



**NetApp® SANtricity® Cloud Connector 1.0**

# User Guide

June 2017 | 215-12357\_A0  
[doccomments@netapp.com](mailto:doccomments@netapp.com)



# Table of Contents

|                                                                                                |           |
|------------------------------------------------------------------------------------------------|-----------|
| <b>About this Guide</b>                                                                        | <b>1</b>  |
| Overview of the SANtricity Cloud Connector                                                     | 1         |
| Overview of the REST API                                                                       | 1         |
| <b>General Concepts</b>                                                                        | <b>3</b>  |
| Sources and Targets                                                                            | 3         |
| Volume Data Consistency                                                                        | 3         |
| REST API                                                                                       | 3         |
| Encryption                                                                                     | 3         |
| Terminology                                                                                    | 4         |
| <b>System Setup</b>                                                                            | <b>5</b>  |
| Host Hardware Requirements                                                                     | 5         |
| Compatible Storage Arrays and Controller Firmware                                              | 5         |
| Compatible Operating Systems                                                                   | 5         |
| Installation                                                                                   | 5         |
| Installing the SANtricity Cloud Connector on a Linux Operating System in graphical mode        | 5         |
| Installing the SANtricity Cloud Connector on a Linux Operating System in console mode          | 11        |
| <b>Initial Configuration</b>                                                                   | <b>15</b> |
| Source and Target Root File Paths                                                              | 15        |
| SANtricity Cloud Connector Credential Login Management                                         | 16        |
| Proxy Server Login Management                                                                  | 16        |
| Configuring the config.json file                                                               | 17        |
| Webserver Port Number                                                                          | 17        |
| State Storage Type                                                                             | 17        |
| S3 Target                                                                                      | 18        |
| Web Services Proxy Connection and Login Information                                            | 18        |
| Job Engine Configuration                                                                       | 18        |
| <b>Backup and Restore Jobs</b>                                                                 | <b>20</b> |
| Backup Jobs                                                                                    | 20        |
| Mapping Snapshots to the SANtricity Cloud Connector Host for Backup                            | 20        |
| Configuring and processing a Block-Based Backup Job to an S3 Target                            | 20        |
| Restore Jobs                                                                                   | 21        |
| Mapping volumes to the SANtricity Cloud Connector Host for Restore                             | 21        |
| Configuring and processing a Block-Based Restore Job to an E-Series target                     | 21        |
| <b>Uninstalling the SANtricity Cloud Connector</b>                                             | <b>23</b> |
| Uninstalling the SANtricity Cloud Connector on a Linux Operating System through graphical mode | 23        |
| Uninstalling the SANtricity Cloud Connector on a Linux Operating System through console mode   | 23        |



## About this Guide


This guide describes general concepts, setup, configuration, and jobs associated with the SANtricity Cloud Connector.

### Overview of the SANtricity Cloud Connector

The SANtricity Cloud Connector is a host-based Linux application that allows you to perform full block-based backup and restore of E-Series volumes to Amazon Simple Storage Service (S3). Available for installation on RedHat and SUSE Linux platforms, the SANtricity Cloud Connector is a packaged solution (.bin file). Once installed, you can configure the SANtricity Cloud Connector to perform backup and restore jobs for E-Series volumes to your existing Amazon S3 account. All jobs performed through the SANtricity Cloud Connector utilize REST-based APIs. Data utilized by the SANtricity Cloud Connector application is encrypted using 128-bit Advanced Encryption Standard (AES) and hashed through SHA-256. Configuration data for the SANtricity Cloud Connector application is stored in a single `config.json` file.

### Overview of the REST API

REST API is an application programming interface designed for experienced developers. Actions performed through the REST API are applied upon execution and without user prompts or confirmation dialogs. The REST API is URL-based, and the accompanying API documentation is completely interactive. Each URL contains a description of the corresponding operation and the ability to perform the action directly through the API documentation. The API documentation is accessible by navigating to </docs/index.html> through a browser.



Each URL endpoint presented within the API documentation has a corresponding POST, DELETE, or GET option. These URL endpoint options, or more properly known as HTTP verbs, are the actions available to the developer through the API documentation.

Data within the REST API is encoded through JavaScript Object Notation (JSON). The structured JSON data from the REST API can be easily parsed by programming languages (C, C++, cURL, Java, Python, Perl, etc.) JSON is a simple key-value pair based encoding with support for list and subject objects.

Objects start and end with curly braces (i.e., { }) while lists start and end with brackets (i.e., [ ]). JSON understands values that are strings, numbers, and Booleans. Numbers are floating point values. The API documentation provides a JSON template for each applicable URL operation, allowing the developer to simply enter parameters under a properly formatted JSON command.

# General Concepts

## Sources and Targets

Host base volume data (i.e., source) backs up to Amazon Simple Storage Service S3 (i.e., the target). Volume data is comprised of snapshots of the volume mapped to the SANtricity Cloud Connector host. The backup process reads data from the volume and saves it to S3. Data is copied to S3 with AES 128-bit encryption.

The backup process is comprised of the three following tasks:

- Volume IO bound tasks
- CPU bound tasks
- Network IO tasks

Volume IO bound tasks read data from the volume. CPU bound tasks perform all CPU-based processes. Network IO tasks write data to S3.

The SANtricity Cloud Connector backup and restore operations must occur on a host that is IO-attached to the E-Series array containing the volumes to be backed up and/or restored. The IO-attached host can be one dedicated for SANtricity Cloud Connector use or one already being used to access the storage for other applications.

Any IO-attached host running the SANtricity Cloud Connector must also run DM-MP multipath and have the multipath-tools package installed. The SANtricity Cloud Connector discovery process relies on the multipath tools package for discovery and recognition of the volume(s) to backup or restore.

## Volume Data Consistency

There is an indirect association between data chunks and the record of which data is included in a volume backup. The backup process loads data chunks to S3 and tracks which chunks are needed for the new volume backup. The current backup is recorded in S3 in a volume record. Once all data chunks are backed up to S3, the record of volume data is updated.

## REST API

The REST API is the programmatic access to register a host agent. Through the control REST API, you can manage the data process between the cloud storage and host agent. The REST API provides access points for the host agent to query for commands to run and for management of the backup and restore processes.

## Encryption

AES 128-bit encryption and SHA-256 hashing is utilized for S3.

All data transmitted to S3 is encrypted. Each piece of data contains a key derived from a SHA-256 hash combination of a user pass phrase and a piece of data unique to the object being encrypted. CBC is the cipher mode with random IV, which is packaged with the encrypted data.

The pass phrase from the user is a customer global pass phrase. The term phrase implies a string and indeed, it can be a simple phrase from the user converted to a byte array based on an UTF-8 encoding of the pass phrase.

The user passphrase is defined at the host agent-level and never leaves the customer premises. It is never stored or transmitted to the cloud. It is up to the user to keep this data safe on their host systems.

## Terminology

- **Volume** – An E-Series volume
- **Data Chunk (aka Chunk)** – A segment of data read from the volume and written to the cloud.
- **Chunk Size** – The size of the data stripe written to the object store. The size is measured in MegaBytes (Mb) and defaults to 64Mb (1024 \* 1024 \* 64). No base 10 rounding is allowed.
- **Block and Block Size** – The block size for the volume.
- **Image-Based Backup** – A process that reads the raw data blocks from a volume and backs them up to a file known as an image. All of the data blocks on the volume are backed up, including empty blocks, blocks occupied by deleted files, blocks associated with partitioning, and filesystem metadata. Image backups have the advantage of being performed against a volume regardless of the partitioning scheme or filesystems on the volume.
- **SCC Host** – The host where SANtricity Cloud Connector is installed. SANtricity Cloud Connector backup and restore operations must occur on a host, which is IO-attached to the E-Series array containing the volume to be backed up and/or restored. The IO-attached host can be one dedicated for SANtricity Cloud Connector use or one already being used to access the storage for other applications.

# System Setup

## Host Hardware Requirements

Before installing the SANtricity Cloud Connector, verify your system meets the following host hardware requirements:

- At least 5GB of memory - 4GB for the maximum configured heap size
- At least 250MB of free disk space is required for the software installation

An installation of SANtricity Web Services 2.1 is required on the server intended for use with the SANtricity Cloud Connector. For information on installing the SANtricity Web Services Proxy, refer to the [NetApp SANtricity Web Services Proxy 2.1 Installation Guide](#).

## Compatible Storage Arrays and Controller Firmware

For a complete and up-to-date listing of all compatible storage arrays and firmware for the SANtricity Cloud Connector, refer to the [NetApp Interoperability Matrix Tool](#).

## Compatible Operating Systems

The SANtricity Cloud Connector 1.0 application is compatible with and supported on the following operating systems:

| Operating System                    | Version | Architecture |
|-------------------------------------|---------|--------------|
| Red Hat Enterprise Linux (RHEL)     | 7.x     | 64 bit       |
| SUSE Linux Enterprise Server (SLES) | 12.x    | 64 bit       |

## Installation

The Cloud Connector packaged solution (.bin file) is available for RedHat and SUSE Linux platforms only. During the installation process, you must specify the listening port number for the SANtricity Cloud Connector. Once installed, the SANtricity Cloud Connector runs as a daemon process.

**NOTE:** You must install the SANtricity Cloud Connector at a different port location as your current Web Services Proxy installation.

### Installing the SANtricity Cloud Connector on a Linux Operating System in graphical mode

Perform the following to install the SANtricity Cloud Connector on a Linux Operating System through graphical mode:

1. Download the SANtricity Cloud Connector installation file to the desired IO host location.
2. Open a terminal window.
3. Through the terminal window, navigate to the directory file containing the SANtricity Cloud Connector installation file.


4. Run the following command under the terminal window to initiate the SANtricity Cloud Connector installation process:

```
./cloudconnector-xxxx.bin -i gui
```

**NOTE:** The xxxx portion of the above command designates the version number of the application.


The installer window is displayed.

5. After reviewing the Introduction statement under the installer window, click **Next** to proceed with the install process.



The License Agreement for NetApp, Inc. Software is displayed within the installer window.


6. To accept the terms of the License Agreement, select **I accept the terms of the License Agreement** and click **Next**.



The Choose Install screen is displayed within the Installer window. The Where Would You Like to Install field displays the following default install folder:


`/opt/netapp/santricity_cloud_connector`

7. To proceed with the default install folder location specified under the Choose Install screen for the SANtricity Cloud Connector, click **Next**.
  - a. If needed, you can change the default install location for the SANtricity Cloud Connector by entering the absolute path to the new desired folder location under the provided field within the installer window.



An Enter the Non SSL Jetty Port Number screen is displayed. A default value of 8080 is assigned to the Non SSL port.

8. To proceed with the default Non SSL port number for the Jetty Web Service, click **Next**.
  - a. If needed, you can change the default Non SSL port number by entering the new desired port number value under the provided field within the installer window.




The Pre-Installation Summary screen is displayed.

9. Review the displayed Pre-Installation Summary and click **Install** to proceed with installation of the SANtricity Cloud Connector.



The installation of the SANtricity Cloud Connector begins and a Webserver Daemon Setup prompt is displayed. Additional setup under the `config.json` file is required to run the webserver daemon. For more information on configuring the `config.json` file, refer to [Configuring the config.json file](#).

10. Click **OK** to acknowledge the Webserver Daemon Setup prompt and proceed with the installation of the SANtricity Cloud Connector.



The Installation Complete message is displayed.

11. Click **Done** to exit the SANtricity Cloud Connector installer.



## Installing the SANtricity Cloud Connector on a Linux Operating System in console mode

Perform the following to install the SANtricity Cloud Connector on a Linux Operating System through console mode:

1. Download the SANtricity Cloud Connector installation file to the desired IO host location.
2. Open a terminal window.
3. Through the terminal window, navigate to the directory file containing the SANtricity Cloud Connector installation file.
4. Run the following command under the terminal window to initiate the SANtricity Cloud Connector installation process:

```
./cloudconnector-xxxx.bin -i console
```

**NOTE:** The `xxxx` portion of the above command designates the version number of the application.

The installation process for the SANtricity Cloud Connector is initialized.

```
=====
NetApp SANtricity Cloud Connector                               (created with InstallAnywhere)
=====

Preparing CONSOLE Mode Installation...

=====
Introduction
-----

InstallAnywhere will guide you through the installation of NetApp SANtricity
Cloud Connector.

It is strongly recommended that you quit all programs before continuing with
this installation.

Respond to each prompt to proceed to the next step in the installation.  If
you want to change something on a previous step, type 'back'.

You may cancel this installation at any time by typing 'quit'.

PRESS <ENTER> TO CONTINUE: █
```

5. Press **Enter** under the installer window to proceed with the installation process.

The End User License Agreement for NetApp, Inc. Software is displayed within the installer window.

**NOTE:** To cancel the installation process at any time, type `quit` under the installer window.

6. Press **Enter** to proceed through each portion of the End User License Agreement.

The License Agreement acceptance statement is displayed under the installer window.

7. To accept the terms of the End User License Agreement and proceed with the installation of the SANtricity Cloud Connector, enter `Y` and press **Enter** under the installer window.

```
PRESS <ENTER> TO CONTINUE:

* VLAN tags
* iSCSI session information
* Personal email address (not required, voluntary only)
* Relay server address (not required, voluntary only)

NetApp recommends that all transmissions of AutoSupport should be through the
use of https or secure SMTP. Once the data has been received by NetApp,
industry standard access protection procedures are followed to prevent
unauthorized access (both outside and within NetApp) to any information
received.

If you choose to provide the voluntary information during setup or through
subsequent modification of AutoSupport configuration, it will be sent to
NetApp in the body of AutoSupport messages. If you do not wish this
information to be transmitted, please do not set up AutoSupport. However, your
support experience may be impacted by excluding this information because it is
often needed to expedite your problem resolution. Please be assured NetApp
employs the highest level of security to prevent unauthorized access or use of
this data, and complies with all relevant data protection laws.

Copyright (C) 1999-2013 NetApp, Inc. All Rights Reserved.

DO YOU ACCEPT THE TERMS OF THIS LICENSE AGREEMENT? (Y/N): Y
```


A Choose Install Folder message with the following default install folder for the SANtricity Cloud Connector is displayed:

```
/opt/netapp/santricity_cloud_connector
```

**NOTE:** If you do not accept the terms of the End User Agreement, type **N** and press **Enter** to terminate the installation process for the SANtricity Cloud Connector.

8. To proceed with the default install folder location for the SANtricity Cloud Connector, press **Enter**.
  - a. If needed, you can change the default install location for the SANtricity Cloud Connector by entering the absolute path to the new desired folder location under the provided field within the installer window.

An Enter the Non SSL Jetty Port Number message is displayed. A default value of 8080 is assigned to the Non SSL port.



```
=====
Choose Install Folder
=====

Please choose a destination folder for this installation.

Where would you like to install?

  Default Install Folder: /opt/netapp/santricity_cloud_connector
ENTER AN ABSOLUTE PATH, OR PRESS <ENTER> TO ACCEPT THE DEFAULT
:

=====

Enter the Non SSL Jetty Port Number
=====

Enter the Non SSL port number that the Jetty Web Service should run on.

The Non SSL default value is 8080.

Non SSL Jetty Port: (Default: 8080): [ ]
```

9. To proceed with the default Non SSL port number for the Jetty Web Service, press **Enter**.
  - a. If needed, you can change the default Non SSL port number by entering the new desired port number value under the provided field within the installer window.

The Pre-Installation Summary for the SANtricity Cloud Connector is displayed.

```
=====
Pre-Installation Summary
=====

Please Review the Following Before Continuing:

Product Name:
  NetApp SANtricity Cloud Connector

Install Folder:
  /opt/netapp/santricity_cloud_connector

Disk Space Information (for Installation Target):
  Required:  154.76 MegaBytes
  Available: 36,676.95 MegaBytes

PRESS <ENTER> TO CONTINUE: █
```

10. Review the displayed Pre-Installation Summary and press **Enter** to proceed with installation of the SANtricity Cloud Connector.

The installation of the SANtricity Cloud Connector begins and a Webserver Daemon Setup prompt is displayed. Additional setup under the `config.json` file is required to run the webserver daemon. For more information on configuring the `config.json` file, refer to [Configuring the config.json file](#).

11. Press **Enter** to acknowledge the Webserver Daemon Setup prompt and proceed with the installation of the SANtricity Cloud Connector.

```
=====
Webserver Daemon Setup
=====

The webserver daemon was setup to run but additional changes are required to
the config.json file before it can be used.

Afterwards you can interact with it using
systemctl start|stop|restart|status cloud_connector.service

PRESS <ENTER> TO ACCEPT THE FOLLOWING (OK):

=====

Installation Complete
=====

Congratulations. NetApp SANtricity Cloud Connector has been successfully
installed to:

  /opt/netapp/santricity_cloud_connector

PRESS <ENTER> TO EXIT THE INSTALLER: █
```

The Installation Complete message is displayed.

12. Press **Enter** to exit the SANtricity Cloud Connector installer.

## Initial Configuration

Upon successful installation, initial setup of the SANtricity Cloud Connector application is performed through the `config.json` file located under the install directory.

After the installation process, you must enter your login information for S3 and Web Services as clear text under the `config.json` file. Once the application starts, your user-sensitive information is encrypted. For more information on entering login information under the `config.json` file, refer to the [SANtricity Cloud Connector Credential Login Management](#) section.

### Source and Target Root File Paths

Once installed, the SANtricity Cloud Connector application files are available under the folder location specified during the installation process. The following table details the directory layout of the SANtricity Cloud Connector on a host setup:

| Directory                   | Files                                 | Description                                                                                |
|-----------------------------|---------------------------------------|--------------------------------------------------------------------------------------------|
| SCC_ROOT_DIR                | <code>s3backup.cmd</code>             | Startup script                                                                             |
|                             | <code>s3backup.jar</code>             | Startup jar file                                                                           |
|                             | <code>hostagentweb.war</code>         |                                                                                            |
|                             | <code>config.json</code>              | Configuration file (JSON format)                                                           |
| SCC_ROOT_DIR/jre            |                                       | Location of the Java Runtime environment.                                                  |
| SCC_ROOT_DIR/lib            |                                       | Location of all third-party JAR libraries used by the solution.                            |
| SCC_ROOT_DIR/samplescripts  |                                       | Location of sample scripts, which can be used to perform various tasks.                    |
| SCC_ROOT_DIR/templates      |                                       | Location of templates, which can be used as input to the REST API endpoints.               |
| SCC_ROOT_DIR/templates/jobs |                                       | Location of job-related templates for backup/restore.                                      |
|                             | <code>blockBackupToS3Target</code>    | Template file for performing an image-based backup of an E-Series volume to an S3 target.  |
|                             | <code>blockRestoreFromS3Target</code> | Template file for performing an image-based restore of an E-Series volume to an S3 target. |
| SCC_ROOT_DIR/web            |                                       | Future Location of files and directories related to the Web GUI interface.                 |

| Directory                              | Files | Description                                                                                                          |
|----------------------------------------|-------|----------------------------------------------------------------------------------------------------------------------|
| SCC_ROOT_DIR/working                   |       | Location of files, which need to persist across a running instance of SANtricity Cloud Connector.                    |
| SCC_ROOT_DIR/working/logs              |       | Location of any log files                                                                                            |
| SCC_ROOT_DIR/working/webapps           |       | When SANtricity Cloud Connector is started, the hostagentweb.war is self-extracted and its contents are placed here. |
| SCC_ROOT_DIR/uninstall_cloud_connector |       | Location of InstallAnywhere files needed to perform a removal of the solution.                                       |

## SANtricity Cloud Connector Credential Login Management

Credential information for logging into Web Services and S3 accounts is encrypted by the SANtricity Cloud Connector application and then decrypted instantaneously when needed. To ensure the security of login credentials, be aware of the following when setting up and using the SANtricity Cloud Connector:

- **Initial configuration** – Upon successful installation of the SANtricity Cloud Connector, you must specify Web Server and S3 account credential information in clear text under the `config.json` file. When the SANtricity Cloud Connector starts up, the application reads the `config.json` file into memory, automatically encrypt the credential information, and then writes the information back to disk.
- **Post-configuration** – When the SANtricity Cloud Connector starts up, the application reads the configuration file into memory and the credential information is verified to ensure it has been encrypted. If not, the credentials are automatically encrypted and written back to disk.

The following are common scenarios where the encrypted credentials may not allow access to either Web Services or S3 accounts:

- **Updated Credentials** – If you have changed the login credentials to Web Services or S3 accounts, you must edit the `config.json` file with the updated information in clear text and then restart the SANtricity Cloud Connector application. The SANtricity Cloud Connector will then encrypt the new login credentials.
- **Host Server Updates** – The key used by the SANtricity Cloud Connector to encrypt credentials relies partially on characteristics of the host where the SANtricity Cloud Connector application is installed. If these characteristics have changed substantially, the encrypted credentials will fail. The SANtricity Cloud Connector will attempt to determine if the failure is due to invalid credentials or a host server update. In either case, the `config.json` file must be updated with information in clear text and then restarted. Upon restart, the SANtricity Cloud Connector application encrypts the updated credentials.

## Proxy Server Login Management

Connection to the SANtricity Web Services is required for the SANtricity Cloud Connector. The connection can be through a proxy server. Login and connection information for the SANtricity Web Services Proxy

must be entered and maintained under the `config.json` file. Refer to [Web Services Proxy Connection and Login Information](#) for further details.

## Configuring the config.json file

Prior to use, you must configure the SANtricity Cloud Connector application through the `config.json` file located under the `<SCC_ROOT_DIR>` folder. Initial configuration of the `config.json` file for the SANtricity Cloud Connector requires the following information:

- Port number where the SANtricity Cloud Connector web server listens for requests
- State Storage Type
- Encryption master key
- Amazon S3 account credentials
- Web Services proxy URL and login information
- Target configuration
- Job Engine configuration

## Webserver Port Number

To configure the port number where the SANtricity Cloud Connector webserver listens for requests, enter the desired port number under the `port` value and the install folder location of the SANtricity Cloud Connector under the `path` value for `ROOT`.

### **EXAMPLE**

```
"webserver": {
  "port": <SCC_PORT>,

  "apps": [
    {
      "context": "ROOT",
      "path": "<SCC_ROOT_DIR>hostagentweb.war"
    },
  ]
}
```

## State Storage Type

**NOTE:** For the 1.0 release of the SANtricity Cloud Connector, only S3 state storage is configurable for use with the application.

Before utilizing the SANtricity Cloud Connector, you must configure the state storage type as S3 under the `config.json` file. Under the `type` value, enter `s3` to specify the state storage type for the SANtricity Cloud Connector application. Under the `bucket` value, enter the bucket name for the S3 storage. For S3 storage, you must also specify the `bucket` and `outputBucket` values. The bucket specified under the `outputBucket` value must exist on S3.

### **Example**

```
"stateStorage": {
  "type": "<s3>",
  "bucket": "<s3-bucketname>"
},
"outputBucket": "<s3-bucketname>",
```

## S3 Target

The SANtricity Cloud Connector's backup and restore capabilities for S3 storage requires knowledge of the S3 target configuration. Under the `s3TargetConfig` value, enter the access ID, access key, end point, and path style access for the S3 target. For the 1.0 release of the SANtricity Cloud Connector, default values under the `maxClients`, `s3Endpoint`, and `usePathStyleAccess` should not be modified.

**NOTE:** You must modify the data under the `s3TargetConfig` value before any backup and restore operations for the S3 storage will occur.

### Example

```
"s3TargetConfig": {
  "accessKeyId": "<your_access_id>",
  "secretAccessKey": "<your_access_key>",
  "maxClients": 6,
  "s3Endpoint": "<your_s3_endpoint>",
  "usePathStyleAccess": true
},
```

## Web Services Proxy Connection and Login Information

The SANtricity Cloud Connector application requires an installation of the Web Services Proxy. Connection and login information for the Web Services Proxy used in conjunction with the SANtricity Cloud Connector application must be entered under the `wsproxy` value.

**NOTE:** You must modify the `wsproxy` value before any backup and restore operations will occur.

### Example

```
"wsproxy": {
  "url": "<web_services_url>",
  "uid": "rw",
  "pwd": "rw"
},
```

**NOTE:** When updating the `pwd` value, enter in the cleartext password. The SANtricity Cloud Connector will detect the cleartext password exists, automatically encrypt the entry, and write the encrypted password back to the `config.json` file upon the application's next initialization. For information on configuring login information for the Web Services Proxy, refer to the [Web Services Proxy User Guide](#).

## Job Engine Configuration

**NOTE:** Default settings under `jobEngineConfig` should not be modified unless specifically requested by NetApp Support.

A dedicated Job Engine handles the encryption/decryption and compression/decompression process for the SANtricity Cloud Connector. Located under the `jobEngineConfig` section of the `config.json` file, the Job Engine utilizes multiple threads and jobs to perform given backups or restores more efficiently. Each job is composed of the following key elements:

- **Reader thread** – Reads the source data. Chunks of data are read from the source and placed into a processing queue, where the processing thread(s) retrieve data.
- **Processing thread(s)** – Containing one or more threads, performs the encryption/decryption and compression/decompression. Once a data chunk is processed, the item is placed into a writer queue, where the writer thread(s) retrieve data.
- **Writer thread(s)** – Containing one or more threads, writes data to the destination.

When large chunk sizes are used (e.g., 64MB), the amount of memory consumed on a host can be substantial (up to 100's GBs of RAM). To keep usage down, the job engine provides a set of configuration parameters, which limits how many threads and jobs are allowed to run simultaneously.

The following parameters for the Job Engine are available for configuration:

- `writerThreadsPerJob` – Specifies the total number of writer threads allowed per job.
- `readerThreadsPerJob` – Specifies the total number of reader threads allowed per job.
- `processingThreadsPerJob` – Specifies the total number of processing threads allowed per job.
- `writerQueueSize` – Specifies the maximum number of requests allowed in the writer queue. Any processing thread attempting to write to a full queue will block until at least one request in the queue has been removed (i.e., pulled off the queue by a writer thread).
- `processingQueueSize` – Specifies the maximum number of requests allowed in the processing queue. If the reader thread attempts to write to a full queue, it will block until at least one request in the queue has been removed (i.e., pulled off the queue by a processing thread).
- `chunkSizeInMB` – Specifies the chunk size in MB. The engine will automatically multiply this value by (1024 \* 1024) to arrive at the true byte count.
- `MaxRunningJobs` – Specifies the maximum number of jobs (backup or restore) that are allowed to run simultaneously.

Refer to the following for the default settings under `jobEngineConfig`:

```
"jobEngineConfig": {
  "writerThreadsPerJob" : 5,
  "readerThreadsPerJob" : 1,
  "processingThreadsPerJob" : 5,
  "writerQueueSize" : 5,
  "processingQueueSize" : 5,
  "chunkSizeInMB" : 64,
  "maxRunningJobs" : 2
},
```

# Backup and Restore Jobs

## Backup Jobs

The SANtricity Cloud Connector uses the concept of jobs to perform the actual backup/restore of an E-Series volume. Templates for backup jobs to S3 targets are available under the `<SCC_ROOT_DIR>/templates/jobs` folder. The SANtricity Cloud Connector application utilizes backup data in the form of snapshots of an E-Series volume. Once you map a snapshot of an E-Series volume for backup to the SANtricity Cloud Connector host, you can use the `blockBackupToS3Target` template to enter the backup volume data and process the backup job.

### Mapping Snapshots to the SANtricity Cloud Connector Host for Backup

The backup process for the SANtricity Cloud Connector application requires the creation of application consistent snapshots of the volume(s) data being backed up mapped to the host. The creation and mapping of these snapshot is required before any backup operations through the SANtricity Cloud Connector application can occur.

Perform the following to map a created snapshot volume of the E-Series Volume to the SANtricity Cloud Connector host for backup:

1. Under the SANtricity Storage Manager, create a snapshot volume of the E-Series volume you want to map to the SANtricity Cloud Connector host.
2. Map the created the snapshot volume to the SANtricity Cloud Connector host in order for the application to recognize the volume.

**NOTE:** For detailed information on how to create a snapshot volume and map the volume to a host, refer to the *SANtricity Online Help*.

3. Retrieve the snapshot volume and array World Wide Name data for the created snapshot volume.
4. Apply the snapshot volume and array World Wide Name data for the created snapshot volume to the `blockBackupToS3Target` template file under the `/jobs` folder. Refer to the [Configuring and processing block-based backup job to an S3 Target](#) for detailed information on how to access, configure, and process the backup job.

### Configuring and processing a Block-Based Backup Job to an S3 Target

A block-based backup job (Image Backup) from an E-Series Storage volume to an S3-compatible target is configurable through the `blockBackupToS3Target` template file under the `/jobs` folder. Perform the following to create a new backup job to an S3-comptaible target:

1. Go to the `<SCC_ROOT_DIR>/templates/jobs` folder.
2. Open the `blockBackupToS3Target` template.
3. Under the `blockBackupToS3Target` template, specify the backup parameters for the S3-compatible target.

#### Example

```
{  
  "jobType": "BACKUP",
```

```

    "description": "Template file demonstrating backup to an S3 target.",
    "sourceType": "eSeries",
    "eSeriesSource" : {
        "arrayWWN": "<Enter_Array_WWN_Here>",
        "volumeWWN": "<Enter_Volume_WWN_Here>",
        "volumeName": "<Optional_For_Reference_Only>"
    },
    "targetType": "s3",
}

```

**NOTE:** E-Series volume WWN information can be obtained through the /eSeries endpoint under the SANtricity Cloud Connector API documentation.

4. Once complete, submit the blockBackupToS3Target template using the POST /jobs HTTP method and URL to process the backup job.

## Restore Jobs

The SANtricity Cloud Connector uses the concept of jobs to perform the actual backup/restore of an E-Series volume. Templates for restore jobs from S3 targets are available under available <SCC\_ROOT\_DIR>/templates/jobs folder. Before performing a backup, you must identify which E-Series volume will be used for the operation. Once you map an E-Series volume for restore to the SANtricity Cloud Connector host, you can use the blockBackupToS3Target template to enter the restore volume data and process the backup job.

### Mapping volumes to the SANtricity Cloud Connector Host for Restore

Before performing a restore from S3, you must decide if the data will be restored to an existing volume on the array or if a new array will be created. Whichever option is selected, it is critical the location of the restore volume have storage capacity as large as the original volume backup.

Perform the following to map an E-Series Volume to the SANtricity Cloud Connector host for restore:

1. Under the SANtricity Storage Manager, select the desired E-Series volume target for the restore job.

-Or-

Create a new E-Series volume for the restore job.

2. Map the selected restore volume to the SANtricity Cloud Connector.

**NOTE:** For detailed information on how to create and map a volume to a host, refer to the *SANtricity Online Help*.

3. Retrieve the volume and array World Wide Name data for the mapped restore volume.
4. Apply the volume and array World Wide Name data for the mapped restore volume to the blockRestoreFromS3Target template file under the /jobs folder. Refer to the [Configuring and processing a Block-Based Restore Job to an E-Series target](#) for detailed information on how to access, configure, and process the restore job.

### Configuring and processing a Block-Based Restore Job to an E-Series target

A block-based restore job (Image Restore) from an S3-comptabile target to an E-Series storage volume target is configurable through the blockRestoreFromS3Target template file under the /jobs folder. Perform the following to create a new restore job for a S3-compatible target:

1. Go to the <SCC\_ROOT\_DIR>/templates/jobs folder.
2. Open the blockRestoreFromS3Target template.
3. Under the blockRestoreFromS3Target template, specify the restore parameters for the S3-compatible target.

#### **Example**

```
{
  "jobType": "RESTORE",
  "description": "Template file demonstrating restore from an S3 source.",
  "sourceType": "s3BackedUpData",
  "s3BackedUpDataSource": {
    "id": <TBD>
    "version": <TBD>
  },
  "targetType": "eseries",
  "eSeriesTarget" : {
    "arrayWWN": "<Enter_Array_WWN_Here>",
    "volumeWWN": "<Enter_Volume_WWN_Here>",
    "volumeName": "<Optional_For_Reference_Only>"
  }
}
```

4. Once complete, submit the blockRestoreFromS3Target template using the POST /jobs HTTP method and URL to process the restore job.

# Uninstalling the SANtricity Cloud Connector

## Uninstalling the SANtricity Cloud Connector on a Linux Operating System through graphical mode

Perform the following to uninstall the SANtricity Cloud Connector on a Linux Operating System:

1. From a terminal window, navigate to the directory containing the SANtricity Cloud Connector uninstall file.

The uninstall file for the SANtricity Cloud Connector is available at the following default directory location:

```
/opt/netapp/santricity_cloud_connector/uninstall_cloud_connector
```

2. From the directory containing the SANtricity Cloud Connector uninstall file, run the following command:

```
./uninstall_cloud_connector -i gui
```

The uninstall process for the SANtricity Cloud Connector is initialized.

3. Click **Uninstall** under the uninstall window to proceed with uninstalling the SANtricity Cloud Connector.



The uninstall process is completed and the SANtricity Cloud Connector application is uninstalled under the Linux operating system.

## Uninstalling the SANtricity Cloud Connector on a Linux Operating System through console mode

Perform the following to uninstall the SANtricity Cloud Connector on a Linux Operating System:

1. From a terminal window, navigate to the directory containing the SANtricity Cloud Connector uninstall file.

The uninstall file for the SANtricity Cloud Connector is available at the following default directory location:


```
/opt/netapp/santricity_cloud_connector/uninstall_cloud_connector
```

2. From the directory containing the SANtricity Cloud Connector uninstall file, run the following command:

```
./uninstall_cloud_connector -i console
```

The uninstall process for the SANtricity Cloud Connector is initialized.

3. Press **Enter** under the uninstall window to proceed with uninstalling the SANtricity Cloud Connector.



```
=====
NetApp SANtricity Cloud Connector                (created with InstallAnywhere)
=====

Preparing CONSOLE Mode Uninstallation...

=====

Uninstall NetApp SANtricity Cloud Connector
=====

About to uninstall...

NetApp SANtricity Cloud Connector

This will remove features installed by InstallAnywhere.  It will not remove
files and folders created after the installation.

PRESS <ENTER> TO CONTINUE: █
```

The uninstall process is completed and the SANtricity Cloud Connector application is uninstalled under the Linux operating system.

## Copyright information

Copyright © 1994–2017 NetApp, Inc. All rights reserved. Printed in the U.S.

No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp. The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

## Trademark information

Active IQ, AltaVault, Arch Design, ASUP, AutoSupport, Campaign Express, Clustered Data ONTAP, Customer Fitness, Data ONTAP, DataMotion, Fitness, Flash Accel, Flash Cache, Flash Pool, FlexArray, FlexCache, FlexClone, FlexGroup, FlexPod, FlexScale, FlexShare, FlexVol, FPolicy, Fueled by SolidFire, GetSuccessful, Helix Design, LockVault, Manage ONTAP, MetroCluster, MultiStore, NetApp, NetApp Insight, OnCommand, ONTAP, ONTAPI, RAID DP, RAID-TEC, SANscreen, SANshare, SANtricity, SecureShare, Simplicity, Simulate ONTAP, Snap Creator, SnapCenter, SnapCopy, SnapDrive, SnapIntegrator, SnapLock, SnapManager, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapValidator, SnapVault, SolidFire, SolidFire Helix, StorageGRID, SyncMirror, Tech OnTap, Unbound Cloud, and WAFL and other names are trademarks or registered trademarks of NetApp, Inc., in the United States, and/or other countries. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such. A current list of NetApp trademarks is available on the web.

<http://www.netapp.com/us/legal/netapptmlist.aspx>

## How to send your comments

You can help us to improve the quality of our documentation by sending us your feedback. You can receive automatic notification when production-level (GA/FCS) documentation is initially released or important changes are made to existing production-level documents.

If you have suggestions for improving this document, send us your comments by email. [doccomments@netapp.com](mailto:doccomments@netapp.com). To help us direct your comments to the correct division, include in the subject line the product name, version, and operating system.

If you want to be notified automatically when production-level documentation is released or important changes are made to existing production-level documents, follow Twitter account @NetAppDoc.

You can also contact us in the following ways:

- NetApp, Inc., 495 East Java Drive, Sunnyvale, CA 94089 U.S.
- Telephone: +1 (408) 822-6000
- Fax: +1 (408) 822-4501
- Support telephone: +1 (888) 463-8277





**NetApp**