



Vendor-specific data source reference

OnCommand Insight

NetApp
June 10, 2024

Table of Contents

- Vendor-specific data source reference 1
 - 3PAR InServ data source 1
 - Amazon AWS EC2 data source 2
 - Brocade Enterprise Fabric Connectivity Manager data source 5
 - Brocade FC Switch data source 7
 - Brocade Sphereon/Intrepid Switch data source 9
 - Cisco FC Switch Firmware (SNMP) data source 10
 - EMC Celerra data source 13
 - EMC CLARiiON (NaviCLI) data source 15
 - EMC Data Domain data source 16
 - EMC ECC StorageScope data source 18
 - Dell EMC ECS data source 18
 - EMC Isilon data source 20
 - Dell EMC PowerStore data source 21
 - EMC RecoverPoint data source 23
 - EMC Solutions Enabler with SMI-S Performance data source 24
 - EMC VNX data source 27
 - EMC VNXe data source 28
 - EMC VPLEX data source 30
 - EMC XtremIO data source 32
 - Fujitsu Eternus data source 33
 - Hitachi Content Platform (HCP) data source 34
 - HDS HiCommand Device Manager data source 35
 - Hitachi Ops Center data collector 41
 - HDS NAS (HNAS) data source 43
 - HP CommandView AE data source 44
 - HP EVA Storage data source 47
 - HPE Nimble data source 49
 - Huawei OceanStor data source 50
 - IBM Cleversafe data source 52
 - IBM DS data source 52
 - IBM PowerVM data source 54
 - IBM SVC data source 55
 - IBM Tivoli Monitoring data source 58
 - IBM TotalStorage DS4000 data source 59
 - IBM XIV data source 60
 - Infinidat InfiniBox data source 61
 - Microsoft Azure compute data source 62
 - Azure NetApp Files data source 64
 - Microsoft Hyper-V data source 64
 - NetApp Clustered Data ONTAP data source 66
 - NetApp Clustered Data ONTAP for Unified Manager data source 69
 - NetApp Data ONTAP operating in 7-Mode data source 70

NetApp E-Series data source	73
NetApp Host and VM File Systems data source	75
NetApp SolidFire data source	77
NetApp StorageGRID data source	79
OpenStack data source	80
Oracle ZFS data source	81
Pure Storage FlashArray data source	83
QLogic FC Switch data source	85
Red Hat (RHEV) data source	86
Violin Flash Memory Array data source	86
VMware vSphere data source	88

Vendor-specific data source reference

The configuration details vary depending on the vendor and model of the data source being added.

If a vendor's data source requires advanced Insight configuration instructions, such as special requirements and specific commands, that information is included in this section.

3PAR InServ data source

OnCommand Insight uses the 3PAR InServ (Firmware 2.2.2+, SSH) data source to discover inventory for HP 3PAR StoreServ storage arrays.

Terminology

OnCommand Insight acquires the following inventory information from the 3PAR InServ data source. For each asset type acquired by Insight, the most common terminology used for this asset is shown. When viewing or troubleshooting this data source, keep the following terminology in mind:

Vendor/Model Term	Insight Term
Physical Disk	Disk
Storage System	Storage
Controller Node	Storage Node
Common Provisioning Group	Storage Pool
Virtual Volume	Volume



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

Requirements

- IP address or FQDN of the InServ cluster
- For inventory, read-only user name and password to the InServ Server.
- For performance, read-write user name and password to the InServ Server.
- Port requirements: 22 (inventory collection), 5988 or 5989 (performance collection) [Note: 3PAR Performance is supported for InServ OS 3.x+]
- For performance collection confirm that SMI-S is enabled by logging into the 3PAR array via SSH.

Configuration

Field	Description
-------	-------------

Cluster IP	IP address or fully-qualified domain name of the InServ cluster
User Name	User name for the InServ Server
Password	Password used for the InServ Server
SMI-S Host IP	IP address of the SMI-S Provider Host
SMI-S User Name	User name for the SMI-S Provider Host
SMI-S Password	Password used for the SMI-S Provider Host

Advanced Configuration

Field	Description
Inventory Poll Interval (min)	Interval between inventory polls (default 40 minutes)
Exclude Devices	Comma-separated list of device IPs to exclude
SSH Process Wait Timeout (sec)	SSH process timeout (default 60 seconds)
Number of SSH Retries	Number of SSH retry attempts
SSH Banner Wait Timeout (sec)	SSH banner wait timeout (default 20 seconds)
SMI-S Port	Port used by SMI-S Provider Host
Protocol	Protocol used to connect to the SMI-S provider
SMI-S namespace	SMI-S namespace
Performance Poll Interval (sec)	Interval between performance polls (default 300 seconds)
Number of SMI-S Connection Retries	Number of SMI-S connection retry attempts

Amazon AWS EC2 data source

OnCommand Insight uses this data source to discover inventory and performance for Amazon AWS EC2.

Pre-requisites:

In order to collect data from Amazon EC2 devices, you must have the following information:

- You must have the IAM Access Key ID
- You must have the Secret Access Key for your Amazon EC2 cloud account
- You must have the "list organization" privilege
- Port 433 HTTPS
- EC2 Instances can be reported as a Virtual Machine, or (less naturally) a Host. EBS Volumes can be reported as both a VirtualDisk used by the VM, as well as a DataStore providing the Capacity for the VirtualDisk.

Access keys consist of an access key ID (for example, AKIAIOSFODNN7EXAMPLE) and a secret access key (for example, wJalrXUtnFEMI/K7MDENG/bPxrFiCYEXAMPLEKEY). You use access keys to sign programmatic requests that you make to EC@ if you use the Amazon EC2 SDKs, REST, or Query API operations. These keys are provided with your contract from Amazon.

How to configure this data source

To configure the Amazon AWS EC2 data source, you will need the AWS IAM Access Key ID and Secret Access Key for your AWS account.

Fill in the data source fields according to the tables below:

Configuration:

Field	Description
AWS Region	Choose AWS region
IAM Role	For use only when acquired on an AU in AWS. See below for more information on IAM Roles.
AWS IAM Access Key ID	Enter AWS IAM Access Key ID. Required if you do not use IAM Role.
AWS IAM Secret Access Key	Enter AWS IAM Secret Access Key. Required if you do not use IAM Role.
I understand AWS will bill me for API requests	Check this to verify your understanding that AWS bills you for API requests made by Insight polling

Advanced Configuration:

Field	Description
Include Extra Regions	Specify additional regions to include in polling.

Cross Account Role	Role for accessing resources in different AWS accounts.
Inventory Poll Interval (min)	Interval between inventory polls (default 60 minutes)
HTTP connection and socket timeout (sec)	HTTP connection timeout (default 300 seconds)
Include AWS tags	Check this to enable support for AWS tags in Insight annotations
Performance Poll Interval (sec)	Interval between performance polls (default 1800 seconds)

Mapping AWS tags to Insight annotations

The AWS EC2 data source includes an option that allows you to populate Insight annotations with tags configured on AWS. The annotations must be named exactly as the AWS tags. Insight will always populate same-named text-type annotations, and will make a "best attempt" to populate annotations of other types (number, boolean, etc). If your annotation is of a different type and the data source fails to populate it, it may be necessary to remove the annotation and re-create it as a text type.

Note that AWS is case-sensitive, while Insight is case-insensitive. So if you create an annotation named "OWNER" in Insight, and tags named "OWNER", "Owner", and "owner" in AWS, all of the AWS variations of "owner" will map to Insight's "OWNER" annotation.

Related Information:

[Managing Access Keys for IAM Users](#)

Include Extra Regions

In the AWS Data Collector **Advanced Configuration** section, you can set the **Include extra regions** field to include additional regions, separated by comma or semi-colon. By default, this field is set to **us-.***, which collects on all US AWS regions. To collect on *all* regions, set this field to **.***.

If the **Include extra regions** field is empty, the data collector will collect on assets specified in the **AWS Region** field as specified in the **Configuration** section.

Collecting from AWS Child Accounts

Insight supports collection of child accounts for AWS within a single AWS data collector. Configuration for this collection is performed in the AWS environment:

- You must configure each child account to have an AWS Role that allows the primary account ID to access EC2 details from the children account.
- Each child account must have the role name configured as the same string
- Enter this role name string into the Insight AWS Data Collector **Advanced Configuration** section, in the **Cross Account Role** field.

Best Practice: It is highly recommended to assign the AWS predefined AmazonEC2ReadOnlyAccess policy to

the ECS primary account. Also, the user configured in the data source should have at least the predefined *AWSOrganizationsReadOnlyAccess* policy assigned, in order to query AWS.

Please see the following for information on configuring your environment to allow Insight to collect from AWS child accounts:

[Tutorial: Delegate Access Across AWS Accounts Using IAM Roles](#)

[AWS Setup: Providing Access to an IAM User in Another AWS Account That You Own](#)

[Creating a Role to Delegate Permissions to an IAM User](#)

IAM Roles

When using *IAM Role* security, you must ensure that the role you create or specify has the appropriate permissions needed to access your resources.

For example, if you create an IAM role named *InstanceEc2ReadOnly*, you must set up the policy to grant EC2 read-only list access permission to all EC2 resources for this IAM role. Additionally, you must grant STS (Security Token Service) access so that this role is allowed to assume roles cross accounts.

After you create an IAM role, you can attach it when you create a new EC2 instance or any existing EC2 instance.

After you attach the IAM role *InstanceEc2ReadOnly* to an EC2 instance, you will be able to retrieve the temporary credential through instance metadata by IAM role name and use it to access AWS resources by any application running on this EC2 instance.



IAM role can be used only when the Acquisition Unit is running in an AWS instance.

Brocade Enterprise Fabric Connectivity Manager data source

OnCommand Insight uses the Brocade Enterprise Fabric Connectivity Manager (EFCM) data source to discover inventory for Brocade EFCM switches. Insight supports EFCM versions 9.5, 9.6, and 9.7.

Requirements



This data collector is not available starting with OnCommand Insight 7.3.11.

- Network address or fully-qualified domain name for the EFCM server
- EFCM version must be 9.5, 9.6, or 9.7
- IP address of the EFCM server
- Read-only username and password for the EFCM server
- Validated access to the Connectrix switch by Telnet from the Insight server, using the read-only username and password over port 51512

Configuration

Field	Description
EFC server	IP address or fully-qualified domain name of the EFC Server
User Name	User name for the switch
Password	Password used for the switch

Advanced configuration

Field	Description
Inventory Poll Interval (min)	Interval between inventory polls (default 15 minutes)
Fabric Name	Fabric name to be reported by the EFCM data source. Leave blank to report the fabric name as WWN.
Communication Port	Port used for communication with the switch
Enable Trapping	Select to enable acquisition upon receiving an SNMP trap from the device. If you select enable trapping, you must also activate SNMP.
Minimum Time Between Traps (sec)	Minimum time between acquisition attempts triggered by traps (default 15 seconds)
Inactive Zonesets	Comma-separated list of inactive Zonesets on which to perform acquisition, in addition to performing acquisition on the active zone sets
NIC to Use	Specify which network interface should be used on the RAU when reporting on SAN devices
Exclude Devices	Comma-separated list of unit names to include or exclude from polling
Use the EFCM switch nickname as the Insight switch name	Select to use the EFCM switch nickname as the Insight switch name
Performance Poll Interval (sec)	Interval between performance polls (default 300 seconds)

Brocade FC Switch data source

OnCommand Insight uses the Brocade FC Switch (SSH) data source to discover inventory for Brocade or rebranded switch devices running Factored Operating System (FOS) firmware 4.2 and later. Devices in both FC switch and Access Gateway modes are supported.

Terminology

OnCommand Insight acquires the following inventory information from the Brocade FC Switch data source. For each asset type acquired by Insight, the most common terminology used for this asset is shown. When viewing or troubleshooting this data source, keep the following terminology in mind:

Vendor/Model Term	Insight Term
Switch	Switch
Port	Port
Virtual Fabric, Physical Fabric	Fabric
Zone	Zone
Logical Switch	Logical Switch
LSAN Zone	IVR Zone



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

Requirements

- The Acquisition Unit (local or remote) will initiate connections to TCP Port 22 on Brocade switches to collect inventory data. The AU will also initiate connections to UDP port 161 for collection of performance data.
- There must be IP connectivity to all switches in the fabric. If you select the Discover all switches in the fabric check box, OCI identifies all the switches in the fabric; however, it needs IP connectivity to these additional switches to discover them.
- The same account is needed globally across all switches in the fabric. You can use PuTTY (open source terminal emulator) to confirm access.
- If the Perform license is installed, ports 161 and 162 must be open to all switches in the fabric for SNMP performance polling.
- SNMP read-only Community String

Configuration

Field	Description
Switch IP	IP address or fully-qualified domain name of the switch
User Name	User name for the switch
Password	Password used for the switch
SNMP Version	SNMP version
SNMP Community String	SNMP read-only community string used to access the switch
SNMP User Name	SNMP version protocol user name (applies only to SNMP v3)
SNMP Password	SNMP version protocol password (applies only to SNMP v3)

Advanced configuration

Field	Description
Fabric Name	Fabric name to be reported by the data source. Leave blank to report the fabric name as WWN.
Exclude Devices	Comma-separated list of device IDs to exclude from polling
Inventory Poll Interval (min)	Interval between inventory polls (default 15 minutes)
Timeout (sec)	Connection timeout (default 30 seconds)
Banner Wait Timeout (sec)	SSH banner wait timeout (default 5 seconds)
Admin Domains Active	Select if using Admin Domains
Retrieve MPR Data	Select to acquire routing data from your multiprotocol router (MPR)
Enable Trapping	Select to enable acquisition upon receiving an SNMP trap from the device. If you select enable trapping, you must also activate SNMP.
Minimum Time Between Traps (sec)	Minimum time between acquisition attempts triggered by traps (default 10 seconds)

Discover all switches in the fabric	Select to discover all switches in the fabric
Choose Favoring HBA vs. Zone Aliases	Choose whether to favor HBA or zone aliases
Performance Poll Interval (sec)	Interval between performance polls (default 300 seconds)
SNMP Auth Protocol	SNMP authentication protocol (SNMP v3 only)
SNMP Privacy Protocol	SNMP privacy protocol (SNMP v3 only)
SNMP Privacy Password	SNMP privacy password (SNMP v3 only)
SNMP Retries	Number of SNMP retry attempts
SNMP Timeout (ms)	SNMP timeout (default 5000 ms)

Brocade Sphereon/Intrepid Switch data source

OnCommand Insight uses the Brocade Sphereon/Intrepid Switch (SNMP) data source to discover inventory for Brocade Sphereon or Intrepid switches.

Requirements



This data collector not available starting with OnCommand Insight 7.3.11.

- There must be IP connectivity to all switches in the fabric. If you select the Discover all switches in the fabric check box, OCI identifies all the switches in the fabric; however, it needs IP connectivity to these additional switches to discover them.
- Read-only community string if using SNMP V1 or SNMP V2.
- HTTP access to the switch to obtain zoning information.
- Access validation by running the `snmpwalk` utility to the switch (see `<install_path>\bin\`).

Configuration

Field	Description
Sphereon Switch	IP address or fully-qualified domain name of the switch
SNMP Version	SNMP version
SNMP Community	SNMP read-only community string used to access the switch

User Name	SMI-S user name for the switch (SNMP v3 only)
Password	SMI-S password for the switch (SNMP v3 only)

Advanced configuration

Field	Description
Inventory Poll Interval (min)	Interval between inventory polls (default 15 minutes)
SNMP Auth Protocol	SNMP authentication protocol (SNMPv3 only)
SNMP Privacy Protocol	SNMP privacy protocol (SNMPv3 only)
SNMP Privacy Password	SNMP privacy password
SNMP Number of Retries	Number of SNMP retry attempts
SNMP Timeout (ms)	SNMP timeout (default 5000 ms)
Fabric Name	Fabric name to be reported by the data source. Leave blank to report the fabric name as WWN.
Enable Trapping	Select to enable acquisition upon receiving an SNMP trap from the device. If you select enable trapping, you must also activate SNMP.
Minimum Time Between Ttraps (seconds)	Minimum time between acquisition attempts triggered by traps (default 10 seconds)
Performance Poll Interval (sec)	Interval between performance polls (default 300 seconds)

Cisco FC Switch Firmware (SNMP) data source

OnCommand Insight uses the Cisco FC Switch Firmware 2.0+ (SNMP) data source to discover inventory for Cisco MDS Fibre Channel switches as well as a variety of Cisco Nexus FCoE switches on which the FC service is enabled. Additionally, you can discover many models of Cisco devices running in NPV mode with this data source.

Terminology

OnCommand Insight acquires the following inventory information from the Cisco FC Switch data source. For each asset type acquired by Insight, the most common terminology used for this asset is shown. When viewing or troubleshooting this data source, keep the following terminology in mind:

Vendor/Model Term	Insight Term
Switch	Switch
Port	Port
VSAN	Fabric
Zone	Zone
Logical Switch	Logical Switch
Name Server Entry	Name Server Entry
Inter-VSAN Routing (IVR) Zone	IVR Zone



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

Requirements

- An IP address of one switch in the fabric or individual switches
- Chassis discovery, to enable fabric discovery
- If using SNMP V2, read-only community string
- Port 161 is used to access the device
- Access validation using the `snmpwalk` utility to the switch (see `<install_path>\bin\`)

Configuration

Field	Description
Cisco Switch IP	IP address or fully-qualified domain name of the switch
SNMP Version	SNMP version v2 or later is required for performance acquisition
SNMP Community String	SNMP read-only community string used to access the switch (not applicable for SNMP v3)
User Name	User name for the switch (SNMP v3 only)
Password	Password used for the switch (SNMPv3 only)

Advanced configuration

Field	Description
Inventory Poll Interval (min)	Interval between inventory polls (default 40 minutes)
SNMP Auth Protocol	SNMP authentication protocol (SNMPv3 only)
SNMP Privacy Protocol	SNMP privacy protocol (SNMPv3 only)
SNMP Privacy Password	SNMP privacy password
SNMP Retries	Number of SNMP retry attempts
SNMP Timeout (ms)	SNMP timeout (default 5000 ms)
Enable Trapping	Select to enable trapping. If you enable trapping, you must also activate SNMP notifications.
Minimum Time Between Traps (sec)	Minimum time between acquisition attempts triggered by traps (default 10 seconds)
Discover All Fabric Switches	Select to discover all switches in the fabric
Exclude Devices	Comma-separated list of device IPs to exclude from polling
Include Devices	Comma-separated list of device IPs to include in polling
Check Device Type	Select to accept only those devices that explicitly advertise themselves as Cisco devices

Primary Alias Type	<p>Provide a first preference for resolution of the alias. Choose from the following:</p> <ul style="list-style-type: none"> • Device Alias <p>This is a user-friendly name for a port WWN (pWWN) that can be used in all configuration commands, as required. All switches in the Cisco MDS 9000 Family support Distributed Device Alias Services (device aliases).</p> <ul style="list-style-type: none"> • None <p>Do not report any alias</p> <ul style="list-style-type: none"> • Port Description <p>A description to help identify the port in a list of ports</p> <ul style="list-style-type: none"> • Zone Alias (all) <p>A user-friendly name for a port that can be used only for zoning configuration</p> <ul style="list-style-type: none"> • Zone Alias (only active) <p>A user-friendly name for a port that can be used only for the active configuration. This is the default.</p>
Secondary Alias Type	Provide a second preference for resolution of the alias
Tertiary Alias Type	Provide a third preference for resolution of the alias
Enable SANTap Proxy Mode Support	Select if your Cisco switch is using SANTap in proxy mode. If you are using EMC RecoverPoint, then you are probably using SANTap.
Performance Poll Interval (sec)	Interval between performance polls (default 300 seconds)

EMC Celerra data source

The Celerra (SSH) data source collects inventory information from Celerra storage. For configuration, this data source requires the IP address of the storage processors and a *read-only* user name and password.

Terminology

OnCommand Insight acquires the following inventory information from the EMC Celerra data source. For each asset type acquired by Insight, the most common terminology used for this asset is shown. When viewing or troubleshooting this data source, keep the following terminology in mind:

Vendor/Model Term	Insight Term
Celerra Network Server	Storage
Celerra Meta Volume / Celerra Storage Pool	Storage Pool
File System	Internal Volume
Data Mover	Controller
File System mounted on a data Mover	File Share
CIFS and NFS Exports	Share
Disk Volume	Backend LUN



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

Requirements

- The IP address of the storage processor
- Read-only user name and password
- SSH port 22

Configuration

Field	Description
Address of Celerra	IP address or fully-qualified domain name of the Celerra device
User Name	Name used to log in to the Celerra device
Password	Password used to log in to the Celerra device

Advanced configuration

Field	Description
-------	-------------

Inventory Poll Interval (minutes)	Interval between inventory polls (default 20 minutes)
SSH Process Wait Timeout (sec)	SSH process timeout (default 600 seconds)
Number of Retries	Number of inventory retry attempts
SSH Banner Wait Timeout (sec)	SSH banner wait timeout (default 20 seconds)

EMC CLARiiON (NaviCLI) data source

Before configuring this data source, make sure that the EMC Navisphere CLI is installed on the target device and on the Insight server. The Navisphere CLI version must match the firmware version on the controller. For performance data collection, statistics logging must be turned on.

Navisphere Command Line Interface syntax

```
naviseccli.exe -h <IP address> -user <user> -password <password> -scope <scope,use 0 for global scope> -port <use 443 by default> command
```

Terminology

OnCommand Insight acquires the following inventory information from the EMC CLARiiON data source. For each asset type acquired by Insight, the most common terminology used for this asset is shown. When viewing or troubleshooting this data source, keep the following terminology in mind:

Vendor/Model Term	Insight Term
Disk	Disk
Storage	Storage
Storage Processor	Storage Node
Thin Pool, RAID Group	Storage Pool
LUN	Volume



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

Requirements

- An IP address of each CLARiiON storage processor

- Read-only Navisphere username and password to the CLARiiON arrays
- NaviCLI must be installed on the Insight server/RAU
- Access validation: Run NaviCLI from the Insight server to each array using the above username and password.
- NaviCLI version should correspond with the newest FLARE code on your array
- For performance, statistics logging must be turned on.
- Port requirements: 80, 443

Configuration

Field	Description
CLARiiON storage	IP address or fully-qualified domain name of the CLARiiON Storage
User Name	Name used to log into the CLARiiON storage device.
Password	Password used to log into the CLARiiON storage device.
CLI Path to navicli.exe path or naviseccli.exe path	Full path to the <code>navicli.exe</code> OR <code>naviseccli.exe</code> executable

Advanced configuration

Field	Description
Inventory Poll Interval (min)	Interval between inventory polls (default 40 minutes)
Use Secure Client (naviseccli)	Select to use secure client (navseccli)
Scope	The secure client scope. The default is Global.
CLARiiON CLI Port	Port used for CLARiiON CLI
Inventory External Process Timeout (sec)	External process timeout (default 1800 seconds)
Performance Poll Interval (sec)	Interval between performance polls (default 300 seconds)
Performance External process timeout (sec)	External process timeout (default 1800 seconds)

EMC Data Domain data source

This data source collects storage and configuration information from EMC Data Domain

deduplication storage systems. To add the data source, you must use specific configuration instructions and commands and be aware of data source requirements and usage recommendations.

Terminology

OnCommand Insight acquires the following inventory information from the EMC Data Domain data source. For each asset type acquired by Insight, the most common terminology used for this asset is shown. When viewing or troubleshooting this data source, keep the following terminology in mind:

Vendor/Model Term	Insight Term
Disk	Disk
Array	Storage
Port	Port
Filesys	Internal Volume
Mtree	QTree
Quota	Quota
NFS and CIFS share	FileShare



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

Requirements

- IP address of the Data Domain device
- Read-only user name and password to the Data Domain storage
- SSH port 22

Configuration

Field	Description
IP address	The IP address or fully-qualified domain name of the Data Domain storage array
User name	The user name for the Data Domain storage array
Password	The password for the Data Domain storage array

Advanced configuration

Field	Description
Inventory Poll Interval (min)	Interval between inventory polls (default 20 minutes)
SSH Process Wait Timeout (sec)	SSH process timeout (default 180 seconds)
SSH Port	SSH service port

EMC ECC StorageScope data source

The EMC ECC StorageScope device has three types of data sources: 5.x, 6.0, and 6.1.

Configuration



This data collector is no longer available starting with OnCommand Insight 7.3.11.

Field	Description
ECC server	IP address or fully-qualified domain name of the ECC Server
User Name	User name for the ECC server
Password	Password for the ECC server

Advanced configuration

Field	Description
ECC Port	Port used for the ECC server
Inventory Poll Interval (min)	Interval between inventory polls (default 30 minutes)
Protocol to Connect to Database	Protocol Used to Connect to the Database
Query File System Information	Select to retrieve details for WWN Aliases and File Systems.

Dell EMC ECS data source

This data collector acquires inventory and performance data from EMC ECS storage systems. For configuration, the data collector requires an IP address of the ECS server and an administrative level domain account.

Terminology

OnCommand Insight acquires the following inventory information from the EMC ECS data source. For each asset type acquired by Insight, the most common terminology used for this asset is shown. When viewing or troubleshooting this data source, keep the following terminology in mind:

Vendor/Model Term	Insight Term
Cluser	Storage
Tenant	Storage Pool
Bucket	Internal Volume
Disk	Disk



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

Requirements

- An IP address of the ECS Management Console
- Administrative level domain account for the ECS system
- Port 443 (HTTPS). Requires outbound connectivity to TCP port 443 on the ECS system.
- For performance, read-only username and password for ssh/scp access.
- For performance, port 22 is required.

Configuration

Field	Description
ECS Host	IP addresses or fully-qualified domain names of the ECS system
ECS Host Port	Port used for communication with ECS Host
ECS Vendor ID	Vendor ID for ECS
Password	Password used for ECS

Advanced configuration

Field	Description
Inventory Poll Interval (minutes)	Interval between inventory polls. The default is 360 minutes.

EMC Isilon data source

The Isilon SSH data source collects inventory and performance from EMC Isilon scale-out NAS storage.

Terminology

OnCommand Insight acquires the following inventory information from the EMC Isilon data source. For each asset type acquired by Insight, the most common terminology used for this asset is shown. When viewing or troubleshooting this data source, keep the following terminology in mind:

Vendor/Model Term	Insight Term
Drive	Disk
Cluster	Storage
Node	Storage Node
File System	Internal Volume



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

Requirements

- Administrator permissions to the Isilon storage
- Validated access by using `telnet` to port 22

Configuration

Field	Description
IP address	The IP address or fully-qualified domain name of the Isilon cluster
User name	The user name for the Isilon cluster
Password	The password for the Isilon cluster

Advanced configuration

Field	Description
Inventory Poll Interval (min)	Interval between inventory polls (default 20 minutes)

Performance Poll Interval (sec)	Interval between performance polls (default 300 seconds)
SSH Process Wait Timeout	SSH process timeout (default 60 seconds)
SSH Port	SSH service port

Running CLI Commands

Starting with OnCommand Insight version 7.3.11 and Service Pack 9, the EMC Isilon data source contains an enhancement that will result in Insight running more CLI commands. If you are using a non-root user within your data source, you will likely have configured a "sudoers" file to grant that user account the ability to run specific CLI commands via SSH.

In order for Insight to understand EMC's "Access Zones" feature, Insight will now additionally run the following new CLI commands:

- `sudo isi zone zones list --format json -verbose`
- `sudo isi zone zones list`

Insight parses the output of these commands, and runs more instances of existing commands, to obtain the logical configuration of objects like qtrees, quotas and NAS shares/exports that reside in non-default Access Zones. Insight now reports those objects for non-default Access Zones as the result of this enhancement. As Insight obtains that data by running existing commands (with different options) no sudoers file change is required in order for those to work; it is only with the introduction of the new commands above that the change is required.

Please update your sudoers file to allow your Insight service account to run those commands before upgrading to this Insight release. Failure to do so will result in your Isilon data sources failing.

"File System" statistics

Beginning with OnCommand Insight 7.3.12, the EMC Isilon data collector introduces "File System" statistics on the node object for EMC Isilon. The existing node statistics reported by OnCommand Insight are "disk" based - i.e, for IOPs and throughput of a storage node, what are the disks in this node doing in aggregate? But for workloads where reads are cached in memory and/or compression is in use, the file system workload may be substantially higher than what actually hits the disks - a data set that compresses 5:1 could therefore have a "File System Read Throughput" value 5x the storage node Read Throughput, as the latter measures the reads off of disk, which expand 5x when the node uncompresses the data to service the client's read request.

Dell EMC PowerStore data source

The Dell EMC PowerStore data collector gathers inventory information from Dell EMC PowerStore storage. For configuration, the data collector requires the IP address of the storage processors and a read-only user name and password.

Terminology

OnCommand Insight acquires the following inventory information from the EMC Data Domain data source. For each asset type acquired by Insight, the most common terminology used for this asset is shown. When viewing

or troubleshooting this data source, keep the following terminology in mind:

Vendor/Model Term	Insight Term
host	host
host_volume_mapping	host_volume_mapping
hardware (it has Drives under "extra_details" object): Drives	Disk
Appliance	StoragePool
Cluster	Storage Array
Node	StorageNode
fc_port	Port
volume	Volume
InternalVolume	file_system
Filesys	Internal Volume
Mtree	QTree
Quota	Quota
NFS and CIFS share	FileShare



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

Requirements

- IP address or fully-qualified domain name of storage processor
- Read-only user name and password

Parent Serial Number explained

Traditionally Insight is capable of reporting the storage array serial number, or the individual storage node serial numbers. However, some storage array architectures do not cleanly align to this. A PowerStore cluster can be comprised of 1-4 appliances, and each appliance has 2 nodes. If the appliance itself has a serial number, that serial number is neither the serial number for the cluster nor the nodes.

The attribute "Parent Serial Number" on the storage node object is populated appropriately for Dell/EMC PowerStore arrays when the individual nodes sit inside an intermediate appliance/enclosure that is just part of

a larger cluster.

Configuration

Field	Description
PowerStore gateway(s)	IP addresses or fully-qualified domain names of PowerStore storage
User Name	User name for PowerStore
Password	Password used for PowerStore

Advanced configuration

Field	Description
HTTPS Port	Default is 443
Inventory Poll Interval (minutes)	Interval between inventory polls. The default is 60 minutes.

OnCommand Insight's PowerStore performance collection makes use of PowerStore's 5-minute granularity source data. As such, Insight polls for that data every five minutes, and this is not configurable.

EMC RecoverPoint data source

The EMC RecoverPoint data source collects inventory information from EMC recoverPoint storage. For configuration, the data source requires the IP address of the storage processors and a *read-only* user name and password.

The EMC RecoverPoint data source gathers the volume-to-volume replication relationships that RecoverPoint coordinates across other storage arrays. OnCommand Insight shows a storage array for each RecoverPoint cluster, and collects inventory data for nodes and storage ports on that cluster. No storage pool or volume data is collected.

Requirements

- IP address or fully-qualified domain name of storage processor
- Read-only user name and password
- REST API access via port 443
- SSH access via PuTTY

Configuration

Field	Description
-------	-------------

Address of RecoverPoint	IP address or fully-qualified domain name of RecoverPoint cluster
User Name	User name for the RecoverPoint cluster
Password	Password for the RecoverPoint cluster

Advanced configuration

Field	Description
TCP Port	TCP Port used to connect to Recoverpoint cluster
Inventory Poll Interval (minutes)	Interval between inventory polls (default 20 minutes)
Excluded Clusters	Comma-separated list of cluster IDs or names to exclude when polling

EMC Solutions Enabler with SMI-S Performance data source

OnCommand Insight discovers Symmetrix storage arrays by using Solutions Enabler `symcli` commands in conjunction with an existing Solutions Enabler server in your environment. The existing Solutions Enabler server has connectivity to the Symmetrix storage array through access to gatekeeper volumes. Administrator permissions are required to access this device.

Terminology

OnCommand Insight acquires the following inventory information from the EMC Solutions Enabler data source. For each asset type acquired by Insight, the most common terminology used for this asset is shown. When viewing or troubleshooting this data source, keep the following terminology in mind:

Vendor/Model Term	Insight Term
Disk	Disk
Disk Group	Disk Group
Storage Array	Storage
Director	Storage Node
Device Pool, Storage Resource Pool (SRP)	Storage Pool
Device, TDev	Volume



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

Requirements

Before configuring this data source, you should ensure that the OnCommand Insight server has TCP connectivity to port 2707 on the existing Solutions Enabler server. OnCommand Insight discovers all the Symmetrix arrays that are “Local” to this server, as seen in “symcfg list” output from that server.

- The EMC Solutions Enabler (CLI) with SMI-S provider application must be installed and the version must match or be earlier than the version running on the Solutions Enabler Server.
- A properly configured `{installdir}\EMC\SYMAPI\config\netcnfg` file is required. This file defines service names for Solutions Enabler servers, as well as the access method (SECURE / NOSECURE /ANY).
- If you require read/write latency at the storage node level, the SMI-S Provider must communicate with a running instance of the UNISPHERE for VMAX application.
- Administrator permissions on the Solutions Enabler (SE) Server
- Read-only user name and password to the SE software
- Solutions Enabler Server 6.5X requirements:
 - SMI-S provider 3.3.1 for SMIS-S V1.2 installed
 - After install, run `\Program Files\EMC\SYMCLI\bin>stordaemon start storsrvd`
- The UNISPHERE for VMAX application must be running and collecting statistics for the Symmetrix VMAX storage arrays that are managed by the SMI-S Provider installation
- Access validation: Verify that the SMI-S provider is running: `telnet <se_server\> 5988`

Configuration



If SMI-S user authentication is not enabled, the default values in the OnCommand Insight data source are ignored.


Having symauth enabled on Symmetrix arrays might inhibit the ability of OnCommand Insight to discover them. OnCommand Insight Acquisition runs as the SYSTEM user on the OnCommand Insight / Remote Acquisition Unit server that is communicating with the Solutions Enabler server. If `hostname\SYSTEM` does not have symauth privileges, OnCommand Insight fails to discover the array.

The EMC Solutions Enabler Symmetrix CLI data source includes support for device configuration for thin provisioning and Symmetrix Remote Data Facility (SRDF).

Definitions are supplied for Fibre Channel and Switch Performance packages.

Field	Description
Service Name	Service name as specified in netcnfg file
Full path to CLI	Full path to the Symmetrix CLI

Advanced configuration

Field	Description
Inventory Poll Interval (min)	Interval between inventory polls (default 40 minutes)
Choose 'Exclude' or 'Include' to specify a list	Specify whether to include or exclude the array list below when collecting data
Inventory Exclude Devices	Comma-separated list of device IDs to include or exclude
Connection Caching	<p>Choose connection caching method:</p> <ul style="list-style-type: none"> • LOCAL means that the OnCommand Insight Acquisition service is running on the Solutions Enabler server, which has Fibre Channel connectivity to the Symmetrix arrays you seek to discover, and has access to gatekeeper volumes. This might be seen in some Remote Acquisition Unit (RAU) configurations. • REMOTE_CACHED is the default and should be used in most cases. This uses the NETCNFG file settings to connect using IP to the Solutions Enabler server, which must have Fibre Channel connectivity to the Symmetrix arrays you seek to discover, and has access to Gatekeeper volumes. • In the event that REMOTE_CACHED options make CLI commands fail, use the REMOTE option. Keep in mind that it will slow down the acquisition process (possibly to hours or even days in extreme cases). The NETCNFG file settings are still used for an IP connection to the Solutions Enabler server that has Fibre Channel connectivity to the Symmetrix arrays being discovered. <div style="border: 1px solid #ccc; padding: 10px; margin-top: 10px;">  <p>This setting does not change OnCommand Insight behavior with respect to the arrays listed as REMOTE by the "symcfg list" output. OnCommand Insight gathers data only on devices shown as LOCAL by this command.</p> </div>
CLI Timeout (sec)	CLI process timeout (default 7200 seconds)
SMI-S Host IP	IP address of the SMI-S Provider Host
SMI-S Port	Port used by SMI-S Provider Host

Protocol	Protocol used to connect to the SMI-S provider
SMI-S Namespace	Interoperability namespace that the SMI-S provider is configured to use
SMI-S User Name	User name for the SMI-S Provider Host
SMI-S Password	User name for the SMI-S Provider Host
Performance Polling Interval (sec)	Interval between performance polls (default 1000 seconds)
Performance Filter Type	Specify whether to include or exclude the array list below when collecting performance data
Performance Filter Device List	Comma-separated list of device IDs to include or exclude
RPO Polling Interval (sec)	Interval between RPO polls (default 300 seconds)

EMC VNX data source

For configuration, the EMC VNX (SSH) data source requires the IP address of the Control Station and a *read-only* username and password.

Configuration

Field	Description
VNX IP	IP address or fully-qualified domain name of the VNX Control Station
VNX User Name	User name for the VNX Control Station
VNX Password	Password for the VNX Control Station

Requirements

- An IP address of the Control Station
- Read-only username and password.
- Access validation: Verify SSH access via PuTTY.

Advanced configuration

Field	Description
-------	-------------

Inventory Poll Interval (min)	Interval between inventory polls (default 40 minutes)
VNX SSH Process Wait Timeout (sec)	VNX SSH process timeout (default 600 seconds)
Celerra Command Retry Attempts	Number of Celerra command retry attempts
CLARiiON External Process Timeout for Inventory (sec)	CLARiiON external process timeout for inventory(default 1800 seconds)
Performance Poll Interval (sec)	Interval between performance polls (default 300 seconds)
CLARiiON External Process Timeout for Performance (sec)	CLARiiON external process timeout for performance (default 1800 seconds)

EMC VNXe data source

The EMC VNXe data source provides inventory support for EMC VNXe and Unity unified storage arrays.

This data source is CLI-based and requires that you install the Unisphere for VNXe CLI (uemcli.exe) on the acquisition unit that the VNXe data source resides on. uemcli.exe uses HTTPS as the transport protocol, so the acquisition unit must be able to initiate HTTPS connections to the VNXe/Unity arrays. You must have at least a read-only user for use by the data source.

Terminology

OnCommand Insight acquires the following inventory information from the EMC VNXe data source. For each asset type acquired by Insight, the most common terminology used for this asset is shown. When viewing or troubleshooting this data source, keep the following terminology in mind:

Vendor/Model Term	Insight Term
Disk	Disk
Storage Array	Storage
Processor	Storage Node
Storage Pool	Storage Pool
General iSCSI Block info, VMWare VMFS	Volume
Shared Folder	Internal Volume

CIFS Share, NFS Share, Share from VMWare NFS datastore	Share
Replication Remote System	Synchronization
iSCSI Node	iSCSI Target Node
iSCSI Initiator	iSCSI Target Initiator



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

Requirements

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

- The VNXe data collector is CLI based; you must install the Unisphere for VNXe CLI, (uemcli.exe) onto the acquisition unit where your VNXe data collector resides.
- uemcli.exe uses HTTPS as the transport protocol, so the acquisition unit will need to be able to initiate HTTPS connections to the VNXe.
- You must have at least a read-only user for use by the data source.
- IP address of the managing Solutions enabler server.
- HTTPS on Port 443 is required
- The EMC VNXe data collector provides NAS and iSCSI support for inventory; fibre channel volumes will be discovered, but Insight does not report on FC mapping, masking, or storage ports.

Configuration

Field	Description
VNXe Storage	IP address or fully-qualified domain name of the VNXe device
User Name	User name for the VNXe device
Password	Password for the VNXe device
Full path to the uemcli executable	Full path to the uemcli.exe executable

Advanced configuration

Field	Description
Inventory Poll Interval (min)	Interval between inventory polls (default 40 minutes)

VNXe CLI Port	Port used for the VNXe CLI
Inventory External Process Timeout (sec)	External process timeout (default 1800 seconds)

EMC VPLEX data source

For configuration, this data source requires an IP address of the VPLEX server and an administrative level domain account.

Terminology

OnCommand Insight acquires the following inventory information from the EMC VPLEX data source. For each asset type acquired by Insight, the most common terminology used for this asset is shown. When viewing or troubleshooting this data source, keep the following terminology in mind:

Vendor/Model Term	Insight Term
Cluster	Storage
Engine	Storage Node
Device, System Extend	Backend Storage Pool
Virtual Volume	Volume
Front-End Port, Back-End Port	Port
Distributed Device	Storage Synchronization
Storage View	Volume Map, Volume Mask
Storage Volume	Backend LUN
ITLs	Backend Path



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

Requirements

- An IP address of the VPLEX server
- Administrative level domain account for the VPLEX server
- Port 443 (HTTPS). Requires outbound connectivity to TCP port 443 on the VPLEX management station.
- For performance, read-only username and password for ssh/scp access.

- For performance, port 22 is required.
- Validate access: Verify by using `telnet` to port 443. For a port other than the default port, with any browser use `HTTPS://<ip>:<port>`

Configuration

Field	Description
IP address of VPLEX Management Console	IP address or fully-qualified domain name of the VPLEX Management Console
User Name	User name for VPLEX CLI
Password	Password used for VPLEX CLI
Performance Remote IP Address of VPLEX Management Console	Performance Remote IP address of the VPLEX Management Console
Performance Remote User Name	Performance Remote user name of VPLEX Management Console
Performance Remote Password	Performance Remote Password of VPLEX Management Console

Advanced configuration

Field	Description
Communication Port	Port used for VPLEX CLI
Inventory Poll Interval (min)	Interval between inventory polls (default 20 minutes)
Connection timeout (sec)	Connection timeout (default 60 seconds)
Number of Retries	Number of inventory retry attempts
Performance Poll Interval (sec)	Interval between performance polls (default 600 seconds)
Performance SSH Process Wait Timeout (sec)	SSH process timeout (default 600 seconds)
SSH Banner Wait Timeout (sec)	SSH banner wait timeout (default 20 seconds)
Number of Retries	Number of performance retry attempts

EMC XtremIO data source

To configure the EMC XtremIO (HTTP) data source, you must have the XtremIO Management Server (XMS) Host address and an account with administrator privileges.

Terminology

OnCommand Insight acquires the following inventory information from the EMC XtremIO data source. For each asset type acquired by Insight, the most common terminology used for this asset is shown. When viewing or troubleshooting this data source, keep the following terminology in mind:

Vendor/Model Term	Insight Term
Disk (SSD)	Disk
Cluster	Storage
Controller	Storage Node
Volume	Volume
LUN Map	Volume Map
Initiator, Target	Volume Mask



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

Requirements

- An IP address of each XtremIO Management Server
- An account with Administrator privileges
- Access to port 443 (HTTPS)

Configuration

Field	Description
XMS Host	IP address or fully-qualified domain name of the XtremIO Management Server
User name	User name for the XtremIO Management Server
Password	Password for the XtremIO Management Server

Advanced configuration

Field	Description
TCP port	TCP Port used to connect to XTremIO Management Server (default 443)
Inventory poll interval (min)	Interval between inventory polls (default 60 minutes)
Connection timeout (sec)	Connection timeout (default 60 seconds)
Performance poll interval(sec)	Interval between performance polls (default 300 seconds)

Fujitsu Eternus data source

The Fujitsu Eternus data source requires the IP address of the storage. It cannot be comma delimited.

Terminology

OnCommand Insight acquires the following inventory information from the Fujitsu Eternus data source. For each asset type acquired by Insight, the most common terminology used for this asset is shown. When viewing or troubleshooting this data source, keep the following terminology in mind:

Vendor/Model Term	Insight Term
Disk	Disk
Storage	Storage
Thin Pool, Flexible Tier Pool, Raid Group	Storage Pool
Standard Volume, Snap Data Volume (SDV), Snap Data Pool Volume (SDPV), Thin Provisioning Volume (TPV)	Volume



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

Requirements

- An IP address of the Eternus storage, which cannot be comma delimited
- SSH Administration-level user name and password

- Port 22
- Ensure that the page scroll is disabled. (clienv-show-more-scroll disable)

Configuration

Field	Description
IP Address of Eternus Storage	IP address of the Eternus storage
User Name	User name for Eternus storage
Password	Password used for the sternus

Advanced configuration

Field	Description
Inventory Poll Interval (min)	Interval between inventory polls (default 20 minutes)
SSH Process Wait Timeout (sec)	SSH process timeout (default 600 seconds)

Hitachi Content Platform (HCP) data source

This data collector supports the Hitachi Content Platform (HCP) using the HCP Management API.

Terminology

OnCommand Insight acquires the following inventory information from the HCP data source. For each asset type acquired by Insight, the most common terminology used for this asset is shown. When viewing or troubleshooting this data source, keep the following terminology in mind:

Vendor/Model Term	Insight Term
HCP Cluster	Storage
Tenant	Storage Pool
Namespace	Internal Volume
Node	Node



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

Inventory Requirements

- IP address of the HCP server
- Read-only user name and password for the HCP software and peer privileges

Configuration

Field	Description
HCP Host	IP address or fully-qualified domain name of the HCP host
HCP Port	Default is 9090
HCP user ID	User name for the HCP host
HCP Password	Password used for the HCP host
HCP Authentication Type	Choose HCP_LOCAL or ACTIVE_DIRECTORY

Advanced configuration

Field	Description
Inventory Poll Interval (min)	Interval between inventory polls (default 60 minutes)
Performance Polling Interval (sec)	Interval between performance polls (default 900 seconds)

HDS HiCommand Device Manager data source

The HDS HiCommand and HiCommand Lite data sources support the HiCommand Device Manager server. OnCommand Insight communicates with the HiCommand Device Manager server using the standard HiCommand API.

Terminology

OnCommand Insight acquires the following inventory information from the HDS HiCommand and HiCommand Lite data sources. For each asset type acquired by Insight, the most common terminology used for this asset is shown. When viewing or troubleshooting this data source, keep the following terminology in mind:

Vendor/Model Term	Insight Term
PDEV	Disk
Journal Pool	Disk Group

Storage Array	Storage
Port Controller	Storage Node
Array Group, DP Pool	Storage Pool
Logical Unit, LDEV	Volume



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

Inventory Requirements

- IP address of the HiCommand Device Manager server
- Read-only user name and password for the HiCommand Device Manager software and peer privileges
- Port requirements: 2001 (http) or 2443 (https)
- Validate access:
 - Log in to the HiCommand Device Manager software using peer user name and password.
 - Verify access to the HiCommand Device Manager API: `telnet <HiCommand Device_Manager_server_ip> 2001`

Performance Requirements

- HDS USP, USP V, and VSP performance
 - Performance Monitor must be licensed.
 - Monitoring switch must be enabled.
 - The Export Tool (`Export.exe`) must be copied to the OnCommand Insight Server.
 - The Export Tool version must match the microcode version of the target array.
- HDS AMS performance
 - Performance Monitor needs to be licensed.
 - The Storage Navigator Modular 2 (SNM2) CLI utility needs to be installed on the OnCommand Insight Server.
 - You must register all AMS, WMS, SMS storage arrays whose performance needs to be acquired by OnCommand Insight by using the following command:


```
auunitaddauto.exe -ip<IP address of Controller0>IP address of Controller1>
```
 - You must ensure that all the arrays that you registered are listed in the output of this command: `auunitref.exe`.

Configuration

Field	Description
-------	-------------

HiCommand Server	IP address or fully-qualified domain name of the HiCommand Device Manager server
User Name	User name for the HiCommand Device Manager server.
Password	Password used for the HiCommand Device Manager server.
Devices - VSP G1000 (R800), VSP (R700), HUS VM (HM700) and USP storages	<p>Device list for VSP G1000 (R800), VSP (R700), HUS VM (HM700) and USP storages. Each storage requires:</p> <ul style="list-style-type: none"> • Array's IP: IP address of the storage • User Name: User name for the storage • Password: Password for the storage • Folder Containing Export Utility JAR Files: The folder containing the Export utility <code>.jar</code> files
SNM2Devices - WMS/SMS/AMS Storages	<p>Device list for WMS/SMS/AMS storages. Each storage requires:</p> <ul style="list-style-type: none"> • Array's IP: IP address of the storage • Storage Navigator CLI Path: SNM2 CLI path • Account Authentication Valid: Select to choose valid account authentication • User Name: User name for the storage • Password: Password for the storage
Choose Tuning Manager for Performance	Choose Tuning Manager for performance and override other performance options
Tuning Manager Host	IP address or fully-qualified domain name of tuning manager
Tuning Manager Port	Port used for Tuning Manager
Tuning Manager Username	User name for Tuning Manager
Tuning Manager Password	password for Tuning Manager



In HDS USP, USP V, and VSP, any disk can belong to more than one array group.

Advanced configuration

Field	Description
HiCommand Server Port	Port used for the HiCommand Device Manager
HTTPs Enabled	Select to enable HTTPs
Inventory Poll Interval (min)	Interval between inventory polls (default 40 minutes)
Choose 'Exclude' or 'Include' to specify a list	Specify whether to include or exclude the array list below when collecting data
Exclude or Include Devices	Comma-separated list of device ID's or array names to include or exclude
Query Host Manager	Select to query host manager
HTTP Timeout (sec)	HTTP connection timeout (default 60 seconds)
Performance Polling Interval (sec)	Interval between performance polls (default 300 seconds)
Export timeout in seconds	Export utility timeout (default 300 seconds)

Hitachi Ops Center data collector

This data collector uses Hitachi Ops Center's integrated suite of applications to access inventory and performance data of multiple storage devices. For inventory and capacity discovery, your Ops Center installation must include both the "Common Services" and "Administrator" components. For performance collection, you must additionally have "Analyzer" deployed.

Terminology

OnCommand Insight acquires the following inventory information from this 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	OnCommand Insight Term
Storage Systems	Storage
Volume	Volume
Parity Groups	Storage Pool(RAID), Disk Groups
Disk	Disk
Storage Pool	Storage Pool(Thin, SNAP)

Vendor/Model Term	OnCommand Insight Term
External Parity Groups	Storage Pool(Backend), Disk Groups
Port	Storage Node → Controller Node → Port
Host Groups	Volume Mapping and Masking
Volume Pairs	Storage Synchronization

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

Inventory Requirements

You must have the following in order to collect inventory data:

- IP address or hostname of the Ops Center server hosting the "Common Services" component
- Root/sysadmin user account and password that exist on all servers hosting Ops Center components. HDS did not implement REST API support for usage by LDAP/SSO users until Ops Center 10.8+

Performance requirements

The following requirements must be met in order to collect performance data:

- The HDS Ops Center "Analyzer" module must be installed
- Storage arrays must be feeding the Ops Center "Analyzer" module

Configuration

Field	Description
Hitachi Ops Center IP Address	IP address or fully-qualified domain name of the Ops Center server hosting the "Common Services" component
User Name	User name for the Ops Center server.
Password	Password used for the Ops Center server.

Advanced configuration

Field	Description
Connection Type	HTTPS (port 443) is the default
Override TCP Port	Specify the port to use if not the default
Inventory Poll Interval (min)	Interval between inventory polls. The default is 40.
Choose 'Exclude' or 'Include' to specify a list	Specify whether to include or exclude the array list below when collecting data.
Filter device List	Comma-separated list of device serial numbers to include or exclude

Performance Poll Interval (sec)	Interval between performance polls. The default is 300.
---------------------------------	---

HDS Storage

Terms applying to objects or references that you might find on HDS storage asset landing pages.

HDS Storage Terminology

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

- Name — comes directly from HDS HiCommand Device Manager’s “name” attribute via the GetStorageArray XML API call
- Model - comes directly from HDS HiCommand Device Manager’s “arrayType” attribute via the GetStorageArray XML API call
- Vendor — HDS
- Family - comes directly from HDS HiCommand Device Manager’s “arrayFamily” attribute via the GetStorageArray XML API call
- IP — this is the management IP address of the array, not an exhaustive list of all IP addresses on the array
- Raw Capacity — a base2 value representing the sum of the total capacity of all disks in this system, regardless of disk role.

HDS Storage Pool

Terms applying to objects or references that you might find on HDS storage pool asset landing pages.

HDS Storage Pool Terminology

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

- Type: The value here will be one of:
 - RESERVED — if this pool is dedicated for purposes other than data volumes, i.e, journaling, snapshots
 - Thin Provisioning — if this is a HDP pool
 - Raid Group — you will not likely see these for a few reasons:

OCI takes a strong stance to avoid double counting capacity at all costs. On HDS, one typically needs to build Raid Groups from disks, create pool volumes on those Raid Groups, and construct pools (often HDP, but could be special purpose) from those pool volumes. If OCI reported both the underlying Raid Groups as is, as well as the Pools, the sum of their raw capacity would vastly exceed the sum of the disks.

Instead, OCI’s HDS HiCommand data collector arbitrarily shrinks the size of Raid Groups by the capacity of pool volumes. This may result in OCI not reporting the Raid Group at all. Additionally, any resulting Raid Groups are flagged in a way such that they are not visible in the OCI WebUI, but they do flow into the OCI Data Warehouse (DWH). The purpose of these decisions is to avoid UI clutter for

things that most users do not care about — if your HDS array has Raid Groups with 50MB free, you probably cannot use that free space for any meaningful outcome.

- Node - N/A, as HDS pools are not tied to any one specific node
- Redundancy - the RAID level of the pool. Possibly multiple values for a HDP pool comprised of multiple RAID types
- Capacity % - the percent used of the pool for data usage, with the used GB and total logical GB size of the pool
- Over-committed Capacity - a derived value, stating “the logical capacity of this pool is oversubscribed by this percentage by virtue of the sum of the logical volumes exceeding the logical capacity of the pool by this percentage”
- Snapshot - shows the capacity reserved for snapshot usage on this pool

HDS Storage Node

Terms applying to objects or references that you might find on HDS storage node asset landing pages.

HDS Storage Node Terminology

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

- Name — The name of the Front-end director (FED) or Channel Adapter on monolithic arrays, or the name of the controller on a modular array. A given HDS array will have 2 or more Storage Nodes
- Volumes — The Volume table will show any volume mapped to any port owned by this storage node

Hitachi Ops Center data collector

This data collector uses Hitachi Ops Center’s integrated suite of applications to access inventory and performance data of multiple storage devices. For inventory and capacity discovery, your Ops Center installation must include both the "Common Services" and "Administrator" components. For performance collection, you must additionally have "Analyzer" deployed.

Terminology

OnCommand Insight acquires the following inventory information from this 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	OnCommand Insight Term
Storage Systems	Storage
Volume	Volume
Parity Groups	Storage Pool(RAID), Disk Groups
Disk	Disk

Vendor/Model Term	OnCommand Insight Term
Storage Pool	Storage Pool(Thin, SNAP)
External Parity Groups	Storage Pool(Backend), Disk Groups
Port	Storage Node → Controller Node → Port
Host Groups	Volume Mapping and Masking
Volume Pairs	Storage Synchronization

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

Inventory Requirements

You must have the following in order to collect inventory data:

- IP address or hostname of the Ops Center server hosting the "Common Services" component
- Root/sysadmin user account and password that exist on all servers hosting Ops Center components. HDS did not implement REST API support for usage by LDAP/SSO users until Ops Center 10.8+

Performance requirements

The following requirements must be met in order to collect performance data:

- The HDS Ops Center "Analyzer" module must be installed
- Storage arrays must be feeding the Ops Center "Analyzer" module

Configuration

Field	Description
Hitachi Ops Center IP Address	IP address or fully-qualified domain name of the Ops Center server hosting the "Common Services" component
User Name	User name for the Ops Center server.
Password	Password used for the Ops Center server.

Advanced configuration

Field	Description
Connection Type	HTTPS (port 443) is the default
Override TCP Port	Specify the port to use if not the default
Inventory Poll Interval (min)	Interval between inventory polls. The default is 40.
Choose 'Exclude' or 'Include' to specify a list	Specify whether to include or exclude the array list below when collecting data.

Filter device List	Comma-separated list of device serial numbers to include or exclude
Performance Poll Interval (sec)	Interval between performance polls. The default is 300.

HDS NAS (HNAS) data source

The HDS NAS (HNAS) data source is an inventory and configuration data source to support discovery of HDS NAS clusters. Insight supports discovering NFS and CIFS shares, file systems (Insight Internal Volumes), and spans (Insight Storage Pools).

This data source is SSH based, so the acquisition unit that will host it needs to be able to initiate SSH sessions to TCP 22 on the HNAS itself, or the Systems Management Unit (SMU) that the cluster is connected to.

Terminology

OnCommand Insight acquires the following inventory information from the HNAS data source. For each asset type acquired by Insight, the most common terminology used for this asset is shown. When viewing or troubleshooting this data source, keep the following terminology in mind:

Vendor/Model Term	Insight Term
Tier	Disk Group
Cluster	Storage
Node	Storage Node
Span	Storage Pool
File System	Internal Volume



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

Requirements

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

- Device IP address
- Port 22, SSH protocol
- Username and password - privilege level: Supervisor
- NOTE: This data collector is SSH based, so the AU that hosts it must be able to initiate SSH sessions to TCP 22 on the HNAS itself, or the Systems Management Unit (SMU) that the cluster is connected to.



This data collector is SSH based, so the AU that hosts it must be able to initiate SSH sessions to TCP 22 on the HNAS itself, or the Systems Management Unit (SMU) that the cluster is connected to.

Configuration

Field	Description
HNAS Host	IP address or fully-qualified domain name of HNAS Management Host
User Name	User name for HNAS CLI
Password	Password used for HNAS CLI

Advanced configuration

Field	Description
Inventory Poll Interval (min)	Interval between inventory polls (default 30 minutes)
SSH Banner Wait Timeout (sec)	SSH banner wait timeout (default 15 seconds)
SSH Command Timeout (sec)	SSH command timeout (default 30 seconds)

HP CommandView AE data source

The HP CommandView Advanced Edition (AE) and CommandView AE CLI/SMI (AE Lite) data sources support inventory and performance from a CommandView (also referred to as HiCommand) Device Manager server.

Terminology

OnCommand Insight acquires the following inventory information from the HP CommandView AE and AE Lite data sources. For each asset type acquired by Insight, the most common terminology used for this asset is shown. When viewing or troubleshooting this data source, keep the following terminology in mind:

Vendor/Model Term	Insight Term
PDEV	Disk
Journal Pool	Disk Group
Storage Array	Storage
Port Controller	Storage Node

Array Group, DP Pool	Storage Pool
Logical Unit, LDEV	Volume



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

Inventory Requirements

- IP address of the HiCommand Device Manager server
- Read-only user name and password for the CommandView AE software and peer privileges
- The CommandView AE Lite version of the device manager has only the CLI licensed
- Port requirement: 2001

Performance Requirements

- HDS USP, USP V, and VSP performance
 - Performance Monitor must be licensed.
 - Monitoring switch must be enabled.
 - The Export Tool (`Export.exe`) must be copied to the OnCommand Insight Server.
 - The Export Tool version must match the microcode version of the target array.
- HDS AMS performance
 - Performance Monitor needs to be licensed.
 - The Storage Navigator Modular 2 (SNM2) CLI utility needs to be installed on the OnCommand Insight Server.
 - You must register all AMS, WMS, SMS storage arrays whose performance needs to be acquired by OnCommand Insight by using the following command:

```
auunitaddauto.exe -ip<IP address of Controller0>IP address of Controller1>
```

- You must ensure that all the arrays that you registered are listed in the output of this command:
`auunitref.exe.`

Configuration

Field	Description
HiCommand Server	IP address or fully-qualified domain name of the HiCommand Device Manager server
User Name	User name for the HiCommand Device Manager server.

Password	Password used for the HiCommand Device Manager server.
Devices - USP, USP V, VSP/R600 Storages	Device list for VSP G1000 (R800), VSP (R700), HUS VM (HM700) and USP storages. Each storage requires: <ul style="list-style-type: none"> • Array's IP: IP address of the storage • User Name: User name for the storage • Password: Password for the storage • Folder Containing Export Utility JAR Files: The folder containing the Export utility .jar files
SNM2Devices - WMS/SMS/AMS Storages	Device list for WMS/SMS/AMS storages. Each storage requires: <ul style="list-style-type: none"> • Array's IP: IP address of the storage • Storage Navigator CLI Path: SNM2 CLI path • Account Authentication Valid: Select to choose valid account authentication • User Name: User name for the storage • Password: Password for the storage
Choose Tuning Manager for Performance	Choose Tuning Manager for performance and override other performance options
Tuning Manager Host	IP address or fully-qualified domain name of tuning manager
Tuning Manager Port	Port used for Tuning Manager
Tuning Manager Username	User name for Tuning Manager
Tuning Manager Password	password for Tuning Manager



In HDS USP, USP V, and VSP, any disk can belong to more than one array group.

Advanced configuration

Field	Description
HiCommand Server Port	Port used for the HiCommand Device Manager
HTTPs Enabled	Select to enable HTTPs

Inventory Poll Interval (min)	Interval between inventory polls (default 40 minutes)
Choose 'Exclude' or 'Include' to specify a list	Specify whether to include or exclude the array list below when collecting data
Exclude or Include Devices	Comma-separated list of device ID's or array names to include or exclude
Query Host Manager	Select to query host manager
HTTP Timeout (sec)	HTTP connection timeout (default 60 seconds)
Performance Polling Interval (sec)	Interval between performance polls (default 300 seconds)
Export timeout in seconds	Export utility timeout (default 300 seconds)

HP EVA Storage data source

For configuration, The EVA Storage (SSSU) data source requires the IP address of the Command View (CV) server and a *read-only* username and password to the CV software. The user must be defined in the CV software.

Terminology

OnCommand Insight acquires the following inventory information from the HP EVA data source. For each asset type acquired by Insight, the most common terminology used for this asset is shown. When viewing or troubleshooting this data source, keep the following terminology in mind:

Vendor/Model Term	Insight Term
Disk	Disk
Disk Group	Disk Group (not modeled)
Storage Cell	Storage
Virtual Disk	Storage Pool
Virtual Disk	Volume



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

Inventory Requirements

- IP address of the CV server
- Read-only username and password to the CV software. The user must be defined in the CV software.
- Third-party software installed on the OnCommand Insight Server/RAU: `sssu.exe`. The `sssu.exe` version should correspond to the CV version.
- Access validation: Run `sssu.exe` commands using username and password.

Performance Requirements

The HP StorageWorks Command View EVA software suite must be installed on the OnCommand Insight Server. Alternatively, you can install a Remote Acquisition Unit (RAU) on the EVA server:

1. Install HP StorageWorks Command View EVA Software Suite on the OnCommand Insight Server, or install a Remote Acquisition Unit on the Command View EVA server.
2. Locate the `evaperf.exe` command. For example, `c:\Program Files\Hewlett-Packard\EVA Performance Monitor\`
3. Using the IP of the Command View server, perform these steps:
 - a. Run this command where 860 is the default port `Evaperf.exe server <Command View Server IP\> 860 <username\>`
 - b. Enter the Command View server password at the password prompt.

This should return a command line prompt and nothing else.

4. Verify the setup by running `evaperf.exe ls`.

You should see a list of arrays or controllers managed by the Command View server. Each line shows a controller on an EVA array.

Configuration

Field	Description
CommandView Server	IP address or fully-qualified domain name of the EVA Storage Manager
User Name	User name for the Command View manager. The name must be defined in Command View.
Password	Password used for the Command View manager.
Performance User Name	For performance, the user name for the Command View manager. The name must be defined in Command View.
Performance Password	For performance, the password used for the Command View manager.

Advanced configuration

Field	Description
Inventory Poll Interval (min)	Interval between inventory polls (default 40 minutes)
CLI Home	Full pathname to the CLI home directory where <code>ssu.exe</code> is located
Inventory Exclude Devices	Comma-separated list of device names to include
Performance Poll Interval (sec)	Interval between performance polls (default 300 seconds)
Performance CLI Home	For Array Performance, full pathname to the CLI home directory where <code>ssu.exe</code> is located. To validate access, run <code>ssu.exe</code>
Command Timeout (sec)	<code>evaperf</code> command wait timeout (default 600 seconds)
Performance Exclude Devices	Comma-separated list of device names to exclude from collecting performance data

HPE Nimble data source

The HPE Nimble data collector supports inventory and performance data for HPE Nimble storage arrays.

Terminology

OnCommand Insight acquires the following inventory information from the HPE Nimble data source. For each asset type acquired by Insight, the most common terminology used for this asset is shown. When viewing or troubleshooting this data source, keep the following terminology in mind:

Vendor/Model Term	Insight Term
Array	Storage
Disk	Disk
Pool	Storage Pool
Volume	Volume
Initiator	Storage Host Alias

Controller	Storage Node
Fibre Channel Interface	Controller



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

Requirements

- The array must be installed and configured, and reachable from the client through its fully qualified domain name (FQDN) or array management IP address.
- The array must be running NimbleOS 2.3.x or later.
- You must have a valid user name and password to the array.
- Port 5392 must be open on the array.

Configuration

Field	Description
Array Management IP Address	Fully qualified domain name (FQDN) or array management IP address.
User Name	User name for the Nimble array
Password	Password for the Nimble array

Advanced configuration

Field	Description
Port	Port used by Nimble REST API. The default is 5392.
Inventory Poll Interval (min)	Interval between inventory polls (default 60 minutes)

Note: The default performance poll interval is 300 seconds and can not be changed. This is the only interval supported by Nimble.

Huawei OceanStor data source

OnCommand Insight uses the Huawei OceanStor (REST/HTTPS) data source to discover inventory for Huawei OceanStor storage.

Terminology

OnCommand Insight acquires the following inventory and performance information from the Huawei OceanStor. For each asset type acquired by OnCommand Insight, 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	OnCommand Insight Term
Storage Pool	Storage Pool
File System	Internal Volume
Controller	Storage Node
FC Port (Mapped)	Volume Map
Host FC Initiator (Mapped)	Volume Mask
NFS/CIFS Share	Share
Share	iSCSI Target Node
iSCSI Link Initiator	iSCSI Initiator Node
Disk	Disk
LUN	Volume

Requirements

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

- Device IP
- Credentials to access OceanStor device manager
- Port 8088 must be available

Configuration

Field	Description
OceanStor Host IP Address	IP address or fully-qualified domain name of the OceanStor Device Manager
User Name	Name used to log into the OceanStor Device Manager
Password	Password used to log into the OceanStor Device Manager

Advanced configuration

Field	Description
TCP Port	TCP Port used to connect to OceanStor Device Manager (default 8088)
Inventory Poll Interval (min)	Interval between inventory polls (default 60 minutes)
Connection Timeout (sec)	Connection timeout (default 60 seconds)

IBM Cleversafe data source

This data source collects inventory and performance data for IBM Cleversafe.

Requirements

The following are requirements for configuring this data source:

- Manager IP Address or Host Name
- A username and password for same
- Port 9440

Configuration

Field	Description
Cleversafe manager Host Name or IP Address	Host IP address of the CleverSafe device
User Name	Name used to log into the Cleversafe
Password	Password used to log into the Cleversafe

Advanced configuration

Field	Description
Inventory Poll Interval (min)	Default is 60 minutes
HTTP Connection Timeout)	Default is 60 seconds

IBM DS data source

The IBM DS (CLI) data source supports DS6xxx and DS8xxx devices only. DS3xxx, DS4xxx, and DS5xxx devices are supported by the NetApp E-Series data source. You should refer to the Insight data source support matrix for supported models and firmware

versions.

Terminology

OnCommand Insight acquires the following inventory information from the IBM DS data source. For each asset type acquired by Insight, the most common terminology used for this asset is shown. When viewing or troubleshooting this data source, keep the following terminology in mind:

Vendor/Model Term	Insight Term
Disk Drive Module	Disk
Storage Image	Storage
Extent Pool	Storage Pool
Fixed Block Volume	Volume



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

Requirements

- IP address of each DS array
- Storage Display Name is optional and cosmetic only
- Read-only username and password on each DS array
- Third-party software installed on the Insight server: IBM dscli
- Access validation: Run `dscli` commands using the username and password
- Port requirements: 80, 443, & 1750

Configuration

Field	Description
DS storage	IP address or fully-qualified domain name of the DS Storage Host
User Name	Name used for the DS CLI
Password	Password used for the DS CLI
Executable dscli.exe Path	Full path to the <code>dscli.exe</code> utility.

Advanced configuration

Field	Description
Inventory Poll Interval (min)	Interval between inventory polls (default 40 minutes)
Storage Display Name	Name of the IBM DS storage array
Inventory Exclude Devices	Comma-separated list of device serial numbers to exclude from inventory collection
Performance Poll Interval (sec)	Interval between performance polls (default 300 seconds)
Performance Filter Type	Include: Data collected only from devices on list. Exclude: No data from these devices is collected
Performance Filter Device List	Comma-separated list of device IDs to include or exclude from performance collection

IBM PowerVM data source

The IBM PowerVM (SSH) data source collects information about virtual partitions running on IBM POWER hardware instances managed by a hardware management console (HMC). For configuration, this data source requires the user name to log in to the HMC through SSH, and the view-level permission on HMC configurations.

Terminology

OnCommand Insight acquires the following inventory information from the IBM PowerVM data source. For each asset type acquired by Insight, the most common terminology used for this asset is shown. When viewing or troubleshooting this data source, keep the following terminology in mind:

Vendor/Model Term	Insight Term
hdisk	Virtual Disk
Managed System	Host
LPAR, VIO Server	Virtual Machine
Volume Group	Data Store
Physical Volume	LUN



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

Requirements

- IP address of the Hardware Management Console (HMC)
- User name and password that provide access to HMC through SSH
- Port requirement SSH-22
- View permission on all management systems and logical partition security domains

The user must also have View permission on HMC configurations and the ability to collect VPD information for the HMC console security grouping. The user must also be allowed Virtual IO Server Command access under the Logical Partition security grouping. It is a best practice to start from a role of an operator and then remove all roles. Read-only users on the HMC do not have privileges to run proxied commands on AIX hosts.

- IBM best practice is to have the devices monitored by two or more HMCs. Be aware that this may cause OnCommand Insight to report duplicated devices, therefore it is highly recommended to add redundant devices to the "Exclude Devices" list in the Advanced Configuration for this data collector.

Configuration

Field	Description
Hardware Management Console (HMC) Address	IP address or fully-qualified domain name of the PowerVM Hardware Management Console
HMC User	User name for the Hardware Management Console
Password	Password used for the Hardware Management Console

Advanced configuration

Field	Description
Inventory Poll Interval (min)	Interval between inventory polls (default 20 minutes)
SSH Port	Port used for SSH to the PowerVM
SSH Process Wait Timeout (sec)	SSH process timeout (default 600 seconds)
Number of Retries	Number of inventory retry attempts
Exclude Devices	Comma-separated list of device IDs or display names to exclude

IBM SVC data source

The IBM SVC data source collects inventory and performance data using SSH,

supporting a variety of devices that run the SVC operating system. The list of supported devices includes models such as the SVC, the v7000, the v5000, and the v3700. Refer to the Insight data source support matrix for supported models and firmware versions.

Terminology

OnCommand Insight acquires the following inventory information from the IBM SVC data source. For each asset type acquired by Insight, the most common terminology used for this asset is shown. When viewing or troubleshooting this data source, keep the following terminology in mind:

Vendor/Model Term	Insight Term
Drive	Disk
Cluster	Storage
Node	Storage Node
Mdisk Group	Storage Pool
Vdisk	Volume
Mdisk	Backend LUN



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

Inventory Requirements

- IP address of each SVC cluster
- Port 22 available
- Public and private key pair that you either generate with Insight or reuse a keypair already in use on your SVC

If you are reusing an existing keypair, you must convert them from Putty format to OpenSSH format.

- Public key installed on the SVC cluster
- Private key needs to be identified in the data source
- Access validation: Open `ssh` session to the SVC cluster using the private key



No third-party software needs to be installed.

Performance Requirements

- SVC Console, which is mandatory for any SVC cluster and required for the SVC discovery foundation package.
- Administrative access level required only for copying performance data files from cluster nodes to the

config node.



Because this access level is not required for the SVC foundation discovery package, the SVC foundation user might not work successfully.

- Port 22 required
- A private and public SSH key must be generated for this user, and the private key stored so that it is accessible from the Acquisition Unit. If the SVC foundation user has the proper permissions, then the same user and key works. The same SSH key can be used for inventory and performance data.
- Enable data collection by connecting to the SVC cluster by SSH and running: `svctask startstats -interval 1`



Alternatively, enable data collection using the SVC management user interface.

Parent Serial Number explained

Traditionally Insight is capable of reporting the storage array serial number, or the individual storage node serial numbers. However, some storage array architectures do not cleanly align to this. An SVC cluster can be comprised of 1-4 appliances, and each appliance has 2 nodes. If the appliance itself has a serial number, that serial number is neither the serial number for the cluster nor the nodes.

The attribute "Parent Serial Number" on the storage node object is populated appropriately for IBM SVC arrays when the individual nodes sit inside an intermediate appliance/enclosure that is just part of a larger cluster.

Configuration

Field	Description
Cluster/s IP	IP address of fully-qualified domain name for the SVC storage
Choose 'Password' or 'OpenSSH Key File' to specify credential type	The credential type used to connect to the device via SSH
Inventory User Name	User name for the SVC CLI
Inventory Password	Password for the SVC CLI
Full Path to Inventory Private Key	Full path to the Inventory private key file
Performance User Name	User name for the SVC CLI for performance collection
Performance Password	Password for the SVC CLI for performance collection
Full Path to Performance Private Key	Full path to the Performance private key file

Advanced configuration

Field	Description
Inventory Poll Interval (min)	Interval between inventory polls (default 40 minutes)
Exclude Devices	Comma-separated list of device IDs to exclude from inventory collection
SSH Process Wait Timeout (sec)	SSH process timeout (default 200 seconds)
Performance Poll Interval (sec)	Interval between performance polls (default 300 seconds)
Performance Exclude Devices	Comma-separated list of device IDs to exclude from performance collection
Performance SSH Process Wait Timeout (sec)	SSH process timeout (default 200 seconds)
To clean up dumped stats files	Select to clean up dumped stats files

IBM Tivoli Monitoring data source

This data source is used solely for File System Utilization. It communicates directly with the Tivoli Monitoring Database, also known as the Tivoli Monitoring Data Warehouse. Oracle and DB2 databases are supported.

Oracle error message



This data collector is no longer available starting with OnCommand Insight 7.3.11.

If the specified SID results in the error message containing "ORA-12154" on attempting to connect, double-check your Oracle DB network service configuration. If the access configuration specifies a fully qualified hostname (for example, "NAMES.DEFAULT_DOMAIN"), try inserting the fully qualified service name in the SID field. A simple example would be that the connection to SID `testdb` is failing and your Oracle configuration specifies a domain of `company.com`. The following string can be used instead of the base SID to try to connect: `testdb.company.com`.

Configuration

Field	Description
Tivoli Monitoring Database IP	IP address or fully-qualified domain name of the Tivoli Monitoring server
User Name	User name for the Tivoli Monitoring server

Password	Password for the Tivoli Monitoring server
----------	---

Advanced configuration

Field	Description
Tivoli Monitoring Database Port	Port used for Tivoli monitoring database
Oracle SID or DB2 Database Name	Oracle listener service ID or DB2 database name
Inventory Poll Interval (min)	Interval between inventory polls (default 60 minutes)
Database Driver to Use	Choose Database Driver to use
Protocol Used to Connect to the Database	Protocol Used to Connect to the Database
Database Schema	Enter Database Schema

IBM TotalStorage DS4000 data source

This data source collects inventory and performance information. There are two possible configurations (firmware 6.x and 7.x+), and they both have the same values. The API collects the volume data statistics.

Configuration

Field	Description
Comma Separated List of Array SANtricity Controller IPs	IP addresses or fully-qualified domain names of controllers, separated by commas

Requirements

- IP address of each DS5 or FAStT array
- Access validation: Ping the IP address of both controllers on each array.

Advanced configuration

Field	Description
Inventory Poll Interval (min)	Interval between inventory polls (default 30 minutes)
Performance Poll Interval (up to 3600 seconds)	Interval between performance polls (default 300 seconds)

IBM XIV data source

IBM XIV (CLI) data source inventory is performed by using the XIV command-line interface. XIV performance is accomplished by making SMI-S calls to the XIV array, which runs a SMI-S provider on port 5989.

Terminology

OnCommand Insight acquires the following inventory information from the IBM XIV data source. For each asset type acquired by Insight, the most common terminology used for this asset is shown. When viewing or troubleshooting this data source, keep the following terminology in mind:

Vendor/Model Term	Insight Term
Disk	Disk
Storage System	Storage
Storage Pool	Storage Pool
Volume	Volume



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

Requirements

- Port requirement: TCP port 7778
- IP address of the XIV management interface
- Read-only user name and password
- The XIV CLI must be installed on the Insight server or RAU
- Access validation: Log in to the XIV user interface from the Insight server using the user name and password.

Configuration

Field	Description
IP Address	IP address or fully-qualified domain name for the XIV storage
User Name	User name for the XIV storage
Password	Password for the XIV storage
Full path to XIV CLI directory	Full path to the XIV CLI directory

Advanced configuration

Field	Description
Inventory Poll Interval (min)	Interval between inventory polls (default 40 minutes)
CLI Process Wait Timeout (ms)	CLI process timeout (default 7200000 ms)
SMI-S Host IP	IP address of the SMI-S Provider Host
SMI-S Port	Port used by SMI-S Provider Host
SMI-S Protocol	Protocol used to connect to the SMI-S provider
SMI-S Namespace	SMI-S namespace
Username	User name for the SMI-S Provider Host
Password	Password for the SMI-S Provider Host
Performance Poll Interval (sec)	Interval between performance polls (default 300 seconds)
Number of SMI-S Connection Retries	Number of SMI-S connection retry attempts

Infinidat InfiniBox data source

The Infinidat InfiniBox (HTTP) data source is used to collect information from the Infinidat InfiniBox storage. You must have access to the InfiniBox Management Node.

Terminology

OnCommand Insight acquires the following inventory information from the InfiniBox data source. For each asset type acquired by Insight, the most common terminology used for this asset is shown. When viewing or troubleshooting this data source, keep the following terminology in mind:

Vendor/Model Term	Insight Term
Drive	Disk
InfiniBox	Storage
Node	Storage Node
Pool	Storage Pool

Volume	Volume
FC Port	Port
Filesystem	Internal Volume
Filesystem	FileShare
Filesystem Exports	Share



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

Configuration

Field	Description
InfiniBox Host	IP address or fully-qualified domain name of the InfiniBox Management Node
User Name	User name for InfiniBox Management Node
Password	Password for the InfiniBox Management Node

Advanced configuration

Field	Description
TCP Port	TCP Port used to connect to InfiniBox Server (default 443)
Inventory Poll Interval (min)	Interval between inventory polls (default 60 minutes)
Connection Timeout	Connection timeout (default 60 seconds)

Microsoft Azure compute data source

OnCommand Insights uses the Azure compute data collector to acquire inventory and performance data from Azure compute instances.

Requirements

You need the following information to configure this data collector:

- Port requirement: 443 HTTPS

- Azure Management Rest IP (management.azure.com)
- Azure Service Principal Application (Client) ID (user account)
- Azure Service Principal Authentication key (user password)

You need to set up an Azure account for Insight discovery. Once the account is properly configured and you register the application in Azure, you will have the credentials required to discover the Azure instance with Insight. The following link describes how to set up the account for discovery: <https://docs.microsoft.com/en-us/azure/active-directory/develop/howto-create-service-principal-portal>

Configuration

Enter data into the data source fields according to the table below:

Field	Description
Azure Service Principal Application (Client) ID (Reader role required)	Sign-in ID to Azure. Requires Reader Role access.
Azure tenant ID	Microsoft tenant ID
Azure Service Principal Authentication Key	Login authentication key
I understand Microsoft bills me for API requests	Check this to verify your understanding that Microsoft bills you for API requests made by Insight polling.

Advanced Configuration

Enter data into the data source fields according to the table below:

Field	Description
Inventory Poll Interval (min)	The default is 60
Choose 'Exclude' or 'Include' to Apply to Filter VMs by Tags	Specify whether to include or exclude VM's by Tags when collecting data. If 'Include' is selected, the Tag Key field can not be empty.
Tag Keys and Values on which to Filter VMs	Click + Filter Tag to choose which VMs (and associated disks) to include/exclude by filtering for keys and values that match keys and values of tags on the VM. Tag Key is required, Tag Value is optional. When Tag Value is empty, the VM is filtered as long as it matches the Tag Key.
Performance Poll Interval (sec)	The default is 300

Azure NetApp Files data source

This data source acquires inventory and performance data for Azure NetApp Files (ANF).

Requirements

The following are requirements for configuring this data source:

- Port requirement: 443 HTTPS
- Azure Management Rest IP (management.azure.com)
- Azure Service Principal Application (Client) ID (user account)
- Azure Service Principal authentication key (user password)
- You need to set up an Azure account for Cloud Insights discovery.

Once the account is properly configured and you register the application in Azure, you will have the credentials required to discover the Azure instance with Cloud Insights. The following link describes how to set up the account for discovery:

<https://docs.microsoft.com/en-us/azure/active-directory/develop/howto-create-service-principal-portal>

Configuration

Field	Description
Azure Service Principal Application (Client) ID	Sign-in ID to Azure
Azure Tenant ID	Azure Tenant ID
Azure Service Principal Authentication Key	Login authentication key
I understand Microsoft bills me for API requests	Check this to verify your understanding that Microsoft bills you for API requests made by Insight polling.

Advanced configuration

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

Microsoft Hyper-V data source

For configuration, the Microsoft Hyper-V data source requires the IP address or the resolvable DNS name for the physical host (hypervisor). This data source uses Powershell (previously used WMI).

Terminology

OnCommand Insight acquires the following inventory information from the Hyper-V data source. For each asset type acquired by Insight, the most common terminology used for this asset is shown. When viewing or troubleshooting this data source, keep the following terminology in mind:

Vendor/Model Term	Insight Term
Virtual hard Disk	Virtual Disk
Host	Host
Virtual Machine	Virtual Machine
Cluster Shared Volumes (CSV), Partition Volume	Data Store
Internet SCSI Device, Multi Path SCSI LUN	LUN
Fiber Channel Port	Port



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

Requirements

- The Hyper-V requires port 5985 opened for data collection and remote access/management.
- IP address of Clustering group node
- Local Administrator user & password on the hypervisor
- Administrative-level user account
- Port requirements: Port 135 and Dynamic TCP ports assigned 1024-65535 for Windows 2003 and older and 49152-65535 for Windows 2008.
- DNS resolution must succeed, even if the data collector is pointed at only an IP address.
- Each Hyper-V hypervisor must have “Resource Metering” turned on for every VM, on every host. This allows each hypervisor to have more data available for Cloud Insights on each guest. If this is not set, fewer performance metrics are acquired for each guest. More information on Resource metering can be found in the microsoft documentation:

[Hyper-V Resource Metering Overview](#)

[Enable-VMResourceMetering](#)

Configuration

Field	Description
-------	-------------

Physical Host IP Address	The IP address or fully-qualified domain name for the physical host (hypervisor)
User Name	Administrator user name for the hypervisor
Password	Password for the hypervisor
NT Domain	The DNS name used by the nodes in the cluster

Advanced configuration

Field	Description
Inventory Poll Interval (min)	Interval between inventory polls (default 20 minutes)
Connection Timeout (ms)	Connection timeout (default 60000 ms)

NetApp Clustered Data ONTAP data source

This data source should be used for storage systems using Clustered Data ONTAP, and requires an administrator account used for read-only API calls.

Terminology

OnCommand Insight acquires the following inventory information from the Clustered Data ONTAP data source. For each asset type acquired by Insight, the most common terminology used for this asset is shown. When viewing or troubleshooting this data source, keep the following terminology in mind:

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



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

Requirements

- Administrator account used for read-only API calls
- Target IP is the cluster management LIF
- Username (with read-only role name to ontapi application to the default Vserver) and password to log into NetApp cluster
- Port requirements: 80 or 443
- License requirements: FCP license and mapped/masked volumes required for discovery

Configuration

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

Advanced configuration

Field	Description
Inventory Poll Interval (min)	Interval between inventory polls (default 20 minutes)
Performance Poll Interval (sec)	Interval between performance polls (default 300 seconds)

Clustered Data ONTAP Storage

Terms applying to objects or references that you might find on NetApp Clustered Data ONTAP storage asset landing pages.

Clustered Data ONTAP Storage Terminology

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

- 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 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, OCI is sourcing this value directly, but this is often not the case. In lieu of the array offering this up, OCI 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 Insight data source as part of inventory reporting.

Clustered Data ONTAP Storage Pool

Terms applying to objects or references that you might find on NetApp Clustered Data ONTAP storage pool asset landing pages.

Clustered Data ONTAP Storage Pool Terminology

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

- 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.

Clustered Data ONTAP Storage Node

Terms applying to objects or references that you might find on NetApp Clustered Data ONTAPs storage node asset landing pages.

Clustered Data ONTAP Storage Node Terminology

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

- 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.

NetApp Clustered Data ONTAP for Unified Manager data source

This data source collects ONTAP 8.1.x data from the Unified Manager (UM) 6.0+ database. Using this data source, Insight discovers all clusters configured and populated in UM. For efficiency, Insight does not call ZAPIs on the cluster itself. Performance is not supported in this data source.

Configuration



This data collector is no longer available starting with OnCommand Insight 7.3.11.

Field	Description
Unified Manager IP	IP address or fully-qualified domain name of the Unified Manager

User Name	User name for the Unified Manager
Password	Password for the Unified Manager
Port	Port used for communication with the Unified Manager (default 3306)

Advanced configuration

Field	Description
Inventory Poll Interval (min) Interval	Interval between inventory polls (default 15 minutes)
Exclude Clusters	Comma-separated list of cluster IPs to exclude

NetApp Data ONTAP operating in 7-Mode data source

For storage systems using Data ONTAP software operating in 7-Mode, you should use the ONTAPI data source, which uses the CLI to obtain capacity numbers.

Terminology

OnCommand Insight acquires the following inventory information from the NetApp Data ONTAP 7-Mode data source. For each asset type acquired by Insight, the most common terminology used for this asset is shown. When viewing or troubleshooting this data source, keep the following terminology in mind:

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



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

Requirements

- IP address of the FAS storage controller and partner
- Port 443
- User name and password for the controller and the partner
- 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 Filer	IP address or fully-qualified domain name for the NetApp Filer
User Name	User name for the NetApp Filer
Password	password for the NetApp Filer
Address of HA Partner Filer in Cluster	IP address or fully-qualified domain name for the HA Partner Filer
User Name of HA Partner Filer in Cluster	User name for the NetApp HA Partner Filer
Password of HA Partner Filer in Cluster	password for the NetApp HA Partner Filer

Advanced configuration

Field	Description
Inventory Poll Interval (min)	Interval between inventory polls (default 20 minutes)
Connection Type	Choose connection type
Connection Port	Port used for NetApp API

Performance Poll Interval (sec)	Interval between performance polls (default 300 seconds)
---------------------------------	--

Storage systems connection

As an alternative to using the default administrative user for this data source, you can configure a user with administrative rights directly on the NetApp storage systems so that this data source 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 OnCommand Insight to execute all NetApp storage API commands. This command is required to use the ZAPI.
 - "login-http-admin": Use this to allow OnCommand Insight to connect to the NetApp storage via HTTP. This command is required to use the ZAPI.
 - "security-api-vfiler": Use this to allow OnCommand Insight 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,
or
login-http-admin,api-*,security-api-vfile,cli-*
```

NetApp E-Series data source

The NetApp E-Series data source collects inventory and performance information. There are two possible configurations (firmware 6.x and firmware 7.x+), and they both have the same values.

Terminology

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

Vendor/Model Term	Insight Term
Drive	Disk
Volume Group	Disk Group
Storage Array	Storage
Controller	Storage Node
Volume Group	Storage Pool
Volume	Volume



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

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)	Interval between inventory polls (default 30 minutes)

Performance Poll Interval (up to 3600 seconds)	Interval between performance polls (default 300 seconds)
--	--

E-Series Storage

Terms applying to objects or references that you might find on NetApp E-Series storage asset landing pages.

E-Series Storage Terminology

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

- Model — model name of the device.
- 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 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. Insight calculates an IOPs-weighted average derived from the volumes in the storage.
- Throughput — the array’s total host facing throughput. Insight sums the volumes’ throughput to derive this value.
- Management — this may contain a hyperlink for the management interface of the device. Created programmatically by the Insight data source as part of inventory reporting.

E-Series Storage Pool

Terms applying to objects or references that you might find on NetApp E-Series storage pool asset landing pages.

E-Series Storage Pool Terminology

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

- 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 by using efficiency technologies you have allocated a sum total of 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.
- Throughput — the sum throughput of all the disks contributing capacity to this storage pool.

E-Series Storage Node

Terms applying to objects or references that you might find on NetApp E-Series storage node asset landing pages.

E-Series Storage Node Terminology

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

- 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 — Utilization is not currently available for NetApp E-Series.
- IOPS — 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. Insights calculates an IOPs weighted average from volumes that belong exclusively to this node.
- Throughput — a number representing the host driven throughput on this controller. Calculated by summing all the throughput for volumes that belong exclusively to this node.
- Processors — CPU count.

NetApp Host and VM File Systems data source

You can use the NetApp Host and VM File Systems data source to retrieve file system details and storage resource mappings for all Microsoft Windows host and VM (virtual

machine) file systems and for all supported Linux VMs (those that are virtually mapped only) existing in the Insight server that are annotated with the configured Compute Resource Group (CRG).

General Requirements

- This feature must be purchased separately.

You can contact your Insight representative for assistance.

- You should check the Insight support matrix to verify that your host or virtual machine operating system is supported.

To verify that links from file systems to storage resources are created, check that the relevant storage or virtualization vendor type and version report the volume or virtual disk identification data required.

Microsoft Windows Requirements

- This data source uses Window Management Instrumentation (WMI) data structures to retrieve data.

This service must be operational and available remotely. In particular, port 135 must be accessible and must be opened if behind a firewall.

- Windows domain users must have the appropriate permissions to access WMI structures.
- Administrator permissions are required.
- Dynamic TCP ports assigned 1024-65535 for Windows 2003 and older
- Ports 49152—65535 for Windows 2008



As a general rule, when trying to use a firewall between Insight, an AU, and this data source, you should consult with your Microsoft team to identify the ports they believe will be required.

Linux Requirements

- This data source uses a Secure Shell (SSH) connection to execute commands on Linux VMs.

The SSH service must be operational and available remotely. In particular, port 22 must be accessible and must be opened if behind a firewall.

- SSH users must have sudo permissions to execute read-only commands on Linux VMs.

You must use the same password to log in to SSH and to answer any sudo password challenge.

Usage Recommendations

- You should annotate a group of hosts and virtual machines that have common operating system credentials using the same Compute Resource Group annotation.

Each group has an instance of this data source discovering file system details from those hosts and virtual machines.

- If you have an instance of this data source for which the success rate is low (for example, OnCommand Insight is discovering file system details for only 50 of 1000 hosts and virtual machines in a group), you should move the hosts and virtual machines for which discovery is successful into a separate Compute Resource Group.

Configuration

Field	Description
User Name	Operating system user with appropriate rights to retrieve file system data For Windows operating system users, this must include the domain prefix.
Password	Password for the operating system user
Compute Resource Group	Annotation value used to flag host and virtual machines for the data source discovers file systems.A blank value indicates that the data source discovers file systems for all hosts and virtual machines not currently annotated with any Compute Resource Group.

Advanced configuration

Field	Description
Inventory poll interval (min)	Interval between inventory polls (default 360 minutes)

NetApp SolidFire data source

The NetApp SolidFire data source supports both iSCSI and Fibre Channel SolidFire configurations, for both inventory and performance collection.

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

Terminology

OnCommand Insight acquires the following inventory information from the NetApp SolidFire data source. For each asset type acquired by Insight, the most common terminology used for this asset is shown. When viewing or troubleshooting this data source, keep the following terminology in mind:

Vendor/Model Term	Insight Term
Drive	Disk
Cluster	Storage

Node	Storage Node
Volume	Volume
Fibre Channel Port	Port
Volume Access Group, LUN Assignment	Volume Map
iSCSI Session	Volume Mask



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

Requirements

The following are requirements for configuring this data source:

- Management Virtual IP Address
- 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
Inventory Poll Interval (min)	Interval between inventory polls (default 60 minutes)
TCP Port	TCP Port used to connect to SolidFire Server (default 443)
Connection Timeout (sec)	Connection timeout (default 60 seconds)
Performance Poll Interval (sec)	Interval between performance polls (default 300 seconds)

Troubleshooting

When SolidFire reports an error it is displayed in OnCommand Insight 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.

NetApp StorageGRID data source

This data source collects inventory and performance data for StorageGRID.

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 (MVIP)	Host IP address of the StorageGRID
User Name	Name used to log into the StorageGRID
Password	Password used to log into the StorageGRID

Advanced configuration

Field	Description
Inventory Poll Interval (min)	Interval between inventory polls (default 60 minutes)
Performance Poll Interval (sec)	Interval between performance polls (default 900 seconds)

OpenStack data source

The OpenStack (REST API / KVM) data source collects information about OpenStack hardware instances. This data source collects inventory data for all OpenStack instances, and optionally, VM performance data.

Requirements

The following are requirements for configuring the OpenStack data source.

- IP address of the OpenStack controller
- OpenStack admin role credentials and sudo access to the Linux KVM hypervisor are recommended.



If you are not using an admin account or admin equivalent privileges, you can still acquire data from the data source. You will need to modify the policy configuration file (i.e. `etc/nova/policy.json`) to allow users with non-admin role to call the API:

- `"os_compute_api:os-availability-zone:detail": ""`
- `"os_compute_api:os-hypervisors": ""`
- `os_compute_api:servers:detail:get_all_tenants": ""`
- For performance collection the OpenStack Ceilometer module must be installed and configured. Configuring the Ceilometer is done by editing the `nova.conf` file for each hypervisor and then restart the Nova Compute service on each hypervisor. The option name changes for different releases of OpenStack:
 - Icehouse
 - Juno
 - Kilo
 - Liberty
 - Mitaka
 - Newton
 - Ocata
- For CPU stats, `"compute_monitors=ComputeDriverCPUMonitor"` needs to be turned on in `/etc/nova/nova.conf` on compute nodes.
- Port requirements:
 - 5000 for http and 13000 for https, for the Keystone service
 - 22 for KVM SSH
 - 8774 for Nova Compute Service
 - 8776 for Cinder Block Service
 - 8777 for Ceilometer Performance Service
 - 9292 for Glance Image Service



The port binds to the specific service, and the service may run on the controller or another host in larger environments.

Configuration

Field	Description
OpenStack Controller IP Address	IP address or fully-qualified domain name of the OpenStack Controller
OpenStack Administrator	User name for an OpenStack Admin
OpenStack Password	Password used for the OpenStack Admin
OpenStack Administrator Tenant	OpenStack Administrator Tenant
KVM Sudo User	KVM Sudo User name
Choose 'Password' or 'OpenSSH Key File' to specify credential type	The credential type used to connect to the device via SSH
Full Path to Inventory Private Key	Full Path to Inventory Private Key
KVM Sudo Password	KVM Sudo Password

Advanced configuration

Field	Description
Enable hypervisor inventory discovery through SSH	Check this to enable hypervisor inventory discovery through SSH
OpenStack Admin URL port	OpenStack Admin URL port
Use HTTPS	Check to use secure HTTP
HTTP Connection Timeout (sec)	Timeout for HTTP connection (default 300 seconds)
SSH Port	Port used for SSH
SSH Process Wait Timeout (sec)	SSH process timeout (default 30 seconds)
SSH Process Retries	Number of inventory retry attempts
Inventory Poll Interval (min)	Interval between inventory polls (default 20 minutes)

Oracle ZFS data source

The Oracle ZFS data source supports inventory and performance collection.

Terminology

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

Vendor/Model Term	Insight Term
Disk (SDD)	Disk
Cluster	Storage
Controller	Storage Node
LUN	Volume
LUN Map	Volume Map
Initiator, Target	Volume mask
Share	Internal Volume



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

Requirements

The following are requirements for configuring this data source:

- Host names for the ZFS Controller-1 and the ZFS Controller-2
- Administrator user name and credentials
- Port requirement: 215 HTTP/HTTPS

Configuration

ZFS Controller-1 Hostname	Host name for storage controller 1
ZFS Controller-2 Hostname	Host name for storage controller 2
User name	User name for the storage system administrator user account
Password	Password for the administrator user account

Advanced configuration

Field	Description
TCP port	TCP Port used to connect to ZFS (default 215)
Connection Type	HTTP or HTTPS
Inventory poll interval	Inventory poll interval (default 60 minutes)
Connection Timeout	Default is 60 seconds
Performance Poll Interval (sec)	Interval between performance polls (default 300 seconds)

Troubleshooting

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

Problem:	Try This:
"Invalid login credentials"	validate Zfs user account and password
"Configuration error" with error message "REST Service is disabled"	Verify REST service is enabled on this device.
"Configuration error " with error message "User unauthorized for command"	Likely due to certain roles (for example, 'advanced_analytics') are not included for the configured user <userName>.Possible Solution: <ul style="list-style-type: none">• Correct the Analytics (statistic) scope for the user \${user} with the read only role:- From the Configuration → Users screen, put your mouse over the role and double click to allow editing• Select "Analytics" from the Scope drop down menu. A list of the possible properties appears.• Click the top most check box and it will select all three properties.- Click the Add button on the right side.• Click the Apply button at the top right of the pop-up window. The pop-up window will close.

Pure Storage FlashArray data source

The Pure Storage FlashArray (HTTP) data source is used to collect information from the Pure Storage Flash Array. Insight supports both inventory and performance collection.

Terminology

OnCommand Insight acquires the following inventory information from the Pure Storage FlashArray data source. For each asset type acquired by Insight, the most common terminology used for this asset is shown. When viewing or troubleshooting this data source, keep the following terminology in mind:

Vendor/Model Term	Insight Term
Drive (SSD)	Disk
Array	Storage
Controller	Storage Node
Volume	Volume
Port	Port
LUN Map (Host, Host Group, Target Port)	Volume Map, Volume Mask



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

Requirements

- Storage system IP address
- User name and password for the Administrator account of the Pure storage system.
- Port requirement: HTTP/HTTPS 80/443

Configuration

Field	Description
FlashArray Host	IIP address or fully-qualified domain name of FlashArray Management Server
User Name	User name for the FlashArray Management Server
Password	Password for the FlashArray Management Server

Advanced configuration

Field	Description
Connection Type	Management Server

TCP Port	TCP Port used to connect to FlashArray Server (default 443)
Connection Timeout (sec)	Connection timeout (default 60 seconds)
Inventory Poll Interval (min)	Interval between inventory polls (default 60 minutes)
Performance Poll Interval (sec)	Interval between performance polls (default 300 seconds)

QLogic FC Switch data source

For configuration, the QLogic FC Switch (SNMP) data source requires the network address for the FC Switch device, specified as an IP address, and an SNMP *read-only* community string used to access the device.

Configuration

Field	Description
SANSurfer Switch	IP address or fully-qualified domain name for the SANSurfer switch
SNMP version	SNMP version
SNMP community	SNMP Community String
User Name	User name for the SANSurfer switch
Password	Password for the SANSurfer switch

Advanced configuration

Field	Description
Inventory Poll Interval (min)	Interval between inventory polls (default 15 minutes)
SNMP Auth Protocol	SNMP authentication protocol (SNMPv3 only)
SNMP Retries	Number of SNMP retry attempts
SNMP Timeout (ms)	SNMP timeout (default 5000 ms)
Enable Trapping	Select to enable trapping

Minimum Time Between Traps (sec)	Minimum time between acquisition attempts triggered by traps (default 10 seconds)
Fabric Name	Fabric name to be reported by the data source. Leave blank to report the fabric name as WWN.
Performance Poll Interval (sec)	Interval between performance polls (default 300 seconds)

Red Hat (RHEV) data source

The Red Hat Enterprise Virtualization (REST) data source collects information about RHEV instances via HTTPS.

Requirements

- IP address of the RHEV server over port 443 via REST API
- Read-only username and password
- RHEV Version 3.0+

Configuration

Field	Description
RHEV Server IP Address	IP address or fully-qualified domain name of the RHEV server
User Name	User name for the RHEV server
Password	Password used for the RHEV server

Advanced configuration

Field	Description
HTTPS Communication Port	Port used for HTTPS communication to RHEV
Inventory Poll Interval (min)	Interval between inventory polls (default 20 minutes)
Connection timeout (sec)	Connection timeout (default 60 seconds)

Violin Flash Memory Array data source

The Violin 6000-Series Flash Memory Array (HTTP) data source collects network information for analysis and validation from Violin 6000-series flash memory arrays.

Terminology



This data collector is no longer available starting with OnCommand Insight 7.3.11.

OnCommand Insight acquires the following inventory information from the Violin 6000-Series Flash Memory Array data source. For each asset type acquired by Insight, the most common terminology used for this asset is shown. When viewing or troubleshooting this data source, keep the following terminology in mind:

Vendor/Model Term	Insight Term
Violin Intelligent Memory Module (VIMM)	Disk
Container	Storage
Memory Gateway	Storage Node
LUN	Volume
Initiator, Initiator Group, Target	Volume Map, Volume Mask



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

Requirements

- You need a read-only user name and password to the storage.
- Validate access with a web browser using the storage IP address.

Configuration

Field	Description
IP address or FQDN of Violin Memory Array Main Gateway	IP address or fully-qualified domain name of the Violin Memory Array Main Gateway
User Name	User name for the Violin Memory Array Main Gateway
Password	Password for the Violin Memory Array Main Gateway

Advanced configuration

Field	Description
Communication Port	Port used for communication with Violin array
HTTPS Enabled	Select to use HTTPS

Inventory Poll Interval (min)	Interval between inventory polls (default 20 minutes)
Connection timeout (sec)	Connection timeout (default 60 seconds)
Performance Poll Interval (sec)	Interval between performance polls (default 300 seconds)

VMware vSphere data source

The VMware vSphere (Web Services) data source collects ESX Host information and requires *read-only* privileges on all objects within the Virtual Center.

Terminology

OnCommand Insight acquires the following inventory information from the VMware vSphere data source. For each asset type acquired by Insight, the most common terminology used for this asset is shown. When viewing or troubleshooting this data source, keep the following terminology in mind:

Vendor/Model Term	Insight Term
Virtual Disk	Disk
Host	Host
Virtual Machine	Virtual Machine
Data Store	Data Store
LUN	LUN
Fiber Channel Port	Port



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

Requirements

- IP address of the Virtual Center server
- Read-only username and password in Virtual Center
- Read-only privileges on all objects within the Virtual Center.
- SDK access on the Virtual Center server
- Port requirements: http-80 https-443
- Validate access by logging in to Virtual Center Client using your user name and password and verifying that the SDK is enabled by entering `telnet <vc_ip> 443`.

Configuration

Field
Description
Virtual Center Address
Network address for the Virtual Center or vSphere server, specified as an IP (<i>nnn.nnn.nnn.nnn</i> format) address or as a host name that can be resolved through DNS.
User Name
User name for the VMware server.
Password
Password for the VMware server.

Advanced configuration

Field	Description
Inventory Poll Interval (min)	Interval between inventory polls (default 20 minutes)
Connection Timeout (ms)	Connection timeout (default 60000 ms)
Filter VMs by	Choose how to filter VMs
Choose 'Exclude' or 'Include' to specify a list	Specify whether to include or exclude the VM list below when collecting data
List of VMs to filter (Comma Separated, or Semicolon Separated If Comma Is Used in the Value)	Comma-separated or semicolon-separated list of VMs to include or exclude from polling
Number of Retries for Requests to vCenter	Number of vCenter Request retry attempts
Communication Port	Port used for VMware server
Performance Poll Interval (sec)	Interval between performance polls (default 300 seconds)

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.