

Washington Systems Center - Storage

TDMF and zDMF - Migration Product Updates

- Rebecca Levesque – (21st Century Software CEO/President)
- Richard Pace – (21st Century Software TDMF/zDMF Lab)
- John Hogan – (21st Century Software TDMF/zDMF Lab)

Accelerate with IBM Storage Webinars

The Free IBM Storage Technical Webinar Series Continues in 2019...

Washington Systems Center – Storage experts cover a variety of technical topics.

Audience: Clients who have or are considering acquiring IBM Storage solutions. Business Partners and IBMers are also welcome.

To automatically receive announcements of upcoming Accelerate with IBM Storage webinars, Clients, Business Partners and IBMers are welcome to send an email request to accelerate-join@hursley.ibm.com.

Located in the Accelerate with IBM Storage Blog:
<https://www.ibm.com/developerworks/mydeveloperworks/blogs/accelerate/?lang=en>

Also, check out the WSC YouTube Channel here:
https://www.youtube.com/channel/UCNuks0go01_ZrVVF1jgOD6Q

2019 Upcoming Webinars:

July 30 – IBM Cyber Security Update

Register Here: <https://ibm.webex.com/ibm/onstage/g.php?MTID=e698ab6a71e858efcd253f6650ba0b428>

August 6 – IBM Storage SAN b-type Extension: Native IP vs FCIP

Register Here: <https://ibm.webex.com/ibm/onstage/g.php?MTID=eba8b985837a2454480877deb0224114f>

August 22 - LinuxONE Servers and IBM Storage Synergies

Register Here: <https://ibm.webex.com/ibm/onstage/g.php?MTID=e246d7bbd6b0af257384b0e93c9032eec>



Business Problem to be Solved

Why customers need non-disruptive data migration

Infrastructure and Operations leadership are no longer "business as usual". It's about driving change and disruption through a product-aligned, business-embedded operating model that enables agility and innovation.

- How do you see your role?
- How are you supporting the business drivers within the IT infrastructure?
- What information are you using to make those decisions?



Source: Gartner Leadership Vision for 2019: Infrastructure and Operations Leader

What's the problem?

- Companies are spending \$1.2-\$2.5 billion annually due to unplanned application outages
- IT is held accountable to maintain availability while managing increasingly complex workloads
- There is no margin of error for outages and no tolerance for extended outages



What are our clients telling us?

Data security concerns
are driving the need to
be more granular...
volume-level
encryption isn't
enough anymore

A large international bank
needs a non-disruptive way to
get data to a second location....

They need any
easier way to do
DevOps

They have to have
easier options to
encrypt data without
an extended outage

TDMF and zDMF can
do all of this and
more....

Move mainframe data effectively with optimal application availability


TDMF and zDMF solutions for IBM Z provide local or global data migration for storage attached to IBM Z mainframes across multivendor environments with continuous business operations.

TDMF z/OS:

- Enables non-disruptive data migration at the volume (physical track) level
- OVA feature of TDMF provides access to replicated offline volume data
- z/VM migrations made simpler through non-disruptive movement

zDMF:

- Enables non-disruptive or minimally-disruptive data migration at the data set (logical) level
- Data Set level migration with encryption of data sets (pervasive encryption)



TDMF Overview

Highlighting recent enhancements

TDMF Data Migration Solutions Features

- **Non-disruptively SWAP** from old device to new device
- **Point in time volume copies** - for full volume and application backups; volume copies for testing
- **Offline Volume Access (OVA) facility** - access offline volume copies for testing; synchronized copy of a volume
- **Switch back facility** – Monitor swapped volumes for updates and have migration switched back to the original volumes

TDMF Data Migration Solutions Features

- **Vendor agnostic** - works with all the leading DASD manufacturers with awareness of individual replication technologies
- **Pacing** - allows control of the number of active volumes and how many I/Os TDMF can issue during a migration
- **TCP/IP Migrations (replication/copy)** - performed across long distances for Data Center/LPAR consolidation, DR sites
- **Hyperswap aware** - TDMF dynamically disables/enables Hyperswap via commands to migrate GDPS/PPRC volumes

TDMF Data Migration Solutions Features

- **GDPS/xDR proxy** - support on z/Linux running under z/VM using the z/VM Agent for TDMF
- **Migration between different sized volumes** - small to large/large to small (rules apply) with dynamic ICKDSF invocation
- **Use of latest DS8K security features** - prevent 'missing Agents', SPID Fence and Soft Fence
- **Non-disruptive migration of attached z/VM volumes** without a z/VM outage
- **Support for migrating z/Linux Guest volumes** under z/VM

Session Objectives

Identify recent TDMF software enhancements introduced in TDMF releases 5.2.0 through 5.8.0:

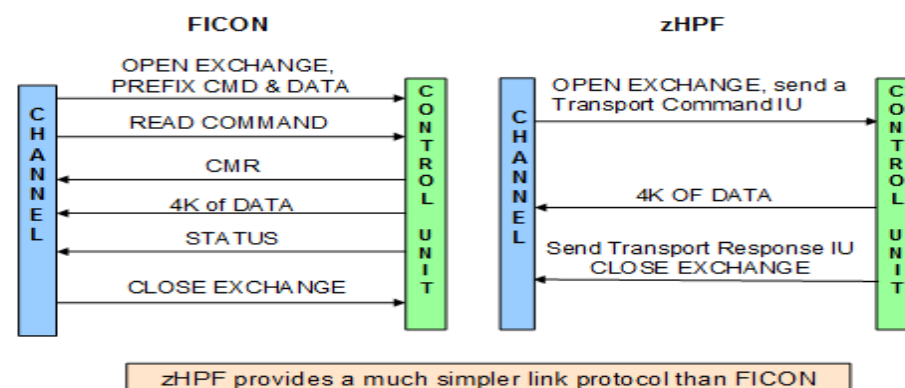
- Support for High Performance FICON (zHPF) data transfer protocol
- Support for Basic HyperSwap and GDPS/PPRC HyperSwap
- TDMF Virtual Storage Constraint Relief (VSCR) in support of migrations with a larger number of volumes, thereby reducing ECSA requirements, using common storage “above the bar”.
- Support for IOS Soft Fence & SPID Fence Feature
- Support for 1 TB EAV Devices
- Removal of 512 maximum volume limit per each TDMF Session
- Successful Session initialization time increased to 60 minutes
- Volume Swapping Improvements
- I/O Pacing Improvements
- zVM Agent

TDMF Support for High Performance FICON (zHPF) - 1 of 3

What is zHPF?

- zHPF is an extension to FICON architecture designed to improve the execution of small block I/O requests. zHPF streamlines the FICON architecture and reduces the overhead on the channel processors, control unit ports, switch ports, and links by improving the way channel programs are written and processed.
- zHPF improves upon FICON by providing a Transport Control Word (TCW) that facilitates the processing of an I/O request by the channel and the control unit. The TCW has a capability that enables multiple channel commands to be sent to the control unit as a single entity instead of being sent as separate commands as with FICON CCWs.

Link Protocol Comparison for a 4KB READ



TDMF Support for High Performance FICON (zHPF) - 2 of 3

The Evolution of High Performance FICON for System z


- First introduced in October of 2008: Initial zHPF announcement with FICON Express4 and FICON Express2 channels
 - Maximum of 31k zHPF IO/sec, 2.2 times the FICON protocol
 - Single track limit for zHPF data transfers
- In July 2009 FICON Express8 channels were introduced
 - Maximum of 52k zHPF IO/sec
 - 64k byte limit for zHPF data transfers
- In July 2010 additional support was introduced
 - Extension to multi-tracks of zHPF data transfers
- In July 2011 with the new FICON Express8S channel
 - Introduced a hardware data router for more efficient zHPF data transfers
 - FICON Express8S channel became the first channel with hardware specifically designed to support zHPF.
- In March 2015 with the new FICON Express16S channel
 - FICON Express16S channel became the first channel to support 16Gbps line speed.

TDMF Support for High Performance FICON (zHPF) - 3 of 3

zHPF Exploitation

- zHPF can be turned on or off. For z/OS exploitation, there is a parameter in the IECIOSxx member of SYS1.PARMLIB (ZHPF=YES/NO) and on the SETIOS command to control whether zHPF is enabled or disabled. The default is ZHPF=NO.
- zHPF-capable channels and devices support both FICON and zHPF protocols simultaneously. The Media Manager component of DFSMS™ detects whether the device supports zHPF or not and builds the appropriate channel programs.
- TDMF introduced support for High Performance FICON for the System z (zHPF) data transfer protocol beginning in 2009 with TDMF 5.2.0.

I/O Driver Benchmark maximum IO/sec Small Blocksize



TDMF Support for Basic HyperSwap

To quiesce and swap a volume in a Basic HyperSwap environment, the HyperSwap function itself must first be blocked - preventing both planned and unplanned HyperSwaps. Once the HyperSwap volumes in the session are swapped and terminated, the block is removed.

When a volume group that contains Basic HyperSwap volumes is about to be quiesced, TDMF blocks HyperSwap processing, by using the IOSHXBLK macro.

- The blocking interface module (IOSHMBLI) must be present in SYS1.LINKLIB and active on all systems in the sysplex.
- Ensure that only one TDMF session migrates HyperSwap volumes at any one point in time.

All HyperSwap volumes in a migration session must be grouped so that HyperSwap needs to be blocked only one time - when all the volumes are ready to be quiesced and synchronized. You can achieve this result by specifying OPTion(SINGLEGroup) on the SESSION control statement, which causes the volumes to be migrated in a group called SINGLE.

TDMF Control Statements for HyperSwap in a GDPS Environment

For a swap migration in a GDPS/ PPRC HyperSwap environment, the Master system will need to communicate with a Netview started task for HyperSwap commands and responses. There must be a TDMF Master or Agent system running on every active GDPS LPAR, including the controlling system(s).

The following example migrates six volumes as a single group (a requirement for HyperSwap-managed volumes), relabeling each source volume as \$\$uuuu, where uuuu is the device number. The optional **HYPERSW DISABLE** NetView command will be issued instead of **HYPERSW OFF**.

The Netview started task name for the TDMF internally issued **MODIFY operator** command is CNMPROC and the authorized Netview console name to be used is CONSTDMF, which is also the TDMF default.

TDMF HyperSwap Volume Migration in a GDPS Environment

Session Example

```
//STEP1 EXEC PGM=GTDMAIN,PARM=MASTER,TIME=1440
//STEPLIB DD DISP=SHR,DSN=HLQ.IBM.HGTD580.GTDLLIB
//GTDKEY DD DISP=SHR,DSN=HLQ.IBM.HGTD580.GTDLLIB
//SYSCOM DD DISP=SHR,DSN=HLQ.IBM.HGTD580.SYSCOM
//SYSPRINT DD SYSOUT=*
//DSFPRINT DD SYSOUT=*
//SYSIN DD *
SESSION
MASTER(TDM1)
AGENT(TDM2 TDM3 KSY1)
NETVIEW(CNMPROC) => NETVIEW started task name
CONSOLE(CONSTDMF) => Console variable to communicate with NETVIEW
OPTIONS(
PACING(NORMAL)
RELABEL($$)
HSWAPDISABLE => deactivates HyperSwap instead of setting it off
CHECKTARGET
CONCURRENT(03 ACTIVE)
FASTCOPY
SINGLEGROUP)
*
MIGRATE SRC000 TGT100
MIGRATE SRC004 TGT104
MIGRATE SRC008 TGT108
MIGRATE SRC00C TGT10C
MIGRATE SRC002 TGT102
MIGRATE SRC006 TGT106
//
```

Suspending the HyperSwap Environment

When a volume group that contains HyperSwap volumes is about to be quiesced, TDMF suspends the GDPS/PPRC HyperSwap environment. TDMF does this by passing the appropriate GDPS commands to NetView with an internal **MODIFY** command (for example, **MODIFY netview_proc,HYPERSW OFF**).

There are different ways for TDMF to suspend the HyperSwap environment:

- Using the **HYPERSW OFF** command.

This method is always chosen by TDMF if the HSWAPDisable option is not requested. The command disables HyperSwap and causes GDPS to unprepare all the HyperSwap managed devices. The subsequent **HYPERSW ON** command might take several minutes.

- Using the **HYPERSW DISABLE** command.

TDMF bypasses the **HYPERSW OFF** command if **OPTion(HSWAPDisable)** was specified on the **SESSION** control statement. This is the recommended method to suspend HyperSwap.


If the **HYPERSW OFF** command is not run, the subsequent **HYPERSW ON** command does not need to prepare the GDPS/PPRC HyperSwap managed devices. This results in a much faster return to normal operation.

HYPERSW ON is likely to take several minutes as the devices must be “prepared” again and this function also includes identification of the PPRC paths.

TDMF Virtual Storage Constraint Relief (VSCR)

Implemented with TDMF 5.4.0 and higher


- To support migrations of a larger number of volumes, TDMF was changed to reduce its ECSA requirements.
- Most TDMF requests for common storage will be satisfied by utilization of 64 bit common storage. Although the amount of storage required will not change, using common storage “above the bar” will provide enough common storage to satisfy all of TDMF’s needs from an area that will have little, if any, impact on other applications. The default 64 bit common area size is 64G.



Why is there a need for a Soft Fence state ?

What potential exposure does it address ?

- After a HyperSwap, or an unplanned site switch, or a volume migration to a new address, potential exposures exist to systems which might inadvertently use the “old” volume. Common scenarios include “old” Metro Mirror volumes or ‘other’ old mirrored volumes such as prior TDMF source volumes.
- Potential data integrity exposures include:
 - exposure to reading down-level data from the old MM primaries
 - exposure to data loss if updates occur on the wrong (old) mirrored primaries meaning data would be lost when the MM pairs are re-established in the reverse direction
 - exposure to bringing the “old” version of the volume online
- Example exposure scenarios
 - GDPS/PPRC provides IPL protection. However, if GDPS is not used to IPL a system image, the IPL could be performed using the former MM primaries
 - Systems outside of the Sysplex might continue to access (and perhaps update) the old MM primary volumes following an unplanned HyperSwap
 - After a TDMF volume migration, the wrong volumes (old source) could be varied online



What is the Purpose of a Soft Fence ?

- IBM Soft Fence will allow a host system to put a volume into a "soft fenced" state and also take it out of the "soft fenced" state. When a volume is in the soft fence state the disk subsystem will prevent all reads and writes to the duplicate volume from any host system.
- A fence state can be cleared using ICKDSF. If the device is offline it cannot be brought online without clearing the fence state first. This may be an issue when IPLing systems that expect access to volumes involved in a migration.
- TDMF supports this function by resetting a possible Soft Fence state for target devices at the beginning of a migration session and by setting the Soft Fence state for migration source devices as soon as they have been swapped with their associated target devices. Supporting source devices are put into the Soft Fence state by default in order to prevent any unintended and undetected access to original source volumes after they have been migrated as this may cause data integrity issues.



DEVSERV Query Device Command to Display Soft Fence Status for a Single Device

DEVSERV QD (Query Device) command to display Soft Fence status
DS QD,devn

DS QD,devn

IEE459I 07.08.18 DEVSERV QDASD 285

UNIT VOLSER SCUTYPE DEVTYPE CYL SSID SCU-SERIAL DEV-SERIAL EFC

07802 JZ7802 2107961 2107900 30051 0078 0175-W2201 0175-W2201 SOF

**** 1 DEVICE(S) MET THE SELECTION CRITERIA

**** 1 DEVICE(S) FAILED EXTENDED FUNCTION CHECKING

The Extended Function Checking (EFC) value of SOF indicates that the device is in a Soft Fence state.

ICKDSF Command to Clear a SOFT or SPID Fence

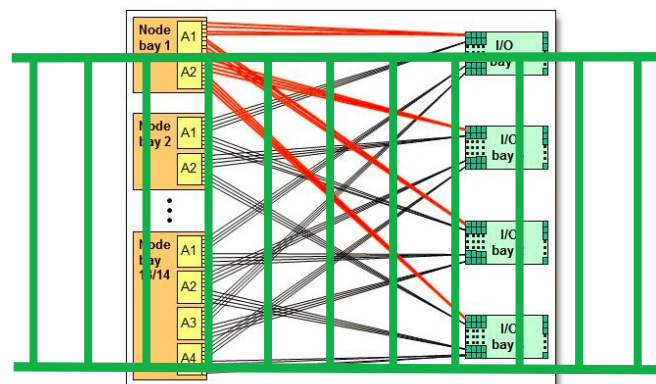
The following example shows the CONTROL command and JCL used to reset a fence status after the path and the device have been repaired. The DD statement identified in the input stream identifies an online DASD volume.

```
//jobname      JOB
//stepname     EXEC PGM=ICKDSF
//SYSPRINT    DD  SYSOUT=A
//DDCARD      DD  UNIT=3380,DISP=OLD,VOL=SER=ABCDEF
//SYSIN       DD  *
              CONTROL CLEARFENCE SOFT DDNAME(DDCARD)
              CONTROL CLEARFENCE SPID DDNAME(DDCARD)
/*
```


Why is there a need for a SPID Fence state ?

What potential exposure does it address ?

- As soon as TDMF z/OS has determined that all required TDMF Agents are running and before entering the Copy phase , source and target volumes are put into the Set Path group ID (SPID) Fence state if the function is available on z/OS and the storage subsystem.
- This prevents any system that did not already have access at the time the migration was started to establish an access path to any of the migration volumes (SPID fenced offline devices cannot be brought online).
- The SPID Fence function can be available on DS8700, DS8800, DS8870 and DS8886 disk subsystems. The SPID Fence state is removed as soon as the migration for a volume pair completes.
- TDMF utilizes this function to prevent Agentless systems from accessing a migration device.
- A SPID fence state can be cleared using ICKDSF.



TDMF Session Options for SOFT and SPID Fence Control

- **SPIDFENCE/NOSPIDFENCE**

- Description: SPID fencing action for source and target volumes.
- Function: Specifies the SPID fencing action to be taken on source and target volumes during the migration.
- Default: Source and target volumes will be put into the SPID fence state (if supported by operating system and storage controller).

- **FENcsource/NOFENcsource**

- Description: Soft fencing action for source volumes
- Function: Specifies the soft fencing action to be taken on source volumes after a successful MIGRATE.
- Default: Source volumes will be put into the Soft Fence state (if supported by operating system and storage controller).


- **UNFENcetarget/NOUNFENcetarget**

- Description: Soft fencing action for target volumes
- Function: Specifies the soft fencing action to be taken on target volumes prior to starting a migration. UNFENcetarget removes the target volumes from the soft fence state prior to the beginning of a migration. NOUNFENcetarget does not remove the target volume from the soft fence state prior to the beginning of a migration.
- Default: Target volumes are taken out of the Soft Fence state (UNFENcetarget).

TDMF Support for Up To 1TB Extended Address EAV Devices

- TDMF 5.5.0 is capable of migrating data sets on Extended Address Volumes (EAVs), which are supported on z/OS 1.10. EAVs have a max number of cylinders per volume to 1,182,006 cylinders (1 TB) of addressable storage.
- NOTE: EAV volumes do not currently support the following data set types in the extended addressing areas:
 - Page data sets
 - VSAM data sets with IMBED or KEYRANGE attributes. These cannot be created but they might have been inherited from prior physical migrations or copies.

* NOTE: Volume images are not to scale



Maximum Volume Pair Limit per Each TDMF Session


- The number of **MIGRATE** or **REPLICATE** statements that define the volume pairs to be migrated or replicated was increased in TDMF 5.6.0 and higher.
- **Maximum Number of volume pairs per TDMF Master Session TDMF 5.5.0 and lower:**
 - A maximum of up to 512 volume pairs per session are allowed, however, the number of LPARs times the number of volume pairs can not exceed 2048

$$(\text{number-of-LPARs} \times \text{volumes} \leq 2048)$$

- **Maximum Number of volume pairs per TDMF Master Session TDMF 5.6.0 and higher:**
 - The maximum number of volume pairings per TDMF session is limited by the system resources, in particular real storage frames. Fixed storage frame requirements rise as the number of volume pairings increase.

TDMF Session Start-Up Timer Increased to 60 Minutes

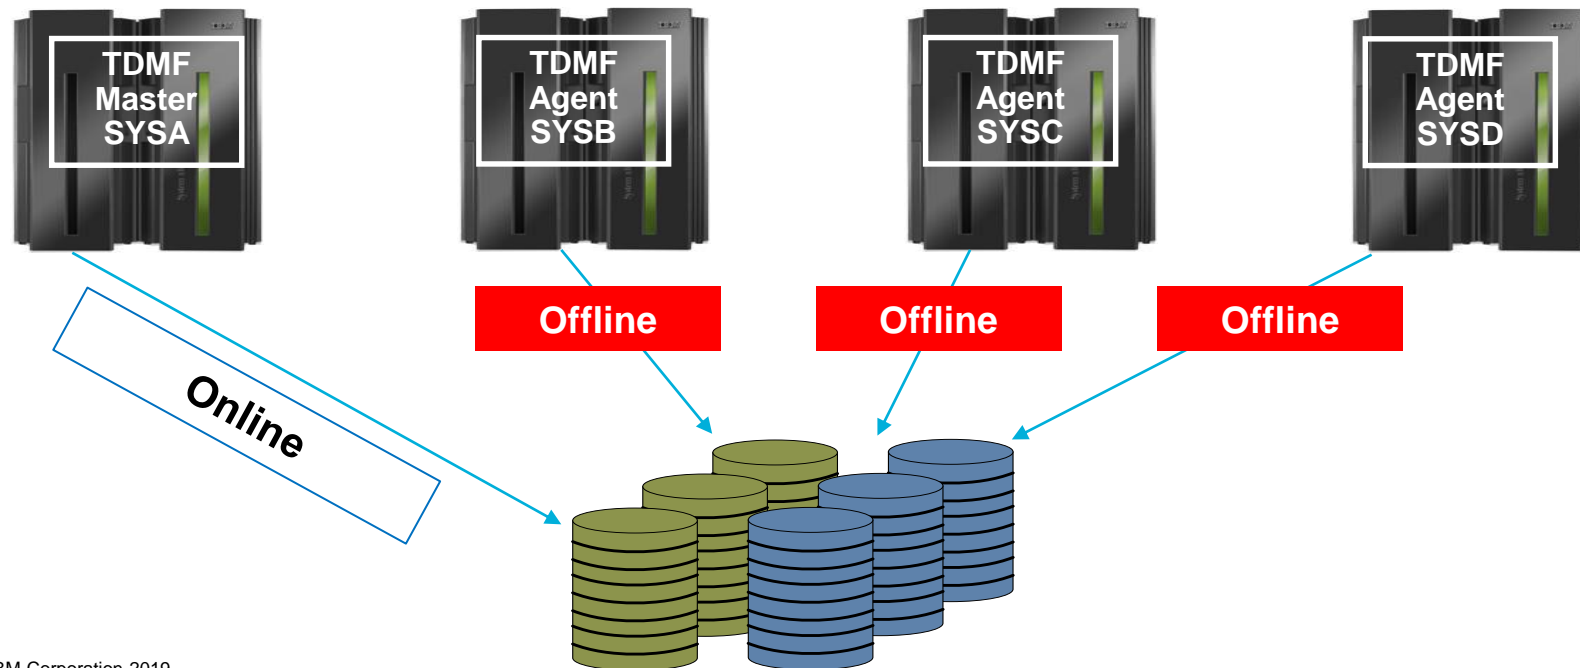
- With TDMF 5.6.0 and higher, there is a 60-minute window after any TDMF job is submitted whereby the other batch jobs must be submitted. If after 60 minutes, the Master or any attached Agent job has not properly initiated for any reason, the session is terminated.
- The previous TDMF Session Start-up Timer was a 15-minute timer which existed with all TDMF version/s and releases prior to 5.6.0 (5.5.0 and lower).
- **Important:** If this situation occurs, the entire session must be restarted.
- **Note:** it does not matter what order the Master and Agent jobs are submitted in, as long as they are all submitted within 60 minutes of each other.



TDMF Volume Swap Communication Reduction

Improved Swap Performance

- This TDMF 5.6.0 enhancement will improve TDMF volume swap performance in multi-system migration sessions where one or more migration source volumes are offline to a host system running a TDMF agent.
- In the above-mentioned scenario, communication between TDMF Master and Agent host systems is minimized. The reduction in Master/Agent communication results in a reduction in I/O, wait and processor times and improved TDMF session run times.



TDMF Dynamic I/O Pacing Improvements


- This TDMF 5.6.0 enhancement is designed to improve performance in the Dynamic Volume Pacing feature.
- The purpose of this enhancement is to maximize TDMF performance while minimizing the migration's impact on the production environment.
- TDMF I/O handling has been modified with this enhancement to:
 - Adjust the number of tracks per TDMF I/O depending on the percentage of available fixable real storage frames
 - Optimize I/O handling commensurate with I/O service times



TDMF Toleration Support for zHyperLink

- IBM DS8880 zHyperLink Technology is a new mainframe attachment for storage. zHyperLink is the first new mainframe input/output (I/O) channel link technology since FICON.
- zHyperLink is a short-distance mainframe attach link designed for up to 10x lower latency than High Performance FICON (zHPF).
- zHyperLink will drastically improve response time and throughput for some of IBM Z's most demanding workloads, lowering database transactional latency provides opportunities for advancements in security, personalization, and mobile access.
- This technology can be installed on existing IBM DS8880 storage systems.
- zHyperLink is intended to speed up DB2 for z/OS transaction processing and improve active log throughput.

Data Migration Agent for z/VM Environments (TDMF 5.7.0)




With the TDMF z/VM agent the devices can be swapped while z/VM applications and guest systems remain online with continuous access to data.

- The swap migration process is controlled from a TDMF z/OS master system.
- The z/VM agent keeps track of disk updates that happen during the migration session
- The changes are passed to z/OS in order to execute the z/VM updates.
- All this happens while z/VM applications remain online.
- With a TDMF agent for z/VM it is now also possible to migrate disks of Linux guest systems

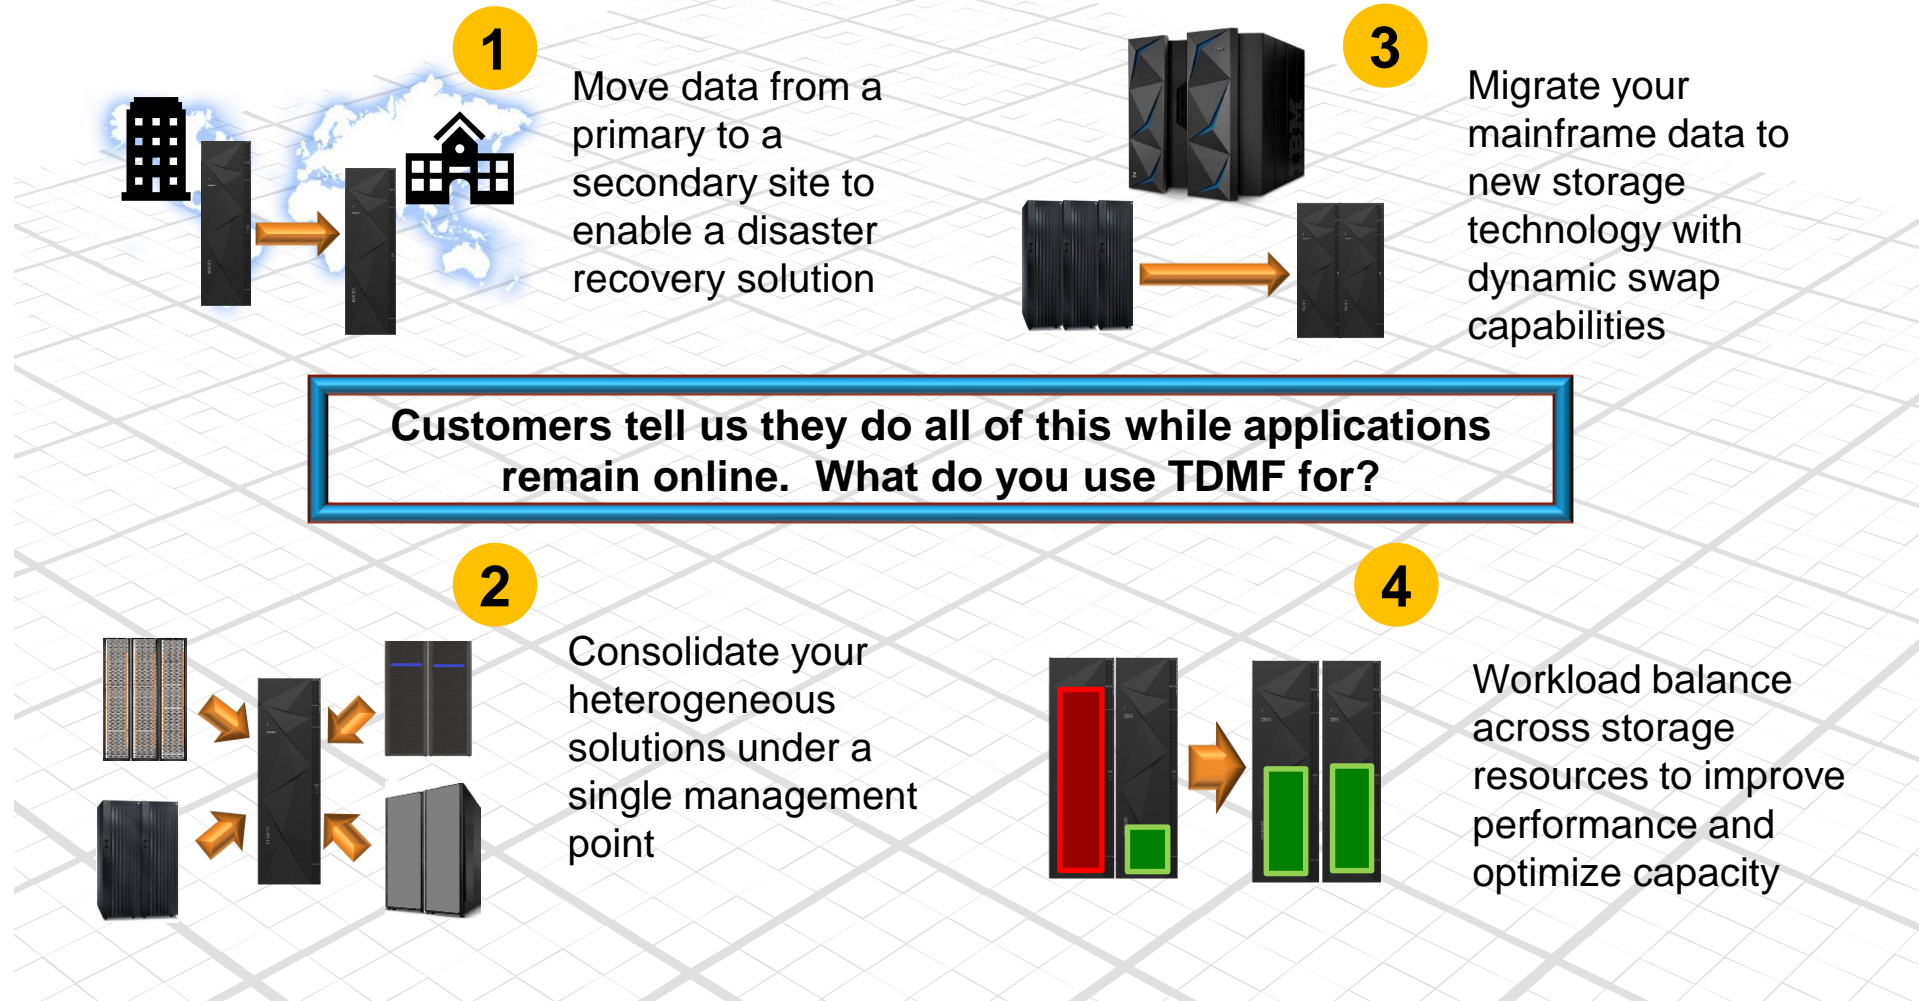
The z/VM agent provides continuous availability with transparent swap capabilities for z/OS and z/VM.

More about the zVM Agent for TDMF



- The z/VM Agent for TDMF z/OS is implemented as a VM server virtual machine that can run multiple concurrent migration sessions.
- Each migration session uses a single dedicated COMMDS that is stored on an MVS formatted volume (containing a valid VTOC) which is shared among all host systems that participate in a migration.
- A TDMF Agent Session on z/VM runs as a thread (GTDVTDMF) in the TDMF Server virtual machine running zCMS, whereas, a TDMF z/OS Agent runs as a batch job.
- A migration session will be initiated by issuing a SESSION command to the TDMF agent application which is then executed as a thread (sub-task) in the TDMF Server virtual machine
- The SESSION command tells the z/VM Agent where to find the SYSCOM dataset associated with the migration session. A TDMF diagnose code is used to interact with the VM CP.
- The TDMF master still runs on z/OS, therefore, Linux volumes to be migrated must be ECKD and formatted according to the IBM label partitioning scheme, z/OS compatible disk layout (CDL)

Data Migration Solutions for IBM z/OS – Use Cases



Roadmap & Strategy

IBM Z Volume Migration looking forward

TDMF Migration Solutions Roadmap

DELIVERED

- Support for zHyperlink - Increase volume pairs allowed per session
- Support for online secondary volume
- Performance improvement for COMMDS I/O
- Improve cleanup utility
- Support TDMF as master on GDPS/zOS Proxy
- Support for z/VM Agent for TDMF on z/VM 7.1

- Transport Mode (ZHPF) for TDMF reads/writes
- Increased number of volume pairs for remote TCP/IP migrations
- TDMF Automation and Reporting Tools

- Volume stacking
- PAV exploitation
- Exploitation of Flashcopy for TDMF copy phase
- Volume replication to Cloud
- Coupling Facility for COMMDS
- Virtual SYSCOM
- Exploitation of z/EDC hardware compression
- Improved user interface
- TDMF as started task
- z/Linux Agent
- z/VM TDMF master

zDMF Overview

zDMF Data Migration Solutions Features today

- **Volume consolidation** – move data sets from multiple small volumes to a single larger volume; helps with UCB limits
- **Multi-volume data set consolidation** - convert multi-volume data set to single volume data set
- **Data set extent consolidation** – move a data set with many extents to data set with fewer extents or single extent; reduce volume fragmentation
- **Move non-EAV to EAV** – move track-managed data sets to cylinder-managed space based on Extended Address Space allocation rules

zDMF Data Migration Solutions Features today

- **Move non-SMS data set to SMS data set** – honors SMS ACS rules for data set placement
- **Move SMS data set to non-SMS**
- **Data set placement** – quickly resolve contention issues by moving problem data sets to other volumes
- **Facilitate storage tiering** - move data sets to different storage types – solid state, enterprise, nearline

zDMF Data Migration Solutions Features today

- **Non-disruptive data set moves** - in many cases data sets can be moved while open without requiring any application outage/downtime.
- **Early completion of unallocated data sets** – unallocated data sets bypass zDMF mirror phase and are migrated and completed immediately upon completion of copy
- **Reduced application downtime for data set moves** - When application downtime is required to complete the move, the downtime is significantly reduced compared to the downtime required when using traditional tools such as DFSMSdss, FDR, et. al.

What zDMF 3.4 does for Pervasive Encryption

zDMF 3.4 Release (GA December 2018) aligns with IBM's Pervasive Encryption strategy to protect mainframe data by enabling clients to transparently encrypt data sets.

Data protection and compliance are business imperatives

Clients are relying on encryption to protect their data, but implementing encryption can be a complex process



Delivered December 2018

- zDMF encrypts Extended Format data sets while they are in use by mirroring I/O to an encrypted version of the data set
- zDMF obtains the KEYLABEL for data sets to be encrypted using the same search order as DFSMS
- zDMF control cards to encrypt data sets are easier and faster than coding DFDSS or IDCAMS control cards to copy and encrypt data sets


Coming 3rd Quarter 2019

- zDMF will have the ability to convert basic/large format data sets to extended format data sets and encrypt them at the same time giving the customer an easy method for encrypting basic data sets or converting to extended format data sets
- zDMF will have the ability to switch a data set from one encryption key to another



Access control – Segregation of duties

- **Data owners** who ***must*** access content will need authority access to the data set ***as well as*** access to the encryption key label
- **Storage administrators** who only ***manage*** the data sets need access to the data set ***but not*** access to the key label (thus protecting access to the content)
- Different keys can be used to protect different data sets – ideal for multiple tenants or data set specific policies.
- Prevent administrators from accessing the content.
- Many utilities can process data preserving encrypted form:
 - COPY, DUMP and RESTORE
 - Migrate/Recall, Backup/Recover, Dump/Data Set Restore
 - PPRC, XRC, FlashCopy®, Concurrent Copy, etc.



***Limit access to data in clear! Remove certain roles from compliance scope
by controlling access to the data through SAF permissions.***

How does zDMF solve the problem?

Once the data set encryption key has been determined it is stored in RACF, the SMS data class or JCL/IDCAMS control card.

The data set must then be recreated to get encrypted. Naturally some data sets have recreation as part of their batch cycles. All other VSAM and sequential data sets will need to be recreated.

Without zDMF:

1. Stop any applications using the data set(s)
2. Unload the data set(s) up or find a current unloaded version
3. Delete the data set(s)
4. Define the data set(s)
5. Reload the data set(s) from the unloaded copy
6. Restart the application



How does data set encryption work without zDMF?

1. Stop any applications using the data set(s)
2. Unload the data set(s) or find a current unloaded

```
//REPROS EXEC PGM=IDCAMS
//SYSPRINT DD SYSOUT=*
//OUT DD DSN=JHSRC.ENCTWX02.C0000001.REPRO,DISP=(,CATLG),
//      STORCLAS=LDMFTSC2,SPACE=(CYL,(100,100),RLSE),
//      DCB=(BLKSIZE=0,LRECL=32760,RECFM=FB,DSORG=PS)
//SYSIN DD *
REPRO -
  INDATASET('JHSRC.ENCTWX02.C0000001') -
  OUTFILE(OUT)
```

3. Delete the data set(s)
4. Redefine the data set(s)

```
//DELDEF EXEC PGM=IDCAMS
//SYSPRINT DD SYSOUT=*
//SYSIN DD *
DELETE JHSRC.ENCTWX02.C0000001
DEFINE CL (NAME(JHSRC.ENCTWX02.C0000001) -
  SHR(2,3) -
  RECSZ(32760 32760) -
  STORCLAS(LDMFTSC1) -
  DATACLAS(JHEXTF) -
  CYL(150 10) -
  CISZ(4096) -
  FREESPACE(00 00) -
  KEYS(8 0))
```

5. Reload the data set(s) from the unloaded copy

```
//RESTORE EXEC PGM=IDCAMS
//SYSPRINT DD SYSOUT=*
//IN DD DSN=JHSRC.ENCTWX02.C0000001.REPRO,DISP=SHR
//SYSIN DD *
REPRO -
  INFILE(IN) -
  OUTDATASET('JHSRC.ENCTWX02.C0000001')
```

6. Restart the applications



The application
is unavailable
while the data
set is unloaded,
redefined and
reloaded

zDMF solution that works

With zDMF

1. Activate a zDMF group with the data set. It will find the key name from RACF, SMS or its own control card
2. Anytime after the zDMF group reaches the I/O MIRROR phase, stop the application
3. DIVERT the zDMF group which takes less than a second per data set
4. Restart the application(s)




Note: DFSMSdss can't do this, it retains a track image

What it means to the applications

The value of zDMF is

- Application doesn't have to be down while the data is being copied and encrypted
- Don't have to code or find unload JCL for data sets
- Don't have to code or find allocation JCL/IDCAMS to define data set
- Don't have to code or find restore JCL for data sets



Data set types

- Data set types supported by data set encryption
 - VSAM extended format
 - Sequential extended format
- Data set types NOT supported by data set encryption
 - VSAM non-extended format
 - Sequential basic format
 - Sequential large format
 - PDS/PDSE
 - BDAM
 - Sequential tape data sets



Why you encrypt your data

This is a screen shot of what your data looks like on disk. Anyone in the installation with DASD admin or SYSPROG type authority can see the data. It's at CYL 0078 track 00 on volume TD5E5D. This information obtained from a LISTVTOC.

```
*** TRACK(CCHH) 00780000      RO DATA 0000000000000000
      COUNT 0078000001000070
0000 D1D6C8D5 4040C4D6 C54040F4 F3F2F140 40D4C1D7 D3C54040 C1E5C540 40C3C1D9 *JOHN..DOE..4321..MAPLE..AVE..CAR*
0020 C5C6D9C5 C54040F9 F1F2F0F7 4040F4F3 F5F3F4F5 F7F8F3F4 40F2F3F7 60F8F860 *EFREE..91207..4353457834.237-88-*
0040 F4F4F4F4 40404040 F0F0F0F6 F0F0F1F3 00500000 00000000 00001000 00000000 *4444....00060013.&.....*
0060 00000000 00000000 0000C8DD 065A5AA5 *.....H!!v*
```

This is a screen shot of the same data encrypted.

```
*** TRACK(CCHH) 00780000      RO DATA 0000000000000000
      COUNT 0078000001000070
0000 60D83FAC 6563B147 8B8F31BC 83D089BB 3CE7A77E 1797D9BF 74819EFA CF1B6E93 *-Q.....c}i..Xx=.pR..a.3..>l*
0020 3548CDF4 CC0BCFEA BB6C8B00 335CEF00 7DD5B54C 102E0EA4 1687BDFF C57E8B63 *...4.....%...*...'N.<...u.g..E=..*
0040 4A1B8F1C 6130E3FB 4FE3CA06 B5D8A214 00500000 00000000 00001400 00000000 *¢.../.T.|T...Qs..&.....*
0060 00000000 00000000 0000BC5F E65A5AA5 *.....-W!!v*
```

zDMF 3.4 Data Set Encryption Demo

zDMF Installation

- ZDMF has 6 installation data sets; 5 of them contain the executable code to run the product and one is a sample library
- The libraries use IBM standard installation names (LLIB, ELIB etc..)
- The load library must be APF authorized
- The zDMF server must have access to the encryption keys and to the DFSMS facility that allows data set encryption
- The other data sets are used by the ISPF user interface except for the samplib data set

```
JHH60.IBM.HGZD340.SGZDELIB  
JHH60.IBM.HGZD340.SGZDLLIB  
JHH60.IBM.HGZD340.SGZDMLIB  
JHH60.IBM.HGZD340.SGZDPLIB  
JHH60.IBM.HGZD340.SGZDTLIB
```

zDMF Address Space

- ZDMF has a server address space that must run on every system that has access to the data being migrated
- The zDMF server proc GZDPROC should be added to a system proclib
- The server proc points to the initialization parms, and the load library.

```
BROWSE      GZD.V340REL.SAMPLIB(GZDZPROC) - 01.20      Line 0000000
Command ==> S
/ZDMF      EXEC PGM=GZDMAIN,COND=(0,LT),REGION=0M,TIME=NOLIMIT,
/           PARM=('CFG=&CHLQ..&VER..SGZDSAMP(&MEM)',
/           'START=&START')
/*
/*
/*           Following is the zDMF Load Library
/*
/STEPLIB   DD  DISP=SHR,DSN=&CHLQ..&VER..SGZDLLIB
/*
/*
/*           Following is the zDMF Security Library
/*
/GZDKEY     DD  DISP=SHR,DSN=&CHLQ..&VER..SGZDLLIB
/*
```

- The server JCL for running zDMF as an STC is member GZDZPROC in SAMPLB
- The STEPLIB and GZDKEY point to your loadlib

zDMF Configuration

The CFG data set contains the zDMF configuration parms:

```
*
*   Sample Start Up Configuration File
*
*
*   4 Digit Server ID
SUBNAME=ZDJ4
*
*   2 Digit Command Prefix
CPFX=J4
*
*   The ZDMF Data Base File
DB=JHH60.ZDMF.V340
*
*   High Level Qualifiers for Target Extent Control Files
ZDPFX=JHH60.ZDMF.V340EXT
*
```

Change the DB(database HLQ), SUBNAME(subsystem name) and CPFX (z/OS command prefix) and the Target HLQ parameters as follows:

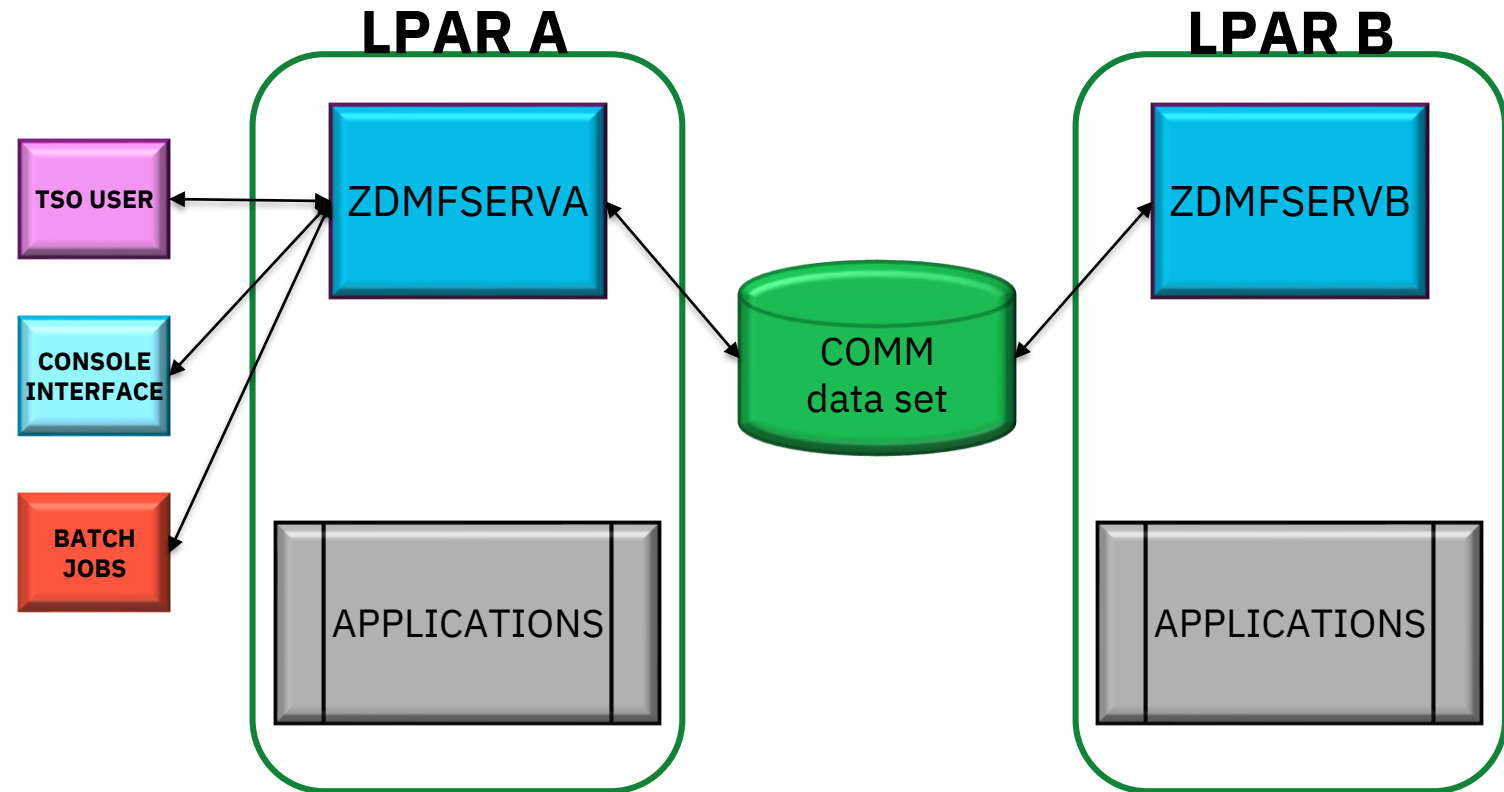
1. DB HLQ – Your data base contains group definitions and is a communication data set between multiple systems
2. SUBNAME – Is any valid available subsystem name on your system
3. CPFX – A 2-4 character z/OS command prefix to communicate to your server
4. ZDPFX – An HLQ used to allocate 'extent map' data sets on your target volumes

zDMF Server Overview

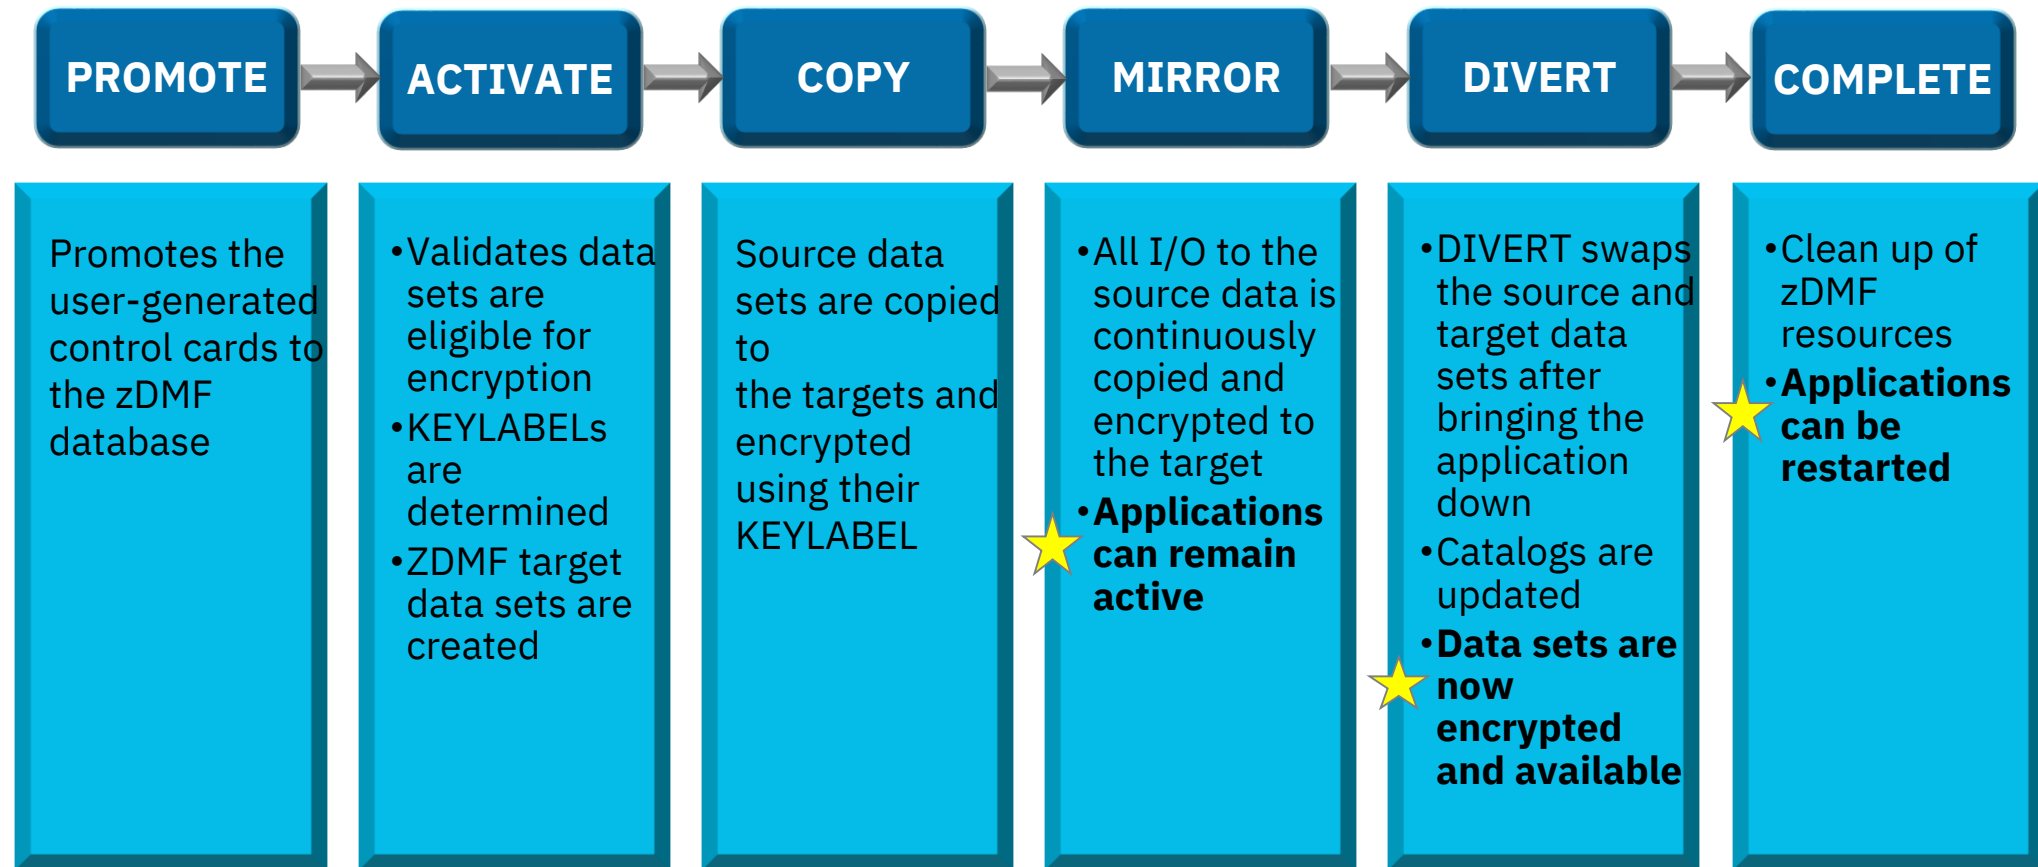
There is a zDMF server for each LPAR in the sysplex

zDMF Server functions are:

- Install/Monitor tracks
- Activate/terminate groups
- Communicate data set extent changes
- Communicate group status changes



The 6 zDMF Phases



zDMF PROMOTE Phase

ZDMF uses a 'GROUP' to identify data sets to be migrated/and encrypted and their target volumes/storage class.

```
//PROMOTE EXEC PGM=GZDBAT,PARM='CPFX=J4,DEBUG=NO',COND=(7,LE)
//STEPLIB DD DISP=SHR,DSN=JHH60.IBM.HGZD340.SGZDLLIB
//SYSPRINT DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//SYSIN DD *
PROMOTE DLM=ZZ
GROUP (ENCTWX02) -
  DELETE_EXISTING_TARGET_DATASETS (YES) -
  EARLY_DATA_SET_COMPLETION (NO) -
SET -
  SPHERE (YES) -
  TOLERATE_SOURCE_VALIDATE_FAILURE (NO) -
  MIGRATE_ONLY_SPECIFIED_VOLUMES (NO) -
  ENCRYPT(YES) -
  KEYLABEL (TDM.AES256.DATA) -
  SOURCE (DSN (JHSRC.ENCTWX02.C00000001)) -
  TOLERATE_TARGET_CREATE_FAILURE (NO) -
  TARGET (DSN (JHTGT) -
    STORAGECLASS(LDMFTSC1))
ZZ
```

To point to your zDMF server all that is required is the command prefix. In this example J4 is the command prefix used. This JOB will store the GROUP in your ZDMF data base where you can later ACTIVATE it to start the migration/encryption process.

After PROMOTING the GROUP, it can be displayed with the Display Groups command, or with the ISPF interface:

```
RESPONSE=TDM1
GZD1100I Command received from userid: JHH61
GZD1101I D G
NAME      1ST-PAIR  DSNS  STATUS      OWNER
ENCTWX02  0000/0000    0  NEVER ACTIVE
ENCTWX04  0000/0000    0  NEVER ACTIVE
```

```
Subsystem J4  Groups  Row 1 to 19 of 30
Command ==> _  Scroll ==> CSR
Group          TDM1
Data Set      Status
Extents
- ENCTWX02      Inactive
- ENCTWX04      Inactive
Activate, Divert, and Terminate reports available
```

zDMF ACTIVATE Phase

The ACTIVATE phase determines where to migrate the data sets and obtains KEYLABELS via RACF, user definition or SMS Dataclas. During this phase a TARGET data set is created that is a clone of the SOURCE data set except that the target will contain the encryption definitions in the VVDS, VTOC, etc. The TARGET data set uses the TARGET DSN HLQ and generates a name that includes the GROUP name for a data set.

```
JHH61      00000290  J4 ACT ENCTWX02
STC01196  00000090  GZD1100I Command received from userid: JHH61 310
          310 00000090  GZD1101I ACT ENCTWX02
          310 00000090  GZD3530I Group ENCTWX02 ACTIVATED
STC01196  00000090  GZD2024I ENCTWX02 is now ACTIVE (00->80) at CMD_ACTIVATE+0284
STC01196  00000090  GZD0169I GROUP ENCTWX02 WAITING FOR TDM2/ZDJ4
STC01196  00000090  GZD1101I GROUP (ENCTWX02) -
STC01196  00000090  GZD1101I DELETE_EXISTING_TARGET_DATASETS (YES) -
STC01196  00000090  GZD1101I EARLY_DATA_SET_COMPLETION (NO)
STC01196  00000090  GZD1101I SET -
STC01196  00000090  GZD1101I SPHERE (YES) -
STC01196  00000090  GZD1101I TOLERATE_SOURCE_VALIDATE_FAILURE (NO) -
STC01196  00000090  GZD1101I MIGRATE_ONLY_SPECIFIED_VOLUMES (NO) -
STC01196  00000090  GZD1101I ENCRYPT(YES) -
STC01196  00000090  GZD1101I KEYLABEL (TDM.AES256.DATA) -
STC01196  00000090  GZD1101I SOURCE (DSN (JHSRC.ENCTWX02.C00000001)) -
STC01196  00000090  GZD1101I TOLERATE_TARGET_CREATE_FAILURE (NO) -
STC01196  00000090  GZD1101I TARGET (DSN (JHTGT) -
STC01196  00000090  GZD1101I STORAGECLASS(LDMFTSC1))
STC01196  00000090  GZD0118I Command parse successful
```

This is the result of an ACTivate command in the SYSLOG. The GROUP is activated and TARGET data sets are allocated. An OPEN is done for the TARGET data set using the encryption key therefore the zDMF server must have access to that encryption key.

```
STC01196  00000090  GZD2024I ENCTWX02 is now ACT-PND (80->C0) at SYNC_PHASE1+0292
STC01196  00000090  GZD0170I Heartbeat task is now activating group ENCTWX02 on 2 system(s).
STC01196  00000090  GZD2024I ENCTWX02 is now MIR-PND (C0->C8) at SYNS_RESUME+091A
```

zDMF ACTIVATE Phase

Messages in the ACTIVATE phase describe the steps taken during activation of the group

```

02/11/2019 03:01:23.221 GZD4160I DATA SET JHSRC.RACF.ENCR0003.C0000001
                                TO GROUP ENCR0003 VIA SOURCE DATASET
                                STATEMENT JHSRC.RACF.ENCR0003.C0000001
02/11/2019 03:01:23.226 GZD4212I FOUND A DATAKEY IN DFP SEGEMENT FOR
                                DSN=JHSRC.RACF.ENCR0003.C0000001
                                RACF_PROFILE=JHSRC.RACF.ENCR0003.C*
                                KEYLABEL=TDM.AES256.DATA.KEY3
02/11/2019 03:01:23.255 GZD4148I SMS VOLUME TD5349 IN STORAGE GROUP
                                ADDED TO TARGET VOLUME LIST
02/11/2019 03:01:23.256 GZD4146I ACS ROUTINE MSG= IGD01008I USER
                                STORAGE CLASS LDMFTSC1 RETAINED.
02/11/2019 03:01:23.372 GZD4166I THE TARGET DATA SET
                                JHTGT.ENCR0003.D9042.T0301232.S00001
                                FOR SOURCE DATA SET
                                JHSRC.RACF.ENCR0003.C0000001 ALLOCATED
02/11/2019 03:01:23.372 GZD4153I THE ACTIVATION ECSA CONTROL BLOCK
                                PHASE IS BEING ENTERED FOR ENCR0003
02/11/2019 03:01:23.376 GZD4154I ACTIVATION PASSING CONTROL TO HEARTBEAT
                                TASK FOR GROUP ENCR0003
    
```

1. Determine what source data sets are to be included in the group



2. Find the KEYLABEL



3. Find the target volumes using SMS ACS call



4. Allocate the target data set



5. ACTIVATE the group

zDMF ACTIVATE Phase

A J4 D DSN ENCTWX02 z/OS command will show the data set details. You can see the encryption key name used by the SOURCE and TARGET volumes and their corresponding extent locations on those volumes.

```

RESPONSE=TDM1
GZD1100I Command received from userid: JHH61
GZD1101I  D DSN
DSNAME SOURCE->TARGET          GROUP
EXT  SRC/TGT      cccCCCCCH  STATUS  TRKS/%  DSCB-CCHHR
JHSRC.ENCTWX02.C0000001      ENCTWX02
-> JHTGT.ENCTWX02.D8255.T2148205.S00001
*** VSAM Cluster
KEYLABEL=TDM.AES256.DATA
JHSRC.ENCTWX02.C0000001.DATA      ENCTWX02
-> JHTGT.ENCTWX02.D8255.T2148205.S00001.DATA
 1 C632(TD2727) 00006660 MIRROR      2250  0000000702
  C647(TD277B) 0000B2E0              100%  0000000529
 2 C632(TD2727) 00006070 MIRROR      750
  C647(TD277B) 0000BC40              100%
 3 C632(TD2727) 00006430 MIRROR      150
  C647(TD277B) 0000BF60              100%
 4 C632(TD2727) 00008E10 MIRROR      300  000000080B
  C647(TD277B) 0000C000              100%  0000000607
JHSRC.ENCTWX02.C0000001.INDEX      ENCTWX02
-> JHTGT.ENCTWX02.D8255.T2148205.S00001.INDX
 1 C632(TD2727) 000048C3 MIRROR      20  0000000806
    
```

At this point a LISTCAT of the TARGET data set shows a valid encrypted data set.

```

CLUSTER ----- JHTGT.ENCTWX02.D8255.T2148205.S00001
IN-CAT --- JH.UCAT.TGT.JH5607
HISTORY
  DATASET-OWNER----(NULL)      CREATION-----2018.255
  RELEASE-----2            EXPIRATION-----0000.000
SMSDATA
  STORAGECLASS ---LDMFTSC1      MANAGEMENTCLASS---(NULL)
  DATACLASS -----JHEXTF      LBACKUP ---0000.000.0000
  CA-RECLAIM-----YES
  EATTR-----NULL
  BWO STATUS-----00000000      BWO TIMESTAMP---00000 00:00:00.0
  BWO-----NULL
RLSDATA
  LOG -----NULL            RECOVERY REQUIRED --(NO)      FRLOG
  VSAM QUIESCED -----NO      RLS IN USE -----NO      LOGREP
  LOGSTREAMID-----NULL
  RECOVERY TIMESTAMP LOCAL-----X'0000000000000000'
  RECOVERY TIMESTAMP GMT-----X'0000000000000000'
ENCPTIONDATA
  DATA SET ENCRYPTION----YES
  DATA SET KEY LABEL----TDM.AES256.DATA
PROTECTION-PSWD-----NULL      RACF-----NO
ASSOCIATIONS
  DATA----JHTGT.ENCTWX02.D8255.T2148205.S00001.DATA
  INDEX---JHTGT.ENCTWX02.D8255.T2148205.S00001.INDX
    
```

zDMF COPY Phase

The next zDMF phase of migration is COPY this is where the SOURCE data is copied to the TARGET volume and the blocks/CIs are encrypted during COPY. The length of this phase is dependent upon the amount of data being migrated and for small data sets can pass without notice.

```
GZD2024I ENCTWX02 is now MIR-PND (C0->C8) at SYNS_RESUME+091A
GZD0215I Synchronization for TD2727 (C632 -> C647) completed (3,470
tracks).
GZD2024I ENCTWX02 is now MIRROR (C8->88) at IS_GROUP_MIRACT+0136
```

Note that COPY starts with MIR-PND and ends when the GROUP moves to the MIRROR phase.

After the COPY phase if the EARLY_dataset_COMPLETE option is set to YES any data sets in the GROUP that are not in use are DIVERTed and COMPLETED.

```
Subsystem J4  Groups  Row 1 to 19 of 31
Command ==>  Scroll ==> CSR
Group  TDM1
Data Set  Status
Extents
- ENCTWX02  Owning System: TDM1  Moved= 0% Copy
Activate report available
- ENCTWX04  Inactive
Activate, Divert, and Terminate reports available
```

zDMF MIRROR Phase


During this phase write I/O to the SOURCE data set(s) extents are mirrored to the TARGET data set and extents that represent encrypted data sets have their TARGET I/O buffers encrypted.

A DFDSS print of the tracks that comprise the TARGET data set will show encrypted data with an MMSX that has the encrypt bit set.

```
.      Subsystem J4  Groups      Row 1 to 19 of 31 .
. Command ==>      Scroll ==> CSR .
.      Group      TDM1 .
.      Data Set   Status .
.      Extents .
. - ENCTWX02  Owning System: TDM1      Moved=100% Mirror .
.   Activate report available .
. .
. - ENCTWX04      Inactive .
.   Activate, Divert, and Terminate reports available .
```

zDMF MIRROR Phase

During this phase write I/O to the SOURCE data set(s) extents are mirrored to the TARGET data set and extents that represent encrypted data sets have their TARGET I/O buffers encrypted.



zDMF DIVERT Phase

During the DIVERT phase the SOURCE and TARGET data sets are 'swapped' by catalog, VVDS and VTOC updates. Note: *** Currently the SOURCE data sets must be closed and unallocated before DIVERT can take place ***. This will be changed in a future release.

If a data set is in use, you will see the

```
J4 DIVERT ENCTWX02
GZD2065E DIVERT delay - ENCRYPT dsntype: JHSRC.ENCTWX02.C0000001
GZD1100I Command received from userid: JHH61 538
GZD1101I DIVERT ENCTWX02
GZD0169I Group ENCTWX02 waiting for TDM2/ZDJ4
```

Once the data set(s) is closed an

```
J4 DIVERT ENCTWX02
GZD1100I Command received from userid: JHH61 565
GZD1101I DIVERT ENCTWX02
GZD3530I GROUP ENCTWX02 Divert process starting
GZD2024I ENCTWX02 is now DIV-PND (88->C4) at CHECK_GROUP_STATUS+0256
GZD0577I Heartbeat task is now diverting group ENCTWX02 on 2 system(s).
GZD4000I DIVERT for GROUP ENCTWX02 started.
GZD4001I DIVERT for GROUP ENCTWX02 ended 0 Data Sets failed DIVERT
GZD2024I ENCTWX02 is now DIVERT (C4->84) at ES_DIVERT+0364
```

In ISPF you will see:

```
Subsystem J4 Groups Row 1 to 19 of 31
> CSR
- Group: ENCTWX02 TDM1
  GZD2065E DIVERT DELAY - ENCRYPT DSNTYPE: JHSRC.ENCTWX02.C0000001 tatus
ENCTWX02 Owning System: TDM1 Moved=100% Mirror
Activate report available
```

Since there are no allocations to the data set(s) the GROUP will enter the zDMF complete phase where the zDMF resources will be freed.

```
GZD2024I ENCTWX02 is now CMP-ALLP (84->C2) at LOG_OUR_STATUS+0812
GZD2024I ENCTWX02 is now COMPLETE (C2->02) at TERMINATE_ALL+021C
GZD0589I Heartbeat task setting group ENCTWX02 to a status of terminated
```

zDMF DIVERT Phase

At this point the SOURCE data set(s) now reside on the TARGET volumes and a LISTCAT of the SOURCE data sets(s) show that it is encrypted.

```
CLUSTER ----- JHSRC.ENCTWX02.C0000001
IN-CAT --- JH.UCAT.SRC.JH5606
HISTORY
  DATASET-OWNER----- (NULL)      CREATION-----2018.255
  RELEASE-----2      EXPIRATION-----0000.000
SMSDATA
  STORAGECLASS ---LDMFTSC1      MANAGEMENTCLASS--- (NULL)
  DATACLASS -----JHEXTF      LBACKUP ---0000.000.0000
  CA-RECLAIM----- (YES)
  EATTR----- (NULL)
  BWO STATUS-----00000000      BWO TIMESTAMP---00000 00:00:00.0
  BWO----- (NULL)
RLSDATA
  LOG ----- (NULL)      RECOVERY REQUIRED -- (NO)      FRLOG --
  VSAM QUIESCED ----- (NO)      RLS IN USE ----- (NO)      LOGREPLI
  LOGSTREAMID----- (NULL)
  RECOVERY TIMESTAMP LOCAL-----X'0000000000000000'
  RECOVERY TIMESTAMP GMT-----X'0000000000000000'
ENCRYPTIONDATA
  DATA SET ENCRYPTION---- (YES)
  DATA SET KEY LABEL----TDM.AES256.DATA
PROTECTION-PSWD----- (NULL)      RACF----- (NO)
ASSOCIATIONS
  DATA-----JHSRC.ENCTWX02.C0000001.DATA
  INDEX-----JHSRC.ENCTWX02.C0000001.INDEX
```

zDMF Multi-LPAR Sysplex Operation

A zDMF server needs to be running on every system that can do I/O to any data set being migrated/encrypted. Other than the STC or JOB name, the server JCL is the same for each system in the SYSPLEX. ZDMF will detect and report when a zDMF server has terminated and will also report on individual servers during GROUP phase changes such as ACTIVATE and DIVERT.

Example of a D HOSTS command.

```
RESPONSE=TDM1
GZD1100I Command received from userid: JHH61
GZD1101I  D HOSTS
CPUID      SMFID SSID CMD-PFX  HEARTBEAT
322CE73906 TDM2 ZDJ4 J4      01:49:00 09/13/18
052CE73906 TDM1 ZDJ4 J4      01:49:05 09/13/18
```

Example of a D HOSTS after a server has been brought down. At this point GROUPs will not ACTIVATE or DIVERT until the server is either removed 'J4 REMOVE TDM2' or brought back up.

```
GZD2027I Server ZDJ4TDM2 now considered IDLE
J4 D HOSTS
GZD1100I Command received from userid: JHH61 608
GZD1101I  D HOSTS
CPUID      SMFID SSID CMD-PFX  HEARTBEAT
322CE73906 TDM2 ZDJ4 J4      01:51:05 09/13/18 IDLE
052CE73906 TDM1 ZDJ4 J4      01:52:15 09/13/18
```

zDMF ISPF Panels

To start the ISPF interface EXEC 'xxx.IBM.HGZD340.SGZDELIB(GZDZDMF)'. On your first entry you will be prompted for user session options. Your command prefix is the connection from ISPF to your server.

The HELP panels have most any detail needed. Note that PF5 will save your settings.

```

User Session Options
Command ==> 09/13/18 01:33:31
zDMF Command Prefix . . . . . J4
Group Definition Data Set . . . JHH60.LDMF.CNTL
Use Log Data Set . . . . . Y (Y/N) Retain 07 generations of logs (01-30)
  Use Browse or View . . . . . V (B/V)
  Log Data Set Name . . . . . JHH60.V33X.LOG
Messages with Local or GMT time. L (L/G)
Show Command Messages . . . . . N (Y/N)
Show Command Diagnostic Info . . N (Y/N)
Create ICKDSF TRKFMT Statements. N (Y/N)
Early Data Set Completion . . . N (Y/N)
zDMF Load Library . . . . . JHH60.IBM.HGZD340.SGZDLLIB
zDMF Rexx Library . . . . . JHH60.IBM.HGZD340.SGZDELIB
zDMF Panel Library . . . . . JHH60.IBM.HGZD340.SGZDPLIB
zDMF Table Library . . . . . JHH60.IBM.HGZD340.SGZDTLIB
zDMF Message Library . . . . . JHH60.IBM.HGZD340.SGZDMLIB
zDMF Security Library . . . . . JHH60.IBM.HGZD340.SGZDLLIB
      User model JOB CARD for Z option clean up JCL
//JOBNAME JOB CLASS=A
//* CLEAN UP JCL JOBCARD

F1=Help F3=Exit F5=Save Settings F12 = Cancel
```

zDMF ISPF Panels

The ISPF main panel is shown below:

```
.      z/OS Dataset Mobility Facility      .
.  Function Number      09/13/18 01:35:29  .
.  or Command ==> _      J4              .
.      zDMF 3.4.0 Functions      TDM1    .
.                                  .
.      1 Manage Groups          .
.      2 Interact with Promoted Groups .
.      3 Display Host Messages  .
.      4 System Change Summary  .
.      5 Security               .
.      6 Set User Session Options .
.      7 Display Installation Options .
.      8 Message Help           .
.      9 Monitor Command Line Help .
.     10 View or Browse Log Data Set .
.     11 SMF Reporting          .
.                                  .
.      PF3=Exit                .
```

Option 2 is where most of the work gets done by ACTIVATING, DIVERTing and monitoring GROUPS. Don't be fooled by Option 1, it just allows you to create JCL/GROUP definitions in a data set and then PROMOTE those GROUPs to your database.

zDMF ISPF Panels

The Option 2 panel is shown below.

```

. Subsystem J4 Groups Row 1 to 19 of 32
. Command ==> _ Scroll ==> CSR
.
. Group TDM1
. Data Set Status
. Extents
.
. ENCTWX02 Error
. Reason Code 1026 Incomplete group
. Enter group command M for more information
.
. ENCTWX04 Inactive
. Activate, Divert, and Terminate reports available
.
. ENCTWX05 Complete
. Activate, Divert, and Terminate reports available

```

The help screens for the Option 2 panel are shown below:

```

- z/OS Dataset Mobility Facility
PF3=End PF7=Page Up PF8=Page Down More: +
PF 1 = Help
PF 4 = Set User Display Options
PF 6 = Set Group and/or Data Set Filters
PF 10 = Toggles Filtering between On and Off
       Initial panel display has Filtering Off
PF 11 = Switches to option 1, manage groups

       Command codes to select zDMF functions
Group Line Data Set Line
A Activate Group
B Display Simulation Report
D Deactivate Group
G Display Group Information
J Show Data Sets and Job
  Names Allocated to them
L Group Object Detail List
M Display Group Activation Messages
N Job Names Allocated to
  all Data Set(s) in Group
R Resume Group
S Suspend Group
E List Extent Object Details
G Display Data Set Information
J Job Names Allocated to this Data Set
L List Data Set Object Details

```

zDMF - Non-extended to Extended Format conversion example

The new zDMF GROUP parameter CONVERT_TO_EXTENDED_FORMAT(YES/NO) can be used with or without encryption to convert non-extended format data sets to extended format.

```
GROUP (BASIC12) -  
  DELETE_EXISTING_TARGET_DATASETS (YES) -  
  EARLY_DATA_SET_COMPLETION (NO)  
SET -  
  SPHERE (YES) -  
  TOLERATE_SOURCE_VALIDATE_FAILURE (NO) -  
  MIGRATE_ONLY_SPECIFIED_VOLUMES (NO) -  
  CONVERT_TO_EXTENDED_FORMAT(YES) -  
  ENCRYPT(YES) -  
  SOURCE (DSN (JHSRC.BASIC12.C00000001)) -  
  TOLERATE_TARGET_CREATE_FAILURE (NO) -  
  TARGET (DSN (JHTGT) -  
    STORAGECLASS(LDMFTSC1))
```

Note: Converting a data set to extended format adds a 32 byte suffix to the data set. Some data sets may have a blocksize that would prevent the addition of 32 bytes to their physical blocksize and those data sets will not qualify for this feature

zDMF – How to plan for Pervasive Encryption

Before you encrypt a single data set:

1. You must be committed to Pervasive Encryption (at least for some applications)
2. You must have identified a 'pilot' application to try

An example scenario for encrypting a group of data sets follows.

1. The pilot application General Ledger is selected with a data set name mask of PROD.GL.**.
2. The KEYLABEL (encryption key name) must be defined to RACF or SMS). Once this happens data sets that recreated will start to be encrypted.
3. There are 600 PROD.GL data sets identified in a generated list.
 - 100 are DB2 related - (use DB2 RE-ORG to encrypt)
 - 100 are GDGs that get encrypted as they are created
 - 100 are other sequential and VSAM data sets that get recreated, and encrypted, during normal application processing
4. That leaves 300 data sets that are eligible for encryption with zDMF.
5. Assuming 1 hour for each data set to: unload it, delete it, define it, reload it, monitor the jobs, ensure you have access.....

Spend 300 staff hours, or use zDMF....DFSMSdss can't do this

Roadmap & Strategy

IBM Z Data set Migration looking forward

zDMF Data Migration Solutions Roadmap



DELIVERED

- Dynamic data set level encryption
- Support for moving multi-volume data sets to EAV Extended Address
- Toleration support for zHyperlink

- Key rotation
- Basic/large format conversion to extended format
- Smart data set selector (native support or interface with other 21st products)

- Support additional data set types for auto-completion (no application stop/start required)
- Dynamic data set compression option
- Move catalogs non-disruptively

Key Takeaways

Answers to user challenges



TDMF and zDMF are the **gold standard** tools for **non-disruptive** movement of data

TDMF provides **more than just migration** capabilities; other use cases bring **day-to-day value**

zDMF data set encryption capabilities **reduce effort and service interruption** required to protect data



Further information

- IBM Transparent Data Migration Facility z/OS: <https://www.ibm.com/us-en/marketplace/transparent-data-migration-facility>
- z/OS Data Set Mobility Facility: <https://www.ibm.com/us-en/marketplace/zos-data-set-mobility-facility>

Notices and disclaimers

- © 2019 International Business Machines Corporation. No part of this document may be reproduced or transmitted in any form without written permission from IBM.
- **U.S. Government Users Restricted Rights — use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM.**
- Information in these presentations (including information relating to products that have not yet been announced by IBM) has been reviewed for accuracy as of the date of initial publication and could include unintentional technical or typographical errors. IBM shall have no responsibility to update this information. **This document is distributed “as is” without any warranty, either express or implied. In no event, shall IBM be liable for any damage arising from the use of this information, including but not limited to, loss of data, business interruption, loss of profit or loss of opportunity.** IBM products and services are warranted per the terms and conditions of the agreements under which they are provided.
- IBM products are manufactured from new parts or new and used parts.
In some cases, a product may not be new and may have been previously installed. Regardless, our warranty terms apply.”
- **Any statements regarding IBM's future direction, intent or product plans are subject to change or withdrawal without notice.**
- Performance data contained herein was generally obtained in a controlled, isolated environments. Customer examples are presented as illustrations of how those
- customers have used IBM products and the results they may have achieved. Actual performance, cost, savings or other results in other operating environments may vary.
- References in this document to IBM products, programs, or services does not imply that IBM intends to make such products, programs or services available in all countries in which IBM operates or does business.
- Workshops, sessions and associated materials may have been prepared by independent session speakers, and do not necessarily reflect the views of IBM. All materials and discussions are provided for informational purposes only, and are neither intended to, nor shall constitute legal or other guidance or advice to any individual participant or their specific situation.
- It is the customer's responsibility to insure its own compliance with legal requirements and to obtain advice of competent legal counsel as to the identification and interpretation of any relevant laws and regulatory requirements that may affect the customer's business and any actions the customer may need to take to comply with such laws. IBM does not provide legal advice or represent or warrant that its services or products will ensure that the customer follows any law.

Notices and disclaimers continued

- Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products about this publication and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products. IBM does not warrant the quality of any third-party products, or the ability of any such third-party products to interoperate with IBM's products. **IBM expressly disclaims all warranties, expressed or implied, including but not limited to, the implied warranties of merchantability and fitness for a purpose.**
- The provision of the information contained herein is not intended to, and does not, grant any right or license under any IBM patents, copyrights, trademarks or other intellectual property right.
- IBM, the IBM logo, ibm.com and [names of other referenced IBM products and services used in the presentation] are trademarks of International Business Machines Corporation, registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at: www.ibm.com/legal/copytrade.shtml