

DFSMS Exploitation of the z/OS zEnterprise Data Compression (zEDC)

Barbara McDonald, IBM - bawhite@us.ibm.com

Glenn Wilcock, IBM – wilcock@us.ibm.com

March 2, 2014
Session 17108



SHARE is an independent volunteer-run information technology association
that provides **education**, **professional networking** and **industry influence**.



Agenda

- **DFSMS Compression with zEDC**
 - Compression Overview
 - Why zEDC?
 - Identifying Candidates for zEDC
 - Configuring and Deploying
 - QSAM / BSAM Exploitation
 - DSS / HSM Exploitation

For additional information, refer to the following SHARE presentations:

- [16138: DFSMS Exploitation of z/OS zEnterprise Data Compression](#) (Pittsburgh) - Cecilia Carranza Lewis, IBM
carranc@us.ibm.com
- [16130: z/OS zEnterprise Data Compression Usage and Configuration for DSS and HSM](#) (Pittsburgh) - Robert Gensler, IBM
rgensle@us.ibm.com

Complete your session evaluations online at www.SHARE.org/Seattle-Eval

Compression Overview



Every day over 2000 petabytes of data are created

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

Compression solves problems in the enterprise

- Improves the effective throughput of data over storage and communication networks
- Allows more data to remain online for increased business value
- Reduces the amount of data for encryption operations
- Typically improves batch turnaround
- Make storage technology including Flash Memory more affordable
- Reduces the CPU and storage required for offline (backup/archive) data

Complete your session evaluations online at www.SHARE.org/Seattle-Eval

Other Use Case Examples & Related Sessions



Other Example Use Cases

- **SMF Archived Data** can be stored compressed to increase the amount of data kept online up to 4X
 - **zSecure** output size of Access Monitor and UNLOAD files reduced up to 10X and CKFREEZE files reduced by up to 4X
 - Up to 5X more ***XML*** data can be stored in sequential files
 - *The IBM Employee Directory* was stored in up to 3X less space
 - **z/OS SVC and Stand Alone DUMPs** can be stored in up to 5X less space
- **16635: MVS Core Technologies Project (MVSE) Opening and WSC Hot Topics**
 - Monday, 11:15 AM-12:15 PM
 - **17035: MQ for z/OS, Using and Abusing New Hardware and the New V8 Features [z/OS]**
 - Wednesday, 10:00 AM-11:00 AM
 - **16805: System z Batch Network Analyzer (zBNA) Tool – Because Batch is Back!**
 - Thursday, 10:00 AM-11:00 AM
 - **16816: RMF: The Latest and Greatest**
 - Thursday, 3:15 PM-4:15 PM

Complete your session evaluations online at www.SHARE.org/Seattle-Eval



IBM z Enterprise Data Compression

New data compression offering that can reduce resource usage



What is it?

- ✓ *zEDC Express is an IO adapter that does high performance industry standard compression*
- ✓ *Used by z/OS Operating System components, IBM Middleware and ISV products*
- ✓ *Applications can use zEDC via industry standard APIs (zlib and Java)*
- ✓ *Each zEDC Express sharable across 15 LPARs, up to 8 devices per CEC.*
- ✓ *Raw throughput up to 1 GB/s per zEDC Express Hardware Adapter*

What Changes?

It is time to revisit your decisions about compression.

- **Disk Savings:** Many people are already getting value from CMPSC compression and software compression today
- **Performance:** High throughput alternative to existing System Z compression for large or active files.
- **Industry Standard:** Low cost compressed data exchange across all platforms
- **Pervasive:** Standard APIs allow quick adoption by middleware products running on System Z


What is the Value?

New sources of customer value

- **QSAM/BSAM** compression can save up to 4x disk space and in some cases shorten elapsed time, reducing batch windows.
- **DFSMSdss** exploitation reduces the number of tape cartridges by up to 19% when performing full volume dumps to tape.
- **DFSMSHsm** exploitation reduces CPU by up to 80% for Migration to ML1 disk, and reduces disk space by up to 58% (compared to using the COMPACT option).
- **Business Partner Data Exchange** can have higher throughput with lower CPU cost
- **Managed File Transfer** saves up to 4x link bandwidth, and up to 80% elapsed time.
- **ISV Products** delivery expanded customer value
- **Java for z/OS R7.1** accelerates common compression classes used by applications and middleware
- **High availability** during concurrent update

Complete your session evaluations online at www.SHARE.org/Seattle-Eval

Compression Coprocessor (CMPSC) vs zEDC



SHARE
Educate • Network • Influence

Using the right hardware compression accelerator for each of your workloads

Compression Coprocessor

z Enterprise Data Compression

On Chip In every IBM eServer™ zSeries® today (and tomorrow) Mature: Decades of use by Access Methods and DB2® Work is performed jointly by CPU and Coprocessor Propriety Compression Format	PCIe Adapter New with IBM zEnterprise® EC12 GA2 and IBM zEnterprise BC12 Mature: Industry Standard with decades of software support Work is performed by the PCIe Adapter Standards Compliant (RFC1951)	
Use Cases		
<u>Small object compression</u> ▪ Rows in a database	<u>Large Sequential Data</u> ▪ QSAM/BSAM Online Sequential Data ▪ Objects stored in a data base	<u>Industry Standard Data</u> ▪ Cross Platform Data Exchange
<u>Users</u> ▪ VSAM for better disk utilization ▪ DB2 for lower memory usage ▪ The majority of customers are currently compressing their DB2 rows	<u>Users</u> ▪ QSAM/BSAM for better disk utilization and batch elapsed time improvements ▪ SMF for increased availability and online storage reduction ▪ DFSMSdss for better disk and tape utilization for backup data ▪ DFSMSHsm for improved CPU, disk reduction ▪ DFSMSdss for improved dump to tape	<u>Users</u> ▪ Java for high throughput standard compression via java.util.zip ▪ Encryption Facility for z/OS for better industry data exchange ▪ IBM Sterling Connect: Direct® for z/OS for better throughput and link utilization ▪ ISV support for increased client value


IBM System z Batch Network Analyzer

Helping determine if you have files that are candidates for zEDC



- **IBM System z Batch Network Analyzer 1.4.2**

- A free, Microsoft Windows-based “as is” tool to analyze batch windows using SMF data
- Available to Customers, Business Partners and IBMers
- Replaces the old BWATOOL
- PC based; provides graphical and text reports
 - Including Gantt charts and support for Alternate Processors



- **Available Now on TechDocs**

- <http://w3.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/PRS5132>

- **zBNA identifies zEDC Compression Candidates**

- Post-process customer provided SMF records, to identify jobs and data sets which are zEDC compression candidates across a specified time window, typically a batch window
- Help estimate utilization of a zEDC feature and help size number of features needed
- Generate a list of data sets by job which already do hardware compression and may be candidates for zEDC
- Generate lists of data sets by job which might be zEDC candidates but are not in extended format

Don't forget System z Batch Network Analyzer (zBNA) Tool – Because Batch is Back! (Session 16805)

Complete your session evaluations online at www.SHARE.org/Seattle-Eval



Configuring and Deploying zEDC



- **Operating system requirements**

- Requires z/OS 2.1 (w/PTFs) and the new zEDC Express for z/OS feature
 - PTFs for BSAM/QSAM exploitation (APAR OA42195 / UA72749)
 - PTFs for DSS for sequential data set support (OA42198 / UA72755)
 - PTFs for DSS exploitation (APAR OA42238 / UA74782)
 - PTFs for HSM exploitation (APAR OA42243 / UA74836)
 - PTFs for IOS (APAR OA45054 / UA74718)
- z/OS V1.13 and V1.12 offer software decompression support only
 - PTFs for BSAM/QSAM co-existence (APAR OA43863 / UA72750 (R12) UA72751 (R13))
 - PTFs for DSS sequential data set co-existence (OA43817/ UA72769 (R12) UA72768 (R13))
 - PTFs for DSS/HSM co-existence (APAR OA42238 / UA74780 (R12) UA74781 (R13))
- zEDC Express for z/OS feature must be enabled in an IFAPRDxx PARMLIB member.

Complete your session evaluations online at www.SHARE.org/Seattle-Eval



Configuring and Deploying zEDC



- **Server requirements**
 - Exclusive to 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
 - This provides up to 4GB/s of compression/decompression
 - Provides high availability during concurrent update (half devices unavailable during update)
 - Recommended minimum configuration per server is two features
 - Steps for installing EDC Express in an existing zEC12/zBC12
 - Apply z/OS Service; Hot plug a zEDC Express adapter; IODF updates and Dynamic Activate

Additional Configuration Considerations




Considerations for allocating a *non-VSAM compressed format data set*

- DFSMS does not require the zEDC Express feature to be available in order to allocate a zEDC compressed format data set
 - DFSMS cannot differentiate between a system where the feature is not configured and a system where the feature is configured but not available.
- If not on a zEC12/zBC12, there are times where DFSMS may still allocate a zEDC compressed format data set.
- *In these cases,*
 - *BSAM/QSAM will write data non-compressed;*
 - *BSAM/QSAM will continue to invoke zEDC to decompress existing compressed data, however, zEDC will use software decompression.*

zEDC should be installed on all systems accessing compressed data*

**** For the full zEDC benefit, zEDC should be active on ALL systems that might access or share compressed format data sets. This eliminates instances where software inflation would be used when zEDC is not available.***

BSAM/QSAM zEDC Compression – Value!



* 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.

Complete your session evaluations online at www.SHARE.org/Seattle-Eval

QSAM/BSAM Data Set Compression with zEDC

*Improved compression ratios and performance**



- **Reduce the cost of keeping your sequential data online**
 - zEDC compresses data up to 4X, saving up to 75% of your sequential disk space.
 - That's as much as 2X better than generic or tailored compression
 - Capture new business opportunities due to lower cost of keeping data online.
- **Better I/O elapsed time for sequential access**
 - Potentially run batch workloads faster than either uncompressed or QSAM/BSAM current compression.
- **Sharply lower CPU cost over existing compression**
 - Enables more pervasive use of compression.
 - Up to 80% reduced CPU cost compared to tailored and generic compression options.
 - CPU cost for zEDC is ~0.1sec/GB in testing on a zEC12
- **Simple Enablement**
 - Use policy to enable zEDC.

** 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.*


Complete your session evaluations online at www.SHARE.org/Seattle-Eval



QSAM/BSAM Parallel Batch Jobs

For parallel QSAM/BSAM workloads that begin to show queuing on a single adapter, adding multiple adapters brings the queue delay back to zero.

The percent usage and throughput (GB/s) per adapter is highlighted in the bar chart.



* 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.

Complete your session evaluations online at www.SHARE.org/Seattle-Eval

QSAM/BSAM zEDC – Setup

- **zEDC setup is similar to that used for existing types of compression (generic and tailored compression) are requested.**
 - It can be selected at either or both the data class level or system level.
 - **Data class level**
 - In addition to existing Tailored (T) and Generic (G) values, new **zEDC Required (ZR)** and **zEDC Preferred (ZP)** values will be available on the COMPACTION option in data class.
 - When COMPACTION=Y in data class, the system level is used.
 - **System level**
 - In addition to existing TAILORED and GENERIC values, new **zEDC Required (ZEDC_R)** and **zEDC Preferred (ZEDC_P)** values will be available on the COMPRESS parameter found in TGDSMSxx member of SYS1.PARMLIB.
 - **Activated using SET SMS=xx or at IPL**
 - Data class continues to take precedence over system level. The default continues to be GENERIC.
- **QSAM/BSAM Conversion**
 - Existing QSAM/BSAM data sets (whether compressed or not) must be copied to a new target data set allocated with zEDC compression.
 - No utility available to perform a conversion without de-compressing source and re-compressing target.
 - Normal tools can be used to perform the copy, for example IEBGENER or REPRO
 - DSS supports COPY, DUMP, RESTORE, DEFrag, CONSOLIDATE, RELEASE, PRINT

Request via SMS Data Class

Panel Utilities Scroll Help

DGTDCDC9 DATA CLASS DEFINE/ALTER Page 3 of 6
Command ==>

SCDS Name . . . : USER6.MYSCDS
Data Class Name : DC93

To DEFINE/ALTER Data Class, Specify:

Compaction	—	(Y, N, T, G, ZR , ZP or blank)
Spanned / <u>Nonspanned</u>		(S, N or blank)
System Determined <u>Blocksize</u>	N	(Y or N)
EATTR		(O, N or blank)

Use ENTER to Perform Verification; Use UP/DOWN Command to View other Panels;
Use HELP Command for Help; Use END Command to Save and Exit; CANCEL to Exit.

- **New Compaction values**

- **ZP: Prefer zEDC compression.** The system will not fail the allocation request but rather create either a tailored compressed data set if the zEDC function is not supported by the system or create a non-compressed extended format data set if the minimum allocation amount requirement (5MB, or 8MB Primary if no Secondary) is not met.
- **ZR: Require zEDC compression.** The system will fail the allocation request if the zEDC function is not supported by the system or the minimum allocation amount requirement (5MB, or 8MB Primary if no Secondary) is not met.

Complete your session evaluations online at www.SHARE.org/Seattle-Eval

Request via SMS Parmlib

- To request the use of zEDC compression at the system level (COMPACTION=Y in Data Class) when creating new compressed format data sets, new values are defined for the COMPRESS parameter found in IGDSMSxx member of SYS1.PARMLIB.

—

— **COMPRESS(TAILORED|GENERIC|zEDC_R|zEDC_P)**

- zEDC_R tells the system to fail the allocation request if the zEDC function is not supported by the system or the minimum allocation amount requirement (5MB, or 8MB Primary if no Secondary) is not met.
- zEDC_P tells the system to not fail the allocation request but rather create either a tailored compressed data set if the zEDC function is not supported by the system or create a non-compressed extended format data set if the minimum allocation amount requirement (5MB, or 8MB Primary if no Secondary) is not met.
- *Therefore, when the default system level is modified to specify zEDC_R or zEDC_P, all new allocation requests using a data class with COMPACTION=Y will request zEDC compression.*

QSAM/BSAM zEDC – Setup cont.



- **Use of Extended Format Version 2**
 - zEDC Compressed Format data sets are created as Version 2 data sets ...regardless of the user's specification in DataClass, JCL or SYS1.PARMLIB.
 - Extended format V2 data sets are new in V2.1.
 - Created to allow DFSMSdss support for FlashCopy® when copying sequential, non-striped, multivolume EF V2 data sets
 - The new zEDC compression for new extended format data sets is **Optional**
 - All previous compression options are still supported.

Identifying zEDC Compressed Format Data Set

Dictionary token identifies type of compression

- **First two bytes of token**

– Generic Token	.10.	.000		0000	X' 4000'
– Tailored Token	.11.	.xxx		0000	X' 6x00'
– zEDC Token	.11.	.000		0001	X' 6001'
– Rejection Token 1...					X' 8000'

Note: The above is not an API....but we'll let you know if this changes

- **Dictionary token displayed in**

- LISTCAT
- DCOLLECT

- **New one-byte Compression Type found in**

- SMF 14/15 records
- ****May also be available via additional APIs...stay tuned****

Complete your session evaluations online at www.SHARE.org/Seattle-Eval

Vendor Considerations



- New DFA Fields**

- For vendors who may want to use new zEDC compressed format data sets within their product, new fields were added to the DFA (mapped via IAHDFA) to indicate DFSMS zEDC SW installed

60 (3C)	Bit string	4	DFAFEAT9	Features byte 9
	1... ..		DFAJ3AA	JES3_ALLOC_ASSIST ENABLED
	.1.. ..		DFAMEMUX	Reserved
	...1.		DFAPDSEG	PDSE Generation support is installed
	...1		DFAZEDCCMP	zEDC Compression support is installed
	 xxxx			Reserved
...				
...				
80 (50)	Unsigned	1	DFACMPTYPEDEF	Default compression type (found in IGDSMSxx COMPRESS)
		0	DFACMPTYPEGEN	Generic compression
		1	DFACMPTYPETLRD	Tailored compression
		2	DFACMPTYPEzEDCR	zEDC compression required
		3	DFACMPTYPEzEDCP	zEDC compression preferred
81 (51)	Character	31		Reserved

Complete your session evaluations online at www.SHARE.org/Seattle-Eval

SMF



- **Use SMF records to determine compression ratio and type**
 - SMF Type 14/15 contains existing compressed/uncompressed bytes to calculate compression ratio.
 - SMF Type 14/15 defines new bits to identify zEDC compressed format data set, as well as indicate if zEDC Express is not available for compression/decompression during this OPEN.

SMF cont.

Compressed Format Data Set Section (Type 1)

This describes the information acquired for compressed format data sets.

Record Type 14

Offsets	Name	Length	Format Description
4(4)	SMF14XF1	1 binary	Bit Meaning When Set 0 Compressed format data set size values invalid. 1 Compression of the data set has been rejected. 2 SMF14zEDCNAC. At times, zEDC hardware not available for compression thus data written uncompressed. 3 SMF14zEDCNAD. At times, zEDC hardware not available for decompression thus data decompressed using software 4-7 Reserved...
6(6)	SMF14CDL	8 binary	Number of bytes of compressed data read or written since this open.
14 (E)	SMF14UDL	8 binary	Number of bytes of data read or written since this open (data length prior to compression).
22 (16)	SMF14CDS	8 binary	Size of the compressed format data set (number of compressed user data bytes).
30 (1E)	SMF14UDS	8 binary	Size of the compressed format data set (number of uncompressed user data bytes).

(Highlighted fields are new with zEDC SPE)

Complete your session evaluations online at www.SHARE.org/Seattle-Eval

SMF cont.

Compressed Format Data Set Section (Type 1)

This describes the information acquired for compressed format data sets.

Record Type 14

Offsets	Name	Length	Format Description
80 x'50'	SMF14CMPTYPE	1 binary	Compression Type
			Meaning When Set
		0	SMF14CMPTYPENA Not compressed format or Unknown
		1	SMF14CMPTYPEGEN Generic Compression
		2	SMF14CMPTYPETLRD Tailored Compression
		3	SMF14CMPTYPEZEDC zEDC Compression

(Highlighted fields are new with zEDC SPE)

Complete your session evaluations online at www.SHARE.org/Seattle-Eval

DFSMSdfp Coexistence



- Coexistence PTFs for DFSMS components to support access to zEDC compressed format data sets.
 - DFSMS coexistence PTFs will allow a user on V1R12 and V1R13 to read/write zEDC compressed format data sets
 - Writes will write user data non-compressed
 - Reads of compressed data will use software decompression

Complete your session evaluations online at www.SHARE.org/Seattle-Eval



DSS and HSM Support for zEDC

- **DFSMSdss**

- Supports DUMP, RESTORE, COPY, DEFRAG, CONSOLIDATE, RELEASE, and PRINT of zEDC compressed format data sets
- Supports new compression keyword for DUMP


- **DFSMShsm**

- Supports MIGRATION, RECALL, BACKUP, RECOVER, FULL-VOLUME DUMP, RECOVER FROMDUMP, FRBACKUP, FRRECOV, ABACKUP, ARECOVER of zEDC compressed format data sets
- Uses the DFSMSdss zEDC support for certain functions

ML1 zEDC Compression – Value!

Up to 80% CPU Reduction
Up to 58% Space Reduction


½ the data for SSM
½ the data to Recycle



** 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.*

Complete your session evaluations online at www.SHARE.org/Seattle-Eval

ML1 zEDC Compression – Value!



* 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.

Complete your session evaluations online at www.SHARE.org/Seattle-Eval

DFSMSdss and DFSMSHsm Data Set Compression with zEDC



Cost effectively migrate and recall data with zEDC to get CPU and storage savings

- Reduce the amount of space required for backup data (*)
- When migrating data to ML1 Disk using zEDC, use up to 58% less disk space compared to using DFSMSHsm with the COMPACT keyword.¹

Better I/O elapsed time for backup/restore/migration/recall

- Reduces amount of data transferred on the channel

Sharply lower CPU cost over existing compression

- Enables more pervasive use of compression
- When migrating data to ML1 Disk using zEDC, use up to 80% less CPU compared to using DFSMSHsm with the COMPACT keyword.¹
- When using DFSMSHsm with zEDC, recalling data from ML1 Disk uses up to 69% less CPU as compared to using DFSMSHsm with the COMPACT option.¹

Reduce media used

- Customers using DFSMSdss with zEDC performing full volume dumps to tape can reduce the number of tape cartridges by almost 19%.¹

New opportunity for savings

- Reduce CPU cost while maintaining end to end encryption of backups using the DFSMSdss feature of the Encryption Facility

Simple Enablement

- New DFSMSdss DUMP keyword and DFSMSHsm SETSYS commands to enable zEDC.

(*) When backing up/archiving uncompressed data

Complete your session evaluations online at www.SHARE.org/Seattle-Eval



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. ¹Measurements for comparisons were completed as part of a formal performance evaluation on a dedicated, isolated test system.

DFSMSdss 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
- **The use of zEDC for backups can be restricted using a new facility class profile**
 - STGADMIN.ADR.DUMP.ZCOMPRESS
- **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

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

Combining DSS 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

Application Programming Considerations for DSS

- 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 for DSS

- 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

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

DFSMSHsm zEDC

- **DFSMSHsm will use the DFSMSdss zEDC support in**
 - Migrate/Recall
 - Backup/Recover
 - Dump, FRBACKUP Dump & Recover
- **SETSYS ZCOMPRESS**
 - ALL | NONE
 - DASDBACKUP (NO | YES), DASDMIGRATE (NO | YES)
 - TAPEBACKUP (NO | YES), TAPEMIGRATE (NO | YES)
- **DEFINE DUMPCLASS**
 - ZCOMPRESS(NO | YES)
- **COMPACTPERCENT works in conjunction with ZCOMPRESS as it does for COMPACT**

Combining SETSYS Options

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

DFSMShsm zEDC

- **DFSMShsm will always call DFSMSdss with the ZCOMPRESS(PREFERRED) option**
 - If zEDC is not available, then the other HSM compression options will be used
- **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

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.

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)))

DSS / HSM Coexistence



- **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

Thank You!



Complete your session evaluations online at www.SHARE.org/Seattle-Eval



Trademarks



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

AIX*	DFSMSdpp	DS6000	IBM*	MQSeries*	Redbooks*	System Storage	Tivoli*	z/OS*
BladeCenter*	DFSMSdss	DS8000*	IBM eServer	MVS	REXX	System x*	WebSphere*	zSeries*
BookManager*	DFSMSHsm	Easy Tier	IBM logo*	OS/390*	RMF	System z	z10 BC	
DataPower*	DFSMSrmm	FICON*	IMS	Parallel Sysplex*	SYSREXX	System z9*	z10 EC	
DB2*	DFSORT	FlashCopy*	InfinBand*	PR/SM	RMF	System z10	z/Architecture*	
DFSMS	Domino*	HiperSockets	Language Environment*	RACF*	SYSREXX	System z10 Business Class	zEnterprise*	

* Registered trademarks of IBM Corporation

The following are trademarks or registered trademarks of other companies.

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.

Cell Broadband Engine is a trademark of Sony Computer Entertainment, Inc. in the United States, other countries, or both and is used under license therefrom.

Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

IT Infrastructure Library is a registered trademark of the Central Computer and Telecommunications Agency which is now part of the Office of Government Commerce.

ITIL is a registered trademark, and a registered community trademark of the Office of Government Commerce, and is registered in the U.S. Patent and Trademark Office.

Java and all Java based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

Linear Tape-Open, LTO, the LTO Logo, Ultrium, and the Ultrium logo are trademarks of HP, IBM Corp. and Quantum in the U.S. and

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

OpenStack is a trademark of OpenStack LLC. The OpenStack trademark policy is available on the [OpenStack website](#).

TEALEAF is a registered trademark of Tealeaf, an IBM Company.

Windows Server and the Windows logo are trademarks of the Microsoft group of countries.

Worklight is a trademark or registered trademark of Worklight, an IBM Company.

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

* Other product and service names might be trademarks of IBM or other 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.

All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

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 information provides only general descriptions of the types and portions of workloads that are eligible for execution on Specialty Engines (e.g. zIIPs, zAAPs, and IFLs) ("SEs"). IBM authorizes customers to use IBM SE only to execute the processing of Eligible Workloads of specific Programs expressly authorized by IBM as specified in the "Authorized Use Table for IBM Machines" provided at www.ibm.com/systems/support/machine_warranties/machine_code/aut.html ("AUT"). No other workload processing is authorized for execution on an SE. IBM offers SE at a lower price than General Purpose Systems (e.g. zSeries). Processors because customers are authorized to use SEs only to process certain types and/or amounts of workloads as specified by IBM in the AUT.

Complete your session evaluations online at www.SHARE.org/Seattle-Eval



Trademarks and Disclaimers (continued)



NOTES:

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.

All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

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 are suggested US list prices and are subject to change without notice. Starting price may not include a hard drive, operating system or other features. Contact your IBM representative or Business Partner for the most current pricing in your geography.

Any proposed use of claims in this presentation outside of the United States must be reviewed by local IBM country counsel prior to such use.

The information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM makes no representation or warranty regarding third-party products or services including those designated as ServerProven, ClusterProven or BladeCenter Interoperability Program products. Support for these third-party (non-IBM) products is provided by non-IBM Manufacturers.

IBM may have patents or pending patent applications covering subject matter in this document. The furnishing of this document does not give you any license to these patents. Send license inquiries, in writing, to IBM Director of Licensing, IBM Corporation, New Castle Drive, Armonk, NY 10504-1785 USA.

Complete your session evaluations online at www.SHARE.org/Seattle-Eval

