



Configure your Virtual Storage Console for VMware vSphere environment

VSC, VASA Provider, and SRA 9.7

NetApp
June 11, 2024

This PDF was generated from <https://docs.netapp.com/us-en/vsc-vasa-provider-sra-97/deploy/reference-esx-host-values-set-by-vsc-for-vmware-vsphere.html> on June 11, 2024. Always check docs.netapp.com for the latest.

Table of Contents

- Configure your Virtual Storage Console for VMware vSphere environment 1
 - Configure ESXi server multipathing and timeout settings 1
 - Regenerate an SSL certificate for Virtual Storage Console 6
 - Requirements for registering VSC in multiple vCenter Servers environment 7
 - Configure the VSC preferences files 8
 - Enable datastore mounting across different subnets 9
 - Access the maintenance console options of the virtual appliance for VSC, VASA Provider, and SRA 10
 - Change the administrator password 12
 - Configure high availability for the virtual appliance for VSC, VASA Provider, and SRA 12
 - MetroCluster configurations supported by the virtual appliance for VSC, VASA Provider, and SRA 14

Configure your Virtual Storage Console for VMware vSphere environment

(VSC) supports numerous environments. Some of the features in these environments might require additional configuration.

You might have to perform some of the following tasks to configure your ESXi hosts, guest operating systems, and VSC:

- Verifying your ESXi host settings, including the UNMAP settings
- Adding timeout values for guest operating systems
- Regenerating the VSC SSL certificate
- Creating storage capability profiles and threshold alarms
- Modifying the preferences file to enable the mounting of datastores across different subnets

Configure ESXi server multipathing and timeout settings

Virtual Storage Console for VMware vSphere checks and sets the ESXi host multipathing settings and HBA timeout settings that work best with storage systems.

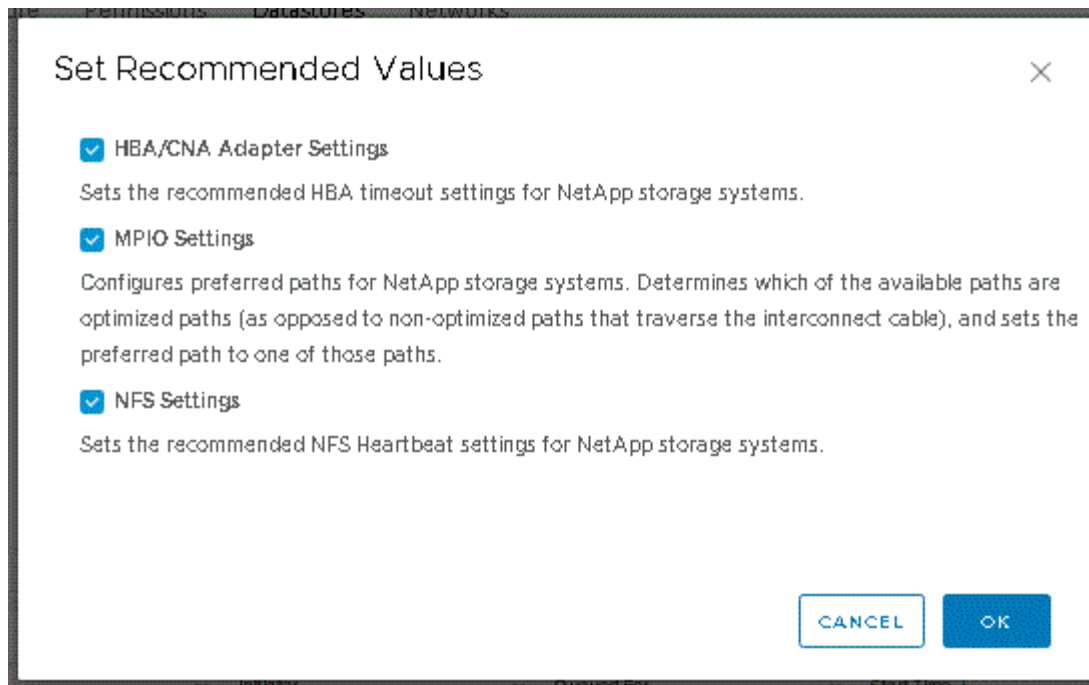
About this task

This process might take a long time, depending on your configuration and system load. The task progress is displayed in the **Recent Tasks** panel. As the tasks are completed, the host status Alert icon is replaced by the Normal icon or the Pending Reboot icon.

Steps

1. From the VMware vSphere Web Client **Home** page, click **vCenter > Hosts**.
2. Right-click a host, and then select **Actions > NetApp VSC > Set Recommended Values**.
3. In the **NetApp Recommended Settings** dialog box, select the values that work best with your system.

The standard, recommended values are set by default.



4. Click **OK**.

ESXi host values set using Virtual Storage Console for VMware vSphere

You can set timeouts and other values on the ESXi hosts using Virtual Storage Console for VMware vSphere to ensure best performance and successful failover. The values that Virtual Storage Console (VSC) sets are based on internal testing.

You can set the following values on an ESXi host:

ESXi advanced configuration

- **VMFS3.HardwareAcceleratedLocking**

You should set this value to 1.

- **VMFS3.EnableBlockDelete**

You should set this value to 0.

NFS settings

- **Net.TcpipHeapSize**

If you are using vSphere 6.0 or later, you should set this value to 32.

- **Net.TcpipHeapMax**

If you are using vSphere 6.0 or later, you should set this value to 1536.

- **NFS.MaxVolumes**

If you are using vSphere 6.0 or later, you should set this value to 256.

- **NFS41.MaxVolumes**

If you are using vSphere 6.0 or later, you should set this value to 256.

- **NFS.MaxQueueDepth**

If you are using the vSphere 6.0 or later version of ESXi host, then you should set this value to 128 or higher to avoid queuing bottlenecks.

For vSphere versions prior to 6.0, you should set this value to 64.

- **NFS.HeartbeatMaxFailures**

You should set this value to 10 for all NFS configurations.

- **NFS.HeartbeatFrequency**

You should set this value to 12 for all NFS configurations.

- **NFS.HeartbeatTimeout**

You should set this value to 5 for all NFS configurations.

FC/FCoE settings

- **Path selection policy**

You should set this value to “RR” (round robin) when FC paths with ALUA are used.

You should set this value to “FIXED” for all other configurations.

Setting this value to “RR” helps to provide load balancing across all of the active/optimized paths. The value “FIXED” is used for older, non-ALUA configurations and helps to prevent proxy I/O.

- **Disk.QFullSampleSize**

You should set this value to 32 for all configurations. Setting this value helps to prevent I/O errors.

- **Disk.QFullThreshold**

You should set this value to 8 for all configurations. Setting this value helps prevent I/O errors.

- **Emulex FC HBA timeouts**

Use the default value.

- **QLogic FC HBA timeouts**

Use the default value.

iSCSI settings

- **Path selection policy**

You should set this value to “RR” for all iSCSI paths.

Setting this value to “RR” helps to provide load balancing across all of the active/optimized paths.

- **Disk.QFullSampleSize**

You should set this value to 32 for all configurations. Setting this value helps to prevent I/O errors.

- **Disk.QFullThreshold**

You should set this value to 8 for all configurations. Setting this value helps prevent I/O errors.

Configure guest operating system scripts

The ISO images of the guest operating system (OS) scripts are mounted on the Virtual Storage Console for VMware vSphere server. To use the guest OS scripts to set the storage timeouts for virtual machines, you must mount the scripts from the vSphere Client.

Operating System Type	60-second timeout settings	190-second timeout settings
Linux	<code>https://<appliance_ip>:8143/vsc/public/writable/linux_gos_timeout-install.iso</code>	<code>https://<appliance_ip>:8143/vsc/public/writable/linux_gos_timeout_190-install.iso</code>
Windows	<code>https://<appliance_ip>:8143/vsc/public/writable/windows_gos_timeout.iso</code>	<code>https://<appliance_ip>:8143/vsc/public/writable/windows_gos_timeout_190.iso</code>
Solaris	<code>https://<appliance_ip>:8143/vsc/public/writable/solaris_gos_timeout-install.iso</code>	<code>https://<appliance_ip>:8143/vsc/public/writable/solaris_gos_timeout_190-install.iso</code>

You should install the script from the copy of the VSC instance that is registered to the vCenter Server that manages the virtual machine. If your environment includes multiple vCenter Servers, you should select the server that contains the virtual machine for which you want to set the storage timeout values.

You should log in to the virtual machine, and then run the script to set the storage timeout values.

Set timeout values for Windows guest operating systems

The guest operating system (OS) timeout scripts set the SCSI I/O timeout settings for Windows guest operating systems. You can specify either a 60-second timeout or a 190-second timeout. You must reboot the Windows guest OS for the settings to take effect.

Before you begin

You must have mounted the ISO image containing the Windows script.

Steps

1. Access the console of the Windows virtual machine, and log in to an account with Administrator privileges.
2. If the script does not automatically start, open the CD drive, and then run the `windows_gos_timeout.reg` script.

The Registry Editor dialog is displayed.

3. Click **Yes** to continue.

The following message is displayed: The keys and values contained in `D:\windows_gos_timeout.reg` have been successfully added to the registry.

4. Reboot the Windows guest OS.
5. Unmount the ISO image.

Set timeout values for Solaris guest operating systems

The guest operating system (OS) timeout scripts set the SCSI I/O timeout settings for Solaris 10. You can specify either a 60-second timeout or a 190-second timeout.

Before you begin

You must have mounted the ISO image containing the Solaris script.

Steps

1. Access the console of the Solaris virtual machine, and log in to an account with root privileges.
2. Run the `solaris_gos_timeout-install.sh` script.

For Solaris 10, a message similar to the following is displayed:

```
Setting I/O Timeout for /dev/s-a - SUCCESS!
```

3. Unmount the ISO image.

Set timeout values for Linux guest operating systems

The guest operating system (OS) timeout scripts set the SCSI I/O timeout settings for versions 4, 5, 6, and 7 of Red Hat Enterprise Linux and versions 9, 10, and 11 of SUSE Linux Enterprise Server. You can specify either a 60-second timeout or a 190-second timeout. You must run the script each time you upgrade to a new version of Linux.

Before you begin

You must have mounted the ISO image containing the Linux script.

Steps

1. Access the console of the Linux virtual machine, and log in to an account with root privileges.
2. Run the `linux_gos_timeout-install.sh` script.

For Red Hat Enterprise Linux 4 or SUSE Linux Enterprise Server 9, a message similar to the following is displayed:

```
Restarting udev... this may take a few seconds.
```

```
Setting I/O Timeout (60s) for /dev/sda - SUCCESS!
```

For Red Hat Enterprise Linux 5, Red Hat Enterprise Linux 6, and Red Hat Enterprise Linux 7 a message similar to the following is displayed:

```
patching file /etc/udev/rules.d/50-udev.rules
```

```
Hunk #1 succeeded at 333 (offset 13 lines).
```

```
Restarting udev... this may take a few seconds.
```

```
Starting udev: [ OK ]
```

```
Setting I/O Timeout (60s) for /dev/sda - SUCCESS!
```

For SUSE Linux Enterprise Server 10 or SUSE Linux Enterprise Server 11, a message similar to the following is displayed:

```
patching file /etc/udev/rules.d/50-udev-default.rules
```

```
Hunk #1 succeeded at 114 (offset 1 line).
```

```
Restarting udev ...this may take a few seconds.
```

```
Updating all available device nodes in /dev: done
```

3. Unmount the ISO image.

Regenerate an SSL certificate for Virtual Storage Console

The SSL certificate is generated when you install (VSC). The distinguished name (DN)

that is generated for the SSL certificate might not be a common name (CN) that the client machines recognize. By changing the keystore and private key passwords, you can regenerate the certificate and create a site-specific certificate.

About this task

You can enable remote diagnostic using the maintenance console and generate site-specific certificate.

[NetApp Knowledgebase Answer 1075654: Virtual Storage Console 7.x: Implementing CA signed certificates](#)

Steps

1. Log in to the maintenance console.
2. Enter 1 to access the Application Configuration menu.
3. In the Application Configuration menu, enter 3 to stop the VSC service.
4. Enter 7 to regenerate SSL certificate.

Requirements for registering VSC in multiple vCenter Servers environment

If you are using Virtual Storage Console for VMware vSphere in an environment where a single VMware vSphere HTML5 client is managing multiple vCenter Server instances, you must register one instance of VSC with each vCenter Server so that there is a 1:1 pairing between VSC and the vCenter Server. Doing this enables you to manage all of the servers running vCenter 6.0 or later in both linked mode and non-linked mode from a single vSphere HTML5 client.



If you want to use VSC with a vCenter Server, then you must have set up or registered one VSC instance for every vCenter Server instance that you want to manage. Each registered VSC instance must be of the same version.

Linked mode is installed automatically during the vCenter Server deployment. Linked mode uses Microsoft Active Directory Application Mode (ADAM) to store and synchronize data across multiple vCenter Server systems.

Using the vSphere HTML5 client to perform VSC tasks across multiple vCenter Servers requires the following:

- Each vCenter Server in the VMware inventory that you want to manage must have a single VSC server registered with it in a unique 1:1 pairing.

For example, you can have VSC server A registered to vCenter Server A, VSC server B registered to vCenter Server B, VSC server C registered to vCenter Server C, and so on.

You **cannot** have VSC server A registered to both vCenter Server A and vCenter Server B.

If a VMware inventory includes a vCenter Server that does not have a VSC server registered to it, but there are one or more vCenter Servers that are registered with VSC, then you can view the instances of VSC and perform VSC operations for the vCenter Servers that have VSC registered.

- You must have the VSC-specific View privilege for each vCenter Server that is registered to the single sign-on (SSO).

You must also have the correct RBAC permissions.

When you are performing a task that requires you to specify a vCenter Server, the **vCenter Server** drop-down box displays the available vCenter Servers in alphanumeric order. The default vCenter Server is always the first server in the drop-down list.

If the location of the storage is known (for example, when you use the **Provisioning** wizard and the datastore is on a host managed by a specific vCenter Server), the vCenter Server list is displayed as a read-only option. This happens only when you use the right-click option to select an item in the vSphere Web Client.

VSC warns you when you attempt to select an object that it does not manage.

You can filter storage systems based on a specific vCenter Server from the VSC summary page. A summary page appears for every VSC instance that is registered with a vCenter Server. You can manage the storage systems that are associated with a specific VSC instance and vCenter Server, but you should keep the registration information for each storage system separate if you are running multiple instances of VSC.

Configure the VSC preferences files

The preferences files contain settings that control Virtual Storage Console for VMware vSphere operations. Under most circumstances, you do not have to modify the settings in these files. It is helpful to know which preference files (VSC) uses.

VSC has several preference files. These files include entry keys and values that determine how VSC performs various operations. The following are some of the preference files that VSC uses:

```
/opt/netapp/vscserver/etc/kamino/kaminoprefs.xml
```

```
/opt/netapp/vscserver/etc/vsc/vscPreferences.xml
```

You might have to modify the preferences files in certain situations. For example, if you use iSCSI or NFS and the subnet is different between your ESXi hosts and your storage system, you have to modify the preferences files. If you do not modify the settings in the preferences file, datastore provisioning fails because VSC cannot mount the datastore.

Set IPv4 or IPv6

There is a new option added to the preference file `kaminoprefs.xml` that you can set to enable support for IPv4 or IPv6 for all storage systems added to VSC.

- The `default.override.option.provision.mount.datastore.address.family` parameter has been added to the `kaminoprefs.xml` preference file to set a preferred data LIF protocol for datastore provisioning.

This preference is applicable for all of the storage systems added to VSC.

- The values for the new option are `IPv4`, `IPv6`, and `NONE`.
- By default the value is set to `NONE`.

Value	Description
NONE	<ul style="list-style-type: none"> Provisioning happens using the same IPv6 or IPv4 address type of data LIF as the type of cluster or management LIF used for adding the storage. If the same IPv6 or IPv4 address type of data LIF is not present in the , then the provisioning happens through the other type of data LIF, if available.
IPv4	<ul style="list-style-type: none"> Provisioning happens using the IPv4 data LIF in the selected . If the does not have an IPv4 data LIF, then the provisioning happens through the IPv6 data LIF, if it is available in the .
IPv6	<ul style="list-style-type: none"> Provisioning happens using the IPv6 data LIF in the selected . If the does not have an IPv6 data LIF, then the provisioning happens through the IPv4 data LIF, if it is available in the .

Enable datastore mounting across different subnets

If you use iSCSI or NFS and the subnet is different between your ESXi hosts and your storage system, you have to modify the Virtual Storage Console for VMware vSphere preferences files. If you do not modify the preferences file, then datastore provisioning fails because (VSC) cannot mount the datastore.

About this task

When datastore provisioning fails, VSC logs the following error messages:

```
Unable to continue. No ip addresses found when cross-referencing kernel ip
addresses and addresses on the controller.
```

```
Unable to find a matching network to NFS mount volume to these hosts.``
```

Steps

1. Log in to your vCenter Server instance.
2. Launch the maintenance console using your unified appliance virtual machine.

[Access the maintenance console options of the virtual appliance for VSC, VASA Provider, and SRA](#)

3. Enter 4 to access the **Support and Diagnostics** option.
4. Enter 2 to access the **Access Diagnostic Shell** option.
5. Enter `vi /opt/netapp/vscserver/etc/kamino/kaminoprefs.xml` to update the

kaminoprefs.xml file.

6. Update the kaminoprefs.xml file.

If you use...	Do this...
iSCSI	Change the value of the entry key <code>default.allow.iscsi.mount.networks</code> from ALL to the value of your ESXi host networks.
NFS	Change the value of the entry key <code>default.allow.nfs.mount.networks</code> from ALL to the value of your ESXi host networks.

The preferences file includes sample values for these entry keys.



The value “ALL” does not mean all networks. “ALL” value enables all of the matching networks, between the host and the storage system, to be used for mounting datastores. When you specify host networks, then you can enable mounting only across the specified subnets.

7. Save and close the kaminoprefs.xml file.

Access the maintenance console options of the virtual appliance for VSC, VASA Provider, and SRA

You can manage your application, system, and network configurations by using the maintenance console of the virtual appliance for Virtual Storage Console (VSC), VASA Provider, and Storage Replication Adapter (SRA). You can change your administrator password and maintenance password. You can also generate support bundles, set different log levels, view and manage TLS configurations, and start remote diagnostics.


Before you begin

You must have installed VMware tools after deploying the virtual appliance for VSC, VASA Provider, and SRA.

About this task

- You must use “maint” as the user name and the password you configured during deployment to log in to the maintenance console of the virtual appliance for VSC, VASA Provider, and SRA.
- You must set a password for the “diag” user while enabling remote diagnostics.

Steps

1. Access the **Summary** tab of your deployed virtual appliance.
2. Click  to start the maintenance console.

You can access the following maintenance console options:

- **Application Configuration**

The following options are available:

- Display server status summary
- Start Virtual Storage Console service
- Stop Virtual Storage Console service
- Start VASA Provider and SRA service
- Stop VASA Provider and SRA service
- Change 'administrator' user password
- Re-generate certificates
- Hard reset keystore and certificates
- Hard reset database
- Change LOG level for Virtual Storage Console service
- Change LOG level for VASA Provider and SRA service
- Display TLS configuration
- Enable TLS protocol
- Disable TLS protocol

◦ **System Configuration**

The following options are available:

- Reboot virtual machine
- Shutdown virtual machine
- Change 'maint' user password
- Change time zone
- Change NTP server

You can provide an IPv6 address for your NTP server.

- Enable/Disable SSH Access
- Increase jail disk size (/jail)
- Upgrade
- Install VMware Tools

◦ **Network Configuration**

The following options are available:

- Display IP address settings
- Change IP address settings

You can use this option to change the IP address post deployment to IPv6.

- Display domain name search settings
- Change domain name search settings

- Display static routes
- Change static routes

You can use this option to add an IPv6 route.

- Commit changes
- Ping a host

You can use this option to ping to an IPv6 host.

- Restore default settings

- **Support and Diagnostics**

The following options are available:

- Generate support bundle
- Access diagnostic shell
- Enable remote diagnostic access

Related information

[VSC and VASA Provider log files](#)

Change the administrator password

You can change the administrator password of the virtual appliance for VSC, VASA Provider, and SRA post deployment using the maintenance console.

Steps

1. From the vCenter Server, open a console to the virtual appliance for VSC, VASA Provider, and SRA.
2. Log in as the maintenance user.
3. Enter **1** in the maintenance console to select **Application Configuration** .
4. Enter **6** to select **Change 'administrator' user password**.
5. Enter a password with minimum eight characters and maximum 63 characters.
6. Enter **y** in the confirmation dialog box.

Configure high availability for the virtual appliance for VSC, VASA Provider, and SRA

The virtual appliance for Virtual Storage Console (VSC), VASA Provider, and Storage Replication Adapter (SRA) supports a (HA) configuration to help provide uninterrupted functionality of VSC, VASA Provider, and SRA during failure.

The virtual appliance for VSC, VASA Provider, and SRA relies on the VMware vSphere (HA) feature and vSphere fault tolerance (FT) feature to provide . (HA) solution provides for rapid recovery from outages caused by:

- Host failure
- Network failure
- Virtual machine failure (Guest OS failure)
- Application (VSC, VASA Provider, and SRA) crash

No additional configuration is required on the virtual appliance to provide . Only the vCenter Server and ESXi hosts must be configured with the VMware vSphere HA feature or the vSphere FT feature based on their requirements. Both HA and FT require clustered hosts together with shared storage. FT has additional requirements and limitations.

In addition to the VMware vSphere HA solution and vSphere FT solution, the virtual appliance also helps keep the VSC, VASA Provider, and SRA services running at all times. The virtual appliance watchdog process periodically monitors all three services, and restarts them automatically when any kind of failure is detected. This helps to prevent application failures.



vCenter HA is not supported by virtual appliance for VSC, VASA Provider, and SRA.

VMware vSphere HA

You can configure your vSphere environment where the virtual appliance for Virtual Storage Console (VSC), VASA Provider, and Storage Replication Adapter (SRA) is deployed for (HA). The VMware HA feature provides failover protection from hardware failures and operating system failures in virtual environments.

The VMware HA feature monitors virtual machines to detect operating system failures and hardware failures. When a failure is detected, the VMware HA feature restarts the virtual machines on the other physical servers in the resource pool. Manual intervention is not required when a server failure is detected.

The procedure to configure VMware HA depend on the version of your vCenter Server. For example, you can use the following reference link and select the required vCenter Server version to view the steps to configure VMware HA.

[VMware vSphere Documentation: Creating and Using vSphere HA Clusters](#)

VMware vSphere Fault Tolerance

The VMware vSphere Fault Tolerance (FT) feature provides (HA) at a higher level and enables you to protect virtual machines without any loss of data or connections. You must enable or disable vSphere FT for the virtual appliance for VSC, VASA Provider, and SRA from your vCenter Server.

Ensure your vSphere license supports FT with the number of vCPUs needed for the virtual appliance in your environment (at least 2 vCPUs; 4 vCPUs for large scale environments).

vSphere FT enables virtual machines to operate continuously even during server failures. When vSphere FT is enabled on a virtual machine, a copy of the primary virtual machine is automatically created on another host (the secondary virtual machine) that is selected by Distributed Resource Scheduler (DRS). If DRS is not enabled, the target host is selected from the available hosts. vSphere FT operates the primary virtual machine and secondary virtual machine in lockstep mode, with each mirroring the execution state of the primary virtual machine to the secondary virtual machine.

When there is a hardware failure that causes the primary virtual machine to fail, the secondary virtual machine immediately picks up where the primary virtual machine stopped. The secondary virtual machine continues to run without any loss of network connections, transactions, or data.

Your system must meet the CPU requirements, virtual machine limit requirements, and licensing requirements for configuring vSphere FT for your vCenter Server instance.

The procedure to configure HA depend on the version of your vCenter Server. For example, you can use the following reference link and select the required vCenter Server version to view the steps to configure HA.

[VMware vSphere Documentation: Fault Tolerance Requirements, Limits, and Licensing](#)

MetroCluster configurations supported by the virtual appliance for VSC, VASA Provider, and SRA

The virtual appliance for Virtual Storage Console (VSC), VASA Provider, and Storage Replication Adapter (SRA) supports environments that use MetroCluster IP and FC configurations for ONTAP. Most of this support is automatic. However, you might notice a few differences when you use a MetroCluster environment with VSC and VASA Provider.

MetroCluster configurations and VSC

You must ensure that VSC discovers the storage system controllers at the primary site and the secondary site. Typically, VSC automatically discovers storage controllers. If you are using a cluster management LIF, then it is a good practice to verify that VSC has discovered the clusters at both sites. Otherwise, you can manually add the storage controllers to VSC. You can also modify the user name and password pairs that VSC uses to connect to the storage controllers.

When a switchover occurs, the on the secondary site take over. These have the “-mc” suffix appended to their names. If a switchover operation occurs while you are performing operations such as provisioning a datastore, the name of the where the datastore resides is changed to include the “-mc” suffix. This suffix is dropped when the switchback occurs, and the on the primary site resume control.



If you have added direct with MetroCluster configuration to VSC, then after switchover, the change in the SVM name (the addition of the “-mc” suffix) is not reflected. All other switchover operations continue to execute normally.

When a switchover or switchback occurs, VSC might take a few minutes to automatically detect and discover the clusters. If this happens while you are performing a VSC operation such as provisioning a datastore, you might experience a delay.

MetroCluster configurations and VASA Provider

VASA Provider automatically supports environments that use MetroCluster configurations. The switchover is transparent in VASA Provider environments. You cannot add direct to VASA Provider.



VASA Provider does not append the “-mc” suffix to the names of the on the secondary site after a switchover.

MetroCluster configurations and SRA

SRA does not support MetroCluster configurations.

Copyright information

Copyright © 2024 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

Trademark information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.