



Global Distributed Storage Engineering

CLARiiON CX Series Storage Admin Guide

GDSE-03011-TM
Technical Manual

1.2

1. Version Control

Version	Date	Author	Reviewed	Comments
0.1	19-June-03	Ludwig Derodel	GDSE	RFC
1.0	03-Oct-03	Ludwig Derodel	-	Final
1.1	26-Jan-04	Ludwig Derodel	GDSE	Final
1.2	30-Aug-05	Ludwig Derodel	GDSE	Updated "privileged user" section

2. Legal Notice

Information contained herein is for Internal Use and may be used only for business purposes authorized by Citigroup.

CITI, CITIBANK, CITICORP, CITIGROUP and the Umbrella Device are trademarks and service marks of Citicorp and are used and registered throughout the world. Citibank, Citicorp, Citigroup and their subsidiaries also claim rights in certain other trademarks and service marks contained in these pages.

3. Contact

Please address any comments or questions regarding the information contained in this document via email to the “*GT Global Distributed Storage Engineering” distribution list or alternatively to gdse@citigroup.com.

4. Latest Versions

It is recommended that the latest version of the document be used from the GDSE website <http://gdse.ny.ssmb.com>. Newer versions may include certification on additional platforms, or comments on issues discovered after deployment.

5. Table Of Contents

1. VERSION CONTROL	2
2. LEGAL NOTICE	3
3. CONTACT.....	3
4. LATEST VERSIONS	3
5. TABLE OF CONTENTS	4
6. INTRODUCTION.....	6
Overview	6
6.1. CX400:.....	6
6.2. CX600:.....	6
7. PREREQUISITES	7
7.1. CHECKLIST.....	7
7.2. SUPPORTED VERSIONS	7
8. INSTALLATION	8
9. MANAGEMENT STATION SETUP	8
9.1. THE HARDWARE	8
9.2. THE SOFTWARE.....	8
9.2.1. Terminal Services.....	8
9.2.2. Local Group Creation.....	9
9.2.3. Configuring Terminal Services	10
9.2.4. Configuring Off Array Management Server	10
9.2.5. Navisphere Security and Domains	11
9.2.6. IIS.....	11
9.2.7. Navisphere Analyzer	11
10. INITIAL ARRAY SETUP	12
10.1. HARDWARE.....	12
10.2. SOFTWARE	12
10.2.1. Connecting to the CLARiiON CX via Navisphere Manager	12
10.2.2. Managing Privileged Users	13
10.2.3. SPE Configuration / Initialization	13
10.2.4. Storage System Settings	14
10.2.5. ALPA.....	14
10.2.6. Logging	15
10.2.7. High/Low Watermark and Cache Settings.....	16
10.2.8. SP Port Speeds.....	17
10.2.9. Burn-in Testing	17

11. RAID GROUPS	18
11.1. DISK SELECTION	18
12. LUNS	20
12.1. CREATING LUNS	20
13. STORAGE GROUPS	22
13.1. CREATING STORAGE GROUPS	22
13.2. CHANGING LUN NAMES	22
13.3. LUN SELECTION.....	23
14. HOST REGISTRATION	25
14.1. NO HBA DRIVER LOADED	25
14.2. NAVISPHERE AGENT INSTALLED AND RUNNING.....	26
14.3. HOST CONNECTED TO A STORAGE GROUP.....	27
15. SNAPVIEW	28
15.1. CLONES.....	28
CLONE STATES.....	28
15.1.1. <i>Creating A Clone Group</i>	29
15.2. ADDING A CLONE TO A CLONE GROUP	30
15.3. SNAPSHOT.....	31
15.3.1. <i>Configuring the Snapshot cache</i>	31
15.3.2. <i>Creating a Snapshot session</i>	32
15.3.3. <i>Creating a Snapshot</i>	32
15.3.4. <i>Activating the snapshot</i>	33
16. MONITORING.....	35
16.1. CREATING A PORTAL	35
16.2. ADDING A STORAGE SYSTEM.....	35
16.3. SNMP TEMPLATE.....	36
16.4. E-MAIL TEMPLATE	38
16.5. ASSIGNING TEMPLATES TO STORAGE SYSTEMS	39
16.6. VERIFYING MONITORING SETUP	41
16.7. DIAL-UP ACCESS	42
17. REFERENCES.....	42
17.1. NAVISPHERE MANAGER PACKAGE.....	43
17.2. TROUBLESHOOTING	44
17.3. TROUBLESHOOTING PERSISTENT BINDING.....	45

6. Introduction

Successful storage administration involves a deep understanding of the goals that need to be accomplished, the tools available, and have a working process flow. This document covers all of these areas in detail with respect to the EMC Clariion CX Series Storage System. For a better understanding of the system and the Best Practices that are supported from the Global Distributed Storage Engineering group, please refer to the documents in [CLARiiON CX Series Best Practices](#)

Overview

6.1. CX400:

- o 2 Front End 2Gb ports per Storage Processor
- o Dual Back End 2Gb loops
 - o 30 drives per loop maximum
- o Dual 2Gb CMI channels
- o Storage Processors configured with 1GB of memory each
- o Single low-power 800MHz Pentium III processor
- o 15 drives (12 volt) maximum per enclosure
- o 60 drives maximum per array
- o 60 RAID groups
- o 32 LUNs per RAID group/512 LUNs per array
 - o Minimum LUN size 1MB
 - o Maximum LUN size 2TB
 - o 256 Luns per initiator

6.2. CX600:

- o 4 Front End 2gb ports per Storage Processor
- o Dual Back End 2Gb loops
 - o 120 drives per loop
- o Dual 2Gb CMI channels
- o Storage Processors will support 2GB or 4GB of memory
- o Dual 2Ghz Intel Processors
- o 15 drives per DAE
- o 240 drives maximum per array
- o 240 RAID groups per array
- o 240 RAID groups per array
- o 32 LUNs per RAID group/1024 LUNs per array
 - o Minimum LUN size 1MB
 - o Maximum LUN size 2TB
 - o 256 Luns per Storage Group.

7. Prerequisites

7.1. Checklist

The following checklist outlines the in-order procedures to proper setup, initialization and management for the Clariion CX Series.

Activity	Where	Who is responsible
1) Management Station Setup	On the same subnet as the Clariion and switches	Storage Management Integration Infrastructure
2) Initial Array Setup	Location where Clariion is being installed	EMC Personnel Storage Management
3) Burn-in Testing	From Management Station	Storage Management
4) Create RAID Group	In Navisphere Manager	Storage Management
5) Storage Request	From System Administrator	System Administrator
6) Build Host	Where host is located	System Administrator
7) Install HBA Driver	Host	System Administrator
8) Install Navisphere Agent	Host	System Administrator
9) Install Powerpath	Host	System Administrator
10) Allocate Storage	In Navisphere Manager on Clariion	Storage Management
11) Verify Visibility	Host	System Administrator

7.2. Supported Versions

All references to supported software and firmware revisions are in the Clariion Best Practices Document. The supported revisions will be covered in detail, complete with historical support versions for reference

8. Installation

The installation process is sectioned into the following components:

Management Station Setup
 Initial Array Setup
 RAID Groups
 LUNS
 Storage Groups
 Host Registration
 SnapView
 Monitoring

9. Management Station Setup

9.1. The Hardware

The [Proliant DL380G2 1.266Mhz \(2U\)](#) is the supported management station and should be installed before attempting to install and manage the Clariion Environment. This station is needed to configure storage system settings and manage the Clariion Environment once initialized.

Note: This system communicates with the Clariion via out-of-band techniques ONLY. There are NO HBA(s) installed on this system and it is never allocated storage from the storage system.

9.2. The Software

Install the software in order as listed in Table 1 Management Station Software. Software revisions should be received from the Clariion CX Series Best Practices Document.

Table 1 Management Station Software Installation Order


Operating System
Off-Array Mgmt Server
Off-Array Mgmt UI
Navisphere Analyzer
NaviCLI
SnapView
IIS (for FTP)

9.2.1. Terminal Services

In order to provide a higher level of security and audit capabilities, access to the terminal server of the management station shall be based on NT Domain style accounts. By default, only an Administrator can login to a terminal session. This needs to be reconfigured to allow only authorized users the ability to login.

9.2.2. Local Group Creation


First, create a local group on the management station containing either the global groups and/or the users that should be members.



1. Create a group with a descriptive name.
2. Add only those users that need access to Navisphere Manager in here then Apply.

9.2.3. Configuring Terminal Services

Now assign the local group to the terminal services configuration so those users can manage the Clariion environment.



1. Open Terminal Services Configuration under Administrative Tools
2. Right-click on the RDP -TCP Connection and go to properties
3. Add the local group created above in Local Group Creation.
4. Add User Access to the Allow column for the group

Idle timeout should be set to 1 hour and end the session, not disconnect.

9.2.4. Configuring Off Array Management Server

Off Array Management Server is required to allow Event Monitor work properly in our environment. There are 2 components: Off Array Mgmt Server and UI. Installation is done via self-extracting executables.

Verify that a web server is running on 127.0.0.1 after these 2 components are installed. Set the IP address of the web server by logging into <http://127.0.0.1/setup.html> and setting the IP address to the non-routable interface address.

9.2.5. Navisphere Security and Domains

Security should be initialized (This will create 1 Global Admin Account) when first attempting to connect to the web server on the Management Station.

Select the Master Node for the Domain to be 127.0.0.1 under "File->Setup Domain->Select Master"

Other Accounts should be created under "Tools->Security->User Management"
After arrays are initialized and have IP connectivity (See Initial Array Setup on pg 12), those systems can be added to the domain. under "File->Setup Domain->Configure Domain"

Verify through Domain Status under "File->Setup Domain"

9.2.6. IIS

Install IIS, but only enable FTP, NOT Web Services (otherwise machine will crash)

Configure FTP home directory and disallow guest account access

9.2.7. Navisphere Analyzer

Install Navisphere Analyzer (reference Best Practices for version)

Enable Data Logging using Logger Control (Reference Best Practices for settings)

10. Initial Array Setup

Before installing the array, 2 IP connections are needed and registered in Whitepages for name resolution.

10.1. Hardware

This process will most likely be done with EMC assistance, following these best practices should allow Storage Management the ability to oversee the process with confidence.

Use the following recommendations:

Rackmount	Determined by Hardware Group
Power	Determined by Hardware Group
Cooling	Determined by Hardware Group
Enclosure Addresses	Refer to Figure 2-2 Changing the EA on the DAE in the <i>CX Series Hardware Reference Manual</i> .
Cabling	Refer to Figure 2-7 CX600-Sample SPE and DAE Cabling in a Single-Cabinet Installation for EMC recommended cabling in the <i>Clariion CX Series Hardware Reference Manual</i> .
Switches Connections	Determined by design on growing basis. See McData Best Practices for new deployment and <i>Brocade 3800 Best Practices for legacy environment</i> .
IP Connectivity	Connect each SP to the network for out-of-band management access.


10.2. Software

10.2.1. Connecting to the CLARiiON CX via Navisphere Manager

In Navisphere Manager on the Management Station go under *File->Setup Domain->Configure Domain* and enter the IP Address for SPA and click on the '>' button. Repeat for SPB of the Array and click on Apply. Storage System should show up under the Storage Systems icon Software.

10.2.2. Managing Privileged Users

To create a more secure environment for CLARiiON management, the use of privileged users are dealt with. They are managed under the Agent tab in each of the SP Properties windows.



By default there are no entries listed, which gives complete OPEN access to modify the settings on the Clariion. In order to restrict the CLARiiON environment, a **dummie** user id (i.e. test) MUST be added in each of the SPs. This will prevent anyone from making changes using navicli. However, monitoring functions can still be done using navicli.

Use caution when entering these entries as misconfiguration can leave the Clariion without any NaviCLI capabilities and a call to EMC will be required.

Entries should be of the form:

[user@hostname](#)

Note that this is no reference to the NT domain and the machine name is NOT fully qualified.

10.2.3. SPE Configuration / Initialization

This will be done by an EMC professional, but be supervised by a Storage Management team member.

Initialization	Follow the instructions in Initializing the Clariion CX Storage System in the <i>CX Series Rackmount Model Hardware Reference Manual</i> using the following recommendations.
----------------	--

Hot Spare	Disk 14
-----------	---------

For burn-in testing purposes of the hot spare disk, initially create disk 14 as an “individual disk” and bind a single LUN to test the disk. Upon success, unbind LUN -> Destroy the RAID Group and create it as a Hot Spare.

Software Installation on the SPs should be at the levels referred to in the Clariion CX Series Best Practices document. An EMC professional should perform this function with supervision. Verification should be made that these levels are also in the Active state via Navisphere Manager.

10.2.4. Storage System Settings

The following settings should be set using Navisphere Manager on the Management Station.


- ALPA
- Logging
- High/Low Watermark and Cache Settings
- SP Port Speeds

10.2.5. ALPA

The Arbitrated Loop Physical Address should be set to the target that we want presented to the hosts connected. Both ports from an SP should be the same address.


Each FC4700 SP should also be assigned unique target Ids. This aids in identifying which Clariion a host might be hooked up to.

Setting the ALPA is found under the ALPA tab in the SP Properties windows under the Storage Tab.




10.2.6. Logging

1. Statistics logging should be checked for both SPs.




10.2.7. High/Low Watermark and Cache Settings

1. Page size should be 8k. Except for Exchange, which should be 4K.
2. 80/60 should be the High/Low watermark settings respectively for CX600 and
 - o 60/40 High/Low for CX400.
3. Watermark Processing should be enabled.
4. Read and Write Caching should be enabled for both SPs.
5. Write cache can be set to a maximum of 2048MB the rest should be allocated to Read cache. Set to maximum when CX600 is configured with 4GB per SP.
6. Read cache settings should be set for 40 (80-100) MB per SP in a CX400 and remainder should be set for write cache.



10.2.8. SP Port Speeds

Since the only approved switch is the Brocade 3800, the speed settings should be set to 2 for all ports. The settings on the switch can be left to auto-negotiate.



10.2.9. Burn-in Testing

Burn-in testing must be performed for any new install and additional DAE installs. During testing, a fair number of disk failures occurred due to poorly seated drives.

A relatively quick procedure will test all drives. Create a single RAID 5 RAID Group with all disks in a DAE and bind a single small LUN on the RAID group to exercise all disks. This should suffice to catch poorly seated drives.


After successful testing, unbind the LUN and destroy the RAID Group. Then the disks are available for general use.

11. RAID Groups


A RAID Group is a set of disks that will be used to bind one or more LUNs. Each LUN that is bound in a RAID Group will be distributed equally in that group.

Creating the RAID Groups is a simple task that is accomplished by Right clicking on the Storage Tab and going to Create RAID Group.

Click on the Advanced button in the popup, which takes you the creation screen. Select the correct RAID Group type and give it an ID.



1. Set the disk selection to Manual and click on Select. This brings up a popup for disk selection.



11.1. Disk Selection

Depending on the RAID Group type, disk selection order is of critical importance.

RAID 1/0 selection is in mirrored pairs. Clariion Best Practices states that RAID1/0 should always span 2 DAEs and buses. Each mirrored pair should span both DAEs and backend buses, such that no mirrored pair sits on the same DAE. Disks should be selected in the order stated by the best practice (Orders should alternate DAEs and buses). (i.e. 0_0_1, 1_0_1 = Bus_Enc_Disk)

RAID 5 allocation is a little more lenient, but as a best practice, it would be easier to select the range of disks from each DAE and pull them over at one time.

Note: Make sure that the right side on the window shows the alternating bus/DAE configuration. A “flarecons” output from EMC will always confirm the correct layout of a RAID 1/0 configuration on the array.

12. LUNs

A logical unit (LUN) is a grouping of one or more disks into one span of disk storage space. A LUN looks like an individual disk to the host operating system. A LUN has RAID type and other properties that define it.


A LUN is carved out of a single RAID Group.

All LUNs in a given RAID Group are of the same RAID type. Either and only one SP can own each LUN in a RAID GROUP at a given timeOptional: Provide step-by-step instructions for updating firmware by an SA.

12.1. Creating LUNs


Creating LUNs are accomplished through a process called binding. To bind a LUN, right click on the Storage Tab in Navisphere Manager and select Bind LUN...

This will bring up a Bind LUN popup.




1. Select the RAID Type that you would like to bind the LUN in.
2. Select the RAID Group for the LUN to be a member of. If there are no RAID Groups available for selection, go back to Creating RAID Group and create a RAID Group first
3. Select the LUN ID.
4. Enable the Read and Write Caches, and disable Auto Assign
5. Set the Default Owner of the LUN to SPA if the LUN ID is even, otherwise set it to SPB.
6. Set the size of the LUN appropriately.
7. Click on Apply to bind the LUN. There should be a popup stating successful initiation of the bind operation

8. While the LUN is binding, Navisphere will go into a transition state, which is denoted by a blinking icon in the top left hand corner of the window and a T over the icons of the affected items in storage system.



Notice the “T”ransition icons

To monitor progress of the binding LUN, examine the properties of the LUN like in the following screen:



The Percent Bound is currently at 80. When it hits 100, the LUN is completely bound. Navisphere will return to normal status after all LUNs have been bound.

13. Storage Groups


13.1. Creating Storage Groups

Storage Groups answers the question “What hosts can see what LUNs?”.

Before moving on, the host should be registered properly. See *Host Registration*

To create a Storage Group, right click on the Storage Tab in Navisphere Manager and select Create Storage Groups...

This brings up the Create Storage Group Popup Window.




Change the name of the Storage Group to the hostname that this Storage Group will correspond to.

The Sharing selection should be set to Dedicated (unless more than 1 host shall see this storage at the same time)..


13.2. Changing LUN Names

Due to the nature of how Navisphere maps Clariion LUNs to Host-side LUNs, it is necessary to rename the LUNs to follow the standard set forth in the Best Practices Document.




Renaming the LUN Name from the original name to the standard name helps to ensure Proper LUN Ordering on the Host side.

13.3. LUN Selection




Click on LUN 000 <storage group>, move it over to the right side and click on Apply.



Click on the rest of the LUNs that are supposed to be assigned to this host and click on the right arrow to move them over to Selected LUNs and click on Apply. This is done to ensure proper LUN ordering on the Host side.

Click on the Connect Hosts... button to bring up the Connect hosts popup



Select the host that should access these LUNs and move it to Hosts to be Connected and click on Apply and OK.

14. Host Registration

Before the Storage Management team can give storage to a host, the host must be registered.

This means:

- The host must be built
- The HBA(s) must be installed
- The driver must be loaded
- Navisphere Agent installed and running


If the proceeding is done, the machine should be properly registered in Navisphere.

For example, let's go through the process of allocating storage for a machine called uelabsrvmp1 that will have 2 HBAs.

First, let's see if it's registered already. This is done via the connectivity status popup.

14.1. No HBA Driver Loaded


In the following popup, there is no uelabsrvmp1 registered. This will coincide with the HBA driver not even being loaded.



There is NO entry for lbuilds4.ny.smb.com at this time

14.2. Navisphere Agent Installed and Running

The following picture represents the server with Navisphere Agent installed and running. It shows the machine as registered. Also, the hostname is now known. At this point, the machine is completely registered. The reason the icons are still blue is because the machine has not been assigned to a storage group.




The entries exist for the newly recognized host. Notice that they are still NOT Registered and the hostname is HBA uid.

Note: Failover Mode and Array Commpath settings are set to 1/Enabled
The name and IP address of server would need to be entered manually when registering for the first time.

14.3. Host connected to a Storage Group

Finally, the following picture shows a machine that has been properly registered and has been assigned to a storage group. Note how the icons are now green and the Accessible Storage Group name is now the name of the Storage Group created for this host (uelab).



15. SnapView

SnapView is a storage-system-based software application that allows you to create a copy of a LUN by using either clones or snapshots. SnapView technology was not certified on the FC4700-2 for two main reasons. The primary reason was in the event of a SP reboot or power failure, the snapshot session would be lost. The secondary was the inability to provide a complete copy of the source lun for instant recovery.

Those two issues have been address in SnapView version 02_10_02. The SnapView software now has 2 features, clones and snapshots.

15.1. Clones

A clone is a complete copy of a source LUN. A source LUN would need to be part of a Clone group in order to create a clone. The software assigns a clone ID to each clone added to the clone group. The ID remains until the clone is removed from its group. There are 6 possible clones states see table below


Clone States

Clone State	Description	Cause of State	Permitted Operation
Reverse Synchronizing	A clone is in the process of copying its data to its source LUN	- Manually administered with Navisphere or CLI	Fracture
Reverse Out-of-Sync	A failure during a reverse synchronization process which makes the source unusable	A reverse synchronization operation failed to complete successfully.	- Reverse synchronize - Remove - Fracture
Synchronized	A clone is a byte-for-byte copy of its source.	- After synchronization - A clone is added to Clone Group without setting "Initial Sync Required"	- Fractured - Remove
Synchronizing	A clone is in the process of copying data from its source LUN	- Automatically started when you add a clone to a clone Group with the Initial Sync Required properly selected.	Fracture
Out-of-Sync	A clone was in the process of synchronizing but failed. Therefore clone is unusable.	A synchronization operation failed to completed successfully	- Synchronize - Remove - Fracture
Consistent	A clone is in a synchronized state that receives host I/O from the source.	A clone is fractured while in a Synchronized state. A clone is unfractured and has yet to transition reached a Synchronized state.	- Fracture - Remove - Synchronize - Reverse Synchronize

In order for Clones to be configured on a CX array, there are a number of steps that need to done.

1. Install and configure SnapView

2. Create and allocated 2 Clone Private LUNs (CPL). Each CPL must be a 250000 blocks must be bound to each SP. The CPL LUNs are use to record changes to data chunks of all fractured clones. Actual data is not stored on the CPLs. Therefore increasing the size of the CPL does not provide any performance improvements. (see Figure below)




3. Enable **Protected Restore** feature across the array by selecting the **Allow Protected Restore** option in the Clone Features Properties dialog box.

15.1.1. Creating A Clone Group

Right-click the storage system icon from the client managing the storage system, in the **Enterprise Storage** dialog box and select **Create Clone Group**. The dialog box below will open.

1. Notice that standard descriptive name should always be used.
2. The quiescent threshold should stay at the default 60 sec.
3. Select the lun to be cloned. Only one lun can be selected at a time. If a server has more than one lun a clone group must be created for each lun.




15.2. Adding a Clone to a Clone Group

From the **Storage** tab of the **Enterprise Storage** dialog box, navigate to the **Clones** icon, and then the clone group name icon to which you want to add a clone.

Right-click the **Clone Group Name**, and then click **Add Clone**. The dialog box below will open.

1. Select the clone from the clone list.



The following should always be selected:

Initial Sync Required – To insure that a full synchronization is done upon creation of clone (process may take from a few minutes to a few hours).

Use Protected Restore – To insure that a Clone Copy is not accidentally overwritten after a reverse sync.

After the Clone is fully synchronized, the clone must be assigned to its own storage group and allocated to the clone server. The process to allocated clones in storage groups is the same as any regular luns.


Important: Although a CX array supports 73GB, 146GB and 250GB disk drives, it should be noted that if a clone lun is created out of 250GB ATA drives it will affect performance on a source lun created from a 73GB and 146GB disk drives.

15.3. Snapshot

Unlike the clone, a snapshot allows you to access point-in-time data of a source LUN. The data is copied from the source LUN to the snapshot in the following manner: Copy-On-First-Write (COFW) and occurs when the production first modifies a data chunk on the source LUN.


15.3.1. Configuring the Snapshot cache

The snapshot size should follow that describe in the Clariion CX Best Practice document. The following picture represents the allocation of cache luns to each SP.



Note: The size of the cache lun can be of any size but 50GB is the maximum recommended lun size.


15.3.2. Creating a Snapshot session



1. Session name – should be descriptive and follow Best Practices
2. Optional mode “**persistent**” should always be selected. Persistent mode allows a session that can withstand the following failures and trespass.
 - a. SP Reboot or failure
 - b. Array reboot or power failure
 - c. Host I/O trespass to the peer SP
3. Select the source LUN.
4. Click Apply for changes to take effect

15.3.3. Creating a Snapshot


In order to add a snapshot, you must select the source lun and then right-click, select snapview and **create snapshot**. The snapshot lun can also be added to a storage group in the same window. The figure below show the storage group “Snap_server”.



15.3.4. Activating the snapshot

Activating a snap lun will make the lun will make that lun accessible to the snapshot host. In order to activate the snap lun, simply click on the lun and select add snapshot. Once the snapshot has been activated, the snapshot host cannot access the snapshot data.

See picture below for a screenshot.



1. Select a session in which you need to activate the snapshot

Once the host is done with the snapshot, it can be deactivated to prevent the snapshot server from using the snap luns.

Note: In both clone and snapshot scenarios, the source lun should be quiesced in order to flush all outstanding I/Os to the disks before a fracture or activating the snapshot.

16. Monitoring

Enterprise Monitoring of the Clariion Storage System is accomplished via Event Monitor.

16.1. Creating a Portal

Considering the environment with a private non-routable IP segment for the management of the storage system and associated FC switches, a “portal” is required for being able to send notifications out to the corporate LAN from the management segment.

This is accomplished by going to **Tools->Portal Configuration**. Right Clicking on **Portals** and going to **Add Portal**


Adding a portal by pulling over the management station to the right. Click on OK and then Apply.

16.2. Adding a Storage System

Adding a Storage system to the Portal is the next stop. This is accomplished by right clicking on the Management Station in the Portal Configuration window and selecting Add Storage System. (see picture below for illustration)

In the popup:

1. Add the private network IP address of the management station in the Select Systems section.
2. Click on the right arrow
3. The Hostname and IP address should appear in the lower right hand window labeled “Selected Systems”.
4. Apply the change.




Creating the Response Templates


The Response templates define the events that trigger notifications. Separate templates for SNMP and E-mail are provided due to slightly different behaviors for notification.

16.3. SNMP Template

Create an SNMP template with the following properties




1. General Events
2. Basic Array Feature Events
3. Error and Critical Event Severity
4. Log to System Log



1. SNMP Management Host should be pointed to the IP of the Nerve Center detailed from EMS
2. The Community string should be set to **public**.

16.4. E-Mail Template

Create a template for E-mail notification with the following properties



1. Use General Events
2. Select Basic Array Feature Events
3. Select Error, Critical and Warning Event Severities
4. Check Log to System Log and Single Notification for Multiple Events

Template Properties

General E-Mail Paging SNMP

Please specify the parameters for e-mail notification:

1 E-Mail to: gdse@citigroup.com

2 Cc:

3 From: gdse@citigroup.com

4 Subject: Clariion Notification emcsrv11

SMTP Server: mailhub-nyc3.ny.ssmc.com

5 ☒ Use Default Message

Time Stamp %T% Event Number %N%
 Severity %SEV% Host %H%
 Storage Array %SUB% %SP% Device %D%
 Description %E%

Test Clear

OK Apply Cancel Help


1. The E-mail to field should be directed to the appropriate team for notification
2. The From field should be set to the group responsible
3. The Subject line should read Clariion Notification <Management Station Name>
4. The SMTP Server should point to a valid Mailhub in the nearest location.

Note: Do no attempt to use our MX Record infrastructure. It is necessary to point to a specific mailhub

5. Use the Default Message format for E-mail notification

16.5. Assigning templates to storage systems

Once the SNMP and EMAIL templates have been created, assign them to all storage systems to be monitored.



Note: A template can be assigned to more than one storage system and multiple templates can be assigned to a single storage system.

1. Expand the IP Address of the Management Station
2. Right Click on the Storage System to be monitored
3. Select the templates to monitor

Repeat for both SNMP and EMAIL templates and all storage systems to be monitored.


Important: The Event Monitor configuration wizard can also be used to create new templates as well as the steps described above.

Also always make sure to create or modify the **agent.config** text file under the Navisphere Agent folder of off-array manager to include the following entries containing the portal IP address and the IP address of all the SPs.

```
clarDescr
clarContact
poll 60
baud 9600
eventlog 2048
user system@192.168.93.1
user system@192.168.93.60
user system@192.168.93.61
```

16.6. Verifying Monitoring Setup

In order to verify the configuration, expand the Monitors tab and look for the correct templates under each of the storage systems.



1. The Management Station is viewed upon as a "Storage System" that was created through the Portal.
2. The Physical Clariion
3. The Template successfully assigned to this storage system.

Testing of the templates can be performed to check end-to-end notification is working

The following is an excerpt of a test SNMP trap sent to Netcool. This should verify that traps are successfully getting up to the Netcool displays.

Alias	Summary	Category	
.120	Error Code: 2002 .Error Message : Response Test Request	Clarion Status	Dis

Similar tests can be run for E-mail notifications. The following is the output of a test e-mail sent from the Clariion.

--sample e-mail

Time Stamp 10/16/02 12:05:51 **Event Number 2002**

Severity Information Host emcsrv11

Storage Array N/A N/A Device N/A

Description Response **Test Request**

16.7. Dial-Up Access

Dial-up access is configured using On-Alert as stated in the best practice. If Dial-up access is required, and EMC Customer Engineering will handle the configuration.

17. References

Document	Location
Clariion CX Series Evaluation	http://san.ny.ssmb.com/GDSE_docs/GDSE-03011-C.pdf
Clariion CX Series Best Practices Guide	http://san.ny.ssmb.com/GDSE_docs/GDSE-03011-B.pdf
Clariion CX Series Rackmount Model Hardware Reference Manual	
EMC Navisphere Manager 6.x Administrator's Guide	

17.1. Navisphere Manager Package

In an effort to facilitate the installation and configuration of the management server, the software package has been bundle in one Windows Package. The package installation follows all GDS standards where the install directory is always c:\Apps_svr. The windows package can be obtained from GDSGOLD per request only. The request must be sent to gtusdisstorinte@citigroup.com.

The package contains the following component:

Navisphere Management Server

Navisphere CLI

SnapView (which contains all the snapview *ndus*)

Navisphere Management UI

Base Codes (which contains all the base *ndus*)

The procedure to install the package is as follows.

Install the Navisphere Agent package. (Wrap ID: nvsa6410)

Install the Navisphere Server Bundle package (Wrap ID: nvsb6410)

Install all relevant *ndus* on the array. (located in the folder Base Codes under the install directory)

17.2. Troubleshooting

Replacing HBAs Connected to Shared CLARiiON Systems

A shared CLARiiON storage system is a storage system with Access Logix software.

Follow these steps to replace an HBA connected to a shared storage system:

1. Stop all I/O to the storage systems connected to the HBA you are replacing.

2. Remove the connection path to the old HBA from each storage system's persistent memory, using these CLI commands:

- a. **navicli port -list**

- b. **navicli -removeHBA**

Refer to the CLI reference manual for information on this command.

3. Power down the server.

4. Remove the HBA you want to replace and install the new HBA.

5. Power up the server.

6. After the HBA is registered with the Host Agent running on the server, reconnect any Storage Groups that you disconnected.

The HBA is registered with the Host Agent if the Connectivity Status dialog box in Manager shows Yes in the **Registered** column for the HBA.

7. In any switch zones that include the HBA, change the World Wide Port Name (WWPN) of the old HBA to the WWPN of the new HBA.

9. Reboot the host with the **-r** (reconfigure) option:

- a. **reboot -- -r**

17.3. Troubleshooting Persistent Binding

Persistent binding on the Clariion platform could be a very confusing task if not done step-by-step. First, there are a number of nomenclatures that I need to be stated.

- All the SP sports have the following common entries in their WWPN:

SP A (port 0) = 50060160XXXXXXXXX

SP A (port 1) = 50060161XXXXXXXXX

SP A (port 2) = 50060162XXXXXXXXX (CX600 only)

SP A (port 2) = 50060163XXXXXXXXX (CX600 only)


SP B (port 0) = 50060168XXXXXXXXX

SP B (port 1) = 50060169XXXXXXXXX

SP B (port 2) = 5006016aXXXXXXXXX (CX600 only)

SP B (port 2) = 5006016bXXXXXXXXX (CX600 only)

The picture below will demonstrate how a Solaris Host with 2 HBAs should be connected and zoned to 4 ports of a CLARiiON array.



As shown in the picture above each HBA is zone to one port of SP A and one port of SP B. This configuration is done to insure that a fabric failure will not route all the traffic to one SP. The zone entries should match the persistent binding entries. For example if lpfc0 is zone to sp A port 0, the lpfc.conf entries should also demonstrate that lpfc0 is bind to sp A port 0. See example below.

```
/kernel/drv/lpfc.conf
```

```
fcplib-WWPN="500060160106005d9:lpfc0t0",
           "500060169106005d9:lpfc0t1",
           "500060161106005d9:lpfc1t0",
           "500060168106005d9:lpfc1t1";
```

```
/kernel/drv/sd.conf
```

```
name="sd" parent="lpfc" target=0 lun=0 hba="lpfc0";
name="sd" parent="lpfc" target=1 lun=0 hba="lpfc0";
name="sd" parent="lpfc" target=0 lun=0 hba="lpfc1";
name="sd" parent="lpfc" target=1 lun=0 hba="lpfc1";
```

The sections highlighted in red shows that lpfc0 is being bound to sp B which is target 0 and sp B which is target 1.