# CA Cloud Storage for System z

## Best Practices

Release 1.1.00

# Contact CA Technologies

**Contact CA Support**

For your convenience, CA Technologies provides one site where you can access the information that you need for your Home Office, Small Business, and Enterprise CA Technologies products. At http://ca.com/support, you can access the following resources:

- Online and telephone contact information for technical assistance and customer services
- Information about user communities and forums
- Product and documentation downloads
- CA Support policies and guidelines
- Other helpful resources appropriate for your product

**Providing Feedback About Product Documentation**

If you have comments or questions about CA Technologies product documentation, you can send a message to techpubs@ca.com.

To provide feedback about CA Technologies product documentation, complete our short customer survey which is available on the CA Support website at http://ca.com/docs.

**Best Practices Guide Process**

These best practices are based on customer experience reported through interviews with development, technical support, and technical services. Therefore, many of these best practices are a collaborative effort stemming from customer feedback.

To continue to build on this process, we encourage you to share common themes of product use that might benefit other users. Please consider sharing your best practices with us.

To share your best *practices, c*ontact us at techpubs@ca.com and preface your email subject line with "Best Practices for product name" so that we can easily identify and categorize them.

# Contents


## Chapter 1: About this Guide 7

Audience ..... 7
Conventions ..... 7

## Chapter 2: Cloud Storage for System z Best Practices 9

Limited-Use CA Vtape License Requirement ..... 9
Share Your Best Practices ..... 9
Manage Riverbed SteelStore Gateway Behavior ..... 10
Implement Disk Storage Features ..... 11
Set Up Numbering Scheme for CTC Control Units and Devices ..... 12
Track Disk Space Usage ..... 15
Jumbo Frames ..... 16
Planning for Disaster Recovery ..... 16
- CA Cloud Storage for System z Virtual Tape System for z/OS Best Practice ..... 18
- CA Cloud Storage for System z Connector Best Practice ..... 18
- Disk Storage Appliances Best Practice ..... 18

# Chapter 1: About this Guide

With CA Cloud Storage for System z, you can create a virtual tape environment from any combination of mainframe and Linux resources. You can add or upgrade resources to take advantage of advancements in capacity and performance. And it scales to meet your needs without the high costs that are associated with hardware solutions.

This section contains the following topics:

Audience (see page 7)
Conventions (see page 7)

## Audience

Users of this guide should be experienced mainframe technicians with knowledge of their mainframe tape systems, tape related software, and security configuration.

## Conventions

The following conventions are used throughout this guide to document features, functions, and other aspects of the system:

- Variable text, most commonly used for data set names and console commands, is entered in italics.

  For example, VVE_SCRATCH=*volser* where *volser* is the Virtual Volume VOLSER.

- Commands are entered in uppercase and lower-case. The uppercase portion is the minimum number of characters that must be entered for CA Vtape to recognize the command. The lower-case portion is provided for clarity.
- Features, functions, and components of CA Cloud Storage for System z are capitalized. These include, for example: Virtual Volumes and Virtual Devices.

# Chapter 2: Cloud Storage for System z Best Practices

This section contains information about how to make your use of CA Cloud Storage for System z successful. Consider these best practices before implementing CA Cloud Storage for System z. The best practice recommendations include how to:

- Improve product performance and efficiency
- Improve reliability and recoverability from failure
- Understand product design characteristics
- Reduce the costs for product usage

Considering these best practices lets you lower the time for downloading and reduce your cost using AWS Glacier. Keep your costs under control by tracking your disk space use and for maximum performance and security, develop a disaster recovery plan, implement data replication, and other device features.

This section contains the following topics:

Limited-Use CA Vtape License Requirement (see page 9)
Share Your Best Practices (see page 9)
Manage Riverbed SteelStore Gateway Behavior (see page 10)
Implement Disk Storage Features (see page 11)
Set Up Numbering Scheme for CTC Control Units and Devices (see page 12)
Track Disk Space Usage (see page 15)
Jumbo Frames (see page 16)
Planning for Disaster Recovery (see page 16)

## Limited-Use CA Vtape License Requirement

New and existing CA Vtape customers require the Limited-Use CA Vtape License to implement CA Cloud Storage for System z.

## Share Your Best Practices

We encourage you to share common themes of product use that might benefit other users. Consider sharing your best practices with us.

To share your best practices, contact us at techpubs@ca.com. Preface your email subject line with "Best Practices for C*A Cloud Storage for System z*" so that we can easily identify and categorize them.

# Manage Riverbed SteelStore Gateway Behavior

The Riverbed SteelStore gateway supports Amazon S3 and AWS Glacier public cloud storage. Both are low-cost storage services that provide secure, durable, and flexible storage for data backup and archive.

The data that is stored on Amazon S3 can be retrieved without delay. The data that is stored on AWS Glacier is lower cost, but it requires a four to five hour delay before retrieving data. Glacier also has a monthly limit of data you can retrieve. After you reach that limit, you are charged for retrieving data from Glacier.

You can use the Riverbed SteelStore gateway Command-Line Interface (CLI) to modify the default behavior. You can get the data back quicker but it can raise your cost.

**Follow these steps:**

1. Log in to the gateway to access the CLI.
2. Enter the enable and configure terminal commands.

   ```
   ssh -l admin <IP Address>
   Riverbed SteelStore
   admin@<IP Address>'s password:
   Last login: Fri Jan 31 20:03:26 2014 from nnn.nnn.nnn.nnn
   enable
   configure terminal
   ```

3. Change the default from one hour to one minute to send an unfilled package.

   The Riverbed SteelStore gateway groups data into packages to minimize the number of transmissions between the gateway and the cloud storage provider. The default behavior is to wait one hour before sending an unfilled package. If this replication interval does not satisfy your Recovery Time Objective (RTO), you can overwrite this value. The following example shows how to change the default:

   ```
   rfsctl exec "-S -w replicator.max_unfilled_package_age=60"
   replicator.max_unfilled_package_age: 3600 -> 60
   ```

4. Overwrite the default behavior for replicating to AWS Glacier using the following commands:

   ```
   rfsctl exec "-S -w rfs_fuse.ca_cs4z_mode=true"
   rfs_fuse.ca_cs4z_mode: false -> true
   ```

5. Increase the Riverbed SteelStore throttle limit.

   By default, the Riverbed SteelStore gateway has a restore throttle of five percent for data that is retrieved from Glacier. This throttle keeps retrievals below the no-cost limit. Therefore, you can use the SteelStore gateway to restore five percent of total cloud usage in a month. The throttle is enforced on an hourly basis. Hourly data retrieval is limited to (five percent of the total cloud use)/(hours per month). You can increase the five percent restore throttle limit up to 100 percent or completely disable it by setting the limit to zero. This change can incur data retrieval charges. To change this behavior, use the following command:

   ```
   rfsctl exec "-S -w prepop.restore_percent_limit_per_hour=0"
   prepop.restore_percent_limit_per_hour: 5 -> 0
   ```

6. Exit the CLI interface by typing exit two times when your changes are complete.

# Implement Disk Storage Features

You can configure CA Cloud Storage for System z with various disk storage appliances. Different appliance devices can have varied characteristics or capabilities that require your consideration.

If your device has the following features, we recommend that you implement them for maximum performance and security.

**Data deduplication**

Data deduplication provides compression by tracking common blocks of data across files that are stored on the appliance. If the virtual tape drive on the MVS Host does not pre-encrypt or compress the data, high compression rates are achieved.

**Data encryption and key management**

Data encryption and key management protect your data from unauthorized access. To achieve the best data deduplication and compression rates, while minimizing host cycles, use a disk storage device that supports encryption. Doing so encrypts the data while it is on-premise and lets you manage the encryption keys.

**Data replication**

Data replication asynchronously copies files onto a secondary device at a remote location using TCPIP. You can use data deduplication to reduce network traffic. When copying asynchronously, consider how much delay there can be before the data is replicated to a secondary device.

Many public cloud providers have sophisticated methods of internal replication to ensure the data integrity of client data. We recommend that you have replicated storage devices to avoid data loss due to catastrophic failure. For more information, see Disaster Recovery.

**Snapshot copy**

Snapshot copy lets you create file copies on the same appliance. These copies are not intended for disaster recovery but they can be used to:

- Recover a file that is accidentally modified or deleted
- Create a point-in-time backup

  A point-in-time backup can be used for some other local purpose while real-time updates continue against the production files.

Snapshot copy can affect the disk access performance when the device is making or accessing the copies. Snapshot copy also increases the number of files that are written to the appliance.

Even on systems without this feature, CA Cloud Storage for System z keeps scratched Virtual Volumes for three days. As Virtual Volumes are scratched, their data files are not immediately deleted on the Linux file system. Instead the data file is renamed to volser.scr-yyyymmdd-hhmmss. You can recover these scratched data files by manually renaming the file to volser.vve.

# Set Up Numbering Scheme for CTC Control Units and Devices

CA Cloud Storage for System z uses Channel-to-Channel (CTC) devices to transfer tape data between your z/OS LPARs and Linux on System z. CTC connectivity provides an efficient and flexible approach to moving data between LPARs in your data center. They are efficient because they use few general processor cycles and are flexible in how you connect them to your LPARs. For example, you can use CTC devices to connect LPARs on the same central processor complex (CPC) or between LPARs on different CPCs. You can also use a FICON director.

The CTC devices are implemented in IBM licensed internal code and do not require dedicated resources. For example, CTC devices use channel path identifiers (CHPIDs) which can be shared between LPARs and the CHPIDs used for CTC communication. They can also be used for other device types, like DASD. Using a FICON director lets you define different device types on the same CHPID.

Plan your physical topology before setting up the CTC connectivity. We recommended that you implement a numbering scheme for your CTC control units and devices.

The following recommendation is similar to the one that comes from IBM System z ESCON and FICON Channel-to-Channel Reference SB10-7034. You do not need to define the send and receive addresses as recommended in the IBM Reference Guide. FICON can send and receive using the same CTC address.

Assign a CTC image ID to every logical partition or basic mode CPC you want to establish CTC communication to. This CTC image ID is only for managing your CTC definition and not used by CTC communication. Create the CTC image ID as a two-digit hexadecimal number in the range X'00' through X'FF' holding a unique value within your complex.

The ID identifies a specific image in your complex that every other image uses to address that specific image.

Identify the control unit and device with four hexadecimal digits that:

1. Identify the control unit or device.
2. Display the CTC image ID.
3. Display the CTC image ID.
4. Use only for addressing.

### Sample Configuration

The sample configuration that follows shows CTC communications between three LPARs existing on the same CPC connected by two CHPIDs using two FICON Directors.

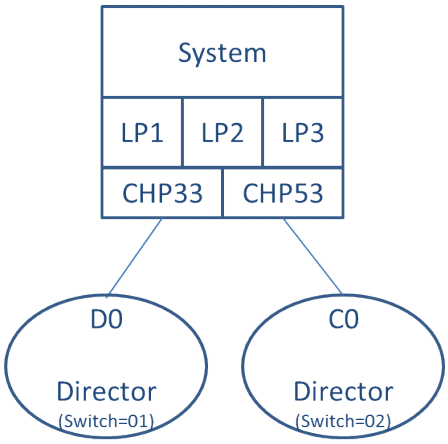


### LPAR Definitions

```
RESOURCE PARTITION=((LP1,1),(LP2,2),(LP3,3))
```

### CHPID MACROS

```
CHPID PATH=(33),TYPE=FC,SWITCH=01,PART=(LP1,LP2,LP3),SHARED
CHPID PATH=(53),TYPE=FC,SWITCH=02,PART=(LP1,LP2,LP3),SHARED
```

### CU MACROS

```
CNTLUNIT CUNUMBR=4010,PATH=33,LINK=D0,UNIT=FCTC,
      UNITADD=((00,2)),CUADD=1
IODEVICE ADDRESS=(4010,2),CUNUMBR=4010,UNIT=FCTC,
      UNITADD=00,NOTPART=LP1
CNTLUNIT CUNUMBR=4020,PATH=33,LINK=D0,UNIT=FCTC,
      UNITADD=((00,2)),CUADD=2
IODEVICE ADDRESS=(4020,2),CUNUMBR=4020,UNIT=FCTC,
      UNITADD=00,NOTPART=LP2
CNTLUNIT CUNUMBR=4030,PATH=33,LINK=D0,UNIT=FCTC,
      UNITADD=((00,2)),CUADD=3
IODEVICE ADDRESS=(4030,2),CUNUMBR=4030,UNIT=FCTC,
      UNITADD=00,NOTPART=LP3
CNTLUNIT CUNUMBR=5010,PATH=53,LINK=C0,UNIT=FCTC,
      UNITADD=((00,2)),CUADD=4
IODEVICE ADDRESS=(5010,2),CUNUMBR=5010,UNIT=FCTC,
      UNITADD=00,NOTPART=LP1
CNTLUNIT CUNUMBR=5020,PATH=53,LINK=C0,UNIT=FCTC,
      UNITADD=((00,2)),CUADD=5
IODEVICE ADDRESS=(5020,2),CUNUMBR=5020,UNIT=FCTC,
      UNITADD=00,NOTPART=LP2
CNTLUNIT CUNUMBR=5030,PATH=53,LINK=C0,UNIT=FCTC,
      UNITADD=((00,2)),CUADD=6
IODEVICE ADDRESS=(5030,2),CUNUMBR=5030,UNIT=FCTC,
      UNITADD=00,NOTPART=LP3
Cloud Storage for System z CTC Usage
```

### Addresses Example

If LPAR1 and LPAR2 are z/OS and they want to communicate to Linux on System z running on LPAR3, the following definitions are used:

- LP1 and LP2 specify addresses 4030-4031 and 5030-5031 to communicate to Linux on System z in LP3.
- Linux on System z specifies 4010-4011 and 5010-5011 when communicating to LP1.
- Linux on System z specifies 4020-4021 and 5020-5021 when communicating to LP2.

# Track Disk Space Usage

CA Cloud Storage for System z records the daily average terabytes usage per month in a file in the /var/lib/cacloud/vault_01 directory. These records make tracking how much disk space you are using easy.

To display or email a report about usage, use the following command:

```
cacloud vault_stats -—mail [email@addresses]
```

**--mail**

Routes the report to an email address

The report is as follows:

```
CA Cloud Storage for System z

Virtual Volume Vault Average Disk Usage

Day Month XX Time EST Year

Month       #Events          Avg_TB       Base2_Avg_TB

2013.10         3             0.05             0.05

2013.11        23             2.72             2.47

2013.12        21             2.44             2.22

2014.01        23             3.08             2.80

2014.02        10             0.86             0.78
```

**#Events**

The number of log records found.

**Avg_TB**

The computed average number of terabytes used in the month, which is calculated as follows:

```
(sum(log_records)/count(log_records))/10^12
```

**Base2_Avg_TB**

The computed average number of terabytes using a base 2 conversion factor for terabytes, which is calculated as follows:

```
(sum(log_records)/count(log_records))/2^40
```

To avoid disk usage billing errors, follow these recommendations:

- Share the control file for multiple CA Vtape subsystems communicating with the same Linux Connector.
- Define the same mount points for an appliance (or appliances) when running multiple Linux Connectors in a CA Cloud Storage for System z Complex. Shared mount points ensure that reports reflect correctly the average disk use for Virtual Volumes.

# Jumbo Frames

Jumbo frames provide more network throughput through the Open System Adaptor or any network adaptor by using a more efficient block size. We recommend using jumbo frames to help reduce your CA Cloud Storage for System z Linux CPU consumption and increase your network throughput.

See the *Linux Install and Configuration Guide* for more information.

# Planning for Disaster Recovery

A disaster recovery plan documents processes and procedures necessary to continue business operations if adverse conditions occur. Make your planning and testing to recover from different levels of disaster a central part of your design. A disaster can range from short outages within your primary data center to a permanent loss of your primary data center.

Planning and testing for outages ensures that your business systems are recoverable. Test, monitor, and review these processes and procedures periodically; business systems often change over time.

Consider the following two production services when planning and testing:

**Recovery Point Objective (RPO)**

A point in time data must be recovered from backup storage so that the data is in a consistent state.

**Recovery Time Objective (RTO)**

The amount of time that is required to restore your data to your desired RPO.

Recovery includes the following scenarios:

**Media Failure Recovery**

A media failure recovery can occur when the z/OS DASD holding the CA Vtape control files fails. You typically restore the control files from a backup copy and the recovery time is typically short.

**Cold Disaster Recovery**

A cold recovery site provides an empty building space. You restore the appliances and infrastructure before restoring your data. This recovery type can require significant time to get your systems back in production.

**Warm Disaster Recovery**

A warm recovery model provides a building space with two appliances in peer replication (primary and secondary configuration) mode. Extra hardware and software can be required before restoring backup data. This recovery type requires less time than a cold disaster recovery, but can cost more.

**Hot or Active-Active Disaster Recovery**

A hot recovery model is the most expensive yet fastest way to get your systems back in the event of an interruption. Hardware and operating systems are kept in sync and in place between data centers. In an active-active disaster recovery setup, there is synchronous data replication between the primary and secondary sites, with no delayed resiliency. While CA Cloud Storage for System z can participate in a hot or active-active disaster recovery scenario, it is beyond the scope of your tape system.

Include the following components in your disaster recovery plans for CA Cloud Storage for System z:

- CA Cloud Storage for System z Virtual Tape System for z/OS
- CA Cloud Storage for System z Connector
- Disk Storage Appliances

## CA Cloud Storage for System z Virtual Tape System for z/OS Best Practice

To recover your CA Vtape VTS environment, you can use hardware mirroring like IBM GDPS for your infrastructure control files. The infrastructure control files for CA Vtape VTS include the Global, BSDS, and Local VCAT data sets. The infrastructure control files can include your ICF catalogs and tape management control files. A mirrored environment lets you have the same address spaces and configuration as your primary site, for communicating with your CA Cloud Storage for System z Connector (Cloud Connector).

If you have not mirrored your files, follow the procedure in the *z/OS Administration Guide to* back up and recover your CA Vtape VTS control files.

## CA Cloud Storage for System z Connector Best Practice

To recover your Cloud Connector environment use hardware mirroring (like IBM GDPS) for your VM operating system and Linux as guest of z/VM. A mirrored environment lets you have minimal changes to the configuration files when you resume normal operations.

An alternative approach to hardware mirroring is to maintain a physical tape backup of your Cloud Connector environment using a product like IBM ADRDSSU. Then at the recovery site you can restore your Cloud Connector environment from the physical tape.

## Disk Storage Appliances Best Practice

When selecting disk storage appliances to use for storing Virtual Volume files, select appliances that support off-site replication. You can recover a lost Virtual Volume only if it was replicated.

The major disk storage vendors support replication; however, the implementation and topologies that various vendors use can be different. How these differences can affect Recovery Time Objectives (RTO) is important.

Disk storage appliances that are tested with CA Cloud Storage for System z include EMC Data Domain, NetApp, and Riverbed SteelStore.

We recommend that you implement a bidirectional replication strategy to facilitate both disaster recovery testing and an actual disaster recovery event. See the *z/OS Administration Guide* for more information.

## EMC Data Domain Appliance

CA Cloud Storage for System z can store Virtual Volume data using the EMC Data Domain disk storage appliance in a private cloud configuration. Consider the replication methods and topologies that the storage appliances use in your disaster planning. Verify that the replication methods meet your recovery time objectives for resuming business systems and processes after site or hardware failures.

EMC Data Domain supports disaster recovery replication methods and topologies, including:

- Directory Replication
- Managed File Replication
- MTree Replication
- Collection Replication

We recommend using MTree replication because this method is a good candidate for consistent point-in-time snapshots of data that is replicated to a secondary appliance.

The z/OS *Administration Guide* provides an example of preparing and testing EMC Data Domain appliances for disaster recovery.

## NetApp Appliance

CA Cloud Storage for System z can store Virtual Volume data using the NetApp disk storage appliance in a private cloud configuration. Consider the replication methods and topologies that the storage appliances use in your disaster planning. Verify that the methods meet your recovery time objectives for resuming business services and processes after site or hardware failures.

NetApp supports replication methods and topologies for disaster recovery, including:

- Synchronous and asynchronous replication
- Snapshot Management including
  - SnapRestore
  - SnapMirror
  - SnapVault

We recommend using SnapMirror because it lets you schedule Virtual Volume data replication to the secondary appliance. This method also provides point-in-time data snapshots as the data replicates to the secondary appliance.

For information about bidirectional replication and an example of preparing and testing NetApp appliances for disaster recovery, see the z/OS *Administration Guide*.

## Riverbed SteelStore Appliance

CA Cloud Storage for System z can store Virtual Volume data using the Riverbed SteelStore disk storage appliance in a public cloud configuration. The Riverbed SteelStore appliance is typically configured to replicate files to your cloud provider. You can optionally configure SteelStore to perform an asynchronous replication to a dedicated secondary appliance at a remote location. This configuration is known as a warm site configuration.

If you do not replicate files to a secondary appliance at the disaster recovery site, cold site recovery of the Riverbed SteelStore appliance is required. You acquire and install a new appliance, import the configuration data, then repopulate by downloading backups from the cloud provider.

A warm site recovery configuration uses two SteelStore appliances. The primary site replicates to both the cloud and to a secondary appliance at the disaster recovery site. This method improves the recovery time by eliminating delays that are associated with installing and repopulating the appliance.

We recommended that you use two appliances in a warm site configuration.

See disaster recovery processes and procedures in the z/OS *Administration Guide*.