

Are You Leveraging the Latest DFSMSHsm Enhancements? Orlando 2026

Robert Gensler
STSM DFSMSHsm Product Owner
rgensle@us.ibm.com

Agenda

z/OS 3.2 and continuous delivery update

Cloud storage update

- *S3 support using direct to cloud and TCT*
 - *Enabling advanced cloud object storage capabilities (object lock, tiering, etc.)*
- *Multiple endpoint support for TS7700 advanced object store*
- *Recovery site support*
- *TLS 1.3 support*

Summary

For all

Additional Granularity for FIXCDS DISPLAY command ([ZOS-I-2518](#))

Expanded support for system symbols in ARCCMDxx ([ZOS-I-3281](#))

HOSTID support for HSEND command ([ZOS-I-1203](#))

Require a valid backup before performing migration for extent reduction ([ZOS-I-2268](#))

Fast Replication

Avoid FlashCopy failures during MIGRATE MOVE of large data sets in Metro Mirror environments (OA65984)

Improvements to FRBACKUP error handling behavior ([ZOS-I-1304](#)/ [ZOS-I-1930](#)) – coming soon with OA68845 on z/OS 3.2 and z/OS 3.1

Cloud storage support

Command migration honors management class L2 days non-usage (OA68200) – z/OS 3.2 ([ZOS-I-3555](#))

Allow HSM to configure the type of container that will create in the cloud for TCT (OA66198) ([ZOS-I-3666](#))

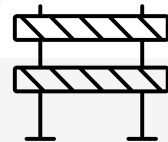


Improve TCT resiliency when using TS7700 Advanced object store (OA66197) ([ZOS-I-3602](#))

Improve Audit processing when using transparent cloud tiering (OA66772) ([ZOS-I-3959](#))

FIXCDS command enhancement

Grant storage admins access to display CDS records while still restricting access to potentially destructive commands

- If the `STGADMIN.ARC.FIXCDS.DISPLAY` facility class profile is not defined, then access to all `FIXCDS` commands are checked against `STGADMIN.ARC.FIXCDS`.
- If the `STGADMIN.ARC.FIXCDS.DISPLAY` facility class profile is defined, then access to the `FIXCDS DISPLAY` command is checked against `STGADMIN.ARC.FIXCDS.DISPLAY` and the rest of the `FIXCDS` commands are checked against `STGADMIN.ARC.FIXCDS`

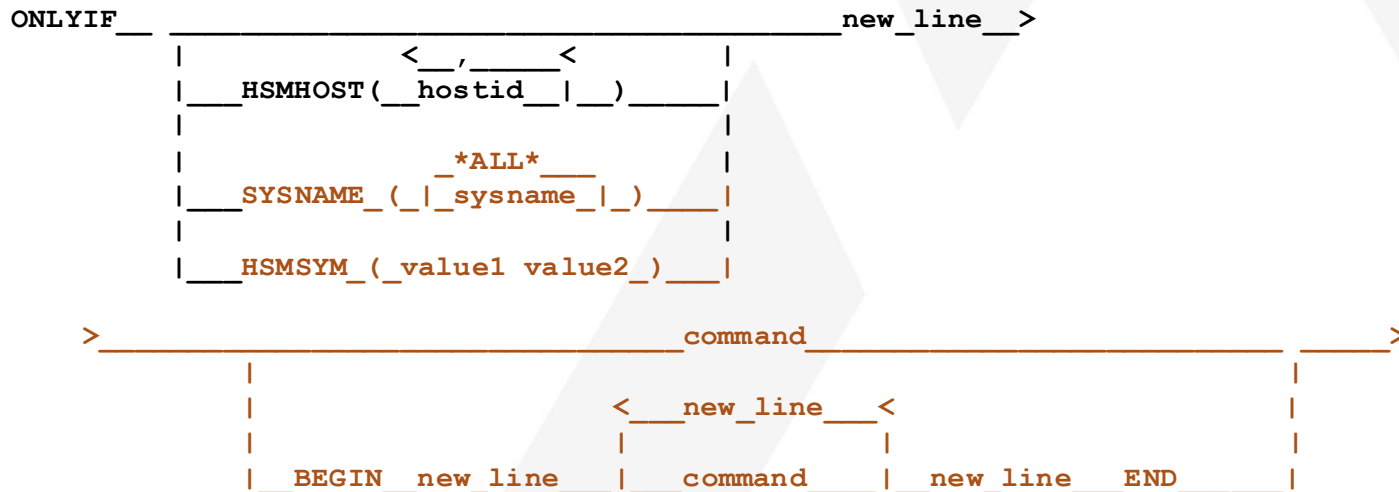


`STGADMIN.ARC.FIXCDS.DISPLAY`

`STGADMIN.ARC.FIXCDS`

ONLYIF updates

Use system name and user defined symbols in your DFSMSHsm configuration making it easier to share a single ARCCMDxx member among different DFSMSHsm servers in different environments



```

ONLYIF SYSNAME(*ALL*)
BEGIN
  /* options common to all environments */
END
ONLYIF HSMSYM(&HSMSYS PROD)
BEGIN
  SETSYS HSMPLEXNAME(PROD)
  SETSYS .....
END
ONLYIF HSMSYM(&HSMSYS DR)
  SETSYS HSMPLEXNAME(DR)
  
```

HSEND command enhancement

Leverage the new HOST parameter on the HSENDCMD command to direct requests to different DFSMSHsm servers (main or auxiliary) on a system.

New Syntax:

```
><---HSENDCMD ----- COMMAND---><
|---- NOWAIT ----|
|----- WAIT -----| | --- HOST(hostid) ---- |
```

Require valid backup before extent reduction

Ensure DFSMSHsm has a valid backup before performing extent reduction

PATCH .MGCB.+168 BIT(....1...), disabled by default

ARC0734I:ACTION= REDUCED FRVOL=volser1 TOVOL=volser2 | TRACKS=tracks |RC=45, REASON= 117,
AGE=days, DSN=dsname

Recommend periodically run report on which data sets are eligible for extent reduction but skipped because there is no valid backup and evaluate backup policy for those data sets

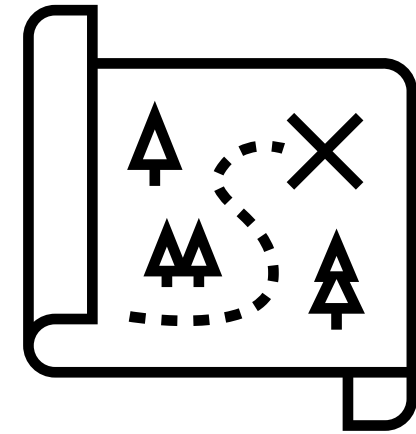
- *Set either SETSYS ACTLOGMSGLVL or PATCH .MGCB.+26 x'FF' to enable ARC0734I messages for data sets that do not meet selection criteria*
- *Enable the following patch to cut an FSR for ARC0734I errors in data set eligibility checking during primary space management.*

PATCH .MGCB.+EF BITS(... 1....)

MIGRATE MOVE

Move SMS managed data sets from one set of disk volumes to another. Extends storage tiering solution built on class transition attributes in management class

MOVE parameter supported on the **MIGRATE VOLUME**, **STORAGEGROUP**, and **DATASETNAME** command



Use the management class serialization error exit to enable movement of data sets that are open by certain applications.

DFSMSdss will call **Db2**, **CICS**, **zFS** or a **user exit** to close open data set, FlashCopy it, and reopen it.

MIGRATE MOVE and Thin Provisioning - OA65984

Avoid FlashCopy failures when moving (COPY w/ DELETE) of large data sets in Metro Mirror environments

By default, DFSMShsm requests that DFSMSdss **preallocate the target extent space during a FlashCopy establish** on thin-provisioned (ESE) volumes.

- Guarantees that the FlashCopy background copy will not fail due to the extent pool running out of physical capacity, but it does elongate the time it takes for the FlashCopy establish to complete.
- Assumption is you are over provisioned, but if you are not, it is safe to set...
PATCH .MGCB.+168 BITS(.....0.)

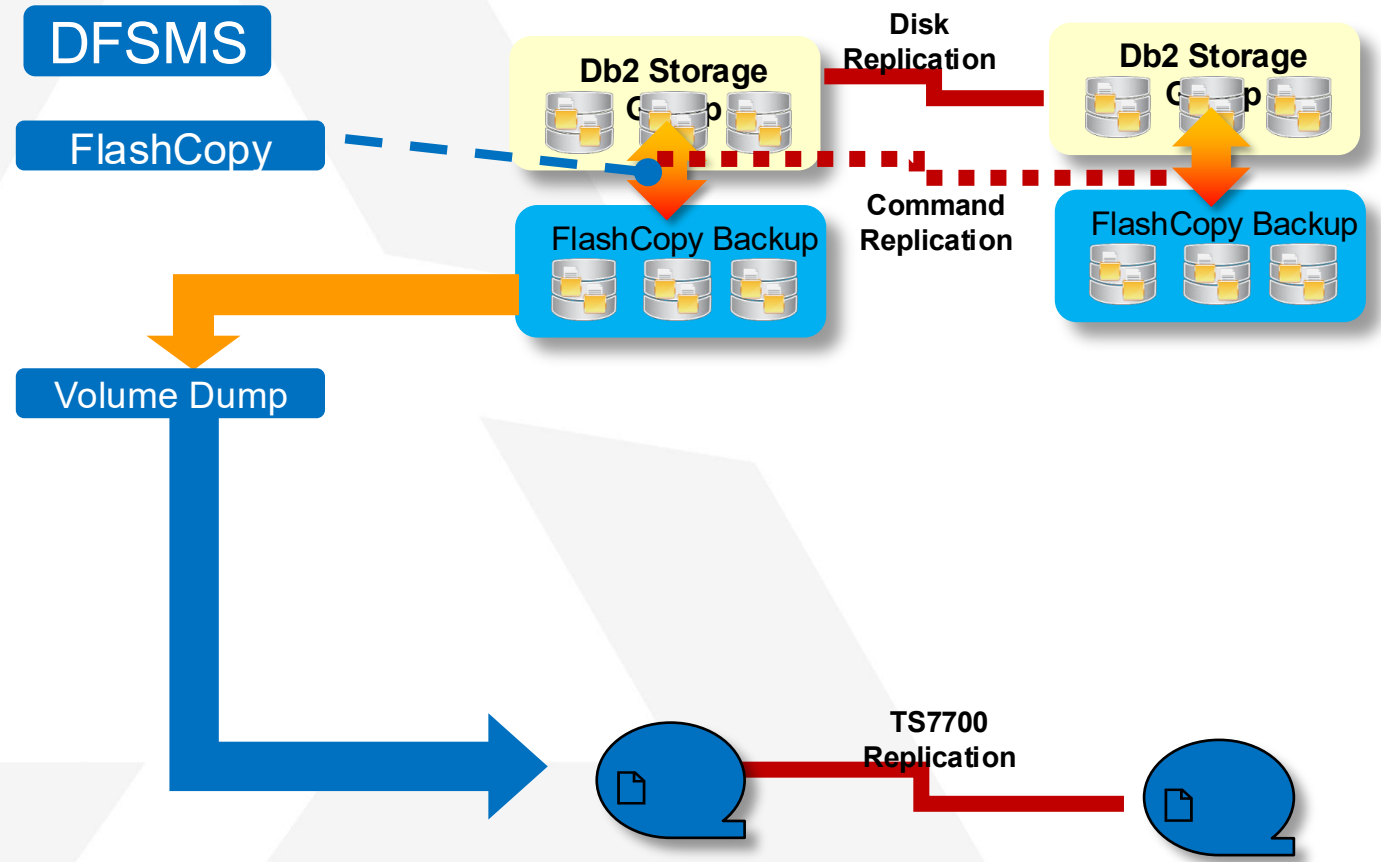
Db2 BACKUP and RESTORE SYSTEM

DFSMSHsm managed PiT backups of entire Db2 for z/OS subsystems on disk, tape and cloud

Integration with Db2 for z/OS to provide transactionally consistent recoveries

Recover entire copy pool, individual volumes and individual data sets

Data set recovery not supported for backups to cloud



Fast replication failure notification - OA68845

Db2 for z/OS subsystem-wide lock (SYSPITR) held while HSM performs an FRBACKUP

If one of the volume backups fail, HSM cleans up target volumes from previously successful backups in the copy pool

May cause other Db2 utilities to timeout with ABEND222

New "FRBACKUP failure notification" will return to Db2 utilities after failure and cleanup target volumes in the background

To enable, issue the following PATCH command in the DFSMSHsm startup parmlib member ARCCMDxx or dynamically after DFSMSHsm has started:

```
PATCH .FRGCB.+80 BITS (. . . . 1 . . . .)
```

Fast replication failure notification - OA68845

Expect an ARC1857I when an FRBACKUP fails and FlashCopy targets need to be re-initialized so they can be used for a subsequent FRBACKUP request

**ARC1857I FAST REPLICATION BACKUP HAS FAILED FOR COPY POOL *cpname*
VERSION *vern*. APPLICABLE BACKUP VOLUMES WILL BE INITIALIZED**

When fast replication backup to a disk backup version fails, DFSMSHsm may initialize the applicable backup volumes as part of the cleanup process to ensure the backup volumes are ready for use by subsequent fast replication backups

Backup volume initialization can be performed when the failed backup version is in NOCOPY mode or when VTOC tracks are in FlashCopy relationships.

Standard ARC1802I message is issued when cleanup has completed.

Fast replication failure notification - OA68845

Expect an ARC1625I, REAS=29 when another FRBACKUP is retried if a failed FRBACKUP of the same cypool is still being cleaned up

ARC1625I FAST REPLICATION COMMAND REJECTED, REAS=reason_code

Explanation

29

The specified copy pool is being processed by DFSMSHsm for a prior request. The previous FRBACKUP request has failed, and the associated backup volumes are being initialized.

Programmer response

29

The FRBACKUP COPYPOOL command can be reissued after DFSMSHsm has completed the backup volume initialization for the copy pool. ARC1802I indicates fast replication backup has completed.



CLOUD STORAGE ENHANCEMENTS

Why use cloud storage

- Transparent cloud tiering provides you a more predicable cost of data management
 - Cost is fixed per data set and is no longer tied to the size of the data set
- Cloud storage in general simplifies management
 - Once data is written it only needs to be touched when read or expired.
 - Unlike virtual tape solutions, DFSMSHsm no longer continually rewrites data from tape to tape as the percentage of valid data on a tape. No more recycle!

Credential and configuration management

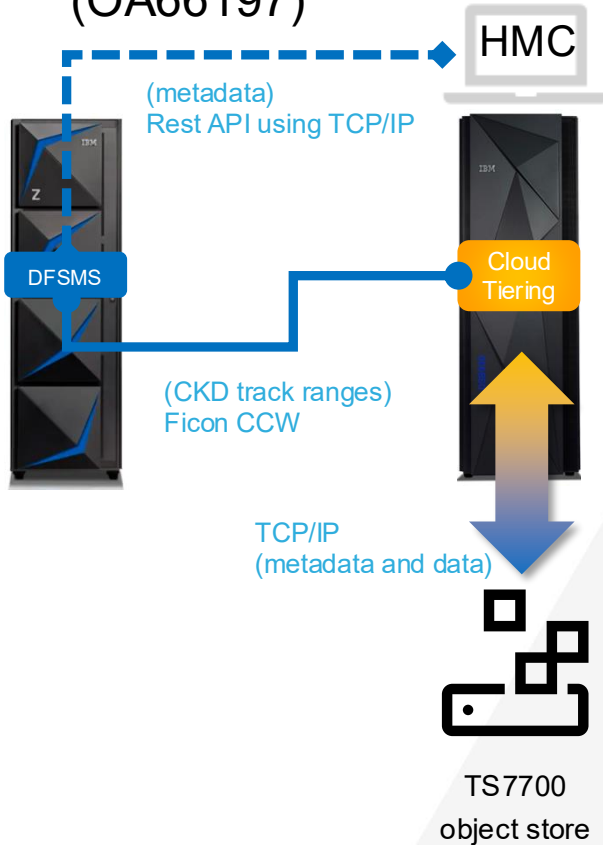
DFSMSHsm now fully leverages DFSMSdfp's cloud data access component

- ***credential and configuration management of cloud object storage including TS7700 advanced object store***
- ***APIs for cloud object storage***

SMS network configuration	CDA configuration file	CDA parameter
Endpoint	provider.json	"host": ... List of preferred and backup hosts
Port Number	provider.json	"port": ...
SSL Version	provider.json	"sslVersion": ...
SSL Key	provider.json	"sslKey" ...
Provider	provider.json	"authentication": { "model": ... },
Identity	Managed using GDKAUTP ISPF panel	---

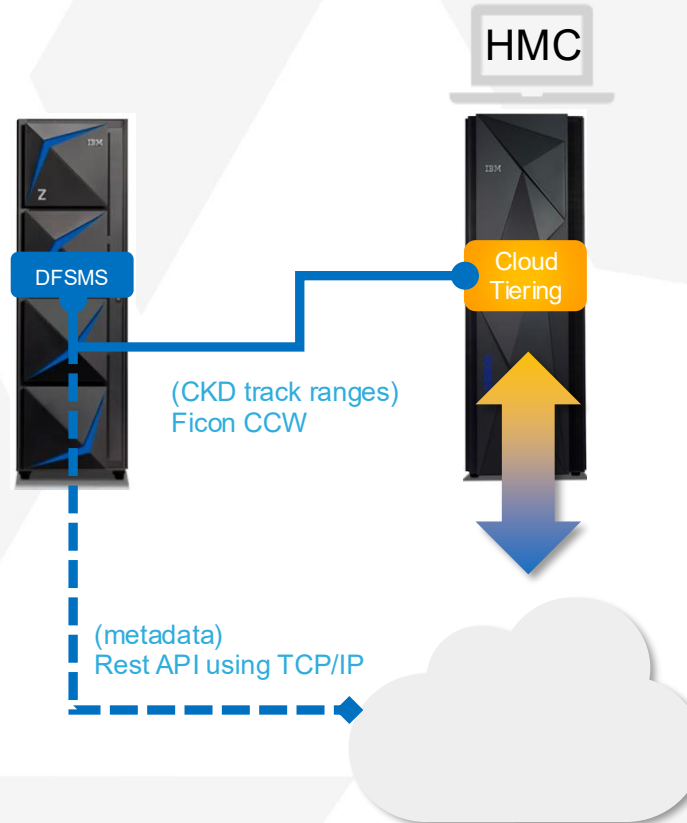
Supported cloud storage configurations

TCT using TS7700 advanced object store (OA66197)



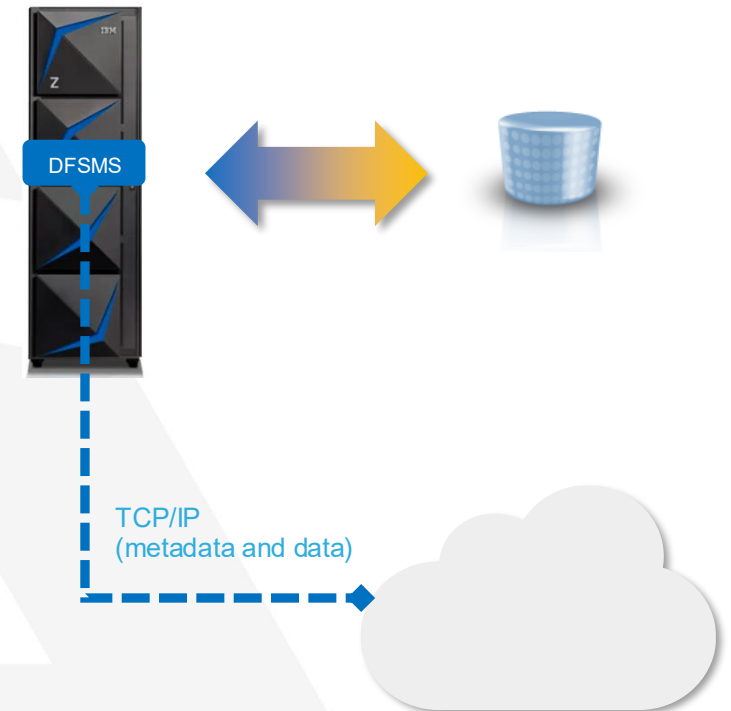
MIPs reduction / Requires TCT / Compression done by DS8000 / No more recycle

Cloud object storage using TCT ((OA66198)



MIPs reduction / Requires TCT / Compression done by DS8000 (10.11.30.0) / No more recycle

Cloud object storage directly without using TCT (OA66198)



No MIPs reduction / No TCT requirement / Compression in DFSMS/ No more recycle

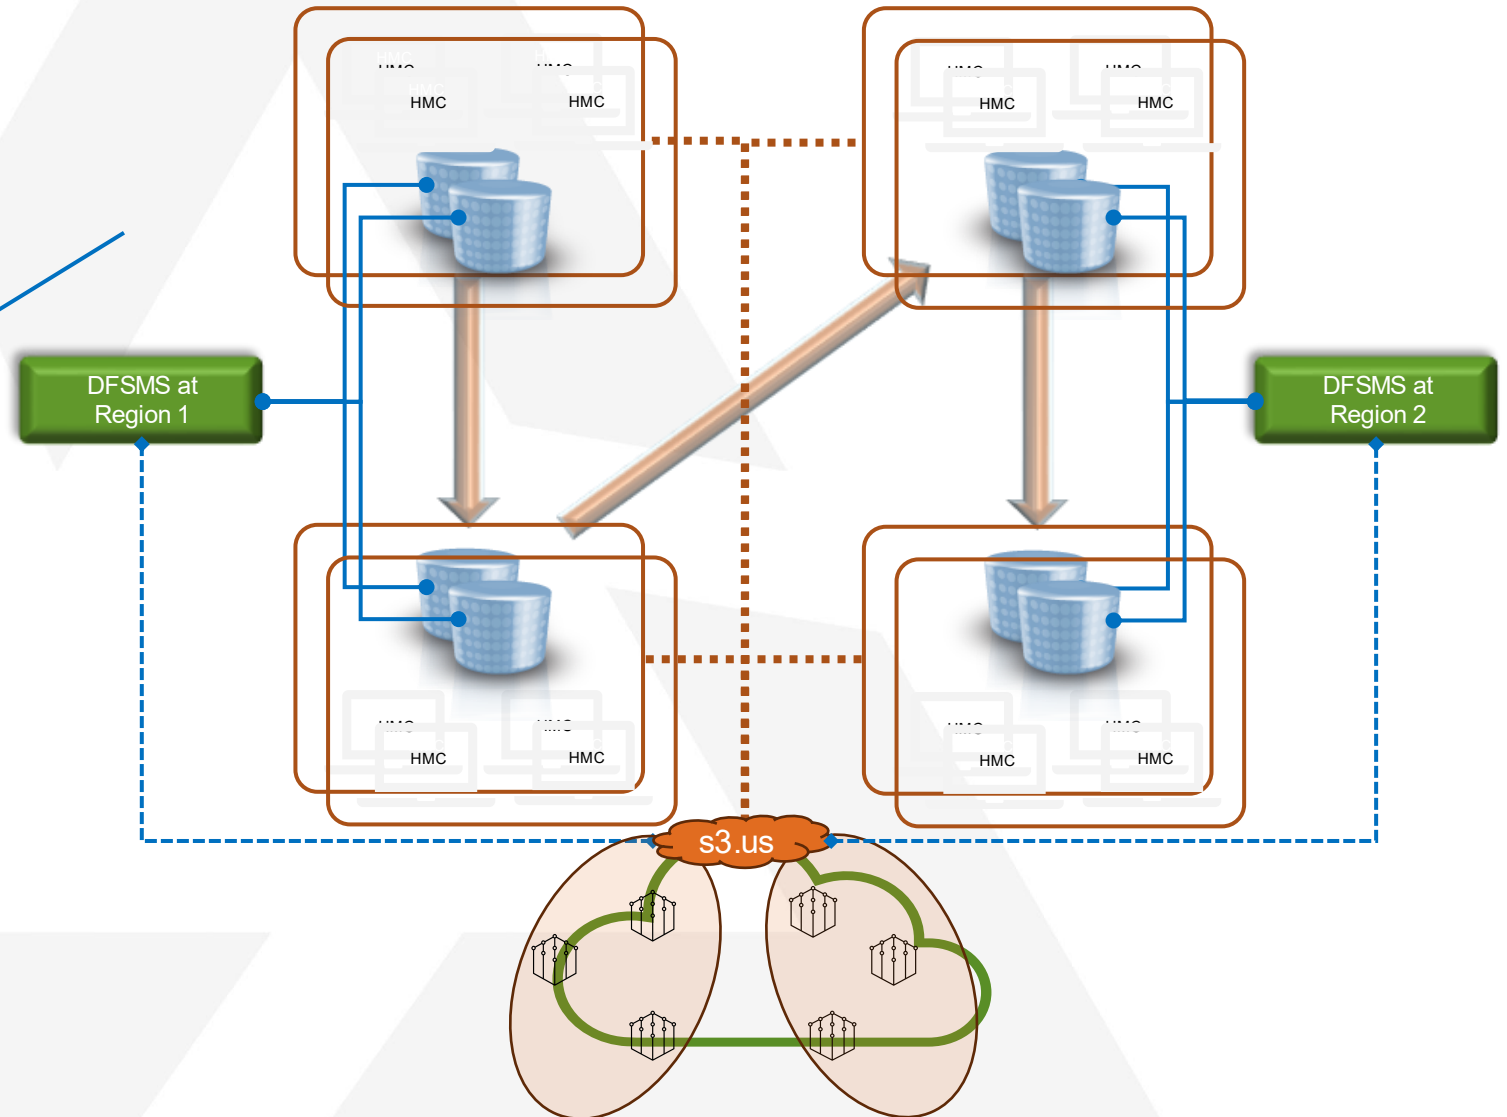
TCT using cloud object storage – OA66198

Improve overall scalability and resiliency of TCT when targeting S3 compatible object stores.

One configuration can cover multiple replication sites and regions

Credential and configuration managed by Cloud Data Access

All control plane commands issued directly to object store using CDA instead of using DS8000 HMC as a pass through.



Leverage object lock for DFSMShsm buckets

Object lock prevents permanent deletion of migration copies and dump versions in cloud object storage

- Update CREATEBUCKET schema to define lifecycle configuration rules that
 - create a bucket with object lock enabled
 - set a lifecycle configuration to delete after n number of days
- Update WRITE operations to sign payloads
- Deleted objects must be recovered in object store before recalling/recovering using DFSMShsm

Assigning storage class to containers

[ZOS-I-3666](#) Using CDA for cloud configuration, enables support for assigning an object storage class to the containers DFSMSHsm creates

- *Smart Tier – good for dynamic workloads*
 - *cloud storage provider automatically tiers data to optimize cost based on monthly usage patterns*
- *Standard – good for active workloads*
- *Vault – good for data accessed less than once a month*
- *Cold Vault – good for data accessed every 90days or less*

```

{
  "name": "CREATEBUCKET",
  "apiEndpoint": "<HOST><GDK_OBJECT_NAME>",
  "httpMethod": "PUT",
  "requestParameters": [
    {
      "mechanism": "HEADER",
      "descriptor": "x-amz-date: <DATE_ISO_8601>"
    },
    {
      "mechanism": "MESSAGE_BODY",
      "descriptor": "<GDK_SCHEMA>",
      "contentType": "application/xml",
      "schema": {
        "elementName": "CreateBucketConfiguration",
        "type": "object",
        "item": [
          {
            "elementName": "LocationConstraint",
            "type": "string",
            "content": "us-smart"
          }
        ]
      }
    }
  ]
}

```

<https://cloud.ibm.com/docs/cloud-object-storage?topic=cloud-object-storage-classes>

Automatically cleanup incomplete multipart upload

- Define lifecycle configuration rule to automatically cleanup incomplete multi-part uploads
- Ensures empty container deletion can successfully cleanup empty and inactive dump or migration containers

```
{
  "name": "CREATEBUCKET",
  "apiEndpoint": "<HOST><GDK_OBJECT_NAME>",
  "httpMethod": "PUT",
  "requestParameters": [
    {
      "mechanism": "HEADER",
      "descriptor": "x-amz-date: <DATE_ISO_8601>"
    },
    {
      "mechanism": "MESSAGE_BODY",
      "descriptor": "<GDK_SCHEMA>",
      "contentType": "application/xml",
      "schema": {
        "elementName": "AbortIncompleteMultipartUpload",
        "type": "object",
        "item": [
          {
            "elementName": "DaysAfterInitiation",
            "type": "integer",
            "content": "1"
          }
        ]
      }
    }
  ]
},
```

Automatically cleanup incomplete multipart upload – existing buckets

Use IBM cloud command line interface to set lifecycle configuration on existing buckets to ensure empty container deletion can successfully cleanup empty and inactive dump or migration containers

List existing buckets

```
ibmcloud cos buckets
```

For each existing bucket

- Get lifecycle configuration and store in json file
- Update json file to set abort incomplete multi part upload property
- Put lifecycle configuration for bucket

For each bucket

```
ibmcloud cos bucket-lifecycle-configuration-get \  
--bucket <bucket name> \  
--region <bucket_region> \  
--output json
```

```
ibmcloud cos bucket-lifecycle-configuration-put \  
--bucket <bucket name> \  
--lifecycle-configuration file:///lifecycle.json \  
--region <bucket_region> \  

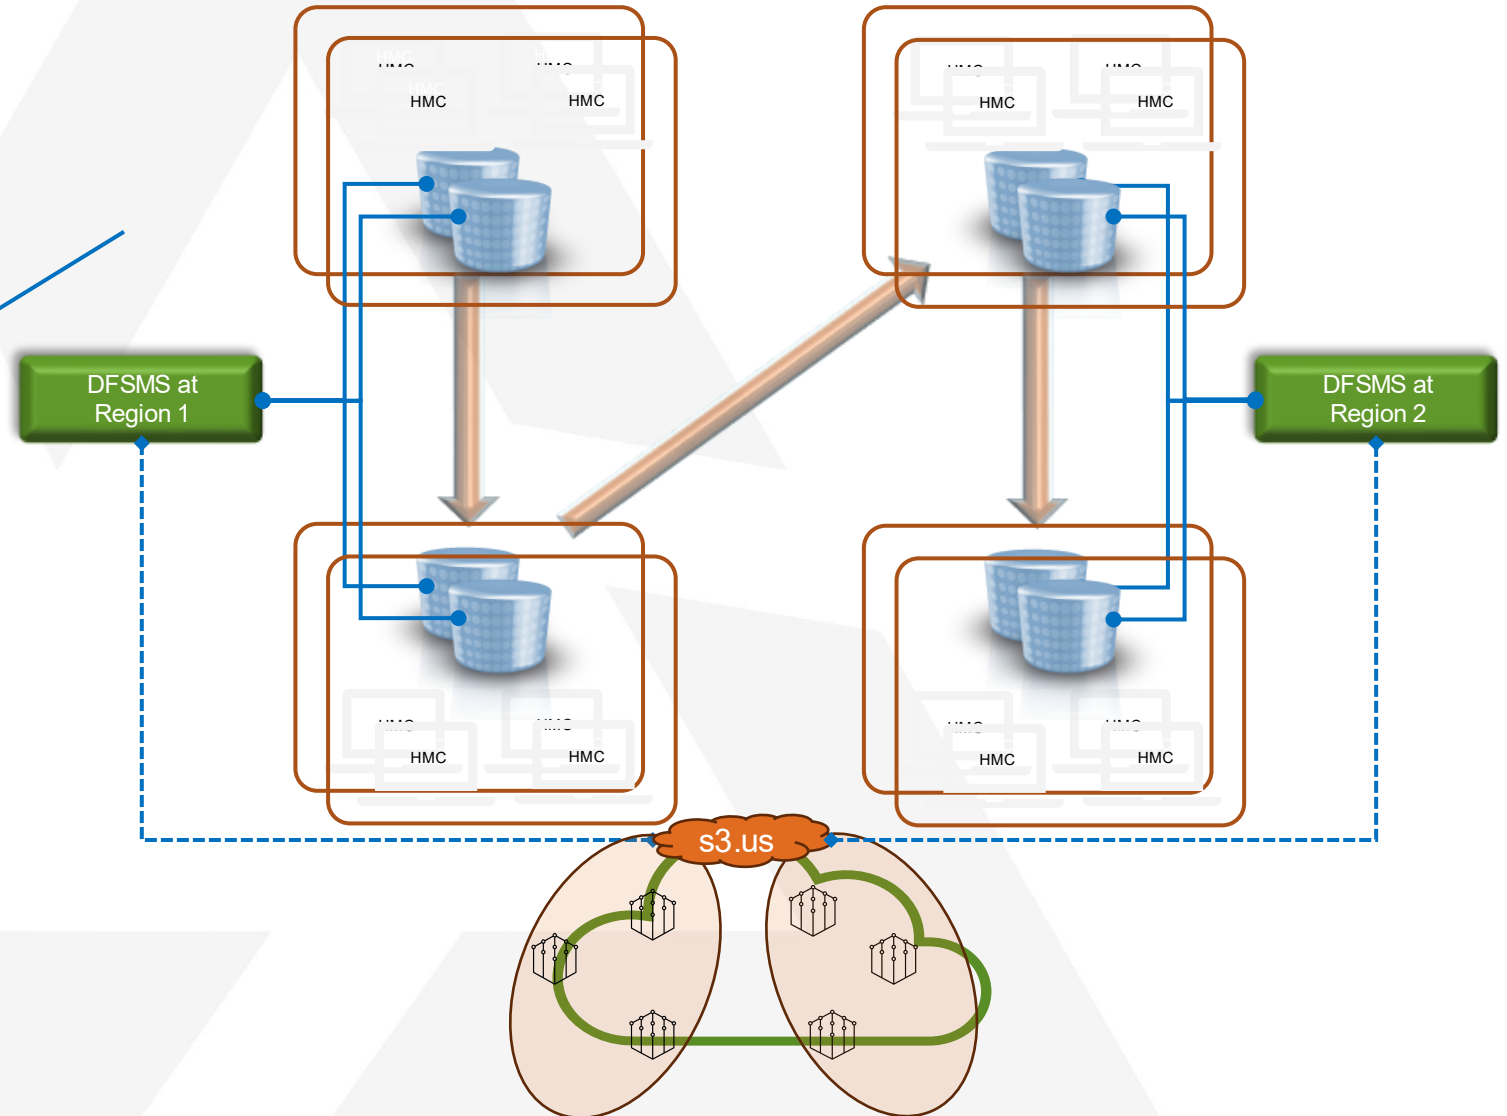
```

Directly using cloud object storage

Use cloud storage for migration and dump even if you do not have TCT capability

Automatically leverage regional or cross-regional bucket resiliency

DFSMSHsm calls DFSMSDss to read data from disk and the data is written directly to cloud object storage



Direct to cloud - Compression

Migration

Update to **SETSYS ZCOMPRESS** for zEDC compression support.

A new optional sub-parameter, **CLOUDMIGRATE(NO | YES)**, to indicate whether HSM should use zEDC compression during migration for direct cloud connections.

SETSYS ZCOMPRESS(CLOUDMIGRATE(YES))

QUERY SETSYS – updated ARC0340I

```
ARC0340I COMPACTION OPTIONS ARE: TAPEMIGRATION=NO, DASDMIGRATION=NO,
ARC0340I (CONT.) TAPEBACKUP=NO, DASDBACKUP=NO, TAPEHARDWARECOMPACT=NO,
ARC0340I (CONT.) ZCOMPRESS OPTIONS ARE: TAPEMIGRATE=NO, DASDMIGRATE=NO,
ARC0340I (CONT.) TAPEBACKUP=NO, DASDBACKUP=NO, CLOUDMIGRATE=YES
```

Compression is skipped for data sets already compressed, encrypted, or both

COMPACTPERCENT is ignored for compressed data sets migrated using the direct to cloud support. **Coming soon with OA67620**

Direct to cloud – Compression

DUMP and FRBACKUP DUMP

Existing sub-parameter, in DEFINE DUMPCLASS

(ZCOMPRESS(NO | YES)) indicate whether HSM should use zEDC compression

LIST DUMPCLASS(dc1001)

```

CLASS=DCL001      UNIT=***** REUSE=N  RESTORE=N  RESET=N  DISABLE=N  REQD=Y  AVAILABLE=Y
DAY=**  FREQ=007  RETPD=000030  TAPEEXPDT=*****  VTOCCOPIES=***  RECRST=YES  ZCOMP=YES
DISP=*****
HWCOMP=NO  ENCRYPT=NONE  ENCRYPT_TYPE=*****  ICOUNT=*****
ENC_KEY=*****
MINSTACK=001  STACK=001
DUDEVCAT=CLOUD  CLOUD=cloudname

```

Direct to cloud – DUMPIO support

DUMPIO: Specifying the DFSMSdss DASD I/O buffering technique

```
SETSYS DUMPIO(n,m)
```

n - buffering technique for BACKVOL command with the DUMP sub-parameter

m – buffering technique for migration

```
SETSYS DUMPIO(4)
```

Note: DUMPIO applies to migration and full volume dump to direct cloud connections. It is ignored for clouds configured to use TCT.

Keep in mind

Today if you activate a new SMS configuration, DFSMShsm automatically detects that configuration change on the next operation that needs to use cloud object storage

When updating the CDA provider.json file or when you update the credentials using GDKAUTHP, you must issue SETSYS CLOUD(NAME(...) CDAPROVIDER) on one hsm host to pickup that change to the configuration in the json files on disk.

SETSYS CLOUD(NAME(...) REFRESH) on the other hsm hosts

S3 support requirements

Available on z/OS 3.2 and z/OS 3.1 with PTF for OA66198

Includes both TCT and direct support

Conditioning PTF for OA66024 available on z/OS 3.1 and z/OS 2.5 for systems that do not have OA66198 installed but may be part of an HSMplex

HSMplex using CRQ/CVQ may be asked to recall or recover data migrated or dumped using the S3 support.

Fails gracefully

Requires CDA enablement

PATCH .MCVT.+5D7(1.....) / enable TCT FVD and CDA support */*

Disabled by default

Needs to be installed on all HSM's in your plex before enabling

Only read at DFSMSdsm startup

DSSXMMODE

When migrating, dumping, or restoring data using the direct to cloud support, DFSMSdss will always be loaded in its own address space.

Sample provider files

CDA DIRECT

```
{
  "name": "DFSMSIBMCOS",
  "host": "s3.us.cloud-object-storage.appdomain.cloud",
  "port": "443",
  "region": "us-standard",
  "httpMechanism": "HTTPS",
  "sslVersion": "TLSV12",
  "sslCiphers": "C013C014C027C028C02FC030",
  "encodeUrlChars": " &$@=;:+,?\"",
  "errorUrlChars": "\\{~}%`~>Ý~<#|",
  "authentication": {
    "model": "AWS4"
  },
  "supportedOperations": {
    {
      "name": "GETOBJECT",
      "apiEndpoint": "<HOST><GDK_OBJECT_NAME>",
      "httpMethod": "GET",
      "requestParameters": {
```

CDA TCT configuration

```
{
  "name": "DFSMSIBMCOS",
  "host": "s3.us.cloud-object-storage.appdomain.cloud",
  "port": "443",
  "region": "us-standard",
  "httpMechanism": "HTTPS",
  "tctType": "TCT",
  "sslVersion": "TLSV12",
  "sslCiphers": "C013C014C027C028C02FC030",
  "encodeUrlChars": " &$@=;:+,?\"",
  "errorUrlChars": "\\{~}%`~>Ý~<#|",
  "localCodePage": "IBM-1047",
  "remoteCodePage": "IS08859-1",
  "authentication": {
    "model": "AWS4"
  },
  "supportedOperations": {
    {
      "name": "GETOBJECT",
```

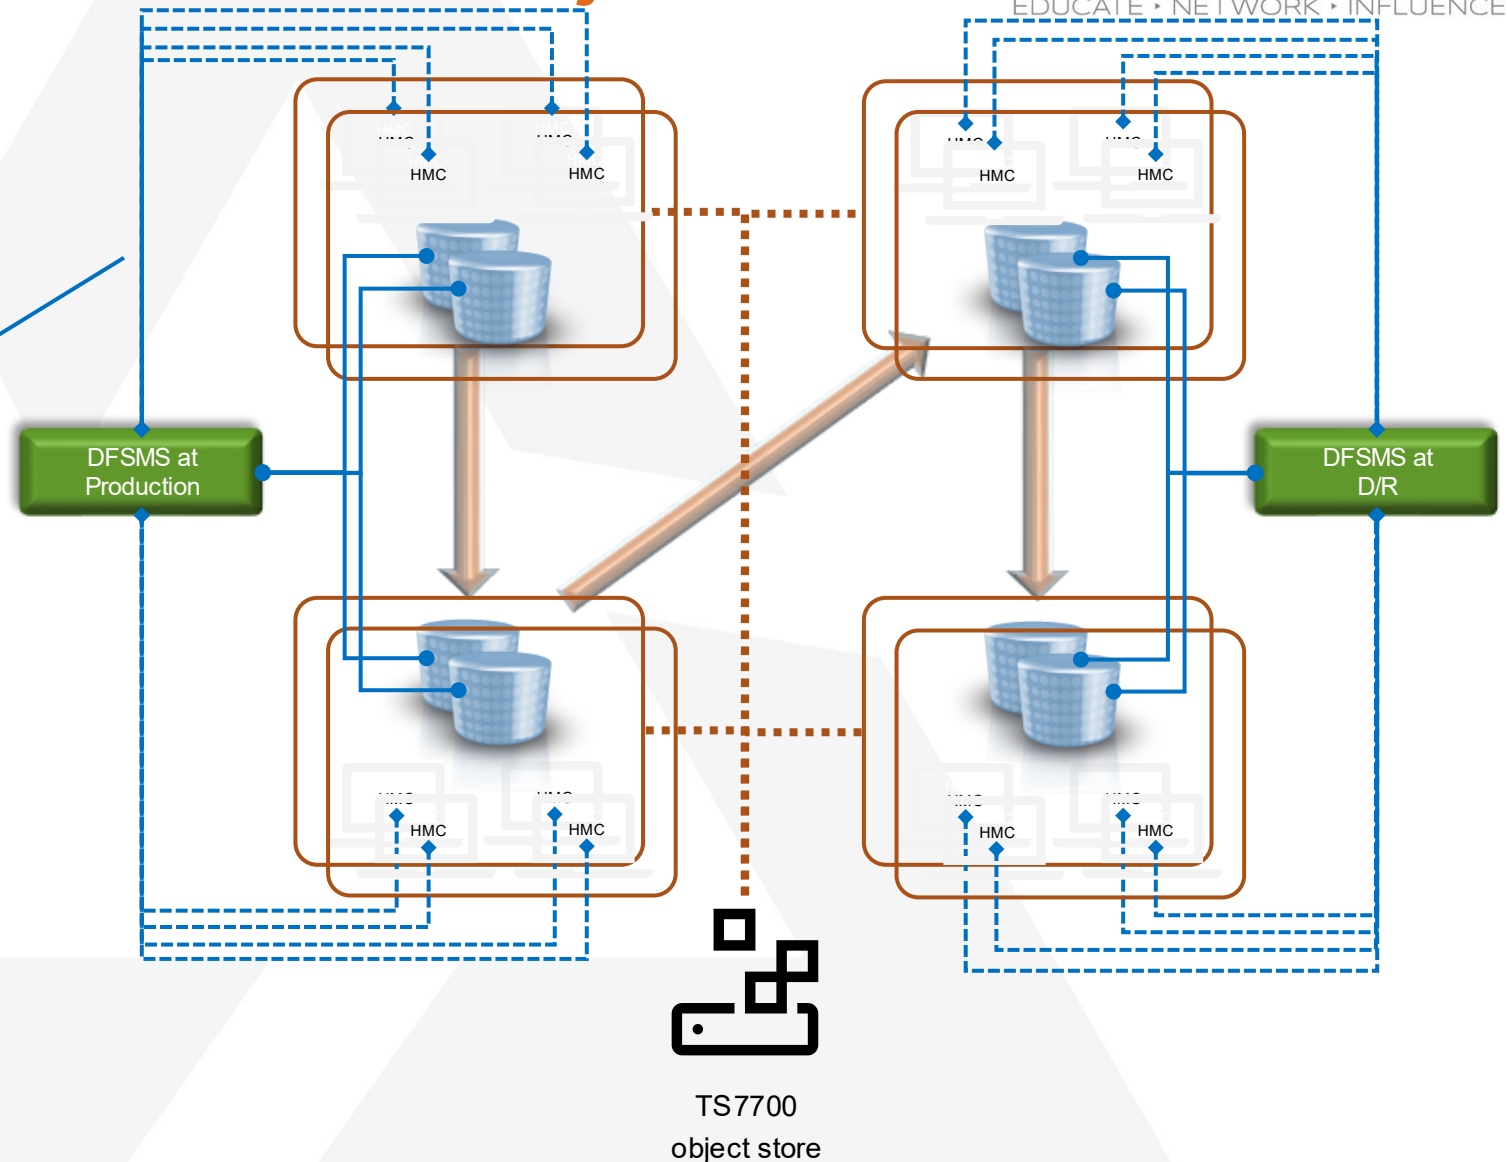

Multiple endpoint support for TS7700 Object Store

DFSMSHsm enables the use of the new DFSMSdfp Cloud Data Access component to configure TCT using TS7700 advanced object store

Endpoints can be defined as either preferred or backup endpoints so that one configuration can cover multiple sites and regions.

Metro Mirror primary and secondaries are weighted equally

Available on z/OS 3.2 and z/OS 3.1 with PTFs for OA66197



JSON provider for TS7700 advanced object store

```
{
  "host": {
    "preferred": [
      "ds8k1-hmc1-rs1-regionA",
      "ds8k1-hmc2-rs1-regionA",
      "ds8k1-hmc1-rs2-regionA",
      "ds8k1-hmc2-rs2-regionA"],
    "backup": [
      "ds8k1-hmc1-rs3-regionB",
      "ds8k1-hmc2-rs3-regionB",
      "ds8k1-hmc1-rs4-regionB",
      "ds8k1-hmc2-rs4-regionB"]
  }
  "sslKey": "*AUTH*/*",
  "port": "8452",
  "sslVersion": "TLSV13",
  "tctType": "TAPE-OBJECT",
  "tctIdentity": "ds8k-hmc-username"
}
```

Define new symbol TCTREGION symbol in IEASYMxx to control the usage of preferred and backup endpoints

PATCH MCVT.+24F BITS(.....1) default is OFF

Can use both PREFERRED and BACKUP, but start giving out based on TCTREGION

TCTREGION	MCVT+x24F	
N/S	Off (default)	Use only preferred endpoints
N/S	On	Use preferred then backup endpoints
PROD	Off (default)	Use only preferred endpoints
PROD	On	Use preferred then backup endpoints
DR	Off (default)	Use only backup endpoints
DR	On	Use backup then preferred endpoints

Multiple endpoint support – new QUERY parameter

NETWORK: Displaying cloud endpoint statistics known to DFSMSHsm

Explanation: NETWORK is an optional parameter requesting DFSMSHsm to output informational statistical data about known cloud endpoints.

Defaults: None

09.21.49 f dfhsm,query network

```
09.21.49 STC00141  ARC0300I  **OPER**  ISSUED===>QUERY NETWORK
09.21.49 STC00141  ARC0101I  QUERY NETWORK COMMAND STARTING ON HOST=2
09.21.49 STC00141  ARC1590I  CLOUD cloudname ENDPOINTS AND STATISITCS
09.21.49 STC00141  ARC1591I  ENDPOINT=https://ep1.ibm.com,
ARC1591I (CONT.) PORT=8452,
09.21.49 STC00141  ARC1591I  (CONT.) PREFERENCE=PREFERRED, STATUS=UNKNOWN, ACTIVE CONNECTIONS=000000000000005
09.21.49 STC00141  ARC1591I
ARC1591I (CONT.) ENDPOINT=https://ep2.ibm.com,
ARC1591I (CONT.) PORT=8452,
09.21.49 STC00141  ARC1591I  (CONT.) PREFERENCE=PREFERRED, STATUS=UNKNOWN, ACTIVE CONNECTIONS=000000000000000
09.21.49 STC00141  ARC0101I  QUERY NETWORK COMMAND COMPLETED ON HOST=2
```

Use SMF119 to measure the connections to the endpoints

Endpoint	Connections
Endpoint 1	1000
Endpoint 2	100
Endpoint 3	100
Endpoint 4	100

TS7700 Advanced Object Store - Compression

Specify ENABLE(TCTCOMPRESSION)
in DEVSUPxx

TCTCOMPRESSION

Enables or disables the transparent
cloud tiering compression feature.

ENABLE(TCTCOMPRESSION) |
DISABLE(TCTCOMPRESSION)

```
DISPLAY DEVSUP  
IEA253I DISPLAY DEVSUP Start of Report 375  
...  
375 00000010  ENABLE (TCTCOMPRESSION)  
...  
375 00000010  DISPLAY DEVSUP End of Report
```

Note: COMPACTPERCENT is *supported for TCT to TS7700 Advanced Object Store*

TS7700 multiple endpoint support - Installation and Enablement

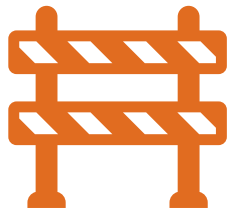
Available on z/OS 3.2 and z/OS 3.1 with ++PTF for OA66197

Enable support in DFSMSHsm ARCCMDxx

```
PATCH .MCVT.+5D7(..1.....) /* enable multiple endpoint support */
```

Disabled by default

Only read at DFSMSHsm startup



Requires CDA enablement

```
PATCH .MCVT.+5D7(1.....) /* enable CDA support */
```

Disabled by default

Needs to be installed on all HSM's in your plex before enabling

Only read at DFSMSHsm startup

<https://public.dhe.ibm.com/eserver/zseries/zos/DFSMS/HSM/OA66197>

TLS 1.3 support

Available on z/OS 3.2 and z/OS 3.1 with PTF for OA66412

```
"region": "us-standard",  
"httpMechanism": "HTTPS",  
"sslVersion": "TLSV13",  
"sslCiphers": "C013C014C027C028C02FC030",  
"encodeUrlChars": " &$@=::+.?"
```



CONFIGURATION GOTCHAS

PLEXNAME

- If you are going to use cloud storage you must set unique PLEXNAME
- **Unique containers will be created automatically for Cloud dump objects every 92 days**
SYSZARC.HSMplexname.MIG.yyyyddd.suffix
- *Where*
 - **SYSZARC** Constant
 - **HSMplexname** HSMplex name, 'PLEX0' is the default suffix
Ensure that you have unique HSMplex names across your sysplexes!
 - **Function** Currently either MIG or DMP
 - **yyyyddd** Four-digit year followed by Julian day
 - **suffix** Random suffix

Requirements to use CDA

- Configure CDA for the userID the DFSMShsm started task is running under (requires OMVS segment)
 1. gdkkeyf.json
 1. Create gdk/ directory in the DFSMShsm userIDs home directory and set the permissions to 770.
 2. Copy /usr/lpp/dfsms/gdk/samples/gdkkeyf.json to ~/gdk/gdkkeyf.json and set the permissions to 660.
 2. config.json
 1. Copy /usr/lpp/dfsms/gdk/samples/gdkconfig.json to ~/gdk/config.json and set the permissions to 660.
 3. provider.json
 1. Create a providers/ directory in the DFSMShsm userIDs home directory and set permissions to 770.
 2. Create your cloudname.json and set permissions to 660. There are several samples provided in /usr/lpp/dfsms/gdk/samples/providers
- Note permissions are different than what CDA documents because we are not setting this up for ourselves but for the userID hsm server runs under. The userID that is configuring CDA for DFSMShsm must be part of group that has write access to the hsm userID home directory and the files created underneath it

[Terminal emulator needs host code page set to IBM 1047 not CP 037](#)

<https://www.ibm.com/docs/en/zos/3.1.0?topic=configuration-user-quick-start>

Associating DFSMSdss cross memory server with hsm ID

```
SETR GENERIC(STARTED)
RDEFINE STARTED ARC*.* STDATA(USER(DFHSM))
SETR RACLIST(STARTED) REFRESH
```

```
CNZ4106I 10.53.35 DISPLAY ACTIVITY 357
  JOBS      M/S      TS USERS      SYSAS      INITS      ACTIVE/MAX VTAM      OAS
00018      00249      00019      00111      00230      00019/00150      00253
  ARC1MIGR ARC1MIGR IEFPROC  NSW *    A=01F8      PER=NO      SMC=000
  PGN=N/A   DMN=N/A   AFF=NONE
  CT=000.011S  ET=299.373S
  WUID=STC48584 USERID=DFHSM
  WKL=W STC      SCL=S STCDEF P=1
  RGP=N/A      SRVR=NO QSC=NO
  ADDR SPACE ASTE=7FCE1E00
```

<https://www.ibm.com/docs/en/zos/3.2.0?topic=task-configuring-dfsmsdsm-invoke-dfsmsdss-as-started>

```
ADR049E (001)-STEND(01), 2026.022 10:40:54 DFSMSDSS FUNCTION TASK ABEND RECOVERY ROUTINE WAS ENTERED. USER ABEND
CODE=4093 REASON CODE=00000090
```

```
ARC1001I dsname MIGRATE FAILED, RC=0068, REAS=0049
ARC1268I MIGRATION FAILED DUE TO AN ERROR IN DFSMSDSS
```

<https://www.ibm.com/docs/en/zos/3.2.0?topic=codes-u4093-xffd>

Requirements to use CDA – Security Admin

Follow the System administrator configuration quick-start steps 1 and 2 to:

1. Configure the CSFKEYS general resource class to protect the keylabel created by CDA

/ Define a generic label with UACC(NONE) so default access is NONE */*

RDEFINE CSFKEYS GDK.** UACC(NONE) ICSF(SYMCPACFWRAP(YES) SYMCPACFRET(YES))

RDEFINE CSFKEYS ARC. UACC(NONE) ICSF(SYMCPACFWRAP(YES) SYMCPACFRET(YES))** ← DFSMSHsm specific required when using TS7700 advanced object store

/ Define a generic label specific to the CDAUSER */*

RDEFINE CSFKEYS GDK.DFHSM.** UACC(NONE) ICSF(SYMCPACFWRAP(YES) SYMCPACFRET(YES))

RDEFINE CSFKEYS ARC.<cloudname> UACC(NONE) ICSF(SYMCPACFWRAP(YES) SYMCPACFRET(YES)) ← DFSMSHsm specific ← DFSMSHsm specific required when using TS7700 advanced object store

/ Permit the CDAUSER to their keylabels */*

PERMIT GDK.DFHSM.** CLASS(CSFKEYS) ID(DFHSM) ACCESS(UPDATE)

PERMIT ARC.<cloudname> CLASS(CSFKEYS) ID(DFHSM) ACCESS(UPDATE) ← DFSMSHsm specific

SETROPTS RACLIST(CSFKEYS) CLASSACT(CSFKEYS) REFRESH

2. PERMIT read access to the following CSFSERV class resources
CSFKGN, CSFRNGL, CSFKRD, CSFKRC2, CSFOWH (CSFKTB2 if ALLOW_NO_CEX)
CSFRNG and CSFIQA

<https://www.ibm.com/docs/en/zos/3.1.0?topic=configuration-system-administrator-quick-start>

Using CDA credential store without crypto express

If the system does not have any Crypto Express configured as CCA co-processor you may see an error when saving the resource authorization

ERROR: encryptKeys: Unable to generate a key. CSNBKGN rc: 12, rsn:0000

Set allow-no-CEX to true in config.json, then CDA will store the key in ICSF in the clear

If you don't have CEX adapters and set allow-no-CEX to true but still see this CSNBKGN rc: 12, rsn:0000

Check the format of your CKDS, it could be using fixed length record format (LRECL=252) check your CKDS with LISTCAT

ARC1584I SETSYS CLOUD – NO CRYPTO EXPRESS CARD IS INSTALLED

To request allowing encryption key to be stored in the clear in the ICSF CKDS issued the following PATCH commend:

```
PATCH .MCVT.+597 BITS (. . . . . 1 .)
```

The patch to revert this option is:

```
PATCH .MCVT.+597 BITS (. . . . . 0 .)
```

```
contents of config.json
{
  ...
  "allow-no-CEX": true,
  ...
}
```

ICSF configuration gotchas

Migrations and recalls work on one host but not another in the plex

ARC1583I ICSF INTERNAL ERROR

PDA analysis shows an ICSF RC8 RSN x'271C' calling ICSF

A key label was supplied for a key identifier parameter. This label is the label of a key in the in-storage CKDS or PKDS. A key record with that label (and the specific type if required by the ICSF callable service) could not be found.

ICSF configuration had sharing defaulted to

`SYSPLEXCKDS(NO,FAIL(NO))/SYSPLEXPADS(NO,FAIL(NO))`

Recommendation from ICSF development

`SYSPLEXCKDS(YES,FAIL(YES))/SYSPLEXPADS(YES,FAIL(YES))`

Display current values

----- ICSF - Installation Option Display - Row 31 to 37 of 37

COMMAND ==>

SCROLL ==> CSR

Active CKDS: SYS1.CSF.CSFCKDS

Active PKDS: SYS1.CSF.CSFPKDS

Active TKDS: SYS1.CSF.CSF TKDS

OPTION

CURRENT VALUE

...

SYSPLEXCKDS	Sysplex consistency for CKDS updates	YES, FAIL (YES)
SYSPLEXPKDS	Sysplex consistency for PKDS updates	YES, FAIL (YES)
SYSPLEXTKDS	Sysplex consistency for TKDS updates	YES, FAIL (YES)

ICSF will join the ICSF sysplex group SYSICSF and this system will participate in sysplex-wide consistency for CKDS data.

<https://www.ibm.com/docs/en/zos/3.1.0?topic=start-parameters-in-installation-options-data-set>

<https://www.ibm.com/docs/en/zos/3.1.0?topic=status-displaying-installation-options>



RECOVERY SITE SUPPORT

DFSMShsm recovery site support

DFSMShsm cloud storage recovery site support preserves access to migration cloud objects at a recovery site regardless of any HSM activity at the primary site.

Support is available on z/OS 3.2 and z/OS V3.1 with the PTFs for APAR OA63678

```
ONLYIF SYSNAME(*ALL*)
BEGIN
  /* options common to all environments */
END
ONLYIF HSMSYM(&HSMSYS PROD)
BEGIN
  SETSYS HSMPLEXNAME(PROD)
  SETSYS . . . .
END
ONLYIF HSMSYM(&HSMSYS DR)
  SETSYS HSMPLEXNAME(DR)
```

SETSIS MIGRATIONCLEANUPDAYS

```
> _____ >
|
|          |----10-----| |----30----| |-----3-----| |-----0-----|
|  MIGRATIONCLEANUPDAYS( __|_recalldays_|__|_statdays_|__|_reconnectdays_|_|_clouddays_|) _____|
```

For clouddays, substitute a decimal number from 0 to 999 for the number of days that DFSMSHsm should keep old cloud migration copies around after a modified version of the same data set has been remigrated to the cloud.

SETSIS default: 0

Before starting DR test, set clouddays to a non-zero value. Consider factors when determining what to set clouddays to

- *Duration of disaster recovery testing*
- *How far back in time you would realistically recover back to and resume business operations after a real disaster*
- *As clouddays gets larger so does the size of the MCDS, potentially. ~300bytes per data set that was recalled and remigrated.*

Functional Changes

New Migrations

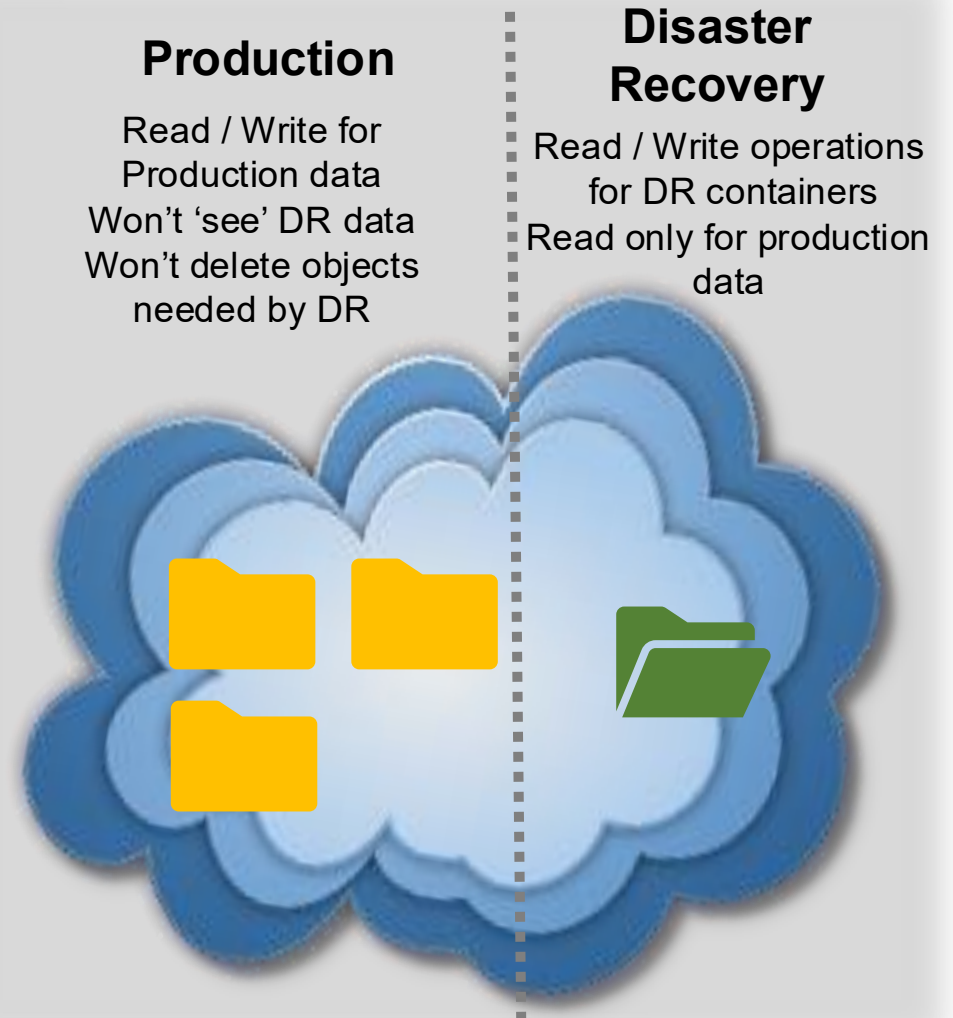
- *Since PLEXNAME is different between primary and DR hosts migrated copies go to different containers*

Re-migrations

- *Data sets eligible for reconnection are processed as they are today*
- *Clouddays zero/not specified*
 - *Cloud migration scratch and then re-migrate data set*
- *Clouddays greater than zero*
 - *Defer cloud migration scratch to secondary space management*
 - *Create retained MCDS records saving cloud information from prior migration*
 - *Create new MCDS records describing the new migration*

Recall

- *Primary HSMplex only sees data sets it migrated, it does not know about data sets migrated from recovery plex*
- *Recovery plex can recall data sets from primary HSMplex and data sets it has migrated as part of the D/R test*



Recommendation...

When not in a DR test window

PATCH .MGCB.+113 BITS(...0....) and set clouddays = 1 in MIGRATIONCLEANUPDAYS SETSYS

Turning the PATCH off and setting clouddays = 1 causes DFSMShsm to

- move the deletes of old migration copies out of primary space management completely.*
- gets actual delete of objects out of HDELETE path*

Deletion of old migration copies now occurs only during secondary space management and not during migration or hdelete

<https://www.ibm.com/support/pages/node/6468023>

New HSM Badge

z/OS DFSMS Academy - DFSMSHsm Explorer

DFSMSHsm Basics: An overview of storage/availability management, DFSMSHsm functions and architecture.

DFSMSHsm Strategy & What's New: A brief history of HSM. A slightly deeper dive into its purpose and more of how the functions work together as a solution to accomplish this purpose.

DFSMSHsm Control Data Sets: An overview of the Control Data Sets that DFSMSHsm uses to maintain an inventory of the data sets and volumes that it manages, which is fundamental to all HSM's functions.

<https://www.credly.com/org/ibm/badge/z-os-dfsms-academy-dfsmsHsm-explorer>



Experience more with IBM



Visit us at the IBM Booth #113

After a full day of technical sessions, take a break with us!

Connect with our experts, snap a photo with the z17 Plexi or the latest Telum II, and get an up-close look at our Spyre Accelerator.

Come back each day for fresh topics and demos at our expert stations.

Think 2026

Join 5000+ senior business and technology leaders who are seizing the AI revolution to unlock unprecedented growth and productivity at **Think 2026**.


Find out more information using the QR code below.



IBM Digital Asset Haven

IBM Digital Asset Haven is the operational backbone for financial institutions and regulated enterprises entering the digital asset economy.

Find out more information using the QR code below.



Your feedback is important!

Submit a session evaluation for each session you attend:

www.share.org/evaluation

