



Plug-in information required to configure Snap Creator

Snap Creator Framework

NetApp
April 09, 2021

This PDF was generated from https://docs.netapp.com/us-en/snap-creator-framework/administration/reference_archive_log_plug_in.html on April 09, 2021. Always check docs.netapp.com for the latest.

Table of Contents

- Plug-in information required to configure Snap Creator 1
 - Archive log plug-in 4
 - Citrix XenServer plug-in 5
 - DB2 plug-in 6
 - IBM Domino plug-in 7
 - MaxDB plug-in 8
 - MySQL plug-in 10
 - Oracle plug-in 12
 - Red Hat KVM plug-in guidelines 14
 - SAP HANA plug-in 16
 - SnapManager for Microsoft Exchange plug-in 16
 - SnapManager for Microsoft SQL Server plug-in 17
 - Sybase ASE plug-in 18
 - VMware VIBE plug-in 21
- Using the plug-in framework to create custom plug-ins 26
- Configuring Snap Creator for multilevel application quiesce operations when using hypervisor plug-ins . . . 27

Plug-in information required to configure Snap Creator

Snap Creator supports the following built-in (or native) plug-ins: Oracle, DB2, MySQL, Sybase ASE, Domino, SnapManager for Microsoft SQL Server, SnapManager for Microsoft Exchange, MaxDB, VMware (vSphere and vCloud Director), Red Hat KVM, Citrix XenServer, and SAP HANA. Community plug-ins are not included in the package and must be downloaded separately.

The following table lists and describes the plug-in parameters and settings:

Parameters	Setting	Description
APP_NAME	oracle	db2
mysql	domino	vibe
smsql	sme	sybase
maxdb	kvm	xen
hana<plug-in>	Determines which application is being backed up. Snap Creator has built-in support for the listed applications. You can either use APP_NAME or configure APP_QUIESCE_CMDXX, APP_UNQUIESCE_CMDXX, and PRE_EXIT_CMDXX. If the application is not directly supported in Snap Creator, you can use a plug-in or run your own application quiesce or unquiesce commands or scripts. <PLUG-IN>: Copy the plug-in to the /path_to_scServer	scAgent/plug-ins directory and specify the plug-in in the APP_NAME parameter. Commands or Scripts: ---- APP_QUIESCE_CMD01=path_to_quiesceCMD APP_UNQUIESCE_CMD01=path_to_unquiesceCMD PRE_EXIT_CMD01=path_to_unquiesceCMD ----
APP_IGNORE_ERROR	(Y	N)
Determines whether Snap Creator should ignore errors from application plug-ins. This is useful when you want to back up multiple databases and do not want to stop the backup if the quiesce or unquiesce operations of one database fails.	APP_DEFINED_BACKUP	(Y

Parameters	Setting	Description
N)	The application plug-in is expected to perform the entire backup operation including quiescing, creating a Snapshot copy, and unquiescing. The built-in plug-ins do not support this kind of backup.	APP_DEFINED_RESTORE
(Y	N)	Enables application-based restore operations. In the event of a restore operation, Snap Creator sends a request to the application plug-in and the plug-in handles the request.
APP_DEFINED_MOUNT	(Y	N)
The built-in mount abilities of Snap Creator are ignored. Instead, the plug-in is responsible for all mount activities including volume or LUN clone creation. The built-in plug-ins do not support this type of mount.	APP_DEFINED_UMOUNT	(Y
N)	The built-in unmount abilities of Snap Creator are ignored. Instead, the plug-in is responsible for handling all unmount activities including volume or LUN clone deletion. The built-in plug-ins do not support this type of unmount operation.	APP_AUTO_DISCOVERY
(Y	N)	Enables application automatic discovery. Snap Creator sends a discover request to the application plug-in and the plug-in is then responsible for the discovery of the storage configuration. This can be done dynamically or made persistent using the APP_CONF_PERSISTENCE parameter, if the information is to be saved to the configuration file.
APP_CONF_PERSISTENCE	(Y	N)

Parameters	Setting	Description
Enables automatic discovery to be persistent, which means changes are dynamically updated in configuration file.	APP_DEFINED_CLONE	(Y
N)	The built-in cloning abilities of Snap Creator are ignored. Instead, the plug-in is responsible for handling all clone activities, including volume or LUN clone creation and deletion. The built-in plug-ins do not support this type of clone.	FS_NAME
plug-in	Determines which plug-in is being used for file system operations.	JAVA_HOME
Text	This setting points to the Java Virtual Machine (JVM) that should be used for executing .class and .jar files.	JVM_ARGS
Text	This setting controls the JVM settings when native Java .class or .jar files are executed. The default setting is -Xms32M -Xmx128M.	JAVA_CLASSPATH
Text	This setting defines the Java classpath. By default, plug-ins/native is configured and can be completed using this environment variable, which is appended to the default.	META_DATA_VOLUME
	Enables a Snapshot copy of the specified volume to be created after the unquiesce operation. This can be valuable for certain plug-ins for which the Snapshot copy of data must be created at different times. The parameter must not only specify the volume but the controller as well (for example, <code>controller1:volume1,volume2;controller2:volume3,volume4;controller3:volume5,volume6</code>).	PERL_HOME

Parameters	Setting	Description
Text	This setting points to the Perl interpreter that should be used for executing .pl files.	PERL_OPTS
Text	This setting controls the PERL interpreter settings when native Perl files are executed. Options for additional settings include directories (-I) that can be passed to the Perl interpreter.	PYTHON_HOME
Text	This setting points to the Python interpreter that should be used for executing .py files.	PYTHON_OPTS
Text	This setting controls the Python interpreter settings when native Python files are executed.	VALIDATE_VOLUMES

Archive log plug-in

The Archive log plug-in works with Snap Creator Archive logs and not with the logs of any application or database.

The following table lists the Archive log plug-in parameters, provides their settings, and describes them:

Parameter	Setting	Description
ARCHIVE_LOG_ENABLE	(Y	N
policy:age)	Enables archive log management (deletion of the old archive logs).	ARCHIVE_LOG_RETENTION
number_of_days	Specifies the number of days the archive logs are retained. This setting must be equal to or greater than NTAP_SNAPSHOT_RETENTIONS.	ARCHIVE_LOG_DIR
change_info_directory/logs	Specifies the path to the directory that contains the archive logs.	ARCHIVE_LOG_EXT
file_extension	Specifies the file extension of the archive logs. For example, if the archive logs are 10192091019.log, specify this parameter setting to LOG.	ARCHIVE_LOG_RECURSIVE_SE ARCH

Parameter	Setting	Description
(Y	N)	Enables the management of archive logs within subdirectories. If the archive logs are located under subdirectories, you should use this parameter.

Citrix XenServer plug-in

Snap Creator supports the backup and restore of Citrix XenServer virtual machines (VMs) by using the Citrix XenServer plug-in.



For the latest information about support or compatibility matrices, see the Interoperability Matrix.

The Citrix XenServer plug-in supports Windows and XenServer.

Consider the following when you use the Citrix XenServer plug-in:

- Active IQ Unified Manager server as a proxy is not supported.
- Mount, unmount, and backup operations using Open Systems SnapVault, and archive log management, are not supported.
- Volume restore operations are not supported; only application-defined restore operations are supported.
- Deleted VMs can be restored.
- Snap Creator Agent must be installed on the host where XenCenter is installed, and Snap Creator Server must not be installed on XenServer.
- The SC_AGENT_TIMEOUT value should be greater than the default value: 600 or higher.
- If the value of APP_DEFINED_RESTORE is Y, then the SnapVault restore operation using the GUI is not supported.
- If the pool master goes down in a server pool, then the Snap Creator configuration file should be modified manually with the new master server for further activity.
- XenServer tools must be installed on all the VMs.
- For Fibre Channel in a storage area network (SAN) environment, the plink.exe tool must be installed on a host where Snap Creator Agent is installed, and the plink.exe path must be added to the system environment variable.

For detailed information about how to add the plink.exe path to the system environment variable, refer to the *Snap Creator Framework Installation Guide*.

- VM pause and unpause operations are performed serially.

For multiple VMs, the duration of VM in the pause state during backup operation depends on the number of VMs.

- Automatic discovery of volumes is supported.

Supported Citrix XenServer configurations

The following Citrix XenServer configurations are supported:

- SAN
 - Supports guest virtual machines with one virtual disk image (VDI) per storage repository.
 - Supports data disks with one VDI per storage repository
- NAS
 - Supports guest VMs installed on NFS mounts.
 - Supports data disks on NFS mounts.

Parameters

The following table lists and describes the Citrix XenServer plug-in parameters:

Parameter	Setting	Description
XEN_VMS	host IP:VM#	Lists virtual machines of a particular host, separated by a slash (/). For example: 10.10.10.192:VM1/VM2/VM3
XEN_HOST_USERS	host IP:username/password	Lists Xen hosts and the corresponding user name and password.
XEN_BIN_PATH	For example: c:\Program Files\Citrix\XenCenter\xe.exe	Specifies the path of the XenServer executable (xe). The XenCenter console is required for importing and exporting the VM metadata.
XEN_METADATA_PATH	For example: c:\scmetadata	Specifies the path on the server to which you can download the virtual machine metadata.
XEN_RESTORE_VMS	For example: xenserver1:vm1,vm2;xenserver2:vm1,vm2	Contains the VMs that must be restored. This parameter is required only during a restore operation.

Related information

[Interoperability Matrix Tool: mysupport.netapp.com/matrix](https://mysupport.netapp.com/matrix)

DB2 plug-in

The DB2 plug-in uses the db2 command to communicate with the database.

The following table lists the DB2 plug-in parameters, provides the parameter settings, and describes the

parameters.

Parameter	Setting	Description
APP_NAME	db2	Provides the application name.
DB2_DATABASES	db_name:user_name	Lists the DB2 databases and the user name. Multiple databases and user names can be specified as a semicolon-separated list: for example, <code>db1:user1;db2:user2</code> .
DB2_CMD	path_to_db2cli_cmd	Specifies the path that is used to initialize the database connection so that further commands can be executed on the database. • UNIX-based environment: <code>db2_install_directory/sqlllib/bin/db2</code> For example: <code>/home/db2inst1/sqlllib/bin/db2</code> • Windows: <code>db2_install_directory\SQLLIB\BIN\db2cmd.exe</code> For example: <code>C:\Program Files\IBM\SQLLIB\BIN\db2cmd.exe</code> If a path is not specified, then <code>sqlllib/db2</code> is used as the path.

Note: The DB2 plug-in handles Write Anywhere File Layout (WAFL) operations by default. If you want to back up a consistency group backup with the DB2 plug-in, you should set the parameter to `NTAP_CONSISTENCY_GROUP_WAFL_SYNC` parameter to N. If you set this parameter to Y, additional and redundant synchronizing operations are performed.

For the latest information about support or to view compatibility matrices, see the Interoperability Matrix.

Related information

[Interoperability Matrix Tool: mysupport.netapp.com/matrix](https://mysupport.netapp.com/matrix)

IBM Domino plug-in

The IBM Domino plug-in for the Snap Creator Framework offers a complete backup and

recovery solution for Domino databases on NetApp storage. With the IBM Domino plug-in, you can back up databases efficiently and restore them as needed without taking database servers offline. The plug-in uses IBM-provided APIs to ensure application consistency.

With key NetApp data protection technologies tightly integrated in the Snap Creator Framework, you can use the IBM Domino plug-in to:

- Create application-consistent Snapshot copies on primary storage
- Replicate Snapshot copies to secondary storage for disaster recovery and archiving

Integrated technologies include Snapshot, SnapMirror, and SnapVault.

Related information

[Snap Creator Framework 4.1.2 IBM Domino Plug-in Operations Guide](#)

MaxDB plug-in

The MaxDB plug-in automates backup and restore operations on MaxDB databases.



For latest information about support or to view compatibility matrices, see the Interoperability Matrix.

The MaxDB plug-in provides the following features:

- A centralized framework to back up, restore, and clone MaxDB databases
- Integration with the database and provision of application consistency
- Utilization of Snapshot technology to create point-in-time copies of the database
- Utilization of SnapRestore to restore a previous Snapshot copy, and therefore an application-consistent database, in seconds, regardless of the capacity or number of files
- Utilization of FlexClone technology to create fast, space-efficient clones of databases based on Snapshot copy backups

The following table lists the MaxDB plug-in parameters, provides their settings, and describes them:

Parameter	Setting	Description
APP_NAME	maxdb	Specifies the application name.
XUSER_ENABLE	(Y N)	
Enables or disables the use of an xuser for MaxDB so that a password is not required for the database user.	HANDLE_LOGWRITER	(Y N)

Parameter	Setting	Description
N)	Executes suspend logwriter (N) or resume logwriter (Y) operations.	DBMCLICMD
path_to_dbmcli_cmd	<p>Specifies the path to the MaxDB dbmcli command.If not set, dbmcli on the search path is used.</p> <div style="border: 1px solid #ccc; padding: 5px; margin-top: 10px;">  <p>If in a Windows environment, the path needs to be contained within double-quotes ("...").</p> </div>	SQLCLICMD
path_to_sqlcli_cmd	Specifies the path for the MaxDB sqlcli command.If not set, sqlcli is used on the search path.	MAXDB_UPDATE_HIST_LOG
(Y	N)	Instructs the MaxDB backup program whether or not to update the MaxDB history log.
MAXDB_DATABASES	db_name:user_name/password	<p>Lists databases to be backed up with the user name and password.Multiple databases and user names can be specified using a comma-separated list: for example,</p> <p><code>db1:user1/password,db2:use r2/password.</code></p>
MAXDB_CHECK_SNAPSHOT_DIRECTORY	<p>Example:</p> <p><code>SID1:directory[,directory. ..];[SID2:directory[,direc tory...]</code></p>	<p>Checks that a Snap Creator Snapshot copy operation is successful and ensures that the Snapshot copy is created. This applies to NFS only. The directory must point to the location that contains the .snapshot directory.</p> <p>Multiple directories can be included in a comma-separated list. Multiple databases can be specified as a semicolon-separated list.</p> <p>In MaxDB 7.8 and later versions, the database backup request is marked Failed in the backup history.</p>

Parameter	Setting	Description
MAXDB_BACKUP_TEMPLATES	template_name Example: <code>na_snap</code>	Specifies a backup template for each database. The template must already exist and be an external type of backup template. To enable Snapshot copy integration for MaxDB 7.8 and later, you must have MaxDB background server functionality and already configured MaxDB backup template.
MAXDB_BG_SERVER_PREFIX	bg_server_prefix Example: <code>na_bg</code>	Specifies the prefix for the background server name. If the MAXDB_BACKUP_TEMPLATES parameter is set, you must also set the MAXDB_BG_SERVER_PREFIX parameter. If you do not set the prefix, the default value <code>na_bg_DATABASE</code> is used.

Related information

Interoperability Matrix Tool: mysupport.netapp.com/matrix

MySQL plug-in

The MySQL plug-in uses Net-MySQL to communicate with the database.


The MySQL plug-in does not support backup and restore operations for multiple databases. The Interoperability Matrix contains the latest information about support and compatibility.

For Snap Creator configurations, the MySQL database user must have the LOCK TABLES privilege granted, along with other privileges, such as SUPER and RELOAD.

The following table lists the MySQL plug-in parameters, provides their settings, and describes them:

Parameter	Setting	Description
APP_NAME	mysql	Application name

Parameter	Setting	Description
MYSQL_DATABASES	db_name:user_name/password	List of MySQL databases, the user name, and the password. You can specify multiple databases with user names and passwords as a semicolon-separated list, for example, <code>db1:user1/pwd1;db2:user2/pwd2.</code>
HOST	host_name	Name of the host where the databases are located. Note: VALIDATE_VOLUMES functions properly only if HOST=localhost. If HOST=IP_address, then VALIDATE_VOLUMES will not discover the MySQL database.
PORTS	db_name:port_number	List of databases and the ports they are listening on, for example, <code>db1:3307;db2:3308.</code>
MASTER_SLAVE	(Y	N)

Parameter	Setting	Description
<p>Specifies the backup database environment: If set to Y, backs up the master database.</p> <p>If set to N, either backs up the slave database or the Master-Slave configuration is not used.</p> <div style="border: 1px solid #ccc; padding: 5px; margin-top: 10px;">  <p>Snap Creator Framework backup deletes the required MySQL binary log files on the MySQL server with the Master_Slave option enabled. The MySQL (master) database backup using Snap Creator Framework removes all binary log files and leaves only a single empty binary log file with reset numbering (.000001). Because of this issue, the slave server fails to start up after the backup operation.</p> </div>	PURGE_BINARY_LOG	(Y

Related information

[Interoperability Matrix Tool: mysupport.netapp.com/matrix](https://mysupport.netapp.com/matrix)

Oracle plug-in

The Oracle plug-in uses SQL*Plus to communicate with the database and quiesce & unquiesce oracle database for backup.

The Snap Creator Oracle plug-in supports Oracle Automatic Storage Management (offline or online backups) and online backup of a Real Application Clusters (RAC) database configuration. The Interoperability Matrix

contains the latest information about support and compatibility.

To use C Shell (CSH) with the Oracle plug-in on UNIX or Linux platforms, the following conditions should be met:

- Snap Creator Agent must be started as the Oracle user, instead of the root user.
- The Oracle user must have the proper environmental variables (ORACLE_HOME and ORACLE_SID) set for the plug-in driven SQL*Plus commands to work.

This can be accomplished using a shell startup file, such as ~/.cshrc.

The following table lists the Oracle plug-in parameters, provides their settings, and describes them:

Parameter	Setting	Description
APP_NAME	Oracle	Application name
ORACLE_DATABASES	db_name:user_name	List of Oracle databases and user names Multiple databases and user names can be specified as a semicolon-separated list, for example, <code>db1:user1;db2:user2.</code>
SQLPLUS_CMD		Path to the sqlplus command
CNTL_FILE_BACKUP_DIR		Path to the directory where the user should store the backup control files The Oracle user must have appropriate permissions to this directory.
ORA_TEMP		Path to the directory to store the temporary file, for example, /tmp. The Oracle user must have appropriate permissions to this directory.
ARCHIVE_LOG_ONLY	(Y	N
policy:<Y	N>)	Informs the Oracle plug-in to perform only a switch log operation This setting is useful if you are handling archive log backups separately from data backups.
ORACLE_HOME		Path to the Oracle home directory

Parameter	Setting	Description
ORACLE_HOME_SID		Path to the Oracle home directory for a given system identifier (SID)When backing up multiple databases, it might be important to specify more than one Oracle home directory.
ORACLE_EXPORT_PARAMETERS	(Y N)	
The ORACLE_HOME and ORACLE_SID environment parameters are exported by using the export command. This is applicable only in UNIX or a Linux-based environment.	ORACLE_BACKUPMODE	
Option to configure offline or online backups by using the Snap Creator policy. The default option is online. ORACLE_BACKUPMODE=hourly: online,daily:offline	ORACLE_SHUTDOWNABORT	(Y

Related information

Interoperability Matrix Tool: mysupport.netapp.com/matrix

Red Hat KVM plug-in guidelines

Kernel-based Virtual Machine (KVM) is a virtualization solution for the Linux kernel. Snap Creator uses the KVM plug-in to back up and restore the guest virtual machines.



For the latest information about support or for compatibility matrices, see the Interoperability Matrix.

The KVM plug-in supports guest operating systems such as Linux, Windows, and Solaris.

The plug-in internally uses virsh commands.

You must consider the following when you use the KVM plug-in:

- Active IQ Unified Manager server as a proxy is not supported.
- Mount, unmount, and backup operations using Open Systems SnapVault, and archive log management, are not supported.
- In a storage area network (SAN) environment, the Linux Host Utilities (LHU) kit is required to collect information about LUNs and volumes from the storage controller.

The LHU kit should be installed on a KVM hypervisor, which is the same location as the Snap Creator

Agent.



If the LHM kit is not installed and the configuration is a mix of network attached storage and storage area network environments, then the backup and restore operations work only on a Network File System (NFS).

- The KVM plug-in supports only the Linux version of the Snap Creator 4.0 and 4.1 Agent.
- Volume restore is not supported; only application-defined restore operations are supported.
- The deleted virtual machines (VMs) cannot be restored.
- The storage controller IPs and host should be either in `/etc/hosts` on the KVM host or in a Domain Name System (DNS).
- Only one KVM host per configuration is supported.
- If a virtual machine is created by using an ISO repository, then to perform any action, you should disconnect this repository from the virtual machine through the Virt Manager console in CD-ROM options.
- The `SC_AGENT_TIMEOUT` value should be more than the default value: 600 or higher.
- The volumes are automatically discovered using automatic detection.

You cannot see the discovered destination volumes if the destination is not in a SnapMirror relationship. You should use `dpstatus` to check the status of the SnapMirror relationship. If a SnapMirror relationship does not exist, you must first create the SnapMirror relationship.

- If the value of `APP_DEFINED_RESTORE` is Y, then the SnapVault restore operation using the GUI is not supported.
- When creating a SnapMirror and SnapVault configuration by using the GUI, the volume details must be entered manually because the SnapMirror and SnapVault relationship is not detected automatically.
- Data disks mapped to the VMs are not backed up.
- VM suspend and resume operations are performed serially.

For multiple VMs, the duration of VM in suspend state during backup operation depends on number of VMs.

Supported KVM configurations

- SAN: Supports guest virtual machines installed on a raw multipath device (LUN with multiple paths).
- NAS: Supports guest virtual machines installed on NFS volumes.



Configurations with multiple virtual machines installed on a single multipath device are not supported.

Guest virtual machines installed on LVM or on an image file in the SAN environment are not supported.

The following table describes the KVM plug-in parameters:

Parameter	Setting	Description
KVM_RESTORE_VM_LIST	Example: VM1, VM2	Specifies the list of VMs to be restored. This parameter is required only during restore.
KVM_VM_MAPPING	Example: VM1:s_c1:/vol/vol1/lun1;VM 2:s_c2:/vol/vol2/lun2;	(Required) Specifies the mapping between the VM and its associated storage controller, LUN, or file path. This parameter is updated dynamically during the discovery process.
KVM_VM_CONFIG_DIR	Default: /etc/libvirt/qemu	(Optional) Specifies the path to the directory where all the XML configuration files of the VM are stored.
KVM_CMD_RETRY_COUNT	Default: 3	(Optional) Specifies the number of times you rerun the command when running it fails in the KVM plug-in.

Related information

[Interoperability Matrix Tool: mysupport.netapp.com/matrix](https://mysupport.netapp.com/matrix)

SAP HANA plug-in

The SAP HANA plug-in enables you to create backups and perform point-in-time recovery of SAP HANA databases based on storage Snapshot copies.

This plug-in uses the SAP HANA hdbsql client to execute SQL commands to provide database consistency and to manage the SAP HANA backup catalog. The plug-in is supported for both SAP Certified Hardware Appliances and Tailored Datacenter Integration (TDI) programs.

The plug-in is installed as part of the Snap Creator Agent on a host that has access to the SAP HANA database nodes.

Related information

[Snap Creator Framework 4.3.3 SAP HANA Plug-in Operations Guide](#)

SnapManager for Microsoft Exchange plug-in

The SnapManager for Microsoft Exchange plug-in is used to centralize backup operations for Microsoft Exchange Server through Snap Creator. Using this plug-in, you can configure tens or hundreds of SnapManager for Exchange servers through Snap Creator, allowing you to view all of your SnapManager for Exchange backup jobs and statuses from a single interface.

Unlike the other plug-ins, the SnapManager plug-ins (SnapManager for Microsoft SQL Server and SnapManager for Microsoft Exchange) use PowerShell to communicate with an existing installation of SnapManager. The SnapManager plug-ins require that the SnapManager products are already installed and operating. The SnapManager plug-ins use the new-backup Powershell cmdlet to create a backup through SnapManager. All cloning and restore actions should continue to be driven through the SnapManager interface.



For the latest information about support or to view compatibility matrices, see the Interoperability Matrix.

The SnapManager for Microsoft Exchange plug-in requires Snap Creator Agent to be installed on the same host as SnapManager for Microsoft Exchange. You should set the SC_AGENT_TIMEOUT value to 900 or higher.

The following table provides SnapManager for Microsoft Exchange plug-in parameters, provides their settings, and describes them:

Parameter	Setting	Description
SME_PS_CONF	Example: "C:\Program Files\NetApp\SnapManager for Exchange\smeShell.psc1"	Specifies the path to the PowerShell configuration file for SnapManager for Microsoft Exchange.
SME_BACKUP_OPTIONS	Example: <code>Server 'EX2K10-DAG01' -GenericNaming -ManagementGroup 'Standard' -NoTruncateLogs \$False -RetainBackups 8 -StorageGroup 'dag01_db01' -BackupCopyRemoteCCRNNode \$False</code>	Specifies the SnapManager for Microsoft Exchange backup options. Snap Creator uses a PowerShell cmdlet for a new backup.
SME_SERVER_NAME	Example: <code>EX2K10-DAG01</code>	Specifies the SnapManager for Microsoft Exchange server name.
SME_32bit	(Y	N)

Related information

[Interoperability Matrix Tool: mysupport.netapp.com/matrix](https://mysupport.netapp.com/matrix)

SnapManager for Microsoft SQL Server plug-in

The SnapManager for Microsoft SQL Server plug-in is used to centralize backup operations for Microsoft SQL Server through Snap Creator. Using this plug-in, you can configure tens or hundreds of SnapManager for Microsoft SQL servers through Snap Creator, allowing you to view all of your SnapManager for Microsoft SQL backup jobs and statuses from a single interface.

Unlike the other plug-ins, the SnapManager plug-ins (SnapManager for Microsoft SQL Server and SnapManager for Microsoft Exchange) use PowerShell to communicate with an existing installation of SnapManager. The SnapManager plug-ins require that the SnapManager products are already installed and operating. The SnapManager plug-ins use the new-backup Powershell cmdlet to create a backup through SnapManager. All cloning and restore actions should continue to be driven through the SnapManager interface.



For the latest information about support or to view compatibility matrices, see the Interoperability Matrix.

When using the SnapManager for Microsoft SQL Server plug-in, you must be aware of the following considerations:

- Snap Creator Agent must be installed on the same host as SnapManager for Microsoft SQL Server. You should set the SC_AGENT_TIMEOUT value to 900 or higher.
- The Powershell should be installed in the Windows Powershell default installation location (for example, C:\WINDOWS\system32\WindowsPowerShell\v1.0).

The following table provides the SnapManager for Microsoft SQL Server plug-in parameters, provides their settings, and describes them:

Parameter	Setting	Description
SMSQL_PS_CONF	Example: "C:\Program Files\NetApp\SnapManager for SQL Server\smsqlShell.psc1"	Specifies the path to the PowerShell configuration file for the SnapManager for Microsoft SQL Server.
SMSQL_BACKUP_OPTIONS	Example: -svr 'SQL' -d 'SQL\SHAREPOINT', '1', 'WSS_Content' -RetainBackups 7 -lb -bksif -RetainSnapofSnapInfo 8 -trlog -gen -mgmt standard	Specifies the SnapManager for Microsoft SQL Server backup options. Snap Creator uses a PowerShell cmdlet for new backup.
SMSQL_SERVER_NAME	Example: SQL\SHAREPOINT	Specifies the SnapManager for Microsoft SQL Server name.
SMSQL_32bit	(Y	N)

Related information

[Interoperability Matrix Tool: mysupport.netapp.com/matrix](https://mysupport.netapp.com/matrix)

Sybase ASE plug-in

The Sybase ASE plug-in uses the isql command to interact with the Sybase database.



For latest information about support or to view compatibility matrices, see the Interoperability Matrix.

The following table lists the Sybase plug-in parameters, provides their settings, and describes them:

Parameter	Setting	Description
SYBASE_USER	user_name	Specifies the operating system user who can run the isql command. This parameter is required for UNIX. This parameter is required if the user running the Snap Creator Agentstart and stop commands (usually the root user) and the user running the isql command are different.
SYBASE_SERVER	data_server_name	Specifies the Sybase data server name (-S option on isql command).For example: <code>p_test</code>
SYBASE_DATABASES	db_name:user_name/password	Lists the databases within the instance to back up. The master database is added; for example: <code>DBAtest2:sa/53616c7404351e</code> .If a database named +ALL is used, then database automatic discovery is used, and the sybsyntax, sybssystemdb, sybssystemprocs, and tempdb databases are excluded. For example: <code>+ALL:sa/53616c71a6351e</code> Encrypted passwords are supported if the NTAP_PWD_PROTECTION parameter is set.
SYBASE_DATABASES_EXCLUDE	db_name	Allows databases to be excluded if the +ALL construct is used. You can specify multiple databases by using a semicolon-separated list.For example, <code>pubs2;test_db1</code>
SYBASE_TRAN_DUMP	db_name:directory_path	Enables you to perform a Sybase transaction dump after creating a Snapshot copy.For example: <code>pubs2:/sybasedumps/pubs2</code> You must specify each database that requires a transaction dump.

Parameter	Setting	Description
SYBASE_TRAN_DUMP_FORMAT	%S_%D_%T.cmn	Enables you to specify the dump naming convention. The following keys can be specified: <ul style="list-style-type: none"> • %S = instance name from SYBASE_SERVER • %D = database from SYBASE_DATABASES • %T = unique timestamp Here is an example: %S_%D_%T.log
SYBASE_TRAN_DUMP_COMPRESS	(Y	N)
Enables or disables native Sybase transaction dump compression.	SYBASE_ISQL_CMD	Example: /opt/sybase/OCS-15__0/bin/isql
Defines the path to the isql command.	SYBASE	Example: /sybase
Specifies the location of the Sybase installation.	SYBASE_LOGDIR	Example: /usr/local/ntap/scServer/logs
Defines the directory where Snap Creator logs are placed.	SYBASE_MANIFEST	Example: DBAtest2:/t_inf_nzl_devs/
Specifies the databases for which the manifest file should be created, along with the location where the manifest file should be placed.	SYBASE_MANIFEST_FORMAT	%S__%D_%T.manifest Example: %S_%D_%T.manifest
Enables you to specify the manifest file naming convention. The following keys can be specified: <ul style="list-style-type: none"> • %S = Instance name from SYBASE_SERVER • %D = database from SYBASE_DATABASES • %T = unique timestamp, which is the same as that used for Snapshot copy naming 	SYBASE_MANIFEST_DELETE	(Y

Parameter	Setting	Description
N)	Allows the manifest to be deleted after the Snapshot copy has been created. The manifest file should be captured in the Snapshot copy so that it is always available with the backup.	SYBASE_EXCLUDE_TEMPDB
(Y	N)	Enables automatic exclusion of user-created temporary databases.

Related information

[Interoperability Matrix Tool: mysupport.netapp.com/matrix](https://mysupport.netapp.com/matrix)

VMware VIBE plug-in

Snap Creator supports the backup of VMware virtual machines and vApps through the VMware VIBE plug-in. The VMware plug-in is an integrated plug-in for both virtual machines with vSphere and vApps with vCloud Director.

You must consider the following when you use the VMware VIBE plug-in:

- The VMware plug-in is supported only on Windows and Linux.

If you are using a non-Windows or non-Linux Snap Creator Server, you need a Snap Creator Windows or Linux agent to run the VMware plug-in.

- Unified Manager server as a proxy is not supported.
- Mount, unmount, and backup operations using Open Systems SnapVault and archive log management are not supported.
- VMware high availability (HA) with the VMware plug-in is not tested and is not supported.
- VMware vCenter Linked Mode with the VMware plug-in is not tested and is not supported.
- The VMware plug-in does not support raw device mapping (RDM).
- The volumes are discovered using automatic detection.

You cannot view a discovered destination volume if it is not in a SnapMirror relationship. You can use the `dpstatus` command to check the SnapMirror relationship. If a SnapMirror relationship does not exist, you must first create the SnapMirror relationship.


- Before you perform restore operations, you must delete all of the VMware snapshot copies.
- After the restore operations are complete, you must run a Snap Creator backup of the restored virtual machines and vApps so that the new environment is cleaned up and all VMware snapshot copies are removed.

If the VMware plug-in cannot clean up VMware snapshot copies and displays an error, you must remove the VMware snapshot copies manually. The VMware plug-in does not guarantee 100% VMware snapshot copy removal. This is a known VMware issue.

- The VMware plug-in supports only 32-bit Snap Creator with a 32-bit Linux system and 64-bit Snap Creator with a 64-bit Linux system.
- The deleted virtual machines cannot be restored.
- The volume restore operation is not supported; only application-defined restore operations are supported.
- The value of the SC_AGENT_TIMEOUT parameter should be set to 1800 or higher.
- The default value of the VIBE_VMWARE_snapshot parameter (VMware snapshot option) is N.
- If the value of APP_DEFINED_RESTORE is Y, then the SnapVault restore operation using the graphical user interface (GUI) is not supported.
- While creating a SnapMirror and SnapVault configuration by using the GUI, you must manually enter the SnapMirror and SnapVault parameters because the SnapMirror and SnapVault relationship is not detected automatically.
- The VMware plug-in discovers the ISO-mounted path as a datastore.

The following table lists the VMware VIBE plug-in parameters, provides the parameter settings, and describes the parameters.

Parameter	Setting	Description
VIBE_DYNAMIC_VOLUMES_UPDATE	Y or N Default: not set	If this parameter is set to N, dynamic volume update is not performed, which means you have to set the VOLUMES, SNAPVAULT_VOLUMES, SNAPSHOTMIRROR_VOLUMES, and NTAP_DFM_DATA_SET parameters manually.
VIBE_NOPING	Default: N	Specifies that Internet Control Message Protocol (ICMP) is not used to ping VMware plug-in or the storage controllers.
VIBE_VCLOUD_IPADDR	N/A	Specifies the IP address or the host name of the vCloud Director that is used for logging in to (vCloud only).

Parameter	Setting	Description
VIBE_VCLOUD_USER	N/A	<p>Specifies the user name to be used for logging in to the vCloud Director (vCloud only). You must set @org or @system (top-level vCloud database).</p> <div style="border: 1px solid #ccc; padding: 10px; margin: 10px 0;">  <p>The vCloud Director system administrator user name must be used to perform the backup and restore operations. These operations fail if the organization administrator credentials or any other user credentials are used.</p> </div> <p>Example: <code>administrator@system</code></p>
VIBE_VCLOUD_PASSWD	N/A	<p>Specifies the password that is associated with the specified VIBE_VCLOUD_USER (vCloud only).</p>
VIBE_VCENTER_USER	N/A	<p>Specifies the user name to be used for logging in to vCenter.</p>
VIBE_VCENTER_PASSWD	N/A	<p>Specifies the password that is associated with the specified VIBE_VCENTER_USER.</p>
VIBE_VCLOUD_NAMES	N/A	<p>Lists the organization, virtual data center, and vApp object names that should be backed up (vCloud only). Example: <code>ORG:VDC1,VDC2:VAPP1,VAPP2; ORG2:VDC3:;ORG3::VAPP6</code></p>

Parameter	Setting	Description
VIBE_VSPHERE_NAMES	N/A	Lists the datastores and virtual machines that should be backed up per vCenter (vSphere only). Example: <code>VCENTER1:DS1:VM1;VCENTER2;DS2,DS3;;VCENTER3::VM4</code>
VIBE_TRIM_VSPHERE_NAMES	N/A	Lists the virtual machines that should be removed from backup per vCenter (vSphere only). Example: <code>VCENTER1:VM99;VCENTER2:VM5,VM12</code>
VIBE_RESTORE_INTERVAL	Default: 30 seconds	Specifies the time between each restore check.
VIBE_RESTORE_TIME	Default: 3600 seconds	Specifies the total time to wait for a complete restore operation to finish.
VIBE_VMWARE_SNAPSHOT	Default: N	Creates a VMware snapshot copies during backup.
VIBE_IGNORE_EXPORTFS=Y or N	Default: N	You must manually add this parameter to the Snap Creator VIBE configuration file. When the value is set to Y, Data ONTAP operating in 7-Mode configurations ignores any exportfs values on the controller. Instead, Data ONTAP maps the volume export path as /vol/datastore_name, where a datastore name is specified for backup. Older environments using vFiler units might use this methodology because the exportfs information of individual datastores is not available from a vFiler unit. Instead, a configuration needs to map the path based on queries to vfiler0.

Related information

[Interoperability Matrix Tool: mysupport.netapp.com/matrix](https://mysupport.netapp.com/matrix)

Requirements for vCloud vApp backup and restore operations using the VMware plug-in

Snap Creator supports the backup of vCloud vApp through the VMware plug-in. vApp and virtual machine backup copies are made by the VMware plug-in through the vCloud Director API and vSphere API, which are invoked on the VMware vCloud Director and VMware vCenter server, respectively.

For vApp backup and restore operations to be successful, you must provide the following details in the configuration file:

- vCloud IP and credentials
- vCloud organizations, virtual data centers (vDCs), and vApp names



If more than one vCenter is attached to vCloud, then the password for the all vCenter servers should be same.

You must consider the following when performing the vCloud backup and restore operations:

- The backup and restore processes for both VMware and vCloud are very similar except for the discovery process, in which vCloud backups require additional discovery of the vCloud Director metadata using representational state transfer (REST) APIs.
- You should provide details of the vCloud with the organizations, vDCs, and vApps to be backed up.
- If a vDC is listed, all the vApps in the vDC are backed up.
- vCloud module discovers virtual machines associated with any vApp that must be backed up and puts them on a backup list.
- If a vApp selected for backup is contained within an organization or a vDC that is also selected for backup, the vApp is backed up only once.



For Virtual Machine File System (VMFS) restore operations using the VMware plug-in, there must be enough space in the volume to create a LUN clone that is equal to the size of the LUN.

Virtual machine backup and restore by using the VMware plug-in

Snap Creator supports the backup of VMware virtual machines through the VMware plug-in. Virtual machine backups are taken through the vSphere API on the VMware vCenter server.

For virtual machine backup, you must provide the following details in the configuration file:

- vCenter IP or host name and credentials
- vSphere virtual machines and datastore names



Snap Creator discovers vCenter only if vCenter is configured on the default port (443).

For the restore operation, you should provide the backup parameters and the Snapshot copy name.

Consider the following when performing the VMware backup and restore processes:

- If a virtual machine is listed and is not available, the plug-in displays an error message. It will not be able to restore a lost virtual machine even if it is backed up.
- If a datastore is listed, all the virtual machines in the datastore are backed up.
- Only the listed virtual machines or virtual machines located in the datastores specified are backed up.
- If a virtual machine selected for backup is contained within a datastore that is also selected for backup, it will be backed up only once.
- The VMware plug-in supports both Network File System (NFS) and VMware Virtual Machine File System (VMFS) datastores.
 - Virtual machine restores on an NFS datastore use Single File SnapRestore (SFSR) on the storage system, which avoids host copies.
 - To restore a virtual machine on a VMFS datastore, perform the following steps:
 - i. Use FlexClone or LUN clone of the LUN contained in a specific restore Snapshot copy.
 - ii. Map it to the cluster.
 - iii. Use vCenter API calls to copy contents from the Snapshot copy of the VMFS datastore to the original VMFS datastore.

Using the plug-in framework to create custom plug-ins

The plug-in framework enables you to create your own plug-ins for Snap Creator or reuse the existing plug-ins. The plug-in provides Snap Creator with the steps for handling backup and recovery of a given application.

The plug-in framework defines the following methods:

- quiesce - Method for handling quiesce for a given application plug-in
- unquiesce - Method for handling unquiesce for a given application plug-in
- discover - Method for handling discovery of storage objects for a given application plug-in
- scdump - Method for handling collection of support information, database, operating system, and SnapDrive
- restore - Method for handling restore for a given application plug-in
- restore_pre - Method for handling prerestore operations for a given application plug-in (can use built-in restore CLI of Snap Creator)
- restore_post - Method for handling post-restore operations for a given application plug-in (can use built-in restore CLI of Snap Creator)
- clone_pre - Method for handling preclone operations for a given application plug-in
- clone_post - Method for handling post-clone operations for a given application plug-in
- describe - Method for describing what a plug-in implements. This is optional for Perl plug-ins but required for native plug-ins under plug-ins/native.
- clone_all - Method for handling cloning for a given application plug-in (cannot use built-in cloning interface)
- clone_cleanup - Method for handling cleanup if a clone operation fails
- restore_cleanup - Method for handling cleanup if a restore operation fails



The plug-in framework supports Perl, PowerShell, Java, UNIX Shell, and Python for programming. NOTE: The plug-in framework enables you to implement objects and functions that exist within the Snap Creator.

+ For example, error handling is available, which means the plug-in can use the standard implementation Snap Creator uses. This reduces the work required to create a plug-in.

- Perl plug-ins are installed under /plug-ins/PLUG-IN-name/PLUG-IN.pm.
- Native plug-ins are installed under /plug-ins/native/plug-in.sh,plug-in.java.plug-in.bat, and so on.
- The plug-in must be installed where it is supposed to run. This can be Snap Creator Server or Snap Creator Agent depending on the set up of Snap Creator.

For more information about the plug-in framework, plug-ins, and a plug-in user guides, see the [Snap Creator Framework Discussions Community forum](#).

Configuring Snap Creator for multilevel application quiesce operations when using hypervisor plug-ins

When you are using the hypervisor (also known as “virtualization”) plug-ins (VMware (VMware vCloud Director and VMware vSphere), KVM, and Citrix XenServer) and want to perform a multilevel application quiesce and backup operation, you need to configure Snap Creator for this type of setup.

This configuration allows you to quiesce an application that resides on a virtual machine, quiesce the virtual machine, and then make a Snapshot copy.

During this process, you will create a hypervisor and application configuration by specifying a parent configuration file with one or more child configuration files. This parent configuration file contains the typical configuration file information such as retention policy, SVM details, and information for the hypervisor plug-in. Each child configuration file contains the details necessary to perform Snap Creator Quiesce and Unquiesce actions using the specific application plug-in.

1. Create a new profile.
2. Create a new configuration.
 - a. On the Configuration page, enter a name for the configuration file.

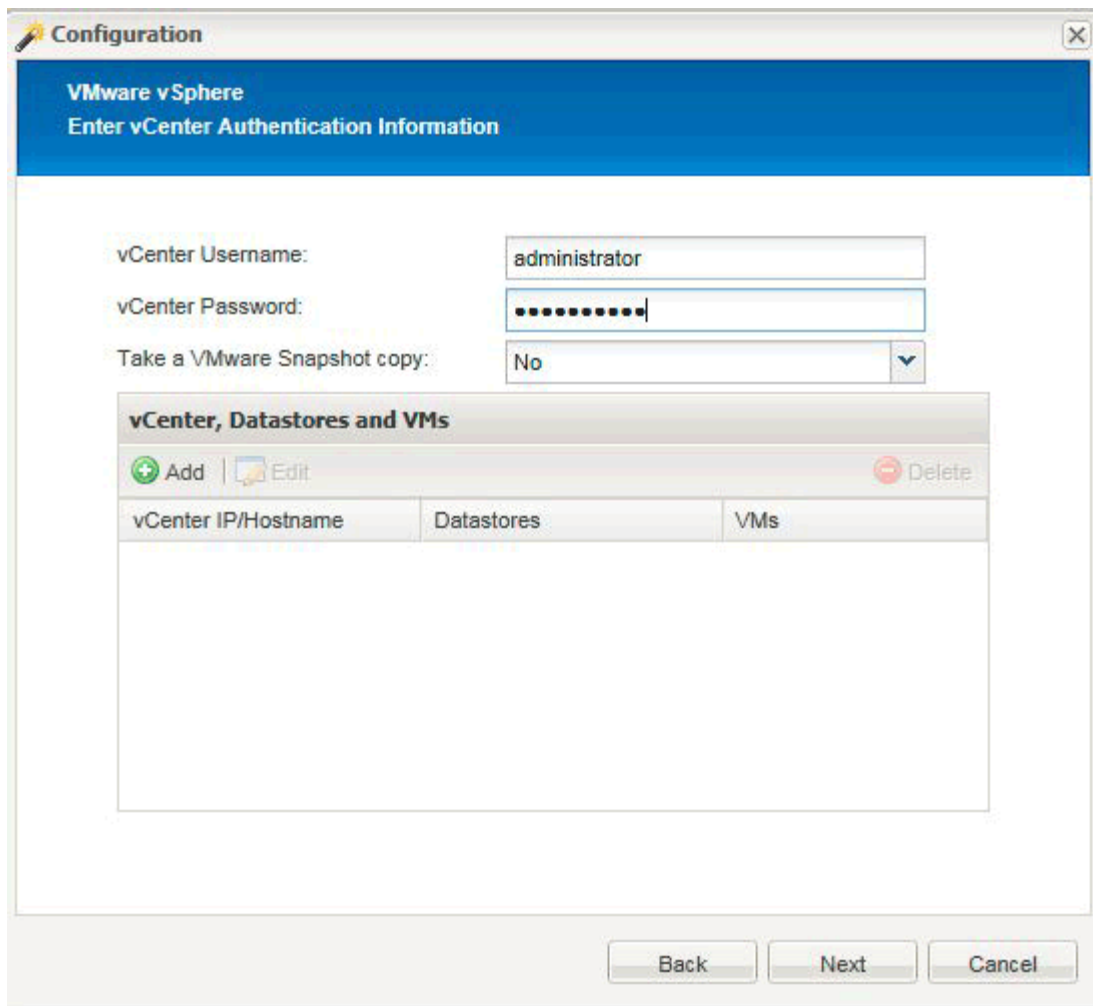


By default, password encryption is enabled to prevent passwords from being displayed in clear text in the configuration file.

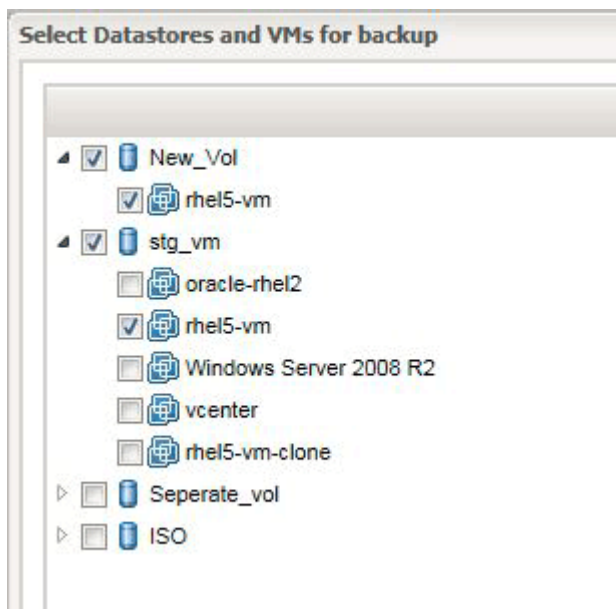
- b. On the Plug-In Type page, select **Virtualization plug-in**.
- c. On the Virtualization Plug-In page, select the plug-in to configure.
- d. On the plug-in parameters page, provide the configuration details associated with the selected plug-in option.

In the following example, VMware vSphere is the selected Virtualization plug-in. The wizard screens that display depend on your selection.

- i. Provide the appropriate information and click **Add**.



- ii. On the New vCenter page, provide the vCenter IP and Hostname, and click **Add**.
- iii. Select the applicable datastores and virtual machines for backup.



- iv. Verify the details you entered are correct.
- v. On the Agent Configuration page, provide the VMware agent details, which are the details of the

system where you have installed the agent.



The Port is the port on which the agent is listening.

Click **Test agent connection** to make sure that the agent is running.

- vi. On the Hypervisor + App Backup page, select **Yes** because both the hypervisor and application-consistent backups are required.
- vii. On the Hypervisor + App configuration page, provide the parent configuration name for the hypervisor and application configuration.
- viii. On the Plug-in Type page, select **Application plug-in**.
- ix. On the Application Plug-ins page, select the application to be used for backup and restore.
- x. Provide the details for the selected application.
- xi. On the Agent Configuration page, provide the application Snap Creator Agent details, which are the details of the application or database host on which you have installed the agent.



Typically, the host is a virtual machine being backed up that has an application running on it.

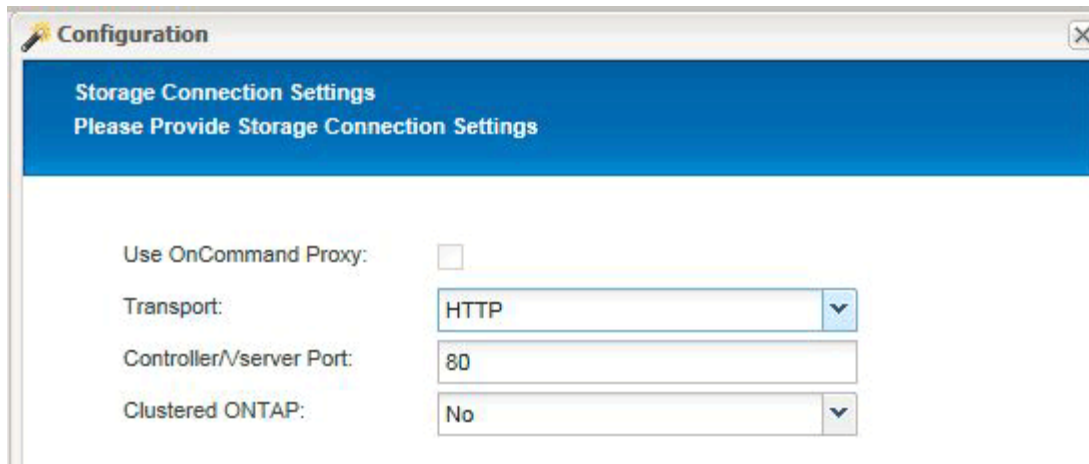
Click **Test agent connection** to make sure that the agent is running.

- xii. On the Summary page, verify the information and click **Finish**.
- xiii. On the Hypervisor + App page, you have the following options:
 - To add additional applications to this configuration, click **Add** and repeat steps vii through xii in this example.
 - To delete applications from this configuration, select the item and click **Delete**.
 - To continue with the main Configuration wizard, click **Next**.



If you have multiple applications listed, you have the option to reorder this list by moving an application up or down in the list. Applications are backed up serially, so if an application needs to be quiesced before another one in the list, you need to place the applications in the proper sequence.

- e. On the Storage Connection Settings page, provide the following information:
 - For the **Transport** option, select **HTTPS**.
 - For the **Controller/Vserver Port** option, leave the default setting (443).
 - For the **Storage Type** option, select **Clustered ONTAP**.
 - For the **Storage Type** option, select **Yes**.



Configuration

Storage Connection Settings
Please Provide Storage Connection Settings

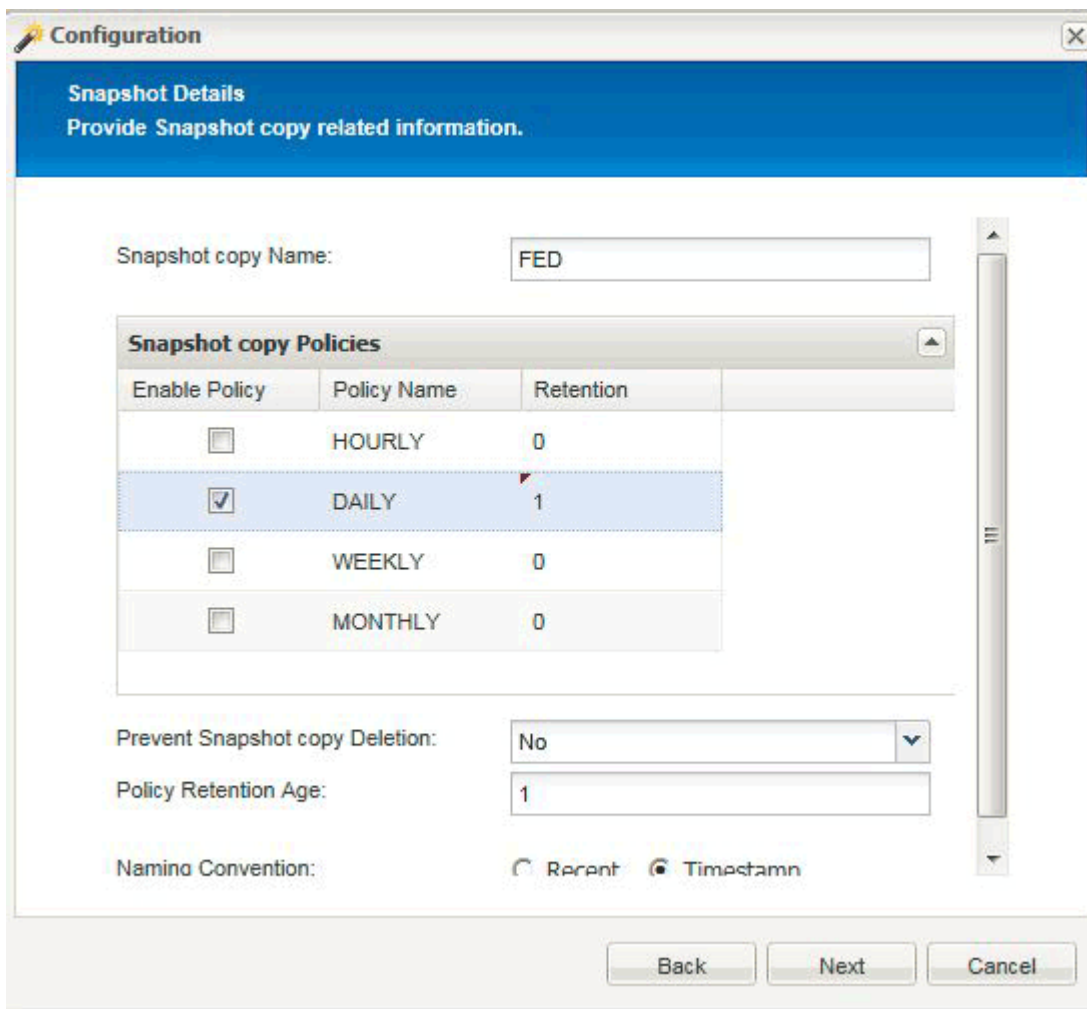
Use OnCommand Proxy:

Transport: HTTP

Controller/Vserver Port: 80

Clustered ONTAP: No

- f. On the New Controller/Vserver page, provide the controller IP address, username, and password.
- g. Provide the Snapshot copy details.



Configuration

Snapshot Details
Provide Snapshot copy related information.

Snapshot copy Name: FED

Snapshot copy Policies		
Enable Policy	Policy Name	Retention
<input type="checkbox"/>	HOURLY	0
<input checked="" type="checkbox"/>	DAILY	1
<input type="checkbox"/>	WEEKLY	0
<input type="checkbox"/>	MONTHLY	0

Prevent Snapshot copy Deletion: No

Policy Retention Age: 1

Namino Convention: Recent Timestamn

Back Next Cancel

- h. On the Snapshot Details Continued page, do not select the **Consistency Group** option.
- i. On the Data Protection page, do not select either of the **Data Transfer** options.
- j. Verify the information on the Summary page and click **Finish**.

Related information

Creating profiles

Creating configuration files

Copyright Information

Copyright © 2021 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.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

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.