



NetApp

Cloud Insights

NetApp
June 25, 2024

This PDF was generated from https://docs.netapp.com/us-en/cloudinsights/task_dc_na_cloud_volumes_ontap.html on June 25, 2024. Always check docs.netapp.com for the latest.

Table of Contents

- NetApp 1
 - NetApp Cloud Volumes ONTAP data collector 1
 - NetApp Cloud Volumes Services for AWS data collector 1
 - NetApp ONTAP Data Management Software data collector 2
 - NetApp ONTAP REST data collector 9
 - NetApp Data ONTAP operating in 7-Mode data collector 15
 - NetApp E-Series Legacy Santricity API data collector 19
 - NetApp E-Series REST data collector 22
 - Configuring the NetApp HCI Management server data collector 25
 - NetApp SolidFire All-Flash Array data collector 27
 - NetApp StorageGRID data collector 29

NetApp

NetApp Cloud Volumes ONTAP data collector

This data collector supports inventory collection from Cloud Volumes ONTAP configurations.

Configuration

Field	Description
NetApp Management IP Address	IP address for Cloud Volumes ONTAP
User Name	User name for Cloud Volumes ONTAP
Password	Password for the above user

Advanced configuration

Field	Description
Connection Type	HTTPS recommended. Also shows default port.
Override Communication Port	Port to use if not default.
Inventory Poll Interval (min)	Default is 60 minutes.
Inventory Concurrent Thread Count	Number of concurrent threads.
Force TLS for HTTPS	Force TLS over HTTPS
Automatically Lookup Netgroups	Automatically Lookup Netgroups
Netgroup Expansion	Select Shell or File
HTTP read timeout seconds	Default is 30 seconds
Force responses as UTF-8	Force responses as UTF-8
Performance Poll Interval (min)	Default is 900 seconds.
Performance Concurrent Thread Count	Number of concurrent threads.
Advanced Counter Data Collection	Check this to have Cloud Insights collect the advanced metrics from the list below.

Troubleshooting

Additional information on this Data Collector may be found from the [Support](#) page or in the [Data Collector Support Matrix](#).

NetApp Cloud Volumes Services for AWS data collector

This data collector supports inventory collection from NetApp Cloud Volumes Services for AWS configurations.

Configuration

Field	Description
Cloud Volumes Region	Region of the NetApp Cloud Volumes Services for AWS
API Key	Cloud Volumes API key
Secret Key	Cloud Volumes secret key

Advanced configuration

Field	Description
Inventory Poll Interval (min)	Default is 60 minutes

Troubleshooting

Some things to try if you encounter problems with this data collector:

Problem:	Try this:
I received an error similar to this one: 'Failed to execute request: Connect to <AWS region endpoint>:8080 [<AWS region endpoint>/AWS region endpoint IP>] failed: connect timed out: GET https://<AWS Region Endpoint FQDN>:8080/v1/Storage/IPRanges HTTP/1.1'	<p>The proxy used by Cloud Insights to communicate with the Acquisition Unit does not communicate between Cloud Insights and the Data Collector itself. Here are a few things you can try:</p> <p>Ensure that the acquisition unit is able to resolve the fqdn and reach the required port. Confirm that a proxy is not required to reach the specified endpoint in the error message. Curl can be used to test the communication between the acquisition unit and the endpoint. Make sure that you are not using a Proxy for this test.</p> <p>Example:</p> <pre>root@acquisitionunit# curl -s -H accept:application/json -H "Content-type: application/json" -H api-key:<api key used in the data collector credentials> -H secret-key:<secret key used in the data collector credentials> -X GET https://<AWS Regional Endpoint>:8080/v1/Storage/IPRanges</pre> <p>See this NetApp KB article.</p>

Additional information on this Data Collector may be found from the [Support](#) page or in the [Data Collector Support Matrix](#).

NetApp ONTAP Data Management Software data collector

This data collector acquires inventory and performance data from storage systems

running ONTAP using read-only API calls from an ONTAP account. This data collector also creates a record in the cluster application registry to accelerate support.

Terminology

Cloud Insights acquires inventory and performance data from the ONTAP data collector. For each asset type acquired, the most common terminology used for the asset is shown. When viewing or troubleshooting this data collector, keep the following terminology in mind:

Vendor/Model Term	Cloud Insights Term
Disk	Disk
Raid Group	Disk Group
Cluster	Storage
Node	Storage Node
Aggregate	Storage Pool
LUN	Volume
Volume	Internal Volume

ONTAP Data Management Terminology

The following terms apply to objects or references that you might find on ONTAP Data Management storage asset landing pages. Many of these terms apply to other data collectors as well.

Storage

- Model – A comma-delimited list of the unique, discrete node model names within this cluster. If all the nodes in the clusters are the same model type, just one model name will appear.
- Vendor – same Vendor name you would see if you were configuring a new data source.
- Serial number – The array serial number. On cluster architecture storage systems like ONTAP Data Management, this serial number may be less useful than the individual “Storage Nodes” serial numbers.
- IP – generally will be the IP(s) or hostname(s) as configured in the data source.
- Microcode version – firmware.
- Raw Capacity – base 2 summation of all the physical disks in the system, regardless of their role.
- Latency – a representation of what the host facing workloads are experiencing, across both reads and writes. Ideally, Cloud Insights is sourcing this value directly, but this is often not the case. In lieu of the array offering this up, Cloud Insights is generally performing an IOPs-weighted calculation derived from the individual internal volumes’ statistics.
- Throughput – aggregated from internal volumes.
Management – this may contain a hyperlink for the management interface of the device. Created programmatically by the Cloud Insights data source as part of inventory reporting.

Storage Pool

- Storage – what storage array this pool lives on. Mandatory.
- Type – a descriptive value from a list of an enumerated list of possibilities. Most commonly will be

“Aggregate” or “RAID Group”.

- Node – if this storage array’s architecture is such that pools belong to a specific storage node, its name will be seen here as a hyperlink to its own landing page.
- Uses Flash Pool – Yes/No value – does this SATA/SAS based pool have SSDs used for caching acceleration?
- Redundancy – RAID level or protection scheme. RAID_DP is dual parity, RAID_TP is triple parity.
- Capacity – the values here are the logical used, usable capacity and the logical total capacity, and the percentage used across these.
- Over-committed capacity – If by using efficiency technologies you have allocated a sum total of volume or internal volume capacities larger than the logical capacity of the storage pool, the percentage value here will be greater than 0%.
- Snapshot – snapshot capacities used and total, if your storage pool architecture dedicates part of its capacity to segments areas exclusively for snapshots. ONTAP in MetroCluster configurations are likely to exhibit this, while other ONTAP configurations are less so.
- Utilization – a percentage value showing the highest disk busy percentage of any disk contributing capacity to this storage pool. Disk utilization does not necessarily have a strong correlation with array performance – utilization may be high due to disk rebuilds, deduplication activities, etc in the absence of host driven workloads. Also, many arrays’ replication implementations may drive disk utilization while not showing as internal volume or volume workload.
- IOPS – the sum IOPs of all the disks contributing capacity to this storage pool.
Throughput – the sum throughput of all the disks contributing capacity to this storage pool.

Storage Node

- Storage – what storage array this node is part of. Mandatory.
- HA Partner – on platforms where a node will fail over to one and only one other node, it will generally be seen here.
- State – health of the node. Only available when the array is healthy enough to be inventoried by a data source.
- Model – model name of the node.
- Version – version name of the device.
- Serial number – The node serial number.
- Memory – base 2 memory if available.
- Utilization – On ONTAP, this is a controller stress index from a proprietary algorithm. With every performance poll, a number between 0 and 100% will be reported that is the higher of either WAFL disk contention, or average CPU utilization. If you observe sustained values > 50%, that is indicative of undersizing – potentially a controller/node not large enough or not enough spinning disks to absorb the write workload.
- IOPS – Derived directly from ONTAP ZAPI calls on the node object.
- Latency – Derived directly from ONTAP ZAPI calls on the node object.
- Throughput – Derived directly from ONTAP ZAPI calls on the node object.
- Processors – CPU count.

Requirements

The following are requirements to configure and use this data collector:

- You must have access to an Administrator account configured for read-only API calls.
- Account details include username and password.
- Port requirements: 80 or 443
- Account permissions:
 - Read only role name to ontapi application to the default Vserver
 - You may require additional optional write permissions. See the Note About Permissions below.
- ONTAP License requirements:
 - FCP license and mapped/masked volumes required for fibre-channel discovery

Permission Requirements for Collecting ONTAP Switch Metrics

Cloud Insights has the ability to collect ONTAP cluster switch data as an option in the collector's [Advanced Configuration](#) settings. In addition to enabling this on the Cloud Insights collector, you must also **configure the ONTAP system** itself to provide [switch information](#), and ensure the correct [permissions](#) are set, in order to allow the switch data to be sent to Cloud Insights.

Configuration

Field	Description
NetApp Management IP	IP address or fully-qualified domain name of the NetApp cluster
User Name	User name for NetApp cluster
Password	Password for NetApp cluster

Advanced configuration

Field	Description
Connection type	Choose HTTP (default port 80) or HTTPS (default port 443). The default is HTTPS
Override Communication Port	Specify a different port if you do not want to use the default
Inventory Poll Interval (min)	Default is 60 minutes.
For TLS for HTTPS	Only allow TLS as protocol when using HTTPS
Automatically Lookup Netgroups	Enable the automatic netgroup lookups for export policy rules
Netgroup Expansion	Netgroup Expansion Strategy. Choose <i>file</i> or <i>shell</i> . The default is <i>shell</i> .
HTTP read timeout seconds	Default is 30

Field	Description
Force responses as UTF-8	Forces data collector code to interpret responses from the CLI as being in UTF-8
Performance Poll Interval (sec)	Default is 900 seconds.
Advanced Counter Data Collection	Enable ONTAP integration. Select this to include ONTAP Advanced Counter data in polls. Choose the desired counters from the list.
Cluster Switch Metrics	Allow Cloud Insights to collect cluster switch data. Note that in addition to enabling this on the Cloud Insights side, you must also configure the ONTAP system to provide switch information , and ensure the correct permissions are set, in order to allow the switch data to be sent to Cloud Insights. See "A Note About Permissions" below.

ONTAP Power Metrics

Several ONTAP models provide power metrics for Cloud Insights that can be used for monitoring or alerting. The lists of supported and unsupported models below are not comprehensive but should provide some guidance; in general, if a model is in the same family as one on the list, the support should be the same.

Supported Models:

A200
 A220
 A250
 A300
 A320
 A400
 A700
 A700s
 A800
 A900
 C190
 FAS2240-4
 FAS2552
 FAS2650
 FAS2720
 FAS2750
 FAS8200
 FAS8300
 FAS8700
 FAS9000

Unsupported Models:

FAS2620
 FAS3250
 FAS3270
 FAS500f
 FAS6280

FAS/AFF 8020
FAS/AFF 8040
FAS/AFF 8060
FAS/AFF 8080

A Note About Permissions

Since a number of Cloud Insights' ONTAP dashboards rely on advanced ONTAP counters, you must enable **Advanced Counter Data Collection** in the data collector Advanced Configuration section.

You should also ensure that write permission to the ONTAP API is enabled. This typically requires an account at the cluster level with the necessary permissions.

To create a local account for Cloud Insights at the cluster level, log in to ONTAP with the Cluster management Administrator username/password, and execute the following commands on the ONTAP server:

1. Before you begin, you must be signed in to ONTAP with an *Administrator* account, and *diagnostic-level commands* must be enabled.
2. Create a read-only role using the following commands.

```
security login role create -role ci_readonly -cmddirname DEFAULT -access  
readonly  
security login role create -role ci_readonly -cmddirname security  
-access readonly  
security login role create -role ci_readonly -access all -cmddirname  
{cluster application-record create}
```

3. Create the read-only user using the following command. Once you have executed the create command, you will be prompted to enter a password for this user.

```
security login create -username ci_user -application ontapi  
-authentication-method password -role ci_readonly
```

If AD/LDAP account is used, the command should be

```
security login create -user-or-group-name DOMAIN\aduser/adgroup  
-application ontapi -authentication-method domain -role ci_readonly
```

If you are collecting cluster switch data:

```
security login rest-role create -role ci_readonly -api  
/api/network/ethernet -access readonly
```

The resulting role and user login will look something like the following. Your actual output may vary:

```

Role Command/ Access
Vserver Name Directory Query Level
-----
cluster1 ci_readonly DEFAULT read only
cluster1 ci_readonly security readonly

```

```

cluster1::security login> show
Vserver: cluster1
Authentication Acct
UserName      Application  Method      Role Name    Locked
-----
ci_user       ontapi      password    ci_readonly  no

```



If ONTAP access control is not set correctly, then Cloud Insights REST calls may fail, resulting in gaps in data for the device. For example, if you have enabled it on the Cloud Insights collector but have not configured the permissions on the ONTAP, acquisition will fail. Additionally, if the role is previously defined on the ONTAP and you are adding the Rest API abilities, ensure that *http* is added to the role.

Troubleshooting

Some things to try if you encounter problems with this data collector:

Inventory

Problem:	Try this:
Receive 401 HTTP response or 13003 ZAPI error code and ZAPI returns "Insufficient privileges" or "not authorized for this command"	Check username and password, and user privileges/permissions.
Cluster version is < 8.1	Cluster minimum supported version is 8.1. Upgrade to minimum supported version.
ZAPI returns "cluster role is not cluster_mgmt LIF"	AU needs to talk to cluster management IP. Check the IP and change to a different IP if necessary
Error: "7 Mode filers are not supported"	This can happen if you use this data collector to discover 7 mode filer. Change IP to point to cdot cluster instead.
ZAPI command fails after retry	AU has communication problem with the cluster. Check network, port number, and IP address. User should also try to run a command from command line from the AU machine.
AU failed to connect to ZAPI via HTTP	Check whether ZAPI port accepts plaintext. If AU tries to send plaintext to an SSL socket, the communication fails.

Problem:	Try this:
Communication fails with SSLException	AU is attempting to send SSL to a plaintext port on a filer. Check whether the ZAPI port accepts SSL, or use a different port.
Additional Connection errors: ZAPI response has error code 13001, "database is not open" ZAPI error code is 60 and response contains "API did not finish on time" ZAPI response contains "initialize_session() returned NULL environment" ZAPI error code is 14007 and response contains "Node is not healthy"	Check network, port number, and IP address. User should also try to run a command from command line from the AU machine.

Performance

Problem:	Try this:
"Failed to collect performance from ZAPI" error	This is usually due to perf stat not running. Try the following command on each node: <pre>> system node systemshell -node * -command "spmctl -h cmd -stop; spmctl -h cmd -exec"</pre>

Additional information may be found from the [Support](#) page or in the [Data Collector Support Matrix](#).

NetApp ONTAP REST data collector

This data collector acquires inventory, EMS logs, and performance data from storage systems running ONTAP 9.14.1 and higher using REST API calls. For Ontap systems on earlier releases, use the ZAPI-based "NetApp ONTAP Data Management Software" collector type.



The ONTAP REST collector may be used as a replacement for the previous ONTAPI-based collector. As such, there may be differences in the metrics that are collected or reported. For more information about the differences between ONTAPI and REST, see the [ONTAP 9.14.1 ONTAPI-to-REST mapping](#) documentation.

Requirements

The following are requirements to configure and use this data collector:

- You must have access to a user account with the required level of access. Note that Admin permissions are required if creating a new REST user/role.
 - Functionally, Cloud Insights primarily makes read requests, but some write permissions are required for Cloud Insights to register with the ONTAP array. See the *Note About Permissions* immediately below.

- ONTAP version 9.14.1 or higher.
- Port requirements: 443

A Note About Permissions

Since a number of Cloud Insights' ONTAP dashboards rely on advanced ONTAP counters, you should keep **Enable Advanced Counter Data Collection** enabled in the data collector Advanced Configuration section.

To create a local account for Cloud Insights at the cluster level, log in to ONTAP with the Cluster management Administrator username/password, and execute the following commands on the ONTAP server:

1. Before you begin, you must be signed in to ONTAP with an *Administrator* account, and *diagnostic-level commands* must be enabled.
2. Retrieve the name of the vservers that is of type *admin*. You will use this name in subsequent commands.

```
vserver show -type admin
```

3. Create a role using the following commands:

```
security login rest-role create -role {role name} -api /api -access
readonly
security login rest-role create -role {role name} -api
/api/cluster/agents -access all
vserver services web access create -name spi -role {role name} -vserver
{vserver name as retrieved above}
security login create -user-or-group-name {username} -application http
-authentication-method password -role {role name}
```

4. Create the read-only user using the following command. Once you have executed the create command, you will be prompted to enter a password for this user.

```
security login create -username ci_user -application http
-authentication-method password -role ci_readonly
```

If AD/LDAP account is used, the command should be

```
security login create -user-or-group-name DOMAIN\aduser/adgroup
-application http -authentication-method domain -role ci_readonly
```

The resulting role and user login will look something like the following. Your actual output may vary:

```
security login rest-role show -vserver <vserver name> -role restRole
```

Vserver	Role Name	API	Access Level
<vserver name>	restRole	/api	readonly
		/api/cluster/agents	all

2 entries were displayed.

```
security login show -vserver <vserver name> -user-or-group-name restUser
```

Vserver: <vserver name>

User/Group	Authentication	Acct	Second
Name	Application Method	Role Name	Locked Method
restUser	http password	restRole	no none

Migration

To migrate from a previous ONTAP (ontapi) data collector to the newer ONTAP REST collector, do the following:

1. Add the REST Collector. It is recommended to enter information for a different user than the one configured for the previous collector. For example, use the user noted in the Permissions section above.
2. Pause the previous collector, so it doesn't continue to collect data.
3. Let the new REST collector acquire data for at least 30 minutes. Ignore any data during this time that does not appear "normal".
4. After the rest period, you should see your data stabilize as the REST collector continues to to acquire.

You can use this same process to return to the previous collector, should you wish.

Configuration

Field	Description
ONTAP management IP Address	IP address or fully-qualified domain name of the NetApp cluster. Must be Cluster Management IP/FQDN.
ONTAP REST User Name	User name for NetApp cluster
ONTAP REST Password	Password for NetApp cluster

Advanced configuration

Field	Description
Inventory Poll Interval (min)	Default is 60 minutes.
Performance Poll Interval (sec)	Default is 60 seconds.
Advanced Counter Data Collection	Select this to include ONTAP Advanced Counter data in polls. Enabled by default.
Enable EMS Event Collection	Select this to include ONTAP EMS log event data. Enabled by default.
EMS Poll Interval (sec)	Default is 60 seconds.

Terminology

Cloud Insights acquires inventory, logs and performance data from the ONTAP data collector. For each asset type acquired, the most common terminology used for the asset is shown. When viewing or troubleshooting this data collector, keep the following terminology in mind:

Vendor/Model Term	Cloud Insights Term
Disk	Disk
Raid Group	Disk Group
Cluster	Storage
Node	Storage Node
Aggregate	Storage Pool
LUN	Volume
Volume	Internal Volume
Storage Virtual Machine/Vserver	Storage Virtual Machine

ONTAP Data Management Terminology

The following terms apply to objects or references that you might find on ONTAP Data Management storage asset landing pages. Many of these terms apply to other data collectors as well.

Storage

- Model – A comma-delimited list of the unique, discrete node model names within this cluster. If all the nodes in the clusters are the same model type, just one model name will appear.
- Vendor – same Vendor name you would see if you were configuring a new data source.
- Serial number – The array UUID
- IP – generally will be the IP(s) or hostname(s) as configured in the data source.
- Microcode version – firmware.
- Raw Capacity – base 2 summation of all the physical disks in the system, regardless of their role.
- Latency – a representation of what the host facing workloads are experiencing, across both reads and

writes. Ideally, Cloud Insights is sourcing this value directly, but this is often not the case. In lieu of the array offering this up, Cloud Insights is generally performing an IOPs-weighted calculation derived from the individual internal volumes' statistics.

- Throughput – aggregated from internal volumes.
Management – this may contain a hyperlink for the management interface of the device. Created programmatically by the Cloud Insights data source as part of inventory reporting.

Storage Pool

- Storage – what storage array this pool lives on. Mandatory.
- Type – a descriptive value from a list of an enumerated list of possibilities. Most commonly will be “Aggregate” or “RAID Group”.
- Node – if this storage array's architecture is such that pools belong to a specific storage node, its name will be seen here as a hyperlink to its own landing page.
- Uses Flash Pool – Yes/No value – does this SATA/SAS based pool have SSDs used for caching acceleration?
- Redundancy – RAID level or protection scheme. RAID_DP is dual parity, RAID_TP is triple parity.
- Capacity – the values here are the logical used, usable capacity and the logical total capacity, and the percentage used across these.
- Over-committed capacity – If by using efficiency technologies you have allocated a sum total of volume or internal volume capacities larger than the logical capacity of the storage pool, the percentage value here will be greater than 0%.
- Snapshot – snapshot capacities used and total, if your storage pool architecture dedicates part of its capacity to segments areas exclusively for snapshots. ONTAP in MetroCluster configurations are likely to exhibit this, while other ONTAP configurations are less so.
- Utilization – a percentage value showing the highest disk busy percentage of any disk contributing capacity to this storage pool. Disk utilization does not necessarily have a strong correlation with array performance – utilization may be high due to disk rebuilds, deduplication activities, etc in the absence of host driven workloads. Also, many arrays' replication implementations may drive disk utilization while not showing as internal volume or volume workload.
- IOPS – the sum IOPs of all the disks contributing capacity to this storage pool.
Throughput – the sum throughput of all the disks contributing capacity to this storage pool.

Storage Node

- Storage – what storage array this node is part of. Mandatory.
- HA Partner – on platforms where a node will fail over to one and only one other node, it will generally be seen here.
- State – health of the node. Only available when the array is healthy enough to be inventoried by a data source.
- Model – model name of the node.
- Version – version name of the device.
- Serial number – The node serial number.
- Memory – base 2 memory if available.
- Utilization – On ONTAP, this is a controller stress index from a proprietary algorithm. With every performance poll, a number between 0 and 100% will be reported that is the higher of either WAFL disk

contention, or average CPU utilization. If you observe sustained values > 50%, that is indicative of undersizing – potentially a controller/node not large enough or not enough spinning disks to absorb the write workload.

- IOPS – Derived directly from ONTAP REST calls on the node object.
- Latency – Derived directly from ONTAP REST calls on the node object.
- Throughput – Derived directly from ONTAP REST calls on the node object.
- Processors – CPU count.

ONTAP Power Metrics

Several ONTAP models provide power metrics for Cloud Insights that can be used for monitoring or alerting. The lists of supported and unsupported models below are not comprehensive but should provide some guidance; in general, if a model is in the same family as one on the list, the support should be the same.

Supported Models:

A200
A220
A250
A300
A320
A400
A700
A700s
A800
A900
C190
FAS2240-4
FAS2552
FAS2650
FAS2720
FAS2750
FAS8200
FAS8300
FAS8700
FAS9000

Unsupported Models:

FAS2620
FAS3250
FAS3270
FAS500f
FAS6280
FAS/AFF 8020
FAS/AFF 8040
FAS/AFF 8060
FAS/AFF 8080

Troubleshooting

Some things to try if you encounter problems with this data collector:

Problem:	Try this:
<p>When attempting to create an ONTAP REST data collector, an error like the following is seen: Configuration: 10.193.70.14: ONTAP rest API at 10.193.70.14 is not available: 10.193.70.14 failed to GET /api/cluster: 400 Bad Request</p>	<p>This is likely due to an older ONTAP array) for example, ONTAP 9.6) which has no REST API capabilities. ONTAP 9.14.1 is the minimum ONTAP version supported by the ONTAP REST collector. "400 Bad Request" responses should be expected on pre-REST ONTAP releases.</p> <p>For ONTAP versions that do support REST but are not 9.14.1 or later, you may see the following similar message: Configuration: 10.193.98.84: ONTAP rest API at 10.193.98.84 is not available: 10.193.98.84: ONTAP rest API at 10.193.98.84 is available: cheryl5-cluster-2 9.10.1 a3cb3247-3d3c-11ee-8ff3-005056b364a7 but is not of minimum version 9.14.1.</p>
<p>I see empty or "0" metrics where the ONTAP ontapi collector shows data.</p>	<p>ONTAP REST does not report metrics that are used internally on the ONTAP system only. For example, system aggregates will not be collected by ONTAP REST, only SVM's of type "data" will be collected.</p> <p>Other examples of ONTAP REST metrics that may report zero or empty data:</p> <p>InternalVolumes: REST no longer reports vol0. Aggregates: REST no longer reports aggr0. Storage: most metrics are a rollup of the Internal Volume metrics, and will be impacted by the above. Storage Virtual Machines: REST no longer reports SVM's of type other than 'data' (e.g. 'cluster', 'mgmt', 'node').</p> <p>You may also notice a change in the appearance of graphs that do have data, due to the change in default performance polling period from 15 minutes to 5 minutes. More frequent polling means more data points to plot.</p>

Additional information may be found from the [Support](#) page or in the [Data Collector Support Matrix](#).

NetApp Data ONTAP operating in 7-Mode data collector

For storage systems using Data ONTAP software operating in 7-Mode, you use the 7-mode data collector, which uses the CLI to obtain capacity and performance data.

Terminology

Cloud Insights acquires the following inventory information from the NetApp 7-mode data collector. For each asset type acquired, the most common terminology used for this asset is shown. When viewing or troubleshooting this data collector, keep the following terminology in mind:



This data collector is [deprecated](#).

Vendor/Model Term	Cloud Insights Term
Disk	Disk
Raid Group	Disk Group
Filer	Storage
Filer	Storage Node
Aggregate	Storage Pool
LUN	Volume
Volume	Internal Volume

Note: These are common terminology mappings only and might not represent every case for this data collector.

Requirements

You need the following to configure and use this data collector:

- IP addresses of the FAS storage controller and partner.
- Port 443
- A custom admin level username and password for controller and partner controller with the following role capabilities for 7-Mode:
 - "api-*": Use this to allow OnCommand Insight to execute all NetApp storage API commands.
 - "login-http-admin": Use this to allow OnCommand Insight to connect to the NetApp storage via HTTP.
 - "security-api-vfiler": Use this to allow OnCommand Insight to execute NetApp storage API commands to retrieve vFiler unit information.
 - "cli-options": Use this to read storage system options.
 - "cli-lun": Access these commands for managing LUNs. Displays the status (LUN path, size, online/offline state, and shared state) of the given LUN or class of LUNs.
 - "cli-df": Use this to display free disk space.
 - "cli-ifconfig": Use this to display interfaces and IP addresses.

Configuration

Field	Description
Address of storage system	IP address or fully-qualified domain name for the NetApp storage system
User Name	User name for the NetApp storage system
Password	Password for the NetApp storage system
Address of HA Partner in Cluster	IP address or fully-qualified domain name for the HA Partner

Field	Description
User Name of HA Partner in Cluster	User name for the HA partner
Password of HA Partner Filer in Cluster	Password for the HA Partner

Advanced configuration

Field	Description
Inventory Poll Interval (min)	Interval between inventory polls. The default is 20 minutes.
Connection Type	HTTPS or HTTP, also displays the default port
Override Connection Port	If blank, use the default port in the Connection Type field, otherwise enter the connection port to use
Performance Poll Interval (sec)	Interval between performance polls. The default is 300 seconds.

Storage systems connection

As an alternative to using the default administrative user for this data collector, you can configure a user with administrative rights directly on the NetApp storage systems so that this data collector can acquire data from NetApp storage systems.

Connecting to NetApp storage systems requires that the user, who is specified when acquiring the main pfiler (on which the storage system exist), meet the following conditions:

- The user must be on vfiler0 (root filer/pfiler).

Storage systems are acquired when acquiring the main pfiler.

- The following commands define the user role capabilities:

- "api-*": Use this to allow Cloud Insights to execute all NetApp storage API commands.

This command is required to use the ZAPI.

- "login-http-admin": Use this to allow Cloud Insights to connect to the NetApp storage via HTTP. This command is required to use the ZAPI.
- "security-api-vfiler": Use this to allow Cloud Insights to execute NetApp storage API commands to retrieve vFiler unit information.
- "cli-options": For "options" command and used for partner IP and enabled licenses.
- "cli-lun": Access these command for managing LUNs. Displays the status (LUN path, size, online/offline state, and shared state) of the given LUN or class of LUNs.
- "cli-df": For "df -s", "df -r", "df -A -r" commands and used to display free space.
- "cli-ifconfig": For "ifconfig -a" command and used for getting filer IP address.
- "cli-rdfile": For "rdfile /etc/netgroup" command and used for getting netgroups.
- "cli-date": For "date" command and used to get full date for getting Snapshot copies.
- "cli-snap": For "snap list" command and used for getting Snapshot copies.

If cli-date or cli-snap permissions are not provided, acquisition can finish, but Snapshot copies are not reported.

To acquire a 7-Mode data source successfully and generate no warnings on the storage system, you should use one of the following command strings to define your user roles. The second string listed here is a streamlined version of the first:

- login-http-admin,api-*,security-api-vfile,cli-rdfile,cli-options,cli-df,cli-lun,cli-ifconfig,cli-date,cli-snap,_,
- login-http-admin,api-*,security-api-vfile,cli-

Troubleshooting

Some things to try if you encounter problems with this data collector:

Inventory

Problem:	Try this:
Receive 401 HTTP response or 13003 ZAPI error code and ZAPI returns “Insufficient privileges” or “not authorized for this command”	Check username and password, and user privileges/permissions.
“Failed to execute command” error	Check whether the user has the following permission on the device: <ul style="list-style-type: none"> • api-* • cli-date • cli-df • cli-ifconfig • cli-lun • cli-operations • cli-rdfile • cli-snap • login-http-admin • security-api-vfiler Also check if the ONTAP version is supported by Cloud Insights and verify if the credentials used match device credentials
Cluster version is < 8.1	Cluster minimum supported version is 8.1. Upgrade to minimum supported version.
ZAPI returns "cluster role is not cluster_mgmt LIF"	AU needs to talk to cluster management IP. Check the IP and change to a different IP if necessary
Error: “7 Mode filers are not supported”	This can happen if you use this data collector to discover 7 mode filer. Change IP to point to cdot filer instead.
ZAPI command fails after retry	AU has communication problem with the cluster. Check network, port number, and IP address. User should also try to run a command from command line from the AU machine.
AU failed to connect to ZAPI	Check IP/port connectivity and assert ZAPI configuration.

Problem:	Try this:
AU failed to connect to ZAPI via HTTP	Check whether ZAPI port accepts plaintext. If AU tries to send plaintext to an SSL socket, the communication fails.
Communication fails with SSLException	AU is attempting to send SSL to a plaintext port on a filer. Check whether the ZAPI port accepts SSL, or use a different port.
Additional Connection errors: ZAPI response has error code 13001, “database is not open” ZAPI error code is 60 and response contains “API did not finish on time” ZAPI response contains “initialize_session() returned NULL environment” ZAPI error code is 14007 and response contains “Node is not healthy”	Check network, port number, and IP address. User should also try to run a command from command line from the AU machine.
Socket timeout error with ZAPI	Check filer connectivity and/or increase timeout.
“C Mode clusters are not supported by the 7 Mode data source” error	Check IP and change the IP to a 7 Mode cluster.
“Failed to connect to vFiler” error	Check that the acquiring user capabilities include the following at a minimum: api- security-api-vfiler login-http-admin Confirm that filer is running minimum ONTAPI version 1.7.

Additional information may be found from the [Support](#) page or in the [Data Collector Support Matrix](#).

NetApp E-Series Legacy Santricity API data collector

The NetApp E-Series Legacy Santricity API data collector gathers inventory and performance data. The collector supports firmware 7.x+ using the same configurations and reporting the same data.

Terminology

Cloud insight acquires the following inventory information from the NetApp E-Series data collector. For each asset type acquired, the most common terminology used for this asset is shown. When viewing or troubleshooting this data collector, keep the following terminology in mind:

Vendor/Model Term	Cloud Insights Term
Disk	Disk

Vendor/Model Term	Cloud Insights Term
Volume Group	Disk Group
Storage Array	Storage
Controller	Storage Node
Volume Group	Storage Pool
Volume	Volume

Note: These are common terminology mappings only and might not represent every case for this data collector.

E-Series Terminology (Landing Page)

The following terms apply to objects or references that you might find on NetApp E-Series asset landing pages. Many of these terms apply to other data collectors as well.

Storage

- Model – model name of the device.
- Vendor – same Vendor name you would see if you were configuring a new datasource
- Serial number – The array serial number. On cluster architecture storage systems like NetApp Clustered Data Ontap, this serial number may be less useful than the individual “Storage Nodes” serial numbers
- IP – generally will be the IP(s) or hostname(s) as configured in the data source
- Microcode version – firmware
- Raw Capacity – base 2 summation of all the physical disks in the system, regardless of their role
- Latency – a representation of what the host facing workloads are experiencing, across both reads and writes. Ideally, Cloud Insights is sourcing this value directly, but this is often not the case. In lieu of the array offering this up, Cloud Insights is generally performing an IOPs-weighted calculation derived from the individual volumes’ statistics.
- Throughput – the array’s total host facing throughput. Ideally sourced directly from the array, if unavailable, Cloud Insights is summing the volumes’ throughput to derive this value
- Management – this may contain a hyperlink for the management interface of the device. Created programmatically by the Cloud Insights datasource as part of inventory reporting

Storage Pool

- Storage – what storage array this pool lives on. Mandatory
- Type – a descriptive value from a list of an enumerated list of possibilities. Most commonly will be “Thin Provisioning” or “RAID Group”
- Node – if this storage array’s architecture is such that pools belong to a specific storage node, its name will be seen here as a hyperlink to its own landing page
- Uses Flash Pool – Yes/No value
- Redundancy – RAID level or protection scheme. E-Series reports “RAID 7” for DDP pools
- Capacity – the values here are the logical used, usable capacity and the logical total capacity, and the percentage used across these. These value both include E-Series “preservation” capacity, resulting both in

numbers and the percentage being higher than what the E-Series own user interface may show

- Over-committed capacity – If via efficiency technologies you have allocated a sum total of volume or internal volume capacities larger than the logical capacity of the storage pool, the percentage value here will be greater than 0%.
- Snapshot – snapshot capacities used and total, if your storage pool architecture dedicates part of its capacity to segments areas exclusively for snapshots
- Utilization – a percentage value showing the highest disk busy percentage of any disk contributing capacity to this storage pool. Disk utilization does not necessarily have a strong correlation with array performance – utilization may be high due to disk rebuilds, deduplication activities, etc in the absence of host driven workloads. Also, many arrays' replication implementations may drive disk utilization while not showing as volume workload.
- IOPS – the sum IOPs of all the disks contributing capacity to this storage pool. If disk IOPs is not available on a given platform, this value will be sourced from the sum of volume IOPs for all the volumes sitting on this storage pool
- Throughput – the sum throughput of all the disks contributing capacity to this storage pool. If disk throughput is not available on a given platform, this value will be sourced from the sum of volume throughput for all the volumes sitting on this storage pool

Storage Node

- Storage – what storage array this node is part of. Mandatory
- HA Partner – on platforms where a node will fail over to one and only one other node, it will generally be seen here
- State – health of the node. Only available when the array is healthy enough to be inventoried by a data source
- Model – model name of the node
- Version – version name of the device.
- Serial number – The node serial number
- Memory – base 2 memory if available
- Utilization – Generally a CPU utilization number, or in the case of NetApp Ontap, a controller stress index. Utilization is not currently available for NetApp E-Series
- IOPS – a number representing the host driven IOPs on this controller. Ideally sourced directly from the array, if unavailable, it will be calculated by summing all the IOPs for volumes that belong exclusively to this node.
- Latency – a number representing the typical host latency or response time on this controller. Ideally sourced directly from the array, if unavailable, it will be calculated by performing an IOPs weighted calculation from volumes that belong exclusively to this node.
- Throughput – a number representing the host driven throughput on this controller. Ideally sourced directly from the array, if unavailable, it will be calculated by summing all the throughput for volumes that belong exclusively to this node.
- Processors – CPU count

Requirements

- The IP address of each controller on the array
- Port requirement 2463

Configuration

Field	Description
Comma-separated list of Array SANtricity Controller IPs	IP addresses and/or fully-qualified domain names for the array controllers

Advanced configuration

Field	Description
Inventory Poll Interval (min)	Default is 30 minutes
Performance Poll Interval up to 3600 seconds	Default is 300 seconds

Troubleshooting

Additional information on this data collector may be found from the [Support](#) page or in the [Data Collector Support Matrix](#).

NetApp E-Series REST data collector

The NetApp E-Series REST data collector gathers inventory and performance data. The collector supports firmware 7.x+ using the same configurations and reporting the same data. The REST collector monitors the encryption status of storage pools as well as the encryption status of related disks and volumes, and it provides storage node CPU Utilization as performance counters—functionality not provided in the legacy SANtricity E-Series collector.

Terminology

Cloud insight acquires the following inventory information from the NetApp E-Series, using REST. For each asset type acquired, the most common terminology used for this asset is shown. When viewing or troubleshooting this data collector, keep the following terminology in mind:

Vendor/Model Term	Cloud Insights Term
Disk	Disk
Volume Group	Disk Group
Storage Array	Storage
Controller	Storage Node
Volume Group	Storage Pool
Volume	Volume

Note: These are common terminology mappings only and might not represent every case for this data collector.

Requirements

- The IP address of each controller on the array
- This collector supports only E-Series model arrays with **native REST API capabilities**. The E-Series org ships an off-array, installable REST API distribution for older E-Series arrays: this collector does not support that scenario. Users with older arrays should continue to use Cloud Insights' [E-Series SANtricity API](#) collector.
- The "E-Series Controller IP Addresses" field supports a comma delimited string of 2 IP/hostnames; the collector will intelligently try the second IP/hostname if the first is inaccessible.
- HTTPS port: default is 8443.

Configuration

Field	Description
E-Series Controller IP Addresses	comma-separated IP addresses and/or fully-qualified domain names for the array controllers

Advanced configuration

Field	Description
Inventory Poll Interval (min)	Default is 30 minutes
Performance Poll Interval up to 3600 seconds	Default is 300 seconds

E-Series Terminology (Landing Page)

The following terms apply to objects or references that you might find on NetApp E-Series asset landing pages. Many of these terms apply to other data collectors as well.

Storage

- Model – model name of the device.
- Vendor – same Vendor name you would see if you were configuring a new datasource
- Serial number – The array serial number. On cluster architecture storage systems like NetApp Clustered Data Ontap, this serial number may be less useful than the individual "Storage Nodes" serial numbers
- IP – generally will be the IP(s) or hostname(s) as configured in the data source
- Microcode version – firmware
- Raw Capacity – base 2 summation of all the physical disks in the system, regardless of their role
- Latency – a representation of what the host facing workloads are experiencing, across both reads and writes. Ideally, Cloud Insights is sourcing this value directly, but this is often not the case. In lieu of the array offering this up, Cloud Insights is generally performing an IOPs-weighted calculation derived from the individual volumes' statistics.
- Throughput – the array's total host facing throughput. Ideally sourced directly from the array, if unavailable, Cloud Insights is summing the volumes' throughput to derive this value
- Management – this may contain a hyperlink for the management interface of the device. Created programmatically by the Cloud Insights datasource as part of inventory reporting

Storage Pool

- Storage – what storage array this pool lives on. Mandatory
- Type – a descriptive value from a list of an enumerated list of possibilities. Most commonly will be “Thin Provisioning” or “RAID Group”
- Node – if this storage array’s architecture is such that pools belong to a specific storage node, its name will be seen here as a hyperlink to its own landing page
- Uses Flash Pool – Yes/No value
- Redundancy – RAID level or protection scheme. E-Series reports “RAID 7” for DDP pools
- Capacity – the values here are the logical used, usable capacity and the logical total capacity, and the percentage used across these. These value both include E-Series “preservation” capacity, resulting both in numbers and the percentage being higher than what the E-Series own user interface may show
- Over-committed capacity – If via efficiency technologies you have allocated a sum total of volume or internal volume capacities larger than the logical capacity of the storage pool, the percentage value here will be greater than 0%.
- Snapshot – snapshot capacities used and total, if your storage pool architecture dedicates part of its capacity to segments areas exclusively for snapshots
- Utilization – a percentage value showing the highest disk busy percentage of any disk contributing capacity to this storage pool. Disk utilization does not necessarily have a strong correlation with array performance – utilization may be high due to disk rebuilds, deduplication activities, etc in the absence of host driven workloads. Also, many arrays’ replication implementations may drive disk utilization while not showing as volume workload.
- IOPS – the sum IOPs of all the disks contributing capacity to this storage pool. If disk IOPs is not available on a given platform, this value will be sourced from the sum of volume IOPs for all the volumes sitting on this storage pool
- Throughput – the sum throughput of all the disks contributing capacity to this storage pool. If disk throughput is not available on a given platform, this value will be sourced from the sum of volume throughput for all the volumes sitting on this storage pool

Storage Node

- Storage – what storage array this node is part of. Mandatory
- HA Partner – on platforms where a node will fail over to one and only one other node, it will generally be seen here
- State – health of the node. Only available when the array is healthy enough to be inventoried by a data source
- Model – model name of the node
- Version – version name of the device.
- Serial number – The node serial number
- Memory – base 2 memory if available
- Utilization – Generally a CPU utilization number, or in the case of NetApp Ontap, a controller stress index. Utilization is not currently available for NetApp E-Series
- IOPS – a number representing the host driven IOPs on this controller. Ideally sourced directly from the array, if unavailable, it will be calculated by summing all the IOPs for volumes that belong exclusively to this node.
- Latency – a number representing the typical host latency or response time on this controller. Ideally

sourced directly from the array, if unavailable, it will be calculated by performing an IOPs weighted calculation from volumes that belong exclusively to this node.

- Throughput – a number representing the host driven throughput on this controller. Ideally sourced directly from the array, if unavailable, it will be calculated by summing all the throughput for volumes that belong exclusively to this node.
- Processors – CPU count

Troubleshooting

Additional information on this data collector may be found from the [Support](#) page or in the [Data Collector Support Matrix](#).

Configuring the NetApp HCI Management server data collector

The NetApp HCI Management server data collector collects NetApp HCI Host information and requires read-only privileges on all objects within the Management server.

This data collector acquires from the **NetApp HCI Management server only**. To collect data from the storage system, you must also configure the [NetApp SolidFire](#) data collector.

Terminology

Cloud Insights acquires the following inventory information from this data collector. For each asset type acquired, the most common terminology used for the asset is shown. When viewing or troubleshooting this data collector, keep the following terminology in mind:

Vendor/Model Term	Cloud Insights Term
Virtual disk	Disk
Host	Host
Virtual machine	Virtual machine
Data store	Data store
LUN	Volume
Fibre channel port	Port

These are common terminology mappings only and might not represent every case for this data collector.

Requirements

The following information is required to configure this data collector:

- IP address of the NetApp HCI Management server
- Read-only username and password for the NetApp HCI Management server
- Read only privileges on all objects in the NetApp HCI Management server.
- SDK access on the NetApp HCI Management server – normally already set up.

- Port requirements: http-80 https-443
- Validate access:
 - Log into the NetApp HCI Management server using above username and password
 - Verify SDK enabled: telnet <vc_ip> 443

Setup and connection

Field	Description
Name	Unique name for the data collector
Acquisition unit	Name of acquisition unit

Configuration

Field	Description
NetApp HCI Storage Cluster MVIP	Management Virtual IP Address
SolidFire Management Node (mNode)	Management Node IP Address
User name	User name used to access the NetApp HCI Management server
Password	Password used to access the NetApp HCI Management server
VCenter User Name	User name for VCenter
VCenter Password	Password for VCenter

Advanced configuration

In the advanced configuration screen, check the **VM Performance** box to collect performance data. Inventory collection is enabled by default.

The following fields can be configured:

Field	Description
Inventory poll interval (min)	Default is 20
Filter VMs by	Select CLUSTER, DATACENTER, or ESX HOST
Choose 'Exclude' or 'Include' to Specify a List	Specify Whether to Include or Exclude VMs
Filter Device List	List of VMs to filter (comma separated, or semicolon separated if comma is used in the value) for Filtering by ESX_HOST, CLUSTER, and DATACENTER Only
Performance poll interval (sec)	Default is 300

Troubleshooting

Some things to try if you encounter problems with this data collector:

Inventory

Problem:	Try this:
Error: Include list to filter VMs cannot be empty	If Include List is selected, please list valid DataCenter, Cluster, or Host names to filter VMs
Error: Failed to instantiate a connection to VirtualCenter at IP	Possible solutions: * Verify credentials and IP address entered. * Try to communicate with Virtual Center using Infrastructure Client. * Try to communicate with Virtual Center using Managed Object Browser (e.g MOB).
Error: VirtualCenter at IP has non-conform certificate that JVM requires	Possible solutions: * Recommended: Re-generate certificate for Virtual Center by using stronger (e.g. 1024-bit) RSA key. * Not Recommended: Modify the JVM java.security configuration to leverage the constraint <code>jdk.certpath.disabledAlgorithms</code> to allow 512-bit RSA key. See JDK 7 update 40 release notes at " http://www.oracle.com/technetwork/java/javase/7u40-relnotes-2004172.html "

Additional information may be found from the [Support](#) page or in the [Data Collector Support Matrix](#).

NetApp SolidFire All-Flash Array data collector

The NetApp SolidFire All-Flash Array data collector supports inventory and performance collection from both iSCSI and Fibre Channel SolidFire configurations.

The SolidFire data collector utilizes the SolidFire REST API. The acquisition unit where the data collector resides needs to be able to initiate HTTPS connections to TCP port 443 on the SolidFire cluster management IP address. The data collector needs credentials capable of making REST API queries on the SolidFire cluster.

Terminology

Cloud Insights acquires the following inventory information from the NetApp SolidFire All-Flash Array data collector. For each asset type acquired by Cloud Insights, the most common terminology used for this asset is shown. When viewing or troubleshooting this data collector, keep the following terminology in mind:

Vendor/Model Term	Cloud Insights Term
Drive	Disk
Cluster	Storage
Node	Storage Node
Volume	Volume
Fibre channel port	Port
Volume Access Group, LUN Assignment	Volume Map

Vendor/Model Term	Cloud Insights Term
iSCSI Session	Volume Mask

Note: These are common terminology mappings only and might not represent every case for this data collector.

Requirements

The following are requirements for configuring this data collector:

- Management Virtual IP Address
- Read-only username and credentials
- Port 443

Configuration

Field	Description
Management Virtual IP Address (MVIP)	Management Virtual IP address of the SolidFire Cluster
User Name	Name used to log into the SolidFire cluster
Password	Password used to log into the SolidFire cluster

Advanced configuration

Field	Description
Connection Type	Choose connection type
Communication Port	Port used for NetApp API
Inventory Poll Interval (min)	Default is 20 minutes
Performance Poll Interval (sec)	Default is 300 seconds

Troubleshooting

When SolidFire reports an error it is displayed in Cloud Insights as follows:

An error message was received from a SolidFire device while trying to retrieve data. The call was <method> (<parameterString>). The error message from the device was (check the device manual): <message>

Where:

- The <method> is an HTTP method, such as GET or PUT.
- The <parameterString> is a comma separated list of parameters that were included in the REST call.
- The <message> is whatever the device returned as the error message.

Additional information on this Data Collector may be found from the [Support](#) page or in the [Data Collector Support Matrix](#).

NetApp StorageGRID data collector

The NetApp StorageGRID data collector supports inventory and performance collection from StorageGRID configurations.



StorageGRID is metered at a different Raw TB to Managed Unit rate. Every 40 TB of unformatted StorageGRID capacity is charged as 1 [Managed Unit \(MU\)](#).

Terminology

Cloud Insights acquires the following inventory information from the NetApp StorageGRID collector. For each asset type acquired, the most common terminology used for this asset is shown. When viewing or troubleshooting this data collector, keep the following terminology in mind:

Vendor/Model Term	Cloud Insights Term
StorageGRID	Storage
Node	Node
Tenant	Storage Pool
Bucket	Internal Volume

Requirements

The following are requirements for configuring this data source:

- StorageGRID Host IP Address
- A username and password for a user that has had the Metric Query and Tenant Access roles assigned
- Port 443

Configuration

Field	Description
StorageGRID Host IP Address	Management Virtual IP address of the StorageGRID appliance
User Name	Name used to log into the StorageGRID appliance
Password	Password used to log into the StorageGRID appliance

Advanced configuration

Field	Description
Inventory Poll Interval (min)	Default is 60 minutes
performance Poll Interval (sec)	Default is 900 seconds

Single Sign-On (SSO)

The [StorageGRID](#) firmware versions have corresponding API versions; 3.0 API and newer versions support single sign-on (SSO) login.

Firmware version	API version	Support single sign on (SSO)
11.1	2	No
11.2	3.0	Yes
11.5	3.3	Yes

Troubleshooting

Additional information on this Data Collector may be found from the [Support](#) page or in the [Data Collector Support Matrix](#).

Copyright information

Copyright © 2024 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.