



IBM Systems & Technology Group

MVSE Project Opening Topics

First! Beantown Edition
Session 13574

John Eells
IBM Poughkeepsie
eells@us.ibm.com
13 August 2013



Permission is granted to SHARE Inc. to publish this presentation paper in the SHARE Inc. proceedings; IBM retains the right to distribute copies of this presentation to whomever it chooses.

Trademarks

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

BladeCenter*	FICON*	Parallel Sysplex*	WebSphere*
DataPower*	IBM*	RMF	zEnterprise*
DFSMSdss	IBM logo*	System z*	z/OS*
DFSMSHsm	InfinBand*		

* 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.

Any information contained in this document regarding Specialty Engines ("SEs") and SE eligible workloads 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). 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 SEs 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.

IBM zEnterprise EC12 (zEC12) System Functions and Features

Five hardware models
Hexa-core 5.5 GHz processor chips
Up to 101 processors configurable as CPs, zAAPs, zIIPs, IFLs, ICFs, or optional SAPs (up to 64-way on z/OS V1.10, 100-way on z/OS V1.11 and higher)
Second generation out of order design
Improvements to pre-fetch instructions
Improved processor cache design
Up to 3TB of Redundant Array of Independent Memory (RAIM) – same as z196
Twice the HSA versus IBM zEnterprise 196 (z196) (32 GB vs. 16 GB)
Decimal-Floating-Point Zoned-Conversion Facility
Flash Express (Storage Class Memory-SCM)
1 MB Pageable Large Pages
Dynamic reconfiguration support for Flash Express
2 GB Large Page Support
Optional PLPA, COMMON page data sets
Crypto Express4S cryptographic coprocessors and accelerators
New support for IBM Enterprise PKCS #11 (EP11) coprocessor
DUKPT for MAC and Data Encryption, Europay, Mastercard, and Visa (EMV) CCA enhancements
New and enhanced instructions
IBM zAware
OSA-Express4S and OSA-Express5S (GbE LX and SX, 10 GbE LR and SR, and 1000BASE-T)



(z/OS® support in blue)
(2013 support in red)

IBM zEnterprise® EC12 (zEC12) Hardware: Processor, Memory and System Structure
Wednesday 8:00
zEC12 Hardware: I/O Subsys, I/O and Parallel Sysplex® Features, and Installation Planning
Wednesday 9:30
z/OS Software Support for zEC12 Server
Wednesday 11:00

FICON Express8S
24K subchannels for FICON channels
IBM zEnterprise Data Compression (zEDC) capability using zEDC Express
RDMA (Remote Direct Memory Access) support for z/OS over Converged Enhanced Ethernet RoCE)
Parallel Sysplex InfiniBand® (PSIFB) Coupling Links
High Performance FICON for IBM System z®
CPU Measurement Facility
CFCC Level 18 and 19 enhancements
Transactional Execution Facility
Runtime Instrumentation Facility
Exploitation of new hardware instructions – XL C/C++ ARCH(10) and TUNE(10)
CCA 4.4 and other enhancements: RKX Key Export Wrap, UDX Reduction/Simplification, additional EP11 algorithms, expanded EMV support, AP Configuration simplification
Optional Non Raised Floor
Optional water cooling and DC Power
Optional overhead Power and I/O cabling
IBM zEnterprise BladeCenter® Extension (zBX)Model 003 support of: <ul style="list-style-type: none"> IBM WebSphere® DataPower® Integration Appliance XI50 for zEnterprise Select IBM BladeCenter PS701 Express blades or IBM BladeCenter HX5 blades
Unified Resource Manager (zManager) enhancements

IBM zEnterprise BC12 (zBC12) System Functions and Features

2 Models – H06, H13
Hexa-core 4.2 GHz processor chips
Up to 13 processors configurable as CPs, zAAPs, zIPs, IFLs, ICFs, or optional SAPs
Second generation out of order design
Improvements to pre-fetch instructions
Improved processor cache design
Up to 496 GB RAIM
16 GB HSA separately managed
Up to 6 CPs at 26 capacity points
Decimal-Floating-Point Zoned-Conversion Facility
Flash Express (Storage Class Memory-SCM)
1 MB Pageable Large Pages
Dynamic reconfiguration support for Flash Express
2 GB Large Page Support
Optional PLPA, COMMON page data sets
Crypto Express4S cryptographic coprocessors and accelerators
New support for IBM Enterprise PKCS #11 (EP11) coprocessor
DUKPT for MAC and Data Encryption, Europay, Mastercard, and Visa (EMV) CCA enhancements
New and enhanced instructions
IBM zAware
OSA-Express4S and OSA-Express5S (GbE LX and SX, 10 GbE LR and SR, and 1000BASE-T)



(z/OS support in blue + red)

FICON Express8S
24K subchannels for FICON® channels
IBM zEnterprise Data Compression (zEDC) capability using zEDC Express
RDMA (Remote Direct Memory Access) support for z/OS over Converged Enhanced Ethernet RoCE)
Parallel Sysplex InfiniBand (PSIFB) Coupling Links
High Performance FICON for System z
CPU Measurement Facility
CFCC Level 18 and 19 enhancements
Transactional Execution Facility
Runtime Instrumentation Facility
Exploitation of new hardware instructions – XL C/C++ ARCH(10) and TUNE(10)
CCA 4.4 and other enhancements: RKX Key Export Wrap, UDX Reduction/Simplification, additional EP11 algorithms, expanded EMV support, AP Configuration simplification
Non-raised floor option available
Overhead Cabling and DC Power Options
zBX Model 003 support of: <ul style="list-style-type: none"> ▪ IBM WebSphere DataPower Integration Appliance XI50 for zEnterprise ▪ Select IBM BladeCenter PS701 Express blades or IBM BladeCenter HX5 blades
zManager enhancements

Three Ways to Compress (and Decompress) on z/OS


- New zEDC Express adapter for zEC12 and zBC12 and zEnterprise Data Compression (zEDC) for z/OS V2.1
 - Compression work is offloaded to the card
 - Minimal CP cycles consumed
 - zlib-based, industry-standard deflate compression
 - Data can be inflated anywhere zlib processing is available
- Compression coprocessor-based instructions
 - Dictionary-based compression
 - Can be inflated on a System z processor
 - All compression consumes apparent CP cycles
 - Compression done on the coprocessor, but accounted for as CP busy time since the CP is unavailable until the coprocessor is done
- Using software compression
 - CPU-intensive
 - Much slower
 - Data can be inflated on anything supporting the same algorithm



zEnterprise Data Compression

- At general availability:
 - New card and corresponding z/OS feature
 - Support for industry standard zlib compression
 - zlib library in z/OS V2.1
 - SMF data compression

- Planned*:
 - BSAM/QSAM data compression for Extended Format Data Sets
 - DFSMSdss™ data compression
 - Java™ and IBM Encryption Facility for z/OS support




Why Compress?

- Much better use of disk space
 - No disk compression available for CKD storage
 - Compression rates vary with data...
 - But internal testing shows us ~4X compression for SMF data
 - And, I expect ~2X compression for Extended Format BSAM/QSAM data but have not tested yet
 - Think about doubling the amount of active BSAM/QSAM data in L0 of the storage hierarchy!
 - Better use of fast-but-expensive, SSD-based “disk” storage

- (Note: No advantage for tape—LZ compression is used in the tape controllers already)

Why Compress?

- More effective FICON and network bandwidth
 - Compression done before the channel sees the data
 - Fewer bytes to flow over FICON
 - Effective FICON data rates >1 GB/Sec expected
 - Compressed data requires fewer network packets
 - Effective data rate & link capacity expected to be increased



SMF Data Compression

■ For SMF data written to log streams

- We expect about a 4:1 compression ratio for SMF data
- Designed to significantly increase SMF recording rates
- Can specify that all SMF data or SMF data written to selected log streams be compressed
- New SMFPRMxx COMPRESS keyword on LSNAME and DEFAULTLSNAME
- New PERMFIX subparameter of COMPRESS to balance fix/unfix overhead with available real memory

■ Corresponding IFASMFDL support

- Automatic inflation on z/OS V2.1 with feature and HW support
- SOFTINFLATE parameter for software-based decompression
 - For z/OS V1.12, z/OS V1.13, and the PTF for APAR OA41156
 - For z/OS V2.1 when zEDC is not available

z/OS

MVS System Management Facilities

(SMF)

Planned Compression Support*

- **Extended Format BSAM and QSAM Compression**
 - New support for Compressed Format data sets planned for 1Q2014
 - In addition to generic (DBBLIB) and tailored (supply a dictionary) compression
 - New COMPACTION option in DATACLAS definition
 - New values on COMPRESS parameter in IGDSMSxx

- **DFSMSdss data compression**
 - Planned for DUMP, COPY, and when DFSMSdss is used as the data mover by DFSMSHsm™ for 3Q2014
 - When a disk output data set is used

- **Java and IBM Encryption Facility for z/OS support**
 - Planned for future updates of IBM 31-bit and 64-bit SDK for z/OS Java Technology Edition, Version 7 (5655-W43 and 5655-W44)
 - Support for IBM Encryption Facility for z/OS (5655-P97) planned to coincide with the Java support

What You'll Need to Use zEDC

■ New Hardware and z/OS features:

- zEDC Express adapter for zEC12 and zBC12
- zEnterprise Data Compression (zEDC) for z/OS V2.1
- For software inflation of compressed data, the PTF for APAR OA41156 on z/OS V1.12 and z/OS V1.13
- zlib on other platforms where you want to process compressed data

■ SMF and RMF™ support



- SMF14 and SMF15 records show compression ratios
- SMF14CDS has the size of the compressed-format data set
- SMF14UDS is the uncompressed size
- New SMF14CMPTYPEzEDC field
- SMF 74 subtype 9 records created by RMF include new PCIe, zEDC Express data
- RMF Monitor I PCIe Activity Report:
 - I/O queue and execution time
 - Compressed and uncompressed data transfer rates
 - Number of compression and decompression requests

IBM System z Batch Network Analyzer

- Helping determine if you have files that are candidates for zEDC: the IBM System z Batch Network Analyzer
 - A free, “as is” tool to analyze batch windows
 - Available to Customers, Business Partners and IBMers
 - Replaces the old BWATOOL
 - PC based, graphical and text reports
 - Including Gantt charts and support for Alternate Processors
- Available from NA Advanced Technical Support
 - <http://w3.ibm.com/support/techdocs/atmastr.nsf/WebIndex/PRS5126>
- New news...zBNA planned to help identify zEDC Compression Candidates
 - SMF postprocessing 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 a list of data sets by job which may be zEDC candidates but are not in extended format
 - Target availability – 4Q 2013*



Measurements




■ SMF and RMF support

- SMF14 and SMF15 records show compression ratios
- SMF14CDS has the size of the compressed-format data set
- SMF14UDS is the uncompressed size
- New SMF14CMPTYPEzEDC field
- SMF 74 subtype 9 records created by RMF include new PCIe, zEDC Express data
- RMF Monitor I PCIE Activity Report:
 - I/O queue and execution time
 - Compressed and uncompressed data transfer rates
 - Number of compression and decompression requests

The SHARE MVSE

“Top 40”...er, “39”



The MVSE project had “a lot” of z/OS requirements

- A lot of them were old
- Some were duplicates
- Some were as relevant as ever
- Some had been rendered moot by things unimagined when they were written
- A lot of them had been satisfied!
- A lot of them had not

- Many at SHARE complained....




SHARE Helps Out!

- There were two ways to look at this
 - It was a “target rich environment”; or,
 - The target was “too diffuse to hit”
- There were 800+ MVSE requirements...so “too diffuse to hit” was our view!
- We asked for SHARE’s help in focusing the effort
- SHARE came through!
- Special thanks to these people for boiling down the ocean:
 - Brad Carson
 - Tom Conley
 - Ed Jaffe
 - Robert Jenkins
 - Sam Knutson
 - Mary Anne Matyaz
 - Cheryl Watson
 - Dave Whitney (IBM)
 - (Apologies to anyone I’ve missed!)



The Result:

- Original intent was the “MVSE Top 40”
 - ...with apologies to Casey Kasem...
 - ...but there was a large n-way tie for #40...
- So, we got the “Top 39”
- Due to timing, one was already satisfied in z/OS V1.13
 - PARM('AGGRGROW') Should be Default for zFS Mount
 - (Hey, we like the easy ones!)



The Result:

- In addition, we think we hit another twelve or so in z/OS V2.1:

1. Dynamic system symbol changes
2. Separate wait time limits (partial)
3. Dynamically add and delete MCS consoles
4. Consoles support for HMC 3270 emulator
5. D PPT command
6. Improve IEF212I message
7. Dynamically modify VLF
8. SETSMF without PROMPT at IPL time
9. IEBCOPY partially-qualified member names
10. Multiple ISPF logical screens for SMP/E
11. TSO Logon Failure information (partial?)
12. REXX: Should Support Major Access Methods (partial)

That's ~33%!



**For more detail on these
Top 39ers:**

**What's New in z/OS 2.1 -
Beantown Edition
Monday 1:30**



The Future Runs on System z

Optimize your z/OS environment

