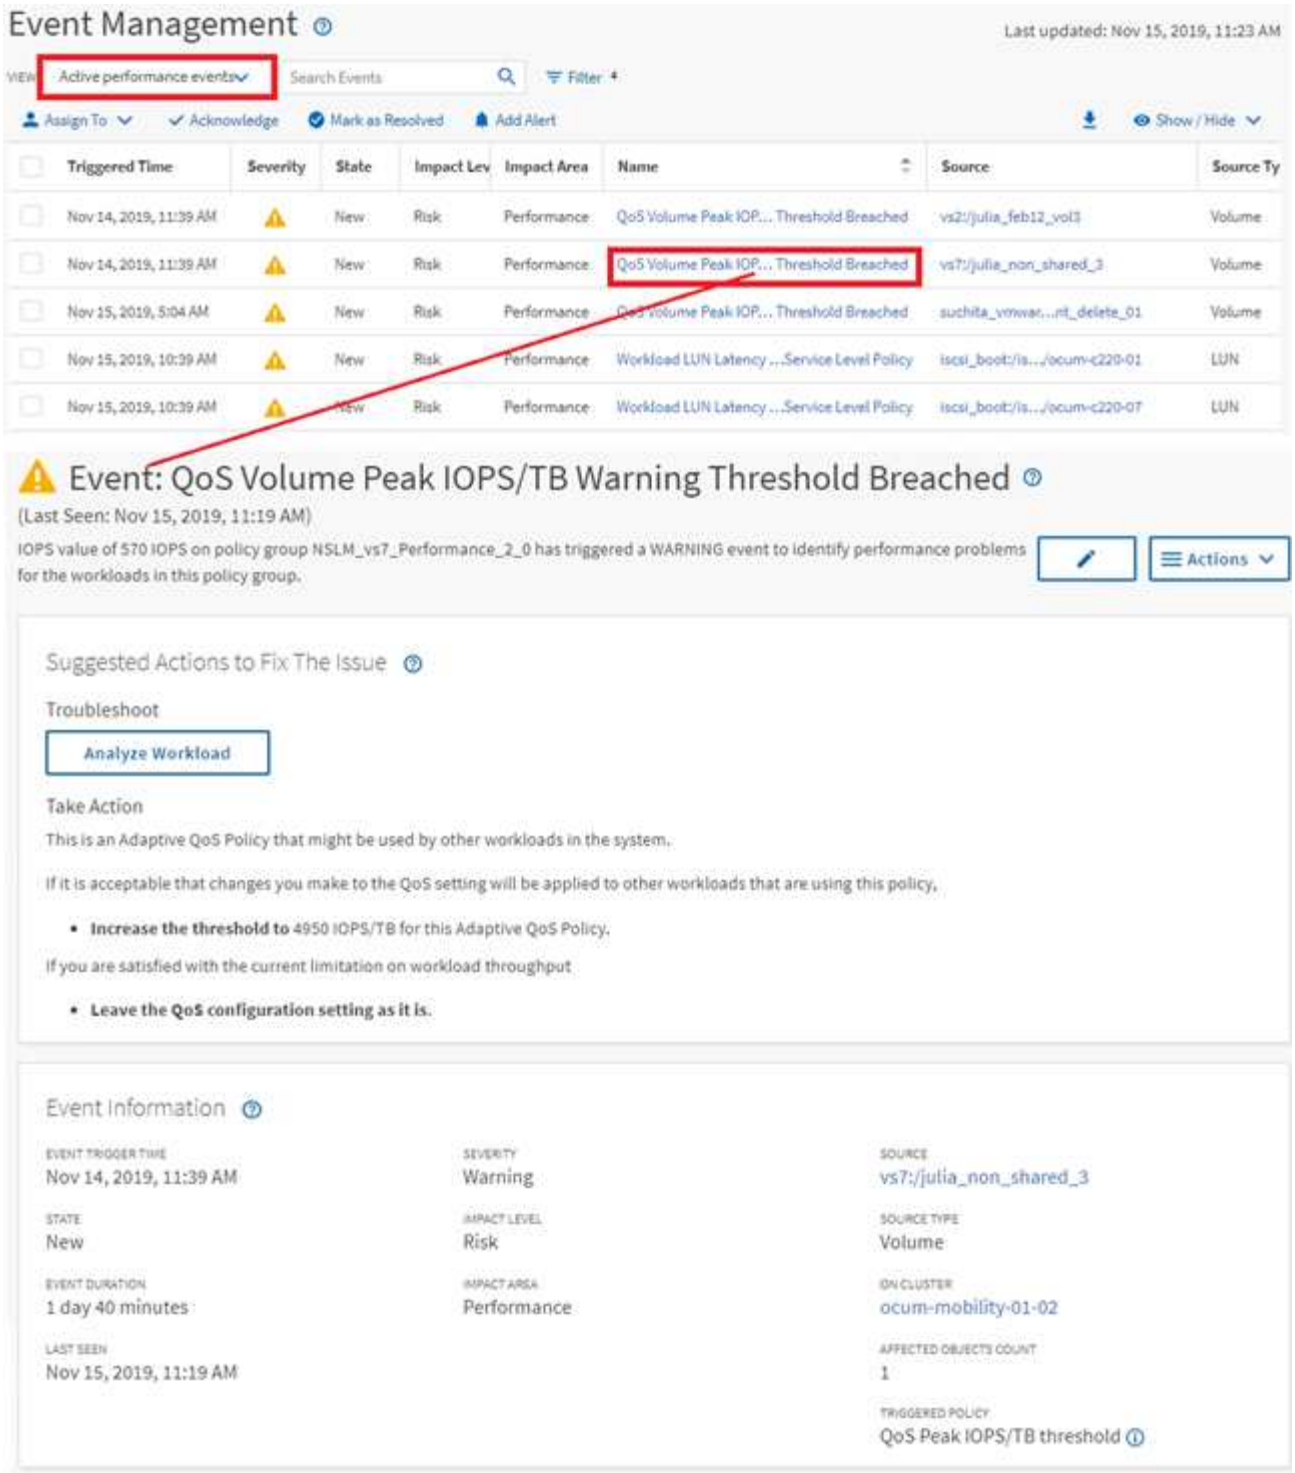
The Unified Manager event detail pages provide you with an in-depth look at any performance event. This is beneficial when investigating performance events, when troubleshooting, and when fine-tuning system performance.

Depending on the type of performance event, you might see one of two types of event detail pages:

- Event details page for user-defined and system-defined threshold policy events
- Event details page for dynamic threshold policy events

This is one example of an event investigation navigation.

1. In the left navigation pane, click **Event Management**.
2. From the View menu, click **Active performance events**.
3. Click the name of the event that you want to investigate and the Event details page is displayed.
4. View the Description of the event and review the Suggested Actions (where available) to view more details about the event that may help you resolve the issue. You can click the **Analyze Workload** button to display detailed performance charts to help further analyze the issue.



Search for storage objects in Active IQ Unified Manager

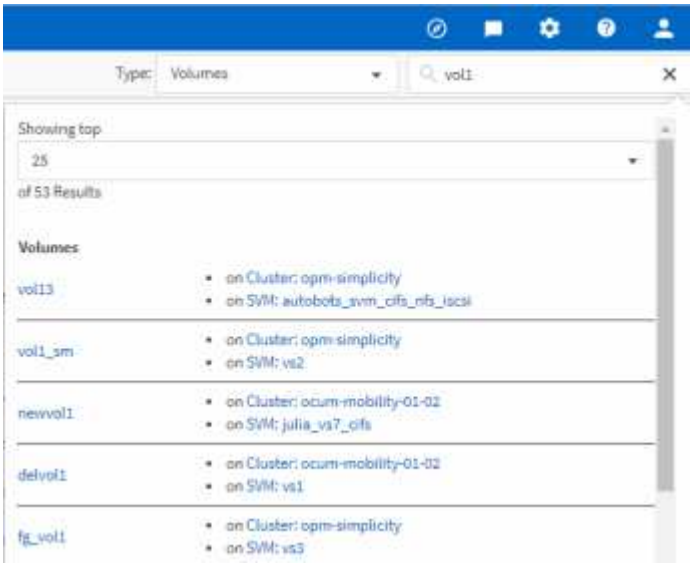
To quickly access a specific object, you can use the **Search all Storage Objects** field at the top of the menu bar. This method of global search across all objects enables you to quickly locate specific objects by type. Search results are sorted by storage object type and you can filter them using the drop-down menu. A valid search must contain at least three characters.

The global search displays the total number of results, but only the top 25 search results are accessible.

Because of this, the global search functionality can be thought of as a shortcut tool for finding specific items if you know the items you want to quickly locate. For complete search results, you can use the search in the object inventory pages and its associated filtering functionality.

You can click the drop-down box and select **All** to simultaneously search across all objects and events. Alternatively, you can click the drop-down box to specify the object type. Type a minimum of three characters of the object or event name into the **Search all Storage Objects** field, and then press **Enter** to display the search results, such as:

- Clusters: cluster names
- Nodes: node names
- Aggregates: aggregate names
- SVMs: SVM names
- Volumes: volume names
- LUNs: LUN paths



LIFs and ports are not searchable in the global search bar.

In this example, the drop-down box has the Volume object type selected. Typing “vol” into the **Search all Storage Objects** field displays a list of all volumes whose names contain these characters. For object searches, you can click any search result to navigate to that object’s Performance Explorer page. For event searches, clicking an item in the search result navigates to the Event Details page.

Filter inventory page content in Active IQ Unified Manager

You can filter inventory page data in Unified Manager to quickly locate data based on specific criteria. You can use filtering to narrow the contents of the Unified Manager pages to show only the results in which you are interested. This provides a very efficient method of displaying only the data in which you are interested.

Use **Filtering** to customize the grid view based on your preferences. Available filter options are based on the object type being viewed in the grid. If filters are currently applied, the number of applied filters displays at the right of the Filter button.

Three types of filter parameters are supported.

Parameter	Validation
String (text)	The operators are **contains**, **starts with**, **ends with**, and **does not contain**.
Number	The operators are **greater than**, **less than**, **in the last**, and **between**.
Enum (text)	The operators are **is** and **is not**.

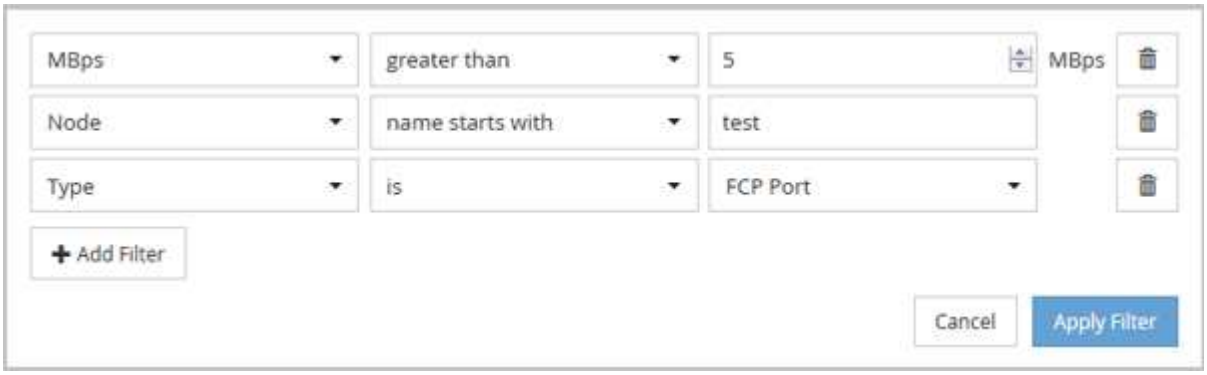
The Column, Operator, and Value fields are required for each filter; the available filters reflect the filterable columns on the current page. The maximum number of filters you can apply is four. Filtered results are based on combined filter parameters. Filtered results apply to all pages in your filtered search, not just the page currently displayed.

You can add filters using the Filtering panel.

1. At the top of the page, click the **Filter** button. The Filtering panel displays.
2. Click the left drop-down list and select an object; for example, *Cluster*, or a performance counter.
3. Click the center drop-down list, and select the operator you want to use.
4. In the last list, select or enter a value to complete the filter for that object.
5. To add another filter, click **+Add Filter**. An additional filter field displays. Complete this filter using the process described in the preceding steps. Note that upon adding your fourth filter, the **+Add Filter** button no longer displays.
6. Click **Apply Filter**. The filter options are applied to the grid and the number of filters is displayed to the right of the Filter button.
7. Use the Filtering panel to remove individual filters by clicking the trash icon at the right of the filter to be removed.
8. To remove all filters, click **Reset** at the bottom of the filtering panel.

Filtering example

The illustration shows the Filtering panel with three filters. The **+Add Filter** button displays when you have fewer than the maximum of four filters.



After clicking **Apply Filter**, the Filtering panel closes, applies your filters, and shows the number of filters applied (3).

Monitor cluster performance from the Dashboard

Monitor cluster performance from the Dashboard in Active IQ Unified Manager

The Unified Manager Dashboard provides a few panels that display high-level performance status of all clusters being monitored by this instance of Unified Manager. It enables you to assess the overall performance of the managed clusters, and to quickly note, locate, or assign for resolution any specific events identified.

Review performance panels on the Active IQ Unified Manager Dashboard

The Unified Manager Dashboard provides a few panels that display high-level performance status for all the clusters that are being monitored in your environment. You can choose to view status for all clusters or for an individual cluster.

In addition to showing performance information, most of the panels also display the number of active events in that category, and the number of new events added over the previous 24 hours. This information helps you decide which clusters you may need to analyze further to resolve reported events. Clicking on the events displays the top few events and provides a link to the Event Management inventory page filtered to show the events in that category.

The following panels provide performance status.

- **Performance Capacity panel**

  When viewing all clusters, this panel displays the performance capacity value for each cluster (averaged over the previous 1 hour) and the number of days until performance capacity reaches the upper limit (based on daily growth rate). Clicking the bar chart takes you to the Nodes inventory page for that cluster. Note that the Nodes inventory page displays the performance capacity averaged over the previous 72 hours, so this value may not match the Dashboard value.

  When viewing a single cluster, this panel displays the cluster performance capacity, total IOPS, and total throughput values.

- **Workload IOPS panel**

  When workload active management is enabled, and when viewing a single cluster, this panel displays the total number workloads that are currently running in a certain range of IOPS.

- **Workload Performance panel**

  When workload active management is enabled, this panel displays the total number of conforming and non-conforming workloads that are assigned to each defined Performance Service Level. Clicking a bar chart takes you to the workloads assigned to that policy in the Workloads page.

- **Usage Overview panel**

  When viewing all clusters, you can choose to view clusters sorted by highest IOPS or throughput (MB/s).

  When viewing a single cluster, you can choose to view workloads on that cluster sorted by highest IOPS or throughput (MB/s).

Interpret performance banner messages in Active IQ Unified Manager

Unified Manager may display banner messages on the Notifications page (from the Notification bell) to alert you to status issues for a particular cluster.

Banner message	Description	Resolution
`No performance data is being collected from cluster cluster_name. Restart Unified Manager to correct this issue.`	The Unified Manager collection service has stopped and no performance data is being collected from any clusters.	Restart Unified Manager to correct this issue. If this does not correct the issue, contact technical support.
`More than x hour(s) of historical data is being collected from cluster cluster_name. Current data collections will start after all historical data is collected.`	A data continuity collection cycle is currently running to retrieve performance data outside of the real-time cluster performance collection cycle.	No action is required. Current performance data will be collected after the data continuity collection cycle is completed. A data continuity collection cycle runs when a new cluster is added or when Unified Manager has been unable to collect current performance data for some reason.

Change the performance statistics collection interval in Active IQ Unified Manager

The default collection interval for performance statistics is 5 minutes. You can change this interval to 10 or 15 minutes if you find that collections from large clusters are not finishing within the default time. This setting affects the collection of statistics from all clusters that this instance of Unified Manager is monitoring.

Before you begin

You must have a user ID and password authorized to log in to the maintenance console of the Unified Manager server.

The issue of performance statistics collections not finishing on time is indicated by the banner messages `Unable to consistently collect from cluster <cluster_name>` or `Data collection is taking too long on cluster <cluster_name>`.

You should change the collection interval only when required because of a statistics collections issue. Do not change this setting for any other reason.

> Changing this value from the default setting of 5 minutes can affect the number and frequency of performance events that Unified Manager reports. For example, system-defined performance thresholds trigger events when the policy is exceeded for 30 minutes. When using 5-minute collections, the policy must be exceeded for six consecutive collections. For 15-minute collections the policy must be exceeded for only two collection periods.

A message at the bottom of the Cluster Setup page indicates the current statistical data collection interval.

Steps

1. Log in using SSH as the maintenance user to the Unified Manager host.

The Unified Manager maintenance console prompts are displayed.

2. Type the number of the menu option labeled **Performance Polling Interval Configuration**, and then press Enter.
3. If prompted, enter the maintenance user password again.
4. Type the number for the new polling interval that you want to set, and then press Enter.

If you changed the Unified Manager collection interval to 10 or 15 minutes, and you have a current connection to an external data provider (such as Graphite), you must change the data provider transmit interval so that it is equal to, or greater, than the Unified Manager collection interval.

Troubleshoot workloads using the workload analyzer

Troubleshoot workloads with the workload analyzer in Active IQ Unified Manager

The workload analyzer provides a way to view important health and performance criteria for a single workload on a single page to assist in troubleshooting. By viewing all current and past events for a workload you can get a better idea why the workload may be having a performance or capacity issue now.

Using this tool can also help you determine if storage is the cause of any performance issues for an application or if the issue is caused by a networking or other related issue.

You can initiate this functionality from a variety of places in the user interface:

- From the Workload Analysis selection on the left-navigation menu
- From the Event details page by clicking the **Analyze Workload** button
- From any workload inventory page (volume, LUN, workload, NFS share, or SMB/CIFS share), by clicking the more icon , then **Analyze Workload**
- From the Virtual Machines page by clicking the **Analyze Workload** button from any Datastore object

When you launch the tool from the left-navigation menu, you can enter the name of any workload that you want to analyze and select the time range for which you want to troubleshoot. When you launch the tool from any of the workload or virtual machine inventory pages, the name of the workload is filled in automatically, and the workload's data is presented with the default 2 hour time range. When you launch the tool from the Event details page, the name of the workload is filled in automatically, and the data of 10 days is displayed.

Understand workload analyzer data in Active IQ Unified Manager

The workload analyzer page displays information about any current events that could be affecting the workload, recommendations to potentially fix the issue causing the event, and charts for analyzing performance and capacity history.

At the top of the page you specify the name of the workload (volume or LUN) that you want to analyze and the timeframe over which you want to see statistics. You can change the timeframe at any point if you want to view a shorter or longer period of time.

The other areas of the page display the analysis results and the performance and capacity charts.

Workload charts for LUNs do not provide the same level of statistics as those charts for volumes, so you will notice differences when analyzing these two types of workloads.

- **Events summary area**

  Displays a brief overview of the number and types of events that have occurred over the timeframe. When there are events from different impact areas (for example, performance and capacity), this information is displayed so you can select details for the type of event you are interested in. Click the event type to view a list of the event names.

  If there is only one event during the timeframe, then a list of recommendations to fix the issue is listed for some events.

- **Event Timeline**

  Shows all occurrences of events during the specified timeframe. Hover your cursor over each event to view the event name.

  If you arrived at this page by clicking the **Analyze Workload** button from the Event details page, the icon for the selected event appears larger so that you can identify the event.

- **Performance charts area**

  Displays charts for latency, throughput (both IOPS and MB/s), and utilization (for both the node and aggregate) based on the timeframe you selected. You can click the View performance details link to display the Performance Explorer page for the workload in case you want to perform further analysis.

  - **Latency** displays the latency for the workload over the selected timeframe. The chart has three views that enable you to see:
    - **Total** latency
    - **Breakdown** latency (broken out by reads, writes, and other processes)
    - **Cluster Components** latency (broken out by cluster component)

See Cluster components and why they can be in contention for a description of the cluster components that are displayed here.
Throughput displays both IOPS and MB/s throughput for the workload over the selected timeframe. The chart has four views that enable you to see:
* **Total** throughput
*** Breakdown throughput (broken out by reads, writes, and other processes)**
* **Cloud Throughput** (the MB/s being used to write data to and read data from the cloud; for those workloads that are tiering capacity to the cloud)
*** IOPS with Forecast (a prediction of what the upper and lower IOPS throughput values were expected to be over the timeframe)**
This chart also displays Quality of Service (QoS) maximum and minimum throughput threshold settings, if configured, so you can see where the system may be limiting the throughput intentionally with QoS policies.
Utilization displays utilization for both the aggregate and node on which the workload is running over the selected timeframe. From here you can see if your aggregate or node are overutilized, possibly causing high latency. When analyzing FlexGroup volumes there are multiple nodes and multiple aggregates listed on the utilization charts.

- **Capacity chart area**

Displays charts for data capacity and Snapshot capacity for the past one month for the workload.

For volumes, you can click the View capacity details link to display the Health Details page for the workload in case you want to perform further analysis. LUNs do not provide this link because there is no Health Details page for LUNs.

- **Capacity View** displays the total available space allocated for the workload and the logical used space (after all NetApp optimizations).
- **Snapshot View** displays the total space reserved for Snapshot copies, and the amount of space currently being used. Note that LUNs do not provide a Snapshot View.
- **Cloud Tier View** displays how much capacity is being used in the local performance tier and how much is being used in the cloud tier.
  These charts include an estimate of the amount of time remaining before the capacity is full for this workload. This information is based on historical usage and requires a minimum of 10 days of data. When less than 30 days of capacity remain, Unified Manager identifies the storage as “almost full”.

Understand when to use the workload analyzer in Active IQ Unified Manager

You would typically use the workload analyzer to troubleshoot a latency issue reported by a user, to more thoroughly analyze a reported event or alert, or to explore a workload that you see is operating abnormally.

In the case where users have contacted you to say that the application they are using is running very slowly, you can check the latency, throughput, and utilization charts for the workload over which the application is running to see if storage is the cause of the performance issue. You can use the capacity chart as well to see if capacity is low because an ONTAP system in which capacity is over 85% used can cause performance issues. These charts will help you determine if the issue is caused by storage or by a networking or other related issue.

In the case where Unified Manager has generated a performance event and you want to review the cause of the issue more thoroughly, you can launch the workload analyzer from the Event details page by clicking the **Analyze Workload** button to research some of the latency, throughput, and capacity trends for the workload.

In the case where you notice a workload that appears to be operating abnormally when viewing any workload inventory page (volume, LUN, workload, NFS share, or SMB/CIFS share), you can click the more icon , then **Analyze Workload** to open the Workload Analysis page to examine the workload further.

Use the workload analyzer in Active IQ Unified Manager

There are many ways to start the workload analyzer from the user interface. Here we describe launching the tool from the left-navigation pane.

Steps

1. In the left navigation pane, click **Workload Analysis**.

   The Workload Analysis page is displayed.

2. If you know the workload name, enter the name. If you are not sure of the full name, enter a minimum of 3 characters and the system displays a list of workloads that match the string.
3. Select the time range if you want to view statistics for longer than the default 2 hours and click **Apply**.

4. View the Summary area to see the events that have occurred during the timeframe.
5. View the performance and capacity charts to see when any of the metrics are abnormal and see if any events align with the abnormal entry.

Monitor cluster performance from the Performance Cluster Landing page

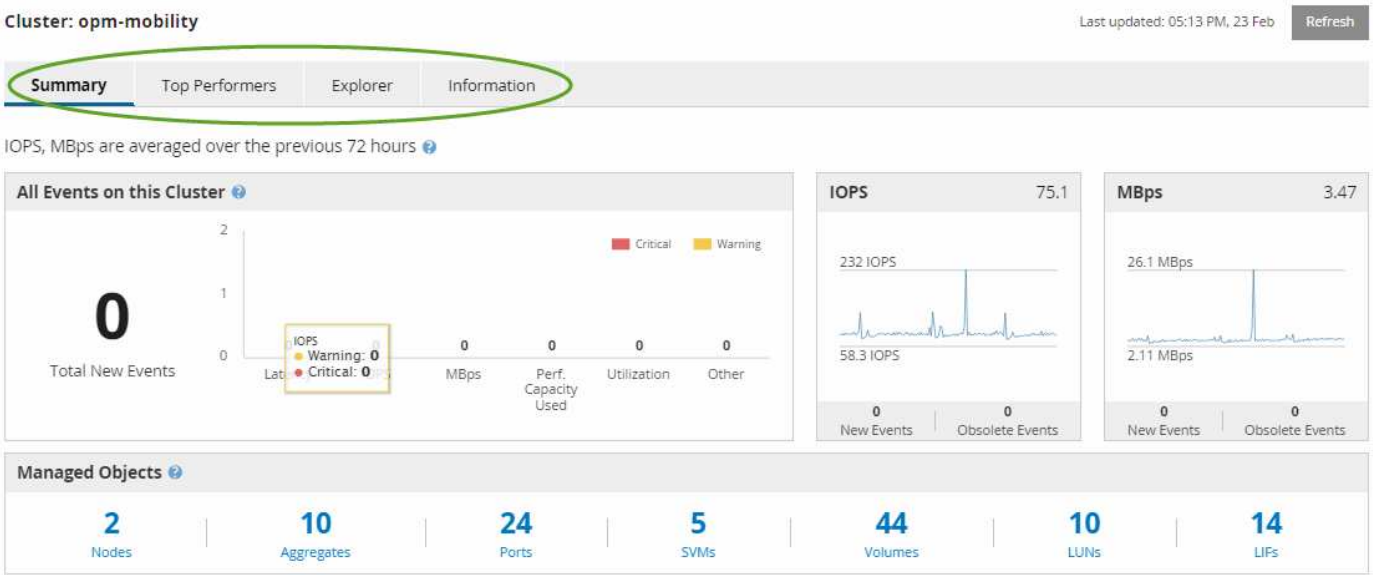
Monitor cluster performance from the Performance Cluster Landing page in Active IQ Unified Manager

The Performance Cluster Landing page displays the high-level performance status of a selected cluster that is being monitored by an instance of Unified Manager. This page enables you to assess the overall performance of a specific cluster, and to quickly note, locate, or assign for resolution any cluster-specific events that are identified.

Understand the Performance Cluster Landing page in Active IQ Unified Manager

The Performance Cluster Landing page provides a high-level performance overview of a selected cluster, with an emphasis on the performance status of the top 10 objects within the cluster. Performance issues are displayed at the top of the page, in the All Events on this Cluster panel.

The Performance Cluster Landing page provides a high-level overview of each cluster that is managed by an instance of Unified Manager. This page provides you with information about events and performance, and enables you to monitor and troubleshoot clusters. The following image shows an example of the Performance Cluster Landing page for the cluster called opm-mobility:



The event count on the Cluster Summary page may not match the event count on the Performance Event Inventory page. This is because the Cluster Summary page can show one event each in the Latency and Utilization bars when combination threshold policies have been breached, whereas the Performance Event Inventory page shows only one event when a combination policy has been breached.

If a cluster was removed from being managed by Unified Manager, the status **Removed** is displayed at the right of the cluster name at the top of the page.

Performance Cluster Landing page

Review the Performance Cluster Landing page in Active IQ Unified Manager

The Performance Cluster Landing page displays the high-level performance status of a selected cluster. The page enables you to access complete details of each performance counter for the storage objects on the selected cluster.

The Performance Cluster Landing page includes four tabs that separate the cluster details into four areas of information:

- Summary page
  - Cluster Events pane
  - MB/s and IOPS performance charts
  - Managed Objects pane
- Top Performers page
- Explorer page
- Information page

Performance Cluster Summary page

Review the Performance Cluster Summary page in Active IQ Unified Manager

The Performance Cluster Summary page provides a summary of the active events, IOPS performance, and MB/s performance for a cluster. This page also includes the total count of the storage objects in the cluster.

Review the Cluster performance events pane in Active IQ Unified Manager

The Cluster performance events pane displays performance statistics and all active events for the cluster. This is most helpful when monitoring your clusters and all cluster-related performance and events.

All Events on this Cluster pane

The All Events on this Cluster pane displays all active cluster performance events for the preceding 72 hours. The Total Active Events is displayed at the far left; this number represents the total of all New and Acknowledged events for all storage objects in this cluster. You can click the Total Active Events link to navigate to the Events Inventory page, which is filtered to display these events.

The Total Active Events bar graph for the cluster displays the total number of active critical and warning events:

- Latency (total for nodes, aggregates, SVMs, volumes, LUNs, and namespaces)
- IOPS (total for clusters, nodes, aggregates, SVMs, volumes, LUNs, and namespaces)
- MB/s (total for clusters, nodes, aggregates, SVMs, volumes, LUNs, namespaces, ports, and LIFs)

- Performance Capacity Used (total for nodes and aggregates)
- Utilization (total for nodes, aggregates, and ports)
- Other (cache miss ratio for volumes)

The list contains active performance events triggered from user-defined threshold policies, system-defined threshold policies, and dynamic thresholds.

Graph data (vertical counter bars) is displayed in red () for critical events, and yellow () for warning events. Position your cursor over each vertical counter bar to view the actual type and number of events. You can click **Refresh** to update the counter panel data.

You can show or hide critical and warning events in the Total Active Events performance graph by clicking the **Critical** and **Warning** icons in the legend. If you hide certain event types, the legend icons are displayed in gray.

Counter panels

The counter panels display cluster activity and performance events for the preceding 72 hours, and includes the following counters:

- **IOPS counter panel**

  IOPS indicates the operating speed of the cluster in number of input/output operations per second. This counter panel provides a high-level overview of the cluster's IOPS health for the preceding 72-hour period. You can position your cursor over the graph trend line to view the IOPS value for a specific time.

- **MB/s counter panel**

  MB/s indicates how much data has been transferred to and from the cluster in megabytes per second. This counter panel provides a high-level overview of the cluster's MB/s health for the preceding 72-hour period. You can position your cursor over the graph trend line to view the MB/s value for a specific time.

The number at the top right of the chart in the gray bar is the average value from the last 72-hour period. Numbers shown at the bottom and top of the trend line graph are the minimum and maximum values for the last 72-hour period. The gray bar below the chart contains the count of active (new and acknowledged) events and obsolete events from the last 72-hour period.

The counter panels contain two types of events:

- **Active**

  Indicates that the performance event is currently active (new or acknowledged). The issue causing the event has not corrected itself or has not been resolved. The performance counter for the storage object remains above the performance threshold.

- **Obsolete**

  Indicates that the event is no longer active. The issue causing the event has corrected itself or has been resolved. The performance counter for the storage object is no longer above the performance threshold.

For **Active Events**, if there is one event, you can position your cursor over the event icon and click the event number to link to the appropriate Event Details page. If there is more than one event, you can click **View all Events** to display the Events Inventory page, which is filtered to show all events for the selected object counter type.

Review the Managed Objects pane in Active IQ Unified Manager

The Managed Objects pane in the Performance Summary tab provides a top-level overview of the storage object types and counts for the cluster. This pane enables you to track the status of the objects in each cluster.

The managed objects count is point-in-time data as of the last collection period. New objects are discovered at 15-minute intervals.

Clicking the linked number for any object type displays the object performance inventory page for that object type. The object inventory page is filtered to show only the objects on this cluster.

The managed objects are:

- **Nodes**

  A physical system in a cluster.

- **Aggregates**

  A set of multiple redundant array of independent disks (RAID) groups that can be managed as a single unit for protection and provisioning.

- **Ports**

  A physical connection point on nodes that is used to connect to other devices on a network.

- **Storage VMs**

  A virtual machine providing network access through unique network addresses. An SVM might serve data out of a distinct namespace, and is separately administrable from the rest of the cluster.

- **Volumes**

  A logical entity holding accessible user data through one or more of the supported access protocols. The count includes both FlexVol volumes and FlexGroup volumes; it does not include FlexGroup constituents.

- **LUNs**

  The identifier of a Fibre Channel (FC) logical unit or an iSCSI logical unit. A logical unit typically corresponds to a storage volume, and is represented within a computer operating system as a device.

- **Network Interfaces**

  A logical network interface representing a network access point to a node. The count includes all interface types.

View the Top Performers page in Active IQ Unified Manager

The Top Performers page displays the storage objects that have the highest performance or the lowest performance, based on the performance counter you select. For example, in the Storage VMs category, you can display the SVMs that have the highest IOPS, or the highest latency, or the lowest MB/s. This page also shows if any of the top performers have any active performance events (New or Acknowledged).

The Top Performers page displays a maximum of 10 of each object. Note that the Volume object includes both FlexVol volumes and FlexGroup volumes.

- **Time Range**

  You can select a time range for viewing the top performers; the selected time range applies to all storage objects. Available time ranges:

  - Last Hour
  - Last 24 Hours
  - Last 72 Hours (default)
  - Last 7 Days

- **Metric**

  Click the **Metric** menu to select a different counter. Counter options are unique to the object type. For example, available counters for the **Volumes** object are **Latency**, **IOPS**, and **MB/s**. Changing the counter reloads the panel data with the top performers based on the selected counter.

  Available counters:

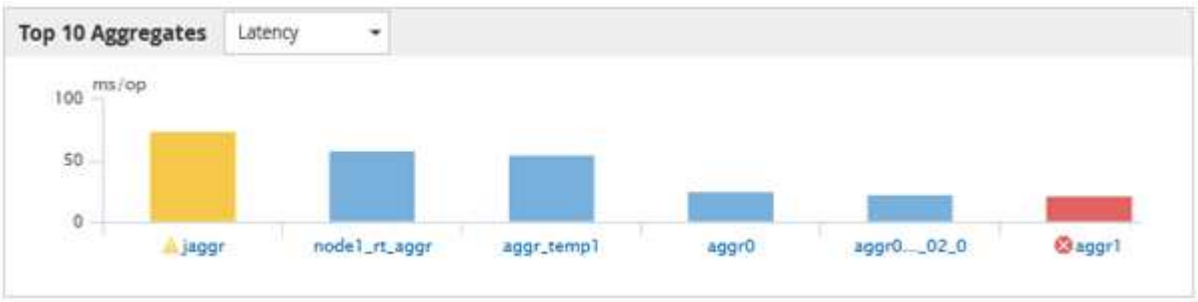
  - Latency
  - IOPS
  - MB/s
  - Performance Capacity Used (for nodes and aggregates)
  - Utilization (for nodes and aggregates)

- **Sort**

  Click the **Sort** menu to select an ascending or descending sort for the selected object and counter. The options are **Highest to lowest** and **Lowest to highest**. These options enable you to view the objects with the highest performance or the lowest performance.

- **Counter bar**

  The counter bar in the graph shows the performance statistics for each object, represented as a bar for that item. The bar graphs are color-coded. If the counter is not breaching a performance threshold, the counter bar is displayed in blue. If a threshold breach is active (a new or acknowledged event), the bar is displayed in the color for the event: warning events are displayed in yellow ( ), and critical events are displayed in red ( ). Threshold breaches are further indicated by severity event indicator icons for warning and critical events.



For each graph, the X axis displays the top performers for the selected object type. The Y axis displays

units applicable to the selected counter. Clicking the object name link below each vertical bar graph element navigates to the Performance Landing page for the selected object.

- **Severity Event indicator**

  The **Severity Event** indicator icon is displayed at the left of an object name for active critical ( ) or warning ( ) events in the top performers graphs. Click the **Severity Event** indicator icon to view:

  - **One event**

    Navigates to the Event details page for that event.

  - **Two or more events**

    Navigates to the Event inventory page, which is filtered to display all events for the selected object.

- **Export button**

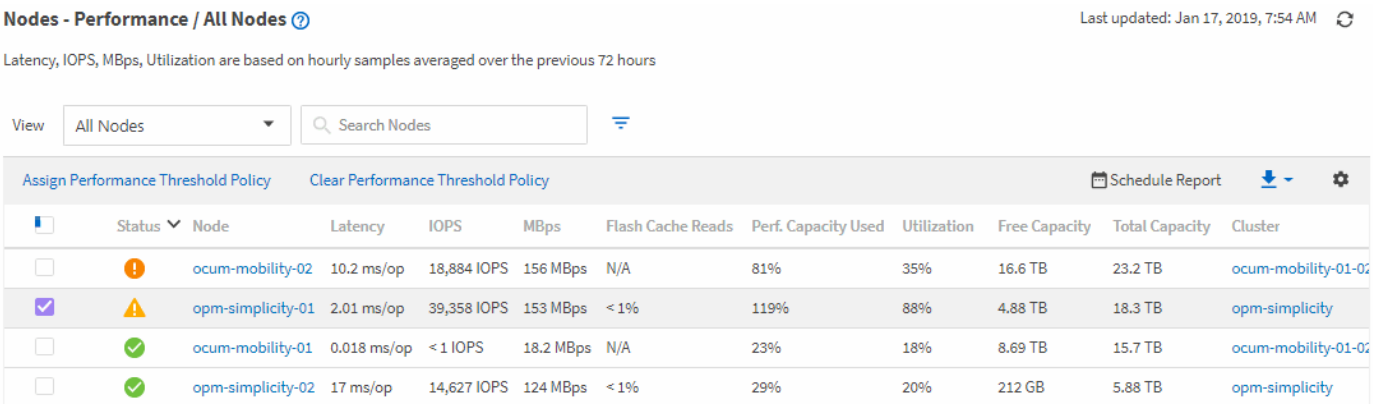
  Creates a `.csv` file that contains the data that appears in the counter bar. You can choose to create the file for the single cluster you are viewing or for all clusters in the data center.

Monitor performance using the Performance Inventory pages

Monitor object performance using Performance Inventory pages in Active IQ Unified Manager

The object inventory performance pages display performance information, performance events, and object health for all objects within an object type category. This provides you with an at-a-glance overview of the performance status of each object within a cluster, for example, for all nodes or all volumes.

Object inventory performance pages provide a high-level overview of object status, enabling you to assess the overall performance of all objects and compare object performance data. You can refine the content of object inventory pages by searching, sorting, and filtering. This is beneficial when monitoring and managing object performance, because it enables you to quickly locate objects with performance issues and to begin the troubleshooting process.



By default, objects on the performance inventory pages are sorted based on object performance criticality. Objects with new critical performance events are listed first, and objects with warning events are listed second.

This provides an immediate visual indication of issues that must be addressed. All performance data is based on a 72-hour average.

You can easily navigate from the object inventory performance page to an object details page by clicking the object name in the object name column. For example, on the Performance/All Nodes inventory page, you would click a node object in the **Nodes** column. The object details page provides in-depth information and detail about the selected object, including side-by-side comparison of active events.

View Performance inventory pages for all storage objects in Active IQ Unified Manager

You use the Performance inventory pages to see a summary of performance information about each of the available storage objects, such as clusters, aggregates, volumes, and so on. You can link to the Performance object detail pages to view detailed information for a particular object.

By default, objects in the view pages are sorted based on event criticality. Objects with critical events are listed first, and objects with warning events are listed second. This provides an immediate visual indication of issues that must be addressed.

You can export data from these pages to a comma-separated values (`.csv`) file, Microsoft Excel file (`.xlsx`), or (`.pdf`) document by using the **Reports** button, and then use the exported data to build reports. Additionally, you can customize the page and then schedule a report to be created and emailed on a regular basis by using the **Scheduled Reports** button.

All the fields on these pages can be used in custom views and in reports. Some of the fields are linked to related pages enabling a more detailed view.

Performance: All Clusters view

The Performance: All Clusters view displays an overview of the performance events, data, and configuration information for each cluster that is monitored by an instance of Unified Manager. This page enables you to monitor the performance of your clusters, and to troubleshoot performance issues and threshold events.

You can assign performance threshold policies to, or clear threshold policies from, any object on the object inventory pages using the **Assign Performance Threshold Policy** and **Clear Performance Threshold Policy** buttons.

The following are some important fields in the Performance: All Clusters view.

- Cluster FQDN: The fully qualified domain name (FQDN) of the cluster.
- IOPS: The input/output operations per second on the cluster.
- MB/s: The throughput on the cluster, measured in MiB per second.
- Capacity fields: Free and total capacity in GiB.
- Host Name or IP Address: The host name or IP address (IPv4 or IPv6) of the cluster management LIF.
- OS Version: The version of ONTAP software that is installed on the cluster.

> If different versions of ONTAP software are installed on the nodes in the cluster, the lowest version number is listed. You can view the ONTAP version that is installed on each node from the Performance: All Nodes view.

- Threshold Policy: The user-defined performance threshold policy, or policies, that are active on this storage object. You can position your cursor over policy names containing an ellipsis (…) to view the full policy name or the list of assigned policy names. The Assign Performance Threshold Policy and Clear Performance Threshold Policy buttons remain disabled until you select one or more objects by clicking the check boxes located at the far left.

Performance: All Volumes view

The Performance: All Volumes view displays an overview of the performance events, counter data, and configuration information for each FlexVol volume and FlexGroup volume that is being monitored by an instance of Unified Manager. This enables you to quickly monitor the performance of your volumes, and to troubleshoot performance issues and threshold events.

If you want to analyze the latency and throughput of a specific object, click the more options button ⋮ then **Analyze Workload** and you can view performance and capacity charts on the Workload Analysis page. You can view the details on System Manager, given you have valid credentials for System Manager.

ⓘ For data protection (DP) volumes, only counter values for user-generated traffic are displayed. Root volumes are not displayed on this page.

The following are some important fields in the Performance: All Volumes view.

- Style: Either FlexVol or FlexGroup.
- Latency: For FlexVol volumes, this is the average response time of the volume for all I/O requests, expressed in milliseconds per operation. For FlexGroup volumes, this is the average latency of all constituent volumes.
- IOPS/TB: The number of input/output operations processed per second based on the total space that is being consumed by the workload, in terabytes. This counter measures how much performance can be delivered by a given amount of storage capacity.
- IOPS: For FlexVol volumes, this is the number of input/output operations per second for the volume. For FlexGroup volumes, this is the sum of IOPS for all constituent volumes.
- MB/s: For FlexVol volumes, this is the throughput on the volume, measured in megabytes per second. For FlexGroup volumes, this is the sum of MB/s for all constituent volumes.
- Capacity fields: Free and total capacity in GiB.

See the following links for more information:

- Assigning performance threshold policies to storage objects
- Removing performance threshold policies from storage objects
- Types of workloads monitored by Unified Manager
- Viewing the QoS policy group settings applied to specific volumes or LUNs
- Understanding the Unified Manager recommendations to tier data to the cloud
- Viewing performance charts to compare volumes or LUNs that are in the same QoS policy group

Performance: All Aggregates view

The Performance: All Aggregates view displays an overview of the performance events, data, and configuration information for each aggregate that is monitored by an instance of Unified Manager. This page enables you to monitor the performance of your aggregates, and to troubleshoot performance issues and

threshold events.

The following are some important fields in the Performance: All Aggregates view.

- Type: The type of aggregate:
  - HDD
  - Hybrid. Combines HDDs and SSDs, but Flash Pool has not been enabled.
  - Hybrid (Flash Pool). Combines HDDs and SSDs, and Flash Pool has been enabled.
  - SSD
  - SSD (FabricPool). Combines SSDs and a cloud tier
  - HDD (FabricPool). Combines HDDs and a cloud tier
  - VMDisk (SDS). Virtual disks within a virtual machine
  - VMDisk (FabricPool). Combines virtual disks and a cloud tier
  - LUN
- Inactive Data Reporting: Whether the inactive data reporting capability is enabled or disabled on this aggregate. When enabled, volumes on this aggregate display the amount of cold data in the Performance: All Volumes view. The value in this field is "N/A" when the version of ONTAP does not support inactive data reporting.
- Threshold Policy: The user-defined performance threshold policy, or policies, that are active on this storage object. You can position your cursor over policy names containing an ellipsis (…) to view the full policy name or the list of assigned policy names. The Assign Performance Threshold Policy and Clear Performance Threshold Policy buttons remain disabled until you select one or more objects by clicking the check boxes located at the far left.
  See the following links for more information:
- Assigning performance threshold policies to storage objects
- Removing performance threshold policies from storage objects

Performance: All Nodes view

The Performance: All Nodes view displays an overview of the performance events, data, and configuration information for each node that is being monitored by an instance of Unified Manager. This enables you to quickly monitor the performance of your nodes, and to troubleshoot performance issues and threshold events.

Flash Cache Reads returns the percentage of read operations on the node that are satisfied by cache, instead of being returned from the disk. Flash Cache data is displayed only for nodes, and only when a Flash Cache
module is installed in the node.

In the **Reports** menu, the **Hardware Inventory Report** option is provided when Unified Manager, and the clusters it is managing, are installed in a site with no external network connectivity. This button generates a .csv file that contains a complete list of cluster and node information; such as hardware model numbers and serial numbers, disk types and counts, installed licenses, and more. This reporting functionality is helpful for contract renewal within secure sites that are not connected to the NetApp Active IQ platform.
You can assign performance threshold policies to, or clear threshold policies from, any object on the object inventory pages using the **Assign Performance Threshold Policy** and **Clear Performance Threshold Policy** buttons.

See the following links for more information:

- Assigning performance threshold policies to storage objects
- Removing performance threshold policies from storage objects
- Generating a hardware inventory report for contract renewal

Performance: All Storage VMs view

The Performance: All Storage VMs view displays an overview of the performance events, data, and configuration information for each storage virtual machine (SVM) that is being monitored by an instance of Unified Manager. This enables you to quickly monitor the performance of your SVMs, and to troubleshoot performance issues and threshold events. The Latency field on this page reports the average response time for all I/O requests, expressed in milliseconds per operation.

The SVMs that are listed on this page include only Data and Cluster SVMs. Unified Manager does not use or display Admin or Node SVMs.

See the following links for more information:

- Assigning performance threshold policies to storage objects
- Removing performance threshold policies from storage objects

Performance: All LUNs view

The Performance: All LUNs view displays an overview of the performance events, data, and configuration information for each LUN that is being monitored by an instance of Unified Manager. This enables you to quickly monitor the performance of your LUNs, and to troubleshoot performance issues and threshold events.

If you want to analyze the latency and throughput of a specific object, click the more icon ⋮, then **Analyze Workload** and you can view performance and capacity charts on the **Workload Analysis** page.

See the following links for more information:

- Monitoring LUNs in a Consistency Group relationship
- Provisioning LUNs
- Assigning performance threshold policies to storage objects
- Removing performance threshold policies from storage objects
- Viewing volumes or LUNs that are in the same QoS policy group.
- Viewing the QoS policy group settings applied to specific volumes or LUNs
- Provisioning LUNs by using APIs

Performance: All NVMe Namespaces view

The Performance: All NVMe Namespaces view displays an overview of the performance events, data, and configuration information for each NVMe Namespace that is being monitored by an instance of Unified Manager. This enables you to quickly monitor the performance and health of your Namespaces, and to troubleshoot issues and threshold events.

The following information, among others is reported:
The current state of the Namespace.
* Offline - Read or write access to the Namespace is not allowed.
* Online - Read and write access to the Namespace is allowed.

* NVFail - The Namespace was automatically taken offline due to an NVRAM failure.
* Space Error - The Namespace has run out of space.

See the following links for more information:

- Assigning performance threshold policies to storage objects
- Removing performance threshold policies from storage objects

Performance: All Network Interfaces view

The Performance: All Network Interfaces view displays an overview of the performance events, data, and configuration information for each network interface (LIF) that is being monitored by this instance of Unified Manager. This page enables you to quickly monitor the performance of your interfaces, and to troubleshoot performance issues and threshold events.
The following are some important fields in the Performance: All Network Interfaces view.

- IOPS: The input/output operations per second. IOPS is not applicable to NFS LIFs and CIFS LIFs, and is displayed as N/A for these types.
- Latency: The average response time for all I/O requests, expressed in milliseconds per operation. Latency is not applicable to NFS LIFs and CIFS LIFs, and is displayed as N/A for these types.
- Home Location: The home location for the interface, displayed as node name and port name, separated by a colon (:). If the location is displayed with an ellipsis (…), you can position your cursor over the location name to view the full location.
- Current Location: The current location for the interface, displayed as node name and port name, separated by a colon (:). If the location is displayed with an ellipsis (…), you can position your cursor over the location name to view the full location.
- Role: The interface role: Data, Cluster, Node Management, or Intercluster.

The interfaces that are listed on this page include Data LIFs, Cluster LIFs, Node Management LIFs, and intercluster LIFs. Unified Manager does not use or display System LIFs.

Performance: All Ports view

The Performance: All Ports view displays an overview of the performance events, data, and configuration information for each port that is being monitored by an instance of Unified Manager. This enables you to quickly monitor the performance of your ports, and to troubleshoot performance issues and threshold events.
For a port role, the network port function is displayed, either Data or Cluster. FCP ports cannot have a role, and the role
is displayed as N/A.

Performance counter values are displayed for physical ports only. Counter values are not displayed for VLANs or interface groups.

See the following links for more information:

- Assigning performance threshold policies to storage objects
- Removing performance threshold policies from storage objects

Performance: QoS Policy Groups view

The QoS Policy Groups view displays the QoS policy groups available on the clusters that Unified Manager is monitoring. This includes traditional QoS policies, adaptive QoS policies, and QoS policies assigned by using Performance Service Levels.

The following are some important fields in the Performance: QoS Policy Groups view.

- QoS Policy Group: The name of the QoS policy group.
  For NetApp Service Level Manager (NSLM) 1.3 policies that have been imported into Unified Manager 9.7 or greater, the name displayed here includes the SVM name and other information that is not in the name when the Performance Service Levelwas defined in NSLM. For example, the name "NSLM_vs6_Performance_2_0" means this is the NSLM system-defined "Performance" PSL policy created on SVM "vs6" with an expected latency of "2 ms/op".
- SVM: The storage VM (SVM) to which the QoS policy group belongs. You can click the storage VM name to navigate to that storage VM's details page. Note that this field is blank if the QoS policy has been created on the Admin storage VM as this storage VM type represents the cluster.
- Min Throughput: The minimum throughput, in IOPS, that the policy group will be guaranteed to provide. For adaptive policies this is the minimum expected IOPS per TB allocated to the volume or LUN, based on the storage object allocated size.
- Max Throughput: The throughput, in IOPS and/or MB/s, that the policy group must not exceed. When this field is blank it means the max throughout defined in ONTAP is infinite. For adaptive policies this is the maximum (peak) possible IOPS per TB allocated to the volume or LUN, based on either the storage object allocated size or the storage object used size.
- Absolute Minimum IOPS: For adaptive policies this is the absolute minimum IOPS value that is used as an override
  when the expected IOPS is less than this value.
- Block Size: The block size specified for the QoS adaptive policy.
- Min Allocation: Whether "allocated space" or "used space" is used to determine the maximum throughput (peak) IOPS.
- Expected Latency: The expected average latency for storage input/output operations.
- Shared: For traditional QoS policies, whether the throughput values defined in the policy group are shared among multiple objects.
- Associated Objects: The number of workloads that are assigned to the QoS policy group. You can click the expand button ( ) next to the QoS Policy Group Name to view more details about the policy group.
- Allocated Capacity: The amount of space that the objects that are in the QoS policy group are currently using.
- Associated Objects: The number of workloads that are assigned to the QoS policy group, separated into volumes
  and LUNs. You can click the number to navigate to a page that provides more details about the selected volumes or LUNs.

For more information see the topics under Managing performance using QoS policy group information.

Refine Performance inventory page contents

Refine performance inventory page content in Active IQ Unified Manager

The inventory pages for performance objects contain tools to help you refine object

inventory data content, enabling you to locate specific data quickly and easily.

Information contained within the Performance object inventory pages can be extensive, often spanning multiple pages. This kind of comprehensive data is excellent for monitoring, tracking, and improving performance; however, locating specific data requires tools to enable you to quickly locate the data for which you are looking. Therefore, the Performance object inventory pages contain functionality for searching, sorting, and filtering. Additionally, searching and filtering can work together to further narrow your results.

Search on Object Inventory Performance pages in Active IQ Unified Manager

You can search strings on Object Inventory Performance pages. Use the **Search** field located at the top right of the page to quickly locate data based on either object name or policy name. This enables you to quickly locate specific objects and their associated data, or to quickly locate policies and view associated policy object data.

Step

1. Perform one of the following options, based on your search requirements:

To locate this…	Type this…
A specific object	The object name into the **Search** field, and click **Search**. The object for which you searched and its related data is displayed.
A user-defined performance threshold policy	All or part of the policy name into the **Search** field, and click **Search**. The objects assigned to the policy for which you searched are displayed.

Sort data on the Object Inventory Performance pages in Active IQ Unified Manager

You can sort all data on Object Inventory Performance pages by any column in ascending or descending order. This enables you to quickly locate object inventory data, which is helpful when examining performance or beginning a troubleshooting process.

The selected column for sorting is indicated by a highlighted column heading name and an arrow icon indicating the sorting direction at the right of the name. An up arrow indicates ascending order; a down arrow indicates descending order. The default sort order is by **Status** (event criticality) in descending order, with the most critical performance events listed first.

Step

1. You can click a column name to toggle the sort order of the column in ascending or descending order.

   The Object Inventory Performance page contents are sorted in ascending or descending order, based on the selected column.

Filter data on Object Inventory Performance pages in Active IQ Unified Manager

You can filter data in the Object Inventory Performance pages to quickly locate data based on specific criteria. You can use filtering to narrow the contents of the Object Inventory Performance pages to show only the results you have specified. This provides a very efficient method of displaying only the performance data in which you are interested.

You can use the Filtering panel to customize the grid view based on your preferences. Available filter options are based on the object type being viewed in the grid. If filters are currently applied, the number of applied filters displays at the right of the Filter button.

Three types of filter parameters are supported.

Parameter	Validation
String (text)	The operators are **contains**, **starts with**, **ends with**, and **does not contain**.
Number	The operators are **greater than**, **less than**, **in the last**, and **between**.
Enum (text)	The operators are **is** and **is not**.

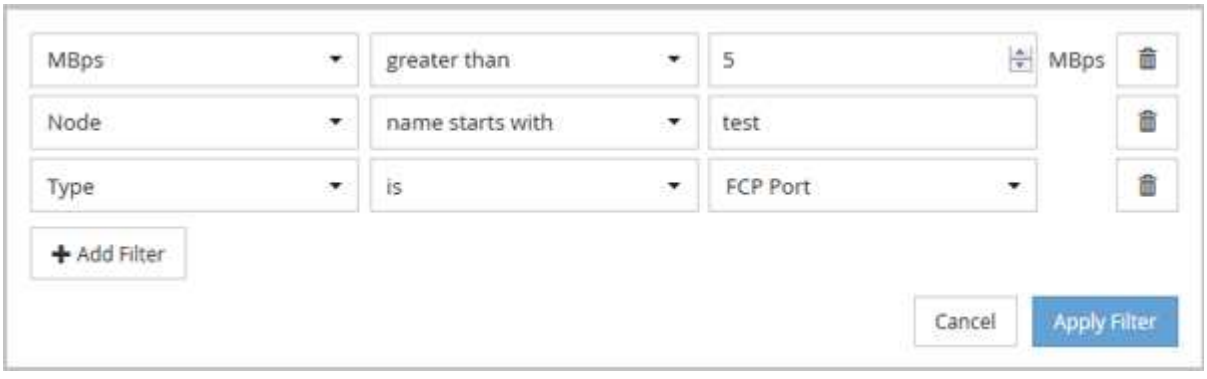
The Column, Operator, and Value fields are required for each filter; the available filters reflect the filterable columns on the current page. The maximum number of filters you can apply is four. Filtered results are based on combined filter parameters. Filtered results apply to all pages in your filtered search, not just the page currently displayed.

You can add filters using the Filtering panel.

1. At the top of the page, click the **Filter** button. The Filtering panel displays.
2. Click the left drop-down list and select an object; for example, *Cluster*, or a performance counter.
3. Click the center drop-down list, and select the operator you want to use.
4. In the last list, select or enter a value to complete the filter for that object.
5. To add another filter, click **+Add Filter**. An additional filter field displays. Complete this filter using the process described in the preceding steps. Note that upon adding your fourth filter, the **+Add Filter** button no longer displays.
6. Click **Apply Filter**. The filter options are applied to the grid and the number of filters is displayed to the right of the Filter button.
7. Use the Filtering panel to remove individual filters by clicking the trash icon at the right of the filter to be removed.
8. To remove all filters, click **Reset** at the bottom of the filtering panel.

Filtering example

The illustration shows the Filtering panel with three filters. The **+Add Filter** button displays when you have fewer than the maximum of four filters.



After clicking **Apply Filter**, the Filtering panel closes, applies your filters, and shows the number of filters applied ().

Review Active IQ Unified Manager recommendations for tiering data to the cloud

The Performance: All Volumes view displays information related to the size of the user data stored on the volume that is inactive (cold). In some cases, Unified Manager identifies certain volumes that would benefit by tiering the inactive data to the cloud tier (cloud provider or StorageGRID) of a FabricPool-enabled aggregate.

> FabricPool was introduced in ONTAP 9.2, so if you are using a version of ONTAP software prior to 9.2, the Unified Manager recommendation to tier data requires upgrading your ONTAP software. Additionally, the **`auto`** tiering policy was introduced in ONTAP 9.4, and the **`all`** tiering policy was introduced in ONTAP 9.6, so if the recommendation is to use the auto tiering policy, you must upgrade to ONTAP 9.4 or greater.

The following three fields on Performance: All Volumes view provide information about whether you can improve your storage system's disk utilization and save space on the performance tier by moving inactive data to the cloud tier.

- **Tiering Policy**

  The tiering policy determines whether the data on the volume remains on the performance tier or whether some of the data is moved from the performance tier to the cloud tier.

  The value in this field indicates the tiering policy set on the volume, even if the volume does not currently reside on a FabricPool aggregate. The tiering policy takes effect only when the volume is on a FabricPool aggregate.

- **Cold Data**

  The cold data displays the size of the user data stored on the volume that is inactive (cold).

  A value is displayed here only when using ONTAP 9.4 or greater software because it requires that the aggregate on which the volume is deployed has the **`inactive data reporting parameter`** set to **`enabled`**, and that the minimum number of cooling days threshold has been met (for volumes that use the **`snapshot-only`** or **`auto`** tiering policy). Otherwise the value is listed as "N/A".

- **Cloud Recommendation**

  After enough information has been captured about the data activity on the volume, Unified Manager may determine there is no action required, or that you could save space on the performance tier by tiering inactive data to the cloud tier.

> The Cold Data field is updated every 15 minutes, but the Cloud Recommendation field is updated every 7 days when the cold data analysis is performed on the volume. Therefore, the exact amount of cold data may differ between the fields. The Cloud Recommendation field displays the date when the analysis was run.

When Inactive Data Reporting is enabled, the Cold Data field displays the exact amount of inactive data. Without the inactive data reporting capability Unified Manager uses performance statistics to determine if data is inactive on a volume. The amount of inactive data is not displayed in the Cold Data field in this case, but it is

displayed when you hover your cursor over the word **Tier** to view the cloud recommendation.

The cloud recommendations you will see are:

- **Learning**. Not enough data has been collected to make a recommendation.
- **Tier**. Analysis has determined that the volume contains inactive (cold) data and that you should configure the volume to move that data to the cloud tier. In some cases this may require that you move the volume to a FabricPool-enabled aggregate first. In other cases where the volume is already on a FabricPool aggregate, you just have to change the tiering policy.
- **No Action**. Either the volume has very little inactive data, the volume is already set to the “auto” tiering policy on a FabricPool aggregate, or the volume is a data protection volume. This value is also displayed when the volume is offline or when it is being used in a MetroCluster configuration.

To move a volume, or to change the volume tiering policy or the aggregate inactive data reporting settings, use ONTAP System Manager, the ONTAP CLI commands, or a combination of these tools.

If you are logged in to Unified Manager with the Application Administrator or Storage Administrator role, the **Configure Volume** link is available in the cloud recommendation when you hover your cursor over the word **Tier**. Click this button to open the Volumes page in System Manager to make the recommended change.

Monitor performance using the Performance Explorer pages

Monitor performance using the Performance Explorer pages in Active IQ Unified Manager

The Performance Explorer pages display detailed information about the performance of each object in a cluster. The page provides a detailed view into the performance of all cluster objects, enabling you to select and compare the performance data of specific objects across various time periods.

You can also assess the overall performance of all objects, and compare object performance data in a side-by-side format.

Use root objects for performance monitoring in Active IQ Unified Manager

The root object is the baseline against which other object comparisons are made. This enables you to view and compare the data from other objects to the root object, providing performance data analysis that helps you to troubleshoot and improve object performance.

The root object name displays at the top of the Comparing pane. Additional objects display below the root object. Although there is no limit to the number of additional objects you can add to the Comparing pane, only one root object is allowed. Data for the root object automatically displays in the graphs in the Counter Charts pane.

You cannot change the root object; it is always set to the object page you are viewing. For example, if you open the Volume Performance Explorer page of Volume1, then Volume1 is the root object and cannot be changed. If you want to compare against a different root object, then you must click the link for an object and open its landing page.

Events and Thresholds are displayed only for root objects.

Filter correlated objects in the Active IQ Unified Manager grid

Filtering enables you to display a smaller, more well-defined subset of objects in the grid. For example, if you have 25 volumes in the grid, filtering enables you to view only those volumes that have throughput less than 90 MBps, or latency greater than 1 ms/op.

Specify a time range for correlated objects

Specify a time range for correlated objects in Active IQ Unified Manager

The Time Range selector on the Performance Explorer page enables you to specify the time range for object data comparison. Specifying a time range refines the contents of the Performance Explorer pages to show only the object data within the time range you have specified.

Refining the time range provides an efficient method of displaying only the performance data in which you are interested. You can select a predefined time range or specify a custom time range. The default time range is the preceding 72 hours.

Select a predefined time range in Active IQ Unified Manager

Selecting a predefined time range is a quick and efficient way for you to customize and focus data output when viewing cluster object performance data. When selecting a predefined time range, data for up to 13 months is available.

Steps

1. At the top right of the **Performance Explorer** page, click **Time Range**.
2. From the right side of the **Time Range Selection** panel, select a predefined time range.
3. Click **Apply Range**.

Specify a custom time range in Active IQ Unified Manager

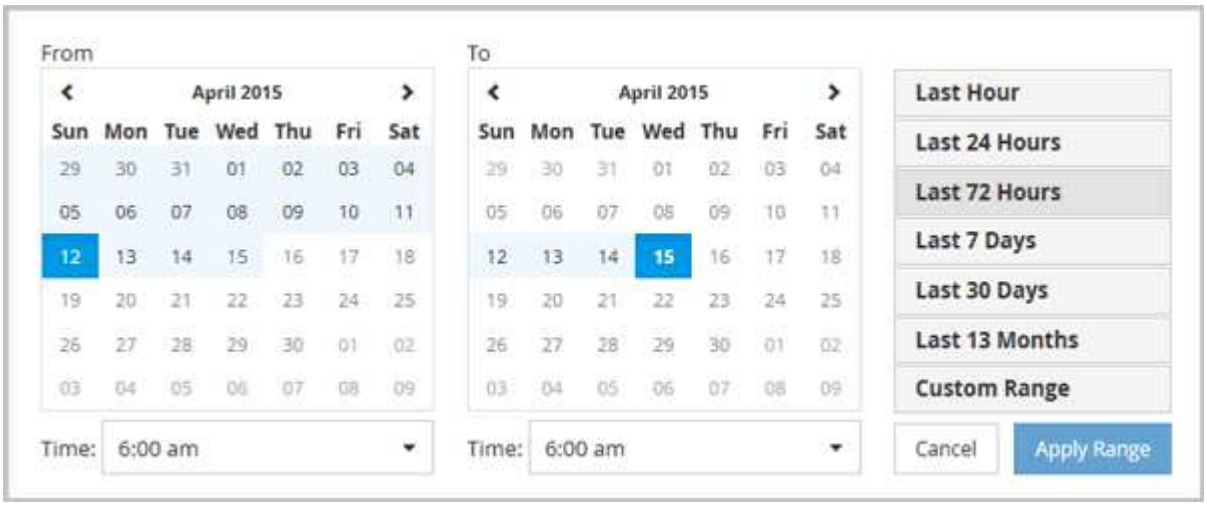
The Performance Explorer page enables you to specify the date and time range for your performance data. Specifying a custom time range provides greater flexibility than using predefined time ranges when refining cluster object data.

You can select a time range between one hour and 390 days. 13 months equals 390 days because each month is counted as 30 days. Specifying a date and time range provides more detail and enables you to zoom in on specific performance events or series of events. Specifying a time range also aids in troubleshooting potential performance issues, as specifying a date and time range displays data surrounding the performance event in finer detail. Use the **Time Range** control to select predefined date and time ranges, or specify your own custom date and time range of up to 390 days. Buttons for predefined time ranges vary from the **Last Hour** through the **Last 13 Months**.

Selecting the **Last 13 Months** option or specifying a custom date range greater than 30 days displays a dialog box alerting you that performance data displayed for a period greater than 30 days is charted using hourly averages and not 5-minute data polling. Therefore, a loss of timeline visual granularity might occur. If you click

the **Do not show again** option in the dialog box, the message does not appear when you select the **Last 13 Months** option or specify a custom date range greater than 30 days. Summary data also applies on a smaller time range, if the time range includes a time/date that is more than 30 days from today.

When selecting a time range (either custom or predefined), time ranges of 30 days or fewer are based on 5-minute interval data samples. Time ranges greater than 30 days are based on one-hour interval data samples.



1. Click the **Time Range** drop-down box and the Time Range panel displays.
2. To select a predefined time range, click one of the **Last…** buttons at the right of the **Time Range** panel. When selecting a predefined time range, data for up to 13 months is available. The predefined time range button you selected is highlighted, and the corresponding days and time display in the calendars and time selectors.
3. To select a custom date range, click the start date in the **From** calendar on the left. Click **<** or **>** to navigate forward or backward in the calendar. To specify the end date, click a date in the **To** calendar on the right. Note that the default end date is today unless you specify a different end date. The **Custom Range** button at the right of the Time Range panel is highlighted, indicating that you have selected a custom date range.
4. To select a custom time range, click the **Time** control below the **From** calendar and select the start time. To specify the end time, click the **Time** control below the **To** calendar on the right and select the end time. The **Custom Range** button at the right of the Time Range panel is highlighted, indicating that you have selected a custom time range.
5. Optionally, you can specify the start and end times when selecting a predefined date range. Select the predefined date range as previously described, then select the start and end times as previously described. The selected dates are highlighted in the calendars, your specified start and end times display in the **Time** controls, and the **Custom Range** button is highlighted.
6. After selecting the date and time range, click **Apply Range**. The performance statistics for that time range display in the charts and in the Events timeline.

Define correlated objects for comparison graphing in Active IQ Unified Manager

You can define a list of correlated objects for data and performance comparison in the Counter Chart pane. For example, if your storage virtual machine (SVM) is experiencing a performance issue, you can compare all volumes in the SVM to identify which volume might be causing the issue.

You can add any object in the correlated objects grid to the Comparing and Counter Chart panes. This enables you to view and compare data of multiple objects and with the root object. You can add and remove objects to

and from the correlated objects grid; however, the root object in the Comparing pane is not removable.

Adding many objects to the Comparing pane may have a negative impact on performance. To maintain performance, you should select a limited number of charts for data comparison.

Steps

1. In the objects grid, locate the object that you want to add, and click the **Add** button.

   The **Add** button turns gray, and the object is added to the additional objects list in the Comparing pane. The object's data is added to the graphs in the Counter Charts panes. The color of the object's eye icon ( ) matches the color of the object's data trend line in the graphs.

2. **Optional:** Hide or show data for selected objects:

To do this…	Take this action…
Hide a selected object	Click the selected object's eye icon ( ) in the Comparing pane. The object's data is hidden, and the eye icon for that object turns gray.
Show a hidden object	Click the gray eye icon of the selected object in the Comparing pane. The eye icon returns to its original color, and the object data is added back into the graphs in the Counter Charts pane.

3. **Optional:** Remove selected objects from the **Comparing** pane:

To do this…	Take this action…
Remove a selected object	Hover over the selected object's name in the Comparing pane to show the remove object button (**X**), and then click the button. The object is removed from the Comparing pane, and its data is cleared from the counter charts.
Remove all selected objects	Click the remove all object's button (**X**) at the top of the Comparing pane. All selected objects and their data are removed, leaving only the root object.

Understand counter charts in Active IQ Unified Manager

Charts in the Counter Charts pane enable you to view and compare performance data for the root object and for objects you have added from the correlated objects grid. This can help you understand performance trends and isolate and resolve performance issues.

Counter charts displayed by default are Events, Latency, IOPS, and MBps. Optional charts that you can choose to display are Utilization, Performance Capacity Used, Available IOPS, IOPS/TB, and Cache Miss Ratio. Additionally, you can choose to view total values or breakdown values for the Latency, IOPS, MBps, and Performance Capacity Used charts.

The Performance Explorer displays certain counter charts by default; whether the storage object supports them all or not. When a counter is not supported, the counter chart is empty and the message `Not applicable`

`for <object>` is displayed.

The charts display performance trends for the root object and for all objects you have selected in the Comparing pane. Data in each chart is arranged as follows:

- **X axis**

  Displays the specified time period. If you have not specified a time range, the default is the preceding 72-hour period.

- **Y axis**

  Displays counter units unique to the selected object, or objects.

Trend line colors match the color of the object name as displayed in the Comparing pane. You can position your cursor over a point on any trend line to view details for time and value for that point.

If you want to investigate a specific period of time within a chart, you can use one of the following methods:

- Use the **<** button to expand the Counter Charts pane to span the width of the page.
- Use the cursor (when it transitions to a magnifying glass) to select a portion of the timeframe in the chart to focus and enlarge that area. You can click Reset Chart Zoom to return the chart to the default timeframe.
- Use the **Zoom View** button to display a large single counter chart that contains expanded details and threshold indicators.

Occasionally, gaps in the trend lines display. Gaps mean that either Unified Manager failed to collect performance data from the storage system or that Unified Manager might have been down.

Understand the types of performance counter charts in Active IQ Unified Manager

There are standard performance charts that display the counter values for the selected storage object. Each of the Breakdown counter charts display the total values separated out into read, write, and other categories. Furthermore, some Breakdown counter charts display additional detail when the chart is displayed in Zoom view.

The following table shows the available performance counter charts.

Available charts	Chart description
Events	Displays critical, error, warning, and information events in correlation with the statistical charts for the root object. Health events display in addition to performance events to provide a complete picture of the reasons performance may be affected.
Latency - Total	Number of milliseconds required to respond to application requests. Note that the average latency values are I/O weighted.

Available charts	Chart description
Latency - Breakdown	The same information shown in Latency Total, but with the performance data separated into read, write, and other latency. This chart option applies only when the selected object is an SVM, node, aggregate, volume, LUN, or namespace.
Latency - Cluster Components	The same information shown in Latency Total, but with the performance data separated into latency by cluster component. This chart option applies only when the selected object is a volume.
IOPS - Total	Number of input/output operations processed per second. When displayed for a node, selecting “Total” displays the IOPS for data moving through this node that may reside on the local or the remote node, and selecting “Total (Local)” displays the IOPS for data that resides only on the current node.
IOPS - Breakdown	The same information shown in IOPS Total, but with the performance data separated into read, write, and other IOPS. This chart option applies only when the selected object is an SVM, node, aggregate, volume, LUN, or namespace. When displayed in Zoom view, the volumes chart displays QoS minimum and maximum throughput values, if configured in ONTAP. When displayed for a node, selecting “Breakdown” displays the IOPS breakdown for data moving through this node that may reside on the local or the remote node, and selecting “Breakdown (Local)” displays the IOPS breakdown for data that resides only on the current node.
IOPS - Protocols	The same information shown in IOPS Total, but the performance data is separated into individual charts for CIFS, NFS, FCP, NVMe, and iSCSI protocol traffic. This chart option applies only when the selected object is an SVM.
IOPS/TB - Total	Number of input/output operations processed per second based on the total space that is being consumed by the workload, in terabytes. Also called I/O density, this counter measures how much performance can be delivered by a given amount of storage capacity. When displayed in Zoom view the volumes chart displays QoS expected and peak throughput values, if configured in ONTAP. This chart option applies only when the selected object is a volume.
MB/s - Total	Number of megabytes of data transferred to and from the object per second.

Available charts	Chart description
MB/s - Breakdown	The same information shown in the MB/s chart, but with the throughput data separated into disk reads, Flash Cache reads, writes, and other. When displayed in Zoom view, the volumes chart displays QoS maximum throughput values, if configured in ONTAP. This chart option applies only when the selected object is an SVM, node, aggregate, volume, LUN, or namespace. Flash Cache data is displayed only for nodes, and only when a Flash Cache module is installed in the node.
Performance Capacity Used - Total	Percentage of performance capacity that is being consumed by the node or aggregate.
Performance Capacity Used - Breakdown	Performance Capacity Used data separated into user protocols and system background processes. Additionally, the amount of free performance capacity is shown.
Available IOPS - Total	Number of input/output operations per second that are currently available (free) on this object. This number is the result of subtracting the currently used IOPS from the total IOPS that Unified Manager calculates that the object can perform. This chart option applies only when the selected object is a node or aggregate.
Utilization - Total	Available resource percentage of the object that is being used. Utilization indicates node utilization for nodes, disk utilization for aggregates, and bandwidth utilization for ports. This chart option applies only when the selected object is a node, aggregate, or port.
Cache Miss Ratio - Total	Percentage of read requests from client applications that are returned from the disk instead of being returned from the cache. This chart option applies only when the selected object is a volume.

Select performance charts to display in Active IQ Unified Manager

The Choose charts drop-down list enables you to select the types of performance counter charts to display in the Counter Charts pane. This enables you to view specific data and counters, based on your performance requirements.

Steps

1. In the **Counter Charts** pane, click the **Choose charts** drop-down list.
2. Add or remove charts:

To…	Do this…
Add or remove individual charts	Click the check boxes next to the charts you want to show or hide
Add all charts	Click **Select All**
Remove all charts	Click **Unselect All**

Your chart selections are displayed in the Counter Charts pane. Note that as you add charts, the new charts are inserted into the Counter Charts pane to match the order of the charts listed in the Choose charts drop-down list. Selecting additional charts might require additional scrolling.

Expand the Counter Charts pane in Active IQ Unified Manager

You can expand the Counter Charts pane so that the charts are larger and more readable.

After you have defined the comparison objects and the time range for counters, you can view a larger Counter Charts pane. You use the **<** button in the middle of the Performance Explorer window to expand the pane.

Step

1. Expand or reduce the **Counter Charts** pane.

To…	Do this…
Expand the Counter Charts pane to fit the width of the page	Click the **<** button
Reduce the Counter Charts pane to the right half of the page	Click the **>** button

Change the Counter Charts focus to a shorter period in Active IQ Unified Manager

You can use your mouse to reduce the time range to focus on a specific period of time in the Counter Chart pane or in the Counter Charts Zoom View window. This enables you to see a more granular and microscopic view of any part of the timeline of performance data, events, and thresholds.

Before you begin

The cursor must have changed to a magnifying glass to indicate that this functionality is active.

When using this feature, which alters the timeline to display values that correspond to the more granular display, the time and date range on the **Time Range** selector does not change from the original values for the chart.

Steps

1. To zoom into a specific period of time, click using the magnifying glass and drag the mouse to highlight the area that you want to see in detail.

   The counter values for the time period you select fills the counter chart.

2. To return to the original period of time as set in the **Time Range** selector, click the **Reset Chart Zoom**

button.

The counter chart displays in its original state.

View event details in the Active IQ Unified Manager Events Timeline

You can view all events and their related details in the Events Timeline pane of Performance Explorer. This is a quick and efficient method of viewing all the health and performance events that occurred on the root object during a specified time range, which can be helpful when troubleshooting performance issues.

The Events Timeline pane shows critical, error, warning, and informational events that occurred on the root object during the selected time range. Each event severity has its own timeline. Single and multiple events are represented by an event dot on the timeline. You can position your cursor over an event dot to see the event details. To increase the visual granularity of multiple events, you can decrease the time range. This spreads out multiple events into single events, enabling you to separately view and investigate each event.

Each performance event dot on the Events Timeline lines up vertically with a corresponding spike in the counter charts trend lines that are displayed below the Events Timeline. This provides a direct visual correlation between events and overall performance. Health events are displayed on the timeline as well, but these types of events do not necessarily line up with a spike in one of the performance charts.

Steps

1. On the **Events Timeline** pane, position the cursor over an event dot on a timeline to view a summary of the event or events at that event point.

   A pop-up dialog displays information about the event types, the date and time when the events occurred, the state, and the event duration.

2. View full event details for one event or multiple events:

To do this…	Click this…
View details for a single event	**View Event Detail** in the pop-up dialog.
View details for multiple events	**View Event Details** in the pop-up dialog. Clicking a single event on the multiple events dialog displays the appropriate Event Details page.

Counter Charts Zoom View

Learn about the zoom view in Counter Charts in Active IQ Unified Manager

The Counter Charts provide a Zoom View that enables you to zoom in on performance details over your specified time period. This enables you to see performance details and events with much higher granularity, which is beneficial when troubleshooting performance issues.

When displayed in Zoom View, some of the breakdown charts provide additional information than what appears when the chart is not in Zoom View. For example, the IOPS, IOPS/TB, and MBps Breakdown chart

Zoom View pages display QoS policy values for volumes and LUNs if they have been set in ONTAP.

> For system-defined performance threshold policies, only the “Node resources over-utilized” and “QoS throughput limit breached” policies are available from the **Policies** list. The other system-defined threshold policies are not available at this time.

Display the Counter Charts zoom view in Active IQ Unified Manager

The Counter Charts Zoom View provides a finer level of detail for the selected counter chart and its associated timeline. This magnifies the counter chart data, enabling you to have a sharper view into performance events and their underlying causes.

You can display the Counter Charts Zoom View for any counter chart.

Steps

1. Click **Zoom View** to open the selected chart a new browser window.
2. If you are viewing a Breakdown chart and then click **Zoom View** the Breakdown chart is shown in Zoom View. You can select **Total** while in Zoom View if you want to change the view option.

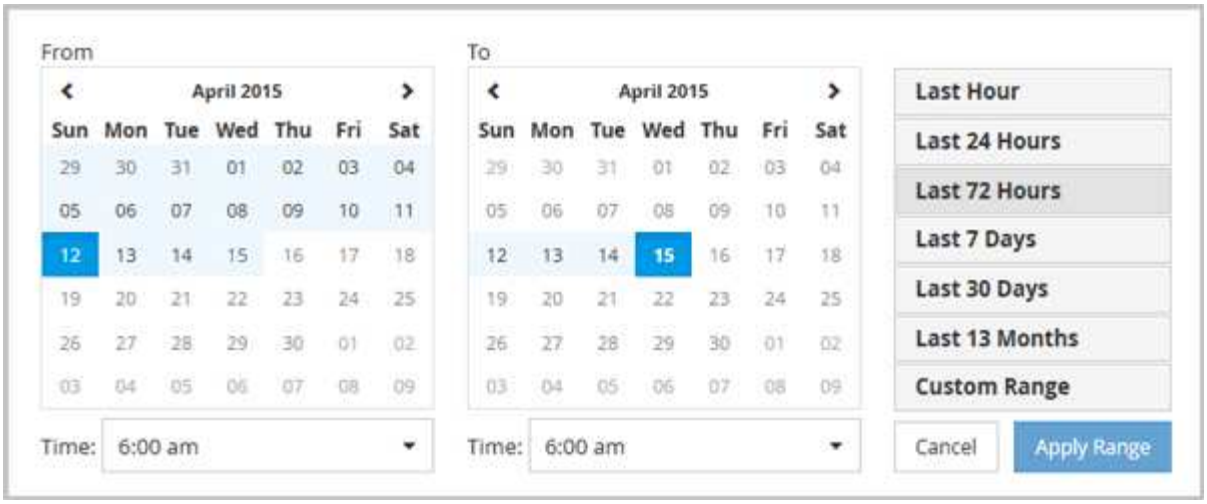
Specify the time range in Zoom View in Active IQ Unified Manager

The **Time Range** control in the Counter Charts Zoom View window enables you to specify a date and time range for the selected chart. This enables you to quickly locate specific data based on either a preset time range or your own custom time range.

You can select a time range between one hour and 390 days. 13 months equals 390 days because each month is counted as 30 days. Specifying a date and time range provides more detail and enables you to zoom in on specific performance events or series of events. Specifying a time range also aids in troubleshooting potential performance issues, as specifying a date and time range displays data surrounding the performance event in finer detail. Use the **Time Range** control to select predefined date and time ranges, or specify your own custom date and time range of up to 390 days. Buttons for predefined time ranges vary from the **Last Hour** through the **Last 13 Months**.

Selecting the **Last 13 Months** option or specifying a custom date range greater than 30 days displays a dialog box alerting you that performance data displayed for a period greater than 30 days is charted using hourly averages and not 5-minute data polling. Therefore, a loss of timeline visual granularity might occur. If you click the **Do not show again** option in the dialog box, the message does not appear when you select the **Last 13 Months** option or specify a custom date range greater than 30 days. Summary data also applies on a smaller time range, if the time range includes a time/date that is more than 30 days from today.

When selecting a time range (either custom or predefined), time ranges of 30 days or fewer are based on 5-minute interval data samples. Time ranges greater than 30 days are based on one-hour interval data samples.



1. Click the **Time Range** drop-down box and the Time Range panel displays.
2. To select a predefined time range, click one of the **Last…** buttons at the right of the **Time Range** panel. When selecting a predefined time range, data for up to 13 months is available. The predefined time range button you selected is highlighted, and the corresponding days and time display in the calendars and time selectors.
3. To select a custom date range, click the start date in the **From** calendar on the left. Click **<** or **>** to navigate forward or backward in the calendar. To specify the end date, click a date in the **To** calendar on the right. Note that the default end date is today unless you specify a different end date. The **Custom Range** button at the right of the Time Range panel is highlighted, indicating that you have selected a custom date range.
4. To select a custom time range, click the **Time** control below the **From** calendar and select the start time. To specify the end time, click the **Time** control below the **To** calendar on the right and select the end time. The **Custom Range** button at the right of the Time Range panel is highlighted, indicating that you have selected a custom time range.
5. Optionally, you can specify the start and end times when selecting a predefined date range. Select the predefined date range as previously described, then select the start and end times as previously described. The selected dates are highlighted in the calendars, your specified start and end times display in the **Time** controls, and the **Custom Range** button is highlighted.
6. After selecting the date and time range, click **Apply Range**. The performance statistics for that time range display in the charts and in the Events timeline.

Select performance thresholds in Counter Charts Zoom View in Active IQ Unified Manager

Applying thresholds in the Counter Charts Zoom View provides a detailed view of occurrences of performance threshold events. This enables you to apply or remove thresholds, and immediately view the results, which can be helpful while deciding whether troubleshooting should be your next step.

Selecting thresholds in the Counter Charts Zoom View enables you to view precise data about performance threshold events. You can apply any threshold that appears under the **Policies** area of the Counter Charts Zoom View.

Only one policy at a time can be applied to the object in the Counter Charts Zoom View.

Step

1. Select or deselect the that is associated with a policy.

The selected threshold is applied to the Counter Charts Zoom View. Critical thresholds are displayed as a red line; warning thresholds are displayed as a yellow line.

View volume latency by cluster component in Active IQ Unified Manager

You can view detailed latency information for a volume by using the Volume Performance Explorer page. The Latency - Total counter chart shows total latency on the volume, and the Latency - Breakdown counter chart is useful for determining the impact of read and write latency on the volume.

Additionally, the Latency - Cluster Components chart shows a detailed comparison of the latency of each cluster component to help determine how each component contributes to the total latency on the volume. The following cluster components are displayed:

- Network
- QoS Limit Max
- QoS Limit Min
- Network Processing
- Cluster Interconnect
- Data Processing
- Aggregate Operations
- Volume Activation
- MetroCluster Resources
- Cloud Latency
- Sync SnapMirror

Steps

1. In the **Volume Performance Explorer** page for your selected volume, from the Latency chart, select **Cluster Components** from the drop-down menu.

   The Latency - Cluster Components chart is displayed.

2. To view a larger version of the chart, select **Zoom View**.

   The cluster component comparative chart is displayed. You can restrict the comparison by deselecting or selecting the  that is associated with each cluster component.

3. To view the specific values, move your cursor into the chart area to see the popup window.

View SVM IOPS traffic by protocol in Active IQ Unified Manager

You can view detailed IOPS information for an SVM by using the Performance/SVM Explorer page. The IOPS - Total counter chart shows total IOPS usage on the SVM, and the IOPS - Breakdown counter chart is useful for determining the impact of read, write, and other IOPS on the SVM.

Additionally, the IOPS - Protocols chart shows a detailed comparison of the IOPS traffic for each protocol that is being used on the SVM. The following protocols are available:

- CIFS
- NFS
- FCP
- iSCSI
- NVMe-FC

Steps

1. In the **Performance/SVM Explorer** page for your selected SVM, from the IOPS chart, select **Protocols** from the drop-down menu.

   The IOPS - Protocols chart is displayed.

2. To view a larger version of the chart, select **Zoom View**.

   The IOPS advanced protocol comparative chart is displayed. You can restrict the comparison by deselecting or selecting the that is associated with a protocol.

3. To view the specific values, move your cursor into the chart area of either chart to see the popup window.

View volume and LUN latency charts to verify performance guarantee in Active IQ Unified Manager

You can view the volumes and LUNs that you have subscribed to the “Performance Guarantee” program to verify that latency has not exceeded the level you have been guaranteed.

The latency performance guarantee is a millisecond per operation value that should not be exceeded. It is based on an hourly average, not on the default five minute performance collection period.

Steps

1. In the **Performance: All Volumes** view or **Performance: All LUNs** view, select the volume or LUN that you are interested in.
2. In the **Performance Explorer** page for your selected volume or LUN, choose **Hourly Average** from the **View statistics in** selector.

   The horizontal line in the Latency chart will show a smoother line as the five-minute collections are replaced with the hourly average.

3. If you have other volumes on the same aggregate that are under the performance guarantee, you can add those volumes to view their latency value in the same chart.

View the performance for All SAN Array clusters in Active IQ Unified Manager

You can use the Performance: All Clusters view to display the performance status of your All SAN Array clusters.

Before you begin

You must have the Operator, Application Administrator, or Storage Administrator role.

You can view overview information for All SAN Array clusters in the Performance: All Clusters view, and details

in the Cluster / Performance Explorer page.

Steps

1. In the left navigation pane, click **Storage** > **Clusters**.
2. Make sure that the “Personality” column is displayed in the **Health: All Clusters** view, or add it using the **Show / Hide** control.

   This column displays “All SAN Array” for your All SAN Array clusters.

3. To view information about the performance in those clusters, select the **Performance: All Clusters** view.

   View the performance information for the All SAN Array cluster.

4. To view detailed information about performance in those clusters, click the name of an All SAN Array cluster.
5. Click the **Explorer** tab.
6. On the **Cluster / Performance Explorer** page, select **Nodes on this Cluster** from the **View and Compare** menu.

   You can compare the performance statistics of both nodes on this cluster to make sure the load is almost identical on both nodes. If there are large discrepancies between the two nodes you can add the second node to the charts and compare the values over a longer timeframe to identify any configuration issues.

View node IOPS based on local workloads in Active IQ Unified Manager

The node IOPS counter chart can highlight where operations are only passing through the local node using a network LIF to perform read/write operations on volumes on a remote node. The IOPS - “Total (Local)” and “Breakdown (Local)” charts display the IOPS for data that resides in local volumes only on the current node.

The “Local” versions of these counter charts are similar to the node charts for Performance Capacity and Utilization because they also show only the statistics for data that resides on local volumes.

By comparing the “Local” versions of these counter charts to the regular Total versions of these counter charts you can see if there is a lot of traffic moving through the local node to access volumes on the remote node. This situation could cause performance issues, possibly indicated by high utilization on the node, if there are too many operations passing through the local node to reach a volume on a remote node. In these cases you may want to move a volume to the local node, or to create a LIF on the remote node where traffic from hosts accessing that volume can be connected.

Steps

1. In the **Performance/Node Explorer** page for your selected node, from the IOPS chart, select **Total** from the drop-down menu.

   The IOPS - Total chart is displayed.

2. Click **Zoom View** to display a larger version of the chart in a new browser tab.
3. Back in the **Performance/Node Explorer** page, from the IOPS chart, select **Total (Local)** from the drop-down menu.

   The IOPS - Total (Local) chart is displayed.

4. Click **Zoom View** to display a larger version of the chart in a new browser tab.
5. View both of the charts next to each other and identify areas where the IOPS values appear to be quite different.
6. Move your cursor over these areas to compare the local and total IOPS for a specific point in time.

Components of the Object Landing pages

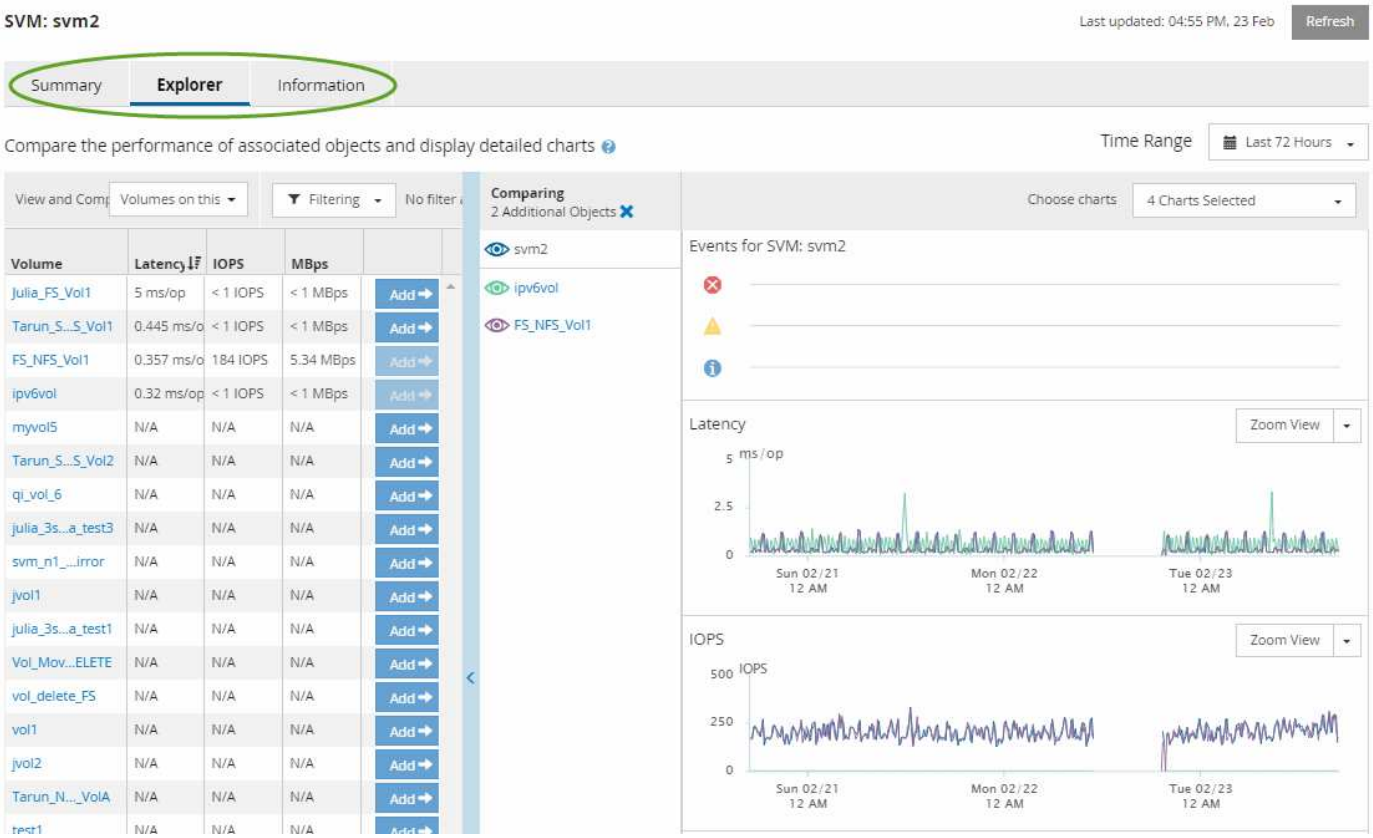
Review the components of the Object Landing pages in Active IQ Unified Manager

The Object Landing pages provide details about all critical, warning, and informational events. They provide a detailed view into the performance of all cluster objects, enabling you to select and compare individual objects across various time periods.

The Object Landing pages enable you to examine the overall performance of all objects, and to compare object performance data in a side-by-side format. This is beneficial when assessing performance and when troubleshooting events.

The data displayed in the counter summary panels and in the Counter Charts are based on a five-minute sampling interval. The data displayed in the objects inventory grid in the left side of the page is based on a one-hour sampling interval.

The following image shows an example of an Object Landing page displaying the Explorer information:



Depending on the storage object that is being viewed, the Object Landing page can have the following tabs that provide performance data about the object:

- Summary

Displays three or four counter charts containing the events and performance per object for the preceding 72-hour period, including a trend line that shows the high and low values during that period.

- Explorer

  Displays a grid of storage objects that are related to the current object, which enables you to compare the performance values of the current object with those of the related objects. This tab includes up to eleven counter charts and a time range selector, which enable you to perform a variety of comparisons.

- Information

  Displays values for non-performance configuration attributes about the storage object, including the installed version of ONTAP software, HA partner name, and number of ports and LIFs.

- Top Performers

  For clusters: Displays the storage objects that have the highest performance or the lowest performance, based on the performance counter that you select.

- Failover Planning

  For nodes: Displays the estimate of the performance impact on a node if the HA partner of the node fails.

- Details

  For volumes: Displays detailed performance statistics for all I/O activity and operations for the selected volume workload. This tab is available for FlexVol volumes, FlexGroup volumes, and constituents of FlexGroups.

Review the Performance Summary page in Active IQ Unified Manager

The Summary page displays counter charts that contain details about the events and performance per object for the preceding 72-hour period. This data is not automatically refreshed, but is current as of the last page load. The charts in the Summary page answer the question *Do I need to look further?*

Charts and counter statistics

The summary charts provide a quick, high-level overview for the last 72-hour period, and help you to identify possible issues that require further investigation.

The Summary page counter statistics are displayed in graphs.

You can position your cursor over the trend line in a graph to view the counter values for a particular point in time. The summary charts also display the total number of active critical and warning events for the preceding 72-hour period for the following counters:

- **Latency**

  Average response time for all I/O requests; expressed in milliseconds per operation.

  Displayed for all object types.

- **IOPS**

Average operating speed; expressed in input/output operations per second.

Displayed for all object types.

- **MB/s**

  Average throughput; expressed in megabytes per second.

  Displayed for all object types.

- **Performance Capacity Used**

  Percentage of performance capacity that is being consumed by a node or aggregate.

  Displayed for nodes and aggregates only.

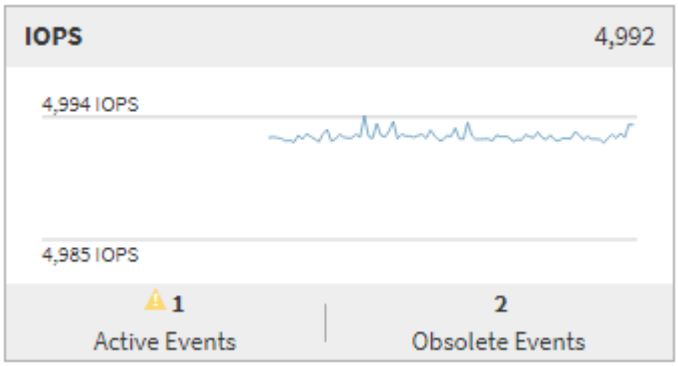
- **Utilization**

  Percentage of object utilization for nodes and aggregates, or bandwidth utilization for ports.

  Displayed for nodes, aggregates, and ports only.

Positioning the cursor over the event count for Active events shows the type and number of events. Critical events are displayed in red ( ), and warning events are displayed in yellow ( ).

The number at the top right of the chart in the gray bar is the average value from the last 72-hour period. Numbers shown at the bottom and top of the trend line graph are the minimum and maximum values for the last 72-hour period. The gray bar below the chart contains the count of active (new and acknowledged) events and obsolete events from the last 72-hour period.



- **Latency counter chart**

  The Latency counter chart provides a high-level overview of the object latency for the preceding 72-hour period. Latency refers to the average response time for all I/O requests; expressed in milliseconds per operation, the service time, wait time, or both experienced by a data packet or block in the cluster storage component under consideration.

  Top (counter value): The number in the header displays the average for the preceding 72-hour period.

  Middle (performance graph): The number at the bottom of the graph displays the lowest latency, and the number at the top of the graph displays the highest latency for the preceding 72-hour period. Position your cursor over the graph trend line to view the latency value for a specific time.

  Bottom (events): On hover, the pop-up displays the details of the events. Click the **Active Events** link

below the graph to navigate to the Events Inventory page to view complete event details.

- **IOPS counter chart**

  The IOPS counter chart provides a high-level overview of the object IOPS health for the preceding 72-hour period. IOPS indicates the speed of the storage system in number of input/output operations per second.

  Top (counter value): The number in the header displays the average for the preceding 72-hour period.

  Middle (performance graph): The number at the bottom of the graph displays the lowest IOPS, and the number at the top of the graph displays the highest IOPS for the preceding 72-hour period. Position your cursor over the graph trend line to view the IOPS value for a specific time.

  Bottom (events): On hover, the pop-up displays the details of the events. Click the **Active Events** link below the graph to navigate to the Events Inventory page to view complete event details.

- **MB/s counter chart**

  The MB/s counter chart displays the object MB/s performance, and indicates how much data has been transferred to and from the object in megabytes per second. The MB/s counter chart provides a high-level overview of the object's MB/s health for the preceding 72-hour period.

  Top (counter value): The number in the header displays the average number of MB/s for the preceding 72-hour period.

  Middle (performance graph): The value at the bottom of the graph displays the lowest number of MB/s, and the value at the top of the graph displays the highest number of MB/s for the preceding 72-hour period. Position your cursor over the graph trend line to view the MB/s value for a specific time.

  Bottom (events): On hover, the pop-up displays the details of the events. Click the **Active Events** link below the graph to navigate to the Events Inventory page to view complete event details.

- **Performance Capacity Used counter chart**

  The Performance Capacity Used counter chart displays the percentage of performance capacity that is being consumed by the object.

  Top (counter value): The number in the header displays the average used performance capacity for the preceding 72-hour period.

  Middle (performance graph): The value at the bottom of the graph displays the lowest used performance capacity percentage, and the value at the top of the graph displays the highest used performance capacity percentage for the preceding 72-hour period. Position your cursor over the graph trend line to view the used performance capacity value for a specific time.

  Bottom (events): On hover, the pop-up displays the details of the events. Click the **Active Events** link below the graph to navigate to the Events Inventory page to view complete event details.

- **Utilization counter chart**

  The Utilization counter chart displays the object utilization percentage. The Utilization counter chart provides a high-level overview of the percentage of the object or bandwidth utilization for the preceding 72-hour period.

  Top (counter value): The number in the header displays the average utilization percentage for the preceding 72-hour period.

Middle (performance graph): The value at the bottom of the graph displays the lowest utilization percentage, and the value at the top of the graph displays the highest utilization percentage for the preceding 72-hour period. Position your cursor over the graph trend line to view the utilization value for a specific time.

Bottom (events): On hover, the pop-up displays the details of the events. Click the **Active Events** link below the graph to navigate to the Events Inventory page to view complete event details.

Events

The events history table, where applicable, lists the most recent events that occurred on that object. Clicking the event name displays details of the event on the Event Details page.

Review the Performance Explorer page in Active IQ Unified Manager

The Performance Explorer page enables you to compare the performance of similar objects in a cluster—for example, all the volumes in a cluster. This is beneficial when troubleshooting performance events and fine-tuning object performance. You can also compare objects with the root object, which is the baseline against which other object comparisons are made.

You can click the **Switch to Health View** button to display the Health details page for this object. In some cases you can learn important information about the storage configuration settings for this object that may help when troubleshooting an issue.

The Performance Explorer page displays a list of cluster objects and their performance data. This page displays all the cluster objects of the same type (for example, volumes and their object-specific performance statistics) in a tabular format. This view provides an efficient overview of cluster object performance.

If “N/A” appears in any cell of the table, it means that a value for that counter is not available because there is no I/O on that object at this time.

The Performance Explorer page contains the following components:

- **Time Range**

  Enables you to select a time range for the object data.

  You can choose a predefined range, or specify your own custom time range.

- **View and Compare**

  Enables you to select which type of correlated object is displayed in the grid.

  The options available depend on the root object type and its available data. You can click the View and Compare drop-down list to select an object type. The object type that you select is displayed in the list.

- **Filtering**

  Enables you to narrow the amount of data you receive, based on your preferences.

  You can create filters that apply to the object data—for example, IOPS greater than 4. You can add up to four simultaneous filters.

- **Comparing**

  Displays a list of the objects that you have selected for comparison with the root object.

  Data for the objects in the Comparing pane is displayed in the Counter Charts.

- **View Statistics In**

  For volume and LUNs, enables you to select whether the statistics are displayed after each collection cycle (default 5 minutes), or whether the statistics are shown as an hourly average. This functionality enables you to view the latency chart in support of the NetApp "Performance Guarantee" program.

- **Counter Charts**

  Displays graphed data for each object performance category.

  Typically, only three or four charts are displayed by default. The Choose charts component enables you to display additional charts, or hide specific charts. You can also choose to show or hide the Events Timeline.

- **Events Timeline**

  Displays performance and health events occurring across the timeline that you selected in the Time Range component.

Manage performance using QoS policy group information

Manage performance using QoS policy group information in Active IQ Unified Manager

Unified Manager enables you to view the quality of service (QoS) policy groups that are available on all the clusters you are monitoring. The policies may have been defined using ONTAP software (System Manager or the ONTAP CLI) or by Unified Manager Performance Service Level policies. Unified Manager also displays which volumes and LUNs have a QoS policy group assigned.

For more information on adjusting QoS settings, see Performance management overview

Control workload throughput with storage QoS in Active IQ Unified Manager

You can create a Quality of Service (QoS) policy group to control the I/O per second (IOPS) or throughput (MB/s) limit for the workloads it contains. If the workloads are in a policy group with no set limit, such as the default policy group, or the set limit does not meet your needs, you can increase the limit or move the workloads to a new or existing policy group that has the desired limit.

“Traditional” QoS policy groups can be assigned to individual workloads; for example, a single volume or LUN. In this case the workload can use the full throughput limit. QoS policy groups also can be assigned to multiple workloads; in which case the throughput limit is “shared” among the workloads. For example, a QoS limit of 9,000 IOPS assigned to three workloads would restrict the combined IOPS from exceeding 9,000 IOPS.

“Adaptive” QoS policy groups can also be assigned to individual workloads or multiple workloads. However, even when assigned to multiple workloads, each workload gets the full throughput limit instead of sharing the

throughput value with other workloads. Additionally, adaptive QoS policies automatically adjust the throughput setting based on the volume size, per workload, thereby maintaining the ratio of IOPS to terabytes as the size of the volume changes. For example, if the peak is set to 5,000 IOPS/TB in an adaptive QoS policy, a 10 TB volume will have a throughput maximum of 50,000 IOPS. If the volume is resized later to 20 TB, adaptive QoS adjusts the maximum to 100,000 IOPS.

Starting with ONTAP 9.5 you can include the block size when defining an adaptive QoS policy. This effectively converts the policy from an IOPS/TB threshold to a MB/s threshold for cases when workloads are using very large block sizes and ultimately using a large percentage of throughput.

For shared group QoS policies, when the IOPS or MB/s of all workloads in a policy group exceeds the set limit, the policy group throttles the workloads to restrict their activity, which can decrease the performance of all workloads in the policy group. If a dynamic performance event is generated by policy group throttling, the event description displays the name of the policy group involved.

In the Performance: All Volumes view, you can sort the affected volumes by IOPS and MB/s to see which workloads have the highest usage that might have contributed to the event. In the Performance/Volumes Explorer page, you can select other volumes, or LUNs on the volume, to compare to the affected workload IOPS or MBps throughput usage.

By assigning the workloads that are overusing the node resources to a more restrictive policy group setting, the policy group throttles the workloads to restrict their activity, which can reduce the use of the resources on that node. However, if you want the workload to be able to use more of the node resources, you can increase the value of the policy group.

You can use System Manager, the ONTAP commands, or Unified Manager Performance Service Levels to manage policy groups, including the following tasks:

- Creating a policy group
- Adding or removing workloads in a policy group
- Moving a workload between policy groups
- Changing the throughput limit of a policy group
- Moving a workload to a different aggregate and/or node

View QoS policy groups available on all clusters in Active IQ Unified Manager

You can display a list of all the QoS policy groups available on the clusters that Unified Manager is monitoring. This includes traditional QoS policies, adaptive QoS policies, and QoS policies managed by Unified Manager Performance Service Level policies.

Steps

1. In the left navigation pane, click **Storage** > **QoS Policy Groups**.

   The Performance: Traditional QoS Policy Groups view is displayed by default.

2. View the detailed configuration settings for each available traditional QoS policy group.
3. Click the expand button ( ) next to the QoS Policy Group name to view more details about the policy group.
4. In the View menu, select one of the additional options to view all the adaptive QoS policy groups or to view all the QoS policy groups that were created using Unified Manager Performance Service Levels.

View volumes or LUNs within the same QoS policy group in Active IQ Unified Manager

You can display a list of the volumes and LUNs that have been assigned to the same QoS policy group.

In the case of traditional QoS policy groups that are “shared” among multiple volumes, this can be helpful to see if certain volumes are overusing the throughput defined for the policy group. It can also help you decide if you can add other volumes to the policy group without a negative affect to the other volumes.

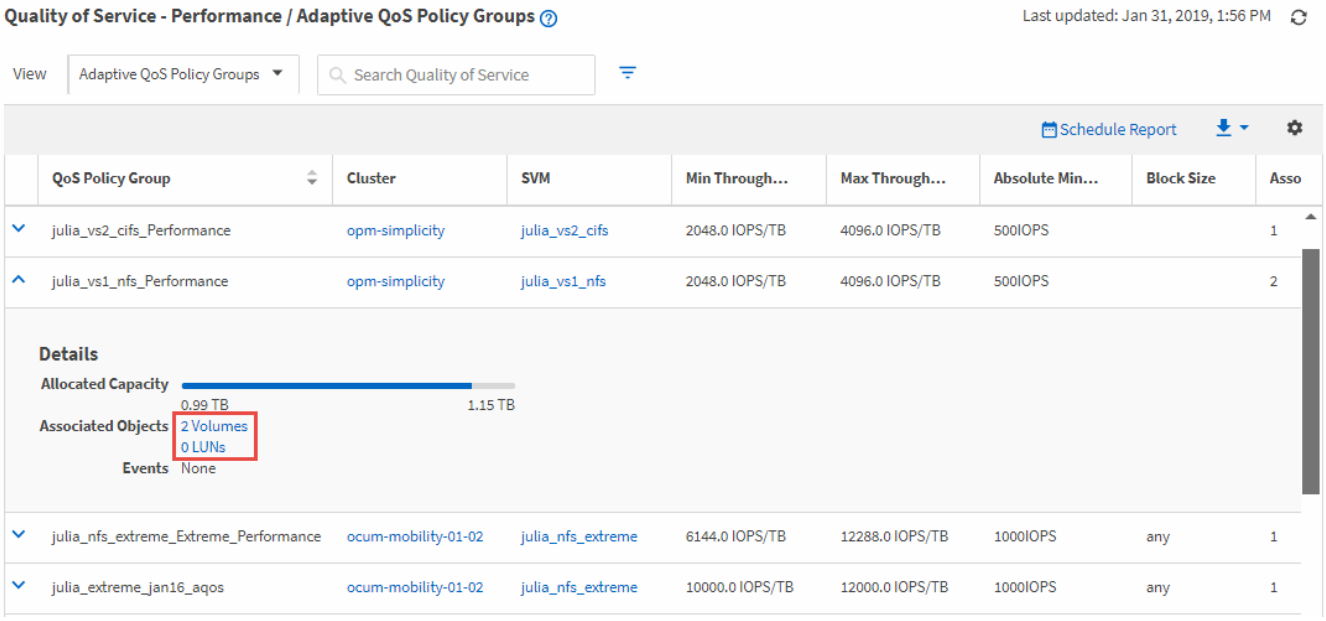
In the case of adaptive QoS policies and Unified Manager Performance Service Levels policies, this can be helpful to view all the volumes or LUNs that are using a policy group so that you can see which objects would be affected if you changed the configuration settings for the QoS policy.

Steps

1. In the left navigation pane, click **Storage** > **QoS Policy Groups**.

   The Performance: Traditional QoS Policy Groups view is displayed by default.

2. If you are interested in traditional policy group, stay on this page. Otherwise, select one of the additional View options to display all the adaptive QoS policy groups or all the QoS policy groups that were created by Unified Manager Performance Service Levels.
3. In the QoS policy that you are interested in, click the expand button ( ) next to the QoS Policy Group name to view more
   details.



4. Click the Volumes or the LUNs link to view the objects using this QoS policy.

   The Performance inventory page for Volumes or LUNs is displayed with the sorted list of objects that are using the QoS policy.

View QoS policy group settings for volumes or LUNs in Active IQ Unified Manager

You can view the QoS policy groups that have been applied to your volumes and LUNs, and you can link to the Performance/QoS Policy Groups view to display the detailed

configuration settings for each QoS policy.

The steps to view the QoS policy that is applied to a volume are shown below. The steps to view this information for a LUN are similar.

Steps

1. In the left navigation pane, click **Storage** > **Volumes**.

   The Health: All Volumes view is displayed by default.

2. In the View menu, select **Performance: Volumes in QoS Policy Group**.
3. Locate the volume that you want to review and scroll to the right until you see the **QoS Policy Group** column.
4. Click the QoS Policy Group name.

   The corresponding Quality of Service page is displayed depending on whether it is a traditional QoS policy, an adaptive QoS policy, or a QoS policy that was created using Unified Manager Performance Service Levels.

5. View the detailed configuration settings for the QoS policy group.
6. Click the expand button ( ) next to the QoS Policy Group name to view more details about the policy group.

Compare volumes or LUNs within the same QoS policy group in Active IQ Unified Manager

You can view the volumes and LUNs that are in the same QoS policy groups and then compare the performance on a single IOPS, MB/s, or IOPS/TB chart to identify any issues.

The steps to compare the performance of volumes in the same QoS policy group are shown below. The steps to view this information for a LUN are similar.

Steps

1. In the left navigation pane, click **Storage** > **Volumes**.

   The Health: All Volumes view is displayed by default.

2. In the View menu, select **Performance: Volumes in QoS Policy Group**.
3. Click the name of the volume that you want to review.

   The Performance Explorer page is displayed for the volume.

4. In the View and Compare menu, select **Volumes in same QoS Policy Group**.

   The other volumes that share the same QoS policy are listed in the table below.

5. Click the **Add** button to add those volumes to the charts so that you can compare the IOPS, MB/s, IOPS/TB, and other performance counters for all the selected volumes in the charts.

   You can change the time range to view the performance over different time intervals other than the default of 72 hours.

View QoS policies in throughput charts with Active IQ Unified Manager

You can view the ONTAP-defined quality of service (QoS) policy settings that have been applied to a volume or LUN in the Performance Explorer and Workload Analysis IOPS, IOPS/TB, and MB/s charts. The information displayed in the charts is different depending on the type of QoS policy that has been applied to the workload.

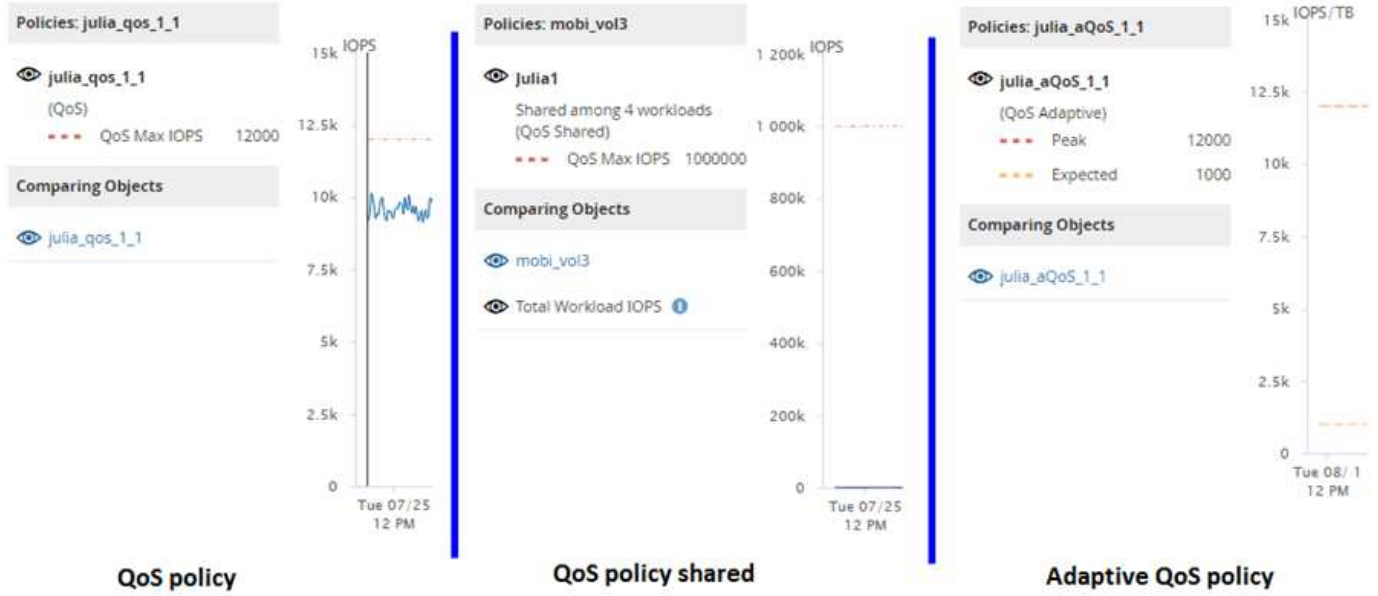
A throughput maximum (or “peak”) setting defines the maximum throughput that the workload can consume, and thereby limits the impact on competing workloads for system resources. A throughput minimum (or “expected”) setting defines the minimum throughput that must be available to the workload so that a critical workload meets minimum throughput targets regardless of demand by competing workloads.

Shared and non-shared QoS policies for IOPS and MB/s use the terms “minimum” and “maximum” to define the floor and ceiling. Adaptive QoS policies for IOPS/TB, which were introduced in ONTAP 9.3, use the terms “expected” and “peak” to define the floor and ceiling.

While ONTAP enables you to create these two types of QoS policies, depending on how they are applied to workloads there are three ways that the QoS policy will be displayed in the performance charts.

Type of policy	Functionality	Indicator in Unified Manager interface
QoS shared policy assigned to a single workload, or QoS non-shared policy assigned to a single workload or multiple workloads	Each workload can consume the specified throughput setting	Displays “(QoS)”
QoS shared policy assigned to multiple workloads	All workloads share the specified throughput setting	Displays “(QoS Shared)”
Adaptive QoS policy assigned to a single workload or multiple workloads	Each workload can consume the specified throughput setting	Displays “(QoS Adaptive)”

The following figure shows an example of how the three options are shown in the counter charts.



When a normal QoS policy that has been defined in IOPS appears in the IOPS/TB chart for a workload, ONTAP converts the IOPS value to an IOPS/TB value and Unified Manager displays that policy in the IOPS/TB chart along with the text “QoS, defined in IOPS”.

When an adaptive QoS policy that has been defined in IOPS/TB appears in the IOPS chart for a workload, ONTAP converts the IOPS/TB value to an IOPS value and Unified Manager displays that policy in the IOPS chart along with the text “QoS Adaptive - Used, defined in IOPS/TB” or “QoS Adaptive - Allocated, defined in IOPS/TB” depending on how the peak IOPS allocation setting is configured. When the allocation setting is set to “allocated-space”, the peak IOPS is calculated based on the size of the volume. When the allocation setting is set to “used-space”, the peak IOPS is calculated based on the amount of data stored in the volume, taking into account storage efficiencies.



The IOPS/TB chart displays performance data only when the logical capacity used by the volume is greater than or equal to 128 GB. Gaps are displayed in the chart when the used capacity falls below 128 GB during the selected timeframe.

View workload QoS minimum and maximum settings in Active IQ Unified Manager

You can view the ONTAP-defined quality of service (QoS) policy settings on a volume or LUN in the Performance Explorer charts. A throughput maximum setting limits the impact of competing workloads on system resources. A throughput minimum setting ensures that a critical workload meets minimum throughput targets regardless of demand by competing workloads.

QoS throughput “minimum” and “maximum” IOPS and MB/s settings are displayed in the counter charts only if they have been configured in ONTAP. Throughput minimum settings are available only on systems running ONTAP 9.2 or later software, only on AFF systems, and they can be set only for IOPS at this time.

Adaptive QoS policies are available starting with ONTAP 9.3 and are expressed using IOPS/TB instead of IOPS. These policies automatically adjust the QoS policy value based on the volume size, per workload, thereby maintaining the ratio of IOPS to terabytes as the size of the volume changes. You can apply an adaptive QoS policy group to volumes only. The QoS terminology “expected” and “peak” are used for adaptive QoS policies instead of minimum and maximum.

Unified Manager generates warning events for QoS policy breaches when workload throughput has exceeded the defined QoS maximum policy setting during each performance collection period for the previous hour. Workload throughput may exceed the QoS threshold for only a short period of time during each collection period, but Unified Manager displays the “average” throughput during the collection period on the chart. For this reason you may see QoS events while the throughput for a workload might not have crossed the policy threshold shown in the chart.

Steps

1. In the **Performance Explorer** page for your selected volume or LUN, perform the following actions to view the QoS ceiling and floor settings:

If you want to…	Do this…
View the IOPS ceiling (the QoS max)	In the IOPS Total or Breakdown chart, click **Zoom View**.
View the MB/s ceiling (the QoS max)	In the MB/s Total or Breakdown chart, click **Zoom View**.

If you want to…	Do this…
View the IOPS floor (the QoS min)	In the IOPS Total or Breakdown chart, click **Zoom View**.
View the IOPS/TB ceiling (the QoS peak)	For volumes, in the IOPS/TB chart, click **Zoom View**.
View the IOPS/TB floor (the QoS expected)	For volumes, in the IOPS/TB chart, click **Zoom View**.

The dashed, horizontal line indicates the maximum or minimum throughput value set in ONTAP. You can also view when changes to the QoS values were implemented.

2. To view the specific IOPS and MB/s values compared to the QoS setting, move your cursor into the chart area to see the popup window.

If you notice that certain volumes or LUNs have very high IOPS or MB/s and are stressing system resources, you can use System Manager or the ONTAP CLI to adjust the QoS settings so that these workloads do not affect the performance of other workloads.

For more information on adjusting QoS settings, see Performance management overview

Manage performance using performance capacity and available IOPS information

Manage performance capacity and available IOPS in Active IQ Unified Manager

Performance capacity indicates how much throughput you can get out of a resource without surpassing the useful performance of that resource. When viewed using existing performance counters, performance capacity is the point at which you get the maximum utilization from a node or aggregate before latency becomes an issue.

Unified Manager collects performance capacity statistics from nodes and aggregates in each cluster. *Performance capacity used* is the percentage of performance capacity that is currently being used, and *performance capacity free* is the percentage of performance capacity that is still available.

While performance capacity free provides a percentage of the resource that is still available, *available IOPS* tells you the number of IOPS that can be added to the resource before reaching the maximum performance capacity. By using this metric, you can be sure that you can add workloads of a predetermined number of IOPS to a resource.

Monitor the performance capacity information has the following benefits:

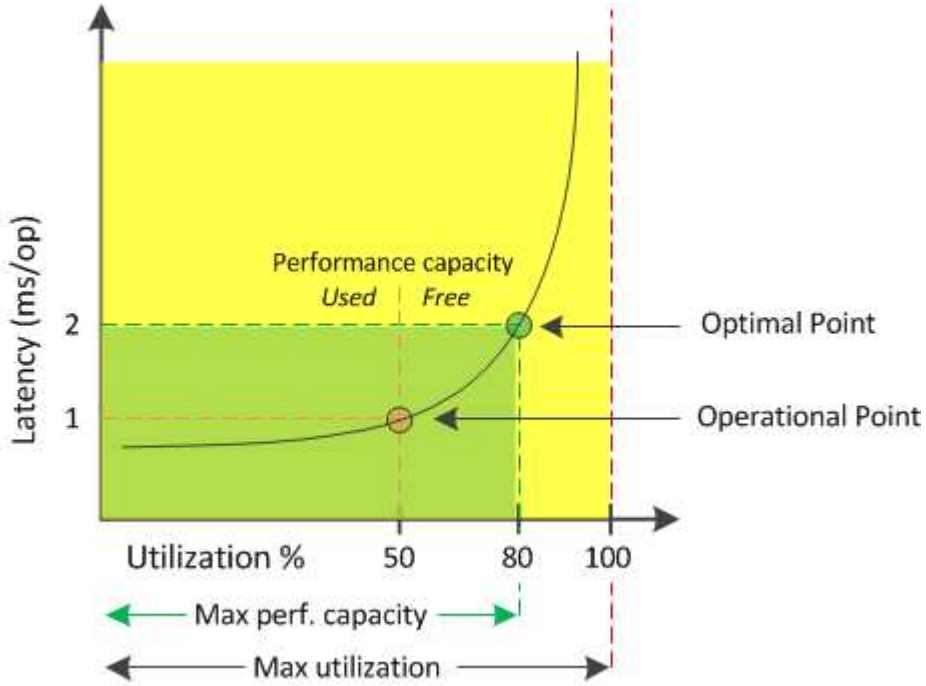
- Assists with workflow provisioning and balancing.
- Helps you prevent overloading a node or pushing its resources beyond the optimal point, thus reducing the need to troubleshoot.
- Helps you determine with greater precision where additional storage equipment might be needed.

Understand the performance capacity used counter in Active IQ Unified Manager

The performance capacity used counter helps you to identify whether the performance of

a node or an aggregate is reaching a point where the performance might degrade if the workloads increase. It can also show you if a node or aggregate is currently being overused during specific periods of time. Performance capacity used is similar to utilization, but the former provides more insight about the available performance capabilities in a physical resource for a specific workload.

The optimal used performance capacity is the point at which a node or an aggregate has optimal utilization and latency (response time), and is being used efficiently. A sample latency versus utilization curve is shown for an aggregate in the following figure.



In this example, the *operational point* identifies that the aggregate is currently operating at 50% utilization with latency of 1.0 ms/op. Based on the statistics captured from the aggregate, Unified Manager determines that additional performance capacity is available for this aggregate. In this example, the *optimal point* is identified as the point when the aggregate is at 80% utilization with latency of 2.0 ms/op. Therefore, you can add more volumes and LUNs to this aggregate so that your systems are used more efficiently.

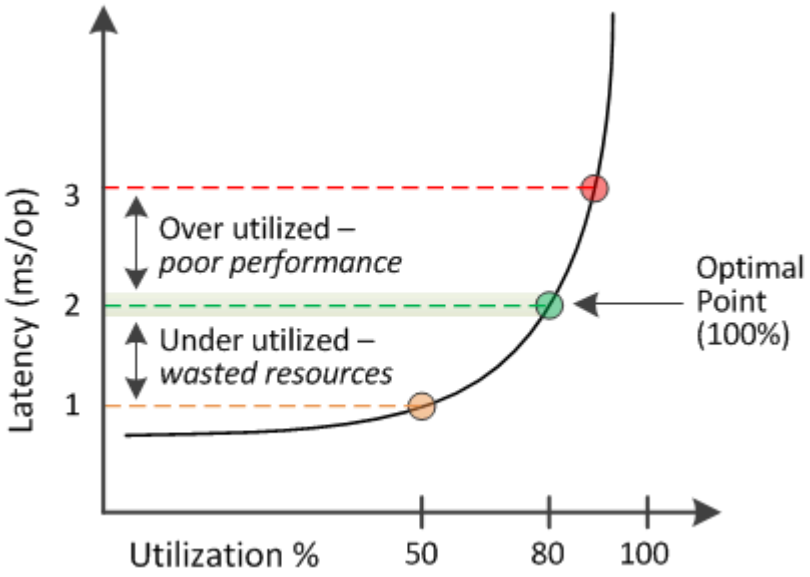
The performance capacity used counter is expected to be a larger number than the “utilization” counter because performance capacity adds in the impact on latency. For example, if a node or aggregate is 70% used, the performance capacity value may be in the 80% to 100% range, depending on the latency value.

In some cases, however, the utilization counter may be higher on the Dashboard page. This is normal because the dashboard refreshes the current counter values at each collection period; it does not display averages over a period of time like the other pages in the Unified Manager user interface. The performance capacity used counter is best used as an indicator of performance averaged over a period of time, whereas the utilization counter is best used for determining the instantaneous usage of a resource.

Use the performance capacity used value in Active IQ Unified Manager

The performance capacity used value helps you identify the nodes and aggregates that are currently being overutilized or underutilized. This enables you to redistribute workloads in order to make your storage resources more efficient.

The following figure shows the latency versus utilization curve for a resource and identifies, with colored dots, three areas where the current operational point could be located.



- A performance capacity used percentage equal to 100 is at the optimal point.

  Resources are being used efficiently at this point.

- A performance capacity used percentage above 100 indicates that the node or aggregate is overutilized, and that workloads are receiving sub-optimal performance.

  No new workloads should be added to the resource, and the existing workloads may need to be redistributed.

- A performance capacity used percentage below 100 indicates that the node or aggregate is underutilized, and that resources are not being used effectively.

  More workloads can be added to the resource.

Unlike utilization, the performance capacity used percentage can be above 100%. There is no maximum percentage, but resources will typically be in the 110% to 140% range when they are being overutilized. Higher percentages would indicate a resource with serious issues.

Understand available IOPS in Active IQ Unified Manager

The available IOPS counter identifies the remaining number of IOPS that can be added to a node or an aggregate before the resource reaches its limit.

The total IOPS that a node can provide is based on the physical characteristics of the node—for example, the number of CPUs, the CPU speed, and the amount of RAM. The total IOPS that an aggregate can provide is based on the physical properties of the disks—for example, a SATA, SAS, or SSD disk.

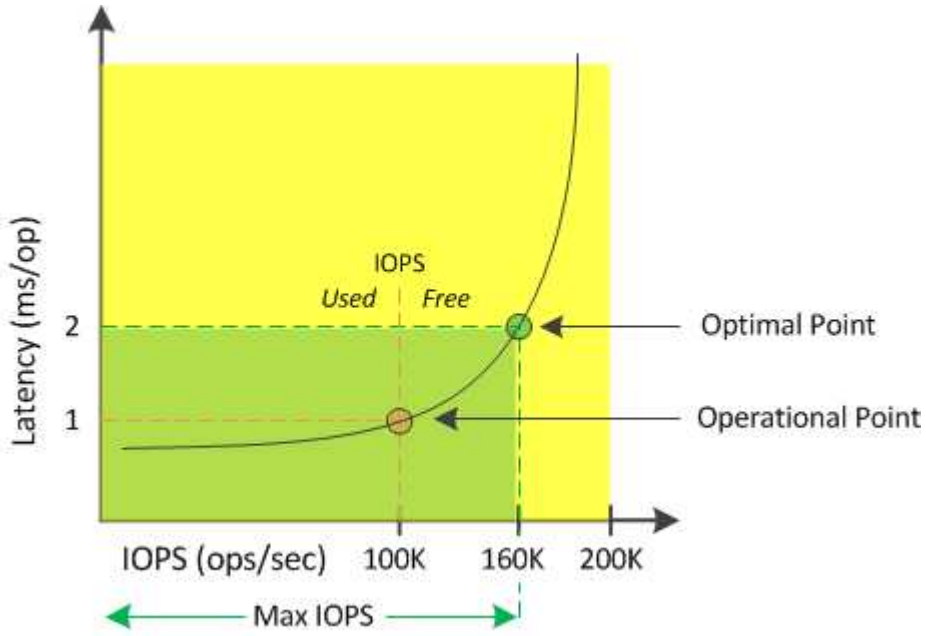
The total IOPS of all the volumes in an aggregate might not match the total IOPS of the aggregate. This is discussed in the following knowledge base article: KB Why does the sum of all volume IOPs in an aggregate not match the aggregate IOPs?

While the performance capacity free counter provides the percentage of a resource that is still available, the

available IOPS counter tells you the exact number of IOPS (workloads) can be added to a resource before reaching the maximum performance capacity.

For example, if you are using a pair of FAS2520 and FAS8060 storage systems, a performance capacity free value of 30% means that you have some free performance capacity. However, that value does not provide visibility into how many more workloads you can deploy to those nodes. The available IOPS counter may show that you have 500 available IOPS on the FAS8060, but only 100 available IOPS on the FAS2520.

A sample latency versus IOPS curve for a node is shown in the following figure.



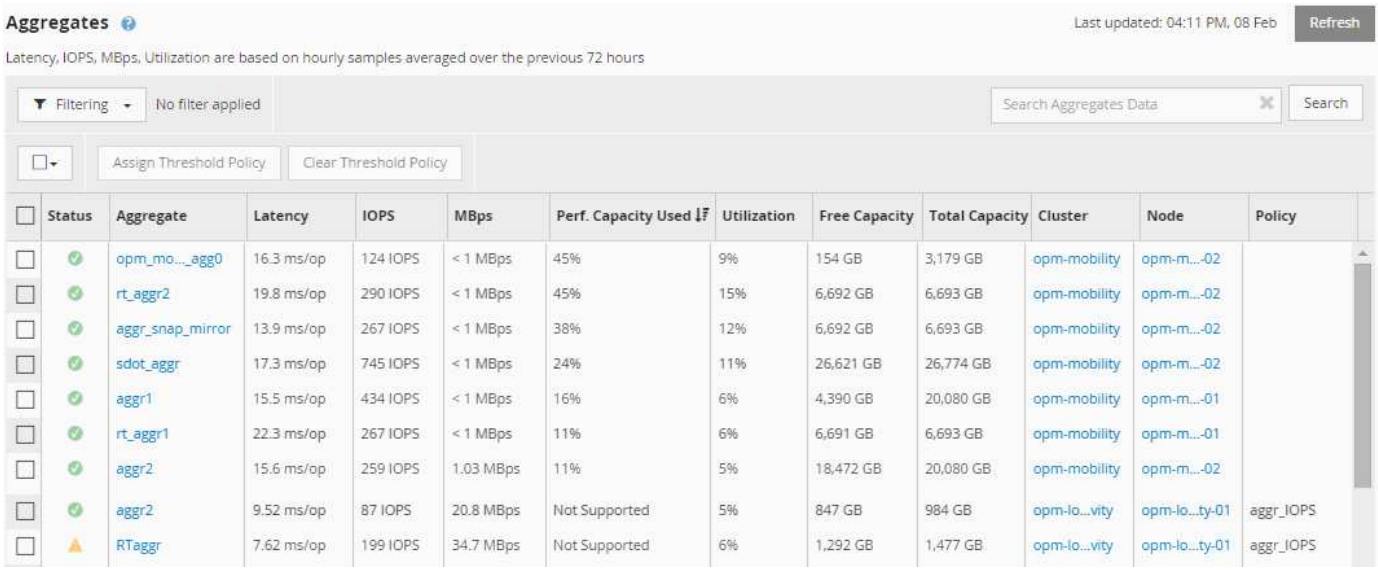
The maximum number of IOPS that a resource can provide is the number of IOPS when the performance capacity used counter is at 100% (the optimal point). The operational point identifies that the node is currently operating at 100K IOPS with latency of 1.0 ms/op. Based on the statistics captured from the node, Unified Manager determines that the maximum IOPS for the node is 160K, which means that there are 60K free or available IOPS. Therefore, you can add more workloads to this node so that your systems are used more efficiently.

When there is minimal user activity in the resource, the available IOPS value is calculated assuming a generic workload based on approximately 4,500 IOPS per CPU core. This is because Unified Manager lacks the data to accurately estimate the characteristics of the workload being served.

View node and aggregate performance capacity used values in Active IQ Unified Manager

You can monitor the performance capacity used values for all nodes or for all aggregates in a cluster, or you can view details for a single node or aggregate.

Performance capacity used values appear in the Dashboard, Performance Inventory pages, Top Performers page, Create Threshold Policy page, Performance Explorer pages, and in detail charts. For example, the Performance: All Aggregates page provides a column Performance Capacity Used to view the performance capacity used value for all aggregates.



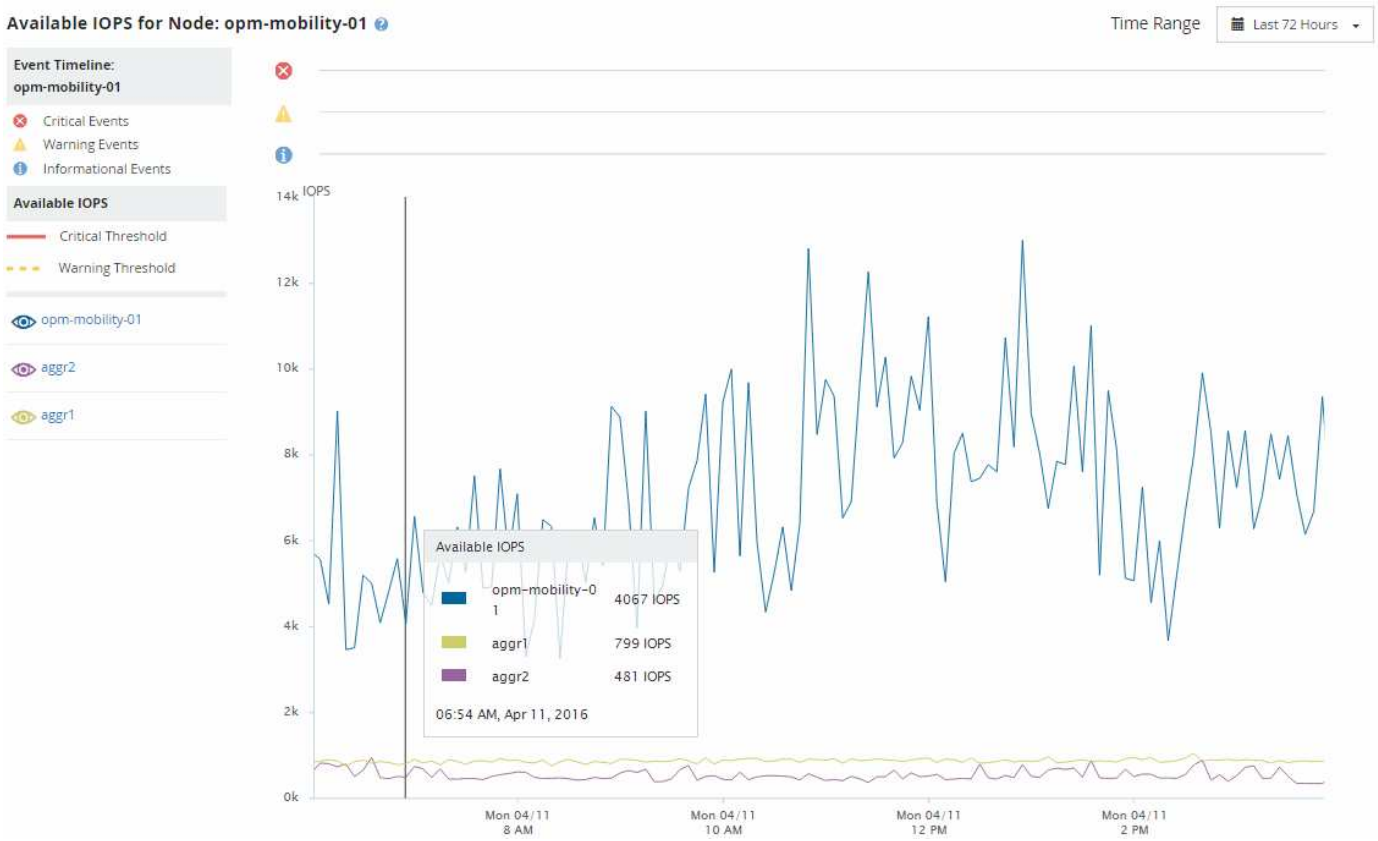
Monitor the performance capacity used counter enables you to identify the following:

- Whether any nodes or aggregates on any clusters have a high performance capacity used value
- Whether any nodes or aggregates on any clusters have active performance capacity used events
- The nodes and aggregates that have the highest and lowest performance capacity used value in a cluster
- Latency and utilization counter values in conjunction with nodes or aggregates that have high performance capacity used values
- How the performance capacity used values for nodes in an HA pair will be affected if one of the nodes fails
- The busiest volumes and LUNs on an aggregate that has a high performance capacity used value

View available IOPS values for nodes and aggregates in Active IQ Unified Manager

You can monitor the available IOPS values for all nodes or for all aggregates in a cluster, or you can view details for a single node or aggregate.

Available IOPS values appear in the Performance Inventory pages and in the Performance Explorer page charts for nodes and aggregates. For example, when viewing a node in the Node/Performance Explorer page, you can select the “Available IOPS” counter chart from the list so you can compare the available IOPS values for the node and multiple aggregates on that node.



Monitor the available IOPS counter enables you to identify:

- The nodes or aggregates that have the greatest available IOPS values to help determine where future workloads can be deployed.
- The nodes or aggregates that have the smallest available IOPS values to identify the resources you should monitor for potential future performance issues.
- The busiest volumes and LUNs on an aggregate that has a small available IOPS value.

View performance capacity counter charts in Active IQ Unified Manager

You can view performance capacity used charts for nodes and aggregates on the Performance Explorer page. This enables you to view detailed performance capacity data for the selected nodes and aggregates for a specific timeframe.

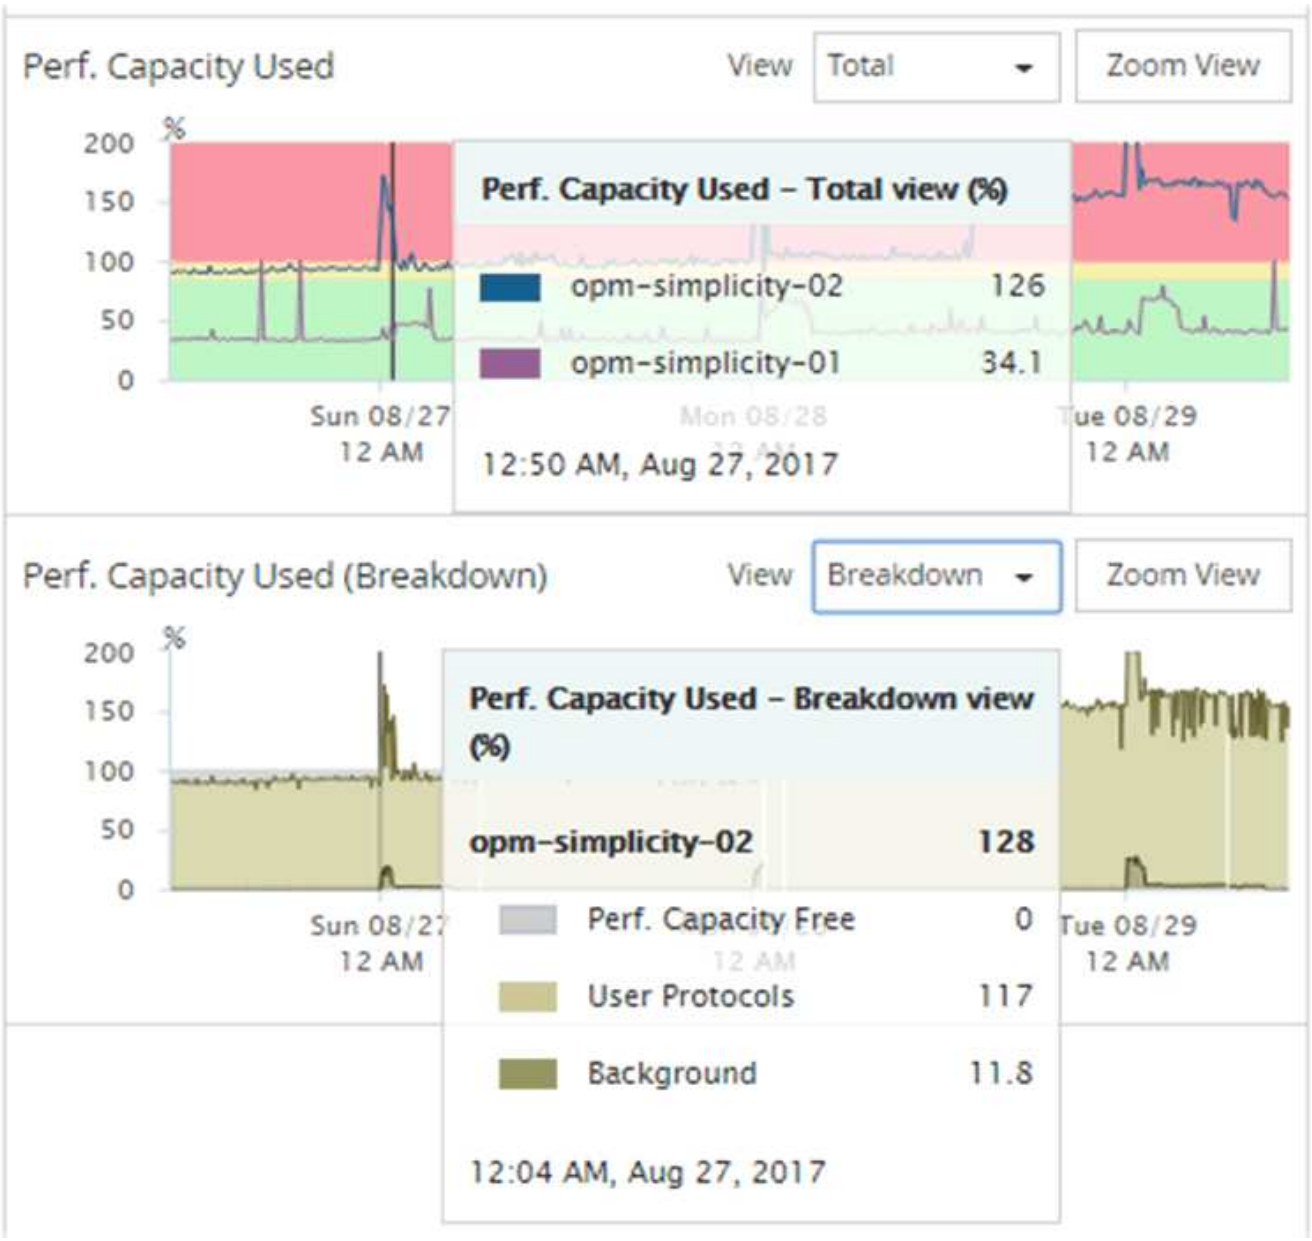
The standard counter chart displays the performance capacity used values for the selected nodes or aggregates. The Breakdown counter chart displays the total performance capacity values for the root object separated into usage based on user protocols versus background system processes. Additionally, the amount of free performance capacity is also shown.

Because some background activities associated with system and data management are identified as user workloads and categorized as user protocols, the user protocols percentage may appear artificially high when those processes run. These processes typically run around midnight when cluster usage is low. If you see a spike in user protocol activity around midnight, verify if cluster backup jobs or other background activities are configured to run at that time.

Steps

1. Select the **Explorer** tab from a node or aggregate **Landing** page.
2. In the **Counter Charts** pane, click **Choose charts**, and then select the **Perf. Capacity Used** chart.
3. Scroll down until you can view the chart.

   The colors of the standard chart show when the object is in the optimal range (yellow), when the object is underutilized (green), and when the object is overutilized (red). The Breakdown chart shows detailed performance capacity details for the root object only.



4. If you want to view either chart in a full size format, click **Zoom View**.

   In this manner you can open multiple counter charts in a separate windows to compare performance capacity used values with IOPS or MBps values over the same timeframe.

Configure performance capacity used thresholds in Active IQ Unified Manager

You can create user-defined performance threshold policies so that events are triggered

when the performance capacity used value for a node or aggregate exceeds the defined performance capacity used threshold setting.

Additionally, nodes can be configured with a “Performance capacity used takeover” threshold policy. This threshold policy totals the performance capacity used statistics for both nodes in an HA pair to determine whether either node would lack sufficient capacity if the other node fails. Because the workload during failover is the combination of the two partner nodes’ workloads, the same performance capacity used takeover policy can be applied to both nodes.



This performance capacity used equivalency is generally true between nodes. However, if there is significantly more cross-node traffic destined for one of the nodes through its failover partner, the total performance capacity used when running all workloads on one partner node versus the other partner node could be slightly different depending on which node has failed.

The performance capacity used conditions can also be used as secondary performance threshold settings to create a combination threshold policy when defining thresholds for LUNs and volumes. The performance capacity used condition is applied to the aggregate or node on which the volume or LUN resides. For example, you can create a combination threshold policy using the following criteria:

Storage object	Performance counter	Warning threshold	Critical threshold	Duration
Volume	Latency	15 ms/op	25 ms/op	20 minutes
Aggregate	Performance capacity used	80%	95%	

Combination threshold policies cause an event to be generated only when both conditions are breached for the entire duration.

Use the performance capacity used counter in Active IQ Unified Manager

Typically, organizations want to operate with a performance capacity used percentage below 100 so that resources are being efficiently used while reserving some additional performance capacity to support peak period demands. You can use threshold policies to customize when alerts are sent for high performance capacity used values.

You can establish specific goals based on your performance requirements. For example, financial services firms might reserve more performance capacity to guarantee the timely execution of trades. These companies might want to set performance capacity used thresholds in the 70-80 percent range. Manufacturing companies with smaller margins might choose to reserve less performance capacity if they are willing to risk performance to better manage IT costs. These companies might set performance capacity used thresholds in the 85-95 percent range.

When the performance capacity used value exceeds the percentage set in a user-defined threshold policy, Unified Manager sends an alert email and adds the event to the Event Inventory page. This enables you to manage potential problems before they impact performance. These events can also be used as indicators that you need to make workload moves and changes within your nodes and aggregates.

Understand and use the Node Failover Planning page

Use the Node Failover Planning page in Active IQ Unified Manager

The Performance/Node Failover Planning page estimates the performance impact on a node if the node's high-availability (HA) partner node fails. Unified Manager bases the estimates on the historical performance of the nodes in the HA pair.

Estimating the performance impact of a failover helps you to plan in the following scenarios:

- If a failover consistently degrades the takeover node's estimated performance to an unacceptable level, you can consider taking corrective actions to reduce the performance impact due to a failover.
- Before initiating a manual failover to perform hardware maintenance tasks, you can estimate how the failover affects the performance of the takeover node in order to determine the best time to perform the task.

Determine corrective actions from the Performance/Node Failover Planning page in Active IQ Unified Manager

Based on the information that is displayed in the Performance/Node Failover Planning page, you can take actions to ensure that a failover does not cause the performance of an HA pair to drop below an acceptable level.

For example, to reduce the estimated performance impact of a failover, you can move some volumes or LUNs from a node in the HA pair to other nodes in the cluster. Doing so ensures that the primary node can continue to deliver acceptable performance after a failover.

Review the components of the Performance/Node Failover Planning page in Active IQ Unified Manager

The components of the Performance/Node Failover Planning page are displayed in a grid and in the Comparing pane. These sections enable you to assess the impact of a node failover on the performance of the takeover node.

Performance statistics grid

The Performance/Node Failover Planning page displays a grid containing statistics for latency, IOPS, utilization, and performance capacity used.

> Latency and IOPS values displayed in this page and in the Performance/Node Performance Explorer page might not match because different performance counters are used to calculate the values to predict node failover.

In the grid, each node is assigned one of the following roles:

- Primary

  The node that takes over for the HA partner when the partner fails. The root object is always the Primary node.

- Partner

  The node that fails in the failover scenario.

- Estimated Takeover

  The same as the Primary node. Performance statistics displayed for this node show the takeover node's performance after it takes over the failed partner.

> Although the workload of the takeover node is equivalent to the combined workloads of both nodes after a failover, the statistics for the Estimated Takeover node are not the sum of the statistics of the Primary node and the Partner node. For example, if the latency of the Primary node is 2 ms/op and the latency of the Partner node is 3 ms/op, the Estimated Takeover node might have a latency of 4 ms/op. This value is a calculation that Unified Manager performs.

You can click the name of the Partner node if you want it to become the root object. After the Performance/Node Performance Explorer page is displayed, you can click the **Failover Planning** tab to see how performance changes in this node failure scenario. For example, if Node1 is the Primary node and Node2 is the Partner node, you can click Node2 to make it the Primary node. In this way, you can see how the estimated performance changes depending on which node fails.

Comparing pane

The following list describes the components displayed in the Comparing pane by default:

- **Events charts**

  They are displayed in the same format as those in the Performance/Node Performance Explorer page. They pertain to the Primary node only.

- **Counter charts**

  They display historical statistics for the performance counter shown in the grid. In each chart, the graph for the Estimated Takeover node shows the estimated performance if a failover had occurred at any given time.

  For example, suppose the Utilization chart shows 73% for the Estimated Takeover node at 11 a.m. on February 8. If a failover had occurred at that time, the utilization of the takeover node would have been 73%.

  The historical statistics help you find the optimal time for initiating a failover, minimizing the possibility of overloading the takeover node. You can schedule a failover only at times when the predicted performance of the takeover node is acceptable.

By default, statistics for both the root object and the partner node are displayed in the Comparing pane. Unlike in the Performance/Node Performance Explorer page, this page does not display the **Add** button for you to add objects for statistics comparison.

You can customize the Comparing pane in the same way as you do in the Performance/Node Performance Explorer page. The following list shows examples of customizing the charts:

- Click a node name to show or hide the node's statistics in the Counter charts.
- Click **Zoom View** to display a detailed chart for a particular counter in a new window.

Create threshold policies for node failover planning in Active IQ Unified Manager

You can create a node threshold policy so that you can be notified in the

Performance/Node Failover Planning page when a potential failover would degrade the performance of the takeover node to an unacceptable level.

The system-defined performance threshold policy named “Node HA pair over-utilized” generates a warning event if the threshold is breached for six consecutive collection periods (30 minutes). The threshold is considered breached if the combined performance capacity used of the nodes in an HA pair exceeds 200%.

The event from the system-defined threshold policy alerts you to the fact that a failover will cause the latency of the takeover node to increase to an unacceptable level. When you see an event that is generated by this policy for a particular node, you can navigate to the Performance/Node Failover Planning page for that node to view the predicted latency value due to a failover.

In addition to using this system-defined threshold policy, you can create threshold policies by using the “Performance Capacity Used - Takeover” counter, and then apply the policy to selected nodes. Specifying a threshold lower than 200% enables you to receive an event before the threshold for the system-defined policy is breached. You can also specify the minimum period of time for which the threshold is exceeded to less than 30 minutes if you want to be notified before the system-defined policy event is generated.

For example, you can define a threshold policy to generate a warning event if the combined performance capacity used of the nodes in an HA pair exceeds 175% for more than 10 minutes. You can apply this policy to Node1 and Node2, which form an HA pair. After receiving a warning event notification for either Node1 or Node2, you can view the Performance/Node Failover Planning page for that node to assess the estimated performance impact on the takeover node. You can take corrective actions to avoid overloading the takeover node if a failover does happen. If you take action when the combined performance capacity used of the nodes is under 200%, the takeover node’s latency does not reach an unacceptable level even if a failover happens during this time.

Use the Performance Capacity Used Breakdown chart for failover planning in Active IQ Unified Manager

The detailed Performance Capacity Used - Breakdown chart shows the performance capacity used for the Primary node and the Partner node. It also shows the amount of free performance capacity on the Estimated Takeover node. This information helps you determine whether you might have a performance issue if the partner node fails.

In addition to showing the total performance capacity used for the nodes, the Breakdown chart breaks the values for each node into user protocols and background processes.

- User protocols are the I/O operations from user applications to and from the cluster.
- Background processes are the internal system processes involved with storage efficiency, data replication, and system health.

This additional level of detail enables you to determine whether a performance issue is caused by user application activity or background system processes, such as deduplication, RAID reconstruct, disk scrubbing, and SnapMirror copies.

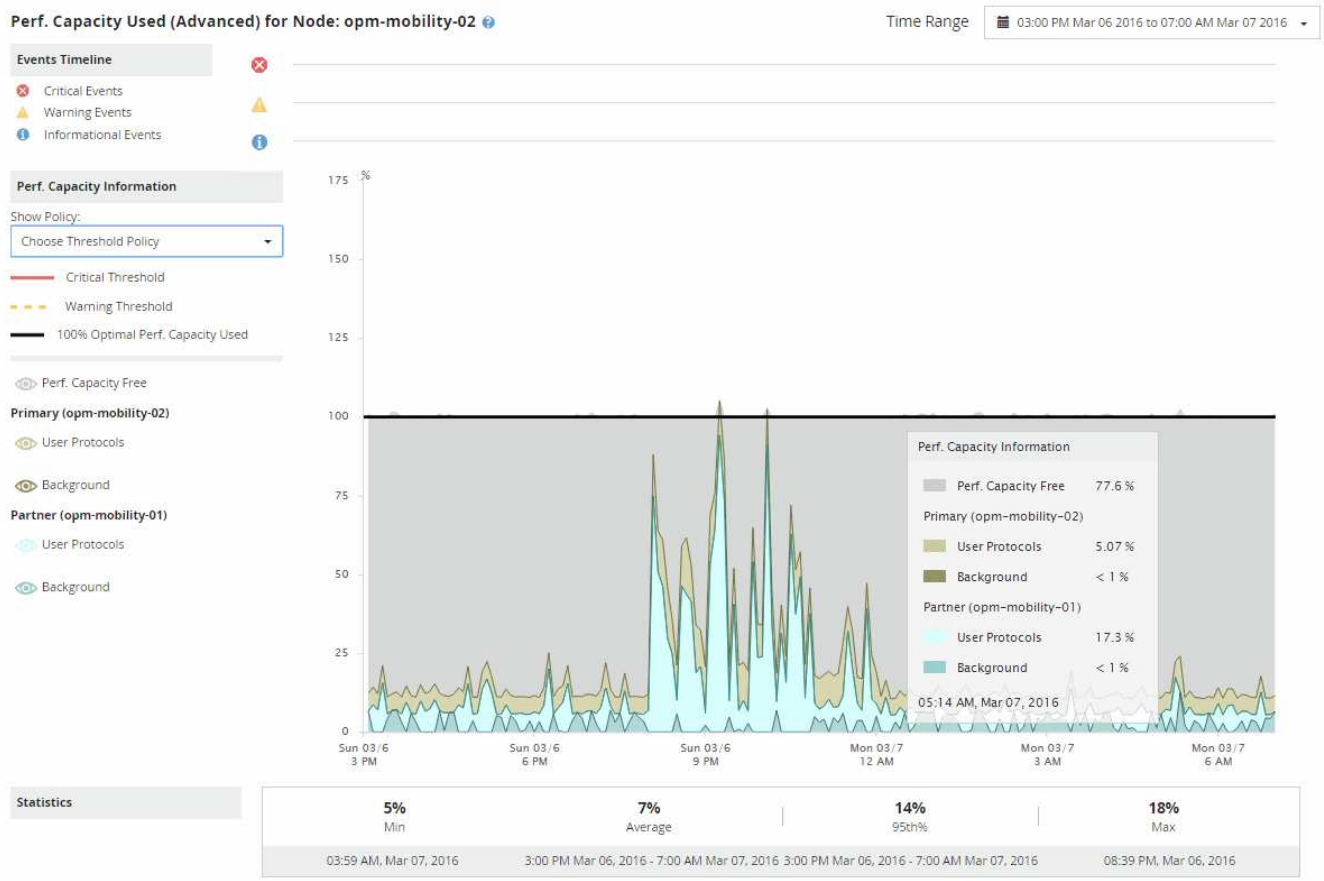
Steps

1. Go to the **Performance/Node Failover Planning** page for the node that will serve as the Estimated Takeover node.
2. From the **Time Range** selector, choose the period of time for which the historical statistics are displayed in the counter grid and counter charts.

The counter charts with statistics for the Primary node, Partner node, and Estimated Takeover node are displayed.

3. From the **Choose charts** list, select **Perf. Capacity Used**.
4. In the **Perf. Capacity Used** chart, select **Breakdown** and click **Zoom View**.

   The detailed chart for Perf. Capacity Used is displayed.



5. Move the cursor over the detailed chart to view the performance capacity used information in the popup window.

   The Perf. Capacity Free percentage is the performance capacity available on the Estimated Takeover node. It indicates how much performance capacity is left on the takeover node after a failover. If it is 0%, a failover will cause the latency to increase to an unacceptable level on the takeover node.

6. Consider taking corrective actions to avoid a low performance capacity free percentage.

   If you plan to initiate a failover for node maintenance, choose a time to fail the partner node when the performance capacity free percentage is not at 0.

Collect data and monitor workload performance

Collect data and monitor workload performance in Active IQ Unified Manager

Unified Manager collects and analyzes workload activity every 5 minutes to identify performance events, and it detects configuration changes every 15 minutes. It retains a

maximum of 30 days of 5-minute historical performance and event data, and it uses this data to forecast the expected latency range for all monitored workloads.

Unified Manager must collect a minimum of 3 days of workload activity before it can begin its analysis and before the latency forecast for I/O response time can be displayed on the Workload Analysis page and in the Event details page. While this activity is being collected, the latency forecast does not display all changes occurring from workload activity. After collecting 3 days of activity, Unified Manager adjusts the latency forecast every 24 hours at 12:00 a.m., to reflect workload activity changes and establish a more accurate dynamic performance threshold.

During the first 4 days that Unified Manager is monitoring a workload, if more than 24 hours have passed since the last data collection, the latency charts will not display the latency forecast for that workload. Events detected prior to the last collection are still available.

Daylight savings time (DST) changes the system time, which alters the latency forecast of performance statistics for monitored workloads. Unified Manager immediately begins to correct the latency forecast, which takes approximately 15 days to complete. During this time you can continue to use Unified Manager, but, since Unified Manager uses the latency forecast to detect dynamic events, some events might not be accurate. Events detected prior to the time change are not affected.

Monitor workload types with Active IQ Unified Manager

You can use Unified Manager to monitor the performance of two types of workloads: user-defined and system-defined.

- ***User-defined workloads***

  The I/O throughput from applications to the cluster. These are processes involved in read and write requests. A volume, LUN, NFS share, SMB/CIFS share, and a workload is a user-defined workload.

  Unified Manager only monitors the workload activity on the cluster. It does not monitor the applications, the clients, or the paths between the applications and the cluster.

  If one or more of the following is true for a workload, it cannot be monitored by Unified Manager:

  - It is a data protection (DP) copy in read-only mode. (DP volumes are monitored for user-generated traffic.)
  - It is an offline data clone.
  - It is a mirrored volume in a MetroCluster configuration.

- ***System-defined workloads***

  The internal processes involved with storage efficiency, data replication, and system health, including:

  - Storage efficiency, such as deduplication
  - Disk health, which includes RAID reconstruct, disk scrubbing, and so on
  - Data replication, such as SnapMirror copies
  - Management activities
  - File system health, which includes various WAFL activities

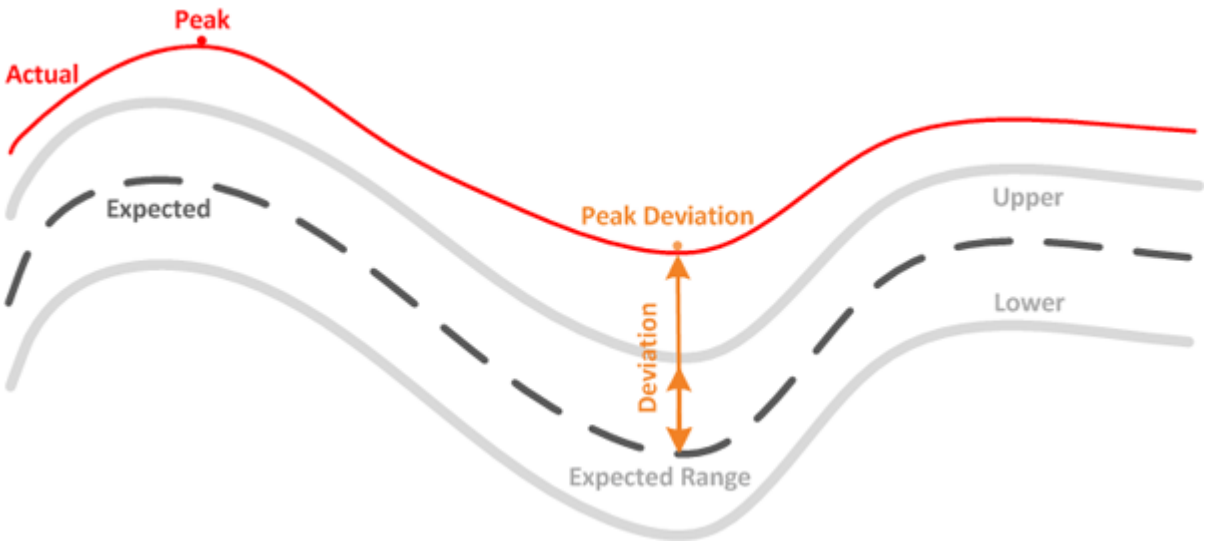
- File system scanners, such as WAFL scan
- Copy offload, such as offloaded storage efficiency operations from VMware hosts
- System health, such as volume moves, data compression, and so on
- Unmonitored volumes

Performance data for system-defined workloads is displayed in the GUI only when the cluster component used by these workloads is in contention. For example, you cannot search for the name of a system-defined workload to view its performance data in the GUI.

Understand workload performance measurement values in Active IQ Unified Manager

Unified Manager measures the performance of workloads on a cluster based on historical and expected statistical values, which form the latency forecast of values for the workloads. It compares the actual workload statistical values to the latency forecast to determine when workload performance is too high or too low. A workload that is not performing as expected triggers a dynamic performance event to notify you.

In the following illustration, the actual value, in red, represents the actual performance statistics in the time frame. The actual value has crossed the performance threshold, which is the upper bounds of the latency forecast. The peak is the highest actual value in the time frame. The deviation measures the change between the expected values (the forecast) and the actual values, while the peak deviation indicates the largest change between the expected values and the actual values.



The following table lists the workload performance measurement values.

Measurement	Description
Activity	The percentage of the QoS limit used by the workloads in the policy group. If Unified Manager detects a change to a policy group, such as adding or removing a volume or changing the QoS limit, the actual and expected values might exceed 100% of the set limit. If a value exceeds 100% of the set limit it is displayed as >100%. If a value is less than 1% of the set limit it is displayed as <1%.
Actual	The measured performance value at a specific time for a given workload.
Deviation	The change between the expected values and the actual values. It is the ratio of the actual value minus the expected value to the upper value of the expected range minus the expected value. A negative deviation value indicates that workload performance is lower than expected, while a positive deviation value indicates that workload performance is higher than expected.
Expected	The expected values are based on the analysis of historical performance data for a given workload. Unified Manager analyzes these statistical values to determine the expected range (latency forecast) of values.
Latency Forecast (Expected Range)	The latency forecast is a prediction of what the upper and lower performance values are expected to be at a specific time. For the workload latency, the upper values form the performance threshold. When the actual value crosses the performance threshold, Unified Manager triggers a dynamic performance event.
Peak	The maximum value measured over a period of time.
Peak Deviation	The maximum deviation value measured over a period of time.
Queue Depth	The number of pending I/O requests that are waiting at the interconnect component.

Measurement	Description
Utilization	For the network processing, data processing, and aggregate components, the percentage of busy time to complete workload operations over a period of time. For example, the percentage of time for the network processing or data processing components to process an I/O request or for an aggregate to fulfill a read or write request.
Write Throughput	The amount of write throughput, in Megabytes per second (MB/s), from workloads on a local cluster to the partner cluster in a MetroCluster configuration.

Understand how Active IQ Unified Manager calculates the expected performance range

The latency forecast is a prediction of what the upper and lower performance values are expected to be at a specific time. For the workload latency, the upper values form the performance threshold. When the actual value crosses the performance threshold, Unified Manager triggers a dynamic performance event.

For example, during regular business hours between 9:00 a.m. and 5:00 p.m., most employees might check their email between 9:00 a.m. and 10:30 a.m. The increased demand on the email servers means an increase in workload activity on the back-end storage during this time. Employees might notice slow response time from their email clients.

During the lunch hour between 12:00 p.m. and 1:00 p.m. and at the end of the work day after 5:00 p.m., most employees are likely away from their computers. The demand on the email servers typically decreases, also decreasing the demand on back-end storage. Alternatively, there could be scheduled workload operations, such as storage backups or virus scanning, that start after 5:00 p.m. and increase activity on the back-end storage.

Over several days, the increase and decrease in workload activity determines the expected range (latency forecast) of activity, with upper and lower boundaries for a workload. When the actual workload activity for an object is outside the upper or lower boundaries, and remains outside the boundaries for a period of time, this might indicate that the object is being overused or underused.

How the latency forecast is formed

Unified Manager must collect a minimum of 3 days of workload activity before it can begin its analysis and before the latency forecast for I/O response time can be displayed in the GUI. The minimum required data collection does not account for all changes occurring from workload activity. After collecting the first 3 days of activity, Unified Manager adjusts the latency forecast every 24 hours at 12:00 a.m. to reflect workload activity changes and establish a more accurate dynamic performance threshold.

Daylight savings time (DST) changes the system time, which alters the latency forecast of performance statistics for monitored workloads. Unified Manager immediately begins to correct the latency forecast, which takes approximately 15 days to complete. During this time you can continue to use Unified Manager, but, since Unified Manager uses the latency forecast to detect dynamic events, some events might not be accurate. Events detected prior to the time change are not affected.

Use the latency forecast for performance analysis in Active IQ Unified Manager

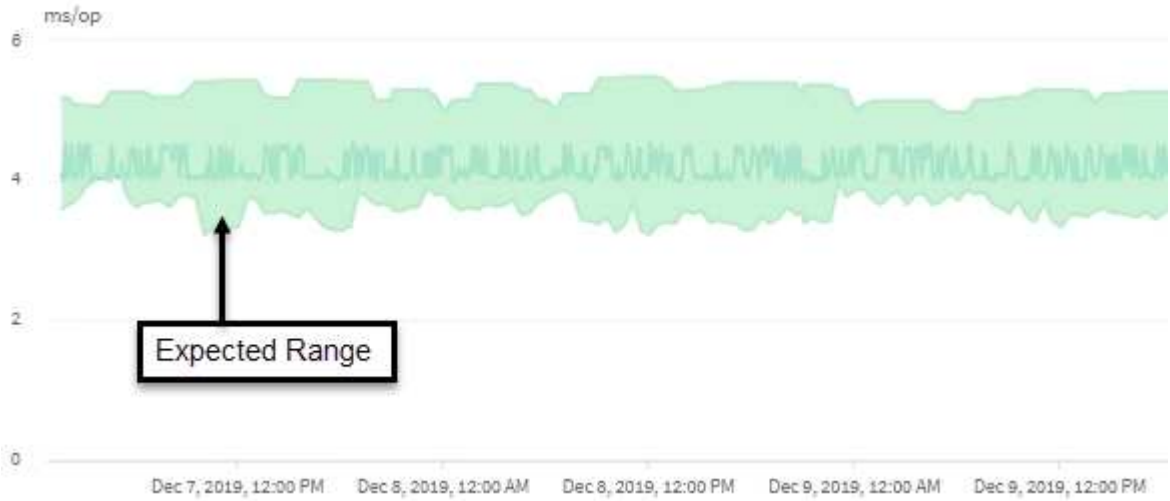
Unified Manager uses the latency forecast to represent the typical I/O latency (response time) activity for your monitored workloads. It alerts you when the actual latency for a workload is above the upper bounds of the latency forecast, which triggers a dynamic performance event, so that you can analyze the performance issue and take corrective action for resolving it.

The latency forecast sets the performance baseline for the workload. Over time, Unified Manager learns from past performance measurements to forecast the expected performance and activity levels for the workload. The upper boundary of the expected range establishes the dynamic performance threshold. Unified Manager uses the baseline to determine when the actual latency is above or below a threshold, or outside the bounds of their expected range. The comparison between the actual values and the expected values creates a performance profile for the workload.

When the actual latency for a workload exceeds the dynamic performance threshold, due to contention on a cluster component, the latency is high and the workload performs more slowly than expected. The performance of other workloads that share the same cluster components might also be slower than expected.

Unified Manager analyzes the threshold crossing event and determines whether the activity is a performance event. If the high workload activity remains consistent for a long period of time, such as several hours, Unified Manager considers the activity to be normal and dynamically adjusts the latency forecast to form the new dynamic performance threshold.

Some workloads might have consistently low activity, where the latency forecast for latency does not have a high rate of change over time. To minimize the number of events during analysis of performance events, Unified Manager triggers an event only for low-activity volumes whose operations and latencies are much higher than expected.



In this example, the latency for a volume has a latency forecast, in gray, of 3.5 milliseconds per operation (ms/op) at its lowest and 5.5 ms/op at its highest. If the actual latency, in blue, suddenly increases to 10 ms/op, due to an intermittent spike in network traffic or contention on a cluster component, it is then above the latency forecast and has exceeded the dynamic performance threshold.

When network traffic has decreased, or the cluster component is no longer in contention, the latency returns within the latency forecast. If the latency remains at or above 10 ms/op for a long period of time, you might need to take corrective action to resolve the event.

Use workload latency to identify performance issues in Active IQ Unified Manager

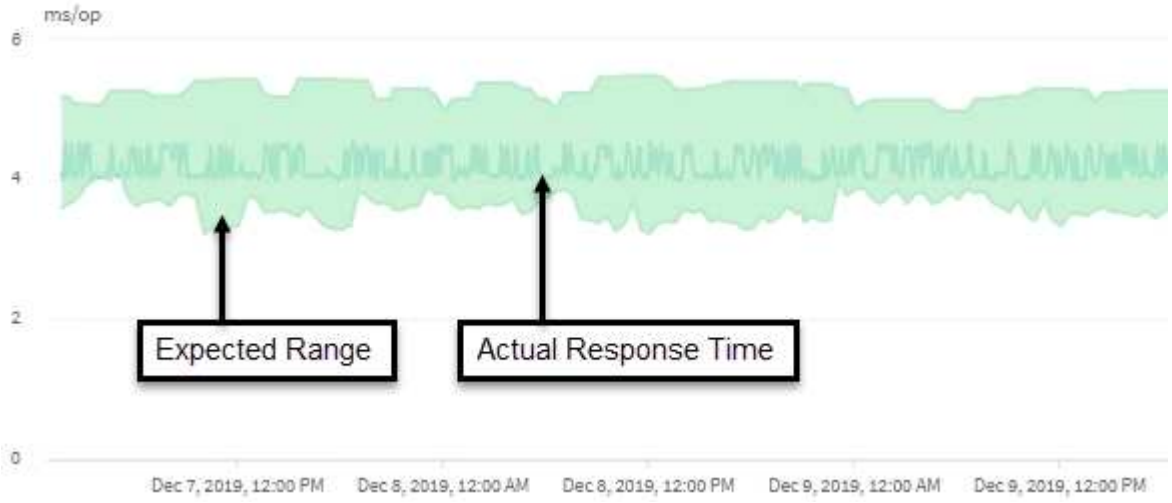
The workload latency (response time) is the time it takes for a volume on a cluster to respond to I/O requests from client applications. Unified Manager uses the latency to detect and alert you to performance events.

A high latency means that requests from applications to a volume on a cluster are taking longer than usual. The cause of the high latency could be on the cluster itself, due to contention on one or more cluster components. High latency could also be caused by issues outside of the cluster, such as network bottlenecks, issues with the client hosting the applications, or issues with the applications themselves.

Unified Manager only monitors the workload activity on the cluster. It does not monitor the applications, the clients, or the paths between the applications and the cluster.

Operations on the cluster, such as making backups or running deduplication, that increase their demand of cluster components shared by other workloads can also contribute to high latency. If the actual latency exceeds the dynamic performance threshold of the expected range (latency forecast), Unified Manager analyzes the event to determine whether it is a performance event that you might need to resolve. The latency is measured in milliseconds per operation (ms/op).

On the Latency Total chart in the Workload Analysis page, you can view an analysis of the latency statistics to see how the activity of individual processes, such as read and write requests, compares to the overall latency statistics. The comparison helps you determine which operations have the highest activity or whether specific operations have abnormal activity that is impacting the latency for a volume. When analyzing performance events, you can use the latency statistics to determine whether an event was caused by an issue on the cluster. You can also identify the specific workload activities or cluster components that are involved in the event.



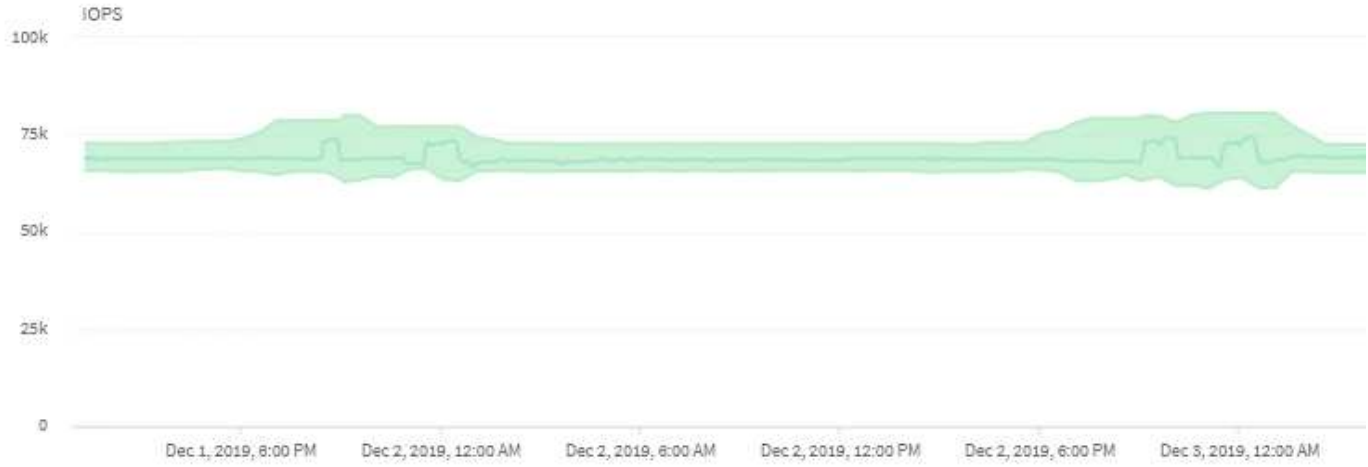
This example shows the Latency chart . The actual response time (latency) activity is a blue line and the latency forecast (expected range) is green.

There can be gaps in the blue line if Unified Manager was unable to gather data. This can occur because the cluster or volume was unreachable, Unified Manager was turned off during that time, or the collection was taking longer than the 5 minute collection period.

Assess how cluster operations can affect workload latency in Active IQ Unified Manager

Operations (IOPS) represent the activity of all user-defined and system-defined workloads on a cluster. The IOPS statistics help you determine whether cluster processes, such as making backups or running deduplication, are impacting workload latency (response time) or might have caused, or contributed to, a performance event.

When analyzing performance events, you can use the IOPS statistics to determine whether a performance event was caused by an issue on the cluster. You can identify the specific workload activities that might have been the main contributors to the performance event. IOPS are measured in operations per second (ops/sec).



This example shows the IOPS chart. The actual operations statistics is a blue line and the IOPS forecast of operations statistics is green.

In some cases where a cluster is overloaded, Unified Manager might display the message `Data collection is taking too long on Cluster` *`cluster_name`*. This means that not enough statistics have been collected for Unified Manager to analyze. You need to reduce the resources the cluster is using so that statistics can be collected.

Monitor MetroCluster performance with Active IQ Unified Manager

Unified Manager enables you to monitor the write throughput between clusters in a MetroCluster configuration to identify workloads with a high amount of write throughput.

If these high-performing workloads are causing other volumes on the local cluster to have high I/O response times, Unified Manager triggers performance events to notify you.

Unified Manager treats the clusters in a MetroCluster configuration as individual clusters. It does not distinguish between clusters that are partners or correlate the write throughput from each cluster.

When a local cluster in a MetroCluster configuration mirrors its data to its partner cluster, the data is written to NVRAM and then transferred over the interswitch links (ISLs) to the remote aggregates. Unified Manager analyzes the NVRAM to identify the workloads whose high write throughput is overutilizing the NVRAM, placing the NVRAM in contention.

Workloads whose deviation in response time has exceeded the performance threshold are called *victims* and workloads whose deviation in write throughput to the NVRAM is higher than usual, causing the contention, are called *bullies*. Because only the write requests are mirrored to the partner cluster, Unified Manager does not analyze read throughput.

You can view the throughput of any of the clusters in a MetroCluster configuration by analyzing the workloads of the corresponding LUNs and volumes from the following screens. You can filter the results by the cluster. From the left navigation pane:

- **Storage > Clusters > Performance: All Clusters** view. See
- **Storage > Volumes > Performance: All Volumes** view.
- **Storage > LUNs > Performance: All LUNs** view.
- **Workload Analysis > All workloads**

Related information

Performance event analysis and notification

Performance event analysis for a MetroCluster configuration

Roles of workloads involved in a performance event

Identifying victim workloads involved in a performance event

Identifying bully workloads involved in a performance event

Identifying shark workloads involved in a performance event

Understand performance events and alerts

Understand performance events and alerts in Active IQ Unified Manager

Performance events are incidents related to workload performance on a cluster. They help you identify workloads with slow response times. Together with health events that occurred at the same time, you can determine the issues that might have caused, or contributed to, the slow response times.

When Unified Manager detects multiple occurrences of the same event condition for the same cluster component, it treats all occurrences as a single event, not as separate events.

You can configure alerts to send email notification automatically when performance events of certain severity types occur.

Identify sources of performance events in Active IQ Unified Manager

Performance events are issues related to workload performance on a cluster. They help you identify storage objects with slow response times, also known as high latency. Together with other health events that occurred at the same time, you can determine the issues that might have caused, or contributed to, the slow response times.

Unified Manager receives performance events from the following sources:

- **User-defined performance threshold policy events**

  Performance issues based on custom threshold values that you have set. You configure performance threshold policies for storage objects; for example, aggregates and volumes, so that events are generated when a threshold value for a performance counter has been breached.

  You must define a performance threshold policy and assign it to a storage object to receive these events.

- **System-defined performance threshold policy events**

  Performance issues based on threshold values that are system-defined. These threshold policies are included with the installation of Unified Manager to cover common performance problems.

  These threshold policies are enabled by default, and you might see events shortly after adding a cluster.

- **Dynamic performance threshold events**

  Performance issues that are the result of failures or errors in an IT infrastructure, or from workloads overutilizing cluster resources. The cause of these events might be a simple issue that corrects itself over a period of time or that can be addressed with a repair or configuration change. A dynamic threshold event indicates that the workloads on an ONTAP system are slow due to other workloads with high usage of shared cluster components.

  These thresholds are enabled by default, and you might see events after three days of collecting data from a new cluster.

Prioritize performance events by severity in Active IQ Unified Manager

Each performance event is associated with a severity type to help you prioritize the events that require immediate corrective action.

- **Critical**

  A performance event occurred that might lead to service disruption if corrective action is not taken immediately.

  Critical events are sent from user-defined thresholds only.

- **Warning**

  A performance counter for a cluster object is out of normal range and should be monitored to make sure it does not reach the critical severity. Events of this severity do not cause service disruption, and immediate corrective action might not be required.

  Warning events are sent from user-defined, system-defined, or dynamic thresholds.

- **Information**

  The event occurs when a new object is discovered, or when a user action is performed. For example, when any storage object is deleted or when there are any configuration changes, the event with severity type Information is generated.

  Information events are sent directly from ONTAP when it detects a configuration change.

For more information, see the following links:

- What happens when an event is received
- What information is contained in an alert email
- Adding alerts
- Adding alerts for performance events

Monitor configuration changes detected by Active IQ Unified Manager

Unified Manager monitors your clusters for configuration changes to help you determine whether a change might have caused or contributed to a performance event. The Performance Explorer pages display a change event icon ( ) to indicate the date and time when the change was detected.

You can review the performance charts in the Performance Explorer pages and in the Workload Analysis page to see whether the change event impacted the performance of the selected cluster object. If the change was detected at or around the same time as a performance event, the change might have contributed to the issue, which caused the event alert to trigger.

Unified Manager can detect the following change events, which are categorized as Informational events:

- A volume moves between aggregates.

  Unified Manager can detect when the move is in progress, completed, or failed. If Unified Manager is down during a volume move, when it is back up it detects the volume move and displays a change event for it.

- The throughput (MB/s or IOPS) limit of a QoS policy group that contains one or more monitored workloads changes.

  Changing a policy group limit can cause intermittent spikes in the latency (response time), which might also trigger events for the policy group. The latency gradually returns to normal and any events caused by the spikes become obsolete.

- A node in an HA pair takes over or gives back the storage of its partner node.

  Unified Manager can detect when the takeover, partial takeover, or giveback operation has been completed. If the takeover is caused by a panicked node, Unified Manager does not detect the event.

- An ONTAP upgrade or revert operation is completed successfully.

  The previous version and new version are displayed.

Understand system-defined performance threshold policies in Active IQ Unified Manager

Unified Manager provides some standard threshold policies that monitor cluster performance and generate events automatically. These policies are enabled by default, and they generate warning or information events when the monitored performance thresholds are breached.

System-defined performance threshold policies are not enabled on Cloud Volumes ONTAP, ONTAP Edge, or ONTAP Select systems.

If you are receiving unnecessary events from any system-defined performance threshold policies, you can disable the events for individual policies from the Event Setup page.

Cluster threshold policies

The system-defined cluster performance threshold policies are assigned, by default, to every cluster being monitored by Unified Manager:

- **Cluster load imbalance**

  Identifies situations in which one node is operating at a much higher load than other nodes in the cluster, and therefore potentially affecting workload latencies.

  It does this by comparing the performance capacity used value for all nodes in a cluster to see if any node has exceeded the 30% threshold value for more than 24 hours. This is a warning event.

- **Cluster capacity imbalance**

  Identifies situations in which one aggregate has a much higher used capacity than other aggregates in the cluster, and therefore potentially affecting space required for operations.

  It does this by comparing the used capacity value for all aggregates in the cluster to see if there is a difference of 70% between any aggregates. This is a warning event.

Node threshold policies

The system-defined node performance threshold policies are assigned, by default, to every node in the clusters being monitored by Unified Manager:

- **Performance Capacity Used Threshold Breached**

  Identifies situations in which a single node is operating above the bounds of its operational efficiency, and therefore potentially affecting workload latencies.

  It does this by looking for nodes that are using more than 100% of their performance capacity for more than 12 hours. This is a warning event.

- **Node HA pair over-utilized**

  Identifies situations in which nodes in an HA pair are operating above the bounds of the HA pair operational efficiency.

  It does this by looking at the performance capacity used value for the two nodes in the HA pair. If the combined performance capacity used of the two nodes exceeds 200% for more than 12 hours, then a controller failover will impact workload latencies. This is an informational event.

- **Node disk fragmentation**

  Identifies situations in which a disk or disks in an aggregate are fragmented, slowing key system services and potentially affecting workload latencies on a node.

  It does this by looking at certain read and write operation ratios across all aggregates on a node. This policy might also be triggered during SyncMirror resynchronization or when errors are found during disk scrub operations. This is a warning event.

The “Node disk fragmentation” policy analyzes HDD-only aggregates; Flash Pool, SSD, and FabricPool aggregates are not analyzed.

Aggregate threshold policies

The system-defined aggregate performance threshold policy is assigned by default to every aggregate in the clusters being monitored by Unified Manager:

- **Aggregate disks over-utilized**

  Identifies situations in which an aggregate is operating above the limits of its operational efficiency, thereby potentially affecting workload latencies. It identifies these situations by looking for aggregates where the disks in the aggregate are more than 95% utilized for more than 30 minutes. This multicondition policy then performs the following analysis to help determine the cause of the issue:

  - Is a disk in the aggregate currently undergoing background maintenance activity?

    Some of the background maintenance activities a disk could be undergoing are disk reconstruction, disk scrub, SyncMirror resynchronization, and reparity.

  - Is there a communications bottleneck in the disk shelf Fibre Channel interconnect?
  - Is there too little free space in the aggregate?
    A warning event is issued for this policy only if one (or more) of the three subordinate policies are also considered breached. A performance event is not triggered if only the disks in the aggregate are more than 95% utilized.

The “Aggregate disks over-utilized” policy analyzes HDD-only aggregates and Flash Pool (hybrid) aggregates; SSD and FabricPool aggregates are not analyzed.

Workload latency threshold policies

The system-defined workload latency threshold policies are assigned to any workload that has a configured Performance Service Level policy that has a defined “expected latency” value:

- **Workload Volume/LUN Latency Threshold Breached as defined by Performance Service Level**

  Identifies volumes (file shares) and LUNs that have exceeded their “expected latency” limit, and that are affecting workload performance. This is a warning event.

  It does this by looking for workloads that have exceeded the expected latency value for 30% of the time during the previous hour.

QoS threshold policies

The system-defined QoS performance threshold policies are assigned to any workload that has a configured ONTAP QoS maximum throughput policy (IOPS, IOPS/TB, or MB/s). Unified Manager triggers an event when the workload throughput value is 15% less than the configured QoS value:

- **QoS Max IOPS or MB/s threshold**

  Identifies volumes and LUNs that have exceeded their QoS maximum IOPS or MB/s throughput limit, and that are affecting workload latency. This is a warning event.

When a single workload is assigned to a policy group, it does this by looking for workloads that have exceeded the maximum throughput threshold defined in the assigned QoS policy group during each collection period for the previous hour.

When multiple workloads share a single QoS policy, it does this by adding the IOPS or MB/s of all workloads in the policy and checking that total against the threshold.

- **QoS Peak IOPS/TB or IOPS/TB with Block Size threshold**

  Identifies volumes that have exceeded their adaptive QoS peak IOPS/TB throughput limit (or IOPS/TB with Block Size limit), and that are affecting workload latency. This is a warning event.

  It does this by converting the peak IOPS/TB threshold defined in the adaptive QoS policy into a QoS maximum IOPS value based on the size of each volume, and then it looks for volumes that have exceeded the QoS max IOPS during each performance collection period for the previous hour.

  This policy is applied to volumes only when the cluster is installed with ONTAP 9.3 and later software.

  When the “block size” element has been defined in the adaptive QoS policy, the threshold is converted into a QoS maximum MB/s value based on the size of each volume. Then it looks for volumes that have exceeded the QoS max MB/s during each performance collection period for the previous hour.

  This policy is applied to volumes only when the cluster is installed with ONTAP 9.5 and later software.

Analyze performance events and notifications in Active IQ Unified Manager

Performance events notify you about I/O performance issues on a workload caused by contention on a cluster component. Unified Manager analyzes the event to identify all workloads involved, the component in contention, and whether the event is still an issue that you might need to resolve.

Unified Manager monitors the I/O latency (response time) and IOPS (operations) for volumes on a cluster. When other workloads overuse a cluster component, for example, the component is in contention and cannot perform at an optimal level to meet workload demands. The performance of other workloads that are using the same component might be impacted, causing their latencies to increase. If the latency crosses the dynamic performance threshold, Unified Manager triggers a performance event to notify you.

Event analysis

Unified Manager performs the following analyses, using the previous 15 days of performance statistics, to identify the victim workloads, bully workloads, and the cluster component involved in an event:

- Identifies victim workloads whose latency has crossed the dynamic performance threshold, which is the upper boundary of the latency forecast:
  - For volumes on HDD or Flash Pool hybrid aggregates (local tier), events are triggered only when the latency is greater than 5 milliseconds (ms) and the IOPS are more than 10 operations per second (ops/sec).
  - For volumes on all-SSD aggregates or FabricPool aggregates (cloud tier), events are triggered only when the latency is greater than 1 ms and the IOPS are more than 100 ops/sec.

- Identifies the cluster component in contention.

  If the latency of victim workloads at the cluster interconnect is greater than 1 ms, Unified Manager treats this as significant and triggers an event for the cluster interconnect.

- Identifies the bully workloads that are overusing the cluster component and causing it to be in contention.
- Ranks the workloads involved, based on their deviation in utilization or activity of a cluster component, to determine which bullies have the highest change in usage of the cluster component and which victims are the most impacted.

An event might occur for only a brief moment and then correct itself after the component it is using is no longer in contention. A continuous event is one that reoccurs for the same cluster component within a five-minute interval and remains in the active state. For continuous events, Unified Manager triggers an alert after detecting the same event during two consecutive analysis intervals.

When an event is resolved, it remains available in Unified Manager as part of the record of past performance issues for a volume. Each event has a unique ID that identifies the event type and the volumes, cluster, and cluster components involved.

A single volume can be involved in more than one event at the same time.

Event state

Events can be in one of the following states:

- **Active**

  Indicates that the performance event is currently active (new or acknowledged). The issue causing the event has not corrected itself or has not been resolved. The performance counter for the storage object remains above the performance threshold.

- **Obsolete**

  Indicates that the event is no longer active. The issue causing the event has corrected itself or has been resolved. The performance counter for the storage object is no longer above the performance threshold.

Event notification

The events are displayed on the Dashboard page and on many other pages in the user interface, and alerts for those events are sent to specified email addresses. You can view detailed analysis information about an event and get suggestions for resolving it on the Event details page and on the Workload Analysis page.

Event interaction

On the Event details page and on the Workload Analysis page, you can interact with events in the following ways:

- Moving the mouse over an event displays a message that shows the date and time when the event was detected.

  If there are multiple events for the same time period, the message shows the number of events.

- Clicking a single event displays a dialog box that shows more detailed information about the event,

including the cluster components that are involved.

The component in contention is circled and highlighted red. You can click **View full analysis** to view the full analysis on the Event details page. If there are multiple events for the same time period, the dialog box shows details about the three most recent events. You can click an event to view the event analysis on the Event details page.

Understand how Active IQ Unified Manager determines performance impact for an event

Unified Manager uses the deviation in activity, utilization, write throughput, cluster component usage, or I/O latency (response time) for a workload to determine the level of impact to workload performance. This information determines the role of each workload in the event and how they are ranked on the Event details page.

Unified Manager compares the last analyzed values for a workload to the expected range (latency forecast) of values. The difference between the values last analyzed and the expected range of values identifies the workloads whose performance was most impacted by the event.

For example, suppose a cluster contains two workloads: Workload A and Workload B. The latency forecast for Workload A is 5-10 milliseconds per operation (ms/op) and its actual latency is usually around 7 ms/op. The latency forecast for Workload B is 10-20 ms/op and its actual latency is usually around 15 ms/op. Both workloads are well within their latency forecast. Due to contention on the cluster, the latency of both workloads increases to 40 ms/op, crossing the dynamic performance threshold, which is the upper bounds of the latency forecast, and triggering events. The deviation in latency, from the expected values to the values above the performance threshold, for Workload A is around 33 ms/op, and the deviation for Workload B is around 25 ms/op. The latency of both workloads spike to 40 ms/op, but Workload A had the bigger performance impact because it had the higher latency deviation at 33 ms/op.

On the Event details page, in the System Diagnosis section, you can sort workloads by their deviation in activity, utilization, or throughput for a cluster component. You can also sort workloads by latency. When you select a sort option, Unified Manager analyzes the deviation in activity, utilization, throughput, or latency since the event was detected from the expected values to determine the workload sort order. For the latency, the red dots (🔴) indicate a performance threshold crossing by a victim workload, and the subsequent impact to the latency. Each red dot indicates a higher level of deviation in latency, which helps you identify the victim workloads whose latency was impacted the most by an event.

Understand cluster components and contention in Active IQ Unified Manager

You can identify cluster performance issues when a cluster component goes into contention. The performance of workloads that use the component slow down and their response time (latency) for client requests increases, which triggers an event in Unified Manager.

A component that is in contention cannot perform at an optimal level. Its performance has declined, and the performance of other cluster components and workloads, called *victims*, might have increased latency. To bring a component out of contention, you must reduce its workload or increase its ability to handle more work, so that the performance can return to normal levels. Because Unified Manager collects and analyzes workload performance in five-minute intervals, it detects only when a cluster component is consistently overused. Transient spikes of overusage that last for only a short duration within the five-minute interval are not detected.

For example, a storage aggregate might be under contention because one or more workloads on it are

competing for their I/O requests to be fulfilled. Other workloads on the aggregate can be impacted, causing their performance to decrease. To reduce the amount of activity on the aggregate, there are different steps you can take, such as moving one or more workloads to a less busy aggregate or node, to lessen the overall workload demand on the current aggregate. For a QoS policy group, you can adjust the throughput limit, or move workloads to a different policy group, so that the workloads are no longer being throttled.

Unified Manager monitors the following cluster components to alert you when they are in contention:

- **Network**

  Represents the wait time of I/O requests by the external networking protocols on the cluster. The wait time is time spent waiting for “transfer ready” transactions to finish before the cluster can respond to an I/O request. If the network component is in contention, it means high wait time at the protocol layer is impacting the latency of one or more workloads.

- **Network Processing**

  Represents the software component in the cluster involved with I/O processing between the protocol layer and the cluster. The node handling network processing might have changed since the event was detected. If the network processing component is in contention, it means high utilization at the network processing node is impacting the latency of one or more workloads.

  When using an All SAN Array cluster in an active-active configuration, the network processing latency value is displayed for both nodes so you can verify the nodes are sharing the load equally.

- **QoS Limit Max**

  Represents the throughput maximum (peak) setting of the storage Quality of Service (QoS) policy group assigned to the workload. If the policy group component is in contention, it means all workloads in the policy group are being throttled by the set throughput limit, which is impacting the latency of one or more of those workloads.

- **QoS Limit Min**

  Represents the latency to a workload that is being caused by QoS throughput minimum (expected) setting assigned to other workloads. If the QoS minimum set on certain workloads use the majority of the bandwidth to guarantee the promised throughput, other workloads will be throttled and see more latency.

- **Cluster Interconnect**

  Represents the cables and adapters with which clustered nodes are physically connected. If the cluster interconnect component is in contention, it means high wait time for I/O requests at the cluster interconnect is impacting the latency of one or more workloads.

- **Data Processing**

  Represents the software component in the cluster involved with I/O processing between the cluster and the storage aggregate that contains the workload. The node handling data processing might have changed since the event was detected. If the data processing component is in contention, it means high utilization at the data processing node is impacting the latency of one or more workloads.

- **Volume Activation**

  Represents the process that tracks the usage of all active volumes. In large environments where more than 1000 volumes are active, this process tracks how many critical volumes need to access resources through

the node at the same time. When the number of concurrent active volumes exceeds the recommended maximum threshold, some of the non-critical volumes will experience latency as identified here.

- **MetroCluster Resources**

  Represents the MetroCluster resources, including NVRAM and interswitch links (ISLs), used to mirror data between clusters in a MetroCluster configuration. If the MetroCluster component is in contention, it means high write throughput from workloads on the local cluster or a link health issue is impacting the latency of one or more workloads on the local cluster. If the cluster is not in a MetroCluster configuration, this icon is not displayed.

- **Aggregate or SSD Aggregate Ops**

  Represents the storage aggregate on which the workloads are running. If the aggregate component is in contention, it means high utilization on the aggregate is impacting the latency of one or more workloads. An aggregate consists of all HDDs, or a mix of HDDs and SSDs (a Flash Pool aggregate), or a mix of HDDs and a cloud tier (a FabricPool aggregate). An “SSD Aggregate” consists of all SSDs (an all-flash aggregate), or a mix of SSDs and a cloud tier (a FabricPool aggregate).

- **Cloud Latency**

  Represents the software component in the cluster involved with I/O processing between the cluster and the cloud tier on which user data is stored. If the cloud latency component is in contention, it means that a large amount of reads from volumes that are hosted on the cloud tier are impacting the latency of one or more workloads.

- **Sync SnapMirror**

  Represents the software component in the cluster involved with replicating user data from the primary volume to the secondary volume in a SnapMirror Synchronous relationship. If the sync SnapMirror component is in contention, it means that the activity from SnapMirror Synchronous operations are impacting the latency of one or more workloads.

Identify workload roles in performance events in Active IQ Unified Manager

Unified Manager uses roles to identify the involvement of a workload in a performance event. The roles include victims, bullies, and sharks. A user-defined workload can be a victim, bully, and shark at the same time.

Role	Description
Victim	A user-defined workload whose performance has decreased due to other workloads, called bullies, that are over-using a cluster component. Only user-defined workloads are identified as victims. Unified Manager identifies victim workloads based on their deviation in latency, where the actual latency, during an event, has greatly increased from its latency forecast (expected range).

Role	Description
Bully	A user-defined or system-defined workload whose over-use of a cluster component has caused the performance of other workloads, called victims, to decrease. Unified Manager identifies bully workloads based on their deviation in usage of a cluster component, where the actual usage, during an event, has greatly increased from its expected range of usage.
Shark	A user-defined workload with the highest usage of a cluster component compared to all workloads involved in an event. Unified Manager identifies shark workloads based on their usage of a cluster component during an event.

Workloads on a cluster can share many of the cluster components, such as aggregates and the CPU for network and data processing. When a workload, such as a volume, increases its usage of a cluster component to the point that the component cannot efficiently meet workload demands, the component is in contention. The workload that is over-using a cluster component is a bully. The other workloads that share those components, and whose performance is impacted by the bully, are the victims. Activity from system-defined workloads, such as deduplication or Snapshot copies, can also escalate into “bullying”.

When Unified Manager detects an event, it identifies all workloads and cluster components involved, including the bully workloads that caused the event, the cluster component that is in contention, and the victim workloads whose performance has decreased due to the increased activity of bully workloads.

If Unified Manager cannot identify the bully workloads, it only alerts on the victim workloads and the cluster component involved.

Unified Manager can identify workloads that are victims of bully workloads, and also identify when those same workloads become bully workloads. A workload can be a bully to itself. For example, a high-performing workload that is being throttled by a policy group limit causes all workloads in the policy group to be throttled, including itself. A workload that is a bully or a victim in an ongoing performance event might change its role or no longer be a participant in the event.

Manage performance thresholds

Manage performance thresholds in Active IQ Unified Manager

Performance threshold policies enable you to determine the point at which Unified Manager generates an event to inform system administrators about issues that could be impacting workload performance. These threshold policies are known as *user-defined* performance thresholds.

This release supports user-defined, system-defined, and dynamic performance thresholds. With dynamic and system-defined performance thresholds, Unified Manager analyzes the workload activity to determine the appropriate threshold value. With user-defined thresholds, you can define the upper performance limits for many performance counters and for many storage objects.

System-defined performance thresholds and dynamic performance thresholds are set by Unified

Manager and are not configurable. If you are receiving unnecessary events from any system-defined performance threshold policies, you can disable individual policies from the Event Setup page.

Use user-defined performance threshold policies in Active IQ Unified Manager

You set performance threshold policies on storage objects (for example, on aggregates and volumes) so that an event can be sent to the storage administrator to inform the administrator that the cluster is experiencing a performance issue.

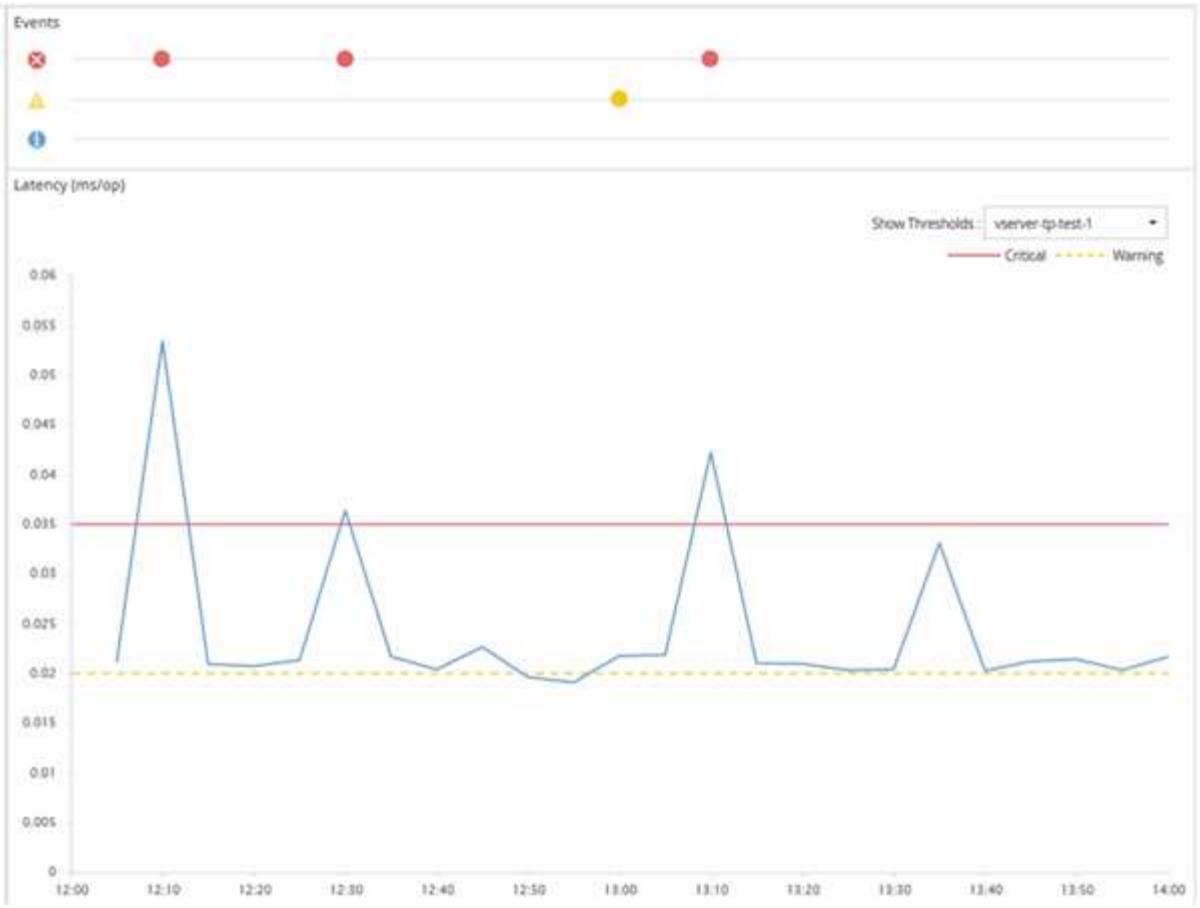
You create a performance threshold policy for a storage object by:

- Selecting a storage object
- Selecting a performance counter associated with that object
- Specifying values that define the performance counter upper limits that are considered warning and critical situations
- Specifying a time period that defines how long the counter must exceed the upper limit

For example, you can set a performance threshold policy on a volume so that you receive a critical event notification whenever IOPS for that volume exceeds 750 operations per second for 10 consecutive minutes. This same threshold policy can also specify that a warning event be sent when IOPS exceeds 500 operations per second for 10 minutes.

The current release provides thresholds that send events when a counter value exceeds the threshold setting. You cannot set thresholds that send events when a counter value falls below a threshold setting.

An example counter chart is shown here, indicating that a warning threshold (yellow icon) was breached at 1:00, and that a critical threshold (red icon) was breached at 12:10, 12:30, and 1:10:



A threshold breach must occur continuously for the specified duration. If the threshold dips below the limit values for any reason, a subsequent breach is considered the start of a new duration.

Some cluster objects and performance counters enable you to create a combination threshold policy that requires two performance counters to exceed their maximum limits before an event is generated. For example, you can create a threshold policy using the following criteria:

Cluster object	Performance counter	Warning threshold	Critical threshold	Duration
Volume	Latency	10 milliseconds	20 milliseconds	15 minutes
Aggregate	Utilization	65%	85%	

Threshold policies that use two cluster objects cause an event to be generated only when both conditions are breached. For example, using the threshold policy defined in the table:

If volume latency is averaging…	And aggregate disk utilization is…	Then…
15 milliseconds	50%	No event is reported.
15 milliseconds	75%	A Warning event is reported.
25 milliseconds	75%	A Warning event is reported.
25 milliseconds	90%	A Critical event is reported.

Respond to breached performance threshold policies in Active IQ Unified Manager

When a counter value exceeds its defined performance threshold value for the amount of time specified in the duration, the threshold is breached and an event is reported.

The event causes the following actions to be initiated:

- The event is displayed in the Dashboard, the Performance Cluster Summary page, the Events page, and the object-specific Performance Inventory page.
- (optional) An email alert about the event can be sent to one or more email recipients, and an SNMP trap can be sent to a trap receiver.
- (optional) A script can be executed to automatically modify or update storage objects.

The first action is always executed. You configure whether the optional actions are performed in the Alert Setup page. You can define unique actions depending on whether a Warning or a Critical threshold policy is breached.

After a performance threshold policy breach has occurred on a storage object, no further events are generated for that policy until the counter value goes below the threshold value, at which point the duration resets for that limit. While the threshold continues to be exceeded, the end time of the event is continually updated to reflect that this event is ongoing.

A threshold event captures, or freezes, the information related to severity and policy definition so that unique threshold information displays with the event, even if the threshold policy is modified in the future.

Understand performance counters tracked using thresholds in Active IQ Unified Manager

Some common performance counters, such as IOPS and MB/s, can have thresholds set for all storage objects. There are other counters that can have thresholds set for only certain storage objects.

Available performance counters

Storage object	Performance counter	Description
Cluster	IOPS	Average number of input/output operations the cluster processes per second.
	MB/s	Average number of megabytes of data transferred to and from this cluster per second.

Storage object	Performance counter	Description
Node	IOPS	Average number of input/output operations the node processes per second.
	MB/s	Average number of megabytes of data transferred to and from this node per second.
	Latency	Average number of milliseconds the node takes to respond to application requests.
	Utilization	Average percentage of the node's CPU and RAM that is being used.
	Performance Capacity Used	Average percentage of performance capacity that is being consumed by the node.
	Performance Capacity Used - Takeover	Average percentage of performance capacity that is being consumed by the node, plus the performance capacity of its partner node.
Aggregate	IOPS	Average number of input/output operations the aggregate processes per second.
	MB/s	Average number of megabytes of data transferred to and from this aggregate per second.
	Latency	Average number of milliseconds the aggregate takes to respond to application requests.
	Utilization	Average percentage of the aggregate's disks that are being used.
	Performance Capacity Used	Average percentage of performance capacity that is being consumed by the aggregate.
Storage VM	IOPS	Average number of input/output operations the SVM processes per second.
	MB/s	Average number of megabytes of data transferred to and from this SVM per second.
	Latency	Average number of milliseconds the SVM takes to respond to application requests.

Storage object	Performance counter	Description
Volume	IOPS	Average number of input/output operations the volume processes per second.
	MB/s	Average number of megabytes of data transferred to and from this volume per second.
	Latency	Average number of milliseconds the volume takes to respond to application requests.
	Cache miss ratio	Average percentage of read requests from client applications that are returned from the volume instead of being returned from cache.
LUN	IOPS	Average number of input/output operations the LUN processes per second.
	MB/s	Average number of megabytes of data transferred to and from this LUN per second.
	Latency	Average number of milliseconds the LUN takes to respond to application requests.
Namespace	IOPS	Average number of input/output operations the namespace processes per second.
	MB/s	Average number of megabytes of data transferred to and from this namespace per second.
	Latency	Average number of milliseconds the namespace takes to respond to application requests.
Port	Bandwidth utilization	Average percentage of the port's available bandwidth that is being used.
	MB/s	Average number of megabytes of data transferred to and from this port per second.
Network Interface (LIF)	MB/s	Average number of megabytes of data transferred to and from this LIF per second.

Understand combination threshold policy objects and counters in Active IQ Unified Manager

Only some performance counters can be used together in combination policies. When

primary and secondary performance counters are specified, both performance counters must exceed their maximum limits before an event is generated.

Primary storage object and counter	Secondary storage object and counter
Volume Latency	Volume IOPS
	Volume MB/s
	Aggregate Utilization
	Aggregate Performance Capacity Used
	Node Utilization
	Node Performance Capacity Used
	Node Performance Capacity Used - Takeover
LUN Latency	LUN IOPS
	LUN MB/s
	Aggregate Utilization
	Aggregate Performance Capacity Used
	Node Utilization
	Node Performance Capacity Used
	Node Performance Capacity Used - Takeover

When a volume combination policy is applied to a FlexGroup volume, instead of to a FlexVol volume, only the “Volume IOPS” and “Volume MB/s” attributes can be selected as the secondary counter. If the threshold policy contains one of the node or aggregate attributes, then the policy will not be applied to the FlexGroup volume, and you will receive an error message describing this case. This is because FlexGroup volumes can exist on more than one node or aggregate.

Create user-defined performance threshold policies in Active IQ Unified Manager

You create performance threshold policies for storage objects so that notifications are sent when a performance counter exceeds a specific value. The event notification identifies that the cluster is experiencing a performance issue.

Before you begin

You must have the Application Administrator role.

You create performance threshold policies by entering the threshold values on the Create Performance Threshold Policy page. You can create new policies by defining all the policy values in this page, or you can make a copy of an existing policy and change the values in the copy (called *cloning*).

Valid threshold values are 0.001 through 10,000,000 for numbers, 0.001-100 for percentages, and 0.001-200 for Performance Capacity Used percentages.

The current release provides thresholds that send events when a counter value exceeds the threshold setting. You cannot set thresholds that send events when a counter value falls below a threshold setting.

Steps

1. In the left navigation pane, select **Event Thresholds** > **Performance**.

   The Performance Thresholds page is displayed.

2. Click the appropriate button depending on whether you want to build a new policy or if you want to clone a similar policy and modify the cloned version.

To…	Click…
Create a new policy	**Create**
Clone an existing policy	Select an existing policy and click **Clone**

   The Create Performance Threshold Policy page or Clone Performance Threshold Policy page is displayed.

3. Define the threshold policy by specifying the performance counter threshold values you want to set for specific storage objects:

   a. Select the storage object type and specify a name and description for the policy.

   b. Select the performance counter to be tracked and specify the limit values that define Warning and Critical events.

      You must define at least one Warning or one Critical limit. You do not need to define both types of limits.

   c. Select a secondary performance counter, if required, and specify the limit values for Warning and Critical events.

      Including a secondary counter requires that both counters exceed the limit values before the threshold is breached and an event is reported. Only certain objects and counters can be configured using a combination policy.

   d. Select the duration of time for which the limit values must be breached for an event to be sent.

   When cloning an existing policy, you must enter a new name for the policy.

4. Click **Save** to save the policy.

   You are returned to the Performance Thresholds page. A success message at the top of the page confirms that the threshold policy was created and provides a link to the Inventory page for that object type so that you can apply the new policy to storage objects immediately.

If you want to apply the new threshold policy to storage objects at this time, you can click the **Go to object_type now** link to go to the Inventory page.

Assign performance threshold policies to storage objects in Active IQ Unified Manager

You assign a user-defined performance threshold policy to a storage object so that Unified Manager reports an event if the value of the performance counter exceeds the policy setting.

Before you begin

You must have the Application Administrator role.

The performance threshold policy, or policies, that you want to apply to the object must exist.

You can apply only one performance policy at a time to an object, or to a group of objects.

You can assign a maximum of three threshold policies to each storage object. When assigning policies to multiple objects, if any of the objects already has the maximum number of policies assigned, Unified Manager performs the following actions:

- Applies the policy to all of the selected objects that have not reached their maximum
- Ignores the objects that have reached the maximum number of policies
- Displays a message that the policy was not assigned to all objects

Steps

1. From the Performance inventory page of any storage object, select the object or objects to which you want to assign a threshold policy:

To assign thresholds to…	Click…
A single object	The check box at the left of that object.
Multiple objects	The check box at the left of each object.
All objects on the page	The ☐▾ drop-down box, and choose **Select all objects on this page**.
All objects of the same type	The ☐▾ drop-down box, and choose **Select all objects**.

   You can use the sorting and filtering functionality to refine the list of objects on the inventory page to make it easier to apply threshold policies to many objects.

2. Make your selection, and then click **Assign Performance Threshold Policy**.

   The Assign Performance Threshold Policy page is displayed, showing a list of threshold policies that exist for that specific type of storage object.

3. Click each policy to display the details of the performance threshold settings to verify that you have selected the correct threshold policy.
4. After you have selected the appropriate threshold policy, click **Assign Policy**.

   A success message at the top of the page confirms that the threshold policy was assigned to the object or objects, and provides a link to the Alerting page so that you can configure alert settings for this object and policy.

If you want to have alerts sent over email, or as an SNMP trap, to notify you that a particular performance event has been generated, you must configure the alert settings in the Alert Setup page.

View performance threshold policies in Active IQ Unified Manager

You can view all of the currently defined performance threshold policies from the Performance Thresholds page.

The list of threshold policies is sorted alphabetically by policy name, and it includes policies for all types of storage objects. You can click a column header to sort the policies by that column. If you are looking for a specific policy, use the filter and search mechanisms to refine the list of threshold policies that appear in the inventory list.

You can hover your cursor over the Policy Name and the Condition name to see the configuration details of the policy. Additionally, you can use the provided buttons to create, clone, edit, and delete user-defined threshold policies.

Step

1. In the left navigation pane, select **Event Thresholds** > **Performance**.

   The Performance Thresholds page is displayed.

Edit user-defined performance threshold policies in Active IQ Unified Manager

You can edit the threshold settings for existing performance threshold policies. This can be useful if you find that you are receiving too many or too few alerts for certain threshold conditions.

Before you begin

You must have the Application Administrator role.

You cannot change the policy name or the type of storage object that is being monitored for existing threshold policies.

Steps

1. In the left navigation pane, select **Event Thresholds** > **Performance**.

   The Performance Thresholds page displays.

2. Select the threshold policy that you want to change and click **Edit**.

   The Edit Performance Threshold Policy page is displayed.

3. Make your changes to the threshold policy and click **Save**.

   You are returned to the Performance Thresholds page.

After they are saved, changes are updated immediately on all storage objects that use the policy.

Depending on the type of changes that you made to the policy, you may want to review the alert settings configured for the objects that use the policy in the Alert Setup page.

Remove performance threshold policies from storage objects in Active IQ Unified Manager

You can remove a user-defined performance threshold policy from a storage object when you no longer want Unified Manager to monitor the value of the performance counter.

Before you begin

You must have the Application Administrator role.

You can remove only one policy at a time from a selected object.

You can remove a threshold policy from multiple storage objects by selecting more than one object in the list.

Steps

1. From the **inventory** page of any storage object, select one or more objects that have at least one performance threshold policy applied.

To clear thresholds from…	Do this…
A single object	Select the check box at the left of that object.
Multiple objects	Select the check box at the left of each object.
All objects on the page	Click  in the column header.

2. Click **Clear Performance Threshold Policy**.

   The Clear Threshold Policy page displays, showing a list of threshold policies that are currently assigned to the storage objects.

3. Select the threshold policy you want to remove from the objects and click **Clear Policy**.

   When you select a threshold policy, the details of the policy display so that you can confirm that you have selected the appropriate policy.

Understand performance threshold policy changes in Active IQ Unified Manager

If you adjust the counter value or duration of an existing performance threshold policy, the policy change is applied to all storage objects that use the policy. The new setting takes place immediately, and Unified Manager begins to compare performance counter values to the new threshold settings for all newly collected performance data.

If any active events exist for objects that are using the changed threshold policy, the events are marked as obsolete, and the threshold policy begins monitoring the counter as a newly defined threshold policy.

When viewing the counter on which the threshold has been applied in the Counter Charts Detailed View, the critical and warning threshold lines reflect the current threshold settings. The original threshold settings do not appear on this page even if you view historical data when the old threshold setting was in effect.

Because older threshold settings do not appear in the Counter Charts Detailed View, you might see historical events that appear below the current threshold lines.

Manage threshold policies when moving objects in Active IQ Unified Manager

Because performance threshold policies are assigned to storage objects, if you move an object, all assigned threshold policies remain attached to the object after the move is completed. For example, if you move a volume or LUN to a different aggregate, the threshold policies are still active for the volume or LUN on the new aggregate.

If a secondary counter condition exists for the threshold policy (a combination policy)--for example, if an additional condition is assigned to an aggregate or a node—the secondary counter condition is applied to the new aggregate or node to which the volume or LUN has been moved.

If any new active events exist for objects that are using the changed threshold policy, the events are marked as obsolete, and the threshold policy begins monitoring the counter as a newly defined threshold policy.

A volume move operation causes ONTAP to send an informational change event. A change event icon appears in the Events timeline on the Performance Explorer page and the Workload Analysis page to indicate the time when the move operation was completed.

If you move an object to a different cluster, the user-defined threshold policy is removed from the object. If required, you must assign a threshold policy to the object after the move operation is completed. Dynamic and system-defined threshold policies, however, are applied automatically to an object after it has moved to a new cluster.

Threshold policy functionality during HA takeover and giveback

When a takeover or giveback operation occurs in a high-availability (HA) configuration, objects that are moved from one node to the other node retain their threshold policies in the same manner as in the manual move operations. Because Unified Manager checks for cluster configuration changes every 15 minutes, the impact of the switchover to the new node is not identified until the next poll of the cluster configuration.

If both a takeover and giveback operation occur within the 15-minute configuration change collection period, you might not see the performance statistics move from one node to the other node.

Threshold policy functionality during aggregate relocation

If you move an aggregate from one node to another node using the `aggregate relocation start` command, both single and combination threshold policies are retained on all objects, and the node portion of the threshold policy is applied to the new node.

Threshold policy functionality during MetroCluster switchover

Objects that move from one cluster to another cluster in a MetroCluster configuration do not retain their user-defined threshold policy settings. If required, you can apply threshold policies on the volumes and LUNs that have moved to the partner cluster. After an object has moved back to its original cluster, the user-defined threshold policy is reapplied automatically.

For information, see Volume behavior during switchover and switchback.

Analyze performance events

Analyze performance events in Active IQ Unified Manager

You can analyze performance events to identify when they were detected, whether they are active (new or acknowledged) or obsolete, the workloads and cluster components involved, and the options for resolving the events on your own.

Display information about performance events in Active IQ Unified Manager

You can use the Event Management inventory page to view a list of all the performance events on the clusters being monitored by Unified Manager. By viewing this information you can determine the most critical events and then drill down to detailed information to

determine the cause of the event.

Before you begin

- You must have the Operator, Application Administrator, or Storage Administrator role.

The list of events is sorted by detected time, with the most recent events listed first. You can click a column header to sort the events based on that column. For example, you can sort by the Status column to view events by severity. If you are looking for a specific event, or for a specific type of event, you can use the filter and search mechanisms to refine the list of events that appear in the list.

Events from all sources are displayed on this page:

- User-defined performance threshold policy
- System-defined performance threshold policy
- Dynamic performance threshold

The Event Type column lists the source of the event. You can select an event to view details about the event in the Event details page.

Steps

1. In the left navigation pane, click **Event Management**.
2. From the View menu, select **Active performance events**.

   The page displays all New and Acknowledged Performance events that have been generated in the past 7 days.

3. Locate an event that you want to analyze and click the event name.

   The details page for the event displays.

> You can also display the details page for an event by clicking the event name link from the Performance Explorer page and from an alert email.

Analyze events from user-defined performance thresholds

Analyze events from user-defined performance thresholds in Active IQ Unified Manager

Events generated from user-defined thresholds indicate that a performance counter for a certain storage object, for example, an aggregate or volume, has crossed the threshold you defined in the policy. This indicates that the cluster object is experiencing a performance issue.

You use the Event details page to analyze the performance event and take corrective action, if necessary, to return performance back to normal.

Respond to user-defined performance threshold events in Active IQ Unified Manager

You can use Unified Manager to investigate performance events caused by a performance counter crossing a user-defined warning or critical threshold. You can also use Unified Manager to check the health of the cluster component to see whether recent

health events detected on the component contributed to the performance event.

Before you begin

- You must have the Operator, Application Administrator, or Storage Administrator role.
- There must be new or obsolete performance events.

Steps

1. Display the **Event details** page to view information about the event.
2. Review the **Description**, which describes the threshold breach that caused the event.

   For example, the message “Latency value of 456 ms/op has triggered a WARNING event based on threshold setting of 400 ms/op” indicates that a latency warning event occurred for the object.

3. Hover your cursor over the policy name to display details about the threshold policy that triggered the event.

   This includes the policy name, the performance counter being evaluated, the counter value that must be breached to be considered a critical or warning event, and the duration by which the counter must exceed the value.

4. Make a note of the **Event Trigger Time** so you can investigate whether other events might have occurred at the same time that could have contributed to this event.
5. Follow one of the options below to further investigate the event, to determine whether you need to perform any actions to resolve the performance problem:

Option	Possible investigation actions
Click the Source object name to display the Explorer page for that object.	This page enables you to view the object details and compare this object with other similar storage objects to see whether other storage objects have a performance issue around the same time. For example, to see whether other volumes on the same aggregate are also having a performance issue.
Click the cluster name to display the Cluster Summary page.	This page enables you to view the details for the cluster on which this object resides to see whether other performance issues have occurred around the same time.

Analyze events from system-defined performance thresholds

Analyze events from system-defined performance thresholds in Active IQ Unified Manager

Events generated from system-defined performance thresholds indicate that a performance counter, or set of performance counters, for a certain storage object has crossed the threshold from a system-defined policy. This indicates that the storage object, for example, an aggregate or node, is experiencing a performance issue.

You use the Event details page to analyze the performance event and take corrective action, if necessary, to return performance back to normal.

System-defined threshold policies are not enabled on Cloud Volumes ONTAP, ONTAP Edge, or ONTAP Select systems.

Respond to system-defined performance threshold events in Active IQ Unified Manager

You can use Unified Manager to investigate performance events caused by a performance counter crossing a system-defined warning threshold. You can also use Unified Manager to check the health of the cluster component to see whether recent events detected on the component contributed to the performance event.

Before you begin

- You must have the Operator, Application Administrator, or Storage Administrator role.
- There must be new or obsolete performance events.

Steps

1. Display the **Event details** page to view information about the event.
2. Review the **Description**, which describes the threshold breach that caused the event.

   For example, the message “Node utilization value of 90 % has triggered a WARNING event based on threshold setting of 85 %” indicates that a node utilization warning event occurred for the cluster object.

3. Make a note of the **Event Trigger Time** so you can investigate whether other events might have occurred at the same time that could have contributed to this event.
4. Under **System Diagnosis**, review the brief description of the type of analysis the system-defined policy is performing on the cluster object.

   For some events a green or red icon is displayed next to the diagnosis to indicate whether an issue was found in that particular diagnosis. For other types of system-defined events counter charts display the performance for the object.

5. Under **Suggested Actions**, click the **Help me do this** link to view the suggested actions you can perform to try and resolve the performance event on your own.

Respond to QoS policy group performance events in Active IQ Unified Manager

Unified Manager generates QoS policy warning events when workload throughput (IOPS, IOPS/TB, or MBps) has exceeded the defined ONTAP QoS policy setting and workload latency is becoming affected. These system-defined events provide the opportunity to correct potential performance issues before many workloads are affected by latency.

Before you begin

- You must have the Operator, Application Administrator, or Storage Administrator role.
- There must be new, acknowledged, or obsolete performance events.

Unified Manager generates warning events for QoS policy breaches when workload throughput has exceeded the defined QoS policy setting during each performance collection period for the previous hour. Workload throughput may exceed the QoS threshold for only a short period of time during each collection period, but Unified Manager displays only the “average” throughput during the collection period on the chart. For this reason you may receive QoS events while the throughput for a workload might not have crossed the policy threshold shown in the chart.

You can use System Manager or the ONTAP commands to manage policy groups, including the following tasks:

- Creating a new policy group for the workload
- Adding or removing workloads in a policy group
- Moving a workload between policy groups
- Changing the throughput limit of a policy group
- Moving a workload to a different aggregate or node

Steps

1. Display the **Event details** page to view information about the event.
2. Review the **Description**, which describes the threshold breach that caused the event.

   For example, the message “IOPS value of 1,352 IOPS on vol1_NFS1 has triggered a WARNING event to identify potential performance problems for the workload” indicates that a QoS Max IOPS event occurred on volume vol1_NFS1.

3. Review the **Event Information** section to see more details about when the event occurred and how long the event has been active.

   Additionally, for volumes or LUNs that are sharing the throughput of a QoS policy you can see the names of the top three workloads that are consuming the most IOPS or MBps.

4. Under the **System Diagnosis** section, review the two charts: one for total average IOPS or MBps (depending on the event), and one for latency. When arranged this way you can see which cluster components are most affecting latency when the workload approached the QoS max limit.

   For a shared QoS policy event, the top three workloads are shown in the throughput chart. If more than three workloads are sharing the QoS policy, then additional workloads are added together in an “Other workloads” category. Additionally, the Latency chart shows the average latency on all workloads that are part of the QoS policy.

   Note that for adaptive QoS policy events that the IOPS and MBps charts show IOPS or MBps values that ONTAP has converted from the assigned IOPS/TB threshold policy based on the size of the volume.

5. Under the **Suggested Actions** section, review the suggestions and determine which actions you should perform to avoid an increase in latency for the workload.

   If required, click the **Help** button to view more details about the suggested actions you can perform to try and resolve the performance event.

Understand adaptive QoS events with a defined block size in Active IQ Unified Manager

Adaptive QoS policy groups automatically scale a throughput ceiling or floor based on the volume size, maintaining the ratio of IOPS to TBs as the size of the volume changes. Starting with ONTAP 9.5 you can specify the block size in the QoS policy to effectively apply a MB/s threshold at the same time.

Assigning an IOPS threshold in an adaptive QoS policy places a limit only on the number of operations that occur in each workload. Depending on the block size that is set on the client that generates the workloads, some IOPS include much more data and therefore place a much larger burden on the nodes that process the operations.

The MB/s value for a workload is generated using the following formula:

```
MB/s = (IOPS * Block Size) / 1000
```

If a workload is averaging 3,000 IOPS and the block size on the client is set to 32 KB, then the effective MB/s for this workload is 96. If this same workload is averaging 3,000 IOPS and the block size on the client is set to 48 KB, then the effective MB/s for this workload is 144. You can see that the node is processing 50% more data when the block size is larger.

Let's look at the following adaptive QoS policy that has a defined block size and how events are triggered based on the block size that is set on the client.

Create a policy and set the peak throughput to 2,500 IOPS/TB with a block size of 32KB. This effectively sets the MB/s threshold to 80 MB/s ((2500 IOPS * 32KB) / 1000) for a volume with 1 TB used capacity. Note that Unified Manager generates a Warning event when the throughput value is 10% less than the defined threshold. Events are generated under the following situations:

Used Capacity	**Event is generated when throughput exceeds this number of …**	
	IOPS	MB/s
1 TB	2,250 IOPS	72 MB/s
2 TB	4,500 IOPS	144 MB/s
5 TB	11,250 IOPS	360 MB/s

If the volume is using 2TB of the available space, and the IOPS is 4,000, and the QoS block size is set to 32KB on the client, then the MB/ps throughput is 128 MB/s ((4,000 IOPS * 32 KB) / 1000). No event is generated in this scenario because both 4,000 IOPS and 128 MB/s are below the threshold for a volume that is using 2 TB of space.

If the volume is using 2TB of the available space, and the IOPS is 4,000, and the QoS block size is set to 64KB on the client, then the MB/s throughput is 256 MB/s ((4,000 IOPS * 64 KB) / 1000). In this case the 4,000 IOPS does not generate an event, but the MB/s value of 256 MB/s is above the threshold of 144 MB/s and an event is generated.

For this reason, when an event is triggered based on a MB/s breach for an adaptive QoS policy that includes the block size, a MB/s chart is displayed in the System Diagnosis section of the Event details page. If the event is triggered based on an IOPS breach for the adaptive QoS policy, an IOPS chart is displayed in the System Diagnosis section. If a breach occurs for both IOPS and MB/s you will receive two events.

For more information on adjusting QoS settings, see Performance management overview.

Respond to node resource overutilization events in Active IQ Unified Manager

Unified Manager generates node resources overutilized warning events when a single node is operating above the bounds of its operational efficiency, and therefore potentially affecting workload latencies. These system-defined events provide the opportunity to correct potential performance issues before many workloads are affected by latency.

Before you begin

- You must have the Operator, Application Administrator, or Storage Administrator role.

- There must be new or obsolete performance events.

Unified Manager generates warning events for node resources overutilized policy breaches by looking for nodes that are using more than 100% of their performance capacity for more than 30 minutes.

You can use System Manager or the ONTAP commands to correct this type of performance issue, including the following tasks:

- Creating and applying a QoS policy to any volumes or LUNs that are overusing system resources
- Reducing the QoS maximum throughput limit of a policy group to which workloads have been applied
- Moving a workload to a different aggregate or node
- Increasing capacity by adding disks to the node, or by upgrading to a node with a faster CPU and more RAM

Steps

1. Display the **Event details** page to view information about the event.
2. Review the **Description**, which describes the threshold breach that caused the event.

   For example, the message “Perf. Capacity Used value of 139% on simplicity-02 has triggered a WARNING event to identify potential performance problems in the data processing unit.” indicates that performance capacity on node simplicity-02 is overused and affecting node performance.

3. Under the **System Diagnosis** section, review the three charts: one for performance capacity used on the node, one for average storage IOPS being used by the top workloads, and one for latency on the top workloads. When arranged in this way you can see which workloads are the cause of the latency on the node.

   You can view which workloads have QoS policies applied, and which do not, by moving your cursor over the IOPS chart.

4. Under the **Suggested Actions** section, review the suggestions and determine which actions you should perform to avoid an increase in latency for the workload.

   If required, click the **Help** button to view more details about the suggested actions you can perform to try and resolve the performance event.

Respond to cluster imbalance performance events in Active IQ Unified Manager

Unified Manager generates cluster imbalance warning events when one node in a cluster is operating at a much higher load than other nodes, and therefore potentially affecting workload latencies. These system-defined events provide the opportunity to correct potential performance issues before many workloads are affected by latency.

Before you begin

You must have the Operator, Application Administrator, or Storage Administrator role.

Unified Manager generates warning events for cluster imbalance threshold policy breaches by comparing the performance capacity used value for all nodes in the cluster to see if there is a load difference of 30% between any nodes.

These steps help you identify the following resources so that you can move high-performing workloads to a lower utilized node:

- The nodes on the same cluster that are less utilized
- The aggregates on the new node that are the least utilized
- The highest-performing volumes on the current node

Steps

1. Display the **Event** details page to view information about the event.
2. Review the **Description**, which describes the threshold breach that caused the event.

   For example, the message “The performance capacity used counter indicates a load difference of 62% between the nodes on cluster Dallas-1-8 and has triggered a WARNING event based on the system threshold of 30%” indicates that performance capacity on one of the nodes is overused and affecting node performance.

3. Review the text in the **Suggested Actions** to move a high-performing volume from the node with the high performance capacity used value to a node with the lowest performance capacity used value.
4. Identify the nodes with the highest and lowest performance capacity used value:
   a. In the **Event Information** section, click the name of the source cluster.
   b. In the **Cluster / Performance Summary** page, click **Nodes** in the **Managed Objects** area.
   c. In the **Nodes** inventory page, sort the nodes by the **Performance Capacity Used** column.
   d. Identify the nodes with the highest and lowest performance capacity used value and write down those names.
5. Identify the volume using the most IOPS on the node that has the highest performance capacity used value:
   a. Click the node with the highest performance capacity used value.
   b. In the **Node / Performance Explorer** page, select **Aggregates on this Node** from the **View and Compare** menu.
   c. Click the aggregate with the highest performance capacity used value.
   d. In the **Aggregate / Performance Explorer** page, select **Volumes on this Aggregate** from the **View and Compare** menu.
   e. Sort the volumes by the **IOPS** column, and write down the name of the volume using the most IOPS, and the name of the aggregate where the volume resides.
6. Identify the aggregate with the lowest utilization on the node that has the lowest performance capacity used value:
   a. Click **Storage** > **Aggregates** to display the **Aggregates** inventory page.
   b. Select the **Performance: All Aggregates** view.
   c. Click the **Filter** button and add a filter where “Node” equals the name of the node with the lowest performance capacity used value that you wrote down in step 4.
   d. Write down the name of the aggregate that has the lowest performance capacity used value.
7. Move the volume from the overloaded node to the aggregate you identified as having low utilization on the new node.

   You can perform the move operation by using ONTAP System Manager, OnCommand Workflow Automation, ONTAP commands, or a combination of these tools.

After a few days, check to see whether you are receiving the same cluster imbalance event from this cluster.

Analyze events from dynamic performance thresholds

Analyse events from dynamic performance thresholds in Active IQ Unified Manager

Events generated from dynamic thresholds indicate that the actual response time (latency) for a workload is too high, or too low, compared to the expected response time range. You use the Event details page to analyze the performance event and take corrective action, if necessary, to return performance back to normal.

Dynamic performance thresholds are not enabled on Cloud Volumes ONTAP, ONTAP Edge, or ONTAP Select systems.

Identify victim workloads in dynamic performance events in Active IQ Unified Manager

In Unified Manager, you can identify which volume workloads have the highest deviation in response time (latency) caused by a storage component in contention. Identifying these workloads helps you understand why the client applications accessing them have been performing slower than usual.

Before you begin

- You must have the Operator, Application Administrator, or Storage Administrator role.
- There must be new, acknowledged, or obsolete dynamic performance events.

The Event details page displays a list of the user-defined and system-defined workloads, ranked by the highest deviation in activity or usage on the component or most impacted by the event. The values are based on the peaks that Unified Manager identified when it detected and last analyzed the event.

Steps

1. Display the **Event details** page to view information about the event.
2. In the Workload Latency and Workload Activity charts, select **Victim Workloads**.
3. Hover your cursor over the charts to view the top user-defined workloads that are affecting the component, and the name of the victim workload.

Identify bully workloads in dynamic performance events in Active IQ Unified Manager

In Unified Manager, you can identify which workloads have the highest deviation in usage for a cluster component in contention. Identifying these workloads helps you understand why certain volumes on the cluster have slow response times (latency).

Before you begin

- You must have the Operator, Application Administrator, or Storage Administrator role.
- There must be new, acknowledged, or obsolete dynamic performance events.

The Event details page displays a list of the user-defined and system-defined workloads ranked by the highest usage of the component or most impacted by the event. The values are based on the peaks that Unified Manager identified when it detected and last analyzed the event.

Steps

1. Display the Event details page to view information about the event.
2. In the Workload Latency and Workload Activity charts, select **Bully Workloads**.
3. Hover your cursor over the charts to view the top user-defined bully workloads that are affecting the component.

Identify shark workloads in a dynamic performance event in Active IQ Unified Manager

In Unified Manager, you can identify which workloads have the highest deviation in usage for a storage component in contention. Identifying these workloads helps you determine if these workloads should be moved to a less-utilized cluster.

Before you begin

- You must have the Operator, Application Administrator, or Storage Administrator role.
- There are new, acknowledged, or obsolete performance dynamic event.

The Event details page displays a list of the user-defined and system-defined workloads ranked by the highest usage of the component or most impacted by the event. The values are based on the peaks that Unified Manager identified when it detected and last analyzed the event.

Steps

1. Display the **Event details** page to view information about the event.
2. In the Workload Latency and Workload Activity charts, select **Shark Workloads**.
3. Hover your cursor over the charts to view the top user-defined workloads that are affecting the component, and the name of the shark workload.

Performance event analysis for a MetroCluster configuration

Analyze MetroCluster performance events with Active IQ Unified Manager

You can use Unified Manager to analyze a performance event for a MetroCluster configuration. You can identify the workloads involved in the event and review the suggested actions for resolving it.

MetroCluster performance events might be due to *bully* workloads that are over-utilizing the interswitch links (ISLs) between the clusters, or due to link health issues. Unified Manager monitors each cluster in a MetroCluster configuration independently, without consideration of performance events on a partner cluster.

Performance events from both clusters in the MetroCluster configuration are also displayed on the Unified Manager Dashboard page. You can also view the Health pages of Unified Manager to check the health of each cluster and to view their relationship.

Analyze a MetroCluster performance event on a cluster in Active IQ Unified Manager

You can use Unified Manager to analyze the cluster in a MetroCluster configuration on which a performance event was detected. You can identify the cluster name, event detection time, and the *bully* and *victim* workloads involved.

Before you begin

- You must have the Operator, Application Administrator, or Storage Administrator role.

- There must be new, acknowledged, or obsolete performance events for a MetroCluster configuration.
- Both clusters in the MetroCluster configuration must be monitored by the same instance of Unified Manager.

Steps

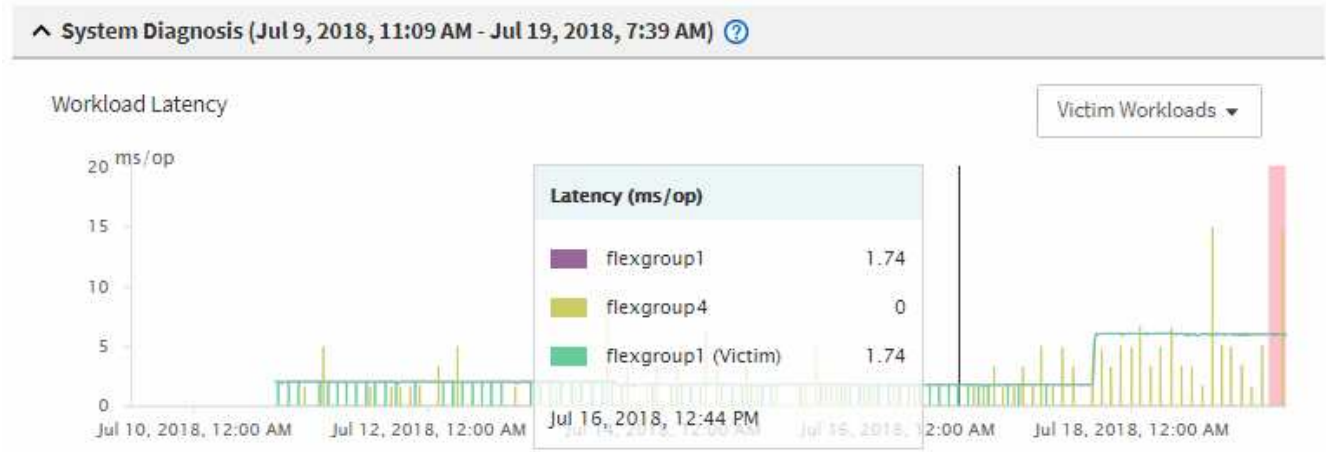
1. Display the **Event details** page to view information about the event.
2. Review the event description to see the names of the workloads involved and the number of workloads involved.

   In this example, the MetroCluster Resources icon is red, indicating that the MetroCluster resources are in contention. You position your cursor over the icon to display a description of the icon.



3. Make a note of the cluster name and the event detection time, which you can use to analyze performance events on the partner cluster.
4. In the charts, review the *victim* workloads to confirm that their response times are higher than the performance threshold.

   In this example, the victim workload is displayed in the hover text. The Latency charts display, at a high-level, a consistent latency pattern for the victim workloads involved. Even though the abnormal latency of the victim workloads triggered the event, a consistent latency pattern might indicate that the workloads are performing within their expected range, but that a spike in I/O increased the latency and triggered the event.



   If you recently installed an application on a client that accesses these volume workloads and that application sends a high amount of I/O to them, you might be anticipating their latencies to increase. If the latency for the workloads returns within the expected range, the event state changes to obsolete, and remains in this state for more than 30 minutes, you can probably ignore the event. If the event is ongoing, and remains in the new state, you can investigate it further to determine whether other issues caused the event.

5. In the Workload Throughput chart, select **Bully Workloads** to display the bully workloads.

The presence of bully workloads indicates that the event might have been caused by one or more workloads on the local cluster overutilizing the MetroCluster resources. The bully workloads have a high deviation in write throughput (MB/s).

This chart displays, at a high-level, the write throughput (MB/s) pattern for the workloads. You can review the write MB/s pattern to identify abnormal throughput, which might indicate that a workload is over-utilizing the MetroCluster resources.

If no bully workloads are involved in the event, the event might have been caused by a health issue with the link between the clusters or a performance issue on the partner cluster. You can use Unified Manager to check the health of both clusters in a MetroCluster configuration. You can also use Unified Manager to check for and analyze performance events on the partner cluster.

Analyze a MetroCluster performance event on a remote cluster in Active IQ Unified Manager

You can use Unified Manager to analyze dynamic performance events on a remote cluster in a MetroCluster configuration. The analysis helps you determine whether an event on the remote cluster caused an event on its partner cluster.

Before you begin

- You must have the Operator, Application Administrator, or Storage Administrator role.
- You must have analyzed a performance event on a local cluster in a MetroCluster configuration and obtained the event detection time.
- You must have checked the health of the local cluster and its partner cluster involved in the performance event and obtained the name of the partner cluster.

Steps

1. Log in to the Unified Manager instance that is monitoring the partner cluster.
2. In the left navigation pane, click **Events** to display the event list.
3. From the **Time Range** selector, select **Last Hour**, and then click **Apply Range**.
4. In the **Filtering** selector, select **Cluster** from the left drop-down menu, type the name of the partner cluster in the text field, and then click **Apply Filter**.

   If there are no events for the selected cluster over the last hour, this indicates that the cluster has not experienced any performance issues during the time that the event was detected on its partner.

5. If the selected cluster has events detected over the last hour, compare the event detection time to the event detection time for the event on the local cluster.

   If these events involve bully workloads causing contention on the data processing component, one or more of these bullies might have caused the event on the local cluster. You can click the event to analyze it and review the suggested actions for resolving it on the Event details page.

   If these events do not involve bully workloads, they did not cause the performance event on the local cluster.

Respond to a dynamic QoS throttling event in Active IQ Unified Manager

You can use Unified Manager to investigate a performance event caused by a Quality of Service (QoS) policy group throttling workload throughput (MB/s). The throttling increased the response times (latency) of volume workloads in the policy group. You can use the

event information to determine whether new limits on the policy groups are needed to stop the throttling.

Before you begin

- You must have the Operator, Application Administrator, or Storage Administrator role.
- There must be new, acknowledged, or obsolete performance events.

Steps

1. Display the **Event details** page to view information about the event.
2. Read the **Description**, which displays the name of the workloads impacted by the throttling.

   The description can display the same workload for the victim and bully, because the throttling makes the workload a victim of itself.

3. Record the name of the volume, using an application such as a text editor.

   You can search on the volume name to locate it later.

4. In the Workload Latency and Workload Utilization charts, select **Bully Workloads**.
5. Hover your cursor over the charts to view the top user-defined workloads that are affecting the policy group.

   The workload at the top of the list has the highest deviation and caused the throttling to occur. The activity is the percentage of the policy group limit used by each workload.

6. In the **Suggested Actions** area, click the **Analyze Workload** button for the top workload.
7. In the Workload Analysis page, set the Latency chart to view all Cluster Components, and the Throughput chart to view Breakdown.

   The breakdown charts are displayed under the Latency chart and the IOPS chart.

8. Compare the QoS Limits in the **Latency** chart to see what amount of throttling impacted the latency at the time of the event.

   The QoS policy group has a maximum throughput of 1,000 operations per second (op/sec), which the workloads in it cannot collectively exceed. At the time of the event, the workloads in the policy group had a combined throughput of over 1,200 op/sec, which caused the policy group to throttle its activity back to 1,000 op/sec.

9. Compare the **Reads/writes latency** values to the **Reads/writes/other** values.

   Both charts show a high number of read requests with high latency, but the number of requests and amount of latency for write requests is low. These values help you determine whether there is a high amount of throughput or number of operations that increased the latency. You can use these values when deciding to put a policy group limit on the throughput or operations.

10. Use ONTAP System Manager to increase the current limit on the policy group to 1,300 op/sec.
11. After a day, return to Unified Manager and enter the workload that you recorded in Step 3 in the **Workload Analysis** page.
12. Select the Throughput Breakdown chart.

The Reads/writes/other chart is displayed.

13. At the top of the page, point your cursor to the change event icon ( ) for the policy group limit change.
14. Compare the **Reads/writes/other** chart to the **Latency** chart.

    The read and write requests are the same, but the throttling has stopped and the latency has decreased.

Respond to a performance event caused by disk failure in Active IQ Unified Manager

You can use Unified Manager to investigate a performance event caused by workloads overutilizing an aggregate. You can also use Unified Manager to check the health of the aggregate to see if recent health events detected on the aggregate contributed to the performance event.

Before you begin

- You must have the Operator, Application Administrator, or Storage Administrator role.
- There must be new, acknowledged, or obsolete performance events.

Steps

1. Display the **Event details** page to view information about the event.
2. Read the **Description**, which describes the workloads involved in the event and the cluster component in contention.

   There are multiple victim volumes whose latency was impacted by the cluster component in contention. The aggregate, which is in the middle of a RAID reconstruct to replace the failed disk with a spare disk, is the cluster component in contention. Under Component in Contention, the Aggregate icon is highlighted red and the name of the aggregate is displayed in parentheses.

3. In the Workload Utilization chart, select **Bully Workloads**.
4. Hover your cursor over the chart to view the top bully workloads that are affecting the component.

   The top workloads with the highest peak utilization since the event was detected are displayed at the top of the chart. One of the top workloads is the system-defined workload Disk Health, which indicates a RAID reconstruct. A reconstruct is the internal process involved with rebuilding the aggregate with the spare disk. The Disk Health workload, along with other workloads on the aggregate, likely caused the contention on the aggregate and the associated event.

5. After confirming that the activity from the Disk Health workload caused the event, wait for approximately 30 minutes for the reconstruction to finish and for Unified Manager to analyze the event and detect whether the aggregate is still in contention.
6. Refresh the **Event details**.

   After the RAID reconstruction is complete, check that the State is obsolete, indicating that the event is resolved.

7. In the Workload Utilization chart, select **Bully Workloads** to view the workloads on the aggregate by peak utilization.
8. In the **Suggested Actions** area, click the **Analyze Workload** button for the top workload.
9. In the **Workload Analysis** page, set the Time Range to display the last 24 hours (1 day) of data for the selected volume.

In the Event Timeline, a red dot (🔴) indicates when the disk failure event occurred.

10. In the Node and Aggregate Utilization chart, hide the line for the Node statistics so that just the Aggregate line remains.
11. Compare the data in this chart to the data at the time of the event in the **Latency** chart.

    At the time of the event, the Aggregate Utilization shows a high amount of read and write activity, caused by the RAID reconstruction processes, which increased the latency of the selected volume. A few hours after the event occurred, both the reads and writes and the latency have decreased, confirming that the aggregate is no longer in contention.

Respond to a performance event caused by HA takeover in Active IQ Unified Manager

You can use Unified Manager to investigate a performance event caused by high data processing on a cluster node that is in a high-availability (HA) pair. You can also use Unified Manager to check the health of the nodes to see whether any recent health events detected on the nodes contributed to the performance event.

Before you begin

- You must have the Operator, Application Administrator, or Storage Administrator role.
- There must be new, acknowledged, or obsolete performance events.

Steps

1. Display the **Event details** page to view information about the event.
2. Read the **Description**, which describes the workloads involved in the event and the cluster component in contention.

   There is one victim volume whose latency was impacted by the cluster component in contention. The data processing node, which took over all workloads from its partner node, is the cluster component in contention. Under Component in Contention, the Data Processing icon is highlighted red and the name of the node that was handling data processing at the time of the event is displayed in parentheses.

3. In the **Description**, click the name of the volume.

   The Volume Performance Explorer page is displayed. At the top of the page, in the Events time line, a change event icon (🔵) indicates the time that Unified Manager detected the start of the HA takeover.

4. Point your cursor to the change event icon for the HA takeover and details about the HA takeover are displayed in hover text.

   In the Latency chart, an event indicates that the selected volume crossed the performance threshold due to high latency around the same time as the HA takeover.

5. Click **Zoom View** to display the Latency chart on a new page.
6. In the View menu, select **Cluster Components** to view the total latency by cluster component.
7. Point your mouse cursor to the change event icon for the start of the HA takeover and compare the latency for data processing to the total latency.

   At the time of the HA takeover, there was a spike in data processing from the increased workload demand on the data processing node. The increased CPU utilization drove up the latency and triggered the event.

8. After fixing the failed node, use ONTAP System Manager to perform an HA giveback, which moves the workloads from the partner node to the fixed node.
9. After the HA giveback is complete, after the next configuration discovery in Unified Manager (approximately 15 minutes), find the event and workload that triggered by the HA takeover in the **Event Management** inventory page.

   The event triggered by the HA takeover now has a state of obsolete, which indicates that the event is resolved. The latency at the data processing component has decreased, which has decreased the total latency. The node that the selected volume is now using for data processing has resolved the event.

Resolve performance events

Resolve performance events in Active IQ Unified Manager

You can use the suggested actions to try and resolve performance events on your own. The first three suggestions are always displayed, and the actions under the fourth suggestion are specific to the type of event displayed.

The **Help me do this** links provide additional information for each suggested action, including instructions for performing a specific action. Some of the actions may involve using Unified Manager, ONTAP System Manager, OnCommand Workflow Automation, ONTAP CLI commands, or a combination of these tools.

Confirm that latency is within the expected range in Active IQ Unified Manager

When a cluster component is in contention, volume workloads that use it might have decreased response time (latency). You can review the latency of each victim workload on the component in contention to confirm that its actual latency is within its expected range. You can also click a volume name to view the historical data for the volume.

If the performance event is in the obsolete state, the latency of each victim involved in the event might have returned within its expected range.

Assess how configuration changes affect workload performance in Active IQ Unified Manager

Configuration changes on the cluster, such as a failed disk, HA failover, or a moved volume, could negatively impact volume performance and cause increased latency.

In Unified Manager, you can review the Workload Analysis page to see when a recent configuration change occurred and compare it to the operations and latency (response time) to see whether there was a change in activity for the selected volume workload.

The performance pages of Unified Manager can only detect a limited number of change events. The health pages provide alerts for other events caused by configuration changes. You can search for the volume in Unified Manager to see the event history.

Improve workload performance by adjusting client-side settings in Active IQ Unified Manager

You can check your client workloads, such as applications or databases, that are sending

I/O to volumes involved in a performance event to determine if a client-side change might correct the event.

When the clients that are connected to volumes on a cluster increase their I/O requests, the cluster must work harder to meet the demand. If you know which clients have a high number of I/O requests to a particular volume on the cluster, you can improve cluster performance by adjusting the number of clients accessing the volume or decreasing the amount of I/O to the volume. You can also apply or increase a limit on the QoS policy group of which the volume is a member.

You can investigate clients and their applications to determine whether the clients are sending more I/O than usual, which might be causing contention on a cluster component. On the Event details page, the System Diagnosis section displays the top volume workloads using the component in contention. If you know which client is accessing a particular volume, you can go to the client to determine whether the client hardware or an application is not operating as expected or is doing more work than usual.

In a MetroCluster configuration, write requests to a volume on a local cluster are mirrored to a volume on the remote cluster. Keeping the source volume on the local cluster in sync with the destination volume on the remote cluster can also increase the demand of both clusters in the MetroCluster configuration. By reducing write requests to these mirrored volumes, the clusters can perform fewer sync operations, which reduces the performance impact on other workloads.

Understand how client or network issues affect performance in Active IQ Unified Manager

When the clients that are connected to volumes on a cluster increase their I/O requests, the cluster must work harder to meet the demand. The increased demand on the cluster can put a component in contention, increase the latency of workloads that use it, and trigger an event in Unified Manager.

On the Event details page, the System Diagnosis section displays the top volume workloads using the component in contention. If you know which client is accessing a particular volume, you can go to the client to determine whether the client hardware or an application is not operating as expected or is doing more work than usual. You might need to contact your client administrator or application vendor for assistance.

You can check your network infrastructure to determine whether there are hardware issues, bottlenecks, or competing workloads that might have caused I/O requests between the cluster and connected clients to perform slower than expected. You might need to contact your network administrator for assistance.

Analyze high-activity volumes within a QoS policy group in Active IQ Unified Manager

You can review the workloads in the Quality of Service (QoS) policy group with the highest change in activity to determine whether more than one workload caused the event. You can also see whether other workloads are still exceeding the set throughput limit or whether they are back within their expected range of activity.

On the Event details page, in the System Diagnosis section, you can sort the workloads by peak deviation in activity to display the workloads with the highest change in activity at the top of the table. These workloads might be the “bullies” whose activity exceeded the set limit and might have caused the event.

You can navigate to the Workload Analysis page for each volume workload to review its IOPS activity. If the workload has periods of very high operations activity, it might have contributed to the event. You can change

the policy group settings for the workload or move the workload to a different policy group.

You can use ONTAP System Manager or the ONTAP CLI commands to manage policy groups, as follows:

- Create a policy group.
- Add or remove workloads in a policy group.
- Move a workload between policy groups.
- Change the throughput limit of a policy group.

Move LIFs to improve performance in Active IQ Unified Manager

Moving logical interfaces (LIFs) to a less busy port can help improve load balancing, assist with maintenance operations and performance tuning, and reduce indirect access.

Indirect access can reduce system efficiency. It occurs when a volume workload is using different nodes for network processing and data processing. To reduce indirect access, you can rearrange LIFs, which involves moving LIFs to use the same node for network processing and data processing. You can configure load balancing to have ONTAP automatically move busy LIFs to a different port or you can move a LIF manually.

Benefits	Considerations
• Improve load balancing. • Reduce indirect access.	When moving a LIF connected to CIFS shares, clients accessing the CIFS shares are disconnected. Any read or write requests to the CIFS shares are disrupted.

You use the ONTAP commands to configure load balancing. For more information, see the ONTAP networking documentation.

You use ONTAP System Manager and the ONTAP CLI commands to move LIFs manually.

Run storage efficiency operations at less busy times

Schedule storage efficiency operations during less busy times in Active IQ Unified Manager

You can modify the policy or schedule that handles storage efficiency operations to run when the impacted volume workloads are less busy.

Storage efficiency operations can use a high amount of cluster CPU resources and become a bully to the volumes on which the operations are being run. If the victim volumes have high activity at the same time when the storage efficiency operations are run, their latency can increase and trigger an event.

On the Event details page, the System Diagnosis section displays workloads in the QoS policy group by peak deviation in activity to identify the bully workloads. If you see “storage efficiency” displayed near the top of the table, these operations are bullying the victim workloads. By modifying the efficiency policy or schedule to run when these workloads are less busy, you can prevent the storage efficiency operations from causing contention on a cluster.

You can use ONTAP System Manager to manage efficiency policies. You can use the ONTAP commands to manage efficiency policies and schedules.

Understand storage efficiency in Active IQ Unified Manager

Storage efficiency enables you to store the maximum amount of data for the lowest cost and accommodates rapid data growth while consuming less space. NetApp strategy for storage efficiency is based on the built-in foundation of storage virtualization and unified storage provided by its core ONTAP operating system and Write Anywhere File Layout (WAFL) file system.

Storage efficiency includes using technologies such as thin provisioning, Snapshot copy, deduplication, data compression, FlexClone, thin replication with SnapVault and volume SnapMirror, RAID-DP, Flash Cache, Flash Pool aggregate, and FabricPool-enabled aggregates which help to increase storage utilization and decrease storage costs.

The unified storage architecture allows you to efficiently consolidate a storage area network (SAN), network-attached storage (NAS), and secondary storage on a single platform.

High-density disk drives, such as serial advanced technology attachment (SATA) drives configured within Flash Pool aggregate or with Flash Cache and RAID-DP technology, increase efficiency without affecting performance and resiliency.

A FabricPool-enabled aggregate includes an all SSD aggregate or HDD aggregate (starting with ONTAP 9.8) as the local performance tier and an object store that you specify as the cloud tier. Configuring FabricPool helps you manage which storage tier (the local tier or the cloud tier) data should be stored based on whether the data is frequently accessed.

Technologies such as thin provisioning, Snapshot copy, deduplication, data compression, thin replication with SnapVault and volume SnapMirror, and FlexClone offer better savings. You can use these technologies individually or together to achieve maximum storage efficiency.

Understand how adding disks affects data reallocation in Active IQ Unified Manager

You can add disks to an aggregate to increase the storage capacity and the performance of that aggregate. After adding the disks, you will see an improvement in read performance only after reallocating the data across the disks you added.

You can use these instructions when Unified Manager has received aggregate events triggered by dynamic thresholds or by system-defined performance thresholds:

- When you have received a dynamic threshold event, on the Event details page, the cluster component icon that represents the aggregate in contention is highlighted red.

  Beneath the icon, in parentheses, is the name of the aggregate, which identifies the aggregate to which you can add disks.

- When you have received a system-defined threshold event, on the Event details page, the event description text lists the name of the aggregate that is having the problem.

  You can add disks and reallocate data on this aggregate.

The disks you add to the aggregate must already exist in the cluster. If the cluster does not have extra disks available, you might need to contact your administrator or purchase more disks. You can use ONTAP System Manager or the ONTAP commands to add disks to an aggregate.

You should reallocate data when using HDD and Flash Pool aggregates only. Do not reallocate data on SSD or FabricPool aggregates.

Improve workload performance by enabling Flash Cache on a node in Active IQ Unified Manager

You can improve workload performance by enabling Flash Cache™ intelligent data caching on each node in the cluster.

A Flash Cache module, or Performance Acceleration Module PCIe-based memory module, optimizes the performance of random read-intensive workloads by functioning as an intelligent external read cache. This hardware works in tandem with the WAFL External Cache software component of ONTAP.

In Unified Manager, on the Event details page, the cluster component icon that represents the aggregate in contention is highlighted red. Beneath the icon, in parentheses, is the name of the aggregate, which identifies the aggregate. You can enable Flash Cache on the node on which the aggregate resides.

You can use ONTAP System Manager or the ONTAP commands to see whether Flash Cache is installed or enabled, and to enable it if not already enabled. The following command indicates whether Flash Cache is enabled on a specific node: **`cluster::> run local options flexscale.enable`**

For more information about Flash Cache and the requirements for using it, see the following technical report:

Technical Report 3832: Flash Cache Best Practices Guide

Improve workload performance by enabling Flash Pool on a storage aggregate in Active IQ Unified Manager

You can improve workload performance by enabling the Flash Pool feature on an aggregate. A Flash Pool is an aggregate that incorporates both HDDs and SSDs. The HDDs are used for primary storage and the SSDs provide a high-performance read and write cache to boost aggregate performance.

In Unified Manager, the Event details page displays the name of the aggregate in contention. You can use ONTAP System Manager or the ONTAP commands to see whether Flash Pool is enabled for an aggregate. If you have SSDs installed, you can use the command-line interface to enable it. If you have SSDs installed, you can run the following command on the aggregate to see whether Flash Pool is enabled: **`cluster::> storage aggregate show -aggregate aggr_name -field hybrid-enabled`**

In this command, *`aggr_name`* is the name of the aggregate, such as the aggregate in contention.

For more information about Flash Pool and the requirements for using it, see the *Clustered Data ONTAP Physical Storage Management Guide*.

Check MetroCluster configuration health in Active IQ Unified Manager

You can use Unified Manager to review the health of the clusters in a MetroCluster configuration over IP or FC. The health status and events help you determine whether there are hardware or software issues that might be impacting the performance of your workloads.

If you configure Unified Manager to send email alerts, you can check your email for any health issues on the local or remote cluster that might have contributed to a performance event. In the Unified Manager GUI, you can select **Event Management** to see a list current events and then use the filters to display MetroCluster configuration events only.

For more information, see Checking the health of clusters in a MetroCluster configuration

Verify MetroCluster configurations in Active IQ Unified Manager

You can prevent performance issues for mirrored workloads in a MetroCluster over FC and IP configurations by ensuring that the MetroCluster configuration is set up correctly. You can also improve workload performance by changing the configuration or upgrading software or hardware components.

Refer to MetroCluster Documentation for the instructions of setting up the clusters in the MetroCluster configuration, including the Fibre Channel (FC) switches, cables, and inter-switch links (ISLs). It also helps you configure the MetroCluster software so that the local and remote clusters can communicate with mirror volume data.
For information specific to your MetroCluster over IP setup, see Install a MetroCluster IP configuration.

You can compare your MetroCluster configuration to the requirements in MetroCluster Documentation to determine whether changing or upgrading components in your MetroCluster configuration might improve workload performance. This comparison can help you answer the following questions:

- Are the controllers appropriate for your workloads?
- Do you need to upgrade your ISL bundles to a larger bandwidth to handle more throughput?
- Can you adjust the buffer-to-buffer credits (BBC) on your switches to increase the bandwidth?
- If your workloads have high write throughput to solid state drive (SSD) storage, do you need to upgrade your FC-to-SAS bridges to accommodate the throughput?

Related information

- For information about replacing or upgrading MetroCluster components, see the MetroCluster Documentation.
- For information about upgrading controllers see Upgrading controllers in a MetroCluster FC configuration using switchover and switchback and Upgrading controllers in a MetroCluster IP configuration using switchover and switchback

Move workloads to a different aggregate in Active IQ Unified Manager

You can use Unified Manager to help identify an aggregate that is less busy than the aggregate where your workloads currently reside, and then you can move selected volumes or LUNs to that aggregate. Moving high performing workloads to a less busy aggregate, or an aggregate with flash storage enabled, allows the workload to perform more efficiently.

Before you begin

- You must have the Operator, Application Administrator, or Storage Administrator role.
- You must have recorded the name of the aggregate that is currently having a performance issue.

- You must have recorded the date and time at which the aggregate received the event.
- Unified Manager must have collected and analyzed a month or more of performance data.

These steps help you identify the following resources so that you can move high-performing workloads to a lower utilized aggregate:

- The aggregates on the same cluster that are less utilized
- The highest-performing volumes on the current aggregate

Steps

1. Identify the aggregate in the cluster that is the least utilized:
   a. From the **Event** details page, click the name of the cluster on which the aggregate resides.

      The cluster details are displayed in the Performance/Cluster Landing page.

   b. On the **Summary** page, click **Aggregates** from the **Managed Objects** pane.

      The list of aggregates on this cluster are displayed.

   c. Click the **Utilization** column to sort the aggregates by least utilized.

      You can also identify those aggregates that have the greatest **Free Capacity**. This provides a list of potential aggregates to which you might want to move workloads.

   d. Write down the name of the aggregate to which you want to move the workloads.
2. Identify the high-performing volumes from the aggregate that received the event:
   a. Click the aggregate that is having the performance issue.

      The aggregate details are displayed in the Performance/Aggregate Explorer page.

   b. From the **Time Range** selector, select **Last 30 Days**, and then click **Apply Range**.

      This enables you to view a longer performance history period than the default 72 hours. You want to move a volume that is using a lot of resources on a consistent basis, not just over the past 72 hours.

   c. From the **View and Compare** control, select **Volumes on this Aggregate**.

      A list of FlexVol volumes and FlexGroup constituent volumes on this aggregate are displayed.

   d. Sort the volumes by highest MB/s, and then by highest IOPS, to see the highest performing volumes.
   e. Write down the names of the volumes that you want to move to a different aggregate.
3. Move the high-performing volumes to the aggregate you identified as having low utilization.

   You can perform the move operation by using ONTAP System Manager, OnCommand Workflow Automation, ONTAP commands, or a combination of these tools.

After a few days, check to see whether you are receiving the same type of events from this node or aggregate.

Move workloads to a different node in Active IQ Unified Manager

You can use Unified Manager to help identify an aggregate on a different node that is less

busy than the node on which your workloads are currently running, and then you can move selected volumes to that aggregate. Moving high-performing workloads to an aggregate on a less busy node allows the workloads on both nodes to perform more efficiently.

Before you begin

- You must have the Operator, Application Administrator, or Storage Administrator role.
- You must have recorded the name of the node that is currently having a performance issue.
- You must have recorded the date and time at which the node received the performance event.
- Unified Manager must have collected and analyzed performance data for a month or longer.

This procedure helps you to identify the following resources so that you can move high-performing workloads to a lower utilized node:

- The nodes on the same cluster that have the greatest free performance capacity
- The aggregates on the new node that have the greatest free performance capacity
- The highest-performing volumes on the current node

Steps

1. Identify a node in the cluster that has the greatest free performance capacity:
   a. On the **Event Details** page, click the name of the cluster on which the node resides.

      The cluster details are displayed in the Performance/Cluster Landing page.

   b. On the **Summary** tab, click **Nodes** from the **Managed Objects** pane.

      The list of nodes on this cluster are displayed.

   c. Click the **Performance Capacity Used** column to sort the nodes by least percentage used.

      This provides a list of potential nodes to which you might want to move workloads.

   d. Write down the name of the node to which you want to move the workloads.
2. Identify an aggregate on the new node that is the least utilized:
   a. In the left navigation pane, click **Storage** > **Aggregates** and select **Performance** > **All Aggregates** from the View menu.

      The Performance: All Aggregates view is displayed.

   b. Click **Filtering**, select **Node** from the left drop-down menu, type the name of the node in the text field, and then click **Apply Filter**.

      The Performance: All Aggregates view is redisplayed with the list of aggregates that are available on this node.

   c. Click the **Performance Capacity Used** column to sort the aggregates by least used.

      This provides a list of potential aggregates to which you might want to move workloads.

   d. Write down the name of the aggregate to which you want to move the workloads.

3. Identify the high-performing workloads from the node that received the event:
   a. Return to the **Event Details** page for the event.
   b. In the **Affected Volumes** field, click the link for the number of volumes.

      The Performance: All Volumes view is displayed with a filtered list of the volumes on that node.

   c. Click the **Total Capacity** column to sort the volumes by the largest allocated space.

      This provides a list of potential volumes that you may want to move.

   d. Write down the names of the volumes that you want to move, and the names of the current aggregates on which they reside.
4. Move the volumes to the aggregates that you identified as having greatest free performance capacity on the new node.

   You can perform the move operation by using ONTAP System Manager, OnCommand Workflow Automation, ONTAP commands, or a combination of these tools.

After a few days, you can check whether you are receiving the same type of events from this node or aggregate.

Move workloads to an aggregate on a different node in Active IQ Unified Manager

You can use Unified Manager to help identify an aggregate on a different node that is less busy than the node where your workloads are currently running, and then you can move selected volumes to that aggregate. Moving high-performing workloads to an aggregate on a less busy node allows workloads on both nodes to perform more efficiently.

Before you begin

- You must have the Operator, Application Administrator, or Storage Administrator role.
- You must have recorded the name of the node that is currently having a performance issue.
- You must have recorded the date and time at which the node received the performance event.
- Unified Manager must have collected and analyzed a month or more of performance data.

These steps help you identify the following resources so that you can move high-performing workloads to a lower utilized node:

- The nodes on the same cluster that are less utilized
- The aggregates on the new node that are the least utilized
- The highest-performing volumes on the current node

Steps

1. Identify a node in the cluster that is the least utilized:
   a. From the **Event** details page, click the name of the cluster on which the node resides.

      The cluster details are displayed in the Performance/Cluster Landing page.

   b. On the **Summary** page, click **Nodes** from the **Managed Objects** pane.

The list of nodes on this cluster are displayed.

c. Click the **Utilization** column to sort the nodes by least utilized.

You can also identify those nodes that have the greatest **Free Capacity**. This provides a list of potential nodes to which you might want to move workloads.

d. Write down the name of the node to which you want to move the workloads.

2. Identify an aggregate on the new node that is the least utilized:

a. In the left navigation pane, click **Storage** > **Aggregates** and select **Performance** > **All Aggregates** from the View menu.

The Performance: All Aggregates view is displayed.

b. Click **Filtering**, select **Node** from the left drop-down menu, type the name of the node in the text field, and then click **Apply Filter**.

The Performance: All Aggregates view is redisplayed with the list of aggregates that are available on this node.

c. Click the **Utilization** column to sort the aggregates by least utilized.

You can also identify those aggregates that have the greatest **Free Capacity**. This provides a list of potential aggregates to which you might want to move workloads.

d. Write down the name of the aggregate to which you want to move the workloads.

3. Identify the high-performing workloads from the node that received the event:

a. Return to the **Event** details page for the event.

b. In the **Affected Volumes** field, click the link for the number of volumes.

The Performance: All Volumes view is displayed with a filtered list of the volumes on that node.

c. Click the **Total Capacity** column to sort the volumes by the largest allocated space.

This provides a list of potential volumes that you may want to move.

d. Write down the names of the volumes that you want to move, and the names of the current aggregates on which they reside.

4. Move the volumes to the aggregates you identified as having low utilization on the new node.

You can perform the move operation by using ONTAP System Manager, OnCommand Workflow Automation, ONTAP commands, or a combination of these tools.

After a few days, check to see whether you are receiving the same type of events from this node or aggregate.

Move workloads to a node in a different HA pair in Active IQ Unified Manager

You can use Unified Manager to help identify an aggregate on a node in a different high-availability (HA) pair that has more free performance capacity than the HA pair where your workloads are currently running. Then you can move selected volumes to aggregates on the new HA pair.

Before you begin

- You must have the Operator, Application Administrator, or Storage Administrator role.
- Your cluster must consist of a minimum of two HA pairs

  You cannot use this remediation process if you have only one HA pair in your cluster.

- You must have recorded the names of the two nodes in the HA pair that are currently having a performance issue.
- You must have recorded the date and time at which the nodes received the performance event.
- Unified Manager must have collected and analyzed performance data for a month or longer.

Moving high-performing workloads to an aggregate on a node with more free performance capacity allows workloads on both nodes to perform more efficiently. This procedure helps you to identify the following resources so that you can move high-performing workloads to a node that has more free performance capacity on a different HA pair:

- The nodes in a different HA pair on the same cluster that have the greatest free performance capacity
- The aggregates on the new nodes that have the greatest free performance capacity
- The highest-performing volumes on the current nodes

Steps

1. Identify the nodes that are part of a different HA pair on the same cluster:
   a. On the **Event Details** page, click the name of the cluster on which the nodes reside.

      The cluster details are displayed in the Performance/Cluster Landing page.

   b. On the **Summary** page, click **Nodes** from the **Managed Objects** pane.

      The list of nodes on this cluster is displayed in the Performance: All Nodes view.

   c. Write down the names of the nodes that are in different HA pairs from the HA pair that is currently having a performance issue.
2. Identify a node in the new HA pair that has the greatest free performance capacity:
   a. On the **Performance: All Nodes** view, click the **Performance Capacity Used** column to sort the nodes by least percentage used.

      This provides a list of potential nodes to which you might want to move workloads.

   b. Write down the name of the node on a different HA pair to which you want to move the workloads.
3. Identify an aggregate on the new node that has the greatest free performance capacity:
   a. On the **Performance: All Nodes** view, click the node.

      The node details are displayed in the Performance/Node Explorer page.

   b. In the **View and Compare** menu, select **Aggregates on this Node**.

      The aggregates on this node are displayed in the grid.

   c. Click the **Performance Capacity Used** column to sort the aggregates by least used.

This provides a list of potential aggregates to which you might want to move workloads.

d. Write down the name of the aggregate to which you want to move the workloads.

4. Identify the high-performing workloads from the nodes that received the event:
   a. Return to the **Event** details page for the event.
   b. In the **Affected Volumes** field, click the link for the number of volumes for the first node.

      The Performance: All Volumes view is displayed with a filtered list of the volumes on that node.

   c. Click the **Total Capacity** column to sort the volumes by the largest allocated space.

      This provides a list of potential volumes that you might want to move.

   d. Write down the names of the volumes that you want to move, and the names of the current aggregates on which they reside.
   e. Perform steps 4c and 4d for the second node that was part of this event to identify possible volumes that you want to move from that node as well.
5. Move the volumes to the aggregates that you identified as having greatest free performance capacity on the new node.

   You can perform the move operation by using ONTAP System Manager, OnCommand Workflow Automation, ONTAP commands, or a combination of these tools.

After a few days, you can check whether you are receiving the same type of events from this node or aggregate.

Move workloads to another node in a different HA pair in Active IQ Unified Manager

You can use Unified Manager to help identify an aggregate on a node in a different HA pair that is less busy than the HA pair where your workloads are currently running. Then you can move selected volumes to aggregates on the new HA pair. Moving high-performing workloads to an aggregate on a less busy node allows workloads on both nodes to perform more efficiently.

Before you begin

- You must have the Operator, Application Administrator, or Storage Administrator role.
- Your cluster must consist of a minimum of two HA pairs; you cannot use this remediation process if you have only one HA pair in your cluster.
- You must have recorded the names of the two nodes in the HA pair that are currently having the performance issue.
- You must have recorded the date and time at which the nodes received the performance event.
- Unified Manager must have collected and analyzed a month or more of performance data.

These steps help you identify the following resources so that you can move high-performing workloads to a lower utilized node on a different HA pair:

- The nodes in a different HA pair on the same cluster that are less utilized
- The aggregates on the new nodes that are the least utilized

- The highest-performing volumes on the current nodes

Steps

1. Identify the nodes that are part of a different HA pair on the same cluster:
   a. In the left navigation pane, click **Storage** > **Clusters** and select **Performance** > **All Clusters** from the View menu.

      The Performance: All Clusters view is displayed.

   b. Click the number in the **Node Count** field for the current cluster.

      The Performance: All Nodes view is displayed.

   c. Write down the names of the nodes that are in different HA pairs from the HA pair that is currently having the performance issue.
2. Identify a node in the new HA pair that is the least utilized:
   a. Click the **Utilization** column to sort the nodes by least utilized.

      You can also identify those nodes that have the greatest **Free Capacity**. This provides a list of potential nodes to which you might want to move workloads.

   b. Write down the name of the node to which you want to move the workloads.
3. Identify an aggregate on the new node that is the least utilized:
   a. In the left navigation pane, click **Storage** > **Aggregates** and select **Performance** > **All Aggregates** from the View menu.

      The Performance: All Aggregates view is displayed.

   b. Click **Filtering**, select **Node** from the left drop-down menu, type the name of the node in the text field, and then click **Apply Filter**.

      The Performance: All Aggregates view is redisplayed with the list of aggregates that are available on this node.

   c. Click the **Utilization** column to sort the aggregates by least utilized.

      You can also identify those aggregates that have the greatest **Free Capacity**. This provides a list of potential aggregates to which you might want to move workloads.

   d. Write down the name of the aggregate to which you want to move the workloads.
4. Identify the high-performing workloads from the nodes that received the event:
   a. Return to the **Event** details page for the event.
   b. In the **Affected Volumes** field, click the link for the number of volumes for the first node.

      The Performance: All Volumes view is displayed with a filtered list of the volumes on that node.

   c. Click the **Total Capacity** column to sort the volumes by the largest allocated space.

      This provides a list of potential volumes that you might want to move.

   d. Write down the names of the volumes that you want to move, and the names of the current aggregates on which they reside.

e. Perform steps 4c and 4d for the second node that was part of this event to identify possible volumes that you want to move from that node as well.

5. Move the volumes to the aggregates you identified as having low utilization on the new node.

   You can perform the move operation by using ONTAP System Manager, OnCommand Workflow Automation, ONTAP commands, or a combination of these tools.

After a few days, check to see whether you are receiving the same type of events from this node or aggregate.

Use QoS policy settings to prioritize node workloads in Active IQ Unified Manager

You can set a limit on a QoS policy group to control the I/O per second (IOPS) or MBps throughput limit for the workloads it contains. If workloads are in a policy group with no set limit, such as the default policy group, or the set limit does not meet your needs, you can increase the set limit or move the workloads to a new or existing policy group that has the desired limit.

If a performance event on a node is caused by workloads overusing the node resources, the event description on the Event details page displays a link to the list of volumes involved. In the Performance/Volumes page, you can sort the affected volumes by IOPS and MBps to see which workloads have the highest usage that might have contributed to the event.

By assigning the volumes that are overusing the node resources to a more restrictive policy group setting, the policy group throttles the workloads to restrict their activity, which can reduce the use of the resources on that node.

You can use ONTAP System Manager or the ONTAP commands to manage policy groups, including the following tasks:

- Creating a policy group
- Adding or removing workloads in a policy group
- Moving a workload between policy groups
- Changing the throughput limit of a policy group

Remove inactive volumes and LUNs in Active IQ Unified Manager

When aggregate free space has been identified as an issue, you can search for unused volumes and LUNs and delete them from the aggregate. This can help to alleviate the low disk space issue.

If a performance event on an aggregate is caused by low disk space, there are a few ways you can determine which volumes and LUNs are no longer being used.

To identify unused volumes:

- On the Event details page, the **Affected Objects Count** field provides a link that displays the list of affected volumes.

  Click the link to display the volumes on the Performance: All Volumes view. From there you can sort the affected volumes by **IOPS** to see which volumes have not been active.

To identify unused LUNs:

1. From the Event details page, write down the name of the aggregate on which the event occurred.
2. In the left navigation pane, click **Storage** > **LUNs** and select **Performance** > **All LUNs** from the View menu.
3. Click **Filtering**, select **Aggregate** from the left drop-down menu, type the name of the aggregate in the text field, and then click **Apply Filter**.
4. Sort the resulting list of affected LUNs by **IOPS** to view the LUNs that are not active.

After you have identified the unused volumes and LUNs, you can use ONTAP System Manager or the ONTAP commands to delete those objects.

Understand how adding disks affects aggregate layout in Active IQ Unified Manager

You can add disks to an aggregate to increase the storage capacity and the performance of that aggregate. After adding the disks, you only see an improvement in performance after reconstructing the aggregate.

When you receive a system-defined threshold event on the Event details page, the event description text lists the name of the aggregate that is having the problem. You can add disks and reconstruct data on this aggregate.

The disks you add to the aggregate must already exist in the cluster. If the cluster does not have extra disks available, you might need to contact your administrator or purchase more disks. You can use ONTAP System Manager or the ONTAP commands to add disks to an aggregate.

Technical Report 3838: Storage Subsystem Configuration Guide

Set up a connection between a Unified Manager server and an external data provider

Set up a connection between Active IQ Unified Manager and an external data provider

A connection between a Unified Manager server and an external data provider enables you to send cluster performance data to an external server so that storage managers can chart the performance metrics using third-party software.

A connection between a Unified Manager server and an external data provider is established through the menu option labeled “External Data Provider” in the maintenance console.

Send Active IQ Unified Manager performance data to an external server

Unified Manager collects a variety of performance data from all the clusters that it is monitoring. You can send specific groups of data to an external server.

Depending on the performance data that you want to chart, you can choose to send one of the following groups of statistics:

Statistics group	Data included	Details
Performance Monitor	High-level performance statistics for the following objects: • LUNs • Volumes	This group provides total IOPS or latency for all LUNs and volumes in all monitored clusters. This group provides the smallest number of statistics.
Resource Utilization	Resource utilization statistics for the following objects: • Nodes • Aggregates	This group provides utilization statistics for the node and aggregate physical resources in all monitored clusters. It also provides the statistics collected in the Performance Monitor group.
Drill Down	Low-level read/write and per-protocol statistics for all tracked objects: • Nodes • Aggregates • LUNs • Volumes • Disks • LIFs • Ports/NICs	This group provides read/write and per-protocol breakdowns for all seven tracked object types in all monitored clusters. It also provides the statistics collected in the Performance Monitor group and in the Resource Utilization group. This group provides the largest number of statistics.

If the name of a cluster, or cluster object, is changed on the storage system, both the old and the new objects will contain performance data on the external server (called the “metric_path”). The two objects are not correlated as the same object. For example, if you change the name of a volume from “volume1_acct” to “acct_vol1”, you will see old performance data for the old volume, and new performance data for the new volume.

See Knowledge Base article 30096 for the list of all performance counters that can be sent to an external data provider.

Unified Manager performance counters that can be exported to an External Data Provider

Set up Graphite to receive performance data from Active IQ Unified Manager

Graphite is an open software tool for gathering and graphing performance data from computer systems. Your Graphite server and software must be configured correctly to receive statistical data from Unified Manager.

NetApp does not test or verify specific versions of Graphite or other third-party tools.

The Graphite server does not receive performance data for Volumes from Unified Manager.

After you have installed Graphite according to the installation instructions, you need to make the following changes to support statistical data transfer from Unified Manager:

- In the `/opt/graphite/conf/carbon.conf` file, the maximum number of files that can be created on the Graphite server per minute must be set to *200* (**`MAX_CREATES_PER_MINUTE = 200`**).

  Depending on the number of clusters in your configuration and the statistics objects you have selected to send, there might be thousands of new files that need to be created initially. At 200 files per minute it might take 15 minutes or longer before all metric files are initially created. After all the unique metric files have been created, this parameter is no longer relevant.

- If you are running Graphite on a server deployed using an IPv6 address, the value for LINE_RECEIVER_INTERFACE in the `/opt/graphite/conf/carbon.conf` file must be changed from "0.0.0.0" to "::" (**`LINE_RECEIVER_INTERFACE = ::`**)

- In the `/opt/graphite/conf/storage-schemas.conf` file, the `retentions` parameter must be used to set the frequency to 5 minutes and the retention period to the number of days that is relevant for your environment.

  The retention period can be as long as what your environment allows, but the frequency value must be set to 5 minutes for at least one retention setting. In the following example, a section is defined for Unified Manager using the `pattern` parameter, and the values set the initial frequency to 5 minutes and the retention period to 100 days: **`[OPM]`**

  ```
  pattern = ^netapp-performance\..

  retentions = 5m:100d
  ```

  > If the default vendor tag is changed from "netapp-performance" to something different, that change must be reflected in the `pattern` parameter as well.

> If the Graphite server is unavailable when the Unified Manager server is attempting to send performance data, the data is not sent and there will be a gap in collected data.

Configure a connection between Active IQ Unified Manager and an external data provider

Unified Manager can send cluster performance data to an external server. You can specify the type of statistical data that is sent, and the interval at which data is sent.

Before you begin

- You must have a user ID authorized to log in to the maintenance console of the Unified Manager server.
- You must have the following information about the external data provider:
  - Server name or IP address (IPv4 or IPv6)
  - Server default port (if not using default port 2003)
- You must have configured the remote server and third-party software so that it can receive statistical data from the Unified Manager server.
- You must know which group of statistics you want to send:
  - PERFORMANCE_INDICATOR: Performance monitor statistics

- RESOURCE_UTILIZATION: Resource utilization and Performance monitor statistics
- DRILL_DOWN: All statistics

- You must know the time interval at which you want to transmit statistics: 5, 10, or 15 minutes

  By default, Unified Manager collects statistics at 5-minute intervals. If you set the transmit interval to 10 (or 15) minutes, the amount of data that is sent during each transmission is two (or three) times larger than when using the default 5-minute interval.

  If you change the Unified Manager performance collection interval to 10 or 15 minutes, you must change the transmit interval so that it is equal to, or larger, than the Unified Manager collection interval.

You can configure a connection between one Unified Manager server and one external data provider server.

Steps

1. Log in as the maintenance user to the maintenance console of the Unified Manager server.

   The Unified Manager maintenance console prompts are displayed.

2. In the maintenance console, type the number of the **External Data Provider** menu option.

   The External Server Connection menu is displayed.

3. Type the number of the **Add/Modify Server Connection** menu option.

   The current server connection information is displayed.

4. When prompted, type `y` to continue.
5. When prompted, enter the IP address or name of the destination server and the server port information (if different from the default port 2003).
6. When prompted, type `y` to verify that the information you entered is correct.
7. Press any key to return to the External Server Connection menu.
8. Type the number of the **Modify Server Configuration** menu option.

   The current server configuration information is displayed.

9. When prompted, type `y` to continue.
10. When prompted, enter the type of statistics to send, the time interval at which the statistics are sent, and whether you want to enable the transmission of statistics now:

For..	Enter…
Statistics group ID	`0` - PERFORMANCE_INDICATOR (default) `1` - RESOURCE_UTILIZATION `2` - DRILL_DOWN

For..	Enter…
Vendor tag	A descriptive name for the folder where the statistics will be stored on the external server. “netapp-performance” is the default name, but you can enter another value. By using dotted notation you can define a hierarchical folder structure. For example, by entering **`stats.performance.netapp`** the statistics will be located in **stats** > **performance** > **netapp**.
Transmit interval	**`5`** (default), **`10`**, or **`15`** minutes
Enable/disable	**`0`** - Disable **`1`** - Enable (default)

11. When prompted, type **`y`** to verify that the information you entered is correct.
12. Press any key to return to the External Server Connection menu.
13. Type **`x`** to exit the maintenance console.

After you have configured the connection, the selected performance data is sent to the destination server at the time interval you specified. It takes a few minutes before the metrics start to appear in the external tool. You might need to refresh your browser to see the new metrics in the metric hierarchy.

Monitor and manage cluster health

Introduction to Active IQ Unified Manager health monitoring

Learn about Active IQ Unified Manager health monitoring

Active IQ Unified Manager (formerly OnCommand Unified Manager) helps you to monitor a large number of systems running ONTAP software through a centralized user interface. The Unified Manager server infrastructure delivers scalability, supportability, and enhanced monitoring and notification capabilities.

The key capabilities of Unified Manager include monitoring, alerting, managing availability and capacity of clusters, managing protection capabilities, and bundling of diagnostic data and sending it to technical support.

You can use Unified Manager to monitor your clusters. When issues occur in the cluster, Unified Manager notifies you about the details of such issues through events. Some events also provide you with a remedial action that you can take to rectify the issues. You can configure alerts for events so that when issues occur, you are notified through email, and SNMP traps.

You can use Unified Manager to manage storage objects in your environment by associating them with annotations. You can create custom annotations and dynamically associate clusters, storage virtual machines (SVMs), and volumes with the annotations through rules.

You can also plan the storage requirements of your cluster objects using the information provided in the capacity and health charts, for the respective cluster object.

Physical and logical capacity

Unified Manager makes use of the concepts of physical and logical space used for ONTAP storage objects.

- Physical capacity: Physical space refers to the physical blocks of storage used in the volume. “Physical used capacity” is typically smaller than logical used capacity due to the reduction of data from storage efficiency features (such as deduplication and compression).
- Logical capacity: Logical space refers to the usable space (the logical blocks) in a volume. Logical space refers to how theoretical space can be used, without accounting for results of deduplication or compression. “Logical space used” is physical space used plus the savings from storage efficiency features (such as deduplication and compression) that have been configured. This measurement often appears larger than the physical used capacity because this does not reflect the data compression and other reductions in the physical space. Thus, the total logical capacity could be higher than the provisioned space.

Capacity measurement units

Unified Manager calculates storage capacity based on binary units of 1024 (2^{10}) bytes. In ONTAP 9.10.0 and earlier, these units were displayed as KB, MB, GB, TB, and PB. Beginning with ONTAP 9.10.1, they are displayed in Unified Manager as KiB, MiB, GiB, TiB, and PiB.



The units used for throughput continue to be kilobytes per second (Kbps), Megabytes per second (Mbps), Gigabytes per second (Gbps), or Terabytes per second (Tbps) and so forth, for all releases of ONTAP.

Capacity unit displayed in Unified Manager for ONTAP 9.10.0 and earlier	Capacity unit displayed in Unified Manager for ONTAP 9.10.1	Calculation	Value in bytes
KB	KiB	1024	1024 bytes
MB	MiB	1024 * 1024	1,048,576 bytes
GB	GiB	1024 * 1024 * 1024	1,073,741,824 bytes
TB	TiB	1024 * 1024 * 1024 * 1024	1,099,511,627,776 bytes

Understand Active IQ Unified Manager health monitoring features

Unified Manager is built on a server infrastructure that delivers scalability, supportability, and enhanced monitoring and notification capabilities. Unified Manager supports monitoring of systems running ONTAP software.

Unified Manager includes the following features:

- Discovery, monitoring, and notifications for systems that are installed with ONTAP software:
  - Physical objects: nodes, disks, disk shelves, SFO pairs, ports, and Flash Cache
  - Logical objects: clusters, storage virtual machines (SVMs), aggregates, volumes, LUNs, namespaces, qtrees, LIFs, Snapshot copies, junction paths, NFS shares, SMB shares, user and group quotas, QoS policy groups, and initiator groups
  - Protocols: CIFS, NFS, FC, iSCSI, NVMe, and FCoE
  - Storage efficiency: SSD aggregates, Flash Pool aggregates, FabricPool aggregates, deduplication, and compression
  - Protection: SnapMirror relationships (synchronous and asynchronous) and SnapVault relationships
- Viewing the cluster discovery and monitoring status
- MetroCluster over FC and IP configurations: viewing and monitoring the configuration, issues, and connectivity status of the cluster components. MetroCluster switches and bridges for MetroCluster over FC configurations
- Enhanced alerts, events, and threshold infrastructure
- LDAP, LDAPS, SAML authentication, and local user support
- RBAC (for a predefined set of roles)
- AutoSupport and support bundle
- Enhanced dashboard to show capacity, availability, protection, and performance health of the environment
- Volume move interoperability, volume move history, and junction path change history
- Scope of Impact area that graphically displays the resources that are impacted for events such as Some Failed Disks, MetroCluster Aggregate Mirroring Degraded, and MetroCluster Spare Disks Left Behind events
- Possible Effect area that displays the effect of the MetroCluster events
- Suggested Corrective Actions area that displays the actions that can be performed to address events such as Some Failed Disks, MetroCluster Aggregate Mirroring Degraded, and MetroCluster Spare Disks Left Behind events

- Resources that Might be Impacted area that displays the resources that might be impacted for events such as for the Volume Offline event, the Volume Restricted event, and the Thin-Provisioned Volume Space At Risk event
- Support for SVMs with FlexVol or FlexGroup volumes
- Support for monitoring node root volumes
- Enhanced Snapshot copy monitoring, including computing reclaimable space and deleting Snapshot copies
- Annotations for storage objects
- Report creation and management of storage object information such as physical and logical capacity, utilization, space savings, performance, and related events
- Integration with OnCommand Workflow Automation to execute workflows

  WFA download

Use Active IQ Unified Manager interfaces to manage storage system health

These sections contain information about the two user interfaces that Active IQ Unified Manager provides for troubleshooting data storage capacity, availability, and protection issues. The two UIs are the Unified Manager web UI and the maintenance console.

If you want to use the protection features in Unified Manager, you must also install and configure OnCommand Workflow Automation (WFA).

Unified Manager web UI

The Unified Manager web UI enables an administrator to monitor and troubleshoot cluster issues relating to data storage capacity, availability, and protection.

These sections describe some common workflows that an administrator can follow to troubleshoot storage capacity, data availability, or protection issues displayed in the Unified Manager web UI.

Maintenance console

The Unified Manager maintenance console enables an administrator to monitor, diagnose, and address operating system issues, version upgrade issues, user access issues, and network issues related to the Unified Manager server itself. If the Unified Manager web UI is unavailable, the maintenance console is the only form of access to Unified Manager.

You can use this information for accessing the maintenance console and using it to resolve issues related to the functioning of the Unified Manager server.

Manage and monitor clusters and cluster object health

Understand cluster monitoring

Understand cluster monitoring in Active IQ Unified Manager

You can add clusters to the Unified Manager database to monitor clusters for availability, capacity, and other details, such as CPU usage, interface statistics, free disk space, qtree

usage, and chassis environmental.

Events are generated if the status is abnormal or when a predefined threshold is breached. If configured to do so, Unified Manager sends a notification to a specified recipient when an event triggers an alert.

Understand node root volumes in Active IQ Unified Manager

You can monitor the node root volume using Unified Manager. The best practice is that the node root volume should have sufficient capacity to prevent the node from going down.

When the used capacity of the node root volume exceeds 80 percent of the total node root volume capacity, the Node Root Volume Space Nearly Full event is generated. You can configure an alert for the event to get a notification. You can take appropriate actions to prevent the node from going down by using either ONTAP System Manager or the ONTAP CLI.

> The functionality of monitoring node root volumes is not available, if clusters run ONTAP 9.14.1 version or later.

Understand events and thresholds for node root aggregates in Active IQ Unified Manager

You can monitor the node root aggregate by using Unified Manager. The best practice is to thickly provision the root volume in the root aggregate to prevent the node from halting.

By default, capacity and performance events are not generated for root aggregates. Also, the threshold values used by Unified Manager are not applicable to the node root aggregates. Only a technical support representative can modify the settings for these events to be generated. When the settings are modified by the technical support representative, the capacity threshold values are applied to the node root aggregate.

You can take appropriate actions to prevent the node from halting by using either ONTAP System Manager or the ONTAP CLI.

> The functionality of monitoring node root aggregates is not available, if clusters run ONTAP 9.14.1 version or later.

Understand quorum and epsilon in ONTAP clusters in Active IQ Unified Manager

Quorum and epsilon are important measures of cluster health and function that together indicate how clusters address potential communications and connectivity challenges.

Quorum is a precondition for a fully functioning cluster. When a cluster is in quorum, a simple majority of nodes are healthy and can communicate with each other. When quorum is lost, the cluster loses the ability to accomplish normal cluster operations. Only one collection of nodes can have quorum at any one time because all of the nodes collectively share a single view of the data. Therefore, if two non-communicating nodes are permitted to modify the data in divergent ways, it is no longer possible to reconcile the data into a single data view.

Each node in the cluster participates in a voting protocol that elects one node master; each remaining node is a secondary. The master node is responsible for synchronizing information across the cluster. When quorum is formed, it is maintained by continual voting. If the master node goes offline and the cluster is still in quorum, a new master is elected by the nodes that remain online.

Because there is the possibility of a tie in a cluster that has an even number of nodes, one node has an extra fractional voting weight called epsilon. If the connectivity between two equal portions of a large cluster fails, the group of nodes containing epsilon maintains quorum, assuming that all of the nodes are healthy. For example, the following illustration shows a four-node cluster in which two of the nodes have failed. However, because one of the surviving nodes holds epsilon, the cluster remains in quorum even though there is not a simple majority of healthy nodes.



Epsilon is automatically assigned to the first node when the cluster is created. If the node that holds epsilon becomes unhealthy, takes over its high-availability partner, or is taken over by its high-availability partner, then epsilon is automatically reassigned to a healthy node in a different HA pair.

Taking a node offline can affect the ability of the cluster to remain in quorum. Therefore, ONTAP issues a warning message if you attempt an operation that will either take the cluster out of quorum or else put it one outage away from a loss of quorum. You can disable the quorum warning messages by using the cluster quorum-service options modify command at the advanced privilege level.

In general, assuming reliable connectivity among the nodes of the cluster, a larger cluster is more stable than a smaller cluster. The quorum requirement of a simple majority of half the nodes plus epsilon is easier to maintain in a cluster of 24 nodes than in a cluster of two nodes.

A two-node cluster presents some unique challenges for maintaining quorum. Two-node clusters use cluster HA, in which neither node holds epsilon; instead, both nodes are continuously polled to ensure that if one node fails, the other has full read-write access to data, as well as access to logical interfaces and management functions.

View the cluster list and details in Active IQ Unified Manager

You can use the Health: All Clusters view to view your inventory of clusters. The Capacity: All Clusters view enables you to view summarized information about storage capacity and utilization in all clusters.

Before you begin

You must have the Operator, Application Administrator, or Storage Administrator role.

You can also view details for individual clusters such as the cluster health, capacity, configuration, LIFs, nodes, and disks in that cluster by using the Cluster / Health details page.

The details in the Health: All Clusters view, Capacity: All Clusters view, and the Cluster / Health details page help you plan your storage. For example, before provisioning a new aggregate, you can select a specific cluster from the Health: All Clusters view and obtain capacity details to determine if the cluster has the required space.

Steps

1. In the left navigation pane, click **Storage** > **Clusters**.
2. In the View menu, select **Health: All Clusters** view to view health information, or **Capacity: All Clusters**

view to view details about storage capacity and utilization in all clusters.

3. Click the name of a cluster to view complete details of the cluster in the **Cluster / Health** details page.

Related information

- Cluster / Health details page
- Performance: All Clusters view
- Monitoring MetroCluster configurations
- Viewing security status for clusters and Storage VMs
- What security criteria are being evaluated

Check the health of clusters in a MetroCluster configuration in Active IQ Unified Manager

You can use Active IQ Unified Manager (Unified Manager) to check the operational health of clusters, and their components, in MetroCluster over FC and MetroCluster over IP configurations. If the clusters were involved in a performance event detected by Unified Manager, the health status can help you determine whether a hardware or software issue contributed to the event.

Before you begin

- You must have the Operator, Application Administrator, or Storage Administrator role.
- You must have analyzed a performance event for a MetroCluster configuration and obtained the name of the cluster involved.
- Both clusters in the MetroCluster configuration over FC and IP must be monitored by the same instance of Unified Manager.

Determine cluster health in MetroCluster over FC configuration

Follow these steps for determining cluster health in a MetroCluster over FC configuration.

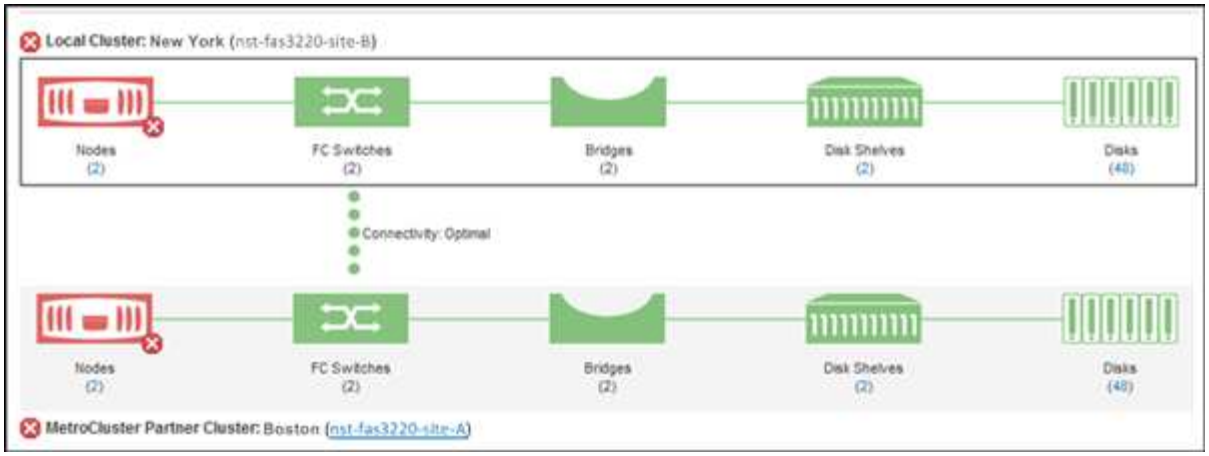
Steps

1. In the left navigation pane, click **Event Management** to display the event list.
2. In the filter panel, select all MetroCluster filters under the **Source Type** category. You see all events raised in your environment for all MetroCluster configurations.
3. Next to a MetroCluster event, click the name of the cluster.

   If no MetroCluster events are displayed, you can use the Search bar to search for the name of the cluster involved in the event related to you MetroCluster over FC configuration.

   The Health: All Clusters view is displayed with detailed information about the event.

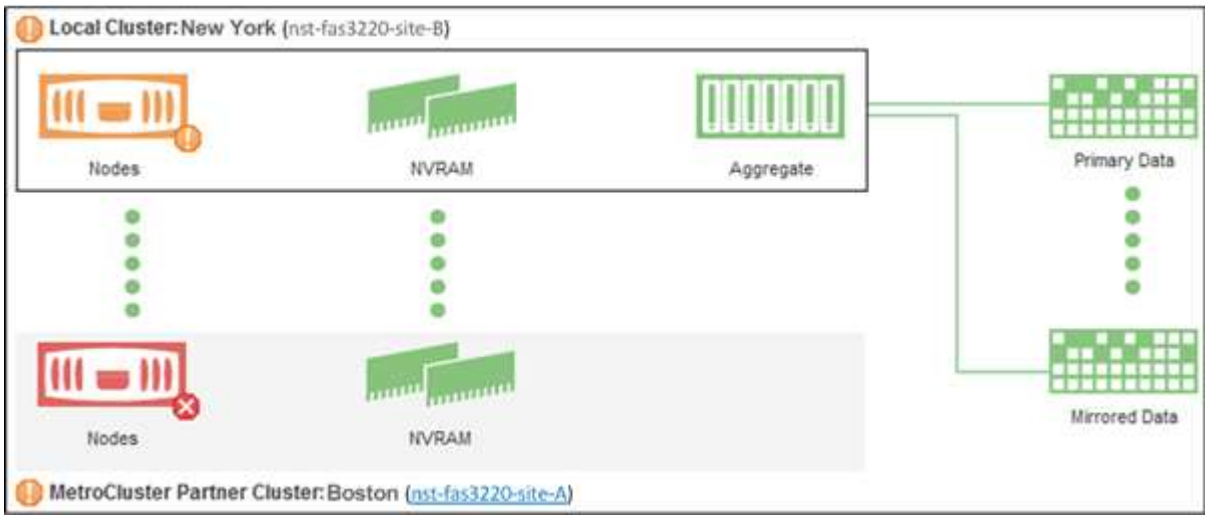
4. Select the **MetroCluster Connectivity** tab to display the health of the connection between the selected cluster and its partner cluster.



In this example, the names and the components of the local cluster and its partner cluster are displayed. A yellow or red icon indicates a health event for the highlighted component. The Connectivity icon represents the link between the clusters. You can point your mouse cursor to an icon to display event information or click the icon to display the events. A health issue on either cluster might have contributed to the performance event.

Unified Manager monitors the NVRAM component of the link between the clusters. If the FC Switches icon on the local or partner cluster or the Connectivity icon is red, a link health issue might have caused the performance event.

5. Select the **MetroCluster Replication** tab.



In this example, if the NVRAM icon on the local or partner cluster is yellow or red, a health issue with the NVRAM might have caused the performance event. If there are no red or yellow icons on the page, a performance issue on the partner cluster might have caused the performance event.

Determine cluster health in MetroCluster over IP configuration

Follow these steps for determining cluster health in a MetroCluster over IP configuration.

Steps

1. In the left navigation pane, click **Event Management** to display the event list.
2. In the filter panel, under the **Source Type** category, select the `MetroCluster Relationship` filter. You see all events raised in your environment for all MetroCluster configurations.

If you cannot see the reported MetroCluster events, you can use the Search bar to search by the name of the cluster involved in the event related to you MetroCluster over IP configuration.

3. Next to the relevant MetroCluster event, click the name of the cluster. The Clusters page is displayed with the details of that cluster.
   For information about determining health issues, see Monitor connectivity issues in MetroCluster over IP configuration.

View the health and capacity status of All SAN Array clusters in Active IQ Unified Manager

You can use the Cluster inventory pages to display the health and capacity status of your All SAN Array clusters.

Before you begin

You must have the Operator, Application Administrator, or Storage Administrator role.

You can view overview information for All SAN Array clusters in the Health: All Clusters view and Capacity: All Clusters view. Additionally, you can view details in the Cluster / Health details page.

Steps

1. In the left navigation pane, click **Storage** > **Clusters**.
2. Make sure that the “Personality” column is displayed in the **Health: All Clusters** view, or add it using the **Show / Hide** control.

   This column displays “All SAN Array” for your All SAN Array clusters.
3. Review the information.
4. To view information about storage capacity in those clusters, select the Capacity: All Clusters view.
5. To view detailed information about health and storage capacity in those clusters, click the name of an All SAN Array cluster.

   View the details in the Health, Capacity, and Nodes tabs in the Cluster / Health details page

View the node list and details in Active IQ Unified Manager

You can use the Health: All Nodes view to view the list of nodes in your clusters. You can use the Cluster / Health details page to view detailed information about nodes that are part of the cluster that is monitored.

Before you begin

You must have the Operator, Application Administrator, or Storage Administrator role.

You can view details such as the node state, cluster that contains the node, aggregate capacity details (used and total), and raw capacity details (usable, spare, and total). You can also obtain information about HA pairs, disks shelves, and ports.

Steps

1. In the left navigation pane, click **Storage** > **Nodes**.

2. On the **Health: All Nodes** view, click the node whose details you want to view.

   The detailed information for the selected node is displayed in the Cluster / Health details page. The left pane displays the list of HA pairs. By default, the HA Details is open, which displays HA state details and events related to the selected HA pair.

3. To view other details about the node, perform the appropriate action:

To view…	Click…
Details about the disk shelves	**Disk Shelves**.
Port-related information	**Ports**.

For more information, see:

- Performance: All Nodes view
- Viewing node and aggregate available IOPS values
- Viewing node and aggregate performance capacity used values

Generate a hardware inventory report for contract renewal in Active IQ Unified Manager

You can generate a report that contains a complete list of cluster and node information; such as hardware model numbers and serial numbers, disk types and counts, installed licenses, and more. This report is helpful for contract renewal within secure sites (“dark” sites) that are not connected to the NetAppActive IQ platform.

Before you begin

You must have the Operator, Application Administrator, or Storage Administrator role.

Steps

1. In the left navigation pane, click **Storage** > **Nodes**.
2. Go to the **Health: All Nodes** view or **Performance: All Nodes** view.
3. Select **Reports** > *** >** Hardware Inventory Report*.

   The hardware inventory report is downloaded as a .csv file with complete information as of the current date.

4. Provide this information to your NetApp support contact for contract renewal.

View the Storage VM list and details in Active IQ Unified Manager

From the Health: All Storage VMs view, you can monitor your inventory of storage virtual machines (SVMs). You can use the Storage VM / Health details page to view detailed information about SVMs that are monitored.

Before you begin

You must have the Operator, Application Administrator, or Storage Administrator role.

You can view SVM details, such as the capacity, efficiency, and configuration of an SVM. You can also view

information about the related devices and related alerts for that SVM.

Steps

1. In the left navigation pane, click **Storage** > **Storage VMs**.
2. Choose one of the following ways to view the SVM details:
   - To view information about the health of all SVMs in all clusters, in the View menu, select Health: All Storage VMs view.
   - To view the complete details, click the Storage VM name.

     You can also view the complete details by clicking **View Details** in the minimal details dialog box.
3. View the objects related to the SVM by clicking **View Related** in the minimal details dialog box.

Related information

- Storage VM: Health details page
- Performance: All Storage VMs view
- Security: Anti-ransomware view
- Viewing security status for clusters and Storage VMs
- Relationship: All Relationships view

View the aggregate list and details in Active IQ Unified Manager

From the Health: All Aggregates view, you can monitor your inventory of aggregates. The Capacity: All Aggregates view enables you to view information about the capacity and utilization of aggregates in all clusters.

Before you begin

You must have the Operator, Application Administrator, or Storage Administrator role.

You can view details such as aggregate capacity and configuration, and disk information from the Aggregate / Health details page. You can use these details before you configure the threshold settings if required.

Steps

1. In the left navigation pane, click **Storage** > **Aggregates**.
2. Choose one of the following ways to view the aggregate details:
   - To view information about the health of all aggregates in all clusters, in the View menu, select Health: All Aggregates view.
   - To view information about the capacity and utilization of all aggregates in all clusters, in the View menu, select Capacity: All Aggregates view.
   - To view the complete details, click the aggregate name.

     You can also view the complete details by clicking **View Details** in the minimal details dialog box.
3. View the objects related to the aggregate by clicking **View Related** from the minimal details dialog box.

Related information

- Aggregate / Health details page
- Performance: All Aggregates view
- Create a report to view aggregates reaching full capacity in Active IQ Unified Manager

View FabricPool capacity information in Active IQ Unified Manager

You can view FabricPool capacity information for clusters, aggregates, and volumes on the Capacity and Performance inventory and details pages for these objects. These pages also display FabricPool mirror information when a mirror tier has been configured.

These pages display information such as the available capacity on the local performance tier and on the cloud tier, how much capacity is being used in both tiers, which aggregates are attached to a cloud tier, and which volumes are implementing the FabricPool features by moving certain information to the cloud tier.

When a cloud tier is mirrored to another cloud provider (the “mirror tier”) then both cloud tiers are displayed in the Aggregate / Health details page.

Steps

1. Perform one of the following:

To view capacity information for…	Do this…
Clusters	a. On the Capacity: All Clusters view, click a cluster. b. On the Cluster / Health details page, click the **Configuration** tab. The display shows the names of any cloud tiers to which this cluster is connected.
Aggregates	a. On the Capacity: All Aggregates view, click an aggregate where the Type field indicates “SSD (FabricPool)” or “HDD (FabricPool)”. b. On the Aggregate / Health details page, click the **Capacity** tab. The display shows the total capacity used in the cloud tier. c. Click the **Disk Information** tab. The display shows the name of the cloud tier and the capacity used. d. Click the **Configuration** tab. The display shows the name of the cloud tier and other detailed information about the cloud tier.

To view capacity information for…	Do this…
Volumes	a. On the Capacity: All Volumes view, click a volume where a policy name appears in the “Tiering Policy” field. b. On the Volume / Health details page, click the **Configuration** tab. The display shows the name of the FabricPool tiering policy assigned to the volume.

2. In the **Workload Analysis** page you can select “Cloud Tier View” in the **Capacity Trend** area to see the capacity being used in the local Performance Tier and in the Cloud Tier over the previous month.

For more information on FabricPool aggregates, see Disks and aggregates overview.

View storage pool details in Active IQ Unified Manager

You can view the details of the storage pool to monitor the storage pool health, total and available cache, and used and available allocations.

Before you begin

You must have the Operator, Application Administrator, or Storage Administrator role.

Steps

1. In the left navigation pane, click **Storage** > **Aggregates**.
2. Click an aggregate name.

   The details of the selected aggregate are displayed.
3. Click the **Disk Information** tab.

   Detailed disk information is displayed.

   The Cache table is displayed only when the selected aggregate is using a storage pool.
4. In the Cache table, move the pointer over the name of the required storage pool.

   The details of the storage pool are displayed.

View the volume list and details in Active IQ Unified Manager

From the Health: All Volumes view, you can monitor your inventory of volumes. The Capacity: All Volumes view enables you to view information about the capacity and utilization of volumes in a cluster.

Before you begin

You must have the Operator, Application Administrator, or Storage Administrator role.

You can also use the Volume / Health details page to view detailed information about volumes that are monitored, including the capacity, efficiency, configuration, and protection of the volumes. You can also view

information about the related devices and related alerts for a specific volume.

Steps

1. In the left navigation pane, click **Storage** > **Volumes**.
2. Choose one of the following ways to view the volume details:
   - To view detailed information about the health of volumes in a cluster, in the View menu, select Health: All Volumes view.
   - To view detailed information about the capacity and utilization of volumes in a cluster, in the View menu, select Capacity: All Volumes view.
   - To view the complete details, click the volume name.

     You can also view the complete details by clicking **View Details** in the minimal details dialog box.
3. **Optional:** View the objects related to the volume by clicking **View Related** from the minimal details dialog box.

Related information

- Volume: Health details page
- Performance: All Volumes view
- Security: Anti-ransomware view
- Viewing volume protection relationships
- Creating a report to view available volume capacity charts

View details about NFS shares in Active IQ Unified Manager

You can view details about all NFS shares, such as its status, the path associated with the volume (FlexGroup volumes or FlexVol volumes), access levels of clients to the NFS shares, and the export policy defined for the volumes that are exported. Use the Health: All NFS Shares view to see all NFS shares on all monitored clusters, and use the Storage VM / Health details page to view all NFS shares on a specific storage virtual machine (SVM).

Before you begin

- NFS license must be enabled on the cluster.
- Network interfaces serving the NFS shares must be configured.
- You must have the Operator, Application Administrator, or Storage Administrator role.

Step

1. In the left navigation pane, follow the steps below depending on whether you want to view all NFS shares or just the NFS shares for a particular SVM.

To…	Follow these steps…
View all NFS shares	Click **Storage** > **NFS Shares**

To…	Follow these steps…
View NFS shares for single SVM	a. Click **Storage** > **Storage VMs** b. Click the SVM for which you want to view the NFS shares details. c. In the Storage VM / Health details page, click the **NFS Shares** tab.

For more information, see Provisioning file share volumes and Provisioning CIFS and NFS file shares by using APIs.

View details about SMB/CIFS shares in Active IQ Unified Manager

You can view details about all SMB/CIFS shares, such as the share name, junction path, containing objects, security settings, and export policies defined for the share. Use the Health: All SMB Shares view to see all SMB shares on all monitored clusters, and use the Storage VM / Health details page to view all SMB shares on a specific storage virtual machine (SVM).

Before you begin

- CIFS license must be enabled on the cluster.
- Network interfaces serving the SMB/CIFS shares must be configured.
- You must have the Operator, Application Administrator, or Storage Administrator role.

Shares in folders are not displayed.

Step

1. In the left navigation pane, follow the steps below depending on whether you want to view all SMB/CIFS shares or just the shares for a particular SVM.

To…	Follow these steps…
View all SMB/CIFS shares	Click **Storage** > **SMB Shares**
View SMB/CIFS shares for single SVM	a. Click **Storage** > **Storage VMs** b. Click the SVM for which you want to view the SMB/CIFS share details. c. In the Storage VM / Health details page, click the **SMB Shares** tab.

For more information, see Provisioning CIFS and NFS file shares by using APIs.

View the list of Snapshot copies in Active IQ Unified Manager

You can view the list of Snapshot copies for a selected volume. You can use the list of Snapshot copies to calculate the amount of disk space that can be reclaimed if one or more Snapshot copies are deleted, and you can delete the Snapshot copies if required.

Before you begin

- You must have the Operator, Application Administrator, or Storage Administrator role.
- The volume containing the Snapshot copies must be online.

Steps

1. In the left navigation pane, click **Storage** > **Volumes**.
2. In the **Health: All Volumes** view, select the volume that contains the Snapshot copies you want to view.
3. In the **Volume / Health** details page, click the **Capacity** tab.
4. In **Details** pane of the **Capacity** tab, in the Other Details section, click the link next to **Snapshot Copies**.

   The number of Snapshot copies is a link that displays the list of Snapshot copies.

Related information

Health/Volumes page

Delete Snapshot copies in Active IQ Unified Manager

You can delete a Snapshot copy to conserve space or to free disk space, or you can delete the Snapshot copy if it is no longer required.

Before you begin

You must have the Application Administrator or Storage Administrator role.

The volume must be online.

To delete a Snapshot copy that is busy or locked, you must have released the Snapshot copy from the application that was using it.

- You cannot delete the base Snapshot copy in a parent volume if a FlexClone volume is using that Snapshot copy.

  The base Snapshot copy is the Snapshot copy that is used to create the FlexClone volume and displays the status `Busy` and Application Dependency as `Busy, Vclone` in the parent volume.

- You cannot delete a locked Snapshot copy that is used in a SnapMirror relationship.

  The Snapshot copy is locked and is required for the next update.

Steps

1. In the left navigation pane, click **Storage** > **Volumes**.
2. In the **Health: All Volumes** view, select the volume that contains the Snapshot copies you want to view.

   The list of Snapshot copies is displayed.

3. In the **Volume / Health** details page, click the **Capacity** tab.
4. In **Details** pane of the **Capacity** tab, in the Other Details section, click the link next to **Snapshot Copies**.

   The number of Snapshot copies is a link that displays the list of Snapshot copies.

5. In the **Snapshot Copies** view, select the Snapshot copies you want to delete, and then click **Delete Selected**.

Calculate reclaimable space for Snapshot copies in Active IQ Unified Manager

You can calculate the amount of disk space that can be reclaimed if one or more Snapshot copies are deleted.

Before you begin

- You must have the Operator, Application Administrator, or Storage Administrator role.
- The volume must be online.
- The volume must be a FlexVol volume; this capability is not supported with FlexGroup volumes.

Steps

1. In the left navigation pane, click **Storage** > **Volumes**.
2. In the **Health: All Volumes** view, select the volume that contains the Snapshot copies you want to view.

   The list of Snapshot copies is displayed.

3. In the **Volume / Health** details page, click the **Capacity** tab.
4. In **Details** pane of the **Capacity** tab, in the Other Details section, click the link next to **Snapshot Copies**.

   The number of Snapshot copies is a link that displays the list of Snapshot copies.

5. In the **Snapshot Copies** view, select the Snapshot copies for which you want to calculate the reclaimable space.
6. Click **Calculate**.

   The reclaimable space (in percentage, and KB, MB, GB, and so on) on the volume is displayed.

7. To recalculate the reclaimable space, select the required Snapshot copies and click **Recalculate**.

Understand cluster object windows and dialog boxes in Active IQ Unified Manager

You can view all your clusters and cluster objects from the respective storage object page. You can also view the details from the corresponding storage object details page. You can now launch System Manager user interface from the following STORAGE and PROTECTION sections of the INVENTORY.

- Cluster Inventory, Cluster Health, and Cluster Performance pages
- Aggregate Inventory, Aggregate Health, and Aggregate Performance pages
- Volume Inventory, Volume Health, and Volume Performance pages
- Node Inventory and Node Performance pages
- StorageVM Inventory, StorageVM Health, and StorageVM Performance pages
- Protection relationship pages

Common Unified Manager health workflows and tasks

Understand common Active IQ Unified Manager health workflows and tasks

Some common administrative workflows and tasks associated with Unified Manager include selecting the storage clusters that are to be monitored; diagnosing conditions that adversely affect data availability, capacity, and protection; restoring lost data; configuring and managing volumes; and bundling and sending diagnostic data to technical support (when necessary).

Unified Manager enables storage administrators to view a dashboard, assess the overall capacity, availability, and protection health of the managed storage clusters, and then quickly identify, locate, diagnose, and assign for resolution any specific issues that might arise.

The most important issues related to a cluster, storage virtual machine (SVM), volume, or FlexGroup volume that affect the storage capacity or data availability of your managed storage objects are displayed in the system health graphs and events on the Dashboard page. When critical issues are identified, this page provides links to support appropriate troubleshooting workflows.

Unified Manager can also be included in workflows that include related manageability tools—such as OnCommand Workflow Automation (WFA)--to support the direct configuration of storage resources.

Common workflows related to the following administrative tasks are described in this document:

- Diagnosing and managing availability issues

  If hardware failure or storage resource configuration issues cause the display of data availability events in the Dashboard page, storage administrators can follow the embedded links to view connectivity information about the affected storage resource, view troubleshooting advice, and assign issue resolution to other administrators.

- Configuring and monitoring performance incidents

  The Administrator can monitor and manage the performance of the storage system resources that are being monitored. See the Introduction to Active IQ Unified Manager performance monitoring for more information.

- Diagnosing and managing volume capacity issues

  If volume storage capacity issues are displayed in the Dashboard page, storage administrators can follow the embedded links to view the current and historical trends related to the storage capacity of the affected volume, view troubleshooting advice, and assign issue resolution to other administrators.

- Configuring, monitoring, and diagnosing protection relationship issues

  After creating and configuring protection relationships, storage administrators can view the potential issues related to protection relationships, the current state of the protection relationships, the current and historical protection job success information about the affected relationships, and troubleshooting advice. See the Learn about managing and monitoring protection relationships for more information.

- Creating backup files and restoring data from backup files.
- Associating storage objects with annotations

  By associating storage objects with annotations, storage administrators can filter and view the events that are related to the storage objects, which enables storage administrators to prioritize and resolve the issues that are associated with the events.

- Using REST APIs to help manage your clusters by viewing the health, capacity, and performance information captured by Unified Manager. See Getting started with Active IQ Unified Manager REST APIs for more information.
- Sending a support bundle to technical support

  Storage administrators can retrieve and send a support bundle to technical support by using the maintenance console. Support bundles must be sent to technical support when the issue requires more detailed diagnosis and troubleshooting than what an AutoSupport message provides.

Monitor and troubleshoot data availability

Monitor and troubleshoot data availability in Active IQ Unified Manager

Unified Manager monitors the reliability with which authorized users can access your stored data, alerts you to conditions that block or impede that access, and enables you to diagnose those conditions and assign and track their resolution.

The availability workflow topics in this section describe examples of how a storage administrator can use the Unified Manager web UI to discover, diagnose, and assign for resolution hardware and software conditions that adversely affect data availability.

Scan for and resolve storage failover interconnect link down conditions

Scan for and resolve storage failover interconnect link down conditions in Active IQ Unified Manager

This workflow provides an example of how you might scan for, evaluate, and resolve downed storage failover interconnect link conditions. In this scenario, you are an administrator using Unified Manager to scan for storage failover risks before starting an ONTAP version upgrade on your nodes.

Before you begin

You must have the Operator, Application Administrator, or Storage Administrator role.

If storage failover interconnections between HA pair nodes fail during a nondisruptive upgrade attempt, the upgrade fails. Therefore, common practice is for the administrator to monitor and confirm storage failover reliability on the cluster nodes targeted for upgrade before the start of an upgrade.

Steps

1. In the left navigation pane, click **Event Management**.
2. In the **Event Management** inventory page, select **Active Availability events**.
3. At the top of the **Event Management** inventory page **Name** column, click and enter `*failover` in the text box to limit the event to display to storage failover-related events.

   All past events related to storage failover conditions are displayed.

   In this scenario, the Unified Manager displays the event, “Storage Failover Interconnect One or More Links Down” in its Availability Incidents section.

4. If one or more events related to storage failover are displayed on the **Event Management** inventory page, perform the following steps:

a. Click the event title link to display event details for that event.

   In this example, you click the event title “Storage Failover Interconnect One or More Links Down”.

   The Event details page for that event is displayed.

b. On the Event details page, you can perform one or more of the following tasks:
   - Review the error message in the Cause field and evaluate the issue.
   - Assign the event to an administrator.
   - Acknowledge the event.

Related information

Event details page

Unified Manager user roles and capabilities

Perform corrective action for a storage failover interconnect link failure in Active IQ Unified Manager

When you display the Event details page of a storage failover-related event, you can review the summary information of the page to determine the urgency of the event, possible cause of the issue, and possible resolution to the issue.

Before you begin

You must have the Operator, Application Administrator, or Storage Administrator role.

In this example scenario, the event summary provided on the Event details page contains the following information about the storage failover interconnect link down condition:

```
Event: Storage Failover Interconnect One or More Links Down

Summary

Severity: Warning
State: New
Impact Level: Risk
Impact Area: Availability
Source: aardvark
Source Type: Node
Acknowledged By:
Resolved By:
Assigned To:
Cause: At least one storage failover interconnected link
       between the nodes aardvark and bonobo is down.
       RDMA interconnect is up (Link0 up, Link1 down)
```

The example event information indicates that a storage failover interconnect link, Link1, between HA pair nodes aardvark and bonobo is down, but that Link0 between Apple and Boy is active. Because one link is

active, the remote dynamic memory access (RDMA) is still functioning and a storage failover job can still succeed.

However, to ensure against both links failing and storage failover protection being totally disabled, you decide to further diagnose the reason for Link1 going down.

Steps

1. From the **Event** details page, you can click the link to the event specified in the Source field to obtain further details of other events that might be related to the storage failover interconnection link down condition.

   In this example, the source of the event is the node named aardvark. Clicking that node name displays the HA Details for the affected HA pair, aardvark and bonobo, on the Nodes tab of the Cluster / Health details page, and displays other events that recently occurred on the affected HA pair.

2. Review the **HA Details** for more information relating to the event.

   In this example, the relevant information is in the Events table. The table shows the “Storage Failover Connection One or More Link Down” event, the time the event was generated, and, again, the node from which this event originated.

Using the node location information in the HA Details, request or personally complete a physical inspection and repair of the storage failover issue on the affected HA pair nodes.

Related information

Event details page

Unified Manager user roles and capabilities

Resolve volume offline issues

Resolve volume offline issues in Active IQ Unified Manager

This workflow provides an example of how you might evaluate and resolve a volume offline event that Unified Manager might display in the Event Management inventory page. In this scenario, you are an administrator using Unified Manager to troubleshoot one or more volume offline events.

Before you begin

You must have the Operator, Application Administrator, or Storage Administrator role.

Volumes might be reported offline for several reasons:

- The SVM administrator has deliberately taken the volume offline.
- The volume’s hosting cluster node is down and storage failover to its HA pair partner has failed also.
- The volume’s hosting storage virtual machine (SVM) is stopped because the node hosting the root volume of that SVM is down.
- The volume’s hosting aggregate is down due to simultaneous failure of two RAID disks.

You can use the Event Management inventory page and the Cluster/Health, Storage VM/Health, and Volume/Health details pages to confirm or eliminate one or more of these possibilities.

Steps

1. In the left navigation pane, click **Event Management**.
2. In the **Event Management** inventory page, select **Active Availability events**.
3. Click the hypertext link displayed for the Volume Offline event.

   The Event details page for the availability incident is displayed.

4. On that page, check the notes for any indication that the SVM administrator has taken the volume in question offline.
5. On the **Event** details page, you can review the information for one or more of the following tasks:
   - Review the information displayed in the Cause field for possible diagnostic guidance.

     In this example, the information in the Cause field informs you only that the volume is offline.

   - Check the Notes and Updates area for any indication that the SVM administrator has deliberately taken the volume in question offline.
   - Click the source of the event, in this case the volume that is reported offline, to get more information about that volume.
   - Assign the event to an administrator.
   - Acknowledge the event or, if appropriate, mark it as resolved.

Perform diagnostic actions for volume offline conditions in Active IQ Unified Manager

After navigating to the Volume / Health details page of a volume reported to be offline, you can search for additional information helpful to diagnosing the volume offline condition.

Before you begin

You must have the Operator, Application Administrator, or Storage Administrator role.

If the volume that is reported offline was not taken offline deliberately, that volume might be offline for several reasons.

Starting at the offline volume's Volume / Health details page, you can navigate to other pages and panes to confirm or eliminate possible causes:

- Click **Volume / Health** details page links to determine if the volume is offline because its host node is down and storage failover to its HA pair partner has failed also.

  See Determining if a volume offline condition is caused by a down node.

- Click **Volume / Health** details page links to determine if the volume is offline and its host storage virtual machine (SVM) is stopped because the node hosting the root volume of that SVM is down.

  See Determining if a volume is offline and SVM is stopped because a node is down.

- Click **Volume / Health** details page links to determine if the volume is offline because of broken disks in its host aggregate.

  See Determining if a volume is offline because of broken disks in an aggregate.

Related information

Unified Manager user roles and capabilities

Determine if a volume is offline because its host node is down in Active IQ Unified Manager

You can use the Unified Manager web UI to confirm or eliminate the possibility that a volume is offline because its host node is down and that storage failover to its HA pair partner is unsuccessful.

Before you begin

You must have the Operator, Application Administrator, or Storage Administrator role.

To determine if the volume offline condition is caused by failure of the hosting node and subsequent unsuccessful storage failover, perform the following actions:

Steps

1. Locate and click the hypertext link displayed under SVM in the **Related Devices** pane of the offline volume's **Volume / Health** details page.

   The Storage VM / Health details page displays information about the offline volume's hosting storage virtual machine (SVM).

2. In the **Related Devices** pane of the **Storage VM / Health** details page, locate and click hypertext link displayed under Volumes.

   The Health: All Volumes view displays a table of information about all the volumes hosted by the SVM.

3. On the **Health: All Volumes** view State column header, click the filter symbol , and then select the option **Offline**.

   Only the SVM volumes that are in offline state are listed.

4. On the Health: All Volumes view, click the grid symbol , and then select the option **Cluster Nodes**.

   You might need to scroll in the grid selection box to locate the **Cluster Nodes** option.

   The Cluster Nodes column is added to the volumes inventory and displays the name of the node that hosts each offline volume.

5. On the **Health: All Volumes** view, locate the listing for the offline volume and, in its Cluster Node column, click the name of its hosting node.

   The Nodes tab on the Cluster / Health details page displays the state of the HA pair of nodes to which the hosting node belongs. The state of the hosting node and the success of any cluster failover operation is indicated in the display.

After you confirm that the volume offline condition exists because its host node is down and storage failover to the HA pair partner has failed, contact the appropriate administrator or operator to manually restart the down node and fix the storage failover problem.

Determine if a volume is offline because its SVM is stopped in Active IQ Unified Manager

You can use the Unified Manager web UI to confirm or eliminate the possibility that a

volume is offline because its host storage virtual machine (SVM) is stopped due to the node hosting the root volume of that SVM being down.

Before you begin

You must have the Operator, Application Administrator, or Storage Administrator role.

To determine if the volume offline condition is caused its host SVM being stopped because the node hosting the root volume of that SVM is down, perform the following actions:

Steps

1. Locate and click the hypertext link displayed under the SVM in the **Related Devices** pane of the offline volume's **Volume / Health** details page.

   The Storage VM / Health details page displays the "running" or the "stopped" status of the hosting SVM. If the SVM status is running, then the volume offline condition is not caused by the node hosting the root volume of that SVM being down.

2. If the SVM status is stopped, then click **View SVMs** to further identify the cause of the hosting SVM being stopped.
3. On the **Health: All Storage VMs** view SVM column header, click the filter symbol and then type the name of the stopped SVM.

   The information for that SVM is shown in a table.

4. On the **Health: All Storage VMs** view, click and then select the option **Root Volume**.

   The Root Volume column is added to the SVM inventory and displays the name of the root volume of the stopped SVM.

5. In the Root Volume column, click the name of the root volume to display the **Storage VM / Health** details page for that volume.

   If the status of the SVM root volume is (Online), then the original volume offline condition is not caused because the node hosting the root volume of that SVM is down.

6. If the status of the SVM root volume is (Offline), then locate and click the hypertext link displayed under Aggregate in the Related Devices pane of the SVM root volume's Volume / Health details page.
7. Locate and click the hypertext link displayed under Node in the **Related Devices** pane of the Aggregate's **Aggregate / Health** details page.

   The Nodes tab on the Cluster / Health details page displays the state of the HA pair of nodes to which the SVM root volume's hosting node belongs. The state of the node is indicated in the display.

After you confirm that the volume offline condition is caused by that volume's host SVM offline condition, which itself is caused by the node that hosts the root volume of that SVM being down, contact the appropriate administrator or operator to manually restart the down node.

Determine if a volume is offline because of broken disks in Active IQ Unified Manager

You can use the Unified Manager web UI to confirm or eliminate the possibility that a volume is offline because RAID disk problems have taken its host aggregate offline.

Before you begin

You must have the Operator, Application Administrator, or Storage Administrator role.

To determine if the volume offline condition is caused by RAID disk problems that are taking the hosting aggregate offline, perform the following actions:

Steps

1. Locate and click the hypertext link displayed under Aggregate in the **Related Devices** pane of the **Volume / Health** details page.

   The Aggregate / Health details page displays the online or offline status of the hosting aggregate. If the aggregate status is online, then RAID disk problems are not the cause of the volume being offline.

2. If the aggregate status is offline, then click **Disk Information** and look for broken disk events in the **Events** list on the **Disk Information** tab.
3. To further identify the broken disks, click the hypertext link displayed under Node in the **Related Devices** pane.

   The Cluster / Health details page is displayed.

4. Click **Disks**, and then select **Broken** in the **Filters** pane to list all disks in the broken state.

   If the disks in the broken state caused the offline state of the host aggregate, the name of the aggregate is displayed in the Impacted Aggregate column.

After confirming that the volume offline condition is caused by broken RAID disks and the consequent offline host aggregate, contact the appropriate administrator or operator to manually replace the broken disks and put the aggregate back online.

Resolve capacity issues

Resolve capacity issues in Active IQ Unified Manager

This workflow provides an example of how you can resolve a capacity issue. In this scenario, you are an administrator or operator and you access the Unified ManagerDashboard page to see if any of the monitored storage objects have capacity issues. You want to determine the possible cause of and resolution to the problem.

Before you begin

You must have the Operator, Application Administrator, or Storage Administrator role.

On the Dashboard page, you look for a “Volume Space Full” error event in the Capacity panel under the events drop-down list.

Steps

1. In the **Capacity** panel of the **Dashboard** page, click the name of the Volume Space Full error event.

   The Event details page for the error is displayed.

2. From the **Event** details page, you can perform one or more of the following tasks:
   - Review the error message in the Cause field and click the suggestions under Suggested Remedial Actions to review descriptions of possible remediations.
   - Click the object name, in this case a volume, in the Source field to get details about the object.

- Look for notes that might have been added about this event.
- Add a note to the event.
- Assign the event to another user.
- Acknowledge the event.
- Mark the event as resolved.

Related information

Event details page

Perform suggested remedial actions for a full volume in Active IQ Unified Manager

After receiving a “Volume Space Full” error event, you review the suggested remedial actions on the Event details page and decide to perform one of the suggested actions.

Before you begin

You must have the Application Administrator or Storage Administrator role.

A user with any role can perform all of the tasks in this workflow that use Unified Manager.

In this example, you have seen a Volume Space Full error event on the Unified ManagerEvent Management inventory page and have clicked the name of the event.

Possible remedial actions you might perform for a full volume include the following:

- Enabling autogrow, deduplication, or compression on the volume
- Resizing or moving the volume
- Deleting or moving data from the volume

Although all of these actions must be performed from either ONTAP System Manager or the ONTAP CLI, you can use Unified Manager to find information you might need to determine which actions to take.

Steps

1. From the **Event** details page, you click the volume name in the Source field to view details about the affected volume.
2. On the **Volume / Health** details page, you click **Configuration** and see that deduplication and compression are already enabled on the volume.

   You decide to resize the volume.
3. In the **Related Devices** pane, you click the name of the hosting aggregate to see if the aggregate can accommodate a larger volume.
4. On the **Aggregate / Health** details page, you see that the aggregate hosting the full volume does have enough uncommitted capacity, so you use ONTAP System Manager to resize the volume, giving it more capacity.

Related information

Event details page

Manage health thresholds

Learn about storage capacity health thresholds in Active IQ Unified Manager

A storage capacity health threshold is the point at which the Unified Manager server generates events to report any capacity problem with storage objects. You can configure alerts to send notification whenever such events occurs.

The storage capacity health thresholds for all aggregates, volumes, and qtrees are set to default values. You can change the settings as required for an object or a group of objects.

Configure global health threshold settings

Configure global health threshold settings in Active IQ Unified Manager

You can configure global health threshold conditions for capacity, growth, Snapshot reserve, quotas, and inodes to monitor your aggregate, volume, and qtree size effectively. You can also edit the settings for generating events for exceeding lag thresholds.

Global health threshold settings apply to all objects with which they are associated, such as aggregates, volumes, and so forth. When thresholds are crossed, an event is generated and, if alerts are configured, an alert notification is sent. Threshold defaults are set to recommended values, but you can modify them to generate events at intervals to meet your specific needs. When thresholds are changed, events are generated or obsoleted in the next monitoring cycle.

Global health threshold settings are accessible from the Event Thresholds section of the left-navigation menu. You can also modify threshold settings for individual objects, from the inventory page or the details page for that object.

- For information, see Configuring global aggregate health threshold values.

  You can configure the health threshold settings for capacity, growth, and Snapshot copies for all aggregates to track any threshold breach.

- For information, see Configuring global volume health threshold values.

  You can edit the health threshold settings for capacity, Snapshot copies, qtree quotas, volume growth, overwrite reserve space, and inodes for all volumes to track any threshold breach.

- For information, see Configuring global qtree health threshold values.

  You can edit the health threshold settings for capacity for all qtrees to track any threshold breach.

- For information, see Editing lag health threshold settings for unmanaged protection relationships.

  You can increase or decrease the warning or error lag time percentage so that events are generated at intervals that are more appropriate to your needs.

Configure global aggregate health threshold values in Active IQ Unified Manager

You can configure global health threshold values for all aggregates to track any threshold breach. Appropriate events are generated for threshold breaches and you can take preventive measures based on these events. You can configure the global values based

on the best practice settings for thresholds that apply to all monitored aggregates.

Before you begin

You must have the Application Administrator or Storage Administrator role.

When you configure the options globally, the default values of the objects are modified. However, if the default values have been changed at the object level, the global values are not modified.

The threshold options have default values for better monitoring, however, you can change the values to suit the requirements of your environment.

When Autogrow is enabled on volumes that reside on the aggregate, the aggregate capacity thresholds are considered breached based on the maximum volume size set by autogrow, not based on the original volume size.

Health threshold values are not applicable to the root aggregate of the node.

Steps

1. In the left navigation pane, click **Event Thresholds** > **Aggregate**.
2. Configure the appropriate threshold values for capacity, growth, and Snapshot copies.
3. Click **Save**.

Related information

Adding users

Configure global volume health threshold values in Active IQ Unified Manager

You can configure the global health threshold values for all volumes to track any threshold breach. Appropriate events are generated for health threshold breaches, and you can take preventive measures based on these events. You can configure the global values based on the best practice settings for thresholds that apply to all monitored volumes.

Before you begin

You must have the Application Administrator or Storage Administrator role.

Most of the threshold options have default values for better monitoring. However, you can change the values to suit the requirements of your environment.

Note that when Autogrow is enabled on a volume that capacity thresholds are considered breached based on the maximum volume size set by autogrow, not based on the original volume size.

The default value of 1000 Snapshot copies is applicable only to FlexVol volumes when the ONTAP version is 9.4 or greater, and to FlexGroup volumes when the ONTAP version is 9.8 and greater. For clusters installed with older versions of ONTAP software, the maximum number is 250 Snapshot copies per volume. For these older versions, Unified Manager interprets this number 1000 (and any number between 1000 and 250) as 250; meaning you will continue to receive events when the number of Snapshot copies reaches 250. If you wish to set this threshold to less than 250 for these older versions, you must set the threshold to 250 or lower here, in the Health: All Volumes view, or in the Volume / Health details page.

Steps

1. In the left navigation pane, click **Event Thresholds** > **Volume**.
2. Configure the appropriate threshold values for capacity, Snapshot copies, qtree quotas, volume growth, and inodes.
3. Click **Save**.

Related information

Adding users

Configure global qtree health threshold values in Active IQ Unified Manager

You can configure the global health threshold values for all qtrees to track any threshold breach. Appropriate events are generated for health threshold breaches, and you can take preventive measures based on these events. You can configure the global values based on the best practice settings for thresholds that apply to all monitored qtrees.

Before you begin

You must have the Application Administrator or Storage Administrator role.

The threshold options have default values for better monitoring, however, you can change the values to suit the requirements of your environment.

Events are generated for a qtree only when a Qtree quota or a Default quota has been set on the qtree. Events are not generated if the space defined in a User quota or Group quota has exceeded the threshold.

Steps

1. In the left navigation pane, click **Event Thresholds** > **Qtree**.
2. Configure the appropriate capacity threshold values.
3. Click **Save**.

Configure lag thresholds for unmanaged protection relationships in Active IQ Unified Manager

You can edit the global default lag warning and error health threshold settings for unmanaged protection relationships so that events are generated at intervals that are appropriate to your needs.

Before you begin

You must have the Application Administrator or Storage Administrator role.

The lag time must be no more than the defined transfer schedule interval. For example, if the transfer schedule is hourly, then the lag time must not be more than one hour. The lag threshold specifies a percentage that the lag time must not exceed. Using the example of one hour, if the lag threshold is defined as 150%, then you will receive an event when the lag time is more than 1.5 hours.

The settings described in this task are applied globally to all unmanaged protection relationships. The settings cannot be specified and applied exclusively to one unmanaged protection relationship.

Steps

1. In the left navigation pane, click **Event Thresholds** > **Relationship**.

2. Increase or decrease the global default warning or error lag time percentage as required.
3. To disable a warning or error event from being triggered from any lag threshold amount, uncheck the box next to **Enabled**.
4. Click **Save**.

Related information

Adding users

Edit individual aggregate health threshold settings in Active IQ Unified Manager

You can edit the health threshold settings for aggregate capacity, growth, and Snapshot copies of one or more aggregates. When a threshold is crossed, alerts are generated and you receive notifications. These notifications help you to take preventive measures based on the event generated.

Before you begin

You must have the Application Administrator or Storage Administrator role.

Based on changes to the threshold values, events are generated or obsoleted in the next monitoring cycle.

When Autogrow is enabled on volumes that reside on the aggregate, the aggregate capacity thresholds are considered breached based on the maximum volume size set by autogrow, not based on the original volume size.

Steps

1. In the left navigation pane, click **Storage** > **Aggregates**.
2. In the **Health: All Aggregates** view, select one or more aggregates and then click **Edit Thresholds**.
3. In the **Edit Aggregate Thresholds** dialog box, edit the threshold settings of one of the following: capacity, growth, or Snapshot copies by selecting the appropriate check box and then modifying the settings.
4. Click **Save**.

Related information

Adding users

Edit individual volume health threshold settings in Active IQ Unified Manager

You can edit the health threshold settings for volume capacity, growth, quota, and space reserve of one or more volumes. When a threshold is crossed, alerts are generated and you receive notifications. These notifications help you to take preventive measures based on the event generated.

Before you begin

You must have the Application Administrator or Storage Administrator role.

Based on changes to the threshold values, events are generated or obsoleted in the next monitoring cycle.

Note that when Autogrow is enabled on a volume that capacity thresholds are considered breached based on the maximum volume size set by autogrow, not based on the original volume size.

The default value of 1000 Snapshot copies is applicable only to FlexVol volumes when the ONTAP version is 9.4 or greater, and to FlexGroup volumes when the ONTAP version is 9.8 and greater. For clusters installed with older versions of ONTAP software, the maximum number is 250 Snapshot copies per volume. For these older versions, Unified Manager interprets this number 1000 (and any number between 1000 and 250) as 250; meaning you will continue to receive events when the number of Snapshot copies reaches 250. If you wish to set this threshold to less than 250 for these older versions, you must set the threshold to 250 or lower here, in the Health: All Volumes view, or in the Volume / Health details page.

Steps

1. In the left navigation pane, click **Storage** > **Volumes**.
2. In the **Health: All Volumes** view, select one or more volumes and then click **Edit Thresholds**.
3. In the **Edit Volume Thresholds** dialog box, edit the threshold settings of one of the following: capacity, Snapshot copies, qtree quota, growth, or inodes by selecting the appropriate check box and then modifying the settings.
4. Click **Save**.

Related information

Adding users

Edit individual qtree health threshold settings in Active IQ Unified Manager

You can edit the health threshold settings for qtree capacity for one or more qtrees. When a threshold is crossed, alerts are generated and you receive notifications. These notifications help you to take preventive measures based on the event generated.

Before you begin

You must have the Application Administrator or Storage Administrator role.

Based on changes to the threshold values, events are generated or obsoleted in the next monitoring cycle.

Steps

1. In the left navigation pane, click **Storage** > **Qtrees**.
2. In the **Capacity: All Qtrees** view, select one or more qtrees and then click **Edit Thresholds**.
3. In the **Edit Qtree Thresholds** dialog box, change the capacity thresholds for the selected qtree or qtrees and click **Save**.

You can also set individual qtree thresholds from the Qtrees tab on the Storage VM / Health details page.

Manage cluster security objectives

Manage cluster security objectives in Active IQ Unified Manager

Unified Manager provides a dashboard that identifies how secure your ONTAP clusters, storage virtual machines (SVMs), and volumes are based on recommendations defined in the *NetApp Security Hardening Guide for ONTAP 9*.

The goal of the security dashboard is to show any areas where your ONTAP clusters do not align with the NetApp recommended guidelines so that you can fix these potential issues. In most cases you will fix the issues using ONTAP System Manager or the ONTAP CLI. Your organization may not follow all of the recommendations, so in some cases you will not need to make any changes.

See the NetApp Security Hardening Guide for ONTAP 9 (TR-4569) for detailed recommendations and resolutions.

In addition to reporting security status, Unified Manager also generates security events for any cluster or SVM that has security violations. You can track these issues in the Event Management inventory page and you can configure alerts for these events so that your storage administrator is notified when new security events occur.

For more information, see What security criteria are being evaluated.

What security criteria is being evaluated

Understand how Active IQ Unified Manager evaluates security criteria

In general, security criteria for your ONTAP clusters, storage virtual machines (SVMs), and volumes are being evaluated against the recommendations defined in the *NetApp Security Hardening Guide for ONTAP 9*.

Some of the security checks include:

- whether a cluster is using a secure authentication method, such as SAML
- whether peered clusters have their communication encrypted
- whether a storage VM has its audit log enabled
- whether your volumes have software or hardware encryption enabled

See the topics on compliance categories and the NetApp Security Hardening Guide for ONTAP 9 for detailed information.

> Upgrade events that are reported from the Active IQ platform are also considered security events. These events identify issues where the resolution requires you to upgrade ONTAP software, node firmware, or operating system software (for security advisories). These events are not displayed in the Security panel, but they are available from the Event Management inventory page.

For more information, see Managing cluster security objectives.

Understand cluster compliance categories in Active IQ Unified Manager

This table describes the cluster security compliance parameters that Unified Manager evaluates, the NetApp recommendation, and whether the parameter affects the overall determination of the cluster being complaint or not complaint.

Having non-compliant SVMs on a cluster will affect the compliance value for the cluster. So in some cases you may need to fix a security issues with an SVM before your cluster security is seen as compliant.

Note that not every parameter listed below appears for all installations. For example, if you have no peered clusters, or if you have disabled AutoSupport on a cluster, then you will not see the Cluster Peering or AutoSupport HTTPS Transport items in the UI page.

Parameter	Description	Recommendation	Affects Cluster Compliance
Global FIPS	Indicates if Global FIPS (Federal Information Processing Standard) 140-2 compliance mode is enabled or disabled. When FIPS is enabled, TLSv1 and SSLv3 are disabled, and only TLSv1.1 and TLSv1.2 are allowed.	Enabled	Yes
Telnet	Indicates if Telnet access to the system is enabled or disabled. NetApp recommends Secure Shell (SSH) for secure remote access.	Disabled	Yes
Insecure SSH Settings	Indicates if SSH uses insecure ciphers, for example ciphers beginning with *cbc.	No	Yes
Login Banner	Indicates if the Login banner is enabled or disabled for users accessing the system.	Enabled	Yes
Cluster Peering	Indicates if communication between peered clusters is encrypted or unencrypted. Encryption must be configured on both the source and destination clusters for this parameter to be considered compliant.	Encrypted	Yes
Network Time Protocol	Indicates if the cluster has one or more configured NTP servers. For redundancy and best service NetApp recommends that you associate at least three NTP servers with the cluster.	Configured	Yes

Parameter	Description	Recommendation	Affects Cluster Compliance
OCSP	Beginning with 9.14.1, Active IQ Unified Manager provides Online Certificate Status Protocol (OCSP) status information at the storage virtual machine (SVM, formerly known as Vserver) level. This means that OCSP validation is applied to all SSL/TLS connections made to the SVM and ensures the integrity and validity of the certificates used in these connections.	Enabled	No
Remote Audit Logging	Indicates if log forwarding (Syslog) is encrypted or not encrypted.	Encrypted	Yes
AutoSupport HTTPS Transport	Indicates if HTTPS is used as the default transport protocol for sending AutoSupport messages to NetApp support.	Enabled	Yes
Default Admin User	Indicates if the Default Admin User (built-in) is enabled or disabled. NetApp recommends locking (disabling) any unneeded built-in accounts.	Disabled	Yes
SAML Users	Indicates if SAML is configured. SAML enables you to configure multi-factor authentication (MFA) as a login method for single sign-on.	No	No
Active Directory Users	Indicates if Active Directory is configured. Active Directory and LDAP are the preferred authentication mechanisms for users accessing clusters.	No	No

Parameter	Description	Recommendation	Affects Cluster Compliance
LDAP Users	Indicates if LDAP is configured. Active Directory and LDAP are the preferred authentication mechanisms for users managing clusters over local users.	No	No
Certificate Users	Indicates if a certificate user is configured to log into the cluster.	No	No
Local Users	Indicates if local users are configured to log into the cluster.	No	No
Remote Shell	Indicates if RSH is enabled. For security reasons, RSH should be disabled. The Secure Shell (SSH) for secure remote access is preferred.	Disabled	Yes
MD5 in Use	Indicates if ONTAP user accounts use less-secure MD5 Hash function. The MD5 Hashed user accounts migration to the more secure cryptographic hash function like SHA-512 is preferred.	No	Yes
Certificate Issuer Type	Indicates the type of digital certificate used.	CA-Signed	No

Understand Storage VM compliance categories in Active IQ Unified Manager

This table describes the storage virtual machine (SVM) security compliance criteria that Unified Manager evaluates, the NetApp recommendation, and whether the parameter affects the overall determination of the SVM being complaint or not complaint.

Parameter	Description	Recommendation	Affects SVM Compliance
Audit Log	Indicates if Audit logging is enabled or disabled.	Enabled	Yes
Insecure SSH Settings	Indicates if SSH uses insecure ciphers, for example ciphers beginning with `cbc*`.	No	Yes

Parameter	Description	Recommendation	Affects SVM Compliance
Login Banner	Indicates if the Login banner is enabled or disabled for users accessing SVMs on the system.	Enabled	Yes
LDAP Encryption	Indicates if LDAP Encryption is enabled or disabled.	Enabled	No
NTLM Authentication	Indicates if NTLM Authentication is enabled or disabled.	Enabled	No
LDAP Payload Signing	Indicates if LDAP Payload Signing is enabled or disabled.	Enabled	No
CHAP Settings	Indicates if CHAP is enabled or disabled.	Enabled	No
Kerberos V5	Indicates if Kerberos V5 authentication is enabled or disabled.	Enabled	No
NIS Authentication	Indicates if the use of NIS authentication is configured.	Disabled	No
FPolicy Status Active	Indicates if FPolicy is created or not.	Yes	No
SMB Encryption Enabled	Indicates if SMB -Signing & Sealing is not enabled.	Yes	No
SMB Signing Enabled	Indicates if SMB -Signing is not enabled.	Yes	No

Understand volume compliance categories in Active IQ Unified Manager

This table describes the volume encryption parameters that Unified Manager evaluates to determine whether the data on your volumes is adequately protected from being accessed by unauthorized users.

Note that the volume encryption parameters do not affect whether the cluster or storage VM is considered compliant.

Parameter	Description
Software Encrypted	Displays the number of volumes that are protected using NetApp Volume Encryption (NVE) or NetApp Aggregate Encryption (NAE) software encryption solutions.
Hardware Encrypted	Displays the number of volumes that are protected using NetApp Storage Encryption (NSE) hardware encryption.

Parameter	Description
Software and Hardware Encrypted	Displays the number of volumes that are protected by both software and hardware encryption.
Not Encrypted	Displays the number of volumes that are not encrypted.

Learn about cluster and SVM compliance status in Active IQ Unified Manager

Clusters and storage virtual machines (SVMs) are considered not compliant when any of the security criteria that is being evaluated against the recommendations defined in the *NetApp Security Hardening Guide for ONTAP 9* are not met. Additionally, a cluster is considered not compliant when any SVM is flagged as being not compliant.

The status icons in the security cards have the following meanings in relation to their compliance:

- - The parameter is configured as recommended.
- - The parameter is not configured as recommended.
- - Either the functionality is not enabled on the cluster, or the parameter is not configured as recommended, but this parameter does not contribute to the compliance of the object.

Note that volume encryption status does not contribute to whether the cluster or SVM are considered compliant.

View security status for clusters and Storage VMs in Active IQ Unified Manager

Active IQ Unified Manager enables you to view the security status of the storage objects in your environment from different points in the interface. You can collect and analyze information and reports based on defined parameters, and detect suspicious behavior or unauthorized system changes on the monitored clusters and storage VMs.

For the security recommendations, see the NetApp Security Hardening Guide for ONTAP 9

View object level security status on Security page

As a system administrator, you can use the **Security** page to get visibility into the security strength of your ONTAP clusters and storage VMs at the data center and site levels. The supported objects are cluster, storage VMs, and volumes. Follow these steps:

Steps

1. In the left navigation pane, click **Dashboard**.
2. Depending on whether you want to view security status for all monitored clusters or for a single cluster, select **All Clusters** or select a single cluster from the drop-down menu.
3. Click the right-arrow in the **Security** panel. The Security page is displayed.

Clicking the bar charts, counts, and `View Reports` links takes you to the Volumes, Clusters, or Storage VMs page for you to view the corresponding details or generate reports, as required.

The Security page displays the following panels:

- **Cluster Compliance**: the security status (number of clusters that are compliant or not compliant) of all the clusters in a data center
- **Storage VM Compliance**: the security status (number of storage VMs that are compliant or not compliant) for all the storage VMs in your data center
- **Volume Encryption**: the volume encryption status (number of volumes that are encrypted or not encrypted) of all the volumes in your environment
- **Volume Anti-ransomware Status**: the security status (number of volumes with anti-ransomware enabled or disabled) of all the volumes in your environment
- **Cluster Authentication and Certificates**: the number of clusters using each type of authentication method, such as SAML, Active Directory, or through certificates and local authentication. The panel also displays the number of clusters whose certificates have either expired or are about to expire in 60 days.

View security details of all clusters on the Clusters page

The **Clusters / Security** details page enables you to view the security compliance status at a cluster level.

Steps

1. In the left navigation pane, click **Storage > Clusters**.
2. Select **View > Security > All Clusters**.

Default security parameters, such as Global FIPS, Telnet, insecure SSH settings, login banner, network time protocol, AutoSupport HTTPS Transport, and the status of cluster certificate expiration are displayed.

You can click the ⋮ more options button and choose to view the security details on the **Security** page of Unified Manager or on System Manager. You should have valid credentials for viewing the details on System Manager.

If a cluster has an expired certificate, you can click `expired` under **Cluster Certificate Validity**, and renew it from System Manager (9.10.1 and later). You cannot click `expired` if the System Manager instance is of a release earlier than 9.10.1.

View security details of all clusters from the storage VMs page

The **Storage VMs / Security** details page enables you to view the security compliance status at a storage VM level.

Steps

1. In the left navigation pane, click **Storage > Storage VMs**.
2. Select **View > Security > All Storage VMs**. A list of clusters with the security parameters is displayed.

You can have a default view of the storage VMs' security compliance by checking the security parameters, such as storage VMs, cluster, login banner, audit log, and insecure SSH settings.

You can click the ⋮ more options button and choose to view the security details on the **Security** page of Unified Manager or on System Manager. You should have valid credentials for viewing the details on System Manager.

For anti-ransomware security details for volumes and storage VMs, see Viewing the anti-ransomware status of all volumes and Storage VMs.

View security events that might require software or firmware updates in Active IQ Unified Manager

There are certain security events that have an impact area of “Upgrade”. These events are reported from the Active IQ platform, and they identify issues where the resolution requires you to upgrade ONTAP software, node firmware, or operating system software (for security advisories).

Before you begin

You must have the Operator, Application Administrator, or Storage Administrator role.

You may want to perform immediate corrective action for some of these issues, whereas other issues may be able to wait until your next scheduled maintenance. You can view all of these events and assign them to users who can resolve the issues. Additionally, if there are certain security upgrade events that you do not want to be notified about, this list can help you identify those events so that you can disable them.

Steps

1. In the left navigation pane, click **Event Management**.

   By default, all Active (New and Acknowledged) events are displayed on the Event Management inventory page.

2. From the View menu, select **Upgrade events**.

   The page displays all active upgrade security events.

View how clusters manage user authentication on all clusters in Active IQ Unified Manager

The Security page displays the types of authentication being used to authenticate users on each cluster, and the number of users who are accessing the cluster using each type. This enables you to verify that user authentication is being performed securely as defined by your organization.

Steps

1. In the left navigation pane, click **Dashboard**.
2. At the top of the dashboard, select **All Clusters** from the drop-down menu.
3. Click the right-arrow in the **Security** panel and the **Security** page is displayed.
4. View the **Cluster Authentication** card to see the number of users who are accessing the system using each authentication type.
5. View the **Cluster Security** card to view the authentication mechanisms being used to authenticate users on each cluster.

If there are some users accessing the system using an insecure method, or using a method that is not recommended by NetApp, you can disable the method.

View the encryption status of all volumes in Active IQ Unified Manager

You can view a list of all the volumes and their current encryption status so you can determine whether the data on your volumes is adequately protected from being accessed by unauthorized users.

Before you begin

You must have the Operator, Application Administrator, or Storage Administrator role.

The types of encryption that can be applied to a volume are:

- Software - Volumes that are protected using NetApp Volume Encryption (NVE) or NetApp Aggregate Encryption (NAE) software encryption solutions.
- Hardware - Volumes that are protected using NetApp Storage Encryption (NSE) hardware encryption.
- Software and Hardware - Volumes that are protected by both software and hardware encryption.
- None - Volumes that are not encrypted.

Steps

1. In the left navigation pane, click **Storage** > **Volumes**.
2. In the View menu, select **Health** > **Volumes Encryption**
3. In the **Health: Volumes Encryption** view, sort on the **Encryption Type** field, or use the Filter to display volumes that have a specific encryption type, or that are not encrypted (Encryption Type of “None”).

View the anti-ransomware status of all volumes and storage VMs in Active IQ Unified Manager

You can view a list of all volumes and storage VMs (SVMs) and their current anti-ransomware status so you can determine whether the data on your volumes and SVMs are adequately protected from ransomware attacks.

Before you begin

You must have the Operator, Application Administrator, or Storage Administrator role.

For more information on the different anti-ransomware statuses, see ONTAP: Enable anti-ransomware.

View security details of all volumes with anti-ransomware detection

Steps

1. In the left navigation pane, click **Storage** > **Volumes**.
2. In the View menu, select **Health** > **Security** > **Anti-ransomware**
3. In the **Security: Anti-ransomware** view, you can sort by the various fields or use the Filter.

> Anti-ransomware is not supported for offline volumes, restricted volumes, SnapLock volumes, FlexGroup volumes, FlexCache volumes, SAN-only volumes, volumes of stopped storage VMs, root volumes of storage VMs, or data protection volumes.

View security details of all storage VMs with anti-ransomware detection

Steps

1. In the left navigation pane, click **Storage > Storage VMs**.
2. Select **View > Security > Anti-ransomware**. A list of SVMs with the anti-ransomware status is displayed.

> Anti-ransomware monitoring is not supported on storage VMs that do not have NAS protocol enabled.

View all active security events in Active IQ Unified Manager

You can view all the active security events and then assign each of them to a user who can resolve the issue. Additionally, if there are certain security events that you do not want to receive, this list can help you identify the events that you want to disable.

Before you begin

You must have the Operator, Application Administrator, or Storage Administrator role.

Steps

1. In the left navigation pane, click **Event Management**.

   By default, New and Acknowledged events are displayed on the Event Management inventory page.

2. From the View menu, select **Active security events**.

   The page displays all New and Acknowledged Security events that have been generated in the past 7 days.

Add alerts for security events in Active IQ Unified Manager

You can configure alerts for individual security events just like any other events received by Unified Manager. Additionally, if you want to treat all security events alike and have email sent to the same person, you can create a single alert to notify you when any security events are triggered.

Before you begin

You must have the Application Administrator or Storage Administrator role.

The example below shows how to create an alert for the “Telnet Protocol Enabled” security event. This will send an alert if Telnet access is configured for remote administrative access into the cluster. You can use this same methodology to create alerts for all security events.

Steps

1. In the left navigation pane, click **Storage Management** > **Alert Setup**.
2. In the **Alert Setup** page, click **Add**.
3. In the **Add Alert** dialog box, click **Name**, and enter a name and description for the alert.
4. Click **Resources** and select the cluster or cluster on which you want to enable this alert.
5. Click **Events** and perform the following actions:
   a. In the Event Severity list, select **Warning**.
   b. In the Matching Events list, select **Telnet Protocol Enabled**.
6. Click **Actions** and then select the name of the user who will receive the alert email in the **Alert these users** field.
7. Configure any other options on this page for notification frequency, issuing SNMP taps, and executing a script.
8. Click **Save**.

Disable specific security events in Active IQ Unified Manager

All events are enabled by default. You can disable specific events to prevent the generation of notifications for those events that are not important in your environment. You can enable events that are disabled if you want to resume receiving notifications for them.

Before you begin

You must have the Application Administrator or Storage Administrator role.

When you disable events, the previously generated events in the system are marked obsolete, and the alerts that are configured for these events are not triggered. When you enable events that are disabled, the notifications for these events are generated starting with the next monitoring cycle.

Steps

1. In the left navigation pane, click **Storage Management** > **Event Setup**.
2. In the **Event** Setup page, disable or enable events by choosing one of the following options:

If you want to…	Then do this…
Disable events	a. Click **Disable**. b. In the Disable Events dialog box, select the **Warning** severity. This is the category for all security events. c. In the Matching Events column, select the security events that you want to disable, and then click the right arrow to move those events to the Disable Events column. d. Click **Save and Close**. e. Verify that the events that you disabled are displayed in the list view of the Event Setup page.
Enable events	a. From the list of disabled events, select the check box for the event, or events, that you want to reenable. b. Click **Enable**.

View security events in Active IQ Unified Manager

Security events provide you with information about the security status of ONTAP clusters, storage virtual machines (SVMs), and volumes based on parameters defined in the *NetApp Security Hardening Guide for ONTAP 9*. These events notify you of potential issues so that you can evaluate their severity and fix the issue if necessary.

Security events are grouped by source type and include the event and trap name, impact level, and severity. These events appear in the cluster and storage VM event categories.

Manage backup and restore operations

Manage backup and restore operations in Active IQ Unified Manager

You can create backups of Active IQ Unified Manager and use the restore feature to restore the backup to the same (local) system or a new (remote) system in case of a system failure or data loss.

There are three backup and restore methods depending on the operating system on which you have installed Unified Manager, and based on the number of clusters and nodes being managed:

Operating System	Size of Deployment	Recommended Backup Method
VMware vSphere	Any	VMware snapshot of the Unified Manager virtual appliance
Red Hat Enterprise Linux	Small	Unified Manager MySQL database dump
	Large	NetApp Snapshot of Unified Manager database
Microsoft Windows	Small	Unified Manager MySQL database dump
	Large	NetApp Snapshot of Unified Manager database with iSCSI protocol

These different methods are described in the sections that follow.

Backup and restore for Active IQ Unified Manager on a virtual appliance

The backup and restore model for Unified Manager when installed on a virtual appliance is to capture and restore an image of the full virtual application.

The following tasks enable you to complete a backup of the virtual appliance:

1. Power off the VM and take a VMware snapshot of the Unified Manager virtual appliance.
2. Make a NetApp Snapshot copy on the datastore to capture the VMware snapshot.

   If the datastore is not hosted on a system running ONTAP software, follow the storage vendor guidelines to create a backup of the VMware snapshot.

3. Replicate the NetApp Snapshot copy, or snapshot equivalent, to alternate storage.
4. Delete the VMware snapshot.

You should implement a backup schedule using these tasks to ensure that the Unified Manager virtual appliance is protected if issues arise.

To restore the VM, you can use the VMware snapshot you created to restore the VM to the backup point-in-time state.

Backup and restore using a MySQL database dump

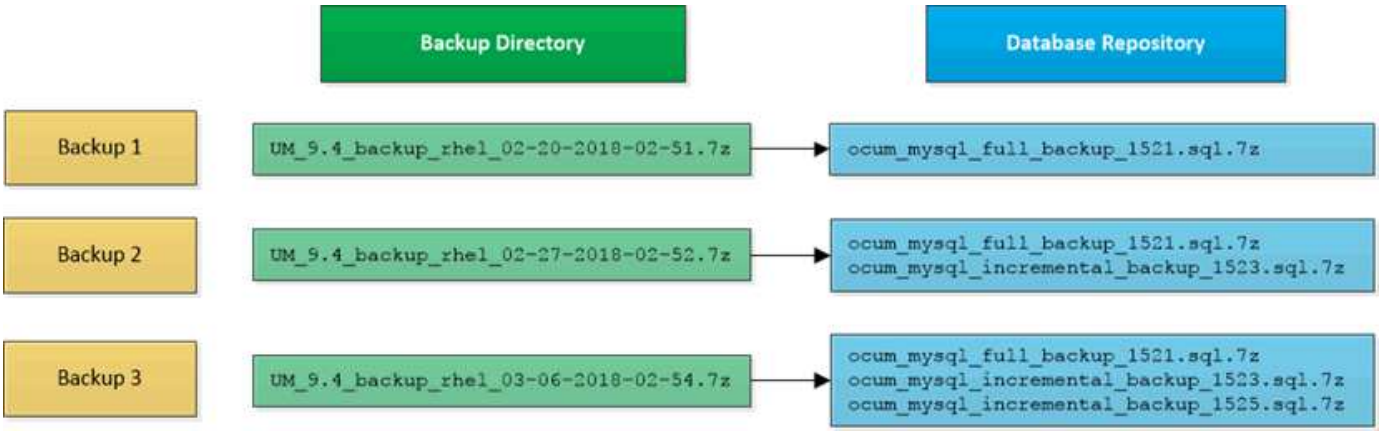
Backup and restore using a MySQL database dump in Active IQ Unified Manager

A MySQL database dump backup is a copy of the Active IQ Unified Manager database and configuration files that you can use in case of a system failure or data loss. You can schedule a backup to be written to a local destination or to a remote destination. It is highly recommended that you define a remote location that is external to the Active IQ Unified Manager host system.

MySQL database dump is the default backup mechanism when Unified Manager is installed on a Linux and Windows server. However, if Unified Manager is managing a large number of cluster and nodes, or if your MySQL backups are taking many hours to complete, you can back up using Snapshot copies. This functionality is available on Red Hat Enterprise Linux and Windows.

A database dump backup consists of a single file in the backup directory and one or more files in the database repository directory. The file in the backup directory is very small because it contains only a pointer to the files located in the database repository directory that are required to recreate the backup.

The first time you generate a database backup a single file is created in the backup directory and a full backup file is created in the database repository directory. The next time you generate a backup a single file is created in the backup directory and an incremental backup file is created in the database repository directory that contains the differences from the full backup file. This process continues as you create additional backups, up to the maximum retention setting, as shown in the following figure.



Do not rename or remove any of the backup files in these two directories or any subsequent restore operation will fail.

If you write your backup files to the local system, you should initiate a process to copy the backup files to a remote location so they will be available in case you have a system issue that requires a complete restore.

Before beginning a backup operation, Active IQ Unified Manager performs an integrity check to verify that all the required backup files and backup directories exist and are writable. It also checks that there is enough space on the system to create the backup file.

Configure database dump backup settings in Active IQ Unified Manager

You can configure the Unified Manager database dump backup settings to set the

database backup path, retention count, and backup schedule. You can enable daily or weekly scheduled backups. By default, scheduled backups are disabled, but you should set a backup schedule.

Before you begin

- You must have the Operator, Application Administrator, or Storage Administrator role.
- You must have a minimum of 150 GB of space available in the location you define as the backup path.

  It is recommended that you use a remote location that is external to the Unified Manager host system.
- When Unified Manager is installed on a Linux system, and using MySQL backup, ensure that the following permissions and ownerships are set on the backup directory.

  Permissions: 0750, Ownership: jboss:maintenance
- When Unified Manager is installed on a Windows system, and using MySQL backup, ensure that only the administrator has access to the backup directory.

More time is required the first time a backup is performed than for subsequent backups because the first backup is a full backup. A full backup can be over 1 GB and can take three to four hours. Subsequent backups are incremental and require less time.

- If you find the number of incremental backup files to be too large for the space you have allocated for backups, you can take a full backup periodically to replace the old backup and its incremental files. As another option, you can take a backup by using Snapshot copies.
- Backup taken during initial 15 days of a new cluster addition might not be accurate enough to get the historical performance data.

Steps

1. In the left navigation pane, click **General** > **Database Backup**.
2. In the **Database Backup** page, click **Backup Settings**.
3. Configure the appropriate values for a backup path, retention count, and schedule.

   The default value for retention count is 10; you can use 0 for creating unlimited backups.
4. Select the **Scheduled Daily** or **Scheduled Weekly** button, and then specify the schedule details.
5. Click **Apply**.

Database dump backup files are created based on the schedule. You can see the available backup files in the Database Backup page.

Learn about database restore in Active IQ Unified Manager

A MySQL database restore is the process of restoring an existing Unified Manager backup file to the same or a different Unified Manager server. You perform the restore operation from the Unified Manager maintenance console.

If you are performing a restore operation on the same (local) system, and the backup files are all stored locally, you can run the restore option using the default location. If you are performing a restore operation on a different Unified Manager system (a remote system), you must copy the backup file, or files, from secondary

storage to the local disk before running the restore option.

During the restore process, you are logged out of Unified Manager. You can log in to the system after the restore process is complete.

If you are restoring the backup image to a new server, after the restore operation completes you need to generate a new HTTPS security certificate and restart the Unified Manager server. You will also need to reconfigure SAML authentication settings, if they are required, when restoring the backup image to a new server.

Old backup files cannot be used to restore an image after Unified Manager has been upgraded to a newer version of software. To save space, all old backup files, except the newest file, are removed automatically when you upgrade Unified Manager.

Related information

Generating an HTTPS security certificate

Enabling SAML authentication

Authentication with Active Directory or OpenLDAP

Restore a MySQL database backup on a Linux system in Active IQ Unified Manager

If data loss or data corruption occurs, you can restore Unified Manager to the previous stable state with minimum loss of data. You can restore the Unified Manager database to a local or remote Red Hat Enterprise Linux system by using the Unified Manager maintenance console.

Before you begin

- You must have the root user credentials for the Linux host on which Unified Manager is installed.
- You must have a user ID and password authorized to log in to the maintenance console of the Unified Manager server.
- You must have copied the Unified Manager backup file and the contents of the database repository directory to the system on which you will perform the restore operation.

  It is recommended that you copy the backup file to the default directory `/data/ocum`-backup. The database repository files must be copied to the `/database-dumps-repo` subdirectory under the `/ocum-backup` directory.

- The backup files must be of `.7z` type.

The restore feature is platform-specific and version-specific. You can restore a Unified Manager backup only on the same version of Unified Manager. You can restore a Linux backup file or a virtual appliance backup file to a Red Hat Enterprise Linux system.

If the backup folder name contains a space, you must include the absolute path or relative path in double quotation marks.

Steps

1. If you are performing a restore onto a new server, after installing Unified Manager do not launch the UI or

configure any clusters, users, or authentication settings when the installation is complete. The backup file populates this information during the restore process.

2. Using Secure Shell, connect to the IP address or fully qualified domain name of the Unified Manager system.
3. Log in to the system with the maintenance user (umadmin) name and password.
4. Enter the command `maintenance_console` and press Enter.
5. In the maintenance console **Main Menu**, enter the number for the **Backup Restore** option.
6. Enter the number for the **Restore MySQL Backup**.
7. When prompted, enter the absolute path of the backup file.

```
Bundle to restore from: /data/ocum-
backup/UM_9.8.N151113.1348_backup_rhel_02-20-2020-04-45.7z
```

After the restore operation is complete, you can log in to Unified Manager.

After you restore the backup, if the OnCommand Workflow Automation server does not work, perform the following steps:

1. On the Workflow Automation server, change the IP address of the Unified Manager server to point to the latest machine.
2. On the Unified Manager server, reset the database password if the acquisition fails in step 1.

Restore a MySQL database backup on Windows in Active IQ Unified Manager

In case of data loss or data corruption, you can use the restore feature to restore Unified Manager to the previous stable state with minimal loss. You can restore the Unified Manager MySQL database to a local Windows system or a remote Windows system by using the Unified Manager maintenance console.

Before you begin

- You must have Windows administrator privileges.
- You must have copied the Unified Manager backup file and the contents of the database repository directory to the system on which you will perform the restore operation.

  It is recommended that you copy the backup file to the default directory `\ProgramData\NetApp\OnCommandAppData\ocum\backup`. The database repository files must be copied to the `\database_dumps_repo` subdirectory under the `\backup` directory.

- The backup files must be of `.7z` type.

The restore feature is platform-specific and version-specific. You can restore a Unified Manager MySQL backup only on the same version of Unified Manager, and a Windows backup can be restored only on a Windows platform.

If the folder names contain a space, you must include the absolute path or relative path of the backup file in double quotation marks.

Steps

1. If you are performing a restore onto a new server, after installing Unified Manager do not launch the UI or configure any clusters, users, or authentication settings when the installation is complete. The backup file populates this information during the restore process.
2. Log in to the Unified Manager system with administrator credentials.
3. Launch PowerShell or command prompt as a Windows administrator.
4. Enter the command `maintenance_console` and press Enter.
5. In the maintenance console **Main Menu**, enter the number for the **Backup Restore** option.
6. Enter the number for the **Restore MySQL Backup**.
7. When prompted, enter the absolute path of the backup file.

   ```
   Bundle to restore from:
   \ProgramData\NetApp\OnCommandAppData\ocum\backup\UM_9.8.N151118.2300_bac
   kup_windows_02-20-2020-02-51.7z
   ```

   After the restore operation is complete, you can log in to Unified Manager.

After you restore the backup, if the OnCommand Workflow Automation server does not work, perform the following steps:

1. On the Workflow Automation server, change the IP address of the Unified Manager server to point to the latest machine.
2. On the Unified Manager server, reset the database password if the acquisition fails in step 1.

Backup and restore using NetApp Snapshots

Backup and restore using NetApp Snapshots in Active IQ Unified Manager

A NetApp Snapshot copy creates a point-in-time image of the Unified Manager database and configuration files that you can use to restore in case of a system failure or data loss. You schedule a Snapshot copy to be written to a volume on one of your ONTAP clusters periodically so that you always have a current copy.

This functionality is not available for Active IQ Unified Manager installed on a virtual appliance.

Configure backup on Linux in Active IQ Unified Manager

If your Active IQ Unified Manager is installed on a Linux machine, then you can decide to configure your backup and restore using NetApp Snapshots.

Snapshot copies take very little time, usually just a few minutes, and the Unified Manager database is locked for a very short timeframe, so there is very little disruption to your installation. The image consumes minimal storage space and incurs negligible performance overhead because it records only changes to files since the last Snapshot copy was made. Because the Snapshot is created on an ONTAP cluster, you can take advantage of other NetApp features such as SnapMirror to create secondary protection, if needed.

Before beginning a backup operation, Unified Manager performs an integrity check to verify that the destination

system is available.

- You can restore a Snapshot copy only on the same version of Active IQ Unified Manager.

  For example, if you created a backup on Unified Manager 9.18, the backup can be restored only on Unified Manager 9.18 systems.

- If there is any change in the Snapshot configuration, it might cause the snapshot to be invalid.

Configure Snapshot copy location in Active IQ Unified Manager

You can configure the volume where Snapshot copies will be stored on one of your ONTAP clusters using ONTAP System Manager or using the ONTAP CLI.

Before you begin

The cluster, storage VM, and volume must meet the following requirements:

- Cluster requirements:
  - ONTAP 9.3 or greater must be installed
  - It should be geographically close to the Unified Manager server
  - It can be monitored by Unified Manager, but it is not required
- Storage VM requirements:
  - The name switch and name mapping must be set to use “files”
  - Local users created to correspond with client-side users
  - Make sure All Read/Write access is selected
  - Make sure that Superuser Access is set to “any” in the export policy
  - NFS for NetApp Snapshot for Linux
  - NFSv4 must be enabled on the NFS server and NFSv4 ID domain specified on the client and storage VM
  - The volume should be at least double the size of the Unified Manager/opt/netapp/data directory

    Use the command du -sh /opt/netapp/data/ to check the current size.
- Volume requirements:
  - The volume should be at least double the size of the Unified Manager /opt/netapp/data directory
  - The security style must be set to UNIX
  - The local snapshot policy must be disabled
  - Volume autosize should be enabled
  - The performance service level should be set to a policy with high IOPS and low latency, such as “Extreme”

For detailed steps to create the NFS volume, see How to configure NFSv4 in ONTAP 9 and the ONTAP 9 NFS Configuration Express Guide.

Specify the destination location for Snapshot copies in Active IQ Unified Manager

You should configure the destination location for Active IQ Unified Manager Snapshot copies on a volume that you have already configured in one of your ONTAP clusters. You should use maintenance console to define the location.

- You must have the root user credentials for the Linux host on which Active IQ Unified Manager is installed.
- You must have a user ID and password authorized to log in to the maintenance console of the Unified Manager server.
- You must have the Cluster Management IP address, the name of the storage VM, the name of the volume, and the storage system user name and password.
- You must have mounted the volume to the Active IQ Unified Manager host, and you must have the mount path.

Steps

1. Use Secure Shell to connect to the IP address or FQDN of the Active IQ Unified Manager system.
2. Log in to the system with the maintenance user (umadmin) name and password.
3. Enter the command `maintenance_console` and press Enter.
4. In the maintenance console **Main Menu**, enter the number for the **Backup Restore** option.
5. Enter the number for **Configure NetApp Snapshot Backup**.
6. Enter the number to configure NFS.
7. Review the information that you will need to provide and then enter the number for **Enter Backup Configuration Details**.
8. To identify the volume where the Snapshot will be written, enter the IP address of the Cluster Management interface, the name of the storage VM, the name of the volume, LUN name, the storage system user name and password, and the mount path.
9. Verify this information and enter `y`.

   The system performs the following tasks:

   - Establishes the connection to the cluster
   - Stops all the services
   - Creates a new directory in the volume and copies the Active IQ Unified Manager database configuration files
   - Deletes the files from Active IQ Unified Manager and creates a symlink to the new database directory
   - Restarts all the services
10. Exit the maintenance console and launch the Active IQ Unified Manager interface to create a schedule for the Snapshot copy if you have not already done this.

Configure backup on Windows in Active IQ Unified Manager

Active IQ Unified Manager supports backup and restore using NetApp Snapshots on Windows operating system with the help of LUN using iSCSI protocol.

Snapshot based backup can be taken while all Unified Manager services are running. A consistent state of database is captured as part of the Snapshot as the backup puts a global read lock on the entire database that prevents any concurrent write. For your Unified Manager system installed on Windows OS to perform backup

and restore using NetApp Snapshots, you should first configure Unified Manager backup to Snapshot based using the maintenance console.

Before you configure Unified Manager for creating Snapshot copies, you should perform the following configuration tasks.

- Configure ONTAP cluster
- Configure Windows host machine

Configure a backup location for Active IQ Unified Manager on Windows

You should configure the volume for storing Snapshot copies after backing up Unified Manager on Windows.

Before you begin

The cluster, storage VM, and volume must meet the following requirements:

- Cluster requirements:
  - ONTAP 9.3 or greater must be installed
  - It should be geographically close to the Unified Manager server
  - It is monitored by Unified Manager
- Storage VM requirements:
  - iSCSI connectivity on ONTAP cluster
  - iSCSI protocol must be enabled for the configured machine
  - You should have a dedicated volume and LUN for backup configuration. The selected volume should contain only one LUN and nothing else.
  - The size of the LUN should be at least twice the data size expected to be handled in the 9.9 Active IQ Unified Manager.

    This sets the same size requirement on volume as well.

  - Make sure All Read/Write access is selected
  - Make sure that Superuser Access is set to “any” in the export policy
- Volume and LUN requirements:
  - The volume should be at least double the size of the Unified Manager MySQL data directory.
  - The security style must be set to Windows
  - The local snapshot policy must be disabled
  - Volume autosize should be enabled
  - The performance service level should be set to a policy with high IOPS and low latency, such as “Extreme”

Configure an ONTAP cluster for backup and restore in Active IQ Unified Manager

You need to perform few pre-configuration steps on ONTAP clusters before you can back up and restore Active IQ Unified Manager using Snapshot copy on Windows systems.

You can configure ONTAP cluster using either the command prompt or System Manager user interface. The configuration of ONTAP cluster involves configuring Data LIFs to be available to be assigned as iSCSI LIFs to

the storage VM. The next step is to configure an iSCSI enabled storage VM using the System Manager user interface. You will need to configure a static network route for this storage VM to control how LIFs use the network for outbound traffic.

> You should have a dedicated volume and a LUN for backup configuration. The selected volume should include only one LUN. The size of the LUN should be at least twice the data size expected to be handled by Active IQ Unified Manager.

You need to perform the following configuration:

Steps

1. Configure a iSCSI enabled storage VM or use an existing storage VM that has the same configuration.
2. Configure a network route for the configured storage VM.
3. Configure a volume of appropriate capacity and a single LUN inside ensuring that the volume is dedicated only for this LUN.

> In a scenario when the LUN is created on System Manager, unmapping the LUN might delete the igroup, and restore might fail. To avoid this scenario, ensure that while creating a LUN, it is created explicitly and is not deleted when the LUN is unmapped.

4. Configure an initiator group in the storage VM.
5. Configure a port set.
6. Integrate the igroup with the portset.
7. Map the LUN to the igroup.

Configure a Windows host machine in Active IQ Unified Manager

You need to configure your Windows host machine before you can use NetApp Snapshot to back up and restore Active IQ Unified Manager.
To start the Microsoft iSCSI initiator on a Windows host machine, type “iscsi” in the search bar and click **iSCSI Initiator**.

Before you begin

You should clean up any previous configurations on the host machine.

If you are trying to start the iSCSI initiator on a fresh installation of Windows, you are prompted for confirmation, and on your confirmation, the iSCSI Properties dialog box is displayed. If it is an existing Windows installation, then the iSCSI Properties dialog box displayed with a target that is either inactive or trying to connect. So, you will need to ensure that all the previous configurations on the Windows host are removed.

Steps

1. Clean up any previous configurations on the host machine.
2. Discover the target portal.
3. Connect to the target portal.
4. Connect using multipath to the target portal.
5. Discover both the LIFs.
6. Discover the LUN configured in the Windows machine as a device.

7. Configure the discovered LUN as a new volume drive in Windows.

Specify the destination location for Snapshot copies on Windows in Active IQ Unified Manager

You should configure the destination location for Active IQ Unified Manager Snapshot copies on a volume that you have already configured in one of your ONTAP clusters. You should use maintenance console to define the location.

- You must have the administrator privilege for Windows host on which Active IQ Unified Manager is installed.
- You must have a user ID and password authorized to log in to the maintenance console of the Unified Manager server.
- You must have the Cluster Management IP address, the name of the storage VM, the name of the volume, LUN name, and the storage system user name and password.
- You must have mounted the volume as a network drive to the Active IQ Unified Manager host, and you must have the mount drive.

Steps

1. Using Power Shell, connect to the IP address or fully qualified domain name of the Active IQ Unified Manager system.
2. Log in to the system with the maintenance user (umadmin) name and password.
3. Enter the command `maintenance_console` and press Enter.
4. In the maintenance console **Main Menu**, enter the number for the **Backup Restore** option.
5. Enter the number for **Configure NetApp Snapshot Backup**.
6. Enter the number to configure iSCSI.
7. Review the information that you will need to provide and then enter the number for **Enter Backup Configuration Details**.
8. To identify the volume where the Snapshot will be written, enter the IP address of the Cluster Management interface, the name of the storage VM, the name of the volume, LUN name, the storage system user name and password, and the mount drive.
9. Verify this information and enter `y`.

   The system performs the following tasks:

   - Storage VM is validated
   - Volume is validated
   - Mount drive and status is validated
   - LUN existence and status
   - Network drive existence
   - Existence of recommend space (more than twice of mysql data directory) at mounted volume is validated
   - LUN path corresponding to the dedicated LUN in the volume
   - igroup name
   - GUID of the volume where the network drive is mounted

- iSCSI initiator used to communicate with ONTAP

10. Exit the maintenance console and launch the Active IQ Unified Manager interface to create a schedule for Snapshot copies.

Configure a Snapshot backup from the maintenance console in Active IQ Unified Manager

To take Active IQ Unified Manager backup by using Snapshot copy, you should perform a few configuration steps from the maintenance console.

Before you begin

You should have the following details for your system:

- Cluster IP address
- Storage VM name
- Volume name
- LUN name
- Mount path
- Storage system credentials

Steps

1. Access the maintenance console of Unified Manager.
2. Enter 4 to select **Backup Restore**.
3. Enter 2 to select **Backup and Restore using NetApp Snapshot**.

   If you want to change the backup configuration, then enter 3 for selecting **Update NetApp Snapshot Backup Configuration**. You can only update the password.

4. From the menu, enter 1 to select the **Configure NetApp Snapshot Backup**.
5. Enter 1 to provide the required information.
6. Provide the username and password for the maintenance console, and then provide the confirmation that LUN is mounted on host.

   The process then verifies that the data directory, LUN path, storage VM, volumes, space availability, drive, and so on provided by you are correct. The operations that proceed in the background are:

   - Services are stopped
   - Database directory is moved to mounted storage
   - Database directory is deleted and symlinks are established
   - Services are restarted
     After the configuration completes in the Active IQ Unified Manager interface, the backup type is modified to NetApp Snapshot and reflects in the user interface as Database backup (Snapshot based).

Before beginning a backup operation, you must check whether there is any change in the Snapshot configuration because it might cause the snapshot to be invalid. Suppose you configured backup in G drive and Snapshot taken. You later reconfigured the backup to E drive and data is saved to E drive as per the new configuration. If you try to restore Snapshot taken while it was in G drive, it fails with error that G drive does not exist.

Define a backup schedule for Linux and Windows in Active IQ Unified Manager

You can configure the schedule at which Unified Manager Snapshot copies are created by using the Unified Manager UI.

Before you begin

- You must have the Operator, Application Administrator, or Storage Administrator role.
- You must have configured the settings for creating Snapshot copies from the maintenance console to identify the destination where the snapshots will be created.

Snapshot copies are created in just a few minutes and the Unified Manager database is locked only for few seconds.

Backup taken during initial 15 days of a new cluster addition might not be accurate enough to get the historical performance data.

Steps

1. In the left navigation pane, click **General** > **Database Backup**.
2. In the **Database Backup** page, click **Backup Settings**.
3. Enter the maximum number of Snapshot copies that you want to retain in the **Retention Count** field.

   The default value for retention count is 10. The maximum number of Snapshot copies is determined by the version of ONTAP software on the cluster. You can leave this field blank to implement the maximum value regardless of ONTAP version.
4. Select the **Scheduled Daily** or **Scheduled Weekly** button, and then specify the schedule details.
5. Click **Apply**.

Snapshot copies are created based on the schedule. You can see the available backup files in the Database Backup page.

Because of the importance of this volume and the snapshots, you may want to create one or two alerts for this volume so you are notified when either:

- The volume space is 90% full. Use the event **Volume Space Full** to set up the alert.

  You can add capacity to the volume using ONTAP System Manager or the ONTAP CLI so that the Unified Manager database does not run out of space.
- The number of snapshots is close to reaching the maximum number. Use the event **Too Many Snapshot Copies** to set up the alert.

  You can delete older snapshots using ONTAP System Manager or the ONTAP CLI so that there is always room for new Snapshot copies.

You configure alerts in the Alert Setup page.

Restore Active IQ Unified Manager by using Snapshot copies

If data loss or data corruption occurs, you can restore Unified Manager to the previous stable state with minimum loss of data. You can restore the Unified Manager Snapshot database to a local or remote operating system by using the Unified Manager

maintenance console.

Before you begin

- You must have the root user credentials for the Linux host and administrative privileges for Windows host machine on which Unified Manager is installed.
- You must have a user ID and password authorized to log in to the maintenance console of the Unified Manager server.

The restore feature is platform-specific and version-specific. You can restore a Unified Manager backup only on the same version of Unified Manager.

Steps

1. Connect to the IP address or fully qualified domain name of the Unified Manager system.
    - Linux: Secure Shell
    - Windows: Power Shell
2. Log in to the system with the root user credentials.
3. Enter the command `maintenance_console` and press Enter.
4. In the maintenance console **Main Menu**, enter 4 for the **Backup Restore** option.
5. Enter 2 for selecting **Backup and Restore using NetApp Snapshot**.

    If you are performing a restore onto a new server, after installing Unified Manager do not launch the UI or configure any clusters, users, or authentication settings when the installation is complete. Enter 1 for selecting **Configure NetApp Snapshot Backup** and configure the settings for Snapshot copies as they are on the original system.

6. Enter 3 for selecting **Restore using NetApp Snapshot**.
7. Select the Snapshot copy from which you want to restore Unified Manager. Press **Enter**.
8. After the restore process is complete, log in to the Unified Manager user interface.

After you restore the backup, if the Workflow Automation server does not work, perform the following steps:

1. On the Workflow Automation server, change the IP address of the Unified Manager server to point to the latest machine.
2. On the Unified Manager server, reset the database password if the acquisition fails in step 1.

Modify the backup type in Active IQ Unified Manager

If you want to change the type of backup for your Active IQ Unified Manager system, then you can use the maintenance console options. The **Unconfigure NetApp Snapshot Backup** option enables you to fall back to the MySQL based backup.

Before you begin

You must have a user ID and password authorized to log in to the maintenance console of the Unified Manager server.

Steps

1. Access the maintenance console.
2. Select 4 from the **Main Menu** for backup and restore.

3. Select 2 from the **Backup and Restore Menu**.
4. Select 4 for **Unconfigure NetApp Snapshot Backup**.

   The actions that are performed are displayed which are, stop the services, break the symlink, move the data from storage to directory, and then start the services again.

   After the backup method is modified, the backup mechanism is changed from Snapshot copy to default MySQL backup. This change appears in the Database Backup section of the General settings.

Understand on-demand backup for Active IQ Unified Manager

You can use the Active IQ Unified Manager user interface to generate on demand backup whenever required. The on-demand backup enables you to instantaneously create a backup using the existing backup method. The on-demand backup does not differentiate between MySQL or NetApp Snapshot based backup.

You can perform on-demand backup using the **Backup Now** button on the Database Backup page. The on-demand backup does not depend on the schedules that you have configured for Active IQ Unified Manager.

Migrate an Active IQ Unified Manager virtual appliance to a Linux system

You can restore a Unified Manager MySQL database dump backup from a virtual appliance to a Red Hat Enterprise Linux system if you want to change the host operating system on which Unified Manager is running.

Before you begin

- On the virtual appliance:
  - You must have the Operator, Application Administrator, or Storage Administrator role.
  - You must know the name of the Unified Manager maintenance user for the restore operation.
- On the Linux system:
  - You must have installed Unified Manager on a Linux server following the instructions in Installing Unified Manager on Linux systems.
  - The version of Unified Manager on this server must be the same as the version on the virtual appliance from which you are using the backup file.
  - Do not launch the UI or configure any clusters, users, or authentication settings on the Linux system after installation. The backup file populates this information during the restore process.
  - You must have the root user credentials for the Linux host.

These steps describe how to create a backup file on the virtual appliance, copy the backup files to the Red Hat Enterprise Linux, and then restore the database backup to the new system.

Steps

1. On the virtual appliance, click **Management** > **Database Backup**.
2. In the **Database Backup** page, click **Backup Settings**.
3. Change the backup path to /jail/support.
4. In the Schedule section, select **Scheduled Daily**, and enter a time a few minutes past the current time so that the backup is created shortly.

5. Click **Apply**.
6. Wait a few hours for the backup to be generated.

   A full backup can be over 1 GB and can take three to four hours to complete.

7. Log in as the root user to the Linux host on which Unified Manager is installed and copy the backup files from /support on the virtual appliance using SCP.`root@<rhel_server>:/# scp -r admin@<vapp_server_ip_address>:/support/* .`

   ```
   root@ocum_rhel-21:/# scp -r admin@10.10.10.10:/support/* .
   ```

   Make sure you have copied the .7z backup file and all the .7z repository files in the /database-dumps-repo subdirectory.

8. At the command prompt, restore the backup: `um backup restore -f /<backup_file_path>/<backup_file_name>`

   ```
   um backup restore -f /UM_9.7.N151113.1348_backup_unix_02-12-2019-04-16.7z
   ```

9. After the restore operation completes, log in to the Unified Manager web UI.

You should perform the following tasks:

- Generate a new HTTPS security certificate and restart the Unified Manager server.
- Change the backup path to the default setting for your Linux system (/data/ocum-backup), or to a new path of your choice, because there is no /jail/support path on the Linux system.
- Reconfigure both sides of your Workflow Automation connection, if WFA is being used.
- Reconfigure SAML authentication settings, if you are using SAML.

After you have verified that everything is running as expected on your Linux system, you can shut down and remove the Unified Manager virtual appliance.

Manage scripts

Manage scripts in Active IQ Unified Manager

You can use scripts to automatically modify or update multiple storage objects in Unified Manager. The script is associated with an alert. When an event triggers an alert, the script is executed. You can upload custom scripts and test their execution when an alert is generated.

The ability to upload scripts to Unified Manager and run them is enabled by default. If your organization does not want to allow this functionality because of security reasons, you can disable this functionality from **Storage Management** > **Feature Settings**.

Understand how scripts work with alerts in Active IQ Unified Manager

You can associate an alert with your script so that the script is executed when an alert is raised for an event in Unified Manager. You can use the scripts to resolve issues with storage objects or identify which storage objects are generating the events.

When an alert is generated for an event in Unified Manager, an alert email is sent to the specified recipients. If you have associated an alert with a script, the script is executed. You can get the details of the arguments passed to the script from the alert email.

If you have created a custom script and associated it with an alert for a specific event type, actions are taken based on your custom script for that event type, and the **Fix it** actions are not available by default on the Management Actions page or Unified Manager dashboard.

The script uses the following arguments for execution:

- `-eventID`
- `-eventName`
- `-eventSeverity`
- `-eventSourceID`
- `-eventSourceName`
- `-eventSourceType`
- `-eventState`
- `-eventArgs`

You can use the arguments in your scripts and gather related event information or modify storage objects.

Example for obtaining arguments from scripts

```
print "$ARGV[0] : $ARGV[1]\n"
print "$ARGV[7] : $ARGV[8]\n"
```

When an alert is generated, this script is executed and the following output is displayed:

```
-eventID : 290
-eventSourceID : 4138
```

Add scripts in Active IQ Unified Manager

You can add scripts in Unified Manager, and associate the scripts with alerts. These scripts are executed automatically when an alert is generated, and enable you to obtain information about storage objects for which the event is generated.

Before you begin

- You must have created and saved the scripts that you want to add to the Unified Manager server.
- The supported file formats for scripts are Perl, Shell, PowerShell, Python, and `.bat` files.

Platform on which Unified Manager is installed	Supported languages
VMware	Perl and Shell scripts

Platform on which Unified Manager is installed	Supported languages
Linux	Perl, Python, and Shell scripts
Windows	PowerShell, Perl, Python, and .bat scripts

  - For Perl scripts, Perl must be installed on the Unified Manager server. For VMware installations, Perl 5 is installed by default and scripts will support only what Perl 5 supports. If Perl was installed after Unified Manager, you must restart the Unified Manager server.
  - For PowerShell scripts, the appropriate PowerShell execution policy must be set on the Windows server so that the scripts can be executed.

If your script creates log files to track the alert script progress, you must make sure that the log files are not created anywhere within the Unified Manager installation folder.

- You must have the Application Administrator or Storage Administrator role.

You can upload custom scripts and gather event details about the alert.

If you do not see this capability available in the user interface it is because the functionality has been disabled by your administrator. If required, you can enable this functionality from **Storage Management** > **Feature Settings**.

Steps

1. In the left navigation pane, click **Storage Management** > **Scripts**.
2. In the **Scripts** page, click **Add**.
3. In the **Add Script** dialog box, click **Browse** to select your script file.
4. Enter a description for the script that you select.
5. Click **Add**.

Delete scripts in Active IQ Unified Manager

You can delete a script from Unified Manager when the script is no longer required or valid.

Before you begin

- You must have the Application Administrator or Storage Administrator role.
- The script must not be associated with an alert.

Steps

1. In the left navigation pane, click **Storage Management** > **Scripts**.
2. In the **Scripts** page, select the script that you want to delete, and then click **Delete**.
3. In the **Warning** dialog box, confirm the deletion by clicking **Yes**.

Test script execution in Active IQ Unified Manager

You can verify that your script is executed correctly when an alert is generated for a storage object.

- You must have the Application Administrator or Storage Administrator role.
- You must have uploaded a script in the supported file format to Unified Manager.

Steps

1. In the left navigation pane, click **Storage Management** > **Scripts**.
2. In the Scripts page, add your test script.
3. In the left navigation pane, click **Storage Management** > **Alert Setup**.
4. In the **Alert Setup** page, perform one of the following actions:

To…	Do this…
Add an alert	a. Click **Add**. b. In the Actions section, associate the alert with your test script.
Edit an alert	a. Select an alert, and then click **Edit**. b. In the Actions section, associate the alert with your test script.

5. Click **Save**.
6. In the **Alert Setup** page, select the alert that you added or modified, and then click **Test**.

   The script is executed with the “-test” argument, and a notification alert is sent to the email addresses that were specified when the alert was created.

Manage and monitor groups

Understand groups

Learn about groups in Active IQ Unified Manager

A group is a dynamic collection of heterogeneous storage objects (clusters, SVMs, or volumes). You can create groups in Unified Manager to easily manage a set of storage objects. The members in a group might change, depending on the storage objects that are monitored by Unified Manager at a point in time.

- Each group has a unique name.
- You must configure a minimum of one group rule for each group.
- You can associate a group with more than one group rule.
- Each group can include multiple types of storage objects such as clusters, SVMs, or volumes.
- Storage objects are dynamically added to a group based on when a group rule is created or when Unified Manager completes a monitoring cycle.
- You can simultaneously apply actions on all the storage objects in a group such as setting thresholds for volumes.

Understand how group rules work for groups in Active IQ Unified Manager

A group rule is a criterion that you define to enable storage objects (volumes, clusters, or

SVMs) to be included in a specific group. You can use condition groups or conditions for defining group rule for a group.

- You must associate a group rule to a group.
- You must associate an object type for a group rule; only one object type is associated for a group rule.
- Storage objects are added or removed from the group after each monitoring cycle or when a rule is created, edited, or deleted.
- A group rule can have one or more condition groups, and each condition group can have one or more conditions.
- Storage objects can belong to multiple groups based on group rules you create.

Conditions

You can create multiple condition groups, and each condition group can have one or more conditions. You can apply all the defined condition groups in a group rule for groups in order to specify which storage objects are included in the group.

Conditions within a condition group are executed using logical AND. All the conditions in a condition group must be met. When you create or modify a group rule, a condition is created that applies, selects, and groups only those storage objects that satisfy all conditions in the condition group. You can use multiple conditions within a condition group when you want to narrow the scope of which storage objects to include in a group.

You can create conditions with storage objects by using the following operands and operator and specifying the required value.

Storage object type	Applicable operands
Volume	• Object name • Owning cluster name • Owning SVM name • Annotations
SVM	• Object name • Owning cluster name • Annotations
Cluster	• Object name • Annotations

When you select annotation as an operand for any storage object, the “Is” operator is available. For all other operands, you can select either “Is” or “Contains” as operator.

- Operand

  The list of operands in Unified Manager changes based on the selected object type. The list includes the object name, owning cluster name, owning SVM name, and annotations that you define in Unified Manager.

- Operator

  The list of operators changes based on the selected operand for a condition. The operators supported in

Unified Manager are “Is” and “Contains”.

When you select the “Is” operator, the condition is evaluated for exact match of operand value to the value provided for the selected operand.

When you select the “Contains” operator, the condition is evaluated to meet one of the following criteria:

  - The operand value is an exact match to the value provided for the selected operand
  - The operand value contains the value provided for the selected operand

- Value

  The value field changes based on the operand selected.

Example of a group rule with conditions

Consider a condition group for a volume with the following two conditions:

- Name contains “vol”
- SVM name is “data_svm”

This condition group selects all volumes that include “vol” in their names and that are hosted on SVMs with the name “data_svm”.

Condition groups

Condition groups are executed using logical OR, and then applied to storage objects. The storage objects must satisfy one of the condition groups to be included in a group. The storage objects of all the condition groups are combined. You can use condition groups to increase the scope of storage objects to include in a group.

Example of a group rule with condition groups

Consider two condition groups for a volume, with each group containing the following two conditions:

- Condition group 1
  - Name contains “vol”
  - SVM name is “data_svm”
    Condition group 1 selects all volumes that include “vol” in their names and that are hosted on SVMs with the name “data_svm”.
- Condition group 2
  - Name contains “vol”
  - The annotation value of data-priority is “critical”
    Condition group 2 selects all volumes that include “vol” in their names and that are annotated with the data-priority annotation value as “critical”.

When a group rule containing these two condition groups is applied on storage objects, then the following storage objects are added to a selected group:

- All volumes that include “vol” in their names and that are hosted on the SVM with the name “data_svm”.
- All volumes that include “vol” in their names and that are annotated with the data-priority annotation value “critical”.

Understand how group actions work on storage objects in Active IQ Unified Manager

A group action is an operation that is performed on all the storage objects in a group. For example, you can configure volume threshold group action to simultaneously change the volume threshold values of all volumes in a group.

Groups support unique group action types. You can have a group with only one volume health threshold group action type. However, you can configure a different type of group action, if available, for the same group. The rank of a group action determines the order in which the action is applied to storage objects. The details page of a storage object provides information about which group action is applied on the storage object.

Example of unique group actions

Consider a volume A that belongs to groups G1 and G2, and the following volume health threshold group actions are configured for these groups:

- `Change_capacity_threshold` group action with rank 1, for configuring the capacity of the volume
- `Change_snapshot_copies` group action with rank 2, for configuring the Snapshot copies of the volume

The `Change_capacity_threshold` group action always takes priority over the `Change_snapshot_copies` group action and is applied to volume A. When Unified Manager completes one cycle of monitoring, the health threshold related events of volume A are re-evaluated per the `Change_capacity_threshold` group action. You cannot configure another volume threshold type of group action for either G1 or G2 group.

Add groups in Active IQ Unified Manager

You can create groups to combine clusters, volumes, and storage virtual machines (SVMs) for ease of management.

Before you begin

You must have the Application Administrator or Storage Administrator role.

You can define group rules to add or remove members from the group and to modify group actions for the group.

Steps

1. In the left navigation pane, click **Storage Management** > **Groups**.
2. In the **Groups** tab, click **Add**.
3. In the **Add Group** dialog box, enter a name and description for the group.
4. Click **Add**.

Edit groups in Active IQ Unified Manager

You can edit the name and description of a group that you created in Unified Manager.

Before you begin

You must have the Application Administrator or Storage Administrator role.

When you edit a group to update the name, you must specify a unique name; you cannot use an existing group name.

Steps

1. In the left navigation pane, click **Storage Management** > **Groups**.
2. In the **Groups** tab, select the group that you want to edit, and then click **Edit**.
3. In the **Edit Group** dialog box, change the name, description, or both for the group.
4. Click **Save**.

Delete groups in Active IQ Unified Manager

You can delete a group from Unified Manager when the group is no longer required.

Before you begin

- None of the storage objects (clusters, SVMs, or volumes) must be associated with any group rule that is associated with the group that you want to delete.
- You must have the Application Administrator or Storage Administrator role.

Steps

1. In the left navigation pane, click **Storage Management** > **Groups**.
2. In the **Groups** tab, select the group that you want to delete, and then click **Delete**.
3. In the **Warning** dialog box, confirm the deletion by clicking **Yes**.

   Deleting a group does not delete the group actions that are associated with the group. However, these group actions will be unmapped after the group is deleted.

Add group rules in Active IQ Unified Manager

You can create group rules for a group to dynamically add storage objects such as volumes, clusters, or storage virtual machines (SVMs) to the group. You must configure at least one condition group with at least one condition to create a group rule.

Before you begin

You must have the Application Administrator or Storage Administrator role.

Storage objects that are currently monitored are added as soon as the group rule is created. New objects are added only after the monitoring cycle is completed.

Steps

1. In the left navigation pane, click **Storage Management** > **Groups**.
2. In the **Group Rules** tab, click **Add**.
3. In the **Add Group Rule** dialog box, specify a name for the group rule.
4. In the **Target Object Type** field, select the type of storage object that you want to group.
5. In the **Group** field, select the required group for which you want to create group rules.
6. In the **Conditions** section, perform the following steps to create a condition, a condition group, or both:

To create….	Do this…
A condition	a. Select an operand from the list of operands. b. Select either **Contains** or **Is** as the operator. c. Enter a value, or select a value from the available list.
A condition group	a. Click **Add Condition Group** b. Select an operand from the list of operands. c. Select either **Contains** or **Is** as the operator. d. Enter a value, or select a value from the available list. e. Click **Add condition** to create more conditions if required, and repeat steps a through d for each condition.

7. Click **Add**.

Example for creating a group rule

Perform the following steps in the Add Group Rule dialog box to create a group rule, including configuring a condition and adding a condition group:

Steps

1. Specify a name for the group rule.
2. Select the object type as storage virtual machine (SVM).
3. Select a group from the list of groups.
4. In the Conditions section, select **Object Name** as the operand.
5. Select **Contains** as the operator.
6. Enter the value as `svm_data`.
7. Click **Add condition group**.
8. Select **Object Name** as the operand.
9. Select **Contains** as the operator.
10. Enter the value as `vol`.
11. Click **Add condition**.
12. Repeat steps 8 through 10 by selecting **data-priority** as the operand in step 8, **Is** as the operator in step 9, and **critical** as the value in step 10.
13. Click **Add** to create the condition for the group rule.

Edit group rules in Active IQ Unified Manager

You can edit group rules to modify the condition groups and the conditions within a condition group to add or remove storage objects to or from a specific group.

Before you begin

You must have the Application Administrator or Storage Administrator role.

Steps

1. In the left navigation pane, click **Storage Management** > **Groups**.
2. In the **Group Rules** tab, select the group rule that you want to edit, and then click **Edit**.
3. In the **Edit Group Rule** dialog box, change the group rule name, associated group name, condition groups, and conditions as required.

   You cannot change the target object type for a group rule.

4. Click **Save**.

Delete group rules in Active IQ Unified Manager

You can delete a group rule from Active IQ Unified Manager when the group rule is no longer required.

Before you begin

You must have the Application Administrator or Storage Administrator role.

When a group rule is deleted, the associated storage objects will be removed from the group.

Steps

1. In the left navigation pane, click **Storage Management** > **Groups**.
2. In the **Group Rules** tab, select the group rule that you want to delete, and then click **Delete**.
3. In the **Warning** dialog box, confirm the deletion by clicking **Yes**.

Add group actions in Active IQ Unified Manager

You can configure group actions that you want to apply to storage objects in a group. Configuring actions for a group enables you to save time, because you do not have to add these actions to each object individually.

Before you begin

You must have the Application Administrator or Storage Administrator role.

Steps

1. In the left navigation pane, click **Storage Management** > **Groups**.
2. In the **Group Actions** tab, click **Add**.
3. In the **Add Group Action** dialog box, enter a name and description for the action.
4. From the **Group** menu, select a group for which you want to configure the action.
5. From the **Action Type** menu, select an action type.

   The dialog box expands, enabling you to configure the selected action type with required parameters.

6. Enter appropriate values for the required parameters to configure a group action.
7. Click **Add**.

Edit group actions in Active IQ Unified Manager

You can edit the group action parameters that you configured in Unified Manager, such as the group action name, description, associated group name, and parameters of the action type.

Before you begin

You must have the Application Administrator or Storage Administrator role.

Steps

1. In the left navigation pane, click **Storage Management** > **Groups**.
2. In the **Group Actions** tab, select the group action that you want to edit, and then click **Edit**.
3. In the **Edit Group Action** dialog box, change the group action name, description, associated group name, and parameters of the action type, as required.
4. Click **Save**.

Configure volume health thresholds for groups in Active IQ Unified Manager

You can configure group-level volume health thresholds for capacity, Snapshot copies, qtree quotas, growth, and inodes.

Before you begin

You must have the Application Administrator or Storage Administrator role.

The volume health threshold type of group action is applied only on volumes of a group.

Steps

1. In the left navigation pane, click **Storage Management** > **Groups**.
2. In the **Group Actions** tab, click **Add**.
3. Enter a name and description for the group action.
4. From the **Group** drop-down box, select a group for which you want to configure group action.
5. Select **Action Type** as the volume health threshold.
6. Select the category for which you want to set the threshold.
7. Enter the required values for the health threshold.
8. Click **Add**.

Delete group actions in Active IQ Unified Manager

You can delete a group action from Unified Manager when the group action is no longer required.

Before you begin

You must have the Application Administrator or Storage Administrator role.

When you delete the group action for the volume health threshold, global thresholds are applied to the storage objects in that group. Any object-level health thresholds that are set on the storage object are not impacted.

Steps

1. In the left navigation pane, click **Storage Management** > **Groups**.
2. In the **Group Actions** tab, select the group action that you want to delete, and then click **Delete**.
3. In the **Warning** dialog box, confirm the deletion by clicking **Yes**.

Reorder group actions in Active IQ Unified Manager

You can change the order of the group actions that are to be applied to the storage objects in a group. Group actions are applied to storage objects sequentially based on their rank. The lowest rank is assigned to the group action that you configured last. You can change the rank of the group action depending on your requirements.

Before you begin

You must have the Application Administrator or Storage Administrator role.

You can select either a single row or multiple rows, and then perform multiple drag-and-drop operations to change the rank of group actions. However, you must save the changes for the re-prioritization to be reflected in the group actions grid.

Steps

1. In the left navigation pane, click **Storage Management** > **Groups**.
2. In the **Group Actions** tab, click **Reorder**.
3. In the **Reorder Group Actions** dialog box, drag and drop the rows to rearrange the sequence of group actions as required.
4. Click **Save**.

Prioritize storage object events using annotations

Understand more about annotations

Learn about annotations in Active IQ Unified Manager

An annotation is a text string (the name) that is assigned to another text string (the value). Each annotation name-value pair can be dynamically associated with storage objects using annotation rules. When you associate storage objects with predefined annotations, you can filter and view the events that are related to them. You can apply annotations to clusters, volumes, and storage virtual machines (SVMs).

Each annotation name can have multiple values; each name-value pair can be associated with a storage object through rules.

For example, you can create an annotation named “data-center” with the values “Boston” and “Canada”. You can then apply the annotation “data-center” with the value “Boston” to volume v1. When an alert is generated for any event on a volume v1 that is annotated with “data-center”, the generated email indicates the location of the volume, “Boston”, and this enables you to prioritize and resolve the issue.

Understand how annotation rules work in Active IQ Unified Manager

An annotation rule is a criterion that you define to annotate storage objects (volumes, clusters, or storage virtual machines (SVMs)). You can use either condition groups or

conditions for defining annotation rules.

- You must associate an annotation rule to an annotation.
- You must associate an object type for an annotation rule; only one object type can be associated for an annotation rule.
- Unified Manager adds or removes annotations from storage objects after each monitoring cycle or when a rule is created, edited, deleted, or reordered.
- An annotation rule can have one or more condition groups, and each condition group can have one or more conditions.
- Storage objects can have multiple annotations. An annotation rule for a particular annotation can also use different annotations in the rule conditions to add another annotation to already annotated objects.

Conditions

You can create multiple condition groups, and each condition group can have one or more conditions. You can apply all the defined condition groups in an annotation rule of an annotation in order to annotate storage objects.

Conditions within a condition group are executed using logical AND. All the conditions in a condition group must be met. When you create or modify an annotation rule, a condition is created that applies, selects, and annotates only those storage objects that meet all the conditions in the condition group. You can use multiple conditions within a condition group when you want to narrow the scope of which storage objects to annotate.

You can create conditions with storage objects by using the following operands and operator and specifying the required value.

Storage object type	Applicable operands
Volume	• Object name • Owning cluster name • Owning SVM name • Annotations
SVM	• Object name • Owning cluster name • Annotations
Cluster	• Object name • Annotations

When you select annotation as an operand for any storage object, the “Is” operator is available. For all other operands, you can select either “Is” or “Contains” as operator. When you select the “Is” operator, the condition is evaluated for an exact match of the operand value with the value provided for the selected operand. When you select the “Contains” operator, the condition is evaluated to meet one of the following criteria:

- The operand value is an exact match to the value of the selected operand.
- The operand value contains the value provided for the selected operand.

Example of an annotation rule with conditions

Consider an annotation rule with one condition group for a volume with the following two conditions:

- Name contains “vol”
- SVM name is “data_svm”

This annotation rule annotates all volumes that include “vol” in their names and that are hosted on SVMs with the name “data_svm” with the selected annotation and the annotation type.

Condition groups

Condition groups are executed using logical OR, and then applied to storage objects. The storage objects must meet the requirements of one of the condition groups to be annotated. The storage objects that meet the conditions of all the condition groups are annotated. You can use condition groups to increase the scope of storage objects to be annotated.

Example of an annotation rule with condition groups

Consider an annotation rule with two condition groups for a volume; each group contains the following two conditions:

- Condition group 1
  - Name contains “vol”
  - SVM name is “data_svm”
    This condition group annotates all volumes that include “vol” in their names and that are hosted on SVMs with the name “data_svm”.
- Condition group 2
  - Name contains “vol”
  - The annotation value of data-priority is “critical”
    This condition group annotates all volumes that include “vol” in their names and that are annotated with the data-priority annotation value as “critical”.

When an annotation rule containing these two condition groups is applied on storage objects, then the following storage objects are annotated:

- All volumes that include “vol” in their names and that are hosted on SVM with the name “data_svm”.
- All volumes that include “vol” in their names and that are annotated with the data-priority annotation value as “critical”.

Understand predefined annotation values in Active IQ Unified Manager

Data-priority is a predefined annotation that has the values Mission critical, High, and Low. These values enable you to annotate storage objects based on the priority of data that they contain. You cannot edit or delete the predefined annotation values.

- **Data-priority:Mission critical**

  This annotation is applied to storage objects that contain mission-critical data. For example, objects that contain production applications can be considered as mission critical.

- **Data-priority:High**

  This annotation is applied to storage objects that contain high-priority data. For example, objects that are hosting business applications can be considered high priority.

- **Data-priority:Low**

  This annotation is applied to storage objects that contain low-priority data. For example, objects that are on secondary storage, such as backup and mirror destinations, might be of low priority.

Add annotations dynamically in Active IQ Unified Manager

When you create custom annotations, Unified Manager dynamically associates clusters, storage virtual machines (SVMs), and volumes with the annotations by using rules. These rules automatically assign the annotations to storage objects.

Before you begin

You must have the Application Administrator or Storage Administrator role.

Steps

1. In the left navigation pane, click **Storage Management** > **Annotations**.
2. In the **Annotations** page, click **Add Annotation**.
3. In the **Add Annotation** dialog box, type a name and description for the annotation.
4. Optional: In the **Annotation Values** section, click **Add** to add values to the annotation.
5. Click **Save**.

Add values to annotations in Active IQ Unified Manager

You can add values to annotations, and then associate storage objects with a particular annotation name-value pair. Adding values to annotations helps you to manage storage objects more effectively.

Before you begin

You must have the Application Administrator or Storage Administrator role.

You cannot add values to predefined annotations.

Steps

1. In the left navigation pane, click **Storage Management** > **Annotations**.
2. In the **Annotations** page, select the annotation to which you want to add a value and then click **Add** in the **Values** section.
3. In the **Add Annotation Value** dialog box, specify a value for the annotation.

   The value that you specify must be unique for the selected annotation.

4. Click **Add**.

Delete annotations in Active IQ Unified Manager

You can delete custom annotations and their values when they are no longer required.

Before you begin

- You must have the Application Administrator or Storage Administrator role.
- The annotation values must not be used in other annotations or group rules.

Steps

1. In the left navigation pane, click **Storage Management** > **Annotations**.
2. In the **Annotations** tab, select the annotation that you want to delete.

   The details of the selected annotation are displayed.
3. Click **Actions** > **Delete** to delete the selected annotation and its value.
4. In the warning dialog box, click **Yes** to confirm the deletion.

View the annotation list and details in Active IQ Unified Manager

You can view the list of annotations that are dynamically associated with clusters, volumes, and storage virtual machines (SVMs). You can also view details such as the description, created by, created date, values, rules, and the objects associated with the annotation.

Steps

1. In the left navigation pane, click **Storage Management** > **Annotations**.
2. In the **Annotations** tab, click the annotation name to view the associated details.

Delete values from annotations in Active IQ Unified Manager

You can delete values associated with custom annotations when that value no longer applies to the annotation.

Before you begin

- You must have the Application Administrator or Storage Administrator role.
- The annotation value must not be associated with any annotation rules or group rules.

You cannot delete values from predefined annotations.

Steps

1. In the left navigation pane, click **Storage Management** > **Annotations**.
2. In the annotations list in the **Annotations** tab, select the annotation from which you want to delete a value.
3. In the **Values** area of the **Annotations** tab, select the value you want to delete, and then click **Delete**.
4. In the **Warning** dialog box, click **Yes**.

   The value is deleted and no longer displayed in the list of values for the selected annotation.

Create annotation rules in Active IQ Unified Manager

You can create annotation rules that Unified Manager uses to dynamically annotate storage objects such as volumes, clusters, or storage virtual machines (SVMs).

Before you begin

You must have the Application Administrator or Storage Administrator role.

Storage objects that are currently monitored are annotated as soon as the annotation rule is created. New objects are annotated only after the monitoring cycle is completed.

Steps

1. In the left navigation pane, click **Storage Management** > **Annotations**.
2. In the **Annotation Rules** tab, click **Add**.
3. In the **Add Annotation Rule** dialog box, specify a name for the annotation rule.
4. In the **Target Object Type** field, select the type of storage object that you want to annotate.
5. In the **Apply Annotation** fields, select the annotation and annotation value that you want to use.
6. In the Conditions section, perform the appropriate action to create a condition, a condition group, or both:

To create…	Do this…
A condition	a. Select an operand from the list of operands. b. Select either **Contains** or **Is** as the operator. c. Enter a value, or select a value from the available list.
A condition group	a. Click **Add Condition Group**. b. Select an operand from the list of operands. c. Select either **Contains** or **Is** as the operator. d. Enter a value, or select a value from the available list. e. Click **Add condition** to create more conditions if required, and repeat steps a through d for each condition.

7. Click **Add**.

Example of creating an annotation rule

Perform the following steps in the Add Annotation Rule dialog box to create an annotation rule, including configuring a condition and adding a condition group:

Steps

1. Specify a name for the annotation rule.
2. Select the target object type as storage virtual machine (SVM).
3. Select an annotation from the list of annotations, and specify a value.
4. In the Conditions section, select **Object Name** as the operand.

5. Select **Contains** as the operator.
6. Enter the value as `svm_data`.
7. Click **Add condition group**.
8. Select **Object Name** as the operand.
9. Select **Contains** as the operator.
10. Enter the value as `vol`.
11. Click **Add condition**.
12. Repeat steps 8 through 10 by selecting **data-priority** as the operand in step 8, **Is** as the operator in step 9, and **mission-critical** as the value in step 10.
13. Click **Add**.

Add annotations manually to individual storage objects in Active IQ Unified Manager

You can manually annotate selected volumes, clusters, and SVMs without using annotation rules. You can annotate a single storage object or multiple storage objects, and specify the required name-value pair combination for the annotation.

Before you begin

You must have the Application Administrator or Storage Administrator role.

Steps

1. Navigate to the storage objects you want to annotate:

To add annotation to…	Do this…
Clusters	a. Click **Storage** > **Clusters**. b. Select one or more clusters.
Volumes	a. Click **Storage** > **Volumes**. b. Select one or more volumes.
SVMs	a. Click **Storage** > **SVMs**. b. Select one or more SVMs.

2. Click **Annotate** and select a name-value pair.
3. Click **Apply**.

Edit annotation rules in Active IQ Unified Manager

You can edit annotation rules to modify the condition groups and conditions within the condition group to add annotations to or remove annotations from storage objects.

Before you begin

You must have the Application Administrator or Storage Administrator role.

Annotations are dissociated from storage objects when you edit the associated annotation rules.

Steps

1. In the left navigation pane, click **Storage Management** > **Annotations**.
2. In the **Annotation Rules** tab, select the annotation rule you want to edit, and then click **Actions** > **Edit**.
3. In the **Edit Annotation Rule** dialog box, change the rule name, annotation name and value, condition groups, and conditions as required.

   You cannot change the target object type for an annotation rule.

4. Click **Save**.

Configure conditions for annotation rules in Active IQ Unified Manager

You can configure one or more conditions to create annotation rules that Unified Manager applies on the storage objects. The storage objects that satisfy the annotation rule are annotated with the value specified in the rule.

Before you begin

You must have the Application Administrator or Storage Administrator role.

Steps

1. In the left navigation pane, click **Storage Management** > **Annotations**.
2. In the **Annotation Rules** tab, click **Add**.
3. In the **Add Annotation Rule** dialog box, enter a name for the rule.
4. Select one object type from the Target Object Type list, and then select an annotation name and value from the list.
5. In the **Conditions** section of the dialog box, select an operand and an operator from the list and enter a condition value, or click **Add Condition** to create a new condition.
6. Click **Save and Add**.

Example of configuring a condition for an annotation rule

Consider a condition for the object type SVM, where the object name contains “svm_data”.

Perform the following steps in the Add Annotation Rule dialog box to configure the condition:

Steps

1. Enter a name for the annotation rule.
2. Select the target object type as SVM.
3. Select an annotation from the list of annotations and a value.
4. In the **Conditions** field, select **Object Name** as the operand.
5. Select **Contains** as the operator.
6. Enter the value as `svm_data`.
7. Click **Add**.

Delete annotation rules in Active IQ Unified Manager

You can delete annotation rules from Active IQ Unified Manager when the rules are no longer required.

Before you begin

You must have the Application Administrator or Storage Administrator role.

When you delete an annotation rule, the annotation is disassociated and removed from the storage objects.

Steps

1. In the left navigation pane, click **Storage Management** > **Annotations**.
2. In the **Annotation Rules** tab, select the annotation rule that you want to delete, and then click **Delete**.
3. In the **Warning** dialog box, click **Yes** to confirm the deletion.

Reorder annotation rules in Active IQ Unified Manager

You can change the order in which Unified Manager applies annotation rules to storage objects. Annotation rules are applied to storage objects sequentially based on their rank. When you configure an annotation rule, the rank is least. But you can change the rank of the annotation rule depending on your requirements.

Before you begin

You must have the Application Administrator or Storage Administrator role.

You can select either a single row or multiple rows and perform many drag-and-drop operations to change the rank of annotation rules. However, you must save the changes for the reprioritization to be displayed in the Annotation Rules tab.

Steps

1. In the left navigation pane, click **Storage Management** > **Annotations**.
2. In the **Annotation Rules** tab, click **Reorder**.
3. In the **Reorder Annotation Rule** dialog box, drag and drop single or multiple rows to rearrange the sequence of the annotation rules.
4. Click **Save**.

   You must save the changes for the reorder to be displayed.

Send a support bundle through web UI and maintenance console

Send a support bundle through the web UI and maintenance console in Active IQ Unified Manager

You should send a support bundle when the issue that you have requires more detailed diagnosis and troubleshooting than an AutoSupport message provides. You can send a support bundle to technical support using the Unified Manager web Ui and maintenance console.

Unified Manager stores a maximum of two full support bundles and three lightweight support bundles at one time.

Related information

Unified Manager user roles and capabilities

Send AutoSupport messages to technical support in Active IQ Unified Manager

The AutoSupport page enables you to send predefined and on-demand AutoSupport messages to your technical support team to ensure a correct operation of your environment, and to assist you in maintaining the integrity of your environment. AutoSupport is enabled by default and it should not be disabled, for you to receive the benefits of NetAppActive IQ.

You can send diagnostic system information and detailed data about the Unified Manager server in a message as and when required, schedule a message to be sent periodically, or even generate and send support bundles to the technical support team.

> A user with a storage administrator role can generate and send on-demand AutoSupport messages and support bundles to technical support. However, only an administrator or maintenance user can enable or disable periodic AutoSupport and configure the HTTP settings as described in the Setting up HTTP proxy server section. In an environment that needs to use an HTTP proxy server, the configuration should be complete before a storage administrator can send on-demand AutoSupport messages and support bundles to technical support.

Send on-demand AutoSupport messages

You can generate and send an on-demand message to technical support, or to a specified email recipient, or to both.

Steps

1. Navigate to **General** > **AutoSupport**, and perform one or both of the following actions:
2. If you want to send the AutoSupport message to technical support, select the **Send to Technical Support** check box.
3. If you want to send the AutoSupport message to a specific email recipient, select the **Send to Email Recipient** check box, and enter the email address of the recipient.
4. Click **Save**.
5. Click **Generate and Send AutoSupport**.

Enable periodic AutoSupport

You can send specific, predefined messages to technical support for issue diagnosis and resolution periodically. This functionality is enabled by default. If disabled, an administrator or maintenance user can enable the settings.

Steps

1. Navigate to **General** > **AutoSupport**.
2. In the Periodic AutoSupport section, select the **Enable Sending AutoSupport Data Periodically to Active IQ** check box.
3. If required, define the name, port, and authentication information for the HTTP proxy server as described in Setting up HTTP proxy server section.
4. Click **Save**.

Upload on-demand support bundle

You can generate and send a support bundle to technical support based on the requirement for troubleshooting. Unified Manager stores only the two most recently generated support bundles. Older support bundles are deleted from the system.

Because some types of support data can use a large amount of cluster resources or take a long time to complete, when you select the full support bundle, you can include or exclude specific data types to reduce the support bundle size. You also have the option to create a lightweight support bundle that contains just 30 days of logs and configuration database records — it excludes performance data, acquisition recording files, and server heap dump.

Steps

1. Navigate to **General** > **AutoSupport**.
2. In the On-Demand Support Bundle section, click **Generate and Send Support Bundle**.
3. To send a light support bundle to technical support, in the Generate and Send Support Bundle pop-up, select the**Generate light support bundle** check box.
4. Alternately, to send a full support bundle, select the **Generate full support bundle** check box. Select the specific data types to include or exclude in the support bundle.

   Even if you do not select any data type, the support bundle is still generated with other Unified Manager data.

5. Select the **Send the bundle to technical support** check box to generate and send the bundle to technical support. If you do not select this check box, the bundle is generated and stored locally in the Unified Manager server. The generated support bundle is available for later use in the /support directory on VMware systems, in `/opt/netapp/data/support/` on Linux systems, and in `ProgramData\NetApp\OnCommandAppData\ocum\support` on Windows systems.
6. Click **Send**.

Set up HTTP proxy server

You can designate a proxy to provide the internet access in order to send AutoSupport content to support if your environment does not provide direct access from the Unified Manager server. This section is available for only administrator and maintenance users.

- **Use HTTP proxy**

  Check this box to identify the server being used as the HTTP proxy.

  Enter the host name or IP address of the proxy server, and the port number used to connect to the server.

- **Use authentication**

  Check this box if you need to provide authentication information to access the server being used as the HTTP proxy.

  Enter the user name and the password required to authenticate with the HTTP proxy.

HTTP proxies that provide only Basic Authentication are not supported.

Access the maintenance console in Active IQ Unified Manager

If the Unified Manager user interface is not in operation, or if you need to perform functions that are not available in the user interface, you can access the maintenance console to manage your Unified Manager system.

Before you begin

You must have installed and configured Unified Manager.

After 15 minutes of inactivity, the maintenance console logs you out.

When installed on VMware, if you have already logged in as the maintenance user through the VMware console, you cannot simultaneously log in using Secure Shell.

Step

1. Follow these steps to access the maintenance console:

On this operating system…	Follow these steps…
VMware	a. Using Secure Shell, connect to the IP address or fully qualified domain name of the Unified Manager virtual appliance. b. Log in to the maintenance console using your maintenance user name and password.
Linux	a. Using Secure Shell, connect to the IP address or fully qualified domain name of the Unified Manager system. b. Log in to the system with the maintenance user (umadmin) name and password. c. Enter the command `maintenance_console` and press Enter.
Windows	a. Log in to the Unified Manager system with administrator credentials. b. Launch PowerShell as a Windows administrator. c. Enter the command `maintenance_console` and press Enter.

The Unified Manager maintenance console menu is displayed.

Generate and upload a support bundle from Active IQ Unified Manager

You can generate a support bundle, containing diagnostic information, so that you can send it to technical support for troubleshooting help.

Starting with Unified Manager 9.8, if your Unified Manager server is connected to the internet, you can also upload the support bundle to NetApp from the maintenance console.

Before you begin

You must have access to the maintenance console as the maintenance user.

Because some types of support data can use a large amount of cluster resources or take a long time to complete, when you select the full support bundle you can specify data types to include or exclude to reduce the support bundle size. You also have the option to create a lightweight support bundle that contains just 30 days of logs and configuration database records — it excludes performance data, acquisition recording files, and server heap dump.

Unified Manager stores only the two most recently generated support bundles. Older support bundles are deleted from the system.

Steps

1. In the maintenance console **Main Menu**, select **Support/Diagnostics**.
2. Select **Generate Light Support Bundle** or **Generate Support Bundle** depending on the level of details you want to have in the support bundle.
3. If you choose the full support bundle, select or deselect the following data types to include or exclude in the support bundle:

   - **database dump**

     A dump of the MySQL Server database.

   - **heap dump**

     A snapshot of the state of the main Unified Manager server processes. This option is disabled by default and should be selected only when requested by customer support.

   - **acquisition recordings**

     A recording of all communications between Unified Manager and the monitored clusters.

   If you deselect all data types, the support bundle is still generated with other Unified Manager data.

4. Type `g`, and then press Enter to generate the support bundle.

   Since the generation of a support bundle is a memory intensive operation, you are prompted to verify that you are sure you want to generate the support bundle at this time.

5. Type `y`, and then press Enter to generate the support bundle.

   If you do not want to generate the support bundle at this time, type `n`, and then press Enter.

6. If you included database dump files in the full support bundle, you are prompted to specify the time period for which you want performance statistics included. Including performance statistics can take a lot of time and space, so you can also dump the database without including performance statistics:

   a. Enter the starting date in the format YYYYMMDD.

      For example, enter `20210101` for January 1, 2021. Enter `n` if you do not want performance statistics to be included.

   b. Enter the number of days of statistics to include, beginning from 12 a.m. on the specified starting date.

      You can enter a number from 1 through 10.

If you are including performance statistics, the system displays the period of time for which performance statistics will be collected.

7. After the support bundle is created you are prompted whether you want to upload it to NetApp. Type `y`, and then press Enter.

   You are prompted to enter your support case number.

8. If you have a case number already, enter the number and press Enter. Otherwise just press Enter.

   The support bundle is uploaded to NetApp.

If your Unified Manager server is not connected to the internet, or if you are unable to upload the support bundle for any other reason, then you can retrieve it and send it manually. You can retrieve it using an SFTP client or by using UNIX or Linux CLI commands. On Windows installations you can use Remote Desktop (RDP) to retrieve the support bundle.

The generated support bundle resides in the /support directory on VMware systems, in /opt/netapp/data/support/ on Linux systems, and in ProgramData\NetApp\OnCommandAppData\ocum\support on Windows systems.

Related information

Unified Manager user roles and capabilities

Retrieve the support bundle using a Windows client in Active IQ Unified Manager

If you are a Windows user, you can download and install a tool to retrieve the support bundle from your Unified Manager server. You can send the support bundle to technical support for a more detailed diagnosis of an issue. Filezilla or WinSCP are examples of tools you can use.

Before you begin

You must be the maintenance user to perform this task.

You must use a tool that supports SCP or SFTP.

Steps

1. Download and install a tool to retrieve the support bundle.
2. Open the tool.
3. Connect to your Unified Manager management server over SFTP.

   The tool displays the contents of the /support directory and you can view all existing support bundles.

4. Select the destination directory for the support bundle you want to copy.
5. Select the support bundle you want to copy and use the tool to copy the file from the Unified Manager server to your local system.

Retrieve the support bundle using a UNIX or Linux client in Active IQ Unified Manager

If you are a UNIX or Linux user, you can retrieve the support bundle from your vApp by using the command-line interface (CLI) on your Linux client server. You can use either

SCP or SFTP to retrieve the support bundle.

Before you begin

You must be the maintenance user to perform this task.

You must have generated a support bundle using the maintenance console and have the support bundle name available.

Steps

1. Access the CLI through Telnet or the console, using your Linux client server.
2. Access the `/support` directory.
3. Retrieve the support bundle and copy it to the local directory using the following command:

If you are using…	Then use the following command…
SCP	`scp <maintenance-user>@<vApp-name-or-ip>:/support/support_bundle_file_name.7z <destination-directory>`
SFTP	`sftp <maintenance-user>@<vApp-name-or-ip>:/support/support_bundle_file_name.7z <destination-directory>`

   The name of the support bundle is provided to you when you generate it using the maintenance console.

4. Enter the maintenance user password.

Examples

The following example uses SCP to retrieve the support bundle:

```
        `$ scp
admin@10.10.12.69:/support/support_bundle_20160216_145359.7z .`
Password: `<maintenance_user_password>`
support_bundle_20160216_145359.7z   100%  119MB  11.9MB/s   00:10
```

The following example uses SFTP to retrieve the support bundle:

```
        `$ sftp
admin@10.10.12.69:/support/support_bundle_20160216_145359.7z .`
Password: `<maintenance_user_password>`
Connected to 10.228.212.69.
Fetching /support/support_bundle_20130216_145359.7z to
./support_bundle_20130216_145359.7z
/support/support_bundle_20160216_145359.7z
```

Send a support bundle to technical support from Active IQ Unified Manager

When an issue requires more detailed diagnosis and troubleshooting information than an AutoSupport message provides, you can send a support bundle to technical support.

Before you begin

You must have access to the support bundle to send it to technical support.

You must have a case number generated through the technical support web site.

Steps

1. Log in to the NetApp Support Site.
2. Upload the file.

   How to upload a file to NetApp

Tasks and information related to several workflows

Understand cluster component contention in Active IQ Unified Manager

You can identify cluster performance issues when a cluster component goes into contention. The performance of workloads that use the component slow down and their response time (latency) for client requests increases, which triggers an event in Unified Manager.

A component that is in contention cannot perform at an optimal level. Its performance has declined, and the performance of other cluster components and workloads, called *victims*, might have increased latency. To bring a component out of contention, you must reduce its workload or increase its ability to handle more work, so that the performance can return to normal levels. Because Unified Manager collects and analyzes workload performance in five-minute intervals, it detects only when a cluster component is consistently overused. Transient spikes of overusage that last for only a short duration within the five-minute interval are not detected.

For example, a storage aggregate might be under contention because one or more workloads on it are competing for their I/O requests to be fulfilled. Other workloads on the aggregate can be impacted, causing their performance to decrease. To reduce the amount of activity on the aggregate, there are different steps you can take, such as moving one or more workloads to a less busy aggregate or node, to lessen the overall workload demand on the current aggregate. For a QoS policy group, you can adjust the throughput limit, or move workloads to a different policy group, so that the workloads are no longer being throttled.

Unified Manager monitors the following cluster components to alert you when they are in contention:

- **Network**

  Represents the wait time of I/O requests by the external networking protocols on the cluster. The wait time is time spent waiting for “transfer ready” transactions to finish before the cluster can respond to an I/O request. If the network component is in contention, it means high wait time at the protocol layer is impacting the latency of one or more workloads.

- **Network Processing**

  Represents the software component in the cluster involved with I/O processing between the protocol layer and the cluster. The node handling network processing might have changed since the event was detected. If the network processing component is in contention, it means high utilization at the network processing

node is impacting the latency of one or more workloads.

When using an All SAN Array cluster in an active-active configuration, the network processing latency value is displayed for both nodes so you can verify the nodes are sharing the load equally.

- **QoS Limit Max**

  Represents the throughput maximum (peak) setting of the storage Quality of Service (QoS) policy group assigned to the workload. If the policy group component is in contention, it means all workloads in the policy group are being throttled by the set throughput limit, which is impacting the latency of one or more of those workloads.

- **QoS Limit Min**

  Represents the latency to a workload that is being caused by QoS throughput minimum (expected) setting assigned to other workloads. If the QoS minimum set on certain workloads use the majority of the bandwidth to guarantee the promised throughput, other workloads will be throttled and see more latency.

- **Cluster Interconnect**

  Represents the cables and adapters with which clustered nodes are physically connected. If the cluster interconnect component is in contention, it means high wait time for I/O requests at the cluster interconnect is impacting the latency of one or more workloads.

- **Data Processing**

  Represents the software component in the cluster involved with I/O processing between the cluster and the storage aggregate that contains the workload. The node handling data processing might have changed since the event was detected. If the data processing component is in contention, it means high utilization at the data processing node is impacting the latency of one or more workloads.

- **Volume Activation**

  Represents the process that tracks the usage of all active volumes. In large environments where more than 1000 volumes are active, this process tracks how many critical volumes need to access resources through the node at the same time. When the number of concurrent active volumes exceeds the recommended maximum threshold, some of the non-critical volumes will experience latency as identified here.

- **MetroCluster Resources**

  Represents the MetroCluster resources, including NVRAM and interswitch links (ISLs), used to mirror data between clusters in a MetroCluster configuration. If the MetroCluster component is in contention, it means high write throughput from workloads on the local cluster or a link health issue is impacting the latency of one or more workloads on the local cluster. If the cluster is not in a MetroCluster configuration, this icon is not displayed.

- **Aggregate or SSD Aggregate Ops**

  Represents the storage aggregate on which the workloads are running. If the aggregate component is in contention, it means high utilization on the aggregate is impacting the latency of one or more workloads. An aggregate consists of all HDDs, or a mix of HDDs and SSDs (a Flash Pool aggregate), or a mix of HDDs and a cloud tier (a FabricPool aggregate). An “SSD Aggregate” consists of all SSDs (an all-flash aggregate), or a mix of SSDs and a cloud tier (a FabricPool aggregate).

- **Cloud Latency**

Represents the software component in the cluster involved with I/O processing between the cluster and the cloud tier on which user data is stored. If the cloud latency component is in contention, it means that a large amount of reads from volumes that are hosted on the cloud tier are impacting the latency of one or more workloads.

- **Sync SnapMirror**

  Represents the software component in the cluster involved with replicating user data from the primary volume to the secondary volume in a SnapMirror Synchronous relationship. If the sync SnapMirror component is in contention, it means that the activity from SnapMirror Synchronous operations are impacting the latency of one or more workloads.

View the Volume / Health details page in Active IQ Unified Manager

You can use the Volume / Health details page to view detailed information about a selected volume, such as capacity, storage efficiency, configuration, protection, annotation, and events generated. You can also view information about the related objects and related alerts for that volume.

You must have the Application Administrator or Storage Administrator role.

Command buttons

The command buttons enable you to perform the following tasks for the selected volume:

- **Switch to Performance View**

  Enables you to navigate to the Volume / Performance details page.

- **Actions**
  - Add Alert

    Enables you to add an alert to the selected volume.

  - Edit Thresholds

    Enables you to modify the threshold settings for the selected volume.

  - Annotate

    Enables you to annotate the selected volume.

  - Protect

    Enables you to create either SnapMirror or SnapVault relationships for the selected volume.

  - Relationship

    Enables you to execute the following protection relationship operations:

    - Edit

      Launches the Edit Relationship dialog box which enables you to change existing SnapMirror policies, schedules, and maximum transfer rates for an existing protection relationship.

- Abort

  Aborts transfers that are in progress for a selected relationship. Optionally, it enables you to remove the restart checkpoint for transfers other than the baseline transfer. You cannot remove the checkpoint for a baseline transfer.

- Quiesce

  Temporarily disables scheduled updates for a selected relationship. Transfers that are already in progress must complete before the relationship is quiesced.

- Break

  Breaks the relationship between the source and destination volumes and changes the destination to a read-write volume.

- Remove

  Permanently deletes the relationship between the selected source and destination. The volumes are not destroyed and the Snapshot copies on the volumes are not removed. This operation cannot be undone.

- Resume

  Enables scheduled transfers for a quiesced relationship. At the next scheduled transfer interval, a restart checkpoint is used, if one exists.

- Resynchronize

  Enables you to resynchronize a previously broken relationship.

- Initialize/Update

  Enables you to perform a first-time baseline transfer on a new protection relationship, or to perform a manual update if the relationship is already initialized.

- Reverse Resync

  Enables you to reestablish a previously broken protection relationship, reversing the function of the source and destination by making the source a copy of the original destination. The contents on the source are overwritten by the contents on the destination, and any data that is newer than the data on the common Snapshot copy is deleted.

◦ Restore

Enables you to restore data from one volume to another volume.
For information, see Restoring data using the Volume / Health details page.



The Restore button and the Relationship operation buttons are not available for volumes that are in synchronous protection relationships.

- **View Volumes**

Enables you to navigate to the Health: All Volumes view.

Capacity tab

The Capacity tab displays details about the selected volume, such as its physical capacity, logical capacity, threshold settings, quota capacity, and information about any volume move operation:

- **Capacity Physical**

  Details the physical capacity of the volume:

  - Snapshot Overflow

    Displays the data space that is consumed by the Snapshot copies.

  - Used

    Displays the space used by data in the volume.

  - Warning

    Indicates that the space in the volume is nearly full. If this threshold is breached, the Space Nearly Full event is generated.

  - Error

    Indicates that the space in the volume is full. If this threshold is breached, the Space Full event is generated.

  - Unusable

    Indicates that the Thin-Provisioned Volume Space At Risk event is generated and that the space in the thinly provisioned volume is at risk because of aggregate capacity issues. The unusable capacity is displayed only for thinly provisioned volumes.

  - Data graph

    Displays the total data capacity and the used data capacity of the volume.

    If autogrow is enabled, the data graph also displays the space available in the aggregate. The data graph displays the effective storage space that can be used by data in the volume, which can be one of the following:

    - Actual data capacity of the volume for the following conditions:
      - Autogrow is disabled.
      - Autogrow-enabled volume has reached the maximum size.
      - Autogrow-enabled thickly provisioned volume cannot grow further.
    - Data capacity of the volume after considering the maximum volume size (for thinly provisioned volumes and for thickly provisioned volumes when the aggregate has space for the volume to reach maximum size)
    - Data capacity of the volume after considering the next possible autogrow size (for thickly provisioned volumes that have an autogrow percentage threshold)

  - Snapshot copies graph

    This graph is displayed only when the used Snapshot capacity or the Snapshot reserve is not zero.

Both the graphs display the capacity by which the Snapshot capacity exceeds the Snapshot reserve if the used Snapshot capacity exceeds the Snapshot reserve.

- **Capacity Logical**

  Displays the logical space characteristics of the volume. The logical space indicates the real size of the data that is being stored on disk without applying the savings from using ONTAP storage efficiency technologies.

  - Logical Space Reporting

    Displays if the volume has logical space reporting configured. The value can be Enabled, Disabled, or Not applicable. “Not applicable” is displayed for volumes on older versions of ONTAP or on volumes that do not support logical space reporting.

  - Used

    Displays the amount of logical space that is being used by data in the volume, and the percentage of logical space used based on the total data capacity.

  - Logical Space Enforcement

    Displays whether logical space enforcement is configured for thinly provisioned volumes. When set to Enabled, the logical used size of the volume cannot be greater than the currently set physical volume size.

- **Autogrow**

  Displays whether the volume automatically grows when it is out of space.

- **Space Guarantee**

  Displays the FlexVol volume setting control when a volume removes free blocks from an aggregate. These blocks are then guaranteed to be available for writes to files in the volume. The space guarantee can be set to one of the following:

  - None

    No space guarantee is configured for the volume.

  - File

    Full size of sparsely written files (for example, LUNs) is guaranteed.

  - Volume

    Full size of the volume is guaranteed.

  - Partial

    The FlexCache volume reserves space based on its size. If the FlexCache volume’s size is 100 MB or more, the minimum space guarantee is set to 100 MB by default. If the FlexCache volume’s size is less than 100 MB, the minimum space guarantee is set to the FlexCache volume’s size. If the FlexCache volume’s size is grown later, the minimum space guarantee is not incremented.

The space guarantee is Partial when the volume is of type Data-Cache.

- **Details (Physical)**

  Displays the physical characteristics of the volume.

- **Total Capacity**

  Displays the total physical capacity in the volume.

- **Data Capacity**

  Displays the amount of physical space used by the volume (used capacity) and the amount of physical space that is still available (free capacity) in the volume. These values are also displayed as a percentage of the total physical capacity.

  When the Thin-Provisioned Volume Space At Risk event is generated for thinly provisioned volumes, the amount of space used by the volume (used capacity) and the amount of space that is available in the volume but cannot be used (unusable capacity) because of aggregate capacity issues is displayed.

- **Snapshot Reserve**

  Displays the amount of space used by the Snapshot copies (used capacity) and amount of space available for Snapshot copies (free capacity) in the volume. These values are also displayed as a percentage of the total snapshot reserve.

  When the Thin-Provisioned Volume Space At Risk event is generated for thinly provisioned volumes, the amount of space used by the Snapshot copies (used capacity) and the amount of space that is available in the volume but cannot be used for making Snapshot copies (unusable capacity) because of aggregate capacity issues is displayed.

- **Volume Thresholds**

  Displays the following volume capacity thresholds:

  - Nearly Full Threshold

    Specifies the percentage at which a volume is nearly full.

  - Full Threshold

    Specifies the percentage at which a volume is full.

- **Other Details**
  - Autogrow Max Size

    Displays the maximum size up to which the volume can automatically grow. The default value is 120% of the volume size on creation. This field is displayed only when autogrow is enabled for the volume.

  - Qtree Quota Committed Capacity

    Displays the space reserved in the quotas.

  - Qtree Quota Overcommitted Capacity

Displays the amount of space that can be used before the system generates the Volume Qtree Quota Overcommitted event.

- Fractional Reserve

  Controls the size of the overwrite reserve. By default, the fractional reserve is set to 100, indicating that 100 percent of the required reserved space is reserved so that the objects are fully protected for overwrites. If the fractional reserve is less than 100 percent, the reserved space for all the space-reserved files in that volume is reduced to the fractional reserve percentage.

- Snapshot Daily Growth Rate

  Displays the change (in percentage, or in KB, MB, GB, and so on) that occurs every 24 hours in the Snapshot copies in the selected volume.

- Snapshot Days to Full

  Displays the estimated number of days remaining before the space reserved for the Snapshot copies in the volume reaches the specified threshold.

  The Snapshot Days to Full field displays a Not Applicable value when the growth rate of the Snapshot copies in the volume is zero or negative, or when there is insufficient data to calculate the growth rate.

- Snapshot Autodelete

  Specifies whether Snapshot copies are automatically deleted to free space when a write to a volume fails because of lack of space in the aggregate.

- Snapshot Copies

  Displays information about the Snapshot copies in the volume.

  The number of Snapshot copies in the volume is displayed as a link. Clicking the link opens the Snapshot Copies on a Volume dialog box, which displays details of the Snapshot copies.

  The Snapshot copy count is updated approximately every hour; however, the list of Snapshot copies is updated at the time that you click the icon. This might result in a difference between the Snapshot copy count displayed in the topology and the number of Snapshot copies listed when you click the icon.

- **Volume Move**

  Displays the status of either the current or the last volume move operation that was performed on the volume, and other details, such as the current phase of the volume move operation which is in progress, source aggregate, destination aggregate, start time, end time, and estimated end time.

  Also displays the number of volume move operations that are performed on the selected volume. You can view more information about the volume move operations by clicking the **Volume Move History** link.

Configuration tab

The Configuration tab displays details about the selected volume, such as the export policy, RAID type, capacity and storage efficiency related features of the volume:

- **Overview**
  - Full Name

Displays the full name of the volume.

- Aggregates

  Displays the name of the aggregate on which the volume resides, or the number of aggregates on which the FlexGroup volume resides.

- Tiering Policy

  Displays the tiering policy set for the volume; if the volume is deployed on a FabricPool-enabled aggregate. The policy can be None, Snapshot Only, Backup, Auto, or All.

- Storage VM

  Displays the name of the SVM that contains the volume.

- Junction Path

  Displays the status of the path, which can be active or inactive. The path in the SVM to which the volume is mounted is also displayed. You can click the **History** link to view the most recent five changes to the junction path.

- Export Policy

  Displays the name of the export policy that is created for the volume. You can click the link to view details about the export policies, authentication protocols, and access enabled on the volumes that belong to the SVM.

- Style

  Displays the volume style. The volume style can be FlexVol or FlexGroup.

- Type

  Displays the type of the selected volume. The volume type can be Read-write, Load-sharing, Data-Protection, Data-cache, or Temporary.

- RAID Type

  Displays the RAID type of the selected volume. The RAID type can be RAID0, RAID4, RAID-DP, or RAID-TEC.

  Multiple RAID types may display for FlexGroup volumes because the constituent volumes for FlexGroups can be on aggregates of different types.

- SnapLock Type

  Displays the SnapLock Type of the aggregate that contains the volume.

- SnapLock Expiry

  Displays the expiry date of SnapLock volume.

- **Capacity**
  - Thin Provisioning

Displays whether thin provisioning is configured for the volume.

- Autogrow

  Displays whether the flexible volume grows automatically within an aggregate.

- Snapshot Autodelete

  Specifies whether Snapshot copies are automatically deleted to free space when a write to a volume fails because of lack of space in the aggregate.

- Quotas

  Specifies whether the quotas are enabled for the volume.

- **Efficiency**
  - Compression

    Specifies whether compression is enabled or disabled.

  - Deduplication

    Specifies whether deduplication is enabled or disabled.

  - Deduplication Mode

    Specifies whether the deduplication operation enabled on a volume is a manual, scheduled, or policy-based operation. If the mode is set to Scheduled, the operation schedule is displayed, and if the mode is set to a policy, the policy name is displayed.

  - Deduplication Type

    Specifies the type of deduplication operation running on the volume. If the volume is in a SnapVault relationship, the type displayed is SnapVault. For any other volume, the type is displayed as Regular.

  - Storage Efficiency Policy

    Specifies the name of the storage efficiency policy that has been assigned through Unified Manager to this volume. This policy can control the compression and deduplication settings.

- **Protection**
  - Snapshot Copies

    Specifies whether automatic Snapshot copies are enabled or disabled.

Protection tab

The Protection tab displays protection details about the selected volume, such as lag information, relationship type, and topology of the relationship.

- **Summary**

  Displays protection relationships (SnapMirror, SnapVault, or Storage VM DR) properties for a selected volume. For any other relationship type, only the Relationship Type property is displayed. If a primary volume is selected, only the Managed and Local Snapshot copy Policy are displayed. Properties displayed

for SnapMirror and SnapVault relationships include the following:

- Source Volume

  Displays the name of the selected volume's source if the selected volume is a destination.

- Lag Status

  Displays the update or transfer lag status for a protection relationship. The status can be Error, Warning, or Critical.

  The lag status is not applicable for synchronous relationships.

- Lag Duration

  Displays the time by which the data on the mirror lags behind the source.

- Last Successful Update

  Displays the date and time of the most recent successful protection update.

  The last successful update is not applicable for synchronous relationships.

- Storage Service Member

  Displays either Yes or No to indicate whether or not the volume belongs to and is managed by a storage service.

- Version Flexible Replication

  Displays either Yes, Yes with backup option, or None. Yes indicates that SnapMirror replication is possible even if source and destination volumes are running different versions of ONTAP software. Yes with backup option indicates the implementation of SnapMirror protection with the ability to retain multiple versions of backup copies on the destination. None indicates that Version Flexible Replication is not enabled.

- Relationship Capability

  Indicates the ONTAP capabilities available to the protection relationship.

- Protection Service

  Displays the name of the protection service if the relationship is managed by a protection partner application.

- Relationship Type

  Displays any relationship type, including Asynchronous Mirror, Asynchronous Vault, Asynchronous MirrorVault, StrictSync, and Sync.

- Relationship State

  Displays the state of the SnapMirror or SnapVault relationship. The state can be Uninitialized, SnapMirrored, or Broken-Off. If a source volume is selected, the relationship state is not applicable and is not displayed.

- Transfer Status

  Displays the transfer status for the protection relationship. The transfer status can be one of the following:

  - Aborting

    SnapMirror transfers are enabled; however, a transfer abort operation that might include removal of the checkpoint is in progress.

  - Checking

    The destination volume is undergoing a diagnostic check and no transfer is in progress.

  - Finalizing

    SnapMirror transfers are enabled. The volume is currently in the post-transfer phase for incremental SnapVault transfers.

  - Idle

    Transfers are enabled and no transfer is in progress.

  - In-Sync

    The data in the two volumes in the synchronous relationship are synchronized.

  - Out-of-Sync

    The data in the destination volume is not synchronized with the source volume.

  - Preparing

    SnapMirror transfers are enabled. The volume is currently in the pre-transfer phase for incremental SnapVault transfers.

  - Queued

    SnapMirror transfers are enabled. No transfers are in progress.

  - Quiesced

    SnapMirror transfers are disabled. No transfer is in progress.

  - Quiescing

    A SnapMirror transfer is in progress. Additional transfers are disabled.

  - Transferring

    SnapMirror transfers are enabled and a transfer is in progress.

  - Transitioning

    The asynchronous transfer of data from the source to the destination volume is complete, and the transition to synchronous operation has started.

        - Waiting

          A SnapMirror transfer has been initiated, but some associated tasks are waiting to be queued.

    - Max Transfer Rate

      Displays the maximum transfer rate for the relationship. The maximum transfer rate can be a numerical value in either kilobytes per second (Kbps), Megabytes per second (Mbps), Gigabytes per second (Gbps), or Terabytes per second (Tbps). If No Limit is displayed, the baseline transfer between relationships is unlimited.

    - SnapMirror Policy

      Displays the protection policy for the volume. DPDefault indicates the default Asynchronous Mirror protection policy, XDPDefault indicates the default Asynchronous Vault policy, and DPSyncDefault indicates the default Asynchronous MirrorVault policy. StrictSync indicates the default Synchronous Strict protection policy, and Sync indicates the default Synchronous policy. You can click the policy name to view details associated with that policy, including the following information:

        - Transfer priority
        - Ignore access time setting
        - Tries limit
        - Comments
        - SnapMirror labels
        - Retention settings
        - Actual Snapshot copies
        - Preserve Snapshot copies
        - Retention warning threshold
        - Snapshot copies with no retention settings
          In a cascading SnapVault relationship where the source is a data protection (DP) volume, only the rule “sm_created” applies.

    - Update Schedule

      Displays the SnapMirror schedule assigned to the relationship. Positioning your cursor over the information icon displays the schedule details.

    - Local Snapshot Policy

      Displays the Snapshot copy policy for the volume. The policy is Default, None, or any name given to a custom policy.

    - Protected By

      Displays the type of protection used for the selected volume. For example, if a volume is protected by Consistency Group and SnapMirror volume relationships, this field displays both SnapMirror and Consistency Group. This field also provides a link that redirects you to the Relationships page to view the unified relationship status. The link is only applicable to constituent relationships.

    - Consistency Group

      For volumes protected by SnapMirror active sync relationships, this column displays the Consistency

Group of the volume.

- **Views**

  Displays the protection topology of the selected volume. The topology includes graphical representations of all volumes that are related to the selected volume. The selected volume is indicated by a dark gray border, and lines between volumes in the topology indicate the protection relationship type. The direction of the relationships in the topology are displayed from left to right, with the source of each relationship on the left and the destination on the right.

  Double bold lines specify an Asynchronous Mirror relationship, a single bold line specifies an Asynchronous Vault relationship, double single lines specify an Asynchronous MirrorVault relationship, and a bold line and non-bold line specifies a Synchronous relationship. The table below indicates if the Synchronous relationship is StrictSync or Sync.

  Right-clicking a volume displays a menu from which you can choose either to protect the volume or restore data to it. Right-clicking a relationship displays a menu from which you can choose to either edit, abort, quiesce, break, remove, or resume a relationship.

  The menus will not display in the following instances:

  - If RBAC settings do not allow this action, for example, if you have only operator privileges
  - If the volume is in a synchronous protection relationship
  - When the volume ID is unknown, for example, when you have an intercluster relationship and the destination cluster has not yet been discovered
    Clicking another volume in the topology selects and displays information for that volume. A question mark ( ) in the upper-left corner of a volume indicates that either the volume is missing or that it has not yet been discovered. It might also indicate that the capacity information is missing. Positioning your cursor over the question mark displays additional information, including suggestions for remedial action.

  The topology displays information about volume capacity, lag, Snapshot copies, and last successful data transfer if it conforms to one of several common topology templates. If a topology does not conform to one of those templates, information about volume lag and last successful data transfer is displayed in a relationship table under the topology. In that case, the highlighted row in the table indicates the selected volume, and, in the topology view, bold lines with a blue dot indicate the relationship between the selected volume and its source volume.

Topology views include the following information:

- Capacity

  Displays the total amount of capacity used by the volume. Positioning your cursor over a volume in the topology displays the current warning and critical threshold settings for that volume in the Current Threshold Settings dialog box. You can also edit the threshold settings by clicking the **Edit Thresholds** link in the Current Threshold Settings dialog box. Clearing the **Capacity** check box hides all capacity information for all volumes in the topology.

- Lag

  Displays the lag duration and the lag status of the incoming protection relationships. Clearing the **Lag** check box hides all lag information for all volumes in the topology. When the **Lag** check box is dimmed, then the lag information for the selected volume is displayed in the relationship table below the topology, as well as the lag information for all related volumes.

- Snapshot

  Displays the number of Snapshot copies available for a volume. Clearing the **Snapshot** check box hides all Snapshot copy information for all volumes in the topology. Clicking a Snapshot copy icon () displays the Snapshot copy list for a volume. The Snapshot copy count displayed next to the icon is updated approximately every hour; however, the list of Snapshot copies is updated at the time that you click the icon. This might result in a difference between the Snapshot copy count displayed in the topology and the number of Snapshot copies listed when you click the icon.

- Last Successful Transfer

  Displays the amount, duration, time, and date of the last successful data transfer. When the **Last Successful Transfer** check box is dimmed, then the last successful transfer information for the selected volume is displayed in the relationship table below the topology, as well as the last successful transfer information for all related volumes.

  - **History**

    Displays in a graph the history of incoming SnapMirror and SnapVault protection relationships for the selected volume. There are three history graphs available: incoming relationship lag duration, incoming relationship transfer duration, and incoming relationship transfer size. History information is displayed only when you select a destination volume. If you select a primary volume, the graphs are empty, and the message No data found is displayed. If the volumes are protected by Consistency Group and SnapMirror synchronous relationships, information for the relationship transfer duration and relationship transfer size is not displayed.

  You can select a graph type from the drop-down list at the top of the History pane. You can also view details for a specific time period by selecting either 1 week, 1 month, or 1 year. History graphs can help you identify trends: for example, if large amounts of data are being transferred at the same time of the day or week, or if the lag warning or lag error threshold is consistently being breached, you can take the appropriate action. Additionally, you can click the **Export** button to create a report in CSV format for the chart that you are viewing.

Protection history graphs display the following information:

- **Relationship Lag Duration**

  Displays seconds, minutes, or hours on the vertical (y) axis, and displays days, months, or years on the horizontal (x) axis, depending on the selected duration period. The upper value on the y axis indicates the maximum lag duration reached in the duration period shown in the x axis. The horizontal orange line on the graph depicts the lag error threshold, and the horizontal yellow line depicts the lag warning threshold. Positioning your cursor over these lines displays the threshold setting. The horizontal blue line depicts the lag duration. You can view the details for specific points on the graph by positioning your cursor over an area of interest.

- **Relationship Transfer Duration**

  Displays seconds, minutes, or hours on the vertical (y) axis, and displays days, months, or years on the horizontal (x) axis, depending on the selected duration period. The upper value on the y axis indicates the maximum transfer duration reached in the duration period shown in the x axis. You can view the details of specific points on the graph by positioning your cursor over the area of interest.

  This chart is not available for volumes that are in synchronous protection relationships.

- **Relationship Transferred Size**

Displays bytes, kilobytes, megabytes, and so on, on the vertical (y) axis depending on the transfer size, and displays days, months, or years on the horizontal (x) axis depending on the selected time period. The upper value on the y axis indicates the maximum transfer size reached in the duration period shown in the x axis. You can view the details for specific points on the graph by positioning your cursor over an area of interest.



This chart is not available for volumes that are in synchronous protection relationships.

History area

The History area displays graphs that provide information about the capacity and space reservations of the selected volume. Additionally, you can click the **Export** button to create a report in CSV format for the chart that you are viewing.

Graphs might be empty and the message No data found displayed when the data or the state of the volume remains unchanged for a period of time.

You can select a graph type from the drop-down list at the top of the History pane. You can also view details for a specific time period by selecting either 1 week, 1 month, or 1 year. History graphs can help you identify trends—for example, if the volume usage is consistently breaching the Nearly Full threshold, you can take the appropriate action.

History graphs display the following information:

- **Volume Capacity Used**

  Displays the used capacity in the volume and the trend in how volume capacity is used based on the usage history, as line graphs in bytes, kilobytes, megabytes, and so on, on the vertical (y) axis. The time period is displayed on the horizontal (x) axis. You can select a time period of a week, a month, or a year. You can view the details for specific points on the graph by positioning your cursor over a particular area. You can hide or display a line graph by clicking the appropriate legend. For example, when you click the Volume Used Capacity legend, the Volume Used Capacity graph line is hidden.

- **Volume Capacity Used vs Total**

  Displays the trend in how volume capacity is used based on the usage history, as well as the used capacity, total capacity, and details of the space savings from deduplication and compression, as line graphs, in bytes, kilobytes, megabytes, and so on, on the vertical (y) axis. The time period is displayed on the horizontal (x) axis. You can select a time period of a week, a month, or a year. You can view the details for specific points on the graph by positioning your cursor over a particular area. You can hide or display a line graph by clicking the appropriate legend. For example, when you click the Trend Capacity Used legend, the Trend Capacity Used graph line is hidden.

- **Volume Capacity Used (%)**

  Displays the used capacity in the volume and the trend in how volume capacity is used based on the usage history, as line graphs, in percentage, on the vertical (y) axis. The time period is displayed on the horizontal (x) axis. You can select a time period of a week, a month, or a year. You can view the details for specific points on the graph by positioning your cursor over a particular area. You can hide or display a line graph by clicking the appropriate legend. For example, when you click the Volume Used Capacity legend, the Volume Used Capacity graph line is hidden.

- **Snapshot Capacity Used (%)**

  Displays the Snapshot reserve and Snapshot warning threshold as line graphs, and the capacity used by

the Snapshot copies as an area graph, in percentage, on the vertical (y) axis. The Snapshot overflow is represented with different colors. The time period is displayed on the horizontal (x) axis. You can select a time period of a week, a month, or a year. You can view the details for specific points on the graph by positioning your cursor over a particular area. You can hide or display a line graph by clicking the appropriate legend. For example, when you click the Snapshot Reserve legend, the Snapshot Reserve graph line is hidden.

Events list

The Events list displays details about new and acknowledged events:

- **Severity**

  Displays the severity of the event.

- **Event**

  Displays the event name.

- **Triggered Time**

  Displays the time that has elapsed since the event was generated. If the time elapsed exceeds a week, the timestamp when the event was generated is displayed.

Related Annotations pane

The Related Annotations pane enables you to view annotation details associated with the selected volume. The details include the annotation name and the annotation values that are applied to the volume. You can also remove manual annotations from the Related Annotations pane.

Related Devices pane

The Related Devices pane enables you to view and navigate to the SVMs, aggregates, qtrees, LUNs, and Snapshot copies that are related to the volume:

- **Storage Virtual Machine**

  Displays the capacity and the health status of the SVM that contains the selected volume.

- **Aggregate**

  Displays the capacity and the health status of the aggregate that contains the selected volume. For FlexGroup volumes, the number of aggregates that comprise the FlexGroup is listed.

- **Volumes in the Aggregate**

  Displays the number and capacity of all the volumes that belong to the parent aggregate of the selected volume. The health status of the volumes is also displayed, based on the highest severity level. For example, if an aggregate contains ten volumes, five of which display the Warning status and the remaining five display the Critical status, then the status displayed is Critical. This component does not appear for FlexGroup volumes.

- **Qtrees**

  Displays the number of qtrees that the selected volume contains and the capacity of qtrees with quota that the selected volume contains. The capacity of the qtrees with quota is displayed in relation to the volume

data capacity. The health status of the qtrees is also displayed, based on the highest severity level. For example, if a volume has ten qtrees, five with Warning status and the remaining five with Critical status, then the status displayed is Critical.

- **NFS Shares**

  Displays the number and status of the NFS shares associated with the volume.

- **SMB Shares**

  Displays the number and status of the SMB/CIFS shares.

- **LUNs**

  Displays the number and total size of all the LUNs in the selected volume. The health status of the LUNs is also displayed, based on the highest severity level.

- **User and Group Quotas**

  Displays the number and status of the user and user group quotas associated with the volume and its qtrees.

- **FlexClone Volumes**

  Displays the number and capacity of all the cloned volumes of the selected volume. The number and capacity are displayed only if the selected volume contains any cloned volumes.

- **Parent Volume**

  Displays the name and capacity of the parent volume of a selected FlexClone volume. The parent volume is displayed only if the selected volume is a FlexClone volume.

Related Groups pane

The Related Groups pane enables you to view the list of groups associated with the selected volume.

Related Alerts pane

The Related Alerts pane enables you to view the list of alerts that are created for the selected volume. You can also add an alert by clicking the Add Alert link or edit an existing alert by clicking the alert name.

View the Storage VM / Health details page in Active IQ Unified Manager

You can use the Storage VM / Health details page to view detailed information about the selected storage VM, such as its health, capacity, configuration, data policies, logical interfaces (LIFs), LUNs, qtrees, user, user group quotas, and protection details . You can also view information about the related objects and related alerts for the storage VM.

You can monitor only data storage VM.

Command buttons

The command buttons enable you to perform the following tasks for the selected storage VM:

- **Switch to Performance View**

  Enables you to navigate to the Storage VM / Performance details page.

- **Actions**
  - Add Alert

    Enables you to add an alert to the selected storage VM.

  - Annotate

    Enables you to annotate the selected storage VM.

- **View Storage VMs**

  Enables you to navigate to the Health: All Storage VMs view.

Health tab

The Health tab displays detailed information about data availability, data capacity, and protection issues of various objects such as volumes, aggregates, NAS LIFs, SAN LIFs, LUNs, protocols, services, NFS shares, and CIFS shares.

You can click the graph of an object to view the filtered list of objects. For example, you can click the volume capacity graph that displays warnings to view the list of volumes that have capacity issues with severity as warning.

- **Availability Issues**

  Displays, as a graph, the total number of objects, including objects that have availability issues and objects that do not have any availability-related issues. The colors in the graph represent the different severity levels of the issues. The information below the graph provides details about availability issues that can impact or have already impacted the availability of data in the storage VM. For example, information is displayed about the NAS LIFs and the SAN LIFs that are down and volumes that are offline.

  You can also view information about the related protocols and services that are currently running, and the number and status of NFS and CIFS shares.

- **Capacity Issues**

  Displays, as a graph, the total number of objects, including objects that have capacity issues and objects that do not have any capacity-related issues. The colors in the graph represent the different severity levels of the issues. The information below the graph provides details about capacity issues that can impact or have already impacted the capacity of data in the storage VM. For example, information is displayed about aggregates that are likely to breach the set threshold values.

- **Protection Issues**

  Provides a quick overview of storage VM protection-related health by displaying, as a field dialog box, the total number of relationships, including relationships that have protection issues and relationships that do not have any protection-related issues. You can also view the status of the storage VM DR relationship for the selected storage VM. The storage VM DR relationships events are displayed here and clicking on the events takes you to the event details page. When unprotected volumes exist, clicking on the link takes you to the Health: All Volumes view where you can view a filtered list of the unprotected volumes on the storage VM. The colors in the graph represent the different severity levels of the issues. Clicking a graph takes you

to the Relationship: All Relationships view, where you can view a filtered list of protection relationship details. The information below the graph provides details about protection issues that can impact or have already impacted the protection of data in the storage VM. For example, information is displayed about volumes that have a Snapshot copy reserve that is almost full or about SnapMirror relationship lag issues.

Capacity tab

The Capacity tab displays detailed information about the data capacity of the selected SVM.

The following information is displayed for an Storage VM with FlexVol volume or FlexGroup volume:

- **Capacity**

  The Capacity area displays details about the used and available capacity allocated from all volumes:

  - Total Capacity

    Displays the total capacity of the Storage VM.

  - Used

    Displays the space used by data in the volumes that belong to the Storage VM.

  - Guaranteed Available

    Displays the guaranteed available space for data that is available for volumes in the Storage VM.

  - Unguaranteed

    Displays the available space remaining for data that is allocated for thinly provisioned volumes in the Storage VM.

- **Volumes with Capacity Issues**

  The Volumes with Capacity Issues list displays, in tabular format, details about the volumes that have capacity issues:

  - Status

    Indicates that the volume has a capacity-related issue of an indicated severity.

    You can move the pointer over the status to view more information about the capacity-related event or events generated for the volume.

    If the status of the volume is determined by a single event, you can view information such as the event name, time and date when the event was triggered, the name of the administrator to whom the event is assigned, and the cause of the event. You can use the **View Details** button to view more information about the event.

    If the status of the volume is determined by multiple events of the same severity, the top three events are displayed with information such as the event name, time and date when the events were triggered, and the name of the administrator to whom the event is assigned. You can view more details about each of these events by clicking the event name. You can also click the **View All Events** link to view the list of generated events.

    A volume can have multiple events of the same severity or different severities. However,

only the highest severity is displayed. For example, if a volume has two events with severities of Error and Warning, only the Error severity is displayed.

- Volume

  Displays the name of the volume.

- Used Data Capacity

  Displays, as a graph, information about the volume capacity usage (in percentage).

- Days to Full

  Displays the estimated number of days remaining before the volume reaches full capacity.

- Thin Provisioned

  Displays whether space guarantee is set for the selected volume. Valid values are Yes and No.

- Aggregates

  For FlexVol volumes, displays the name of the aggregate that contains the volume. For FlexGroup volumes, displays the number of aggregates that are used in the FlexGroup.

Configuration tab

The Configuration tab displays configuration details about the selected storage VM, such as its cluster, root volume, the type of volumes it contains (FlexVol volumes), policies, and protection created on the storage VM:

- **Overview**
  - Cluster

    Displays the name of the cluster to which the storage VM belongs.

  - Allowed Volume Type

    Displays the type of volumes that can be created in the storage VM. The type can be FlexVol or FlexVol/FlexGroup.

  - Root Volume

    Displays the name of the root volume of the storage VM.

  - Allowed Protocols

    Displays the type of protocols that can be configured on the storage VM. Also, indicates if a protocol is up (), down (), or is not configured ().

- **Data Network Interfaces**
  - NAS

    Displays the number of NAS interfaces that are associated with the storage VM. Also, indicates if the interfaces are up () or down ().

- SAN

  Displays the number of SAN interfaces that are associated with the storage VM. Also, indicates if the interfaces are up ( ) or down ( ).

- FC-NVMe

  Displays the number of FC-NVMe interfaces that are associated with the Storage VM. Also, indicates if the interfaces are up ( ) or down ( ).

- **Management Network Interfaces**
  - Availability

    Displays the number of management interfaces that are associated with the Storage VM. Also, indicates if the management interfaces are up ( ) or down ( ).

- **Policies**
  - Snapshots

    Displays the name of the Snapshot policy that is created on the Storage VM.

  - Export Policies

    Displays either the name of the export policy if a single policy is created or displays the number of export policies if multiple policies are created.

- **Protection**
  - Storage VM DR

    Displays whether the selected storage VM is protected, destination, or unprotected and the name of the destination on which the storage VM is protected. If the selected storage VM is destination, then the details of source storage VM are displayed. In case of fan-out, this field displays the number of total destination storage VMs on which the storage VM is protected. The count link takes you to the storage VM relationship grid filtered on source storage VM.

  - Protected Volumes

    Displays the number of protected volumes on the selected storage VM out of the total volumes. If you are viewing a destination storage VM, then the number link is for the destination volumes of the selected storage VM.

  - Unprotected Volumes

    Displays the number of unprotected volumes on the selected storage VM.

- **Services**
  - Type

    Displays the type of service that is configured on the storage VM. The type can be Domain Name System (DNS) or Network Information Service (NIS).

  - State

    Displays the state of the service, which can be Up ( ), Down ( ), or Not Configured ( ).

- Domain Name

  Displays the fully qualified domain names (FQDNs) of the DNS server for the DNS services or NIS server for the NIS services. When the NIS server is enabled, the active FQDN of the NIS server is displayed. When the NIS server is disabled, the list of all the FQDNs are displayed.

- IP Address

  Displays the IP addresses of the DNS or NIS server. When the NIS server is enabled, the active IP address of the NIS server is displayed. When the NIS server is disabled, the list of all the IP addresses are displayed.

Network Interfaces tab

The Network Interfaces tab displays details about the data network interfaces (LIFs) that are created on the selected storage VM:

- **Network Interface**

  Displays the name of the interface that is created on the selected storage VM.

- **Operational Status**

  Displays the operational status of the interface, which can be Up (), Down (), or Unknown (). The operational status of an interface is determined by the status of its physical ports.

- **Administrative Status**

  Displays the administrative status of the interface, which can be Up (), Down (), or Unknown (). The administrative status of an interface is controlled by the storage administrator to make changes to the configuration or for maintenance purposes. The administrative status can be different from the operational status. However, if the administrative status of an interface is Down, the operational status is Down by default.

- **IP Address / WWPN**

  Displays the IP address for Ethernet interfaces and the World Wide Port Name (WWPN) for FC LIFs.

- **Protocols**

  Displays the list of data protocols that are specified for the interface, such as CIFS, NFS, iSCSI, FC/FCoE, FC-NVMe, and FlexCache.

- **Role**

  Displays the interface role. The roles can be Data or Management.

- **Home Port**

  Displays the physical port to which the interface was originally associated.

- **Current Port**

  Displays the physical port to which the interface is currently associated. If the interface is migrated, the current port might be different from the home port.

- **Port Set**

  Displays the port set to which the interface is mapped.

- **Failover Policy**

  Displays the failover policy that is configured for the interface. For NFS, CIFS, and FlexCache interfaces, the default failover policy is Next Available. Failover policy is not applicable for FC and iSCSI interfaces.

- **Routing Groups**

  Displays the name of the routing group. You can view more information about the routes and the destination gateway by clicking the routing group name.

  Routing groups are not supported for ONTAP 8.3 or later and therefore a blank column is displayed for these clusters.

- **Failover Group**

  Displays the name of the failover group.

Qtrees tab

The Qtrees tab displays details about qtrees and their quotas. You can click the **Edit Thresholds** button if you want to edit the health threshold settings for qtree capacity for one or more qtrees.

Use the **Export** button to create a comma-separated values (.csv) file containing the details of all the monitored qtrees. When exporting to a CSV file you can choose to create a qtrees report for the current storage VM, for all storage VMs in the current cluster, or for all storage VMs for all clusters in your data center. Some additional qtrees fields appear in the exported CSV file.

- **Status**

  Displays the current status of the qtree. The status can be Critical (), Error (), Warning (), or Normal ().

  You can move the pointer over the status icon to view more information about the event or events generated for the qtree.

  If the status of the qtree is determined by a single event, you can view information such as the event name, time and date when the event was triggered, the name of the administrator to whom the event is assigned, and the cause of the event. You can use **View Details** to view more information about the event.

  If the status of the qtree is determined by multiple events of the same severity, the top three events are displayed with information such as the event name, time and date when the events were triggered, and the name of the administrator to whom the event is assigned. You can view more details about each of these events by clicking the event name. You can also use **View All Events** to view the list of generated events.

  > A qtree can have multiple events of the same severity or different severities. However, only the highest severity is displayed. For example, if a qtree has two events with severities of Error and Warning, only the Error severity is displayed.

- **Qtree**

  Displays the name of the qtree.

- **Cluster**

  Displays the name of the cluster containing the qtree. Appears only in the exported CSV file.

- **Storage Virtual Machine**

  Displays the storage virtual machine (SVM) name containing the qtree. Appears only in the exported CSV file.

- **Volume**

  Displays the name of the volume that contains the qtree.

  You can move the pointer over the volume name to view more information about the volume.

- **Quota Set**

  Indicates whether a quota is enabled or disabled on the qtree.

- **Quota Type**

  Specifies if the quota is for a user, user group, or a qtree. Appears only in the exported CSV file.

- **User or Group**

  Displays the name of the user or user group. There will be multiple rows for each user and user group. When the quota type is qtree or if the quota is not set, then the column is empty. Appears only in the exported CSV file.

- **Disk Used %**

  Displays the percentage of disk space used. If a disk hard limit is set, this value is based on the disk hard limit. If the quota is set without a disk hard limit, the value is based on the volume data space. If the quota is not set or if quotas are off on the volume to which the qtree belongs, then “Not applicable” is displayed in the grid page and the field is blank in the CSV export data.

- **Disk Hard Limit**

  Displays the maximum amount of disk space allocated for the qtree. Unified Manager generates a critical event when this limit is reached and no further disk writes are allowed. The value is displayed as “Unlimited” for the following conditions: if the quota is set without a disk hard limit, if the quota is not set, or if quotas are off on the volume to which the qtree belongs.

- **Disk Soft Limit**

  Displays the amount of disk space allocated for the qtree before a warning event is generated. The value is displayed as “Unlimited” for the following conditions: if the quota is set without a disk soft limit, if the quota is not set, or if quotas are off on the volume to which the qtree belongs. By default, this column is hidden.

- **Disk Threshold**

  Displays the threshold value set on the disk space. The value is displayed as “Unlimited” for the following conditions: if the quota is set without a disk threshold limit, if the quota is not set, or if quotas are off on the volume to which the qtree belongs. By default, this column is hidden.

- **Files Used %**

Displays the percentage of files used in the qtree. If the file hard limit is set, this value is based on the file hard limit. No value is displayed if the quota is set without a file hard limit. If the quota is not set or if quotas are off on the volume to which the qtree belongs, then “Not applicable” is displayed in the grid page and the field is blank in the CSV export data.

- **File Hard Limit**

  Displays the hard limit for the number of files permitted on the qtrees. The value is displayed as “Unlimited” for the following conditions: if the quota is set without a file hard limit, if the quota is not set, or if quotas are off on the volume to which the qtree belongs.

- **File Soft Limit**

  Displays the soft limit for the number of files permitted on the qtrees. The value is displayed as “Unlimited” for the following conditions: if the quota is set without a file soft limit, if the quota is not set, or if quotas are off on the volume to which the qtree belongs. By default, this column is hidden.

User and Group Quotas tab

Displays details about the user and user group quotas for the selected storage VM. You can view information such as the status of the quota, name of the user or user group, soft and hard limits set on the disks and files, amount of disk space and number of files used, and the disk threshold value. You can also change the email address associated with a user or user group.

- **Edit Email Address command button**

  Opens the Edit Email Address dialog box, which displays the current email address of the selected user or user group. You can modify the email address. If the **Edit Email Address** field is blank, the default rule is used to generate an email address for the selected user or user group.

  If more than one user has the same quota, the names of the users are displayed as comma-separated values. Also, the default rule is not used to generate the email address; therefore, you must provide the required email address for notifications to be sent.

- **Configure Email Rules command button**

  Enables you to create or modify rules to generate an email address for the user or user group quotas that are configured on the storage VM. A notification is sent to the specified email address when there is a quota breach.

- **Status**

  Displays the current status of the quota. The status can be Critical ( ), Warning ( ), or Normal ( ).

  You can move the pointer over the status icon to view more information about the event or events generated for the quota.

  If the status of the quota is determined by a single event, you can view information such as the event name, time and date when the event was triggered, the name of the administrator to whom the event is assigned, and the cause of the event. You can use **View Details** to view more information about the event.

  If the status of the quota is determined by multiple events of the same severity, the top three events are displayed with information such as the event name, time and date when the events were triggered, and the name of the administrator to whom the event is assigned. You can view more details about each of these events by clicking the event name. You can also use **View All Events** to view the list of generated events.

A quota can have multiple events of the same severity or different severities. However, only the highest severity is displayed. For example, if a quota has two events with severities of Error and Warning, only the Error severity is displayed.

- **User or Group**

  Displays the name of the user or user group. If more than one user has the same quota, the names of the users are displayed as comma-separated values.

  The value is displayed as “Unknown” when ONTAP does not provide a valid user name because of SecD errors.

- **Type**

  Specifies if the quota is for a user or a user group.

- **Volume or Qtree**

  Displays the name of the volume or qtree on which the user or user group quota is specified.

  You can move the pointer over the name of the volume or qtree to view more information about the volume or qtree.

- **Disk Used %**

  Displays the percentage of disk space used. The value is displayed as “Not applicable” if the quota is set without a disk hard limit.

- **Disk Hard Limit**

  Displays the maximum amount of disk space allocated for the quota. Unified Manager generates a critical event when this limit is reached and no further disk writes are allowed. The value is displayed as “Unlimited” if the quota is set without a disk hard limit.

- **Disk Soft Limit**

  Displays the amount of disk space allocated for the quota before a warning event is generated. The value is displayed as “Unlimited” if the quota is set without a disk soft limit. By default, this column is hidden.

- **Disk Threshold**

  Displays the threshold value set on the disk space. The value is displayed as “Unlimited” if the quota is set without a disk threshold limit. By default, this column is hidden.

- **Files Used %**

  Displays the percentage of files used in the qtree. The value is displayed as “Not applicable” if the quota is set without a file hard limit.

- **File Hard Limit**

  Displays the hard limit for the number of files permitted on the quota. The value is displayed as “Unlimited” if the quota is set without a file hard limit.

- **File Soft Limit**

Displays the soft limit for the number of files permitted on the quota. The value is displayed as “Unlimited” if the quota is set without a file soft limit. By default, this column is hidden.

- **Email Address**

  Displays the email address of the user or user group to which notifications are sent when there is a breach in the quotas.

NFS Shares tab

The NFS Shares tab displays information about NFS shares such as its status, the path associated with the volume (FlexGroup volumes or FlexVol volumes), access levels of clients to the NFS shares, and the export policy defined for the volumes that are exported. NFS shares will not be displayed in the following conditions: if the volume is not mounted or if the protocols associated with the export policy for the volume do not contain NFS shares.

- **Status**

  Displays the current status of the NFS shares. The status can be Error ( ) or Normal ( ).

- **Junction Path**

  Displays the path to which the volume is mounted. If an explicit NFS exports policy is applied to a qtree, the column displays the path of the volume through which the qtree can be accessed.

- **Junction Path Active**

  Displays whether the path to access the mounted volume is active or inactive.

- **Volume or Qtree**

  Displays the name of the volume or qtree to which the NFS export policy is applied. If an NFS export policy is applied to a qtree in the volume, the column displays both the names of the volume and the qtree.

  You can click the link to view details about the object in the respective details page. If the object is a qtree, links are displayed for both the qtree and the volume.

- **Volume State**

  Displays the state of the volume that is being exported. The state can be Offline, Online, Restricted, or Mixed.

  - Offline

    Read or write access to the volume is not allowed.

  - Online

    Read and write access to the volume is allowed.

  - Restricted

    Limited operations, such as parity reconstruction, are allowed, but data access is not allowed.

  - Mixed

The constituents of a FlexGroup volume are not all in the same state.

- **Security Style**

  Displays the access permission for the volumes that are exported. The security style can be UNIX, Unified, NTFS, or Mixed.

  - UNIX (NFS clients)

    Files and directories in the volume have UNIX permissions.

  - Unified

    Files and directories in the volume have a unified security style.

  - NTFS (CIFS clients)

    Files and directories in the volume have Windows NTFS permissions.

  - Mixed

    Files and directories in the volume can have either UNIX permissions or Windows NTFS permissions.

- **UNIX Permission**

  Displays the UNIX permission bits in an octal string format, which is set for the volumes that are exported. It is similar to the UNIX style permission bits.

- **Export Policy**

  Displays the rules that define the access permission for volumes that are exported. You can click the link to view details about the rules associated with the export policy such as the authentication protocols and the access permission.

SMB Shares tab

Displays information about the SMB shares on the selected storage VM. You can view information such as the status of the SMB share, share name, path associated with the storage VM, the status of the junction path of the share, containing object, state of the containing volume, security data of the share, and export policies defined for the share. You can also determine whether an equivalent NFS path for the SMB share exists.

Shares in folders are not displayed in the SMB Shares tab.

- **View User Mapping command button**

  Launches the User Mapping dialog box.

  You can view the details of user mapping for the storage VM.

- **Show ACL command button**

  Launches the Access Control dialog box for the share.

  You can view user and permission details for the selected share.

- **Status**

  Displays the current status of the share. The status can be Normal () or Error ().

- **Share Name**

  Displays the name of the SMB share.

- **Path**

  Displays the junction path on which the share is created.

- **Junction Path Active**

  Displays whether the path to access the share is active or inactive.

- **Containing Object**

  Displays the name of the containing object to which the share belongs. The containing object can be a volume or a qtree.

  By clicking the link, you can view details about the containing object in the respective Details page. If the containing object is a qtree, links are displayed for both qtree and volume.

- **Volume State**

  Displays the state of the volume that is being exported. The state can be Offline, Online, Restricted, or Mixed.

  - Offline

    Read or write access to the volume is not allowed.

  - Online

    Read and write access to the volume is allowed.

  - Restricted

    Limited operations, such as parity reconstruction, are allowed, but data access is not allowed.

  - Mixed

    The constituents of a FlexGroup volume are not all in the same state.

- **Security**

  Displays the access permission for the volumes that are exported. The security style can be UNIX, Unified, NTFS, or Mixed.

  - UNIX (NFS clients)

    Files and directories in the volume have UNIX permissions.

  - Unified

Files and directories in the volume have a unified security style.

- NTFS (CIFS clients)

  Files and directories in the volume have Windows NTFS permissions.

- Mixed

  Files and directories in the volume can have either UNIX permissions or Windows NTFS permissions.

- **Export Policy**

  Displays the name of the export policy applicable to the share. If an export policy is not specified for the storage VM, the value is displayed as Not Enabled.

  You can click the link to view details about the rules associated with the export policy, such as access protocols and permissions. The link is disabled if the export policy is disabled for the selected storage VM.

- **NFS Equivalent**

  Specifies whether there is an NFS equivalent for the share.

SAN tab

Displays details about LUNs, initiator groups, and initiators for the selected storage VM. By default, the LUNs view is displayed. You can view details about the initiator groups in the Initiator Groups tab and details about initiators in the Initiators tab.

- **LUNs tab**

  Displays details about the LUNs that belong to the selected storage VM. You can view information such as the LUN name, LUN state (online or offline), the name of the file system (volume or qtree) that contains the LUN, the type of host operating system, the total data capacity and serial number of the LUN. The LUN Performance column provides a link to the LUN/Performance details page.

  You can also view information whether thin provisioning is enabled on the LUN and if the LUN is mapped to an initiator group. If it is mapped to an initiator, you can view the initiator groups and initiators that are mapped to the selected LUN.

- **Initiator Groups tab**

  Displays details about initiator groups. You can view details such as the name of the initiator group, the access state, the type of host operating system that is used by all the initiators in the group, and the supported protocol. When you click the link in the access state column, you can view the current access state of the initiator group.

  - **Normal**

    The initiator group is connected to multiple access paths.

  - **Single Path**

    The initiator group is connected to a single access path.

  - **No Paths**

There is no access path connected to the initiator group.

You can view whether initiator groups are mapped to all the interfaces or specific interfaces through a port set. When you click the count link in the Mapped interfaces column, either all interfaces are displayed or specific interfaces for a port set are displayed. Interfaces that are mapped through the target portal are not displayed. The total number of initiators and LUNs that are mapped to an initiator group is displayed.

You can also view the LUNs and initiators that are mapped to the selected initiator group.

- **Initiators tab**

  Displays the name and type of the initiator and the total number of initiator groups mapped to this initiator for the selected storage VM.

```
initiator groups that are mapped to the selected initiator group.
```

Related Annotations pane

The Related Annotations pane enables you to view the annotation details associated with the selected storage VM. Details include the annotation name and the annotation values that are applied to the storage VM. You can also remove manual annotations from the Related Annotations pane.

Related Devices pane

The Related Devices pane enables you to view the cluster, aggregates, and volumes that are related to the storage VM:

- **Cluster**

  Displays the health status of the cluster to which the storage VM belongs.

- **Aggregates**

  Displays the number of aggregates that belong to the selected storage VM. The health status of the aggregates is also displayed, based on the highest severity level. For example, if an storage VM contains ten aggregates, five of which display the Warning status and the remaining five display the Critical status, then the status displayed is Critical.

- **Assigned Aggregates**

  Displays the number of aggregates that are assigned to an storage VM. The health status of the aggregates is also displayed, based on the highest severity level.

- **Volumes**

  Displays the number and capacity of the volumes that belong to the selected storage VM. The health status of the volumes is also displayed, based on the highest severity level. When there are FlexGroup volumes in the storage VM, the count also includes FlexGroups; it does not include FlexGroup constituents.

Related Groups pane

The Related Groups pane enables you to view the list of groups associated with the selected storage VM.

Related Alerts pane

The Related Alerts pane enables you to view the list of alerts that are created for the selected storage VM. You can also add an alert by clicking the **Add Alert** link or edit an existing alert by clicking the alert name.

View the Cluster / Health details page in Active IQ Unified Manager

The Cluster / Health details page provides detailed information about a selected cluster, such as health, capacity, and configuration details. You can also view information about the network interfaces (LIFs), nodes, disks, related devices, and related alerts for the cluster.

The status next to the cluster name, for example (Good), represents the communication status; whether Unified Manager can communicate with the cluster. It does not represent the failover status or overall status of the cluster.

Command buttons

The command buttons enable you to perform the following tasks for the selected cluster:

- **Switch to Performance View**

  Enables you to navigate to the Cluster / Performance details page.

- **Actions**
  - Add Alert: Opens the Add Alert dialog box, which enables you to add an alert to the selected cluster.
  - Rediscover: Initiates a manual refresh of the cluster, which enables Unified Manager to discover recent changes to the cluster.

    If Unified Manager is paired with OnCommand Workflow Automation, the rediscovery operation also reacquires cached data from WFA, if any.

    After the rediscovery operation is initiated, a link to the associated job details is displayed to enable tracking of the job status.

  - Annotate: Enables you to annotate the selected cluster.
- **View Clusters**

  Enables you to navigate to the Health: All Clusters view.

Health tab

Displays detailed information about the data availability and data capacity issues of various cluster objects such as nodes, SVMs, and aggregates. Availability issues are related to the data-serving capability of the cluster objects. Capacity issues are related to the data-storing capability of the cluster objects.

You can click the graph of an object to view a filtered list of the objects. For example, you can click the SVM capacity graph that displays warnings to view a filtered list of SVMs. This list contains SVMs that have volumes or qtrees that have capacity issues with a severity level of Warning. You can also click the SVMs availability graph that displays warnings to view the list of SVMs that have availability issues with a severity level of Warning.

Availability Issues

Graphically displays the total number of objects, including objects that have availability issues and objects that do not have any availability-related issues. The colors in the graph represent the different severity levels of the issues. The information below the graph provides details about availability issues that can impact or have already impacted the availability of data in the cluster. For example, information is displayed about disk shelves that are down and aggregates that are offline.



The data displayed for the SFO bar graph is based on the HA state of the nodes. The data displayed for all other bar graphs is calculated based on the events generated.

Capacity Issues

Graphically displays the total number of objects, including objects that have capacity issues and objects that do not have any capacity-related issues. The colors in the graph represent the different severity levels of the issues. The information below the graph provides details about capacity issues that can impact or have already impacted the capacity of data in the cluster. For example, information is displayed about aggregates that are likely to breach the set threshold values.

Capacity tab

Displays detailed information about the capacity of the selected cluster.

Capacity

Displays the data capacity graph about the used capacity and available capacity from all allocated aggregates:

- Logical Space Used

  The real size of the data that is being stored on all aggregates on this cluster without applying the savings from using ONTAP storage efficiency technologies. This does not include Snapshot copies.

- Data Reduction

  Displays the ratio without Snapshot copies and with two significant digits, for example, 1.8 to 1. This ratio is based on the configured ONTAP storage efficiency settings.

- Used

  The physical capacity that is used by data on all aggregates. This does not include the capacity that is used for parity, right-sizing, and reservation.

- Available

  Displays the capacity available for data.

- Spares

  Displays the storable capacity available for storage in all the spare disks.

- Provisioned

  Displays the capacity that is provisioned for all the underlying volumes.

Details

Displays detailed information about the used and available capacity. The calculation excludes the root aggregate data.

- Total Capacity

  Displays the total capacity of the cluster. This does not include the capacity that is assigned for parity.

- Used

  Displays the capacity that is used by data. This does not include the capacity that is used for parity, right-sizing, and reservation.

- Available

  Displays the capacity available for data.

- Provisioned

  Displays the capacity that is provisioned for all the underlying volumes.

- Spares

  Displays the storable capacity available for storage in all the spare disks.

Cloud Tier

Displays the total cloud tier capacity used, and the capacity used for each connected cloud tier for FabricPool-enabled aggregates on the cluster. A FabricPool can be either licensed or unlicensed.

Physical Capacity Breakout by Disk Type

The Physical Capacity Breakout by Disk Type area displays detailed information about the disk capacity of the various types of disks in the cluster. By clicking the disk type, you can view more information about the disk type from the Disks tab.

- Total Usable Capacity

  Displays the available capacity and spare capacity of the data disks.

- HDD

  Graphically displays the used capacity and available capacity of all the HDD data disks in the cluster. The dotted line represents the spare capacity of the data disks in the HDD.

- Flash
  - SSD Data

    Graphically displays the used capacity and available capacity of the SSD data disks in the cluster.

  - SSD Cache

    Graphically displays the storable capacity of the SSD cache disks in the cluster.

- SSD Spare

  Graphically displays the spare capacity of the SSD, data, and cache disks in the cluster.

- Unassigned Disks

  Displays the number of unassigned disks in the cluster.

Aggregates with Capacity Issues list

Displays in tabular format details about the used capacity and available capacity of the aggregates that have capacity risk issues.

- Status

  Indicates that the aggregate has a capacity-related issue of a certain severity.

  You can move the pointer over the status to view more information about the event or events generated for the aggregate.

  If the status of the aggregate is determined by a single event, you can view information such as the event name, time and date when the event was triggered, the name of the administrator to whom the event is assigned, and the cause of the event. You can click the **View Details** button to view more information about the event.

  If the status of the aggregate is determined by multiple events of the same severity, the top three events are displayed with information such as the event name, time and date when the events are triggered, and the name of the administrator to whom the event is assigned. You can view more details about each of these events by clicking the event name. You can also click the **View All Events** link to view the list of generated events.

  An aggregate can have multiple capacity-related events of the same severity or different severities. However, only the highest severity is displayed. For example, if an aggregate has two events with severity levels of Error and Critical, only the Critical severity is displayed.

- Aggregate

  Displays the name of the aggregate.

- Used Data Capacity

  Graphically displays information about the aggregate capacity usage (in percentage).

- Days to Full

  Displays the estimated number of days remaining before the aggregate reaches full capacity.

Configuration tab

Displays details about the selected cluster, such as IP address, contact, and location:

Cluster Overview

- Management Interface

Displays the cluster-management LIF that Unified Manager uses to connect to the cluster. The operational status of the interface is also displayed.

- Host Name or IP Address

  Displays the FQDN, short name, or the IP address of the cluster-management LIF that Unified Manager uses to connect to the cluster.

- FQDN

  Displays the fully qualified domain name (FQDN) of the cluster.

- OS Version

  Displays the ONTAP version that the cluster is running. If the nodes in the cluster are running different versions of ONTAP, then the earliest ONTAP version is displayed.

- Contact

  Displays details about the administrator whom you should contact in case of issues with the cluster.

- Location

  Displays the location of the cluster.

- Personality

  Identifies if this is an All SAN Array configured cluster.

Remote Cluster Overview

Provides details about the remote cluster in a MetroCluster configuration. This information is displayed only for MetroCluster configurations.

- Cluster

  Displays the name of the remote cluster. You can click the cluster name to navigate to the details page of the cluster.

- Host name or IP Address

  Displays the FQDN, short name, or IP address of the remote cluster.

- Location

  Displays the location of the remote cluster.

MetroCluster Overview

Provides details about the local cluster in a MetroCluster over FC or MetroCluster over IP configurations. This information is displayed only for MetroCluster over FC or IP configurations.

- Type

  Displays whether the MetroCluster type is two-node or four-node. For MetroCluster over IP, only four-node is supported.

- Configuration

  Displays the MetroCluster configuration over FC and IP, which can have the following values:

For FC

- Stretch Configuration with SAS cables
- Stretch Configuration with FC-SAS bridge
- Fabric Configuration with FC switches

  For a four-node MetroCluster, only Fabric Configuration with FC switches is supported.

For IP

- IP configuration with Ethernet switches (L2 or L3, depending on how the cluster is configured)
  - Automated Unplanned Switch Over (AUSO)

    Displays whether automated unplanned switchover is enabled for the local cluster. By default, AUSO is enabled for all clusters in a two-node MetroCluster configuration in Unified Manager. You can use the command-line interface to change the AUSO setting. This is supported for only MetroCluster over FC.

  - Switch-Over mode

    Displays the switch-over mode for the MetroCluster over IP configuration. The available values are: `Active`, `Negotiated Switchover`, and `Automatic Unplanned Switchover`.

Nodes

- Availability

  Displays the number of nodes that are up ( ) or down ( ) in the cluster.

- OS Versions

  Displays the ONTAP versions that the nodes are running as well as the number of nodes running a particular version of ONTAP. For example, 9.6 (2), 9.3 (1) specifies that two nodes are running ONTAP 9.6, and one node is running ONTAP 9.3.

Storage Virtual Machines

- Availability

  Displays the number of SVMs that are up ( ) or down ( ) in the cluster.

Network Interfaces

- Availability

  Displays the number of non-data LIFs that are up ( ) or down ( ) in the cluster.

- Cluster-Management Interfaces

  Displays the number of cluster-management LIFs.

- Node-Management Interfaces

  Displays the number of node-management LIFs.

- Cluster Interfaces

  Displays the number of cluster LIFs.

- Intercluster Interfaces

  Displays the number of intercluster LIFs.

Protocols

- Data Protocols

  Displays the list of licensed data protocols that are enabled for the cluster. The data protocols include iSCSI, CIFS, NFS, NVMe, and FC/FCoE.

Protection

- Mediators

  Displays whether the cluster supports mediators and the connectivity status of the mediator. It indicates whether the mediator is configured, and if configured, it displays the status of the mediators.

  - Not Applicable

    Displays when the cluster doesn't support mediators.

  - Not Configured

    Displays when the cluster supports mediators, but the mediator is not configured.

  - IP Address

    Displays when the cluster supports mediators and the mediator is configured. The mediator status is indicated by color. The color green indicates the mediator status is reachable. The color red indicates the mediator status is unreachable.

Cloud Tiers

Lists the names of the cloud tiers to which this cluster is connected. It also lists the type (Amazon S3, Microsoft Azure Cloud, IBM Cloud Object Storage, Google Cloud Storage, Alibaba Cloud Object Storage, or StorageGRID), and the states of the cloud tiers (Available or Unavailable).

MetroCluster Connectivity tab

Displays the issues and connectivity status of the cluster components in the MetroCluster over FC configuration. A cluster is displayed in a red box when the disaster recovery partner of the cluster has issues.



The MetroCluster Connectivity tab is displayed only for clusters that are in a MetroCluster over FC configuration.

You can navigate to the details page of a remote cluster by clicking the name of the remote cluster. You can

also view the details of the components by clicking the count link of a component. For example, clicking the count link of the node in the cluster displays the node tab in the details page of the cluster. Clicking the count link of the disks in the remote cluster displays the disk tab in the details page of the remote cluster.

> When managing an eight-node MetroCluster configuration, clicking the count link of the Disk Shelves component displays only the local shelves of the default HA pair. Also, there is no way to display the local shelves on the other HA pair.

You can move the pointer over the components to view the details and the connectivity status of the clusters in case of any issue and to view more information about the event or events generated for the issue.

If the status of the connectivity issue between components is determined by a single event, you can view information such as the event name, time and date when the event was triggered, the name of the administrator to whom the event is assigned, and the cause of the event. The View Details button provides more information about the event.

If status of the connectivity issue between components is determined by multiple events of the same severity, the top three events are displayed with information such as the event name, time and date when the events are triggered, and the name of the administrator to whom the event is assigned. You can view more details about each of these events by clicking the event name. You can also click the **View All Events** link to view the list of generated events.

MetroCluster Replication tab

Displays the status of the data that is being replicated in a MetroCluster over FC configuration. You can use the MetroCluster Replication tab to ensure data protection by synchronously mirroring the data with the already peered clusters. A cluster is displayed in a red box when the disaster recovery partner of the cluster has issues.

> The MetroCluster Replication tab is displayed only for clusters that are in a MetroCluster over FC configuration.

In a MetroCluster environment, you can use this tab to verify the logical connections and peering of the local cluster with the remote cluster. You can view the objective representation of the cluster components with their logical connections. This helps to identify the issues that might occur during mirroring of metadata and data.

In the MetroCluster Replication tab, local cluster provides the detailed graphical representation of the selected cluster and MetroCluster partner refers to the remote cluster.

Network Interfaces tab

Displays details about all the non-data LIFs that are created on the selected cluster.

Network Interface

Displays the name of the LIF that is created on the selected cluster.

Operational Status

Displays the operational status of the interface, which can be Up (), Down (), or Unknown (). The operational status of a network interface is determined by the status of its physical ports.

Administrative Status

Displays the administrative status of the interface, which can be Up (), Down (), or Unknown (). You can control the administrative status of an interface when you make changes to the configuration or during maintenance. The administrative status can be different from the operational status. However, if the administrative status of a LIF is Down, the operational status is Down by default.

IP Address

Displays the IP address of the interface.

Role

Displays the role of the interface. Possible roles are Cluster-Management LIFs, Node-Management LIFs, Cluster LIFs, and Intercluster LIFs.

Home Port

Displays the physical port to which the interface was originally associated.

Current Port

Displays the physical port to which the interface is currently associated. After LIF migration, the current port might be different from the home port.

Failover Policy

Displays the failover policy that is configured for the interface.

Routing Groups

Displays the name of the routing group. You can view more information about the routes and the destination gateway by clicking the routing group name.

Routing groups are not supported for ONTAP 8.3 or later and therefore a blank column is displayed for these clusters.

Failover Group

Displays the name of the failover group.

Nodes tab

Displays information about nodes in the selected cluster. You can view detailed information about the HA pairs, disk shelves, and ports:

HA Details

Provides a pictorial representation of the HA state and the health status of the nodes in the HA pair. The health status of the node is indicated by the following colors:

- **Green**

  The node is in a working condition.

- **Yellow**

  The node has taken over the partner node or the node is facing some environmental issues.

- **Red**

  The node is down.

You can view information about the availability of the HA pair and take required action to prevent any risks. For example, in the case of a possible takeover operation, the following message is displayed: Storage failover possible.

You can view a list of the events related to the HA pair and its environment, such as fans, power supplies, NVRAM battery, flash cards, service processor, and connectivity of disk shelves. You can also view the time when the events were triggered.

You can view other node-related information, such as the model number.

If there are single-node clusters, you can also view details about the nodes.

Disk Shelves

Displays information about the disk shelves in the HA pair.

You can also view events generated for the disk shelves and the environmental components, and the time when the events were triggered.

- **Shelf ID**

  Displays the ID of the shelf where the disk is located.

- **Component Status**

  Displays environmental details of the disk shelves, such as power supplies, fans, temperature sensors, current sensors, disk connectivity, and voltage sensors. The environmental details are displayed as icons in the following colors:

  - **Green**

    The environmental components are in working properly.

  - **Grey**

    No data is available for the environmental components.

  - **Red**

    Some of the environmental components are down.

- **State**

  Displays the state of the disk shelf. The possible states are Offline, Online, No status, Initialization required, Missing, and Unknown.

- **Model**

Displays the model number of the disk shelf.

- **Local Disk Shelf**

  Indicates whether the disk shelf is located on the local cluster or the remote cluster. This column is displayed only for clusters in a MetroCluster configuration.

- **Unique ID**

  Displays the unique identifier of the disk shelf.

- **Firmware Version**

  Displays the firmware version of the disk shelf.

Ports

Displays information about the associated FC, FCoE, and Ethernet ports. You can view details about the ports and the associated LIFs by clicking the port icons.

You can also view the events generated for the ports.

You can view the following port details:

- Port ID

  Displays the name of the port. For example, the port names can be e0M, e0a, and e0b.

- Role

  Displays the role of the port. The possible roles are Cluster, Data, Intercluster, Node-Management, and Undefined.

- Type

  Displays the physical layer protocol used for the port. The possible types are Ethernet, Fibre Channel, and FCoE.

- WWPN

  Displays the World Wide Port Name (WWPN) of the port.

- Firmware Rev

  Displays the firmware revision of the FC/FCoE port.

- Status

  Displays the current state of the port. The possible states are Up, Down, Link Not Connected, or Unknown ().

  You can view the port-related events from the Events list. You can also view the associated LIF details, such as LIF name, operational status, IP address or WWPN, protocols, name of the SVM associated with the LIF, current port, failover policy and failover group.

Disks tab

Displays details about the disks in the selected cluster. You can view disk-related information such as the number of used disks, spare disks, broken disks, and unassigned disks. You can also view other details such as the disk name, disk type, and the owner node of the disk.

Disk Pool Summary

Displays the number of disks, which are categorized by effective types (FCAL, SAS, SATA, MSATA, SSD, NVMe SSD, SSD CAP, Array LUN, and VMDISK), and the state of the disks. You can also view other details, such as the number of aggregates, shared disks, spare disks, broken disks, unassigned disks, and unsupported disks. If you click the effective disk type count link, disks of the selected state and effective type are displayed. For example, if you click the count link for the disk state Broken and effective type SAS, all disks with the disk state Broken and effective type SAS are displayed.

Disk

Displays the name of the disk.

RAID Groups

Displays the name of the RAID group.

Owner Node

Displays the name of the node to which the disk belongs. If the disk is unassigned, no value is displayed in this column.

State

Displays the state of the disk: Aggregate, Shared, Spare, Broken, Unassigned, Unsupported or Unknown. By default, this column is sorted to display the states in the following order: Broken, Unassigned, Unsupported, Spare, Aggregate, and Shared.

Local Disk

Displays either Yes or No to indicate whether the disk is located on the local cluster or the remote cluster. This column is displayed only for clusters in a MetroCluster configuration.

Position

Displays the position of the disk based on its container type: for example, Copy, Data, or Parity. By default, this column is hidden.

Impacted Aggregates

Displays the number of aggregates that are impacted due to the failed disk. You can move the pointer over the count link to view the impacted aggregates and then click the aggregate name to view details of the aggregate. You can also click the aggregate count to view the list of impacted aggregates in the Health: All Aggregates view.

No value is displayed in this column for the following cases:

- For broken disks when a cluster containing such disks is added to Unified Manager

- When there are no failed disks

Storage Pool

Displays the name of the storage pool to which the SSD belongs. You can move the pointer over the storage pool name to view details of the storage pool.

Storable Capacity

Displays the disk capacity that is available for use.

Raw Capacity

Displays the capacity of the raw, unformatted disk before right-sizing and RAID configuration. By default, this column is hidden.

Type

Displays the types of disks: for example, ATA, SATA, FCAL, or VMDISK.

Effective Type

Displays the disk type assigned by ONTAP.

Certain ONTAP disk types are considered equivalent for the purposes of creating and adding to aggregates, and spare management. ONTAP assigns an effective disk type for each disk type.

Spare Blocks Consumed %

Displays in percentage the spare blocks that are consumed in the SSD disk. This column is blank for disks other than SSD disks.

Rated Life Used %

Displays in percentage an estimate of the SSD life used, based on the actual SSD usage and the manufacturer’s prediction of SSD life. A value greater than 99 indicates that the estimated endurance has been consumed, but may not indicate SSD failure. If the value is unknown, then the disk is omitted.

Firmware

Displays the firmware version of the disk.

RPM

Displays the revolutions per minute (RPM) of the disk. By default, this column is hidden.

Model

Displays the model number of the disk. By default, this column is hidden.

Vendor

Displays the name of the disk vendor. By default, this column is hidden.

Shelf ID

Displays the ID of the shelf where the disk is located.

Bay

Displays the ID of the bay where the disk is located.

Related Annotations pane

Enables you to view the annotation details associated with the selected cluster. The details include the annotation name and the annotation values that are applied to the cluster. You can also remove manual annotations from the Related Annotations pane.

Related Devices pane

Enables you to view device details that are associated with the selected cluster.

The details include properties of the device that is connected to the cluster such as the device type, size, count, and health status. You can click on the count link for further analysis on that particular device.

You can use MetroCluster Partner pane to obtain count and also details on the remote MetroCluster partner along with its associated cluster components such as nodes, aggregates, and SVMs. The MetroCluster Partner pane is displayed only for clusters in a MetroCluster configuration.

The Related Devices pane enables you to view and navigate to the nodes, SVMs, and aggregates that are related to the cluster:

MetroCluster Partner

Displays the health status of the MetroCluster partner. Using the count link, you can navigate further and obtain information about the health and capacity of the cluster components.

Nodes

Displays the number, capacity, and health status of the nodes that belong to the selected cluster. Capacity indicates the total usable capacity over available capacity.

Storage Virtual Machines

Displays the number of SVMs that belong to the selected cluster.

Aggregates

Displays the number, capacity, and the health status of the aggregates that belong to the selected cluster.

Related Groups pane

Enables you to view the list of groups that includes the selected cluster.

Related Alerts pane

The Related Alerts pane enables you to view the list of alerts for the selected cluster. You can also add an alert by clicking the Add Alert link or edit an existing alert by clicking the alert name.

Related information

Volumes page
Viewing the cluster list and details

View the Aggregate / Health details page in Active IQ Unified Manager

You can use the Aggregate / Health details page to view detailed information about the selected aggregate, such as the capacity, disk information, configuration details, and events generated. You can also view information about the related objects and related alerts for that aggregate.

Command buttons

When monitoring a FabricPool-enabled aggregate, the committed and overcommitted values on this page are relevant only to the local, or performance tier, capacity. The amount of space available in the cloud tier is not reflected in the overcommitted values. Similarly, the aggregate threshold values are relevant only to the local performance tier.

The command buttons enable you to perform the following tasks for the selected aggregate:

- **Switch to Performance View**

  Enables you to navigate to the Aggregate / Performance details page.

- **Actions**
  - Add Alert

    Enables you to add an alert to the selected aggregate.

  - Edit Thresholds

    Enables you to modify the threshold settings for the selected aggregate.

- **View Aggregates**

  Enables you to navigate to the Health: All Aggregates view.

Capacity tab

The Capacity tab displays detailed information about the selected aggregate, such as its capacity, thresholds, and daily growth rate.

By default, capacity events are not generated for root aggregates. Also, the threshold values used by Unified Manager are not applicable to node root aggregates. Only a technical support representative can modify the settings for these events to be generated. When the settings are modified by a technical support representative, the threshold values are applied to the node root aggregate.

- **Capacity**

  Displays the data capacity graph and the Snapshot copies graph, which display capacity details about the aggregate:

- Logical Space Used

  The real size of the data that is being stored on the aggregate without applying the savings from using ONTAP storage efficiency technologies.

- Used

  The physical capacity used by data in the aggregate.

- Overcommitted

  When space in the aggregate is overcommitted, the chart displays a flag with the overcommitted amount.

- Warning

  Displays a dotted line at the location where the warning threshold is set; meaning space in the aggregate is nearly full. If this threshold is breached, the Space Nearly Full event is generated.

- Error

  Displays a solid line at the location where the error threshold is set; meaning space in the aggregate is full. If this threshold is breached, the Space Full event is generated.

- Snapshot Copies graph

  This graph is displayed only when the used Snapshot capacity or the Snapshot reserve is not zero.

Both of the graphs display the capacity by which the Snapshot capacity exceeds the Snapshot reserve if the used Snapshot capacity exceeds the Snapshot reserve.

- **Cloud Tier**

  Displays the space used by data in the cloud tier for FabricPool-enabled aggregates. A FabricPool can be either licensed or unlicensed.

  When the cloud tier is mirrored to another cloud provider (the “mirror tier”) then both cloud tiers are displayed here.

- **Details**

  Displays detailed information about capacity.

  - Total Capacity

    Displays the total capacity in the aggregate.

  - Data Capacity

    Displays the amount of space used by the aggregate (used capacity) and the amount of available space in the aggregate (free capacity).

  - Snapshot Reserve

    Displays the used and free Snapshot capacity of the aggregate.

- Overcommitted Capacity

  Displays the aggregate overcommitment. Aggregate overcommitment enables you to provide more storage than is actually available from a given aggregate, as long as not all of that storage is currently being used. When thin provisioning is in use, the total size of volumes in the aggregate can exceed the total capacity of the aggregate.

  > If you have overcommitted your aggregate, you must monitor its available space carefully and add storage as required to avoid write errors due to insufficient space.

- Cloud Tier

  Displays the space used by data in the cloud tier for FabricPool-enabled aggregates. A FabricPool can be either licensed or unlicensed. When the cloud tier is mirrored to another cloud provider (the mirror tier) then both cloud tiers are displayed here

- Total Cache Space

  Displays the total space of the solid-state drives (SSDs) or allocation units that are added to a Flash Pool aggregate. If you have enabled Flash Pool for an aggregate but have not added any SSDs, then the cache space is displayed as 0 KB.

  > This field is hidden if Flash Pool is disabled for an aggregate.

- Aggregate Thresholds

  Displays the following aggregate capacity thresholds:

  - Nearly Full Threshold

    Specifies the percentage at which an aggregate is nearly full.

  - Full Threshold

    Specifies the percentage at which an aggregate is full.

  - Nearly Overcommitted Threshold

    Specifies the percentage at which an aggregate is nearly overcommitted.

  - Overcommitted Threshold

    Specifies the percentage at which an aggregate is overcommitted.

- Other Details: Daily Growth Rate

  Displays the disk space used in the aggregate if the rate of change between the last two samples continues for 24 hours.

  For example, if an aggregate uses 10 GB of disk space at 2 pm and 12 GB at 6 pm, the daily growth rate (GB) for this aggregate is 2 GB.

- Volume Move

  Displays the number of volume move operations that are currently in progress:

        - Volumes Out

          Displays the number and capacity of the volumes that are being moved out of the aggregate.

          You can click the link to view more details, such as the volume name, aggregate to which the volume is moved, status of the volume move operation, and the estimated end time.

        - Volumes In

          Displays the number and remaining capacity of the volumes that are being moved into the aggregate.

          You can click the link to view more details, such as the volume name, aggregate from which the volume is moved, status of the volume move operation, and the estimated end time.

        - Estimated used capacity after volume move

          Displays the estimated amount of used space (as a percentage, and in KB, MB, GB, and so on) in the aggregate after the volume move operations are complete.

- **Capacity Overview - Volumes**

  Displays graphs that provide information about the capacity of the volumes contained in the aggregate. The amount of space used by the volume (used capacity) and the amount of available space (free capacity) in the volume is displayed. When the Thin-Provisioned Volume Space At Risk event is generated for thinly provisioned volumes, the amount of space used by the volume (used capacity) and the amount of space that is available in the volume but cannot be used (unusable capacity) because of aggregate capacity issues is displayed.

  You can select the graph you want to view from the drop-down lists. You can sort the data displayed in the graph to display details such as the used size, provisioned size, available capacity, fastest daily growth rate, and slowest growth rate. You can filter the data based on the storage virtual machines (SVMs) that contain the volumes in the aggregate. You can also view details for thinly provisioned volumes. You can view the details of specific points on the graph by positioning your cursor over the area of interest. By default, the graph displays the top 30 filtered volumes in the aggregate.

Disk Information tab

Displays detailed information about the disks in the selected aggregate, including the RAID type and size, and the type of disks used in the aggregate. The tab also graphically displays the RAID groups, and the types of disks used (such as SAS, ATA, FCAL, SSD, or VMDISK). You can view more information, such as the disk's bay, shelf, and rotational speed, by positioning your cursor over the parity disks and data disks.

- **Data**

  Graphically displays details about dedicated data disks, shared data disks, or both. When the data disks contain shared disks, graphical details of the shared disks are displayed. When the data disks contain dedicated disks and shared disks, graphical details of both the dedicated data disks and the shared data disks are displayed.

    - **RAID Details**

      RAID details are displayed only for dedicated disks.

        - Type

Displays the RAID type (RAID0, RAID4, RAID-DP, or RAID-TEC).

- Group Size

  Displays the maximum number of disks allowed in the RAID group.

- Groups

  Displays the number of RAID groups in the aggregate.

- **Disks Used**
  - Effective Type

    Displays the types of data disks (for example, ATA, SATA, FCAL, SSD, or VMDISK) in the aggregate.

  - Data Disks

    Displays the number and capacity of the data disks that are assigned to an aggregate. Data disk details are not displayed when the aggregate contains only shared disks.

  - Parity Disks

    Displays the number and capacity of the parity disks that are assigned to an aggregate. Parity disk details are not displayed when the aggregate contains only shared disks.

  - Shared Disks

    Displays the number and capacity of the shared data disks that are assigned to an aggregate. Shared disk details are displayed only when the aggregate contains shared disks.

- **Spare Disks**

  Displays the disk effective type, number, and capacity of the spare data disks that are available for the node in the selected aggregate.

  When an aggregate is failed over to the partner node, Unified Manager does not display all of the spare disks that are compatible with the aggregate.

- **SSD Cache**

  Provides details about dedicated cache SSD disks and shared cache SSD disks.

  The following details for the dedicated cache SSD disks are displayed:

  - **RAID Details**
    - Type

      Displays the RAID type (RAID0, RAID4, RAID-DP or RAID-TEC).

    - Group Size

      Displays the maximum number of disks allowed in the RAID group.

- Groups

  Displays the number of RAID groups in the aggregate.

- **Disks Used**
  - Effective Type

    Indicates that the disks used for cache in the aggregate are of type SSD.

  - Data Disks

    Displays the number and capacity of the data disks that are assigned to an aggregate for cache.

  - Parity Disks

    Displays the number and capacity of the parity disks that are assigned to an aggregate for cache.

- **Spare Disks**

  Displays the disk effective type, number, and capacity of the spare disks that are available for the node in the selected aggregate for cache.

  When an aggregate is failed over to the partner node, Unified Manager does not display all of the spare disks that are compatible with the aggregate.

Provides the following details for the shared cache:

- **Storage Pool**

  Displays the name of the storage pool. You can move the pointer over the storage pool name to view the following details:

  - Status

    Displays the status of the storage pool, which can be healthy or unhealthy.

  - Total Allocations

    Displays the total allocation units and the size in the storage pool.

  - Allocation Unit Size

    Displays the minimum amount of space in the storage pool that can be allocated to an aggregate.

  - Disks

    Displays the number of disks used to create the storage pool. If the disk count in the storage pool column and the number of disks displayed in the Disk Information tab for that storage pool do not match, then it indicates that one or more disks are broken and the storage pool is unhealthy.

  - Used Allocation

    Displays the number and size of the allocation units used by the aggregates. You can click the aggregate name to view the aggregate details.

        - Available Allocation

          Displays the number and size of the allocation units available for the nodes. You can click the node name to view the aggregate details.

    - **Allocated Cache**

      Displays the size of the allocation units used by the aggregate.

    - **Allocation Units**

      Displays the number of allocation units used by the aggregate.

    - **Disks**

      Displays the number of disks contained in the storage pool.

    - **Details**
        - Storage Pool

          Displays the number of storage pools.

        - Total Size

          Displays the total size of the storage pools.

- **Cloud Tier**

  Displays the name of the cloud tier, if you have configured a FabricPool-enabled aggregate, and shows the total space used. When the cloud tier is mirrored to another cloud provider (the mirror tier) then the details for both cloud tiers are displayed here

Configuration tab

The Configuration tab displays details about the selected aggregate, such as its cluster node, block type, RAID type, RAID size, and RAID group count:

- **Overview**
    - Node

      Displays the name of the node that contains the selected aggregate.

    - Block Type

      Displays the block format of the aggregate: either 32-bit or 64-bit.

    - RAID Type

      Displays the RAID type (RAID0, RAID4, RAID-DP, RAID-TEC or Mixed RAID).

    - RAID Size

      Displays the size of the RAID group.

    - RAID Groups

Displays the number of RAID groups in the aggregate.

- SnapLock Type

  Displays the SnapLock Type of the aggregate.

- **Cloud Tier**

  If this is a FabricPool-enabled aggregate, the details for the cloud tier are displayed. Some fields are different depending on the storage provider. When the cloud tier is mirrored to another cloud provider (the “mirror tier”) then both cloud tiers are displayed here.

  - Provider

    Displays the name of the storage provider, for example, StorageGRID, Amazon S3, IBM Cloud Object Storage, Microsoft Azure Cloud, Google Cloud Storage, or Alibaba Cloud Object Storage.

  - Name

    Displays the name of the cloud tier when it was created by ONTAP.

  - Server

    Displays the FQDN of the cloud tier.

  - Port

    The port being used to communicate with the cloud provider.

  - Access Key or Account

    Displays the access key or account for the cloud tier.

  - Container Name

    Displays the bucket or container name of the cloud tier.

  - SSL

    Displays whether SSL encryption is enabled for the cloud tier.

History area

The History area displays graphs that provide information about the capacity of the selected aggregate. Additionally, you can click the **Export** button to create a report in CSV format for the chart that you are viewing.

You can select a graph type from the drop-down list at the top of the History pane. You can also view details for a specific time period by selecting either 1 week, 1 month, or 1 year. History graphs can help you identify trends: for example, if the aggregate usage is consistently breaching the Nearly Full threshold, you can take the appropriate action.

History graphs display the following information:

- **Aggregate Capacity Used (%)**

  Displays the used capacity in the aggregate and the trend in how aggregate capacity is used based on the

usage history as line graphs, in percentage, on the vertical (y) axis. The time period is displayed on the horizontal (x) axis. You can select a time period of a week, a month, or a year. You can view the details for specific points on the graph by positioning your cursor over a particular area. You can hide or display a line graph by clicking the appropriate legend. For example, when you click the Capacity Used legend, the Capacity Used graph line is hidden.

- **Aggregate Capacity Used vs Total Capacity**

  Displays the trend in how aggregate capacity is used based on the usage history, as well as the used capacity and the total capacity, as line graphs, in bytes, kilobytes, megabytes, and so on, on the vertical (y) axis. The time period is displayed on the horizontal (x) axis. You can select a time period of a week, a month, or a year. You can view the details for specific points on the graph by positioning your cursor over a particular area. You can hide or display a line graph by clicking the appropriate legend. For example, when you click the Trend Capacity Used legend, the Trend Capacity Used graph line is hidden.

- **Aggregate Capacity Used (%) vs Committed (%)**

  Displays the trend in how aggregate capacity is used based on the usage history, as well as the committed space as line graphs, as a percentage, on the vertical (y) axis. The time period is displayed on the horizontal (x) axis. You can select a time period of a week, a month, or a year. You can view the details for specific points on the graph by positioning your cursor over a particular area. You can hide or display a line graph by clicking the appropriate legend. For example, when you click the Space Committed legend, the Space Committed graph line is hidden.

Events list

The Events list displays details about new and acknowledged events:

- **Severity**

  Displays the severity of the event.

- **Event**

  Displays the event name.

- **Triggered Time**

  Displays the time that has elapsed since the event was generated. If the time elapsed exceeds a week, the timestamp for when the event was generated is displayed.

Related Devices pane

The Related Devices pane enables you to view the cluster node, volumes, and disks that are related to the aggregate:

- **Node**

  Displays the capacity and the health status of the node that contains the aggregate. Capacity indicates the total usable capacity over available capacity.

- **Aggregates in the Node**

  Displays the number and capacity of all the aggregates in the cluster node that contains the selected aggregate. The health status of the aggregates is also displayed, based on the highest severity level. For example, if a cluster node contains ten aggregates, five of which display the Warning status and the

remaining five of which display the Critical status, then the status displayed is Critical.

- **Volumes**

  Displays the number and capacity of FlexVol volumes and FlexGroup volumes in the aggregate; the number does not include FlexGroup constituents. The health status of the volumes is also displayed, based on the highest severity level.

- **Resource Pool**

  Displays the resource pools related to the aggregate.

- **Disks**

  Displays the number of disks in the selected aggregate.

Related Alerts pane

The Related Alerts pane enables you to view the list of alerts that are created for the selected aggregate. You can also add an alert by clicking the Add Alert link or edit an existing alert by clicking the alert name.

Related information

Viewing storage pool details

Protect and restore data

Create and troubleshoot protection relationships

Types of SnapMirror protection in Active IQ Unified Manager

Depending on the deployment of your data storage topology, Unified Manager enables you to configure multiple types of SnapMirror protection relationships. All variations of SnapMirror protection offer failover disaster recovery protection, but offer differing capabilities in performance, version flexibility, and multiple backup copy protection.

Traditional SnapMirror Asynchronous protection relationships

Traditional SnapMirror Asynchronous protection provides block replication mirror protection between source and destination volumes.

In traditional SnapMirror relationships, mirror operations execute faster than they would in alternative SnapMirror relationships because the mirror operation is based on block replication. However, traditional SnapMirror protection requires that the destination volume run under the same or later minor version of ONTAP software as the source volume within the same major release (for example, version 8.x to 8.x, or 9.x to 9.x). Replication from a 9.1 source to a 9.0 destination is not supported because the destination is running an earlier major version.

SnapMirror Asynchronous protection with version-flexible replication

SnapMirror Asynchronous protection with version-flexible replication provides logical replication mirror protection between source and destination volumes, even if those volumes are running under different versions of ONTAP 8.3 or later software (for example, version 8.3 to 8.3.1, or 8.3 to 9.1, or 9.2.2 to 9.2).

In SnapMirror relationships with version-flexible replication, mirror operations do not execute as quickly as they would in traditional SnapMirror relationships.

Because of slower execution, SnapMirror with version-flexible replication protection is not suitable to implement in either of the following circumstances:

- The source object contains more than 10 million files to protect.
- The recovery point objective for the protected data is two hours or less. (That is, the destination must always contain mirrored, recoverable data that is no more than two hours older than data at the source.)

In either of the listed circumstances, the faster block-replication based execution of default SnapMirror protection is required.

SnapMirror Asynchronous protection with version-flexible replication and backup option

SnapMirror Asynchronous protection with version-flexible replication and backup option provides mirror protection between source and destination volumes and the capability to store multiple copies of the mirrored data at the destination.

The storage administrator can specify which Snapshot copies are mirrored from source to destination and can also specify how long to retain those copies at the destination, even if they are deleted at the source.

In SnapMirror relationships with version-flexible replication and backup option, mirror operations do not

execute as quickly as they would in traditional SnapMirror relationships.

SnapMirror Unified Replication (mirror and vault)

SnapMirror unified replication allows you to configure disaster recovery and archiving on the same destination volume. As with SnapMirror, unified data protection performs a baseline transfer the first time you invoke it. A baseline transfer under the default unified data protection policy “MirrorAndVault” makes a Snapshot copy of the source volume, then transfers that copy and the data blocks it references to the destination volume. Like SnapVault, unified data protection does not include older Snapshot copies in the baseline.

SnapMirror Synchronous protection with strict synchronization

SnapMirror Synchronous protection with “strict” synchronization ensures that the primary and secondary volumes are always a true copy of each other. If a replication failure occurs when attempting to write data to the secondary volume, then the client I/O to the primary volume is disrupted.

SnapMirror Synchronous protection with regular synchronization

SnapMirror Synchronous protection with “regular” synchronization does not require that the primary and secondary volume are always a true copy of each other; thereby ensuring availability of the primary volume. If a replication failure occurs when attempting to write data to the secondary volume, the primary and secondary volumes fall out of sync and client I/O will continue to the primary volume.

> The Restore button and the Relationship operation buttons are not available when monitoring synchronous protection relationships from the Health: All Volumes view or the Volume / Health details page.

SnapMirror active sync

The SnapMirror active sync feature is available with ONTAP 9.8 and later, and you can use it to protect applications with LUNs, enabling applications to fail over transparently, ensuring business continuity in case of a disaster.

It allows you to discover and monitor the synchronous SnapMirror relationships for Consistency Groups (CGs) available on clusters and storage virtual machines from Unified Manager. SnapMirror active sync is supported on AFF clusters or All SAN Array (ASA) clusters, where the primary and secondary clusters can be either AFF or ASA. SnapMirror active sync protects applications with iSCSI or FCP LUNs.

When you view the volumes and LUNs protected by the SnapMirror active sync relationship, you can get a unified view for protection relationships, Consistency Groups in volume inventory, view protection topology for Consistency Group relationships, view historical data for Consistency Group relationships up to a year. You can also download the report. You can also view the summary of Consistency Group relationships, search support for Consistency Group relationships, and gain information about volumes that are protected by the Consistency Group.

On the Relationships page, you can also sort, filter, and extend protection on the source and destination storage objects and their relationship that are protected by the Consistency Group.

To know more about SnapMirror active sync, refer to ONTAP 9 Documentation for Snapmirror active sync (formerly SM-BC).

Set up protection relationships in Unified Manager

Set up protection relationships in Active IQ Unified Manager

There are several steps that you must perform to use Unified Manager and OnCommand Workflow Automation to set up SnapMirror and SnapVault relationships to protect your data.

Before you begin

- You must have the Application Administrator or Storage Administrator role.
- You must have established peer relationships between two clusters or two storage virtual machines (SVMs).
- OnCommand Workflow Automation must be integrated with Unified Manager:
  - Set up OnCommand Workflow Automation.
  - Verifying Unified Manager data source caching in Workflow Automation.

Steps

1. Depending on the type of protection relationship you want to create, do one of the following:
   - Create a SnapMirror protection relationship.
   - Create a SnapVault protection relationship.
2. If you want to create a policy for the relationship, depending on the relationship type you are creating, do one of the following:
   - Create a SnapVault policy.
   - Create a SnapMirror policy.
3. Create a SnapMirror or SnapVault schedule.

Configure a connection between Workflow Automation and Active IQ Unified Manager

You can configure a secure connection between OnCommand Workflow Automation (WFA) and Unified Manager. Connecting to Workflow Automation enables you to use protection features such as SnapMirror and SnapVault configuration workflows, as well as commands for managing SnapMirror relationships.

Before you begin

- The installed version of Workflow Automation must be 5.1.1P6 or greater.

  The “WFA pack for managing Clustered Data ONTAP” is included in WFA 5.1.1P6 so there is no need to download this pack from the NetAppStorage Automation Store and install it separately onto your WFA server as was required in the past. WFA download

- You must have the name of the database user that you created in Unified Manager to support WFA and Unified Manager connections.

  This database user must have been assigned the Integration Schema user role.

- You must be assigned either the Administrator role or the Architect role in Workflow Automation.

- You must have the host address, port number 443, user name, and password for the Workflow Automation setup.
- You must have the Application Administrator or Storage Administrator role.

Steps

1. In the left navigation pane, click **General** > **Workflow Automation**.
2. In the **Database User** area of the **Workflow Automation page**, select the name, and enter the password for the database user that you created to support Unified Manager and Workflow Automation connections.
3. In the **Workflow Automation Credentials** area of the page, enter the host name or IP address (IPv4 or IPv6), and the user name and password for the Workflow Automation setup.

   You must use the Unified Manager server port (port 443).

4. Click **Save**.
5. If you use a self-signed certificate, click **Yes** to authorize the security certificate.

   The Workflow Automation page displays.

6. Click **Yes** to reload the web UI, and add the Workflow Automation features.

Related information

NetApp Documentation: OnCommand Workflow Automation (current releases)

Verify Active IQ Unified Manager data source caching in Workflow Automation

You can determine whether Unified Manager data source caching is working correctly by checking if data source acquisition is successful in Workflow Automation. You might do this when you integrate Workflow Automation with Unified Manager to ensure that Workflow Automation functionality is available after the integration.

Before you begin

You must be assigned either the Administrator role or the Architect role in Workflow Automation to perform this task.

Steps

1. From the Workflow Automation UI, select **Execution** > **Data Sources**.
2. Right-click the name of the Unified Manager data source, and then select **Acquire Now**.
3. Verify that the acquisition succeeds without errors.

   Acquisition errors must be resolved for Workflow Automation integration with Unified Manager to succeed.

How reinstalling or upgrading OnCommand Workflow Automation affects Active IQ Unified Manager

Before reinstalling or upgrading OnCommand Workflow Automation, you must first remove the connection between OnCommand Workflow Automation and Unified Manager, and ensure that all OnCommand Workflow Automation currently running or scheduled jobs are stopped.

You must also manually delete Unified Manager from OnCommand Workflow Automation.

After you reinstall or upgrade OnCommand Workflow Automation, you must set up the connection with Unified Manager again.

Remove OnCommand Workflow Automation setup from Active IQ Unified Manager

You can remove the OnCommand Workflow Automation setup from Unified Manager when you no longer want to use Workflow Automation.

Before you begin

You must have the Application Administrator or Storage Administrator role.

Steps

1. In the left navigation pane, click **General** > **Workflow Automation** in the left Setup menu.
2. In the **Workflow Automation** page, click **Remove Setup**.

Perform a protection relationship failover and failback

Perform a protection relationship failover and failback in Active IQ Unified Manager

When a source volume in your protection relationship is disabled because of a hardware failure or a disaster, you can use the protection relationship features in Unified Manager to make the protection destination read/write accessible and fail over to that volume until the source is online again; then, you can fail back to the original source when it is available to serve data.

Before you begin

- You must have the Application Administrator or Storage Administrator role.
- You must have set up OnCommand Workflow Automation to perform this operation.

Steps

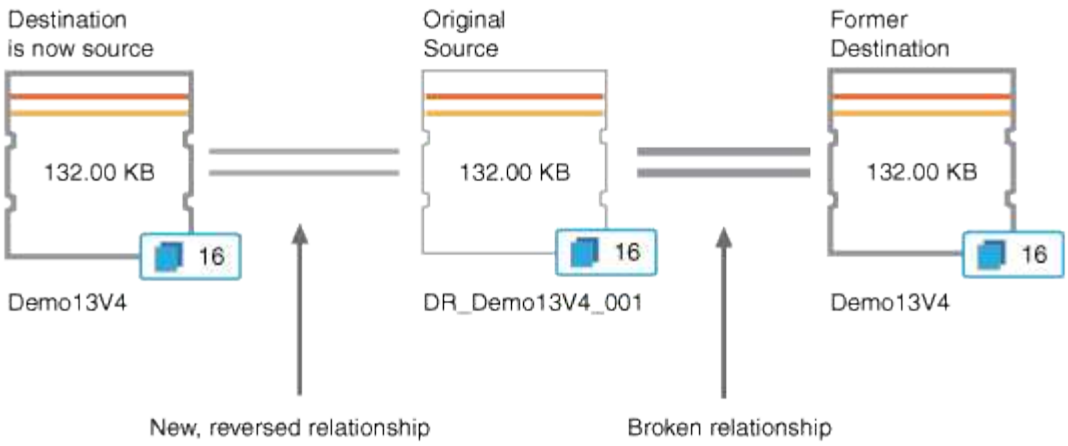
1. Break the SnapMirror relationship.

   You must break the relationship before you can convert the destination from a data protection volume to a read/write volume, and before you can reverse the relationship.

2. Reverse the protection relationship.

   When the original source volume is available again, you might decide to reestablish the original protection relationship by restoring the source volume. Before you can restore the source, you must synchronize it with the data written to the former destination. You use the reverse resync operation to create a new protection relationship by reversing the roles of the original relationship and synchronizing the source volume with the former destination. A new baseline Snapshot copy is created for the new relationship.

   The reversed relationship looks similar to a cascaded relationship:



3. Break the reversed SnapMirror relationship.

   When the original source volume is resynchronized and can again serve data, use the break operation to break the reversed relationship.

4. Remove the relationship.

   When the reversed relationship is no longer required, you should remove that relationship before reestablishing the original relationship.

5. Resynchronize the relationship.

   Use the resynchronize operation to synchronize data from the source to the destination and to reestablish the original relationship.

Break a SnapMirror protection relationship from the Active IQ Unified Manager Volume/Health page

You can break a protection relationship from the Volume / Health details page and stop data transfers between a source and destination volume in a SnapMirror relationship. You might break a relationship when you want to migrate data, for disaster recovery, or for application testing. The destination volume is changed to a read-write volume. You cannot break a SnapVault relationship.

Before you begin

- You must have the Application Administrator or Storage Administrator role.
- You must have set up Workflow Automation.

Steps

1. In the **Protection** tab of the **Volume / Health** details page, select from the topology the SnapMirror relationship you want to break.
2. Right-click the destination and select **Break** from the menu.

   The Break Relationship dialog box is displayed.

3. Click **Continue** to break the relationship.
4. In the topology, verify that the relationship is broken.

Reverse protection relationships from the Volume/ Health details page in Active IQ Unified Manager

When a disaster disables the source volume in your protection relationship, you can use the destination volume to serve data by converting it to read/write while you repair or replace the source. When the source is again available to receive data, you can use the reverse resynchronization operation to establish the relationship in the reverse direction, synchronizing the data on the source with the data on the read/write destination.

Before you begin

- You must have the Application Administrator or Storage Administrator role.
- You must have set up Workflow Automation.
- The relationship must not be a SnapVault relationship.
- A protection relationship must already exist.
- The protection relationship must be broken.
- Both the source and destination must be online.
- The source must not be the destination of another data protection volume.
- When you perform this task, data on the source that is newer than the data on the common Snapshot copy is deleted.
- Policies and schedules created on the reverse resynchronization relationship are the same as those on the original protection relationship.

  If policies and schedules do not exist, they are created.

Steps

1. From the **Protection** tab of the **Volume / Health** details page, locate in the topology the SnapMirror relationship on which you want to reverse the source and destination, and right-click it.
2. Select **Reverse Resync** from the menu.

   The Reverse Resync dialog box is displayed.

3. Verify that the relationship displayed in the **Reverse Resync** dialog box is the one for which you want to perform the reverse resynchronization operation, and then click **Submit**.

   The Reverse Resync dialog box is closed and a job link is displayed at the top of the Volume / Health details page.

4. **Optional:** Click **View Jobs** on the **Volume / Health** details page to track the status of each reverse resynchronization job.

   A filtered list of jobs is displayed.

5. **Optional:** Click the **Back** arrow on your browser to return to the **Volume / Health** details page.

   The reverse resynchronization operation is finished when all job tasks are completed successfully.

Remove a protection relationship from the Volume/ Health details page in Active IQ Unified Manager

You can remove a protection relationship to permanently delete an existing relationship between the selected source and destination: for example, when you want to create a

relationship using a different destination. This operation removes all metadata and cannot be undone.

Before you begin

- You must have the Application Administrator or Storage Administrator role.
- You must have set up Workflow Automation.

Steps

1. In the **Protection** tab of the **Volume / Health** details page, select from the topology the SnapMirror relationship you want to remove.
2. Right-click the name of the destination and select **Remove** from the menu.

   The Remove Relationship dialog box is displayed.
3. Click **Continue** to remove the relationship.

   The relationship is removed from the Volume / Health details page.

Resynchronize protection relationships from the Volume/Health details page in Active IQ Unified Manager

You can resynchronize data on a SnapMirror or SnapVault relationship that was broken and then the destination was made read/write so that data on the source matches the data on the destination. You might also resynchronize when a required common Snapshot copy on the source volume is deleted causing SnapMirror or SnapVault updates to fail.

Before you begin

- You must have the Application Administrator or Storage Administrator role.
- You must have set up OnCommand Workflow Automation.

Steps

1. From the **Protection** tab of the **Volume / Health** details page, locate in the topology the protection relationship that you want to resynchronize and right-click it.
2. Select **Resynchronize** from the menu.

   Alternatively, from the **Actions** menu, select **Relationship** > **Resynchronize** to resynchronize the relationship for which you are currently viewing the details.

   The Resynchronize dialog box is displayed.
3. In the **Resynchronization Options** tab, select a transfer priority and the maximum transfer rate.
4. Click **Source Snapshot Copies**; then, in the **Snapshot Copy** column, click **Default**.

   The Select Source Snapshot Copy dialog box is displayed.
5. If you want to specify an existing Snapshot copy rather than transferring the default Snapshot copy, click **Existing Snapshot Copy** and select a Snapshot copy from the list.
6. Click **Submit**.

You are returned to the Resynchronize dialog box.

7. If you selected more than one source to resynchronize, click **Default** for the next source for which you want to specify an existing Snapshot copy.
8. Click **Submit** to begin the resynchronization job.

   The resynchronization job is started, you are returned to the Volume / Health details page and a jobs link is displayed at the top of the page.

9. **Optional:** Click **View Jobs** on the **Volume / Health details** page to track the status of each resynchronization job.

   A filtered list of jobs is displayed.

10. **Optional:** Click the **Back** arrow on your browser to return to the **Volume / Health** details page.

    The resynchronization job is finished when all job tasks successfully complete.

Resolve a protection job failure

Resolve a protection job failure in Active IQ Unified Manager

This workflow provides an example of how you might identify and resolve a protection job failure from the Unified Manager dashboard.

Before you begin

Because some tasks in this workflow require that you log in using the Administrator role, you must be familiar with the roles required to use various functionality.

In this scenario, you access the Dashboard page to see if there are any issues with your protection jobs. In the Protection Incident area, you notice that there is a Job Terminated incident, showing a Protection Job Failed error on a volume. You investigate this error to determine the possible cause and potential resolution.

Steps

1. In the Protection Incidents panel of the Dashboard Unresolved Incidents and Risks area, you click the **Protection job failed** event.

   The linked text for the event is written in the form `object_name:/object_name - Error Name`, such as `cluster2_src_svm:/cluster2_src_vol2 - Protection Job Failed`.

   The Event details page for the failed protection job displays.

2. Review the error message in the Cause field of the **Summary** area to determine the problem and evaluate potential corrective actions.

   See Identifying the problem and performing corrective actions for a failed protection job.

Identify the problem and perform corrective actions for a failed protection job in Active IQ Unified Manager

You review the job failure error message in the Cause field on the Event details page and

determine that the job failed because of a Snapshot copy error. You then proceed to the Volume / Health details page to gather more information.

Before you begin

You must have the Application Administrator role.

The error message provided in the Cause field on the Event details page contains the following text about the failed job:

```
Protection Job Failed. Reason: (Transfer operation for
relationship 'cluster2_src_svm:cluster2_src_vol2->cluster3_dst_svm:
managed_svc2_vol3' ended unsuccessfully. Last error reported by
Data ONTAP: Failed to create Snapshot copy 0426cluster2_src_vol2snap
on volume cluster2_src_svm:cluster2_src_vol2. (CSM: An operation
failed due to an ONC RPC failure.)
Job Details
```

This message provides the following information:

- A backup or mirror job did not complete successfully.

  The job involved a protection relationship between the source volume `cluster2_src_vol2` on the virtual server `cluster2_src_svm` and the destination volume `managed_svc2_vol3` on the virtual server named `cluster3_dst_svm`.

- A Snapshot copy job failed for `0426cluster2_src_vol2snap` on the source volume `cluster2_src_svm:/cluster2_src_vol2`.

In this scenario, you can identify the cause and potential corrective actions of the job failure. However, resolving the failure requires that you access either the System Manager web UI or the ONTAP CLI commands.

Steps

1. You review the error message and determine that a Snapshot copy job failed on the source volume, indicating that there is probably a problem with your source volume.

   Optionally, you could click the **Job Details** link at the end of the error message, but for the purposes of this scenario, you choose not to do that.

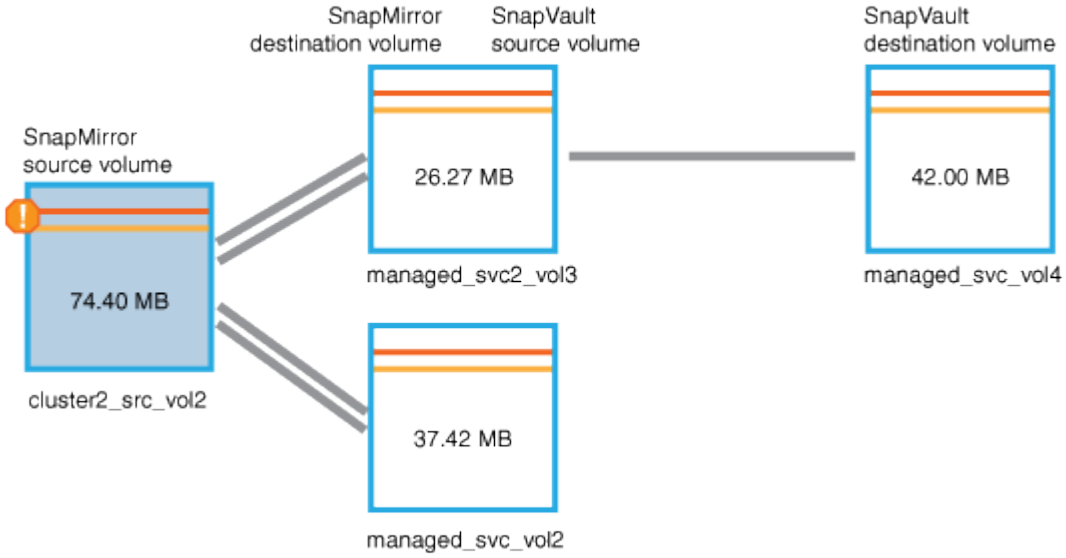
2. You decide that you want to try to resolve the event, so you do the following:
   a. Click the **Assign To** button and select **Me** from the menu.
   b. Click the **Acknowledge** button so that you do not continue to receive repeat alert notifications, if alerts were set for the event.
   c. Optionally, you can also add notes about the event.
3. Click the **Source** field in the **Summary** pane to see details about the source volume.

   The **Source** field contains the name of the source object: in this case, the volume on which the Snapshot copy job was scheduled.

The Volume / Health details page displays for `cluster2_src_vol2`, showing the content of the Protection tab.

4. Looking at the protection topology graph, you see an error icon associated with the first volume in the topology, which is the source volume for the SnapMirror relationship.

   You also see the horizontal bars in the source volume icon, indicating the warning and error thresholds set for that volume.



5. You place your cursor over the error icon to see the pop-up dialog box that displays the threshold settings and see that the volume has exceeded the error threshold, indicating a capacity issue.
6. Click the **Capacity** tab.

   Capacity information about volume `cluster2_src_vol2` displays.

7. In the **Capacity** panel, you see that there is an error icon in the bar graph, again indicating that the volume capacity has surpassed the threshold level set for the volume.
8. Below the capacity graph, you see that volume autogrow has been disabled and that a volume space guarantee has been set.

   You could decide to enable autogrow, but for the purposes of this scenario, you decide to investigate further before making a decision about how to resolve the capacity problem.

9. You scroll down to the **Events** list and see that Protection Job Failed, Volume Days Until Full, and Volume Space Full events were generated.
10. In the **Events** list, you click the **Volume Space Full** event to get more information, having decided that this event seems most relevant to your capacity issue.

    The Event details page displays the Volume Space Full event for the source volume.

11. In the **Summary** area, you read the Cause field for the event: `The full threshold set at 90% is breached. 45.38 MB (95.54%) of 47.50 MB is used.`
12. Below the Summary area, you see Suggested Corrective Actions.



The Suggested Corrective Actions display only for some events, so you do not see this area

for all types of events.

You click through the list of suggested actions that you might perform to resolve the Volume Space Full event:

- Enable autogrow on this volume.
- Resize the volume.
- Enable and run deduplication on this volume.
- Enable and run compression on this volume.

13. You decide to enable autogrow on the volume, but to do so, you must determine the available free space on the parent aggregate and the current volume growth rate:
    a. Look at the parent aggregate, `cluster2_src_aggr1`, in the **Related Devices** pane.

    You can click the name of the aggregate to get further details about the aggregate.

    You determine that the aggregate has sufficient space to enable volume autogrow.

    b. At the top of the page, look at the icon indicating a critical incident and review the text below the icon.

    You determine that "Days to Full: Less than a day | Daily Growth Rate: 5.4%".

14. Go to System Manager or access the ONTAP CLI to enable the `volume autogrow` option.

    Make note of the names of the volume and aggregate so you have them available when enabling autogrow.

15. After resolving the capacity issue, return to the Unified Manager **Event** details page and mark the event as resolved.

Resolve lag issues in Active IQ Unified Manager

This workflow provides an example of how you might resolve a lag issue. In this scenario, you are an administrator or operator accessing the Unified ManagerDashboard page to see if there are any problems with your protection relationships and, if they exist, to find solutions.

Before you begin

You must have the Application Administrator or Storage Administrator role.

In the Dashboard page, you look at the Unresolved Incidents and Risks area and see a SnapMirror Lag error in the Protection pane under Protection Risks.

Steps

1. In the **Protection** pane on the **Dashboard** page, locate the SnapMirror relationship lag error and click it.

   The Event details page for the lag error event is displayed.

2. From the **Event** details page you can perform one or more of the following tasks:
   - Review the error message in the Cause field of the Summary area to determine if there is any

suggested corrective action.

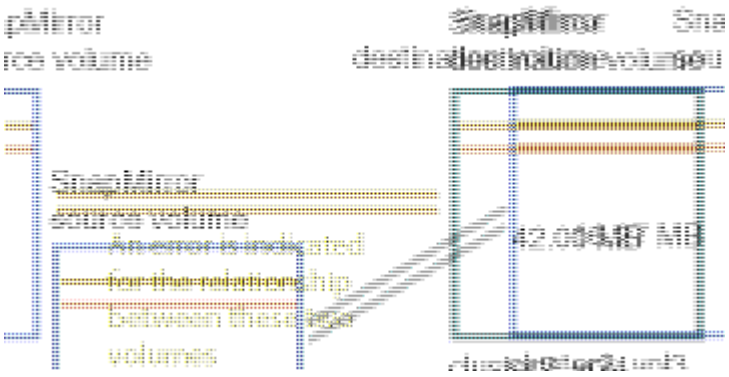
- Click the object name, in this case a volume, in the Source field of the Summary area to get details about the volume.
- Look for notes that might have been added about this event.
- Add a note to the event.
- Assign the event to a specific user.
- Acknowledge or resolve the event.

3. In this scenario, you click the object name (in this case, a volume) in the Source field of the **Summary** area to get details about the volume.

   The Protection tab of the Volume / Health details page is displayed.

4. In the **Protection** tab, you look at the topology diagram.

   You note that the volume with the lag error is the last volume in a three-volume SnapMirror cascade. The volume you selected is outlined in dark gray, and a double orange line from the source volume indicates a SnapMirror relationship error.



5. Click each of the volumes in the SnapMirror cascade.

   As you select each volume, the protection information in the Summary, Topology, History, Events, Related Devices, and Related Alerts areas changes to display details relevant to the selected volume.

6. You look at the **Summary** area and position your cursor over the information icon in the **Update Schedule** field for each volume.

   In this scenario, you note that the SnapMirror policy is DPDefault, and the SnapMirror schedule updates hourly at five minutes after the hour. You realize that all of the volumes in the relationship are attempting to complete a SnapMirror transfer at the same time.

7. To resolve the lag issue, you modify the schedules for two of the cascaded volumes so that each destination begins a SnapMirror transfer after its source has completed a transfer.

Manage and monitor protection relationships

Learn about managing and monitoring protection relationships in Active IQ Unified Manager

Active IQ Unified Manager enables you to create protection relationships, to monitor and

troubleshoot SnapMirror and SnapVault relationships on managed clusters, and to restore data when it is overwritten or lost.

For SnapMirror operations there are two replication types:

- Asynchronous

  Replication from the primary to the secondary volume is determined by a schedule.

- Synchronous

  Replication is performed simultaneously on the primary and secondary volume.

You can perform up to 10 protection jobs simultaneously with no performance impact. You might experience some performance impact when you run between 11 and 30 jobs simultaneously. Running more than 30 jobs simultaneously is not recommended.

View volume protection status in Active IQ Unified Manager

The Data Protection page presents a holistic view of the data protection details for all the protected volumes in a single cluster, or all clusters in a data center.

Steps

1. In the left navigation pane, click **Dashboard**.
2. Depending on whether you want to view data protection status for all monitored clusters or for a single cluster, select **All Clusters** or select a single cluster from the drop-down menu.
3. Click the right-arrow at the top of the Data Protection panel. The **Data Protection** page is displayed.

Depending on whether you selected a single or all clusters in your data center, this page displays the data protection status of the volumes protected by Snapshot copies or SnapMirror policies, and also displays the count of the unprotected volumes.

Selecting a cluster from the **Individual Cluster** list displays the Snapshot protection and SnapMirror relationship status of the protected volumes in that cluster.

Clicking the events on this page takes you to the Event details page. You can click the **View All** link to view all active protection events in the Event Management inventory page. You can hover your mouse to view the respective counts and legends. You can click:

- The bar charts for unprotected volumes and volumes protected by Snapshot copies to go to the Volumes page and view the details.
- The bar charts for all relationships to go to the Relationships page, where the details are filtered by the source cluster.

View protection status of volumes protected by Snapshot copies

Snapshot Copies Overview: An overview of the volumes protected by Snapshot copies, such as:

- The total number of volumes protected and not protected by Snapshot copies.
- The total number of volumes that are using or exceeding the reserve space for the Snapshot copies.

Snapshot Copies Analysis details the following information:

- Individual events for Snapshot copies, including the events raised in the last 24 hours.
- Detailed chart for volumes protected and not protected by Snapshot copies.
- Volumes using, not using, and breaching the reserved Snapshot copy capacity.
- Break-up of the volume counts in terms of the number of their Snapshot copies.

Points to note for Snapshot copies

- For counting the volumes protected by Snapshot copies, both source and destination volumes are considered.
- The number of Snapshot copies returned is for only the volumes that are online and available.
- The graph range for the number of Snapshot copies is dynamic. It is generated based on the number of snapshots present in the selected cluster.
- For considering a volume as protected, the schedule for the Snapshot copy of the volume should be enabled.
- The value for the reserve space for Snapshot copies is important for viewing the amount of disk space utilized or for calculating the space that can be reclaimed if one or more Snapshot copies are deleted.

View protection status for SnapMirror relationships

SnapMirror Overview: An overview of the volumes protected by SnapMirror policies, such as:

- The number of volumes protected by the respective SnapMirror policies, such as, volume SnapMirror relationships, storage VM disaster recovery (SVM-DR), and their combinations.
- The total number of SnapMirror relationships experiencing recovery point objective (RPO) lag based on the lag status.

SnapMirror Analysis details the following information:

- Individual events raised for SnapMirror relationships, including the events raised in the last 24 hours
- The number of volumes protected by each type of SnapMirror policy.
- The number of relationships protected by the SnapMirror relationship types, such as Asynchronous Mirror, Asynchronous Vault, Asynchronous MirrorVault, StricxtSync, SnapMirror active sync consistency group, and Sync.
- The number of healthy and unhealthy relationships.
- Break-up of the volume relationship count. You can toggle the charts by RPO lag time and status.
- Lag thresholds for unmanaged relationship. You can click the settings icon ( ) to configure the lag threshold settings. For more information, see Configuring lag threshold settings for unmanaged protection relationships.

Points to note for SnapMirror relationships

- For counting SnapMirror relationships, the source volumes, which are enabled for reading and writing, are counted. Destination and root volumes are not considered.
- For SnapMirror relationship, the events are displayed for the source cluster.
- The SnapMirror relationship count includes the number of volumes with sources and destinations on the same or different clusters.

- The lag duration of the RPO lag in data replication is based on the SnapMirror relation. The status is categorized as `ok`, `warning`, or `error`, based on the relationship threshold set. You can look at the status to determine whether the parameters are working as expected, or whether you need to address any issue.
- If a volume has multiple SnapMirror relationships, each relationship type is counted for the RPO lag.
- Volume relationships are considered unhealthy if there are issues in data replication between the source and destination, for example, when the relationship is broken.

View clusters protected by MetroCluster configuration

The **MetroCluster Protection** panel on the **Data Protection** page displays the number of clusters protected or not protected by MetroCluster over FC or IP configuration at your site. Clicking the bar charts on this panel takes you to the Clusters page where the cluster details are filtered based on the protected or unprotected clusters. For information about monitoring your MetroCluster configuration, see Monitoring MetroCluster configurations.

View volume protection relationships in Active IQ Unified Manager

From the Relationship: All Relationships view, and from the Volume Relationships page, you can view the status of existing volume SnapMirror and SnapVault relationships. You can also examine details about protection relationships, including transfer and lag status, source and destination details, schedule and policy information, and so on.

Before you begin

You must have the Application Administrator or Storage Administrator role.

You can also initiate relationship commands from this page.

Steps

1. In the left navigation pane, click **Storage** > **Volumes**.
2. From the View menu, select **Relationship** > **All Relationships**.

   The Relationship: All Relationships view is displayed.

3. Choose one of the following ways to view the volume protection details:
   - To view current information about all the volume relationships, remain on the default **All Relationships** page.
   - To view detailed information about the volume transfer trends over a period of time, in the View menu, select Relationship: Last 1 month Transfer Status view.
   - To view detailed information about the volume transfer activity on a day to day basis, in the View menu, select Relationship: Last 1 month Transfer Rate view.

   The volume transfer views display information for volumes in asynchronous relationships only - volumes in synchronous relationships are not shown.

Monitor Consistency Group LUN relationships in Active IQ Unified Manager

If your ONTAP environment supports SnapMirror active sync to protect applications with LUNs, you can view and monitor those LUNs on Active IQ Unified Manager.

SnapMirror active sync ensures zero recovery time objective (RTO) during failover in SAN environments. In a typical deployment supporting SnapMirror active sync, the LUNs on volumes are protected by Consistency Group relationships.

These primary and secondary LUNs are composite LUNs, or a replica LUN pair with the same UUID and serial number. The I/O operations (both read and write) are multiplexed across the source and destination sites on these composite LUNs, ensuring transparency.

For viewing composite LUNs, both the primary and secondary clusters with the LUNs that are a part of the Consistency Group relationship should be added and discovered on Unified Manager. Only iSCSI and FCP LUNs are supported.

For information about SnapMirror active sync, see ONTAP 9 Documentation for SnapMirror active sync (formerly SM-BC).

To view composite LUNs in your environment follow these steps:

Steps

1. In the left navigation pane, click **Storage** > **LUNs**.
2. From the View menu, select **Relationship** > **All LUNs**.

   The Relationship: All LUNs view is displayed.

You can view the LUN details, such as the LUN name, volume, storage VM hosting the LUN, cluster, Consistency Group, and the partner LUN. You can click each of these components to drill down to a detailed view. Clicking the Consistency Group takes you to the Relationships page.

Clicking the partner LUN enables you to view its configuration details on the SAN tab of the Storage VM Details page for the storage VM on which the partner LUN is hosted. Information such as the initiators and initiator groups and other aspects of the partner LUN is displayed.

You can perform the standard grid-level functions of sorting, filtering, generating and uploading reports for the protected LUNs in your environment.

Create a SnapVault relationship from the Health: All Volumes view in Active IQ Unified Manager

You can use the Health: All Volumes view to create SnapVault relationships for one or more volumes on the same Storage VM to enable data backups for protection purposes.

Before you begin

- You must have the Application Administrator or Storage Administrator role.
- You must have set up Workflow Automation.

The **Protect** menu does not display in the following instances:

- If RBAC settings do not allow this action: for example, if you have only operator privileges
- When the volume ID is unknown: for example, when you have an intercluster relationship and the destination cluster has not yet been discovered

Steps

1. In the left navigation pane, click **Storage** > **Volumes**.

2. In the **Health: All Volumes** view, select a volume you want to protect and click **Protect**.

   Alternatively, to create multiple protection relationships on the same storage virtual machine (SVM), select one or more volumes in the Health: All Volumes view, and click **Protect** on the toolbar.

3. Select **SnapVault** from the menu.

   The Configure Protection dialog box is launched.

4. Click **SnapVault** to view the **SnapVault** tab and to configure the secondary volume information.
5. Click **Advanced** to set deduplication, compression, autogrow, and space guarantee as needed, and then click **Apply**.
6. Complete the **Destination Information** area and the **Relationship Settings** area in the **SnapVault** tab.
7. Click **Apply**.

   You are returned to the Health: All Volumes view.

8. Click the protection configuration job link at the top of the **Health: All Volumes** view.

   If you are creating only one protection relationship, the Job details page is displayed; however, if you are creating more than one protection relationship, a filtered list of all the jobs associated with the protection operation is displayed.

9. Do one of the following:
   - If you have only one job, click **Refresh** to update the tasks list and task details associated with the protection configuration job and to determine when the job is complete.
   - If you have more than one job:
     i. Click a job in the jobs list.
     ii. Click **Refresh** to update the tasks list and task details associated with the protection configuration job and to determine when the job is complete.
     iii. Use the **Back** button to return to the filtered list and view another job.

Create a SnapVault protection relationship from the Volume/Health details page in Active IQ Unified Manager

You can create a SnapVault relationship using the Volume / Health details page so that data backups are enabled for protection purposes on volumes.

Before you begin

- You must have the Application Administrator or Storage Administrator role.
- You must have set up Workflow Automation to perform this task.

The **Protect** menu does not display in the following instances:

- If RBAC settings do not allow this action: for example, if you have only operator privileges
- When the volume ID is unknown: for example, when you have an intercluster relationship and the destination cluster has not yet been discovered

Steps

1. In the **Protection** tab of the **Volume / Health** details page, right-click a volume in the topology view that you want to protect.
2. Select **Protect** > **SnapVault** from the menu.

   The Configure Protection dialog box is launched.
3. Click **SnapVault** to view the **SnapVault** tab and to configure the secondary resource information.
4. Click **Advanced** to set deduplication, compression, autogrow, and space guarantee as needed, and then click **Apply**.
5. Complete the **Destination Information** area and the **Relationship Settings** area in the **Configure Protection** dialog box.
6. Click **Apply**.

   You are returned to the Volume / Health details page.
7. Click the protection configuration job link at the top of the **Volume / Health** details page.

   The Job details page is displayed.
8. Click **Refresh** to update the tasks list and task details associated with the protection configuration job and to determine when the job is complete.

   When the job tasks are complete, the new relationships are displayed in the Volume / Health details page topology view.

Create a SnapMirror relationship from the Health: All Volumes view in Active IQ Unified Manager

Using the Health: All Volumes view enables you to create several SnapMirror protection relationships at one time by selecting more than one volume on the same storage VM.

Before you begin

- You must have the Application Administrator or Storage Administrator role.
- You must have set up Workflow Automation.

The **Protect** menu does not display in the following instances:

- If RBAC settings do not allow this action: for example, if you have only operator privileges
- When the volume ID is unknown: for example, when you have an intercluster relationship and the destination cluster has not yet been discovered

Steps

1. In the **Health: All Volumes** view, select a volume that you want to protect.

   Alternatively, to create multiple protection relationships on the same SVM, select one or more volumes in the Health: All Volumes view, and click **Protect** > **SnapMirror** on the toolbar.

   The Configure Protection dialog box is displayed.
2. Click **SnapMirror** to view the **SnapMirror** tab and to configure the destination information.

3. Click **Advanced** to set the space guarantee, as needed, and then click **Apply**.
4. Complete the **Destination Information** area and the **Relationship Settings** area in the **SnapMirror** tab.
5. Click **Apply**.

   You are returned to the Health: All Volumes view.

6. Click the protection configuration job link at the top of the **Health: All Volumes view**.

   If you are creating only one protection relationship, the Job details page is displayed; however, if you are creating more than one protection relationship, a list of all the jobs associated with the protection operation is displayed.

7. Do one of the following:
   - If you have only one job, click **Refresh** to update the tasks list and task details associated with the protection configuration job and to determine when the job is complete.
   - If you have more than one job:
     i. Click a job in the jobs list.
     ii. Click **Refresh** to update the tasks list and task details associated with the protection configuration job and to determine when the job is complete.
     iii. Use the **Back** button to return to the filtered list and view another job.

Depending on the destination SVM you specified during configuration or on the options you enabled in your Advanced settings, the resulting SnapMirror relationship might be one of several possible variations:

- If you specified a destination SVM that runs under the same or a newer version of ONTAP compared to that of the source volume, a block-replication-based SnapMirror relationship is the default result.
- If you specified a destination SVM that runs under the same or a newer version of ONTAP compared to that of the source volume, but you enabled version-flexible replication in the Advanced settings, a SnapMirror relationship with version-flexible replication is the result.
- If you specified a destination SVM that runs under an earlier version of ONTAP than that of the source volume, and the earlier version supports version-flexible replication, a SnapMirror relationship with version-flexible replication is the automatic result.

Create a SnapMirror relationship from the Volume Health details page in Active IQ Unified Manager

You can use the Volume / Health details page to create a SnapMirror relationship so that data replication is enabled for protection purposes. SnapMirror replication enables you to restore data from the destination volume in the event of data loss on the source.

Before you begin

- You must have the Application Administrator or Storage Administrator role.
- You must have set up Workflow Automation.

The **Protect** menu does not display in the following instances:

- If RBAC settings do not allow this action: for example, if you have only operator privileges
- When the volume ID is unknown: for example, when you have an intercluster relationship and the destination cluster has not yet been discovered

You can perform up to 10 protection jobs simultaneously with no performance impact. You might experience some performance impact when you run between 11 and 30 jobs simultaneously. Running more than 30 jobs simultaneously is not recommended.

Steps

1. In the **Protection tab** of the **Volume / Health** details page, right-click in the topology view the name of a volume that you want to protect.
2. Select **Protect** > **SnapMirror** from the menu.

   The Configure Protection dialog box is displayed.
3. Click **SnapMirror** to view the **SnapMirror** tab and to configure the destination information.
4. Click **Advanced** to set the space guarantee, as needed, and then click **Apply**.
5. Complete the **Destination Information** area and the **Relationship Settings** area in the **Configure Protection** dialog box.
6. Click **Apply**.

   You are returned to the Volume / Health details page.
7. Click the protection configuration job link at the top of the **Volume / Health** details page.

   The job's tasks and details are displayed in the Job details page.
8. In the **Job** details page, click **Refresh** to update the tasks list and task details associated with the protection configuration job and to determine when the job is complete.
9. When the job tasks are complete, click **Back** on your browser to return to the **Volume / Health** details page.

   The new relationship is displayed in the Volume / Health details page topology view.

Depending on the destination SVM you specified during configuration or on the options you enabled in your Advanced settings, the resulting SnapMirror relationship might be one of several possible variations:

- If you specified a destination SVM that runs under the same or a newer version of ONTAP compared to that of the source volume, a block-replication-based SnapMirror relationship is the default result.
- If you specified a destination SVM that runs under the same or a newer version of ONTAP compared to that of the source volume, but you enabled version-flexible replication in the Advanced settings, a SnapMirror relationship with version-flexible replication is the result.
- If you specified a destination SVM that runs under an earlier version of ONTAP, or a version that is higher than that of the source volume and the earlier version supports version-flexible replication, a SnapMirror relationship with version-flexible replication is the automatic result.

Create a SnapMirror relationship with version-flexible replication in Active IQ Unified Manager

You can create a SnapMirror relationship with version-flexible replication. Version-flexible replication enables you to implement SnapMirror protection even if source and destination volumes run under different versions of ONTAP.

Before you begin

- You must have the Application Administrator or Storage Administrator role.
- You must have set up Workflow Automation.
- The source and destination SVMs must each have a SnapMirror license enabled.
- The source and destination SVMs must each run under a version of ONTAP software that supports version-flexible replication.

SnapMirror with version-flexible replication enables you to implement SnapMirror protection even in heterogeneous storage environments in which not all storage is running under one version of ONTAP; however, mirror operations performed under SnapMirror with version-flexible replication do not execute as quickly as they would under traditional block replication SnapMirror.

Steps

1. Display the **Configure Protection** dialog box for the volume that you want to protect.
   - If you are viewing the Protection tab of the Volume / Health details page, right-click in the topology view that has the name of a volume that you want to protect and select **Protect** > **SnapMirror** from the menu.
   - If you are viewing the Health: All Volumes view, locate a volume that you want to protect and right-click it; then select **Protect** > **SnapMirror** from the menu.
     The Configure Protection dialog box is displayed.
2. Click **SnapMirror** to view the **SnapMirror** tab.
3. Complete the **Destination Information** area and the **Relationship Settings** area in the **Configure Protection** dialog box.

   If you specify a destination SVM that runs under an earlier version of ONTAP than the source volume you are protecting, and if that earlier version supports version-flexible replication, this task automatically configures SnapMirror with version-flexible replication.

4. If you specify a destination SVM that runs under the same version of ONTAP as that of the source volume, but you still want to configure SnapMirror with version-flexible replication, click **Advanced** to enable version-flexible replication and then click **Apply**.
5. Click **Apply**.

   You are returned to the Volume / Health details page.

6. Click the protection configuration job link at the top of the **Volume / Health** details page.

   The jobs tasks and details are displayed in the Job details page.

7. In the **Job** details page, click **Refresh** to update the tasks list and task details associated with the protection configuration job and to determine when the job is complete.
8. When the job tasks are complete, click **Back** on your browser to return to the **Volume / Health** details page.

   The new relationship is displayed in the Volume / Health details page topology view.

Create SnapMirror relationships with version-flexible replication and back option capability in Active IQ Unified Manager

You can create a SnapMirror relationship with version-flexible replication and backup option capability. Backup option capability enables you to implement SnapMirror

protection and also retain multiple versions of backup copies at the destination location.

Before you begin

- You must have the Application Administrator or Storage Administrator role.
- You must have set up Workflow Automation.
- The source and destination SVMs must each have a SnapMirror license enabled.
- The source and destination SVMs must each have a SnapVault license enabled.
- The source and destination SVMs must each run under a version of ONTAP software that supports version-flexible replication.

Configuring SnapMirror with backup option capability enables you to protect your data with SnapMirror disaster recovery capabilities, such as volume failover ability, and at the same time provide SnapVault capabilities, such as multiple backup copy protection.

Steps

1. Display the **Configure Protection** dialog box for the volume that you want to protect.
   - If you are viewing the Protection tab of the Volume / Health details page, right-click in the topology view the name of a volume that you want to protect and select **Protect** > **SnapMirror** from the menu.
   - If you are viewing the Health: All Volumes view, locate a volume you want to protect and right-click it; then select **Protect** > **SnapMirror** from the menu.
     The Configure Protection dialog box is displayed.
2. Click **SnapMirror** to view the **SnapMirror** tab.
3. Complete the **Destination Information** area and the **Relationship Settings** area in the **Configure Protection** dialog box.
4. Click **Advanced** to display the **Advanced Destination Settings** dialog box.
5. If the **Version-Flexible Replication** check box is not already selected, select it now.
6. Select the **With backup option** check box to enable backup option capability; then click **Apply**.
7. Click **Apply**.

   You are returned to the Volume / Health details page.

8. Click the protection configuration job link at the top of the **Volume / Health** details page.

   The jobs tasks and details are displayed in the Job details page.

9. In the **Job** details page, click **Refresh** to update the task list and task details associated with the protection configuration job and to determine when the job is complete.
10. When the job tasks are complete, click **Back** on your browser to return to the **Volume / Health** details page.

    The new relationship is displayed in the Volume / Health details page topology view.

Configure destination efficiency settings in Active IQ Unified Manager

You can configure destination efficiency settings such as deduplication, compression, autogrow, and space guarantee on a protection destination using the Advanced Destination Settings dialog box. You use these settings when you want to maximize

space utilization on a destination or secondary volume.

Before you begin

You must have the Application Administrator or Storage Administrator role.

By default, efficiency settings match those of the source volume, except for compression settings in a SnapVault relationship, which are disabled by default.

Steps

1. Click either the **SnapMirror** tab or the **SnapVault** tab in the **Configure Protection** dialog box, depending on the type of relationship you are configuring.
2. Click **Advanced** in the **Destination Information** area.

   The Advanced Destination Settings dialog box is opened.
3. Enable or disable the efficiency settings for deduplication, compression, autogrow, and space guarantee, as required.
4. Click **Apply** to save your selections and return to the **Configure Protection** dialog box.

Create SnapMirror and SnapVault schedules in Active IQ Unified Manager

You can create basic or advanced SnapMirror and SnapVault schedules to enable automatic data protection transfers on a source or primary volume so that transfers take place more frequently or less frequently, depending on how often the data changes on your volumes.

Before you begin

- You must have the Application Administrator or Storage Administrator role.
- You must have already completed the Destination Information area in the Configure Protection dialog box.
- You must have set up Workflow Automation to perform this task.

Steps

1. From the **SnapMirror** tab or **SnapVault** tab of the **Configure Protection** dialog box, click the **Create Schedule** link in the **Relationship Settings** area.

   The Create Schedule dialog box is displayed.
2. In the **Schedule Name** field, type the name you want to give to the schedule.
3. Select one of the following:
   - **Basic**

     Select if you want to create a basic interval-style schedule.
   - **Advanced**

     Select if you want to create a cron-style schedule.
4. Click **Create**.

   The new schedule is displayed in the SnapMirror Schedule or SnapVault Schedule drop-down list.

Extend protection using cascade or fanout relationships in Active IQ Unified Manager

You can extend protection from an existing relationship by creating either a fanout from the source volume or a cascade from the destination volume of an existing relationship. You might do this when you need to copy data from one site to many sites or to provide additional protection by creating more backups.

You can extend protection to volumes using consistency group, which is a container that holds several volumes so that you can manage all of the volumes as one entity. You can view the SnapMirror active sync Consistency Group and the synchronous consistency group relationship in the Relationships page of Unified Manager.

Before you begin

- You must have the Application Administrator or Storage Administrator role.
- You must have set up Workflow Automation.

Steps

1. Click **Protection** > **Relationships**.
   Alternatively, you view the relationships from the Volume details page.
2. From the **Volume Relationships** page, select the SnapMirror relationship from which you want to extend protection.
3. On the action bar, click **Extend Protection**.
4. In the menu, select either **From Source** or **From Destination**, depending on whether you are creating a fanout relationship from the source or a cascade relationship from the destination.
5. Select either **With SnapMirror** or **With SnapVault** depending on the type of protection relationship you are creating.

   The **Configure Protection** dialog box is displayed.

   This can be achieved from the unified relationship / Volume Relationship and Volume / Health details page.

6. Complete the information as indicated in the **Configure Protection** dialog box.

Edit protection relationships from the Volume Relationships page in Active IQ Unified Manager

You can edit existing protection relationships to change the maximum transfer rate, the protection policy, or the protection schedule. You might edit a relationship to decrease the bandwidth used for transfers, or to increase the frequency of scheduled transfers because data is changing often.

Before you begin

You must have the Application Administrator or Storage Administrator role.

The selected volumes must be protection relationship destinations. You cannot edit relationships when source volumes, load-sharing volumes, or volumes that are not the destination of a SnapMirror or SnapVault relationship are selected.

Steps

1. From the **Volume Relationships** page, select in the volumes list one or more volumes in the same SVM for which you want to edit relationship settings, and then select **Edit** from the toolbar.

   The Edit Relationship dialog box is displayed.

2. In the **Edit Relationship** dialog box, edit the maximum transfer rate, protection policy, or protection schedule, as needed.
3. Click **Apply**.

   The changes are applied to the selected relationships.

Edit protection relationships from the Volume/ Health details page in Active IQ Unified Manager

You can edit existing protection relationships to change the current maximum transfer rate, protection policy, or protection schedule. You might edit a relationship to decrease the bandwidth used for transfers, or to increase the frequency of scheduled transfers because data is changing often.

Before you begin

- You must have the Application Administrator or Storage Administrator role.
- You must have installed and configured Workflow Automation.

The selected volumes must be protection relationship destinations. You cannot edit relationships when source volumes, load-sharing volumes, or volumes that are not the destination of a SnapMirror or SnapVault relationship are selected.

Steps

1. From the **Protection** tab of the **Volume / Health** details page, locate in the topology the protection relationship you want to edit and right-click it.
2. Select **Edit** from the menu.

   Alternatively, from the **Actions** menu, select **Relationship** > **Edit** to edit the relationship for which you are currently viewing the details.

   The **Edit Relationship** dialog box is displayed.

3. In the Edit Relationship dialog box, edit the maximum transfer rate, protection policy, or protection schedule, as needed.
4. Click **Apply**.

   The changes are applied to the selected relationships.

Create a SnapMirror policy to maximize transfer efficiency in Active IQ Unified Manager

You can create a SnapMirror policy to specify the SnapMirror transfer priority for protection relationships. SnapMirror policies enable you to maximize transfer efficiency from the source to the destination by assigning priorities so that lower-priority transfers

are scheduled to run after normal-priority transfers.

Before you begin

- You must have the Application Administrator or Storage Administrator role.
- You must have set up Workflow Automation.
- This task assumes that you have already completed the Destination Information area in the Configure Protection dialog box.

Steps

1. From the **SnapMirror** tab of the **Configure Protection** dialog box, click the **Create Policy** link in the **Relationship Settings** area.

   The Create SnapMirror Policy dialog box is displayed.

2. In the **Policy Name** field, type a name you want to give the policy.
3. In the **Transfer Priority** field, select the transfer priority you want to assign to the policy.
4. In the **Comment** field, enter an optional comment for the policy.
5. Click **Create**.

   The new policy is displayed in the SnapMirror Policy drop-down list.

Create a SnapVault policy to maximize transfer efficiency in Active IQ Unified Manager

You can create a new SnapVault policy to set the priority for a SnapVault transfer. You use policies to maximize the efficiency of transfers from the primary to the secondary in a protection relationship.

Before you begin

- You must have the Application Administrator or Storage Administrator role.
- You must have set up Workflow Automation.
- You must have already completed Destination Information area in the Configure Protection dialog box.

Steps

1. From the **SnapVault** tab of the **Configure Protection** dialog box, click the **Create Policy** link in the **Relationship Settings** area.

   The SnapVault tab is displayed.

2. In the **Policy Name** field, type the name that you want to give the policy.
3. In the **Transfer Priority** field, select the transfer priority that you want to assign to the policy.
4. **Optional:** In the **Comment** field, enter a comment for the policy.
5. In the **Replication Label** area, add or edit a replication label, as necessary.
6. Click **Create**.

   The new policy is displayed in the Create Policy drop-down list.

Abort a transfer from the Volume Relationships page in Active IQ Unified Manager

You can abort an active data protection transfer when you want to stop a SnapMirror replication that is in progress. You can also clear the restart checkpoint for transfers subsequent to the baseline transfer. You might abort a transfer when it conflicts with another operation, such as a volume move.

You cannot abort volumes relationships that are protected by the Consistency Group.

Before you begin

- You must have the Application Administrator or Storage Administrator role.
- You must have set up Workflow Automation.

The abort action does not display in the following instances:

- If RBAC settings do not allow this action: for example, if you have only operator privileges
- When the volume ID is unknown: for example, when you have an intercluster relationship and the destination cluster has not yet been discovered

You cannot clear the restart checkpoint for a baseline transfer.

Steps

1. To abort transfers for one or more protection relationships, from the **Volume Relationships** page, select one or more volumes and, on the toolbar, click **Abort**.

   The Abort Transfer dialog box is displayed.

2. If you want to clear the restart checkpoint for a transfer that is not a baseline transfer, select **Clear Checkpoints**.
3. Click **Continue**.

   The Abort Transfer dialog box is closed, and the status of the abort job displays at the top of the Volume Relationships page, along with a link to the job details.

4. **Optional:** Click the **View details** link to go to the **Job** details page for additional details and to view job progress.

Abort a transfer from the Volume/Health details page in Active IQ Unified Manager

You can abort an active data protection transfer when you want to stop a SnapMirror replication that is in progress. You can also clear the restart checkpoint for a transfer if it is not a baseline transfer. You might abort a transfer when it conflicts with another operation, such as a volume move.

You cannot abort volumes relationships that are protected by the Consistency Group.

Before you begin

- You must have the Application Administrator or Storage Administrator role.
- You must have set up Workflow Automation.

The abort action does not display in the following instances:

- If RBAC settings do not allow this action: for example, if you have only operator privileges
- When the volume ID is unknown: for example, when you have an intercluster relationship and the destination cluster has not yet been discovered

You cannot clear the restart checkpoint for a baseline transfer.

Steps

1. In the **Protection** tab of the **Volume / Health** details page, right-click the relationship in the topology view for the data transfer you want to abort and select **Abort**.

   The Abort Transfer dialog box is displayed.

2. If you want to clear the restart checkpoint for a transfer that is not a baseline transfer, select **Clear Checkpoints**.
3. Click **Continue**.

   The Abort Transfer dialog box is closed, and the status of the abort operation displays at the top of the Volume / Health details page along with a link to the job details.

4. **Optional:** Click the **View details** link to go to the **Job** details page for additional details and to view job progress.
5. Click each job task to view its details.
6. Click the Back arrow on your browser to return to the **Volume / Health** details page.

   The abort operation is finished when all job tasks successfully complete.

Quiesce a protection relationship from the Volume Relationships page in Active IQ Unified Manager

From the Volume Relationships page, you can quiesce a protection relationship to temporarily prevent data transfers from occurring. You might quiesce a relationship when you want to create a Snapshot copy of a SnapMirror destination volume that contains a database, and you want to ensure that its contents are stable during the Snapshot copy operation.

Before you begin

- You must have the Application Administrator or Storage Administrator role.
- You must have set up Workflow Automation.

The quiesce action does not display in the following instances:

- If RBAC settings do not allow this action; for example, if you have only operator privileges
- When the volume ID is unknown; for example, when you have an intercluster relationship and the destination cluster has not yet been discovered
- When you have not paired Workflow Automation and Unified Manager

Steps

1. To quiesce transfers for one or more protection relationships, from the **Volume Relationships** page, select one or more volumes and, on the toolbar, click **Quiesce**.

   The Quiesce dialog box is displayed.

2. Click **Continue**.

   The status of the quiesce job is displayed at the top of the Volume / Health details page, along with a link to the job details.

3. Click the **View details** link to go to the **Job** details page for additional details and job progress.
4. **Optional:** Click the **Back** arrow on your browser to return to the **Volume Relationships** page.

   The quiesce job is finished when all job tasks are successfully completed.

Quiesce a protection relationship from the Volume/ Health details page in Active IQ Unified Manager

You can quiesce a protection relationship to temporarily prevent data transfers from occurring. You might quiesce a relationship when you want to create a Snapshot copy of a SnapMirror destination volume that contains a database, and you want to ensure that its contents are stable during the Snapshot copy.

Before you begin

- You must have the Application Administrator or Storage Administrator role.
- You must have set up Workflow Automation.

The quiesce action does not display in the following instances:

- If RBAC settings do not allow this action, for example, if you have only operator privileges
- When the volume ID is unknown, for example, when you have an intercluster relationship and the destination cluster has not yet been discovered
- When you have not paired Workflow Automation and Unified Manager

Steps

1. In the **Protection** tab of the **Volume / Health** details page, right-click the relationship in the topology view for the protection relationship that you want to quiesce.
2. Select **Quiesce** from the menu.
3. Click **Yes** to continue.

   The status of the quiesce job is displayed at the top of the Volume / Health details page, along with a link to the job details.

4. Click the **View details** link to go to the **Job** details page for additional details and job progress.
5. **Optional:** Click the Back arrow on your browser to return to the **Volume / Health** details page.

   The quiesce job is finished when all job tasks are successfully completed.

Break a SnapMirror relationship from the Active IQ Unified Manager Volume Relationships page

You can break a protection relationship to stop data transfers between a source volume and a destination volume in a SnapMirror relationship. You might break a relationship when you want to migrate data, for disaster recovery, or for application testing. The destination volume is changed to a read/write volume. You cannot break a SnapVault relationship.

Before you begin

- You must have the Application Administrator or Storage Administrator role.
- You must have set up Workflow Automation.

Steps

1. From the **Volume Relationships** page, select one or more volumes with protection relationships for which you want to stop data transfers and, on the toolbar, click **Break**.

   The Break Relationship dialog box is displayed.

2. Click **Continue** to break the relationship.
3. In the **Volume Relationships** page, verify in the **Relationship State** column that the relationship is broken.

   The Relationship State column is hidden by default, so you might need to select it in the show/hide column list .

Remove a protection relationship from the Volume Relationships page in Active IQ Unified Manager

From the Volume Relationships page, you can remove a protection relationship to permanently delete an existing relationship between the selected source and destination: for example, when you want to create a relationship using a different destination. This operation removes all metadata and cannot be undone.

Before you begin

- You must have the Application Administrator or Storage Administrator role.
- You must have set up Workflow Automation.

Steps

1. From the **Volume Relationships** page, select one or more volumes with protection relationships you want to remove and, on the toolbar, click **Remove**.

   The Remove Relationship dialog box is displayed.

2. Click **Continue** to remove the relationship.

   The relationship is removed from the Volume Relationships page.

Resume scheduled transfers on a quiesced relationship from the Volume Relationships page in Active IQ Unified Manager

After you have quiesced a relationship to stop scheduled transfers from occurring, you can use **Resume** to re-enable scheduled transfers so that data on the source or primary volume is protected. Transfers resume from a checkpoint, if one exists, at the next scheduled transfer interval.

Before you begin

- You must have the Application Administrator or Storage Administrator role.
- You must have set up Workflow Automation.

You can select no more than 10 quiesced relationships on which to resume transfers.

Steps

1. From the Volume **Relationships** page, select one or more volumes with quiesced relationships, and, on the toolbar, click **Resume**.
2. In the **Resume** dialog box, click **Continue**.

   You are returned to the Volume Relationships page.
3. To view the related job tasks and to track their progress, click the job link that is displayed at the top of the **Volume Relationships** page.
4. Do one of the following:
   - If only one job is displayed, in the Job details page click **Refresh** to update the tasks list and task details associated with the protection configuration job and to determine when the job is complete.
   - If more than one job is displayed,
     i. In the Jobs page, click the job for which you want to view the details.
     ii. In the Job details page, click **Refresh** to update the tasks list and task details associated with the protection configuration job and to determine when the job is complete.
     After the jobs finish, data transfers are resumed at the next scheduled transfer interval.

Resume transfers on a quiesced relationship from the Volume / Health details page in Active IQ Unified Manager

After you have quiesced a relationship to stop scheduled transfers from occurring, you can use **Resume** on the Volume / Health details page to reenable scheduled transfers so that data on the source or primary volume is protected. Transfers resume from a checkpoint, if one exists, at the next scheduled transfer interval.

Before you begin

- You must have the Application Administrator or Storage Administrator role.
- You must have set up Workflow Automation.

Steps

1. In the **Protection** tab of the **Volume / Health** details page, right-click in the topology view a quiesced relationship that you want to resume.

Alternatively, select **Resume** from the **Actions** > **Relationship** menu.

2. In the **Resume** dialog box, click **Continue**.

   You are returned to the Volume / Health details page.

3. To view the related job tasks and to track their progress, click the job link that is displayed at the top of the **Volume / Health** details page.
4. In the **Job** details page, click **Refresh** to update the tasks list and task details associated with the protection configuration job and to determine when the job is complete.

   After the jobs are complete, data transfers are resumed at the next scheduled transfer interval.

Initialize or update protection relationships from the Volume Relationships page in Active IQ Unified Manager

From the Volume Relationships page, you can perform a first-time baseline transfer on a new protection relationship, or update a relationship if it is already initialized and you want to perform a manual, unscheduled incremental update to transfer immediately.

You cannot initialize or update volumes protected by Consistency Groups.

Before you begin

- You must have the Application Administrator or Storage Administrator role.
- You must have set up OnCommand Workflow Automation.

Steps

1. From the **Volume Relationships** page, right-click a volume and select one or more volumes with relationships that you want to update or initialize, and then, on the toolbar, click **Initialize/Update**.

   The **Initialize/Update** dialog box is displayed.

2. In the **Transfer Options** tab, select a transfer priority and the maximum transfer rate.
3. Click **Source Snapshot Copies**; then, in the **Snapshot Copy** column, click **Default**.

   The Select Source Snapshot Copy dialog box is displayed.

4. If you want to specify an existing Snapshot copy rather than transferring the default Snapshot copy, click **Existing Snapshot Copy** and select a Snapshot copy from the list.
5. Click **Submit**.

   You are returned to the **Initialize/Update** dialog box.

6. If you selected more than one source to initialize or update, click **Default** for the next source for which you want to specify an existing Snapshot copy.
7. Click **Submit** to begin the initialization or update job.

   The initialization or update job is started, you are returned to the Volume Relationships page, and a jobs link is displayed at the top of the page.

8. **Optional:** Click **View Jobs** on the **Health: All Volumes** view to track the status of each initialization or

update job.

A filtered list of jobs is displayed.

9. **Optional:** Click each job to see its details.
10. **Optional:** Click the **Back** arrow on your browser to return to the **Volume Relationships** page.

The initialization or update operation is finished when all tasks successfully finish.

Initialize or update protection relationships from the Volume/Health details page in Active IQ Unified Manager

You can perform a first-time baseline transfer on a new protection relationship, or update a relationship if it is already initialized and you want to perform a manual, unscheduled incremental update to transfer data immediately.

NOTE: You cannot initialize or update volumes protected by Consistency Groups.

Before you begin

- You must have the Application Administrator or Storage Administrator role.
- You must have set up OnCommand Workflow Automation.

Steps

1. From the **Protection** tab of the **Volume / Health** details page, locate in the topology the protection relationship that you want to initialize or update, and then right-click it.
2. Select **Initialize/Update** from the menu.

   Alternatively, from the **Actions** menu, select **Relationship** > **Initialize/Update** to initialize or update the relationship for which you are currently viewing the details.

   The Initialize/Update dialog box is displayed.

3. In the **Transfer Options** tab, select a transfer priority and the maximum transfer rate.
4. Click **Source Snapshot Copies**; then, in the **Snapshot Copy** column, click **Default**.

   The Select Source Snapshot Copy dialog box is displayed.

5. If you want to specify an existing Snapshot copy rather than transferring the default Snapshot copy, click **Existing Snapshot Copy** and select a Snapshot copy from the list.
6. Click **Submit**.

   You are returned to the Initialize/Update dialog box.

7. If you selected more than one source to initialize or update, click **Default** for the next read/write source for which you want to specify an existing Snapshot copy.

   You cannot select a different Snapshot copy for data protection volumes.

8. Click **Submit** to begin the initialization or update job.

   The initialization or update job is started, you are returned to the Volume / Health details page, and a jobs link is displayed at the top of the page.

9. **Optional:** Click **View Jobs** on the **Volume / Health** details page to track the status of each initialization or update job.

   A filtered list of jobs is displayed.

10. **Optional:** Click each job to see its details.
11. **Optional:** Click the Back arrow on your browser to return to the **Volume / Health** details page.

    The initialization or update operation is finished when all job tasks successfully complete.

Resynchronize protection relationships from the Volume Relationships page in Active IQ Unified Manager

From the Volume Relationships page, you can resynchronize a relationship either to recover from an event that disabled your source volume or when you want to change the current source to a different volume.

Before you begin

- You must have the Application Administrator or Storage Administrator role.
- You must have set up Workflow Automation.

Steps

1. From the **Volume Relationships** page, select one or more volumes with quiesced relationships and, from the toolbar, click **Resynchronize**.

   The Resynchronize dialog box is displayed.

2. In the **Resynchronization Options** tab, select a transfer priority and the maximum transfer rate.
3. Click **Source Snapshot Copies**; then, in the **Snapshot Copy** column, click **Default**.

   The Select Source Snapshot Copy dialog box is displayed.

4. If you want to specify an existing Snapshot copy rather than transferring the default Snapshot copy, click **Existing Snapshot Copy** and select a Snapshot copy from the list.
5. Click **Submit**.

   You are returned to the Resynchronize dialog box.

6. If you selected more than one source to resynchronize, click **Default** for the next source for which you want to specify an existing Snapshot copy.
7. Click **Submit** to begin the resynchronization job.

   The resynchronization job is started, you are returned to the Volume Relationships page, and a jobs link is displayed at the top of the page.

8. **Optional:** Click **View Jobs** on the **Volume Relationships** page to track the status of each resynchronization job.

   A filtered list of jobs is displayed.

9. **Optional:** Click the **Back** arrow on your browser to return to the **Volume Relationships** page.

The resynchronization operation is finished when all job tasks successfully finish.

Reverse protection relationships from the Volume Relationships page in Active IQ Unified Manager

When a disaster disables the source volume in your protection relationship, you can use the destination volume to serve data by converting it to a read/write volume while you repair or replace the source. When the source is again available to receive data, you can use the reverse resynchronization operation to establish the relationship in the reverse direction, synchronizing the data on the source with the data on the read/write destination.

Before you begin

- You must have the Application Administrator or Storage Administrator role.
- You must have set up Workflow Automation.
- The relationship must not be a SnapVault relationship.
- A protection relationship must already exist.
- The protection relationship must be broken.
- Both the source and destination must be online.
- The source must not be the destination of another data protection volume.
- When you perform this task, data on the source that is newer than the data on the common Snapshot copy is deleted.
- Policies and schedules created on reverse resynchronization relationships are the same as those on the original protection relationship.

  If policies and schedules do not exist, they are created.

Steps

1. From the **Volume Relationships** page, select one or more volumes with relationships that you want to reverse, and, on the toolbar, click **Reverse Resync**.

   The Reverse Resync dialog box is displayed.

2. Verify that the relationships displayed in the **Reverse Resync** dialog box are the ones for which you want to perform the reverse resynchronization operation, and then click **Submit**.

   The reverse resynchronization operation is started, you are returned to the Volume Relationships page, and a jobs link is displayed at the top of the page.

3. **Optional:** Click **View Jobs** on the **Volume Relationships** page to track the status of each reverse resynchronization job.

   A filtered list of jobs related to this operation is displayed.

4. **Optional:** Click the **Back** arrow on your browser to return to the **Volume Relationships** page.

   The reverse resynchronization operation is finished when all job tasks successfully complete.

Restore data using the Volume and Volume/Health details pages in Active IQ Unified Manager

You can restore overwritten or deleted files, directories, or an entire volume from a Snapshot copy by using the restore feature on the Volume and Volume / Health details pages.

Before you begin

You must have the Application Administrator or Storage Administrator role.

Note the following points:

- You can't restore a directory or a volume from a snapshot in any of the following scenarios:
  - The **Yes, preserve all snapshot copies** option is selected.
  - You are using an LDAP account to monitor the cluster.
  - You are logged in as a domain user.
  - Your account is configured for SAML authentication.
  - You are using the NDMP protocol.

    For more details, see the Knowledge Base article How to configure NDMP authentication in 'SVM-scope' mode.
- You can't restore NTFS file streams.
- The restore option is not available when:
  - The volume ID is unknown: for example, when you have an intercluster relationship and the destination cluster has not yet been discovered.
  - The volume is configured for SnapMirror Synchronous replication.

Steps

1. From the left navigation pane, go to **Storage > Volumes**.
2. Select the volume and click the **Restore** button. Alternately, click on the volume to go to **Volume / Health details > Actions > Restore**. The Restore dialog box is displayed. For information on this page, see Restore dialog box.
3. Select the volume and Snapshot copy from which you want to restore data, if different from the default.
4. Select the items you want to restore, for example, the source LUN.

   You can restore the entire volume, or you can specify folders and files you want to restore.
5. Select the location to which you want the selected items restored: either **Original Location** or **Alternate Existing Location**.
6. If you select an alternate existing location, do one of the following:
   - In the Restore Path text field, type the path of the location to which you want to restore the data and then click **Select Directory**.
   - Click **Browse** to launch the Browse Directories dialog box and complete the following steps:
     i. Select the destination cluster, storage VM (SVM), and volume to which you want to restore.

ii. In the Name table, select a directory name that has to be restored.

iii. Click **Select Directory**.

7. Click **Restore**.

   The restore process begins. A job is created at the backend to complete the restore process.

8. If you want to view the job progress, from the left navigation pane, navigate to **Protection > Jobs** to view the status of the restore job from the list of jobs.

> If a restore operation fails between Cloud Volumes ONTAP HA clusters with an NDMP error, you may need to add an explicit AWS route in the destination cluster so that the destination can communicate with the source system's cluster management LIF. You perform this configuration step by using the NetApp Console.

Learn about resource pools in Active IQ Unified Manager

Resource pools are groups of aggregates that are created by a storage administrator using Unified Manager to provide provisioning to partner applications for backup management.

You might pool your resources based on attributes such as performance, cost, physical location, or availability. By grouping related resources into a pool, you can treat the pool as a single unit for monitoring and provisioning. This simplifies the management of these resources and allows for a more flexible and efficient use of the storage.

During secondary storage provisioning, Unified Manager determines the most suitable aggregate in the resource pool for protection using the following criteria:

- The aggregate is a data aggregate (not a root aggregate) and it is ONLINE.
- The aggregate is on a destination cluster node whose ONTAP version is the same or greater than the source cluster major version.
- The aggregate has the largest available space of all the aggregates in the resource pool.
- After provisioning the destination volume, the aggregate space is within the nearly-full and nearly overcommitted threshold defined for the aggregate (global or local threshold, whichever is applicable).
- The number of FlexVol volumes on the destination node must not exceed the platform limit.

Create resource pools in Active IQ Unified Manager

You can use the Create Resource Pool dialog box to group aggregates for provisioning purposes.

Before you begin

You must have the Application Administrator or Storage Administrator role.

Steps

Resource pools can contain aggregates from different clusters, but the same aggregate cannot belong to different resource pools.

1. In the left navigation pane, click **Protection** > **Resource Pools**.

2. In the **Resource Pools page**, click **Create**.
3. Follow the instructions in the **Create Resource Pool** dialog box to provide a name and description and to add aggregates as members to the resource pool you want to create.

Edit resource pools in Active IQ Unified Manager

You can edit an existing resource pool when you want to change the resource pool name and the description.

Before you begin

You must have the Application Administrator or Storage Administrator role.

The **Edit** button is enabled only when one resource pool is selected. If more than one resource pool is selected, the **Edit** button is disabled.

Steps

1. In the left navigation pane, click **Protection** > **Resource Pools**.
2. Select one resource pool from the list.
3. Click **Edit**.

   The Edit Resource Pool window is displayed.

4. Edit the resource pool name and description as needed.
5. Click **Save**.

   The new name and description are displayed in the resource pool list.

View resource pool inventory in Active IQ Unified Manager

You can use the Resource Pools page to view the resource pool inventory and to monitor the remaining capacity for each resource pool.

Before you begin

You must have the Application Administrator or Storage Administrator role.

Step

1. In the left navigation pane, click **Protection** > **Resource Pools**.

   The resource pool inventory is displayed.

Add resource pool members in Active IQ Unified Manager

A resource pool consists of a number of member aggregates. You can add aggregates to existing resource pools to increase the amount of space available for secondary volume provisioning.

Before you begin

You must have the Application Administrator or Storage Administrator role.

You can add no more than 200 aggregates to a resource pool at one time. Aggregates shown in the

Aggregates dialog box do not belong to any other resource pool.

Steps

1. In the left navigation pane, click **Protection** > **Resource Pools**.
2. Select a resource pool from the **Resource Pools** list.

   The resource pool members are displayed in the area below the resource pool list.

3. In the resource pool member area, click **Add**.

   The Aggregates dialog box is displayed.

4. Select one or more aggregates.
5. Click **Add**.

   The dialog box is closed and the aggregates are displayed in the member list for the selected resource pool.

Remove aggregates from resource pools in Active IQ Unified Manager

You can remove aggregates from an existing resource pool: for example, when you want to use an aggregate for some other purpose.

Before you begin

You must have the Application Administrator or Storage Administrator role.

Resource pool members are displayed only when a resource pool is selected.

Steps

1. In the left navigation pane, click **Protection** > **Resource Pools**.
2. Select the resource pool from which you want to remove member aggregates.

   The list of member aggregates is displayed in the Members pane.

3. Select one or more aggregates.

   The **Remove** button is enabled.

4. Click **Remove.**

   A warning dialog box is displayed.

5. Click **Yes** to continue.

   The selected aggregates are removed from the Members pane.

Delete resource pools in Active IQ Unified Manager

You can delete resource pools when they are no longer needed. For example, you might want to redistribute the member aggregates from one resource pool to several other resource pools, making the original resource pool obsolete.

Before you begin

You must have the Application Administrator or Storage Administrator role.

The **Delete** button is enabled only when at least one resource pool is selected.

Steps

1. In the left navigation pane, click **Protection** > **Resource Pools**.
2. Select the resource pool you want to delete.
3. Click **Delete**.

   The resource pool is removed from the resource pool list and its aggregates are removed from the members list.

Monitor Storage VM Disaster Recovery protection relationships

Learn about storage VM disaster recovery relationships in Active IQ Unified Manager

Active IQ Unified Manager supports monitoring of storage VM disaster recovery relationships which provides disaster recovery at the granularity of a storage VM level. The storage VM disaster recovery enables the recovery of data present in the constituent volumes of the storage VM and the recovery of storage VM configuration.

A storage VM DR relationship is created from the source storage VM to the destination storage VM to provide asynchronous disaster recovery. You can select either to replicate all or subset of the storage VM configuration (excluding network and protocol configuration) along with the data volumes based on the cluster setup.

After the storage VM disaster recovery relationship is configured, when the source storage VM becomes unavailable due to either hardware failure or environmental disaster, the destination storage VM is started, that provides access to data with least disruption. Similarly, when the source storage VM becomes available, it is resynchronized with the destination storage VM and then, the source restarts to provide data. You can use SnapMirror commands to configure and manage storage VM disaster recovery relationship.

Monitor storage VMs using the Relationships page in Active IQ Unified Manager

You can monitor your storage VM disaster recovery relationships from the Relationships page in the PROTECTION section of the INVENTORY. By default, the Relationships page lists only the top-level relationships as the constituent relationships filter is applied.

Before you begin

You must have the Application Administrator or Storage Administrator role.

You use filters to view the storage VM disaster recovery relationships.

Steps

1. In the left navigation pane, click **PROTECTION** > **Relationships**.

   The page displays all type of relationships: volume, consistency group, and storage VM relationships.

2. Click **Filter**, and then select **Relationship Object Type** and **Storage VM** to view only storage VM disaster recovery relationships.

3. Click **Apply Filter**.

   You should clear the constituent relationships filter to view all the protection relationships.

   The page displays only storage VM disaster recovery relationships.

View protection relationships from storage VMs page in Active IQ Unified Manager

Using the Storage VMs page, you can view the status of existing storage VMs' disaster recovery relationships.

Before you begin

You must have the Application Administrator or Storage Administrator role.

You can also examine details of the protection relationships, including transfer and lag status, source, and destination detail. You can schedule reports or download existing reports in the format that you require. The **Show/Hide** button enables you to add the required columns to the reports as they are not displayed by default.

Steps

1. In the left navigation pane, click **STORAGE** > **Storage VMs**.
2. From the **VIEW** menu, select **Relationship** > **All Relationships**.

   The Relationship: All Relationships view is displayed with all the configured storage VMs.

View storage VMs based on protection status in Active IQ Unified Manager

You can use the Storage VMs page of the Inventory to view all the storage VMs in Active IQ Unified Manager and filter the storage VMs based on their protection status.

Before you begin

You must have the Application Administrator or Storage Administrator role.

A new column Protection Role is added to the storage VMs view that provides information on whether the storage VM is protected or unprotected.

If a source cluster is not added to Active IQ Unified Manager, then all the information related to that cluster is unavailable in the grids.

Steps

1. In the left navigation pane, click **STORAGE** > **Storage VMs**.
2. From the **VIEW** menu, select **Health** > **All Storage VMs**.

   The Health: All Storage VMs is displayed.

3. Click **Filter** to view one of the following storage VMs.

To view	Filter value
Protected storage VMs	**Protection Role** is **Protected**
Unprotected storage VMs	**Protection Role** is **Unprotected**

You cannot view both the protected and unprotected storage VMs at the same time. You will need to clear the existing filter to reapply a new filter option.

4. Click **Apply Filter**.

   The Unsaved view displays all the storage VMs that are either protected or unprotected by storage VM disaster recovery based on your filter selections.

Manage storage VM associations in Active IQ Unified Manager

Storage virtual machine (storage VM) associations are mappings from a source storage VM to a destination storage VM that are used by partner applications for resource selection and secondary volume provisioning.

Associations are created between a source storage VM and a destination storage VM, regardless of whether the destination storage VM is a secondary destination or a tertiary destination. You cannot use a secondary destination storage VM as a source to create an association with a tertiary destination storage VM.

As an Application Administrator or Storage Administrator, you can view the storage VM associations in your environment on the **Protection** > **Storage VM Associations** page.

You can associate SVMs in three ways:

- **Associate any storage VM**: You can create an association between any primary source storage VM and one or more destination SVMs. This means that all existing SVMs that currently require protection, as well as any SVMs that are created in the future, are associated with the specified destination SVMs. For example, you might want applications from several different sources at different locations to be backed up to one or more destination SVMs in one location.
- **Associate a particular storage VM**: You can create an association between a specific source storage VM and one or more specific destination SVMs. For example, if you are providing storage services to many clients whose data must be separate from one another, you can choose this option to associate a specific source storage VM to a specific destination storage VM that is assigned to only that client.
- **Associate with an external storage VM**: You can create an association between a source storage VM and an external flexible volume of a destination storage VM.

Create storage VM associations

The Create Storage Virtual Machine Associations wizard enables partner protection applications to associate a source storage VM with a destination storage VM for use with SnapMirror and SnapVault relationships. Partner applications use these associations at the time of initial provisioning of destination volumes to determine which resources to select.dd

Before you begin

- The storage VM you are associating must already exist.
- You must have the Application Administrator or Storage Administrator role.

For any source storage VM and relationship type, you can choose only one destination storage VM on each destination cluster.

Changing associations using the delete and create functions affects only future provisioning operations. It does not move existing destination volumes.

Steps

1. In the left navigation pane, click **Protection** > **Storage VM Associations**.
2. On the **Storage VM Associations** page, click **Create**.

   The **Create Storage Virtual Machine Associations** wizard is launched.

3. Select one of the following sources:
   - **Any**

     Choose this option when you want to create an association between any primary storage VM source to one or more destination storage VM. This means that all existing storage VM that currently require protection, as well as any storage VM that are created in the future, are associated with the specified destination storage VM. For example, you might want applications from several different sources at different locations backed up to one or more destination storage VM in one location.

   - **Single**

     Choose this option when you want to select a specific source storage VM associated with one or more destination storage VM. For example, if you are providing storage services to many clients whose data must be separate from one another, choose this option to associate a specific storage VM source to a specific storage VM destination that is assigned only to that client.

   - **None (External)**

     Choose this option when you want to create an association between a source storage VM and an external flexible volume of a destination storage VM.

4. Select one or both of the protection relationship types you want to create:
   - **SnapMirror**
   - **SnapVault**
5. Click **Next**.
6. Select one or more storage VM protection destination.
7. Click **Finish**.

Delete storage VM associations

You can delete storage VM associations for partner applications to remove the secondary provisioning relationship between source and destination storage VM; for example, you might do this when the destination storage VM is full and you want to create a new storage VM protection associations.

Before you begin

You must have the Application Administrator or Storage Administrator role.

The **Delete** button is disabled until at least one storage VM association is selected. Changing associations using the delete and create functions affects only future provisioning operations; it does not move existing destination volumes.

Steps

1. In the left navigation pane, click **Protection** > **Storage VM Associations**.
2. Select at least one storage VM association.

The **Delete** button is enabled.

3. Click **Delete.**

   A warning dialog box is displayed.

4. Click **Yes** to continue.

   The selected storage VM association is removed from the list.

SVM and resource pool requirements to support storage services in Active IQ Unified Manager

You can better ensure conformance in partner applications if you observe some SVM association and resource pool requirements that are specific to storage services: for example, when you associate SVM and create resource pools in Unified Manager to support a protection topology in a storage service provided by a partner application.

Some applications partner with the Unified Manager server to provide services that automatically configure and execute SnapMirror or SnapVault backup protection between source volumes and protection volumes in secondary or tertiary locations. To support these protection storage services, you must use Unified Manager to configure the necessary SVM associations and resource pools.

To support storage service single-hop or cascaded protection, including replication from a SnapMirror source or SnapVault primary volume to either destination SnapMirror or to SnapVault backup volumes that reside in secondary or tertiary locations, observe the following requirements:

- SVM associations must be configured between the SVM containing the SnapMirror source or SnapVault primary volume and any SVM on which either a secondary volume or a tertiary volume resides.
  - For example, to support a protection topology in which source volume Vol_A resides on SVM_1, and SnapMirror secondary destination volume Vol_B resides on SVM_2, and tertiary SnapVault backup volume Vol_ C resides on SVM_3, you must use the Unified Manager web UI to configure a SnapMirror association between SVM_1 and SVM_2 and a SnapVault backup association between SVM_1 and SVM_3.

    In this example, any SnapMirror association or SnapVault backup association between SVM_2 and SVM_3 is not necessary and is not used.

  - To support a protection topology in which both source volume Vol_A and SnapMirror destination volume Vol_B reside on SVM_1, you must configure a SnapMirror association between SVM_1 and SVM_1.
- The resource pools must include cluster aggregate resources available to the associated SVMs.

  You configure resource pools through the Unified Manager web UI and then assign through the partner application the storage service secondary target and tertiary target nodes.

Learn about jobs in Active IQ Unified Manager

A job is a series of tasks that you can monitor using Unified Manager. Viewing jobs and their associated tasks enables you to determine if they have completed successfully.

Jobs are initiated when you create SnapMirror and SnapVault relationships, when you perform any relationship

operation (break, edit, quiesce, remove, resume, resynchronize, and reverse resync), when you perform data restoration tasks, when you log in to a cluster, and so on.

When you initiate a job, you can use the Jobs page and the Job details page to monitor the job and the progress of the associated job tasks.

Monitor jobs in Active IQ Unified Manager

You can use the Jobs page to monitor job status and to view job properties such as storage service type, state, submitted time, and completed time to determine whether or not a job has successfully completed.

Before you begin

You must have the Application Administrator or Storage Administrator role.

Steps

1. In the left navigation pane, click **Protection** > **Jobs**.

   The Jobs page is displayed.

2. View the **State** column to determine the status of those jobs currently running.
3. Click on a job name to view details about that particular job.

   The Job details page is displayed.

View job details in Active IQ Unified Manager

After you start a job, you can track its progress from the Job details page and monitor the associated tasks for possible errors.

Before you begin

You must have the Application Administrator or Storage Administrator role.

Steps

1. In the left navigation pane, click **Protection** > **Jobs**.
2. In the Jobs page, click a job name in the **Name** column to display the list of tasks associated with the job.
3. Click on a task to display additional information in the **Task Details** pane and the **Task Messages** pane to the right of the task list.

Abort jobs in Active IQ Unified Manager

You can use the Jobs page to abort a job if it is taking too long to finish, is encountering too many errors, or is no longer needed. You can abort a job only if its status and type allow it. You can abort any running job.

Before you begin

You must have the Application Administrator or Storage Administrator role.

Steps

1. In the left navigation pane, click **Protection** > **Jobs**.
2. From the list of jobs, select one job, and then click **Abort**.
3. At the confirmation prompt, click **Yes** to abort the selected job.

Retry a failed protection job in Active IQ Unified Manager

After you have taken measures to fix a failed protection job, you can use **Retry** to run the job again. Retrying a job creates a new job using the original job ID.

Before you begin

You must have the Application Administrator or Storage Administrator role.

You can retry only one failed job at a time. Selecting more than one job disables the **Retry** button. Only jobs of the type Protection Configuration and Protection Relationship Operation can be retried.

Steps

1. In the left navigation pane, click **Protection** > **Jobs**.
2. From the list of jobs, select a single failed Protection Configuration or Protection Relationship Operation type job.

   The **Retry** button is enabled.

3. Click **Retry**.

   The job is restarted.

Description of Protection relationships windows and dialog boxes

Learn about the resource pools page in Active IQ Unified Manager

The Resource Pools page displays existing resource pools and their members, and enables you to create, monitor, and manage resource pools for provisioning purposes.

Command buttons

The command buttons enable you to perform the following tasks:

- **Create**

  Launches the Create Resource Pool dialog box, which you can use to create resource pools.

- **Edit**

  Enables you to edit the name and description of the resource pools that you create.

- **Delete**

  Enables you to delete one or more resource pools.

Resource Pools list

The Resource Pools list displays (in tabular format) the properties of existing resource pools.

- **Resource Pool**

  Displays the name of the resource pool.

- **Description**

  Describes the resource pool.

- **SnapLock Type**

  Displays the SnapLock type that is being used by the aggregates in the resource pool. Valid values for SnapLock type are Compliance, Enterprise, and Non-SnapLock. A resource pool can contain aggregates of only one SnapLock type.

- **Total Capacity**

  Displays the total capacity (in MB, GB, and so on) of the resource pool.

- **Used Capacity**

  Displays the amount of space (in MB, GB, and so on) that is used in the resource pool.

- **Available Capacity**

  Displays the amount of space (in MB, GB, and so on) that is available in the resource pool.

- **Used %**

  Displays the percentage of space that is used in the resource pool.

Members list command buttons

The Members list command buttons enable you to perform the following tasks:

- **Add**

  Enables you to add members to the resource pool.

- **Delete**

  Enables you to delete one or more members from the resource pool.

Members list

The Members list displays (in tabular format) the resource pool members and their properties when a resource pool is selected.

- **Status**

  Displays the current status of the member aggregate. The status can be Critical (), Error (), Warning (), or Normal ().

- **Aggregate Name**

  Displays the name of the member aggregate.

- **State**

  Displays the current state of the aggregate, which can be one of the following:

  - Offline

    Read or write access is not allowed.

  - Online

    Read and write access to the volumes that are hosted on this aggregate is allowed.

  - Restricted

    Limited operations (such as parity reconstruction) are allowed, but data access is not allowed.

  - Creating

    The aggregate is being created.

  - Destroying

    The aggregate is being destroyed.

  - Failed

    The aggregate cannot be brought online.

  - Frozen

    The aggregate is (temporarily) not serving requests.

  - Inconsistent

    The aggregate has been marked corrupted; you should contact technical support.

  - Iron Restricted

    Diagnostic tools cannot be run on the aggregate.

  - Mounting

    The aggregate is in the process of mounting.

  - Partial

    At least one disk was found for the aggregate, but two or more disks are missing.

  - Quiescing

    The aggregate is being quiesced.

  - Quiesced

    The aggregate is quiesced.

- Reverted

  The revert of an aggregate is completed.

- Unmounted

  The aggregate has been unmounted.

- Unmounting

  The aggregate is being taken offline.

- Unknown

  The aggregate is discovered, but the aggregate information is not yet retrieved by the Unified Manager server.

By default, this column is hidden.

- **Cluster**

  Displays the name of the cluster to which the aggregate belongs.

- **Node**

  Displays the name of the node on which the aggregate resides.

- **Total Capacity**

  Displays the total capacity (in MB, GB, and so on) of the aggregate.

- **Used Capacity**

  Displays the amount of space (in MB, GB, and so on) that is used in the aggregate.

- **Available Capacity**

  Displays the amount of space (in MB, GB, and so on) that is available in the aggregate.

- **Used %**

  Displays the percentage of space that is used in the aggregate.

- **Disk Type**

  Displays the RAID configuration type, which can be one of the following:

  - RAID0: All the RAID groups are of type RAID0.
  - RAID4: All the RAID groups are of type RAID4.
  - RAID-DP: All the RAID groups are of type RAID-DP.
  - RAID-TEC: All the RAID groups are of type RAID-TEC.
  - Mixed RAID: The aggregate contains RAID groups of different RAID types (RAID0, RAID4, RAID-DP, and RAID-TEC).
    By default, this column is hidden.

Use the Create Resource Pool dialog box in Active IQ Unified Manager

You can use the Create Resource Pool dialog box to name and describe a new resource pool and to add aggregates to and delete aggregates from that resource pool.

Resource Pool Name

The text boxes enable you to add the following information to create a resource pool:

Enables you to specify a resource pool name.

Description

Enables you to describe a resource pool.

Members

Displays the members of the resource pool. You can also add and delete members.

Command buttons

The command buttons enable you to perform the following tasks:

- **Add**

  Opens the Aggregates dialog box so that you can add aggregates from a specific cluster to the resource pool. You can add aggregates from different clusters, but the same aggregates cannot be added to more than one resource pool.

- **Remove**

  Enables you to remove selected aggregates from the resource pool.

- **Create**

  Creates the resource pool. This button is not enabled until information has been entered in the Resource Pool Name or Description fields.

- **Cancel**

  Discards the changes and closes the Create Resource Pool dialog box.

Use the Edit Resource Pool dialog box in Active IQ Unified Manager

You can use the Edit Resource Pool dialog box to change the name and description of an existing resource pool. For example, if the original name and description is inaccurate or incorrect, you can change them so they are more precise.

Text boxes

The text boxes enable you to change the following information for the selected resource pool:

- **Resource Pool Name**

  Enables you to enter a new name.

- **Description**

  Enables you to enter a new description.

Command buttons

The command buttons enable you to perform the following tasks:

- **Save**

  Saves the changes to the resource pool name and description.

- **Cancel**

  Discards the changes and closes the Edit Resource Pool dialog box.

Use the aggregates dialog box in Active IQ Unified Manager

You can use the Aggregates dialog box to select the aggregates that you want to add to your resource pool.

Command buttons

The command buttons enable you to perform the following tasks:

- **Add**

  Adds the selected aggregates to the resource pool. The Add button is not enabled until at least one aggregate is selected.

- **Cancel**

  Discards the changes, and closes the Aggregates dialog box.

Aggregates list

The Aggregates list displays (in tabular format) the names and properties of monitored aggregates.

- **Status**

  Displays the current status of a volume. The status can be Critical (), Error (), Warning (), or Normal ().

  You can move the pointer over the status to view more information about the event or events generated for the volume.

- **Aggregate Name**

  Displays the name of the aggregate.

- **State**

  Displays the current state of the aggregate, which can be one of the following:

  - Offline

Read or write access is not allowed.

- Restricted

  Limited operations (such as parity reconstruction) are allowed, but data access is not allowed.

- Online

  Read and write access to the volumes that are hosted on this aggregate is allowed.

- Creating

  The aggregate is being created.

- Destroying

  The aggregate is being destroyed.

- Failed

  The aggregate cannot be brought online.

- Frozen

  The aggregate is (temporarily) not serving requests.

- Inconsistent

  The aggregate has been marked corrupted; you should contact technical support.

- Iron Restricted

  Diagnostic tools cannot be run on the aggregate.

- Mounting

  The aggregate is in the process of mounting.

- Partial

  At least one disk was found for the aggregate, but two or more disks are missing.

- Quiescing

  The aggregate is being quiesced.

- Quiesced

  The aggregate is quiesced.

- Reverted

  The revert of an aggregate is completed.

- Unmounted

The aggregate is offline.

- Unmounting

  The aggregate is being taken offline.

- Unknown

  The aggregate is discovered, but the aggregate information is not yet retrieved by the Unified Manager server.

- **Cluster**

  Displays the name of the cluster on which the aggregate resides.

- **Node**

  Displays the name of the storage controller that contains the aggregate.

- **Total Capacity**

  Displays the total data size (in MB, GB, and so on) of the aggregate. By default, this column is hidden.

- **Committed Capacity**

  Displays the total space (in MB, GB, and so on) that is committed for all the volumes in the aggregate. By default, this column is hidden.

- **Used Capacity**

  Displays the amount of space (in MB, GB, and so on) that is used in the aggregate.

- **Available Capacity**

  Displays the amount of space (in MB, GB, and so on) that is available for data in the aggregate. By default, this column is hidden.

- **Available %**

  Displays the percentage of space that is available for data in the aggregate. By default, this column is hidden.

- **Used %**

  Displays the percentage of space that is used by data in the aggregate.

- **RAID Type**

  Displays the RAID type of the selected volume. The RAID type can be RAID0, RAID4, RAID-DP, RAID-TEC, or Mixed RAID.

Learn about the jobs page in Active IQ Unified Manager

The Jobs page enables you to view the current status and other information about all partner application protection jobs that are currently running, as well as jobs that have

completed. You can use this information to see which jobs are still running and whether a job has succeeded or failed.

Command buttons

The command buttons enable you to perform the following tasks:

- **Abort**

  Aborts the selected job. This option is available only if the selected job is running.

- **Retry**

  Restarts a failed job of type Protection Configuration or Protection Relationship Operation. You can retry only one failed job at a time. If more than one failed job is selected, the **Retry** button is disabled. You cannot retry failed storage service jobs.

- **Refresh**

  Refreshes the list of jobs and the information associated with them.

Jobs list

The Jobs list displays, in tabular format, a list of the jobs that are in progress. By default, the list displays only the jobs generated within the past week. You can use column sorting and filtering to customize which jobs are displayed.

- **Status**

  Displays the current status of a job. The status can be Error (❗) or Normal (✅).

- **Job Id**

  Displays the identification number of the job. By default, this column is hidden.

  The job identification number is unique and is assigned by the server when it starts the job. You can search for a particular job by entering the job identification number in the text box provided by the column filter.

- **Name**

  Displays the name of the job.

- **Type**

  Displays the job type. The job types are as follows:

  - **Cluster Acquisition**

    A Workflow Automation job is rediscovering a cluster.

  - **Protection Configuration**

    A protection job is initiating Workflow Automation workflows, such as cron schedules, SnapMirror policy creation, and so on.

- **Protection Relationship Operation**

  A protection job is running SnapMirror operations.

- **Protection Workflow Chain**

  A Workflow Automation job is executing multiple workflows.

- **Restore**

  A restore job is running.

- **Cleanup**

  The job is cleaning up storage service member artifacts that are no longer needed for restore purposes.

- **Conform**

  The job is checking the configuration of storage service members to ensure that they conform.

- **Destroy**

  The job is destroying a storage service.

- **Import**

  The job is importing unmanaged storage objects into an existing storage service.

- **Modify**

  The job is modifying attributes of an existing storage service.

- **Subscribe**

  The job is subscribing members to a storage service.

- **Unsubscribe**

  The job is unsubscribing members from a storage service.

- **Update**

  A protection update job is running.

- **WFA Configuration**

  A Workflow Automation job is pushing cluster credentials and synchronizing database caches.

- **State**

Displays the running state of the job. State options are as follows:

- **Aborted**

  The job has been aborted.

- **Aborting**

  The job is in the process of aborting.

- **Completed**

  The job has finished.

- **Running**

  The job is running.

- **Submitted Time**

  Displays the time the job was submitted.

- **Duration**

  Displays the amount of time the job took to complete. This column is displayed by default.

- **Completed Time**

  Displays the time the job finished. By default, this column is hidden.

Learn about the job details page in Active IQ Unified Manager

The Job details page enables you to view status and other information about specific protection job tasks that are running, that are queued, or that have completed. You can use this information to monitor protection job progress and to troubleshoot job failures.

Job summary

The job summary displays the following information:

- Job ID
- Type
- State
- Submitted Time
- Completed Time
- Duration

Command buttons

The command buttons enable you to perform the following tasks:

- **Refresh**

  Refreshes the task list and the properties associated with each task.

- **View Jobs**

  Returns you to the Jobs page.

Job tasks list

The Job tasks list displays in a table all the tasks associated with a specific job and the properties related to each task.

- **Started Time**

  Displays the day and time the task started. By default, the most recent tasks are displayed at the top of the column and older tasks are displayed at the bottom.

- **Type**

  Displays the type of task.

- **State**

  The state of a particular task:

  - **Completed**

    The task has finished.

  - **Queued**

    The task is about to run.

  - **Running**

    The task is running.

  - **Waiting**

    A job has been submitted and some associated tasks are waiting to be queued and executed.

- **Status**

  Displays the task status:

  - **Error (❗)**

    The task failed.

  - **Normal (✅)**

    The task succeeded.

  - **Skipped (↪)**

    A task failed, resulting in subsequent tasks being skipped.

- **Duration**

  Displays the elapsed time since the task began.

- **Completed Time**

  Displays the time the task completed. By default, this column is hidden.

- **Task ID**

  Displays the GUID that identifies an individual task for a job. The column can be sorted and filtered. By default, this column is hidden.

- **Dependency order**

  Displays an integer representing the sequence of tasks in a graph, with zero being assigned to the first task. By default, this column is hidden.

- **Task Details pane**

  Displays additional information about each job task, including the task name, task description, and, if the task failed, a reason for the failure.

- **Task Messages pane**

  Displays messages specific to the selected task. Messages might include a reason for the error and suggestions for resolving it. Not all tasks display task messages.

Use the Advanced Secondary Settings dialog box in Active IQ Unified Manager

You can use the Advanced Secondary Settings dialog box to enable version-flexible replication, multiple copy backup, and space-related settings on a secondary volume. You might use the Advanced Secondary Settings dialog box when you want to change enable or disable the current settings.

Space-related settings maximize the amount of data being stored, including the following: deduplication, data compression, autogrow, and space guarantee.

The dialog box includes the following fields:

- **Enable Version-Flexible Replication**

  Enables SnapMirror with version-flexible replication. Version-flexible replication enables SnapMirror protection of a source volume even if the destination volume is running under an earlier version of ONTAP than that of the source volume.

  - Enable Backup

    If version-flexible replication is enabled, also enables multiple Snapshot copies of the SnapMirror source data to be transferred to and retained at the SnapMirror destination.

- **Enable Deduplication**

  Enables deduplication on the secondary volume in a SnapVault relationship so that duplicate data blocks are eliminated to achieve space savings. You might use deduplication when space savings are at least 10 percent and when data overwrite rate is not rapid. Deduplication is often used for virtualized environments, file shares, and backup data. This setting is disabled by default. When enabled, this operation is initiated after each transfer.

  - Enable Compression

    Enables transparent data compression. You might use compression when space savings are at least 10 percent, when the potential overhead is acceptable, and when there are sufficient system resources

for compression to complete during nonpeak hours. In a SnapVault relationship, this setting is disabled by default. Compression is available only when deduplication is selected.

- Compress Inline

  Enables immediate space savings by compressing data before writing data to disk. You might use inline compression when your system has no more than 50 percent utilization during peak hours, and when the system can accommodate new writes and additional CPU during peak hours. This setting is available only when “Enable Compression” is selected.

- **Enable Autogrow**

  Enables you to automatically grow the destination volume when the free space percentage is below the specified threshold, as long as space is available on the associated aggregate.

- **Maximum Size**

  Sets the maximum percentage to which a volume can grow. The default is 20 percent greater than the source volume size. A volume does not grow automatically if the current size is greater than or equal to the maximum autogrow percentage. This field is enabled only when the autogrow setting is enabled.

- **Increment Size**

  Specifies the percentage increment by which the volume automatically grows before reaching the maximum percentage of the source volume.

- **Space Guarantee**

  Ensures that enough space is allocated on the secondary volume so that data transfers always succeed. The space guarantee setting can be one of the following:

  - File
  - Volume
  - None

    For example, you might have a 200 GB volume that contains files totaling 50 GB; however, those files hold only 10 GB of data. Volume guarantee allocates 200 GB to the destination volume, regardless of contents on the source. File guarantee allocates 50 GB to ensure that enough space is reserved for files on the source; selecting None in this scenario means that only 10 GB is allocated on the destination for the actual space used by file data on the source.

  The space guarantee is set to Volume by default.

Command buttons

The command buttons enable you to perform the following tasks:

- **Apply**

  Saves the selected efficiency settings and applies them when you click **Apply** in the Configure Protection dialog box.

- **Cancel**

  Discards your selections and closes the Advanced Destination Settings dialog box.

Use the Advanced Destination Settings dialog box in Active IQ Unified Manager

You can use the Advanced Destination Settings dialog box to enable space guarantee settings on a destination volume. You might select advanced settings when space guarantee is disabled on the source, but you want it enabled on the destination. The settings for deduplication, compression, and autogrow in a SnapMirror relationship are inherited from the source volume and cannot be changed.

Space Guarantee

Ensures that enough space is allocated on the destination volume so that data transfers always succeed. The space guarantee setting can be one of the following:

- File
- Volume
- None

For example, you might have a 200-GB volume that contains files totaling 50 GB; however, those files hold only 10 GB of data. Volume guarantee allocates 200 GB to the destination volume, regardless of contents on the source. File guarantee allocates 50 GB to ensure that enough space is reserved for source files on the destination; selecting **None** in this scenario means that only 10 GB is allocated on the destination for the actual space used by file data on the source.

The space guarantee is set to Volume by default.

Use the Restore dialog box in Active IQ Unified Manager

You can use the Restore dialog box to restore data to a volume from a specific Snapshot copy.

Restore from

The Restore from area enables you to specify from where you want to restore data.

- **Volume**

  Specifies the volume from which you want to restore data. By default, the volume on which you initiated the restore action is selected. You can select a different volume from the drop-down list that contains all the volumes with protection relationships to the volume on which you initiated the restore action.

- **Snapshot copy**

  Specifies which Snapshot copy you want to use to restore data. By default, the most recent Snapshot copy is selected. You can also select a different Snapshot copy from the drop-down list. The Snapshot copy list changes according to which volume is selected.

- **List maximum of 995 files and directories**

  By default a maximum of 995 objects are shown in the list. You can deselect this checkbox if you want to view all objects within the selected volume. This operation may take some time if the number of items is very large.

Select items to restore

The Select items to restore area enables you to select either the entire volume or specific files and folders you want to restore. You can select a maximum of 10 files, folders, or a combination of both. When the maximum number of items is selected, the item selection check boxes are disabled.

- **Path field**

  Displays the path to the data you want to restore. You can either navigate to the folder and files you want to restore, or you can type the path. This field is empty until you select or type a path. Clicking  after you have chosen a path moves you up one level in the directory structure.

- **Folders and files list**

  Displays the contents of the path you entered. By default, the root folder is initially displayed. Clicking a folder name displays the contents of the folder.

You can select items to restore as follows:

- When you enter the path with a particular file name specified in the path field, the specified file is displayed in the Folders and Files.
- When you enter a path without specifying a particular file, the contents of the folder are displayed in the Folders and Files list, and you can select up to 10 files, folders, or a combination of both to restore.

  If a folder contains more than 995 items, a message displays to indicate there are too many items to display, and if you proceed with the operation all items in the specified folder are restored. You can deselect the “List maximum of 995 files and directories” checkbox if you want to view all objects within the selected volume.



You cannot restore NTFS file streams.

Restore to

The Restore to area enables you to specify where you want to restore the data.

- **Original Location in Volume_Name**

  Restores the selected data to the directory on the source from which the data was originally backed up.

- **Alternate Location**

  Restores the selected data to a new location:

  - Restore Path

    Specifies an alternate path for restoring the selected data. The path must already exist. You can use the **Browse** button to navigate to the location where you want the data restored, or you can enter the path manually using the format `cluster://svm/volume/path`.

  - Preserve directory hierarchy

    When checked, preserves the structure of the original file or directory. For example, if the source is /A/B/C/myFile.txt and the destination is /X/Y/Z, Unified Manager restores the data using the following directory structure on the destination: /X/Y/Z/A/B/C/myFile.txt.

Command buttons

The command buttons enable you to perform the following tasks:

- **Cancel**

  Discards your selections and closes the Restore dialog box.

- **Restore**

  Applies your selections and begins the restore process.

Use the Browse Directories dialog box in Active IQ Unified Manager

You can use the Browse Directories dialog box when you want to restore data to a directory on a cluster and SVM that is different from the original source. The original source cluster and volume are selected by default.

The Browse Directories dialog box enables you to select the cluster, SVM, volume, and directory path to which you want data restored.

- **Cluster**

  Lists the available cluster destinations to which you can restore. By default, the cluster of the original source volume is selected.

- **SVM drop-down list**

  Lists the available SVM available for the selected cluster. By default, the SVM of the original source volume is selected.

- **Volume**

  Lists all of the read/write volumes in a selected SVM. You can filter the volumes by name and by space available. The volume with the most space is listed first, and so on, in descending order. By default, the original source volume is selected.

- **File path text box**

  Enables you to type the file path to which you want data restored. The path you enter must already exist.

- **Name**

  Displays the names of the available folders for the selected volume. Clicking a folder in the Name list displays the subfolders, if any exist. Files contained in the folders are not displayed. Clicking  after you have selected a folder moves you up one level in the directory structure.

Command buttons

The command buttons enable you to perform the following tasks:

- **Select Directory**

  Applies your selections and closes the Browse Directories dialog box. If no directory is selected, this button is disabled.

- **Cancel**

  Discards your selections and closes the Browse Directories dialog box.

Use the Configure Protection dialog box in Active IQ Unified Manager

You can use the Configure Protection dialog box to create SnapMirror and SnapVault relationships for all read, write, and data protection volumes on clusters to ensure that the data on a source volume or primary volume is replicated.

Source tab

- **Topology view**

  Displays a visual representation of the relationship that you are creating. The source in the topology is highlighted by default.

- **Source Information**

  Displays details about the selected source volumes, including the following information:

  - Source cluster name
  - Source SVM name
  - Cumulative volume total size

    Displays the total size of all the source volumes that are selected.

  - Cumulative volume used size

    Displays the cumulative volume used size for all the selected source volumes.

  - Source volume

    Displays the following information in a table:

    - Source Volume

      Displays the names of the selected source volumes.

    - Type

      Displays the volume type.

    - SnapLock Type

      Displays the SnapLock type of the volume. The options are Compliance, Enterprise, and Non-SnapLock.

    - Snapshot Copy

      Displays the Snapshot copy that is used for the baseline transfer. If the source volume is read/write, the value of Default in the Snapshot copy column indicates that a new Snapshot copy is created by default, and is used for the baseline transfer. If the source volume is a data protection volume, the value of Default in the Snapshot copy column indicates that no new Snapshot copy is created, and

all existing Snapshot copies are transferred to the destination. Clicking the Snapshot copy value displays a list of Snapshot copies from which you can select an existing Snapshot copy to use for the baseline transfer. You cannot select a different default Snapshot copy if the source type is data protection.

SnapMirror tab

Enables you to specify a destination cluster, storage virtual machine (SVM), and aggregate for a protection relationship, as well as a naming convention for destinations while creating a SnapMirror relationship. You can also specify a SnapMirror policy and schedule.

- **Topology view**

  Displays a visual representation of the relationship that you are creating. The SnapMirror destination resource in the topology is highlighted by default.

- **Destination Information**

  Enables you to select the destination resources for a protection relationship:

  - Advanced link

    Launches the Advanced Destination Settings dialog box when you are creating a SnapMirror relationship.

  - Cluster

    Lists the clusters that are available as protection destination hosts. This field is required.

  - storage virtual machine (SVM)

    Lists the SVMs that are available on the selected cluster. A cluster must be selected before the SVM list is populated. This field is required.

  - Aggregate

    Lists the aggregates that are available on the selected SVM. A cluster must be selected before the Aggregate list is populated. This field is required. The Aggregate list displays the following information:

    - Rank

      When multiple aggregates satisfy all the requirements for a destination, the rank indicates the priority in which the aggregate is listed, according to the following conditions:

      A. An aggregate that is located on a different node than the source volume node is preferred to enable fault domain separation.
      B. An aggregate on a node with fewer volumes is preferred to enable load balancing across nodes in a cluster.
      C. An aggregate that has more free space than other aggregates is preferred to enable capacity balancing.
      A rank of 1 means that the aggregate is the most preferred according to the three criteria.

    - Aggregate Name

      Name of the aggregate

        - Available Capacity
        - Amount of space that is available on the aggregate for data
        - Resource Pool

          Name of the resource pool to which the aggregate belongs

    - Naming Convention

      Specifies the default naming convention that is applied to the destination volume. You can accept the naming convention that is provided, or you can create a custom one. The naming convention can have the following attributes: %C, %M, %V, and %N, where %C is the cluster name, %M is the SVM name, %V is the source volume, and %N is the topology destination node name.

      The naming convention field is highlighted in red if your entry is invalid. Clicking the “Preview Name” link displays a preview of the naming convention that you entered, and the preview text updates dynamically as you type a naming convention in the text field. A suffix between 001 and 999 is appended to the destination name when the relationship is created, replacing the nnn that displays in the preview text, with 001 being assigned first, 002 assigned second, and so on.

- **Relationship Settings**

  Enables you to specify the maximum transfer rate, SnapMirror policy, and schedule that the protection relationship uses:

    - Max Transfer Rate

      Specifies the maximum rate at which data is transferred between clusters over the network. If you choose not to use a maximum transfer rate, the baseline transfer between relationships is unlimited.

    - SnapMirror Policy

      Specifies the ONTAP SnapMirror policy for the relationship. The default is DPDefault.

    - Create Policy

      Launches the Create SnapMirror Policy dialog box, which enables you to create and use a new SnapMirror policy.

    - SnapMirror Schedule

      Specifies the ONTAP SnapMirror policy for the relationship. Available schedules include None, 5min, 8hour, daily, hourly, and weekly. The default is None, indicating that no schedule is associated with the relationship. Relationships without schedules have no lag status values unless they belong to a storage service.

    - Create Schedule

      Launches the Create Schedule dialog box, which enables you to create a new SnapMirror schedule.

SnapVault tab

Enables you to specify a secondary cluster, SVM, and aggregate for a protection relationship, as well as a naming convention for secondary volumes while creating a SnapVault relationship. You can also specify a SnapVault policy and schedule.

- **Topology view**

  Displays a visual representation of the relationship that you are creating. The SnapVault secondary resource in the topology is highlighted by default.

- **Secondary Information**

  Enables you to select the secondary resources for a protection relationship:

  - Advanced link

    Launches the Advanced Secondary Settings dialog box.

  - Cluster

    Lists the clusters that are available as secondary protection hosts. This field is required.

  - storage virtual machine (SVM)

    Lists the SVMs that are available on the selected cluster. A cluster must be selected before the SVM list is populated. This field is required.

  - Aggregate

    Lists the aggregates that are available on the selected SVM. A cluster must be selected before the Aggregate list is populated. This field is required. The Aggregate list displays the following information:

    - Rank

      When multiple aggregates satisfy all the requirements for a destination, the rank indicates the priority in which the aggregate is listed, according to the following conditions:

      A. An aggregate that is located on a different node than the primary volume node is preferred to enable fault domain separation.
      B. An aggregate on a node with fewer volumes is preferred to enable load balancing across nodes in a cluster.
      C. An aggregate that has more free space than other aggregates is preferred to enable capacity balancing.
      A rank of 1 means that the aggregate is the most preferred according to the three criteria.

    - Aggregate Name

      Name of the aggregate

    - Available Capacity
    - Amount of space that is available on the aggregate for data
    - Resource Pool

      Name of the resource pool to which the aggregate belongs

  - Naming Convention

    Specifies the default naming convention that is applied to the secondary volume. You can accept the naming convention that is provided, or you can create a custom one. The naming convention can have

the following attributes: %C, %M, %V, and %N, where %C is the cluster name, %M is the SVM name, %V is the source volume, and %N is the topology secondary node name.

The naming convention field is highlighted in red if your entry is invalid. Clicking the “Preview Name” link displays a preview of the naming convention that you entered, and the preview text updates dynamically as you type a naming convention in the text field. If you type an invalid value, the invalid information displays as red question marks in the preview area. A suffix between 001 and 999 is appended to the secondary name when the relationship is created, replacing the nnn that displays in the preview text, with 001 being assigned first, 002 assigned second, and so on.

- **Relationship Settings**

  Enables you to specify the maximum transfer rate, SnapVault policy, and SnapVault schedule that the protection relationship uses:

  - Max Transfer Rate

    Specifies the maximum rate at which data is transferred between clusters over the network. If you choose not to use a maximum transfer rate, the baseline transfer between relationships is unlimited.

  - SnapVault Policy

    Specifies the ONTAP SnapVault policy for the relationship. The default is XDPDefault.

  - Create Policy

    Launches the Create SnapVault Policy dialog box, which enables you to create and use a new SnapVault policy.

  - SnapVault Schedule

    Specifies the ONTAP SnapVault schedule for the relationship. Available schedules include None, 5min, 8hour, daily, hourly, and weekly. The default is None, indicating that no schedule is associated with the relationship. Relationships without schedules have no lag status values unless they belong to a storage service.

  - Create Schedule

    Launches the Create Schedule dialog box, which enables you to create a SnapVault schedule.

Command buttons

The command buttons enable you to perform the following tasks:

- **Cancel**

  Discards your selections, and closes the Configure Protection dialog box.

- **Apply**

  Applies your selections, and begins the protection process.

Use the Create Schedule dialog box in Active IQ Unified Manager

The Create Schedule dialog box enables you to create a basic or advanced protection

schedule for SnapMirror and SnapVault relationship transfers. You might create a new schedule to increase the frequency of data transfers due to frequent data updates, or you might create a less frequent schedule when data changes infrequently.

Schedules cannot be configured for SnapMirror Synchronous relationships.

- **Destination Cluster**

  The name of the cluster you selected in the SnapVault tab or SnapMirror tab of the Configure Protection dialog box.

- **Schedule Name**

  The name you provide for the schedule. Schedule names can consist of the characters A through Z, a through z, 0 through 9, as well as any of the following special characters: ! @ # $ % ^ & * ( ) _ -. Schedule names may not include the following characters: < >.

- **Basic or Advanced**

  The schedule mode you want to use.

  Basic mode includes the following elements:

  - Repeat

    How often a scheduled transfer occurs. Choices include hourly, daily, and weekly.

  - Day

    When a repeat of weekly is selected, the day of the week a transfer occurs.

  - Time

    When Daily or Weekly is selected, the time of day a transfer occurs.

  Advanced mode includes the following elements:

  - Months

    A comma-separated numerical list representing the months of the year. Valid values are 0 through 11, with zero representing January, and so on. This element is optional. Leaving the field blank implies that transfers occur every month.

  - Days

    A comma-separated numerical list representing the day of the month. Valid values are 1 through 31. This element is optional. Leaving the field blank implies that a transfer occurs every day of the month.

  - Weekdays

    A comma-separated numerical list representing the days of the week. Valid values are 0 through 6, with 0 representing Sunday, and so on. This element is optional. Leaving the field blank implies that a transfer occurs every day of the week. If a day of the week is specified but a day of the month is not specified, a transfer occurs only on the specified day of the week and not every day.

  - Hours

A comma-separated numerical list representing the number of hours in a day. Valid values are 0 through 23, with 0 representing midnight. This element is optional.

- Minutes

  A comma-separated numerical list representing the minutes in an hour. Valid values are 0 through 59. This element is required.

Use the Create SnapMirror Policy dialog box in Active IQ Unified Manager

The Create SnapMirror Policy dialog box enables you to create a policy to set the priority for SnapMirror transfers. You use policies to maximize the efficiency of transfers from the source to the destination.

- **Destination Cluster**

  The name of the cluster you selected in the SnapMirror tab of the Configure Protection dialog box.

- **Destination SVM**

  The name of the SVM you selected in the SnapMirror tab of the Configure Protection dialog box.

- **Policy Name**

  The name you provide for the new policy. Policy names can consist of the characters A through Z, a through z, 0 through 9, period (.), hyphen (-), and underscore (_).

- **Transfer Priority**

  The priority at which a transfer runs for asynchronous operations. You can select either Normal or Low. Transfer relationships with policies that specify a normal transfer priority run before those with policies that specify a low transfer priority.

- **Comment**

  An optional field in which you can add comments about the policy.

- **Transfer Restart**

  Indicates what restart action to take when a transfer is interrupted by an abort operation or any type of failure, such as a network outage. You can select one of the following:

  - Always

    Specifies that a new Snapshot copy is created before restarting a transfer, then, if one exists, the transfer is restarted from a checkpoint, followed by an incremental transfer from the newly created Snapshot copy.

  - Never

    Specifies that interrupted transfers are never restarted.

Command buttons

The command buttons enable you to perform the following tasks:

- **Cancel**

  Discards the selections and closes the Configure Protection dialog box.

- **Apply**

  Applies your selections and begins the protection process.

Use the Create SnapVault Policy dialog box in Active IQ Unified Manager

The Create SnapVault Policy dialog box enables you to create a policy to set the priority for SnapVault transfers. You use policies to maximize the efficiency of transfers from the primary to the secondary volume.

- **Destination Cluster**

  The name of the cluster that you selected in the SnapVault tab of the Configure Protection dialog box.

- **Destination SVM**

  The name of the SVM that you selected in the SnapVault tab of the Configure Protection dialog box.

- **Policy Name**

  The name you provide for the new policy. Policy names can consist of the characters A through Z, a through z, 0 through 9, period (.), hyphen (-), and underscore (_).

- **Transfer Priority**

  The priority at which the transfer is run. You can select either Normal or Low. Transfer relationships with policies that specify a normal transfer priority run before those with policies that specify a low transfer priority. The default setting is Normal.

- **Comment**

  An optional field in which you can a add comment of up to 255 characters about the SnapVault policy.

- **Ignore Access Time**

  Specifies whether incremental transfers are ignored for files that have only their access time changed.

- **Replication Label**

  Lists in a table the rules associated with Snapshot copies selected by ONTAP that have a specific replication label in a policy. The following information and actions are also available:

  - Command buttons

    The command buttons enable you to perform the following actions:

    - Add

      Enables you to create a Snapshot copy label and retention count.

    - Edit Retention Count

Enables you to change the retention count for an existing Snapshot copy label. The retention count must be a number between 1 and 251. The sum of all retention counts for all rules cannot exceed 251.

- Delete

  Enables you to delete an existing Snapshot copy label.

- Snapshot Copy Label

  Displays the Snapshot copy label. If you select one or more volumes with the same local Snapshot copy policy, an entry for each label in the policy is displayed. If you select multiple volumes that have two or more local Snapshot copy policies, the table displays all labels from all policies

- Schedule

  Displays the schedule associated with each Snapshot copy label. If a label has more than one schedule associated with it, the schedules for that label are displayed in a comma-separated list. If you select multiple volumes with the same label but with different schedules, the schedule displays “Various” to indicate that more than one schedule is associated with the selected volumes.

- Destination Retention Count

  Displays the number of Snapshot copies with the specified label that are retained on the SnapVault secondary. Retention counts for labels with multiple schedules displays the sum of retention counts of each label and schedule pair. If you select multiple volumes with two or more local Snapshot copy policies, the retention count is empty.

Edit a protection relationship in Active IQ Unified Manager

You can edit an existing protection relationship to change the maximum transfer rate, the protection policy, or the protection schedule.

Destination Information

- **Destination Cluster**

  The name of the selected destination cluster.

- **Destination SVM**

  The name of the selected SVM

- **Relationship Settings**

  Enables you to specify the maximum transfer rate, SnapMirror policy, and schedule that the protection relationship uses:

  - Max Transfer Rate

    Specifies the maximum rate at which baseline data is transferred between clusters over the network. When selected, network bandwidth is limited to the value you specify. You can enter a numerical value and then select either kilobytes per second (KBps), megabytes per second (MBps), gigabytes per second (GBps), or terabytes per second (TBps). The maximum transfer rate that you specify must be greater than 1 KBps and less than 4 TBps. If you choose not to use a maximum transfer rate, the

baseline transfer between relationships is unlimited. If the primary cluster and the secondary cluster are the same, this setting is disabled.

- SnapMirror Policy

  Specifies the ONTAP SnapMirror policy for the relationship. The default is DPDefault.

- Create Policy

  Launches the Create SnapMirror Policy dialog box, which enables you to create and use a new SnapMirror policy.

- SnapMirror Schedule

  Specifies the ONTAP SnapMirror policy for the relationship. Available schedules include None, 5min, 8hour, daily, hourly, and weekly. The default is None, indicating that no schedule is associated with the relationship. Relationships without schedules have no lag status values unless they belong to a storage service.

- Create Schedule

  Launches the Create Schedule dialog box, which enables you to create a new SnapMirror schedule.

Command buttons

The command buttons enable you to perform the following tasks:

- **Cancel**

  Discards the selections and closes the Configure Protection dialog box.

- **Submit**

  Applies your selections and closes the Edit Relationship dialog box.

Use the Initialize/Update dialog box in Active IQ Unified Manager

The Initialize/Update dialog box enables you to perform a first-time baseline transfer on a new protection relationship, or to update a relationship if it is already initialized and you want to perform a manual, unscheduled, incremental update.

Transfer Options tab

The Transfer Options tab enables you to change the initialization priority of a transfer and to change the bandwidth used during transfers.

- **Transfer Priority**

  The priority at which the transfer is run. You can select either Normal or Low. Relationships with policies that specify a normal transfer priority run before those that specify a low transfer priority. Normal is selected by default.

- **Max Transfer Rate**

  Specifies the maximum rate at which data is transferred between clusters over the network. If you choose

not to use a maximum transfer rate, the baseline transfer between relationships is unlimited. If you select more than one relationship with different maximum transfer rates, you can specify one of the following maximum transfer rate settings:

- Use values specified during individual relationship setup or edit

  When selected, initialization and update operations use the maximum transfer rate specified at the time of each relationship's creation or edit. This field is available only when multiple relationships with different transfer rates are being initialized or updated.

- Unlimited

  Indicates that there is no bandwidth limitation on transfers between relationships. This field is available only when multiple relationships with different transfer rates are being initialized or updated.

- Limit bandwidth to

  When selected, network bandwidth is limited to the value you specify. You can enter a numerical value and then select either kilobytes per second (KBps), Megabytes per second (MBps), Gigabytes per second (GBps), or Terabytes per second (TBps). The maximum transfer rate that you specify must be greater than 1 KBps and less than 4 TBps.

Source Snapshot Copies tab

The Source Snapshot Copies tab displays the following information about the source Snapshot copy that is used for the baseline transfer:

- **Source Volume**

  Displays the names of the corresponding source volumes.

- **Destination Volume**

  Displays the names of the selected destination volumes.

- **Source Type**

  Displays the volume type. The type can be either Read/write or Data Protection.

- **Snapshot Copy**

  Displays the Snapshot copy that is used for the data transfer. Clicking the Snapshot copy value displays the Select Source Snapshot Copy dialog box, in which you can select a specific Snapshot copy for your transfer, depending on the type of protection relationship that you have and the operation that you are performing. The option to specify a different Snapshot copy is not available for data protection type sources.

Command buttons

The command buttons enable you to perform the following tasks:

- **Cancel**

  Discards your selections and closes the Initialize/Update dialog box.

- **Submit**

  Saves your selections and starts the initialize or update job.

Use the resynchronize dialog box in Active IQ Unified Manager

The Resynchronize dialog box enables you to resynchronize data on a SnapMirror or SnapVault relationship that was previously broken and then the destination was made a read/write volume. You might also resynchronize when a required common Snapshot copy on the source volume is deleted causing SnapMirror or SnapVault updates to fail.

Resynchronization Options tab

The Resynchronization Options tab enables you to set the transfer priority and the maximum transfer rate for the protection relationship that you are resynchronizing.

- **Transfer Priority**

  The priority at which the transfer is run. You can select either Normal or Low. Relationships with policies that specify a normal transfer priority run before those with policies that specify a low transfer priority.

- **Max Transfer Rate**

  Specifies the maximum rate at which data is transferred between clusters over the network. When selected, network bandwidth is limited to the value that you specify. You can enter a numerical value and then select either kilobytes per second (KBps), megabytes per second (MBps), gigabytes per second (GBps), or TBps. If you choose not to use a maximum transfer rate, the baseline transfer between relationships is unlimited.

Source Snapshot Copies tab

The Source Snapshot Copies tab displays the following information about the source Snapshot copy that is used for the baseline transfer:

- **Source Volume**

  Displays the names of the corresponding source volumes.

- **Destination Volume**

  Displays the names of the selected destination volumes.

- **Source Type**

  Displays the volume type: either Read/write or Data Protection.

- **Snapshot Copy**

  Displays the Snapshot copy that is used for the data transfer. Clicking the Snapshot copy value displays the Select Source Snapshot Copy dialog box, in which can select a specific Snapshot copy for your transfer, depending on the type of protection relationship you have and the operation you are performing.

Command buttons

- **Submit**

  Begins the resynchronization process and closes the Resynchronize dialog box.

- **Cancel**

  Cancels your selections and closes the Resynchronize dialog box.

Use the Select Source Snapshot Copy dialog box in Active IQ Unified Manager

You use the Select Source Snapshot Copy dialog box to select a specific Snapshot copy to transfer data between protection relationships, or you select the default behavior, which varies depending on whether you are initializing, updating, or resynchronizing a relationship, and whether the relationship is a SnapMirror or SnapVault.

Default

Enables you to select the default behavior for determining which Snapshot copy is used for initialize, update, and resynchronize transfers for SnapVault and SnapMirror relationships.

If you are performing a SnapVault transfer, the default behavior for each operation is as follows:

Operation	Default SnapVault behavior when source is read/write	Default SnapVault behavior when source is Data Protection (DP)
Initialize	Creates a new Snapshot copy and transfers it.	Transfers the last exported Snapshot copy.
Update	Transfers only labeled Snapshot copies, as specified in the policy.	Transfers the last exported Snapshot copy.
Resynchronize	Transfers all labeled Snapshot copies created after the newest common Snapshot copy.	Transfers the newest labeled Snapshot copy.

If you are performing a SnapMirror transfer, the default behavior for each operation is as follows:

Operation	Default SnapMirror behavior	Default SnapMirror behavior when relationship is second hop in a SnapMirror to SnapMirror cascade
Initialize	Creates a new Snapshot copy and transfers it and all Snapshot copies created prior to the new Snapshot copy.	Transfers all Snapshot copies from the source.
Update	Creates a new Snapshot copy and transfers it and all Snapshot copies created prior to the new Snapshot copy.	Transfers all Snapshot copies.

Operation	Default SnapMirror behavior	Default SnapMirror behavior when relationship is second hop in a SnapMirror to SnapMirror cascade
Resynchronize	Creates a new Snapshot copy and then transfers all Snapshot copies from the source.	Transfers all Snapshot copies from the secondary volume to the tertiary volume, and deletes any data added after creation of the newest common Snapshot copy.

Existing Snapshot Copy

Enables you to select an existing Snapshot copy from the list if Snapshot copy selection is allowed for that operation.

- **Snapshot Copy**

  Displays the existing Snapshot copies from which you can select for a transfer.

- **Date Created**

  Displays the date and time the Snapshot copy was created. Snapshot copies are listed from most recent to least recent, with the most recent at the top of the list.

If you are performing a SnapVault transfer and you want to select an existing Snapshot copy to transfer from a source to a destination, the behavior for each operation is as follows:

Operation	SnapVault behavior when specifying a Snapshot copy	SnapVault behavior when specifying a Snapshot copy in a cascade
Initialize	Transfers the specified Snapshot copy.	Source Snapshot copy selection is not supported for data protection volumes.
Update	Transfers the specified Snapshot copy.	Source Snapshot copy selection is not supported for data protection volumes.
Resynchronize	Transfers the selected Snapshot copy.	Source Snapshot copy selection is not supported for data protection volumes.

If you are performing a SnapMirror transfer and you want to select an existing Snapshot copy to transfer from a source to a destination, the behavior for each operation is as follows:

Operation	SnapMirror behavior when specifying a Snapshot copy	SnapMirror behavior when specifying a Snapshot copy in a cascade
Initialize	Transfers all Snapshot copies on the source, up to the specified Snapshot copy.	Source Snapshot copy selection is not supported for data protection volumes.

Operation	SnapMirror behavior when specifying a Snapshot copy	SnapMirror behavior when specifying a Snapshot copy in a cascade
Update	Transfers all Snapshot copies on the source, up to the specified Snapshot copy.	Source Snapshot copy selection is not supported for data protection volumes.
Resynchronize	Transfers all Snapshot copies from the source, up to the selected Snapshot copy, and then deletes any data added after creation of the newest common Snapshot copy.	Source Snapshot copy selection is not supported for data protection volumes.

Command buttons

The command buttons enable you to perform the following tasks:

- **Submit**

  Submits your selections and closes the Select Source Snapshot Copy dialog box.

- **Cancel**

  Discards your selections and closes the Select Source Snapshot Copy dialog box.

Reverse resync in Active IQ Unified Manager

When you have a protection relationship that is broken because the source volume is disabled and the destination is made a read/write volume, reverse resynchronization enables you to reverse the direction of the relationship so that the destination becomes the new source and the source becomes the new destination.

When a disaster disables the source volume in your protection relationship, you can use the destination volume to serve data by converting it to read/write, while you repair or replace the source, update the source, and reestablish the relationship. When you perform a reverse resync operation, data on the source that is newer than the data on the common Snapshot copy is deleted.

Before reverse resync

Displays the source and destination of a relationship before a reverse resync operation.

- **Source Volume**

  The name and location of the source volume before a reverse resync operation.

- **Destination Volume**

  The name and location of the destination volume before a reverse resync operation.

After reverse resync

Displays what the source and destination of a relationship is after a reserve resync operation.

- **Source Volume**

The name and location of the source volume after a reverse resync operation.

- **Destination Volume**

  The name and location of the destination volume after a reverse resync operation.

Command buttons

The command buttons enable you to perform the following actions:

- **Submit**

  Begins the reverse resynchronization process.

- **Cancel**

  Closes the Reverse Resync dialog box without initiating a reverse resync operation.

Display the Relationship: All relationships view in Active IQ Unified Manager

The Relationship: All Relationships view displays information about protection relationships on the storage system.

By default, when you access the Relationships page, the report that is displayed includes the top-level protection relationships for both volumes and storage VMs. The controls along the top of the page enable you to select a particular view, perform searches to locate specific objects, create and apply filters to narrow the list of displayed data, add/remove/reorder columns on the page, and export the data on the page to a .csv, .pdf, or .xlsx file. After you have customized the page, you can save the results as a custom view and then schedule a report of this data to be generated and emailed on a regular basis. By default, when you select the **Relationships** menu, the report displayed includes protection relationships for both volumes and storage VMs in your datacenter. You can use the **Filter** option to view only selected storage systems like only volumes or only storage VMs. The same report is displayed in the Storage page and only for the selected storage entity. If you want to view either volume or storage VM relationships, you can access either the **Storage** > **Volumes** > **Relationship: All Relationships** page or access **Protection** > **Relationships** > **Relationship: All Relationships**, and use the **Relationship Object Type** option in the **Filter** to filter out only volumes or storage VMs data.

The Relationships page that lists all the protection relationships has the link **View in System Manager** for the Destination cluster that allows you to view the same objects in ONTAP System Manager.

- **Status**

  Displays the current status of the protection relationship.

  The status can be one of Error ( ), Warning ( ), or OK ( ).

- **Source Storage VM**

  Displays the name of the source SVM. You can view more details about the source SVM by clicking the SVM name.

  When a SVM exists on the cluster but has not yet been added to the Unified Manager inventory, or that the SVM was created after the cluster's last refresh, this field will be empty. You must ensure that the SVM exists, or perform a rediscovery on the cluster to refresh the list of resources.

- **Source**

  Displays either the source volume or source storage VM being protected based on your selection. You can view more details about the source volume or storage VM by clicking the volume or storage VM name.

  If the message `Resource-key not discovered` is displayed, this might indicate that the volume exists on the cluster but has not yet been added to the Unified Manager inventory, or that the volume was created after the cluster's last refresh. You must ensure that the volume exists, or perform a rediscovery on the cluster to refresh the list of resources.

- **Destination Storage VM**

  Displays the name of the destination SVM. You can view more details about the destination SVM by clicking the SVM name.

- **Destination**

  Displays the name of the destination volume or storage VM based on your selection. You can view more details about the destination volume or storage VM by clicking the respective object name.

- **Relationship Object Type**

  Displays the type of object used in the relationship, such as storage VM, volume, and Consistency Group. For objects in a Consistency Relationship, the relationship source and destinations display the Consistency Group, and clicking them takes you to the LUNs page to view the relationship.

- **Policy**

  Displays the name of the protection policy for SnapMirror relationship. You can click the policy name to view details associated with that policy, including the following information:

  - Transfer Priority

    Specifies the priority at which a transfer runs for asynchronous operations. The transfer priority is Normal or Low. Normal priority transfers are scheduled before low priority transfers. The default is Normal.

  - Ignore Access Time

    Applies only to SnapVault relationships. This specifies whether incremental transfers ignore files which have only their access time changed. The values are either True or False. The default is False.

  - When Relationship is Out of Sync

    Specifies the action ONTAP performs when a synchronous relationship is not able to be synchronized. StrictSync relationships will restrict access to the primary volume if there is a failure to synchronize with the secondary volume. Sync relationships do not restrict access to the primary if there is a failure to synchronize with the secondary.

  - Tries Limit

    Specifies the maximum number of times to attempt each manual or scheduled transfer for a SnapMirror relationship. The default is 8.

  - Comments

Provides a text field for comments for specific to the selected policy.

- SnapMirror Label

  Specifies the SnapMirror label for the first schedule associated with the Snapshot copy policy. The SnapMirror label is used by the SnapVault subsystem when you back up Snapshot copies to a SnapVault destination.

- Retention Setting

  Specifies how long backups are kept, based on the time or the number of backups.

- Actual Snapshot Copies

  Specifies the number of Snapshot copies on this volume that match the specified label.

- Preserve Snapshot Copies

  Specifies the number of SnapVault Snapshot copies that are not deleted automatically even if the maximum limit for the policy is reached. The values are either True or False. The default is False.

- Retention Warning Threshold

  Specifies the Snapshot copy limit at which a warning is sent to indicate that the maximum retention limit is nearly reached.

- **Lag Duration**

  Displays the amount of time that the data on the mirror lags behind the source.

  The lag duration should be close to, or equal to, 0 seconds for StrictSync relationships.

- **Lag Status**

  Displays the lag status for managed relationships, and for unmanaged relationships that have a schedule associated with that relationship. Lag status can be:

  - Error

    The lag duration is greater than or equal to the lag error threshold.

  - Warning

    The lag duration is greater than or equal to the lag warning threshold.

  - OK

    The lag duration is within normal limits.

  - Not Applicable

    The lag status is not applicable for synchronous relationships because a schedule cannot be configured.

- **Last Successful Update**

Displays the time of the last successful SnapMirror or SnapVault operation.

The last successful update is not applicable for synchronous relationships.

- **Constituent Relationships**

  Displays whether there are any volumes in the selected object.

- **Relationship Type**

  Displays the relationship type used to replicate a volume. Relationship types include:

  - Asynchronous Mirror
  - Asynchronous Vault
  - Asynchronous MirrorVault
  - StrictSync
  - Sync

- **Transfer Status**

  Displays the transfer status for the protection relationship. The transfer status can be one of the following:

  - Aborting

    SnapMirror transfers are enabled; however, a transfer abort operation that might include removal of the checkpoint is in progress.

  - Checking

    The destination volume is undergoing a diagnostic check and no transfer is in progress.

  - Finalizing

    SnapMirror transfers are enabled. The volume is currently in the post-transfer phase for incremental SnapVault transfers.

  - Idle

    Transfers are enabled and no transfer is in progress.

  - In-Sync

    The data in the two volumes in the synchronous relationship are synchronized.

  - Out-of-Sync

    The data in the destination volume is not synchronized with the source volume.

  - Preparing

    SnapMirror transfers are enabled. The volume is currently in the pre-transfer phase for incremental SnapVault transfers.

  - Queued

SnapMirror transfers are enabled. No transfers are in progress.

- Quiesced

  SnapMirror transfers are disabled. No transfer is in progress.

- Quiescing

  A SnapMirror transfer is in progress. Additional transfers are disabled.

- Transferring

  SnapMirror transfers are enabled and a transfer is in progress.

- Transitioning

  The asynchronous transfer of data from the source to the destination volume is complete, and the transition to synchronous operation has started.

- Waiting

  A SnapMirror transfer has been initiated, but some associated tasks are waiting to be queued.

- **Last Transfer Duration**

  Displays the time taken for the last data transfer to complete.

  The transfer duration is not applicable for StrictSync relationships because the transfer should be simultaneous.

- **Last Transfer Size**

  Displays the size, in bytes, of the last data transfer.

  The transfer size is not applicable for StrictSync relationships.

- **Mediators**

  Displays the mediator status.

  - Not Applicable

    If the cluster does not support SnapMirror active sync.

  - Not Configured

    If it is not configured, or if it is configured but only the destination cluster is added and the source cluster is not added in Unified Manager.

  - Mediator IP Address

    If it is configured, and both source and destination clusters are added in Unified Manager.

- **State**

  Displays the state of the SnapMirror or SnapVault relationship. The state can be Uninitialized,

SnapMirrored, or Broken-Off. If a source volume is selected, the relationship state is not applicable and is not displayed.

- **Relationship Health**

  Displays the relationship heath of the cluster.

- **Unhealthy Reason**

  The reason the relationship is in an unhealthy state.

- **Transfer Priority**

  Displays the priority at which a transfer runs. The transfer priority is Normal or Low. Normal priority transfers are scheduled before low priority transfers.

  The transfer priority is not applicable for synchronous relationships because all transfers are treated with the same priority.

- **Schedule**

  Displays the name of the protection schedule assigned to the relationship.

  The schedule is not applicable for synchronous relationships.

- **Version Flexible Replication**

  Displays either Yes, Yes with backup option, or None.

- **Source Cluster**

  Displays the FQDN, short name, or IP address of the source cluster for the SnapMirror relationship.

- **Source Cluster FQDN**

  Displays the name of the source cluster for the SnapMirror relationship.

- **Source Node**

  Displays the name of the source node name link for the SnapMirror relationship of a volume and displays the SnapMirror relationship node count link when the object is a Storage VM or Consistency Group.

In the custom view, when you click the node name link, you can view and extend protection for storage objects on which the volumes of those Consistency Groups belong to a SnapMirror active sync relationship.

When you click the node count link, it takes you to the node page with respective nodes associated with that relationship. When the node count is 0, there is no value displayed as there are no nodes associated with the relationship.

- **Destination Node**

  Displays the name of the destination node name link for the SnapMirror relationship of a volume and displays the SnapMirror relationship node count link when the object is a Storage VM or Consistency Group.

  When you click the node count link, it takes you to the node page with respective nodes associated with

that relationship. When the node count is 0, there is no value displayed as there are no nodes associated with the relationship.

- **Destination Cluster**

  Displays the name of the destination cluster for the SnapMirror relationship.

- **Destination Cluster FQDN**

  Displays the FQDN, short name, or IP address of the destination cluster for the SnapMirror relationship.

- **Protected By**

  Displays the different relationships. In this column, you can view volume and consistency group relationships for clusters and storage virtual machines order, including:

  - SnapMirror
  - Storage VM DR
  - SnapMirror, Storage VM DR
  - Consistency Group
  - SnapMirror, Consistency Group.

Related information

- For information about **Relationship: MetroCluster** view, see Monitoring MetroCluster configurations.
- For information about **Relationship: Last 1 month Transfer Status** view, see Relationship: Last 1 month Transfer Status view.
- For information about **Relationship: All Relationships** view, see Relationship: Last 1 month Transfer Rate view.

Display the Last 1 month Transfer Status view in Active IQ Unified Manager

The Relationship: Last 1 month Transfer Status view enables you to analyze the transfer trends over a period of time for volumes and Storage VMs in asynchronous relationships. This page also displays whether the transfer was a success or a failure.

The controls along the top of the page enable you to perform searches to locate specific objects, create, and apply filters to narrow the list of displayed data, add/remove/reorder columns on the page, and export the data on the page to a `.csv`, `.pdf`, or `.xlsx` file. After you have customized the page, you can save the results as a custom view and then schedule a report of this data to be generated and emailed on a regular basis. You can use the **Filter** option to view only selected storage systems like only volumes or only Storage VMs. The same report is displayed in the Storage page and only for the selected storage entity. For example, if you want to view volume relationships you can access either the Relationship: Last 1 Month Transfer Status report for the Storage VMs either from the **Storage** > **Storage VMs** > **Relationship: Last 1 Month Transfer Status** menu or from **Protection** > **Relationships** > **Relationship: Last 1 Month Transfer Status** menu, and use the **Filter** to only view data for volumes.

- **Source Volume**

  Displays the source volume name.

- **Destination Volume**

  Displays the destination volume name.

- **Operation Type**

  Displays the type of volume transfer.

- **Operation Result**

  Displays whether volume transfer was successful.

- **Transfer Start Time**

  Displays the volume transfer start time.

- **Transfer End Time**

  Displays the volume transfer end time.

- **Transfer Duration**

  Displays the time taken (in hours) to complete the volume transfer.

- **Transfer Size**

  Displays the size (in MB) of the transferred volume.

- **Source SVM**

  Displays the storage virtual machine (SVM) name.

- **Source Cluster**

  Displays the source cluster name.

- **Destination SVM**

  Displays the destination SVM name.

- **Destination Cluster**

  Displays the destination cluster name.

Related information

- For information about **Relationship: All Relationships** view, see Relationship: All Relationships view.
- For information about **Relationship:MetroCluster** view, see Monitoring MetroCluster configurations.
- For information about **Relationship: All Relationships** view, see Relationship: Last 1 month Transfer Rate view.

Display the Relationship: Last 1 month Transfer Rate view in Active IQ Unified Manager

The Relationship: Last 1 month Transfer Rate view enables you to analyze the amount of data volume that is transferred on a day-to-day basis for volumes in asynchronous

relationships. This page also provides details about daily transfers and the time required to complete the transfer operation for volumes and Storage VMs.

The controls along the top of the page enable you to perform searches to locate specific objects, create and apply filters to narrow the list of displayed data, add/remove/reorder columns on the page, and export the data on the page to a .csv, .pdf, or .xlsx file. After you have customized the page, you can save the results as a custom view and then schedule a report of this data to be generated and emailed on a regular basis. For example, if you want to view volume relationships, you can access either the **Storage** > **Volumes** > **Relationship: Last 1 Month Transfer Rate** menu or access **Protection** > **Relationships** > **Relationships:Last 1 Month Transfer Rate** menu, and use the **Filter** to only view data for volumes.

- **Total Transfer Size**

  Displays the total size of the volume transfer in gigabytes.

- **Day**

  Displays the day on which the volume transfer was initiated.

- **End Time**

  Displays the volume transfer end time with date.

Related information

- For information about **Relationship:MetroCluster** view, see Monitoring MetroCluster configurations.
- For information about **Relationship: Last 1 month Transfer Status** view, see Relationship: Last 1 month Transfer Status view.
- For information about **Relationship: All Relationships** view, see Relationship: Last 1 month Transfer Rate view

Generate custom reports

Unified Manager reporting

Learn about Active IQ Unified Manager reporting

Active IQ Unified Manager (formerly OnCommand Unified Manager) provides the ability to view, customize, download, and schedule reports for your ONTAP storage systems. The reports can provide details about the storage system capacity, health, performance, security, and protection relationships.

The new Unified Manager reporting and scheduling functionality introduced in Active IQ Unified Manager 9.6 replaces the previous reporting engine that was retired in Unified Manager version 9.5.

Reporting provides different views of your network, offering actionable intelligence on capacity, health, performance, security, and protection data. You can customize your views by showing, hiding, and rearranging columns, filtering data, sorting data, and searching the results. You can save custom views for reuse, download them as reports, and schedule them as recurring reports to distribute through email.

You can download views in Microsoft® Excel format and customize them. You can use advanced Excel features, such as complex sorts, layered filters, pivot tables, and charts. When satisfied with the resulting Excel report, you can upload the Excel file for use each time the report is scheduled and shared.

In addition to generating reports from the user interface, you can extract health, security, and performance data from Unified Manager using these additional methods:

- Using Open Database Connectivity (ODBC) and ODBC tools to directly access the database for cluster information
- Executing Unified Manager REST APIs to return the information you are interested in reviewing

From this release of Active IQ Unified manager, following enhancements are made to the reports:

- Email is sent for a report as per the configured schedule. Even when you generate an on-demand report, you will receive an email.
- The file name of the report and metadata of the report includes the hostname from where the report was generated.
  Even if any one changes the filename, still you can identify the hostname from where the report was generated due to this enhancement.

Access points for generating reports

Explore three ways to access reporting capabilities in Active IQ Unified Manager

You can gather information in Unified Manager about your clusters to create reports from the UI, MySQL database queries, and REST APIs.

These sections describe Unified Manager reporting and scheduling through the UI.

There are three ways you can access the reporting capabilities provided by Unified Manager:

- Extracting data directly from the inventory pages in the UI.
- Using Open Database Connectivity (ODBC) and ODBC tools to access all the available objects.
- Executing Unified Manager REST APIs to return the information that you want to review.

These sections describe Unified Manager reporting and scheduling through the UI.

Access databases for custom reporting in Active IQ Unified Manager

Unified Manager uses a MySQL database to store data from the clusters that it is monitoring. Data is persisted into various schemas in the MySQL database.

All table data from the following databases are available:

Database	Description
netapp_model	Data about the objects on ONTAP controllers.
netapp_model_view	Data about the objects on ONTAP controllers, suitable for report tool consumption.
netapp_performance	Cluster specific performance counters.
ocum	Unified Manager application data and information to support UI filtering, sorting, and the calculation of some derived fields.
ocum_report	Data for inventory configuration and capacity-related information.
ocum_report_birt	Views for inventory configuration and capacity-related data, suitable for report tool consumption.
opm	Performance configuration settings and threshold information.
scalemonitor	Data about the Unified Manager application health and performance issues.
vmware_model	VMware object data for datastores hosted on NetApp storage.
vmware_model_view	Views for VMware object data for datastores hosted on NetApp storage, suitable for report tool consumption.
vmware_performance	VMware performance counter data for datastores hosted on NetApp storage.

A reporting user — a Database user with the Report Schema role — is able to access the data in these tables. This user has read-only access to reporting and other database views directly from the Unified Manager database. Note that this user does not have permission to access any tables that contain user data or cluster credential information.

Use Active IQ Unified Manager REST APIs for reporting

You can use REST APIs to help manage your clusters by viewing the health, capacity, performance, and security information captured by Unified Manager.

REST APIs are exposed through the Swagger web page. You can access the Swagger web page to display the Unified Manager REST API documentation, as well as to manually issue an API call. From the Unified Manager web UI, in the Menu Bar, click the **Help** button and then select **API Documentation**. For information about Unified Manager REST APIs, see Getting started with Active IQ Unified Manager REST APIs.

You must have the Operator, Storage Administrator, or Application Administrator role to access the REST APIs.

Understand reports in Active IQ Unified Manager

Reports display detailed information about storage, network, quality of service, and protection relationships, helping you to identify and solve potential problems before they occur.

When you customize a view, you can save it with a unique name for future use. You can schedule a report based on that view to run on a regular basis and share it with others. You can also download the view in Excel to customize it using advanced Excel features, then upload that file back into Unified Manager . If you schedule a report using that view, it will use the Excel file you uploaded to create robust reports you can share.

You can manage all reports that have been scheduled from the Report Schedules page.

You must have the Application Administrator or Storage Administrator role to manage reports.

You can download reports as comma-separated values (CSV), Excel, or PDF files.

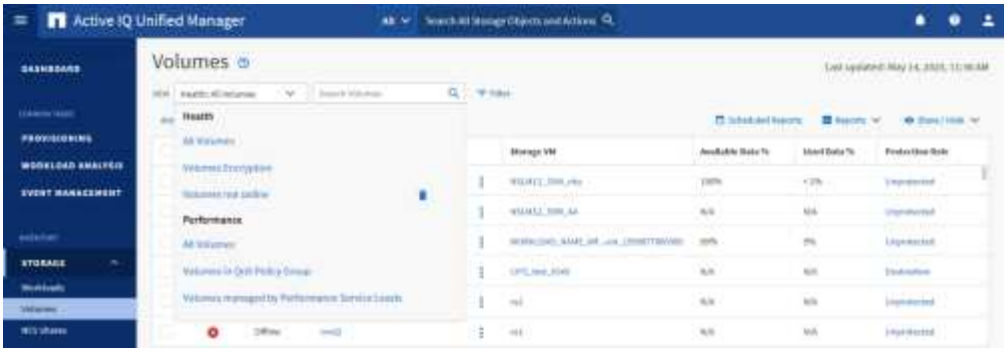
Understand the view and report relationship in Active IQ Unified Manager

Views and inventory pages become reports when you download or schedule them.

You can customize and save views and inventory pages for reuse. Almost everything you can view in Unified Manager can be saved, reused, customized, scheduled, and shared as a report.

In the view drop down, items with the delete icon are existing custom views that you or another user have created. Items without an icon are default views provided with Unified Manager. Default views cannot be modified or deleted.

- If you delete a custom view from the list, it also deletes any Excel files or scheduled reports that use that view.
- If you change a custom view, reports that use that view will reflect the change the next time the report is generated and sent by email according to the report schedule. When changing views, make sure your changes work with any associated Excel customizations used for the reports. If needed, you can update the Excel file by downloading it, making the required changes, and uploading it as a new Excel customization for the view.



Only users with the Application Administrator or Storage Administrator role can see the delete icon, change or delete a view, or change or delete a scheduled report.

Learn about types of reports in Active IQ Unified Manager

This table provides a comprehensive list of the views and inventory pages that are available as reports that you can customize, download, and schedule.

Active IQ Unified Manager reports

Type	Storage or network object
Capacity	Clusters Aggregates Volumes Qtrees
Health	Clusters Nodes Aggregates Storage VMs Volumes SMB/CIFS shares NFS shares

Type	Storage or network object
Performance	Clusters Nodes Aggregates Storage VMs Volumes LUNs NVMe namespaces Network Interfaces (LIFs) Ports
Quality of Service	Traditional QoS policy groups Adaptive QoS policy groups Performance Service Level policy groups
Volume protection relationships (available from the Volumes page)	All relationships Last 1 month transfer status Last 1 month transfer rate
Security	Storage VMs Clusters

Understand reporting limitations in Active IQ Unified Manager

There are some limitations with the new Active IQ Unified Manager reporting functionality of which you should be aware.

Existing reports from previous versions of Unified Manager

You can only edit the schedule and recipients for existing reports that were created and imported (as .rptdesign files) in Unified Manager 9.5 and earlier releases. If you customized any of the standard reports that were provided with Unified Manager 9.5 or earlier, these custom reports are not imported into the new reporting tool.

If you need to edit existing reports imported from .rptdesign files, do one of the following and remove the imported report:

- create a new view and schedule a report from that view (preferred)
- hover over the report, copy the SQL, and pull the data using an external tool

The default views can be generated as reports without the need for any customization. You can use the new reporting solution to recreate any custom reports.

Schedule and report relationship

You can create many different schedules with any combination of recipients for each saved report. However, you cannot reuse the schedule for multiple reports.

Report protection

Any user with the appropriate permissions can edit or delete reports. There is no way to prevent other users from removing or making changes to saved views or schedules.

Event reports

Although you can customize the event view and download the resulting report in CSV format, you cannot schedule recurring event reports for generation and distribution.

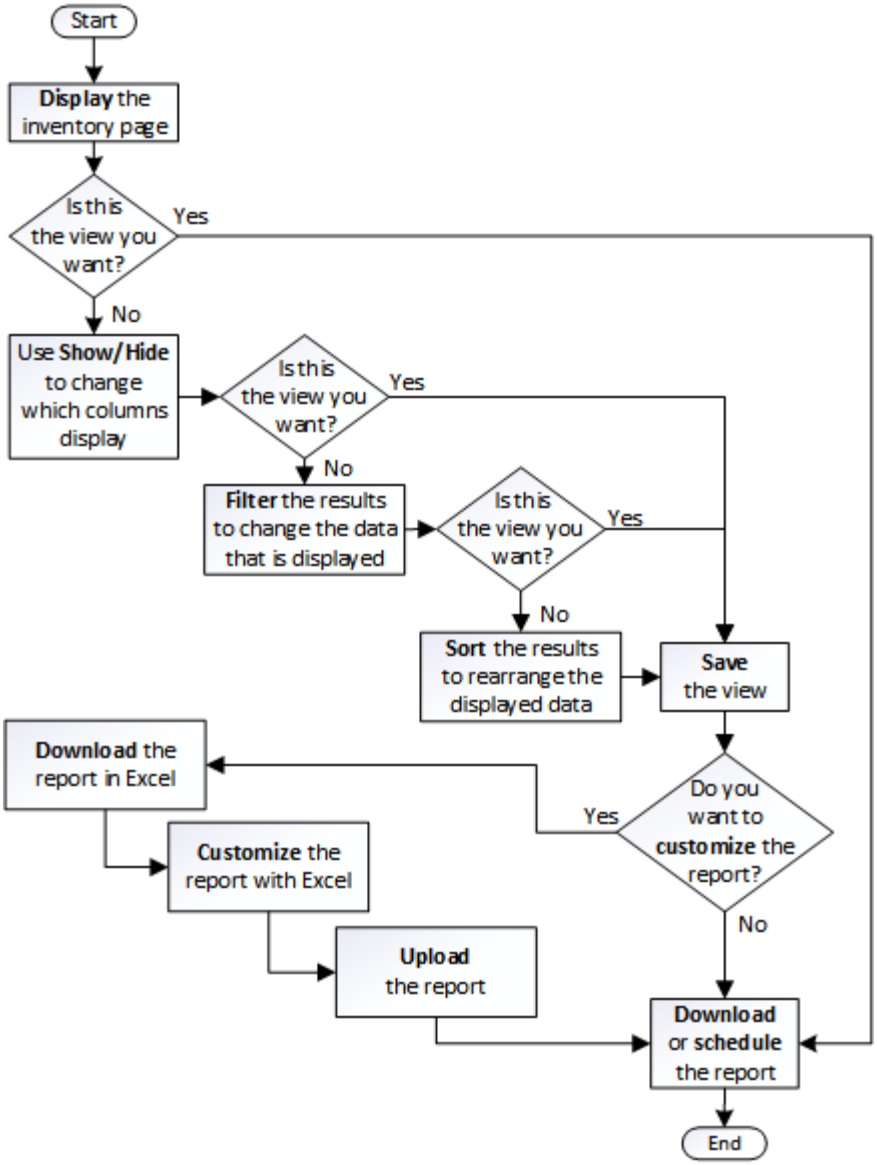
Report attachments

Reports cannot be sent in the body of an email. Instead, reports are only sent as PDF, Excel, or CSV attachments.

Work with reports

Understand report workflows in Active IQ Unified Manager

Decision tree that describes the report workflow.



Get started with reporting in Active IQ Unified Manager

Create a sample custom report to experience exploring views and scheduling reports. This quick start report finds a list of volumes that you might want to move to the cloud tier because there is a fair amount of inactive (cold) data. You will open the Performance: All Volumes view, customize the view using filters and columns, save the custom view as a report, and schedule the report to share once a week.

Before you begin

- You must have the Application Administrator or Storage Administrator role.
- You must have configured FabricPool aggregates and have volumes on those aggregates.

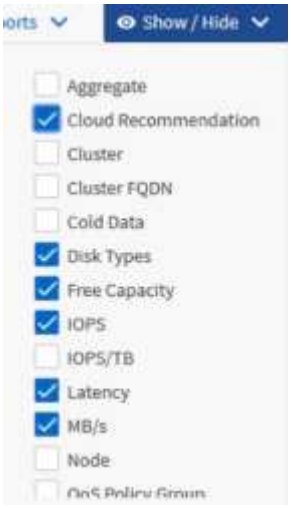
Follow the steps below to:

- Open the default view
- Customize the columns by filtering and sorting the data

- Save the view
- Schedule a report to be generated for the custom view

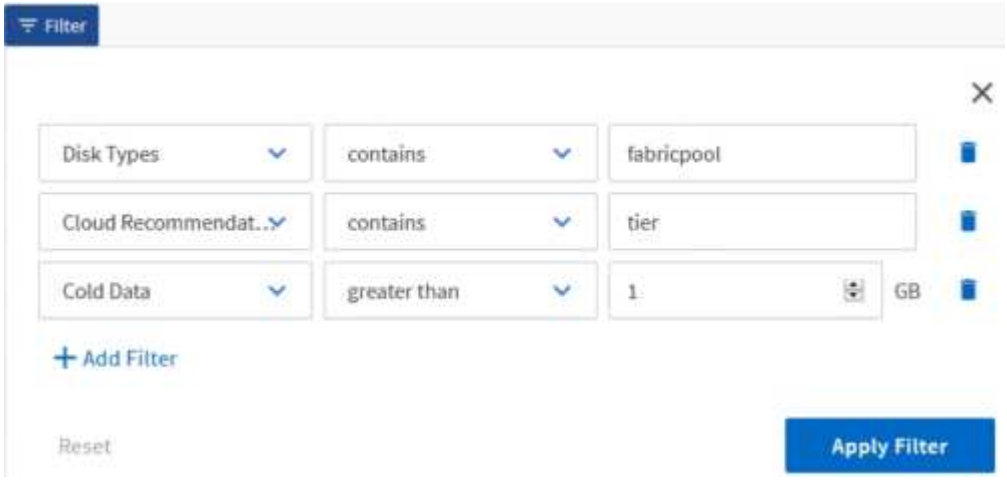
Steps

1. In the left navigation pane, click **Storage** > **Volumes**.
2. In the View menu, select **Performance** > **All Volumes**.
3. Click **Show/Hide** to make sure the “Disk Types” column appears in the view.



Add or remove other columns to create a view that contains the fields that are important for your report.

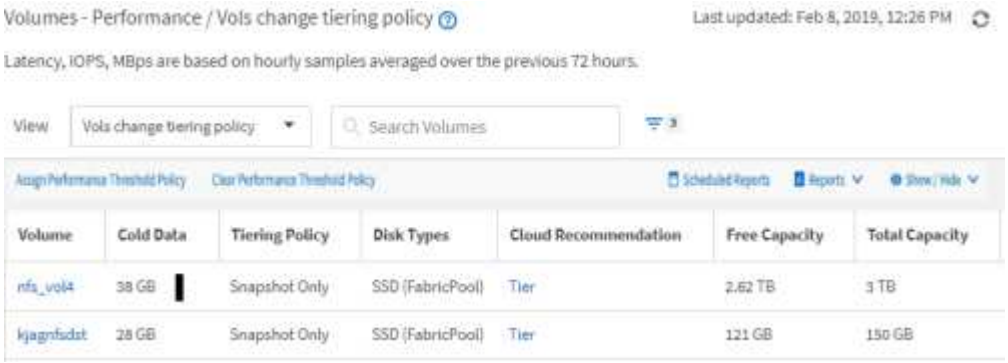
4. Drag the “Disk Types” column next to the “Cloud Recommendation” column.
5. Click the filter icon to add the following three filters, and then click **Apply Filter**:
   - Disk Types contains FabricPool
   - Cloud Recommendation contains tier
   - Cold Data greater than 10 GB



Note that each filter is joined with a logical AND so that all volumes returned must meet all the criteria. You can add a maximum of five filters.

6. Click the top of the Cold Data column to sort the results so that the volumes with the most cold data appear at the top of the view.

7. When the view is customized, the view name is Unsaved View. Name the view to reflect what the view is showing, for example “Vols change tiering policy”. When done, click the check mark or press **Enter** to save the view with the new name.



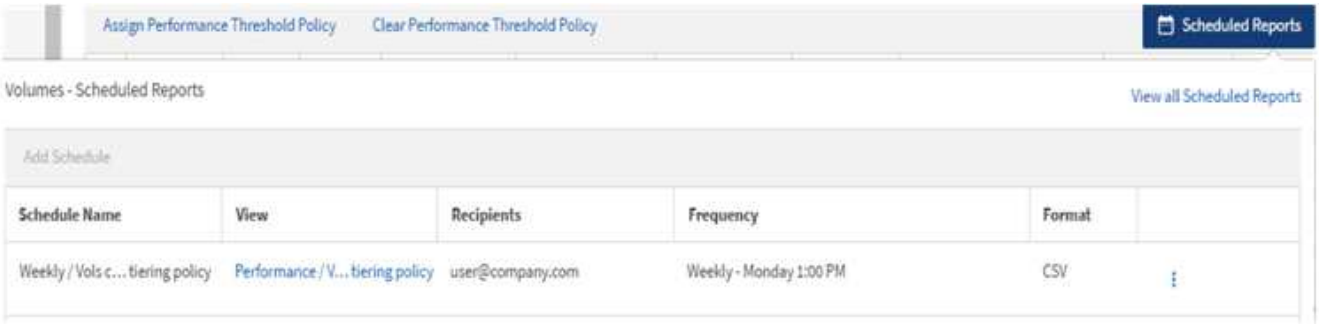
8. Download the report as a **CSV**, **Excel**, or **PDF** file to see the output before you schedule or share it.

   Open the file with an installed application, such as Microsoft Excel (CSV or Excel) or Adobe Acrobat (PDF), or save the file.

   > You can further customize your report using complex filters, sorts, pivot tables, or charts by downloading the view as an Excel file. After you open the file in Excel, use the advanced features to customize your report. When satisfied, upload the Excel file. This file, with its customizations, is applied to the view when the report is run.

   For more information on customizing reports using Excel, see *Sample Microsoft Excel reports*.

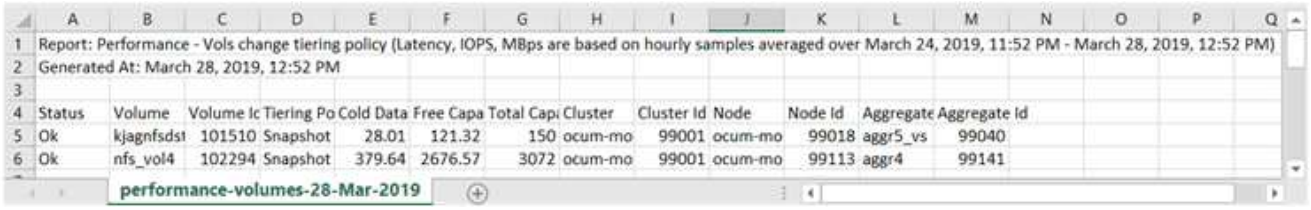
9. Click the **Scheduled Reports** button on the inventory page. All scheduled reports relating to the object, in this case volumes, appear in the list.



10. Click **Add Schedule** to add a new row to the Report Schedules page so you can define the schedule characteristics for the new report.
11. Enter a name for the report and complete the other report fields, then click the check mark (✓) at the end of the row.

    The report is sent immediately as a test. After that, the report generates and is sent by email to the recipients listed using the specified frequency.

    The following sample report is in CSV format:



	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q
1	Report: Performance - Vols change tiering policy (Latency, IOPS, MBps are based on hourly samples averaged over March 24, 2019, 11:52 PM - March 28, 2019, 12:52 PM)																
2	Generated At: March 28, 2019, 12:52 PM																
3																	
4	Status	Volume	Volume Ic	Tiering Po	Cold Data	Free Capa	Total Capi	Cluster	Cluster Id	Node	Node Id	Aggregate	Aggregate Id				
5	Ok	kjagnfsdst	101510	Snapshot	28.01	121.32	150	ocum-mo	99001	ocum-mo	99018	aggr5_vs	99040				
6	Ok	nfs_vol4	102294	Snapshot	379.64	2676.57	3072	ocum-mo	99001	ocum-mo	99113	aggr4	99141				

performance-volumes-28-Mar-2019

The following sample report is in PDF format:



Report: Performance - Vols change tiering policy (Latency, IOPS, MBps are based on hourly samples averaged over March 24, 2019, 11:51 PM - March 28, 2019, 12:51 PM)
Generated At: March 28, 2019, 12:51 PM

Status	Volume	Tiering Policy	Cold Data (GB)	Free Capacity (GB)	Total Capacity (GB)	Cluster	Node	Aggregate
Ok								
Ok								

Based on the results shown in the report, you might want to use ONTAP System Manager or the ONTAP CLI to change the tiering policy to “auto” or “all” for certain volumes to offload more cold data to the cloud tier.

Search for a scheduled report in Active IQ Unified Manager

You can search for scheduled reports by name, view name, object type, or recipients.

1. In the left navigation pane, click **Storage Management** > **Report Schedules**.
2. Use the **Search Scheduled Reports** text field.

To find reports by …	Try …
Schedule name	Type part of the report schedule name.
View name	Type part of the report view name. Default views and custom views appear in the view list.
Recipient	Type part of the email address.
File type	Type "PDF", "CSV", or "XLSX".

3. You can click a column heading to sort reports in ascending or descending order by that column, such as schedule name or format.

Customize reports

Customize reports in Active IQ Unified Manager

There are many ways you can customize views so that you can create a report that contains all the information you need to manage your ONTAP clusters.

Start with a default inventory page or a custom view, then customize it by adding or removing columns, changing the column order, filtering the data, or sorting on a specific column in ascending or descending order.

Starting in Unified Manager 9.8, you can also download the view in Excel to customize it using advanced features. When done, upload the customized Excel file. If you schedule a report using that view, it uses the customized Excel file to create robust reports you can share.

For more information on customizing reports using Excel, see *Sample Microsoft Excel reports*.

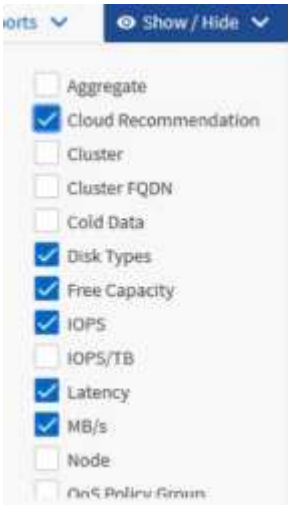
You must have the Application Administrator or Storage Administrator role to manage reports.

Customize columns in a report in Active IQ Unified Manager

Use **Show/Hide** to choose the columns you want to use in your report. Drag the columns on the inventory page to rearrange them.

Steps

1. Click **Show/Hide** to add or remove columns.



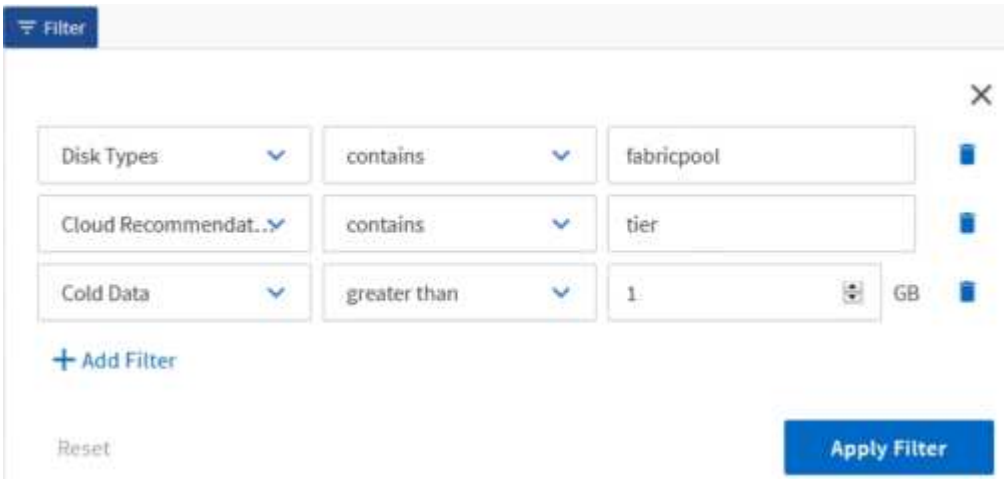
2. On the inventory page, drag columns to rearrange them in the order you want in your report.
3. Name the unsaved view to save your changes.

Filter data in Active IQ Unified Manager

Filter the data to make sure the results match your report requirements. Filtering enables you to display only the data in which you are interested.

Steps

1. Click the filter icon to add filters to focus the results you want to view, and then click **Apply Filter**.



2. Name the unsaved view to save your changes.

Sort data in Active IQ Unified Manager

To sort the results, click a column and indicate ascending or descending order. Sorting data prioritizes the information you need for the report.

Steps

1. Click the top of a column to sort the results so that the most important information appears at the top of the view.
2. Name the unsaved view to save your changes.

Search to refine your view in Active IQ Unified Manager

After you have the view you want, you can further refine the results using the Search field to focus on the results that you want to include in the report.

Steps

1. Open the custom or default view that you want to use as the basis of your report.
2. Type in the Search field to refine the data listed in the view. You can enter partial data in any of the displayed columns. For example, if you want to search for nodes that include "US_East" in the name, you can refine the full list of nodes.

   The results of your search are saved in the custom view and used in the resulting scheduled report.
3. Name the unsaved view to save your changes.

Use Excel to customize your report in Active IQ Unified Manager

After you have saved the view, you can download it in Excel Workbook format (.xlsx). When you open the Excel file, you can use advanced Excel features to customize your report.

Before you begin

You can only upload an Excel Workbook file with the .xlsx extension.

For example, some advanced Excel features you can use in your report include:

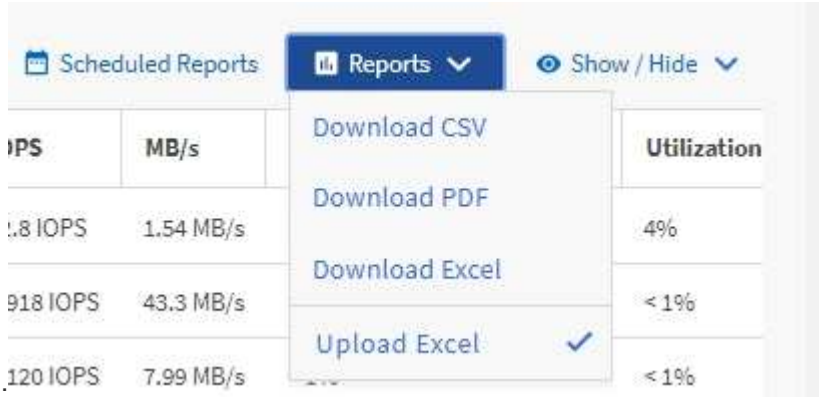
- Multi-column sort
- Complex filtering
- Pivot tables
- Charts

- The downloaded Excel file uses the default file name for the view, not your saved name.
  - The format is `<View Area>-<Day>-<Month>-<Year>-<Hour>-<Minute>-<Second>.xlsx`.
  - For example, a custom saved view named `Volumes-not online` has a file name of `health-volumes-05-May-2020-19-18-00.xlsx` if saved at that day and time.
- You can add sheets to the Excel file, but do not change existing sheets.
  - Do not change the existing sheets, data and info. Instead, copy the data to a new page

that you create.

- One exception to the above rule is that you can create formulas on the “data” page. Use the data page formulas to create charts on new pages.
- Do not name a new sheet data or info.

- If a customized Excel file exists, there is a check mark beside the **Reports** > **Upload Excel** menu item. When you download the Excel file, the version with the customizations is used.



Steps

1. Open the default, custom, or saved view that you want to use as the basis of your report.
2. Select **Reports** > **Download Excel**.
3. Save the file.
   The file is saved to your downloads folder.
4. Open the saved file in Excel.
   Do not move the file to a new location, or if you do your work in another location, save the file back to the original location using the original file name before uploading the file.
5. Customize the file using Excel features, such as complex sorts, layered filters, pivot tables, or charts. For more information, see the Microsoft® Excel documentation.
6. Select **Reports** > **Upload Excel** and select the file that you modified.
   The most recently downloaded file is uploaded from the same file location.
7. Send yourself a test report using the **Scheduled Reports** feature.

Download reports in Active IQ Unified Manager

You can download reports and save the data to a local or network drive as a comma-separated values (CSV) file, a Microsoft Excel (.XLSX) file, or a PDF file. You can open CSV and XLSX files with spreadsheet applications, such as Microsoft Excel, and PDF files with readers such as Adobe Acrobat.

Steps

1. Click the **Reports** button to download the report as one of the following:

Choose	To…
Download CSV	Save the report as a comma-separated values (CSV) file.
Download PDF	Save the report as a .pdf file.

Choose	To…
Download Excel	Save the report as a Microsoft Excel (XLSX) file.

Schedule reports

Schedule reports in Active IQ Unified Manager

After you have a view you want to reuse and share as a report, you can schedule it using Active IQ Unified Manager. You can manage scheduled reports, changing the recipients and distribution frequency for each report schedule.

You can schedule most views or inventory pages in Unified Manager. Exceptions are events, which are reports you can download as CSV files, but you cannot schedule events for regeneration and sharing. You also cannot download or schedule the dashboards, favorites, or configuration pages.

Starting in Active IQ Unified Manager 9.8, you can download views in Microsoft® Excel format and customize them. You can use advanced Excel features such as complex sorts, layered filters, pivot tables, and charts. When satisfied with the resulting Excel report, you can upload the Excel file for use each time the report is scheduled and shared.

You can schedule the built in views or views you customize. You can choose which file type to send, either CSV, PDF, or XSLX. When you schedule a report for the first time, you can download it and assign yourself as the only recipient to see the report as your recipients will see it.

Schedule a report

Schedule a report in Active IQ Unified Manager

After you have a view or Excel file that you want to schedule for regular generation and distribution, you can schedule the report.

Before you begin

- You must have the Application Administrator or Storage Administrator role.
- You must have configured the SMTP server settings in the **General** > **Notifications** page so that the reporting engine can send reports as email attachments to the list of recipients from the Unified Manager server.
- The email server must be configured to allow attachments to be sent with the generated emails.

Use the following steps to test and schedule a report to be generated for a view. Select or customize the view you want to use. The following procedure uses a network view that shows the performance of your network interfaces, but you can use any view you want.

Steps

1. Open your view. This example uses the default network view that shows LIF performance. In the left navigation pane, click **Network > Network Interfaces**.
2. Customize the view as needed using the built-in Unified Manager features.
3. After you customized the view, you can provide a unique name in the **View** field and click the check mark to save it.



4. You can use the advanced features of Microsoft® Excel to customize your report. For details, see Using Excel to customize your report.
5. To see the output before you schedule or share it:

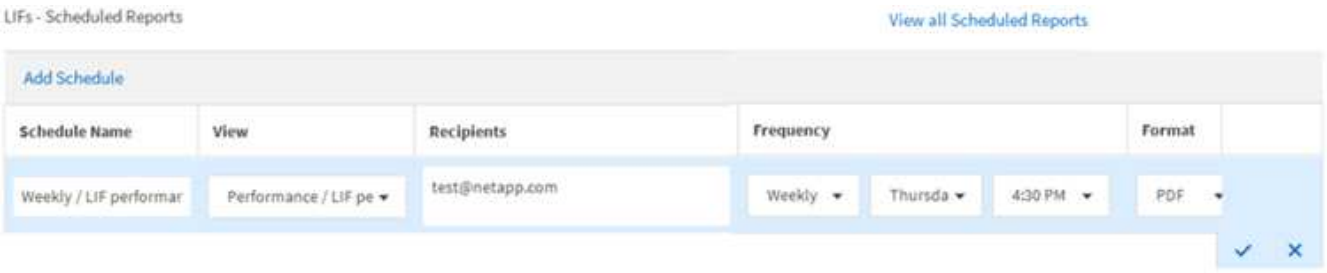
Option	Description
If you used Excel to customize the report	View the existing downloaded Excel file.
If you did not use Excel to customize the report	Download the report as a **CSV**, **PDF** or **XLSX** file.

   Open the file with an installed application, such as Microsoft Excel (CSV/XSLX) or Adobe Acrobat (PDF).

6. If you are satisfied with the report, click **Scheduled Reports**.
7. In the Report Schedules page, click **Add Schedule**.
8. Accept the default name, which is a combination of the view name and the frequency, or customize the **schedule name**.
9. To test the scheduled report the first time, only add yourself as the **recipient**. When satisfied, add the email addresses for all report recipients.
10. Specify how frequently the report will be generated and sent to the recipients. You can choose **Daily**, **Weekly**, or **Monthly**.
11. Select the format, either **PDF**, **CSV**, or **XSLX**.

    For reports where you used Excel to customize the content, always select **XSLX**.

12. Click the checkmark (✓) to save the report schedule.



   The report is sent immediately as a test. After that, the report generates and is sent by email to the recipients listed using the scheduled frequency.

Schedule imported .rptdesign reports in Active IQ Unified Manager

You can schedule existing reports that were created and imported in an earlier release of Unified Manager.

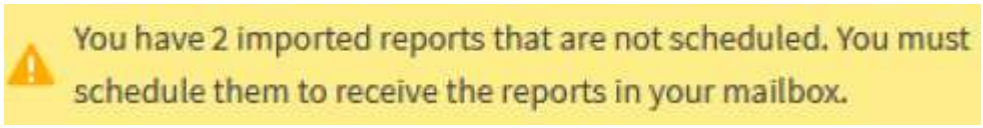
Scheduling imported reports requires the following:

- imported BIRT designed .rptdesign file reports in an earlier Unified Manager release
- applicable when upgrading to Unified Manager 9.6 GA or later

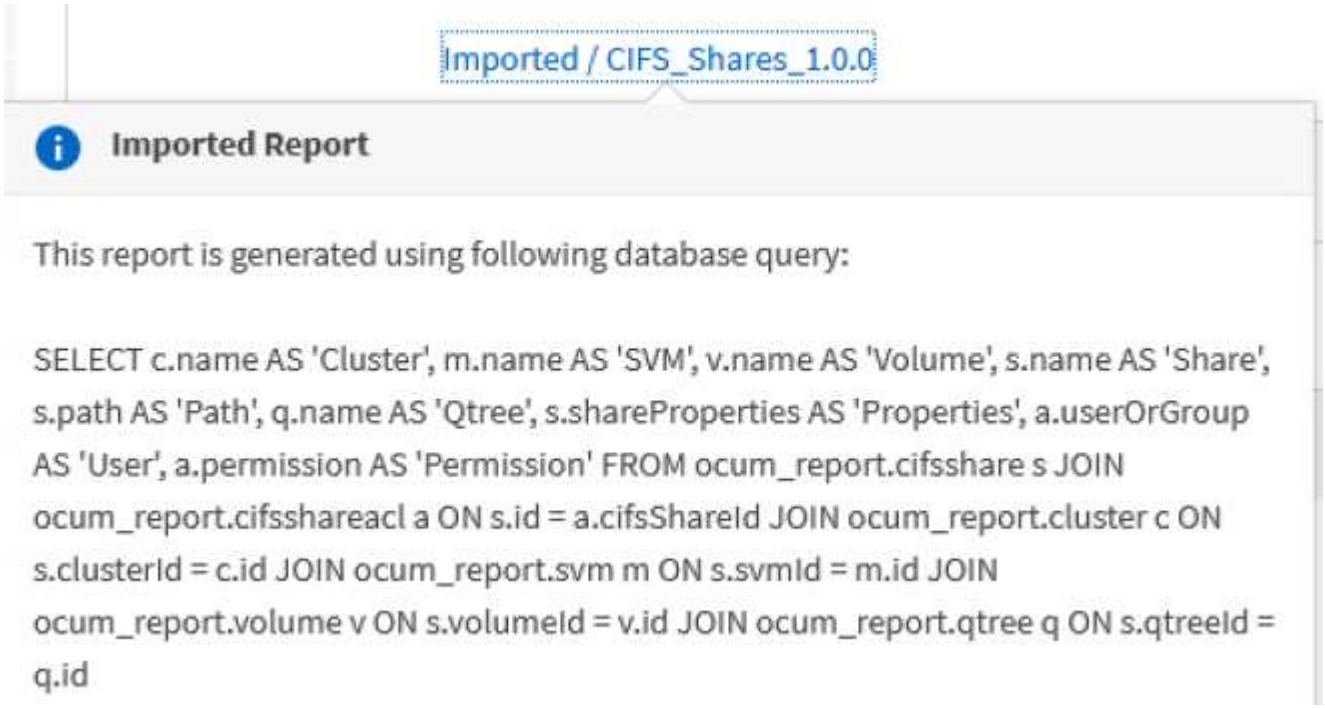
After upgrading to Unified Manager 9.6 GA or later, the Report Schedules page lists the imported reports. You can edit the schedule for these reports to specify the recipient email addresses, frequency, and format (PDF or CSV). Otherwise these reports cannot be edited or viewed in the Unified Manager UI.

Steps

1. Open the Report Schedules page. If you have imported reports, a message appears.



2. Click the **View** name to display the SQL query that is being used to generate the report.



3. Click the more icon ⋮, click **Edit**, define the report schedule details, and save the report.

You can also delete any unwanted reports from the more icon ⋮.

Manage report schedules

Manage report schedules in Active IQ Unified Manager

You can manage your report schedules from the Report Schedules page. You can view, modify, or delete existing schedules.

Before you begin

You cannot schedule new reports from the Report Schedules page. You can only add scheduled reports from the object inventory pages.

- You must have the Application Administrator or Storage Administrator role.

Steps

1. In the left navigation pane, click **Storage Management** > **Report Schedules**.
2. On the Report Schedules page:

If you want to…	Then…
View an existing schedule	Scroll through the list of existing reports using the scroll bars and page controls.
Edit an existing schedule	a. Click the more icon ⋮ for the schedule you want to use. b. Click **Edit**. c. Make the necessary changes. d. Click the check mark to save your changes.
Delete an existing schedule	a. Click the more icon ⋮ for the schedule you want to use. b. Click **Delete**. c. Confirm your decision.

Edit scheduled reports in Active IQ Unified Manager

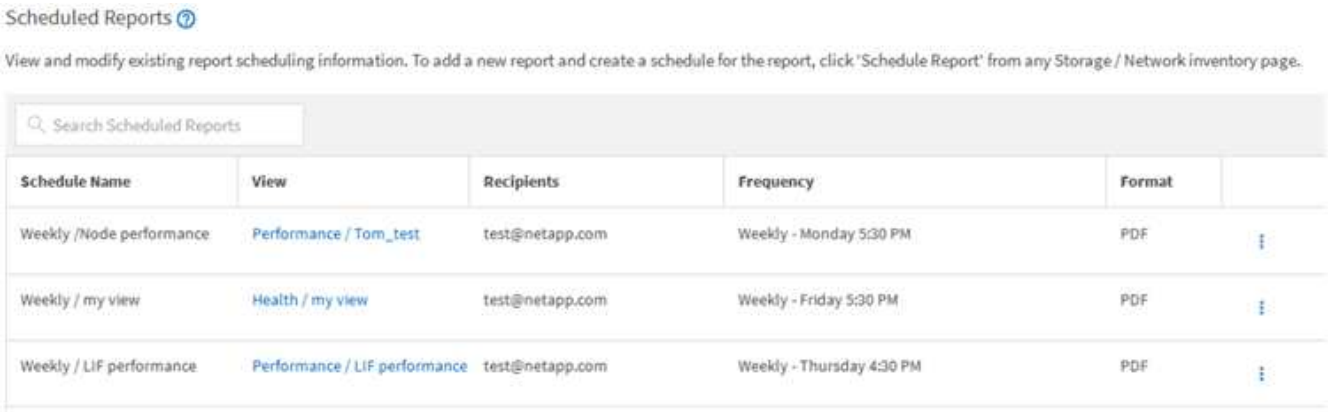
After reports are scheduled, you can edit them on the Report Schedules page.

Before you begin

- You must have the Application Administrator or Storage Administrator role.

Steps

1. In the left navigation pane, click **Storage Management** > **Report Schedules**.



If you have the appropriate permissions you can alter any report and its schedule in the system.

2. Click the more icon ⋮ for the schedule you want to change.
3. Click **Edit**.

4. You can change the **Schedule Name**, **Recipient** list, **Frequency**, and **Format** for the report schedule.
5. When done, click the check mark to save your changes.

Delete scheduled reports in Active IQ Unified Manager

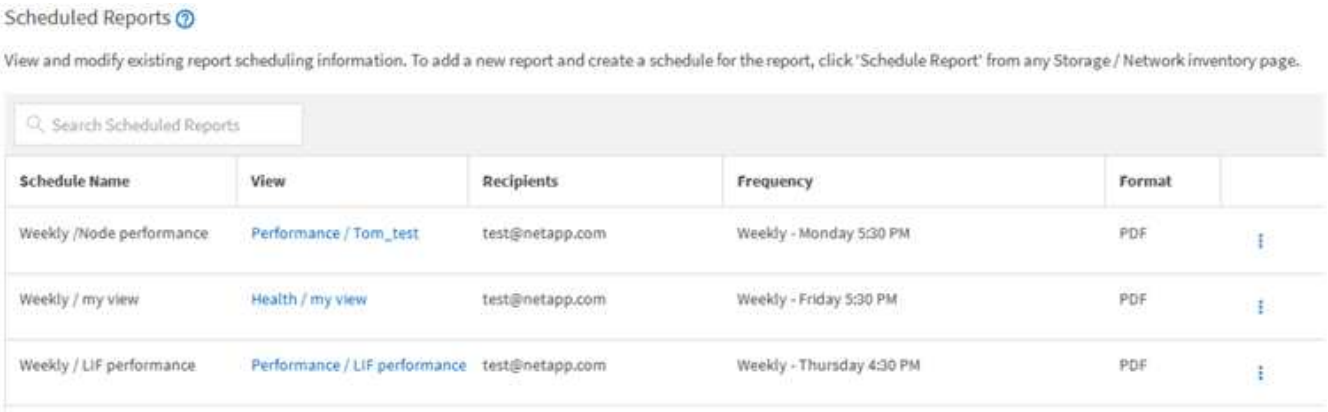
After reports are scheduled, you can delete them from the Report Schedules page.

Before you begin

- You must have the Application Administrator or Storage Administrator role.

Steps

1. In the left navigation pane, click **Storage Management** > **Report Schedules**.



If you have the appropriate permissions you can remove any report and its schedule in the system.

2. Click the more icon ⋮ for the schedule you want to remove.
3. Click **Delete**.
4. Confirm your decision.

   The scheduled report is removed from the list and will no longer be generated and distributed on the set schedule.

If you delete a custom view from the inventory page, any custom Excel files or scheduled reports that use that view are also deleted.

Sample custom reports

Understand sample custom reports in Active IQ Unified Manager

These sample custom reports are commonly used to help you identify potential problems and respond to potential issues before problems occur.

The list of reports in this section is not exhaustive and will grow over time. You can suggest custom reports to add to this section by providing documentation feedback.

You must have the Application Administrator or Storage Administrator role to manage reports.

Customize cluster storage reports

Create a report to view capacity by cluster model in Active IQ Unified Manager

You can create a report to analyze storage capacity and utilization of clusters based on storage system model.

Before you begin

- You must have the Application Administrator or Storage Administrator role.

Use the following steps to create a custom view that displays capacity by cluster model, and then schedule a report to be generated for that view.

Steps

1. In the left navigation pane, click **Storage** > **Clusters**.
2. In the View menu, select **Capacity** > **All Clusters**.
3. Select **Show/Hide** to remove any columns, such as “Cluster FQDN” and “OS Version”, that you do not want in the report.
4. Drag the “Total Raw Capacity”, the “Model/Family”, and the three aggregate columns near the “Cluster” column.
5. Click the top of the “Model/Family” column to sort the results by cluster type.
6. Save the view with a specific name that reflects what the view is showing, for example, “Capacity by Cluster Model”.
7. Click the **Scheduled Reports** button on the inventory page.
8. Click **Add Schedule** to add a new row to the **Report Schedules** page so that you can define the schedule characteristics for the new report.
9. Enter a name for the report schedule and complete the other report fields, then click the check mark (✓) at the end of the row.

   The report is sent immediately as a test. After that, the report generates and is sent by email to the recipients listed using the specified frequency.

Based on the results shown in the report, you might want to add more capacity to certain clusters, or upgrade older cluster models.

Identify clusters with unallocated LUN capacity in Active IQ Unified Manager

You can create a report to find the clusters with the most unallocated LUN capacity, greater than .5 TB, to help identify where you can add additional workloads.

Before you begin

- You must have the Application Administrator or Storage Administrator role.

Use the following steps to create a custom view that displays clusters with the most unallocated LUN capacity, and then schedule a report to be generated for that view.

Steps

1. In the left navigation pane, click **Storage** > **Clusters**.
2. In the View menu, select **Capacity** > **All Clusters**.
3. Select **Show/Hide** to remove any columns you do not want in the report.
4. Drag the “Unallocated LUN Capacity” column near the “HA Pair” column.
5. Click the filter icon, add the following filter, and then click **Apply Filter**:
   - Unallocated LUN capacity greater than 0.5 TB
6. Click the top of the “Unallocated LUN Capacity” column to sort the results by the greatest amount of unallocated LUN capacity.
7. Save the view with a specific name that reflects what the view is showing, for example “Most unallocated LUN capacity” and click the check mark (✓).
8. Click the **Scheduled Reports** button on the inventory page.
9. Click **Add Schedule** to add a new row to the Report Schedules page so that you can define the schedule characteristics for the new report.
10. Enter a name for the report schedule and complete the other report fields, then click the check mark (✓) at the end of the row.

    The report is sent immediately as a test. After that, the report generates and is sent by email to the recipients listed using the specified frequency.

Based on the results shown in the report, you might want to use the unallocated LUN capacity of the cluster.

Create a report to view HA pairs with the most available capacity in Active IQ Unified Manager

You can create a report to find the high availability (HA) pairs with the most capacity to provision new volumes and LUNs.

Before you begin

- You must have the Application Administrator or Storage Administrator role.

Use the following steps to create a custom view that displays HA pairs sorted by the most available capacity to provision new volumes and LUNs, and then schedule a report to be generated for that view.

Steps

1. In the left navigation pane, click **Storage** > **Clusters**.
2. In the View menu, select **Capacity** > **All Clusters**.
3. Select **Show/Hide** to remove any columns you do not want in the report.
4. Drag the “Aggregate Unused Capacity” column near the “HA Pair” column.
5. Click the filter icon, add the following filter, and then click **Apply Filter**:
   - Aggregate Unused Capacity greater than 0.5 TB
6. Click the top of the “Aggregate Unused Capacity” column to sort the results by the greatest amount of unused aggregate capacity.
7. Save the view with a specific name that reflects what the view is showing, for example, “Least used aggregate capacity”, and click the check mark (✓).
8. Click the **Scheduled Reports** button on the inventory page.
9. Click **Add Schedule** to add a new row to the Report Schedules page so that you can define the schedule

characteristics for the new report.

10. Enter a name for the report schedule and complete the other report fields, then click the check mark (✓) at the end of the row.

    The report is sent immediately as a test. After that, the report generates and is sent by email to the recipients listed using the specified frequency.

Based on the results shown in the report, you might want to balance HA pairs based on aggregate capacity.

Create a report to view nodes running older versions of ONTAP in Active IQ Unified Manager

You can create a report to display the version of ONTAP software that is installed on all cluster nodes so that you can see which nodes you should upgrade.

Before you begin

- You must have the Application Administrator or Storage Administrator role.

Use the following steps to create a custom view that displays nodes running older versions of ONTAP, and then schedule a report to be generated for that view.

Steps

1. In the left navigation pane, click **Storage** > **Nodes**.
2. Select **Show/Hide** to remove any columns that you do not want in the report.
3. Drag the “OS Version” column near the “Node” column.
4. Click the top of the “OS Version” column to sort the results by oldest version of ONTAP.
5. Save the view with a specific name that reflects what the view is showing, for example, “Nodes by ONTAP version”.
6. Click the **Scheduled Reports** button on the inventory page.
7. Click **Add Schedule** to add a new row to the Report Schedules page so that you can define the schedule characteristics for the new report.
8. Enter a name for the report schedule and complete the other report fields, then click the check mark (✓) at the end of the row.

   The report is sent immediately as a test. After that, the report generates and is sent by email to the recipients listed using the specified frequency.

Based on the results shown in the report, you might want to upgrade nodes running older versions of ONTAP.

Customize aggregate capacity reports

Create a report to view aggregates reaching full capacity in Active IQ Unified Manager

You can create a report to find the aggregates that are reaching full capacity so that you can add more capacity or move workloads to other aggregates.

Before you begin

- You must have the Application Administrator or Storage Administrator role.

Use the following steps to create a custom view that displays aggregates reaching full capacity, then schedule

a report to be generated for that view.

Steps

1. In the left navigation pane, click **Storage** > **Aggregates**.
2. In the View menu, select **Capacity** > **All Aggregates**.
3. Select **Show/Hide** to remove any columns you do not want in the report.
4. Click the filter icon, add the following filter, and then click **Apply Filter**:
   - Days to Full less than 45 Days
5. Click the top of the “Days to Full” column to sort the results by the least number of days remaining to reach full capacity.
6. Save the view with a specific name that reflects what the view is showing, for example, “Days to full aggregate capacity”, and click the check mark (✓).
7. Click the **Scheduled Reports** button on the inventory page.
8. Click **Add Schedule** to add a new row to the **Report Schedules** page so that you can define the schedule characteristics for the new report.
9. Enter a name for the report schedule and complete the other report fields, then click the check mark (✓) at the end of the row.

   The report is sent immediately as a test. After that, the report generates and is sent by email to the recipients listed using the specified frequency.

Based on the results shown in the report, you might want to increase storage on aggregates reaching full capacity. Additionally, you might want to increase the Days Until Full capacity threshold to more than the default 7 days so that you receive events that provide more time to react to space getting low on aggregates.

Create a report to view aggregates that are 80% or more full in Active IQ Unified Manager

You can create a report to highlight the aggregates that are 80% or more full.

Before you begin

- You must have the Application Administrator or Storage Administrator role.

Use the following steps to create a custom view that displays aggregates that are 80% or more full, then schedule a report to be generated for that view.

Steps

1. In the left navigation pane, click **Storage** > **Aggregates**.
2. In the View menu, select **Capacity** > **All Aggregates**.
3. Select **Show/Hide** to remove any columns you do not want in the report.
4. Drag the “Available Data %” and the “Used Data %” columns near the “Aggregate” column.
5. Click the filter icon, add the following filters, and then click **Apply Filter**:
   - Used Data % is greater than 80%
6. Click the top of the “Used Data %” column to sort the results by capacity percent.
7. Save the view with a specific name that reflects what the view is showing, for example, “Aggregates nearing full”, and click the check mark (✓) .
8. Click the **Scheduled Reports** button on the inventory page.

9. Click **Add Schedule** to add a new row to the Report Schedules page so that you can define the schedule characteristics for the new report.
10. Enter a name for the report schedule and complete the other report fields, then click the check mark (✓) at the end of the row.

    The report is sent immediately as a test. After that, the report generates and is sent by email to the recipients listed using the specified frequency.

Based on the results shown in the report, you might want to move some data from certain aggregates.

Create a report to view aggregates that are overcommitted in Active IQ Unified Manager

You can create a report to analyze storage capacity and use of aggregates, and to view aggregates that are overcommitted.

Before you begin

- You must have the Application Administrator or Storage Administrator role.

Use the following steps to create a custom view that displays aggregates that are exceeding the overcommitted threshold, then schedule a report to be generated for that view.

Steps

1. In the left navigation pane, click **Storage** > **Aggregates**.
2. In the View menu, select **Capacity** > **All Aggregates**.
3. Select **Show/Hide** to remove any columns that you do not want in the report.
4. Drag the “Overcommitted Capacity %” column near the “Aggregate” column.
5. Click the filter icon, add the following filters, and then click **Apply Filter**:
   - Overcommitted Capacity % is greater than 100%
6. Click the top of the “Overcommitted Capacity %” column to sort the results by capacity percent.
7. Save the view with a specific name that reflects what the view is showing, for example, “Aggregates overcommitted”, and click the check mark (✓).
8. Click the **Scheduled Reports** button on the inventory page.
9. Click **Add Schedule** to add a new row to the Report Schedules page so that you can define the schedule characteristics for the new report.
10. Enter a name for the report schedule and complete the other report fields, then click the check mark (✓) at the end of the row.

    The report is sent immediately as a test. After that, the report generates and is sent by email to the recipients listed using the specified frequency.

Based on the results shown in the report, you might want to add some more capacity to aggregates, or move some data from certain aggregates.

Customize volume capacity reports

Identify volumes nearing capacity with Snapshot Autodelete disabled in Active IQ Unified Manager

You can create a report that contains the list of volumes that are approaching full capacity

with the Snapshot Autodelete feature disabled. The results can help identify volumes where you might want to configure Snapshot Autodelete.

Before you begin

- You must have the Application Administrator or Storage Administrator role.

Follow the steps below to create a custom view that displays the required columns in the correct order, and then schedule a report to be generated for that view.

Steps

1. In the left navigation pane, click **Storage** > **Volumes**.
2. In the View menu, select **Capacity** > **All Volumes**.
3. Select **Show/Hide** to remove any columns you do not want in the report.
4. Drag and drop the “Snapshot Autodelete” and “Days To Full” columns near the “Available Data Capacity” column.
5. Click the filter icon, add the following two filters, and then click **Apply Filter**:
   - Days To Full less than 30 days
   - Snapshot Autodelete is Disabled
6. Click the top of the **Days To Full** column so that the volumes with the fewest remaining days appear at the top of the list.
7. Save the view with a specific name that reflects what the view is showing, for example “Vols near capacity”.
8. Click the **Scheduled Reports** button on the inventory page.
9. Enter a name for the report schedule and complete the other report fields, then click the check mark (✓) at the end of the row.

   The report is sent immediately as a test. After that, the report generates and is sent by email to the recipients listed using the specified frequency.

Based on the results shown in the report, you might want to enable Snapshot Autodelete on the volumes or find a way to increase the available space.

Identify volumes with thin provisioning disabled in Active IQ Unified Manager

When a volume is not thin provisioned it takes up the full amount of space on the disk as defined when the volume was created. Identifying volumes that have thin provisioning disabled helps you decide whether you want to enable thin-provisioning on certain volumes.

Before you begin

- You must have the Application Administrator or Storage Administrator role.

Follow the steps below to create a custom view that displays the required columns in the correct order, and then schedule a report to be generated for that view.

Steps

1. In the left navigation pane, click **Storage** > **Volumes**.
2. In the View menu, select **Capacity** > **All Volumes**.

3. Select **Show/Hide** to remove any columns you do not want in the report.
4. Drag and drop the “Used Data %” and the “Thin Provisioned” columns near the “Available Data Capacity” column.
5. Click the filter icon, add the following filter, **Thin provisioned** is **No**, and then click **Apply Filter**.
6. Click the top of the “Used Data %” column to sort the results so that the volumes with the highest percentage appear at the top of the list.
7. Save the view with a name to reflect what the view is showing, for example “Vols no thin provisioning”.
8. Click the **Scheduled Reports** button on the inventory page.
9. Click **Add Schedule** to add a new row to the **Report Schedules** page so you can define the schedule characteristics for the new report.
10. Enter a name for the report schedule and complete the other report fields, then click the checkmark (✓) at the end of the row.

    The report is sent immediately as a test. After that, the report generates and is sent by email to the recipients listed using the specified frequency.

Based on the results shown in the report, you might want to enable thin provisioning on certain volumes.

Identify volumes that should move data to the cloud tier in Active IQ Unified Manager

You can create a report that contains the list of volumes that currently reside on FabricPool aggregates, that have a cloud recommendation of Tier, and that have a large amount of cold data. This report can help you decide if you should change the tiering policy for certain volumes to “auto” or “all” to offload more cold (inactive) data to the cloud tier.

Before you begin

- You must have the Application Administrator or Storage Administrator role.
- You must have configured FabricPool aggregates and have volumes on those aggregates.

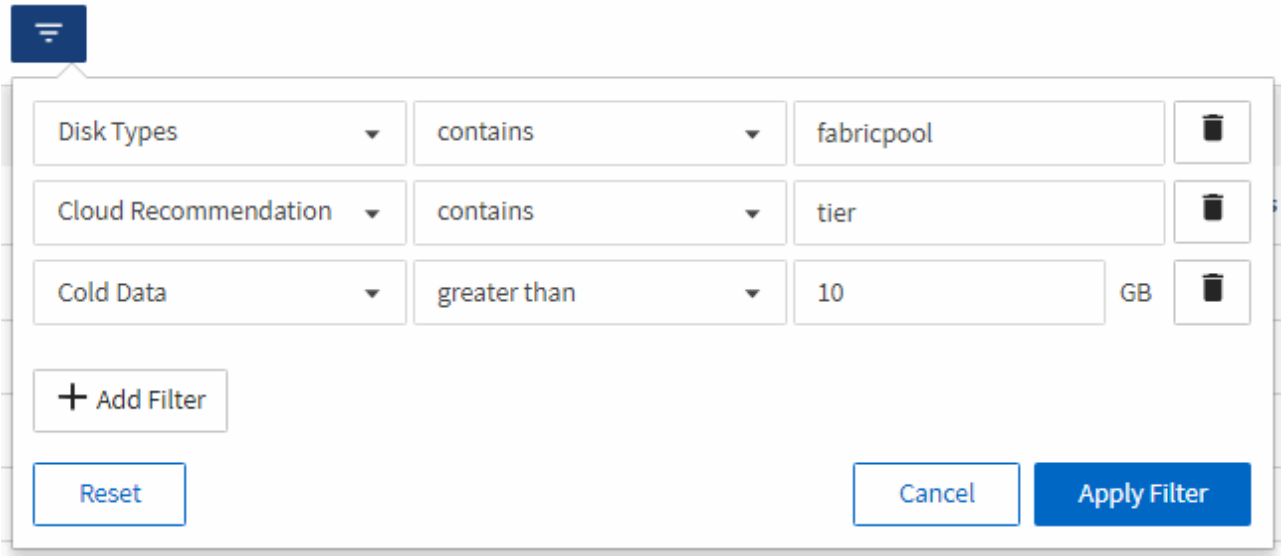
Follow the steps below to create a custom view that displays the required columns in the correct order, and then schedule a report to be generated for that view.

Steps

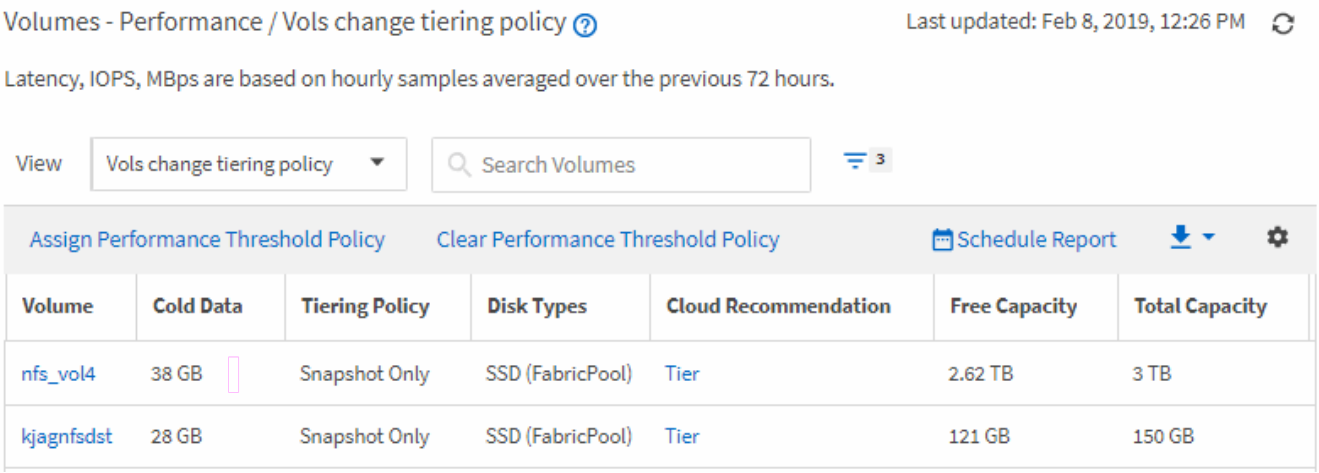
1. In the left navigation pane, click **Storage** > **Volumes**.
2. In the View menu, select **Performance** > **All Volumes**.
3. In the column chooser, make sure that the “Disk Type” column appears in the view.

   Add or remove other columns to create a view that is important for your report.

4. Drag and drop the “Disk Type” column near the “Cloud Recommendation” column.
5. Click the filter icon, add the following three filters, and then click **Apply Filter**:
   - Disk Type contains FabricPool
   - Cloud Recommendation contains tier
   - Cold Data greater than 10 GB



6. Click the top of the Cold Data column so that the volumes with the most cold data appear at the top of the view.
7. Save the view with a name to reflect what the view is showing, for example “Vols change tiering policy”.



8. Click the **Scheduled Reports** button on the inventory page.
9. Click **Add Schedule** to add a new row to the Report Schedules page so that you can define the schedule characteristics for the new report.
10. Enter a name for the report schedule and complete the other report fields, then click the check mark (✓) at the end of the row.

    The report is sent immediately as a test. After that, the report generates and is sent by email to the recipients listed using the specified frequency.

Based on the results shown in the report, you might want to use System Manager or the ONTAP CLI to change the tiering policy to “auto” or “all” for certain volumes to offload more cold data to the cloud tier.

Customize Qtree capacity reports

Create a report to view qtrees that are nearly full in Active IQ Unified Manager

You can create a report to analyze storage capacity and utilization of qtrees, and to view

qtrees that are nearly full.

Before you begin

- You must have the Application Administrator or Storage Administrator role.

Use the following steps to create a custom view that displays qtrees that are nearly full, and then schedule a report to be generated for that view.

Steps

1. In the left navigation pane, click **Storage** > **Qtrees**.
2. Select **Show/Hide** to remove any columns you do not want in the report.
3. Drag the “Disk Used %” column near the “Qtrees” column.
4. Click the filter icon, add the following filters, and then click **Apply Filter**:
   - Disk Used % is greater than 75%
5. Click the top of the “Disk Used %” column to sort the results by capacity percent.
6. Save the view with a specific name that reflects what the view is showing, for example, “Qtrees nearing full”, and click the check mark (✓).
7. Click the **Scheduled Reports** button on the inventory page.
8. Click **Add Schedule** to add a new row to the **Report Schedules** page so you can define the schedule characteristics for the new report.
9. Enter a name for the report schedule and complete the other report fields, then click the check mark (✓) at the end of the row.

   The report is sent immediately as a test. After that, the report generates and is sent by email to the recipients listed using the specified frequency.

Based on the results shown in the report, you might want to adjust the disk hard and soft limits (if set), or balance data across the qtrees.

Customize NFS share reports

Customize NFS share reports in Active IQ Unified Manager

You can customize NFS share reports to analyze information about NFS export policies and rules for volumes on your storage systems. For example, you can customize reports to display volumes with inaccessible mount paths and volumes with the default export policy.

Create a report to view volumes that have an inaccessible mount path in Active IQ Unified Manager

You can create a report to find volumes that have an inaccessible mount path.

Before you begin

- You must have the Application Administrator or Storage Administrator role.

Use the following steps to create a custom view for volumes that have an inaccessible mount path, and then schedule a report to be generated for that view.

Steps

1. In the left navigation pane, click **Storage** > **NFS Shares**.
2. Select **Show/Hide** to remove any columns that you do not want in the report.
3. Click the filter icon, add the following filter, and then click **Apply Filter**:
   - Mount Path Active is No
4. Save the view with a specific name that reflects what the view is showing, for example “Volumes with an inaccessible mount path” and click the check mark (✓).
5. Click the **Scheduled Reports** button on the inventory page.
6. Click **Add Schedule** to add a new row to the Report Schedules page so you can define the schedule characteristics for the new report.
7. Enter a name for the report schedule and complete the other report fields, then click the check mark (✓) at the end of the row.

   The report is sent immediately as a test. After that, the report generates and is sent by email to the recipients listed using the specified frequency.

Based on the results shown in the report, you might want to correct the inaccessible mount paths.

Create a report to view volumes using the default export policy in Active IQ Unified Manager

You can create a report to find volumes that are using the default export policy.

Before you begin

- You must have the Application Administrator or Storage Administrator role.

Use the following steps to create a custom view for volumes that are using the default export policy, and then schedule a report to be generated for that view.

Steps

1. In the left navigation pane, click **Storage** > **NFS Shares**.
2. Select **Show/Hide** to remove any columns that you do not want in the report.
3. Drag the “Export Policy” column near the “Volume” column.
4. Click the filter icon, add the following filter, and then click **Apply Filter**:
   - Export Policy contains default
5. Save the view with a specific name that reflects what the view is showing, for example “Volumes with a default export policy” and click the check mark (✓).
6. Click the **Scheduled Reports** button on the inventory page.
7. Click **Add Schedule** to add a new row to the Report Schedules page so you can define the schedule characteristics for the new report.
8. Enter a name for the report schedule and complete the other report fields, then click the check mark (✓) at the end of the row.

   The report is sent immediately as a test. After that, the report generates and is sent by email to the recipients listed using the specified frequency.

Based on the results shown in the report, you might want to configure a custom export policy.

Customize storage VM reports

Customize storage VM reports in Active IQ Unified Manager

You can create storage VM reports to analyze volume information and to view overall health and storage availability. For example, you can create reports to display SVMs reaching the maximum volume count and to analyze stopped SVMs.

Create a report to view storage VMs reaching the maximum volume limit in Active IQ Unified Manager

You can create a report to find SVMs that are reaching the maximum volume limit.

Before you begin

- You must have the Application Administrator or Storage Administrator role.

Use the following steps to create a custom view that displays storage VMs that are reaching the maximum volume limit, and then schedule a report to be generated for that view.

Steps

1. In the left navigation pane, click **Storage** > **Storage VMs**.
2. Select **Show/Hide** to remove any columns you do not want in the report.
3. Drag the “Volume Count” and “Maximum Allowed Volumes” columns near the “Storage VM” column.
4. Click the top of the “Maximum Allowed Volumes” column to sort the results by the highest number of volumes.
5. Save the view with a specific name that reflects what the view is showing, for example “SVMs reaching max volumes” and click the check mark (✓).
6. Click the **Scheduled Reports** button on the inventory page.
7. Click **Add Schedule** to add a new row to the **Report Schedules** page so that you can define the schedule characteristics for the new report.
8. Enter a name for the report schedule and complete the other report fields, then click the check mark (✓) at the end of the row.

   The report is sent immediately as a test. After that, the report generates and is sent by email to the recipients listed using the specified frequency.

Based on the results shown in the report, you might want to balance the volumes assigned to storage VMs or, if possible, use ONTAP System Manager to change the maximum allowed volumes.

Create a report to view stopped storage VMs in Active IQ Unified Manager

You can create a report to display a list of all stopped SVMs.

Before you begin

- You must have the Application Administrator or Storage Administrator role.

Use the following the steps to create a custom view that displays stopped storage VMs, and then schedule a report to be generated for that view.

Steps

1. In the left navigation pane, click **Storage** > **Storage VMs**.
2. In the View menu, select **Health** > **All Storage VMs**.
3. Select **Show/Hide** to remove any columns that you do not want in the report.
4. Drag the “State” column near the “Storage VM” column.
5. Click the filter icon, add the following filter, and then click **Apply Filter**:
   - State is Stopped
6. Save the view with a specific name that reflects what the view is showing, for example “Stopped SVMs” and click the check mark (✓).
7. Click the **Scheduled Reports** button on the inventory page.
8. Click **Add Schedule** to add a new row to the **Report Schedules** page so that you can define the schedule characteristics for the new report.
9. Enter a name for the report schedule and complete the other report fields, then click the check mark (✓) at the end of the row.

   The report is sent immediately as a test. After that, the report generates and is sent by email to the recipients listed using the specified frequency.

Based on the results shown in the report, you might want to investigate why the SVM is stopped to see if you should restart the stopped SVMs.

Customize volume relationship reports

Customize volume relationship reports in Active IQ Unified Manager

The Volume Relationships Inventory report enables you to analyze the storage inventory details in a cluster, understand the degree of protection that is required for volumes, and filter the volume details based on source of failure, pattern, and schedules.

Create a report to group volume relationships by source of failure in Active IQ Unified Manager

You can create a report that groups volumes by the reason the relationship is in an unhealthy state.

Before you begin

- You must have the Application Administrator or Storage Administrator role.

Use the following steps to create a custom view that groups volumes by source of failure, and then schedule a report to be generated for that view.

Steps

1. In the left navigation pane, click **Storage** > **Volumes**.
2. In the View menu, select **Relationship** > **All Relationships**.
3. Select **Show/Hide** to make sure the “Relationship Health” and “Unhealthy Reason” columns appear in the view.

   Add or remove other columns to create a view that is important for your report.

4. Drag the “Relationship Health” and “Unhealthy Reason” columns near the “State” column.
5. Click the filter icon, add the following filter, and then click **Apply Filter**:
    - Relationship health is bad
6. Click the top of the “Unhealthy Reason” column to group the volume relationships by source of failure.
7. Save the view with a specific name that reflects what the view is showing, for example, “Vol relationships by failure”.
8. Click the **Scheduled Reports** button on the inventory page.
9. Enter a name for the report schedule and complete the other report fields, then click the check mark (✓) at the end of the row.

    The report is sent immediately as a test. After that, the report generates and is sent by email to the recipients listed using the specified frequency.

Based on the results shown in the report, you can investigate the source and impact of each type of failure.

Create a report to group volume relationships by issue in Active IQ Unified Manager

You can create a report that groups volume relationships by issue.

Before you begin

- You must have the Application Administrator or Storage Administrator role.

Use the following steps to create a custom view that groups volume relationships by issue, and then schedule a report to be generated for that view.

Steps

1. In the left navigation pane, click **Storage** > **Volumes**.
2. In the View menu, select **Relationship** > **All Relationships**.
3. Select **Show/Hide** to remove any columns you do not want in the report.
4. Drag the “Unhealthy Reason” column near the “State” column.
5. Click the top of the “Unhealthy Reason” column to group the volumes by issue.
6. Save the view with a specific name that reflects what the view is showing, for example, “Vol relationships by issue”.
7. Click the **Scheduled Reports** button on the inventory page.
8. Enter a name for the report schedule and complete the other report fields, then click the check mark (✓) at the end of the row.

    The report is sent immediately as a test. After that, the report generates and is sent by email to the recipients listed using the specified frequency.

Based on the results shown in the report, you can investigate the source and impact of each type of issue.

Create a report to view volume transfer trends at time intervals in Active IQ Unified Manager

You can create a report that displays volume transfer trends at specific time intervals.

Before you begin

- You must have the Application Administrator or Storage Administrator role.

Use the following steps to create a custom view for volumes at specific time intervals, and then schedule a report to be generated for that view.

Steps

1. In the left navigation pane, click **Storage** > **Volumes**.
2. In the View menu, select **Relationship** > **Last 1 month Transfer Status**.
3. Select **Show/Hide** to remove any columns you do not want in the report.
4. Drag the Transfer Duration column near the “Operational Result” column.
5. Click the filter icon, add the following filter, and then click **Apply Filter**:
   - Transfer End Time in the last 7 Days
6. Click the top of the “Transfer Duration” column to sort the volumes by time interval.
7. Save the view with a specific name that reflects what the view is showing, for example “Volumes by duration”.
8. Click the **Scheduled Reports** button on the inventory page.
9. Enter a name for the report schedule, set the frequency as **Weekly**, and complete the other report fields, then click the check mark (✓) at the end of the row.

   The report is sent immediately as a test. After that, the report generates and is sent by email to the recipients listed using the specified frequency.

Based on the results shown in the report, you can investigate the transfer time intervals.

Create a report to view failed or successful volume transfers in Active IQ Unified Manager

You can create a report that displays the status of volume transfers. You can view both failed and successful volume transfers in this report.

Before you begin

- You must have the Application Administrator or Storage Administrator role.

Use the following steps to create a custom view to show which transfers failed and which were successful, and then schedule a report to be generated for that view.

Steps

1. In the left navigation pane, click **Storage** > **Volumes**.
2. In the View menu, select **Relationship** > **Last 1 month Transfer Status**.
3. Select **Show/Hide** to remove any columns that you do not want in the report.
4. Drag the “Operation Result” column near the “State” column.
5. Click the top of the “Operation Result” column to sort the volumes by the status.
6. Save the view with a specific name that reflects what the view is showing, for example “Volumes by transfer status”.
7. Click the **Scheduled Reports** button on the inventory page.
8. Enter a name for the report schedule and complete the other report fields, then click the check mark (✓) at the end of the row.

The report is sent immediately as a test. After that, the report generates and is sent by email to the recipients listed using the specified frequency.

Based on the results shown in the report, you can investigate the transfer status.

Create a report to view volume transfers based on transfer size in Active IQ Unified Manager

You can create a report to view volume transfers based on transfer size.

Before you begin

- You must have the Application Administrator or Storage Administrator role.

Use the following steps to create a custom view for volume transfers based on transfer size, and then schedule a report to be generated for that view.

Steps

1. In the left navigation pane, click **Storage** > **Volumes**.
2. In the View menu, select **Relationship** > **Last 1 month Transfer Rate**.
3. Click the top of the “Total Transfer Size” column to sort the volume transfers by size.
4. Save the view with a specific name that reflects what the view is showing, for example “Volumes by transfer size”.
5. Click the **Scheduled Reports** button on the inventory page.
6. Enter a name for the report schedule and complete the other report fields, then click the check mark (✓) at the end of the row.

   The report is sent immediately as a test. After that, the report generates and is sent by email to the recipients listed using the specified frequency.

Based on the results shown in the report, you can investigate the volume relationships by transfer size.

Create a report to view volume transfers grouped by day in Active IQ Unified Manager

You can create a report to view volume transfers grouped by day.

Before you begin

- You must have the Application Administrator or Storage Administrator role.

Use the following steps to create a custom view for volume transfers grouped by day, and then schedule a report to be generated for that view.

Steps

1. In the left navigation pane, click **Storage** > **Volumes**.
2. In the View menu, select **Relationship** > **Last 1 month Transfer Rate**.
3. Click the top of the “Day” column to sort the volume transfers by day.
4. Save the view with a specific name that reflects what the view is showing, for example “Volume transfers by day”.
5. Click the **Scheduled Reports** button on the inventory page.
6. Enter a name for the report schedule and complete the other report fields, then click the check mark (✓) at the end of the row.

The report is sent immediately as a test. After that, the report generates and is sent by email to the recipients listed using the specified frequency.

Based on the results shown in the report, you can investigate the volume transfers by day.

Customize volume performance reports

Create a report to identify volumes with high cold data on a non-FabricPool aggregate in Active IQ Unified Manager

You can create a report to display volumes with a high amount of cold data on a non-FabricPool aggregate. This can help you identify volumes that should be moved to a FabricPool aggregate.

Before you begin

- You must have the Application Administrator or Storage Administrator role.

Use the following steps to create a custom view for volumes with a high amount of cold data on a non-FabricPool-enabled aggregate, and then schedule a report to be generated for that view.

Steps

1. In the left navigation pane, click **Storage** > **Volumes**.
2. In the View menu, select **Performance** > **All Volumes**.
3. Select **Show/Hide** to make sure the “Disk Type” column appears in the view.

   Add or remove other columns to create a view that is important for your report.
4. Drag the “Disk Type” column near the “Cold Data” column.
5. Click the filter icon, add the following filter, and then click **Apply Filter**:
   - Cold Data greater than 100 GB
   - Disk Type contains SSD
6. Click the top of the “Disk Type” column to sort the volumes by disk type so that the disk type SSD (FabricPool) is at the bottom.
7. Save the view with a specific name that reflects what the view is showing, for example “Cold data vols not FabricPool”.
8. Click the **Scheduled Reports** button on the inventory page.
9. Enter a name for the report schedule and complete the other report fields, then click the check mark (✓) at the end of the row.

   The report is sent immediately as a test. After that, the report generates and is sent by email to the recipients listed using the specified frequency.

Based on the results shown in the report, you can find volumes that are good candidates to be moved to FabricPool aggregates.

Sample Microsoft Excel reports

Undestand sample Microsoft Excel reports in Active IQ Unified Manager

These sample Microsoft Excel reports are intended to introduce reporting options available using Excel advanced features.

The advanced functionality of Excel can create a wide array of reports specific to your needs. For full information on how to use Excel, see the product documentation.

You must have the Application Administrator or Storage Administrator role to manage reports.

Create a report to show an aggregate capacity table and chart in Active IQ Unified Manager

You can create a report to analyze capacity in an Excel file by using summed totals and the clustered column chart format.

Before you begin

- You must have the Application Administrator or Storage Administrator role.

Use the following steps to open a Health: All Aggregates view, download the view in Excel, create an available capacity chart, upload the customized Excel file, and schedule the final report.

Steps

1. In the left navigation pane, click **Storage** > **Aggregates**.
2. Select **Reports** > **Download Excel**.



   Depending on your browser, you might need to click **OK** to save the file.

3. If needed, click **Enable Editing**.
4. In Excel, open the downloaded file.
5. Create a new sheet (⊕) after the `data` sheet and name it **Total Data Capacity**.
6. Add the following columns on the new Total Data Capacity sheet:
   a. Total Data Capacity (GB)
   b. Committed Capacity (GB)

c. Used Data Capacity (GB)

d. Available Data Capacity (GB)

7. In the first row of each column, enter the following formula, making sure that it references the data sheet (data!), and references the correct column and row specifiers for the data captured (Total Data Capacity pulls data from column E, rows 2 through 20).

   a. =SUM(data!E$2:data!E$20)

   b. =SUM(data!F$2:data!F$50)

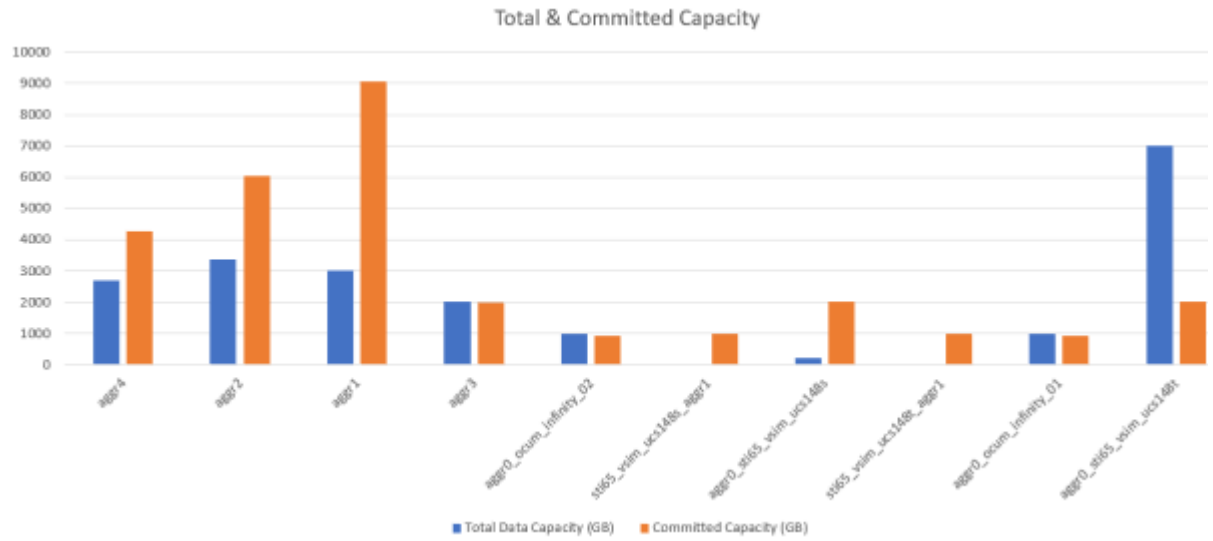
   c. =SUM(data!G$2:data!G$50)

   d. =SUM(data!H$2:data!H$50)

   The formula totals each column based on the current data.

Total Data Capacity (GB)	Committed Capacity (GB)	Used Data Capacity (GB)	Available Data Capacity (GB)
5380.31	6892.47	11764.27	3911.03

1. On the data sheet, select the **Total Data Capacity (GB)** and **Committed Capacity (GB)** columns.
2. Select **Recommended Charts** from the **Insert** menu and select the **Clustered Column** chart.
3. Right-click the chart and select **Move chart** to move the chart to the `Total Data Capacity` sheet.
4. Using the **Design** and **Format** menus, which are available when the chart is selected, you can customize how the chart looks.
5. When satisfied, save the file with your changes. Do not change the file name or location.

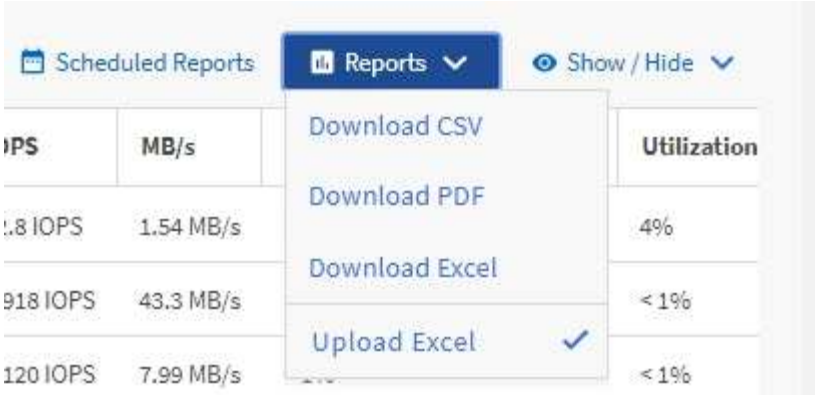


6. In Unified Manager, select **Reports** > **Upload Excel**.

   Make sure you are in the same view where you downloaded the Excel file.

7. Select the Excel file you have modified.
8. Click **Open**.
9. Click **Submit**.

   A check mark appears beside the **Reports** > **Upload Excel** menu item.



10. Click **Scheduled Reports**.
11. Click **Add Schedule** to add a new row to the Report Schedules page so that you can define the schedule characteristics for the new report.

    Select the **XLSX** format for the report.

12. Enter a name for the report schedule and complete the other report fields, then click the check mark (✓) at the end of the row.

    The report is sent immediately as a test. After that, the report generates and is sent by email to the recipients listed using the specified frequency.

Based on the results shown in the report, you might want to investigate how to best use the available capacity across your network.

Show aggregate total versus available capacity charts in Active IQ Unified Manager

You can create a report to analyze storage total and committed capacity in an Excel chart format.

Before you begin

- You must have the Application Administrator or Storage Administrator role.

Use the following steps to open a Health: All Aggregates view, download the view in Excel, create a total and committed capacity chart, upload the customized Excel file, and schedule the final report.

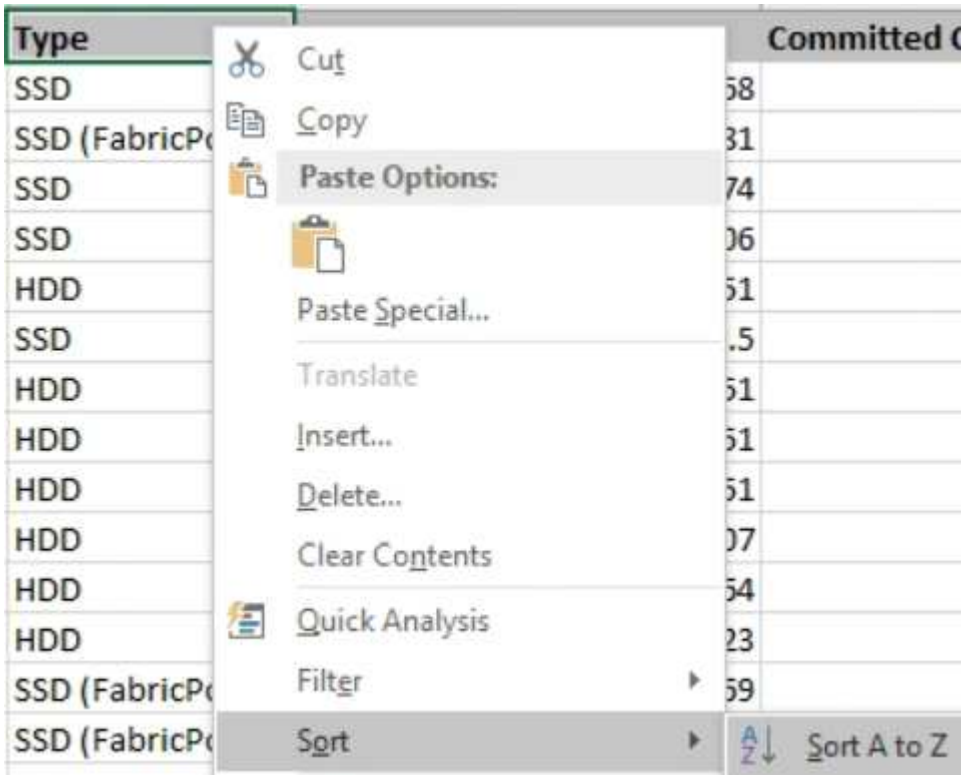
Steps

1. In the left navigation pane, click **Storage** > **Aggregates**.
2. Select **Reports** > **Download Excel**.



Depending on your browser, you might need to click **OK** to save the file.

3. In Excel, open the downloaded file.
4. If needed, click **Enable Editing**.
5. On the data sheet, right-click the Type column and select **Sort** > **Sort A to Z**.



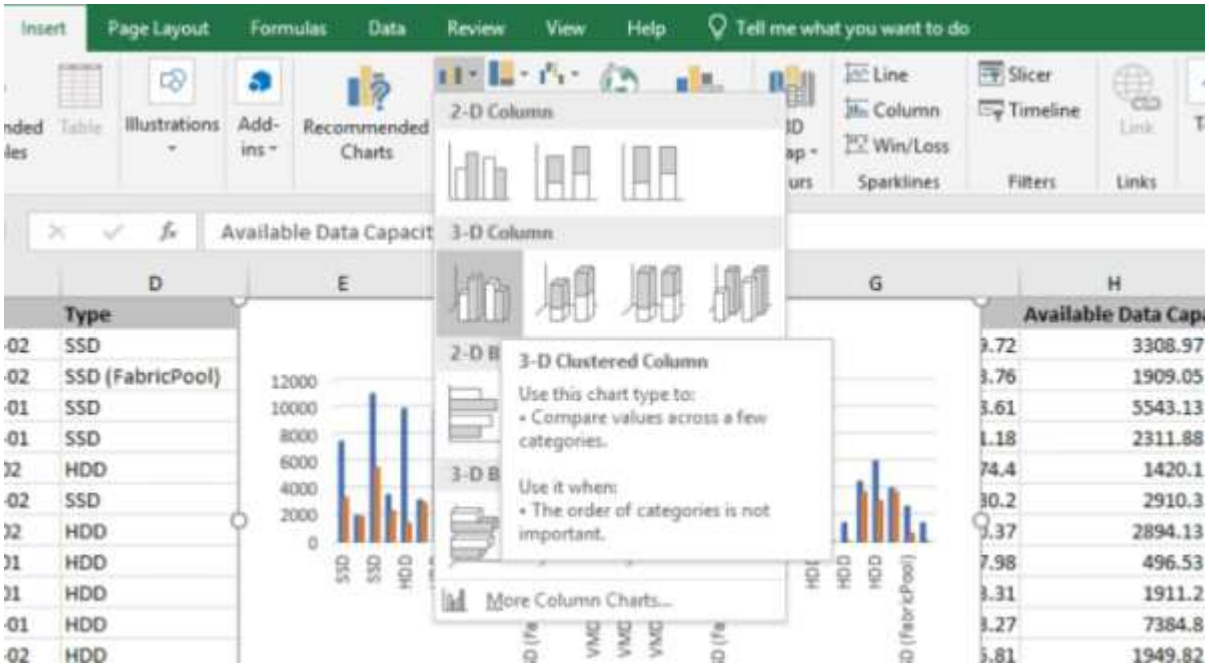
This will arrange your data by storage type, such as:

- HDD
- Hybrid
- SSD
- SSD (FabricPool)

6. Select the `Type`, `Total Data Capacity`, and `Available Data Capacity` columns.

7. On the **Insert** menu, select a `3-D column` chart.

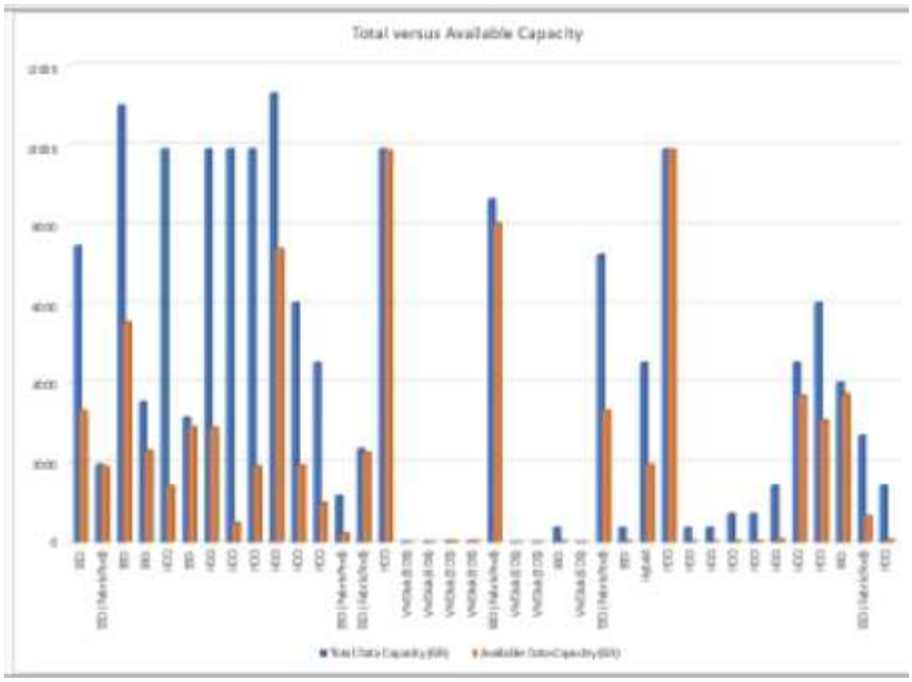
   The chart appears on the data sheet.



8. Right-click the chart and select **Move Chart**.
9. Select **New sheet** and name the sheet **Total Storage Charts**.

   Make sure the new sheet appears after the info and data sheets.

10. Name the chart title **Total versus Available Capacity**.
11. Using the **Design** and **Format** menus, available when the chart is selected, you can customize how the chart looks.
12. When satisfied, save the file with your changes. Do not change the file name or location.

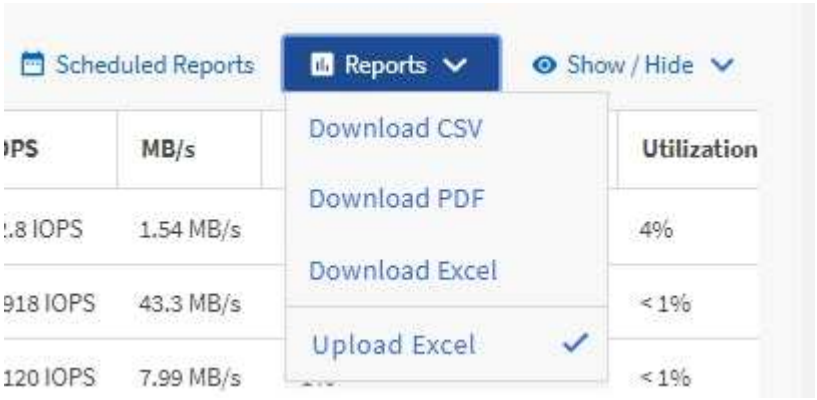


13. In Unified Manager, select **Reports** > **Upload Excel**.

    Make sure you are in the same view where you downloaded the Excel file.

14. Select the Excel file you have modified.
15. Click **Open**.
16. Click **Submit**.

    A check mark appears beside the **Reports** > **Upload Excel** menu item.



17. Click **Scheduled Reports**.
18. Click **Add Schedule** to add a new row to the **Report Schedules** page so that you can define the schedule characteristics for the new report.

    Select the **XLSX** format for the report.

19. Enter a name for the report schedule and complete the other report fields, then click the check mark (✓) at the end of the row.

The report is sent immediately as a test. After that, the report generates and is sent by email to the recipients listed using the specified frequency.

Based on the results shown in the report, you might want to balance the load on your aggregates.

Create a report to view available volume capacity charts in Active IQ Unified Manager

You can create a report to analyze available volume capacity in an Excel chart.

Before you begin

- You must have the Application Administrator or Storage Administrator role.

Use the following steps to open a Health: All Volumes view, download the view in Excel, create an available capacity chart, upload the customized Excel file, and schedule the final report.

Steps

1. In the left navigation pane, click **Storage** > **Volumes**.
2. Select **Reports** > **Download Excel**.



   Depending on your browser, you might need to click **OK** to save the file.

3. If needed, click **Enable Editing**.
4. In Excel, open the downloaded file.
5. On the `data` sheet, select the data that you want to use in the `Volume` and `Available Data` % columns.
6. On the **Insert** menu, select a `3-D piechart`.

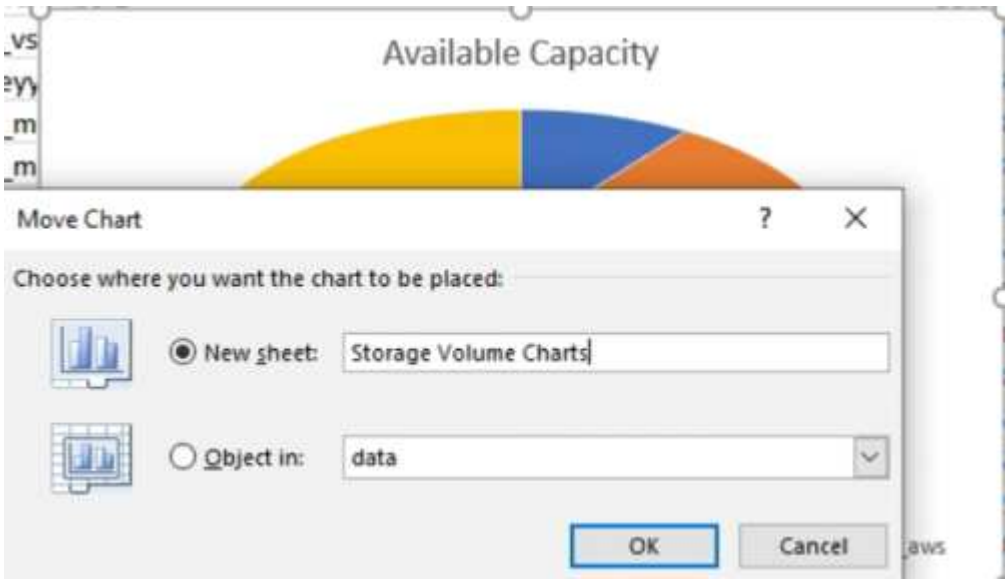
   The chart shows which volumes have the largest available space. The chart appears on the data sheet.

   Depending on your network configuration, selecting the entire columns or too many data rows might make your pie chart unreadable. This sample uses the 3-D pie chart, but you can use any chart type. Use the chart that best displays the data you want to capture.

7. Name the chart title **Available Capacity**.
8. Right-click the chart and select **Move Chart**.

9. Select **New sheet** and name the sheet **Storage Volume Charts**.

   Make sure the new sheet appears after the info and data sheets.

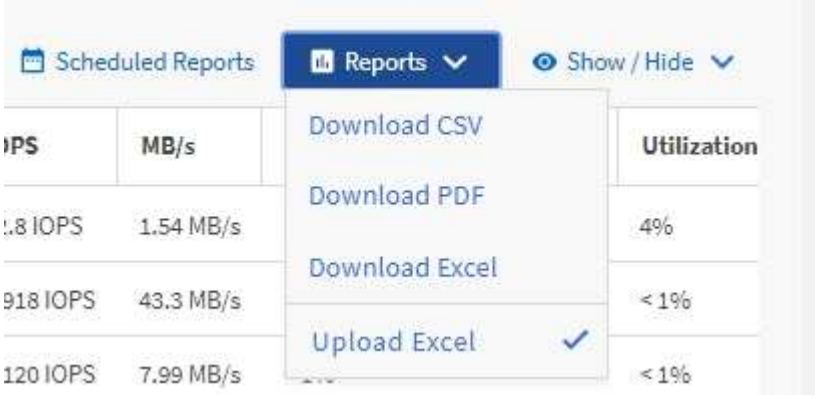


10. Using the **Design** and **Format** menus, available when the chart is selected, you can customize how the chart looks.
11. When satisfied, save the file with your changes.
12. In Unified Manager, select **Reports** > **Upload Excel**.

    Make sure you are in the same view where you downloaded the Excel file.

13. Select the Excel file you have modified.
14. Click **Open**.
15. Click **Submit**.

    A check mark appears beside the **Reports** > **Upload Excel** menu item.



16. Click **Scheduled Reports**.
17. Click **Add Schedule** to add a new row to the **Report Schedules** page so that you can define the schedule characteristics for the new report.
18. Enter a name for the report schedule and complete the other report fields, then click the check mark (✓) at

the end of the row.

Select the **XLSX** format for the report.

The report is sent immediately as a test. After that, the report generates and is sent by email to the recipients listed using the specified frequency.

Based on the results shown in the report, you might want to balance the load on your volumes.

Create a report to view aggregates with the most available IOPS in Active IQ Unified Manager

This report shows which aggregates have the most available IOPS per aggregate type on which you can provision new workloads.

Before you begin

- You must have the Application Administrator or Storage Administrator role.

Use the following steps to open a Health: All Volumes view, download the view in Excel, create an available capacity chart, upload the customized Excel file, and schedule the final report.

Steps

1. In the left navigation pane, click **Storage** > **Aggregates**.
2. Select **Performance: All Aggregates** from the **View** drop-down.
3. Select **Show/Hide** to show the `Available IOPS` column and hide the `Cluster FQDN`, `Inactive Data Reporting`, and `Threshold Policy` columns.
4. Drag and drop the `Available IOPS` and `Free Capacity` columns next to the `Type` column.
5. Name and save the custom view `Available IOPS Per Aggr`.
6. Select **Reports** > **Download Excel**.



Depending on your browser, you might need to click **OK** to save the file.

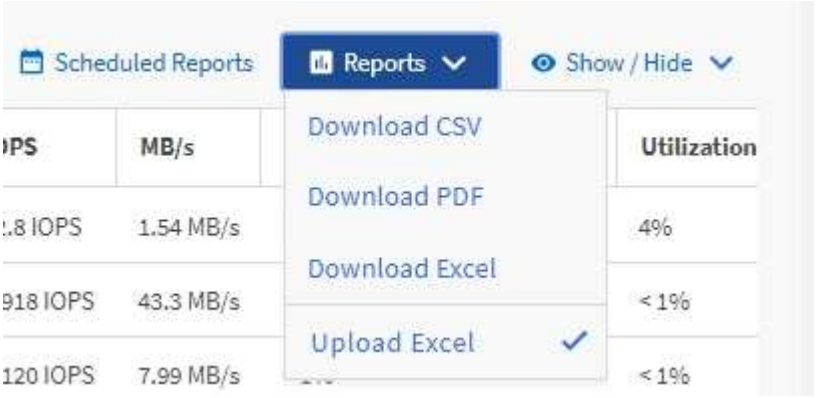
7. If needed, click **Enable Editing**.
8. In Excel, open the downloaded file.

9. On the data sheet, click the little triangle at the top-left of the sheet to select the entire sheet.
10. On the **Data** ribbon, select **Sort** from the `Sort & Filter area.`
11. Set the following sort levels:
    a. Specify the **Sort by** as `Available IOPS` (IOPS), the **Sort On** as `Cell Values,` and the **Order** as `Largest to Smallest.`
    b. Click **Add Level**.
    c. Specify the **Sort by** as `Type`, the **Sort On** as `Cell Values,` and the **Order** as `Z to A.`
    d. Click **Add Level**.
    e. Specify the **Sort by** as `Free Capacity (GB),` the **Sort On** as `Cell Values,` and the **Order** as `Largest to Smallest.`
    f. Click **OK**.
12. Save and close the Excel file.
13. In Unified Manager, select **Reports** > **Upload Excel**.

    Make sure you are in the same view where you downloaded the Excel file.

14. Select the Excel file you have modified, in this case `performance-aggregates-<date>.xlsx.`
15. Click **Open**.
16. Click **Submit**.

    A check mark appears beside the **Reports** > **Upload Excel** menu item.



17. Click **Scheduled Reports**.
18. Click **Add Schedule** to add a new row to the Report Schedules page so that you can define the schedule characteristics for the new report.
19. Enter a name for the report schedule and complete the other report fields, then click the check mark (✓) at the end of the row.

    Select the **XLSX** format for the report.

    The report is sent immediately as a test. After that, the report generates and is sent by email to the recipients listed using the specified frequency.

Based on the results shown in the report, you might want to provision new workloads on the aggregates that

have the greatest available IOPS.

Manage storage using REST APIs

Getting started with Active IQ Unified Manager

Get started with Active IQ Unified Manager REST APIs

Active IQ Unified Manager provides a set of APIs to manage your storage resources on the supported storage systems through a RESTful web service interface for any third-party integration.

In these topics, you will find information about Unified Manager APIs, sample workflows to resolve specific issues, and some sample codes. Using this information, you can create RESTful clients of NetApp Manageability software solutions for managing NetApp systems. The APIs are based on the Representational State Transfer (REST) architectural style. All the four REST operations Create, Read, Update, and Delete (also known as CRUD) are supported.

See Active IQ Unified Manager for additional resources and details about the benefits of the Active IQ Unified Manager REST API.

Identify intended users of Active IQ Unified Manager REST APIs

The topics here are intended for developers creating applications that interface with the Active IQ Unified Manager software through REST APIs.

Storage administrators and architects, can refer to this information to gain a basic understanding of how the Unified Manager REST APIs can be used to build client applications to manage and monitor NetApp storage systems.

You should use this information if you want to use the storage provider, ONTAP cluster, and management administration APIs for managing your storage.

> You must have one of the following roles: Operator, Storage Administrator, or Application Administrator. You must know the IP address or fully qualified domain name of the Unified Manager server on which you want to execute the REST APIs.

Manage Active IQ Unified Manager API access and categories

The Active IQ Unified Manager APIs enable you to manage and provision storage objects in your environment. You can also access the Unified Manager web UI to perform some of these functions.

Construct a URL to directly access REST APIs

You can access the REST APIs directly through a programming language, such as Python, C#, C++, JavaScript, and so forth. Enter the host name or IP address and the URL to access the REST APIs in the format

https://<hostname>/api

> The default port is 443. You can configure the port as required by your environment.

Access the online API documentation page

You can access the *API Documentation* reference content page that is packaged along with the product to display the API documentation, as well as to manually issue an API call (on the interface, for example, Swagger). You can access this documentation on clicking the **Menu Bar** > **Help button** > **API Documentation**

Alternatively, enter the host name or IP address and the URL to access the REST API page in the format

https://<hostname>/docs/api/

Categories

The API calls are organized into functionally based on the areas or categories. To locate a specific API, click the applicable API category.

The REST APIs provided with Unified Manager help you to perform administrative, monitoring, and provisioning functions. The APIs are grouped under the following categories.

- **datacenter**

  This category contains the APIs that help you in datacenter storage management and analytics using tools, such as Work Flow Automation and Ansible. The REST APIs under this category provide information about the clusters, nodes, aggregates, volumes, LUNs, file shares, namespaces, and other elements in your data center.

- **management-server**

  The APIs under the **management-server** category contain the `jobs, system,` and `events` APIs. Jobs are operations that are scheduled for asynchronous execution related to managing of storage objects or workloads on Unified Manager. The `events` API returns events in your data center, and the `system` API returns the Unified Manager instance details.

- **storage-provider**

  This category contains all of the provisioning APIs required for managing and provisioning file shares, LUNs, Performance Service Levels, and Storage Efficiency Policies. The APIs also enable you to configure access endpoints, Active Directories, as well as assign Performance Service Levels and Storage Efficiency Policies on storage workloads.

- **administration**

  This category contains the APIs used for running administrative tasks, such as maintaining backup settings, viewing trust store certificates for the Unified Manager datasources, and managing ONTAP clusters as datasources for Unified Manager.

- **gateway**

  Unified Manager enables you to invoke ONTAP REST APIs through the APIs under the gateway category and manage the storage objects in your data center.

- **security**

  This category contains APIs for managing Unified Manager users.

Understand REST services offered in Active IQ Unified Manager

You should be aware of the REST services and operations offered, before you start using the Active IQ Unified Manager APIs.

The provisioning and administrative APIs that are used for configuring the API server support the read (GET) or write (POST, PATCH, DELETE) operations. The following are some examples of the GET, PATCH, POST, and DELETE operations that are supported by the APIs:

- Example for GET: `GET /datacenter/cluster/clusters` retrieves cluster details in your data center. The maximum number of records that is returned by the `GET` operation is 1000.

  The APIs enable you to filter, sort, and order the records by supported attributes.

- Example for POST: `POST /datacenter/svm/svms` creates a custom Storage Virtual Machine (SVM).
- Example for PATCH: `PATCH /datacenter/svm/svms/{key}` modifies the properties of an SVM, using its unique key.
- Example for DELETE: `DELETE /storage-provider/access-endpoints/{key}` deletes an access endpoint from a LUN, SVM, or file share by using its unique key.

The REST operations that can be performed by using the APIs depend on the role of the Operator, Storage Administrator, or Application Administrator user.

User role	Supported REST method
Operator	Read-only access to data. Users with this role can run all GET requests.
Storage Administrator	Read access to all data. Users with this role can run all GET requests. Additionally, they have write access (to run PATCH, POST, and DELETE requests) to perform specific activities, such as managing, storage service objects, and storage management options.
Application Administrator	Read and write access to all data. Users with this role can run GET, PATCH, POST, and DELETE requests for all functions.

For more information about all the REST operations, see the *Online API documentation*.

Understand API version numbers in Active IQ Unified Manager

The REST API URIs in Active IQ Unified Manager specifies a version number. For example, `/v2/datacenter/svm/svms`. The version number `v2` in `/v2/datacenter/svm/svms` indicates the API version used in a specific release. The version number minimizes the impact of API changes on the client software by sending back a response that the client can process.

The numerical part of this version number is incremental with respect to releases. URIs with a version number provide a consistent interface that maintains backward compatibility in future releases. You also find the same

APIs without a version, for example `/datacenter/svm/svms`, that indicate the base APIs without a version. The base APIs are always the latest version of the APIs.

> On the top right corner of your Swagger interface, you can select the version of the API to use. The highest version is selected by default. It is recommended that you use the highest version of a particular API (with respect to the incremental integer) available in your Unified Manager instance.

For all requests, you must explicitly request the API version that you want to use. When the version number is specified, the service does not return response elements that your application is not designed to handle. In REST requests, you should include the version parameter. The earlier versions of the APIs are eventually deprecated after a few releases. In this release, the `v1` version of the APIs is deprecated.

Understand the ONTAP storage resource model in Active IQ Unified Manager

The storage resources in ONTAP can be broadly classified into *physical storage resources* and *logical storage resources.* To effectively manage your ONTAP systems using the APIs provided in Active IQ Unified Manager, you must understand the storage resource model and the relationship between various storage resources.

- **Physical storage resources**

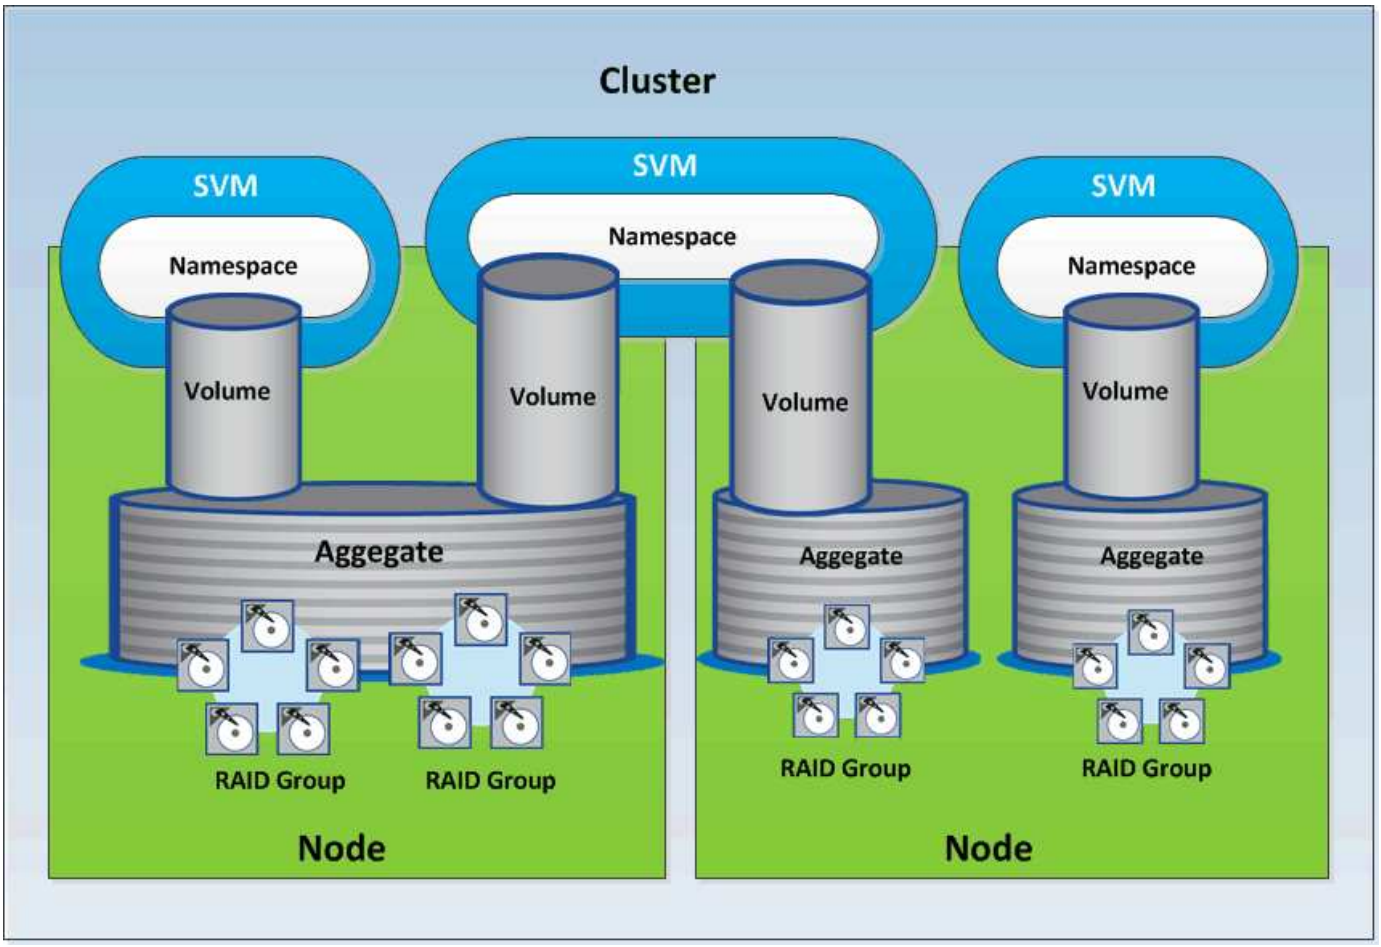
  Refers to the physical storage objects provided by ONTAP. Physical storage resources include disks, clusters, storage controllers, nodes, and aggregates.

- **Logical storage resources**

  Refers to the storage resources provided by ONTAP that are not tied to a physical resource. These resources are associated with a storage virtual machine (SVM, formerly known as Vserver), and they exist independently of any specific physical storage resource such as a disk, array LUN, or aggregate.

  Logical storage resources include volumes of all types and qtrees, as well as the capabilities and configurations you can use with these resources, such as Snapshot copies, deduplication, compression, and quotas.

The following illustration shows the storage resources in a 2-node cluster:



REST API access and authentication in Active IQ Unified Manager

Access and authenticate REST APIs in Active IQ Unified Manager

The Active IQ Unified Manager REST API is accessible by using any REST client or programming platform that can issue HTTP requests with a basic HTTP authentication mechanism.

A sample request and response:

- **Request**

```
GET
https://<IP
address/hostname>:<port_number>/api/v2/datacenter/cluster/clusters
```

- **Response**

```
{
  "records": [
```

```
    {
      "key": "4c6bf721-2e3f-11e9-a3e2-
00a0985badbb:type=cluster,uuid=4c6bf721-2e3f-11e9-a3e2-00a0985badbb",
      "name": "fas8040-206-21",
      "uuid": "4c6bf721-2e3f-11e9-a3e2-00a0985badbb",
      "contact": null,
      "location": null,
      "version": {
        "full": "NetApp Release Dayblazer__9.5.0: Thu Jan 17 10:28:33
UTC 2019",
        "generation": 9,
        "major": 5,
        "minor": 0
      },
      "isSanOptimized": false,
      "management_ip": "10.226.207.25",
      "nodes": [
        {
          "key": "4c6bf721-2e3f-11e9-a3e2-
00a0985badbb:type=cluster_node,uuid=12cf06cc-2e3a-11e9-b9b4-
00a0985badbb",
          "uuid": "12cf06cc-2e3a-11e9-b9b4-00a0985badbb",
          "name": "fas8040-206-21-01",
          "_links": {
            "self": {
              "href": "/api/datacenter/cluster/nodes/4c6bf721-2e3f-11e9-
a3e2-00a0985badbb:type=cluster_node,uuid=12cf06cc-2e3a-11e9-b9b4-
00a0985badbb"
            }
          },
          "location": null,
          "version": {
            "full": "NetApp Release Dayblazer__9.5.0: Thu Jan 17
10:28:33 UTC 2019",
            "generation": 9,
            "major": 5,
            "minor": 0
          },
          "model": "FAS8040",
          "uptime": 13924095,
          "serial_number": "701424000157"
        },
        {
          "key": "4c6bf721-2e3f-11e9-a3e2-
00a0985badbb:type=cluster_node,uuid=1ed606ed-2e3a-11e9-a270-
00a0985bb9b7",
```

```
          "uuid": "1ed606ed-2e3a-11e9-a270-00a0985bb9b7",
          "name": "fas8040-206-21-02",
          "_links": {
            "self": {
              "href": "/api/datacenter/cluster/nodes/4c6bf721-2e3f-11e9-
a3e2-00a0985badbb:type=cluster_node,uuid=1ed606ed-2e3a-11e9-a270-
00a0985bb9b7"
            }
          },
          "location": null,
          "version": {
            "full": "NetApp Release Dayblazer__9.5.0: Thu Jan 17
10:28:33 UTC 2019",
            "generation": 9,
            "major": 5,
            "minor": 0
          },
          "model": "FAS8040",
          "uptime": 14012386,
          "serial_number": "701424000564"
        }
      ],
      "_links": {
        "self": {
          "href": "/api/datacenter/cluster/clusters/4c6bf721-2e3f-11e9-
a3e2-00a0985badbb:type=cluster,uuid=4c6bf721-2e3f-11e9-a3e2-
00a0985badbb"
        }
      }
    },
```

- *IP address/hostname* is the IP address or the fully qualified domain name (FQDN) of the API server.
- Port 443

  443 is the default HTTPS port. You can customize the HTTPS port, if required.

To issue HTTP requests from a web browser, you have to use REST API browser plugins. You can also access the REST API by using scripting platforms such as cURL and Perl.

Authentication

Unified Manager supports the basic HTTP authentication scheme for APIs. For secure information flow (request and response), the REST APIs are accessible only over HTTPS. The API server provides a self-signed SSL certificate to all clients for server verification. This certificate can be replaced by a custom certificate (or CA certificate).

You must configure user access to the API server for invoking the REST APIs. The users can be local users (user profiles stored in the local database) or LDAP users (if you have configured the API server to

authenticate over LDAP). You can manage user access by logging in to the Unified Manager Administration Console user interface.

Learn about Active IQ Unified Manager HTTP status codes

While running the APIs or troubleshooting issues, you should be aware of the various HTTP status codes and error codes that are used by Active IQ Unified Manager APIs.

The following table lists the error codes related to authentication:

HTTP status code	Status code title	Description
200	OK	Returned on successful execution of synchronous API calls.
201	Created	Creation of new resources by synchronous calls, such as configuration of Active Directory.
202	Accepted	Returned on successful execution of asynchronous calls for provisioning functions, such as creating LUNs and files shares.
400	Invalid request	Indicates input validation failure. User has to correct the inputs, for example, valid keys in a request body.
401	Unauthorized request	You are not authorized to view the resource/Unauthorized.
403	Forbidden request	Accessing the resource you were trying to reach is forbidden.
404	Resource not found	The resource you were trying to reach is not found.
405	Method Not Allowed	Method not allowed.
429	Too Many Requests	Returned when the user sends too many requests within a specific time.
500	Internal server error	Internal server error. Failed to get the response from server. This internal server error may or may not be permanent. For example, if you run a `GET` or `GET ALL` operation and receive this error, it is recommended that you repeat this operation for a minimum of five retries. If it is a permanent error, then the status code returned continues to be 500. If the operation succeeds, the status code returned is 200.

Follow recommended practices for using Active IQ Unified Manager APIs

When using the APIs in Active IQ Unified Manager, you should follow certain recommended practices.

- All response content type must be in the following format for a valid execution:

```
application/json
```

- The API version number is not related to the product version number. You should use the latest version of the API available for your Unified Manager instance. For more information about Unified Manager API versions, see the “REST API versioning in Active IQ Unified Manager” section.
- While updating array values using a Unified Manager API, you must update the entire string of values. You cannot append values to an array. You can only replace an existing array.
- You can use filter operators, such as pipe (|) and wild card (*) for all query parameters, except for double values, for example, IOPS and performance in the metrics APIs.
- Avoid querying objects by using a combination of the filter operators wild card (*) and pipe (|). It might retrieve an incorrect number of objects.
- When using values for filter, ensure that the value does not contain any `?` character. This is to mitigate risks of SQL injection.
- Note that the `GET` (all) request for any API returns a maximum of 1000 records. Even if you run the query by setting the `max_records` parameter to a value higher than 1000, only 1000 records are returned.
- For performing administrative functions, it is recommended that you use the Unified Manager UI.

Troubleshoot with system logs in Active IQ Unified Manager

System logs enable you to analyze the causes of failure and troubleshooting issues that may arise while running the APIs.

Retrieve the logs from the following location for troubleshooting issues related to the API calls.

Log location	Use
`/var/log/ocie/access_log.log`	Contains all API call details, such as the user name of the user invoking the API, start time, execution time, status, and URL. You can use this log file to check the frequently-used APIs, or troubleshoot any GUI workflow. You can also use it to scale analysis, based on the execution time.
`/var/log/ocum/ocumserver.log`	Contains all API execution logs. You can use this log file to troubleshoot and debug the API calls.

Log location	Use
`/var/log/ocie/server.log`	Contains all Wildfly server deployments and start/stop service related logs. You can use this log file to find the root cause of any issues occurring during the start, stop, or deployment of the Wildfly server.
`/var/log/ocie/au.log`	Contains acquisition unit related logs. You can use this log file when you have created, modified, or deleted any objects in ONTAP but they do not get reflected for the Active IQ Unified Manager REST APIs.

Understand how Active IQ Unified Manager uses Job objects for asynchronous processing

Active IQ Unified Manager provides the `jobs` API that retrieves information about the Jobs performed while running other APIs. you must know how asynchronous processing works using the Job object.

Some of the API calls, particularly those that are used for adding or modifying resources, can take longer to complete than other calls. Unified Manager processes these long-running requests asynchronously.

Asynchronous requests described using Job object

After making an API call that runs asynchronously, the HTTP response code 202 indicates the request has been successfully validated and accepted, but not yet completed. The request is processed as a background task which continues to run after the initial HTTP response to the client. The response includes the Job object anchoring the request, including its unique identifier.

Query the Job object associated with an API request

The Job object returned in the HTTP response contains several properties. You can query the state property to determine if the request completed successfully. A Job object can be in one of the following states:

- `NORMAL`
- `WARNING`
- `PARTIAL_FAILURES`
- `ERROR`

There are two techniques you can use when polling a Job object to detect a terminal state for the task, either success or failure:

- Standard polling request: The current Job state is returned immediately.
- Long polling request: When the job state moves to `NORMAL, ERROR,` or `PARTIAL_FAILURES.`

Steps in an asynchronous request

You can use the following high-level procedure to complete an asynchronous API call:

1. Issue the asynchronous API call.
2. Receive an HTTP response 202 indicating successful acceptance of the request.
3. Extract the identifier for the Job object from the response body.
4. Within a loop, wait for the Job object to reach the terminal state `NORMAL, ERROR,` or `PARTIAL_FAILURES.`
5. Verify the terminal state of the Job and retrieve the Job result.

Learn about the Hello API server sample program in Active IQ Unified Manager

The *Hello API server* is a sample program that demonstrates how to invoke a REST API in Active IQ Unified Manager using a simple REST client. The sample program provides you basic details about the API server in the JSON format (the server supports only `application/json` format).

The URI used is: `https://<hostname>/api/datacenter/svm/svms`. This sample code takes the following input parameters:

- The API server IP address or FQDN
- Optional: Port number (default: 443)
- User name
- Password
- Response format (`application/json`)

To invoke REST APIs, you can also use other scripts such as Jersey and RESTEasy to write a Java REST client for Active IQ Unified Manager. You should be aware of the following considerations about the sample code:

- Uses an HTTPS connection to Active IQ Unified Manager to invoke the specified REST URI
- Ignores the certificate provided by Active IQ Unified Manager
- Skips the host name verification during the handshake
- Uses `javax.net.ssl.HttpsURLConnection` for a URI connection
- Uses a third-party library (`org.apache.commons.codec.binary.Base64`) for constructing the Base64 encoded string used in the HTTP basic authentication

To compile and execute the sample code, you must use Java compiler 1.8 or later.

```
import java.io.BufferedReader;
import java.io.InputStreamReader;
import java.net.URL;
import java.security.SecureRandom;
import java.security.cert.X509Certificate;
import javax.net.ssl.HostnameVerifier;
```

```
import javax.net.ssl.HttpsURLConnection;
import javax.net.ssl.SSLContext;
import javax.net.ssl.SSLSession;
import javax.net.ssl.TrustManager;
import javax.net.ssl.X509TrustManager;
import org.apache.commons.codec.binary.Base64;

public class HelloApiServer {

    private static String server;
    private static String user;
    private static String password;
    private static String response_format = "json";
    private static String server_url;
    private static String port = null;

    /*
     * * The main method which takes user inputs and performs the *
necessary steps
     * to invoke the REST URI and show the response
     */ public static void main(String[] args) {
        if (args.length < 2 || args.length > 3) {
            printUsage();
            System.exit(1);
        }
        setUserArguments(args);
        String serverBaseUrl = "https://" + server;
        if (null != port) {
            serverBaseUrl = serverBaseUrl + ":" + port;
        }
        server_url = serverBaseUrl + "/api/datacenter/svm/svms";
        try {
            HttpsURLConnection connection =
getAllTrustingHttpsUrlConnection();
            if (connection == null) {
                System.err.println("FATAL: Failed to create HTTPS
connection to URL: " + server_url);
                System.exit(1);
            }
            System.out.println("Invoking API: " + server_url);
            connection.setRequestMethod("GET");
            connection.setRequestProperty("Accept", "application/" +
response_format);
            String authString = getAuthorizationString();
            connection.setRequestProperty("Authorization", "Basic " +
authString);
```

```
            if (connection.getResponseCode() != 200) {
                System.err.println("API Invocation Failed : HTTP error
code : " + connection.getResponseCode() + " : "
                        + connection.getResponseMessage());
                System.exit(1);
            }
            BufferedReader br = new BufferedReader(new
InputStreamReader((connection.getInputStream())));
            String response;
            System.out.println("Response:");
            while ((response = br.readLine()) != null) {
                System.out.println(response);
            }
            connection.disconnect();
        } catch (Exception e) {
            e.printStackTrace();
        }
    }

    /* Print the usage of this sample code */ private static void
printUsage() {
        System.out.println("\nUsage:\n\tHelloApiServer <hostname> <user>
<password>\n");
        System.out.println("\nExamples:\n\tHelloApiServer localhost admin
mypassword");
        System.out.println("\tHelloApiServer 10.22.12.34:8320 admin
password");
        System.out.println("\tHelloApiServer 10.22.12.34 admin password
");
        System.out.println("\tHelloApiServer 10.22.12.34:8212 admin
password \n");
        System.out.println("\nNote:\n\t(1) When port number is not
provided, 443 is chosen by default.");
    }

    /* * Set the server, port, username and password * based on user
inputs. */ private static void setUserArguments(
            String[] args) {
        server = args[0];
        user = args[1];
        password = args[2];
        if (server.contains(":")) {
            String[] parts = server.split(":");
            server = parts[0];
            port = parts[1];
        }
```

```
    }

    /*
     * * Create a trust manager which accepts all certificates and * use
this trust
     * manager to initialize the SSL Context. * Create a
HttpsURLConnection for this
     * SSL Context and skip * server hostname verification during SSL
handshake. * *
     * Note: Trusting all certificates or skipping hostname verification *
is not
     * required for API Services to work. These are done here to * keep
this sample
     * REST Client code as simple as possible.
     */ private static HttpsURLConnection
getAllTrustingHttpsUrlConnection()    {        HttpsURLConnection conn =
null;        try {            /* Creating a trust manager that does not
validate certificate chains */            TrustManager[]
trustAllCertificatesManager = new                    TrustManager[]{new
X509TrustManager(){
     public X509Certificate[] getAcceptedIssuers(){return  null;}
     public void checkClientTrusted(X509Certificate[]
certs, String authType){}
     public void checkServerTrusted(X509Certificate[]
certs, String authType){}            }};            /* Initialize the
SSLContext with the all-trusting trust manager */
     SSLContext sslContext = SSLContext.getInstance("TLS");
sslContext.init(null, trustAllCertificatesManager, new
SecureRandom());
HttpsURLConnection.setDefaultSSLSocketFactory(sslContext.getSocketFactory(
));            URL url = new URL(server_url);            conn =
(HttpsURLConnection) url.openConnection();            /* Do not perform an
actual hostname verification during SSL Handshake.            Let all
hostname pass through as verified.*/
conn.setHostnameVerifier(new HostnameVerifier() {                public
boolean verify(String host, SSLSession                        session) {
return true;                }            });        } catch (Exception e)
{            e.printStackTrace();        }        return conn;    }

    /*
     * * This forms the Base64 encoded string using the username and
password *
     * provided by the user. This is required for HTTP Basic
Authentication.
     */ private static String getAuthorizationString() {
        String userPassword = user + ":" + password;
```

```
        byte[] authEncodedBytes =
Base64.encodeBase64(userPassword.getBytes());
        String authString = new String(authEncodedBytes);
        return authString;
    }

}
```

Unified Manager REST APIs

Understand categories for Active IQ Unified Manager REST APIs

The REST APIs for Active IQ Unified Manager are listed in this section, based on their categories.

You can view the online documentation page from your Unified Manager instance that includes the details of every REST API call. This document does not repeat the details of the online documentation. Each API call listed or described in this document includes only the information you need to locate the call on the documentation page. After locating a specific API call, you can review the complete details of that call, including the input parameters, output formats, HTTP status codes, and request processing type.

The following information is included for each API call within a workflow to help locate the call on the documentation page:

- Category

  The API calls are organized on the documentation page into functionally-related areas or categories. To locate a specific API call, scroll down to the bottom of the page, and then click the applicable API category.

- HTTP verb (call)

  The HTTP verb identifies the action performed on a resource. Each API call is executed through a single HTTP verb.

- Path

  The path determines the specific resource which the action uses to as a part of performing a call. The path string is appended to the core URL to form the complete URL identifying the resource.

Manage storage objects in a data center using APIs in Active IQ Unified Manager

The REST APIs under the `datacenter` category enable you to manage the storage objects in your data center, such as clusters, nodes, aggregates, storage VMs, volumes, LUNs, file shares, and namespaces. These APIs are available for querying the configuration of the objects, while some of them enable you to perform the operations of adding, deleting, or modifying these objects.

Most of these APIs are GET calls that provide cross-cluster aggregation with filtering, sorting and pagination support. On running these APIs, they return data from the database. Therefore, the newly created-objects need to be discovered by the next acquisition cycle for appearing in the response.

If you want to query the details of a specific object, you need to enter the unique ID of that object to view its details. For example,
For the metrics and analytics information of the storage objects, see Viewing performance metrics.

```
curl -X GET "https://<hostname>/api/datacenter/cluster/clusters/4c6bf721-2e3f-11e9-a3e2-00a0985badbb" -H "accept: application/json" -H "Authorization: Basic <Base64EncodedCredentials>"
```

The CURL commands, examples, requests, and responses to the APIs, are available on your Swagger API interface. You can filter and order the results by specific parameters as indicated on Swagger. These APIs enable you to filter the results for specific storage objects, such as cluster, volume, or storage VM.

APIs for storage objects in your data center

HTTP verb	Path	Description
GET	/datacenter/cluster/clusters /datacenter/cluster/clusters/{key}	You can use this method to view the details of the ONTAP clusters across the data center. The API returns information, such as the IPv4 or IPv6 address of the cluster, information about the node, such as node health, performance capacity, and High Availability (HA) pair, and indicates whether the cluster is All SAN Array.
GET	/datacenter/cluster/licensing/licenses /datacenter/cluster/licensing/licenses/{key}	Returns the details of the licenses installed on the clusters in your data center. You can filter your results based on the required criteria. Information, such as the license key, cluster key, expiry date, and license scope is returned. You can enter a license key to retrieve the details of a specific license.
GET	/datacenter/cluster/nodes /datacenter/cluster/nodes/{key}	You can use this method to view the details of the nodes in the data center. You can view information about the cluster, node health, performance capacity, and High Availability (HA) pair for the node.
GET	/datacenter/protocols/cifs/shares /datacenter/protocols/cifs/shares/{key}	You can use this method to view the details of the CIFS shares in the data center. Apart from cluster, SVM, and volume details, information about Access Control List (ACL) is also returned.

HTTP verb	Path	Description
GET	/datacenter/protocols/nfs/export-policies /datacenter/protocols/nfs/export-policies/{key}	You can use this method to view the details of the export policies for the supported NFS services. You can query the export policies for a cluster or storage VM and reuse the export policy key for provisioning NFS file shares. For more information about the assigning and reusing export policies on workloads, see “Provisioning CIFS and NFS file shares”.
GET	/datacenter/storage/aggregates /datacenter/storage/aggregates/{key}	You can use this method to view the collection of aggregates in the data center or a specific aggregate for provisioning workloads on them or monitoring. Information, such as cluster and node details, performance capacity used, available and used space, and storage efficiency is returned.
GET	/datacenter/storage/luns /datacenter/storage/luns/{key}	You can use this method to view the collection of LUNs in the entire data center. You can view information about the LUN, such as cluster and SVM details, QoS policies, and igroups.
GET	/datacenter/storage/qos/policies /datacenter/storage/qos/policies/{key}	You can use this method to view the details of all the QoS policies applicable for the storage objects in the data center. Information, such as the cluster and SVM details, the fixed or adaptive policy details, and number of objects applicable for that policy is returned.
GET	/datacenter/storage/qtrees /datacenter/storage/qtrees/{key}	You can use this method to view the qtree details across the data center for all FlexVol volumes or FlexGroup volumes. Information, such as the cluster and SVM details, FlexVol volume, and export policy are returned.

HTTP verb	Path	Description
GET	/datacenter/storage/volumes /datacenter/storage/volumes/{key}	You can use this method to view the collection of volumes in the data center. Information about the volumes, such as SVM and cluster details, QoS and export policies, whether the volume is of type read-write, data-protection, or load-sharing, is returned. For FlexVol and FlexClone volumes, you can view the information about the respective aggregates. For a FlexGroup volume, the query returns the list of constituent aggregates.
GET POST DELETE PATCH	/datacenter/protocols/san/igroups /datacenter/protocols/san/igroups/{key}	You can assign initiator groups (igroups) authorized to access particular LUN targets. If there is an existing igroup, you can assign it. You can also create igroups and assign them to the LUNs. You can use these methods to query, create, delete, and modify igroups respectively. Points to note: • POST: While creating an igroup, you can designate the storage VM on which you want to assign access. • DELETE: You need to provide the igroup key as an input parameter to delete a particular igroup. If you have already assigned an igroup to a LUN, you cannot delete that igroup. • PATCH: You need to provide the igroup key as an input parameter to modify a particular igroup. You must also enter the property that you want to update, along with its value.

HTTP verb	Path	Description

`GET`
 	
`POST`
 	
`DELETE`
 	
`PATCH`
	`/datacenter/svm/svms`
	
	`/datacenter/svm/svms/{key}`
		You can use these methods to view, create, delete, and modify Storage Virtual Machines (storage VMs).
		
		• `POST`: You must enter the storage VM object that you want to create as an input parameter. You can create a custom storage VM, and then assign required properties to it.
		• `DELETE`: You need to provide the storage VM key to delete a particular storage VM.
		• `PATCH`: You need to provide the storage VM key to modify a particular storage VM. You also need to enter the properties that you want to update, along with their values.

Points to note:

If you have enabled SLO-based workload provisioning on your environment, while creating the storage VM, ensure that it supports all of the protocols required for provisioning LUNs and file shares on them, for example, CIFS or SMB, NFS, FCP, and iSCSI. The provisioning workflows might fail if the storage VM does not support the required services. It is recommended that the services for the respective types of workloads are also enabled on the storage VM.

If you have enabled SLO-based workload provisioning on your environment, you cannot delete that storage VM on which storage workloads have been provisioned. When you delete a storage VM on which a CIFS or SMB server has been configured, this API also deletes the CIFS or SMB server, along with the local Active Directory configuration. However, the CIFS or SMB server name continues to be in the Active Directory configuration that you must delete manually from the Active Directory server.

APIs for network elements in your data center

The following APIs in the datacenter category retrieve information about the ports and network interfaces in your environment, specifically the FC ports, FC interfaces, ethernet ports, and IP interfaces.

HTTP verb	Path	Description
GET	/datacenter/network/ethernet/ports /datacenter/network/ethernet/ports/{key}	Retrieves information about all the ethernet ports in your data center environment. With a port key as an input parameter, you can view the information of that specific port. Information, such as cluster details, broadcast domain, port details, such as state, speed, and type, and whether the port is enabled, is retrieved.
GET	/datacenter/network/fc/interfaces /datacenter/network/fc/interfaces/{key}	You can use this method to view the details of the FC interfaces in your data center environment. With an interface key as an input parameter, you can view the information of that specific interface. Information, such as cluster details, home node details, and home port details, is retrieved.
GET	/datacenter/network/fc/ports /datacenter/network/fc/ports/{key}	Retrieves information about all the FC ports used in the nodes in your data center environment. With a port key as an input parameter, you can view the information of that specific port. Information, such as cluster details, port description, supported protocol, and the state of the port is retrieved.
GET	/datacenter/network/ip/interfaces /datacenter/network/ip/interfaces/{key}	You can use this method to view the details of the IP interfaces in your data center environment. With an interface key as an input parameter, you can view the information of that specific interface. Information, such as cluster details, IPspace details, home node details, whether failover is enabled, is retrieved.

Access ONTAP APIs through proxy access in Active IQ Unified Manager

The gateway APIs provide you with the advantage of using the Active IQ Unified Manager credentials to run ONTAP REST APIs and managing storage objects. These APIs are available when the API Gateway feature is enabled from the Unified Manager web UI.

Unified Manager REST APIs support only a select set of actions to be performed on the Unified Manager datasources, that is ONTAP clusters. You can avail the other features through ONTAP APIs. The gateway APIs allow Unified Manager to be a pass-through interface for tunneling all API requests to be performed on ONTAP clusters that it manages, without logging in to each data center cluster individually. It performs as a single point of management for running the APIs across the ONTAP clusters managed by your Unified Manager instance.

The API Gateway feature allows Unified Manager to be a single control plane from which you can manage multiple ONTAP clusters, without logging in to them individually. The gateway APIs enable you to remain logged in to Unified Manager and manage the ONTAP clusters by running ONTAP REST API operations.

All users can run a query by using the GET operation. Application Administrators can run all ONTAP REST operations.

The gateway acts as a proxy to tunnel the API requests by maintaining the header and body requests in the same format as in the ONTAP APIs. You can use your Unified Manager credentials and execute the specific operations to access and manage the ONTAP clusters without passing individual cluster credentials. It continues to manage the cluster authentication and cluster management, but redirects the API requests to run directly on the specific cluster. The response returned by the APIs is the same as the response returned by the respective ONTAP REST APIs executed directly from ONTAP.

HTTP verb	Path (URL)	Description
`GET`	`/gateways`	This GET method retrieves the list of all the clusters managed by Unified Manager that support ONTAP REST calls. You can verify the cluster details and choose to run other methods based on the cluster UUID or universal unique identifier (UUID). The gateway APIs retrieve only those clusters supported by ONTAP 9.5 or later, and added to Unified Manager over HTTPS.

HTTP verb	Path (URL)	Description

`GET`
 	
`POST`
 	
`DELETE`
 	
`PATCH`
 	
`OPTIONS` (not available on Swagger)
 	
`HEAD` (not available on Swagger)
	`/gateways/{uuid}/{path}`
	
	The value for {uuid} must be replaced with the cluster UUID on which the REST operation is to be performed. Also, ensure that the UUID is of the cluster supported by ONTAP 9.5 or later, and added to Unified Manager over HTTPS. {path} must be replaced by the ONTAP REST URL. You must remove `/api/` from the URL.
		This is a single point proxy API, supporting POST, DELETE, PATCH operations and GET for all the ONTAP REST APIs. No restrictions apply on any of the API as long as it is supported by ONTAP. The tunnelling or proxy functionality cannot be disabled.
		
		The `OPTIONS` method returns all the operations supported by an ONTAP REST API. For example, if an ONTAP API supports only the `GET` operation, running the `OPTIONS` method by using this gateway API returns `GET` as the response. This method is not supported on Swagger, but can be performed on other API tools.
		
		The `OPTIONS` method determines whether a resource is available. This operation can be used to view the metadata about a resource in the HTTP response headers. This method is not supported on Swagger, but can be performed on other API tools.

Understand the API Gateway tunneling

The gateway APIs enable you to manage ONTAP objects through Unified Manager. Unified Manager manages the clusters and authentication details and redirects the requests to the ONTAP REST endpoint. The gateway API transforms the URL and Hypermedia as the Engine of Application State (HATEOAS) links in the header and response body with the API gateway base URL. The gateway API acts as the proxy base URL to which you append the ONTAP REST URL and execute the required ONTAP REST endpoint.

For an ONTAP API to run successfully through API gateway, the API must be supported by that version of the ONTAP cluster on which it is being run. Running an API that is not supported on the ONTAP cluster does not return any results.

In this example, the gateway API (proxy base URL) is: `/gateways/{uuid}/`

The ONTAP API taken is: `/storage/volumes`. You need to add the ONTAP API REST URL as the value for the path parameter.

While adding the path, ensure that you have removed the “/” symbol at the beginning of the URL. For the API `/storage/volumes`, add `storage/volumes`.

The appended URL is: `/gateways/{uuid}/storage/volumes`

On running the `GET` operation, the generated URL is the following:
`GEThttps://<hostname\>/api/gateways/<cluster_UUID\>/storage/volumes`

The `/api` tag of the ONTAP REST URL is removed in the appended URL and that for the gateway API is retained.

Sample cURL command

```
curl -X GET "https://<hostname>/api/gateways/1cd8a442-86d1-11e0-ae1c-
9876567890123/storage/volumes" -H "accept: application/hal+json" -H
"Authorization: Basic <Base64EncodedCredentials>"
```

The API returns the list of storage volumes in that cluster. The response format is the same as you receive when you run the same API from ONTAP. The status codes returned are the ONTAP REST status codes.

Set API scope

All APIs have a context set within the scope of the cluster. APIs that operate on the basis of storage VMs also have the cluster as the scope, that is, the API operations are performed on a particular storage VM within a managed cluster. When you run the `/gateways/{uuid}/{path}` API, ensure that you enter the cluster UUID (Unified Manager datasource UUID) for the cluster on which you run the operation. For setting the context to a particular storage VM within that cluster, enter the storage VM key as the X-Dot-SVM-UUID parameter or the storage VM name as the X-Dot-SVM-Name parameter. The parameter is added as the filter in the string header and the operation is run within the scope of that storage VM inside that cluster.

Sample cURL command

```
curl -X GET "https://<hostname>/api/gateways/e4f33f90-f75f-11e8-9ed9-
00a098e3215f/storage/volume" -H "accept: application/hal+json" -H "X-Dot-
SVM-UUID: d9c33ec0-5b61-11e9-8760-00a098e3215f"
-H "Authorization: Basic <Base64EncodedCredentials>"
```

For more information about using ONTAP REST APIs, see
ONTAP REST API Automation

Perform administrative tasks using APIs in Active IQ Unified Manager

You can use the APIs under the `administration` category to modify backup settings, verify the backup file information and cluster certificates, and also manage ONTAP clusters as Active IQ Unified Manager datasources.



You must have the Application Administrator role for running these operations. You can also use the Unified Manager web UI for configuring these settings.

HTTP verb	Path	Description
GET PATCH	/admin/backup-settings /admin/backup-settings	You can use the GET method to view the settings of the backup schedule configured in Unified Manager by default. You can verify the following: • Whether the schedule is enabled or disabled • Frequency of the backup scheduled (daily or weekly) • Time of the backup • Maximum number of backup files that should be retained in the application The time of the backup is in server time zone. The database backup settings are available on Unified Manager by default, and you cannot create a backup schedule. However, you can use the PATCH method to modify the default settings.
GET	/admin/backup-file-info	A backup dump file is generated every time the backup schedule is modified for Unified Manager. You can use this method to verify whether the backup file is generated according to the modified backup settings, and whether the information on the file matches the modified settings.
GET	/admin/datasource-certificate	You can use this method to view the datasource (cluster) certificate from the trust store. Validating the certificate is required before adding an ONTAP cluster as a Unified Manager datasource.

HTTP verb	Path	Description
GET POST PATCH DELETE	/admin/datasources/clusters /admin/datasources/clusters/{key}	You can use the GET method to retrieve the details of the datasources (ONTAP clusters) managed by Unified Manager. You can also add a new cluster to Unified Manager as a datasource. For adding a cluster, you must know its host name, user name, and password. For modifying and deleting a cluster managed as a datasource by Unified Manager, use the ONTAP cluster key.

Manage users in Active IQ Unified Manager using APIs

You can use the APIs in the security category to control user access to selected cluster objects in Active IQ Unified Manager. You can add local users or database users. You can also add remote users or groups that belong to an authentication server. Based on the privileges of the roles that you assign to the users, they can manage the storage objects or view the data in Unified Manager.

You must have the Application Administrator role for running these operations. You can also use the Unified Manager web UI for configuring these settings.

The APIs under the security category use the users parameter, that is the user name, and not the key parameter as the unique identifier for the user entity.

HTTP verb	Path	Description
GET POST	/security/users /security/users	You can use these methods to get the details of users or add new user to Unified Manager. You can add specific roles to the users based on their user types. While adding users, you must provide passwords for the local user, maintenance user, and database user.
GET PATCH DELETE	/security/users/{name}	The GET method enables you to retrieve all the details of a user, such as the name, email address, role, authorization type. The PATCH method enables you to update the details. The DELETE method enables you to remove the user.

View performance metrics using APIs in Active IQ Unified Manager

Active IQ Unified Manager provides you with a set of APIs under the `/datacenter` category that enable you to view the performance data of your clusters and storage objects in a data center. These APIs retrieve the performance data of the different storage objects such as clusters, nodes, LUNs, volumes, aggregates, storage VMs, FC interfaces, FC ports, Ethernet ports, and IP interfaces.

The `/metrics` and `/analytics` APIs give different views of the performance metrics, using which, you can drill down to different levels of details for the following storage objects in your data center:

- clusters
- nodes
- storage VMs
- aggregates
- volumes
- LUNs
- FC interfaces
- FC ports
- Ethernet ports
- IP interfaces

The following table draws a comparison between the `/metrics` and `/analytics` APIs with respect to the particulars of the performance data retrieved.

Metrics	Analytics
Performance details for a single object. For example, the `/datacenter/cluster/clusters/{key}/metrics` API requires the cluster key to be entered as the path parameter for retrieving the metrics for that specific cluster.	Performance details for multiple objects of the same type in a data center. For example, the `/datacenter/cluster/clusters/analytics` API retrieves the collective metrics of all the clusters in a data center.
Performance metrics sample for a storage object based on the time interval parameter for retrieval.	The high-level aggregated value of performance for a certain type of storage object for a certain period (above 72 hours).
Basic details of the object are retrieved, such as details of a node or cluster.	No specific details are retrieved.
Accumulated counters, such as minimum, maximum, 95th percentile, and the average performance values over a period of time are retrieved for a single object, such as read, write, total, and other counters.	A single aggregated value is displayed for all the objects of the same type.

Metrics	Analytics
The time range and sample data is based on the following schedule: The time range for the data. Examples can be 1h, 12h, 1d, 2d, 3d, 15d, 1w, 1m, 2m, 3m, 6m. You get 1 hour samples if the range is more than 3 days (72 hr) else it is 5 minutes samples. The period for each time range is as follows: • 1h: Metrics over the most recent hour sampled over 5 minutes. • 12h: Metrics over the most recent 12 hours sampled over 5 minutes. • 1d: Metrics over the most recent day sampled over 5 minutes. • 2d: Metrics over the most recent 2 days sampled over 5 minutes. • 3d: Metrics over the most recent 3 days sampled over 5 minutes. • 15d: Metrics over the most recent 15 days sampled over 1 hour. • 1w: Metrics over the most recent week sampled over 1 hour. • 1m: Metrics over the most recent month sampled over 1 hour. • 2m: Metrics over the most recent 2 months sampled over 1 hour. • 3m: Metrics over the most recent 3 months sampled over 1 hour. • 6m: Metrics over the most recent 6 months sampled over 1 hour. Available values : 1h, 12h, 1d, 2d, 3d, 15d, 1w, 1m, 2m, 3m, 6m Default value : 1h	Above 72 hours. The duration over which this sample is calculated are represented in the ISO-8601 standard format.

Output sample for metrics APIs

For example, the `/datacenter/cluster/nodes/{key}/metrics` API retrieves the following details (among others) for a node:



The 95 percentile in the summary value indicates that 95% of the samples collected for the period have a counter value less than the value specified as the 95 percentile.

```
{
     "iops": {
       "local": {
```

```
        "other": 100.53,
        "read": 100.53,
        "total": 100.53,
        "write": 100.53
      },
      "other": 100.53,
      "read": 100.53,
      "total": 100.53,
      "write": 100.53
    },
    "latency": {
      "other": 100.53,
      "read": 100.53,
      "total": 100.53,
      "write": 100.53
    },
    "performance_capacity": {
      "available_iops_percent": 0,
      "free_percent": 0,
      "system_workload_percent": 0,
      "used_percent": 0,
      "user_workload_percent": 0
    },
    "throughput": {
      "other": 100.53,
      "read": 100.53,
      "total": 100.53,
      "write": 100.53
    },
    "timestamp": "2018-01-01T12:00:00-04:00",
    "utilization_percent": 0
  }
],
"start_time": "2018-01-01T12:00:00-04:00",
"summary": {
  "iops": {
    "local_iops": {
      "other": {
        "95th_percentile": 28,
        "avg": 28,
        "max": 28,
        "min": 5
      },
      "read": {
        "95th_percentile": 28,
        "avg": 28,
```

```
          "max": 28,
          "min": 5
        },
        "total": {
          "95th_percentile": 28,
          "avg": 28,
          "max": 28,
          "min": 5
        },
        "write": {
          "95th_percentile": 28,
          "avg": 28,
          "max": 28,
          "min": 5
        }
      },
```

Output sample for analytics APIs

For example, the `/datacenter/cluster/nodes/analytics` API retrieves the following values (among others) for all the nodes:

```
{     "iops": 1.7471,
      "latency": 60.0933,
      "throughput": 5548.4678,
      "utilization_percent": 4.8569,
      "period": 72,
      "performance_capacity": {
        "used_percent": 5.475,
        "available_iops_percent": 168350
      },
      "node": {
        "key": "37387241-8b57-11e9-8974-
00a098e0219a:type=cluster_node,uuid=95f94e8d-8b4e-11e9-8974-00a098e0219a",
        "uuid": "95f94e8d-8b4e-11e9-8974-00a098e0219a",
        "name": "ocum-infinity-01",
        "_links": {
          "self": {
            "href": "/api/datacenter/cluster/nodes/37387241-8b57-11e9-8974-
00a098e0219a:type=cluster_node,uuid=95f94e8d-8b4e-11e9-8974-00a098e0219a"
          }
        }
      },
      "cluster": {
        "key": "37387241-8b57-11e9-8974-
00a098e0219a:type=cluster,uuid=37387241-8b57-11e9-8974-00a098e0219a",
```

```
        "uuid": "37387241-8b57-11e9-8974-00a098e0219a",
        "name": "ocum-infinity",
        "_links": {
          "self": {
            "href": "/api/datacenter/cluster/clusters/37387241-8b57-11e9-
8974-00a098e0219a:type=cluster,uuid=37387241-8b57-11e9-8974-00a098e0219a"
          },
      "_links": {
        "self": {
          "href": "/api/datacenter/cluster/nodes/analytics"
        }
      }
    },
```

List of the available APIs

The following table describes the `/metrics` and `/analytics` APIs in details.

> The IOPS and performance metrics returned by these APIs are double values, for example `100.53`. Filtering these float values by the pipe (|) and wildcard (*) characters is not supported.

HTTP Verb	Path	Description
`GET`	`/datacenter/cluster/clusters/{key}/metrics`	Retrieves performance data (sample and summary) for a cluster specified by the input parameter of the cluster key. Information, such as the cluster key and UUID, time range, IOPS, throughput, and the number of samples is returned.
`GET`	`/datacenter/cluster/clusters/analytics`	Retrieves high-level performance metrics for all the clusters in a data center. You can filter your results based on the required criteria. Values, such as aggregated IOPS, throughput, and the period of collection (in hours) are returned.
`GET`	`/datacenter/cluster/nodes/{key}/metrics`	Retrieves performance data (sample and summary) for a node specified by the input parameter of the node key. Information, such as the node UUID, time range, summary of the IOPS, throughput, latency, and performance, the number of samples collected, and percentage utilized is returned.

HTTP Verb	Path	Description
GET	/datacenter/cluster/nodes/analytics	Retrieves high-level performance metrics for all the nodes in a data center. You can filter your results based on the required criteria. Information, such as node and cluster keys, and values, such as aggregated IOPS, throughput, and the period of collection (in hours) are returned.
GET	/datacenter/storage/aggregates/{key}/metrics	Retrieves performance data (sample and summary) for an aggregate specified by the input parameter of the aggregate key. Information, such as the time range, summary of the IOPS, latency, throughput, and performance capacity, the number of samples collected for each counter, and percentage utilized is returned.
GET	/datacenter/storage/aggregates/analytics	Retrieves high-level performance metrics for all the aggregates in a data center. You can filter your results based on the required criteria. Information, such as aggregate and cluster keys, and values, such as aggregated IOPS, throughput, and the period of collection (in hours) are returned.
GET	/datacenter/storage/luns/{key}/metrics /datacenter/storage/volumes/{key}/metrics	Retrieves performance data (sample and summary) for a LUN or a file share (volume) specified by the input parameter of the LUN or volume key. Information, such as the summary of the minimum, maximum, and average of the read, write, and total IOPS, latency, and throughput, and the number of samples collected for each counter is returned.
GET	/datacenter/storage/luns/analytics /datacenter/storage/volumes/analytics	Retrieves high-level performance metrics for all the LUNs or volumes in a data center. You can filter your results based on the required criteria. Information, such as storage VM and cluster keys, and values, such as aggregated IOPS, throughput, and the period of collection (in hours) are returned.

HTTP Verb	Path	Description
GET	/datacenter/svm/svms/{key}/metrics	Retrieves performance data (sample and summary) for a storage VM specified by the input parameter of the storage VM key. Summary of the IOPS based on each supported protocol, such as nvmf, fcp, iscsi, and nfs, throughput, latency, and the number of samples collected are returned.
GET	/datacenter/svm/svms/analytics	Retrieves high-level performance metrics for all the storage VMs in a data center. You can filter your results based on the required criteria. Information, such as storage VM UUID, aggregated IOPS, latency, throughput, and the period of collection (in hours) are returned.
GET	/datacenter/network/ethernet/ports/{key}/metrics	Retrieves the performance metrics for a specific ethernet port specified by the input parameter of the port key. When an interval (time range) is provided from the supported range, the API returns the accumulated counters, such as minimum, maximum, and the average performance values over the period of time.
GET	/datacenter/network/ethernet/ports/analytics	Retrieves the high-level performance metrics for all the ethernet ports in your data center environment. Information, such as the cluster and node key and UUID, throughput, period of collection, and utilization percent for the ports is returned. You can filter the result by the available parameters, such as the port key, utilization percent, cluster and node name and UUID, and so forth.

HTTP Verb	Path	Description
GET	/datacenter/network/fc/interfaces/{key}/metrics	Retrieves the performance metrics for a specific network FC interface specified by the input parameter of the interface key. When an interval (time range) is provided from the supported range, the API returns the accumulated counters, such as minimum, maximum, and the average performance values over the period of time.
GET	/datacenter/network/fc/interfaces/analytics	Retrieves the high-level performance metrics for all the ethernet ports in your data center environment. Information, such as the cluster and FC interface key and UUID, throughput, IOPS, latency, and storage VM is returned. You can filter the result by the available parameters, such as the cluster and FC interface name and UUID, storage VM, throughput, and so forth.
GET	/datacenter/network/fc/ports/{key}/metrics	Retrieves the performance metrics for a specific FC port specified by the input parameter of the port key. When an interval (time range) is provided from the supported range, the API returns the accumulated counters, such as minimum, maximum, and the average performance values over the period of time.
GET	/datacenter/network/fc/ports/analytics	Retrieves the high-level performance metrics for all the FC ports in your data center environment. Information, such as the cluster and node key and UUID, throughput, period of collection, and utilization percent for the ports is returned. You can filter the result by the available parameters, such as the port key, utilization percent, cluster and node name and UUID, and so forth.

HTTP Verb	Path	Description
GET	/datacenter/network/ip/interfaces/{key}/metrics	Retrieves the performance metrics for a network IP interface as specified by the input parameter of the interface key. When an interval (time range) is provided from the supported range, the API returns information, such as the number of samples, accumulated counters, throughput, and the number of packets received and transmitted.
GET	/datacenter/network/ip/interfaces/analytics	Retrieves the high-level performance metrics for all the network IP interfaces in your data center environment. Information, such as the cluster and IP interface key and UUID, throughput, IOPS, and latency is returned. You can filter the result by the available parameters, such as the cluster and IP interface name and UUID, IOPS, latency, throughput, and so forth.

View jobs and system details using APIs in Active IQ Unified Manager

You can use the `jobs` API under the `management-server` category to view the execution details of asynchronous operations. The `system` API under the `management-server` category enables you to view the instance details in your Active IQ Unified Manager environment.

View Jobs

In Active IQ Unified Manager, operations, such as adding and modifying resources are performed by synchronous and asynchronous API invocations. Invocations that are scheduled for asynchronous execution can be tracked by a Job object created for that invocation. Each Job object has a unique key for identification. Each Job object returns the Job object URI for you to access and track the progress of the job. You can use this API for retrieving the details of each execution.

By using this API, you can query all the Job objects for you data center, including historical data. Querying all the jobs, by default, returns the details of the last 20 jobs triggered through the web UI and API interface. Use the inbuilt filters to view specific Jobs. You can also use the Job key to query the details of a specific job and run the next set of operations on the resources.

Category	HTTP verb	Path	Description
management-server	GET	/management-server/jobs	Returns the job details of all the jobs. Without any sort order, the last submitted Job object is returned on top.

Category	HTTP verb	Path	Description
management-server	GET	`/management-server/jobs/{key}` Enter the job key of the Job object to view the specific details of that job.	Returns the details of the specific Job object.

View system details

By using the `/management-server/system` API, you can query the instance-specific details of your Unified Manager environment. The API returns information about the product and services, such as the version of Unified Manager installed on your system, UUID, vendor name, host OS, and the name, description, and status of the services running on the Unified Manager instance.

Category	HTTP verb	Path	Description
management-server	GET	`/management-server/system`	No input parameter is required for running this API. The system details of the current Unified Manager instance are returned by default.

Manage events and alerts using APIs in Active IQ Unified Manager

The `events`, `alerts`, and `scripts` APIs under the `management-server` category enable you to manage the events, alerts, and the scripts associated with the alerts in your Active IQ Unified Manager environment.

View and modify events

Unified Manager receives the events that are generated on ONTAP for the clusters monitored and managed by Unified Manager. By using these APIs, you can view the events generated for your clusters, and resolve and update them.

By running the `GET` method for the `/management-server/events` API, you can query the events in your data center, including historical data. Use the inbuilt filters, such as name, impact level, impact area, severity, state, resource name, and resource type, to view specific events. The resource type and area parameters return information about the storage object on which the event has occurred, and the impact area returns the information about the issue for which the event is raised, such as availability, capacity, configuration, security, protection, and performance.

By running the PATCH operation for this API, you can enable the resolution workflow for the event. You can assign an event to yourself or another user and acknowledge the receipt of the event. On performing the steps on the resources to resolve the issue that triggered the event, you can use this API to mark the event as resolved.

For more information about events, see Learn about Active IQ platform events.

Category	HTTP verb	Path	Description
management-server	GET	`/management-server/events` `/management-server/events/{key}`	When you run the Get ALL method, the response body consists of the event details of the all the events in your data center. When you retrieve the event details by a specific key, you can view the details of a specific event and run the next set of operations on the resources. The response body consists of the details of that event.
management-server	PATCH	`management-server/events/{key}`	Run this API to assign an event or change the state to acknowledged or resolved. You can also use this method to assign the event to yourself or another user. It is a synchronous operation.

Manage alerts

Events are automatically and continuously generated. Unified Manager generates an alert only when an event meets certain filter criteria. You can select the events for which alerts should be generated. By using the `/management-server/alerts` API, you can configure alerts to send notifications automatically when specific events or events of certain severity types occur.

For more information about alerts, see Learn about alerts.

Category	HTTP verb	Path	Description
management-server	GET	`/management-server/alerts` `/management-server/alerts/{key}`	Query all the existing alerts in your environment or a specific alert, by using the alert key. You can view the information about the alerts generated in your environment such as the alert description, action, email ID to which the notification is sent, event and severity.

Category	HTTP verb	Path	Description
management-server	POST	`/management-server/alerts`	This method enables you to add alerts for specific events. You must add the alert name, the physical or logical resource, or event on which the alert is applicable, whether the alert is enabled, and whether you are issuing SNMP traps. You can add additional details for which you want to generate the alert, such as the action, notification email ID, the script details, in case you are adding an alert script, and so forth.
management-server	PATCH and DELETE	`management-server/events/{key}`	You can use these methods to modify and delete specific alerts. You can modify different attributes, such as description, name, and enabling and disabling the alert. You can delete an alert when the alert is no longer required.

While selecting a resource for adding an alert, note that selecting a cluster as the resource does not automatically select the storage objects within that cluster. For example, if you create an alert for all the critical events for all clusters, you will receive alerts only for cluster critical events. You will not receive alerts for critical events on nodes, aggregates, and so forth.

Manage scripts

By using the `/management-server/scripts` API, you can also associate an alert with a script that is executed when an alert is triggered.
You can use scripts to automatically modify or update multiple storage objects in Unified Manager. The script is associated with an alert. When an event triggers an alert, the script is executed. You can upload custom scripts and test their execution when an alert is generated. You can associate an alert with your script so that the script is executed when an alert is raised for an event in Unified Manager.

For more information about scripts, see Manage scripts.

Category	HTTP verb	Path	Description
management-server	GET	`/management-server/scripts`	Use this API to query all the existing scripts in your environment. Use the standard filter and order by operations to view only specific scripts.
management-server	POST	`/management-server/scripts`	Use this API to add a description for the script and upload the script file associated with an alert.

Manage workloads

View storage workloads using APIs in Active IQ Unified Manager

The APIs listed here enable you to view a consolidated list of storage workloads for all of the ONTAP clusters in your data center. The APIs also provide a summary view of the number of the storage workloads provisioned in your Active IQ Unified Manager environment, and their capacity and performance (IOPS) statistics.

View storage workloads

You can use the following method to view all the storage workloads in all the clusters in your data center. For information about filtering the response based on specific columns, see the API reference documentation available in your Unified Manager instance.

Category	HTTP verb	Path
storage-provider	GET	`/storage-provider/workloads`

View storage workloads summary

You can use the following method to assess the used capacity, available capacity, used IOPS, available IOPS, and number of storage workloads managed by each Performance Service Level. The storage workloads displayed can be for any LUN, NFS file share, or CIFS share. The API provides a storage workloads overview, an overview of the storage workloads provisioned by the Unified Manager, a data center overview, an overview of the total, used, and available space and IOPS in the data center, in terms of the assigned Performance Service Levels. The information received in response to this API is used to populate the dashboard in the Unified Manager UI.

Category	HTTP verb	Path
storage-provider	GET	`/storage-provider/workloads-summary`

Manage access endpoints using APIs in Active IQ Unified Manager

You need to create access endpoints or logical interfaces (LIFs), which are required for provisioning Storage Virtual Machines (SVMs), LUNs, and file shares. You can view,

create, modify, and delete the access endpoints for the SVMs, LUNs, or file shares in your Active IQ Unified Manager environment.

View access endpoints

You can view a list of the access endpoints in your Unified Manager environment by using the following method. To query a list of access endpoints of a particular SVM, LUN, or file share, you need to enter the unique identifier for the SVM, LUN, or file share. You can also enter the unique access endpoint key to retrieve the details of the particular access endpoint.

Category	HTTP verb	Path

storage-provider	GET	`/storage-provider/access-endpoints`
		
		`/storage-provider/access-endpoints/{key}`

Add access endpoints

You can create custom access endpoints and assign required properties to it. You must enter the details of the access endpoint that you want to create as the input parameters. You can use this API, or the System Manager or ONTAP CLI to create an access endpoint on each node. Both IPv4 and IPv6 addresses are supported for access endpoints creation.



You must configure your SVM with a minimum number of access endpoints per node for successful provisioning of LUNs and file shares. You should configure your SVM with at least two access endpoints per node, one supporting CIFS and/or NFS protocol, another supporting iSCSI or FCP protocol.

Category	HTTP verb	Path
storage-provider	POST	`/storage-provider/access-endpoints`

Delete access endpoints

You can delete a specific access endpoint by using the following method. You need to provide the access endpoint key as an input parameter to delete a particular access endpoint.

Category	HTTP verb	Path
storage-provider	DELETE	`/storage-provider/access-endpoints/{key}`

Modify access endpoints

You can modify an access endpoint and update its properties by using the following method. You need to provide the access endpoint key to modify a particular access endpoint. You also need to enter the property that you want to update, along with its value.

Category	HTTP verb	Path
storage-provider	PATCH	`/storage-provider/access-endpoints/{key}`

Manage Active Directory mappings using APIs in Active IQ Unified Manager

You can use the APIs listed here to manage Active Directory mappings on the SVM that are required for provisioning CIFS shares on the SVMs. Active Directory mappings need to be configured for mapping the SVMs with ONTAP.

View Active Directory mappings

You can view the configuration details of the Active Directory mappings for an SVM by using the following method. For viewing the Active Directory mappings on an SVM, you need to enter the SVM key. For querying the details of a particular mapping, you must enter the mapping key.

Category	HTTP verb	Path

storage-provider	GET	`/storage-provider/active-directories-mappings`
		
		`/storage-provider/active-directories-mappings/{key}`

Add Active Directory mapping

You can create Active Directory mappings on an SVM by using the following method. You must enter the mapping details as the input parameters.

Category	HTTP verb	Path
storage-provider	POST	`/storage-provider/active-directories-mappings`

Manage file shares using APIs in Active IQ Unified Manager

You can use the `/storage-provider/file-shares` API to view, add, modify, and delete the CIFS and NFS file share volumes in your data center environment.

Before provisioning the file share volumes, ensure that the SVM has been created and provisioned with the supported protocols. If you are assigning Performance Service Levels (PSLs) or Storage Efficiency Policies (SEPs), while provisioning, the PSLs or SEPs should be created before creating the file shares.

View file shares

You can use the following method to view the file share volumes available in your Unified Manager environment. When you have added an ONTAP cluster as a datasource on Active IQ Unified Manager, the storage workloads for those clusters are automatically added to your Unified Manager instance. This API retrieves the file shares automatically and manually added to your Unified Manager instance. You can view the details of a specific file share by running this API with the file share key.

Category	HTTP verb	Path

storage-provider	GET	`/storage-provider/file-shares`
		
		`/storage-provider/file-shares/{key}`

Add file shares

You can use the following method to add CIFS and NFS file shares in your SVM. You must enter the details of the file share that you want to create, as the input parameters. You cannot use this API for adding FlexGroup volumes.

Category	HTTP verb	Path
storage-provider	POST	`/storage-provider/file-shares`

Depending on whether the access control list (ACL) parameters or the export policy parameters are provided, CIFS shares or NFS file shares are created. If you do not provide the values for the ACL parameters, CIFS shares are not created, and NFS shares are created by default, providing access to all.

Creating data-protection volumes: When you add file shares to your SVM, the type of the volume that is mounted, by default, is `rw` (read-write). For creating data-protection (DP) volumes, specify `dp` as the value for the `type` parameter.

Delete file shares

You can use the following method to delete a specific file share. You need to enter the file share key as an input parameter to delete a particular file share.

Category	HTTP verb	Path
storage-provider	DELETE	`/storage-provider/file-shares/{key}`

Modify file shares

You can use the following method to modify a file share and update its properties.

You need to provide the file share key to modify a particular file share. Additionally, you need to enter the property that you want to update, along with its value.

Note that you can update only one property at a single invocation of this API. For multiple updates, you need to run this API as many times.

Category	HTTP verb	Path
storage-provider	PATCH	`/storage-provider/file-shares/{key}`

Manage LUNs using APIs in Active IQ Unified Manager

You can use the `/storage-provider/luns` API to view, add, modify, and delete the LUNs in your data center environment.

Before provisioning the LUNs, ensure that the SVM has been created and provisioned with the supported protocols. If you are assigning Performance Service Levels (PSLs) or Storage Efficiency Policies (SEPs), while provisioning, the PSLs or SEPs should be created before creating the LUN.

View LUNs

You can use the following method to view the LUNs in your Unified Manager environment. When you have added an ONTAP cluster as a datasource on Active IQ Unified Manager, the storage workloads for those clusters are automatically added to your Unified Manager instance. This API retrieves all the LUNs automatically and manually added to your Unified Manager instance. You can view the details of a specific LUN by running this API with the LUN key.

Category	HTTP verb	Path

storage-provider	GET	`/storage-provider/luns`
		
		`/storage-provider/luns/{key}`

Add LUNs

You can use the following method to add LUNs to your SVMs.

Category	HTTP verb	Path
storage-provider	POST	`/storage-provider/luns`

In your cURL request, if you provide a value for the optional parameter volume_name_tag in the input, then that value is used while naming the volume during the LUN creation. This tag enables searching the volume easily. If you provide the volume key in the request, the tagging is skipped.

Delete LUNs

You can use the following method to delete a specific LUN. You need to provide the LUN key to delete a particular LUN.

If you have created a volume in ONTAP and then provisioned LUNs through Unified Manager on that volume, when you delete all of the LUNs by using this API, the volume also gets deleted from the ONTAP cluster.

Category	HTTP verb	Path
storage-provider	DELETE	`/storage-provider/luns/{key}`

Modify LUNs

You can use the following method to modify a LUN and update its properties. You need to provide the LUN key to modify a particular LUN. You also need to enter the LUN property that you want to update, along with its value. For updating LUN arrays by using this API, you should review the recommendations in "Recommendations for using the APIs".



You can update only one property at a single invocation of this API. For multiple updates, you need to run this API as many times.

Category	HTTP verb	Path
storage-provider	PATCH	`/storage-provider/luns/{key}`

Manage Performance Service Levels using APIs in Active IQ Unified Manager

You can view, create, modify, and delete Performance Service Levels by using the storage provider APIs for on your Active IQ Unified Manager.

View Performance Service Levels

You can use the following method to view the Performance Service Levels for assigning them to storage workloads. The API lists all of the system-defined and user-created Performance Service Levels, and retrieves the attributes of all of the Performance Service Levels. If you want to query a specific Performance Service Level, you need to enter the unique ID of the Performance Service Level to retrieve its details.

Category	HTTP verb	Path

storage-provider	GET	`/storage-provider/performance-service-levels`
		
		`/storage-provider/performance-service-levels/{key}`

Add Performance Service Levels

You can use the following method to create custom Performance Service Levels and assign them to your storage workloads if the system-defined Performance Service Levels do not meet the required service level objectives (SLOs) for the storage workloads. Enter the details for the Performance Service Level that you want to create. For the IOPS properties, ensure that you enter valid range of values.

Category	HTTP verb	Path
storage-provider	POST	`/storage-provider/performance-service-levels`

Delete Performance Service Levels

You can use the following method to delete a specific Performance Service Level. You cannot delete a

Performance Service Level if it is assigned to a workload, or if it is the only available Performance Service Level. You need to provide the unique ID of the Performance Service Level as an input parameter to delete a particular Performance Service Level.

Category	HTTP verb	Path
storage-provider	DELETE	`/storage-provider/performance-service-levels/{key}`

Modify Performance Service Levels

You can use the following method to modify a Performance Service Level and update its properties. You cannot modify a Performance Service Level that is system-defined or is assigned to a workload. You need to provide the unique ID of the to modify a particular Performance Service Level. You must also enter the IOPS property that you want to update, along with a valid value.

Category	HTTP verb	Path
storage-provider	PATCH	`/storage-provider/performance-service-levels/{key}`

View aggregate capabilities based on Performance Service Levels

You can use the following method to query the aggregate capabilities based on Performance Service Levels. This API returns the list of aggregates available in your data center and indicates the capabilities in terms of the Performance Service Levels that can be supported in those aggregates. While provisioning workloads on a volume, you can view the capability of an aggregate to support a particular Performance Service Level, and provision workloads based on that capability. Your ability to specify the aggregate is available only when you are provisioning a workload by using APIs. This functionality is not available on the Unified Manager web UI.

Category	HTTP verb	Path

storage-provider	GET	`/storage-provider/aggregate-capabilities`
		
		`/storage-provider/aggregate-capabilities/{key}`

Manage Storage Efficiency Policies using APIs in Active IQ Unified Manager

You can view, create, modify, and delete Storage Efficiency Policies by using the storage provider APIs.

Note the following points:

- It is not mandatory to assign a Storage Efficiency Policy while creating a workload on Unified Manager.
- You cannot unassign a Storage Efficiency Policy from a workload after a policy is assigned to it.

- If a workload has some storage settings specified on ONTAP volumes, such as deduplication and compression, those settings can be overwritten by the settings specified in the Storage Efficiency Policy that you apply when you add the storage workloads on Unified Manager.

View Storage Efficiency Policies

You can use the following method to view the Storage Efficiency Policies before assigning them to storage workloads. This API lists all of the system-defined and user-created Storage Efficiency Policies, and retrieves the attributes of all of the Storage Efficiency Policies. If you want to query a specific Storage Efficiency Policy, you need to enter the unique ID of the policy to retrieve its details.

Category	HTTP verb	Path

storage-provider	GET	`/storage-provider/storage-efficiency-policies`
		
		`/storage-provider/storage-efficiency-policies/{key}`

Add Storage Efficiency Policies

You can use the following method to create custom Storage Efficiency Policies, and assign them to your storage workloads if the system-defined policies do not meet the provisioning requirements for your storage workloads. Enter the details of the Storage Efficiency Policy that you want to create, as input parameters.

Category	HTTP verb	Path
storage-provider	POST	`/storage-provider/storage-efficiency-policies`

Delete Storage Efficiency Policies

You can use the following method to delete a specific Storage Efficiency Policy. You cannot delete a Storage Efficiency Policy if it is assigned to a workload, or if it is the only available Storage Efficiency Policy. You need to provide the unique ID of the Storage Efficiency Policy as an input parameter to delete a particular Storage Efficiency Policy.

Category	HTTP verb	Path
storage-provider	DELETE	`/storage-provider/storage-efficiency-policies/{key}`

Modify Storage Efficiency Policies

You can use the following method to modify a Storage Efficiency Policy and update its properties. You cannot modify a Storage Efficiency Policy that is system-defined or is assigned to a workload. You need to provide the unique ID of the Storage Efficiency Policy to modify a particular Storage Efficiency Policy. Additionally, you need to provide the property that you want to update, along with its value.

Category	HTTP verb	Path
storage-provider	PATCH	`/storage-provider/storage-efficiency-policies/{key}`

Common workflows for storage management

Understand common API workflows for storage management in Active IQ Unified Manager

The common workflows provide client application developers with examples of how Active IQ Unified Manager APIs can be called by a client application to execute common storage management functions. This section contains some of these sample workflows.

The workflows describe some of the commonly used storage management use cases along with sample codes for you to use. Each of the tasks is described using a workflow process consisting of one or more API calls.

Understand the API calls used in the workflows

You can view the online documentation page from your Unified Manager instance that includes the details of every REST API call. This document does not repeat the details of the online documentation. Each API call used in the workflow samples in this document includes only the information you need to locate the call on the documentation page. After locating a specific API call, you can review the complete details of the call, including the input parameters, output formats, HTTP status codes, and request processing type.

The following information is included for each API call within a workflow to help locate the call on the documentation page:

- Category: The API calls are organized on the documentation page into functionally related areas or categories. To locate a specific API call, scroll to the bottom of the page and click the applicable API category.
- HTTP verb (call): The HTTP verb identifies the action performed on a resource. Each API call is executed through a single HTTP verb.
- Path: The path determines the specific resource which the action applies to as part of performing a call. The path string is appended to the core URL to form the complete URL identifying the resource.

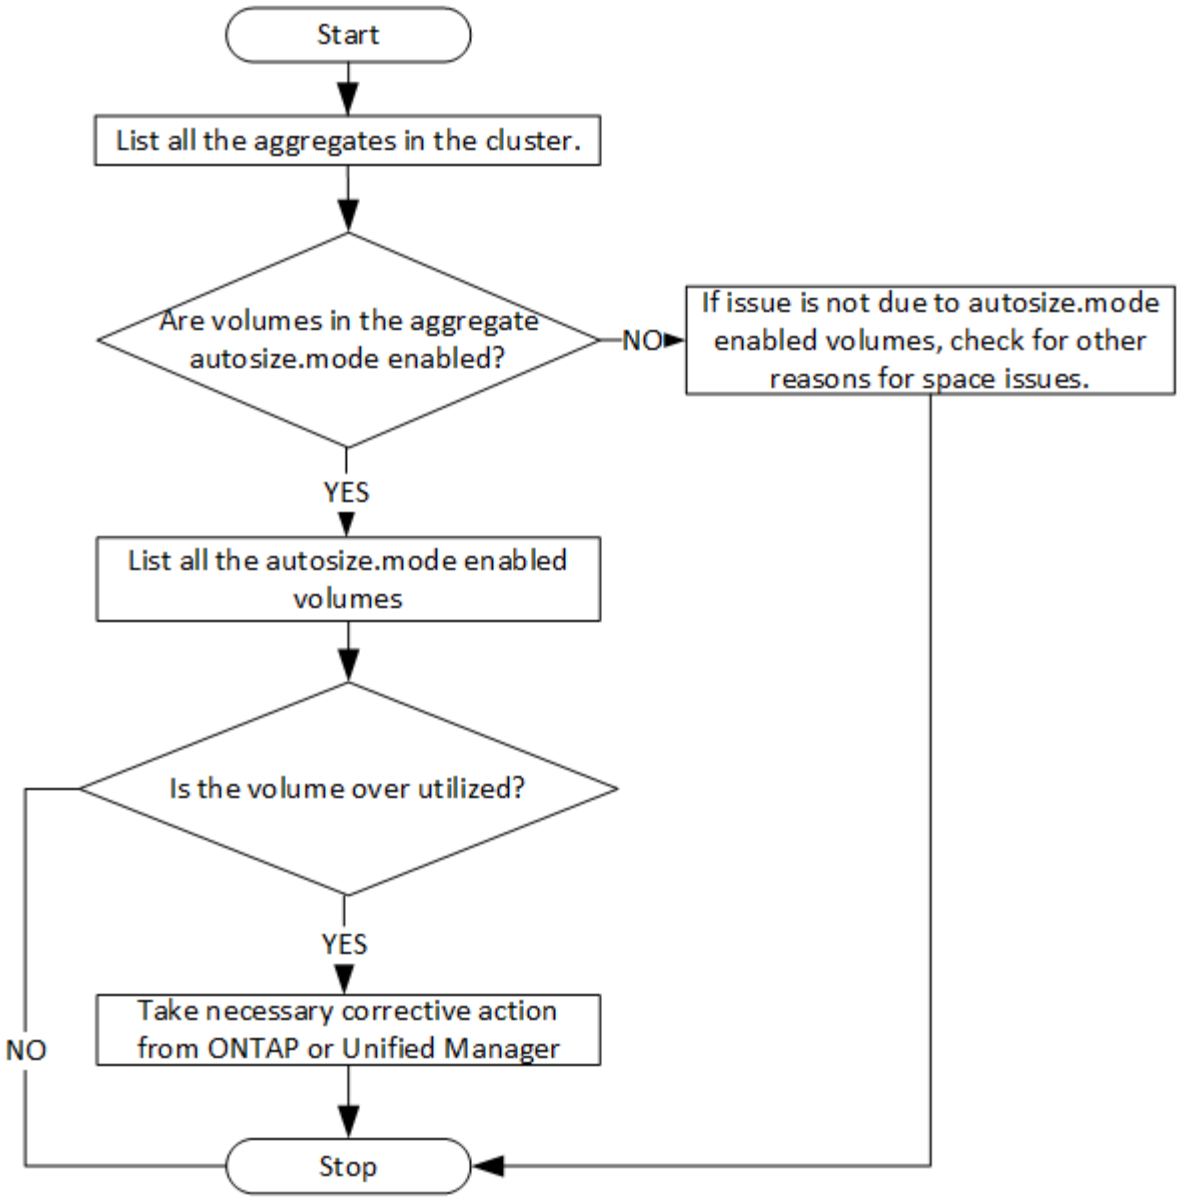
Determine space issues in aggregates using Active IQ Unified Manager APIs

You can use the data center APIs in Active IQ Unified Manager to monitor the availability and utilization of space in your volumes. You can determine space issues in your volume and identify storage resources that are overutilized or underutilized.

The data center APIs for aggregates retrieve the relevant information about available and used space, and space saving efficiency settings. You can also filter the retrieved information based on specified attributes.

One method to determine any lack of space in your aggregates is to verify whether there are volumes in your environment with autosize-mode enabled. You should then identify which volumes are being over-utilized and perform any corrective actions.

The following flowchart illustrates the process of retrieving information about volumes with autosize-mode enabled:



This flow assumes that the clusters have already been created in ONTAP and added to Unified Manager.

1. Obtain the cluster key, unless you know the value:

Category	HTTP verb	Path
datacenter	GET	`/datacenter/cluster/clusters`

2. Using the cluster key as the filter parameter, query the aggregates on that cluster.

Category	HTTP verb	Path
datacenter	GET	`/datacenter/storage/aggregates`

3. From the response, analyze the space usage of the aggregates and determine which aggregates have space issues. For each aggregate with space issue, obtain the aggregate key from the same JSON output.

4. Using each aggregate key, filter all the volumes that have the value for the autosize.mode parameter as `grow`.

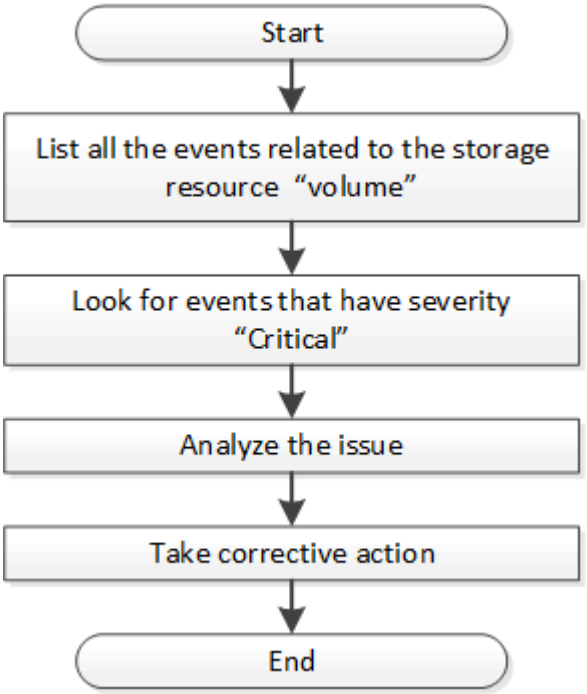
Category	HTTP verb	Path
datacenter	GET	`/datacenter/storage/volumes`

5. Analyze which volumes are being over-utilized.
6. Perform any necessary corrective action, such as moving the volume across aggregates, to address the space issues in your volume. You can perform these actions from ONTAP or Unified Manager web UI.

Determine storage object issues using Active IQ Unified Manager event APIs

When a storage object in your data center crosses a threshold, you get a notification about that event. Using this notification, you can analyze the issue and take corrective action by using the `events` APIs.

This workflow takes the example of a volume as the resource object. You can use the `events` APIs to retrieve the list of events related to a volume, analyze the critical issues for that volume, and then take corrective actions to rectify the issue.



Follow these steps to determine the issues in your volume before taking remedial steps.

Steps

1. Analyze the critical Active IQ Unified Manager events notifications for the volumes in your data center.
2. Query all the events for the volumes by using the following parameters in the /management-server/events API:
   `"resource_type": "volume"`
   `"severity": "critical"`

Category	HTTP verb	Path
management-server	GET	/management-server/events

3. View the output and analyze the issues in the specific volumes.
4. Perform the necessary actions by using the Unified Manager REST APIs or web UI to resolve the issues.

Troubleshoot ONTAP volumes using Active IQ Unified Manager gateway APIs

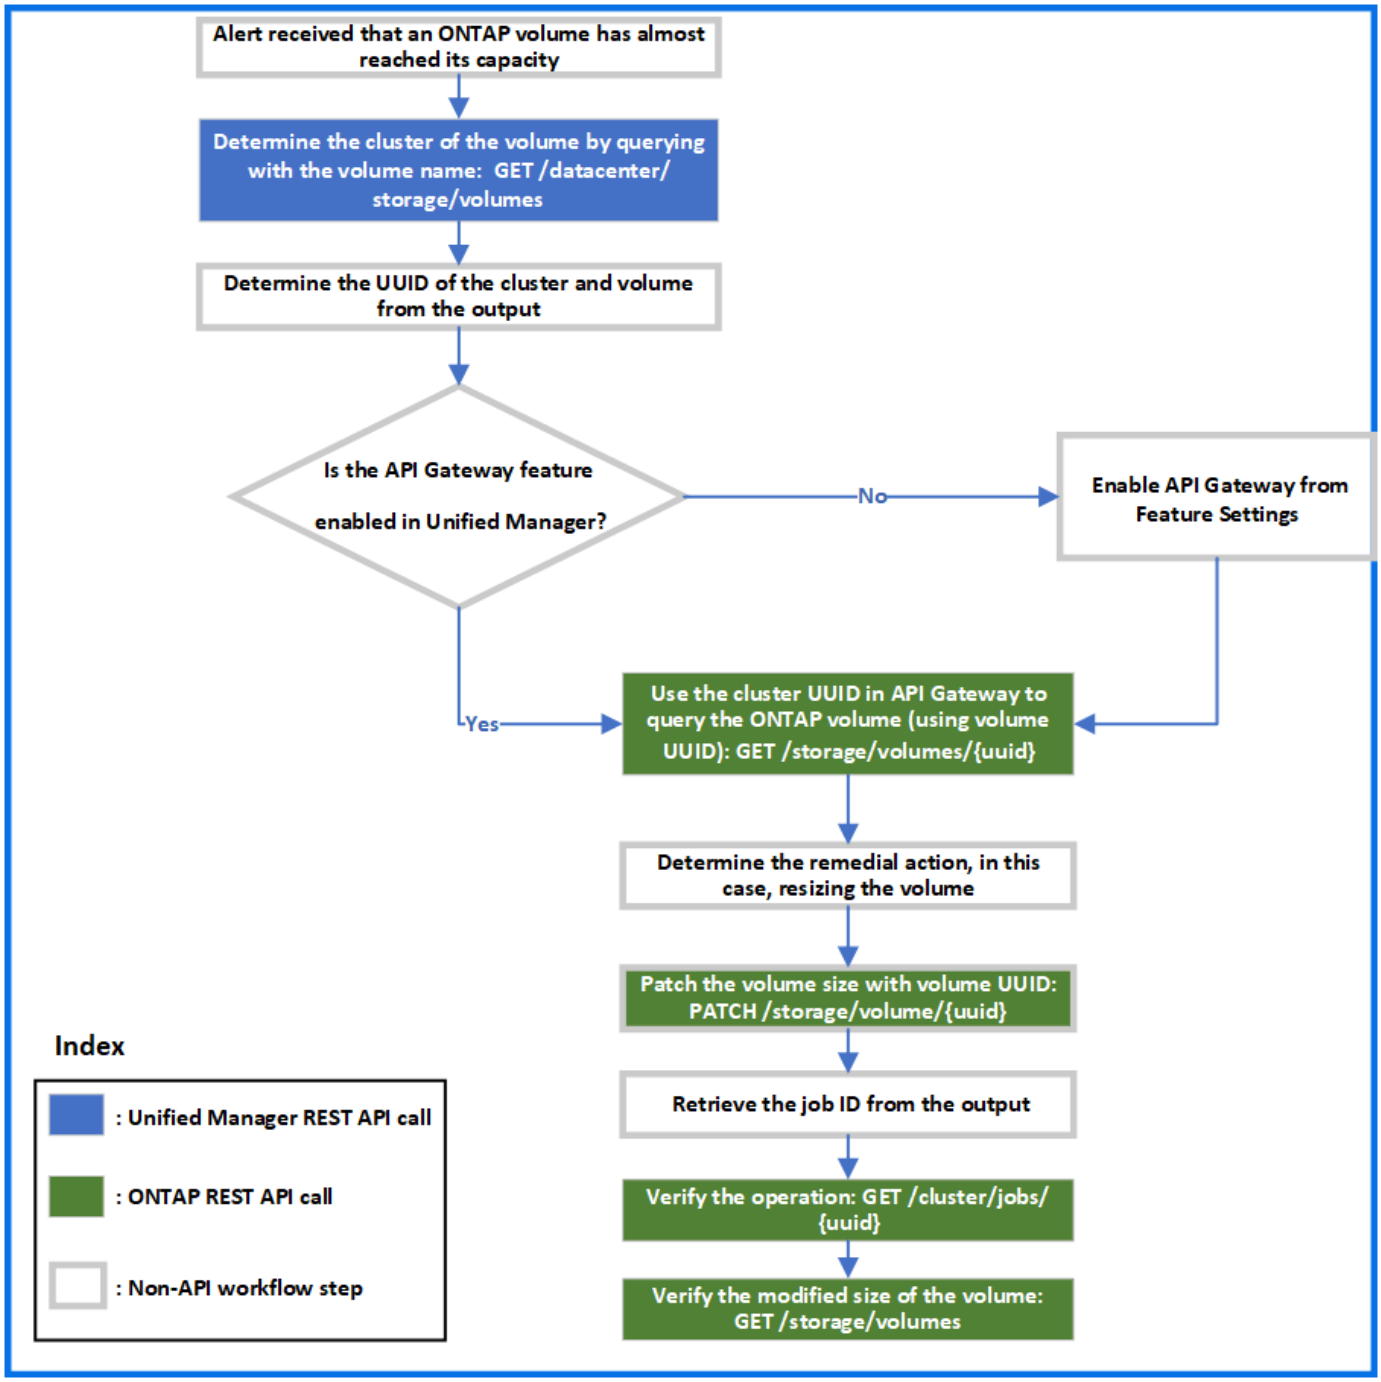
The gateway APIs act as a gateway to invoke ONTAP APIs to query information about your ONTAP storage objects and take remedial measures to address the reported issues.

This workflow takes up a sample use case in which an event is raised when an ONTAP volume almost reaches its capacity. The workflow also demonstrates how to address this issue by invoking a combination of Active IQ Unified Manager and ONTAP REST APIs.

Before running the workflow steps, ensure the following:

- You are aware of the gateway APIs and how they are used. For information, see Accessing ONTAP APIs through proxy access.
- You are aware of the usage of ONTAP REST APIs. For information about using ONTAP REST APIs, see ONTAP Automation documentation.
- You are an Application Administrator.
- The cluster on which you want to run the REST API operations is supported by ONTAP 9.5 or later, and the cluster is added to Unified Manager over HTTPS.

The following diagram illustrates each step in the workflow for troubleshooting the issue of ONTAP volume capacity use.



The workflow covers the invocation points of both the Unified Manager and ONTAP REST APIs.

1. Note the volume name from the event notifying the volume capacity utilization.
2. By using the volume name as the value in the name parameter, query the volume by running the following Unified Manager API.

Category	HTTP verb	Path
datacenter	GET	`/datacenter/storage/volumes`

3. Retrieve the cluster UUID and volume UUID from the output.

4. On the Unified Manager web UI, navigate to **General** > **Feature Settings** > **API Gateway** to verify whether the API Gateway feature is enabled. Unless it is enabled, the APIs under the gateway category are not available for you to invoke. Enable the feature if it is disabled.
5. Use the cluster UUID to run the ONTAP API `/storage/volumes/{uuid}` through API gateway. The query returns the volume details when the volume UUID is passed as the API parameter.

   For running the ONTAP APIs through API gateway, the Unified Manager credentials are passed internally for authentication, and you do not need to run an additional authentication step for individual cluster access.

Category	HTTP verb	Path
Unified Manager: gateway ONTAP: storage	GET	Gateway API: `/gateways/{uuid}/{path}` ONTAP API: `/storage/volumes/{uuid}`

   In /gateways/{uuid}/{path}, the value for {uuid} must be replaced with the cluster UUID on which the REST operation is to be performed. {path} must be replaced by the ONTAP REST URL /storage/volumes/{uuid}.

   The appended URL is: `/gateways/{cluster_uuid}/storage/volumes/{volume_uuid}`

   On running the GET operation, the generated URL is: `GEThttps://<hostname\>/api/gateways/<cluster_UUID\>/storage/volumes/{volume_uuid}`

   Sample cURL command

```
curl -X GET "https://<hostname>/api/gateways/1cd8a442-86d1-11e0-ae1c-
9876567890123/storage/volumes/028baa66-41bd-11e9-81d5-00a0986138f7"
-H "accept: application/hal+json" -H "Authorization: Basic
<Base64EncodedCredentials>"
```

6. From the output, determine the size, usage, and remedial measure to be taken. In this workflow, the remedial measure taken is resizing the volume.
7. Use the cluster UUID and run the following ONTAP API through the API gateway to resize the volume. For information about the input parameters for the gateway and ONTAP APIs, see step 5.

Category	HTTP verb	Path
Unified Manager: gateway ONTAP: storage	PATCH	Gateway API: `/gateways/{uuid}/{path}` ONTAP API: `/storage/volumes/{uuid}`

   Along with the cluster UUID and volume UUID, you must enter a value for the size parameter for resizing the volume. Ensure to enter the value *in bytes*. For example, if you

want to increase the size of a volume from 100 GB to 120 GB, enter the value for parameter size at the end of the query: `-d {\"size\": 128849018880}"`

Sample cURL command

```
curl -X PATCH "https://<hostname>/api/gateways/1cd8a442-86d1-11e0-ae1c-
9876567890123/storage/volumes/028baa66-41bd-11e9-81d5-00a0986138f7" -H
    "accept: application/hal+json" -H "Authorization: Basic
<Base64EncodedCredentials>" -d
    {\"size\": 128849018880}"
```

The JSON output returns a Job UUID.

8. Verify whether the job ran successfully by using the Job UUID. Use the cluster UUID and Job UUID to run the following ONTAP API through the API gateway. For information about the input parameters for the gateway and ONTAP APIs, see step 5.

Category	HTTP verb	Path

Unified Manager: gateway
 	
ONTAP: cluster
	GET	Gateway API:
		`/gateways/{uuid}/{path}`
		
		ONTAP API:
		`/cluster/jobs/{uuid}`

The HTTP codes returned are the same as the ONTAP REST API HTTP status codes.

9. Run the following ONTAP API to query the details of the resized volume. For information about the input parameters for the gateway and ONTAP APIs, see step 5.

Category	HTTP verb	Path

Unified Manager: gateway
 	
ONTAP: storage
	GET	Gateway API:
		`/gateways/{uuid}/{path}`
		
		ONTAP API:
		`/storage/volumes/{uuid}`

The output displays an increased volume size of 120 GB.

Workflows for workload management

Provision and modify storage workloads in Active IQ Unified Manager

Using Active IQ Unified Manager, you can provision and modify storage workloads (LUNs, NFS file shares, and CIFS shares). Provisioning consists of multiple steps, from the creation of the Storage Virtual Machine (SVM) to applying Performance Service Level and Storage Efficiency Policies on the storage workloads. Modifying workloads consist of the steps for modifying specific parameters and enabling additional features on them.

The following workflows are described:

- Workflow for provisioning Storage Virtual Machines (SVMs) on Unified Manager.

  this workflow is required to be performed before provisioning LUNs or file shares on Unified Manager.

- Provisioning file shares.
- Provisioning LUNs.
- Modifying LUNs and file shares (by using the example for updating the Performance Service Level parameter for the storage workloads).
- Modifying an NFS file share to support CIFS protocol
- Modifying workloads to upgrade QoS to AQoS

For each provisioning workflow (LUN and file shares), ensure you must have completed the workflow for verifying the SVMs on the clusters.

You must also read the recommendations and limitations before using each API in the workflows. The relevant details of the APIs are available in their individual sections listed in the related concepts and references.

Verify SVMs on ONTAP clusters using Active IQ Unified Manager APIs

Before provisioning file shares or LUNs, you must verify whether the clusters have Storage Virtual Machines (SVMs) created on them.

The workflow assumes that ONTAP clusters to have been added to Unified Manager, and the cluster key has been obtained. Clusters should have the required licenses for provisioning LUNs and file shares on them.

1. Verify whether the cluster has an SVM created.

Category	HTTP verb	Path
datacenter	GET	`/datacenter/svm/svms` `/datacenter/svm/svms/{key}`

Sample cURL

```
curl -X GET "https://<hostname>/api/datacenter/svm/svms" -H "accept:
application/json" -H "Authorization: Basic <Base64EncodedCredentials>"
```

2. If the SVM key is not returned, then create the SVM. For creating the SVMs, you require the cluster key on which you provision the SVM. You also need to specify the SVM name. Follow these steps.

Category	HTTP verb	Path
datacenter	GET	`/datacenter/cluster/clusters` `/datacenter/cluster/clusters/{key}`

Get the cluster key.

Sample cURL

```
curl -X GET "https://<hostname>/api/datacenter/cluster/clusters" -H
"accept: application/json" -H "Authorization: Basic
<Base64EncodedCredentials>"
```

3. From the output, get the cluster key, and then use it as an input for creating the SVM.

> While creating the SVM, ensure that it supports all the protocols required for provisioning LUNs and file shares on them, for example, CIFS, NFS, FCP, and iSCSI. The provisioning workflows might fail if the SVM does not support the required services. It is recommended that the services for the respective types of workloads are also enabled on the SVM.

Category	HTTP verb	Path
datacenter	POST	`/datacenter/svm/svms`

Sample cURL

Enter the SVM object details as input parameters.

```
curl -X POST "https://<hostname>/api/datacenter/svm/svms" -H "accept:
application/json" -H "Content-Type: application/json" -H "Authorization:
Basic <Base64EncodedCredentials>" "{ \"aggregates\": [ { \"_links\": {},
\"key\": \"1cd8a442-86d1,type=objecttype,uuid=1cd8a442-86d1-11e0-ae1c-
9876567890123\",
\"name\": \"cluster2\", \"uuid\": \"02c9e252-41be-11e9-81d5-
00a0986138f7\" } ],
\"cifs\": { \"ad_domain\": { \"fqdn\": \"string\", \"password\":
\"string\",
\"user\": \"string\" }, \"enabled\": true, \"name\": \"CIFS1\" },
\"cluster\": { \"key\": \"1cd8a442-86d1-11e0-ae1c-
123478563412,type=object type,uuid=1cd8a442-86d1-11e0-ae1c-
9876567890123\" },
\"dns\": { \"domains\": [ \"example.com\", \"example2.example3.com\" ],
\"servers\": [ \"10.224.65.20\", \"2001:db08:a0b:12f0::1\" ] },
\"fcp\": { \"enabled\": true }, \"ip_interface\": [ { \"enabled\": true,
\"ip\": { \"address\": \"10.10.10.7\", \"netmask\": \"24\" },
\"location\": { \"home_node\": { \"name\": \"node1\" } }, \"name\":
```

```
\"dataLif1\" } ], \"ipspace\": { \"name\": \"exchange\" },
\"iscsi\": { \"enabled\": true }, \"language\": \"c.utf_8\",
\"ldap\": { \"ad_domain\": \"string\", \"base_dn\": \"string\",
\"bind_dn\": \"string\", \"enabled\": true, \"servers\": [ \"string\" ]
},
\"name\": \"svm1\", \"nfs\": { \"enabled\": true },
\"nis\": { \"domain\": \"string\", \"enabled\": true,
\"servers\": [ \"string\" ] }, \"nvme\": { \"enabled\": true },
\"routes\": [ { \"destination\": { \"address\": \"10.10.10.7\",
\"netmask\": \"24\" }, \"gateway\": \"string\" } ],
\"snapshot_policy\": { \"name\": \"default\" },
\"state\": \"running\", \"subtype\": \"default\"}"
```

The JSON output displays a Job object key that you can use to verify the SVM that you created.

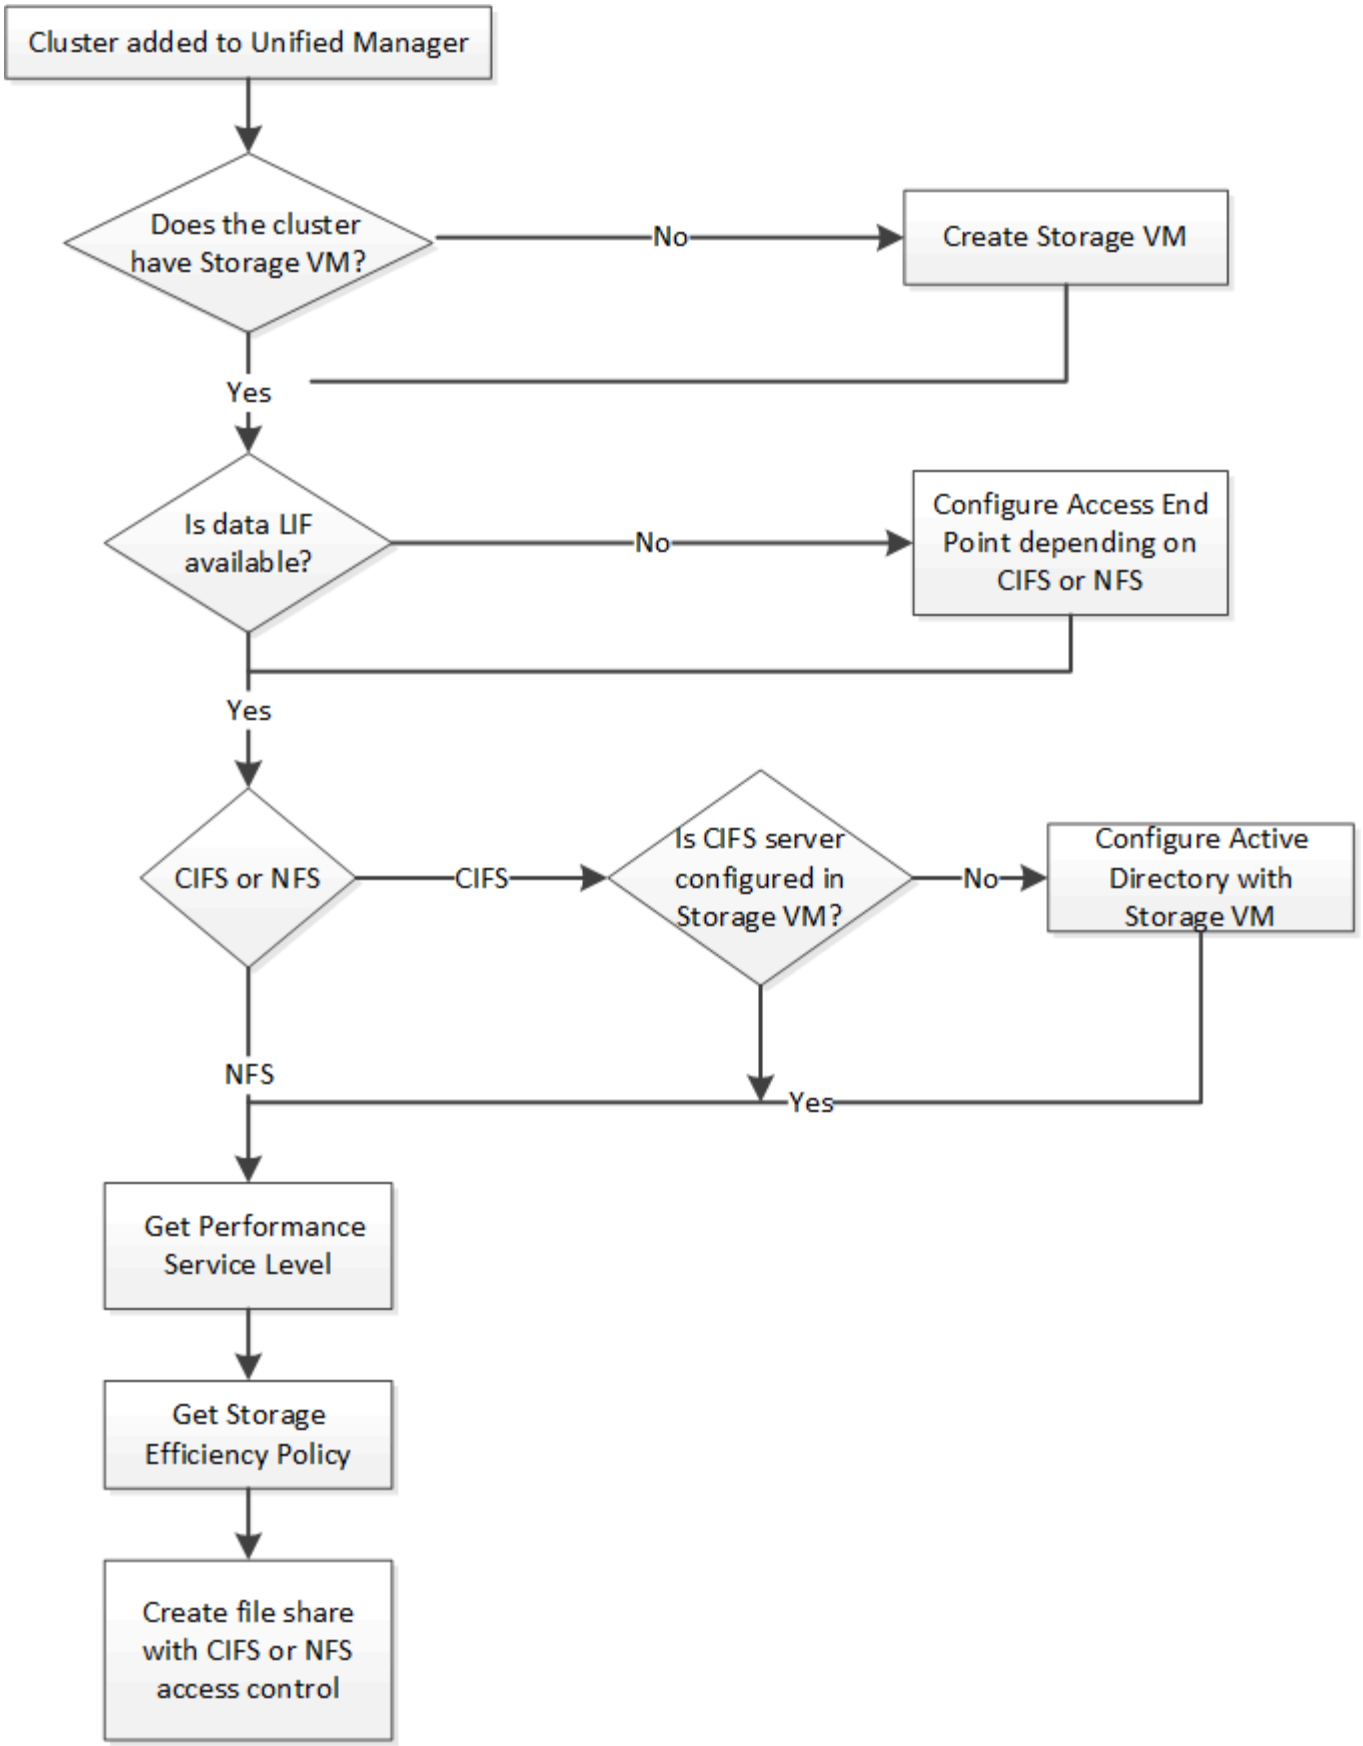
4. Verify the SVM creation by using the job object key for query. If the SVM is created successfully, the SVM key is returned in the response.

Category	HTTP verb	Path
management-server	GET	`/management-server/jobs/{key}`

Provision CIFS and NFS file shares using Active IQ Unified Manager APIs

You can provision CIFS shares and NFS file shares on your Storage Virtual Machines (SVMs) by using the provisioning APIs provided as a part of Active IQ Unified Manager. This provisioning workflow details the steps for retrieving the keys of the SVMs, Performance Service Levels, and Storage Efficiency Policies before creating the file shares.

The following diagram illustrates each step in a file share provisioning workflow. It includes provisioning both CIFS shares and NFS file shares.



Ensure the following:

- ONTAP clusters have been added to Unified Manager, and the cluster key has been obtained.
- SVMs have been created on the clusters.
- The SVMs support CIFS and NFS services. Provisioning file shares might fail if the SVMs do not support the required services.
- The FCP port is online for port provisioning.

1. Determine whether Data LIFs or access endpoints are available on the SVM on which you want to create the CIFS share. Get the list of available access endpoints on the SVM:

Category	HTTP verb	Path
storage-provider	GET	`/storage-provider/access-endpoints` `/storage-provider/access-endpoints/{key}`

Sample cURL

```
curl -X GET "https://<hostname>/api/storage-provider/access-endpoints?resource.key=7d5a59b3-953a-11e8-8857-00a098dcc959" -H "accept: application/json" -H "Authorization: Basic <Base64EncodedCredentials>"
```

2. If your access endpoint is available on the list, obtain the access endpoint key, else create the access endpoint.

Ensure that you create access endpoints that have the CIFS protocol enabled on them. Provisioning CIFS shares fails unless you have created an access endpoint with the CIFS protocol enabled on it.

Category	HTTP verb	Path
storage-provider	POST	`/storage-provider/access-endpoints`

Sample cURL

You must enter the details of the access endpoint that you want to create, as input parameters.

```
curl -X POST "https://<hostname>/api/storage-provider/access-endpoints" -H "accept: application/json" -H "Content-Type: application/json" -H "Authorization: Basic <Base64EncodedCredentials>"
{ \"data_protocols\": \"nfs\",
\"fileshare\": { \"key\": \"cbd1757b-0580-11e8-bd9d-00a098d39e12:type=volume,uuid=f3063d27-2c71-44e5-9a69-a3927c19c8fc\" },
\"gateway\": \"10.132.72.12\",
\"ip\": { \"address\": \"10.162.83.26\",
\"ha_address\": \"10.142.83.26\",
```

```
\"netmask\": \"255.255.0.0\" },
\"lun\": { \"key\": \"cbd1757b-0580-11e8-bd9d-
00a098d39e12:type=lun,uuid=d208cc7d-80a3-4755-93d4-5db2c38f55a6\" },
\"mtu\": 15000, \"name\": \"aep1\",
\"svm\": { \"key\": \"cbd1757b-0580-11e8-bd9d-
00a178d39e12:type=vserver,uuid=1d1c3198-fc57-11e8-99ca-00a098d38e12\" },
\"vlan\": 10}"
```

The JSON output displays a Job object key that you can use to verify the access endpoint that you created.

3. Verify the access endpoint:

Category	HTTP verb	Path
management-server	GET	`/management-server/jobs/{key}`

4. Determine whether you have to create a CIFS share or an NFS file share. For creating CIFS shares, follow these substeps:
   a. Determine whether the CIFS server is configured on your SVM, that is determine whether an Active Directory mapping is created on the SVM.

Category	HTTP verb	Path
storage-provider	GET	`/storage-provider/active-directories-mappings`

   b. If the Active Directory mapping is created, take the key, else create the Active Directory mapping on the SVM.

Category	HTTP verb	Path
storage-provider	POST	`/storage-provider/active-directories-mappings`

Sample cURL

You must enter the details for creating the Active Directory mapping, as the input parameters.

```
curl -X POST "https://<hostname>/api/storage-provider/active-
directories-mappings" -H "accept: application/json" -H "Content-Type:
application/json" -H "Authorization: Basic
<Base64EncodedCredentials>"
{ \"_links\": {},
\"dns\": \"10.000.000.000\",
\"domain\": \"example.com\",
\"password\": \"string\",
```

```
\"svm\": { \"key\": \"9f4ddea-e395-11e9-b660-
005056a71be9:type=vserver,uuid=191a554a-f0ce-11e9-b660-005056a71be9\"
},
\"username\": \"string\"}"
```

This is a synchronous call and you can verify the creation of the Active Directory mapping in the output. In case of an error, the error message is displayed for you to troubleshoot and rerun the request.

5. Obtain the SVM key for the SVM on which you want to create the CIFS share or the NFS file share, as described in the *Verifying SVMs on clusters* workflow topic.
6. Obtain the key for the Performance Service Level by running the following API and retrieving the key from the response.

Category	HTTP verb	Path
storage-provider	GET	`/storage-provider/performance-service-levels`

You can retrieve the details of the system-defined Performance Service Levels by setting the `system_defined` input parameter to `true`. From the output, obtain the key of the Performance Service Level that you want to apply on the file share.

7. Optionally, obtain the Storage Efficiency Policy key for the Storage Efficiency Policy that you want to apply on the file share by running the following API and retrieving the key from the response.

Category	HTTP verb	Path
storage-provider	GET	`/storage-provider/storage-efficiency-policies`

8. Create the file share. You can create a file share that supports both CIFS and NFS by specifying the access control list and export policy. The following substeps provide information if you want to create a file share for supporting only one of the protocols on the volume. You can also update an NFS file share to include the access control list after you have created the NFS share. For information, see the *Modifying storage workloads* topic.

   a. For creating only a CIFS share, gather the information about access control list (ACL). For creating the CIFS share, provide valid values for the following input parameters. For each user group that you assign, an ACL is created when a CIFS/SMB share is provisioned. Based on the values you enter for ACL and Active Directory mapping, the access control and mapping are determined for the CIFS share when it is created.

      A cURL command with sample values

```
{
  "access_control": {
    "acl": [
      {
```

```
        "permission": "read",
        "user_or_group": "everyone"
      }
    ],
    "active_directory_mapping": {
      "key": "3b648c1b-d965-03b7-20da-61b791a6263c"
    },
```

b. For creating only an NFS file share, gather the information about the export policy. For creating the NFS file share, provide valid values for the following input parameters. Based on your values, the export policy is attached with the NFS file share when it is created.

> While provisioning the NFS share, you can either create an export policy by providing all the required values or provide the export policy key and reuse an existing export policy. If you want to reuse an export policy for the storage VM, you need to add the export policy key. Unless you know the key, you can retrieve the export policy key by using the `/datacenter/protocols/nfs/export-policies` API. For creating a new policy, you must enter the rules as displayed in the following sample. For the entered rules, the API tries to search for an existing export policy by matching the host, storage VM, and rules. If there is an existing export policy, it is used. Otherwise a new export policy is created.

A cURL command with sample values

```
"export_policy": {
      "key": "7d5a59b3-953a-11e8-8857-
00a098dcc959:type=export_policy,uuid=1460288880641",
      "name_tag": "ExportPolicyNameTag",
      "rules": [
        {
          "clients": [
            {
              "match": "0.0.0.0/0"
            }
```

After configuring access control list and export policy, provide the valid values for the mandatory input parameters for both CIFS and NFS file shares:

> Storage Efficiency Policy is an optional parameter for creating file shares.

Category	HTTP verb	Path
storage-provider	POST	`/storage-provider/file-shares`

The JSON output displays a Job object key that you can use to verify the file share that you created.
. Verify the file share creation by using the Job object key returned in querying the job:

Category	HTTP verb	Path
management-server	GET	`/management-server/jobs/{key}`

At the end of the response, you see the key of the file share created.

```
    ],
    "job_results": [
        {
            "name": "fileshareKey",
            "value": "7d5a59b3-953a-11e8-8857-
00a098dcc959:type=volume,uuid=e581c23a-1037-11ea-ac5a-00a098dcc6b6"
        }
    ],
    "_links": {
        "self": {
            "href": "/api/management-server/jobs/06a6148bf9e862df:-
2611856e:16e8d47e722:-7f87"
        }
    }
}
```

1. Verify the creation of the file share by running the following API with the returned key:

Category	HTTP verb	Path
storage-provider	GET	`/storage-provider/file-shares/{key}`

Sample JSON output

You can see that the POST method of `/storage-provider/file-shares` internally invokes all the APIs required for each of the functions and creates the object. For example, it invokes the `/storage-provider/performance-service-levels/` API for assigning the Performance Service Level on the file share.

```
{
    "key": "7d5a59b3-953a-11e8-8857-
00a098dcc959:type=volume,uuid=e581c23a-1037-11ea-ac5a-00a098dcc6b6",
    "name": "FileShare_377",
    "cluster": {
        "uuid": "7d5a59b3-953a-11e8-8857-00a098dcc959",
        "key": "7d5a59b3-953a-11e8-8857-
00a098dcc959:type=cluster,uuid=7d5a59b3-953a-11e8-8857-00a098dcc959",
        "name": "AFFA300-206-68-70-72-74",
        "_links": {
```

```
            "self": {
                "href": "/api/datacenter/cluster/clusters/7d5a59b3-953a-
11e8-8857-00a098dcc959:type=cluster,uuid=7d5a59b3-953a-11e8-8857-
00a098dcc959"
            }
        }
    },
    "svm": {
        "uuid": "b106d7b1-51e9-11e9-8857-00a098dcc959",
        "key": "7d5a59b3-953a-11e8-8857-
00a098dcc959:type=vserver,uuid=b106d7b1-51e9-11e9-8857-00a098dcc959",
        "name": "RRT_ritu_vs1",
        "_links": {
            "self": {
                "href": "/api/datacenter/svm/svms/7d5a59b3-953a-11e8-
8857-00a098dcc959:type=vserver,uuid=b106d7b1-51e9-11e9-8857-
00a098dcc959"
            }
        }
    },
    "assigned_performance_service_level": {
        "key": "1251e51b-069f-11ea-980d-fa163e82bbf2",
        "name": "Value",
        "peak_iops": 75,
        "expected_iops": 75,
        "_links": {
            "self": {
                "href": "/api/storage-provider/performance-service-
levels/1251e51b-069f-11ea-980d-fa163e82bbf2"
            }
        }
    },
    "recommended_performance_service_level": {
        "key": null,
        "name": "Idle",
        "peak_iops": null,
        "expected_iops": null,
        "_links": {}
    },
    "space": {
        "size": 104857600
    },
    "assigned_storage_efficiency_policy": {
        "key": null,
        "name": "Unassigned",
        "_links": {}
```

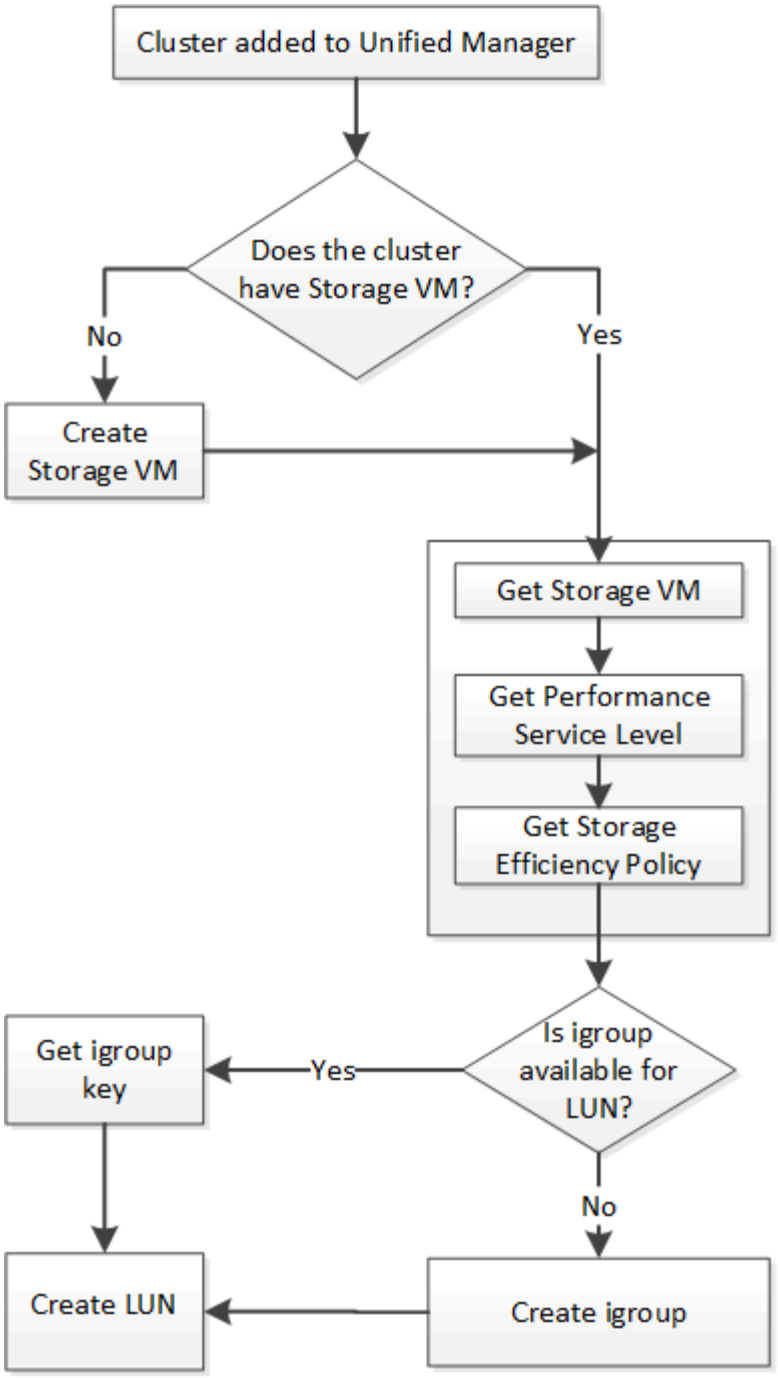
```
    },
    "access_control": {
        "acl": [
            {
                "user_or_group": "everyone",
                "permission": "read"
            }
        ],
        "export_policy": {
            "id": 1460288880641,
            "key": "7d5a59b3-953a-11e8-8857-
00a098dcc959:type=export_policy,uuid=1460288880641",
            "name": "default",
            "rules": [
                {
                    "anonymous_user": "65534",
                    "clients": [
                        {
                            "match": "0.0.0.0/0"
                        }
                    ],
                    "index": 1,
                    "protocols": [
                        "nfs3",
                        "nfs4"
                    ],
                    "ro_rule": [
                        "sys"
                    ],
                    "rw_rule": [
                        "sys"
                    ],
                    "superuser": [
                        "none"
                    ]
                },
                {
                    "anonymous_user": "65534",
                    "clients": [
                        {
                            "match": "0.0.0.0/0"
                        }
                    ],
                    "index": 2,
                    "protocols": [
                        "cifs"
```

```
                    ],
                    "ro_rule": [
                        "ntlm"
                    ],
                    "rw_rule": [
                        "ntlm"
                    ],
                    "superuser": [
                        "none"
                    ]
                }
            ],
            "_links": {
                "self": {
                    "href": "/api/datacenter/protocols/nfs/export-
policies/7d5a59b3-953a-11e8-8857-
00a098dcc959:type=export_policy,uuid=1460288880641"
                }
            }
        }
    },
    "_links": {
        "self": {
            "href": "/api/storage-provider/file-shares/7d5a59b3-953a-
11e8-8857-00a098dcc959:type=volume,uuid=e581c23a-1037-11ea-ac5a-
00a098dcc6b6"
        }
    }
}
```


Provision LUNs using Active IQ Unified Manager APIs

You can provision LUNs on your Storage Virtual Machines (SVMs) by using the provisioning APIs provided as a part of Active IQ Unified Manager. This provisioning workflow details the steps for retrieving the keys of the SVMs, Performance Service Levels, and Storage Efficiency Policies before creating the LUN.

The following diagram illustrates the steps in a LUN provisioning workflow.



This workflow assumes that the ONTAP clusters have been added to Unified Manager, and the cluster key has been obtained. The workflow also assumes that the SVMs have already been created on the clusters.

1. Obtain the SVM key for the SVM on which you want to create the LUN, as described in the *Verifying SVMs on clusters* workflow topic.
2. Obtain the key for the Performance Service Level by running the following API and retrieving the key from the response.

Category	HTTP verb	Path
storage-provider	GET	`/storage-provider/performance-service-levels`

You can retrieve the details of the system-defined Performance Service Levels by setting the `system_defined` input parameter to `true`. From the output, obtain the key of the Performance Service Level that you want to apply on the LUN.

3. Optionally, obtain the Storage Efficiency Policy key for the Storage Efficiency Policy that you want to apply on the LUN by running the following API and retrieving the key from the response.

Category	HTTP verb	Path
storage-provider	GET	`/storage-provider/storage-efficiency-policies`

4. Determine if initiator groups (igroups) have been created to grant access to the LUN target that you want to create.

Category	HTTP verb	Path
datacenter	GET	`/datacenter/protocols/san/igroups` `/datacenter/protocols/san/igroups/{key}`

You must enter the parameter value for indicating the SVM for which the igroup has authorized access. Additionally, if you want to query a particular igroup, enter the igroup name (key) as an input parameter.

5. In the output, if you can find the igroup that you want to grant access to, obtain the key. Otherwise create the igroup.

Category	HTTP verb	Path
datacenter	POST	`/datacenter/protocols/san/igroups`

You must enter the details of the igroup that you want to create, as the input parameters. This is a synchronous call and you can verify the igroup creation in the output. In case of an error, a message is displayed for you to troubleshoot and rerun the API.

6. Create the LUN.

Category	HTTP verb	Path
storage-provider	POST	`/storage-provider/luns`

For creating the LUN, ensure that you have added the retrieved values as mandatory input parameters.

Storage Efficiency Policy is an optional parameter for creating LUNs.

Sample cURL

You must enter all the details of the LUN that you want to create, as the input parameters.

The JSON output displays a Job object key that you can use to verify the LUN that you created.

7. Verify the LUN creation by using the Job object key returned in querying the Job:

Category	HTTP verb	Path
management-server	GET	`/management-server/jobs/{key}`

At the end of the response, you see the key of the LUN created.

8. Verify the creation of the LUN by running the following API with the returned key:

Category	HTTP verb	Path
storage-provider	GET	`/storage-provider/luns/{key}`

Sample JSON output

You can see that the POST method of `/storage-provider/luns` internally invokes all the APIs required for each of the functions and creates the object. For example, it invokes the `/storage-provider/performance-service-levels/` API for assigning the Performance Service Level on the LUN.

== Troubleshooting steps for failure in LUN creation or mapping

On completing this workflow, you might still see a failure in your LUN creation. Even if the LUN is created successfully, the LUN mapping with the igroup might fail due to an unavailability of a SAN LIF or access endpoint on the node on which you create the LUN. In case of a failure, you can see the following message:

```
The nodes <node_name> and <partner_node_name> have no LIFs configured with
the iSCSI or FCP protocol for Vserver <server_name>. Use the access-
endpoints API to create a LIF for the LUN.
```

Follow these troubleshooting steps to work around this failure.

1. Create an access endpoint supporting ISCSI/FCP protocol on the SVM on which you tried creating the LUN.

Category	HTTP verb	Path
storage-provider	POST	`/storage-provider/access-endpoints`

Sample cURL

You must enter the details of the access endpoint that you want to create, as the input parameters.

Ensure that in the input parameter you have added the address to indicate the home node of the LUN and the ha_address to indicate the partner node of the home node. When you run this operation, it creates access endpoints on both the home node and the partner node.

2. Query the job with the Job object key returned in the JSON output to verify that it has run successfully to add the access endpoints on the SVM and that the iSCSI/FCP services have been enabled on the SVM.

Category	HTTP verb	Path
management-server	GET	`/management-server/jobs/{key}`

Sample JSON output

At the end of the output, you can see the key of the access endpoints created. In the following output, the "name": "accessEndpointKey" value indicates the access endpoint created on the home node of the LUN, for which the key is 9c964258-14ef-11ea-95e2-00a098e32c28. The "name": "accessEndpointHAKey" value indicates the access endpoint created on the partner node of the home node, for which the key is 9d347006-14ef-11ea-8760-00a098e3215f.

3. Modify the LUN to update the igroup mapping. For more information about workflow modification, see “Modifying storage workloads”.

Category	HTTP verb	Path
storage-provider	PATCH	`/storage-provider/lun/{key}`

In the input, specify the igroup key with which you want to update the LUN mapping, along with the LUN key.

Sample cURL

The JSON output displays a Job object key that you can use to verify whether the mapping is successful.

4. Verify the LUN mapping by querying with the LUN key.

Category	HTTP verb	Path
storage-provider	GET	`/storage-provider/luns/{key}`

Sample JSON output

In the output you can see the LUN has been successfully mapped with the igroup (key d19ec2fa-fec7-11e8-b23d-00a098e32c28) with which it was initially provisioned.

Modify storage workloads

Modify storage workloads using APIs in Active IQ Unified Manager

Modifying storage workloads consists of updating LUNs or file shares with missing parameters, or changing the existing parameters.

This workflow takes the example of updating Performance Service Levels for LUNs and file shares.

> The workflow assumes that the LUN or file share has been provisioned with Performance Service Levels.

Modify file shares

While modifying a file share, you can update the following parameters:

- Capacity or size.
- Online or offline setting.
- Storage Efficiency Policy.
- Performance Service Level.
- Access control list (ACL) settings.
- Export policy settings. You can also delete export policy parameters and revert the default (empty) export policy rules on the file share.

> During a single API run, you can update only one parameter.

This procedure describes adding a Performance Service Level to a file share. You can use the same procedure for updating any other file share property.

1. Obtain the CIFS share or NFS file share key of the file share that you want to update. This API queries all the file shares on your data center. Skip this step if you already know the file share key.

Category	HTTP verb	Path
storage-provider	GET	`/storage-provider/file-shares`

2. View the details of the file share by running the following API with the file share key that you obtained.

Category	HTTP verb	Path
storage-provider	GET	`/storage-provider/file-shares/{key}`

   View the details of the file share in the output.

```
"assigned_performance_service_level": {
        "key": null,
        "name": "Unassigned",
        "peak_iops": null,
        "expected_iops": null,
        "_links": {}
      },
```

3. Obtain the key for the Performance Service Level that you want to assign on this file share. Currently no

policy is assigned to it.

Category	HTTP verb	Path
Performance Service Levels	GET	`/storage-provider/performance-service-levels`

You can retrieve the details of the system-defined Performance Service Levels by setting the `system_defined` input parameter to `true`. From the output, obtain the key of the Performance Service Level that you want to apply to the file share.

4. Apply the Performance Service Level on the file share.

Category	HTTP verb	Path
Storage Provider	PATCH	`/storage-provider/file-shares/{key}`

In the input, you must specify only the parameter that you want to update, along with the file share key. In this case, it is the key of the Performance Service Level.

Sample cURL

```
curl -X POST "https://<hostname>/api/storage-provider/file-shares"`-H
"accept: application/json" -H "Authorization: Basic
<Base64EncodedCredentials>" -d
"{
\"performance_service_level\": { \"key\": \"1251e51b-069f-11ea-980d-
fa163e82bbf2\" },
}"
```

The JSON output displays a Job object that you can use to verify the whether the access endpoints on the home and partner nodes have been created successfully.

5. Verify whether the Performance Service Level has been added to the file share by using the Job object key displayed in your output.

Category	HTTP verb	Path
Management Server	GET	`/management-server/jobs/{key}`

If you query by the ID of the Job object, you see whether the file share is updated successfully. In case of a failure, troubleshoot the failure and run the API again. On successful creation, query the file share to see the modified object:

Category	HTTP verb	Path
storage-provider	GET	`/storage-provider/file-shares/{key}`

View the details of the file share in the output.

```
"assigned_performance_service_level": {
        "key": "1251e51b-069f-11ea-980d-fa163e82bbf2",
        "name": "Value",
        "peak_iops": 75,
        "expected_iops": 75,
        "_links": {
            "self": {
                "href": "/api/storage-provider/performance-service-
levels/1251e51b-069f-11ea-980d-fa163e82bbf2"
            }
        }
```

Update LUNs

While updating a LUN, you can modify the following parameters:

- Capacity or size
- Online or offline setting
- Storage Efficiency Policy
- Performance Service Level
- LUN map

During a single API run, you can update only one parameter.

This procedure describes adding a Performance Service Level to a LUN. You can use the same procedure for updating any other LUN property.

1. Obtain the LUN key of the LUN that you want to update. This API returns details of all the LUNS in your data center. Skip this step if you already know the LUN key.

Category	HTTP verb	Path
Storage Provider	GET	`/storage-provider/luns`

2. View the details of the LUN by running the following API with the LUN key that you obtained.

Category	HTTP verb	Path
Storage Provider	GET	`/storage-provider/luns/{key}`

View the details of the LUN in the output. You can see that there is no Performance Service Level assigned to this LUN.

Sample JSON output

```
  "assigned_performance_service_level": {
        "key": null,
        "name": "Unassigned",
        "peak_iops": null,
        "expected_iops": null,
        "_links": {}
      },
```

3. Obtain the key for the Performance Service Level that you want to assign to the LUN.

Category	HTTP verb	Path
Performance Service Levels	GET	`/storage-provider/performance-service-levels`



You can retrieve the details of the system-defined Performance Service Levels by setting the `system_defined` input parameter to `true`. From the output, obtain the key of the Performance Service Level that you want to apply on the LUN.

4. Apply the Performance Service Level on the LUN.

Category	HTTP verb	Path
Storage Provider	PATCH	`/storage-provider/lun/{key}`

In the input, you must specify only the parameter that you want to update, along with the LUN key. In this case it is the key of the Performance Service Level.

Sample cURL

```
curl -X PATCH "https://<hostname>/api/storage-provider/luns/7d5a59b3-
953a-11e8-8857-00a098dcc959" -H "accept: application/json" -H "Content-
Type: application/json" H "Authorization: Basic
<Base64EncodedCredentials>" -d
"{ \"performance_service_level\": { \"key\": \"1251e51b-069f-11ea-980d-
fa163e82bbf2\" }"
```

The JSON output displays a Job object key that you can use to verify the LUN that you updated.

5. View the details of the LUN by running the following API with the LUN key that you obtained.

Category	HTTP verb	Path
Storage Provider	GET	`/storage-provider/luns/{key}`

View the details of the LUN in the output. You can see that the Performance Service Level is assigned to this LUN.

Sample JSON output

```
     "assigned_performance_service_level": {
        "key": "1251e51b-069f-11ea-980d-fa163e82bbf2",
        "name": "Value",
        "peak_iops": 75,
        "expected_iops": 75,
        "_links": {
            "self": {
                "href": "/api/storage-provider/performance-service-
levels/1251e51b-069f-11ea-980d-fa163e82bbf2"
            }
```

Modify an NFS file share to support CIFS using Active IQ Unified Manager APIs

You can modify an NFS file share to support CIFS protocol. During file share creation, it is possible to specify both access control list (ACL) parameters and export policy rules for the same file share. However, if you want to enable CIFS on the same volume where you created an NFS file share, you can update the ACL parameters on that file share to support CIFS.

Before you begin

1. An NFS file share must have been created with only the export policy details. For information, see *Managing file shares* and *Modifying storage workloads*.
2. You must have the file share key to run this operation. For information about viewing file share details and retrieving the file share key by using the Job ID, see *Provisioning CIFS and NFS file shares*.

This is applicable for an NFS file share that you created by adding only export policy rules and not ACL parameters. You modify the NFS file share to include the ACL parameters.

Steps

1. On the NFS file share, perform a `PATCH` operation with the ACL details for allowing CIFS access.

Category	HTTP verb	Path
storage-provider	PATCH	`/storage-provider/file-shares`

Sample cURL

Based on the access privileges you assign to the user group, as displayed in the following sample, an ACL is created and assigned to the file share.

```
{
  "access_control": {
```

```
    "acl": [
      {
        "permission": "read",
        "user_or_group": "everyone"
      }
    ],
    "active_directory_mapping": {
      "key": "3b648c1b-d965-03b7-20da-61b791a6263c"
    }
```

Sample JSON output

The operation returns the Job ID of the Job that runs the update.

2. Verify whether the parameters have been added correctly by querying the file share details for the same file share.

Category	HTTP verb	Path
storage-provider	GET	`/storage-provider/file-shares/{key}`

Sample JSON output

```
"access_control": {
        "acl": [
            {
                "user_or_group": "everyone",
                "permission": "read"
            }
        ],
        "export_policy": {
            "id": 1460288880641,
            "key": "7d5a59b3-953a-11e8-8857-
00a098dcc959:type=export_policy,uuid=1460288880641",
            "name": "default",
            "rules": [
                {
                    "anonymous_user": "65534",
                    "clients": [
                        {
                            "match": "0.0.0.0/0"
                        }
                    ],
                    "index": 1,
                    "protocols": [
                        "nfs3",
```

```
                        "nfs4"
                    ],
                    "ro_rule": [
                        "sys"
                    ],
                    "rw_rule": [
                        "sys"
                    ],
                    "superuser": [
                        "none"
                    ]
                },
                {
                    "anonymous_user": "65534",
                    "clients": [
                        {
                            "match": "0.0.0.0/0"
                        }
                    ],
                    "index": 2,
                    "protocols": [
                        "cifs"
                    ],
                    "ro_rule": [
                        "ntlm"
                    ],
                    "rw_rule": [
                        "ntlm"
                    ],
                    "superuser": [
                        "none"
                    ]
                }
            ],
            "_links": {
                "self": {
                    "href": "/api/datacenter/protocols/nfs/export-
policies/7d5a59b3-953a-11e8-8857-
00a098dcc959:type=export_policy,uuid=1460288880641"
                }
            }
        }
    },
    "_links": {
        "self": {
            "href": "/api/storage-provider/file-shares/7d5a59b3-953a-
```

```
11e8-8857-00a098dcc959:type=volume,uuid=e581c23a-1037-11ea-ac5a-
00a098dcc6b6"
        }
    }
```

You can see the ACL assigned along with the export policy to the same file share.

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

Information about third-party copyright and licenses used in this product.

Notice for Active IQ Unified Manager 9.18

Copyright information

Copyright © 2026 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.