

IBM Storage Ceph Plugin for vSphere

User Guide 2024-05-20



Important

DRAFT BETA PROVIDED WITH NO SUPPORT: Do not upgrade a production installation to a beta release. Upgrade support is not provided between beta releases. No automated upgrade path is provided from a beta release to the final GA release. The cluster will need to be rebuilt for production use.

Last updated: 2024-05-20

Contents

- Important..... i**

- Figures..... v**

- Tables..... vii**

- What's new..... 9**
- Trademarks..... 11**
- Overview..... 13**
- Compatibility and minimum requirements..... 15**
- Restrictions and limitations..... 19**
- Installing..... 21**
 - Installation overview..... 21
 - Deploying the appliance..... 22
 - Registering vSphere plug-in into vCenter 23
 - Registering with linked vCenters..... 24
 - Checking the plug-in status..... 24
 - Unregistering the vSphere Plugin..... 25
 - Forced unregistration..... 25
- Managing storage systems in the plug-in dashboard..... 27**
 - Adding a storage system to the plug-in..... 27
 - Removing a storage system from the plug-in..... 28
 - Editing storage system details..... 28
 - Connecting with IBM Storage Insights..... 29
 - Manually refreshing the storage system and vSphere objects 29
- Managing VMFS datastores..... 31**
 - Creating a VMFS datastore..... 31
 - Attaching a VMFS datastore to a virtual machine..... 32
 - Removing a VMFS datastore..... 32
 - Expanding a VMFS datastore..... 33
- Managing datastore snapshots (Technology Preview)..... 35**
 - Copying snapshots to a new datastore..... 35
 - Deleting a snapshot..... 36
 - Taking a manual snapshot of a VMFS datastore..... 36
- Additional IBM Storage information for vSphere objects (Technology Preview)..... 39**
 - VMFS Datastore summary panel..... 39
 - vSphere Cluster summary and configuration..... 39
 - ESXi host summary and configuration..... 40
 - Creating ESXi host definitions..... 41
- Troubleshooting..... 43**
 - Collecting a snap for support engagement..... 43
 - vSphere task names misreported..... 44
 - Storage devices refresh timeout..... 44
 - Miscellaneous troubleshooting..... 46
- Related information and publications..... 47**

Figures

- 1. vSphere architecture..... 14
- 2. Including beta fixes for download..... 21

Tables

1. vSphere user permission requirements.....	17
--	----

What's new

Learn about the dates and nature of updates to the published information of the IBM Storage Ceph Plugin for vSphere.

Date	What's new
7 June 2024	<p data-bbox="596 409 1300 443">Initial Release (IBM Storage Ceph Plugin for vSphere 1.0.0)</p> <p data-bbox="596 457 1458 520">The version information was added in IBM® part of the initial release and includes the following features:</p> <ul data-bbox="596 535 1252 779" style="list-style-type: none"><li data-bbox="596 535 1024 569">• Deployment of the vSphere Plugin<li data-bbox="596 579 1089 613">• Registration of the plug-in to vCenter(s)<li data-bbox="596 623 748 657">• Dashboard<li data-bbox="596 667 1247 701">• Adding compatible storage systems for management<li data-bbox="596 711 943 745">• Managing VMFS datastores<li data-bbox="596 756 1247 789">• Managing datastore snapshots (Technology Preview)

Trademarks

IBM, the IBM logo, and ibm.com[®] are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at Copyright and trademark information at www.ibm.com/legal/copytrade.shtml.

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.

VMware, the VMware logo, VMware Cloud Foundation, VMware Cloud Foundation Service, VMware vCenter Server, and VMware vSphere are registered trademarks or trademarks of VMware, Inc. or its subsidiaries in the United States and/or other jurisdictions.

Other product and service names might be trademarks of IBM or other companies.

Overview

The IBM Storage Ceph Plugin for vSphere allows management and control of storage volumes from within the VMware vSphere Client. The plug-in and associated workflow is integrated into the vSphere user interface. The interface allows for provisioning and management of VMFS datastores across multiple storage systems, all from within a single interface.

IBM Storage Ceph Plugin for vSphere is a VMware vCenter add-on that simplifies storage workflow for daily operations, such as:

- Adding an IBM Storage Ceph NVMe Gateway.
- Creating and attaching a new NVMe namespace.
- Removing a retired storage from vSphere clusters.

The plug-in supports all block storage devices within the IBM Storage Ceph 7.1 or higher cluster.

The IBM Storage Ceph Plugin for vSphere is based on the VMware Remote plug-in architecture and is supported on VMware vSphere 7.0 and higher.

Plug-in architecture

The plug-in architecture consists of two parts, which are delivered as a single OVF/OVA appliance:

- VMware PhotonOS - the underlying operating system of the appliance.

Note: Photon OS is a minimalist Linux operating system from VMware. For more information, see [Introduction to Photon OS · VMware Photon OS 4.0 Documentation](#).

- IBM Storage Ceph Plugin for vSphere application - a collection of scripts and Docker containers to facilitate integration into both vSphere and IBM Storage environments.

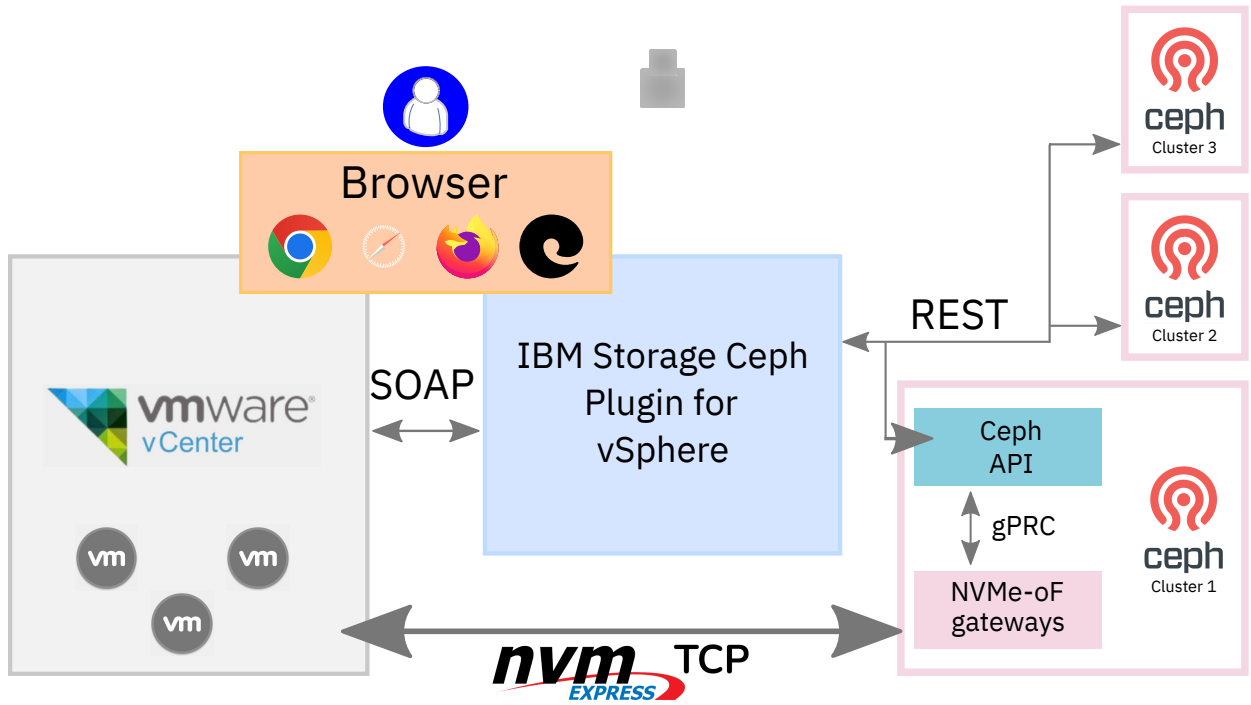


Figure 1. vSphere architecture

Compatibility and minimum requirements

The IBM Storage Ceph Plugin for vSphere is based on the VMware remote plug-in architecture and is distributed and deployed by using an Open Virtualization Appliance (OVA) bundle.

Appliance requirements

The virtual machine appliance requires the following resources:

- 2 vCPUs and 4 GB of memory
- 20 GB of datastore space - can be deployed on a vSphere Virtual Machine File System (VMFS) or vVol datastore
- Red Hat Enterprise Linux 9.2 or higher.
- Following are the networking (Ethernet, TCPIP) requirements:
 - One IP address (IPv4), Static or DHCP
 - Gateway
 - DNS
 - Netmask

Hardware requirements

- Ready Nodes specifications
 - 2 Intel Xeon Gold 6438N processors (32C/64T 2.0 GHz)
 - 512 GB DDR4
 - 1 BOSS controller with 2x 460 GB M.2 SSD RAID1
 - 1 Dual port 10 GbE or 25 GbE network adapter
 - NVMe SSDs per node - 8x, 16x, 24x
 - NVMe SSD capacities - 3.84 TB, 7.68 TB, or 15.36 TB
 - Dual-port Intel 810-C 100 GbE network adapters - 1, 2
 - SR optical transceivers are included
- Rack - 2RU per IBM Storage Ready Node
- Networking
 - Management switch - 2 × 1 GbE Base-T ports per IBM Storage Ready Node
 - High-speed switches
 - 2x 100 GbE QSFP28 ports per IBM Storage Ready Node to High Speed Switch 1
 - 2x 100 GbE QSFP28 ports per IBM Storage Ready Node to High Speed Switch
 - 2 SR optical transceivers

VMware environment requirements

- VMware 7 update 3 and higher or vSphere 8.0 to support NVMe over TCP.
- All browsers and platforms are supported by vCenter.
- NVMe over TCP connectivity between vCenter, the appliance, and the storage system.
- Fully qualified domain name (FQDN) or IP address of vCenters.
- Credentials for a vCenter user account with administrative permissions.

Storage cluster requirements

- IBM Storage Ceph 7.1 or later cluster with at least four nodes.
- Replication factor - 3
- Estimated data durability - 9-nines
- Ceph Dashboard installed with the credentials
- An rbd pool
- NVMe-oF configuration, with minimum of:
 - One gateway deployed
 - One subsystem created
 - One listener added

Note: A host must be added before adding the listener.

For more information about NVMe-oF configuration, see **Ceph NVMe-oF gateway** within the [IBM Storage Ceph documentation](#) for your IBM Storage Ceph version.

The following are a couple of example cluster configurations:

- 4 nodes with 8 × 3.84 TB NVMe SSDs = 120 TB raw, 40 TB usable
- 4 nodes with 24 × 15.36 TB NVMe SSDs = 1474 TB raw, 490 TB usable

IBM Storage Ceph Plugin for vSphere requirements

IBM Storage Ceph 7.1 or later.

Access and system information requirements

The following are access and system information that is needed for registering the plug-in:

- IP address or FQDN of the storage system with port number 8443.
- Credentials for a user account on the storage system with administrative permissions.

Note: Create a dedicated user ID with an Administrator role for the plug-in. Do not use either superuser or any user account with the Security Administrator role.

Network requirements

- BMC - The BMC network is used to interact with server management interfaces such as Redfish, IPMI, or KVM console.
 - 1 × 1 GbE link to server BMC port
 - 1 × 27 CIDR range per rack
 - Used for BMC / server management
- Node - The Node network is used by cephadm to configure and manage the IBM Storage Ceph cluster.
 - 1 × 1 GbE link
 - 1 × 27 CIDR range per rack
 - Used for cephadm cluster management
- NVMe/TCP - The NVMe/TCP network is used to carry traffic between NVMe/TCP clients (hosts) and IBM Storage Ceph NVMe/TCP Gateways.
 - 2 × 100 GbE links
 - Distinct IP address and upstream switch
 - 2 × 27 CIDR ranges per rack
 - Used for IBM Storage Ceph NVMe/TCP gateways

- Cluster - The Cluster network is used to carry traffic between IBM Storage Ceph components.
 - 2 × 100 GbE MLAG / cLAG
 - 1 × 27 CIDR range per rack
 - Bonded that use LACP
 - Layer 3 + 4 xmit_hash_policy
 - Switch transmits hashing set to src-ip, src-port, proto, dst-ip, dst-port
 - Used for Ceph client_network and cluster_network

Operating system requirements

Red Hat Enterprise Linux (RHEL) version 9.2 or later must be installed to the BOSS RAID1 volume on IBM Storage Ready Nodes.

vSphere user permission requirements

Use the vSphere user permission requirements table to ensure the vSphere user has the correct permissions for IBM Storage Ceph Plugin for vSphere operations.

<i>Table 1. vSphere user permission requirements</i>	
vSphere user privileges	Permissions to use IBM Storage Ceph Plugin for vSphere workflow
<ul style="list-style-type: none"> • Tasks <ul style="list-style-type: none"> – Create task – Update task 	<ul style="list-style-type: none"> • Create, Expand, or Delete Datastore • Take or Delete Snapshot • Copy to New Datastore • Add Host
<ul style="list-style-type: none"> • Plug-in <ul style="list-style-type: none"> – Manage plug-ins 	<ul style="list-style-type: none"> • Add, Edit, or Delete Storage System
<ul style="list-style-type: none"> • Host <ul style="list-style-type: none"> – Configuration <ul style="list-style-type: none"> - Storage Partition Configuration – Inventory <ul style="list-style-type: none"> - Add Host to Cluster - Modify Cluster – Local Operations <ul style="list-style-type: none"> - Add Host to vCenter 	<ul style="list-style-type: none"> • Add Host • Create, Expand, Attach, Detach, or Delete Datastore • Copy to New Datastore
<ul style="list-style-type: none"> • Datastore <ul style="list-style-type: none"> – Rename Datastore 	<ul style="list-style-type: none"> • Copy to New Datastore

Restrictions and limitations

Learn about the restrictions and limitations with IBM Storage Ceph Plugin for vSphere. Limitations are functionality restrictions that are part of the predefined system design and capabilities in a particular version.

MVP restrictions and limitations

- Only a single Ceph cluster has been tested.
- Only a single RBD pool can be used.
 - The pool has a predefined name of `rbd`.

Note: To have a pool with a name other than `rbd`, manually set the name in the `plugin-config.json` file. In the JSON file, add:

```
"rbd_pool" : "NEW_NAME"
```

- Only a single NVMe subsystem with a predefined name can be used.

The NQN has the predefined name of `nqn.2016-06.io.spdk:cnode1`

Note: To have a subsystem with an NQN other than the predefined `nqn.2016-06.io.spdk:cnode1`, manually set the name in the `plugin-config.json` file. In the JSON file, add:

```
"subsystem_NQN" : "NEW_NQN"
```

General restrictions and limitations

Multiple I/O group systems

On storage systems with multiple I/O groups, ESXi hosts must have connectivity to all I/O groups. When creating new VMFS datastores with the plug-in, if a host is defined in a subset of I/O groups and volumes are created in I/O groups where the host is not configured, volume mappings are created. However, the mapping does not get detected by the host. Therefore, the **Create VMFS Datastore** task fails.

ESXi host connectivity

Only TCP/IP and NVMe connectivity between the ESXi hosts and the storage systems are supported by the plug-in.

Multiple vCenters must be in linked mode

The plug-in does not support multiple vCenters to be registered to the same plug-in appliance unless they are configured as linked-mode vCenters. When multiple vCenters that are not in linked-mode are used, multiple instances of the vCenter plug-in appliance must be used.

Installing

The IBM Storage Ceph Plugin for vSphere is built upon the VMware remote plug-in architecture and is distributed as a VM appliance. The plug-in appliance is delivered as an OVA and is deployed into vSphere as a standard OVA/OVF template. After deployment, you are required to register the plug-in into the vCenter environment with the appliance CLI.

Downloading the IBM Storage Ceph Plugin for vSphere OVA from Fix Central for beta

To install the IBM Storage Ceph Plugin for vSphere during beta release, your Fix Central must be set to include beta fixes.

1. Go to <https://www.ibm.com/support/fixcentral/ac/options>.
2. Select **Include beta fixes**.

User options

Choose your options.

- Include entitled fixes available to IBM staff
- Include preview fixes
- Include preview products (requires preview fixes)
- Include beta fixes**
- Include beta products (requires beta fixes)

Submit

Figure 2. Including beta fixes for download

3. Click **Submit**.
4. In the **Find Product** tab, select **IBM Storage Ceph** from the **Product selector** and **Release 7.1**.
5. Click **Continue**.
6. Download the **IBM Storage Ceph Plugin for vSphere 1.0.0-Beta** fix pack.

Installation overview

Use these instructions are for a quick installation overview of the IBM Storage Ceph Plugin for vSphere.

About this task

The instructions here are for high-level installation instructions only. For more in-depth instructions, follow the detailed installation steps.

Procedure

1. Download the Open Virtual Appliance (OVA) installation package.

[IBM Fix Central](#)

Note: For beta release download instructions, see [“Downloading the IBM Storage Ceph Plugin for vSphere OVA from Fix Central for beta”](#) on page 21.

2. Deploy the appliance.

From the vSphere Client, deploy the **OVF template** within a VMware cluster.

Note: In the **Select networks** step, select a network that allows communication between vCenter and the management interfaces for IBM Storage Systems.

3. Power on the appliance after the OVA is successfully deployed.

- By default, the appliance is configured with a firewall and therefore does not respond to ping (ICMP) requests.
- Initial login has a default password of IBMplugin and requires the user to change the password before continuing.

4. Register the vSphere Plugin into vCenter.

- a) Log in to the vSphere Plugin appliance as root.
- b) Register the vSphere Plugin with the vCenter instance or instances.

The registration command displays the thumbprint of the vSphere instance for verification.

```
ibm-plugin register -u <vCenter Username> -v <IP_or_FQDN_of_vCenter>
```

Note: To link multiple linked-mode vCenters to a single plug-in appliance, register each vCenter independently.

Enter the vCenter password when prompted.

- c) To complete the plug-in activation, refresh the vSphere Client browser.

5. Check the plug-in status to see which vCenters are registered to the appliance.

```
ibm-plugin status
```

The **ibm-plugin status** command can be run at any time.

If the plug-in is not registered to any vCenter, the **Plugin Registered** output displays as `False`. For example:

```
root@IBM-production-ova[~]#ibm-plugin status
Plugin Version: 1.0.0
Plugin Build: 2024_05_01-1120
Plugin Registered: False
```

Deploying the appliance

Learn to deploy the IBM Storage Ceph Plugin for vSphere appliance into an existing vSphere environment.

Before you begin

Download the IBM Storage Ceph Plugin for vSphere OVA or plug-in upgrade package from [IBM Fix Central](#).

Note: For beta release download instructions, see [“Downloading the IBM Storage Ceph Plugin for vSphere OVA from Fix Central for beta”](#) on page 21.

Procedure

1. From the vSphere Client, right-click on the VMware cluster that the appliance will be deployed to and select **Deploy OVF Template**.
2. In the **Deploy OVF Template** wizard, on the **Select an OVF Template** page, select **Local file > Upload Files**.
3. From the local file system, select all the IBM Storage Ceph Plugin for vSphere OVA files.

4. On the **Select a name and folder** page, provide a name for the virtual machine, select a deployment location, and click **Next**.
5. On the **Select a compute resource** page, select a VMware cluster or specific ESXi host on that the virtual machine (VM) will run on and click **Next**.
6. Verify the details are correct on the **Review details** page.
7. Read and accept appropriate license agreements in the **License Agreements** step.
8. On the **Select storage** page, select a compatible datastore, where the appliance will be stored.
9. On the **Select networks** page, select the appropriate destination network to use for the plug-in.

Note: You must select a network that allows communication between vCenter and the management interfaces for IBM Storage Ceph.

10. Pre-configure the networking details from within the **Customize Template** step.
Networking details that are provided here are automatically configured in the operating system of the virtual machine upon deployment.
 - a) Under Network Type, select either **DHCP** or **Static** configuration.
 - b) If Static IP configuration is selected, complete the necessary information:
 - Hostname (optional)
 - IP address
 - Gateway
 - Netmask
 - DNS Server(s)
 - DNS Domain (optional)

11. On the **Ready to complete** page, review the final list of chosen settings.

12. Click **Finish** to deploy the OVA with defined customizations.

13. After the OVA is successfully deployed, power on the appliance.

After the initial power-on and configuration, the plug-in appliance should be available on the network.

Note: By default, the appliance is configured with a firewall and therefore does not respond to ping (ICMP) requests.

14. You can now SSH into the appliance.

Note: Initial log in has a default password of `IBMplugin` and requires the user to change the password before continuing.

```
ssh root@IP_ADDRESS_OF_PLUGIN_VM
```

Registering vSphere plug-in into vCenter

Register the IBM Storage Ceph Plugin for vSphere into your vSphere environment.

Before you begin

Be sure that all of the deployment steps that are specified in [“Deploying the appliance” on page 22](#) are completed successfully.

Procedure

1. Log in to the vSphere plug-in appliance as `root`.
2. Register the plug-in with vCenter(s).

```
ibm-plugin register -u vCENTER_USERNAME -v IP_OR_FQDN_OF_vCENTER
```

The registration command displays the thumbprint of the vSphere instance for verification. You are then be prompted to enter the vCenter password.

Note: The hostname that is specified for the vCenter instance must either be the Fully Qualified Domain Name (FQDN) or IP address. Registering the plug-in with a short hostname cause errors when using the plug-in.

What to do next

Once the plug-in has been successfully registered, a banner message appears in the vSphere Client. Click **REFRESH BROWSER** on the banner to complete the plug-in activation.

Note: Refreshing the browser is necessary for the plug-in activation.

In addition, a number of vSphere tasks are started. Initially, **Download plug-in** and **Deploy plug-in** tasks appear, and then a final task **Populating information from IBM Storage Systems** begins. This population process discovers existing host and datastore information from the vSphere environment and correlate this with any host definitions or volumes present in any storage systems registered in the plug-in dashboard. This population task is triggered whenever the VM appliance is restarted, registered within a vCenter, or when IBM Storage Ceph system is added to the dashboard. This process can take a few minutes, during which time some inconsistencies or inaccuracies with the information is displayed in the summary panels and dashboard.

Registering with linked vCenters

The plug-in appliance supports multiple vCenters in linked-mode configuration. This allows multiple linked-mode vCenters to share a single plug-in appliance. However, in this configuration, each vCenter must be registered independently. Storage systems need to be registered into the plug-in only once to be available from all linked vCenters.

Note: While the plug-in supports multiple vCenters in linked-mode, it does not support multiple vCenters that are not linked. If multiple vCenters are not linked, each vCenter needs its own deployment of the plug-in appliance.

Each vCenter in the linked-mode configuration must be registered separately to the plug-in by using the **ibm-plugin** command with **register** parameter. Use the **ibm-plugin status** to view successful registration with all linked-mode vCenters. For more information, see [“Checking the plug-in status” on page 24](#).

```
Plugin Version: 1.0.0.0
Plugin Build: 2024_03_07-1120
Plugin Registered: True
Registered vCenters:
+-----+-----+
| FQDN                | Username                |
+-----+-----+
| test-vcsa-1.ssd.hursley.ibm.com | administrator@vsphere.local |
| test-vcsa-2.ssd.hursley.ibm.com | administrator@vsphere.local |
+-----+-----+
Network Settings:
+-----+-----+
| DHCP      | no                      |
| Address   | x.x.x.x/x              |
| Gateway   | x.x.x.x                |
| DNS-1     | x.x.x.x                |
| DNS-2     | x.x.x.x                |
| Hostname  | xx-xxx-xx              |
+-----+-----+
```

Checking the plug-in status

Check the plug-in status to verify vCenter registration and version information.

Use the **ibm-plugin** command with the **status** parameter to check which vCenter(s) are registered to the appliance at any time. Using the **ibm-plugin** command also allows for display of the current version of the plug-in that is available in the appliance.

```
ibm-plugin status
```

Linked vCenters are shown together in the registered vCenters section.

For example,

```
root@IBM-production-ova [ ~ ]# ibm-plugin status
Plugin Version: 1.0.0.0
Plugin Build: 2024_03_07-1120
Plugin Registered: True
Registered vCenters:
+-----+-----+
| FQDN                               | Username                               |
+-----+-----+
| test-vcsa-1.ssd.hursley.ibm.com    | administrator@vsphere.local          |
| test-vcsa-2.ssd.hursley.ibm.com    | administrator@vsphere.local          |
+-----+-----+
Network Settings:
+-----+-----+
| DHCP      | no                                     |
| Address   | x.x.x.x/x                             |
| Gateway   | x.x.x.x                               |
| DNS-1     | x.x.x.x                               |
| DNS-2     | x.x.x.x                               |
| Hostname  | xx-xxx-xx                             |
+-----+-----+
```

If the plug-in is not registered to any vCenter, the **Plugin Registered** boolean shows as *False*. For example,

```
root@IBM-production-ova [ ~ ]# ibm-plugin status
Plugin Version: 1.0.0.0
Plugin Build: 2024_03_07-1120
Plugin Registered: False
```

Unregistering the vSphere Plugin

Unregistering the IBM Storage Ceph Plugin for vSphere removes the plug-in from vCenter and removes the vCenter registration from the plug-in appliance.

Use the **ibm-plugin** command to unregister the plug-in with any vCenter provided.

```
ibm-plugin unregister -u <vCENTER_USERNAME> -v <IP_or_FQDN_OF_vCENTER>
```

The following output confirms successful unregistration:

Plugin successfully unregistered from vCenter.

Similarly to registration, unregistering a plug-in from vCenter displays a successful undeployment message and appropriate completion in the **Recent Tasks** display.

Important: Click **REFRESH BROWSER** to complete the removal of the plug-in from the vCenter user interface.

Forced unregistration

Unregistration requires communication between the plug-in appliance and vCenter. In the case that the vCenter is unable to be communicated with, you can force the unregistration of the plug-in from the appliance.

When attempting to unregister using the **ibm-plugin unregister** command and the vCenter is not able to communicate, the following message appears:

```
root@ibm-storage-plugin [ ~ ]# ibm-plugin unregister -u administrator@vsphere.local -v test-
vcsa.ssd.hursley.ibm.com
vCenter Password:
Cannot connect to registered vCenter instance at test-vcsa.ssd.hursley.ibm.com.
Do you want to remove this vCenter instance from the local plugin config? [Y/n]
```

Selecting yes (Y) forces the unregistration at the local appliance level, but does not remove the plug-in from vCenter. Manual removal of the plug-in from vCenter is required to complete this operation.

Manually removing the plug-in from vCenter, when running vCenter 8.0 or higher

Plug-ins can now be removed from within the main vSphere Client within environments running vCenter 8.0 or higher.

1. From the Shortcuts menu in the vSphere Client, go to **Administration > Solutions > Client Plugins**.
2. Select the vSphere Plugin plug-in.
3. Select the specific plug-in to be removed, and click **REMOVE**.
4. The **Remove Instance** displays, confirming the removal of the plug-in. Click **YES**.
5. When prompted, refresh the browser to complete the plug-in removal.

Managing storage systems in the plug-in dashboard

Use the vSphere Plugin to manage and control storage objects directly from within the VMware vSphere Client.

You can add, remove, and edit storage systems from the dashboard.

- Add a storage system to the plug-in by going to **IBM Storage > Add Storage System**. For more information, see [“Adding a storage system to the plug-in” on page 27](#).
- To delete a storage system from the plug-in, from the dashboard select the storage system that you want to remove and click the delete icon. For more information, see [“Removing a storage system from the plug-in” on page 28](#).
- To edit storage system details, from the dashboard select the storage system that you want to edit and click the pencil icon. For more information, see [“Editing storage system details” on page 28](#).

Adding a storage system to the plug-in

Use the IBM Storage Ceph Plugin for vSphere to register compatible IBM storage systems.

Before you begin

Add an IBM Storage Ceph storage system that meets the minimum requirements for the IBM Storage Ceph Plugin for vSphere. For a full list of compatible storage systems, see [“Compatibility and minimum requirements” on page 15](#).

Procedure

Use the following procedure to add a storage device to the vSphere Plugin.

1. From the vSphere Client, go to the plug-in dashboard from the vSphere Client main menu and click **IBM Storage**.
2. From the dashboard, click **Add Storage System**.
The **Add Storage System** wizard opens.
3. In the **Connect to Storage** step of the wizard, complete the mandatory fields

URL

Fill in as `[https://]<IP_ADDRESS | FQDN>:PORT`

Note: HTTPS is optional and is automatically removed. Entering an IP address or FQDN and a port number are mandatory.

Username

This is the username for the administrator account. Use the Ceph Dashboard username.

Password

Use the Ceph Dashboard password.

Note:

- The following user roles are allowed: SecurityAdmin, Administrator, RestrictedAdmin, and 3SiteAdmin, but it is recommended to use a user account assigned with the Administrator role.
 - Object-Based Access Control is supported with Ownership Groups. In multi-tenancy environments or configurations where you want restricted access, consider creating Ownership Groups within the storage system and configure object ownership on the specific hosts and child pools.
4. Click **VALIDATE**.
The validate step checks proper TCP/IP connectivity and authentication credentials to the storage system. When validated, an optional alias field is available.

Note: The alias field is designed to simplify administration of multiple storage systems by allowing a friendly name to help identify specific storage clusters or resources. This field can be left blank.

5. Click **NEXT**

6. From the **Review** step, make sure that all parameters are correct and click **ADD** to add the storage system to the plug-in.

The wizard automatically closes on success, refreshing the dashboard to display the newly added storage system.

After a storage system has been added to the plug-in dashboard, a vSphere task **Populating information from IBM Storage Systems** displays. The population process discovers host and datastore information from the vSphere environment. The process also correlates the information with any host definitions or volumes present in storage systems that are registered in the plug-in dashboard.

The population task is triggered in any of the following situations:

- When the VM appliance restarts.
- When the VM appliance gets registered within a vCenter.
- When adding an IBM storage system from within the plug-in dashboard.
- When initiating a manual refresh of an IBM storage system within the plug-in dashboard.

The population process can take a few minutes. During the process, notice inconsistencies or inaccuracies can occur with the information that is displayed in the IBM Storage summary panes and dashboard.

Removing a storage system from the plug-in

If the IBM Storage Ceph Plugin for vSphere no longer needs to manage a specific storage system, it can be removed from the plug-in.

Procedure

1. From the vSphere Plugin Dashboard, select the storage system to remove and click the delete icon.
2. Confirm the removal of the storage system from the plug-in.

Note: Removing a storage system from the plug-in does not delete any data, volumes, logical configuration, or settings on the storage system. Removing the storage system removes only the ability to provision any subsequent storage volumes or manage existing VMFS datastores with the vSphere Plugin. Any existing datastores remain on the storage system and accessible from vSphere.

Editing storage system details

After a storage system is added to the IBM Storage Ceph Plugin for vSphere, authentication, and alias details can be modified.

Procedure

1. From the vSphere Plugin Dashboard, select the storage system to edit and click the pencil icon.
2. In the revalidate step of the wizard, check the content and complete the following fields:

- **Username**
- **Password**

3. Click **VALIDATE**.

The plug-in attempts to authenticate with the storage system to ensure that the credentials are correct.

4. Optional: Once the storage system is authenticated, the alias name is available to update. Existing aliases are prefilled. This field can be left empty.

5. Click **EDIT** to complete the changes.

Note: If new pools were added to the storage system after the vSphere Plugin registration, the changes are not reflected directly in the pool details on the plug-in dashboard. To use newly created pools with the vSphere Plugin, you need to edit the storage system and select the additional pools.

When the changes are saved successfully, the wizard closes and the dashboard automatically refreshes to reflect the new changes to the storage system.

Connecting with IBM Storage Insights

Connect a storage system with a registered instance of IBM Storage Insights from the vSphere Plugin dashboard.

Before you begin

The IBM Storage Ceph storage system must be enabled for IBM Storage Insights before connecting by using the vSphere Plugin. For more information, see **Dashboard > Configuring IBM Call Home and IBM Storage Insights** within the [IBM Storage Ceph documentation](#).

Procedure

1. From the vSphere Plugin dashboard, select the storage system.
2. Click the **IBM Storage Insights** icon.

Note: The **IBM Storage Insights** icon is only enabled when first configured through the IBM Storage Ceph storage system.

Clicking the **IBM Storage Insights** icon opens a separate browser tab to the registered instance of IBM Storage Insights.

Manually refreshing the storage system and vSphere objects

Manually refresh the dashboard when storage system or vSphere objects get out of sync with the vSphere plug-in. The manual refresh rescans all the objects that the plug-in manages and restores the sync status.

Procedure

- From the vSphere Plugin dashboard, click the **Refresh** icon.

Note: Refreshing the dashboard is a background, asynchronous process. During the process, the vCenter task, *Populating information from IBM Storage Systems* populates in the recent task list. During this brief refresh period, vCenter information might show out of date information.

The manual rescan gathers an inventory of vSphere related objects, queries the storage system for storage related objects, and refreshes the database that the vSphere Plugin manages.

Managing VMFS datastores

Use the vSphere Plugin to manage and control storage objects from within the VMware vSphere Client.

You can create, add, remove, and expand VMFS datastores from the dashboard. Manage VMFS datastores from the vSphere plug-in by going to **Actions > IBM Storage** and then selecting **Create VMFS Datastore**, **Delete VMFS datastore**, or **Expand VMFS Datastore**.

To attach an existing VMFS datastore to a virtual machine, select the datastore and click **Actions > New Virtual Machine**.

Creating a VMFS datastore

Use the IBM Storage Ceph Plugin for vSphere to automate the creation of vSphere Virtual Machine File System (VMFS) datastores.

Before you begin

Identify the vSphere cluster for which you want to create the new VMFS datastore(s).

Note: In multi-IO-group configurations, ensure that hosts are configured and zoned to all IO-groups on the storage system - volumes created in IO-groups that are not accessible to the host is not accessible.

About this task

This workflow initially evaluates storage connectivity to ensure all hosts in the selected vSphere cluster have sufficient access to the storage system. If any hosts in the selected vSphere cluster are deemed not to be correctly defined on the storage system, or if conditions exist that prevent the host status being reported as online by the underlying storage system, an alert is displayed. Host connectivity depends on if the host NQN of the given ESXi has been added to the host list of the NVME-oF subsystem. Where possible, consider using host clusters to simplify management and storage provisioning for clustered host environments.

Volume to be created:

- The name specified in the workflow is used both when creating the volumes on the storage system and for the VMFS datastore(s).
- The specified capacity defines the size of the volume created in the storage system. Note that a small area of the volume is allocated to the VMFS partition and so the available capacity for the VMFS datastore may differ from the underlying volume's capacity.

The appropriate host volume mappings are created for all members of the vSphere cluster.

By default, all volumes are created as Thin Provisioned or Space Efficient volumes on the storage system. If this is not the preferred behavior, consider assigning a Provisioning Policy to the storage pool on the storage system. This Provisioning Policy defines how volumes are created by default and overrides the default behavior.

Finally, a VMFS datastore is created from each volume, and the storage adapters for each vSphere cluster member is rescanned to correctly mount the new datastore(s).

Note: Version VMFS6 is automatically used for all datastores created by the plug-in. In accordance with best practices, a single volume is created for each VMFS datastore requested.

Procedure

1. From the vSphere Client, select the vSphere Cluster that requires new VMFS datastores, and click **Actions > IBM Storage**.
2. From the IBM Storage plug-in submenu, select **Create VMFS Datastore**.

3. Define a **Name** for the new VMFS datastore.
The same name is used when creating the volume on the storage system.
4. Specify the required **Size**.
Size is specified in either TB or GB.
5. Click **NEXT**.
6. In the **Storage System Selection** page, select the storage system for which the datastore will be created on.

Note: It can take a few moments while all ESXi hosts in the VMware cluster are cross-referenced with host objects and assigned ports. Depending on the number of hosts in the vSphere cluster, the cross-referencing can take up to 1 minute. Any discrepancies that might exist between the host objects or port assignments are indicated. A hyperlink to the management interface for the selected storage system is available for easy access to the element manager.

7. Review the summary and click **Create**.
A progress bar will appear within the view, and vSphere tasks will be generated for each step in the workflow.
8. Monitor the **Recent Tasks** view in vSphere to confirm the task is completed successfully.
After all tasks are successfully completed, the VMFS datastore(s) will be available for use.

Attaching a VMFS datastore to a virtual machine

Use the IBM Storage Ceph Plugin for vSphere to attach an existing VMFS datastore to a virtual machine (VM).

Procedure

1. From the vSphere Client, select the vSphere Cluster that requires new VMFS datastores, and click **Actions > New Virtual Machine**.
The **New Virtual Machine** wizard displays.
2. Proceed through the wizard, filling in the fields.
Click **NEXT** to continue through the wizard steps.
 - a) In **Select a name and folder**, enter a virtual machine name and select a location for which datacenter and host virtual machine needs to be deployed.
 - b) In **Select storage**, select the datastore.
 - c) In **Ready to complete**, review the virtual machine information and click **FINISH**.

Removing a VMFS datastore

Use the IBM Storage Ceph Plugin for vSphere to automate the removal of vSphere Virtual Machine File System (VMFS) datastores.

Before you begin

In the vSphere Client inventory, identify the VMFS datastore that you wish to delete and identify any remaining Virtual Machines or VM Templates that exist on this datastore. Remove these VMs or VM Templates from the vSphere inventory or migrate them to a different datastore.

Note: If any VM or VM Templates exist on the datastore, you cannot delete the datastore with the plug-in.

About this task

This task initiates a workflow to delete a VMFS datastore and the underlying volume from the storage system. All data is lost upon volume deletion.

Note: To prevent active volumes or host mappings from being deleted inadvertently, the system supports a global setting *Volume Protection*. This setting prevents these objects from being deleted if the system detects that they had recent I/O activity. If the underlying storage system is configured with *Volume Protection*, an informational message is displayed.

Removing a VMFS datastore only deletes the underlying NVMe-oF namespace. The associated RBD volumes remain. After datastore deletion, to restore data from the RBD volumes or to remove the volumes from the storage system, contact your storage administrator.

Procedure

1. From the vSphere Client, select the vSphere Cluster that requires new VMFS datastores, and click **Actions > IBM Storage**.
2. Select **Delete VMFS Datastore**.
3. Review the summary panel and click **Delete**.
4. Monitor the **Recent Tasks** view in vSphere to confirm the task was completed successfully.

Expanding a VMFS datastore

Use the IBM Storage Ceph Plugin for vSphere to automate the expansion of capacity for vSphere VMFS datastores.

Before you begin

In the vSphere Client inventory, identify the VMFS datastore that you want to expand and identify the associated storage system and pool in which it resides. Review the available capacity of this pool to ensure that there is sufficient space to store the increased volume capacity.

About this task

This task initiates a workflow that increases the available capacity of the selected VMFS datastore.

Initially, to increase the backing volume this workflow issues a volume expansion command to the underlying storage system for the volume, which is hosting the VMFS datastore.

Second, a storage rescan is initiated on the connected hosts to ensure that all hosts correctly identify the increase in volume capacity.

Finally, the VMFS partition is expanded to take advantage of the additional available capacity and the task complete notification shows in vCenter to report a successful completion

Procedure

1. Select the VMFS datastore that requires expansion and from the **Actions** menu (or by right click menu), select **IBM Storage**.
2. From the IBM Storage Ceph Plugin for vSphere submenu, select **Expand VMFS Datastore**.
3. Enter a value to increase the datastore by, select the appropriate unit, for example, TB or GB.
4. Review the capacity consumption bar to visualize the updated projected consumption.
5. Click **Expand**.
6. Monitor the **Recent Tasks** view in vSphere to confirm the task is completed successfully.

Managing datastore snapshots (Technology Preview)

Use the IBM Storage Ceph Plugin for vSphere to view and manage a selected datastore's snapshots on the storage system.

Important: Technology Preview features are not supported with IBM production service level agreements (SLAs), might not be functionally complete, and IBM does not recommend using them for production. These features provide early access to upcoming product features, enabling customers to test functionality and provide feedback during the development process.

Use these steps to get to the datastore snapshot information

1. In the vSphere Client inventory, identify the VMFS datastore on which you want to manage a datastore snapshot.
2. From the **Configure** tab, go to **Datastore Snapshots**.

If the datastore is part of a volume group with snapshots, the following information is displayed:

- Snapshot Name - The name of the snapshot on the storage system.
- Safeguarded - Whether the snapshot was marked as safeguarded when it was captured.
- Time Created - Time stamp when a snapshot was taken.

From this view, in addition to viewing snapshot information, you can complete the following actions:

- [“Copying snapshots to a new datastore” on page 35](#)
- [“Deleting a snapshot” on page 36](#)
- [“Taking a manual snapshot of a VMFS datastore” on page 36](#)

Copying snapshots to a new datastore

Use the IBM Storage Ceph Plugin for vSphere to automate copying an existing Datastore's snapshot into a new Datastore.

Before you begin

In the vSphere Client inventory, identify the VMFS Datastore that you want to copy to a new Datastore. To know more about taking a new manual snapshot of a Datastore, see [“Taking a manual snapshot of a VMFS datastore” on page 36](#).

About this task

This task triggers a workflow that creates a copy (thin-clone or full clone) from an existing Datastore's snapshot.

Thin Clone

Thin clone is a copy for which only the changed data that is written to the source populates the copy. A thin clone retains a dependency on the source Datastore.

Clone

This clone type is also populated from a volume's snapshot after copy. But unlike thin clone, it becomes independent of its source volume after the background process finishes copying all required data from the source image to the clone.

Assuming all hosts report online, the new Datastore is mapped to the same host or host cluster as the source Datastore.

Procedure

1. From the vSphere Client, select the VMFS Datastore on which you want to perform the copy operation.
2. From the **Configure** tab, go to **Datastore Snapshots**.
3. Select the desired snapshot from the datagrid. If the list is empty, click **Take Snapshot** to take a manual snapshot.
4. Once the snapshot is selected, click the **Clone** icon next to Take Snapshot.
5. By default, the new Datastore name is populated by appending the current timestamp to the source Datastore's name. If required, declare the new Datastore's name.
6. Based on clone requirement, click the **Thin Clone** or **Clone** card option.
7. Click **Copy** to confirm selections and to initiate the process.
8. Monitor the **Recent Tasks** view in vSphere to confirm the task is completed successfully.

After all tasks are successfully completed, the new VMFS datastore will be available for use.

Note: For a full clone, you do not have to wait for the background copy to complete to use the new Datastore.

Deleting a snapshot

Use the IBM Storage Ceph Plugin for vSphere to automate the removal of snapshots of a virtual machine File System (VMFS) Datastore.

Before you begin

In the vSphere Client inventory, identify the VMFS Datastore of which you want to remove snapshots.

About this task

This task triggers a workflow to remove a Datastore's volume group snapshot. Safeguarded snapshots cannot be deleted by using the IBM Storage Ceph Plugin for vSphere. To delete safeguarded snapshots, contact your system administrator.

Note: If the selected snapshot has corresponding thin cloned volumes, the snapshot transitions into a dependent deleting state and is not be deleted from the storage system.

Procedure

1. From the vSphere Client, select the **VMFS Datastore** of which you want to remove snapshots.
2. Navigate to the Configure tab, and scroll to **Datastore Snapshots**.
3. Select the snapshot to be deleted from the datagrid.
4. After the snapshot is selected, click the **Trash** icon.
5. Click **Delete** to confirm the deletion of the snapshot.
6. Monitor the **Recent Tasks** view in vSphere to confirm the task is completed successfully.

Taking a manual snapshot of a VMFS datastore

Use the IBM Storage Ceph Plugin for vSphere to automate the process of triggering a snapshot of a virtual machine File System (VMFS) datastore.

Before you begin

In the vSphere Client inventory, identify the VMFS datastore of which you want to take snapshot.

Note: The Volume Group Snapshot feature in IBM Storage Ceph Plugin for vSphere facilitates the VMFS Datastore Snapshot feature. This feature is only available on supported platforms that run 8.5.1.0 or newer.

About this task

This task identifies the storage system that presents the volume that backs the VMFS datastore and creates a volume group snapshot.

If the volume is not part of a volume group on the storage system, a new volume group is created with the same name as the volume. Later, the volume is placed within that volume group and a volume group snapshot is taken. If the volume is already part of a volume group, then a snapshot is taken of the volume group. If multiple volumes are already configured in the volume group, an informational message is displayed in vSphere and all volumes are snapshotted together.

In all cases, define an optional snapshot name to help identify the reason for the snapshot.

Procedure

1. Select the VMFS datastore that requires a snapshot.
2. Under **Actions**, select **IBM Storage**.
Alternatively, navigate to the Configure tab on the selected Datastore, and scroll to **Datastore Snapshots**.
3. Select **Take Snapshot**.
4. Optional: Define a name to help identify the snapshot in the future.
5. Click **Take Snapshot** to take a snapshot.
6. Monitor the **Recent Tasks** view in vSphere to confirm the task is completed successfully.

Additional IBM Storage information for vSphere objects (Technology Preview)

Use the vSphere Plugin to view additional information for any associated ESXi hosts, vSphere cluster, or VMFS datastores taken from IBM storage systems.

Important: Technology Preview features are not supported with IBM production service level agreements (SLAs), might not be functionally complete, and IBM does not recommend using them for production. These features provide early access to upcoming product features, enabling customers to test functionality and provide feedback during the development process.

VMFS Datastore summary panel

Use the IBM Storage Ceph Plugin for vSphere to display additional information about a selected Datastore and the underlying volume from the storage system.

In the vSphere inventory, identify the VMFS Datastore you wish to explore.

To see a summary view of the Datastore and the information about the underlying storage volume, navigate to the **Summary** tab and scroll to the **IBM Storage card**.

- **Storage System** - The name of the storage system where the volume was created.
- **Pool Name** - The name of the pool where the volume was created.
- **Namespace UUID** - The UUID of the underlying NVMe-oF namespace.
- **Volume Name** - The name of the underlying volume.

vSphere Cluster summary and configuration

Use the vSphere plug-in to view information about the vSphere cluster's representation on the IBM Storage Ceph Plugin for vSphere dashboard.

Viewing vSphere cluster summary information

In the vSphere Client inventory, identify the vSphere cluster you wish to explore.

To see a summary view, navigate to the **Summary** tab within the vSphere cluster and scroll to the **IBM Storage card**. This view shows if the vSphere cluster is represented on each registered storage system, if it is configured the host cluster name and host count is displayed. To see more details, or refresh the connection between the vSphere cluster and the storage system, click the **Go to Host Connections** link at the bottom of the card. Alternatively, navigate to the **Configure** tab within the vSphere cluster, and scroll to **IBM Storage > Host Connections**.

Note: When launching this view, a discovery task is triggered to refresh up-to-date information from each registered storage system. The ESXi host's WWPNs reported by vSphere is compared to both existing host definitions and any visible host ports on the storage system an SAN fabric. Depending on the size of your environment, this can take some time.

Host connections panel

The datagrid on the **Host Connections** view shows the representation of the vSphere cluster on every registered storage system.

Note: It is not a requirement to use host clusters to perform actions using the plug-in. However, if using host clusters, every ESXi host must be defined in the same host cluster on the storage system.

Columns in the datagrid:

- **Storage System Name** - The name of the storage system taken directly from the system.

- **Alias** - (Optional) The user defined alias name for the storage system.
- **Storage system status** - The status of the storage system as seen on the main dashboard.
- Cluster represented on the storage system -
 - Yes (green label) - Every ESXi host in the vSphere cluster is defined and in the same host cluster on the storage system.
 - No (green label) - Every ESXi host in the vSphere cluster is defined, and none of them are in a host cluster.
 - No (yellow label) - Either an ESXi host is not defined on the storage system, or the hosts are part of 2 or more different clusters on the storage system.

Note: If there is no alias for any of the storage systems the alias column is automatically hidden. You can alter column's visibility using the icon in the bottom left of the datagrid.

To view further details, expand the details pane using the double arrow (>>) on the left-hand side of each row. The detail pane displays the following information in an accordion for each ESXi host in the selected vSphere cluster:

- Host name on the storage system - If the ESXi host is defined on the selected storage system, this is what it is identified by.
- Status on the storage system - If the ESXi host is defined on the selected storage system, this reports the status of the host. Otherwise it shows as "Undefined".
- Host Cluster Name - If the host is defined and part of a host cluster, this will show the name of the host cluster on the selected storage system.
- Site Name - If the host is defined and is associated with a site on the selected storage system.
- vVol Supported - If the host is defined with a host type which supports "vVol".
- Port Count - The number of adapter ports associated with the defined host.

If there is a mismatch between the environment in the vSphere Client and the selected storage system, an error is raised with a link to the storage system to resolve.

ESXi host summary and configuration

Use the vSphere plug-in to display information about the ESXi host's representation on the IBM Storage Ceph Plugin for vSphere dashboard.

Viewing host summary information

In the vSphere Client inventory, identify the ESXi host that you want to explore.

To view storage information for an ESXi server, navigate to the summary tab for the selected ESXi host and scroll to the **IBM Storage card**. This view shows the ESXi host's status that is reported from each storage system that is registered in the plug-in dashboard. To see more details or refresh the connection, click the **Go to Host Connections** link at the bottom of the card. Alternatively, navigate to the **Configure** tab within the vSphere cluster, and scroll to **IBM Storage > Host Connections**.

Note: When launching this view, a discovery task is triggered to refresh up-to-date information from each registered storage system. The ESXi host's WWPNs reported by vSphere will be compared to both existing host definitions and any visible host ports on the storage system an SAN fabric. Depending on the size of your environment, this can take some time.

Host connections panel

The datagrid on the **Host Connections** view shows the visibility of the selected ESXi host to every storage system registered in the plug-in dashboard.

- **Storage System Name** - The name of the storage system taken directly from the system.
- **Alias** - The user defined alias name for the storage system.

- **Storage system status** - The status of the storage system as reported in the main dashboard.
- **Storage system's visibility of ESXi host** - If the ESXi host is defined on the storage system, the associated host status will be displayed. Otherwise it will report "Undefined".

To define the **Undefined** ESXi host on the storage system, the following actions can be performed from within the ESXi Host configuration panel:

- [Creating Host definition from a selected ESXi host](#)

Note: If there is no alias for the IBM Storage Ceph system, the alias column is automatically hidden. You can alter column's visibility using the icon in the bottom left of the datagrid.

To view further details for a given host, expand the details pane using the double arrow (>>) on the left-hand side of each row.

The detail pane displays the following information:

- **ESXi host name** - The selected vSphere ESXi hostname as reported in vSphere.
- **Host name on the storage system** - If the ESXi host is defined on the selected storage system, this is what it is identified by.
- **Status on the storage system** - The status is determined based on the list of hosts that are added to the subsystem. If the given ESXi host is on the list, the status is `Online`, otherwise, the status is `Undefined`.
- **Host Cluster Name** - If the host is defined and part of a host cluster, this will show the name of the host cluster on the selected storage system.
- **Site Name** - If the host is defined and is associated with a site on the selected storage system.
- **vVol Supported** - If the host is defined, and has a type "adminlun or vVol". This is required when using vSphere Virtual Volumes (vVol).
- **Adapter identifiers** - The associated NVMe and subsystem NQN, as reported by vSphere.

If the host is not defined on the selected storage system, you can define that by clicking on **Create Host** that is visible on the **Detail** pane. See [Creating Host definition from a selected ESXi host](#)

Creating ESXi host definitions

Learn how to create host definitions for selected ESXi servers on IBM Storage Ceph from within the vSphere Client.

Before you begin

In the vSphere Client inventory, identify the ESXi server that you wish to create. Ensure any storage adapter drivers have been installed on the ESXi server and that the storage adapter is visible in the **Configure > Storage Adapters** view in the vSphere Client. Ensure that NVMe/TCP connections are completed.

About this task

This process evaluates any configured storage adapter on the selected ESXi host, and displays any storage adapter WWPNs that are visible in the output of the **lsfcportcandidate** command on the storage system. Only host storage adapter WWPNs that are visible from the storage system is available to create a host definition. Hosts are created with the vVol host type by default to enable support for vSphere Virtual Volumes.

During this process you have to enter the name of the host object to be created on the storage system which, is auto-populated with the name as it appears in the vSphere Client. However, if the ESXi host is registered in vSphere with an IP address (rather than fully qualified domain name (FQDN) or hostname), you are prompted to manually define an appropriate name. This is due to host object naming conventions within the storage system.

Note: When launching this view, a discovery task is triggered to refresh up-to-date information from each registered storage system. The ESXi host's WWPNs reported by vSphere will be compared to both existing host definitions and any visible host ports on the storage system an SAN fabric. Depending on the size of your environment, this can take some time.

Procedure

1. Select the ESXi host in the vSphere Client inventory
2. Select the **Configure** tab.
3. Select **IBM Storage > Host Connections > .**
4. Select the IBM Storage Ceph system that you wish to create the host definition.
5. Select **Add Host**.
6. Define a name for the new host object on the storage system (if required).
7. Select the WWPNs from the list that you want to use for the host definition on the storage system.
8. Click **Add**.
9. Monitor the **Recent Tasks** view in vSphere to confirm the task is completed successfully.

Troubleshooting

Use this information to help detect and solve problems that you might encounter when using the vSphere Plugin.

Collecting a snap for support engagement

The vSphere Plugin maintains log files that record different types of events.

Collect logs to help troubleshoot issues when using the vSphere Plugin or when contacting IBM Support.

Generating a snap log bundle

Access the Command Line Interface (CLI) of the plug-in appliance by either connecting with SSH to the appliance, or accessing the plug-in appliance's console view from within vSphere Client. Use the **ibm-plugin** command to review plug-in status, register/unregister vCenters and generate a snap log bundle.

```
root@ibm-storage-plugin [ ~ ]# ibm-plugin -h
usage: ibm-plugin [-h] {status,snap,register,unregister} ...
```

Run this command to register, unregister or get the status of a remote vCenter Plugin.

```
positional arguments:
{status,snap,register,unregister}
status Get plugin version and registration status
snap Create a snap containing relevant files for debug
register Register the plugin to vCenter
unregister Unregister the plugin from vCenter
```

```
options:
-h, --help show this help message and exit
```

Use the **ibm-plugin snap** command to generate the snap bundle when engaging with IBM support.

```
root@ibm-storage-plugin [ ~ ]# ibm-plugin snap
Snap file created at /tmp/snap.ibm-plugin.230601.161526.tgz
```

Use, for example, **cp** to copy the file to a local machine.

Copying the plug-in snap file to a IBM Storage System

Another way of accessing the log bundle would be to copy the snap file directly from the plug-in appliance to a IBM Storage Ceph storage system using the **scp** command. Download the file from the storage system using the web browser.

To copy the snap file, use a user account with access to the IBM Storage system (in this example the user account is **ibmadmin**):

```
scp /tmp/<snap file name>.tgz ibmadmin@<FQDN or IP of storage system>:/dumps/
```

```
root@ibm-storage-plugin [ ~ ]# scp /tmp/snap.ibm-plugin.230601.161526.tgz ibmadmin@vvolsftw-af7.ssd.hursley.ibm.com:/dumps/
(ibmadmin@vvolsftw-af7.ssd.hursley.ibm.com) Password:
snap.ibm-
plugin.230601.161526.tgz
      100% 8051 8.8MB/s 00:00
root@ibm-storage-plugin [ ~ ]#
```

Once the file has been copied,

1. Access the management interface of the IBM storage system from a web browser, and select **Settings > Support > Support Package**.
2. Select **Download Support Package > Download Existing Package**.
3. Filter the list of files by entering the filename generated by the **ibm-plugin snap** command, e.g. **snap.ibm-plugin** and click **Download**.

Reviewing logs locally on the plug-in appliance

To review the log files locally on the plug-in appliance:

1. Navigate to `/opt/persistent` directory.
2. Use the **ls | grep log** Linux command.

vSphere task names misreported

The IBM Storage Ceph Plugin for vSphere introduces a number of dedicated tasks into vSphere. These tasks detail the progress of storage actions and provide extra visibility of storage related tasks to the vSphere administrator.

In certain circumstances, the task name does not get correctly reported in the vSphere client and instead, the task ID is shown.

Browser caching

It may be possible to correct these issues by either simply refreshing the browser window, or logging out of the vSphere Client session, and re-authenticating in a new window.

Restarting the vSphere client service in vCenter

If restarting the local browser session has not resolved the issue, it might be required to restart the vSphere Client Service from the vCenter Server Appliance Management Interface (VAMI).

Restarting the service via VAMI:

1. In a web browser, go to the vCenter Server Appliance Management Interface on port 5480, for example. `https://appliance-IP-address-or-FQDN:5480`.
2. Log in as root. The default root password is the password that you set while deploying the vCenter Server Appliance.
3. Select **Services**.
4. Select the **VMware vSphere Client** service and click **Restart**.

Note: If the issue is still not resolved, a restart of the vCenter Server Appliance might be required.

Storage devices refresh timeout

Use this information to help with general troubleshooting for storage devices refresh timeouts for IBM Storage Ceph Plugin for vSphere operations.

In certain circumstances, vCenter can encounter storage devices refresh timeouts for IBM plug-in operations.

Rescan timeouts can arise in vSphere environments of varying sizes. The most likely reasons can be any of the following:

1. Inaccessible datastore - If objects like virtual machines, templates, and snapshots are stored on the VMFS datastore but the underlying volume is removed from the storage system.
2. In large-scale vSphere deployments, the ESXi host or vSphere Cluster has access to many volumes on the storage system. However, only on a few volumes have datastores that are created, and others remain unused. Some of the unused volumes on the storage system might be offline.

3. Stale iSCSI adapters.
4. Hardware incompatibility.

Note: A timeout occurs due to multiple reasons. The list does not cover the entire range of possible causes.

Following are the list of possible errors that are observed because of rescan timeout:

1. Permanent Device Loss (PDL)
 - a. Matches for "offset due to Permanent Device Loss (PDL)" in /var/log/vmkernel.log.
 - b. A datastore is shown as **unavailable** in the Storage view.
 - c. A storage adapter indicates the Operational State of the device as Lost Communication.
 - d. All paths to the device are marked as Dead.
2. Error "state in doubt; requested fast path state update"
 - a. Matches for "state in doubt; requested fast path state update" in /var/log/vmkernel.log
 - b. Rescan timeouts can appear when the Host Bus Adapter (HBA) driver aborts a command for various reasons such as:
 - i) The command took longer than the timeout period to complete.
 - ii) General overload on the array
 - iii) Fabric issues (Bad ISL, outdated firmware, bad fabric cable/GBIC)

Check hardware incompatibility within vSphere client.

1. In the vSphere Client, navigate to the **ESXi** host.
2. Click the **Updates** tab.
3. Under **Hosts**, check **Hardware Compatibility**.

To resolve the errors listed, take the following steps:

1. In the vSphere Client, navigate to the **ESXi** host.
2. Click the **Configure** tab.
3. Under **Storage**, click **Storage Devices**. All volumes available to the ESXi host are listed in the Storage Devices table. The total number of volumes available to the host is displayed directly under the table. If the unused volumes are many or have inaccessible datastores, the storage refresh or a rescan most likely results in a timeout error for any operation. A recent task error message is raised: "Timeout in storage rescan happened, device inventory cleanup might be required on all vSphere hosts in the cluster"
4. To speed up the rescan operation, delete the volume from the storage system, which permanently detaches the devices that the ESXi host does not see as a volume.



Warning: Before deleting any volumes from the storage system, check with the Storage Administrator as deletion is a permanent action.

5. In the **Storage Devices** table, look for the Datastore column. If it says **Not Consumed**, then it means that there is no VMFS datastore on it.
6. Check the identifier of a device (in the format naa.*) in the **Name** column or by selecting device in the Storage Devices table. Copy the number after naa. This ID must match the VDisk UUID present on the storage system.

Note: To learn more about how to check VDisk UUID refer to your product documentation.

7. Delete the volume on the storage system if no VMFS datastore is created on the volume.



Warning: Before deleting any volumes from the storage system, check with the Storage Administrator as deletion is a permanent action.

8. Perform a rescan on all ESXi hosts that had visibility of the volume. The device is automatically removed from the Storage Adapters.

Note: Even if the devices are detached from the ESXi host, the volumes remains accessible to the ESXi host as long as the volumes are not removed from the storage system.

9. Repeat steps these steps for all unused volumes.

Miscellaneous troubleshooting

IBM Storage Ceph Plugin for vSphere is built with Photon OS Linux 4.0.

Displaying logs for IBM Storage Ceph Plugin for vSphere

From the plugin virtual machine, run the following command:

```
tail -F /opt/persistent/vsphere-plugin.log
```

Displaying logs for IBM Storage Ceph

From the Ceph node that is running an active ceph-mgr, run the following command:

```
journalctl -f _COMM=ceph-mgr
```

Pinging the virtual machine

Use this information to help with general troubleshooting, or alternatively refer to Photon OS documentation. <https://vmware.github.io/photon/docs-v4/>

By default, the plug-in VM appliance does not respond to ping (ICMP) requests as it is protected by a security policy complete with firewall. This firewall drops all ping requests by default. However, you can enable ping responses:

1. SSH into the plug-in appliance as root or connect to the VM console within the vSphere Client.
2. Run **iptables -A OUTPUT -p icmp -j ACCEPT** command.
3. Run **iptables -A INPUT -p icmp -j ACCEPT** command.

Changing network configuration

Change the network configuration on the VM appliance.

Note: The status IP configuration might need to be updated manually.

1. Go to the vSphere instance that the virtual machine was deployed on.
2. Log in and within the inventory go to the virtual machine.
3. From the virtual machine **Summary** tab, click **Launch Web Console**.

A console in a new tab launches.

4. From the console, log in as root by using the virtual machine credentials.
5. Edit the network config file found in: `/etc/systemd/network/`.
The config file is called `10-gosc-eth0.network`.
6. Edit by using the vi or vim command **vi /etc/systemd/network/10-gosc-eth0.network**.
7. Complete the network changes by using the following command: **systemctl restart systemd-networkd**.
8. Verify that the virtual machine can be reached on its new network settings (for example, through SSH).

Note: If network setting is changed, reregistration might be needed. To reregister the vSphere Plugin, by first unregistering and then registering the plug-in.

- a. [“Unregistering the vSphere Plugin” on page 25](#)
- b. [“Registering vSphere plug-in into vCenter ” on page 23](#)

Related information and publications

You can find additional information and publications that are related to the IBM Storage Ceph Plugin for vSphere at the following websites:

- [IBM Storage Ceph documentation](#)
- [VMware vSphere Documentation](#)
- [VMware Documentation](#)
- [VMware Product Support](#)
- [VMware Knowledge Base](#)

