

zEnterprise Data Compression Usage and Configuration


DFSMSdss/DFSMSHsm

Robert Gensler
rgensle@us.ibm.com
IBM

August 6, 2014
Session: 16130



#SHAREorg



Legal Disclaimer



NOTICES AND DISCLAIMERS

Copyright © 2013 by International Business Machines Corporation.

No part of this document may be reproduced or transmitted in any form without written permission from IBM Corporation.

Product information and data has been reviewed for accuracy as of the date of initial publication. Product information and data is subject to change without notice. This document could include technical inaccuracies or typographical errors. IBM may make improvements and/or changes in the product(s) and/or programs(s) described herein at any time without notice.

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. Consult your local IBM representative or IBM Business Partner for information about the product and services available in your area.

Any reference to an IBM Program Product in this document is not intended to state or imply that only that program product may be used. Any functionally equivalent program, that does not infringe IBM's intellectual property rights, may be used instead. It is the user's responsibility to evaluate and verify the operation of any non-IBM product, program or service.

THE INFORMATION PROVIDED IN THIS DOCUMENT IS DISTRIBUTED "AS IS" WITHOUT ANY WARRANTY, EITHER EXPRESS OR IMPLIED. IBM EXPRESSLY DISCLAIMS ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT. IBM shall have no responsibility to update this information. IBM products are warranted according to the terms and conditions of the agreements (e.g., IBM Customer Agreement, Statement of Limited Warranty, International Program License Agreement, etc.) under which they are provided. IBM is not responsible for the performance or interoperability of any non-IBM products discussed herein.



Trademarks

The following are trademarks of the International Business Machines Corporation in the United States and/or other countries.

BookManager*	Enterprise Storage Server*	IP PrintWay	RMF
CICS*	ES/9000*	Language Environment*	S/370
DB2*	FlashCopy*	Lotus*	S/390*
DB2 Universal Database	GDPS*	Multiprise*	Tivoli*
developerWorks*	HiperSockets	MVS	TotalStorage*
DFSMSdftp	IBM*	Notes*	WebSphere*
DFSMSdss	IBM eServer	OS/390*	z/Architecture
DFSMSHsm	IBM e(logo)server*	Parallel Sysplex*	z/OS*
DFSMSrmm	IBM logo*	RACF*	zSeries*
DFSORT	IMS	RAMAC*	
Domino	InfoPrint*		

Intel is a trademark of the Intel Corporation in the United States and other countries.

Java and all Java-related trademarks and logos are trademarks or registered trademarks of Sun Microsystems, Inc., in the United States and other countries.

Microsoft, Windows and Windows NT are registered trademarks of Microsoft Corporation.

UNIX is a registered trademark of The Open Group in the United States and other countries.

* All other products may be trademarks or registered trademarks of their respective companies.

Notes:

Performance is in Internal Throughput Rate (ITR) ratio based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput improvements equivalent to the performance ratios stated here.

IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply.

All customer examples cited or described in this presentation are presented as illustrations of the manner in which some customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics will vary depending on individual customer configurations and conditions.

This publication was produced in the United States. IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the product or services available in your area.

Information about non-IBM products is obtained from the manufacturers of those products or their published announcements. IBM has not tested those products and cannot confirm the 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.

Prices subject to change without notice. Contact your IBM representative or Business Partner for the most current pricing in your geography.

This presentation and the claims outlined in it were reviewed for compliance with US law. Adaptations of these claims for use in other geographies must be reviewed by the local country counsel for compliance with local laws.

Agenda

- Compression overview
- Support for new zEDC format data sets
- DFSMSdss exploitation of zEDC
- DFSMShsm exploitation of zEDC

Compression Overview



- Data encoding has been an area of research and development even before magnetic storage devices were commercially available
 - Claude E. Shannon first articulated the theory of data compression in his paper “The Mathematical Theory of Communication” in 1948
 - Four years later in 1952 David A. Huffman published his work on optimal encoding “A Method for the Construction of Minimum-Redundancy Codes”

Compression Overview

Data Compression will become pervasive

- I/O throughput is struggling to keep up with increasingly data driven applications
- Batch workloads are accessing more data from disk and network connections
- Business opportunities can be lost due to cost prohibitive nature of keeping data online

Data needs to be shared across platforms

- Data is being exchanged among business partners
- Compression can substantially reduce the amount of data transferred
- Industry standard formats need to be used for transparent peer to peer communication

Solves problems in the enterprise

- Improves the effective throughput of data over storage and communication networks
- Allows more data to remain online
- Less data to perform encryption against
- Makes high cost/byte storage technology such as flash memory more affordable

IBM zEnterprise Data Compression (zEDC)

What is it?

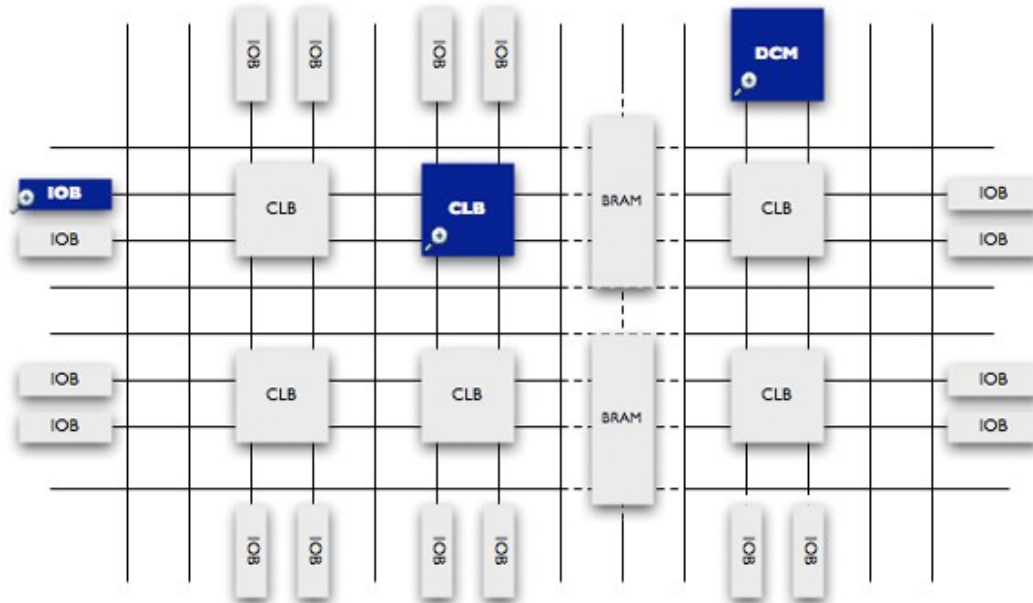
- A combined software (z/OS V2R1 plus PTFs) and hardware (zEDC Express) solution designed to help reduce resource consumption, storage utilization and optimize cross platform exchange of data

How is it different?

- Performance: Reduced CPU overhead
- Efficient: Optimized algorithms scan text to locate the re-use of phrases and refers back to earlier references
- Industry Standard: Compatible with open zlib based compression – widely used across all platforms
- Economical: Reduced storage requirements and improved effective bandwidth without significant CPU overhead

zEDC Overview

- Field Programmable Gate Arrays (FPGAs) are programmable semiconductor devices that are based around a matrix of configurable logic blocks (CLBs) connected via programmable interconnects.
 - As opposed to Application Specific Integrated Circuits (ASICs) where the device is custom built for the particular design, FPGAs can be programmed to the desired application or functionality requirements.



<http://www.xilinx.com/company/gettingstarted/index.htm>

zEDC Configuration

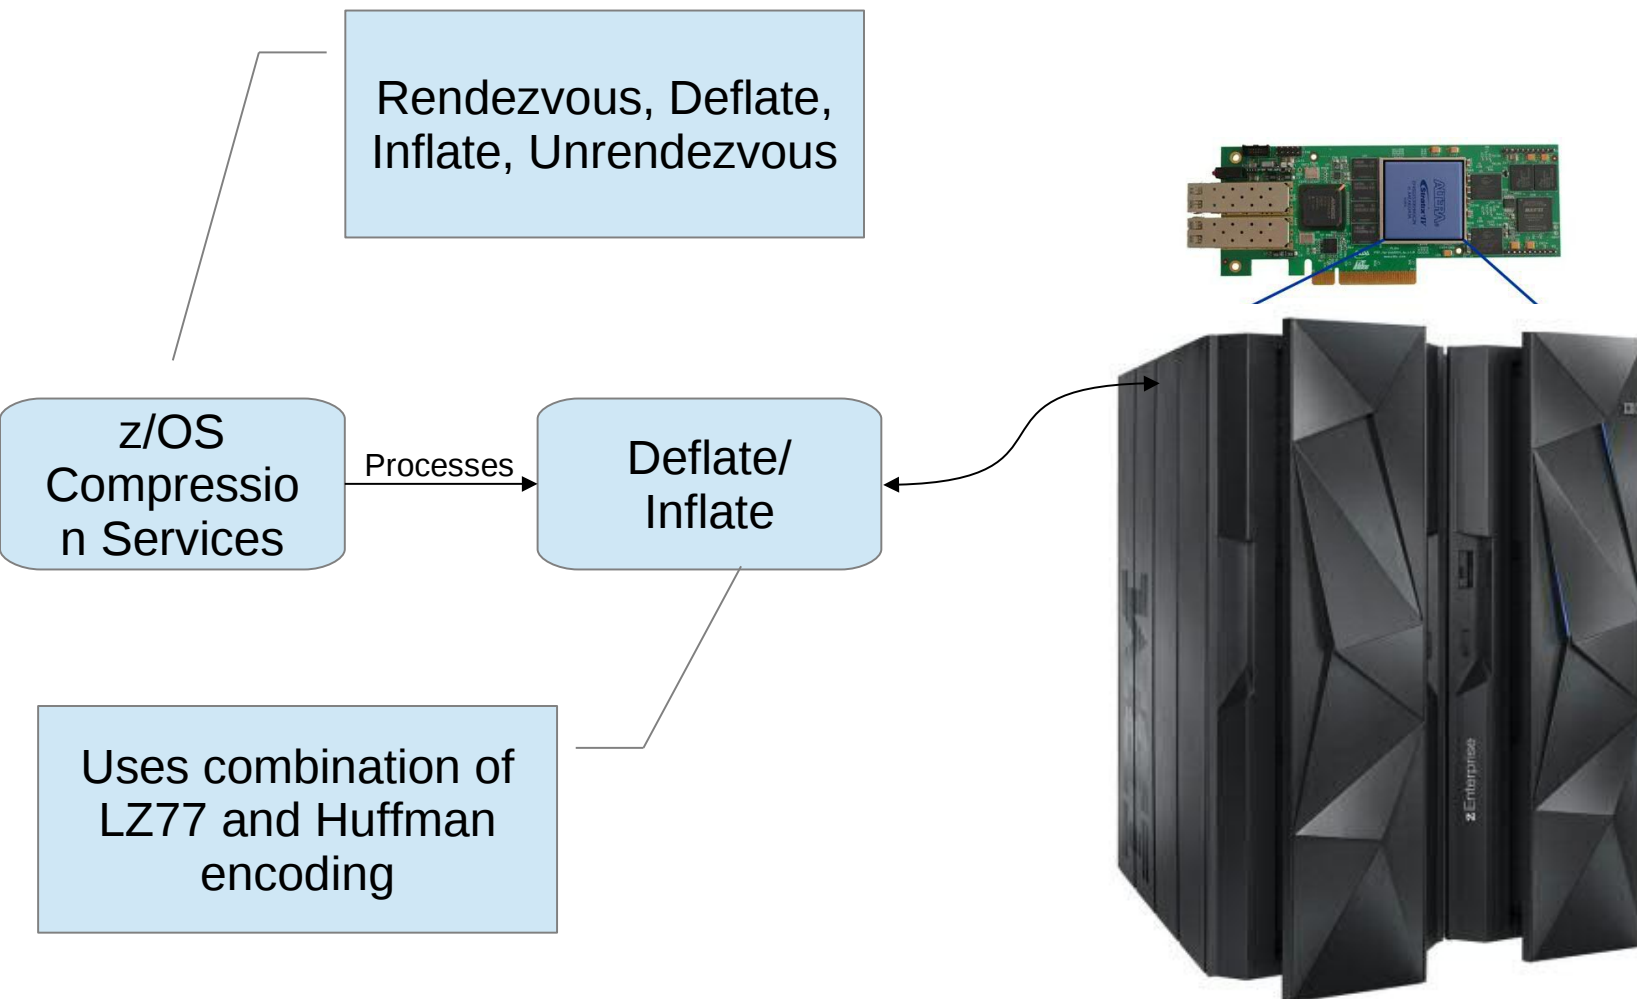
- **Operating system requirements**
 - Requires z/OS V2R1
 - *w/ PTFs for new zEDC Express for z/OS feature*
 - PTFs for zEDC sequential data set support
 - *Coexistence PTFs for z/OS V1.13 and V1.12*
 - PTFs for DFSMSdss/DFSMSHsm zEDC exploitation*
 - *Coexistence PTFs for z/OS V1.13 and V1.12**
- Coexistence PTFs offers software decompression support only
- *Planned availability 3Q2014

zEDC Configuration

- **Server requirements**

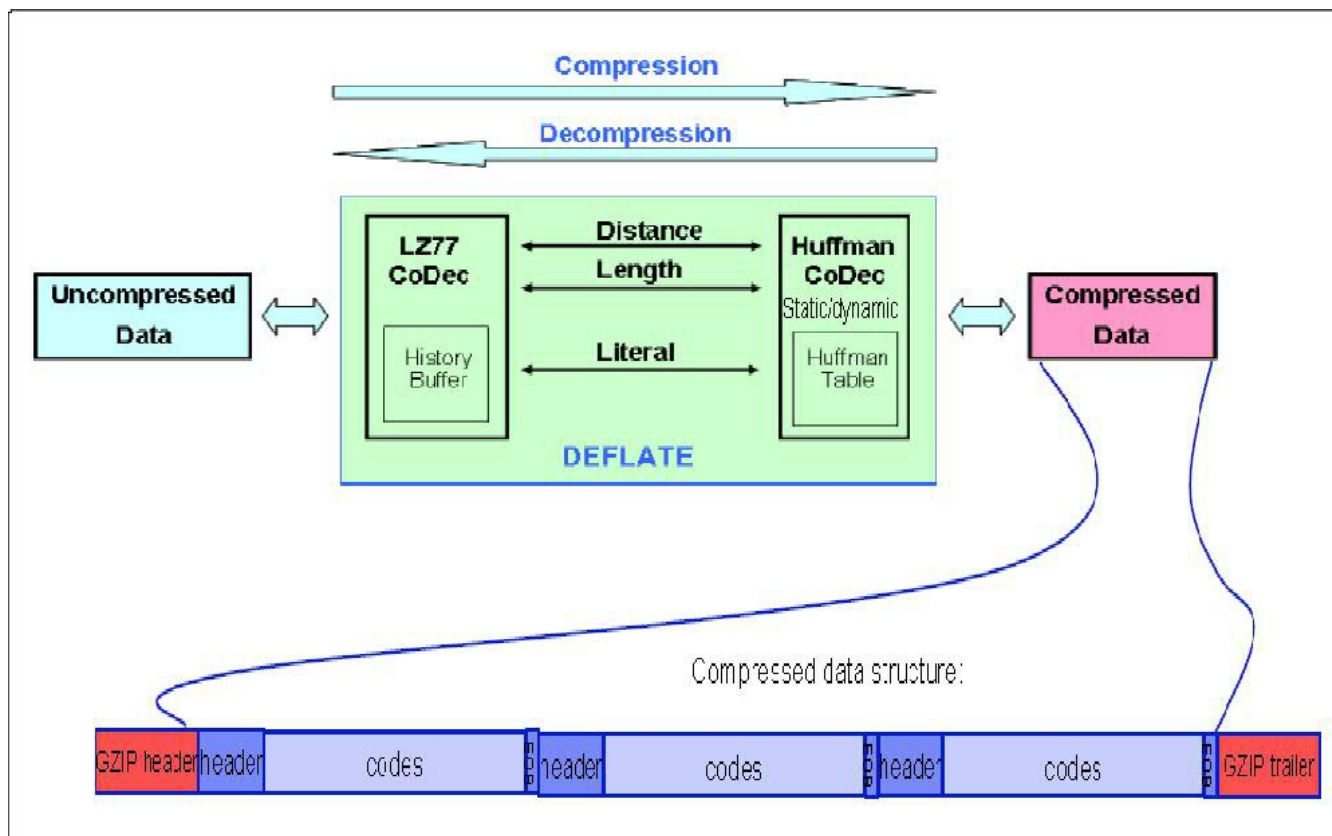
- Available on zEC12 and zBC12
- New zEDC Express feature for PCIE I/O drawer (FC#0420)
 - Each feature can be shared across up to 15 LPARs
 - Up to 8 features available on zEC12 or zBC12
- Recommended high availability configuration per server is four features
 - *Provides 4GB/s of compression/decompression*
 - *Provides high availability during concurrent update*
 - *Half the devices unavailable during update*
 - *Recommended minimum configuration per server is two features*

Constructing the zEDC Stack



zEDC Overview

- Refer to Dale Riedy's and Anthony Sofia's 2014 SHARE in Ahaheim presentation “zEDC: What is it and how do I use it?”



Agenda

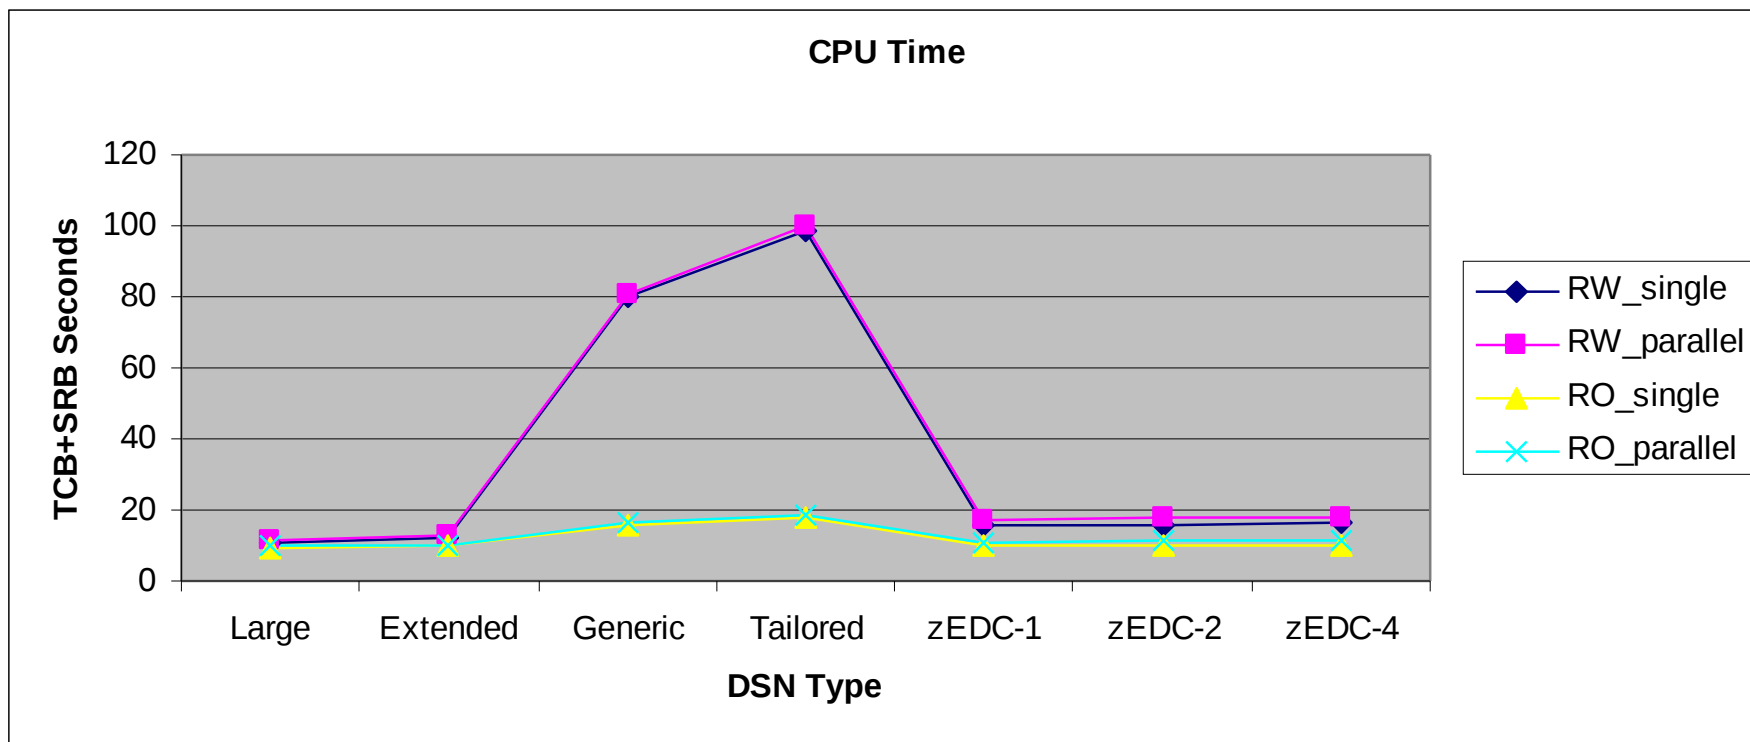
- Compression overview
- Support for new zEDC format data sets
- DFSMSdss exploitation of zEDC
- DFSMShsm exploitation of zEDC

zEDC Sequential Data Set Value

- **z/OS DFSMS (BSAM/QSAM) introduces a new type of compression (zEDC) for non-VSAM extended format data sets.**
- **Value: For customers who don't use BSAM/QSAM compression today**
 - Take advantage of the disk space savings available through zEDC compression with minimal CPU overhead.
- **Value: For customers who currently do use BSAM/QSAM compression**
 - The CPU cost of compressing BSAM/QSAM data can be reduced when using zEDC compression compared to existing BSAM/QSAM compression options.
 - Note that the disk space savings may vary depending on the type of compression used.

zEDC Sequential Data Set Value

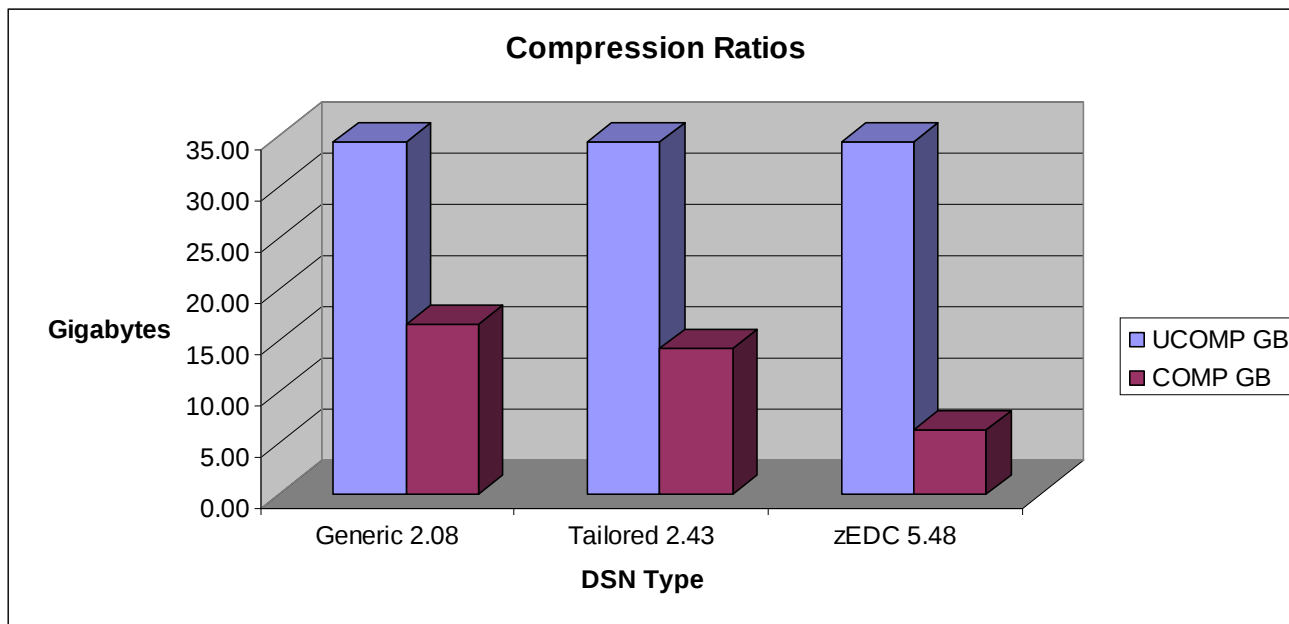
zEDC CPU time comparison



Disclaimer: Based on projections and/or measurements completed in a controlled environment. Results may vary by customer based on individual workload, configuration and software levels.

zEDC Sequential Data Set Value

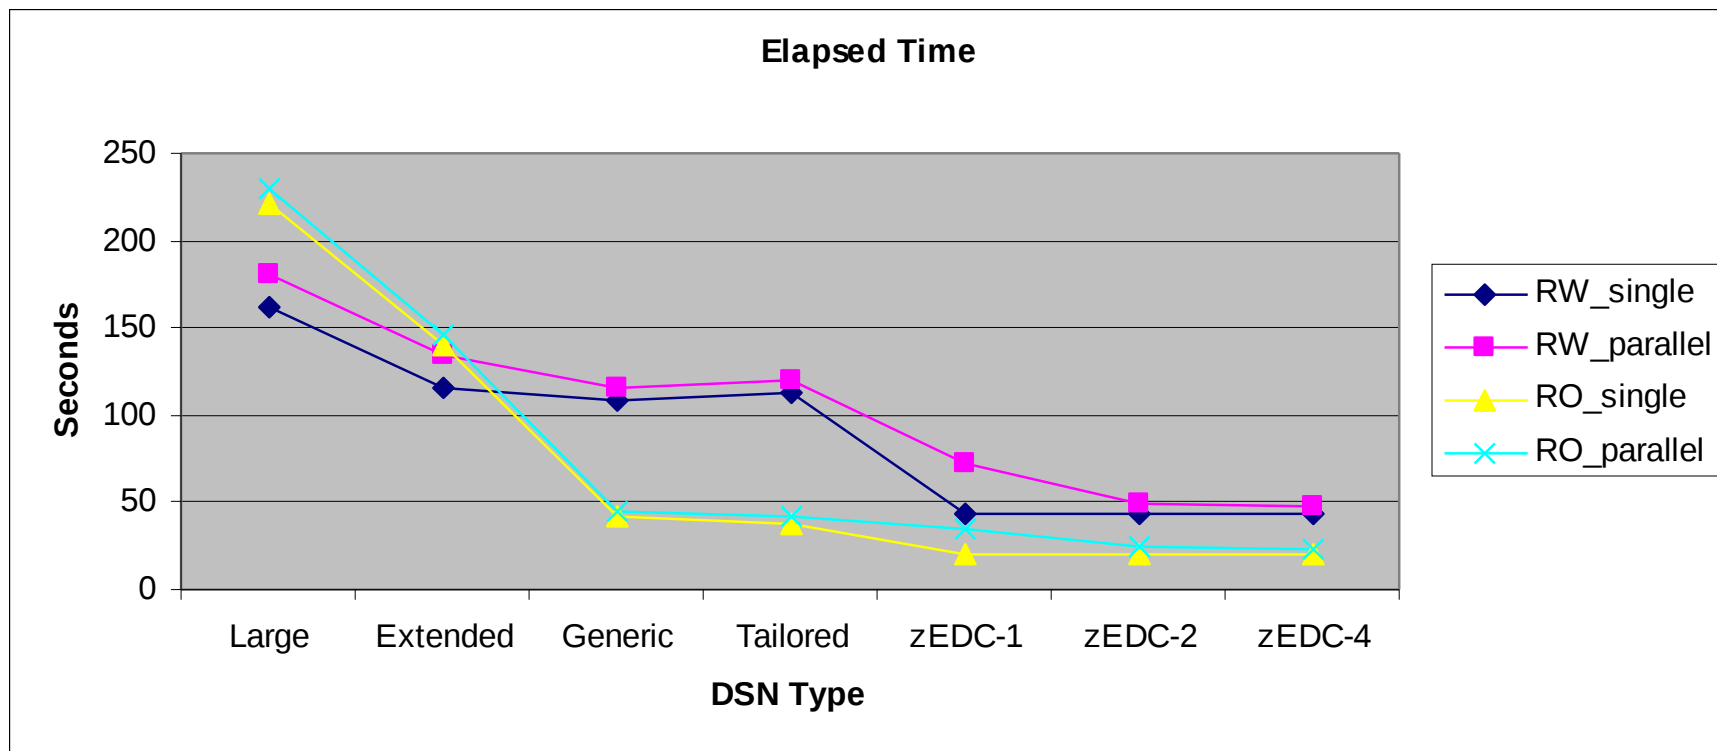
zEDC Compression ratios compared to generic and tailored compression



Disclaimer: Based on projections and/or measurements completed in a controlled environment. Results may vary by customer based on individual workload, configuration and software levels.

zEDC Sequential Data Set Value

zEDC elapsed time comparison



Disclaimer: Based on projections and/or measurements completed in a controlled environment. Results may vary by customer based on individual workload, configuration and software levels.

Supporting New zEDC Sequential Data Sets

- DFSMSdss supports new zEDC sequential data sets in all of its functions
 - COPY, DUMP, RESTORE, CONSOLIDATE, DEFRAG, PRINT, RELEASE
- DFSMShsm supports new zEDC sequential data sets in all of its functions
 - MIGRATION, RECALL, BACKUP, RECOVER, FULL-VOLUME DUMP, RECOVER FROMDUMP, FRBACKUP, FRRECOV, ABACKUP, ARECOVER

Supporting New zEDC Sequential Data Sets

- REBLOCK not supported for zEDC data sets
 - DFSMSdss ADRREBLK installation exit is not called for zEDC data sets
- When copying or restoring a data set the compression type of the source is carried forward to the target

Target exists	Target format	Source format	Results
No	N/A	zEDC	zEDC
Yes	generic	zEDC	zEDC
Yes	tailored	zEDC	zEDC

Coexisting with z/OS V1R12 and z/OS V1R13

- DFSMSdss
 - supports restoring zEDC format sequential data sets
- DFSMSHsm
 - supports RECALL, RECOVER, RECOVER from DUMP, FRRECOVER from DUMP, and ARECOVER
- Data is not decompressed on the restore so there is no additional CPU overhead
 - Read access to the data on the down-level systems will require software inflate
 - Writes will be performed without compressing the data

Agenda

- Compression overview
- Support for new zEDC format data sets
- DFSMSdss exploitation of zEDC
- DFSMSHsm exploitation of zEDC

Comparing Compression Options

- DFSMSDss currently supports two forms of compression
 - COMPRESS - Straight forward compression algorithm using no special purpose facilities
 - *Value: minimal CPU overhead*
 - HWCOMPRESS - Uses special purpose CMPSC instruction of System Z processor
 - *Value: Typically better compression ratio*
- zEDC provides compression with minimal CPU overhead and competitive compression ratios

Manually Creating Backups Using zEDC

- New keyword, ZCOMPRESS, provided for the DUMP command
 - Accepted for all FULL, TRACKS, physical and logical DATASET backups to DASD and tape
- Has three sub-parameters
 - REQUIRED – backup must be created using zEDC
 - *If zEDC accelerators cannot be used then the backup should fail*
 - PREFERRED – the backup should be created using zEDC
 - *If zEDC accelerators cannot be used then the backup should continue anyway*
 - NONE – backup should not be created using zEDC

Combining Compression Options

- ZCOMPRESS(PREFERRED) can be specified in conjunction with COMPRESS or HWCOMPRESS
 - *DUMP FULL INDY(VOL001) OUTDD(DDUMPOUT)
ZCOMPRESS(PREF) COMPRESS*
- DFSMSDss will attempt to use zEDC accelerators but if none are available it will revert to using the algorithm provided by the COMPRESS or HWCOMPRESS keyword
 - A backup is either compressed entirely with zEDC accelerators or entirely with the COMPRESS or HWCOMPRESS algorithms

Realizing the Variations

ZCOMP	HWCOMPRESS	COMPRESS	Results
Pref	No	No	Attempt to use zEDC. If not available continue creating backup uncompressed
Pref	No	Yes	Attempt to use zEDC. If not available continue creating backup using COMPRESS
Pref	Yes	No	Attempt to use zEDC. If not available continue creating backup using HWCOMPRESS
Pref	Yes	Yes	Invalid Combination

System Programmer Considerations

- The use of zEDC can be overridden using the Installation Options Exit Routine (ADRUIXIT)
 - Can disable/enable usage of zEDC
 - Can override ZCOMPRESS sub-parameters
- If combinations of compression are overridden in ADRUIXIT
 - zEDC will take precedence

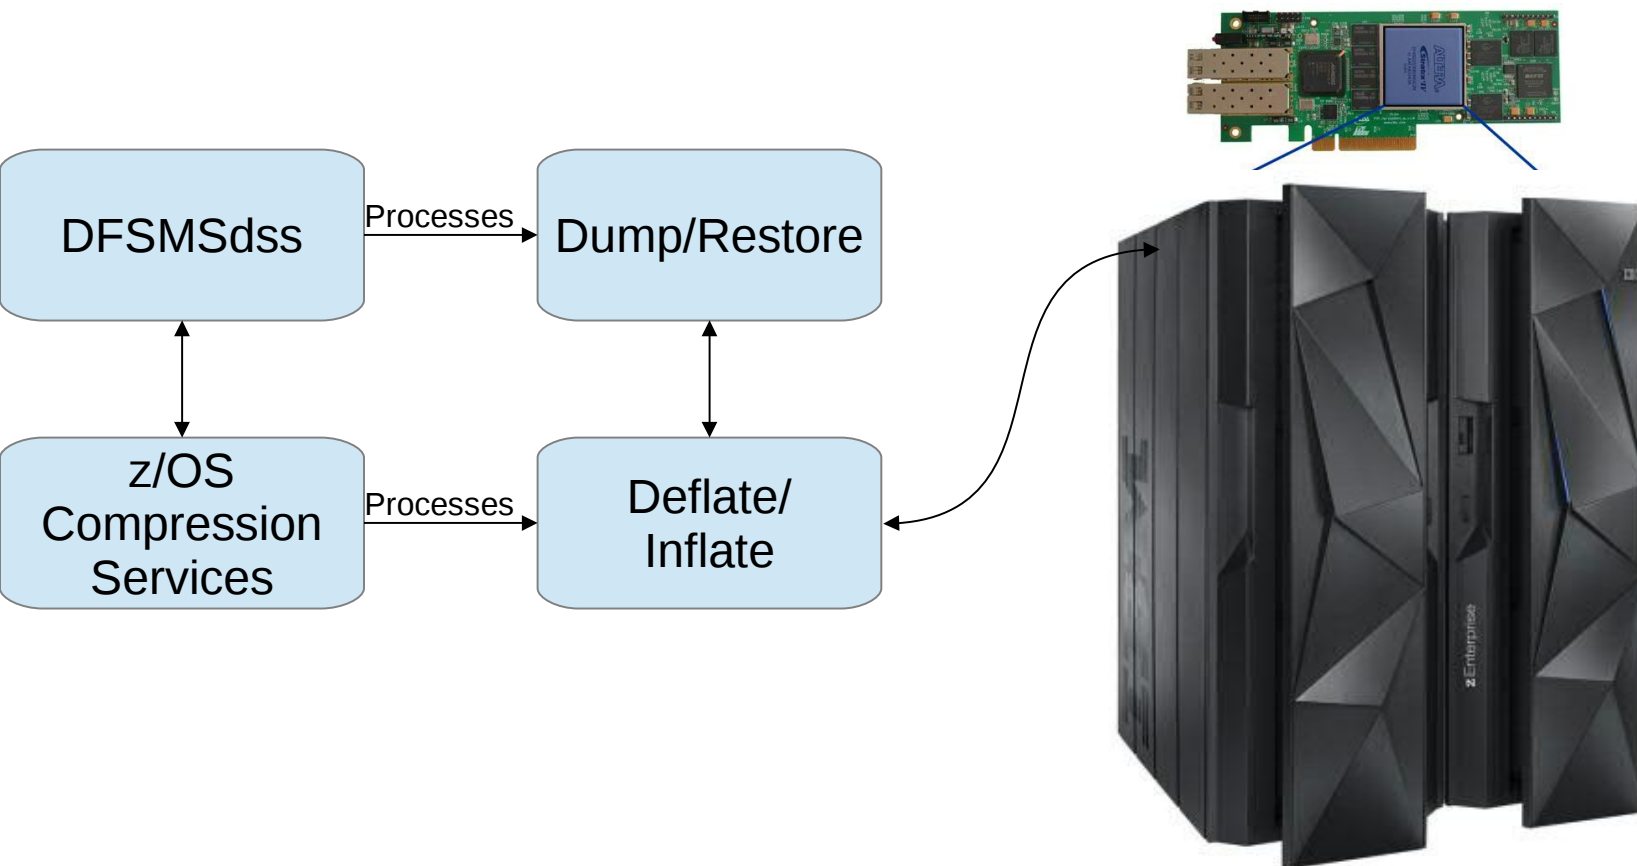
System Programmer Considerations

- The use of zEDC for backups can be restricted using a new facility class profile
 - STGADMIN.ADR.DUMP.ZCOMPRESS
- If the profile does not exist then anyone can use zEDC accelerators when creating backups
- Can use the profile for simple fixed cost charge-back schemes

Manually Recovering Backups

- Backups created using zEDC accelerators are automatically decompressed
 - No new keywords for the RESTORE command
- Use either zEDC accelerators or software inflate
- Any user can restore a backup created using zEDC
 - No facility class profile to restrict the use of zEDC during RESTORE

Constructing the zEDC Stack



Usage Considerations

- Avoid specifying any sort of compression options on a dump command when
 - backing up zEDC data sets individually
 - *DFSMSdss will attempt to compress the data of an already compressed data set*
 - the output dump data set is a zEDC format data set
 - *DFSMSdss will compress user data blocks and then pass them to BSAM which will also try to compress them*
- When using ZCOMPRESS avoid creating backup on tape that is also performing compression

Application Programming Considerations

- DFSMSdss informs applications that zEDC accelerators will be used when creating a backup
 - Exit 0 – Function Startup
 - *New bit ei00zcomp indicates whether or not zEDC accelerators will be used*
- Compression ratio presented to Applications that invoke DFSMSdss
 - Exit 14 – Function ending
 - *New 1 byte field ei14zcsv indicates compression ratio for the backup*
 - *Valid values 0-99*

Application Programming Considerations

- New feature bit added to ADRMCLVL
 - On return from call to ADRMCLVL DFSMSDss places feature related information in general register 14
 - *Byte 0, bit 1 – when on indicates the PTFs for DFSMSDss zEDC support are installed*
- See the 'How to determine DFSMSDss version, release and modification level' section of the Storage Administration manual

Agenda

- Compression overview
- Support for new zEDC format data sets
- DFSMSdss exploitation of zEDC
- DFSMShsm exploitation of zEDC

Compression in DFSMShsm Today

- Currently, DFSMShsm does not use DFSMSdss to compress user data during migration and backup
 - Uses its own host based compression algorithm to compress user data
- DFSMShsm does use DFSMSdss COMPRESS or HWCOMPRESS to compress user data during full-volume dump

Exploiting zEDC Support in DFSMSHsm

- DFSMSHsm will use the DFSMSdss zEDC support in
 - Migrate/Recall
 - Backup/Recover
 - Full Volume DUMP
 - Recover and FRRECOV from DUMP

Exploiting zEDC Support in DFSMShsm

- DFSMShsm will call DFSMSdss with the ZCOMPRESS(PREFERRED) option
- DFSMShsm will not use zEDC Services during migration or backup functions when DFSMShsm is the data mover
 - Partitioned Data Sets will utilize the standard DFSMShsm compaction methodology in place.

Configuring the Use of zEDC (SETSYS)

- SETSYS ZCOMPRESS
 - ALL | NONE
 - DASDBACKUP (NO | YES)
 - DASDMIGRATE (NO | YES)
 - TAPEBACKUP (NO | YES)
 - TAPEMIGRATE (NO | YES)
- COMPACTPERCENT works in conjunction with ZCOMPRESS as it does for COMPACT

Combining SETSYS Options

- If zEDC services are not available at the time of the backup or migration then DFSMShsm will look at the values specified in the COMPACT SETSYS parameter

ZCOMPRESS	COMPACT	Results
None	None	DFSMShsm will create backup or migrate data set without using any form of compression
None	All	DFSMShsm will create backup or migrate data set by using its current form of compression
All	None	DFSMShsm will attempt to use zEDC services to compress backup. If the services are unavailable then backup will be uncompressed
All	All	DFSMShsm will attempt to use zEDC services to compress backup. If the services are unavailable then backup will be compressed using its current form of compression

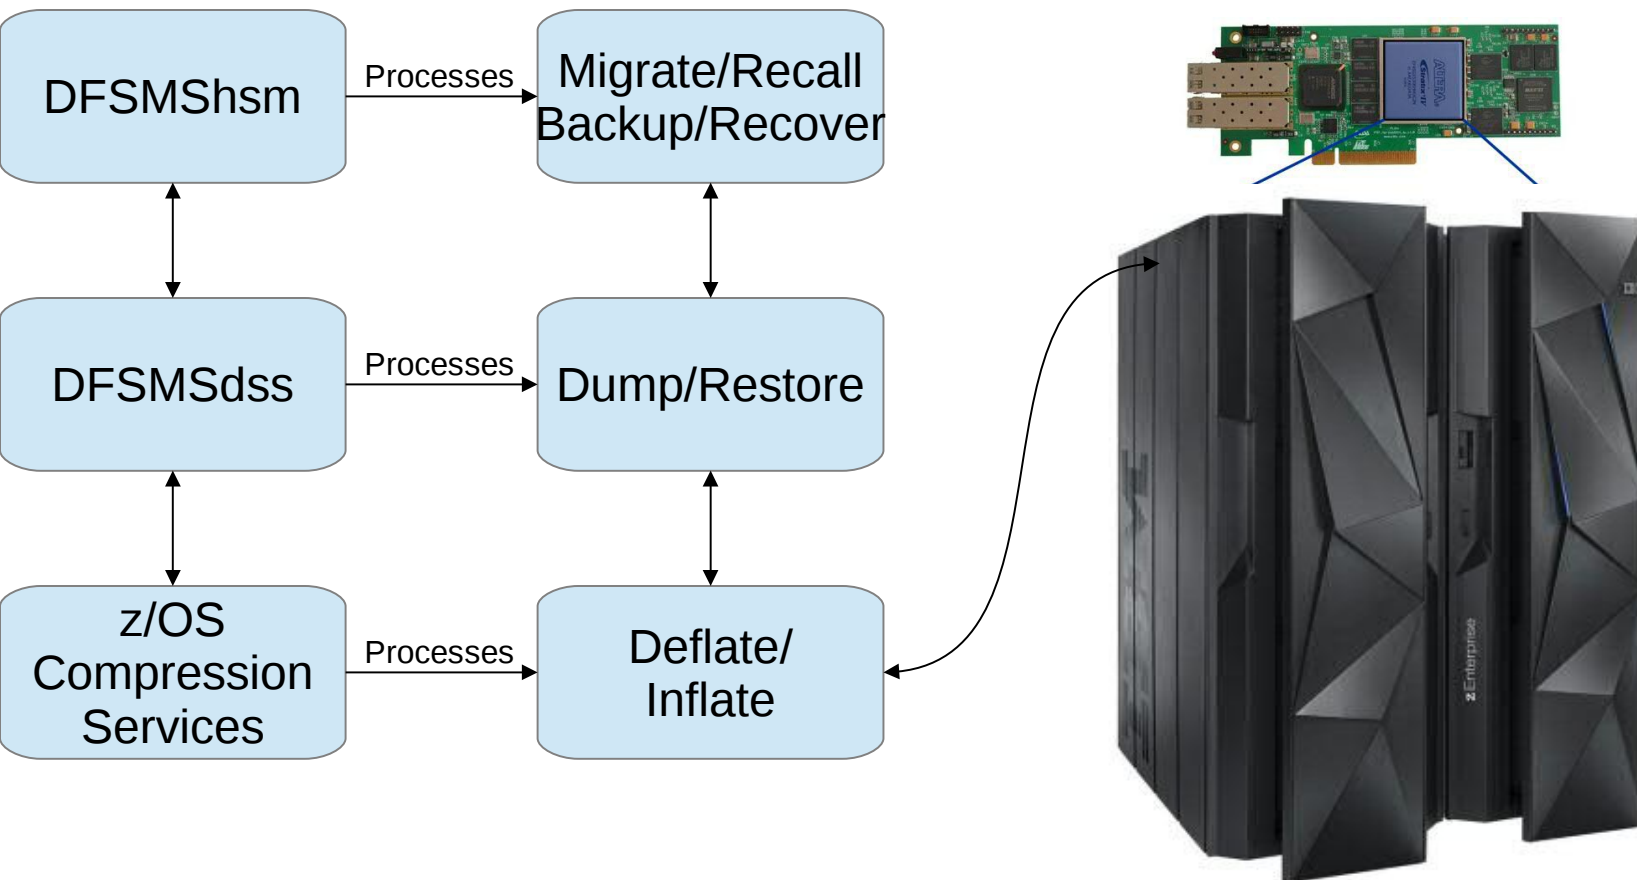
Controlling ZCOMPRESS for Volume Dumps

- DEFINE DUMPCLASS(ZCOMPRESS(NO | YES))
 - Valid for BACKVOL and FRBACKUP when DUMP is specified
- If zEDC hardware is available, the DFSMSdss is invoked using the ZCOMPRESS(PREFERRED) option.
 - In the case of a zEDC hardware failure, the dump may or may not be compressed depending on the other DUMPCLASS options

Controlling ZCOMPRESS for Volume Dumps

- If the ZCOMPRESS keyword is specified through its dump class and
 - the patch to use COMPRESS is specified
 - *DFSMSHsm specifies both the ZCOMPRESS(PREFERRED) and the COMPRESS keywords in the DFSMSdss DUMP command.*
 - the patch to use HWCOMPRESS is specified
 - *DFSMSHsm specifies both the ZCOMPRESS(PREFERRED) and the HWCOMPRESS keywords in the DFSMSdss DUMP command.*

Constructing the zEDC Stack



Overriding SETSYS for Individual Data Sets

- ARCMBEXT – volume and data set backup
 - Can bypass compression for a particular data set when
 - *SETSYS(ZCOMPRESS(ALL))*
 - *SETSYS(ZCOMPRESS(DASDBACKUP(YES)))*
 - *SETSYS(ZCOMPRESS(TAPEBACKUP(YES)))*
- ARCMDEXT – volume and data set migration
 - Can bypass compression for a particular data set when
 - *SETSYS(ZCOMPRESS(ALL))*
 - *SETSYS(ZCOMPRESS(DASDBACKUP(YES)))*
 - *SETSYS(ZCOMPRESS(TAPEBACKUP(YES)))*

Recovering Data using zEDC

- Since DFSMSdss is the data mover, DFSMSHsm will use zEDC services to automatically decompress data
 - Even if the use of zEDC has been disabled using SETSYS or DEFINE DUMPCLASS

Coexisting on V1R12 and V1R13

- DFSMSdss
 - Will allow a user V1R12 and V1R13 release to restore backups created using zEDC services
 - *Software inflate is used*
- DFSMSHsm
 - Will allow a user on a V1R12 and V1R13 release to RECALL, RECOVER, RECOVER from DUMP, or FRRECOV from DUMP data sets migrated, backed up or dumped using zEDC Services on V2R1
 - *Will leverage the coexistence support provided by DFSMSdss*

References

- PTFs for zEDC exploitation or software decompression have a fix category of IBM.Function.zEDC
- DFSMSDss APARs
 - zEDC format sequential data set support
 - *OA42198 – UA72755 HDZ2210*
 - *OA43817 – UA72769 HDZ1D10, UA72768 HDZ1C10*
 - Partial Release Reporting error (OA45229)
 - zEDC exploitation
 - OA42238 – PTFs targeted for 3Q2014
 - will contain PTFs for HDZ2210, HDZ1D10, HDZ1C10

References

- PTFs for zEDC exploitation or software decompression have a fix category of IBM.Function.zEDC
- DFSMSHsm APARs
 - zEDC format sequential data set support
 - *Leverages support provided by DFSMSDss*
 - zEDC exploitation
 - OA42243 – PTFs targeted for 3Q2014
 - will contain PTFs for HDZ2210, HDZ1D10, HDZ1C10

Related Sessions

- **Session 16138: DFSMS Exploitation of z/OS zEnterprise Data Compression (Barb McDonald from IBM)**
 - Monday August 4, 2014: 1:30PM-2:30PM Room 305
- **Session 15671: System z Batch Network Analyzer (zBNA) Tool Hands-on Lab**
 - Thursday August 7, 2014: 4:15PM-5:15PM Room 301
- **Session 15709: System Z Performance: More Than You Think (Steve Grabarits from IBM)**
 - Wednesday August 6, 2014: 3:00PM-4:00PM Room 303

References

- DFSMSdss Storage Administration
- DFSMSHsm Storage Administration
- DFSMS Installation Exits
- zEnterprise Data Compression FAQ
 - <https://www-304.ibm.com/support/docview.wss?uid=tss1fq131484&aid=1>
- Configuring zEDC on z/OS (a cheat sheet) by Anthony Sofia
 - https://www-304.ibm.com/connections/blogs/systemz/entry/configuring_zedc_on_zos?lang=en_us

System Z Social Media Channels

Top Facebook pages related to System z:

IBM System z
IBM Academic Initiative System z
IBM Master the Mainframe Contest
IBM Destination z
Millennial Mainframer
IBM Smarter Computing

Top LinkedIn groups related to System z:

System z Advocates
SAP on System z
IBM Mainframe- Unofficial Group
IBM System z Events
Mainframe Experts Network
System z Linux
Enterprise Systems
Mainframe Security Gurus

Twitter profiles related to System z:

IBM System z
IBM System z Events
IBM DB2 on System z
Millennial Mainframer
Destination z
IBM Smarter Computing

YouTube accounts related to System z:

IBM System z
Destination z
IBM Smarter Computing

Top System z blogs to check out:

- Mainframe Insights
- Smarter Computing
- Millennial Mainframer
- Mainframe & Hybrid Computing
- The Mainframe Blog
- Mainframe Watch Belgium
- Mainframe Update
- Enterprise Systems Media Blog
- Dancing Dinosaur
- DB2 for z/OS
- IBM Destination z
- DB2utor



zEnterprise Data Compression Usage and Configuration

DFSMSdss/DFSMSHsm

Robert Gensler
rgensle@us.ibm.com
IBM

August 6, 2014
Session: 16130



#SHAREorg

