DFSMS Exploitation of the z/OS zEnterprise Data Compression (zEDC): BSAM/QSAM Usage and Configuration

Barbara McDonald, Cecilia Carranza Lewis
IBM - bawhite@us.ibm.com
IBM – carranc@us.ibm.com

August 4, 2014
Session 16138

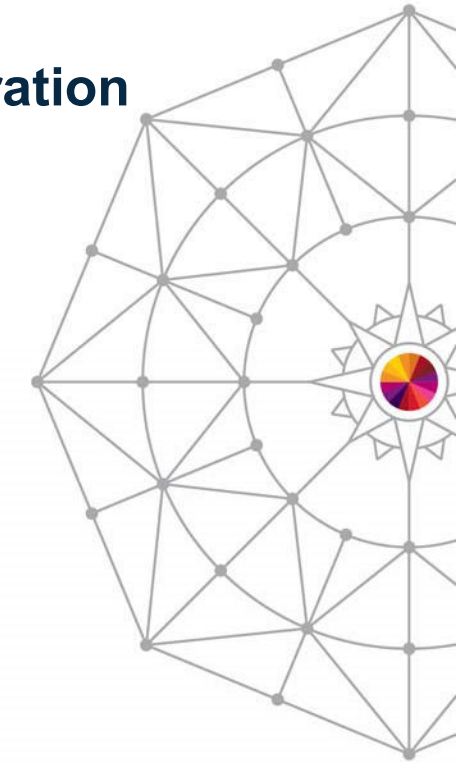

#SHAREorg

Agenda

- **BSAM/QSAM Compression with zEDC**
  - Why zEDC?
  - Configuring and Deploying
  - Identifying Candidates
  - Setup and Use
  - Coexistence

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.
- **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
- Improved availability with **SMF** compression.

ISVs Exploit zEDC

zEDC was expressly created using industry standard APIs to encourage ISVs to leverage its high-speed compression value in applications ISVs create. With access to zEDC, ISV applications are more valuable to end users.

ALEBRA	▪ **Alebra – Parallel Data Mover (PDM)** – Uses zEDC compression in lieu of Software-based compression to provide excellent qualities of service.
ASE	▪ **ASE** – OMCS – Takes SLIKZIP and SLIKSFTP performance to a whole new level
PKWARE	▪ **PKWARE** – PKZIP and SecureZIP v15 – Accelerated deflate compression and automatic detection of zEDC
software AG	**Software AG** – Entire Net-Work – High performance transaction processing

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

- **15572: MVS Core Technologies Project (MVSE) Opening and WSC Hot Topics**
  - Monday, 11:15 AM-12:15 PM
- **15674: Exploiting System z Innovation for Mainframe-based Managed File Transfer (MFT) with IBM Sterling Connect:Direct for z/OS**
  - Monday, 1:30 PM-2:30 PM
- **16130: z/OS zEnterprise Data Compression Usage and Configuration for DSS and HSM**
  - Wednesday, 10:00 AM-11:00 AM
- **15706: System z Batch Network Analyzer (zBNA) Tool – Because Batch is Back!**
  - Wednesday, 10:00 AM-11:00 AM
- **15671: System z Batch Network Analyzer (zBNA) Tool Hands-on Lab**
  - Thursday, 4:15 PM-5:15 PM
- **15709: System z Performance: More Than You Think**
  - Wednesday, August 6, 2014: 3:00 PM-4:00 PM

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.*

BSAM/QSAM zEDC Compression – Value!

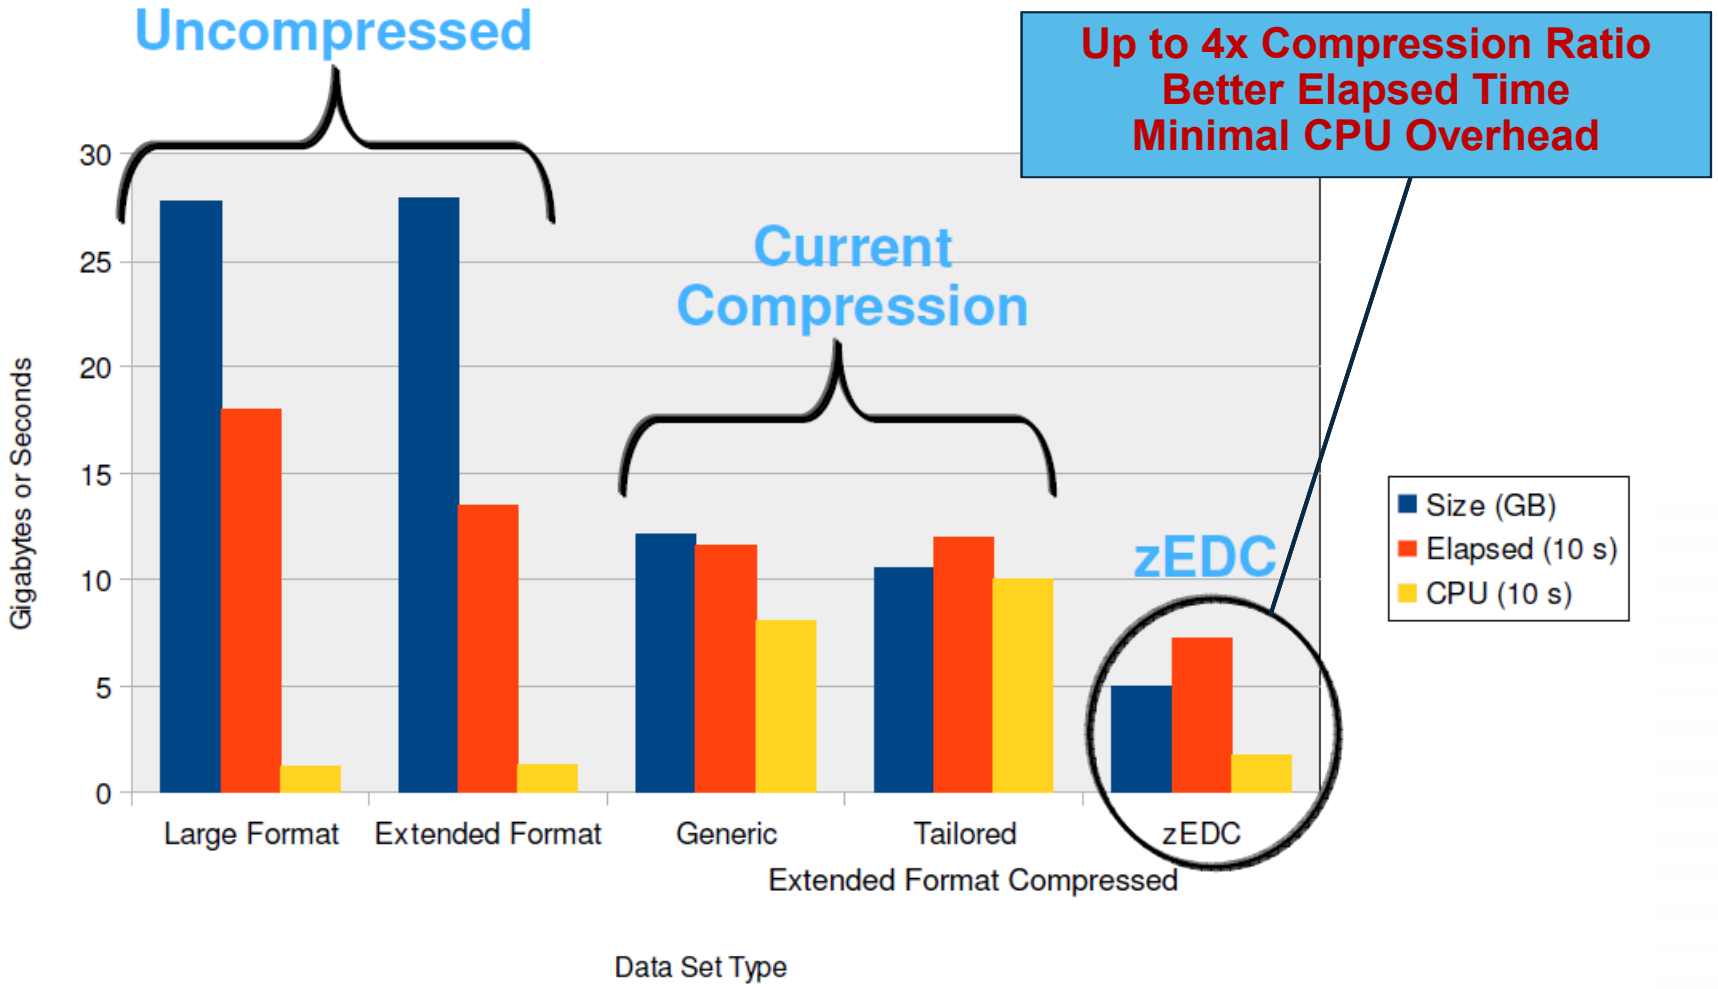

* **Disclaimer:** *Based on projections and/or measurements completed in a controlled environment. Results may vary by customer based on individual workload, configuration and software levels.*

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.

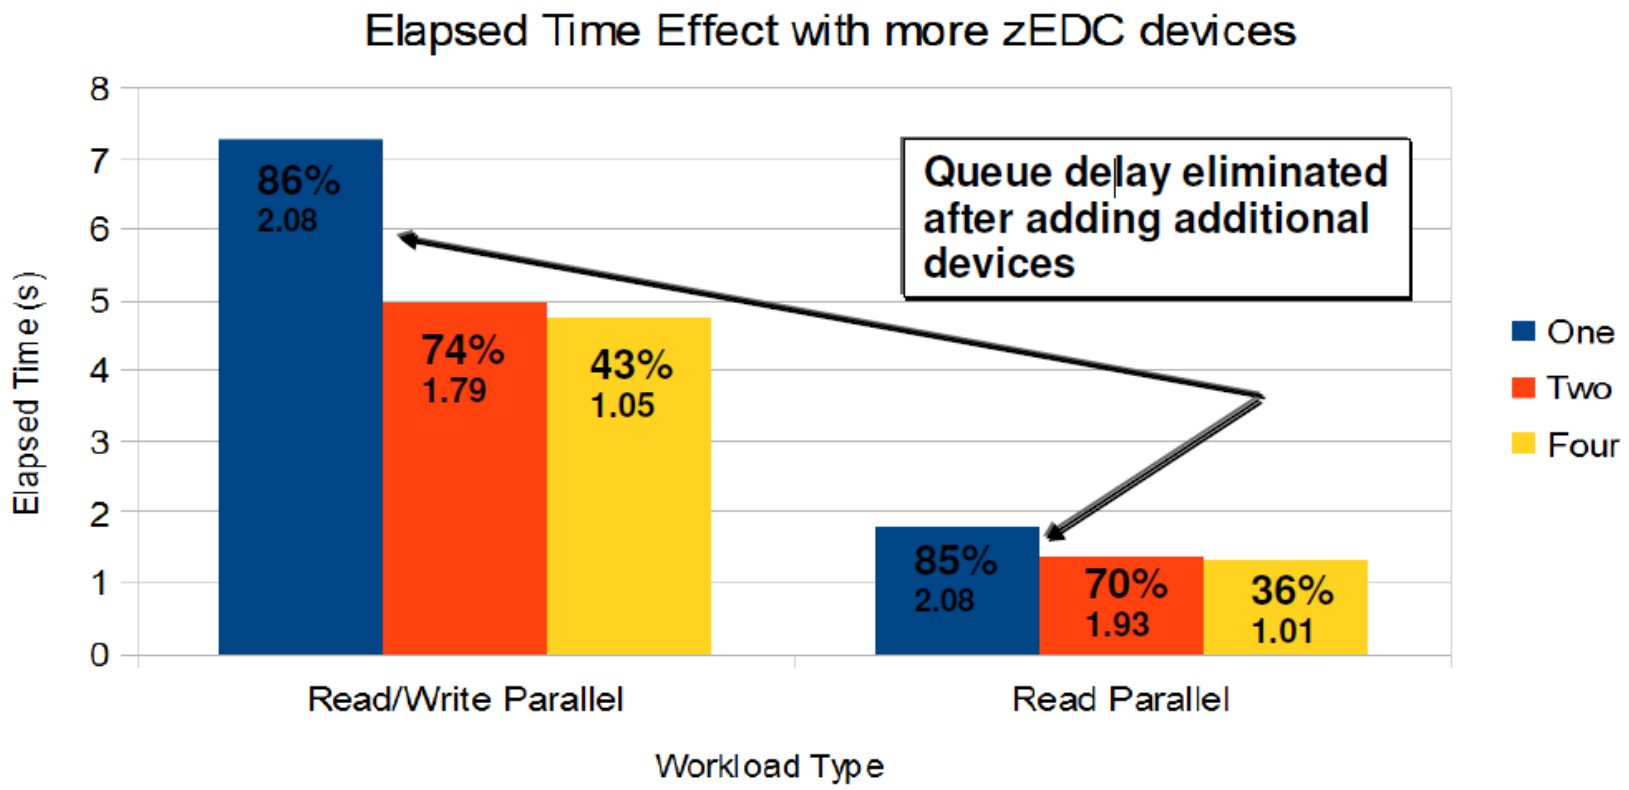

Configuring and Deploying zEDC

- **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)
    - Publication Updates available with z/OS 1Q14 Refresh
  - z/OS V1.13 and V1.12 offer software decompression support only
    - PTFs for BSAM/QSAM co-existence (APAR OA43863 / UA72750 (R12) UA72751 (R13)
  - zEDC Express for z/OS feature must be enabled in an IFAPRDxx PARMLIB member.
- **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.***

Compression Coprocessor (CMPSC) vs zEDC

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

Small object compression	Large Sequential Data	Industry Standard Data
▪ Rows in a database	▪ QSAM/BSAM Online Sequential Data ▪ Objects stored in a data base	▪ Cross Platform Data Exchange
Users	**Users**	**Users**
▪ VSAM for better disk utilization ▪ DB2 for lower memory usage ▪ The majority of customers are currently compressing their DB2 rows	▪ QSAM/BSAM for better disk utilization and batch elapsed time improvements ▪ SMF for increased availability and online storage reduction	▪ 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 Hands-on Lab, Thursday at 4:15PM (Session 15671)

- 0

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 IGDSMSxx 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,

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.

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 / Nonspanned  . . . . .         (S, N or blank)
  System Determined Blocksize . . 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.

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.***

DFA -
New fields 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

DFSMS allocation determines compressed format

- DFSMS allocation processing determines if a data set can be allocated as compressed format.
  - The type of compression to be used for the data set is not determined until the first OPEN for output of the data set.
- The following table summarizes the system behavior during SMS allocation processing for a new data set based on system levels and the user's allocation request.

z/OS Level	z/OS V2.1			
Processor Level	zEC12/zBC12		Pre-zEC12/zBC12	
Meets minimum compression space requirements[1]	Meet space requirements	Does not meet space requirements	Meet space requirements	Does not meet space requirements
Request: zEDC Required	•Allocation successful •Create as compressed format (V2)	•Allocation fails with IGD17168I message	•Allocation fails with IGD17168I message	•Allocation fails with IGD17168I message
Request: zEDC Preferred	•Allocation successful •Create as compressed format (V2)	•Allocation successful •Create as non-compressed extended format	•Allocation successful •Create as compressed format (V2)	•Allocation successful •Create as non-compressed extended format

1 - Minimum compression space requirements are 5MB, or 8MB Primary if no Secondary

SMS Scheduling System Selection for JES3 Environment

- In a JES3 environment, SMS identifies a list of one or more target systems where the job should be scheduled.
  - The table below identifies how SMS would select systems (based on priority) when the data set is to be allocated with zEDC compression (Required or Preferred)
- A system is capable of BSAM/QSAM zEDC Compression:
  - *zEC12 or zBC12*

Priority	System Description Based on zEDC Capability
1	**System capable of BSAM/QSAM zEDC compression.** **Devices are available on the system**
2	**System capable of BSAM/QSAM zEDC compression.** **No devices currently available on the system, however devices were available during this IPL**
3	**System capable of BSAM/QSAM zEDC compression.** **No devices currently available on the system, and none have been available during this IPL**
4	**System not capable of BSAM/QSAM zEDC compression.**

DFSMS OPEN Processing Determines Compression Type

- The first OPEN of the data set for output determines the compression type for the data set based on data class and PARMLIB specifications.
  - For a zEDC request, it must also determine level of the system. When running on a multi-system sysplex, it is possible for the data set to be allocated on one system but opened on a different system.
- The following table summarizes the system behavior during OPEN processing for a new data set based on system levels and the zEDC compression request.

z/OS Level	z/OS V2.1			
Processor Level	zEC12/zBC12		Pre-zEC12/zBC12	
DS1COMPR Flag set during allocation	DS1COMPR=on	DS1COMPR=off	DS1COMPR=on	DS1COMPR=off
Request: zEDC Required	•Create as zEDC compressed format (V2)	•N/A (Note 1)	•Create as zEDC compressed format (V2) (Note 2)	•N/A (Note 1)
Request: zEDC Preferred	•Create as zEDC compressed format (V2)	•Create as non-compressed extended format	•Create as Tailored compressed format (V2)	•Create as non-compressed extended format

Notes:

1. Since the allocation determined that the minimum compression space was not met, OPEN will not check the data class for a compression type.
2. Since the allocation was successful for this data set (DS1COMPR=on), the system will avoid failing OPEN but instead create the data set in zEDC compressed format. In this situation, all data will be written non-compressed.

DFSMS Allocation Determines Compressed Format (Coexistence)

- DFSMS allocation processing determines if a data set can be allocated as compressed format.
  - The type of compression to be used for the data set is not determined until the first OPEN for output of the data set.
- The following table summarizes the system behavior during SMS allocation processing for a new data set based on system levels and the user's allocation request.

z/OS Level	**z/OS V1R12 or V1R13**	
Meets minimum compression space requirements:5MB Primary if no Secondary	Meets space requirement	Does not meet space requirement
DataClass specifies **COMPACTION other than 'N'** **(See Note1)**	Allocation successful. Create as compressed format (DS1COMPR=on)	Allocation successful. Create as non-compressed extended format (DS1COMPR=off
DataClass specifies **COMPACTION = N or not specified**	Allocation successful. Create as non-compressed extended format (DS1COMPR=off)	Allocation successful. Create as non-compressed extended format (DS1COMPR=off)

Note1: On downlevel systems, SMS allocation does not differentiate between the different COMPACTION options.

DFSMS Open Processing Determines Type of Compression (Coexistence)

- The first OPEN of the data set for output determines the compression type for the data set based on data class and PARMLIB specifications.
  - For a zEDC request, it must also determine level of the system. When running on a multi-system sysplex, it is possible for the data set to be allocated on one system but opened on a different system.
- The following table summarizes the system behavior during OPEN processing for a new data set based on system levels and the zEDC compression request.

z/OS Level	**z/OS V1R12 or V1R13**	
	DS1COMPR=on	DS1COMPR=off
zEDC Request (see Note 3)		
zEDC Required	Create as zEDC compressed format (EF V2) **(See Note2)**	n/a **(See Note1)**
zEDC Preferred	Create as tailored compressed format	n/a
N/A	N/A	Create as non-compressed extended format

- Note1: Since the allocation set DS1COMPR=off, OPEN will not check the data class for a compression type.
- Note2: Since the allocation was successful for this data set, the system will avoid failing OPEN but instead create the data set in zEDC compressed format (which causes the data set to be V2). In this situation, all data will be written non-compressed on a V1R12 or V1R13 system. However, if written from a V2.1 system capable of zEDC compression, the data will be written compressed. (Also, for this case, the EF V2 specification is set by the system during the first OPEN for output instead of SMS allocation time and only for zEDC compressed format. Note that the restriction that a user cannot request a to create a new EF V2 data set on downlevel systems continues to exist).
- Note3: Since new values for zEDC compression are not available for the COMPRESS parameter found in IGDSMSxx member of SYS1.PARMLIB on downlevel systems, only DataClass will be used to determine if zEDC compressed format is requested.

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****

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)

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)

Compatible with Existing Compression Types

- **zEDC Compressed format data sets are essentially compatible with other compressed format data sets**
  - **DFSMSdfp**
    - Supports BSAM and QSAM access for zEDC compressed format data sets
      - zEDC compressed format data sets will be defined as extended format V2 (EF V2) data sets, regardless of the user's specification in DataClass, JCL or SYS1.PARMLIB.
        - » Note: FEOV is not supported on output for V2 data sets. When FEOV is issued on output to a V2 data set, abend 737-48 is issued.
  - **DFSMSdss**
    - Supports DUMP, RESTORE, COPY, DEFRAG, CONSOLIDATE, and PRINT of zEDC compressed format data sets
      - REBLOCK keyword is not supported on COPY for zEDC compressed format data sets
      - Note: As today, when copying or restoring compressed format data sets the type of compression used is carried along from the source. This is true whether the preallocated target was usable, or had to be scratched and reallocated. Also, DFSMSdss does not support copying a compressed format data set to a non-compressed format data set or vice versa.
  - **DFSMShsm**
    - Supports Migrate/Recall of zEDC compressed format data sets (DSS is the DataMover)

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

DFSMSdss Coexistence

- **Coexistence PTFs for DFSMSdss (COPY/DUMP/RESTORE)**
  - DFSMSdss will allow a user on V1R12 and V1R13 to RESTORE a compressed format sequential data set when the form of compression used was zEDC compression.
    - Information indicating the data set is a compressed format data set will be preserved during the RESTORE.
  - DFSMSdss will fail logical data set COPY and DUMP operations of extended format data sets in the zEDC compressed format.
    - A new reason code will be added to the existing ADR778E indicating a compressed format data set compressed with the zEDC form of compression is not supported on this release.
  - DFSMSdss will fail logical data set COPY and RESTORE operations when a pre-allocated output data set is compressed format with the zEDC form of compression.
    - A new reason code will be added to the existing ADR285E indicating a pre-allocated compressed format data set compressed with the zEDC form of compression is not supported on this release.

Thank You!

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

VSAM Compression

- Introduced in DFSMS/MVS 1.2 (KSDS only)
- Only KSDS can be compressed
  - Data component is compressed
    - Index is not eligible for compression
  - Every record has a compression prefix
  - Spanned/non-spanned KSDS compression
    - Same as current rules for CI occupancy
  - Offset plus length of prime key not compressed
  - Freespace maintained as today
    - Same specification will provide room for more records in compressed format
  - Updates allowed
  - Uses Dictionary Building Block (DBB) based dictionaries
    - Tailored compression only for SAM data sets

Dictionary Building Block (DBB) Compression

- Host based
- Ziv-Lempel Technique
- Dictionary Building Blocks
- 8K-64K sample
- Dictionary token in catalog

Why Use VSAM Compression?

- Makes more user data available to the application per an I/O request
  - Reduced number of I/O requests reducing channel utilization
- Reduced disk space usage
- Improved controller cache usage
- More efficient use of processor storage
- Relief from 4GB limit

Allowing VSAM Compression

- Primary allocation of at least 5 MB or 8 MB if no secondary allocation is specified
- Maximum record length specified must be at least key offset plus key length plus 40 bytes
- Must be SMS managed extended format data set

Requesting VSAM Compression

```
SCDS Name . . . : MY.TEST.SCDS
Data Class Name : DCCOMP
To DEFINE Data Class, Specify:
  Retpd or Expdt  . . . . . .          (0 to 9999, YYYY/MM/DD or blank)
  Volume Count  . . . . . . . 1        (1 to 59 or blank)
    Add'l Volume Amount . . .          (P=Primary, S=Secondary or blank)
  Imbed . . . . . . . . . . .          (Y, N or blank)
  Replicate . . . . . . . . .          (Y, N or blank)
  CIsize Data . . . . . . . .          (1 to 32768 or blank)
  % Freespace CI  . . . . . .          (0 to 100 or blank)
              CA  . . . . . .          (0 to 100 or blank)
  Shareoptions Xregion  . . .          (1 to 4 or blank)
               Xsystem  . . .          (3, 4 or blank)
  Compaction  . . . . . . . .          (Y, N, T, G or blank)
  Media Interchange
    Media Type  . . . . . . .          (1, 2, 3, 4 or blank)
    Recording Technology  . .          (18, 36, 128, 256 or blank)
```

Is It Compressed?

- **Catalog**
  - Uncompressed size
  - Compressed size
  - Compression flag
  - Compression dictionary token
  - Compression characteristic data
- **VTOC**
  - Compression indicator (DS1COMPR)
- **SMF**
  - Record type 64
- **Also....**
  - ISMF
  - DCOLLECT
  - LISTCAT
  - IEHLIST

Trademarks

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

AIX*	DFSMSdfp	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 Processors/Central 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.