

EMC Centera Universal Access Data Protection and Disaster Recovery Best Practices

This document describes the models for using EMC Centera Universal Access (CUA) with a customer application and Centera for the purpose of Enhanced Availability and disaster recovery. It also includes guidelines for failover testing in a pre-production customer environment.

Published December 2005

Copyright © 2005, 2006 EMC Corporation. All rights reserved.

EMC believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED “AS IS.” EMC CORPORATION MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Use, copying, and distribution of any EMC software described in this publication requires an applicable software license.

EMC², EMC, EMC ControlCenter, AlphaStor, ApplicationXtender, Catalog Solution, Celerra, CentraStar, CLARAlert, CLARiiON, ClientPak, Connectrix, Co-StandbyServer, Dantz, Direct Matrix Architecture, DiskXtender, Documentum, EmailXtender, EmailXtract, HighRoad, Legato, Legato NetWorker, Navisphere, OpenScale, PowerPath, RepliStor, ResourcePak, Retrospect, Smarts, SnapShotServer, SnapView/IP, SRDF, Symmetrix, TimeFinder, VisualSAN, VSAM Assist, Xtender, Xtender Solutions, and where information lives are registered trademarks and EMC Developers Program, EMC OnCourse, EMC Proven, EMC Snap, EMC Storage Administrator, Access Logix, ArchiveXtender, Authentic Problems, Automated Resource Manager, AutoStart, AutoSwap, AVALONidm, C-Clip, Celerra Replicator, Centera, CLARevent, Codebook Correlation Technology, Common Information Model, CopyCross, CopyPoint, DatabaseXtender, Direct Matrix, DiskXtender 2000, EDM, E-Lab, EmailXaminer, Enginuity, eRoom, FarPoint, FLARE, Global File Virtualization, Graphic Visualization, InfoMover, Invista, MirrorView, NetWin, NetWorker, OnAlert, Powerlink, PowerSnap, Rainfinity, RecoverPoint, RepliCare, SafeLine, SAN Advisor, SAN Copy, SAN Manager, SDMS, SnapImage, SnapSure, SnapView, StorageScope, SupportMate, SymmAPI, SymmEnabler, Symmetrix DMX, UltraPoint, Viewlets, VisualSRM, and WebXtender are trademarks of EMC Corporation. All other trademarks used herein are the property of their respective owners.

Part Number H1587.1

Table of Contents

Introduction	4
Centera Universal Access (CUA) 3.6 Models	4
CUA Data Availability Protection	4
Software RAID	4
CUA writeback	5
Automated CUA database backup	5
CUA read failover	6
Basic CUA failure recovery	6
Recovery from disk failure	7
Recovery from motherboard failure	7
Recovery from a site failure	7
CUA Enhanced Availability	8
Requirements and restrictions	8
Configuration and initialization	8
Monitoring	9
Failover	9
Rules for Enhanced Availability CUA Setups	10
Single-site Enhanced Availability CUA pairs	10
Dual-site Enhanced Availability CUA pairs	11
Redundant Dual-site Enhanced Availability CUA pairs	11
Appendix A: Validating CUA Failover in a Test Environment	12
Introduction	12
Sources	13
Additional steps prior to validation	13
Validating CUA failover in a test environment	13
Next steps	17

Introduction

EMC® Centera™ Universal Access (CUA) provides several levels of data protection; from software RAID that secures content from local disk failures to database synchronization across a customer's wide area network (WAN) to protect content from a total CUA failure. The level of data protection required depends on the customer's requirements and also on the Centera topology deployed in the customer's organization.

This document concentrates on the different deployments of Centera Universal Access to provide varying levels of data protection. While designed to educate Centera Universal Access customers on the Enhanced Availability capabilities of CUA, this document is not intended to replace the standard documentation. For detailed descriptions of basic CUA functions please refer to the Centera Universal Access Monitor and Administration Guide (P/N 069001199 Rev A09).

Topics covered in this document include:

- CUA features relating to data protection and disaster recovery
- Solution topologies
- Solution limitations
- Implementation rules and considerations

The basic CUA functions described in this document are the same across the Centera node and Dell server hardware platforms unless specifically indicated.

Centera Universal Access (CUA) 3.6 Models

As of January, 2006, the current version of Centera Universal Access is version 3.6. Centera Universal Access comes in two model versions:

1. CNRCUAISW (Basic CUA Server Internal to Centera)
2. CNRCUAESW (Basic CUA Server External to Centera)

These models contain safeguards to protect content written to the CUA file system and recovery procedures to rebuild a CUA in the event that the CUA should require repair or replacement. Basic CUA models may also fail over reads to a Replica Centera if the customer has deployed Replicated Centera clusters.

In addition to basic data protection, these models also allow the customer to specify a retention period per directory, rendering the content in that directory non-modifiable and non-erasable for the duration of the retention period. Finally, these models support Enhanced Availability when configured on two CUA server platforms deployed as an Enhanced Availability pair. In this configuration, a CUA server designated as an Active CUA by the customer synchronizes its database with a Standby CUA, either in the same local area network (LAN) or across a wide area network (WAN).

CUA Data Availability Protection

Centera Universal Access is designed to protect content within the local system by minimizing single points of hardware failure and by using Centera as a backup. The following sections describe the means Centera Universal Access uses to protect local content.

Software RAID

CUA uses software RAID to protect content written by an application to the CUA file system. CUA software hosted on the Dell PowerEdge 2650 and 2850 servers and the Centera Generation 3 and 4 node platforms use software RAID to stripe data across the internal drives in the server enclosure. This configuration protects the CUA's system, database, and file system content in case of a single drive failure.

CUA writeback

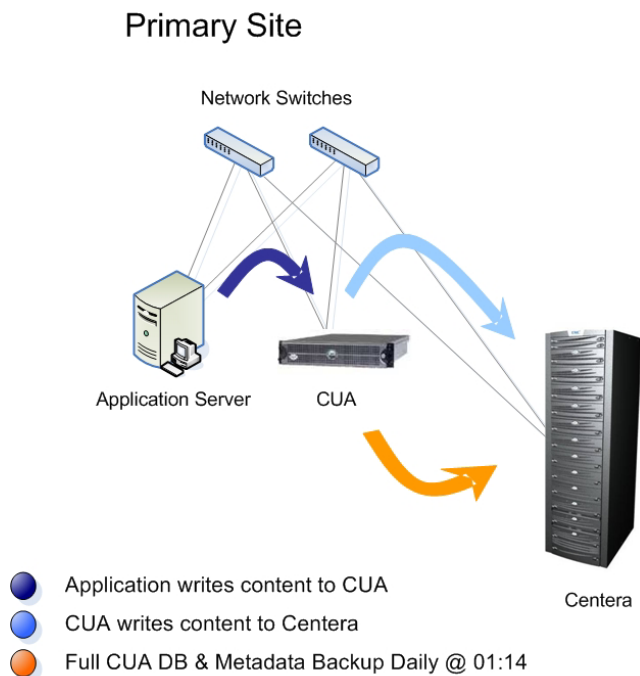
CUA uses a file system and an internal database to manage content in Centera on behalf of non-integrated applications. When an application writes content to CUA, it uses a network file system protocol such as NFS or CIFS, or the Internet protocol FTP. The application will use these same protocols, or possibly HTTP, to retrieve content from CUA. After the content is written to the CUA file system, CUA writes the content to Centera and stores the resulting Centera-generated content address in an internal database.¹ This process is called *writeback*. During writeback CUA uses the Centera API to create a content descriptor file (CDF), which contains all pertinent file system attributes plus information about the content itself. CUA stores the same information in its database.

Automated CUA database backup

The internal CUA database maintains information about the content CUA writes to Centera. Each record in the CUA database contains the same information as the CDF the CUA writes to Centera with the application's content, so there is a one-to-one correlation between the CDF content written by CUA in Centera and the content of CUA's internal database. As a precaution against potential failures, CUA writes an image of its internal database to Centera on a nightly basis and sends the content address of that image to an e-mail address defined by the CUA

administrator. CUA keeps seven days worth of database images in Centera and rotates them out using a first-in, first-out (FIFO) model.

In the event of a CUA failure, the CUA database may be recovered from Centera using a built-in CUA recovery process. This process is described in the later section titled **Single CUA Recovery**.



Single-site CUA configuration

¹ CUA versions 3.5 and later write files to Centera sooner than previous versions. In previous versions of CUA, files could be written to Centera up to 45 minutes after they were written to CUA. CUA 3.5 writes a file to Centera after one minute providing the file has not been modified in that time. The file may still be opened, as the writeback policy is based on the last modification time. During writeback, a retention value of 0 is set on the C-Clip™. If the user updates the file 25 minutes later, the file will be written back again 10 minutes after the last update.

Example The user writes a file at T1, CUA will write it back to Centera at T1+1minute.

1. If the user updates the file at T1+25minutes, CUA will write it back again at T1+35minutes.
2. If the user updates the file at T1+45minutes and at T1+50minutes, CUA will write it back again at T1+60minutes.

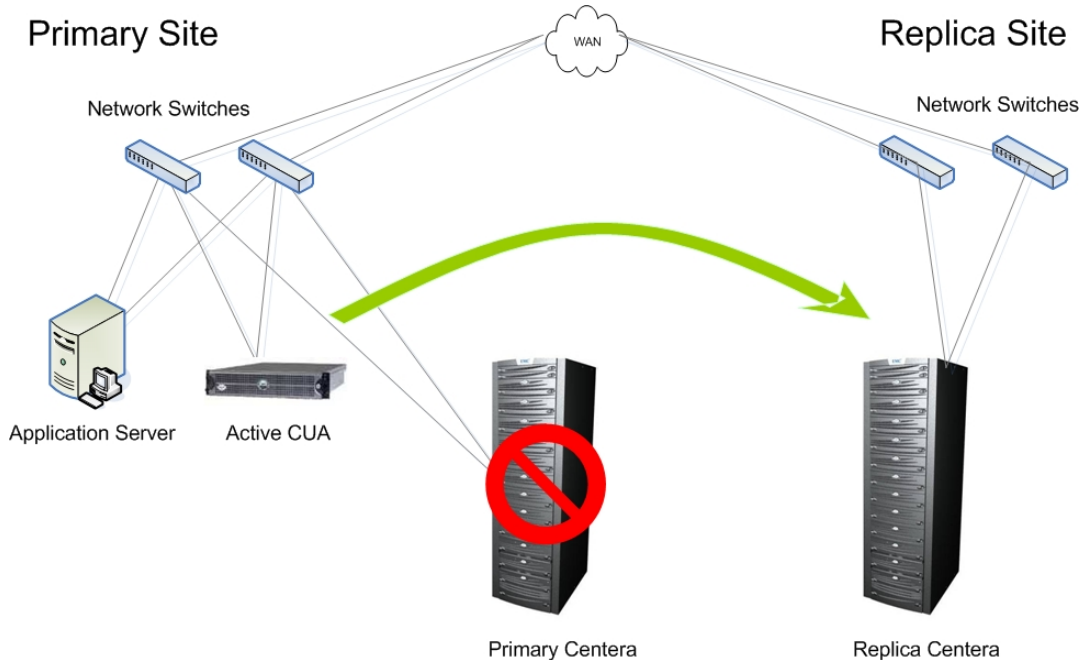
When a retention policy is specified, the CUA updates the C-Clip and sets the retention. Once a file has not been modified for 30 minutes, the C-Clip will be re-written setting the retention value specified by the policy. (The blob itself will not be re-written.)

CUA read failover

Centera Universal Access is a Centera-integrated application and so makes use of the Centera application failover model. Customers can use this capability in a Replicated Centera environment to ensure that an Active CUA continues to deliver content to the customer application even if the Primary Centera is not available.

In a Replicated Centera environment, a Primary Centera Replicates newly acquired content to a Replica Centera, typically located in a different site. Applications connecting to the Primary Centera via the Centera SDK automatically acquire the IP addresses of the Replica Centera. Though the Replica IP addresses are not explicitly revealed to the application, the IP addresses are available for certain Centera SDK functions to fail over to the Replica Centera if the Primary Centera is no longer available.

In the case of the Centera SDK read function, CUA connects to a Primary Centera cluster to deliver content to a customer application. Because the Primary Centera is replicating to a secondary Replica Centera cluster in another location, CUA acquires the IP addresses of the Replica Centera. If the Primary Centera becomes unavailable, read requests by CUA automatically fail over to the Replica, and CUA continues to deliver the requested content to the application. This process is same whether a single CUA or Enhanced Availability pair is used.



- In the event of a lost connection to the Primary Centera, the Active CUA automatically fails over to the Replica Centera to deliver content requested by the Application Server

Replicated Centera environment with a single CUA and read failover

Basic CUA failure recovery

This scenario assumes a single Active CUA that addresses a Centera cluster. The Centera cluster might be replicated, but that possibility is not important in this scenario.

Recovery from disk failure

CUA uses software RAID to protect content against a single disk failure. However, the CUA hardware platforms have other single points of failure. The Dell 2650 and 2850 servers do have redundant processors, but not redundant motherboards. In addition, CUA software recognizes only one of the two Ethernet cards in the Dell 2650 and the Centera Gen3 node. If these components should fail, the hardware must be replaced.

In the event of a failure of a single drive the CUA will continue to function on the remaining drives. If the CUA software is hosted on a Dell server, the customer should contact Dell Customer Service to replace the failed drive. In the U.S., the terms of the Dell Silver Level Service agreement provides for the dispatch of a technician within four hours of determining the hard drive failure. Internationally, Dell service response time varies by region. See www.dell.com for details by region. If the CUA software is hosted on a Centera Gen 3 or Gen 4 node, EMC Customer Service will replace the failed drive.

Recovery from motherboard failure

A motherboard failure on a CUA will disable it from further operation. Any applications that use CUA as a conduit to the Centera cluster will no longer be able to communicate with Centera.

In the case of a single-drive failure, restoring full operations on CUA is achieved by replacing the failed drive. In the case of a more catastrophic failure, such as the loss of the motherboard, the entire CUA server must be replaced. The following steps are needed to perform disaster recovery:

1. A replacement CUA must be installed. It must be a clean install of the same or newer version of CUA software that was on the failed CUA. There cannot be any other files loaded onto the replacement CUA. Ideally, the replacement should be an identical hardware configuration as the failed CUA.
2. A restore is executed on the replacement CUA against the Centera cluster using the content address from the latest nightly backup performed by the failed CUA. This will restore the files and content addresses from the Centera cluster.
3. Upon completion of this restore operation, CUA will perform a time-based query against the Centera cluster to restore any files written to Centera since the last backup².

Upon completion of the full restore process, applications may now begin to access Centera via the newly available CUA.

In some failure conditions, e.g., motherboard failure, the drives from a failed CUA server may be swapped into a new server. In this case the restore process is not required, though a new CUA license may be required for older CUA versions.

Recovery from a site failure

A customer concerned with comprehensive data protection should implement Centera replication to safeguard Centera content from loss in the event of a catastrophic failure at the Primary site. In the event of a catastrophic failure at the Primary site the following steps are needed to perform disaster recovery:

1. The replacement CUA must be a clean install of the same or newer version of CUA software that was running at the Primary site. There cannot be any other files loaded onto the replacement CUA. Ideally, the replacement CUA should be an identical hardware configuration to the Primary.

² As of CentraStar® version 2.4, replication of any object from the Primary Centera to its Replica is delayed by 10 minutes. This delay will result in fewer files available to CUA as the result of the CUA's query.

2. A restore is executed on the replacement CUA against the Replica Centera using the content address from the latest nightly backup performed by the Active CUA. This will restore the files and content addresses from the Replica Centera that have been replicated from the Primary Centera. (Keep in mind that files stored on the Primary Centera and not yet replicated will be missing on the Replica Centera.)
3. Upon completion of this restore operation, the replacement CUA will perform a time-based query against the Replica Centera to restore any files written to the Primary Centera and acquired by the Replica since the last backup.

CUA Enhanced Availability

Centera Universal Access uses an Enhanced Availability model to improve the availability model between CUA and Centera. The term Enhanced Availability (EA) is used to indicate an Active/Passive (Standby) CUA pair, versus an Active/Active pair typically associated with high availability (HA). Additionally, in EA operator intervention—in the form of a CLI command—is required to initiate the EA failover process.

In a CUA EA configuration, upon loss of heartbeat between the Active and Standby CUAs, an alert is sent to the administrator so that he or she may activate the failover process.

The Centera Universal Access Enhanced Availability configuration employs a second Standby CUA that is dedicated to tracking the progress of the Active CUA. In the event of a complete failure of the Active CUA, availability can be restored by transitioning the Standby CUA to the Active state. This is accomplished by entering a command via the management CLI (CUA setup utility). When the Active CUA is installed, its operation mode is designated as Active. A second CUA may be added at any point in the future to assume the role of Standby.

Immediately after configuration, the Standby CUA is in the initializing state until it has synchronized with the Active CUA. After synchronization is complete, the Standby CUA enters the ready state. A transition of Standby CUA in the ready state to the Active state will immediately provide fully qualified name access to all objects stored by the previously Active CUA.

Requirements and restrictions

- The Active and Standby CUAs must be the same hardware platform, that is, either a Dell 2650 or 2850s, or Centera Gen 3 or Gen 4 nodes. This is to maintain consistency of the disk drive formats that host the internal CUA databases.
- The Active and Standby CUAs must point to the same Centera. If this is not done, then the two systems will not be able to access the CAs that the other has stored.
- The Active and Standby CUAs are using Gigabit Ethernet LAN access. This is required because of the high volume of data traffic. All traffic, ingestion, writeback, response to read requests, retrieval from Centera, and the heartbeat are using the same connection.
- The Standby CUA is not available to the application server for read/write access while in Standby mode.
- Files on the Active CUA that are not yet written to Centera are not available on the Standby. If the files have not been written back to Centera they only reside on the Active CUA's disks.

Configuration and initialization

Configuration will be handled via a CLI (command line interface), similar to the Centera CLI. The installation engineer will enter the information provided on the pre-site qualification form. This data includes:

- Hostname (Active and Standby must each have a unique hostname.)

- IP Information (Active and Standby must each have a unique IP address.)
- Centera Information (identical for Active and Standby)
- E-mail Configuration (identical for Active and Standby)

In addition to the configuration settings listed above, there is a menu item called “Configure Enhanced Availability.” This menu requires the following information:

- Name of the Active CUA
- IP of the Active CUA
- Name of the Standby CUA
- IP of the Standby CUA

Once the CUAs have been configured, they now need to be initialized. The Active CUA is running with NFS mount points and CIFS shares usable in read/write configurations. When the Active CUA is initially configured, it starts a listener process that enables the Standby CUA to communicate with the Active CUA. When the Standby CUA is initially configured, it performs the steps listed below:

1. Turns off mount points of NFS and CIFS shares.
2. Sets its state to “initializing.”
3. Locates the Active CUA’s database backup in the object creation log and performs a full restore.
4. Starts the Standby Centera Universal Access daemon, and starts tracking the Active CUA object creation log.

At this point, the Standby CUA is synchronized with the Active CUA and it sets its state to “ready.” The Standby continues to track the object creation log, and updates its local XMD database.

Monitoring

A Web-based GUI shows name, IP, and state information about the two CUAs. This information includes the names of the Active and Standby CUAs and an indication of the state of the Standby CUA.

Failover

Failover is a manual operation initiated by the administrator locally or remotely over an IP connection. On the Standby CUA the administrator runs the management CLI and chooses the option “Initiate CUA failover.” If the previously Active CUA is up and can be contacted, an exception is raised and failover does not occur. Otherwise, this CUA will become the Active node and the NFS mount points and CIFS shares become available. CIFS and NFS clients need to manually re-establish connectivity with this CUA, as it will have a different IP address and hostname.

During the failover process the new Active CUA may be configured with the IP address and hostname of the old Active CUA. Any clients that were connected to the old Active CUA should be rebooted to access the new Active CUA.

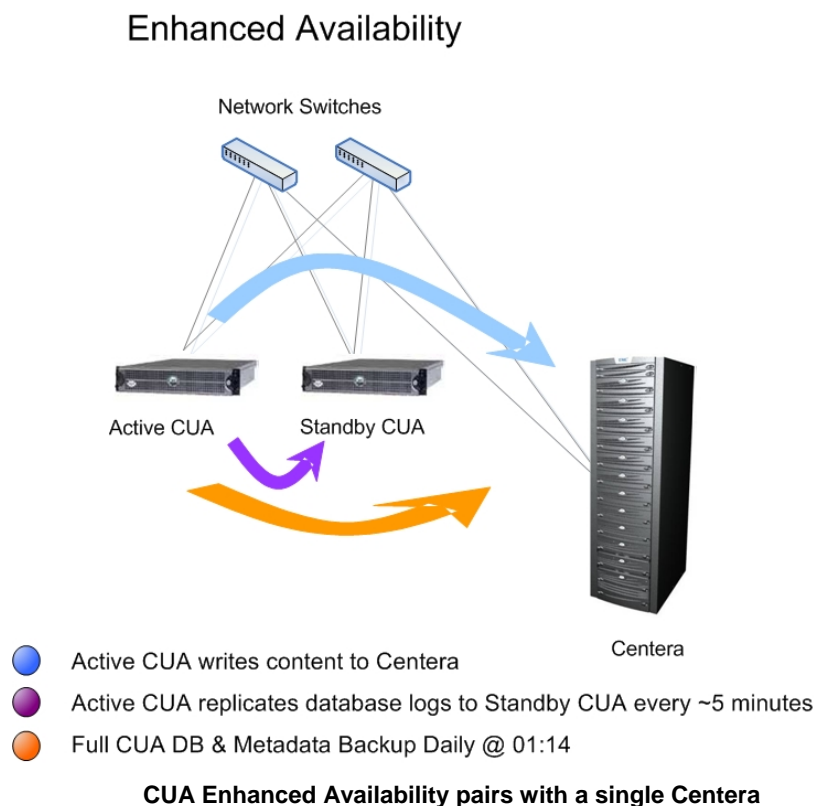
A time-based query will be performed against the Centera to pick up any files that were committed since the last update from the old Active CUA. Under this methodology, the only files that won’t be available will be ones that were written to the old Active CUA, but not yet written back to the Centera. Those files may be found in the “Lost and Found” directory on the old Active CUA after failover is complete.

Rules for Enhanced Availability CUA Setups

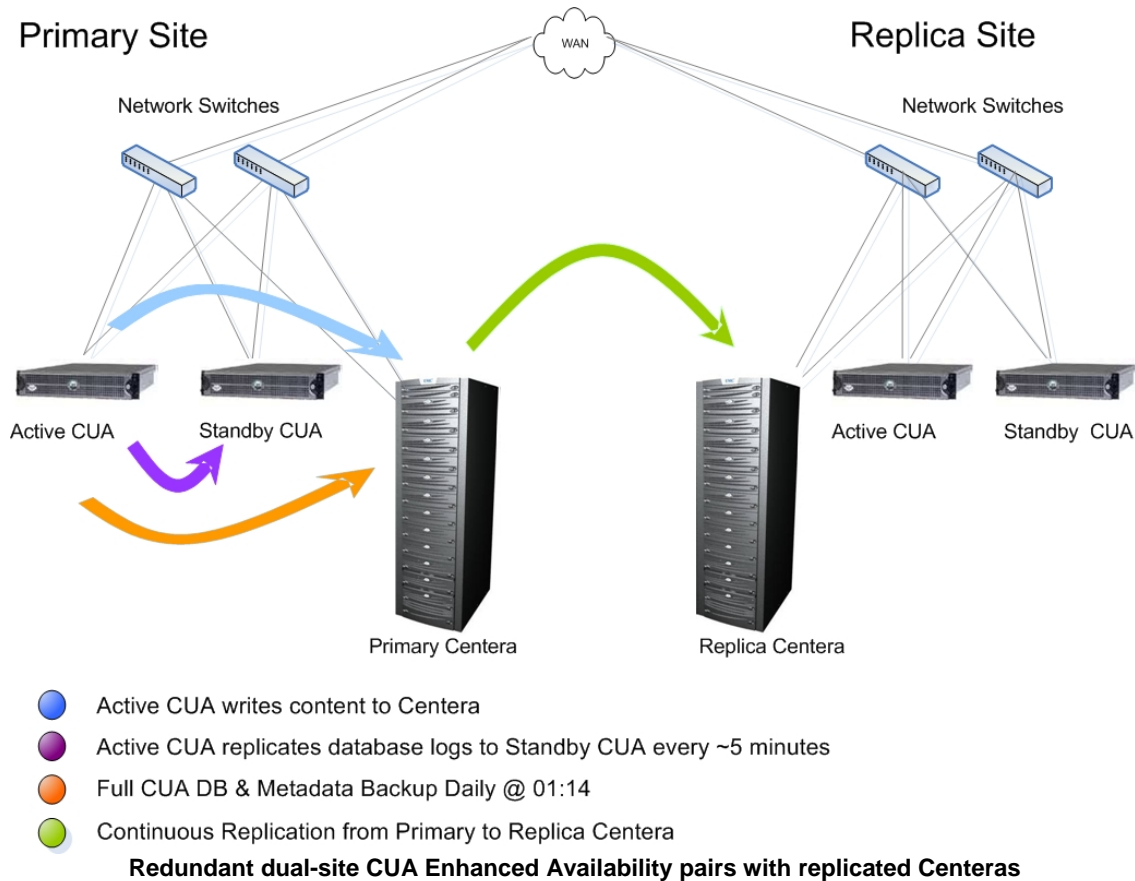
Single-site Enhanced Availability CUA pairs

CUA may be deployed with a non-replicated Centera cluster as an Enhanced Availability pair. Both the Active and Standby CUA are configured to address the single Centera cluster. In the event of a failure of the Active CUA, the Standby is promoted to Active status while the former Active CUA is repaired or replaced. At this point the former Active CUA is in the role of Standby. Later, the current Active CUA can be failed-back to the Standby, making it the Active CUA once again.

Note that a single-site deployment of CUA Enhanced Availability pairs will not protect content in the loss of the data center. If the Centera is lost, so is the content that CUA manages. To protect against this type of disaster, a replicated Centera configuration should be used.



CUA at the Replica site is restored, it synchronizes with its Standby CUA to guard against another potential single CUA failure.



Appendix A: Validating CUA Failover in a Test Environment

Introduction

This document describes how to validate a successful Centera Universal Access failover, and fail back again, in a test environment.

In a typical customer environment, Centera Universal Access's failover capabilities may never be used. Nevertheless, many customers test CUA failover as part of their overall acceptance testing prior to a CUA purchase, or as part of routine disaster recovery operations. Though the CUA failover operations are accurately described in the *CUA Monitor and Administration Guide*, there are areas where CUA failover operations, and subsequent restoration (also called *failback*) operations, can lead to unintended results. This document provides a step-by-step guide to a successful demonstration of CUA failover and failback operation in a test environment

Typically, a CUA failover and failback operation is performed to demonstrate the following high-level operations in an Enhanced Availability CUA environment:

1. Successful failover from an Active CUA to a Standby CUA
2. Validation that the previous Active and Standby CUAs have exchanged roles
3. Failback to the original Active and Standby configurations
4. Validation that the failback operation was successfully implemented and that a production environment may now be configured with confidence.

This document shows, in detail, the steps required for a successful implementation of failover testing:

1. Shut down the Active CUA
2. Verify that the Active CUA is not available
3. Transition the Standby CUA to an Active state
4. Use the IP Address of the old Active CUA
5. Configure the new Active CUA for Enhanced Availability
6. Reconfigure the old Active CUA as the Standby
7. Shut down the new Active CUA
8. Verify that the new Active CUA is not available
9. Transition the new Standby CUA to an Active state
10. Use the IP Address of the new Active CUA

In a test where only a single failover is required, follow steps 1 through 6 only.

This document should be used as an addition to, and not a replacement for, the *CUA Monitoring and Administration Guide* or the *CUA Service Guide*. These guides contain all CUA configuration instructions necessary for a successful implementation of CUA Enhanced Availability. For example, this document includes instructions to reinstall CUA software, but does not include step-by-step instructions for the installation of CUA software. For a successful test of CUA failover, keep the current *CUA Monitor and Administration Guide* handy.

Sources

CUA 3.5 Monitor and Administration Guide, P/N 069001199

Additional steps prior to validation

Before beginning the validation of CUA Enhanced Availability, you may want to perform additional configuration steps to ensure a smooth validation process:

- When the failover completes successfully, the system sends an e-mail notification to the address configured under E-Mail Alert Settings in the service administrative account. Set up this e-mail address and check it as a further indication that failover has completed successfully. See Section 2, *Initial CUA Configuration, E-mail Alert Settings* in the *CUA Monitor and Administration Guide* for detailed instructions on configuring the E-mail Alert Settings.
- If a profile other than the **Anonymous Profile** is being used, ensure that the **pea file** is copied to the Standby CUA. See Section 5, *Configuration Options, Using Access Profiles* in the *CUA Monitor and Administration Guide* for more information.

Validating CUA failover in a test environment

To transition the Standby CUA to an Active state in a test environment, complete the following instructions, and check off Step 1 after completing the numbered sub-steps:

___ Step 1: Shut down the Active CUA

1. Log in to the Active CUA using an SSH client and log in using the gwsetup account and password.
2. Select **s** at the **Setup Utility** menu to shut down the Active CUA.

If the Active CUA is running and reachable the software will refuse to transition the Standby CUA to an Active state. This is necessary in a test environment. In a real disaster, the CUA would most likely be unavailable due to a network or component failure. Complete the following instructions, and check off Step 2 after completing the numbered sub-steps:

— Step 2: Verify that the Active CUA is not available

1. Connect to the Standby CUA **Setup Utility** using an SSH client and log in using the gwsetup account and password.
2. Select **X** at the **Setup Utility** menu to enter the Bash Shell.
3. Type the following command: `tail /var/log/storigen_nms`
4. Look for output similar to the following:
Oct 28 09:34:39 CUA2 nmslib[619]: [63] [28-Oct-2004 9:34:39.614] [sev 3] [SCAGD] Active CUA (CUA1.centera.lab.emc.com) is [not responding]; pid 6530
Oct 28 09:37:14 CUA2 nmslib[619]: [64] [28-Oct-2004 9:37:14.184] [sev 4] [SCAGD] Active CUA (CUA1.centera.lab.emc.com) is in an [unreachable] state; pid 6530
5. Only proceed with the failover procedure below once you see [unreachable] in the output for the Active CUA.
6. **Note:** If you do not see [unreachable] in the output, enter the command `tail /var/log/storigen_nms` periodically until [unreachable] is displayed confirming that the Active CUA has shutdown.

Now it is time to transition the Standby CUA to an Active state. Complete the following instructions, and check off Step 3 after completing the numbered sub-steps:

— Step 3: Transition the Standby CUA to an Active state

1. Connect to the Standby CUA **Setup Utility** using an SSH client and log in using the gwsetup account and password.
2. Select **V** at the **Setup Utility** menu to start the failover process.

Failover causes the Standby CUA to attempt to initialize itself as the Active CUA.

When the failover completes, the Standby CUA transitions to the Active state and NFS exports and CIFS shares become available.

Note that any files that had not been written from the failed CUA to the Centera at the time of failure will not be available on the newly transitioned CUA.

At this time you may boot up the old Active CUA.

Because the new Active CUA has a different IP address to the failed CUA, CIFS and NFS clients will need to re-establish connectivity with the new Active CUA. If you would like the new Active CUA to use the same IP address as the old Active CUA, complete the following instructions, and check off Step 4 after completing the numbered sub-steps:

— Step 4: Use the IP Address of the old Active CUA

1. Log in to the Active CUA using an SSH client and log in using the gwsetup account and password.
2. Select **C** at the **Setup Utility** menu to display the **Configuration** menu.
3. Select **C** at the **Configuration** menu to begin the configuration process.
4. Enter the original Active CUA's hostname and IP address when prompted.
5. Modify any other specific settings such as the **SMTP From** and **SMTP ReplyTo** fields.
6. Press **Enter** to step through the rest of the options.
7. Reboot the system when prompted.

After the reboot, the new Active CUA will have the failed CUA's IP address and hostname. Any clients that were connected to the failed CUA should be rebooted to access the new Active CUA.

Note that if you configure the new Active CUA to use the old Active CUA's hostname/IP address, you must configure the old Active CUA with either a new hostname/IP address or the old Standby CUA's hostname/IP

address. Otherwise, you will have two CUA systems contending for the same hostname/IP address when the new Standby CUA comes online. Read the instructions in Step 6 carefully to ensure that there is no conflict between the new Active and old Active CUAs on the network.

Do not attempt to setup any new Standby CUA without first configuring the new Active CUA for Enhanced Availability. To configure the new Active CUA for Enhanced Availability, complete the following instructions, and check off Step 5 after completing the numbered sub-steps:

— Step 5: Configure the new Active CUA for Enhanced Availability

1. Connect to the Active CUA using an SSH client and log in with the gwsetup name and password.
2. Select **C** at the **Setup Utility** menu to open the **Configuration** menu.
3. Select **V** at the **Configuration** menu to begin the configuration and display the following prompts:
Configure Enhanced Availability
Active Gateway: [node.company.com]
Standby Gateway:
Do you want to change this configuration? (y=yes, n=no) [n]: y
Is this system an Active CUA? (y=yes, n=no) [y]: y
Will there be a Standby CUA? (y=yes, n=no) [n]: y
Standby CUA name or address:
Enter y at the first three prompts and then enter the hostname or IP address of the Standby CUA and press **Enter** on your keyboard.
4. After the message Active System configuration completed appears, press the Enter key on your keyboard.
5. Select **Q** to exit the **Configuration** menu and return to the main menu.
6. Select **X** to enter a **Bash Shell**.
7. Type the following command: tail /var/log/storigen_nms
8. Look for output similar to the following to ensure that the Active CUA is configured correctly:
Oct 28 09:32:05 CUA1 nmslib[603]: [75] [28-Oct-2004 9:32:05.560] [sev 2]
[SFSBACKUP] Backup complete, CA is 3PADPTQ8KHOCIE3NGHJ48T4S5VG

Note: Re-execute Step 7 if you do not see the Backup Complete message.

As failover is complete, and the new Active CUA is configured, the old Active CUA can now be reconfigured to become the Standby CUA. To reconfigure the old Active CUA as the Standby CUA, complete the following instructions, and check off Step 6 after completing the numbered sub-steps:

— Step 6: Reconfigure the old Active CUA as the Standby

When the old Active CUA was rebooted after Step 3, it registered that it was no longer the Active CUA and carried out the following activities:

- Removed the default TCP/IP gateway address.
- Restarted the network.
- Cleaned up /sfs/gateway_cifs and /sfs/gateway_nfs by moving any files that were not yet written to Centera from /sfs/gateway_nfs and /sfs/gateway_cifs to /sfs/serviceinfo/lost+found/.

At this point, the old Active CUA is not connected to the network. On Centera Gen 3 and Gen 4 systems, if the CUA is running version 3.0.5 or above, you can connect to the CUA by setting your laptop IP address to 10.255.0.2, connecting to the eth1 port of the CUA with a cross-over cable, and logging in through 10.255.0.1 using an SSH client. For all other systems, you must plug in a monitor and keyboard to the back of the CUA to reconfigure it.

Since the default TCP/IP gateway address was removed from the old Active CUA during the failover operation, you must enter the default TCP/IP gateway address through the **Configuration** menu and reboot the CUA. Make sure that there is no conflict with the new Active CUA's TCP/IP address. If you used the IP address of the old Active CUA for the new Active CUA as part of Step 4, then you must use a different IP address than the one originally used by this system.

1. Connect to the new Standby CUA using an SSH client and log in with the gwsetup name and password.
2. Select **C** at the **Setup Utility** menu to enter the **Configuration** menu.
3. Select **V** at the **Configuration** menu to begin configuration and display the following prompts:
 Configure Enhanced Availability
 Active Gateway: node.company.com
 Standby Gateway:
 Do you want to change this configuration? (y=yes, n=no) [n]: y
 Is this system an Active CUA? (y=yes, n=no) [y]: n
 Enter the name or IP address of the Active CUA:
 Enter y at the first prompt, n at the second and then complete the configuration by entering the hostname or IP Address of the Active CUA.
4. After the message Active System configuration completed appears, press the Return key on your keyboard.
5. Select **Q** to exit the **Configuration** menu and return to the main menu.
6. Select **X** to enter a **Bash Shell**.
7. Type the following command: tail /var/log/storigen_nms
8. Look for output similar to the following:
 Oct 28 09:27:17 CUA2 nmslib[619]: [60] [28-Oct-2004 9:27:17.047] [sev 2]
 [SFSRESTORE] Restore complete from CA 3PADPTQ8KHOCie3NGHJ48T4S5VG

Note: Re-execute Step 7 if you do not see the Restore complete message.

Now that the new Active and new Standby CUAs are configured successfully, it is time to transition the new Standby CUA to an Active state. This process involves a second failover and is sometimes called “failback,” as the original Active CUA, which is now the new Standby CUA, transitions back into its original Active role. To transition the new Standby CUA to an Active state, complete the following instructions, and check off Step 7 after completing the numbered sub-steps:

— Step 7: Shut down the new Active CUA

1. Log in to the new Active CUA using an SSH client and log in using the gwsetup account and password.
2. Select **S** at the **Setup Utility** menu to shut down the new Active CUA.

If the new Active CUA is running and reachable the software will refuse to transition the new Standby CUA to an Active state. This is necessary in a test environment. In a real disaster, the CUA would most likely be unavailable due to a network or component failure. Complete the following instructions, and check off Step 8 after completing the numbered sub-steps:

— Step 8: Verify that the new Active CUA is not available

1. Connect to the new Standby CUA **Setup Utility** using an SSH client and log in using the gwsetup account and password.
2. Select **X** at the **Setup Utility** menu to enter the Bash Shell.
3. Type the following command: tail /var/log/storigen_nms
4. Look for output similar to the following:
 Oct 28 09:34:39 CUA2 nmslib[619]: [63] [28-Oct-2004 9:34:39.614] [sev 3] [SCAGD] Active CUA (CUA1.centera.lab.emc.com) is [not responding]; pid 6530
 Oct 28 09:37:14 CUA2 nmslib[619]: [64] [28-Oct-2004 9:37:14.184] [sev 4] [SCAGD] Active CUA (CUA1.centera.lab.emc.com) is in an [unreachable] state; pid 6530
5. Only proceed with the failover procedure below once you see [unreachable] in the output for the new Active CUA.
6. **Note:** If you do not see [unreachable] in the output, enter the command
 tail /var/log/storigen_nms periodically until [unreachable] is displayed confirming that the new Active CUA has shut down.

Now it is time to transition the new Standby CUA to an Active state. Complete the following instructions, and check off Step 9 after completing the numbered sub-steps:

___ **Step 9: Transition the new Standby CUA to an Active state**

1. Connect to the new Standby CUA **Setup Utility** using an SSH client and log in using the gwsetup account and password.
2. Select **V** at the **Setup Utility** menu to start the failover process.

Failover causes the new Standby CUA to attempt to initialize itself as the Active CUA.

When the failover completes, the new Standby CUA transitions to the Active state and NFS exports and CIFS shares become available.

Because the new Active CUA has a different IP address to the failed CUA, CIFS and NFS clients will need to re-establish connectivity with the new Active CUA. If you would like the new Active CUA to use the same IP address as the old Active CUA, complete the following instructions, and check off Step 10 after completing the numbered sub-steps:

___ **Step 10: Use the IP Address of the new Active CUA**

1. Log in to the Active CUA using an SSH client and log in using the gwsetup account and password.
2. Select **C** at the **Setup Utility** menu to display the **Configuration** menu.
3. Select **C** at the **Configuration** menu to begin the configuration process.
4. Enter the new Active CUA's hostname and IP Address when prompted.
5. Modify any other specific settings such as the **SMTP From** and **SMTP ReplyTo** fields.
6. Press **Enter** to step through the rest of the options.
7. Reboot the system when prompted.

After the reboot, the Active CUA will have the failed CUA's IP address and hostname. Any clients that were connected to the failed CUA should be rebooted to access the Active CUA.

Congratulations! You have successfully tested CUA failover operations thoroughly in your test environment.

Next steps

If your test environment is also your production environment, you **must** reinstall CUA software on both the Active and Standby CUA systems and reconfigure each system for Enhanced Availability before moving into a production phase. This step is critical due to the state of the CUA Enhanced Availability pair at this completed stage of failover testing. The current Active CUA is no longer participating in Enhanced Availability. The other CUA is in a modified Standby state due to prior failover operations and reinstallation of the CUA software is **required**.