Cisco Videoscape Distribution Suite Transparent Caching Troubleshooting Guide

Release 5.7.3

March 2016

Cisco Systems, Inc.
www.cisco.com

Cisco has more than 200 offices worldwide.
Addresses, phone numbers, and fax numbers
are listed on the Cisco website at
www.cisco.com/go/offices.

Text Part Number: OL-29792-05

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

The following information is for FCC compliance of Class A devices: This equipment has been tested and found to comply with the limits for a Class A digital device, pursuant to part 15 of the FCC rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses, and can radiate radio-frequency energy and, if not installed and used in accordance with the instruction manual, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference, in which case users will be required to correct the interference at their own expense.

The following information is for FCC compliance of Class B devices: This equipment has been tested and found to comply with the limits for a Class B digital device, pursuant to part 15 of the FCC rules. These limits are designed to provide reasonable protection against harmful interference in a residential installation. This equipment generates, uses and can radiate radio frequency energy and, if not installed and used in accordance with the instructions, may cause harmful interference to radio communications. However, there is no guarantee that interference will not occur in a particular installation. If the equipment causes interference to radio or television reception, which can be determined by turning the equipment off and on, users are encouraged to try to correct the interference by using one or more of the following measures:

- Reorient or relocate the receiving antenna.
- Increase the separation between the equipment and receiver.
- Connect the equipment into an outlet on a circuit different from that to which the receiver is connected.
- Consult the dealer or an experienced radio/TV technician for help.

Modifications to this product not authorized by Cisco could void the FCC approval and negate your authority to operate the product.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB’s public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED “AS IS” WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: www.cisco.com/go/trademarks. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

CONTENTS

Preface

Revised: March 2016, OL-29792-05

The *Cisco Videoscape Distribution Suite Transparent Caching Troubleshooting Guide* provides information on troubleshooting and fixing issues with a VDS TC installation. This guide includes the following chapters:

- Cache Engine Replacement Procedure for a C-Series Installation with a Fresh Install of VDS TC 5.7.3
- Cache Engine Replacement Procedure in a 5.7.3 C-Series Installation Upgraded from Release 5.2.1
- Cache Engine Replacement Procedure in a 5.7.3 C-Series Installation Upgraded from Release 5.6.1
- Cache Engine Replacement Procedure for a Blade Server Installation with a Fresh Install of VDS TC 5.7.3
- Cache Engine Replacement Procedure for a 5.7.3 Blade Server Installation Upgraded from Release 5.2.1
- Cache Engine Replacement Procedure for a 5.7.3 Blade Server Installation Upgraded from Release 5.6.1
- Management Server Replacement Procedure for a System with a Fresh Install of VDS TC 5.7.3
- Management Server Replacement Procedure for a VDS TC 5.7.3 System Upgraded from VDS TC 5.2.1
- Management Server Replacement Procedure for a VDS TC 5.7.3 System Upgraded from VDS TC 5.6.1
- Storage Disk Replacement Procedure for NetApp Storage Enclosure
- Storage Disk Replacement Procedure for IBM Storage Enclosure
- VDS TC Integrated Appliance Cache Disk Replacement
- System Disk Replacement in a VDS TC C-Series Server Installation
- VDS TC Integrated Appliance Cache Disk Replacement
- System Disk Replacement in a VDS TC C-Series Server Installation
- System Disk Replacement in a VDS TC Blade Server Installation
- VDS TC Troubleshooting Files
- Troubleshooting Storage Disk Problems
- Troubleshooting Network Connectivity Issues

Obtaining Documentation and Submitting a Service Request

For information on obtaining documentation, submitting a service request, and gathering additional information, see the monthly *What's New in Cisco Product Documentation*, which also lists all new and revised Cisco technical documentation, at:

http://www.cisco.com/en/US/docs/general/whatsnew/whatsnew.html

Subscribe to the *What's New in Cisco Product Documentation* as an RSS feed and set content to be delivered directly to your desktop using a reader application. The RSS feeds are a free service. Cisco currently supports RSS Version 2.0.

CHAPTER 1

Cache Engine Replacement Procedure for a C-Series Installation with a Fresh Install of VDS TC 5.7.3

Note Only use this chapter if you are replacing a Cache Engine in a VDS TC C-Series system that had a fresh install of VDS TC 5.7.3 and was *not* upgraded to 5.7.3. If you are replacing a Cache Engine in a VDS TC C-Series system that was upgraded to VDS TC 5.7.3 from VDS TC 5.2.1, refer to Chapter 2, “Cache Engine Replacement Procedure in a 5.7.3 C-Series Installation Upgraded from Release 5.2.1”. If you are replacing a Cache Engine in a VDS TC C-Series system that was upgraded to VDS TC 5.7.3 from VDS TC 5.6.1, refer to Chapter 3, “Cache Engine Replacement Procedure in a 5.7.3 C-Series Installation Upgraded from Release 5.6.1”.

This chapter discusses how to determine that a cache engine has failed in a VDS TC C-Series installation and the steps that are required to replace it.

When a cache engine fails, the volume_selection_algorithm_time setting in the cluster_conf.xml file controls when the volumes of a failed cache engine will be redistributed (mounted) to the remaining cache engines. By default this value is 24 hours, which means that by default, if a cache engine fails, its volumes will not be redistributed (mounted) to the remaining cache engines until 24 hours after the cache engine fails.

It is recommended that you leave the default set to 24 hours, however if you need to change this setting, add the following text to the cluster_conf file in the <service> section under the <memory> subsection. The *seconds* parameters is the number of seconds that the system should wait after a cache engine fails before its volumes are redistributed (mounted) to the remaining cache engines:

```
<io>
    <volume_selection_algorithm_time>seconds</volume_selection_algorithm_time>
</io>
```

For example, to have the system wait 30 minutes after a cache engine fails before redistributing the volumes of the failed cache engine enter the following:

```
<io>
    <volume_selection_algorithm_time>1800</volume_selection_algorithm_time>
</io>
```

For more information on how to make changes to the configuration file and how to apply these changes, please refer to the *Cisco Videoscape Distribution Suite Transparent Caching Software Configuration Guide* available at http://www.cisco.com/c/dam/en/us/td/docs/video/videoscape/distribution_suite/vds/v5_7_3/VDS-TC_5_7_3_software_config_guide.pdf.

Symptoms of a Failed VDS TC C-Series Cache Engine

To determine whether a cache engine has failed in a VDS TC C-Series Server installation, log into VDS TC Manager. By default, the Status Dashboard displays. If you are already logged into VDS TC Manager GUI, choose **Status > Dashboard**.

If a cache engine has failed it will appear with a red box surrounding it, as shown with the first cache engine in the following figure. If a cache engine cannot bootup, it will appear with a grey box surrounding it.

Figure 1-1 Failed Cache Engine

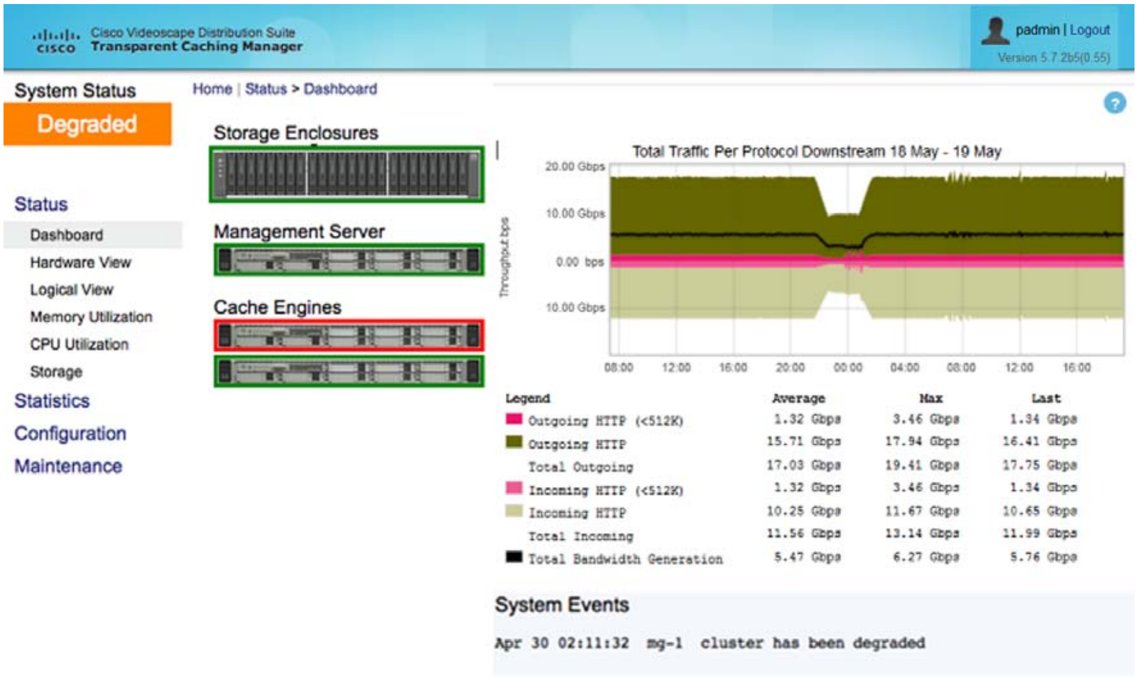

Legend	Average	Max	Last
🟥 Outgoing HTTP (<512K)	1.32 Gbps	3.46 Gbps	1.34 Gbps
🟫 Outgoing HTTP	15.71 Gbps	17.94 Gbps	16.41 Gbps
Total Outgoing	17.03 Gbps	19.41 Gbps	17.75 Gbps
🟪 Incoming HTTP (<512K)	1.32 Gbps	3.46 Gbps	1.34 Gbps
🟨 Incoming HTTP	10.25 Gbps	11.67 Gbps	10.65 Gbps
Total Incoming	11.56 Gbps	13.14 Gbps	11.99 Gbps
■ Total Bandwidth Generation	5.47 Gbps	6.27 Gbps	5.76 Gbps

System Events

```
Apr 30 02:11:32  mg-1  cluster has been degraded
```

Replacing a Cache Engine and Running the Installation Tasks

The first step in replacing a cache engine in a VDS TC C-Series cluster installation is to physically replace the cache engine and perform the installation steps on the new cache engine. Follow this procedure to perform these steps:

Before You Begin

To replace a failed cache engine you must request a new VDS TC license from Cisco. To request a new license, you must have the serial number of the failed cache engine and the serial number of the new cache engine. After you have gathered this information, contact Cisco support for a new license file.

Caution Do not continue until you have the license file.

Procedure

Note Before removing the failed cache engine, make note of what number corresponds with the cache engine.

Step 1 Using the Cisco UCS C220M3 documentation, unmount the failed cache engine and mount the new cache engine.

Step 2 In the *Cisco Videoscape Distribution Suite Transparent Caching Software C-Series Cluster Installation Guide*, available at http://www.cisco.com/c/dam/en/us/td/docs/video/videoscape/distribution_suite/vds/v5_7_3/VDS-TC_5_7_3_CSeries_swInstallGuide.pdf, refer to Chapter 1, “Prerequisites and C-Series Cluster Physical Installation” for connecting the network cables to the new cache engine.

Step 3 Configure the new VDS TC cache engine by performing all of the tasks in Chapter 4, “Cache Engine Servers Installation” in the *Cisco Videoscape Distribution Suite Transparent Caching Software C-Series Cluster Installation Guide.*

Note When the “Enter Cache Engine Number” dialog box appears, enter the number that corresponds with the cache engine that you are replacing.

Figure 1-2 ***Enter Cache Engine Number***

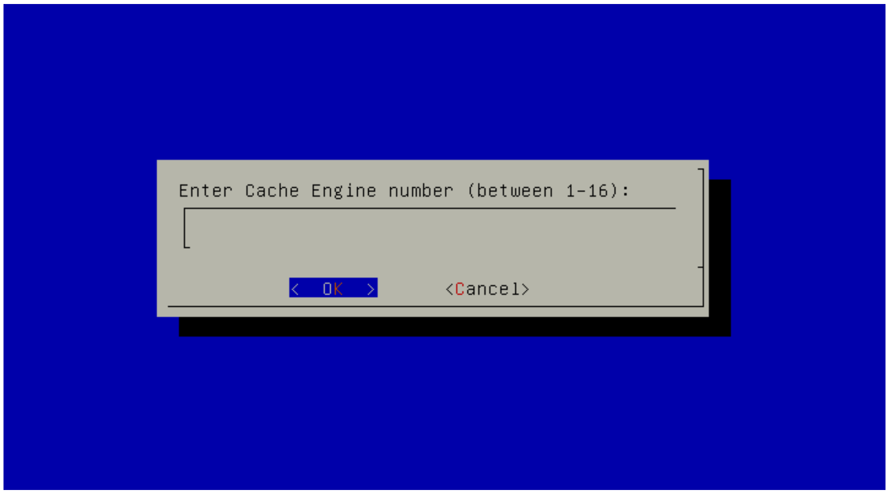

```
Enter Cache Engine number (between 1-16):
```

Updating the SSH Keys and Connecting the Cache Engine to the Storage Enclosure

After completing the “Cache Engine Servers Installation” chapter, you must update the grid SSH keys on the VDS TC management server and configure the new cache engine to connect to the storage enclosure. All of these steps are performed on the VDS TC management server.

Follow this procedure to perform these steps on the VDS TC management server:

Note The following steps are performed on the VDS TC management server.

Procedure

Step 1 On the PBR routers, stop redirecting traffic to the VDS TC solution.

Step 2 Close any open VDS TC Manager windows and ensure that no other administrators are connected to the VDS TC Manager before proceeding.

Step 3 Using SSH software, such as Putty, open an SSH connection to the IP address of the VDS TC management server.

Step 4 Log into the system using the username **padmin** and the password provided by Cisco.

Step 5 Enter the command **su admin** to log into the CLI of the VDS TC management server. The default password is the serial number of the VDS TC management server.

Step 6 From the VDS TC prompt, enter the **enable** command. When prompted, enter the Enable mode password and press **Enter**. You are now logged into Enable mode and the Enable prompt, console#, should appear.

Step 7 From the Enable mode prompt, enter the command **oper service stop** to stop the caching service.

Step 8 Wait a few minutes and then enter the command **show status**. Check to see if the Device State shows “stopped.” For example:

```
console# show status
Operational state Device state Administrative state
disabledstoppedunlocked
```

Caution Do not proceed to the next step until the Device State shows “stopped.” You may need to repeat the **show status** command several times before you see this status. Wait several minutes between executions of the **show status** command.

Step 9 Enter the **exit** command to exit Enabled mode.

Step 10 Enter the **exit** command to exit the VDS TC CLI.

Step 11 Enter the command **su root** to change to the root user. Enter the password for the root user when prompted. This password is provided by Cisco.

Step 12 Enter the command **cd /opt/pang/useful/**.

Step 13 To reset the SSH keys, enter the command **./replace_server_keys.sh -servers** #, where # is the total number of C-Series cache engines in the VDS TC installation. For example, if you are replacing a cache engine in a VDS TC solution that has 16 cache engines, you would enter the command **./replace_server_keys.sh -servers 16**.

Step 14 Enter the command **cd /opt/pang/useful/configure_storages** to change to the storage configuration script directory.

Step 15 Enter the command **./configure_storages.py -s** <*no_of_storage_enclosures*> **-b** <*no_of_blades_in_grid*> **-c** <*the_replaced_CE-ID_number*>. For example, to replace CE-2 in a VDS TC C-Series installation that has two cache engines, and one storage enclosure, enter the command **./configure_storages.py -s 1 -b 2 -c 2**.

Note The configure storages script will format the data disks. To configure the storages without formatting the disks enter the command **./configure_storages.py -s** <*no_of_storage_enclosures*> **-b** <*no_of_blades_in_grid*> **-c** <*the_replaced_CE-ID_number*> **-x**.

Step 16 When the warning message about the disk format appears, enter **y** to continue the storages configuration.

```
*********** WARNING! **********
All storages are about to be reconfigured. All data on the data disks will
be lost.
Do you want to continue? (y/n)y
```

Note You may be prompted to enter the password several times.

When the script is finished running, you should see the word “COMPLETE”. After running this script, the new VDS TC cache engine will be connected to the storage enclosures.

Step 17 Use SSH to connect to each CE and enter the command **iscsiadm -m session** to check the connections:

```
ssh ce-<cache engine number>
iscsiadm -m session
```

Your output should be similar to the following:

```
ƒÞ 8 * sessions per storage, for example for 1 storage:
tcp: [17] 10.11.14.100:3260,1 iqn.1992-
01.com.Avago:2365.60080e50002f4e82000000005063c74a
tcp: [18] 10.11.14.101:3260,2 iqn.1992-
01.com.Avago:2365.60080e50002f4e82000000005063c74a
tcp: [19] 10.11.16.100:3260,1 iqn.1992-
01.com.Avago:2365.60080e50002f4e82000000005063c74a
tcp: [20] 10.11.16.101:3260,2 iqn.1992-
01.com.Avago:2365.60080e50002f4e82000000005063c74a
tcp: [21] 10.11.15.100:3260,1 iqn.1992-
01.com.Avago:2365.60080e50002f4e82000000005063c74a
tcp: [22] 10.11.15.101:3260,2 iqn.1992-
01.com.Avago:2365.60080e50002f4e82000000005063c74a
tcp: [23] 10.11.17.101:3260,2 iqn.1992-
01.com.Avago:2365.60080e50002f4e82000000005063c74a
tcp: [24] 10.11.17.100:3260,1 iqn.1992-
01.com.Avago:2365.60080e50002f4e82000000005063c74a
```

Updating the NetApp Storage Enclosure

Note This process is only for the NetApp storage enclosure. If your system uses an IBM storage enclosure, refer to Updating the IBM Storage Enclosure section.

After you have configured the VDS TC cache engine to connect with the storage enclosure, you must update the host port identifiers on the NetApp storage enclosure for the new cache engine. Follow this procedure to perform this task:

Procedure

Step 1 From the SSH connection to the VDS TC management server, enter the following commands to start the VNC server:

a. **rm /tmp/.X1-lock** (This step is optional. It is only required if there is a leftover lock file from a previous vncserver instance.)

b. **rm /tmp/.X11-unix/X1** (This step is optional. It is only required if there is a leftover lock file from a previous vncserver instance.)

c. **vncserver :1**

Step 2 Use a VNC client to connect from your local computer to the VDS TC management server.

Figure 1-3 VNC Viewer

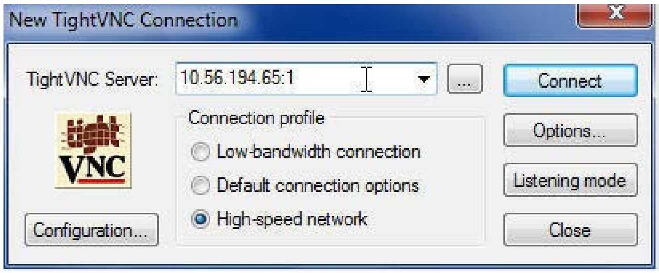

TightVNC Server: 10.56.194.65:1

Connection profile

○ Low-bandwidth connection
○ Default connection options
● High-speed network

Step 3 From the VNC console window, enter the command **cd /opt/SMgr/client** to change folders.

Step 4 Enter the command **./SMclient** to start the SANtricity ™ ES Storage Manager software.

Note You must perform the following steps on all of the storage subsystems.

Step 5 Right-click the storage subsystem that you are configuring and choose **Manage Storage Array**. Click **No** in any pop-up windows that appear.

Figure 1-4 SANtricity Storage Manager

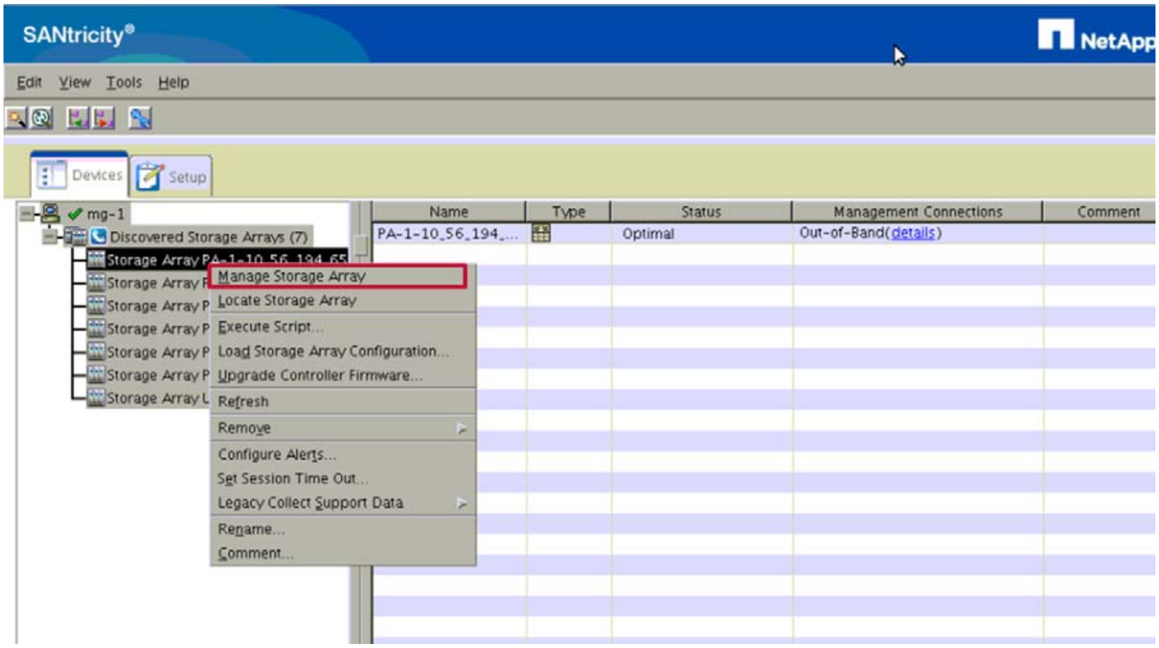

Name	Type	Status	Management Connections	Comment
PA-1-10_56_194_...		Optimal	Out-of-Band(details)	

Step 6 In the Subsystem Management window that appears, click the **Host Mappings** tab.

Step 7 In the navigation pane on the left, expand the **Host Group** folder.

Step 8 Right-click the cache engine that was replaced, for example Host CE-1, and choose **Manage Host Port Identifiers** from the pop-up menu.

Figure 1-5 Manage Host Port Identifiers

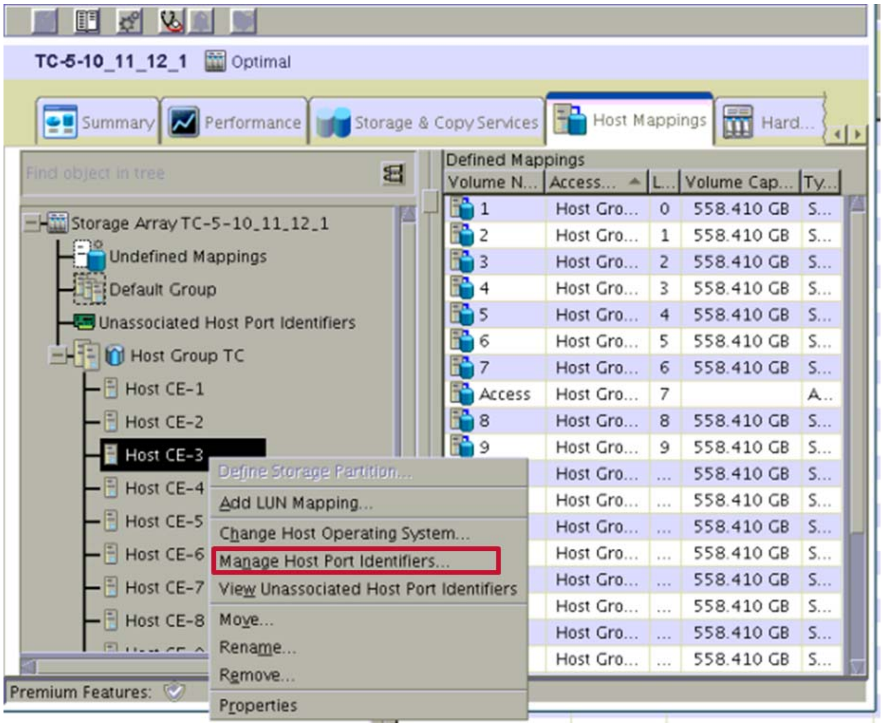

Defined Mappings

Volume N...	Access...	L...	Volume Cap...	Ty...
1	Host Gro...	0	558.410 GB	S...
2	Host Gro...	1	558.410 GB	S...
3	Host Gro...	2	558.410 GB	S...
4	Host Gro...	3	558.410 GB	S...
5	Host Gro...	4	558.410 GB	S...
6	Host Gro...	5	558.410 GB	S...
7	Host Gro...	6	558.410 GB	S...
Access	Host Gro...	7		A...
8	Host Gro...	8	558.410 GB	S...
9	Host Gro...	9	558.410 GB	S...
	Host Gro...	...	558.410 GB	S...
	Host Gro...	...	558.410 GB	S...
	Host Gro...	...	558.410 GB	S...
	Host Gro...	...	558.410 GB	S...
	Host Gro...	...	558.410 GB	S...
	Host Gro...	...	558.410 GB	S...
	Host Gro...	...	558.410 GB	S...
	Host Gro...	...	558.410 GB	S...

Step 9 From the Manage Host Port Identifiers window that appears, in the Host Port Identifier Information table, ensure that the old iqn is selected and click **Replace**.

Figure 1-6 Replace IQN

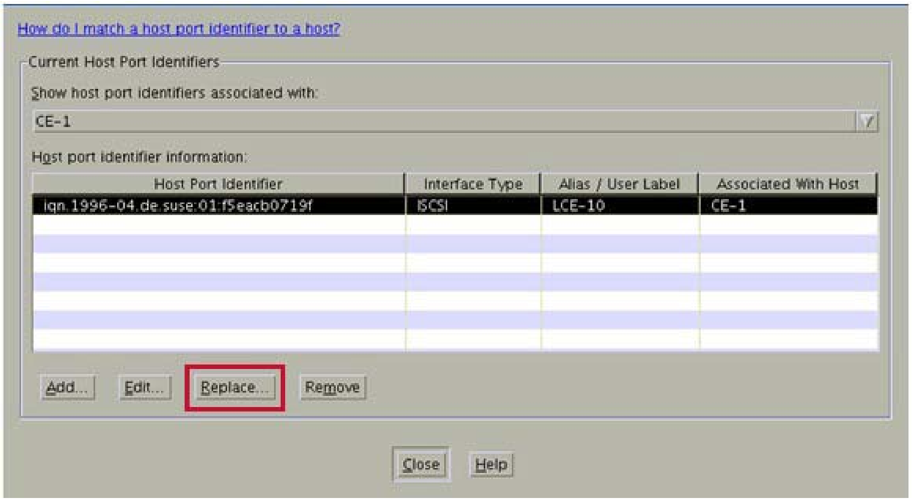
How do I match a host port identifier to a host?
Current Host Port Identifiers
Add...
Edit...
Replace...
Remove
Close
Help

Show host port identifiers associated with: CE-1

Host port identifier information:

Host Port Identifier	Interface Type	Alias / User Label	Associated With Host
iqn.1996-04.de.suse:01:f5eacb0719f	iSCSI	LCE-10	CE-1

Step 10 From the Replace Host Port Identifier window that appears, configure the following:

- From the Choose a Host Interface Type drop-down list, choose **ISCSI**.
- Click the **Replace by Selecting a Known Unassociated Host Port Identifier** radio button.

- From the Known Unassociated Host Port Identifier drop-down list, choose the iqn for the *newly* installed cache engine. (It should be the only one in the list.)

Step 11 Click **Replace**.

Figure 1-7 Replace Host Port Identifier

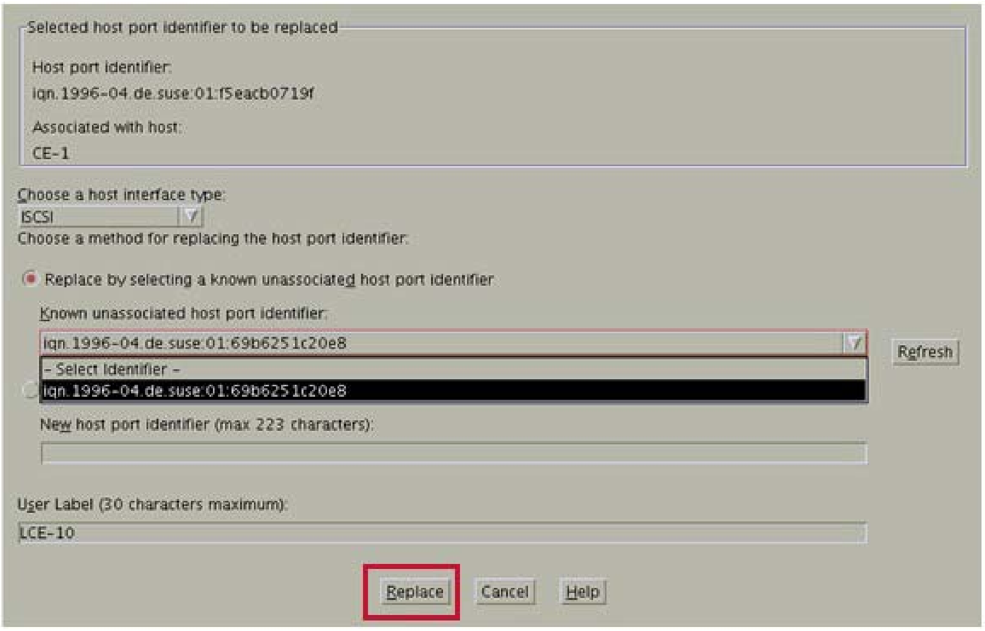

Selected host port identifier to be replaced

Host port identifier:
iqn.1996-04.de.suse:01:f5eacb0719f

Associated with host:
CE-1

Choose a host interface type:
ISCSI

Choose a method for replacing the host port identifier:

● Replace by selecting a known unassociated host port identifier

Known unassociated host port identifier:
iqn.1996-04.de.suse:01:69b6251c20e8

- Select Identifier -
iqn.1996-04.de.suse:01:69b6251c20e8

New host port identifier (max 223 characters):

User Label (30 characters maximum):
LCE-10

Step 12 From the Manage Host Port Identifiers window, click **Close**.

Step 13 Repeat Step 5 through Step 12 on any remaining storage subsystems.

Step 14 Close the Subsystem Management window and exit the SANtricity Storage Manager.

Step 15 Close the VNC client.

Updating the IBM Storage Enclosure

Note This process is only for the IBM storage enclosure. If your system uses the NetApp E2724 storage enclosure, refer to Updating the NetApp Storage Enclosure section.

After you have configured the VDS TC cache engine to connect with the storage enclosure, you must update the host port identifiers on the IBM storage enclosure for the new cache engine. Follow this procedure to perform this task:

Procedure

Step 1 From the SSH connection to the VDS TC management server, enter the following commands to start the VNC server:

a. **rm /tmp/.X1-lock** (This step is optional. It is only required if there is a leftover lock file from a previous vncserver instance.)

b. **rm /tmp/.X11-unix/X1** (This step is optional. It is only required if there is a leftover lock file from a previous vncserver instance.)

c. **vncserver :1**

Step 2 Use a VNC client to connect from your local computer to the VDS TC management server.

Figure 1-8 VNC Viewer

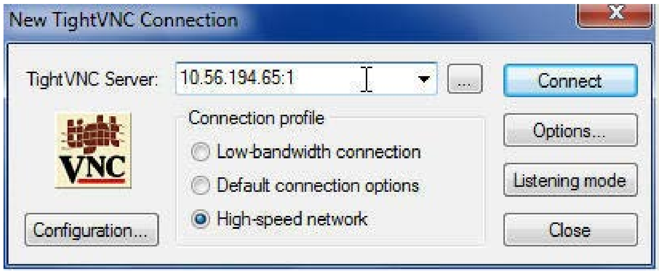

New TightVNC Connection

TightVNC Server: 10.56.194.65:1

Connection profile

○ Low-bandwidth connection

○ Default connection options

● High-speed network

Step 3 From the VNC console window, enter the command **/opt/IBM_DS/client/SMclient** to start the IBM Storage Manager.

Note You must perform the following steps on all of the storage subsystems.

Step 4 Right-click the storage subsystem that you are configuring and choose **Manage Storage Subsystem**. Click **No** in any pop-up windows that appear.

Figure 1-9 IBM DS Storage Manager

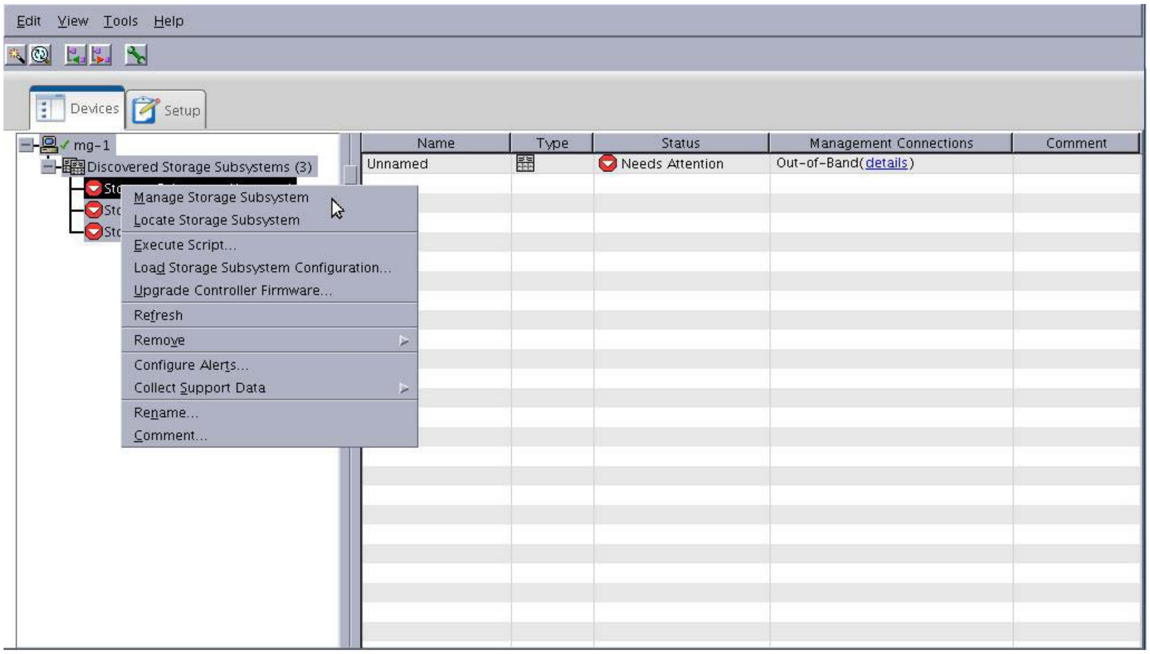

Name	Type	Status	Management Connections	Comment
Unnamed		Needs Attention	Out-of-Band(details)	

Step 5 In the Subsystem Management window that appears, click the **Host Mappings** tab.

Step 6 In the navigation pane on the left expand the **Host Group** folder.

Step 7 Right-click the cache engine that was replaced, for example Host CE-1, and choose **Manage Host Port Identifiers** from the pop-up menu.

Figure 1-10 Manage Host Port Identifiers

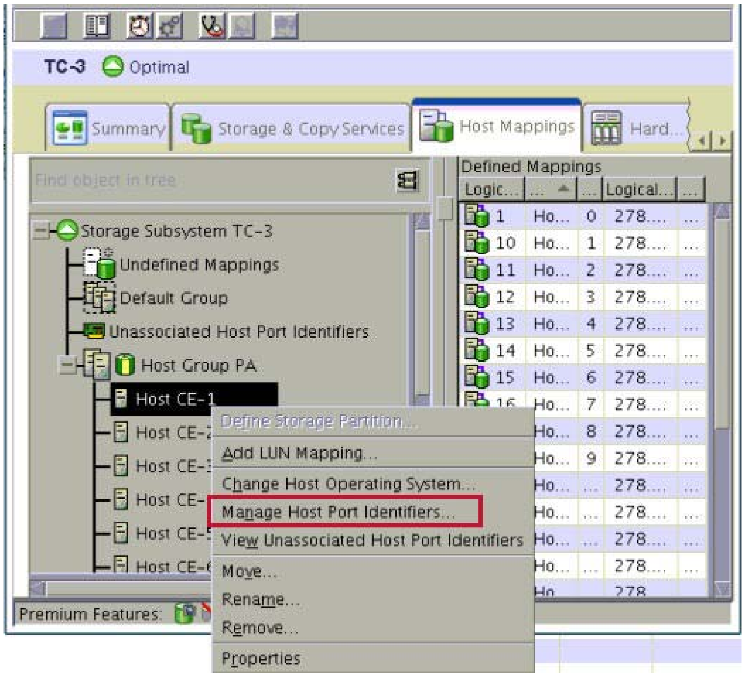

Logic...	...	...	Logical...	...
1	Ho...	0	278....	...
10	Ho...	1	278....	...
11	Ho...	2	278....	...
12	Ho...	3	278....	...
13	Ho...	4	278....	...
14	Ho...	5	278....	...
15	Ho...	6	278....	...
16	Ho...	7	278....	...
	Ho...	8	278....	...
	Ho...	9	278....	...
	Ho...	...	278....	...
	Ho...	...	278....	...
	Ho...	...	278....	...
	Ho...	...	278....	...
	Ho		278	

Step 8 From the Manage Host Port Identifiers window that appears, in the Host Port Identifier Information table, ensure that the old iqn is selected and click **Replace**.

Figure 1-11 Replace IQN

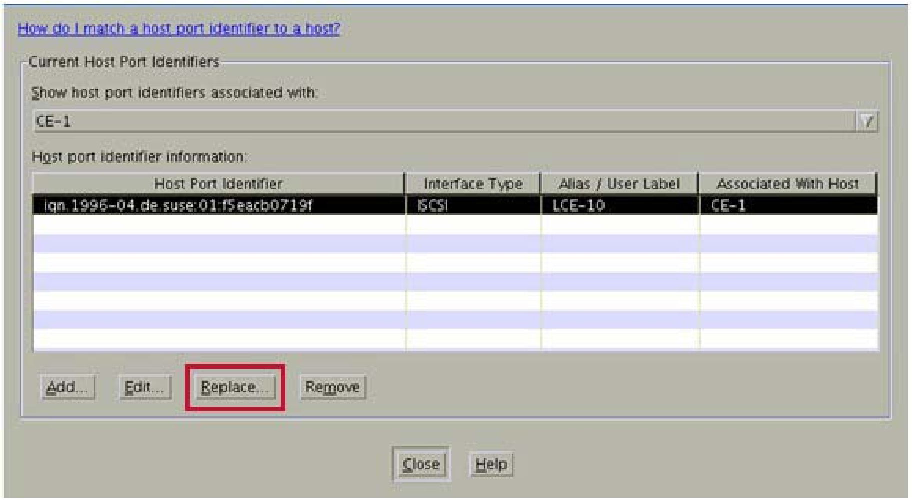
How do I match a host port identifier to a host?
Current Host Port Identifiers
Add...
Edit...
Replace...
Remove
Close
Help

Show host port identifiers associated with: CE-1

Host port identifier information:

Host Port Identifier	Interface Type	Alias / User Label	Associated With Host
iqn.1996-04.de.suse:01:f5eacb0719f	ISCSI	LCE-10	CE-1

Step 9 From the Replace Host Port Identifier window that appears, configure the following:

- From the Choose a Host Interface Type drop-down list, choose **ISCSI**.
- Click the **Replace by Selecting a Known Unassociated Host Port Identifier** radio button.

- From the Known Unassociated Host Port Identifier drop-down list, choose the iqn for the *newly* installed cache engine. (It should be the only one in the list.)

Step 10 Click **Replace**.

Figure 1-12 Replace Host Port Identifier

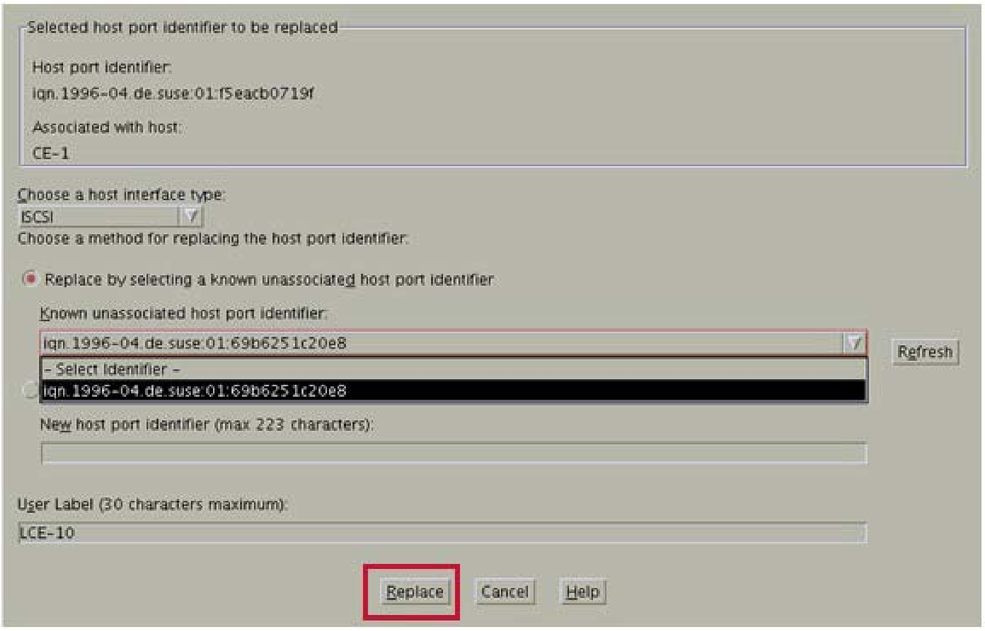

Selected host port identifier to be replaced

Host port identifier:
iqn.1996-04.de.suse:01:f5eacb0719f

Associated with host:
CE-1

Choose a host interface type:
ISCSI

Choose a method for replacing the host port identifier:

● Replace by selecting a known unassociated host port identifier

Known unassociated host port identifier:
iqn.1996-04.de.suse:01:69b6251c20e8

- Select Identifier -
iqn.1996-04.de.suse:01:69b6251c20e8

○ New host port identifier (max 223 characters):

User Label (30 characters maximum):
LCE-10

Step 11 From the Manage Host Port Identifiers window, click **Close**.

Step 12 Repeat Step 4 through Step 11 on any remaining storage subsystems.

Step 13 Close the Subsystem Management window and exit the IBM Storage Manager.

Step 14 Close the VNC client.

Finishing the Cache Engine Replacement

Follow this procedure to finish configuring and installing the new cache engine:

Note You will be re-running the storage script that you ran in the Updating the SSH Keys and Connecting the Cache Engine to the Storage Enclosure section.

Note You must be logged in with root permissions to perform the following steps.

Procedure

Step 1 From the SSH connection to the VDS TC management server, enter the command **cd /opt/pang/useful/configure_storages** to change to the storage configuration script directory.

Step 2 Enter the command **./configure_storages.py -s** <*no_of_storage_enclosures*> **-b** <*no_of_blades*> **-c** <*the_replaced_CE-ID_number*>. For example, to replace CE-2 in a VDS TC C-Series installation that has two cache engines, and one storage enclosure, enter the command **./configure_storages_py -s 1 -b 2 -c 2**.

Step 3 When the warning message about the disk format appears, enter **y** to continue the storages configuration.

```
*********** WARNING! **********
All storages are about to be reconfigured. All data on the data disks will
be lost.
Do you want to continue? (y/n)y
```

Step 4 Follow the on screen instructions. If prompted, enter the password for the root user that was provided by Cisco. You should see the word "COMPLETE" when the script is finished running.

Note You may be prompted to enter the password several times.

Note Ensure that the VDS TC image upgrade file and the *new* VDS TC license that you received from Cisco are located on the VDS TC management server in the /tftpboot folder before performing these steps. Please note that you must receive a new license from Cisco support that includes the serial number of the replacement cache engine.

Step 5 Enter the command **su admin** to log into the VDS TC CLI. The CLI prompt console> appears.

Step 6 Enter the command **enable** to enter Enable mode. At the password prompt that appears, enter the enable password. The CLI prompt console# appears.

Step 7 To upgrade the software on the new cache engine, enter the command **upgrade server** *CE_number* **127.0.0.1 VDS-TC_FI_5.7.3b54_Server_Cluster.tar.gz**, where *CE_number* is the cache engine that you replaced.

Step 8 To import the new license that includes the serial number of the new cache engine, enter the command **license import 127.0.0.1** *filename*, where *filename* is the name of the new license file.

Note You must be in Enable mode, *not* Configuration mode to import a new license.

Step 9 Enter the command **license activate** to apply and activate the license.

Step 10 To start the application on the replaced cache engine, enter the command **oper server** *#*, where *#* is the number of the cache engine that was replaced. For example, if you replaced CE-2, enter the command **oper server 2**. You should see the following output:

```
console# oper server 2
oper server 2# start
Starting server 2
service is started on server 2
```

Step 11 After the application has started on the new cache engine, enter the command **oper service start** to start the VDS TC service.

Step 12 Enter the command **show cluster-bus-ip** to check if the spread communication among cluster members is set to broadcast.

```
console# show cluster-bus-ip
Processing...
MG-1: [Broadcast]
CE-1: [Broadcast]
```

```
CE-2: [Broadcast]
CE-3: [Broadcast]
CE-4: [Broadcast]
CE-5: [Broadcast]
CE-6: [Broadcast]
CE-7: [Broadcast]
CE-8: [Broadcast]
OK. All the CEs are configured to work with Broadcast IP.
```

Step 13 If the communication is set to broadcast, continue on to Step 16. If the communication is set to *multicast* instead of *broadcast*, enter the command **config** to switch to Configuration mode and then enter the command **cluster-bus-ip broadcast**. For example:

```
configuration# cluster-bus-ip broadcast
Processing...
All the CEs are configured to work with Broadcast IP.
```

Step 14 Enter the command **exit** to exit the Configuration mode.

Step 15 Enter the command **show cluster-bus-ip** to confirm that the management server and all cache engines are configured to use broadcast. For example:

```
console# show cluster-bus-ip
Processing...
MG-1: [Broadcast]
CE-1: [Broadcast]
CE-2: [Broadcast]
CE-3: [Broadcast]
CE-4: [Broadcast]
CE-5: [Broadcast]
CE-6: [Broadcast]
CE-7: [Broadcast]
CE-8: [Broadcast]
OK. All the CEs are configured to work with Broadcast IP.
```

Step 16 Wait a few minutes and then enter the command **show status**. Do not proceed to the next step until you see a Device Status of "Started." For example:

```
console# show status
Operational state Device state Administrative state
enabled started unlocked
```

Step 17 On the PBR routers, start redirecting traffic to the VDS TC solution.

CHAPTER 2

Cache Engine Replacement Procedure in a 5.7.3 C-Series Installation Upgraded from Release 5.2.1

Note Only use this chapter if you are replacing a Cache Engine in a VDS TC C-Series system that was upgraded to VDS TC 5.7.3 from VDS TC 5.2.1. If you are replacing a Cache Engine in a VDS TC C-Series system that had a fresh install of VDS TC 5.7.3 and was *not* upgraded to 5.7.3, refer to Chapter 1, "Cache Engine Replacement Procedure for a C-Series Installation with a Fresh Install of VDS TC 5.7.3". If you are replacing a Cache Engine in a VDS TC C-Series system that was upgraded to VDS TC 5.7.3 from VDS TC 5.6.1, refer to Chapter 3, "Cache Engine Replacement Procedure in a 5.7.3 C-Series Installation Upgraded from Release 5.6.1".

To determine whether the system has been upgraded and from what version, from the Management Server enter the command **grep "upgrade system to" /var/log/peerapp/peerapp_system***. This will display a list of all previous versions.

This chapter discusses how to determine that a cache engine has failed in a VDS TC C-Series installation and the steps that are required to replace it.

When a cache engine fails, the volume_selection_algorithm_time setting in the cluster_conf.xml file controls when the volumes of a failed cache engine will be redistributed (mounted) to the remaining cache engines. By default this value is 24 hours, which means that by default, if a cache engine fails, its volumes will not be redistributed (mounted) to the remaining cache engines until 24 hours after the cache engine fails.

It is recommended that you leave the default set to 24 hours, however if you need to change this setting, add the following text to the cluster_conf file in the <service> section under the <memory> subsection. The *seconds* parameters is the number of seconds that the system should wait after a cache engine fails before its volumes are redistributed (mounted) to the remaining cache engines:

```
<io>
    <volume_selection_algorithm_time>seconds</volume_selection_algorithm_time>
</io>
```

For example, to have the system wait 30 minutes after a cache engine fails before redistributing the volumes of the failed cache engine enter the following:

```
<io>
    <volume_selection_algorithm_time>1800</volume_selection_algorithm_time>
</io>
```

For more information on how to make changes to the configuration file and how to apply these changes, please refer to the *Cisco Videoscape Distribution Suite Transparent Caching Software Configuration Guide* available at http://www.cisco.com/c/dam/en/us/td/docs/video/videoscape/distribution_suite/vds/v5_7_3/VDS-TC_5_7_3_software_config_guide.pdf.

Symptoms of a Failed VDS TC C-Series Cache Engine

To determine whether a cache engine has failed in a VDS TC C-Series Server installation, log into VDS TC Manager. By default, the Status Dashboard displays. If you are already logged into VDS TC Manager GUI, choose **Status > Dashboard**.

If a cache engine has failed it will appear with a red box surrounding it, as shown with the first cache engine in the following figure. If a cache engine cannot bootup, it will appear with a grey box surrounding it.

Figure 2-1 Failed Cache Engine

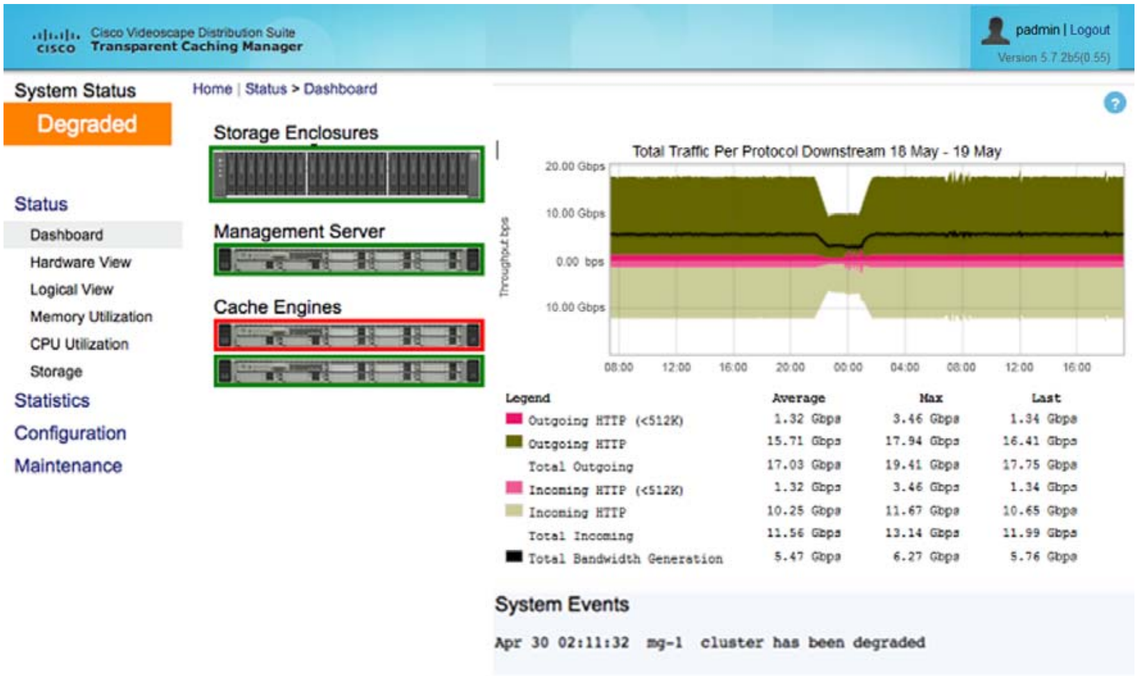

Legend	Average	Max	Last
🟥 Outgoing HTTP (<512K)	1.32 Gbps	3.46 Gbps	1.34 Gbps
🟩 Outgoing HTTP	15.71 Gbps	17.94 Gbps	16.41 Gbps
Total Outgoing	17.03 Gbps	19.41 Gbps	17.75 Gbps
🟥 Incoming HTTP (<512K)	1.32 Gbps	3.46 Gbps	1.34 Gbps
🟩 Incoming HTTP	10.25 Gbps	11.67 Gbps	10.65 Gbps
Total Incoming	11.56 Gbps	13.14 Gbps	11.99 Gbps
■ Total Bandwidth Generation	5.47 Gbps	6.27 Gbps	5.76 Gbps

System Events

```
Apr 30 02:11:32  mg-1  cluster has been degraded
```

Replacing a Cache Engine and Running the Installation Tasks

The first step in replacing a cache engine in a VDS TC C-Series cluster installation is to physically replace the cache engine and perform the installation steps on the new cache engine. Follow this procedure to perform these steps:

Obtaining Required Files

Before you replace the VDS TC cache engine, you must obtain the following:

- **The VDS TC 5.2.0 ISO image file:** VDS-TC_Installer-5.2.0b123-5.2.0b124-ISO-5.2.0b23-Cisco.iso
- **The updated kernel load for VDS TC 5.2.1:** linux-2.6.27.19-5-llpf_10--KERNEL_LOAD_1414071783-r44244.tar.
- **The VDS TC 5.7.3 ISO image file:** VDS-TC_Installer_5.7.3b53-5.7.3b54-ISO-7-Cisco.iso
- **A new VDS TC license from Cisco:** To replace a failed cache engine you must request a new VDS TC license from Cisco. To request a new license, you must have the serial number of the failed cache engine and the serial number of the new cache engine. After you have gathered this information, contact Cisco support for a new license file.

Caution Do not continue until you have theses files.

Procedure

Note Before removing the failed cache engine, make note of what number corresponds with the cache engine.

Step 1 Using the Cisco UCS C220M3 documentation, unmount the failed cache engine and mount the new cache engine.

Step 2 Confirm that the correct firmware is installed on the Cache Engine for VDS TC 5.7.3:

- Cisco UCS C220-M3S:
  - Firmware version: 2.0(6f)
  - Image file: ucs-c220-huu-2.0.6f.iso

Step 3 In the *Cisco Videoscape Distribution Suite Transparent Caching Software C-Series Cluster Installation Guide*, available at http://www.cisco.com/c/dam/en/us/td/docs/video/videoscape/distribution_suite/vds/v5_7_3/VDS-TC_5_7_3_CSeries_swInstallGuide.pdf, refer to Chapter 1, “Prerequisites and C-Series Cluster Physical Installation” for connecting the network cables to the new cache engine.

Step 4 On the VDS TC Manager, confirm that the VDS-TC_Installer-5.2.0b123-5.2.0b124-ISO-5.2.0b23-Cisco.iso file is in the /opt/pang/iso folder. If it is not, use SFTP software, such as WinSCP, to copy the file to the /opt/pang/iso folder.

Step 5 Configure the new VDS TC cache engine by performing *all* of the tasks in Chapter 4, “Cache Engine Servers Installation” in the *Cisco Videoscape Distribution Suite Transparent Caching Software C-Series Cluster Installation Guide* for VDS TC Release 5.2.0.

Note Make sure you are using the installation guide for VDS TC Release 5.2.0 and that you perform all of the Steps in Chapter 4, including the “Running the Post Installation Scripts” section.

Note When the “Enter Cache Engine Number” dialog box appears, enter the number that corresponds with the cache engine that you are replacing.

Figure 2-2 Enter Cache Engine Number

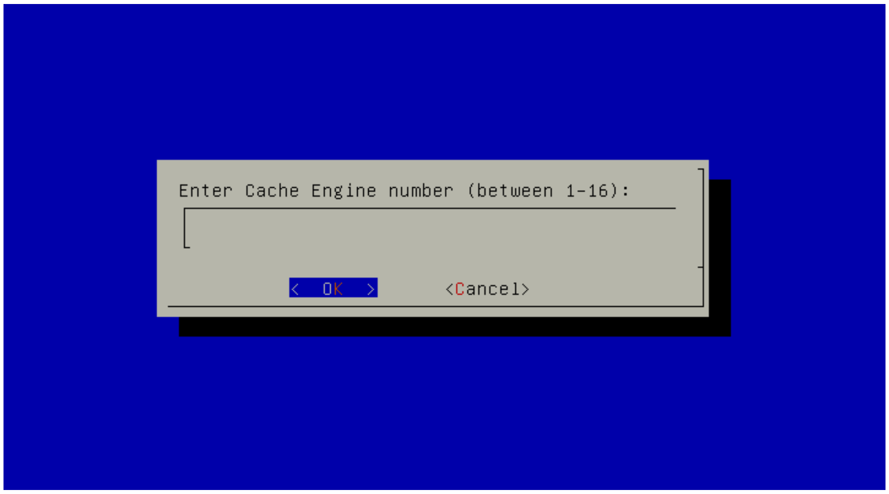

Enter Cache Engine number (between 1-16):

Updating the SSH Keys and Connecting the Cache Engine to the Storage Enclosure

After completing the "Cache Engine Servers Installation" chapter, you must update the grid SSH keys on the VDS TC management server and configure the new cache engine to connect to the storage enclosure. All of these steps are performed on the VDS TC management server.

Follow this procedure to perform these steps on the VDS TC management server:

Note The following steps are performed on the VDS TC management server.

Procedure

Step 1 On the PBR routers, stop redirecting traffic to the VDS TC solution.

Step 2 Close any open VDS TC Manager windows and ensure that no other administrators are connected to the VDS TC Manager before proceeding.

Step 3 Using SSH software, such as Putty, open an SSH connection to the IP address of the VDS TC management server.

Step 4 Log into the system using the username **padmin** and the password provided by Cisco.

Step 5 Enter the command **su admin** to log into the CLI of the VDS TC management server. The default password is the serial number of the VDS TC management server.

Step 6 From the VDS TC prompt, enter the **enable** command. When prompted, enter the Enable mode password and press **Enter**. You are now logged into Enable mode and the Enable prompt, console#, should appear.

Step 7 From the Enable mode prompt, enter the command **oper service stop** to stop the caching service.

Step 8 Wait a few minutes and then enter the command **show status**. Check to see if the Device State shows "stopped." For example:

```
console# show status
```

```
Operational state Device state Administrative state
disabledstoppedunlocked
```

Caution Do not proceed to the next step until the Device State shows “stopped.” You may need to repeat the **show status** command several times before you see this status. Wait several minutes between executions of the **show status** command.

Step 9 Enter the **exit** command to exit Enabled mode.

Step 10 Enter the **exit** command to exit the VDS TC CLI.

Step 11 Enter the command **su root** to change to the root user. Enter the password for the root user when prompted. This password is provided by Cisco.

Step 12 Enter the command **cd /opt/pang/useful/**.

Step 13 To reset the SSH keys, enter the command **./replace_server_keys.sh -servers** *#*, where # is the total number of C-Series cache engines in the VDS TC installation. For example, if you are replacing a cache engine in a VDS TC solution that has 16 cache engines, you would enter the command **./replace_server_keys.sh -servers 16**.

Step 14 Enter the command **cd /opt/pang/useful/configure_storages** to change to the storage configuration script directory.

Step 15 Enter the command **./configure_storages.py -s** <*no_of_storage_enclosures*> **-b** <*no_of_blades_in_grid*> **-c** <*the_replaced_CE-ID_number*>. For example, to replace CE-2 in a VDS TC C-Series installation that has two cache engines, and one storage enclosure, enter the command **./configure_storages.py -s 1 -b 2 -c 2**.

Note The configure storages script will format the data disks. To configure the storages without formatting the disks enter the command **./configure_storages.py -s** <*no_of_storage_enclosures*> **-b** <*no_of_blades_in_grid*> **-c** <*the_replaced_CE-ID_number*> **-x**.

Step 16 When the warning message about the disk format appears, enter **y** to continue the storages configuration.

```
*********** WARNING! **********
All storages are about to be reconfigured. All data on the data disks will
be lost.
Do you want to continue? (y/n)y
```

Note You may be prompted to enter the password several times.

When the script is finished running, you should see the word “COMPLETE”. After running this script, the new VDS TC cache engine will be connected to the storage enclosures.

Step 17 Use SSH to connect to each CE and enter the command **iscsiadm -m session** to check the connections:

```
ssh ce-<cache engine number>
iscsiadm -m session
```

Your output should be similar to the following:

```
ƒÞ 8 * sessions per storage, for example for 1 storage:
tcp: [17] 10.11.14.100:3260,1 iqn.1992-
01.com.Avago:2365.60080e50002f4e82000000005063c74a
tcp: [18] 10.11.14.101:3260,2 iqn.1992-
01.com.Avago:2365.60080e50002f4e82000000005063c74a
```

```
tcp: [19] 10.11.16.100:3260,1 iqn.1992-
01.com.Avago:2365.60080e50002f4e82000000005063c74a
tcp: [20] 10.11.16.101:3260,2 iqn.1992-
01.com.Avago:2365.60080e50002f4e82000000005063c74a
tcp: [21] 10.11.15.100:3260,1 iqn.1992-
01.com.Avago:2365.60080e50002f4e82000000005063c74a
tcp: [22] 10.11.15.101:3260,2 iqn.1992-
01.com.Avago:2365.60080e50002f4e82000000005063c74a
tcp: [23] 10.11.17.101:3260,2 iqn.1992-
01.com.Avago:2365.60080e50002f4e82000000005063c74a
tcp: [24] 10.11.17.100:3260,1 iqn.1992-
01.com.Avago:2365.60080e50002f4e82000000005063c74a
```

Updating the NetApp Storage Enclosure

Note This process is only for the NetApp storage enclosure. If your system uses an IBM storage enclosure, refer to Updating the IBM Storage Enclosure section.

After you have configured the VDS TC cache engine to connect with the storage enclosure, you must update the host port identifiers on the NetApp storage enclosure for the new cache engine. Follow this procedure to perform this task:

Procedure

Step 1 From the SSH connection to the VDS TC management server, enter the following commands to start the VNC server:

a. **rm /tmp/.X1-lock** (This step is optional. It is only required if there is a leftover lock file from a previous vncserver instance.)

b. **rm /tmp/.X11-unix/X1** (This step is optional. It is only required if there is a leftover lock file from a previous vncserver instance.)

c. **vncserver :1**

Step 2 Use a VNC client to connect from your local computer to the VDS TC management server.

Figure 2-3 VNC Viewer

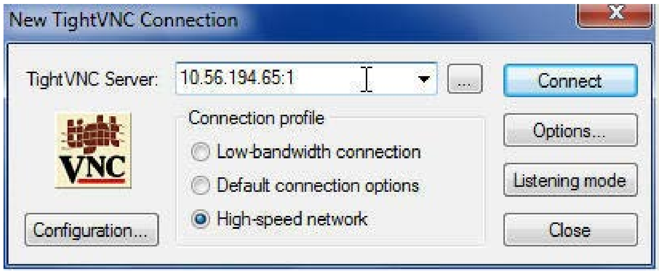

New TightVNC Connection

TightVNC Server: 10.56.194.65:1

Connection profile

○ Low-bandwidth connection

○ Default connection options

● High-speed network

Step 3 From the VNC console window, enter the command **cd /opt/SMgr/client** to change folders.

Step 4 Enter the command **./SMclient** to start the SANtricity ™ ES Storage Manager software.

Note You must perform the following steps on all of the storage subsystems.

Step 5 Right-click the storage subsystem that you are configuring and choose **Manage Storage Array**. Click **No** in any pop-up windows that appear.

Figure 2-4 SANtricity Storage Manager

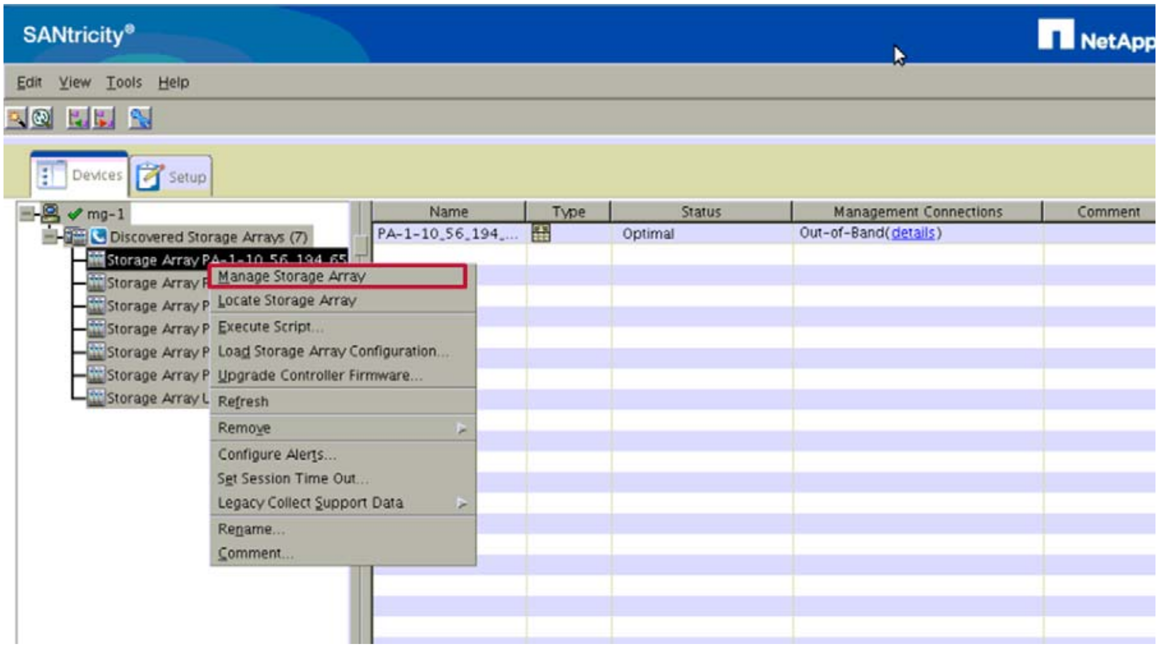

Name	Type	Status	Management Connections	Comment
PA-1-10.56.194....		Optimal	Out-of-Band(details)	

Step 6 In the Subsystem Management window that appears, click the **Host Mappings** tab.

Step 7 In the navigation pane on the left, expand the **Host Group** folder.

Step 8 Right-click the cache engine that was replaced, for example Host CE-1, and choose **Manage Host Port Identifiers** from the pop-up menu.

Figure 2-5 Manage Host Port Identifiers

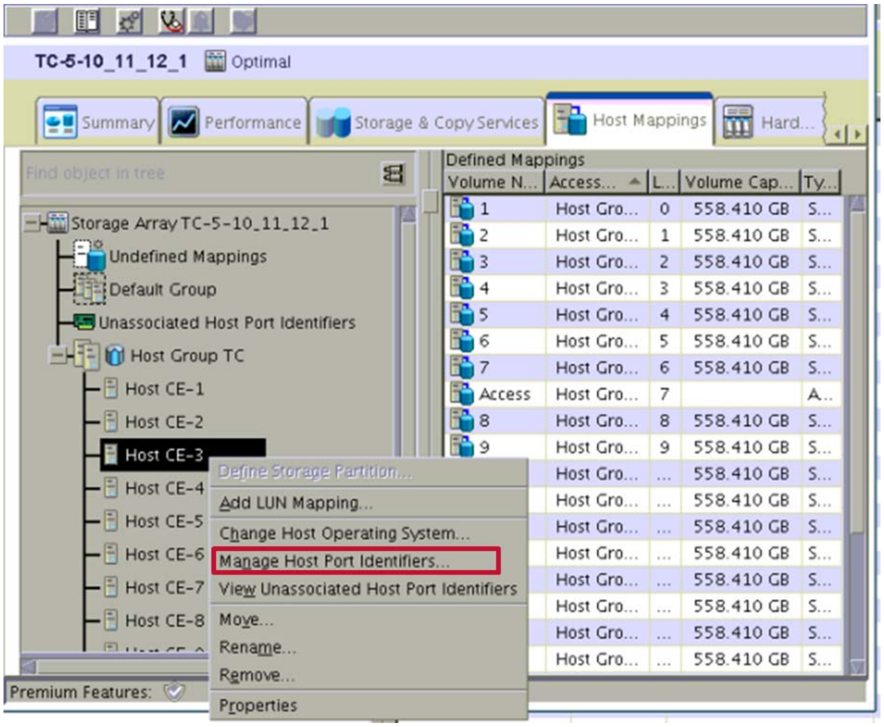

Volume N...	Access...	L...	Volume Cap...	Ty...
1	Host Gro...	0	558.410 GB	S...
2	Host Gro...	1	558.410 GB	S...
3	Host Gro...	2	558.410 GB	S...
4	Host Gro...	3	558.410 GB	S...
5	Host Gro...	4	558.410 GB	S...
6	Host Gro...	5	558.410 GB	S...
7	Host Gro...	6	558.410 GB	S...
Access	Host Gro...	7		A...
8	Host Gro...	8	558.410 GB	S...
9	Host Gro...	9	558.410 GB	S...
	Host Gro...	...	558.410 GB	S...
	Host Gro...	...	558.410 GB	S...
	Host Gro...	...	558.410 GB	S...
	Host Gro...	...	558.410 GB	S...
	Host Gro...	...	558.410 GB	S...
	Host Gro...	...	558.410 GB	S...
	Host Gro...	...	558.410 GB	S...
	Host Gro...	...	558.410 GB	S...

Step 9 From the Manage Host Port Identifiers window that appears, in the Host Port Identifier Information table, ensure that the old iqn is selected and click **Replace**.

Figure 2-6 Replace IQN

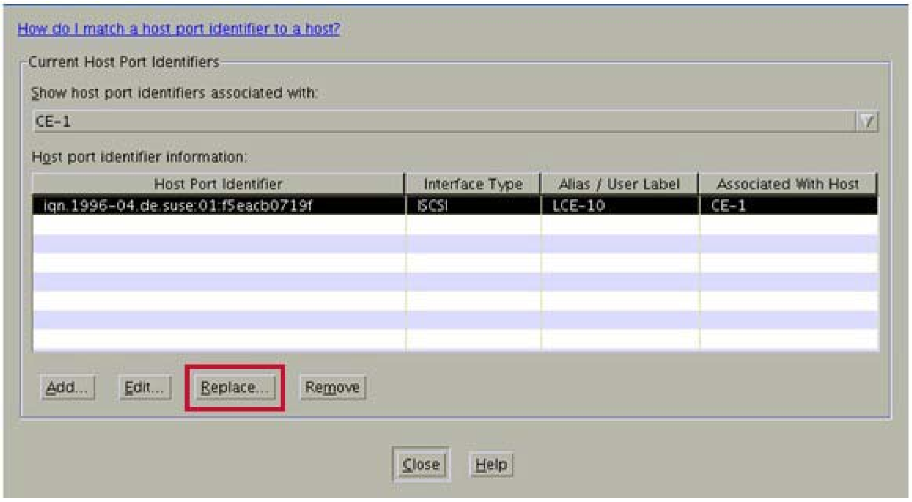

How do I match a host port identifier to a host?
Current Host Port Identifiers
Add...
Edit...
Replace...
Remove
Close
Help

Show host port identifiers associated with:

CE-1

Host port identifier information:

Host Port Identifier	Interface Type	Alias / User Label	Associated With Host
iqn.1996-04.de.suse:01:f5eacb0719f	iSCSI	LCE-10	CE-1

Step 10 From the Replace Host Port Identifier window that appears, configure the following:

- From the Choose a Host Interface Type drop-down list, choose **ISCSI**.
- Click the **Replace by Selecting a Known Unassociated Host Port Identifier** radio button.

- From the Known Unassociated Host Port Identifier drop-down list, choose the iqn for the *newly* installed cache engine. (It should be the only one in the list.)

Step 11 Click **Replace**.

Figure 2-7 Replace Host Port Identifier

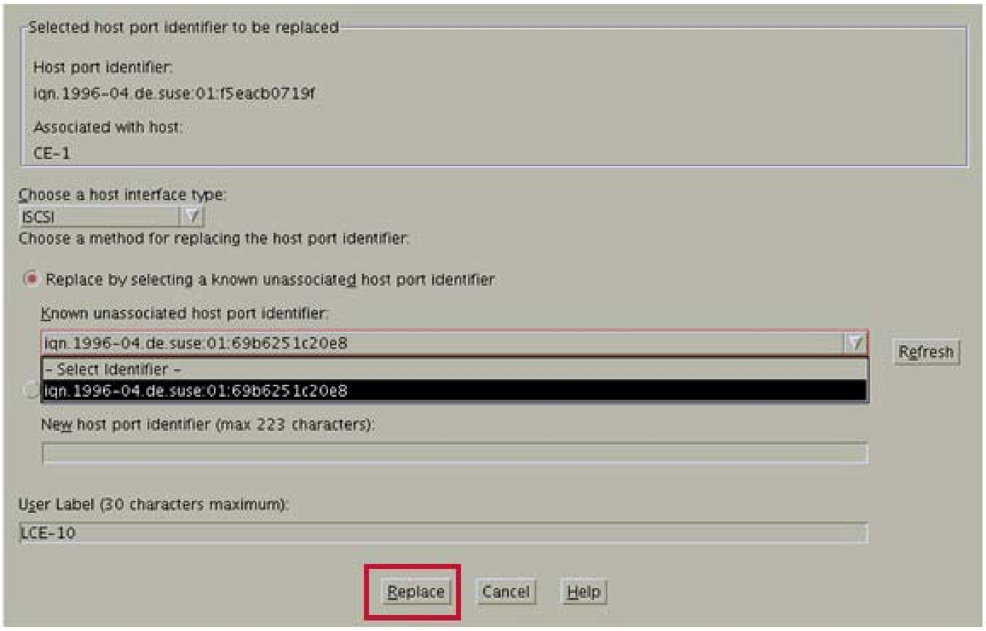

Selected host port identifier to be replaced

Host port identifier:
iqn.1996-04.de.suse:01:f5eacb0719f

Associated with host:
CE-1

Choose a host interface type:
iSCSI

Choose a method for replacing the host port identifier:

● Replace by selecting a known unassociated host port identifier

Known unassociated host port identifier:
iqn.1996-04.de.suse:01:69b6251c20e8
- Select Identifier -
iqn.1996-04.de.suse:01:69b6251c20e8

○ New host port identifier (max 223 characters):

User Label (30 characters maximum):
LCE-10

Step 12 From the Manage Host Port Identifiers window, click **Close**.

Step 13 Repeat Step 5 through Step 12 on any remaining storage subsystems.

Step 14 Close the Subsystem Management window and exit the SANtricity Storage Manager.

Step 15 Close the VNC client.

Updating the IBM Storage Enclosure

Note This process is only for the IBM storage enclosure. If your system uses the NetApp E2724 storage enclosure, refer to Updating the NetApp Storage Enclosure section.

After you have configured the VDS TC cache engine to connect with the storage enclosure, you must update the host port identifiers on the IBM storage enclosure for the new cache engine. Follow this procedure to perform this task:

Procedure

Step 1 From the SSH connection to the VDS TC management server, enter the following commands to start the VNC server:

a. **rm /tmp/.X1-lock** (This step is optional. It is only required if there is a leftover lock file from a previous vncserver instance.)

b. **rm /tmp/.X11-unix/X1** (This step is optional. It is only required if there is a leftover lock file from a previous vncserver instance.)

c. **vncserver :1**

Step 2 Use a VNC client to connect from your local computer to the VDS TC management server.

Figure 2-8 VNC Viewer

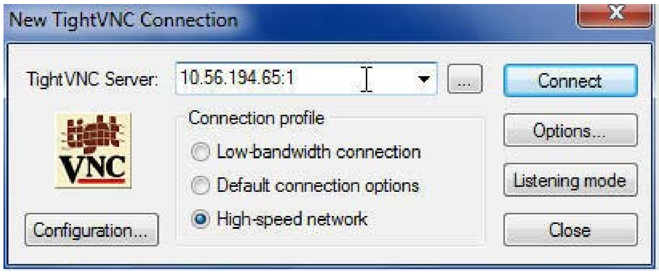

New TightVNC Connection

TightVNC Server: 10.56.194.65:1

Connection profile

○ Low-bandwidth connection
○ Default connection options
● High-speed network

Step 3 From the VNC console window, enter the command **/opt/IBM_DS/client/SMclient** to start the IBM Storage Manager.

Note You must perform the following steps on all of the storage subsystems.

Step 4 Right-click the storage subsystem that you are configuring and choose **Manage Storage Subsystem**. Click **No** in any pop-up windows that appear.

Figure 2-9 IBM DS Storage Manager

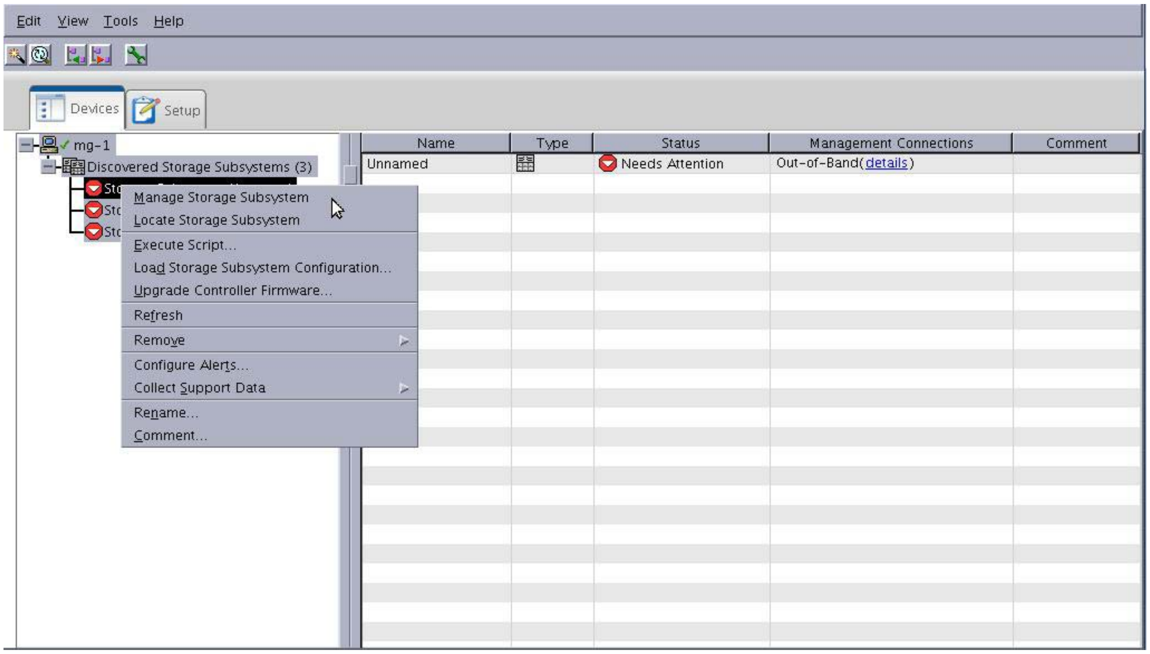

Name	Type	Status	Management Connections	Comment
Unnamed		Needs Attention	Out-of-Band(details)	

Step 5 In the Subsystem Management window that appears, click the **Host Mappings** tab.

Step 6 In the navigation pane on the left expand the **Host Group** folder.

Step 7 Right-click the cache engine that was replaced, for example Host CE-1, and choose **Manage Host Port Identifiers** from the pop-up menu.

Figure 2-10 Manage Host Port Identifiers

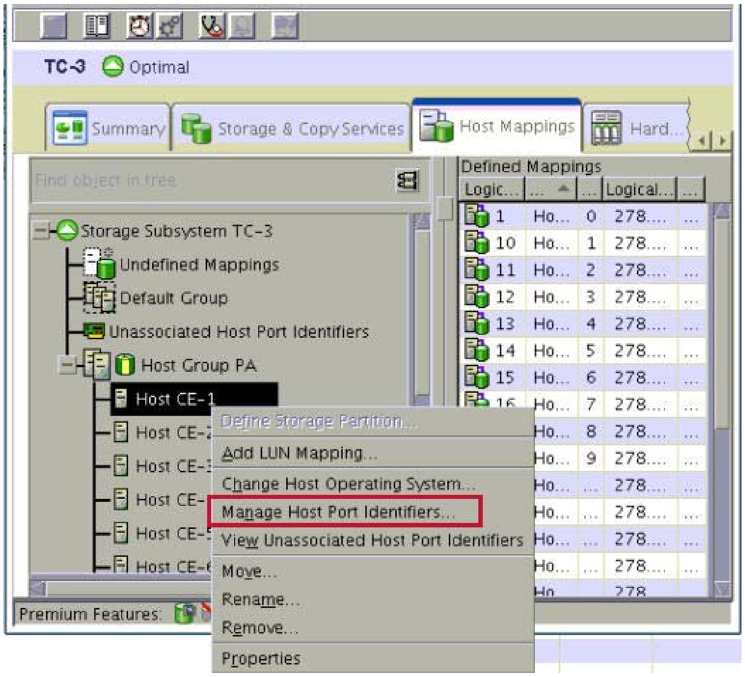

Defined Mappings

Logic...	...	...	Logical...	...
1	Ho...	0	278....	...
10	Ho...	1	278....	...
11	Ho...	2	278....	...
12	Ho...	3	278....	...
13	Ho...	4	278....	...
14	Ho...	5	278....	...
15	Ho...	6	278....	...
16	Ho...	7	278....	...
	Ho...	8	278....	...
	Ho...	9	278....	...
	Ho...	...	278....	...
	Ho...	...	278....	...
	Ho...	...	278....	...
	Ho...	...	278....	...
	Ho		278	

Step 8 From the Manage Host Port Identifiers window that appears, in the Host Port Identifier Information table, ensure that the old iqn is selected and click **Replace**.

Figure 2-11 Replace IQN

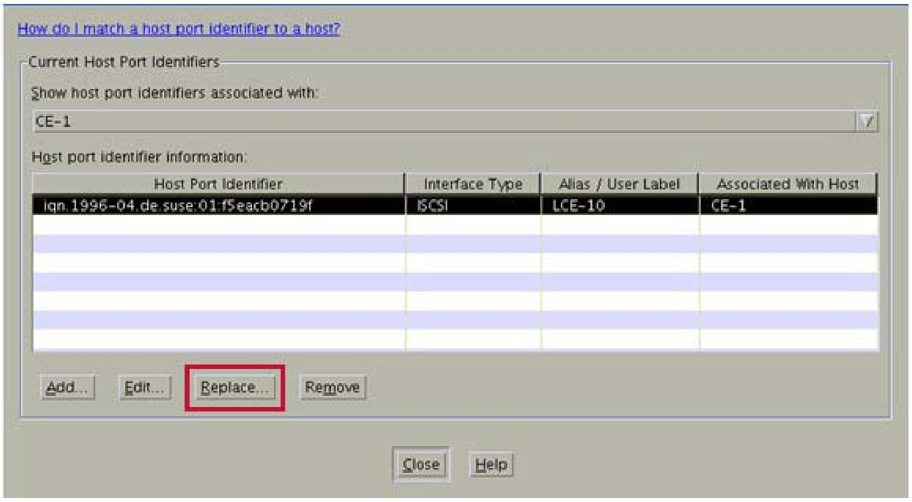
How do I match a host port identifier to a host?

Current Host Port Identifiers

Show host port identifiers associated with:

CE-1

Host port identifier information:

Host Port Identifier	Interface Type	Alias / User Label	Associated With Host
iqn.1996-04.de.suse:01:f5eacb0719f	ISCSI	LCE-10	CE-1

Step 9 From the Replace Host Port Identifier window that appears, configure the following:

- From the Choose a Host Interface Type drop-down list, choose **ISCSI**.
- Click the **Replace by Selecting a Known Unassociated Host Port Identifier** radio button.

- From the Known Unassociated Host Port Identifier drop-down list, choose the iqn for the *newly* installed cache engine. (It should be the only one in the list.)

Step 10 Click **Replace**.

Figure 2-12 Replace Host Port Identifier

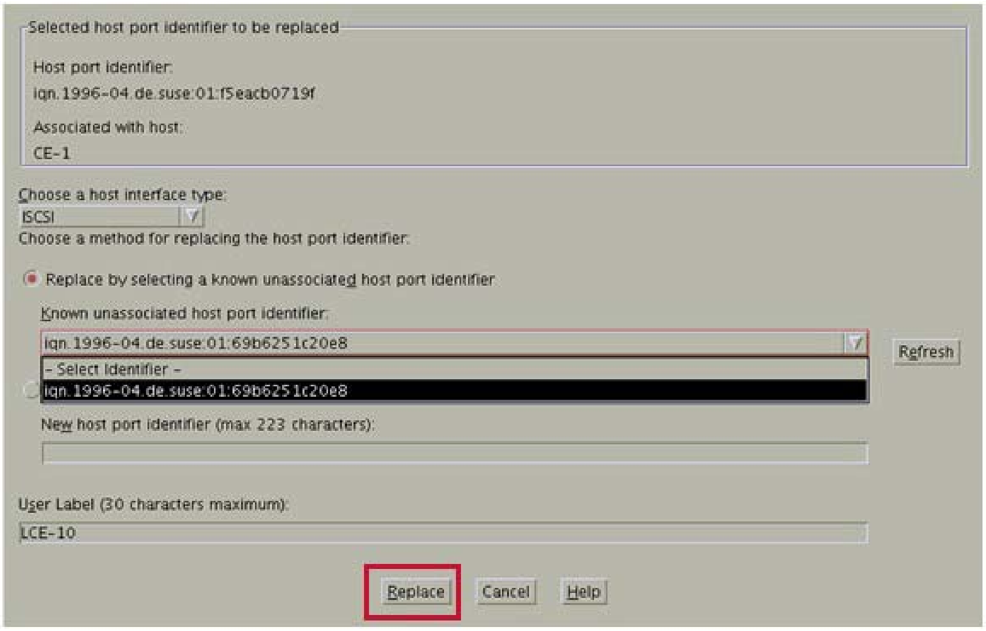

Selected host port identifier to be replaced

Host port identifier:
iqn.1996-04.de.suse:01:f5eacb0719f

Associated with host:
CE-1

Choose a host interface type:
ISCSI

Choose a method for replacing the host port identifier:

● Replace by selecting a known unassociated host port identifier

Known unassociated host port identifier:
iqn.1996-04.de.suse:01:69b6251c20e8
- Select Identifier -
iqn.1996-04.de.suse:01:69b6251c20e8

New host port identifier (max 223 characters):

User Label (30 characters maximum):
LCE-10

Step 11 From the Manage Host Port Identifiers window, click **Close**.

Step 12 Repeat Step 4 through Step 11 on any remaining storage subsystems.

Step 13 Close the Subsystem Management window and exit the IBM Storage Manager.

Step 14 Close the VNC client.

Finishing the Cache Engine Replacement

Follow this procedure to finish configuring and installing the new cache engine:

Note You will be re-running the storage script that you ran in the Updating the SSH Keys and Connecting the Cache Engine to the Storage Enclosure section.

Note You must be logged in with root permissions to perform the following steps.

Procedure

Step 1 From the SSH connection to the VDS TC management server, enter the command **cd /opt/pang/useful/configure_storages** to change to the storage configuration script directory.

Step 2 Enter the command **./configure_storages.py -s** *<no_of_storage_enclosures>* **-b** *<no_of_blades>* **-c** *<the_replaced_CE-ID_number>*. For example, to replace CE-2 in a VDS TC C-Series installation that has two cache engines, and one storage enclosure, enter the command **./configure_storages_py -s 1 -b 2 -c 2**.

Step 3 When the warning message about the disk format appears, enter **y** to continue the storages configuration.

```
*********** WARNING! **********
All storages are about to be reconfigured. All data on the data disks will
be lost.
Do you want to continue? (y/n)y
```

Step 4 Follow the on screen instructions. If prompted, enter the password for the root user that was provided by Cisco. You should see the word “COMPLETE” when the script is finished running.

Note You may be prompted to enter the password several times.

Note Ensure that the VDS TC image upgrade file and the *new* VDS TC license that you received from Cisco are located on the VDS TC management server in the /tftpboot folder before performing these steps. Please note that you must receive a new license from Cisco support that includes the serial number of the replacement cache engine.

Step 5 Using SFTP software, such as WinSCP, connect to the management IP address that was assigned to the VDS TC server. Log in using the user name **padmin** and the password that was provided by Cisco and do the following:

a. Copy the updated kernel load file, linux-2.6.27.19-5-llpf_10--KERNEL_LOAD_1414071783-r44244.tar, to the **/tmp** folder.

b. Copy the VDS TC 5.7.3 ISO image file, VDS-TC_Installer_5.7.3b53-5.7.3b54-ISO-7-Cisco.iso, to the **/opt/pang/iso** folder.

c. Copy the new VDS TC license that you received from Cisco to the **/tftpboot** folder.

d. Close the SFTP software.

Step 6 From an SSH connection to the VDS TC management server logged in as padmin, enter the command **su root**. When prompted, enter the password that was provided by Cisco. Perform the following steps to extract the FI file for VDS TC 5.7.3:

a. Enter the command **mkdir /opt/pang/iso/5.7.3_iso** to create a new folder for the 5.7.3 files.

b. Enter the command **cd /opt/pang/iso** to change folders.

c. Enter the command **mount VDS-TC_Installer_5.7.3b53-5.7.3b54-ISO-7-Cisco.iso -o loop /opt/pang/iso/5.7.3_iso** to extract the VDS-TC_FI_5.7.3b54_Server_Cluster.tar.gz file.

d. Enter the command **cd /** to change to the root folder and then enter the command **umount /opt/pang/iso/5.7.3_iso** to unmount the iso image.

e. Enter the command **rm /opt/pang/iso/VDS-TC_Installer_5.7.3b53-5.7.3b54-ISO-7-Cisco.iso** to delete the iso file.

f. Enter the command **cd /opt/pang/iso/5.7.3_iso/install-tools/PeerApp_FI-GA**.

g. Enter the command **cp VDS-TC_FI_5.7.3b54_Server_Cluster.tar.gz /tftpboot**.

Step 7 Next you will update the Kernel on the cache engine. Enter the command **cd /tmp**.

Step 8 To copy the updated kernel installation package file to the new cache engine, enter the command **scp linux-2.6.27.19-5-llpf_10--KERNEL_LOAD_1414071783-r44244.tar root@ce-#:/tmp**, where # is the number of the cache engine that you replaced. For example, if you replaced cache engine 1, enter the following:

```
scp linux-2.6.27.19-5-llpf_10--KERNEL_LOAD_1414071783-r44244.tar root@ce-1:/tmp
```

Step 9 From the VDS TC management server, enter the command **ssh ce-*X***, where *X* is the number of the cache engine that you replaced.

Step 10 Enter the command **cd /tmp** to change to the /tmp folder.

Step 11 Enter the command **tar -xvf linux-2.6.27.19-5-llpf_10--KERNEL_LOAD_1414071783-r44244.tar** to extract the kernel files.

Step 12 Enter the command **cd linux-2.6.27.19-5-llpf_10--KERNEL_LOAD_1414071783** to change to the new folder that was created when extracting the files in Step 6.

Step 13 Enter the command **./linux-2.6.27.19-5-llpf_10--INSTALL_PHASE-1.sh** to run the first phase of the kernel installation.

Step 14 Each time you see the following prompt, press **Enter**:

```
Press [Enter] key...
```

Note You will see this prompt several times.

Step 15 When you see the following prompt, press **Enter**. This will cause the system to reboot for the first time. When the system reboots, your SSH connection to the cache engine will close.

```
Going to REBOOT, Press [Enter] key...
```

Step 16 After the cache engine reboots, from the VDS TC management server, enter the command **ssh ce-*X*** to reconnect to the cache engine, where *X* is the number of the cache engine to which are reconnecting.

Step 17 Enter the command **cd /tmp**.

Step 18 Enter the command **cd linux-2.6.27.19-5-llpf_10--KERNEL_LOAD_1414071783**.

Step 19 Enter the command **./linux-2.6.27.19-5-llpf_10--INSTALL_PHASE-2_Cisco.sh** to run the second phase of the kernel installation.

Step 20 Each time you see the following prompt, press **Enter**:

```
Press [Enter] key...
```

Step 21 When you see the following prompt, press **Enter**. This will cause the system to reboot. When the system reboots, your SSH connection to the cache engine will close.

```
Going to REBOOT, Press [Enter] key...
```

Step 22 After the cache engine reboots, from the VDS TC management server, enter the command **ssh ce-*X***, where *X* is the number of the cache engine to which you are reconnecting.

Step 23 To confirm the kernel update, enter the command **uname –a** and verify that the date Tue Aug 19 16:27:56 GMT 2014 appears in the output, as shown in the following example.

```
Linux ce-1 2.6.27.19-llpf_10-5-default #26 SMP Tue Aug 19 16:27:56 GMT 2014 x86_64 x86_64
x86_64 GNU/Linux
```

Step 24 Enter **logout** to close the SSH connection to the cache engine.

Step 25 Enter the command **su admin** to log into the VDS TC CLI. The CLI prompt console> appears.

Step 26 Enter the command **enable** to enter Enable mode. At the password prompt that appears, enter the enable password. The CLI prompt console# appears.

Step 27 Upgrade the software on the new cache engine by entering the command **upgrade server** *CE_number* **127.0.0.1 VDS-TC_FI_5.7.3b54_Server_Cluster.tar.gz**, where *CE_number* is the cache engine you replaced.

Step 28 To import the new license that includes the serial number of the new cache engine, enter the command **license import 127.0.0.1** *filename*, where *filename* is the name of the new license file.

Note You must be in Enable mode, *not* Configuration mode to import a new license.

Step 29 Enter the command **license activate** to apply and activate the license.

Step 30 To start the application on the replaced cache engine, enter the command **oper server** *#*, where *#* is the number of the cache engine that was replaced. For example, if you replaced CE-2, enter the command **oper server 2**. You should see the following output:

```
console# oper server 2
oper server 2# start
Starting server 2
service is started on server 2
```

Step 31 After the application has started on the new cache engine, enter the command **oper service start** to start the VDS TC service.

Step 32 Enter the command **show cluster-bus-ip** to check if the spread communication among cluster members is set to broadcast.

```
console# show cluster-bus-ip
Processing...
MG-1: [Broadcast]
CE-1: [Broadcast]
CE-2: [Broadcast]
CE-3: [Broadcast]
CE-4: [Broadcast]
CE-5: [Broadcast]
CE-6: [Broadcast]
CE-7: [Broadcast]
CE-8: [Broadcast]
OK. All the CEs are configured to work with Broadcast IP.
```

Step 33 If the communication is set to broadcast, continue on to Step 36. If the communication is set to *multicast* instead of *broadcast*, enter the command **config** to switch to Configuration mode and then enter the command **cluster-bus-ip broadcast**. For example:

```
configuration# cluster-bus-ip broadcast
Processing...
All the CEs are configured to work with Broadcast IP.
```

Step 34 Enter the command **exit** to exit the Configuration mode.

Step 35 Enter the command **show cluster-bus-ip** to confirm that the management server and all cache engines are configured to use broadcast. For example:

```
console# show cluster-bus-ip
Processing...
MG-1: [Broadcast]
CE-1: [Broadcast]
CE-2: [Broadcast]
CE-3: [Broadcast]
CE-4: [Broadcast]
```

```
CE-5: [Broadcast]
CE-6: [Broadcast]
CE-7: [Broadcast]
CE-8: [Broadcast]
OK. All the CEs are configured to work with Broadcast IP.
```

Step 36 Wait a few minutes and then enter the command **show status**. Do not proceed to the next step until you see a Device Status of “Started.” For example:

```
console# show status
Operational state Device state Administrative state
enabled started unlocked
```

Step 37 On the PBR routers, start redirecting traffic to the VDS TC solution.

CHAPTER 3

Cache Engine Replacement Procedure in a 5.7.3 C-Series Installation Upgraded from Release 5.6.1

Note Only use this chapter if you are replacing a Cache Engine in a VDS TC C-Series system that was upgraded to VDS TC 5.7.3 from VDS TC 5.6.1. If you are replacing a Cache Engine in a VDS TC C-Series system that has a fresh install of VDS TC 5.7.3 and was *not* upgraded to 5.7.3, refer to Chapter 1, "Cache Engine Replacement Procedure for a C-Series Installation with a Fresh Install of VDS TC 5.7.3". If you are replacing a Cache Engine in a VDS TC C-Series system that was upgraded to VDS TC 5.7.3 from VDS TC 5.2.1, refer to Chapter 2, "Cache Engine Replacement Procedure in a 5.7.3 C-Series Installation Upgraded from Release 5.2.1".

To determine whether the system has been upgraded and from what version, from the Management Server enter the command **grep "upgrade system to" /var/log/peerapp/peerapp_system***. This will display a list of all previous versions.

This chapter discusses how to determine that a cache engine has failed in a VDS TC C-Series installation and the steps that are required to replace it.

When a cache engine fails, the volume_selection_algorithm_time setting in the cluster_conf.xml file controls when the volumes of a failed cache engine will be redistributed (mounted) to the remaining cache engines. By default this value is 24 hours, which means that by default, if a cache engine fails, its volumes will not be redistributed (mounted) to the remaining cache engines until 24 hours after the cache engine fails.

It is recommended that you leave the default set to 24 hours, however if you need to change this setting, add the following text to the cluster_conf file in the <service> section under the <memory> subsection. The *seconds* parameters is the number of seconds that the system should wait after a cache engine fails before its volumes are redistributed (mounted) to the remaining cache engines:

```
<io>
    <volume_selection_algorithm_time>seconds</volume_selection_algorithm_time>
</io>
```

For example, to have the system wait 30 minutes after a cache engine fails before redistributing the volumes of the failed cache engine enter the following:

```
<io>
    <volume_selection_algorithm_time>1800</volume_selection_algorithm_time>
</io>
```

For more information on how to make changes to the configuration file and how to apply these changes, please refer to the *Cisco Videoscape Distribution Suite Transparent Caching Software Configuration Guide* available at http://www.cisco.com/c/dam/en/us/td/docs/video/videoscape/distribution_suite/vds/v5_7_3/VDS-TC_5_7_3_software_config_guide.pdf.

Symptoms of a Failed VDS TC C-Series Cache Engine

To determine whether a cache engine has failed in a VDS TC C-Series Server installation, log into VDS TC Manager. By default, the Status Dashboard displays. If you are already logged into VDS TC Manager GUI, choose **Status > Dashboard**.

If a cache engine has failed it will appear with a red box surrounding it, as shown with the first cache engine in the following figure. If a cache engine cannot bootup, it will appear with a grey box surrounding it.

Figure 3-1 Failed Cache Engine

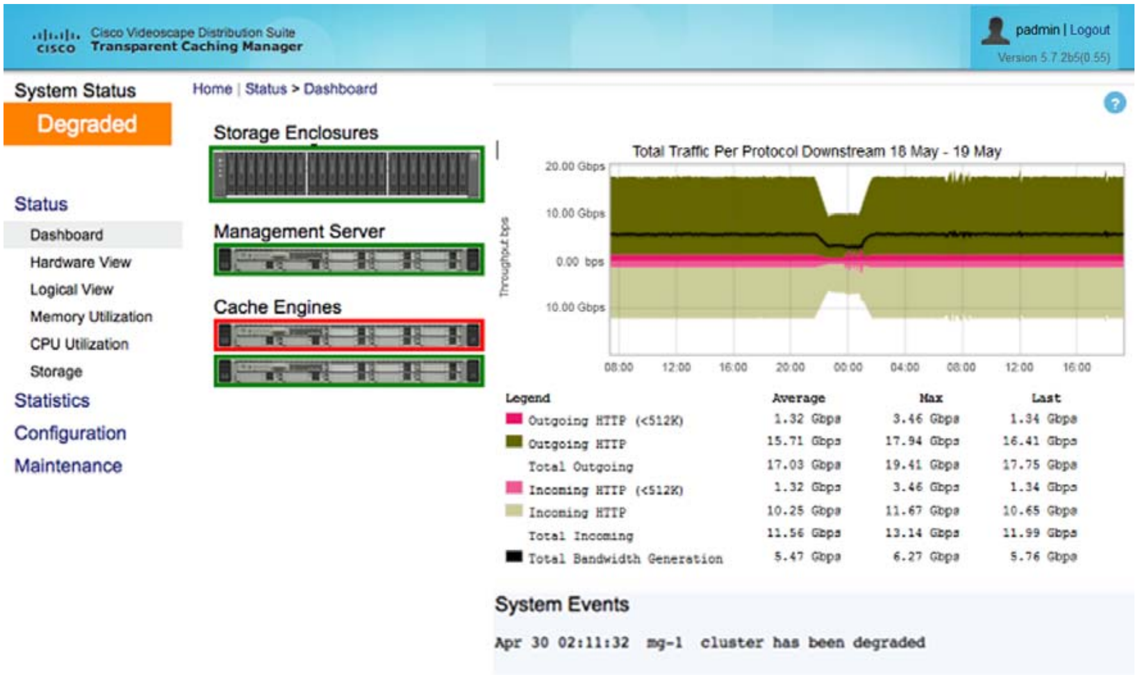

Legend	Average	Max	Last
🟥 Outgoing HTTP (<512K)	1.32 Gbps	3.46 Gbps	1.34 Gbps
🟫 Outgoing HTTP	15.71 Gbps	17.94 Gbps	16.41 Gbps
Total Outgoing	17.03 Gbps	19.41 Gbps	17.75 Gbps
🟪 Incoming HTTP (<512K)	1.32 Gbps	3.46 Gbps	1.34 Gbps
🟨 Incoming HTTP	10.25 Gbps	11.67 Gbps	10.65 Gbps
Total Incoming	11.56 Gbps	13.14 Gbps	11.99 Gbps
■ Total Bandwidth Generation	5.47 Gbps	6.27 Gbps	5.76 Gbps

System Events

```
Apr 30 02:11:32  mg-1  cluster has been degraded
```

Replacing a Cache Engine and Running the Installation Tasks

The first step in replacing a cache engine in a VDS TC C-Series cluster installation is to physically replace the cache engine and perform the installation steps on the new cache engine. Follow this procedure to perform these steps:

Obtaining Required Files

Before you replace the VDS TC cache engine, you must obtain the following:

- The VDS TC 5.6.1 ISO image file, VDS-TC_Installer_5.6.1b57-5.6.1b58-ISO-15-Cisco.iso
- The VDS TC 5.7.3 ISO image file, VDS-TC_Installer_5.7.3b53-5.7.3b54-ISO-7-Cisco.iso
- A new VDS TC license from Cisco. To replace a failed cache engine you must request a new VDS TC license from Cisco. To request a new license, you must have the serial number of the failed cache engine and the serial number of the new cache engine. After you have gathered this information, contact Cisco support for a new license file.

Caution Do not continue until you have theses files.

Procedure

Note Before removing the failed cache engine, make note of what number corresponds with the cache engine.

Step 1 Using the Cisco UCS C220M3 documentation, unmount the failed cache engine and mount the new cache engine.

Step 2 Confirm that the correct firmware is installed on the Cache Engine for VDS TC 5.7.3:

- Cisco UCS C220-M3S:
  - Firmware version: 2.0(6f)
  - Image file: ucs-c220-huu-2.0.6f.iso

Step 3 In the *Cisco Videoscape Distribution Suite Transparent Caching Software C-Series Cluster Installation Guide*, available at http://www.cisco.com/c/dam/en/us/td/docs/video/videoscape/distribution_suite/vds/v5_7_3/VDS-TC_5_7_3_CSeries_swInstallGuide.pdf, refer to Chapter 1, “Prerequisites and C-Series Cluster Physical Installation” for connecting the network cables to the new cache engine.

Step 4 On the VDS TC Manager, confirm that the VDS-TC_Installer_5.6.1b57-5.6.1b58-ISO-15-Cisco.iso file is in the /opt/pang/iso folder. If it is not, use SFTP software, such as WinSCP, to copy the file to the /opt/pang/iso folder.

Step 5 Configure the new VDS TC cache engine by performing all of the tasks in Chapter 4, “Cache Engine Servers Installation” in the *Cisco Videoscape Distribution Suite Transparent Caching Software C-Series Cluster Installation Guide* for VDS TC Release 5.6.1.

Note Make sure you are using the installation guide for VDS TC Release 5.6.1 and that you perform all of the Steps in Chapter 4, including the “Running the Post Installation Scripts” section.

Note When the “Enter Cache Engine Number” dialog box appears, enter the number that corresponds with the cache engine that you are replacing.

Figure 3-2 Enter Cache Engine Number

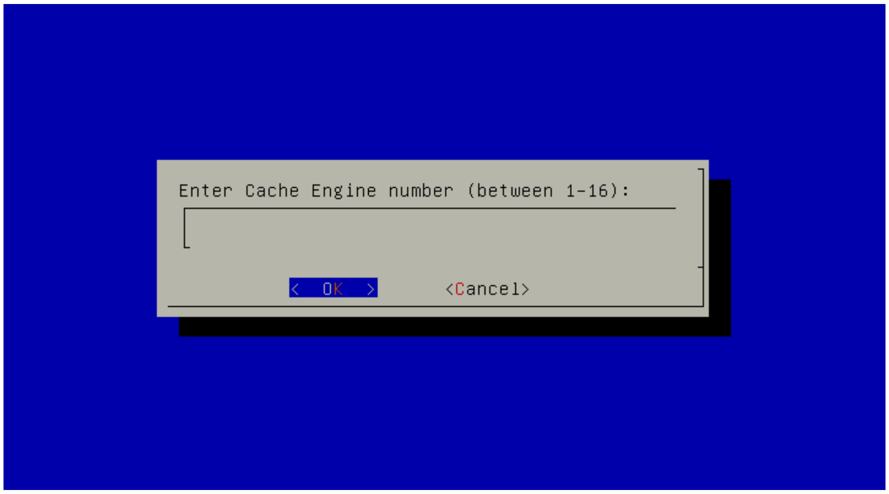

Enter Cache Engine number (between 1-16):

Updating the SSH Keys and Connecting the Cache Engine to the Storage Enclosure

After completing the "Cache Engine Servers Installation" chapter, you must update the grid SSH keys on the VDS TC management server and configure the new cache engine to connect to the storage enclosure. All of these steps are performed on the VDS TC management server.

Follow this procedure to perform these steps on the VDS TC management server:

Note The following steps are performed on the VDS TC management server.

Procedure

Step 1 On the PBR routers, stop redirecting traffic to the VDS TC solution.

Step 2 Close any open VDS TC Manager windows and ensure that no other administrators are connected to the VDS TC Manager before proceeding.

Step 3 Using SSH software, such as Putty, open an SSH connection to the IP address of the VDS TC management server.

Step 4 Log into the system using the username **padmin** and the password provided by Cisco.

Step 5 Enter the command **su admin** to log into the CLI of the VDS TC management server. The default password is the serial number of the VDS TC management server.

Step 6 From the VDS TC prompt, enter the **enable** command. When prompted, enter the Enable mode password and press **Enter**. You are now logged into Enable mode and the Enable prompt, console#, should appear.

Step 7 From the Enable mode prompt, enter the command **oper service stop** to stop the caching service.

Step 8 Wait a few minutes and then enter the command **show status**. Check to see if the Device State shows "stopped." For example:

```
console# show status
```

```
Operational state Device state Administrative state
disabledstoppedunlocked
```

Caution Do not proceed to the next step until the Device State shows “stopped.” You may need to repeat the **show status** command several times before you see this status. Wait several minutes between executions of the **show status** command.

Step 9 Enter the **exit** command to exit Enabled mode.

Step 10 Enter the **exit** command to exit the VDS TC CLI.

Step 11 Enter the command **su root** to change to the root user. Enter the password for the root user when prompted. This password is provided by Cisco.

Step 12 Enter the command **cd /opt/pang/useful/**.

Step 13 To reset the SSH keys, enter the command **./replace_server_keys.sh -servers** *#*, where # is the total number of C-Series cache engines in the VDS TC installation. For example, if you are replacing a cache engine in a VDS TC solution that has 16 cache engines, you would enter the command **./replace_server_keys.sh -servers 16**.

Step 14 Enter the command **cd /opt/pang/useful/configure_storages** to change to the storage configuration script directory.

Step 15 Enter the command **./configure_storages.py -s** <*no_of_storage_enclosures*> **-b** <*no_of_blades_in_grid*> **-c** <*the_replaced_CE-ID_number*>. For example, to replace CE-2 in a VDS TC C-Series installation that has two cache engines, and one storage enclosure, enter the command **./configure_storages.py -s 1 -b 2 -c 2**.

Note The configure storages script will format the data disks. To configure the storages without formatting the disks enter the command **./configure_storages.py -s** <*no_of_storage_enclosures*> **-b** <*no_of_blades_in_grid*> **-c** <*the_replaced_CE-ID_number*> **-x**.

Step 16 When the warning message about the disk format appears, enter **y** to continue the storages configuration.

```
*********** WARNING! **********
All storages are about to be reconfigured. All data on the data disks will
be lost.
Do you want to continue? (y/n)y
```

Note You may be prompted to enter the password several times.

When the script is finished running, you should see the word “COMPLETE”. After running this script, the new VDS TC cache engine will be connected to the storage enclosures.

Step 17 Use SSH to connect to each CE and enter the command **iscsiadm -m session** to check the connections:

```
ssh ce-<cache engine number>
iscsiadm -m session
```

Your output should be similar to the following:

```
ƒÞ 8 * sessions per storage, for example for 1 storage:
tcp: [17] 10.11.14.100:3260,1 iqn.1992-
01.com.Avago:2365.60080e50002f4e82000000005063c74a
tcp: [18] 10.11.14.101:3260,2 iqn.1992-
01.com.Avago:2365.60080e50002f4e82000000005063c74a
```

```
tcp: [19] 10.11.16.100:3260,1 iqn.1992-
01.com.Avago:2365.60080e50002f4e82000000005063c74a
tcp: [20] 10.11.16.101:3260,2 iqn.1992-
01.com.Avago:2365.60080e50002f4e82000000005063c74a
tcp: [21] 10.11.15.100:3260,1 iqn.1992-
01.com.Avago:2365.60080e50002f4e82000000005063c74a
tcp: [22] 10.11.15.101:3260,2 iqn.1992-
01.com.Avago:2365.60080e50002f4e82000000005063c74a
tcp: [23] 10.11.17.101:3260,2 iqn.1992-
01.com.Avago:2365.60080e50002f4e82000000005063c74a
tcp: [24] 10.11.17.100:3260,1 iqn.1992-
01.com.Avago:2365.60080e50002f4e82000000005063c74a
```

Updating the NetApp Storage Enclosure

Note This process is only for the NetApp storage enclosure. If your system uses an IBM storage enclosure, refer to Updating the IBM Storage Enclosure section.

After you have configured the VDS TC cache engine to connect with the storage enclosure, you must update the host port identifiers on the NetApp storage enclosure for the new cache engine. Follow this procedure to perform this task:

Procedure

Step 1 From the SSH connection to the VDS TC management server, enter the following commands to start the VNC server:

a. **rm /tmp/.X1-lock** (This step is optional. It is only required if there is a leftover lock file from a previous vncserver instance.)

b. **rm /tmp/.X11-unix/X1** (This step is optional. It is only required if there is a leftover lock file from a previous vncserver instance.)

c. **vncserver :1**

Step 2 Use a VNC client to connect from your local computer to the VDS TC management server.

Figure 3-3 VNC Viewer

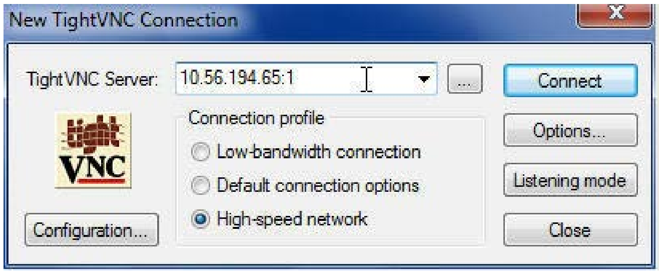

TightVNC Server: 10.56.194.65:1

● High-speed network

Step 3 From the VNC console window, enter the command **cd /opt/SMgr/client** to change folders.

Step 4 Enter the command **./SMclient** to start the SANtricity ™ ES Storage Manager software.

Note You must perform the following steps on all of the storage subsystems.

Step 5 Right-click the storage subsystem that you are configuring and choose **Manage Storage Array**. Click **No** in any pop-up windows that appear.

Figure 3-4 SANtricity Storage Manager

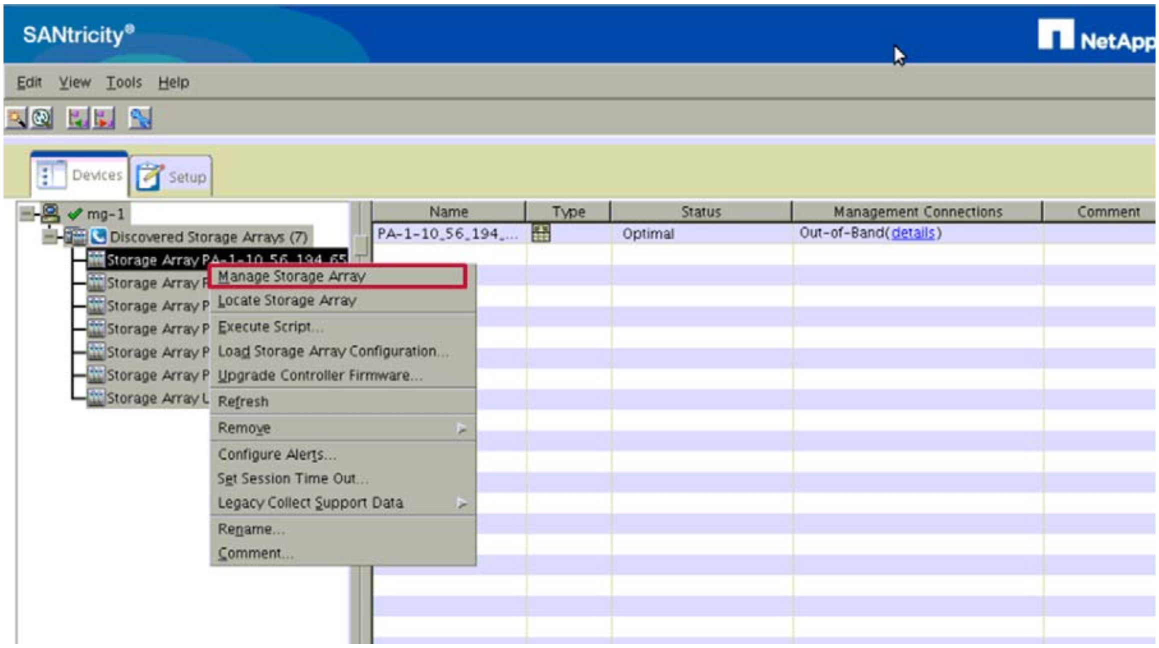

Name	Type	Status	Management Connections	Comment
PA-1-10.56.194....	▦	Optimal	Out-of-Band(details)	

Step 6 In the Subsystem Management window that appears, click the **Host Mappings** tab.

Step 7 In the navigation pane on the left, expand the **Host Group** folder.

Step 8 Right-click the cache engine that was replaced, for example Host CE-1, and choose **Manage Host Port Identifiers** from the pop-up menu.

Figure 3-5 Manage Host Port Identifiers

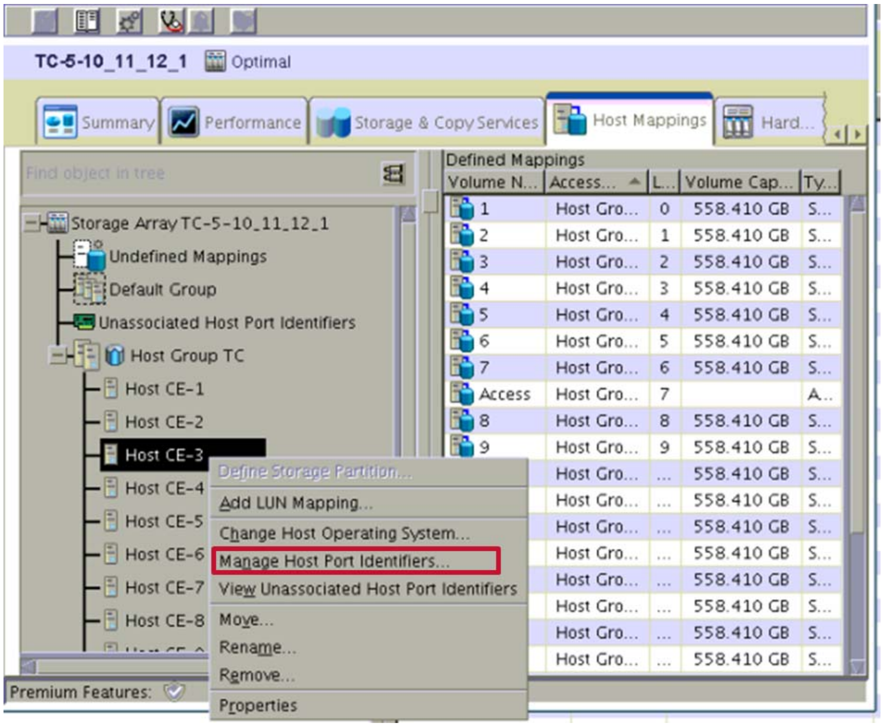

Volume N...	Access...	L...	Volume Cap...	Ty...
1	Host Gro...	0	558.410 GB	S...
2	Host Gro...	1	558.410 GB	S...
3	Host Gro...	2	558.410 GB	S...
4	Host Gro...	3	558.410 GB	S...
5	Host Gro...	4	558.410 GB	S...
6	Host Gro...	5	558.410 GB	S...
7	Host Gro...	6	558.410 GB	S...
Access	Host Gro...	7		A...
8	Host Gro...	8	558.410 GB	S...
9	Host Gro...	9	558.410 GB	S...
	Host Gro...	...	558.410 GB	S...
	Host Gro...	...	558.410 GB	S...
	Host Gro...	...	558.410 GB	S...
	Host Gro...	...	558.410 GB	S...
	Host Gro...	...	558.410 GB	S...
	Host Gro...	...	558.410 GB	S...
	Host Gro...	...	558.410 GB	S...
	Host Gro...	...	558.410 GB	S...

Step 9 From the Manage Host Port Identifiers window that appears, in the Host Port Identifier Information table, ensure that the old iqn is selected and click **Replace**.

Figure 3-6 Replace IQN

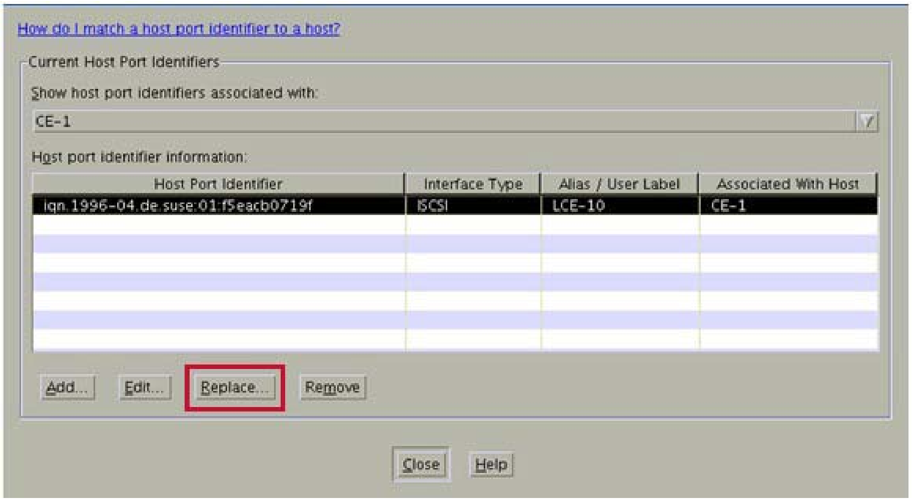
How do I match a host port identifier to a host?
Current Host Port Identifiers
Add...
Edit...
Replace...
Remove
Close
Help

Show host port identifiers associated with:

CE-1

Host port identifier information:

Host Port Identifier	Interface Type	Alias / User Label	Associated With Host
iqn.1996-04.de.suse:01:f5eacb0719f	iSCSI	LCE-10	CE-1

Step 10 From the Replace Host Port Identifier window that appears, configure the following:

- From the Choose a Host Interface Type drop-down list, choose **ISCSI**.
- Click the **Replace by Selecting a Known Unassociated Host Port Identifier** radio button.

- From the Known Unassociated Host Port Identifier drop-down list, choose the iqn for the *newly* installed cache engine. (It should be the only one in the list.)

Step 11 Click **Replace**.

Figure 3-7 Replace Host Port Identifier

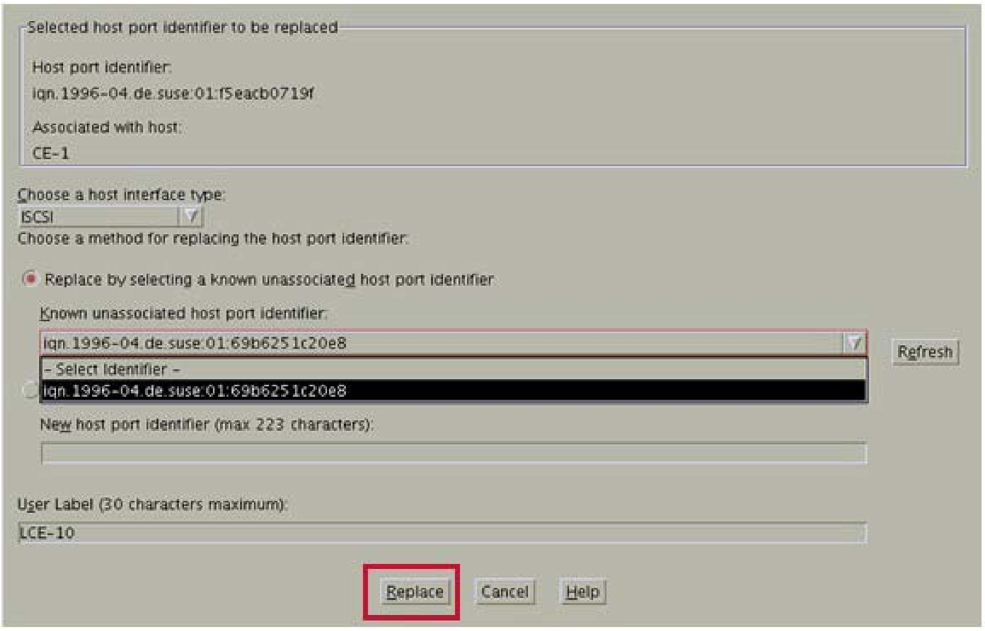

Selected host port identifier to be replaced

Host port identifier:
iqn.1996-04.de.suse:01:f5eacb0719f

Associated with host:
CE-1

Choose a host interface type:
iSCSI

Choose a method for replacing the host port identifier:

● Replace by selecting a known unassociated host port identifier

Known unassociated host port identifier:
iqn.1996-04.de.suse:01:69b6251c20e8

- Select Identifier -
iqn.1996-04.de.suse:01:69b6251c20e8

○ New host port identifier (max 223 characters):

User Label (30 characters maximum):
LCE-10

Step 12 From the Manage Host Port Identifiers window, click **Close**.

Step 13 Repeat Step 5 through Step 12 on any remaining storage subsystems.

Step 14 Close the Subsystem Management window and exit the SANtricity Storage Manager.

Step 15 Close the VNC client.

Updating the IBM Storage Enclosure

Note This process is only for the IBM storage enclosure. If your system uses the NetApp E2724 storage enclosure, refer to Updating the NetApp Storage Enclosure section.

After you have configured the VDS TC cache engine to connect with the storage enclosure, you must update the host port identifiers on the IBM storage enclosure for the new cache engine. Follow this procedure to perform this task:

Procedure

Step 1 From the SSH connection to the VDS TC management server, enter the following commands to start the VNC server:

a. **rm /tmp/.X1-lock** (This step is optional. It is only required if there is a leftover lock file from a previous vncserver instance.)

b. **rm /tmp/.X11-unix/X1** (This step is optional. It is only required if there is a leftover lock file from a previous vncserver instance.)

c. **vncserver :1**

Step 2 Use a VNC client to connect from your local computer to the VDS TC management server.

Figure 3-8 VNC Viewer

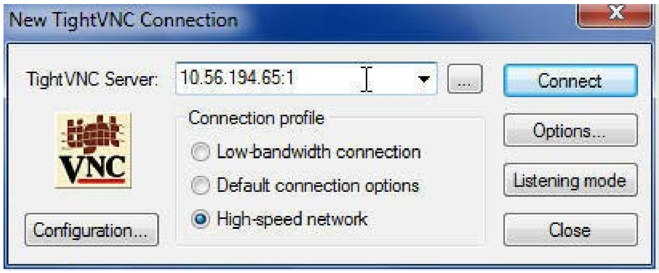

TightVNC Server: 10.56.194.65:1

● High-speed network

Step 3 From the VNC console window, enter the command **/opt/IBM_DS/client/SMclient** to start the IBM Storage Manager.

Note You must perform the following steps on all of the storage subsystems.

Step 4 Right-click the storage subsystem that you are configuring and choose **Manage Storage Subsystem**. Click **No** in any pop-up windows that appear.

Figure 3-9 IBM DS Storage Manager

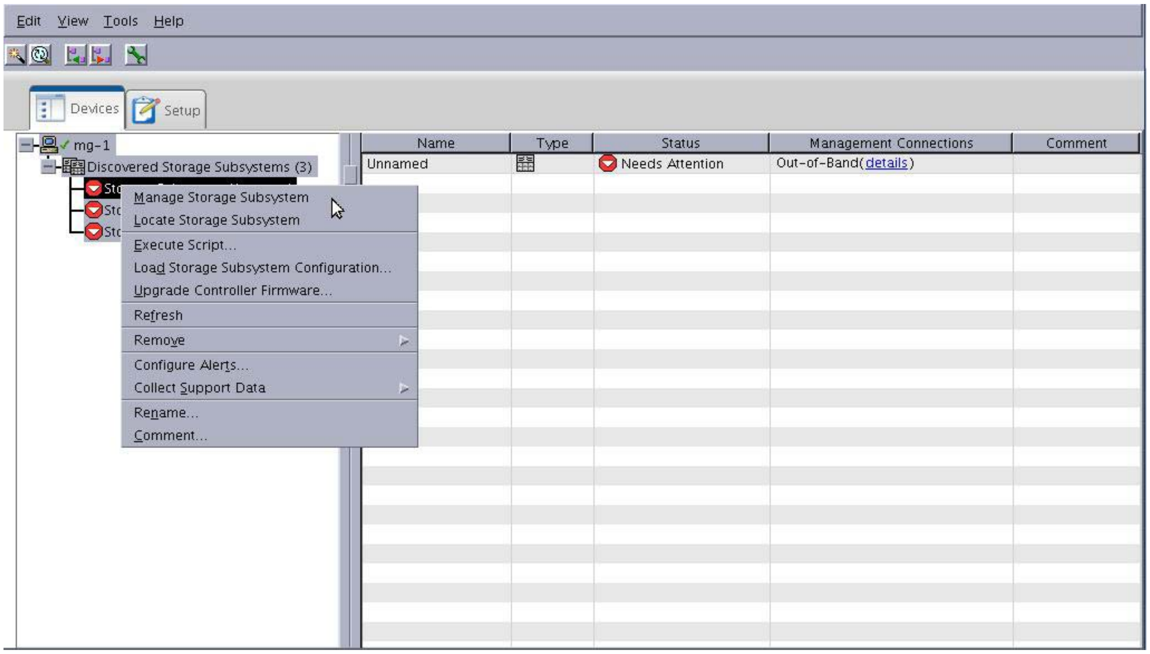

Name	Type	Status	Management Connections	Comment
Unnamed		Needs Attention	Out-of-Band(details)	

Step 5 In the Subsystem Management window that appears, click the **Host Mappings** tab.

Step 6 In the navigation pane on the left expand the **Host Group** folder.

Step 7 Right-click the cache engine that was replaced, for example Host CE-1, and choose **Manage Host Port Identifiers** from the pop-up menu.

Figure 3-10 Manage Host Port Identifiers

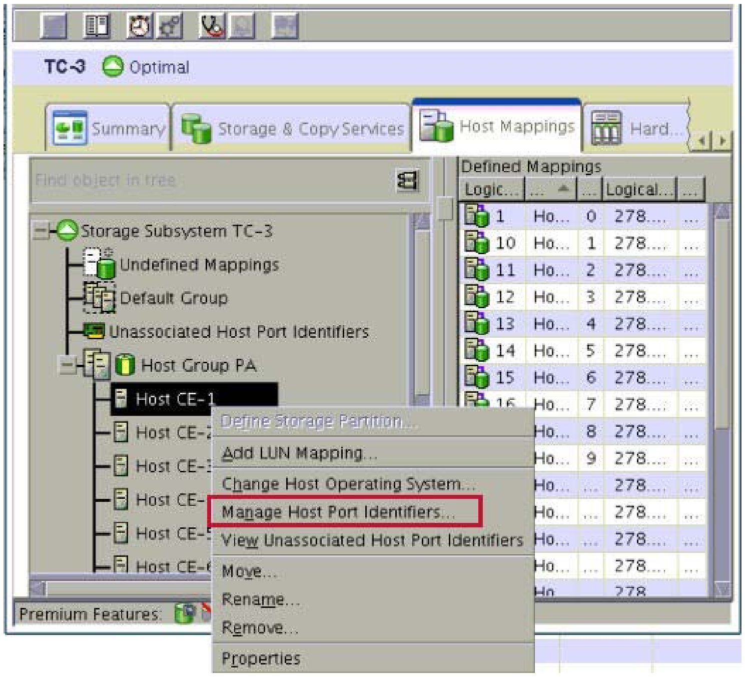

Logic...	...	...	Logical...	...
1	Ho...	0	278....	...
10	Ho...	1	278....	...
11	Ho...	2	278....	...
12	Ho...	3	278....	...
13	Ho...	4	278....	...
14	Ho...	5	278....	...
15	Ho...	6	278....	...
16	Ho...	7	278....	...
	Ho...	8	278....	...
	Ho...	9	278....	...
	Ho...	...	278....	...
	Ho...	...	278....	...
	Ho...	...	278....	...
	Ho...	...	278....	...
	Ho		278	

Step 8 From the Manage Host Port Identifiers window that appears, in the Host Port Identifier Information table, ensure that the old iqn is selected and click **Replace**.

Figure 3-11 Replace IQN

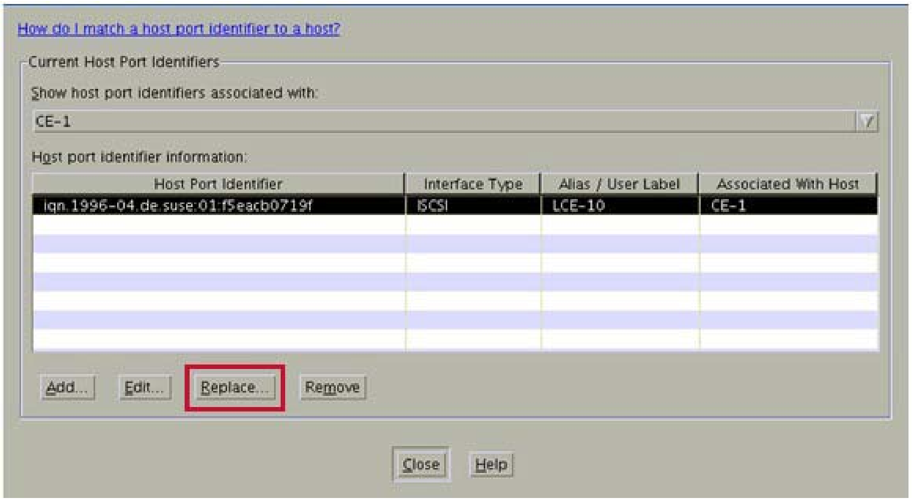
How do I match a host port identifier to a host?

Current Host Port Identifiers

Show host port identifiers associated with:

CE-1

Host port identifier information:

Host Port Identifier	Interface Type	Alias / User Label	Associated With Host
iqn.1996-04.de.suse:01:f5eacb0719f	ISCSI	LCE-10	CE-1

Step 9 From the Replace Host Port Identifier window that appears, configure the following:

- From the Choose a Host Interface Type drop-down list, choose **ISCSI**.
- Click the **Replace by Selecting a Known Unassociated Host Port Identifier** radio button.

- From the Known Unassociated Host Port Identifier drop-down list, choose the iqn for the *newly* installed cache engine. (It should be the only one in the list.)

Step 10 Click **Replace**.

Figure 3-12 Replace Host Port Identifier

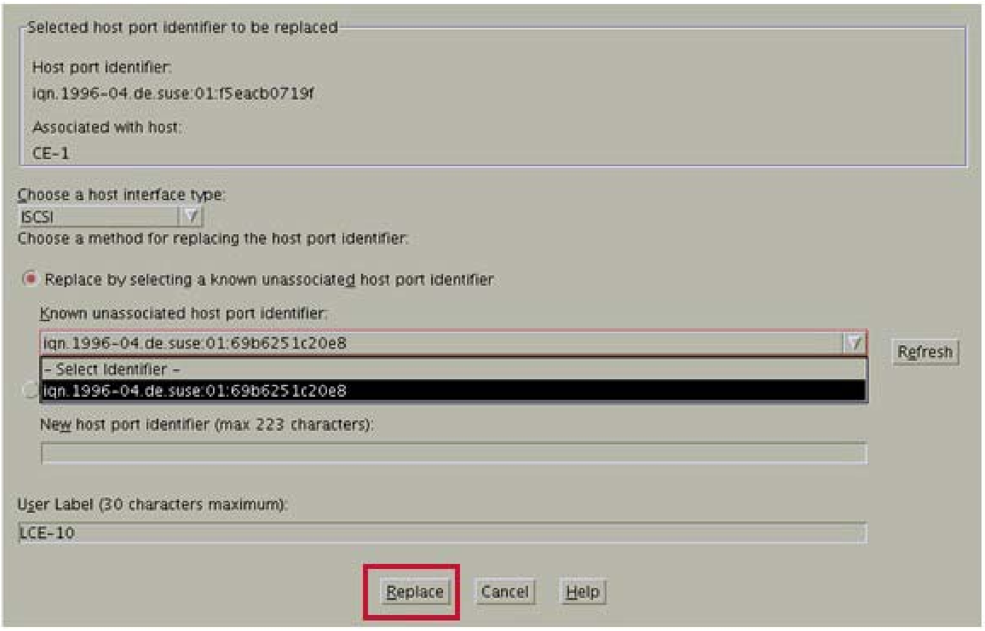

Selected host port identifier to be replaced

Host port identifier:
iqn.1996-04.de.suse:01:f5eacb0719f

Associated with host:
CE-1

Choose a host interface type:
ISCSI

Choose a method for replacing the host port identifier:

● Replace by selecting a known unassociated host port identifier

Known unassociated host port identifier:
iqn.1996-04.de.suse:01:69b6251c20e8
- Select Identifier -
iqn.1996-04.de.suse:01:69b6251c20e8

○ New host port identifier (max 223 characters):

User Label (30 characters maximum):
LCE-10

Step 11 From the Manage Host Port Identifiers window, click **Close**.

Step 12 Repeat Step 4 through Step 11 on any remaining storage subsystems.

Step 13 Close the Subsystem Management window and exit the IBM Storage Manager.

Step 14 Close the VNC client.

Finishing the Cache Engine Replacement

Follow this procedure to finish configuring and installing the new cache engine:

Note You will be re-running the storage script that you ran in the Updating the SSH Keys and Connecting the Cache Engine to the Storage Enclosure section.

Note You must be logged in with root permissions to perform the following steps.

Procedure

Step 1 From the SSH connection to the VDS TC management server, enter the command **cd /opt/pang/useful/configure_storages** to change to the storage configuration script directory.

Step 2 Enter the command **./configure_storages.py -s** <*no_of_storage_enclosures*> **-b** <*no_of_blades*> **-c** <*the_replaced_CE-ID_number*>. For example, to replace CE-2 in a VDS TC C-Series installation that has two cache engines, and one storage enclosure, enter the command **./configure_storages_py -s 1 -b 2 -c 2**.

Step 3 When the warning message about the disk format appears, enter **y** to continue the storages configuration.

```
*********** WARNING! **********
All storages are about to be reconfigured. All data on the data disks will
be lost.
Do you want to continue? (y/n)y
```

Step 4 Follow the on screen instructions. If prompted, enter the password for the root user that was provided by Cisco. You should see the word "COMPLETE" when the script is finished running.

Note You may be prompted to enter the password several times.

Step 5 Using SFTP software, such as WinSCP, connect to the management IP address that was assigned to the VDS TC server. Log in using the user name **padmin** and the password that was provided by Cisco. Copy the VDS TC 5.7.3 ISO image file, VDS-TC_Installer_5.7.3b53-5.7.3b54-ISO-7-Cisco.iso, to the **/opt/pang/iso** folder.

Step 6 Copy the new VDS TC license that you received from Cisco to the **/tftpboot** folder.

Step 7 Close the SFTP software.

Step 8 From the SSH connection to the VDS TC management server logged in as root, perform the following steps to extract the FI file for VDS TC 5.7.3:

a. Enter the command **mkdir /opt/pang/iso/5.7.3_iso** to create a new folder for the 5.7.3 files.

b. Enter the command **cd /opt/pang/iso** to change folders.

c. Enter the command **mount VDS-TC_Installer_5.7.3b53-5.7.3b54-ISO-7-Cisco.iso -o loop /opt/pang/iso/5.7.3_iso** to extract the VDS-TC_FI_5.7.3b54_Server_Cluster.tar.gz file.

d. Enter the command **cd /** to change to the root folder and then enter the command **umount /opt/pang/iso/5.7.3_iso** to unmount the iso image.

e. Enter the command **rm /opt/pang/iso/VDS-TC_Installer_5.7.3b53-5.7.3b54-ISO-7-Cisco.iso** to delete the iso file.

f. Enter the command **cd /opt/pang/iso/5.7.3_iso/install-tools/PeerApp_FI-GA**.

g. Enter the command **cp VDS-TC_FI_5.7.3b54_Server_Cluster.tar.gz /tftpboot**.

Step 9 Enter the command **su admin** to log into the VDS TC CLI. The CLI prompt console> appears.

Step 10 Enter the command **enable** to enter Enable mode. At the password prompt that appears, enter the enable password. The CLI prompt console# appears.

Step 11 Upgrade the software on the new cache engine by entering the command **upgrade server** *CE_number* **127.0.0.1 VDS-TC_FI_5.7.3b54_Server_Cluster.tar.gz**, where *CE_number* is the cache engine you replaced.

Step 12 To import the new license that includes the serial number of the new cache engine, enter the command **license import 127.0.0.1** *filename*, where *filename* is the name of the new license file.

Note You must be in Enable mode, *not* Configuration mode to import a new license.

Step 13 Enter the command **license activate** to apply and activate the license.

Step 14 To start the application on the replaced cache engine, enter the command **oper server** *#*, where # is the number of the cache engine that was replaced. For example, if you replaced CE-2, enter the command **oper server 2**. You should see the following output:

```
console# oper server 2
oper server 2# start
Starting server 2
service is started on server 2
```

Step 15 After the application has started on the new cache engine, enter the command **oper service start** to start the VDS TC service.

Step 16 Enter the command **show cluster-bus-ip** to check if the spread communication among cluster members is set to broadcast.

```
console# show cluster-bus-ip
Processing...
MG-1: [Broadcast]
CE-1: [Broadcast]
CE-2: [Broadcast]
CE-3: [Broadcast]
CE-4: [Broadcast]
CE-5: [Broadcast]
CE-6: [Broadcast]
CE-7: [Broadcast]
CE-8: [Broadcast]
OK. All the CEs are configured to work with Broadcast IP.
```

Step 17 If the communication is set to broadcast, continue on to Step 20. If the communication is set to *multicast* instead of *broadcast*, enter the command **config** to switch to Configuration mode and then enter the command **cluster-bus-ip broadcast**. For example:

```
configuration# cluster-bus-ip broadcast
Processing...
All the CEs are configured to work with Broadcast IP.
```

Step 18 Enter the command **exit** to exit the Configuration mode.

Step 19 Enter the command **show cluster-bus-ip** to confirm that the management server and all cache engines are configured to use broadcast. For example:

```
console# show cluster-bus-ip
Processing...
MG-1: [Broadcast]
CE-1: [Broadcast]
CE-2: [Broadcast]
CE-3: [Broadcast]
CE-4: [Broadcast]
CE-5: [Broadcast]
CE-6: [Broadcast]
CE-7: [Broadcast]
CE-8: [Broadcast]
OK. All the CEs are configured to work with Broadcast IP.
```

Step 20 Wait a few minutes and then enter the command **show status**. Do not proceed to the next step until you see a Device Status of "Started." For example:

```
console# show status
Operational state Device state Administrative state
enabled started unlocked
```

Step 21 On the PBR routers, start redirecting traffic to the VDS TC solution.

CHAPTER 4

Cache Engine Replacement Procedure for a Blade Server Installation with a Fresh Install of VDS TC 5.7.3

Note Only use this chapter if you are replacing a Cache Engine in a VDS TC Blade Server system that had a fresh install of VDS TC 5.7.3 and was *not* upgraded to 5.7.3. If you are replacing a Cache Engine in a VDS TC Blade Server system that was upgraded to VDS TC 5.7.3 from VDS TC 5.2.1, refer to Chapter 5, “Cache Engine Replacement Procedure for a 5.7.3 Blade Server Installation Upgraded from Release 5.2.1”. If you are replacing a Cache Engine in a VDS TC Blade Server system that was upgraded to VDS TC 5.7.3 from VDS TC 5.6.1, refer to Chapter 6, “Cache Engine Replacement Procedure for a 5.7.3 Blade Server Installation Upgraded from Release 5.6.1”.

This chapter discusses how to determine that a cache engine has failed in a VDS TC Blade Server installation and the steps that are required to replace it.

When a cache engine fails, the volume_selection_algorithm_time setting in the cluster_conf.xml file controls when the volumes of a failed cache engine will be redistributed (mounted) to the remaining cache engines. By default this value is 24 hours, which means that by default, if a cache engine fails, its volumes will not be redistributed (mounted) to the remaining cache engines until 24 hours after the cache engine fails.

It is recommended that you leave the default set to 24 hours, however if you need to change this setting, add the following text to the cluster_conf file in the <service> section under the <memory> subsection. The *seconds* parameters is the number of seconds that the system should wait after a cache engine fails before its volumes are redistributed (mounted) to the remaining cache engines:

```
<io>
    <volume_selection_algorithm_time>seconds</volume_selection_algorithm_time>
</io>
```

For example, to have the system wait 30 minutes after a cache engine fails before redistributing the volumes of the failed cache engine enter the following:

```
<io>
    <volume_selection_algorithm_time>1800</volume_selection_algorithm_time>
</io>
```

For more information on how to make changes to the configuration file and how to apply these changes, please refer to the *Cisco Videoscape Distribution Suite Transparent Caching Software Configuration Guide* available at http://www.cisco.com/c/dam/en/us/td/docs/video/videoscape/distribution_suite/vds/v5_7_3/VDS-TC_5_7_3_software_config_guide.pdf.

Symptoms of a Failed VDS TC Blade Server Cache Engine

To determine whether a cache engine has failed in a VDS TC Blade Server installation, log into VDS TC Manager GUI. By default, the Status > Dashboard window is displayed. If you are already logged into VDS TC Manager, choose **Status** > **Dashboard**.

If a cache engine has failed it will appear with a gray box surrounding it.

Figure 4-1 Failed Cache Engine

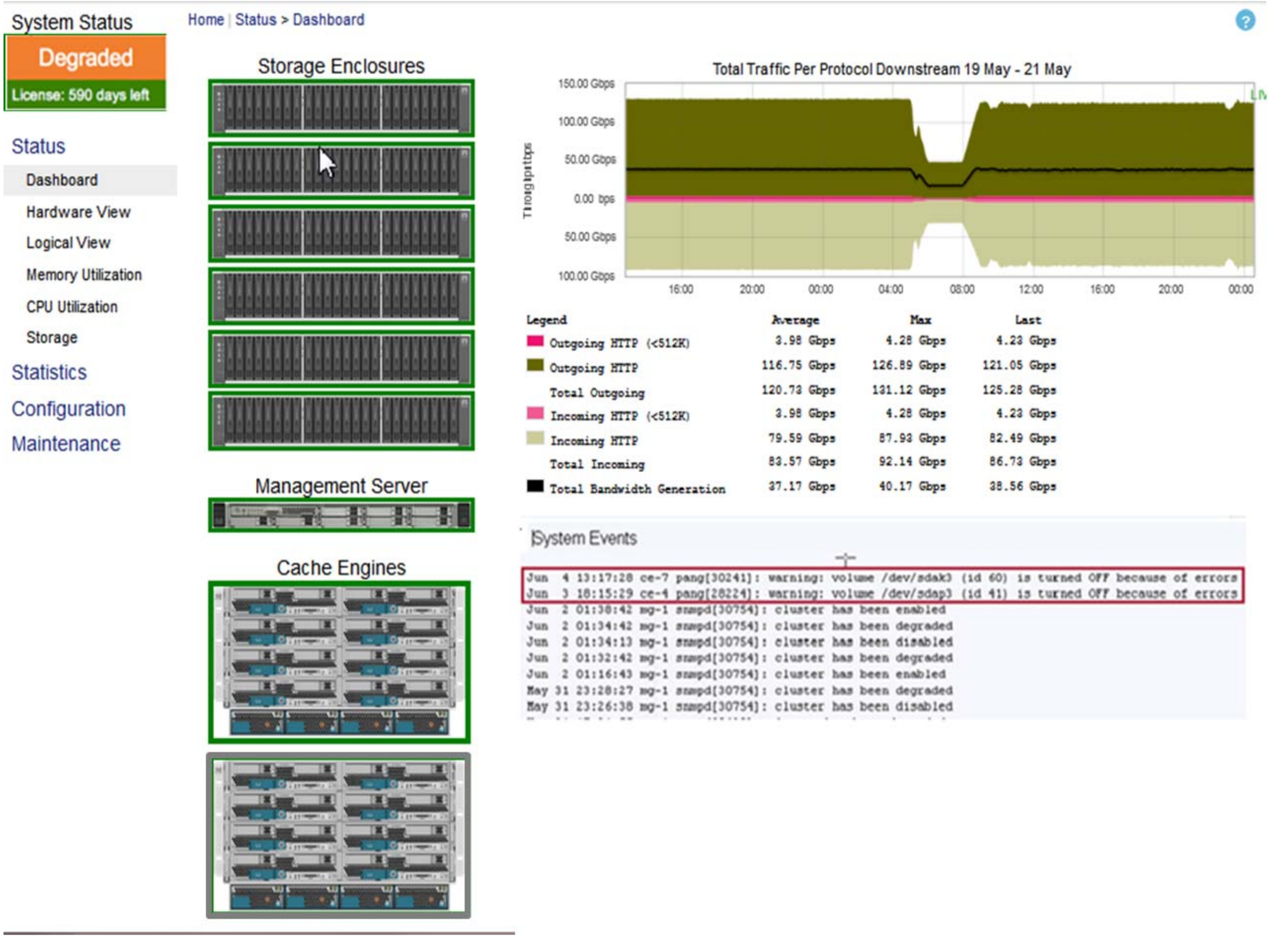

Legend	Average	Max	Last
🟥 Outgoing HTTP (<512K)	3.98 Gbps	4.28 Gbps	4.23 Gbps
🟫 Outgoing HTTP	116.75 Gbps	126.89 Gbps	121.05 Gbps
Total Outgoing	120.73 Gbps	131.12 Gbps	125.28 Gbps
🟥 Incoming HTTP (<512K)	3.98 Gbps	4.28 Gbps	4.23 Gbps
🟨 Incoming HTTP	79.59 Gbps	87.93 Gbps	82.49 Gbps
Total Incoming	83.57 Gbps	92.14 Gbps	86.73 Gbps
■ Total Bandwidth Generation	37.17 Gbps	40.17 Gbps	38.56 Gbps

System Events

```
Jun  4 13:17:28 ce-7 pang[30241]: warning: volume /dev/sdak3 (id 60) is turned OFF because of errors
Jun  3 18:15:29 ce-4 pang[28224]: warning: volume /dev/sdap3 (id 41) is turned OFF because of errors
Jun  2 01:38:42 mg-1 snmpd[30754]: cluster has been enabled
Jun  2 01:34:42 mg-1 snmpd[30754]: cluster has been degraded
Jun  2 01:34:13 mg-1 snmpd[30754]: cluster has been disabled
Jun  2 01:32:42 mg-1 snmpd[30754]: cluster has been degraded
Jun  2 01:16:43 mg-1 snmpd[30754]: cluster has been enabled
May 31 23:28:27 mg-1 snmpd[30754]: cluster has been degraded
May 31 23:26:38 mg-1 snmpd[30754]: cluster has been disabled
```

Replacing a Cache Engine and Running the Installation Tasks

The first step in replacing a cache engine in a VDS TC Blade Server cluster installation is to physically replace the cache engine and perform the installation steps on the new cache engine. Follow this procedure to perform these steps:

Before You Begin

To replace a failed cache engine you must request a new VDS TC license from Cisco. To request a new license, you must have the serial number of the failed cache engine and the serial number of the new cache engine. After you have gathered this information, contact Cisco support for a new license file.

Caution Do not continue until you have the license file.

Procedure

Note Before removing the failed cache engine, make note of what number corresponds with the cache engine.

Step 1 Using the Cisco UCS 5108 Server Chassis Installation Guide available at http://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/hw/chassis-install-guide/ucs5108_install.html, unmount the failed Blade Server cache engine and mount the new Blade Server cache engine.

Step 2 In the *Cisco Videoscape Distribution Suite Transparent Caching Software Blade Server Cluster Installation Guide*, available at http://www.cisco.com/c/dam/en/us/td/docs/video/videoscape/distribution_suite/vds/v5_7_3/VDS-TC_5_7_3_BladeSrvr_swInstallGuide.pdf, refer to Chapter 1, “Prerequisites and Blade Server Cluster Physical Installation” for connecting the network cables to the new cache engine.

Note Complete the following steps for the cache engine that you are replacing.

Step 1 Use VNC to connect from your local computer to the VDS TC Management Server using the IP address that you configured, adding :1 to the end, for example 10.56.194.65:1. When prompted to enter a password, enter the root password that was provided by Cisco. If you encounter problems opening a VNC client connection to the VDS TC Management Server, you may need to restart VNC Server 1 on the VDS TC Management Server. Follow these steps to restart this server:

a. From an SSH connection to the VDS TC Management server where you are logged in with root permissions, enter the command **rm /tmp/.X11-unix/X1**.

b. Enter the command **vncserver :1** to start the VNC Server.

Figure 4-2 VNC Viewer

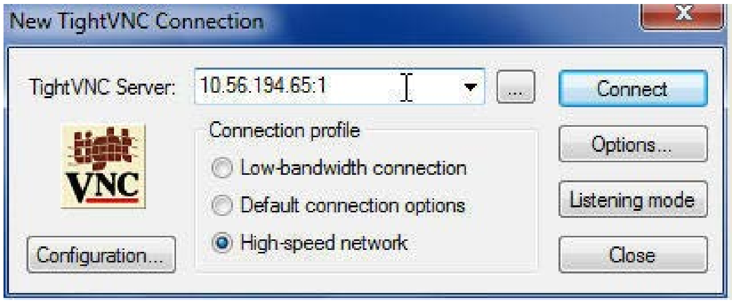

TightVNC Server: 10.56.194.65:1

Connection profile

○ Low-bandwidth connection
○ Default connection options
● High-speed network

Note If you have problems restarting the VNC Server 1 instance, you can also start another VNC Server instance by entering the command **vncserver :2**.

Step 2 From the VNC console, enter the command **firefox &**.

Figure 4-3 VNC Console

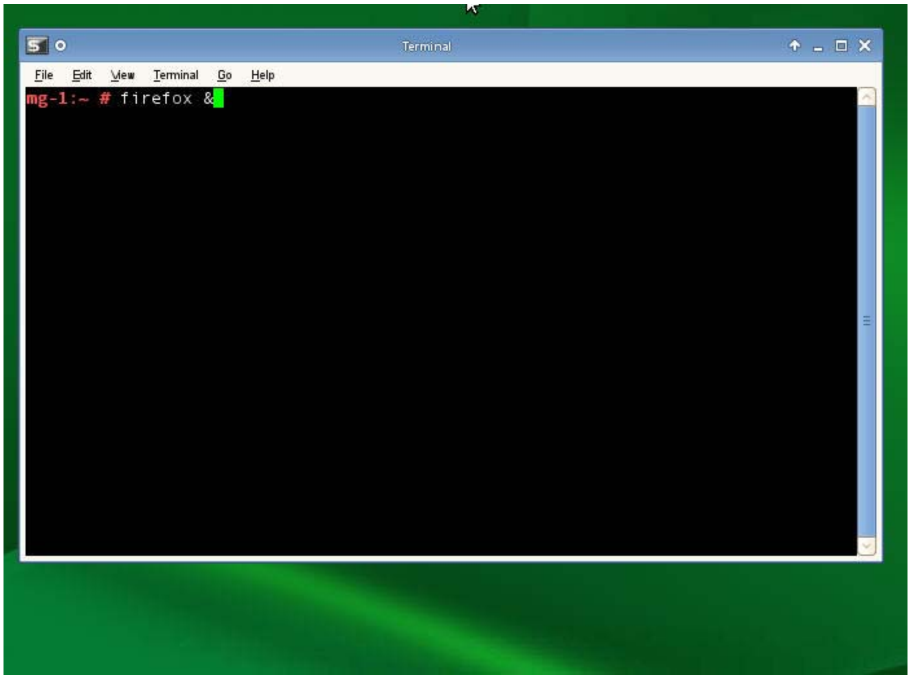

```
mg-1:~ # firefox &
```

Step 3 In the Firefox window that opens, enter the Management port IP address for the fabric interconnect. In a redundant installation, enter the Cluster (virtual) IP address.

Step 4 The Cisco UCS Manager window appears. From this window, click **Launch KVM Manager**. If you receive a message that says "To view this page ensure that Adobe Flash Player version 10.0.0 or greater is installed." perform the following steps:

a. Close Firefox

b. From the terminal window, enter the command **su –**

c. Enter the command **pkill -9 firefox**

d. Enter the command **cd /root**

e. Enter the command **mv .mozilla .mozilla.bak**

f. Enter the command **firefox &** to restart Firefox.

g. In the Firefox window that opens, enter the Management port IP address for the fabric interconnect. In a redundant installation, enter the Cluster (virtual) IP address.

h. Repeat Step 4.

Figure 4-4 Cisco UCS Manager

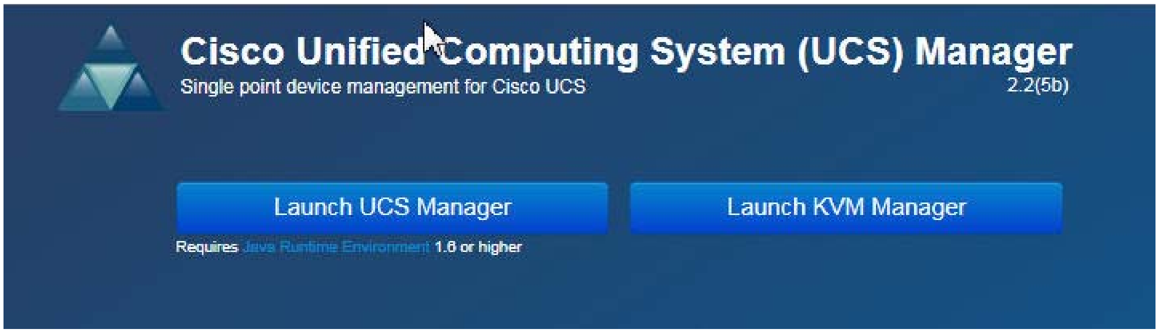
Cisco Unified Computing System (UCS) Manager
Single point device management for Cisco UCS
2.2(5b)
Launch UCS Manager
Launch KVM Manager
Requires Java Runtime Environment 1.6 or higher

Step 5 The UCS - KVM Launch Manager login window appears. Enter the username **admin** and the password that was configured during the initial system setup of the fabric interconnect (the default password is **password**), and then click **OK**.

Figure 4-5 UCS KVM Launch Manager Login Window

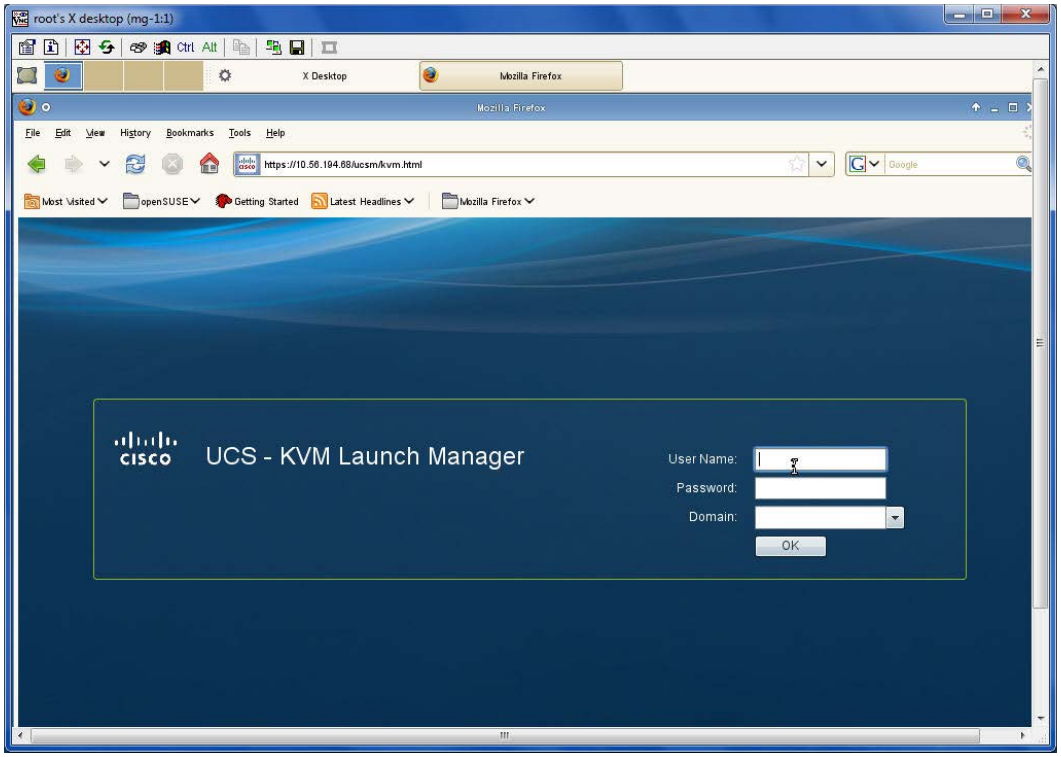

Step 6 The Service Profiles window appears, listing all of the cache engine profiles. From this window, in the row for the cache engine you are replacing, in the Launch KVM column click **Launch** to launch the KVM console.

Figure 4-6 Service Profiles Window

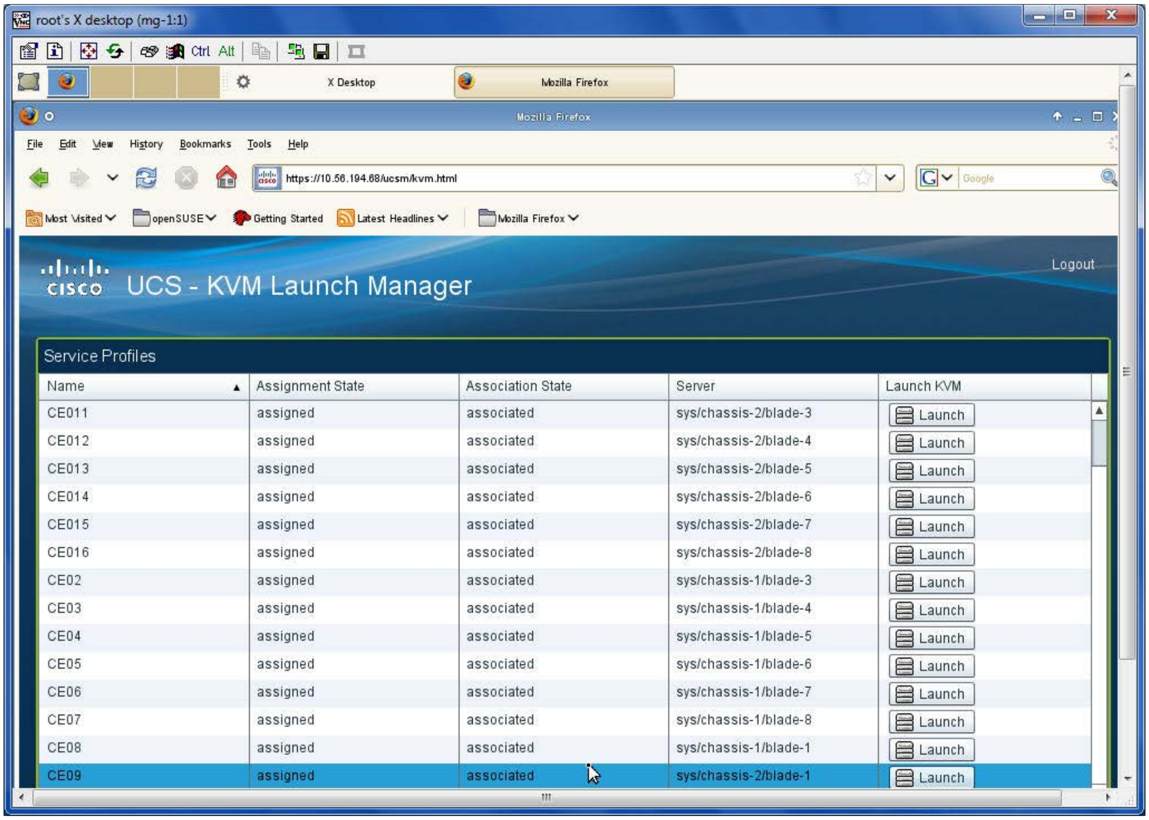
root's X desktop (mg-1:1)
Ctrl Alt
X Desktop
Mozilla Firefox
Mozilla Firefox
File Edit View History Bookmarks Tools Help
https://10.56.194.68/ucsm/kvm.html
Google
Most Visited openSUSE Getting Started Latest Headlines Mozilla Firefox
Logout

UCS - KVM Launch Manager

Service Profiles

Name	Assignment State	Association State	Server	Launch KVM
CE011	assigned	associated	sys/chassis-2/blade-3	Launch
CE012	assigned	associated	sys/chassis-2/blade-4	Launch
CE013	assigned	associated	sys/chassis-2/blade-5	Launch
CE014	assigned	associated	sys/chassis-2/blade-6	Launch
CE015	assigned	associated	sys/chassis-2/blade-7	Launch
CE016	assigned	associated	sys/chassis-2/blade-8	Launch
CE02	assigned	associated	sys/chassis-1/blade-3	Launch
CE03	assigned	associated	sys/chassis-1/blade-4	Launch
CE04	assigned	associated	sys/chassis-1/blade-5	Launch
CE05	assigned	associated	sys/chassis-1/blade-6	Launch
CE06	assigned	associated	sys/chassis-1/blade-7	Launch
CE07	assigned	associated	sys/chassis-1/blade-8	Launch
CE08	assigned	associated	sys/chassis-1/blade-1	Launch
CE09	assigned	associated	sys/chassis-2/blade-1	Launch

Step 7 Click **OK** to open the .jnlp file with Java(TM) Web Start Launcher.

Step 8 From the Warning Security window, click **Continue**.

Step 9 If you receive a window that asks if you want to run this application, click **Run**.

Step 10 From the KVM Console window, make sure the **KVM Console** tab is selected and choose **Virtual Media > Activate Virtual Devices**.

Step 11 If you see the Unencrypted Virtual Media Session window, choose **Accept this session** and click **Apply**.

Figure 4-7 Unencrypted Virtual Media Session

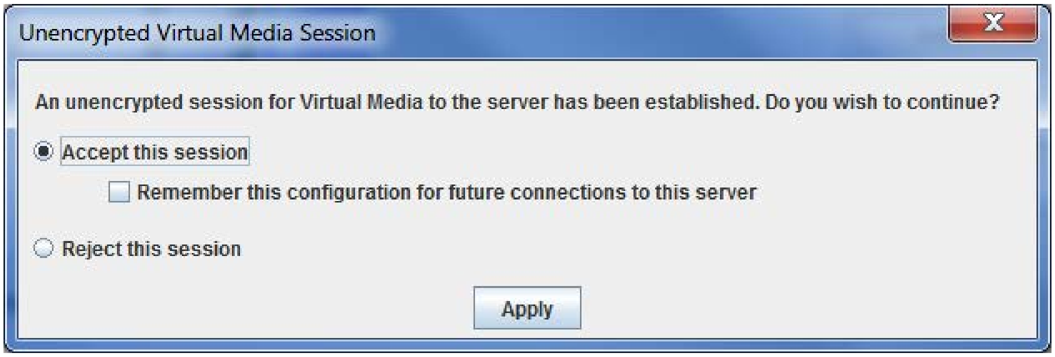
Unencrypted Virtual Media Session

An unencrypted session for Virtual Media to the server has been established. Do you wish to continue?

● Accept this session

□ Remember this configuration for future connections to this server

○ Reject this session

Step 12 Choose **Virtual Media > Map CD/DVD**.

Figure 4-8 Map CD/DVD

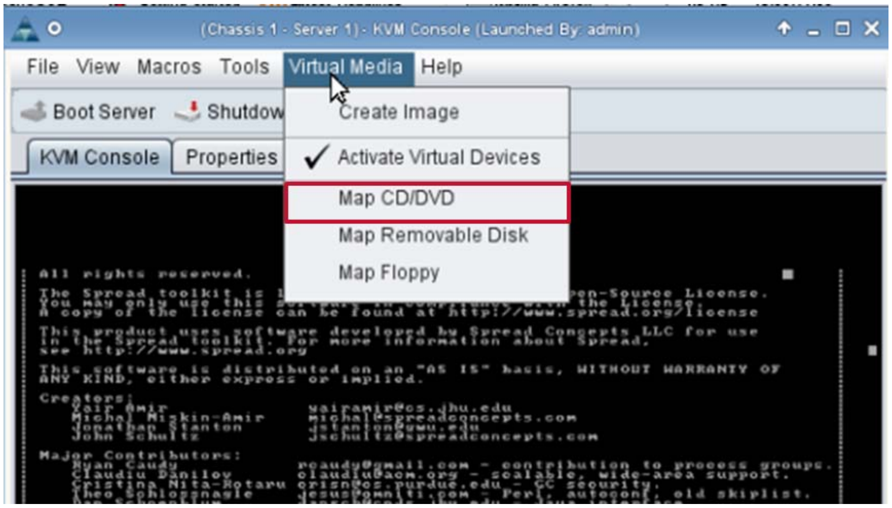

```
All rights reserved.
The Spread toolkit is l                              pen-Source License.
You may only use this                                 the License.
A copy of the license can be found at http://www.spread.org/license
This product uses software developed by Spread Concepts LLC for use
in the Spread toolkit. For more information about Spread,
see http://www.spread.org
This software is distributed on an "AS IS" basis, WITHOUT WARRANTY OF
ANY KIND, either express or implied.
Creators:
    Yair Amir              yairamir@cs.jhu.edu
    Michal Miskin-Amir     michal@spreadconcepts.com
    Jonathan Stanton       jstanton@gwu.edu
    John Schultz           jschultz@spreadconcepts.com
Major Contributors:
    Ryan Caudy            rcaudy@gmail.com - contribution to process groups.
    Claudiu Danilov       claudiu@acm.org - scalable, wide-area support.
    Cristina Nita-Rotaru  crisn@cs.purdue.edu - GC security.
    Theo Schlossnagle     jesus@omniti.com - Perl, autoconf, old skiplist.
```

Step 13 In the window that opens, browse to the **/opt/pang/iso/** folder to find the ISO installer image and click **Open**.

Figure 4-9 Map Device

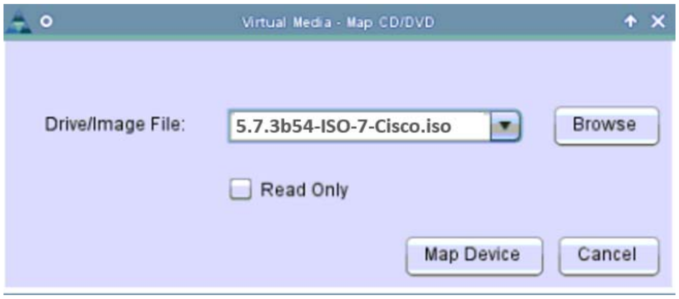

Drive/Image File: 5.7.3b54-ISO-7-Cisco.iso

□ Read Only

Step 14 Check the **Read Only** check box and click **Map Device**.

Step 15 Click **Reset** under the menu bar to reboot the cache engine.

Figure 4-10 Reset the Cache Engine

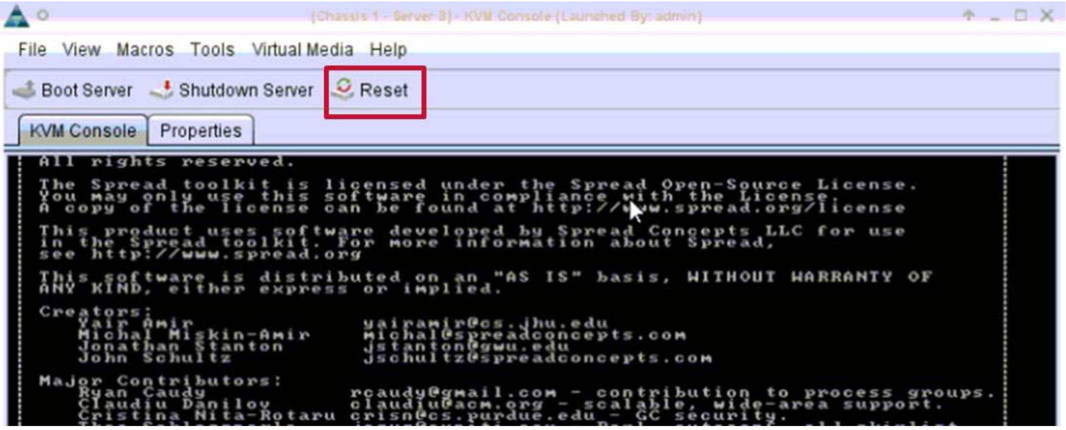
(Chassis 1 - Server 3) - KVM Console (Launched By: admin)
File View Macros Tools Virtual Media Help
Boot Server
Shutdown Server
Reset
KVM Console
Properties

```
All rights reserved.
The Spread toolkit is licensed under the Spread Open-Source License.
You may only use this software in compliance with the License.
A copy of the license can be found at http://www.spread.org/license
This product uses software developed by Spread Concepts LLC for use
in the Spread toolkit. For more information about Spread,
see http://www.spread.org
This software is distributed on an "AS IS" basis, WITHOUT WARRANTY OF
ANY KIND, either express or implied.
Creators:
    Yair Amir               yairamir@cs.jhu.edu
    Michal Miskin-Amir      michal@spreadconcepts.com
    Jonathan Stanton        jstanton@gwu.edu
    John Schultz            jschultz@spreadconcepts.com
Major Contributors:
    Ryan Caudy             rcaudy@gmail.com - contribution to process groups.
    Claudiu Danilov        claudiu@acm.org - scalable, wide-area support.
    Cristina Nita-Rotaru crisn@cs.purdue.edu - GC security.
```

Step 16 From the Reset Server Warning window that appears, click **OK** to reset the server.

Figure 4-11 Reset Server Warning Window

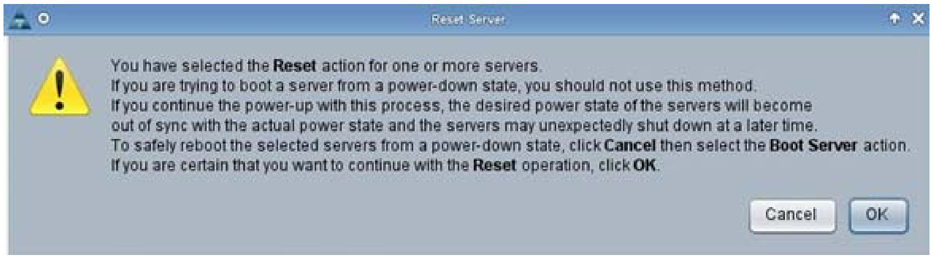

You have selected the **Reset** action for one or more servers.
If you are trying to boot a server from a power-down state, you should not use this method.
If you continue the power-up with this process, the desired power state of the servers will become
out of sync with the actual power state and the servers may unexpectedly shut down at a later time.
To safely reboot the selected servers from a power-down state, click **Cancel** then select the **Boot Server** action.
If you are certain that you want to continue with the **Reset** operation, click **OK**.

Step 17 From the Reset Server Service window that appears, choose **Power Cycle** and click **OK**.

Figure 4-12 Reset Server Service

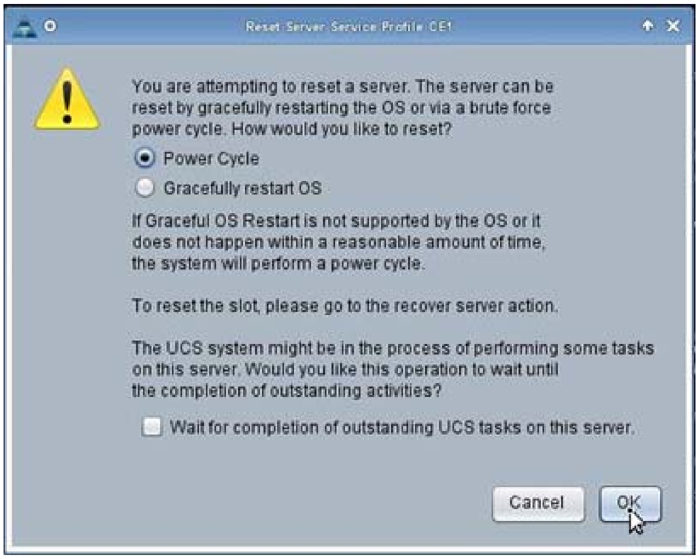

You are attempting to reset a server. The server can be reset by gracefully restarting the OS or via a brute force power cycle. How would you like to reset?

● Power Cycle

○ Gracefully restart OS

If Graceful OS Restart is not supported by the OS or it does not happen within a reasonable amount of time, the system will perform a power cycle.

To reset the slot, please go to the recover server action.

The UCS system might be in the process of performing some tasks on this server. Would you like this operation to wait until the completion of outstanding activities?

□ Wait for completion of outstanding UCS tasks on this server.

Step 18 From the KVM Console window you will see the VDS TC Cache Engine reboot. When you see the prompt to press F6 to enter the boot menu, click the KVM Console window to make sure it has focus and press **F6**. The Please Select Boot Device dialog box appears.

Note You may need to press **F6** several times to see the Please Select Boot Device dialog box.

Figure 4-13 System Reboot Window

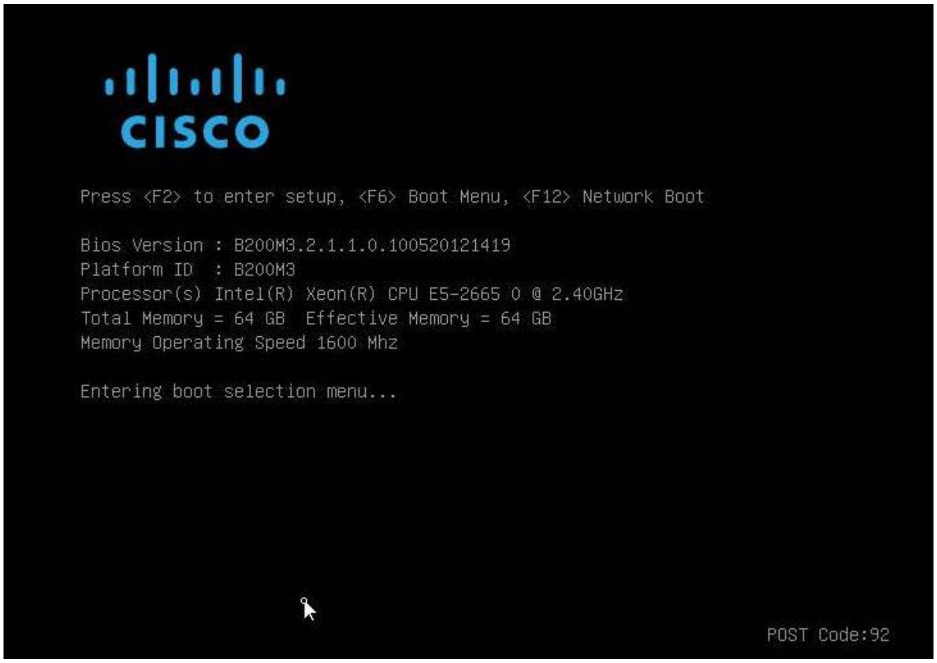

```
Press <F2> to enter setup, <F6> Boot Menu, <F12> Network Boot

Bios Version : B200M3.2.1.1.0.100520121419
Platform ID  : B200M3
Processor(s) Intel(R) Xeon(R) CPU E5-2665 0 @ 2.40GHz
Total Memory = 64 GB  Effective Memory = 64 GB
Memory Operating Speed 1600 Mhz

Entering boot selection menu...

                                                  POST Code:92
```

Step 19 Using the up and down arrows, choose **Cisco vKVM-Mapped vDVD1.22**. The system will reboot from the virtual ISO image.

Step 20 When the VDS-TC Installer environment displays in the KVM console window, choose **Installer - No Serial**.

Figure 4-14 VDS TC Installer Window

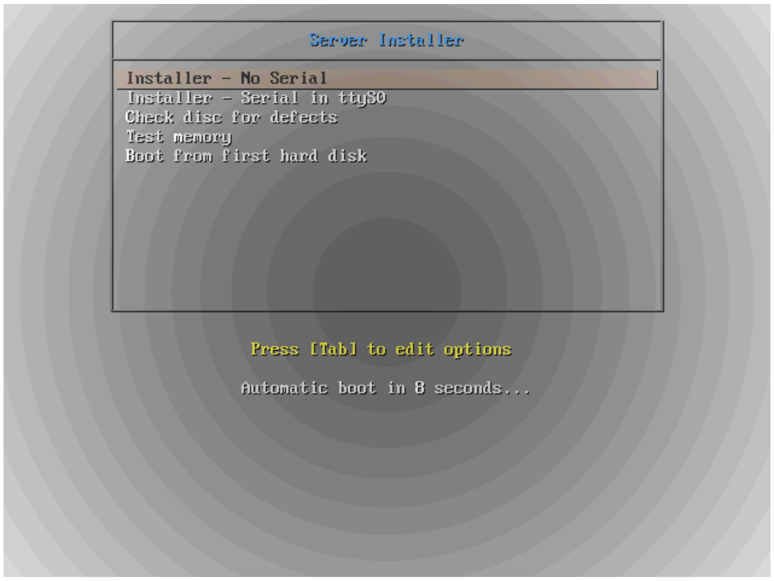

```
Server Installer

Installer - No Serial
Installer - Serial in ttyS0
Check disc for defects
Test memory
Boot from first hard disk

Press [Tab] to edit options

Automatic boot in 8 seconds...
```

Step 21 When the Welcome screen appears, enter **Y** to continue the installation process.

Figure 4-15 Welcome Screen

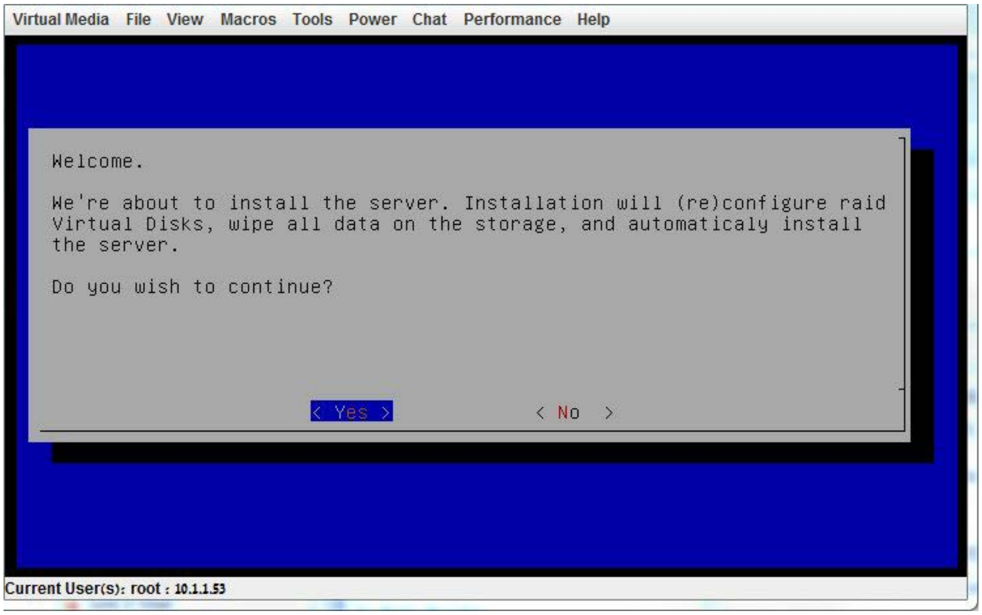

Virtual Media File View Macros Tools Power Chat Performance Help

```
Welcome.

We're about to install the server. Installation will (re)configure raid
Virtual Disks, wipe all data on the storage, and automaticaly install
the server.

Do you wish to continue?

< Yes >                < No  >
```

Current User(s): root : 10.1.1.53

Step 22 When the Choose the Appliance Deployment window appears, choose **Cache Engine** and select **OK**.

Figure 4-16 Choose the Appliance Deployment Window

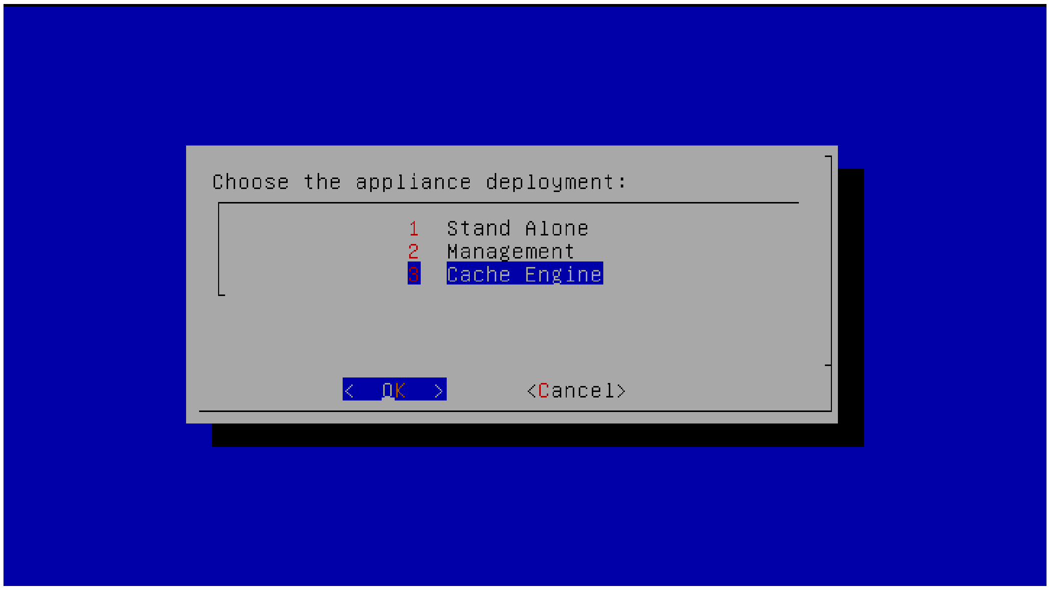

Choose the appliance deployment:

```
1  Stand Alone
2  Management
3  Cache Engine
```

Note When the "Enter Cache Engine Number" dialog box appears, enter the number that corresponds with the cache engine that you are replacing.

Figure 4-17 Enter Cache Engine Number

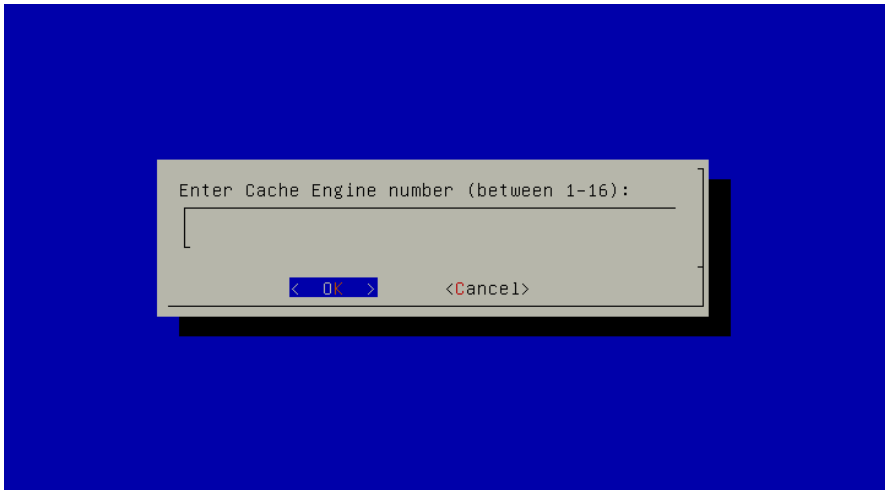

Enter Cache Engine number (between 1-16):

Step 23 From the Active-Active Fabric Interconnect dialog box, choose **yes**.

Figure 4-18 Active-Active Fabric Interconnect

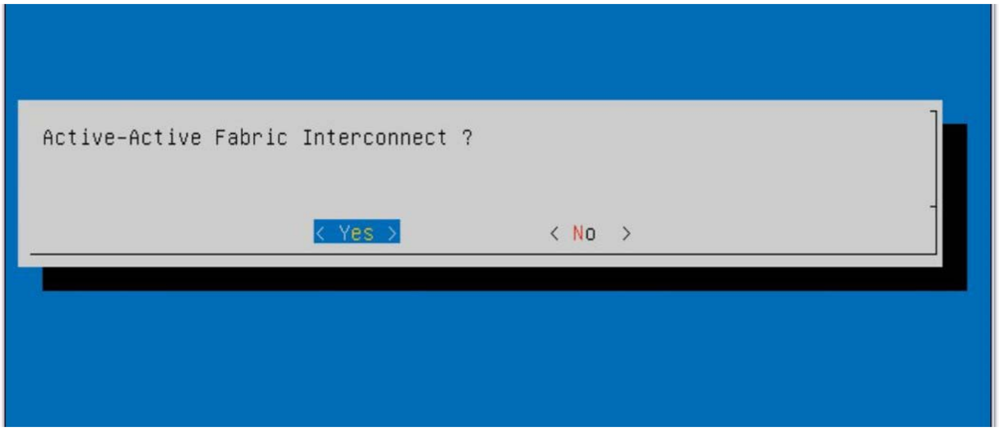

```
Active-Active Fabric Interconnect ?
```

Step 24 Confirm the Cache Engine configuration. If everything is correct, choose **Yes**. If there is a mistake, choose **No** and repeat Step 20 through Step 23.

Figure 4-19 Configuration Confirmation

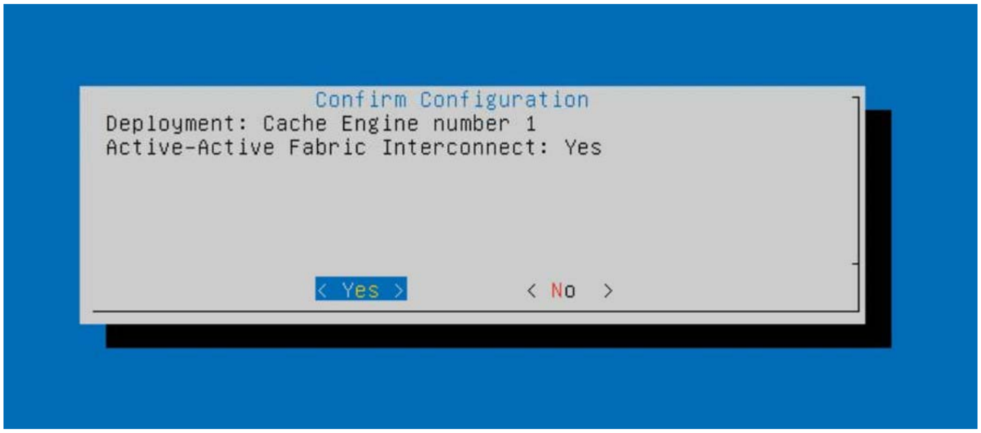

```
Confirm Configuration
Deployment: Cache Engine number 1
Active-Active Fabric Interconnect: Yes
```

Step 25 From the KVM window you will see the installation begin. You must wait for this installation process to complete before you continue.

Note This installation should take about 40 minutes.

Figure 4-20 Cache Engine Installation Process

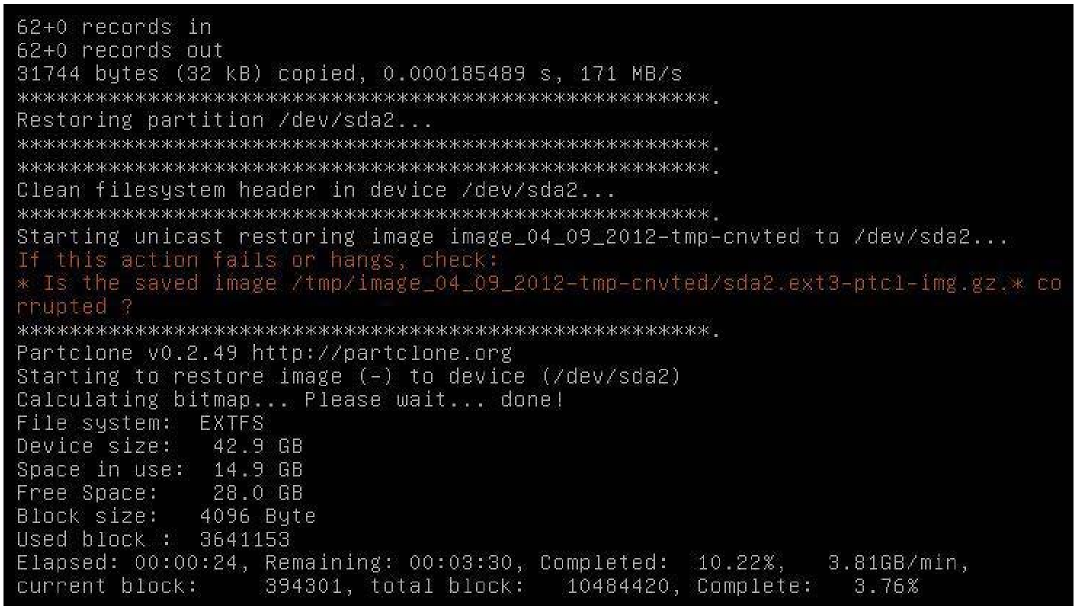

```
62+0 records in
62+0 records out
31744 bytes (32 kB) copied, 0.000185489 s, 171 MB/s
*****************************************************.
Restoring partition /dev/sda2...
*****************************************************.
*****************************************************.
Clean filesystem header in device /dev/sda2...
*****************************************************.
Starting unicast restoring image image_04_09_2012-tmp-cnvted to /dev/sda2...
If this action fails or hangs, check:
* Is the saved image /tmp/image_04_09_2012-tmp-cnvted/sda2.ext3-ptcl-img.gz.* co
rrupted ?
*****************************************************.
Partclone v0.2.49 http://partclone.org
Starting to restore image (-) to device (/dev/sda2)
Calculating bitmap... Please wait... done!
File system:  EXTFS
Device size:   42.9 GB
Space in use:  14.9 GB
Free Space:    28.0 GB
Block size:   4096 Byte
Used block :  3641153
Elapsed: 00:00:24, Remaining: 00:03:30, Completed:  10.22%,   3.81GB/min,
current block:     394301, total block:   10484420, Complete:   3.76%
```

Step 26 When the installation completes, you must unmap the ISO image file. When the “Installation completed. Press <Enter> to reboot.” message appears, press **Enter**.

Figure 4-21 Installation Completed Window

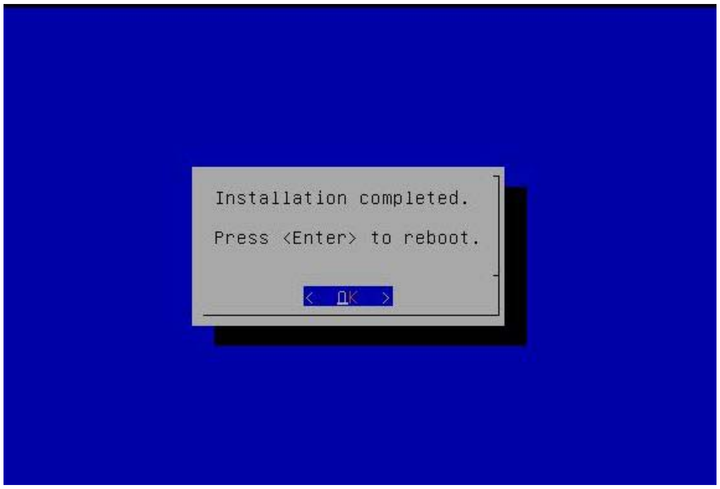

Installation completed.

Press <Enter> to reboot.

Note *Wait* for the system to reboot and display the BIOS splash screen, as shown in Figure 4-13 “System Reboot Window”, before proceeding to the next step.

Step 27 After the system reboots, from the KVM Console window, choose **Virtual Media** > **Activate Virtual Devices** to deactivate the virtual devices and unmap the image file.

Figure 4-22 Deactivate Virtual Device

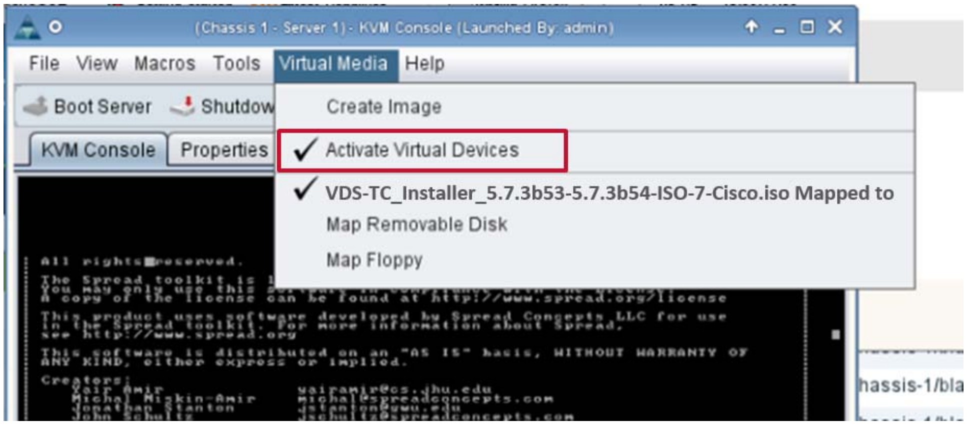
(Chassis 1 - Server 1) - KVM Console (Launched By: admin)
File View Macros Tools Virtual Media Help
Boot Server Shutdow
KVM Console Properties
Create Image
✓ Activate Virtual Devices
✓ VDS-TC_Installer_5.7.3b53-5.7.3b54-ISO-7-Cisco.iso Mapped to
Map Removable Disk
Map Floppy
hassis-1/bla

```
All rights reserved.
The Spread toolkit is l
You may only use this
A copy of the license can be found at http://www.spread.org/license
This product uses software developed by Spread Concepts LLC for use
in the Spread toolkit. For more information about Spread,
see http://www.spread.org
This software is distributed on an "AS IS" basis, WITHOUT WARRANTY OF
ANY KIND, either express or implied.
Creators:
    Yair Amir              yairamir@cs.jhu.edu
    Michal Miskin-Amir     michal@spreadconcepts.com
    Jonathan Stanton       jstanton@gwu.edu
    John Schultz           jschultz@spreadconcepts.com
```

Step 28 From the Close Virtual Media dialog box, click **Yes** to confirm the unmapping of the image file. This unmounts the ISO image.

Figure 4-23 Close Virtual Media

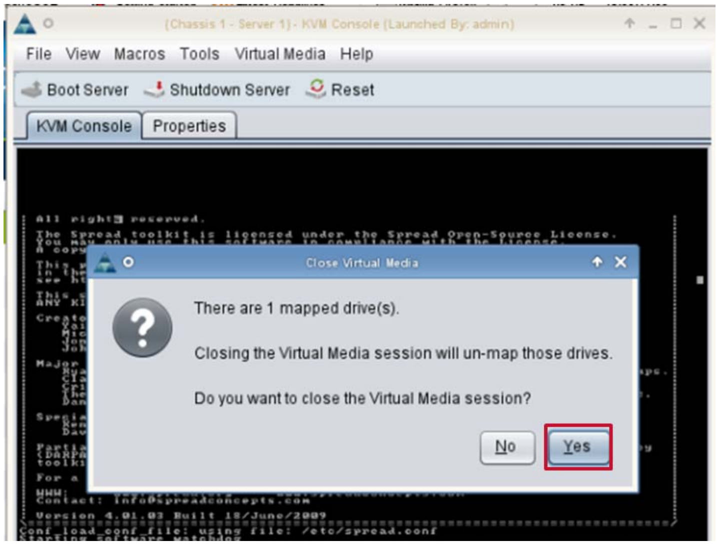
(Chassis 1 - Server 1) - KVM Console (Launched By: admin)
File View Macros Tools Virtual Media Help
Boot Server Shutdown Server Reset
KVM Console Properties
Close Virtual Media
No
Yes

```
All right reserved.
The Spread toolkit is licensed under the Spread Open-Source License.
Contact: info@spreadconcepts.com
Version 4.01.03 Built 18/June/2009
conf_load_conf_file: using file: /etc/spread.conf
```

There are 1 mapped drive(s).

Closing the Virtual Media session will un-map those drives.

Do you want to close the Virtual Media session?

Step 29 Click **Reset** under the menu bar to reboot the cache engine.

Figure 4-24 Reset the Cache Engine

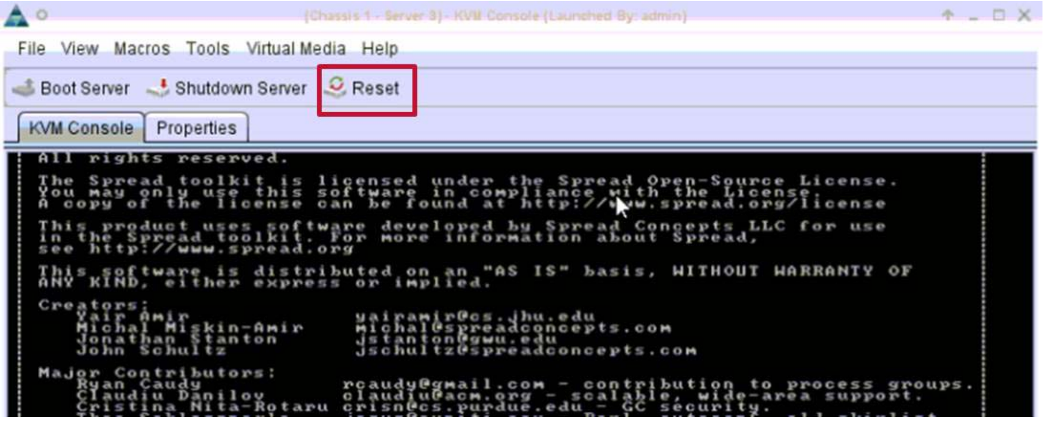
(Chassis 1 - Server 3) - KVM Console (Launched By: admin)
File View Macros Tools Virtual Media Help
Boot Server
Shutdown Server
Reset
KVM Console
Properties

```
All rights reserved.

The Spread toolkit is licensed under the Spread Open-Source License.
You may only use this software in compliance with the License.
A copy of the license can be found at http://www.spread.org/license

This product uses software developed by Spread Concepts LLC for use
in the Spread toolkit. For more information about Spread,
see http://www.spread.org

This software is distributed on an "AS IS" basis, WITHOUT WARRANTY OF
ANY KIND, either express or implied.

Creators:
    Yair Amir              yairamir@cs.jhu.edu
    Michal Miskin-Amir     michal@spreadconcepts.com
    Jonathan Stanton       jstanton@gwu.edu
    John Schultz           jschultz@spreadconcepts.com

Major Contributors:
    Ryan Caudy            rcaudy@gmail.com - contribution to process groups.
    Claudiu Danilov       claudiu@acm.org - scalable, wide-area support.
    Cristina Nita-Rotaru  crisn@cs.purdue.edu - GC security.
```

Step 30 From the Reset Server Warning window that appears, click **OK** to reset the server.

Figure 4-25 Reset Server Warning Window

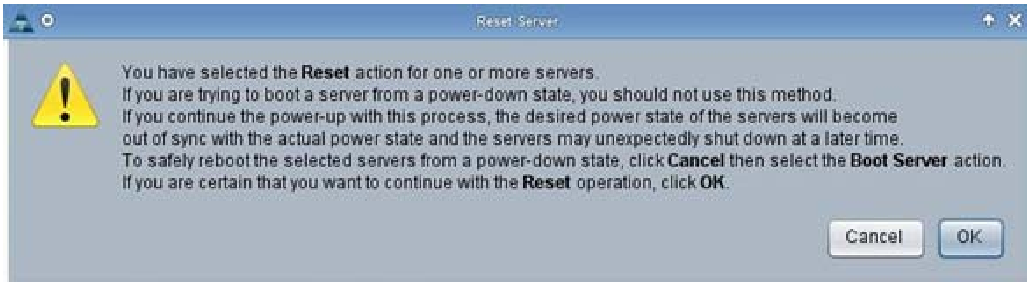

You have selected the **Reset** action for one or more servers.
If you are trying to boot a server from a power-down state, you should not use this method.
If you continue the power-up with this process, the desired power state of the servers will become
out of sync with the actual power state and the servers may unexpectedly shut down at a later time.
To safely reboot the selected servers from a power-down state, click **Cancel** then select the **Boot Server** action.
If you are certain that you want to continue with the **Reset** operation, click **OK**.

Step 31 From the Reset Server Service window that appears, choose **Power Cycle** and click **OK**.

Figure 4-26 Reset Server Service

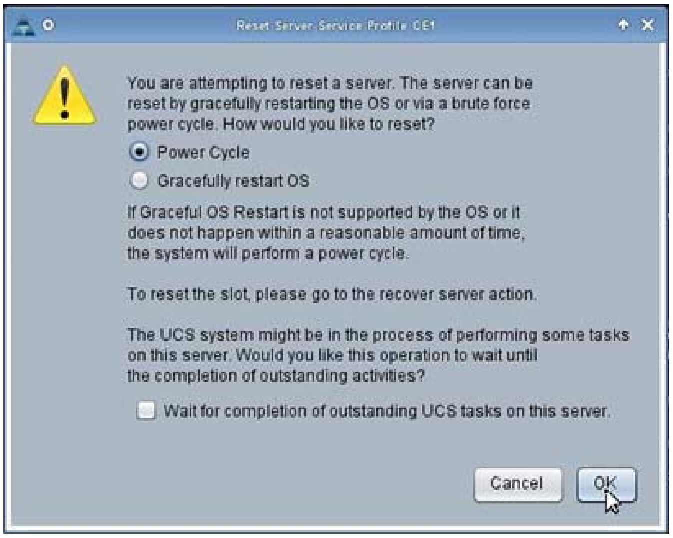

You are attempting to reset a server. The server can be reset by gracefully restarting the OS or via a brute force power cycle. How would you like to reset?

● Power Cycle

○ Gracefully restart OS

If Graceful OS Restart is not supported by the OS or it does not happen within a reasonable amount of time, the system will perform a power cycle.

To reset the slot, please go to the recover server action.

The UCS system might be in the process of performing some tasks on this server. Would you like this operation to wait until the completion of outstanding activities?

□ Wait for completion of outstanding UCS tasks on this server.

Updating the SSH Keys and Connecting the Cache Engine to the Storage Enclosure

After completing the "Running the Post Installation Scripts" section in the "Cache Engine Servers Installation" chapter, you must update the grid SSH keys on the VDS TC management server and configure the new cache engine to connect to the storage enclosure. All of these steps are performed on the VDS TC management server.

Follow this procedure to perform these steps on the VDS TC management server:

Note The following steps are performed on the VDS TC management server.

Procedure

Step 1 On the PBR routers, stop redirecting traffic to the VDS TC solution.

Step 2 Close any open VDS TC Manager windows and ensure that no other administrators are connected to the VDS TC Manager before proceeding.

Step 3 Using SSH software, such as Putty, open an SSH connection to the IP address of the VDS TC management server.

Step 4 Log into the system using the username **padmin** and the password provided by Cisco.

Step 5 Enter the command **su admin** to log into the CLI of the VDS TC management server. The default password is the serial number of the VDS TC management server.

Step 6 From the VDS TC prompt, enter the **enable** command. When prompted, enter the Enable mode password and press **Enter**. You are now logged into Enable mode and the Enable prompt, console#, should appear.

Step 7 From the Enable mode prompt, enter the command **oper service stop** to stop the caching service.

Step 8 Wait a few minutes and then enter the command **show status**. Check to see if the Device State shows "stopped." For example:

```
console# show status
Operational state Device state Administrative state
disabledstoppedunlocked
```

Caution Do not proceed to the next step until the Device State shows "stopped." You may need to repeat the **show status** command several times before you see this status. Wait several minutes between executions of the **show status** command.

Step 9 Enter the **exit** command to exit Enabled mode.

Step 10 Enter the **exit** command to exit the VDS TC CLI.

Step 11 Enter the command **su root** to change to the root user. Enter the password for the root user when prompted. This password is provided by Cisco.

Step 12 Enter the command **cd /opt/pang/useful/**.

Step 13 To reset the SSH keys, enter the command **./replace_server_keys.sh -servers** #, where # is the total number of Blade Server cache engines in the VDS TC installation. For example, if you are replacing a cache engine in a VDS TC solution that has 16 cache engines, you would enter the command **./replace_server_keys.sh -servers 16**.

Step 14 Enter the command **cd /opt/pang/useful/configure_storages** to change to the storage configuration script directory.

Step 15 Enter the command **./configure_storages.py -s** <*no_of_storage_enclosures*> **-b** <*no_of_blades*> **-c** <*the_replaced_CE-ID_number*>. For example, to replace CE-2 in a VDS TC Blade Server installation that has two cache engines, and one storage enclosure, enter the command **./configure_storages.py -s 1 -b 2 -c 2**.

Note The configure storages script will format the data disks. To configure the storages without formatting the disks enter the command **./configure_storages.py -s** <*no_of_storage_enclosures*> **-b** <*no_of_blades_in_grid*> **-c** <*the_replaced_CE-ID_number*> **-x**.

Step 16 When the warning message about the disk format appears, enter **y** to continue the storages configuration.

```
*********** WARNING! **********
All storages are about to be reconfigured. All data on the data disks will
be lost.
Do you want to continue? (y/n)y
```

Step 17 Follow the on screen instructions. If prompted, enter the password for the root user that was provided by Cisco.

Note You may be prompted to enter the password several times.

When the script is finished running, you should see the word “COMPLETE”. After running this script, the new VDS TC cache engine will be connected to the storage enclosures.

Step 18 Use SSH to connect to each CE and enter the command **iscsiadm -m session** to check the connections:

```
ssh ce-<cache engine number>
iscsiadm -m session
```

If you are using 1Gb connections to the storage enclosures, your output should be similar to the following:

Note The following output will look different if you are using 10Gb connections to the storage enclosures.

```
ƒÞ 8 * sessions per storage, for example for 1 storage:
tcp: [17] 10.11.14.100:3260,1 iqn.1992-
01.com.Avago:2365.60080e50002f4e82000000005063c74a
tcp: [18] 10.11.14.101:3260,2 iqn.1992-
01.com.Avago:2365.60080e50002f4e82000000005063c74a
tcp: [19] 10.11.16.100:3260,1 iqn.1992-
01.com.Avago:2365.60080e50002f4e82000000005063c74a
tcp: [20] 10.11.16.101:3260,2 iqn.1992-
01.com.Avago:2365.60080e50002f4e82000000005063c74a
tcp: [21] 10.11.15.100:3260,1 iqn.1992-
01.com.Avago:2365.60080e50002f4e82000000005063c74a
tcp: [22] 10.11.15.101:3260,2 iqn.1992-
01.com.Avago:2365.60080e50002f4e82000000005063c74a
tcp: [23] 10.11.17.101:3260,2 iqn.1992-
01.com.Avago:2365.60080e50002f4e82000000005063c74a
tcp: [24] 10.11.17.100:3260,1 iqn.1992-
01.com.Avago:2365.60080e50002f4e82000000005063c74a
```

Updating the NetApp Storage Enclosure

Note This process is only for the NetApp storage enclosure. If your system uses an IBM storage enclosure, refer to the Updating the IBM Storage Enclosure section.

After you have configured the VDS TC cache engine to connect with the storage enclosure, you must update the host port identifiers on the NetApp storage enclosure for the new cache engine. Follow this procedure to perform this task:

Procedure

Step 1 From the SSH connection to the VDS TC management server, enter the following commands to start the VNC server:

a. **rm /tmp/.X1-lock** (This step is optional. It is only required if there is a leftover lock file from a previous vncserver instance.)

b. **rm /tmp/.X11-unix/X1** (This step is optional. It is only required if there is a leftover lock file from a previous vncserver instance.)

c. **vncserver :1**

Step 2 Use a VNC client to connect from your local computer to the VDS TC management server.

Figure 4-27 VNC Viewer

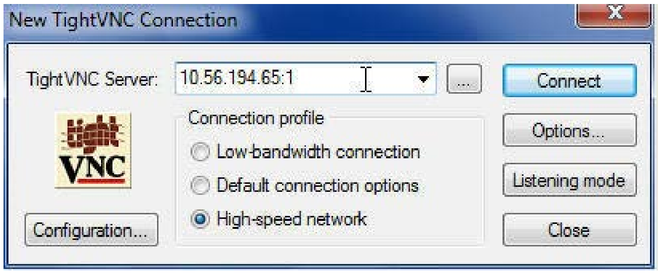

New TightVNC Connection

TightVNC Server: 10.56.194.65:1

Connection profile

○ Low-bandwidth connection
○ Default connection options
● High-speed network

Step 3 From the VNC console window, enter the command **cd /opt/SMgr/client** to change folders.

Step 4 Enter the command **./SMclient** to start the SANtricity ™ ES Storage Manager software.

Note You must perform the following steps on all of the storage subsystems.

Step 5 Right-click the storage subsystem that you are configuring and choose **Manage Storage Array**. Click **No** in any pop-up windows that appear.

Figure 4-28 SANtricity Storage Manager

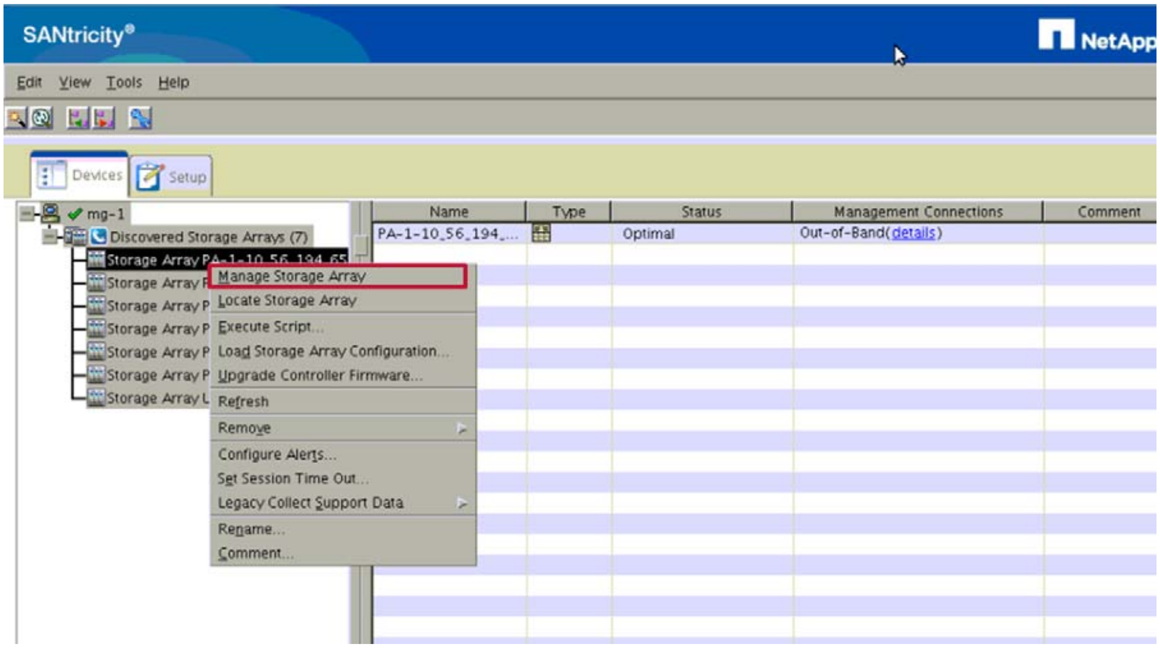

Step 6 In the Subsystem Management window that appears, click the **Host Mappings** tab.

Step 7 In the navigation pane on the left, expand the **Host Group** folder.

Step 8 Right-click the cache engine that was replaced, for example Host CE-1, and choose **Manage Host Port Identifiers** from the pop-up menu.

Figure 4-29 Manage Host Port Identifiers

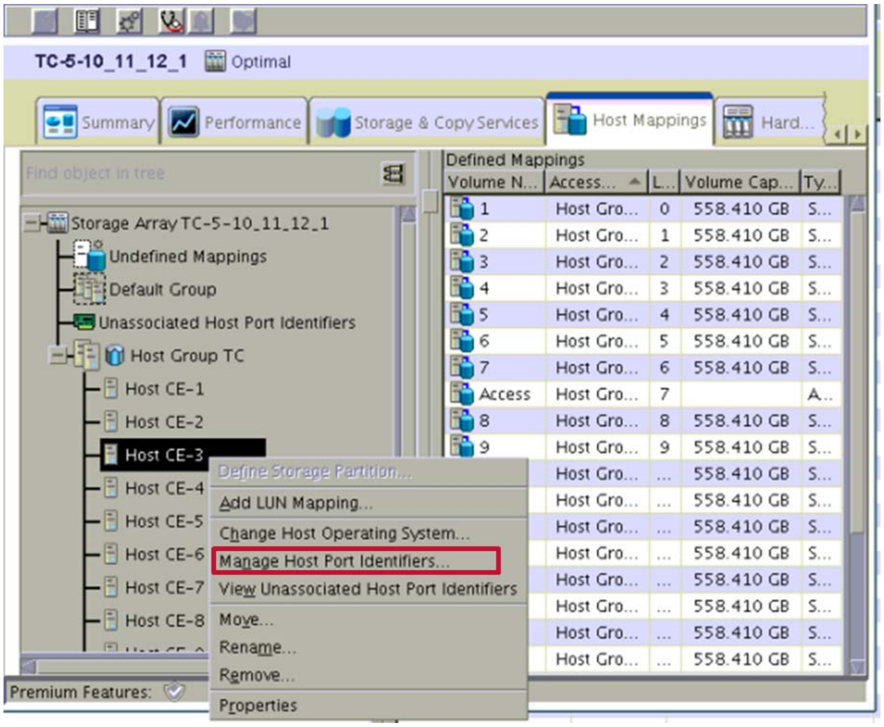

Defined Mappings

Volume N...	Access...	L...	Volume Cap...	Ty...
1	Host Gro...	0	558.410 GB	S...
2	Host Gro...	1	558.410 GB	S...
3	Host Gro...	2	558.410 GB	S...
4	Host Gro...	3	558.410 GB	S...
5	Host Gro...	4	558.410 GB	S...
6	Host Gro...	5	558.410 GB	S...
7	Host Gro...	6	558.410 GB	S...
Access	Host Gro...	7		A...
8	Host Gro...	8	558.410 GB	S...
9	Host Gro...	9	558.410 GB	S...
	Host Gro...	...	558.410 GB	S...
	Host Gro...	...	558.410 GB	S...
	Host Gro...	...	558.410 GB	S...
	Host Gro...	...	558.410 GB	S...
	Host Gro...	...	558.410 GB	S...
	Host Gro...	...	558.410 GB	S...
	Host Gro...	...	558.410 GB	S...
	Host Gro...	...	558.410 GB	S...

Step 9 From the Manage Host Port Identifiers window that appears, in the Host Port Identifier Information table, ensure that the old iqn is selected and click **Replace**.

Figure 4-30 Replace IQN

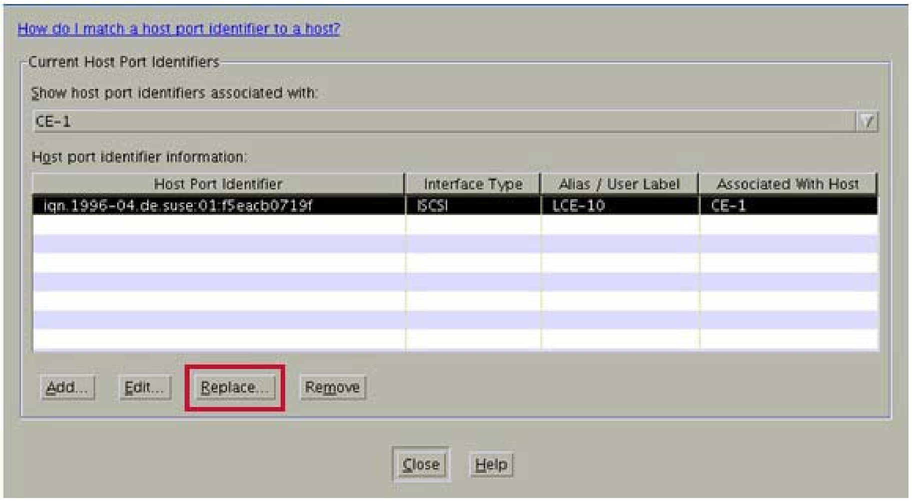
How do I match a host port identifier to a host?
Current Host Port Identifiers
Add...
Edit...
Replace...
Remove
Close
Help

Show host port identifiers associated with: CE-1

Host port identifier information:

Host Port Identifier	Interface Type	Alias / User Label	Associated With Host
iqn.1996-04.de.suse:01:f5eacb0719f	ISCSI	LCE-10	CE-1

Step 10 From the Replace Host Port Identifier window that appears, configure the following:

- From the Choose a Host Interface Type drop-down list, choose **ISCSI**.
- Click the **Replace by Selecting a Known Unassociated Host Port Identifier** radio button.

- From the Known Unassociated Host Port Identifier drop-down list, choose the iqn for the *newly* installed cache engine. (It should be the only one in the list.)

Step 11 Click **Replace**.

Figure 4-31 Replace Host Port Identifier

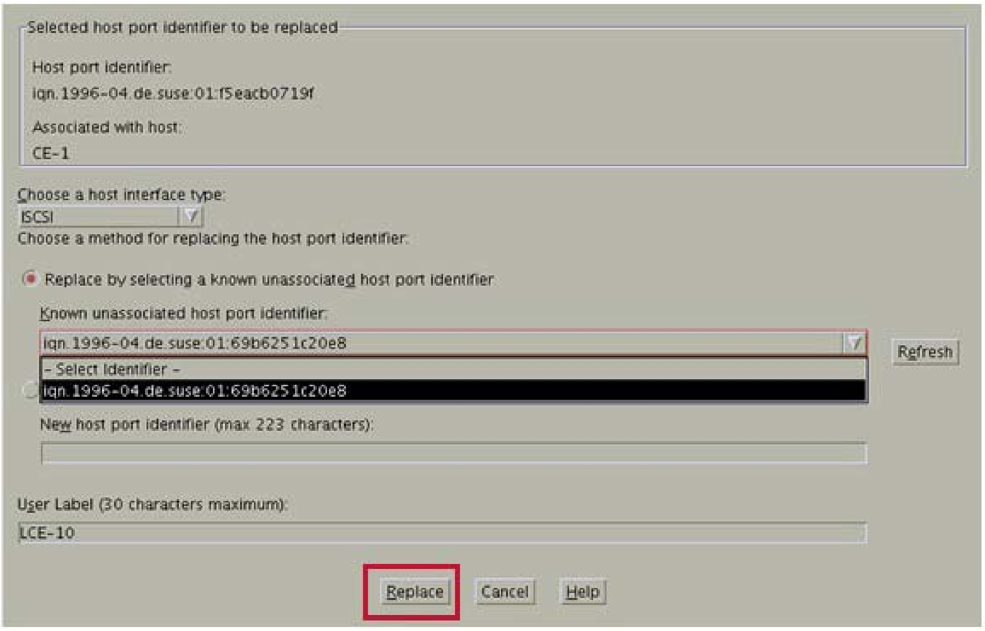

Selected host port identifier to be replaced

Host port identifier:
iqn.1996-04.de.suse:01:f5eacb0719f

Associated with host:
CE-1

Choose a host interface type:
ISCSI

Choose a method for replacing the host port identifier:

● Replace by selecting a known unassociated host port identifier

Known unassociated host port identifier:
iqn.1996-04.de.suse:01:69b6251c20e8
- Select Identifier -
iqn.1996-04.de.suse:01:69b6251c20e8

New host port identifier (max 223 characters):

User Label (30 characters maximum):
LCE-10

Step 12 From the Manage Host Port Identifiers window, click **Close**.

Step 13 Repeat Step 5 through Step 12 on any remaining storage subsystems.

Step 14 Close the Subsystem Management window and exit the SANtricity Storage Manager.

Step 15 Close the VNC client.

Updating the IBM Storage Enclosure

Note This process is only for the IBM storage enclosure. If your system uses the NetApp E2724 storage enclosure, refer to the Updating the NetApp Storage Enclosure section.

After you have configured the VDS TC cache engine to connect with the storage enclosure, you must update the host port identifiers on the IBM storage enclosure for the new cache engine. Follow this procedure to perform this task:

Procedure

Step 1 From the SSH connection to the VDS TC management server, enter the following commands to start the VNC server:

a. **rm /tmp/.X1-lock** (This step is optional. It is only required if there is a leftover lock file from a previous vncserver instance.)

b. **rm /tmp/.X11-unix/X1** (This step is optional. It is only required if there is a leftover lock file from a previous vncserver instance.)

c. **vncserver :1**

Step 2 Use a VNC client to connect from your local computer to the VDS TC management server.

Figure 4-32 VNC Viewer

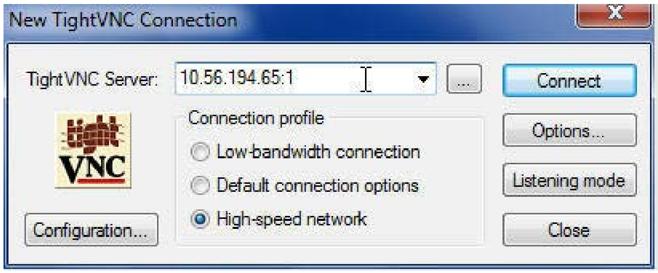

TightVNC Server: 10.56.194.65:1

Connection profile

○ Low-bandwidth connection

○ Default connection options

● High-speed network

Step 3 From the VNC console window, enter the command **/opt/IBM_DS/client/SMclient** to start the IBM Storage Manager.

Note You must perform the following steps on all of the storage subsystems.

Step 4 Right-click the storage subsystem that you are configuring and choose **Manage Storage Subsystem**. Click **No** in any pop-up windows that appear.

Figure 4-33 IBM DS Storage Manager

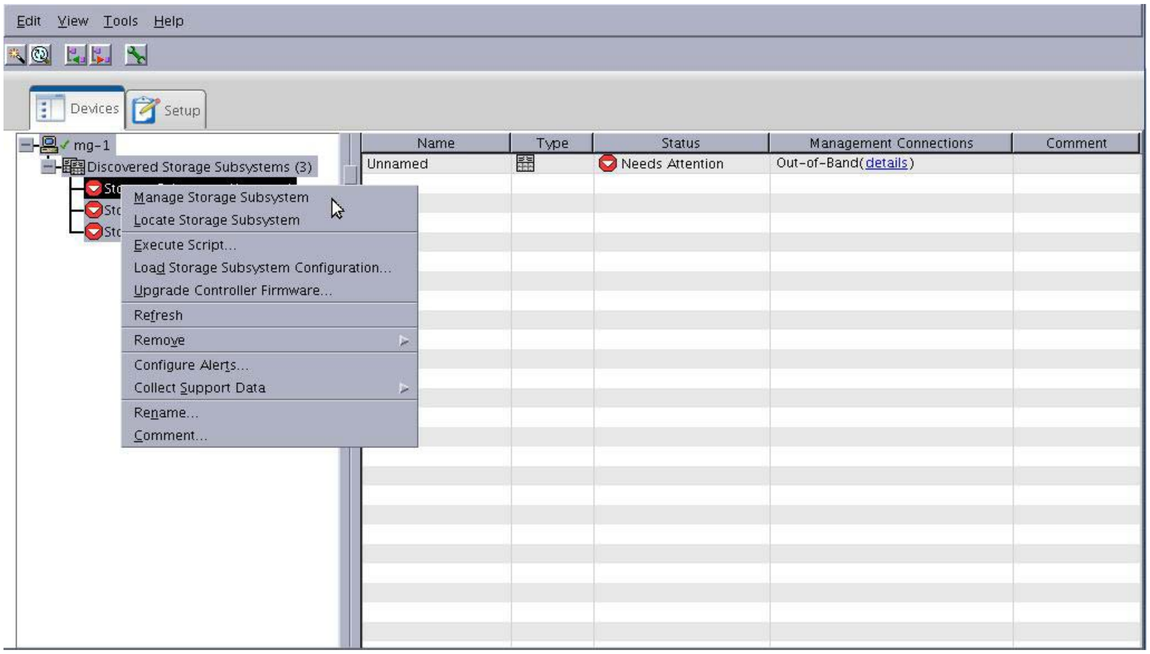

Name	Type	Status	Management Connections	Comment
Unnamed		Needs Attention	Out-of-Band(details)	

Step 5 In the Subsystem Management window that appears, click the **Host Mappings** tab.

Step 6 In the navigation pane on the left expand the **Host Group** folder.

Step 7 Right-click the cache engine that was replaced, for example Host CE-2, and choose **Manage Host Port Identifiers** from the pop-up menu.

Figure 4-34 Manage Host Port Identifiers

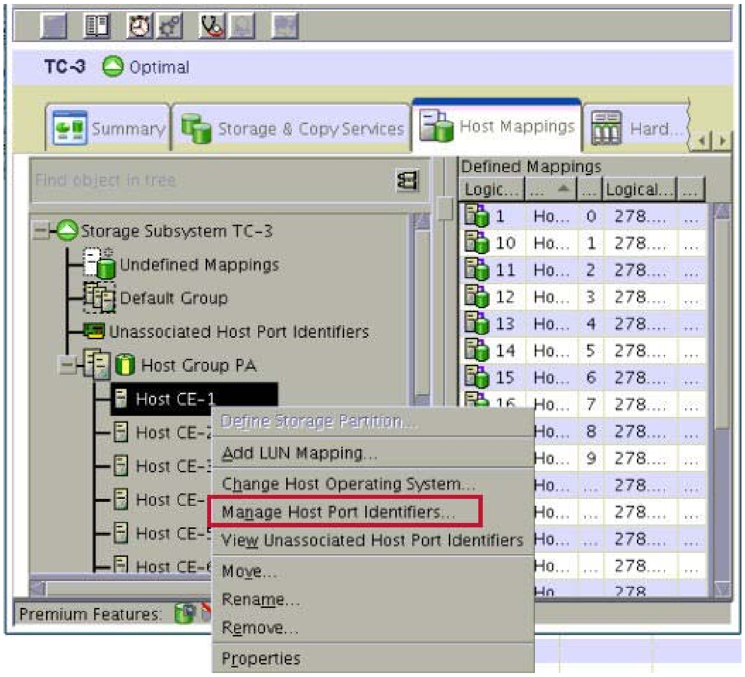

Logic...	...	...	Logical...	...
1	Ho...	0	278....	...
10	Ho...	1	278....	...
11	Ho...	2	278....	...
12	Ho...	3	278....	...
13	Ho...	4	278....	...
14	Ho...	5	278....	...
15	Ho...	6	278....	...
16	Ho...	7	278....	...
	Ho...	8	278....	...
	Ho...	9	278....	...
	Ho...	...	278....	...
	Ho...	...	278....	...
	Ho...	...	278....	...
	Ho...	...	278....	...
	Ho		278	

Step 8 From the Manage Host Port Identifiers window that appears, in the Host Port Identifier Information table, ensure that the iqn is selected and click **Replace**.

Figure 4-35 Replace IQN

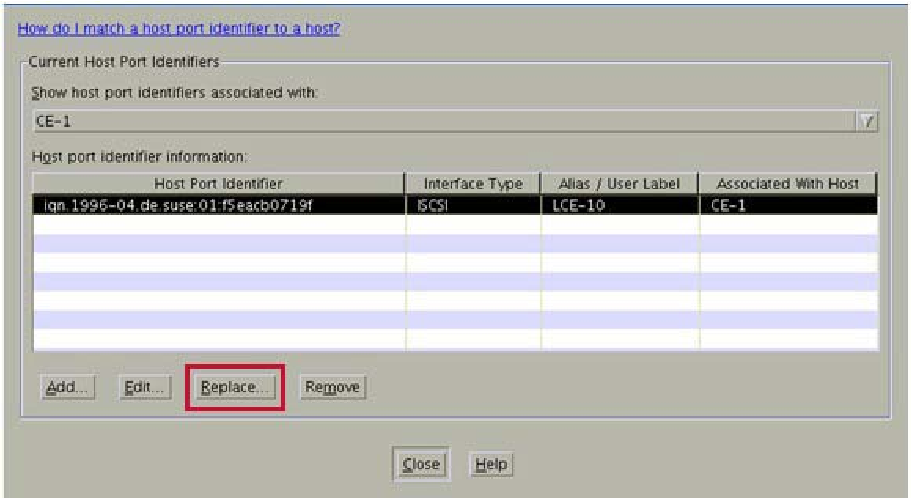
How do I match a host port identifier to a host?
Current Host Port Identifiers
Add...
Edit...
Replace...
Remove
Close
Help

Show host port identifiers associated with: CE-1

Host port identifier information:

Host Port Identifier	Interface Type	Alias / User Label	Associated With Host
iqn.1996-04.de.suse:01:f5eacb0719f	ISCSI	LCE-10	CE-1

Step 9 From the Replace Host Port Identifier window that appears, configure the following:

- From the Choose a Host Interface Type drop-down list, choose **ISCSI**.
- Click the **Replace by Selecting a Known Unassociated Host Port Identifier** radio button.

- From the Known Unassociated Host Port Identifier drop-down list, choose the iqn for the newly installed cache engine. (It should be the only one in the list.)

Step 10 Click **Replace**.

Figure 4-36 Replace Host Port Identifier

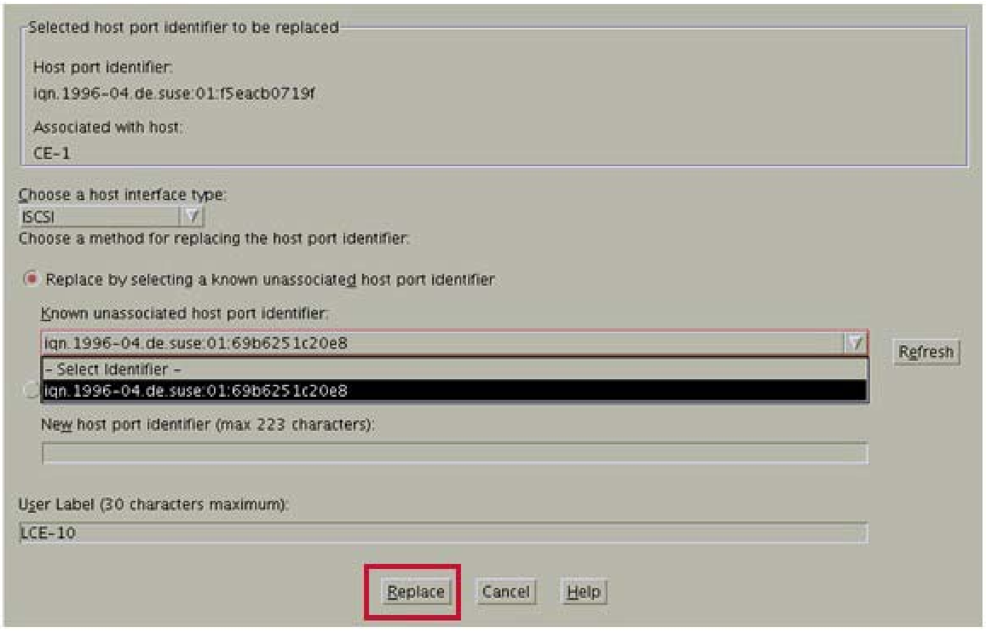

Selected host port identifier to be replaced

Host port identifier:
iqn.1996-04.de.suse:01:f5eacb0719f

Associated with host:
CE-1

Choose a host interface type:
ISCSI

Choose a method for replacing the host port identifier:

● Replace by selecting a known unassociated host port identifier

Known unassociated host port identifier:
iqn.1996-04.de.suse:01:69b6251c20e8

- Select Identifier -
iqn.1996-04.de.suse:01:69b6251c20e8

New host port identifier (max 223 characters):

User Label (30 characters maximum):
LCE-10

Step 11 From the Manage Host Port Identifiers window, click **Close**.

Step 12 Repeat Step 4 through Step 11 on any remaining storage subsystems.

Step 13 Close the Subsystem Management window and exit the IBM Storage Manager.

Step 14 Close the VNC client.

Finishing the Cache Engine Replacement

Follow this procedure to finish configuring and installing the new cache engine:

Note You will be re-running the storage script that you ran in the Updating the SSH Keys and Connecting the Cache Engine to the Storage Enclosure section.

Note You must be logged in with root permissions to perform the following steps.

Procedure

Step 1 From the SSH connection to the VDS TC management server, enter the command **cd /opt/pang/useful/configure_storages** to change to the storage configuration script directory.

Step 2 Enter the command **./configure_storages.py -s** *no_of_storage_enclosures* **-b** *no_of_blades* **-c** *the_replaced_CE-ID_number*. For example, to replace CE-2 in a VDS TC Blade Server installation that has two cache engines, and one storage enclosure, enter the command **./configure_storages.py -s 1 -b 2 -c 2**.

Step 3 When the warning message about the disk format appears, enter **y** to continue the storages configuration.

```
*********** WARNING! **********
All storages are about to be reconfigured. All data on the data disks will
be lost.
Do you want to continue? (y/n)y
```

Step 4 Follow the on screen instructions. If prompted, enter the password for the root user that was provided by Cisco. You should see the word “COMPLETE” when the script is finished running.

Note You may be prompted to enter the password several times.

Note Ensure that the VDS TC image upgrade file and the new VDS TC license that you received from Cisco are located on the VDS TC management server in the /tftpboot folder before performing these steps.

Step 5 Enter the command **su admin** to log into the VDS TC CLI. The CLI prompt console> appears.

Step 6 Enter the command **enable** to enter Enable mode. At the password prompt that appears, enter the enable password. The CLI prompt console# appears.

Step 7 Upgrade the software on the new cache engine by entering the command **upgrade server** *CE_number* **127.0.0.1 VDS-TC_FI_5.7.3b54_Server_Cluster.tar.gz**, where *CE_number* is the cache engine you replaced.

Step 8 To import the new license that includes the serial number of the new cache engine, enter the command **license import 127.0.0.1** *filename*, where *filename* is the name of the new license file.

Note You must be in Enable mode, *not* Configuration mode to import a new license.

Step 9 Enter the command **license activate** to apply and activate the license.

Step 10 To start the application on the replaced cache engine, enter the command **oper server** *#*, where *#* is the number of the cache engine that was replaced. For example, if you replaced CE-2, enter the command **oper server 2**. You should see the following output:

```
console# oper server 2
oper server 2# start
Starting server 2
service is started on server 2
```

Step 11 After the application has started on the new cache engine, enter the command **oper service start** to start the VDS TC service.

Step 12 Wait a few minutes and then enter the command **show status**. Do not proceed to the next step until you see a Device Status of “Started.” For example:

```
console# show status
Operational state Device state Administrative state
enabled started unlocked
```

Step 13 On the PBR routers, start redirecting traffic to the VDS TC solution.

Cache Engine Replacement Procedure for a 5.7.3 Blade Server Installation Upgraded from Release 5.2.1

Note Only use this chapter if you are replacing a Cache Engine in a VDS TC Blade Server system that was upgraded to VDS TC 5.7.3 from VDS TC 5.2.1. If you are replacing a Cache Engine in a VDS TC Blade Server system that has a fresh install of VDS TC 5.7.3 and was *not* upgraded to 5.7.3, refer to Chapter 4, "Cache Engine Replacement Procedure for a Blade Server Installation with a Fresh Install of VDS TC 5.7.3". If you are replacing a Cache Engine in a VDS TC Blade Server system that was upgraded to VDS TC 5.7.3 from VDS TC 5.6.1, refer to Chapter 6, "Cache Engine Replacement Procedure for a 5.7.3 Blade Server Installation Upgraded from Release 5.6.1".

To determine whether the system has been upgraded and from what version, from the Management Server enter the command **grep "upgrade system to" /var/log/peerapp/peerapp_system***. This will display a list of all previous versions.

This chapter discusses how to determine that a cache engine has failed in a VDS TC Blade Server installation and the steps that are required to replace it.

When a cache engine fails, the volume_selection_algorithm_time setting in the cluster_conf.xml file controls when the volumes of a failed cache engine will be redistributed (mounted) to the remaining cache engines. By default this value is 24 hours, which means that by default, if a cache engine fails, its volumes will not be redistributed (mounted) to the remaining cache engines until 24 hours after the cache engine fails.

It is recommended that you leave the default set to 24 hours, however if you need to change this setting, add the following text to the cluster_conf file in the <service> section under the <memory> subsection. The *seconds* parameters is the number of seconds that the system should wait after a cache engine fails before its volumes are redistributed (mounted) to the remaining cache engines:

```
<io>
    <volume_selection_algorithm_time>seconds</volume_selection_algorithm_time>
</io>
```

For example, to have the system wait 30 minutes after a cache engine fails before redistributing the volumes of the failed cache engine enter the following:

```
<io>
    <volume_selection_algorithm_time>1800</volume_selection_algorithm_time>
</io>
```

For more information on how to make changes to the configuration file and how to apply these changes, please refer to the *Cisco Videoscape Distribution Suite Transparent Caching Software Configuration Guide* available at http://www.cisco.com/c/dam/en/us/td/docs/video/videoscape/distribution_suite/vds/v5_7_3/VDS-TC_5_7_3_software_config_guide.pdf.

Symptoms of a Failed VDS TC Blade Server Cache Engine

To determine whether a cache engine has failed in a VDS TC Blade Server installation, log into VDS TC Manager GUI. By default, the Status > Dashboard window is displayed. If you are already logged into VDS TC Manager, choose **Status** > **Dashboard**.

If a cache engine has failed it will appear with a gray box surrounding it.

Figure 5-1 Failed Cache Engine

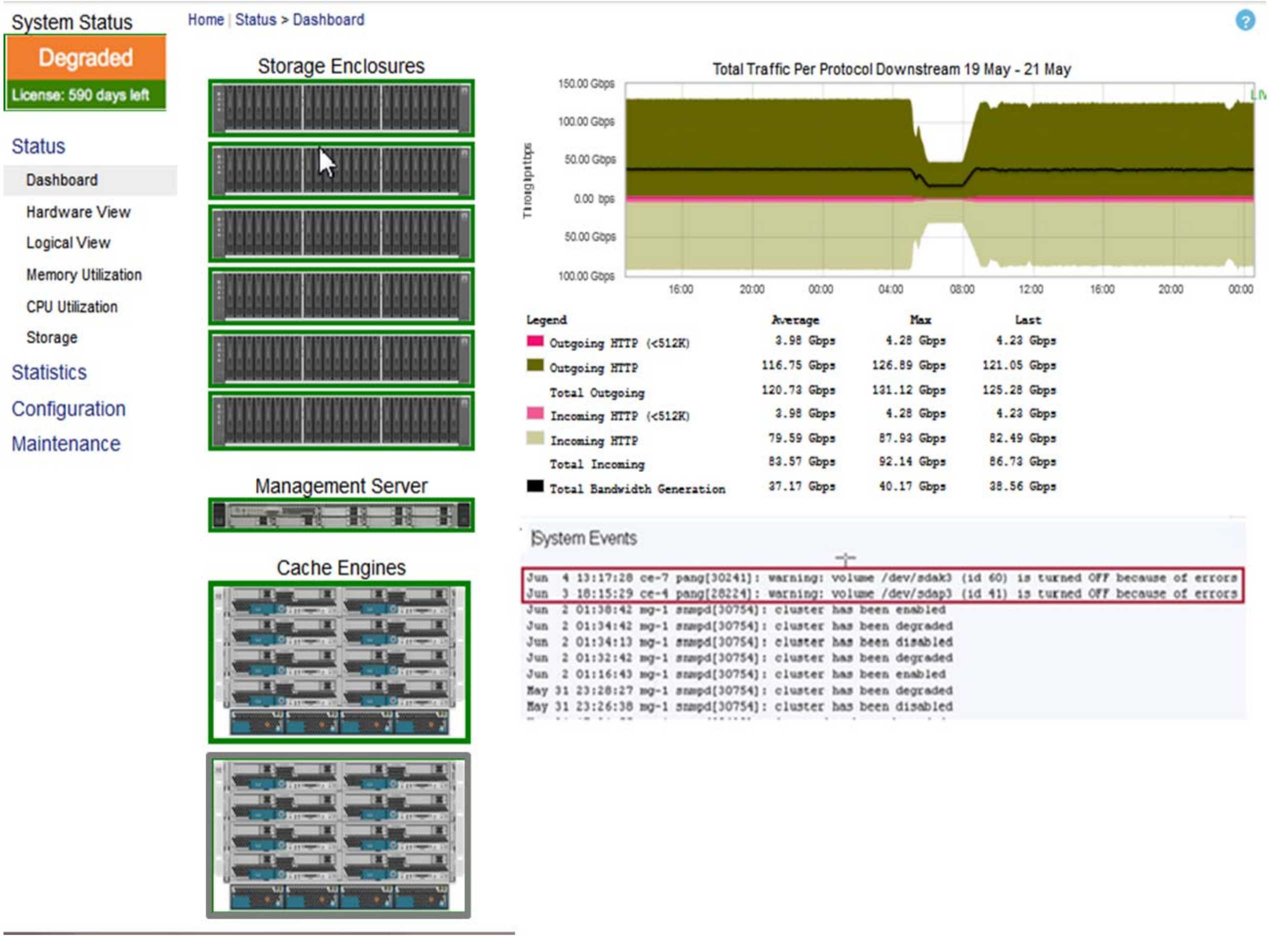

Legend	Average	Max	Last
🟥 Outgoing HTTP (<512K)	3.98 Gbps	4.28 Gbps	4.23 Gbps
🟫 Outgoing HTTP	116.75 Gbps	126.89 Gbps	121.05 Gbps
Total Outgoing	120.73 Gbps	131.12 Gbps	125.28 Gbps
🟥 Incoming HTTP (<512K)	3.98 Gbps	4.28 Gbps	4.23 Gbps
🟨 Incoming HTTP	79.59 Gbps	87.93 Gbps	82.49 Gbps
Total Incoming	83.57 Gbps	92.14 Gbps	86.73 Gbps
■ Total Bandwidth Generation	37.17 Gbps	40.17 Gbps	38.56 Gbps

```
Jun  4 13:17:28 ce-7 pang[30241]: warning: volume /dev/sdak3 (id 60) is turned OFF because of errors
Jun  3 18:15:29 ce-4 pang[28224]: warning: volume /dev/sdap3 (id 41) is turned OFF because of errors
Jun  2 01:38:42 mg-1 snmpd[30754]: cluster has been enabled
Jun  2 01:34:42 mg-1 snmpd[30754]: cluster has been degraded
Jun  2 01:34:13 mg-1 snmpd[30754]: cluster has been disabled
Jun  2 01:32:42 mg-1 snmpd[30754]: cluster has been degraded
Jun  2 01:16:43 mg-1 snmpd[30754]: cluster has been enabled
May 31 23:28:27 mg-1 snmpd[30754]: cluster has been degraded
May 31 23:26:38 mg-1 snmpd[30754]: cluster has been disabled
```

Obtaining and Uploading the Required Files

Obtaining Required Files

Before you replace the VDS TC cache engine, you must obtain the following:

- **The VDS TC 5.2.0 ISO image file:** VDS-TC_Installer-5.2.0b123-5.2.0b124-ISO-5.2.0b23-Cisco.iso
- **Updated kernel load file for VDS TC 5.2.1:** linux-2.6.27.19-5-llpf_10--KERNEL_LOAD_1414071783-r44244.tar. This file is published on cisco.com in the same location as the VDS TC 5.2.1 files.
- **The VDS TC 5.7.3 ISO image file:** VDS-TC_Installer_5.7.3b53-5.7.3b54-ISO-7-Cisco.iso
- **The Cisco_network_updater_v_1.8.tgz file:** Obtain this file from Cisco. This file should be available in the same location as the VDS TC 5.7.3 upgrade image file. This file is needed to correctly configure the cache engines for the Active- Active configuration that is supported by VDS TC 5.7.3 on a Blade Server installation.
- **A new VDS TC license from Cisco:** To replace a failed cache engine you must request a new VDS TC license from Cisco. To request a new license, you must have the serial number of the failed cache engine and the serial number of the new cache engine. After you have gathered this information, contact Cisco support for a new license file.

Copying the Files to the VDS TC Management Server

Procedure

Step 1 Using SFTP software, such as WinSCP, connect to the VDS TC Management Server using the IP address that you assigned during the “Configuring the Boot Image” process. Log in using the user name **padmin** and the password that was provided by Cisco.

Step 2 Create the folder **/opt/pang/iso**.

Step 3 Copy the VDS TC 5.2.0 ISO file, VDS-TC_Installer-5.2.0b123-5.2.0b124-ISO-5.2.0b23-Cisco.iso, into the **/opt/pang/iso** folder.

Step 4 Copy the VDS TC 5.7.3 ISO image file, VDS-TC_Installer_5.7.3b53-5.7.3b54-ISO-7-Cisco.iso, into the **/opt/pang/iso** folder.

Note You need the ISO image for both VDS TC 5.2.0 and VDS TC 5.7.3 to complete the Cache Engine replacement.

Step 5 Copy the Cisco_network_updater_v_1.8.tgz file to the **/tmp** folder.

Step 6 Copy the linux-2.6.27.19-5-llpf_10--KERNEL_LOAD_1406831120_rev41398.tar file to the **/tmp** folder.

Step 7 Copy the new VDS TC license that you received from Cisco to the **/tftpboot** folder.

Step 8 Close the SFTP software.

Extracting the VDS TC 5.7.3 Files

Step 1 From the SSH connection to the VDS TC management server logged in as root, enter the command **mkdir /opt/pang/iso/5.7.3_iso** to create a new folder for the 5.7.3 files.

Step 2 Enter the command **cd /opt/pang/iso** to change folders.

Step 3 Enter the command **mount VDS-TC_Installer_5.7.3b53-5.7.3b54-ISO-7-Cisco.iso -o loop /opt/pang/iso/5.7.3_iso** to extract the VDS-TC_FI_5.7.3b54_Server_Cluster.tar.gz file.

Step 4 Enter the command **cd /** to change to the root folder and then enter the command **umount /opt/pang/iso/5.7.3_iso** to unmount the iso image.

Step 5 Enter the command **rm /opt/pang/iso/VDS-TC_Installer_5.7.3b53-5.7.3b54-ISO-7-Cisco.iso** to delete the iso file.

Step 6 Enter the command **cd /opt/pang/iso/5.7.3_iso/install-tools/PeerApp_FI-GA**.

Step 7 Enter the command **cp VDS-TC_FI_5.7.3b54_Server_Cluster.tar.gz /tftpboot**.

Replacing a Cache Engine and Running the Installation Tasks

The first step in replacing a cache engine in a VDS TC Blade Server cluster installation is to physically replace the cache engine and perform the installation steps on the new cache engine. Follow this procedure to perform these steps:

Before You Begin

To replace a failed cache engine you must request a new VDS TC license from Cisco. To request a new license, you must have the serial number of the failed cache engine and the serial number of the new cache engine. After you have gathered this information, contact Cisco support for a new license file.

Caution Do not continue until you have the license file.

Procedure

Note Before removing the failed cache engine, make note of what number corresponds with the cache engine.

Step 1 Using the Cisco UCS 5108 Server Chassis Installation Guide available at http://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/hw/chassis-install-guide/ucs5108_install.html, unmount the failed Blade Server cache engine and mount the new Blade Server cache engine.

Step 2 Confirm that the correct firmware is installed on the Cache Engine for VDS TC 5.7.3:

- Cisco UCS-B200-M3:
  - Firmware version: 2.2.5b
  - Image file: ucs-k9-bundle-b-series.2.2.5b.B.bin

Step 3 In the *Cisco Videoscape Distribution Suite Transparent Caching Software Blade Server Cluster Installation Guide*, available at http://www.cisco.com/c/dam/en/us/td/docs/video/videoscape/distribution_suite/vds/v5_7_3/VDS-TC_5_7_3_BladeSrvr_swInstallGuide.pdf, refer to Chapter 1, "Prerequisites and Blade Server Cluster Physical Installation" for connecting the network cables to the new cache engine.

Note Complete the following steps for the cache engine that you are replacing.

Step 1 Use VNC to connect from your local computer to the VDS TC Management Server using the IP address that you configured, adding :1 to the end, for example 10.56.194.65:1. When prompted to enter a password, enter the root password that was provided by Cisco. If you encounter problems opening a VNC client connection to the VDS TC Management Server, you may need to restart VNC Server 1 on the VDS TC Management Server. Follow these steps to restart this server:

a. From an SSH connection to the VDS TC Management server where you are logged in with root permissions, enter the command **rm /tmp/.X11-unix/X1**.

b. Enter the command **vncserver :1** to start the VNC Server.

Figure 5-2 VNC Viewer

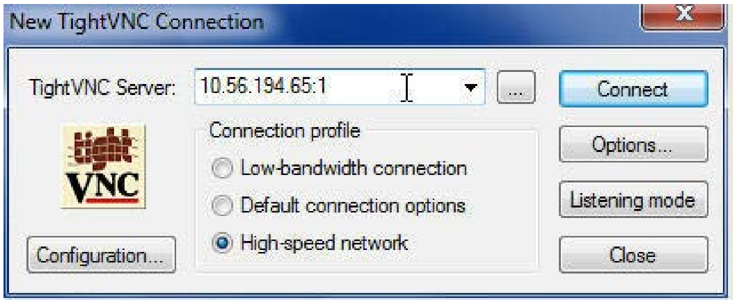

New TightVNC Connection

TightVNC Server: 10.56.194.65:1

Connection profile
○ Low-bandwidth connection
○ Default connection options
● High-speed network

Note If you have problems restarting the VNC Server 1 instance, you can also start another VNC Server instance by entering the command **vncserver :2**.

Step 2 From the VNC console, enter the command **firefox &**.

Figure 5-3 VNC Console

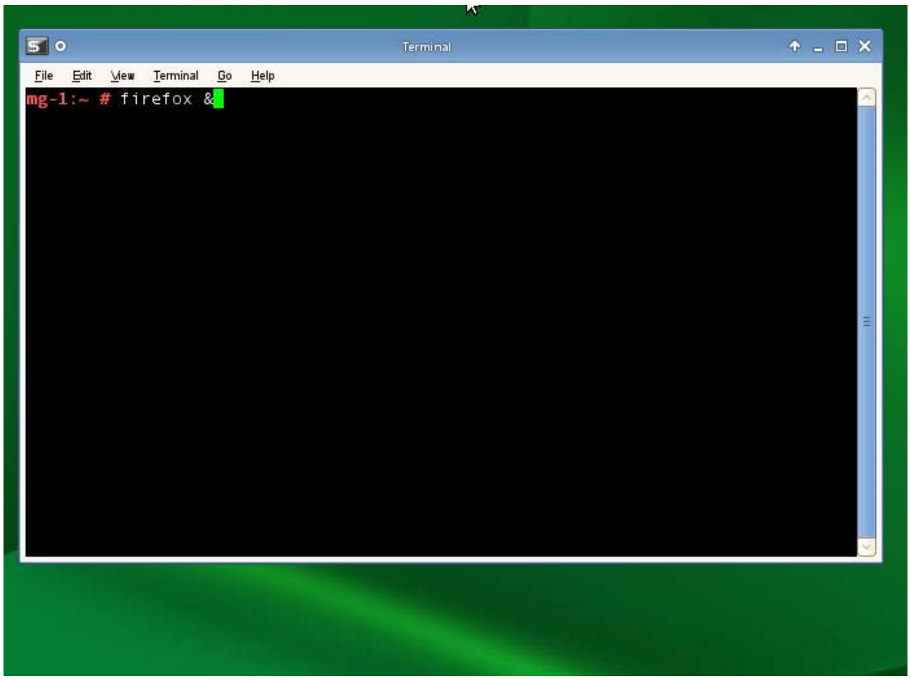

```
mg-1:~ # firefox &
```

Step 3 In the Firefox window that opens, enter the Management port IP address for the fabric interconnect. In a redundant installation, enter the Cluster (virtual) IP address.

Step 4 The Cisco UCS Manager window appears. From this window, click **Launch KVM Manager**. If you receive a message that says “To view this page ensure that Adobe Flash Player version 10.0.0 or greater is installed.” perform the following steps:

a. Close Firefox

b. From the terminal window, enter the command **su –**

c. Enter the command **pkill -9 firefox**

d. Enter the command **cd /root**

e. Enter the command **mv .mozilla .mozilla.bak**

f. Enter the command **firefox &** to restart Firefox.

g. In the Firefox window that opens, enter the Management port IP address for the fabric interconnect. In a redundant installation, enter the Cluster (virtual) IP address.

h. Repeat Step 4.

Figure 5-4 Cisco UCS Manager

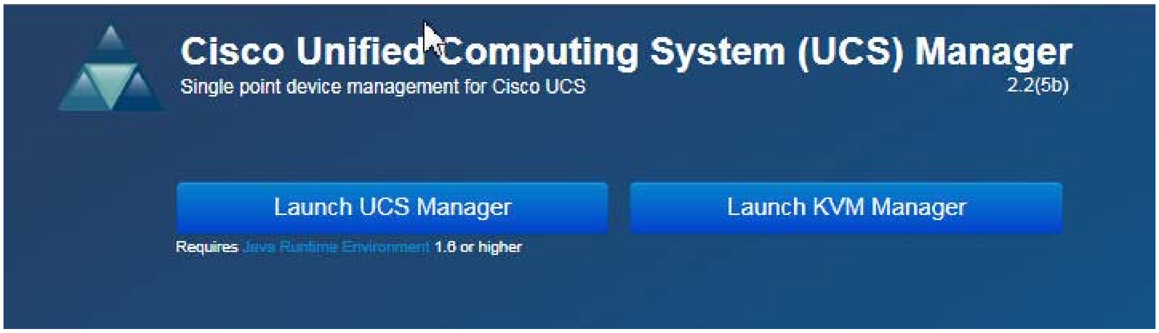
Cisco Unified Computing System (UCS) Manager
Single point device management for Cisco UCS
2.2(5b)
Launch UCS Manager
Launch KVM Manager
Requires Java Runtime Environment 1.6 or higher

Step 5 The UCS - KVM Launch Manager login window appears. Enter the username **admin** and the password that was configured during the initial system setup of the fabric interconnect (the default password is **password**), and then click **OK**.

Figure 5-5 UCS KVM Launch Manager Login Window

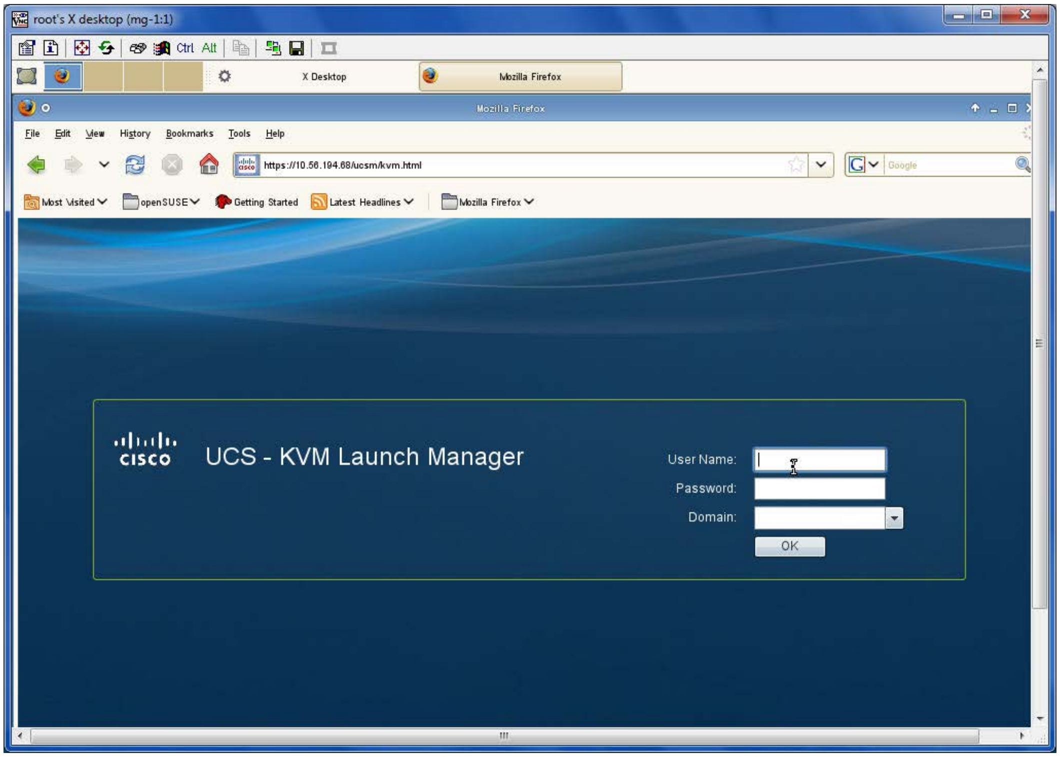

Step 6 The Service Profiles window appears, listing all of the cache engine profiles. From this window, in the row for the cache engine you are replacing, in the Launch KVM column click **Launch** to launch the KVM console.

Figure 5-6 Service Profiles Window

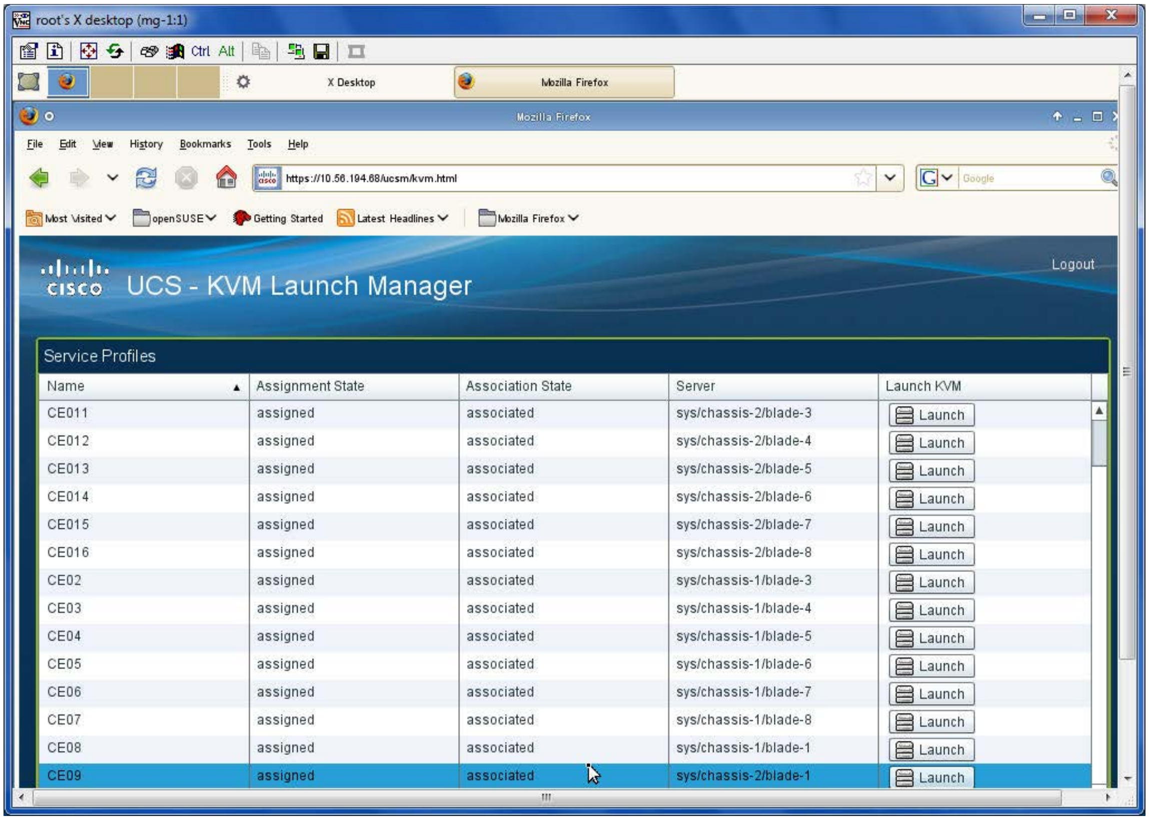

Service Profiles

Name	Assignment State	Association State	Server	Launch KVM
CE011	assigned	associated	sys/chassis-2/blade-3	Launch
CE012	assigned	associated	sys/chassis-2/blade-4	Launch
CE013	assigned	associated	sys/chassis-2/blade-5	Launch
CE014	assigned	associated	sys/chassis-2/blade-6	Launch
CE015	assigned	associated	sys/chassis-2/blade-7	Launch
CE016	assigned	associated	sys/chassis-2/blade-8	Launch
CE02	assigned	associated	sys/chassis-1/blade-3	Launch
CE03	assigned	associated	sys/chassis-1/blade-4	Launch
CE04	assigned	associated	sys/chassis-1/blade-5	Launch
CE05	assigned	associated	sys/chassis-1/blade-6	Launch
CE06	assigned	associated	sys/chassis-1/blade-7	Launch
CE07	assigned	associated	sys/chassis-1/blade-8	Launch
CE08	assigned	associated	sys/chassis-1/blade-1	Launch
CE09	assigned	associated	sys/chassis-2/blade-1	Launch

Step 7 Click **OK** to open the .jnlp file with Java(TM) Web Start Launcher.

Step 8 From the Warning Security window, click **Continue**.

Step 9 If you receive a window that asks if you want to run this application, click **Run**.

Step 10 From the KVM Console window, make sure the **KVM Console** tab is selected and choose **Virtual Media > Activate Virtual Devices**.

Step 11 If you see the Unencrypted Virtual Media Session window, choose **Accept this session** and click **Apply**.

Figure 5-7 Unencrypted Virtual Media Session

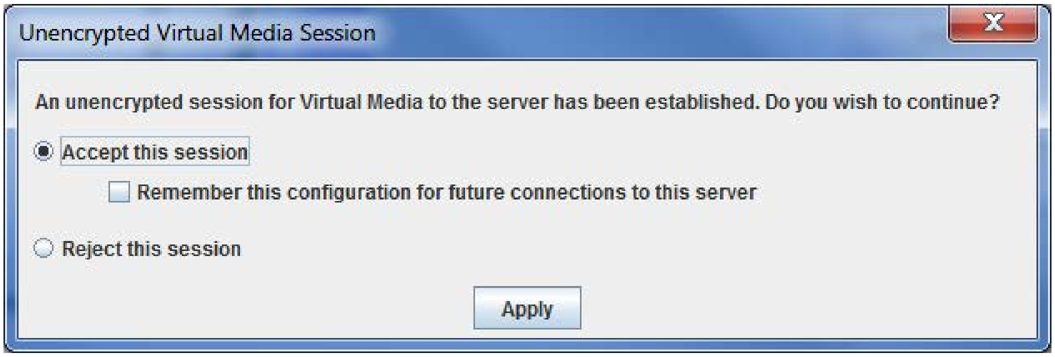

Unencrypted Virtual Media Session

An unencrypted session for Virtual Media to the server has been established. Do you wish to continue?

● Accept this session

□ Remember this configuration for future connections to this server

○ Reject this session

Step 12 Choose **Virtual Media > Map CD/DVD**.

Figure 5-8 Map CD/DVD

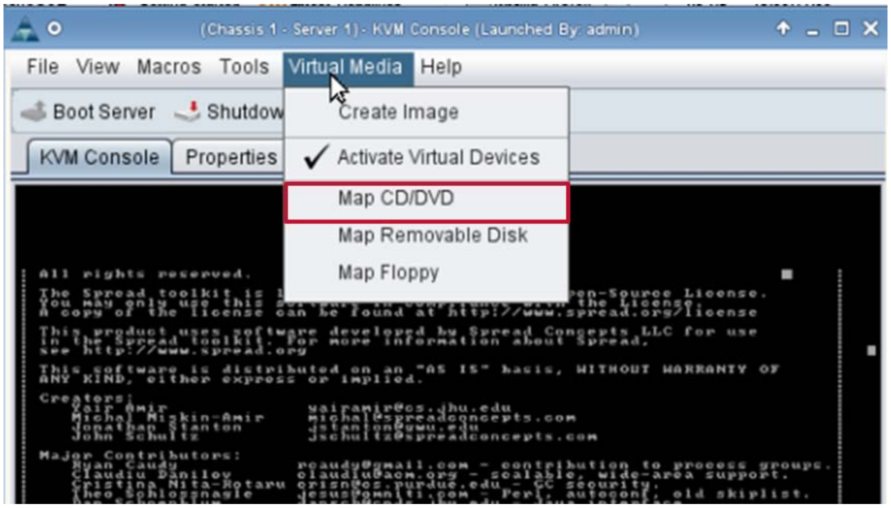

```
All rights reserved.
The Spread toolkit is l                                  pen-Source License.
You may only use this                                    the License.
A copy of the license can be found at http://www.spread.org/license
This product uses software developed by Spread Concepts LLC for use
in the Spread toolkit. For more information about Spread,
see http://www.spread.org
This software is distributed on an "AS IS" basis, WITHOUT WARRANTY OF
ANY KIND, either express or implied.
Creators:
    Yair Amir              yairamir@cs.jhu.edu
    Michal Miskin-Amir     michal@spreadconcepts.com
    Jonathan Stanton       jstanton@gwu.edu
    John Schultz           jschultz@spreadconcepts.com
Major Contributors:
    Ryan Caudy            rcaudy@gmail.com - contribution to process groups.
    Claudiu Danilov       claudiu@acm.org - scalable, wide-area support.
    Cristina Nita-Rotaru  crisn@cs.purdue.edu - GC security.
    Theo Schlossnagle     jesus@omniti.com - Perl, autoconf, old skiplist.
```

Step 13 In the window that opens, browse to the **/opt/pang/iso/** folder to find the ISO installer image for VDS TC 5.2.0, VDS-TC_Installer-5.2.0b123-5.2.0b124-ISO-5.2.0b23-Cisco.iso, and click **Open**.

Figure 5-9 Map Device

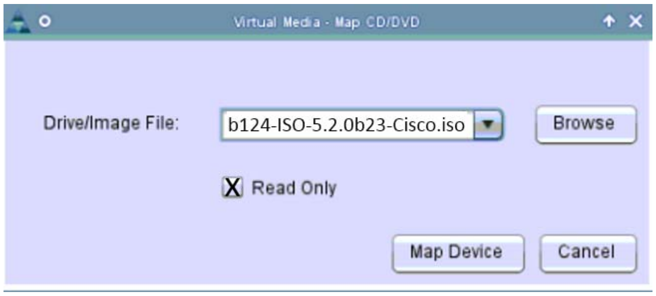

Drive/Image File: b124-ISO-5.2.0b23-Cisco.iso

☒ Read Only

Step 14 Check the **Read Only** check box and click **Map Device**.

Step 15 Click **Reset** under the menu bar to reboot the cache engine.

Figure 5-10 Reset the Cache Engine

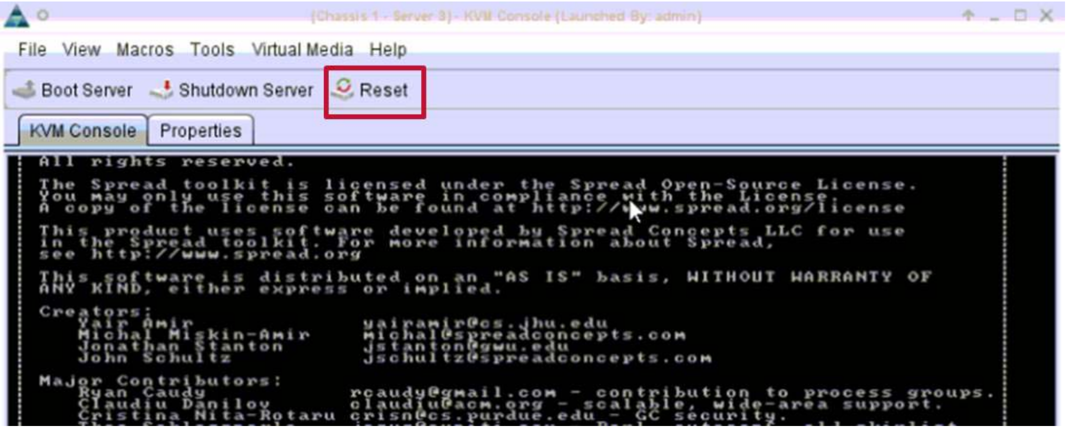
{Chassis 1 - Server 3} - KVM Console (Launched By: admin)
File View Macros Tools Virtual Media Help
Boot Server
Shutdown Server
Reset
KVM Console
Properties

```
All rights reserved.

The Spread toolkit is licensed under the Spread Open-Source License.
You may only use this software in compliance with the License.
A copy of the license can be found at http://www.spread.org/license

This product uses software developed by Spread Concepts LLC for use
in the Spread toolkit. For more information about Spread,
see http://www.spread.org

This software is distributed on an "AS IS" basis, WITHOUT WARRANTY OF
ANY KIND, either express or implied.

Creators:
    Yair Amir              yairamir@cs.jhu.edu
    Michal Miskin-Amir     michal@spreadconcepts.com
    Jonathan Stanton       jstanton@gwu.edu
    John Schultz           jschultz@spreadconcepts.com

Major Contributors:
    Ryan Caudy            rcaudy@gmail.com - contribution to process groups.
    Claudiu Danilov       claudiu@acm.org - scalable, wide-area support.
    Cristina Nita-Rotaru crisn@cs.purdue.edu - GC security.
```

Step 16 From the Reset Server Warning window that appears, click **OK** to reset the server.

Figure 5-11 Reset Server Warning Window

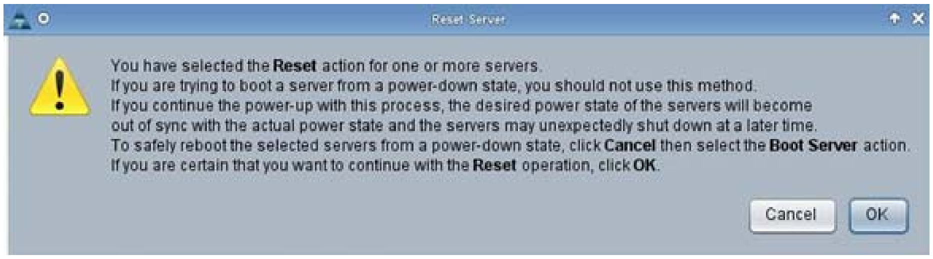

You have selected the **Reset** action for one or more servers.
If you are trying to boot a server from a power-down state, you should not use this method.
If you continue the power-up with this process, the desired power state of the servers will become
out of sync with the actual power state and the servers may unexpectedly shut down at a later time.
To safely reboot the selected servers from a power-down state, click **Cancel** then select the **Boot Server** action.
If you are certain that you want to continue with the **Reset** operation, click **OK**.

Step 17 From the Reset Server Service window that appears, choose **Power Cycle** and click **OK**.

Figure 5-12 Reset Server Service

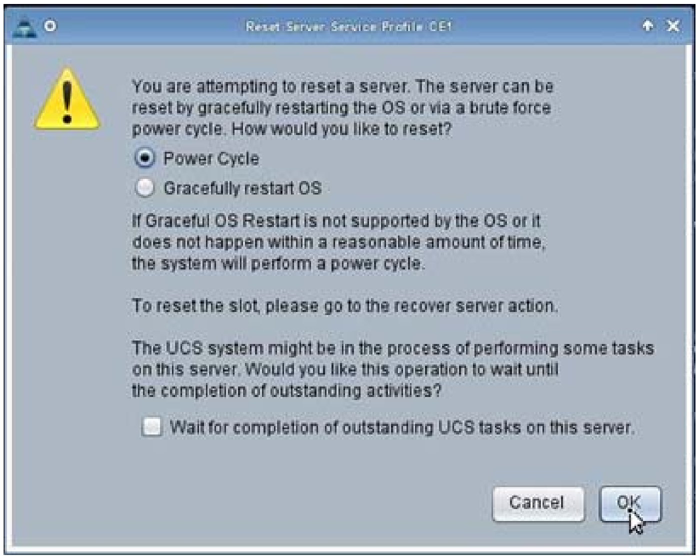

You are attempting to reset a server. The server can be reset by gracefully restarting the OS or via a brute force power cycle. How would you like to reset?

● Power Cycle
○ Gracefully restart OS

If Graceful OS Restart is not supported by the OS or it does not happen within a reasonable amount of time, the system will perform a power cycle.

To reset the slot, please go to the recover server action.

The UCS system might be in the process of performing some tasks on this server. Would you like this operation to wait until the completion of outstanding activities?

□ Wait for completion of outstanding UCS tasks on this server.

Step 18 From the KVM Console window you will see the VDS TC Cache Engine reboot. When you see the prompt to press F6 to enter the boot menu, click the KVM Console window to make sure it has focus and press **F6**. The Please Select Boot Device dialog box appears.

Note You may need to press **F6** several times to see the Please Select Boot Device dialog box.

Figure 5-13 System Reboot Window

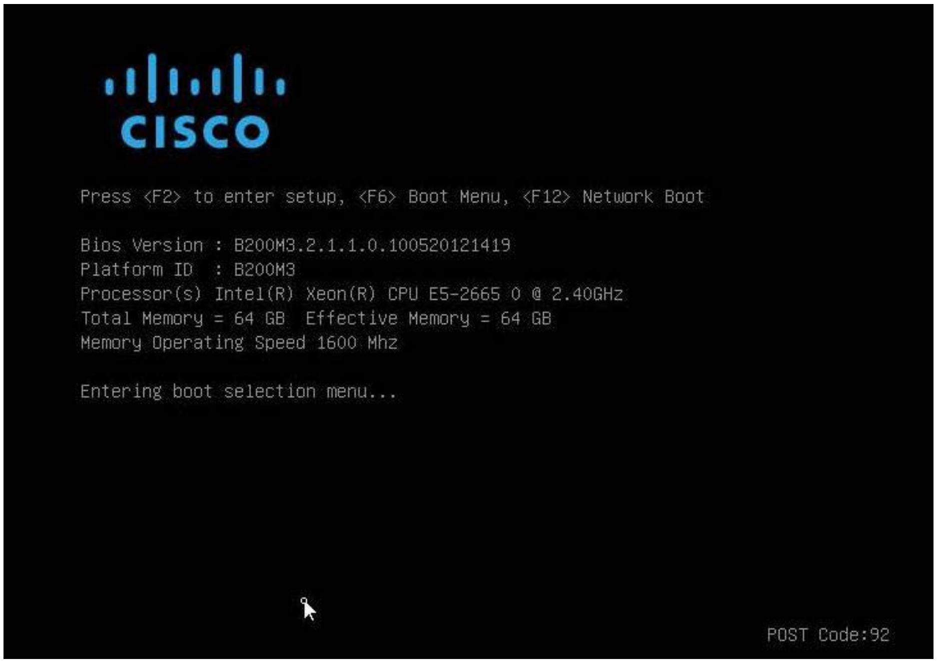

```
Press <F2> to enter setup, <F6> Boot Menu, <F12> Network Boot

Bios Version : B200M3.2.1.1.0.100520121419
Platform ID  : B200M3
Processor(s) Intel(R) Xeon(R) CPU E5-2665 0 @ 2.40GHz
Total Memory = 64 GB  Effective Memory = 64 GB
Memory Operating Speed 1600 Mhz

Entering boot selection menu...

                                              POST Code:92
```

Step 19 Using the up and down arrows, choose **Cisco vKVM-Mapped vDVD1.22**. The system will reboot from the virtual ISO image.

Step 20 When the VDS-TC Installer environment displays in the KVM console window, choose **Installer - No Serial**.

Figure 5-14 VDS TC Installer Window

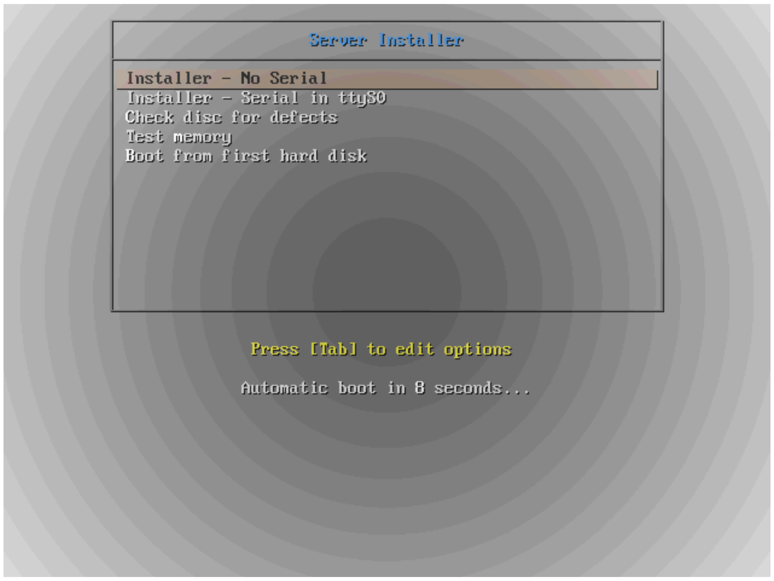

```
Server Installer

Installer - No Serial
Installer - Serial in ttyS0
Check disc for defects
Test memory
Boot from first hard disk

Press [Tab] to edit options

Automatic boot in 8 seconds...
```

Step 21 When the Welcome screen appears, enter **Y** to continue the installation process.

Figure 5-15 Welcome Screen

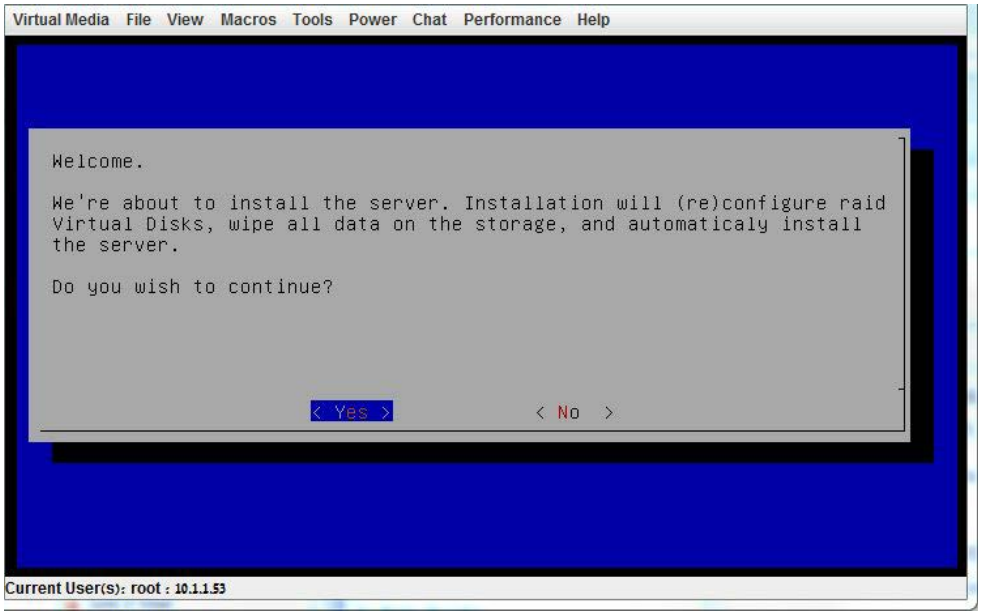

Virtual Media File View Macros Tools Power Chat Performance Help

```
Welcome.

We're about to install the server. Installation will (re)configure raid
Virtual Disks, wipe all data on the storage, and automaticaly install
the server.

Do you wish to continue?

< Yes >                 < No  >
```

Current User(s): root : 10.1.1.53

Step 22 When the Choose the Appliance Deployment window appears, choose **Cache Engine** and select **OK**.

Figure 5-16 Choose the Appliance Deployment Window

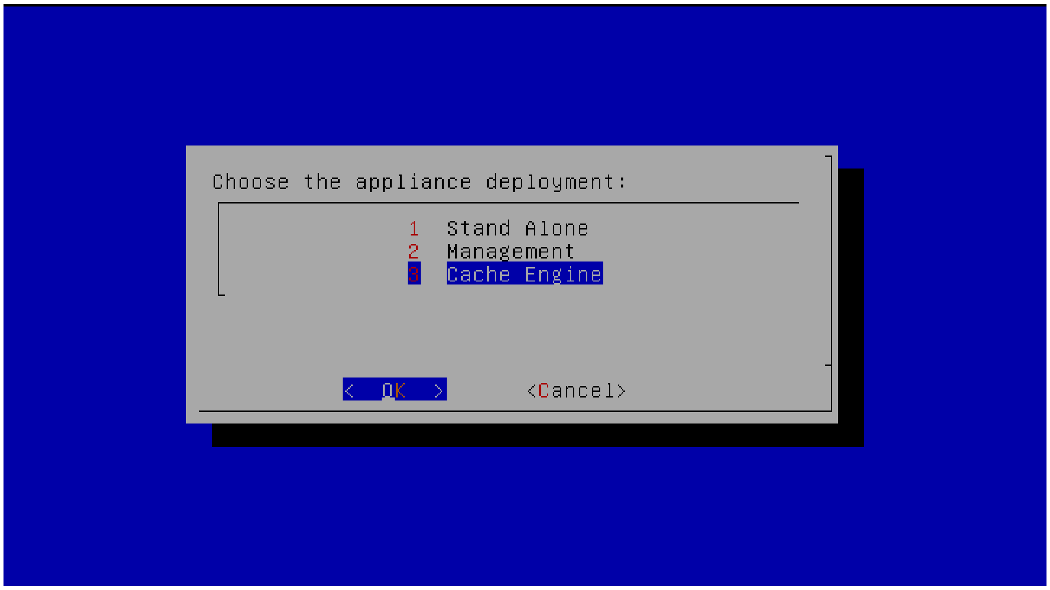

```
Choose the appliance deployment:
        1  Stand Alone
        2  Management
        3  Cache Engine
```

Step 23 When the "Enter Cache Engine Number" dialog box appears, enter the number that corresponds with the cache engine that you are replacing.

Figure 5-17 Enter Cache Engine Number

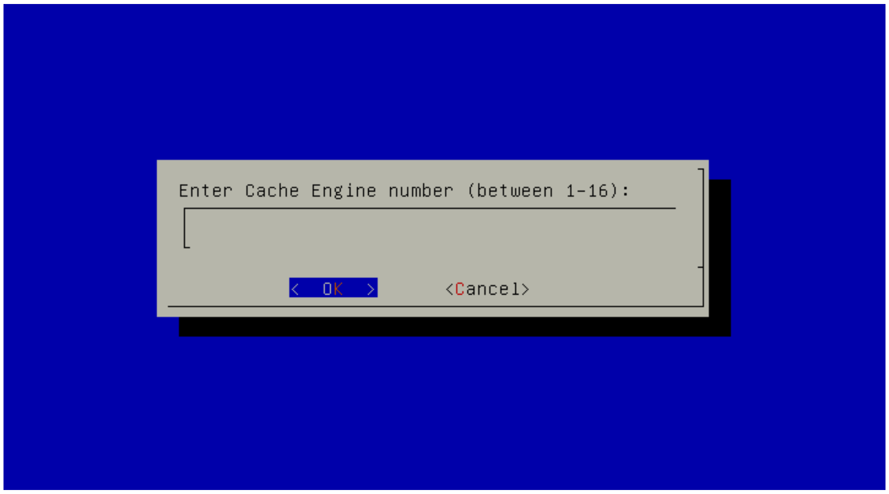

```
Enter Cache Engine number (between 1-16):
```

Step 24 Confirm the Cache Engine configuration. If everything is correct, choose **Yes**. If there is a mistake, choose **No** and repeat Step 20 through Step 23.

Figure 5-18 Configuration Confirmation

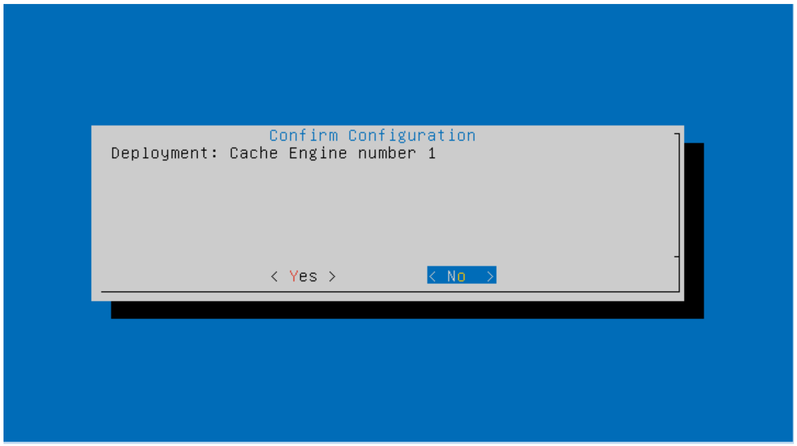

Confirm Configuration
Deployment: Cache Engine number 1

Step 25 From the KVM window you will see the installation begin. You must wait for this installation process to complete before you continue.

Note This installation should take about 40 minutes.

Figure 5-19 Cache Engine Installation Process

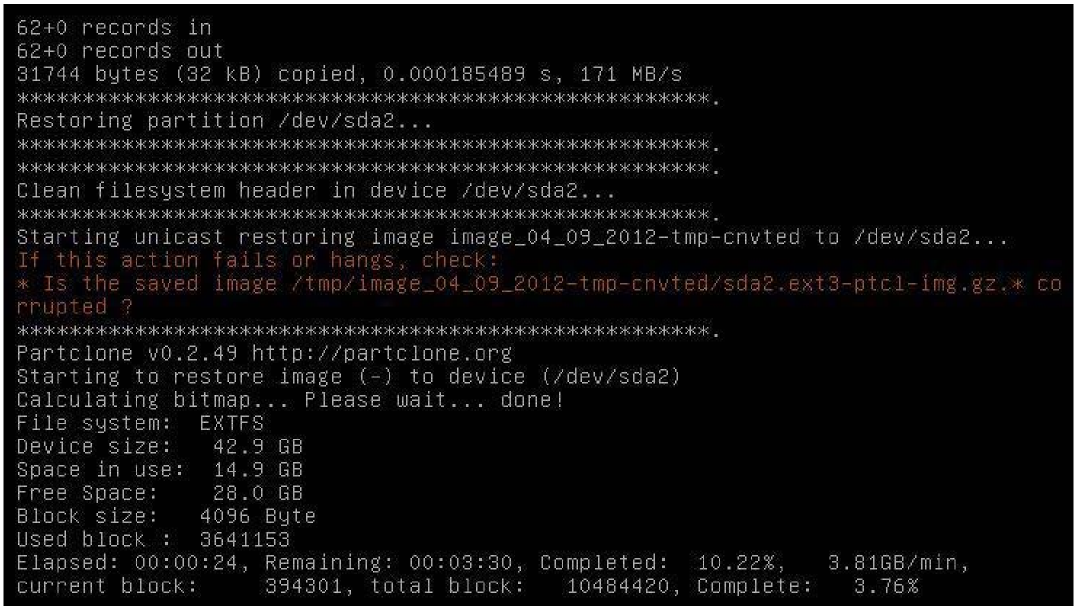

```
62+0 records in
62+0 records out
31744 bytes (32 kB) copied, 0.000185489 s, 171 MB/s
*****************************************************.
Restoring partition /dev/sda2...
*****************************************************.
*****************************************************.
Clean filesystem header in device /dev/sda2...
*****************************************************.
Starting unicast restoring image image_04_09_2012-tmp-cnvted to /dev/sda2...
If this action fails or hangs, check:
* Is the saved image /tmp/image_04_09_2012-tmp-cnvted/sda2.ext3-ptcl-img.gz.* co
rrupted ?
*****************************************************.
Partclone v0.2.49 http://partclone.org
Starting to restore image (-) to device (/dev/sda2)
Calculating bitmap... Please wait... done!
File system:  EXTFS
Device size:   42.9 GB
Space in use:  14.9 GB
Free Space:    28.0 GB
Block size:   4096 Byte
Used block :  3641153
Elapsed: 00:00:24, Remaining: 00:03:30, Completed:  10.22%,   3.81GB/min,
current block:     394301, total block:   10484420, Complete:   3.76%
```

Step 26 When the installation completes, you must unmap the ISO image file. When the “Installation completed. Press <Enter> to reboot.” message appears, press **Enter**.

Figure 5-20 Installation Completed Window

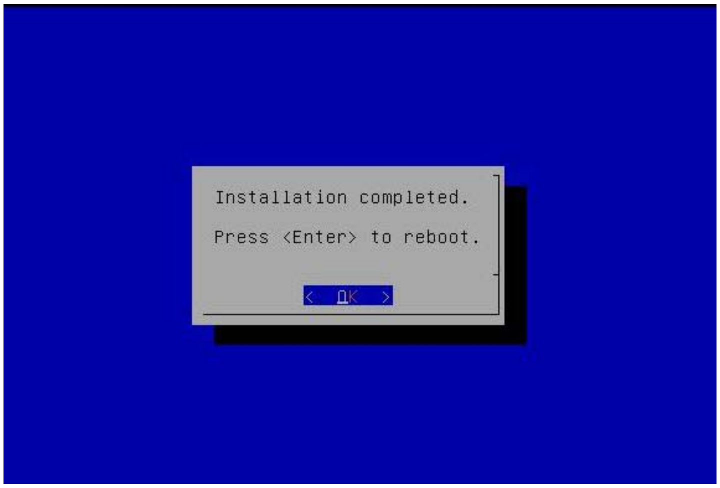

```
Installation completed.

Press <Enter> to reboot.

< OK >
```

Note Wait for the system to reboot and display the BIOS splash screen, as shown in Figure 5-13 "System Reboot Window", before proceeding to the next step.

Step 27 After the system reboots, from the KVM Console window, choose **Virtual Media** > **Activate Virtual Devices** to deactivate the virtual devices and unmap the image file.

Figure 5-21 Deactivate Virtual Device

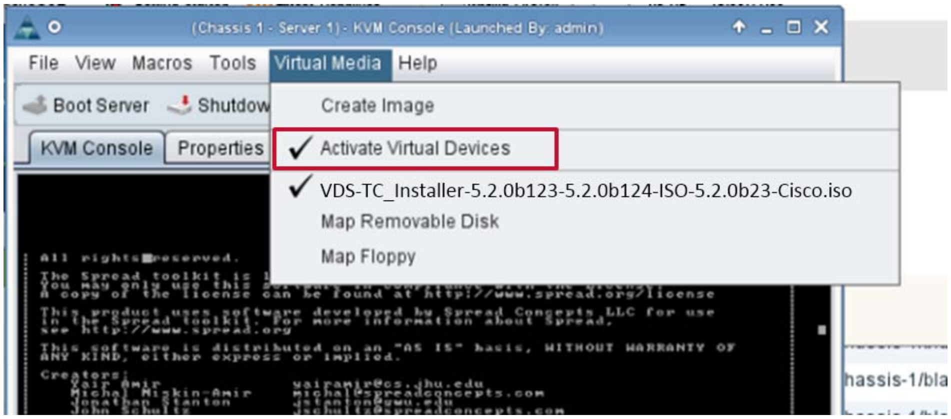
(Chassis 1 - Server 1) - KVM Console (Launched By: admin)
File
View
Macros
Tools
Virtual Media
Help
Boot Server
Shutdow
KVM Console
Properties
Create Image
✓ Activate Virtual Devices
✓ VDS-TC_Installer-5.2.0b123-5.2.0b124-ISO-5.2.0b23-Cisco.iso
Map Removable Disk
Map Floppy
hassis-1/bla

```
All rights reserved.
The Spread toolkit is l
You may only use this
A copy of the license can be found at http://www.spread.org/license
This product uses software developed by Spread Concepts LLC for use
in the Spread toolkit. For more information about Spread,
see http://www.spread.org
This software is distributed on an "AS IS" basis, WITHOUT WARRANTY OF
ANY KIND, either express or implied.
Creators:
    Yair Amir              yairamir@cs.jhu.edu
    Michal Miskin-Amir     michal@spreadconcepts.com
    Jonathan Stanton       jstanton@gwu.edu
    John Schultz           jschultz@spreadconcepts.com
```

Step 28 From the Close Virtual Media dialog box, click **Yes** to confirm the unmapping of the image file. This unmounts the ISO image.

Figure 5-22 Close Virtual Media

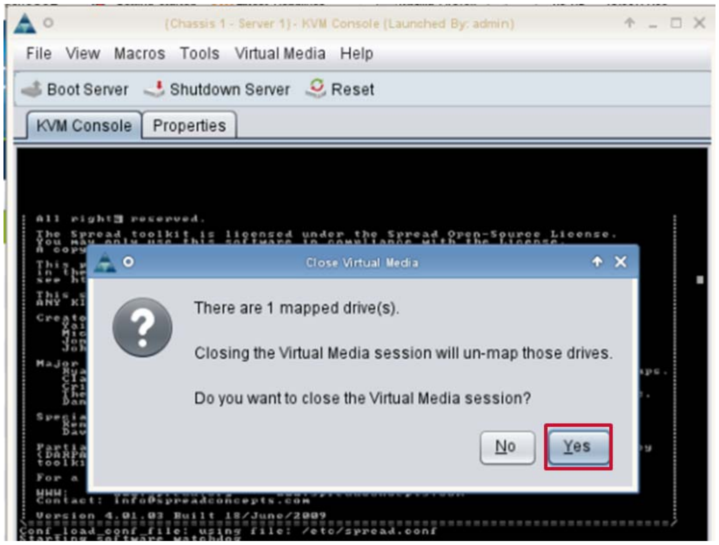
(Chassis 1 - Server 1) - KVM Console (Launched By: admin)
File View Macros Tools Virtual Media Help
Boot Server
Shutdown Server
Reset
KVM Console
Properties
Close Virtual Media

There are 1 mapped drive(s).

Closing the Virtual Media session will un-map those drives.

Do you want to close the Virtual Media session?

```
All rights reserved.
The Spread toolkit is licensed under the Spread Open-Source License.
Contact: info@spreadconcepts.com
Version 4.01.03 Built 18/June/2009
Conf_load_conf_file: using file: /etc/spread.conf
```

Step 29 Click **Reset** under the menu bar to reboot the cache engine.

Figure 5-23 Reset the Cache Engine

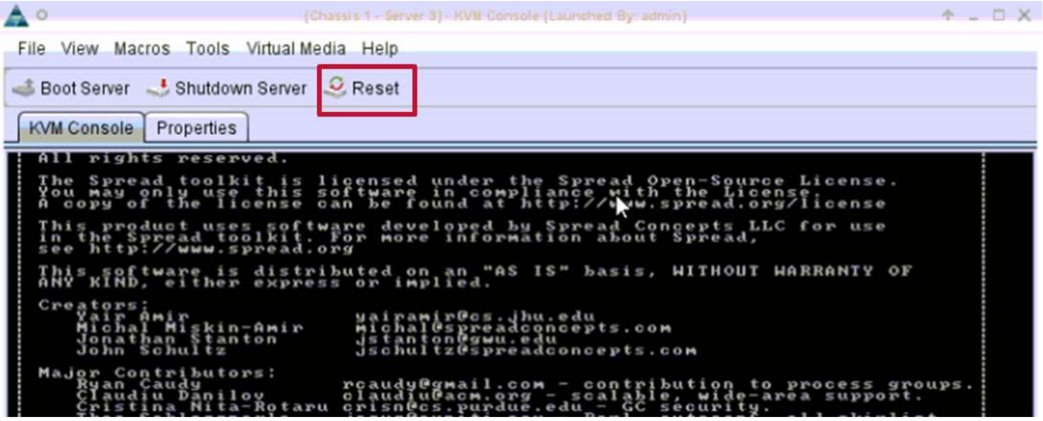
(Chassis 1 - Server 3) - KVM Console (Launched By: admin)
File View Macros Tools Virtual Media Help
Boot Server
Shutdown Server
Reset
KVM Console
Properties

```
All rights reserved.

The Spread toolkit is licensed under the Spread Open-Source License.
You may only use this software in compliance with the License.
A copy of the license can be found at http://www.spread.org/license

This product uses software developed by Spread Concepts LLC for use
in the Spread toolkit. For more information about Spread,
see http://www.spread.org

This software is distributed on an "AS IS" basis, WITHOUT WARRANTY OF
ANY KIND, either express or implied.

Creators:
    Yair Amir                yairamir@cs.jhu.edu
    Michal Miskin-Amir       michal@spreadconcepts.com
    Jonathan Stanton         jstanton@gwu.edu
    John Schultz             jschultz@spreadconcepts.com

Major Contributors:
    Ryan Caudy              rcaudy@gmail.com - contribution to process groups.
    Claudiu Danilov         claudiu@acm.org - scalable, wide-area support.
    Cristina Nita-Rotaru crisn@cs.purdue.edu - GC security.
```

Step 30 From the Reset Server Warning window that appears, click **OK** to reset the server.

Figure 5-24 Reset Server Warning Window

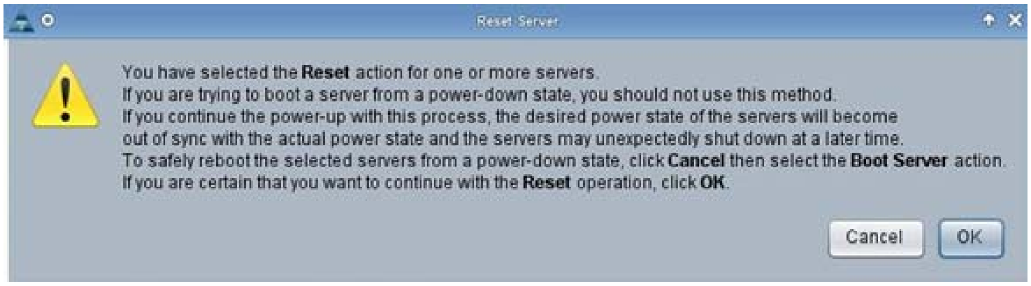

Reset Server

You have selected the **Reset** action for one or more servers.
If you are trying to boot a server from a power-down state, you should not use this method.
If you continue the power-up with this process, the desired power state of the servers will become out of sync with the actual power state and the servers may unexpectedly shut down at a later time.
To safely reboot the selected servers from a power-down state, click **Cancel** then select the **Boot Server** action.
If you are certain that you want to continue with the **Reset** operation, click **OK**.

Step 31 From the Reset Server Service window that appears, choose **Power Cycle** and click **OK**.

Figure 5-25 Reset Server Service

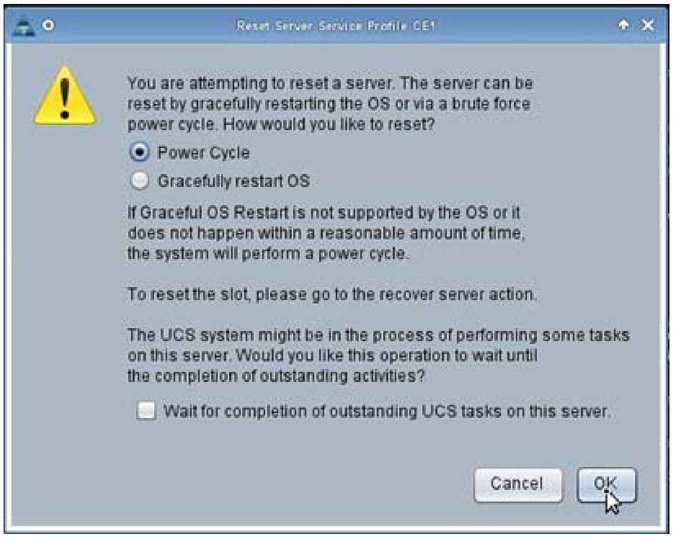

Reset Server Service Profile CE1

You are attempting to reset a server. The server can be reset by gracefully restarting the OS or via a brute force power cycle. How would you like to reset?

● Power Cycle
○ Gracefully restart OS

If Graceful OS Restart is not supported by the OS or it does not happen within a reasonable amount of time, the system will perform a power cycle.

To reset the slot, please go to the recover server action.

The UCS system might be in the process of performing some tasks on this server. Would you like this operation to wait until the completion of outstanding activities?

□ Wait for completion of outstanding UCS tasks on this server.

Step 32 After the cache engine reboots, from the KVM Console window on the KVM tab, log in using the **root** user and the password that was provided by Cisco.

Step 33 Enter the command **cd /opt/pang/utilities/CE/Multipath_Driver**

Step 34 Enter the command **tar -zxvf rdac-LINUX-09.03.0C05.0652-source.tar.gz**

Step 35 Enter the command **cd linuxrdac-09.03.0C05.0652**

Step 36 Enter the command **make clean**

Step 37 Enter the command **make**

Step 38 Enter the command **make install**

Step 39 If you are prompted about whether the new MPP driver should manage the iSCSI storages, answer yes to continue.

Step 40 Enter the command **reboot -f** to reboot the cache engine.

Step 41 After the cache engine reboots, log back into the KVM console as the root user, using the password that was provided by Cisco.

Step 42 To verify that the installation of the multipath driver was successful, enter the command **chkconfig mpp**. If the installation was successful, this command should return the following output:

```
mpp on
```

Step 43 From the KVM console connection to the Cache Engine, enter the following commands to temporarily update the network interface:

a. **vconfig add eth0 50**

b. **ifconfig eth0 10.11.12.*X*/24**
where *X* is the number of the cache engine that you are replacing plus 1. For example, if you are replacing cache engine 3, you would enter **ifconfig eth0 10.11.12.4/24**.

Note Please pay close attention to the IP address that you are assigning. If it is not entered correctly, it can cause a problem replacing the cache engine.

Step 44 From the KVM console connection to the Cache Engine, ping the IP address of the management server. If this ping fails, repeat Step 43. If it fails again, contact your Cisco Support Engineer.

Updating the SSH Keys and Connecting the Cache Engine to the Storage Enclosure

After completing the "Running the Post Installation Scripts" section in the "Cache Engine Servers Installation" chapter, you must update the grid SSH keys on the VDS TC management server and configure the new cache engine to connect to the storage enclosure. All of these steps are performed on the VDS TC management server.

Follow this procedure to perform these steps on the VDS TC management server:

Note The following steps are performed on the VDS TC management server.

Procedure

Step 1 On the PBR routers, stop redirecting traffic to the VDS TC solution.

Step 2 Close any open VDS TC Manager windows and ensure that no other administrators are connected to the VDS TC Manager before proceeding.

Step 3 Using SSH software, such as Putty, open an SSH connection to the IP address of the VDS TC management server.

Step 4 Log into the system using the username **padmin** and the password provided by Cisco.

Step 5 Enter the command **su admin** to log into the CLI of the VDS TC management server. The default password is the serial number of the VDS TC management server.

Step 6 From the VDS TC prompt, enter the **enable** command. When prompted, enter the Enable mode password and press **Enter**. You are now logged into Enable mode and the Enable prompt, console#, should appear.

Step 7 From the Enable mode prompt, enter the command **oper service stop** to stop the caching service.

Step 8 Wait a few minutes and then enter the command **show status**. Check to see if the Device State shows "stopped." For example:

```
console# show status
Operational state Device state Administrative state
disabledstoppedunlocked
```

Caution Do not proceed to the next step until the Device State shows “stopped.” You may need to repeat the **show status** command several times before you see this status. Wait several minutes between executions of the **show status** command.

Step 9 Enter the **exit** command to exit Enabled mode.

Step 10 Enter the **exit** command to exit the VDS TC CLI.

Step 11 Enter the command **su root** to change to the root user. Enter the password for the root user when prompted. This password is provided by Cisco.

Step 12 Enter the command **cd /opt/pang/useful/**.

Step 13 To reset the SSH keys, enter the command **./replace_server_keys.sh -servers** #, where # is the total number of Blade Server cache engines in the VDS TC installation. For example, if you are replacing a cache engine in a VDS TC solution that has 16 cache engines, you would enter the command **./replace_server_keys.sh -servers 16**.

Step 14 Enter the command **cd /opt/pang/useful/configure_storages** to change to the storage configuration script directory.

Step 15 Enter the command **./configure_storages.py -s** <*no_of_storage_enclosures*> **-b** <*no_of_blades*> **-c** <*the_replaced_CE-ID_number*>. For example, to replace CE-2 in a VDS TC Blade Server installation that has two cache engines, and one storage enclosure, enter the command **./configure_storages.py -s 1 -b 2 -c 2**.

Note The configure storages script will format the data disks. To configure the storages without formatting the disks enter the command **./configure_storages.py -s** <*no_of_storage_enclosures*> **-b** <*no_of_blades_in_grid*> **-c** <*the_replaced_CE-ID_number*> **-x**.

Step 16 When the warning message about the disk format appears, enter **y** to continue the storages configuration.

```
*********** WARNING! **********
All storages are about to be reconfigured. All data on the data disks will
be lost.
Do you want to continue? (y/n)y
```

Step 17 Follow the on screen instructions. If prompted, enter the password for the root user that was provided by Cisco.

Note You may be prompted to enter the password several times.

When the script is finished running, you should see the word “COMPLETE”. After running this script, the new VDS TC cache engine will be connected to the storage enclosures.

Step 18 Use SSH to connect to each CE and enter the command **iscsiadm -m session** to check the connections:

```
ssh ce-<cache engine number>
iscsiadm -m session
```

If you are using 1Gb connections to the storage enclosures, your output should be similar to the following:

Note The following output will look different if you are using 10Gb connections to the storage enclosures.

```
ƒÞ 8 * sessions per storage, for example for 1 storage:
tcp: [17] 10.11.14.100:3260,1 iqn.1992-
01.com.Avago:2365.60080e50002f4e82000000005063c74a
tcp: [18] 10.11.14.101:3260,2 iqn.1992-
01.com.Avago:2365.60080e50002f4e82000000005063c74a
tcp: [19] 10.11.16.100:3260,1 iqn.1992-
01.com.Avago:2365.60080e50002f4e82000000005063c74a
tcp: [20] 10.11.16.101:3260,2 iqn.1992-
01.com.Avago:2365.60080e50002f4e82000000005063c74a
tcp: [21] 10.11.15.100:3260,1 iqn.1992-
01.com.Avago:2365.60080e50002f4e82000000005063c74a
tcp: [22] 10.11.15.101:3260,2 iqn.1992-
01.com.Avago:2365.60080e50002f4e82000000005063c74a
tcp: [23] 10.11.17.101:3260,2 iqn.1992-
01.com.Avago:2365.60080e50002f4e82000000005063c74a
tcp: [24] 10.11.17.100:3260,1 iqn.1992-
01.com.Avago:2365.60080e50002f4e82000000005063c74a
```

Upload Network Updater File to the Cache Engine

Step 1 Using SSH software, such as Putty, open an SSH connection to the management IP address of the VDS TC management server.

Step 2 Log into the system using the username **padmin** and the password that was provided by Cisco.

Step 3 Enter the command **su root** and enter the password that was provided by Cisco for the root account.

Step 4 Enter the command **cd /tmp**.

Step 5 Enter the command **ls Cisco_network_updater_v_1.8.tgz** to confirm that the Cisco_network_updater_v_1.8.tgz file is in the tftpboot folder. If it is not, refer to the Copying the Files to the VDS TC Management Server section to upload the file.

Step 6 Enter the command **scp Cisco_network_updater_v_1.8.tgz ce-#:/tmp** where # is the number of the cache engine that you are replacing. For example, if you are replacing CE-3, enter the command **scp Cisco_network_updater_v_1.8.tgz ce-3:/tmp**

Update the Kernel

Step 1 From an SSH connection to the VDS TC server, log in as root with the password provided by Cisco and enter the command **cd /tmp**.

Step 2 To copy the updated kernel installation package file to the new cache engine, enter the command s**cp linux-2.6.27.19-5-llpf_10--KERNEL_LOAD_1414071783-r44244.tar root@ce-#:/tmp**, where **#** is the number of the cache engine that you replaced. For example, if you replaced cache engine 1, enter the following:

```
scp linux-2.6.27.19-5-llpf_10--KERNEL_LOAD_1414071783-r44244.tar root@ce-1:/tmp
```

Step 3 From the VDS TC management server, enter the command **ssh ce-*X***, where ***X*** is the number of the cache engine that you replaced. For example, if you replaced CE-3, enter **ssh ce-3**.

Step 4 Enter the command **cd /tmp** to change to the /tmp folder.

Step 5 Enter the command **tar -xvf linux-2.6.27.19-5-llpf_10--KERNEL_LOAD_1414071783-r44244.tar** to extract the kernel files.

Step 6 Enter the command **cd linux-2.6.27.19-5-llpf_10--KERNEL_LOAD_1414071783** to change to the new folder that was created when extracting the files in Step 5.

Step 7 Enter the command **./linux-2.6.27.19-5-llpf_10--INSTALL_PHASE-1.sh** to run the first phase of the kernel installation.

Step 8 Each time you see the following prompt, press **Enter**:

```
Press [Enter] key...
```

Note You will see this prompt several times.

Step 9 When you see the following prompt, press **Enter**. This will cause the system to reboot for the first time. When the system reboots, the SSH connection to the VDS TC server will close.

```
Going to REBOOT, Press [Enter] key...
```

Step 10 After the cache engine reboots, from the VDS TC management server, enter the command **ssh ce-*X*** to reconnect to the cache engine, where ***X*** is the number of the cache engine to which are reconnecting.

Step 11 Enter the command **cd /tmp**.

Step 12 Enter the command **cd linux-2.6.27.19-5-llpf_10--KERNEL_LOAD_1414071783**.

Step 13 Enter the command **./linux-2.6.27.19-5-llpf_10--INSTALL_PHASE-2_Cisco.sh** to run the second phase of the kernel installation.

Step 14 Each time you see the following prompt, press **Enter**:

```
Press [Enter] key...
```

Step 15 When you see the following prompt, press **Enter**. This will cause the system to reboot. When the system reboots, the SSH connection to the VDS TC server will close.

```
Going to REBOOT, Press [Enter] key...
```

Step 16 After the cache engine reboots, from the VDS TC management server, enter the command **ssh ce-*X*** to reconnect to the cache engine, where ***X*** is the number of the cache engine to which are reconnecting.

Step 17 To confirm the kernel update, enter the command **uname –a** and verify that the date Tue Aug 19 16:27:56 GMT 2014 appears in the output, as shown in the following example.

```
Linux ce-1 2.6.27.19-llpf_10-5-default #26 SMP Tue Aug 19 16:27:56 GMT 2014 x86_64 x86_64
x86_64 GNU/Linux
```

Step 18 Enter **logout** to close the SSH connection to the cache engine.

Updating the NetApp Storage Enclosure

Note This process is only for the NetApp storage enclosure. If your system uses an IBM storage enclosure, refer to the Updating the IBM Storage Enclosure section.

After you have configured the VDS TC cache engine to connect with the storage enclosure, you must update the host port identifiers on the NetApp storage enclosure for the new cache engine. Follow this procedure to perform this task:

Procedure

Step 1 From the SSH connection to the VDS TC management server, enter the following commands to start the VNC server:

a. **rm /tmp/.X1-lock** (This step is optional. It is only required if there is a leftover lock file from a previous vncserver instance.)

b. **rm /tmp/.X11-unix/X1** (This step is optional. It is only required if there is a leftover lock file from a previous vncserver instance.)

c. **vncserver :1**

Step 2 Use a VNC client to connect from your local computer to the VDS TC management server.

Figure 5-26 VNC Viewer

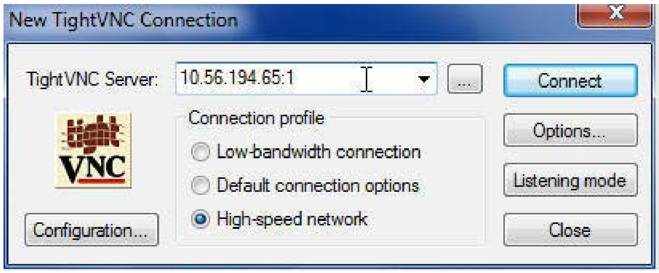

New TightVNC Connection

TightVNC Server: 10.56.194.65:1

Connection profile
○ Low-bandwidth connection
○ Default connection options
● High-speed network

Step 3 From the VNC console window, enter the command **cd /opt/SMgr/client** to change folders.

Step 4 Enter the command **./SMclient** to start the SANtricity ™ ES Storage Manager software.

Note You must perform the following steps on all of the storage subsystems.

Step 5 Right-click the storage subsystem that you are configuring and choose **Manage Storage Array**. Click **No** in any pop-up windows that appear.

Figure 5-27 SANtricity Storage Manager

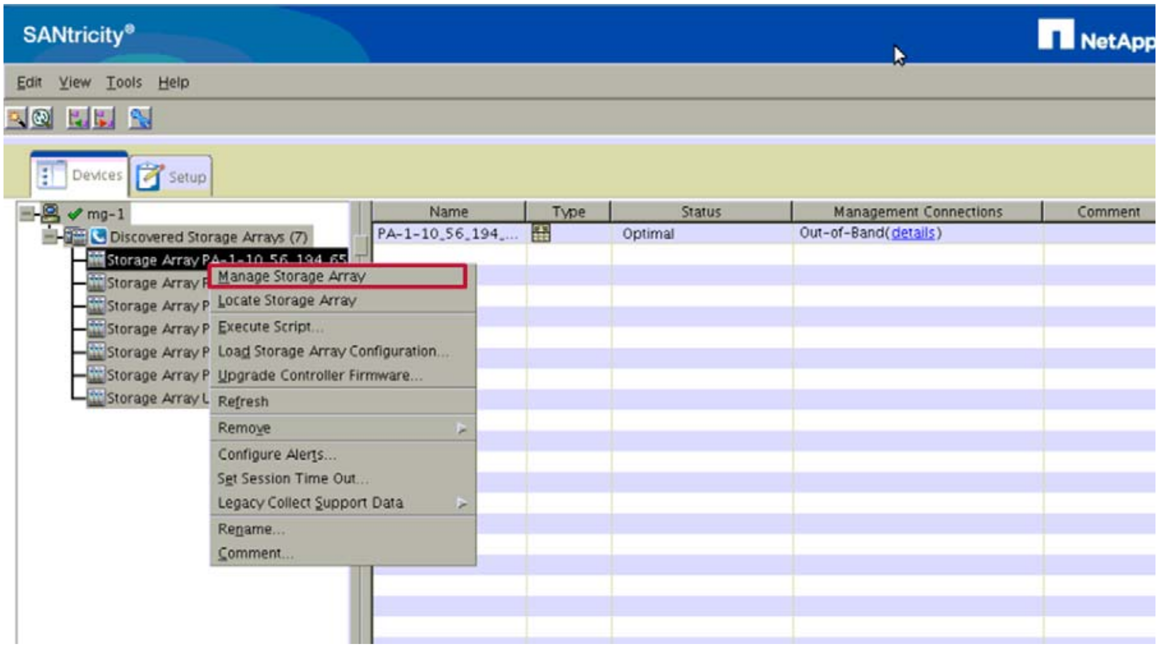

Name	Type	Status	Management Connections	Comment
PA-1-10.56.194....	▦	Optimal	Out-of-Band(details)	

Step 6 In the Subsystem Management window that appears, click the **Host Mappings** tab.

Step 7 In the navigation pane on the left, expand the **Host Group** folder.

Step 8 Right-click the cache engine that was replaced, for example Host CE-1, and choose **Manage Host Port Identifiers** from the pop-up menu.

Figure 5-28 Manage Host Port Identifiers

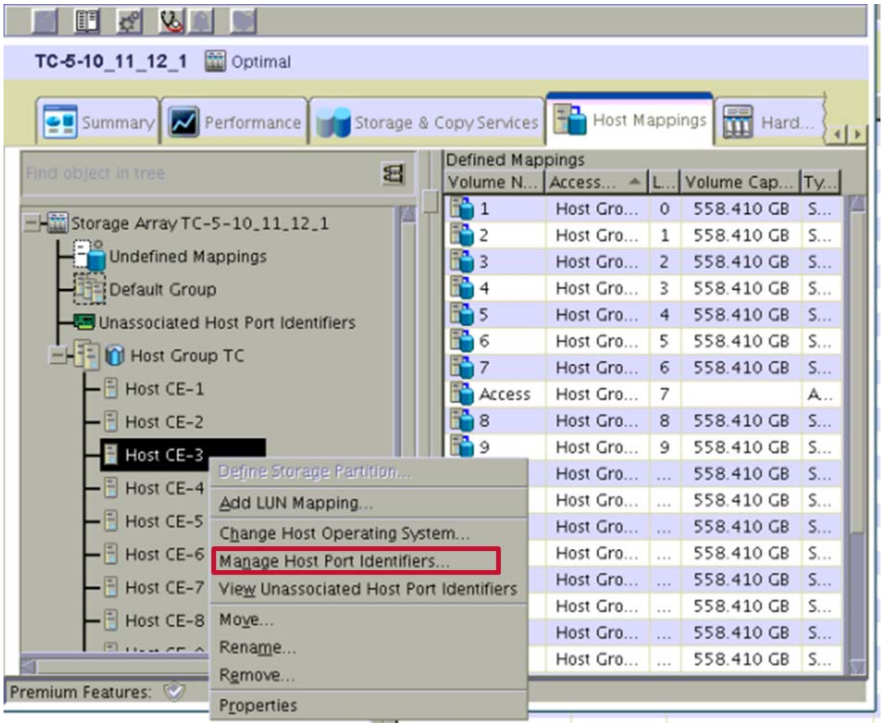

Defined Mappings

Volume N...	Access...	L...	Volume Cap...	Ty...
1	Host Gro...	0	558.410 GB	S...
2	Host Gro...	1	558.410 GB	S...
3	Host Gro...	2	558.410 GB	S...
4	Host Gro...	3	558.410 GB	S...
5	Host Gro...	4	558.410 GB	S...
6	Host Gro...	5	558.410 GB	S...
7	Host Gro...	6	558.410 GB	S...
Access	Host Gro...	7		A...
8	Host Gro...	8	558.410 GB	S...
9	Host Gro...	9	558.410 GB	S...
	Host Gro...	...	558.410 GB	S...
	Host Gro...	...	558.410 GB	S...
	Host Gro...	...	558.410 GB	S...
	Host Gro...	...	558.410 GB	S...
	Host Gro...	...	558.410 GB	S...
	Host Gro...	...	558.410 GB	S...
	Host Gro...	...	558.410 GB	S...
	Host Gro...	...	558.410 GB	S...

Step 9 From the Manage Host Port Identifiers window that appears, in the Host Port Identifier Information table, ensure that the old iqn is selected and click **Replace**.

Figure 5-29 Replace IQN

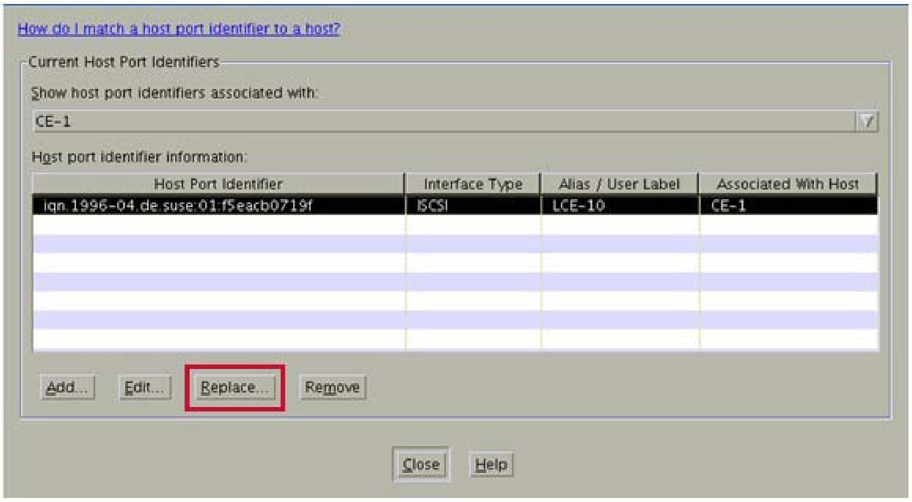
How do I match a host port identifier to a host?
Current Host Port Identifiers
Add...
Edit...
Replace...
Remove
Close
Help

Show host port identifiers associated with: CE-1

Host port identifier information:

Host Port Identifier	Interface Type	Alias / User Label	Associated With Host
iqn.1996-04.de.suse:01:f5eacb0719f	iSCSI	LCE-10	CE-1

Step 10 From the Replace Host Port Identifier window that appears, configure the following:

- From the Choose a Host Interface Type drop-down list, choose **ISCSI**.
- Click the **Replace by Selecting a Known Unassociated Host Port Identifier** radio button.

- From the Known Unassociated Host Port Identifier drop-down list, choose the iqn for the *newly* installed cache engine. (It should be the only one in the list.)

Step 11 Click **Replace**.

Figure 5-30 Replace Host Port Identifier

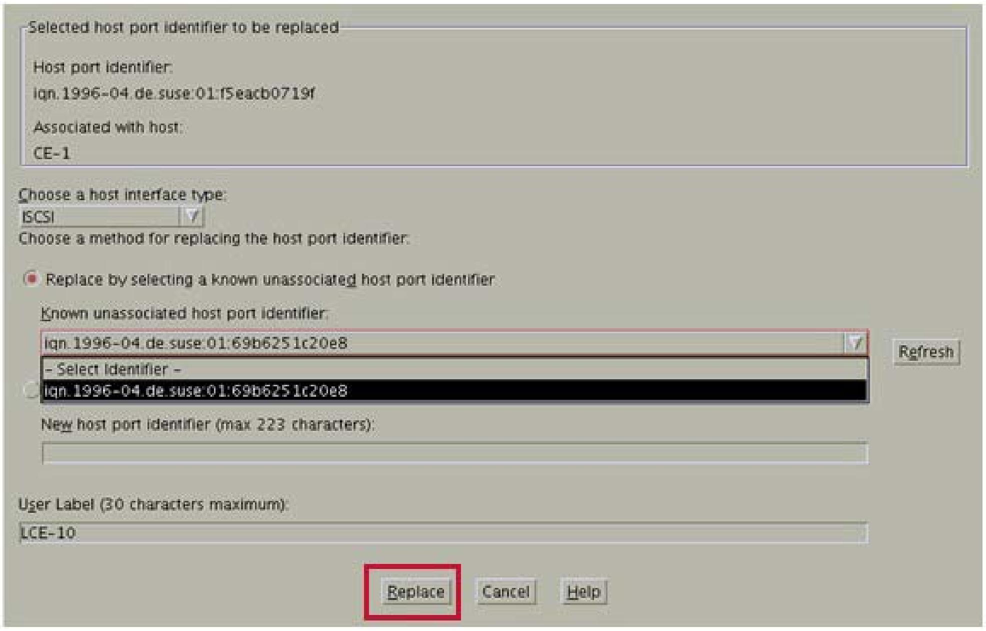

Selected host port identifier to be replaced

Host port identifier:
iqn.1996-04.de.suse:01:f5eacb0719f

Associated with host:
CE-1

Choose a host interface type:
ISCSI

Choose a method for replacing the host port identifier:

● Replace by selecting a known unassociated host port identifier

Known unassociated host port identifier:
iqn.1996-04.de.suse:01:69b6251c20e8
- Select Identifier -
iqn.1996-04.de.suse:01:69b6251c20e8

○ New host port identifier (max 223 characters):

User Label (30 characters maximum):
LCE-10

Step 12 From the Manage Host Port Identifiers window, click **Close**.

Step 13 Repeat Step 5 through Step 12 on any remaining storage subsystems.

Note Do not continue until you have completed Step 13.

Step 14 Close the Subsystem Management window and exit the SANtricity Storage Manager.

Step 15 Close the VNC client.

Updating the IBM Storage Enclosure

Note This process is only for the IBM storage enclosure. If your system uses the NetApp E2724 storage enclosure, refer to the Updating the NetApp Storage Enclosure section.

After you have configured the VDS TC cache engine to connect with the storage enclosure, you must update the host port identifiers on the IBM storage enclosure for the new cache engine. Follow this procedure to perform this task:

Procedure

Step 1 From the SSH connection to the VDS TC management server, enter the following commands to start the VNC server:

a. **rm /tmp/.X1-lock** (This step is optional. It is only required if there is a leftover lock file from a previous vncserver instance.)

b. **rm /tmp/.X11-unix/X1** (This step is optional. It is only required if there is a leftover lock file from a previous vncserver instance.)

c. **vncserver :1**

Step 2 Use a VNC client to connect from your local computer to the VDS TC management server.

Figure 5-31 VNC Viewer

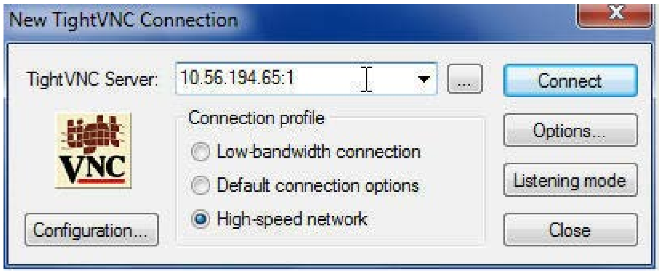

Step 3 From the VNC console window, enter the command **/opt/IBM_DS/client/SMclient** to start the IBM Storage Manager.

Note You must perform the following steps on all of the storage subsystems.

Step 4 Right-click the storage subsystem that you are configuring and choose **Manage Storage Subsystem**. Click **No** in any pop-up windows that appear.

Figure 5-32 IBM DS Storage Manager

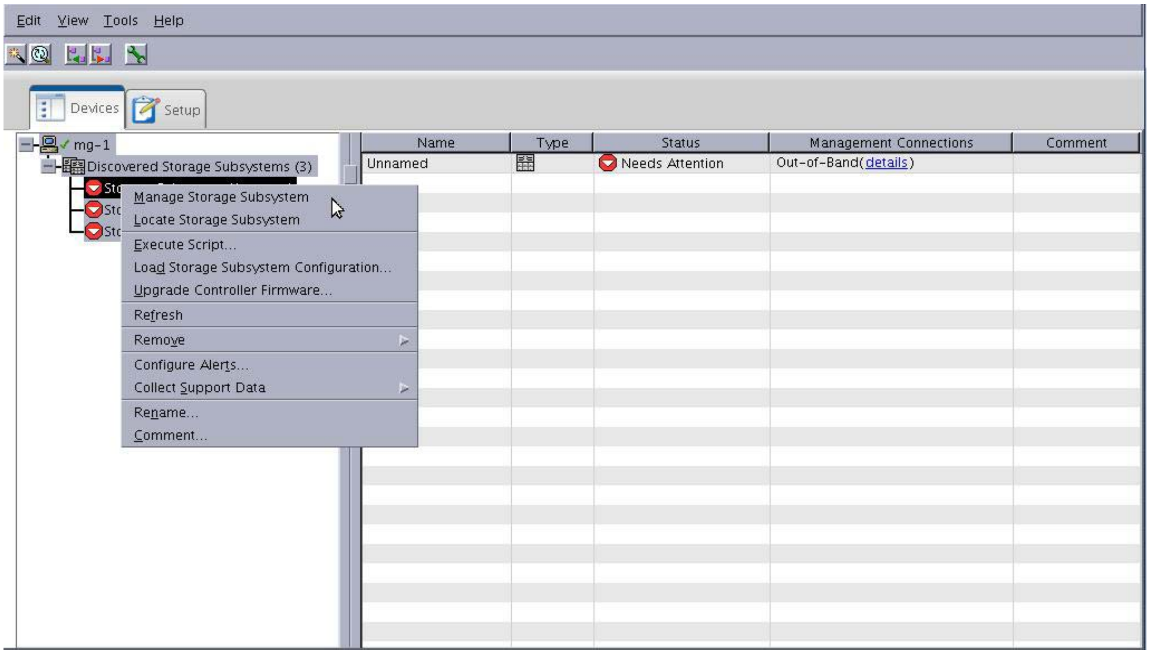

Name	Type	Status	Management Connections	Comment
Unnamed	▦	Needs Attention	Out-of-Band(details)	

Step 5 In the Subsystem Management window that appears, click the **Host Mappings** tab.

Step 6 In the navigation pane on the left expand the **Host Group** folder.

Step 7 Right-click the cache engine that was replaced, for example Host CE-2, and choose **Manage Host Port Identifiers** from the pop-up menu.

Figure 5-33 Manage Host Port Identifiers

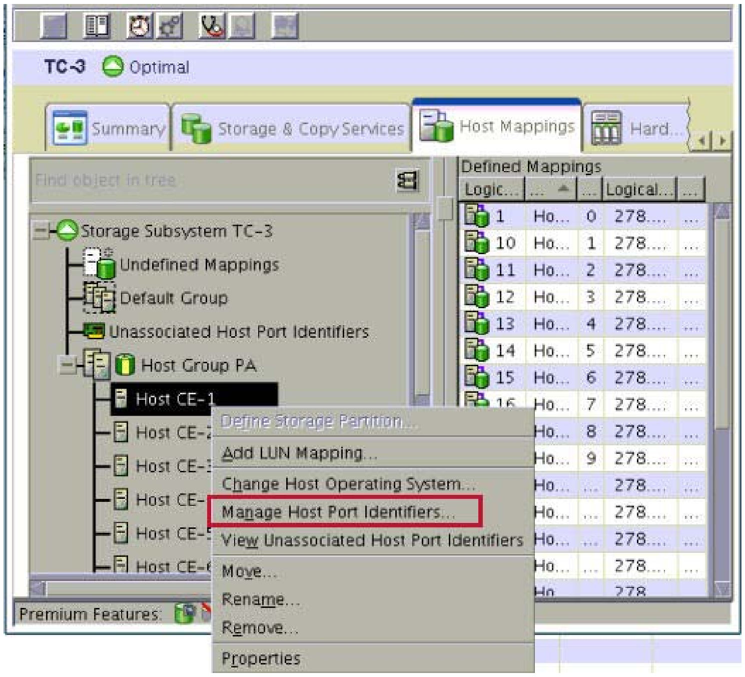

Logic...	...	...	Logical...	...
1	Ho...	0	278....	...
10	Ho...	1	278....	...
11	Ho...	2	278....	...
12	Ho...	3	278....	...
13	Ho...	4	278....	...
14	Ho...	5	278....	...
15	Ho...	6	278....	...
16	Ho...	7	278....	...
	Ho...	8	278....	...
	Ho...	9	278....	...
	Ho...	...	278....	...
	Ho...	...	278....	...
	Ho...	...	278....	...
	Ho...	...	278....	...
	Ho		278	

Step 8 From the Manage Host Port Identifiers window that appears, in the Host Port Identifier Information table, ensure that the iqn is selected and click **Replace**.

Figure 5-34 Replace IQN

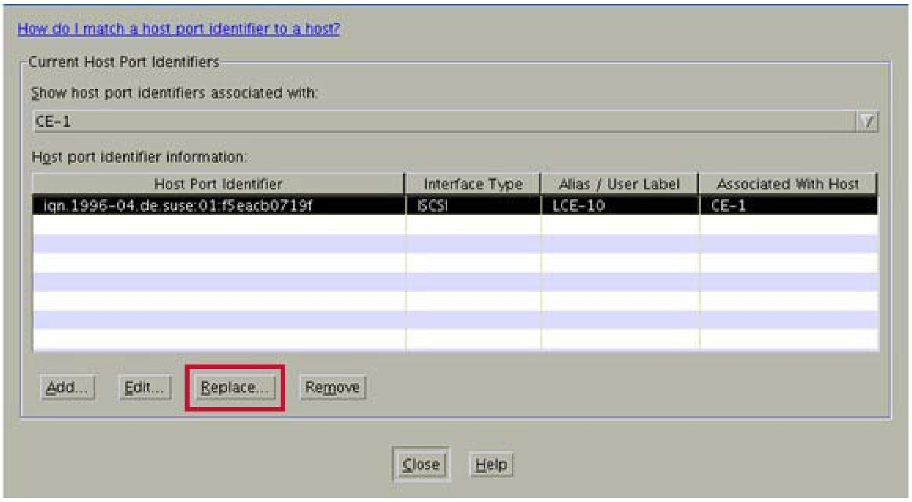

How do I match a host port identifier to a host?

Current Host Port Identifiers

Show host port identifiers associated with:

CE-1

Host port identifier information:

Host Port Identifier	Interface Type	Alias / User Label	Associated With Host
iqn.1996-04.de.suse:01:f5eacb0719f	ISCSI	LCE-10	CE-1

Step 9 From the Replace Host Port Identifier window that appears, configure the following:

- From the Choose a Host Interface Type drop-down list, choose **ISCSI**.
- Click the **Replace by Selecting a Known Unassociated Host Port Identifier** radio button.
- From the Known Unassociated Host Port Identifier drop-down list, choose the iqn for the newly installed cache engine. (It should be the only one in the list.)

Step 10 Click **Replace**.

Figure 5-35 Replace Host Port Identifier

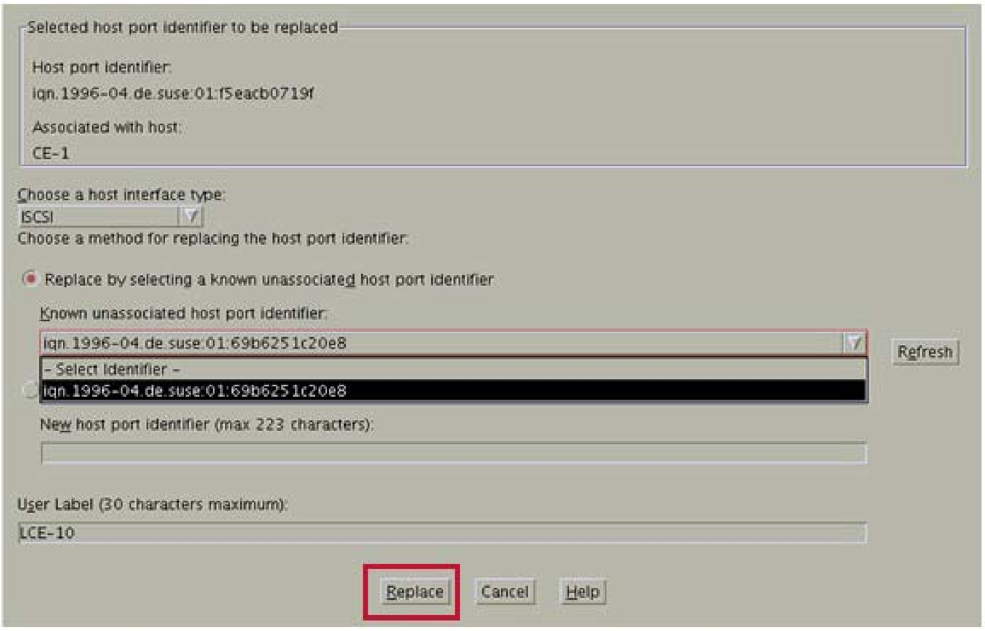

Selected host port identifier to be replaced

Host port identifier:
iqn.1996-04.de.suse:01:f5eacb0719f

Associated with host:
CE-1

Choose a host interface type:
ISCSI

Choose a method for replacing the host port identifier:

● Replace by selecting a known unassociated host port identifier

Known unassociated host port identifier:
iqn.1996-04.de.suse:01:69b6251c20e8
- Select Identifier -
iqn.1996-04.de.suse:01:69b6251c20e8

New host port identifier (max 223 characters):

User Label (30 characters maximum):
LCE-10

Step 11 From the Manage Host Port Identifiers window, click **Close**.

Step 12 Repeat Step 4 through Step 11 on any remaining storage subsystems.

Note Do not continue until you have completed Step 12.

Step 13 Close the Subsystem Management window and exit the IBM Storage Manager.

Step 14 Close the VNC client.

Finishing the Cache Engine Replacement

Follow this procedure to finish configuring and installing the new cache engine:

Note You will be re-running the storage script that you ran in the Updating the SSH Keys and Connecting the Cache Engine to the Storage Enclosure section.

Note You must be logged in with root permissions to perform the following steps.

Procedure

Step 1 From the SSH connection to the VDS TC management server, enter the command **cd /opt/pang/useful/configure_storages** to change to the storage configuration script directory.

Step 2 Enter the command **./configure_storages.py -s** *no_of_storage_enclosures* **-b** *no_of_blades* **-c** *the_replaced_CE-ID_number*. For example, to replace CE-2 in a VDS TC Blade Server installation that has two cache engines, and one storage enclosure, enter the command **./configure_storages.py -s 1 -b 2 -c 2**.

Step 3 When the warning message about the disk format appears, enter **y** to continue the storages configuration.

```
*********** WARNING! **********
All storages are about to be reconfigured. All data on the data disks will
be lost.
Do you want to continue? (y/n)y
```

Step 4 Follow the on screen instructions. If prompted, enter the password for the root user that was provided by Cisco. You should see the word “COMPLETE” when the script is finished running.

Note You may be prompted to enter the password several times.

Note Ensure that the VDS-TC_FI_5.7.3b54_Server_Cluster.tar.gz file and the new VDS TC license that you received from Cisco are located on the VDS TC management server in the /tftpboot folder before performing these steps.

Step 5 Enter the command **su admin** to log into the VDS TC CLI. The CLI prompt console> appears.

Step 6 Enter the command **enable** to enter Enable mode. At the password prompt that appears, enter the enable password. The CLI prompt console# appears.

Step 7 Upgrade the software on the new cache engine by entering the command **upgrade server** *CE_number* **127.0.0.1 VDS-TC_FI_5.7.3b54_Server_Cluster.tar.gz**, where *CE_number* is the cache engine you replaced.

Step 8 Use VNC to connect from your local computer to the VDS TC Management Server using the IP address that you configured, adding :1 to the end, for example 10.56.194.65:1. When prompted to enter a password, enter the root password that was provided by Cisco. If you encounter problems opening a VNC client connection to the VDS TC Management Server, you may need to restart VNC Server 1 on the VDS TC Management Server. Follow these steps to restart this server:

a. From an SSH connection to the VDS TC Management server where you are logged in with root permissions, enter the command **rm /tmp/.X11-unix/X1**.

b. Enter the command **vncserver :1** to start the VNC Server.

Figure 5-36 VNC Viewer

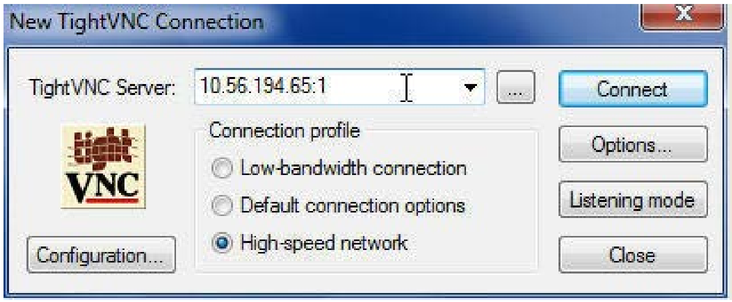

New TightVNC Connection

TightVNC Server: 10.56.194.65:1

Connection profile
○ Low-bandwidth connection
○ Default connection options
● High-speed network

Note If you have problems restarting the VNC Server 1 instance, you can also start another VNC Server instance by entering the command **vncserver :2**.

Step 9 From the VNC console terminal window, enter the command **java -version** to check the current version of java. If the version is *not* 1.8.0_66 or later, perform the following steps to download and install the latest version of java:

a. Log out of any sessions that you have open to the UCS Manager.

b. Go to **www.oracle.com** to download the most current version of java. From the main page that appears, click **Downloads**. From the Top Downloads section, click **Java SE**. From the window that appears, in the Java section, click **Java SE**. From the Java SE Downloads page, click the **JRE Download** button. From the Java SE Runtime Environments page, click the **Accept License Agreement** radio button and download the most current Linux X64 version with the extension of tar.gz.

c. Using SFTP software, such as WinSCP, connect to the VDS TC Management Server and copy the file that you downloaded in Step a to the /tftpboot folder.

d. Enter the command **/opt/pang/useful/update_java.sh /tftpboot/***filename* where *filename* is the name of the file you copied to the /tftpboot folder in Step b. For example:

```
/opt/pang/useful/update_java.sh /tftpboot/jre-8u72-linux-x64.gz
```

e. Enter the command **java -version** to confirm that the java version has been updated. For example:

```
mg-1: # java -version
java version "1.8.0_72"
Java(TM) SE Runtime Environment (build 1.8.0_72-b17)
Java HotSpot(TM) 64-Bit Server VM (build 25.72-b17, mixed mode)
```

Step 10 From the VNC console, enter the command **firefox &**.

Figure 5-37 VNC Console

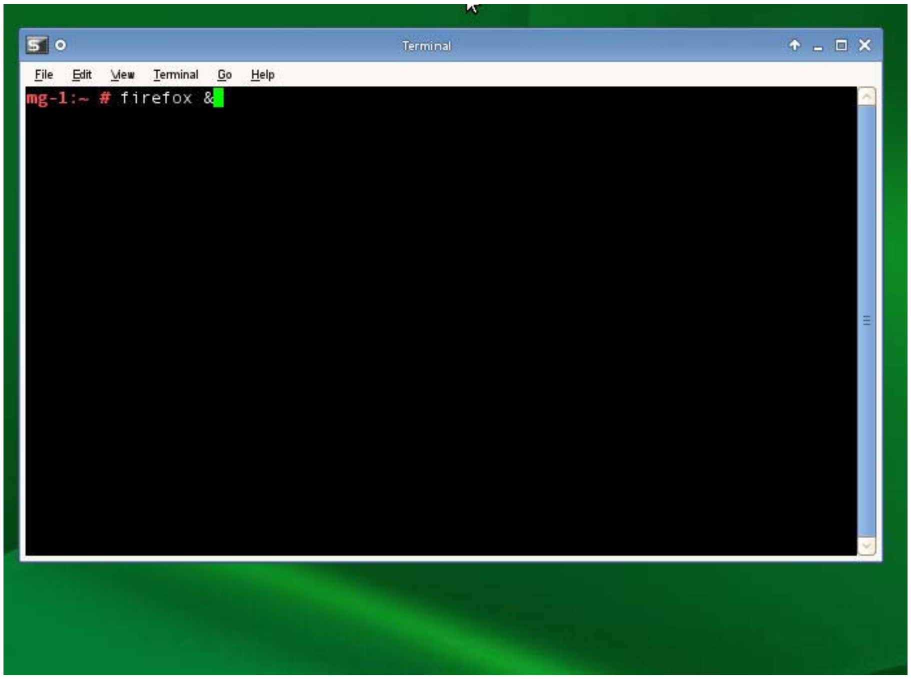

```
mg-1:~ # firefox &
```

Step 11 In the Firefox window that opens, enter the Management port IP address for the fabric interconnect. In a redundant installation, enter the Cluster (virtual) IP address.

Step 12 The Cisco UCS Manager window appears. From this window, click **Launch KVM Manager**. If you receive a message that says "To view this page ensure that Adobe Flash Player version 10.0.0 or greater is installed." perform the following steps:

a. Close Firefox.

b. From the terminal window, enter the command **su –**

c. Enter the command **pkill -9 firefox**

d. Enter the command **cd /root**

e. Enter the command **mv .mozilla .mozilla.bak**

f. Enter the command **firefox &** to restart Firefox.

g. In the Firefox window that opens, enter the Management port IP address for the fabric interconnect. In a redundant installation, enter the Cluster (virtual) IP address.

h. Repeat Step 12.

Figure 5-38 Cisco UCS Manager

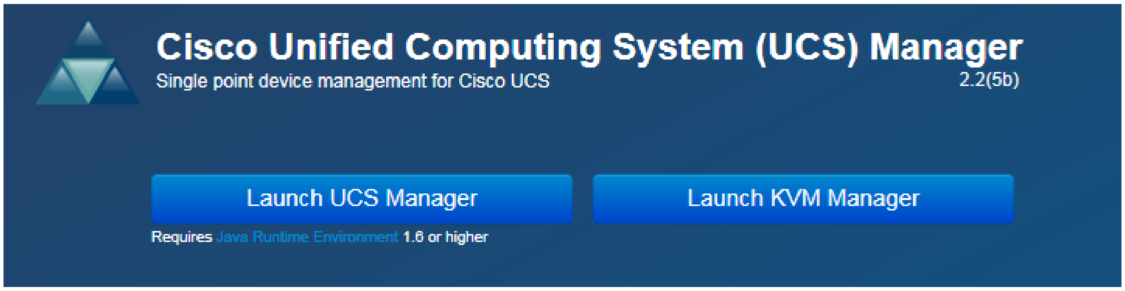

Cisco Unified Computing System (UCS) Manager

Single point device management for Cisco UCS

2.2(5b)

Requires Java Runtime Environment 1.6 or higher

Step 13 The UCS - KVM Launch Manager login window appears. Enter the username **admin** and the password that was configured during the initial system setup of the fabric interconnect (the default password is **password**), and then click **OK**. The UCS - KVM Launch Manager window appears.

Step 14 Perform the following steps to configure the new network settings on the cache engine that you replaced:

a. From the KVM Launch Manager window in the Launch KVM column, click the **Launch** button for **CE-*X*** where *X* is the number of the cache engine that you replaced. When prompted, log in as **root** with the password that was provided by Cisco.

b. Enter the command **cd /tmp**.

c. Enter the command **tar –zxvf Cisco_network_updater_v_1.8.tgz** to extract the package files.

d. Enter the command **cd Cisco_Network_Updater**.

e. Enter the command **./update_network_settings.sh**.

f. From the Welcome window, choose **Yes**, and press **Enter**.

Figure 5-39 Welcome Window

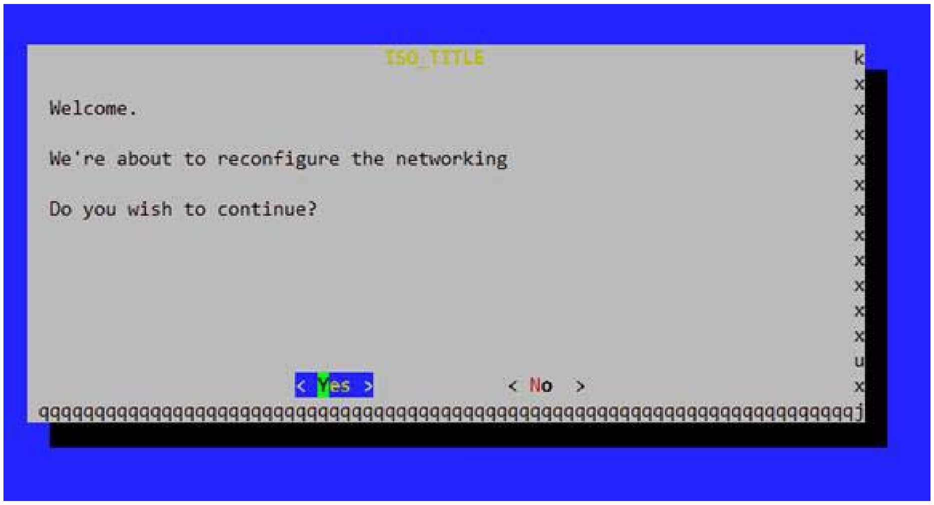

```
ISO_TITLE

Welcome.

We're about to reconfigure the networking

Do you wish to continue?

< Yes >                  < No  >
```

g. From the Appliance Deployment window, choose **Cache Engine**, choose **OK**, and then press **Enter**.

Figure 5-40 Appliance Deployment Window

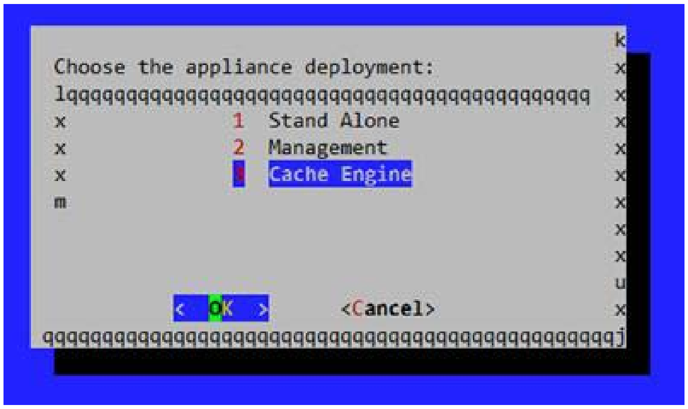

```
                                                     k
  Choose the appliance deployment:                   x
  lqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqq      x
  x              1  Stand Alone                      x
  x              2  Management                       x
  x              3  Cache Engine                     x
  m                                                  x
                                                     x
                                                     x
                                                     u
           <  OK  >        <Cancel>                  x
 qqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqj
```

h. In the Enter Cache Engine Number window, enter the number of the Cache Engine that you replaced, choose **OK**, and then press **Enter**. For example, if you are configuring the first cache engine, enter 1.

Figure 5-41 Cache Engine Window

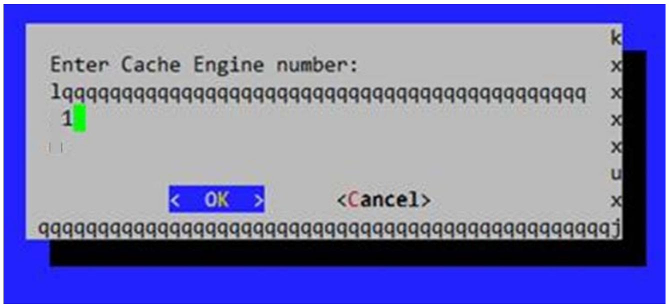

```
                                                     k
  Enter Cache Engine number:                         x
  lqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqq     x
   1                                                 x
                                                     x
                                                     u
           <  OK  >        <Cancel>                  x
 qqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqj
```

i. From the Active-Active Fabric Interconnect window, choose **Yes** and press **Enter**.

Figure 5-42 Active-Active Fabric Interconnect Window

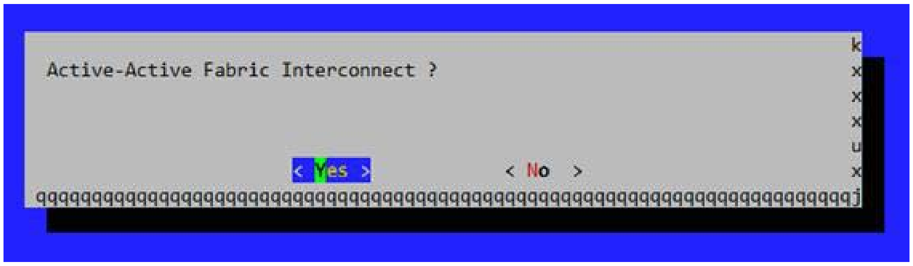

```
                                                                      k
  Active-Active Fabric Interconnect ?                                 x
                                                                      x
                                                                      x
                                                                      u
                     < Yes >             < No  >                      x
 qqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqj
```

j. From the Confirm Configuration window, confirm the Cache Engine configuration. If everything is correct, choose **Yes** and press **Enter**. If there is a mistake, choose **No** and repeat Step f through Step i.

k. Enter the command **reboot -f** to reboot the cache engine. This will close your SSH connection to the cache engine.

Step 15 To import the new license that includes the serial number of the new cache engine, enter the command **license import 127.0.0.1** *filename*, where *filename* is the name of the new license file.

Note You must be in Enable mode, *not* Configuration mode to import a new license.

Step 16 Enter the command **license activate** to apply and activate the license.

Step 17 To start the application on the replaced cache engine, enter the command **oper server** *#*, where # is the number of the cache engine that was replaced. For example, if you replaced CE-2, enter the command **oper server 2**. You should see the following output:

```
console# oper server 2
oper server 2# start
Starting server 2
service is started on server 2
```

Step 18 After the application has started on the new cache engine, enter the command **oper service start** to start the VDS TC service.

Step 19 Wait a few minutes and then enter the command **show status**. Do not proceed to the next step until you see a Device Status of “Started.” For example:

```
console# show status
Operational state Device state Administrative state
enabled started unlocked
```

Step 20 On the PBR routers, start redirecting traffic to the VDS TC solution.

CHAPTER 6

Cache Engine Replacement Procedure for a 5.7.3 Blade Server Installation Upgraded from Release 5.6.1

Note Only use this chapter if you are replacing a Cache Engine in a VDS TC Blade Server system that was upgraded to VDS TC 5.7.3 from VDS TC 5.6.1. If you are replacing a Cache Engine in a VDS TC Blade Server system that has a fresh install of VDS TC 5.7.3 and was *not* upgraded to 5.7.3, refer to Chapter 4, "Cache Engine Replacement Procedure for a Blade Server Installation with a Fresh Install of VDS TC 5.7.3". If you are replacing a Cache Engine in a VDS TC Blade Server system that was upgraded to VDS TC 5.7.3 from VDS TC 5.2.1, refer to Chapter 5, "Cache Engine Replacement Procedure for a 5.7.3 Blade Server Installation Upgraded from Release 5.2.1".

To determine whether the system has been upgraded and from what version, from the Management Server enter the command **grep "upgrade system to" /var/log/peerapp/peerapp_system***. This will display a list of all previous versions.

This chapter discusses how to determine that a cache engine has failed in a VDS TC Blade Server installation and the steps that are required to replace it.

When a cache engine fails, the volume_selection_algorithm_time setting in the cluster_conf.xml file controls when the volumes of a failed cache engine will be redistributed (mounted) to the remaining cache engines. By default this value is 24 hours, which means that by default, if a cache engine fails, its volumes will not be redistributed (mounted) to the remaining cache engines until 24 hours after the cache engine fails.

It is recommended that you leave the default set to 24 hours, however if you need to change this setting, add the following text to the cluster_conf file in the <service> section under the <memory> subsection. The *seconds* parameters is the number of seconds that the system should wait after a cache engine fails before its volumes are redistributed (mounted) to the remaining cache engines:

```
<io>
    <volume_selection_algorithm_time>seconds</volume_selection_algorithm_time>
</io>
```

For example, to have the system wait 30 minutes after a cache engine fails before redistributing the volumes of the failed cache engine enter the following:

```
<io>
    <volume_selection_algorithm_time>1800</volume_selection_algorithm_time>
</io>
```

For more information on how to make changes to the configuration file and how to apply these changes, please refer to the *Cisco Videoscape Distribution Suite Transparent Caching Software Configuration Guide* available at http://www.cisco.com/c/dam/en/us/td/docs/video/videoscape/distribution_suite/vds/v5_7_3/VDS-TC_5_7_3_software_config_guide.pdf.

Symptoms of a Failed VDS TC Blade Server Cache Engine

To determine whether a cache engine has failed in a VDS TC Blade Server installation, log into VDS TC Manager GUI. By default, the Status > Dashboard window is displayed. If you are already logged into VDS TC Manager, choose **Status** > **Dashboard**.

If a cache engine has failed it will appear with a gray box surrounding it.

Figure 6-1 Failed Cache Engine

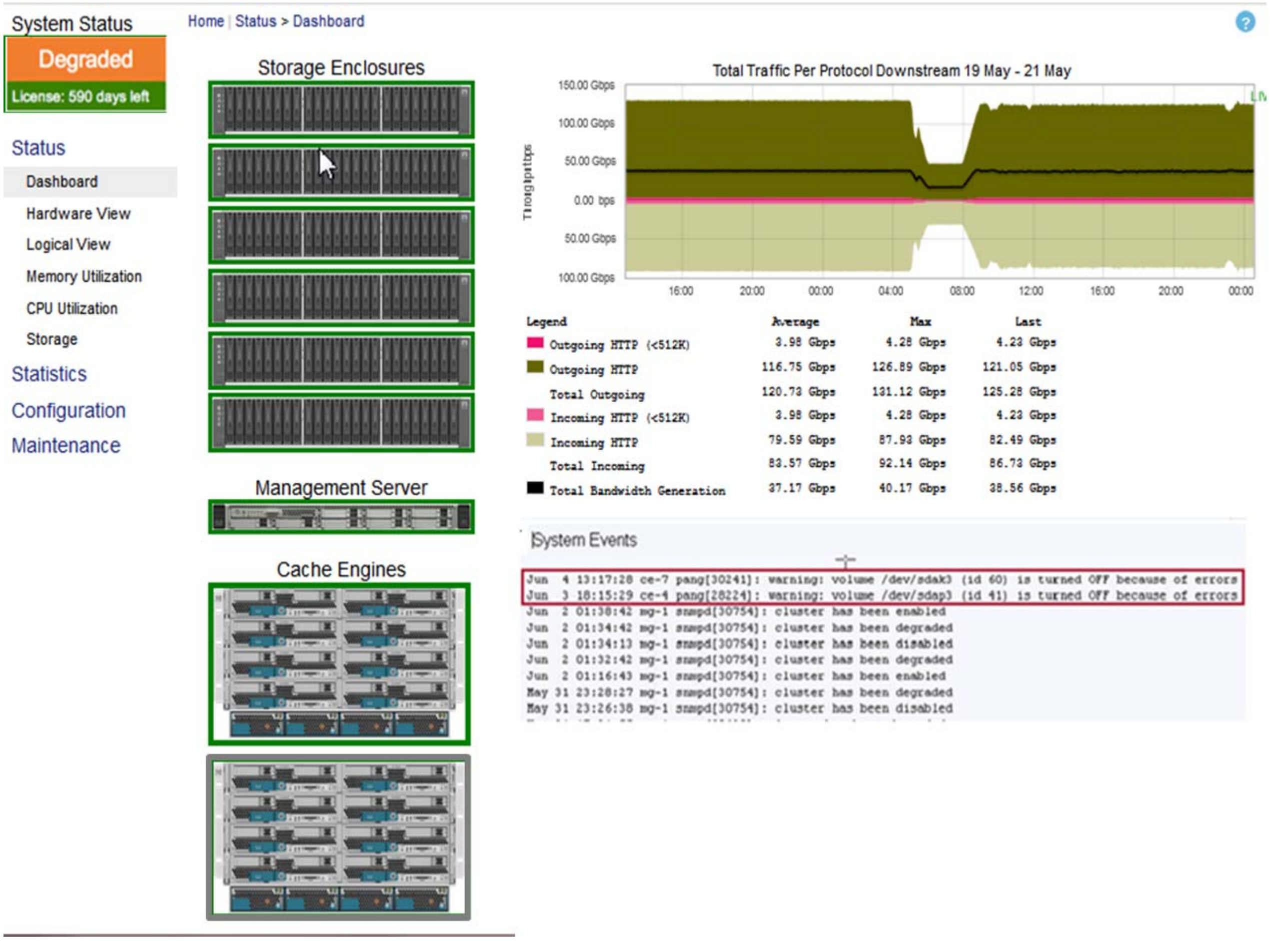

Legend	Average	Max	Last
🟥 Outgoing HTTP (<512K)	3.98 Gbps	4.28 Gbps	4.23 Gbps
🟫 Outgoing HTTP	116.75 Gbps	126.89 Gbps	121.05 Gbps
Total Outgoing	120.73 Gbps	131.12 Gbps	125.28 Gbps
🟥 Incoming HTTP (<512K)	3.98 Gbps	4.28 Gbps	4.23 Gbps
🟨 Incoming HTTP	79.59 Gbps	87.93 Gbps	82.49 Gbps
Total Incoming	83.57 Gbps	92.14 Gbps	86.73 Gbps
■ Total Bandwidth Generation	37.17 Gbps	40.17 Gbps	38.56 Gbps

```
Jun  4 13:17:28 ce-7 pang[30241]: warning: volume /dev/sdak3 (id 60) is turned OFF because of errors
Jun  3 18:15:29 ce-4 pang[28224]: warning: volume /dev/sdap3 (id 41) is turned OFF because of errors
Jun  2 01:38:42 mg-1 snmpd[30754]: cluster has been enabled
Jun  2 01:34:42 mg-1 snmpd[30754]: cluster has been degraded
Jun  2 01:34:13 mg-1 snmpd[30754]: cluster has been disabled
Jun  2 01:32:42 mg-1 snmpd[30754]: cluster has been degraded
Jun  2 01:16:43 mg-1 snmpd[30754]: cluster has been enabled
May 31 23:28:27 mg-1 snmpd[30754]: cluster has been degraded
May 31 23:26:38 mg-1 snmpd[30754]: cluster has been disabled
```

Obtaining and Uploading the Required Files

Obtaining Required Files

Before you replace the VDS TC cache engine, you must obtain the following:

- **The VDS TC 5.6.1 ISO image file:** VDS-TC_Installer_5.6.1b57-5.6.1b58-ISO-15-Cisco.iso
- **The VDS TC 5.7.3 ISO image file:** VDS-TC_Installer_5.7.3b53-5.7.3b54-ISO-7-Cisco.iso
- **A new VDS TC license from Cisco:** To replace a failed cache engine you must request a new VDS TC license from Cisco. To request a new license, you must have the serial number of the failed cache engine and the serial number of the new cache engine. After you have gathered this information, contact Cisco support for a new license file.

Copying the Files to the VDS TC Management Server

Procedure

Step 1 Using SFTP software, such as WinSCP, connect to the VDS TC Management Server using the IP address that you assigned during the “Configuring the Boot Image” process. Log in using the user name **padmin** and the password that was provided by Cisco.

Step 2 Enter the command **cd /opt/pang/iso** to change to the /opt/pang/iso folder.

Step 3 Copy the VDS TC 5.6.1 ISO file, VDS-TC_Installer_5.6.1b57-5.6.1b58-ISO-15-Cisco.iso, into the **/opt/pang/iso** folder.

Step 4 Copy the VDS TC 5.7.3 ISO image file, VDS-TC_Installer_5.7.3b53-5.7.3b54-ISO-7-Cisco.iso, into the **/opt/pang/iso** folder.

Note You need the ISO image for both VDS TC 5.6.1 and VDS TC 5.7.3 to complete the Cache Engine replacement.

Step 5 Copy the new VDS TC license that you received from Cisco to the **/tftpboot** folder.

Step 6 Close the SFTP software.

Extracting the VDS TC 5.7.3 Files

Step 1 From the SSH connection to the VDS TC management server logged in as root, enter the command **mkdir /opt/pang/iso/5.7.3_iso** to create a new folder for the 5.7.3 files.

Step 2 Enter the command **cd /opt/pang/iso** to change folders.

Step 3 Enter the command **mount VDS-TC_Installer_5.7.3b53-5.7.3b54-ISO-7-Cisco.iso -o loop /opt/pang/iso/5.7.3_iso** to extract the VDS-TC_FI_5.7.3b54_Server_Cluster.tar.gz file.

Step 4 Enter the command **cd /** to change to the root folder and then enter the command **umount /opt/pang/iso/5.7.3_iso** to unmount the iso image.

Step 5 Enter the command **rm /opt/pang/iso/VDS-TC_Installer_5.7.3b53-5.7.3b54-ISO-7-Cisco.iso** to delete the iso file.

Step 6 Enter the command **cd /opt/pang/iso/5.7.3_iso/install-tools/PeerApp_FI-GA**.

Step 7 Enter the command **cp VDS-TC_FI_5.7.3b54_Server_Cluster.tar.gz /tftpboot**.

Replacing a Cache Engine and Running the Installation Tasks

The first step in replacing a cache engine in a VDS TC Blade Server cluster installation is to physically replace the cache engine and perform the installation steps on the new cache engine. Follow this procedure to perform these steps:

Before You Begin

To replace a failed cache engine you must request a new VDS TC license from Cisco. To request a new license, you must have the serial number of the failed cache engine and the serial number of the new cache engine. After you have gathered this information, contact Cisco support for a new license file.

Caution Do not continue until you have the license file.

Procedure

Note Before removing the failed cache engine, make note of what number corresponds with the cache engine.

Step 1 Using the Cisco UCS 5108 Server Chassis Installation Guide available at http://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/hw/chassis-install-guide/ucs5108_install.html, unmount the failed Blade Server cache engine and mount the new Blade Server cache engine.

Step 2 Confirm that the correct firmware is installed on the Cache Engine for VDS TC 5.7.3:

- Cisco UCS-B200-M3:
  - Firmware version: 2.2.5b
  - Image file: ucs-k9-bundle-b-series.2.2.5b.B.bin

Step 3 In the *Cisco Videoscape Distribution Suite Transparent Caching Software Blade Server Cluster Installation Guide*, available at http://www.cisco.com/c/dam/en/us/td/docs/video/videoscape/distribution_suite/vds/v5_7_3/VDS-TC_5_7_3_BladeSrvr_swInstallGuide.pdf, refer to Chapter 1, “Prerequisites and Blade Server Cluster Physical Installation” for connecting the network cables to the new cache engine.

Note Complete the following steps for the cache engine that you are replacing.

Step 1 Use VNC to connect from your local computer to the VDS TC Management Server using the IP address that you configured, adding :1 to the end, for example 10.56.194.65:1. When prompted to enter a password, enter the root password that was provided by Cisco. If you encounter problems opening a VNC client connection to the VDS TC Management Server, you may need to restart VNC Server 1 on the VDS TC Management Server. Follow these steps to restart this server:

a. From an SSH connection to the VDS TC Management server where you are logged in with root permissions, enter the command **rm /tmp/.X11-unix/X1**.

b. Enter the command **vncserver :1** to start the VNC Server.

Figure 6-2 VNC Viewer

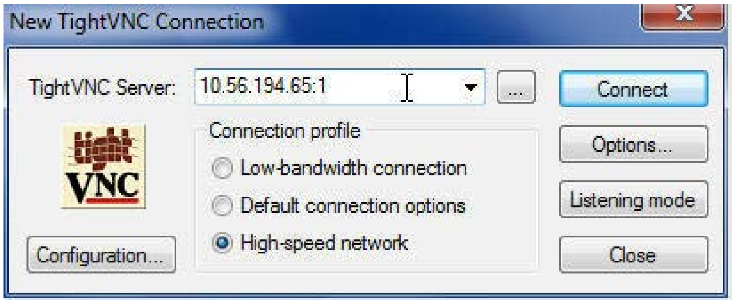

New TightVNC Connection

TightVNC Server: 10.56.194.65:1

Connection profile

○ Low-bandwidth connection

○ Default connection options

● High-speed network

Note If you have problems restarting the VNC Server 1 instance, you can also start another VNC Server instance by entering the command **vncserver :2**.

Step 2 From the VNC console, enter the command **firefox &**.

Figure 6-3 VNC Console

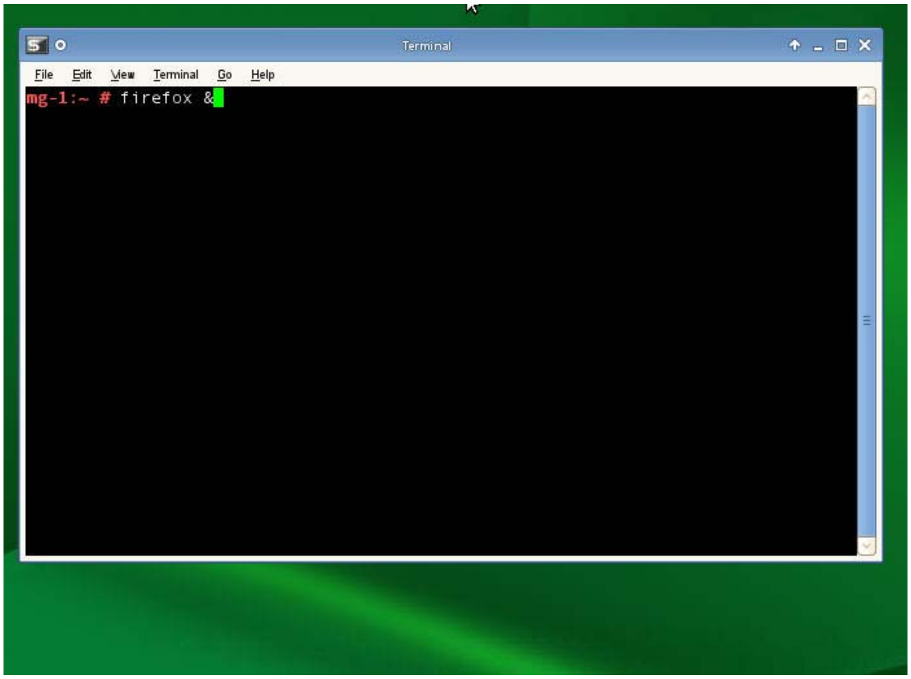

```
mg-1:~ # firefox &
```

Step 3 In the Firefox window that opens, enter the Management port IP address for the fabric interconnect. In a redundant installation, enter the Cluster (virtual) IP address.

Step 4 The Cisco UCS Manager window appears. From this window, click **Launch KVM Manager**. If you receive a message that says “To view this page ensure that Adobe Flash Player version 10.0.0 or greater is installed.” perform the following steps:

a. Close Firefox

b. From the terminal window, enter the command **su –**

c. Enter the command **pkill -9 firefox**

d. Enter the command **cd /root**

e. Enter the command **mv .mozilla .mozilla.bak**

f. Enter the command **firefox &** to restart Firefox.

g. In the Firefox window that opens, enter the Management port IP address for the fabric interconnect. In a redundant installation, enter the Cluster (virtual) IP address.

h. Repeat Step 4.

Figure 6-4 Cisco UCS Manager

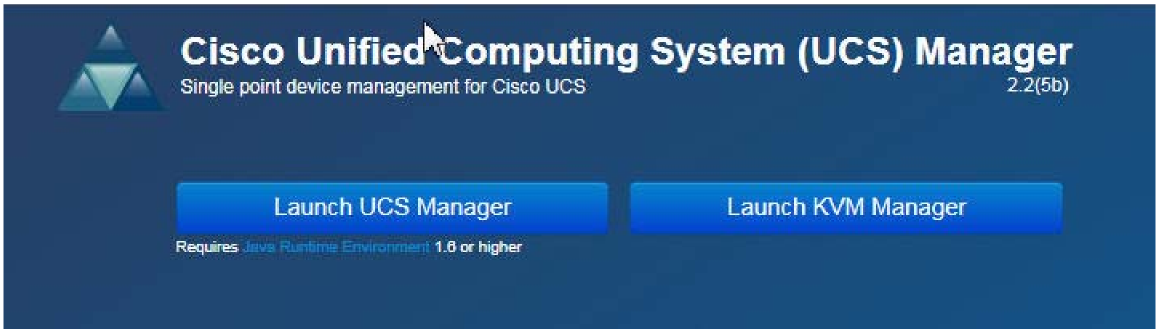
Cisco Unified Computing System (UCS) Manager
Single point device management for Cisco UCS
2.2(5b)
Launch UCS Manager
Launch KVM Manager
Requires Java Runtime Environment 1.6 or higher

Step 5 The UCS - KVM Launch Manager login window appears. Enter the username **admin** and the password that was configured during the initial system setup of the fabric interconnect (the default password is **password**), and then click **OK**.

Figure 6-5 UCS KVM Launch Manager Login Window

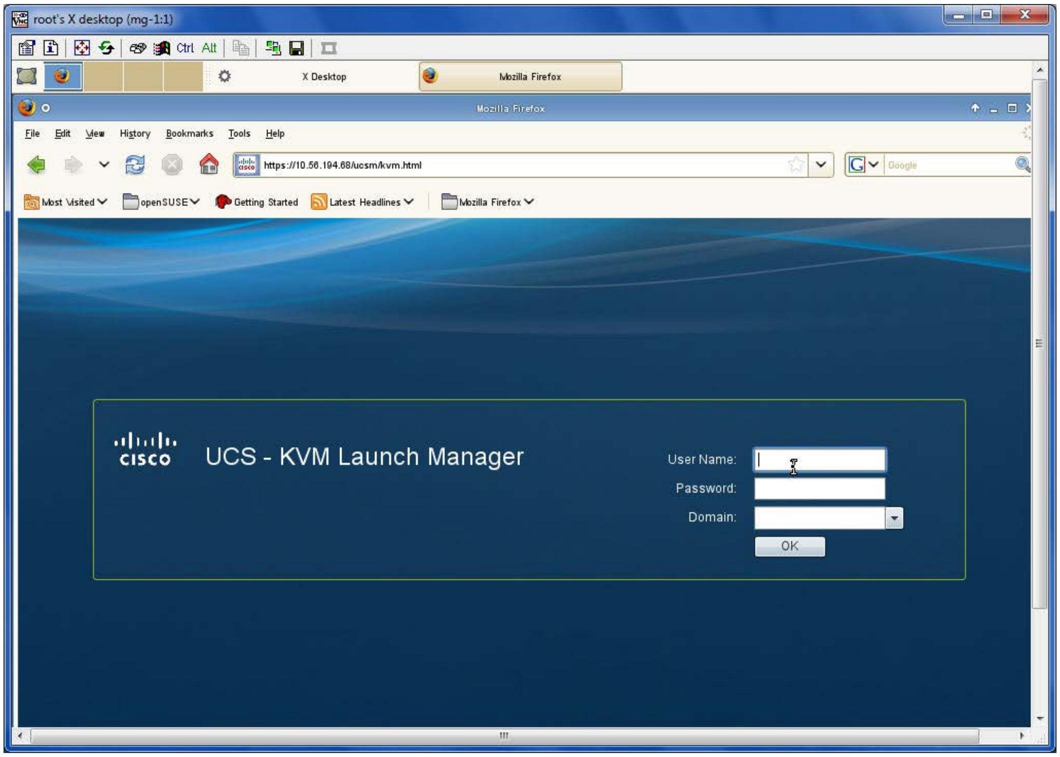

Step 6 The Service Profiles window appears, listing all of the cache engine profiles. From this window, in the row for the cache engine you are replacing, in the Launch KVM column click **Launch** to launch the KVM console.

Figure 6-6 Service Profiles Window

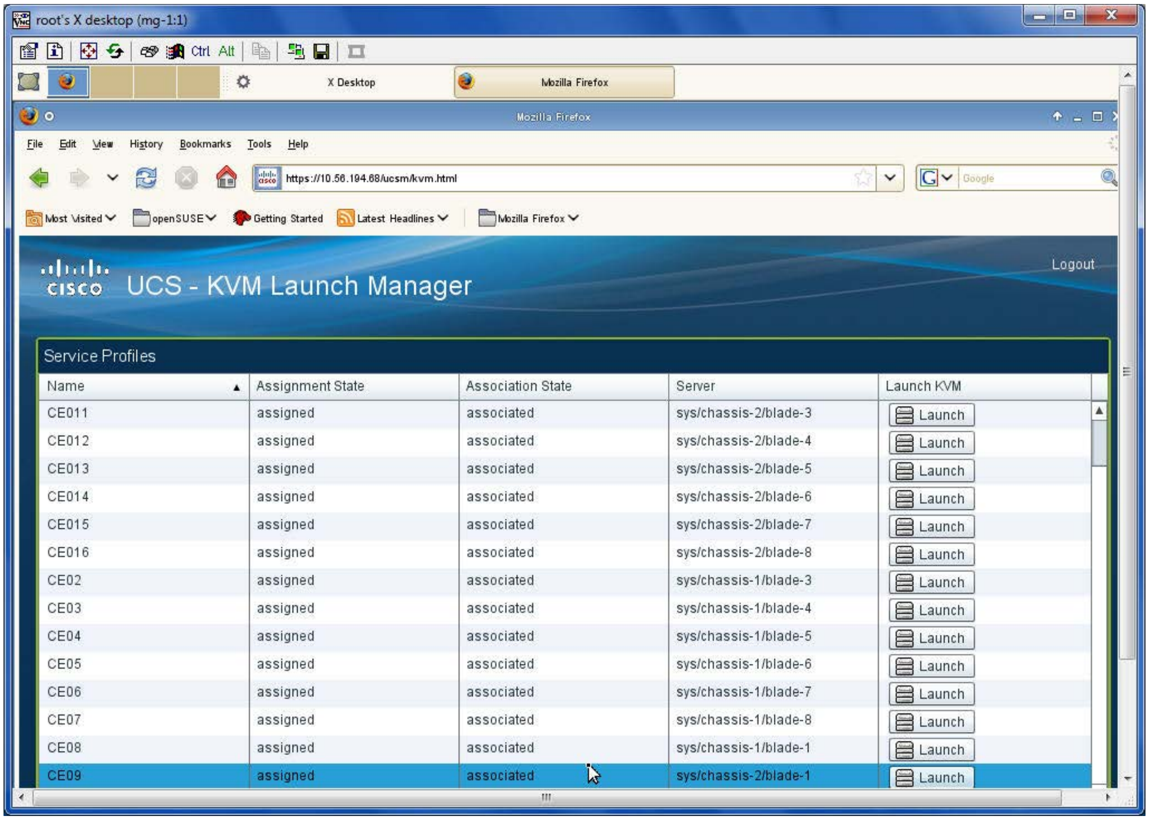

https://10.56.194.68/ucsm/kvm.html

UCS - KVM Launch Manager

Service Profiles

Name	Assignment State	Association State	Server	Launch KVM
CE011	assigned	associated	sys/chassis-2/blade-3	Launch
CE012	assigned	associated	sys/chassis-2/blade-4	Launch
CE013	assigned	associated	sys/chassis-2/blade-5	Launch
CE014	assigned	associated	sys/chassis-2/blade-6	Launch
CE015	assigned	associated	sys/chassis-2/blade-7	Launch
CE016	assigned	associated	sys/chassis-2/blade-8	Launch
CE02	assigned	associated	sys/chassis-1/blade-3	Launch
CE03	assigned	associated	sys/chassis-1/blade-4	Launch
CE04	assigned	associated	sys/chassis-1/blade-5	Launch
CE05	assigned	associated	sys/chassis-1/blade-6	Launch
CE06	assigned	associated	sys/chassis-1/blade-7	Launch
CE07	assigned	associated	sys/chassis-1/blade-8	Launch
CE08	assigned	associated	sys/chassis-1/blade-1	Launch
CE09	assigned	associated	sys/chassis-2/blade-1	Launch

Step 7 Click **OK** to open the .jnlp file with Java(TM) Web Start Launcher.

Step 8 From the Warning Security window, click **Continue**.

Step 9 If you receive a window that asks if you want to run this application, click **Run**.

Step 10 From the KVM Console window, make sure the **KVM Console** tab is selected and choose **Virtual Media > Activate Virtual Devices**.

Step 11 If you see the Unencrypted Virtual Media Session window, choose **Accept this session** and click **Apply**.

Figure 6-7 Unencrypted Virtual Media Session

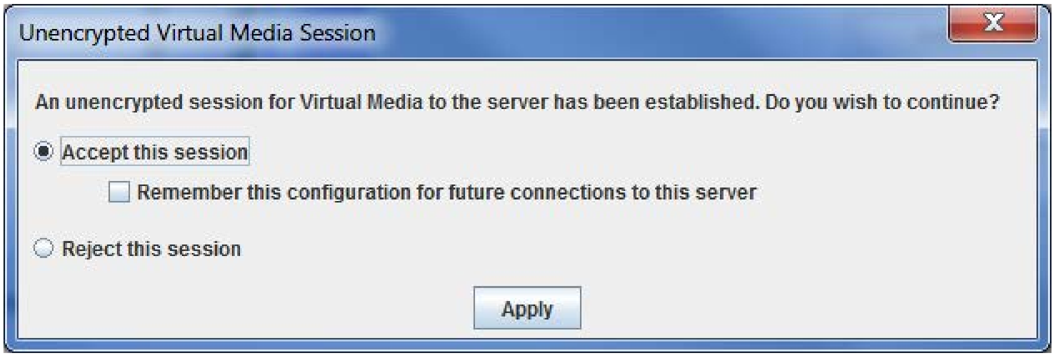

Unencrypted Virtual Media Session

An unencrypted session for Virtual Media to the server has been established. Do you wish to continue?

● Accept this session

□ Remember this configuration for future connections to this server

○ Reject this session

Apply

Step 12 Choose **Virtual Media** > **Map CD/DVD**.

Figure 6-8 Map CD/DVD

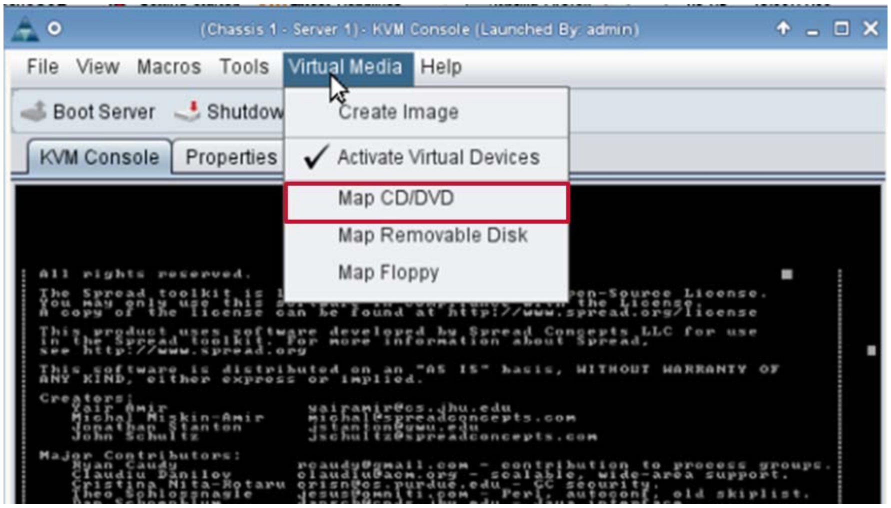

```
All rights reserved.
The Spread toolkit is l                                   pen-Source License.
You may only use this                                     the License.
A copy of the license can be found at http://www.spread.org/license
This product uses software developed by Spread Concepts LLC for use
in the Spread toolkit. For more information about Spread,
see http://www.spread.org
This software is distributed on an "AS IS" basis, WITHOUT WARRANTY OF
ANY KIND, either express or implied.
Creators:
    Yair Amir                 yairamir@cs.jhu.edu
    Michal Miskin-Amir        michal@spreadconcepts.com
    Jonathan Stanton          jstanton@gwu.edu
    John Schultz              jschultz@spreadconcepts.com
Major Contributors:
    Ryan Caudy               rcaudy@gmail.com - contribution to process groups.
    Claudiu Danilov          claudiu@acm.org - scalable, wide-area support.
    Cristina Nita-Rotaru     crisn@cs.purdue.edu - GC security.
    Theo Schlossnagle        jesus@omniti.com - Perl, autoconf, old skiplist.
```

Step 13 In the window that opens, browse to the **/opt/pang/iso/** folder to find the ISO installer image for VDS TC 5.6.1, VDS-TC_Installer_5.6.1b57-5.6.1b58-ISO-15-Cisco.iso, and click **Open**.

Figure 6-9 Map Device

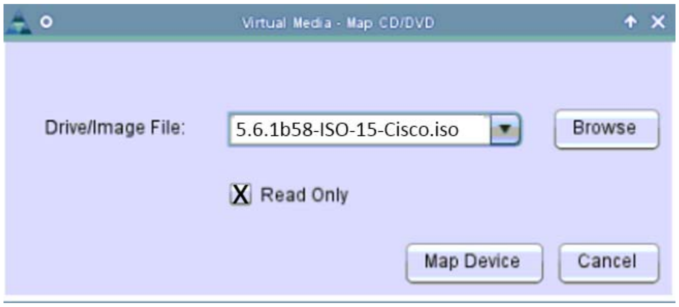

Drive/Image File: 5.6.1b58-ISO-15-Cisco.iso

☒ Read Only

Step 14 Check the **Read Only** check box and click **Map Device**.

Step 15 Click **Reset** under the menu bar to reboot the cache engine.

Figure 6-10 Reset the Cache Engine

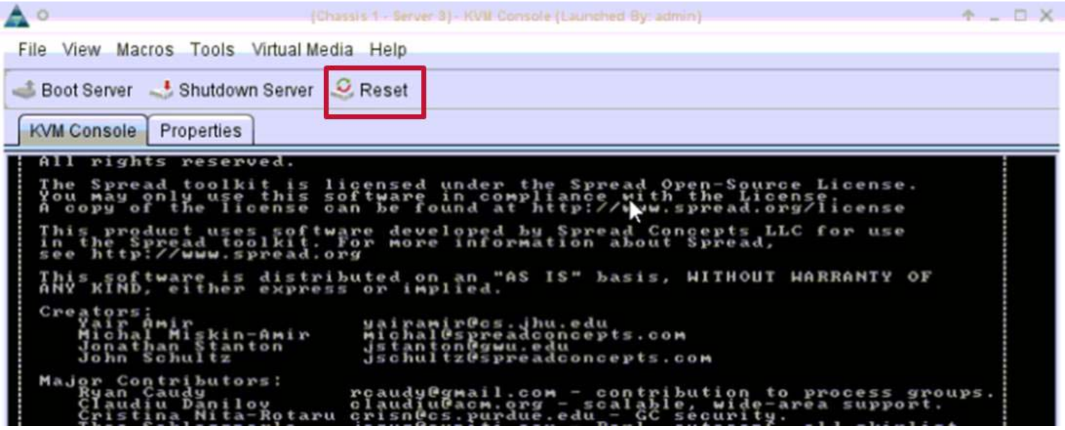

```
All rights reserved.

The Spread toolkit is licensed under the Spread Open-Source License.
You may only use this software in compliance with the License.
A copy of the license can be found at http://www.spread.org/license

This product uses software developed by Spread Concepts LLC for use
in the Spread toolkit. For more information about Spread,
see http://www.spread.org

This software is distributed on an "AS IS" basis, WITHOUT WARRANTY OF
ANY KIND, either express or implied.

Creators:
    Yair Amir              yairamir@cs.jhu.edu
    Michal Miskin-Amir     michal@spreadconcepts.com
    Jonathan Stanton       jstanton@gwu.edu
    John Schultz           jschultz@spreadconcepts.com

Major Contributors:
    Ryan Caudy            rcaudy@gmail.com - contribution to process groups.
    Claudiu Danilov       claudiu@acm.org - scalable, wide-area support.
    Cristina Nita-Rotaru crisn@cs.purdue.edu - GC security.
```

Step 16 From the Reset Server Warning window that appears, click **OK** to reset the server.

Figure 6-11 Reset Server Warning Window

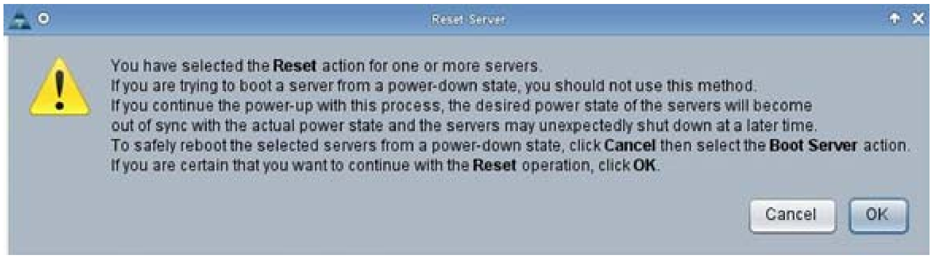

You have selected the **Reset** action for one or more servers.
If you are trying to boot a server from a power-down state, you should not use this method.
If you continue the power-up with this process, the desired power state of the servers will become
out of sync with the actual power state and the servers may unexpectedly shut down at a later time.
To safely reboot the selected servers from a power-down state, click **Cancel** then select the **Boot Server** action.
If you are certain that you want to continue with the **Reset** operation, click **OK**.

Step 17 From the Reset Server Service window that appears, choose **Power Cycle** and click **OK**.

Figure 6-12 Reset Server Service

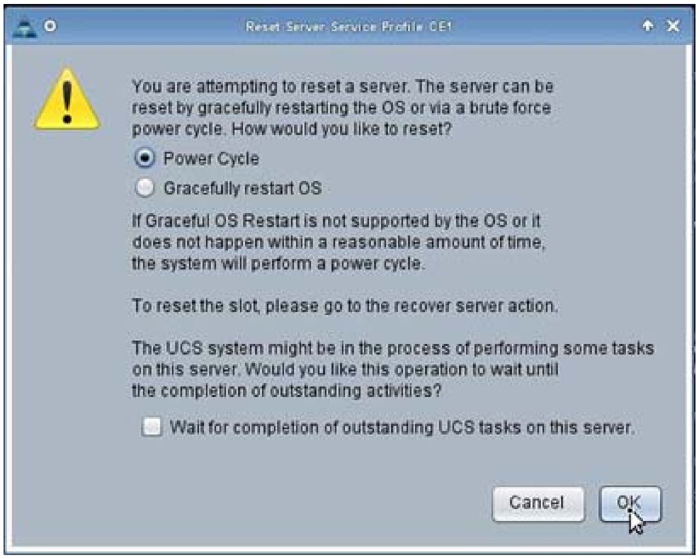

You are attempting to reset a server. The server can be reset by gracefully restarting the OS or via a brute force power cycle. How would you like to reset?

● Power Cycle
○ Gracefully restart OS

If Graceful OS Restart is not supported by the OS or it does not happen within a reasonable amount of time, the system will perform a power cycle.

To reset the slot, please go to the recover server action.

The UCS system might be in the process of performing some tasks on this server. Would you like this operation to wait until the completion of outstanding activities?

□ Wait for completion of outstanding UCS tasks on this server.

Step 18 From the KVM Console window you will see the VDS TC Cache Engine reboot. When you see the prompt to press F6 to enter the boot menu, click the KVM Console window to make sure it has focus and press **F6**. The Please Select Boot Device dialog box appears.

Note You may need to press **F6** several times to see the Please Select Boot Device dialog box.

Figure 6-13 System Reboot Window

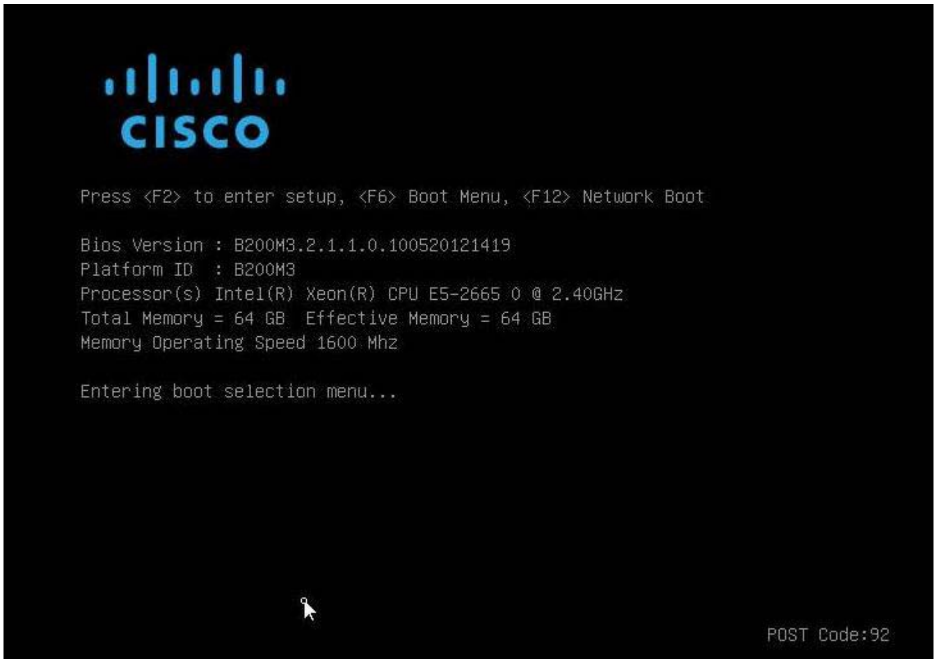

```
Press <F2> to enter setup, <F6> Boot Menu, <F12> Network Boot

Bios Version : B200M3.2.1.1.0.100520121419
Platform ID  : B200M3
Processor(s) Intel(R) Xeon(R) CPU E5-2665 0 @ 2.40GHz
Total Memory = 64 GB  Effective Memory = 64 GB
Memory Operating Speed 1600 Mhz

Entering boot selection menu...

                                              POST Code:92
```

Step 19 Using the up and down arrows, choose **Cisco vKVM-Mapped vDVD1.22**. The system will reboot from the virtual ISO image.

Step 20 When the VDS-TC Installer environment displays in the KVM console window, choose **Installer - No Serial**.

Figure 6-14 VDS TC Installer Window

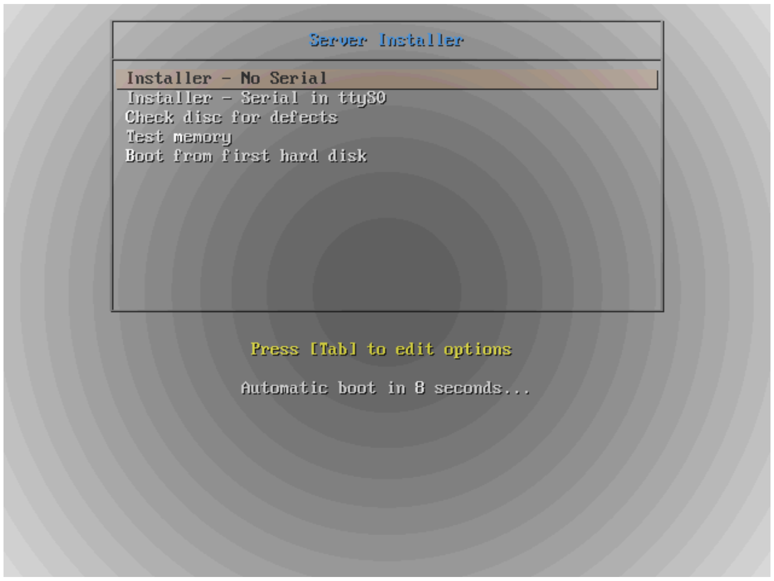

```
Server Installer

Installer - No Serial
Installer - Serial in ttyS0
Check disc for defects
Test memory
Boot from first hard disk

Press [Tab] to edit options

Automatic boot in 8 seconds...
```

Step 21 When the Welcome screen appears, enter **Y** to continue the installation process.

Figure 6-15 Welcome Screen

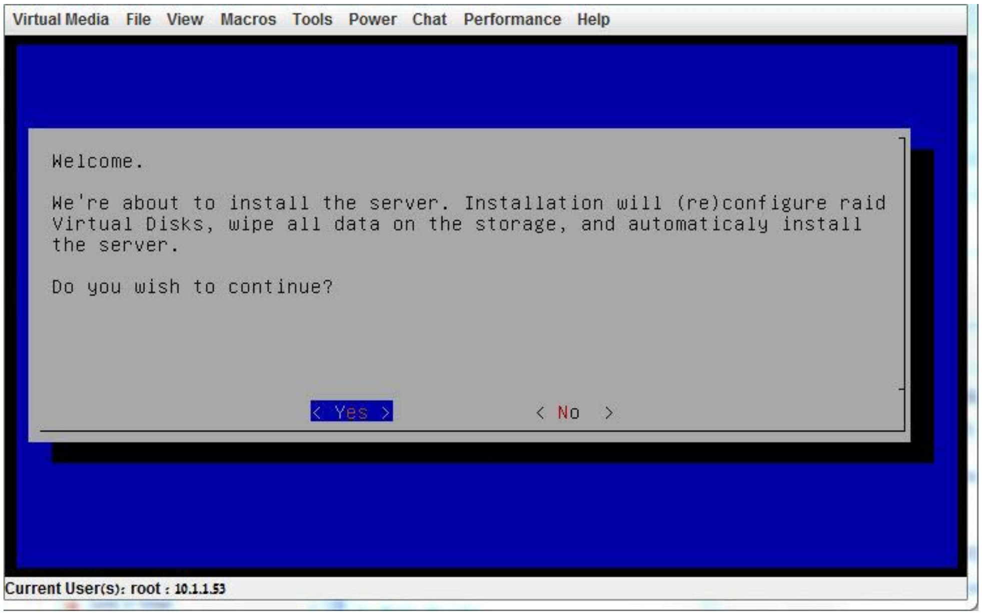

Virtual Media File View Macros Tools Power Chat Performance Help

```
Welcome.

We're about to install the server. Installation will (re)configure raid
Virtual Disks, wipe all data on the storage, and automaticaly install
the server.

Do you wish to continue?

< Yes >                 < No  >
```

Current User(s): root : 10.1.1.53

Step 22 When the Choose the Appliance Deployment window appears, choose **Cache Engine** and select **OK**.

Figure 6-16 Choose the Appliance Deployment Window

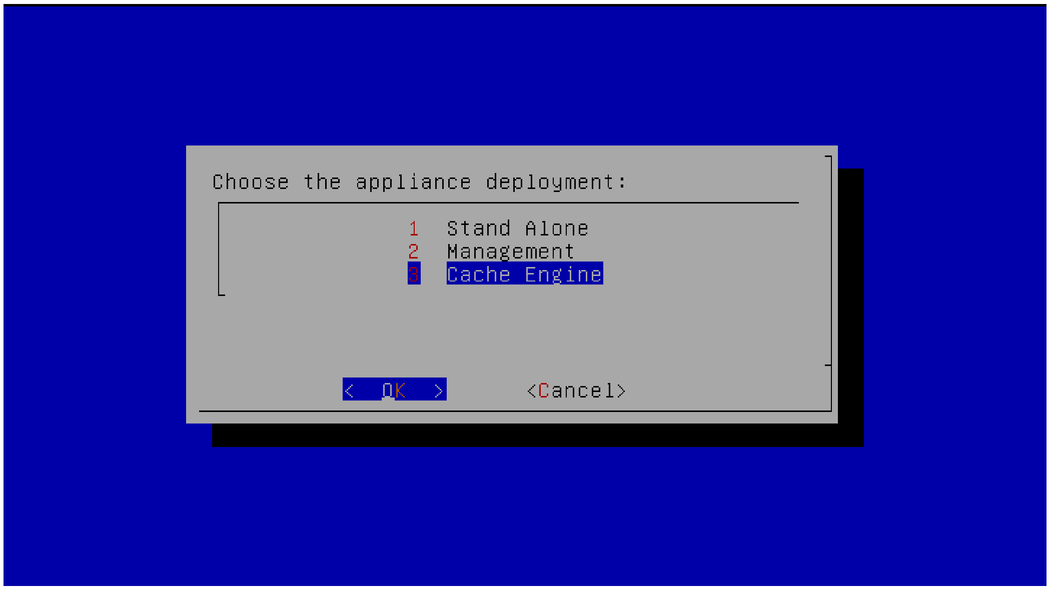

```
Choose the appliance deployment:

          1  Stand Alone
          2  Management
          3  Cache Engine
```

Step 23 When the "Enter Cache Engine Number" dialog box appears, enter the number that corresponds with the cache engine that you are replacing.

Figure 6-17 Enter Cache Engine Number

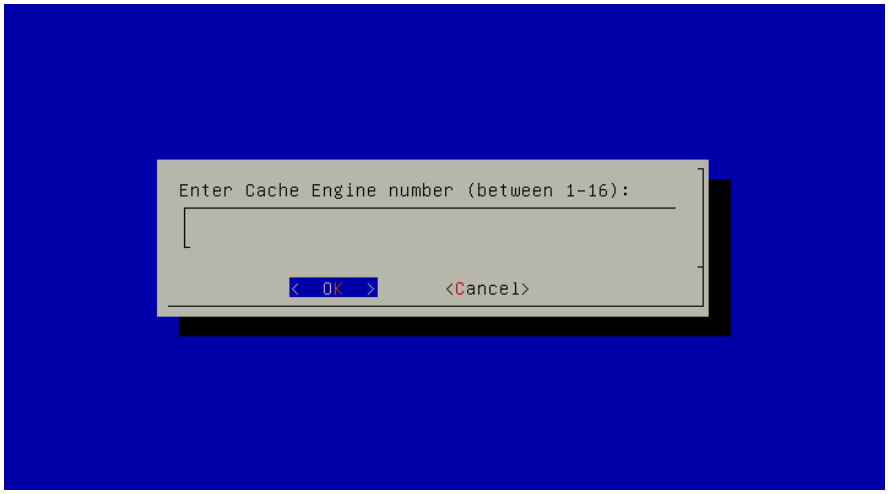

```
Enter Cache Engine number (between 1-16):
```

Step 24 From the Active-Active Fabric Interconnect dialog box, choose **yes**.

Figure 6-18 Active-Active Fabric Interconnect

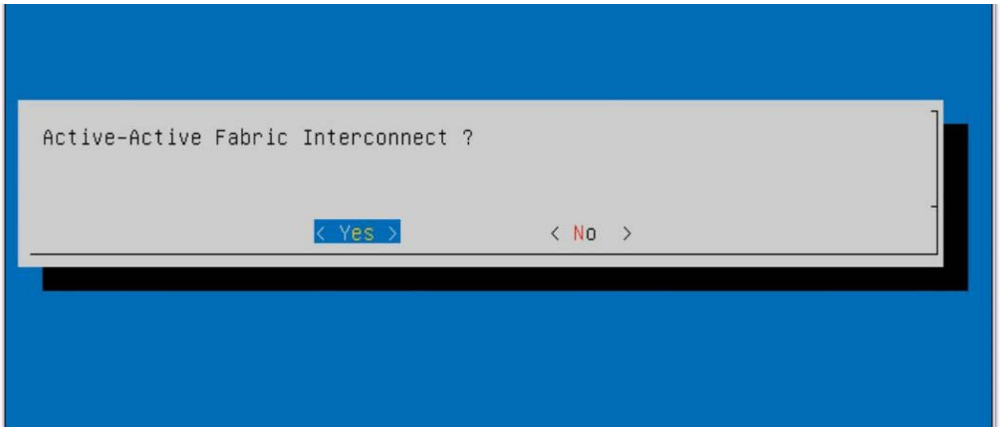

Active-Active Fabric Interconnect ?

Step 25 Confirm the Cache Engine configuration. If everything is correct, choose **Yes**. If there is a mistake, choose **No** and repeat Step 21 through Step 24.

Figure 6-19 Configuration Confirmation

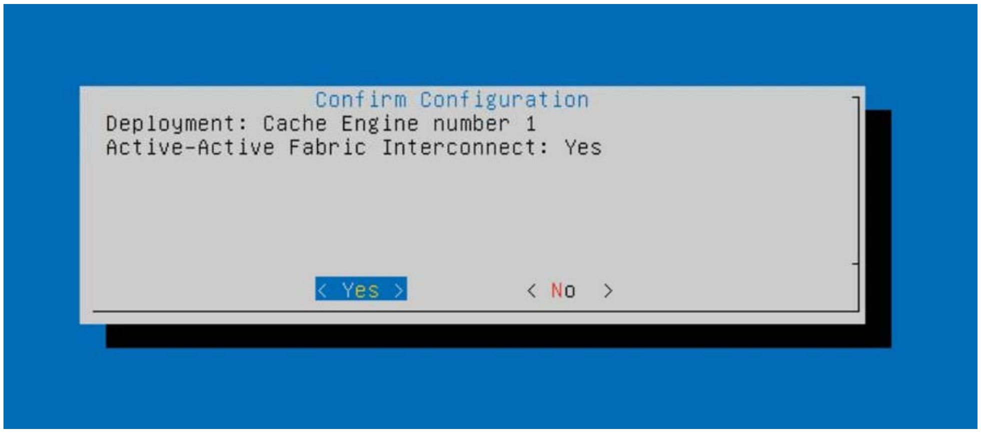

Confirm Configuration
Deployment: Cache Engine number 1
Active-Active Fabric Interconnect: Yes

Step 26 From the KVM window you will see the installation begin. You must wait for this installation process to complete before you continue.

Note This installation should take about 40 minutes.

Figure 6-20 Cache Engine Installation Process

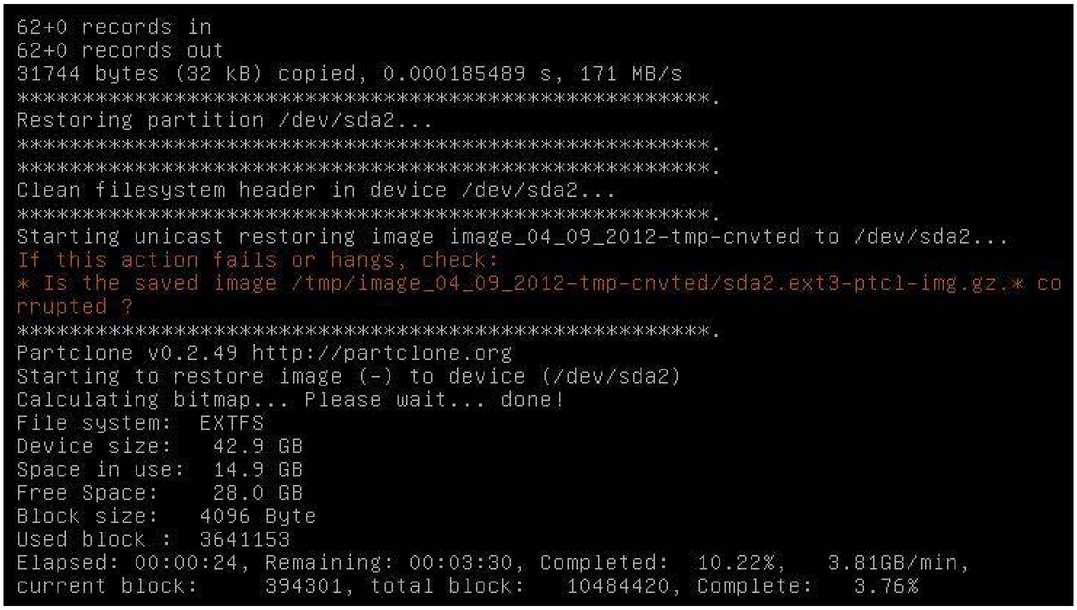

```
62+0 records in
62+0 records out
31744 bytes (32 kB) copied, 0.000185489 s, 171 MB/s
******************************************************.
Restoring partition /dev/sda2...
******************************************************.
******************************************************.
Clean filesystem header in device /dev/sda2...
******************************************************.
Starting unicast restoring image image_04_09_2012-tmp-cnvted to /dev/sda2...
If this action fails or hangs, check:
* Is the saved image /tmp/image_04_09_2012-tmp-cnvted/sda2.ext3-ptcl-img.gz.* co
rrupted ?
******************************************************.
Partclone v0.2.49 http://partclone.org
Starting to restore image (-) to device (/dev/sda2)
Calculating bitmap... Please wait... done!
File system:  EXTFS
Device size:   42.9 GB
Space in use:  14.9 GB
Free Space:    28.0 GB
Block size:   4096 Byte
Used block :  3641153
Elapsed: 00:00:24, Remaining: 00:03:30, Completed:  10.22%,   3.81GB/min,
current block:     394301, total block:   10484420, Complete:   3.76%
```

Step 27 When the installation completes, you must unmap the ISO image file. When the “Installation completed. Press <Enter> to reboot.” message appears, press **Enter**.

Figure 6-21 Installation Completed Window

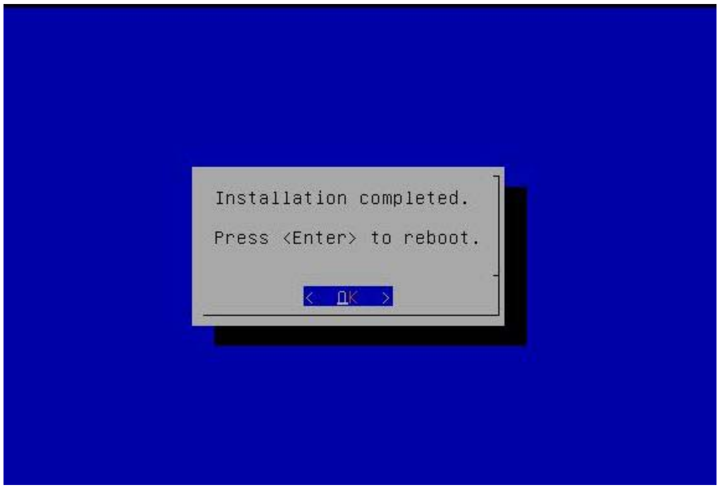

```
Installation completed.

Press <Enter> to reboot.
```

Note *Wait* for the system to reboot and display the BIOS splash screen, as shown in Figure 6-13 “System Reboot Window”, before proceeding to the next step.

Step 28 After the system reboots, from the KVM Console window, choose **Virtual Media** > **Activate Virtual Devices** to deactivate the virtual devices and unmap the image file.

Figure 6-22 Deactivate Virtual Device

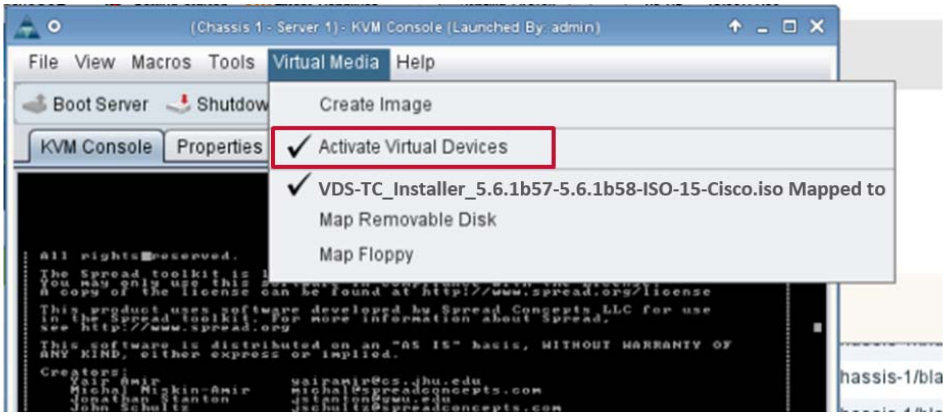
(Chassis 1 - Server 1) - KVM Console (Launched By: admin)
File
View
Macros
Tools
Virtual Media
Help
Boot Server
Shutdow
KVM Console
Properties
Create Image
✓ Activate Virtual Devices
✓ VDS-TC_Installer_5.6.1b57-5.6.1b58-ISO-15-Cisco.iso Mapped to
Map Removable Disk
Map Floppy
hassis-1/bla

```
All rights reserved.
The Spread toolkit is l
You may only use this
A copy of the license can be found at http://www.spread.org/license
This product uses software developed by Spread Concepts LLC for use
in the Spread toolkit. For more information about Spread,
see http://www.spread.org
This software is distributed on an "AS IS" basis, WITHOUT WARRANTY OF
ANY KIND, either express or implied.
Creators:
    Yair Amir            yairamir@cs.jhu.edu
    Michal Miskin-Amir   michal@spreadconcepts.com
    Jonathan Stanton     jstanton@gwu.edu
    John Schultz         jschultz@spreadconcepts.com
```

Step 29 From the Close Virtual Media dialog box, click **Yes** to confirm the unmapping of the image file. This unmounts the ISO image.

Figure 6-23 Close Virtual Media

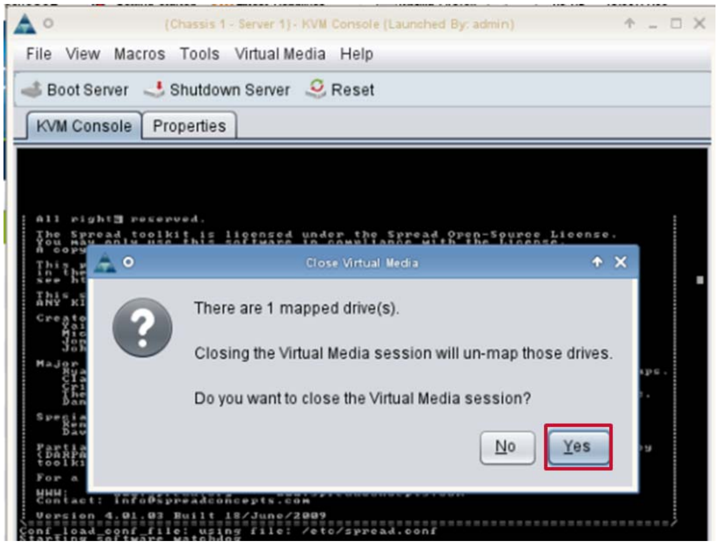

There are 1 mapped drive(s).

Closing the Virtual Media session will un-map those drives.

Do you want to close the Virtual Media session?

```
All rights reserved.
The Spread toolkit is licensed under the Spread Open-Source License.
Contact: info@spreadconcepts.com
Version 4.01.03 Built 18/June/2009
Conf_load_conf_file: using file: /etc/spread.conf
```

Step 30 Click **Reset** under the menu bar to reboot the cache engine.

Figure 6-24 *Reset the Cache Engine*

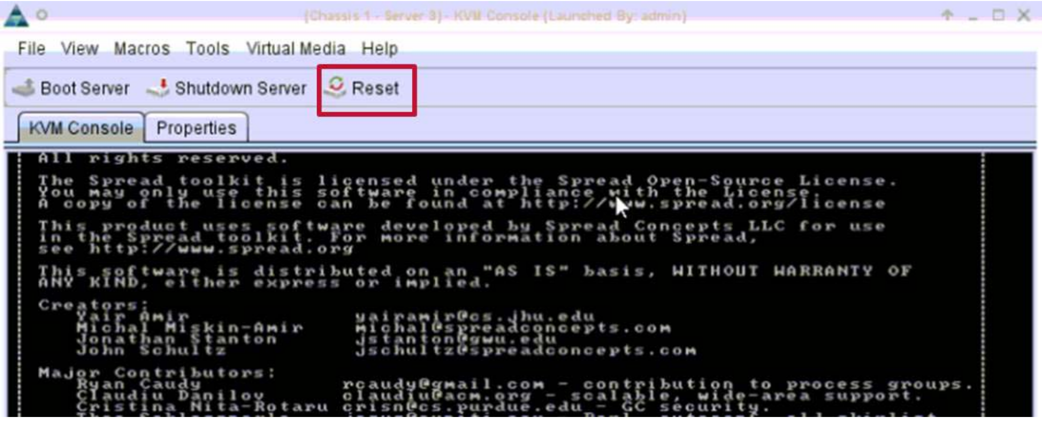
[Chassis 1 - Server 3] - KVM Console (Launched By: admin)
File View Macros Tools Virtual Media Help
Boot Server
Shutdown Server
Reset
KVM Console
Properties

```
All rights reserved.

The Spread toolkit is licensed under the Spread Open-Source License.
You may only use this software in compliance with the License.
A copy of the license can be found at http://www.spread.org/license

This product uses software developed by Spread Concepts LLC for use
in the Spread toolkit. For more information about Spread,
see http://www.spread.org

This software is distributed on an "AS IS" basis, WITHOUT WARRANTY OF
ANY KIND, either express or implied.

Creators:
    Yair Amir              yairamir@cs.jhu.edu
    Michal Miskin-Amir     michal@spreadconcepts.com
    Jonathan Stanton       jstanton@gwu.edu
    John Schultz           jschultz@spreadconcepts.com

Major Contributors:
    Ryan Caudy            rcaudy@gmail.com - contribution to process groups.
    Claudiu Danilov       claudiu@acm.org - scalable, wide-area support.
    Cristina Nita-Rotaru crisn@cs.purdue.edu - GC security.
```

Step 31 From the Reset Server Warning window that appears, click **OK** to reset the server.

Figure 6-25 *Reset Server Warning Window*

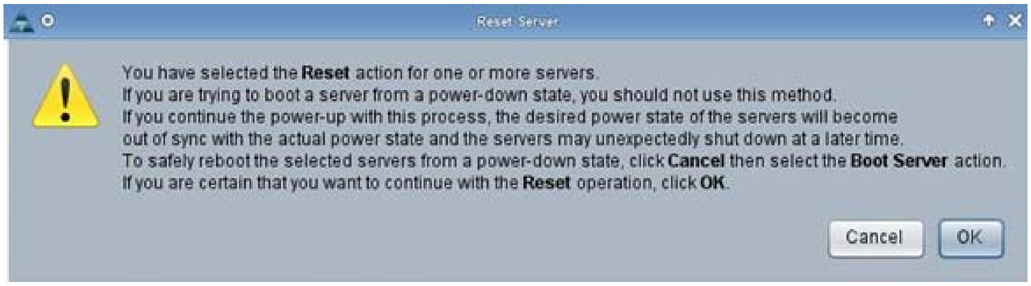

You have selected the **Reset** action for one or more servers.
If you are trying to boot a server from a power-down state, you should not use this method.
If you continue the power-up with this process, the desired power state of the servers will become
out of sync with the actual power state and the servers may unexpectedly shut down at a later time.
To safely reboot the selected servers from a power-down state, click **Cancel** then select the **Boot Server** action.
If you are certain that you want to continue with the **Reset** operation, click **OK**.

Step 32 From the Reset Server Service window that appears, choose **Power Cycle** and click **OK**.

Figure 6-26 *Reset Server Service*

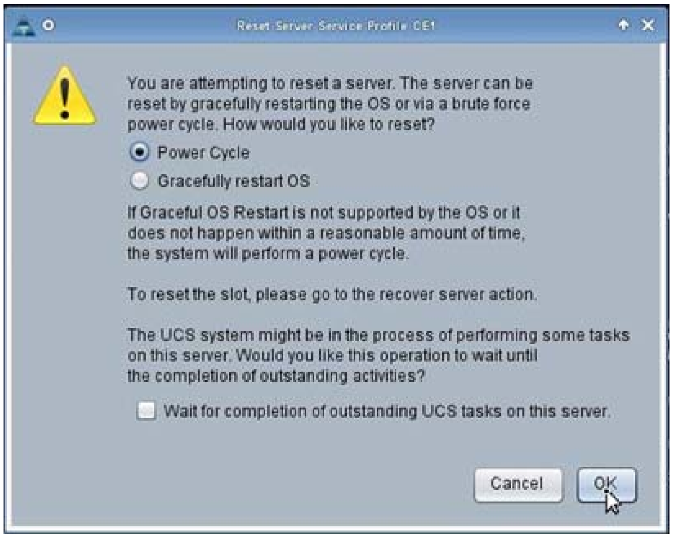

You are attempting to reset a server. The server can be reset by gracefully restarting the OS or via a brute force power cycle. How would you like to reset?

● Power Cycle

○ Gracefully restart OS

If Graceful OS Restart is not supported by the OS or it does not happen within a reasonable amount of time, the system will perform a power cycle.

To reset the slot, please go to the recover server action.

The UCS system might be in the process of performing some tasks on this server. Would you like this operation to wait until the completion of outstanding activities?

□ Wait for completion of outstanding UCS tasks on this server.

Step 33 After the cache engine reboots, from the KVM Console window on the KVM tab, log in using the **root** user and the password that was provided by Cisco.

Step 34 Enter the command **cd /opt/pang/utilities/CE/Multipath_Driver**

Step 35 Enter the command **tar -zxvf rdac-LINUX-09.03.0C05.0652-source.tar.gz**

Step 36 Enter the command **cd linuxrdac-09.03.0C05.0652**

Step 37 Enter the command **make clean**

Step 38 Enter the command **make**

Step 39 Enter the command **make install**

Step 40 If you are prompted about whether the new MPP driver should manage the iSCSI storages, answer yes to continue.

Step 41 Enter the command **reboot -f** to reboot the cache engine.

Step 42 After the cache engine reboots, log back into the KVM console as the root user, using the password that was provided by Cisco.

Step 43 To verify that the installation of the multipath driver was successful, enter the command **chkconfig mpp**. If the installation was successful, this command should return the following output:

```
mpp on
```

Step 44 From the KVM console connection to the Cache Engine, enter the following commands to temporarily update the network interface:

a. **vconfig add eth0 50**

b. **ifconfig eth0 10.11.12.*X*/24**
where *X* is the number of the cache engine that you are replacing plus 1. For example, if you are replacing cache engine 3, you would enter **ifconfig eth0 10.11.12.4/24**.

Note Please pay close attention to the IP address that you are assigning. If it is not entered correctly, it can cause a problem replacing the cache engine.

Step 45 From the KVM console connection to the Cache Engine, ping the IP address of the management server. If this ping fails, repeat Step 44. If it fails again, contact your Cisco Support Engineer.

Updating the SSH Keys and Connecting the Cache Engine to the Storage Enclosure

After completing the “Running the Post Installation Scripts” section in the “Cache Engine Servers Installation” chapter, you must update the grid SSH keys on the VDS TC management server and configure the new cache engine to connect to the storage enclosure. All of these steps are performed on the VDS TC management server.

Follow this procedure to perform these steps on the VDS TC management server:

Note The following steps are performed on the VDS TC management server.

Procedure

Step 1 On the PBR routers, stop redirecting traffic to the VDS TC solution.

Step 2 Close any open VDS TC Manager windows and ensure that no other administrators are connected to the VDS TC Manager before proceeding.

Step 3 Using SSH software, such as Putty, open an SSH connection to the IP address of the VDS TC management server.

Step 4 Log into the system using the username **padmin** and the password provided by Cisco.

Step 5 Enter the command **su admin** to log into the CLI of the VDS TC management server. The default password is the serial number of the VDS TC management server.

Step 6 From the VDS TC prompt, enter the **enable** command. When prompted, enter the Enable mode password and press **Enter**. You are now logged into Enable mode and the Enable prompt, console#, should appear.

Step 7 From the Enable mode prompt, enter the command **oper service stop** to stop the caching service.

Step 8 Wait a few minutes and then enter the command **show status**. Check to see if the Device State shows "stopped." For example:

```
console# show status
Operational state Device state Administrative state
disabledstoppedunlocked
```

Caution Do not proceed to the next step until the Device State shows "stopped." You may need to repeat the **show status** command several times before you see this status. Wait several minutes between executions of the **show status** command.

Step 9 Enter the **exit** command to exit Enabled mode.

Step 10 Enter the **exit** command to exit the VDS TC CLI.

Step 11 Enter the command **su root** to change to the root user. Enter the password for the root user when prompted. This password is provided by Cisco.

Step 12 Enter the command **cd /opt/pang/useful/**.

Step 13 To reset the SSH keys, enter the command **./replace_server_keys.sh -servers** #, where # is the total number of Blade Server cache engines in the VDS TC installation. For example, if you are replacing a cache engine in a VDS TC solution that has 16 cache engines, you would enter the command **./replace_server_keys.sh -servers 16**.

Step 14 Enter the command **cd /opt/pang/useful/configure_storages** to change to the storage configuration script directory.

Step 15 Enter the command **./configure_storages.py -s** <*no_of_storage_enclosures*> **-b** <*no_of_blades*> **-c** <*the_replaced_CE-ID_number*>. For example, to replace CE-2 in a VDS TC Blade Server installation that has two cache engines, and one storage enclosure, enter the command **./configure_storages.py -s 1 -b 2 -c 2**.

Note The configure storages script will format the data disks. To configure the storages without formatting the disks enter the command **./configure_storages.py -s** <*no_of_storage_enclosures*> **-b** <*no_of_blades_in_grid*> **-c** <*the_replaced_CE-ID_number*> **-x**.

Step 16 When the warning message about the disk format appears, enter **y** to continue the storages configuration.

```
*********** WARNING! **********
All storages are about to be reconfigured. All data on the data disks will
be lost.
Do you want to continue? (y/n)y
```

Step 17 Follow the on screen instructions. If prompted, enter the password for the root user that was provided by Cisco.

Note You may be prompted to enter the password several times.

When the script is finished running, you should see the word “COMPLETE”. After running this script, the new VDS TC cache engine will be connected to the storage enclosures.

Step 18 Use SSH to connect to each CE and enter the command **iscsiadm -m session** to check the connections:

```
ssh ce-<cache engine number>
iscsiadm -m session
```

If you are using 1Gb connections to the storage enclosures, your output should be similar to the following:

Note The following output will look different if you are using 10Gb connections to the storage enclosures.

```
ƒÞ 8 * sessions per storage, for example for 1 storage:
tcp: [17] 10.11.14.100:3260,1 iqn.1992-
01.com.Avago:2365.60080e50002f4e82000000005063c74a
tcp: [18] 10.11.14.101:3260,2 iqn.1992-
01.com.Avago:2365.60080e50002f4e82000000005063c74a
tcp: [19] 10.11.16.100:3260,1 iqn.1992-
01.com.Avago:2365.60080e50002f4e82000000005063c74a
tcp: [20] 10.11.16.101:3260,2 iqn.1992-
01.com.Avago:2365.60080e50002f4e82000000005063c74a
tcp: [21] 10.11.15.100:3260,1 iqn.1992-
01.com.Avago:2365.60080e50002f4e82000000005063c74a
tcp: [22] 10.11.15.101:3260,2 iqn.1992-
01.com.Avago:2365.60080e50002f4e82000000005063c74a
tcp: [23] 10.11.17.101:3260,2 iqn.1992-
01.com.Avago:2365.60080e50002f4e82000000005063c74a
tcp: [24] 10.11.17.100:3260,1 iqn.1992-
01.com.Avago:2365.60080e50002f4e82000000005063c74a
```

Updating the NetApp Storage Enclosure

Note This process is only for the NetApp storage enclosure. If your system uses an IBM storage enclosure, refer to the Updating the IBM Storage Enclosure section.

After you have configured the VDS TC cache engine to connect with the storage enclosure, you must update the host port identifiers on the NetApp storage enclosure for the new cache engine. Follow this procedure to perform this task:

Procedure

Step 1 From the SSH connection to the VDS TC management server, enter the following commands to start the VNC server:

a. **rm /tmp/.X1-lock** (This step is optional. It is only required if there is a leftover lock file from a previous vncserver instance.)

b. **rm /tmp/.X11-unix/X1** (This step is optional. It is only required if there is a leftover lock file from a previous vncserver instance.)

c. **vncserver :1**

Step 2 Use a VNC client to connect from your local computer to the VDS TC management server.

Figure 6-27 VNC Viewer

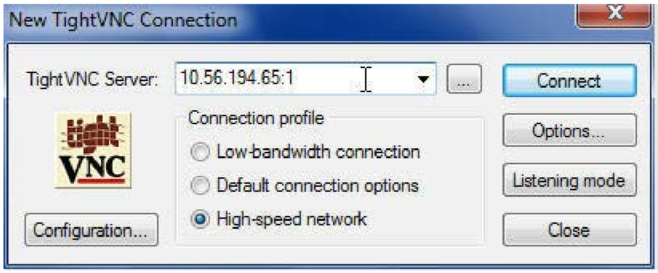

New TightVNC Connection

TightVNC Server: 10.56.194.65:1

Connection profile

○ Low-bandwidth connection
○ Default connection options
● High-speed network

Step 3 From the VNC console window, enter the command **cd /opt/SMgr/client** to change folders.

Step 4 Enter the command **./SMclient** to start the SANtricity ™ ES Storage Manager software.

Note You must perform the following steps on all of the storage subsystems.

Step 5 Right-click the storage subsystem that you are configuring and choose **Manage Storage Array**. Click **No** in any pop-up windows that appear.

Figure 6-28 SANtricity Storage Manager

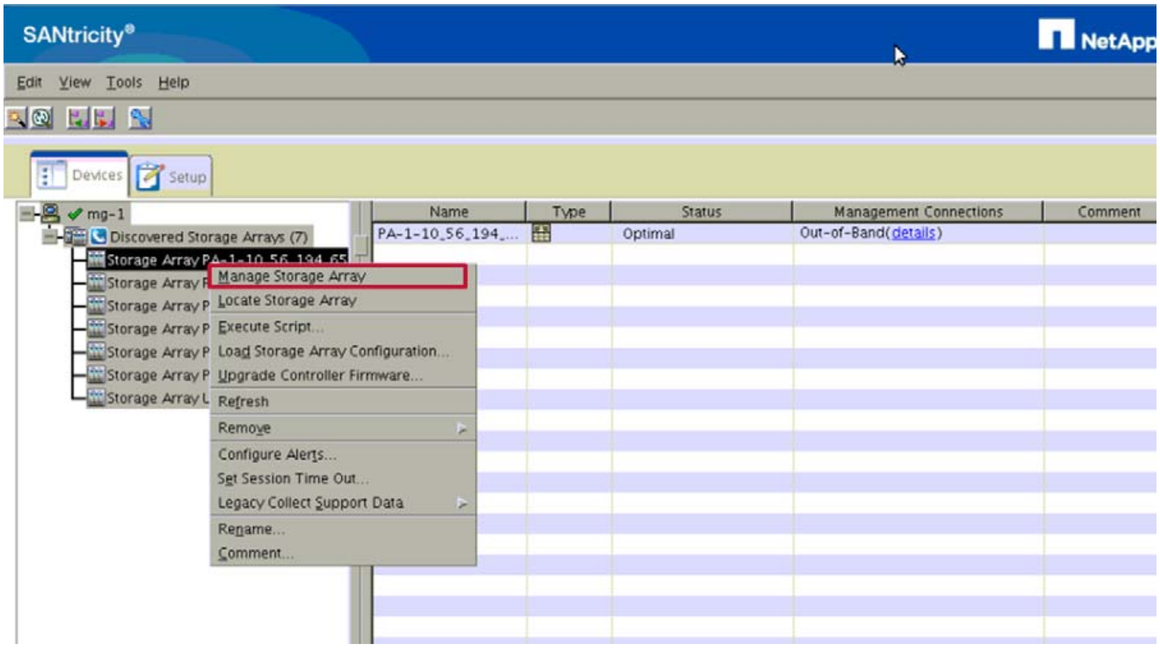

Name	Type	Status	Management Connections	Comment
PA-1-10_56_194_...		Optimal	Out-of-Band(details)	

Step 6 In the Subsystem Management window that appears, click the **Host Mappings** tab.

Step 7 In the navigation pane on the left, expand the **Host Group** folder.

Step 8 Right-click the cache engine that was replaced, for example Host CE-1, and choose **Manage Host Port Identifiers** from the pop-up menu.

Figure 6-29 Manage Host Port Identifiers

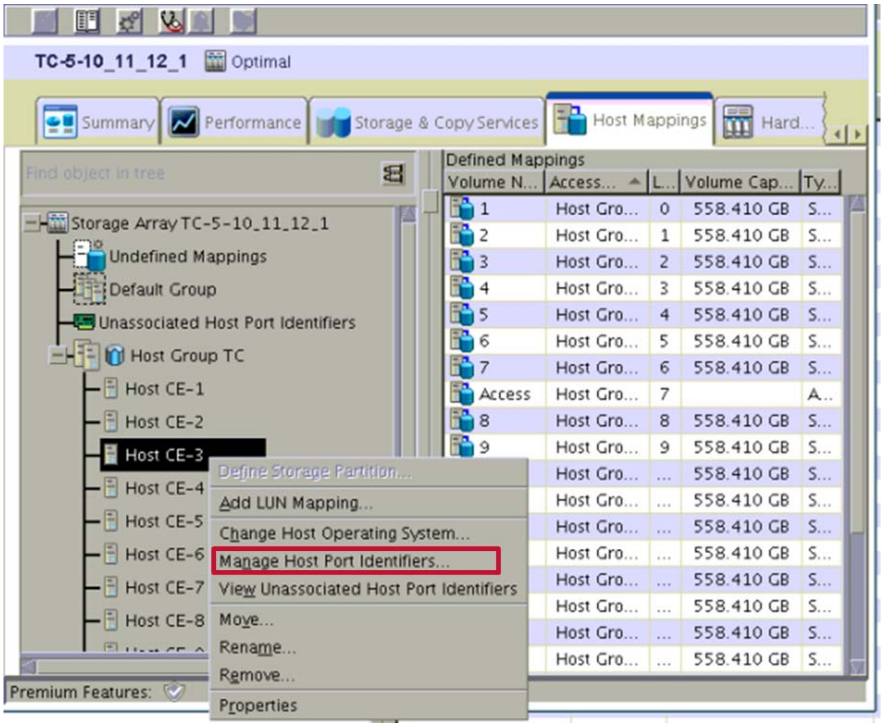

Defined Mappings

Volume N...	Access...	L...	Volume Cap...	Ty...
1	Host Gro...	0	558.410 GB	S...
2	Host Gro...	1	558.410 GB	S...
3	Host Gro...	2	558.410 GB	S...
4	Host Gro...	3	558.410 GB	S...
5	Host Gro...	4	558.410 GB	S...
6	Host Gro...	5	558.410 GB	S...
7	Host Gro...	6	558.410 GB	S...
Access	Host Gro...	7		A...
8	Host Gro...	8	558.410 GB	S...
9	Host Gro...	9	558.410 GB	S...
	Host Gro...	...	558.410 GB	S...
	Host Gro...	...	558.410 GB	S...
	Host Gro...	...	558.410 GB	S...
	Host Gro...	...	558.410 GB	S...
	Host Gro...	...	558.410 GB	S...
	Host Gro...	...	558.410 GB	S...
	Host Gro...	...	558.410 GB	S...
	Host Gro...	...	558.410 GB	S...

Step 9 From the Manage Host Port Identifiers window that appears, in the Host Port Identifier Information table, ensure that the old iqn is selected and click **Replace**.

Figure 6-30 Replace IQN

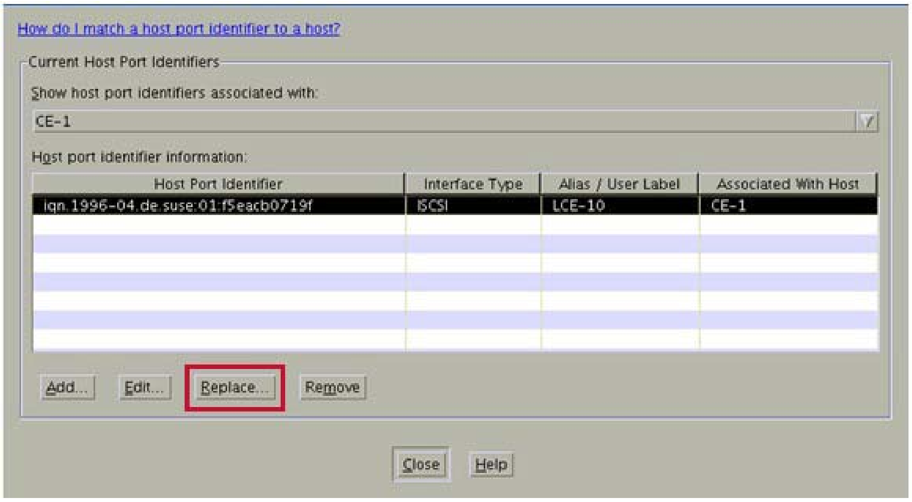
How do I match a host port identifier to a host?
Current Host Port Identifiers
Add...
Edit...
Replace...
Remove
Close
Help

Show host port identifiers associated with: CE-1

Host port identifier information:

Host Port Identifier	Interface Type	Alias / User Label	Associated With Host
iqn.1996-04.de.suse:01:f5eacb0719f	ISCSI	LCE-10	CE-1

Step 10 From the Replace Host Port Identifier window that appears, configure the following:

- From the Choose a Host Interface Type drop-down list, choose **ISCSI**.
- Click the **Replace by Selecting a Known Unassociated Host Port Identifier** radio button.

- From the Known Unassociated Host Port Identifier drop-down list, choose the iqn for the *newly* installed cache engine. (It should be the only one in the list.)

Step 11 Click **Replace**.

Figure 6-31 Replace Host Port Identifier

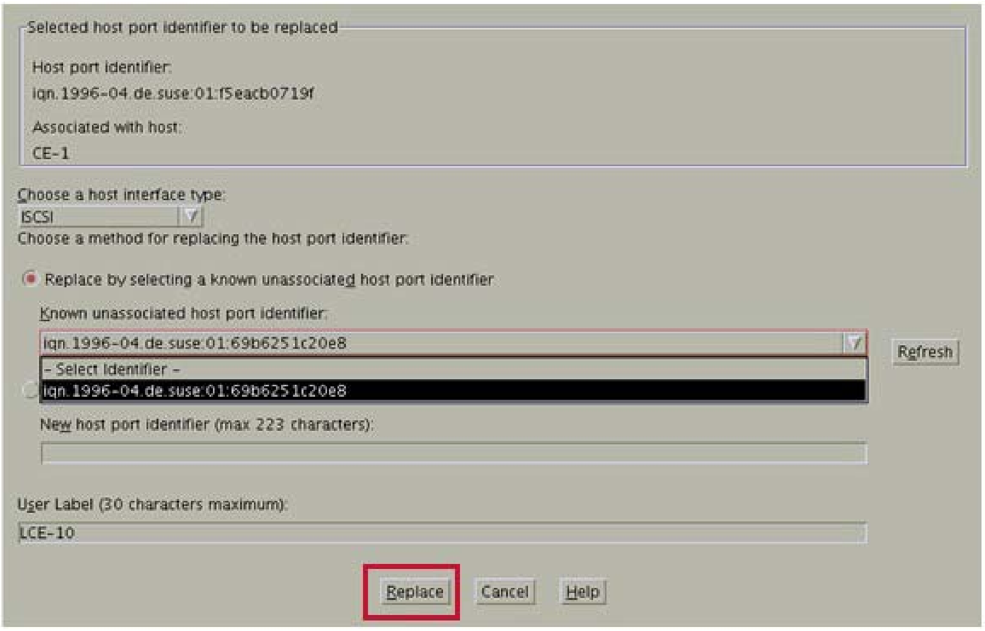

Selected host port identifier to be replaced

Host port identifier:
iqn.1996-04.de.suse:01:f5eacb0719f

Associated with host:
CE-1

Choose a host interface type:
ISCSI

Choose a method for replacing the host port identifier:

● Replace by selecting a known unassociated host port identifier

Known unassociated host port identifier:
iqn.1996-04.de.suse:01:69b6251c20e8
- Select Identifier -
iqn.1996-04.de.suse:01:69b6251c20e8

New host port identifier (max 223 characters):

User Label (30 characters maximum):
LCE-10

Step 12 From the Manage Host Port Identifiers window, click **Close**.

Step 13 Repeat Step 5 through Step 12 on any remaining storage subsystems.

Note Do not continue until you have completed Step 13.

Step 14 Close the Subsystem Management window and exit the SANtricity Storage Manager.

Step 15 Close the VNC client.

Updating the IBM Storage Enclosure

Note This process is only for the IBM storage enclosure. If your system uses the NetApp E2724 storage enclosure, refer to the Updating the NetApp Storage Enclosure section.

After you have configured the VDS TC cache engine to connect with the storage enclosure, you must update the host port identifiers on the IBM storage enclosure for the new cache engine. Follow this procedure to perform this task:

Procedure

Step 1 From the SSH connection to the VDS TC management server, enter the following commands to start the VNC server:

a. **rm /tmp/.X1-lock** (This step is optional. It is only required if there is a leftover lock file from a previous vncserver instance.)

b. **rm /tmp/.X11-unix/X1** (This step is optional. It is only required if there is a leftover lock file from a previous vncserver instance.)

c. **vncserver :1**

Step 2 Use a VNC client to connect from your local computer to the VDS TC management server.

Figure 6-32 VNC Viewer

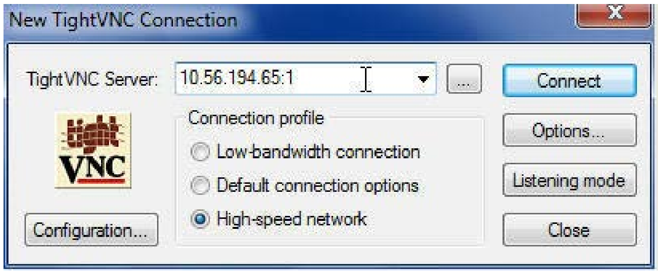

New TightVNC Connection

TightVNC Server: 10.56.194.65:1

Connection profile

○ Low-bandwidth connection
○ Default connection options
● High-speed network

Step 3 From the VNC console window, enter the command **/opt/IBM_DS/client/SMclient** to start the IBM Storage Manager.

Note You must perform the following steps on all of the storage subsystems.

Step 4 Right-click the storage subsystem that you are configuring and choose **Manage Storage Subsystem**. Click **No** in any pop-up windows that appear.

Figure 6-33 IBM DS Storage Manager

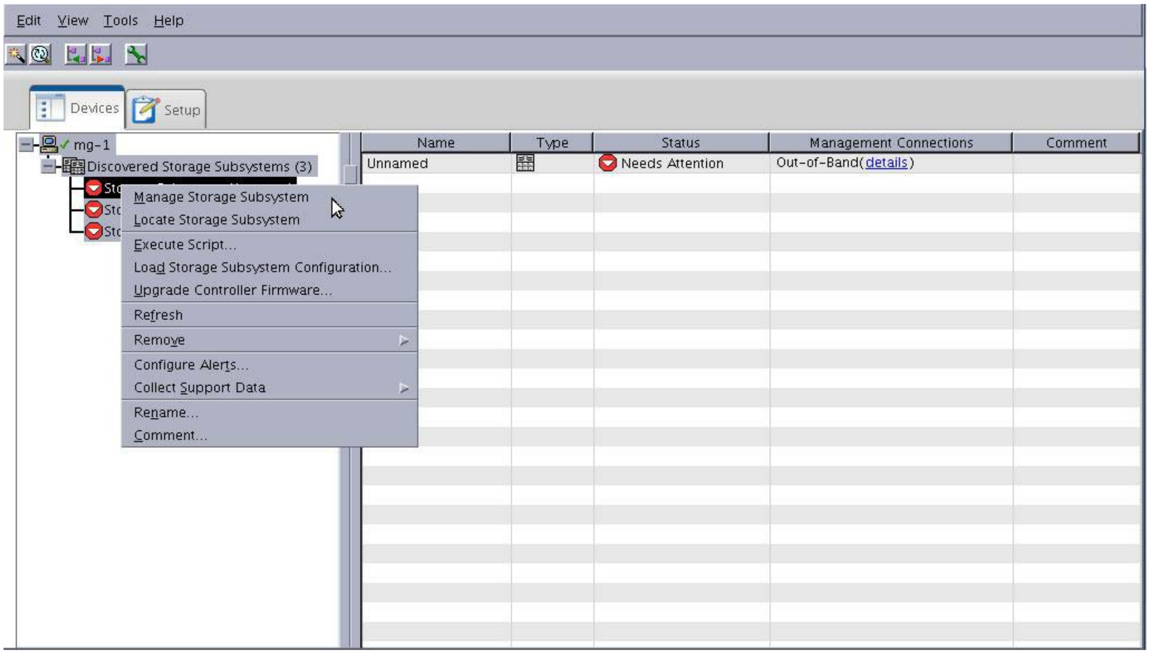

Name	Type	Status	Management Connections	Comment
Unnamed	▦	Needs Attention	Out-of-Band(details)	

Step 5 In the Subsystem Management window that appears, click the **Host Mappings** tab.

Step 6 In the navigation pane on the left expand the **Host Group** folder.

Step 7 Right-click the cache engine that was replaced, for example Host CE-2, and choose **Manage Host Port Identifiers** from the pop-up menu.

Figure 6-34 Manage Host Port Identifiers

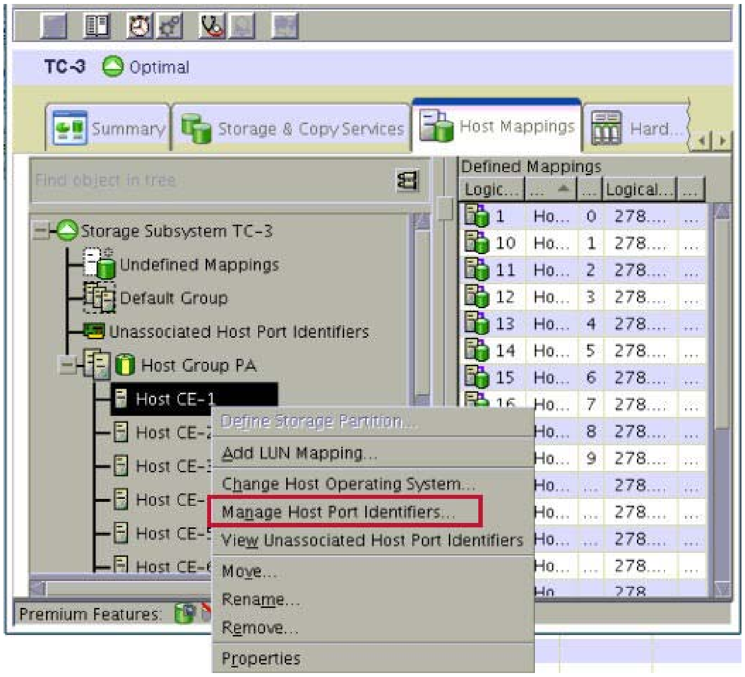

Logic...	...	...	Logical...	...
1	Ho...	0	278....	...
10	Ho...	1	278....	...
11	Ho...	2	278....	...
12	Ho...	3	278....	...
13	Ho...	4	278....	...
14	Ho...	5	278....	...
15	Ho...	6	278....	...
16	Ho...	7	278....	...
	Ho...	8	278....	...
	Ho...	9	278....	...
	Ho...	...	278....	...
	Ho...	...	278....	...
	Ho...	...	278....	...
	Ho...	...	278....	...
	Ho		278	

Step 8 From the Manage Host Port Identifiers window that appears, in the Host Port Identifier Information table, ensure that the iqn is selected and click **Replace**.

Figure 6-35 Replace IQN

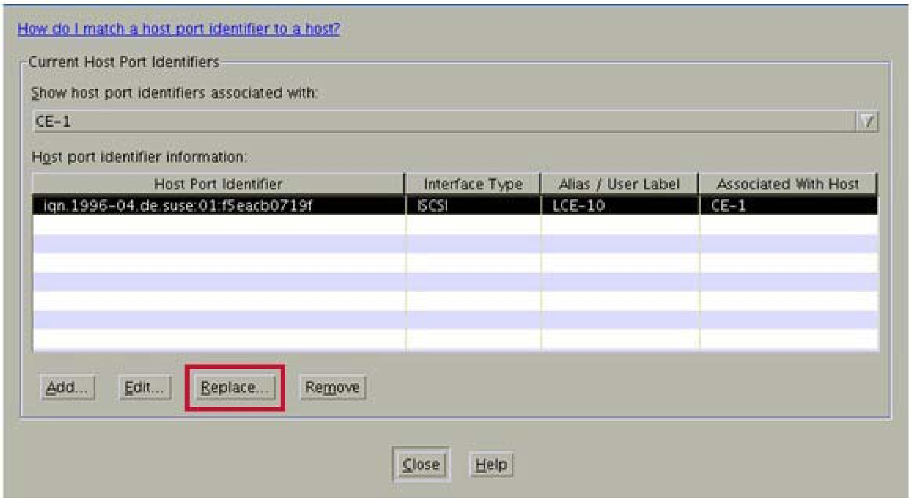

How do I match a host port identifier to a host?

Current Host Port Identifiers

Show host port identifiers associated with: CE-1

Host port identifier information:

Host Port Identifier	Interface Type	Alias / User Label	Associated With Host
iqn.1996-04.de.suse:01:f5eacb0719f	ISCSI	LCE-10	CE-1

Step 9 From the Replace Host Port Identifier window that appears, configure the following:

- From the Choose a Host Interface Type drop-down list, choose **ISCSI**.
- Click the **Replace by Selecting a Known Unassociated Host Port Identifier** radio button.
- From the Known Unassociated Host Port Identifier drop-down list, choose the iqn for the newly installed cache engine. (It should be the only one in the list.)

Step 10 Click **Replace**.

Figure 6-36 Replace Host Port Identifier

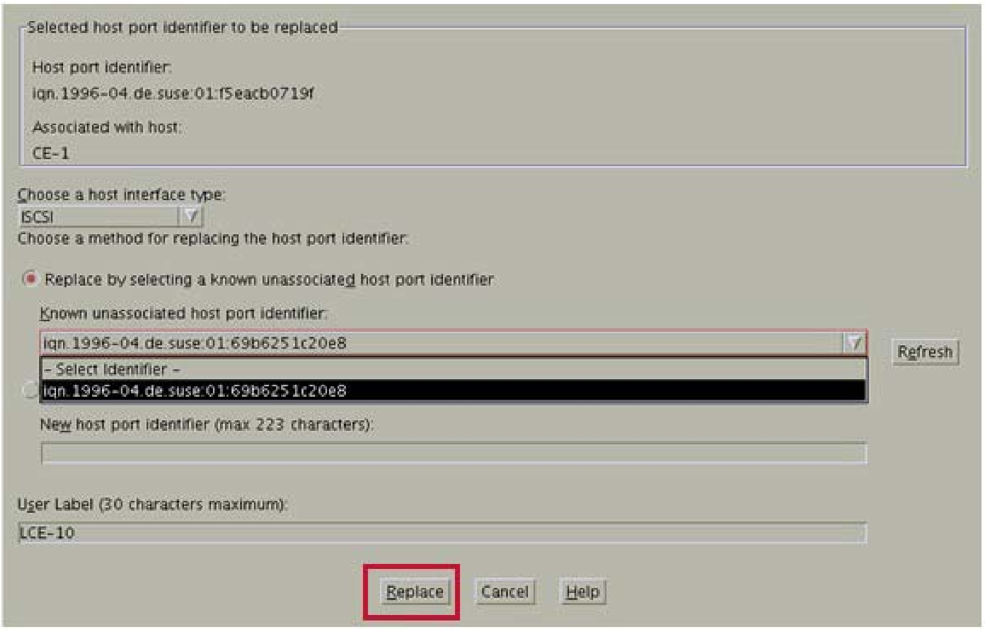

Selected host port identifier to be replaced

Host port identifier:
iqn.1996-04.de.suse:01:f5eacb0719f

Associated with host:
CE-1

Choose a host interface type:
ISCSI

Choose a method for replacing the host port identifier:

● Replace by selecting a known unassociated host port identifier

Known unassociated host port identifier:
iqn.1996-04.de.suse:01:69b6251c20e8

- Select Identifier -
iqn.1996-04.de.suse:01:69b6251c20e8

New host port identifier (max 223 characters):

User Label (30 characters maximum):
LCE-10

Step 11 From the Manage Host Port Identifiers window, click **Close**.

Step 12 Repeat Step 4 through Step 11 on any remaining storage subsystems.

Note Do not continue until you have completed Step 12.

Step 13 Close the Subsystem Management window and exit the IBM Storage Manager.

Step 14 Close the VNC client.

Finishing the Cache Engine Replacement

Follow this procedure to finish configuring and installing the new cache engine:

Note You will be re-running the storage script that you ran in the Updating the SSH Keys and Connecting the Cache Engine to the Storage Enclosure section.

Note You must be logged in with root permissions to perform the following steps.

Procedure

Step 1 From the SSH connection to the VDS TC management server, enter the command **cd /opt/pang/useful/configure_storages** to change to the storage configuration script directory.

Step 2 Enter the command **./configure_storages.py -s** *no_of_storage_enclosures* **-b** *no_of_blades* **-c** *the_replaced_CE-ID_number*. For example, to replace CE-2 in a VDS TC Blade Server installation that has two cache engines, and one storage enclosure, enter the command **./configure_storages.py -s 1 -b 2 -c 2**.

Step 3 When the warning message about the disk format appears, enter **y** to continue the storages configuration.

```
*********** WARNING! **********
All storages are about to be reconfigured. All data on the data disks will
be lost.
Do you want to continue? (y/n)y
```

Step 4 Follow the on screen instructions. If prompted, enter the password for the root user that was provided by Cisco. You should see the word “COMPLETE” when the script is finished running.

Note You may be prompted to enter the password several times.

Note Ensure that the VDS-TC_FI_5.7.3b54_Server_Cluster.tar.gz file and the new VDS TC license that you received from Cisco are located on the VDS TC management server in the /tftpboot folder before performing these steps.

Step 5 Enter the command **su admin** to log into the VDS TC CLI. The CLI prompt console> appears.

Step 6 Enter the command **enable** to enter Enable mode. At the password prompt that appears, enter the enable password. The CLI prompt console# appears.

Step 7 Upgrade the software on the new cache engine by entering the command **upgrade server** *CE_number* **127.0.0.1 VDS-TC_FI_5.7.3b54_Server_Cluster.tar.gz**, where *CE_number* is the cache engine you replaced.

Step 8 To import the new license that includes the serial number of the new cache engine, enter the command **license import 127.0.0.1** *filename*, where *filename* is the name of the new license file.

Note You must be in Enable mode, *not* Configuration mode to import a new license.

Step 9 Enter the command **license activate** to apply and activate the license.

Step 10 To start the application on the replaced cache engine, enter the command **oper server** *#*, where *#* is the number of the cache engine that was replaced. For example, if you replaced CE-2, enter the command **oper server 2**. You should see the following output:

```
console# oper server 2
oper server 2# start
Starting server 2
service is started on server 2
```

Step 11 After the application has started on the new cache engine, enter the command **oper service start** to start the VDS TC service.

Step 12 Wait a few minutes and then enter the command **show status**. Do not proceed to the next step until you see a Device Status of “Started.” For example:

```
console# show status
Operational state Device state Administrative state
enabled started unlocked
```

Step 13 On the PBR routers, start redirecting traffic to the VDS TC solution.

CHAPTER 7

Management Server Replacement Procedure for a System with a Fresh Install of VDS TC 5.7.3

Note Only use this chapter if you are replacing a Management Server in a VDS TC system that had a fresh install of VDS TC 5.7.3 and was *not* upgraded to 5.7.3. If you are replacing a Management Server in a VDS TC system that was upgraded to VDS TC 5.7.3 from VDS TC 5.2.1, refer to Chapter 8, "Management Server Replacement Procedure for a VDS TC 5.7.3 System Upgraded from VDS TC 5.2.1". If you are replacing a Management Server in a VDS TC that was upgraded to VDS TC 5.7.3 from VDS TC 5.6.1, refer to Chapter 9, "Management Server Replacement Procedure for a VDS TC 5.7.3 System Upgraded from VDS TC 5.6.1".

This chapter discusses how to replace a VDS TC management server. It contains the following sections:

- Prerequisites
- Replacing the VDS TC Management Server
- Post Installation

Prerequisites

Before you replace the VDS TC management server, you must obtain the following:

- The VDS TC 5.7.3 ISO image

Note For more information refer to the "Cisco Videoscape Distribution Suite Transparent Caching Software C-Series Cluster Installation Guide" or the "Cisco Videoscape Distribution Suite Transparent Caching Software Blade Server Cluster Installation Guide".

Backup Files

Before you perform a VDS TC management server replacement, you need to backup the PALicense.xml and cluster_conf.xml files. Follow these steps to backup these files:

Note **Important:** The following procedure for backing up files apply only if the VDS TC management server is still accessible. If the management server is not accessible, please contact Cisco support for the VDS License and cluster_conf.xml files.

Step 1 Using SFTP software, such as WinSCP, connect to the VDS TC management server.

Step 2 Copy the following files to a location other than the VDS TC management server:

- /opt/pang/mgmt/config/PALicense.xml
- /opt/pang/mgmt/config/cluster_conf.xml

Replacing the VDS TC Management Server

Procedure

Step 1 Disconnect the old VDS TC management server and replace it with the new management server. Next, connect the new management server to the cache engines. For more information, refer to Chapter 1 “Prerequisites and C-Series Cluster Physical Installation” in the “Cisco Videoscape Distribution Suite Transparent Caching Software C-Series Cluster Installation Guide” for a C-Series Cluster installation, and Chapter 1 “Prerequisites and Blade Server Cluster Physical Installation” chapter in the “Cisco Videoscape Distribution Suite Transparent Caching Software Blade Server Cluster Installation Guide” for a Blade Server cluster installation.

Step 2 Use one of the following chapters in the “Cisco Videoscape Distribution Suite Transparent Caching Software C-Series Cluster Installation Guide” or “Cisco Videoscape Distribution Suite Transparent Caching Software Blade Server Cluster Installation Guide” to configure the BIOS and operating system of the management server and install the ISO image:

- For a VDS TC C-Series cluster installation, follow the steps in Chapter 3“Management Server Installation” in the “Cisco Videoscape Distribution Suite Transparent Caching Software C-Series Cluster Installation Guide.”
- For a VDS TC Blade Server cluster installation, follow the steps in Chapter 3 “Management Server Installation” in the “Cisco Videoscape Distribution Suite Transparent Caching Software Blade Server Cluster Installation Guide.”

Note At this point, perform *only* the steps in Chapter 3 of the VDS TC Installation Guides. Do *not* configure the cache storage or install the VDS TC cluster software until instructed to do so in later steps.

Note When prompted, use the old VDS TC management server Network parameters and address values.

Step 3 Using SSH software, such as Putty, open an SSH connection to the management IP address of the VDS TC management server. Log in as a root user.

Step 4 To update the cluster SSH keys, enter the following commands:

- **cd /opt/pang/useful/**

- **./replace_server_keys.sh -servers** <*no_ of_ces*>, where *no_of_ces* is the number of cache engines in the cluster. For example, for a cluster installation that has four cache engines, enter the command **./replace_server_keys.sh -servers 4**.

Step 5 Using either the "Cisco Videoscape Distribution Suite Transparent Caching Software C-Series Cluster Installation Guide" or the "Cisco Videoscape Distribution Suite Transparent Caching Software Blade Server Cluster Installation Guide", install and configure the Storage Manager using the following sections:

- For a VDS TC C-Series cluster installation:
  - Perform the steps in the *Installing the Storage Manager on the Management Server* section of Chapter 5 "Configuring the Cache Storage".
  - Perform Step 1 through Step 6 in the *Configuring the First Storage Array* section of Chapter 5"Configuring the Cache Storage". In Step 6, choose **Automatic** for the method for the addition of the storage subsytem.
- For a VDS TC Blade Server cluster installation:
  - Perform the steps in the *Installing the Storage Manager on the Management Server* section of Chapter 5 "Configuring the Cache Storage".
  - Perform Step 1 through Step 6 in the *Configuring the First Storage Array* section of Chapter 5 "Configuring the Cache Storage". In Step 6, choose **Automatic** for the method for the addition of the storage subsytem.

Step 6 From the management server, enter the command **scp root@ce-1:/etc/hosts /etc/hosts**.

Step 7 Obtain a copy of the VDS TC Cluster software package (FI) with the same version number as the Cache Engines.

Step 8 From the management server run the following commands:

- **cd /tftpboot/**
- **mkdir Installer_ver5.7.3**
- **cd Installer_ver5.7.3**

Step 9 Back on the management server, enter the command ../**tar -zxvf VDS-TC_FI_5.7.3b54_Server_Cluster.tar.gz** to extract the files.

Note If you are connected remotely to the management server, use the **screen** command. This ensures that if the session is disconnected during the installation as result of a Management IP set operation, you will be able to reconnect and resume the session with the **screen -r** command.

Step 10 Edit the GA_installer.rc file to set the "UPGRADE CACHE SOFTWARE" parameter to "none" and the GRACE_UPGRADE parameter to "yes". Use the command vi GA_installer.rc to edit this file.

For example in a VDS TC 16 CE solution:

vi GA_installer.rc

```
#MANAGEMENT_SOFTWARE {no, yes}
MANAGEMENT_SOFTWARE=yes

#UPGRADEL_CACHE_SOFTWARE {all, none, space separated list(ce-1 ce-2 ce-3
ce-4 ce-5 ce-6 ce-7 ce-8 ce-9 ce-10 ce-11 ce-12 ce-13 ce-14 ce-15 ce-16)}
UPGRADEL_CACHE_SOFTWARE=none

UPGRADE_TAR_BALL=<filename>.tar.gz
MANAGEMENT_ENVIRONMENT=yes
```

```
GRACE_UPGRADE=no
FIRST_INSTALL=yes
```

Step 11 After you are finished editing the GA_installer.rc file, in the vi editor enter **:x** and press **Enter** to save and exit the file.

Step 12 Next, perform the following steps to run the ./GA_installer.sh script:

a. Enter the command **./GA_installer.sh** to run the installation script.

b. At the Enter primary name server IP prompt, enter the IP address of your primary DNS server. Enter **y** to confirm the entry.

Note If a DNS server is unavailable, enter 0.0.0.0 as the IP address of the DNS server to prevent system functionality from being negatively affected.

c. At the Enter secondary name server IP prompt, enter the IP address of your secondary DNS server. Enter **y** to confirm the entry.

Note If a secondary DNS server is unavailable, leave this entry blank.

Figure 7-1 Installation Script Dialog

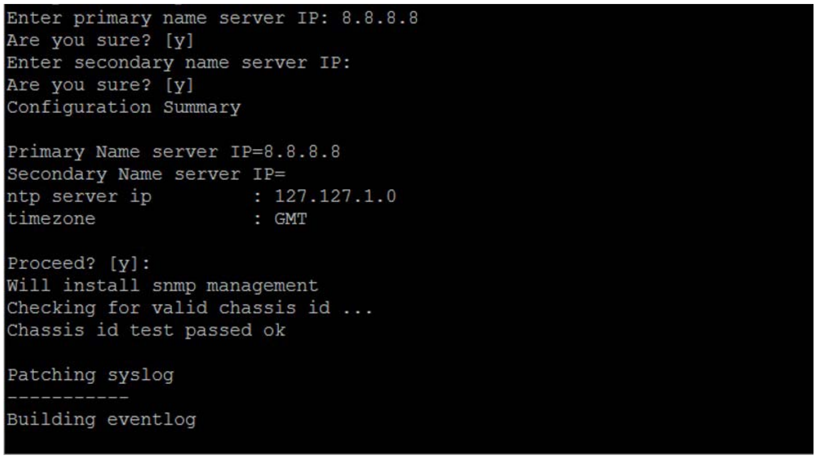

```
Enter primary name server IP: 8.8.8.8
Are you sure? [y]
Enter secondary name server IP:
Are you sure? [y]
Configuration Summary

Primary Name server IP=8.8.8.8
Secondary Name server IP=
ntp server ip         : 127.127.1.0
timezone              : GMT

Proceed? [y]:
Will install snmp management
Checking for valid chassis id ...
Chassis id test passed ok

Patching syslog
-----------
Building eventlog
```

d. The GA Installer script displays a summary of the configuration. Review this configuration. If the configuration is correct, press **Y** to proceed. If the configuration is not correct, press **N** to exit the installation and start over.

e. When asked if you want to change the management IPv4 settings, enter **N** unless you need to change the IP address information that was configured during the original VDS TC Management Server installation. If you enter **Y**, enter information at the following prompts:

- Management IPv4 address ():
- Management IPv4 netmask ():
- Default IPv4 gateway ():

Note To accept the current setting, press **Enter**.

f. When asked if you want to change the management IPv6 settings, enter **N**.

Note While the VDS TC software contains the infrastructure to configure IPv6 on the management server, this functionality is not currently supported.

Step 13 After the installation script is finished running, enter the command **cd /opt/pang/useful** to change directories.

Step 14 Enter the command **./management_machine_replacement.sh** {*list_of_ces*} where *list_of_ces* is a list of all existing cache engines separated by spaces. For example, if you are replacing the management server for a VDS-TS installation with eight cache engines, enter:

```
./management_machine_replacement.sh ce-1 ce-2 ce-3 ce-4 ce-5 ce-6 ce-7 ce-8
```

Post Installation

Note This procedure is used to restore the previously backed up files. If the backup files are not available, skip Step 1, and proceed to Step 2.

Procedure

Step 1 Using SFTP software, such as WinSCP, connect to the VDS TC management server using the user name **padmin** and the password that was provided by Cisco.

Step 2 Copy the files that you backed up in the Prerequisites section (PALicense.xml and cluster_conf.xml) to the /tftpboot folder on the management server.

Step 3 Log into the CLI of the VDS TC system and enter Enable mode.

Step 4 Enter the command **config** to switch to Configuration mode.

Step 5 To import the cluster_conf.xml configuration file, enter the command **import 127.0.0.1 cluster_conf.xml**.

Step 6 **Important!** If Step 5 fails, run the following commands from the VDS TC management server:

a. **cd /opt/pang/mgmt/bin/pang_rescue**

b. Enter the command **passwords**.

c. Enter the command **Exit**.

d. Repeat Step 5.

Note If Step 5 fails a second time, contact your Cisco Support Engineer.

Step 7 To apply the imported configuration file, enter the command **apply**. If the “Service stop required to modify NTP related parameters. Stop service ? [y|n]” message appears, enter **y**. For example:

```
configuration# apply
applying configuration...
Service stop required to modify NTP related parameters. Stop service ? [y|n] :y
```

Step 8 Enter the command **exit** to exit the Configuration mode.

Installing the License

After importing the configuration file you need to install the license file.

Procedure

Step 1 To import the license file to the system from the VDS TC Enable mode CLI, enter the command **license import 127.0.0.1** *filename*, where *filename* is the license file name. When you enter this command, the details of the license will be displayed. Verify the features of the license and the expiration date.

Note You must be in Enable mode, *not* Configuration mode to import a new license.

Step 2 To activate the imported license, enter the command **license activate**. Enter **Y** when prompted to confirm the activation.

Step 3 Enter the command **oper service start** to start the VDS TC service.

Step 4 To confirm that the VDS TC service has successfully started, enter the command **show eventlog** and look for the message "PANG started".

Note If the post installation steps fail, run the following commands on the VDS TC management server:

- **/opt/pang/mgmt/bin/pang_rescue**
- **Type: passwords**
- **Exit**
- Repeat the post installation steps.

CHAPTER 8

Management Server Replacement Procedure for a VDS TC 5.7.3 System Upgraded from VDS TC 5.2.1

Note Only use this chapter if you are replacing a Management Server in a VDS TC system that was upgraded to VDS TC 5.7.3 from VDS TC 5.2.1. If you are replacing a Management Server in a VDS TC system that had a fresh install of VDS TC 5.7.3 and was *not* upgraded to 5.7.3, refer to Chapter 7, "Management Server Replacement Procedure for a System with a Fresh Install of VDS TC 5.7.3". If you are replacing a Management Server in a VDS TC that was upgraded to VDS TC 5.7.3 from VDS TC 5.6.1, refer to Chapter 9, "Management Server Replacement Procedure for a VDS TC 5.7.3 System Upgraded from VDS TC 5.6.1".

To determine whether the system has been upgraded and from what version, from the Management Server enter the command **grep "upgrade system to" /var/log/peerapp/peerapp_system***. This will display a list of all previous versions.

This chapter discusses how to replace a VDS TC management server. It contains the following sections:

- Prerequisites
- Replacing the VDS TC Management Server
- Post Installation

Prerequisites

Before you replace the VDS TC management server, you must obtain the following files:

- The VDS TC 5.2.0 ISO image file: VDS-TC_Installer-5.2.0b123-5.2.0b124-ISO-5.2.0b23-Cisco.iso
- The updated kernel load for VDS TC 5.2.1: linux-2.6.27.19-5-llpf_10--KERNEL_LOAD_1414071783-r44244.tar.
- The VDS TC 5.7.3 ISO image file: VDS-TC_Installer_5.7.3b53-5.7.3b54-ISO-7-Cisco.iso

Note For more information refer to the “Cisco Videoscape Distribution Suite Transparent Caching Software C-Series Cluster Installation Guide” or the “Cisco Videoscape Distribution Suite Transparent Caching Software Blade Server Cluster Installation Guide”.

Backup Files

Before you perform a VDS TC management server replacement, you need to backup the PALicense.xml and cluster_conf.xml files. Follow these steps to backup these files:

Note **Important:** The following procedure for backing up files apply only if the VDS TC management server is still accessible. If the management server is not accessible, please contact Cisco support for the VDS License and cluster_conf.xml files.

Step 1 Using SFTP software, such as WinSCP, connect to the VDS TC management server.

Step 2 Copy the following files to a location other than the VDS TC management server:

- /opt/pang/mgmt/config/PALicense.xml
- /opt/pang/mgmt/config/cluster_conf.xml

Replacing the VDS TC Management Server

Procedure

Step 1 Disconnect the old VDS TC management server and replace it with the new management server. Next, connect the new management server to the cache engines. For more information, refer to Chapter 1 “Prerequisites and C-Series Cluster Physical Installation” in the “Cisco Videoscape Distribution Suite Transparent Caching Software C-Series Cluster Installation Guide” for a C-Series Cluster installation, and Chapter 1 “Prerequisites and Blade Server Cluster Physical Installation” chapter in the “Cisco Videoscape Distribution Suite Transparent Caching Software Blade Server Cluster Installation Guide” for a Blade Server cluster installation.

Step 2 Confirm that the correct firmware is installed on the Management Server for VDS TC 5.7.3:

- Cisco UCS C220-M3S:
  - Firmware version: 2.0(6f)
  - Image file: ucs-c220-huu-2.0.6f.iso
- Cisco UCS-C240-M3S:
  - Firmware version: 2.0(6f)
  - Image file: ucs-c240-huu-2.0.6f.iso

Step 3 Use one of the following chapters in the “Cisco Videoscape Distribution Suite Transparent Caching Software C-Series Cluster Installation Guide” or “Cisco Videoscape Distribution Suite Transparent Caching Software Blade Server Cluster Installation Guide” to configure the BIOS:

- For a VDS TC C-Series cluster installation, follow the steps in Chapter 3“Management Server Installation” in the “Cisco Videoscape Distribution Suite Transparent Caching Software C-Series Cluster Installation Guide” for Release 5.7.3.
- For a VDS TC Blade Server cluster installation, follow the steps in Chapter 3 “Management Server Installation” in the “Cisco Videoscape Distribution Suite Transparent Caching Software Blade Server Cluster Installation Guide” for Release 5.7.3.

Step 4 Use one of the following chapters in the “Cisco Videoscape Distribution Suite Transparent Caching Software C-Series Cluster Installation Guide” or “Cisco Videoscape Distribution Suite Transparent Caching Software Blade Server Cluster Installation Guide” to configure the operating system of the management server and install the VDS TC 5.2.0 ISO image:

- For a VDS TC C-Series cluster installation, follow the steps in Chapter 3“Management Server Installation” in the “Cisco Videoscape Distribution Suite Transparent Caching Software C-Series Cluster Installation Guide” for VDS TC Release 5.2.0.
- For a VDS TC Blade Server cluster installation, follow the steps in Chapter 3 “Management Server Installation” in the “Cisco Videoscape Distribution Suite Transparent Caching Software Blade Server Cluster Installation Guide” for VDS TC Release 5.2.0.

Note Make sure you are using the installation guide for VDS TC Release 5.2.0.

Note At this point, perform *only* the steps in Chapter 3 of the VDS TC Installation Guides. Do *not* configure the cache storage or install the VDS TC cluster software until instructed to do so in later steps.

Note When prompted, use the old VDS TC management server Network parameters and address values.

Step 5 Using SSH software, such as Putty, open an SSH connection to the management IP address of the VDS TC management server. Log in as a root user.

Step 6 To update the cluster SSH keys, enter the following commands:

- **cd /opt/pang/useful/**
- **./replace_server_keys.sh -servers** <*no_ of_ces*>, where *no_of_ces* is the number of cache engines in the cluster. For example, for a cluster installation that has four cache engines, enter the command **./replace_server_keys.sh -servers 4**.

Step 7 Using either the “Cisco Videoscape Distribution Suite Transparent Caching Software C-Series Cluster Installation Guide” or the “Cisco Videoscape Distribution Suite Transparent Caching Software Blade Server Cluster Installation Guide”, install and configure the Storage Manager using the following sections:

- For a VDS TC C-Series cluster installation:
  - Perform the steps in the *Installing the Storage Manager on the Management Server* section of Chapter 5 “Configuring the Cache Storage”.
  - Perform Step 1 through Step 6 in the *Configuring the First Storage Array* section of Chapter 5 “Configuring the Cache Storage.” In Step 6, choose **Automatic** for the method for the addition of the storage subsytem.
- For a VDS TC Blade Server cluster installation:

- Perform the steps in the *Installing the Storage Manager on the Management Server* section of Chapter 5 “Configuring the Cache Storage”.
- Perform Step 1 through Step 6 in the *Configuring the First Storage Array* section of Chapter 5 “Configuring the Cache Storage.” In Step 6, choose **Automatic** for the method for the addition of the storage subsytem.

Step 8 From the management server enter the command **scp root@ce-1:/etc/hosts /etc/hosts**

Step 9 Using SFTP software, such as WinSCP, connect to the management IP address that was assigned to the VDS TC server. Log in using the user name **padmin** and the password that was provided by Cisco. Copy the updated kernel load file, linux-2.6.27.19-5-llpf_10--KERNEL_LOAD_1414071783-r44244.tar, to the **/tmp** folder. Close the SFTP software.

Step 10 From the SSH connection to the VDS TC management server logged in as root, enter the command **cd /tmp**.

Step 11 Enter the command **tar -xvf linux-2.6.27.19-5-llpf_10--KERNEL_LOAD_1414071783-r44244.tar** to extract the kernel files.

Step 12 Enter the command **cd linux-2.6.27.19-5-llpf_10--KERNEL_LOAD_1414071783** to change to the new folder that was created when extracting the files in Step 11.

Step 13 Enter the command **./linux-2.6.27.19-5-llpf_10--INSTALL_PHASE-1.sh** to run the first phase of the kernel installation.

Step 14 Each time you see the following prompt, press **Enter**:

```
Press [Enter] key...
```

Note You will see this prompt several times.

Step 15 When you see the following prompt, press **Enter**. This will cause the system to reboot for the first time.

```
Going to REBOOT, Press [Enter] key...
```

Step 16 After the management server reboots, log in as padmin and enter the command **su root**. When prompted, enter the password that was provided by Cisco.

Step 17 Enter the command **cd /tmp**.

Step 18 Enter the command **cd linux-2.6.27.19-5-llpf_10--KERNEL_LOAD_1414071783**.

Step 19 Enter the command **./linux-2.6.27.19-5-llpf_10--INSTALL_PHASE-2_Cisco.sh** to run the second phase of the kernel installation.

Step 20 Each time you see the following prompt, press **Enter**:

```
Press [Enter] key...
```

Step 21 When you see the following prompt, press **Enter**. This will cause the system to reboot.

```
Going to REBOOT, Press [Enter] key...
```

Step 22 After the management server reboots, log in as padmin and enter the command **su root**. When prompted, enter the password that was provided by Cisco.

Step 23 To confirm the kernel update, enter the command **uname –a** and verify that the date Tue Aug 19 16:27:56 GMT 2014 appears in the output, as shown in the following example.

```
Linux ce-1 2.6.27.19-llpf_10-5-default #26 SMP Tue Aug 19 16:27:56 GMT 2014 x86_64 x86_64
x86_64 GNU/Linux
```

Step 24 Using SFTP software, such as WinSCP, connect to the management IP address that was assigned to the VDS TC server. Log in using the user name **padmin** and the password that was provided by Cisco. Copy the VDS TC 5.7.3 ISO image file, VDS-TC_Installer_5.7.3b53-5.7.3b54-ISO-7-Cisco.iso, to the **/opt/pang/iso** folder. Close the SFTP software.

Step 25 From the Management Server logged in as root, perform the following steps to extract the FI file for VDS TC 5.7.3:

a. Enter the command **mkdir /opt/pang/iso/5.7.3_iso** to create a new folder for the 5.7.3 files.

b. Enter the command **cd /opt/pang/iso** to change folders.

c. Enter the command **mount VDS-TC_Installer_5.7.3b53-5.7.3b54-ISO-7-Cisco.iso -o loop /opt/pang/iso/5.7.3_iso** to extract the VDS-TC_FI_5.7.3b54_Server_Cluster.tar.gz file.

d. Enter the command **cd /** to change to the root folder and then enter the command **umount /opt/pang/iso/5.7.3_iso** to unmount the iso image.

e. Enter the command **rm /opt/pang/iso/VDS-TC_Installer_5.7.3b53-5.7.3b54-ISO-7-Cisco.iso** to delete the iso file.

Step 26 From the management server run the following commands:

- **cd /tftpboot/**
- **mkdir Installer_ver5.7.3**
- **cd Installer_ver5.7.3**

Step 27 Enter the command **cp /opt/pang/iso/5.7.3_iso/install-tools/PeerApp_FI-GA/VDS-TC_FI_5.7.3b54_Server_Cluster.tar.gz /tftpboot** to copy the VDS-TC_FI_5.7.3b54_Server_Cluster.tar.gz file to the /tftpboot folder.

Step 28 Enter the command **tar -zxvf ../VDS-TC_FI_5.7.3b54_Server_Cluster.tar.gz** to extract the files.

Note If you are connected remotely to the management server, use the **screen** command. This ensures that if the session is disconnected during the installation as result of a Management IP set operation, you will be able to reconnect and resume the session with the **screen -r** command.

Step 29 Edit the GA_installer.rc file to set the “UPGRADE CACHE SOFTWARE” parameter to “none” and the GRACE_UPGRADE parameter to “yes”. Use the command vi GA_installer.rc to edit this file.

For example in a VDS TC 16 CE solution:

vi GA_installer.rc

```
#MANAGEMENT_SOFTWARE {no, yes}
MANAGEMENT_SOFTWARE=yes

#UPGRADEL_CACHE_SOFTWARE {all, none, space separated list(ce-1 ce-2 ce-3
ce-4 ce-5 ce-6 ce-7 ce-8 ce-9 ce-10 ce-11 ce-12 ce-13 ce-14 ce-15 ce-16)}
UPGRADEL_CACHE_SOFTWARE=none

UPGRADE_TAR_BALL=<filename>.tar.gz
MANAGEMENT_ENVIRONMENT=yes
GRACE_UPGRADE=no
FIRST_INSTALL=yes
```

Step 30 After you are finished editing the GA_installer.rc file, in the vi editor enter **:x** and press **Enter** to save and exit the file.

Step 31 Next, perform the following steps to run the ./GA_installer.sh script:

a. Enter the command **./GA_installer.sh** to run the installation script.

b. At the Enter primary name server IP prompt, enter the IP address of your primary DNS server. Enter **y** to confirm the entry.

Note If a DNS server is unavailable, enter 0.0.0.0 as the IP address of the DNS server to prevent system functionality from being negatively affected.

c. At the Enter secondary name server IP prompt, enter the IP address of your secondary DNS server. Enter **y** to confirm the entry.

Note If a secondary DNS server is unavailable, leave this entry blank.

Figure 8-1 Installation Script Dialog

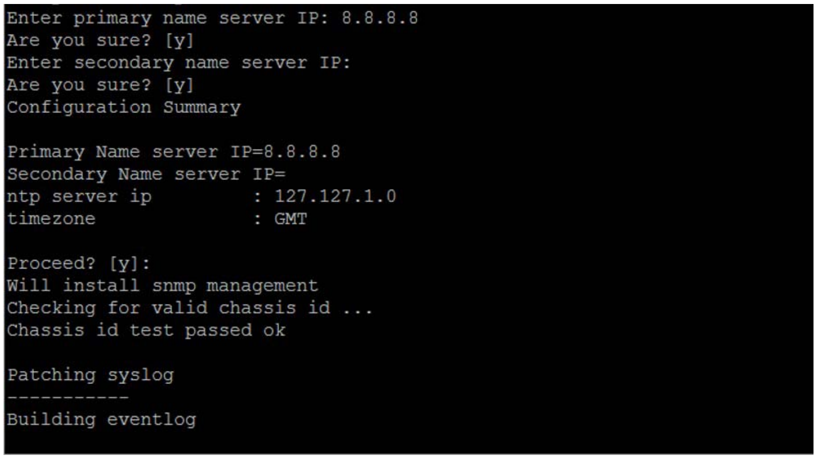

```
Enter primary name server IP: 8.8.8.8
Are you sure? [y]
Enter secondary name server IP:
Are you sure? [y]
Configuration Summary

Primary Name server IP=8.8.8.8
Secondary Name server IP=
ntp server ip         : 127.127.1.0
timezone              : GMT

Proceed? [y]:
Will install snmp management
Checking for valid chassis id ...
Chassis id test passed ok

Patching syslog
-----------
Building eventlog
```

d. The GA Installer script displays a summary of the configuration. Review this configuration. If the configuration is correct, press **Y** to proceed. If the configuration is not correct, press **N** to exit the installation and start over.

e. When asked if you want to change the management IPv4 settings, enter **N** unless you need to change the IP address information that was configured during the original VDS TC Management Server installation. If you enter **Y**, enter information at the following prompts:

- Management IPv4 address ():
- Management IPv4 netmask ():
- Default IPv4 gateway ():

Note To accept the current setting, press **Enter**.

f. When asked if you want to change the management IPv6 settings, enter **N**.

Note While the VDS TC software contains the infrastructure to configure IPv6 on the management server, this functionality is not currently supported.

Step 32 After the installation script is finished running, enter the command **cd /opt/pang/useful** to change directories.

Step 33 Enter the command **./management_machine_replacement.sh** {*list_of_ces*} where *list_of_ces* is a list of all existing cache engines separated by spaces. For example, if you are replacing the management server for a VDS-TS installation with eight cache engines, enter:

```
./management_machine_replacement.sh ce-1 ce-2 ce-3 ce-4 ce-5 ce-6 ce-7 ce-8
```

Post Installation

Note This procedure is used to restore the previously backed up files. If the backup files are not available, skip Step 1, and proceed to Step 2.

Procedure

Step 1 Using SFTP software, such as WinSCP, connect to the VDS TC management server using the user name **padmin** and the password that was provided by Cisco.

Step 2 Copy the files that you backed up in the Prerequisites section (PALicense.xml and cluster_conf.xml) to the /tftpboot folder on the management server.

Step 3 Log into the CLI of the VDS TC system and enter Enable mode.

Step 4 Enter the command **config** to switch to Configuration mode.

Step 5 To import the cluster_conf.xml configuration file, enter the command **import 127.0.0.1 cluster_conf.xml**.

Step 6 **Important!** If Step 5 fails, run the following commands from the VDS TC management server:

a. **cd /opt/pang/mgmt/bin/pang_rescue**

b. Enter the command **passwords**.

c. Enter the command **Exit**.

d. Repeat Step 5.

Note If Step 5 fails a second time, contact your Cisco Support Engineer.

Step 7 To apply the imported configuration file, enter the command **apply**. If the “Service stop required to modify NTP related parameters. Stop service ? [y|n]” message appears, enter **y**. For example:

```
configuration# apply
applying configuration...
Service stop required to modify NTP related parameters. Stop service ? [y|n] :y
```

Step 8 Enter the command **exit** to exit the Configuration mode.

Installing the License

After importing the configuration file you need to install the license file.

Procedure

Step 1 To import the license file to the system from the VDS TC Enable mode CLI, enter the command **license import 127.0.0.1** *filename*, where *filename* is the license file name. When you enter this command, the details of the license will be displayed. Verify the features of the license and the expiration date.

Note You must be in Enable mode, *not* Configuration mode to import a new license.

Step 2 To activate the imported license, enter the command **license activate**. Enter **Y** when prompted to confirm the activation.

Step 3 Enter the command **oper service start** to start the VDS TC service.

Step 4 To confirm that the VDS TC service has successfully started, enter the command **show eventlog** and look for the message “PANG started”.

Note If the post installation steps fail, run the following commands on the VDS TC management server:

- **/opt/pang/mgmt/bin/pang_rescue**
- **Type: passwords**
- **Exit**
- Repeat the post installation steps.

Management Server Replacement Procedure for a VDS TC 5.7.3 System Upgraded from VDS TC 5.6.1

Note Only use this chapter if you are replacing a Management Server in a VDS TC system that was upgraded to VDS TC 5.7.3 from VDS TC 5.6.1. If you are replacing a Management Server in a VDS TC system that had a fresh install of VDS TC 5.7.3 and was *not* upgraded to 5.7.3, refer to Chapter 7, "Management Server Replacement Procedure for a System with a Fresh Install of VDS TC 5.7.3". If you are replacing a Management Server in a VDS TC that was upgraded to VDS TC 5.7.3 from VDS TC 5.2.1, refer to Chapter 8, "Management Server Replacement Procedure for a VDS TC 5.7.3 System Upgraded from VDS TC 5.2.1".

To determine whether the system has been upgraded and from what version, from the Management Server enter the command **grep "upgrade system to" /var/log/peerapp/peerapp_system***. This will display a list of all previous versions.

This chapter discusses how to replace a VDS TC management server. It contains the following sections:

- Prerequisites
- Replacing the VDS TC Management Server
- Post Installation

Prerequisites

Before you replace the VDS TC management server, you must obtain the following:

- The VDS TC 5.6.1 ISO image file, VDS-TC_Installer_5.6.1b57-5.6.1b58-ISO-15-Cisco.iso
- The VDS TC 5.7.3 ISO image file, VDS-TC_Installer_5.7.3b53-5.7.3b54-ISO-7-Cisco.iso

Note For more information refer to the "Cisco Videoscape Distribution Suite Transparent Caching Software C-Series Cluster Installation Guide" or the "Cisco Videoscape Distribution Suite Transparent Caching Software Blade Server Cluster Installation Guide".

Backup Files

Before you perform a VDS TC management server replacement, you need to backup the PALicense.xml and cluster_conf.xml files. Follow these steps to backup these files:

Note **Important:** The following procedure for backing up files apply only if the VDS TC management server is still accessible. If the management server is not accessible, please contact Cisco support for the VDS License and cluster_conf.xml files.

Step 1 Using SFTP software, such as WinSCP, connect to the VDS TC management server.

Step 2 Copy the following files to a location other than the VDS TC management server:

- /opt/pang/mgmt/config/PALicense.xml
- /opt/pang/mgmt/config/cluster_conf.xml

Replacing the VDS TC Management Server

Procedure

Step 1 Disconnect the old VDS TC management server and replace it with the new management server. Next, connect the new management server to the cache engines. For more information, refer to Chapter 1 “Prerequisites and C-Series Cluster Physical Installation” in the “Cisco Videoscape Distribution Suite Transparent Caching Software C-Series Cluster Installation Guide” for a C-Series Cluster installation, and Chapter 1 “Prerequisites and Blade Server Cluster Physical Installation” chapter in the “Cisco Videoscape Distribution Suite Transparent Caching Software Blade Server Cluster Installation Guide” for a Blade Server cluster installation.

Step 2 Confirm that the correct firmware is installed on the Management Server for VDS TC 5.7.3:

- Cisco UCS C220-M3S:
  - Firmware version: 2.0(6f)
  - Image file: ucs-c220-huu-2.0.6f.iso
- Cisco UCS-C240-M3S:
  - Firmware version: 2.0(6f)
  - Image file: ucs-c240-huu-2.0.6f.iso

Step 3 Use one of the following chapters in the “Cisco Videoscape Distribution Suite Transparent Caching Software C-Series Cluster Installation Guide” or “Cisco Videoscape Distribution Suite Transparent Caching Software Blade Server Cluster Installation Guide” to configure the BIOS:

- For a VDS TC C-Series cluster installation, follow the steps in Chapter 3 “Management Server Installation” in the “Cisco Videoscape Distribution Suite Transparent Caching Software C-Series Cluster Installation Guide” for Release 5.7.3.
- For a VDS TC Blade Server cluster installation, follow the steps in Chapter 3 “Management Server Installation” in the “Cisco Videoscape Distribution Suite Transparent Caching Software Blade Server Cluster Installation Guide” for Release 5.7.3.

Step 4 Use one of the following chapters in the "Cisco Videoscape Distribution Suite Transparent Caching Software C-Series Cluster Installation Guide" or "Cisco Videoscape Distribution Suite Transparent Caching Software Blade Server Cluster Installation Guide" to configure the operating system of the management server and install the VDS TC 5.6.1 ISO image:

- For a VDS TC C-Series cluster installation, follow the steps in Chapter 3"Management Server Installation" in the "Cisco Videoscape Distribution Suite Transparent Caching Software C-Series Cluster Installation Guide" for VDS TC Release 5.6.1.
- For a VDS TC Blade Server cluster installation, follow the steps in Chapter 3 "Management Server Installation" in the "Cisco Videoscape Distribution Suite Transparent Caching Software Blade Server Cluster Installation Guide" for VDS TC Release 5.6.1.

Note Make sure you are using the installation guide for VDS TC Release 5.6.1.

Note At this point, perform *only* the steps in Chapter 3 of the VDS TC Installation Guides. Do *not* configure the cache storage or install the VDS TC cluster software until instructed to do so in later steps.

Note When prompted, use the *old* VDS TC management server Network parameters and address values.

Step 5 Using SSH software, such as Putty, open an SSH connection to the management IP address of the VDS TC management server. Log in as a root user.

Step 6 To update the cluster SSH keys, enter the following commands:

- **cd /opt/pang/useful/**
- **./replace_server_keys.sh -servers** <*no_ of_ces*>, where *no_of_ces* is the number of cache engines in the cluster. For example, for a cluster installation that has four cache engines, enter the command **./replace_server_keys.sh -servers 4**.

Step 7 Using either the "Cisco Videoscape Distribution Suite Transparent Caching Software C-Series Cluster Installation Guide" or the "Cisco Videoscape Distribution Suite Transparent Caching Software Blade Server Cluster Installation Guide", install and configure the Storage Manager using the following sections:

- For a VDS TC C-Series cluster installation:
  - Perform the steps in the *Installing the Storage Manager on the Management Server* section of Chapter 5 "Configuring the Cache Storage".
  - Perform Step 1 through Step 6 in the *Configuring the First Storage Array* section of Chapter 5 "Configuring the Cache Storage." In Step 6, choose **Automatic** for the method for the addition of the storage subsytem.
- For a VDS TC Blade Server cluster installation:
  - Perform the steps in the *Installing the Storage Manager on the Management Server* section of Chapter 5 "Configuring the Cache Storage".
  - Perform Step 1 through Step 6 in the *Configuring the First Storage Array* section of Chapter 5 "Configuring the Cache Storage." In Step 6, choose **Automatic** for the method for the addition of the storage array.

Step 8 From the management server enter the command **scp root@ce-1:/etc/hosts /etc/hosts**

Step 9 Using SFTP software, such as WinSCP, connect to the management IP address that was assigned to the VDS TC server. Log in using the user name **padmin** and the password that was provided by Cisco. Copy the VDS TC 5.7.3 ISO image file, VDS-TC_Installer_5.7.3b53-5.7.3b54-ISO-7-Cisco.iso, to the **/opt/pang/iso** folder. Close the SFTP software.

Step 10 From the Management Server logged in as root, perform the following steps to extract the FI file for VDS TC 5.7.3:

a. Enter the command **mkdir /opt/pang/iso/5.7.3_iso** to create a new folder for the 5.7.3 files.

b. Enter the command **cd /opt/pang/iso** to change folders.

c. Enter the command **mount VDS-TC_Installer_5.7.3b53-5.7.3b54-ISO-7-Cisco.iso -o loop /opt/pang/iso/5.7.3_iso** to extract the VDS-TC_FI_5.7.3b54_Server_Cluster.tar.gz file.

d. Enter the command **cd /** to change to the root folder and then enter the command **umount /opt/pang/iso/5.7.3_iso** to unmount the iso image.

e. Enter the command **rm /opt/pang/iso/VDS-TC_Installer_5.7.3b53-5.7.3b54-ISO-7-Cisco.iso** to delete the iso file.

Step 11 From the management server run the following commands:

- **cd /tftpboot/**
- **mkdir Installer_ver5.7.3**
- **cd Installer_ver5.7.3**

Step 12 Enter the command **cp /opt/pang/iso/5.7.3_iso/install-tools/PeerApp_FI-GA/VDS-TC_FI_5.7.3b54_Server_Cluster.tar.gz /tftpboot** to copy the VDS-TC_FI_5.7.3b54_Server_Cluster.tar.gz file to the /tftpboot folder.

Step 13 Enter the command **tar -zxvf ../VDS-TC_FI_5.7.3b54_Server_Cluster.tar.gz** to extract the files.

Note If you are connected remotely to the management server, use the **screen** command. This ensures that if the session is disconnected during the installation as result of a Management IP set operation, you will be able to reconnect and resume the session with the **screen -r** command.

Step 14 Edit the GA_installer.rc file to set the “UPGRADE CACHE SOFTWARE” parameter to “none” and the GRACE_UPGRADE parameter to “yes”. Use the command vi GA_installer.rc to edit this file.

For example in a VDS TC 16 CE solution:

vi GA_installer.rc

```
#MANAGEMENT_SOFTWARE {no, yes}
MANAGEMENT_SOFTWARE=yes

#UPGRADEL_CACHE_SOFTWARE {all, none, space separated list(ce-1 ce-2 ce-3
ce-4 ce-5 ce-6 ce-7 ce-8 ce-9 ce-10 ce-11 ce-12 ce-13 ce-14 ce-15 ce-16)}
UPGRADEL_CACHE_SOFTWARE=none

UPGRADE_TAR_BALL=<filename>.tar.gz
MANAGEMENT_ENVIRONMENT=yes
GRACE_UPGRADE=no
FIRST_INSTALL=yes
```

Step 15 After you are finished editing the GA_installer.rc file, in the vi editor enter **:x** and press **Enter** to save and exit the file.

Step 16 Next, perform the following steps to run the ./GA_installer.sh script:

a. Enter the command **./GA_installer.sh** to run the installation script.

b. At the Enter primary name server IP prompt, enter the IP address of your primary DNS server. Enter **y** to confirm the entry.

Note If a DNS server is unavailable, enter 0.0.0.0 as the IP address of the DNS server to prevent system functionality from being negatively affected.

c. At the Enter secondary name server IP prompt, enter the IP address of your secondary DNS server. Enter **y** to confirm the entry.

Note If a secondary DNS server is unavailable, leave this entry blank.

Figure 9-1 Installation Script Dialog

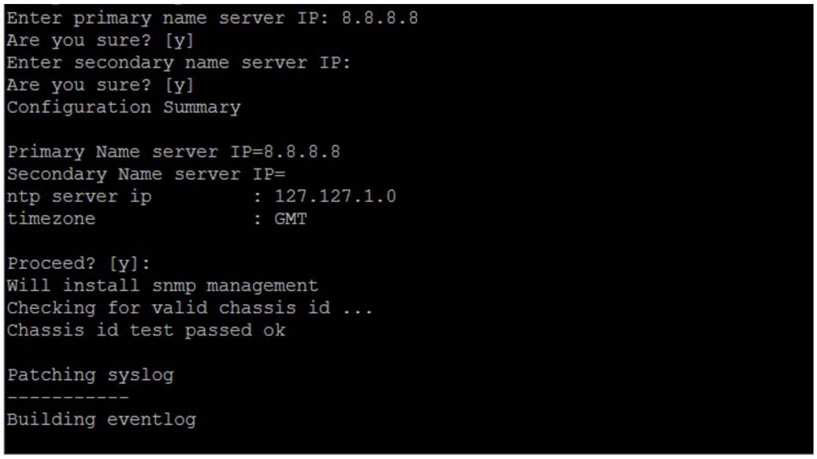

```
Enter primary name server IP: 8.8.8.8
Are you sure? [y]
Enter secondary name server IP:
Are you sure? [y]
Configuration Summary

Primary Name server IP=8.8.8.8
Secondary Name server IP=
ntp server ip         : 127.127.1.0
timezone              : GMT

Proceed? [y]:
Will install snmp management
Checking for valid chassis id ...
Chassis id test passed ok

Patching syslog
-----------
Building eventlog
```

d. The GA Installer script displays a summary of the configuration. Review this configuration. If the configuration is correct, press **Y** to proceed. If the configuration is not correct, press **N** to exit the installation and start over.

e. When asked if you want to change the management IPv4 settings, enter **N** unless you need to change the IP address information that was configured during the original VDS TC Management Server installation. If you enter **Y**, enter information at the following prompts:

- Management IPv4 address ():
- Management IPv4 netmask ():
- Default IPv4 gateway ():

Note To accept the current setting, press **Enter**.

f. When asked if you want to change the management IPv6 settings, enter **N**.

Note While the VDS TC software contains the infrastructure to configure IPv6 on the management server, this functionality is not currently supported.

Step 17 After the installation script is finished running, enter the command **cd /opt/pang/useful** to change directories.

Step 18 Enter the command **./management_machine_replacement.sh** {*list_of_ces*} where *list_of_ces* is a list of all existing cache engines separated by spaces. For example, if you are replacing the management server for a VDS-TS installation with eight cache engines, enter:

```
./management_machine_replacement.sh ce-1 ce-2 ce-3 ce-4 ce-5 ce-6 ce-7 ce-8
```

Post Installation

Note This procedure is used to restore the previously backed up files. If the backup files are not available, skip Step 1, and proceed to Step 2.

Procedure

Step 1 Using SFTP software, such as WinSCP, connect to the VDS TC management server using the user name **padmin** and the password that was provided by Cisco.

Step 2 Copy the files that you backed up in the Prerequisites section (PALicense.xml and cluster_conf.xml) to the /tftpboot folder on the management server.

Step 3 Log into the CLI of the VDS TC system and enter Enable mode.

Step 4 Enter the command **config** to switch to Configuration mode.

Step 5 To import the cluster_conf.xml configuration file, enter the command **import 127.0.0.1 cluster_conf.xml**.

Step 6 **Important!** If Step 5 fails, run the following commands from the VDS TC management server:

a. **cd /opt/pang/mgmt/bin/pang_rescue**

b. Enter the command **passwords**.

c. Enter the command **Exit**.

d. Repeat Step 5.

Note If Step 5 fails a second time, contact your Cisco Support Engineer.

Step 7 To apply the imported configuration file, enter the command **apply**. If the “Service stop required to modify NTP related parameters. Stop service ? [y|n]” message appears, enter **y**. For example:

```
configuration# apply
applying configuration...
Service stop required to modify NTP related parameters. Stop service ? [y|n] :y
```

Step 8 Enter the command **exit** to exit the Configuration mode.

Installing the License

After importing the configuration file you need to install the license file.

Procedure

Step 1 To import the license file to the system from the VDS TC Enable mode CLI, enter the command **license import 127.0.0.1** *filename*, where *filename* is the license file name. When you enter this command, the details of the license will be displayed. Verify the features of the license and the expiration date.

Note You must be in Enable mode, *not* Configuration mode to import a new license.

Step 2 To activate the imported license, enter the command **license activate**. Enter **Y** when prompted to confirm the activation.

Step 3 Enter the command **oper service start** to start the VDS TC service.

Step 4 To confirm that the VDS TC service has successfully started, enter the command **show eventlog** and look for the message “PANG started”.

Note If the post installation steps fail, run the following commands on the VDS TC management server:

- **/opt/pang/mgmt/bin/pang_rescue**
- **Type: passwords**
- **Exit**
- Repeat the post installation steps.

CHAPTER 10

Storage Disk Replacement Procedure for NetApp Storage Enclosure

This chapter discusses how to determine that a disk has failed in the VDS TC NetApp SAN storage enclosure and the steps that are required to replace the disk.

Note If you purchased your NetApp storage enclosure from a NetApp partner and not from Cisco, please contact the NetApp partner for help with the storage disk replacement.

Note If you are having problems with a NetApp storage controller, you must open a support case. Replacing a NetApp storage controller is not a customer procedure; it must be performed by an authorized NetApp technician.

If you purchased the NetApp storage enclosure from a NetApp partner and not from Cisco, please contact the NetApp partner to open a case. If you purchased the NetApp storage enclosure from Cisco as part of the VDS TC solution, please contact Cisco TAC to open a case.

Symptoms of a Failed VDS TC Storage Disk

One symptom of a failed storage disk is if the following message frequently appears in the System Events section of the Status > Dashboard window in the VDS TC Manager:

“warning: Volume /dev/*volume_name* (*volume_id*) is turned OFF because of errors”

For example:

```
“warning: Volume /dev/sdak3 (id 60) is turned OFF because of errors”
```

Figure 10-1 SNMP Traps Message

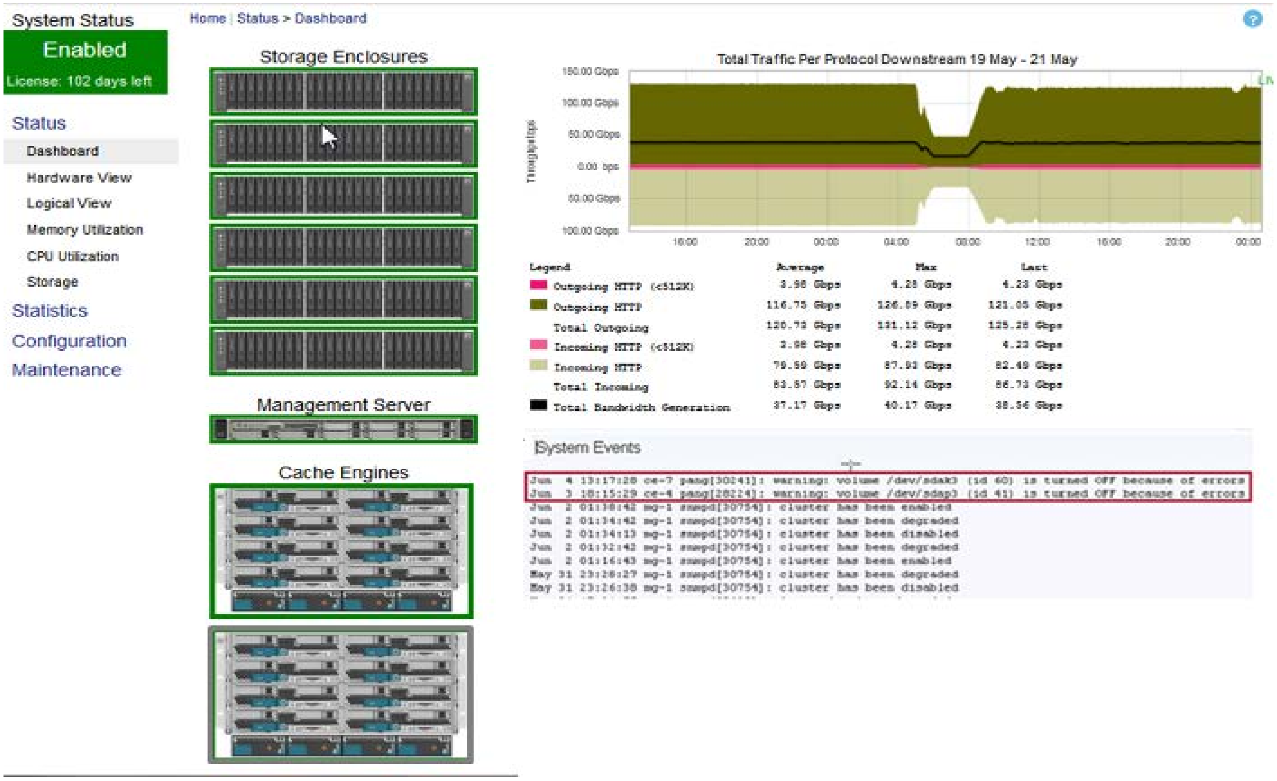

Legend	Average	Max	Last
🟥 Outgoing HTTP (<512K)	3.98 Gbps	4.28 Gbps	4.23 Gbps
🟫 Outgoing HTTP	116.75 Gbps	126.89 Gbps	121.05 Gbps
Total Outgoing	120.73 Gbps	131.12 Gbps	125.28 Gbps
🟥 Incoming HTTP (<512K)	3.98 Gbps	4.28 Gbps	4.23 Gbps
🟨 Incoming HTTP	79.59 Gbps	87.93 Gbps	82.49 Gbps
Total Incoming	83.57 Gbps	92.14 Gbps	86.73 Gbps
■ Total Bandwidth Generation	37.17 Gbps	40.17 Gbps	38.56 Gbps

System Events

```
Jun  4 13:17:28 ce-7 pang[30241]: warning: volume /dev/sdak3 (id 60) is turned OFF because of errors
Jun  3 18:15:29 ce-4 pang[28224]: warning: volume /dev/sdap3 (id 41) is turned OFF because of errors
Jun  2 01:38:42 mg-1 snmpd[30754]: cluster has been enabled
Jun  2 01:34:42 mg-1 snmpd[30754]: cluster has been degraded
Jun  2 01:34:13 mg-1 snmpd[30754]: cluster has been disabled
Jun  2 01:32:42 mg-1 snmpd[30754]: cluster has been degraded
Jun  2 01:16:43 mg-1 snmpd[30754]: cluster has been enabled
May 31 23:28:27 mg-1 snmpd[30754]: cluster has been degraded
May 31 23:26:38 mg-1 snmpd[30754]: cluster has been disabled
```

Please note that this error message can also be caused by problems other than a failed disk. To see more information about the failure and to determine if it is from a failed storage disk, follow these steps:

Step 1 Log into Cisco VDS TC Manager. From the main window that appears, choose **Status** > **Storage**.

Note The status of the storage controllers will appear as Operational, even if one of the disks being managed has failed.

Figure 10-2 Status > Storage Window

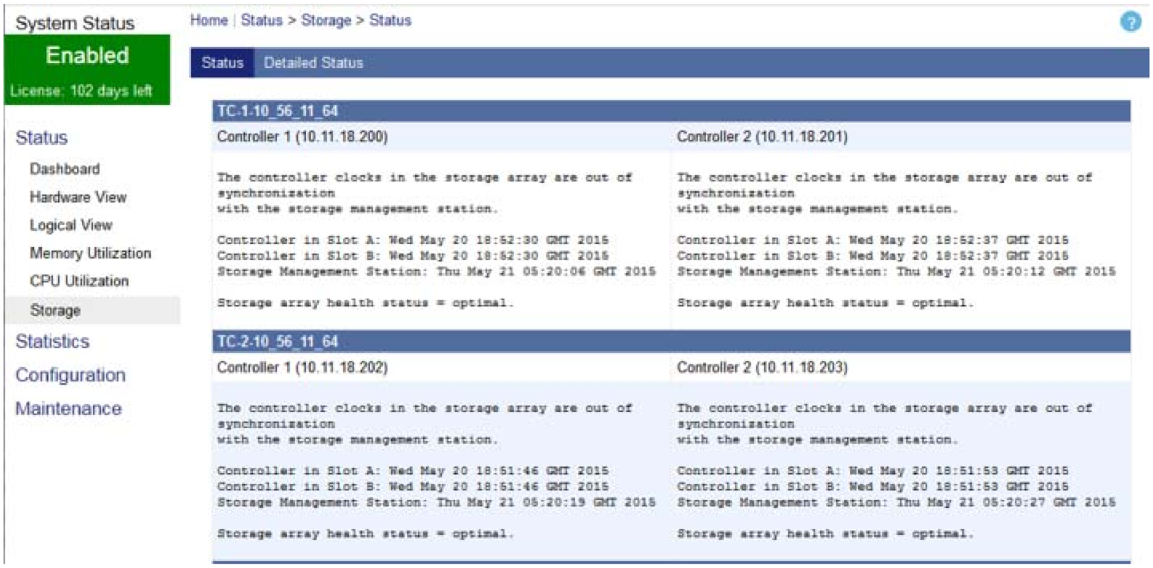

TC-1-10_56_11_64

Controller 1 (10.11.18.200)

```
The controller clocks in the storage array are out of
synchronization
with the storage management station.

Controller in Slot A: Wed May 20 18:52:30 GMT 2015
Controller in Slot B: Wed May 20 18:52:30 GMT 2015
Storage Management Station: Thu May 21 05:20:06 GMT 2015

Storage array health status = optimal.
```

Controller 2 (10.11.18.201)

```
The controller clocks in the storage array are out of
synchronization
with the storage management station.

Controller in Slot A: Wed May 20 18:52:37 GMT 2015
Controller in Slot B: Wed May 20 18:52:37 GMT 2015
Storage Management Station: Thu May 21 05:20:12 GMT 2015

Storage array health status = optimal.
```

TC-2-10_56_11_64

Controller 1 (10.11.18.202)

```
The controller clocks in the storage array are out of
synchronization
with the storage management station.

Controller in Slot A: Wed May 20 18:51:46 GMT 2015
Controller in Slot B: Wed May 20 18:51:46 GMT 2015
Storage Management Station: Thu May 21 05:20:19 GMT 2015

Storage array health status = optimal.
```

Controller 2 (10.11.18.203)

```
The controller clocks in the storage array are out of
synchronization
with the storage management station.

Controller in Slot A: Wed May 20 18:51:53 GMT 2015
Controller in Slot B: Wed May 20 18:51:53 GMT 2015
Storage Management Station: Thu May 21 05:20:27 GMT 2015

Storage array health status = optimal.
```

Step 2 From the Status > Storage window, click the **Detailed Status** tab to see detailed information about the status of the storage controllers and information about any possible drives that are having problems.

Step 3 When the detailed storage status information appears, look for information that appears in yellow boxes. This information will provide details about the storage subsytem and drives that are affected.

Figure 10-3 Detailed Storage Status Output

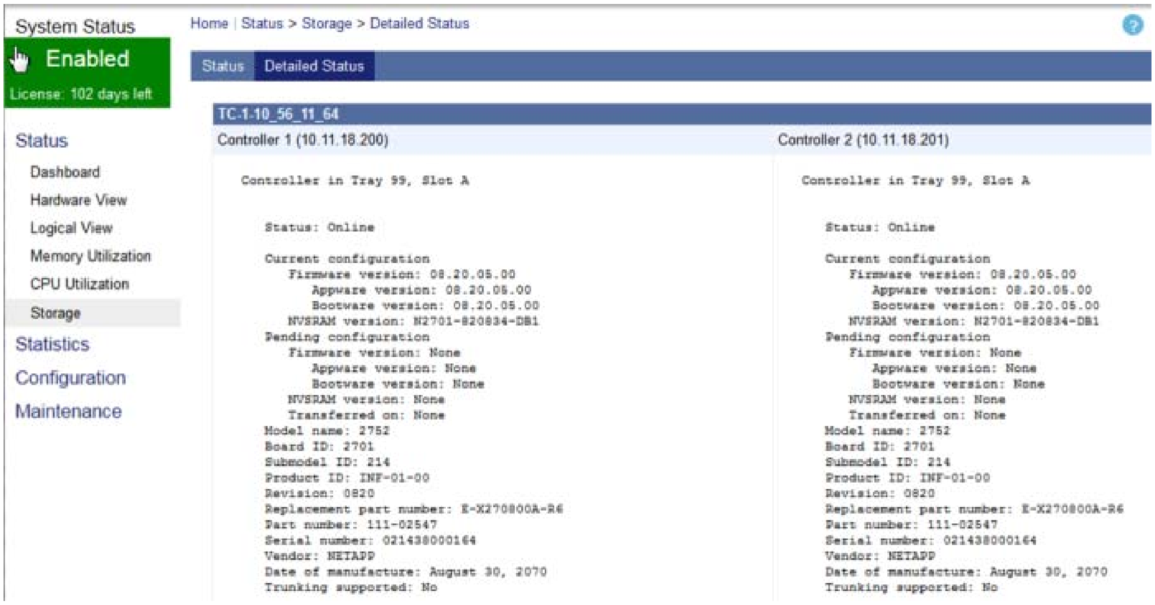

TC-1-10_56_11_64

Controller 1 (10.11.18.200)

```
Controller in Tray 99, Slot A


   Status: Online

   Current configuration
      Firmware version: 08.20.05.00
         Appware version: 08.20.05.00
         Bootware version: 08.20.05.00
      NVSRAM version: N2701-820834-DB1
   Pending configuration
      Firmware version: None
         Appware version: None
         Bootware version: None
      NVSRAM version: None
      Transferred on: None
   Model name: 2752
   Board ID: 2701
   Submodel ID: 214
   Product ID: INF-01-00
   Revision: 0820
   Replacement part number: E-X270800A-R6
   Part number: 111-02547
   Serial number: 021438000164
   Vendor: NETAPP
   Date of manufacture: August 30, 2070
   Trunking supported: No
```

Controller 2 (10.11.18.201)

```
Controller in Tray 99, Slot A


   Status: Online

   Current configuration
      Firmware version: 08.20.05.00
         Appware version: 08.20.05.00
         Bootware version: 08.20.05.00
      NVSRAM version: N2701-820834-DB1
   Pending configuration
      Firmware version: None
         Appware version: None
         Bootware version: None
      NVSRAM version: None
      Transferred on: None
   Model name: 2752
   Board ID: 2701
   Submodel ID: 214
   Product ID: INF-01-00
   Revision: 0820
   Replacement part number: E-X270800A-R6
   Part number: 111-02547
   Serial number: 021438000164
   Vendor: NETAPP
   Date of manufacture: August 30, 2070
   Trunking supported: No
```

Note The Volume Usage window in VDS TC Manager, available at Statistics > Storage, may show the Volume State as active even if the drive has failed, as seen in the screenshots below.

Figure 10-4 Volume Usage Tab

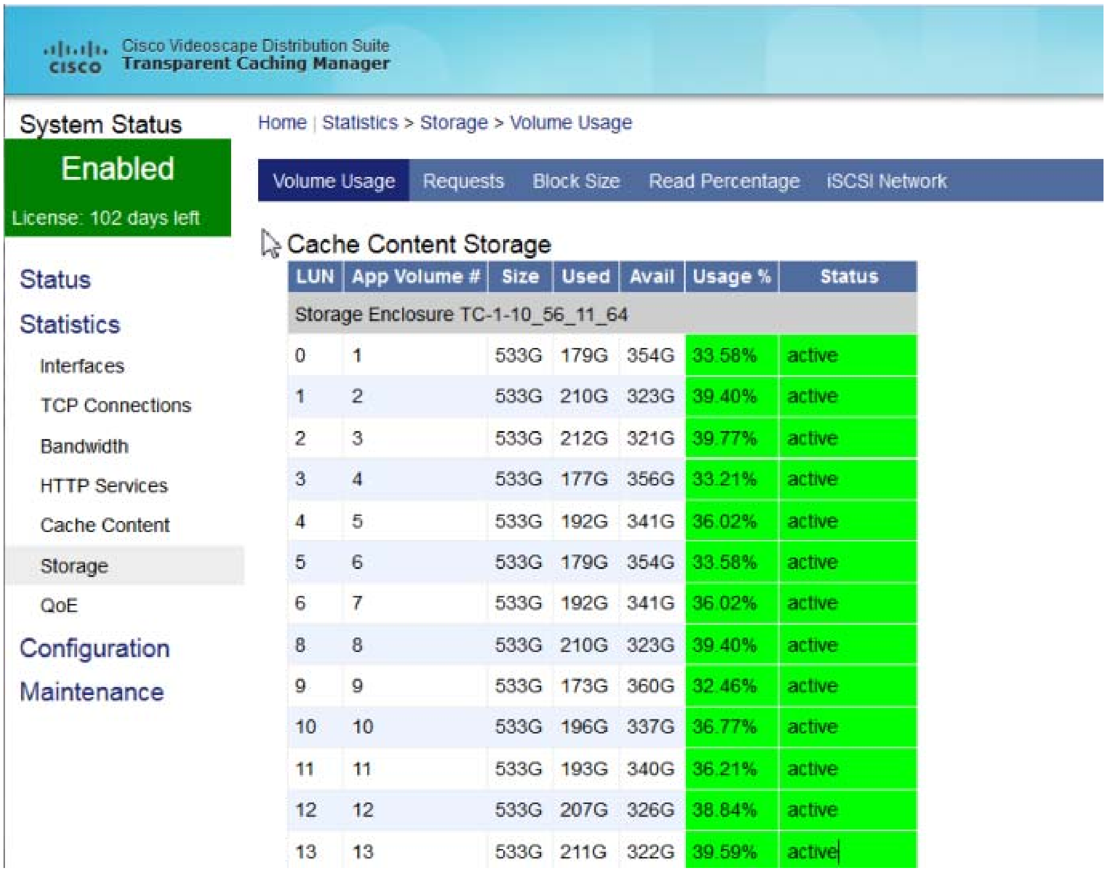

Cache Content Storage

LUN	App Volume #	Size	Used	Avail	Usage %	Status
Storage Enclosure TC-1-10_56_11_64						
0	1	533G	179G	354G	33.58%	active
1	2	533G	210G	323G	39.40%	active
2	3	533G	212G	321G	39.77%	active
3	4	533G	177G	356G	33.21%	active
4	5	533G	192G	341G	36.02%	active
5	6	533G	179G	354G	33.58%	active
6	7	533G	192G	341G	36.02%	active
8	8	533G	210G	323G	39.40%	active
9	9	533G	173G	360G	32.46%	active
10	10	533G	196G	337G	36.77%	active
11	11	533G	193G	340G	36.21%	active
12	12	533G	207G	326G	38.84%	active
13	13	533G	211G	322G	39.59%	active

Replacing a Storage Disk

Follow this procedure to replace a storage disk in the VDS TC NetApp SAN storage enclosure:

Procedure

Step 1 Using SSH software, open an SSH connection to the management IP address of the VDS TC management server.

Step 2 Log into the system using the username **padmin** and the password that was provided by Cisco.

Step 3 Enter the command **su root** to change to the root user. Enter the password for the root user when prompted. This password is provided by Cisco.

Step 4 To start the VNC server, enter the following commands:

a. **rm /tmp/.X1-lock** (This step is optional. It is only required if there is a leftover lock file from a previous vncserver instance.)

b. **rm /tmp/.X11-unix/X1** (This step is optional. It is only required if there is a leftover lock file from a previous vncserver instance.)

c. **vncserver :1**

Step 5 Use a VNC client to connect from your local computer to the VDS TC management server.

Figure 10-5 VNC Viewer

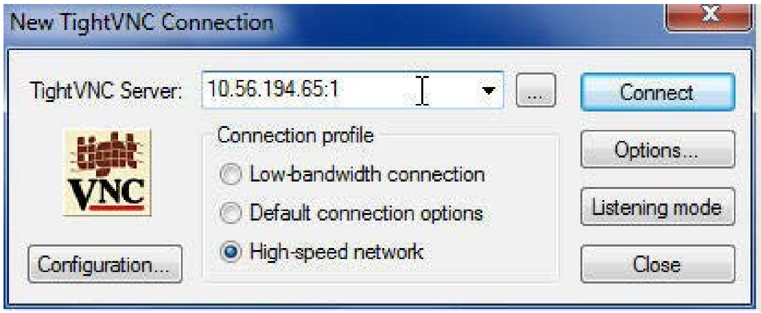

TightVNC Server: 10.56.194.65:1

Step 6 From the VNC console window, enter the command **/opt/SMgr/client/SMclient** to start the SANtricity ™ ES Storage Manager software.

Step 7 Right-click the storage array for which you are replacing a disk and choose **Manage Storage Array**. Click **No** in any pop-up windows that appear.

Figure 10-6 NetApp SANtricity Storage Manager

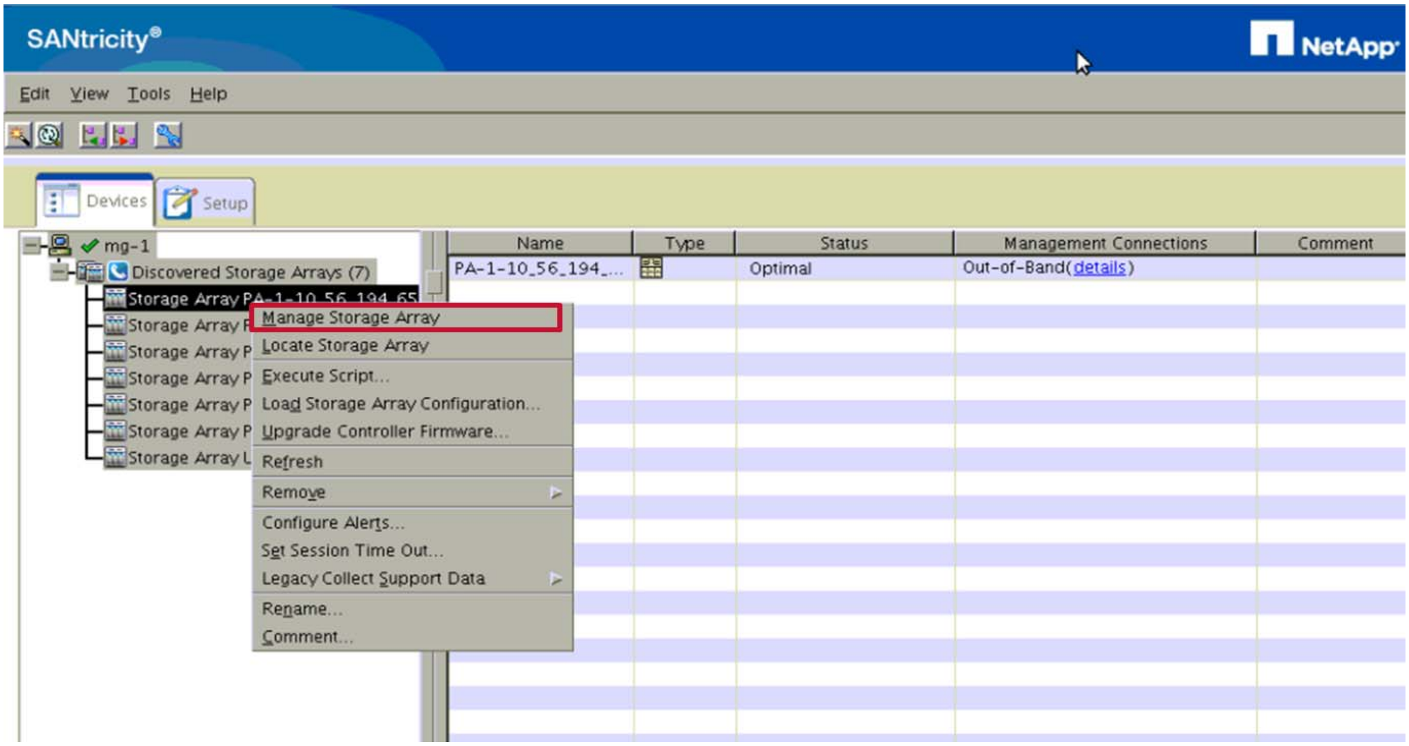

Name	Type	Status	Management Connections	Comment
PA-1-10_56_194_...		Optimal	Out-of-Band(details)	

Step 8 From the Subsystem Management window that appears, click the **Host Mappings** tab.

Step 9 Ensure that Storage Array is selected in the navigation pane on the left. In the Defined Mappings table, make note of all of the LUN numbers that are currently in use, including the LUN number of the failed drive. You will need this information in a later step.

Step 10 From the Array Management window that appears, click the **Storage & Copy Services** tab

Step 11 In the navigation pane on the left, expand the Volume Groups folder and locate the faulty Volume Group. Click + to expand the Volume Group and confirm that it contains a faulty volume. In the example, Volume Group 9 has a failed volume.

Note A red drive icon with an X next to it indicates a failed drive, as shown in the example.

Figure 10-7 NetApp SANtricity Storage Manager - Failed Volume

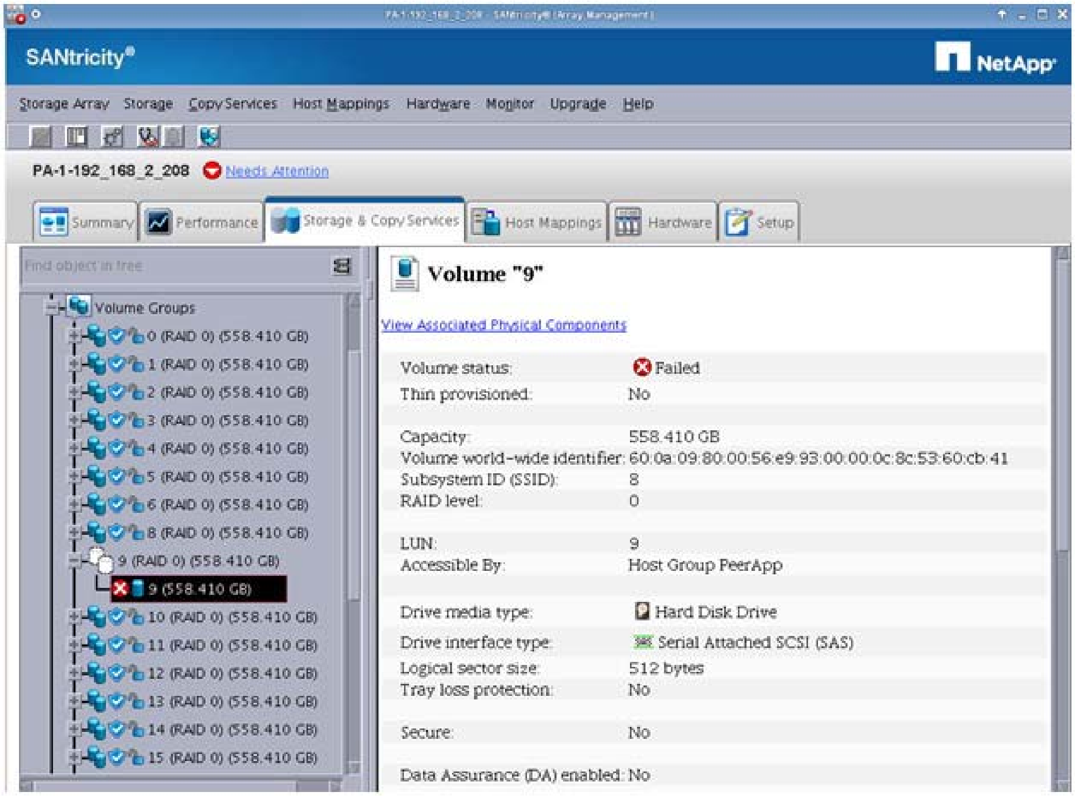

Volume "9"

Volume status:	☒ Failed
Thin provisioned:	No
Capacity:	558.410 GB
Volume world-wide identifier:	60:0a:09:80:00:56:e9:93:00:00:0c:8c:53:60:cb:41
Subsystem ID (SSID):	8
RAID level:	0
LUN:	9
Accessible By:	Host Group PeerApp
Drive media type:	Hard Disk Drive
Drive interface type:	Serial Attached SCSI (SAS)
Logical sector size:	512 bytes
Tray loss protection:	No
Secure:	No
Data Assurance (DA) enabled:	No

Step 12 Right-click the faulty Volume Group (*not* the volume) and choose **Delete**.

Figure 10-8 Delete Faulty Volume Group

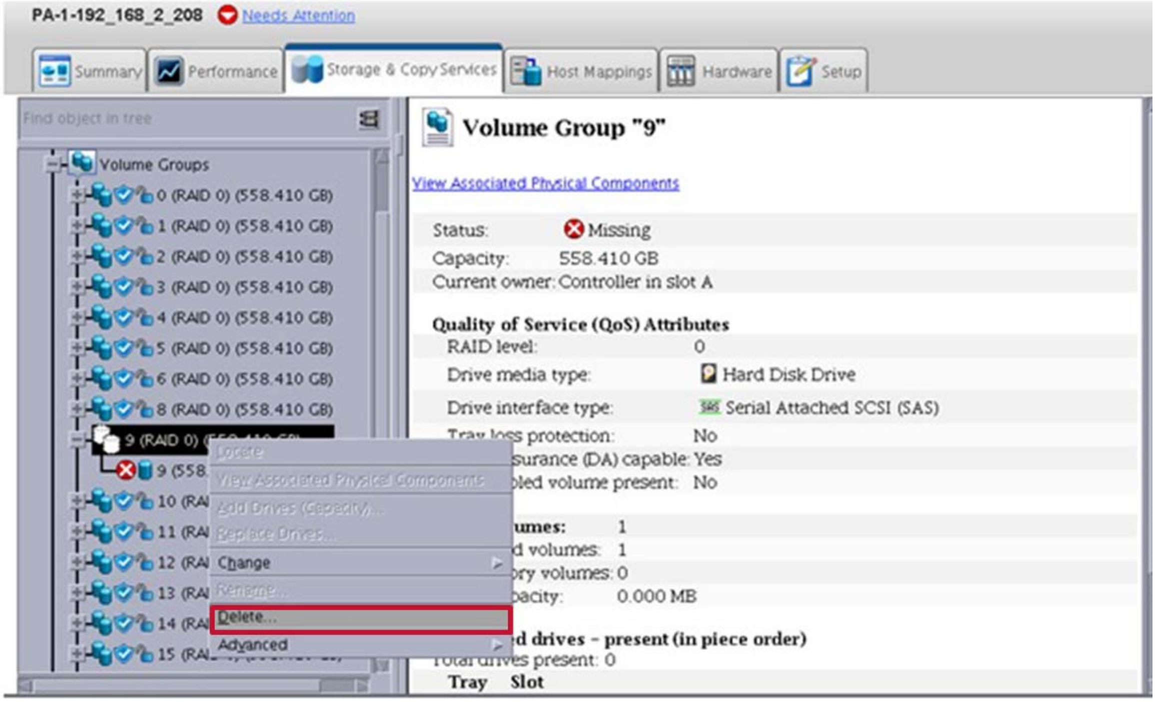

Volume Group "9"

View Associated Physical Components

Status: Missing
Capacity: 558.410 GB
Current owner: Controller in slot A

Quality of Service (QoS) Attributes

RAID level: 0
Drive media type: Hard Disk Drive
Drive interface type: Serial Attached SCSI (SAS)
Tray loss protection: No
surance (DA) capable: Yes
bled volume present: No

umes: 1
d volumes: 1
ory volumes: 0
pacity: 0.000 MB

d drives – present (in piece order)

Total drives present: 0

Tray Slot

Step 13 From the SANtricity Delete Volume Groups window that appears, ensure that the faulty volume group is selected and click **Delete**.

Step 14 From the Confirm Delete Volumes window, enter **Yes** and then click **OK**.

Figure 10-9 Confirm Delete Arrays Window

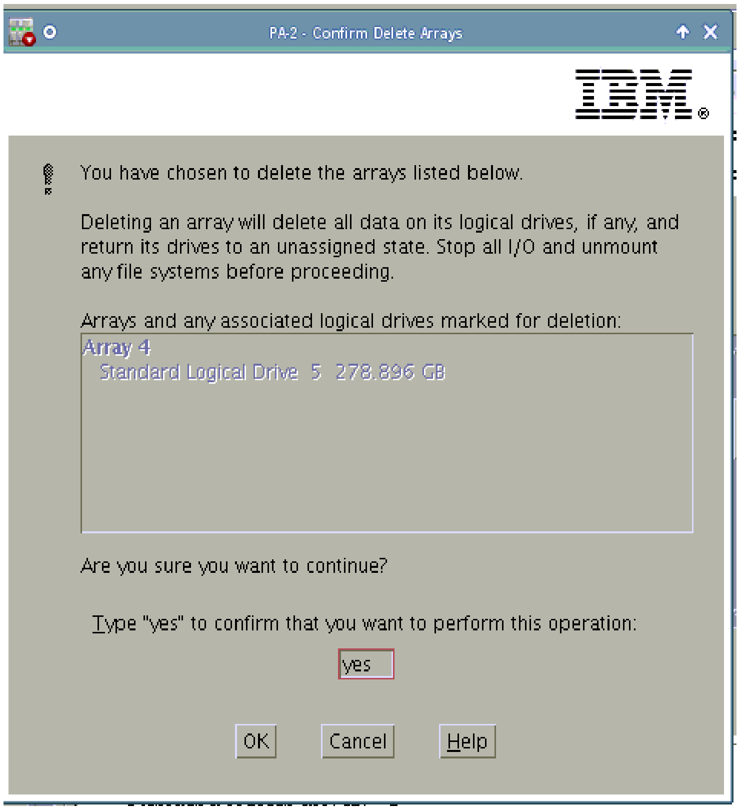

You have chosen to delete the arrays listed below.

Deleting an array will delete all data on its logical drives, if any, and return its drives to an unassigned state. Stop all I/O and unmount any file systems before proceeding.

Arrays and any associated logical drives marked for deletion:

Array 4
Standard Logical Drive 5 278.896 GB

Are you sure you want to continue?

Type "yes" to confirm that you want to perform this operation: yes

Step 15 From the SSH session to the VDS TC management server, enter the command **su admin** to log into the VDS TC CLI. The CLI prompt console> appears.

Step 16 Enter the command **enable** to enter Enable mode. At the password prompt that appears, enter the enable password. The CLI prompt console# appears.

Step 17 Enter the command **cache volume remove** to remove the faulty volume. After the command finds the volume to remove, it will display the prompt “Are you sure?”. Enter **Y** at this prompt to continue to remove the drive. For example:

```
console# cache volume remove
Searching volume(s) to remove...
Volume 43(/dev/sdu1) is going to be removed from CMDB. Are you sure? [y|N]
: y


Done.
Volume 43 is removed.
. . .
```

Step 18 Following the directions for the NetApp E2724 storage enclosure, insert the replacement drive.

Step 19 From the SANtricity Storage Manager window in the VNC client window, click the **Storage & Copy Services** tab.

Step 20 Under the Storage Array in the Navigation Pane on the left, right-click the **Total Unconfigured Capacity** and choose **Create Volume Group**.

Figure 10-10 Create Volume Group

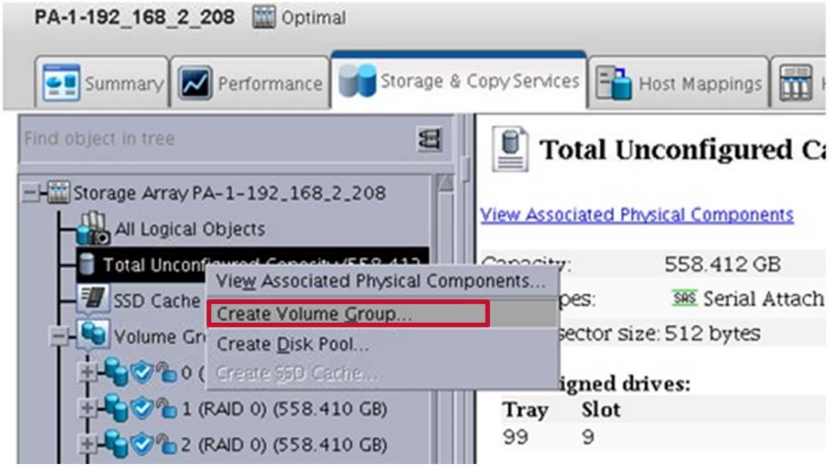

Step 21 From the initial Wizard window, click **Next**.

Figure 10-11 Initial Wizard Window

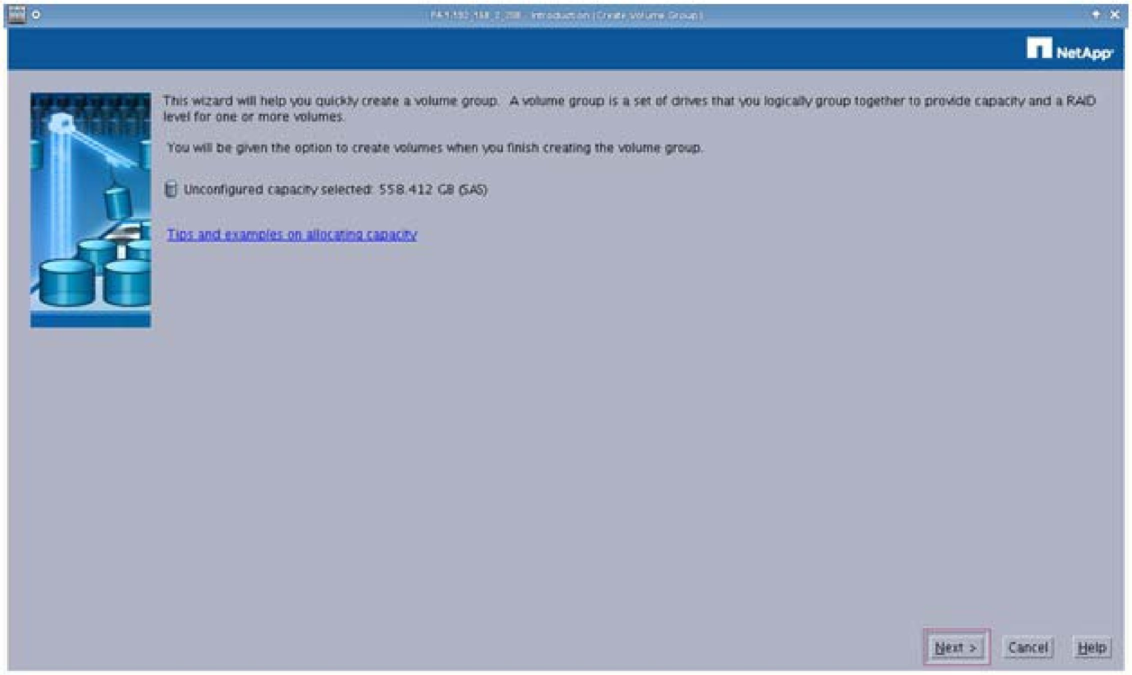

This wizard will help you quickly create a volume group. A volume group is a set of drives that you logically group together to provide capacity and a RAID level for one or more volumes.

You will be given the option to create volumes when you finish creating the volume group.

Unconfigured capacity selected: 558.412 GB (SAS)

Tips and examples on allocating capacity

Step 22 In the next window that appears, do *not* change the name in the Volume Group Name field. You will use the default Volume Group Name. In the Drive Selection Choices section of the window, click the **Manual (Advanced)** radio button and click **Next**.

Figure 10-12 Volume Group Name and Drive Selection

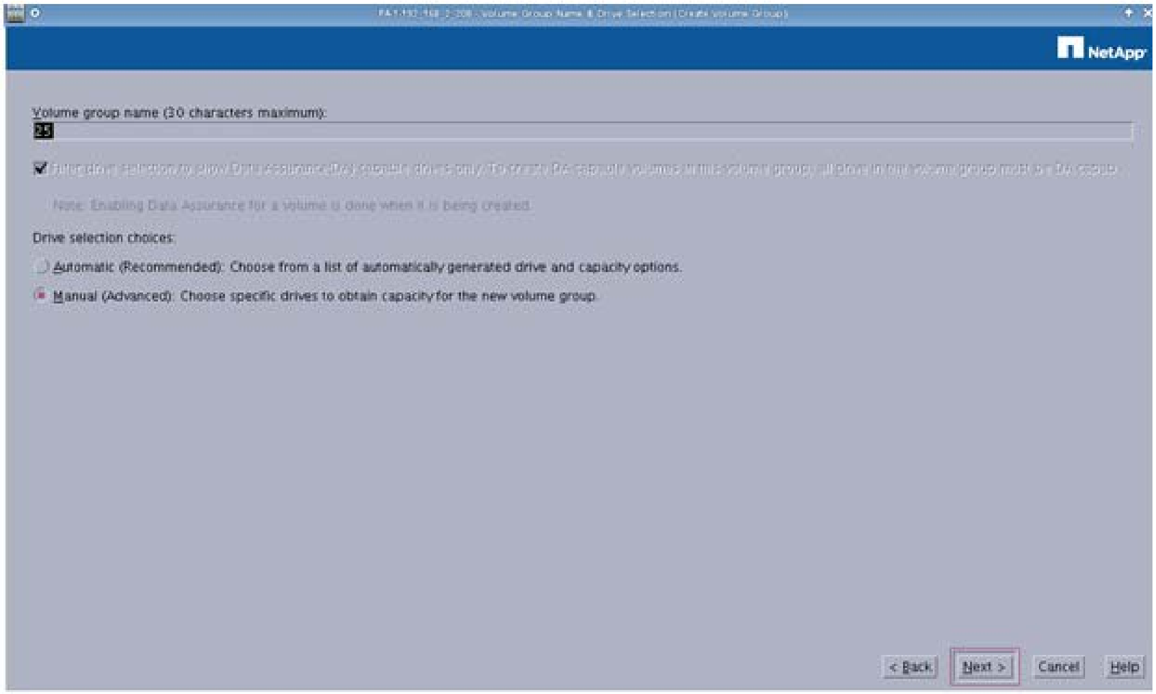

PA1-192.168.2.208 - Volume Group Name & Drive Selection (Create Volume Group)

Volume group name (30 characters maximum):
25

☑ Filter drive selection to show Data Assurance (DA) capable drives only. To create DA capable volumes in this volume group, all drives in the volume group must be DA capable.

Note: Enabling Data Assurance for a volume is done when it is being created.

Drive selection choices:

○ Automatic (Recommended): Choose from a list of automatically generated drive and capacity options.

● Manual (Advanced): Choose specific drives to obtain capacity for the new volume group.

Step 23 From the next window that appears, from the RAID Level drop-down list, choose **RAID 0**. In the Unselected Drives section, choose the drive that is listed and click **Add** to add the drive to the Selected Drives section.

Figure 10-13 RAID Level

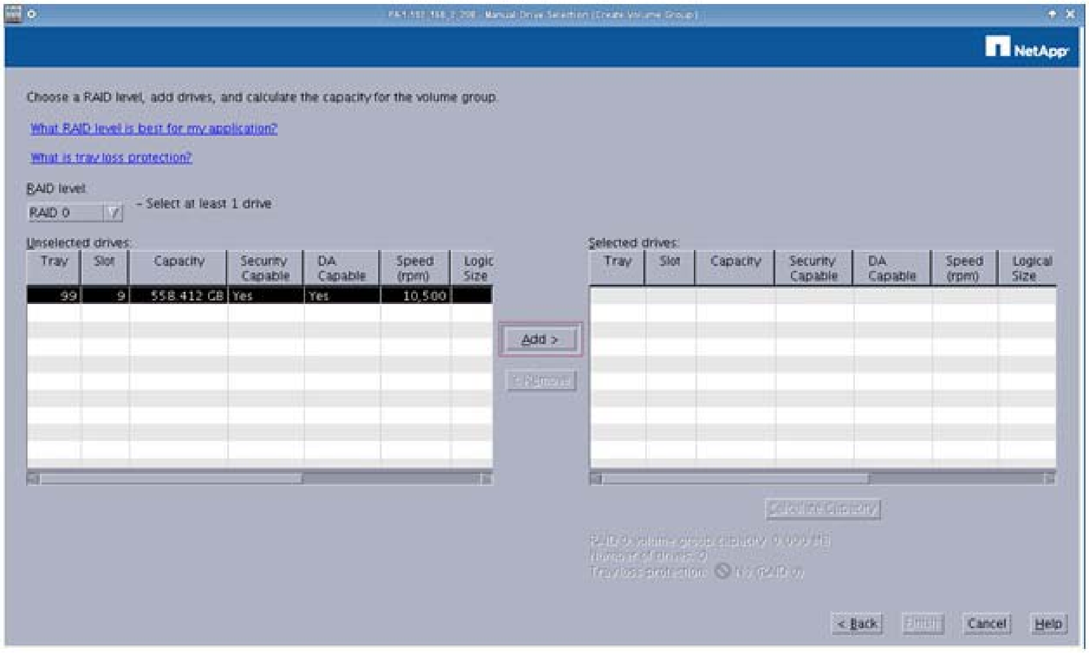

Choose a RAID level, add drives, and calculate the capacity for the volume group.

What RAID level is best for my application?

What is tray loss protection?

RAID level: RAID 0 – Select at least 1 drive

Unselected drives:

Tray	Slot	Capacity	Security Capable	DA Capable	Speed (rpm)	Logic Size
99	9	558.412 GB	Yes	Yes	10,500	

Selected drives:

Tray	Slot	Capacity	Security Capable	DA Capable	Speed (rpm)	Logical Size

RAID 0 volume group capacity: 0.000 TB
Number of drives: 0
Tray loss protection: No (RAID 0)

Step 24 Click **Calculate Capacity** and then click **Finish**.

Figure 10-14 Calculate Capacity

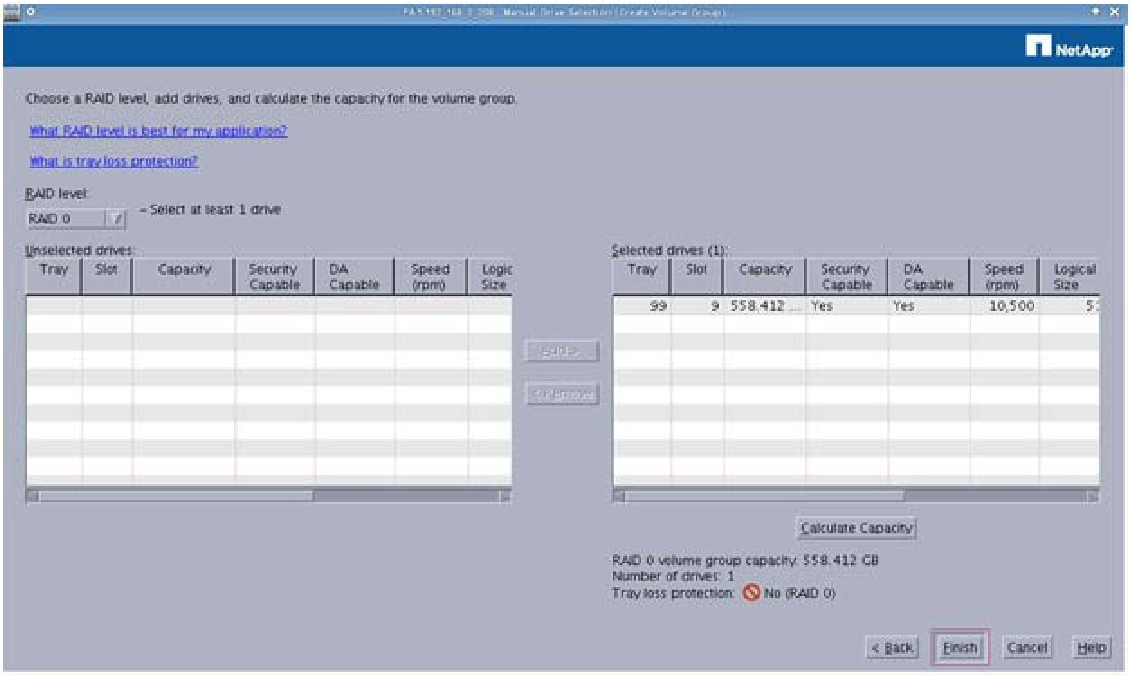

PA-1-192_168_2_208 - Manual Drive Selection (Create Volume Group)

Choose a RAID level, add drives, and calculate the capacity for the volume group.

What RAID level is best for my application?

What is tray loss protection?

RAID level:

RAID 0 – Select at least 1 drive

Unselected drives:

Tray	Slot	Capacity	Security Capable	DA Capable	Speed (rpm)	Logic Size

Selected drives (1):

Tray	Slot	Capacity	Security Capable	DA Capable	Speed (rpm)	Logical Size
99	9	558.412 ...	Yes	Yes	10,500	5

RAID 0 volume group capacity: 558.412 GB
Number of drives: 1
Tray loss protection: 🚫 No (RAID 0)

Step 25 From the Volume Group was Successfully Created window, click **Yes** and then **OK**.

Figure 10-15 Volume Group Successfully Created

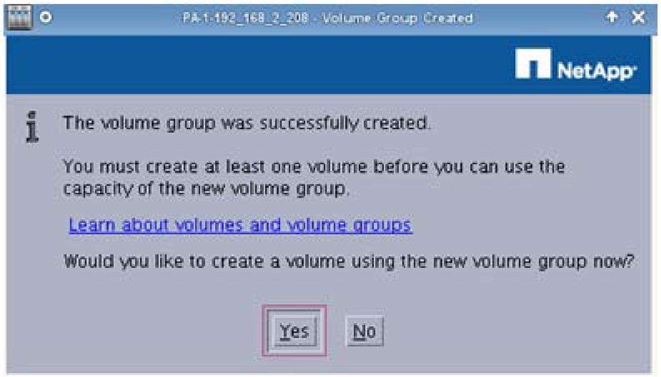

PA-1-192_168_2_208 - Volume Group Created

The volume group was successfully created.

You must create at least one volume before you can use the capacity of the new volume group.

Learn about volumes and volume groups

Would you like to create a volume using the new volume group now?

Step 26 From the Data Assurance Not Supported window, click

Figure 10-16 Data Assurance Not Supported Window

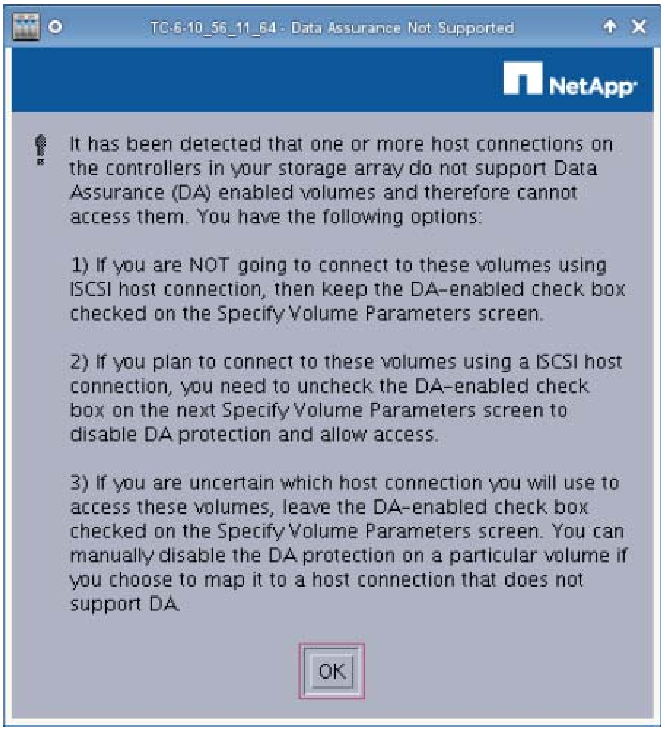

It has been detected that one or more host connections on the controllers in your storage array do not support Data Assurance (DA) enabled volumes and therefore cannot access them. You have the following options:

1) If you are NOT going to connect to these volumes using ISCSI host connection, then keep the DA-enabled check box checked on the Specify Volume Parameters screen.

2) If you plan to connect to these volumes using a ISCSI host connection, you need to uncheck the DA-enabled check box on the next Specify Volume Parameters screen to disable DA protection and allow access.

3) If you are uncertain which host connection you will use to access these volumes, leave the DA-enabled check box checked on the Specify Volume Parameters screen. You can manually disable the DA protection on a particular volume if you choose to map it to a host connection that does not support DA.

Step 27 From the Create Volume window, set the following values:

- **a.** In the New Volume Capacity text box, enter the number displayed in the Free Capacity field, which is the field right above the New Logical Drive Capacity text box.
- **b.** Do *not* change the Volume Name. (Leave the default.)
- **c.** From the Map to Host drop-down list choose **Map Later**.
- **d.** Uncheck the Enable data assurance (DA) protection on the new volume check box to disable this option.
- **e.** In the Volume I/O Characteristics Type drop-down list, choose **Custom**.
- **f.** Make sure that the Enable Dynamic Cache Read Prefetch check box is *unchecked*.
- **g.** From the Segment Size drop-down list, choose **256KB**.
- **h.** Click **Finish**.

Figure 10-17 Create Volume

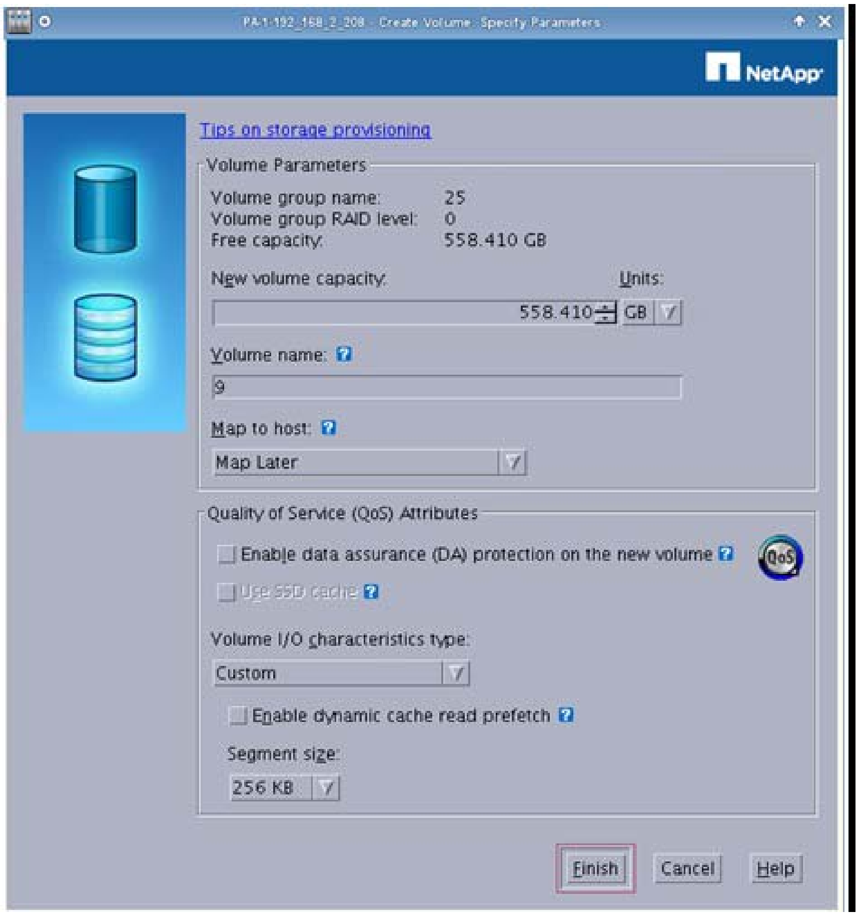

PA-1-192_168_2_208 - Create Volume - Specify Parameters

Tips on storage provisioning

Volume Parameters

Volume group name: 25
Volume group RAID level: 0
Free capacity: 558.410 GB

New volume capacity: 558.410
Units: GB

Volume name: 9

Map to host: Map Later

Quality of Service (QoS) Attributes

□ Enable data assurance (DA) protection on the new volume

□ Use SSD cache

Volume I/O characteristics type: Custom

□ Enable dynamic cache read prefetch

Segment size: 256 KB

Step 28 In the Subsystem Management window, from the Storage & Copy Services tab, right-click the newly added volume and choose **Change** > **Cache Settings**.

Figure 10-18 Change Cache Settings

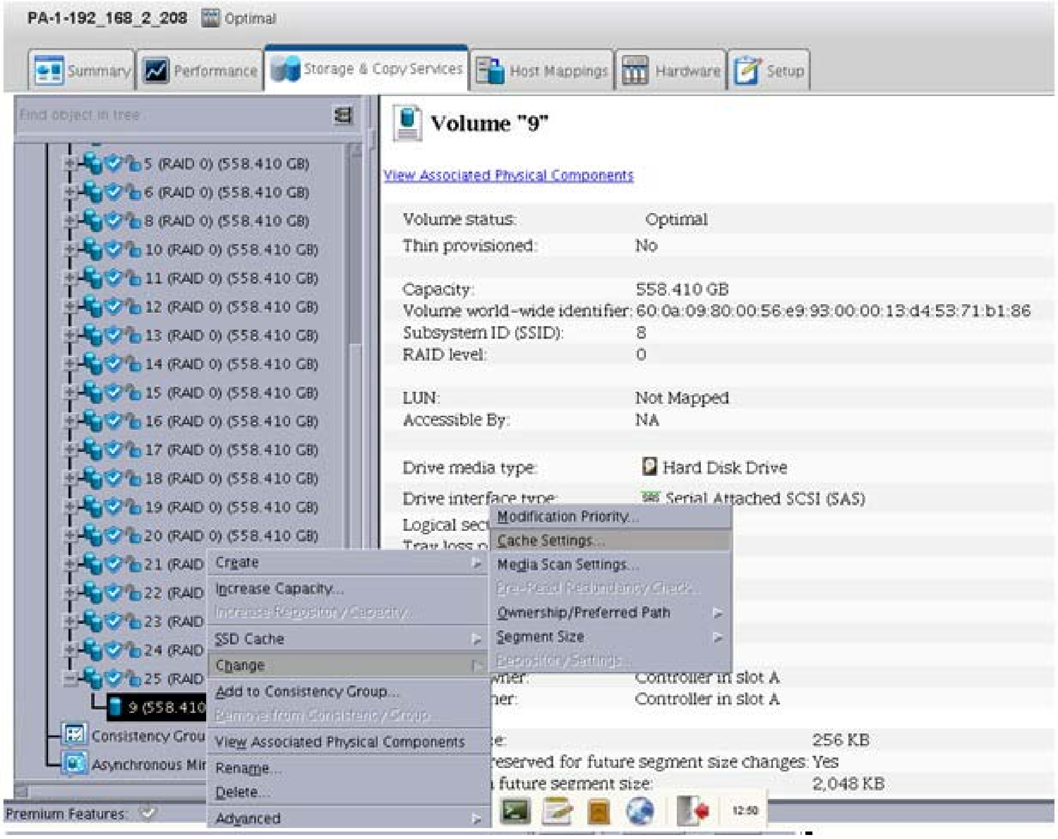

Volume "9"

Volume status: Optimal
Thin provisioned: No

Capacity: 558.410 GB
Volume world-wide identifier: 60:0a:09:80:00:56:e9:93:00:00:13:d4:53:71:b1:86
Subsystem ID (SSID): 8
RAID level: 0

LUN: Not Mapped
Accessible By: NA

Drive media type: Hard Disk Drive
Drive interface type: Serial Attached SCSI (SAS)
Logical sec
Tray loss p
wner: Controller in slot A
ner: Controller in slot A

e: 256 KB
eserved for future segment size changes: Yes
future segment size: 2,048 KB

Step 29 In the Change Cache Settings window, make sure the newly added volume is selected in the Select Volumes section and set the following values:

a. Enable Read Caching check box: **checked**

b. Enable Dynamic Cache Read Prefetch check box: **unchecked**

c. Enable Write Caching check box: **checked**

d. Enable Write Caching without Batteries check box: **unchecked**

e. Enable Write Caching with Mirroring check box: **unchecked**

Figure 10-19 Change Cache Settings Values

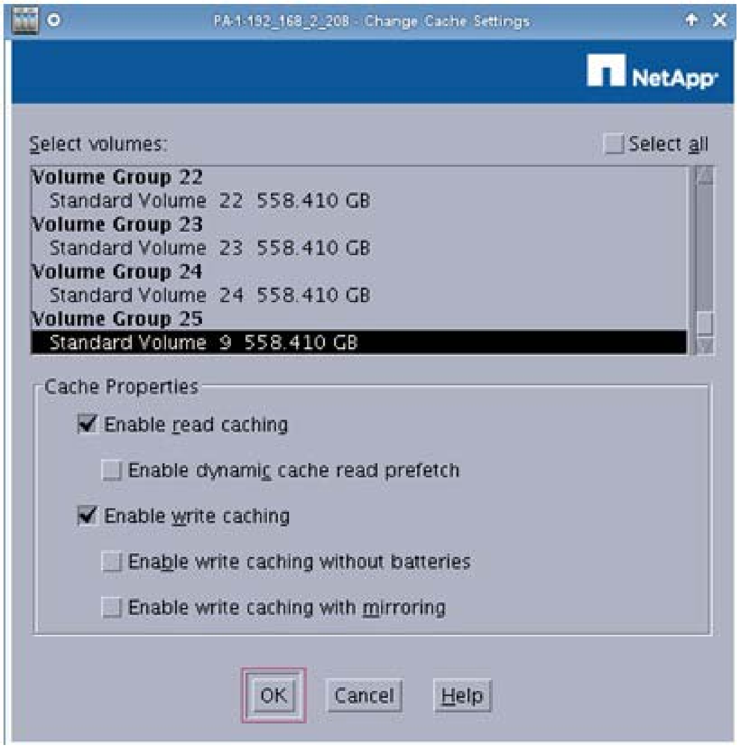

PA-1-192_168_2_208 - Change Cache Settings

Select volumes: □ Select all

Volume Group 22
Standard Volume 22 558.410 GB
Volume Group 23
Standard Volume 23 558.410 GB
Volume Group 24
Standard Volume 24 558.410 GB
Volume Group 25
Standard Volume 9 558.410 GB

Cache Properties

☑ Enable read caching

□ Enable dynamic cache read prefetch

☑ Enable write caching

□ Enable write caching without batteries

□ Enable write caching with mirroring

Step 30 Click **OK**, click **Yes**, and then click **OK** to complete changing the cache settings.

Step 31 From the Subsystem Management window, click the **Host Mappings** tab.

Step 32 Open the **Undefined Mapping** folder, right-click the unmapped LUN, and choose **Add LUN Mapping**. The unmapped LUN will have the format of "#=LUN?". In the example the unmapped LUN shows as "9 = LUN?".

Figure 10-20 Add LUN Mapping

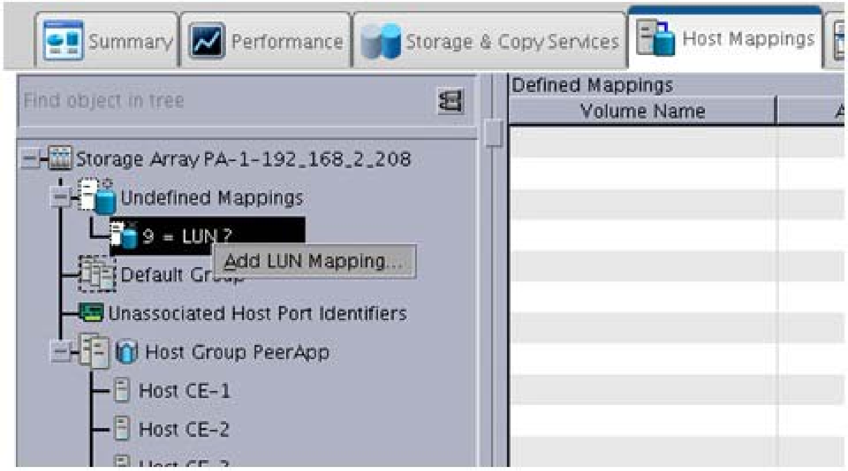

Step 33 From the Define Additional Mapping window, from the Host Group or Host drop-down list, choose **Host Group TC**. In the Logical Unit Number drop-down list, choose a number that has *NEVER* been assigned to any disk before. Use the information that you gathered in Step 9.

Warning **Do not choose the same LUN number that was previously assigned to the disk that you replaced, or any other disk.**

Step 34 Click **Add** and then click **Close**.

Figure 10-21 Define Additional Mapping

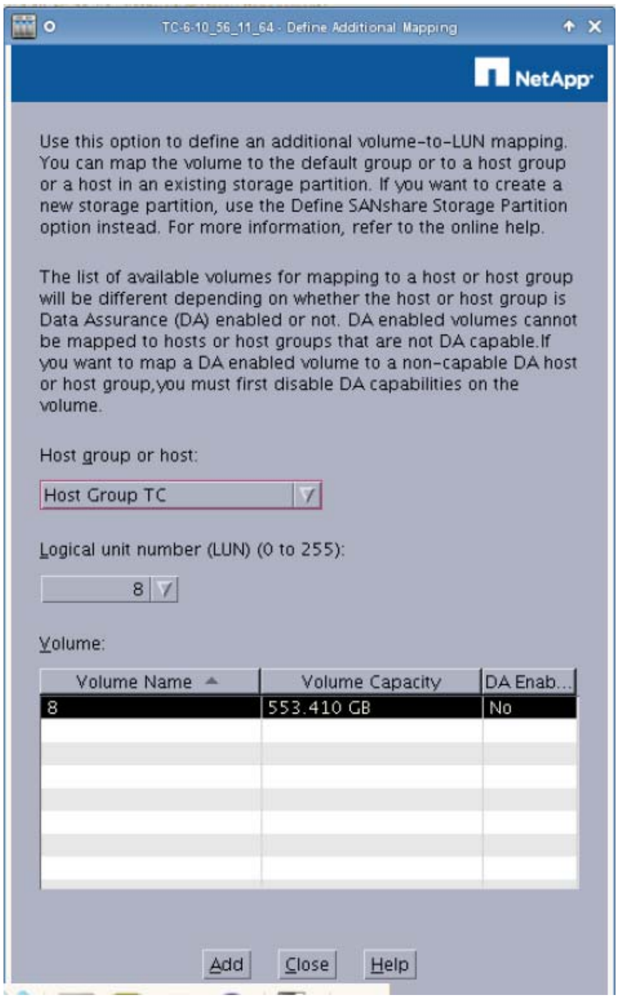

TC-6-10_56_11_64 - Define Additional Mapping

Use this option to define an additional volume-to-LUN mapping. You can map the volume to the default group or to a host group or a host in an existing storage partition. If you want to create a new storage partition, use the Define SANshare Storage Partition option instead. For more information, refer to the online help.

The list of available volumes for mapping to a host or host group will be different depending on whether the host or host group is Data Assurance (DA) enabled or not. DA enabled volumes cannot be mapped to hosts or host groups that are not DA capable.If you want to map a DA enabled volume to a non-capable DA host or host group,you must first disable DA capabilities on the volume.

Host group or host:

Host Group TC

Logical unit number (LUN) (0 to 255):

8

Volume:

Volume Name	Volume Capacity	DA Enab...
8	553.410 GB	No

Step 35 From the VDS TC CLI in enable mode, enter the command **cache volume insert**. When prompted to Insert disk, enter **Y**. This process will format the drive.

```
console# cache volume insert
Searching for new disks....
Insert disk /dev/sdaw? [y/n] y


disk /dev/sdaw is inserted at index 43
Updating blades...
    ce-1 is updated
    ce-2 is updated
    ce-3 is updated
    ce-4 is updated

Procedure complete
```

Step 36 From the Cisco VDS TC Manager, choose **Statistics** > **Storage**. Confirm that the volume that you replaced shows “active”.

Step 37 Choose **Status** > **Storage** and then click the **Detailed Status** tab to see detailed information about the status of the storage controllers and information about any possible drives that are having problems. When the detailed storage status information appears, confirm that no new yellow boxes appear.

CHAPTER 11

Storage Disk Replacement Procedure for IBM Storage Enclosure

This chapter discusses how to determine that a disk has failed in the VDS TC IBM SAN storage enclosure and the steps that are required to replace the disk.

Note If you purchased your IBM storage enclosure from an IBM partner and not from Cisco, please contact the IBM partner for help with the storage disk replacement.

Note If you are having problems with an IBM storage controller, you must open a support case. Replacing an IBM storage controller is not a customer procedure; it must be performed by an authorized IBM technician.

If you purchased the IBM storage enclosure from an IBM partner and not from Cisco, please contact the IBM partner to open a case. If you purchased the IBM storage enclosure from Cisco as part of the VDS TC solution, please contact Cisco TAC to open a case.

Symptoms of a Failed VDS TC Storage Disk

One symptom of a failed storage disk is if the following message frequently appears in the SNMP traps section of the Status > Dashboard window in the VDS TC Manager:

"warning: Volume /dev/*volume_name* (*volume_id*) is turned OFF because of errors"

For example:

```
"warning: Volume /dev/sdak3 (id 60) is turned OFF because of errors"
```

Figure 11-1 SNMP Traps Message

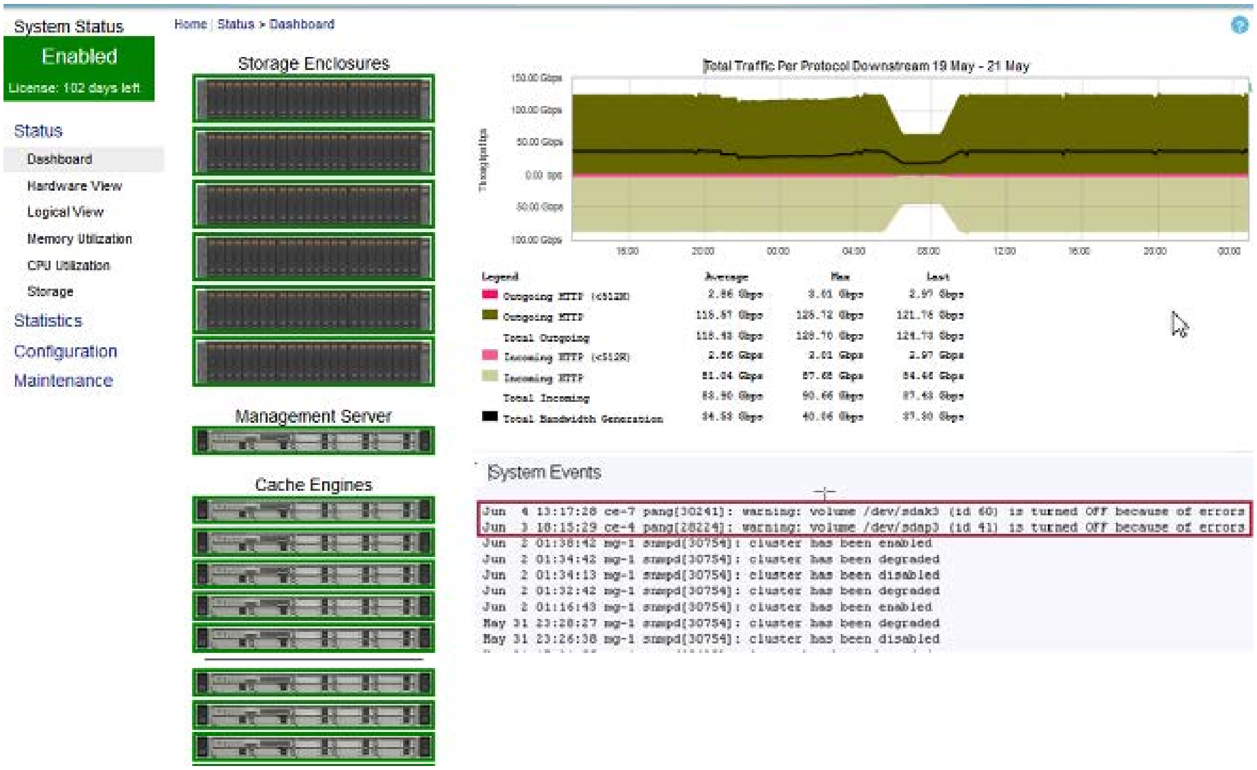

Legend	Average	Max	Last
🟥 Outgoing HTTP (<512K)	2.86 Gbps	3.01 Gbps	2.97 Gbps
🟫 Outgoing HTTP	118.57 Gbps	125.72 Gbps	121.76 Gbps
Total Outgoing	118.43 Gbps	128.70 Gbps	124.73 Gbps
🟥 Incoming HTTP (<512K)	2.86 Gbps	3.01 Gbps	2.97 Gbps
🟩 Incoming HTTP	81.04 Gbps	87.68 Gbps	84.46 Gbps
Total Incoming	83.90 Gbps	90.66 Gbps	87.43 Gbps
■ Total Bandwidth Generation	34.53 Gbps	40.06 Gbps	37.30 Gbps

```
Jun  4 13:17:28 ce-7 pang[30241]: warning: volume /dev/sdak3 (id 60) is turned OFF because of errors
Jun  3 10:15:29 ce-4 pang[28224]: warning: volume /dev/sdap3 (id 41) is turned OFF because of errors
Jun  2 01:38:42 mg-1 snmpd[30754]: cluster has been enabled
Jun  2 01:34:42 mg-1 snmpd[30754]: cluster has been degraded
Jun  2 01:34:13 mg-1 snmpd[30754]: cluster has been disabled
Jun  2 01:32:42 mg-1 snmpd[30754]: cluster has been degraded
Jun  2 01:16:43 mg-1 snmpd[30754]: cluster has been enabled
May 31 23:28:27 mg-1 snmpd[30754]: cluster has been degraded
May 31 23:26:38 mg-1 snmpd[30754]: cluster has been disabled
```

Please note that this error message can also be caused by problems other than a failed disk. To see more information about the failure and to determine if it is from a failed storage disk, follow these steps:

Step 1 Log into Cisco VDS TC Manager. From the main window that appears, choose **Status** > **Storage**.

Note The status of the storage controllers will appear as Operational, even if one of the disks being managed has failed.

Figure 11-2 Status > Storage Window

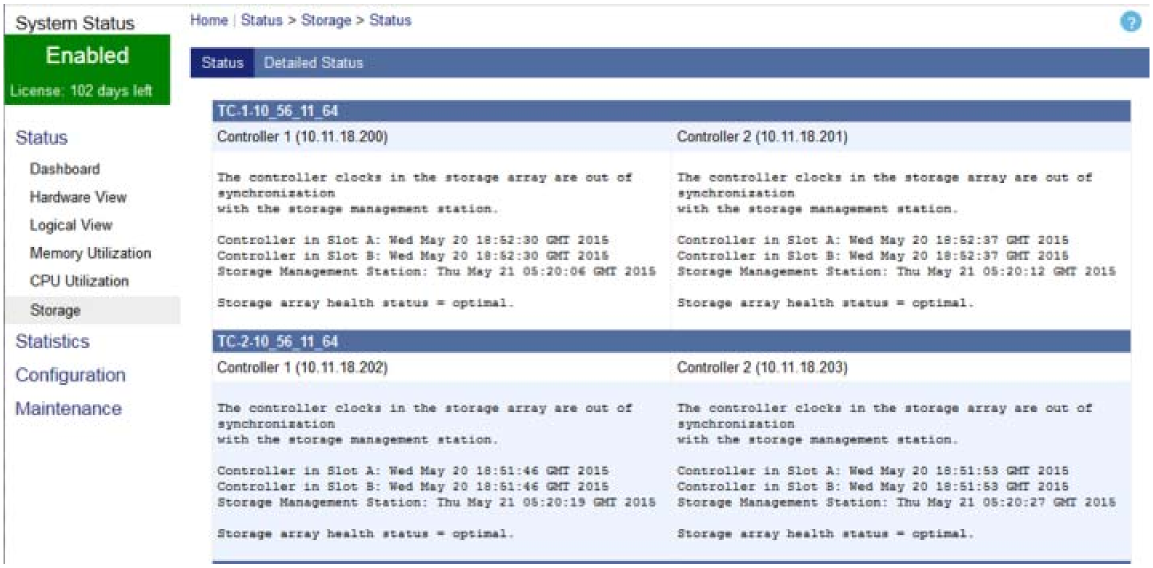

TC-1-10_56_11_64

Controller 1 (10.11.18.200)

```
The controller clocks in the storage array are out of
synchronization
with the storage management station.

Controller in Slot A: Wed May 20 18:52:30 GMT 2015
Controller in Slot B: Wed May 20 18:52:30 GMT 2015
Storage Management Station: Thu May 21 05:20:06 GMT 2015

Storage array health status = optimal.
```

Controller 2 (10.11.18.201)

```
The controller clocks in the storage array are out of
synchronization
with the storage management station.

Controller in Slot A: Wed May 20 18:52:37 GMT 2015
Controller in Slot B: Wed May 20 18:52:37 GMT 2015
Storage Management Station: Thu May 21 05:20:12 GMT 2015

Storage array health status = optimal.
```

TC-2-10_56_11_64

Controller 1 (10.11.18.202)

```
The controller clocks in the storage array are out of
synchronization
with the storage management station.

Controller in Slot A: Wed May 20 18:51:46 GMT 2015
Controller in Slot B: Wed May 20 18:51:46 GMT 2015
Storage Management Station: Thu May 21 05:20:19 GMT 2015

Storage array health status = optimal.
```

Controller 2 (10.11.18.203)

```
The controller clocks in the storage array are out of
synchronization
with the storage management station.

Controller in Slot A: Wed May 20 18:51:53 GMT 2015
Controller in Slot B: Wed May 20 18:51:53 GMT 2015
Storage Management Station: Thu May 21 05:20:27 GMT 2015

Storage array health status = optimal.
```

Step 2 From the Status > Storage window, click the **Detailed Status** tab to see detailed information about the status of the storage controllers and information about any possible drives that are having problems.

Step 3 When the detailed storage status information appears, look for information that appears in yellow boxes. This information will provide details about the storage subsytem and drives that are affected.

Figure 11-3 Detailed Storage Status Output

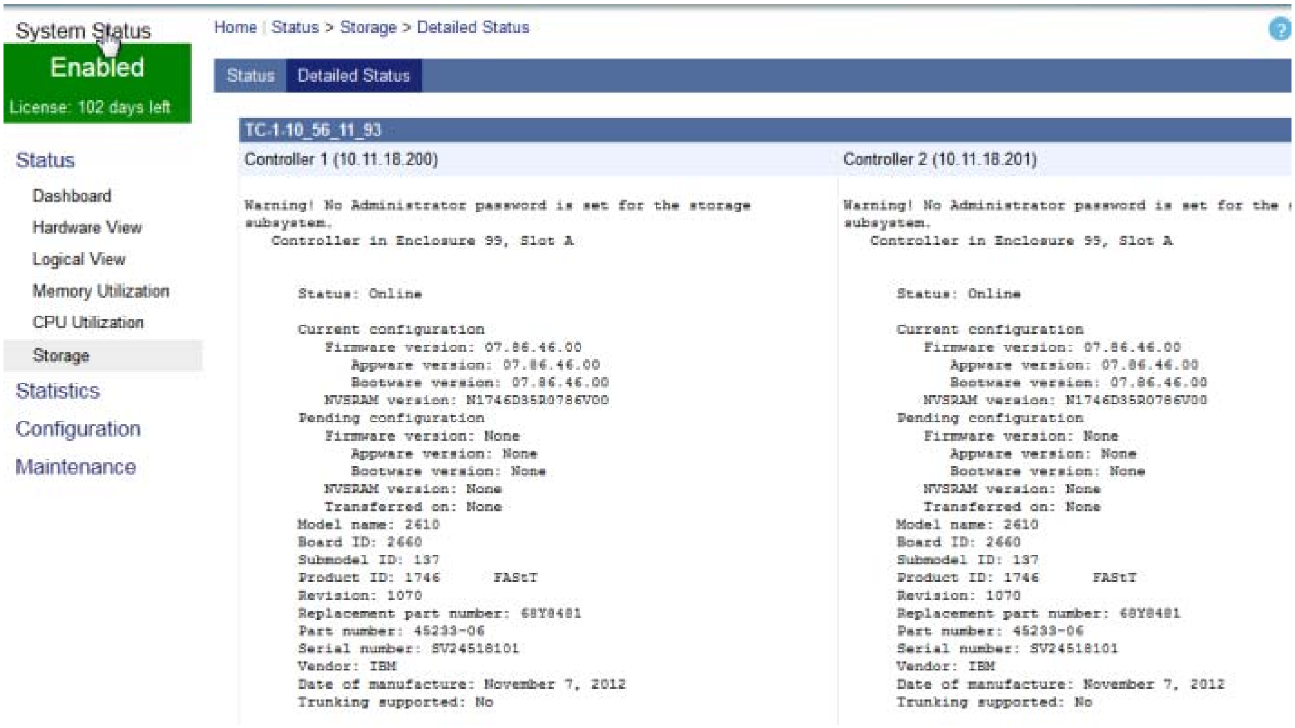

TC-1-10_56_11_93

Controller 1 (10.11.18.200)

```
Warning! No Administrator password is set for the storage
subsystem.
   Controller in Enclosure 99, Slot A


      Status: Online

      Current configuration
         Firmware version: 07.86.46.00
            Appware version: 07.86.46.00
            Bootware version: 07.86.46.00
         NVSRAM version: N1746D35R0786V00
      Pending configuration
         Firmware version: None
            Appware version: None
            Bootware version: None
         NVSRAM version: None
         Transferred on: None
      Model name: 2610
      Board ID: 2660
      Submodel ID: 137
      Product ID: 1746      FAStT
      Revision: 1070
      Replacement part number: 68Y8481
      Part number: 45233-06
      Serial number: SV24518101
      Vendor: IBM
      Date of manufacture: November 7, 2012
      Trunking supported: No
```

Controller 2 (10.11.18.201)

```
Warning! No Administrator password is set for the
subsystem.
   Controller in Enclosure 99, Slot A


      Status: Online

      Current configuration
         Firmware version: 07.86.46.00
            Appware version: 07.86.46.00
            Bootware version: 07.86.46.00
         NVSRAM version: N1746D35R0786V00
      Pending configuration
         Firmware version: None
            Appware version: None
            Bootware version: None
         NVSRAM version: None
         Transferred on: None
      Model name: 2610
      Board ID: 2660
      Submodel ID: 137
      Product ID: 1746      FAStT
      Revision: 1070
      Replacement part number: 68Y8481
      Part number: 45233-06
      Serial number: SV24518101
      Vendor: IBM
      Date of manufacture: November 7, 2012
      Trunking supported: No
```

Note The Volume Usage window in VDS TC Manager, available at Statistics > Storage, may show the Volume State as active even if the drive has failed, as seen in the screenshots below.

Figure 11-4 Volume Usage Tab

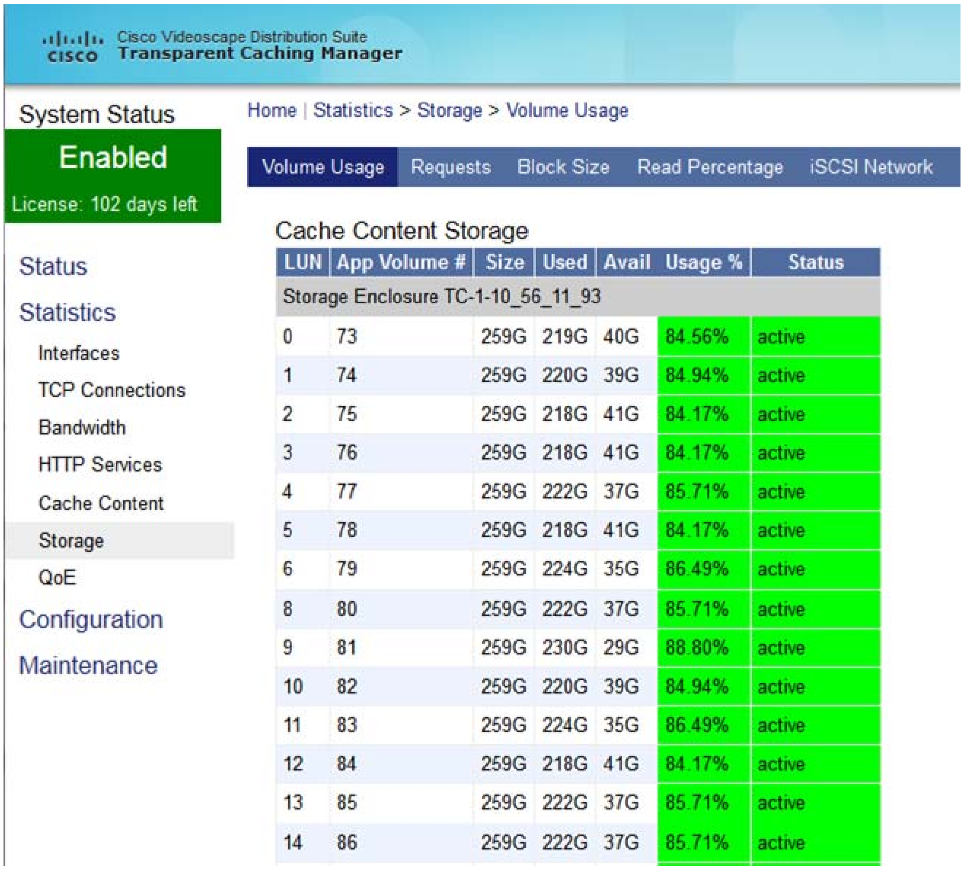

Cache Content Storage

LUN	App Volume #	Size	Used	Avail	Usage %	Status
Storage Enclosure TC-1-10_56_11_93						
0	73	259G	219G	40G	84.56%	active
1	74	259G	220G	39G	84.94%	active
2	75	259G	218G	41G	84.17%	active
3	76	259G	218G	41G	84.17%	active
4	77	259G	222G	37G	85.71%	active
5	78	259G	218G	41G	84.17%	active
6	79	259G	224G	35G	86.49%	active
8	80	259G	222G	37G	85.71%	active
9	81	259G	230G	29G	88.80%	active
10	82	259G	220G	39G	84.94%	active
11	83	259G	224G	35G	86.49%	active
12	84	259G	218G	41G	84.17%	active
13	85	259G	222G	37G	85.71%	active
14	86	259G	222G	37G	85.71%	active

Replacing a Storage Disk

Follow this procedure to replace a storage disk in the VDS TC IBM SAN storage enclosure:

Procedure

Step 1 Using SSH software, open an SSH connection to the management IP address of the VDS TC management server.

Step 2 Log into the system using the username **padmin** and the password that was provided by Cisco.

Step 3 Enter the command **su root** to change to the root user. Enter the password for the root user when prompted. This password is provided by Cisco.

Step 4 To start the VNC server, enter the following commands:

  a. **rm /tmp/.X1-lock** (This step is optional. It is only required if there is a leftover lock file from a previous vncserver instance.)

  b. **rm /tmp/.X11-unix/X1** (This step is optional. It is only required if there is a leftover lock file from a previous vncserver instance.)

  c. **vncserver :1**

Step 5 Use a VNC client to connect from your local computer to the VDS TC management server.

Figure 11-5 VNC Viewer

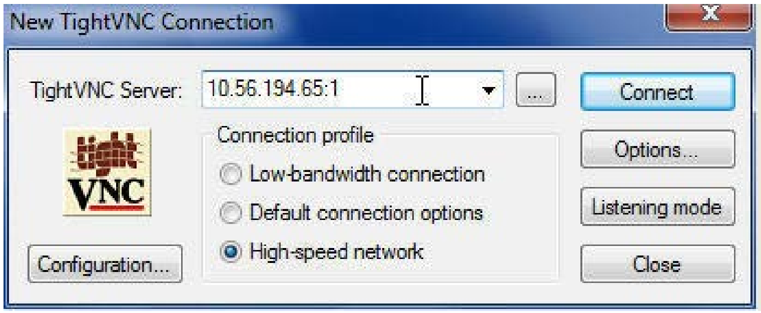

TightVNC Server: 10.56.194.65:1

Connection profile

○ Low-bandwidth connection
○ Default connection options
● High-speed network

Step 6 From the VNC console window, enter the command **/opt/IBM_DS/client/SMclient** to start the IBM Storage Manager.

Step 7 Right-click the storage subsystem for which you are replacing a disk and choose **Manage Storage Subsystem**. Click **No** in any pop-up windows that appear.

Figure 11-6 IBM DS Storage Manager

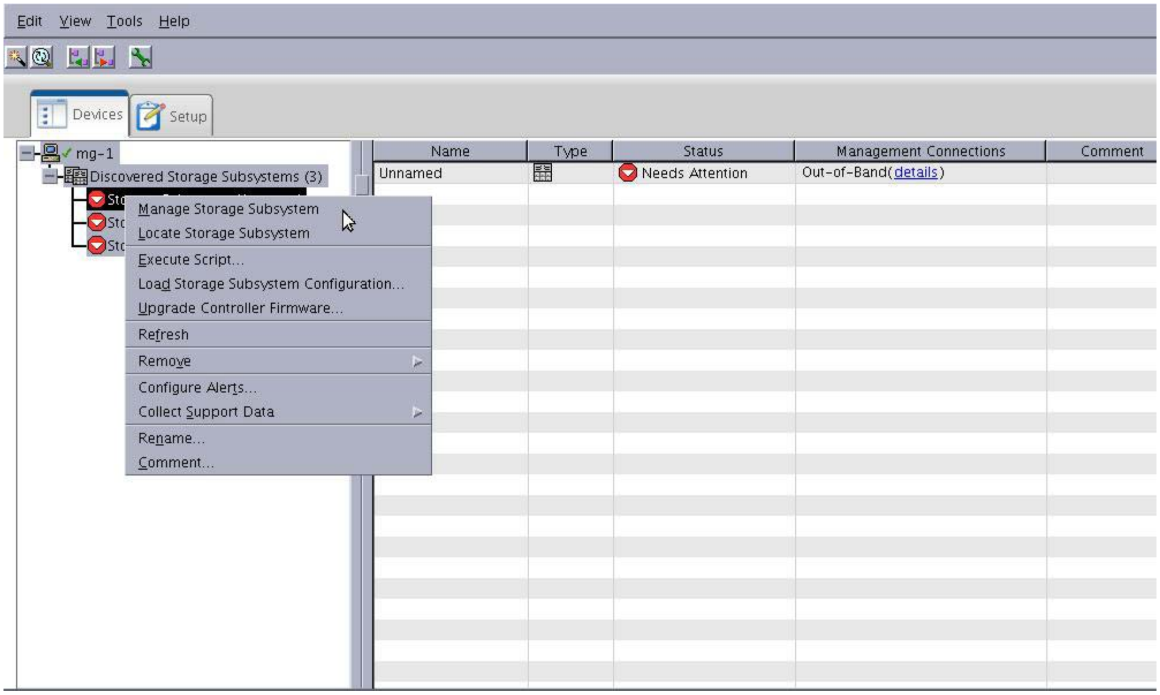

Step 8 From the Subsystem Management window that appears, click the **Host Mappings** tab.

Step 9 Ensure that Storage Subsystem is selected in the navigation pane on the left. In the Defined Mappings table, make note of all of the LUN numbers that are currently in use, including the LUN number of the failed drive. You will need this information in a later step.

Figure 11-7 Exiting LUN Numbers

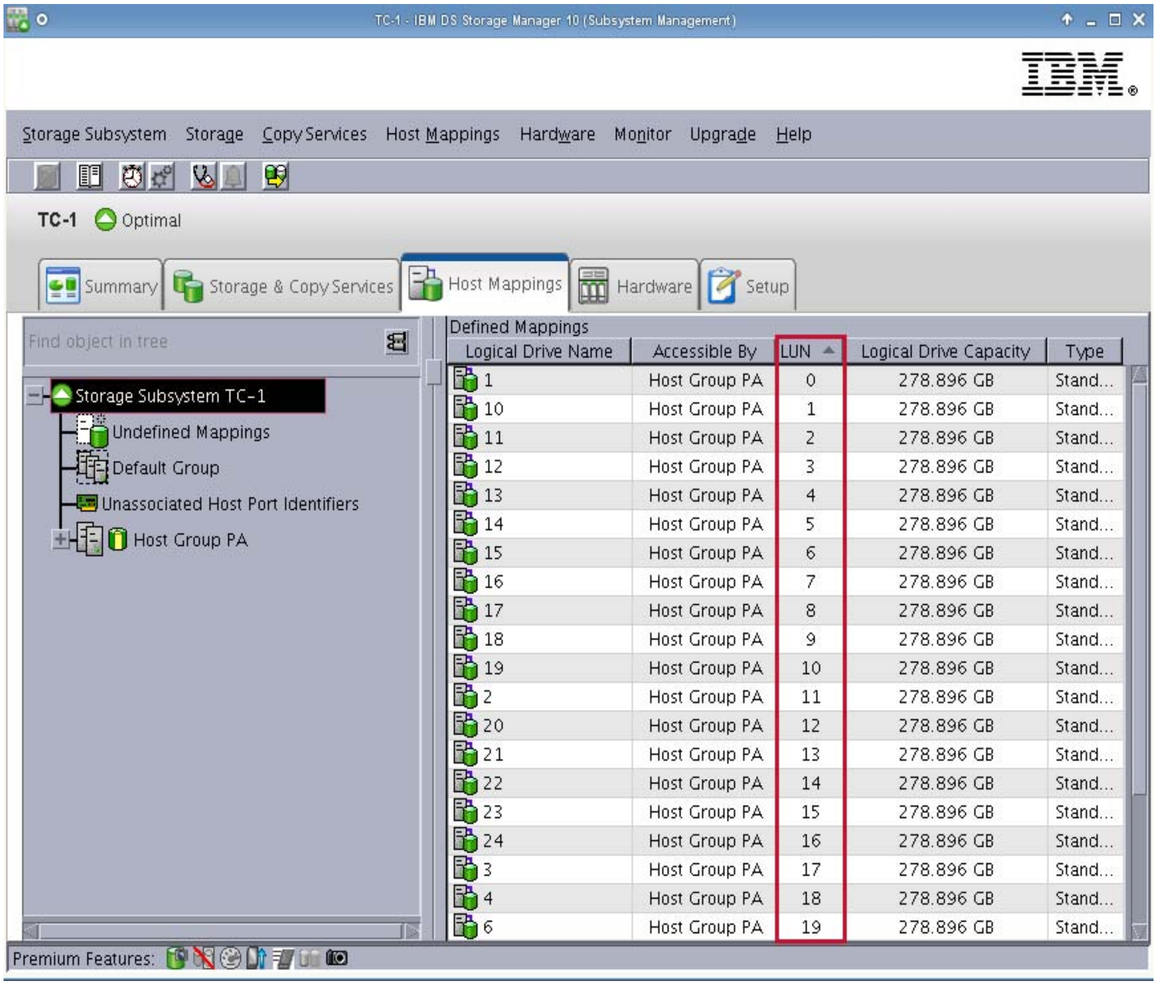

Defined Mappings

Logical Drive Name	Accessible By	LUN	Logical Drive Capacity	Type
1	Host Group PA	0	278.896 GB	Stand...
10	Host Group PA	1	278.896 GB	Stand...
11	Host Group PA	2	278.896 GB	Stand...
12	Host Group PA	3	278.896 GB	Stand...
13	Host Group PA	4	278.896 GB	Stand...
14	Host Group PA	5	278.896 GB	Stand...
15	Host Group PA	6	278.896 GB	Stand...
16	Host Group PA	7	278.896 GB	Stand...
17	Host Group PA	8	278.896 GB	Stand...
18	Host Group PA	9	278.896 GB	Stand...
19	Host Group PA	10	278.896 GB	Stand...
2	Host Group PA	11	278.896 GB	Stand...
20	Host Group PA	12	278.896 GB	Stand...
21	Host Group PA	13	278.896 GB	Stand...
22	Host Group PA	14	278.896 GB	Stand...
23	Host Group PA	15	278.896 GB	Stand...
24	Host Group PA	16	278.896 GB	Stand...
3	Host Group PA	17	278.896 GB	Stand...
4	Host Group PA	18	278.896 GB	Stand...
6	Host Group PA	19	278.896 GB	Stand...

Step 10 In the Subsystem Management window that appears, click the **Storage & Copy Services** tab.

Step 11 In the navigation pane on the left, expand the Arrays folder and locate the array with the faulty drive. Expand the array to confirm the presence of a faulty drive.

Note A red drive icon with an X indicates a failed drive, as shown in the example.

Figure 11-8 Subsystem Management Array

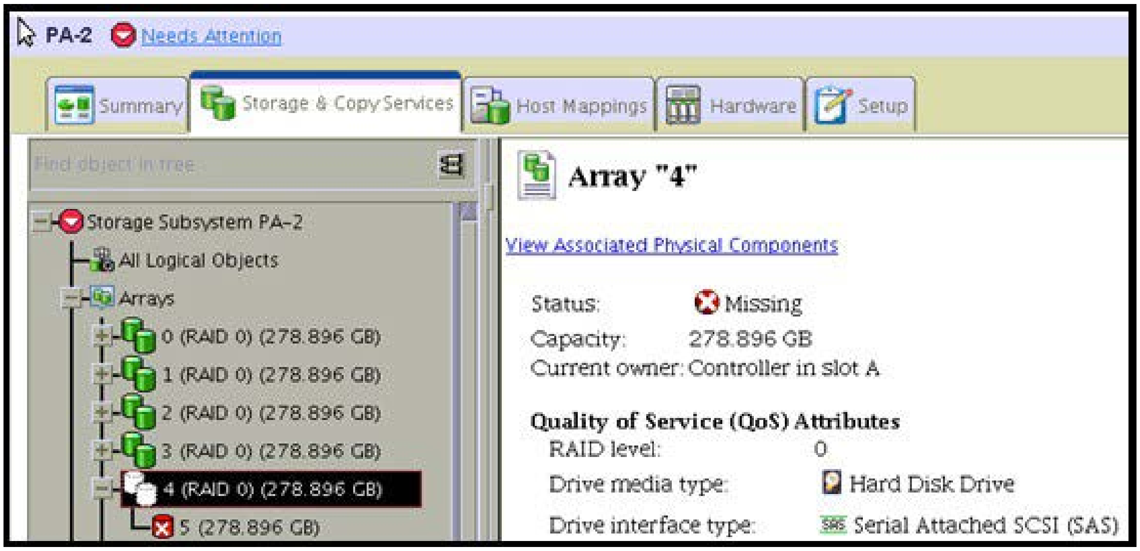

Status: Missing
Capacity: 278.896 GB
Current owner: Controller in slot A

Quality of Service (QoS) Attributes
RAID level: 0
Drive media type: Hard Disk Drive
Drive interface type: Serial Attached SCSI (SAS)

Step 12 Right-click the faulty array (*not* the logical drive), and choose **Delete**.

Figure 11-9 Delete Array

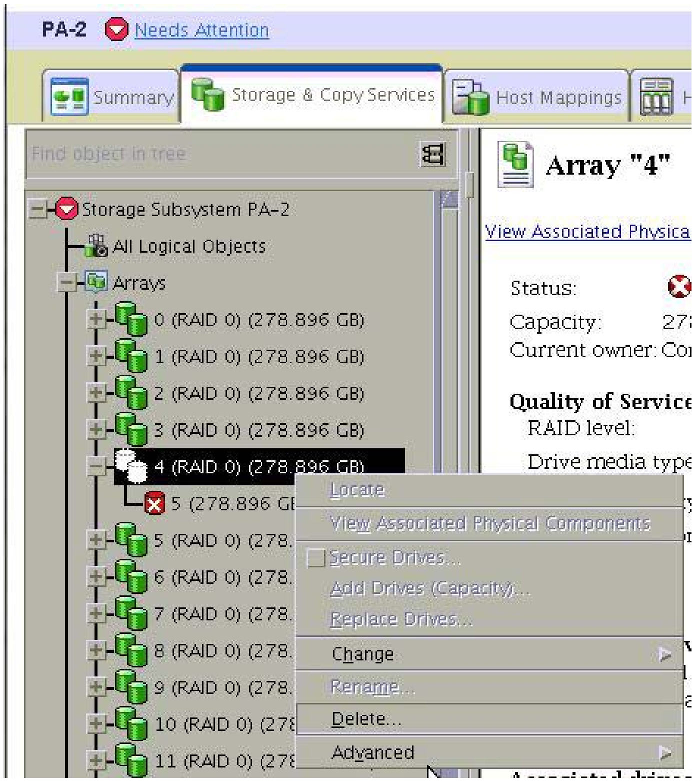

Step 13 From the IBM Delete Arrays window, click **Delete**.

Figure 11-10 Delete Arrays Window

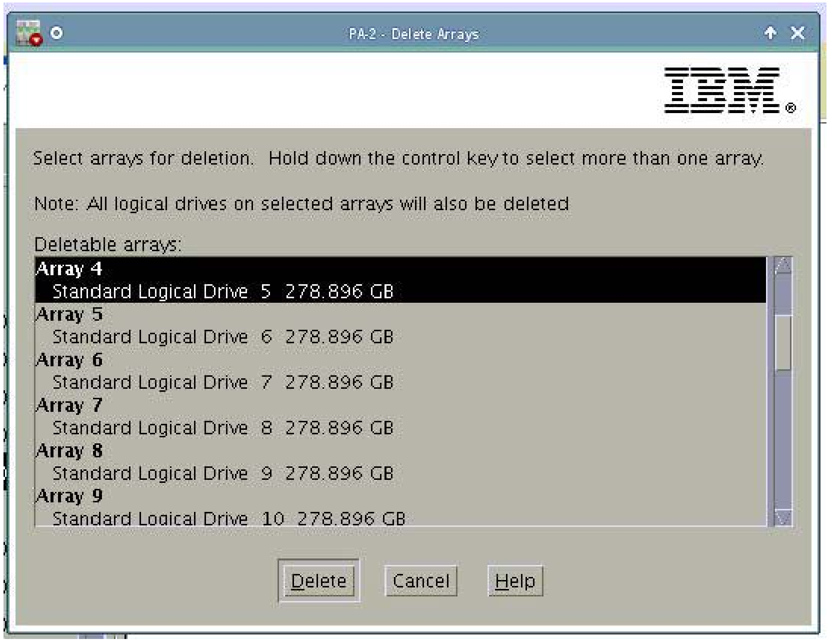

Select arrays for deletion. Hold down the control key to select more than one array.

Note: All logical drives on selected arrays will also be deleted

Deletable arrays:

Array 4
Standard Logical Drive 5 278.896 GB
Array 5
Standard Logical Drive 6 278.896 GB
Array 6
Standard Logical Drive 7 278.896 GB
Array 7
Standard Logical Drive 8 278.896 GB
Array 8
Standard Logical Drive 9 278.896 GB
Array 9
Standard Logical Drive 10 278.896 GB

Step 14 From the Confirm Delete Arrays window, enter **Yes** and then click **OK**.

Figure 11-11 Confirm Delete Arrays Window

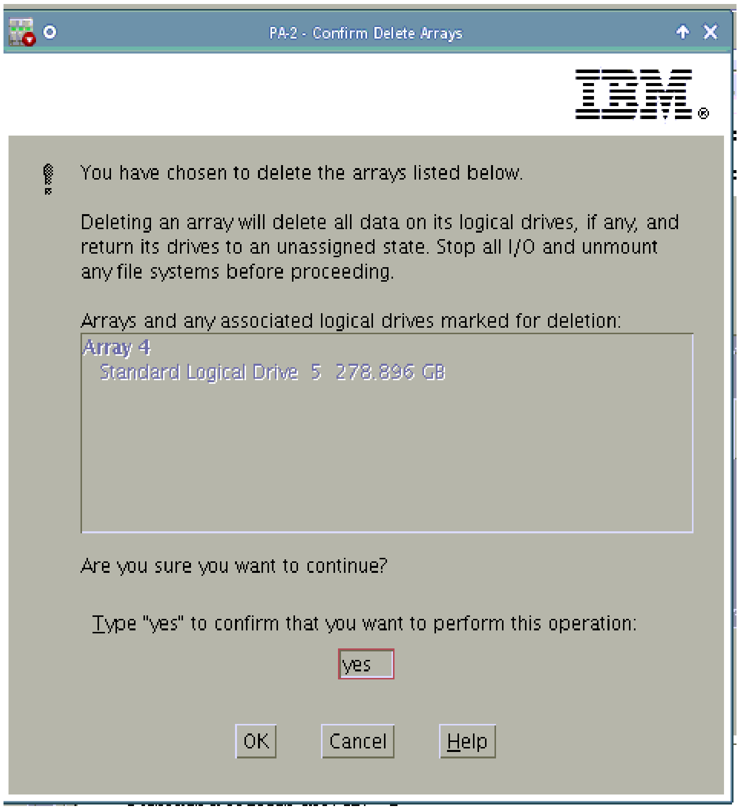

You have chosen to delete the arrays listed below.

Deleting an array will delete all data on its logical drives, if any, and return its drives to an unassigned state. Stop all I/O and unmount any file systems before proceeding.

Arrays and any associated logical drives marked for deletion:

Array 4
  Standard Logical Drive 5 278.896 GB

Are you sure you want to continue?

Type "yes" to confirm that you want to perform this operation:

yes

Step 15 From the SSH session to the VDS TC management server, enter the command **su admin** to log into the VDS TC CLI. The CLI prompt console> appears.

Step 16 Enter the command **enable** to enter Enable mode. At the password prompt that appears, enter the enable password. The CLI prompt console# appears.

Step 17 Enter the command **cache volume remove** to remove the faulty volume. After the command finds the volume to remove, it will display the prompt “Are you sure?”. Enter **Y** at this prompt to continue to remove the drive. For example:

```
console# cache volume remove
Searching volume(s) to remove...
Volume 43(/dev/sdu1) is going to be removed from CMDB. Are you sure? [y|N]
: y


Done.
Volume 43 is removed.
. . .
```

Step 18 Following the directions for the IBM DS3524 storage enclosure, insert the replacement drive and initialize the drive. The new disk will appear as “not initialized.”

Note When you initialize a drive, all data on the drive is lost.

Step 19 From the IBM Storage Manager window in the VNC client window, click the **Storage & Copy Services** tab.

Step 20 Under the Storage Subsystem in the Navigation Pane on the left, right-click the **Total Unconfigured Capacity** and choose **Create Array**.

Figure 11-12 Create Array

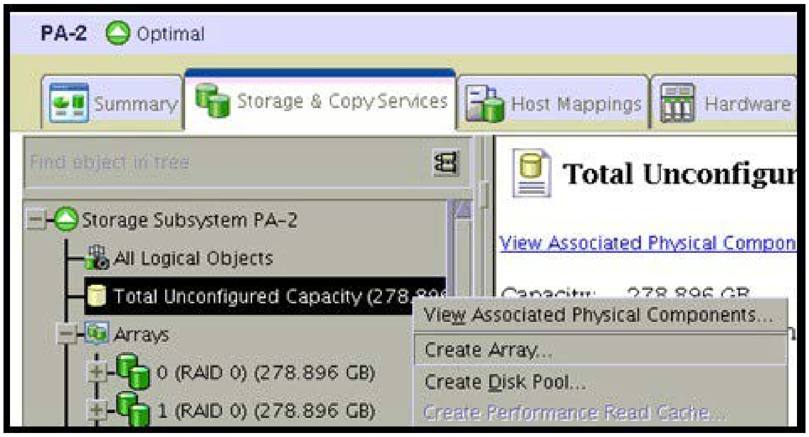

Step 21 From the initial Wizard window, click **Next**.

Figure 11-13 Initial Wizard Window

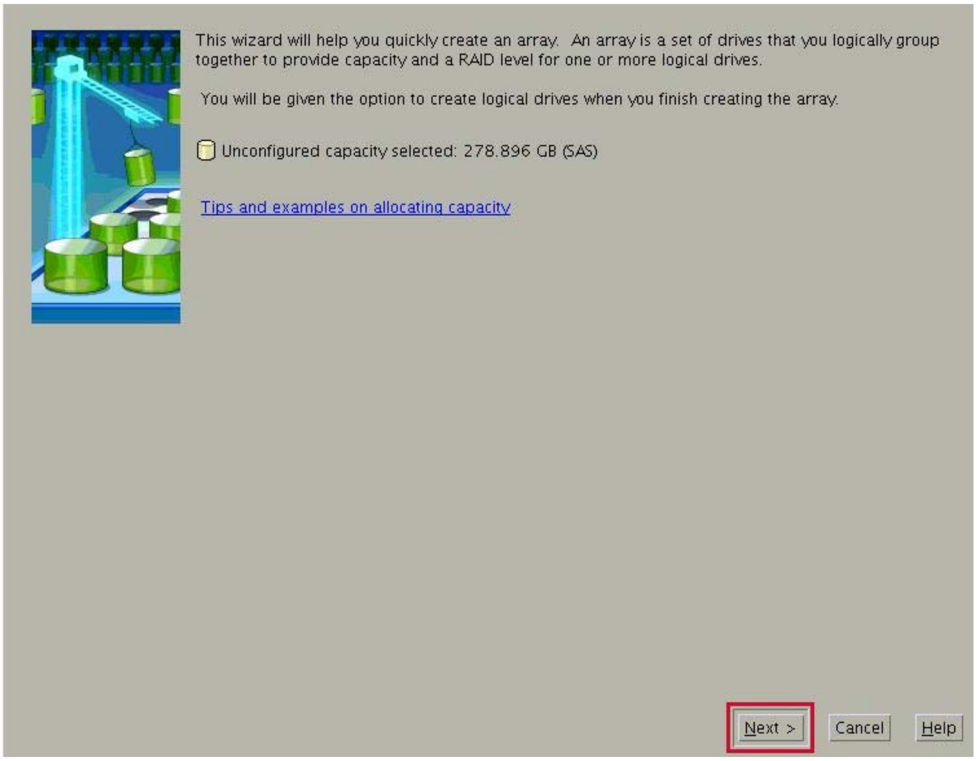

This wizard will help you quickly create an array. An array is a set of drives that you logically group together to provide capacity and a RAID level for one or more logical drives.

You will be given the option to create logical drives when you finish creating the array.

Unconfigured capacity selected: 278.896 GB (SAS)

Tips and examples on allocating capacity

Step 22 In the next window that appears, do *not* change the name in the Array Name field. You will use the default Array Name. In the Drive Selection Choices section of the window, click the **Manual (Advanced)** radio button and click **Next**.

Figure 11-14 Array Name and Drive Selection

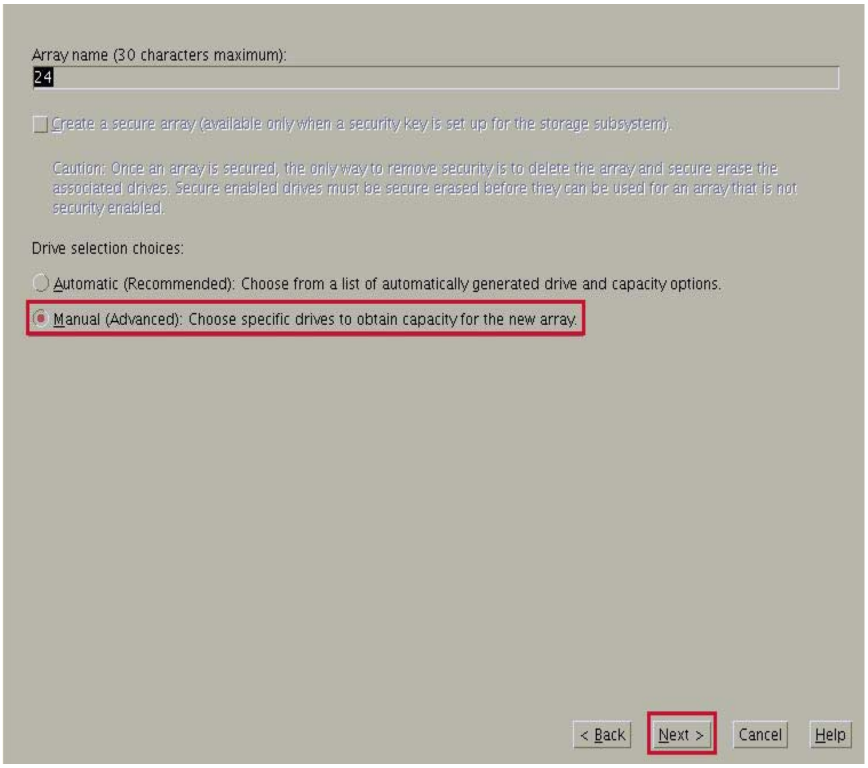

Array name (30 characters maximum):

24

□ Create a secure array (available only when a security key is set up for the storage subsystem).

Caution: Once an array is secured, the only way to remove security is to delete the array and secure erase the associated drives. Secure enabled drives must be secure erased before they can be used for an array that is not security enabled.

Drive selection choices:

○ Automatic (Recommended): Choose from a list of automatically generated drive and capacity options.

● Manual (Advanced): Choose specific drives to obtain capacity for the new array.

Step 23 From the next window that appears, from the RAID Level drop-down list, choose **RAID 0**. In the Unselected Drives section, choose the drive that is listed and click **Add** to add the drive to the Selected Drives section.

Figure 11-15 RAID Level

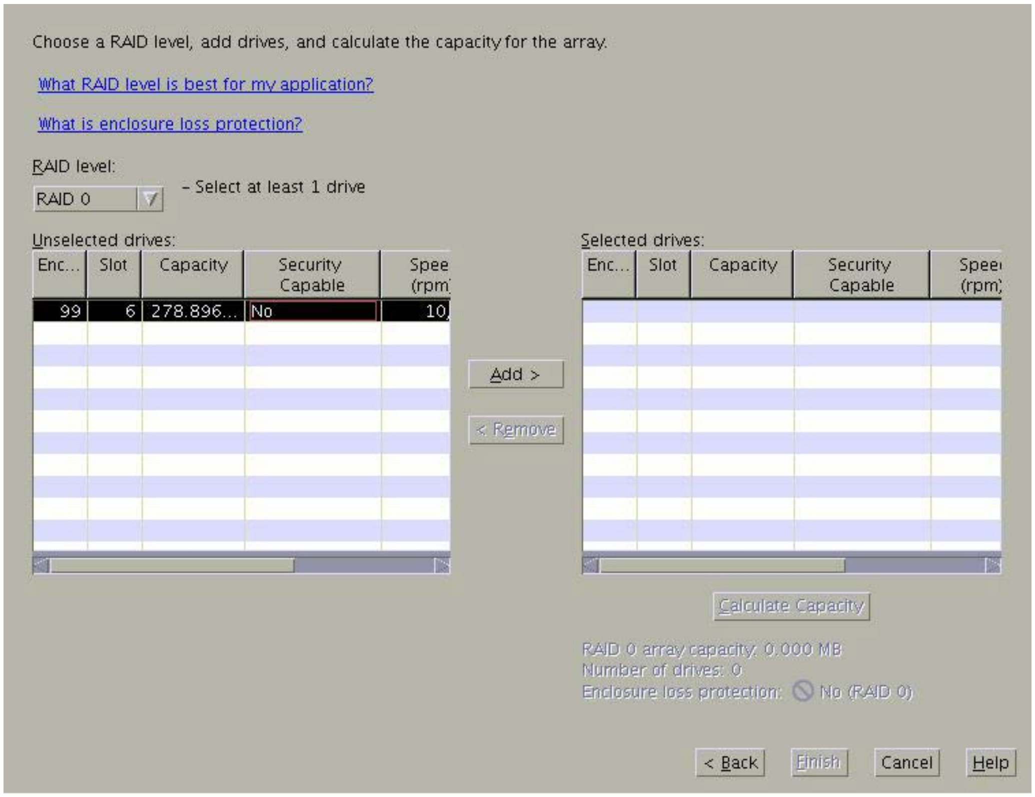

Choose a RAID level, add drives, and calculate the capacity for the array.

What RAID level is best for my application?

What is enclosure loss protection?

RAID level: RAID 0 – Select at least 1 drive

Unselected drives:

Enc...	Slot	Capacity	Security Capable	Spee (rpm)
99	6	278.896...	No	10,

Selected drives:

Enc...	Slot	Capacity	Security Capable	Spee (rpm)

RAID 0 array capacity: 0.000 MB
Number of drives: 0
Enclosure loss protection: No (RAID 0)

Step 24 Click **Calculate Capacity** and then click **Finish**.

Figure 11-16 Calculate Capacity

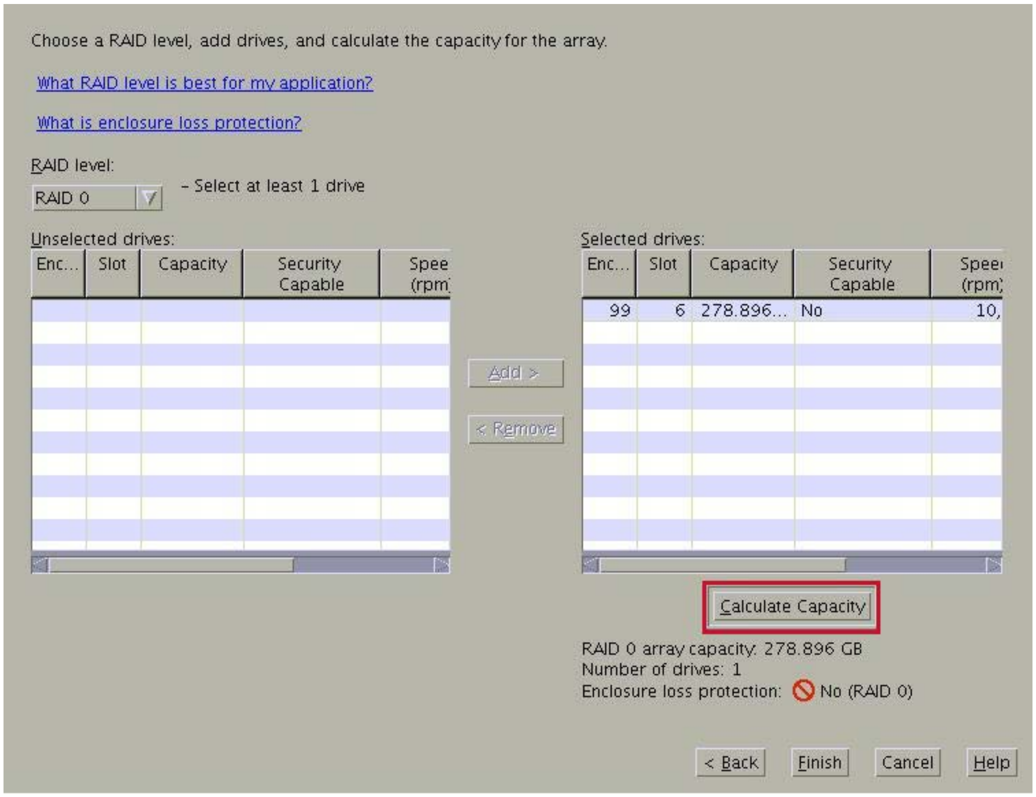

Choose a RAID level, add drives, and calculate the capacity for the array.

What RAID level is best for my application?

What is enclosure loss protection?

RAID level: RAID 0 - Select at least 1 drive

Unselected drives:

Enc...	Slot	Capacity	Security Capable	Spee (rpm

Selected drives:

Enc...	Slot	Capacity	Security Capable	Spee (rpm)
99	6	278.896...	No	10,

RAID 0 array capacity: 278.896 GB
Number of drives: 1
Enclosure loss protection: ⊘ No (RAID 0)

Step 25 From the Array was Successfully Created window, click **Yes**.

Figure 11-17 Array Successfully Created

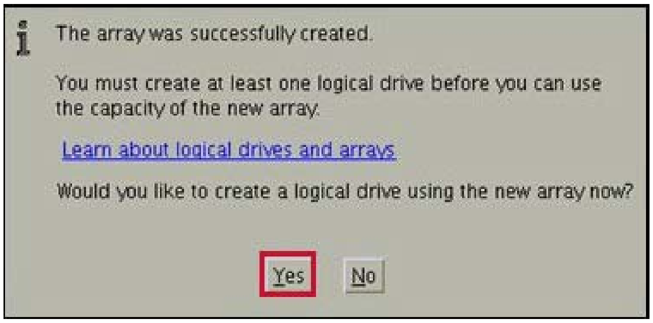

The array was successfully created.

You must create at least one logical drive before you can use the capacity of the new array.

Learn about logical drives and arrays

Would you like to create a logical drive using the new array now?

Step 26 When the PI (Protection Information) Not Supported window appears, click **OK**.

Figure 11-18 PI (Protection Information) Not Supported

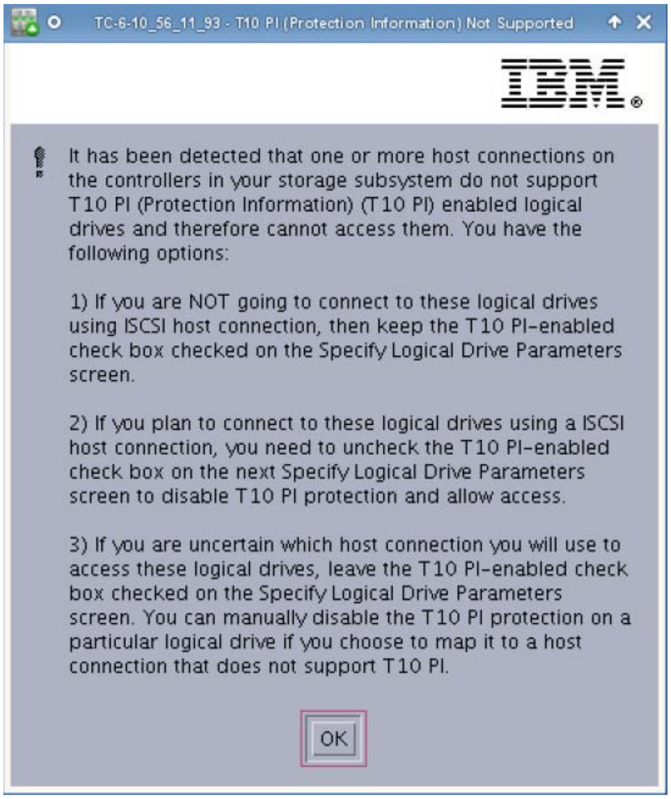
TC-6-10_56_11_93 - T10 PI (Protection Information) Not Supported
IBM.

It has been detected that one or more host connections on the controllers in your storage subsystem do not support T10 PI (Protection Information) (T10 PI) enabled logical drives and therefore cannot access them. You have the following options:

1) If you are NOT going to connect to these logical drives using ISCSI host connection, then keep the T10 PI-enabled check box checked on the Specify Logical Drive Parameters screen.

2) If you plan to connect to these logical drives using a ISCSI host connection, you need to uncheck the T10 PI-enabled check box on the next Specify Logical Drive Parameters screen to disable T10 PI protection and allow access.

3) If you are uncertain which host connection you will use to access these logical drives, leave the T10 PI-enabled check box checked on the Specify Logical Drive Parameters screen. You can manually disable the T10 PI protection on a particular logical drive if you choose to map it to a host connection that does not support T10 PI.

Step 27 From the Create Logical Drive window, set the following values:

a. In the New Logical Drive Capacity text box, enter the number displayed in the Free Capacity field, which is the field right above the New Logical Drive Capacity text box.

b. Do *not* change the Logical Drive Name.

c. From the Map to Host drop-down list choose **Map Later**.

d. In the Logical Drive I/O Characteristics Type drop-down list, choose **Custom**.

e. Make sure that the Enable Dynamic Cache Read Prefetch check box is *unchecked*.

f. From the Segment Size drop-down list, choose **256KB**

g. Click **Finish**.

Figure 11-19 Create Logical Drive

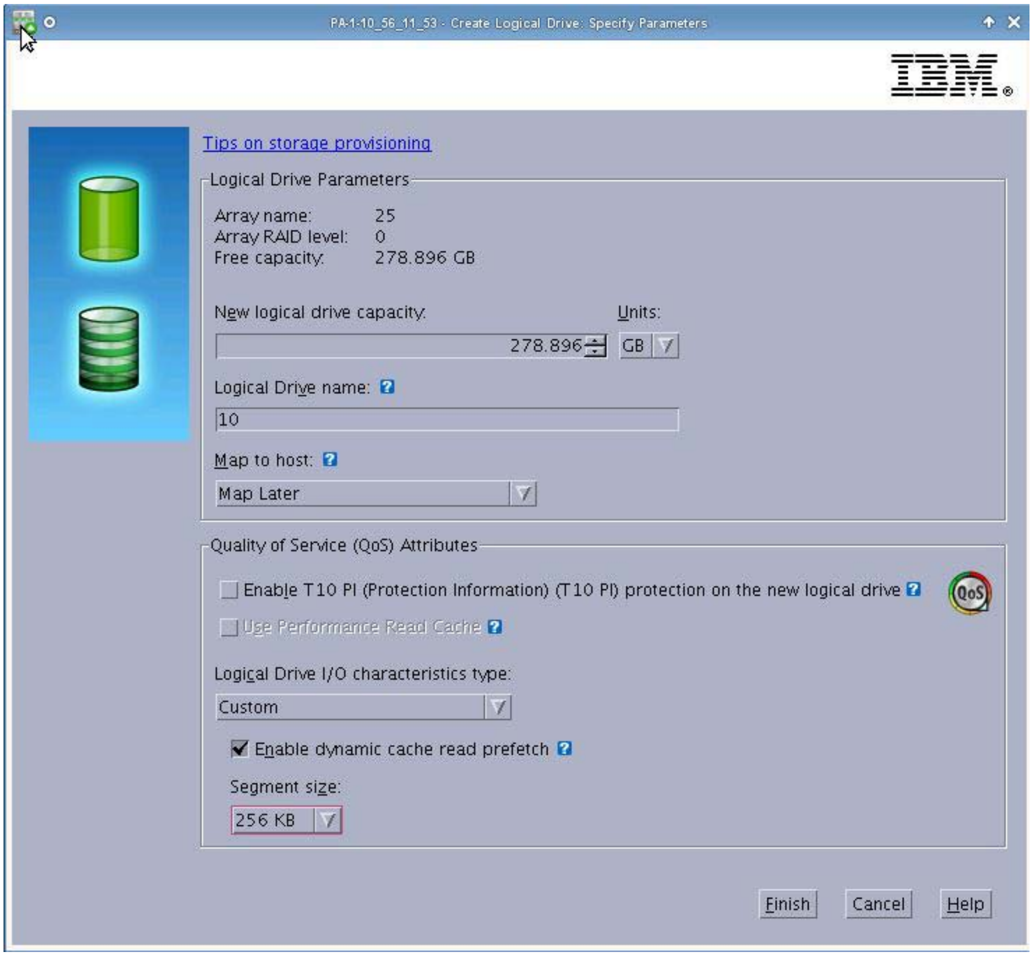

PA-1-10_56_11_53 - Create Logical Drive: Specify Parameters
IBM
Tips on storage provisioning

Logical Drive Parameters

Array name: 25
Array RAID level: 0
Free capacity: 278.896 GB

New logical drive capacity: 278.896
Units: GB

Logical Drive name: 10

Map to host: Map Later

Quality of Service (QoS) Attributes

□ Enable T10 PI (Protection Information) (T10 PI) protection on the new logical drive

□ Use Performance Read Cache

Logical Drive I/O characteristics type: Custom

☑ Enable dynamic cache read prefetch

Segment size: 256 KB

Step 28 In the Subsystem Management window, from the Storage & Copy Services tab, right-click the newly added logical drive and choose **Change** > **Cache Settings**.

Figure 11-20 Change Cache Settings

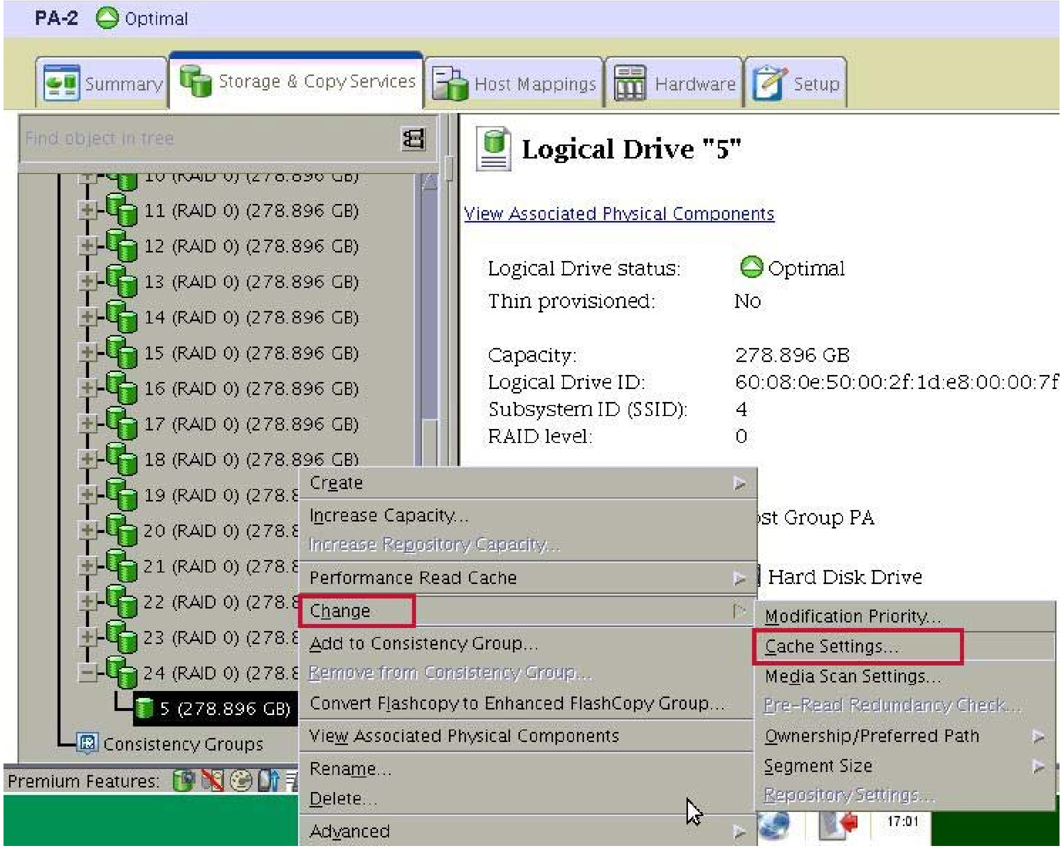

PA-2 Optimal

Logical Drive "5"

View Associated Physical Components

Logical Drive status: Optimal
Thin provisioned: No

Capacity: 278.896 GB
Logical Drive ID: 60:08:0e:50:00:2f:1d:e8:00:00:7f
Subsystem ID (SSID): 4
RAID level: 0

Step 29 In the Change Cache Settings window, make sure the newly added logical drive is selected in the Select Logical Drive section and set the following values:

a. Enable Read Caching check box: **unchecked**

b. Enable Dynamic Cache Read Prefetch check box: **unchecked**

c. Enable Write Caching check box: **checked**

d. Enable Write Caching without Batteries check box: **unchecked**

e. Enable Write Caching with Mirroring check box: **unchecked**

Figure 11-21 Change Cache Settings Values

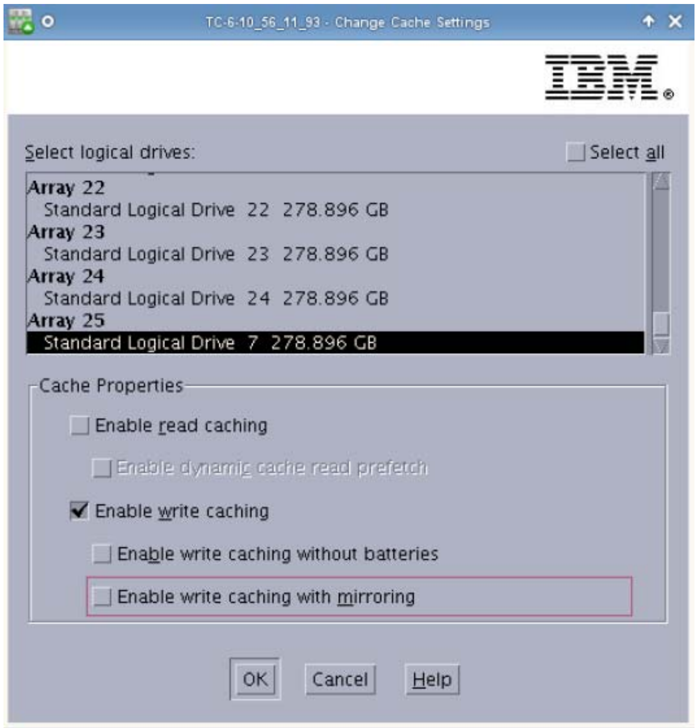

TC-6-10_56_11_93 - Change Cache Settings

Select logical drives: □ Select all

Array 22
  Standard Logical Drive 22 278.896 GB
Array 23
  Standard Logical Drive 23 278.896 GB
Array 24
  Standard Logical Drive 24 278.896 GB
Array 25
  Standard Logical Drive 7 278.896 GB

Cache Properties

□ Enable read caching

□ Enable dynamic cache read prefetch

☑ Enable write caching

□ Enable write caching without batteries

□ Enable write caching with mirroring

Step 30 Click **OK**, click **Yes**, and then click **OK** to complete changing the cache settings.

Step 31 From the Subsystem Management window, click the **Host Mappings** tab.

Step 32 Open the **Undefined Mapping** folder, right-click the unmapped LUN, and choose **Add LUN Mapping**. The unmapped LUN will have the format of "#=LUN?". In the example the unmapped LUN shows as "5 = LUN?".

Figure 11-22 Add LUN Mapping

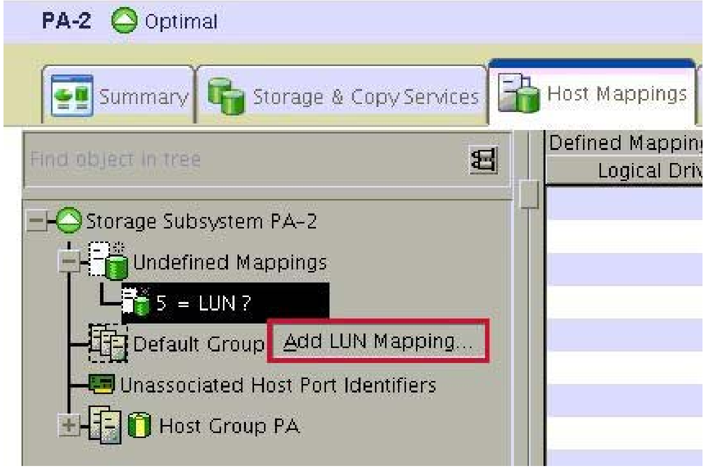

Step 33 From the Add LUN Mapping window, in the Host Group or Host drop-down list, choose **Host Group PA**. In the Logical Unit Number drop-down list, choose a number that has *NEVER* been assigned to any disk before. Use the information that you gathered in Step 9.

Warning **Do not choose the same LUN number that was previously assigned to the disk that you replaced, or any other disk.**

Step 34 Click **Add** and then click **Close**.

Figure 11-23 Add LUN Mapping

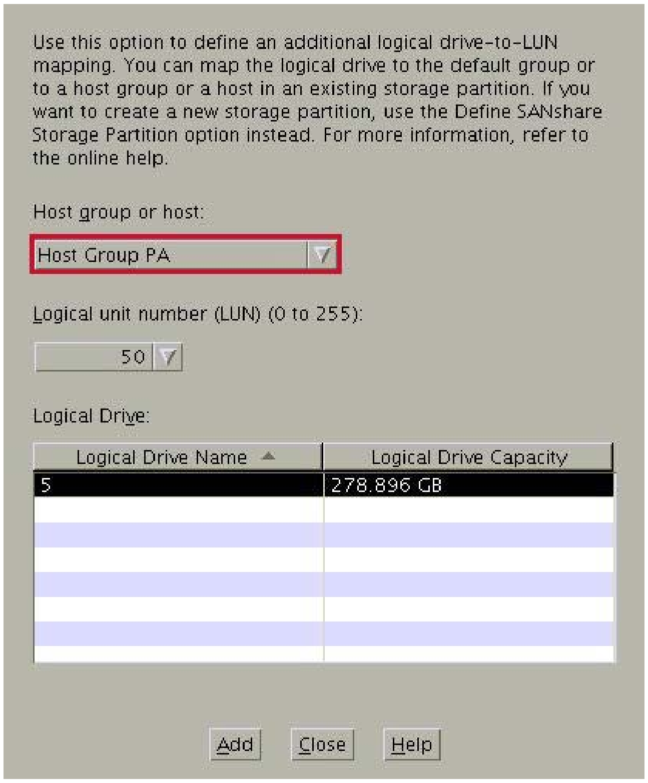

Use this option to define an additional logical drive-to-LUN mapping. You can map the logical drive to the default group or to a host group or a host in an existing storage partition. If you want to create a new storage partition, use the Define SANshare Storage Partition option instead. For more information, refer to the online help.

Host group or host: Host Group PA

Logical unit number (LUN) (0 to 255): 50

Logical Drive:

Logical Drive Name	Logical Drive Capacity
5	278.896 GB

Step 35 From the VDS TC CLI in enable mode, enter the command **cache volume insert**. When prompted to Insert disk, enter **Y**. This process will format the drive.

```
console# cache volume insert Searching for new disks.... Insert disk /dev/sdaw? [y/n] y


disk /dev/sdaw is inserted at index 43
Updating blades...
ce-1 is updated ce-2 is updated ce-3 is updated ce-4 is updated



Procedure complete
```

Step 36 From the Cisco VDS TC Manager, choose **Statistics** > **Storage**. Confirm that the volume that you replaced shows “active”.

Note It may take up to 10 minutes for the volume to appear.

Step 37 Choose **Status** > **Storage** and then click the **Detailed Status** tab to see detailed information about the status of the storage controllers and information about any possible drives that are having problems. When the detailed storage status information appears, confirm that no new yellow boxes appear.

CHAPTER 12

Server Cluster NetApp Storage Enclosure Replacement Procedure

This chapter discusses the steps that are required to replace a server cluster NetApp storage enclosure.

Note If you purchased your NetApp storage enclosure from a NetApp partner and not from Cisco, please contact the NetApp partner for help with the storage enclosure replacement.

Note If you are having problems with a NetApp storage controller, you must open a support case. Replacing a NetApp storage controller is not a customer procedure; it must be performed by an authorized NetApp technician.

If you purchased the NetApp storage enclosure from a NetApp partner and not from Cisco, please contact the NetApp partner to open a case. If you purchased the NetApp storage enclosure from Cisco as part of the VDS TC solution, please contact Cisco TAC to open a case.

Replacing a Storage Enclosure

Note Before performing the steps in this chapter, you must determine whether the storage enclosure needs to be replaced. If you purchased the NetApp storage enclosure from a NetApp partner, contact this partner to help you determine if the storage enclosure needs to be replaced. If you purchased the NetApp enclosure as part of the Cisco VDS TC solution, contact Cisco TAC to have them help you determine if the storage enclosure needs to be replaced.

Follow this procedure to replace a storage enclosure in the VDS TC cluster installation:

Disconnect a Faulty Storage Enclosure

Follow this procedure to disconnect a faulty storage enclosure:

Procedure

Step 1 As soon as the storage enclosure failure is detected, disconnect the iSCSI data cables from the faulty enclosure to restore the system to nominal operation.

Step 2 On the PBR routers, stop redirecting traffic to the VDS TC solution.

Step 3 Follow these steps to stop the VDS TC service:

a. Log into the VDS TC CLI. The CLI prompt console> appears.

b. Enter the command **enable** to enter Enable mode. At the password prompt that appears, enter the enable password. The CLI prompt console# appears.

c. Enter the command **oper service stop** to stop the VDS TC service. When the “Are you sure (y/n)” prompt appears, enter **Y**.

d. Wait several minutes and then enter the command **show status** to confirm that the VDS TC service has stopped. The Operational State should show “disabled”, the Device State should show “stopped”. The following example is a partial output from a VDS TC Cluster installation:

```
console# show status
Cluster state: disabled

Server Slot    Status        Operational state    Device state   Administrative state
ce-1           powered on    disabled             stopped        unlocked
ce-2           powered on    disabled             stopped        unlocked
ce-3           powered on    disabled             stopped        unlocked
ce-4           powered on    disabled             stopped        unlocked
ce-5           powered on    disabled             stopped        unlocked
ce-6           powered on    disabled             stopped        unlocked
ce-7           powered on    disabled             stopped        unlocked
ce-8           powered on    disabled             stopped        unlocked
ce-9           powered on    disabled             stopped        unlocked
```

Note If you do not see an Operational State of “disabled” and a Device State of “stopped” on all of the cache engines, wait several more minutes, and repeat the **show status** command. Do not proceed to Step 4 until you see an Operational State of “disabled” and a Device State of “stopped” on all of the cache engines.

Step 4 Follow these steps to reboot the cache engines. You must perform these steps on all of the cache engines:

a. From an SSH connection to the VDS TC management server, enter the command **ssh root@ce-#**, where # is the number of the cache engine you need to reboot. Enter the password for the root user when prompted.

b. Enter the command **echo b > /proc/sysrq-trigger** to reboot the cache engine.

Step 5 After all of the cache engines have rebooted, perform the following steps on *all* of the cache engines:

a. Enter the **yast** command.

b. From the YaST Control Center, choose **Network Services** in the left pane, choose **iSCSI Initiator** from the right pane, and then press **Enter**.

Note You may need to scroll down in the window to see the iSCSI Initiator option.

Figure 12-1 iSCSI Initiator

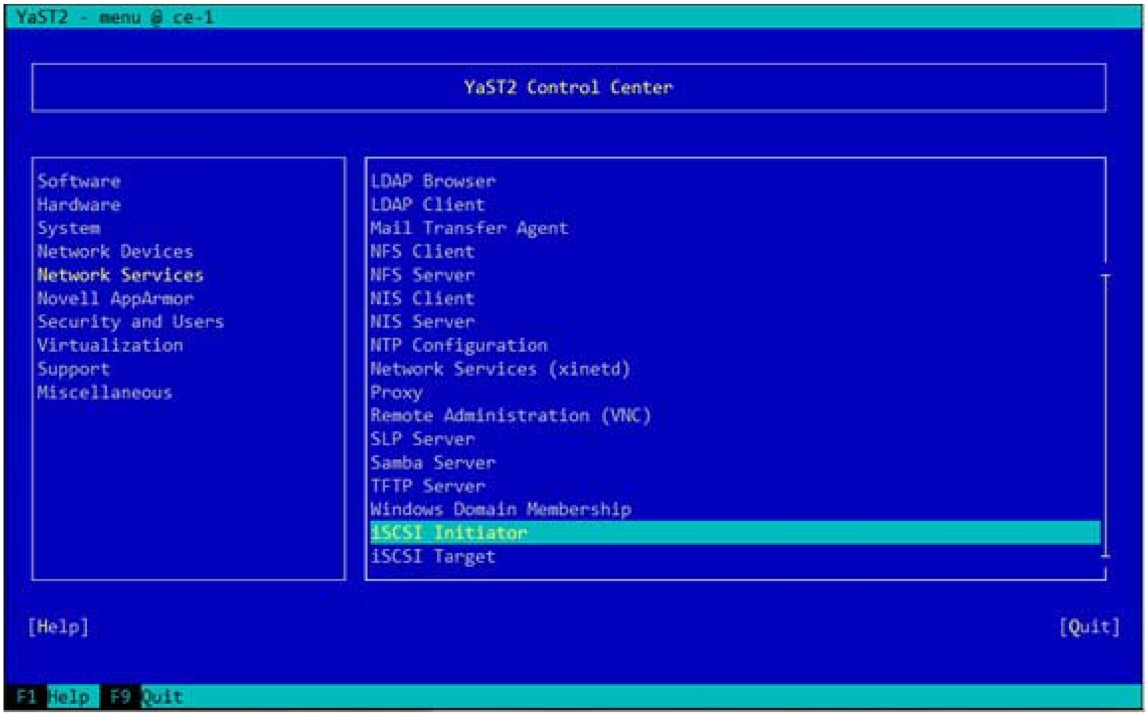

c. From the iSCSI Initiator Overview window, press **Alt-D** to select Discovered Targets.

Figure 12-2 Discovered Targets

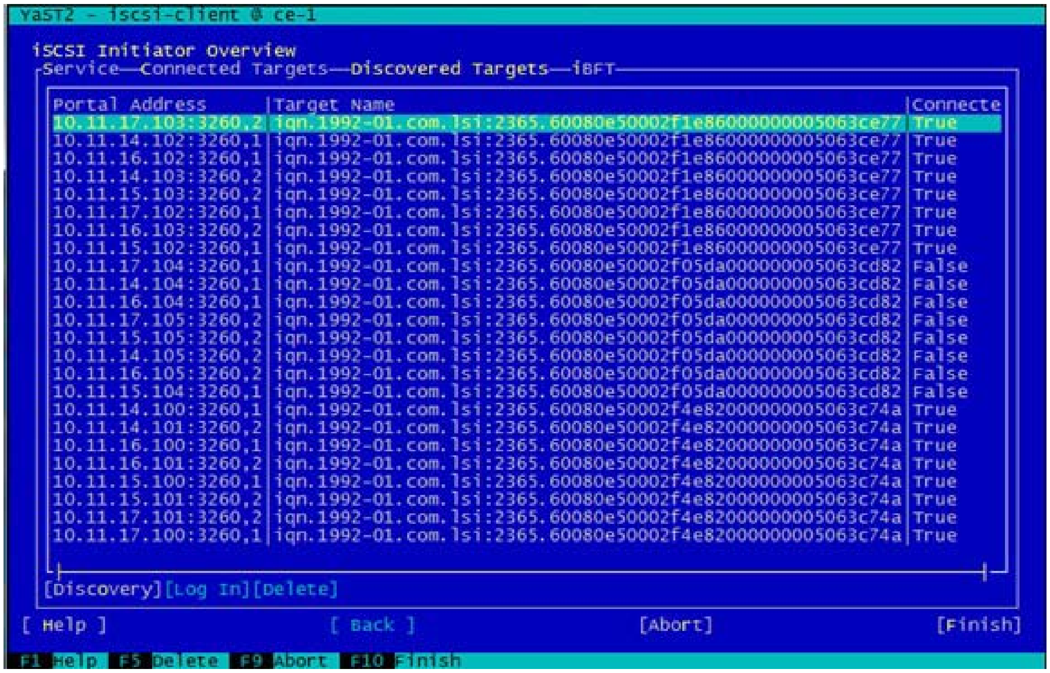

```
YaST2 - iscsi-client @ ce-1

 iSCSI Initiator Overview
 Service—Connected Targets—Discovered Targets—iBFT

  Portal Address     |Target Name                                         |Connecte
  10.11.17.103:3260,2|iqn.1992-01.com.lsi:2365.60080e50002f1e86000000005063ce77|True
  10.11.14.102:3260,1|iqn.1992-01.com.lsi:2365.60080e50002f1e86000000005063ce77|True
  10.11.16.102:3260,1|iqn.1992-01.com.lsi:2365.60080e50002f1e86000000005063ce77|True
  10.11.14.103:3260,2|iqn.1992-01.com.lsi:2365.60080e50002f1e86000000005063ce77|True
  10.11.15.103:3260,2|iqn.1992-01.com.lsi:2365.60080e50002f1e86000000005063ce77|True
  10.11.17.102:3260,1|iqn.1992-01.com.lsi:2365.60080e50002f1e86000000005063ce77|True
  10.11.16.103:3260,2|iqn.1992-01.com.lsi:2365.60080e50002f1e86000000005063ce77|True
  10.11.15.102:3260,1|iqn.1992-01.com.lsi:2365.60080e50002f1e86000000005063ce77|True
  10.11.17.104:3260,1|iqn.1992-01.com.lsi:2365.60080e50002f05da000000005063cd82|False
  10.11.14.104:3260,1|iqn.1992-01.com.lsi:2365.60080e50002f05da000000005063cd82|False
  10.11.16.104:3260,1|iqn.1992-01.com.lsi:2365.60080e50002f05da000000005063cd82|False
  10.11.17.105:3260,2|iqn.1992-01.com.lsi:2365.60080e50002f05da000000005063cd82|False
  10.11.15.105:3260,2|iqn.1992-01.com.lsi:2365.60080e50002f05da000000005063cd82|False
  10.11.14.105:3260,2|iqn.1992-01.com.lsi:2365.60080e50002f05da000000005063cd82|False
  10.11.16.105:3260,2|iqn.1992-01.com.lsi:2365.60080e50002f05da000000005063cd82|False
  10.11.15.104:3260,1|iqn.1992-01.com.lsi:2365.60080e50002f05da000000005063cd82|False
  10.11.14.100:3260,1|iqn.1992-01.com.lsi:2365.60080e50002f4e82000000005063c74a|True
  10.11.14.101:3260,2|iqn.1992-01.com.lsi:2365.60080e50002f4e82000000005063c74a|True
  10.11.16.100:3260,1|iqn.1992-01.com.lsi:2365.60080e50002f4e82000000005063c74a|True
  10.11.16.101:3260,2|iqn.1992-01.com.lsi:2365.60080e50002f4e82000000005063c74a|True
  10.11.15.100:3260,1|iqn.1992-01.com.lsi:2365.60080e50002f4e82000000005063c74a|True
  10.11.15.101:3260,2|iqn.1992-01.com.lsi:2365.60080e50002f4e82000000005063c74a|True
  10.11.17.101:3260,2|iqn.1992-01.com.lsi:2365.60080e50002f4e82000000005063c74a|True
  10.11.17.100:3260,1|iqn.1992-01.com.lsi:2365.60080e50002f4e82000000005063c74a|True

 [Discovery][Log In][Delete]
[ Help ]              [ Back ]              [Abort]              [Finish]
F1 Help  F5 Delete  F9 Abort  F10 Finish
```

d. Choose a target that has a Connected status of “False”. You may need to resize the Discovered Targets window or scroll to the right to see the Connected column. Press **F5** to delete the connected node.

Figure 12-3 Discovered Targets - False Status

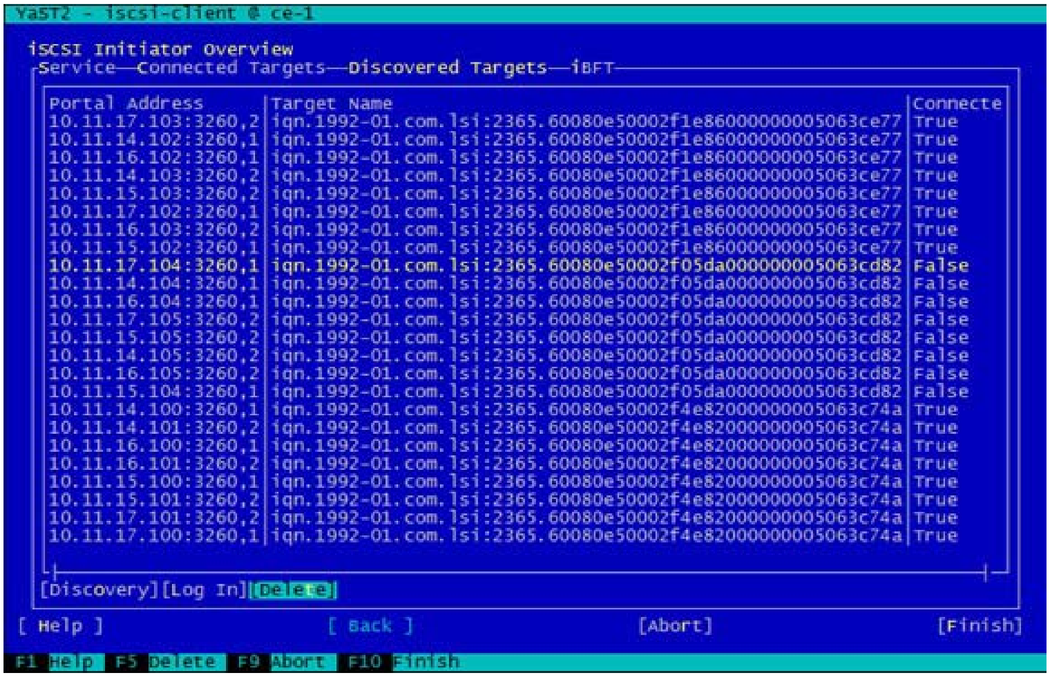

```
YaST2 - iscsi-client @ ce-1

 iSCSI Initiator Overview
 Service─Connected Targets─Discovered Targets─iBFT─
  Portal Address       |Target Name                                         |Connecte
  10.11.17.103:3260,2 |iqn.1992-01.com.lsi:2365.60080e50002f1e8600000000005063ce77|True
  10.11.14.102:3260,1 |iqn.1992-01.com.lsi:2365.60080e50002f1e8600000000005063ce77|True
  10.11.16.102:3260,1 |iqn.1992-01.com.lsi:2365.60080e50002f1e8600000000005063ce77|True
  10.11.14.103:3260,2 |iqn.1992-01.com.lsi:2365.60080e50002f1e8600000000005063ce77|True
  10.11.15.103:3260,2 |iqn.1992-01.com.lsi:2365.60080e50002f1e8600000000005063ce77|True
  10.11.17.102:3260,1 |iqn.1992-01.com.lsi:2365.60080e50002f1e8600000000005063ce77|True
  10.11.16.103:3260,2 |iqn.1992-01.com.lsi:2365.60080e50002f1e8600000000005063ce77|True
  10.11.15.102:3260,1 |iqn.1992-01.com.lsi:2365.60080e50002f1e8600000000005063ce77|True
  10.11.17.104:3260,1 |iqn.1992-01.com.lsi:2365.60080e50002f05da000000005063cd82|False
  10.11.14.104:3260,1 |iqn.1992-01.com.lsi:2365.60080e50002f05da000000005063cd82|False
  10.11.16.104:3260,1 |iqn.1992-01.com.lsi:2365.60080e50002f05da000000005063cd82|False
  10.11.17.105:3260,2 |iqn.1992-01.com.lsi:2365.60080e50002f05da000000005063cd82|False
  10.11.15.105:3260,2 |iqn.1992-01.com.lsi:2365.60080e50002f05da000000005063cd82|False
  10.11.14.105:3260,2 |iqn.1992-01.com.lsi:2365.60080e50002f05da000000005063cd82|False
  10.11.16.105:3260,2 |iqn.1992-01.com.lsi:2365.60080e50002f05da000000005063cd82|False
  10.11.15.104:3260,1 |iqn.1992-01.com.lsi:2365.60080e50002f05da000000005063cd82|False
  10.11.14.100:3260,1 |iqn.1992-01.com.lsi:2365.60080e50002f4e82000000005063c74a|True
  10.11.14.101:3260,2 |iqn.1992-01.com.lsi:2365.60080e50002f4e82000000005063c74a|True
  10.11.16.100:3260,1 |iqn.1992-01.com.lsi:2365.60080e50002f4e82000000005063c74a|True
  10.11.16.101:3260,2 |iqn.1992-01.com.lsi:2365.60080e50002f4e82000000005063c74a|True
  10.11.15.100:3260,1 |iqn.1992-01.com.lsi:2365.60080e50002f4e82000000005063c74a|True
  10.11.15.101:3260,2 |iqn.1992-01.com.lsi:2365.60080e50002f4e82000000005063c74a|True
  10.11.17.101:3260,2 |iqn.1992-01.com.lsi:2365.60080e50002f4e82000000005063c74a|True
  10.11.17.100:3260,1 |iqn.1992-01.com.lsi:2365.60080e50002f4e82000000005063c74a|True

 [Discovery][Log In][Delete]
[ Help ]                  [ Back ]                  [Abort]                  [Finish]
F1 Help  F5 Delete  F9 Abort  F10 Finish
```

e. Repeat Steps d for all of the targets that have a Connected status of “False”. When you are finished, there should be no more targets listed with a Connected status of “False”.

f. Press **F10** to finish the configuration and exit the iSCSI Initiator and then press **F9** to quit the YaST program.

Step 6 Repeat Step 5 for the remaining cache engines.

Removing the Faulty Storage Volumes

After the failed storage iSCSI connections have been removed from all of the cache engines, follow these steps to clean all of the volumes from the cmdb:

Step 1 Log into the VDS TC CLI. The CLI prompt console> appears.

Step 2 Enter the command **enable** to enter Enable mode. At the password prompt that appears, enter the enable password. The CLI prompt console# appears.

Step 3 Enter the command **show volumes** to see the state of the volumes.

```
Licensed volumes : 72
Volume name State Owner
/mnt/vol1 not mounted
/mnt/vol2 not mounted
/mnt/vol3 not mounted
/mnt/vol4 not mounted
/mnt/vol5 not mounted
```

```
/mnt/vol6 not mounted
/mnt/vol7 not mounted
/mnt/vol8 not mounted
/mnt/vol9 not mounted
/mnt/vol10 not mounted
/mnt/vol11 not mounted
/mnt/vol12 not mounted
/mnt/vol13 not mounted
/mnt/vol14 not mounted
/mnt/vol15 not mounted
/mnt/vol16 not mounted
/mnt/vol17 not mounted
/mnt/vol18 not mounted
/mnt/vol19 not mounted
/mnt/vol20 not mounted
/mnt/vol21 not mounted
/mnt/vol22 not mounted
/mnt/vol23 not mounted
/mnt/vol24 not mounted
/mnt/vol25 not mounted
/mnt/vol26 not mounted
/mnt/vol27 not mounted
/mnt/vol28 not mounted
/mnt/vol29 not mounted
/mnt/vol30 not mounted
/mnt/vol31 not mounted
/mnt/vol32 not mounted
/mnt/vol33 not mounted
/mnt/vol34 not mounted
/mnt/vol35 not mounted
/mnt/vol36 not mounted
/mnt/vol37 not mounted
/mnt/vol38 not mounted
/mnt/vol39 not mounted
/mnt/vol40 not mounted
/mnt/vol41 not mounted
/mnt/vol42 not mounted
/mnt/vol43 not mounted
/mnt/vol44 not mounted
/mnt/vol45 not mounted
/mnt/vol46 not mounted
/mnt/vol47 not mounted
/mnt/vol48 not mounted
/mnt/vol49 inactive
/mnt/vol50 inactive
/mnt/vol51 inactive
/mnt/vol52 inactive
/mnt/vol53 inactive
/mnt/vol54 inactive
/mnt/vol55 inactive
/mnt/vol56 inactive
/mnt/vol57 inactive
/mnt/vol58 inactive
/mnt/vol59 inactive
/mnt/vol60 inactive
/mnt/vol61 inactive
/mnt/vol62 inactive
/mnt/vol63 inactive
/mnt/vol64 inactive
/mnt/vol65 inactive
/mnt/vol66 inactive
/mnt/vol67 inactive
/mnt/vol68 inactive
/mnt/vol69 inactive
```

```
/mnt/vol70 inactive
/mnt/vol71 inactive
/mnt/vol72 inactive
```

Step 4 For each volume that has a state of "inactive", enter the command **cache volume remove_content**. When prompted, enter the number of one of the inactive volumes. When the "Are you sure?" message appears, enter **yes**. In this example, volume 49 was identified as "inactive".

```
console# cache volume remove_content
Licensed volumes : 72
Please enter volume number <1-72>
49
Are you sure? This will remove all hashes from volume 49.
[yes|no] no : yes
Removing all the content from volume 49...
Done.
```

Note Repeat Step 4 for any volumes that were marked as "inactive".

Removing the Faulty Storage Subsystem from the SANtricity Storage Manager

Follow these steps to remove the faulty storage subsytem from the SANtricity Storage Manager:

Step 1 From an SSH connection to the VDS TC management server logged in as root, enter the following commands to start the VNC server:

a. **rm /tmp/.X1-lock** (This step is optional. It is only required if there is a leftover lock file from a previous vncserver instance.)

b. **rm /tmp/.X11-unix/X1** (This step is optional. It is only required if there is a leftover lock file from a previous vncserver instance.)

c. **vncserver :1**

Step 2 Use a VNC client to connect from your local computer to the VDS TC management server.

Figure 12-4 VNC Client

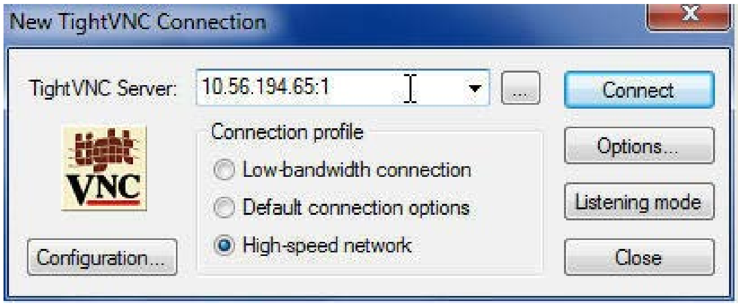

New TightVNC Connection

TightVNC Server: 10.56.194.65:1

Connection profile

○ Low-bandwidth connection

○ Default connection options

● High-speed network

Step 3 From the VNC console window, enter the command **cd /opt/SMgr/client** to change folders.

Step 4 Enter the command **./SMclient** to start the SANtricity ™ ES Storage Manager software.

Step 5 Click the **Devices** tab and in the navigation pane on the left, choose the faulty array.

Step 6 In the Management Connections column, click the Out-of-Band **Details** link.

Figure 12-5 Out-of-Band Details

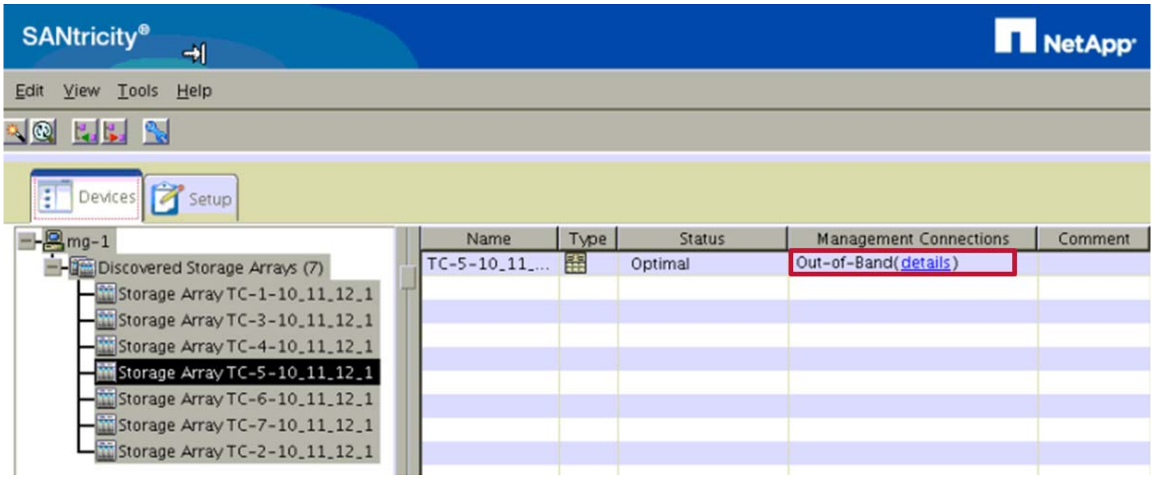

- mg-1
  - Discovered Storage Arrays (7)
    - Storage Array TC-1-10_11_12_1
    - Storage Array TC-3-10_11_12_1
    - Storage Array TC-4-10_11_12_1
    - Storage Array TC-5-10_11_12_1
    - Storage Array TC-6-10_11_12_1
    - Storage Array TC-7-10_11_12_1
    - Storage Array TC-2-10_11_12_1

Name	Type	Status	Management Connections	Comment
TC-5-10_11_...	▦	Optimal	Out-of-Band(details)	

Step 7 The Management Connections window appears. From this window, make note of the IP addresses that are listed. You will need to assign these same IP addresses to the replacement storage enclosure.

Figure 12-6 Management Connections Window

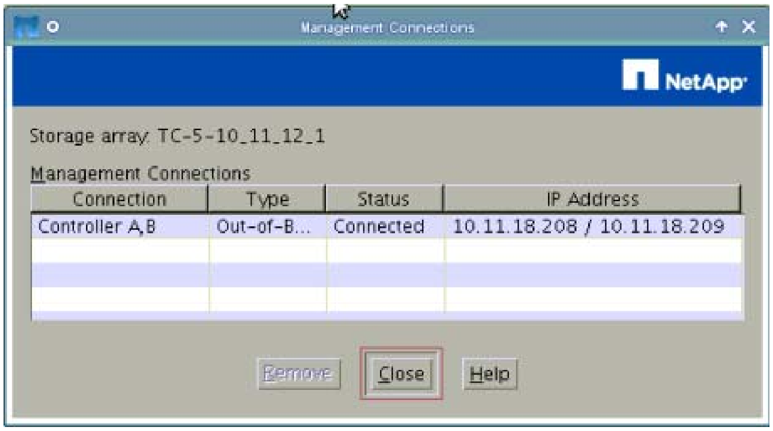

Storage array: TC-5-10_11_12_1

Management Connections

Connection	Type	Status	IP Address
Controller A,B	Out-of-B...	Connected	10.11.18.208 / 10.11.18.209

Step 8 Click **Close** to close the Management Connections window.

Step 9 Right-click the faulty array and choose **Remove** > **Storage Array**.

Figure 12-7 Remove Storage Array

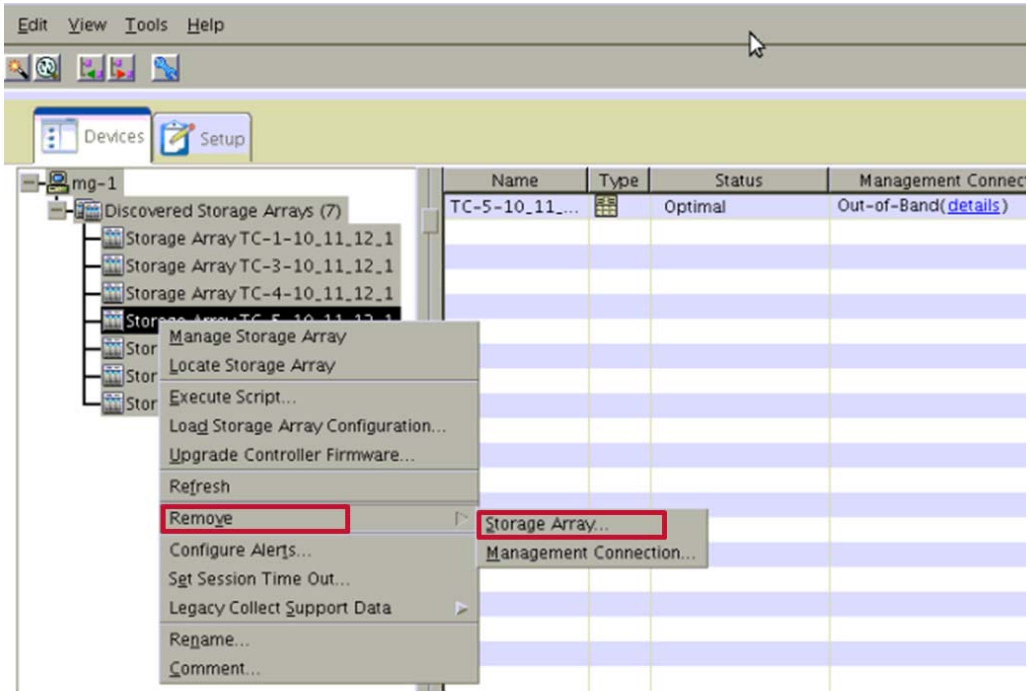
Edit View Tools Help
Devices
Setup
mg-1
Discovered Storage Arrays (7)
Storage Array TC-1-10_11_12_1
Storage Array TC-3-10_11_12_1
Storage Array TC-4-10_11_12_1
Storage Array TC-5-10_11_12_1
Stor
Stor
Stor
Manage Storage Array
Locate Storage Array
Execute Script...
Load Storage Array Configuration...
Upgrade Controller Firmware...
Refresh
Remove
Configure Alerts...
Set Session Time Out...
Legacy Collect Support Data
Rename...
Comment...
Storage Array...
Management Connection...

Name	Type	Status	Management Connec
TC-5-10_11_...		Optimal	Out-of-Band(details)

Step 10 When the Confirm Remove Storage Subsystem dialog box appears, click **Yes**.

Figure 12-8 Confirm Remove Storage Subsystem

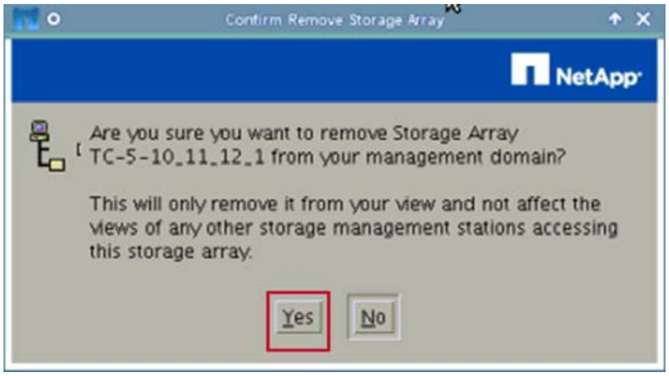

Are you sure you want to remove Storage Array TC-5-10_11_12_1 from your management domain?

This will only remove it from your view and not affect the views of any other storage management stations accessing this storage array.

Connecting the Replacement Storage Enclosure

Follow these steps to connect the replacement storage enclosure to the VDS TC cluster system:

Step 1 Follow the procedures in the storage enclosure installation manual to finish disconnecting and physically unmounting the failed storage enclosure.

Step 2 Follow the procedures in the storage enclosure installation manual to connect the replacement storage enclosure and assign the appropriate IP addresses to the new enclosure. You will need to connect to the default factory assigned IP addresses and change them to the IP addresses that you discovered in Step 7 of the Removing the Faulty Storage Subsystem from the SANtricity Storage Manager section. The following IP addresses are assigned to the different storage enclosures:

- [1]="10.11.18.200 10.11.18.201"
- [2]="10.11.18.202 10.11.18.203"
- [3]="10.11.18.204 10.11.18.205"
- [4]="10.11.18.206 10.11.18.207"
- [5]="10.11.18.208 10.11.18.209"
- [6]="10.11.18.210 10.11.18.211”

Step 3 From the VNC client connection, enter the command **/opt/SMgr/client/SMclient** to run the SANtricity Storage Manager.

Step 4 Click the **Devices** tab and in the navigation pane on the left, right-click the Management Server and choose **Add Storage Array**.

Figure 12-9 Add Storage Array

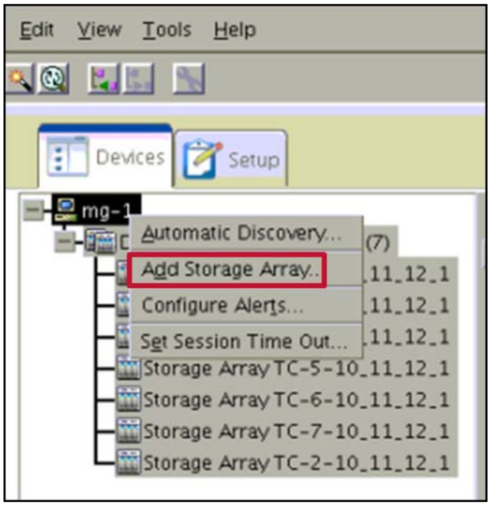

Step 5 From the Add New Storage Array Manual window, in the Controller fields enter the IP addresses that you discovered in Step 7 of the Removing the Faulty Storage Subsystem from the SANtricity Storage Manager section and then click **Add**. These IP addresses should match the IP addresses that you configured on the storage enclosure in Step 2. In the example in this chapter the IP addresses were 10.11.18.204 and 10.11.18.205 for replacing storage enclosure 3.

Figure 12-10 Add New Storage Array

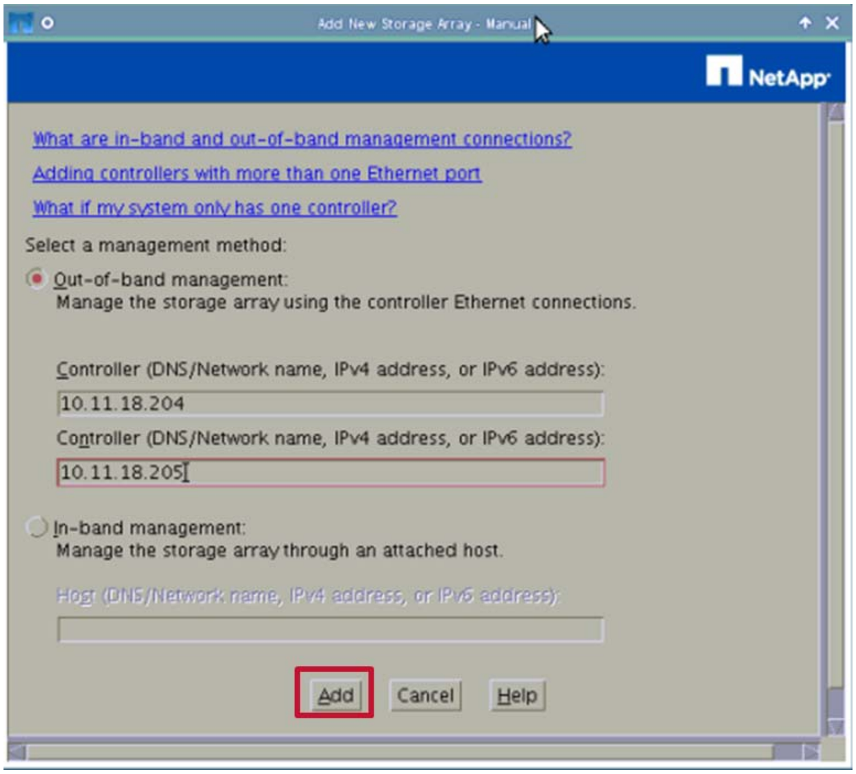

What are in-band and out-of-band management connections?

Adding controllers with more than one Ethernet port

What if my system only has one controller?

Select a management method:

● Out-of-band management:
Manage the storage array using the controller Ethernet connections.

Controller (DNS/Network name, IPv4 address, or IPv6 address): 10.11.18.204

Controller (DNS/Network name, IPv4 address, or IPv6 address): 10.11.18.205

○ In-band management:
Manage the storage array through an attached host.

Host (DNS/Network name, IPv4 address, or IPv6 address):

Step 6 From the Storage Array Added window, click **No**.

Step 7 In the navigation pane on the left, right-click the storage array that you just added and choose **Manage Storage Array**.

Step 8 From the Array Management window, click the **Setup** tab and then click the **Configure iSCSI Host Ports** link.

Figure 12-11 Configure the iSCSI Host Ports

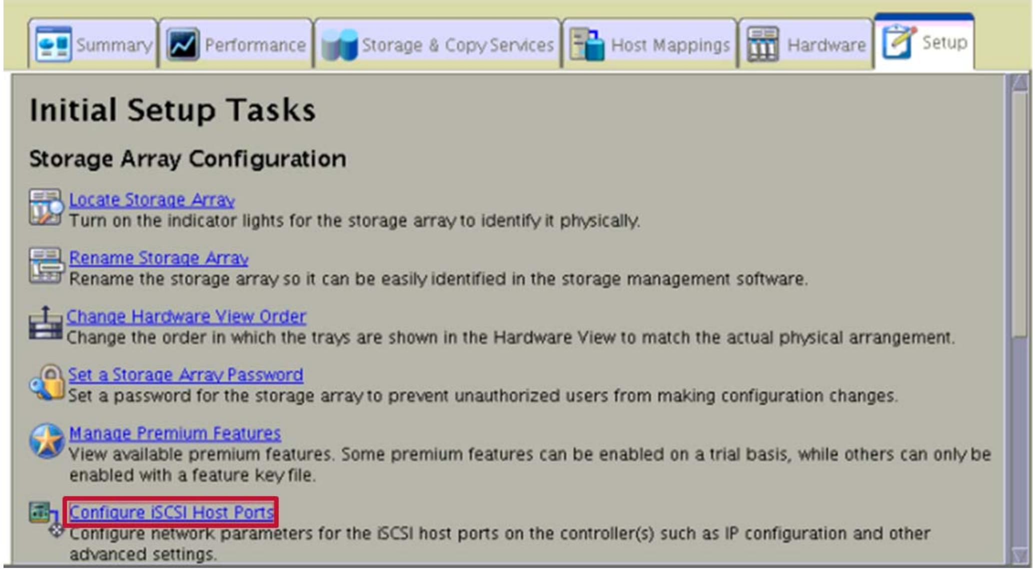

Initial Setup Tasks

Storage Array Configuration

Locate Storage Array
Turn on the indicator lights for the storage array to identify it physically.

Rename Storage Array
Rename the storage array so it can be easily identified in the storage management software.

Change Hardware View Order
Change the order in which the trays are shown in the Hardware View to match the actual physical arrangement.

Set a Storage Array Password
Set a password for the storage array to prevent unauthorized users from making configuration changes.

Manage Premium Features
View available premium features. Some premium features can be enabled on a trial basis, while others can only be enabled with a feature key file.

Configure iSCSI Host Ports
Configure network parameters for the iSCSI host ports on the controller(s) such as IP configuration and other advanced settings.

Step 9 From the Configure iSCSI Ports window, configure the IP addresses on the Controller A and Controller B iSCSI ports. Use the appropriate IP addresses from the table below.

Table 12-1 iSCSI Port Addresses

Enclosure	Port	IP Address
Enclosure 1	Controller A, port 3	10.11.14.100
Enclosure 1	Controller A, port 4	10.11.15.100
Enclosure 1	Controller B, port 3	10.11.14.101
Enclosure 1	Controller B, port 4	10.11.15.101
Enclosure 2	Controller A, port 3	10.11.14.102
Enclosure 2	Controller A, port 4	10.11.15.102
Enclosure 2	Controller B, port 3	10.11.14.103
Enclosure 2	Controller B, port 4	10.11.15.103
Enclosure 3	Controller A, port 3	10.11.14.104
Enclosure 3	Controller A, port 4	10.11.15.104
Enclosure 3	Controller B, port 3	10.11.14.105
Enclosure 3	Controller B, port 4	10.11.15.105
Enclosure 4	Controller A, port 3	10.11.14.106
Enclosure 4	Controller A, port 4	10.11.15.106
Enclosure 4	Controller B, port 3	10.11.14.107
Enclosure 4	Controller B, port 4	10.11.15.107
Enclosure 5	Controller A, port 3	10.11.14.108
Enclosure 5	Controller A, port 4	10.11.15.108

Table 12-1 iSCSI Port Addresses (continued)

Enclosure	Port	IP Address
Enclosure 5	Controller B, port 3	10.11.14.109
Enclosure 5	Controller B, port 4	10.11.15.109
Enclosure 6	Controller A, port 3	10.11.14.110
Enclosure 6	Controller A, port 4	10.11.15.110
Enclosure 6	Controller B, port 3	10.11.14.111
Enclosure 6	Controller B, port 4	10.11.15.111

Step 10 Check the **Enable ICMP Ping Responses** check box and uncheck the *Enable IPv6* check box.

Step 11 Click the **Advanced Port Settings** button. In the Advanced Port Settings window, check the **Enable Jumbo Frames** check box and enter an MTU of **9000**.

Configure the New Storage Enclosure

Follow these steps to configure the new storage enclosure:

Step 1 From an SSH connection to the VDS TC management server logged in as root, enter the command **cd /opt/pang/useful/configure_storages** to change directories.

Step 2 Enter the command **./configure_storages.py -s** <*the # of storage enclosures in the system*> **-b** <*the number of the CEs in the system*> **-n** < *the number of the replaced storage enclosure*>.

For example, for a system with 8 cache engines and 3 storage enclosures, where the third storage enclosure was replaced, enter the following command:

```
./configure_storages_pa.sh -s 3 -b 8 -n 3.
```

Step 3 Follow the on screen instructions. If prompted, enter the password for the root user that was provided by Cisco.

Note You may be prompted to enter the password several times.

Step 4 To perform disk partitioning you must log into CE-1. To log into CE-1, from the SSH connection to the VDS TC management server, enter the command **ssh root@ce-1**.

Step 5 To find the non-partitioned sd devices, enter the command **fdisk -l**. Look for the sd devices that say “doesn't contain a valid partition table”. Make note of these devices. These are the devices that you will need to format.

For example, in the following output sd devices sdax, sday, and sdaz are non-partitioned sd devices because they show “doesn't contain a valid partition table” as part of their output.

```
# fdisk -l
. . .
Disk /dev/sdaw: 299.4 GB, 299462819840 bytes
255 heads, 63 sectors/track, 36407 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes
Disk identifier: 0x00000000

Device Boot Start  End     Blocks  Id  System
```

```
/dev/sdaw1     1   1        8001   83  Linux
/dev/sdaw2     2   2050 16458592+ 83  Linux
/dev/sdaw3  2051   2573  4200997+ 83  Linux
/dev/sdaw4  2574  36407 271771605 83 Linux

Disk /dev/sdax: 299.4 GB, 299462819840 bytes
255 heads, 63 sectors/track, 36407 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes
Disk identifier: 0x00000000

Disk /dev/sdax doesn't contain a valid partition table

Disk /dev/sday: 299.4 GB, 299462819840 bytes
255 heads, 63 sectors/track, 36407 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes
Disk identifier: 0x00000000

Disk /dev/sday doesn't contain a valid partition table

Disk /dev/sdaz: 299.4 GB, 299462819840 bytes
255 heads, 63 sectors/track, 36407 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes
Disk identifier: 0x00000000

Disk /dev/sdaz doesn't contain a valid partition table
```

Step 6 Next you must determine an available volume number. You will use this information to format the drives of the replacement storage array in the next step. To determine available volume numbers, enter the command **/opt/pang/bin/fdisk.sh**. For example, in the following output, there are no gaps in the middle of the volumes. The next available volume number is 49.

```
ce-1:/opt/pang/bin # /opt/pang/bin/fdisk.sh
PeerApp,0001,U /dev/sdz1
PeerApp,0002,U /dev/sdaa1
PeerApp,0003,U /dev/sdab1
PeerApp,0004,U /dev/sdac1
PeerApp,0005,U /dev/sdad1
PeerApp,0006,U /dev/sdae1
PeerApp,0007,U /dev/sdaf1
PeerApp,0008,U /dev/sdag1
PeerApp,0009,U /dev/sdah1
PeerApp,0010,U /dev/sdai1
PeerApp,0011,U /dev/sdaj1
PeerApp,0012,U /dev/sdak1
PeerApp,0013,U /dev/sdal1
PeerApp,0014,U /dev/sdam1
PeerApp,0015,U /dev/sdan1
PeerApp,0016,U /dev/sdao1
PeerApp,0017,U /dev/sdap1
PeerApp,0018,U /dev/sdaq1
PeerApp,0019,U /dev/sdar1
PeerApp,0020,U /dev/sdas1
PeerApp,0021,U /dev/sdat1
PeerApp,0022,U /dev/sdau1
PeerApp,0023,U /dev/sdav1
PeerApp,0024,U /dev/sdaw1
PeerApp,0025,U /dev/sdb1
PeerApp,0026,U /dev/sdc1
PeerApp,0027,U /dev/sdd1
PeerApp,0028,U /dev/sde1
PeerApp,0029,U /dev/sdf1
PeerApp,0030,U /dev/sdg1
```

```
PeerApp,0031,U /dev/sdh1
PeerApp,0032,U /dev/sdi1
PeerApp,0033,U /dev/sdj1
PeerApp,0034,U /dev/sdk1
PeerApp,0035,U /dev/sdl1
PeerApp,0036,U /dev/sdm1
PeerApp,0037,U /dev/sdn1
PeerApp,0038,U /dev/sdo1
PeerApp,0039,U /dev/sdp1
PeerApp,0040,U /dev/sdq1
PeerApp,0041,U /dev/sdr1
PeerApp,0042,U /dev/sds1
PeerApp,0043,U /dev/sdt1
PeerApp,0044,U /dev/sdu1
PeerApp,0045,U /dev/sdv1
PeerApp,0046,U /dev/sdw1
PeerApp,0047,U /dev/sdx1
PeerApp,0048,U /dev/sdy1
```

Note You can use a number available in the middle if there is a gap, or you can use the next available number after the highest volume number.

Step 7 Follow these steps to format all of the non-partitioned sd devices that you discovered in Step 5 one at a time:

a. Enter the command **cd /opt/pang/useful** to change directories.

b. Enter the command **./format_disks.sh -format_one** <*the sd ID*> <*the next available volume number*>, where <*the sd ID*> is the sd id of a non-partitioned sd devices that you discovered in Step 5 and <*the next available volume number*> is an available volume number that you discovered in Step 6.

For example, the following command formats the sd device with the id of sdax and assigns it a volume number of 49:

```
ce-1:/opt/pang/useful # ./format_disks.sh -format_one /dev/sdax 49
Formating: /dev/sdax index for signature 49
disk[49]: /dev/sdax - Erasing boot sector... done
Writing partition table... done
Erasing created partitions... done
Initializing the disk... done
Tuning the filesystem... Marking the disk... done
format_disks.sh: 4.2 finished, see /opt/pang/useful/installog.txt for
details
```

Step 8 Repeat step 7, using the information from Step 5 and Step 6 to format the remainder of the non-partitioned devices and assign them available volume numbers.

Step 9 Reboot all of the cache engines in the grid, except for CE-1, so that they can re-read the formatted disks. Follow these commands to reboot the cache engines:

a. From CE-1, enter the command **exit** to return to the VDS TC management sever.

b. Enter the command **ssh root@ce-***X*, where *X* is the number of a cache engine that you need to reboot.

c. Enter the command **reboot -f** to reboot the cache engine. This will end your SSH session to the cache engine and return you to the VDS TC management server.

d. Repeat Step b and Step c for the remaining cache engines.

Step 10 After you have finished rebooting all of the cache engines, log into each cache engine to confirm that the iSCSI connections were established to the replaced enclosure:

a. Enter the command **ssh root@ce-***X*, where *X* is the number of a cache engine to check.

b. Enter the command **iscsiadm -m session**. If the iSCSI connections have been established, you will see the following results:

```
tcp: [1] 10.11.14.102:3260,1 iqn.1992-08.com.netapp:2752.600a0980005dde250000000053d8deff
tcp: [2] 10.11.14.103:3260,2 iqn.1992-08.com.netapp:2752.600a0980005dde250000000053d8deff
tcp: [3] 10.11.15.103:3260,2 iqn.1992-08.com.netapp:2752.600a0980005dde250000000053d8deff
tcp: [4] 10.11.15.102:3260,1 iqn.1992-08.com.netapp:2752.600a0980005dde250000000053d8deff
tcp: [5] 10.11.14.104:3260,1 iqn.1992-08.com.netapp:2752.600a0980005dde220000000053d8c40e
tcp: [6] 10.11.15.105:3260,2 iqn.1992-08.com.netapp:2752.600a0980005dde220000000053d8c40e
tcp: [7] 10.11.14.105:3260,2 iqn.1992-08.com.netapp:2752.600a0980005dde220000000053d8c40e
tcp: [8] 10.11.15.104:3260,1 iqn.1992-08.com.netapp:2752.600a0980005dde220000000053d8c40e
tcp: [9] 10.11.14.100:3260,1 iqn.1992-08.com.netapp:2752.600a0980005ddc030000000053d8c4b6
tcp: [10] 10.11.14.101:3260,2 iqn.1992-08.com.netapp:2752.600a0980005ddc030000000053d8c4b6
tcp: [11] 10.11.15.100:3260,1 iqn.1992-08.com.netapp:2752.600a0980005ddc030000000053d8c4b6
tcp: [12] 10.11.15.101:3260,2 iqn.1992-08.com.netapp:2752.600a0980005ddc030000000053d8c4b6
```

c. Repeat Steps a and b for the remaining cache engines.

Step 11 Follow these steps to restart the VDS TC service:

a. Log into the VDS TC CLI. The CLI prompt console> appears.

b. Enter the command **enable** to enter Enable mode. At the password prompt that appears, enter the enable password. The CLI prompt console# appears.

c. Enter the command **oper service start** to start the VDS TC service.

d. Wait a few minutes and then enter the command **show status** to confirm that the VDS TC service has started. The Device State should show “started”. For example:

```
console# show status
Cluster state: enabled

Server Slot          Status               Operational state    Device state
Administrative state
ce-1                 powered on           enabled              started
unlocked
ce-2                 powered on           enabled              started
unlocked
ce-3                 powered on           enabled              started
unlocked
ce-4                 powered on           enabled              started
unlocked
ce-5                 powered on           enabled              started
unlocked
ce-6                 powered on           enabled              started
unlocked
ce-7                 powered on           enabled              started
unlocked
ce-8                 powered on           enabled              started
unlocked
ce-9                 powered on           enabled              started
unlocked
ce-10                powered on           enabled              started
unlocked
ce-11                powered on           enabled              started
unlocked
ce-12                powered on           enabled              started
unlocked
ce-13                powered on           enabled              started
unlocked
```

```
ce-14                powered on           enabled              started
unlocked
ce-15                powered on           enabled              started
unlocked
ce-16                powered on           enabled              started
unlocked
```

e. Enter the command **show volumes** to confirm that all of the volumes are mounted. For example:

```
console# show volumes
Licensed volumes : 48
Volume name          State                Owner
/mnt/vol1            mounted              ce-1
/mnt/vol2            mounted              ce-2
/mnt/vol3            mounted              ce-2
/mnt/vol4            mounted              ce-2
/mnt/vol5            mounted              ce-2
/mnt/vol6            mounted              ce-2
/mnt/vol7            mounted              ce-2
/mnt/vol8            mounted              ce-2
/mnt/vol9            mounted              ce-2
/mnt/vol10           mounted              ce-2
/mnt/vol11           mounted              ce-2
/mnt/vol12           mounted              ce-2
/mnt/vol13           mounted              ce-1
/mnt/vol14           mounted              ce-1
/mnt/vol15           mounted              ce-1
/mnt/vol16           mounted              ce-1
/mnt/vol17           mounted              ce-1
/mnt/vol18           mounted              ce-1
/mnt/vol19           mounted              ce-1
/mnt/vol20           mounted              ce-1
/mnt/vol21           mounted              ce-1
/mnt/vol22           mounted              ce-1
/mnt/vol23           mounted              ce-2
/mnt/vol24           mounted              ce-1
```

Step 12 On the PBR routers, start redirecting traffic to the VDS TC solution.

Step 13 After you have begun redirecting traffic to the VDS TC solution, follow these steps to confirm that the volumes on the replaced enclosure are being used:

a. Log into VDS TC Manager and choose **Statistics** > **Storage**.

b. Verify that the status of the volumes show Active and that there is a number other than 0 in the Usage % column.

Figure 12-12 Volume Status and Usage

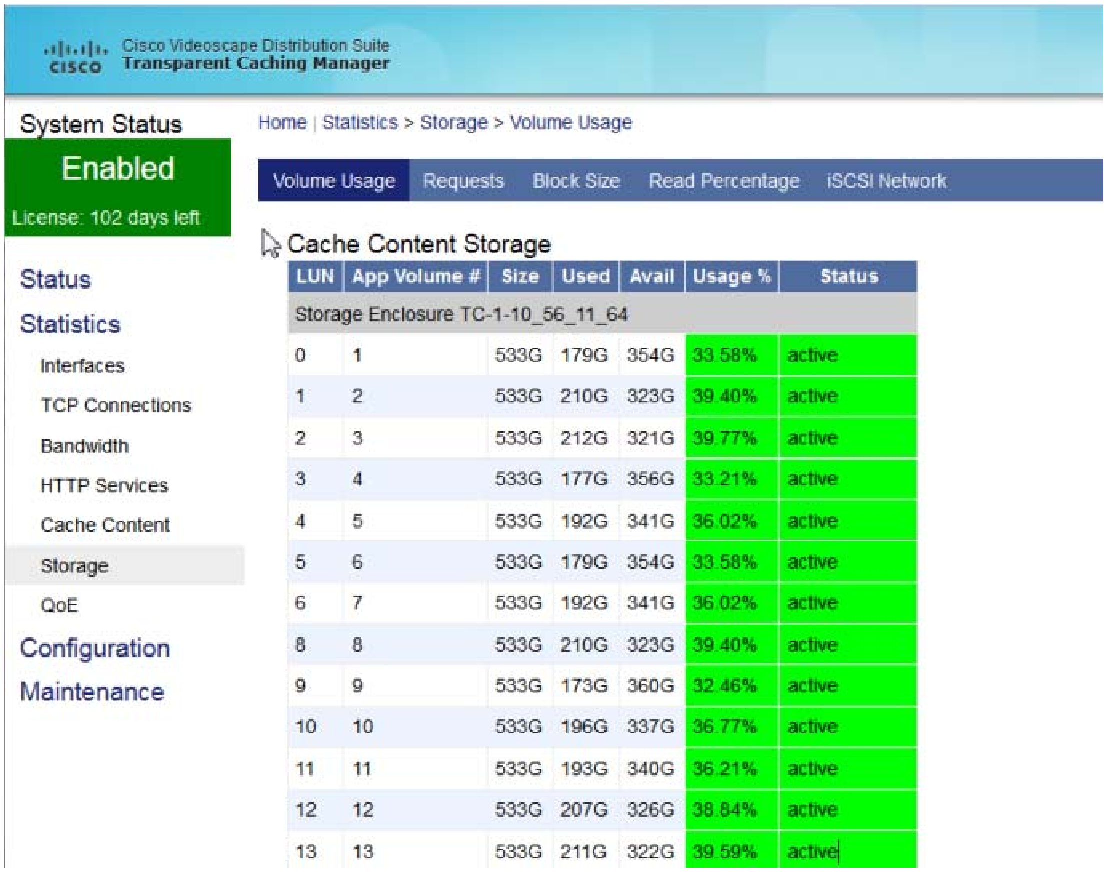

Cache Content Storage

LUN	App Volume #	Size	Used	Avail	Usage %	Status
Storage Enclosure TC-1-10_56_11_64						
0	1	533G	179G	354G	33.58%	active
1	2	533G	210G	323G	39.40%	active
2	3	533G	212G	321G	39.77%	active
3	4	533G	177G	356G	33.21%	active
4	5	533G	192G	341G	36.02%	active
5	6	533G	179G	354G	33.58%	active
6	7	533G	192G	341G	36.02%	active
8	8	533G	210G	323G	39.40%	active
9	9	533G	173G	360G	32.46%	active
10	10	533G	196G	337G	36.77%	active
11	11	533G	193G	340G	36.21%	active
12	12	533G	207G	326G	38.84%	active
13	13	533G	211G	322G	39.59%	active

CHAPTER 13

Server Cluster IBM Storage Enclosure Replacement Procedure

This chapter discusses the steps that are required to replace a server cluster IBM storage enclosure.

Note If you purchased your IBM storage enclosure from an IBM partner and not from Cisco, please contact the IBM partner for help with the storage enclosure replacement.

Note If you are having problems with an IBM storage controller, you must open a support case. Replacing an IBM storage controller is not a customer procedure; it must be performed by an authorized IBM technician.

If you purchased the IBM storage enclosure from an IBM partner and not from Cisco, please contact the IBM partner to open a case. If you purchased the IBM storage enclosure from Cisco as part of the VDS TC solution, please contact Cisco TAC to open a case.

Replacing a Storage Enclosure

Follow this procedure to replace a storage enclosure in the VDS TC cluster installation:

Disconnect a Faulty Storage Enclosure

Follow this procedure to disconnect a faulty storage enclosure:

Procedure

Step 1 As soon as the storage enclosure failure is detected, disconnect the iSCSI data cables from the faulty enclosure to restore the system to nominal operation.

Step 2 On the PBR routers, stop redirecting traffic to the VDS TC solution.

Step 3 Follow these steps to stop the VDS TC service:

a. Log into the VDS TC CLI. The CLI prompt console> appears.

b. Enter the command **enable** to enter Enable mode. At the password prompt that appears, enter the enable password. The CLI prompt console# appears.

c. Enter the command **oper service stop** to stop the VDS TC service. When the “Are you sure (y/n)” prompt appears, enter **Y**.

d. Wait several minutes and then enter the command **show status** to confirm that the VDS TC service has stopped. The Operational State should show “disabled”, the Device State should show “stopped”. The following example is a partial output from a VDS TC Cluster installation:

```
console# show status
Cluster state: disabled

Server Slot    Status        Operational state    Device state   Administrative state
ce-1           powered on    disabled             stopped        unlocked
ce-2           powered on    disabled             stopped        unlocked
ce-3           powered on    disabled             stopped        unlocked
ce-4           powered on    disabled             stopped        unlocked
ce-5           powered on    disabled             stopped        unlocked
ce-6           powered on    disabled             stopped        unlocked
ce-7           powered on    disabled             stopped        unlocked
ce-8           powered on    disabled             stopped        unlocked
ce-9           powered on    disabled             stopped        unlocked
```

Note If you do not see an Operational State of “disabled” and a Device State of “stopped” on all of the cache engines, wait several more minutes, and repeat the **show status** command. Do not proceed to Step 4 until you see an Operational State of “disabled” and a Device State of “stopped” on all of the cache engines.

Step 4 Follow these steps to reboot the cache engines. You must perform these steps on all of the cache engines:

a. From an SSH connection to the VDS TC management server, enter the command **ssh root@ce-***#*, where # is the number of the cache engine you need to reboot. Enter the password for the root user when prompted.

b. Enter the command **echo b** > **/proc/sysrq-trigger** to reboot the cache engine.

Step 5 After all of the cache engines have rebooted, perform the following steps on *all* of the cache engines:

a. Enter the **yast** command.

b. From the YaST Control Center, choose **Network Services** in the left pane, choose **iSCSI Initiator** from the right pane, and then press **Enter**.

Note You may need to scroll down in the window to see the iSCSI Initiator option.

Figure 13-1 iSCSI Initiator

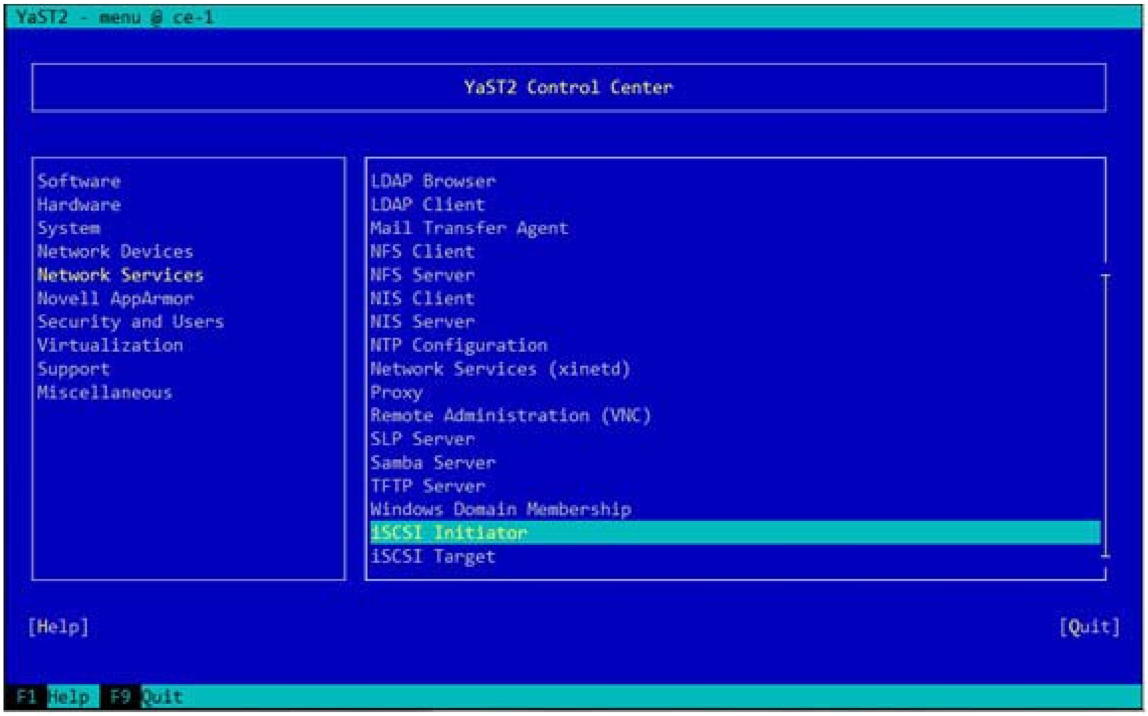

```
YaST2 - menu @ ce-1

                         YaST2 Control Center

Software                   LDAP Browser
Hardware                   LDAP Client
System                     Mail Transfer Agent
Network Devices            NFS Client
Network Services           NFS Server
Novell AppArmor            NIS Client
Security and Users         NIS Server
Virtualization             NTP Configuration
Support                    Network Services (xinetd)
Miscellaneous              Proxy
                           Remote Administration (VNC)
                           SLP Server
                           Samba Server
                           TFTP Server
                           Windows Domain Membership
                           iSCSI Initiator
                           iSCSI Target

[Help]                                                    [Quit]

F1 Help  F9 Quit
```

c. From the iSCSI Initiator Overview window, press **Alt-D** to select Discovered Targets.

Figure 13-2 Discovered Targets

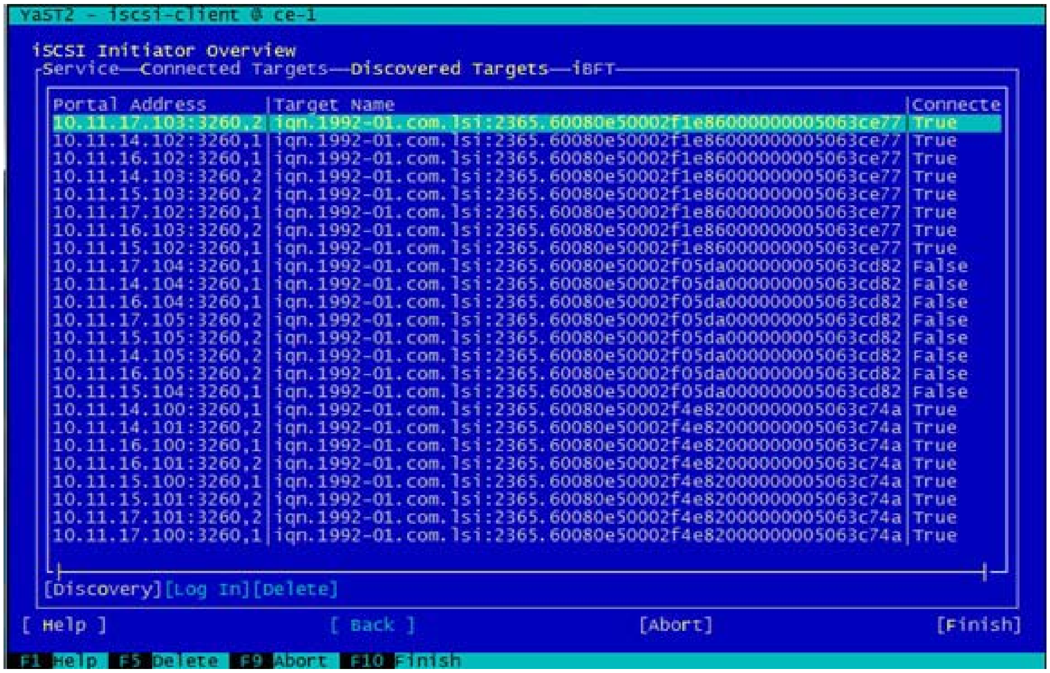

```
YaST2 - iscsi-client @ ce-1

 iSCSI Initiator Overview
 Service—Connected Targets—Discovered Targets—iBFT
  Portal Address       |Target Name                                        |Connecte
  10.11.17.103:3260,2  |iqn.1992-01.com.lsi:2365.60080e50002f1e86000000005063ce77|True
  10.11.14.102:3260,1  |iqn.1992-01.com.lsi:2365.60080e50002f1e86000000005063ce77|True
  10.11.16.102:3260,1  |iqn.1992-01.com.lsi:2365.60080e50002f1e86000000005063ce77|True
  10.11.14.103:3260,2  |iqn.1992-01.com.lsi:2365.60080e50002f1e86000000005063ce77|True
  10.11.15.103:3260,2  |iqn.1992-01.com.lsi:2365.60080e50002f1e86000000005063ce77|True
  10.11.17.102:3260,1  |iqn.1992-01.com.lsi:2365.60080e50002f1e86000000005063ce77|True
  10.11.16.103:3260,2  |iqn.1992-01.com.lsi:2365.60080e50002f1e86000000005063ce77|True
  10.11.15.102:3260,1  |iqn.1992-01.com.lsi:2365.60080e50002f1e86000000005063ce77|True
  10.11.17.104:3260,1  |iqn.1992-01.com.lsi:2365.60080e50002f05da000000005063cd82|False
  10.11.14.104:3260,1  |iqn.1992-01.com.lsi:2365.60080e50002f05da000000005063cd82|False
  10.11.16.104:3260,1  |iqn.1992-01.com.lsi:2365.60080e50002f05da000000005063cd82|False
  10.11.17.105:3260,2  |iqn.1992-01.com.lsi:2365.60080e50002f05da000000005063cd82|False
  10.11.15.105:3260,2  |iqn.1992-01.com.lsi:2365.60080e50002f05da000000005063cd82|False
  10.11.14.105:3260,2  |iqn.1992-01.com.lsi:2365.60080e50002f05da000000005063cd82|False
  10.11.16.105:3260,2  |iqn.1992-01.com.lsi:2365.60080e50002f05da000000005063cd82|False
  10.11.15.104:3260,1  |iqn.1992-01.com.lsi:2365.60080e50002f05da000000005063cd82|False
  10.11.14.100:3260,1  |iqn.1992-01.com.lsi:2365.60080e50002f4e82000000005063c74a|True
  10.11.14.101:3260,2  |iqn.1992-01.com.lsi:2365.60080e50002f4e82000000005063c74a|True
  10.11.16.100:3260,1  |iqn.1992-01.com.lsi:2365.60080e50002f4e82000000005063c74a|True
  10.11.16.101:3260,2  |iqn.1992-01.com.lsi:2365.60080e50002f4e82000000005063c74a|True
  10.11.15.100:3260,1  |iqn.1992-01.com.lsi:2365.60080e50002f4e82000000005063c74a|True
  10.11.15.101:3260,2  |iqn.1992-01.com.lsi:2365.60080e50002f4e82000000005063c74a|True
  10.11.17.101:3260,2  |iqn.1992-01.com.lsi:2365.60080e50002f4e82000000005063c74a|True
  10.11.17.100:3260,1  |iqn.1992-01.com.lsi:2365.60080e50002f4e82000000005063c74a|True

 [Discovery][Log In][Delete]
[ Help ]              [ Back ]              [Abort]              [Finish]
F1 Help  F5 Delete  F9 Abort  F10 Finish
```

d. Choose a target that has a Connected status of “False”. You may need to resize the Discovered Targets window or scroll to the right to see the Connected column. Press **F5** to delete the connected node.

Figure 13-3 Discovered Targets - False Status

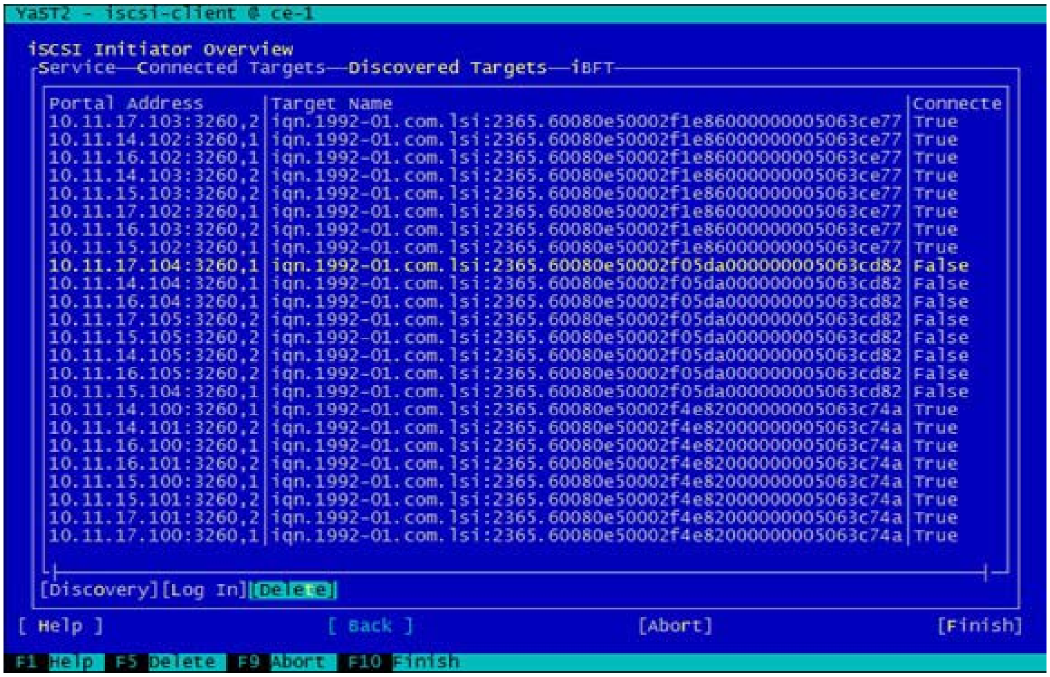

```
YaST2 - iscsi-client @ ce-1

 iSCSI Initiator Overview
 Service─Connected Targets─Discovered Targets─iBFT─

  Portal Address       |Target Name                                        |Connecte
  10.11.17.103:3260,2  |iqn.1992-01.com.lsi:2365.60080e50002f1e86000000005063ce77|True
  10.11.14.102:3260,1  |iqn.1992-01.com.lsi:2365.60080e50002f1e86000000005063ce77|True
  10.11.16.102:3260,1  |iqn.1992-01.com.lsi:2365.60080e50002f1e86000000005063ce77|True
  10.11.14.103:3260,2  |iqn.1992-01.com.lsi:2365.60080e50002f1e86000000005063ce77|True
  10.11.15.103:3260,2  |iqn.1992-01.com.lsi:2365.60080e50002f1e86000000005063ce77|True
  10.11.17.102:3260,1  |iqn.1992-01.com.lsi:2365.60080e50002f1e86000000005063ce77|True
  10.11.16.103:3260,2  |iqn.1992-01.com.lsi:2365.60080e50002f1e86000000005063ce77|True
  10.11.15.102:3260,1  |iqn.1992-01.com.lsi:2365.60080e50002f1e86000000005063ce77|True
  10.11.17.104:3260,1  |iqn.1992-01.com.lsi:2365.60080e50002f05da000000005063cd82|False
  10.11.14.104:3260,1  |iqn.1992-01.com.lsi:2365.60080e50002f05da000000005063cd82|False
  10.11.16.104:3260,1  |iqn.1992-01.com.lsi:2365.60080e50002f05da000000005063cd82|False
  10.11.17.105:3260,2  |iqn.1992-01.com.lsi:2365.60080e50002f05da000000005063cd82|False
  10.11.15.105:3260,2  |iqn.1992-01.com.lsi:2365.60080e50002f05da000000005063cd82|False
  10.11.14.105:3260,2  |iqn.1992-01.com.lsi:2365.60080e50002f05da000000005063cd82|False
  10.11.16.105:3260,2  |iqn.1992-01.com.lsi:2365.60080e50002f05da000000005063cd82|False
  10.11.15.104:3260,1  |iqn.1992-01.com.lsi:2365.60080e50002f05da000000005063cd82|False
  10.11.14.100:3260,1  |iqn.1992-01.com.lsi:2365.60080e50002f4e82000000005063c74a|True
  10.11.14.101:3260,2  |iqn.1992-01.com.lsi:2365.60080e50002f4e82000000005063c74a|True
  10.11.16.100:3260,1  |iqn.1992-01.com.lsi:2365.60080e50002f4e82000000005063c74a|True
  10.11.16.101:3260,2  |iqn.1992-01.com.lsi:2365.60080e50002f4e82000000005063c74a|True
  10.11.15.100:3260,1  |iqn.1992-01.com.lsi:2365.60080e50002f4e82000000005063c74a|True
  10.11.15.101:3260,2  |iqn.1992-01.com.lsi:2365.60080e50002f4e82000000005063c74a|True
  10.11.17.101:3260,2  |iqn.1992-01.com.lsi:2365.60080e50002f4e82000000005063c74a|True
  10.11.17.100:3260,1  |iqn.1992-01.com.lsi:2365.60080e50002f4e82000000005063c74a|True

 [Discovery][Log In][Delete]

[ Help ]                  [ Back ]                 [Abort]                  [Finish]
F1 Help  F5 Delete  F9 Abort  F10 Finish
```

e. Repeat Steps d for all of the targets that have a Connected status of “False”. When you are finished, there should be no more targets listed with a Connected status of “False”.

f. Press **F10** to finish the configuration and exit the iSCSI Initiator and then press **F9** to quit the YaST program.

Step 6 Repeat Step 5 for the remaining cache engines.

Removing the Faulty Storage Volumes

After the failed storage iSCSI connections have been removed from all of the cache engines, follow these steps to clean all of the volumes from the cmdb:

Step 1 Log into the VDS TC CLI. The CLI prompt console> appears.

Step 2 Enter the command **enable** to enter Enable mode. At the password prompt that appears, enter the enable password. The CLI prompt console# appears.

Step 3 Enter the command **show volumes** to see the state of the volumes.

```
Licensed volumes : 72
Volume name State Owner
/mnt/vol1 not mounted
/mnt/vol2 not mounted
/mnt/vol3 not mounted
/mnt/vol4 not mounted
/mnt/vol5 not mounted
```

```
/mnt/vol6 not mounted
/mnt/vol7 not mounted
/mnt/vol8 not mounted
/mnt/vol9 not mounted
/mnt/vol10 not mounted
/mnt/vol11 not mounted
/mnt/vol12 not mounted
/mnt/vol13 not mounted
/mnt/vol14 not mounted
/mnt/vol15 not mounted
/mnt/vol16 not mounted
/mnt/vol17 not mounted
/mnt/vol18 not mounted
/mnt/vol19 not mounted
/mnt/vol20 not mounted
/mnt/vol21 not mounted
/mnt/vol22 not mounted
/mnt/vol23 not mounted
/mnt/vol24 not mounted
/mnt/vol25 not mounted
/mnt/vol26 not mounted
/mnt/vol27 not mounted
/mnt/vol28 not mounted
/mnt/vol29 not mounted
/mnt/vol30 not mounted
/mnt/vol31 not mounted
/mnt/vol32 not mounted
/mnt/vol33 not mounted
/mnt/vol34 not mounted
/mnt/vol35 not mounted
/mnt/vol36 not mounted
/mnt/vol37 not mounted
/mnt/vol38 not mounted
/mnt/vol39 not mounted
/mnt/vol40 not mounted
/mnt/vol41 not mounted
/mnt/vol42 not mounted
/mnt/vol43 not mounted
/mnt/vol44 not mounted
/mnt/vol45 not mounted
/mnt/vol46 not mounted
/mnt/vol47 not mounted
/mnt/vol48 not mounted
/mnt/vol49 inactive
/mnt/vol50 inactive
/mnt/vol51 inactive
/mnt/vol52 inactive
/mnt/vol53 inactive
/mnt/vol54 inactive
/mnt/vol55 inactive
/mnt/vol56 inactive
/mnt/vol57 inactive
/mnt/vol58 inactive
/mnt/vol59 inactive
/mnt/vol60 inactive
/mnt/vol61 inactive
/mnt/vol62 inactive
/mnt/vol63 inactive
/mnt/vol64 inactive
/mnt/vol65 inactive
/mnt/vol66 inactive
/mnt/vol67 inactive
/mnt/vol68 inactive
/mnt/vol69 inactive
```

```
/mnt/vol70 inactive
/mnt/vol71 inactive
/mnt/vol72 inactive
```

Step 4 For each volume that has a state of “inactive”, enter the command **cache volume remove_content**. When prompted, enter the number of one of the inactive volumes. When the “Are you sure?” message appears, enter **yes**. In this example, volume 49 was identified as “inactive”.

```
console# cache volume remove_content
Licensed volumes : 72
Please enter volume number <1-72>
49
Are you sure? This will remove all hashes from volume 49.
[yes|no] no : yes
Removing all the content from volume 49...
Done.
```

Note Repeat Step 4 for any volumes that were marked as “inactive”.

Removing the Faulty Storage Subsystem from the Storage Manager

Follow these steps to remove the faulty storage subsytem from the DS Storage Manager:

Step 1 From an SSH connection to the VDS TC management server logged in as root, enter the following commands to start the VNC server:

a. **rm /tmp/.X1-lock** (This step is optional. It is only required if there is a leftover lock file from a previous vncserver instance.)

b. **rm /tmp/.X11-unix/X1** (This step is optional. It is only required if there is a leftover lock file from a previous vncserver instance.)

c. **vncserver :1**

Step 2 Use a VNC client to connect from your local computer to the VDS TC management server.

Figure 13-4 VNC Client

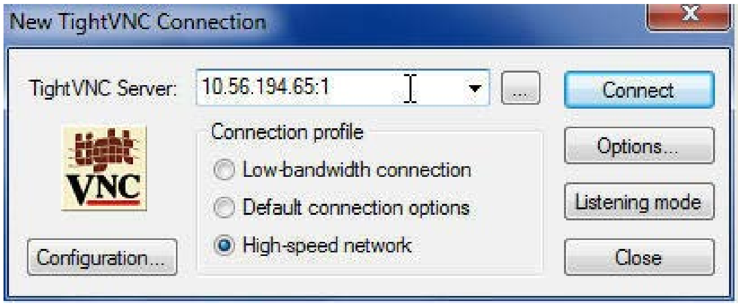

New TightVNC Connection

TightVNC Server: 10.56.194.65:1

Connection profile
○ Low-bandwidth connection
○ Default connection options
● High-speed network

Step 3 From the VNC console window, enter the command **/opt/IBM_DS/client/SMclient** to run the DS Storage Manager.

Step 4 Click the **Devices** tab and in the navigation pane on the left, choose the faulty array. A faulty array will have an icon with a red lightening bolt and a status of “Unresponsive”.

Step 5 In the Management Connections column, click the Out-of-Band **Details** link.

Figure 13-5 Out-of-Band Details

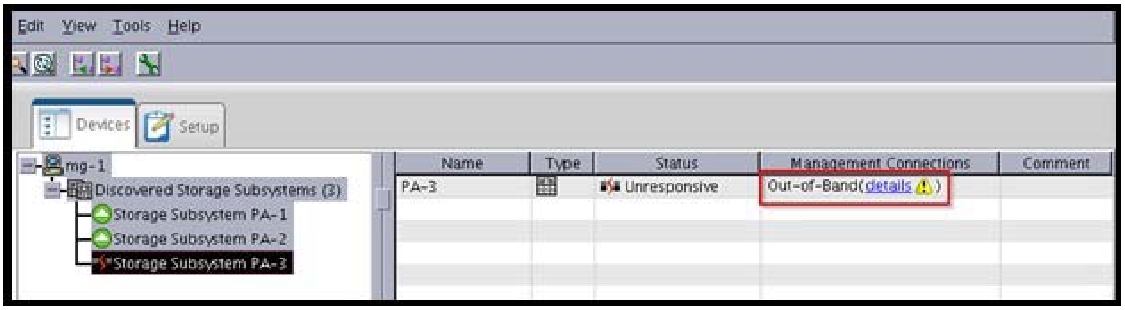

Name	Type	Status	Management Connections	Comment
PA-3		Unresponsive	Out-of-Band(details ⚠)	

Step 6 The Management Connections window appears. From this window, make note of the IP addresses that are listed. You will need to assign these same IP addresses to the replacement storage enclosure.

Figure 13-6 Management Connections Window

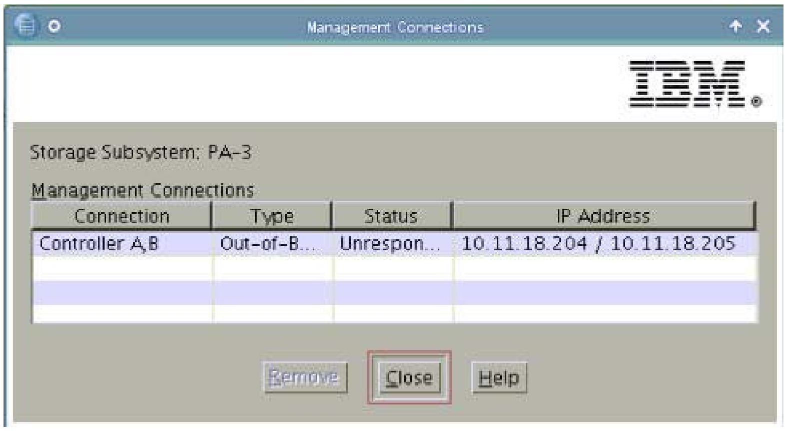

Storage Subsystem: PA-3

Connection	Type	Status	IP Address
Controller A,B	Out-of-B...	Unrespon...	10.11.18.204 / 10.11.18.205

Step 7 Click **Close** to close the Management Connections window.

Step 8 Right-click the faulty array and choose **Remove** > **Storage Subsystem**.

Figure 13-7 Remove Storage Subsystem

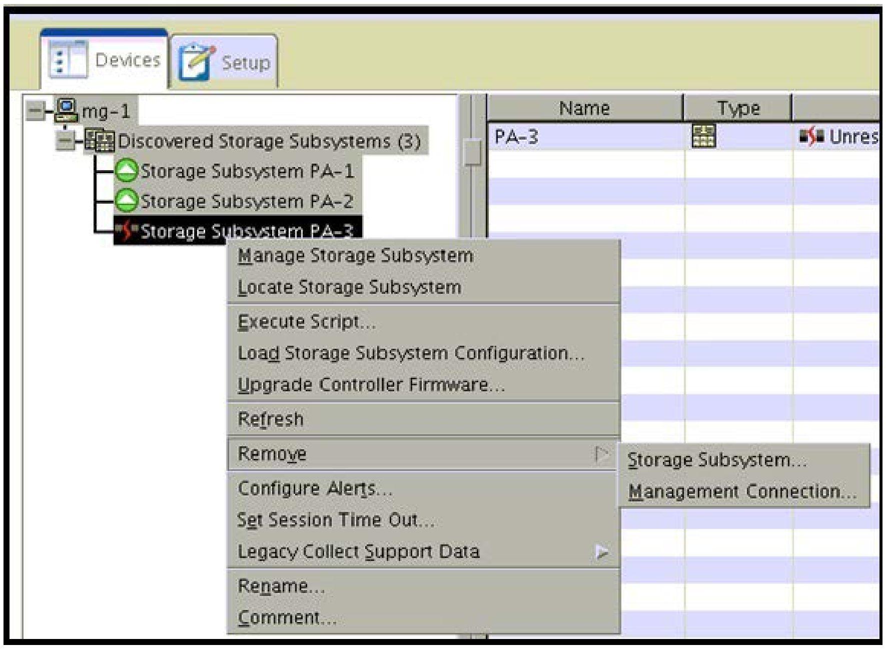

Step 9 When the Confirm Remove Storage Subsystem dialog box appears, click **Yes**.

Figure 13-8 Confirm Remove Storage Subsystem

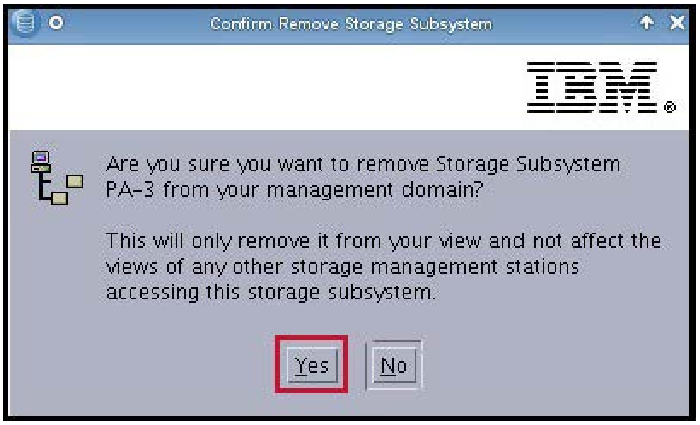

Are you sure you want to remove Storage Subsystem PA-3 from your management domain?

This will only remove it from your view and not affect the views of any other storage management stations accessing this storage subsystem.

Connecting the Replacement Storage Enclosure

Follow these steps to connect the replacement storage enclosure to the VDS TC cluster system:

Step 1 Follow the procedures in the storage enclosure installation manual to finish disconnecting and physically unmounting the failed storage enclosure.

Step 2 Follow the procedures in the storage enclosure installation manual to connect the replacement storage enclosure and assign the appropriate IP addresses to the new enclosure. You will need to connect to the default factory assigned IP addresses and change them to the IP addresses that you discovered in Step 6 of the Removing the Faulty Storage Subsystem from the Storage Manager section. The following IP addresses are assigned to the different storage enclosures:

- [1]="10.11.18.200 10.11.18.201"
- [2]="10.11.18.202 10.11.18.203"
- [3]="10.11.18.204 10.11.18.205"
- [4]="10.11.18.206 10.11.18.207"
- [5]="10.11.18.208 10.11.18.209"

Step 3 From the VNC client connection, enter the command **/opt/IBM_DS/client/SMclient** to run the DS Storage Manager.

Step 4 Click the **Devices** tab and in the navigation pane on the left, right-click the Management Server and choose **Add Storage Subsystem**.

Figure 13-9 Add Storage Subsystem

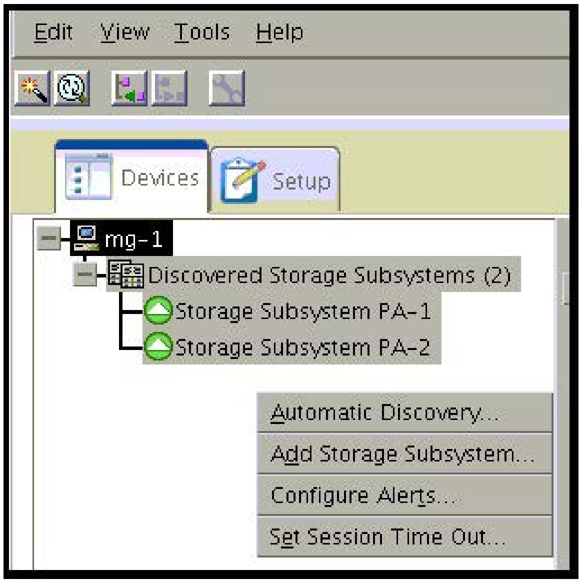

Step 5 From the Add New Storage Subsystem Manual window, in the Controller fields enter the IP addresses that you discovered in Step 6 of the Removing the Faulty Storage Subsystem from the Storage Manager section and then click **Add**. These IP addresses should match the IP addresses that you configured on the storage enclosure in Step 2. In the example in this chapter the IP addresses were 10.11.18.204 and 10.11.18.205 for replacing storage enclosure 3.

Figure 13-10 Add New Storage Subsystem

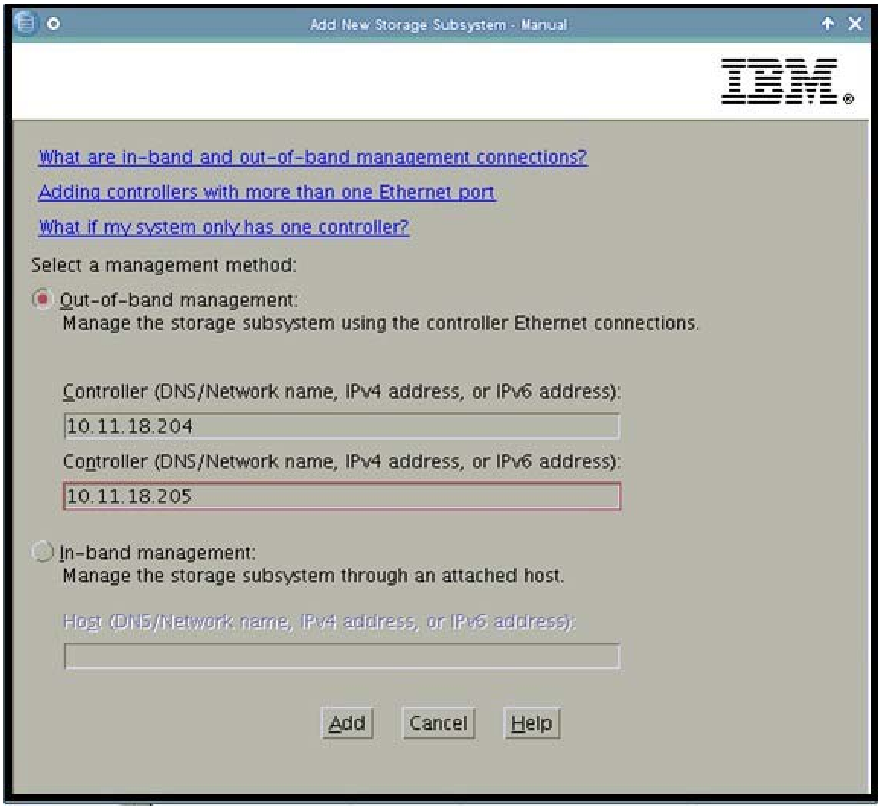

Add New Storage Subsystem - Manual

What are in-band and out-of-band management connections?

Adding controllers with more than one Ethernet port

What if my system only has one controller?

Select a management method:

● Out-of-band management:
Manage the storage subsystem using the controller Ethernet connections.

Controller (DNS/Network name, IPv4 address, or IPv6 address): 10.11.18.204

Controller (DNS/Network name, IPv4 address, or IPv6 address): 10.11.18.205

○ In-band management:
Manage the storage subsystem through an attached host.

Host (DNS/Network name, IPv4 address, or IPv6 address):

Step 6 From the Storage Subsystem Added window, click **No**.

Step 7 In the navigation pane on the left, right-click the storage subsystem that you just added and choose **Manage Storage Subsystem**.

Step 8 From the Subsystem Management window, click the **Setup** tab and then click the **Configure iSCSI Host Ports** link.

Figure 13-11 Configure the iSCSI Host Ports

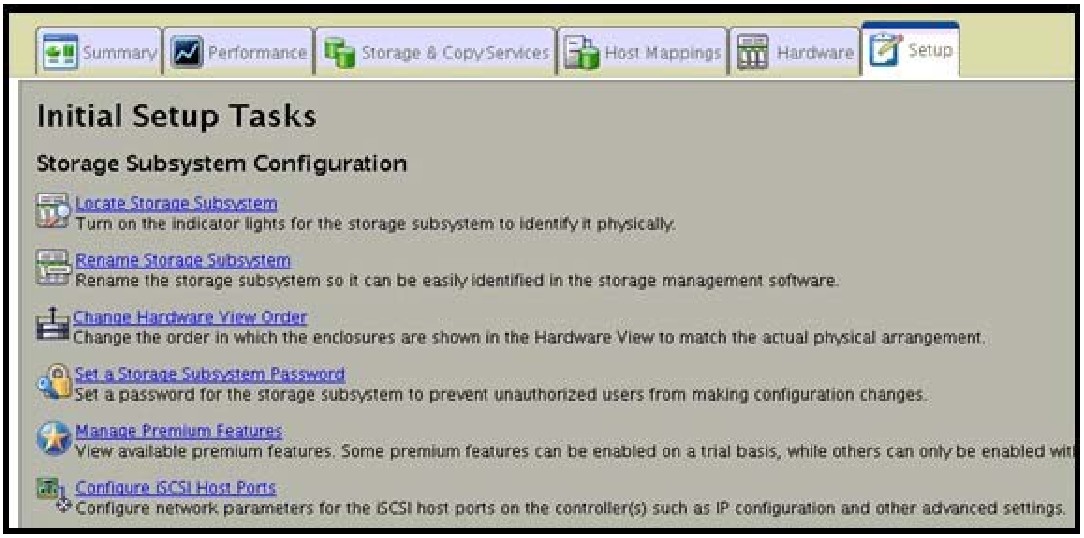

Initial Setup Tasks

Storage Subsystem Configuration

Locate Storage Subsystem
Turn on the indicator lights for the storage subsystem to identify it physically.

Rename Storage Subsystem
Rename the storage subsystem so it can be easily identified in the storage management software.

Change Hardware View Order
Change the order in which the enclosures are shown in the Hardware View to match the actual physical arrangement.

Set a Storage Subsystem Password
Set a password for the storage subsystem to prevent unauthorized users from making configuration changes.

Manage Premium Features
View available premium features. Some premium features can be enabled on a trial basis, while others can only be enabled wit

Configure iSCSI Host Ports
Configure network parameters for the iSCSI host ports on the controller(s) such as IP configuration and other advanced settings.

Step 9 From the Configure iSCSI Ports window, configure the IP addresses on the Controller A and Controller B iSCSI ports. Use the appropriate IP addresses from the appropriate table below, based on which storage enclosure you replaced and whether you have 1GB connections or 10GB connections to the storage enclosures.

Table 13-1 1 GB Installation iSCSI Port Addresses

Enclosure	Port	IP Address
Enclosure 1	Controller A, port 3	10.11.14.100
Enclosure 1	Controller A, port 4	10.11.15.100
Enclosure 1	Controller A, port 5	10.11.16.100
Enclosure 1	Controller A, port 6	10.11.17.100
Enclosure 1	Controller B, port 3	10.11.14.101
Enclosure 1	Controller B, port 4	10.11.15.101
Enclosure 1	Controller B, port 5	10.11.16.101
Enclosure 1	Controller B, port 6	10.11.17.101
Enclosure 2	Controller A, port 3	10.11.14.102
Enclosure 2	Controller A, port 4	10.11.15.102
Enclosure 2	Controller A, port 5	10.11.16.102
Enclosure 2	Controller A, port 6	10.11.17.102
Enclosure 2	Controller B, port 3	10.11.14.103
Enclosure 2	Controller B, port 4	10.11.15.103
Enclosure 2	Controller B, port 5	10.11.16.103
Enclosure 2	Controller B, port 6	10.11.17.103
Enclosure 3	Controller A, port 3	10.11.14.104
Enclosure 3	Controller A, port 4	10.11.15.104

Table 13-1 1 GB Installation iSCSI Port Addresses (continued)

Enclosure	Port	IP Address
Enclosure 3	Controller A, port 5	10.11.16.104
Enclosure 3	Controller A, port 6	10.11.17.104
Enclosure 3	Controller B, port 3	10.11.14.105
Enclosure 3	Controller B, port 4	10.11.15.105
Enclosure 3	Controller B, port 5	10.11.16.105
Enclosure 3	Controller B, port 6	10.11.17.105
Enclosure 4	Controller A, port 3	10.11.14.106
Enclosure 4	Controller A, port 4	10.11.15.106
Enclosure 4	Controller A, port 5	10.11.16.106
Enclosure 4	Controller A, port 6	10.11.17.106
Enclosure 4	Controller B, port 3	10.11.14.107
Enclosure 4	Controller B, port 4	10.11.15.107
Enclosure 4	Controller B, port 5	10.11.16.107
Enclosure 4	Controller B, port 6	10.11.17.107
Enclosure 5	Controller A, port 3	10.11.14.108
Enclosure 5	Controller A, port 4	10.11.15.108
Enclosure 5	Controller A, port 5	10.11.16.108
Enclosure 5	Controller A, port 6	10.11.17.108
Enclosure 5	Controller B, port 3	10.11.14.109
Enclosure 5	Controller B, port 4	10.11.15.109
Enclosure 5	Controller B, port 5	10.11.16.109
Enclosure 5	Controller B, port 6	10.11.17.109

Table 13-2 10 GB Installation iSCSI Port Addresses

Enclosure	Port	IP Address
Enclosure 1	Controller A, port 3	10.11.14.100
Enclosure 1	Controller B, port 3	10.11.14.101
Enclosure 2	Controller A, port 3	10.11.14.102
Enclosure 2	Controller B, port 3	10.11.14.103
Enclosure 3	Controller A, port 3	10.11.14.104
Enclosure 3	Controller B, port 3	10.11.14.105
Enclosure 4	Controller A, port 3	10.11.14.106
Enclosure 4	Controller B, port 3	10.11.14.107
Enclosure 5	Controller A, port 3	10.11.14.108
Enclosure 5	Controller B, port 3	10.11.14.109

Step 10 Check the **Enable ICMP Ping Responses** check box and uncheck the *Enable IPv6* check box.

Step 11 Click the **Advanced Port Settings** button. In the Advanced Port Settings window, check the **Enable Jumbo Frames** check box and enter an MTU of **9000**.

Configure the New Storage Enclosure

Follow these steps to configure the new storage enclosure:

Step 1 From an SSH connection to the VDS TC management server logged in as root, enter the command **cd /opt/pang/useful/configure_storages** to change directories.

Step 2 Enter the command **./configure_storages.py -s** <*the # of storage enclosures in the system*> **-n** < *the number of the replaced storage enclosure*> **-b** <*the number of the CEs in the system*>.

For example, for a system with 8 cache engines and 3 storage enclosures, where the third storage enclosure was replaced, enter the following command:

```
./configure_storages_pa.sh -s 3 -n 3 -b 8.
```

Step 3 Follow the on screen instructions. If prompted, enter the password for the root user that was provided by Cisco.

Note You may be prompted to enter the password several times.

Step 4 From the VNC console window, open the DS Storage Manager if it is not already open. If the storage enclosure that you replaced shows that it is in recovery mode and has a name of “Storage Subsystem Unnamed” perform the following steps:

a. Right-click the **Storage Subsystem unnamed** and choose **Manage Storage Subsystem**.

Figure 13-12 IBM DS Storage Manager

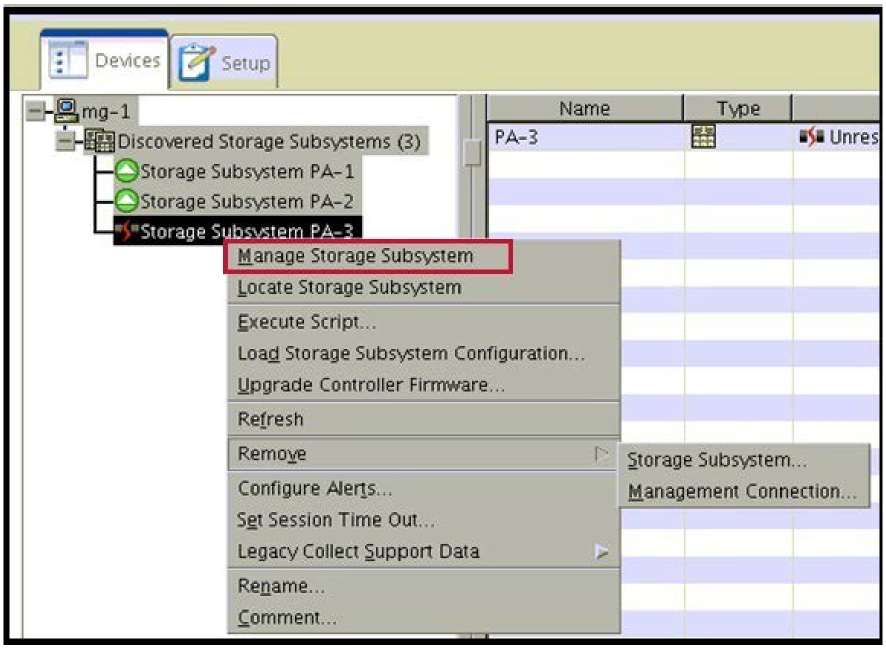

b. From the Unamed Subsystem Management window, choose **Monitor** > **Health** > **Clear Recovery Mode**.

Figure 13-13 Clear Recovery Mode

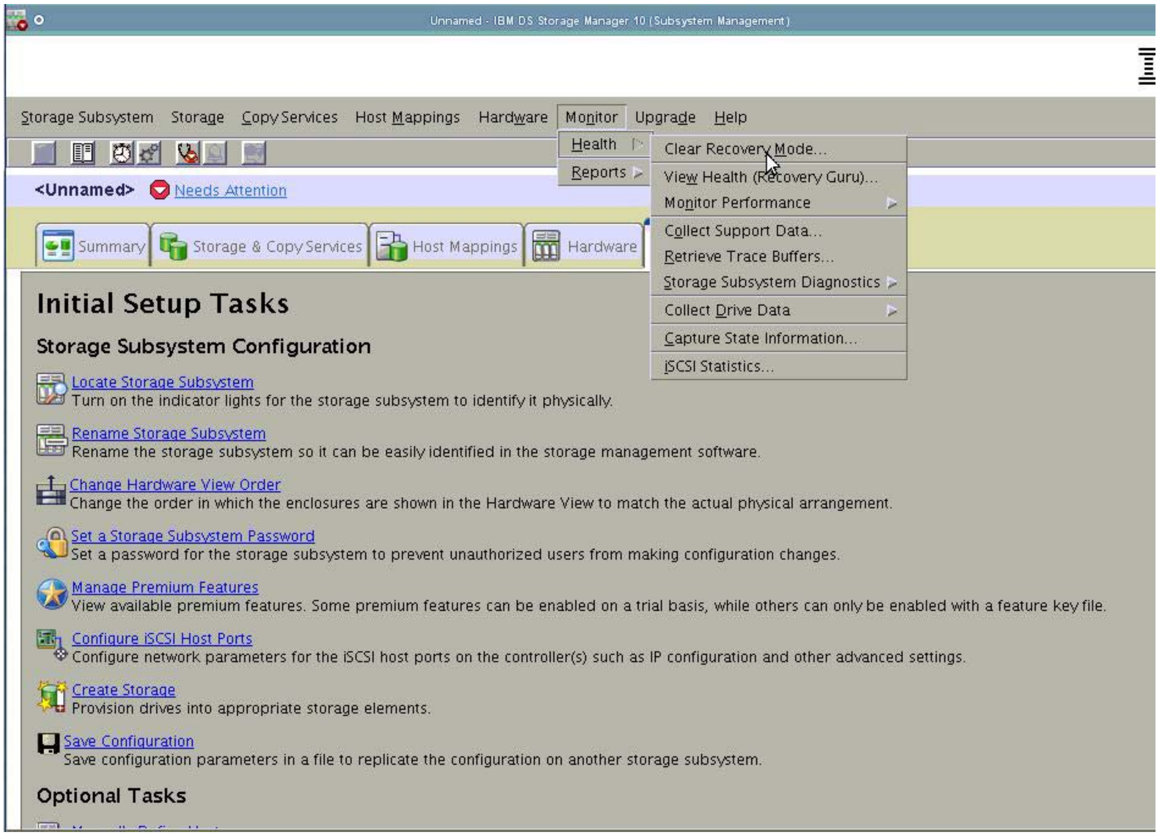

Initial Setup Tasks

Storage Subsystem Configuration

Locate Storage Subsystem
Turn on the indicator lights for the storage subsystem to identify it physically.

Rename Storage Subsystem
Rename the storage subsystem so it can be easily identified in the storage management software.

Change Hardware View Order
Change the order in which the enclosures are shown in the Hardware View to match the actual physical arrangement.

Set a Storage Subsystem Password
Set a password for the storage subsystem to prevent unauthorized users from making configuration changes.

Manage Premium Features
View available premium features. Some premium features can be enabled on a trial basis, while others can only be enabled with a feature key file.

Configure iSCSI Host Ports
Configure network parameters for the iSCSI host ports on the controller(s) such as IP configuration and other advanced settings.

Create Storage
Provision drives into appropriate storage elements.

Save Configuration
Save configuration parameters in a file to replicate the configuration on another storage subsystem.

Optional Tasks

c. From the popup window, enter **clear recovery mode** in the text box.

d. Click **OK**.

Figure 13-14 Clear Recovery Mode

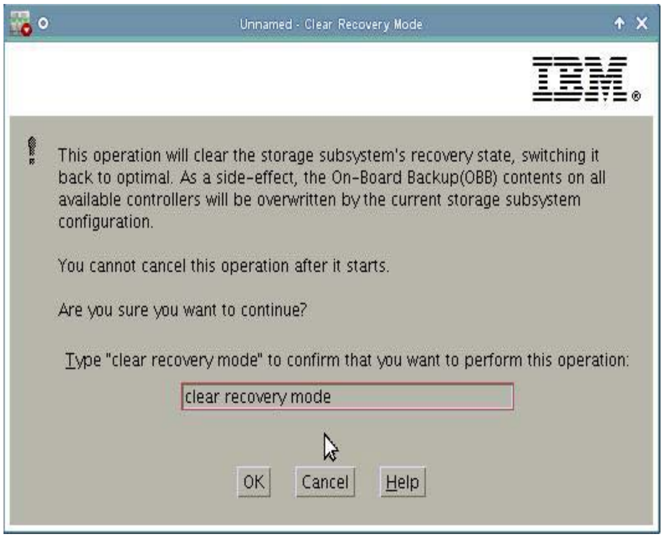

This operation will clear the storage subsystem's recovery state, switching it back to optimal. As a side-effect, the On-Board Backup(OBB) contents on all available controllers will be overwritten by the current storage subsystem configuration.

You cannot cancel this operation after it starts.

Are you sure you want to continue?

Type "clear recovery mode" to confirm that you want to perform this operation:

clear recovery mode

e. From the IBM DS Storage Manager, right-click the unnamed storage subsystem icon and choose **Rename** from the pop-up menu.

Figure 13-15 Rename Storage Subsystem

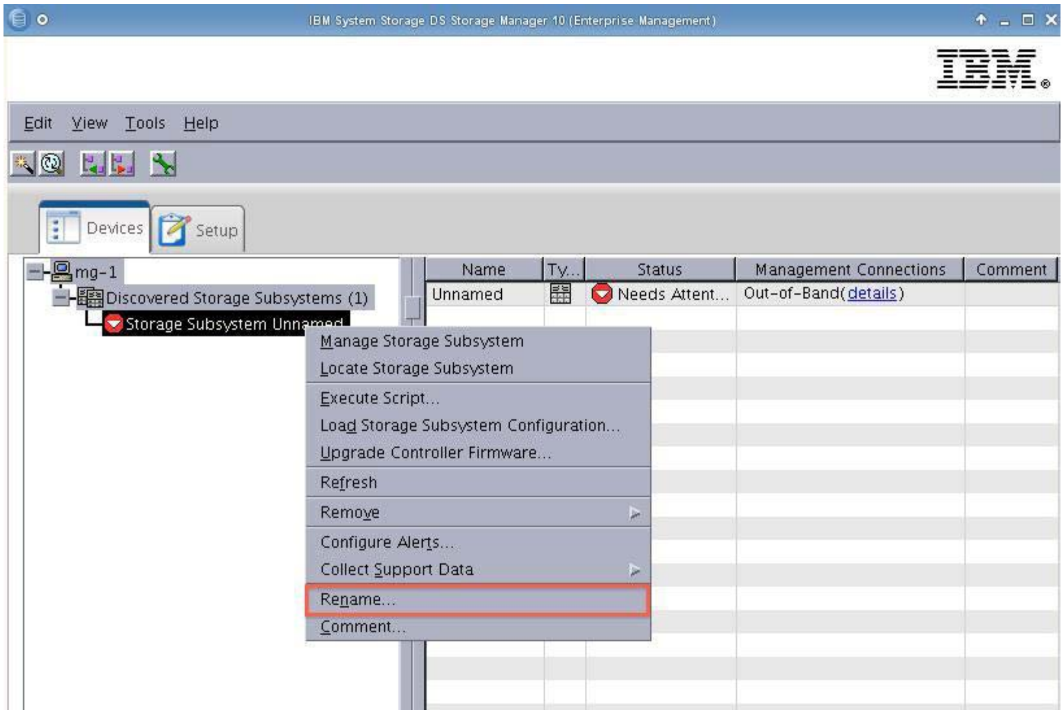

f. In the Rename Storage Subsystem window that appears, enter the name for the storage subsystem, such as Enclosure-1, and click **OK**.

Figure 13-16 Rename Storage Subsystem Dialog Box

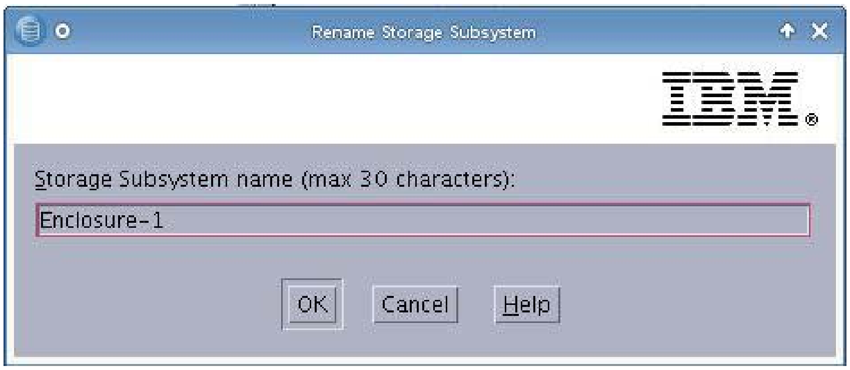

Storage Subsystem name (max 30 characters):

Enclosure-1

g. When the Warning Message appears, click **Yes**.

Figure 13-17 Warning Message

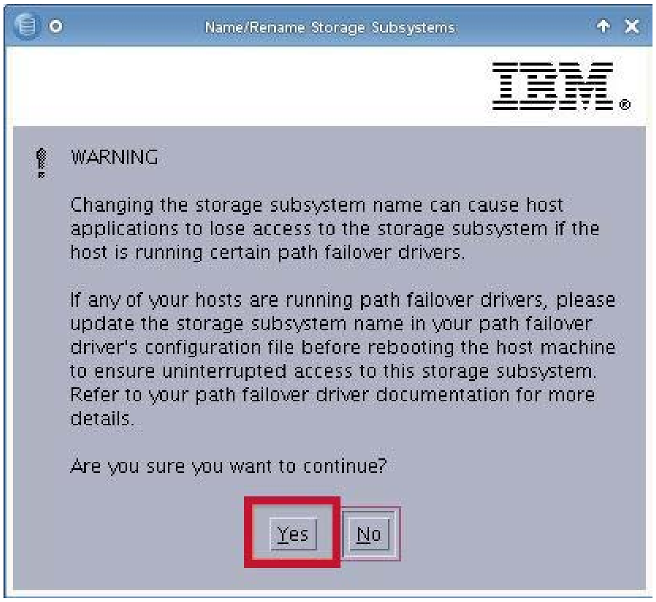

WARNING

Changing the storage subsystem name can cause host applications to lose access to the storage subsystem if the host is running certain path failover drivers.

If any of your hosts are running path failover drivers, please update the storage subsystem name in your path failover driver's configuration file before rebooting the host machine to ensure uninterrupted access to this storage subsystem. Refer to your path failover driver documentation for more details.

Are you sure you want to continue?

Step 5 To perform disk partitioning you must log into CE-1. To log into CE-1, from the SSH connection to the VDS TC management server, enter the command **ssh root@ce-1**.

Step 6 To find the non-partitioned sd devices, enter the command **fdisk -l**. Look for the sd devices that say “doesn't contain a valid partition table”. Make note of these devices. These are the devices that you will need to format.

For example, in the following output sd devices sdax, sday, and sdaz are non-partitioned sd devices because they show “doesn't contain a valid partition table” as part of their output.

```
# fdisk -l
. . .
Disk /dev/sdaw: 299.4 GB, 299462819840 bytes
255 heads, 63 sectors/track, 36407 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes
Disk identifier: 0x00000000

Device Boot Start  End     Blocks  Id  System
/dev/sdaw1     1   1        8001   83  Linux
/dev/sdaw2     2   2050 16458592+ 83  Linux
/dev/sdaw3  2051   2573  4200997+ 83  Linux
/dev/sdaw4  2574  36407 271771605 83 Linux

Disk /dev/sdax: 299.4 GB, 299462819840 bytes
255 heads, 63 sectors/track, 36407 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes
Disk identifier: 0x00000000

Disk /dev/sdax doesn't contain a valid partition table

Disk /dev/sday: 299.4 GB, 299462819840 bytes
255 heads, 63 sectors/track, 36407 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes
Disk identifier: 0x00000000

Disk /dev/sday doesn't contain a valid partition table
Disk /dev/sdaz: 299.4 GB, 299462819840 bytes
255 heads, 63 sectors/track, 36407 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes
Disk identifier: 0x00000000

Disk /dev/sdaz doesn't contain a valid partition table
```

Step 7 Next you must determine an available volume number. You will use this information to format the drives of the replacement storage array in the next step. To determine available volume numbers, enter the command **/opt/pang/bin/fdisk.sh**. For example, in the following output, there are no gaps in the middle of the volumes. The next available volume number is 49.

```
ce-1:/opt/pang/bin # /opt/pang/bin/fdisk.sh
PeerApp,0001,U /dev/sdz1
PeerApp,0002,U /dev/sdaa1
PeerApp,0003,U /dev/sdab1
PeerApp,0004,U /dev/sdac1
PeerApp,0005,U /dev/sdad1
PeerApp,0006,U /dev/sdae1
PeerApp,0007,U /dev/sdaf1
PeerApp,0008,U /dev/sdag1
PeerApp,0009,U /dev/sdah1
PeerApp,0010,U /dev/sdai1
PeerApp,0011,U /dev/sdaj1
PeerApp,0012,U /dev/sdak1
PeerApp,0013,U /dev/sdal1
PeerApp,0014,U /dev/sdam1
PeerApp,0015,U /dev/sdan1
PeerApp,0016,U /dev/sdao1
PeerApp,0017,U /dev/sdap1
PeerApp,0018,U /dev/sdaq1
PeerApp,0019,U /dev/sdar1
PeerApp,0020,U /dev/sdas1
PeerApp,0021,U /dev/sdat1
PeerApp,0022,U /dev/sdau1
PeerApp,0023,U /dev/sdav1
PeerApp,0024,U /dev/sdaw1
PeerApp,0025,U /dev/sdb1
PeerApp,0026,U /dev/sdc1
PeerApp,0027,U /dev/sdd1
PeerApp,0028,U /dev/sde1
PeerApp,0029,U /dev/sdf1
PeerApp,0030,U /dev/sdg1
PeerApp,0031,U /dev/sdh1
PeerApp,0032,U /dev/sdi1
PeerApp,0033,U /dev/sdj1
PeerApp,0034,U /dev/sdk1
PeerApp,0035,U /dev/sdl1
PeerApp,0036,U /dev/sdm1
PeerApp,0037,U /dev/sdn1
PeerApp,0038,U /dev/sdo1
PeerApp,0039,U /dev/sdp1
PeerApp,0040,U /dev/sdq1
PeerApp,0041,U /dev/sdr1
PeerApp,0042,U /dev/sds1
PeerApp,0043,U /dev/sdt1
PeerApp,0044,U /dev/sdu1
PeerApp,0045,U /dev/sdv1
PeerApp,0046,U /dev/sdw1
PeerApp,0047,U /dev/sdx1
PeerApp,0048,U /dev/sdy1
```

Note You can use a number available in the middle if there is a gap, or you can use the next available number after the highest volume number.

Step 8 Follow these steps to format all of the non-partitioned sd devices that you discovered in Step 6 one at a time:

a. Enter the command **cd /opt/pang/useful** to change directories.

b. Enter the command **./format_disks.sh -format_one** <*the sd ID*> <*the next available volume number*>, where <*the sd ID*> is the sd id of a non-partitioned sd devices that you discovered in Step 6 and <*the next available volume number*> is an available volume number that you discovered in Step 7.

For example, the following command formats the sd device with the id of sdax and assigns it a volume number of 49:

```
ce-1:/opt/pang/useful # ./format_disks.sh -format_one /dev/sdax 49
Formating: /dev/sdax index for signature 49
disk[49]: /dev/sdax - Erasing boot sector... done
Writing partition table... done
Erasing created partitions... done
Initializing the disk... done
Tuning the filesystem... Marking the disk... done
format_disks.sh: 4.2 finished, see /opt/pang/useful/installog.txt for
details
```

Step 9 Repeat step 8, using the information from Step 6 and Step 7 to format the remainder of the non-partitioned devices and assign them available volume number.

Step 10 Reboot all of the cache engines in the grid, except for CE-1, so that they can re-read the formatted disks. Follow these commands to reboot the cache engines:

a. From CE-1, enter the command **exit** to return to the VDS TC management sever.

b. Enter the command **ssh root@ce-***X*, where *X* is the number of a cache engine that you need to reboot.

c. Enter the command **reboot -f** to reboot the cache engine. This will end your SSH session to the cache engine and return you to the VDS TC management server.

d. Repeat Step b and Step c for the remaining cache engines.

Step 11 After you have finished rebooting all of the cache engines, log into each cache engine to confirm that the iSCSI connections were established to the replaced enclosure:

a. Enter the command **ssh root@ce-***X*, where *X* is the number of a cache engine to check.

b. Enter the command **iscsiadm -m session**. If the iSCSI connections have been established, you will see the following results:

```
tcp: [1] 10.11.14.102:3260,1 iqn.1992-08.com.netapp:2752.600a0980005dde250000000053d8deff
tcp: [2] 10.11.14.103:3260,2 iqn.1992-08.com.netapp:2752.600a0980005dde250000000053d8deff
tcp: [3] 10.11.15.103:3260,2 iqn.1992-08.com.netapp:2752.600a0980005dde250000000053d8deff
tcp: [4] 10.11.15.102:3260,1 iqn.1992-08.com.netapp:2752.600a0980005dde250000000053d8deff
tcp: [5] 10.11.14.104:3260,1 iqn.1992-08.com.netapp:2752.600a0980005dde220000000053d8c40e
tcp: [6] 10.11.15.105:3260,2 iqn.1992-08.com.netapp:2752.600a0980005dde220000000053d8c40e
tcp: [7] 10.11.14.105:3260,2 iqn.1992-08.com.netapp:2752.600a0980005dde220000000053d8c40e
tcp: [8] 10.11.15.104:3260,1 iqn.1992-08.com.netapp:2752.600a0980005dde220000000053d8c40e
tcp: [9] 10.11.14.100:3260,1 iqn.1992-08.com.netapp:2752.600a0980005ddc030000000053d8c4b6
tcp: [10] 10.11.14.101:3260,2 iqn.1992-08.com.netapp:2752.600a0980005ddc030000000053d8c4b6
tcp: [11] 10.11.15.100:3260,1 iqn.1992-08.com.netapp:2752.600a0980005ddc030000000053d8c4b6
tcp: [12] 10.11.15.101:3260,2 iqn.1992-08.com.netapp:2752.600a0980005ddc030000000053d8c4b6
```

c. Repeat Steps a and b for the remaining cache engines.

Step 12 Follow these steps to restart the VDS TC service:

a. Log into the VDS TC CLI. The CLI prompt console> appears.

b. Enter the command **enable** to enter Enable mode. At the password prompt that appears, enter the enable password. The CLI prompt console# appears.

c. Enter the command **oper service start** to start the VDS TC service.

d. Wait a few minutes and then enter the command **show status** to confirm that the VDS TC service has started. The Device State should show “started”. For example:

```
console# show status
Cluster state: enabled

Server Slot          Status               Operational state    Device state
Administrative state
ce-1                 powered on           enabled              started
unlocked
ce-2                 powered on           enabled              started
unlocked
ce-3                 powered on           enabled              started
unlocked
ce-4                 powered on           enabled              started
unlocked
ce-5                 powered on           enabled              started
unlocked
ce-6                 powered on           enabled              started
unlocked
ce-7                 powered on           enabled              started
unlocked
ce-8                 powered on           enabled              started
unlocked
ce-9                 powered on           enabled              started
unlocked
ce-10                powered on           enabled              started
unlocked
ce-11                powered on           enabled              started
unlocked
ce-12                powered on           enabled              started
unlocked
ce-13                powered on           enabled              started
unlocked
ce-14                powered on           enabled              started
unlocked
ce-15                powered on           enabled              started
unlocked
ce-16                powered on           enabled              started
unlocked
```

e. Enter the command **show volumes** to confirm that all of the volumes show mounted. For example:

```
console# show volumes
Licensed volumes : 48
Volume name          State                Owner
/mnt/vol1            mounted              ce-1
/mnt/vol2            mounted              ce-2
/mnt/vol3            mounted              ce-2
/mnt/vol4            mounted              ce-2
/mnt/vol5            mounted              ce-2
/mnt/vol6            mounted              ce-2
/mnt/vol7            mounted              ce-2
/mnt/vol8            mounted              ce-2
```

```
/mnt/vol9           mounted             ce-2
/mnt/vol10          mounted             ce-2
/mnt/vol11          mounted             ce-2
/mnt/vol12          mounted             ce-2
/mnt/vol13          mounted             ce-1
/mnt/vol14          mounted             ce-1
/mnt/vol15          mounted             ce-1
/mnt/vol16          mounted             ce-1
/mnt/vol17          mounted             ce-1
/mnt/vol18          mounted             ce-1
/mnt/vol19          mounted             ce-1
/mnt/vol20          mounted             ce-1
/mnt/vol21          mounted             ce-1
/mnt/vol22          mounted             ce-1
/mnt/vol23          mounted             ce-2
/mnt/vol24          mounted             ce-1
```

Step 13 On the PBR routers, start redirecting traffic to the VDS TC solution.

Step 14 After you have begun redirecting traffic to the VDS TC solution, follow these steps to confirm that the volumes on the replaced enclosure are being used:

a. Log into VDS TC Manager and choose **Statistics** > **Storage**.

b. Verify that the status of the volumes show Active and that there is a number other than 0 in the Usage % column.

Figure 13-18 Volume Status and Usage

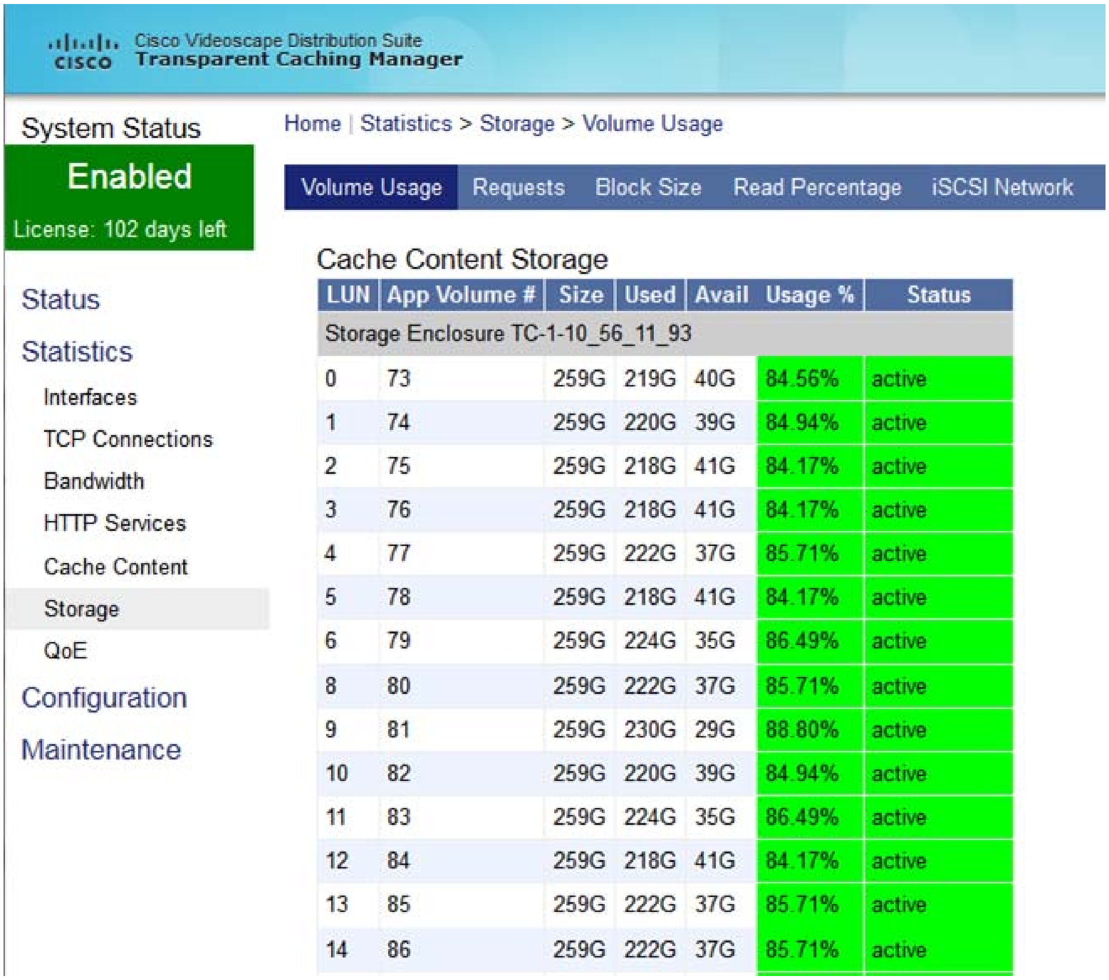

Cache Content Storage

LUN	App Volume #	Size	Used	Avail	Usage %	Status
Storage Enclosure TC-1-10_56_11_93						
0	73	259G	219G	40G	84.56%	active
1	74	259G	220G	39G	84.94%	active
2	75	259G	218G	41G	84.17%	active
3	76	259G	218G	41G	84.17%	active
4	77	259G	222G	37G	85.71%	active
5	78	259G	218G	41G	84.17%	active
6	79	259G	224G	35G	86.49%	active
8	80	259G	222G	37G	85.71%	active
9	81	259G	230G	29G	88.80%	active
10	82	259G	220G	39G	84.94%	active
11	83	259G	224G	35G	86.49%	active
12	84	259G	218G	41G	84.17%	active
13	85	259G	222G	37G	85.71%	active
14	86	259G	222G	37G	85.71%	active

VDS TC Integrated Appliance Cache Disk Replacement

This chapter discusses the steps that are required to replace a cache disk in a VDS TC Integrated Appliance installation.

Symptoms of a Failed VDS TC Storage Disk

One symptom of a failed storage disk is if the following message frequently appears in the SNMP traps section of the Status > Dashboard window in the VDS TC Manager:

“warning: Volume /dev/*volume_name* (*volume_id*) is turned OFF because of errors”

For example:

```
“warning: Volume /dev/sd03 (id 12) is turned OFF because of errors”
```

Figure 14-1 SNMP Traps Message

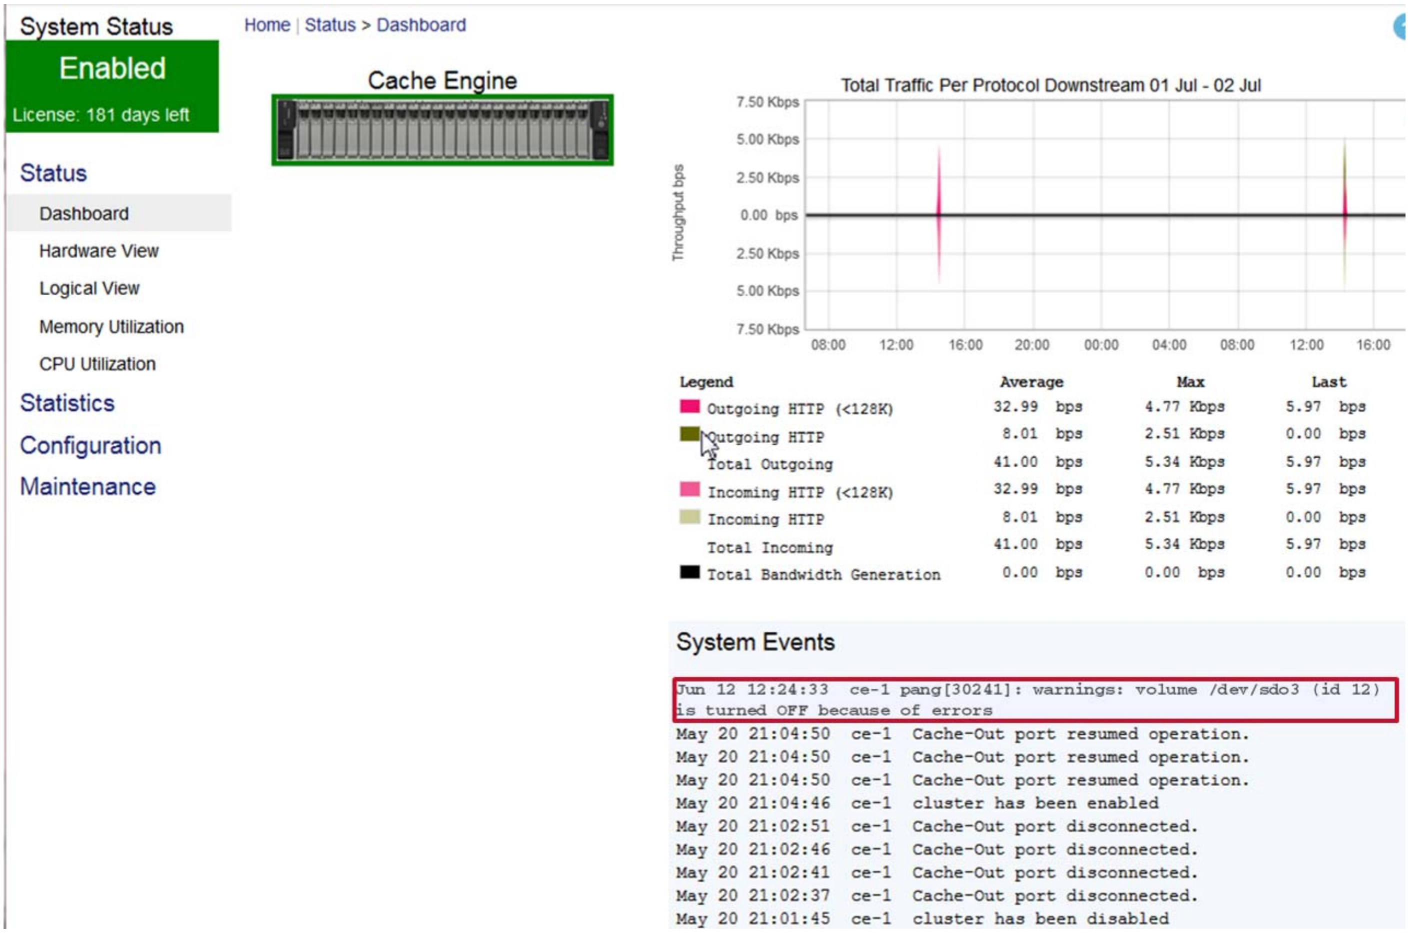

Legend	Average	Max	Last
🟥 Outgoing HTTP (<128K)	32.99 bps	4.77 Kbps	5.97 bps
🟫 Outgoing HTTP	8.01 bps	2.51 Kbps	0.00 bps
Total Outgoing	41.00 bps	5.34 Kbps	5.97 bps
🟥 Incoming HTTP (<128K)	32.99 bps	4.77 Kbps	5.97 bps
🟨 Incoming HTTP	8.01 bps	2.51 Kbps	0.00 bps
Total Incoming	41.00 bps	5.34 Kbps	5.97 bps
⬛ Total Bandwidth Generation	0.00 bps	0.00 bps	0.00 bps

System Events

```
Jun 12 12:24:33  ce-1 pang[30241]: warnings: volume /dev/sdo3 (id 12)
is turned OFF because of errors
May 20 21:04:50  ce-1  Cache-Out port resumed operation.
May 20 21:04:50  ce-1  Cache-Out port resumed operation.
May 20 21:04:50  ce-1  Cache-Out port resumed operation.
May 20 21:04:46  ce-1  cluster has been enabled
May 20 21:02:51  ce-1  Cache-Out port disconnected.
May 20 21:02:46  ce-1  Cache-Out port disconnected.
May 20 21:02:41  ce-1  Cache-Out port disconnected.
May 20 21:02:37  ce-1  Cache-Out port disconnected.
May 20 21:01:45  ce-1  cluster has been disabled
```

Please note that this error message can also be caused by problems other than a failed disk.

Note The Volume Usage window in VDS TC Manager, available at Statistics > Storage, may show the Volume State as active even if the drive has failed, as seen in the screenshots below.

Figure 14-2 Volume Usage Tab

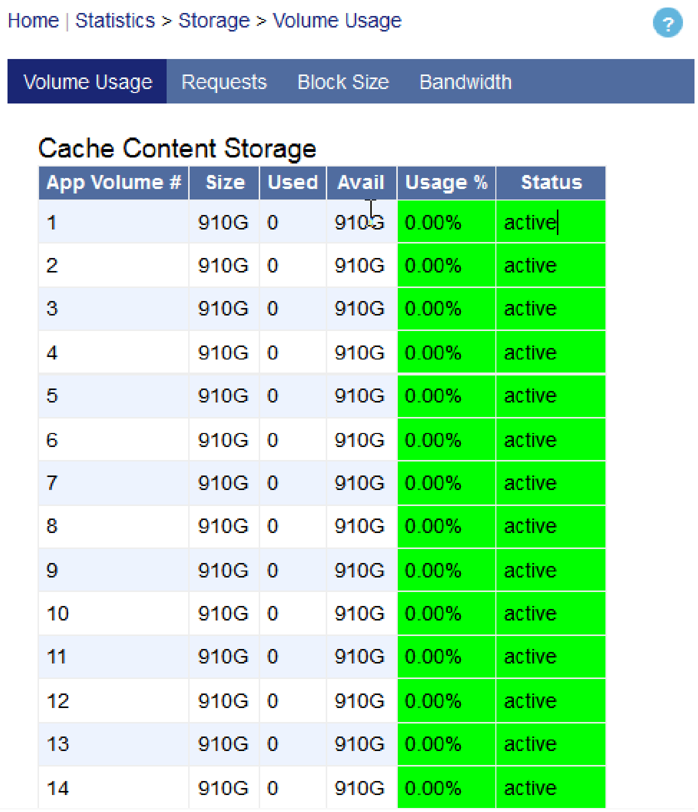

Cache Content Storage

App Volume #	Size	Used	Avail	Usage %	Status
1	910G	0	910G	0.00%	active
2	910G	0	910G	0.00%	active
3	910G	0	910G	0.00%	active
4	910G	0	910G	0.00%	active
5	910G	0	910G	0.00%	active
6	910G	0	910G	0.00%	active
7	910G	0	910G	0.00%	active
8	910G	0	910G	0.00%	active
9	910G	0	910G	0.00%	active
10	910G	0	910G	0.00%	active
11	910G	0	910G	0.00%	active
12	910G	0	910G	0.00%	active
13	910G	0	910G	0.00%	active
14	910G	0	910G	0.00%	active

Integrated Appliance Cache Disk Replacement Procedure

Follow this procedure to replace a cache disk in a VDS TC Integrated Appliance installation:

Installing the Storage Manager

Procedure

Step 1 Perform the following steps to download the latest version of the MegaRaid Storage Manager for Linux 64 bit.

a. Go to http://www.avagotech.com/support.

b. From the main Support page that appears, under Find Documentation and Downloads, click **Support Documents and Downloads**. The Support Document and Downloads window appears.

c. From the Product Family drop-down list choose **All Server Storage**, in the Product drop-down list leave **All** selected, and from the Asset Type drop-down list choose **Management Software and Tools**. Click **Search**.

d. From the search results window that appears, click **Management Software and Tools**.

e. From the list that appears, in the row labeled “Latest MegaRAID Storage Manager (MSM)" for Linux 64-bit, click the download icon.

f. From the Download Agreement window that appears, click **I agree** to continue the download.

Step 2 Use SCP software, such as WinSCP, to copy the file that you download to the /tmp directory on the VDS TC Integrated Appliance.

Step 3 Using SSH software, such as Putty, open an SSH connection to the IP address of the VDS TC Integrated Appliance.

Step 4 Log into the system using the username **padmin** and the password provided by Cisco.

Step 5 Enter the command **su root** to change to the root user. Enter the password for the root user when prompted. This password is provided by Cisco.

Step 6 Enter the command **cd /opt/pang/utilities/RAID** to change directories.

Step 7 To move the file that you downloaded in Step 1 to the correct folder, enter the command **mv /tmp/***LSI_mgr_filename*, where *LSI_mgr_filename* is the name of the file that you copied to the /tmp folder in Step 1.

Step 8 Enter the command **tar -zxvf** *LSI_mgr_filename***.tar.gz**, where *LSI_mgr_filename***.tar.gz** is the name of the file that you downloaded in Step 1, to extract the Storage Manager files.

Step 9 Enter the command **cd disk**.

Step 10 Enter the command **find . -name "*snmp*" -exec rm -v {} \;**

Step 11 To start installing the Storage Manager, enter the command **./install.csh**.

Step 12 When the License Agreement appears, read the agreement and press **Y** to accept the license agreement and continue with the installation.

Figure 14-3 Accept the Storage Manager License Agreement

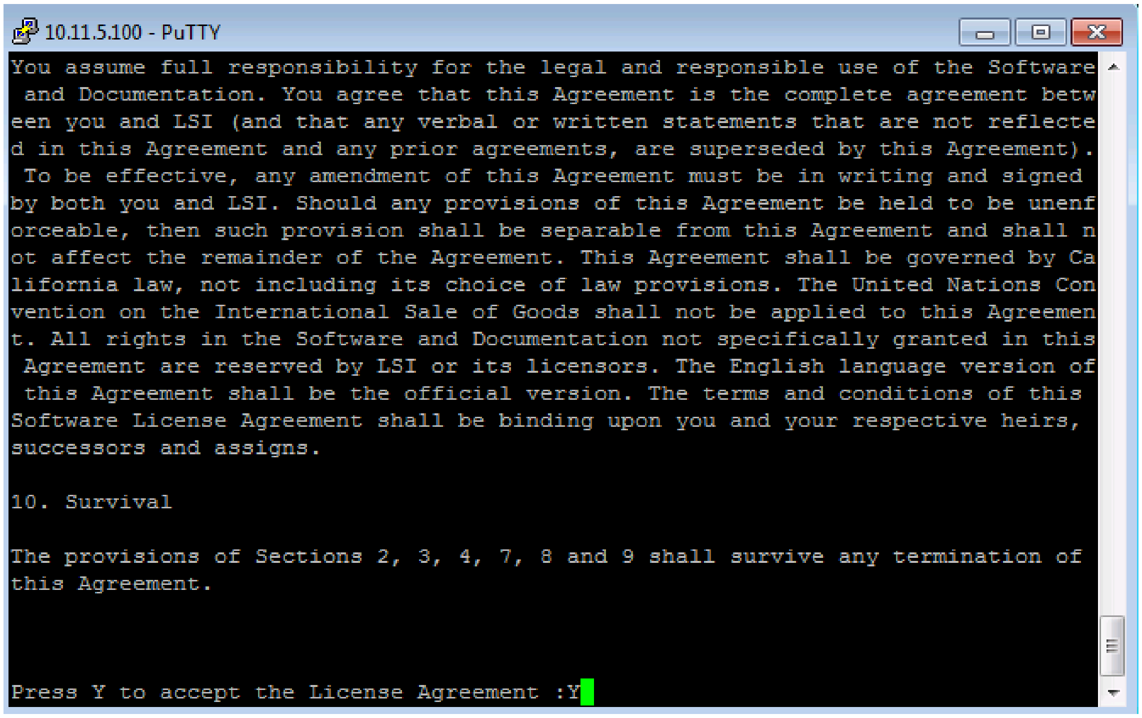

```
You assume full responsibility for the legal and responsible use of the Software
 and Documentation. You agree that this Agreement is the complete agreement betw
een you and LSI (and that any verbal or written statements that are not reflecte
d in this Agreement and any prior agreements, are superseded by this Agreement).
 To be effective, any amendment of this Agreement must be in writing and signed
by both you and LSI. Should any provisions of this Agreement be held to be unenf
orceable, then such provision shall be separable from this Agreement and shall n
ot affect the remainder of the Agreement. This Agreement shall be governed by Ca
lifornia law, not including its choice of law provisions. The United Nations Con
vention on the International Sale of Goods shall not be applied to this Agreemen
t. All rights in the Software and Documentation not specifically granted in this
 Agreement are reserved by LSI or its licensors. The English language version of
 this Agreement shall be the official version. The terms and conditions of this
Software License Agreement shall be binding upon you and your respective heirs,
successors and assigns.

10. Survival

The provisions of Sections 2, 3, 4, 7, 8 and 9 shall survive any termination of
this Agreement.



Press Y to accept the License Agreement :Y
```

Step 13 When asked to choose the Setup Type, enter **4**. This option will only install components required for local configuration.

Figure 14-4 *Setup Type*

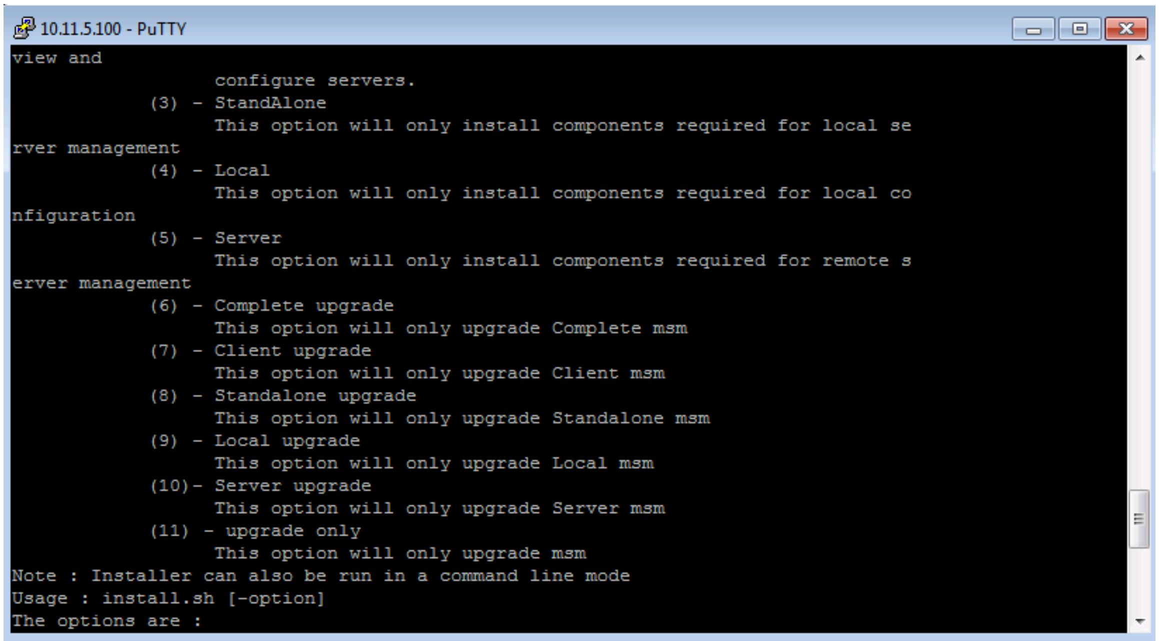

```
view and
                  configure servers.
            (3) - StandAlone
                  This option will only install components required for local se
rver management
            (4) - Local
                  This option will only install components required for local co
nfiguration
            (5) - Server
                  This option will only install components required for remote s
erver management
            (6) - Complete upgrade
                  This option will only upgrade Complete msm
            (7) - Client upgrade
                  This option will only upgrade Client msm
            (8) - Standalone upgrade
                  This option will only upgrade Standalone msm
            (9) - Local upgrade
                  This option will only upgrade Local msm
            (10)- Server upgrade
                  This option will only upgrade Server msm
            (11) - upgrade only
                  This option will only upgrade msm
Note : Installer can also be run in a command line mode
Usage : install.sh [-option]
The options are :
```

Step 14 When prompted to enter the key size, enter **1**.

Figure 14-5 *Enter the Key Size*

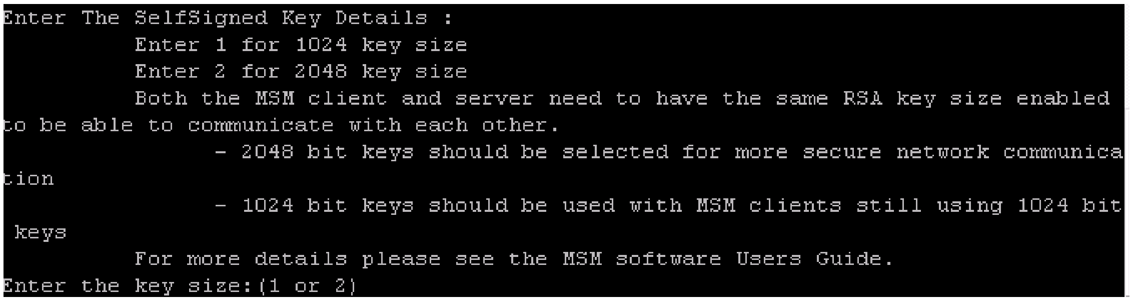

```
Enter The SelfSigned Key Details :
          Enter 1 for 1024 key size
          Enter 2 for 2048 key size
          Both the MSM client and server need to have the same RSA key size enabled
to be able to communicate with each other.
                - 2048 bit keys should be selected for more secure network communica
tion
                - 1024 bit keys should be used with MSM clients still using 1024 bit
 keys
          For more details please see the MSM software Users Guide.
Enter the key size:(1 or 2)
```

Step 15 When prompted to enter the Alert Notifications of Users Choice, enter **1**.

Figure 14-6 Alert Notifications of Users Choice

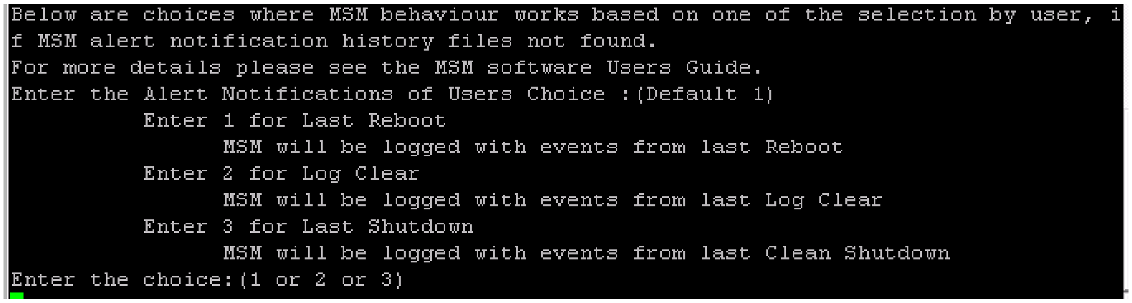

```
Below are choices where MSM behaviour works based on one of the selection by user, i
f MSM alert notification history files not found.
For more details please see the MSM software Users Guide.
Enter the Alert Notifications of Users Choice :(Default 1)
          Enter 1 for Last Reboot
                MSM will be logged with events from last Reboot
          Enter 2 for Log Clear
                MSM will be logged with events from last Log Clear
          Enter 3 for Last Shutdown
                MSM will be logged with events from last Clean Shutdown
Enter the choice:(1 or 2 or 3)
```

Step 16 Wait for the Storage Manager installation to complete and then proceed to the Replacing the Integrated Appliance Cache Disk task.

Replacing the Integrated Appliance Cache Disk

Follow this procedure to replace the failed cache disk in a VDS TC Integrated Appliance installation.

Procedure

Step 1 Enter the following commands to start the VNC server:

a. **rm /tmp/.X1-lock** (This step is optional. It is only required if there is a leftover lock file from a previous vncserver instance.)

b. **rm /tmp/.X11-unix/X1** (This step is optional. It is only required if there is a leftover lock file from a previous vncserver instance.)

c. **vncserver :1**

Step 2 Use a VNC client to connect from your local computer to the VDS TC management server.

Figure 14-7 VNC Client

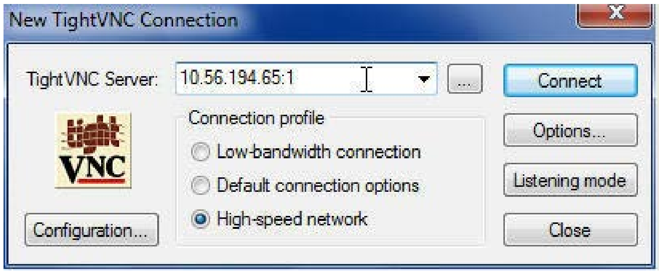

New TightVNC Connection

TightVNC Server: 10.56.194.65:1

Connection profile

○ Low-bandwidth connection

○ Default connection options

● High-speed network

Step 3 From the VNC console window, enter the command **cd /usr/local/MegaRAID\ Storage\ Manager/** and then enter the command **./startupui.sh** to start the MegaRAID Storage Manager.

Note In order to run the MegaRAID Storage Manager, you may need to restart the Vivaldi framework program. To restart the Vivaldi framework program, enter the command /etc/init.d/vivaldiframeworkd start while logged into the VDS TC Integrated Appliance as root. After you are finished using Storage Manager, remember to stop the Vivaldi framework program be entering the /etc/init.d/vivaldiframeworkd stop command.

Step 4 From the Storage Manager login window, login in as root by entering the username **root** and the password for the root user, and then click **Login**.

Figure 14-8 Storage Manager Login Window

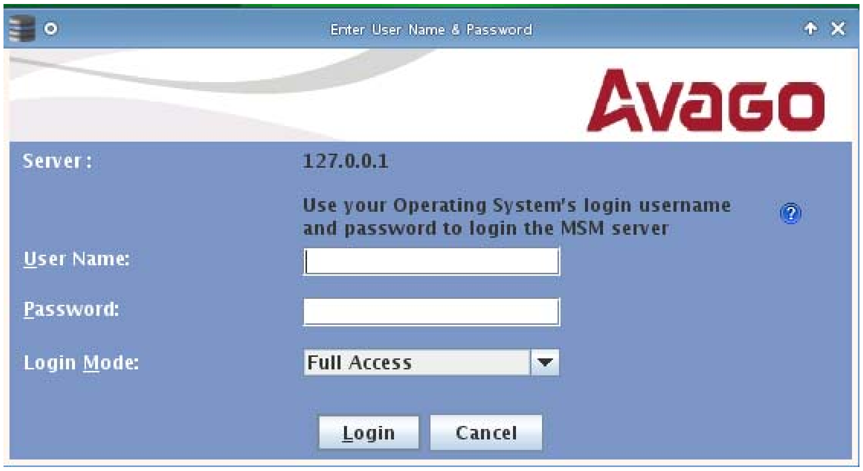

Enter User Name & Password

Server : 127.0.0.1

Use your Operating System's login username and password to login the MSM server

User Name:

Password:

Login Mode: Full Access

Step 5 From the MegaRAID Storage Manager window, click the **Physical** tab. The drive that failed will show a status of Failed.

Step 6 Follow the “Drive Replacement Procedure” in the “Maintaining the Server” chapter of the *Cisco UCS C240 Server Installation and Service Guide* located at http://www.cisco.com/en/US/docs/unified_computing/ucs/c/hw/C240/install/replace.html#wp1220308, to replace the faulty hard drive.

Step 7 After the disk has been replaced, using SSH software log back into the VDS TC Integrated Appliance, using the username **padmin**.

Step 8 Enter the command **su root** to change to the root user. Enter the password for the root user when prompted. This password is provided by Cisco.

Step 9 Enter the following commands to start the VNC server:

a. **rm /tmp/.X1-lock** (This step is optional. It is only required if there is a leftover lock file from a previous vncserver instance.)

b. **rm /tmp/.X11-unix/X1** (This step is optional. It is only required if there is a leftover lock file from a previous vncserver instance.)

c. **vncserver :1**

Step 10 Use a VNC client to connect from your local computer to the VDS TC management server.

Figure 14-9 VNC Client

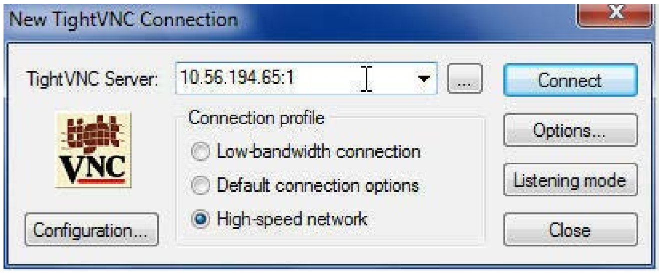
New TightVNC Connection

TightVNC Server: 10.56.194.65:1

Connection profile
○ Low-bandwidth connection
○ Default connection options
● High-speed network

Step 11 From the VNC console window, enter the command **cd /usr/local/MegaRAID\ Storage\ Manager/** and then enter the command **./startupui.sh**.

Step 12 From the Storage Manager login window, login in as root by entering the username **root** and the password for the root user, and then click **Login**.

Figure 14-10 Storage Manager Login Window

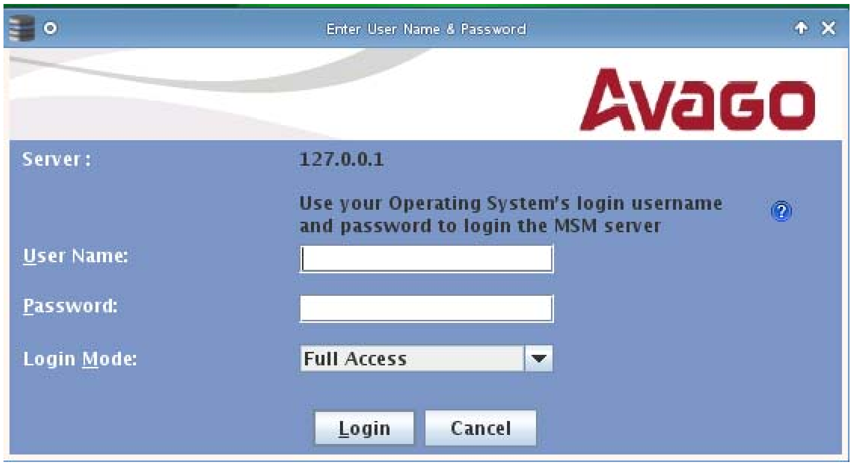
Enter User Name & Password

Server : 127.0.0.1

Use your Operating System's login username and password to login the MSM server

User Name:

Password:

Login Mode: Full Access

Step 13 From the MegaRAID Storage Manager window, click the **Physical** tab. The new disk should show a status of “Unconfigured Good”.

Figure 14-11 Storage Manager Physical Tab

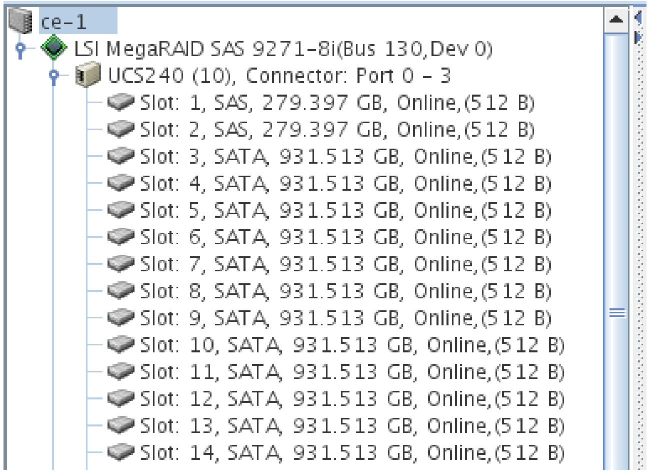

ce-1
LSI MegaRAID SAS 9271-8i(Bus 130,Dev 0)
UCS240 (10), Connector: Port 0 - 3
Slot: 1, SAS, 279.397 GB, Online,(512 B)
Slot: 2, SAS, 279.397 GB, Online,(512 B)
Slot: 3, SATA, 931.513 GB, Online,(512 B)
Slot: 4, SATA, 931.513 GB, Online,(512 B)
Slot: 5, SATA, 931.513 GB, Online,(512 B)
Slot: 6, SATA, 931.513 GB, Online,(512 B)
Slot: 7, SATA, 931.513 GB, Online,(512 B)
Slot: 8, SATA, 931.513 GB, Online,(512 B)
Slot: 9, SATA, 931.513 GB, Online,(512 B)
Slot: 10, SATA, 931.513 GB, Online,(512 B)
Slot: 11, SATA, 931.513 GB, Online,(512 B)
Slot: 12, SATA, 931.513 GB, Online,(512 B)
Slot: 13, SATA, 931.513 GB, Online,(512 B)
Slot: 14, SATA, 931.513 GB, Online,(512 B)

Step 14 To recreate the LUN, right-click the RAID controller and choose **Create Virtual Drive**.

Figure 14-12 Create Virtual Drive

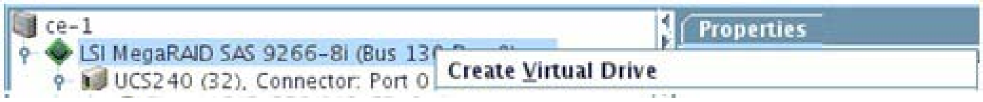

ce-1
LSI MegaRAID SAS 9266-8i (Bus 13
UCS240 (32), Connector: Port 0

Step 15 If Create Virtual Drive option appeared in Step 14, go on to Step 16. If the Create Virtual Drive option did *not* appear in the pop-up menu in Step 14, follow these steps:

a. Right-click the raid controller and choose **Scan Foreign Configuration** and when prompted, clear it.

b. Right-click the raid controller and choose **Preserved Cache** and when prompted, follow the directions on the screen to clear it.

c. Check if there are any unnecessary processes running (if there are, cancel them).

d. Right-click the RAID controller and choose **Create Virtual Drive**.

Step 16 From the Storage Manager window that appears, choose **Advanced** and click **Next**.

Step 17 From the RAID Level drop-down list on the Create Drive Group window, choose **RAID 0**.

Step 18 In the Select Unconfigured Drives pane, choose the drive that is listed and click **Add**.

Figure 14-13 Select Unconfigured Drive

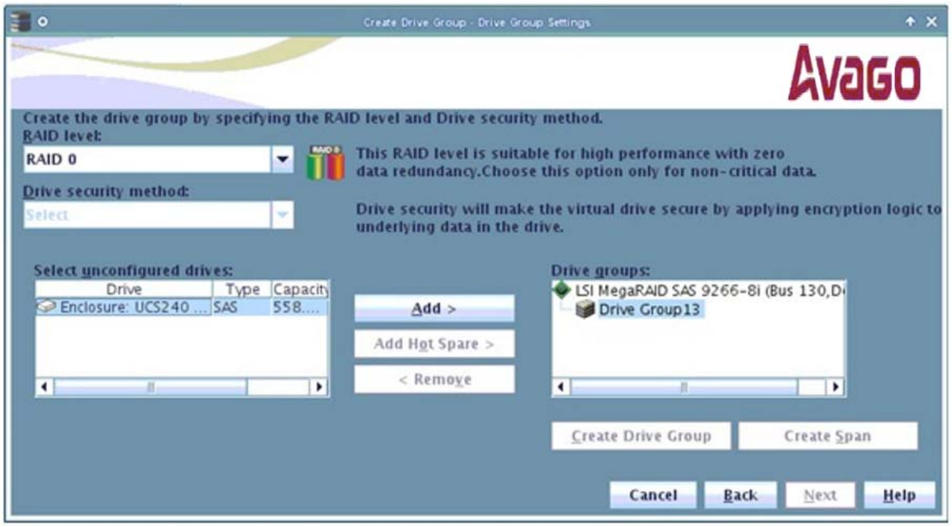
Create Drive Group - Drive Group Settings

Create the drive group by specifying the RAID level and Drive security method.

RAID level: RAID 0

This RAID level is suitable for high performance with zero data redundancy.Choose this option only for non-critical data.

Drive security method: Select

Drive security will make the virtual drive secure by applying encryption logic to underlying data in the drive.

Select unconfigured drives:

Drive	Type	Capacity
Enclosure: UCS240 ...	SAS	558....

Drive groups:

- LSI MegaRAID SAS 9266-8i (Bus 130,D
  - Drive Group13

Step 19 Click **Create Drive Group** and then click **Next**.

Figure 14-14 Create Drive Group

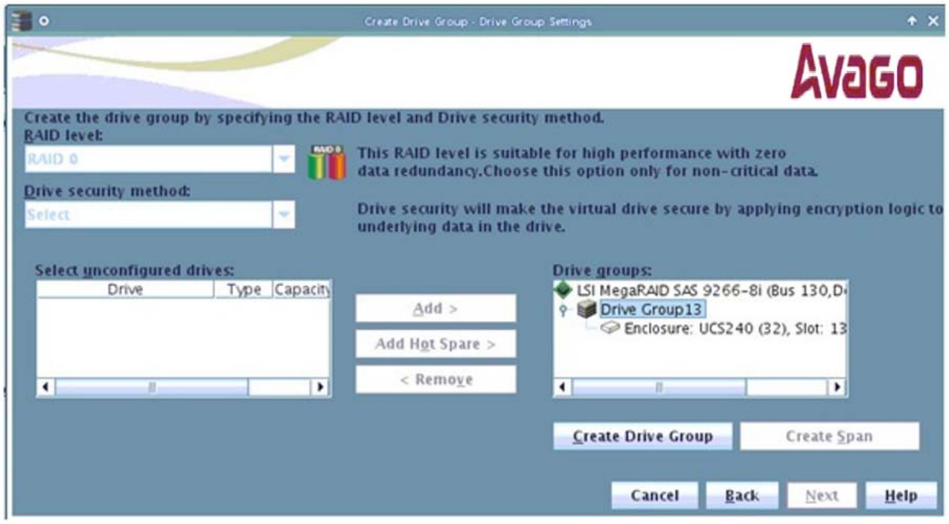
Create Drive Group - Drive Group Settings

Create the drive group by specifying the RAID level and Drive security method.

RAID level: RAID 0

This RAID level is suitable for high performance with zero data redundancy.Choose this option only for non-critical data.

Drive security method: Select

Drive security will make the virtual drive secure by applying encryption logic to underlying data in the drive.

Select unconfigured drives:

Drive	Type	Capacity

Drive groups:

- LSI MegaRAID SAS 9266-8i (Bus 130,D
  - Drive Group13
    - Enclosure: UCS240 (32), Slot: 13

Step 20 From the Parameters Configuration window, configure the following values exactly as they appear:

- Initialization State: **Fast Initialization**
- Strip Size: **256 KB**
- Read Policy: **No Read Ahead**
- Write Policy: **Always Write Back**
- I/O Policy: **Direct IO**
- Access Policy: **Read Write**
- Disk Cache Policy: **Enabled**

Step 21 Click **Create Virtual Drive**.

Figure 14-15 Virtual Drive Configuration Parameters

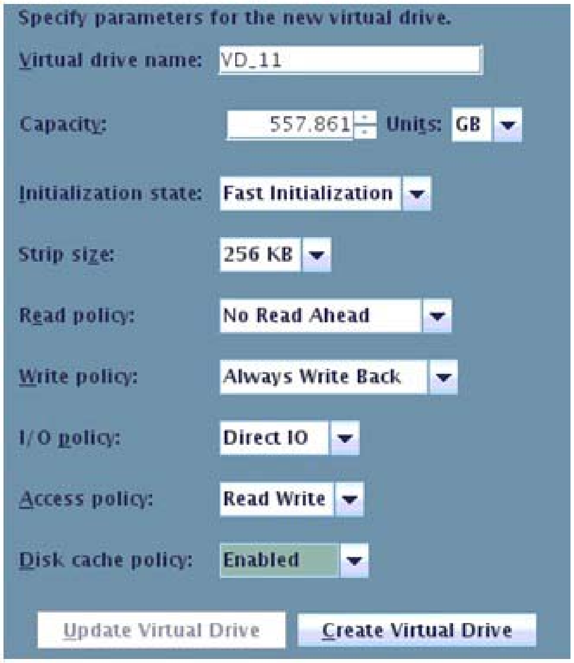

Specify parameters for the new virtual drive.

Virtual drive name: VD_11

Capacity: 557.861 Units: GB

Initialization state: Fast Initialization

Strip size: 256 KB

Read policy: No Read Ahead

Write policy: Always Write Back

I/O policy: Direct IO

Access policy: Read Write

Disk cache policy: Enabled

Step 22 When the Always Write Back Selected dialog box appears, click **Yes** and then click **Next** in the next window.

Figure 14-16 Always Write Back Selected Dialog Box

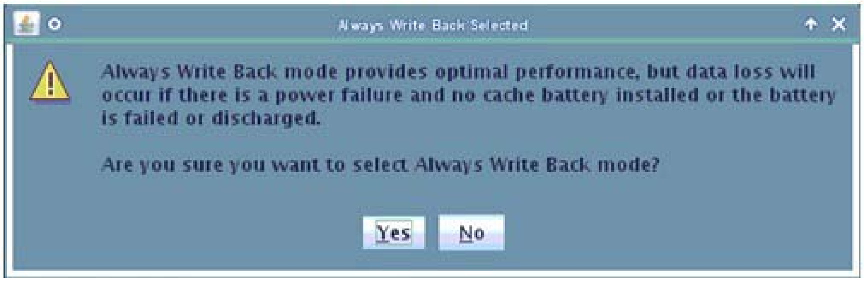

Always Write Back Selected

Always Write Back mode provides optimal performance, but data loss will occur if there is a power failure and no cache battery installed or the battery is failed or discharged.

Are you sure you want to select Always Write Back mode?

Step 23 From the Create Virtual Drive Summary window, review the settings and click **Finish** to finish creating the virtual drive if the settings are correct. If any of the settings are wrong, click **Back** to correct them.

Figure 14-17 Create Virtual Drive Summary

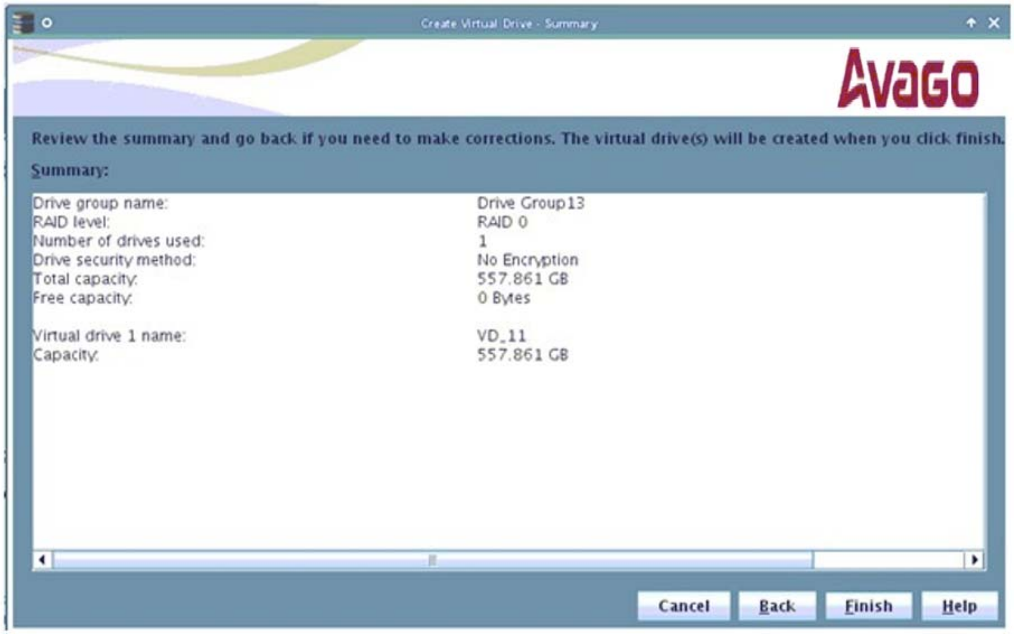

Review the summary and go back if you need to make corrections. The virtual drive(s) will be created when you click finish.

Summary:

Drive group name: Drive Group13
RAID level: RAID 0
Number of drives used: 1
Drive security method: No Encryption
Total capacity: 557.861 GB
Free capacity: 0 Bytes

Virtual drive 1 name: VD_11
Capacity: 557.861 GB

Step 24 Exit the MegaRAID Storage Manager.

Step 25 Close the VNC client.

Step 26 When not using the MegaRaid Storage Manager, it is a best practice to stop the Vivaldi framework program. From the SSH connection to the VDS TC Integrated Appliance while still logged in as root, perform the following steps to stop the Vivaldi framework program and prevent it from automatically starting on a reboot of the system:

a. Enter the command **/etc/init.d/vivaldiframeworkd stop** to stop the Vivaldi framework program.

b. To confirm that the program is stopped, enter the **/etc/init.d/vivaldiframeworkd status** command. If the program is stopped, you will see the message “Framework is stopped...”

c. Enter the command **chkconfig vivaldiframeworkd off** to make sure the Vivaldi framework program is not automatically started on a reboot of the system. To confirm that the program is disabled, enter the command **chkconfig vivaldiframeworkd**. If the program is disabled, you should see the message “vivaldiframeworkd off.”

Note If you need to run the Storage Manager, you will need to restart the Vivaldi framework program. To restart the Vivaldi framework program, enter the command **/etc/init.d/vivaldiframeworkd start** while logged into the VDS TC Integrated Appliance as root. After you are finished using Storage Manager, remember to stop the Vivaldi framework program be entering the **/etc/init.d/vivaldiframeworkd stop** command.

Step 27 From the SSH connection to the VDS TC Integrated Appliance, enter the command **cd /opt/pang/bin/** to change directories and then enter the command **./fdisk.sh** to find the missing drive. (This will represent the drive that you replaced.) For example, based on the following results from the **./fdisk.sh** command, volume #11 is the missing volume because PeerApp,0011,U does not appear:

```
ce-1:/opt/pang/bin # cd /opt/pang/bin/
ce-1:/opt/pang/bin # ./fdisk.sh
PeerApp,0001,U /dev/sdb1
```

```
PeerApp,0002,U /dev/sdc1
PeerApp,0003,U /dev/sdd1
PeerApp,0004,U /dev/sde1
PeerApp,0005,U /dev/sdf1
PeerApp,0006,U /dev/sdg1
PeerApp,0007,U /dev/sdh1
PeerApp,0008,U /dev/sdi1
PeerApp,0009,U /dev/sdj1
PeerApp,0010,U /dev/sdk1
PeerApp,0012,U /dev/sdm1
```

Step 28 Once you have determined the drive that was replaced, you need to remove the hashes for this volume from the database. Enter the **su admin** command to log into the VDS TC CLI. The CLI prompt console> appears.

Step 29 Enter the command **enable** to enter Enable mode. At the password prompt that appears, enter the enable password. The CLI prompt console# appears.

Step 30 Enter the command **cache volume remove_content**. When prompted, enter the volume number of the missing volume that you discovered in Step 27. When prompted to confirm that you want to remove all hashes from the volume, enter **Y**. For example:

```
Please enter volume number <1-11>
11
Are you sure? This will remove all hashes from volume 11.
[yes|no] no : yes
Removing all the content from volume 11...
Done.
console#
```

Step 31 Enter the **exit** command to exit Enable mode, and then enter the **exit** command to exit the VDS TC CLI.

Step 32 Next you need to find the sd device name of the newly inserted disk. Enter the command **fdisk -l | less** to determine this information. For example, based on the following output, /dev/sdl is the likely sd device name of the replaced drive because of the unpartitioned drive of a size 600gb:

```
ce-1:/opt/pang/bin # fdisk -l | less
...
Disk /dev/sdm: 598.9 GB, 598999040000 bytes
255 heads, 63 sectors/track, 72824 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes
Disk identifier: 0x00000000
Device Boot Start End Blocks Id System
/dev/sdm1 1 1 8001 83 Linux
/dev/sdm2 2 2050 16458592+ 83 Linux
/dev/sdm3 2051 2573 4200997+ 83 Linux
/dev/sdm4 2574 72824 564291157+ 83 Linux
...
Disk /dev/sdl: 598.9 GB, 598999040000 bytes
255 heads, 63 sectors/track, 72824 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes
Disk identifier: 0x00000000
...
```

Step 33 Once you have determined the sd device name of the newly installed drive, you must format the new disk using this sd device name.

To format the disk, enter the command **/usr/bin/sudo /opt/pang/useful/format_disks.sh -format_one /dev/***sd_device_name disk_replaced_number*, where *sd_device_name* is the sd device name that you discovered in Step 32 and *disk_replaced_number* is the disk number that you discovered in Step 27.

For example, to format a drive that has an sd device name of /dev/sdl and a volume number of 11, enter the command **/usr/bin/sudo /opt/pang/useful/format_disks.sh -format_one /dev/sdn 12**. The following example output will appear:

```
/usr/bin/sudo /opt/pang/useful/format_disks.sh -format_one /dev/sdn 12
Formatting /dev/sdn (disk index: 12)
format_disks.sh: 4.4 finished, see /root/installog.txt for details
ce-1:~ #
```

Step 34 To finish the process you must restart the VDS TC software and its services. Follow these steps to restart the VDS TC software and services:

a. Log into the VDS TC CLI. The CLI prompt console> appears.

b. Enter the command **enable** to enter Enable mode. At the password prompt that appears, enter the enable password. The CLI prompt console# appears.

c. Enter the command **oper service stop** to stop the VDS TC service. When prompted to confirm the restart, enter **Y**.

d. Enter the command **oper service start** to restart the VDS TC service.

CHAPTER 15

System Disk Replacement in a VDS TC C-Series Server Installation

This chapter discusses how to determine when a VDS TC C-Series Server system disk has failed and the steps that are required to replace the failed system disk.

Symptoms of a Failed VDS TC C-Series Server Failed System Disk

For the VDS TC solution, the two system disks of the VDS TC C-Series Server use a RAID 1 mirror configuration. An alarm sounds on the MegaRAID controller when a virtual drive changes from an optimal state to another state, when a hot spare rebuilds, and for test purposes.

Beep Codes

For a RAID 1 configuration, the MegaRAID controller issues the following beep codes in case of a failure:

- If one of the mirrored drives is lost, the virtual drive state will be “Degraded” and the MegaRAID controller will issue a beep code of 1 second on and 1 second off.
- If both drives are lost in a RAID 1 configuration, the state will be “Offline” and the Mega RAId controller will issue a beep code of 3 seconds on and 1 second off.

SNMP Traps

In addition, when a system disk fails, the Cisco Integrated Management Controller (CIMC) will send out SNMP traps. The following is an example of these SNMP traps in the snmpd log files:

```
Jun  6 16:49:47 sismith-u5 snmptrapd[8790]: 2014-06-06 16:49:47 10.8.16.73 [UDP:
[10.8.16.73]:39094->[10.8.16.1]]: .1.3.6.1.2.1.1.3.0 = Timeticks: (2680730)
7:26:47.30^I.1.3.6.1.6.3.1.1.4.1.0 = OID: .1.3.6.1.4.1.9.9.719.0.2^I.1.3.6.1.4.1.9.9.719.1.1.1.1.1.69 =
INTEGER: 255^I.1.3.6.1.4.1.9.9.719.1.1.1.1.11.69 = STRING: "HDD2_STATUS: Drive Slot sensor,
Drive Presence was deasserted"^I.1.3.6.1.4.1.9.9.719.1.1.1.1.4.69 = OID:
.1.3.6.1.4.1.9.9.719.1.45.2.1.1^I.1.3.6.1.4.1.9.9.719.1.1.1.1.5.69 = STRING:
"sys/rack-unit-1/board/hdd-1"^I.1.3.6.1.4.1.9.9.719.1.1.1.1.10.69 = Hex-STRING: 07 DD 06 06 17 31
29 00 ^I.1.3.6.1.4.1.9.9.719.1.1.1.1.14.69 = Hex-STRING: 07 DD 06 06 17 31 29 00
```

```
^I.1.3.6.1.4.1.9.9.719.1.1.1.1.9.69 = INTEGER: 181^I.1.3.6.1.4.1.9.9.719.1.1.1.1.22.69 = INTEGER:
5^I.1.3.6.1.4.1.9.9.719.1.1.1.1.7.69 = INTEGER: 0^I.1.3.6.1.4.1.9.9.719.1.1.1.1.20.69 = INTEGER:
0^I.1.3.6.1.4.1.9.9.719.1.1.1.1.16.69 = Counter32: 1
```

```
Jun  6 16:49:50 sismith-u5 snmptrapd[8790]: 2014-06-06 16:49:50 10.8.16.73 [UDP:
[10.8.16.73]:39094->[10.8.16.1]]: .1.3.6.1.2.1.1.3.0 = Timeticks: (2681527)
7:26:55.27^I.1.3.6.1.6.3.1.1.4.1.0 = OID: .1.3.6.1.4.1.9.9.719.0.1^I.1.3.6.1.4.1.9.9.719.1.1.1.1.1.14001
= INTEGER: 0^I.1.3.6.1.4.1.9.9.719.1.1.1.1.11.14001 = STRING: "State change on VD 00/0 from
OPTIMAL(3) to DEGRADED(2)"^I.1.3.6.1.4.1.9.9.719.1.1.1.1.4.14001 = OID:
.1.3.6.1.4.1.9.9.719.1.45.8.1.1^I.1.3.6.1.4.1.9.9.719.1.1.1.1.5.14001 = STRING:
"sys/rack-unit-1/board/storage-SAS-4/lun-"^I.1.3.6.1.4.1.9.9.719.1.1.1.1.10.14001 = STRING:
"unknown"^I.1.3.6.1.4.1.9.9.719.1.1.1.1.14.14001 = STRING:
"unknown"^I.1.3.6.1.4.1.9.9.719.1.1.1.1.9.14001 = INTEGER:
884^I.1.3.6.1.4.1.9.9.719.1.1.1.1.22.14001 = INTEGER: 5^I.1.3.6.1.4.1.9.9.719.1.1.1.1.7.14001 =
INTEGER: 0^I.1.3.6.1.4.1.9.9.719.1.1.1.1.20.14001 = INTEGER:
1^I.1.3.6.1.4.1.9.9.719.1.1.1.1.16.14001 = Counter32: 1
```

VDS TC C-Series Server System Disk Replacement Procedure

Follow this procedure to replace a failed system disk in a VDS TC C-Series Server installation:

Note We recommend following industry standard practice of using drives of the same capacity when creating RAID volumes. If drives of different capacities are used, the usable portion of the smallest drive will be used on all drives that make up the RAID volume.

Before You Begin

Step 1 Follow these steps to stop the server VDS TC application:

a. Log into the VDS TC CLI. The CLI prompt console> appears.

b. Enter the command **enable** to enter Enable mode. At the password prompt that appears, enter the enable password. The CLI prompt console# appears.

c. Enter the command **oper service stop** to stop the VDS TC caching service.

d. Wait a few minutes and then enter the command **show status**. Confirm that the Device State shows “stopped.”

Procedure

Step 1 Physically remove the failed drive. Replace a failed HDD or SSD with a drive of the same size, model, and manufacturer. If needed, refer to the service note for your server model. In general, the steps are similar for most models.

Step 2 Insert the replacement drive in the VDS TC C-Series server. Refer to the “Cisco UCS C220 Server Installation and Service Guide” or the “Cisco UCS C240 Server Installation and Service Guide” for information on this process. The replacement drive should start automatically resynchronizing.

Step 3 Wait for the mirror synchronization to complete. You can monitor the status from the Avago Configuration Utility.

Note The time to complete the synchronization can vary depending on the size of the disks in the RAID array.

Step 4 When the mirror synchronization is complete, log into the VDS TC CLI. The CLI prompt console> appears.

Step 5 Enter the command **enable** to enter Enable mode. At the password prompt that appears, enter the enable password. The CLI prompt console# appears.

Step 6 Enter the command **oper service start** to start the VDS TC caching service.

Step 7 Wait a few minutes and then enter the command **show status**.

Step 8 From the MegaRaid Storage Manager, confirm that the replaced drive shows a status of “Online”.

CHAPTER 16

System Disk Replacement in a VDS TC Blade Server Installation

This chapter discusses how to determine when a VDS TC Blade Server system disk has failed and the steps that are required to replace the failed system disk.

Symptoms of a Failed VDS TC Blade Server Failed System Disk

For the VDS TC Blade Server, the two system disks use a RAID 1 mirror configuration. An alarm sounds on the MegaRAID controller when a virtual drive changes from an optimal state to another state, when a hot spare rebuilds, and for test purposes.

Beep Codes

For a RAID 1 configuration, the MegaRAID controller issues the following beep codes in case of a failure:

- If one of the mirrored drives is lost, the virtual drive state will be “Degraded” and the MegaRAID controller will issue a beep code of 1 second on and 1 second off.
- If both drives are lost in a RAID 1 configuration, the state will be “Offline” and the Mega RAID controller will issue a beep code of 3 seconds on and 1 second off.

SNMP Traps

In addition, when a system disk fails, the Cisco Integrated Management Controller (CIMC) will send out SNMP traps. The following is an example of these SNMP traps in the snmpd log files:

```
Jun  6 16:49:47 sismith-u5 snmptrapd[8790]: 2014-06-06 16:49:47 10.8.16.73 [UDP:
[10.8.16.73]:39094->[10.8.16.1]]: .1.3.6.1.2.1.1.3.0 = Timeticks: (2680730)
7:26:47.30^I.1.3.6.1.6.3.1.1.4.1.0 = OID: .1.3.6.1.4.1.9.9.719.0.2^I.1.3.6.1.4.1.9.9.719.1.1.1.1.1.69 =
INTEGER: 255^I.1.3.6.1.4.1.9.9.719.1.1.1.1.11.69 = STRING: "HDD2_STATUS: Drive Slot sensor,
Drive Presence was deasserted"^I.1.3.6.1.4.1.9.9.719.1.1.1.1.4.69 = OID:
.1.3.6.1.4.1.9.9.719.1.45.2.1.1^I.1.3.6.1.4.1.9.9.719.1.1.1.1.5.69 = STRING:
"sys/rack-unit-1/board/hdd-1"^I.1.3.6.1.4.1.9.9.719.1.1.1.1.10.69 = Hex-STRING: 07 DD 06 06 17 31
29 00 ^I.1.3.6.1.4.1.9.9.719.1.1.1.1.14.69 = Hex-STRING: 07 DD 06 06 17 31 29 00
```

^I.1.3.6.1.4.1.9.9.719.1.1.1.1.9.69 = INTEGER: 181^I.1.3.6.1.4.1.9.9.719.1.1.1.1.22.69 = INTEGER: 5^I.1.3.6.1.4.1.9.9.719.1.1.1.1.7.69 = INTEGER: 0^I.1.3.6.1.4.1.9.9.719.1.1.1.1.20.69 = INTEGER: 0^I.1.3.6.1.4.1.9.9.719.1.1.1.1.16.69 = Counter32: 1

- Jun 6 16:49:50 sismith-u5 snmptrapd[8790]: 2014-06-06 16:49:50 10.8.16.73 [UDP: [10.8.16.73]:39094->[10.8.16.1]]: .1.3.6.1.2.1.1.3.0 = Timeticks: (2681527) 7:26:55.27^I.1.3.6.1.6.3.1.1.4.1.0 = OID: .1.3.6.1.4.1.9.9.719.0.1^I.1.3.6.1.4.1.9.9.719.1.1.1.1.1.14001 = INTEGER: 0^I.1.3.6.1.4.1.9.9.719.1.1.1.1.11.14001 = STRING: "**State change on VD 00/0 from OPTIMAL(3) to DEGRADED(2)**"^I.1.3.6.1.4.1.9.9.719.1.1.1.1.4.14001 = OID: .1.3.6.1.4.1.9.9.719.1.45.8.1.1^I.1.3.6.1.4.1.9.9.719.1.1.1.1.5.14001 = STRING: "sys/rack-unit-1/board/storage-SAS-4/lun-"^I.1.3.6.1.4.1.9.9.719.1.1.1.1.10.14001 = STRING: "unknown"^I.1.3.6.1.4.1.9.9.719.1.1.1.1.14.14001 = STRING: "unknown"^I.1.3.6.1.4.1.9.9.719.1.1.1.1.9.14001 = INTEGER: 884^I.1.3.6.1.4.1.9.9.719.1.1.1.1.22.14001 = INTEGER: 5^I.1.3.6.1.4.1.9.9.719.1.1.1.1.7.14001 = INTEGER: 0^I.1.3.6.1.4.1.9.9.719.1.1.1.1.20.14001 = INTEGER: 1^I.1.3.6.1.4.1.9.9.719.1.1.1.1.16.14001 = Counter32: 1

VDS TC Blade Server System Disk Replacement Procedure

Follow this procedure to replace a system disk in a VDS TC Blade Server installation:

Note We recommend following industry standard practice of using drives of the same capacity when creating RAID volumes. If drives of different capacities are used, the usable portion of the smallest drive will be used on all drives that make up the RAID volume.

Before You Begin

Step 1 Follow these steps to stop the server VDS TC application:

a. Log into the VDS TC CLI. The CLI prompt console> appears.

b. Enter the command **enable** to enter Enable mode. At the password prompt that appears, enter the enable password. The CLI prompt console# appears.

c. Enter the command **oper service stop** to stop the VDS TC caching service.

d. Wait a few minutes and then enter the command **show status**. Confirm that the Device State shows “stopped.”

Step 2 Replace a failed HDD or SSD with a drive of the same size, model, and manufacturer. Before changing an HDD in a running system, check the service profile in UCS Manager to make sure that the new hardware configuration is within the parameters allowed by the service profile.

Procedure

Step 1 Connect the KVM dongle to the server with the failed drive.

Step 2 Connect a monitor, keyboard, and mouse to the destination server.

Step 3 Physically replace the failed drive.If needed, refer to the service note for your server model. In general, the steps are similar for most models.

Step 4 Boot the server, using the power switch on the front of the server. If necessary, disable the quiet boot feature and boot again. For more information on disabling the quite boot feature, see the "Cisco UCS Manager B-Series Troubleshooting Guide" at http://www.cisco.com/en/US/partner/docs/unified_computing/ucs/ts/guide/UCSTroubleshooting_chapter_0111.html#t_howtodisablequietboot.

Step 5 Wait for the Avago Configuration Utility banner.

Step 6 To enter the Avago Configuration Utility, press **Ctrl-C**.

Step 7 From the SAS Adapter List screen, choose the SAS adapter used in the server. The Cisco B200 M3 servers have an Avago SAS 2004 RAID controller on the motherboard.

Step 8 Choose **RAID Properties**. The View Array screen appears.

Step 9 Choose **Manage Array**. The Manage Array screen appears.

Step 10 Choose **Activate Array**. When the activation is complete, the RAID status changes to Optimal.

Step 11 On the Manage Array screen, choose **Synchronize Array**.

Step 12 Wait for the mirror synchronization to complete and monitor the progress bar that comes up.

Note The time to complete the synchronization can vary depending on the size of the disks in the RAID array.

Step 13 When the mirror synchronization is complete, press the **ESC** key several times to go back through each of the screens (one at a time) and then exit the Avago Configuration Utility.

Step 14 Choose the reboot option to implement the changes.

Step 15 After the system has rebooted, log into the VDS TC CLI. The CLI prompt console> appears.

Step 16 Enter the command **enable** to enter Enable mode. At the password prompt that appears, enter the enable password. The CLI prompt console# appears.

Step 17 Enter the command **oper service start** to start the VDS TC caching service.

Step 18 Wait a few minutes and then enter the command **show status**.

Step 19 From the MegaRaid Storage Manager, confirm that the replaced drive shows a status of "Online".

VDS TC Troubleshooting Files

From the VDS TC Manager, you can export the system logs for a specific period of time, which will help with troubleshooting. If the system is a Cluster installation, it will export the logs from all servers within that Cluster system, including the Management Server.

Follow these steps to export the system logs:

Step 1 In the Cisco VDS TC Manager window choose **Maintenance > Logs**. The Logs window appears.

Step 2 In the From: field, enter a starting date and in the To: field enter an ending date.

Step 3 Click **Export Logs**. You will see the log file being processed.

Figure 17-1 Log File Processing

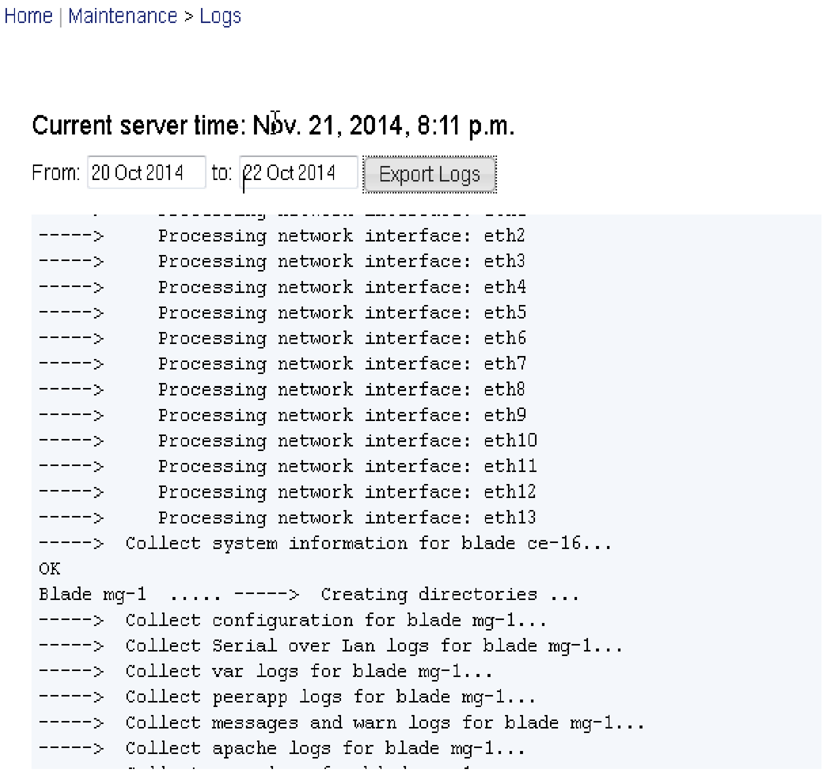

Current server time: Nov. 21, 2014, 8:11 p.m.

From: 20 Oct 2014 to: 22 Oct 2014

```
----->     Processing network interface: eth2
----->     Processing network interface: eth3
----->     Processing network interface: eth4
----->     Processing network interface: eth5
----->     Processing network interface: eth6
----->     Processing network interface: eth7
----->     Processing network interface: eth8
----->     Processing network interface: eth9
----->     Processing network interface: eth10
----->     Processing network interface: eth11
----->     Processing network interface: eth12
----->     Processing network interface: eth13
----->  Collect system information for blade ce-16...
OK
Blade mg-1  ..... ----->  Creating directories ...
----->  Collect configuration for blade mg-1...
----->  Collect Serial over Lan logs for blade mg-1...
----->  Collect var logs for blade mg-1...
----->  Collect peerapp logs for blade mg-1...
----->  Collect messages and warn logs for blade mg-1...
----->  Collect apache logs for blade mg-1...
```

Step 4 After the log file is generated, the Export Logs file popup window appears.

Figure 17-2 Export Log Dialog

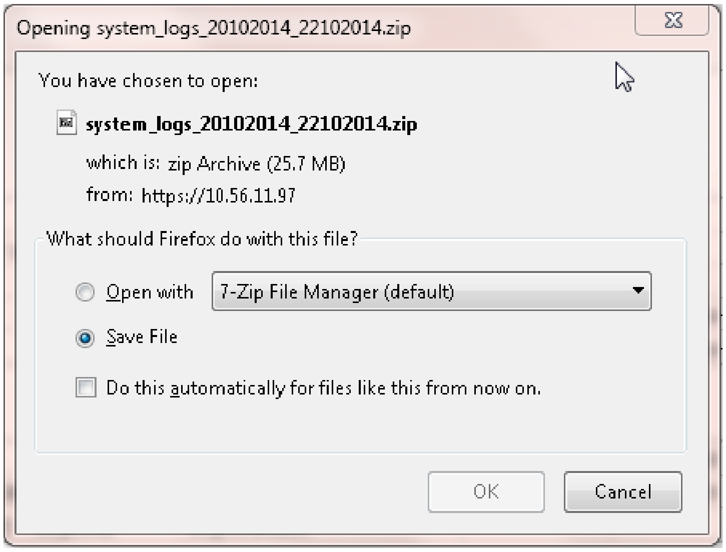

Opening system_logs_20102014_22102014.zip

You have chosen to open:

system_logs_20102014_22102014.zip

which is: zip Archive (25.7 MB)

from: https://10.56.11.97

What should Firefox do with this file?

○ Open with 7-Zip File Manager (default)

● Save File

□ Do this automatically for files like this from now on.

Step 5 Choose to **Open** or **Save** the file. If you choose to save the file, enter the location and file name in the File Save dialog box.

Step 6 Click **OK**.

Step 7 This export saves the following files:

- VDS TC Management Server files:
  - access_log
  - rcapache2.out
  - error_log
  - operpang.conf_<CE_ID>
  - pang.conf_<CE_ID>
  - cluster_conf.xml
  - device.xml
  - groups.xml
  - network.xml
  - policies.xml
  - rules.xml
  - signatures.xml
  - smartfilter.xml
  - subscriber.xml
  - timeplans.xml
  - eth0
  - eth1
  - ifconfig.log

  - peerapp_system<*timestamp*>log
  - peerapp_system_alarm.log
  - snmpd.log
  - spread_mg-1.log
  - cpu_info.log
  - dmesg.log
  - dmidecode.log
  - dstat.log
  - ipmitool.log
  - top.log
  - messages
  - messages-<*time_stamp*>.gz
  - serial_ce-1.log--<*time_stamp*>.gz
  - warn-<*time_stamp*>.gz
- VDS TC Cache Engine files:
  - pang.conf
  - iscsi.txt
  - eth#
  - ifconfig.txt
  - core.#
  - pang.log.<*time_stamp*>.gz
  - osversion.txt
  - version.txt
  - snmpd.log
  - spread_<CE-#>.log
  - cpu_info.log
  - dmesg.log
  - dmidecode.log
  - dstat.log
  - ipmitool_sel_list.txt
  - top.log
  - messages-<*time_stamp*>.gz
  - warn-<*time_stamp*>.gz

CHAPTER 18

Troubleshooting Storage Disk Problems

Recovering an Inactive Volume not Caused by a Hardware Problem

Perform the following steps to determine whether a volume has failed due to a hardware problem and recover the volume if the failure was *not* due to a hardware problem:

Step 1 In VDS TC Manager, check for alarms or errors reported for the inactive volume. To check for this information choose **Status** > **Storage** and click the **Detailed Status** tab. Check for possible drives that are having problems.

Step 2 If a disk is faulty, contact Cisco support. If the inactive disk does not show as faulty, from the VDS TC CLI in enable mode, enter the **cache volume deactivate** command. You will be prompted to enter the volume number to deactivate. For example, if the inactive volume is 29, do the following:

```
console# cache volume deactivate
Licensed volumes: 120
Please enter volume number <1-120> 29
```

Step 3 After the volume is deactivated, log into CE-1. To log into CE-1, from an SSH connection to the VDS TC management server, enter the command **ssh root@ce-1**.

Step 4 To determine the device name of the inactive volume, enter the command **/opt/pang/bin/fdisk.sh|sort**. The following is a portion of the output from this command, showing the device name for volume 29:

```
PeerApp,0029,U /dev/sdac1
```

The device name in this example is /dev/sdac.

Step 5 Enter the command **xfs_repair -v -L -r** *device_name***4** *device_name***3**, where *device_name* is the device name discovered in Step 4.

Step 6 After Step 5 completes successfully, from the VDS TC CLI enter the command **cache volume activate**. You will be prompted to enter the volume number to activate. For example, i the volume you need to activate is volume 29, do the following:

```
console# cache volume activate
Licensed volumes: 120
Please enter volume number <1-120> 29
```

Step 7 Enter the **show volumes** command and confirm that the volume no longer shows a status of “inactive”.

Identifying the Physical Location of a Volume

Follow these steps to identify the physical location of a storage volume:

Step 1 From a connection to the VDS TC management server log into CE-1 using the command **ssh root@ce-1**.

Step 2 To identify the numbering scheme of the volumes, enter the command **/opt/pang/bin/fdisk.sh script**.

Step 3 To identify the mount point for a specific volume, enter the command **/opt/pang/bin/fdisk.sh | grep** *volume_#*, where *volume_#* is the volume # that you are trying to identify. For example, enter the following command to determine the mount point for volume 102:

```
ce-1:~ # /opt/pang/bin/fdisk.sh | grep 0102
```

The output will look like the following:

```
PeerApp,0102,U     /dev/sdae1
```

In this example, the mount point is /dev/sdae1.

Step 4 To identify the storage enclosure and LUN on which this volume is mounted, enter the command /opt/mpp/lsvdev | grep /dev/*mount_point*, where *mount_point* is the mount point that you discovered in Step 3 without the integer. For our example the command would be:

```
ce-1:~ # /opt/mpp/lsvdev | grep /dev/sdae
```

The output for this command would look like the following:

```
PA-2           5     -> /dev/sdae
```

In this output, the first value is the enclosure name (PA-2 in this example) and the second value is the LUN ID (5 in this example.)

Step 5 Use a VNC client to connect from your local computer to the VDS TC management server.

Step 6 From the VNC console window, do one of the following, depending on which storage enclosure you are using:

a. If your VDS TC solution uses an IBM storage enclosure:

- Start the IBM Storage Manager by using a shortcut on the desktop or by entering the command **/opt/IBM_DS/client/SMclient** from a terminal connection.
- Right-click the storage array that you discovered in Step 4 and choose **Manage Storage Subsystem**.
- Click the **Host Mappings** tab. From the Defined Mappings table, identify the physical drive using the LUN ID that was obtained in the Step 4. For example, in the following output, the logical drive name for LUN ID 5 is 14.

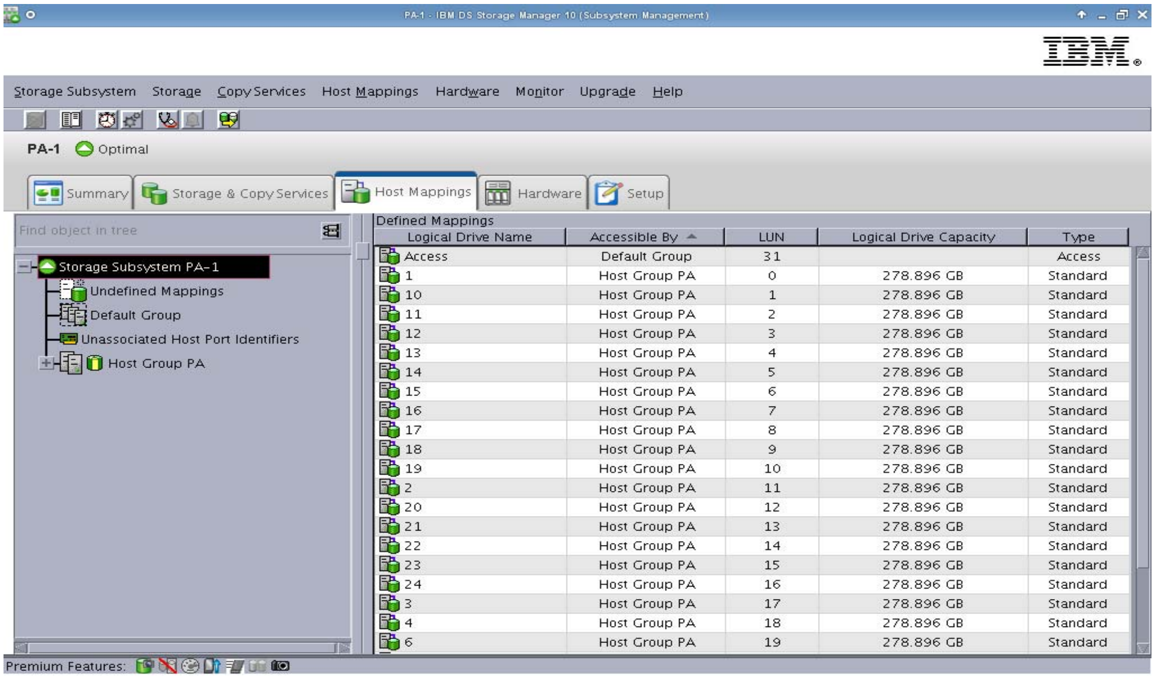

Defined Mappings

Logical Drive Name	Accessible By	LUN	Logical Drive Capacity	Type
Access	Default Group	31		Access
1	Host Group PA	0	278.896 GB	Standard
10	Host Group PA	1	278.896 GB	Standard
11	Host Group PA	2	278.896 GB	Standard
12	Host Group PA	3	278.896 GB	Standard
13	Host Group PA	4	278.896 GB	Standard
14	Host Group PA	5	278.896 GB	Standard
15	Host Group PA	6	278.896 GB	Standard
16	Host Group PA	7	278.896 GB	Standard
17	Host Group PA	8	278.896 GB	Standard
18	Host Group PA	9	278.896 GB	Standard
19	Host Group PA	10	278.896 GB	Standard
2	Host Group PA	11	278.896 GB	Standard
20	Host Group PA	12	278.896 GB	Standard
21	Host Group PA	13	278.896 GB	Standard
22	Host Group PA	14	278.896 GB	Standard
23	Host Group PA	15	278.896 GB	Standard
24	Host Group PA	16	278.896 GB	Standard
3	Host Group PA	17	278.896 GB	Standard
4	Host Group PA	18	278.896 GB	Standard
6	Host Group PA	19	278.896 GB	Standard

a. If your VDS TC solution uses a NetApp storage enclosure:

- Start the SANtricity Storage Manager by using a shortcut on the desktop or by entering the command **/opt/SMgr/client/SMclient** from a terminal connection.
- Right-click the storage array that you discovered in Step 4 and choose **Manage Storage Array**.
- Click the **Host Mappings** tab. From the Defined Mappings table, identify the physical drive using the LUN ID that was obtained in the Step 4. For example, in the following output, the logical drive name for LUN ID 5 is 6.

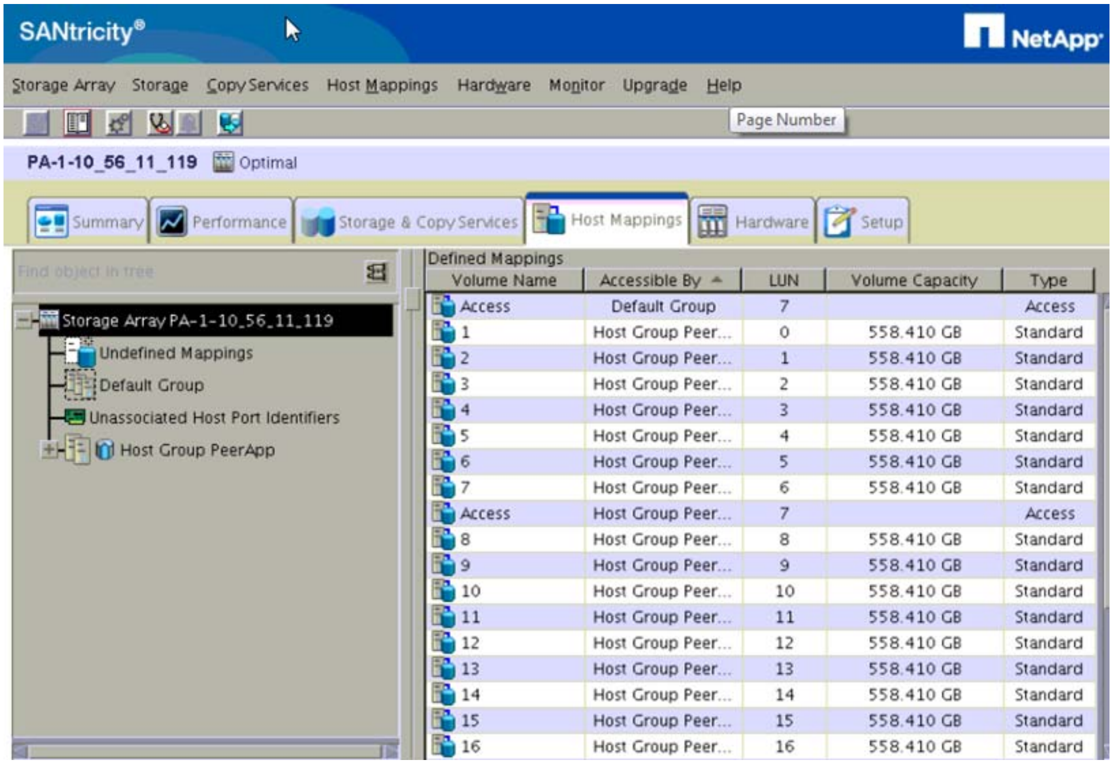

Defined Mappings

Volume Name	Accessible By	LUN	Volume Capacity	Type
Access	Default Group	7		Access
1	Host Group Peer...	0	558.410 GB	Standard
2	Host Group Peer...	1	558.410 GB	Standard
3	Host Group Peer...	2	558.410 GB	Standard
4	Host Group Peer...	3	558.410 GB	Standard
5	Host Group Peer...	4	558.410 GB	Standard
6	Host Group Peer...	5	558.410 GB	Standard
7	Host Group Peer...	6	558.410 GB	Standard
Access	Host Group Peer...	7		Access
8	Host Group Peer...	8	558.410 GB	Standard
9	Host Group Peer...	9	558.410 GB	Standard
10	Host Group Peer...	10	558.410 GB	Standard
11	Host Group Peer...	11	558.410 GB	Standard
12	Host Group Peer...	12	558.410 GB	Standard
13	Host Group Peer...	13	558.410 GB	Standard
14	Host Group Peer...	14	558.410 GB	Standard
15	Host Group Peer...	15	558.410 GB	Standard
16	Host Group Peer...	16	558.410 GB	Standard

CHAPTER 19

Troubleshooting Network Connectivity Issues

VDS TC network connectivity problems can be caused by a variety of issues. This chapter covers some of the key VDS TC specific troubleshooting steps you should perform when encountering network connectivity issues from the VDS TC installation.

Step 1 Confirm that the status of the cache engine Ethernet interfaces is "UP":

- From the CLI of the VDS TC system, enter the **show eth_status** in either Regular mode or Enable mode. The status of all of the interfaces should be "UP", as shown in the following example:

```
console# show eth_status
Blade  eth0 eth1   eth2  eth3  eth4  eth5  eth6  eth7   eth8   eth9 eth10 eth11  eth12
ce-1    UP    UP     UP    UP    UP   UP    UP    UP     UP    UP      UP     UP    UP
ce-2    UP    UP     UP    UP    UP   UP    UP    UP     UP    UP      UP     UP    UP
ce-3    UP    UP     UP    UP    UP   UP    UP    UP     UP    UP      UP     UP    UP
ce-4    UP    UP     UP    UP    UP   UP    UP    UP     UP    UP      UP     UP    UP
ce-5    UP    UP     UP    UP    UP   UP    UP    UP     UP    UP      UP     UP    UP
ce-6    UP    UP     UP    UP    UP   UP    UP    UP     UP    UP      UP     UP    UP
ce-7    UP    UP     UP    UP    UP   UP    UP    UP     UP    UP      UP     UP    UP
ce-8    UP    UP     UP    UP    UP   UP    UP    UP     UP    UP      UP     UP    UP
ce-9    UP    UP     UP    UP    UP   UP    UP    UP     UP    UP      UP     UP    UP
ce-10   UP    UP     UP    UP    UP   UP    UP    UP     UP    UP      UP     UP    UP
ce-11   UP    UP     UP    UP    UP   UP    UP    UP     UP    UP      UP     UP    UP
ce-12   UP    UP     UP    UP    UP   UP    UP    UP     UP    UP      UP     UP    UP
ce-13   UP    UP     UP    UP    UP   UP    UP    UP     UP    UP      UP     UP    UP
ce-14   UP    UP     UP    UP    UP   UP    UP    UP     UP    UP      UP     UP    UP
ce-15   UP    UP     UP    UP    UP   UP    UP    UP     UP    UP      UP     UP    UP
ce-16    UP    UP     UP    UP    UP    UP    UP    UP     UP     UP      UP     UP     UP
```

For any interfaces that are not up, continue check for possible reasons, such as cabling issues.

Step 2 Confirm that the spread communication among cluster members is set to broadcast:

- From the CLI of the VDS TC system, enter the command **show cluster-bus-ip**.

  ```
  console# show cluster-bus-ip
  Processing...
  MG-1: [Broadcast]
  CE-1: [Broadcast]
  CE-2: [Broadcast]
  CE-3: [Broadcast]
  CE-4: [Broadcast]
  CE-5: [Broadcast]
  CE-6: [Broadcast]
  CE-7: [Broadcast]
  CE-8: [Broadcast]
  ```

```
OK. All the CEs are configured to work with Broadcast IP.
```

- If the communication is set to multicast instead of broadcast, enter the command **cluster-bus-ip broadcast**. For example:

```
configuration# cluster-bus-ip broadcast
Processing...
All the CEs are configured to work with Broadcast IP.
```

Step 3 Confirm the correct speed settings of the Ethernet interfaces on the cache engines:

- On each cache engine, while logged in as root, enter the **ethtool** *interface_name* command. For example, to check the speed setting of the eth1 interface on C#-1, enter the following command:

```
ce-1:~ # ethtool eth1
Settings for eth1:
        Supported ports: [ TP ]
        Supported link modes:   10baseT/Half 10baseT/Full
                                100baseT/Half 100baseT/Full
                                1000baseT/Full
        Supports auto-negotiation: Yes
        Advertised link modes:  10baseT/Half 10baseT/Full
                                100baseT/Half 100baseT/Full
                                1000baseT/Full
        Advertised auto-negotiation: Yes
        Speed: 1000Mb/s
        Duplex: Full
        Port: Twisted Pair
        PHYAD: 1
        Transceiver: internal
        Auto-negotiation: on
        Supports Wake-on: pumbg
        Wake-on: g
        Current message level: 0x00000007 (7)
        Link detected: yes
```

Note If an interface shows a speed of 100Mb and auto-negotiated is **on**, you should replace the cable.

Step 4 Confirm that the following interfaces are configured with an MTU of 9K.

- If you are using a Cisco Nexus 7000 Series switch, use the **show interface** command to confirm that all of the following ports are configured with an MTU of 9216 bytes.
  - All ports that connect to the VDS TC cache engines.
  - All ports that connect to iSCSI ports of the storage devices
  - All ports that connect to Fabric Extender (FEX) ports
  - All ports that connect to routers, including the physical ports and the port channel

Note See Appendix A “Cisco VDS TC Blade Server Cabling Guides” in the “Cisco Videoscape Distribution Suite Transparent Caching Software Blade Server Cluster Installation Guide” for specific port information.

- If you are using a Cisco Catalyst 4500X switch, use the **show interface** command to confirm that all of the following ports are configured with an MTU of 9216 bytes.
  - All ports that connect to the VDS TC cache engines

  - All ports between the Cisco Catalyst 4500X switches, including both physical ports and port channels

Note See Appendix A “Cisco VDS TC C-Series Cabling Guides” in the “Cisco Videoscape Distribution Suite Transparent Caching Software C-Series Cluster Installation Guide” for specific port information.

- If you are using a Cisco Catalyst 4948E switch, use the **show interface** command to confirm that all of the following ports are configured with an MTU of 9216 bytes.
  - All ports that connect to iSCSI ports of the storage devices
  - All ports that connect to the Cisco Catalyst 4500X switches, including both physical ports and port channels

Note See Appendix A “Cisco VDS TC C-Series Cabling Guides” in the “Cisco Videoscape Distribution Suite Transparent Caching Software C-Series Cluster Installation Guide” for specific port information.

- If you have a VDS TC Blade Server installation, use the **ifconfig** command to confirm that all vNIC interfaces on all of the cache engines (eth0 through eth12) have an MTU of 9000 bytes.

```
ce-1:~ # ifconfig
eth0      Link encap:Ethernet  HWaddr 00:25:B5:00:01:FF
          inet addr:10.11.12.2  Bcast:10.11.12.255  Mask:255.255.255.0
          UP BROADCAST RUNNING MULTICAST  MTU:9000  Metric:1
          RX packets:2139321620 errors:0 dropped:0 overruns:0 frame:0
          TX packets:799506045 errors:0 dropped:1 overruns:0 carrier:0
          collisions:0 txqueuelen:1000
          RX bytes:1244766315958 (1187101.6 Mb)  TX bytes:220300148495 (210094.5 Mb)

eth1      Link encap:Ethernet  HWaddr 00:25:B5:00:02:8F
          inet addr:10.11.80.2  Bcast:10.11.80.255  Mask:255.255.255.0
          UP BROADCAST RUNNING MULTICAST  MTU:9000  Metric:1
          RX packets:4324769 errors:0 dropped:0 overruns:0 frame:0
          TX packets:45 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:1000
          RX bytes:319167124 (304.3 Mb)  TX bytes:2880 (2.8 Kb)

eth2      Link encap:Ethernet  HWaddr 00:25:B5:00:03:FF
          UP BROADCAST RUNNING MULTICAST  MTU:9000  Metric:1
          RX packets:4324804 errors:0 dropped:0 overruns:0 frame:0
          TX packets:0 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:1000
          RX bytes:319169322 (304.3 Mb)  TX bytes:0 (0.0 b)
```

Output omitted

- If you have a VDS TC C-Series Server installation, use the **ifconfig** command to confirm that all vNIC interfaces on all of the cache engines (eth2 through eth13) have an MTU of 9000 bytes.

```
eth2   Link encap:Ethernet  HWaddr D4:8C:B5:BD:14:1C
       UP BROADCAST RUNNING MULTICAST  MTU:9000  Metric:1
       RX packets:18780767 errors:176 dropped:65528 overruns:0 frame:151
       TX packets:58782317423 errors:0 dropped:0 overruns:0 carrier:0
       collisions:0 txqueuelen:16384
       RX bytes:1648108179 (1571.7 Mb)  TX bytes:83835341083081 (79951611.5 Mb)

eth3   Link encap:Ethernet  HWaddr D4:8C:B5:BD:14:1D
       UP BROADCAST RUNNING MULTICAST  MTU:9000  Metric:1
```

```
          RX packets:84177343726 errors:4842 dropped:200836 overruns:0 frame:4841
          TX packets:14467531 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:16384
          RX bytes:269293757793518 (256818540.3 Mb)  TX bytes:20330131810 (19388.3 Mb)

eth4      Link encap:Ethernet  HWaddr D4:8C:B5:BD:14:1E
          UP BROADCAST RUNNING SLAVE MULTICAST  MTU:9000  Metric:1
          RX packets:21163343784 errors:9633216 dropped:0 overruns:0 frame:9629902
          TX packets:12348910809 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:1000
          RX bytes:82848395163084 (79010386.6 Mb)  TX bytes:4965978640262 (4735926.2 Mb)
```

Output omitted

- All of the iSCSI host ports on the storage devices should have an MTU of 9000 bytes. Follow these steps to use the IBM or NetApp Storage Manager software to confirm this:

Step 1 Using SSH software, such as Putty, open an SSH connection to the IP address that was assigned to the VDS TC Management Server.

Step 2 Log into the system using the username **padmin** and the password provided by Cisco.

Step 3 Enter the command **sudo -l** to change to superuser. Enter the password for the root user when prompted. This password is provided by Cisco.

Step 4 Enter the command **vncserver :1** to start the VNC Server.

Step 5 Use a VNC viewer to connect from your local computer to the VDS TC Management Server, using the management IP address of the VDS TC Management Server. If you are prompted to log in, enter the root user with the password that was provided by Cisco.

Step 6 Enter one of the following commands from the VNC console, based on the type of storage enclosure you are using:

- If you are using a NetApp storage enclosure, in the xterm window enter the command **`/opt/SMgr/client/SMclient`**.
- If you are using an IBM storage enclosure, in the xterm window enter the command **/opt/IBM_DS/client/SMclient**.

Step 7 Choose the first Storage Subsystem if using an IBM storage enclosure or the first Storage Array if using a NetApp storage enclosure:

a. If using an IBM storage enclosure, right-click the storage subsystem and choose **Manage Storage Subsystem**. If using a NetApp storage enclosure, right-click the storage array and choose **Manage Storage Array**.

b. Click the **Hardware** tab.

c. From the menu bar, choose **Hardware** > **Controller** > **Configure** > **iSCSI Ports**.

d. From the iSCSI port drop-down list choose the interface to check.

e. Click **Advanced Port Settings**.

f. In the window that appears, ensure that the Enable jumbo frames check box is checked and that the MTU size is set to **9000**.

g. Click **Cancel**.

h. Repeat Steps d through g for each interface.

Step 8 Repeat Step 7 for each remaining Storage Subsystems or Storage Array.

Step 9 Click **Cancel**.

Step 10 From the Storage Subsystem menu, choose **Exit**.

Step 11 Close the IBM System Storage™ DS Storage Manager.

Step 4 Use the **lspci | grep Eth** command to confirm that only supported NICs are installed in the VDS TC management server and cache engines. For a list of supported NICs, contact your Cisco representative.

```
ce-1:~ # lspci | grep Eth
01:00.0 Ethernet controller: Intel Corporation Device 1521 (rev 01)
01:00.1 Ethernet controller: Intel Corporation Device 1521 (rev 01)
0b:00.0 Ethernet controller: Cisco Systems Inc 10G Ethernet NIC (rev a2)
0c:00.0 Ethernet controller: Cisco Systems Inc 10G Ethernet NIC (rev a2)
0f:00.0 Ethernet controller: Cisco Systems Inc 10G Ethernet NIC (rev a2)
10:00.0 Ethernet controller: Cisco Systems Inc 10G Ethernet NIC (rev a2)
11:00.0 Ethernet controller: Cisco Systems Inc 10G Ethernet NIC (rev a2)
12:00.0 Ethernet controller: Cisco Systems Inc 10G Ethernet NIC (rev a2)
13:00.0 Ethernet controller: Cisco Systems Inc 10G Ethernet NIC (rev a2)
14:00.0 Ethernet controller: Cisco Systems Inc 10G Ethernet NIC (rev a2)
15:00.0 Ethernet controller: Cisco Systems Inc 10G Ethernet NIC (rev a2)
16:00.0 Ethernet controller: Cisco Systems Inc 10G Ethernet NIC (rev a2)
17:00.0 Ethernet controller: Cisco Systems Inc 10G Ethernet NIC (rev a2)
18:00.0 Ethernet controller: Cisco Systems Inc 10G Ethernet NIC (rev a2)
You have new mail in /var/mail/root
```