Maintenance Guide

Version 5.7

Web: www.peerapp.com

PeerApp Inc.

29 Crafts Street
Newton, MA. 02458, USA

Support@peerapp.com

Copyright ©2015 by PeerApp
If you require any assistance, please contact us:
PeerApp Headquarters
PeerApp Inc.
29 Crafts Street
Newton, MA 02458, USA

http://www.peerapp.com
support@peerapp.com

[Document version: 5.7 - 1.07]

Table of Contents

Chapter 1: About this Guide 1

Intended Audience 1

Scope of Guide 1

Important and Notes 1

Chapter 2: Cache Engine Replacement Procedure 2

Overview 2

Replacing a Cache Engine 2

Configuring the Storages 4

Chapter 3: Cache Engine Addition Procedure 8

Overview 8

Adding a Cache Engine 8

Configuring the Storages 10

Chapter 4: External Storage Replacement Procedure 16

Overview 16

Replacing an External Storage 16

Removing the Faulty Storage Volumes 18

Removing the Faulty Storage from the Storage Manager 20

Connecting a New Storage to an Existing System 23

Configure the New Storage 28

Chapter 5: Replacing the Management Server Procedure 29

Overview 29

Prerequisites 29

Backup Files 29

Replacing the Management Server 29

Installing the Management Server 30

Post Installation 31

Chapter 6: Connecting an External Storage to an Existing System 32

Overview 32

Connecting a Server Storage 32

Configure the New Storage 36

Chapter 7: External Storage Disk Replacement Procedure 37
Overview 37
Replacing an External Storage Disk 37

Chapter 8: Hot Swap HDD Replacement for DAS 49
Create the New Virtual Disk (LUN) 49

Chapter 9: UltaBand Clear Cache Content 58
Overview 58
Grid 58
Standalone 58

Chapter 10: Upgrade Procedure from 5.6 to 5.7 for UB5000 60
Prerequisites 60
Upgrading the UltraBand Software 60
Post Upgrade 64

Chapter 11: The Rollback Procedure from 5.7 to 5.6 for UB5000 66
The UltraBand Rollback Procedure 66

Chapter 12: Manual UltraBand Rollback Procedure from 5.7 to 5.6 for UB5000 67

Chapter 13: UBWeb Installation 69
Prerequisites 69
Pre-Installation Verification 69
Installation 69

Chapter 14: Upgrade Procedure to 5.7 for UB6000 and above 77
Prerequisites 77
Upgrading the Firmware, OS and UltraBand Software 77
Downloading the Files 77
Grid Upgrade 78
Standalone/Grid-Ready Upgrade 79

Chapter 1: About this Guide

The PeerApp Maintenance Guide is designed to provide specific information to assist the user in remedying various issues described in the forthcoming chapters.

Intended Audience

This guide is intended for service personnel responsible for assisting in the replacement or addition of vendors software and/or hardware.

Scope of Guide

Chapter 1: About this Guide

Chapter 2: Cache Engine Replacement Procedure

Chapter 3: Cache Engine Addition Procedure

Chapter 4: External Storage Replacement Procedure

Chapter 5: Replacing the Management Server Procedure

Chapter 6: Connecting an External Storage to an Existing System

Chapter 7: External Storage Disk Replacement Procedure

Chapter 8: Hot Swap HDD Replacement for DAS

Chapter 9: UltraBand Clear Caching Content

Chapter 10: Upgrade Procedure from 5.6 to 5.7 for UB5000

Chapter 11: The Rollback Procedure from 5.7 to 5.6 for UB5000

Chapter 12: Manual UltraBand Rollback Procedure from 5.7 to 5.6 for UB5000

Chapter 13: UBWeb Installation

Chapter 14: Upgrade Procedure to 5.7 for UB6000 and above

Important and Notes

IMPORTANT! This provides important information about a procedure, which if not followed may result in faults in the system.

Note: A note provides additional information about the current topic.

Chapter 2: Cache Engine Replacement Procedure

Overview

To determine whether a cache engine has failed in a Server installation, log into the UBView. By default, the Logical Status tab is displayed. If you are already logged into UBView, click the **Logical Status** tab.

If a cache engine has failed it will appear with a gray box surrounding it, as show with the bottom cache engine in the following figure.

Replacing a Cache Engine

To Replace the Cache Engine:

1. Connect a cache engine machine to the grid and install it by using **"Cache Engine Servers Installation"** from the installation guide.

   Note: While installing, when it asks about the Cache Engine number, enter number of the CE that should be replaced.

2. Update the grid ssh keys by running the replace servers key command from the management:

   cd /opt/pang/useful/
   ./replace_server_keys.sh -servers <number of blades in the grid>

   For example in a grid with 16 cache engines run:
   mg-1:~ # cd /opt/pang/useful/
   mg-1:/opt/pang/useful #
   mg-1:/opt/pang/useful # ./replace_server_keys.sh -servers 16
   Version 3.1
   mg-1:/opt/pang/useful #

3. Connecting the new cache engine to the storage.

   On the management use the command:

   ```
   cd /opt/pang/useful/configure_storages
   ```

   Run the configure storage command:

   ```
   ./configure_storages.py –s < number of storage enclosures> -b < number of blades in the grid> -c <the replaced CE-ID number>
   ```

   Note: If an error message appears, ignore it.

   A warning message about the SSD format will be shown, to continue the storages configuration, enter y.

********** WARNING! **********

SSD drives of selected blades will be formatted.

Do you want to continue? (y/n) y

Connect to each CE and check the connections by using:

ssh ce-<cache engine number>

iscsiadm -m session

The expected output per product is:

1G Storages:

- 8 * sessions per storage, for example for 1 storage:

  tcp: [17] 10.11.14.100:3260,1 iqn.1992-01.com.lsi:2365.60080e50002f4e82000000005063c74a

  tcp: [18] 10.11.14.101:3260,2 iqn.1992-01.com.lsi:2365.60080e50002f4e82000000005063c74a

  tcp: [19] 10.11.16.100:3260,1 iqn.1992-01.com.lsi:2365.60080e50002f4e82000000005063c74a

  tcp: [20] 10.11.16.101:3260,2 iqn.1992-01.com.lsi:2365.60080e50002f4e82000000005063c74a

  tcp: [21] 10.11.15.100:3260,1 iqn.1992-01.com.lsi:2365.60080e50002f4e82000000005063c74a

  tcp: [22] 10.11.15.101:3260,2 iqn.1992-01.com.lsi:2365.60080e50002f4e82000000005063c74a

  tcp: [23] 10.11.17.101:3260,2 iqn.1992-01.com.lsi:2365.60080e50002f4e82000000005063c74a

  tcp: [24] 10.11.17.100:3260,1 iqn.1992-01.com.lsi:2365.60080e50002f4e82000000005063c74a

10G Storages:

- 4 * sessions per storage, for example for 1 storage:

  tcp: [1] 10.11.14.100:3260,1 iqn.1984-05.com.dell:powervault.md3600i.6d4ae520009493cd000000004fc7eac9

  tcp: [2] 10.11.14.101:3260,2 iqn.1984-05.com.dell:powervault.md3600i.6d4ae520009493cd000000004fc7eac9

  tcp: [3] 10.11.15.100:3260,1 iqn.1984-05.com.dell:powervault.md3600i.6d4ae520009493cd000000004fc7eac9

  tcp: [4] 10.11.15.101:3260,2 iqn.1984-05.com.dell:powervault.md3600i.6d4ae520009493cd000000004fc7eac9

Configuring the Storages

1. From the management server, start the VNCServer by using the following commands:

   rm /tmp/.X1-lock
   rm /tmp/.X11-unix/X1
   vncserver :1

2. Connect to the server with a VNC Viewer

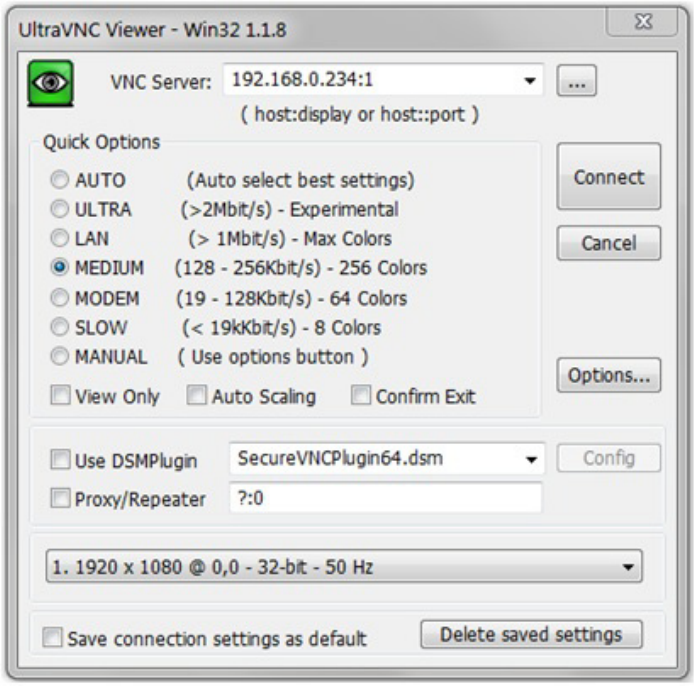

UltraVNC Viewer

3. Run the MD Storage Manager using the command:

   SMclient

4. For each storage enclosure in the Grid,

   go to the "*Host Mappings"* tab and right-click on the replaced cache engine (Host CE-x) and click "*Manage Host Port Identifiers*"

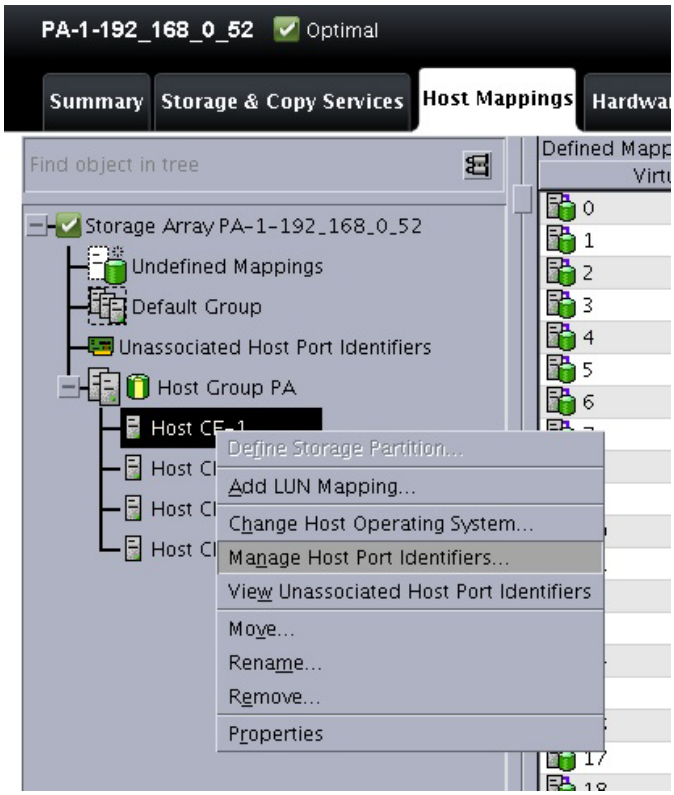

Manage Host Port Identifiers

5 From the Host Port Identifiers,

- Click on **old iqn**
- Click **Replace**

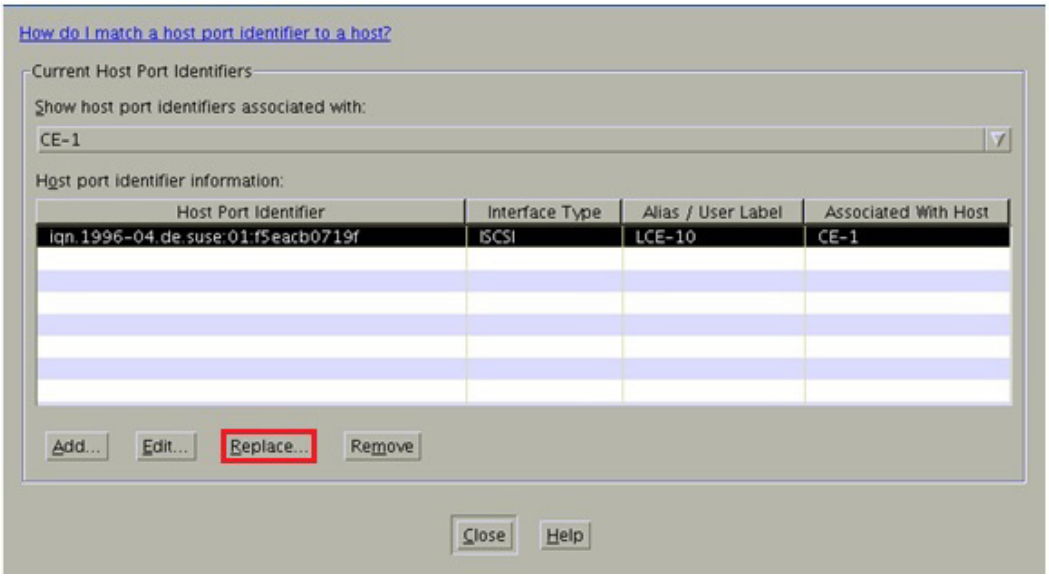

Old iqn

6 From the Host Port Identifier

- Choose **ISCSI** (where applicable)
- Select **Replace by Selecting Known**
- Choose the new *server iqn* from the list

■ Click **Replace**

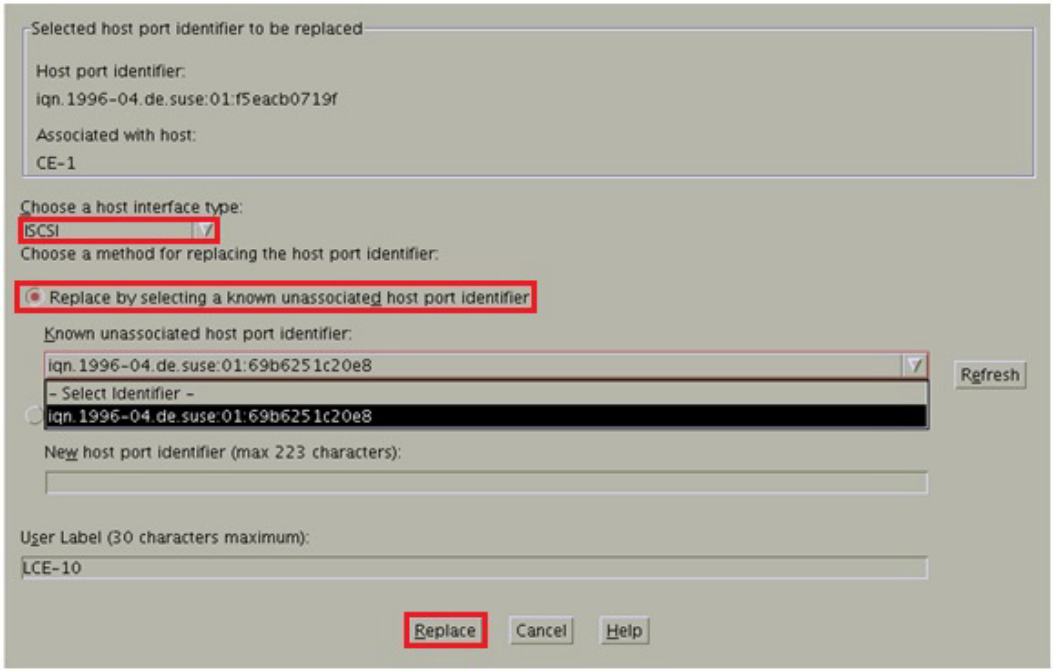

Host Port Identifiers

7 From the Current Host Port Identifiers

■ Click **Close**

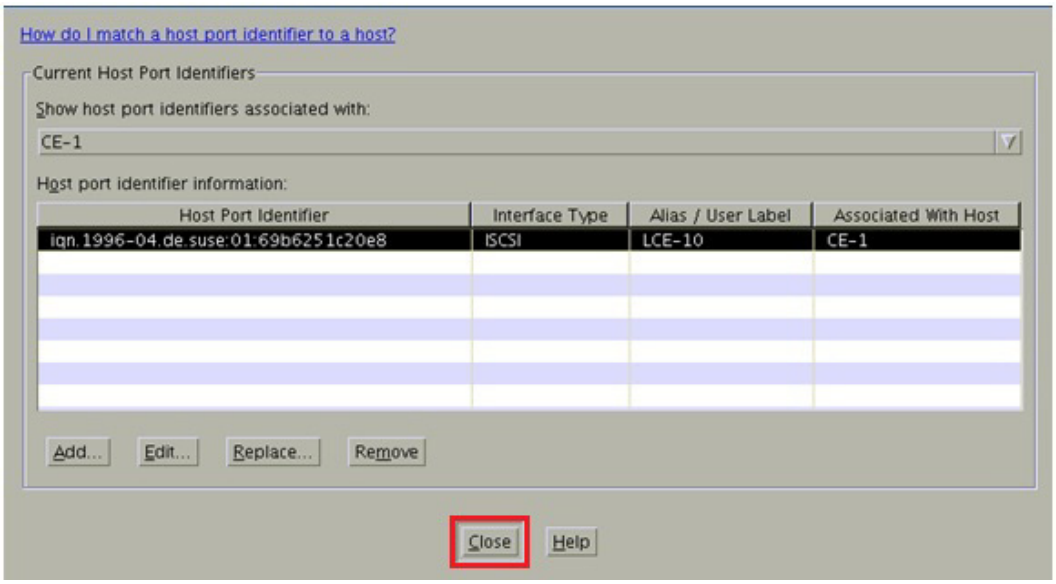

Current Host Port Identifiers

Note: Repeat Step 4-Step 7 for all the storage enclosures.

8 On the management use the command:

cd /opt/pang/useful/configure_storages

Run the configure storage command:

./configure_storages.py -s < number of storage enclosures> -b < number of blades in the grid> -c <the replaced CE-ID number>

9 Software installation:

Copy the UltraBand Software Package to /tftpboot/

Note: The upgrade file name should be in this format (even version number):
`PeerApp_GA_5.7.0bXX_grid.tar.gz`

10 In the management enter to the CLI in enable mode.

- Upgrade the replaced CE by using:
  console# upgrade server <CE number> 127.0.0.1 <software_name>

11 Activate a new license what contains the new CE system ID. Activation license guide is in section "*Installing and Configuring the Software*" in the installation guide (under "post installation steps".

12 Start the application on the replaced CE. Example on CE-2:

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

Procedure completed.

Chapter 3: Cache Engine Addition Procedure

Overview

The following chapter details the procedure for adding a Cache Engine.

Adding a Cache Engine

To Add a Cache Engine:

1. Connect all new cache engine machines to the grid and install them by using **"Cache Engine Servers Installation"** from the installation guide.

   Note: While installing, when it asks about the Cache Engine number, enter number of the CE that should be added.

2. After all the new cache engines have been installed, update the grid ssh keys by running the replace servers key command from the management:

   cd /opt/pang/useful/
   ./replace_server_keys.sh -servers <number of blades in the grid>

   For example in a grid that is expanding to 16 cache engines run the following:
   mg-1:~ # cd /opt/pang/useful/
   mg-1:/opt/pang/useful #
   mg-1:/opt/pang/useful # ./replace_server_keys.sh -servers 16
   Version 3.1
   mg-1:/opt/pang/useful #

3. Connecting the new cache engines to the storage -

   On the management use the command:

   cd /opt/pang/useful/configure_storages

   Run the configure storage command:

   ./configure_storages.py -s < number of storage enclosures> -b < number of blades in the grid> -c <the additional CE-ID number>

   Note: If an error message appears, ignore it.

   A warning message about the SSD format will be shown, to continue the storages configuration, enter y.

   ********** WARNING! **********

   SSD drives of selected blades will be formatted.

   Do you want to continue? (y/n) y

Connect to each CE and check the connections by using:

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

The expected output per product is:

1G Storages:

- 8 * sessions per storage, for example for 1 storage:

  tcp: [17] 10.11.14.100:3260,1 iqn.1992-01.com.lsi:2365.60080e50002f4e82000000005063c74a

  tcp: [18] 10.11.14.101:3260,2 iqn.1992-01.com.lsi:2365.60080e50002f4e82000000005063c74a

  tcp: [19] 10.11.16.100:3260,1 iqn.1992-01.com.lsi:2365.60080e50002f4e82000000005063c74a

  tcp: [20] 10.11.16.101:3260,2 iqn.1992-01.com.lsi:2365.60080e50002f4e82000000005063c74a

  tcp: [21] 10.11.15.100:3260,1 iqn.1992-01.com.lsi:2365.60080e50002f4e82000000005063c74a

  tcp: [22] 10.11.15.101:3260,2 iqn.1992-01.com.lsi:2365.60080e50002f4e82000000005063c74a

  tcp: [23] 10.11.17.101:3260,2 iqn.1992-01.com.lsi:2365.60080e50002f4e82000000005063c74a

  tcp: [24] 10.11.17.100:3260,1 iqn.1992-01.com.lsi:2365.60080e50002f4e82000000005063c74a

10G Storages:

- 4 * sessions per storage, for example for 1 storage:

  tcp: [1] 10.11.14.100:3260,1 iqn.1984-05.com.dell:powervault.md3600i.6d4ae520009493cd000000004fc7eac9

  tcp: [2] 10.11.14.101:3260,2 iqn.1984-05.com.dell:powervault.md3600i.6d4ae520009493cd000000004fc7eac9

  tcp: [3] 10.11.15.100:3260,1 iqn.1984-05.com.dell:powervault.md3600i.6d4ae520009493cd000000004fc7eac9

  tcp: [4] 10.11.15.101:3260,2 iqn.1984-05.com.dell:powervault.md3600i.6d4ae520009493cd000000004fc7eac9

Note: The above script should be executed N times, where N stands for the number of new cache-engines added to the grid. For example, when adding two new cache engines to grid with two CE, execute the command twice, once with “-b 3 -c 3" and afterwards with “-b 4 -c 4".

Configuring the Storages

1 From the management server, start the VNCServer by using the following commands

rm /tmp/.X1-lock
rm /tmp/.X11-unix/X1
vncserver :1

2 Connect to the server with a VNC Viewer.

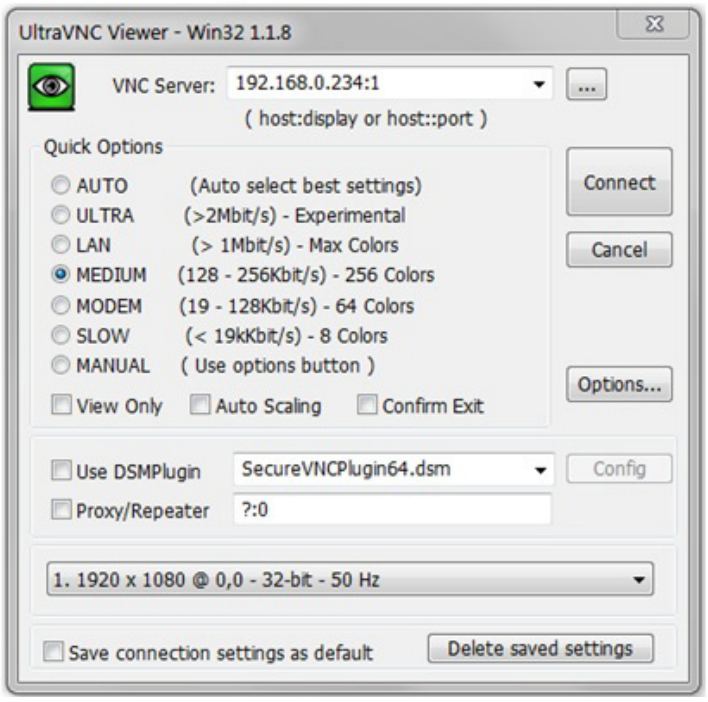

UltraVNC Viewer

3 Run the MD Storage Manager using the command:

SMclient

4 For each storage enclosure in the Grid,

go to the "*Host Mappings"* tab and right-click on the *"Host Group PA"*.

- Select **Define**
- Select **Host**

5 Enter the cache engine name with a dash,

Note: For example it could be CE-1.

- Click **Next**

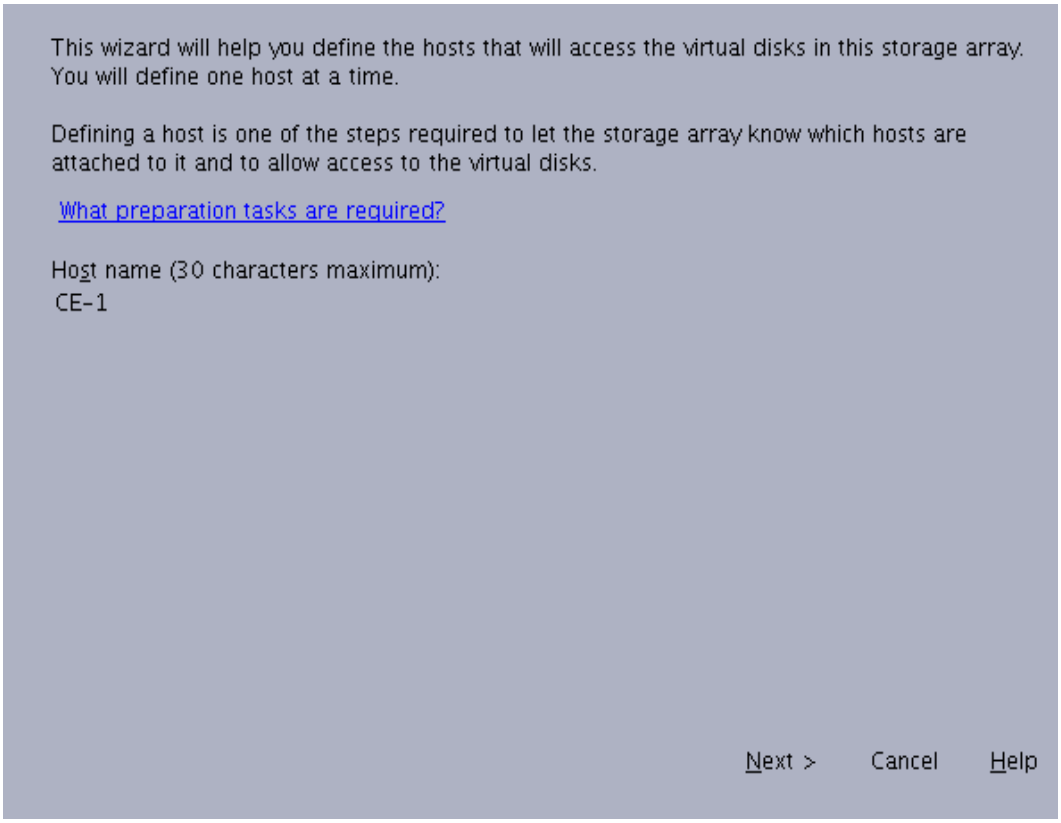

Manage Host Port Identifiers

6 Open a terminal to the management and open a ssh connection to the extended CE to check its iqn by using the command *cat /etc/iscsi/initiatorname.iscs*i

For example:

```
ce-3:~ # cat /etc/iscsi/initiatorname.iscsi
##
## /etc/iscsi/iscsi.initiatorname
##
## Default iSCSI Initiatorname.
##
## DO NOT EDIT OR REMOVE THIS FILE!
## If you remove this file, the iSCSI daemon will not start.
## If you change the InitiatorName, existing access control lists
## may reject this initiator.  The InitiatorName must be unique
## for each iSCSI initiator.  Do NOT duplicate iSCSI InitiatorNames.
InitiatorName=iqn.1996-04.de.suse:01:ff51f84a3bd
```

Note: The iqn is in the last line. It is bold in the example.

7 From the Define Host screen,

- Choose "**ISCSI**" (where applicable),
- Select the option **Add by Selecting Known…**
- Choose the server iqn from the drop-down list as found on step 6,
- In the *User Label* field, enter the cache engine name with a underscore (CE_<ce#>), in the example it will be CE_1,
- Click **Add**
- Click **Next**

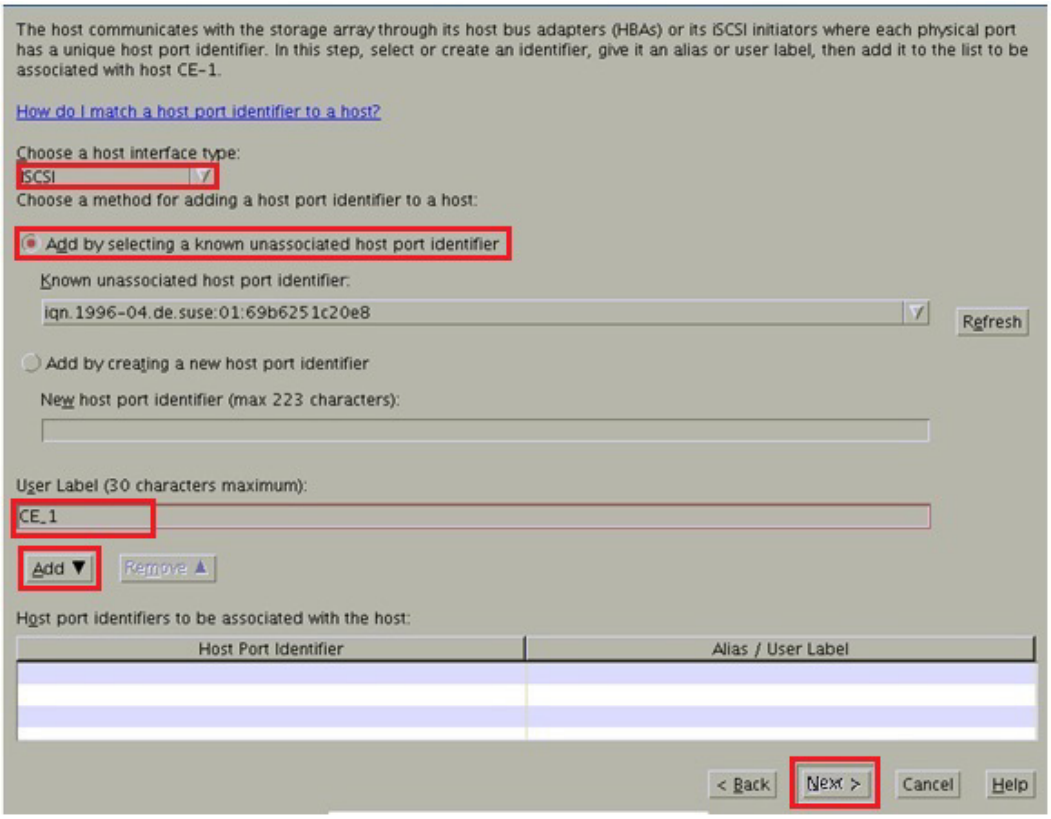

Define Host

8 Under “Host type” (operating system),

- Choose **Linux**
  - **a** In case that there is more than one “Linux” Host Type,
    - Choose: **Linux (MPP\RDAC)**
- Click **Next**

Host Type

9 From the Current Host Definition screen,

- Click **Finish**

Current Host Definition

Note: A popup screen is displayed asking, *"Do you want to define another host?"*. If there are additional CEs to add, answer **Yes**. Repeat the previous Step 5-Step 9.

Note: Repeat Step 4-Step 9 for all the storage enclosures.

10 From the management use the command:

cd /opt/pang/useful/configure_storages

For each additional CE Run the configure storage command:

./configure_storages.py -s < number of storage enclosures> -b < number of blades in the grid> -c <the additional CE-ID number>

Note: For each CE, run this script, each running the "-b" parameter which is the number of the CEs on the new grid and the "-c" parameter is changed for each running CE that is added. For example for an expanded grid of 4 CEs with 1 storage (2 CEs added) run:

./configure_storages.py -s 1 -b 4 -c 3
./configure_storages.py -s 1 -b 4 -c 4

Note: Additional steps: Follow the Administration Guide instructions regarding requesting a license with additional servers and contact your network engineers regarding additional bounce IP addresses and configuration to modify the cluster_conf appropriately.

11 Activate a new license that contains the new CE system ID. Activation license guide is in section "Installing and Configuring the Software" in the installation guide (see "Post Installation Steps").

12 Upload a new cluster_conf.xml with parameters for the new grid structure.

13 Software installation

- Copy the UltraBand Software Package to /tftpboot/

Note: The upgrade file name should be in this format (even version number):
`PeerApp_GA_5.7.0bXX_grid.tar '`

14 From the management,

- Enter to the CLI in enable mode.

15 Upgrade the added CEs by using:

console# upgrade server <CE number> 127.0.0.1 <software_name>

16 Start the application on the added CEs. For example using CE-2:

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

Chapter 4: External Storage Replacement Procedure

Overview

The following chapter details the procedure for replacing the External Storage.

Replacing an External Storage

In Order to Disconnect a Faulty Storage Use the Following Procedure:

1. As soon as the storage enclosure failure is detected, disconnect the iSCSI from the faulty enclosure to restore the system to nominal operation.

   ■ Stop the pang service using CLI.

   If the system was under load when a storage enclosure fails then in all probability all the cache engines will reboot.

   If by some chance any of the cache engines did not reboot manually reboot it via the commands:

   ssh root@ce-x
   echo b > /proc/sysrq-trigger

2. After all the cache engines have been rebooted, on every cache engine, run yast with the command:

   yast

3. From yast, enter the **Network Services**

   ■ Select **iSCSI Initiator**

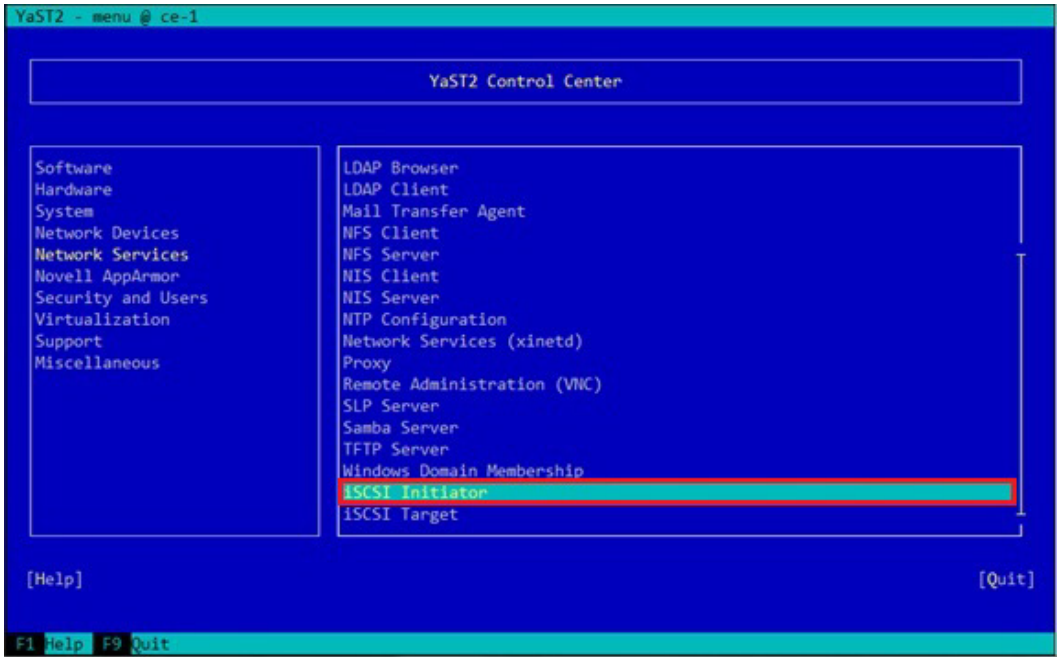

iSCSI Initiator

4 Select **Discovered Targets**,

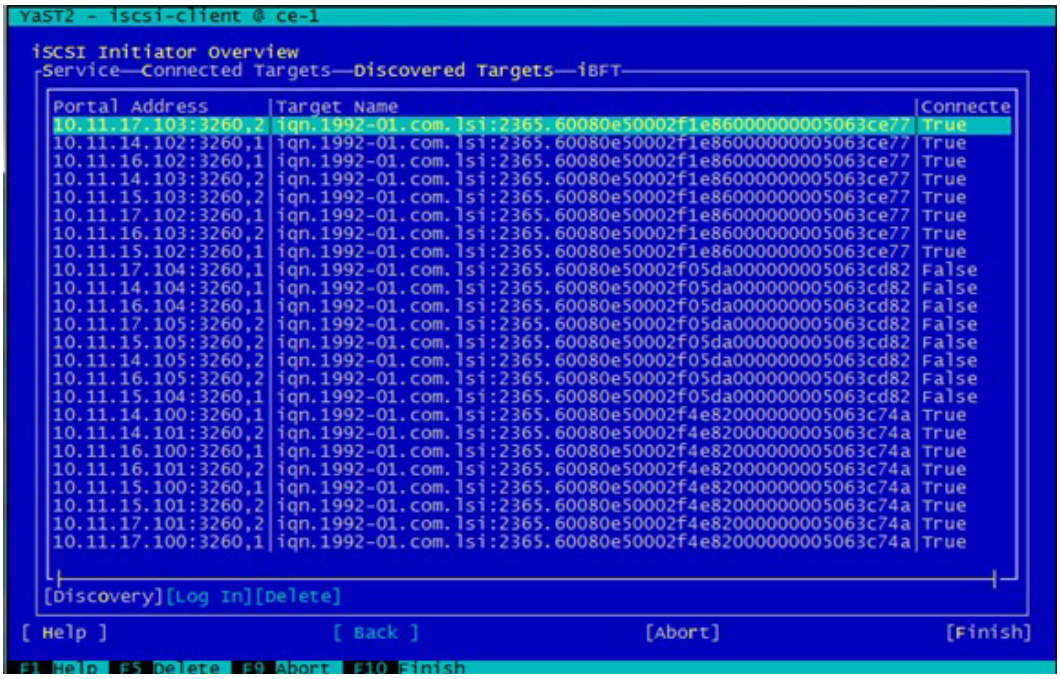

Discovered Targets

5 Select a node that has a "*false*" connection status

- Choose [**Delete**] for each "*false*" connection node.

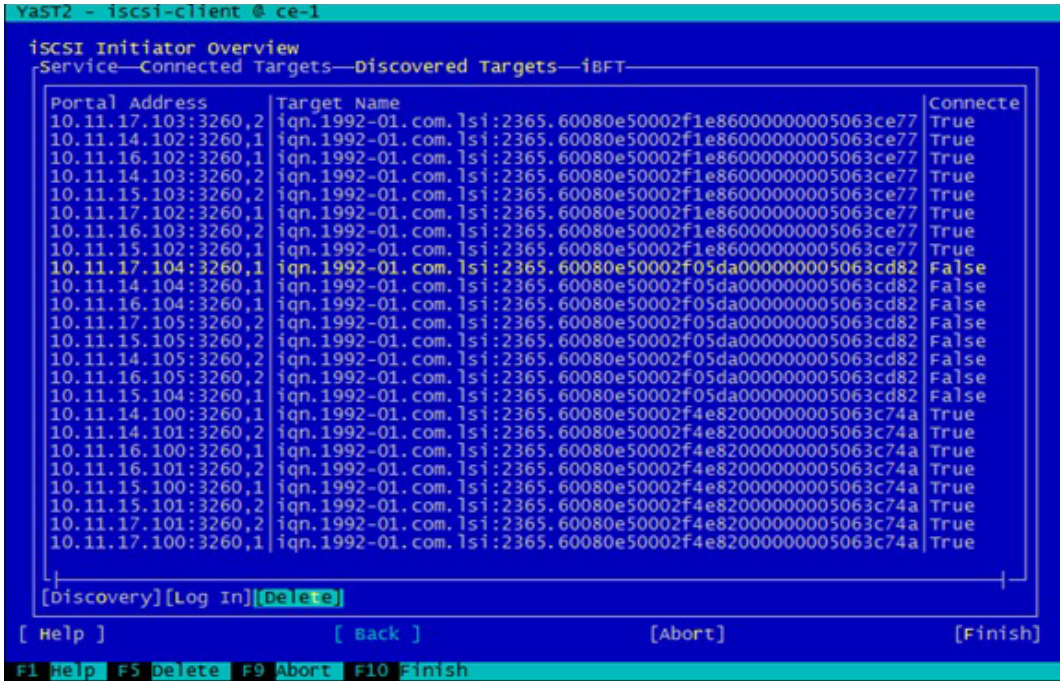

False Connection

Note: Repeat step 5 until there are no more nodes with a connected “false” status.

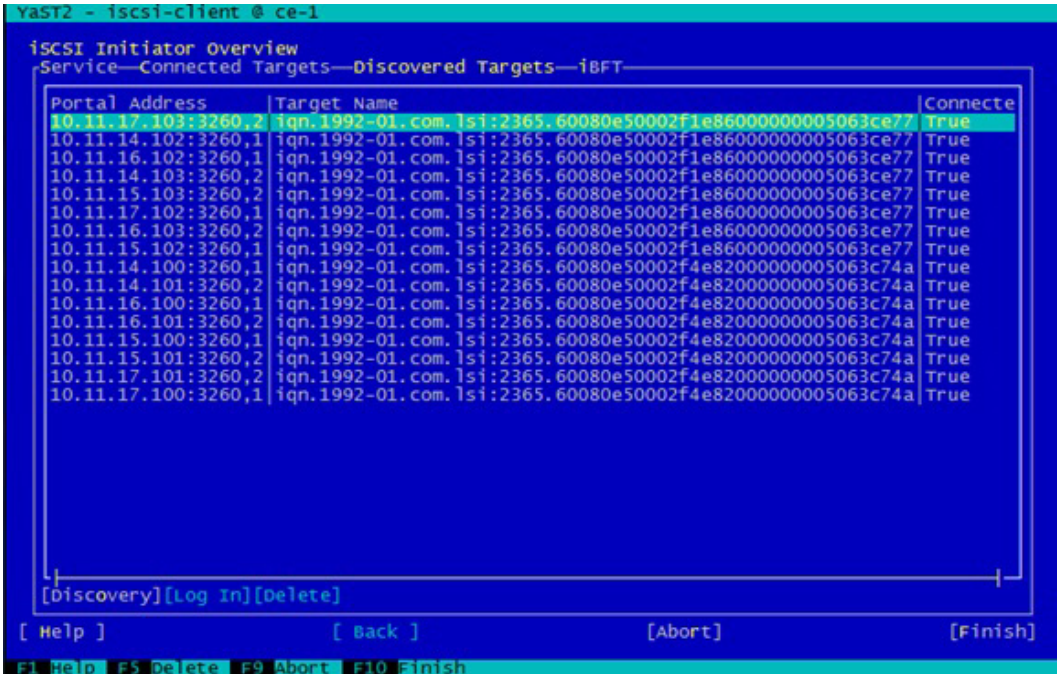

Discovered Targets 2

6 Choose **Finish** and **Quit**.

Note: Repeat steps 2-6 on all the cache engines.

Removing the Faulty Storage Volumes

After the failed storage ISCSI connections have been removed from all the cache engines, clean all volumes from the cmdb.

1. From the management server enter to enable mode in the CLI.
2. Using CLI, enter the command *show volumes*.

```
console# show 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
```

3. For each volume that has inactive status perform the command using CLI: *cache volume remove_content* and when asked, enter each time, one of the number of the inactive volumes. Enter **yes** for "*Are you sure?*".

   For example:

```
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.
. . .
console# cache volume remove_content
Licensed volumes : 72
Please enter volume number <1-72>
72
Are you sure? This will remove all hashes from volume 72.
[yes|no] no : yes
Removing all the content from volume 72...
Done.
```

Removing the Faulty Storage from the Storage Manager

From the Management Server,

1. Start *vncserver*:

```
rm /tmp/.X1-lock
rm /tmp/.X11-unix/X1
vncserver :1
```

2. Connect using **VNC Management Server**.

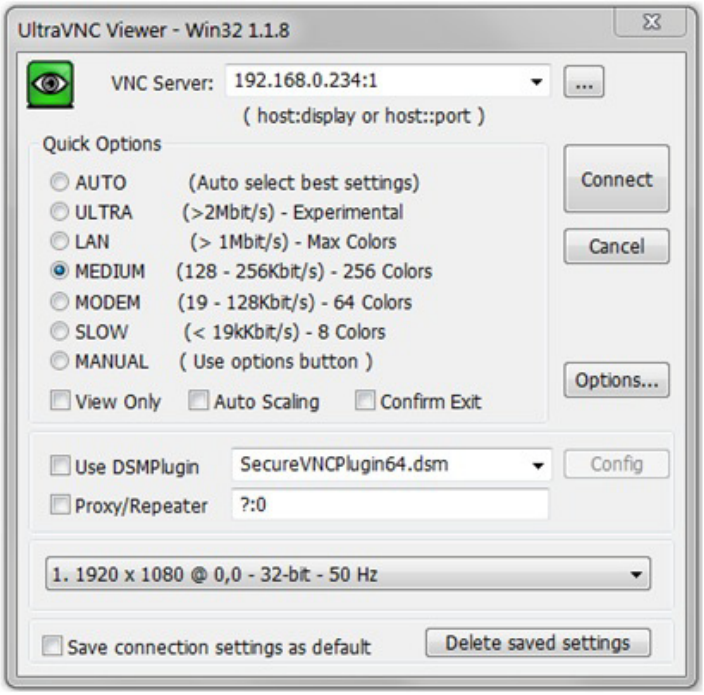

UltraVNC Viewer

3 Run the **ISCSI Storage Manager**

SMclient

4 From the *Devices* tab, locate the **faulty array**

- Select **Out-of-Band details**

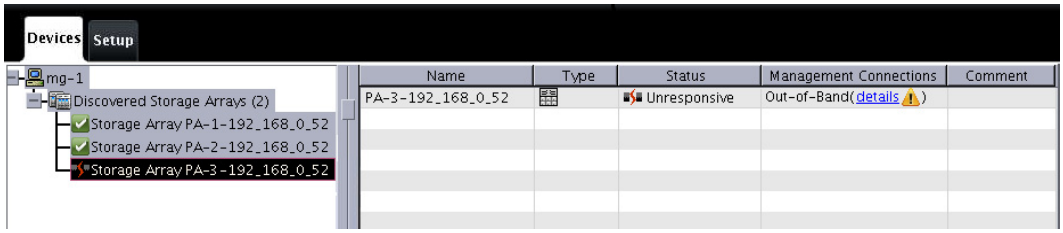

Devices Tab

Note: The IP addresses: Those same IP addresses need to be assigned to the replacement storage.

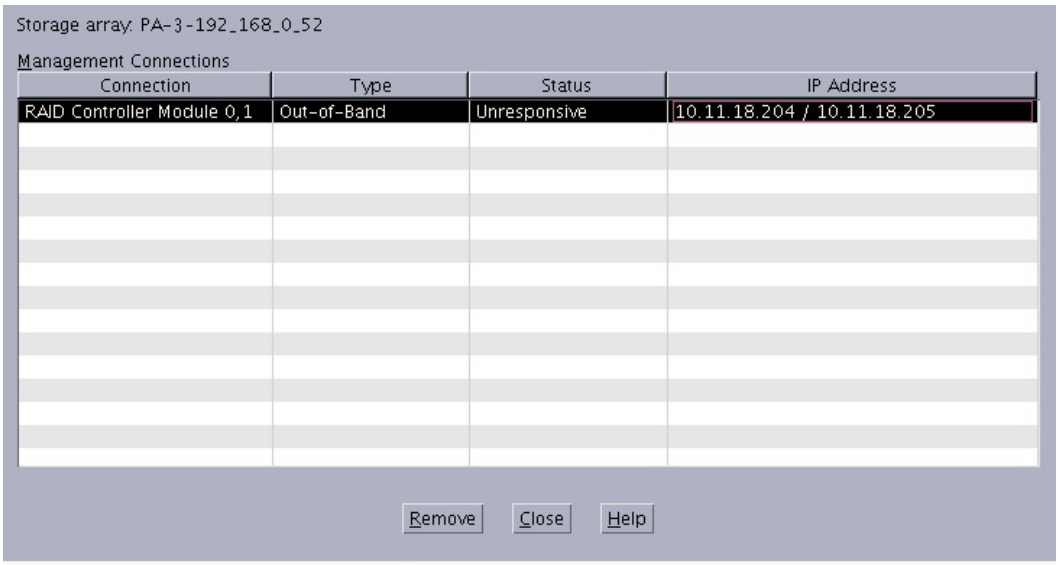

Note the IP Addresses

5 Click **Close**

6 Right-click on the "faulty array" and choose **Remove > Storage Array.**

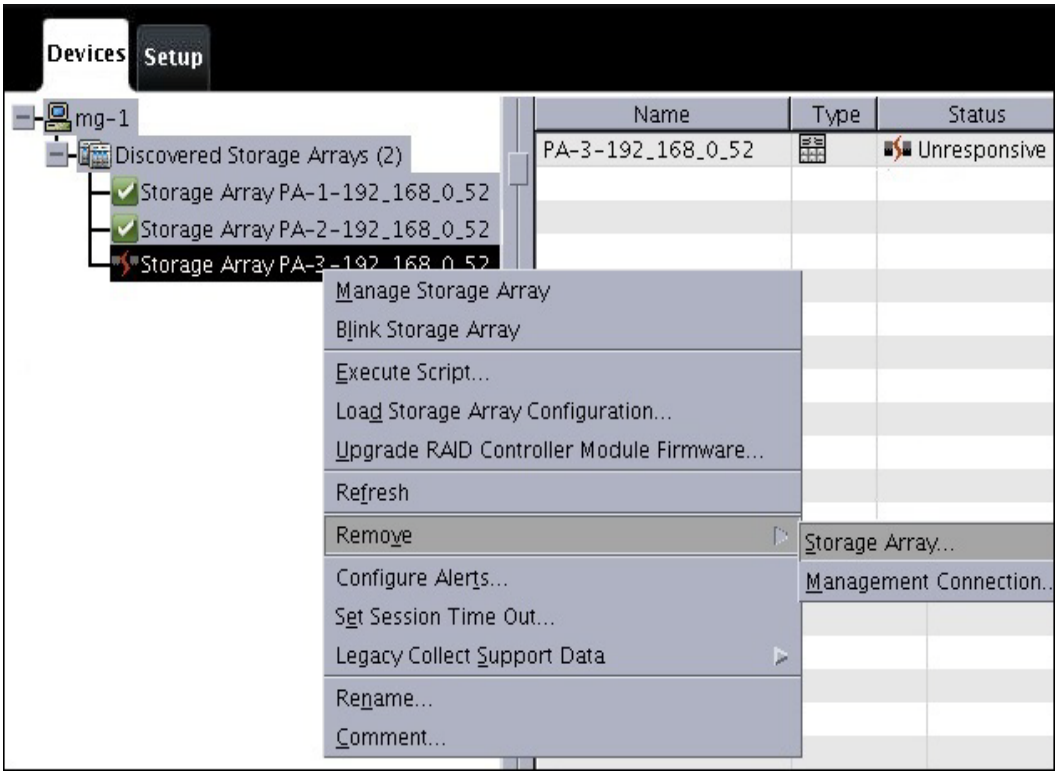

Remove Storage Array

7 From the *Confirm Remove Storage Array*

- Click **Yes**

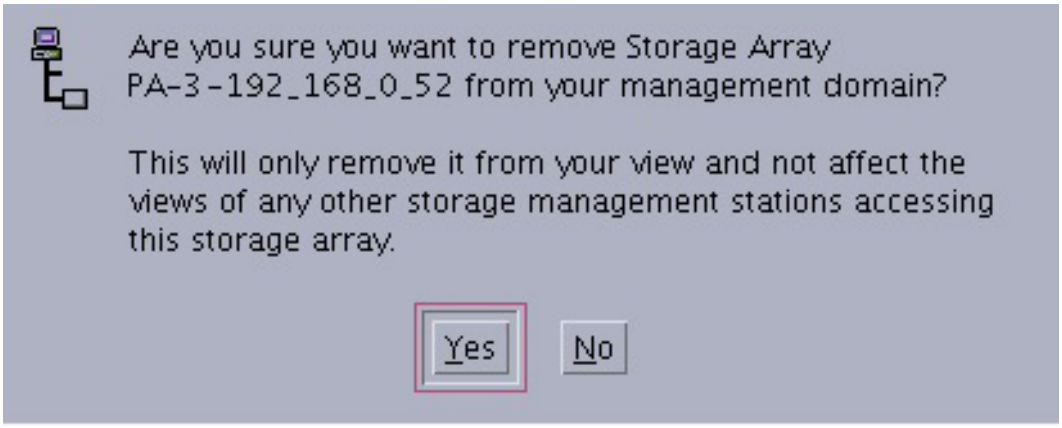

Confirm Remove Storage Array

The system now is back to operational state without the failed storage causing interference with IO operations.

Connecting a New Storage to an Existing System

In Order to Connect a New Storage Addition to an Existing System Use the Following Procedure:

From the *Management Server*,

1. Start the **VNC Server**, and enter,

   rm /tmp/.X1-lock
   rm /tmp/.X11-unix/X1
   vncserver :1

2. Connect using *VNC Management Server*,

   - Run the **ISCSI Storage Manager**

3. From the *Devices* tab,

   - Right-Click and choose **Add Storage Array**

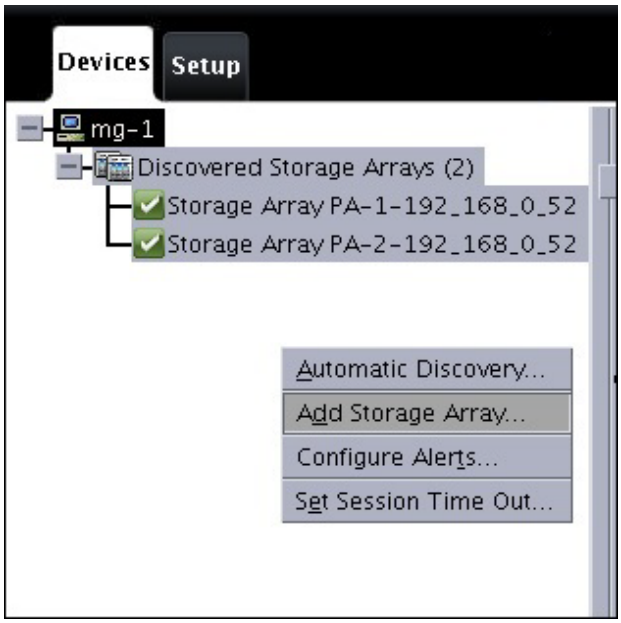

Add Storage Array

4 Follow the procedure in the installations manual that documents how to assign the appropriate IP addresses to a new storage from the factory by connecting to the default factory assigned IP addresses and changing them to the appropriate IP addresses in sequence be it an additional storage or filling a hole in the sequence

Up to 10 storages are supported - the IP addresses are assigned to each storage per the following table:

[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"
[7]="10.11.18.212 10.11.18.213"
[8]="10.11.18.214 10.11.18.215"
[9]="10.11.18.216 10.11.18.217"
[10]="10.11.18.218 10.11.18.218"

5 Enter the IP addresses (in the example this is Storage #3so the IPs are 10.11.18.204 & 10.11.18.205),

- Click **Add**.

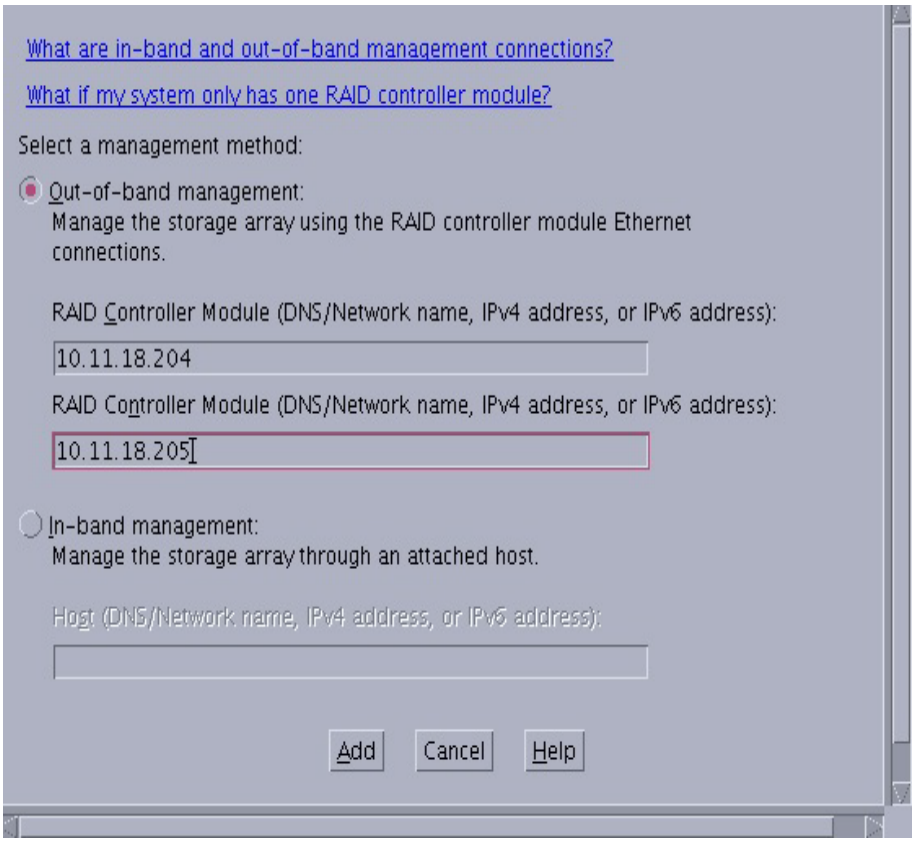

Add New Storage Array

6 From the *Storage Array Added*,

- Click **No**.

7 Manage the *Added Storage Array* by right-clicking,

- Choose **Manage Storage Array**

8 From the *Managing* window,

- Choose the **Setup** tab
- Select **Configure the iSCSI Host Ports**

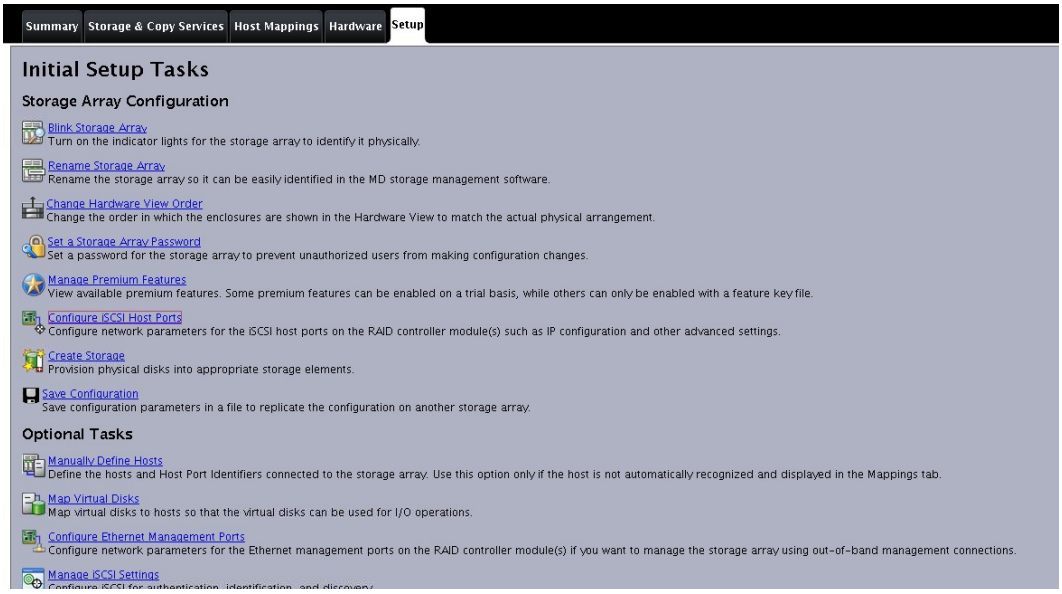

Configure the iSCSI Host Ports

9 Assign the ISCSI IP addresses and the settings (*Enable ICMP PING responses*, *disable IPv6* and in Advanced Port Settings, *enable Jumbo frame of MTU size 9000*) - for more instructions, see the Grid Installation Guide.

To assist with the configuration of the IP addresses, use the following IP address table for the storages with the 1G ISCSI connections. For the 10G ISCI connections, use only the addresses in **bold:**

Storage # 01
[1]="10.11.14.100"
[2]="10.11.15.100"
[3]="10.11.16.100"
[4]="10.11.17.100"
[5]="10.11.14.101"
[6]="10.11.15.101"
[7]="10.11.16.101"
[8]="10.11.17.101"
Storage # 02
[9]="10.11.14.102"
[10]="10.11.15.102"
[11]="10.11.16.102"
[12]="10.11.17.102"
[13]="10.11.14.103"
[14]="10.11.15.103"
[15]="10.11.16.103"
[16]="10.11.17.103"
Storage # 03
[17]="10.11.14.104"
[18]="10.11.15.104"
[19]="10.11.16.104"
[20]="10.11.17.104"
[21]="10.11.14.105"
[22]="10.11.15.105"
[23]="10.11.16.105"
[24]="10.11.17.105"
Storage # 04
[25]="10.11.14.106"
[26]="10.11.15.106"
[27]="10.11.16.106"
[28]="10.11.17.106"

[29]="10.11.14.107"
[30]="10.11.15.107"
[31]="10.11.16.107"
[32]="10.11.17.107"
Storage # 05
[33]="10.11.14.108"
[34]="10.11.15.108"
[35]="10.11.16.108"
[36]="10.11.17.108"
[37]="10.11.14.109"
[38]="10.11.15.109"
[39]="10.11.16.109"
[40]="10.11.17.109"
Storage # 06
[41]="10.11.14.110"
[42]="10.11.15.110"
[43]="10.11.16.110"
[44]="10.11.17.110"
[45]="10.11.14.111"
[46]="10.11.15.111"
[47]="10.11.16.111"
[48]="10.11.17.111"
Storage # 07
[49]="10.11.14.112"
[50]="10.11.15.112"
[51]="10.11.16.112"
[52]="10.11.17.112"
[53]="10.11.14.113"
[54]="10.11.15.113"
[55]="10.11.16.113"
[56]="10.11.17.113"
Storage # 08
[57]="10.11.14.114"
[58]="10.11.15.114"
[59]="10.11.16.114"
[60]="10.11.17.114"
[61]="10.11.14.115"
[62]="10.11.15.115"
[63]="10.11.16.115"
[64]="10.11.17.115"
Storage # 09
[65]="10.11.14.116"
[66]="10.11.15.116"
[67]="10.11.16.116"
[68]="10.11.17.116"
[69]="10.11.14.117"
[70]="10.11.15.117"
[71]="10.11.16.117"
[72]="10.11.17.117"
Storage # 10
[73]="10.11.14.118"
[74]="10.11.15.118"
[75]="10.11.16.118"
[76]="10.11.17.118"
[77]="10.11.14.119"
[78]="10.11.15.119"
[79]="10.11.16.119"
[80]="10.11.17.119"

Configure the New Storage

1 From the *Management*,

■ **Running the Configure Storages Script** using the commands:

cd /opt/pang/useful/configure_storages/
./configure_storages.py -s <the number of the storages in the system> -n < the number of the replaced storage> -b <the number of the CEs in the system>

For example, for a system with 8 cache engines and 3 storages (when the 3rd one is the new added one):

./configure_storages.py -s 3 -n 3 -b 8

A warning message about the disk format will be shown, to continue the storages configuration enter y.

********** WARNING! **********
The selected storage is about to be reconfigured. All data on that storage
will be lost. You'll need to format disks manually.
Do you want to continue? (y/n) y

2 In order to begin using the new volumes, start the Service for Pang.

Chapter 5: Replacing the Management Server Procedure

Overview

The following chapter details the procedure for replacing the management server.

Prerequisites

Before Replacing the Management Server, you must possess the following:

- Disk-on-key with the ISO image
- The software version that is running on the Cache Engines

Note: For further information, refer to the installation manual (Chapter 1: Rack Mounting).

Backup Files

IMPORTANT! The following procedure for backing up files apply only if the Management Server is still accessible. If the management Server is not accessible, please contact customer support for the License and cluster_conf.xml files.

1. Use WinSCP to the management server.
2. Copy and backup these files to a place that is not the management server:
   - */opt/pang/mgmt/config/PALicense.xml*
   - */opt/pang/mgmt/config/cluster_conf.xml*

Replacing the Management Server

First disconnect the old Management Server and replace it with the new Management Server. Next, connect the new Management Server to the Cache Engines.

Note: For more information, refer to the installation manual (Chapter 2: Cabling - Management Server).

Installing the Management Server

1 Configure the BIOS of the server, and install the ISO.

Note: For more information, refer to the installation manual (Chapter 4: Configuring the Management Server).

2 After the ISO installation is completed.

Note: When requested, use the old Management Server Network parameter and address values.

3 Using a PuTTy session, connect as a root user to the management IP.

4 From the management, update the grid ssh keys by running:

- *cd /opt/pang/useful/*
- *./replace_server_keys.sh -servers < number of blades in the grid>*

IMPORTANT! (For example) For a Grid with 4 blades run:
./replace_server_keys.sh -servers 4

5 Refer to the Installation Guide (Chapter 6: Configuring the Cache Storage), and configure the Storages.

6 Obtain a copy of the UltraBand 6000 Grid software package (FI) with the same version number as the Cache Engines.

7 From the management run:

- *cd /tftpboot/*
- *mkdir Installer_ver<version #>*
- *cd Installer_ver<version #>*

8 Copy the Software Package installer to *Installer_ver<version #>*

9 On the archive, run:

- *tar -zxvf <FI_installer_name>*

Note: If you are connected remotely to the 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*.

10 Edit GA_installer.rc - Edit the line "UPGRADE CACHE SOFTWARE" to be none.

For example:
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

11 Run the *./GA_installer.sh*

Note: Refer to the Installation Manual (Chapter 7: Installing and Configuring the Software - Installing the Management Environment and Software section).

12 From the Management (root permissions),

- run the command: cd /opt/pang/useful
- Run the script *management_machine_replacement.sh* with the parameters list, denoting all the existing cache blades, as follows: For a Grid of 8:

./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 for restoring the previously backed up files. If the backup files are not available, skip step 1, and proceed to step 2 to install the License and cluster_conf.xml files obtained from customer support.

1 Use WinSCP to the management server and import the backup files - PALicense.xml and cluster_conf.xml to the folder */tftpboot*

2 Import the cluster_conf.xml and PALicense.xml

Note: Refer to the Installation Guide (Chapter 7: Installing and Configuring the Software - Post Installation Steps).

IMPORTANT! If step 2 fails - Run the following from the Management Server:
-*/opt/pang/mgmt/bin/pang_rescue*
-Type: *passwords*
-Exit
-Repeat step 2

3 Exit from the CLI, (still) as the root user enter the following commands:

cd /etc

scp root@ce-1:/etc/hosts .

sed -i 's/mg-1/mg-1.cosco.com mg-1/g' hosts

Chapter 6: Connecting an External Storage to an Existing System

Overview

The following chapter details the procedure for connecting a new storage addition to an existing system.

Connecting a Server Storage

In Order to Connect a New Storage Addition to an Existing System Use the Following Procedure:

From the *Management Server*,

1. Start the **VNCServer**, and enter,

   rm /tmp/.X1-lock
   rm /tmp/.X11-unix/X1
   vncserver :1

2. Connect to the Management Server using **VNC** and run**:**

   SMclient

3. From the *Devices* tab,

   - Right-Click and choose **Add Storage Array**

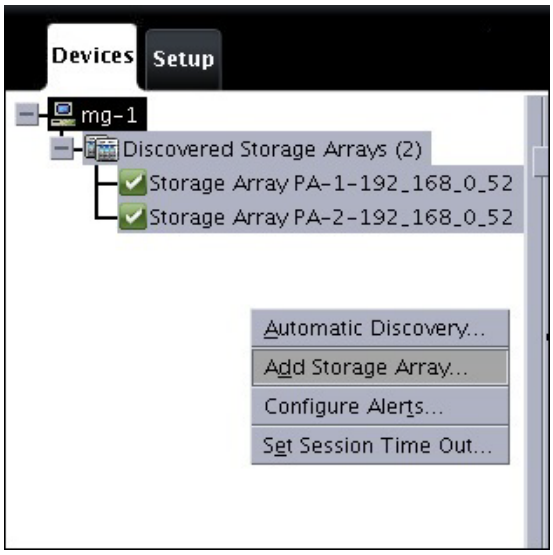

Add Storage Array

4 Follow the procedure in the installations manual that documents how to assign the appropriate IP addresses to a new storage from the factory by connecting to the default factory assigned IP addresses and changing them to the appropriate IP addresses in sequence be it an additional storage or filling a hole in the sequence

Up to 10 storages are supported – the IP addresses are assigned to each storage per the following table:

[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"
[7]="10.11.18.212 10.11.18.213"
[8]="10.11.18.214 10.11.18.215"
[9]="10.11.18.216 10.11.18.217"
[10]="10.11.18.218 10.11.18.218"

5 Enter the IP addresses (in the example this is Storage #3 so the IPs are 10.11.18.204 & 10.11.18.205),

■ Click **Add**.

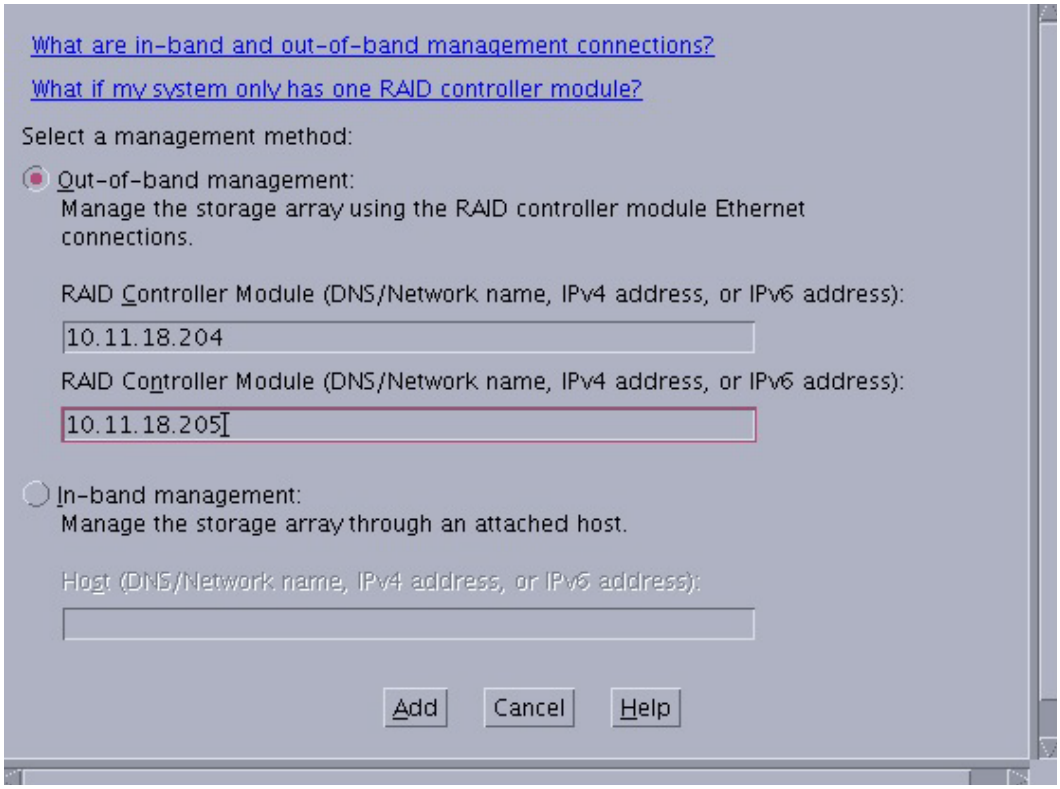

Add New Storage Array

6 From the *Storage Array Added* screen,

■ Click **No**

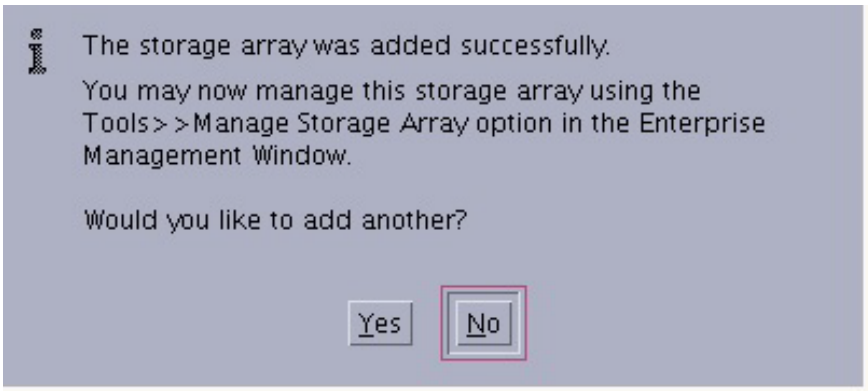

Storage Array Added

7 Manage the *Added Storage Array* by right-clicking it,

- Choose **Manage Storage Array**

8 From the *Managing* window,

- Choose the **Setup** tab
- Select the **Configure ISCSI Host Ports.**

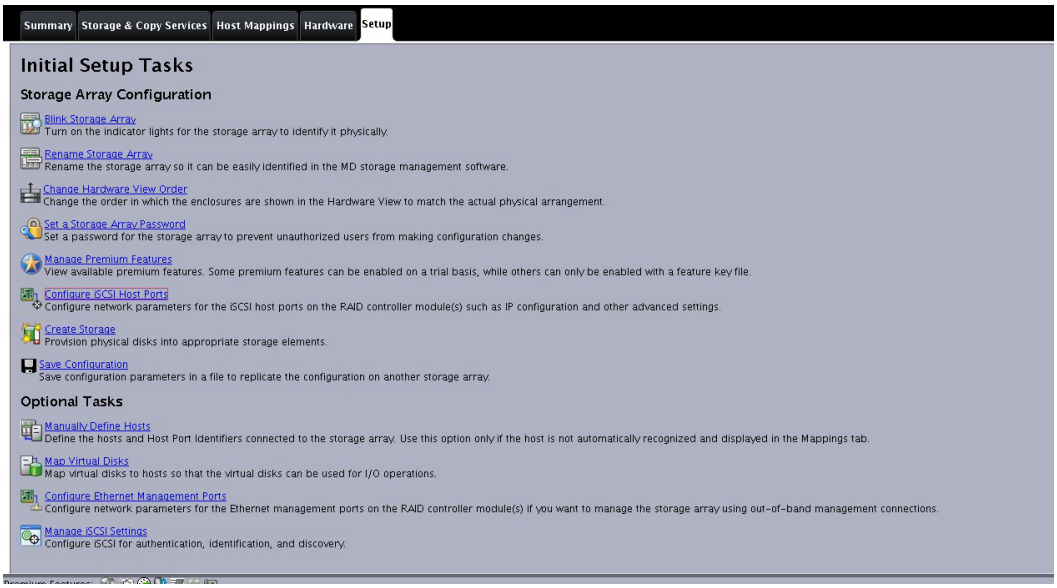

Initial Setup Tasks

9 Assign the iSCSI IP addresses and the settings (enable *ICMP PING responses*, *jumbo & disable IPv6* and in Advanced Port Settings enable *Jumbo frame of MTU size 9000*) - as per the for more instructions use in the installation manual.

To assist with the configuration of the IP addresses, use the following IP address table for the storages with the 1G ISCSI connections. For the 10G ISCI connections, use only the addresses in **bold:**

Storage # 01
[1]="10.11.14.100"
[2]="10.11.15.100"
[3]="10.11.16.100"
[4]="10.11.17.100"
[5]="10.11.14.101"
[6]="10.11.15.101"
[7]="10.11.16.101"
[8]="10.11.17.101"

Storage # 02
[9]="10.11.14.102"
[10]="10.11.15.102"
[11]="10.11.16.102"
[12]="10.11.17.102"
[13]="10.11.14.103"
[14]="10.11.15.103"
[15]="10.11.16.103"
[16]="10.11.17.103"
Storage # 03
[17]="10.11.14.104"
[18]="10.11.15.104"
[19]="10.11.16.104"
[20]="10.11.17.104"
[21]="10.11.14.105"
[22]="10.11.15.105"
[23]="10.11.16.105"
[24]="10.11.17.105"
Storage # 04
[25]="10.11.14.106"
[26]="10.11.15.106"
[27]="10.11.16.106"
[28]="10.11.17.106"
[29]="10.11.14.107"
[30]="10.11.15.107"
[31]="10.11.16.107"
[32]="10.11.17.107"
Storage # 05
[33]="10.11.14.108"
[34]="10.11.15.108"
[35]="10.11.16.108"
[36]="10.11.17.108"
[37]="10.11.14.109"
[38]="10.11.15.109"
[39]="10.11.16.109"
[40]="10.11.17.109"
Storage # 06
[41]="10.11.14.110"
[42]="10.11.15.110"
[43]="10.11.16.110"
[44]="10.11.17.110"
[45]="10.11.14.111"
[46]="10.11.15.111"
[47]="10.11.16.111"
[48]="10.11.17.111"
Storage # 07
[49]="10.11.14.112"
[50]="10.11.15.112"
[51]="10.11.16.112"
[52]="10.11.17.112"
[53]="10.11.14.113"
[54]="10.11.15.113"
[55]="10.11.16.113"
[56]="10.11.17.113"
Storage # 08
[57]="10.11.14.114"
[58]="10.11.15.114"
[59]="10.11.16.114"
[60]="10.11.17.114"
[61]="10.11.14.115"
[62]="10.11.15.115"
[63]="10.11.16.115"

[64]="10.11.17.115"
Storage # 09
[65]="10.11.14.116"
[66]="10.11.15.116"
[67]="10.11.16.116"
[68]="10.11.17.116"
[69]="10.11.14.117"
[70]="10.11.15.117"
[71]="10.11.16.117"
[72]="10.11.17.117"
Storage # 10
[73]="10.11.14.118"
[74]="10.11.15.118"
[75]="10.11.16.118"
[76]="10.11.17.118"
[77]="10.11.14.119"
[78]="10.11.15.119"
[79]="10.11.16.119"
[80]="10.11.17.119"

Configure the New Storage

1 To stop the *Pang Service*, from the CLI,

■ Use *Oper Service Stop* from the CLI

Note: Repeat the following steps for each new storage enclosure

2 From the *Management*,

■ Run the configure storage script:

cd /opt/pang/useful/configure_storages
`./configure_storages.py -s <the number of the storages in the system> -n < the number of the new storage> -b <the number of the CEs in the system>`

For example, for a system with 8 cache engines, and 3 storages (when the 3rd one is the new added one):

./configure_storages.py -s 3 -n 3 -b 8

A warning message about the disk format will be shown, to continue the storages configuration enter y.

********** WARNING! **********
The selected storage is about to be reconfigured. All data on that storage
will be lost. You'll need to format disks manually.
Do you want to continue? (y/n) y

Note: Repeat Step 2 for each new storage enclosure

3 Reboot all of the cache engines in the grid (except ce-1) allowing the re-reading of formatted disks (by using the command reboot *-f)*.

4 Update the `cluster_conf.xml` with the new controller addresses (for more information, contact customer support)

5 In order to begin using the new volumes, start the Service for Pang.

Chapter 7: External Storage Disk Replacement Procedure

Overview

As soon as a disk failure is detected, delete the faulty Disk Group in the storage to cease the I/O errors and restore the system nominal operation.

Replacing an External Storage Disk

To Replace the External Storage Disk:

From the Management Server,

1. Start the VNC Viewer, and enter

   rm /tmp/.X1-lock
   rm /tmp/.X11-unix/X1
   vncserver :1

2. Connect using VNC the Management to the server

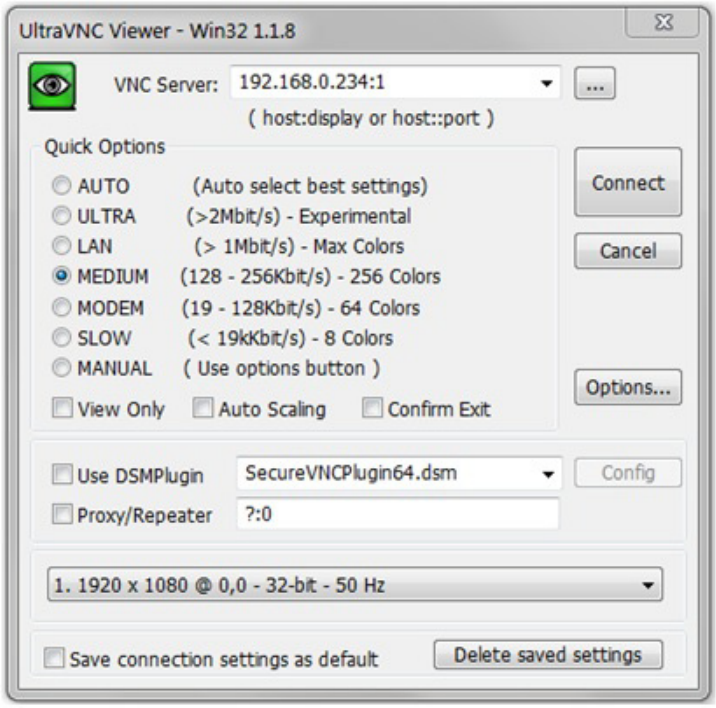

VNC Viewer

3. Run the MD Storage Manager

4 From the storage tab,

- Locate the faulty Disk Group
- Click + to verify that there is a faulty Virtual Disk

Note: If the LUN is 0. You cannot delete it. You must contact customer support for further assistance. If the LUN is any other number, you can delete the disk by using the procedure detailed below.

Note: By clicking the + and then clicking the Virtual Disk, a summary of the drive specifications is displayed on the right pane. Look under the LUN designation for the LUN number.

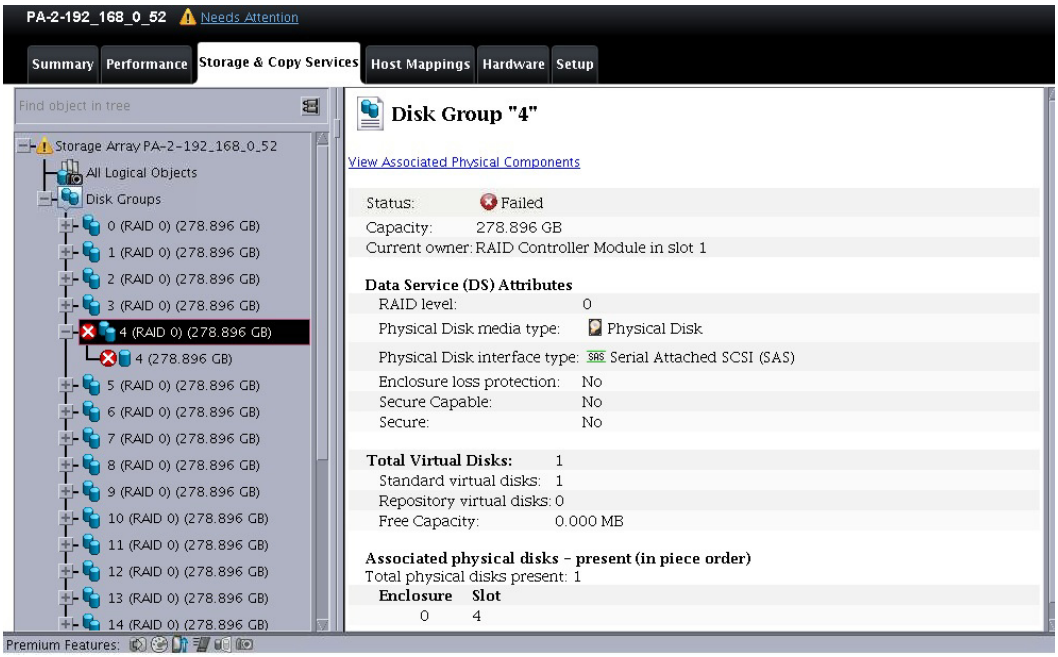

Disk Group 4

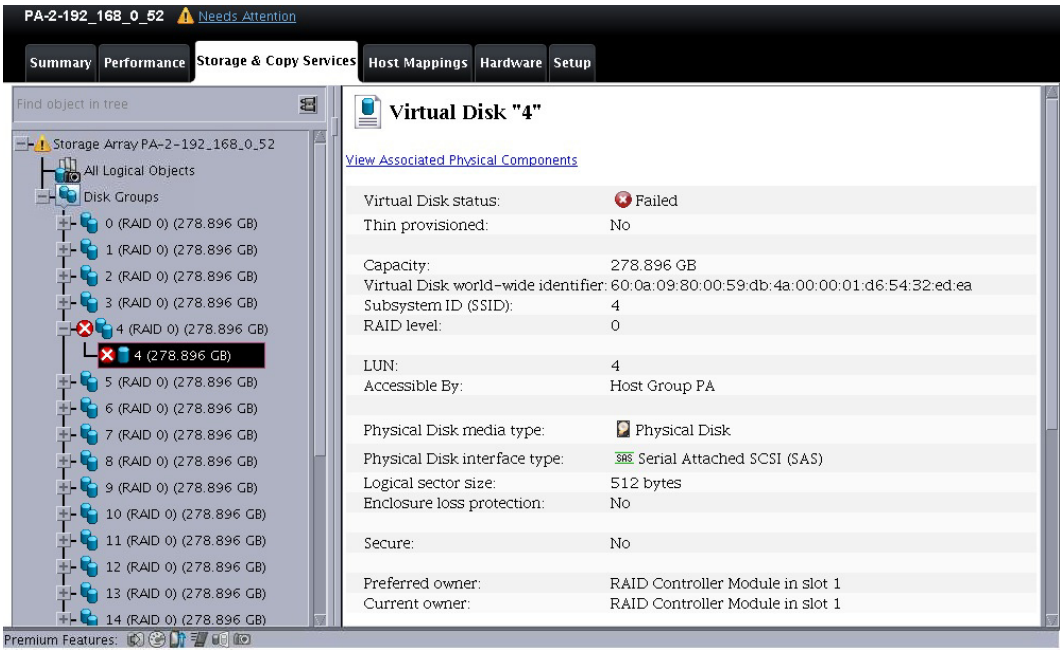

Disk Group 5

5 Right click on the faulty Disk Group (not the Virtual Disk)

- Click **Delete**

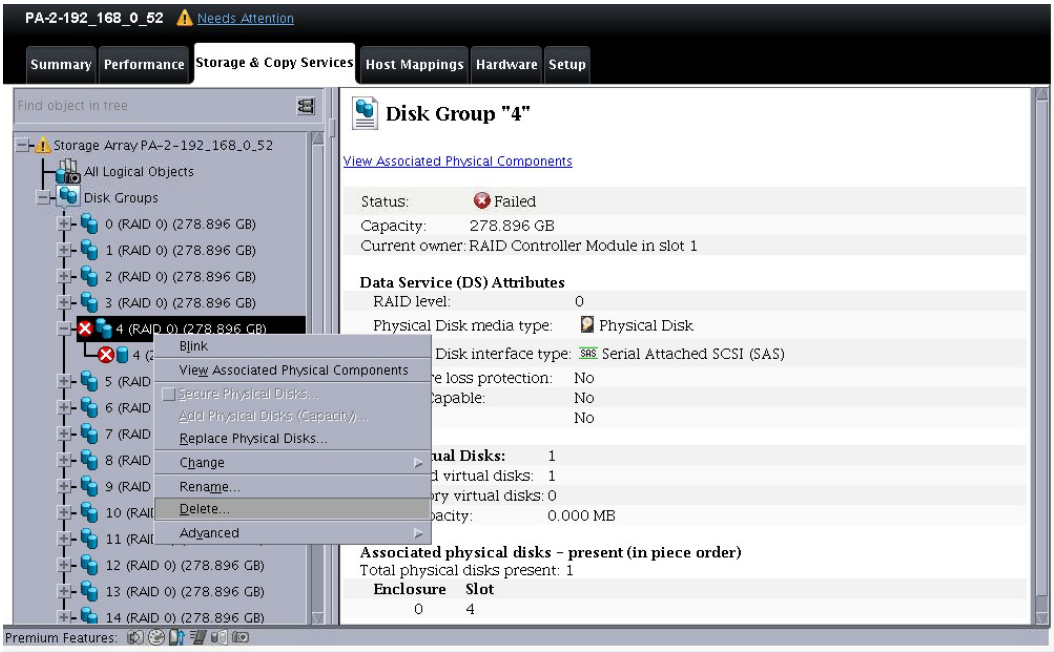

Faulty Disk Group

6 From the Delete Disk Groups window

- Click **Delete**

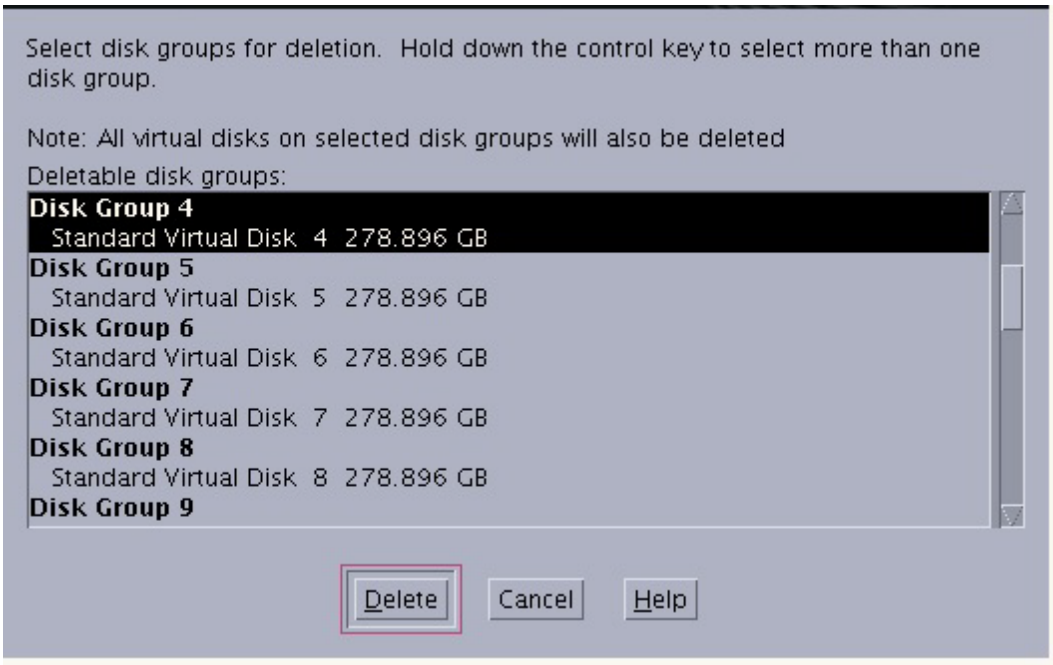

Delete Disk Group

7 From the Confirm Delete Disk Groups window,

- Enter "**Yes**"
- Click **OK**
- Click **OK** (process completed)

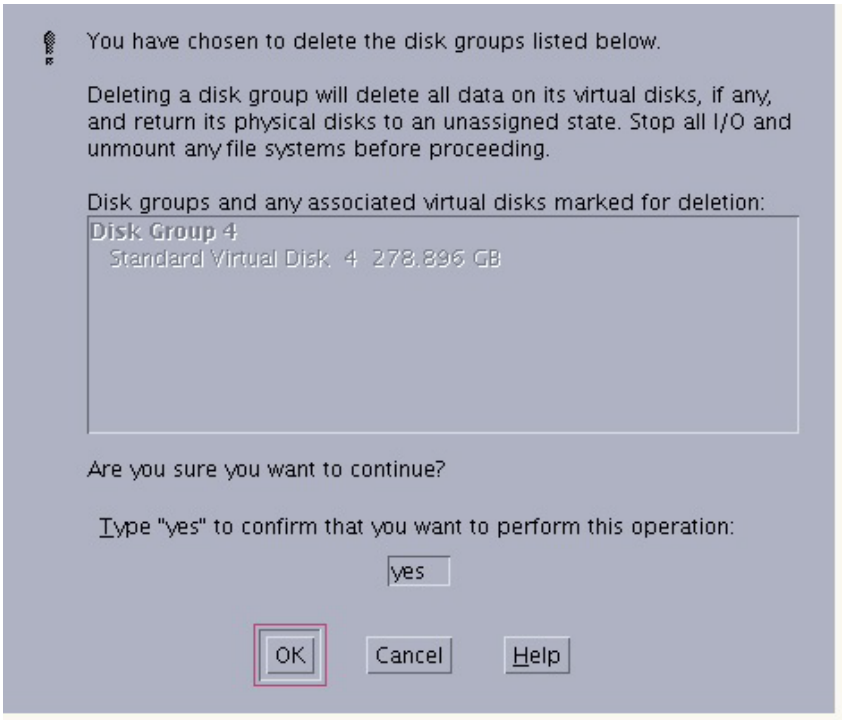

Confirm Delete Disk Group

8 Using CLI, use this procedure to remove the faulted volume:

```
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.
. . .

Note: Wait until the physical device is repaired

9 From the Storage Manager, select the Storage &Copy Services tab,

- Right-click the *Total Unconfigured Capacity*
- Click **Create Disk Group**

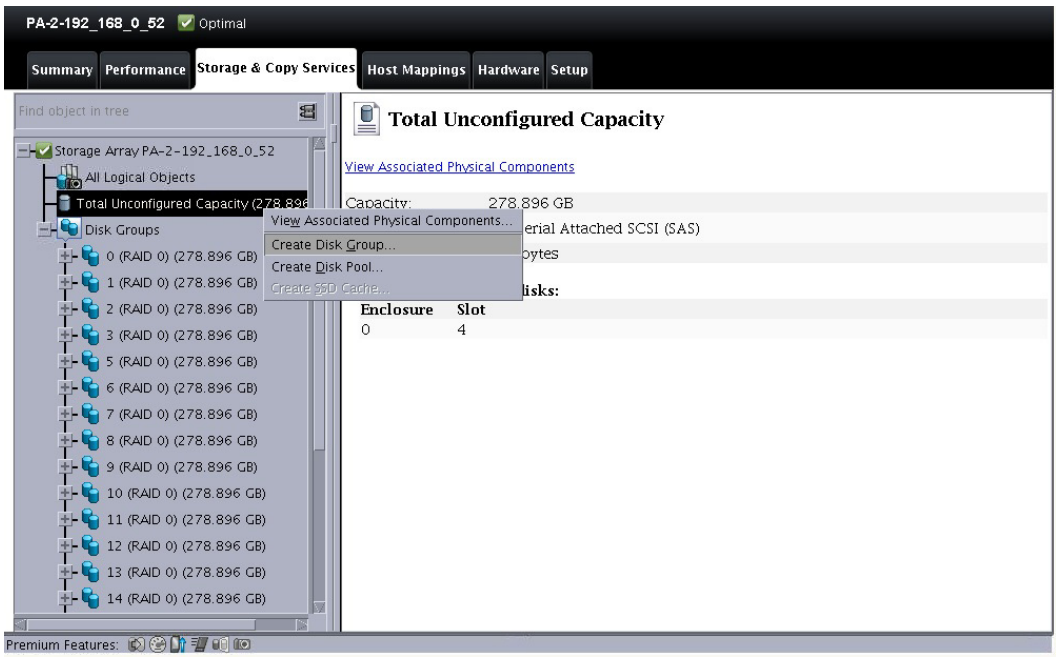

Create Disk Group

10 From the Wizard,

- Click **Next**

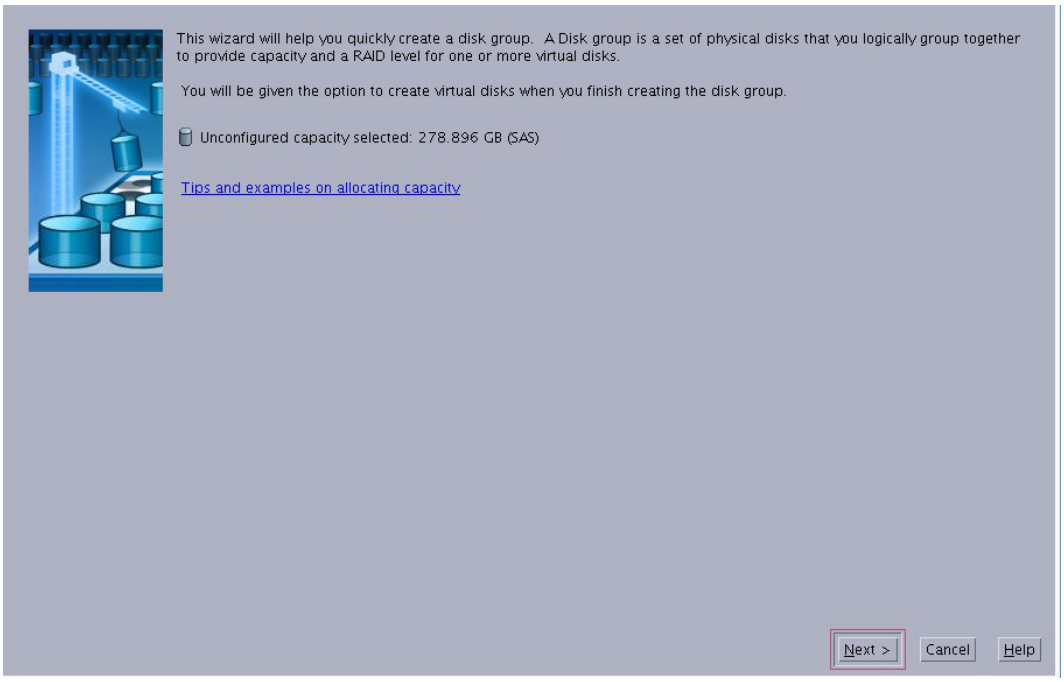

Wizard

11 From the Disk Group Name window,

- **Do Not Change the Disk Group Name** (Leave the Default Name)

12 From the Raid Level window,

- Click **Manual**,
- Click **Next**

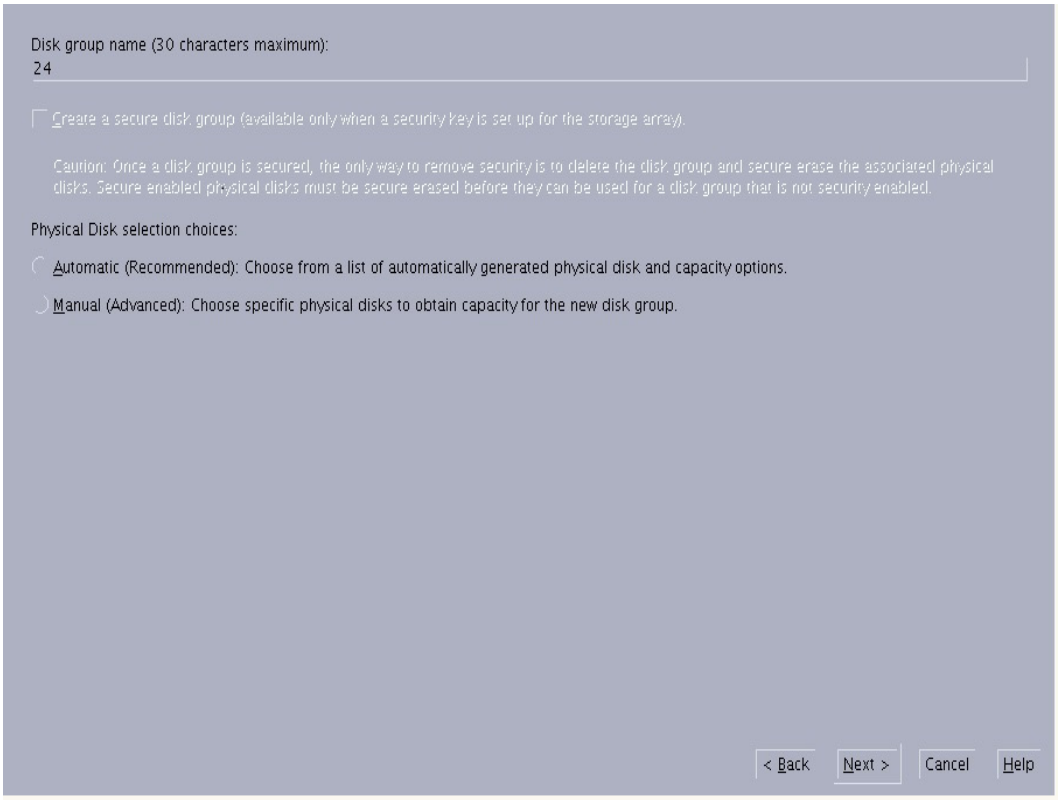

Disk Group Name

13 From the Raid Level window,

- Select **Raid Level 0**,
- Click the **Unassigned Drive**,
- Click **Add,**
- Click **Calculate Capacity**,
- Click **Finish**

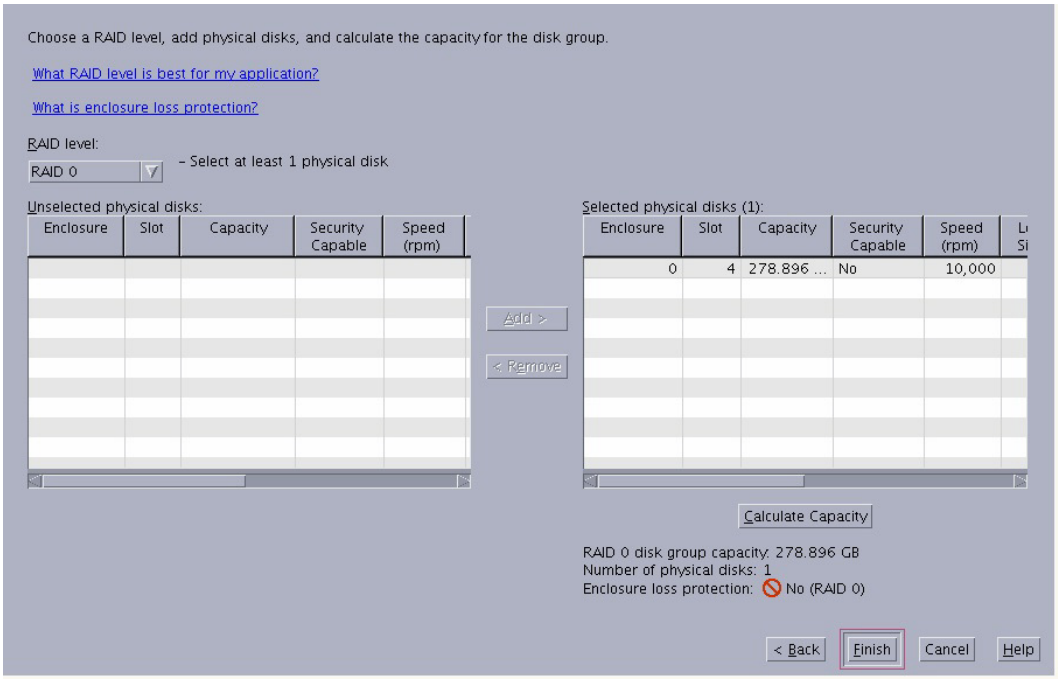

Raid Level 0

14 From the Disk Group was Successfully Created window,

- Click **Yes**

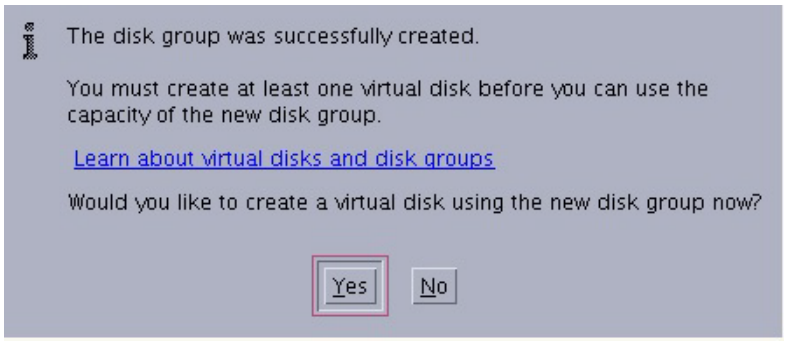

Disk Group Successfully Completed

15 From the New Virtual Disk Capacity,

- Enter the number shown in the free capacity as displayed below,
- **Do Not Change the Virtual Disk Name (Leave the Default)**

16 In the map to host field,

- Choose **Map Later**

17 In the Virtual Disk Characteristics,

- Choose **Custom**
- Clear (disable) the **Dynamic Cache Read Prefetch**

18 In segment size,

- Select **256KB**
- Click **Finish** and **OK**

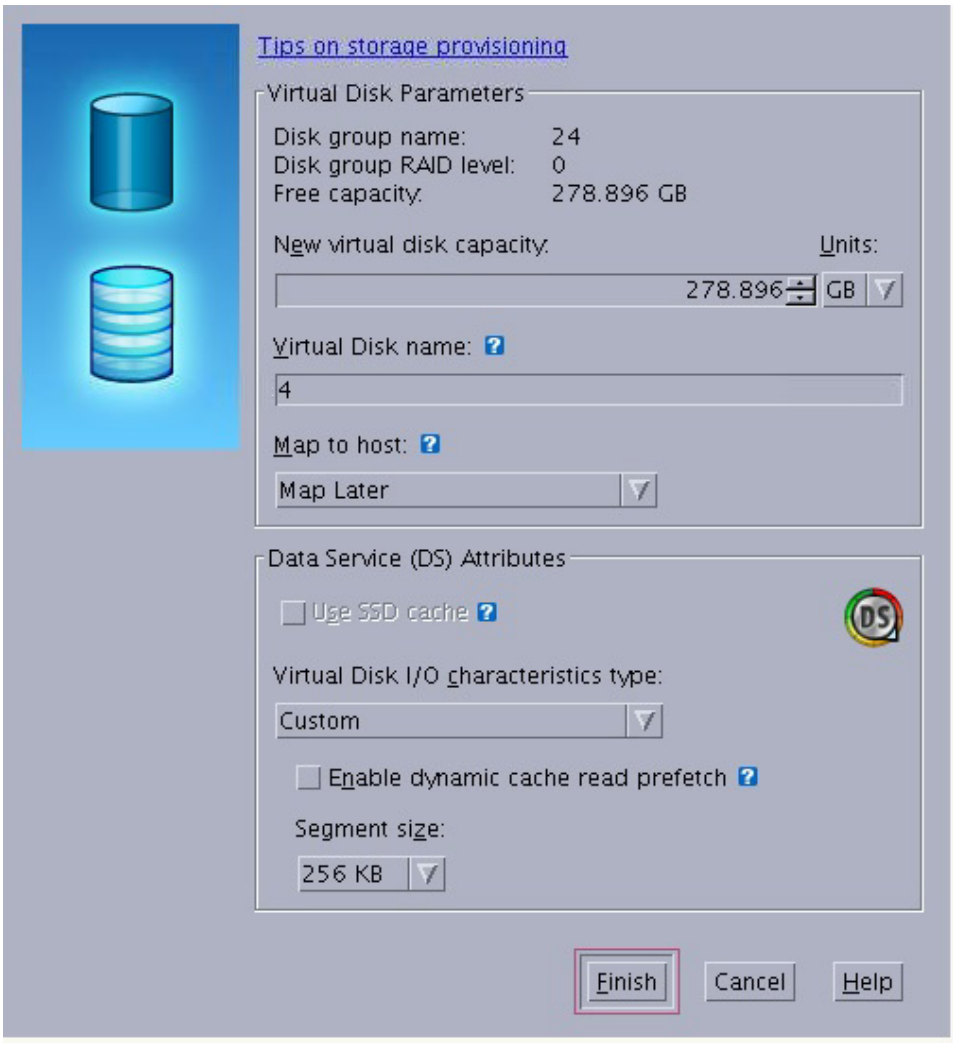

Create Virtual Disk

19 From the Storage & Copy Services,

- Right-click the newly added Virtual Disk (Virtual Disk 5 in the example)
- Select **Change**
  - Select **Cache Settings**

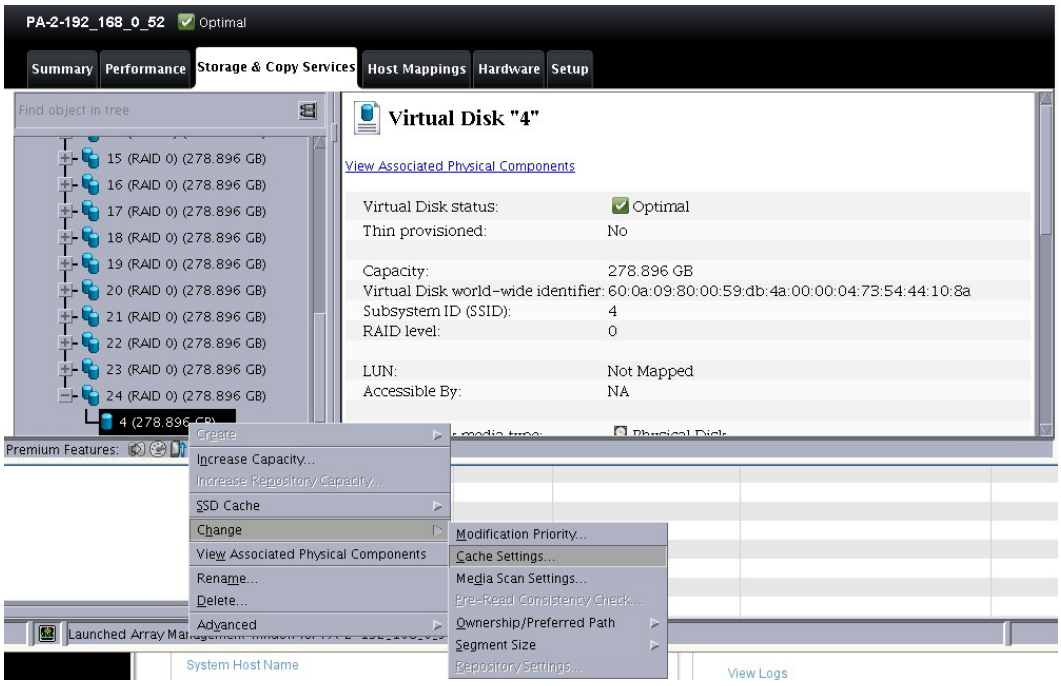

Cache Settings

20 Configure the following cache settings to match exactly as specified in the example above.

For Dell Platform:

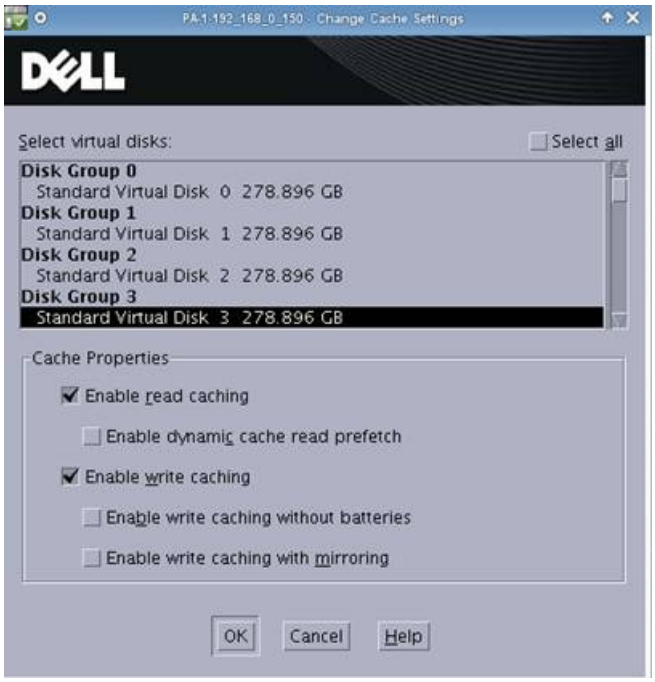

For IBM Platform:

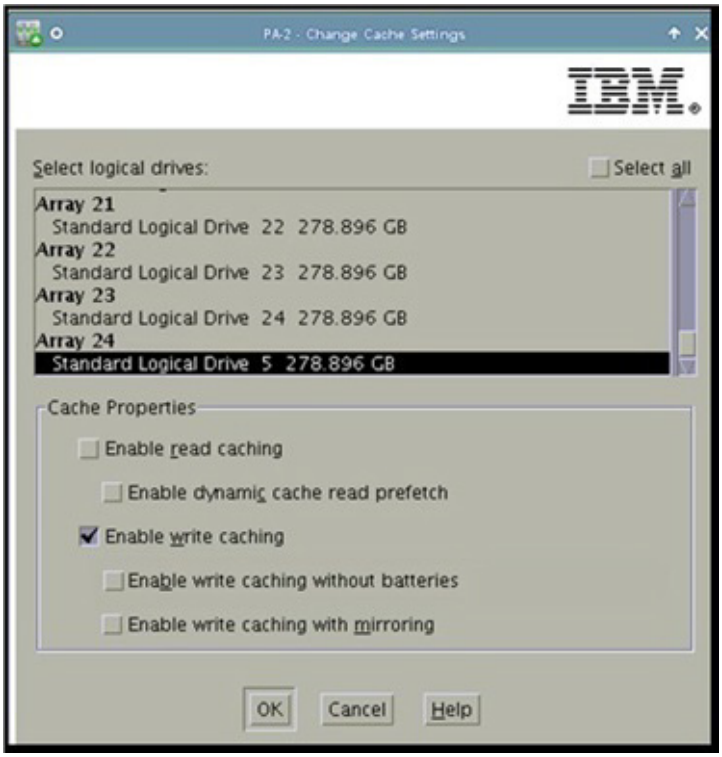

Cache Settings Match

1. From the Host Mappings tab
   - Open the **Undefined Mapping**,
   - Right-click to the **Unmapped LUN**
   - Click **Add Lun Mapping**

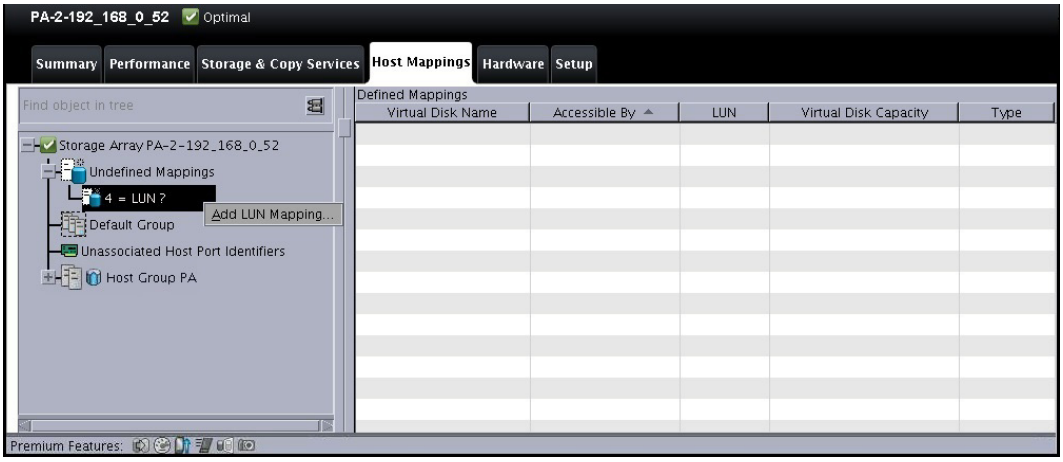

Hosts Mappings tab

2. In Host Group
   - Choose **Host Group PA**
3. In the logical unit number
   - Choose a LUN number that was NEVER assigned to any disk before

Note: Do not choose the same LUN number that was previously assigned to the disk (or any disk).

- Click **Add**,
- Click **Close**

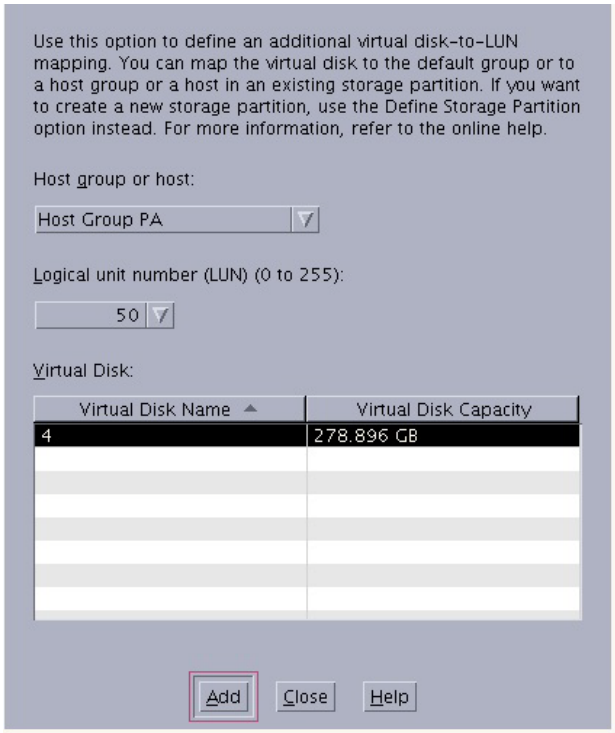

Host Group

4 Using CLI, add the volume as follows:

```
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
```

Chapter 8: Hot Swap HDD Replacement for DAS

The following procedure is intended for the replacement of a hard disk drive using “Hot Swap”.

This procedure requires using the UltraBand CLI commands. For more information about the CLI, see UltraBand CLI Commands in the UltraBand Administration Guide.

Prerequisite

- In the storage array, replace the old hard drive with the new hard drive.

Create the New Virtual Disk (LUN)

Installing the LSI Storage Manager

Note: Download the LSI Storage Manager from: http://www.lsi.com/support.
Use the *WinSCP* to copy the utility file to the server.
On the server installed with ISO 5.6, the utility is located *in /opt/pang/utilities/CE/RAID*

1. From the server, login as **root.**
2. Install the LSI Storage Manager by using the following commands:

```
tar -zxvf 13.04.03.01_Linux(x64)_MSM.tar.gz
cd disk
find . -name "*snmp*" -exec rm -v {} \;
./install.csh
```

3. From the *License Agreement*,
   - Read the *License Agreement*
   - Press **Y** to continue

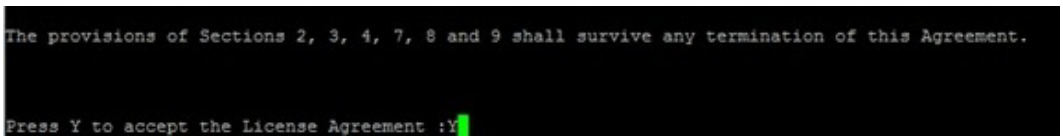

License Agreement

4. From the Setup Menu,
   - Choose **4** from the Setup Type

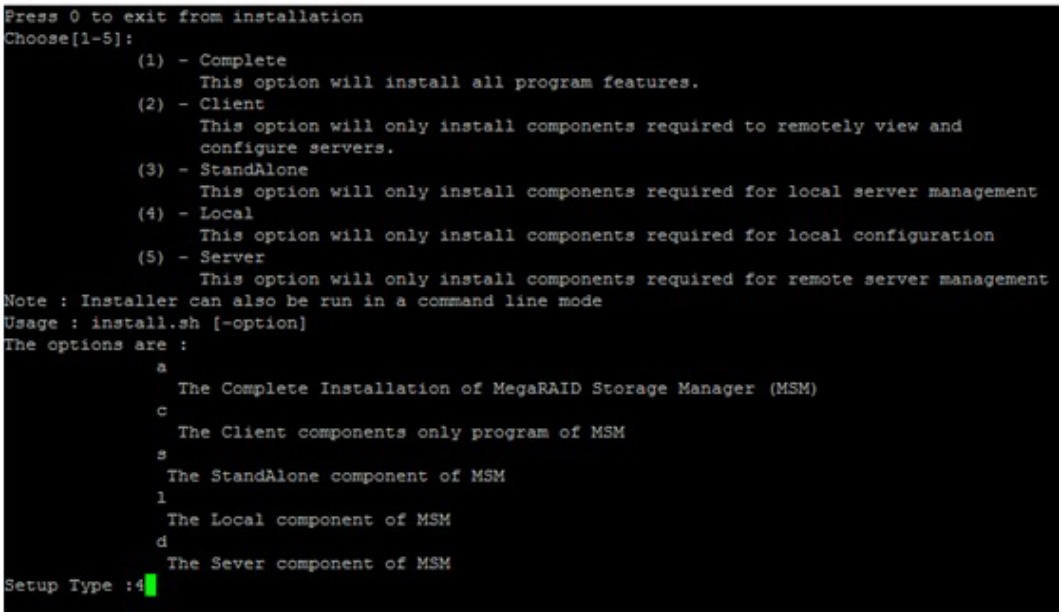

Setup Menu

Wait for the installation to complete.

5 Start the Storage Manager Agent with the following commands:

echo 1 > /proc/sys/vm/overcommit_memory

/etc/init.d/vivaldiframeworkd start

Starting a VNC

6 Start the VNCServer

rm /tmp/.X1-lock
rm /tmp/.X11-unix/X1
vncserver :1

7 Using the VNC Viewer

- Connect to the server using the VNC Viewer

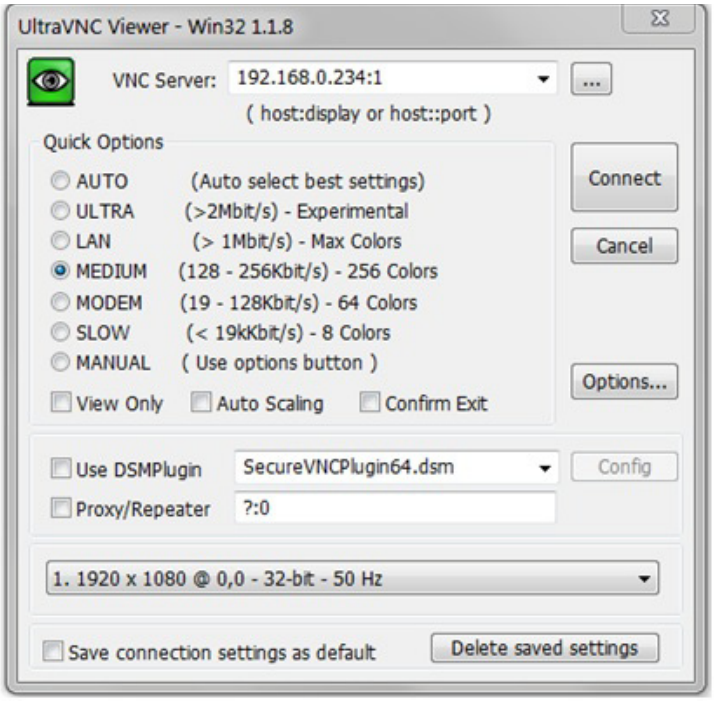

UltraVNC Viewer

8 Start the Megaraid Storage Manager by using the following commands:

```
cd /usr/local/MegaRAID\ Storage\ Manager/
    ./startupui.sh
```

Megaraid Storage Manager

9 From the LSI Login page,

■ Login as root user.

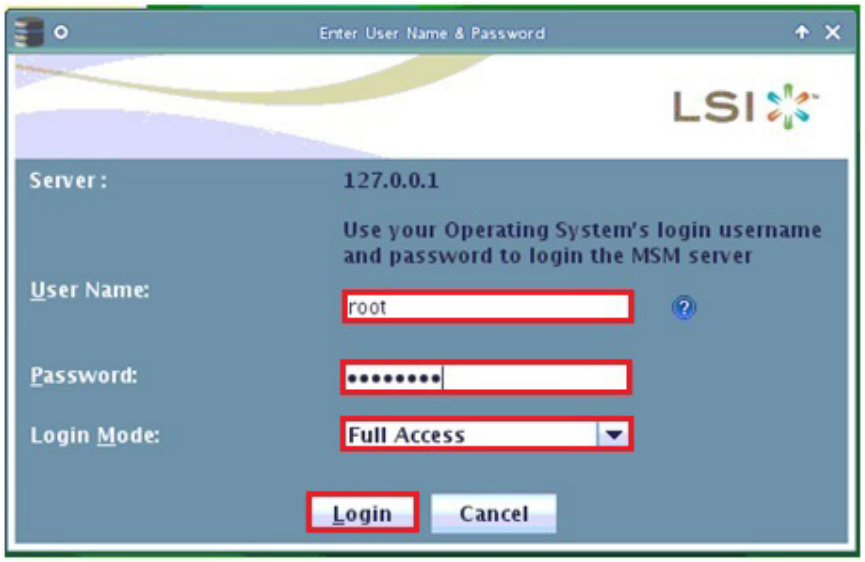

Login LSI

10 Replace the faulty disk (containing the errors) by applying the instructions from the Hardware Vendor.

11 After the disk has been replaced, it will be displayed as an "*Unconfigured Good"* in the Storage Manager "Physical" tab.

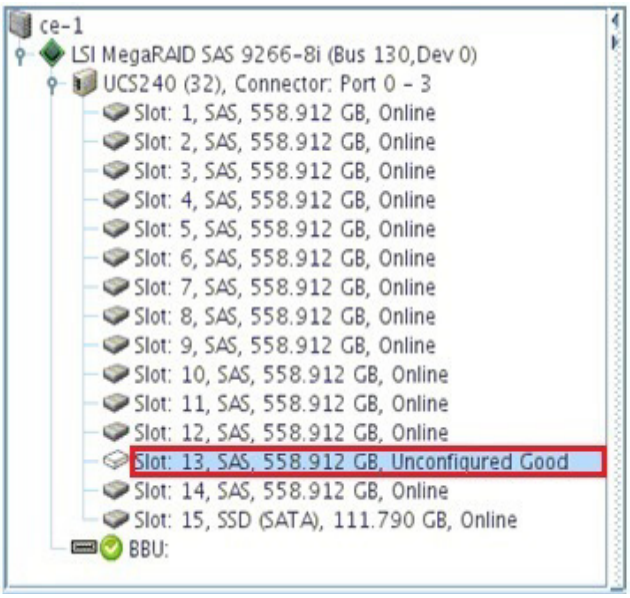

Storage Manager Physical tab

12 Re-create the LUN by right-clicking on the raid controller and choosing **Create Virtual Drive**

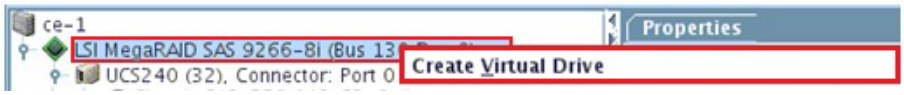

Create Virtue Drive

Note: If the option "*Create Virtual Drive*" is not available, try the following steps:

- Right-click the raid controller and choose **Scan Foreign Configuration** and **clear it**.
- Right-click on the raid controller and choose **Preserved Cache** and clean it (follow the instructions on the screen).
- Check if there are any unnecessary processes running (if there is **-cancel it)**.

13 From the next LSI screen,

- Choose **Advanced**

14 From the LSI Raid screen,

- Select **Raid Level - Raid 0**
- Select the *Unconfigured Drive* and click **Add**

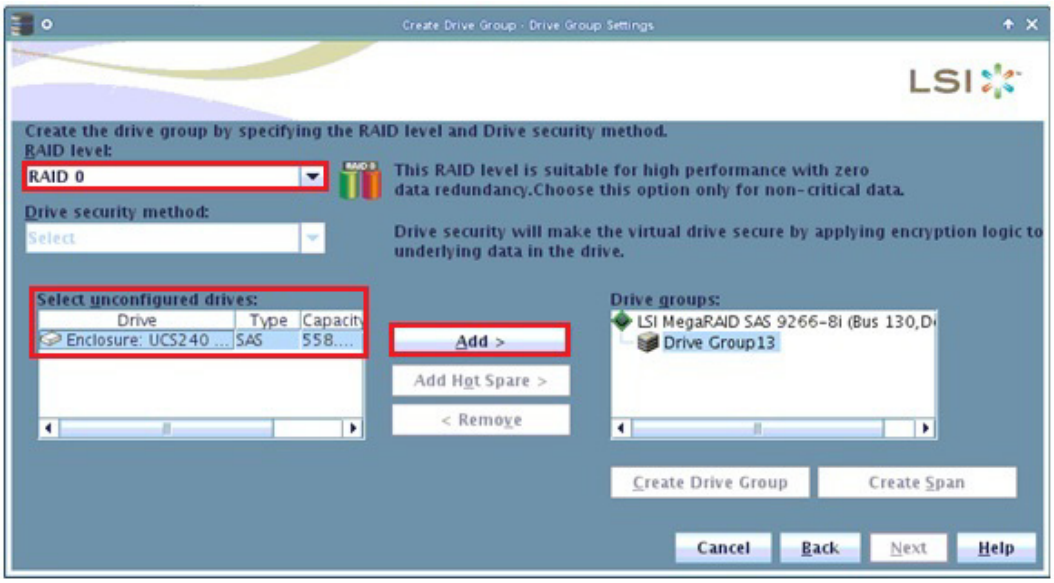

Raid Level - Raid 0

15 Continuing from the LSI screen,

- Click **Create Drive Group**
- Click **Next**.

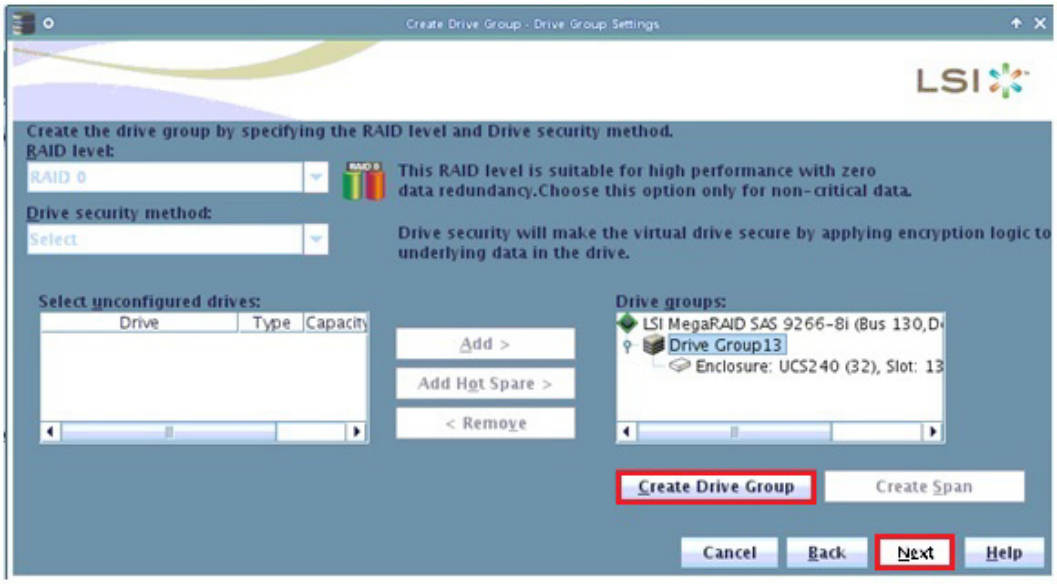

Create Drive Group

16 From the Parameters Configuration screen,

- Configure the parameters **EXACTLY** as specified below: (Except for the following fields: "*Virtual Drive Name*" *and* "*Capacity*")
- Click **Create Virtual Drive**

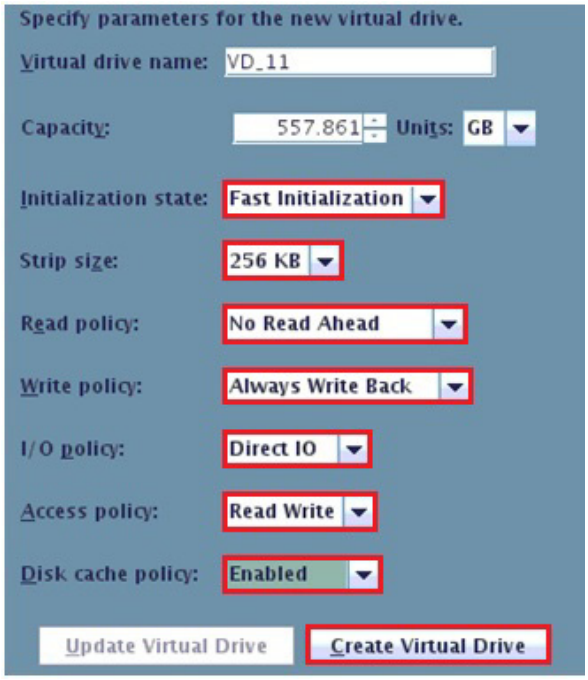

Parameter Configuration

- From the Always Write Back screen,
- Click **Yes**
- Click **Next**

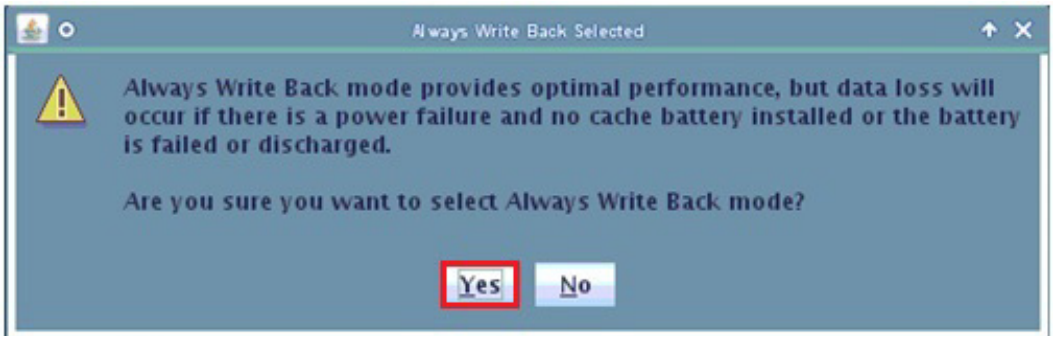

17 From the *LSI Main screen*,

- Click **Finish**

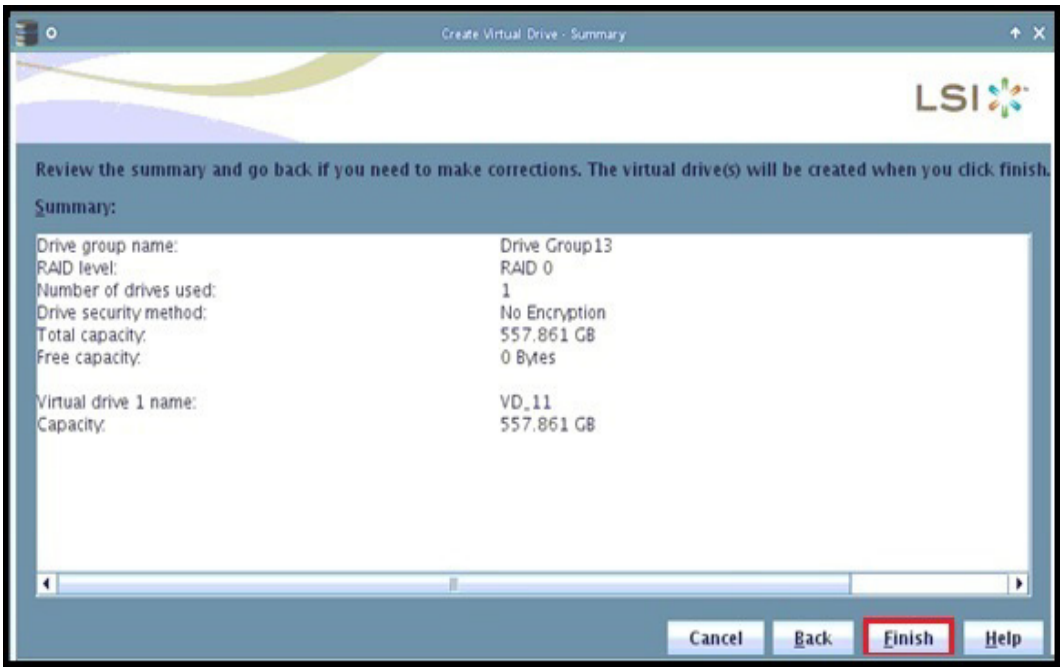

LSI

18 Stop the Storage Manager Agent with the following commands:

```
/etc/init.d/vivaldiframeworkd stop
chkconfig vivaldiframeworkd off
```

echo 2 > /proc/sys/vm/overcommit_memory

IMPORTANT! If the SMagent service exists, you will need to stop it by running the following commands:

```
/etc/init.d/SMagent stop
chkconfig SMagent off
```

19 Find the **volume number** of the drive that failed (using the script */opt/pang/bin/fdish.sh*)

for example:

```
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,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
```

In the example the missing volume is volume #3 because PeerApp,0003,U is missing.

20 The new device should be seen on the Linux as follows:

```
[192.168.0.90]:root@ce-1:~> fdisk -l

Disk /dev/sdb: 119.4 GB, 119453777920 bytes
255 heads, 63 sectors/track, 14522 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes
```

```
Disk identifier: 0x00000000

   Device Boot      Start         End      Blocks   Id  System
/dev/sdb1               1           1        8001   83  Linux
/dev/sdb2               2        2050    16458592+  83  Linux
/dev/sdb3            2051        2573     4200997+  83  Linux
/dev/sdb4            2574       14522    95980342+  83  Linux

Disk /dev/sdc: 299.4 GB, 299439751168 bytes
255 heads, 63 sectors/track, 36404 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes
Disk identifier: 0x96c1d735

   Device Boot      Start         End      Blocks   Id  System
/dev/sdc1               1           1        8001   83  Linux
/dev/sdc2               2        2091    16787925   83  Linux
/dev/sdc3            2092        4181    16787925   83  Linux
/dev/sdc4                        4182                36404          258831247+ 83  Linux


Disk /dev/sda: 146.1 GB, 146163105792 bytes
255 heads, 63 sectors/track, 17769 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes
Disk identifier: 0x00045b2d

   Device Boot      Start         End      Blocks   Id  System
/dev/sda1               1         262     2104483+  82  Linux swap / Solaris
/dev/sda2   *         263        5483    41937682+  83  Linux
/dev/sda3            5484       10051    36692460   83  Linux
/dev/sda4           10052       17769    61994835   83  Linux
```

(*Single Volume*)

```
Disk /dev/sdd: 299.4 GB, 299439751168 bytes
255 heads, 63 sectors/track, 36404 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes
Disk identifier: 0x00000000
Disk /dev/sdd doesn't contain a valid partition table
```

21 Format the volume using the "*single volume*" option with the correct **device name** and the ***volume number*** needed:

```
ce-1:/opt/pang/useful # ./format_disks.sh -format_one /dev/sdd 3
format_disks.sh: 4.4 finished, see /opt/pang/useful/installog.txt for details
```

22 You should now be able to see the new formatted volume:

```
[192.168.0.90]:root@ce-1:/opt/pang/useful> ../bin/fdisk.sh
PeerApp,0001,U    /dev/sdb1
PeerApp,0002,U    /dev/sdc1
PeerApp,0003,U    /dev/sdd1
```

23 Run the following command to remove the old volume from the PANG volume table:

Note: The last number is the volume id

```
[192.168.0.90]:root@ce-1:/opt/pang/bin> ./cmdbutils_rt -I lo -e 3
getopt returned c=I optarg= lo
```

```
getopt returned c=e optarg= 3
14549 pang
spread_interface lo is in ip 127.0.0.1
connect to spread name 4803@127.0.0.1, with private group #cmdb_rt#ce-1
DEBUGG: m_system_update->init() rv=0
delete_volume
delete Volume 3
Deleting Volume 3

Bye.
```

24 Now for the fun part. Run this command to insert the new volume to the PANG volume table:

Note: The last number is the volume id.

```
[192.168.0.90]:root@ce-1:/opt/pang/bin> ./cmdbutils_rt  -I lo -b 3
getopt returned c=I optarg= lo
getopt returned c=b optarg= 3
15408 pang
spread_interface lo is in ip 127.0.0.1
connect to spread name 4803@127.0.0.1, with private group #cmdb_rt#ce-1
DEBUGG: m_system_update->init() rv=0
insert_volume
insert Volume 3
Bye.
```

25 Completed! You can see that the volume exists and is available on the CLI as mounted:

```
console# show volumes
Licensed volumes : 15
Volume name       State          Owner
/mnt/vol1         mounted        ce-1
/mnt/vol2         mounted        ce-1
/mnt/vol3         mounted        ce-1
….
```

Note: You can see the attached pang.log file to understand the mounting procedure. You can use it to check the log if you need to verify your results.

Chapter 9: UltraBand Clear Caching Content

Overview

The following chapter details the procedure for clearing all cache content in a system.

IMPORTANT! *The following procedure entails formatting the storage and deleting the database.*

Grid

Clearing the Cache Content for a Grid platform:

1. Stop the UltraBand service using the following:
   a. From the CLI, Stop the service using the command *"oper service stop"*.
   b. From the CLI, make sure the service is stopped using the command *"show status"*.
2. From the *Management*, login as a root user and run the Configure Storages Script using the commands:

   cd /opt/pang/useful/configure_storages/
   ./configure_storages.py -s <the number of the storages in the system> -b <the number of the CEs in the system>
   For example, for a system with 8 cache engines and 3 storages:
   ./configure_storages.py -s 3 -b 8

   A warning message about the disk format will be shown, to continue the storages configuration enter y.

   ********** WARNING! **********
   The selected storage is about to be reconfigured. All data on that storage
   will be lost. You'll need to format disks manually.
   Do you want to continue? (y/n) y

   Wait until the configure_storages.py script is completed
3. Using the CLI, start the service using the "oper service start" command.

Standalone

Clearing the Cache Content for a Standalone platform:

1. Stop the UltraBand service using the following:

a. From the CLI, Stop the service using the command *"oper service stop"*.

b. From the CLI, make sure the service is stopped using the command *"show status"*.

2. Login as a root user

3. Run the Format Disk script using the commands:

```
cd /opt/pang/useful/
./format_disks.sh -format_all
```

Wait until the format_disks.sh is completed

4. Run the Format SSD script using the commands:

```
cd /opt/pang/bin/
./ssd_auto_install.sh -f
```

Wait until the ssd_auto_install.sh script is completed

5. Using the CLI, start the service using the "oper service start" command.

Chapter 10: Upgrade Procedure from 5.6 to 5.7 for UB5000

The following procedure is intended for upgrading version 5.6 to version 5.7 for UB5000.

IMPORTANT! If you are using a version lower than 5.6, you must first upgrade the system to 5.6 before upgrading to 5.7.

Prerequisites

1. Make sure the previous software version file still exists on /tftpboot in case of Rollback requirement.
2. Optionally, it's recommended to disable the traffic redirection towards the UltraBand system. However, the UltraBand solution capable to forward the live traffic without affecting it, using the bypass mechanism. The bypass mechanism takes an effect immediately when the caching service is stopped

Upgrading the UltraBand Software

To Upgrade the UltraBand Software:

1. Upgrade the system to version 5.7 by running the following CLI commands:
   - **Grid: `upgrade all 127.0.0.1 filename`”**.
   - **Standalone: `upgrade 127.0.0.1 filename`**”.

   Note: The filename is the “.tar.gz” upgrade package.

The output pattern for the Grid upgrade should resemble the following:

```
console# upgrade all 127.0.0.1 PeerApp_GA_5.7.0b114_grid.tar.gz

Log file: /var/log/peerapp/GA_installer_14-10-20-09.59.20.log


Checking for network connectivity...
Contacting ce-2 machine ...
Ok
Contacting ce-1 machine ...
Ok

Testing for a need to backup existing configuration...
```

Detected major version upgrade (to 5.7)
Upgrading to version 5.7 requires a new license file to be activated on the platform once the upgrade process is over.
Without this 5.7 license file, the platform will NOT start.
If you do not have a new license file for version 5.7 - please do not proceed with the upgrade !!
Do you want to proceed with the upgrade [Y/N]Y ?
Are you sure? [Y/N]Y
Please provide license file location information (on an accessible TFTP server):
TFTP server address:
127.0.0.1
Path to 5.7 license file:
PALicense-5.7.xml
File found , validating ...
Valid 5.7 license found , proceeding with this upgrade...
Checking free space for upgrade..
Ok
Checking enough available space on /var for upgrade..
Ok
Install management software might disconnect current session and require reconnect.
Continue installation of management software? [Y/n]y
Will install snmp management
Checking for valid chassis id ...
Chassis id test passed ok

Installing SPREAD environment
Installing SPREAD files
Installing SNMP agent (+PeerApp subagent)
Upgrade net-snmp packages
Preparing packages for installation...
libsensors4-3.0.3-3.6
python-tk-2.6.0-2.19
snmp-mibs-5.7.2-9.1.2
libsnmp30-5.7.2-9.1.2
OpenIPMI-2.0.16-7.1
Updating etc/sysconfig/ipmi...
perl-SNMP-5.7.2-9.1.2
net-snmp-5.7.2-9.1.2
Updating etc/sysconfig/net-snmp...
Stopping any running daemons ..
Installing SNMP agent in /etc/inittab
Installing SNMP trap daemon in /etc/inittab
Installing smartfilter crontab...
starting bpctl
Installing dnsmasq (replacing nscd for dns only)
Starting name service masq caching server - Warning: dnsmasq already running! ..done
Failed installing analyzer files ...
Installing monitoring daemon in crontab
Installing Serial Over Lan
*** Serial Over Lan cannot be autoconfigured on this platform. Make sure it's configured correctly. ***
Starting SOL ..done
Upgrade existing web installation...

Web installation finished successfully
installing policy manager
Backing up PolicyManager ver. 5.7.0.37 database...
Installing tacacs support
Succeeded
Starting upgrade on ce-2..
Checking for valid chassis id ...
Chassis id test passed ok

Stopping watchdog daemons ..

Done ...
Stopping running daemons ..
Installing SPREAD environment
Installing SPREAD files
Installing SNMP agent (+PeerApp subagent)
Stopping any running daemons ..
Installing SNMP agent in /etc/inittab
starting bpctl
Installing Storage Monitor
Cannot find last driver!
Upgrading caching engine ...
Backing up existing configuration
Done
Installing File system Snapshot at Boot
Shutting down irqbalance ..done
Cannot update snmp links!
Installing watchdogs...
Done...

Succeeded

Starting upgrade on ce-1..
Checking for valid chassis id ...
Chassis id test passed ok

Stopping watchdog daemons ..
Done ...
Stopping running daemons ..
Installing SPREAD environment
Installing SPREAD files
Installing SNMP agent (+PeerApp subagent)
Stopping any running daemons ..
Installing SNMP agent in /etc/inittab
starting bpctl
Installing Storage Monitor
Cannot find last driver!
Upgrading caching engine ...
Backing up existing configuration
Done
Installing File system Snapshot at Boot
Shutting down irqbalance ..done
Cannot update snmp links!
Installing watchdogs...
Done...

Succeeded

Starting database and configuration backup , please hold on...
All done
Please remember to import and activate the 5.7 license before starting service.
console#

The output pattern for the Standalone upgrade should resemble the following:

console# upgrade 127.0.0.1 PeerApp_GA_5.7.0b113_standalone.tar.gz

Log file: /var/log/peerapp/GA_installer_14-10-20-08.35.55.log

Testing for a need to backup existing configuration...
Detected major version upgrade (to 5.7)
Upgrading to version 5.7 requires a new license file to be activated on the platform once the upgrade process is over.
Without this 5.7 license file, the platform will NOT start.
If you do not have a new license file for version 5.7 - please do not proceed with the upgrade !!
Do you want to proceed with the upgrade [Y/N]Y ?
Are you sure? [Y/N]Y
Please provide license file location information (on an accessible TFTP server):
TFTP server address:
127.0.0.1
Path to 5.7 license file:
PALicense.xml
File found , validating ...
Valid 5.7 license found , proceeding with this upgrade...
Checking free space for upgrade..
Ok
Checking enough available space on /var for upgrade..
Ok
Starting system installation
Starting software installation...

Doing pre-installation sanity tests
Storage consistency OK
Done
Installing SPREAD environment
Installing SPREAD files
Shutting down CRON daemon..done
Starting CRON daemon..done

Installing SNMP agent (+PeerApp subagent)
Upgrade net-snmp packages
Preparing packages for installation...
libsensors4-3.0.3-3.6
python-tk-2.6.0-2.19
snmp-mibs-5.7.2-9.1.2
libsnmp30-5.7.2-9.1.2
OpenIPMI-2.0.16-7.1
Updating etc/sysconfig/ipmi...
perl-SNMP-5.7.2-9.1.2
net-snmp-5.7.2-9.1.2
Updating etc/sysconfig/net-snmp...
Backing up old logs
Stopping any running daemons ..
Installing SNMP agent in /etc/inittab
Installing SNMP trap daemon in /etc/inittab
Installing smartfilter crontab...
Installing monitoring daemon in crontab
Installing dnsmasq (replacing nscd for dns only)
Starting name service masq caching server - Warning: dnsmasq already running! ..done

Installing Analyzer
Backing up analyzer configuration ...
Failed installing analyzer files ...

Upgrade existing web installation...

Web installation finished successfully
installing policy manager
Backing up PolicyManager ver. 5.7.0.37 database...

```
Installing tacacs support
Succeeded
Starting upgrade...
Starting database and configuration backup , please hold on...
All done
Starting software installation...

Doing pre-installation sanity tests
Storage consistency OK
Done
Stopping watchdog daemons ..
Installing SPREAD environment
Installing SPREAD files
Shutting down CRON daemon..done
Starting CRON daemon..done

Installing SNMP agent (+PeerApp subagent)
Backing up old logs
Stopping any running daemons ..
Installing SNMP agent in /etc/inittab
Installing monitoring daemon in crontab
Installing standalone caching engine...
Upgrading caching engine ...
Installing Storage Monitor
starting bpctl
insmod /lib/modules/2.6.27.19-llpf_10-5-default/kernel/drivers/net/bpctl_mod.ko
Done
Installing File system Snapshot at Boot
Shutting down irqbalance ..done
Shutting down CRON daemon..done
Starting CRON daemon..done


Succeeded



enable patrol on raid controllers
Please remember to import and activate the 5.7 license before starting service.
```

Post upgrade

1. Verify that the version was installed properly by using the command "**`show version`**".

   Note: The user must login again to access the CLI prompt.

2. Import and activate the 5.7 license using the following CLI commands:

   license import 127.0.0.1 ***.xml

   license activate

3. From the enable mode, start PANG by using the "`oper service start`" command.
4. Enable the traffic redirection (if it was disabled before the upgrade).

   The upgrade procedure is completed.

IMPORTANT! Only if the upgrade fails, use The Rollback Procedure from 5.7 to 5.6 for UB5000.

Chapter 11: The Rollback Procedure from 5.7 to 5.6 for UB5000

IMPORTANT! This chapter is **Only** used in case that the upgrade procedure from 5.6 to 5.7 for UB5000 fails.

The UltraBand Rollback Procedure:

1 From the Management Server, run:

/opt/pang/mgmt/config/downgrade_proc.sh

Note: If the rollback failed with the following message: *"Not found suitable backup files, cannot restore version 5.6 on this system. Please refer to the manual downgrade procedure or contact support."* **Stop the UltraBand Rollback Procedure** and See chapter 12 to perform the manual rollback from 5.7 to 5.6.

2 Upgrade the system to the version 5.6, by running either of the following CLI commands:

- **Grid: `upgrade all 127.0.0.1 filename`"**.
- **Standalone: `upgrade 127.0.0.1 filename`**".

Note: The filename is the ".tar.gz" upgrade package.

3 Verify the version is installed properly using the command "show version". You must be logged in to access the CLI prompt for this action.

4 From the enable mode, start PANG by using the "`oper service start`" command.

5 Enable the traffic redirection (if it was disabled before the procedure).

Chapter 12: Manual UltraBand Rollback Procedure from 5.7 to 5.6 for UB5000

IMPORTANT! Use the following procedure **only** if the UltraBand Rollback Procedure has failed with the message, *"Not found suitable backup files, cannot restore version 5.6 on this system. Please refer to the manual downgrade procedure or contact support"*.

IMPORTANT! Running the Manual Rollback Procedure will remove all existing cache data storage and version 5.7 cmbd (the storage will be formatted).

How to Manually Rollback the UltraBand from 5.7 to 5.6:

1. Stop the UltraBand service by running:
   - `oper service stop`
2. Login to management server as root.
3. Clean up storage and local disks from any 5.7 cmdb copies, using the following commands:

Standalone Server:

- `/opt/pang/mgmt/config/wipe_old_database.sh`

Grid:

- `su - install`
- `scp /opt/pang/mgmt/config/wipe_old_database.sh admin@ce-1:~`
- `ssh -l admin ce-1 'sudo ~/wipe_old_database.sh'`
- `exit`

4. Format the storage, using the following command:

Standalone server:

- `/opt/pang/useful/format_disks.sh -format_all`
- `/opt/pang/bin/ssd_auto_install.sh -f`

Grid:

- `cd /opt/pang/useful/configure_storages`
- `./configure_storages.py -s` <the number of the storages in the system> `-b` <the number of the CEs in the system>

  For example, for a system with 8 cache engines and 3 storages:
  `./configure_storages.py -s 3 -b 8`

A warning message about the disk format will be shown, to continue the storages configuration enter y.

********** WARNING! **********
The selected storage is about to be reconfigured. All data on that storage
will be lost. You'll need to format disks manually.
Do you want to continue? (y/n) y

5 Restore previous configurations, using the following commands:

- `cd /opt/5.6_backup`
- `cp -vf *.xml *.xsd *.xsl *.conf /opt/pang/mgmt/config &>/dev/null`

Note: Please pay attention: any changes done to the cluster configuration during a period when 5.7 was installed - **will be lost**.

6 Uninstall *dnsmasq* and replace it with *nscd,* using the following commands:

- `service dnsmasq stop &> /dev/null`
- `/sbin/chkconfig dnsmasq off &> /dev/null`
- `cat /etc/dnsmasq.conf|grep "^server=" |sed 's/server=/nameserver /g' > /etc/resolv.conf`
- `rpm -e dnsmasq &> /dev/null`
- `sed -i 's/^#nsc/nsc/' /etc/inittab`
- `init q`
- `/sbin/chkconfig nscd on &>/dev/null`

7 Rename the backup directory, using the following command:

- `cp -r /opt/5.6_backup /opt/5.6_backup.bak &>/dev/null`

8 Upgrade the system to the version 5.6, by running either of the following CLI commands:

- **Grid: `upgrade all 127.0.0.1 filename`"**.
- **Standalone: `upgrade 127.0.0.1 filename`**".

Note: The filename is the ".tar.gz" upgrade package.

9 Verify the version is installed properly using the command "show version". You must be logged in to access the CLI prompt for this action.

10 From the enable mode, start PANG by using the "`oper service start`" command.

11 Enable the traffic redirection (if it was disabled before the procedure).

Chapter 13: UBWeb Installation

The UBWeb is an additional tool integral to the UBView which was developed by PeerApp. The UBWeb is a web monitoring application for both Standalone and Grid configurations. The UBWeb uses graphs and statistics to help you monitor the UltraBand HTTP and P2P caching performance.

UBWeb provides an immediate presentation of the system overall operational performance, logical and cache statuses, and information regarding browsing, speed, time and activity.

Prerequisites

Pre-Installation Verification

In order for the UBWeb to be properly installed, the user must verify that the system contains the following:

1. A minimum of 24 GB RAM for each cache engine (including standalone).
2. Installing the required SSD drives compatible to the UltraBand product model.
3. A request for a new license file for the specific platform which includes the UBWeb.

   Make sure the expected amount of cache out is not limited by the cache out license.

Installation

Note: Using Standalone and Grid-Ready platform, Management and Cache Engine refer to the same server.

IMPORTANT! For Grid, repeat the following procedure on each Cache Engine.

Installing the LSI Storage Manager on the Cache Engine:

Note: Download the LSI Storage Manager from: http://www.lsi.com/support. First use the *WinSCP* to copy the utility file to the Management server, and then copy the file to the Cache Engine using the SCP command. On the server installed with ISO 5.7, the utility is located in */opt/pang/utilities/CE/RAID*.

1. From the Cache Engine, login as **root.**
2. Install the LSI Storage Manager by using the following commands:
   - *tar -zxvf 13.04.03.01_Linux\(x64\)_MSM.tar.gz*
   - *cd disk*

- *find . -name "*snmp*" -exec rm -v {} \;*
- *./install.csh*

3 From the *License Agreement*,

- Read the *License Agreement*
- Press **Y** to continue

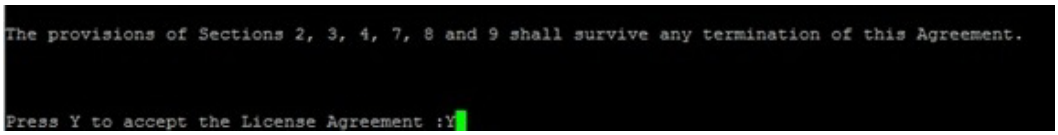

4 From the Setup Menu,

- Choose 4 (Setup Type)

Wait for the installation to complete.

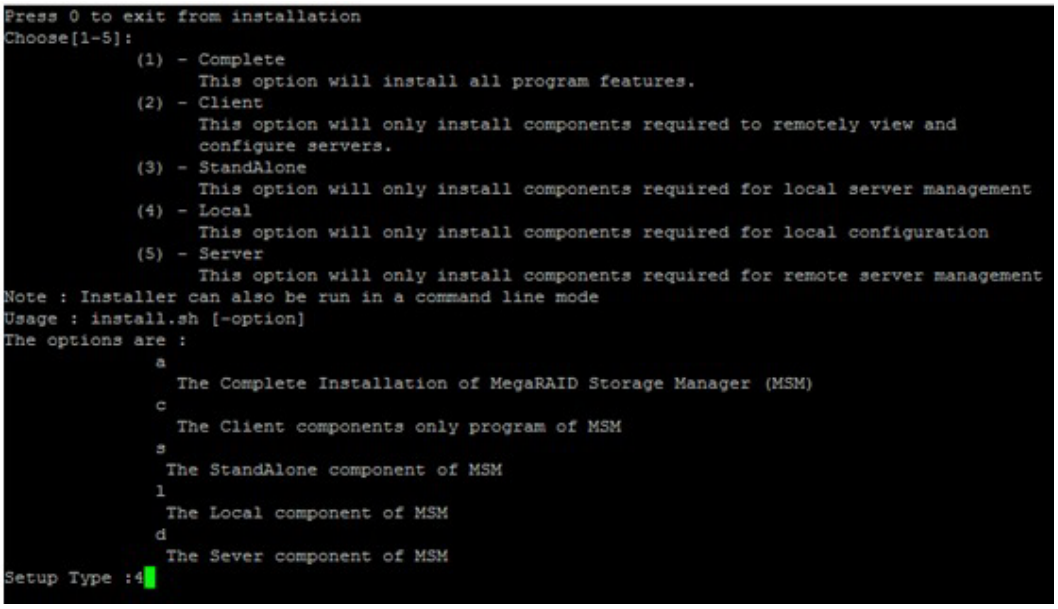

5 Start the Storage Manager Agent with the following commands:

- *echo 1 > /proc/sys/vm/overcommit_memory*
- */etc/init.d/vivaldiframeworkd start*

Starting a VNC

1 From the Management Server,

Start the VNCServer using the commands:

rm /tmp/.X1-lock
rm /tmp/.X11-unix/X1
vncserver :1

2 Using the VNC Viewer

- Connect to the server using the VNC Viewer

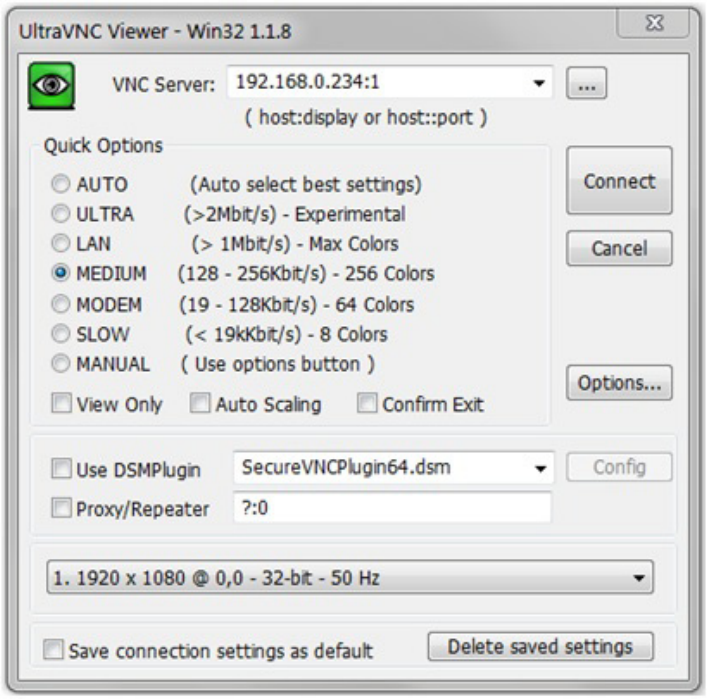

IMPORTANT! For the Grid platform, connect to the Cache Engine using ssh -X
For example: `ssh -X ce-1`

3 Start the Megaraid Storage Manager on the Cache Engine using the following commands:

- `cd /usr/local/MegaRAID\ Storage\ Manager/`
- `./startupui.sh`

4 From the LSI Login page,

- Login as root user.

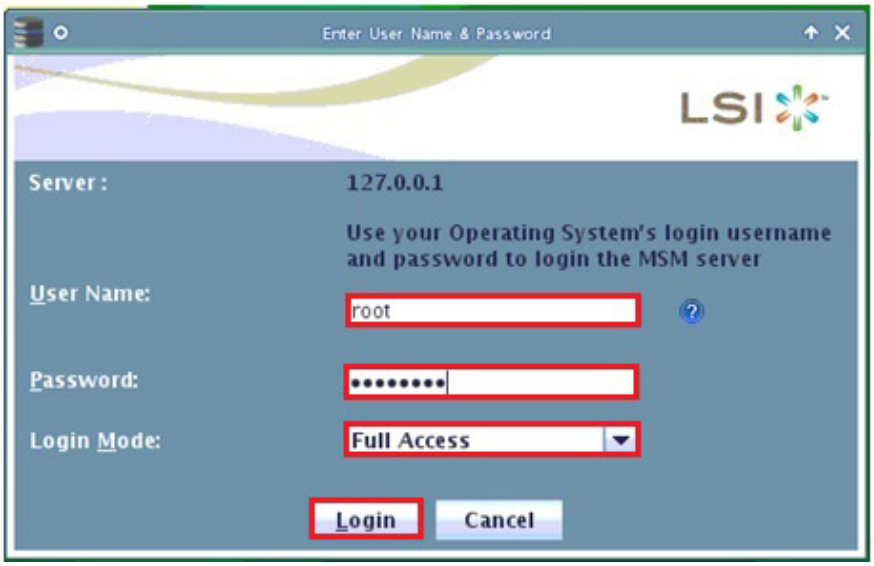

5 After the SSD disk has been physically added, it is displayed in the Storage Manager "Physical" tab as "*Unconfigured Good"*.

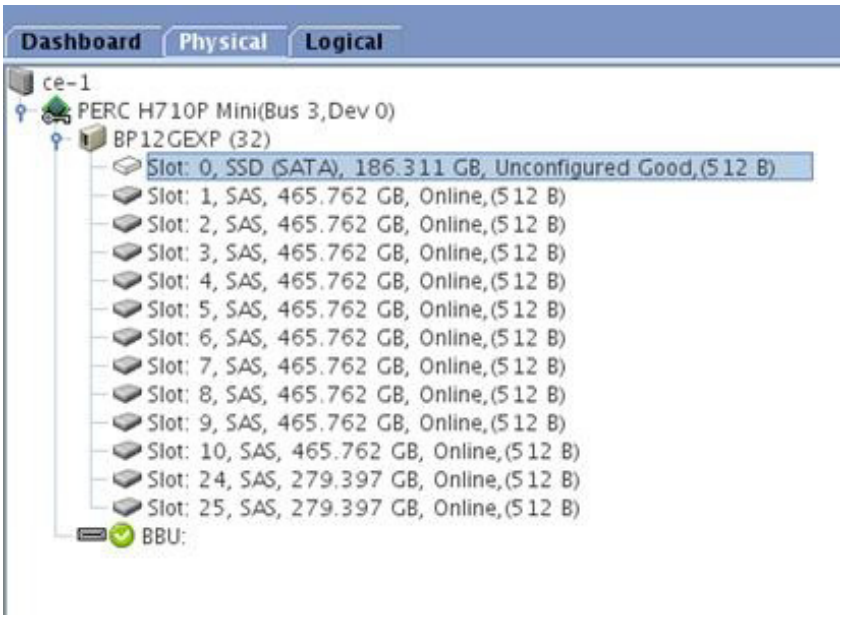

6 Create the LUN by right-clicking on the raid controller and choose **Create Virtual Drive.**

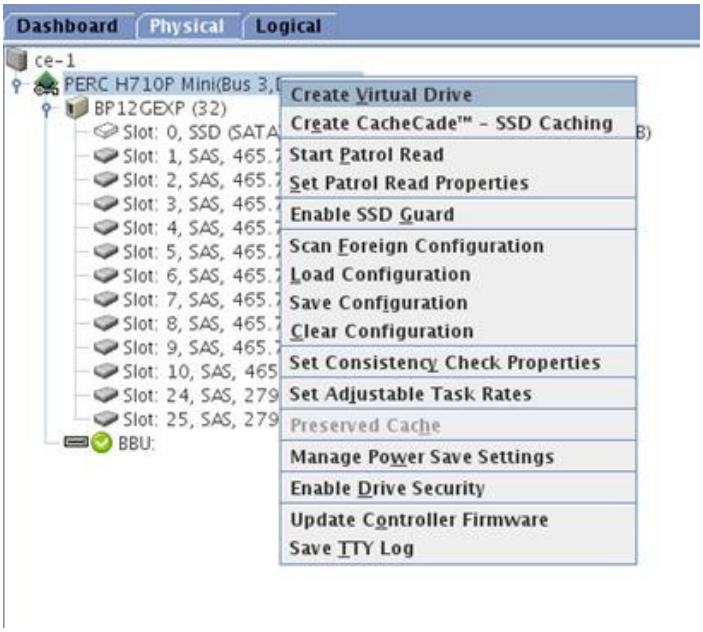

Notes: If the *Create Virtual Drive* option is not available, use the following steps:

- Right-click the raid controller and choose **Scan Foreign Configuration** and clear it.
- Right-click on the raid controller and choose **Preserved Cache** and clean it (follow the instructions on the screen).
- Check if there are any unnecessary processes running (if there is -cancel it).

7 From the LSI screen,

- Choose **Advanced**

8 From the LSI Raid screen,

- Select **Raid Level - Raid 0**
- Select the **Unconfigured Drive** and click **Add**

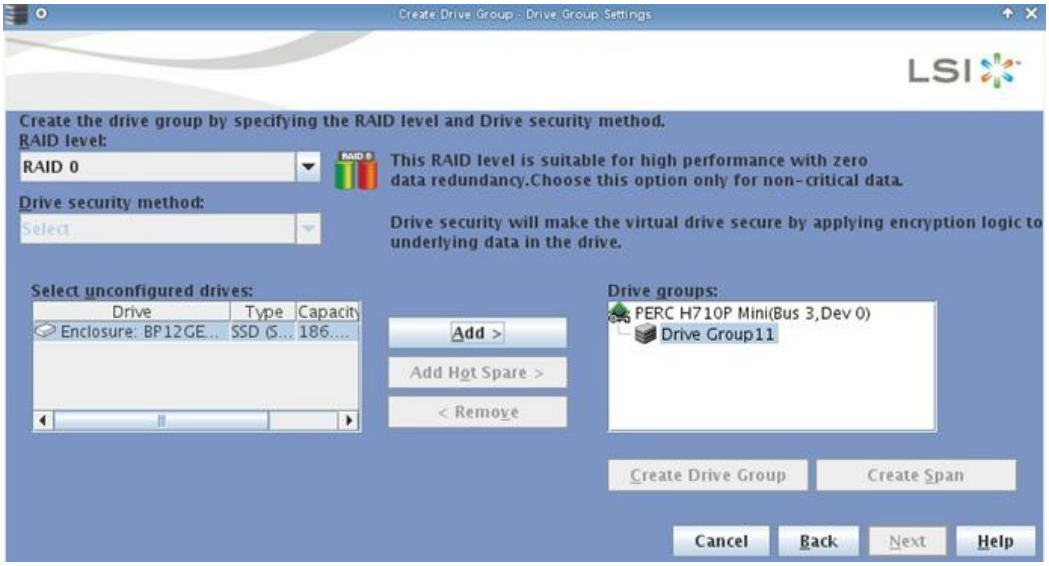

9 From the LSI screen,

- Click **Create Drive Group**
- Click **Next**

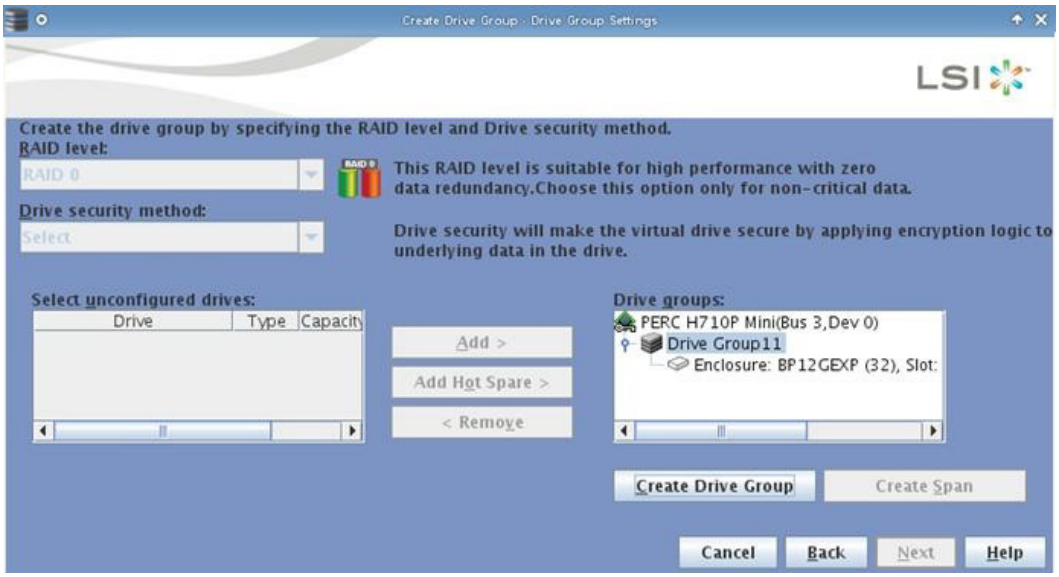

10 From the Parameters Configuration screen,

- Configure the parameters **EXACTLY** as specified below: (Except for the fields: *"Virtual Drive Name" and "Capacity"*)
- Click **Create Virtual Drive**

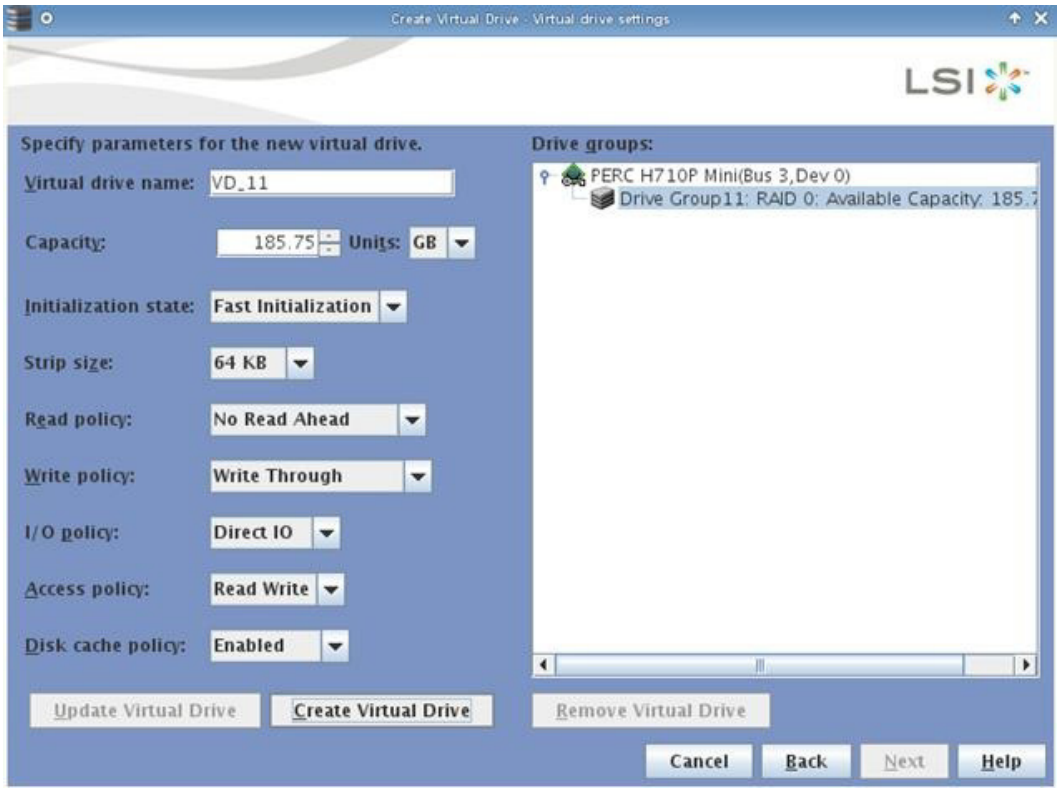

11 From the Write Through Selected screen,

- Click **Yes**

- Click **Next**

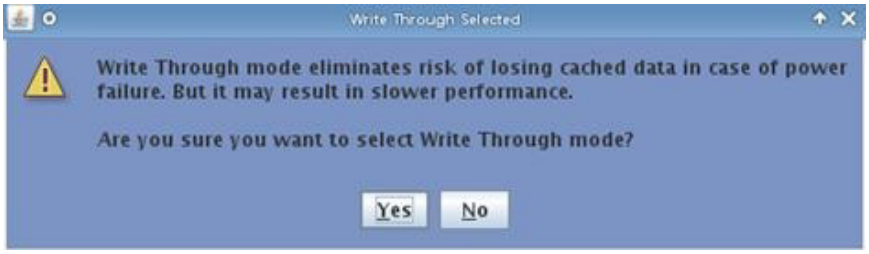

12 From the LSI Main screen,

- Click **Finish**

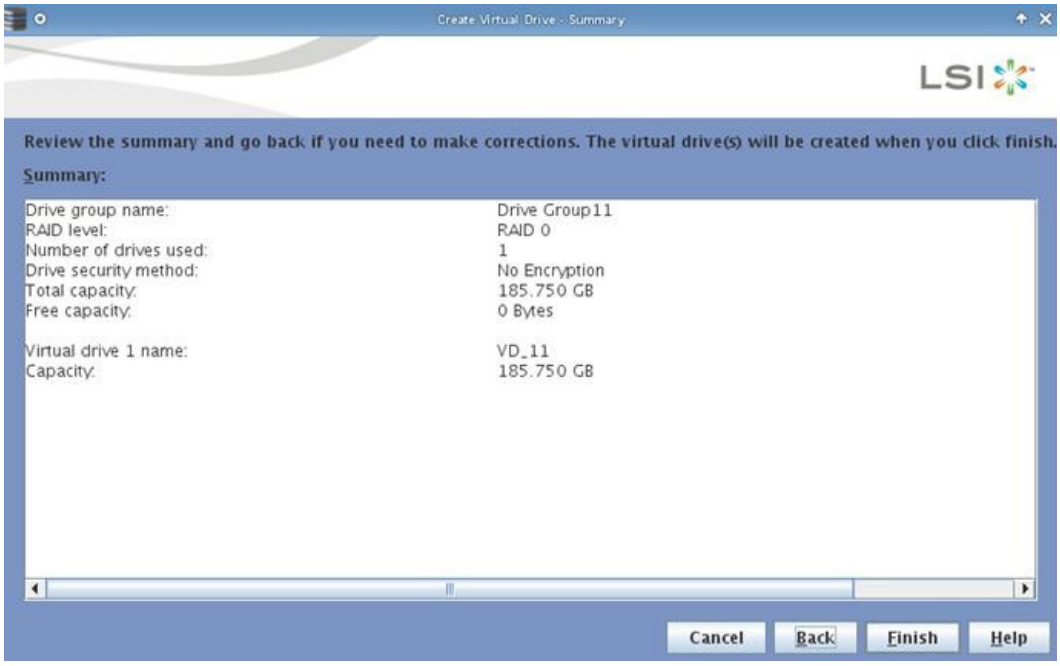

13 Stop the Storage Manager Agent with the following commands:

- `/etc/init.d/vivaldiframeworkd stop`
- `chkconfig vivaldiframeworkd off`
- `echo 2 > /proc/sys/vm/overcommit_memory`

14 Format the SSD device using the command:

- `/opt/pang/bin/ssd_auto_install.sh -f`

15 Verify using the following command: `/opt/pang/bin/fdisk.sh:`

The following, *PeerSSD,024x,u* is displayed below:

IMPORTANT! Before continuing with the Post Installation, for Grid platforms, repeat the Installation procedure on each Cache Engine.

Post Installation

1. From the management, verify the cluster_conf.xml that *enable-web-cache* (protocol) is enabled. The value must be 1 (displayed below)

Cluster_conf.xml:

<protocols>
<enable-bittorent>0</enable-bittorent>
enable-edk>1</enable-edk>
<enable-http>1</enable-http>
<enable-ares>1</enable-ares>
<enable-smartfilter>1</enable-smartfilter>
<enable-web-cache>1</enable-web-cache>
<enable_cache_out_port>1</enable_cache_out_port>
</protocols>

If the value is 0

OR

There is no *enable-web-cache* at the protocols,

- Edit the cluster_conf.xml and insert the changes,
- Import and apply the new configuration via CLI

2. Import and activate the new License via CLI
3. Restart the service via CLI using the commands:
   - oper service stop
   - oper service start

Chapter 14: Upgrade Procedure to 5.7 for UB6000 and above

The following procedure is intended for checking the Firmware versions and upgrading both, the Operation System and UltraBand Software to version 5.7.

IMPORTANT! If you are using a version lower than 5.6, you must first upgrade the system to 5.6 before upgrading to 5.7.

Note: In order to update the Firmware, make sure that you have the Dell Server Update Utility version:14.12.200.69 before starting the procedure. The Utility can be found at http://www.dell.com/support/home/us/en/19/Drivers/DriversDetails?driverId=9CXFR

The process includes the following:

- Checking the Firmware version
- Operation System Upgrade
- Software Upgrade

IMPORTANT! There is no rollback to the old Firmware, Operation System or Software.

Pre-Requisites

It's recommended to disable the traffic redirection towards the UltraBand system.

IMPORTANT! An important thing to remember is that the UltraBand solution is capable to forward the live traffic without affecting it by using the bypass mechanism. The bypass mechanism is automatic when the cache service is stopped.

Note: In order to save backup time and space in the upgrade process, it is recommended to remove irrelevant files from the /tftpboot directory.

1. Check that the grid has stopped.
2. Check that you have a valid license.

Upgrading the Firmwares, OS and UltraBand Software

Downloading the Files:

1 Download to the /tmp directory the following files:

Installer-ISO-5.7.0bx-Application-5.7.0by.iso
Installer-ISO-5.7.0bx-Application-5.7.0by.iso.md5
prepare_iso.sh

2 Prepare for installation by executing:

cd /tmp
./prepare_iso.sh Installer-ISO-5.7.0bx-Application-5.7.0by.iso Installer-ISO-5.7.0bx-Application-5.7.0by.iso.md5

3 In order to survive network disconnect, run the command:

screen

Note: During the installation process, if the SSH session disconnects, reconnect to the server via SSH and run the command:
screen -x

4 Load the Upgrade Tool by running the command:

/opt/iso_upgrade/install-tools/upgrade.py

Grid Upgrade:

The following menu options are displayed:

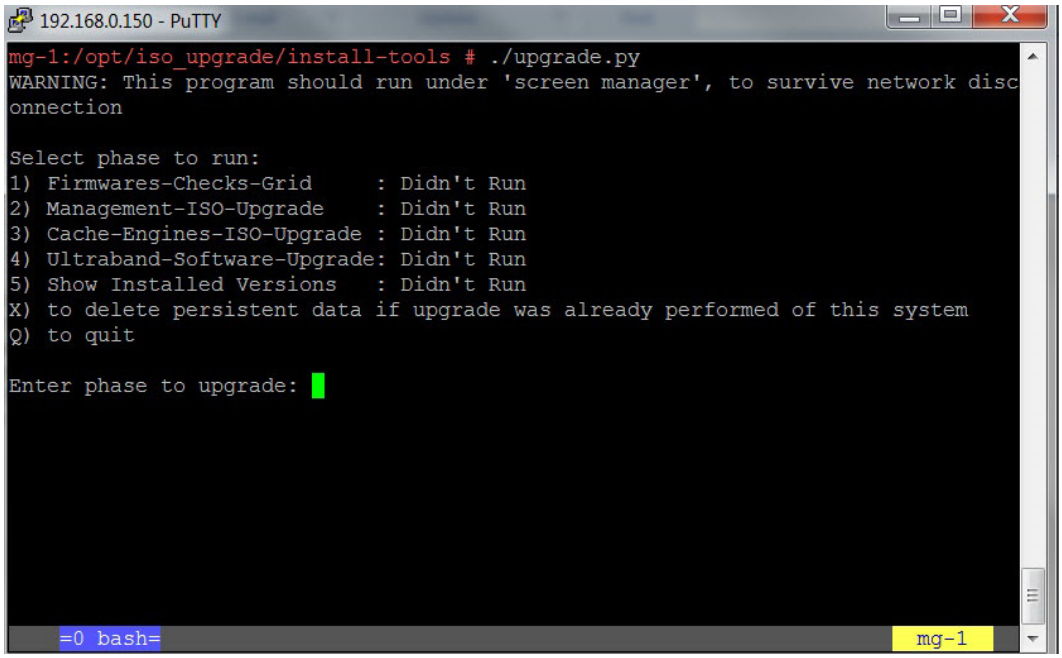

- **Firmwares-Checks-Grid:** Checks the firmware version.
- **Management-ISO-Upgrade:** Upgrades the Operation System for the Management server.
- **Cache-Engines-ISO-Upgrade:** Upgrades the Operation System for the Cache Engines servers.

- **Ultraband-Software-Upgrade:** Upgrades the software.
- **Show Installed Versions:** Determines the Grid system status.

1. To start the upgrade process, enter 1

   The Upgrade tool starts by checking and confirming the Firmware versions. If the Firmwares are not up-to-date, a message is displayed informing, that you should upgrade firmwares before continuing.

   Note: If your Firmware needs to be updated, upgrade the Firmware using SUU 14.12.200.69. If there are any questions, please contact Customer Support.

2. After the Firmware is confirmed, enter 2 for the Management-ISO-Upgrade.

   The management server reboots itself to perform the upgrade. In order to confirm the network connectivity you must ping the management server. When the system is back Online, login to the system using ssh, run the screen command and load the Upgrade Tool using the following command:

   /opt/iso_upgrade/install-tools/upgrade.py
   When the 2nd step is completed, the menu is displayed again,

3. Continue by entering 3, to upgrade all Cache Engines.

   The upgrade.py waits for all Cache-engine to go back online.

4. Enter 4 to run the UltraBand-Software-Upgrade.

   Note: If you do not have a valid license, after the upgrade is completed, the grid will not work.

5. Enter 5 and run the *Show Installed Versions,* confirm that the output is equivalent to the version you just installed.

Standalone/Grid-Ready Upgrade:

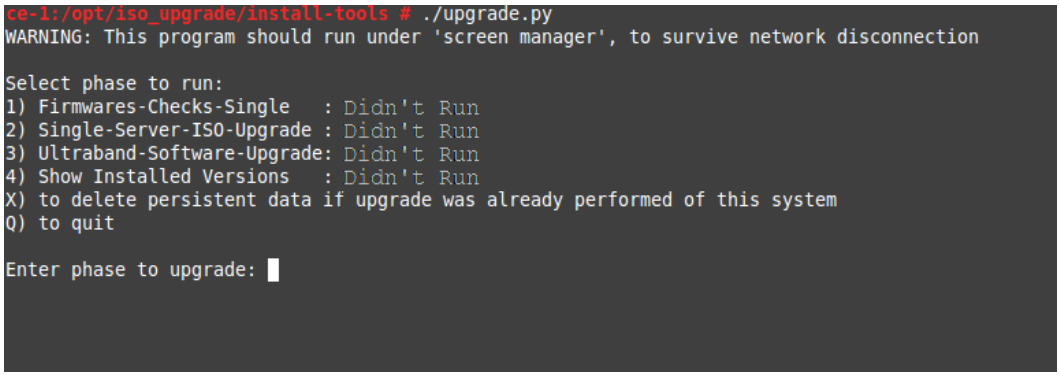

The following menu options are displayed:

- **Firmwares-Checks-Grid:** Checks the firmware version.

  Note: If your Firmware needs to be upgraded, please contact Customer Support.

- **Single-Server-ISO-Upgrade:** Upgrades the Operation System.
- **Ultraband-Software-Upgrade:** Upgrades the software.

■ **Show Installed Versions:** Determines the system status.

1 To start the upgrade process, enter 1

The Upgrade tool starts by checking and confirming the Firmware versions. If the firmwares are not up-to-date, a message is displayed informing, that you should upgrade firmwares before continuing.

2 After the Firmware is confirmed, enter 2 for the Single-Server-ISO-Upgrade.

The server reboots itself to perform the upgrade. In order to confirm the network connectivity you must ping the server. When the system is back Online, login to the system using ssh, run the screen command and load the Upgrade Tool using the following command:

/opt/iso_upgrade/install-tools/upgrade.py

When the 2nd step is completed, the menu is displayed again,

3 Enter 3 to run the UltraBand-Software-Upgrade.

Note: If you do not have a valid license, after the upgrade is completed, the system will not work.

4 Enter 4 and run the *Show Installed Versions,* confirm that the output is equivalent to the version you just installed