Redbooks

ibm.com/redbooks

Enhancing Value to Existing and Future Workloads with IBM z13

Harry Yudenfriend

IBM Academy of Technology

Learn

IBM z Systems

IBM

Point-of-View

Enhancing Value to Existing and Future Workloads with IBM z13

An IBM Redbooks Point-of-View publication from IBM Systems

By Harry Yudenfriend, IBM Fellow

Highlights

Discover how to take advantage of advances in technology to address key demands on today's mainframes:

- Cloud computing, analytics, mobile access, social media, and security (CAMSS) put new demands on businesses and IT for enhanced resilience, performance, and scale.
- The volume of data and increased file sizes push networks to their limits and stress quality of service.
- A powerful, workload-optimized, open, agile, and highly secure IT infrastructure is required to meet customer expectations.

The business environment demands responses from technology

Technologies are rapidly changing the business environment. Cloud computing is increasingly seen as a growth engine for business with new applications and capabilities. The proliferation of mobile devices, social media access, massive amounts of data, and the ever-looming threat of security breaches, make it imperative for organizations to adapt to meet these demands. More and more organizations are relying on real-time insight from structured and unstructured data to make important decisions and to drive customer interactions. Having a robust IT infrastructure is essential in order for organizations to thrive in this new reality.

A recommendation made by the IBM® Institute of Business Value concerning these issues is that *an IT infrastructure must adapt to changing technologies and business requirements, and be prepared to address numerous disruptions that can halt a digital enterprise.*[1]

These new dynamics are reframing the conversations on IT infrastructure. In addition, C-level executives are acknowledging that a robust IT infrastructure is essential. The increased reliance on data as a vital resource combined with the expectations of today's customers are forcing organizations to search out new and innovative ways to use hardware, software, networking, and storage.

The IBM z Systems™ have many distinct advantages in this dynamic business environment:

- z Systems are built on years of IBM leadership in virtualization and can host more virtual machines in a single footprint than any other platform. Its superior management services enables a mainframe's resources to be shared among workloads, achieving the highest possible utilization of the platform.[2]
- z Systems are architected to operate with the highest levels of reliability and availability. Each component is built to be resilient with support for nondisruptive configuration changes and dynamic replacement capabilities.

[1] *The IT infrastructure conversation*, IBM Institute for Business Value, 2014, GBE03612-USEN-01, found at: http://www-01.ibm.com/common/ssi/cgi-bin/ssialias?htmlfid=GBE03612USEN&appname=skmwww#loaded

[2] *Comparing Virtualization Alternatives - What's best for your business? Solitaire Interglobal Report, 2012. found at:* http://public.dhe.ibm.com/common/ssi/ecm/en/zsl03192usen/ZSL03192USEN.PDF

- z Systems have garnered some of the highest marks for security in the industry.
- It has been found that a mainframe-based cloud implementation can cost less, overall, than a cloud built on other, competing platforms.

IBM z Systems I/O plays an important role in supporting workloads

No matter where you start in your IT infrastructure analysis, the ability to access, process, and move data to support your workloads is a key factor. IBM's strategy ensures that IBM mainframes provide the I/O capabilities necessary to handle the volume, velocity, processing power, and resilience needed for the new demands on and for data. The ability to provide consistent and secure interfaces and the ability to get the right data to the right people are essential for today's environment. The z Systems with its forward thinking architecture is positioning itself to handle these demands now and in the future.

The growth in cloud computing, analytics, mobile access, social media, and security (CAMSS) drive additional requirements on z Systems. It must be able to execute transactions with the same service level agreements (SLA) and enterprise class qualities of service that the most demanding of z Systems clients have come to expect. For example, adding new data sources from the cloud into a workflow might increase elapsed times. Mobile access to data on z Systems might add an unpredictable increase in transaction volume resulting in I/O contention. It is critical that middleware (such as, IBM DB2®, IBM CICS®, IBM MQ, and IBM IMS™) can scale to meet these demands. Lowering database transactional latency is critical to mitigate the impact of new data sources and transaction volumes on both the existing and new workloads.

The IBM z13™ is built from the casters up for CAMSS workloads. The I/O enhancements provides reduced transactional latency, thus mitigating increases to transaction response times that might be introduced by adding cloud sourced data to the work flows. These performance improvements also enhance the scalability of IBM z/OS® and middleware, such as DB2, to meet the demands of mobile applications. The increase in the volume of I/O is driving the requirement for further improvements to the already industry leading quality of service (QoS) capabilities of z Systems. These new QoS enhancements affect processing in the following ways:

- Reduce transactional latency to mitigate increases in transaction response times introduced by adding new data sources and workloads.
- Improve resilience by enhancing the quality of the Fibre Channel links to reduce error rates and increasing performance to handle workload spikes and contention that might occur after hardware failures.
- Extend the z/OS workload manager into the storage area network (SAN) fabric to manage the end-to-end work according to client specified policy. This capability includes all read and write operations and the synchronous replication activity generated by write operations using the IBM DS8870 Metro Mirror technology.
- Provide reduced cost for the physical infrastructure with enhanced virtualization to allow sharing more of the enterprise I/O traffic over shared inter switch links (ISLs).
- Provide improved availability with additional flexibility and scale to the I/O configuration.

z13 is the latest generation of the IBM mainframe to deliver substantial value to the I/O infrastructure:

- IBM Fibre Connection (FICON®) Dynamic Routing is a feature that enables use of SAN dynamic routing polices in the fabric to lower costs by allowing clients to share FICON and Metro Mirror (FCP) on the same inter-switch links (ISLs). This feature can reduce the cross site bandwidth requirements and simplify configuration definitions and performance management by providing predictable and repeatable performance.
- Mainframe SAN Fabric Priority with the IBM DS8870, or other storage technologies, extends the z/OS Workload Manager (WLM) to the SAN infrastructure providing improved resilience and autonomic capabilities while enhancing the value of FICON Dynamic Routing. Fabric priority is supported for FICON Express16S and FICON Express8S channels on the IBM z13.
- FICON Express16S with the IBM DS8870 and 16 Gbs host bus adapters, enables the IBM z13 to provide up to 32% reduced I/O service times for DB2 log writes as compared to FICON Express8S on EC12.This difference can reduce DB2 transactional latency by up to 19%. Additionally, I/O bound batch jobs can also reduce elapsed time by up to 32%, thus helping to relieve pressure on the batch window.

- FICON Express16S channels provide Forward Error Correction (FEC) technology. The host bus adapters (HBAs) enable the Fibre Channel link protocol to operate at higher speeds, over longer distances, with reduced power and higher throughput, while retaining the same reliability and robustness that FICON is traditionally known for.
- IBM High Performance FICON for z Systems (zHPF) Extended Distance II feature with the IBM DS8870 or other supporting storage, reduces the impact of distance on zHPF I/O response times by up to 50% at 100KM for large write operations (>64K), providing significant response time improvements after a HyperSwap® for multi-site IBM Parallel Sysplex® environments. FICON Express16S and FICON Express8S support the zHPF Extended Distance II protocol to DS8870 16 Gbs capable ports.
- Scaling to six logical channel subsystems (LCSS) allows for up to 85 client usable LPARs.
- All z13 FICON channels (FICON Express8, FICON Express8S, and FICON Express16S) support up to 32K devices per channel, up from 24K on the EC12.
- A fourth subchannel set for each logical channel subsystem (LCSS) is provided to eliminate single points of failure for storage after a disk failure. It facilitates the exploitation of IBM DS8870 Multi-target Metro Mirror storage replication with IBM Geographically Dispersed Parallel Sysplex™ (IBM GDPS®) and IBM TotalStorage Productivity Center for Replication (IBM TPC-R) HyperSwap.

Improve quality of service

IBM z Systems input/output (I/O) channels provide a pipeline through which data is exchanged between systems or between a system and external devices. Fibre Connection (FICON) is an IBM I/O technology (used by z Systems) built on the Fibre Channel architecture.

Reducing transactional latency mitigates increases to transaction response times introduced by adding workloads. Also, performance improvements might improve the scalability of IBM z/OS and middleware, (such as, DB2) to meet the demands of mobile applications and analytics. The increase in the volume of I/O activity drives the requirement for further improvements to the already industry leading quality of service capabilities of z Systems.

The FICON channel matches data storage and access requirements with the latest technology in servers, control units, and storage devices. The FICON Express16S channels allow faster and more efficient data transfer while at the same time allowing you to use currently installed single mode and multimode fiber optic cabling. FICON uses long wavelength (LX) and short wavelength (SX) transceivers with multimode and single mode fiber optic media for data transmission.

FICON Express16s channels reduce I/O latency. This feature is valuable for reduced transactional latency for databases and shrinking the batch window for I/O bound batch work. IBM performance testing shows that FICON Express16s and DS8870 16 Gbs Fibre Channel ports improves I/O service times between by up to 32% as compared to FICON Express8S running on the IBM EC12.[3]

In support of the IBM z13 Mainframe, the IBM DS8870 storage subsystem also adds enterprise class qualities of service to further provide value for z System clients. The IBM DS8870 has issued a statement of direction (SOD)[4] for 2Q2015 for the support of 16 Gbs host bus adapters to complement the FICON Express16S channel. In addition, QoS features include Forward Error Correction codes for optical links to improve reliability, comprehensive SAN fabric priority controlled by the z/OS Workload Manager and propagated to replication traffic to improve resilience and FICON Dynamic Routing support to help clients reduce cost.

Streamline I/O execution

zHPF is an optional feature that streamlines z Systems I/O execution for some I/O capabilities. Essentially, zHPF uses the hardware assists built into Fibre Channel host bus adapters (HBAs) for the FCP protocol, to transport z Systems I/O commands. At the same time, zHPF preserves the z Systems traditional qualities of service, such as:

- In-band I/O measurements
- End-to-end data integrity protection transparent to applications and middleware
- Workload management
- Self-describing components for recognition of single points of failure

The main value of zHPF is I/O performance. With zHPF, fewer Fibre Channel sequences are needed to execute an I/O operation. I/O latency is reduced and

[3] Fifty percent mix of reads and writes, 256K per I/O operation

[4] Release memorandum - RFA 61029 - 01/14/15 (RFA Guide 29.08), section 2.7.7, FICON Express16S - a new generation for FICON, zHPF, and FCP.

there is less contention at the channel, control unit ports, and in the SAN fabric.

zHPF supports multitrack operations. It allows the channel to operate at rates that fully use the bandwidth of a FICON Express channel. The zHPF device commands fully support multiple tracks of data that can be transferred in a single I/O operation.

The zHPF Extended Distance II feature available with the z13 and IBM DS8870 (2Q2015) improves the I/O service time for executing large write operations (> 64 KB) at long distances. FICON and zHPF both support 100 KM distances. The FCP transport used by zHPF typically requires an interlocked round trip between the channel and control unit for each 64KB to be written. Enhancements to zHPF are provided to help eliminate the multiple interlocked round trips thus reducing the I/O service for large write operations at a distance. Long distance writes can occur, for example, after a Hyperswap event or in a multi-site workload with application systems running in both sites

By addressing limits on the FICON channels, consolidating multiple storage subsystems on the same set of channels (such as, a switch distributed star topology) is possible. This action increases the scale of the I/O configuration and facilitates the use of other I/O technologies by freeing I/O adapter slots. For example, zEDC, RoCEE Express and Flash Express all require I/O adapter slots.

Support for dynamic routing policies in the SAN

With z13, FICON channels are no longer restricted to the use of static SAN routing policies for cascading FICON directors. This enhancement is provided to reduce cost, improve performance, simplify management, and position z Systems for future innovations is SAN technology. The z Systems feature that supports dynamic routing in the SAN is called FICON Dynamic Routing (FIDR). FIDR is designed to be switch vendor agnostic. Clients need to determine which switch vendors support FIDR or only enable FIDR for switches that support FIDR. It is important to understand that dynamic routing policies must be used only with supporting switches and devices. z/OS health checks and HCD function are added to help clients identify devices that support FIDR and switches that are running with FIDR enabled.

The IBM z13 provides an architected means for the host operating system software to specify an I/O priority for the SAN fabric use. This capability allows z/OS to extend the z/OS Workload Manager (WLM)[5] to manage the SAN fabric, thus completing the management of the entire end-to-end flow of an I/O operation.[6] The z/OS operating system manages the flow by first querying the range of supported fabric priorities from each SAN switch in the I/O configuration. From this information, z/OS derives a priority range supported for all the physical switches. WLM then assigns I/O priorities consistent with the client specified goals for the workloads within the supported range.

With the Fabric I/O priority specified by WLM, the write operations are managed by the fabric. For read operations, when the IBM DS8000® sends the data to the host, the DS8870 uses the I/O priority specified by WLM on the original read I/O command. This priority setting enables the fabric to manage the return of the data to the host according to the appropriate priority. As more and more I/O work shares a pool of ISLs, fabric priority is a critical capability for managing performance toward the client goals. The IBM DS8870 is also enhanced to propagate the fabric priority of the write operations that require replication services such as, IBM Metro Mirror to the FCP-based generated I/O activity. This innovation provides a true end-to-end workload management capability for z/OS traffic through the SAN infrastructure.

As more and more I/O work shares a pool of ISLs, fabric priority is a critical capability for managing performance toward the client goals. The IBM DS8870 is also enhanced to propagate the fabric priority of the write operations that require replication services, such as, IBM Metro Mirror to the PPRC FCP-based generated I/O activity. This innovation provides a true end-to-end workload management capability for z/OS traffic through the SAN infrastructure.

Reduce I/O latency and improve I/O resiliency

The primary value of moving to faster link technologies is reduction of I/O latency for “large” read/write operations. The time for DB2 to write to the log can also be quite substantial. A reduction in log write latency improves DB2 transactional latency. As stated

[5] Robert Vaupel, Managing Workloads in z/OS, IBM Systems Magazine, January 2004, found at: http://www.ibmsystemsmag.com/mainframe/administrator/systemsmanagement/Managing-Workloads-in-z-OS/

[6] Harry M. Yudenfriend, IBM System z® Innovations Automatically Define Configurations for Greater Availability, IBM Systems Magazine, July 2013, found at: http://www.ibmsystemsmag.com/mainframe/storage/Data-Management/zdac_yudenfriend/

previously, FICON Express16S with DS8870 16 Gbs links can reduce I/O service times for the DB2 log write up to 32%. For a well-tuned SAP banking workload, 60% of the transactional latency is spent waiting for synchronous I/O. This synchronous I/O is devoted to writes to the DB2 log. A 32% reduction applied to 60% of the latency yields a 19% improvement to the transactional latency, a substantial improvement.

It is important to understand two important facts about the history of the z Systems FICON Express channels and their growing capabilities in terms of increasing I/O rates and link bandwidth. First, the growth in I/O rates has been mostly from High Performance FICON for z Systems (zHPF). Second, the growth in link bandwidth also comes mostly with zHPF exploitation. Both of these trends will continue.

The faster link speeds are more sensitive to the quality of the cabling infrastructure. In the past, if the optical fibers were abused (such as if they were subjected to excessive bending and twisting, leaving the dust covers off, or faulty connections at patch panels), clients might have seen an increase in I/O error rates after transitioning to a faster link speed (for example, 4 Gbs to 8 Gbs). IBM z Systems brings an Enterprise Class of quality to the FICON Express16S channel to help smooth the transition to faster link speeds.

With the FICON Express16s channel, IBM added Forward Error Correction (FEC) capabilities to the Fibre Channel link protocol, using the most advanced FEC coding in the industry. FEC allows FICON channels to operate at higher speeds, over longer distances, with reduced power and higher throughput, while retaining the same reliability and robustness that FICON has traditionally been known for.

With FEC enabled, errors that might occur the faster link speeds will likely be corrected by the error correction technology in the optical transmit/receive ports. Clients will experience fewer I/O errors, thereby easing the transition to the new link technologies and reducing the potential impact to any production workloads by I/O errors. For optimal value, the end-to-end channel to control unit path needs to run at 16 gigabit link speed for the optimal latency reduction. Likewise, each link, the entire path from the channel through the switch ports to the storage subsystem, should be protected by an FEC capable link to minimize the risk of I/O errors.

z Systems provide two tools for testing the I/O configuration:

- IBM System z I/O Exerciser simplifies the chore of exercising the I/O connections in the I/O configuration before bringing up z/OS and running production work. It helps identify possible cabling or definition errors by validating that all the paths defined to each device actually connect to the same physical device. Additionally, a data transfer stress test is performed to help verify the quality of the cabling infrastructure.
- IBM z/OS I/O Exerciser runs as a started task under z/OS. It can be used to test individual device ranges in a controlled fashion. To obtain this tool, go to this web page:

  https://www14.software.ibm.com/webapp/iwm/web/preLogin.do?source=swg-beta-ibmioexzos

When errors are detected, how does the client identify the root cause? In other words, which connection is faulty? The new T11 Read Diagnostic Parameters Extended Link Service (RDP ELS) defines a method for some intelligence in the SAN fabric to retrieve standard counters that describe the optical signal strength (send and receive), error counters, and other critical information for determining the quality of the link. SAN management tooling can be used to display the link information for every architected point in the connection.

In a future release of z/OS, after a link error is detected (for example, IFCC, CC3, reset event, or link incident report), link data returned from Read Diagnostic Parameters is used to differentiate between errors due to failures in the optics versus failures due to dirty or faulty links. Additionally, z/OS health checks will be delivered to detect when the end-to-end link speeds are inconsistent and if all paths to a control unit have inconsistent link speeds. This capability simplifies diagnosing performance problems. Finally, z/OS operator commands will allow clients to display the diagnostic parameters, which includes information such as the optical signal strength, so that clients can pinpoint bad links without having to move around the floor inserting light meters.

What's next: How IBM can help

IBM has a broad range of expertise to help you incorporate existing, new, and emerging workloads into your z Systems environment. IBM can help you with any z Systems issues you have around these workloads.

The z Systems continue to make improvements that support new and dynamic workloads. For more information about z Systems, contact your local IBM representative.

Resources for more information

For more information about the concepts highlighted in the paper, see the following resources:

- *Get More Out of Your IT Infrastructure With IBM z13 I/O Enhancements,* REDP-5134

  http://www.redbooks.ibm.com/abstracts/redp5134.html?Open

- *IBM System z Connectivity Handbook*, SG24-5444

  http://www.redbooks.ibm.com/abstracts/sg245444.html?Open

- *IBM zEnterprise System Technical Introduction*, SG24-8050

  http://www.redbooks.ibm.com/abstracts/sg248050.html?Open

- *System z End-to-End Extended Distance Guide*, SG24-8047

  http://www.redbooks.ibm.com/abstracts/sg248047.html?Open

- *Security on the IBM Mainframe: Volume 1 A Holistic Approach to Reduce Risk and Improve Security*, SG24-7803

  http://www.redbooks.ibm.com/redpieces/abstracts/sg247803.html?Open

- *Understanding Mainframe Economics*, REDP-5127

  http://www.redbooks.ibm.com/abstracts/redp5127.html?Open

- IBM z Systems product page

  http://www-03.ibm.com/systems/z/?lnk=mprSY-sysz-usen

- IBM z Systems Academic Initiative

  http://www-03.ibm.com/systems/z/education/academic/schools_ap.html

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to: *IBM Director of Licensing, IBM Corporation, North Castle Drive, Armonk, NY 10504-1785 U.S.A.*

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of 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.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

Note: The IBM statements regarding its plans, directions, and intent are subject to change or withdrawal without notice at IBM's sole discretion. Information regarding potential future products is intended to outline our general product direction and it should not be relied on in making a purchasing decision. The information mentioned regarding potential future products is not a commitment, promise, or legal obligation to deliver any material, code, or functionality. Information about potential future products may not be incorporated into any contract. The development, release, and timing of any future features or functionality described for our products remain at our sole discretion.

This document, REDP-5135-00, was created or updated on February 5, 2015.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. These and other IBM trademarked terms are marked on their first occurrence in this information with the appropriate symbol (or), indicating US registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at `http://www.ibm.com/legal/copytrade.shtml`

The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:
CICS®
DB2®
DS8000®
FICON®
GDPS®
Geographically Dispersed Parallel Sysplex™
HyperSwap®
IBM®
IMS™
Parallel Sysplex®
Redbooks®
Redbooks (logo) 
System z®
z/OS®

The following terms are trademarks of other companies:

Other company, product, or service names may be trademarks or service marks of others.