



Global Distributed Storage Engineering

Centera 3.0 Certification

GDSE-05024-SG

Solutions Guide

Version 1.3

VERSION CONTROL

Version	Date	Author	Reviewed	Comments
Draft 0.1	21 st October 2005	Stephen Hopkins	Chris Clayton	Peer Review
Draft 0.2	15 th November 2005	Stephen Hopkins	EMC Engineering, Burt Giron	Peer Review
Release 1.0	28 th November 2005	Stephen Hopkins	RFC Review	
Release 1.1	20 th January 2006	Stephen Hopkins		Updated after comments from RFC
Release 1.2	8 th February 2006	Stephen Hopkins		Update ICQ and Status notification. Added multi-cabinet configurations
Release 1.3	4 th April 2006	Stephen Hopkins		Updated usable capacity table, also modified regen buffer policy setting

CONTENTS

Legal Notice	6
Contact	6
Latest Versions	6
Confidentiality	6
1 INTRODUCTION	7
1.1 Intended Audience	7
1.2 Background	7
1.3 Security Classification	8
1.4 Release Status	8
2 CENTERA OVERVIEW	9
2.1 Using Centera Storage	9
2.2 Centera Architecture Overview	11
3 CTI CENTERA CONFIGURATION STANDARDS	13
3.1 Centera hardware expansion	15
3.2 Centera Universal Access	15
3.3 Multi-Cabinet Centera Clusters	15
4 PREPARATION FOR INSTALLATION	17
4.1.1 Replication	18
4.1.2 Storage-Strategy	19
4.1.3 Current RPQ Requirements for CTI Centera	20
4.1.4 Management Server/Infrastructure	22
4.1.5 Running the Centera management interface	23
5 INSTALLATION OF CENTERA	25
6 CTI CONFIGURATION	26
6.1 CTI specific configuration	26
7 UPGRADING AN EXISTING CENTERA	29
7.1 SDK Compatibility with CentraStar 3.0	29

7.2	Changes required in order to start using virtual pools	30
7.3	Migrating old data application to new virtual pools	32
8	MANAGEMENT TASKS	34
8.1	Adding a new application	34
8.1.1	Creating a new virtual pool	34
8.1.2	Creating a new application profile	35
8.1.3	Synchronising the pool and profile data to the remote Centera	36
8.1.4	Updating the replication configuration	37
8.1.5	Creating the PEA file	37
8.2	Shutting down a Centera	39
8.3	Centera replication management	39
8.3.1	Replication topology	40
8.3.2	Replication status	40
8.3.3	Pausing replication	41
8.3.4	Re-starting replication after a pause	41
8.3.5	Replication parking queue	41
8.4	Copying data between Centera (“restore” command) – <u>DO NOT USE</u>	42
8.5	EMC access management	42
8.6	Virtual Pool capacity reporting	43
9	CTI CENTERA CONFIGURATION SUMMARY	46
9.1	Certified Components Summary	46
9.2	Configuration rule Summary	46
9.3	Important Information Summary	48
10	GENERAL INFORMATION	49
10.1	Centera Time Synchronisation	49
10.2	Centera Retention Classes	49
10.3	Error/Heath Reporting	49
10.3.1	Email	51
10.3.2	SNMP	51
10.3.3	API Monitoring	52
11	CENTERA ARCHITECTURE	53
11.1	Single Instancing of Data	53
11.1.1	What happens if a duplicate key is generated	53
11.2	Centera Hardware	54

11.3	Data Access Security	56
11.3.1	Virtual Pools	57
11.3.2	Profiles	57
11.3.3	Access Control Lists	58
11.4	Pool Entry Authorisation (PEA) file	59
11.5	Data protection	61
11.6	Remote copies of data (Replication)	62
11.7	Performance	63
11.7.1	4-Node Centera discussion	63
12	PRODUCT RISK ACCEPTANCE AND ICQ	65
13	APPENDIX A – PERFORMANCE STATS	67
14	APPENDIX B – BLANK EMC CENTERA CHANGE CONTROL REQUEST FORM	69
15	APPENDIX C – EXAMPLE CENTERA HEALTH REPORT	75
16	APPENDIX D – EXAMPLE EMAIL ERROR REPORTS	92
16.1	Replication paused due to incorrect configuration of profile rights	92
16.2	Disk failure on node	93
16.3	Node alert	93
17	APPENDIX E – CENTRASTAR 3.0 ALERTS	95
18	APPENDIX F – CENTERA MIB	103

Diagrams/Tables

Figure 1	Centera Generation 2/3 hardware (Left) Generation 4 hardware (right)	9
Figure 2	Creation of a C-Clip	10
Figure 3	Centera nodes/capacity table	14
Figure 4	Centera management configuration	22
Figure 5	CTI certified replication topology	40
Figure 6	Centera 3.0 alerting sensors	50
Figure 7	Centera Schematic Diagram (Gen 3 Hardware)	56
Figure 8	Application access to replicated Centeras	60
Figure 9	Centera Performance (Gen 4 Hardware, #Threads vs file size)	67
Figure 10	Replication throughput (Mirrored protection 50KB files)	67
Figure 11	Replication throughput (Mirrored protection 10MB files)	68
Figure 12	Centera 3.0 Alerts	101

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.

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.

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 expanded scope, or comments on issues discovered after deployment.

Confidentiality

This document is confidential and no part of it should be shared by EMC with any of its other customers.

1 Introduction

EMC term Centera as Content Addressed Storage (CAS). This is a relatively new class of storage in the same way that Network Addressable Storage (NAS) class appeared a few years ago. Centera from EMC is the first CAS product and is aimed at static data (fixed content) rather than changing data that is still in use and hence changing. The easiest way to differentiate between SAN/locally attached disk, NAS and CAS is that:

- **SAN/locally attached disks** are accessed as blocks of data
- **NAS** storage is accessed as files
- **CAS** storage is accessed via an Application Programmers Interface (API) as objects.

This document is a major update to the previous Centera Best Practice documentation (GDSE-0322-B) and includes new features introduced in Centera 3.0 that allow Centera to be configured and used in new ways.

Note: The Centera operating system is called CentraStar and hence it is more correct to refer to CentraStar 3.0 rather than Centera 3.0.

1.1 Intended Audience

This document is intended to describe the configuration and installation/upgrading of Centera storage. This document provides more detail than is usual in a solution guide due to the lack of technical and best practice information available from EMC at this time.

This document is not intended for application developers wishing to use the Centera Software Developers Kit (SDK). Users of the SDK should use the EMC documents "Centera Programmers Guide" (P/N 069001127) and "Centera API Reference Guide" (P/N 069001185) that deal specifically with the API interface and its best practice.

1.2 Background

Today, Citigroup businesses mostly use tape or optical platters to store compliant or archived data. This is largely a result of regulations that require certain financial documents to be stored on unalterable media, typically referred to as WORM, Write-Once-Read-Many. WORM storage is any kind of storage medium that is un-tamperable. In other words, data can be written to only a single time, but can be read from any number of times. Typically this has been optical disk, where the surface is permanently etched using a laser in order to record information. The leading magnetic storage vendors have begun introducing magnetic storage solutions that are "WORM-Like". The record retention features of these solutions enable them to compete with optical disks as a WORM alternative. In addition, it is widely believed that while optical disks, such as CDs and DVDs, will continue to be popular in the entertainment industry, optical technology's future in the storage industry is nearing the end of its applicability. As a result of these factors, Citigroup businesses have begun to consider replacing their WORM storage requirements from optical to magnetic storage that meets the standards outlined in the US Securities and Exchange Commission (SEC) Interpretive Release on Rule 17a-4. SEC Rule 17a-4 is the most stringent governmental compliance with regards to technology used for record retention requirements.

The SEC Interpretive Release clarified the Commission's views on the standards for broker-dealer firms preserving electronic records and email on disk-based storage media. Specifically, the Interpretive Release affirms that broker-dealers can meet the non-rewriteable and non-erasable requirements of Rule 17a-4 through the employment of a storage solution that uses integrated hardware and software codes to prevent the overwriting, erasure or alteration of the

records. The SEC Interpretive Release also states that storage solutions that only mitigate the risk a record will be overwritten or erased do not meet the requirements of Rule 17a-4.

"Such systems - which may use software applications to protect electronic records, such as authentication and approval policies, passwords or other extrinsic security controls - do not maintain the records in a manner that is non-rewriteable and non-erasable."

The EMC Centera in its Compliance Edition Plus version meets the SEC requirements for WORM-like storage and hence has been certified by Citigroup Technology Infrastructure (CTI) to meet the company's requirement to store compliant data.

1.3 Security Classification

EMC Centera has been classified as Class C infrastructure for its security classification. This classification means that it is suitable to hold information of Confidential or lower security classifications. Any application wishing to store data of Restricted or higher security classification must apply suitable encryption or other CTI approved data protection before this information is stored on the Centera.

1.4 Release Status

The Centera solution in this guide is being released as an engineering ready production build. There are currently pending items of an RA, Standard Build, and/or SBCC that should be completed prior to a production deployment. If a production deployment is needed prior to the completion of these items, the operational team should verify and/or implement any needed processes or procedures as interim controls. For further clarification, speak with your Control Officer for guidance.

2 Centera Overview

CAS storage is intended for static data (also known as fixed content) for example voice recording, scanned documents or other archived data. Its raw performance is not as good as SAN/locally attached disk or NAS. However, when compared to either optical disks or tape CAS introduces significant performance and management benefits.

One of its major strengths is the ability to provide WORM-like storage for record retention and file management applications that specify individual retention periods for each data object stored. Centera provides a new class of storage that can replace optical media for storage of compliant data in an unchangeable format for record or for audit purposes.

Another major benefit of Centera is its ability to scale with minor management, as the data is not accessed via traditional file systems.



Figure 1 Centera Generation 2/3 hardware (Left) Generation 4 hardware (right)

There are three versions of Centera:

- Basic Edition – Retention periods are not enforced
- Governance Edition – Retention periods are enforced but they can be overridden using an audited delete command.
- Compliance Edition Plus (CE+) – Retention periods are enforced, no overrides are possible.

Only the CE+ version is being certified for use within CTI, see 3 “CTI Centera Configuration Standards”.

2.1 Using Centera Storage

It must be stressed that data is sent to the Centera directly by an application such as document management or Email archiving application rather than via the file system/volume manager as in traditional disk storage. The applications communicate to the Centera over a TCP/IP network using an Application Programmer Interface (API) supplied by EMC in their Centera Software Developers Kit (SDK). Use of the API must be built into the application itself. Hence any application using Centera must be specifically designed to work with Centera by the application vendor, or CTI development team if the application is developed in-house.

The application sends the data object to the Centera over the IP network, this object may for example be a picture, Email, document, voice recording. The Centera merely treats it as a piece of data that it needs to store. A piece of information to be stored on a Centera is referred to as a “blob”, this is a distinct bit sequence of the users data and should not be confused with a “Binary Large Object” that is also referred to as a “blob”.

The data stored within a Centera is formed of two components:

- The blob of application data to be stored
- The Content Descriptor File (CDF) that contains the meta-data.

Together the content and the CDF are referred to as a C-Clip.

When a user presents data to the Centera, the system calculates a unique Content Address (CA) for the data based on its content and then stores this reference in a newly created XML based CDF file.

The CA is created by using the M, GM or M++ naming schemes (each of which are described later). This CA can be likened to a fingerprint of the object and statistically this key, which is based on the actual binary contents of the object, will be unique.

The CDF also contains application-specific metadata. The system then calculates another CA for this the CDF, and stores the complete C-Clip package on the Centera. The second CA is a handle for the C-Clip that the system returns to the user, and it is this CA that is used to retrieve the user's data (see diagram below).




Figure 2 Creation of a C-Clip

When the application needs to retrieve the object the application must present the Content Address to the Centera; this is the only way in which the application may access the object. It is therefore critical that these keys be protected by the application by backups, replication etc, as without them you will have great difficulty retrieving your data.

Note: The only way to recover content addresses if you lose them is to scan the Centera for objects created between a pair of date/times, and look at each object in turn to discover if you have the correct piece of data (the scan can also return meta-data associated with each object). This pre-supposes that you can identify the required object when you find it and also that you know when the data was written to the Centera.

On receipt of the content address from the application the object is retrieved from disk and the object returned to the requesting application. If this object is modified and then re-saved to the Centera the contents will have changed, however slightly, and hence the Content Address generated will be different from the original object; this means that a modified object never overwrites the original.

A Centera accepts storage or retrieval instructions from any authorised application using the Centera API calls and hence many different applications may use the same Centera storage unit concurrently.

It is important to understand that as each C-Clip contains both data and a CDF file, the overhead for small and large files is the same. Hence, the overhead as a percentage of the original data can be large for small files. For this reason applications can make use of the Centera function for small amounts of data that allows data to be stored directly in the CDF record itself rather than the CDF record having a pointer to the data and this data being held in another file.

2.2 Centera Architecture Overview

The Centera architecture has been termed Redundant Array of Independent Nodes (RAIN). What this means is that each Centera is made up of many standard nodes with copies of data being held on different nodes within the Centera rather than copies of data being held on different physical disks, as is the case with RAID protection. Whilst the Centera supports both Content Protection Mirrored (CPM) and Content Protection Parity (CPP) (which can be equated to RAID-1 (Mirroring) and RAID-5, respectively) CPP is not certified for use within CTI.

Only CPM protection has been certified for use within CTI.

(See section 11.5 “Data protection” for more details).

In order to protect against the loss of a Centera unit data within a Centera must be replicated to another Centera located in a different physical location. A single Centera may only be replicated to one other Centera. Any objects stored on the a Centera are automatically transmitted to the remote site, giving a total of 4 copies of the data, two local and two remotely. Note that this replication is asynchronous and uses the same network ports as used for application access, hence the source and target Centera must have access to each other over this network port.

Centera will only replicate data that is written after replication is enabled; i.e. data written before replication is enabled will not be automatically replicated. To copy data that was written before replication was enabled requires the data to be copied manually by other mechanisms within the Centera (e.g. the Centera “restore” command).

(See section 11.6 “Remote copies of data (Replication)” for more details).

Each piece of data stored within Centera has a Content Address generated from the data itself; from this it follows that each time the same piece of information is stored it will generate the same Content Address. Hence, whilst all CDF files are unique multiple CDF files may point to the same data object. Any meta-data such as retention periods or application meta-data is placed within the CDF, this enables a multiple references by different applications to the same piece of information.

A Centera never deletes information when its retention period expires, the application must perform a delete to erase any data after the retention period expires.

During deletion the CDF is removed rather than the data itself, the data is removed by the garbage collection processes once all CDF references to a piece of data are deleted. (See section 11.1 “Single Instanting of Data” for more details).

Access to the Centera is by the use of profiles, the Centera security uses data pools and application profiles to segregate data and to grant or deny access rights to the data within those pools. Profiles can be looked upon as usernames used by applications to access the Centera, each application must have its own profile as this profile also determines which pool any data is written to.

The default, “anonymous” profile, must not be used by any application.

From CentraStar 3.0 onwards Centera has the ability to segregate data into virtual pools, prior to CentraStar 3.0 all data on a Centera was stored in a single “default” pool. Each profile has privileges that allow it to read, write, delete etc from a subset of the available virtual pools.

A virtual pool is not a physical construct like a volume but simply a way of indexing the available data whilst at the same time restricting who has access to that data. All the storage in a Centera is available to any virtual pool and hence there is no need to manually apportion storage to any particular pool.

(See section 11.3 “Data Access Security” for further details on profiles and virtual pools)

3 CTI Centera Configuration Standards

Whilst EMC provides many options for Centera, CTI has certified Centera installations conforming to the following configurations:

Centera Edition:	Compliance Edition Plus (CE+) only
Data Protection:	Content Protection Mirrored (CPM) only
Modems:	Modems must not be connected to the Centea

Notes:

Basic and Governance editions of Centera have not been certified for 3.0 (only CE+ has been certified previously) due to issues with the Centera management interface. With 3.0 there is only a single user account (admin), this account has full privileges and there is no audit trail or logging of any type with the Centera product. Due to these issues the Centera management interface ports must be protected and compensating controls introduced to overcome the lack of audit logging. In the Basic and Governance editions access to the management interface is provided via the access node Ethernet ports (i.e. the same ports used for application access). In a CE+ edition (with RPQ) management is provided by network ports on the storage nodes and hence are physically separate from application access ports. This physical separation allows the management ports to be protected as described later in this document.

EMC are currently investigating the way in which the management interface is provided across all editions of Centera, a move to segregated management ports for all editions is expected in a later release of CentraStar. When this new management network configuration or suitable role based administration and logging is introduced by EMC into Centera the certification of Basic and Governance editions will be re-visited.

Previously (CentraStar 2.2SP2) the minimum size of Centera configuration that was certified was 16 nodes with 4 dedicated access nodes and 12 dedicated storage nodes (Using EMC standard Centera notation this can be written as <# access nodes>/<# storage nodes>/<total # nodes> or in this case 4/12/16). There were primarily two reasons for this:

1. In order to provide sufficient access node capacity to both service application IO with sufficient bandwidth and perform replication the EMC minimum specification was 4 access nodes.
2. Centera storage was built with 8 node building blocks, however in order to have 4 access nodes a total of 16 nodes was required (an 8 node configuration could only have 2 access nodes i.e. 2/6/8). This meant that a minimum of 16 nodes was required to meet performance requirements and provide replication.

Between CentraStar 2.2 and CentraStar 3.0 several changes were made to the CentraStar code that dramatically improved the performance of a Centera. One of the main gains in performance was within the replication code, this now transfers batches of data rather than treating each data object as an individual transfer. Another important change was that since CentraStar 2.4 all nodes are now storage nodes, any number of these nodes may also be access nodes. Hence an 8-node Centera cluster can now have 4 access nodes i.e. 4/8/8. A Centera system with only 2 access nodes that is also configured for replication is now a supported EMC configuration.

Any Centera initially installed using CentraStar versions prior to 2.4 may be upgraded to use combined storage/access nodes when the upgrade to CentraStar 3.0 is performed. However extra licenses (CNR8NMANSW, included with new purchases since CentraStar 2.4) are required for this functionality. At least two pure storage nodes must be available

to provide management connectivity via the Ethernet ports on the storage nodes in a CE+ Centera.

With the introduction of the Generation 4 hardware and CentraStar 3.0 the Centera building blocks have been changed to 4 nodes and hence a configurations of 4, 8, 16, 24, 32 etc nodes are now possible.

The 4-node Centera is a reduced performance configuration that should only be used in circumstances where the performance of an application (or applications) can be defined and shown to be of a level that can be supported by the 4-node Centera.

In order to use a 4-node configuration the maximum concurrent mixed Read/Write file IO rate for 100KB files would be:

1. 18 objects per second if both Centera's in a bi-directional pair are being written to by applications.
2. 29 objects per second if only the primary Centera in a bi-directional pair is written to by applications.

If total application IO is above these figures then the 4-node configuration should not be used. 4-node configurations are also only available for Gen-4 hardware node configurations.

The table below gives an overview of the Centera configurations currently certified for use within CTI.

Total # Nodes	Minimum # Access Nodes	# Storage nodes	Gen 2/3 (320GB Disks)		Gen 4 (500GB Disks)		Notes
			Approx Useable Capacity	Rege n Buffer size (# Disks)	Approx Useable Capacity	Regen Buffer size (# Disks)	
4	2	4	1.68TB	2	N/A	N/A	Limited performance configuration only
8	4	8	3.36TB	4	5.61TB	4	Standard minimum system configuration
12	4	12	5.29TB	4	8.81TB	4	
16	4	16	7.20TB	4	12.01TB	4	
20	4	20	9.13TB	4	15.21TB	4	
24	4	24	11.05TB	4	17.62TB	8	
28	4	28	12.97TB	4	20.82TB	8	
32	4	32	14.92TB	4	24.02TB	8	Max number of nodes in a cabinet (Centera can be formed of a maximum of 4 cabinets)

Figure 3 Centera nodes/capacity table

Capacity figures make the following assumptions:

- All systems with 320GB disks are generation 2/3 hardware.
- 320GB Drives have 246GB usable after formatting and internal CentraStar data.
- All systems with 500GB disks are generation 4 hardware.

- 500GB Drives have 410GB usable after formatting and internal CentraStar data.
- The usable capacity **does not** include the space reserved for the regen buffer in the event of hardware failure.
- Generation 2/3 hardware has a max cube size of 32 nodes.
- Generation 4 hardware has a max cube size of 16 nodes.
- Regen buffer of 1 Disk - Hard Stop for smallest cubes having 4 nodes, 2 Disks - HardStop for smallest cube having 8+ Nodes. **(Note: each disk reserved for the regen buffer is mirrored capacity and hence twice the capacity is reserved and unavailable for application use, the actual amount of space reserved is shown in the table)**
- Nodes are added to the same cube until this cube is filled, only then is a new cube started.

For Centera clusters made up of multiple cabinets the total capacity can be calculated by working out the number of nodes per cabinet, looking up the usable capacity for each cabinet and adding the these figures together to provide a total capacity.

The base configuration that should be used in normal circumstances is 4/8/8, if there is any doubt as to the level of performance required the 4/8/8 configuration or higher should be used.

Except on the 2/4/4 configuration extra access nodes above 4 may be configured to increase application IO bandwidth. The only caveat to this is that a minimum of 2 pure storage nodes is required to provide management connectivity.

3.1 Centera hardware expansion

As of September 2005, only Gen 4 hardware is being shipped and hence any expansion of existing Gen 2/3 Centera will be with Gen 4 hardware. Whilst Gen 4 hardware is compatible with existing Gen 2/3 Centera units, due to the requirement for the installation of Gen 4 power equipment into the rack (Gen 4 units have dual power feeds and do not use the existing ATS power unit used by Gen 2/3 units) the maximum number of Centera nodes in a rack is reduced by 4.

Hence an existing 4/12/16 configuration (the minimum configuration previously certified) can only be expanded with an additional 12 Gen 4 nodes added making a total of 24 nodes within the cabinet. If Gen 3 hardware were available the maximum number of nodes in a cabinet would have been 32.

When expanding a Gen 4 Centera cluster each cabinet may contain up to the maximum 32 nodes as only the Gen 4 power equipment is already installed (there is no ATS power switch).

3.2 Centera Universal Access

The Centera Universal Access (CUA) is an appliance for using Centera as a file system; this allows data to be stored onto a Centera accessed via the CUA using NFS or CIFS. If this type of functionality is required then the use of NetApp filers is recommended.

The use of the Centera Universal Access (CUA); previously known as the Centera Access Gateway (CAG), is not certified for use.

3.3 Multi-Cabinet Centera Clusters

The use of multi-cabinet Centera cabinet is permitted up to the current EMC **maximum of 4 cabinets per Centera cluster**. As IO will be spread across all cabinets **a minimum of 4 access nodes must be configured per cabinet** to reduce inter-cabinet traffic within the Centera. As mutli-cabinet configurations will hold substantial amount of data it is important that the correct number of access nodes are configured to support the workload (4 access nodes per cabinet is a minimum) by working with EMC to understand the workload for the cluster and configure an appropriate number of access nodes.

4 Preparation for installation

As part of the pre-installation of a Centera unit EMC will require a Change Control Request Form (CCRF) to be completed. This form includes the following information:

- Installation address
- Centera operational contact information (name, telephone, Email etc)
- Number of storage and access nodes
- Remote Dial-in and Alerting options
- Network connectivity to CTI networks including IP addresses, subnets and IP gateways.
- A list of applications to use the Centera including the version of Centera SDK used.
- Replication configuration information

All Centera access nodes within a Centera cluster should be connected to the same subnet but for resilience purposes should be connected to different physical infrastructure IP switches.

See “14 Appendix B – Blank EMC Centera Change Control Request Form” for a copy of a blank CCRF.

Prior to installation of the Centera by EMC, CTI personnel will need to provide the following information/infrastructure:

- Primary and Secondary (optional) contact information for the Centera
- Number of access nodes to be configured
- A single 100Mb Full duplex or 1Gb Full duplex connection with UTP CAT 5e (or Cat 6) network cable is required for required for each access node configured, the maximum length of this cable between the Centera and its network port is 30m. For the Generation 4 hardware Cat 6 cables are required rather than Cat 5e. An IP address, subnet mask and default gateway will also be required for each access node. (Auto-negotiated link speeds should not be used)
- SMTP Gateway address (two Gateways may be defined)
- Reply address for service Emails (this Email address should be a valid Citigroup Email address that is externally visible e.g. xxxxxx@citigroup.com, the creation of this address requires local BISO sign-off before this address can be created by the Global Operations Messaging Systems Support group). This Email address should point to an internal distribution list specifically created at each site and containing the BAU personnel responsible for managing the Centera at this site. This address can also be used as the Email address to be sent the daily health messages and any Centera fault Emails.
- A name for the Centera cluster
- Installation of Management server(s) and network infrastructure.

As part of the CCRF the following settings should be requested from EMC

CCRF Section	Setting	Comments	CTI Value
4a	Installation of modem on cluster	Modems must not be installed on even a temporary basis.	No
4b	Outbound SMTP (Email) alerting		Yes
4c	OnAlert Workstation		No
6	ConnectEMC mode		On
6	ConnectEMC Reporting Interval		1 Day

CCRF Section	Setting	Comments	CTI Value
6	Additional delay before next report		None
6	Domain to which cluster belongs		citigroup.com
6	Data shredding	Erase data with 7 passes to bit patterns to ensure data cannot be read.	On
6	Garbage collection	Deletion of data blobs after all c-clips pointing to a data blob have been removed.	On
6	Incremental GC	Garbage collection	On
6	Full GC	Garbage collection	On
6	Set M++	MD5 Naming Scheme, M++ uses MD5 plus 120 bits of SHA-256 hash.	Enabled
6	Set protection write parity enabled?	Content Parity Protection (CPP) is not certified for use	No
6	Regen buffer: Mode: Alert or Stop		Stop
6	Limit in disks per cube	Number of disks to be reserved for the regen buffer. The number of disks depends on the size of the smallest cube in the cluster. Smallest Cube Regen 4 nodes 1 Disk 8+ nodes 2 Disks	1 Disk or 2 Disks
6	Enables PING capability	Whilst by default PING should not be enabled if this is required locally then PING may be enabled.	No
6	Security default retention (seconds, minutes, days etc)		Not Used
6	Enable SNMP	If SNMP monitoring is to be used enable.	Yes

4.1.1 Replication

For contingency purposes, all Centera installations must be paired with a remote replica to protect against hardware or site failure. As data is not continually changed on a Centera the requirement for protection in the case of hardware failure or disaster is different to traditional storage. Once an individual piece of data has been saved to a remote location and provided this replica continues to exist it need never be copied again (data integrity is ensured by the disk scrubbing and comparison of the files data with the computed CA value). Whilst enterprise tape backup products are available for Centera, given the typical amount of data held within a Centera and the unchanging nature of the information, the need for tape backup to archive Centera data is not necessary and thus not certified for use within CTI.

EMC support several replication configurations however at this time GDSE is only certifying the use of bi-directional replication between a pair of Centera clusters.

NOTE: Each pool on a local Centera that requires replication must replicate to the same remote Centera. Whilst individual pools may have replication disabled all production data should be configured for replication.

A minimum of 4 access nodes on each Centera cluster should be configured as replication destinations to provide resilience and ensure sufficient I/O throughput for replication between Centera clusters. (Limited performance configurations must have two access nodes configured)

In order to configure the replication the following questions need to be answered:

CCRF Section	Setting	Comments	CTI Value
6	Replication address		Enter IP addresses of the all access nodes of the paired Centera.
6	Replication enabled		Yes
6	Replicated delete	A deletion of data at the primary Centera will cause the replica data to be removed at the remote Centera This is set cluster wide, all applications using this Centera must support replicated deletes before this is enabled. The best practice is for this to be enabled unless specifically required to be disabled.	Yes
6	Replicate incoming replicated objects?	This is used in chain replication A->B->C, this topology is not current certified	No
6	Replicate all pools	Options for replication are essentially replication "on" or "off" and is set at a per pool level. As there is no other mechanism for protecting data in a Centera data in all pools will normally be replicated.	Yes
6	Replication profile name	Name of access profile on the remote Centera to use for replication	This name will be locally defined but an example would be rep_<Centera> (where Centera is the cluster name of this Centera). This access account will be created on the remote Centera.

4.1.2 Storage-Strategy

Data can be stored on a Centera using one of two Storage Strategies: Storage Strategy Capacity or Storage Strategy Performance.

Storage Strategy Capacity ensures single instancing and thus optimises capacity.

Storage Strategy Performance improves the speed of write operations at the cost of single instance storage. When Storage Strategy Performance is enabled, identical content may be stored multiple times. This strategy should be used as a default on all CTI Centera.

Setting the Storage Strategy to Performance does not improve performance substantially for large files (> 1 MB). When setting the Storage Strategy to Performance the default threshold (maximum file size) is 250 KB. Objects larger than this threshold will revert to Storage Strategy Capacity in order to benefit from single instance storage. The threshold can be adjusted to optimise the balance between performance and capacity.

CCRF Section	Setting	Comments	CTI Value
6	Storage-Strategy	Ensure all applications using the Centera can run in performance mode. In "capacity" mode the M++ naming scheme is used. In "performance" mode the GM (or M++ above the threshold) naming scheme is used	Performance
6	If performance – full or partial	Full – Both CDF and blob addresses use M++ naming Scheme. Partial – Only CDF's use M++ content addressing scheme	Full
6	Threshold for storage-strategy	Default is 250KB, depending on the application you may wish to change this value. Below the threshold the GM naming scheme is used, above the threshold the M++ naming scheme is used.	250KB, unless local applications require a different value.

4.1.3 Current RPQ Requirements for CTI Centera

There are two CentraStar 3.0 configuration parameters that do not comply with the operational standards required by CTI. EMC has committed to enhancements in later versions of CentraStar, until such time every new Centera install will require an EMC RPQ to be completed by the local EMC sales team. An RPQ is a non-standard installation request.

The two parameters requiring the RPQ are Email Event and Health Reporting, and Remote Management of Compliance Plus Edition Centera.

Email Monitoring – In order to send Email health reports to EMC from the Centera units via the Citigroup mail routers the "from" address for all Emails must be changed to a Citigroup valid address. By default the "from" address of all Emails sent from a Centera is *ConnectEMC@emc.com* this is not a valid Citigroup address. Emails being sent external to Citigroup without a valid Citigroup "from" address will be dropped at the mail router. The RPQ allows the "from" address to be changed to a valid Citigroup supplied Email address. **RPQ 47393** should be referenced; this RPQ was requested by Alan Davison (EMC London) on Citigroup's behalf and submitted on 22nd September 2005. EMC has stated that the "from" address will be customer configurable via the management interface in the 3.1 release of CentraStar.

Remote Management – In the standard Compliance Plus (CE+) edition of Centera remote management is not permitted, which does not allow CTI to manage their Centera system in a manner that meets both: controls, and is operationally reasonable. Currently, in order to manage a Centera CE+ a computer running the Centera management software must be manually connected to the Ethernet port on a Centera storage node. This system may not be left connected to the Centera and neither may this network connection be connected to the corporate network.

After discussions with EMC it has been agreed that EMC will modify the Centera code to allow remote management of CE+ Centera (the target version for remote management is the 4.0 release of CentraStar). Until such time that remote management is available using physically separate network interfaces in the standard product. The local EMC account should reference **RPQ 49262**, requesting Centera CE+ remote management using the storage nodes' network ports. RPQ 49262 was requested by Joe Malinka (EMC New York) on Citigroup's behalf and submitted on 10th November 2005. The implementation of this RPQ will allow the management network discussed later in this document to be constructed, this infrastructure will also support the proposed changes to the management mechanism in release 4.0 without modification.

NOTE: an RPQ is merely an approval process, an EMC Centera engineer is required to configure each Centera to be installed in order to provide remote management and change the "from" Email address.

The nature of the changes made to the Centera as part of the RPQ process mean that EMC will need to check and possibly re-install these changes after adding new nodes or upgrading the CentraStar code in any way. This will continue until such time that these features are contained within the standard CentraStar releases and the RPQ's are no longer required.

4.1.4 Management Server/Infrastructure

As part of the pre-installation work for the Centera the management infrastructure needs to be created. The management configuration assumes that the RPQ for the use of storage node ports for management traffic on the Centera has been completed and the modifications made to the Centera.




Figure 4 Centera management configuration

The Centera has two sets of Network connections:

Access Nodes: 1Gb/s used for application data transfers

Storage Nodes: 100Mb/s used for management traffic (these are actually 1Gb/s ports but the extra bandwidth is not required for management traffic)

The access nodes are connected directly to the corporate network and are used for all application data traffic using TCP port 3218.

The two storage nodes are connected to the local Common Extranet Platform (CEP) firewall environment. This is a standard network infrastructure for segregating vendor equipment but also allows for the segregation of equipment from the normal corporate network. Each device on the CEP is placed into a VLAN within the CEP and can only communicate with other members of the VLAN or systems on the corporate network via the CEP firewall. Each storage node should be placed in a different VLAN and connected to different CEP hardware to remove single points of failure (storage nodes may share VLAN's with other storage infrastructure such as CLARiiON, DMX service processors or Connectrix manager servers). The CEP firewall should be configured to allow the management servers to access the storage nodes using TCP port 3682.

The Centera management servers should be Sun Unix servers running Solaris 8.0 and built to the current CTI standard with the packages shown below loaded:

PBCInt3.5.2-01_A1	(Powerbroker client package)
CentCLI3.0_A0	(CenteraCLI package)

The Centera package is located on /net/stealth-nj/export1/home1.localhost/sw/PKG.

The Solaris server should have PowerBroker configured to access a PowerBroker server and this server should have its pb.conf file configured to give access to the CTICenteraCLI image to only authorised Centera administrators running on the Centera management servers. PowerBroker should also be configured to provide keystroke logging of log all IO generated by the CenteraCLI software into a file on the PowerBroker server. These log files should be retained to comply with local controls processes. Details of how to install PowerBroker on a client may be found in the GDSE document "PowerBroker Implementation for SAN – Solution Design" this can be found at http://san.ny.smb.com/GDSE_docs/GDSE-05006-S.pdf.

An example of the type of code that is required in the pb.conf on the PowerBroker server to support the Centera is shown below:

```
Centera_mgmt_hosts = {"Centera-mgt1.eur.nsroot.net", "Centera-mgt2.eur.nsroot.net"} ;
Centera_users = {"xx12345", "yy67890"};
Centera_progs = { "CTICenteraCli" };

if ((user in Centera_users) && (basename(command) in Centera_progs) && (submithost in
Centera_mgmt_hosts)) {

runuser = "root";
rungroup = "!g!";
rungroups = {"!G!"};
runcommand = basename(command);

iolog = logmktemp("/var/log/pb/pb."+ user + "." + basename(command) + ".XXXXXX");

setenv("PATH",
"/bin:/usr/bin:/usr/sbin:/export/opt/EMC/Centera/3_0/SystemOperator/bin");
accept;
}
```

4.1.5 Running the Centera management interface

The Centera CLI must be run from the nominated Centera management station, this management station will have the management software installed and more importantly be the only servers permitted through the CEP firewall to contact the Centera management ports. A user should log onto the Centera management server using their standard personal SOE Unix account. The Centera CLI will not be accessible from non-privileged accounts to ensure that the command can only be run via PowerBroker. The command below should be executed to invoke

the CLI management interface (obviously the \$PATH environment variable can be used to include both the PowerBroker and Centera software directories and hence reduce typing).

```
$/export/opt/pb/bin/pbrun /export/opt/EMC/Centera/3_0/SystemOperator/bin/CTICenteraCLI
CenteraCli
```

```
Initializing CenteraCli...
```

```
Username: admin
```

```
Password:
```

```
IP: 146.128.135.117
```

In 3.0 "admin" is the only customer account

The password typed is not echoed

The IP address of one of the management ports

```
Centera interactive shell (Implementation CLI 3.0.609, 2005-05-02)
```

```
Connected to RemoteManager at 146.128.135.117:3682 using FRMP version 2
```

```
Type "help" for available commands.
```

```
Config#
```

From this point Centera CLI commands may be issued.

Command recall and command line editing in the current version of CenteraCLI does not function correctly, the delete key however does function correctly. EMC will fix this bug in a later version of the CenteraCLI.

5 Installation of Centera

Installation of the Centera is performed by EMC in accordance with the CCRF document completed prior to installation. It is the responsibility of CTI to ensure that an appropriate location, power, network and management is available prior to EMC arriving on site.

See the EMC Centera specification document for details of the space, power and environmental requirements. (Currently this document can be found on the EMC website in the location shown http://www.emc.com/pdf/products/centera/centera_ss.pdf)

A Centera cabinet contains 16, 24, or 32 nodes (prior to CentraStar 3.0), after CentraStar 3.0 and the introduction of Gen 4 hardware expansion is in multiples of 4 nodes up to a maximum of 32 nodes per cabinet. The installation procedure for each cluster configuration is the same. There are multiple steps in this procedure:

- Connecting a Centera cluster to the customer LAN.
- Powering up a Centera cluster.
- Booting the Centera cluster.
- Initializing the Centera internal LAN.
- Configuring a Centera cluster (including setting up replication).
- Verifying the Centera cluster.

When the cluster is up and running and has been verified by the EMC Customer Engineer, the cluster will be handed over to CTI.

Note: Network connections should not use “Auto” for Speed or Duplex settings but should be hard coded to the appropriate setting. The cable distance between the Centera and the network switch/router should be no longer than 30m.

6 CTI Configuration

Once EMC have handed over the Centera after the installation the local CTI administrator should check that the Centera is configured as per CTI standards as noted in the CCRF supplied to EMC.

At a minimum the following checks should be made:

Check to be performed	CenteraCLI command
Check required number of storage and access nodes are configured, available and have the version of Centera is set to CE+	<code>show security all</code>
Check access node IP addresses are configured correctly	<code>show ip list</code>
Check data-shredding is enabled and the storage-strategy, performance mode, threshold are set correctly	<code>show features</code>
Check that the regen buffer mode is set to "Hard" and the buffer limit is either 1 or 2 depending on the number of nodes in the smallest cube	<code>show capacity regenerationbuffer</code>
Check that the anonymous profile has been disabled	<code>show profile detail anonymous</code>
Check that the Centera has the correct date/time set, during testing the Centera time has been incorrectly set by EMC during installation on several occasions. Surprisingly there are currently no commands that show the time of the Centera cluster and the only way that an administrator can check the clusters time is indirectly via the health report. EMC will be introducing a CLI command to show the cluster time in a later release (expected to be the CentraStar "G")	To generate a health report from the Centera CLI type: <pre>show report health <local-file.xml></pre> where <local-file.xml> is a file on the local server where the xml formatted health report can be written. Open the XML health file and look for the <FileLastUpdated> tag, this is about 28 lines down from the first line of the file. Within this tag should be a date and time, this is the date/time of the cluster at the time the health report was run.

As soon as the Centera is handed over to CTI the admin password should be changed from its default to a value known only by the local Centera administration team. It should be understood that in order to access the Centera for management purposes administrators will need to access the CenteraCLI binary on the Centera management node(s) from their own SOE account and have their sessions controlled and logged via PowerBroker. Therefore the administrator password only plays a minor part in securing the Centera.

6.1 CTI specific configuration

To ensure that only data with the retention time set to appropriate maximum/minimum values are stored on the Centera the retention time constraints must be set. These set the minimum and maximum retention times that will be accepted, these values are set globally for all profiles and pools in the Centera. If an application tries to store data that meets one of the following criteria the Centera will reject that particular write:

- No retention time is specified (this defaults to infinite on a CE+ Centera)
- Retention time is less than the minimum value set as a constraint globally on the Centera.

- Retention time is greater than the maximum value set as a constraint globally on the Centera.

The actual values for the retention maximum/minimum settings will depend on local applications using the Centera and their particular retention requirements. The minimum value should be set to the lowest valid retention time of all the applications using the Centera, this must be greater than zero to ensure that no application uses a zero retention time as this has been used in the past by applications failing to meet best practice and bypassing the setting of retention times on Centera when data is written. The maximum value is set to the largest valid retention time for any application using the Centera and should not typically exceed 20 years.

The commands below sets the maximum/minimum retention constraints to a Minimum of 1 month and a Maximum of 8 years, both commands must be issued for the change to take effect.

```
Config# set constraint check_retention_present
Mandatory retention period in clip? (true, false) [false]: true

Issue the command?
(yes, no) [no]: y
Config# set constraint check_retention_range
Minimal retention period: 1 month
Maximal retention period: 8 years
Issue the command?
(yes, no) [no]: y
```

```
Config# show constraint list
```

```
Name                Constraint
-----
check_retention_present true
check_retention_range min: 1 month - max: 8 years
-----
```

```
Config#
```

Centera has a security setting that when enabled prevents all EMC accounts (there are several standard accounts on a Centera used by EMC customer service and EMC engineering for maintenance/fault finding purposes) from being used; this is the "lock" setting. All nodes should be set to "locked" to prevent EMC access to the Centera without CTI authorisation. Nodes should only be unlocked when EMC personnel require access to the Centera and local CTI Change Control processes have validated their requirement for access to the Centera.

To show the current state of the security on the nodes issue the command shown below (note the example shown below is for a Governance Edition Centera, ignore fields other than "Locked"):

```
Config# show security all
Default Retention Period: 0 seconds
```

Cube	Node	Roles	Status	Locked	Compliance
1	c001n01	access	on	Yes	CE
1	c001n02	access	on	Yes	CE
1	c001n03	access	on	Yes	CE
1	c001n04	access	on	Yes	CE
1	c001n05	storage	on	Yes	CE
1	c001n06	storage	on	Yes	CE
1	c001n07	storage	on	Yes	CE
1	c001n08	storage	on	Yes	CE
1	c001n09	storage	on	Yes	CE
1	c001n10	storage	on	Yes	CE
1	c001n11	storage	on	Yes	CE
1	c001n12	storage	on	Yes	CE
1	c001n13	storage	on	Yes	CE
1	c001n14	storage	on	Yes	CE
1	c001n15	storage	on	Yes	CE
1	c001n16	storage	on	No	CE

This node is unlocked and EMC's standard accounts are active on this node

Config#

To lock all nodes that may still be unlocked issue the command shown below (this single command locks all nodes):

```
Config# set security lock
Issue the command?
(yes, no) [no]: y
```

Config#

The security status of all nodes should then be checked via “show security all” to ensure that all nodes now show as “Locked”.

7 Upgrading an existing Centera

Like a new Centera install, the upgrade process of the CentraStar software will be performed by EMC professional services. Before the upgrade EMC will require that a CCRF form is completed that details the configuration required after the upgrade, the contents of the CCRF will be the same as the entries detailed in the section 4 “Preparation for installation” section.

There are no pre-requisites for the upgrade of Centera from 2.2 to 3.0 and the upgrade may take place either online or offline. The upgrade from 2.2 to 3.0 takes place in two parts; firstly 2.2 to 2.3.3 and secondly 2.3.3 to 3.0.

If the upgrade takes place online there is no requirement for applications to reduce IO during the migration although some reduction in performance may be experienced during the upgrade especially in clusters with a small number of access nodes. This reduction in performance is due to the requirement to reboot nodes on a rolling basis, hence at any time during the upgrade either a storage or an access node will be offline. As each piece of data is held on at least two nodes at no time is any data unavailable during the upgrade. Each node takes between 10-30 minutes to reboot and hence the total time required for an upgrade depends on the number of nodes in the cluster plus the pre and post install tasks such as checking the health of the Centera unit.

Replication need not be suspended between Centera during the upgrade process although replication may be suspended if preferred. If the Centera upgrade is performed online and replication is suspended all outstanding replication data will be queued until replication is resumed.

If the upgrade takes place offline all nodes are rebooted at once and hence the elapse time for the upgrade is less. Typically EMC will request a 2-4 hour slot to upgrade the Centera offline, during which time the Centera will be unavailable, this time will include all pre and post work as well as the cluster reboot. In a Centera with n nodes the Centera will remain offline until such time as n-1 nodes are back online.

7.1 SDK Compatibility with CentraStar 3.0

All versions of the SDK are compatible with later versions of CentraStar although not all functionality in the newer version of CentraStar will be available to applications running earlier versions of the SDK.

An issue with CentraStar 3.0 has been discovered where an application fails a “clip_exists” API call after an upgrade to CentraStar 3.0. This is due to the clips being in a transient state as they are moved into the default pool and applications not being able to perform the “clip_exists” call against clips that have not yet been processed. Whilst few applications use this API call, EMC should be requested to install CentraStar 3.0 hotfix 1 that fixes this issue. Further information on this issue can be found on the EMC powerlink site and using ETA emc118842. (CentraStar 3.0SP1 and above contain this fix)

Depending on the current version of Centera and Centera SDK used, different post upgrade tasks will be required to use the virtual pool features introduced in CentraStar 3.0. All profiles from CentraStar 3.0 onwards should use their own virtual pools (provided they use SDK 2.3.315 or higher. Applications using a lower version of the SDK should upgrade to version 2.3 or higher if at all possible, otherwise they must continue to use the default pool)

Only applications integrated with SDK 2.3 (Level 315 also known as phase II) or above are able to use virtual pools. Applications that are able to use SDK 2.3.315 should use this SDK version, if this is not possible then the application will need to continue writing data into the default pool.

All new applications designated for Centera storage must use SDK 2.3.315 or above as versions prior to SDK 2.3.315 do not support the use of virtual storage pools.

7.2 Changes required in order to start using virtual pools

The following tasks need to be performed for each access profile in order for existing applications to start using virtual pools:

- Create a new virtual pool to be associated with existing profile.
- Grant the appropriate access rights to this new virtual pool to the existing profile (typically this will be “rweqd”).
- Export the new pool information and import data onto replication target Centera cluster.
- Grant rights to the new pool for the access profile on the replication target Centera
- Grant access rights to the new pool to the access profile used for replication. As bi-directional replication is used the replication access profiles on both Centera clusters (source and target) should be configured with access to their own local pool.
- Ensure that the application is offline when then next stages are performed as otherwise there is the risk that an application will temporarily lose access to some of its data.
- Change the home pool for the profile from “default” to the new virtual pool name.
- If old data (currently in the “default” pool) was written by the application using version 2.3 of the SDK or higher then move the existing data currently stored in the “default” pool by this profile to the new virtual pool.

The migration time from the default pool to the new virtual pool is relative to the number of c-clips to be moved, typically around 5-10 clips per second. For example, a full cluster with 25 million objects, (assuming a 50/50 split between c-clips and blobs) would take approximately one month to complete. Applications can be online during the migration.

Once the migration has finished then the final steps will need to be completed. Depending on the application there will be two possible scenarios:

- If none or only some of the application data could be moved from the “default” pool to the new virtual pool then whilst new data will be written to the home pool the access profile must retain access to the default pool in order to access (or dispose of) old data. As the previous certified version of CentraStar was 2.2SP2 with applications that are likely to use an SDK prior to 2.3.315 it is likely that this will be the case in most upgrades within CTI.
- If all application data has been removed from the “default” pool for this profile then all access rights to the “default” pool should be revoked for this access profile.

If an application currently uses the “anonymous” profile (i.e. does not conform to CTI standards) then a new application specific profile should be created as detailed in 8.1 “Adding a new application” and the application configured to use this profile in order to comply with the CTI standards. As the anonymous profile has been used existing C –Clips cannot be migrated, hence the new profile should retain access to the default pool as with pre-version 2.3 SDK applications.

Note: From version 2.3 of the SDK the access profile name was saved in the C-Clip meta-data by the SDK and hence C-Clips stored by each profile may be identified. Prior to version 2.3 of

the SDK this information was not stored and hence it is not possible to differentiate between C-Clips saved by different access profiles.

Example: The example below is for an application “A” using the profile AppA, with a new virtual pool PoolA. This application has always used v2.3 of the SDK and hence its data should be migrated and access to the “default” pool removed after the migration.

```
Config# create pool PoolA
Pool Mask [rdqeDcw]: rdqecw
Pool Quota [1 GB]: 6GB
Issue the command?
(yes, no) [no]: y
Created pool PoolA with ID b209b196-1dd1-11b2-afa8-8e107366e86f-4.
```

```
Config# set grant PoolA AppA
```

```
Granted Pool Rights for Profile: rdqew
Issue the command?
(yes, no) [no]: y
```

```
The new Effective Rights:
Pool Name      Profile Name      Effective
-----
PoolA          AppA              rdqe--w-
```

```
Config#set grants PoolA rep_centera2
```

```
Granted Pool Rights for Profile: cd
Issue the command?
(yes, no) [no]: y
```

```
The new Effective Rights:
Pool Name      Profile Name      Effective
-----
PoolA          rep_centera2      -d---c--
```

```
Config#export poolprofilessetup
Export complete setup? (yes, no) [no]: n
Export based on pool or profile selection? (pool, profile) [pool]:
Export all Pools? (yes, no) [no]:
Pools to Export: PoolA
Please enter the pathname where the export file should be saved: /var/tmp/poola.txt
```

Exported 1 pool and 2 profiles.

```
Config#
```

On the target of the replication from this Centera cluster, the following commands need to be run

```
Config# import poolprofilessetup
Please enter the pathname for the pool import data: /var/tmp/poola.txt
Found 1 pool and 2 profiles.
```

Step 1: Importing Pools

=====

```
Import pools? (yes, no) [yes]: y
```

```
Pool Name      ID                                     Mask      Profiles
-----
PoolA          b209b196-1dd1-11b2-afa8-8e107366e86f-4  rdqe-cw      2
```

```
Import pool PoolA? (yes, no, abort) [yes]: y
Imported pool PoolA.
```

Step 2: Importing Profiles

=====

```
Import profiles? (yes, no) [yes]: n
```

```
Config#set grants PoolA rep_centera1
```

```
Granted Pool Rights for Profile: cd
Issue the command?
```

(yes, no) [no]: y

The new Effective Rights:

Pool Name	Profile Name	Effective
PoolA	rep_centera1	-d---c--

Config# set grants PoolA AppA

Granted Pool Rights for Profile: rdqew

Issue the command?

(yes, no) [no]: y

The new Effective Rights:

Pool Name	Profile Name	Effective
PoolA	AppA	rdqe--w-

Config#

At this point (or earlier if preferred) the application should be offline. The command below should be run on both source and target Centera to set the pool to use for writes for this access profile.

Config# update profile AppA

Profile Secret [unchanged]:

Enable Profile? (yes, no) [yes]:

Monitor Capability? (yes, no) [no]:

Profile-Metadata Capability? (yes, no) [no]:

Home Pool [default]: PoolA

Issue the command?

(yes, no) [no]: y

Establish a Pool Entry Authorization for application use? (yes, no) [no]:

Config#

Once the home pool of the application has been updated the application may be re-started. All writes will now be placed in the new virtual pool and the next stage of the process will be to migrate existing data in the "default" pool to the new home pool. At this point the application still has access to the "default" pool as its historic data is still located in this pool.

Note: Any queries/reads/deletes etc from a profile that has rights to multiple virtual pools will transparently access all the authorised pools. For example a C-Clip query command (this lists all C-Clips between created between the supplied dates) will list all C-Clips that exist between these dates in any of the pools that the profile has access to.

7.3 Migrating old data application to new virtual pools

Once the new home pool has been created for a profile the migration of all data stored by this profile (using SDK v2.3+) in the "default" pool may begin. **The migration only takes place on the Centera cluster on which the command is issued**, hence the migration commands need to be issued on both the source and target Centera.

Config# create poolmapping AppA PoolA

Issue the command?

(yes, no) [no]: y

Added the mapping to the mapping scratchpad.

Config# migrate poolmapping start

Do you want to start a migration process based on the scratchpad mappings? (yes, no) [no]: y

Capacity information might become incorrect.

Do you want to recalculate it? (yes, no) [yes]:

Issue the command?

(yes, no) [no]: y

Committed the scratchpad mappings and started the migration process.

Config#

As the migration process may take an extended period (a month is possible) the progress of the migration should be checked at regular intervals. The command below shows how the progress of the migration may be monitored (the example below shows a completed migration):

```
Config# show pool migration
Migration Status:          finished

Average Migration Progress: 100%
Completed Online Nodes:    12/12
ETA of slowest Node:       0 hours

Profile Name               Active Pool Name
-----
AppA                       PoolA
-----
```

Config#

Migration cannot physically move data as the C-Clips cannot be re-written, the migration process affects the pool indexes and which virtual pool the objects appear in. If the object is written to the Centera using the SDK 3.0 or greater the pool that an object belongs to is written into the CDF meta data and hence is independent of any mappings.

Any mappings should be left in place after the migration as the active mappings are used to rebuild the virtual pool indexes in the event of any index corruption. If the active mapping are removed and a rebuild occurs data will be displayed back in the default pool.

8 Management Tasks

8.1 Adding a new application

This task is probably the most common task that will need to be performed on a Centera. In order to add a new application to the Centera the following must be accomplished:

- New virtual pools and profiles must be created on both the local and replica Centera clusters,
- The replication setting must also be updated to replicate the new pool.
- Lastly, the Pool Entry Authorisation (PEA) file containing the credentials to access the application profile must be created.

A PEA file is an XML file containing cluster information, a profile name and a secret (a secret is an alphanumeric value much like a password) to be used by an application to gain access to the profile named in the file. The PEA file functions similar to an ssh key file (for more details see 11.4 “Pool Entry Authorisation (PEA) file”).

All applications should be configured to use individual profiles and virtual pools; this has the following benefits:

- Each application only has access to data they have stored. In previous versions of Centera where all data was stored in the same pool an application had the ability to read and perhaps delete (subject to retention checks) any other applications data if it could either find or guess a c-clip address.
- As each application stores its data into a different virtual pool and capacity statistics can be run against each virtual pool it is possible to provide usage information and hence provide charge-back. This allows the Centera to be used in a similar manner to other consolidated network storage within CTI.
- In future versions of CentraStar functionality that is currently provided at the Centera level (i.e. for the whole storage system) is planned to be available at the pool level. This means that in the future it may be possible to:
 - o Replicate pools to different Centera
 - o Provide different retention max/min
 - o Enable/Disable replicated delete

All of these facilities rely on the use of virtual pools and hence applications should use a version of the SDK that allows the use of Virtual pools (SDK 2.3.315 – also known as SDK 2.3 phase II - or above).

Any application profile or virtual pool should be assigned the minimum privileges required in order for the application using it to function correctly (remember a virtual pool will require a “c” clip-copy privilege for the replication profile even though this is not required for the application profile itself)

8.1.1 Creating a new virtual pool

The first stage in creating access for a new application is to create the virtual pool to be the repository for the applications data, initially this pool is only created on the source Centera. Before creating the pool an idea of the possible data growth of the pool will be needed as the pool creation command asks for a setting for the Pool quota, the default being 1GB (setting a value of -1 indicates an infinite quota). This quota is a soft quota and if the amount of data stored in the pool exceeds the value set and error message is sent but data can still be saved to the pool by the application. If a value for the pool quota value is not known then set this value high to ensure that spurious error messages are not generated.

For the Centera CLI issue the “create pool” command as shown below for the creation of the “Finance” virtual pool. Full privileges have been granted to the pool (except privileged delete) to ensure that rights are not denied to either the application or replication when they are created.

```
Config# create pool Finance
Pool Mask [rdqeDcw]: rdqecw
Pool Quota [1 GB]: 300GB
Issue the command?
(yes, no) [no]: y
Created pool Finance with ID b209b196-1dd1-11b2-afa8-8e107366e86f-3.

Config# show pool detail Finance

Centera Pool Detail Report
-----
Generated on 07 September 2005 15:52:04 BST

Pool Name: Finance
Pool ID: b209b196-1dd1-11b2-afa8-8e107366e86f-3
Pool Mask: rdqe-cw
Cluster Mask: rdqe-cw-m-

Pool Quota: 300 GB
Used Pool Capacity: 0 GB
Free Pool Capacity: 300 GB
Number of C-Clips: --
Number of Files: --
Number of scheduled tasks: 0

Granted Rights to Access Profiles: none

Pool Mappings: none

Config#
```

8.1.2 Creating a new application profile

Once the new home virtual pool has been created the new application profile may be created using this pool on the source Centera.

In the example below a profile “Finance” is created using the virtual pool created in the previous example and the PEA file for access to the source Centera is written to a file in a local directory. In this example the application requires the following permissions to the pool:

- Read
- Delete
- Query
- Exists
- Write

Notice that the clip-copy permission is not required as this is only required for replication profiles. The option to let the Centera auto generate the profile secret should be chosen unless there are specific application reasons to use a user chosen value.

```
Config# create profile Finance
Profile Secret [generate]:
Enable Profile? (yes, no) [no]: y
Monitor Capability? (yes, no) [no]: n
Profile-Metadata Capability? (yes, no) [no]: n
Profile Type (access, cluster) [access]:
Home Pool [default]: Finance
Granted Rights for the Profile in the Home Pool [rdqeDcw]: rdqew
Issue the command?
(yes, no) [no]: y
Establish a Pool Entry Authorization for application use? (yes, no) [no]: y
Please enter Pool Authorization creation information: /var/tmp/finance_1.pea
```

Config#

8.1.3 Synchronising the pool and profile data to the remote Centera

Whilst a pool and profile both named Finance now exist on the source Centera, the remote Centera has neither. The pool and profile identifiers on both Centera (source and target) must be identical, the only way to ensure that the identifiers are identical on both systems is to export the data from the source Centera and transfer this to the target Centera.

Simply creating the new profiles and pools on the target Centera with the same name as on the source Centera will not work, as replication will not be possible; replication can only take place between pools with the same identifiers.

Hence, the next stage of the process is to transfer the pool and profile information to the target Centera. This takes place in two stages:

- First, the data from the source Centera is exported to a local file
- Secondly the exported information is imported onto the target Centera and a new PEA generated.

The example below continues with the Finance example used previously; the export command is issued on the source Centera.

```
Config# export poolprofilessetup
Export complete setup? (yes, no) [no]: n
Export based on pool or profile selection? (pool, profile) [pool]: pool
Export all Pools? (yes, no) [no]: n
Pools to Export: Finance
Please enter the pathname where the export file should be saved: /var/tmp/finance.txt
```

Exported 1 pool and 1 profile.

Config#

Once the configuration file has been created this information can be imported onto the target Centera. The export command writes the pool or profile details from the Centera that you are connected to into a file on the system running the CenteraCLI. The import command reads the file (again on the system running the CenteraCLI) and creates the new pools/profiles on the Centera to which you are currently connected. If you always create new pools/profiles on one of the Centera (e.g. the local one) you will do the export on the local Centera and the import onto the remote Centera. Connect to the remote Centera via the CLI and issue the command shown below.

```
Config# import poolprofilessetup
Please enter the pathname for the pool import data: /var/tmp/finance.txt
Found 1 pool and 1 profile.
```

Step 1: Importing Pools

=====

Import pools? (yes, no) [yes]: y

Pool Name	ID	Mask	Profiles
Finance	b209b196-1dd1-11b2-afa8-8e107366e86f-3	rdqe-cw	1

```
Import pool Finance? (yes, no, abort) [yes]: y
Imported pool Finance.
```

Step 2: Importing Profiles

=====

Import profiles? (yes, no) [yes]: y

Profile Name	Home Pool	Type	Enabled	Monitor Cap
Finance	Finance	Access	yes	no

```
Import profile Finance? (yes, no, abort) [yes]: y
Establish a Pool Entry Authorization for application use? (yes, no) [no]: y
Please enter Pool Authorization creation information: /var/tmp/finance_2.pea
Imported profile Finance.
```

Pool Name	Granted Rights
Finance	rdqe--w

```
Import Granted Rights in pool Finance for profile Finance? (yes, no, abort) [yes]: y
Imported Granted Rights settings in pool Finance for profile Finance.
```

```
Config#
```

8.1.4 Updating the replication configuration

Once the virtual pools and profiles are created on both the source and target Centera the replication needs to be configured. Even though the Centera pair may be configured to replicate “all pools” the access profile used to perform the replication must be granted access to the newly created pool before replication can take place. As replication is configured to be bi-directional the replication profiles on both source and target Centera need updating. The replication profile requires the “clip-copy” right to the pool to enable the replication of data and the “delete” right if replicated delete is enabled as part of the replication settings.

The example below shows the target Centera having its replication profile updated (the process to update the replication profile on the source Centera is identical apart from the replication profile name and hence has not been shown). The access profile on the target Centera used by the source Centera is rep_centera1.

```
Config# set grants Finance rep_centera1
```

```
Granted Pool Rights for Profile: dc
Issue the command?
(yes, no) [no]: y
```

```
The new Effective Rights:
```

Pool Name	Profile Name	Effective
Finance	rep_centera1	-d---c--

```
Config#
```

Replication requires the clip-copy and delete privileges (The delete privilege is only required if replicated delete is enabled on the Centera)

Once the access profile has been updated data will be replicated between the two Centera for this virtual pool.

8.1.5 Creating the PEA file

As discussed previously any application wishing to access the Centera must have access credentials for both the source and target Centera. During the creation of the profiles on the source and target Centera PEA files were created, the final stage is to merge the files to produce a PEA file that may be passed to the application that is to use the profile.

The example below takes the PEA files created in the previous examples and creates the final combined PEA file required for application use.

Source Centera PEA file (finance_1.pea)

```
<pea version="1.0.0">
<defaultkey name="Finance">
<credential id="csp1.secret"
enc="base64">0MgKC+mB6srAcr5EwAFvX3IRhyqjevZBR2ARDVKHPAryE5bU2coVo5gpQzj2o3In9meNAfWerNg
QRPYXucTYusGpjN5+0kqsB30iIaIeTyGrfaeKk1ZL/htXwXdl1gnhbAw4g/hodCqCyutbqA550NiXqKz0ZixygEr
3A01uyL+GUie3M+mw1rZ+wH2VQ/xTHwinVS16HBBb176xyyo758zuFsFiuiEL+YYJG450U3NSTNd8wZM7hebtioG
yDyBFwr9XacTnn5/980LZm2EE7wrKeIDABJhdH8zw2uLEL7wscgS1pCabHPu+2XMHm+4Vn3wHbDaCLXBhArgCGFL
```

```
8MeNPgptnttAS1F5zXovwLUaxY5RcrRgH/bHqoX1USYhjCJM0PTSMuYK9dxzGI7A2J+LvihGpHxjiw/OLKA1XSW
7Yu4+Kf1bg+diDxNF1x8PYWeZEEB7ZbPBYSAAwqrLTI5pphxFsw9mdjRZbRouxSUe0TgylREzLHgX7GsQZyUtPbc
eWUq1C6pS/lb3hhXgCtnjPKZMgXsvPYmi1fk0quDX2J0Z18DA5a9bqjzUtahG5PMJIbs5ANrZ5on1EzAhXEU0B2f
3J0cK1u0Azvdq3pr9K9U7w0mkgLcwFepxLFRSNNZ3CVD+bcGTmZe0vH21kFxfjw0fpCQqmpkK8xCf7uW4=</crede
ntial>
</defaultkey>
<key type="cluster" id="b209b196-1dd1-11b2-afa8-8e107366e86f" name="Finance">
<credential id="csp1.secret"
enc="base64">0MgKC+mB6srAcr5EwAFvX3IRhyqjevZBR2ARDVKHPAryE5bU2coVo5gpQzj2o3In9meNAfWerNg
QRPYXUCtYusGpjN5+0kqsB30iIaleTyGrfaeKk1ZL/htXwXdl1gnhbAw4g/hodCqCyutbqA550NiXqKz0ZixygEr
3A01uyL+GUIe3M+mw1rZ+wH2VQ/xTHwinVS16HBBb176xyyo758zuFsFuiEl+YYJG450U3NSTNd8wZM7hebtioG
yDyBFwr9XacTnn5/980LZm2EE7wrKeIDABJhdH8zw2uLEL7wscgS1pCabHPu+2XMHm+4Vn3wHbDaCLXBhArgCGFL
8MeNPgptnttAS1F5zXovwLUaxY5RcrRgH/bHqoX1USYhjCJM0PTSMuYK9dxzGI7A2J+LvihGpHxjiw/OLKA1XSW
7Yu4+Kf1bg+diDxNF1x8PYWeZEEB7ZbPBYSAAwqrLTI5pphxFsw9mdjRZbRouxSUe0TgylREzLHgX7GsQZyUtPbc
eWUq1C6pS/lb3hhXgCtnjPKZMgXsvPYmi1fk0quDX2J0Z18DA5a9bqjzUtahG5PMJIbs5ANrZ5on1EzAhXEU0B2f
3J0cK1u0Azvdq3pr9K9U7w0mkgLcwFepxLFRSNNZ3CVD+bcGTmZe0vH21kFxfjw0fpCQqmpkK8xCf7uW4=</crede
ntial>
</key>
</pea>
```

The contents of this file should be pasted into a new file, in this case finance.pea.

Target Centera PEA file (finance_2.pea)

```
<pea version="1.0.0">
<defaultkey name="Finance">
<credential id="csp1.secret"
enc="base64">0MgKC+mB6srAcr5EwAFvX3IRhyqjevZBR2ARDVKHPAryE5bU2coVo5gpQzj2o3In9meNAfWerNg
QRPYXUCtYusGpjN5+0kqsB30iIaleTyGrfaeKk1ZL/htXwXdl1gnhbAw4g/hodCqCyutbqA550NiXqKz0ZixygEr
3A01uyL+GUIe3M+mw1rZ+wH2VQ/xTHwinVS16HBBb176xyyo758zuFsFuiEl+YYJG450U3NSTNd8wZM7hebtioG
yDyBFwr9XacTnn5/980LZm2EE7wrKeIDABJhdH8zw2uLEL7wscgS1pCabHPu+2XMHm+4Vn3wHbDaCLXBhArgCGFL
8MeNPgptnttAS1F5zXovwLUaxY5RcrRgH/bHqoX1USYhjCJM0PTSMuYK9dxzGI7A2J+LvihGpHxjiw/OLKA1XSW
7Yu4+Kf1bg+diDxNF1x8PYWeZEEB7ZbPBYSAAwqrLTI5pphxFsw9mdjRZbRouxSUe0TgylREzLHgX7GsQZyUtPbc
eWUq1C6pS/lb3hhXgCtnjPKZMgXsvPYmi1fk0quDX2J0Z18DA5a9bqjzUtahG5PMJIbs5ANrZ5on1EzAhXEU0B2f
3J0cK1u0Azvdq3pr9K9U7w0mkgLcwFepxLFRSNNZ3CVD+bcGTmZe0vH21kFxfjw0fpCQqmpkK8xCf7uW4=</crede
ntial>
</defaultkey>
<key type="cluster" id="b7f0e2f0-1dd1-11b2-955b-88d18c53b922" name="Finance">
<credential id="csp1.secret"
enc="base64">0MgKC+mB6srAcr5EwAFvX3IRhyqjevZBR2ARDVKHPAryE5bU2coVo5gpQzj2o3In9meNAfWerNg
QRPYXUCtYusGpjN5+0kqsB30iIaleTyGrfaeKk1ZL/htXwXdl1gnhbAw4g/hodCqCyutbqA550NiXqKz0ZixygEr
3A01uyL+GUIe3M+mw1rZ+wH2VQ/xTHwinVS16HBBb176xyyo758zuFsFuiEl+YYJG450U3NSTNd8wZM7hebtioG
yDyBFwr9XacTnn5/980LZm2EE7wrKeIDABJhdH8zw2uLEL7wscgS1pCabHPu+2XMHm+4Vn3wHbDaCLXBhArgCGFL
8MeNPgptnttAS1F5zXovwLUaxY5RcrRgH/bHqoX1USYhjCJM0PTSMuYK9dxzGI7A2J+LvihGpHxjiw/OLKA1XSW
7Yu4+Kf1bg+diDxNF1x8PYWeZEEB7ZbPBYSAAwqrLTI5pphxFsw9mdjRZbRouxSUe0TgylREzLHgX7GsQZyUtPbc
eWUq1C6pS/lb3hhXgCtnjPKZMgXsvPYmi1fk0quDX2J0Z18DA5a9bqjzUtahG5PMJIbs5ANrZ5on1EzAhXEU0B2f
3J0cK1u0Azvdq3pr9K9U7w0mkgLcwFepxLFRSNNZ3CVD+bcGTmZe0vH21kFxfjw0fpCQqmpkK8xCf7uW4=</crede
ntial>
</key>
</pea>
```

The area of the file highlighted above in the finance_2.pea file should be pasted into the file finance.pea as shown below:

```
<pea version="1.0.0">
<defaultkey name="Finance">
<credential id="csp1.secret"
enc="base64">0MgKC+mB6srAcr5EwAFvX3IRhyqjevZBR2ARDVKHPAryE5bU2coVo5gpQzj2o3In9meNAfWerNg
QRPYXUCtYusGpjN5+0kqsB30iIaleTyGrfaeKk1ZL/htXwXdl1gnhbAw4g/hodCqCyutbqA550NiXqKz0ZixygEr
3A01uyL+GUIe3M+mw1rZ+wH2VQ/xTHwinVS16HBBb176xyyo758zuFsFuiEl+YYJG450U3NSTNd8wZM7hebtioG
yDyBFwr9XacTnn5/980LZm2EE7wrKeIDABJhdH8zw2uLEL7wscgS1pCabHPu+2XMHm+4Vn3wHbDaCLXBhArgCGFL
8MeNPgptnttAS1F5zXovwLUaxY5RcrRgH/bHqoX1USYhjCJM0PTSMuYK9dxzGI7A2J+LvihGpHxjiw/OLKA1XSW
7Yu4+Kf1bg+diDxNF1x8PYWeZEEB7ZbPBYSAAwqrLTI5pphxFsw9mdjRZbRouxSUe0TgylREzLHgX7GsQZyUtPbc
eWUq1C6pS/lb3hhXgCtnjPKZMgXsvPYmi1fk0quDX2J0Z18DA5a9bqjzUtahG5PMJIbs5ANrZ5on1EzAhXEU0B2f
3J0cK1u0Azvdq3pr9K9U7w0mkgLcwFepxLFRSNNZ3CVD+bcGTmZe0vH21kFxfjw0fpCQqmpkK8xCf7uW4=</crede
ntial>
</defaultkey>
<key type="cluster" id="b209b196-1dd1-11b2-afa8-8e107366e86f" name="Finance">
<credential id="csp1.secret"
enc="base64">0MgKC+mB6srAcr5EwAFvX3IRhyqjevZBR2ARDVKHPAryE5bU2coVo5gpQzj2o3In9meNAfWerNg
QRPYXUCtYusGpjN5+0kqsB30iIaleTyGrfaeKk1ZL/htXwXdl1gnhbAw4g/hodCqCyutbqA550NiXqKz0ZixygEr
3A01uyL+GUIe3M+mw1rZ+wH2VQ/xTHwinVS16HBBb176xyyo758zuFsFuiEl+YYJG450U3NSTNd8wZM7hebtioG
yDyBFwr9XacTnn5/980LZm2EE7wrKeIDABJhdH8zw2uLEL7wscgS1pCabHPu+2XMHm+4Vn3wHbDaCLXBhArgCGFL
8MeNPgptnttAS1F5zXovwLUaxY5RcrRgH/bHqoX1USYhjCJM0PTSMuYK9dxzGI7A2J+LvihGpHxjiw/OLKA1XSW
7Yu4+Kf1bg+diDxNF1x8PYWeZEEB7ZbPBYSAAwqrLTI5pphxFsw9mdjRZbRouxSUe0TgylREzLHgX7GsQZyUtPbc
```

```
eWUq1C6pS/lb3hhXgCtNjPKZMgXsvPYmi1fk0quDX2J0Zl8DA5a9bqjzUtahG5PMJIbs5ANrZ5on1EzAhXEU0B2f
3J0cK1u0Azvdq3pr9K9U7w0mkgLcwfEpxLFRSNNZ3CVD+bcGTmZe0vH21kFxjw0fpCQQmpkK8xCf7uW4=</crede
ntial>
</key>
<key type="cluster" id="b7f0e2f0-1dd1-11b2-955b-88d18c53b922" name="Finance">
<credential id="csp1.secret"
enc="base64">0MgKC+mB6srAcr5EwAFvX3IRhyqjevZBR2ARDVKHPAryE5bU2coVo5gpQzj2o3In9meNAfWerNg
QRPYUcTYusGpjN5+0kqsB30iIaleTyGrfaeKk1ZL/htXwXdl1gnhbAw4g/hodCqCyutbqA550NiXqKz0ZixygEr
3A01uyL+GUIe3M+mw1rZ+wH2VQ/xTHwinVS16HBBb176xyyo758zuFsFiuiEl+YYJG450U3NSTNd8wZM7hebtioG
yDyBFwr9XacTnn5/980LZm2EE7wrKeIDABJhdH8zw2uLEL7wcsgS1pCabHPu+2XMHm+4Vn3wHbDaCLXBhArgCGFL
8MeNPgptnttAS1F5zXovWlRUaxY5RcrRgH/bHqoX1USYhjCJM0PTSMuYK9dxzGI7A2J+LvihGpHxjiw/OLKA1XSW
7Yu4+Kf1bg+diDXNF1x8PYWeZEEB7ZbPBYSaAwqrLTi5pphxFSw9mdjRZbRouxsuE0TGylREzLHGx7GsQZyUtPbc
eWUq1C6pS/lb3hhXgCtNjPKZMgXsvPYmi1fk0quDX2J0Zl8DA5a9bqjzUtahG5PMJIbs5ANrZ5on1EzAhXEU0B2f
3J0cK1u0Azvdq3pr9K9U7w0mkgLcwfEpxLFRSNNZ3CVD+bcGTmZe0vH21kFxjw0fpCQQmpkK8xCf7uW4=</crede
ntial>
</key>
</pea>
```

This new file, in this case finance.pea, is now a valid PEA file that can be given to the application team and used as the credentials to access either Centera using the new access profile.

8.2 Shutting down a Centera

In CentraStar 3.0 (and previous versions of Centera) it is not possible for an administrator to shutdown a Centera cluster in a controlled manner. No “shutdown” or equivalent command exists in the CentraStar CLI to initiate a shutdown.

An undocumented shutdown command does exist within the CenteraViewer interface in 3.0. The use of the CenteraViewer is not supported, as its only effective role is to provide a window for the CentraStar CLI, in addition, it disables the ability to log all CLI commands. It is recommended that replication on both the target and source Centera be paused before any shutdown takes place.

CTI has requested the availability of a shutdown command within the CLI, which will be added in the next version of CentraStar.

Until the administrator is given the facility to shutdown a Centera, only an EMC Service Engineer can perform this task.

It is important that a Centera is shutdown in a controlled manner. The CMOS clocks on Centera nodes are only updated when a node is shutdown in a controlled manner. Currently, there is no external time synchronisation daemon for CentraStar, hence, it is important that the internal time of the Centera is protected at all times.

8.3 Centera replication management

There are several commands that allow replication to be managed and monitored. Below are some basic statements regarding replication:

- If there is insufficient space on the remote Centera replication will be automatically paused and an alert generated. Space must be made available on the remote Centera by either installing more capacity or removing objects that are no longer required. (As the capacity of the Centera is monitored and alerts sent when the Centera is nearing capacity)
- Replication can only take place between pools having identical pool id's
- The profile used for replication must have sufficient permissions i.e. clip-copy and delete (if replicated delete is configured), to update the remote pools.

8.3.1 Replication topology




Figure 5 CTI certified replication topology

Figure 5 shows a Bi-directional replication. Data written to cluster A will automatically be replicated to cluster B and data written to cluster B will automatically be replicated to cluster A. In case of a disaster, cluster A will failover to cluster B for read operations and cluster B will failover to cluster A. There is no need for a restore after a disaster. It is also possible for either application server to access cluster A or B in the event that one of the Centera becomes unavailable.

8.3.2 Replication status

The current Replication status can be checked by issuing the “show replication” command as shown below. This command shows the current status of replication, the number of clips still to be replicated, the replication speed and other information. Replication issues can also generate Email, SNMP and Email to EMC alerts depending on the situation..

```
Config# show replication detail
Centra Replication Detail Report
-----
Generated on Friday, October 14, 2005 11:13:02 AM BST

Replication Enabled:
Replication Paused:
Replication Address:

Replicate Delete:
Replicate incoming replicated Objects: no
Replicated Pools:
Profile Name:

Number of C-Clips to be replicated:
Number of Blobs to be replicated:
Number of MB to be replicated:
Number of Parked Entries:

Replication Speed:
Replication Speed:
Ingest:
Replication ETA:
Replication Lag:

Config#
```

Tuesday, August 30, 2005 8:10:36 AM BST
 146.128.135.121:3218,
 146.128.135.122:3218,
 146.128.135.123:3218,
 146.128.135.124:3218
 yes
 all pools
 rep_centra1
 0
 0
 0
 0
 0.00 C-Clip/s
 0.00 Kb/s
 0.00 C-Clip/s
 <1 hour
 <1 hour

Replication Link Status
 IP addresses of remote Centra
 Replication of deletes is enabled
 Profile used on remote Centra
 Replication Queue
 Entries that failed to
 Outgoing replication rate
 Rate of files arriving on replication

8.3.3 Pausing replication

The replication can be temporarily suspended by pausing all replication, this stops data being transmitted to the remote site but clips requiring replication are still tracked ready for their transmission when the link is running once more. Clips waiting to be transmitted are shown in the “show replication detail” in the “Number of C-Clips to be replicated” field. The command to pause replication is shown below.

```
Config# replication pause
Issue the command?
(yes, no) [no]: y
```

```
Config#
```

Once the command has been issued the status of replication in the “show replication detail” will change from “Replication Paused no” to “Replication Paused yes (by user)”.

Replication can also be automatically paused for other reasons such as:

- Insufficient capacity on the target Centera
- If the profile used for replication has insufficient privileges (i.e. it can write to or delete from one of the pools)
- If the credentials used for authenticate to the remote profile are incorrect
- If more that 10,000 clips have been unable to be transmitted and now reside in the “parking” queue.

It should be remembered that in bi-directional replication configurations replication would need to be paused on both Centera to stop all traffic between the two Centera.

8.3.4 Re-starting replication after a pause

If the replication has been automatically paused the any underlying issue(s) should be cleared before replication is resumed.

Replication can be re-started by the command below:

```
Config# replication resume
Issue the command?
(yes, no) [no]: y
```

```
Config#
```

As soon as replication has been un-paused any clips waiting for replication will start to move across the link.

8.3.5 Replication parking queue

Entries only exist in the “parking” queue that have failed to transfer to the remote Centera. Details of the number of clips in the parking queue can be found as shown below:

```
Config# show replication parking
Centera Replication Parking Report
-----
Generated on 20 October 2005 14:23:17 BST

Replication Enabled:          30 August 2005 08:10:36 BST
Replication Paused:           no
Replication Address:          146.128.135.121:3218,
```

```

146.128.135.122:3218,
146.128.135.123:3218,
146.128.135.124:3218
Replicate Delete: yes
Replicate incoming replicated Objects: no
Replicated Pools: all pools
Profile Name: rep_centera1

Number of Parked Entries: 0
Failed Writes: 0
Failed Deletes: 0
Failed Privileged Deletes: 0
Failed Unknown Action: 0

```

Config#

Examples of reasons why objects will be put into the parking queue are shown below:

- The link is down between the source and target Centera
- Missing pool at the remote end
- Date mismatch between Centera clusters stopping a replicated delete (clip still within its retention period on the target Centera)

Items held in the parking queue are re-tried once a day for a maximum of 7 days, after this time only a manual retries can be attempted.

Once the issue (or issues) that caused the clip(s) to be put into the parking queue the transfer may be re-tried by using the command below:

```

Config# replication parking process
Issue the command?
(yes, no) [no]: y

```

Config#

8.4 Copying data between Centera (“restore” command) – DO NOT USE

The “restore” command should not be used due to a system bug that overprotects data on the target Centera (EMC reference ETA emc117139 – “Centera: Using restore to migrate data from one Centera to another causes massive over-protection and potentially a service outage”). Any migrations or movement of data must be carried out by EMC customer services.

Until such time that EMC provides a fix for the “restore” command bug, this section is for information only.

Whilst the “restore” command is not truly replication it can be used to seed a Centera prior to replication being enabled or to copy data to a different Centera.

The “restore” command allows either all clips (“full”) or only clips between a pair of dates (“partial”) to be migrated. It is also possible to select either all pools or an individual the pool to copy. Data is copied from the source Centera to the same pool on the target Centera.

8.5 EMC access management

As discussed previously EMC accounts are only active on nodes shown as “unlocked” by the “show security all” command (see 6.1 “CTI specific configuration”).

Whilst unlocked nodes can be locked via the “set security lock” command, conversely a node (storage nodes only on CE+) may be unlocked to allow EMC access via the “set security unlock” command as shown below:

Config# set security unlock c001n14 ← Centera node to be unlocked
 Warning!: Cluster security can be compromised if you unlock a node
 continue? (yes, no) [no]: y
 Issue the command?
 (yes, no) [no]: y
 Config#

All nodes should be left in the “locked” state, unless EMC service personnel are working on the Centera, thereafter, all nodes should be re-set to the “locked” state as shown below:

Config# set security lock
 Issue the command?
 (yes, no) [no]: y
 Config#

8.6 Virtual Pool capacity reporting

As stated previously one of the major benefits of CentraStar 3.0 is the ability to segregate each application's data in a different virtual pool. This segregation ensures that each application's data can only be accessed by the authorised application also as the capacity in each virtual pool is reported this information can be used for charge-back purposes.

In CentraStar 3.0 capacities are displayed in “marketing” rather than “engineering” GB/KB i.e. multiples of 1000 rather than 1024. This will be corrected by EMC in a later release of CentraStar.

The command “show pool detail” displays capacity and utilisation information. For additional detail, the command “show pool capacity” displays all relevant usage information for all Virtual Pools:

```
Config# show pool capacity
Capacity / Pool      Quota      Used      Free      C-Clips  Files
-----
Legal                500 GB      120 GB    380 GB     5234     5129
Finance              300 GB      240 GB     60 GB    10193     9894
PoolA                 6 GB         0 GB         6 GB         0         0
default              --          92 GB         --     2672     2586
Total                --          452 GB         --    19099    17609
-----
```

Config#

The amount shown as “free” is based on the amount of storage available until the soft quota is breached.

Pool capacity figures are formed of two components: A series of monthly aggregates giving figures for all clips written in a particular month and a set of rolling figures for the current month. Hence a pools capacity figure is made up of the figures for each monthly aggregate that the pool has been in existence plus the rolling figures for the current month. These monthly aggregates are not directly available in reporting.

These capacity figures are updated in one of three ways:

- Real-time Capacity Calculation – Whenever a clip is added or deleted to/from a virtual pool the pool's rolling figures are updated to reflect this information immediately.

- Capacity Initialisation or Recalculation – When a Centera is installed or upgraded to CentraStar 3.0 a process runs to initialise the pool statistics. After the initial run each months aggregate will be automatically created and the pools rolling figures reset ready for the next month.
- Migration recalculation – When a migration occurs the option is given to recalculate the capacity information, when this option is selected a task to recalculate capacity is queued in the same manner as the monthly aggregate record creation. **The recalculation will only take place on the first day of the calendar month, until this time the capacity values will be incorrect.**

Scheduled capacity tasks can be found by using the command “show pool capacitytasks”

```
Config# Show pool capacitytasks
From      To      Status      ETA
-----
01-Oct-2005  31-Oct-2005  scheduled
```

Config#

The example above shows a recalculation for all data written for the month of October scheduled to run. Whilst there is no way for a user to recalculate arbitrary time segments, EMC customer services can execute this operation.

For charge-back purposes it is therefore suggested to use figures gathered from the Centera shortly after the monthly aggregate calculation has completed to ensure as accurate usage figures as possible.

For capacity planning purposes it is also possible to gather information on the total storage used on the Centera cluster using the “show capacity total” command. An example of this command is shown below:

```
Config# Show capacity total
Number of nodes:          16
Number of nodes with storage role: 12

Total Raw Capacity:      20,731 GB (100%)
Used Raw Capacity:       5,271 GB (25%)
Free Raw Capacity:      15,461 GB (75%)

Used Raw Capacity:       5,271 GB (25%)
System Resources:       114 GB (1%)
Offline Capacity:       0 GB (0%)
Spare Capacity:        5,154 GB (25%)
Used Capacity:          3 GB (0%)
  Audit & Metadata:      3 GB (0%)
  Protected User Data:  0 GB (0%)

Used Object Count:      22
```

Config#

The capacity remaining on the Centera cluster (the “Available Capacity” value) can also be gathered by using the “show capacity availability” command as shown below:

```
Config# show capacity availability
Number of nodes:          16
Number of nodes with storage role: 12

Total Raw Capacity:      20,731 GB (100%)
Used Raw Capacity:       5,271 GB (25%)
Free Raw Capacity:      15,461 GB (75%)

System Buffer:           1,341 GB (7%)
Regeneration Buffer:     648 GB (3%)
Available Capacity:     13,471 GB (65%)
```

Total Object Count:	360 M	(100%)
Used Object Count:	22	(0%)
Free Object Count:	360 M	(100%)

Config#

Ultimately a Centera usage will be limited by either total capacity or maximum object count. That is to say, a Centera will be unable to store any more objects if the threshold of maximum objects (object bound), or maximum capacity (capacity bound) has been met.

9 CTI Centera Configuration Summary

9.1 Certified Components Summary

- Only the CE+ version of Centera is being certified for use within CTI (2 “Centera Overview”)
- Only Content Protection Mirrored (CPM) is certified for use within CTI (2.2 “Centera Architecture Overview”)
- The use of the Centera Universal Access (CUA); previously known as the Centera Access Gateway (CAG), is not certified for use. (3.2 “Centera Universal Access”)
- The management network, Solaris management server and PowerBroker server are to provide compensating controls for Centera’s lack of roles based management accounts and auditing and these must be present in all configurations. (4.1.4 “Management Server/Infrastructure”)
- The minimum base configuration for new installs is an 8-node Centera with 4 access/storage nodes and 4 storage nodes. A 4-nodes Centera configuration is approved for situations where limited capacity is required and there is a limited performance requirement as detailed further in this document. If there are any doubts as to the capacity of performance requirements the 8-node configuration applies. (3 “CTI Centera Configuration Standards”)
- The current maximum number of Certified nodes within a Centera cabinet is 32 (Centera clusters with Gen 2/3 and Gen 4 nodes are limited to a maximum of 28 nodes). (3.1 “Centera hardware expansion”)
- A Centera cluster may be formed of a maximum of 4 cabinets (i.e. 128 Nodes) (3.3 Multi-Cabinet Centera Clusters)

9.2 Configuration rule Summary

- Only data of Confidential or lower security classification may be stored on a Centera without additional application security measures, appropriate to the security classification of the data being stored, being applied. (1.3 Security Classification)
- Centera Retention classes should not be used by any application. (10.2 “Centera Retention Classes”)
- A minimum of two pure storage nodes (i.e. no access node function) should exist in a Centera cluster to provide management network connections. (3 “CTI Centera Configuration Standards”)
- All Centera implementations should be configured with bi-directional replication between two Centera in different physical locations. (4.1.1 “Replication”)
- Replication must be configured to use specific access profiles rather than the “anonymous” profile. (2.2 “Centera Architecture Overview”)

- Each Centera cabinet in a multi-cabinet configuration should have a minimum of 4 access nodes configured. (3.3 Multi-Cabinet Centera Clusters)
- All access nodes in a Centera cluster should be configured on the same subnet. (4 Preparation for installation)
- A minimum of 4 access nodes on each Centera cluster should be configured as replication destinations to provide resilience and ensure sufficient replication throughput between Centera clusters. (Limited performance configurations must have two access nodes configured). (4.1.1 “Replication”)
- Network connections should not use “Auto” for Speed or Duplex settings but should be hard coded to the appropriate setting. The cable distance between the Centera and the network switch/router should be no longer than 30m. (5 “Installation of Centera”)
- Maximum/Minimum retention times on the Centera must be set. The minimum value should be set to the lowest valid retention time of all the applications using the Centera, this must be greater than zero to ensure that no application uses a zero retention time and hence avoid having a real retention time stored as part of the meta data (an object with zero retention time may be deleted immediately). The maximum value is set to the largest valid retention time for any application using the Centera. (6.1 “CTI specific configuration”)
- Replicated delete, this is a global setting and should be enabled unless an application using the Centera cannot function with replicated delete enabled. Applications that require different settings for replicated delete cannot use the same Centera. (4.1.1 “Replication”)
- All Centera nodes should be left in the “locked” state to ensure that EMC cannot access the Centera without CTI’s intervention. (6.1 “CTI specific configuration”)
- All applications using a newly installed Centera (i.e. not an upgrade to an existing Centera) with CentraStar 3.0 or higher must use a minimum version of 2.3.315 for the Centera SDK. (7.1 “SDK Compatibility with CentraStar 3.0”)
- Centera upgraded to CentraStar 3.0 should migrate all application data to application specific virtual pools as soon as possible (this will not be possible for data written with an SDK prior to 2.3.315). Applications that cannot use virtual pools (SDK less than v2.3.315) should upgrade to SDK v2.3.315 or greater, as soon as possible. (7.1 “SDK Compatibility with CentraStar 3.0”)
- The “anonymous” profile must be disabled on all CTI Centera. (6 “CTI Configuration”)
- Each application must have a unique access profile for authentication with a nominated “home” virtual pool. (2.2 “Centera Architecture Overview”)
- No access profile whether used for application or replication purposes should be granted more privileges than required for normal functioning of the application/replication i.e. do not grant all privileges to profiles. (8.1.1 “Creating a new virtual pool”)
- Each application must have unique “home” virtual pool, unless the application design requires that it shares its virtual pool with other applications. For example, applications using SDK prior to v2.3.315 cannot use a specific virtual pool, hence, must share the “default” pool) (7.1 “SDK Compatibility with CentraStar 3.0”)

9.3 Important Information Summary

- In order to use combined access/storage nodes on a Centera upgraded from CentraStar 2.2 an additional software license is required (The EMC part number of this license is CNR8NMANSW). (3 “CTI Centera Configuration Standards”)
- Simply creating the new profiles and pools on the target Centera with the same name as on the source Centera will not work, as replication will not be possible; replication can only take place between pools with the same identifiers. Target pools must be created using the import/export commands, this ensures that the pool identifiers are the same on both source and target Centera. (8.3 “Centera replication management”)
- Centera will only replicate data that is written after replication is enabled i.e. data written before replication is enabled will not be automatically replicated. To copy data that was written before replication was enabled requires data to be copied manually by other mechanisms within the Centera (e.g. the Centera “restore” command). (2.2 “Centera Architecture Overview”)
- It is important that a Centera is shutdown in a controlled manner, (only an EMC Centera Service Engineer is allowed to do so). The CMOS clocks on nodes within the Centera that provided the time whilst the Centera is powered off are only updated when a node is shutdown in a controlled manner. Currently there is no external time synchronisation for a Centera it is important that the internal time of the Centera is protected as much as is possible (10.1 “Centera Time Synchronisation”)
- The “restore” command should not be used due to a code bug that overprotects data on the target Centera (EMC reference ETA emc117139 – “Centera: Using restore to migrate data from one Centera to another causes massive over-protection and potentially a service outage”). Any migrations or movement of data must be carried out by EMC customer services. (8.4 “Copying data between Centera (“restore” command) – DO NOT USE”)
- Pool capacity re-calculation (e.g. after a migration) starts on the first day of the next calendar month, until this re-calculation completes the capacity values may be incorrect. (8.6 “Virtual Pool capacity reporting”)
- Centera Gen2/3 nodes are end-of-availability. As such, additional capacity to Gen2/3 will use Gen4 nodes. Gen 4 nodes have dual power outlets, which require installation of a power adapter into an existing Gen2/3 cabinet. The power adapter reduces the number of nodes that can be installed in a Gen 2/3 Centera cabinet to 28 nodes. (3.1 “Centera hardware expansion”)
- Automatically creating a PEA file containing the credentials for both Centera is not possible with the current Centera software. To create a usable PEA file credential files need to be created from both the local and remote Centera and the key section from the remote PEA manually added to the local PEA file to form a PEA file containing both sets of credentials. (8.1.5 “Creating the PEA file”)
- A Centera never deletes information when its retention period expires, the application must perform a delete to erase any data after the retention period expires. (2.2 “Centera Architecture Overview”)

10 General Information

10.1 Centera Time Synchronisation

At this time Centera does not support external time synchronisation. EMC plans to introduce this feature in a later version of CentraStar, initially for Basic and Governance editions and later for CE+.

CentraStar includes an internal NTP-like service that manages time synchronization of all nodes within the cluster. All the Centera nodes will therefore drift at the same rate. Centera time is calculated upon boot-up using the average CMOS clock time of a selected number of access nodes. Therefore, proper shutdown procedures must be adhered to in order to ensure the system time is accurately determined.

CMOS clocks on nodes are only updated by the Centera cluster time during a clean node/cluster shutdown/reboot.

CMOS clocks will be updated every hour in version 3.1 of CentraStar.

10.2 Centera Retention Classes

Centera has the facility to define a named retention class and give this class a retention period. Administrators can choose to assign an application to use a specific retention class rather than setting a specific retention period. Any data stored will make reference the retention class in the meta data where a specific retention time is normally held, the object will be retained until the current time value set for the retention class is exceeded.

Retention classes are not certified for use within CTI and must not be used by any application.

The reason for this is that:

- When using max/min retention policy, the use of retention classes cause API writes to fail (In CentraStar 3.1 this is no longer the case)
- The retention class name is written to the c-clip not the retention period itself, hence changing the retention period associated with a retention class (a CE+ Centera can only have this retention period increased) modifies the retention time for all objects ever written using this retention class. Therefore the retention period for an object cannot be guaranteed.

10.3 Error/Heath Reporting

A Centera can report errors and its status by four mechanisms:

- Email to EMC (encrypted using AES-256 encryption algorithm)
- Email to local CTI administrators
- SNMP trap
- API monitoring

EMC Control Centre monitoring is not currently supported by CTI but is planned for a later release.

All Email and SNMP transfers are made via the principle access node. All alerts are sent via both the Email and SNMP routes if enabled.

CentraStar 3.0 has over 100 sensors actively monitoring the health, capacity and performance of a Centera cluster. Of those sensors, 20 have one or more evaluators defined that can trigger 44 distinct alert events. The table below shows the list of the alerting sensors in CentraStar v3.0.

Each sensor has the following attributes:

- The Sensor Name indicates which component of the Centera system is being monitored;
- The Frequency indicates the frequency with which the sensor's Inspector checks the Centera component, in minutes (m) or hours (h); and
- The Type of sensor: this classification exists for historical reasons – Type B and C sensors are used to monitor a set of hardware related events which have not been fully migrated to the new sensor framework, whereas Type A sensors actually use the full sensor framework as described above. In future releases, all Type B and C sensors will be reformatted to Type A.

#	Sensor Name	Freq	Type
1	Application.Main.ReplicationManager.Main.ParkedReplicationsClipCount	12h	A
2	Application.Main.ReplicationManager.Main.ParkedRestoreClipCount	12h	A
3	Application.Main.ReplicationManager.Main.ReplicationETA	24h	A
4	Application.Main.ReplicationManager.Main.ReplicationSubstate	10m	A
5	Application.Main.SECComplianceAdapter.Main.ComplianceModeRevert	1m	A
6	CapacityReportingManager.pool.<poolid>.capacity	15m	A
7	Foundation.Cluster.PowerManager.Main.PowerRedundancyStatus	Real-time	A
8	Foundation.Cluster.RegenerationManager.Main.ClusterClipIntegrity	1h	A
9	Platform.Main.FPHealthCheck.Main.CPUTemperature	1h	A
10	Platform.Main.FPHealthCheck.Main.VarPartitionUsedSpace	1h	A
11	Server.Product.SuicidePreventionManager.Main.AvailableCapacityHardStopPercent	1h	A
12	Server.Product.SuicidePreventionManager.Main.FreeObjectPercent	1h	A
13	Server.Product.SuicidePreventionManager.Main.RegenerationBufferAlertPercent	1h	A
14	Server.Product.SuicidePreventionManager.Main.RegenerationBufferHardStopPercent	1h	A
15	Hardware.Main.NIC.Main.ExternalNICFailureCount	1m	B
16	Hardware.Main.NIC.Main.InternalNICFailureCount	1m	B
17	Hardware.Main.Switch.Main.CubeSwitchFailureCount	1m	B
18	Hardware.DiskOffline	Real-time	C
19	Hardware.NodeOffline	Real-time	C
20	Hardware.RootSwitchOffline	Real-time	C

Figure 6 Centera 3.0 alerting sensors

A complete list of all alerts for CentraStar 3.0 is shown in "Appendix E – CentraStar 3.0 Alerts".

All alert messages of "Warning" or higher (i.e. "Error", "Critical") except the "CPUTemperature" "Warning" alert and any "PoolReachingQuota" alerts should be escalated to local Centera administrators by the locally selected monitoring mechanism.

The CPU Temperature (Warning) and PoolReachingQuota (any) alerts may, at the discretion of the local administrators, be alerted, but these are not mandatory notifications.

In CentraStar alerts the failure of devices or components that would normally produce an "Error" or "Critical" alert in other devices are frequently notified as a "Warning" by CentraStar e.g. the

failure of a node or disk within a node. The reason for this is the Centera will recover from the loss of a disk or node by rebuilding the data using redundant components. It is important that the failing components are replaced as soon as possible to ensure that the Centera can recover from any subsequent failures.

10.3.1 Email

In addition to specific error alert Emails, a daily health report Email is also sent to both EMC and local system administrators. An example of a health report can be seen in “Appendix C – Example Centera Health Report”. This report shows the status of the Centera cluster and its individual components.

The whilst the Email sent to EMC includes CTI confidential information (such as the access nodes IP addresses, default router and MAC address) this Email does not breach CTI security policy as EMC already have this information (installation CCRF form) and the Email is sent encrypted using an CTI approved encryption algorithm (i.e. AES-256).

During the installation of the Centera EMC will configure the local Email gateway and local Email address (or mail list) to use for Email notification. The Email gateway can be changed via the “set notifications” command and the mail address/ list for local administrators can be changed via the “set owner” CLI command.

10.3.2 SNMP

Prior to CentraStar 2.3 SNMP reporting was only available on Basic and Governance edition Centera. As of CentraStar 2.3 SNMP alerts are available to all editions of Centera.

Centera only supports the use of SNMP Traps (two types heartbeat and problem notification). There is no ability to get information or set configuration via SNMP directly e.g. Get, Getnext, Set.

The heartbeat trap can send three types of messages to the management station, each describing a different level in problem severity:

- Informational: Centera cluster is healthy
- Warning: A non critical error has occurred
- Critical: The Centera cluster is in a critical state

The problem notification trap is sent when a problem occurs on the Centera. Unlike the heartbeat trap, this trap is send only when a problem occurs. Two error levels can be conveyed in this trap:

- Warning: A non critical error has occurred
- Critical: The Centera cluster is in a critical state

The description field of this trap contains a short, concise message stating the nature of the problem. It is the intent that the system administrator will use Centera CLI for problem correlation and resolution.

The community strings and SNMP management station are configured using the “set snmp” command.

```
Config# set snmp
Enable SNMP (yes, no) [no]: y
Management station [169.183.163.206:162]:
Community name [public]: abcd1234
```

```
Heartbeat trap interval [1 minute]:  
Issue the command?  
(yes, no) [no]: y
```

```
Config#
```

A copy of the Centera MIB can be found in “Appendix F – Centera MIB” and a list of the alerts can be found in “Appendix E – CentraStar 3.0 Alerts”.

10.3.3 API Monitoring

It is also possible to monitor the status of a Centera using specific calls in the Application Programming Interface supplied in the SDK. The SDK provides functions that allow an application to both gather performance statistics and retrieve any errors. Any application written to use the SDK will also require a valid application profile with the “monitoring” privilege. At this time CTI has neither written nor certified any tools that use this interface.

11 Centera Architecture

The next section goes into greater detail about the way in which a Centera functions and the way in which it can be configured. This section is included due to the lack of technical and best practice information currently available from EMC.

11.1 Single Instancing of Data

Each piece of data stored within Centera has a Content Address generated from the data itself; from this it follows that each time the same piece of information is stored it will generate the same Content Address. Hence, whilst all CDF files are unique multiple CDF files may point to the same data object. Any meta-data such as retention periods or application meta-data is placed within the CDF, this enables multiple references by different applications to the same information. Deletion by an application removes the CDF, rather than the data itself, the data is later removed by a garbage collection processes once all CDF references to a piece of data are deleted.

Its important to understand that multiple CDFs can be used to increase retention periods, once a CDF has been written, as its WORM-like, it cannot be changed or deleted. Therefore, to increase the retention of an object a new CDF record with a longer retention period is created referencing the same data. The application now has two CDF entries that it must erase when the data expires, one that can be deleted after the original expiry time and a second that can only be deleted when the new, longer, retention period expires. Both these CDF entries must be tracked and managed by the application.

As of Centera 2.3 service pack 3 there are 3 possible naming schemes:

M – This is the original content addressing scheme, here a 128 bit MD5 key is generated from the data and used to form the content address. During 2004 researchers presented a paper that demonstrated that artificially constructed MD5 hash collisions could be generated, however these collisions were generated using data that differed by only 6 bits in a 1024 bit value. Whilst the likelihood of collisions being deliberately generated is small in the extreme EMC switched their standard naming schemes to GM and M++ to preclude this attack.

GM – This is the “performance” addressing scheme, here additional information is appended to the naming scheme such as 70 bit random string and 35 bit timestamp. Whilst this naming scheme will statistically never generate a collision a side effect is that single instancing is effectively disabled.

M++ – This is the “capacity” addressing scheme, here the content address is generated by combining the 128 MD5 hash value and 120 bits of a SHA-256 hash to form the hash value. Whilst single instancing remains in place there is an increase in CPU time required to compute the content address (approximately 50% over the M scheme).

The GM and M++ naming schemes are combined when a Centera is used in the “performance” storage strategy. A threshold is defined at the Centera level (default 250KB), objects below are stored using the GM scheme, and objects above use the M++ scheme.

NOTE: If an object is protected by Content Parity Protection (CPP – see 11.5 Data protection), the threshold becomes 6x250KB or 1.5MB; hence only files above 1.5MB are single instanced in a CPP configuration.

11.1.1 What happens if a duplicate key is generated

The content address, which is based on the hash value produced by the M, GM or M++ naming scheme is used to identify an object, it is statistically possible that a collision (duplicate hash value) will be generated eventually, however it will take a long, long time.

Address Scheme	Number of files beyond which collisions become statistically likely	Saving 10,000 files per second it would take this long before collisions became likely.
M	2^{64}	58 Million years
GM	Not possible	N/A
M++	2^{124}	67 Million, Billion, Billion years

Note: The age of the universe is measured to be around 15 Billion years old.

In the event of duplicate content address occurring, Centera's SDK permits applications to detect the event and modify the object by adding further information to the second object so that it generates a unique address if required. Nevertheless, it is up to the application to ensure that it is indeed an address collision and not the same data being stored twice.

11.2 Centera Hardware

The Centera architecture has been termed Redundant Array of Independent Nodes (RAIN). What this means is that each Centera cluster is made up of many standard nodes. There has been 4 generations of Centera node hardware. Specifications for the current node (Gen 4) and the previous generation (Gen 3) are outlined below:

Generation 3.0 Hardware (Gen 3)

- 2GHz Pentium 4 CPU
- 512MB DDR RAM
- 4 x 320GB IDE Disk Drives (Gives approx 1.2TB raw capacity)
- 2 x 100MBPS Network Interface Cards (for internal connectivity)
- 1 x 1Gb (copper) Network Interface Card (for host connectivity)
- Single power feed

Generation 4.0 Hardware (Gen 4)

- 2.8GHz Pentium 4 CPU
- 1GB DDR RAM
- 4 x 320GB IDE Disk Drives (Gives approx 1.2TB raw capacity) – After Jan 2006 4 x 500GB systems are also shipping.
- 1 x 1Gb Network Interface Cards (for internal connectivity)
- 1 x 1Gb (copper) Network Interface Card (for host connectivity)
- Dual power feeds with integrated AC transfer switch

Only Gen 4 nodes are now being shipped and these require a minimum of CentraStar 3.0.

As computing technology emerges, newer Centera nodes will increase both: processing power, and storage capacity.

Nodes are inter-connected using two independent TCP/IP network switches, located in the Centera cabinet. Each cabinet supports a maximum of 32 nodes. Centera cabinets are connected together by Gigabit Ethernet connections and a maximum of 4 cabinets can be inter-connected to form a single Centera cluster.

Prior to the Gen 4 hardware Centera units were built from 8 node increments, yielding Centera clusters containing 8, 16, 24 or 32 nodes within a single cabinet. The Centera Gen 4 hardware introduces expansion in 4 node increments. EMC refers to a group of nodes within a single cabinet as a “Cube”; each cube is made up of 8 to 32 Gen 2/3 nodes, or 4 to 16 Gen 4 nodes. Groups of Cubes configured together as a unit are referred to as a Cluster.

Nodes can be separated into one of four classes:

Access Node: Access nodes are where applications connect. Access nodes process all I/O transactions between the Centera and applications. More access nodes give the Centera the ability to transfer more data to applications, hence, application performance should scale relative to number of access nodes.

[A pure access node is no longer available for new installs from CentraStar 2.4 onwards, in CentraStar 2.4 a combined Access/Storage node was introduced; however prior to 2.4 this was the only access node available. It is included here as any Centera originally installed using CentraStar versions prior to 2.4 may still have pure access nodes even after being upgraded to CentraStar 3.0]

Storage Node: These nodes hold the data stored and cannot be accessed by external applications except via the access nodes. More storage nodes increase the data capacity of the Centera.

Access/Storage Node: From CentraStar 2.4 onwards a node may be configured to provide both the Access and Storage node functions, this is now the default configuration for an access node. Whilst there is a slight performance degradation of the access node performance (around 5% has been quoted by EMC) the use of combined function nodes allows access to the storage contained on this node that would not be available in a normal storage node. The use of this feature requires a license key from EMC and is included with all new Centera installs since 2.4.

Spare: This is a node which has not been assigned to either storage or access node duties.




Figure 7 Centera Schematic Diagram (Gen 3 Hardware)

Prior to the Gen 3 hardware two separate power sources were installed in the Centera cabinet. Centera Gen1/2 nodes consisted of single power outlets that were distributed evenly to the two power sources. In the event of a power failure to a single power source, half of the Centera nodes lost power. Data was distributed in such a way as to ensure at least one copy of data was available even after the loss of a single power feed.

In Centera Gen 3 hardware an Automatic Transfer Switch (ATS) was added, which feeds power to all Centera nodes and in the event of a single power feed failure ensures that all nodes continue to receive power.

The new Centera Gen 4 hardware has dual power sources to the cabinet and each node also has dual power outlets, hence the ATS is not required for Gen 4 nodes.

Expansion of Centera Gen2/3 cabinets will require the installation of additional network switches within the Centera cabinet as Gen 4 nodes may only have a maximum of 16 nodes in a Cube rather than the 32 node maximum for Gen2/3 nodes. The installation of these extra network switches reduces the number of nodes that can be installed in a Gen 2/3 Centera cabinet to 28 nodes.

11.3 Data Access Security

The Centera Access Security model provides a way to authenticate applications and authorise their access to a Centera when they request a connection. This prevents unauthorised applications from storing or retrieving data from a Centera. The security model operates on the application level, not on the level of individual end users.

This Access security model uses data pools and application profiles to segregate data and to control access rights to the data within those pools.

11.3.1 Virtual Pools

Data within the Centera is stored within pools. Prior to CentraStar 3.0 a single pool was used for all data stored. With the introduction of CentraStar 3.0 there is now a “default” pool and up to 98 user-defined virtual pools (there is also the Cluster pool but this is a superset of all pools and is not normally used to access data). If a Centera is being upgraded to CentraStar 3.0 all existing data will be initially migrated to the “default” pool but may be migrated to specific pools as part of the upgrade process (see 7 Upgrading an existing Centera).

To segregate application data and provide the ability to report charge-back information, all applications should be configured to use their own virtual storage pool. However, there may be occasions when several applications may have to share the same virtual pool, for example, several instances of the same application on different servers. Without this segregation it is possible for an application to gain access (subject to retention policies) to another application’s data, provided that application either knew or could recover (e.g. clip query) an object’s Content Address. Whilst an application may not know the content of a particular object, unauthorised access presents similar security concerns as those found on traditional block and file system storage devices

11.3.2 Profiles

Access to the Centera is gained using application profiles. To use a profile an application must authenticate itself to the Centera by supplying both the profile name and a secret key defined when the profile was created (the secret key may be changed at a later date).

Each profile has a home virtual pool associated with it; all writes made using this profile will be made to this pool. There are three profile types:

- **“Anonymous” profile:** This profile exists to provide compatibility with previous versions of CentraStar when profiles did not exist. Any application failing to provide profile information will use this profile. No authentication is required and at installation this profile has the “default” pool as its home pool. This profile cannot be deleted but on CTI systems should be disabled to ensure all applications authenticate before being given access to relevant data on the Centera.
- **Access profile:** This is the main class of profile. An access profile should be created for each application accessing the Centera. Each access profile should have an application specific virtual pool associated as its home pool. Each access profile may also be granted access rights to other pools in addition to the home pool. Whilst the profile cannot write into these other pools (a profile can only write to its home profile), it can read, search and if required delete C-Clips from them. Each profile should have a “secret” key associated with it, this key is normally created in the form of a Pool Entry Authorisation (PEA) file that is created at pool creation time and contains XML entries used by the SDK on the application server to authenticate to the Centera. The Centera administrator must give this authentication information to the application owner before they can access the Centera. [Access profiles should also be created for the replication of data between Centera]
- **Cluster profile:** This is a special profile with “cluster” pool as its home pool, this profile does not allow normal writing of data and is normally used by applications such as backup/restore and as such it’s not normally used.

Currently, CTI has not certified any enterprise backup utilities for Centera tape backup.

An application using a particular profile only has access to data stored in pools that it has access to and the level of access will depend on the access rights granted to it.

11.3.3 Access Control Lists

Centera Profile access controls are similar to those of traditional file system ACL's. The atomic permissions that form the Access Control Lists (ACL's) are shown below:

- Write (w) – Write C-Clip
- Read (r) – Read C-Clip
- Delete (d) – Delete C-Clip (also required for replicated deletes, see 11.6 “Remote copies of data (Replication)”)
 - Exists (e) – Check C-Clip exists
- Privilege Delete (D) – Delete all copies of a C-Clip, overrules retention period (Only available on Governance Edition Centera's)
- Query (q) Search pool for C-Clips created between two time/dates and return the Content Address of C-Clips found
- Clip-Copy (c) – Needed for replication of C-Clips
- Purge (p) – Remove C-Clip (only cluster profiles)
- Monitor (m) – Retrieve Centera statistics via API calls using this profile
- Profile-Driven Metadata – Allows storing or retrieving per profile automatically added metadata (currently only used by Centera Seek)

ACL's using these permissions are set at creation time on the masks shown below (these ACL's can also be modified at a later date if required):

- Cluster Level – (Cluster Mask, this is equivalent to the root profile in CentraStar versions prior to 3.0)
- Pool Level – Pool Mask set at pool creation time
- Profile Level – Profile Mask set at profile creation time

When an application connects to a Centera its access permissions are generated for each pool that it has access to, this is done by looking at the ACL's stored in the masks in the following order:

Cluster -> Pool-> Profile

A privilege missing from a higher level cannot be granted at a lower level, for example, if write (w) is not present in the Cluster Mask as this is the first ACL checked profile on the Centera no profile would be able to write to any pool.

e.g.

Cluster Mask: rdqe cw
 Pool Mask: rdqe w
 Profile Mask: r qe

Resulting ACL rqe (Read, Query, Exists)

Cluster Mask: rdqe cw
 Pool Mask: rdqe
 Profile Mask: r qe w

Resulting ACL = rqe (Read, Query, Exists)

In the second example the “write” permission is not available despite being requested at the profile mask level as it was denied at the higher pool level.

11.4 Pool Entry Authorisation (PEA) file

An application can authenticate to the Centera in a number of ways but all rely on the “secret” associated with the profile to be used. The authentication information is typically created in the form of the PEA file which is produced by the Centera administrator when the application profile is created.

When an application profile is created the opportunity is given to create the PEA credentials, the command below shows a profile being created using the Centera CLI (The items in **bold** relate to the creation of the profile secret and pea file), where no input values are shown a “Return or NULL” was typed to select the default action.

```
Config# create profile Finance
Profile Secret [generate]:
Enable Profile? (yes, no) [no]: y
Monitor Capability? (yes, no) [no]:
Profile-Metadata Capability? (yes, no) [no]:
Profile Type (access, cluster) [access]:
Home Pool [default]:
Granted Rights for the Profile in the Home Pool [rdqeDcw]: rdqew
Issue the command?
(yes, no) [no]: y
Establish a Pool Entry Authorization for application use? (yes, no) [no]: y
Please enter Pool Authorization creation information: /var/tmp/Finance.pea
```

Privileges requested

File spec on local system to receive pea data

This command creates an application profile called Finance and the associated PEA file called /var/tmp/Finance.pea on the host system from which the CLI command is being run. The secret is auto-generated in this example; a user-supplied secret could also be used. The file created contains XML in a format understood by the Centera SDK and contains the credentials to allow an application to authenticate to the Centera.

Once the credentials have been created they should be given to the relevant application owner, once the file has been given to the application owner, the pea file created (in this case /var/tmp/Finanace.pea) may be moved to a secure area for archive or deleted. If the pea file becomes lost another pea file can be created and the secret changed.

The example in the following diagram shows a cluster setup with one application server running application 1, and two Centera clusters: cluster A (Primary) and cluster B (Target/COB). Replication is configured between cluster A and cluster B. Cluster A is the primary cluster for application 1. If cluster A is not available, the application will then fail-over to cluster B.




Figure 8 Application access to replicated Centeras

To support application authentication and authorization in this configuration, application profiles and their corresponding PEA files were created on both cluster A and cluster B. The content of the two PEA files were merged into a single PEA file, this file must in turn be stored on the application server, the specific location is normally determined by the configuration of the application to use the pea file.

As the replication process uses profiles to gain access to the remote Centera a profile must also be created on the target Centera. Replication profiles are normally created by EMC during the installation but may be created later if required.

Below is an example of the PEA file produced by the create profile command (the secret information is not reproduced here for clarity purpose).

```
<pea version="1.0.0">
  <defaultkey name="Finance">
    <credential id="csp1.secret" enc="base64">
      LOCALSECRETSTUFFHERE</credential>
    </defaultkey>
    <key type="cluster" id="5b6b8e46-1dd2-11b2-8c30-931792b40e1e"
name="Finance">
      <credential id="csp1.secret" enc="base64">
        LOCALSECRETSTUFFHERE </credential>
      </key>
</pea>
```

The PEA file produced only has the credentials for a single Centera. Considering, all CTI Centera are remotely replicated, a PEA file containing credentials for both the local and remote Centera is therefore required. Having both credentials ensures that the application can access both the local and remote copies of the data.

For example, a standard SDK "FPPool_Open" command will expect to be able to access both the local and remote Centera, if it cannot access both Centeras due to incorrect or missing credentials the command will fail.

Automatically creating a PEA file containing the credentials for both Centera is not possible with the current Centera software. To create a usable PEA file, credential files need to be created from both the local and remote Centera, and the key section from the remote Centera PEA file manually added to the PEA file of the local Centera to form a PEA file containing both sets of credentials, as shown below.

```
<pea version="1.0.0">
  <defaultkey name="Finance">
    <credential id="csp1.secret" enc="base64">
      LOCALSECRETSTUFFHERE</credential>
    </defaultkey>
    <key type="cluster" id="5b6b8e46-1dd2-11b2-8c30-931792b40e1e"
name="Finance">
      <credential id="csp1.secret" enc="base64">
        LOCALSECRETSTUFFHERE </credential>
      </key>
      <key type="cluster" id="b11ba322-1dd1-11b2-a069-ebb546c38de2"
name="Finance">
        <credential id="csp1.secret" enc="base64">
          REMOTESECRETSTUFFHERE </credential>
        </key>
      </key>
    </pea>
```

The application must also be able to find where the PEA file is located; this is normally achieved by setting an environment variable CENTERA_PEA_LOCATION that points to the application's PEA file.

It is also possible to supply the credential information directly to the Centera SDK API calls from within the application by supplying the profile name and credentials in the API application call. As this is related to specific application coding further explanation of this method of application authentication is outside the scope of this document.

11.5 Data protection

Centera has two data protection modes, Content Protection Mirrored (CPM) and Content Protection Parity (CPP).

CPP is discussed for information only as CPP is not certified for use within CTI.

CPP segments a data object into 6 parts and stores each one on a different node in the same cluster. CPP calculates a parity fragment from the stored fragments and stores that as the 7th data segment on yet another node. This mechanism is a derivative of RAID-3. Objects under 250KB in size are protected by CPM even if CPP is enabled. CPP typically reduces write performance by around 50%.

CPM stores a complete copy of the data object on a mirror node; hence, the failure of an individual node means that the data is still available. Data is not mirrored locally within each node. CPM enhances faster data regeneration and improves performance during normal operation. Object retrieval is improved because the node with the access role can select the least-loaded node with the storage role from which to retrieve a specific object. **CPM is the default protection scheme and the only scheme certified for use within CTI.**

Should a node or disk fail for any reason then the Centera automatically starts to re-create a copy of the data that was held on the failed node on another node. If a disk is reported as failed this will immediately trigger a regeneration task on every node that has a copy of the objects on the failed disk. If a node is unreachable by the other nodes for more than two hours all objects on that node will be regenerated.

An industry problem for any long-term storage of data is that of media obsolescence, e.g. today's tape or optical disk platter is unlikely to be in use 5,10 or 20 years. Centera access is via published API's over standard TCP/IP connections that provide a layer of abstraction between the application and its data. Each Centera is made up of standard nodes, however these will change in specification over time becoming more powerful and storing more data. The code that is used to replace failed nodes can also be used to migrate data from older to newer nodes as these become available, hence, avoiding the trap of media obsolescence or the need to manually transfer data. All that is required is the ability to connect to the Centera via published interfaces to access the data. This does of course pre-suppose that you have an application available that can make sense of the object when it has been retrieved from Centera and also that you have the Content Address keys of the original application that stored the data.

11.6 Remote copies of data (Replication)

A Centera may also make remote copies of data between two Centeras connected via a TCP/IP network, such as a corporate WAN/LAN. Any objects stored are automatically transmitted to the remote site, giving a total of 4 copies of the data, two local and two remote. Replication is asynchronous and uses the same network TCP port (3682) as application access to the Centera, hence, the source and target Centera must have access to each other over this network port.

Centera will only replicate data that is written after replication is enabled. i.e. data written before replication is enabled will not be automatically replicated. To copy data that was written before replication was enabled requires data to be copied manually by other mechanisms within the Centera (i.e. the Centera "restore" command).

In a replicated environment all data is transmitted to the remote Centera approximately 15 minutes after the data is first written to the primary Centera, a busy network link or heavy load on the Centera access nodes will increase the propagation delay. Using the "show replication detail" command the throughput and replication lag etc can be seen.

Replication may be paused and re-started as required. Replication items are queued and wait their turn to be transmitted. Any data waiting to be replicated whilst the link is down or paused remains in the queue and when the link is re-established the replication process will copy all outstanding data waiting to be replicated. If an object to be sent is larger than 1MB the replication process will check first if this object already exists (this is not performed via the normal API call and hence the "e" profile privilege is not required), if the object already exists then the data object will not be transmitted. If for whatever reason replication of a file fails after several retries, the file will be moved to the "parked" queue.

It is also possible to set the replication process to replicate deletes. This ensures that a delete on the source also takes place on the target Centera. This can only be enabled on a per Centera basis and hence all pools will have replicated deletes either enabled or disabled. **Best practice is to have this setting enabled, as it is more common for applications to require replicated deletes to be enabled.**

It is important that all applications sharing a Centera can cope with the option for replicated deletes, as some applications manage their own deletion of data on the remote Centera whilst others require replicated deletes.

When an application connects to a replicated Centera the application normally opens a connection to both Centera, the primary for normal transactions and the replica in case the primary is not available. It is also possible to connect to the Centera using a "lazy open" API call where a connection is only opened to the primary Centera and a connection to the replica is only created if the primary cannot be contacted. The choice of either a normal open or a lazy

open will be made by the application developer and as such is beyond the scope of this document. As data is replicated to both Centeras, data may be retrieved from either Centera using the same c-clip address.

The major difference between replication and the restore command is that the administrator initiates the restores whereas replication automatically runs on a timed basis.

A restore can either be a “full” or “partial” the restore also requires access to an access profile on the target of the restore via a .pea file, this access profile must have the clip copy “c” privilege for all pools into which it must write.

Full Restore: a “full” copies all clips on the source Centera to the target Centera. In order for the copy to be successful all virtual pools must exist on the target Centera before the restore is initiated. If a copy of the virtual pool containing the clips to be transferred does not exist on the target Centera these clips will not be transferred. The lack of a virtual pool on the target Centera, or the profile’s lack of access to the non-target pools are the only mechanisms available to limit the clips transferred to a particular selection of pools.

Partial Restore: A “partial” restore is similar to a “full” restore except that the user is prompted for a start and end date/time between which all clips will be transferred. No clips outside of this date/time range will be transferred and the transfer is subject to the same requirements regarding virtual pools and access profiles as a “full” restore.

To ensure that a restore starts as soon as possible, clips are transferred as soon as they are identified rather than scanning the whole Centera to build a full list. A side effect of this is that it is not possible to accurately estimate the amount of data to be transferred or the time this will take. Each hour, the next 1,000 clips to be transferred are queried and based on the write-timestamp of the first and last clip CentraStar determines the time slice these 1,000 clips represent. This time slice is the used to extrapolate an estimate of the remaining queue.

11.7 Performance

One of the implications of the Centera architecture is that there is a fixed overhead for a given transaction regardless of the size of the data object that is being transferred. Due to this overhead, throughput depends on the size of objects. Overall performance improves using large objects. The second implication is due to the single Ethernet port on each access node; the transfer of a single data object will not be faster than the maximum throughput of this port. As a compensation for the maximum transfer rate, it is recommend that applications make use of multithreading within the SDK. The architecture is highly parallel and supports many parallel activities. A single thread cannot take advantage of the multiple Access Nodes, and just as importantly, the multiple Storage Nodes. Of course there are limits to multithreading as well. Too much concurrent activity on the system will lead to tasks competing for resources, particularly tasks requiring disk resources on Storage Nodes, as these will create bottlenecks that cause a performance degradation. Since threads are distributed evenly over the available nodes, the number of Access and Storage Nodes influence the number of threads that can be supported. Clearly a 32-node cluster will perform better than a 16-node cluster.

Refer to “Figure 9 Centera Performance (Gen 4 Hardware, #Threads vs file size)” for performance details related to the number of Access Nodes, threads, and object size.

11.7.1 4-Node Centera discussion

The 4-node Centera is a reduced performance configuration that should only be used in circumstances where the performance of an application (or applications) can be defined and shown to be of a level that can be supported by the 4-node Centera.

The 2/4/4 configuration will support a maximum the total IO rate of less than 70 files/second read (100KB files) or 40 files/second write (100KB file). Depending on the file size the number of files/MB/second will vary dramatically, the 100KB size has been taken as a typical file size. Whilst replication runs at a lower priority than host read/writes care should be taken to include replication traffic created by writes in the overall traffic calculations as replication.

A calculated estimate of the impact of replication on traffic rates is shown below.

Assumptions

- Maximum performance figures as shown above
- 80% 20% split for Write/Read (in a normal disk environment this would be reversed with 80% of all traffic being reads but it is assumed that most the of the traffic will be writes due to its fixed content nature)
- Replication takes place 15 minutes after initial write
- Data should be replicated to the remote site as soon as possible (minimise latency of replication)
- IO write rate is continuous and hence replication traffic generated 15 minutes before will generate the same replication traffic as the writes currently being made.
- As replication takes place 15 minutes after initial write there is no caching of data to be replicated.
- Writes rates are identical at both local and replicated sites.

Writes at primary site only

Therefore – Mixed IO Rate = $80\% \text{ of IO Rate [Normal Write traffic]} \times (1/40)$
 $+ 80\% \text{ of IO Rate [Replication traffic outgoing]} \times (1/70)$
 $+ 20\% \text{ of IO Rate [Normal read traffic]} \times (1/70)$

The maximum throughput of Mixed IO would be a rate of 29 100KB files per second for a single replicated configuration.

Write at both primary and replication sites

Therefore – Mixed IO Rate = $80\% \text{ of IO Rate [Normal Write traffic]} \times (1/40)$
 $+ 80\% \text{ of IO Rate [Replication traffic outgoing]} \times (1/70)$
 $+ 20\% \text{ of IO Rate [Normal read traffic]} \times (1/70)$
 $+ 80\% \text{ of IO rate [Replication traffic incoming]} \times (1/40)$

The maximum throughput of Mixed IO would be a rate of 18 100KB files per second for a Bi-directionally replicated configuration.

Hence, for a 4-node Centera that is a bi-directional replicated configuration, the maximum IO rate to the primary Centera would be:

1. 29 (100KB) objects / Second where only the primary Centera is being written to.
2. 18 (100KB) objects / Second where both the primary and replica Centera are being written to

12 Product Risk Acceptance and ICQ

A Product Risk Acceptance (**222316**) has been completed for the EMC Centera product. An ICQ (Infrastructure Compliance Questionnaire) has also been completed for the Centera storage product. The ICQ reference number is **234302**. The ICQ has highlighted security gaps in the Centera and its associated software. This solution guide details the compensating controls for the product but for further details refer to the ICQ entry in Archer (<https://archer.ti.citigroup.net/>).

This product is dependant on an SOE Management server (Solaris), CenteraCLI, a CEP/firewall and PowerBroker. To meet the security requirements the following compensating controls are to be implemented.

Security Issue	Compensating Control
Password Management (CISS 2.1, 2.2.2, 2.4, 2.4.1, 2.5, 2.8) Failed Login Attempts (CISS 2.7) Unique UserIDs (CISS 2.2.1) Warning Banner (CISS 3.7)	SOE Solaris management workstation configured to limit management access to the Centera management team. The SOE management system(s) are the only system(s) able to reach the management ports on the Centera through the CEP/firewall.
Authorisation and Access Controls	Access to the Centera CLI management interface is limited to the Centera management team via PowerBroker. The management server(s) are the only systems able to connect to the management ports of the Centera in the CEP. All access to the CLI is controlled and audited by PowerBroker. Application Access - Applications accessing the Centera are required to authenticate themselves using a challenge-response mechanism based on RFC 2104 HMAC (Hashed Message Authentication Codes), the Centera HMAC implementation uses MD5 hashing as its cryptographic hashing algorithm. Each application is given a 256 byte key that is paired with its access profile, keys are generated using the java SecureRandomClass when application profiles are created on the Centera. This key is used in combination with HMAC to authenticate the application before access to Centera data is granted. It should be noted that using HMAC at no time is this secret key sent across the network, only the HMAC codes derived from the applications key are sent.
Auditing (CISS 3.4)	The Centera does not provide any auditing. The current configuration of the profile, pools and replication settings should be reconciled with expected configuration as part of a local Level-3 process.
Incident Detection/Response	Whilst all management commands are logged via PowerBroker. There is no logging of application authentication failures within the

	Centera. This is detailed as a "Fail" to ICQ item 5.2.1
Role Based Access Control (CISS 6.3, 6.41)	The Centera does not support the use of roles based accounts for management. The security administration role will be preformed as part of the PowerBroker administration/logging process. Centera administrators will not have access to the PowerBroker configuration nor administrator rights on the SOE Solaris management server and hence cannot perform user or logging administration.
Maker Checker Verification (CISS 6.4.2)	A level 3 documented process will need to be implemented to provide maker/checker functions for any administrative actions to the Centera environment.
Remote Vendor Access	The Centera sends alerts and periodic encrypted (AES-256) health reports to EMC but remote access to the Centera is not configured and hence EMC have no remote access to the device.

13 Appendix A – Performance Stats

32-node	Read (files/sec)					Write (files/sec)				
	1	10	20	30	40	1	10	20	30	40
1 KB	39	375	723	414	897	3	33	66	100	116
5 KB	42	430	721	811	1167	3	38	75	109	140
10 KB	44	466	723	845	900	4	38	77	112	139
15 KB	45	427	694	846	891	3	35	74	109	128
20 KB	47	429	833	796	866	3	35	71	102	118
50 KB	45	387	627	689	801	3	36	73	105	130
100 KB	37	297	466	612	688	3	35	69	99	128
200 KB	33	267	426	491	343	3	36	70	99	121

Figure 9 Centera Performance (Gen 4 Hardware, #Threads vs file size)




Figure 10 Replication throughput (Mirrored protection 50KB files)




Figure 11 Replication throughput (Mirrored protection 10MB files)

14 Appendix B – Blank EMC Centera Change Control Request Form

Centera Change Control Form Notes and Rules for Submission

- ALL CCA activities must be accompanied by the latest version of the Centera CCRF and a copy of the Centera Procedure Generator procedure to be followed for this activity. There will be no exceptions to this requirement due to new settings and functionality in the new CentraStar release that addresses data security issues.
- CCRF form must be completely filled out. For new installs complete sections 1-7. For Upgrades or re-configurations complete sections 1-7 & sections 8-9
- For all install, upgrade or re-configurations, CE must also upload the corresponding procedure as provided by the latest Centera Procedure Generator.
- If installing a CUA (CAG) in addition to the Centera, the CE must also include the CUA (CAG) CCRF.
- If CUA is to be installed on a Gen3 node, the minimum number of Centera nodes in the cluster after the CUA is installed is 8. If the customer has only purchased 8 nodes and one node is to be installed as a CUA this is not an approved configuration and the CCA will be rejected.
- If CUA is to be installed on a Gen4 node, the minimum number of Centera nodes in the cluster after the CUA is installed is 4 or 8. IF the customer has only purchased 4 or 8 nodes and one node is to be installed as a CUA this is not an approved configuration and the CCA will be rejected.
- CE must complete Contact information below for ease of contacting CCA submitter. Failure to do so will result in activity rejection.
- Section 2 is a requirement. Failure to answer all the questions in section 2 will result in the activity being rejected. See Primus emc101608.
- Section 2 confusion. Frequently the field is providing data about layered applications. This happens when an application is writing to the CUA or Legato Disk Extender to list some examples. In these cases we ***only*** want information about those applications writing directly to the Centera. We do ***not*** want to know about the PACS application that might be writing to the CUA or Disk Extender.
- Garbage Collection confusion. The responses to the Garbage Collection questions are frequently contradicting one another. There are 3 possible sets of answers: All 3 questions set to yes/enabled, all three questions set to off/disabled, and lastly, the feature of Garbage Collection is on, Full GC is enabled, but Inc GC is disabled. Any other combination will result in the activity being rejected.
- For version 3 hardware and below a cabinet, a rack, and a cube all represent 32 nodes, but with version 4 hardware a cube is 16 nodes and two cubes can be mounted in one cabinet/rack. Please be aware of this as you go through the form.
- Target code level: This must be a complete version number. Phrases like “latest” are ***not*** allowed.

CCRF is updated to support CentraStar 3.0 features and functionality

- o Section 2: Application information
 - Added CUA question.
 - Added statement regarding all applications writing to Centera
- o Section 3: Storage on Access
 - A CE+ cluster must have a minimum of 2 nodes storage only role.
 - Max of 8 Access roles per rack.
 - Changed default installed version to capture actual version of CentraStar to be installed.
- o Section 6:
 - Added replication configuration questions.
 - Added Garbage Collection questions.

Centera Presite and CCRF tasks			
Tasks*	Description	Date Completed	Initials
Task 1	SE completes Section 1 & 2 and obtains approval		
Task 2	SE attends customer handoff meeting to discuss proposed config, provides an approved doc to CE. Section 3 & 4 is completed by CE & SE.		
Task 3	SE obtains customer sign-off of CCRF		
Task 4	CE submits CCRF and receives approvals (CCA process)		
Task 5	CE performs installation		
Task 6	CE communicates details of completed install to SE		

CE Contact information:

CE Name:		CE Cell Phone Number:	
CE Email Address:		CE Acknowledges having read the notes on this page	☐ Yes ☐ No

NOTES: All Fields marked with (\$) require an optional license that must be purchased by the customer. Do not configure without a license.
 All fields marked with (*) are options. Do not answer unless customer requires this functionality.
 Click check box ☐ to change default value. Always check for the latest version of this document:
 EMC: http://www.cs.isus.emc.com/csw22/centera2/forms/Centera_Change_Control.doc
 Partner: https://emc.clariion.com/restricted/centera/forms/Centera_Change_Control.doc
 Configurations greater than 4 racks require an RPQ

Please select the appropriate service event.

- ☐ - New Installation – Only fill in sections 1 – 7
☐ - Upgrade or Reconfiguration – Fill in sections 1 – 9

Section 1: Fill in shipping and install address along with Customer and EMC contacts.

Site Name	Site ID		
Shipping Address			
Installation Address			
	Name	Phone	Email Address
Primary Site Contact			
Secondary Site Contact			
EMC TSG TC			
EMC Centera SE			
Network Administrator			
EMC District Service Mgr			
EMC Customer Engineer			

Section 2: Fill in application supplier details. See bullets on page 1 of this document.

Is there a CUA (CAG)?	☐ Yes ☐ No	
If Yes – What is the Version of the CUA software?		Separate CUA CCRF must be filled in.
If more than one application is writing to the Centera, List all that are writing to Centera (word wrap enabled)		
Application Supplier(s)		Application Name(s)
Application Version(s)		SDK Version(s)
☐ WIN 2003/2000/NT4 ☐ Solaris 2.9/2.8/2.6 ☐ AIX 5.2/5.1/4.3.3 ☐ RedHat 7.1/7.2 ☐ IRIX 6.5.17 ☐ HP-UX 11i/11.0 ☐ z/OS		

Section 3: Provide cluster configuration details.

Cabinets/Racks (1-4)	Cubes (1-8)	Final Node Count (4-128)	Access Role Only	Storage Role Only	Dual Role S/AN
Centera Node Type(s)		Target Code Level Whole version number only		☐ If other than latest, please provide justification in section 7.	
☐ Gen1 ☐ Gen2 ☐ Gen3 ☐ Gen4					
If installing Gen4 hardware in a customer supplied rack an RPQ is required. Please attach RPQ to CCA.					

Section 4: Provide Dial-In and remote alerting configuration.

Remote Dial-In and Alerting Options		
A	Will customer allow a 33.6 Modem attached to cluster? (Dial-In Only) Preferred	☐ Yes ☐ No
B	Will the customer allow SMTP (Email) alerting to EMC? (Outbound Only) Preferred	☐ Yes ☐ No
C	OnAlert Workstation? (Secondary Option). (Not Required if A&B are chosen)	☐ Yes ☐ No
D	If none of the above were selected has the District Service Manager been notified?	☐ Yes ☐ No
E	Modem Phone Number A	*Modem Phone Number B

Section 5: Provide network connectivity details.

Network Connectivity	☐ 100BaseT ☐ 1000BaseT ☐ SFP (Gen4 – Optical)	Copper-to-Fiber Converter?	Transition Networks ☐ Yes ☐ No
Switch Model and Make		Centera Internal Subnet	☐ 10.255 ☐ 172.19

Access Node Configuration (AN= Access Node)							
Subnet Mask:		Gateway IP:		DNS Server:			
☒ Distribute Access Nodes in even pairs across multiple racks. ☐ Alternate Access Node Distribution (Provide comments on the specific configuration in Section #7)							
AN1		AN2		*AN3		*AN4	
*AN5		*AN6		*AN7		*AN8	
*AN1 1		*AN1 1		*AN1 2		*AN12	
*AN1 3		*AN1 4		*AN1 5		*AN16	

Section 6: Enter cluster configuration details.

Function	Parameter	Data (default)
set owner	Administrator Name:	
	Administrator email:	
	Administrator phone:	
	Location of the Cluster:	
	Name of the Cluster:	
	Serial Number: Located on 1 st cube	
set cluster notification (2.3 & > have default EMC alert email addresses enabled)	ConnectEMC Mode: (on, off)	☐ On ☐ Off
	smtp server1:	
	smtp server2:	
	ConnectEMC Reporting Interval?	☐ 1 Day
	Additional delay before next report?	NONE
	Email recipients: (customer only)	
	Reply address:	
	Domain to which the cluster belongs:	
set cluster service	Site ID: (EMC Clarify Site ID)	
	Service ID:	☐ Premium
	Service info: (enter specific notes)	
set cluster replication (\$)	*Replication address: (x.x.x.x:3218)	
	*Replication Enabled?	☐ Yes ☐ No
	If Yes, this Cluster is the: (If bidirectional replication is configured – chose both)	☐ Source ☐ Target
	*Replicate Delete (v2.3 and higher)	☐ Yes ☐ No
	*Replicate incoming replicated Objects?	☐ Yes ☐ No
	*Replicate all Pools?	☐ Yes ☐ No
	*Replication Profile Name: (- = anonymous)	
set security compliance ce (\$)	*Compliance Edition enabled?	☐ Yes ☐ No
set security compliance ce+ (\$)	*Compliance Edition Plus enabled?	☐ Yes ☐ No
set feature	*data-shredding (\$) (CE and CE+ only)	☐ On ☐ Off
	storage-strategy: Check Application Compatibility	☐ Capacity ☐ Performance
	If performance – full or partial	☐ Full ☐ Partial
	Threshold for storage-strategy	☐ 250KB ☐ Other:
	*garbage-collection (see page 1 for valid settings)	☐ On ☐ Off
set GC (See bullet on page 1)	Incremental GC	☐ Enable ☐ Disable

	Full GC	☐ Enable	☐ Disable
set M++	*MD5 Naming Scheme	☐ Enabled	☐ Disabled ☐ Disabled for CC
*set protection write parity (\$)	*Set protection write parity enabled?	☐ Yes	☐ No
	Threshold for parity	☐ 250KB	☐ Other:
	Fallback protection scheme (v3.0)	☐ [-]	☐ m2
*set capacity regenerationbuffer	*Mode: Alert or Stop.	☐ Alert	☐ Stop
	Buffer Threshold: (v2.3/2.4)	☐ 375GB	☐ Other:
	Limit in disks per cube (v3.0)	☐ 1 Disk	☐ Other:
*set ICMP	*Enables PING capability	☐ Yes	☐ No
*set security default retention (\$)	*expressed in seconds, minutes, days, months & years		
*set SNMP	Enable SNMP	Management Station IP	Community Name
	☐ Yes ☐ No		Heartbeat Trap (Minutes)

Section 7: Please provide comments and include special instructions, notes and RPQ number:

Notes:

STOP

Section 8: This section is only required for cluster upgrades and reconfigurations.

14.1.1.1.1 HARDWARE UPGRADES	
Is this an 8 Node Gen2 or Gen3 Storage Upgrade (CNR8TBUP, CNR8NUPG3)? <u>If yes then,</u>	☐ Yes ☐ No
What is the installed software revision? (enter current software revision >)	
Is this a 4 Node Gen4 Storage Upgrade (CNR8TBUP, CNR8NUPG3)? <u>If yes then,</u>	☐ Yes ☐ No
What is the installed software revision? (enter current software revision >)	
Is this adding Gen4 (V4) cube to a Gen2/3 (V2/3) cluster (CNRPDU + CNR4NFBAG4)?	☐ Yes ☐ No
Is this adding Gen4 (V4) cube to a Gen4 cluster (CNRPDU + CNR4NFBAG4)?	☐ Yes ☐ No
Is this adding Gen4 (V4) nodes to a Gen4 cluster (CNR4NFUPG4)?	☐ Yes ☐ No
Is this adding one or more cabinets? <u>If yes then,</u>	☐ Yes ☐ No
What is the installed software revision? (enter current software revision >)	
Is this adding Gen3 (V3) rack to a Gen 1 or 2 or 3 (V1/2/3) cluster?	☐ Yes ☐ No
Is this adding Gen4 (V4) rack to a Gen2/3 (V2/3) cluster (CNRRK + CNR4NFBAG4)?	☐ Yes ☐ No
Is this adding Gen4 (V4) rack to a Gen4 cluster (CNRRK + CNR4NFBAG4)?	☐ Yes ☐ No
Is this adding a 5 th , 6 th , 7 th or 8 th rack	☐ Yes ☐ No
If yes must provide current cluster object count/node and RPQ number	
14.1.1.1.2 SOFTWARE UPGRADE	
Is this a CentraStar upgrade? <u>If yes then,</u>	☐ Yes ☐ No
What is the installed software revision? (enter current software revision >)	
What version do you plan to install? (enter software revision >)	
Garbage Collection	
Is this to change the configuration of GC?	☐ Yes ☐ No
Incremental GC	☐ On ☐ Off
Full GC	☐ On ☐ Off
M++	
MD5 Naming Scheme	☐ Enabled ☐ Disabled ☐ Disabled for CC
UPGRADE LICENSE	
CNR8NODEPUPG RTU License for Basic Operation Parity Protection – for Gen2 HW Upgrade)	☐ Yes ☐ No
CNR8NPUPG3 RTU License for Basic Operation Parity Protection – for Gen3 HW (Upgrade)	☐ Yes ☐ No
CNR4NPUPG4 RTU License for Basic Operation Parity Protection – for Gen4 HW Upgrade)	☐ Yes ☐ No
CNRREPSW RTU License for Replication	☐ Yes ☐ No
CNRREPSWG4 RTU License for Replication Gen4	☐ Yes ☐ No
CNRMANSW RTU License for Storage on Access (2 Nodes per lic. CPM Only)	☐ Yes ☐ No
CNRPANSW RTU License for Storage on Access (3 Nodes per lic. CPP Only)	☐ Yes ☐ No
CNRCELIC RTU Centera CE Gen 2/3	☐ Yes ☐ No
CNRCEPLUSLIC RTU Centera CE+ Gen2/3	☐ Yes ☐ No
CNRCELICG4 RTU Centera Governance Gen4 (CE)	☐ Yes ☐ No
CNRCEPLUSG4 RTU Centera CE+ Gen4	☐ Yes ☐ No
CNRSWRK2RK RTU Centera Multi-rack	☐ Yes ☐ No

14.1.1.1.3 SYSTEM REINITIALIZATION

Is this a Data deletion request? <u>If yes then,</u>	☐ Yes ☐ No
What is the requested date? (Allow 2-3 business days see PRIMUS emc48888)	
Is this a Factory Initialization request? <u>If yes then,</u>	☐ Yes ☐ No
What is the requested date? (Allow 2-3 business days see PRIMUS emc51666)	

SYSTEM POWER DOWN / RELOCATION

Is this a system power down or relocations request? <u>If yes then, see PRIMUS emc66477</u>	☐ Yes ☐ No
---	--

14.1.1.1.4 NODE RECONFIGURATION

Is this a request to change a node(s) from Access to Storage?	☐ Yes ☐ No
Is this a request to add a node(s) role? (dual role S/AN – 2.4 and above)	☐ Yes ☐ No
Is this a request to change a node(s) from Storage to Access? (RPQ for production systems)	☐ Yes ☐ No
Is this a request to change a node(s) from Storage to CAG? (RPQ for production systems)	☐ Yes ☐ No

PLATFORM MIGRATION

Is this migrating CUA from a DELL platform to a Centera node (in a Centera Gen3/4 cluster)?	☐ Yes ☐ No
Is this migrating CUA from a Centera node (in a Centera Gen3/4 cluster) to a DELL platform?	☐ Yes ☐ No

MECHANICAL REPLACEMENT

Is this a Mechanical Replacement?	☐ Yes ☐ No
If Yes, provide details in section 9.	

Section 9: Please provide comments and include special instructions, notes and RPQ number:

Notes:

15 Appendix C – Example Centera Health Report

Please note that the health report shown below does not come from a production Centera running CE+ and is provided to show the type of information provided in a health report. It does not show an example of a report from a production Centera configure to CTI standards.

Centera Health Report

Sequence Number # 42

Report Format Version 1.3

Report Identification

Type: health

Format DTD: health-1.3.dtd

Report Interval: 24 hrs 0 min 0 sec

Creation Date and Time: 10-11-2005 07:53:00.995 UTC

Cluster Identification

Name: Centera2

Serial Number: APM00XXXXXXXXXX

ClusterID: b7f0e2f0-1dd1-11b2-955b-123456789012

ClusterDomain: citigroup.com

License List

Cluster Service Settings

Site ID: 12345678

Customer Name:

Location: 388 Greenwich St. 12th flr Lab

Admin Contact: Local Admin

Admin Email: Local_admin@citigroup.com

Admin Phone: TBD

Modem Phone1: not configured

Modem Phone2: not configured

Active Software
Version: 3.0.0-542-557-8615

Activation Manager:

Distribution Manager:

Regeneration Buffer: 647856340992

Regeneration Buffer
Policy: HardStop

Compliance Mode: Basic

Default Retention
Period: 0

Replication Status: Running

ReplicationOfDeletes : Enabled

Replication Profile: rep_centera2

Replication IP: 146.128.135.117:3218,146.128.135.118:3218,146.128.135.119:3218,146.128.135.120:3218

SNMP Status:

SNMP Community: public

SNMP IP Address:

SNMP Port: 162

SNMP Trap Interval:

Notification: Enabled

Protection Scheme: M2

Protection Threshold: 0

Protection Scheme Alternate: M2

Naming Scheme: Performance

Naming Scheme Mode: Full

Naming Scheme Threshold: 256000

Garbage Collection Status: Enabled

Shredding Status: Disabled

Hardware Components

Cabinet Information

Name	Serial #	Revision #	Active software version	Total Capacity	Sys Resources	Offline Capacity	Used Capacity	Audit MetaData	Protected User Data	Free Capacity	Spare Capacity	Available Capacity	Avail Cap Until Alert	System Buffer	Regeneration Buffer
1				19308 GB	106 GB	0 GB	2 GB	2 GB	0 GB	14399 GB	4800 GB	12686 GB	13621657038848	1191064895488	647856340992

Cabinet 1

Switch Information

Cluster Prefix: 10.255

Name	Type	MAC Addr	Serial #	Revision #	Firmware	Rail	Description	IP Addr
c001sw0	Cabinet	00:00:cd:0f:e0:60	58368265			0	Allied Telesyn AT-RP48 Rapier 48i version 2.4.1-01 01-May-2003	10.255.1.61
c001sw1	Cabinet	00:00:cd:11:c6:7c	58781389			1	Allied Telesyn AT-RP48 Rapier 48i version 2.4.1-01 01-May-2003	10.255.1.62

Port Information

c001sw0

Port Number	Type	Speed Setting	Status
41	trunk		Up
42	trunk		Up
43	trunk		Up
44	trunk		Up
45	trunk		Up
46	trunk		Up
47	trunk		Up
48	trunk		Up

49	uplink		Not installed
50	uplink		Not installed

c001sw1

Port Number	Type	Speed Setting	Status
41	trunk		Up
42	trunk		Up
43	trunk		Up
44	trunk		Up
45	trunk		Up
46	trunk		Up
47	trunk		Up
48	trunk		Up
49	uplink		Not installed
50	uplink		Not installed

ATS Information

Modem Information

Node Information

Node	Serial #	Revision #	Sys ID	Model	Role	Rail	Type	Motherboard
c001n01	40100312		ce299332-1dd1-11b2-936b-df4b46c5f919	118032400-A06	access	B		
c001n02	40100326		c013be62-1dd1-11b2-8854-be05392823ba	118032400-A06	access	A		
c001n03	50400120		d7802a54-1dd1-11b2-a59a-cc03754b8b9c	118032400-A09	access	B		
c001n04	50200097		b725e320-1dd1-11b2-86ed-d0099200a208	118032400-A09	access	A		
c001n05	35200306		9616a16a-1dd1-11b2-955b-88d18c53b922	118032400-A06	bic,storage	B		
c001n06	44300017		bf8fe39e-1dd1-11b2-9883-dd8f8924dc3f	118032400-A09	bic,storage	A		
c001n07	40100353		c92377c2-1dd1-11b2-87d6-eed5dac348c7	118032400-A06	bic,storage	B		
c001n08	40100348		ca4f67fa-1dd1-11b2-88d3-c04025071a92	118032400-A06	bic,storage	A		
c001n09	40100251		c6ec1d38-1dd1-11b2-8d90-eeef07480c60	118032400-A06	bic,storage	B		
c001n10	40100099		c0512540-1dd1-11b2-9669-b1d2139d7510	118032400-A06	bic,storage	A		
c001n11	40100242		ba4c986e-1dd1-11b2-a704-86eb22f11428	118032400-A06	bic,storage	B		
c001n12	40100057		9a484ef0-1dd1-11b2-bb7c-b17769cc2f45	118032400-A06	bic,storage	A		
c001n13	35200240		a7e6c140-1dd1-11b2-b32a-ffef8dba59e0	118032400-A06	bic,storage	B		

c001n14	40100083		c715c912-1dd1-11b2-a849-ed8fd46f9116	118032400-A06	blc,storage	A		
c001n15	35200157		f7f594c6-1dd1-11b2-bf17-d09dbc49dff7	118032400-A06	blc,storage	B		
c001n16	44800149		a2aaab42-1dd1-11b2-becd-e313edf84e38	118032400-A09	blc,storage	A		

Node Health and Service Settings

Node	Node Status	Maintenance Mode	Compliance Mode	SW Ver	Hot Fixes	Platform Ver	Platform Hot Fixes
c001n01	online		basic	3.0.0-542-557-8615			
c001n02	online		basic	3.0.0-542-557-8615			
c001n03	online		basic	3.0.0-542-557-8615			
c001n04	online		basic	3.0.0-542-557-8615			
c001n05	online		basic	3.0.0-542-557-8615			
c001n06	online		basic	3.0.0-542-557-8615			
c001n07	online		basic	3.0.0-542-557-8615			
c001n08	online		basic	3.0.0-542-557-8615			
c001n09	online		basic	3.0.0-542-557-8615			
c001n10	online		basic	3.0.0-542-557-8615			
c001n11	online		basic	3.0.0-542-557-8615			
c001n12	online		basic	3.0.0-542-557-8615			
c001n13	online		basic	3.0.0-542-557-8615			
c001n14	online		basic	3.0.0-542-557-8615			
c001n15	online		basic	3.0.0-542-557-8615			
c001n16	online		basic	3.0.0-542-557-8615			

Node Capacity

Node	Total Capacity	Sys Resources	Offline Capacity	Used Capacity	Audit MetaData	Protected User Data	System Buffer	Available Capacity	Free Capacity	Used Obj Count	Total Obj Count	Free Obj Count
c001n01	1207 GB	7 GB	0 GB	0 GB	0 GB	0 GB	0	0 GB	0 GB	0	0	0
c001n02	1207 GB	7 GB	0 GB	0 GB	0 GB	0 GB	0	0 GB	0 GB	0	0	0
c001n03	1207 GB	7 GB	0 GB	0 GB	0 GB	0 GB	0	0 GB	0 GB	0	0	0
c001n04	1207 GB	7 GB	0 GB	0 GB	0 GB	0 GB	0	0 GB	0 GB	0	0	0
c001n05	1207 GB	7 GB	0 GB	0 GB	0 GB	-1	96750010368	1110 GB	1200 GB	1	30000000	29999999
c001n06	1207 GB	7 GB	0 GB	0 GB	0 GB	-1	111782395904	1096 GB	1200 GB	1	30000000	29999999
c001n07	1207 GB	7 GB	0 GB	0 GB	0 GB	-1	96750010368	1110 GB	1200 GB	0	30000000	30000000

c001n08	1207 GB	7 GB	0 GB	0 GB	0 GB	-1	96750010368	1110 GB	1200 GB	3	30000000	29999997
c001n09	1207 GB	7 GB	0 GB	0 GB	0 GB	-1	96750010368	1110 GB	1200 GB	1	30000000	29999999
c001n10	1207 GB	7 GB	0 GB	0 GB	0 GB	-1	96750010368	1110 GB	1200 GB	0	30000000	30000000
c001n11	1207 GB	7 GB	0 GB	0 GB	0 GB	-1	96750010368	1110 GB	1200 GB	2	30000000	29999998
c001n12	1207 GB	7 GB	0 GB	0 GB	0 GB	-1	96750010368	1110 GB	1200 GB	2	30000000	29999998
c001n13	1207 GB	7 GB	0 GB	0 GB	0 GB	-1	96750010368	1110 GB	1200 GB	4	30000000	29999996
c001n14	1207 GB	7 GB	0 GB	0 GB	0 GB	-1	96750010368	1110 GB	1200 GB	3	30000000	29999997
c001n15	1207 GB	7 GB	0 GB	0 GB	0 GB	-1	96750010368	1110 GB	1200 GB	1	30000000	29999999
c001n16	1207 GB	7 GB	0 GB	0 GB	0 GB	-1	111782395904	1096 GB	1200 GB	0	30000000	30000000

CPU Information

c001n01

Serial #	Revision #	CPU Temp (Celsius)	Case Temp (Celsius)
		36	

c001n02

Serial #	Revision #	CPU Temp (Celsius)	Case Temp (Celsius)
		34	

c001n03

Serial #	Revision #	CPU Temp (Celsius)	Case Temp (Celsius)
		34	24

c001n04

Serial #	Revision #	CPU Temp (Celsius)	Case Temp (Celsius)
		34	25

c001n05

Serial #	Revision #	CPU Temp (Celsius)	Case Temp (Celsius)
		38	

c001n06

Serial #	Revision #	CPU Temp (Celsius)	Case Temp (Celsius)
		40	26

c001n07

Serial #	Revision #	CPU Temp (Celsius)	Case Temp (Celsius)
----------	------------	--------------------	---------------------

		37	
c001n08			
Serial #	Revision #	CPU Temp (Celsius)	Case Temp (Celsius)
		39	
c001n09			
Serial #	Revision #	CPU Temp (Celsius)	Case Temp (Celsius)
		39	
c001n10			
Serial #	Revision #	CPU Temp (Celsius)	Case Temp (Celsius)
		43	
c001n11			
Serial #	Revision #	CPU Temp (Celsius)	Case Temp (Celsius)
		43	
c001n12			
Serial #	Revision #	CPU Temp (Celsius)	Case Temp (Celsius)
		41	
c001n13			
Serial #	Revision #	CPU Temp (Celsius)	Case Temp (Celsius)
		51	
c001n14			
Serial #	Revision #	CPU Temp (Celsius)	Case Temp (Celsius)
		46	
c001n15			
Serial #	Revision #	CPU Temp (Celsius)	Case Temp (Celsius)
c001n16			
Serial #	Revision #	CPU Temp (Celsius)	Case Temp (Celsius)
		40	29

Power Supply Information

NIC Information

c001n01

Name	Config	DNS IP Addr	Duplex	IP Addr	Link Speed	MAC Addr	Revision #	Serial #	Subnet	Type	Status
eth0				10.255.1.1		00:E0:81:60:9A:3A					enabled
eth1				10.255.1.1		00:E0:81:60:9A:39					enabled
eth2	Manual	146.128.201.218	F	146.128.135.121	100	00:0E:0C:06:04:80			255.255.255.0		enabled

c001n02

Name	Config	DNS IP Addr	Duplex	IP Addr	Link Speed	MAC Addr	Revision #	Serial #	Subnet	Type	Status
eth0				10.255.1.2		00:E0:81:60:99:06					enabled
eth1				10.255.1.2		00:E0:81:60:99:05					enabled
eth2	Manual	146.128.201.218	F	146.128.135.122	100	00:0E:0C:06:D2:22			255.255.255.0		enabled

c001n03

Name	Config	DNS IP Addr	Duplex	IP Addr	Link Speed	MAC Addr	Revision #	Serial #	Subnet	Type	Status
eth0				10.255.1.3		00:E0:81:63:81:E4					enabled
eth1				10.255.1.3		00:E0:81:63:81:E3					enabled
eth2	Manual	146.128.201.218	F	146.128.135.123	100	00:0E:0C:6C:27:46			255.255.255.0		enabled

c001n04

Name	Config	DNS IP Addr	Duplex	IP Addr	Link Speed	MAC Addr	Revision #	Serial #	Subnet	Type	Status
eth0				10.255.1.4		00:E0:81:63:81:52					enabled
eth1				10.255.1.4		00:E0:81:63:81:51					enabled
eth2	Manual	146.128.201.218	F	146.128.135.124	100	00:0E:0C:38:8B:60			255.255.255.0		enabled

c001n05

Name	Config	DNS IP Addr	Duplex	IP Addr	Link Speed	MAC Addr	Revision #	Serial #	Subnet	Type	Status
eth0				10.255.1.5		00:E0:81:60:9A:84					enabled
eth1				10.255.1.5		00:E0:81:60:9A:83					enabled
eth2						00:0E:0C:07:42:80					disabled

c001n06

Name	Config	DNS IP Addr	Duplex	IP Addr	Link Speed	MAC Addr	Revision #	Serial #	Subnet	Type	Status
eth0				10.255.1.6		00:E0:81:61:E8:35					enabled
eth1				10.255.1.6		00:E0:81:61:E8:34					enabled

eth2						00:0E:0C:38:A4:D0					disabled
------	--	--	--	--	--	-------------------	--	--	--	--	----------

c001n07

Name	Config	DNS IP Addr	Duplex	IP Addr	Link Speed	MAC Addr	Revision #	Serial #	Subnet	Type	Status
eth0				10.255.1.7		00:E0:81:60:9C:FA					enabled
eth1				10.255.1.7		00:E0:81:60:9C:F9					enabled
eth2						00:0E:0C:07:49:60					disabled

c001n08

Name	Config	DNS IP Addr	Duplex	IP Addr	Link Speed	MAC Addr	Revision #	Serial #	Subnet	Type	Status
eth0				10.255.1.8		00:E0:81:60:9D:B0					enabled
eth1				10.255.1.8		00:E0:81:60:9D:AF					enabled
eth2						00:0E:0C:07:43:B6					disabled

c001n09

Name	Config	DNS IP Addr	Duplex	IP Addr	Link Speed	MAC Addr	Revision #	Serial #	Subnet	Type	Status
eth0				10.255.1.9		00:E0:81:60:9D:5A					enabled
eth1				10.255.1.9		00:E0:81:60:9D:59					enabled
eth2						00:0E:0C:07:3B:AE					disabled

c001n10

Name	Config	DNS IP Addr	Duplex	IP Addr	Link Speed	MAC Addr	Revision #	Serial #	Subnet	Type	Status
eth0				10.255.1.10		00:E0:81:60:9A:64					enabled
eth1				10.255.1.10		00:E0:81:60:9A:63					enabled
eth2						00:0E:0C:07:41:72					disabled

c001n11

Name	Config	DNS IP Addr	Duplex	IP Addr	Link Speed	MAC Addr	Revision #	Serial #	Subnet	Type	Status
eth0				10.255.1.11		00:E0:81:60:9B:04					enabled
eth1				10.255.1.11		00:E0:81:60:9B:03					enabled
eth2						00:0E:0C:07:49:74					disabled

c001n12

Name	Config	DNS IP Addr	Duplex	IP Addr	Link Speed	MAC Addr	Revision #	Serial #	Subnet	Type	Status
eth0				10.255.1.12		00:E0:81:60:9E:6C					enabled
eth1				10.255.1.12		00:E0:81:60:9E:6B					enabled

eth2						00:0E:0C:07:43:C0					disabled
------	--	--	--	--	--	-------------------	--	--	--	--	----------

c001n13

Name	Config	DNS IP Addr	Duplex	IP Addr	Link Speed	MAC Addr	Revision #	Serial #	Subnet	Type	Status
eth0				10.255.1.13		00:E0:81:60:9A:C2					enabled
eth1				10.255.1.13		00:E0:81:60:9A:C1					enabled
eth2						00:0E:0C:07:3C:EA					disabled

c001n14

Name	Config	DNS IP Addr	Duplex	IP Addr	Link Speed	MAC Addr	Revision #	Serial #	Subnet	Type	Status
eth0				10.255.1.14		00:E0:81:60:A0:10					enabled
eth1				10.255.1.14		00:E0:81:60:A0:0F					enabled
eth2						00:0E:0C:07:4B:AC					disabled

c001n15

Name	Config	DNS IP Addr	Duplex	IP Addr	Link Speed	MAC Addr	Revision #	Serial #	Subnet	Type	Status
eth0				10.255.1.15		00:E0:81:60:98:8C					enabled
eth1				10.255.1.15		00:E0:81:60:98:8B					enabled
eth2						00:0E:0C:06:02:AC					disabled

c001n16

Name	Config	DNS IP Addr	Duplex	IP Addr	Link Speed	MAC Addr	Revision #	Serial #	Subnet	Type	Status
eth0				10.255.1.16		00:E0:81:64:7E:4A					enabled
eth1				10.255.1.16		00:E0:81:64:7E:49					enabled
eth2						00:07:E9:2E:0C:42					disabled

Drive Information

c001n01

Index	Serial #	Revision #	Model	Status	Mode	Total Capacity	Sys Resources	Used Capacity	Audit MetaData	Protected User Data	Free Capacity	System Buffer	Available Capacity	Used Obj Count
/dev/cstardb:0				enabled		302 GB	2 GB	0 GB	0 GB	0 GB	300 GB	0	0 GB	0
/dev/cstardb:1				enabled		302 GB	2 GB	0 GB	0 GB	0 GB	300 GB	0	0 GB	0
/dev/cstardc:2				enabled		302 GB	2 GB	0 GB	0 GB	0 GB	300 GB	0	0 GB	0

/dev/cstardd: 3				enabled		302 GB	2 GB	0 GB	0 GB	0 GB	300 GB	0	0 GB	0
--------------------	--	--	--	---------	--	--------	------	------	------	------	--------	---	------	---

c001n02

Index	Serial #	Revision #	Model	Status	Mode	Total Capacity	Sys Resources	Used Capacity	Audit MetaData	Protected User Data	Free Capacity	System Buffer	Available Capacity	Used Obj Count
/dev/cstardd: 0				enabled		302 GB	2 GB	0 GB	0 GB	0 GB	300 GB	0	0 GB	0
/dev/cstardb: 1				enabled		302 GB	2 GB	0 GB	0 GB	0 GB	300 GB	0	0 GB	0
/dev/cstardc:2				enabled		302 GB	2 GB	0 GB	0 GB	0 GB	300 GB	0	0 GB	0
/dev/cstardd: 3				enabled		302 GB	2 GB	0 GB	0 GB	0 GB	300 GB	0	0 GB	0

c001n03

Index	Serial #	Revision #	Model	Status	Mode	Total Capacity	Sys Resources	Used Capacity	Audit MetaData	Protected User Data	Free Capacity	System Buffer	Available Capacity	Used Obj Count
/dev/cstardd: 0				enabled		302 GB	2 GB	0 GB	0 GB	0 GB	300 GB	0	0 GB	0
/dev/cstardb: 1				enabled		302 GB	2 GB	0 GB	0 GB	0 GB	300 GB	0	0 GB	0
/dev/cstardc:2				enabled		302 GB	2 GB	0 GB	0 GB	0 GB	300 GB	0	0 GB	0
/dev/cstardd: 3				enabled		302 GB	2 GB	0 GB	0 GB	0 GB	300 GB	0	0 GB	0

c001n04

Index	Serial #	Revision #	Model	Status	Mode	Total Capacity	Sys Resources	Used Capacity	Audit MetaData	Protected User Data	Free Capacity	System Buffer	Available Capacity	Used Obj Count
/dev/cstardd: 0				enabled		302 GB	2 GB	0 GB	0 GB	0 GB	300 GB	0	0 GB	0
/dev/cstardb: 1				enabled		302 GB	2 GB	0 GB	0 GB	0 GB	300 GB	0	0 GB	0
/dev/cstardc:2				enabled		302 GB	2 GB	0 GB	0 GB	0 GB	300 GB	0	0 GB	0
/dev/cstardd: 3				enabled		302 GB	2 GB	0 GB	0 GB	0 GB	300 GB	0	0 GB	0

c001n05

Index	Serial #	Revision #	Model	Status	Mode	Total Capacity	Sys Resources	Used Capacity	Audit MetaData	Protected User Data	Free Capacity	System Buffer	Available Capacity	Used Obj Count
/dev/cstarda:0				enabled		302 GB	2 GB	0 GB	0 GB	-1	300 GB	32319209472	270 GB	0
/dev/cstardb:1				enabled		302 GB	2 GB	0 GB	0 GB	0 GB	300 GB	2097152	300 GB	1
/dev/cstardc:2				enabled		302 GB	2 GB	0 GB	0 GB	-1	300 GB	32214351872	270 GB	0
/dev/cstardd:3				enabled		302 GB	2 GB	0 GB	0 GB	0 GB	300 GB	32214351872	270 GB	0

c001n06

Index	Serial #	Revision #	Model	Status	Mode	Total Capacity	Sys Resources	Used Capacity	Audit MetaData	Protected User Data	Free Capacity	System Buffer	Available Capacity	Used Obj Count
/dev/hda:0				enabled		302 GB	2 GB	0 GB	0 GB	-1	300 GB	47351595008	256 GB	1
/dev/hdb:1				enabled		302 GB	2 GB	0 GB	0 GB	0 GB	300 GB	2097152	300 GB	0
/dev/hdc:2				enabled		302 GB	2 GB	0 GB	0 GB	-1	300 GB	32214351872	270 GB	0
/dev/hdd:3				enabled		302 GB	2 GB	0 GB	0 GB	0 GB	300 GB	32214351872	270 GB	0

c001n07

Index	Serial #	Revision #	Model	Status	Mode	Total Capacity	Sys Resources	Used Capacity	Audit MetaData	Protected User Data	Free Capacity	System Buffer	Available Capacity	Used Obj Count
/dev/cstarda:0				enabled		302 GB	2 GB	0 GB	0 GB	-1	300 GB	32319209472	270 GB	0
/dev/cstardb:1				enabled		302 GB	2 GB	0 GB	0 GB	0 GB	300 GB	2097152	300 GB	0
/dev/cstardc:2				enabled		302 GB	2 GB	0 GB	0 GB	-1	300 GB	32214351872	270 GB	0
/dev/cstardd:3				enabled		302 GB	2 GB	0 GB	0 GB	0 GB	300 GB	32214351872	270 GB	0

c001n08

Index	Serial #	Revision #	Model	Status	Mode	Total Capacity	Sys Resources	Used Capacity	Audit MetaData	Protected User Data	Free Capacity	System Buffer	Available Capacity	Used Obj Count
-------	----------	------------	-------	--------	------	----------------	---------------	---------------	----------------	---------------------	---------------	---------------	--------------------	----------------

/dev/cstardc:0				enabled		302 GB	2 GB	0 GB	0 GB	-1	300 GB	32319209472	270 GB	1
/dev/cstardb:1				enabled		302 GB	2 GB	0 GB	0 GB	0 GB	300 GB	2097152	300 GB	1
/dev/cstardc:2				enabled		302 GB	2 GB	0 GB	0 GB	-1	300 GB	32214351872	270 GB	1
/dev/cstardd:3				enabled		302 GB	2 GB	0 GB	0 GB	0 GB	300 GB	32214351872	270 GB	0

c001n09

Index	Serial #	Revision #	Model	Status	Mode	Total Capacity	Sys Resources	Used Capacity	Audit MetaData	Protected User Data	Free Capacity	System Buffer	Available Capacity	Used Obj Count
/dev/cstardc:0				enabled		302 GB	2 GB	0 GB	0 GB	-1	300 GB	32319209472	270 GB	1
/dev/cstardb:1				enabled		302 GB	2 GB	0 GB	0 GB	0 GB	300 GB	2097152	300 GB	0
/dev/cstardc:2				enabled		302 GB	2 GB	0 GB	0 GB	-1	300 GB	32214351872	270 GB	0
/dev/cstardd:3				enabled		302 GB	2 GB	0 GB	0 GB	0 GB	300 GB	32214351872	270 GB	0

c001n10

Index	Serial #	Revision #	Model	Status	Mode	Total Capacity	Sys Resources	Used Capacity	Audit MetaData	Protected User Data	Free Capacity	System Buffer	Available Capacity	Used Obj Count
/dev/cstardc:0				enabled		302 GB	2 GB	0 GB	0 GB	-1	300 GB	32319209472	270 GB	0
/dev/cstardb:1				enabled		302 GB	2 GB	0 GB	0 GB	0 GB	300 GB	2097152	300 GB	0
/dev/cstardc:2				enabled		302 GB	2 GB	0 GB	0 GB	-1	300 GB	32214351872	270 GB	0
/dev/cstardd:3				enabled		302 GB	2 GB	0 GB	0 GB	0 GB	300 GB	32214351872	270 GB	0

c001n11

Index	Serial #	Revision #	Model	Status	Mode	Total Capacity	Sys Resources	Used Capacity	Audit MetaData	Protected User Data	Free Capacity	System Buffer	Available Capacity	Used Obj Count
/dev/cstardc:0				enabled		302 GB	2 GB	0 GB	0 GB	-1	300 GB	32319209472	270 GB	1

/dev/cstardb:1				enabled		302 GB	2 GB	0 GB	0 GB	0 GB	300 GB	2097152	300 GB	1
/dev/cstardc:2				enabled		302 GB	2 GB	0 GB	0 GB	-1	300 GB	32214351872	270 GB	0
/dev/cstardd:3				enabled		302 GB	2 GB	0 GB	0 GB	0 GB	300 GB	32214351872	270 GB	0

c001n12

Index	Serial #	Revision #	Model	Status	Mode	Total Capacity	Sys Resources	Used Capacity	Audit MetaData	Protected User Data	Free Capacity	System Buffer	Available Capacity	Used Obj Count
/dev/cstarda:0				enabled		302 GB	2 GB	0 GB	0 GB	-1	300 GB	32319209472	270 GB	1
/dev/cstardb:1				enabled		302 GB	2 GB	0 GB	0 GB	0 GB	300 GB	2097152	300 GB	0
/dev/cstardc:2				enabled		302 GB	2 GB	0 GB	0 GB	-1	300 GB	32214351872	270 GB	1
/dev/cstardd:3				enabled		302 GB	2 GB	0 GB	0 GB	0 GB	300 GB	32214351872	270 GB	0

c001n13

Index	Serial #	Revision #	Model	Status	Mode	Total Capacity	Sys Resources	Used Capacity	Audit MetaData	Protected User Data	Free Capacity	System Buffer	Available Capacity	Used Obj Count
/dev/cstarda:0				enabled		302 GB	2 GB	0 GB	0 GB	-1	300 GB	32319209472	270 GB	1
/dev/cstardb:1				enabled		302 GB	2 GB	0 GB	0 GB	0 GB	300 GB	2097152	300 GB	1
/dev/cstardc:2				enabled		302 GB	2 GB	0 GB	0 GB	-1	300 GB	32214351872	270 GB	1
/dev/cstardd:3				enabled		302 GB	2 GB	0 GB	0 GB	0 GB	300 GB	32214351872	270 GB	1

c001n14

Index	Serial #	Revision #	Model	Status	Mode	Total Capacity	Sys Resources	Used Capacity	Audit MetaData	Protected User Data	Free Capacity	System Buffer	Available Capacity	Used Obj Count
/dev/cstarda:0				enabled		302 GB	2 GB	0 GB	0 GB	-1	300 GB	32319209472	270 GB	1
/dev/cstardb:1				enabled		302 GB	2 GB	0 GB	0 GB	0 GB	300 GB	2097152	300 GB	0

/dev/cstardc:2				enabled		302 GB	2 GB	0 GB	0 GB	-1	300 GB	32214351872	270 GB	1
/dev/cstardd:3				enabled		302 GB	2 GB	0 GB	0 GB	0 GB	300 GB	32214351872	270 GB	1

c001n15

Index	Serial #	Revision #	Model	Status	Mode	Total Capacity	Sys Resources	Used Capacity	Audit MetaData	Protected User Data	Free Capacity	System Buffer	Available Capacity	Used Obj Count
/dev/cstarda:0				enabled		302 GB	2 GB	0 GB	0 GB	-1	300 GB	32319209472	270 GB	1
/dev/cstardb:1				enabled		302 GB	2 GB	0 GB	0 GB	0 GB	300 GB	2097152	300 GB	0
/dev/cstardc:2				enabled		302 GB	2 GB	0 GB	0 GB	-1	300 GB	32214351872	270 GB	0
/dev/cstardd:3				enabled		302 GB	2 GB	0 GB	0 GB	0 GB	300 GB	32214351872	270 GB	0

c001n16

Index	Serial #	Revision #	Model	Status	Mode	Total Capacity	Sys Resources	Used Capacity	Audit MetaData	Protected User Data	Free Capacity	System Buffer	Available Capacity	Used Obj Count
/dev/hda:0				enabled		302 GB	2 GB	0 GB	0 GB	-1	300 GB	47351595008	256 GB	0
/dev/hdb:1				enabled		302 GB	2 GB	0 GB	0 GB	0 GB	300 GB	2097152	300 GB	0
/dev/hdc:2				enabled		302 GB	2 GB	0 GB	0 GB	-1	300 GB	32214351872	270 GB	0
/dev/hdd:3				enabled		302 GB	2 GB	0 GB	0 GB	0 GB	300 GB	32214351872	270 GB	0

Cluster Pool List

Cluster Capacity

Total Capacity: 19308 GB

System Resources: 106 GB

Offline Capacity: 0 GB
Used Capacity: 2 GB
Audit MetaData: 2 GB
Protected User Data: 0 GB
Spare Capacity: 4800 GB
Free Capacity: 14399 GB
Available Capacity: 12686 GB
Avail Cap Until Alert: 13621657038848
Used Object Count: 18
Total Object Count: 360000000
Free Object Count: 359999982
System Buffer: 1191064895488
Regeneration Buffer: 647856340992

16 Appendix D – Example Email error reports

16.1 Replication paused due to incorrect configuration of profile rights



Centera Alert Report

Version 1.1



Alert Identification

Type: Software

Symptom Code: 1.1.1.1.20.03

Format Version: 1.1

Format DTD: alert_email-1.1.dtd

Creation Date and Time: 08-03-2005 01:36:54.216 UTC



Cluster Identification

Name: Centera2

Serial Number: APM00040306018

Active Software Version: 3.0.0-542-557-8615

Cluster ID: b11ba322-1dd1-11b2-a069-ebb546c38de2

Cluster Domain: citigroup.com



Alert Description

Alert Level: ERROR

Alert Description: \$Application.Main.ReplicationManager.Main.ReplicationSubstate@\$CLUSTER:c001n04:7987cdd8-1dd2-11b2-8533-ccc44fa4b380:paused_no_capability

Component: node

Component ID: c001n04

Sub-Component: sensor

Sub-Component ID: \$Application.Main.ReplicationManager.Main.ReplicationSubstate@\$CLUSTER

16.2 Disk failure on node

Centera Alert Report

Version 1.1

Alert Identification

Type: Hardware

Symptom Code: 3.1.5.1.01.01

Format Version: 1.1

Format DTD: alert_email-1.1.dtd

Creation Date and Time: 08-11-2005 09:48:35.069 UTC

Cluster Identification

Name: Centera2

Serial Number: APM00040306018

Active Software Version: 3.0.0-542-557-8615

Cluster ID: b11ba322-1dd1-11b2-a069-ebb546c38de2

Cluster Domain: citigroup.com

Alert Description

Alert Level: WARNING

Alert Description: disk

Component: node

Component ID: c001n16

Sub-Component: disk

Sub-Component ID: /dev/cstarda:0

16.3 Node alert

Centera Alert Report

Version 1.1

Alert Identification

Type: Hardware

Symptom Code: 3.1.6.1.01.01

Format Version: 1.1

Format DTD: alert_email-1.1.dtd

Creation Date and Time: 08-15-2005 10:47:02.726 UTC



Cluster Identification

Name: Centera2

Serial Number: APM00040306018

Active Software Version: 3.0.0-542-557-8615

Cluster ID: b11ba322-1dd1-11b2-a069-ebb546c38de2

Cluster Domain: citigroup.com



Alert Description

Alert Level: WARNING

Alert Description: node

Component: cabinet

Component ID: 1

Sub-Component: node

Sub-Component ID: c001n16

17 Appendix E – CentraStar 3.0 Alerts

The following table shows the full list of Centera alerts, a brief description of the problem reported and a recommendation for resolution.

- Symptom Code uniquely determines an alert event (the first 5 numbers of the symptom code actually identify the sensor triggering the alert).
- Threshold indicates the value at which the Evaluator will issue an alert event
- Severity indicates the severity level of the problem found:
 1. OK (used for improvements only),
 2. Notification,
 3. Warning,
 4. Error, and
 5. Critical.

The resolution discusses some recommended steps can be taken to identify and resolve the problem.

This is normal test

NOTE: If the resolution indicates “EMC Support has been notified of the problem and a case will automatically be opened”, this assumes ConnectEMC is enabled and the email alert was successfully sent to EMC.

Symptom Code	Sensor Name	Threshold (Severity)	Problem Description	Resolution
1.1.1.1.01.01 1.1.1.1.01.02 1.1.1.1.01.03	ParkedReplicationsClipCount	> 10000 (Warning) > 25000 (Error) > 50000 (Critical)	The replication parking on one or more of the nodes in the cluster contains more than 10000/25000/50000 entries. This is typically due to connectivity issues between this cluster and the replica, or authorization and capability issues on the replica. In order to prevent the replication parking to increase, CentraStar has paused the replication process.	EMC Support has been notified of the problem and a case will automatically be opened. Check the replication parking using the CLI or CV to identify the root cause of the problem.
1.1.1.1.05.01 1.1.1.1.05.02 1.1.1.1.05.03	ParkedRestoreClipCount	> 10000 (Warning) > 25000 (Error) > 50000 (Critical)	The restore parking on one or more of the nodes in the cluster contains more than 10000/25000/50000 entries. This is typically due to connectivity issues between this cluster and the replica, or authorization and capability issues on the replica. In order to prevent the restore parking to increase, CentraStar has paused the restore process.	EMC Support has been notified of the problem and a case will automatically be opened. Check the replication parking using the CLI or CV to identify the root cause of the problem.

Symptom Code	Sensor Name	Threshold (Severity)	Problem Description	Resolution
1.1.1.1.03.01 1.1.1.1.03.02 1.1.1.1.03.03	ReplicationETA	>= 24 h (Warning) >= 48 h (Error) >= 120 h (Critical)	The replication process on the cluster will probably take longer than 24/48/120 hours to process the entire replication queue, taking into account write and replication rates over the past 24 hours. This might indicate a problem with replication or a write activity that is higher than average.	EMC Support has been notified of the problem and a case will automatically be opened.
1.1.1.1.20.01	ReplicationSubstate	paused_clusterfull (Error)	CentraStar has paused the replication process on the cluster because the replica cluster has not enough capacity to service the replication requests.	Check the capacity of the target cluster.
1.1.1.1.20.02	ReplicationSubstate	paused_parking_overflow (Error)	The replication parking on this cluster contains more than 10,000 entries on one or more nodes. This is typically due to connectivity issues between this cluster and the replica, or authorization and capability issues on the replica.	Check the replication parking using the CLI or CV to identify the root cause of the problem.
1.1.1.1.20.03	ReplicationSubstate	paused_no_capability (Error)	The replica cluster has refused a replication request because the access profile used for replication does not have the correct capabilities.	Check if the write capability is enabled on the target cluster.

Symptom Code	Sensor Name	Threshold (Severity)	Problem Description	Resolution
1.1.1.1.20.04	ReplicationSubstate	paused_authentication_failure (Error)	The replica cluster has refused a replication request because the access profile used for replication could not be authenticated.	Check the access profile on the target cluster.
1.1.2.1.02.01	ComplianceModeRevert	Average <> Maximum (Critical)	One or more nodes in the cluster do not share the cluster compliance settings.	EMC Support has been notified of the problem and a case will automatically be opened
2.1.2.1.01.01	ClusterClipIntegrity	INCOMPLETE (Error)	CentraStar has determined that one or more C-Clips have become unavailable due to a combination of disks and/or nodes being offline.	EMC Support has been notified of the problem and a case will automatically be opened.
4.1.1.1.01.01 4.1.1.1.01.02	CPUTemperature	> 66 °C (Warning) > 86 °C (Error)	CentraStar has determined that one or more nodes in the cluster have a running temperature higher than 66 °C/86 °C. This will reduce the performance of the affected nodes and if the situation persists, the node will go offline.	EMC Support has been notified of the problem and a case will automatically be opened.
4.1.1.1.02.01	VarPartitionUsedSpace	> 88% (Error)	CentraStar has determined that one or more nodes in the cluster with an internal partition used by CentraStar have limited capacity.	EMC Support has been notified of the problem and a case will automatically be opened.

Symptom Code	Sensor Name	Threshold (Severity)	Problem Description	Resolution
5.2.2.1.03.01 5.2.2.1.03.02 5.2.2.1.03.03	AvailableCapacityHardStop Percent	< 20% (Warning) < 10% (Error) <= 5 % (Critical)	This sensor will be active when no Regeneration Buffer has been defined. The alerts are issued when available capacity on at least one cube in the cluster is less than 20%/10%/5% of the total raw capacity. As there is no regeneration buffer defined, it will soon not be possible to write to the affected cube(s) anymore and disk and node regenerations will not be able to complete.	Contact your EMC sales representative to order more capacity. For the Error and Critical levels, EMC Support will be notified of the problem and a case will automatically be opened.
5.2.2.1.04.01 5.2.2.1.04.02 5.2.2.1.04.03	FreeObjectPercent	<20 % (Warning) <10 % (Error) =0 % (Critical)	The free object count on the cluster is less than 20/10/0% of the allowed object count. It will soon not be possible/is no longer possible to write to the cluster anymore. The maximum object count in the 3.0 release of CentraStar is 30M per node.	Contact your EMC sales representative to order more capacity. For the Error and Critical levels, EMC Support will be notified of the problem and a case will automatically be opened.
5.2.2.1.01.01 5.2.2.1.01.02 5.2.2.1.01.03 5.2.2.1.01.04	RegenerationBufferAlert Percent	< 125 % (Notification) < 100 % (Warning) < 50 % (Error) < 25% (Critical)	This sensor will be active when a Regeneration Buffer has been defined with the Alert Only policy. The available capacity on at least one cube in the cluster is 25% higher than/equal to/50% of/25% of the configured regeneration buffer. As the policy is set to Alert Only it is still possible to write to the affected cube(s).	Contact your EMC sales representative to order more capacity. For the Error and Critical levels, EMC Support will be notified of the problem and a case will automatically be opened.

Symptom Code	Sensor Name	Threshold (Severity)	Problem Description	Resolution
5.2.2.1.02.01 5.2.2.1.02.02 5.2.2.1.02.03	RegenerationBufferHardStopPercent	< 150% (Warning) < 125% (Error) <= 100% (Critical)	This sensor will be active when a Regeneration Buffer has been defined with the Hard Stop policy. The available capacity on at least one cube in the cluster is 50%/25% higher than/less than the configured regeneration buffer. As the policy is set to Hard Stop it will soon/is not possible to write to the affected cube(s) anymore. Disk and node regenerations will still be able to use the capacity reserved in the regeneration buffer.	Contact your EMC sales representative to order more capacity. For the Error and Critical levels, EMC Support will be notified of the problem and a case will automatically be opened.
3.1.3.1.02.01	ExternalNICFailureCount	> 1 (Error)	CentraStar has determined that one or more NICs used for data access by your application are not functioning correctly. This will reduce bandwidth available for data access by your application.	EMC Support has been notified of the problem and a case will automatically be opened.
3.1.3.1.01.01 3.1.3.1.01.02	InternalNICFailureCount	1 (Warning) >1 (Error)	CentraStar has determined that one or more NICs used by the internal Centera network are not functioning correctly.	EMC Support has been notified of the problem and a case will automatically be opened.
3.1.4.1.01.01 3.1.4.1.01.02	CubeSwitchFailureCount	= 1 (Error) > 1 (Critical)	CentraStar has determined that one or more cube switches in the cluster are not functioning correctly. This will reduce the internal bandwidth of the cluster and represents a single point of failure until addressed.	EMC Support has been notified of the problem and a case will automatically be opened.

Symptom Code	Sensor Name	Threshold (Severity)	Problem Description	Resolution
3.1.5.1.01.01	DiskOffline	1 (Warning)	There are one or more disks offline in the cluster.	EMC Support has been notified of the problem and a case will automatically be opened.
3.1.6.01.01.01	NodesOffline	1 (Warning)	There are one or more nodes offline in the cluster.	EMC Support has been notified of the problem and a case will automatically be opened.
3.1.4.1.02.01	RootSwitchFailures	1 (Warning)	There is one root switch in the cluster not functioning correctly.	EMC Support has been notified of the problem and a case will automatically be opened.
1.1.3.1.01.01 1.1.3.1.01.02 1.1.3.1.01.03	PoolReachingQuota	>80% (Warning) >90% (Error) > 100% (Critical)	This alert indicates that one or more pools have almost used up their capacity quota. The pool capacity is virtual capacity and is not related to the physical capacity of a Centera. Reaching a pool quota generates an alert but it does not stop an application from writing to the affected pool(s).	EMC Customer Support Center has been notified of the problem but a case will not be opened.
2.1.3.1.01.01 ¹	NoRedundantPower	> 1 (Warning)	This alert indicates that one or more nodes on a cluster are not receiving redundant power.	EMC Customer Support Center has been notified of the problem but a case will not be opened.
n/a	SDK Alert	No connection	The SDK is unable to establish a connection with Centera.	Please verify the network connection between the application server and Centera and the availability of the Centera.

Figure 12 Centera 3.0 Alerts

¹ This alert is not part of CentraStar 3.0 – it will be part of the 3.0 SP1 service pack.

Draft

18 Appendix F – Centera MIB

```
-- Centera.MIB: EMC BDG Centera Manager MIB, version 2.0.1
-- Copyright © 2003 EMC Corporation. All rights reserved.
--
-- The definition of this MIB derives from RFC 1902.

CENTERA-MIB DEFINITIONS ::= BEGIN

IMPORTS
    enterprises, Counter32, OBJECT-TYPE,
    MODULE-IDENTITY, NOTIFICATION-TYPE
    FROM SNMPv2-SMI;

centera MODULE-IDENTITY
    LAST-UPDATED "200303121252Z"
    ORGANIZATION "EMC BDG"
    CONTACT-INFO
        " Postal: EMC BDG
          Schaliënhoeverdreef 20 i
          B-2800 Mechelen
          Belgium"
    DESCRIPTION
        "The EMC BDG Centera Manager MIB module"
    ::= { emc 5 }

emc    OBJECT IDENTIFIER ::= { enterprises 1139 }

-- Top level groups

notificationTrap    OBJECT IDENTIFIER ::= { centera 0 }
notificationData    OBJECT IDENTIFIER ::= { centera 2 }

-- The notificationData group
-- The members of this group are the OIDs for VarBinds that contain
-- notification data.

genericNotify    OBJECT IDENTIFIER ::= { notificationData 1 }

notifyTimestamp OBJECT-TYPE
    SYNTAX Counter32
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION
        "The timestamp of the notification."
    ::= { genericNotify 1 }

notifyServer OBJECT-TYPE
    SYNTAX OCTET STRING
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION
        "The cluster ID."
    ::= { genericNotify 2 }

notifyDescription OBJECT-TYPE
    SYNTAX OCTET STRING
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION
        "A complete description of the event."
    ::= { genericNotify 3 }

notifySeverity OBJECT-TYPE
    SYNTAX INTEGER {
        informational (1),
        warning (2),
        severe (3)
    }
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION
        "The severity level of the event is indicated by an integer number in the
        range from 1 to 3."
    ::= { genericNotify 4 }
```

-- The Centera SNMP traps
-- The definition of these objects mimics the SNMPv2 convention for
-- sending traps. The enterprise OID gets appended with a 0
-- and then with the specific trap code.

```
trapNotification NOTIFICATION-TYPE
    OBJECTS {
        notifyTimestamp,
        notifyServer,
        notifyDescription,
        notifySeverity
    }
    STATUS current
    DESCRIPTION
    "This trap describes the state of the Centera cluster."
    ::= { notificationTrap 1 }

trapAlarmNotification NOTIFICATION-TYPE
    OBJECTS {
        notifyTimestamp,
        notifyServer,
        notifyDescription,
        notifySeverity
    }
    STATUS current
    DESCRIPTION
    "This trap identifies a problem on the Centera cluster.
     The description can be used to describe the nature of the change."
    ::= { notificationTrap 2 }
```

END