Compression on z Systems

Anthony Sofia (atsofia@us.ibm.com)
Senior Technical Staff Member – z Systems Design and Development

04/26/2017

Explosive growth in data

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 the cost prohibitive nature of keeping data online

Data needs to be shared across different 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

Storage technology, including Flash, more affordable with compression

z Systems Compression Technology Overview

Using the right hardware compression acceleration for each of your workloads

Compression Coprocessor	z Enterprise Data Compression
On Chip – **Improved on z13!** 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

Benefits of Faster CPU Compression with IBM z13

The z13's compression engine reduces CPU needed for compression with CMPSC **up to 50%** compared to the zEC12. This can provide significant CPU performance and response time advantages for applications and middleware using compression.

Our measurements have shown that I/O intensive batch jobs using QSAM/BSAM Generic and Tailored compression can take up to 49% less CPU and up to 40% less elapsed time on z13 than on zEC12.

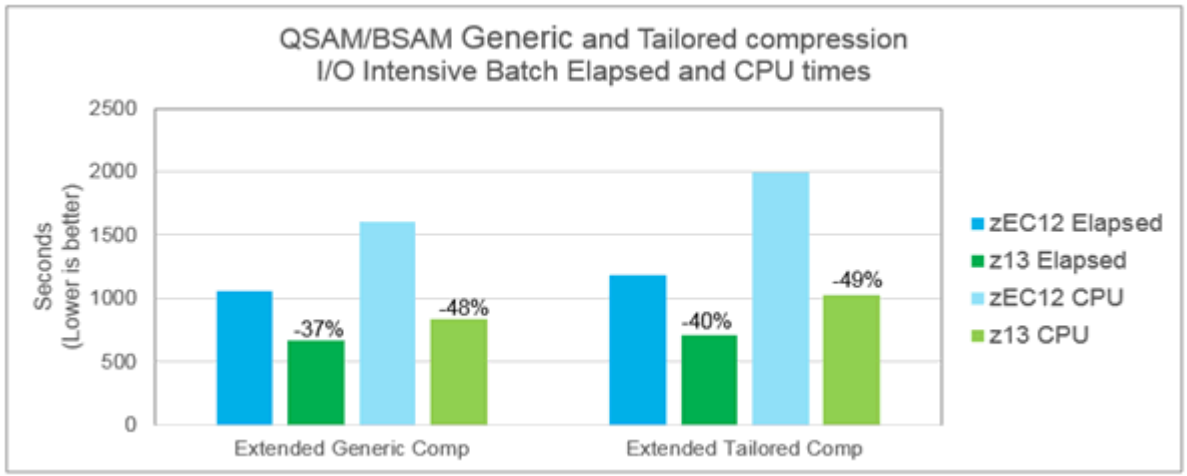

Disclaimer: Measurement results were collected in a controlled measurement environment. Your results may vary.

Benefits of Faster CPU Compression with IBM z13 (cont)

The advantages of compressing DB2 tables are well understood. DB2 compression can significantly reduce the DASD space needed to store data. DB2 compression can also result in improved buffer pool hit ratios, reducing IO operations and helping response time and CPU consumption. DB2 gets significant benefit from improved CPU compression performance on the z13.

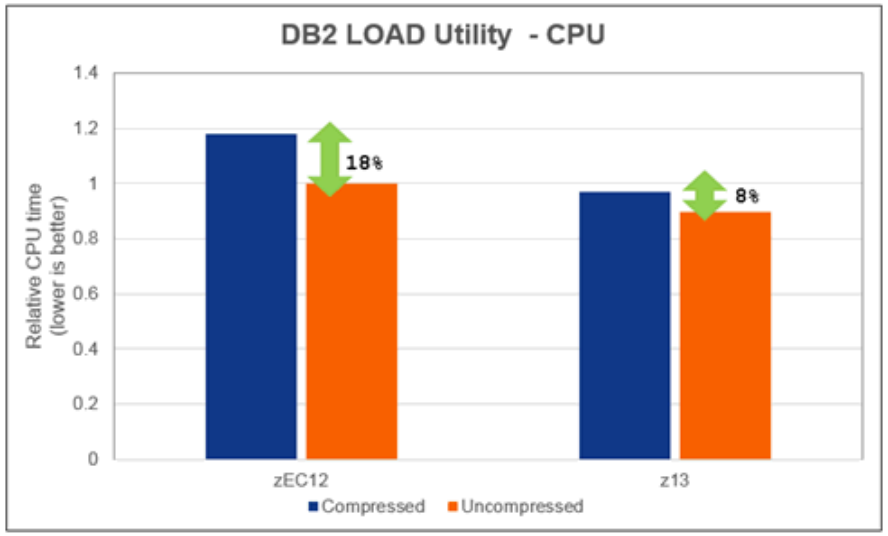

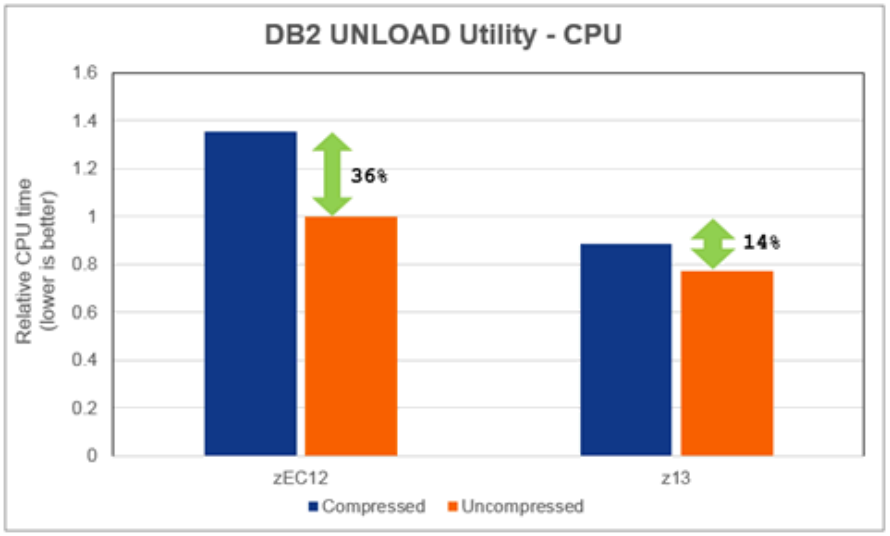

Loading data to compressed tables, the CPU cost of compression is reduced by up to 59% on z13 as compared to zEC12. Given this improvement, the CPU cost of loading to compressed tables vs. uncompressed tables can be as low as 8% on z13.

Unloading data from compressed tables, the CPU overhead of compression is reduced by up to 69% on z13 as compared to zEC12. Given this improvement, the CPU cost of unloading from compressed tables vs. uncompressed tables can be as low as 14% on z13.

Disclaimer: Measurement results were collected in a controlled measurement environment. Your results may vary.

IBM zEnterprise Data Compression (zEDC)

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 vs typical 50 MB a second in software*

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 IBM 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** 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** deliver expanded customer value
- **Java for z/OS V7R1** accelerates common compression classes used by applications and middleware
- **Improved availability** with SMF

Deploying zEDC

- **Operating system requirements**
  - Requires z/OS 2.1 (w/PTFs) and the new zEDC Express for z/OS feature
  - z/OS V1.13 and V1.12 offer software decompression support only
- **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
    - This provides up to 4 GB/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 zEDC Express in an existing zEC12/zBC12
    - Apply z/OS Service; Hot plug a zEDC Express adapter; IODF updates and Dynamic Activate
- **Capacity Planning**
  - The **IBM System z Batch Network Analyzer** tool reports on potential zEDC usage for QSAM/BSAM data sets

    http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/PRS5132

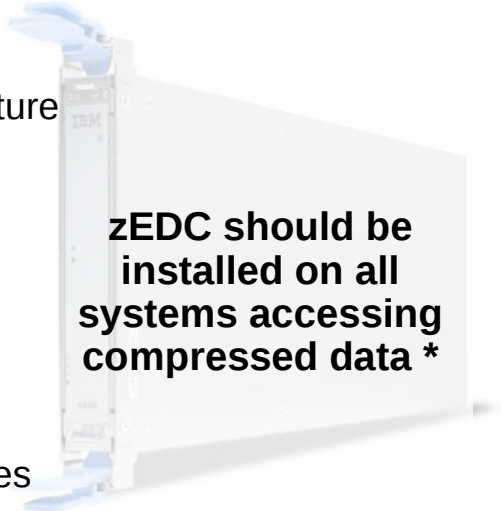

zEDC should be installed on all systems accessing compressed data *

* 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.

zEDC and Resource Groups

- During Resource Group level microcode updates half of the zEDC adapters are brought offline during the update.
- To avoid a temporary single point of failure, four zEDC devices can be deployed across two resource groups

Two Card Configuration

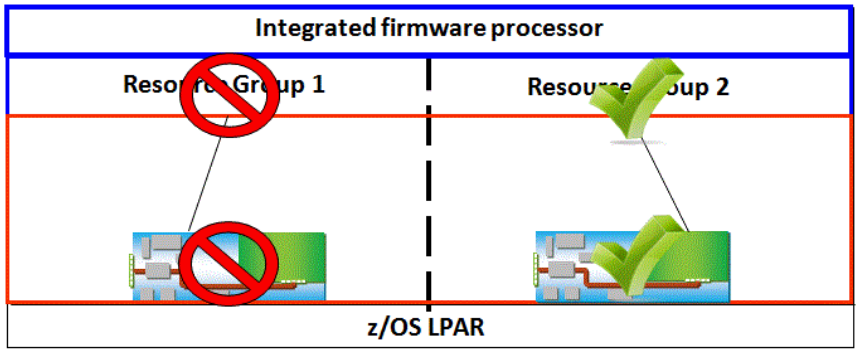

Four Card Configuration

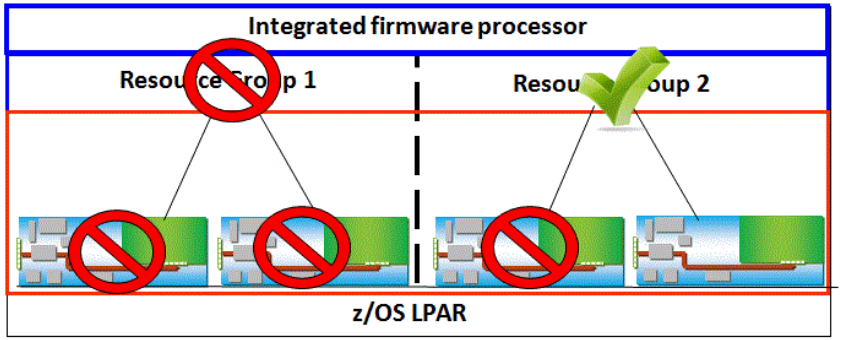

zEDC Product Usage Overview

Roadmap of Growing Value and Commitment to Compression

- Value buy-in from IBM products across the product stack
- z Systems ISVs have been very active with zEDC

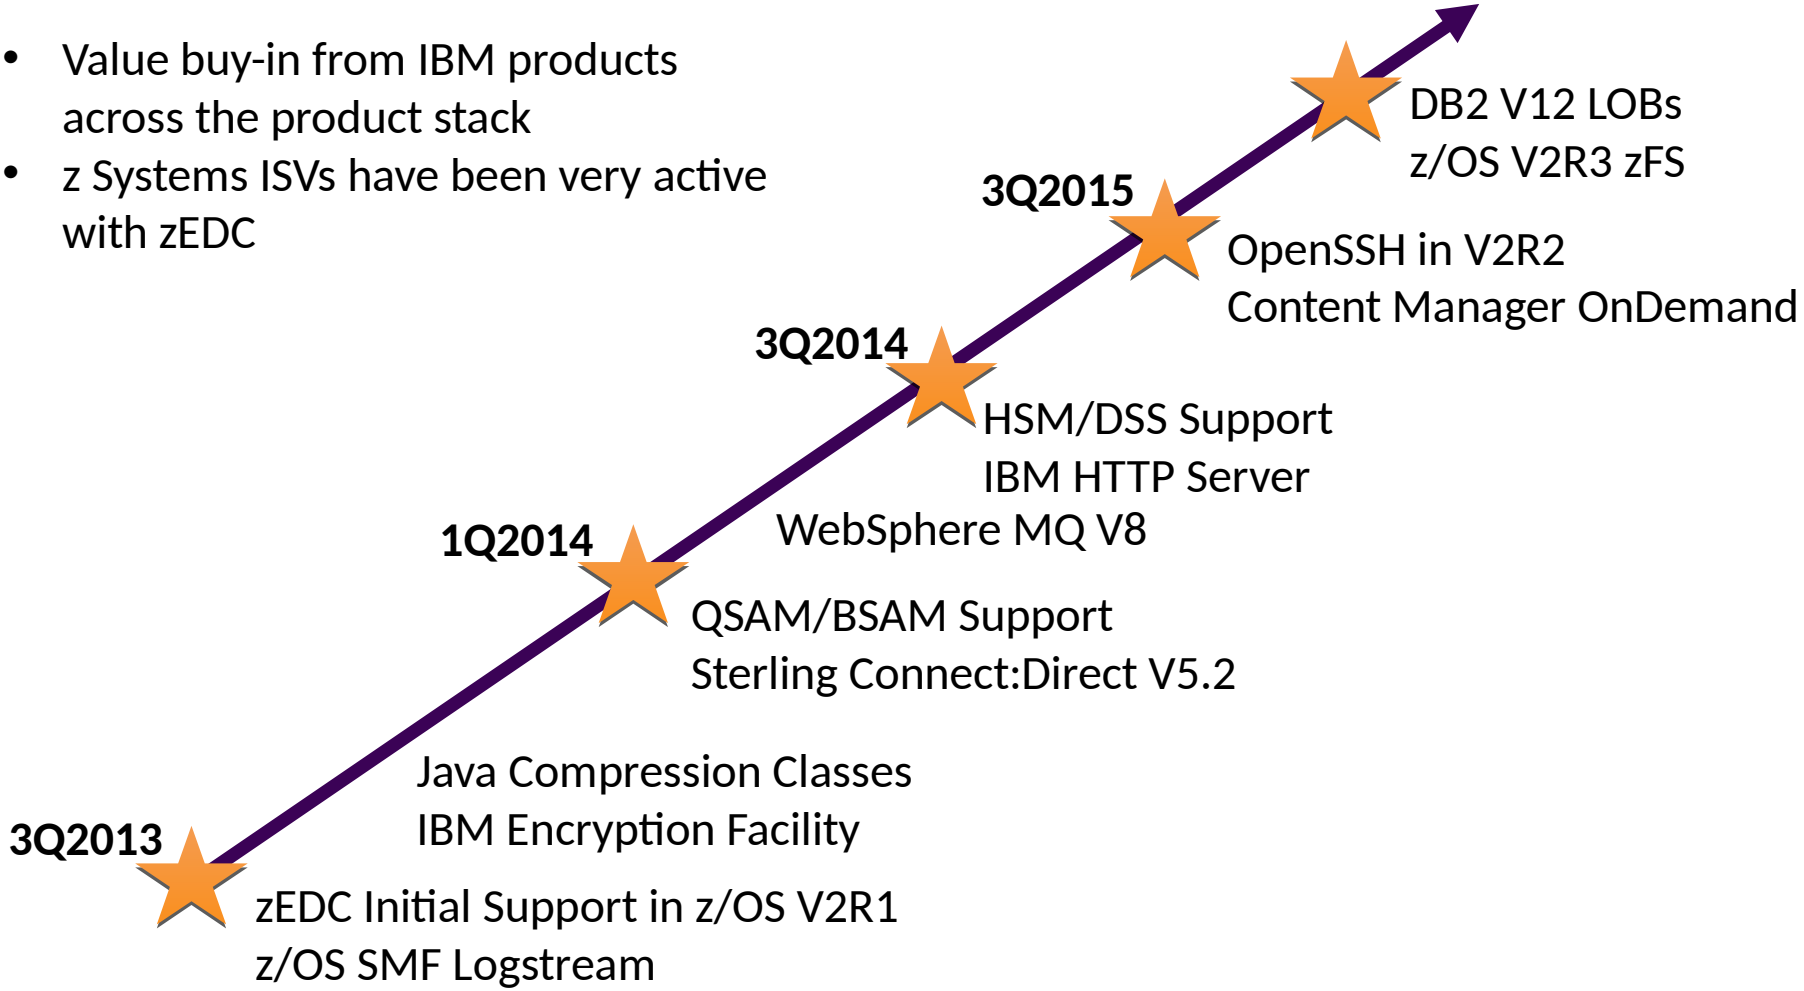

Sequential Data Set Compression with BSAM/QSAM and zEDC

Reduce the cost of keeping your sequential data online

- zEDC compresses data up to 4X, saving up to 75% of your sequential data disk space
- 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

Simple Enablement

- Use a policy to enable the zEDC

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

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.

Data Replication

- Replication technologies which move data in physical format can take advantage of the reduced storage requirements of data compressed with zEDC.
  - Significant amounts of zEDC compressed data can reduce the amount of data transferred as well as the elapsed time to complete the transfer.

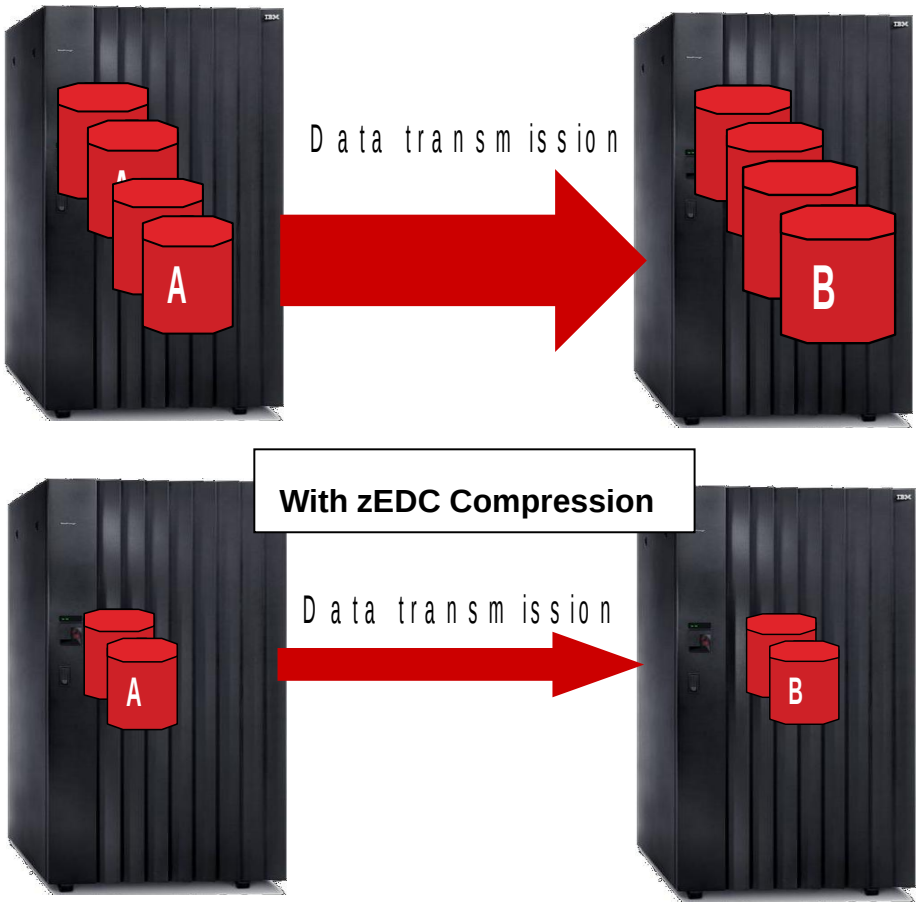

QSAM/BSAM zEDC

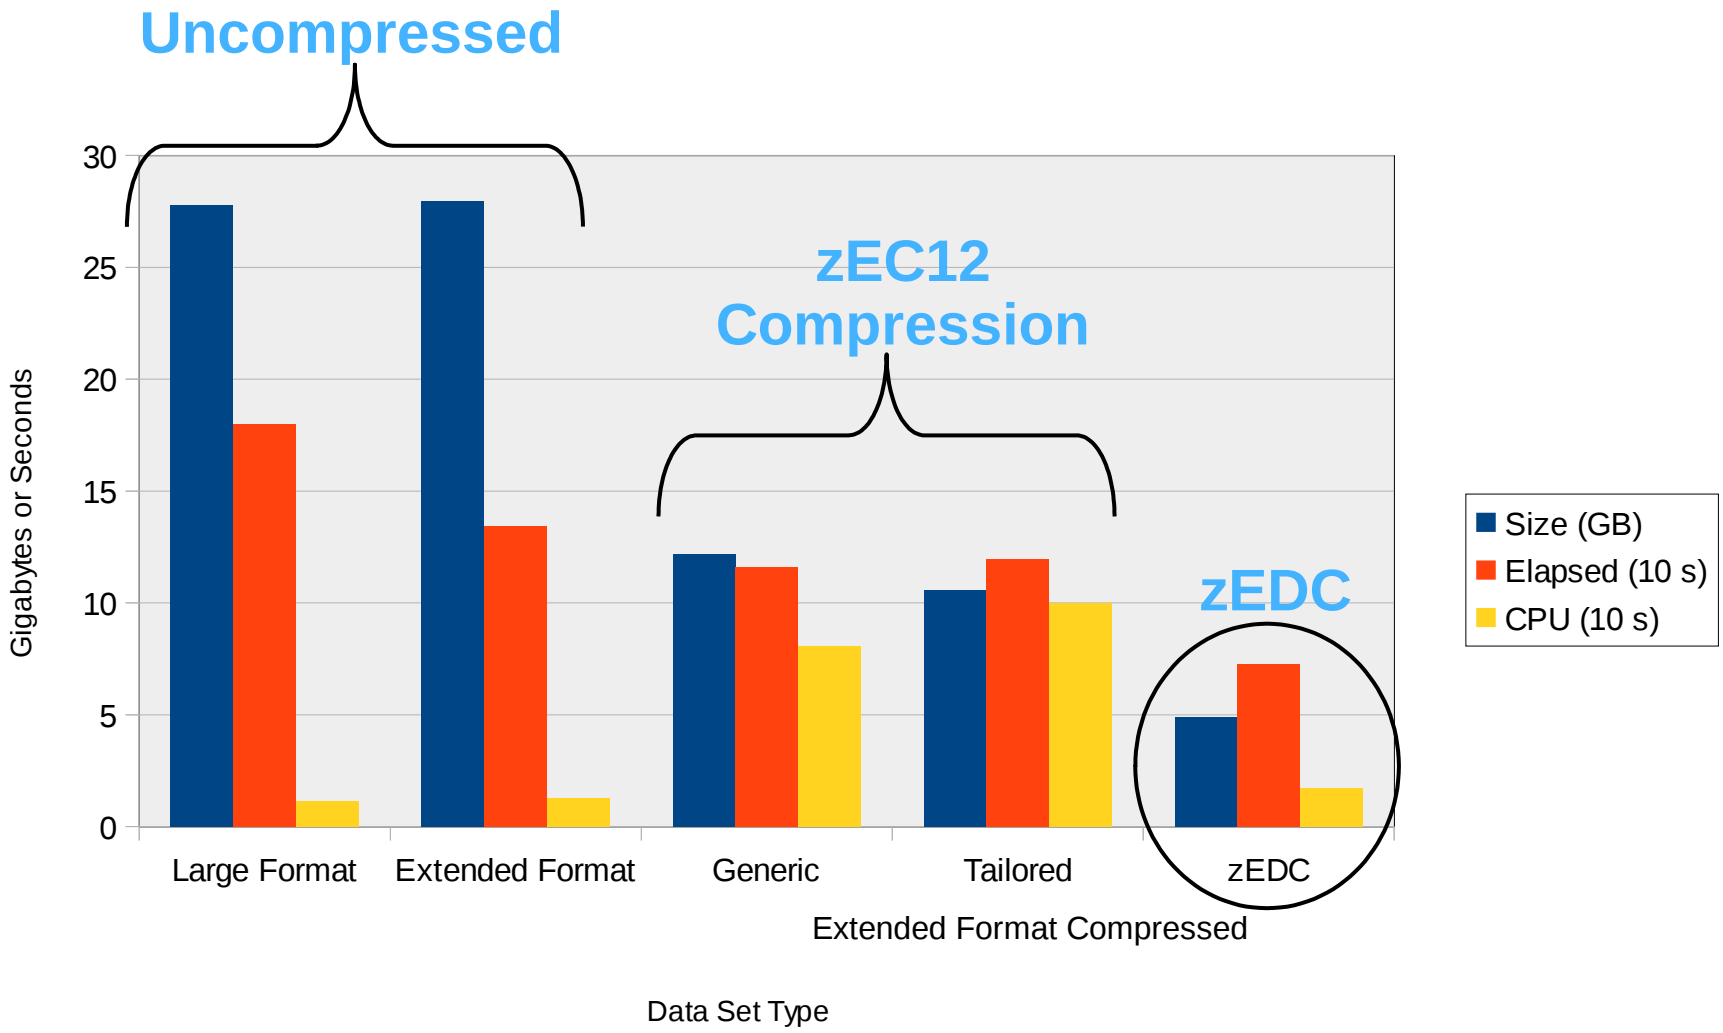

Disclaimer: Based on projections and/or measurements completed in a controlled environment. Results may vary by customer based on individual workload, configuration and software levels.

zEDC with IBM Sterling Connect:Direct for z/OS 5.2

High performance data compression for Managed File Transfer

- Sterling Connect:Direct can automatically leverage zEDC Express Accelerator for file compression and decompression as files are transferred
- Works with various dataset and file types on z/OS
- File transfers can be z/OS to z/OS, or z/OS to Distributed (UNIX®, Linux®, Linux on System z, and Windows)
- Fully compatible with zlib compression used in IBM Sterling Connect: Direct today – no changes required at end points

Up to **80%** reduction in elapsed time to transfer a file from z/OS to z/OS (results vary by dataset type and characteristics of the data) with minimal CPU increase

Significant improvement in CPU time with zEDC over zlib software compression (sender TCB time and receiver TCB time)

Users currently using compression will see a comparable compression ratio

Disclaimer: Results based on internal controlled measurements.
Results may vary by customer based on individual workload, data, configuration and software levels.

Benefits using DFSMShsm Migration/Recall with zEDC

Reduced VTS and Tape Usage

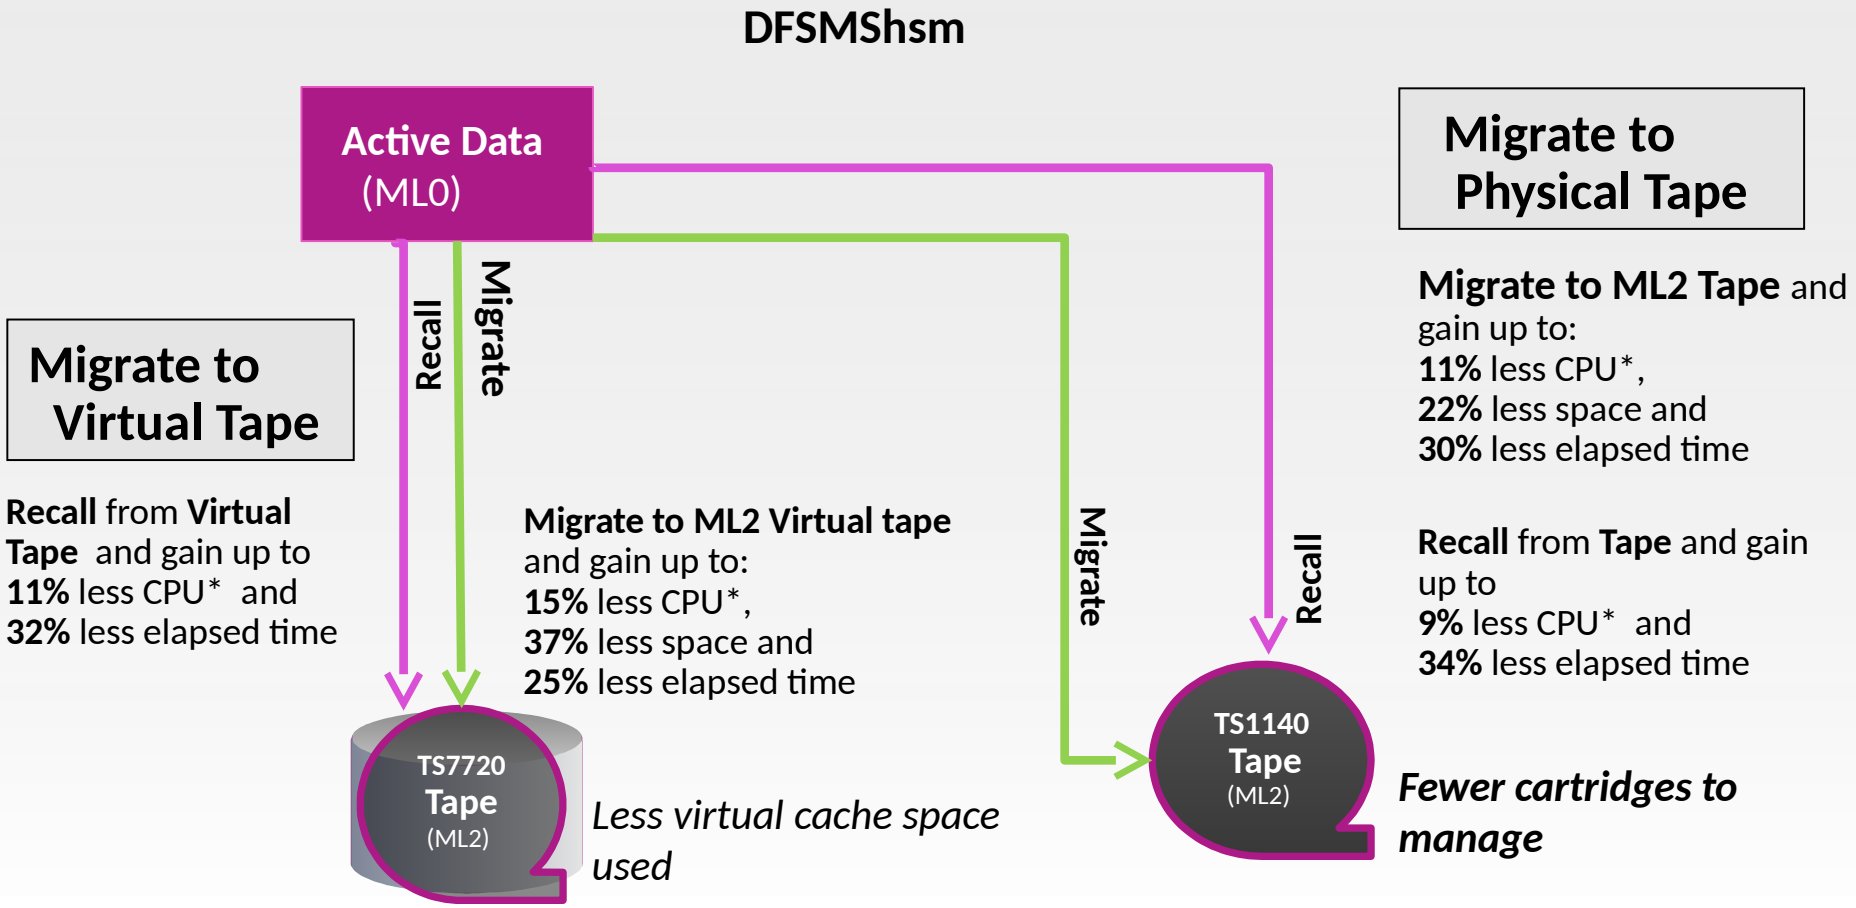

*as compared to using DFSMShsm with no host compaction options on z13

Results measured in an internal lab environment. Your results will vary.

Benefits using DFSMShsm with zEDC for ML1 Migrations

Compress data with less CPU

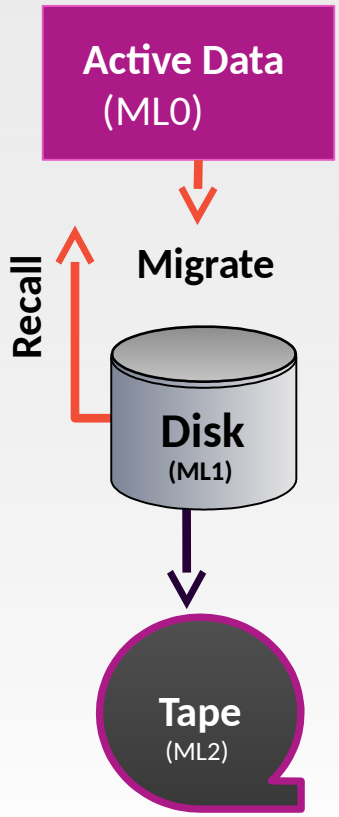

Migrate and use: **Up to 58% less space Up to 80%** less CPU*

Recall from disk with **69%** reduction in CPU*

Side benefit: ***Shorter elapsed time***

*as compared to using DFSMShsm with the COMPACT option on zEC12.

When migrating data to ML1 Disk using zEDC, use up to 58% less disk space and use up to 80% less CPU compared to using DFSMShsm with the COMPACT option. 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 Customers using DFSMSdss with zEDC performing full volume dumps to tape can reduce the number of tape cartridges by up to 19%

Results measured in an internal lab environment. Your results will vary.

zEDC with IBM Content Manager OnDemand V9.5
High data compression of statements and reports

Content Manager OnDemand can automatically leverage the zEDC Express Accelerator when storing data into the archive as well as when data is requested for retrieval

Fully compatible with zlib compression used in IBM Content Manager OnDemand today – no changes required at end points

MIPS reduction of 75% when compared to existing software based compression (results vary by document size and type)

Users currently using compression will see a comparable compression ratio

OnDemand

- **Long term archive for static content**
- **Designed for computer generated output (statements, policies, etc.)**
- **Large data-streams with potentially millions of documents daily**
- **High-performance batch loading**
- **Efficient storage and retrieval**
- **Print data-stream segmentation and indexing**

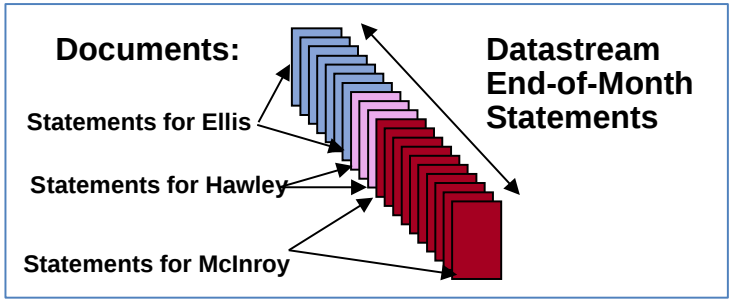

Disclaimer: Results based on internal controlled measurements using IBM Content Manager OnDemand and various document types. Results may vary by customer based on individual workload, data, configuration and software levels.

zEDC with z/OS SMF Logger

Alleviate SMF constraints across the entire life cycle of a record using compression technology

Record Creation

Records created by z/OS, DB2, CICS® and applications are all written to SMF

Record Buffering

zEDC Express can be used to compress SMF data resident in the z/OS System Logger cutting down on logger storage requirements.

Record Extraction

SMF can extract compressed data from logger faster than uncompressed data.

Targeting a compressed QSAM/BSAM data set for long-term archival can also optimize disk space. zEDC, tailored or generic compression can be used depending on the requirements

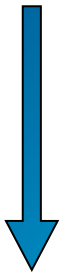

Record Archival

Reading SMF data from a compressed data sets can increase the performance of applications that access that data.

- Store up to **4x** less data in System Logger
- Logger CPU usage reduced by up to **30%**
- Up to **15%** reduction in elapsed time for SMF extraction from Logger
- SMF data stored in zEDC compress BSAM can save up to **4x** in archived SMF data size
- Programs reading SMF data from a zEDC compressed data set can see an elapsed time reduction

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)
Disclaimer: These results are based on projections and measurements completed in a controlled environment. Results may vary by customer based on individual workload, configuration and software levels (Logger)

zEDC with WebSphere MQ for z/OS V8

WebSphere MQSeries® has always provided compression options for message data passed over MQ channels via the COMPMSG attribute.

The existing zlib options are the following:

- ZLIBFAST - Message data compression is performed using the zlib compression technique. A fast compression time is preferred.
- ZLIBHIGH - Message data compression is performed using the zlib compression technique. A high level of compression is preferred.

Starting with WebSphere MQ for z/OS V8 the COMPMSG(ZLIBFAST) attribute will now use zEDC when available to perform compression and decompression of message data.

This support is ideal for channels that handle large, 32KB requests.

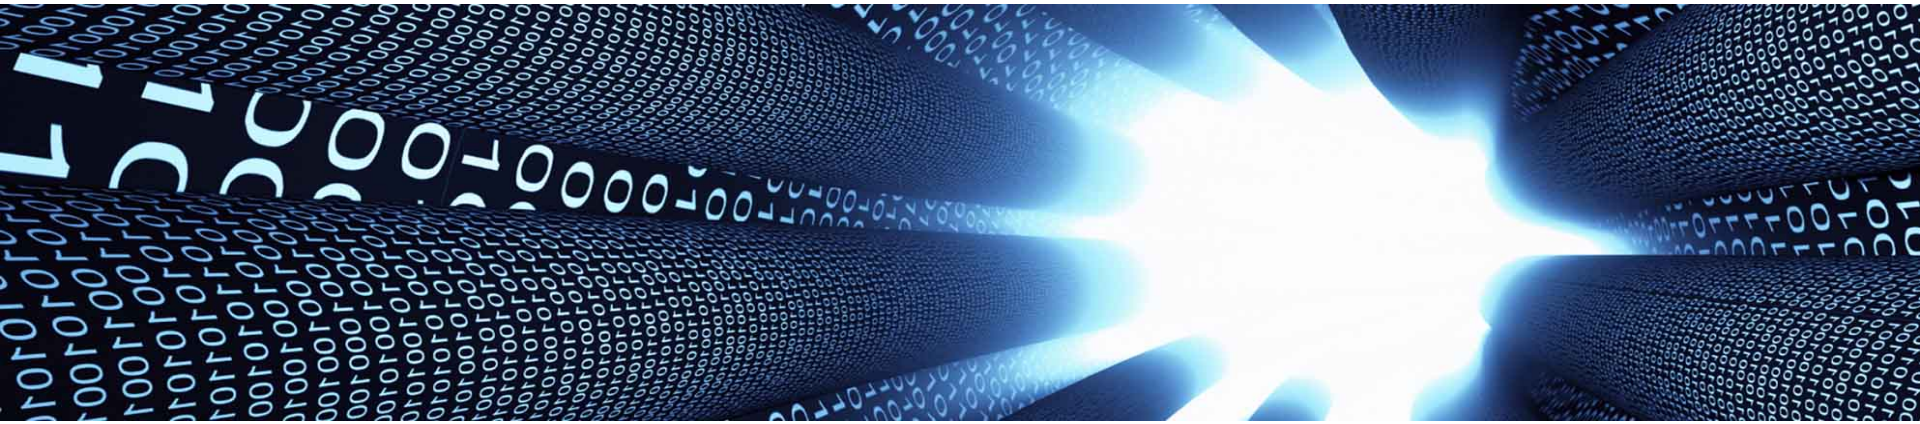

zEDC with Java

Transparent enablement of the *java.util.zip* package enables high throughput compression and decompression for

- Application Business Data Exchange
- HTTP Responses for Web Services
  - Servlet Filters
  - WebSphere® Web Services component
- Large objects that are serialized and stored
- Exploited through standard Java APIs java.util.zip* in the latest releases of Java 7.0.0, and Java V7R1
  - zEDC java.util.zip.Deflater in memory test improved elapsed time up to 55x and CPU time up to 240x when compared to zlib software compression.
  - Java application to compress files using java.util.zip.GZIPOutputStream class
    - Up to 90% reduction in CPU time using zEDC hardware versus zlib software shown in graphs to the right
    - Up to 74% reduction in Elapsed time using zEDC hardware versus zlib software

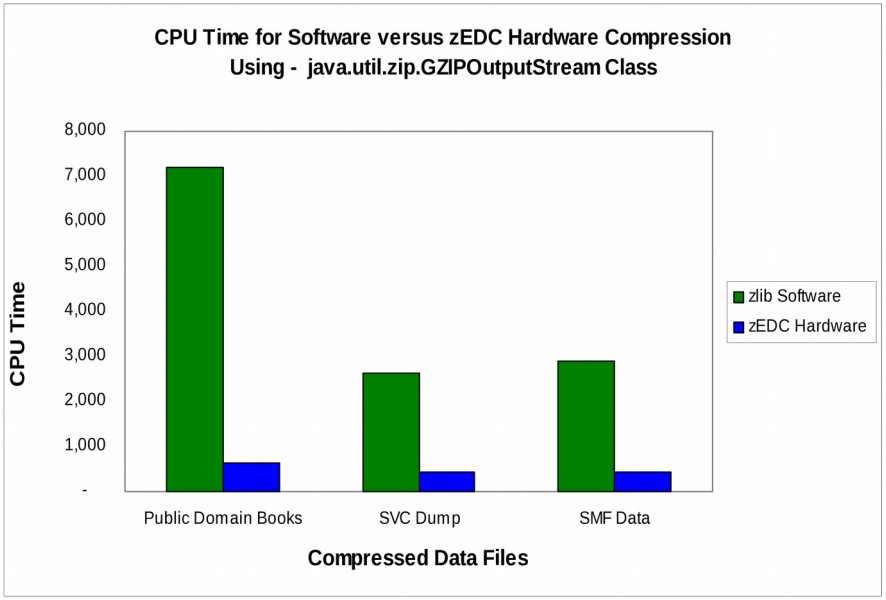

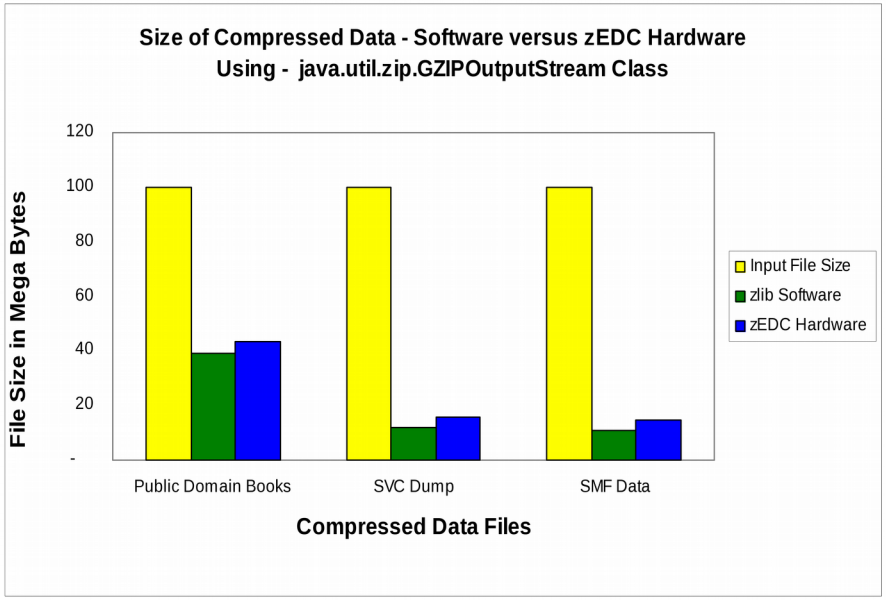

* Not all java.util.zip classes exploit zEDC

Disclaimer: Results are based on internal controlled measurements using java.util.zip.Deflater on data already in memory. Results may vary based on the application's use of java.util.zip classes and other work done by the application

zEDC with Java Application – Encryption Facility for z/OS

Increased throughput and functionality for standard compliant business partner data exchange. The Encryption Facility (EF) for z/OS can now use zEDC to compress and decompress data

1. Before performing an encryption and after performing a decryption
2. As a stand-alone operation NEW

Compressing data with zEDC enables higher throughput than no compression or software compression

zEDC can provide IBM EF users reductions of up to 60% in elapsed time and up to 70% in CPU time for environments where compression is already in use

For IBM EF users not already using compression, compression with zEDC can provide IBM EF users a reduction of up to 44% in elapsed time and up to 46% in CPU times

Results based on files containing public domain books. Results may vary by customer based on individual workload, data, configuration, and software levels.

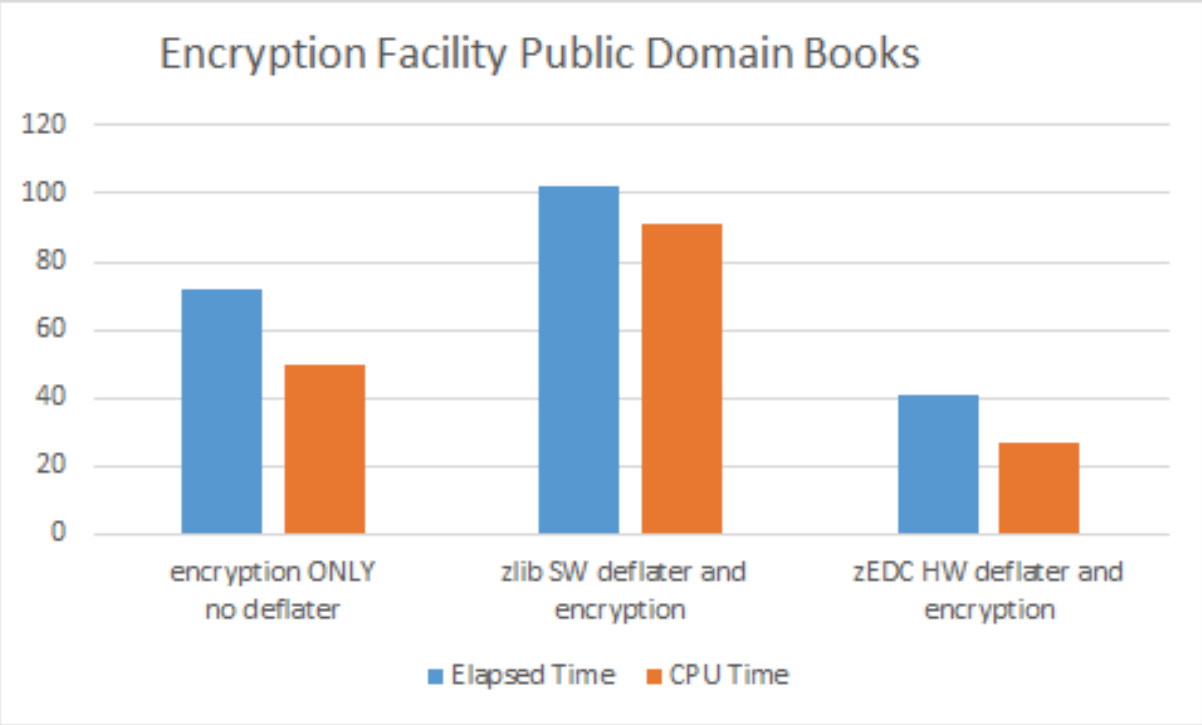

Disclaimer: Results based on internal controlled measurements using IBM Encryption Facility for files containing public domain books. Results may vary by customer based on individual workload, data, configuration and software levels.

Garanti Bank

Headquartered in Istanbul, T. Garanti Bankası A.Ş. is an integrated financial-services group with total assets of $97,6B, providing more than 13,7 million customers with banking products and services*. Garanti runs 2 zEC12 servers with DB2, CICS, MQ and GDPS supporting over 556M transactions and 40K batch jobs per day.

Garanti recently started using the zEnterprise Data Compression (zEDC) feature

Benefits of using zEDC according to our performance measurements

- **Significant improvement in Batch helps Online workloads**
  - Between **17 to 86% CPU decrease** in Batch Jobs CPU time – Reduced CPU can also helps lower software costs
  - On average **Elapsed Time of Batch cut in half** providing the ability to drive **more Batch workload** in less time
- **Huge disk savings compared to using DFSMS compression**
  - Between **18 to 80% data space reduction** – RMF reports indicates a 10:1 Compression Ratio on sequential data (BSAM, QSAM)
  - zEDC compression resulted in lower network traffic and improved **XRC bandwidth utilization** for failover

« If you are using QSAM and BSAM datasets in batch jobs, and their sizes are greater than 5MB, zEDC is the best and simplest way to reduce CPU usage »

***Serhat Balli**, z/OS Principal specialist*

*As of September 30, 2015

SHARE Charts from Winter 2016

- Presented by IBM and Humana
- Pulled a subset of charts showing the larger Humana system
- For full charts see:
  http://s23.a2zinc.net/clients/SHARE/Winter2016/Public/SessionDetails.aspx?FromPage=Sessions.aspx&SessionID=400&nav=true&Role=U%27

zEDC Usage – SMF Logger Only – LDC1

August 4, 2015 – zEC12

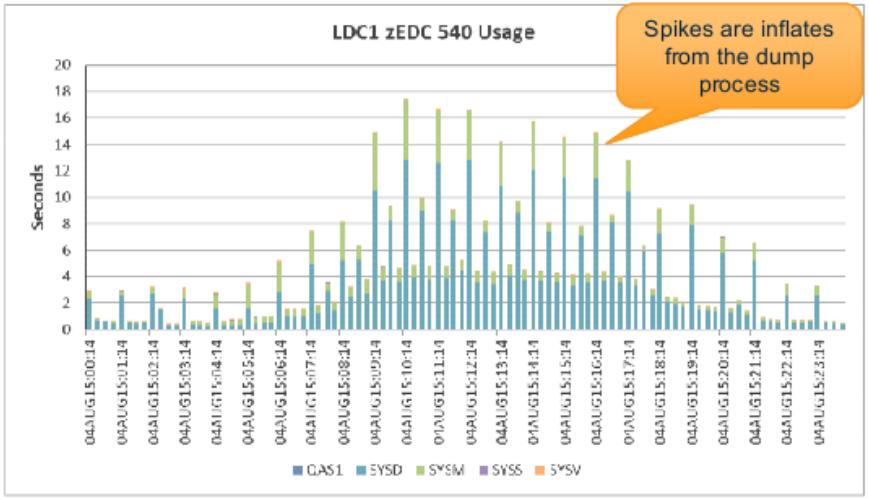

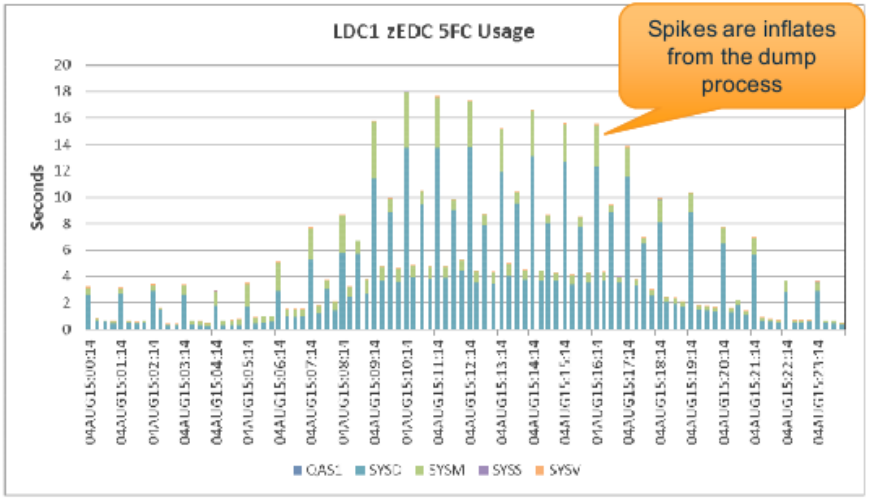

- zEDC card usage in seconds from type 74 subtype 9 records
- 18 out of 900 seconds or 2% busy

Total SMF Deflation

August 4, 2015 – zEC12

	LDC1	LDC2	Total
Raw SMF in Terabytes	0.74	0.38	1.12
Deflated SMF in Terabytes	0.09	0.04	0.13
Deflation Rate			88.28%
Deflation Ratio			8.53 to 1

- LDC1 has twice the SMF of LDC2
- One of the LDC1 LPARs has a heavy DB2 DDF load, producing 101 records

IBM System z Batch Network Analyzer and Compression

Helping determine if you have files that are candidates for zEDC

- **IBM System z Batch Network Analyzer**
  - A free, “as is” tool to analyze batch windows
  - Available to Customers, Business Partners and IBMers
  - PC based, and provides graphical and text reports
    - Including Gantt charts and support for Alternate Processors
- **Available Now on TechDocs**

  https://www-03.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 estimate the number of features needed
  - Generate a list of data sets by job which already do hardware compression and may be candidates for zEDC
  - Generate a list of data sets by job which may be zEDC candidates but are not in extended format

zEDC RMF Reporting

New RMF™ report shows the utilization of each device.

RMF Postprocessor Interval Report : PCIE Activity Report

RMF Version : z/OS V2R1 SMF Data : z/OS V2R1
Start : 02/24/2014-05.48.00 End : 02/24/2014-05.48.44 Interval : 00:45:000 minutes

Hardware Accelerator Activity

Function ID	Time Busy %	Request Execution Time	Std Dev for Request Execution Time	Request Queue Time	Std Dev for Request Queue Time	Request Size	Transfer Rate T
0013	0.689	7.78	0.417	15.0	0.953	24.3	21.5

Hardware Accelerator Compression Activity

Function ID	Compression Request Rate	Compression Throughput	Compression Ratio	Decompression Request Rate	Decompression Throughput	Decompression Ratio	B
0013	885	11.6	1.14	0	0		64

The percent of this interval where this specific zEDC Express device was executing requests

Compression ratio of all requests serviced by zEDC. This will span all usage of the zEDC Express devices

Average request queue time in Microseconds for this device.

ZSP03927-USEN-00

zEnterprise Data Compression Redbook

- **Provides overview of the technology**
- **Covers configuration of SMF, QSAM/BSAM and DFHSM/DFDSS**
- **Examples using zBNA**
- **Available Now on www.redbooks.ibm.com:**

  http://www.redbooks.ibm.com/Redbooks.nsf/RedpieceAbstracts/sg248259.html?Open

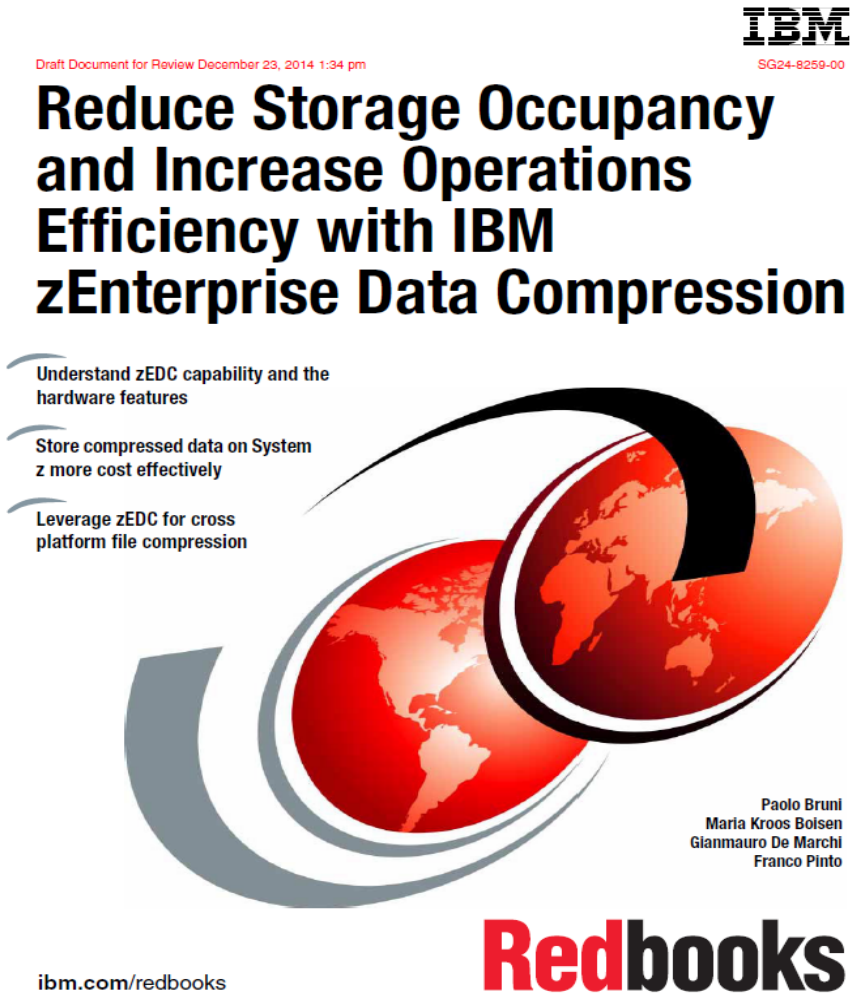

Thank You!