



EMC[®] Cloud Tiering Appliance and Cloud Tiering Appliance/VE

Version 12.0

Getting Started Guide

P/N 300-005-094

Rev 24

Copyright © 2016-2017 Dell Inc. or its subsidiaries All rights reserved.

Published July 2017

Dell 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." DELL 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 DELL SOFTWARE DESCRIBED IN THIS PUBLICATION REQUIRES AN APPLICABLE SOFTWARE LICENSE.

Dell, EMC, and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be the property of their respective owners. Published in the USA.

Dell EMC
Hopkinton, Massachusetts 01748-9103
1-508-435-1000 In North America 1-866-464-7381
<http://www.DellEMC.com>

Chapter 1	Introduction.....	11
	Overview of the Cloud Tiering Appliance	12
	File tiering and migration	12
	Block archiving	13
	Cloud Tiering Appliance/VE (CTA/VE)	13
	High Availability for CTA and CTA/VE.....	13
	CTA implementation with VNX or Unity file primary storage	14
	CTA implementation with NetApp primary storage	15
	CTA implementation with Unity block snapshots	17
	CTA and CTA/VE tasks.....	17
	Using CTA and CTA/VE	19
Chapter 2	Cloud Tiering Appliance Hardware and Port Configurations	21
	Cloud Tiering Appliance 12 qualified hardware.....	22
	Cloud Tiering Appliance High Availability types.....	24
	Contents of the appliance	25
	Appliance diagrams.....	25
	Port details	26
Chapter 3	Installing the Cloud Tiering Appliance.....	27
	Appliance setup	28
	Installing the virtual edition	28
	CTA and CTA/VE for high availability	29
	High availability with VNX or Unity primary storage	30
	High availability with NetApp primary storage.....	30
	Network console management for Gen 8 models	31
	Using CTA CLI commands.....	32
	Rebooting the system	32
	Performing a CD clean install on the appliance	32
	Performing a CD clean install on a generic server.....	33
	Adding the CTA serial number.....	34
	Configuring CTA.....	34
	Configure the CTA network	35
	Configure the hostname, domain, and DNS server	36
	Configure iSCSI IQN targets (Unity block archiving and restore only)	36

Update the hosts file	37
File Encryption	37
Enabling keystore replication.....	37
Configuring and using file encryption.....	37
Graphical user interface.....	38
Command line interface.....	39
Chapter 4 Deploying the Cloud Tiering Appliance.....	41
Deployment process for archiving.....	42
Step 1: Set up CTA	43
Step 2: Configure the archive environment	43
Step 3: Configure the destination	43
Step 4: Define policies.....	44
Step 5: Create a task.....	44
Step 6: Run the simulation task (optional)	44
Step 7: Run the policy task	44
Supported platforms	44
Deploying CTA with VNX.....	45
Prerequisites for using VNX as a file migration source or destination	46
Prerequisites for VNX as an archive source	46
Adding VNX to the CTA configuration.....	47
Configure name resolution for archiving.....	50
Configuring VNX to Centera or cloud archiving on the CTA.....	51
Prerequisite tasks on the VNX Control Station for file migration or archiving .	52
Deploying CTA with NetApp.....	57
Prerequisites for using NetApp as a file migration source.....	57
Prerequisites for using NetApp as an archiving source	57
vFiler configuration	60
Configuring NetApp archiving on the CTA	60
Adding NetApp filer to the CTA configuration.....	60
Verify the FPolicy configuration for CTA	63
Deploying CTA with a VNXe	64
Prerequisites for using VNXe as a file migration destination.....	64
Adding VNXe to the CTA configuration.....	64
Deploying CTA with Unity	65
Prerequisites for Unity as an archive source.....	65
Configure name resolution for archiving.....	66
Configuring Unity to cloud archiving on the CTA	66
Adding a Unity file server to the CTA configuration	67
Adding a Unity block server to the CTA configuration	68
Deploying CTA with Virtustream	70
Deploying CTA with a Windows server.....	71
Adding a Windows server to the CTA configuration.....	72

Installing the EMWS copy agent for CTA	72
Installing and running LGDUP	73
Deploying CTA with Isilon	74
Prerequisite tasks when using Isilon as a CIFS share destination	74
Prerequisite tasks when using Isilon as an NFS export destination	76
Prerequisite task when using Isilon as a file migration destination	76
Adding Isilon to the CTA configuration	77
CTA and Isilon SmartConnect	78
Deploying the CTA with Data Domain	80
Configuring a NAS-based repository	81
Deploying CTA with Amazon S3	82
Adding Amazon S3 to the CTA configuration	82
Deploying CTA with Atmos	83
Adding Atmos to the CTA configuration	83
Installing the SSL certificate on the CTA	84
Deploying CTA with Azure	85
Adding Azure to the CTA configuration	85
Deploying CTA with Centera	86
Deploying CTA with ECS	86
Adding ECS to the CTA configuration	87
Chapter 5 Maintaining the Cloud Tiering Appliance.....	89
Importing a file list archive	90
Adding the primary servers	90
Configure name resolution for archiving.....	90
Configuring Unity to cloud archiving on the CTA	91
Configuring the import provider	92
Configuring the import task.....	92
Importing the file list.....	93
Validating the import file	94
Backing up the configuration	94
Creating a backup dump.....	95
Restoring a backup dump	95
Encryption Key restore in case of catastrophic failures.....	97
Maintaining the database.....	98
Checking database size and disk capacity	98
Performing database maintenance	98
Migrating from CTA to CTA/VE	99
Shutting down and restarting the appliance	100
Chapter 6 Cloud Tiering Appliance System Settings	101
Security hardening	102
Single security database.....	102

Disable root logins.....	104
Strengthen passwords.....	104
Age passwords	105
Configuring the GUI access method	105
STIG hardening.....	105
Enabling STIG hardening.....	105
Disabling STIG hardening.....	106
LDAP client configuration.....	107
Global LDAP settings.....	107
LDAP authentication	108
Configuring LDAP settings.....	109
Certificate management.....	111
Appliance mail delivery settings	111
Log settings.....	112
Configuring log rotation	112
Configuring SCP of rotated log files	112
Configuring alerts	114
Configuring email alerts.....	114
Configuring SNMP alerts	115
Enabling SNMP polling.....	116
System command accounting.....	116
Tracking user command history	117
Tracking user login history	117
Tracking daemon command history	117
Windows domain user	118
Creating a Windows domain user.....	118
Adding an admin user to the local administrator group.....	118
Configuring Windows for Kerberos	119
Configuring Windows 2008 for NTLM	119
SID translator	120
Installing the SID translator	120
Creating the SID translation file	121
Uploading the SID translation file	121
Deleting a SID translation file.....	122
Appendix A Network topology scenarios	123
Advanced network topologies.....	124
Configuring the CTA with bonding	124
Configuring the CTA with two subnets	125
Configuring the CTA with more than two subnets	126
VLAN tagging modes for the CTA/VE	127
ESX Server virtual switch tagging	127
ESX Server virtual guest tagging.....	128

Appendix B Alerts..... 131
Supported SNMP traps 132
CTA alerts 133

Figures

Figure 1.	VNX or Unity file implementation	14
Figure 2.	NetApp implementation	15
Figure 3.	Unity block implementation.....	17
Figure 4.	Rear view of Gen 8 appliance.....	25
Figure 5.	Front view of Gen 8 appliance with bezel removed.....	26
Figure 6.	Cloud Tiering Appliance deployment process	42

Table 1.	Dell R630 G13 hardware qualified for CTA	22
Table 2.	CTA-HA that is based on Dell R630 G13	22
Table 3.	CTA that is based on Intel C1U.....	23
Table 4.	CTA-HA that is based on Intel C1U	24
Table 5.	VMware ESX Server interoperability.....	28
Table 6.	Supported platform matrix for archiving	44
Table 7.	Supported platform matrix for file migration	45
Table 8.	VNX deployment checklist.....	45
Table 9.	NetApp deployment checklist.....	57
Table 10.	Supported SNMP traps	132
Table 11.	CTA alerts	133

Preface

As part of an effort to improve its product lines, EMC periodically releases revisions of its software and hardware. Therefore, some functions described in this document might not be supported by all versions of the software or hardware currently in use.

The product release notes provide the most up-to-date information on product features. Contact your EMC technical support professional if a product does not function properly or does not function as described in this document.

Audience

This document is part of the Cloud Tiering Appliance and Cloud Tiering Appliance/VE documentation set. The documentation is intended for use by:

- Storage management administrators who are new to the Cloud Tiering Appliance and Cloud Tiering Appliance/VE.
- Existing customers who are new to version 12.0.

Related documents

The following publications provide additional information:

- Cloud Tiering Appliance White Papers — Provide best practices for configuring specific file servers and for conducting specific tasks.
- Cloud Tiering Appliance and Cloud Tiering Appliance/VE online help — Provides detailed reference information on the graphical user interface.
- Cloud Tiering Appliance Upgrade Guide — Provides instructions on upgrading to the current CTA version.
- Cloud Tiering Appliance and Cloud Tiering Appliance/VE Release Notes — Provides an overview of new features and lists any limitations.
- CTA man pages — Provide detailed command-line help, as well as overview information. A good starting point is: `man rffm`. PDFs of all man pages are available from:

`/opt/rainfinity/filemanagement/doc`

Where to get help

Support, product, and licensing information can be obtained as follows:

For information about products, software updates, licensing, service, and product documentation, go to Online Support (registration required) at: <https://Support.EMC.com>. After logging in, locate the appropriate **Support by Product** page.

For technical support and service requests, go to Online Support at: <https://Support.EMC.com>. After logging in, locate **Create a service request**. To open a service request, you must have a valid support agreement. Contact your Sales Representative for details about obtaining a valid support agreement or to answer any questions about your account.

Chapter 1 Introduction

This chapter includes the following sections:

Overview of the Cloud Tiering Appliance	12
CTA implementation with VNX or Unity file primary storage.....	14
CTA implementation with NetApp primary storage.....	15
CTA implementation with Unity block snapshots.....	17
CTA and CTA/VE tasks	17
Using CTA and CTA/VE	19

Overview of the Cloud Tiering Appliance

The Cloud Tiering Appliance (CTA) provides a full range of features including the ability to:

- Tier or archive and recall file data
- Migrate files
- Perform orphan file management
- Perform stub file recovery
- Archive and restore block data (Unity only)
- Simulate the potential effect of policies before taking action

The CTA software also includes a robust reporting interface that provides valuable insight into the efficacy of archiving policies.

Depending on the type of task, CTA supports a variety of platforms:

- Supported platform matrix for archiving in [Table 6](#).
- Supported platform matrix for file migration in [Table 7](#).

File tiering and migration

CTA includes technology that optimizes primary file storage by automatically moving inactive files based on policies to secondary storage. Secondary storage could be lower cost drives, such as NL-SAS or SATA drives, or to other platforms, including public and private clouds. Files that are moved appear as if they are on primary storage. File tiering dramatically improves storage efficiency, and backup and restore time. File archiving onto storage with WORM functionality can support additional business requirements such as compliance and retention.

As an example, a CTA may be configured to locate all NAS data that has not been accessed in one year, and move that data to secondary storage. For each file it moves, the appliance will leave behind a small space-saving stub file that points to the real data on the secondary storage device. When a user tries to access the data in its original location on the primary NAS, the user will be transparently provided with the actual data that the stub points to, from secondary storage.

If a multi-tier policy is used, the appliance may be configured to move files from a secondary storage device tier to a tertiary storage device tier. This can be particularly useful in cases where the secondary storage device represents a tier that is smaller, faster, and more expensive to maintain than a larger, slower, and cheaper storage used in the tertiary tier. Once the files are moved, the space-saving stub file on the primary NAS tier would be updated to point to the data's new location on the tertiary storage tier.

In addition to tiering data from primary to secondary storage and leaving a stub behind on the primary server for recall later, the CTA can also permanently migrate files from a source to a destination without leaving a stub.

Block archiving

CTA leverages the Unity snapshot differentials API in order to efficiently take backups of block data to the cloud. You can archive a full copy of a LUN, consistency group or thin clone to the cloud and perform incremental archives using subsequent snapshots. You can also customize the *Common Base Frequency* attribute in the CTA management GUI to decide the spacing of full copies.

Block snapshots that have been archived to the cloud, are independent of the base resource on the source and are not altered by the archiving. Instead they share blocks with the baseline snapshot established by CTA. Once the LUN/snapshot is archived, it can be deleted in the source array to free up space.

Cloud Tiering Appliance/VE (CTA/VE)

The Cloud Tiering Appliance/VE (CTA/VE) is a VMware virtual appliance installed on a VMware ESX/ESXi Server. CTA/VE is provided in an industry-standard virtual appliance distribution that consists of an Open Virtualization Format (OVF) and Virtual Machine Disk (VMDK) file.

Virtual appliances are prebuilt software solutions, comprised of one or more virtual machines that are packaged, updated, maintained, and managed as a unit. Unlike a traditional hardware appliance, these software appliances allow customers to acquire, deploy, and manage pre-integrated solution stacks more quickly and easily.

High Availability for CTA and CTA/VE

The Cloud Tiering Appliance for High Availability (CTA-HA) and the Cloud Tiering Appliance/VE for High Availability (CTA/VE-HA) complement the existing CTA or CTA/VE by ensuring that tiered files can always be recalled, even in the event that the primary appliance goes down. This ensures complete transparency and nondisruptive service for clients. The high availability appliance is not used for any purpose other than recall. For example, it does not perform archiving or orphan file management, nor does it have a graphical user interface.

The CTA-HA is a dedicated machine that runs the VNX, Unity or NetApp callback agents and is delivered preloaded with software. Installation instructions for the CTA-HA differ slightly from the CTA. The CTA/VE-HA provides the same functionality as the CTA-HA but is installed on a virtual appliance like the CTA/VE. A virtual HA can be deployed with a physical CTA and vice versa. In this way, a network with a physical CTA does not need a second piece of hardware to provide HA, it can use a CTA/VE-HA.

When a high-availability appliance is deployed alongside a CTA or CTA/VE, the underlying APIs of VNX, Unity NAS servers and NetApp file servers are leveraged to create a highly available environment for data recall. The VNX, Unity and NetApp implementations differ, as shown in [Figure 1](#) and [Figure 2](#).

High-availability appliances are not needed in all tiered source and destination combinations. The CTA and CTA/VE Interoperability Matrix provides more details.

CTA implementation with VNX or Unity file primary storage

Figure 1 shows the recall architecture of CTA implementation with VNX, or Unity files:

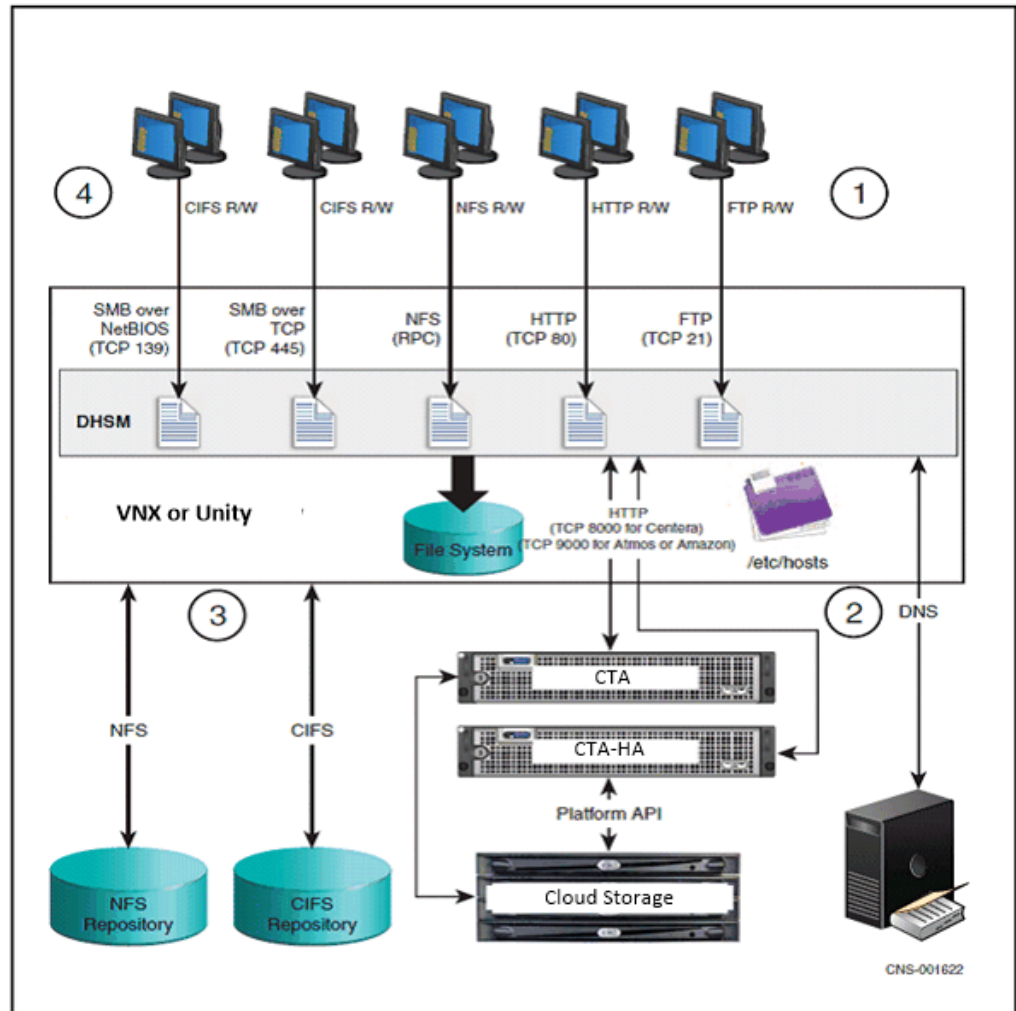


Figure 1. VNX or Unity file implementation

Circled numbers correspond to the following steps that illustrate the recall process in the VNX or Unity implementation:

1. Clients send read/write operations for files that have been archived. These stub files contain information about where the file has been archived to. These operations are intercepted by the DHSM layer on the VNX or Unity system prior to being serviced from the file system.
2. If the file has been archived to EMC Centera® or cloud storage (such as EMC Atmos™, Amazon S3, ECS/S3, or Microsoft Azure), the VNX or Unity resolves the fully qualified domain name (FQDN) stored in the stub file to the IP address of the CTA or CTA-HA.

The storage system then uses HTTP to read the archived data from the appliance, which in turn reads it from cloud storage by using the platform API. If an appliance does not respond to the HTTP read requests, the VNX or Unity system uses an alternate IP address of another appliance configured in DNS. Every

callback server (CTA, CTA-HA, CTA/VE, or CTA/VE-HA) has its IP address associated with a single hostname in DNS. The FQDN uses that hostname, which may have multiple IP addresses associated with it.

3. If the file has been archived to an NFS or CIFS repository, the VNX opens a connection to the repository and reads back the data.
4. The VNX responds to the client operation as usual if the recall was successful, or the client receives a message that the file cannot be opened if the recall fails.

Note When VNX data has been archived to a VNX, VNXe, NetApp, Data Domain, Isilon, or Microsoft Windows repository, the CTA or CTA/VE is not involved in the recall process, and a high availability appliance is not used.

CTA implementation with NetApp primary storage

Figure 2 shows the recall architecture of CTA implementation with a NetApp system:

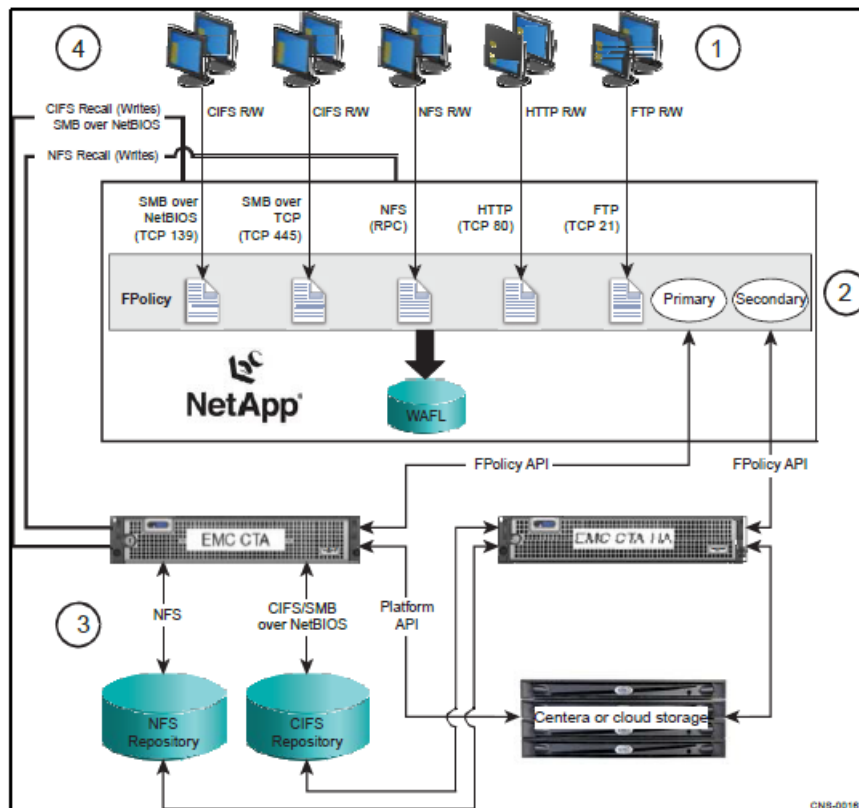


Figure 2. NetApp implementation

Circled numbers correspond to the following steps that illustrate the archive and recall process in the NetApp implementation:

1. Clients send read/write operations for files that have been archived. These operations are intercepted by the FPolicy layer on the NetApp prior to being serviced from the Write Anywhere File Layout (WAFL) filesystem.

2. The NetApp is configured with the following groups:
 - a. A primary group of callback servers, such as a CTA and possibly one or more CTA-HAs.
 - b. A secondary group, such as one or more physical or virtual HAs.

The NetApp will send FPolicy callbacks to servers registered in the primary group in round-robin fashion. If a server does not reply to the callback, it is removed from its group. If there are no servers in the primary group, the callbacks are distributed in a round-robin fashion among the servers in the secondary group.

For CTA/VE, the primary group of callback servers consists of virtual machines such as a CTA/VE and possibly one or more CTA/VE-HAs. The secondary group consists of one or more CTA/VE-HAs.

3. The appliance connects to the filer by using CIFS to read the contents of the stub file. The stub file points to where the file data is stored. The appliance then connects to the NFS repository, CIFS repository, Centera, or cloud storage where the data was archived. It then reads the data by using the native protocol and the file data is written back to the NetApp.
4. The filer responds to the client operation as usual if the recall was successful, or with an "access denied" message if the recall failed.

Note It is a requirement that the software versions of all the appliances match. For example, do not deploy a configuration with a CTA that is running version 8.0 and a CTA-HA that is running version 7.4. While the software does not perform any explicit checks to ensure the versions are compatible, the running of different software versions has not been tested and may result in unexpected behavior.

CTA implementation with Unity block snapshots

Figure 3 shows the architecture of CTA implementation with Unity block snapshots.

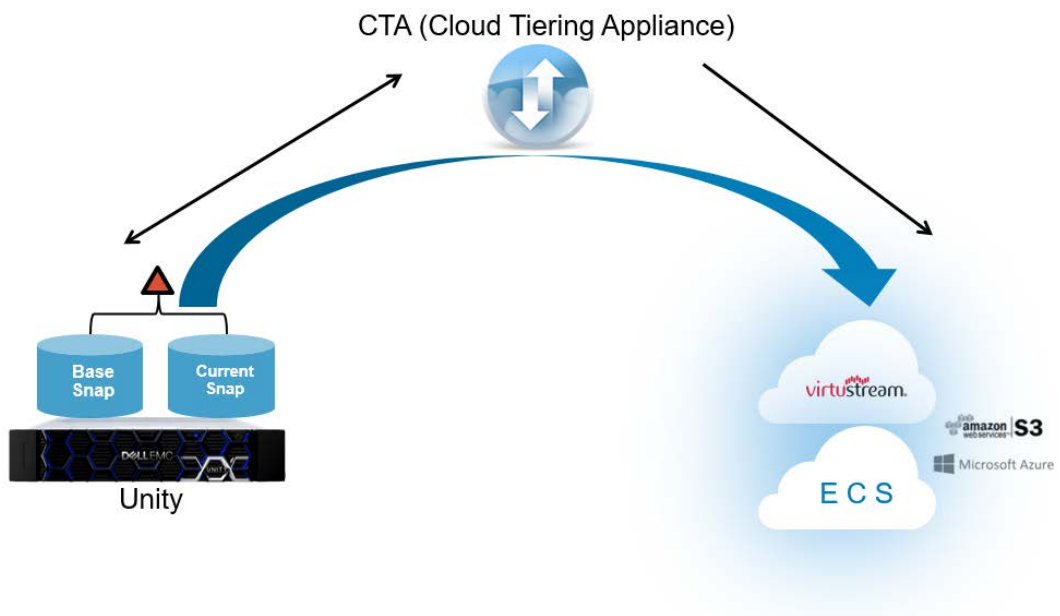


Figure 3. Unity block implementation

CTA and CTA/VE tasks

The CTA or CTA/VE may be used to run several different tasks:

- Tiering or archiving
- Deleting orphaned data
- Auxiliary tasks, such as stub scanning and backup
- Migration tasks such as repository migration and file migration

For archiving, deleting and migrating files, the software leverages a policy engine to select the files. Users can combine and evaluate multiple rules together in a single policy. Several rule types are available.

Before running the archive, delete, or migration task, the running of a simulation allows administrators to review real-time results without executing the task. The results will return:

- Aggregated summary of total files matched
- Total bytes potentially archived

Also, if an optional detailed simulation is run, a list of files that match the policy is saved on the disk for review.

Run a simulation to gain insight into the efficiency of a task before running the task. This practice is notably important for the delete tasks, since these tasks remove data. A report displays results of the task.

Archive tasks may be one of three types:

- Archive (with policy) — Archives regular (non-stub) files. Files are selected for archiving based on the archive policy.
- Multi-tier (with policy) — For this archiving task, regular and stub files are evaluated with the multi-tier policy.
 - a. If a regular file matches the policy, it is archived.
 - b. If a stub file matches the policy, archived data is moved to a different storage tier and the stub is updated to point to the new location.
- Multi-tier stub (with policy) — For this archiving task, only stub files are evaluated with the multi-tier stub policy. If a stub file matches the policy, archived data is moved to a different storage tier and the stub is updated to point to the new location. Otherwise, the archived data remains in the current storage tier.

In addition, CTA can perform an archive task on an imported list of files.

Delete tasks may be one of two types:

- Delete orphan with policy — Deletes orphans on secondary storage that match the delete orphans policy.
- Delete stub with policy — The delete stub task deletes stubs that match the delete stubs policy. Stubs on primary storage and files on the second tier that are no longer under retention or that were defined without any retention period are automatically deleted.

Auxiliary tasks are:

- Scan stubs — When a file is archived, a stub file remains on the source and an entry is added to the CTA database, and maps the name and location of the archived file to its stub. The stub scanning task tracks the stub file information. When a stub has not been detected on a CIFS share or NFS export for 30 or more days, the corresponding archived file is designated as an orphan. If stub files are moved from the original archive location to another location, manually run a stub scanner task on the new location so that the CTA does not consider these files orphans.
- Backup — The backup task performs periodic backups of the CTA configuration and database. Schedule backup tasks as part of a regular maintenance program. Only one backup task may be scheduled.

Migration tasks are:

- Repository Migration — Repository migration moves archived files from one storage tier to another storage tier. Migration can be to a NAS repository, to a Centera, or to an on-premise or service provider cloud such as Atmos or Amazon S3. All stub files that point to this data will be updated to point to the new location.

- **File Migration (with policy)** — The file migration task moves files from a primary to a secondary server that are defined as the source and destination in the task definition. The task evaluates both normal and stub files based on the migrate file policy with a resulting action of either migrate or don't migrate.

The CTA software also has the capability to recover stub files accidentally deleted by client systems. It can even recover prior versions of files archived to any secondary storage destination.

Note Do not duplicate stubs, because CTA does not support them. Repository migration and Orphan deletion tasks will randomly detect and manage the first occurrence of a stub, and will ignore the duplicates. This can lead to data unavailability and data loss.

Using CTA and CTA/VE

Once the appliance has been deployed on the network, the administrator can manage data through the CTA or CTA/VE graphical user interface (GUI) or command line interface (CLI). [Graphical user interface](#) explains how to invoke the GUI. Online help documents all GUI pages.

Technical system details that are not related to the GUI, but are required to configure the CTA or CTA/VE, are provided in the following chapters and appendixes:

- [Chapter 4, Deploying the Cloud Tiering Appliance](#)
- [Chapter 5, Maintaining the Cloud Tiering Appliance](#)
- [Chapter 6, Cloud Tiering Appliance System Settings](#)
- [Appendix A, Network Topology scenarios](#)

Except where expressly stated otherwise, all sections of this guide apply to both the CTA and CTA/VE.

Chapter 2 Cloud Tiering Appliance Hardware and Port Configurations

This chapter contains the following sections:

Cloud Tiering Appliance 12 qualified hardware	22
Appliance diagrams	25
Port details	26

Cloud Tiering Appliance 12 qualified hardware

- Dell R630 G13 server is qualified with CTA 12. User must perform a CD clean install of CTA and CTA-HA after purchasing these servers.
- Table 1 lists the Gen 13 hardware configuration of the Dell R630 server qualified for CTA

Table 1. Dell R630 G13 hardware qualified for CTA

Component	Dell Power Edge R630 G13
Chassis	Chassis with up to 8, 2.5" hard drives, 3 PCIe slots
Size	1U
Power	Dual, Hot-plug, Redundant Power Supply (1+1) 750W
CPUs	Single Intel Xeon E5-2620 v3 2.4GHz. 15M Cache
Disks	Four 1.2TB 10K RPM SAS 12Gbps 2.5" Hot-plug hard drives in a RAID-1 configuration with two hot spares. SATA or Nearline SAS hard drives are an option.
RAID controller	PERC H730 Integrated RAID Controller, 1GB Cache
CD-ROM	Internal SATA DVD-ROM that can read CD or DVD material for system install or upgrade.
Memory	2400-MHz, 16GB
Network interfaces	Two 10/100/1000 Mbps and two 100 Mbps/1 Gbps/10 Gbp
VGA	Two 15-pin B10VGA ports on the front and back panels to connect the system to a VGA display for system console.
Keyboard connector	Two 4-pin, USB 2.0-compliant ports on the front and back panels to connect a standard USB keyboard for system console.
Mouse connector	Two 4-pin, USB 2.0-compliant ports on the front and back panels to connect a standard USB mouse for system console.
Serial port	One standard DB9 serial connector on the back panel for a serial-terminal system.

- Table 2 lists the Gen 13 hardware configuration of the Dell R630 server qualified for CTA-HA

Table 2. CTA-HA that is based on Dell R630 G13

Component	Dell Power Edge R630 G13
Chassis	Chassis with up to 8, 2.5" hard drives, 3 PCIe slots
Size	1U
Power	Dual, Hot-plug, Redundant Power Supply (1+1) 750W

Component	Dell Power Edge R630 G13
CPUs	Single Intel Xeon E5-2620 v3 2.4GHz. 15M Cache
Disks	Two 1.2TB 10K RPM SAS 12Gbps 2.5" Hot-plug hard drives in a RAID-1 configuration. SATA or Nearline SAS hard drives are an option.
RAID controller	PERC H730 Integrated RAID Controller, 1GB Cache
CD-ROM	Internal SATA DVD-ROM that can read CD or DVD material for system install or upgrade.
Memory	2400-MHz, 8GB Total
Network interfaces	Two 10/100/1000 Mbps and two 100 Mbps/1 Gbps/10 Gbp
VGA	Two 15-pin B10VGA ports on the front and back panels to connect the system to a VGA display for system console.
Keyboard connector	Two 4-pin, USB 2.0-compliant ports on the front and back panels to connect a standard USB keyboard for system console.
Mouse connector	Two 4-pin, USB 2.0-compliant ports on the front and back panels to connect a standard USB mouse for system console.
Serial port	One standard DB9 serial connector on the back panel for a serial-terminal system.

When performing CTA/CTA-HA bare metal installations, please refer to the corresponding documentation on the vendor's website for hardware comparable to the Dell R630.

Cloud Tiering Appliance type

The CTA8-APL hardware appliance based on Intel C1U hardware is bundled with VNX.

- [Table 3](#) lists the Gen 8 hardware configuration for the CTA that is based on the Intel C1U hardware.

Table 3. CTA that is based on Intel C1U

Component	CTA8-APL
Chassis	The appliance is based on Intel C1U Kylin node hardware.
Size	1U form factor.
Power	Dual 750 watt. Item (b) in Figure 4 . Power button and identification light is item (f) in Figure 5 .
CPUs	Dual Intel Sandy Bridge Six Core AES 256 HW encryption.
Disks	Four 900 GB 2.5-inch, 10K RPM hard drives in a RAID-5 configuration. Items (b) through (e) in Figure 5 .
RAID controller	Intel RMS25CB080 (Condado Beach).

Component	CTA8-APL
CD-ROM	Read-only DVD that can read CD or DVD material for system upgrades. Item (a) in Figure 5 .
Memory	16 GB
Network interfaces	Four on-board gigabit 10/100/1000 TX Ethernet copper ports with RJ45 connectors. Item (e). Two 10GbE optical ports provided with the Intel mezzanine module. Item (g). One dedicated management port. Item (f). All items in Figure 4 .
VGA	Standard VGA video connector for a system console. Item (a) in Figure 5 .
Keyboard connector	Standard USB keyboard connector for a system console. Item (d) in Figure 5 .
Mouse connector	Standard USB mouse connector for a system console. Item (c) in Figure 5 .

Cloud Tiering Appliance High Availability types

CTA8-HA-APL based on Intel C1U hardware is bundled with VNX. [Table 4](#) lists the Gen 8 hardware configuration for the CTA-HA that is based on the Intel C1U hardware.

Table 4. CTA-HA that is based on Intel C1U

Component	CTA8-HA-APL
Chassis	The appliance is based on Intel C1U Kylin node hardware.
Size	1U form factor.
Power	Dual 750 watt.
CPU	Single Intel Sandy Bridge Six Core AES 256 HW encryption.
Disks	Two 900 GB 2.5-inch, 10K RPM hard drives in a RAID-1 configuration. Items (b) and (c) in Error! Reference source not found .
RAID controller	Intel RMS25CB080 (Condado Beach).
CD-ROM	Read-only DVD that can read CD or DVD material for system upgrades. Item (a) in
Memory	16 GB.
Network interfaces	Four on-board gigabit 10/100/1000 TX Ethernet copper ports with RJ45 connectors. Item (e). Two 10GbE optical ports provided with the Intel mezzanine module. Item (g). One dedicated management port. Item (f). All items in Figure 4 .

Component	CTA8-HA-APL
VGA	Standard VGA video connector for a system console. Item (a) in Figure 4 .
Keyboard connector	Standard USB keyboard connector for a system console. Item (d) in Figure 4
Mouse connector	Standard USB mouse connector for a system console. Item (c) in Figure 4 .

Models CTA8-APL and CTA8-HA-APL support four on-board gigabit Ethernet copper 10/100/1000TX ports, two 10 gigabit Ethernet SR (short range) optical ports, and one dedicated management port (MGMT). [Network console management for Gen 8 models](#) describes how to use the management port.

Contents of the appliance

The CTA or CTA-HA ships with robust, fault-tolerant hardware consistent with the mission-critical application for which it is used.

The following items are included in the appliance package:

- Cloud Tiering Appliance (1U 19-inch rack-mountable Gen 8 model).
- Two universal rails for mounting the appliance on a 19-inch rack.
- Two sets of power cords.
- Copper patch cables for the number of ports on your appliance.
- Media kit with the software recovery DVD.
- One crossover cable for Gen 8 models.

Note The VGA monitor, keyboard, and mouse for a system console are not included.

Appliance diagrams

These diagrams illustrate configurations of the CTA and CTA-HA.




Figure 4. Rear view of Gen 8 appliance



Figure 5. Front view of Gen 8 appliance with bezel removed

Port details

Models CTA8-APL and CTA8-HA-APL support four on-board gigabit Ethernet copper 10/100/1000TX ports, two 10 gigabit Ethernet SR (short range) optical ports, and one dedicated management port (MGMT). [Network console management for Gen 8 models](#) describes how to use the management port.

Chapter 3 Installing the Cloud Tiering Appliance

This chapter contains the following sections:

Appliance setup	28
Installing the virtual edition	28
CTA and CTA/VE for high availability	29
Network console management for Gen 8 models	31
Performing a CD clean install on the appliance	32
Performing a CD clean install on a generic server	33
Adding the CTA serial number	34
Configuring CTA	34
File Encryption	37
Graphical user interface	38
Command line interface	39

Appliance setup

Before an appliance may be used to perform tasks, the appliance and the software must be properly configured:

- If a Cloud Tiering Appliance (CTA) is being deployed, port details that are used to connect the appliance to the network are provided in [Chapter 2, Cloud Tiering Appliance Hardware and Port Configurations](#).
- If a Cloud Tiering Appliance/VE (CTA/VE) is being deployed, follow the instructions in [Installing the virtual edition](#).
- If a Cloud Tiering Appliance for High Availability (CTA-HA) or Cloud Tiering Appliance/VE for High Availability (CTA/VE-HA) is being deployed, [CTA and CTA/VE for high availability](#) describes configuration considerations.
- Follow the instructions in [Network console management for Gen 8 models](#).
- The CTA software is preinstalled on every new appliance. If the software must be reinstalled and no previous CTA information or data needs to be retained, follow the instructions provided in [Performing a CD clean install on the appliance](#).
- To install the appliance on the network, follow instructions provided in [Configuring CTA](#).
- If a CTA-HA or CTA/VE-HA is deployed and encryption is to be used, follow the instructions provided in [File Encryption](#).
- If the system requires security hardening or any other special configuration, [Chapter 6, Cloud Tiering Appliance System Settings](#) provides information for all system settings.

Then proceed to configure the appliance for your environment as described in [Chapter 4, Deploying the Cloud Tiering Appliance](#).

Installing the virtual edition

The CTA/VE and CTA/VE HA are installed on the VMware server. [Table 5](#) shows the interoperability.

Table 5. VMware ESX Server interoperability

Virtual Appliance	VMware ESX Server	Comments
CTA/VE	ESXi 5.0 ESXi 5.1 ESXi 5.5 ESXi 6.0	Four virtual CPUs, 16 GB of RAM, 1 TB of disk space, 2 gigabit virtual interfaces are reserved.
CTA/VE HA	ESXi 5.0 ESXi 5.1 ESXi 5.5 ESXi 6.0	Four virtual CPUs, 4 GB of RAM, 100 GB of disk space, 2 gigabit virtual interfaces are reserved.

Note If VMware snapshots will be used, reserve extra disk space when creating the VMware datastore during installation. To estimate the total amount of disk space to reserve for the datastore, add thirty percent of the size of the CTA/VE VMDK file for each snapshot.

The following example shows the steps to install the CTA/VE or CTA/VE HA virtual appliance using vSphere Client 5.0:

1. Download the CTA .OVA file.
 - For a CTA/VE, the file is: `ctave-<version>.ova`
 - For a CTA/VE HA, the file is: `ctaveha-<version>.ova`
2. Open the vSphere Client.
 - a. Click the **Hosts** tab. To find the appliance with the most free space, consider %CPU and %Memory.
 - b. Select the line with the ESX server for the installation. A summary of the CPU, memory, and data store capacities appears.

If the ESX Server has enough CPU and memory available to install the CTA/VE, per the requirements of [Table 5](#), proceed to the next step.
3. Import the OVA file. Instructions differ depending upon VMware version.
4. Select **File > Deploy OVF Template**.
5. On the **Deploy OVF Template** screen, click **Browse** to locate the OVA file.
6. Click **Next** through several screens until the **Storage** screen appears. Select the destination storage on the appliance.
7. Click **Next** through several screens until the **Ready to Complete** screen appears. Validate the information and click **Finish**.
8. The import may take 3–30 minutes depending on the network connection between the VI Client and the VMware ESX Server. Approximately 1 GB will initially be transferred across the network.

CTA and CTA/VE for high availability

High availability is a feature on both CTA and CTA/VE. Both the CTA-HA and the CTA/VE-HA deliver solutions for a redundancy, which ensure that clients do not experience data unavailability due to failure of the primary appliance.

When using CTA-HA or CTA/VE-HA for recall, callback services are configured on the appliance or the virtual appliance.

- CCD is used to recall files from Centera to VNX.
- ACD is used to recall files from the cloud to VNX and Unity.
- FPolicy callback is used to recall from all secondary devices to NetApp.

No callback is used for any of the NAS repository types, such as VNX, VNXe, Unity, Windows, Isilon, NetApp, or Data Domain when the source is VNX.

This configuration eliminates a single point of failure for the primary callback service and ensures transparent client access to archived data.

To fulfill requirements for high availability, recall operations can be handled by a group of appliances such as CTAs, CTA-HAs, CTA/VEs, or CTA/VE-HAs.

High availability with VNX or Unity primary storage

For VNX primary storage archived to a Centera, Atmos, or VNX, or Unity files archived to Amazon S3 servers, ECS/S3, Azure, or Virtustream, the Data Movers/NAS servers resolve an HTTP fully qualified domain name (FQDN) to the IP addresses of the appliances. If a Data Mover/NAS server identifies multiple IP addresses mapped to the same FQDN, it will select the first address it finds and attempt to send the recall request. If the IP address is not responsive, the Data Mover/NAS server will select subsequent addresses for the FQDN and attempt to send the recall requests to those addresses.

All recall requests generated by a Data Mover/NAS server when resolving the FQDN are sent to a single appliance even if multiple IP addresses are found. Each Data Mover/NAS server can be configured to send recall requests to a preferred appliance which provides coarse-grained load balancing of recall requests at the Data Mover/NAS server level. For more information, see [Deploying CTA with Unity](#).

To ensure that the Data Mover/NAS server will contact CTAs and CTA-HAs for recall, first create a DNS record for the CTA, then add the IP addresses of all CTA-HAs to that record. The same name points to the IP addresses of the CTA and all HA recall devices. [Configure name resolution for archiving](#) provides details on how to use local hostname resolution or DNS to add the HA appliances.

Run **ccdsetup** on all CTA-HAs or CTA/VE-HAs that will process recall requests from the VNX Data Movers. Run **acdsetup** on all CTA-HAs or CTA/VE-HAs that will process recall requests from the VNX Data Movers or Unity NAS servers. These scripts link multiple appliances to process recall requests from a common set of VNX or Unity Data Movers/NAS servers. [Configuring VNX to Centera or cloud archiving on the CTA](#) provides details on **ccdsetup** and **acdsetup**.

No additional appliances are involved in recall when the CTA or CTA/VE archives data from VNX primary storage to NAS repositories serving as secondary storage. The Data Movers use the CIFS and NFS protocols to recall data directly from secondary storage.

High availability with NetApp primary storage

NetApp filers allow FPolicy clients (such as CTA, CTA-HA, CTA/VE, or CTA/VE-HA) to register for callbacks in response to user access to files with specific attributes. When using the appliances, a callback will be generated when a read/write operation occurs to a file with the CIFS offline bit set.

For NetApp primary storage, multiple appliances can register in the primary or secondary FPolicy groups of the filer. In the event that a registered server becomes unresponsive, it is removed from its group. Recall requests will be sent by the filer in a round-robin fashion to the IP addresses registered in the primary group. If there are no responsive IP addresses in the primary group, then the requests are load-balanced across the servers in the secondary group.

Run **fpsetup** on the CTA-HAs or CTA/VE-HAs that will process recall requests. Use this script to link together multiple appliances that will process recall requests that are sent from a common set of NetApp Filers. Later, when configuring NetApp filers, you will have the option to select specific appliances that will register in the primary and secondary groups. [Configuring NetApp archiving on the CTA](#) provides details on **fpsetup**.

Appliances are always involved in recall when the CTA or CTA/VE archives data from NetApp primary storage to any secondary storage location. NetApp filers do not recall data directly from VNX, Unity, Centera, or NetApp storage.

Note If file encryption is not in use, a single CTA-HA or CTA/VE-HA can provide redundancy for multiple CTAs or CTA/VEs. A single CTA or CTA/VE can have multiple CTA-HAs or CTA/VE-HAs registered to provide redundancy. Do not use a CTA to provide redundancy for another CTA or a CTA/VE to provide redundancy for another CTA/VE.

Network console management for Gen 8 models

In the event that CTA or CTA-HA becomes unresponsive, network console management allows users to control the appliance or reboot the system. For Gen 8 models based on Intel C1U hardware, communication is established through the dedicated management port on the back of the appliance which is labeled MGMT in [Error! Reference source not found.](#)

Note For security purposes, network console management should be enabled on a network with access limited to system administrators only.

Use the local console to configure the CTA for network console management and to change the BMC Web Console Management User password.

1. Connect the power cord and power on the appliance.
2. Using a crossover cable, connect one end to your laptop Ethernet port and the other end to the management port on the back of the appliance. The laptop must be running Java Run Time Environment (JRE) version 6.0 update 10 or later.
3. Configure the laptop with static IP 192.168.1.2 and netmask 255.255.255.0.
4. In a web browser with the popup blocker disabled, type the IP address: **192.168.1.1**. The CTA BMC Web Console appears.
5. To login to the console, type the username and password:
 - Username: **root**
 - Password: **rain**

The integrated BMC Web Console appears

To configure remote console management:

1. Select the **Configuration** tab.
2. Select **IPv4 Network**.
 - a. For LAN Channel, select **Intel(R) RMM**.

- b. Change the IP address to a valid network IP address.
- c. Click **Save**.

Note Once remote console management is configured for the network IP address, local console management is no longer available.

To change the BMC Web Console Management User password:

1. Select the **Configuration** tab.
2. Select **Users**.
 - a. Under User Name, select **root**.
 - b. Click **Modify User**. The Modify User page appears.
3. Type and confirm a password for the root user.

Using CTA CLI commands

To control the appliance from the console:

1. Type the network IP address of the BMC Web Console.
2. Login with the root user and password.
3. Select the **Remote Control** tab.
4. Select **Console Redirection** and click **Launch Console**.
5. A window appears with a message asking to open `viewer.jsp`. Click **OK**. A Java applet window appears.
6. Type the CTA username and password to log in.

[Command line interface](#) provides information on the CTA CLI.

Rebooting the system

To reboot the appliance from the console:

1. Type the network IP address of the BMC Web Console.
2. Login with the root user and password.
3. Select the **Remote Control** tab.
4. Select **Server Power Control**.
5. Select **Power Cycle Server** and click **Perform Action**.

Performing a CD clean install on the appliance

CTA and CTA-HA are shipped with the software installed on the appliance. If needed, a CD clean install installs all necessary packages and binary files. To perform a clean install of CTA/VE or CTA/VE-HA see [Installing the virtual edition](#).

Before starting the installation, check to see if the CTA is connected to another appliance for HA, another CTA, or a stand-alone appliance with a callback daemon running. If so, stop all callback daemons with the following commands:

```
fpolicycallback stop
```

```
atmoscallback stop
celerracallback stop
```

To perform a CD clean install on a CTA or CTA-HA:

1. If using a downloaded ISO image:

- a. Run **md5sum** to verify the image integrity.

The output of the md5sum commands is in the README file that is posted to Online Support, with all the downloads. [Where to Get Help](#) provides information on how to access Online Support.

The ISO files are named:

```
cta-11.0.0-xxx.x86_64.iso
ctaha-11.0.0-xxx.x86_64.iso
```

- b. Burn a CD from the ISO image.

2. Insert the disk for CTA or CTA/HA depending on the CTA appliance type.

3. Boot CTA Appliance from the inserted disk.

4. When the appliance boots, choose **Install/Restore_cta**.

5. Choose **Yes** when the **Destroying ALL data on /dev/sda**, continue prompt appears.

The appropriate packages are installed.

A restart occurs after installation completes and the login prompt appears.

6. Log in with username **root** and password **rain**.

7. Use the Cloud Tiering Appliance setup tool menu that appears to configure the time and network settings.

If the CTA will be configured for VNX to cloud archiving, use FileMover Settings as described in step 7 of [Adding VNX to the CTA configuration](#).

Configure the single set of credentials for recall before running **ccdsetup.sh** or **acdsetup.sh** as described in [Configuring VNX to Centera or cloud archiving on the CTA](#).

Performing a CD clean install on a generic server

When performing a bare metal installation of CTA/CTA-HA using the CD clean install process, we recommend using the Dell Power Edge R630 configurations as the baseline for hardware requirements. For configuration information, see [Cloud Tiering Appliance 12 qualified hardware](#).

To perform a CD clean install of CTA or CTA-HA:

1. If using a downloaded ISO image:

- a. Run **md5sum** to verify the image integrity.

The output of the md5sum commands shown in the README file that is posted to Online Support with all the downloads. [Where to get help](#) provides information on how to access Online Support.

The ISO files are named:

```
cta-11.0.0-xxx.x86_64.iso
ctaha-11.0.0-xxx.x86_64.iso
```

- b. Burn a DVD from the ISO image.
2. Insert disc for CTA or CTA/HA depending on the server you are installing.
3. Boot the server from the inserted disc.
4. When the server boots, choose **Install/Restore_cta**.
5. Choose **Yes**, when the **Destroying ALL data on /dev/sda**, continue prompt appears.

The appropriate packages are installed.

A restart occurs after installation completes and the login prompt appears.

6. Log in with username **root** and password **rain**.
7. Use the Cloud Tiering Appliance setup tool menu that appears to configure the time and network settings.

If the CTA will be configured for VNX to Centera or cloud archiving, use FileMover Settings as described in step 7 of [Adding a VNX to the CTA configuration](#).

Configure the single set of credentials for recall before running **ccdsetup.sh** or **acdsetup.sh** as described in [Configuring VNX to Centera or cloud archiving on the CTA](#).

Adding the CTA serial number

Load the CTA serial number on the appliance with the CLI command:

```
rffm setEmcSerialNumber <EMCSERIALNUMBER>
```

where *EMCSERIALNUMBER* is located on a pull out card on the front of the appliance or printed on the media kit label for CTA/VE versions of the product.

For CTA 12.0 systems, which are bundled with Unity, use the Unity serial number. To obtain the Unity serial number, ssh to the management IP as service and run the `svc_diag` command.

Example output of the `svc_diag` command:

```
===== Now executing basic state =====
* **System Serial Number is: VIRT162479MMTZ**
* System Friendly Host Name is: VIRT162479MMTZ-spa
* Current Software version: upc_Unity_4_4_upcBuilder-
4.1.0.8714268-GNOSIS_DEBUG
* UUID: 42377CA0-6DAF-8E3D-4519-BEAE20EDA933
* Unisphere IP address(es): X.X.X.X fe80::260:16fb:291c:32a7
```

Configuring CTA

Before proceeding with the setup, ensure that you have the following information for each appliance:

- IP address
 - Subnet mask
 - Hostname
 - Default gateway IP
 - DNS server IP — If specifying multiple IP addresses, use spaces to separate IP addresses.
 - CTA serial number from CTA packaging
1. Set up the appliance:
 - a. For a CTA or CTA-HA, connect the keyboard and monitor to the appliance. Connect the power cord and power on the appliance. If a keyboard and monitor are not available, connect a PC or laptop using a crossover cable connected to the MGMT port for Gen 8 models.
 - b. For a CTA/VE, power on the appliance.
 2. Log in to the appliance by using the local keyboard and monitor. Type **root** as the login name. Type **rain** as the password.
The Cloud Tiering Appliance setup tool appears. This tool performs basic setup tasks that are not available through the CTA GUI.
 3. Select **Change Cloud Tiering Appliance Password**, and change the password.
 4. Select **Configure Cloud Tiering Appliance Networking**. The network configuration menu appears.
Use the menu to:
 - a. Configure the CTA network.
 - b. Configure the hostname, domain, and DNS server.
 5. Select **Configure date and time (including timezone)**. The current date/time and a list of current NTP servers are displayed.
Use the menu to set the time zone and date for the appliance and configure an NTP server.
- Note** When archiving to the cloud (for example, Atmos, Amazon S3, or Azure), the CTA clock must be accurate and in sync with the cloud service. Therefore it is required to have an NTP server configured when archiving to these destinations.

Configure the CTA network

Configure the CTA network:

1. Select option 1 from the Network Configuration menu. The Cloud Tiering Appliance Network Setup, Main Menu appears.
On the list of available physical interfaces on the appliance, eth0 appears highlighted. To highlight a different interface, use the up arrow and down arrow keys.
2. With eth0 highlighted, press **Enter**. The configuration menu for the eth0 interface appears:

- a. Use the up arrow and down arrow keys to highlight the IP address field. Press **Enter** and type a new IP address value into the **New Value** column. Press **Enter**.
 - b. Repeat the process to provide the subnet mask, gateway, and MTU settings.
3. When the configuration for this interface is complete, press the left arrow key to exit the eth0 interface configuration.
4. To save the interface configuration, select **Yes** and press **Enter**. Note that the changes are saved, but will not be implemented until the Cloud Tiering Appliance Network Setup menu is exited.
5. Press the left arrow key to exit from the **Cloud Tiering Appliance Network Setup, Main Menu**. When prompted, select **Yes** to save your changes.

Configure the hostname, domain, and DNS server

Configure the hostname, domain, and DNS servers:

1. Select option 2 from the Network Configuration Menu. The following menu appears:

```
EMC Cloud Tiering Appliance Setup Tool (Configure Hostname,
Domain and DNS Server(s))
Hostname                = fm
Domain                  =
DNS Server               =
```

```
Do you want to change the configuration [Y/N]?
```

2. Type **Y**. Use the menu to configure the hostname, domain, and DNS servers.

The new hostname, domain, and DNS server information is summarized after all the changes are entered, and you are given the ability to accept or make further changes to these settings. To keep the new settings and return to the network configuration menu, press **Enter**.

3. Verify that the network configuration has been saved and network connectivity can be established properly.

Configure iSCSI IQN targets (Unity block archiving and restore only)

To use the block archiving and block restore features, the CTA need to have iSCSI IQN targets configured.

1. Select option 3 from the Cloud Tiering Appliance setup tool.
2. Select option 3 from the Network Configuration Menu. The following output appears:

```
EMC Cloud Tiering Appliance Setup Tool (Configure iSCSI IQN(s)
for block)
```

```
Auto generating iSCSI IQN for the first time ... iqn.1992-
04.com.emc:cta:fm-0
```

Update the hosts file

If `/etc/hosts` on the CTA requires customization, do not manually edit the file directly.

To customize the `/etc/hosts` file:

1. Edit the `/etc/hosts.local` file.
2. Select option 4 from the Network Configuration Menu to import changes from `/etc/hosts.local` into `/etc/hosts`.

Note If `/etc/hosts` is customized, maintain `/etc/hosts.local` to preserve those customizations and restore `/etc/hosts` if needed.

File Encryption

Archiving to cloud services such as Atmos or Amazon S3 is a means of archiving at rest data. However, there is a risk of compromising sensitive data stored in the cloud unless file encryption is used.

To archive to the cloud with file encryption, a primary CTA or CTA/VE must be deployed with a CTA-HA or CTA/VE-HA. The HA appliance is configured with a keystore that replicates the keystore on the primary appliance. As long as a key is active and the keystores on the HA and primary appliances are in sync, policy-based file level encryption is enabled.

Note If file encryption is active, a single CTA-HA or CTA/VE-HA cannot provide redundancy for multiple CTAs or CTA/VEs.

Enabling keystore replication

The CTA supports encryption of data that is archived to a cloud platform, such as Atmos or Amazon S3. To enable encryption on the CTA, you must enable replication of the encryption keystore on the CTA-HA or CTA/VE-HA. The primary keystore is located on the CTA, and the replica is created on the HA. Before enabling encryption, ensure that NTP is configured on the CTA or CTA/VE and on the HA appliance.

To enable keystore replication on the CTA-HA or CTA/VE-HA, type:

```
krdsetup init_rffm
```

This stores the IP address of the main CTA or CTA/VE and the timestamp of the most recently replicated keystore in the configuration file:

```
/opt/rainfinity/filemanagement/conf/krd.xml
```

Then it starts the keystore replication daemon that automatically keeps the keystore on the high availability appliance in sync with the primary appliance.

Configuring and using file encryption

Once keystore replication has been enabled, the file encryption key can be generated using the GUI on the primary appliance by selecting **Configuration > File Encryption Key**.

To archive to a cloud service with file encryption on, click **Encrypted** when adding a rule to an archive policy. If no key is active, **Encryption** is grayed out.

To migrate secondary storage tier data to a cloud service with file encryption on, click **Encryption** when selecting the destination for the repository migration task.

If an active key has not been replicated to the CTA-HA or CTA/VE-HA, an error appears indicating that file encryption is not enabled.

Note If a system administrator changes the CTA hostname or configures a network interface, attempts to generate an encryption key on the CTA or synchronize a keystore on the CTA-HA might fail. To work around this problem, reboot the CTA.

Graphical user interface

To access the graphical user interface from a web browser:

1. In the navigation field of the web browser, type the IP address of the CTA. The CTA GUI appears.
2. Type the username and password for the default account which are:
 - a. Username: **admin**
 - b. Password: **rain**

The selections appear as follows:

- c. **Dashboard** — Top level view displays summaries of Common Tasks, Alerts, and Reports.
- d. **Schedule** — Displays a list of scheduled tasks that are currently being processed and the status of each task.
- e. **Archives** — Displays lists of archived files, LUNs, and consistency groups, as well as an archived file report. Also provides a search option to find archived files, recover stub files, and delete orphan files.
- f. **Policies** — Provides options that apply to creating and managing policies, including:
 - A list of policies, matching expressions, and NAS destinations.
 - Create new policy.
 - Create new matching expression.
 - Create new NAS destination.
- g. **Configuration** — Provides configuration of servers, repositories, and settings for servers.
- h. **System** — Provides system configuration of users, passwords, logs, alerts, and security.

Command line interface

As an alternative to the GUI, you can use a command line interface to send commands to the filemanagement daemon.

To log in to the CLI by using SSH, the default username and password are:

- Username: **root**
- Password: **rain**

The most commonly used commands are:

- **fmsupportdump** — Creates a dump of the appliance's current state for technical support. Run this command from `/var`.
- **rffm** — Configures the appliance and issues all commands that the GUI interface supports. To see a list of all commands available, type **rffm --help** or to view the man page for more detailed help, type **man rffm**.
- **fmbackup/fmrestore** — Backs up and restores the configuration as described in [Backing up the configuration](#).
- **rssysstat** — Displays statistics about the CTA.

Man pages for the command line tools are stored in the software installation directory. To access the man pages, type **man *command_name***, as in **man rssysstat**.

Chapter 4 Deploying the Cloud Tiering Appliance

This chapter contains the following sections:

Deployment process for archiving	42
Supported platforms.....	44
Deploying CTA with VNX.....	45
Deploying CTA with NetApp.....	57
Deploying CTA with a VNXe	64
Deploying CTA with Unity	65
Deploying CTA with Virtustream.....	70
Deploying CTA with a Windows server.....	71
Deploying CTA with Isilon.....	74
Deploying the CTA with Data Domain.....	80
Configuring a NAS-based repository	81
Deploying CTA with Amazon S3	82
Deploying CTA with Atmos.....	83
Deploying CTA with Azure.....	85
Deploying CTA with Centera	86

Deployment process for archiving

Figure 6 illustrates how the Cloud Tiering Appliance (CTA) or Cloud Tiering Appliance/VE is deployed for archiving.

Note The first time you install and configure the CTA system, click the **Add Task Wizard** in the CTA dashboard to guide you through the steps of configuring your system, archiving, or migration task.

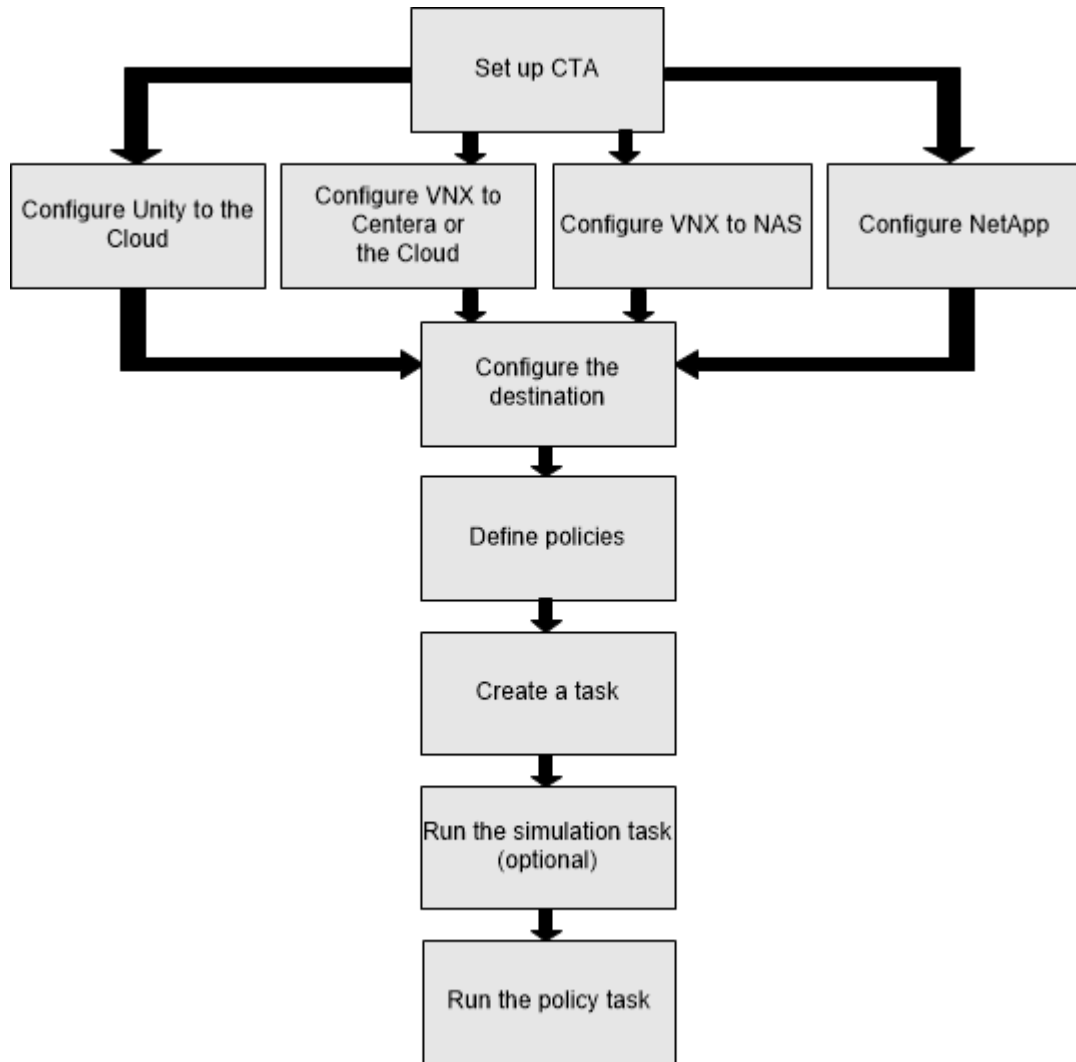


Figure 6. Cloud Tiering Appliance deployment process

Note For VNX to NAS configuration, NAS repositories can be VNX, VNXe, Windows, Isilon, NetApp, or Data Domain.

For details on deploying the CTA or CTA/VE in various environments, see the following sections. Steps in the five boxes at the bottom of the flowchart are performed by using the CTA GUI. The online help describes these steps in detail.

Step 1: Set up CTA

1. Install the CTA/VE, if applicable.
2. Add the CTA serial number.
3. Configure CTA networking.
4. Enable the keystore on the HA for file encryption to the cloud, if applicable.

Step 2: Configure the archive environment

Unity to the Cloud

1. Configure Unity properties.
2. Initialize recall services.
3. Configure name resolution for recall.
4. Configure Unity settings under **Common API Settings**.
5. Configure DHSM.

VNX to Centera or the Cloud

1. Configure VNX properties.
2. Initialize recall services.
3. Configure name resolution for recall.
4. Configure VNX settings under **Common API Settings**.

VNX to NAS

1. Configure VNX properties.
2. Configure VNX settings under **Common API Settings**.
3. Configure DHSM.

NetApp

1. Configure ONTAPI.
2. Configure vFilers, if applicable.
3. Initialize recall services.
4. Configure NetApp properties.

Step 3: Configure the destination

1. Configure NAS repositories.
2. Configure non-NAS repositories.

Step 4: Define policies

1. Create matching expressions and archive destinations.
2. Specify policy type, retention, delayed stubbing, stub retention (as applicable).

Step 5: Create a task

1. Create an archive, delete, or auxiliary task.
2. Select the source, as applicable.

Step 6: Run the simulation task (optional)

1. Select **Run Simulation Now**.
2. Collect real-time results in CTA.
3. Review policy efficacy against real-time results.

Step 7: Run the policy task

1. Determine optimal task scheduling.
2. Select archive conditions or start times (as applicable).
3. Monitor archiving activity for errors.

Supported platforms

Platform support varies depending on the task performed.

Table 6 shows platform support for archiving.

Table 6. Supported platform matrix for archiving

Primary tier	Secondary tier
VNX NetApp	VNX VNXe Unity Windows Isilon NetApp Data Domain Centera Atmos Amazon S3 Microsoft Azure
Unity	Amazon S3 ECS/S3 Microsoft Azure Virtustream

VNX, VNXe, Unity, Windows, Isilon, NetApp, and Data Domain are collectively referred to as NAS repositories, as described in [Configuring a NAS-based repository](#).

[Table 7](#) shows platform support for file migration.

Table 7. Supported platform matrix for file migration

Source	Destination
VNX	VNX VNXe Isilon
NetApp	VNX VNXe Isilon
Windows	VNX VNXe

Note Stub retention for Windows is not supported. The destination must be offline until migration is complete. Otherwise there is a risk of data loss or damage.

CTA does not support server message block (SMB) signing. For CTA to authenticate properly with a CIFS file server, SMB signing must be disabled within the Active Directory domain. If SMB signing is required and enabled for client-server communications, CTA cannot browse CIFS shares on the file server or execute any CIFS archiving tasks.

Deploying CTA with VNX

To use CTA with a VNX Data Mover, first perform configuration steps on the CTA, and then on the VNX Control Station.

VNX supports DART version 7.0 or later.

With VNX as a source, the deployment requirements vary with task type.

Table 8. VNX deployment checklist

Configure:	File Migration	Archiving
System prerequisites	✓	✓
VNX properties	including NDMP	NDMP not required
Callback daemon: CCD, ACD	N/A	✓
Name resolution	N/A	✓
VNX Control Station	✓	✓

Configure:	File Migration	Archiving
DHSM	✓	✓

Prerequisites for using VNX as a file migration source or destination

To migrate files to or from a VNX Data Mover, ensure that the following conditions are satisfied:

- CTA requires access through both CIFS and NFS to complete a file migration (KB Article Number 000464016).
- For NFS, CTA has root access with read/write permission on source and destination exports.
- CIFS user has local administrator access on all source and destination shares and is part of the Backup Operators group.
- If migrating files from a source that is shared and exported, CTA has root access with read/write permission to both shares and exports.
- If migrating dedupe files from VNX to Isilon, set the backup data threshold for VNX File Deduplication and Compression on a Data Mover to 0 before migration:

```
$fs_dedupe -default -set <movername> -backup_data_threshold 0
```

Where *movername* is the name of the Data Mover.

After migration, set `fs_dedupe` back to its original value.

Note Do not run any other backup applications on the Data Mover while migration is running because this is a global parameter that will cause all other backup applications to inflate dedup files.

- Snapshare snapshots are enabled on the source.
- There is a system user for the XML API setup on the VNX Control Station. [Configuring automatically created DHSM connections for file migration or archiving](#) describes how to set up the user.

The system user for the XML API is not needed for VNX as a destination.

- For every server, properties are configured on the CTA for:
 - Control Station IP address for VNX as source
 - NDMP username and password

[Adding VNX to the CTA configuration](#) provides details on configuring a VNX server.

Prerequisites for VNX as an archive source

To archive data from a VNX Data Mover, the CTA requires access to the Common API (TCP port 5080).

To archive NFS data, the CTA needs the following:

- Mount v3 RPC service

- NFS v3 RPC service
- NLM v4 RPC service
- Root access with read/write export permissions for all NFS data that will be archived.

To archive CIFS data, the CTA needs SMB over NetBIOS (TCP port 139).

Direct command line access to the VNX Control Station is not used by the CTA.

When configuring VNX Data Mover properties on the CTA, plan to provide:

- Credentials for a Common API user. This single set of credentials is used for both archive and recall.
- Control Station IP address.
- (For CIFS archiving only) The NetBIOS name of the file server or Kerberos FQDN.
- (For CIFS archiving only) Credentials for Windows domain user.

When archiving to a NAS repository, virus scanning can have the following effects:

- Virus scanning on the destination will negatively impact CTA performance and should be disabled.
- Virus scanning on the destination can also delete files, leaving stub files on primary storage that cannot be opened.
- Virus scanning can be enabled on the source, but scanning on read can cause time out conditions with no reply errors on the source if CTA is the first application to access a file.

To disable virus scanning on the VNX, mount the file system with the *noscan* option.

Adding VNX to the CTA configuration

Note The VNX Properties page is not displayed.

1. In the navigation field of the web browser, type the IP address of the CTA. The CTA GUI appears.
2. Type the username and password for the default account which are:
 - a. Username: **admin**
 - b. Password: **rain**
3. Select **Configuration** and **Server**. The **Server List** appears. Click **Add**.
4. On the **Create Server** page that appears, select **VNX**. Click **OK**.
5. For each step in the File Server wizard, provide the required information indicated with a red asterisk and click **Next**.

Step	Description
Basic File Server Information	<p>Type the NetBIOS server name.</p> <p>To identify the VNX fileserver as a VDM (Virtual Data Mover), select the File Server Is VDM checkbox. Then do one of the following:</p> <ul style="list-style-type: none"> • When the File Server Is VDM checkbox is enabled, ensure that the NFS exports are created on the physical Data Mover and not on the VDM. In this scenario, the interface used by the VDM should not be attached to the VDM. • When the File Server Is VDM checkbox is disabled, ensure that the NFS exports are created on the VDM and not on the physical Data Mover. In this scenario the interface used by the VDM should be attached to the VDM.
IP Addresses	<p>Type the Data Mover IP address:</p> <p>When editing an existing server, click Update to retrieve the IP address from the DNS that is based on the server name.</p> <p>To specify an additional IP address, click Add.</p> <p>To delete an existing IP address, select an IP and click Delete.</p> <p>The Control Station IP is required for file migration transactions and is used to create a snapshot of the source. For archiving, this allows CTA to automatically perform some prerequisite tasks for archiving as described in Configuring automatically created DHSM connections If this field is empty, the CTA takes no action and the prerequisite tasks must be performed manually.</p>
CIFS Specific Settings	<p>This is the Windows domain user to be used by the appliance. The domain user must be a member of the local Administrator group and part of the Backup Operators group on the VNX system. Windows domain user provides more information.</p> <p>Specify either a NetBIOS domain or a Kerberos FQDN. The Kerberos FQDN is a CTA server configuration setting, and if selected, the same FQDN applies to all servers. Before selecting the Kerberos FQDN option, follow steps in Configuring Windows for Kerberos to configure your Fully Qualified Domain on the CTA.</p>

Step	Description
<p>NDMP Specific Settings</p>	<p>For file migration, type the username and password for an NDMP user on the source and destination servers. By default, the port for NDMP traffic is 10000.</p> <p>Note NDMP specific settings do not apply for Windows to VNX migration.</p>
<p>VNX system as Source</p>	<p>This option configures the CTA to archive data from the VNX Data Mover. If multiple appliances are connected to the same VNX Data Mover, only one appliance should be configured with the VNX as Source. This option is required only if the VNX is serving as a source for archiving.</p> <p>Multiple appliances may be configured to archive data from a single VNX Data Mover, but only one CTA or CTA/VE should be used to archive data from a single filesystem.</p>
<p>CCD DNS Name</p> <p>ACD DNS Name</p> <p>This step appears if VNX as source is yes.</p>	<p>The CCD DNS Name is required if archiving to a Centera. For the CCD DNS name, type the FQDN of the VNX Callback DNS entry. Note that the FQDN is case-sensitive.</p> <p>The ACD DNS Name is required if archiving to a cloud storage server. For the ACD DNS name, type the FQDN of the Cloud Callback DNS entry. Note that the FQDN is case-sensitive.</p> <p>The DNS names for the VNX Callback agent and Cloud Callback agent must be distinct. They cannot be the same. If a callback name is not specified, archive tasks to the destination will not start.</p>
<p>Directory Exclusion List</p> <p>This step appears VNX as source is yes.</p>	<p>These are the directories to exclude for all tasks that use scanning. Migration tasks do not scan, so this setting does not apply to file migration. The CTA ignores all system directories such as, etc, lost+found, and ckpt by default.</p> <p>Verify that stub files are not in the excluded directories. CTA will not access the excluded directories and the stub files will become orphans.</p>

6. Review the Summary and click **Finish** to define the VNX file server.
7. Select Configuration and **Common API Settings**. The **Common API Settings** page appears.

Type the username and password for the FileMover Credentials.. The system uses this username and password to create an HTTP connection by using XML API. This

same username and password are used when creating the FileMover user in [Prerequisite tasks on the VNX Control Station for file migration or archiving](#)

The FileMover setting applies to all VNX file servers configured on the CTA. If all VNX file servers are archive targets only, the FileMover setting is not required.

Note Because CTA migrates data using NDMP, it can migrate both NFS and CIFS metadata in a single protocol migration, provided that the NDMP format does not require conversion. For example, when both the source and destination are VNX systems, no conversion is necessary. However to migrate NFS and CIFS metadata between systems that require NDMP format conversion, you must both *export and share* the file system. For example, CTA cannot migrate CIFS data in an NFS migration involving systems other than VNX. This is because after CTA converts the NDMP data stream, it can only migrate CIFS metadata using CIFS protocol.

Configure name resolution for archiving

When the VNX Data Mover needs to establish a connection to the appliance to recall data from a Centera or a cloud service, it tries to resolve the FQDN from the HTTP DHSM connection in its local hosts file. If it cannot be resolved locally, the Data Mover will use DNS.

Note These instructions pertain only to configuring the hosts file on the VNX. [Configuring CTA](#) provides information on how to configure the CTA hosts file.

- To use DNS:
 - a. Create a DNS entry for the callback daemon that points to the appliance.
As a best practice, the CTA hostname and CCD hostname should be different. Using the same hostname for both can raise issues when trying to login to the appliance.
 - b. Create multiple entries by the same name for each callback appliance.
 - c. For each entry that is created, select the checkbox for **Create associated pointer (PTR) record** to ensure that it will be included in the Reverse Lookup Zones list.

Note The VNX FileMover supports DNS HA failover. If the DNS server resolves the callback daemon hostname to multiple IP addresses, the VNX FileMover transparently switches to the server at the next available IP address.

- DNS is the preferred method for name resolution. If using DNS is not feasible, use local hostname resolution:
 - a. Log in to the VNX Control Station as root and mount the Data Mover to edit the local hosts file with **vi**:

```
mount server_2:/ /mnt/source
cd /mnt/source/.etc
vi hosts
```

where *server_2* is the name of your VNX Data Mover.

- b. Edit the hosts file to add one line for each appliance, similar to the following example:

```

10.0.0.1 <rainccd.domain> # CCD on CTA-HA
10.0.0.2 <rainccd.domain> # CCD on CTA
10.0.0.3 <rainccd.domain> # CCD on CTA/VE
10.0.0.1 <rainacd.domain> # ACD on CTA-HA
10.0.0.2 <rainacd.domain> # ACD on CTA
10.0.0.3 <rainacd.domain> # ACD on CTA/VE

```

where:

- *rainccd.domain* is the FQDN that will be used to create the HTTP DHSM connection described in Adding VNX to the CTA configuration (see description for CCD DNS Name).
 - *rainacd.domain* is the FQDN that will be used to create the HTTP DHSM connection described in Adding VNX to the CTA configuration (see description for ACD DNS Name).
- c. Save the file and confirm that the VNX Control Station is not mounted to the Data Mover:

```

cd ~
umount /mnt/source

```

Configuring VNX to Centera or cloud archiving on the CTA

To archive from a VNX to a Centera or cloud storage, configure the callback service so that the CTA is in the recall path.

Configure the callback service to recall from Centera

To configure recall from the Centera:

1. From the console on the CTA, log in as root.
2. Type **!** to escape to the command line and type:

```
ccdsetup init_rffm
```

3. Type **n** when the following message appears:

```

By default the Celerra Callback Daemon will connect to the
CTA service on the local machine.
Do you wish to configure another CTA? (y/n)

```

4. If there is a secondary callback agent such as a CTA-HA, log in on that agent as root, and repeat steps 2 and 3. In step 3, type **y** to provide the IP address and the root password of the CTA.

If an invalid IP address is provided, the CCD.log file located in `/var/log/rainfinity/filemanagement/recall` will fill with errors to indicate that there was no response from the primary agent. To correct the problem, repeat steps 2 through 4 of this procedure.

Configure the callback service to recall from the cloud

To configure recall from the cloud service such as Atmos or Amazon S3:

1. From the console on the CTA, log in as root.
2. Type **!** to escape to the command line and type:

```
acdsetup init_rffm
```

3. Type **n** when the following message appears:

```
By default the Cloud Callback Daemon will connect to the CTA
service on the local machine.
Do you wish to configure another CTA? (y/n)
```

4. If there is a secondary callback agent such as a CTA-HA, log in on that agent as root, and repeat steps 2 and 3. In step 3, type **y** to provide the IP address and root password of the primary callback agent.

If an invalid IP address is provided, the ACD.log file located in /var/log/rainfinity/filemanagement/recall will fill with errors to indicate that there was no response from the primary agent. To correct the problem, repeat steps 2 through 4 of this procedure.

Prerequisite tasks on the VNX Control Station for file migration or archiving

If a VNX has not been configured as a source for file migration or archiving, perform the following steps:

1. Enable filename translation on the VNX Control Station.

The CTA, CTA-HA, or CTA/VE expects that all filenames are derived from the VNX Network Server in UTF-8 format. To preserve filenames correctly:

- a. Log in to the VNX Control Station as **nasadmin**.
- b. Use a text editor to open the file: **/nas/site/locale/xlt.cfg**.
- c. Locate the last line of the file. Typically the last line appears as:

```
:::8859-1.txt: Any thing that didn't match above will
be assumed to be latin-1
```

Add the following line immediately above the last line:

```
::CTA_IP_ADDR::: CTA requires no translation (UTF-8)
```

where *CTA_IP_ADDR* is the IP address of your appliance. Add an entry for each CTA. HA appliances are not added.

For example, if the IP address is 10.10.18.1, type:

```
::10.10.18.1::: CTA requires no translation (UTF-8)
```

To specify a subnet, type:

```
::10.10.18.0,255.255.255.0::: CTA requires no
translation (UTF-8)
```

- d. To update the configuration, type:
`/nas/sbin/uc_config -update xlt.cfg`

- e. To verify the new configuration, type:
`/nas/sbin/uc_config -verify CTA_IP_ADDR -mover ALL`

where *CTA_IP_ADDR* is the IP address of your appliance. Output will appear in the format:

```
server_name : CTA_IP_ADDR is UTF-8
```

2. Create the FileMover user. Log in to the VNX Control Station CLI as root and type the command:

```
/nas/sbin/server_user <data_mover> -add -md5 -passwd <user>
```

For example: `/nas/sbin/server_user server_2 -add -md5 -passwd rffm`

The username and password must match the FileMover Credentials under Common API settings configured on the CTA in Adding VNX to the CTA configuration.

3. Allow the IP addresses of the CTA or CTA/VE to open connections to the FileMover interface. While logged in to the VNX Control Station as an administrator (such as “nasadmin”), run the following command for all IP addresses of all appliances that will perform archiving or service recall requests for the Data Mover:

```
server_http <data_mover> -append dhsm -users <user> -hosts <ip_address>
```

For example: `server_http server_2 -append dhsm -users rffm -hosts 192.168.0.100,192.168.0.101,<CTA_IP_address>`

Note A single VNX Data Mover can be configured as an archiving source with multiple appliances, but more than one CTA or CTA/VE should never be used to archive data from a single filesystem.

4. Enable DHSM (FileMover) for the Data Mover. To enable DHSM and keep it enabled if the Data Mover reboots, run the following command once:

```
server_http <data_mover> -service dhsm -start
```

5. Enable DHSM for specific filesystems that will be used as archiving sources. To enable DHSM and keep it enabled if the Data Mover reboots, run the following command once per filesystem.

```
fs_dhsm -modify <primary_fs> -state enabled
```

For example: `fs_dhsm -modify fileSystem1 -state enabled`

6. Ensure that the DHSM offline attribute is enabled for filesystems that will be used for archiving. To verify that the offline attribute is on, run the command:

```
fs_dhsm -i <fs_name> | grep 'offline attr'
```

- If the offline attribute is on, the following line will appear:

```
offline attr = on
```

- If the offline attribute is off, turn it on with the command:

```
fs_dhsm -m <fs_name> -offline_attr on
```

Caution Once the offline attribute is set to on, it must remain on. Disabling the offline attributes after executing archive tasks can lead to data loss.

Create one or more connections from the Data Mover to the secondary storage locations for each filesystem that will be archived. Each CIFS or NFS repository used to store archived data needs to be configured as a DHSM connection for the VNX filesystem. If data will be archived to a Centera or cloud storage service, a DHSM

connection that uses the HTTP protocol must be configured for that filesystem to the CTA or callback FQDN.

Configuring automatically created DHSM connections for file migration or archiving

The CTA or CTA/VE can automatically create DHSM connections for VNX systems running DART 7.0.

To configure this feature, perform the following steps as root user on the VNX and the appliance:

Run the commands on the VNX as root user, and not as nasadmin.

1. Check to see if the XML API server is running on the VNX, type:

```
ps -ef | grep start_xml_api_server | grep -v grep
```

The following example shows a server that is already running:

```
[root@celerra01 sbin]# ps -ef | grep start_xml_api_server |  
grep -v grep  
root      14821  3226   0 15:41 ?                00:00:00 /bin/sh  
/nas/sbin/start_xml_api_server
```

- If it is running, restart the server by typing:

```
/nas/sbin/hup_api
```

- If it is not running, start the server by typing:

```
/nas/sbin/start_xml_api_server
```

If the server fails to start or restart:

- a. Delete the file `/nas/api/exit_now`.
- b. Delete the file `/nas/api/api_retry`.
- c. Repeat the process to check if the server is running and to start it.

If the XML API server still fails to start, contact VNX support.

2. Create a new system user for the XML API. Use the API GUI on the VNX Control Station.

- For a new user on a VNX running DART 7.x:

- a Log into Unisphere as root. For Scope, select **Global**.

- b Select: Settings > Security > Local Users > Create. The New User screen appears.

- For a new user on a VNX running DART 7.1.x:

- a Log into Unisphere as sysadmin. For scope, select Global.

- b Select: Settings > Security > User Management > Local Users for File > Create. The New User screen appears.

Use the same username and password that was defined for the Common API user in Prerequisite tasks on the VNX Control Station for file migration or archiving.

When a user password is updated or changed on the VNX Control Station, update the FileMover Credentials under Common API settings for the VNX Properties on the appliance as in Adding VNX to the CTA configuration and update the DHSM connection password with the command:

```
fs_dhsm -connection <primary_fs> -modify <cid> -password <new_password>
```

3. Define VNX Data Mover properties on the CTA or CTA/VE. Adding VNX to the CTA configuration describes the following properties in greater detail:
 - a. For Control Station, provide the Control Station IPs.
 - b. For FileMover Settings, type the username and password that were created for the new system user.

If DHSM connections do not exist, the CTA automatically creates the connections before running each archiving task.

Configuring manually created DHSM connections for file migration or archiving

DHSM connections must be created manually if any of the following conditions apply:

- DART 6.0 is being used with an NFS-exported filesystem on a VDM.
- The CTA is not being used to automatically create DHSM connections.

Commands to create the connection for different archiving scenarios are provided as follows:

- When archiving CIFS data to NAS, you archive to a CIFS repository configured on the appliance.

Create a connection to each CIFS repository that will hold archived data. This setting applies to any repository that is part of a multi-tier destination. Log in to the CLI of the VNX Control Station and type the command:

```
fs_dhsm -connection <primary_fs> -create -type cifs -admin '<fqdn>\<domain_administrator>' -secondary '\\<fqdn_of_secondary_server>\<repository_path>' -local_server <local_cifs_server>
```

For example: `fs_dhsm -connection fileSystem1 -create -type cifs -admin 'mydomain.prv\administrator' -secondary '\\oldServer.mydomain.prv\dir1\dir2' -local_server ns80dm1`

Note Use the apostrophe instead of quotation marks to encapsulate the CIFS administrative username and UNC path of the secondary storage location.

- When archiving NFS data to NAS, you archive to an NFS repository configured on the appliance.

Create a connection to each NFS repository that will hold archived data. Log in to the CLI of the VNX Control Station, and type the command:

```
fs_dhsm -connection <primary_fs> -create -type nfsv3 -secondary '<fqdn_of_secondary_server>:/<repository_path>' -proto TCP -useRootCred True
```

For example: `fs_dhsm -connection fileSystem1 -create -type nfsv3 -secondary 'oldServer.mydomain.prv:/dir1/dir2' -proto TCP -useRootCred True`

- When archiving any type of data to an Centera CAS or cloud storage server, recall requests will flow from the Data Mover to CTA, CTA-HA, or CTA/VE.

- c. To create the connection for an Centera, log in to the CLI of the VNX Control Station, and type the command:

```
fs_dhsm -connection <primary_fs> -create -type http -
secondary 'http://<fqdn for CCD>/fmroot' -httpPort 8000 -
cgi n -user <user> -timeout 60
```

For example: `fs_dhsm -connection fileSystem1 -create -type http -secondary 'http://CCD01.mydomain.prv/fmroot' -httpPort 8000 -cgi n -user rffm - timeout 60`

When prompted, type a password for the file mover settings username .

- d. To create the connection for a cloud storage server, log in to the CLI of the VNX Control Station and type the command:

```
fs_dhsm -connection <primary_fs> -create -type http -
secondary 'http://<fqdn for ACD>/fmroot' -httpPort 9000 -
cgi n -user <user> -timeout 60
```

For example: `fs_dhsm -connection fileSystem1 -create -type http -secondary 'http://ACD01.mydomain.prv/fmroot' -httpPort 9000 -cgi n -user rffm -timeout 60`

When prompted, type a password for the 'rffm' user.

These same settings are used in Adding VNX to the CTA configuration.

- e. The FQDN for the callback daemon is used for CCD DNS Name and ACD DNS Name. For more information, see [Adding VNX to the CTA configuration](#). The FQDN must be distinct even if the VNX and cloud callback daemons are running on the same appliance.
 - f. The same user and password credentials are used for FileMover Credentials under Common API Settings in Adding VNX to the CTA configuration.

Regardless of the type of connection (CIFS, NFS, or HTTP), the target of a connection should be specified as a hostname or FQDN in the command:

```
fs_dhsm -connection <primary_fs> -create
```

- When a VNX Data Mover needs to establish a connection to secondary storage, it first attempts to resolve the hostname in the local hosts file. If the name cannot be resolved locally, the Data Mover then issues a DNS query.
- When archiving to NAS from a VNX, a DNS record is required to resolve the FQDN of the secondary storage server to IP addresses if the local hostname resolution of the VNX is not going to be used. A PTR record (reverse DNS) is also required to map the IP addresses of the secondary storage server to the FQDN.

Note The VNX File Level Retention (FLR) enabled filesystems cannot be used as an archiving source.

Deploying CTA with NetApp

To use the CTA with a NetApp filer, first configure the filer, and then configure the appliance.

With NetApp as a source, the deployment requirements vary with task type.

Table 9. NetApp deployment checklist

Configure:	File Migration	Archiving
System prerequisites	✓	✓
NetApp properties	including NDMP	NDMP not required
NetApp vFiler prerequisites	if applicable + ndmpd	if applicable
Callback daemon: FCD	N/A	✓

Prerequisites for using NetApp as a file migration source

To migrate files from a NetApp, ensure that the following conditions are satisfied:

- For NFS, CTA has root access with read/write permission on source exports.
- CIFS user has local administrator access on all source shares and exports.
- If migrating files from a source that is shared and exported, CTA has root access with read/write permission to both shares and exports.
- Create unicode and convert unicode options on the volume are enabled from the NetApp console with the **vol options** command.
- NDMP is enabled from the NetApp console by typing:

```
ndmpd on
options ndmpd.authtype plaintext
```

- For NetApp filers running Data ONTAP 7.2 or later, disable duplicate session detection by typing:

```
options cifs.client.dup-detection off
```

Note CTA works best with duplicate session detection disabled. However, if the network environment requires duplicate session detection, the option can be set to **name**.

- NDMP specific settings are specified for every NetApp filer configured on the CTA.

[Adding NetApp filer to the CTA configuration](#) provides details on configuring a NetApp server.

Prerequisites for using NetApp as an archiving source

To archive any data from a NetApp filer, the CTA or CTA/VE requires access to:

- SMB over NetBIOS (TCP port 139)

- ONTAPI (TCP port 80)

In addition, to archive NFS data, the CTA or CTA/VE will require the following:

- Portmap v2 RPC service (TCP port 111)
- Mount v3 RPC service
- NFS v3 RPC service
- NLM v4 RPC service
- Root and read/write export permissions for all NFS data that will be archived
- inode to pathname mapping is enabled for NFS clients that will access stub files

When configuring a NetApp filer on the CTA or CTA/VE, plan to provide:

- All IP addresses that are used by the filer
- Credentials for local administrator access through both CIFS and ONTAPI
- The NetBIOS name of the filer

Note If a NetApp filer leverages its vScan interface for virus scanning, the IP addresses of the vScan servers must be configured on the appliance as excluded clients on the NetApp FPolicy Special Clients configuration page in the GUI. This allows the virus scanner to scan the stub file upon a recall event. Failure to configure excluded clients properly will lead to recall failures when vScan is used in conjunction with FPolicy.

Local administrator's credentials are required because ONTAPI access is used to send API calls. If a user other than root is specified, the following option must be set:

```
options httpd.admin.hostsequiv.enable on
```

Ensure that the appliance hostname:

- Can be resolved to its IP addresses in the local `/etc/hosts` file of the NetApp filer.
- Maps to a user with privileges to access the ONTAPI interface in the `/etc/hosts.equiv` file on the filer.

Additional configuration prerequisites vary, depending upon the existing network environment.

If using a vFiler, precede all options commands with **vfiler run <vfiler_name>**.

- Set the following NetApp options by typing:

```
options fpolicy.enable on
options cifs.nfs_root_ignore_acl on
options httpd.admin.hostsequiv.enable on
options httpd.admin.enable on
options nfs.tcp.enable on
options nfs.udp.enable on
options cifs.netbios_over_tcp.enable on
```

- If using NFS to archive CIFS shares, type:

```
options wafl.nt_admin_priv_map_to_root on
```

And add the following line in the /etc/usermap.cfg file:

```
<domain_name>\<cifs_user> == root
```

For example: domainABC\userABC == root

- For NetApp filers running Data ONTAP 7.2 or later, disable duplicate session detection by tping:

```
options cifs.client.dup-detection off
```

Note CTA works best with duplicate session detection disabled. However, if the network environment requires duplicate session detection, the option can be set to **name**.

- To properly support stub files, NetApp FPolicy requires a particular CIFS offline bit attribute on the stub files:
 - g. The CIFS protocol must be enabled on the NetApp filer to archive either CIFS or NFS datasets. An active CIFS license must be installed on all file servers that are archive sources.
 - h. NFS-only exports must be shared as well.
- To properly recall stub files, FPolicy must be enabled (**options fpolicy.enable on**) and rfpolicy must be the only screen policy registered for reads and writes. If a policy that monitors stub files on the NetApp filer was previously installed, manually delete it.
- To configure NFS archiving, perform the following steps on the NFS-only source directories:
 - i. Create a share at the qtree or volume level for qtree sources.
 - j. Create a share at the volume level for non-qtrees, that is, those not part of any qtree.
 - k. Add access to only the CTA user.

Note The CTA does not support name clashes on qtrees. For example, QTREE1 against qtree1.

- Configure the NetApp source for UTF-8 language encoding with the command:

```
vol lang <vol> <language_encoding>
```

where *language_encoding* ends with .UTF-8, for example en_US.UTF-8.

- If a virus scanning or backup program runs on your NetApp filer, type the IP address of the machines that host these applications as Excluded Clients on the NetApp FPolicy Special Clients configuration GUI page.
- When archiving to a NAS repository, virus scanning on the destination will negatively impact CTA performance and should be disabled. Virus scanning on the destination could also delete files, leaving stub files on primary storage that cannot be opened. Virus scanning can be enabled on the source

when needed, however scanning on read should not be enabled because it can cause a timeout and no reply errors on the source during archiving.

vFiler configuration

To use NetApp vFilers with the CTA, ensure that:

- The CTA can access both the vFiler and the hosting NetApp filer.
- vFilers and main filers are in IP spaces that can reach each other.

If NetApp vFilers are being used for file migration, the root password is not the same as the NDMP password. Before configuring NDMP for the vFiler on the CTA, use the NetApp CLI to retrieve the NDMP password by typing:

```
ndmpd password root
```

Either use this password, or create a new NDMP username and password when configuring NDMP in [Adding NetApp filer to the CTA configuration](#).

Configuring NetApp archiving on the CTA

To archive from the NetApp filer, configure the FPolicy callback service on the CTA, CTA-HA, or CTA/VE.

1. Type the following:

```
fpsetup init_rffm
```

2. At the prompt that appears, select the interface on which the FPolicy callback daemon should listen for callbacks from NetApp filers. If there is only one interface, it will be selected automatically:
 - a. If this is the primary callback agent in the environment, type **n**.
 - b. If this machine is being configured as the secondary callback agent, type **y**. When prompted, type the IP address and the root password of the primary agent.

Adding NetApp filer to the CTA configuration

1. In the navigation field of the web browser, type the IP address of the CTA. The CTA GUI appears.
2. Type the username and password for the default account which are:
 - a. Username: **admin**
 - b. Password: **rain**
3. Select **Configuration** and **Server**. The **Server List** appears. Click **Add**.
4. On the **Server** page that appears, select **NetApp**. Click **OK**.

5. For each step in the File Server wizard, provide the required information indicated with a red asterisk and click **Next**.

Step	Description
Basic File Server Information	Type the NetBIOS server name.
IP Addresses vFiler Host IP Address	<p>Type the IP address for NetApp filer:</p> <p>When editing an existing server, click Update to retrieve the IP address from the DNS that is based on the server name.</p> <p>To specify an additional IP address, click Add.</p> <p>To delete an existing IP address, select an IP and click Delete. If using a vFiler with</p> <p>OnTAP versions earlier than 7.3, type the IP address of the hosting NetApp filer.</p>
CIFS Specific Settings	<p>This is the Windows domain user to be used by the appliance. To avoid permission issues during archiving, recall, and migration, add this user who is a domain user to the backup operator group. Log in as root on the NetApp and type:</p> <p>useradmin domainuser add <user> -g "Backup Operators"</p> <p>where <i>user</i> is the Windows domain username for login.</p> <p>If this user is not in the domain administrator group, add it to the file server's local administrators group. Log in as root on the NetApp and type:</p> <p>useradmin domainuser add <user> -g administrators</p> <p>where <i>user</i> is the Windows domain username for login.</p> <p>Windows domain user provides more information on administering domain users.</p> <p>Specify either a NetBIOS domain or a Kerberos FQDN. The Kerberos FQDN is a CTA server configuration setting, and if selected, the same FQDN applies to all servers. Before selecting the Kerberos FQDN option, follow steps in Configuring Windows for Kerberos to configure your Fully Qualified Domain on the CTA.</p>

Step	Description
NDMP Specific Settings	<p>For file migration, type the username and password for an NDMP user on the NetApp source. By default, the port for NDMP traffic is 10000.</p> <p>The NDMP user must belong to the backup operators group. To create a user in the backup operators group, log in as root on the NetApp, and type:</p> <p>useradmin user add <user> -g "Backup Operators"</p> <p>where <i>user</i> is the username for login.</p> <p>The command will prompt for a password, but this is not the NDMP password and is not used for the NDMP specific settings.</p> <p>The NDMP password is automatically created when the user is created. Use the NetApp CLI to retrieve the NDMP password by typing:</p> <p>ndmpd password <user></p> <p>where <i>user</i> is a user in the backup operators group. Use this password for the NDMP specific settings.</p> <p>For a NetApp filer, the root username and password can also be used as the NDMP username and password. However for a vFiler, the NDMP password is different from the root password. vFiler configuration provides instructions on how to find the NDMP password for the root user on a vFiler.</p>
NetApp as Source	<p>This option configures the CTA to archive data from the NetApp filer. If multiple appliances are connected to the same NetApp filer, only one appliance should be configured with the NetApp as Source. These options are not required if this NetApp is used as a destination.</p>
NetApp Local Admin This step appears if NetApp as Source is yes .	<p>Type the username and password of a user on the NetApp filer. The user must be a member of the NetApp local administrator's group.</p>
Directory Exclusion List This step appears if NetApp as Source is yes .	<p>These are the directories to exclude for all tasks that use scanning. The CTA ignores all system directories such as, etc, lost+found, and ckpt by default.</p> <p>Verify that stub files are not in the excluded directories. CTA will not access the excluded directories and the stub files will become orphans.</p>

Step	Description
NetApp FPolicy callback agents This step appears if NetApp as Source is yes .	<p>The primary agent recalls all files when it is registered with the NetApp. A secondary agent recalls files when the primary is unavailable.</p> <p>If the FPolicy callback agent is not explicitly configured as a secondary agent, then it is a primary agent and the NetApp file server will load balance between the registered primary agents.</p> <p>If no primary agents respond, then the NetApp filer will contact any of the registered secondary agents. When one of the primary agents is responsive again, the NetApp filer will automatically fail back to the primary agent.</p> <p>For the primary agent, select the agent that is on the same subnet as the NetApp machine. For the secondary agent, select another agent on the same subnet. If no such agent exists, select an agent on the next physically closest subnet. Up to two secondaries are supported.</p> <p>Primary or secondary agents may include CTA-HAs. If the CTA-HA is a primary, the recall is load balanced between the CTA-HA and other primary agents. If the CTA-HA is selected as a secondary, it is only used for recall if the CTA goes down.</p>

6. Review the Summary and click **Finish** to define the NetApp filer.

Verify the FPolicy configuration for CTA

After configuring the NetApp, FPolicy is normally configured automatically.

1. To confirm the FPolicy settings, type:

```
fpolicy
```

To verify that:

- a. FPolicy is enabled — Screen shows "CIFS file policy is enabled".
 - b. rfpolicy is enabled — Screen shows "File policy rfpolicy (file screening) is enabled".
 - c. If an HA has been configured, the IP address is listed.
2. If the settings are not correct, manually configure the fpolicy settings with the commands:

```
fpolicy create rfpolicy screen
fpolicy enable rfpolicy
fpolicy options rfpolicy required on
```

3. If there is a CTA-HA, manually configure it as a secondary FPolicy server with the command:

```
fpolicy options rfpolicy secondary_servers <ip>
```

where *ip* is the IP address of the CTA-HA. If there are multiple CTA-HAs, separate the IP addresses with commas.

Deploying CTA with a VNXe

CTA supports the VNXe server as an archive destination. For file migration, CTA supports the VNXe server as a destination when the source does not contain any stub files.

Prerequisites for using VNXe as a file migration destination

To migrate files to a VNXe, ensure that the following conditions are satisfied:

- For NFS, CTA has root access with read/write permission on destination exports.
- CIFS user has local administrator access on all destination shares.
- VNXe does not support stubs so CTA stub aware migration to a VNXe is not supported. To use a VNXe as a destination, recall all stubs to the source before starting a migration.
- For every server, properties are configured on the CTA for the NDMP username and password

[Adding VNXe to the CTA configuration](#) provides details on configuring a VNX server.

Adding VNXe to the CTA configuration

To configure CTA with a VNXe server:

1. In the navigation field of the web browser, type the IP address of the CTA. The CTA GUI appears.
2. Type the username and password for the default account which are:
 - a. Username: **admin**
 - b. Password: **rain**
3. Select **Configuration** and **Server**. The **Server List** appears. Click **Add**.
4. On the **Create Server** page that appears, select **VNXe**. Click **OK**.
5. For each step in the File Server wizard, provide the required information indicated with a red asterisk and click **Next**.

Step	Description
Basic File Server Information	Type the NetBIOS server name.
IP Addresses	Type the IP address for the VNXe: When editing an existing server, click Update to retrieve the IP address from the DNS that is based on the server name. To specify an additional IP address, click Add . To delete an existing IP address, select an IP and click Delete .

Step	Description
CIFS Specific Settings	<p>This is the Windows domain user to be used by the appliance. The domain user must be a member of the local Administrator group and part of the Backup Operators group on the VNXe. Windows domain user provides more information on administering domain users.</p> <p>Specify either a NetBIOS domain or a Kerberos FQDN. The Kerberos FQDN is a CTA server configuration setting, and if selected, the same FQDN applies to all servers. Before selecting the Kerberos FQDN option, follow steps in Configuring Windows for Kerberos to configure your Fully Qualified Domain on the CTA.</p>
NDMP Specific Settings	<p>For file migration, type the username and password for an NDMP user on the destination servers. By default, the port for NDMP traffic is 10000.</p> <p>If no NDMP user exists on the VNXe, use the VNXe GUI to create a new user with username and password.</p>

Deploying CTA with Unity

To use CTA with Unity, perform the following configuration steps on the CTA.

CTA file archiving for Unity is supported by Unity version 4.1.x and higher. CTA block archiving for Unity is supported by Unity version 4.2.x and higher.

With Unity as a source, CTA currently supports only archiving tasks. File migration from or to Unity is not supported. When performing archive tasks from Unity, the destination can only be cloud (Virtustream, AmazonS3, ECS/S3, or Azure).

Prerequisites for Unity as an archive source

To archive data from Unity, the CTA requires access to the Common API (TCP port 5080).

To archive NFS data, the CTA needs the following:

- Mount v3 RPC service
- NFS v3 RPC service
- NLM v4 RPC service
- Root access with read/write export permissions for all NFS data that will be archived.

To archive CIFS data, the CTA needs SMB over NetBIOS (TCP port 139).

When configuring Unity properties on the CTA, plan to provide:

- Credentials for Management and DHSM. These credentials are used for both archive and recall.
- Management IP address.
- (For CIFS archiving only) The NetBIOS name of the Unity NAS server credentials for the [Windows domain user](#).

When archiving, virus scanning can be enabled on the source, but scanning on read can cause time-out conditions with no reply errors on the source if CTA is the first application to access a file.

Configure name resolution for archiving

When the Unity needs to establish a connection to the appliance to recall file data from a cloud service, it tries to resolve the FQDN from the HTTP DHSM connection in its local hosts file. If it cannot be resolved locally, the NAS server will use DNS.

Note [Configuring CTA](#) provides information on how to configure the CTA hosts file.

To use DNS:

1. Create a DNS entry for the callback daemon that points to the appliance.

As a best practice, the CTA hostname and ACD hostname should be different. Using the same hostname for both can cause issues when trying to log into the appliance.

2. Create multiple entries by the same name for each callback appliance.

For each entry that is created, select the checkbox for **Create associated pointer (PTR) record** to ensure that it will be included in the Reverse Lookup Zones list.

Note The Unity NAS server supports DNS HA failover. If the DNS server resolves the callback daemon hostname to multiple IP addresses, the Unity NAS server transparently switches to the server at the next available IP address.

DNS is the preferred method for name resolution. If using DNS is not feasible, use local hostname resolution.

Configuring Unity to cloud archiving on the CTA

To archive from Unity to cloud storage, configure the cloud callback service so that the CTA is in the recall path.

Configure the callback service to recall from the cloud

To configure recall from the cloud service (Virtustream, Amazon S3, or Azure):

1. From the console on the CTA, log in as root.
2. Type **!** to escape to the command line and type:

```
\opt\rainfinity\filemanagement\bin\acdsetup init_rffm
```

3. Type **n** when the following message appears:

```
By default the Cloud Callback Daemon will connect to the CTA service
on the local machine.
Do you wish to configure another CTA? (y/n)
```

4. If there is a secondary callback agent such as a CTA-HA, log in on that agent as root, and repeat steps 2 and 3. In step 3, type **y** to provide the IP address and root password of the primary callback agent.

If an invalid IP address is provided, the ACD.log file, located in `/var/log/rainfinity/filemanagement/recall`, will contain errors to indicate that

there was no response from the primary agent. To correct the problem, repeat steps 2 through 4 of this procedure.

Configuring automatically created DHSM connections for archiving

The CTA or CTAVE will automatically create DHSM connections for Unity systems running OE version 4.1.x and higher.

When archiving any type of data to cloud storage server, recall requests will flow from the Unity NAS server to CTA, CTA-HA, or CTAVE.

Adding a Unity file server to the CTA configuration

1. In the navigation field of the web browser, type the IP address of the CTA. The CTA GUI appears.
2. Type the username and password for the CTA account:
 - a. Username: **admin**
 - b. Password: **rain**
3. Select **Configuration** and **Server**. The **Server List** appears. Click **Add**.
4. On the **Create Server** page that appears, select **Unity** under **NAS File Servers**. Click **OK**.
5. For each step in the **Create Server wizard**, provide the required information.

Step	Description
Basic File Server Information	Type the NetBIOS server name.
IP Addresses	Type the Unity NAS server IP address: When editing an existing server, click Update to retrieve the IP address from the DNS that is based on the server name. To specify an additional IP address, click Add . To delete an existing IP address, select an IP and click Delete . The Management IP is required for archiving. This allows CTA to automatically perform some prerequisite tasks for archiving, as described in Configuring automatically created DHSM connections for file migration or archiving . If this field is empty, the CTA takes no action, and you must perform the prerequisite tasks manually.
CIFS Specific Settings	This is the Windows domain user to be used by the appliance. The domain user must be a member of the local Administrator group and part of the Backup Operators group on the Unity NAS Server. For more information, see Windows domain user . Specify either a NetBIOS domain or a Kerberos FQDN. The Kerberos FQDN is a CTA server configuration setting, and if selected, the same FQDN applies to all servers. Before selecting the Kerberos FQDN option, follow the steps in Configuring Windows for Kerberos to configure your Fully Qualified Domain on the CTA. The CIFS credential is not required if the Unity system performs only NFS archiving.

Step	Description
Unity as Source	<p>This option configures the CTA to archive data from the Unity NAS server. If multiple appliances are connected to the same Unity NAS server, only one appliance should be configured with the Unity as Source. This option is required only if the Unity is serving as a source for archiving.</p> <p>Multiple appliances can be configured to archive data from a single Unity NAS server, but only one CTA or CTA/VE can be used to archive data from a single file system.</p>
Unity Callback Agent settings This step appears if Unity as source is yes	<p>The ACD DNS Name is required for archiving to a cloud storage server. For the ACD DNS name, type the FQDN of the Cloud Callback DNS entry. If a callback name is not specified, archive tasks to the destination will not start. The FQDN is case-sensitive.</p>
Directory Exclusion List	<p>The directories to exclude for all tasks that use scanning. The CTA ignores all system directories such as, etc, lost+found, and ckpt by default.</p> <p>Verify that stub files are not in the excluded directories. If they are in the excluded directories, CTA will not access the excluded directories, and the stub files will become orphans.</p>

6. Review the Summary and click **Finish** to define the Unity file server.
7. Select Configuration and **Common API Settings**. The Common API Settings page appears.
8. Type the username and password for **Management Credentials** and **DHSM Credentials** under **Unity Settings**. CTA uses these credentials to create a secondary HTTP DHSM connection using the REST API.

The **Unity settings** apply to all Unity servers configured on the CTA. If all Unity servers are archive targets only, the **Unity Settings** are not required.

Adding a Unity block server to the CTA configuration

1. In the navigation field of the web browser, type the IP address of the CTA. The CTA GUI appears.
2. Type the username and password for the default account:
 - a. Username: **admin**
 - b. Password: **rain**
3. Select **Configuration** and **Server**. The **Server List** appears. Click **Add**.
4. On the **Create Server** page that appears, select **Unity** under **Block Server**. Click **OK**.
5. For each step in the **Create Server wizard**, provide the required information.

Step		Description
Basic Block Server Information	Name	Name of the block server (Unity system).
	Server FQDN	Fully Qualified Domain Name (FQDN) of the block server. This should resolve to the Unity system management IP address.
ISCSI Target	ISCSI Targets	Select the ISCSI targets.
Chap Credentials	Use Hex Format Username Password	If the selected ISCSI targets require CHAP credentials, select whether you want to use Hex format for the CHAP password. Then type the CHAP user name and CHAP secret
Mutual Chap Credentials	Username Password	If the ISCSI targets require mutual CHAP credentials, type the CHAP user name and CHAP secret.
Directory Exclusion List		The directories to exclude for all tasks that use scanning. The CTA ignores all system directories such as, etc, lost+found, and ckpt by default. Verify that stub files are not in the excluded directories. If they are in the excluded directories, CTA will not access the excluded directories, and the stub files will become orphans.

- Review the Summary and click **Finish** to define the Unity block server.
- Select Configuration and **Common API Settings**. The Common API Settings page appears.
- Type the username and password for **Management Credentials** and **DHSM Credentials** under **Unity Settings**. CTA uses these credentials to create a secondary HTTP DHSM connection using the REST API.

The **Unity settings** apply to all Unity servers configured on the CTA. If all Unity servers are archive targets only, the **Unity Settings** are not required.

Create a policy for archiving Unity files or snapshots

- Under **Policies**, select **Policy** and click **Add**. The **Policy Wizard** appears.
- Configure policy settings. For archiving Unity files, select **Archive** as the policy type. For archiving Unity snapshots, select **Block Archive** as the policy type.
- To start defining rules for this policy, click **Next**. The **Define Policy Rules Wizard** appears.
- Click **Add Rule**.
- For each expression you want to build, provide the settings and click **Add**.
- When you are done defining expressions, click **Next**.
- Fill in the settings for the archive destination and click **Save Rule**. Then Click **Next**.
- Click **Save Policy**.

Create a task to archive or restore Unity files

Use time-based scheduling to archive or restore a Unity snapshot to and from the cloud (Amazon S3, ECS/S3, Azure, or Virtustream).

1. Under **Schedule**, click **Add**.
2. On the **Schedule Wizard**, select **Archive**.
3. Select the server type, name, protocol, and path of the file server that contains the file to archive or restore.
4. Select the policy to use for archiving.
5. Select the archive destination and fill in the appropriate fields.
6. Fill in the fields for a **Time Based** schedule.
7. Click **Finish**.

Create a task to archive or restore Unity snapshots

Use time-based scheduling to archive or restore a Unity snapshot to and from the cloud (Amazon S3, ECS/S3, Azure, or Virtustream).

1. Under **Schedule**, click **Add**.
2. On the **Schedule Wizard**, select the **Block** tab.
3. Select **Block Archive** or **Block Restore** and click **Next**.
4. Select the block server and the LUN or consistency group that contains the snapshot to archive or restore.
5. Select the snapshot to restore and click **Next**.
6. Select the block server on which to restore the selected snapshot and click **Next**.
7. Fill in the fields for a **Time Based** schedule.
8. Click **Finish**.

Deploying CTA with Virtustream

CTA supports Virtustream cloud storage as an archiving destination.

Note The Virtustream manager should provide you with the following settings: Access Key ID, Secret Access Key, and Bucket Name.

Adding Virtustream to the CTA configuration

Configure the CTA with Virtustream.

1. In the navigation field of the web browser, type the IP address of the CTA. The CTA GUI appears.
2. Type the username and password for the default account:
 - a. Username: **admin**
 - b. Password: **rain**
3. Select **Configuration** and **Server**. The **Server List** appears. Click **Add**.

4. On the **Create Server** page that appears, select **Virtustream** under **Cloud Servers**. Click **OK**.
5. For each step in the File Server wizard, provide the required information and click **Next**.

Step	Description
Basic File Server Information	Type the logical server name. Each file server is configured separately with a unique name and is associated with a single Virtustream bucket.
Web Service Specific Settings	<p>URL — URL corresponding to the Virtustream bucket.</p> <p>Access Key ID — Type the key that CTA uses to gain REST access to the Virtustream bucket.</p> <p>Secret Access Key — Type the secret credential used with the Access Key ID.</p> <p>Bucket Name — Type the Virtustream bucket destination where CTA adds, deletes, or edits objects.</p>

6. Review the Summary and click **Finish** to define the Virtustream.

Note When archiving to the cloud (for example, to Atmos, Virtustream, Amazon S3, or Azure), the CTA clock must be accurate and in sync with the cloud service. Therefore, it is required to have an NTP server configured when archiving to these destinations.

Deploying CTA with a Windows server

CTA supports the Windows 2003 and 2008 servers as a CIFS archiving destination or as a file migration source. If migrating files from a Windows server, the EMCOPY and LGDUP utilities must be installed on the Windows server before running the file migration task.

Prerequisites for using Windows as a file migration source

To migrate files from Windows as a source, ensure that the following conditions are satisfied:

- The amount of free disk space on the Windows source is enough to hold a snapshot of the volume during migration. To estimate the amount of space required, estimate the amount of change to the data since the previous snapshot. For example if 25% of the data has changed, ensure that 25% of the disk space is free.
- The amount of free disk space on the destination is enough to hold the source data. If there is not enough disk space, the migration task will complete without error, but some data will not be migrated.
- CTA supports Windows cluster servers as a file migration source. However, because the Microsoft Volume Snapshot Service (VSS) is not cluster aware, if a failover occurs, the migration will fail.

Adding a Windows server to the CTA configuration

1. In the navigation field of the web browser, type the IP address of the CTA. The CTA GUI appears.
2. Type the username and password for the default account which are:
 - a. Username: **admin**
 - b. Password: **rain**
3. Select **Configuration** and **Server**. The **Server List** appears. Click **Add**.
4. On the **Create Server** page that appears, select **Windows**. Click **OK**.
5. For each step in the File Server wizard, provide the required information indicated with a red asterisk and click **Next**.

Step	Description
Basic File Server Information	Type the NetBIOS server name.
IP Addresses	Type the IP address for the Windows server. When editing an existing server, click Update to retrieve the IP address from the DNS that is based on the server name. To specify an additional IP address, click Add . To delete an existing IP address, select an IP and click Delete .
CIFS Specific Settings	This is the Windows domain user to be used by the appliance. The domain user must be a member of the local Administrator group and part of the Backup Operators group on the Windows server. Windows domain user provides more information on administering domain users. Specify either a NetBIOS domain or a Kerberos FQDN. The Kerberos FQDN is a CTA server configuration setting, and if selected, the same FQDN applies to all servers. Before selecting the Kerberos FQDN option, follow steps in Configuring Windows for Kerberos to configure your Fully Qualified Domain on the CTA.

6. Review the Summary and click **Finish** to define the Windows server.

Note NDMP is not used for file migration from a Windows server as a source and NDMP specific settings are not included as Windows properties.

Installing the EMWS copy agent for CTA

To perform CIFS file migration from a Windows source to a VNX or VNXe destination, install the EMWS copy agent on the Windows server.

1. Log in to the Windows server as administrator with full backup, restore, and security permissions.
2. Go to the EMC Online Support web site:
<http://Support.EMC.com>
3. To download the EMWS copy agent:

- a. For CTA, select **Downloads**. Search for **Cloud Tiering Appliance**. Download **EMC CTA Migration Windows Service (EMWS)**.
 - b. For CTA/VE, select **Downloads**. Search for **Cloud Tiering Appliance/VE**. Download **EMC CTA Migration Windows Service (EMWS)**.
4. To install EMWS type:
- emws_V1_install /user <domain>\<user> /passwd <password>** where *domain* is the local machine name and *user* is an administrator. Both are required.
- a. To update an installation, type: **emws_V1_install**
 - b. To uninstall, type: **emws_V1_install/uninstall**

Installing and running LGDUP

When running a Windows file migration task, CTA compares the local users and local groups on the source with those on the destination. If any users or groups are missing on the destination, the run will fail with a message written to the CTA file migration log.

The LGDUP utility merges the users and groups from the source with the users and groups on the destination. Install the LGDUP utility on the Windows machine and run it to configure the users and groups on the source and destination before running a file migration task.

Note For VNXe, creation of local users is not supported, so CTA only migrates local groups to VNXe. SID translation of local users to VNXe is not supported.

1. Log in to the Windows machine as administrator with full backup, restore, and security permissions.
2. Go to the EMC Online Support web site:
<http://Support.EMC.com>
3. To download the LGDUP utility, select **Downloads**. Search for **VNX**. Download **CIFSTools.zip**.
4. Unzip the file. In the resulting directory, click down to locate the `lgdup.exe` file.

Note The LGDUP utility is a 32-bit command line application. If the Windows server is running a 64-bit operating system, run LGDUP in 32-bit compatibility mode.

5. To merge the local users and local groups database from the Windows source with the destination and replicate privileges to the target:
 - a. For migration to VNX, type:


```
lgdup.exe -v -u -l lgdup_log.txt \\<src_server_name>
\\<dst_server_name>
```
 - b. For migration to VNXe, type:


```
lgdup.exe -v -l lgdup_log.txt \\<src_server_name>
\\<dst_server_name>
```

where *src_server_name* is the name of the Windows NetBIOS name and *dst_server_name* is the name of the VNX or VNXe NetBIOS name.

Deploying CTA with Isilon

CTA supports Isilon as a NAS destination for migration or archiving. To use CTA with an Isilon, first configure the Isilon, and then configure the appliance.

Prerequisite tasks when using Isilon as a CIFS share destination

1. The Isilon must be part of a Windows domain or Active Directory. Verify that:
 - a. Isilon and the CTA CIFS user are in the same domain, such as rain.prv.
 - b. You are not in WorkGroup mode.
2. On a Windows machine that is part of the Isilon domain, create the destination directory that CTA will use as a share. Security on the directory must allow full control to the group Domain Admins and user Administrator. Administrator is the same Windows domain user that will be used by CTA to archive to the Isilon.
 - a. Log in to the Windows machine as administrator.
 - b. Using Windows Explorer, create a directory in **ifs/data**. For example, create **CTADir**.
 - c. For folder CTADir, right-click **Properties**. The properties screen for that directory appears.
 - d. On the **Security** tab, click **Add**. For Windows 7 or 8, click **Edit > Add**. The **Select Users, Computers, or Groups** screen appears.
 - e. For **Enter the object names to select**, type the group **Domain Admins**. Click **Check Names**. If the group name is found, it is underlined.
 - f. Click **OK**. The properties screen reappears.
 - g. For Allow, select **Full Control**. Click **Apply**.
 - h. On the **Security** tab, click **Add**. The **Select Users, Computers, or Groups** screen appears.
 - i. For **Enter the object names to select**, type the username for the Windows domain user to be used by the appliance as described in [Windows domain user](#). Click **Check Names**. If the user name is found, it will be underlined.
 - j. Click **OK**. The properties screen reappears.
 - k. For Allow, select **Full Control**.
 - l. Click **OK**.
3. For CTA to communicate with the Isilon using CIFS protocol, SMB service on the Isilon must be enabled. Check the SMB service settings on the Isilon.
 - a. In the navigation field of the web browser, type the IP address of the Isilon. The Isilon Administration GUI appears.
 - b. Log in to the Isilon with administrator privileges.
 - c. Select **File Sharing > SMB > Settings**.
 - For **Service status**, verify that **Enable** is selected. If not, select it.
 - For **Security mode**, verify that **User** is selected. If not, select it.

If any changes are made, click **Submit**.

4. From the Isilon console, verify that the NetBIOS Name Service (NBNS) is running.

a. Check that the NBNS is enabled with the command:

```
isi services | grep nbns
```

b. If the NBNS is not enabled, type:

```
isi services -a nbns enable
```

c. Even if the NBNS is enabled, it may not be running. To verify that the NBNS is running, check whether the Isilon is listening on port 139 with the command:

```
netstat -an | grep LISTEN | grep 139
```

d. If the Isilon is not listening on port 139, disable and re-enable the NBNS with the commands:

```
isi services -a nbns disable  
isi services -a nbns enable
```

e. Recheck whether the Isilon is listening on port 139:

```
netstat -an | grep LISTEN | grep 139
```

5. From the Isilon Administration GUI, create a CIFS share for the CTA destination directory.

a. Select **File Sharing > SMB > Add Share**.

– For **Share name**, type the name of the share or export created, for example, **CTADir**.

– For **Directory to share**, click **Browse**. Select the path to the directory created in step 2, for example, **/ifs/data/CTADir**. Click **OK**.

– For Directory ACLs, select **Do not change existing permissions**.

b. Under Users and Groups, select **Add**. The Choose Users and Groups screen appears.

– For **From the location**, select the domain, such as **rain.prv** that was referenced in step 1.

– For **Name**, type **Administrator**. This is the same Windows domain user that CTA will use.

– Click **Search**. Search results list the Administrator user of the domain.

– Select the Administrator user and click **Choose**. Under Users and Groups, the Administrator is listed with the checkbox selected.

c. Click **Edit permissions**. The permissions screen appears.

– Select **Run as root**. All permissions are set to Allow.

– Click **OK**.

d. Click **Submit**. The SMB Summary appears with the share added.

Prerequisite tasks when using Isilon as an NFS export destination

For CTA to communicate with the Isilon using NFS protocol, NFS service on the Isilon must be enabled. First check the NFS service settings on the Isilon and then create the NFS export for the CTA destination directory.

Note Before creating the NFS export for the CTA destination directory, you must configure the Isilon server so that it provides CTA with root CIFS access permissions. You can do this running specific Isilon-specific CLI commands that add "run-as-root" permissions for the CTA domain user on every Isilon CIFS share used by CTA system. The following example shows Isilon CLI commands used to do this for a single CIFS share:

```
Isilon v6.5:    > isi smb permission create --
sharename='myShareName' --account-
name='myDomain'\\'myUserName' --account-type=user --
permission-type=allow --permission=run-as-root
```

```
Isilon v7.0:    > isi smb share permission create
'myShareName' 'myDomain'\\'myUserName' --run-as-root
```

1. In the navigation field of the web browser, type the IP address of the Isilon. The Isilon Administration GUI appears.
2. Log in to the Isilon with administrator privileges.
3. Select **File Sharing > NFS > Settings**.
 - a. For **Service status**, verify that **Enable** is selected. If not, select it.
 - b. For **NFSv4 Support**, verify that **Disable** is selected. If not, select it.

If any changes are made, click **Submit**.

4. Select **File Sharing > NFS > Add Export**.
 - a. Under Settings:
 - b. For **Directories**, click **Browse**. Select the path to the directory. Click **OK**.
 - c. For **Permissions**, select **Enable write access** and select **Enable mount access to subdirectories**.
 - d. Under Access Control:
 - e. For **User name**, type **root**.
 - f. For **Group names**, type **everyone**.
5. Click **Submit**. The NFS Summary appears with the export added.

Prerequisite task when using Isilon as a file migration destination

To configure CTA with Isilon as destination for file migration, an NDMP administrator user on the Isilon is required.

1. In the navigation field of the web browser, type the IP address of the Isilon. The Isilon Administration GUI appears.
2. Log in to the Isilon with administrator privileges.

3. Select **File System > Backup > NDMP Settings**.
 - a. Under Service, if the NDMP service is disabled, click **Enable**.
 - b. Under Settings, port number 10000 is the default. This is the same port number for the NDMP Specific Settings configured when [Adding Isilon to the CTA configuration](#).
 - c. In the NDMP Administrators section, click **Add administrator**.
 - d. For username, type **ndmp_user**. This is the same username that CTA will use.
 - e. For password, type a password. This is the same password that CTA will use.
 - f. Retype the password.

Adding Isilon to the CTA configuration

Configure Isilon properties on the CTA.

1. In the navigation field of the web browser, type the IP address of the CTA. The CTA GUI appears.
2. Type the username and password for the default account which are:
 - a. Username: **admin**
 - b. Password: **rain**
3. Select **Configuration** and **Server**. The **Server List** appears. Click **Add**.
4. On the **Create Server** page that appears, select **Isilon**. Click **OK**.
5. For each step in the File Server wizard, provide the required information indicated with a red asterisk and click **Next**.

Step	Description
Basic File Server Information	Type the NetBIOS server name.
IP Addresses	Type the IP address for the Isilon: When editing an existing server, click Update to retrieve the IP address from the DNS that is based on the server name. To specify an additional IP address, click Add . To delete an existing IP address, select an IP and click Delete . If the Isilon cluster is configured using SmartConnect, add all IP addresses configured for the cluster. CTA and Isilon SmartConnect provides additional information about Isilon properties.

When configuring this Isilon cluster as in [Adding Isilon to the CTA configuration](#), the settings are:

NetBIOS server name:	isilon1
IP Addresses:	10.4.0.80, 10.4.0.81, 10.4.0.82
CIFS NetBIOS Domain or CIFS Kerberos FQDN:	mydomain.prv

Configure CTA with all the SmartConnect Pool IPs, but not the SmartConnect Service IP (such as 10.4.0.89). The SmartConnect Service IP address is only used for DNS delegation.

CTA will not be able to correctly load balance its access to the Isilon cluster if the cluster's IP Pool does not match the configuration in CTA. If Isilon nodes or SmartConnect Pool IP addresses are added to or removed from an Isilon cluster, edit the corresponding CTA configuration to add or remove the IPs.

DNS forward and reverse lookup

CTA requires forward DNS (name to IP) and reverse DNS (IP to name) lookups to be configured correctly in your DNS server. Once the Isilon cluster, the CTA, and the DNS server are configured for SmartConnect and DNS delegation, use the CTA CLI to File configuration.

To verify that forward DNS lookup resolves the Isilon cluster name to one of the IPs configured in the SmartConnect IP Pool, at the CTA CLI prompt, type:

```
nslookup <name>
```

If the Isilon cluster is configured using the SmartConnect round-robin option, repeated nslookup name to IP requests should cycle through all the addresses configured in the SmartConnect IP Pool.

For example, if a three node Isilon cluster is configured to use SmartConnect with:

Isilon Cluster Name: isilon1.mydomain.prv

Isilon SmartConnect Service IP: 10.4.0.89

Isilon SmartConnect Pool IPs: 10.4.0.80, 10.4.0.81,
10.4.0.82

nslookup isilon1.mydomain.prv must resolve to one of the IPs: 10.4.0.80, 10.4.0.81, or 10.4.0.82.

If nslookup fails or returns an IP that is not in the SmartConnect IP Pool, the DNS server is not correctly configured to delegate to the Isilon SmartConnect Server IP. See the Isilon documentation to configure forward DNS lookup and use nslookup to test the configuration again.

To verify that reverse DNS lookup resolves all of the IPs configured in the SmartConnect IP Pool to the DNS name of the Isilon cluster, at the CTA CLI prompt, type:

```
nslookup <IP>
```

For example:

- **nslookup 10.4.0.80** must resolve to isilon1.mydomain.prv
- **nslookup 10.4.0.81** must resolve to isilon1.mydomain.prv
- **nslookup 10.4.0.82** must resolve to isilon1.mydomain.prv

If any of these lookups fails to resolve correctly the DNS server's reverse lookup zone is not correctly configured. See the Isilon documentation to configure reverse DNS lookup and use nslookup to test the configuration again.

Note If either the forward DNS or reverse DNS lookup tests fail, do not use an Isilon server configured with SmartConnect in CTA for archiving, repository migration, or file migration.

Deploying the CTA with Data Domain

CTA supports the EMC Data Domain storage product as an NFS destination for archiving and repository migration. CIFS source data can be exported and archived to Data Domain through NFS while clients access the archived files through CIFS. The stub file retains the CIFS attributes such as security settings and ADS.

Note CTA supports Data Domain CIFS protocol starting from Data Domain 5.3. However, CTA requires admin privileges to access CIFS shares and successfully create all files on Data Domain.

Configure Data Domain server properties on the CTA.

1. In the navigation field of the web browser, type the IP address of the CTA. The CTA GUI appears.
2. Type the username and password for the default account which are:
 - a. Username: admin
 - b. Password: **rain**
3. Select **Configuration** and **Server**. The **Server List** appears. Click **Add**.
4. On the **Create Server** page that appears, select **Data Domain**. Click **OK**.

- For each step in the File Server wizard, provide the required information indicated with a red asterisk and click **Next**.

Step	Description
Basic File Server Information	Type the NetBIOS server name.
IP Addresses	Type the IP address for the Data Domain server: When editing an existing server, click Update to retrieve the IP address from the DNS that is based on the server name. To specify an additional IP address, click Add . To delete an existing IP address, select an IP and click Delete .
CIFS Specific Settings	This is the Windows domain user to be used by the appliance. The domain user must be a member of the local Administrator group and part of the Backup Operators group on the VNXe. Windows domain user provides more information on administering domain users. Specify either a NetBIOS domain or a Kerberos FQDN. The Kerberos FQDN is a CTA server configuration setting, and if selected, the same FQDN applies to all servers. Before selecting the Kerberos FQDN option, follow steps in Configuring Windows for Kerberos to configure your Fully Qualified Domain on the CTA. If the Data Domain is an NFS export destination only, this setting is not used.

- Review the Summary and click **Finish** to define the Data Domain.

Configuring a NAS-based repository

Any VNX, VNXe, Unity, Windows, Isilon, NetApp, or Data Domain system can be configured as a NAS-based repository.

Note The appliance must have read/write access to any share or export that may be used as an archive source or destination. In addition, the appliance must have read/write permission for any file that it may archive.

- In the navigation field of the web browser, type the IP address of the CTA. The CTA GUI appears.
- Type the username and password for the default account which are:
 - Username: **admin**
 - Password: **rain**
- Select **Configuration** and **NAS Repository**. The **NAS Repository List** appears. Click **Add**. The **NAS Repository Properties** dialog box appears.

You can either:

- Add an existing server and repository.
- Add a new server and repository.

4. To add an existing server and repository, specify the following:
 - a. Type — Select a type of NAS server from the list.
 - b. Name — Select a file server of that type from the list.

Note The file server must have a proper DNS entry defined that links the file server name with the IP address.
 - c. Protocol — Select NFS or CIFS. The source and repository protocol types must match.
 - d. If the source protocol is CIFS, the NAS repository protocol must be CIFS.
 - e. If the source protocol is NFS, the NAS repository protocol must be NFS.

If the CIFS protocol is selected, use the CIFS user in the filesystem CIFS DHSM connection string for CIFS specific settings when configuring the primary storage on the appliance:
 - f. Path — Click **Browse** to select an existing path.

Once the path is specified, a name in the form of **Repository at <path>** appears in the **Name** field.
 - g. Maximum limit of disk usage — Type a percentage value for disk usage. Default value is 90%.
5. To add a new server, select the Type and click **Add Server**. Follow instructions to add the server as outlined in:
 - a. [Adding VNX to the CTA configuration](#)
 - b. [Adding NetApp filer to the CTA configuration](#)
 - c. [Deploying CTA with a VNXe](#)
 - d. [Deploying CTA with Unity](#)
 - e. [Deploying CTA with a Windows server](#)
 - f. [Deploying CTA with Isilon](#)
 - g. [Deploying the CTA with Data Domain](#)
6. Click **OK** to finish defining the NAS repository.

Deploying CTA with Amazon S3

Amazon Simple Storage Service (S3) is cloud storage. CTA supports Amazon S3 as an archiving destination.

Note The Amazon S3 account manager uses the Amazon Web Services (AWS) Management Console to generate the Web Service Specific Settings which are Access Key ID, Secret Access Key, and Bucket Name.

Adding Amazon S3 to the CTA configuration

Configure the CTA for an Amazon S3.

1. In the navigation field of the web browser, type the IP address of the CTA. The CTA GUI appears.

2. Type the username and password for the default account which are:
 - a. Username: **admin**
 - b. Password: **rain**
3. Select **Configuration** and **Server**. The **Server List** appears. Click **Add**.
4. On the **Create Server** page that appears, select **Amazon S3**. Click **OK**.
5. For each step in the File Server wizard, provide the required information indicated with a red asterisk and click **Next**.

Step	Description
Basic File Server Information	Type the logical server name. Each file server is configured separately with a unique name and is associated with a single Amazon S3 bucket.
Web Service Specific Settings	<p>URL — URL corresponding to the Amazon S3 bucket as displayed in the AWS Management Console.</p> <p>Access Key ID — Type the key that CTA uses to gain REST access to the Amazon S3 bucket. The Amazon S3 account manager generates the key.</p> <p>Secret Access Key — Type the secret credential used with the Access Key ID. The Secret Access Key is generated with the Access Key ID.</p> <p>Bucket Name — Type the Amazon S3 bucket destination where CTA adds, deletes, or edits objects. The Amazon S3 account manager configures the Amazon S3 bucket.</p>

6. Review the Summary and click **Finish** to define the Amazon S3.

Note When archiving to the cloud (for example, Atmos, Amazon S3, or Azure), the CTA clock must be accurate and in sync with the cloud service. Therefore it is required to have an NTP server configured when archiving to these destinations.

Deploying CTA with Atmos

The Atmos cloud storage platform is an on-premise or service provider cloud. CTA supports Atmos as an archiving destination, and as a source or destination for repository migration.

Adding Atmos to the CTA configuration

Configure the CTA for an Atmos (or ViPR as Atmos).

1. In the navigation field of the web browser, type the IP address of the CTA. The CTA GUI appears.
2. Type the username and password for the default account which are:
 - a. Username: **admin**
 - b. Password: **rain**
3. Select **Configuration** and **Server**. The **Server List** appears. Click **Add**.

4. On the **Create Server** page that appears, select **Atmos**. Click **OK**.
5. For each step in the File Server wizard, provide the required information indicated with a red asterisk and click **Next**.

Step	Description
Basic File Server Information	Type the name to identify the Atmos.
Web Service Specific Settings	<p>DNS Name — Specify the name used to resolve the IP addresses in the Atmos cluster.</p> <p>Port — The GUI access method. HTTPS is the default and is typically used when the Atmos is deployed remotely. Select HTTPS or HTTP to specify the communication protocol. The default port number for HTTPS (10080) or HTTP (80) automatically appears. If your Atmos connects to HTTPS or HTTP through a different port, type the number.</p> <p>Username — Type the UID within a given Subtenant on the Atmos. CTA uses this UID to access storage on the cluster. If there is a subtenant, specify the username as: <i><Subtenant_ID> <UID></i>, where <i>Subtenant_ID</i> is an alphanumeric string generated by the Atmos.</p> <p>The UID is not the Tenant Name, the Subtenant Name, or the Subtenant ID.</p> <p>Password — Type the Shared Secret associated with the UID on the Subtenant Information page of the Atmos.</p>

6. Review the Summary and click **Finish** to define the Atmos.

Note When archiving to the cloud (for example, Atmos, Amazon S3, or Azure), the CTA clock must be accurate and in sync with the cloud service. Therefore it is required to have an NTP server configured when archiving to these destinations.

Installing the SSL certificate on the CTA

By default, the Atmos is configured without server-side authentication. However, the Atmos administrator can install an SSL X.509 certificate on the Atmos to provide more security to Atmos clients. If the certificate that is installed on the remote Atmos is available to the CTA, the CTA can use the certificate to authenticate the Atmos before archiving data to it. The CTA checks the DNS name in the certificate with the DNS name of the Atmos configured on the CTA. If the DNS names match, authentication succeeds.

To use the SSL X.509 certificate with the CTA:

- The certificate must be in PEM file format. The Atmos administrator can provide the file.
- On the CTA, the name of the certificate must be `atmos_trusted_CAs.pem`
- Install the certificate under:

```
/opt/rainfinity/filemanagement/conf/ATMOS_certs/<atmos_DNS>
```

where *atmos_DNS* is the DNS name of the Atmos configured on the CTA. The names are case-sensitive and must match exactly. Adding Atmos to the CTA configuration provides details on configuring the DNS name on the CTA.

To run a CTA script that installs the PEM certificate file in the correct directory, type:

```
/opt/rainfinity/filemanagement/bin/atmoscert.pl
```

Deploying CTA with Azure

Microsoft Azure is cloud storage. CTA supports Azure as an archiving destination.

Adding Azure to the CTA configuration

Configure the CTA for an Azure.

1. In the navigation field of the web browser, type the IP address of the CTA. The CTA GUI appears.
2. Type the username and password for the default account which are:
 - a. Username: **admin**
 - b. Password: **rain**
3. Select **Configuration** and **Server**. The **Server List** appears. Click **Add**.
4. On the **Create Server** page that appears, select **Azure**. Click **OK**.
5. For each step in the File Server wizard, provide the required information indicated with a red asterisk and click **Next**.

Step	Description
Basic File Server Information	Type the name to identify the Azure.
Web Service Specific Settings	Container Name — Type the destination where CTA adds, deletes, or edits objects. The Windows Azure Management Portal administrator configures the containers associated with the storage account. Account Name— Type the name of the Azure storage account that CTA uses to gain REST access to the Azure container. The Windows Azure Management Portal administrator creates this account. Access Key — Type the key that CTA uses to gain REST access to the Azure container. When a storage account is created, Windows Azure generates primary and secondary access keys for that account. Either key is valid.

6. Review the Summary and click **Finish** to define the Azure.

Note When archiving to the cloud (for example, Atmos, Amazon S3, or Azure), the CTA clock must be accurate and in sync with the cloud service. Therefore it is required to have an NTP server configured when archiving to these destinations.

Deploying CTA with Centera

CTA supports the Centera as an archiving destination, and as a source or destination for repository migration. Configure the CTA for an Centera.

1. In the navigation field of the web browser, type the IP address of the CTA. The CTA GUI appears.
2. Type the username and password for the default account which are:
 - a. Username: **admin**
 - b. Password: **rain**
3. Select **Configuration** and **Server**. The **Server List** appears. Click **Add**.
4. On the **Create Server** page that appears, select **Centera**. Click **OK**.
5. For each step in the File Server wizard, provide the required information indicated with a red asterisk and click **Next**.

Step	Description
Basic File Server Information	Type the name to identify the Centera.
Access Node	Type the Centera access node name or IP address: To specify an additional Access Node, click Add . To delete an existing Access Node, select a node and click Delete . The Access Node String is automatically generated when the Access Node is added or deleted. You cannot type data directly into the field.
Centera Authentication Type	Select from one of the three choices: User profile — If selected, type the username and password of the Centera profile that is to be used for archiving. PEA file — This option requires that a profile and pool entry authorization (PEA) file was created to access Centera, and that a copy of the PEA file resides on the CTA. If selected, the PEA file is used to authenticate the CTA connection with Centera. Type the path to the file on the local machine or browse for the file. A copy of the file will be stored with the CTA configuration. Anonymous — If selected, no security is used to authenticate with Centera.

6. Review the Summary and click **Finish** to define the Centera.

Deploying CTA with ECS

ECS supports APIs, including ECS-CAS, ECS-S3, EMC Atmos, and NFS. CTA supports ECS as an archiving destination.

Adding ECS to the CTA configuration

Configure the CTA for an ECS.

1. In the navigation field of the web browser, type the IP address of the CTA. The CTA GUI appears.
2. Type the username and password for the default account which are:
 - a. Username: **admin**
 - b. Password: **rain**
3. Select **Configuration** and **Server**. The **Server List** appears. Click **Add**.
4. ECS-CAS can be added to CTA as Centera. Refer to [Deploying CTA with Centera](#) for deployment instructions.
5. ECS/S3 can be added to CTA as Amazon S3 or Virtustream. Refer to [Deploying CTA with Amazon S3](#) for deployment instructions.

Note Manual port configuration changes might be needed on CTA, depending on the ECS port configuration. Please contact support if needed.
6. ECS/ATMOS-REST can be added to CTA as an Atmos server. Refer to [Deploying CTA with Atmos](#) for deployment instructions.

Chapter 5 Maintaining the Cloud Tiering Appliance

This chapter contains the following sections:

Importing a file list archive	90
Backing up the configuration	94
Encryption Key restore in case of catastrophic failures	97
Maintaining the database	98
Migrating from CTA to CTA/VE.....	99
Shutting down and restarting the appliance	100

Importing a file list archive

The CTA can perform archival tasks on lists of files imported from a third-party software provider. To generate and import the file lists properly, the CTA administrator and the third-party software administrator must exchange information as follows:

1. The third-party software administrator gives the CTA administrator the names of the primary servers.
2. The CTA administrator:
 - a. Adds the primary servers as source servers on the CTA as described in [Adding the primary servers](#).
 - b. Configures an import provider with username and password as described in [Configuring the import provider](#).
 - c. Configures an import file task as described in
 - d. [Configuring the import](#) task.

The CTA administrator provides the server, provider, and task information to the third-party software administrator.

3. With settings defined on the CTA, the third-party software generates an XML file containing lists of files on a primary server that are to be archived.

Note For optimum performance during file import, limit the XML file list to 1 million entries.

4. The XML file is transferred to the CTA as described in [Importing the file list](#). Once CTA validates the imported file, the import files archive task is queued to run.

Import files archive tasks are run in the same way as other archive tasks. The archive destination is defined with the CTA policy. The source is defined in the imported XML file list. Online help describes how to schedule and review results of an archive task.

Adding the primary servers

To archive files from the primary servers of the third-party software, configure the file servers as a source on the CTA. A VNX Data Mover, Unity NAS server, or NetApp filer can be a source. Details on how to add the file servers to the CTA configuration are provided in:

- [Adding VNX to the CTA configuration](#)
- [Adding a Unity file server to the CTA configuration](#)
- [Adding NetApp filer to the CTA configuration](#)

The CTA administrator gives the sever names to the third-party software administrator.

Configure name resolution for archiving

When the Unity needs to establish a connection to the appliance to recall file data from a cloud service, it tries to resolve the FQDN from the HTTP DHSM connection in its local hosts file. If it cannot be resolved locally, the NAS server will use DNS.

Note [Configuring CTA](#) provides information on how to configure the CTA hosts file.

To use DNS:

1. Create a DNS entry for the callback daemon that points to the appliance.

As a best practice, the CTA hostname and ACD hostname should be different. Using the same hostname for both can cause issues when trying to log into the appliance.

2. Create multiple entries by the same name for each callback appliance.

For each entry that is created, select the checkbox for **Create associated pointer (PTR) record** to ensure that it will be included in the Reverse Lookup Zones list.

Note The Unity NAS server supports DNS HA failover. If the DNS server resolves the callback daemon hostname to multiple IP addresses, the Unity NAS server transparently switches to the server at the next available IP address.

DNS is the preferred method for name resolution. If using DNS is not feasible, use local hostname resolution.

Configuring Unity to cloud archiving on the CTA

To archive from Unity to cloud storage, configure the cloud callback service so that the CTA is in the recall path.

Configure the callback service to recall from the cloud

To configure recall from the cloud service (Virtustream, Amazon S3, or Azure):

1. From the console on the CTA, log in as root.
2. Type **!** to escape to the command line and type:

```
\opt\rainfinity\filemanagement\bin\acdsetup init_rffm
```

3. Type **n** when the following message appears:

```
By default the Cloud Callback Daemon will connect to the CTA
service on the local machine.
Do you wish to configure another CTA? (y/n)
```

4. If there is a secondary callback agent such as a CTA-HA, log in on that agent as root, and repeat steps 2 and 3. In step 3, type **y** to provide the IP address and root password of the primary callback agent.

If an invalid IP address is provided, the ACD.log file, located in `/var/log/rainfinity/filemanagement/recall`, will contain errors to indicate that there was no response from the primary agent. To correct the problem, repeat steps 2 through 4 of this procedure.

Configuring automatically created DHSM connections for archiving

The CTA or CTAVE will automatically create DHSM connections for Unity systems running firmware version 4.1.x and higher.

When archiving any type of data to cloud storage server, recall requests will flow from the Unity NAS server to CTA, CTA-HA, or CTAVE.

Note File systems with Unity File Level Retention (FLR) enabled cannot be used as an archiving source.

- [Adding a Unity block server to the CTA configuration](#)
- [Adding NetApp filer to the CTA configuration](#)

The CTA administrator gives the server names to the third-party software administrator.

Configuring the import provider

To copy an archive file list to the CTA, the third-party software provider must have permission to write to the CTA. The import provider is a unique Linux user on the CTA that has limited security access to a staging directory. When the provider is configured, CTA creates the staging directory for the imported XML file.

1. Select **Configuration** and **Import Provider**. The **Import Provider List** appears.
2. Click **Add**. The **Import Provider Properties** page appears.
3. Specify the following:
 - a. **Username** — Type a name for the user that is allowed to log into the CTA from the provider. This must be a valid Linux user name. The third-party software uses the name for the provider in the XML file that it creates.
 - b. **Password** — Type a password used to log into the CTA from the provider.
4. Click **Commit** to define the Provider.

The CTA administrator gives the name and the password to the third-party software administrator.

Configuring the import task

For the CTA to operate on the imported file list, the header of the XML file must contain a task name configured on the CTA. The CTA matches the name in the header of the XML file to an import files archive task.

Configure an import files archive task based on an archive policy. Files are archived to the destination defined in the policy.

1. Select **Schedule** and click **Add**.
2. For each step in the Schedule wizard, provide the required information indicated with a red asterisk and click **Next**.

Step	Description
Task Type	On the Archive tab, select one of the archive task types. At the bottom of the screen, check the box to Import file list provided by an external provider.
Import File Task	Type a task name. Select the import provider defined in Configuring the import provider .
Policy	Click Add Policy .

Step	Description
Define Policy Parameters	Select an archive policy type: archive , multi_tier , or multi_tier_stub . Type a policy name.
Define Policy Rules	Click Add Rule .
Define a Policy Rule	To archive all files in the imported list, build an expression with: File Attribute: Size Operator: >= Value: 1 Unit: Bytes Click Add . Leave the default action of Archive .
Select the Archive Destination	Select a destination. Click Save Rule .
Define Policy Rule	The Define Policy Rules page reappears with the new rule added.
Verify Policy Rule	After verifying the policy definition, click Save Policy .
Policy	The Policy page reappears with the new policy added.
Schedule	Select the time-based Schedule Type for the import task: Import , Daily , Weekly , Monthly or Future . The task will not run until the XML file list is imported and validated.

3. After verifying the task definition, click **Finish**.

The CTA administrator gives the task name to the third-party software administrator.

Importing the file list

To transfer the file list to the CTA, the CTA imports the XML file or the import provider copies the XML file to the CTA.

CTA imports the file

If the XML file list is in a location that is accessible by the CTA, you can use the CTA GUI to import the file:

1. Select **Schedule**, and show schedules of type **Import Files**.
2. Highlight the import task defined in
3. [Configuring the import](#) task. Click **Import and select the Import File option**.
4. In the Explorer window that appears, select the XML file and click Open.

Import provider copies the file

If the import provider is copying the XML file to the CTA, it uses the name and password assigned by the CTA administrator in [Configuring the import provider](#) to log in to the CTA. The XML file list is deposited into a staging directory as:

```
/opt/rainfinity/filemanagement/import/providers/<PROVIDER>/import_files_<ID>.xml
```

where *PROVIDER* is the logical name of the provider defined in CTA and *ID* is the unique ID in the XML file.

Validating the import file

Every 30 seconds, the CTA scans for imported file lists. Before launching import file archive task, the CTA checks the imported XML file to confirm that it has matching values for:

- provider name
- task name
- file server names

The CTA also verifies that it can access all shares and exports listed in the XML file.

If any validation checks fail, an error is reported in the import log. To review import logs before running the import file archive task:

1. Select **Schedule**, and show schedules of type **Import Files**.
2. For any import file archive task defined, click **Import and select the Import Logs option**.
3. Each row corresponds to an XML file list import. On any row where the status shows a validation error, click **View Log** to display the log. The text file that appears will list problems with the XML file.

The logs for the XML file list import are also accessible outside of the GUI:

```
/opt/rainfinity/filemanagement/import/providers/<PROVIDER>/log
```

where *PROVIDER* is the logical name of the provider defined in CTA.

Backing up the configuration

The CTA and CTA/VE contain configuration information and critical database tables. The CTA-HA contains no persistent data. If data on a CTA-HA is lost, the CTA-HA software must be reinstalled as described in [Performing a CD clean install on the appliance](#).

If data on a CTA or a CTA/VE is lost, the software must be reinstalled and the last backup copy of the configuration and database tables must be restored. For this reason, backup the CTA or CTA/VE configuration and the critical database tables nightly.

Note Task and simulation log files are not included in a backup. To preserve these files, copy the `/opt/rainfinity/filemanagement/log/fw` directory to secure storage either periodically or before performing a CD clean install.

The backup feature uses the following process:

- Critical CTA system configuration data is backed up to a gzipped tar file (.tgz).

- The tar file is written to a destination on an Centera or NAS repository.

To perform disaster recovery, the CTA uses a catalog file (DBBackup.out) to locate tar files on the destination and to reconstruct the CTA system configuration after a disaster. The catalog file is stored on both the CTA and in a secondary disaster recovery location.

Creating a backup dump

Regular backups may be scheduled to run automatically.

1. Select Configuration and **Backup and Recovery Settings**.

Under File Management Backup Destination, specify:

- a. The number of backups — The default value is 5.
- b. Select Destination — The Centera or NAS repository where the backup files will be stored. The backup files are gzipped tar files (.tgz) that contain CTA system configuration data.

Note Only NFS NAS repositories are listed as backup destinations. CIFS NAS repositories are not supported as backup destinations and are not listed.

With NAS repositories, the data is stored with a directory structure that is similar to that of an archive destination.

- c. Select Disaster Recovery Location — The NFS export where the backup catalog file (DBBackup.out) will be stored. The DBBackup.out file is also stored on the CTA. The file stored on the NFS export is only used if the file stored on the CTA is lost or damaged.

Note Ensure that the selected location is only assigned to one CTA. Multiple CTAs must not use the same disaster recovery location.

2. Select **Schedule** and click Add. The Schedule Wizard appears.

3. On the **Auxiliary** tab, select **Backup**. Click **Next**.

4. For backup tasks, time-based is the only schedule selection. Select the daily, weekly, or monthly recurring time for backups to run.

To perform a nonrecurring backup, or to perform a backup immediately, run the script:

```
fmbackup
```

When the backup is complete, the system returns the message:

```
Done. The backup has been output into /tmp/DUMPFILe.
```

where DUMPFILe is a unique filename generated by the backup script.

Restoring a backup dump

Backups are typically restored after a system failure. To restore a backup, start with a freshly installed appliance. Steps are performed from both the GUI and the command line.

1. Configure networking. Configure the CTA network provides details.
2. Configure the hostname, domain, and DNS servers. Configure the hostname, domain, and DNS server provides details.
3. Configure the destination for the restored backup files (.tgz).
 - a. If the backup files were written to Centera, configure Centera as the destination for the restored files. [Deploying CTA with Centera](#) provides details.
 - b. If the backup files were written to a NAS repository, configure a NAS repository as the destination for the restored files. Configuring a NAS-based repository provides details.
4. Mount the NFS export where the backup catalog file (DBBackup.out) is stored. This is the disaster recovery location described in [Creating a backup dump](#).
5. Copy DBBackup.out to /opt/rainfinity/filemanagement/conf.
6. From the GUI, select **Configuration** and **Backup and Recovery Settings**.
Under Recover File Management, select the .tgz file to restore and click **Restore**. Using information from the DBBackup.out file, the .tgz backup file will be restored to the /var/fmrestore location.
7. The fmrestote script reconstructs the CTA system configuration from the .tgz file restored in step 6.
8. To run the script in a screen session that will not be interrupted, type:

```
screen  
fmrestore <backup_file.tgz>
```

As the restoration occurs, the system will prompt for input to:

- a. Confirm restoration.
- b. Start the FPolicy callback service for a NetApp.
- c. Start the callback daemons for VNX and for the cloud service.

At each prompt, type **y**. When asked if you want to add another server, type **n**.

Note The database tables are restored at a rate of approximately 10 million entries per hour.

If restoring data to the same machine, the CTA automatically restarts at the conclusion of the restoration process. If restoring data to a different machine, the CTA must be manually restarted. Also, original network configuration files, such as /etc/hosts, may need to be manually edited to reflect the new IP and hostname of the new machine.

Typical output of the **fmrestore** script is as follows:

```
[root@fm2 bin]# fmrestore /var/fmbackup_9.0_cta.Thu_26-01-12_16_41.tgz
Expanding /var/fmbackup_9.0_cta.Thu_26-01-12_16_41.tgz in /var...
This will overwrite your configuration and database. Are you sure?
Press 'y' to continue or 'n' to abort now
Y
Stopping CTA GUI ...
Stopping Tomcat server
Tomcat server stopped
Stopping filemanagement daemon ...
Stopping filemanagement daemon watchdog
Stopping filemanagement daemon done

done

Empty the current database...

Removing Unique key constraints ...

Restore configuration and database...
Starting network time protocol daemon (NTPD) done
Warning: backup is missing file: /etc/sysconfig/network

Adding Unique key constraints ...

Updating Database index statistics...

Starting CTA GUI ...
Starting Tomcat server
Tomcat server started

Starting filemanagement daemon ...
Starting rslogd (already running): done

Starting rslogd Monitor (already running): done
Starting rslogd cpu check (already running): done

Starting filemanagement daemon done

Starting filemanagement daemon watchdog done

done

Restore Done..
```

Encryption Key restore in case of catastrophic failures

If CTA experiences a catastrophic failure (that is, a hardware failure) and must be restored from backup, you must take special care to ensure that keys are not lost.

Recovery of the Main CTA

It is possible that keys were generated and replicated after the most recent backup was created. In this case, restoring the keystore from backup would result in keys being lost. If these keys had previously been used to encrypt archived files, the contents of those files will be unrecoverable by the main CTA. Use the `krd` executable

in command line mode by the system administrator to recover the keystore from the CTA-HA, which is presumed to have the most recent copy. The krd will ensure that a local keystore will not be overwritten with a new keystore that contains fewer keys than the original.

Recovery of the CTA-HA

No special action is necessary. The main CTA is presumed to have the most recent copy of the keystore. Use the krdsetup script as if the CTA-HA is being installed for the first time.

Maintaining the database

After archiving millions of files, archiving tasks may become slow as the number of entries in the archival database grows larger. To improve performance, use a CTA vacuum process to clear the database of unused entries and reindex the entries that remain.

The database maintenance process can take several hours. While the process is running, the filemanagement daemon must be halted and the GUI may not be used. System administrators should plan to run database maintenance when the appliance is not needed.

Recalls are not interrupted by database maintenance.

Checking database size and disk capacity

Before performing database maintenance, verify that the maintenance is warranted.

1. Check the current database size. Click the Archived Files tab:
2. Check how much disk space the database is using by typing:

```
du -sh /var/lib/pgsql/data/base
```

The database should use about 1 million archived files per gigabyte. If the database is using more than twice the expected disk space, for example, if 1 million archived files are occupying more than 2 gigabytes, perform database maintenance.

Performing database maintenance

The command line script that starts database maintenance performs the following steps:

1. Stops the filemanagement daemon and GUI.
2. Runs the database vacuum process.
3. Restarts the daemon and the GUI.

Database maintenance can be run immediately as a CTA task or scheduled to run at a specific time.

Note Before running the CTA vacuum process, it is recommended that you backup the CTA system and copy the backup file to an off-host location.

- To start the CTA vacuum process immediately, type:

```
rffm doDBMaintenance
```

- To schedule the CTA vacuum process to start at a later time and to run repeatedly, type:

```
rffm scheduleVacuumTask VacuumStartTime=<StartTime>
VacuumWeekRepetition=<Number_of_weeks>
```

where:

- StartTime* is the date and time when the first vacuum task will run. The format is yyyy-mm-dd hh:mm:ss. After each run, the value is reset to the date and time of the repeat run.
- Number_of_weeks* is the number of weeks between runs. It is specified as an integer and the value should be between 4 and 24. The default is 8.

Every 30 minutes, CTA checks to see if there are any vacuum tasks scheduled.

One hour before the vacuum task is scheduled to start, CTA sends the alert “A VACUUM TASK WILL START WITHIN THE NEXT 1 HOUR”.

For example, to start the CTA vacuum process at 2 a.m. on May 1, 2011 and run the task every four weeks, type:

```
rffm scheduleVacuumTask VacuumStartTime="2011-05-01 02:00:00"
VacuumWeekRepetition=4
```

- At or around 1 a.m. on May 1, 2011, CTA sends the alert that the vacuum process will start within the hour.
 - Following the run, the StartTime is reset to at 2 a.m. on May 29, 2011.
- To delete the CTA vacuum process, type:

```
rffm deleteVacuumSchedule
```

The log of the process is: /var/lib/pgsql/vacuum.log

The output of the process is:
/opt/rainfinity/filemanagement/conf/DBMaintenance.log

Migrating from CTA to CTA/VE

The procedure to migrate from CTA to CTA/VE involves backing up the CTA database and restoring it to a CTA/VE. The migration path consists of the following steps:

Note *Migration* can only occur between systems running the same CTA versions (for example, between two systems running CTA 12.0 to 12.0). If two systems are running different CTA variants, one or both should first be upgraded so that they are running the same variant before migration occurs.

1. Upgrade the FMA to CTA version 7.5 or later.
2. Back up the current FMA configuration to a Centera or NAS repository. If preserving additional custom configuration or task simulation and support logs, back up the files separately.
3. Install CTA/VE version 7.5 or later on the ESX Server and configure networking.
4. Copy the backup file to the CTA/VE.

5. Shutdown the FMA. Remove it from the network.
6. Manually reconfigure the IP address of the CTA/VE to replace the CTA on the network.
7. Restore the backup file on the CTA/VE.
8. Reboot the CTA/VE.

Note The network bonding feature, which is available on the CTA, is not available on the CTA/VE.

If a CTA is being upgraded to CTA/VE, and the VNX or NetApp implementation includes a CTA-HA for recall:

1. Use the CTA-HA to provide recall functionality while upgrading the CTA to CTA/VE.
2. Once recall has been tested and verified for the CTA/VE following the upgrade:
 - a. Use the CTA-HA to provide recall for the CTA/VE, or
 - b. Install a CTA/VE-HA to provide recall for the CTA/VE

Shutting down and restarting the appliance

To shut down and restart a working CTA or CTA/VE:

1. Stop any running tasks.
2. Stop all services with the commands:

```
filemanagement stop
celerracallback stop
atmoscallback stop
fpolicycallback stop
```
3. If an HA appliance is deployed, verify that the HA appliance has taken over file recalls by attempting to open an archived file.
4. Either shut down or reboot the appliance.
 - a. To shut down the appliance, type the command:

```
shutdown -h now
```
 - b. To reboot the appliance, type the command:

```
reboot
```

For the CTA-HA, only the callback services are stopped. The **filemanagement stop** command is not used.

Chapter 6 Cloud Tiering Appliance System Settings

This chapter presents the following topic:

Security hardening	102
Configuring the GUI access method	105
STIG hardening.....	105
LDAP client configuration	107
Certificate management.....	111
Appliance mail delivery settings.....	111
Log settings	112
Configuring alerts	114
System command accounting	116
Windows domain user	118
SID translator	120

Security hardening

By default, security hardening is not enabled.

To configure security hardening:

1. Start the Cloud Tiering Appliance setup tool, type **rfhsetup**.
2. Select **Configure System Security**. A set of security settings options appears.
3. Select **Harden Appliance**.

The default settings for the items that affect the appliance security level are:

- a. Use single security database =no
- b. Disable root logins =no
- c. Strengthen passwords =no
- d. Age passwords =no
- e. Harden to STIG requirements =disabled

When all four settings are “no,” security hardening is disabled and this disabled security level is referred to as the default level.

If any of the settings is set to a non-default value, security hardening is enabled.

Note In addition to the security settings, the GUI access method may also be configured from the **Harden Appliance** menu. By default, the GUI is accessible over both http and https. Enabling **https only** or **redirecting http to https** does not change the appliance setting to hardened.

Single security database

If the single security database setting is enabled, all authentications on the device will go through standard Linux Pluggable Authentication Modules (PAMs). This applies to both GUI and CLI access.

Both the GUI and the CLI provide two types of users:

- Admin users belonging to the wheel group and Rainfinity groups
- Ops users belonging to the Rainfinity group

CLI users are configured independently from the GUI users.

If **rfhsetup** is used to enable the security database setting option, existing GUI users will not be able to log in. Use **rfhsetup** to disable root logins, and then create the new user.

To disable root logins:

1. Start the Cloud Tiering Appliance setup tool, type **rfhsetup**.
2. Select **Configure System Security**. A set of security settings options appears.
3. Select **Harden Appliance**.
4. Select **Configure root login settings**.
5. For **Allow login as root?**, select **N**.

Note Once root login is disabled, you must log into the CTA with the new username and password to use the CLI.

Admin users

An admin user who is a member of the wheel group and logged in through SSH can become a superuser to:

- Create/delete other users
- Run **rfhsetup**

To add an admin user for access from the CLI:

1. Log in to the CTA as **root**.
2. Type the following commands:

```
useradd -G rainfinity,wheel <username>  
passwd <username>
```

Ops users

An ops user belongs to the Rainfinity group.

To add an ops user for access from the CLI:

1. Log in to the CTA as **root**
2. Type the following commands:

```
useradd -G rainfinity <username>  
passwd <username>
```

Linux PAM users

A Linux PAM user is created through the CLI. When a Linux PAM user is logged in to the GUI with the single security database setting enabled, the user's role (admin or ops) is cached for the duration of the session.

If the administrator changes the user's setting while the user is logged in, the user's role will not be refreshed until one of the three following conditions occurs:

- User logs out.
- GUI is restarted.
- Cached user information in the Tomcat server expires due to inactivity.

Adding users with the GUI

To add a new admin or ops user with the GUI:

1. Log in as **admin**.
2. From the **Configuration** tab, select **Cloud Tiering Appliance Users**.
3. Select **Add a New User**. In the **Cloud Tiering User Properties** dialog box that appears:
 - a. Type the name.
 - b. Type a new password.
 - c. Specify the type of user:

- Super User — The admin user.
- Regular User — The ops user.

Note When the single security database setting is disabled, users created through the GUI are allowed to log in through the GUI but not the CLI. In addition, if the single security database setting is enabled, user accounts cannot be created through the GUI. If the user attempts to invoke the configuration page for Cloud Tiering Appliance Users, a warning appears.

Disable root logins

If root logins are disabled, the only way to add new users or to run **rfhsetup** is for an admin user (such as a user who belongs to the wheel group) to log in to the device, and then become a root user.

When the setting to disable root logins is being changed to yes, the CTA checks to ensure that:

- There is at least one admin user other than root who belongs to the wheel group. This user must have a configured password.
- The wheel users are in the local `/etc/group` file. The CTA ignores LDAP users while performing this check because LDAP servers occasionally become unreachable.

Note Configure a small set of admin users locally for each CTA. Most admin and ops users are configured on an LDAP server. In this way, the management of these users scales to large networks.

Strengthen passwords

If the **passwd** command is run with password strengthening enabled, your new password must be at least eight characters long and satisfy the following requirements:

- At least three characters are different from the previous password.
- At least one character is an uppercase letter.
- At least one character is a number.
- At least one character is a special character.

In a clustered environment, run the **passwd** command on both the primary and backup nodes.

Note The root user can change any password including its own to any value, regardless of the password strengthening setting to strengthen it.

Age passwords

If password aging is enabled, every user (except root) who can log in with a shell account will have an aging password. The root user configures:

- When to print a user warning that a password is about to expire.
- The maximum number of days a password can remain valid before it must be changed.
- How often a password may be changed.
- The number of days following password expiration after which the account will be locked. Once an account is locked, only the root user can unlock the account by using the change command to change the age of the password.

Note If a large number of devices are deployed, a central authentication service (such as LDAP) should be used. Password administration through the central site greatly facilitates user scalability, as one user is not required to log in to every deployed CTA to update an aging password.

Configuring the GUI access method

By default, the GUI can be accessed by both HTTP and HTTPS. To change this for the CTA:

1. Start the Cloud Tiering Appliance setup tool, type **rfhsetup**.
2. Select **Configure System Security**. A set of **security settings** options appears.
3. Select **Harden Appliance**.
4. Select **Configure GUI access method**:
 - a. To disable access over HTTP, select **Only enable GUI access over https**.
 - b. To redirect http traffic to HTTPS instead of disabling HTTP, select **Redirect GUI access over http to https**.

STIG hardening

Security Technical Implementation Guide (STIG) is a set of security guidelines issued by the US Department of Defense. These STIG UNIX guidelines define how UNIX/Linux appliances should behave from a security standpoint.

Enabling STIG hardening

The CTA provides an option for hardening the appliance to meet the UNIX STIG Guide (Version 5, Release 1). When STIG hardening is enabled, the security settings change as follows:

- The user must type the root password to gain access to the CTA in single user mode.
- After three consecutive login attempts, the account is disabled.
- Only the root user can reenableView a disabled account.

- The login delay between login prompts increases from 2 to 4 seconds.
- New passwords are required to be a minimum of nine characters in length.
- When changing passwords, the past five passwords cannot be reused as the new password value.
- The root account's home directory will be set to a permission value of 700.
- Man page file permissions will be set to 644.
- User-directories must not contain undocumented startup files with permissions greater than 750 (that is, they must allow write access only for that user).
- The system and default user umask must be set to 077.
- Access to the cron utility will be restricted using the cron.allow and cron.deny files.
- Crontab file permissions above 700 will not be permitted (in the /etc/cron.daily, /etc/cron.hourly, /etc/cron.weekly directories).
- The inetd.conf file permissions will be set to 440.
- Unnecessary accounts, for example, games and news will be deleted.
- sysctl.conf file will be set to 600 permission.

To enable STIG hardening on the CTA and CTA-HA:

1. Start the Cloud Tiering Appliance setup tool, type **rfhsetup**.
2. Select **Configure System Security**.
3. Select **Harden Appliance**.
4. Select **Harden to STIG requirements**.
5. When prompted with Enable changes to conform to STIG Hardening requirements?, Type **y**.

Disabling STIG hardening

When STIG hardening is disabled, the security settings change as follows:

- No password prompt is made prior to connecting in single-user mode.
- User accounts are unlocked, even after three or more failed login attempts.
- The login delay is set to the current default setting, which is less than four seconds at this time.
- When changing passwords, the minimum length must be:
 - a. If password hardening is enabled: eight characters, with at least one lowercase, one uppercase, one digit, and one special character.
 - b. If password hardening and STIG hardening are disabled: the minimum requirements for the new password is that it should be six characters long.
- When STIG hardening is disabled, the user can reuse previously set passwords.

- The /root directory permissions is reset to 750.
- Man page file permissions remain at 644. That is, this STIG hardening change is retained.
- User-directory permissions remain at the value prior to STIG hardening.
- The system and default user umask must be set to 022.
- Unnecessary groups/accounts that are deleted during STIG hardening remain deleted even after STIG hardening is disabled.
- Access to the cron utility is unrestricted using the cron.allow and cron.deny files.

To disable STIG hardening on the CTA:

1. Start the Cloud Tiering Appliance setup tool, type **rfhsetup**.
2. Select **Configure System Security**.
3. Select **Harden Appliance**.
4. Select **Harden to STIG requirements**.
5. When prompted with Enable changes to conform to STIG Hardening requirements?, Type **n**.

STIG hardening is disabled when the appliance hardening level is reset to the default level as follows:

1. Start the Cloud Tiering Appliance setup tool, type **rfhsetup**.
2. Select **Configure System Security**.
3. Select **Remove Appliance Hardening Settings**.

LDAP client configuration

LDAP directory trees represent hierarchical directory information, such as people and phone numbers that belong to an organization. The CTA supports Lightweight Directory Access Protocol (LDAP) for user authentication and authorization.

Global LDAP settings

Global LDAP settings affect all LDAP operations. The following settings impact how the LDAP client on the CTA behaves if the LDAP server does not respond.

Bind type

There are two types of binds:

- **Hard** — The CTA continues to retry the bind attempt until a maximum timeout is reached.
- **Soft** — The CTA attempts to bind once and abort if the server does not respond.

Time limits

There are two types of time limits:

- Search time limit — The amount of time for which the LDAP client waits for an initial response from the server.
- Bind time limit — The amount of time for which the LDAP client attempts to bind.

By default, these time limits are set to 10 seconds to allow the appliance to remain responsive when the LDAP server is down, and to fail over to an alternate authentication mechanism, if another mechanism is configured.

Server type

The CTA LDAP client works with the following types of LDAP servers:

- OpenLDAP
- Active directory with RFC 2307 support

Identity Management for Unix should be installed on the ActiveDirectory server for CTA to be able to authenticate an ActiveDirectory user. The AD user used to log into CTA must have the NIS domain setup correctly on the AD server. The user should be added to the wheel and rainfinity groups to have administrator privileges.

LDAP authentication

When LDAP is configured, LDAP authentication is established through a sequence of events.

- A user connects to the CTA. The user is challenged for to authenticate.
- The CTA LDAP client contacts the LDAP server to validate the user's credentials.
- To validate that the client is trusted, the server attempts to establish a secure communication channel with the client and then authenticate by using a plain-text password or SASL.

The client does the following to establish the secure communication channel:

- a. Requests the server's public key.
- b. Validates that the server's public certificate is signed by a known Certificate Authority (CA).
- c. Encrypts its data using the server's public certificate. Only the private key stored on the server can decrypt this data.

Initial data from the client contains negotiation information that the server and client use to establish a secure communication channel.

Just as the client uses the server's public key to encrypt its first message, the server ensures that the client is authentic by requesting the client's public certificate and validating that it is signed by a known Certificate Authority.

After the secure channel is established, the password is exchanged. If SASL is configured, it can be used instead of a password.

- The server and client can negotiate an encryption scheme to secure all traffic between them.

Once authentication is established and an encryption scheme is optionally selected, the LDAP client requests user authentication.

Configuring LDAP settings

To start LDAP configuration:

1. Start the Cloud Tiering Appliance setup tool, type **rfhsetup**.
2. Select **Configure System Security**.
3. Select **Configure LDAP**.
4. Select **Enable LDAP**.
5. Configure the basic LDAP settings:

- a. Maximum time the LDAP client waits for an initial response from the server

Type a period of time. The client retries after waiting for 2 seconds and thereafter continues retrying after doubling the wait time from the previous retry attempt. The client continues to retry until either the server responds or the configured LDAP search time limit is exceeded. The default time limit is 10 seconds.

- b. LDAP bind policy

Select **soft** or **hard**. The default setting is **hard**, which indicates that the client will retry bind connections to the LDAP server.

- c. Maximum time for which the LDAP client waits for a bind response from the server

Type a period of time. If the bind policy is set to **soft**, this setting has no effect. If the bind policy is set to **hard**, this policy causes a bind retry mechanism to occur.

- d. LDAP server type

Select from the supported server types:

- OpenLDAP — Applies to LDAP servers distributed by OpenLDAP.
- Active Directory with RFC2307 support

Note: If you do not have a Unix Attribute tab in the AD settings, you must install Identify management for Unix, which is available from Microsoft. This generates the Unix Attribute tab.

Other LDAP servers have not been validated for CTA version 7.2 or later.

- e. IP address or hostname for the LDAP server

If using Open LDAP, for which certificate-based authentication is enabled, type the hostname that matches the hostname used in the certificate generation. If an IP address was used in the certificate generation instead of the hostname, type the IP address.

Note Failure to type the proper information will create problems during the LDAP setup. This is one of the most common configuration errors during LDAP setup.

f. LDAP basedn

Type the suffix for your domain name.

g. LDAP binddn

Specify the LDAP binddn for the Active Directory.

h. LDAP bindpw

Password for the specified binddn. This is specific to AD configuration.

i. LDAP CA Certificate Path

Path to the Certificate Authority certificate stored on the CTA.

j. LDAP-enabled SASL

To configure SASL, provide the following information:

- SASL KDC address
- Domain name
- Kerberos principal details
- Kerberos keytab file
- Name server IP address

Note With LDAP enabled, CTA users can log into the CLI.

To log into the GUI:

- The single security database (SSD) setting must be enabled on the CTA.
- Users allowed to log in to the CTA must be added to the **rainfinity** and **wheel** groups in the LDAP server database.

If the GUI is running and LDAP is enabled through `rfhsetup`, the GUI will not recognize LDAP authentication attempts until it is restarted by typing the following command:

```
/opt/rainfinity/filemanagement/bin/fmgui restart
```

To avoid this problem:

1. Enable external authentication (LDAP) before enabling the single security database.
2. Invoke the GUI.

Note If changing the CTA IP address or moving the CTA to a new location, disable remote login and re-enable root login before changing or moving the CTA. After changing or moving the CTA, re-enable LDAP with the IP address or hostname of a LDAP server that is reachable in the new configuration.

Certificate management

When configuring LDAP, TLS, and SSL for authentication, key and certificate files are required. In order for authentication encryption to work correctly, these keys and certificates must be:

- Periodically refreshed
- Correctly located on the appliance

Each certificate has an expiration date. Every week, the CTA checks the validity of each certificate. Certificate warning information is logged into the `/var/log/secure` file, and if the alert is enabled, email is sent when the certificate is due to expire. Once a certificate expiration warning is received, SSL/TLS certificates must be updated.

To update and manage the keys and certificates:

1. Start the Cloud Tiering Appliance setup tool, type **rfhsetup**.
2. Select **Configure System Security**.
3. Select **Certificate Management**.
4. To update either the Certificate Authority (CA) public certificate or the Client key and certificate for use with SSL/TLS:
 - a. Select **Update Certificate**.
 - b. Type **Y**.
 - c. Type the scp path from which the selected certificate or key file will be copied to the CTA.

Appliance mail delivery settings

The CTA supports delivery of alerts through email. To send these alerts, sendmail must be properly configured. A menu is provided within the `rfhsetup` tool. To use this menu:

1. Start the Cloud Tiering Appliance setup tool, type **rfhsetup**.
2. Select **Configure CTA Mail Delivery Settings**. The **Appliance Mail Configuration** menu appears.

Follow the prompts to configure:

- a. **Change Configuration** — When prompted, type **y**.
- b. **Sender's email address** — Type the address that will appear in the From field of the alert emails sent by the CTA. For example, **johndoe@acme.com**.
- c. **SMTP server** — Type the server to which mail should be sent. For example, **mailhub.eng.acme.com**.
- d. **email verification** — Type a recipient email address to which test emails may be sent. For example, **adminjoe@acme.org**. The `rfhsetup` script will attempt to verify the mail configuration by sending two emails.

Wait a few minutes. Check the email account to see if these emails were successfully received.

3. **Mail Test 1** — To confirm the receipt of an email with the subject Mail Test 1, type **y**. Otherwise, type **n**.
4. **Mail Test 2** — To confirm the receipt of an email with the subject Mail Test 2, type **y**. Otherwise, type **n**.

If either of the test emails was received, mail delivery is working and mail setup is done.

If neither test email was received, verify:

- The name of the SMTP server. Check with your system administrator.
- The email address provided for the test email.
- The SMTP server is reachable. Try to ping it.

Log settings

When the security level is set to harden, any event that might affect the security of the system is written to the CTA log files. Use the Cloud Tiering Appliance setup tool to administer and preserve log files.

Configuring log rotation

With log rotation, the user controls the periodic rotation of files.

To configure log rotation:

1. Start the Cloud Tiering Appliance setup tool, type **rfhsetup**.
2. Select **Configure Logging Options**.
3. Select **Configure Log Rotation**.
4. Follow the prompts to configure:
 - a. Log rotation frequency — Daily, weekly, or monthly
 - b. Rotation mode — Size or time
 - c. Max log size (for non-debug files)
 - d. Max debug log size
 - e. Number of copies to keep for each log file

Configuring SCP of rotated log files

Log rotation is the first step in archiving the CTA system logs. These log files are eventually deleted as a part of the normal rotation process. However, in many customer environments, it may be necessary to preserve these files by copying them to a remote server. Use the CTA to create a tar file of these rotated system and the CTA logs, then secure copy them to a remote server.

Configuring the public-private key exchange

Prior to configuring secure copy (SCP) of rotated log files, a public-private key exchange must take place.

To configure the public-private key exchange:

1. Log in to the CTA or CTA-HA as root.
2. Generate the public key by typing `ssh-keygen -t rsa`.
 - a. When prompted, press **Enter** to accept default answers for:
 - b. File in which to save the key, for example, type: `/root/.ssh/id_rsa`
 - c. No passphrase
 - d. Confirm no passphrase
 - e. At the end of the configuration, a message appears acknowledging:
 - f. Your identification is saved in `/root/.ssh/id_rsa`.
 - g. Your public key is saved in `/root/.ssh/id_rsa.pub`.
3. For the external server where the log files will be placed, create a user with write access to the copy directory. Do not use the root user.

Note In the following steps, *server* is the IP address or hostname of the external server, and *user* is the name of the user on the external server which will copy the files.

4. Log in to the CTA or CTA-HA and use SSH to:

```
Create the directory ~/.ssh by typing the command:  
ssh <user>@<server> mkdir -p .ssh
```

- a. Type the user password.
 - b. Append the public key on the CTA or CTA-HA by typing the command:

```
cat /root/.ssh/id_rsa.pub | ssh <user>@<server> 'cat >>  
.ssh/authorized_keys'
```
 - c. Type the user password.
 - d. Set correct permissions by typing the command:

```
ssh <user>@<server> chmod -R 700 .ssh
```
 - e. Type the user password.
5. To verify successful completion, attempt to log in to the external server as **user** from the root account on the CTA by typing:

```
ssh <user>@<server>
```

You should not be prompted for a password.

You can now successfully use SCP without a password to send the rotated log files to your external server.

Configuring SCP of rotated log files by using `rfhsetup`

Once the public-private key exchange is completed, configure scp of rotated log files:

1. Start the Cloud Tiering Appliance setup tool, type `rfhsetup`.
2. Select **Configure Logging Options**.
3. Select **Configure SCP of Rotated Log Files**.
4. Follow the prompts to configure:
 - a. The SCP Remote Address — The IP address or hostname of the external server. This is the external server referenced in [Configuring the public-private key exchange](#).
 - b. The username to whose account the log files will be copied — The name of the user on the external server who will copy the files. Same as the user provided in [Configuring the public-private key exchange](#).
 - c. The full path to the directory at the remote site where the log files should be placed. The user must have write access to this directory.

Following the configuration, the CTA will test SCP by attempting to copy a test file. If this test fails, the SCP settings will be accepted, but SCP is probably not configured properly. Correct the error that is blocking SCP and rerun the Cloud Tiering Appliance setup tool.

Configuring alerts

A CTA can be configured to monitor various system log files and send email to alert whenever an event occurs. CTA sends notifications for SNMP traps and alerts listed in [Appendix B, Alerts](#).

Configuring email alerts

Use the GUI to review and configure the list of email alerts:

1. Click the **Alert Settings** link on the **Configuration** tab.
2. Click the **Edit log alert Pattern** link.

A list of alerts with the various alert settings appears:

- a. Alerts may be individually enabled.
- b. If alerts occur more than once within a specified time period, edit the throttle time to suppress the repeated alerts. A different throttle time may be applied to each alert.

Only admin users can view this configuration page.

To configure email alerts from the command line:

1. Start the Cloud Tiering Appliance setup tool, type `rfhsetup`.
2. Select **Configure Logging Options**.
3. Select **Configure Log Alerts**.
4. Follow the prompts to configure:

- a. When asked to enable alerts, type **y**.
- b. Specify one or more email addresses separated by a space or comma, to receive the alerts.

Configuring SNMP alerts

Use the GUI to configure SNMP alerts:

1. Click the **SNMP Configuration** link on the **Configuration** tab.
2. On the SNMP Settings page that appears, add a notification host. This is the host to which alerts will be sent:
 - a. IP address
 - b. UDP port (typically SNMP uses UDP port 161 for general SNMP messages and UDP port 162 for SNMP trap messages)
 - c. Community string
 - d. Security type
3. Click **Commit**.
4. Click the **Alert Settings** link on the **Configuration** tab.
5. Under Alerts, click **Enable SNMP alerts**.

Only admin users can view this configuration page.

To configure SNMP alerts from the command line:

1. Configure the SNMP Notification Host:
 - a. Start the Cloud Tiering Appliance setup tool, type **rfhsetup**.
 - b. Select **Configure Logging Options**.
 - c. Select **Configure SNMP**.
 - d. Select **Configuration SNMP Notification Hosts**.
 - e. Add the SNMP Notification Hosts:
 - The number of hosts that may be added is unlimited.
 - For each host, specify: IPv4 address, UDP port number, SNMP community string, and SNMP version.
 - The community string must be alphanumeric, and may include dashes and underscores.
2. Enable SNMP alert generation:
 - a. Start the Cloud Tiering Appliance setup tool, type **rfhsetup**.
 - b. Select **Configure Logging Options**.
 - c. Select **Configure Log Alerts**.
 - d. Follow the prompts to configure:
 - When asked to enable alerts, select **Yes**.

- Specify the type of alert delivery. Select either **email only**, **SNMP only**, or **email and SNMP**.

Enabling SNMP polling

Use the GUI to enable SNMP polling:

1. Click the **SNMP Configuration** link on the **Configuration** tab.
2. On the SNMP Settings page that appears:
 - a. Type a community string.
 - b. Select a security type.
 - c. Click **Add**. The community string is added to the Current Community String list.
3. Click **Commit**.

To enable SNMP polling from the command line, configure the SNMP Community String to be used for polling:

1. Start the Cloud Tiering Appliance setup tool, type **rfhsetup**.
2. Select **Configure Logging Options**.
3. Select **Configure SNMP**.
4. Select **Configuration SNMP Community Strings**.
5. Add the SNMP Community Strings.
 - a. The number of strings that may be added is unlimited.
 - b. For each string, specify the SNMP community string and SNMP version.
 - c. The community string must be alphanumeric, and may include dashes and underscores.

Note To poll for SNMP objects without enabling **rfalrtd**, execute the command: **service rfsnmp start** from the root account. This restarts SNMP and no alert history is viewable until the alert daemon is restarted.

System command accounting

The CTA provides the ability to track any command that is successfully executed and launches a new process.

To track command history, the CTA uses the psacct Process Accounting package. This package tracks commands that are entered. In addition to commands, the CTA extends this package to track command arguments.

To enable System Command Accounting on the CTA:

1. Start the Cloud Tiering Appliance setup tool, type **rfhsetup**.
2. Select **Configure Logging Options**.
3. Select **Configure System Command Accounting**.
4. Type **y** to enable system command accounting.

Tracking user command history

After enabling System Command Accounting, admin users can track the list of commands entered on the system with the tool: `/opt/rainfinity/bin/rflastcomm`.

To use this tool, admin users must be a superuser.

Examples of its use are as follows:

- To list the commands entered by all users, use the tool without any options, or:

```
/opt/rainfinity/bin/rflastcomm
```

- To list the commands entered by a specific user, type:

```
/opt/rainfinity/bin/rflastcomm -u <username>
```

- To list commands entered by a user since a start date on 5 p.m. on June 6, 2007, use the tool with the following arguments:

```
/opt/rainfinity/bin/rflastcomm -u <username> -s '2007-06-06 17:00:00'
```

- To track system/daemon/session history, type:

```
/opt/rainfinity/bin/rfquerycshis.sh
```

- For a help menu and additional options, type:

```
/opt/rainfinity/bin/rflastcomm --help
```

Tracking user login history

After enabling System Command Accounting, admin users can track the login history with the tool `/usr/bin/last`.

To run this tool, admin users must **su** as root first.

This tool is part of the standard psacct Process Accounting package. For detailed info on using this tool, type: **man last**.

Tracking daemon command history

To query daemon command history, such as **xmlrpc** commands issued to the daemon from the GUI or through various CTA CLI commands, use the tool: `/opt/rainfinity/bin/rfquerycshis.sh`.

- To obtain the daemon command history, type:

```
/opt/rainfinity/bin/rfquerycshis.sh -t dc
```

- To query the system command history, type:

```
/opt/rainfinity/bin/rfquerycshis.sh -t sc
```

- To query the user login history, type:

```
/opt/rainfinity/bin/rfquerycshis.sh -t ls
```

- To list hardware related messages from the system log files, type:

```
/opt/rainfinity/bin/rfquerycshis.sh -t hw
```

Windows domain user

When a new file server is added to the CTA configuration, CIFS specific settings include the username and password for the Windows domain user to be used by the CTA. Before adding a new CIFS file server, use the instructions in the following sections to set up the Windows domain user:

- [Creating a Windows domain user](#)
- [Adding an admin user to the local administrator group](#)

In addition, for CIFS authentication configure either:

- Fully Qualified Domain settings — For Kerberos support as described in [Configuring Windows for Kerberos](#).
- Domain controller Group Policy Object (GPO) — For NT LAN Manager (NTLM) support as described in [Configuring Windows 2008 for NTLM](#).

Note CIFS specific settings are applicable for all servers except Centera, Atmos, and Amazon S3.

Creating a Windows domain user

To create an administrator in the Windows 2000, 2003, or 2008 domain:

1. Log in to the primary domain controller as the **Domain Administrator**.
2. From the **Start** menu, select **Start > Programs > Administrative Tools > Active Directory Users and Computers**.
3. Right-click **Users**.
4. Select **New > User**. The **New Object — User** dialog box appears:
 - a. In the **Full name** box, type **CTA Administrator**.
 - b. In the **Login name** box, type **fadmin**.
The fadmin login is the CTA Administrator Windows Domain user.
 - c. Type a password.
This password is the fadmin Windows password.
 - d. (Optional) Select **Password Never Expires**.
5. Click **Finish**.

Note If you have NetApp filers but no Windows 2000, 2003, or 2008 servers in your domain, then you must include fadmin in the domain administrator group. Otherwise you will not be able to include the fadmin user in the NetApp filers' administrators group.

Adding an admin user to the local administrator group

The fadmin account must be added to the administrators group on the CIFS file servers that will be involved in CTA archiving. To add a CTA Windows domain user on a NetApp filer or a VNX Data Mover:

1. Log in to the primary domain controller as the **Domain Administrator**.
2. From the **Start** menu, select **Start > Programs > Administrative Tools > Computer Management**. The **MMC** application appears.
3. To start a Computer Management session with the file server:
 - a. From the **Action** menu, select **Connect to another computer**. The **Select Computer** dialog box appears.
 - b. Click **Browse** or type the file server name to select the NetApp, VNX, or Unity system to connect to.
 - c. Click **OK**.
4. To include the fadmin user in the administrator group for the CIFS file server:
 - a. Under **System Tools**, in the folder **Local Users and Groups**, select **Groups**.
 - b. Select **Administrators**. The **Administrators Properties** dialog box appears.
 - c. Click **Add**. The **Select Users or Groups** dialog box appears.
 - Click **Locations**. From the **Locations** menu, select the domain instead of the local computer.
 - Under **Enter the object names to select**, type **fadmin** to add the domain user.
 - d. Click **OK**. The **Administrator's Properties** dialog box reappears with the newly added fadmin user.
 - e. Click **OK**.

Repeat this process for any other file servers that will be involved in CTA archiving.

Configuring Windows for Kerberos

By default, the Windows domain controller supports Kerberos authentication. Before configuring a file server for CIFS authentication with Kerberos, configure the Fully Qualified Domain Settings on the CTA:

1. Log in to the CTA GUI.
2. Select **Configuration** and **Fully Qualified Domain**.
CTA displays the Fully Qualified Domain List.
3. Click **Add** to add a new Fully Qualified Domain.
CTA displays a Fully Qualified Domain Properties dialog box.
4. Enter the **Domain Name** and **IP Address** of the Domain Controller.
5. Click **Commit**.

Configuring Windows 2008 for NTLM

If using NTLM for CIFS authentication instead of Kerberos, confirm that the domain controller is configured for NTLM:

1. Log in to the Windows 2008 domain controller as the **Domain Administrator**.
2. From the **Start** menu, select **Run**. In the **Run** dialogue box that appears, type **gpmmc.msc** and click **OK**. The **Group Policy Management** dialog box appears.

3. Expand the domain. Under **Group Policy Objects**, right-click **Default Domain Policy** and select **Edit**. The **Group Policy Management Editor** appears.
4. Under **Computer Configuration**, select **Policies > Window Settings > Security Settings > Local Policies > Security Options**.
In the list of policies, scroll down to **Network security: LAN Manager Authentication**. Confirm that the policy setting shows that NTLM is configured for authentication.
5. Close the **Group Policy Management Editor**.

SID translator

The CTA is able to translate Security IDs (SID) in the security properties of the files and directories involved in a CIFS file migration task. The capability may be used to assist projects in which the data's group or user association changes from the source to the destination. For example, when the Access Control List (ACL) is defined in terms of local groups on the source file server.

When the data is migrated to the destination server, the ACL should be defined in terms of corresponding local groups. The rules governing this translation are defined in the SID translation files.

Installing the SID translator

The SID translator is packaged as a Microsoft Installer package, known as an MSI file. The SID translator supports systems running Windows XP Professional (SP1 or later) and Windows 2003 Server Standard and Enterprise editions.

.NET Framework 4 must be installed on the Windows server before installing the SID translator utilities. The latest .NET Framework may be downloaded from Microsoft at:

<http://msdn.microsoft.com/en-us/netframework/default.aspx>

To install the SID translator:

1. Verify that your client has the .NET Framework 4 installed.
2. Go to the EMC Online Support web site:
<http://Support.EMC.com>
3. To download the SID translator utilities, select **Downloads**. Search for **Cloud Tiering Appliance**. Download **SIDUtilities<BuildNumber>.msi**, where *<BuildNumber>* is a number similar to 8_0b15.

To start the installation script, run the following file:
SIDUtilities<BuildNumber>.msi

Creating the SID translation file

Use the SID translator to create a SID translation file.

1. Once the utility is installed, the SID translator screen appears:
2. Click **Select** for **The groups and users to migrate** under **Source**. The **Select Groups to Migrate** dialog box appears:
 - a. In the **Source name** box, type the name and click **Query**. The list of local groups on that source appears.
 - b. Select the local groups to migrate. Select only the local groups that are unique to your source. For example, you would not select the Administrators group that also likely exists on other file servers.
 - c. Click **Select**.

The **SID Translator** dialog box reappears with **The Groups to Migrate** listed.

3. Click **Select** for **The container to hold all of the migrated groups** under **Destination**. The **Select Container for Migrated Groups** dialog box appears:
 - a. In the **Destination name** box, type a name and click **Query**. The list of local groups on that destination appears.
 - b. Select the destination for the source and click **Select**.

The **SID Translator** dialog box reappears with the names for the **Source** and the **Destination** listed.

4. Click **Configure** for **Map File Output** under **Options**. The **Map File Output** dialog box appears:
 - a. Type the filename or click **Browse** to select an existing XML file.
 - b. Click **OK**.

The **SID Translator** dialog box reappears with the name of the **Map File Output** listed.

5. Click **Test Run** to test the validity of the setup by creating and removing groups on the destination.
6. Click **Migrate** to create the SID translation file.

Uploading the SID translation file

To upload the SID translation file:

1. Select **Configuration** and **File Migration Settings**. The File Migration Settings page appears.
2. To add a SID translation file:
 - a. Click **Add**.
 - b. Navigate to find the XML file and choose that file. When the file is selected, it is added to the list.

Deleting a SID translation file

To delete a SID translation file:

1. Select **Configuration** and **File Migration Settings**. The File Migration Settings page appears with a list of previously uploaded SID translation files.
2. Select the file to delete and click **Delete**.

Appendix A Network topology scenarios

This appendix presents the following topics:

- Advanced network topologies 124
- VLAN tagging modes for the CTA/VE 127

Advanced network topologies

For many environments, using a single networking interface will satisfy networking requirements.

However, there are cases when more complex topologies are needed.

- Combining ethernet interfaces to form a bonded interface. This topology is used for high availability, to protect the CTA installation from a single point of failure.
- [Configuring the CTA with bonding](#) provides details on how to set up this network topology.
- Using two subnets, one for the NAS primary storage tier, and another for either the NAS/CAS secondary tier or for a management interface.
- [Configuring the CTA with two subnets](#) provides details on how to set up this network topology.
- Using more than two subnets, for example, when there are three teams using a CTA distributed across three different subnets.
- [Configuring the CTA with more than two subnets](#) provides details on how to set up this network topology.

Configuring the CTA with bonding

This configuration applies to the CTA installation and is commonly used when fault tolerance must be built into the networking layer. In this example, eth0+eth1 are combined into a bonded interface that is configured with the balance-rr bonding mode:

1. Start the network configuration menu:
 - a. Type **rfhsetup** from the CTA command prompt to invoke the system setup menu.
 - b. Select **Configure Cloud Tiering Appliance Networking**. The network configuration menu appears.
 - c. Select **Configure Networking**.
2. Add new bond interface:
 - a. Type **A** to add an interface. Use the right arrow key to highlight **Bond**, and press **Enter**.
 - b. When prompted for a name of the new bond, use the up arrow key to autogenerate a name. The name generated is **bond1**. Press **Enter** to complete.
3. Edit new bond setting:
 - a. Use the up and down arrow keys to select the **bond1** interface. Press **Enter** to edit the configuration.
 - b. Specify a value for each item:
 - For Slave, type **eth0 eth1**.

- For Trunking Mode, select **balance-rr**.

Complete other values as needed.

- c. Once the interface configuration is defined, press the left arrow key to exit the current menu. When prompted, select **Yes** to keep the new setting.

Note If a bond interface is active, do not edit the IP addresses of the any Slaves. Remove the bond before editing Slave configurations.

4. Save new settings, exit, and restart network services:
 - a. Press the left arrow key to exit the main menu. When prompted, select **Yes** to save the configuration.
 - b. The setup utility will restart the CTA network services for the new configuration and return to the network configuration menu.

Note This configuration does not apply to CTA/VE.

Configuring the CTA with two subnets

In this example, the CTA is configured for two subnets with two physical ports (eth0, eth1):

1. Start the network configuration menu:
 - a. Type **rfhsetup** from the CTA command prompt to invoke the system setup menu.
 - b. Select **Configure Cloud Tiering Appliance Networking**. The network configuration menu appears.
 - c. Select **Configure Networking**.
2. Edit settings for the physical ports eth0 and eth1:
 - a. Use the up and down arrow keys to select **eth0** and press **Enter**. The configuration menu for the eth0 interface appears.
 - b. Provide information for each item to properly configure the interface.
 - Press **Enter** to edit an item, the press **Enter** again to complete.
 - Press the left arrow key to exit the menu.
 - Select **Yes** to keep new settings.
 - c. Repeat these steps for the eth1 interface.
3. Save new settings, exit, and restart network services:
 - a. Press the left arrow key to exit the main menu. When prompted, select **Yes** to save the configuration.
 - b. The setup utility will restart the CTA network services according to the new configuration and return to the network configuration setup menu.

Configuring the CTA with more than two subnets

In this example, the CTA is configured for more than two subnets with two physical interfaces. This configuration utilizes VLAN tagging and the switch connected to the CTA ethernet ports must be properly configured for tagging. In Cisco terminology, the switchport mode is set to **trunk**, and the required VLANs are allowed on the ports:

1. Start the network configuration menu:
 - a. Type **rfhsetup** from the CTA command prompt to invoke the system setup menu.
 - b. Select **Configure Cloud Tiering Appliance Networking**. The network configuration menu appears.
 - c. Select **Configure Networking**.
 2. Add new bond interface:
 - a. Type **A** to add an interface. Use the right arrow key to select **Bond**, and press **Enter**.
 - b. When prompted for the name of the new interface, press the up arrow key to generate a name. The name generated is **bond1**. Press **Enter** to complete.
 3. Edit the bond configuration:
 - a. Use the up and down arrow keys to select the new bond interface. Press **Enter**. The configuration menu for the interface appears.
 - b. For Slave, type **eth0 eth1**. Complete other values as needed.
 - c. Once the interface configuration is defined, press the left arrow key to exit the current menu. When prompted, select **Yes** to keep the new setting.
- Note** Configuration settings are saved, but are not implemented until the Cloud Tiering Appliance Network Setup menu is exited.
4. Add new VLAN interfaces:
 - a. Type **A** to add an interface. Use the right arrow key to select **VLAN**, and press **Enter**.
 - b. Type a name for the VLAN bond interface. The naming convention is *<interface>.<vlan-ID>*. For example, eth0.5 is a VLAN interface on eth0 with a VLAN ID of 5
 - c. Repeat these steps to create two more VLAN bond interfaces.
 5. Edit the VLAN configuration:
 - a. Use the up and down arrow keys to select the new VLAN interface. Press **Enter**. The configuration menu for the interface appears.
 - b. Provide information for each item to properly configure the interface:
 - Press **Enter** to edit an item, and then press **Enter** again to complete.
 - Press the left arrow key to exit the menu.
 - Select **Yes** to keep the new settings.

- c. Repeat these steps for each new VLAN interface.

Note To delete a bond that is part of a VLAN bond interface, delete the VLAN bond interface before deleting the bond.

6. Save the new settings, exit, and restart network services:
 - a. Press the left arrow key to exit the main menu. When prompted, select **Yes** to save the configuration.
 - b. The setup utility will restart the CTA network services for the new configuration and return to the network configuration menu.

VLAN tagging modes for the CTA/VE

The CTA/VE supports two VLAN tagging modes:

- [ESX Server virtual switch tagging](#)
- [ESX Server virtual guest tagging](#)

ESX Server virtual switch tagging

In the Virtual Switch Tagging (VST) mode, a VLAN ID is assigned to an ESX Server switch port. Untagged layer 2 traffic is sent by using the link between the switch port and the CTA/VE interface. When the switch receives this traffic, it directs it to the configured VLAN.

On the CTA/VE, configure each physical eth1, eth2, eth3 or eth4 port with an IP address, Net Mask, and Default Gateway.

Note When using the VST mode, do not create a VLAN interface.

Configuring the VLAN number on the ESX switchport in VST mode

Virtual switch tagging is enabled when the port group's VLAN ID is set to any number between 1 and 4094, inclusive.

To use VST, create appropriate port groups. Give each port group a label and a VLAN ID. Port group values must be unique on a virtual switch. Once the port group is created, you can use the port group label in the virtual machine configuration.

To configure port group properties:

1. Log in to the VMware VI Client and select the server from the inventory panel.
2. The hardware configuration page for this server appears.
3. On the **Configuration** tab, click **Networking**.
4. Click **Properties for a network**.
5. The vSwitch Properties dialog box appears.
6. On the **Ports** tab, select the port group and click **Edit**.
7. In the Properties dialog box for the port group, click the **General** tab to edit:
 - a. Network Label — This is the name of the port group that you are creating.

Name	IP Address	Network Mask	Up/Down	Comment
eth0			DOWN	Unconfigured
eth1			DOWN	Unconfigured
eth2			DOWN	Unconfigured
eth3			DOWN	Unconfigured

1 of 4 entries displayed

```
Command: [Q]uit [A]dd [R]emove [S]lave [U]p [D]own re[F]resh
[Help] Status: OK
      rfhsetup <- Network configuration -> Interface eth0's
configuration
```

4. Type **A** to add a new interface. Use the left and right arrow keys to select a VLAN interface and press **Enter**.
5. Type a name for the VLAN interface. The naming convention is *<bond>.<vlan-ID>*. For example, to add VLAN ID 20 on eth0, the name will be eth0.20. After typing the name, press **Enter**.

Note The new VLAN bond interface (for example, eth0.20) will be added to the interface list.
6. Use the up and down arrow keys to select the newly created VLAN interface. Press the right arrow key. The eth0.20 VLAN configuration screen appears. Add the IP address, netmask, and gateway.
7. Use the left arrow key to exit the eth0.20 configuration menu and save the configuration.
8. Use the left arrow key to exit the Configure Networking menu and apply the saved configuration.

Appendix B Alerts

This appendix presents the following topics:

Supported SNMP traps	132
CTA alerts	133

Supported SNMP traps

Table 10 lists the SNMP traps for which the CTA will send a notification.

Table 10. Supported SNMP traps

Notification name	MIB where it is defined	SNMP OID
eRAAlertDaemonRestarted	EMC-RAINFINITY-ALERTS-MIB	1.3.6.1.4.1.1139.9.3.2.0.1
eRAAlertsHistoryReset	EMC-RAINFINITY-ALERTS-MIB	1.3.6.1.4.1.1139.9.3.2.0.2
eRARainfinityAlert	EMC-RAINFINITY-ALERTS-MIB	1.3.6.1.4.1.1139.9.3.2.0.4
eRAGenericAlert	EMC-RAINFINITY-ALERTS-MIB	1.3.6.1.4.1.1139.9.3.2.0.5
eRASecurityAlert	EMC-RAINFINITY-ALERTS-MIB	1.3.6.1.4.1.1139.9.3.2.0.3
eRHSTemperatureAlert	EMC-RAINFINITY-HARDWARE-STATUS-MIB	1.3.6.1.4.1.1139.9.3.1.0.1
eRHSFanAlert	EMC-RAINFINITY-HARDWARE-STATUS-MIB	1.3.6.1.4.1.1139.9.3.1.0.2
eRHSPowerSupplyAlert	EMC-RAINFINITY-HARDWARE-STATUS-MIB	1.3.6.1.4.1.1139.9.3.1.0.3
eRHSMemoryAlert	EMC-RAINFINITY-HARDWARE-STATUS-MIB	1.3.6.1.4.1.1139.9.3.1.0.4
eRHSDiskAlert	EMC-RAINFINITY-HARDWARE-STATUS-MIB	1.3.6.1.4.1.1139.9.3.1.0.5
eRHSNICAlert	EMC-RAINFINITY-HARDWARE-STATUS-MIB	1.3.6.1.4.1.1139.9.3.1.0.6

CTA alerts

The CTA alerts are classified by type:

- Rainfinity alerts
- Generic alerts
- Security alerts
- Hardware alerts

Table 11 lists all the CTA alerts.

Table 11. CTA alerts

Index	Pattern name	Description	Type	SNMP OID
001-0001	(sshd telnetd).*session opened	User logged into this system via ssh or telnet.	securityAlert	1.3.6.1.4.1.1139.9.3.2.0.3
001-0002	(sshd telnetd).*session closed	An ssh or telnet user logged out.	securityAlert	1.3.6.1.4.1.1139.9.3.2.0.3
001-0003	Permission denied for illegal user	User is not authenticated and was denied access to the system.	securityAlert	1.3.6.1.4.1.1139.9.3.2.0.3
001-0005	failed to bind to LDAP server	Attempt to bind to the LDAP server failed. Possible causes include a misconfigured LDAP server address or a network connectivity issue. Delays in logging in or executing commands may result if the LDAP server is unavailable.	securityAlert	1.3.6.1.4.1.1139.9.3.2.0.3
001-0006	Log rotation	Log rotation settings have been modified via rfhsetup.	rainfinityAlert	1.3.6.1.4.1.1139.9.3.2.0.4
001-0007	scp of system log files	Attempt to securely copy the system log files has been made. The alert indicates whether the attempt succeeded or failed.	genericAlert	1.3.6.1.4.1.1139.9.3.2.0.5
001-0008	scp of CTA log files	Attempt to securely copy CTA log files has been made. The alert indicates whether the attempt succeeded or failed.	genericAlert	1.3.6.1.4.1.1139.9.3.2.0.5

Index	Pattern name	Description	Type	SNMP OID
001-0009	Invalid user	User is not known and was denied access to the system.	securityAlert	1.3.6.1.4.1.1139.9.3.2.0.3
001-0010	Accepted keyboard-interactive	User is not known and was denied access to the system.	securityAlert	1.3.6.1.4.1.1139.9.3.2.0.3
001-0011	Security level	System's security level has been modified.	securityAlert	1.3.6.1.4.1.1139.9.3.2.0.3
001-0013	Certificate Expire Warning	One certificate will expire soon or has already expired.	securityAlert	1.3.6.1.4.1.1139.9.3.2.0.3
001-0014	authentication failure	User's login attempt via ssh has failed due to an invalid password.	securityAlert	1.3.6.1.4.1.1139.9.3.2.0.3
001-0015	changed password expiry	User's password expiry information is changed. This typically happens if password aging is enabled, modified or disabled.	securityAlert	1.3.6.1.4.1.1139.9.3.2.0.3
001-0016	password changed	User's password has been changed in the local user database.	securityAlert	1.3.6.1.4.1.1139.9.3.2.0.3
001-0017	Log Alerts system enabled	rfaletd has been started.	securityAlert	1.3.6.1.4.1.1139.9.3.2.0.3
001-0018	Log Alerts system disabled	rfaletd has been terminated.	securityAlert	1.3.6.1.4.1.1139.9.3.2.0.3
001-3001	rfhsetup	User has launched the rfhsetup tool.	rainfinityAlert	1.3.6.1.4.1.1139.9.3.2.0.4
002-1001	Temperature Alert	Temperature status has changed or a temperature failure has occurred.	hardwareAlert	1.3.6.1.4.1.1139.9.3.1.0.1
002-1002	Fan Alert	Fan status has changed or a fan failure has occurred.	hardwareAlert	1.3.6.1.4.1.1139.9.3.1.0.2
002-1003	Power Supply Alert	Power supply status has changed or a power supply failure has occurred.	hardwareAlert	1.3.6.1.4.1.1139.9.3.1.0.3

Index	Pattern name	Description	Type	SNMP OID
002-1004	Memory Hardware Alert	Memory hardware status has changed or a memory hardware failure has occurred.	hardwareAlert	1.3.6.1.4.1.1139.9.3.1.0.4
002-1005	Disk Alert	Disk hardware status has changed or a disk failure has occurred.	hardwareAlert	1.3.6.1.4.1.1139.9.3.1.0.5
002-1006	NIC Alert	Network card status has changed or a network card failure has occurred.	hardwareAlert	1.3.6.1.4.1.1139.9.3.1.0.6
002-1007	capacity utilization	Partition capacity utilization exceeds a pre-configured threshold.	genericAlert	1.3.6.1.4.1.1139.9.3.2.0.5
002-3001	Problem starting filemanagement daemon	Problem starting filemanagement daemon	rainfinityAlert	1.3.6.1.4.1.1139.9.3.2.0.4
002-3002	filemanagement daemon stopped	filemanagement daemon has been stopped	rainfinityAlert	1.3.6.1.4.1.1139.9.3.2.0.4
002-3003	filemanagement daemon started	filemanagement daemon has been started	rainfinityAlert	1.3.6.1.4.1.1139.9.3.2.0.4
002-3004	The number of CCD connection in CLOSE_WAIT	The number of CCD connection in CLOSE_WAIT exceeds threshold	rainfinityAlert	1.3.6.1.4.1.1139.9.3.2.0.4
003-0001	Partition is full	A partition is full.	genericAlert	1.3.6.1.4.1.1139.9.3.2.0.5
301-0001	filemanagement daemon enabled	filemanagement daemon has been enabled.	rainfinityAlert	1.3.6.1.4.1.1139.9.3.2.0.4
301-0002	filemanagement daemon disabled	filemanagement daemon has been disabled.	rainfinityAlert	1.3.6.1.4.1.1139.9.3.2.0.4
301-0003	CTA-HA unable to contact Cloud Tiering Appliance	CTA has not been responsive for CTA-HA (as FCD) after a sufficiently long period.	rainfinityAlert	1.3.6.1.4.1.1139.9.3.2.0.4

Index	Pattern name	Description	Type	SNMP OID
301-0005	Automatic creation of a DHSM connection	rfwalker could not automatically create a DHSM connection. Manually create it with command fs_dhsm. See the task summaries for more details.	rainfinityAlert	1.3.6.1.4.1.1139.9.3.2.0.4
301-0006	Archived files count limit exceeded.	Archived files limit has been exceeded.	rainfinityAlert	1.3.6.1.4.1.1139.9.3.2.0.4
301-0007	Could not update capacity values	CTA was unable to update Capacity values for the specified filesystem. This will affect Capacity based archiving on this filesystem.	rainfinityAlert	1.3.6.1.4.1.1139.9.3.2.0.4
301-0008	Database vacuum is due to be performed.	DB Data folder size exceeds 2 times the number of archived files. Perform a database vacuum.	rainfinityAlert	1.3.6.1.4.1.1139.9.3.2.0.4
302-0001	CTA-HA unable to contact Cloud Tiering Appliance	CTA has not been responsive for CTA-HA (as CCD) after a sufficiently long period.	rainfinityAlert	1.3.6.1.4.1.1139.9.3.2.0.4
302-0002	(Centera) had this problem during connect	Centera has not responded to a connection request.	rainfinityAlert	1.3.6.1.4.1.1139.9.3.2.0.4
302-0003	The recall username/password does not seem to be in sync across CTAs and may cause recall problems	CCD has been monitoring multiple CTAs and the recall passwords on all the CTAs are not the same.	rainfinityAlert	1.3.6.1.4.1.1139.9.3.2.0.4
302-0004	has been rejected because a Centera with same name exists for a different CTA IP	Centera with same name has been added on multiple CTAs.	rainfinityAlert	1.3.6.1.4.1.1139.9.3.2.0.4

Index	Pattern name	Description	Type	SNMP OID
302-0005	The digest password or digest username is empty	Recall credentials are empty. If using CCD with multiple CTAs, ensure all the CTAs have credentials configured	rainfinityAlert	1.3.6.1.4.1.1139.9.3.2.0.4
303-0001	GUI user logged in successfully	GUI user has logged in successfully.	rainfinityAlert	1.3.6.1.4.1.1139.9.3.2.0.4
303-0002	GUI login attempt failed	GUI login attempt has failed.	rainfinityAlert	1.3.6.1.4.1.1139.9.3.2.0.4
303-0003	GUI user logged out	GUI user has logged out.	rainfinityAlert	1.3.6.1.4.1.1139.9.3.2.0.4
304-0001	exceeds threshold	NAS Repository usage has exceeded the configured threshold.	rainfinityAlert	1.3.6.1.4.1.1139.9.3.2.0.4
304-0002	Task failed too many operations and terminated early	Task has failed too many operations and has been terminated early.	rainfinityAlert	1.3.6.1.4.1.1139.9.3.2.0.4
304-0003	Failure setting retention for the Atmos object	Failure to set retention for the Atmos object. Verify that the Atmos policy specifies retention.	rainfinityAlert	1.3.6.1.4.1.1139.9.3.2.0.4
305-0001	StubScanner progress	Display number of CFA files converted by the StubScanner periodically. Only displayed after CFA Stubs are found.	rainfinityAlert	1.3.6.1.4.1.1139.9.3.2.0.4
305-0002	StubScanner Complete	Display final number of CFA files converted by the StubScanner.	rainfinityAlert	1.3.6.1.4.1.1139.9.3.2.0.4
305-0003	DXConversion progress	Display number of DX NAS stub files converted periodically. Only displayed while DXConversion is running.	rainfinityAlert	1.3.6.1.4.1.1139.9.3.2.0.4
305-0004	DXConversion Complete	Display final number of DX NAS stub files converted.	rainfinityAlert	1.3.6.1.4.1.1139.9.3.2.0.4

Index	Pattern name	Description	Type	SNMP OID
306-0001	Task completed	Task has completed. Additional details provide task information and statistics.	rainfinityAlert	1.3.6.1.4.1.1139.9.3.2.0.4
701-0001	Unable to open connection to Centera	Either a invalid PEA file has been specified or connectivity to Centera is down.	rainfinityAlert	1.3.6.1.4.1.1139.9.3.2.0.4
701-0002	Archive Warning threshold reached	Warning when Capacity Archive has reached the threshold.	rainfinityAlert	1.3.6.1.4.1.1139.9.3.2.0.4
701-0003	Archive by Capacity started	Notice when Capacity Archive has started.	rainfinityAlert	1.3.6.1.4.1.1139.9.3.2.0.4
701-0004	Tasks remaining in the pending queue will be ignored	Alert when pending queue is not empty when FMD has been shutdown.	rainfinityAlert	1.3.6.1.4.1.1139.9.3.2.0.4
701-0005	Failed to archive backup file	Alert when copying of backup file to backup destination has failed.	rainfinityAlert	1.3.6.1.4.1.1139.9.3.2.0.4
701-0006	Failed to restore backup file	Alert when restoring of backup file from backup destination has failed.	rainfinityAlert	1.3.6.1.4.1.1139.9.3.2.0.4
701-0007	No successful database backups done for the past	Alert when no backup task has run successfully for the past X number of days.	rainfinityAlert	1.3.6.1.4.1.1139.9.3.2.0.4
701-0008	Policy execution caused a repository to exceed its capacity limit	Alert when the execution of a policy has caused a repository to exceed its capacity limit.	rainfinityAlert	1.3.6.1.4.1.1139.9.3.2.0.4
701-0009	File archiving has reached over 80%	Alert when archived files count has reached 80% of 500 million max file limit.	rainfinityAlert	1.3.6.1.4.1.1139.9.3.2.0.4
701-0010	A fatal error has caused a CTA process to terminate	A fatal error has caused a CTA process to terminate. The log lists specific error messages.	rainfinityAlert	1.3.6.1.4.1.1139.9.3.2.0.4

Index	Pattern name	Description	Type	SNMP OID
701-0013	A vacuum task will start within the next 1 hour.	Alert when a vacuum task will start in the next 1 hour	rainfinityAlert	1.3.6.1.4.1.1139.9.3.2.0.4
701-0014	File Migration task reached its file threshold limit	Recursive File Migration task has reached its file threshold limit and will not auto run again. Correct the situation and try again.	rainfinityAlert	1.3.6.1.4.1.1139.9.3.2.0.4
801-0001	Failed to recall file from NetApp	A recall from NetApp failed	rainfinityAlert	1.3.6.1.4.1.1139.9.3.2.0.4
901-0001	Found a duplicate Atmos name and so this configuration will not be pulled down	Duplicate Atmos name rejected by ACD.	rainfinityAlert	1.3.6.1.4.1.1139.9.3.2.0.4
901-0002	The recall username/password does not seem to be in sync across CTAs and may cause recall problems	Password is not in sync across multiple CTAs configured for this ACD	rainfinityAlert	1.3.6.1.4.1.1139.9.3.2.0.4

All alerts are listed in the Log Pattern Index of the GUI.

Use the CLI or the Alert Settings page in the CTA GUI to edit log alert patterns. Configuration options include:

- Status — Alerts can be individually enabled.
- Throttle time — A different throttle time may be applied to each alert pattern. If alerts occur more than once within a specified throttle time, the repeated alerts are suppressed.
- Included in summary — Included alerts appear in the alert summary.

Note To enable e-mail alert messages from the device, alert settings must be configured.