

\* The zEDC Express is an optional feature, exclusive to the zEC12 and zBC12, designed to provide hardware-based acceleration for data compression and decompression for the enterprise, helping to improve cross platform data exchange, reduce CPU consumption, and save disk space.

\* A minimum of one feature can be ordered and a maximum of eight can be installed on the system, on the PCIe I/O drawer.

- Up to two zEDC Express features per I/O domain can be installed.

- A zEDC Express feature can be shared by up to 15 LPARs.

\* Using the new z/OS zEnterprise Data Compression (zEDC) capability with a new I/O feature on the zEC12 or zBC12 – zEDC Express enables customers to do hardware based data compression.

- The combination can help cross-platform data exchange, save storage and reduce CPU consumption.

\* IBM plans that the IBM Encryption Facility for z/OS, will exploit zEnterprise Data Compression (zEDC) for z/OS V2.1, running on zEC12 and zBC12 zEnterprise servers with the zEDC Express adapter when the Java release supporting zEDC becomes available.

- This will complement the software compression support that exists today with Encryption Facility OpenPGP support.

NOTE: Support for data recovery (decompression) in the case that zEDC is not installed, or installed but not available, on the system, is provided via software on z/OS V2R1, and on V1R13 and V1R12 with appropriate PTFs. Software decompression is slow and will use considerable processor resources, thus it is not recommended for production environments.

\* zEDC can help to improve SMF logger compression with less system overhead.

- Corresponding support in the SMF dump program IFASMFDL is designed to support

both hardware-based and software-based decompression.

\* Decompression support is available on z/OS V1.12 and z/OS V1.13) with the PTF for APAR OAA41156.

- This new function is expected to allow higher write rates for SMF data when hardware compression is enabled.

\* RMF support for hardware compression includes SMF Type 74 subtype 9 records and a new Monitor PCIE Activity report intended to provide information about compression activity on the system.

\* z/OS V2.1 leverages use of the industry-standard zlib open source library available for z/OS UNIX System Services.

NOTE: This version of the library supports the sending of compression and decompression requests to the zEDC Express. The z/OS-provided zlib library is provided as a UNIX archive file that can be statically linked into IBM, ISV, or customer applications that currently use zlib, enabling additional exploitation of compression through zEDC Express and expanding potential compression opportunities.

\* zEDC is complementary to existing System z compression technology.

- Smaller records and files best suited for hardware chip compression will still use hardware compression instructions to get best benefit, and larger files that should need a different compression algorithm will be directed to zEDC.

#### How is zEDC different than what's available today?

1. zEDC compression acceleration is an efficient alternative for many types of files whereas on chip compression is optimized for short records such as database rows

2. zEDC is 'dictionary - less' and uses optimized algorithms to deliver high performance. It is specially optimized for larger data files.

3. zEDC is compatible with industry standard, open zlib based compression – used today by Java and other applications. z/OS V2.1 provides the zlib\* library which supports the sending of compression and decompression requests to the zEDC Express.

NOTE: The z/OS provided zlib library is provided as a UNIX® archive file that can be statically linked into IBM, ISV, or customer applications that currently use zlib, enabling additional exploitation of compression through zEDC Express and expanding potential compression opportunities.

\* Adapter support for zEDC is provided by Resource Group (RG) code running on the system Integrated firmware processor (IFP).

- For resilience, there are always two independent RGs on the system, sharing the IFP.

- It is, thus, recommended that a minimum of two zEDC features be installed, one per RG.


- Consider also the total data throughput required and that, in the case of one feature becoming unavailable, the others should be able to absorb the load.

NOTE: For best data throughput and availability, it is suggested that at least two features per RG are installed.

\* Integrated firmware processor (IFP) is a single PU dedicated to the support of the native PCIe features such as 10GbE zEDC Express.

SOD: In a future z/VM deliverable IBM plans to offer z/VM support for guest exploitation of the zEDC Express feature on zEC12 & zBC12.

Every day 2.5 quintillion bytes of data are created



A new compression acceleration capability is available beginning with z/OS V2.1. Up to 24 X throughput improvement with zlib. Up to 118 X in CPU reduction. Up to 4 X Compress Data.

#### Typical Client Use Cases:

- Significant disk savings with trivial CPU cost for large BSAM/QSAM sequential files (Reduced logger overhead allows collection of more SMF data).
- More efficiently store audit data in application logs
- Increase the amount of active data you can keep active by compressing more frequently accessed data
- Reduce the amount of data needed for data migration and backup/restore (results based on customer workloads)
- Transparent acceleration of Java compressed applications
- DFSMS™ planned exploiters TBD

There is one compression coprocessor per zEDC Express feature.

The IBM System z Batch Network Analyzer (zBNA) can be used to assess eligible applicants. It's a free, "as is" tool providing graphical and text reports, including Gantt charts.

SOD: IBM plans for future updates of IBM 31-bit and 64-bit SDK7 for z/OS Java Technology Edition, Version 7 (IBM SDK7 for z/OS Java) to provide exploitation of the zEDC Express feature.

15% reduction in elapsed time for SMF extraction with up to 40% reduction for CPU time. In addition, Logger overhead is reduced by up to 30%.

[The zEDC Express is feature code FC# 0420 installed in PCIe I/O drawer] Optimized for active cross platform data exchange. Enables compression of active and inactive data.

Useful for files that previously used software compression as well.

#### More news...zBNA will identify 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 a list of data sets by job which may be zEDC candidates but are not in extended format Target availability – 4Q 2013.



#### 'zlib' is a software library used for data compression:

zlib compressed data is typically written with a gzip or a zlib wrapper. The wrapper encapsulates the raw DEFLATE data by adding a header and trailer. This provides stream identification and error detection that are not provided by the raw DEFLATE data.

Operating system requirements: Requires z/OS 2.1 and new zEDC for z/OS feature

- z/OS V1.13 and V1.12 offer software decompression support only

- Easy to set up and use – transparent to application software ; Use policy (DATACLASS) to set up compression.

- No changes to access method

Server requirements: Exclusive to zEC12 (with Driver 15E) and zBC12

- New zEDC Express feature for PCIe I/O drawer (FC#0420)One compression coprocessor per zEDC Express feature

- Each feature can be shared across up to 15 LPARs

- Recommended minimum configuration per server is two featuresUp to 8 features available on zEC12 or zBC12

- For best performance, feature is **needed on all systems accessing the compressed data**

Planned exploitation: Hardware exploitation first for log files - SMF records (September 2013) reduced logger overhead allows collection of more SMF data

- All systems sharing sequential BSAM/QSAM extended format (1Q'14)

SOD: Java using standard zlib compression library for compression services. Java applications and middleware can be transparently accelerated by enabling Java for hardware compression

SOD: DFSMS™ planned exploiters TBD.

\* zEDC disk compression extends the value of your DS8000 disk flash purchase by using less flash to hold the same amount of data. zEDC will allow you to more fully utilize the flash space.

\* The ability to compress data is usually provided via a specific product, service or client application.

NOTE: If data flows between such software then it is possible that data is compressed twice without knowledge of the zEDC Express or existing compression technology. Any interactions should be examined to determine the correct location and technology to be used for the compression.

But wait. Doesn't IBM offer a no-charge hardware compression today on every z processor chip?

Ans. YES. Hardware compression is available today on the System z processor chip. It's the CMPSC instruction and is IBM proprietary compression that uses a dictionary to compress data.

DB2 is one of the main use cases for this compression which is optimized for cases where row-wise access to the data is required. DFSMS also utilizes this hardware compression with generic compression available for VSAM and non-VSAM extended format data sets, as well as tailored compression available for non-VSAM extended format data sets.

NOTE: This new compression is 'different'. zEDC is dictionary-less. zEDC scans text and looks for the re-use of a phrase and puts in a 'back' reference. For large files, zEDC can be more efficient. Compression is done on the new zEDC Express feature, having little effect on CPU resources.

#### IBM Mainframe Compression Technologies

| Type  | Optimized for  | Performance Overhead  | Supported data   | Frequency of access post compression   |
|---|--|---|--|--|
| CMPSC compression on System z processor chip  | Row-wise access to data, optimal for DB2 or select DFSMS files                             | On chip, relatively little CPU overhead and less I/O; fast                                  | • DB2 – optimized for row-wise access to data is required<br>• DFSMS files – for VSAM and non-VSAM extended format data sets                     | Often  |
| Other software compression (zlib, or similar) | Most compression uses industry standard today. Used by many file types                     | Higher CPU – software instructions executed. Note: if Java then eligible for zIIP (or zAAP) | Any. De facto standard for almost any type of data   | Often  |
| Tape hardware compression                     | Optimized for use with large files, archival purposes                                      | Performed by the tape subsystem   | Any  | Often to rare – application dependent  |
| Archival / Backup                             | Archive data and data to backup/copy   | CPU overhead, longer wall clock time  | DFSMSshsm, DFSMSdsss   | Designed for active primary data   |
| Real time compression                         | IBM NAS storage  | No performance degradation  | SVC  | Designed for active primary data   |
| zEDC compression                              | Active, for cross platform data exchange. Enables compression of active and inactive data. | Processing on zEDC Express – expect minimal CPU overhead, low latency                       | • SMF through logger * zlib<br>• SODs – Java<br>• SOD - BSAM/QSAM Extended format<br>• SOD - DFSMSshsm, DFSMSdsss<br>• SOD - Encryption Facility | Needs frequent access. Useful for files that previously did software compression via zlib. |

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