CTA version 7.5 Online Help

Table Of Contents

Cloud Tiering Appliance

Cloud Tiering Appliance Help

The EMC Cloud Tiering Appliance (CTA) extends the benefits of location transparency from the tree level to the individual file level. All archived files appear as if they reside on primary storage. When an archived file is accessed, the requested data is automatically and transparently returned from archived storage to primary storage. CTA provides an option for backup to ignore the archived content.

In addition to archiving, you can use CTA to schedule many types of tasks on file servers such as file deletion, file backup, repository migration, and file migration.

The EMC Cloud Tiering Appliance/VE (CTA/VE) is the VMware virtual appliance edition with the same features as the full edition. Throughout the online help, CTA is used to describe both products unless explicitly indicated otherwise.

The following scenarios show users how to perform two common tasks.

Policy-based migration

File level migration transfers files between primary file servers.

- Configure CTA for file migration.
  - Set up the CTA on the network.
  - Configure each Celerra Data Mover.
- Create a file migration policy.
  - Create a migrate_file policy with one rule.
  - Create one catch-all rule to migrate files with size greater than 0 bytes.
- Create and run a file migration task
  - Create a task to migrate files from the Celerra source to the Celerra destination for the migrate file policy. All supported source and destination servers are listed in the supported_platform_matrix_for_file_migration.
  - Run an optional simulation task.
  - Run task.
  - View progress.
  - Perform manual steps to complete the file migration.

Multi-tier archiving

Multi-tier archiving is based on a multi-tier policy.

- Configure CTA for multi-tier archiving
  - Set up the CTA on the network.
  - Configure the Celerra and EMC Centera file servers.
- Create a multi-tier policy
  - Create a multi-tier policy for both normal and stub files based on two rules.
  - Create one rule to archive files older than three months to a NAS repository. Configure the NAS repository when creating the rule.
  - Create another rule to archive files older than six months from the NAS repository to the EMC Centera. Configure the NAS repository when creating the rule.
- Create a multi-tier archiving task
  - Create a task to archive from the Celerra source for the multi-tier policy.
  - Run an optional simulation task.
  - Run task.
  - View progress.
- View archived files

Documentation

All of the Cloud Tiering Appliance documentation is accessible from the table of contents in the left pane. The documentation set contains:

- Online Help. The Help link on any GUI page provides more information on specific CTA settings.
- Getting Started Guide. This link accesses the *EMC Cloud Tiering Appliance and Cloud Tiering Appliance/VE Getting Started Guide*.
- Release notes. The most current version of the *EMC Cloud Tiering Appliance and Cloud Tiering Appliance/VE Release Notes* is available from EMC Powerlink.
- *EMC Cloud Tiering Appliance man pages*

New Features

The Cloud Tiering Appliance replaces the File Management Appliance. The product names are:

- EMC Cloud Tiering Appliance
- EMC Cloud Tiering Appliance/VE

Cloud Tiering Appliance and Cloud Tiering Appliance/VE 7.5 include the following new features:

- Expanded platform support for archiving.

  Multi-tier policy archiving adds platform support for VNX as a primary tier and for VNX, VNXe, and Isilon as secondary and tertiary tier targets.

- Expanded platform support for archiving.

  NAS repository migration adds platform support for VNX, VNXe, and Isilon as source and destination.

- Incremental file migration.

  The file migration task adds snapshot support on source servers. An initial migration is run, and then incremental runs only migrate files changed since the time of the last migration. This reduces server downtime and allows more client write access during migration.

- Expanded platform support for policy-based migration

  File migration adds platform support for VNX and NetApp as source servers and VNX and VNXe as destination servers.

Configuring the Cloud Tiering Appliance

CTA Deployment Overview

The EMC Cloud Tiering Appliance (CTA) plugs directly into your network. For those using the Cloud Tiering Appliance/VE (CTA/VE), it installs on a VMware virtual appliance.

The *Getting Started Guide* provides instructions to configure the CTA or the CTA/VE for your environment. The tasks to perform will vary depending upon your network configuration. Configuration tasks can be performed:

- From the CLI by using the root account.
- From the GUI by users logged in with an administrator account.

Once the appliance is successfully deployed, certain configurations may be changed by using the GUI.

- Edit, Delete, or Add a File Server
- Edit, Delete, or Add a NAS repository

Configure CTA for file migration

The policy-based migration example described in the overview is based on Celerra to Celerra file migration.

To configure Celerras for file migration, refer to the Getting Started Guide. Instructions are provided to:

- Review the prerequisites for the servers.
  - For NFS, CTA has root access with read/write permissions on source and destination exports.
  - CIFS user has administrator access on all shares.
  - If migrating files from a source that is shared and exported, CTA has root access with read/write permission to both shares and exports.
  - Snapsure snapshots must be enabled on the source.
  - There is a system user for the XML API setup on the source.
  - For every Celerra server, properties are configured on the CTA for:
    - Control Station IP address
    - NDMP username and password
- Configure the CTA network.
- Add the source and destination Celerras to the CTA configuration.

To ensure that local users and groups on source servers will have the same security access to migrated files on destination servers, CTA supports SID translation. To use this feature, upload a SID translation file on the File Migration Settings page. During migration, CTA uses the mappings in the file to translate CIFS user SIDs on the source to equivalent SIDs on the destination so that ACLs with SIDs specific to the source will remain consistent after migration. The Getting Started Guide provides information on how to generate a SID translation file.

Once the CTA has been configured, proceed to create_a_file migration policy.

Configure CTA for multi-tier archiving

The multi-tier archiving example described in the overview involves configuring a Celerra that is running DART 6.0 as a source to archive to a NAS repository on the first tier. From the NAS repository, older files are archived to an EMC Centera as a second tier. After the files are moved to the second tier, the stubs remain on the Celerra source, but point to files on the EMC Centera.

To configure the Celerra, refer to the Getting Started Guide. Instructions are provided to:

- Review the prerequisites for using Celerra as an archiving source.
- Add the Celerra to the CTA configuration.
- Configure the CTA for Celerra to EMC Centera archiving so that it includes local hostname resolution.
- Perform the prerequisite tasks on the Celerra Control Station before archiving.
- Configure the automatically created DHSM connections.

To configure the NAS repository and the EMC Centera, refer to the procedures here:

- Edit, Delete, or Add a Server describes how to configure CTA to add a file server. Refer to EMC Centera Properties to configure a new EMC Centera.
- Creating a policy includes steps to create the NAS repository as part of rule creation for the policy.

Once the CTA has been configured and the EMC Centera has been added, proceed to create_a_multi-tier_policy.

Configure an archive file list provider

The CTA can import a list of files to archive from an external third-party software provider. The third-party software administrator generates an XML file that is copied to the CTA. Before the import occurs, the CTA must be configured with:

- Celerra properties, VNX properties, or NetApp properties — The name and properties of the source servers where the files to archive are stored. The third-party software administrator gives the server configuration details to the CTA administrator. When the servers are added, CTA creates the NetBIOS server names.
- Provider properties — A provider name that is allowed to log in to the CTA. When the provider is configured, CTA creates a staging directory to be used for the imported file list.
- Import Files task — The name of the task that operates on the imported list of files to archive.

The CTA administrator gives the server name, provider name, and task name configured on the CTA to the third-party software administrator to use in generating the XML file.

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

For more details on importing a file list archive, refer to the Getting Started Guide

Edit, Delete, or Add a Server

Before archiving to any server, the appliance must be configured with that server. The Getting Started Guide provides instructions to perform the initial configuration. Once the initial configuration is complete, use the GUI to edit, delete, or add a server.

To edit, delete, or add a server

From the **Configuration** tab, select **File Servers**. The **File Server List** page appears.

- To edit a server, select the name of the server and click **Edit**. The properties page for the selected server type will appear.
- To delete a server, select the name of the server and click **Delete**.
- To add a new server:

1. Click **New** and on the **File Server Properties** page that appears.
2. Select the server type from the **Type** list box.

When editing or adding a server, the properties page for the selected server type will appear.

More

Atmos Properties
Celerra Properties
Centera Properties
Data Domain Properties
Isilon properties
NetApp Properties
Windows Server Properties
VNX properties
VNXe properties

Edit, Delete, or Add a NAS repository

Before archiving to a Celerra, VNX, VNXe, Windows, Isilon, NetApp, or Data Domain, a NAS repository must be configured on the appliance. The NAS repository may be configured before the policy is defined as outlined in the Getting Started Guide, or it can be added as part of the policy creation process.

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

To edit, delete, or add a NAS repository before defining a policy

On the **Configuration** tab, select **NAS Repository and NAS group**. The **NAS Repository List and NAS Group List** page appears.

- To edit a NAS repository, select the name of the repository and click **Edit**.
- To delete a NAS repository, select the name of the repository and click **Delete**.
- To add a new NAS repository, click **New**. The **Create New NAS Repository** page.

How do I?

Creating a NAS repository as an archive destination when creating a policy.

Page help
Alert pattern names

Alerts are reported on the **Log Alert Patterns** page.

Index	Pattern name	Description
001-0001	session opened	User logged into the system through SSH or Telnet.
001-0002	session closed	An SSH or Telnet user logged out.
001-0003	authentication failure	User was not authenticated and was denied access to the system.
001-0004	Telnet	Attempt to login by using Telnet. Alert is only generated if the Telnet service is running. This service is disabled by default.
001-0005	failed to bind to LDAP server	Attempt to bind to the LDAP server failed. This failure could be due to a misconfigured LDAP server address, or a network connectivity issue. The user could experience delays in logging in or executing commands if the LDAP server is unavailable.
001-0006	log rotation	Log rotation settings have been modified by using rfhsetup.
001-0007	scp of system log files	An attempt to securely copy the system log files was made. The alert indicates whether the attempt succeeded or failed.
001-0008	scp of Rainfinity log files	An attempt to securely copy the Rainfinity log files was made. Alert will indicates whether the attempt succeeded or failed.
001-0010	accepted password	Login through SSH was accepted with a valid password.
001-0011	security level	System security level was modified.
001-0013	certificate expiration warning	A certificate will expire soon or has already expired.
001-0014	failed password	Login attempt through SSH failed due to an invalid password.
001-0015	changed password expiration	Password expiration information has been changed. This commonly occurs if password aging is enabled, modified, or disabled.
001-0016	password changed	Password has been changed in the local user database.
001-0017	log alerts system enabled	The rfalertd has been started.
001-0018	log alerts system disabled	The rfalertd has been stopped.
001-3001	rfhsetup	The rfhsetup tool started.
002-1001	temperature alert	Alert is sent when reading by a temperature sensor exceeds or drops below a safe threshold.
002-1002	fan alert	Alert is sent when a fan status has changed or a fan failure occurs.
002-1003	power supply alert	Alert is sent when a power supply status has changed or a power supply failure occurs.
002-	memory alert	Alert is sent when a memory hardware status has

1004		changed or a memory hardware failure occurs. NOTE: If a memory hardware failure occurs, the system might shutdown before generating the alert.
002-1005	disk alert	Alert is sent when a disk status has changed, or when a disk failure occurs. This alert is related to the mechanical operation of the hard disk.
002-1006	NIC alert	Alert is sent when a network card status has changed, or when a network card failure (or port failure within that network card) occurs.
002-1007	capacity utilization alert	Disk capacity utilization exceeds the preconfigured threshold of 85%.
002-1008	timezone alert	Time zone has been changed
002-3001	problem starting File Management	CTA daemon is not present
002-3002	File Management stopped	CTA daemon has stopped.
002-3003	File Management started	CTA daemon has started.
002-3004	number of CCD connections in CLOSE_WAIT	When a Celerra recalls files stored on an EMC Centera, the request passes through the CCD on the CTA which creates a TCP connection in CLOSE_WAIT for a few milliseconds. Multiple connections in the CLOSE_WAIT state consume resources on the CTA. If this number continually grows or does not decrease when there is no recall activity, reboot the CTA.
003-0001	partition full	Disk partition is full. This alert is sent when any partition on the system exceeds 99% utilization.
301-0001	FileManagement enabled	CTA daemon has been enabled.
301-0002	FileManagement disabled	CTA daemon has been disabled.
301-0003	FMHA appliance unable to contact File Management Appliance	CTA has not respond to the CTA-HA (as FCD) after a sufficiently long period.
301-0004	no matching DHSM connection	This alert is sent when a Celerra archiving task fails because there is no DHSM connection configured on the Celerra for the specified secondary server. A DHSM connection to the secondary must be created on the Celerra using the fs_dhsm command.
301-0005	automatic creation of DHSM connection failed	The rfwalker cannot automatically create a DHSM connection. This connection might need to be manually created with the command fs_dhsm.
301-0006	ARCHIVE FILES COUNT LIMIT EXCEEDED	The maximum archive file limit is exceeded. The CTA and CTA/VE appliance have a maximum number of supported files that can be archived. Once this limit is reached, jobs to archive additional files will not be allowed. Existing jobs will be allowed to complete.
301-0007	could not update capacity values	This alert is sent when CTA is unable to obtain disk capacity values for primary servers. Restart the CTA daemon. If the alert persists, contact EMC technical support.
302-0001	FMHA alert (CCD)	Cloud Tiering Appliance High Availability (CTA-HA) is unable to contact CTA with Celerra as primary storage.
302-0002	Centera has this problem during connect	The EMC Centera is not responding to a connection request.

303-0001	GUI user logged in successfully	
303-0002	GUI login attempt failed	
303-0003	GUI user logged out	
304-0001	Exceeds threshold	NAS repository exceeds the configured threshold.
305-0001	Stub scanner progress	The CTA displays the number of CFA files converted by the StubScanner.
305-0002	Stub scanner complete	The CTA displays the final number of CFA files converted by a StubScanner run.
305-0003	DX Conversion progress	Display number of DX NAS stub files converted periodically. Only displayed while DX Conversion is running.
305-0004	DX Conversion Complete	Display final number of DX NAS stub file converted.
701-0001	Centera alert	The CTA is unable to open connection to the EMC Centera.
701-0002	Archive warning threshold reached	Warning when the a source reaches the warning capacity threshold.
701-0003	Archive by Capacity started	Notification that archive by capacity has started.
701-0004	Tasks remaining in the pending queue will be ignored	The CTA daemon has been shut down. Pending tasks in the queue will be ignored.
801-0001	Recall failure alert	A recall attempt from archived storage has failed.
901-0001	Found a duplicate Atmos name and so this configuration will not be pulled down.	Duplicate Atmos name rejected by ACD.
901-0002	The recall username and password does not seem to be in sync across FMAs and may cause recall problems.	Password is not in sync across multiple CTAs configured for this ACD.

More:

SNMP Traps

Alert settings

Use this page to configure alerts. The types of events that can trigger an alert notification are listed in the Alert Pattern Names table. The Rainfinity setup tool can also be used to configure alerts.

To access this page

Select the **Configuration** tab, select **Alert Settings**.

Field Descriptions

Field	Setting	Description
Alerts	Enable all alerts	Select to send both SNMP and e-mail alerts when an event occurs.
	Enable SNMP alerts	Select to send an SNMP alert when an event occurs.
	Enable e-mail alerts	Select to send an e-mail alert when an event occurs.
	Disable all alerts	Select to suppress all alert notification.
	New e-mail address	Enter an -mail address. Click **Add**. Alerts will be sent to all e-mail addresses listed. To remove an e-mail address, highlight the address and click **Delete**.
Alert Summary	Enable log alert summary	Select to collect alerts into a single alert summary.
	Summary time in hour	Frequency with which the alert summary is sent.
	Summary size in kilobytes	When the size of the alert summary file is reached, the alert summary is sent.
Edit log pattern		Click this link to Edit log alert Pattern.

To change or reset any of the settings, make a selection and click **Commit.**

Atmos Callback status

Use this page to view the status of every Atmos Callback agent that is listed on this page.

To access this page

From the **Configuration** tab, select **Celerra and Atmos Callback Status.**

Field Descriptions

Field	Description
Callback Agent IP Address	The IP address of the Atmos Callback agent. The IP address that is provided when the Atmos Callback Service is initially configured on the CTA by using the command: **/opt/rainfinity/filemanagement/bin/acdsetup.sh init_rffm** The address is listed in the configuration file: /opt/rainfinity/filemanagement/conf/rain.xml To delete a callback agent, check the box next to the IP address and click **Delete** In addition to deleting the callback agent in the GUI, you must also remove the CTA configuration from the Atmos Callback Daemon (ACD). To remove the IP address of the CTA: `acdsetup.sh rm_rffm <ip_address>` If you do not remove the CTA configuration and you restart the ACD or simply leave it running, it will re-register itself to the CTA and appear in the CTA GUI.
Status	Status of the connection to the port.
Port	Default port number is 9000. If no number is reported, there is no connection.
Last Recall Time	Date and time that a recall was last performed.
Total Recalls In Progress	Total number of recalls currently in progress.
Last Status Update	Date and time that the status was last updated.
Supported FM apps	List of CTAs that the callback agent is configured to monitor.
Warnings	Only one Atmos can be configured as a source on multiple CTAs. If a second CTA appears.

To delete a callback agent, select the checkbox next to the name of the agent and click **Delete**.

Atmos properties

Use this page to set up a new or existing Atmos file server as an archive destination. Atmos is a cloud-optimized storage (COS) product for EMC.

To access this page

1. From the **Configuration** tab, select **File Servers**. The **File Server List** page appears.
2. Do one of the following:

- Select **Atmos** as the type of file server to add.
- Select an Atmos to edit.

Field Descriptions

Field		Description
Basic File Server Information	Name	Logical server name.
Web Service Specific Settings	DNS Name	DNS name used to resolve IP addresses in the Atmos cluster.
	Port	• Select **HTTPS** or **HTTP** to specify the communication protocol for Atmos. • Type the port number through which HTTPS or HTTP connects.
	Username	The UID of a Subtenant on the Atmos. CTA uses this UID to access storage on the cluster. If there is a subtenant, specify the username as: <*Subtenant_ID*>/<*UID*> where *Subtenant_ID* is an alphanumeric string generated by the Atmos. NOTE: The UID is not the Tenant Name, the Subtenant Name, or the Subtenant ID.
	Password	The Shared Secret associated with the UID on the Subtenant Information page of the Atmos. When configuring an Atmos Subtenant user, the format is: <shared_secret>/<UID>

After adding or updating any properties:

- To accept the changes, click **Commit**.
- To cancel the process, click **Cancel**.

How Do I?

Edit or add a file server.

Archived report settings

Use this page to configure the data reported on the **Archived Report** page.

To access this page

- From the **Archived Files** tab, select **Archived Report Settings**.
- From the **Configuration** tab, select **Archived Report Settings**.

Field Descriptions

Field		Description
File Size		Specify the file size groups to present in the Archived Report Pie charts.
	Unit	The file size unit in kilobytes, megabytes, gigabytes, terabytes, or petabytes.
	Group A Group B Group C Group D	Specify the maximum file size for each group.
File Type		Specify the file type extensions to chart. Checked extensions are included in the Archived Report Pie charts.
Archive Time Range	View archive report for last	Type the number of months to chart on the Archive history. Values must be between 1 and 12 months.
	View archive report by date range	Type the start date and end date to chart on the Archive history. Dates must be in the format: yyyy-mm-dd.
Archive Report Cleanup	Cleanup archive report older than	Type the number of days after which the Archive report should be regenerated. Values must be between 1 and 730 days.
File Server Groups		Create or edit file server groups for CTA. To create a new file server group: 1. Type a new Group Name. 2. Select file servers from the **Available File Servers** list, and click the right arrow to add the file server to the **Member File Servers** list. To edit an existing file server group: 1. Select a name from **Configured Groups** and click **Display Members**. The name appears in the **Group Name** field. 2. Use the right and left arrows to change the member file servers in the group.
	Group Name	Name of the file server group.
	Available File Servers	All the available file servers. Use the right and left arrows to move the file servers between available file servers and member file servers.

	Member File Servers	File servers that are members of the file server group.
	Configured Groups	• To display the members of each file server group in the **Member File Servers** list, click **Display Members**. • To highlight the file server group to delete, click **Delete Selected**.

After changing the Archive Report settings:

- To accept the changes, click **Commit**.
- To cancel the process, click **Cancel**.

Backup and recovery settings

Use this page to:

- Restore backup files to the CTA.
- Configure backups for the CTA configuration and database.

The backup task uses the settings from this page.

NOTE: If no recent backup is listed under Available backup files, then the backup probably failed. Errors that occur during a backup run are logged in the system.log file. Check this file for error messages such as "out of disk space on destination".

To restore a backup, use the CLI script **fmrestore** described in the Getting Started Guide.

To access this page

From the **Configuration** tab, select **Backup and Recovery Settings.**

Field Descriptions

Field		Description
Recover File Management	Available backup files	Select an available backup file (.tgz) from the list and click **Restore**. Up to a maximum of 15 backup file names are stored along with the EMC Centera or NAS ID. Any subsequent backups will overwrite older backups. NOTE: Before performing a restore, ensure that no active tasks, such as archiving or stub scanning, are running.
File Management Backup Destination	Number of Backups	Specify the number of backups. The default value is 5.
	Select Destination	Select the **Centera** or **NAS Repository** where the backup files (.tgz) will be stored. Only NFS NAS repositories are listed as backup destinations. CIFS NAS repositories are not supported. NOTE: The backup location pathname must not exceed 170 characters.
	Select Disaster Recovery Location	Select the NFS export where the backup catalog file (DBBackup.out) will be stored. • Select the file server. • Type the path to the repository on the file server. Click **Browse** to find the path. NOTE: The backup location path name must not exceed 170 characters. To save the configuration settings, click **Save Backup Location**.

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

Celerra Callback status

Use this page to view the status of every Celerra callback agent that is listed on this page.

To access this page

From the **Configuration** tab, select **Celerra and Atmos Callback Status.**

To delete a callback agent, select the checkbox next to the name of the agent and click **Delete.**

Field Descriptions

Field	Description
Callback Agent IP Address	This is the IP address that is provided when the Celerra Callback Service is initially configured on the CTA by using the command: **/opt/rainfinity/filemanagement/bin/ccdsetup.sh init_rffm** It is listed in the configuration file: /opt/rainfinity/filemanagement/conf/rain.xml To delete a callback agent, check the box next to the IP address and click **Delete** In addition to deleting the callback agent in the GUI, you must also remove the CTA configuration from the Celerra Callback Daemon (CCD). To remove the IP address of the CTA: `ccdsetup.sh rm_rffm <ip_address>` If you do not remove the CTA configuration and you restart the CCD or simply leave it running, it will re-register itself to the CTA and appear in the CTA GUI.
Status	Shows the status of the connection to the port.
Port	Default port number is 8000. If no number is reported, there is no connection.
Last Recall Time	The date and time that a recall was last performed.
Total Recalls In Progress	Shows the total number of recalls currently in progress.
Last Status Update	The date and time that the status was last updated.
Supported FM apps	A list of CTAs that the Celerra Callback Agent is configured to monitor.
Warnings	Only one Celerra can be configured as a source on a single CTA. If a second CTA is detected with this Celerra designated as an archive source, a warning appears. Celerra as Source provides more information on configuring this Celerra property.

Celerra properties

Use this page to set up a new or existing Celerra Data Mover.

To access this page

1. From the **Configuration** tab, select **File Servers**. The **File Server List** page appears.
2. Do one of the following:

- Click **New** to add a file server. Select **Celerra** as the type of file server to add.
- Select a Celerra to edit.

Field Descriptions

Field		Description
FileMover Settings		Click the link to configure the credentials for archive and recall.
Basic File Server Information	Name	The Celerra server NetBIOS name or hostname. NOTE: This is not the Fully Qualified Domain Name (FQDN) or IP address.
	File server is VDM	Identifies the Celerra as a Virtual Data Mover. Virtual Data Movers support only the CIFS protocol.
IP Addresses	New IP address	Type the Celerra Data Mover IP address. • When editing an existing server, click **Update** to retrieve the IP address from the DNS based on the server name. • To specify an additional IP address, click **Add**. The IP address is added to the list. • To delete an existing IP address, select an address and click **Delete**.
Control Station	IP address	Type the IP address of the Celerra Control Station. • For file migration, this is required for all Celerra servers involved in CIFS transactions. • For archiving, this allows the CTA to automatically perform some preconfiguration steps for archiving. If this field is empty, the CTA takes no action and the preconfiguration steps must be performed manually. NOTE: A system user for the XML API and FileMover API operations must be configured with a valid password on the Celerra Control Station.
CIFS Specific Settings	Username	Type the username of the Microsoft Windows domain user to be used by the CTA. The Windows domain user should be part of the file server's local administrator's group. The CIFS credential is not required if only the Celerra performs NFS archiving.

	Password	Type the Windows domain user password.
	NetBIOS Domain	Type the Windows NetBIOS domain name.
NDMP Specific Settings	Username	Type the username of the NDMP user on the source and destination Celerra Data Movers. File migration uses NDMP settings.
	Password	Type the password for the NDMP user.
	Port	Type the port used for NDMP operations. The Celerra and the CTA use this port for NDMP traffic. Default is 10000.
Celerra as Source	Enable Celerra as source	Configures the CTA to archive data from the Celerra Data Mover. If more than one CTA is connected to the same Celerra Data Mover, configure only one CTA with this option. If more than one CTA is configured to archive data from a single Celerra Data Mover, data loss might occur.
Celerra Callback Agent Settings	CCD DNS Name	Type the FQDN of the Celerra Callback DNS entry. NOTE: The FQDN is case-sensitive. The username and password for the FileMover API and Celerra HTTP authentication credentials are provided on the FileMover Settings page.
Atmos Callback Agent Settings	ACD DNS Name	Type the FQDN of the Atmos Callback DNS entry. NOTE: The FQDN is case-sensitive.
Directory Exclusion List	Exclude Directory	Specify the names of any directories to exclude. These directories are skipped for archiving and stub scanning tasks. Directory exclusion does not apply to migration tasks. The Directory exclusion list help page provides details on how to specify excluded directory names. Verify that stub files are not in the excluded directories. CTA will not access the excluded directories and the stub files will become orphans. • To specify an additional directory to exclude, type the directory name and click **Add**. • To delete an existing directory from the list, select the directory and click **Delete**.

After adding or updating any properties:

- To accept the changes, click **Commit**.
- To cancel the process, click **Cancel**.

How Do I?

Edit or add a file server.

EMC Centera properties

Use this page to set up a new or existing EMC Centera server.

To access this page

1. From the **Configuration** tab, select **File Servers**. The **File Server List** page appears.
2. Do one of the following:

- Select **Centera** as the type of file server to add.
- Select an EMC Centera to edit.

NOTE: When editing an EMC Centera, only the Number of Access Nodes can be updated.

Editing an existing EMC Centera configuration risks data unavailability for stubs that were archived with the original configuration. To change EMC Centera configuration settings other than Number of Access Nodes, delete the existing server and create an EMC Centera server.

Field Descriptions

Field	Description
Name	The logical name used to identify the EMC Centera. Do not change this name after archiving or stub files will not be recallable. NOTE: If multiple FM appliances are deployed in a Celerra environment with multiple EMC Centera or Atmos systems, each EMC Centera or Atmos must have a unique logical name. Otherwise the Celerra or Atmos Callback daemon will not be able to reliably resolve the destination name.
Access Node IP NOTE: This setting does not appear when editing an existing EMC Centera.	Specify the IP address of the EMC Centera access node. • To specify an additional access node IP, click **Add**. The IP address is added to the list and will be added as an entry in the Access Node String field. • To delete and existing node, select a node IP and click **Delete**.
Access Node String	This value is automatically generated when the Access Node IP address is added or deleted. You cannot enter data by typing directly into the field.
Number of Access Nodes	Type the number of access nodes. The default value is 4. Reduce this number if you have fewer than four access nodes.
Authentication	Select from one of three choices: • Anonymous — If selected, no security is used to authenticate with the EMC Centera. • User profile — If selected, the username and password created on the EMC Centera are required. • PEA file — This option requires that a Profile and pool entry authorization (PEA) file was created to access the EMC 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 the EMC Centera. Type the path to the file on the local machine or browse for the file. A copy of the file is stored with the CTA configuration.

After adding or updating any properties:

- To accept the changes, click **Commit**.
- To cancel the process, click **Cancel**.

How Do I?

Edit or add a file server.

Change password

Use this page to change a user password.

To access this page

From the **Configuration** tab, select **Change Password**.

To change a password:

1. Type the old password.
2. Type the new password.
3. Retype the new password to confirm.
4. Click **Commit**.

Command history

Use this page to query commands. Only admin users are authorized to use this page.

To access this page

From the **Configuration** tab, select **Command History**.

To generate a list of commands for a certain criteria, type the criteria and click **Query**.

Field Descriptions

Field	Description
History type	Select **daemon**, **system**, or **session**.
Start time	The start of the period to query. This value is only active for system history types. • In the first textbox, type the date in YYYY-MM-DD format. • In the second textbox, type the time in HH:MM:SS format.
End time	The end of the period to query. This value is only active for system history types. • In the first textbox, type the date in YYYY-MM-DD format. • In the second textbox, type the time in HH:MM:SS format.
Command name	Type the command name. If no command name is specified, the command history for all commands is listed.
User name	Type the name of the user who issued the command. If no user is specified, the command history for all users is listed.

Create NAS group

Use this page to configure a new NAS group. A NAS group is a collection of NAS repositories that are used for archiving.

An archiving task will archive to the first repository in the NAS group until it reaches the specified maximum disk usage limit. It will continue to iterate through all the repositories in the group until the task is complete. If all repositories reach the maximum disk usage limit, the task will send an alert and exit.

To access this page

1. From the **Configuration** tab, select **NAS Repository and NAS Group**. The **NAS Repository List and NAS Group List** page appears.
2. Do one of the following:
   3. Click **New** to add a new NAS group.
   4. Select a NAS group to edit.

Field Descriptions

Field		Description
Name		Type a new NAS group name. A group name can be up to 20 characters.
Archive Destination	Protocol	Select **CIFS** or **NFS**. The list of available repositories may change.
	Repositories	To add a selection to the list, select a repository and click **Add to List**. NOTE: A NAS group may be comprised of a collection of CIFS or NFS repositories. However a mixture of CIFS and NFS repositories in a single NAS group is not allowed. To delete one of the repositories from the list, click **Delete Selected**.

After adding or updating any groups:

- To accept the changes, click **Save Group** to accept the changes.
- To cancel the process, click **Cancel** to cancel the process.

Create or edit a NAS repository

Use this page to add a new or edit an existing NAS repository. A NAS repository can be a Celerra or VNX Data Mover, VNXe, NetApp Filer, Data Domain, Windows server or Isilon.

To access this page

1. From the **Configuration** tab, select **NAS Repository and NAS group**. The **NAS Repository and NAS Group List** page appears.
2. Do one of the following:

   - Click **New** to add a new to add a new NAS repository.
   - Select a NAS repository to edit.

Field Descriptions

Field	Description
Name	The repository name is automatically generated when the repository definition is saved. You cannot enter data by typing directly into the field.
File Server	Select from the list of available file servers. The available file servers are the configured servers listed on the File Server List. The file server must have a proper DNS entry defined that links the file server name with the IP address.
Protocol	Select **CIFS** or **NFS**. The source and repository protocol types must match. • If the source protocol is CIFS, the NAS repository protocol must be CIFS. • If source protocol is NFS, the NAS repository protocol must be NFS.
Path	Type the path to the repository on the file server. Click **Browse** to find the path.
Maximum limit of disk usage	When the capacity on the repository reaches the limit specified, CTA stops archiving to this destination. Default value is 90 percent. NOTE: When editing an existing NAS repository, this is the only value that can be changed.

After adding or updating any repositories:

- To accept the changes, click **Save Repository**.
- To cancel the process, click **Cancel**.

How Do I?

Edit or add a NAS Repository.
Create a NAS repository as an archive destination when creating a policy.

Data Domain properties

Use this page to set up a new or existing Data Domain file server as an archive destination from a Celerra source.

To access this page

1. From the Configuration tab, select **File Servers**. The File Server List page appears.
2. Do one of the following:

- Select **Data Domain** as the type of file server to add.
- Select a Data Domain file server to edit.

Field Descriptions

Field		Description
Basic File Server Information	Name	The logical server name.
IP Addresses	New IP address	Type the IP address of the Data Domain server. • To specify an additional IP address, click **Add**. The IP address is added to the list. • To delete an existing IP address, choose an address and click **Delete**. **Update** is not an available option. Since Data Domain is as NFS only server, it cannot retrieve the IP address from the DNS that is based on the server name.

After adding or updating any properties:

- To accept the changes, click **Commit**.
- To cancel the process, click **Cancel**.

How Do I?

Edit or add a file server.

Directory exclusion list

This option is used to exclude directories for scanning, such as with archiving or stub scanning tasks. Migration tasks do not scan, so this option does not apply for file migration.

Directory names can be specified in two different ways:

- To exclude any directory name from the start directory of the task, specify the name without slashes. For example, if etc is to be excluded, type **etc**.
- To exclude the full directory path from a share, specify the name with slashes. For example, if etc in the home directory is to be excluded, type: **\home\etc**.

NOTE: Enter the leading forward slash or backward slash depending on the protocol.

- For CIFS, use \
- For NFS use, /

To exclude directories for migration tasks, create a file migration policy. Add a rule with a file matching expression that defines the directories to be skipped.

Edit NAS group

Use this page to edit an existing group.

To access this page

Do one of the following:

- From the **Policies** tab, click the down arrow next to a NAS group and select **Edit**.
- From the **Configuration** tab

1. Select **NAS Repository and NAS Group**. The **NAS Repository List and NAS Group List** page appears.
2. Select the name of the group and click **Edit**.

Field Descriptions

Field		Description
Destination Name		The existing destination name appears. A destination name can have a maximum of 20 characters.
Archive Destination	Protocol	The protocol is either **CIFS** or **NFS**. The list of available repositories might change.
	Repositories	To add a selection to the list, select a repository and click **Add to List**. NOTE: A NAS group may comprise of a collection of CIFS or NFS repositories. However, a mixture of CIFS and NFS repositories in a single NAS group is not allowed. To delete one of the repositories from the list, click **Delete Selected**.

After editing a NAS group:

- To save the group. click **Save Group.**
- To cancel editing the group, click **Cancel.**

CTA configuration

The Cloud Tiering Appliance configuration applies to CTA or CTA/VE. Two sets of configuration options are available.

Configuration

Use the CTA configuration options to perform the following tasks:

Task	Reference
Change a user password.	Change Password
Edit, delete, or add a user.	Users
Review log file info, and configure log file rotation and backup settings.	Log Settings
Enable alerts and configure alert summary.	Alert Settings
View all basic, password aging, LDAP authentication, and Radius authentication settings that were set using the Setup Tool.	System Security Settings
Run a query on commands run from the CLI or GUI.	Command History
Add an SNMP community string or notification host.	SNMP Configuration

File Server Configuration

Use the file server configuration options to perform the following tasks:

Task	Reference
Add, edit, or delete file servers.	File Servers
Add, edit, or delete NAS repositories.	NAS Repository and NAS group
Edit the callback agent settings for a Celerra or VNX.	FileMover Settings
Add, edit, or delete a file server that is allowed to provide a list of files to archive.	Import Provider List
Add or delete a SID translation file for file migration.	File Migration Settings
For NetApp Fpolicy Special Clients: • Delete existing Fpolicy Callback Agents. • Add Excluded Clients. • Add Blocked Clients.	NetApp Fpolicy Special Clients
Check the Celerra or Atmos Callback Status.	Celerra and Atmos Callback Status
Configure the results to display on the Archived Report.	Archived Report Settings
Configure the settings for backup and recovery.	Backup and Recovery Settings

File Migration Settings

Use this page to add the optional SID translation file used by the file migration task. A SID translation file maps local users or groups on source primary servers to users or groups on destination primary servers with the same file access after migration.

A SID translator is shipped with the CTA on the product CD. Instructions to install and use the SID translator to generate a SID translation file are provided in the Getting Started Guide.

To access this page

From the **Configuration** tab, select **File Migration Settings.**

Field Descriptions

Field	Description
SID Translation File	To add a SID translation file: • Type the path to the file or click **Browse** to find the path. • Click **Add to list**. The file name appears on the list of Available SID Translation Files. To remove a SID translation file: • Select the name of the file to remove. • Click **Remove from list**.

File server list

Use this page to add a new file server, or to edit or delete an existing one.

To access this page

From the **Configuration** tab, select **File Servers**.

Field Descriptions

Task	Reference
Reconfigure a file server.	Select the file server name and click **Edit**. The **File Server Properties** page appears.
Delete a file server.	Select the file server name and click **Delete**.
Verify user credentials for CIFS on any file server.	Select the file server name and click **Verify**. If successful, a message authenticating the server appears.
Create new file server.	Click **New**. The **File Server Properties** page appears. Select the file server type and enter properties for that file server.

File server type

Select the type of file server on the File Server Properties page. If the file server you are configuring does not appear on the list, then it is not officially supported. Once the type of file server is selected, the property page for that type of file server automatically appears.

CTA supports the following types of file servers.

- Celerra: EMC Celerra Data Mover for NFS or CIFS
- NetApp: Network Appliance file server for NFS or CIFS
- EMC Centera: EMC Centera file server
- Windows: Windows file server
- Atmos: Atmos file server
- Data Domain: Data Domain file server
- VNX: VNX file server
- VNXe: VNXe file server
- Isilon: Isilon file server

FileMover settings

Use this page to define the callback agent settings for a Celerra or VNX Data Mover. If single or multiple CTAs are deployed in a multiple Celerra environment, this single set of credentials is used across all appliances to authenticate with FileMover and Celerra or Atmos callbacks.

The username and password used for the FileMover settings must match the username and password used when creating the DHSM connection. The Getting Started Guide provides instructions on creating the DHSM connection. If the credentials do not match, attempts to recall after archiving will fail.

To access this page

This page can be accessed in multiple ways. From the **Configuration** tab:

- Select **FileMover Settings**.
- Select **File Servers**. The **File Server List** page appears.

1. Do one of the following:
   2. Click **New** to add a file server. Select **Celerra** or **VNX** as the type of file server to add.
   3. Select a Celerra or VNX to edit.

b. On the **Celerra Properties** or **VNX Properties** page, click **FileMover Settings**.

Field Descriptions

Field	Description
Username	Type the username for the FileMover API user.
Password	Type the password for the FileMover API user.

After adding or updating any properties:

- To accept the changes, click **Commit**.
- To cancel the process, click **Cancel**.

How Do I?

Edit or add a file server.

Import Provider List

Use this page to add, edit, or delete a file server that is allowed to provide a list of files to archive.

To access this page

From the **Configuration** tab, select **Import Providers**.

Field Descriptions

Task	Reference
Reconfigure a file server.	Select the file server name and click **Edit**. The **Provider Properties** page appears.
Delete a file server.	Select the file server name and click **Delete**.
Create new file server.	Click **New**. The **Provider Properties** page appears.

Isilon properties

Use this page to set up a new or existing Isilon as a destination file server. Before configuring Isilon properties on the CTA, review the prerequisites for Isilon in the Getting Started Guide.

To access this page

1. From the **Configuration** tab, select **File Servers**. The **File Server List** page appears.
2. Do one of the following:
3. Select **Isilon** as the type of file server to add.
4. Select an existing Isilon file server to edit.

Field Descriptions

Field		Description
Basic File Server Information	Name	The Isilon server NetBIOS name.
IP Addresses	New IP address	Type the Isilon server IP address. • When editing an existing server, click **Update** to retrieve the IP address from the DNS based on the server name. • To specify an additional IP address, click **Add**. The IP address is added to the list. • To delete an existing IP address, select an address and click **Delete**.
CIFS Specific Settings	Username	Type the username of the Microsoft Windows domain user to be used by the CTA. The Windows domain user should be part of the file server's local administrators' group. NOTE: If the Isilon is an NFS export destination only, this setting is not used.
	Password	Type the Windows domain user password.
	NetBIOS Domain	Type the Windows NetBIOS domain name.

After adding or updating any properties:

- To accept the changes, click **Commit**.
- To cancel the process, click **Cancel**.

Related links

Editing or adding a file server.

Log alert patterns

Use this page to view and edit the log alert patterns. The Rainfinity setup tool can also be used to configure the CTA to monitor various system log files and send e-mail to alert whenever an event occurs.

To access this page

1. From the **Configuration** tab, select **Alert Settings**.
2. Select **Edit Log Alert Pattern**.

Field Descriptions

Each log alert is listed with an identifier and a set of configuration options.

Field	Description
Index	A unique alert index identifies the alert type. A complete listing is provided on the **Alert Pattern Names** page.
Alert pattern name	A set of unique alerts are listed.
Status	Select **enabled** or **disabled**. The different types of alerts may be individually enabled.
Current throttle time	Enter throttle time in minutes. A different throttle time may be applied to each alert. If alerts occur more than once within a specified throttle time, the repeated alerts are suppressed.
Restore throttle time	Click **default** to restore individual throttle times to default values. At the bottom of the column, click **all default** to set the restore throttle time to default values for all alerts.
Included in summary	Select **included** or **not included**. Included alerts appear in the alert summary.

To save change, click **commit**.

More:

SNMP Traps

Log settings

Use this page to review and edit current log settings. The Rainfinity setup tool can also be used to administer and preserve log files.

To access this page

From the **Configuration** tab, select **Log Settings**.

Field Descriptions

Field		Description
Log Files Info		Lists the type, location, and names of various log and configuration files.
Log File Rotation Settings	Rotate type	Select **size** or **time**.
	Rotate frequency	If **time** is selected for Rotate type, select **daily**, **weekly**, or **monthly**.
	Log file copies to keep	Specify the number of log file copies to keep.
	Maximum log file size	If **size** is selected for Rotate type, specify the log file size in **KB**, or **MB** for non-debug files.
	Maximum debug log file size	If **size** is selected for Rotate type, specify the debug file size in **KB** or **MB**.
Log File Backup	Copy log files to a remote server	Select this option to copy files to an external server.
	Remote IP address	The IP address of the external server.
	Remote username	The username to whose account the log files are copied. The user must have write access to the copy directory. Do not specify the root user.
	Remote directory	The directory at the remote site where the log files are placed.

To change or reset any of the settings, make a selection and click **Commit.**

NAS repository and NAS group list

Use this page to manage NAS repositories and NAS groups.

- Use the NAS repository list to add a new a NAS repository, or to edit or delete an existing one.
- NAS repositories may be placed together in NAS groups. Use the NAS Group List to add a new NAS group, or to edit or delete an existing one.

To access this page

From the **Configuration** tab, select **NAS Repository and NAS Group.**

Field Descriptions

Select **NAS Repository List** options to configure a NAS repository:

Task	Reference
Reconfigure a NAS repository.	Select the NAS repository name and click **Edit**. The **Create New NAS Repository** page appears.
Delete a NAS repository.	Select the NAS repository name and click **Delete**.
Create new NAS repository.	Click **New**. The **Edit NAS Repository** page appears.

Select **NAS Group List** options to configure a NAS group:

Task	Reference
Reconfigure a NAS group.	Select the NAS group name and click **Edit**. The **Create New NAS group** page appears.
Delete a NAS group.	Select the NAS group name and click **Delete**.
Create new NAS group.	Click **New**. The **Edit NAS group** page appears.

Related links

Editing or adding a NAS Repository.
Creating a NAS repository as an archive destination when creating a policy.

NetApp file servers

CTA communicates with NetApp filers by using native NetApp API (ONTAPI). This interface must be enabled on each NetApp filer.

To enable access by using the custom NetApp API, type:

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

Verify that this host IP is in `/vol/vol0/etc/hosts.equiv` on each file server.

NetApp properties

Use this page to set up a new or existing NetApp Filer.

To access this page

1. From the Configuration tab, select **File Servers**. The File Server List page appears.
2. Do one of the following:

- Select **NetApp** as the type of file server to add.
- Select a NetApp to edit.

Field Descriptions

Field		Description
Basic File Server Information	Name	The NetApp filer NetBIOS name
IP Addresses	New IP address	Type the IP address of the NetApp filer. • When editing a server, click **Update** to retrieve the IP address from the DNS based on the server name. • To specify an additional IP address, click **Add**. The IP address is added to the list. • To delete an IP address, select the address and click **Delete**. NOTE: Ensure that all IP addresses used to access the file server are listed.
vFiler Host IP	Host IP Address	If using a vFiler, type the IP address of the hosting NetApp filer. Instructions for configuring the hosting vFiler are provided in the Getting Started Guide.
CIFS Specific Settings	Username	Type the username of the Microsoft Windows domain user to be used by the CTA. The Windows domain user should be part of the file server's local administrators group.
	Password	Type the Windows domain user password.
	NetBIOS Domain	Type the Windows NetBIOS domain name.
NDMP Specific Settings	Username	Type the username of the NDMP user on the NetApp source. The NDMP user must belong to the Backup Operators group. File migration uses NDMP settings.
	Password	Type the password for the NDMP user. NOTE: For a NetApp vFiler, the NDMP password is different from the root password. The Getting Started Guide provides instructions on how to retrieve an NDMP password for a root user.
	Port	Type the port used for NDMP operations. The NetApp and the CTA use this port for NDMP traffic. Default is 10000.
NetApp as Source	Enable NetApp as source	Configures the CTA to archive data from the NetApp filer. If more than one CTA is connected to the same NetApp filer, configure only one CTA with this option.

		If more than one CTA is configured to archive data from a single NetApp filer, data loss may occur.
NetApp Local Admin	Admin	Type the username of a local NetApp user that is a member of the administrator group on the NetApp filer.
	Password	Type the password of an admin user.
Directory Exclusion List	Exclude Directory	Specify the names of any directories to exclude. These directories are skipped for all archiving and stub scanning tasks. Directory exclusion does not apply to migration tasks. The Directory exclusion list help page provides details on how to specify excluded directory names. Verify that stub files are not in the excluded directories. CTA will not access the excluded directories and the stub files will become orphans. • To specify an additional directory to exclude, type the directory name and click **Add**. • To delete an existing directory from the list, select the directory and click **Delete**.
NetApp FPolicy Callback Agents	Primary Agent's IP address	For the primary agent, select the agent that is on the same subnet as the NetApp machine. The primary agent recalls all files when it is registered with the NetApp.
	Secondary Agent's IP address	For the secondary agent, select another agent on the same subnet. A secondary agent recalls files when the primary is unavailable. If no such agent exists, select an agent on the next physically closest subnet. Up to two secondaries are supported. Secondary agents may include FMHA appliances. • 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. • If no primary agents respond, then the NetApp filer contacts any of the registered secondary agents. When one of the primary agents is responsive again, the NetApp filer automatically fails back to the primary agent.

After adding or updating any properties:

- To accept the changes, click **Commit**.
- To cancel the process, click **Cancel**.

How Do I?

Edit or add a file server.

NetApp FPolicy special clients

Use this page to configure NetApp FPolicy Special Clients. All registered FPolicy callback agents or daemons are listed on this page.

To access this page

From the **Configuration** tab, select **NetApp FPolicy Special Clients**.

Field Descriptions

Field	Description

Configured FPolicy Callback Agents	For each callback agent, the name and IP address is listed.
	• Click View Status to invoke a new window.
	  ◦ Window shows the status of the agent
	  ◦ Window lists the recall status of all NetApp filers registered with this agent.
	• To delete a callback agent, select the **To Delete** checkbox for the agent and click **Commit**.
	In addition to deleting the callback agent in the GUI, you must also remove the CTA configuration from the FPolicy Callback Daemon (FCD). To remove the IP address of the CTA from the FCD, type:
	`fpsetup.sh rm_rffm <ip_address>`
	
	If you do not remove the CTA configuration and you restart the FCD or simply leave the FCD running, it will reregister itself to the CTA and appear in the CTA GUI.

Excluded Clients	Lists the IP addresses of clients that cannot recall. While these clients are excluded from recall, they may read the stub files. For example, by adding the excluded client IP addresses here, you enable the stub awareness feature that enables you to move stubs without recalling the files.
	Three types of input are accepted:
	• Single IP address.
	• DNS hostname or the entry that exists on the DNS. If entered, the CTA goes to the DNS server to get the IP address for the host. NOTE: If there is more than one IP address associated with the DNS hostname, do not use this method of input. Manually type the IP addresses.
	• Wildcard IP range. For example:
	  ◦ 10.1.1.0/24 represents the IPs ranging from 10.1.1.0~10.1.1.254
	  ◦ 10.1.1.128/27 represents the IPs ranging from 10.1.1.128~10.1.1.159.
	• To specify an additional IP address to exclude, type the IP address and click **Add**.
	• To remove an existing IP address from the list of Excluded Clients, select the IP address and click **Delete**.
	
	NOTE: If a backup or virus scanning program such as vScan runs on your

	NetApp filer, add the IP address of the machines that host these applications as excluded clients. This allows the virus scanner to scan the stub file upon a recall event.
Blocked Clients	Lists the IP addresses of clients that are completely blocked from stub file access. • To specify an additional IP address to block, type the IP address and click **Add**. • To remove an existing IP address from the list of blocked clients, select the IP address and click **Delete**.

After adding or updating any properties:

- To accept the changes, click **Commit**.
- To cancel the process, click **Cancel**.

NetApp FPolicy special client status

This page lists the status of each NetApp FPolicy special client.

To access this page

1. From the **Configuration** tab, select **NetApp FPolicy Special Clients**.
2. Select a configured FPolicy Callback agents and click **View Status**.

Field Descriptions

At the top of the page, statistics for the configured FPolicy Callback agent are listed. A list of the NetApp FPolicy special clients appears below.

Field	Description
NetApp Name	NetApp with which FCD is registered for callbacks.
NetApp IP	IP address FCD uses for FPolicy management and status.
Registration Time	The last time that FCD registered for callbacks from this NetApp.
Last Recall Time	The last time that FCD successfully recalled a file.
Last Recall Error Time	The last time that FCD failed to recall a file.
Last Recall Error	Error message that corresponds to the last time that FCD failed to recall a file.
Num Recalls OK	The number of files that were successfully recalled.
Num Recalls Failed	The number of failed attempts to recall files.
Num Recalls in Progress	The number of file recalls currently in progress.

NOTE: All values are recorded since the last time the NetApp FPolicy service was started.

Provider Properties

Use this page to define an external provider that is allowed to copy an archive file list to the CTA. A third-party software administrator controls the external provider. After configuring provider properties, give the provider name and password to the third-party software administrator.

The Getting Started Guide provides details on configuring the CTA prior to using the import file feature. In addition, Imported File List Archive Task Technical Notes are available from EMC Powerlink.

To access this page

From the **Configuration** tab, select **Import Providers**. The **Import Provider List** page appears.

- Click **New** to create a provider or
- Select a provider to edit.

Field Descriptions

Field	Description
Name	Type the name of the user that is allowed to log into the CTA from the provider. This must be a valid Linux user name.
Password	Type the password used to log into the CTA from the provider.

After adding or updating any properties:

- To accept the changes, click **Commit**.
- To cancel the process, click **Cancel**.

How Do I?

Configure an archive file list provider.

Rainfinity setup tool

The Rainfinity Setup Tool is a menu-driven tool, which is accessed from the CLI, and is provided to perform basic setup tasks that are not available through the GUI. To access the setup tool, you must be logged into the CLI as root.

To use the Rainfinity Setup Tool, an admin user (user who belongs to the wheel group) must:

1. Log in to the Cloud Tiering Appliance.

- Log in as the super user
- Type the root password.

2. Type **rfhsetup** and press **Enter**. The **Rainfinity Setup Tool** appears.

More

Getting Started Guide

User properties

Use this page to define the properties for a user.

To access this page

1. From the **Configuration** tab, select **Users**.
2. Do one of the following:
    3. Select a user to edit.
    4. Click **Add a New User**.

Field Descriptions

Field		Description
User Information	Name	If editing a user, this field appears with the current username. Type a new username to add a new user.
	Enabled?	Selected by default. If not selected, then the user cannot log in.
	New Password	Password for the user account.
	New Password Again	Type the new password again.
User Access Rights	Super User	If selected, the user has administrative privileges.
	Regular User	If the user does not have administrative privileges, click the level of user privileges for this user.
Add or Edit		Click **Commit** to add or edit the user.

Users

Use this page to edit, delete, or add a user.

To access this page

From the **Configuration** tab, select **Users**.

Field Descriptions

Select options to configure a user:

Task	Options
Reconfigure a user.	Select the username. The **User Properties** page appears.
Delete a user.	Select a user, and click **Delete**.
Create new user.	Click **Add a New User**. The **User Properties** page appears.

SNMP configuration

Use this page to configure SNMP settings. The Rainfinity setup tool can also be used to configure SNMP community strings and notification hosts.

To access this page

From the **Configuration** tab, select **SNMP Configuration**.

After changing any of the SNMP settings, click **Commit**. Otherwise, the changes will not be applied.

Field Descriptions

Field		Description
Community Strings	New Community String	The text string acts as a password and is used to authenticate messages that are sent between the SNMP manager and the SNMP agent. The string must be alphanumeric and can include dashes and underscores.
	Security Type	Select the security type. CTAs support SNMPv1 and SNMPv2c.
	Add	Click to add the new string to the current community string list. An unlimited number of strings can be added.
Current Community String List		If defined, this list includes all Community Strings with corresponding Security Types. To delete a listed string, click **Delete**.
Notify Target	Notify IP Address	The IP address of the notification host to which SNMP alerts will be sent.
	Notify UDP Port	The UDP port of the notification host.
	Community String	Type a community string. The string must be alphanumeric and can include dashes and underscores.
	Security Type	Select the security type. Cloud Tiering Appliances support SNMPv1 and SNMPv2c.
	Add Target	Click to add the new target to the Current Notification Target List. An unlimited number of targets may be added.
Current Notification Target List		If defined, this lists the IP addresses, UDP ports, community strings, and security types of all notification targets. To delete a listed target, click **Delete Target.**

Related links:

SNMP Traps

SNMP Traps

The CTA can send an e-mail notification for the following 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	1.3.6.1.4.1.1139.9.3.2.0.1
eRARainfinityAlert	EMC-RAINFINITY-ALERTS-MIB	1.3.6.1.4.1.1139.9.3.2.0.1
eRAGenericAlert	EMC-RAINFINITY-ALERTS-MIB	1.3.6.1.4.1.1139.9.3.2.0.1
eRASecurityAlert	EMC-RAINFINITY-ALERTS-MIB	1.3.6.1.4.1.1139.9.3.2.0.1
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

System security settings

This page lists current security settings, but does not allow editing. The *Getting Started Guide* describes system security settings in more detail, including how to use the Rainfinity Setup Tool to reset them.

To access this page

From the **Configuration** tab, select **System Security Settings.**

Field Descriptions

Field		Description
Basic Security Settings	Security level	Set to either **default** or **harden**. The admin user sets the security level by using the CLI.
	Password hardening	If a password command is run with password hardening enabled, the new password must be at least eight characters long and satisfy multiple requirements.
	Disable root user login	If root logins are disabled, the only way to add new users, or to run rfhsetup, is for an admin user to: 1. Log in to the device. 2. Type **su** to become the root user.
	Use single security database	If set to yes, users can use the same username and password to connect to the appliance using the CLI or GUI.
	Last database config time	If set to yes, the last date and time the security settings were changed appear.
Password Aging Settings		If password aging is enabled, every user (except root) who can log in with a shell account will have an aging password.
LDAP Settings		If LDAP authentication is set, user credentials are checked against an LDAP database.
Radius Settings		If RADIUS authentication is set and LDAP authentication fails, user credentials are checked against a RADIUS database.

VNX properties

Use this page to set up a new or existing VNX Data Mover.

To access this page

1. From the **Configuration** tab, select **File Servers**. The **File Server List** page appears.
2. Do one of the following:

- Click **New** to add a file server. Select **VNX** as the type of file server to add.
- Select a VNX to edit.

Field Descriptions

Field		Description
FileMover Settings		Click the link to configure the credentials for archive and recall.
Basic File Server Information	Name	The VNX server NetBIOS name or hostname. NOTE: This is not the Fully Qualified Domain Name (FQDN) or IP address.
	File server is VDM	Identifies the VNX as a Virtual Data Mover. Virtual Data Movers support only the CIFS protocol.
IP Addresses	New IP address	Type the VNX Data Mover IP address. • When editing an existing server, click **Update** to retrieve the IP address from the DNS based on the server name. • To specify an additional IP address, click **Add**. The IP address is added to the list. • To delete an existing IP address, select an address and click **Delete**.
Control Station	IP address	Type the IP address of the VNX Control Station. • For file migration, this is required for all VNX servers involved in CIFS transactions. • For archiving, this allows the CTA to automatically perform some preconfiguration steps for archiving. If this field is empty, the CTA takes no action and the preconfiguration steps must be performed manually. NOTE: A system user for the XML API and FileMover API operations must be configured with a valid password on the VNX Control Station.
CIFS Specific Settings	Username	Type the username of the Microsoft Windows domain user to be used by the CTA. The Windows domain user should be part of the file server's local administrator's group. The CIFS credential is not required if only the VNX performs NFS archiving.
	Password	Type the Windows domain user password.

	NetBIOS Domain	Type the Windows NetBIOS domain name.
NDMP Specific Settings	Username	Type the username of the NDMP user on the source and destination VNX. Offline file migration uses NDMP settings.
	Password	Type the password for the NDMP user.
	Port	Type the port used for NDMP operations. The VNX and the CTA use this port for NDMP traffic. Default is 10000.
VNX as Source	Enable VNX as source	Configures the CTA to archive data from the VNX. If more than one CTA is connected to the same VNX, configure only one CTA with this option. If more than one CTA is configured to archive data from a single VNX, data loss might occur.
VNX Callback Agent Settings	CCD DNS Name	Type the FQDN of the VNX Callback DNS entry. NOTE: The FQDN is case-sensitive. The username and password for the FileMover API and VNX HTTP authentication credentials are provided on the FileMover Settings page.
Atmos Callback Agent Settings	ACD DNS Name	Type the FQDN of the Atmos Callback DNS entry. NOTE: The FQDN is case-sensitive.
Directory Exclusion List	Exclude Directory	Specify the names of any directories to exclude. These directories are skipped for archiving and stub scanning tasks. Directory exclusion does not apply to migration tasks. The Directory exclusion list help page provides details on how to specify excluded directory names. Verify that stub files are not in the excluded directories. CTA will not access the excluded directories and the stub files will become orphans. • To specify an additional directory to exclude, type the directory name and click **Add**. • To delete an existing directory from the list, select the directory and click **Delete**.

After adding or updating any properties:

- To accept the changes, click **Commit**.
- To cancel the process, click **Cancel**.

How Do I?

Edit or add a file server.

VNXe properties

Use this page to set up a new or existing VNXe server.

To access this page

1. From the **Configuration** tab, select **File Servers**. The **File Server List** page appears.
2. Do one of the following:

- Click **New** to add a file server. Select **VNXe** as the type of file server to add.
- Select a VNXe to edit.

Field Descriptions

Field		Description
FileMover Settings		Not applicable because VNXe is not a supported as archive source.
Basic File Server Information	Name	The VNXe server NetBIOS name or hostname. NOTE: This is not the Fully Qualified Domain Name (FQDN) or IP address.
	File server is VDM	Not applicable because VNXe is not supported as a Virtual Data Mover.
IP Addresses	New IP address	Type the VNXe IP address. • When editing an existing server, click **Update** to retrieve the IP address from the DNS based on the server name. • To specify an additional IP address, click **Add**. The IP address is added to the list. • To delete an existing IP address, select an address and click **Delete**.
Control Station	IP address	Not applicable because VNXe is not supported as a Control Station.
CIFS Specific Settings	Username	Type the username of the Microsoft Windows domain user to be used by the CTA. The Windows domain user should be part of the file server's local administrator's group.
	Password	Type the Windows domain user password.
	NetBIOS Domain	Type the Windows NetBIOS domain name.
NDMP Specific Settings	Username	Type the username of the NDMP user on the migration source and destination VNXe. Offline file migration uses NDMP settings.
	Password	Type the password for the NDMP user.
	Port	Type the port used for NDMP operations. The VNXe and the CTA use this port for NDMP traffic. Default is 10000.
VNXe as Source	Enable Celerra as source	Not applicable because VNXe is not a supported as archive source.
Celerra Callback Agent Settings	CCD DNS Name	Not applicable because VNXe is not a supported as archive source.
Atmos Callback	ACD DNS	Not applicable because VNXe is not a supported as

Agent Settings	Name	archive source.
Directory Exclusion List	Exclude Directory	Not applicable because VNXe is not a supported as archive source.

After adding or updating any properties:

- To accept the changes, click **Commit**.
- To cancel the process, click **Cancel**.

How Do I?

Edit or add a file server.

Windows domain user

The Windows domain user is required when a new file server is configured and the CTA requires administrative access to CIFS data sets for archiving and recall purposes. For the CTA, the domain user is typically referred to as fmadmin, but any user with sufficient administrator privileges works.

The domain user is the username used for CIFS-specific properties on:

- Celerra Properties
- VNX Properties
- VNXe Properties
- NetApp Properties
- Windows Properties
- Isilon Properties

The Getting Started Guide provides instructions on how to create or add a Windows domain user.

Windows properties

Use this page to set up a new or existing Windows 2003 or Windows 2008 as a destination file server.

To access this page

1. From the **Configuration** tab, select **File Servers**. The **File Server List** page appears.
2. Do one of the following:
3. Select **Windows** as the type of file server to add.
4. Select an existing Windows file server to edit.

Field Descriptions

Field		Description
Basic File Server Information	Name	The Windows server NetBIOS name.
IP Addresses	New IP address	Type the Windows server IP address. • When editing an existing server, click **Update** to retrieve the IP address from the DNS based on the server name. • To specify an additional IP address, click **Add**. The IP address is added to the list. • To delete an existing IP address, select an address and click **Delete**.
CIFS Specific Settings	Username	Type the username of the Microsoft Windows domain user to be used by the CTA. The Windows domain user should be part of the file server's local administrators' group. Windows servers only support the CIFS protocol.
	Password	Type the Windows domain user password.
	NetBIOS Domain	Type the Windows NetBIOS domain name.

After adding or updating any properties:

- To accept the changes, click **Commit**.
- To cancel the process, click **Cancel**.

Related links

Editing or adding a file server.

Policies

Policies Overview

A policy is a set of one or more user-created rules. For each rule, an archive, delete, or file migration action may be associated. When a policy is applied to an export or share in a Cloud Tiering Appliance environment, the rules are applied to each file in the export or share. If a file falls within the rule, the action is taken.

For example, a rule might be: If a file has not been accessed in more than a year, archive it.

Rules can also be written to prevent archiving.

Rules are created by using file matching expressions that contain AND logic. For example:

- If a file is larger than 5 MB AND its filename has a suffix of .doc AND it has not been accessed in 8 months, archive it.
- If a file has been modified in the past year AND its filename contains the letters ENGR, do not archive it.

The rules in a policy are checked in order, so a logical OR is effected. This means, as soon as a rule resolves to TRUE, the rule action is taken, and no further rules are applied to that file. If there are multiple rules that might resolve to TRUE for a given file, only the first one that resolves to TRUE will be executed, so the order that rules are arranged in a multi-rule policy must be carefully considered.

Using policies

Policies are built upon rules that are built upon expressions. The policy-creation process begins by defining a file matching expression. These expressions are combined to make rules. Rules are file matching expressions with an action to either archive or don't archive, delete or don't delete, migrate or don't migrate. If archive is the action, a destination for the archive is defined with the rule.

Each policy must have at least one rule, but can have more than one. Files are evaluated for action based on the rules defined. File attributes are compared with the file matching expression listed for each rule, starting with the first rule listed which is Rule 0. Once a

file matches a rule, the action is executed and other rules are not checked for that file.

How Do I?

Create a multi-tier policy
Create rules

Create a file migration policy

The policy-based migration example described in the overview involves configuring a policy to migrate files greater than 0 bytes. This catch-all policy will migrate all files under a given share or export.

To create a file migration policy

1. From the **Policies** tab, click the link in the upper right corner to **Create new policy**. The **Create Policy** page appears.
2. Configure policy settings.

Field	Description
Policy Name	Type **Migrate files GR 0 bytes**.
Policy Type	Select **migrate_file**.

3. To define a rule for this policy, click **Add Rule**. The **Add File Matching Criteria to Rule** page appears.
4. Select **New File Matching Expression**, and provide the settings to build the expression. You cannot type entries directly into the **Expression** field.

Field	Description
File Attributes	Select **size**.
Operators	Select **>**.
Attribute Values	Type **0** and select **byte(s)** from the list.

5. Click **Add to Rule**. **size > 0 bytes** appears in the **Expression** field.
6. For **Select Action**, select **Migrate**. Because the task is a migration, the archive destination does not apply and appears greyed out.
7. Click **Save Rule**. The **Edit Policy** page reappears with the new rule listed.
8. Click **Save Policy & Schedule**. The **Create New Task** page appears.

To proceed with the example, see Create, run, and complete a file migration task.

How Do I?

Create rules.

Create a multi-tier policy

The multi-tier archiving example described in the overview involves configuring a multi-tier policy to identify the following:

- Files that are more than three months old for archiving to a NAS repository.
- Files that are more than six months old for archiving to an EMC Centera.

The NAS repository is created as part of the policy rule creation for files last accessed more than three months ago. The EMC Centera destination was previously configured as part of configuring CTA for multi-tier archiving and will be the repository for files that were last accessed more than six months ago.

Create the multi-tier policy with two rules:

- One rule that archives files older than six months to an EMC Centera.
- One rule that archives files older than three months to a NAS repository.

To create a multi-tier policy

1. From the **Policies** tab, click the link in the upper right corner to **Create new policy**. The **Create Policy** page appears.
2. Configure policy settings.

Field	Description
Policy Name	Type **Multi tier old files**.
Policy Type	Select **multi_tier**. This type of policy evaluates both normal and stub files.
Retention Period	Leave the value set to 0 years. A value greater than zero will set a protection that prohibits files from being deleted from the destination until the retention period has expired.
Delayed Period	Leave the value set to 0 days. The file gets stubbed on the source when any active policy matches the file attribute and after the delay stubbing time has expired. The time increment is defined in days. The default value is 0.
Stub Retention	Leave cleared. If selected, stubs are retained on the source for the Retention Period and cannot be deleted until that period has elapsed. Stubs under retention are not subject to multi-tier archiving if the files on the destination are under retention.

3. To define the first rule for this policy, click **Add Rule**. The **Add File Matching Criteria to Rule** page appears.
4. Select **New File Matching Expression**, and provide the settings to build the expression. You cannot type entries directly into the **Expression** field.

Field	Description
File Attributes	Select **last_accessed**.
Operators	Select **>**.
Attribute Values	Type **6** and select **months** from the list.

5. Click **Add to Rule**. **last_accessed > 6 months** appears in the **Expression** field.
6. For **Select Action**, select **Archive**. For **Archive Destination**, select **Centera** and select an EMC Centera from the list.
7. Click **Save Rule**. The **Edit Policy** page reappears with the new rule listed.
8. To define the second rule for this policy, click **Add Rule**. The **Add File Matching Criteria to Rule** page appears.
9. Select **New File Matching Expression**, and provide the settings to build the expression. You cannot type entries directly into the Expression field.

Field	Description
File Attributes	Select **last_accessed**.
Operators	Select **>**.
Attribute Values	Type **3** and select **months** from the list.

10. Click **Add to Rule**. **last_accessed > 3 months** appears in the **Expression** field.
11. For **Select Action**, select **Archive**.
12. For **Archive Destination**, select **Add NAS Repository**.
13. From the list, select a protocol and select a repository. Click **Add to List**.

    If no repository has been defined, click **NAS Repo & Groups**. Follow instructions on create a NAS repository as an archive destination.
14. Click **Save Rule**. The **Edit Policy** page reappears with both rules listed.

    The order of the rules is important. The rules are evaluated from top to bottom. Files that fail the first rule are caught by the second rule. To ensure that files between three and six months old are not archived to the EMC Centera, the greater than six month rule must be the first rule evaluated.
15. Click **Save Policy & Schedule**. The **Create New Task** page appears.

To proceed with the example, see Create a multi-tier archiving task.

How Do I?

Create rules.
Create or edit a NAS repository.

Create a NAS repository as an archive destination

While a rule is being added to a policy, a NAS repository or NAS group can be added as an archive destination. The NAS repository may be a Celerra Data Mover, VNX, VNXe, NetApp Filer, Data Domain, Isilon, or Windows server.

To create a NAS repository as part of policy creation

1. On the **Policies** tab, select **Create new policy**. The **Create Policy** page appears.
2. Click **Add Rule**. The **Add File Matching Criteria to Rule** page appears.
3. Select **Archive** from the **Select Action** list. Click **NAS Repo & Groups**. The **NAS Repository and NAS Group List** page appears.
4. Click **New** to create a NAS repository. The **NAS Repository Configuration** page appears.
5. Configure the NAS repository settings.
6. Click **Save Repository**. The **NAS Repository and NAS Group List** page reappears with the new repository listed.
7. Click **Finish**. The **Add File Matching Expression to Rule** page reappears.
8. For **Archive Destination**, click **Add NAS Repositories**.
9. Select the new NAS repository from the list and click **Add to List**.

Create rules

To derive the full benefit of the CTA file matching feature, combine expressions to form rules. Rules are used to determine which files to archive, or which files not to archive.

The major difference between individual rules and the expressions within them is:

- Multiple rules for a policy are evaluated in a large OR statement, where the first rule to match in the top-down list is the one to take effect. All subsequent rules are skipped. If no rules are matched, no action occurs.
- Every expression within a rule must be matched in order for that rule to qualify as a match. All the expressions are essentially processed in a large AND statement.

Two examples illustrate different ways to accomplish the same task, using one rule or many. While combining all expressions into a single rule might appear to be the simplest approach, maintenance and administration can become more difficult. With separate rules, it is easier to make additions and changes.

To create a single rule to archive all .doc files over 30 days old:

1. On the **Policies** tab, select **Create new policy**. The **Create Policy** page appears.
2. Click **Add Rule**. The **Add Rule** page appears.
3. To build a **New File Matching Expression**, specify the following properties, and click **Add to Rule**:

Field	Value
File Attributes	Select **filename**.
Operators	Select **matches regex**.
Attribute Values	Type **\.doc$**.

4. To build another **New File Matching Expression**, specify the following properties, and click **Add to Rule**:

Field	Value
File Attributes	Select **last_accessed**.
Operators	Select **>**.
Attribute Values	Type **30**. Select **days**.

5. For **Select Action**, select **Archive**, and select an **Archive Destination**.
6. Click **Add to List**.
7. Click **Save Rule**. The **Create Policy** page reappears.
8. Click **Save Policy**.

To create multiple rules to archive all .doc files over 30 days old:

1. On the **Policies** tab, select **Create new policy**. The **Create Policy** page appears.
2. Click **Add Rule**. The **Add Rule** page appears.
3. To build a **New File Matching Expression**, specify the following properties:

Field	Value
File Attributes	Select **last_accessed**.
Operators	Select **>**.
Attribute Values	Type **30**. Select **days**.

1. Click **Add to Rule**.
2. For **Select Action**, select **Don't Archive.**
3. Click **Save Rule**. The **Create Policy** page reappears.
4. Click **Add Rule**. The **Add Rule** page reappears.
5. To build another **New File Matching Expression**, specify the following properties:

Field	Value
File Attributes	Select **filename**.
Operators	Select **matches regex**.
Attribute Values	Type **\.doc$**.

1. Click **Add to Rule**.
2. For **Select Action**, select **Archive.** Select an **Archive Destination**.
3. Click **Save Rule**. The **Create Policy** page reappears.
6. Click **Save Policy**.

Considerations for adding rules

If the task of archiving all .pdf files is added, either the single-rule method or the multi-rule method requires one of the following types of new rules:

- The single-rule method requires a new rule with two expressions, less than 30 days and .pdf.
- The multi-rule method requires one additional rule, a .pdf match after the .doc match. The less than 30 days rule would still apply, and any file without .doc would be compared for .pdf. A file with .pdf would match this new rule and would be archived.

In this example, the difference is small. Only one new rule is required in either case. However, if additional tasks are added several times, for .exe, .txt, .xls, and .ppt files for example, each change requires creating a rule. If a change to the file size is added after creation of these rules, the process for each method becomes significantly different.

- Using the single-rule method, every new rule must be edited to add an expression to specify size greater than 100 KB.

- Using the multi-rule method, add a single rule to the top of the policy that specifies "Don't Archive files with size less than 100 KB." In this way, any file less than 100 KB, regardless of extension, will not be archived and will not require additional processing time.

When considering the possibility of hundreds of rules, this may become a major difference between the single-rule and multi-rule methods.

Guidelines for ordering rules

While it is possible to add multiple expressions under a single rule, the proper ordering of rules provides the most flexibility for expansion and changes.

To filter out the greatest number of files early so they do not pass through rules unnecessarily:

- Order the most general Don't Archive rules so that they are processed first. Examples include less than 30 days, less than 100 kilobytes, all files with NEW in the beginning, and so on.
- Order more specific Archive rules so that they are processed later. Examples include .txt files, .doc files, _old files, and so on

This same guideline applies when modifying or expanding policies.

How Do I?

Create file matching expressions
Regular expressions

Page help
Add file matching criteria to rule

Use this page to create or edit a policy rule by defining:

- File Matching Criteria
- Action type

To access this page

This page can be accessed in multiple ways. From the **Policies** tab:

- Select Create new policy and click **Add Rule**.
- Edit an existing policy and click **Add Rule**.

Field Descriptions

Field		Description
Add File Matching Criteria to Rule	Expression	Combination of individual expressions strung together using &&. You cannot enter data by typing directly into the field. Use Build Expression or Edit Expression to add or remove data.

	Build Expression	Use a combination of these options to add data to the Expression field.
		• New File Matching Expression — To define a new expression, select attributes and operators. The operators and attribute values change depending upon the file attribute selected. Once defined, click **Add to Rule** to write the new expression to the Expression field.
		  • When setting the **last_accessed** file attribute, enter a period greater than three months. A shorter period could slow performance with the constant archiving and retrieving of files due to client access.
		  • If the operator is **matches regex**, Regular Expressions can be used.
		• Saved File Matching Expression — To use a previously saved expression, highlight the expression and click **Add to Rule**.

	Edit Expression	Individual expressions that are added to the Expression field also appear in this list. To remove an individual expression from the Expression field, select the expression and click **Delete Selected**.
Select Action		Once the file selection criteria are set by the Expression, choose the action to take if the criteria are met. For archive, multi-tier, or multi-tier stub policy types:

		• Archive • Don't Archive For delete orphans or delete stubs policy types: • Delete • Don't delete For migrate file policy types: • Migrate • Don't migrate
	Archive Destination (only selectable for archive, multi-tier, or multi-tier stub policy types)	Select an existing destination: • Centera • Add NAS Repositories • Atmos • NAS Groups If no NAS repositories or NAS groups have been defined, click NAS Repo&Groups to add a new repository or group.

After creating a new rule:

- To save the rule, click **Save Rule**.
- To cancel creating the new rule, click **Cancel**.

How Do I?

Create rules

Create file matching expression

Use this page to create a file matching expression.

To access this page

From the **Policies** tab, select **Create new file matching expression.**

Field Descriptions

Field	Description
Name	The name for the expression.
Expression	Combination of individual expressions strung together using &&. You cannot enter data by typing directly into the field. Use Build Expression or Edit Expression to add or remove data.

Build Expression	Use a combination of these options to add data to the Expression field.
	• New File Matching Expression — To define a new expression, select attributes and operators. The operators and attribute values will change depending upon the file attribute selected. Once defined, click **Add to Rule** to write the new expression to the Expression field.
	  • When setting the **last_accessed** file attribute, enter a period greater than three months. A shorter period could slow performance with the constant archiving and retrieving files due to client access.
	  • If the operator is **matches regex**, Regular Expressions can be used.
	• Saved File Matching Expression — To use a previously saved expression, highlight the expression and click **Add to Rule**.

Edit Expression	Individual expressions that are added to the Expression field also appear in this list. To remove an individual expression from the Expression field, highlight the expression and click **Delete Selected**.

After creating a new file matching expression:

- To save the new file matching expression, click **Save File Matching Expression**.
- To cancel creating the new file matching expression, click **Cancel**.

How Do I?

Create rules

Create policy

Use this page to create a new policy.

To access this page

From the **Policies** tab, select **Create new policy.**

Field Descriptions

Field	Description
Policy Name	Type a new policy name.
Policy Type	Select one of the policy types: • archive — Evaluates normal files to be archived. • multi_tier — Evaluates all normal and stub files to be archived. • multi_tier_stub — Evaluates stub files to be archived. • delete_orphans — Automatically deletes orphans. • delete_stubs — Automatically deletes stubs and associated archive files when retention expires. • migrate_file — Evaluates files to be migrated between primary file servers.
Retention Period	Applies to the policy types: archive, multi_tier, or multi_tier_stub. The period of time for which a file is to be retained on the Celerra, NetApp, EMC Centera, and Data Domain destinations. If Stub Retention is selected, the same retention period applies to stub files on the Celerra or VNX source. The increments of time are defined in terms of years, months, weeks, or days. The default value for the retention period is 0 years. NOTE: To retain files, file retention must be enabled on the destination. Platform documentation provides instructions on how to enable file retention.
Delay Period	Applies to the policy types: archive, multi_tier, or multi_tier_stub for both EMC Centera and NAS destinations. Type the number of days after archiving that files should be stubbed on the source. The default setting is 0 days.
Stub Retention	If selected, stubs are retained on the source for the Retention Period. Stubs under retention are not subject to multi-tier archiving. This option applies to archiving from Celerra or VNX to any destination except Atmos. It is not supported for NetApp sources. NOTE: Stubs under retention cannot be modified or deleted and can only be deleted after the retention period has elapsed. Stubs that have been under retention can never be moved or renamed.
Days missing more than	Applies only to the policy type: delete_orphans. Type the period of time to wait after archiving before orphan files are deleted. The default setting is 30 days.
Delete All	Applies only to the policy type: delete_stubs. When selected, all stubs for which:

	• The retention period is 0 • The retention period has expired and that are matched by the policy will be deleted. By default, the selection is cleared.

After creating a new policy:

- To Add Rule, click **Add Rule**. Creating Rules provides more information about using rules.
- To save the policy, click **Save Policy**.
- To save the policy and the schedule, click **Save Policy & Schedule**.
- To cancel the policy creation, click **Cancel**.

Delete stubs policy

The delete_stubs policy deletes stubs and the associated repository files, based on policy rules. For example, a policy may delete all stubs and the files they point to if the files are older than two years.

When creating a delete_stubs policy, a Delete All box is selectable. When this box is selected, all data that is not under retention will be deleted. By default, this setting is not selected.

The delete_stubs policy deletes data and is irreversible. Exercise caution when using this policy. Either backup data or run a simulation before executing the policy task.

Regarding retention, there are three kinds of stubs:

- Stubs without retention — These stubs were created with an archive policy that did not specify a retention period.
- Stubs with unexpired retention — These stubs were created with an archive policy that did specify a retention period that has not yet expired. So these stubs cannot be altered in any way.
- Stubs with expired retention — These stubs were created with an archive policy that specified a retention period that has expired.

Stubs without retention and stubs with expired retention are similar. Both can be deleted by a client, which creates orphans. Both can be replaced by a real file when a full recall occurs. However the policy behaves differently depending on whether Delete All is selected.

- If Delete All is selected, the delete_stubs policy will apply policy rules to stubs without retention.
- If Delete All is not selected, the delete_stubs policy will apply policy rules to stubs with expired retention, but not to stubs without retention.

The delete_stubs policy rules never affect stubs with unexpired retention regardless of the Delete All setting.

Edit file matching expression

Use this page to edit file matching expression based on file attributes.

NOTE: Multiple policies can actively reference the same file matching expression. Before editing and saving the expression, verify that any affected policies will be valid with the change.

To access this page

The **Policies** tab lists the file matching expressions. Click the down arrow next to any file matching expression listed and select **Edit**.

Field Descriptions

Field	Description
Name	The name for the expression.
Expression	Combination of individual expressions strung together by using &&. You cannot enter data by typing directly into the field. Use Build Expression or Edit Expression to add or remove data.

Build Expression	Use a combination of these options to add data to the Expression field.
	• New File Matching Expression — To define a new expression, select attributes and operators. The operators and attribute values will change depending upon the file attribute selected. Once defined, click **Add to Rule** to write the new expression to the Expression field.
	  • When setting the **last_accessed** file attribute, enter a period greater than 3 months. A shorter period could slow performance with the constant archiving and retrieving files due to client access.
	  • If the operator is **matches regex**, Regular Expressions can be used.
	• Saved File Matching Expression — To use a previously saved expression, highlight the expression and click **Add to Rule**.

Edit Expression	Individual expressions that are added to the Expression field also appear in this list. To remove an individual expression from the Expression field, select the expression and click **Delete Selected**.

After creating a new file matching expression:

- To save the file matching expression, click **Save File Matching Expression**.
- To cancel creating the file matching expression, click **Cancel**.

Edit policy

Use this page to edit an existing policy.

To access this page

The **Policies** tab, lists the policies. Click the down arrow next to any policy listed and select **Edit**.

Field Descriptions

Field	Description
Policy Name	The policy name appears.
Policy Type	Policy types are: • archive — Evaluates normal files to be archived. • multi_tier — Evaluates all normal and stub files to be archived. • multi_tier_stub — Evaluates stub files to be archived. • delete_orphans — Automatically deletes orphans. • delete_stubs — Automatically deletes stubs and associated archive files when retention expires. • migrate_file — Evaluates files to be migrated between primary file servers. The policy type cannot be changed after creating rules.
Retention Period	Applies to the policy types: archive, multi_tier, or multi_tier_stub. The period of time for which a file is to be retained on the Celerra, NetApp, EMC Centera, and Data Domain destinations. If Stub Retention is selected, the same retention period applies to stub files on the Celerra or VNX source. The increments of time are defined in terms of years, months, weeks, or days. The default value for the retention period is 0 years.
Delay Period	Applies to the policy types: archive, multi_tier, or multi_tier_stub for both EMC Centera and NAS destinations. Type the number of days after archiving that files should be stubbed on the source. The default setting is 0 days.
Stub Retention	If selected, stubs are retained on the source for the Retention Period. Stubs under retention are not subject to multi-tier archiving. This options applies to archiving from Celerra or VNX to any destination except Atmos. It is not supported for NetApp sources. NOTE: Stubs under retention cannot be modified or deleted and can only be deleted after the retention period has elapsed. Stubs that have been under retention can never be moved or renamed.
Days missing more than	Applies only to the policy type: delete_orphans. Type the period of time to wait after archiving before orphan files are deleted. The default setting is 30 days.
Delete All	Applies only to the policy type: delete_stub. When selected, all stubs will be deleted without retention. By default, the selection is cleared.

After editing a new policy:

- To Add Rule, click **Add Rule**. Creating Rules provides more information about using rules.
- To save the policy, click **Save Policy**.
- To save the policy with a different policy name, click **Save As**
- To save the policy and the schedule, click **Save Policy & Schedule**.
- To cancel the policy edit, click **Cancel**.

File attributes for file matching expressions

File attributes are components of the file matching expression. Attributes vary depending upon the policy type. For example, the last_modified and last_attr_changed attributes do not apply to the delete_orphans or delete_stubs policy types.

Attribute	Description	Example
last_accessed	The last time that a client application program read the file.	Reading a file updates this value. Opening, reading, and closing a file without saving it changes the last_accessed time to the time of the last client read operation. This value maps to the NFS atime and CIFS lastAccessedTime attributes.
last_modified	The last time that a client application program wrote to the file.	Writing to a file updates this value. Opening, editing, saving and closing a file changes the last_modified time to the time of the last client write operation. This value maps to the NFS mtime and CIFS lastModifiedTime attributes.
last_attr_changed	The last time that the file attributes (permissions) were changed, or that a client application program wrote the file.	Updating the attributes of a file using chmod (Unix) or on the Properties page (Windows) changes the last_attr_changed time. Writing to a file also updates this value. This value maps to the NFS ctime and CIFS lastChangedTime attributes.
size	The size of the file.	Files of a particular size are evaluated for the policy.
filename	The name of the file.	Files with certain names or characters are evaluated for the policy. Regular expressions can be implemented.
dirname	The name of the directory.	Directories with specified names or characters are evaluated for the policy. Regular expressions can be implemented.

Multi-tier policy

A multi-tier policy archives files to one tier and then later relocates the files to another tier. Stubs that refer to relocated files are updated to refer to the new destination. This policy selects files based on file matching expressions. The multi-tier policy scans both stub files and normal files. The multi-tier stub policy scans stub files only.

The supported file server matrix for multi-tier policies is:

Primary tier	Secondary tier	Tertiary tier
Celerra VNX	Celerra VNX VNXe Windows Isilon NetApp Data Domain	Celerra VNX VNXe Windows Isilon NetApp Data Domain EMC Centera Atmos
Celerra VNX	Atmos	None
Celerra VNX	EMC Centera	None
NetApp	Celerra VNX VNXe Windows Isilon NetApp Data Domain	Celerra VNX VNXe Windows Isilon NetApp Data Domain EMC Centera Atmos
NetApp	Atmos	None
NetApp	EMC Centera	None

NOTE: Celerra, VNX, VNXe, Windows, Isilon, NetApp, and Data Domain are collectively referred to as NAS repositories.

Related links

Editing or adding a file server.
Editing or adding a NAS Repository.
Creating a NAS repository as an archive destination when creating a policy.

Policies

Use this page to see an overview of all defined policies, file matching expressions, and NAS destinations.

To access this page

Select the **Policies** tab.

Field Descriptions

- All currently defined policies are listed.

Field	Description
Policy Name	The policy name as defined on the Create Policy page.
Policy Type	The policy type may be archive, multi_tier, multi_tier_stub, delete_orphans, delete_stubs, or migrate_file.
Destinations	Destination for the policy, if applicable. For each policy, use the context menu to edit or delete. Policies are not tied to any source. The source is defined by using Create New Task.

- All global file matching Expressions that are currently defined are listed.

Field	Description
Name	The expression name as defined on Create File Matching Expression.
Expression	For each expression, use the context menu to edit, delete, move up, or move down.

- All currently defined NAS destinations are listed.

Field	Description
Name	The destination name as defined on Create NAS Destination .
Expression	For each destination, use the context menu to edit, delete, move up, or move down.

Quick Links

- Create new policy
- Create New File Matching Expression
- Create new NAS destination

Regular expressions

Use special characters and quantifiers to define attribute values for regular expressions, where `matches regex` is selected as the operator type. The implementation will vary, depending upon whether `filename` or `dirname` is selected as the file attribute.

Regular expressions are subject to the following conditions:

- Regular expressions are a means of matching substrings within strings. CTA uses a superset of the POSIX extended standard for regular expressions. POSIX defines certain syntax rules that are common across programming languages such as Perl. For more details, refer to the POSIX man page, or man 7 regex.
- Because of pathname expansion, do not use the same syntax as UNIX glob patterns understood by shells. For example, *.txt will match all files with a .txt extension in a UNIX shell, but is not a valid regular expression.
- When dealing with case-sensitivity, conventions of the particular protocol in question (CIFS or NFS) are followed. This means that for NFS, uppercase and lowercase are matched differently. However, case is not significant for CIFS.
- Regular expressions will return true if paths are matched partially or wholly. Use multiple expressions to narrow the match.

Special Characters

Special characters qualify digits or characters in a regular expression.

Character	Description
.	wildcard
$	end of string
^	beginning of string
\	escape character (precedes a special character)
\d	any digit (0-9)
\D	any nondigit
\w	any word character (a-z, A-Z, 0-9)
\W	any nonword character
\s	any whitespace (Examples: " ", newline, tab)
\S	any nonwhitespace

Quantifiers

Quantifiers indicate the quantity of the preceding characters.

Quantifier	Definition

*	0 or more
+	1 or more
?	0 or 1
{x}	x times
[xyz]	Character class that matches against a list or range: - [aeiou] matches any vowel - [a-z] matches any lowercase letter - [A-Z0-9] any uppercase letter or number

Examples

Examples of how to use attribute values to define expressions for `filename` are provided below.

Match all files:	**Attribute**
with a ".doc" extension	`\.doc$`
with either ".doc" or ".txt" extension	`(\.doc|\.txt)`
containing "name"	`name`
beginning with "New"	`^New`
ending in "old"	`old$`
with "abc" followed by single character and then "xyz"	`abc.xyz`
with at least five characters	`.{5}`
with only five characters	`^.{5}$`
with 5 a's followed by "config"	`a{5}config`
ending in a "2" follwed by ".dat"	`.*2\.dat$`
ending in a "20" to "29" followed by ".dat"	`.*2[0-9]\.dat$`

Examples of how to use attribute values to define expressions for `dirname` are provided below. Use the NFS format with forward slashes in the path name. Backward slashes are not supported.

Description	**File Matching Expression**
Match a single file under a specified directory /dir1/dir2.	`dirname matches singlefile "/dir1/dir2/file1.txt"`
Match all files only in /Home. For example, files in /Home will be matched, but not files in /Home/dir2.	`dirname == "/Home/"`

Match files in directory name dirX, but not in the subdirectories of dirX.	`dirname matches regex ".*/dirX/$"`
	Same as:
	`dirname matches regex "/dirX/$"`

Match files in directory name dirX and in the subdirectories of dirX.	`dirname matches regex ".*/dirX/.*"`
	Same as:
	`dirname matches regex "/dirX/.*"`

Match files in directories starting from dirX. For example, files in /dirX/dir1 and /dirX/dir1/dir2 will be matched, but not files in /dir1/dirX.	`dirname matches regex "^/dirX/.*"`
	Same as:
	`dirname matches regex "^/dirX/"`

Match all files in directories that start with the letter B, in paths starting at the export. Include subdirectories of directories that start with the letter B.	`dirname matches regex ".*/B.*/.*"`
	Same as:
	`dirname matches regex "/B.*/.*"`

NOTE: All directory paths start from share/export, not from the start directory of the task.

For help beyond the definitions and examples provided here, contact EMC Customer Support.

How Do I?

Create rules.
Create file matching expressions.

Scheduling tasks

CTA Tasks Overview

CTA tasks include:

- Archive
- Delete
- Scanning, backup, archive repository migration, file migration

Once a new task is scheduled, the task is listed on the Schedule Summary page. Before running an archive, delete, or auxiliary repository migration task run, run a quick or detailed simulation.

Archive Tasks

The archiving and recall architecture of CTA preserves file data. During archiving, file data and metadata is read by using CIFS or NFS and stored on a secondary storage tier. During recall, only the file data is retrieved from secondary storage, the metadata is not used. In this way, the CTA software supports multiprotocol data.

Note that the stub re-creation feature utilizes the file security and last modified timestamps in the database. The original file security and last modified timestamp from the time of archiving are restored when a stub is re-created. The saved security data is in a single protocol only and therefore re-created stubs could have a different security setting than the original file at the time of archiving. Stub re-creation is an activity that can only be performed by an administrator. The security of re-created stub files should be manually verified as part of this process.

Archive tasks can also operate on file lists imported from an external provider. To use the import files task, CTA configuration data is exchanged with the external provider before the file list is generated and imported.

Delete Tasks

Delete tasks simplify CTA administration by helping to manage the growth of secondary storage. The delete tasks are run against the CTA database to match files to be deleted. Orphan files or stubs that match the policy are automatically deleted.

Auxiliary Tasks

Auxiliary tasks include:

- Stub scanning — A critical task for orphan management. If this task is not automatically added to the schedule table, manually add it.
- Backup — Automatic backups to a NAS or EMC Centera destination can be scheduled as a daily or weekly task.
- Repository migration — Migrates all archived data from one repository to another storage tier. The migration can be to a NAS repository, to an EMC Centera, or an Atmos. Any stub files are updated to reference the new destination.
- File migration with policy — Migrates files between two primary file servers. Source and destination servers are defined with the task. The task evaluates all files based on the migrate_file policy. The resulting action is either migrate or don't migrate.

How Do I?

Run a Simulation
Archive tasks
Import archive file list
Delete orphans or stubs
Stub scanning
Back up and restore files on an CTA
Migrate a repository
File migration

Create, run, and complete a file migration task

The policy-based migration example described in the overview involves performing a scan of files on the source to determine if any fit the rules of the file migration policy. Files that are larger than 0 bytes in size, or all files, will be migrated to the destination primary server.

Before running any migration task, review the file migration prerequisites. Perform a simulation before running the migration task. In this example, the file migration task is created and a simulation is run. If results of the simulation show that the migration will work as expected, the file migration task is run.

If any orphan deletion, stub scanner, or archive tasks are scheduled, disable these tasks, otherwise data loss will occur.

To run a simulation for a file migration task

1. From the **Schedule** tab, click the link at the upper-right corner to **Schedule a new task**.
2. The **Create New Task** page appears. Specify the values for fields that apply to the file migration task:

Field		Description
Select Task Type		Select **Auxiliary.** Select **File Migration with policy Migrate files GR 0 bytes**.
Select Source	File Server	Select the name of a file server in your environment.
	Protocol	Select **CIFS**.
	Source path	Click **Browse** to locate a source path. Multiple source paths can be specified. A CTA CIFS user must have write permission for the source.
Select Destination	File Server	Select the name of a file server in your environment.
	Protocol	Select **CIFS**.
	Destination path	Click **Browse** to locate a destination path.
	Network bandwidth	(optional) Leave blank.
	SID translation file	(optional) Leave blank.
Select Start Time	Time	Select **Run Simulation Now**. Select **Detailed**.
Automatic Recursive File Migration		(optional) Leave blank.
	File Threshold Limit	(optional) Leave blank.

3. Click **Save Task**. The Schedule Summary page reappears with the new schedule listed. The most recently added task is at the bottom of the list. The Detailed Simulation is running.
4. To see the results of:

- Most recent simulation, select Sim Summary.
- Previously run simulations, select Simulation History from the **Status** column list.

The results of the detailed simulation provide a summary of the number of files to be migrated, bytes to be migrated, file ops failed, and a list of the source and destination paths for files that can be migrated.

5. When simulation results show that files on the source fit the rules in the **Migrate files GR 0 bytes** policy, the file migration is ready to be run.

To run a file migration task

1. On the **Schedule** tab, under List options, select Show schedules of type **File Migration**. All tasks defined for the specified schedule type appear with a **Source** and a **Policy Name**.
2. Find the row with the policy name **Migrate files GR 0 bytes**. This was the last task defined and should appear at the bottom of the list. To start the migration now, select **Run Now** from the **Status** column list.
3. To check the job progress, click **View Summary**. A Task Summary window appears. Files that are greater than 0 bytes should start migrating.
    4. Click **View file list** to access a text file with a list of source and destination paths for successfully migrated files. This list is only created if enabled with the CLI command:

        ```
        rffm setFileMigrationJournalingEnable Enable
        ```

    - Click **View log** to access a text file with a chronological record of the migration job including files that failed to migrate and the reason for the failure.

To complete the file migration task

Once the File Threshold Limit is reached, automatic file migration stops. To complete the migration, perform the following manual steps:

1. Disable stub scanning on the source, if applicable.
2. Restrict access to the source.
3. Execute the final incremental run.
4. Review migration logs to ensure that all data was successfully copied and there is no inconsistency between the data sets.
5. Switch client access to the new storage location. This is the new source.
6. Remove the old source from the network.
7. Enable stub scanning on the new source, if applicable.

More

File Migration
Run a simulation

Create a multi-tier archiving task

The multi-tier archiving example described in the overview involves performing a weekly scan of normal and stub files on the source to determine if any fit the rules of the multi-tier policy in Creating a policy. Normal files that are more than three months old will be archived to the NAS repository leaving stubs on the Celerra. Both normal files that are more than six months old and stubs for files that are more than six months old are examined for the offline file location. If the current destination is not the EMC Centera, the offline file moves to the EMC Centera and the stub offline path is updated to reference the new destination.

Before running any archiving task, a simulation should be performed. In this example, the multi-tier archiving task is created and a simulation is run.

To run a simulation for a multi-tier archiving task

1. From the **Schedule** tab, click the link at the upper-right corner to **Schedule a new task**.
2. The **Create New Task** page appears. Specify a value for each of the following fields:

Field		Description
Select Task Type		Select **Multi-tier with policy.** Select the policy **Multi tier old files**.
Select Source	File Server	Select the name of a file server in your environment. NOTE: The NetApp NearStore SnapLock and Celerra CWORM file systems are not supported as file archiving sources. They may only serve as file archiving destinations.
	Protocol	Select **CIFS**.
	Source path	Click **Browse** to locate a source path. Multiple source paths can be specified. A CTA CIFS user must have write permission for the source.
Select Archive Condition	Time	Select **Run Simulation Now**. Select **Detailed**.

3. Click **Save Task**. The Schedule Summary page reappears with the new schedule listed. The most recently added task is at the bottom of the list. The Detailed Simulation is running.
4. To see the results of:

- Most recent simulation, select View Summary.
- Previously run simulations, select Simulation History from the **Status** column list.

The results of the detailed simulation provide a summary of the number files to be archived, bytes to be archived, file ops failed, and a detailed file list.

5. When results show that files on the source fit the rules in the **Multi tier old files** policy, proceed to Run and view job progress.

More

Run a simulation

Interoperability with quotas

Celerra, VNX, and NetApp file servers support quotas that are used to limit the amount of data or number of files a user can write to that server. For example, a user with a 100 MB quota who creates a 100 MB file will not be able to write any additional data. However, after the file is archived, the user will be able to write an additional 100 MB of data.

Before archiving, consider the filesystem quota policy to determine whether quota space needs to be reclaimed after archiving or if usage needs to be restricted regardless of archiving activity.

NOTE: Only byte quotas are affected and file limits remain the same.

Run a Simulation

For archive, delete, repository migration, or file migration tasks, it is recommended that you run a simulation of the task before running the actual task. The results of the simulation will show the data that is to be archived, deleted, or moved.

- A simulation on an archive, delete, or file migration task evaluates the task policy against the source data set rather than the database. Real-time results provide an aggregated summary of the total files matched, total bytes potentially archived, deleted, or migrated, and an optional detailed file list. If the results of the simulation show files on the source that fit the rules in your policy, you are ready to run the task.
- A simulation on a repository migration task provides a list of files that reside on the source. If these confirm the files that you want to move, you are ready to run the task.

NOTE: A simulation is particularly important to run before performing any delete task. Because a delete task removes data, be sure to review the simulation results to confirm the data that is to be removed before running the task.

To run a detailed simulation when defining a task

1. On the **Schedule** tab select **Schedule a new task**. The **Create New Task** page appears.
2. Specify values for:

- Select Task Type — Select an archive, multi tier, multi tier stub, delete orphan, delete stub, repository migration, or file migration task type. For archive, multi tier, multi tier stub, or file migration, select the policy name.
- Select Source — Select a file server, protocol, and source path.
- Select Destination — For file migration, select a file server, protocol, and destination path.
- Select Archive Condition or Select Start Time — Select **Run Simulation now**. For a detailed summary, select **Detailed**.

3. Click **Save Task**. The Schedule Summary page reappears with the new schedule listed. The most recently added task is at the bottom of the list and the simulation is running.

To run a detailed simulation for a defined task

1. On the **Schedule** tab, under List options, select Show schedules of type: archive, multi tier, multi tier stub, delete orphan, delete stub, repository migration or file migration. All tasks defined for the specified type appear with a **Source** and a **Policy Name**.
2. In the row with the appropriate **Source** and the **Policy Name** of interest, select **Run Detailed Sim** from the **Status** column list. The simulation starts running.

Review the results

- Sim Summary in the **Sim Summary** column shows results of the most recent simulation.
- Simulation History from the **Status** column list, shows results of previously run simulations.

A quick simulation provides the number of files and number of bytes that will be affected. In addition to the number of files and bytes, a detailed simulation lists the names of the files that will be affected.

How Do I?

Create a multi-tier archiving task
Run and view job progress

Run and view job progress

Once results of the simulation confirm that the multi-tier policy will work as expected, you are ready to run the archiving task.

To run a multi-tier archiving task

1. On the **Schedule** tab, under List options, select Show schedules of type **Multi Tier**. All tasks defined for the specified schedule type appear with a **Source** and a **Policy Name**.
2. Find the row with the policy name **Multi tier old files**. This was the last task defined and should appear at the bottom of the list. To start the archive now, select **Run Now** from the **Status** column list.
3. To check the job progress, click **View Summary**. A Task Summary window appears. Files that are more than three months old should start archiving to the NAS repository.
4. Click **View file list** to access a text file with a list of successfully archived files.
5. Click **View log** to access a text file with a chronological record of the archive job.

NOTE: If the CTA needs to be shut down before an archiving task is completed, select **Stop** from the **Status** column list. Unpredictable results might occur if the archiving task is not stopped before shutdown.

When archiving multiple files larger than or equal to 1 GB to a NAS repository, error messages may appear that indicate that the disk space on the destination is full even if disk space is available. To work around this problem, run the archive job again. Files that were not archived during the first run are archived during the second run.

5. Once the multi-tier archive task has been run, see View archived files.

Related links

Create a multi-tier archiving task
Run a simulation

Import Archive File List

The import file task archives all files on a list provided by an external provider. A third-party software administrator controls the external provider. After defining the import file task, give the task name to the third-party software administrator. The third-party application will generate an XML file with the task name in the header.

After the third-party software administrator generates the archive file list, the CTA can import the file or the external provider can copy the file to the CTA. Once the CTA has the file list and has matched the name in the header to an import files task, archiving can begin as scheduled.

To schedule an import files task

1. From the **Schedule** tab, click the link at the upper-right corner to **Schedule a new task**.
2. The **Create New Task** page appears. Specify a value for each of the following fields:

Field		Description
Select Task Type		Check **Import.** For **Archive with policy**, **Multi tier with policy**, or **Multi tier stub with policy**, select a policy. Files are archived to the destination defined in the policy. NOTE: To archive all entries in the imported list, use a policy defined with a rule to archive all files with `size >= 1 bytes`.
Import File List	Task Name	Type a task name.
	Provider	Select a provider.
Select Archive Condition	Time	Select the daily, weekly, monthly, or future schedule for the archive task. Set the start time. The run will not occur until the XML file list is imported and validated.

3. Click **Save Task**. A warning box appears reminding the user to use the Import Files option on the Schedule list page. The Schedule Summary page reappears with the new schedule listed in green. The most recently run task is at the top of the list.

To transfer the file list to the CTA

Using the source server name, the provider name, and the task name configured on the CTA, the third-party software administrator generates a list of files for archiving. This XML file is transferred to the CTA in one of two ways:

- The CTA administrator imports the file using Import file.
- The third-party software administrator uses the name and password from the Provider properties page to log in to the CTA and copy the file to a staging directory. The XML file is deposited as:

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

where *PROVIDER* is the name from the Provider properties page and the *ID* is the unique ID in the XML file.

Import and archive logs

Every 30 seconds, the CTA scans for imported file lists and generates an import log. Before archiving begins, check the log for errors. View the log to confirm that all primary exports and shares are accessible. Logs for the XML file list import are also available outside of the GUI in:

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

The logs generated during archiving are available from the task summary. Files in the imported file list that are not found on the primary server are removed from the CTA database with a warning written to the archive log.

Delete orphans or stubs

The delete orphan and delete stub tasks work with the delete_orphans and delete_stubs policy types to manage the growth of secondary storage. By default, the scheduled time of these deletions is every Friday at 6 p.m.

Tasks set by the delete policy will delete data. Run a simulation before running any delete task. **Run a Simulation** describes two ways to run a detailed simulation.

To schedule a delete task

1. From the **Schedule** tab, click **Schedule a new task**. The **Create New Task** page appears.
2. Complete the **Create New Task** page as follows:

Field		Description
Select Task Type		Select **Delete** and select either: • **Delete orphan with policy**— Automatically deletes orphans on secondary storage that match the delete_orphan policy. • **Delete stub with policy**— Automatically deletes stubs on primary storage and archived files on the second tier that match the delete_stubs policy when the retention expires. The policies listed are defined using the **Create Policy** page.
Select Source NOTE: This field does not apply to **Delete orphan with policy**.	File Server	Select the name of a file server in your environment.
	Protocol	Select **NFS** or **CIFS**.
	Source path	Click **Browse** to locate a source path. Multiple source paths may be specified. A CTA CIFS user must have write permission for the source.
Select Start Time		Select the daily, weekly, monthly, or future schedule for the stub scanner. Set the start time. **Run every Friday at 18:00** is the default setting.

3. Click **Save Task**. The **Schedule Summary** page reappears with the new schedule listed in green. The most recently run task appears at the top of the list.

Stub Scanning

The stub scanning task is automatically added to the schedule table when an archiving task is executed on a share or export. The start directory is the top level share or export. By default, the scheduled time of the scan is every Friday at 6 p.m.

In cases where an export or share contains CTA stub files and the file server has never been used as an archiving source, the stub scanning task must be manually added.

Stub scanning is critical for orphan management. If the stub scanning task is not manually added, stubs that are not orphans may be mistakenly identified as orphans and could be deleted and result in data loss.

To schedule a stub scanning task

1. From the **Schedule** tab, click **Schedule a new task**. The **Create New Task** page appears.
2. Complete the **Create New Task** page as follows:

Field		Description
Select Task Type		Select **Auxiliary** and **Scan stubs**.
Select Source	File Server	Select the name of a file server in your environment.
	Protocol	Select **NFS** or **CIFS**.
	Source path	Click **Browse** to locate a source path. Multiple source paths may be specified. A CTA CIFS user must have write permission for the source.
Select Start Time		Select the daily, weekly, or monthly schedule. Set the start time. **Run every Friday at 18:00** is the default setting.

3. Click **Save Task**. The **Schedule Summary** page reappears with the new schedule listed in green. The most recently run task is at the top of the list.

Back up and restore files on a CTA

Regular backups of the CTA configuration and the critical database tables are highly recommended. Either the GUI or CLI scripts may be used to schedule backups and retrieve backup files. A CLI script is used to restore the backup file. The Getting Started Guide provides instructions on how to use the CLI scripts.

To configure automatic backups

1. Use Backup and Recovery Settings to configure the number and location of backup files for your system.

   A backup file containing critical CTA system configuration data is saved as a gzipped tar file (.tgz). The tar file is written to an EMC Centera or an NFS NAS destination.

   Up to 15 backup files are stored on the destination. The 15 backup filenames and their corresponding EMC Centera or NAS IDs are stored in a database file (DBBackup.out). This database file is stored in a standard location on the CTA and in a secondary disaster recovery location. The DBBackup.out file stored in the secondary location is used in the event that the file stored on the CTA is damaged or lost.

   To perform disaster recovery, the CTA uses database information from DBBackup.out to locate tar files on the destination and to reconstruct the CTA system configuration.

2. From the **Schedule** tab, click **Schedule a new task**. The **Create New Task** page appears.
3. Complete the **Create New Task** page as follows:

Field	Description
Select Task Type	Select **Auxiliary** and **Backup**.
Select Start Time	Select the weekly or monthly schedule. Set the start time. **Run every Friday at 18:00** is the default setting.

3. Click **Save Task**. The Schedule Summary page reappears with the new schedule listed in green. The most recently run task is at the top of the list.

To restore a backup

- Use Backup and Recovery Settings to select the backup file to be restored.
- Refer to the Getting Started Guide for instructions on using the **fmrestore** script to restore the backup file on the CTA.

Migrate a Repository

Repository migration moves all archived data from one repository to another storage tier. Migration can be to a NAS repository, to an EMC Centera, or to an Atmos.

Prerequisites

Before you begin, consider these conditions:

- If stub files are restored from backup or they are relocated on the primary filesystem and a migration task is planned:
  - Run a stub scanner task to scan the primary filesystems. This updates the CTA database so that the repository migration task can find stubs and update the offline path for those stub files.
  - Do not copy or restore stub files to multiple locations. This can cause stub files to become inaccessible after a repository migration.
- Do not migrate a primary filesystem and a repository at the same time. Migrate the primary filesystem first and then migrate the repository by following these steps:

1. Migrate the primary filesystem to a primary filesystem destination.
2. Remove the primary filesystem source from the network. The primary filesystem destination becomes the new primary filesystem.
3. Run a stub scanner task on the new primary filesystem.
4. Migrate the repository, if needed.

To schedule a repository migration task

1. From the **Schedule** tab, click **Schedule a new task**. The **Create New Task** page appears.
2. Complete the **Create New Task** page as follows:

Field		Description
Select Task Type		Select **Auxiliary** and **Repository Migration**.
Select Source	Source Repository	Select an EMC Centera, Atmos, or NAS repository in your environment.
	Protocol	Select **CIFS** or **NFS**. Source and destination must have the same protocol.
Select Destination		Select an EMC Centera, Atmos, or NAS repository in your environment. Multiple NAS repositories may be selected. To select multiple NAS repositories, press **Ctrl** while selecting the repository names.
Delay Period		Accept the default of seven days or type a new value. The delay is the period between the day that the stub files on the source are updated with the new destination path, and the day that the original files on the source repository are removed. NOTE: The delay period should be long enough to allow for the backup of stubs files updated with the new destination. If backups or snapshots are not needed, the delay period can be set to 0 days to remove the source repository data

		immediately after updating the source file.
Select Start Time		**Run in future** is the only selection. Set the start date and time when the migration will be run. NOTE: Once the migration task begins, any active archiving tasks that use the source as a target must be stopped and disabled.

3. Click **Save Task**. The Schedule Summary page reappears with the new schedule listed in green. The most recently run task is at the top of the list.

File Migration

The file migration task moves files between two primary file servers that are defined as the source and destination in the task definition. The task evaluates both normal and CTA stub files based on the migrate_file policy with a resulting action of either migrate or don't migrate.

Prerequisites

Before you begin, consider these conditions:

- For NFS, CTA has root access with read/write permissions on source and destination exports.
- CTA as an NFS client does not support NFSv4 on source or destination servers.
- CTA supports incremental NDMP backup for migrations.
  - Full read and write client access is supported during migration.
  - Scheduled downtime is required only for the last incremental run.
  - Destination must be offline until migration is complete, otherwise there is a risk of data loss or damage.
- For Celerra, Snapsure snapshots are enabled on the source.
- For NetApp, NDMP is enabled from the console by typing:

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

- CTA stub files will be recreated without recall on the destination during a migration. This applies to all migration destinations except VNXe. After running a file migration task, run a stub scanner task on the destination to update the CTA database with the new location of the stubs.

Disable all orphan deletion, stub scanner, and archive tasks on the source prior to migration. If any of these three types of tasks are not disabled, data loss will occur.

To schedule a file migration task

1. From the **Schedule** tab, click **Schedule a new task**. The **Create New Task** page appears.
2. Complete the **Create New Task** page as follows:

Field		Description
Select Task Type		Select **Auxiliary** and **File Migration with policy**. Select a migrate_file type policy to use.
Select Source	File server	Select a Celerra, VNX, or NetApp in your environment.
	Protocol	Select **CIFS** or **NFS**.
	Source path	Click **Browse** to locate a source path. Multiple source paths can be specified.
Select Destination	File server	Select a Celerra, VNX, or VNXe in your environment.
	Destination path	Click **Browse** to locate a destination path. Only a single destination path can be specified.
	Network bandwidth	(optional) Limit the bandwidth in mb/sec to control network and CPU activity.

	SID translation file	(optional) Specify the file to map local CIFS users from the source to the destination. This file is added to the CTA configuration on the File Migration Settings page.
Select Start Time		Select the start time for the file migration. **Run now** is the default setting.
Automatic Recursive File Migration	File Threshold Limit	(optional) Select this option and type a File Threshold Limit to automatically trigger an incremental migration when the previous run has finished. The runs will continue as long as the number of successfully migrated files is greater than the File Threshold Limit.

3. Click **Save Task**. The Schedule Summary page reappears with the new schedule listed in green. The most recently run task is at the top of the list.

Incremental copy

If a previous file migration task was interrupted or a set of baseline files were already copied, files to be migrated may exist on the destination. If files existing on the destination are different from the source files, CTA will overwrite them.

Take the source offline before the final incremental migration run. Create, run, and complete a file migration task details the steps required to complete a file migration task.

Supported platform matrix for file migration

The supported file server matrix for file migration is:

Source	Destination
Celerra VNX	Celerra VNX VNXe
NetApp	Celerra VNX VNXe

Page help
Create New Task

Use this page to create a new task.

To access this page

From the Schedule tab, select **Schedule a new task.**

Field Descriptions

Field		Description
Select Task Type		Select a task type: • Archive — Select the task: Archive with policy, Mult-tier with policy, or Multi-tier stub with policy, and then select the policy for which this archive task is being defined. Import is an option for any archive task. • Delete — Select the task: Delete orphan with policy, or Delete stub with policy, and then select the policy for which this delete task is being defined. • Auxiliary — Select the task: Scan stubs, Backup, Repository Migration , or File Migration with policy tasks. The policies listed are defined by using the Create Policy page.
Select Source	File Server	Select a source for the schedule from a list of file server names. • Archive with import selected does not use this option. • Repository Migration uses a different set of options. • File Migration uses the same source options, but different destination options. NOTE: The NetApp NearStore SnapLock and Celerra CWORM filesystems are not supported as file archiving sources. They may only serve as file archiving destinations.
	Protocol	Select **CIFS** or **NFS** to narrow the Source path selection. If archiving to a NAS destination, the protocol must match the protocol originally used to archive the selected source.
	Source path	Click **Browse** to locate a source path. Multiple source paths may be specified. A CTA CIFS user must have write permission for the source.
Select Start Time or Select Archive Conditions	Time	Select this option to schedule a specific or repeating time for a task to run. The start time is an estimation. The task will be activated on or after this time depending on currently running tasks. The start time selections are: • Run on Import — When selected, the archiving task starts automatically every time an XML file is imported to CTA. • Run now • Run Simulation now. For detailed simulation results. select **Detailed**. • Run every day at • Run on specific day, every # weeks at • Run on specific day, every # months at • Run in future

	Capacity used	Only applies for archive task types without **Import** selected. Select this option to force a policy to run when the source volume reaches a given capacity. Two values are specified for archive by capacity. • Warning threshold — Sends an alert when the warning threshold is reached. This optional value must be less than the archive threshold. • Archive threshold — Starts the specified policy when the source volume reaches the specified percentage of used disk space. An alert is also sent when the archive threshold is reached and the archive policy is started. If the policy fails to archived enough files to bring the source used disk space below the archive threshold, the policy will run again one hour after the last policy run finished.
Automatic Recursive File Migration	File Threshold Limit	Only applies for file migration task types. Select this option and type a File Threshold Limit to automatically trigger an incremental migration when the previous run has finished. The runs will continue as long as the number of successfully migrated files is greater than the File Threshold Limit. When the number of successfully migrated files is equal to or less than the limit, CTA will stop running the task automatically and send an alert that the limit has been reached. The final incremental migration must be run manually and offline, with client write access to the source volume blocked. If selected, the two start time options are **Run now** or **Run in future**.

When import is selected for an archive task, no source is selected and the options change.

Field		Description
Import File List	Task Name	Type a task name.
	Providers	Select a provider. The provider supplies a list of files to be archived to the CTA.

After creating a new task:

- To save the task, click **Save Task**.
- To cancel creating the task, click **Cancel**.

Archive tasks

There are three types of archive tasks:

- Archive with policy — Archives normal files that match the archive policy.
- Multi-tier with policy — Archives both normal and stub files that match the multi_tier policy.
- Multi-tier stub with policy — Updates stubs and relocates archived data that match the multi_tier_stub policy.

In addition, import is an option for any archive task.

Archive tasks check files on the source against the policy rules and perform the archive task per the archive conditions.

Before scheduling an archive task:

- Understand the archive conditions that you will apply to your task.
- Consider the filesystem quota policy as described in interoperability with quotas.

After defining an archive task, but before running the task, run a simulation.

Related links

Creating rules
Running a simulation

Import File List

Use this page to import a file list for an import archive file task. The file list is an XML file generated by a third-party software package using the server name of the Celerra, VNX, or NetApp defined on the CTA, the provider name from the Provider properties page, and the task name from the Import file task.

To access this page

1. On the **Schedule** tab, under List options, select Show schedules of type **Import File**. A list of Import file tasks appears.
2. Click the down arrow in the **Status** column of any task and select **Import file**.

Field Descriptions

The Task name is listed.

The Provider is listed.

Field	Description
Import File	Select an import file list to use for archiving. Click **Browse** to locate the XML file. The name and location of the file is provided by the third-party software administrator.

After locating the file:

- To import the file, click **Import**.
- To cancel importing the file, click **Cancel**.

Import Logs

Import logs apply to the import files tasks. One import files task may have several XML file lists associated with the task. Each row of the import log corresponds to a separate XML file for the same task.

To access this page

The **Schedule** tab lists all tasks by source.

1. Under **List Options**, select schedules of the type **Import Files**. The list refreshes to show all Import Files tasks.
2. Click the down arrow in the **Status** column of any source and select **Import Logs**.

Field Descriptions

Field	Description
Log	An ASCII text formatted log that can be used for diagnostic purposes. • To link to open the file in a browser window, click **View**. • To download the text file, click **Download**.
Import Date	Date that the XML file list was imported.
Generation Date	Date that the XML file list was originally generated.
Entries	Number of entries in the XML file.
Loaded	Number of files that are loaded to the database.
Duplicates	Number of files already in the database.
Errors	Number of validation or general errors.

Status	Status values are:
	• loaded — All files are loaded.
	• partial — Some files are loaded.
	• validation error — Validation failure. No attempt to load database. To be valid, CTA checks the XML file to confirm that it can reach all shares and exports, and that it has matching values for:
	  • provider name
	  • task name
	  • file server name
	• general error — Other failure. See import log for details.

List Options

At the bottom of the list, an option enables you to narrow or expand the range of dates for the items displayed. The list begins with the XML file that was most recently imported. Each time you reset the option, click **Refresh List** to redisplay the results.

Quick Links

- Back to main schedule page

Task Summary

Use this page to display the details of a particular task.

To access this page

On the **Schedule** tab, click **View Summary** for any schedule listed. The summary is for a run.

Field Descriptions

Run summary details vary depending on the task type: archive with policy, multi tier with policy, multi tier stub with policy, delete orphan with policy, delete stub with policy, scan stubs, backup, repository migration, and file migration.

Field	Description
Start time	Date and time that the task began.
End time	Date and time that the task ended.
Number of Directories Processed	Number of directories that were processed for the task.
Last Dir Processed	Name of last directory processed in the task.
Number of Files Processed	Number of files that were successfully processed.
Last File Processed	Name of the last file processed in the task.
Number of Files Archived	Number of files written to the destination in an archive task.
Number of Files Migrated	Number of files migrated in a repository or file migration task.
Bytes Archived	Number of bytes written to the destination in an archive task.
Bytes Migrated	Number of bytes migrated in a repository or file migration task.
File Ops Failed	Number of files processed unsuccessfully.
Number of Files Stubbed	Number of files stubbed at the time of the archive. If the Delay Stubbing Period is greater than zero days, this number may not match the Number of Files Archived.
Number of Stub Files Delayed	Number of files to be stubbed after the Delay Stubbing Period has elapsed. If the Delay Stubbing period is greater than zero days, this number should match the Number of Files Archived.
Detailed File List	Click **View file list** or **Download file list** to view or download a list of successfully archived or migrated files. For each file, the following data is presented in columnar format: • Filepath • File size • Last modified time • Last accessed time • Destination • Retention • Destination threshold (for NAS destinations only) The downloaded file list can be saved locally as a text file to import into an Excel spreadsheet for the purposes of customized reporting.
Log	Click **View log** or **Download log** to view or download a

	schedule log in ASCII text format that is typically used for diagnostic purposes.

Edit Task

Use this page to modify the archive conditions or start time for any task.

To access this page

The **Schedule** tab lists all tasks. Click the down arrow in the **Status** column of any task and select **Edit**.

Field Descriptions

The Task Type may not be modified.

Field		Description
Select Source NOTE: Values for the Source may be changed if the scheduled task has not run yet.	File Server	Select a source for the schedule from a list of file server names. NOTE: The NetApp NearStore SnapLock and Celerra CWORM file systems are not supported as file archiving sources. They may only serve as file archiving destinations.
	Protocol	Select **CIFS** or **NFS** to narrow the Source Path selection. If archiving to a NAS destination, the protocol must match the protocol originally used to archive the selected source.
	Source path	Click **Browse** to locate a source path. Multiple source paths may be specified. A CTA CIFS user must have write permission for the source.
Select Start Time or Select Archive Conditions	Time	Select this option to schedule a specific time or repeating time for a policy to run. The start time is an estimation. The task will be activated on or after this time depending on currently running tasks. The start time selections are: • Run on Import (for Import Files schedule only) • Run now • Run Simulation now. For detailed simulation results, select **Detailed**. • Run every day at • Run every specific day of the week • Run every specific day of the month • Run in future
	Capacity used	Select this option to force a policy to run when the source volume reaches a given capacity. Two values are specified for archive by capacity. • Warning threshold — Sends an alert when the warning threshold is reached. This optional value must be less than the archive threshold. • Archive threshold — Starts the specified policy when the source volume reaches the specified

		percentage of used disk space. An alert is also sent when the archive threshold is reached and the archive policy is started. If the policy fails to archived enough files to bring the source used disk space below the archive threshold, the policy will run again one hour after the last policy run finished.
Automatic Recursive File Migration	File Threshold Limit	Only applies for file migration task types. Select this option and type a File Threshold Limit to automatically trigger an incremental migration when the previous run has finished. The runs will continue as long as the number of successfully migrated files is greater than the File Threshold Limit. When the number of successfully migrated files is equal to or less than the limit, CTA will stop running the task automatically and send an alert that the limit has been reached. The final incremental migration must be run manually and offline, with client write access to the source volume blocked. If selected, the two start time options are **Run now** or **Run in future**.

After creating a new task:

- To save the task, click **Save Task**.
- To cancel editing the task, click **Cancel**.

Schedule History

Use this page to see the schedule history of all the run summaries for a particular task. Runs are listed in chronological order starting with the most recently run task at the top of the table.

To access this page

The **Schedule** tab lists all tasks by source. Click the down arrow in the **Status** column of any source and select **History**.

Field Descriptions

Field	Description
Log	An ASCII text formatted log that can be used for diagnostic purposes. • To link to open the file in a browser window, click **View**. • To download the text file, click **Download**.
Archived File List	For archive tasks. This list of files archived during the task is in ASCII text format and may be used for diagnostic purposes. • To link to open the file in a browser window, click **View** . • To download the text file, click **Download.**
Files Archived	For archive tasks. The total number of files that were archived during a particular run.
Total Bytes Archived	For archive tasks. The total number of bytes that were archived during a particular run.
Files Migrated	For file migration tasks. The total number of files that were migrated during a particular run.
Total Bytes Migrated	For file migration tasks. The total number of bytes that were migrated during a particular run.
Deleted File List	For delete tasks. This list of files deleted during the task is in ASCII text format and may be used for diagnostic purposes. • To link to open the file in a browser window, click **View** . • To download the text file, click **Download.**
Files Deleted	For delete tasks. The total number of files that were deleted during a particular run.
Total Bytes Deleted	For delete tasks. The total number of bytes that were deleted during a particular run.
File Operations Failed	For archive and delete tasks. Number of files that matched the criteria but failed to archive. Possible reasons for failure include: permissions problems, and network communication problems.
Start Time	When the task began.
End Time	When the task ended.
Status	For file migration tasks. The status of the migration may be Success, Failed, Still in progress.
Incremental Base Date	For file migration tasks. The date and time when the last run took its migration snapshot. For incremental runs, only files changed since this date and time are considered for migration. If the initial file migration runs to completion, **Full Migration** appears.
Directories Processed	For archive, multi-tier, multi-tier stub, repository migration, and import files tasks. Number of the directories scanned. Knowing how many directories are on the filer is useful in determining the length of a scan.
Last Directory	For archive, repository migration, and import files tasks. Last directory

Processed	scanned. Knowing where the scan stopped is useful if a scan terminated unexpectedly.
Files Processed	For archive, repository migration, and import files tasks. Number of files on the source.
Last File Processed	For archive, delete, repository migration, and import files tasks. The last file processed.

List Options

At the bottom of the schedule history, an option enables you to narrow or expand the range of dates for the items displayed. The schedule history begins with the most recent task performed which is used to find past results. Each time you reset the option, click **Refresh List** to redisplay the results.

Quick Links

- Back to main schedule page

Simulation History

The simulation history applies to the simulation tasks. Tasks are listed in chronological order and start with the most recently run task.

To access this page

The **Schedule** tab lists all tasks. Click the down arrow in the **Status** column of any archive task and select **Simulation History**.

Field Descriptions

Field	Description
Log	An ASCII text formatted log that can be used for diagnostic purposes. • To link to open the file in a browser window, click **View**. • To download the text file, click **Download**.
Archived File List	For archive tasks. This list of files archived during the simulation is in ASCII text format and may be used for diagnostic purposes. • To link to open the file in a browser window, click **View**. • To download the text file, click **Download.**
Files Archived	For archive tasks. The total number of files that would be archived per the simulation run.
Total Bytes Archived	For archive tasks. The total number of bytes that would be archived per the simulation run.
Files Migrated	For file migration tasks. The total number of files that would be migrated per the simulation run.
Total Bytes Migrated	For file migration tasks. The total number of bytes that would be migrated per the simulation run.
Deleted File List	For delete tasks. This list of files that would be deleted is in ASCII text format and may be used for diagnostic purposes. • To link to open the file in a browser window, click **View** . • To download the text file, click **Download.**
Files Deleted	For delete tasks. The total number of files that would be deleted.
Total Bytes Deleted	For delete tasks. The total number of bytes that would be deleted.
File Operations Failed	For archive and delete tasks. Number of files that matched the criteria but would fail to archive. Possible reasons for failure could include: permissions problems, network communication problems, and so forth.
Start Time	When the simulation run began.
End Time	When the simulation run ended.
Status	For file migration tasks. The status of the simulation may be Success, Failed, Still in progress.
Incremental Base Date	For file migration tasks. The date and time when the last simulation run took its migration snapshot. For incremental runs, only files changed since this date and time are considered for migration. If the initial simulation migration runs to completion, **Full Migration** appears.
Directories Processed	For archive, repository migration, and import files tasks. Number of the directories scanned. Knowing how many directories are on the filer is useful in determining the length of a scan.
Last Directory Processed	Last directory scanned. Knowing where the scan stopped is useful if a scan terminated unexpectedly.
Files Processed	For archive, repository migration, and import files tasks. Number of files

	on the source.
Last File Processed	For archive, delete, repository migration, and import files tasks. Last file processed.

List Options

At the bottom of the simulation history, an option enables you to narrow or expand the range of dates for the items displayed. Since the schedule history begins with the most recent task performed, this option is useful in finding past results. Each time you reset the option, click **Refresh List** to redisplay the results.

Quick Links

- Back to main schedule page

Schedule Summary

Use this page to display a list of all tasks that are scheduled to, or have already run.

To access this page

Select the **Schedule** tab.

Schedule Type

In the schedule list, the columns change depending upon the **List Options** selected. When the schedule type is changed, the list automatically refreshes. If the status setting for enabled, disabled, or all is changed, you must click **Refresh List** to redisplay the results.

- Archive, Multi-Tier, Multi-Tier Stub — The archive tasks use policy rules to select the files to archive. Files are copied to the destination and replaced with stub files on the source. When Archive, Multi-Tier, or Multi-Tier Stub is selected, Source, Estimated Start Time, Current Usage %, Archive Capacity % Threshold, Policy Name, Summary, Sim Summary, and Status columns are displayed.
- Stub Scanner — This type of task scans the source looking for stub files for the purpose of orphan management. When it finds a stub file, it updates the database. When Stub Scanner is selected, Source, Estimated Start Time, Summary, and Status columns are displayed.
- Delete Orphan, Delete Stub— The delete tasks use policy rules to select the files to delete. When Delete Orphan or Delete Stub is selected, Source, Estimated Start Time, Policy Name, Summary, Sim Summary, and Status columns are displayed.
- Backup — This type of task runs a backup as configured on the **Backup and recovery settings** page. When Backup is selected, Estimated Start Time, Summary, and Status columns are displayed.
- Repository Migration — This type of task migrates files from the source to the destination at some time in the future. When Repository Migration is selected, Source, Estimated Start Time, Summary, Sim Summary, and Status columns are displayed.
- File Migration — This type of task migrates files from a source primary server to a destination primary server. When File Migration is selected, Source, Estimated Start Time, Policy Name, Summary, Sim Summary, and Status columns are displayed.
- Import Files— This type of schedule is associated with an archive task. When Import Files is selected, Task Name, Provider, Estimated Start Time, Policy Name, Summary, Sim Summary, and Status columns are displayed.

Schedule List

Field	Description
Task Name	For Import Files only. The name of the archive task used for import archive file list.
Provider	For Import Files only. The name of the provider used for import archive file list.
Source	The source to which the schedule is tied.
Estimated Start Time	When the schedule will or has run. For the stub scanner task, the time is

	fixed.
Current Usage %	The current capacity usage for the source.
Archive Capacity % Threshold	The capacity level required to trigger the archive task. If blank, this optional value was not specified.
Policy Name	The policy name applies to the archive, delete, file migration, and import files tasks. All defined policies are listed on Policies.
Summary	Click View Summary to display a summary of the most current run. The Schedule History is a compilation of all run summaries. • If the task failed to archive, delete, or migrate any files, the log file is empty and the task summary displays an empty screen. • If the log file is missing, the task summary displays the message, "Log file not found."
Sim Summary	Click Sim Summary to display a summary of the most current simulation run. The Simulation History is a compilation of all run summaries. • If the task failed to archive, delete, or migrate any files, the log file is empty and the task summary displays an empty screen. • If the log file is missing, the task summary displays the message, "Log file not found."
Status	For each task on the list, select any of the following actions: • Edit — Provides a means to edit the task. • Run now — Runs the task now, irrespective of the schedule. Not applicable for backup tasks. • History — Displays the Schedule History of all the run summaries for a particular task. • Run Quick Sim — Displays the total number of files and the total number of bytes affected. A list of files is not provided. Not applicable for stub scanner or backup task types. • Run Detailed Sim — Displays the total number of files, the total number of bytes, and the names of the files that will be affected. Not applicable for stub scanner or backup task types. • Simulation History — Not applicable for stub scanner or backup task types. Displays the Simulation History. • Enable — Places the task in the run schedule. By default, the status for all tasks is enable. • Disable — Removes the task from the run schedule without deleting the task definition. Not applicable for stub scanner or backup task types. Stub scanner or backup tasks may only be deleted. • Stop — Stops running a task in progress. • Delete — Deletes the task definition. This permanently removes the task from the run schedule. • Import Files — Imports a file list for the import files task. • Import Logs — Displays the logs for the import files task.

NOTES:

- Any currently running tasks appear in green.
- The currently running or most recently completed tasks appear at the top of the list.
- Completed archive tasks may be deleted from the list and the referenced policy may also be deleted.

- If an archive task is stopped by changing the status to **Stop**, the task will stop after a few seconds, but the task color will remain green. Use View Log on the View Summary to verify the task status.

Quick Links

- Schedule a new task

Simulation Summary

Use this page to display the details from a simulation run of a particular task.

To access this page

On the **Schedule** tab, click **Sim Summary** for any of the source directories listed. The summary is for a simulation.

Field Descriptions

Simulation summary details also vary depending on the task type and only apply to: archive with policy, multi tier with policy, multi tier stub with policy, delete orphan with policy, delete stub with policy, repository migration, and file migration with policy.

Field	Description
Start time	Date and time that the task began.
End time	Date and time that the task ended.
Number of Directories Processed	Number of directories that were processed for the task.
Last Dir Processed	Name of last directory processed in the task.
Number of Files Processed	Number of files that were successfully processed.
Last File Processed	Name of the last file processed in the task.
Number of Files to be Archived	Number of files to be written to the destination in an archive task.
Number of Files to be Migrated	Number of files to be migrated in a repository or file migration task.
Bytes to be Archived	Number of bytes to be written to the destination in an archive task.
Bytes to be Migrated	Number of bytes to be migrated in a repository or file migration task.
File Ops Failed	Number of files processed unsuccessfully.
Detailed File List	Click **View file list** or **Download file list** to view or download a list of successfully archived files. For each file, the following data is presented in columnar format: • Filepath • File size • Last modified time • Last accessed time • Destination • Retention • Destination threshold (for NAS destinations only) The downloaded file list can be saved locally as a text file to import into an Excel spreadsheet for the purposes of customized reporting.
Log	Click **View log** or **Download log** to view or download a schedule log in ASCII text format that is typically used for

	diagnostic purposes.

Archived Files

View the archived report

After running an archive task, you can view an archive report for a file server.

To view the archived report

1. Click the **Archived Files** tab. The Archived Report page appears.
2. Under Report options in the **Archived Report** box, a file server name is selected. To view the archive report for a different file server, select that file server name. The display automatically refreshes.

An archived report displays:

- Archive history — Bar chart of total megabytes archived per server, or per server group, over a period of months.
- File size — Pie chart summary of archived data based on file size.
- File type — Pie chart summary of archived data based on file extension.
- Summary table — Summary of values listed by source. Depending on the report option settings, the source may be filtered by file server or file server group.

To change display units, use the Archived Report Settings page.

View archived files

After running an archive task, you can view a summary of all the files that have been archived. Use this summary to administer stub and orphan files.

To list archived files

1. On the **Archived Files** tab, click **View archived files list**.
2. Specify the list options to define the results to display:

Field	Description
Source	Select the name of the source file server.
Protocol	Select **CIFS**.
Path	Type the source path click **Browse** to locate the source path. NOTE: The path accepts the wildcard characters * or ?. For example, to find files archived to the share, myshare, and the directory, dir1, use the path `\myshare\dir1\*`. To find all files that end with .doc in the same location, use the path `\myshare\dir\*.doc`.
Archived from	Leave this option unspecified.
Show orphan files	Leave this option unspecified.

1. Click **Show orphan files**. With the default setting of **30 or more days ago**, the list displays any orphan files detected by the stub scanner that are 30 or more days old.

   To reduce the number of orphan files listed, change the **Show orphan files** setting to a value greater than 30 or more.
2. Click **Get List**. A list of archived files appears with the following information displayed:

Field	Description
File Path	The path relative to the browse path, excluding the share or export name.
Destination	The destination name.
Retention Expiration	The date by which file retention expires. After this date, a file may be deleted from the secondary storage.
Modify Time	The last time that a file was modified before archiving.
Archived Time	The time that a file was archived.

Page help
Archived File List

Use this page to display the archived files for any file servers available from on the **Archived Report** page.

To access this page

From the **Archived Files** tab, select **View archived files list.**

List options

Select list options to display a list of archive results for one particular source.

Field	Description
Source	Specify the host name.
Protocol	Select **NFS** or **CIFS** to narrow the source path selection.
Path	Export or share of the files. Click **Browse** to search for the path on the source. If **Browse** is used, either the NFS or CIFS protocol must be selected. The path accepts the wildcard characters * or ?. For example, to find files archived to the share, `myshare`, and the directory, `dir1`, use the path `\myshare\dir1\*`. To find all files that end with `.doc` in the same location, use the path `\myshare\dir\*.doc`.
Archived from	If values are provided, the list of results is limited to files archived during a particular date and time range.
Show orphan files	• With show orphan files is not selected, the archived files are displayed when **Get List** is selected. The default setting for this option is not selected. • With show orphan files selected, stub files that no longer exist on the source and left orphan files on the destination are displayed when **Get List** is selected. These files were detected by the stub scanner during the period specified. The default period is 30 or more days ago.

To display the list, dlick **Get List**. If no filtering parameters are selected, a complete list of archived files appears.

List

A list of files that meet the filtering criteria appears.

Field	Description
File Path	Path listed is relative to the browse path. It does not include the share or export name.
Destination	Lists the destination name.

Retention Expiration	Lists the date by which file retention expires. After this date, a file could be deleted from the secondary storage. If the Show orphan files option was selected, the Delete Orphan button becomes active. Any file for which the retention time has expired could be deleted by clicking **Delete Orphan**.
Modify Time	The same file on the source could have been archived several times. To help identify the correct file version, the last time that a file was modified before archiving is listed.
Archived Time	The time that a file was archived.

Buttons control files in the list:

Button	Description
Select All	Checks all files displayed in the list.
Clear All	Clears checkboxes for all files displayed in the list.
Recover Stub	Select the checkbox of any file to restore from an archive on the EMC Centera or NAS destination. Then click **Recover Stub** to recover the stub to the source. This option is useful in cases where a file might have been mistakenly deleted. If the stub file still exists on the source, the current date is appended to end of the stub file name. The Show orphan files checkbox can be used to find files that do not have a stub file on the file system. The following conditions apply to stub recovery: • To recover stubs, the full parent directory structure beginning with the file path field must exist on the primary server. • If the stub is being recovered from a NAS destination and if the Retention Period is set to a value greater than zero, it is recovered with read permission only. Select **Recover stub files with write permission** to recover all selected stub files with read-write permission. Edit Policy shows the **Retention Period** setting. • For multiprotocol stub recovery, the recovered metadata is restored with the archive protocol. For example, if a file that is both CIFS and NFS is archived with the CIFS protocol, when the stub is recovered to the source it will be of the CIFS protocol.
Delete Orphan Only active if Show orphan files option is selected.	An archived file is considered an orphan when its stub on the source does not correspond to a file on the destination. A special stub scanner task runs in the background to identify orphan files. Click this button to permanently remove orphan files from the destination for the selected stubs. NOTE: An archived file can only be removed if its Retention Period has expired. For an EMC Centera, a file can only be removed if the retention period has expired on an EMC Centera with compliance mode enabled. Only one user at a time can browse the Archived File List and perform orphan management operations. Database queries scan the largest table of archived files and are very CPU intensive.

Quick Links

- View Stub Recovery Log

- View Orphan Deletion Log

These logs are in ASCII text format and provide detailed information that you can use for diagnostic purposes.

Archived Report

Use this page to display archive information for a specific file server. Reported values are configured with Archived Report Settings.

To access this page

Select the **Archived Files** tab.

Report options

NOTE: The archived statistics presented in the Archived Report are global and not related to the current server overview.

Use these options to control the table display.

Field	Description
Filter by	Select the means by which the results are reported. • File Server — Hostname. If selected, the Source column lists the exports/shares of this file server. • File Server Group — To use this option, the file server group must be first defined with File Server Groups. If selected, the Source column lists the file servers that belong to this group.
Available reports	Values for the most recently saved archive results are displayed. For values of earlier archive tasks, select the name of a different archive report from the list.
Current Usage %	Left unselected, values for the most recently saved archive results are displayed. If selected, values for a current report are generated and displayed. NOTE: For new results, there may be a time delay to process the data.

To update the list, click **Refresh**.

Archive history

A bar chart shows the total megabytes archived for the file server each month over the course of months specified.

Pie charts

- File Size — Percentage of files of a particular file size.
- File Type — Percentage of files with a particular file extension.

Archive summary

A table provides a breakdown of the values reported by source. Groups are as specified for File Size.

Field	Description
Source	Lists exports/shares or servers depending on the Report options selected. If filtering by File Server, exports/shares are listed. If filtering by File Server Group, servers are listed.
Total Archived	Total number of megabytes archived for each source.
# of Files Archived	Total number of files archived for each source.
% for Group A	Percentage of files with sizes in the range for Group A.
% for Group B	Percentage of files with sizes in the range for Group B.
% for Group C	Percentage of files with sizes in the range for Group C.
% for Group D	Percentage of files with sizes in the range for Group D.
% for Other	Percentage of files with sizes beyond the range for Group D.

Quick Links

- View archived files list
- View reports
- Archived Report Settings

Report

Use this page to view detailed reports for any file servers available on the **Archived Report** page.

To access this page

From the **Archived Files** tab, select **View reports.**

Field Descriptions

The table is organized by file server with the shares or exports for each file server grouped together. For each share or export, the following data is provided:

Field	Description
Source	The servers used as archive sources within the period specified.
Data Archived	The aggregate amount of data that was archived for each share or export within the period specified.
Currently Used Space	Space used on the file system or volume since the last archive.
Space Saved (%)	If values are provided, the list of results is limited to files archived during a particular date and time range.
Detailed Reports	Select a report for a particular date and time. Click **View Report** to display the Detailed Report. • If no files were archived, the log file is empty and the Detailed Report displays an empty screen. • If the log file is missing the Detailed Report displays the message, "Log file not found."

List Options

At the bottom of the report, a list of options enables you to edit the reporting period for the data displayed. The default period is one month. Data may also be limited to shares or exports on a single file server. Each time you reset any options, click **Refresh List** to redisplay the results.

Quick Links

- Back to Main Schedule Page

Tools and Diagnostics

Files

Various files are used for configuration, licensing, scripts, and other functions. Some of the most commonly used files are:

File	Description
/opt/rainfinity/filemanagement/conf/rain.xml	CTA configuration file.
/opt/rainfinity/filemanagement/conf/keyfile	Authentication key used to communicate with the daemon.
/opt/rainfinity/filemanagement/conf/nodeid	The ID of this node (must match a NodeInfo entry in rain.xml).
/opt/rainfinity/filemanagement/conf/remotekeyfile*	Keys used for communicating between nodes in a distributed deployment.

In addition, commonly used files are stored together in a few directories:

Directory	Description
/var/rainfinity	This directory stores runtime state information.
/opt/rainfinity/bin /opt/rainfinity/filemanagement/bin	These directories contain the default executables and scripts that the CTA uses.
/opt/rainfinity/filemanagement/doc	This directory contains PDFs of all the man pages.

Logging

The CTA logs entries to several different locations. The information written to log files can be used by EMC Customer Support to diagnose problems that may be encountered.

Main log file

The main log file is:

```
/var/log/rainfinity/filemanagement/system.log
```

The main log file is rotated by the system, and up to four backup copies are saved. The log file size has an upper limit of 50 megabytes.

Log file entries follow a particular format:

```
<date and time> <level> log message
```

The level is one of the following:

Level	Description
EMERG	The most severe conditions that prevent the continuation of operation, such as an immediate system shutdown.
ALERT	System conditions that require immediate attention, such as corrupted system database, insufficient disk space, or lack of file descriptors.
CRITICAL	Mostly serious and nonrecoverable system/application malfunctions, such as failing hardware (hard device errors) or software.
ERROR	Mostly correctable errors, for example errors other than hard device errors. Continuation of operation is possible. Error conditions are automatically recoverable.
WARNING	Warning messages that require attention.
NOTICE	Notices that require attention at a later time. Non-error conditions that might require special handling.
INFO	Informational messages.
DEBUG	Messages for support use in debugging.

In practice, the CRITICAL, ERROR, WARNING, INFO, and DEBUG levels are most commonly used.

Diagnostic log files

Low-level diagnostics for the CTA service produce the log file:

```
/var/log/rainfinity/filemanagement/debug.log
```

Log and error messages that occur when files are being recalled from different destinations are written to log files specific to the platform:

- Recalls from EMC Centera, when archived using Celerra or VNX

```
/var/log/rainfinity/filemanagement/recall/CCD.log
```

- Recalls from Atmos, when archived using Celerra or VNX

```
/var/log/rainfinity/filemanagement/recall/ACD.log
```

- Recalls from NetApp

```
/var/log/rainfinity/filemanagement/recall/FCD.log
```

CTA logs

A log for any task on the Schedule Summary is available from the Task Summary.

Two types of logs are available from the Archived File List.

- Stub Recovery Log

```
/var/log/rainfinity/filemanagement/stubrecovery.log
```

- Orphan Deletion Log

```
/var/log/rainfinity/filemanagement/deleteorphan.log
```

Glossary

A

API: Application programming interface. A source code interface provided bythe computer application to support requests for services.

archiving: Process that walks the share/export and performs policy-based filearchiving.

Atmos Callback Service: CTA callback service to support FileMover recall from Atmos.

C

Celerra Callback Daemon (CCD): CTA callback service to support FileMover recall from the EMC Centera.

Celerra FileMover: HSM implementation used to support offline files on the Celerra.

Centera API: API used to write and read files from the EMC Centera.

Centera content address: Unique key to the saved file on the EMC Centera.

D

DHSM: Distributed Hierarchical Storage Management is the former name forCelerra FileMover.

F

File version: Multiple copies on secondary storage of the same file or path.

FileMover API: API over HTTP exposed by Celerra Data Mover to create stub files.

Fpolicy Callback Daemon (FCD): File Management callback daemon used to support NetApp Fpolicy recallfrom all secondary storage.

Fpolicy server: NetApp Fpolicy server. Provides notification when client accessesstub files.

FQDN: Fully Qualified Domain Name. Used with the Celerra Callback DNS entry.

H

HSM: Hardware security module.

L

LDAP: Lightweight Directory Access Protocol

M

MB: Megabyte, 106 bytes.

N

NAS: Network attached storage.

O

orphan file: Files on the secondary storage with no reference to the primary storage.

P

primary storage: NAS device that exports CIFS or NFS volumes.

R

RADIUS: Remote Authentication Dial In User Service

retention period: Number of days from time of archiving that a file can not be deleted.

S

secondary storage: Data storage that is a backup to primary storage.

SNMP: Simple Network Management Protocol

STIG: Security Tests Implementation Guide

stub file/offline files: Files that appear as normal files on the primary storage but point todata content stored on the secondary storage.

T

TACACS+: Terminal Access Controller Access-Control System Plus

V

VMotion: VMware VMotion technology is virtual machine mobility unique to VMware.

Index

D

E

F