

Research Report

IBM's On-Chip Integrated Accelerator for zEnterprise Data Compression (zEDC)

Executive Summary

A key design goal in the design of central processing units (CPUs) is to offload non-essential work to other processors or specialized adapters. By doing so, the CPU can be focused on its primary task of performing mathematical calculations – churning thread after thread executing real work rather than handling system functions.

On the other hand – by architecting and placing certain system functions in logic close to a CPU where those functions can be processed more expeditiously – stupendous improvements in the performance of those tasks can be achieved; Quality of Service (QoS) can be dramatically improved; and cost for associated hardware (such as storage or peripherals) can be significantly lowered.

So, when IBM announces that it will move a system function onto its CPU and co-locate that function with its microprocessor cores – **IT IS BIG NEWS**. By co-locating a system function on the same processor die, IBM is saying: "It is extremely important to use CPU cycles to execute this particular system function. The cost in CPU cycles is vastly outweighed by the benefits that will be achieved by handling this system function at the CPU level."

What Has IBM Announced?

When IBM introduced its zEC12 system architecture several years ago, the company also introduced a specialized adapter for data compression known as the zEDC Express I/O adapter. The intent of this adapter was to offload the CPU from having to handle bulk compression functions.

A few years later, IBM took a more subtle step with its z14 chip design by optimizing in-core compression acceleration. By doing this, IBM improved start/stop latency (wait time); it improved its compression ratio using Huffman encoding; and it introduced order-preserving compression techniques. These were the first baby-steps in what has now become a major shift in IBM Z processor and system design: extensive on-chip data compression.



When IBM announces that it will move a system function onto its CPU and co-locate that function with its microprocessor cores – IT IS BIG NEWS.

With IBM's newest Z, the z15, IBM has announced that it is discontinuing its zEDC Express I/O adapter approach – instead moving compression facilities on-chip. This new on-chip compression facility is known as “Integrated Accelerator for zEnterprise Data Compression.”

With the latest addition to the IBM Z® family, the z15, IBM has announced that it is discontinuing its zEDC Express I/O adapter approach – instead moving compression facilities on-chip. This new on-chip compression facility is known as “Integrated Accelerator for zEnterprise Data Compression.”

By co-locating compression with its processor, IBM can focus on:

- Substantially reducing the amount of data (the size of files) that are being transferred intersystem or over networks;
- Speeding overall throughput over storage and communications networks;
- Reducing the amount of data that has to be encrypted/decrypted (saving systems overhead);
- Allowing more data to remain online or in the cloud (due to smaller file sizes); and,
- Help reduce storage costs (making the purchase of advanced storage arrays and flash memory more affordable than ever before.)

So far, the preliminary results of this co-location effort have been, as expected, stupendous:

- **The largest IBM z15 with the Integrated Accelerator for zEDC provides up to 17x the total compression throughput of a z14 configured with the maximum number of zEDC Express cards.** (Disclaimer: Measurements were collected in a controlled environment running an IBM-developed workload under z/OS comprised of 100% compression. Individual results may vary. Results are workload dependent.)
- **Up to 8x faster application elapsed time with no additional CPU time using IBM z15 Integrated Accelerator for zEDC compared to z14 zEDC Express for compression/decompression.** (Disclaimer: Measurements were collected in a controlled environment running an IBM-developed workload under z/OS comprised of an equal mix of compression and decompression. Individual results may vary. Results are workload dependent. This claim is based on IBM zlib support, where the elapsed time is greatly reduced by staying in an on-chip synchronous mode as compared with using the zEDC Express I/O adapter.)



- **Compresses up to 260 GB/sec with the Integrated Accelerator for zEDC on the largest IBM z15.** (Disclaimer: Measurements were collected in a controlled environment running an IBM-developed workload under z/OS comprised of 100% compression. Individual results may vary. Results are workload dependent).

The Use Cases Abound

The ways to exploit IBM's new compression facility abound, including the compression of:

- SMF (Systems Management Facility) logstreams – for increased availability, reduced online storage, and for decreased system logger file sizes)
- QSAM (queued sequential access method) and BSAM (basic sequential access method) data sets
- DFSMSdss (Data Facility Storage Management subsystem)/ DFSMSHsm (Data Facility Storage Management) – when backing up and restoring data or migrating and recalling data
- IBM Java V7.0.0 SR7, Java V7R1 runtime, and Java V8.0 environments – for higher throughput using standard compression
- IBM Encryption Facility files – for building industry-standard compressed Open PGP privacy files
- IBM Sterling Connect: Direct for z/OS files – for better throughput and link use
- IBM WebSphere MQ messages – for channel message compression
- Applications that use zlib – for application programs that directly use compression with the zlib open source library application programming interfaces

How Has IBM Implemented Its Integrated Accelerator for zEnterprise Data Compression?

IBM has implemented this accelerator as part of a nest accelerator function where one accelerator unit is dedicated to each processor chip – and shared by the cores on the chip. For IBM Z, using this nesting approach is a new concept in accelerator design – enabling lower latency compression at high bandwidth speeds, problem state execution combined with hardware/firmware interlocks to ensure system responsiveness. Instructions are executed in millicode.

The commands for compression can be found in the z/OS zlib (compression library.) This is an open source C language library that provides compression and decompression. It supports RFC1950 (zlib), RFC1951 (DEFLATE), and RFC1952 (GZIP). And zlib supports a streaming model such that files can be compressed/decompressed in chunks.

The new Integrated Accelerator for zEDC operates in dual modes: 1) a synchronous (simultaneous, at the same time) mode specific to zlib when invoked by an application instruction in a virtual address space. An easy and straightforward approach that does not require special hypervisor support to execute; and 2) an asynchronous (step-at-a-time) mode for large operations under z/OS. Using this approach, an application (such as BSAM or QSAM) issues a request for I/O asynchronous execution, which is carried out by a special assist processor on behalf of the application. This approach enables load balancing of high-compression workloads, offers low latency and high bandwidth performance compared to the zEDC Express adapter, and can be transparently executed by authorized users. Further, it is important to note that these cycles are not run on the CPU, but rather on a special assist processor.

Compatibility

IBM has taken great care to ensure backwards/forward compatibility. When using the new Integrated Accelerator for zEnterprise Data Compression, all z/OS configurations stay the same. Nothing is required to change when z/OS is moved from a z14 to z15. IBM does caution, however, that due to the massive capacity available on the z15, customers ensure that there is enough capacity available on z13 or z14 servers for proper fail-over and disaster recovery handling.

Summary Observations

Simply put: compressing data substantially reduces the amount of data that needs to be moved or transferred. Instead of accessing or moving large files, smaller, compressed files with the same data can be accessed and processed. Is using valuable CPU real estate to handle data compression justifiable? Do the pros (reduced data volume, speed, reduced latency, reduced network traffic, more efficient encryption, more data available online, reduced storage costs) outweigh the cons (the use of CPU cycles)? Our answer to this question would be a resounding: “yes!”

It is clear to us that IBM has now made data compression a top priority – largely because the amount of data input/output subsystems have to deal with are creating processing bottlenecks. Compressing data at the processor chip level will shrink file sizes and help eliminate those bottlenecks. And less space will be used in storage. Both of these are highly desirable from a customer perspective. Customers are also aware that batch workloads are now dealing with more data than ever before. Compressing data will shrink the size of data sets while enabling batch workloads to process more data in less time.

Three years ago, we watched the IBM Z organization change the industry by moving encryption/decryption security functions onto the Z microprocessor... we expect the same result: a major change in how the industry handles compression, with IBM Z again in the leadership role.

Further, compression provides for access to more online data, which is important to customers who have had to take data offline due to storage limitations. Business opportunities are being lost because the data needed is sometimes taken offline. Compression enables more data to be kept online, making more data available to analytics applications and opening new opportunities for businesses.

Finally, consider that data is being exchanged regularly between business partners. Compression shrinks the size of files, speeds network transfer, and reduces storage costs – all desirable outcomes between business partners.

Three years ago, we watched the IBM Z organization change the industry by moving encryption/decryption security functions onto the Z microprocessor – making way for “pervasive encryption.” For relatively little systems overhead, IBM customers are seeing a dramatic improvement in security processing speed. No other system rivals IBM Z in this respect. Now, with the introduction of IBM’s new system (the z15), IBM has done it again. This time IBM has committed microprocessor development resources to moving data compression on-chip – and we expect the same result: a major change in how the industry handles compression, with IBM Z again in the leadership role.



Clabby Analytics LLC
<http://www.clabbyanalytics.com>
Telephone: 001 (843) 297-5150

© 2019 Clabby Analytics
All rights reserved
August, 2019

Clabby Analytics is an independent technology research and analysis organization. Unlike many other research firms, we advocate certain positions – and encourage our readers to find counter opinions – then balance both points-of-view in order to decide on a course of action. Other research and analysis conducted by Clabby Analytics can be found at www.ClabbyAnalytics.com.