Virtual SAN with Photon Controller Setup Guide

Version 1.2

This document supports the version of each product listed and supports all subsequent versions until the document is replaced by a new edition. To check for more recent editions of this document, see http://www.vmware.com/support/pubs.

EN-002523-01

vmware®

You can find the most up-to-date technical documentation on the VMware Web site at:

http://www.vmware.com/support/

The VMware Web site also provides the latest product updates.

If you have comments about this documentation, submit your feedback to:

docfeedback@vmware.com

VMware, Inc.
3401 Hillview Ave.
Palo Alto, CA 94304
www.vmware.com

Contents

About VMware Virtual SAN with Photon Controller

VMware Virtual SAN with Photon Controller Setup Guide describes how to install and configure VMware Virtual SAN on a Photon Controller cluster. This guide explains the requirements of the Photon Controller cluster, how to set up Virtual SAN on the cluster, and how to manage the cluster.

Intended Audience

This information is for software developers and virtualization administrators who are familiar with virtualization technology and data center storage concepts.

Updated Information 1

Virtual SAN with Photon Controller Setup Guide is updated with each release of the product or when necessary.

This table provides the update history of *Virtual SAN with Photon Controller Setup Guide*.

Revision	Description
EN-002523-01	■ Updated architecture diagram. ■ Updated the high-level tasks used to deploy a Photon Controller cluster. ■ Added step to verify cluster configuration on ESXI hosts. ■ Additional minor revisions.
EN-002523-00	Initial release.

Introduction to Virtual SAN with Photon Controller

2

VMware Virtual SAN™ is a distributed software layer within the ESXi hypervisor. It aggregates local or direct-attached capacity devices to create a single storage pool shared across all hosts in the Virtual SAN cluster. Virtual SAN can support the data storage needs for applications running on a Photon™ Platform.

Photon Platform

VMware Photon Platform is a distributed, high-scale container runtime and management platform that supports cloud-native applications. This platform is for large-scale deployments of application containers. You can use Photon Platform to implement a highly available, resilient infrastructure, to support a high degree of change and high-scale operations. The technical advantages of ESXi enable Photon Platform to integrate with Virtual SAN.

Photon Platform provides Kubernetes as a Service for deployment, maintenance, and scaling of applications. Kubernetes is an open-source platform for automating deployment, scaling, and operations of application containers across clusters of hosts, providing a container-centric infrastructure. You can use Kubernetes to manage containerized applications across your cloud hosts.

Photon Platform consists of a control plane called the VMware Photon Controller and a lightweight computing host. Each physical host in a Photon Controller cluster contains an ESXI host, and the Photon Controller is used to manage ESXi hosts.

Photon Controller

VMware Photon Controller is an open-source host controller that you can use to manage hardware, containers, VMs, and host clusters. It provides a distributed, high-scale control plane optimized for cloud-native applications, which include containers and developer frameworks, such as Pivotal Cloud Foundry (PCF).

Photon Controller combines ESXi with an innovative control plane to deliver highly scalable infrastructure as a service. It is designed to handle hundreds to thousands of hosts, and many thousands of tenant VMs.

Virtual SAN

Virtual SAN can provide a common image datastore for a Photon Controller cluster. Virtual SAN is implemented directly in the ESXi hypervisor.

Virtual SAN uses a software-defined approach that creates shared storage for virtual machines and cloud native applications. It virtualizes the local physical storage resources of ESXi hosts and turns them into pools of storage. You can divide and assign these pools to VMs and applications according to their quality of service requirements.

For more information about Virtual SAN, see *Administering VMware Virtual SAN* version 6.2.

Architecture Overview

A Photon Platform site includes a management infrastructure and a cloud infrastructure that contains one or more shared image datastores. Virtual SAN provides a common datastore to support the requirements of applications within a Photon Controller cluster. Virtual SAN can support the storage needs of all hosts in the Photon Controller cluster. You also can configure a Virtual SAN cluster to support only the management hosts or only the cloud hosts.

The management infrastructure of a Photon Platform site consists of one or more ESXi hosts connected to a dedicated management network and a shared datastore. Each management host runs a single management VM, which you can access through a management client. The Virtual SAN management service is a VM running on a single management host. You can use the Virtual SAN management service to create and manage a Virtual SAN cluster.

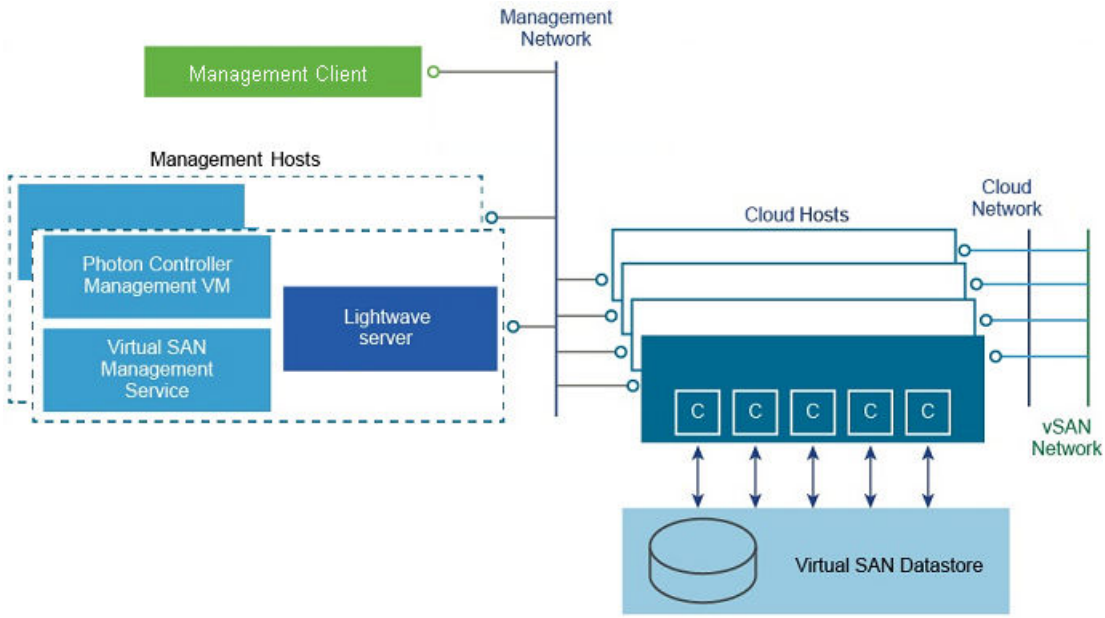

The cloud infrastructure of a Photon Platform site consists of one or more ESXi hosts connected to a dedicated cloud network, and one or more shared image datastores. Administrators provide compute, network, and storage resources for their tenants. Tenants deploy VMs to cloud hosts to run their workloads.

The Photon Controller management hosts can reside inside or outside the Virtual SAN cluster. The Virtual SAN datastore can support the storage needs of the cloud hosts and any management hosts that reside within the Virtual SAN cluster.

Each host in a Photon Controller cluster must allow you to mount at least one storage volume used to hold uploaded images. These images are used as VM boot disks, or to provide the basis for ephemeral or persistent disks created by tenants. Image storage volumes can also be used to store tenant VMs. Virtual SAN can provide shared storage for both the management hosts and cloud hosts in a Photon Controller cluster.

Virtual SAN Management for Photon Platform

You can use the Virtual SAN management service to configure and manage Virtual SAN without vCenter Server®. You can deploy the Virtual SAN management service on an ESXi host. Then you can use RVC commands or the Virtual SAN SDK to configure and manage the Virtual SAN cluster.

Virtual SAN SDK

The Virtual SAN SDK for Java includes documentation, libraries, and code examples that developers can use to build solutions integrated into Virtual SAN. You can use the APIs to automate all aspects of installation, configuration, monitoring, and troubleshooting of Virtual SAN. All Virtual SAN management APIs are supported in a Photon Controller environment.

Ruby vSphere Console

The Ruby vSphere Console (RVC) is an interactive command-line console user interface, based on *RbVmomi Ruby interface to the vSphere API*. You can use RVC commands to create a Virtual SAN cluster, and to manage and troubleshoot the cluster. Most of the Virtual SAN RVC management commands are supported in a Photon Controller environment. RVC is included with the Virtual SAN management service. For more information about RVC, see *VMware Ruby vSphere Console Command Reference for Virtual SAN*.

To use RVC, you run the ssh command to access the IP address of the Virtual SAN management service and log in as a privileged user. No additional configuration is required to begin.

Lightwave

Lightwave™ provides identity, authentication, and authorization functions for Photon Platform. The Lightwave server resides in a VM on a management host.

Cluster Requirements 3

The Photon Controller cluster must meet the requirements for enabling Virtual SAN.

- All capacity devices, drivers, and firmware versions in your Virtual SAN configuration must be certified and listed in the Virtual SAN section of the *VMware Compatibility Guide*.
- A Virtual SAN cluster must contain a minimum of three hosts that contribute capacity to the cluster.
- A host that resides in a Virtual SAN cluster cannot participate in other clusters.

Hardware Requirements

The ESXi hosts in your Photon Controller cluster must meet the Virtual SAN hardware requirements.

All capacity devices, drivers, and firmware versions in your Virtual SAN configuration must be certified and listed in the Virtual SAN section of the *VMware Compatibility Guide*.

Table 3-1. Storage Device Requirements for Virtual SAN Hosts

Storage Component	Requirements
Cache	■ You must have one SAS or SATA solid state disk (SSD) or PCIe flash device. ■ Before the **Number of failures to tolerate**calculation, the flash cache must be at least 10% of the anticipated consumed capacity, without the protection copies. ■ vSphere Flash Read Cache must not use any of the flash devices reserved for Virtual SAN cache. ■ Cache flash devices must not be formatted with VMFS or another file system.
Virtual machine data storage	■ For hybrid group configuration, at least one SAS, NL-SAS, or SATA magnetic disk must be available. ■ For all-flash disk group configuration, at least one SAS or SATA SSD or PCIe flash device must be available.
Storage controllers	You must have one SAS or SATA host bus adapter (HBA), or a RAID controller that is in passthrough mode or RAID 0 mode.

Booting Requirements

The memory requirements for Virtual SAN depend on the number of disk groups and devices that are managed by the ESXi hypervisor.

Memory

Each host must contain a minimum of 32 GB of memory to accommodate the maximum number of disk groups and capacity devices per disk group.

Flash Boot Devices

Booting a Virtual SAN host from a USB device or SD card requires a boot device of at least 4 GB.

If the memory of the ESXi host is greater than 512 GB, you must boot the host from a SATADOM or disk device. Booting a Virtual SAN host from a SATADOM device requires a single-level cell (SLC) device. The boot device must be at least 16 GB.

Booting an ESXi 6.0 host from a USB device or SD card results in the system writing Virtual SAN trace logs to RAM disk. These logs are automatically offloaded to persistent media during a shutdown or system crash (PANIC). This method is the only way to support the handling of Virtual SAN traces when booting an ESXi host from a USB device or SD card. If a power failure occurs, Virtual SAN trace logs are not preserved.

Booting an ESXi 6.0 host from a SATADOM device results in the system writing Virtual SAN trace logs directly to the SATADOM device. The SATADOM device must meet the specifications in this guide.

Software Requirements

The components in your environment must meet the software version requirements for using Virtual SAN.

The ESXi hosts that participate in Virtual SAN clusters must be running ESXi 6.5.0a.

Networking Requirements

The network infrastructure and the networking configuration on the ESXi hosts must meet the minimum networking requirements for Virtual SAN.

Table 3-2. Networking Requirements for Virtual SAN

Networking Component	Requirement
Host bandwidth	Each host must have minimum bandwidth dedicated to Virtual SAN. ■ Dedicated 1 Gbps for hybrid configurations ■ Dedicated or shared 10 Gbps for all-flash configurations
Connection between hosts	Each host in the Virtual SAN cluster, regardless of whether it contributes capacity, must have a VMkernel network adapter for Virtual SAN traffic.
Host network	All hosts in your Virtual SAN cluster must be connected to a Virtual SAN Layer 2 or Layer 3 network.
Multicast	Multicast must be enabled on the physical switches and routers that handle Virtual SAN traffic along the Layer 2 path and (optionally) the Layer 3 path.
IPv4 and IPv6 support	The Virtual SAN network can support IPv4 or IPv6 addresses.

License Requirements

Using Virtual SAN in production environments requires a valid license for Virtual SAN.

Before installing and configuring a Virtual SAN with Photon Controller cluster, review the End User License Agreement. You must accept the license agreement before you can deploy the software OVA.

Deploy Virtual SAN on a Photon Controller Cluster

4

You can deploy the Virtual SAN management service on any management host, and create a Virtual SAN cluster in a Photon Platform environment.

You can configure a Photon Controller cluster before or after you deploy the Virtual SAN management service. You can deploy the Virtual SAN management service on a management host that resides inside or outside the cluster. For more information about Photon Controller installation and configuration, see the Photon Controller documentation.

These tasks are typically used to deploy a Photon Controller cluster with Virtual SAN:

1 Configure a network to support Virtual SAN.

2 Download the Photon Controller installation appliance.

3 Deploy the Photon Controller Management VM, Lightwave server, on Virtual SAN on management ESXi hosts.

4 Deploy the Photon Controller agent on cloud ESXi hosts.

5 Create the Virtual SAN cluster and datastore.

This guide describes the tasks required to deploy the Virtual SAN management service and configure a Virtual SAN cluster in a Photon Platform environment.

This chapter includes the following topics:

Configure the Network

Configure a network to support Virtual SAN traffic on the Photon Controller cluster. Use ESXCLI commands to set the network parameters.

Connect to each ESXi host and configure a port group. Add an IP interface to support Virtual SAN. All hosts participating in Virtual SAN must be connected to a single Layer 2 network and multicast must be enabled.

Prerequisites

Verify that the host cluster meets the networking requirements listed in this guide.

Procedure

1 Add a port group on the virtual switch to support Virtual SAN traffic.

```
esxcli network vswitch standard portgroup add –p VSAN_portgroup –v vSwitch0
```

2 Set the VLAN ID for the *VSAN_portgroup*. A dedicated VLAN carries multicast Virtual SAN traffic among participating hosts.

```
esxcli network vswitch standard portgroup set ––vlan–id 0–4095
```

3 Add an interface to support Virtual SAN traffic.

```
esxcli network ip interface add –p VSAN_portgroup –i vmk1
```

4 (Optional) Enable DHCP on the interface.

```
esxcli network ip interface ipv4 set –t dhcp –i vmk1
```

5 Tag the interface for Virtual SAN traffic.

```
esxcli network ip interface tag add –t VSAN –i vmk1
```

For more information about network requirements for Virtual SAN, refer to *Administering VMware Virtual SAN* version 6.2.

Deploy the Virtual SAN Management Service

You can deploy the Virtual SAN management service when you deploy a Photon Controller management VM.

Use the OVF Tool to install the Photon Controller Installer VM on a management host. Deploy Lightwave, Photon Controller management VM, and the Virtual SAN management service on a management host. Deploy the Photon Controller Agent on cloud hosts in the Photon Controller cluster.

Prerequisites

- Obtain the OVF file that contains the Photon Controller Installer VM.
- Obtain the OVA file that contains the Virtual SAN management service.

Procedure

1 Log in to a host that contains the OVF Tool.

2 At the command-line prompt, deploy the Photon Controller Installer VM.

```
ovftool ––acceptAllEulas ––noSSLVerify ––skipManifestCheck ––X:injectOvfEnv
––overwrite ––powerOffTarget ––powerOn ––diskMode=thin ––net:"NAT"="<network name>"
––datastore="<datastore name>" ––name="Photon Controller Installer" ––
prop:enable_syslog="true"
––prop:admin_password=<administrator password> <path to OVF>/installer–ova.ova
vi://root:<host password>@<host IP>
```

For example:

```
ovftool ––acceptAllEulas ––noSSLVerify ––skipManifestCheck ––X:injectOvfEnv
––overwrite ––powerOffTarget ––powerOn ––diskMode=thin ––net:"NAT"="VM Network"
––datastore="datastore1" ––name="Photon Controller Installer" ––prop:enable_syslog="true"
––prop:admin_password=vmware /Applications/photon/installer–ova.ova
vi://root:password@10.110.120.5
```

3 Log in to the Photon Controller Installer VM and edit the following file: deployment.yml

- In the Compute > Hypervisors section, add the ESXi hosts in your cluster.
- In the lightwave > controllers section, add or delete controllers to match your environment.

- In the photon > cloud section, add the cloud host IDs.
- In the photon > controllers section, add or modify the Photon Controller VMs.

4 Prepare the Virtual SAN management service OVA.

a Create the following target directory in the Photon Controller Installer VM.

```
/var/opt/vmware/photon/controller/appliances/vsan.ova-dir
```

b Extract the files from the Virtual SAN OVA into the target directory.

For example:

```
tar -xvf vsan.ova -C /var/opt/vmware/photon/controller/appliances/vsan.ova-dir/
```

c Create a deployment backup file.

For example:

```
cd /var/opt/vmware/photon/controller/appliances/vsan.ova-dir
          cp vsan.ovf vsan.ova.bak
```

d Verify that the files that compose the Virtual SAN management service are present in the target directory.

For example:

```
ls /var/opt/vmware/photon/controller/appliances/vsan.ova-dir
vsan-disk1.vmdk
vsan.cert
vsan.mf
vsan.ovf
vsan.ovf.bak
```

5 From the Photon Controller Installer VM, run the following command to deploy Lightwave, Photon Controller Management VM, and the Virtual SAN management service. This command also installs the Photon Controller Agent on the cloud hosts.

```
root@lw-photon [ /opt/vmware/bin ]# /opt/vmware/photon/controller/bin/photon-setup platform
install -config deployment-hybrid.yml
```

The Virtual SAN management service is running on the management host. You can issue RVC commands to create and manage a Virtual SAN cluster.

You can access the Virtual SAN management API or RVC at **rvc *"<user>@<domain>"@<management vm>*:8006**.

For example:

```
vcsa-03:~ # rvc "vsanadmin@mydomain.com"@vsan-mgmt-srvr.mydomain.com:8006

Install the "ffi" gem for better tab completion.

WARNING: Nokogiri was built against LibXML version 2.7.6, but has dynamically loaded 2.9.2

The authenticity of host 'vsan-mgmt-srvr.rainpole.com' can't be established.
Public key fingerprint is 8551ba8526ffa1450a1df00c276c1f0e46f3dc5661a1de3983b666cdbc8c2d5d.
Are you sure you want to continue connecting (y/n)? y

Warning: Permanently added 'vsan-mgmt-srvr.mydomain.com' (vim) to the list of known hosts

password:************
0 /
1 vsan-mgmt-srvr.mydomain.com/
```

Configure a Hybrid Virtual SAN Cluster

You can create and configure a hybrid Virtual SAN cluster on the Photon Controller cluster by using RVC commands. In hybrid clusters, flash devices are used for the cache layer and magnetic disks are used for the capacity layer.

The Virtual SAN management service is a VM running on a management host. Use RVC commands to create and configure a hybrid Virtual SAN cluster.

Procedure

1 Connect the Ruby vSphere Console (RVC) to the Virtual SAN management service.

Access the RVC as follows: **rvc *"<user>@<domain>"@<management vm>*:8006**.

For example:

```
vcsa-03:~ # rvc "vsanadmin@mydomain.com"@vsan-mgmt-srvr.mydomain.com:8006

Install the "ffi" gem for better tab completion.

WARNING: Nokogiri was built against LibXML version 2.7.6, but has dynamically loaded 2.9.2

The authenticity of host 'vsan-mgmt-srvr.rainpole.com' can't be established.
Public key fingerprint is 8551ba8526ffa1450a1df00c276c1f0e46f3dc5661a1de3983b666cdbc8c2d5d.
Are you sure you want to continue connecting (y/n)? y

Warning: Permanently added 'vsan-mgmt-srvr.rainpole.com' (vim) to the list of known hosts

password:************
0 /
1 vsan-mgmt-srvr.mydomain.com/
```

2 Navigate to the *computers* directory. The cluster configuration files are in this directory.

```
> cd /10.110.120.130/Global/computers/
```

3 To create the Virtual SAN cluster, run the RVC command, **cluster.create *<cluster_name>***.

```
/10.110.120.130/Global/computers> cluster.create demo
```

This step creates an empty Virtual SAN cluster that contains no hosts.

4 In the *computers* directory, verify that the Virtual SAN cluster exists. Note the cluster name and the cluster ID number.

```
/10.110.120.130/Global/computers> ls
0 demo (cluster): cpu 0 GHz, memory 0 GB
```

5 Add hosts to the cluster. To add one or more hosts, run the RVC command, **cluster.add_host *<cluster_name>* *<host_IP>* *<host_IP>* -u *<username>* -p *<password>***.

```
/10.110.120.130/Global/computers>
cluster.add_host 0 10.160.100.10 10.160.100.11 10.160.100.12 -u root -p password
......Done
......Done
......Done
```

6 Claim all available storage devices on a host by runiing the RVC command, **vsan.host_consume_disks *<host_ID>***.

```
/10.110.120.130/Global/computers/demo/hosts> vsan.host_consume_disks 0
......Done
```

7 Navigate to the *computers* directory, and verify that the cluster claimed the storage devices. To display disk statistics, run the RVC command, **vsan.disks_stats *<cluster_name>***.

```
/10.110.120.130/Global/computers> vsan.disks_stats demo
2016-05-23 09:55:28 +0000: Fetching VSAN disk info from esxi1 (may take a moment) ...
2016-05-23 09:55:28 +0000: Fetching VSAN disk info from esxi2 (may take a moment) ...
2016-05-23 09:55:28 +0000: Fetching VSAN disk info from esxi3 (may take a moment) ...
2016-05-23 09:55:42 +0000: Done fetching VSAN disk infos
+----------------------+--------+-------+------+-----------+--------+----------+---------+
|                      |        |       | Num  | Capacity  |        |          | Status  |
| DisplayName          | Host   | isSSD | Comp | Total     | Used   | Reserved | Health  |
+----------------------+--------+-------+------+-----------+--------+----------+---------+
| mpx.vmhba1:C0:T6:L0  | esxi1  | SSD   | 0    | 50    GB  | 0.00 % | 0.00 %   | OK (v3) |
| mpx.vmhba1:C0:T2:L0  | esxi1  | MD    | 0    | 59.39 GB  | 4.00 % | 3.44 %   | OK (v3) |
+----------------------+--------+-------+------+-----------+--------+----------+---------+
| mpx.vmhba1:C0:T5:L0  | esxi1  | SSD   | 0    | 50 GB     | 0.00 % | 0.00 %   | OK (v3) |
| mpx.vmhba1:C0:T1:L0  | esxi1  | MD    | 0    | 59.39 GB  | 0.40 % | 0.00 %   | OK (v3) |
+----------------------+--------+-------+------+-----------+--------+----------+---------+
| mpx.vmhba1:C0:T5:L0  | esxi2  | SSD   | 0    | 50 GB     | 0.00 % | 0.00 %   | OK (v3) |
| mpx.vmhba1:C0:T1:L0  | esxi2  | MD    | 0    | 59.39 GB  | 0.40 % | 0.00 %   | OK (v3) |
+----------------------+--------+-------+------+-----------+--------+----------+---------+
| mpx.vmhba1:C0:T6:L0  | esxi2  | SSD   | 0    | 50 GB     | 0.00 % | 0.00 %   | OK (v3) |
| mpx.vmhba1:C0:T2:L0  | esxi2  | MD    | 0    | 59.39 GB  | 0.40 % | 0.00 %   | OK (v3) |
+----------------------+--------+-------+------+-----------+--------+----------+---------+
| mpx.vmhba1:C0:T5:L0  | esxi3  | SSD   | 0    | 50 GB     | 0.00 % | 0.00 %   | OK (v3) |
| mpx.vmhba1:C0:T1:L0  | esxi3  | MD    | 0    | 59.39 GB  | 0.40 % | 0.00 %   | OK (v3) |
+----------------------+--------+-------+------+-----------+--------+----------+---------+
| mpx.vmhba1:C0:T6:L0  | esxi3  | SSD   | 0    | 50 GB     | 0.00 % | 0.00 %   | OK (v3) |
| mpx.vmhba1:C0:T2:L0  | esxi3  | MD    | 0    | 59.39 GB  | 0.40 % | 0.00 %   | OK (v3) |
+----------------------+--------+-------+------+-----------+--------+----------+---------+
```

8 Check the cluster health. Navigate to the *computers* directory, and run the following RVC command to display health information: **vsan.health.health_summary <*cluster_name*>**.

For example:

```
/10.110.120.130/Global/computers> vsan.health.health_summary demo
Overall health: green
+-----------------------------------------------------+---------+
| Health check                                        | Result  |
+-----------------------------------------------------+---------+
| Network                                             | Error   |
|   Hosts with connectivity issues                    | Passed  |
|   Virtual SAN cluster partition                     | Passed  |
|   Unexpected Virtual SAN cluster members            | Passed  |
|   Hosts with Virtual SAN disabled                   | Passed  |
|   All hosts have a Virtual SAN vmknic configured    | Passed  |
|   All hosts have matching subnets                   | Passed  |
|   All hosts have matching multicast settings        | Passed  |
|   Multicast assessment based on other checks        | Passed  |
|   VSAN: Basic (unicast) connectivity check          | Passed  |
|   VSAN: MTU check (ping with large packet size)     | Passed  |
```

```
|   vMotion: Basic (unicast) connectivity check       | Passed  |
|   vMotion: MTU check (ping with large packet size)  | Passed  |
|   Active multicast connectivity check               | Passed  |
+-----------------------------------------------------+---------+
| Cluster                                             | Passed  |
|   ESX Virtual SAN Health service installation       | Passed  |
|   Virtual SAN Health Service up-to-date             | Passed  |
...
```

The health summary shows the status of the Virtual SAN health checks and provides an overall health status for the cluster.

9 (Optional) To verify that the Virtual SAN cluster is configured correctly, run the following RVC command on each ESXi host: **esxcli vsan cluster get**.

For example:

```
/10.110.120.130/Global/computers> esxcli vsan cluster get
Cluster Information
   Enabled: true
   Current Local Time: 2017-04-21T03:17:52Z
   Local Node UUID: 58f9767d-0144-4fe5-83ea-02000a8b22b8
   Local Node Type: NORMAL
   Local Node State: BACKUP
   Local Node Health State: HEALTHY
   Sub-Cluster Master UUID: 58f9767e-12f4-c4cb-75df-02000aa6a118
   Sub-Cluster Backup UUID: 58f9767d-0144-4fe5-83ea-02000a8b22b8
   Sub-Cluster UUID: a3e85f52-2078-434b-9e6b-1494ede19bdf
   Sub-Cluster Membership Entry Revision: 2
   Sub-Cluster Member Count: 3
   Sub-Cluster Member UUIDs: 58f9767e-12f4-c4cb-75de-02000ab6a118,
       58f9767d-0144-4fe5-83ea-02000a8a22b8, 58f9767f-d56f-a2f7-ecbd-02000ac3cf44
   Sub-Cluster Membership UUID: 3e77f958-7ab5-5f89-5a9a-02000ab6a118
```

Configure an All-Flash Virtual SAN Cluster

You can create and configure an all-flash Virtual SAN cluster on the Photon Controller cluster by using RVC commands. In all-flash clusters, flash devices are used for the cache tier and the capacity tier.

The Virtual SAN management service is a VM running on a management host. Use RVC commands to create and configure an all-flash Virtual SAN cluster.

Prerequisites

Verify that you have a dedicated or shared 10-Gbps network to support the all-flash Virtual SAN cluster.

Procedure

1 Connect the Ruby vSphere Console (RVC) to the Virtual SAN management service.

Access the RVC as follows: **rvc *"<user>@<domain>"@<management vm>*:8006**.

For example:

```
vcsa-03:~ # rvc "vsanadmin@mydomain.com"@vsan-mgmt-srvr.mydomain.com:8006

Install the "ffi" gem for better tab completion.

WARNING: Nokogiri was built against LibXML version 2.7.6, but has dynamically loaded 2.9.2

The authenticity of host 'vsan-mgmt-srvr.rainpole.com' can't be established.
```

```
Public key fingerprint is 8551ba8526ffa1450a1df00c276c1f0e46f3dc5661a1de3983b666cdbc8c2d5d.
Are you sure you want to continue connecting (y/n)? y

Warning: Permanently added 'vsan-mgmt-srvr.mydomain.com' (vim) to the list of known hosts

password:************
0 /
1 vsan-mgmt-srvr.mydomain.com/
```

2 Navigate to the *computers* directory. The cluster configuration files are in this directory.

```
> cd /10.110.120.140/Global/computers/
```

3 To create the Virtual SAN cluster, run the RVC command, **cluster.create *<cluster_name>***.

```
/10.110.120.140/Global/computers> cluster.create demo
```

This step creates an empty Virtual SAN cluster that contains no hosts.

4 In the *computers* directory, verify that the Virtual SAN cluster exists. Note the cluster name and the cluster ID number.

```
/10.110.120.140/Global/computers> ls
0 demo (cluster): cpu 0 GHz, memory 0 GB
```

5 To add one or more hosts to the cluster, run the RVC command, **cluster.add_host *<cluster_name>* *<host_IP>* *<host_IP>* -u *<username>* -p *<password>***.

```
/10.110.120.140/Global/computers>
cluster.add_host 0 10.160.100.20 10.160.100.21 10.160.100.22 -u root -p password
......Done
......Done
......Done
```

6 Mark the disks for use in the Virtual SAN cluster.

a Navigate to the *hosts* directory and list the hosts. Note the unique ID assigned to each host.

```
/10.110.120.140/Global/computers> cd 1/hosts/
/10.110.170.117/Global/computers/demo/hosts> ls
0 esxi5 (host): cpu 2*1*2.80 GHz, memory 34.00 GB
1 esxi6 (host): cpu 2*1*2.80 GHz, memory 34.00 GB
2 esxi7 (host): cpu 2*1*2.80 GHz, memory 34.00 GB
```

b To display information about the storage devices attached to each host, run the RVC command, **vsan.disks_info *<host_ID>***.

```
/10.110.120.140/Global/computers/demo/hosts> vsan.disks_info 0
/10.110.120.140/Global/computers/demo/hosts> vsan.disks_info ./esxi1
2016-05-23 09:43:39 +0000: Gathering disk information for host esxi1
2016-05-23 09:43:44 +0000: Done gathering disk information
Disks on host esxi5:
+--------------------------+-------+-------+---------------------------------------+
| DisplayName              | isSSD | Size  | State                                 |
+--------------------------+-------+-------+---------------------------------------+
| Local VMware Disk        | SSD   | 50 GB | eligible                              |
| (mpx.vmhba1:C0:T2:L0)    |       |       |                                       |
| VMware Virtual Disk      |       |       | Checksum Enabled: false               |
+--------------------------+-------+-------+---------------------------------------+
| Local VMware Disk        | SSD   | 50 GB | eligible                              |
| (mpx.vmhba1:C0:T1:L0)    |       |       |                                       |
| VMware Virtual Disk      |       |       | Checksum Enabled: false               |
+--------------------------+-------+-------+---------------------------------------+
| Local VMware Disk        | MD    | 20 GB | ineligible (Existing partitions found |
| (mpx.vmhba1:C0:T0:L0)    |       |       |  on disk ‘mpx.vmhba1:C0:T0:L0’.)      |
|                          |       |       |                                       |
| VMware Virtual Disk      |       |       | Partition table:                      |
|                          |       |       | 2: 4.00 GB, type = vfat               |
|                          |       |       | 5: 0.24 GB, type = vfat               |
|                          |       |       | 6: 0.24 GB, type = vfat               |
|                          |       |       | 7: 0.11 GB, type = coredump           |
|                          |       |       | 8: 0.28 GB, type = vfat               |
|                          |       |       | 9: 2.50 GB, type = coredump           |
|                          |       |       |                                       |
|                          |       |       | Checksum Enabled: false               |
+--------------------------+-------+-------+---------------------------------------+
| Local VMware Disk        | SSD   | 50 GB | eligible                              |
| (mpx.vmhba1:C0:T3:L0)    |       |       |                                       |
| VMware Virtual Disk      |       |       |                                       |
|                          |       |       | Checksum Enabled: false               |
+--------------------------+-------+-------+---------------------------------------+
```

c On each host, run the following RVC command to mark a storage device for use as cache or capacity: **vsan.host_claim_disks_differently *<host_ID>* -d *<device_ID>* -c *<device_type>***.

```
/10.110.120.140/Global/computers/demo/hosts> vsan.host_claim_disks_differently 0
                                             -d mpx.vmhba1:C0:T1:L0 -c capacity_flash
2016-05-11 07:22:30 +0000: Gathering disk information for host esxi5
2016-05-11 07:22:38 +0000: Done gathering disk information
Disks on host esxi5:
Claiming Local VMware Disk (mpx.vmhba1:C0:T1:L0) as capacity_flash ...
```

7 To claim all available storage devices on a host, run the RVC command, **vsan.host_consume_disks *<host_ID>***.

```
/10.110.120.140/Global/computers/demo/hosts> vsan.host_consume_disks 0
......Done
```

8 Navigate to the *computers* directory, and verify that the storage devices claimed by the cluster. To display disk statistics, run the RVC command, **vsan.disks_stats *<cluster_name>***.

```
/10.110.120.140/Global/computers> vsan.disks_stats demo
2016-05-23 09:55:28 +0000: Fetching VSAN disk info from esxi5 (may take a moment) ...
2016-05-23 09:55:28 +0000: Fetching VSAN disk info from esxi6 (may take a moment) ...
2016-05-23 09:55:28 +0000: Fetching VSAN disk info from esxi7 (may take a moment) ...
2016-05-23 09:55:42 +0000: Done fetching VSAN disk infos
+----------------------+--------+-------+------+-----------+--------+----------+---------+
|                      |        |       | Num  | Capacity  |        |          | Status  |
| DisplayName          | Host   | isSSD | Comp | Total     | Used   | Reserved | Health  |
+----------------------+--------+-------+------+-----------+--------+----------+---------+
| mpx.vmhba1:C0:T6:L0  | esxi5  | SSD   | 0    | 50    GB  | 0.00 % | 0.00 %   | OK (v3) |
| mpx.vmhba1:C0:T2:L0  | esxi5  | MD    | 0    | 59.39 GB  | 4.00 % | 3.44 %   | OK (v3) |
+----------------------+--------+-------+------+-----------+--------+----------+---------+
| mpx.vmhba1:C0:T5:L0  | esxi5  | SSD   | 0    | 50 GB     | 0.00 % | 0.00 %   | OK (v3) |
| mpx.vmhba1:C0:T1:L0  | esxi5  | MD    | 0    | 59.39 GB  | 0.40 % | 0.00 %   | OK (v3) |
+----------------------+--------+-------+------+-----------+--------+----------+---------+
| mpx.vmhba1:C0:T5:L0  | esxi6  | SSD   | 0    | 50 GB     | 0.00 % | 0.00 %   | OK (v3) |
| mpx.vmhba1:C0:T1:L0  | esxi6  | MD    | 0    | 59.39 GB  | 0.40 % | 0.00 %   | OK (v3) |
+----------------------+--------+-------+------+-----------+--------+----------+---------+
| mpx.vmhba1:C0:T6:L0  | esxi6  | SSD   | 0    | 50 GB     | 0.00 % | 0.00 %   | OK (v3) |
| mpx.vmhba1:C0:T2:L0  | esxi6  | MD    | 0    | 59.39 GB  | 0.40 % | 0.00 %   | OK (v3) |
+----------------------+--------+-------+------+-----------+--------+----------+---------+
| mpx.vmhba1:C0:T5:L0  | esxi7  | SSD   | 0    | 50 GB     | 0.00 % | 0.00 %   | OK (v3) |
| mpx.vmhba1:C0:T1:L0  | esxi7  | MD    | 0    | 59.39 GB  | 0.40 % | 0.00 %   | OK (v3) |
+----------------------+--------+-------+------+-----------+--------+----------+---------+
| mpx.vmhba1:C0:T6:L0  | esxi7  | SSD   | 0    | 50 GB     | 0.00 % | 0.00 %   | OK (v3) |
| mpx.vmhba1:C0:T2:L0  | esxi7  | MD    | 0    | 59.39 GB  | 0.40 % | 0.00 %   | OK (v3) |
+----------------------+--------+-------+------+-----------+--------+----------+---------+
```

Flash devices claimed for the capacity tier are listed as MD.

9 Check the cluster health. Navigate to the *computers* directory, and use the following RVC command to display health information: **vsan.health.health_summary *<cluster_name>***

```
/10.110.120.140/Global/computers> vsan.health.health_summary demo
Overall health: green
+-----------------------------------------------------+---------+
| Health check                                        | Result  |
+-----------------------------------------------------+---------+
| Network                                             | Error   |
|   Hosts with connectivity issues                    | Passed  |
|   Virtual SAN cluster partition                     | Passed  |
|   Unexpected Virtual SAN cluster members            | Passed  |
|   Hosts with Virtual SAN disabled                   | Passed  |
|   All hosts have a Virtual SAN vmknic configured    | Passed  |
|   All hosts have matching subnets                   | Passed  |
|   All hosts have matching multicast settings        | Passed  |
|   Multicast assessment based on other checks        | Passed  |
|   VSAN: Basic (unicast) connectivity check          | Passed  |
|   VSAN: MTU check (ping with large packet size)     | Passed  |
```

```
|   vMotion: Basic (unicast) connectivity check       | Passed  |
|   vMotion: MTU check (ping with large packet size)  | Passed  |
|   Active multicast connectivity check               | Passed  |
+-----------------------------------------------------+---------+
| Cluster                                             | Passed  |
|   ESX Virtual SAN Health service installation       | Passed  |
|   Virtual SAN Health Service up-to-date             | Passed  |
...
```

The health summary shows the status of the Virtual SAN health checks and provides an overall health status for the cluster.

10 (Optional) To verify that the Virtual SAN cluster is configured correctly, run the following RVC command on each ESXi host: **esxcli vsan cluster get**.

```
/10.110.120.140/Global/computers> esxcli vsan cluster get
Cluster Information
   Enabled: true
   Current Local Time: 2017-04-21T03:17:52Z
   Local Node UUID: 58f9767d-0144-4fe5-83ea-02000a8b22b8
   Local Node Type: NORMAL
   Local Node State: BACKUP
   Local Node Health State: HEALTHY
   Sub-Cluster Master UUID: 58f9767e-12f4-c4cb-75df-02000aa6a118
   Sub-Cluster Backup UUID: 58f9767d-0144-4fe5-83ea-02000a8b22b8
   Sub-Cluster UUID: a3e85f52-2078-434b-9e6b-1494ede19bdf
   Sub-Cluster Membership Entry Revision: 2
   Sub-Cluster Member Count: 3
   Sub-Cluster Member UUIDs: 58f9767e-12f4-c4cb-75de-02000ab6a118,
       58f9767d-0144-4fe5-83ea-02000a8a22b8, 58f9767f-d56f-a2f7-ecbd-02000ac3cf44
   Sub-Cluster Membership UUID: 3e77f958-7ab5-5f89-5a9a-02000ab6a118
```

Managing the Virtual SAN Cluster 5

You can perform common management tasks on the Virtual SAN cluster, such as creating and deleting the Virtual SAN cluster.

This information helps you manage a Virtual SAN with Photon Controller. For more information about using RVC commands to manage a Virtual SAN cluster, see the *VMware Ruby vSphere Console Command Reference for Virtual SAN* documentation.

Create and Delete a Virtual SAN Cluster

When you first create a Virtual SAN cluster, the cluster is empty with no host or storage devices.

To create a Virtual SAN cluster, run the RVC command, **`cluster.create <cluster_name>`**.

For example:

```
cluster.create demo
```

To delete a Virtual SAN cluster, run the RVC command, **`cluster.delete <cluster_name>`**.

For example:

```
cluster.delete demo
```

This chapter includes the following topics:

Displaying Information About the Virtual SAN Cluster

Use RVC commands to display information about the performance and health of your Virtual SAN cluster.

Display Cluster Summary

You can display summary and health information about your Virtual SAN cluster.

To display summary information, run the RVC command, **`vsan.cluster_info <cluster_ID>`**.

Display Cluster Health

To display health information, run the RVC command, **`vsan.health.health_summary <cluster_ID>`**.

Managing Hosts in the Virtual SAN Cluster

Use RVC commands to manage hosts in your Virtual SAN cluster. You can display information, add hosts, and remove hosts.

Display Host Information

You can display information about hosts in your Virtual SAN cluster. Omit the host ID to display information about all hosts in the cluster.

To display host information, run the RVC command, **vsan.host_info *<host_ID>***.

Place a Host in Maintenance Mode

Before you shut down, reboot, disconnect, or remove a host that is a member of a Virtual SAN cluster, place the host in maintenance mode. Select a data evacuation mode, such as Ensure Accessibility or Evacuate All Data.

To place a host in maintenance mode, run the RVC command,
vsan.enter_maintenance_mode *<host_ID>* –v *<data evacuation mode>*.

For example:

```
vsan.enter_maintenance_mode 5 –v evacuateAllData
```

To take a host out of maintenance mode, run the RVC command, **vsan.exit_maintenance_mode *<host_ID>***.

For example:

```
host.exit_maintenance_mode 5
```

Add a Host

You can add hosts to your Virtual SAN cluster.

To add one or more hosts, run the RVC command,
cluster.add_host *<cluster_name>* *<host_ID>* *<host_ID>* –u *<username>* –p *<password>*.

For example:

```
cluster.add_host demo 5 –u root –p password
```

Remove a Host

Before you remove a host from your Virtual SAN cluster, place the host in maintenance mode.

To remove the host, run the RVC command,
cluster.remove_host *<cluster_name>* *<host_IP>* –u *<username>* –p *<password>*.

For example:

```
cluster.remove_host 5
```

Managing Storage Devices in the Virtual SAN Cluster

Use RVC and ESXCLI commands to manage capacity disks and cache devices in your Virtual SAN cluster. You can display information, add storage devices, and remove devices.

Display Disk Information

You can display information about the disks on hosts in the Virtual SAN cluster. Omit the host identifier to display disk information for all hosts in the cluster.

To display information about storage devices in the Virtual SAN cluster, run the RVC command, **vsan.disks_info *<host_ID>***.

Display Disk Statistics

You can display statistical information about the disks in the Virtual SAN cluster, including whether or not it is a magnetic disk or solid state drive, how many components reside on the disk, disk capacity, how much is used, if it's health is OK and so on.

To display disk statistics, run the RVC command, **vsan.disks_stats *<host_ID>***.

Add a Disk to a Disk Group

You can add local devices to disk groups on a host. The devices must be the same type as the existing devices in the disk groups, such as flash devices or magnetic disks.

When Virtual SAN is set to claim disks in manual mode, you can add capacity disks to a host in the cluster. To select the disk group where you want to add a capacity disk, enter the name of the disk group's caching device.

To add a disk to a host's disk group, run the ESXCLI command,
esxcli vsan storage add –s *<SSD_ID>* –d *<disk_ID>*.

For example:

```
esxcli vsan storage add –s mpx.vmhba1:C0:T2:L0 –d mpx.vmhba1:C0:T0:L0
```

Remove a Disk from a Disk Group

You can remove a capacity disk from a disk group on a host. Because removing unprotected devices might be disruptive for the Virtual SAN datastore and virtual machines in the datastore, avoid removing devices or disk groups.

When Virtual SAN is set to claim disks in manual mode, you can remove capacity disks from a host in the cluster. To select the disk group where you want to remove a capacity disk, enter the name of the disk group's caching device.

To remove a capacity disk from a host's disk group, run the ESXCLI command,
esxcli vsan storage remove –d *<MD_ID>*.

For example:

```
esxcli vsan storage remove –d mpx.vmhba1:C0:T0:L0
```

Replace a Cache Disk or a Disk Group

When you remove a cache device from a disk group, Virtual SAN removes the entire disk group. Because removing unprotected devices might be disruptive for the Virtual SAN datastore and virtual machines in the datastore, avoid removing devices or disk groups.

Enter the name of the disk group's cache device to remove the disk group. To replace a cache device, remove the cache device and then add the cache and it's associated capacity device.

To remove a cache device, run the ESXCLI command, **`esxcli vsan storage remove –s <SSD_ID>`**. To add the disk group, run the ESXCLI command **`esxcli vsan storage add –s <SSD_ID> –d <MD_ID>`**

For example:

```
esxcli vsan storage add –s naa.55cd2e404b9cfe4b –d naa.5000c50057ae7d3b
```

Security Reference 6

You can use the security features of Virtual SAN to safeguard your environment from attack.

Services, Ports, and External Interfaces

The operation of Virtual SAN depends on certain services, ports, and external interfaces which need to be open or enabled.

- Port 8006 provides an HTTPS end-point for the Virtual SAN Health Service.
- Port 8007 provides an HTTP end-point for local use. This port accepts connections only from IP 127.0.0.1

Certain network ports between ESXi hosts on the Virtual SAN network must be enabled.

- By default, Virtual SAN multicast uses the following addresses and ports:
  - Agent Group Multicast uses IP address 224.2.3.4 and port 23451
  - Master Group Multicast uses IP address: 224.1.2.3 and port 12345
- Virtual SAN uses TCP port 2233 for unicast, on Virtual SAN-enabled network interfaces on the ESXi nodes.

Resources

Resources that must be protected include security-related configuration files and passwords, and the recommended access controls for secure operation.

- The following file contains an encrypted password for *vpxuser*, used to communicate with ESXi hosts:

  `/etc/vmware-vsan-health/clusterstate.data`

Log Files

Log files that contain system messages are available.

The following Virtual SAN log files are located on ESXi hosts:

- `/var/log/clomd.log`
- `/var/log/epd.log`
- `/var/log/osfsd.log`
- The following log records client (SMS) registered certificate in base64 encoded text format:

  `/var/log/vsanvpd.log`

The vsanvpd.log also contains messages about successful or failed client requests. For example:

```
2014-09-10T01:00:21Z vsanSoapServer: run:141:Failed to accept client 10.114.174.101
[30]: SSL_ERROR_SSLerror:14094416:SSL routines:SSL3_READ_BYTES:sslv3 alert certificate unknown
2014-09-10T01:00:22Z vsanSoapServer: run:133:Serve accepted client 10.112.185.33
2014-09-10T01:00:22Z vsanSoapServer: verify_cert_with_store:811:Cannot verify cert with CA
store /etc/vmware/ssl/castore.pem: self signed certificate (18)
```

General operational information and security logs are located in the Virtual SAN Management Server:

```
/var/log/vmware/vsan-health/vmware-vsan-health-service.log
```

Additional Information

Virtual SAN does not create any additional accounts during installation and bootstrap.

Privileges assigned to *vsan-health* user only allow it to manage the life-cycle (start, stop, restart, status) of the Virtual SAN Health Service.

Virtual SAN with Photon Controller is released and distributed as part of Photon Controller. The security update and patch process is the same as for Photon Controller.

Index