



Ilya Krutov

Enhancing Capabilities of the IBM eX5 Portfolio with the Latest Technologies

The IBM® eX5 product portfolio represents the fifth generation of servers built upon IBM Enterprise X-Architecture®. The servers in the eX5 product portfolio are based on Intel Xeon processor E7-8800/4800/2800 product families. These systems are the ideal choice for mission-critical environments supporting high performance databases, enterprise applications, and virtualized workloads.

The eX5 servers feature unique processor, memory, and I/O scaling capabilities, enabling them to handle the most resource demanding enterprise workloads. Built-in advanced reliability, availability, and serviceability (RAS) features allow IBM eX5 servers to run core mission-critical business services that require 24 hours per day, 7 days per week (24x7) availability.

In keeping with its leadership position of providing continuous technology improvement and innovation, IBM refreshes and enhances eX5 offerings by adding support for the most up-to-date technologies.

This IBM Redpaper™ publication describes enhancements in capabilities of the IBM eX5 portfolio with support for these latest hardware and software technologies.

The following technologies are covered:

- ▶ Microsoft Windows Server 2012
- ▶ VMware vSphere 5.1
- ▶ IBM eXFlash with the new Enterprise and Enterprise Value solid-state drives (SSDs)
- ▶ Workload-optimized models with the latest SSDs and High IOPS PCIe SSD adapters
- ▶ 16 Gb Fibre Channel connectivity
- ▶ IBM flash storage offerings

This publication is intended for IT professionals who are involved in the planning and design of the x86-based virtualization, database, and enterprise application solutions.

Executive summary

The IBM eX5 servers use fifth-generation IBM Enterprise X-Architecture, delivering innovation with enhanced reliability and availability features to enable optimal performance for databases, enterprise applications, and virtualized environments.

In keeping with its leadership position providing continuous technology improvement and innovation, IBM refreshes and enhances eX5 offerings by adding support for the most up-to-date technologies, including the following items:

- Enhancements for virtualized environments

Both Windows Server 2012 and vSphere 5.1 hypervisors double the amount of supported memory per physical server compared to their previous versions (Windows Server 2008 R2 and vSphere 4.1 respectively). As a result, the eX5 servers can potentially host twice more virtual machines (VMs) without incurring additional software licensing costs, providing memory was the limiting factor.

In addition, both virtualization environments expanded the amount of memory supported by a single VM to 1 TB, which allows to effectively utilize eX5 memory expansion capabilities to virtualize enterprise workloads that require large amounts of memory.

- Enhancements for databases

Windows Server 2012 supports up to 4 TB of memory allowing to improve overall database performance and scalability by utilizing up to 4 TB of memory in the 8-way x3850 X5 configuration and up to 3 TB in the 4-way x3850 X5 with MAX5 attached.

As of March 18, 2013, IBM System x3850 X5 achieved the best performance result ever on the TCP-E benchmark: 5,457.20 tpsE (transactions per second E) at \$249.58 USD / tpsE.

The x3850 X5 achieved this new level of online transaction processing (OLTP) performance using Microsoft SQL Server 2012 Enterprise Edition and Microsoft Windows Server 2012 Standard Edition. The x3850 X5 was configured with eight Intel Xeon E7-8870 processors at 2.40 GHz with 30 MB shared L3 cache per processor (8 processors/80 cores/160 threads), and 4096 GB of memory.

- IBM eXFlash enhancements

With the introduction of Enterprise Value and next generation Enterprise solid-state drives (SSDs), the IBM eXFlash increases storage I/O performance and capacity by almost two times, with the choice to optimize the cost per IOPS ratio depending on the workload type. These new SSD offerings significantly expand the capabilities of the eX5 servers in supporting resource-demanding enterprise workloads.

- Storage I/O enhancements

The recent acquisition by IBM of Texas Memory Systems® (TMS) extends IBM leadership in flash optimized storage infrastructures. IBM flash storage systems deliver up to 500,000 read IOPS with less than 100 microseconds latency, while providing up to 10 TB of usable data storage just in 1U of rack space.

IBM eX5 servers can leverage high performance capabilities of the IBM flash storage in OLTP and online analytical processing (OLAP) solutions.

With support for the 16 Gb Fibre Channel adapters, eX5 server can help to increase resource utilization and enhance I/O scaling capabilities in virtualized environments by increasing the speed of storage fabric and the number of available I/O slots.

Combined with the reliable high-speed solid-state drive technology and storage tiering, 16 Gb Fibre Channel fabric can help to significantly decrease storage I/O response time to match the processing power of the eX5 servers.

IBM eX5 portfolio overview

The four systems in the eX5 family are IBM System x3850 X5, x3950 X5, x3690 X5, and the IBM BladeCenter® HX5. The eX5 technology is primarily designed around three major workloads: database servers, server consolidation using virtualization services, and enterprise resource planning (application and database) servers. Each system can scale with more memory by adding an IBM MAX5 memory expansion unit to the server, and the x3850 X5, x3950 X5, and HX5 can also be scaled by connecting two servers together to form a single system.

Figure 1 shows the IBM eX5 family.



Figure 1 eX5 family (top to bottom): IBM BladeCenter HX5 (2-node), IBM System x3690 X5, and IBM System x3850 X5 (the IBM System x3950 X5 looks the same as the x3850 X5)

The IBM System x3850 X5 and x3950 X5 are 4U high rack-optimized servers. The x3850 X5 and the workload-optimized x3950 X5 are the new flagship servers of the IBM x86 server family. These systems are designed for maximum utilization, reliability, and performance for computer-intensive and memory-intensive workloads. These servers can be connected together to form a single system with twice the resources, or support memory scaling with the attachment of a MAX5. With the Intel Xeon processor E7 family, the x3850 X5 and x3950 X5 can scale to a two-server plus two-MAX5 configuration.

The IBM System x3690 X5 is a 2U rack-optimized server. This machine brings features and performance to the middle tier and a memory scalability option with MAX5.

The IBM BladeCenter HX5 is a single-wide (30 mm) blade server that follows the same design as all previous IBM blades. The HX5 brings unprecedented levels of capacity to high-density environments. The HX5 is expandable to form either a two-node system with four processors, or a single-node system with the MAX5 memory expansion blade.

When compared to other machines in the IBM System x® portfolio, these systems represent the upper end of the spectrum, are suited for the most demanding x86 tasks, and can handle jobs that previously might have run on other platforms. To assist with selecting the ideal system for a given workload, IBM designed workload-specific models for virtualization and database needs.

Table 1 gives an overview of the processor, memory, and internal storage scalability features of IBM eX5 systems.

Table 1 Scalability features for the eX5 systems

Maximum configurations		x3850 X5 or x3950 X5	x3690 X5	HX5
Processors	One-node	4	2	2
	Two-node	8	Not available	4
Memory	One-node	2 TB (64 DIMMs) ^a	1 TB (32 DIMMs) ^b	512 GB (16 DIMMs)
	One-node with MAX5	3 TB (96 DIMMs) ^a	2 TB (64 DIMMs) ^b	1.25 TB (40 DIMMs)
	Two-node	4 TB (128 DIMMs) ^a	Not available	1 TB (32 DIMMs)
	Two-node with MAX5	6 TB (192 DIMMs) ^a	Not available	Not available
1.8-inch SSDs	One-node	16	24	2
	Two-node	32	Not available	4

a. Requires full processors to install and use all memory.

b. Requires that the memory mezzanine board is installed along with processor 2.

Latest benchmarks

With the introduction of Microsoft Windows Server 2012, IBM eX5 systems can now use even more memory than before (up to 4 TB) to increase the performance of memory-intensive workloads such as transactional databases.

On November 28, 2012, IBM published a world-record performance result for the 4-socket x3850 X5 server running TPC-E benchmark¹. The TPC-E benchmark is designed to enable clients to more objectively measure and compare the performance and price of various OLTP systems.

The IBM System x3850 X5 server achieved 3,218.46 tpsE (transactions per second E) at \$225.30 USD / tpsE. The x3850 X5 achieved this new level of OLTP performance using Microsoft SQL Server 2012 Enterprise Edition and Microsoft Windows Server 2012 Standard Edition. The x3850 X5 was configured with four Intel Xeon E7-4870 processors at 2.40 GHz with 30 MB shared L3 cache per processor (4 processors/40 cores/80 threads), and 2048 GB of memory.

On March 18, 2013, IBM published the best performance result ever achieved on TCP-E benchmark². The IBM System x3850 X5 server achieved 5,457.20 tpsE (transactions per second E) at \$249.58 USD / tpsE. The x3850 X5 achieved this new level of OLTP performance using Microsoft SQL Server 2012 Enterprise Edition and Microsoft Windows Server 2012 Standard Edition. The x3850 X5 was configured with eight Intel Xeon E7-8870 processors at 2.40 GHz with 30 MB shared L3 cache per processor (8 processors/80 cores/160 threads), and 4096 GB of memory.

Based on the TPC-E results, the x3850 X5 provides almost 70% increase in performance when scaling from four to eight sockets under TPC-E transactional workload.

¹ http://public.dhe.ibm.com/eserver/benchmarks/news/newsblurb_x3850X5_tpce_20121128.pdf

² http://public.dhe.ibm.com/eserver/benchmarks/news/newsblurb_x3850X5_8P_tpce_20130318.pdf

Microsoft Windows Server 2012

The following Windows Server 2012 benefits affect the daily IT operations of a company:

- ▶ **Simplified management:** Windows Server 2012 offers an integrated, highly available, and easy to manage multi-server platform with breakthrough efficiency and ubiquitous automation.
- ▶ **Affordable storage:** Reduce storage costs by using SMB based file storage instead of SAN storage. For pooled virtual machine deployment, storage requirements can be lowered for a set of users by using a combination of SMB and local DAS.
- ▶ **Continuous availability:** Cost-effective, highest IT service uptime is designed to protect against a wide range of failure scenarios and help both prevent downtime and speed recovery when needed.

Among many enhancements delivered by the Microsoft Windows Server 2012 family of operating systems, one of the key improvements is scalability for both traditional and virtualized environments.

Table 2 summarizes key scalability features of the Windows Server family.

Table 2 Windows Server family scalability features

Feature	Windows Server 2008 R2	Windows Server 2012	Improvement ratio
Logical processors	256	320	1.25x
Physical memory	2 TB	4 TB	2x

By doubling the amount of supported memory, Windows Server 2012 combined with the 8-socket x3850 X5 (up to 4 TB of memory) or 4-socket x3850 X5 with MAX5 (up to 3 TB of memory) allows to dramatically increase the performance of the enterprise workloads while lowering total cost of ownership. The most recent example of the outcome from scalability improvement is the world record set by x3850 X5 on TPC-E benchmark (see “Latest benchmarks” on page 4 for more information).

Table 3 summarizes key scalability features of the Windows Server Hyper-V family for virtualized environments.

Table 3 Windows Server Hyper-V family scalability features

Feature	Windows Server 2008 R2 Hyper V	Windows Server 2012 Hyper V	Improvement ratio
Host			
Logical processors	256	320	1.25x
Physical memory	2 TB	4 TB	2x
vCPU per host	512	2,048	4x
Virtual machine (VM)			
vCPU per VM	4	64	16x
Memory per VM	64 GB	1 TB	16x
Active VMs per host	384	1,024	2.7x

Based on the numbers described in Table 3 on page 5, you can potentially deploy almost three times more virtual machines on a single x3850 X5 if vCPU was the limiting factor.

For memory demanding virtual machines, these numbers mean that you can potentially deploy twice more virtual machines on 8-way x3850 X5, and 50% more VMs on 4-way x3850 X5 with MAX5. In addition, you can configure VMs with more memory and higher number of vCPUs, allowing to consolidate higher capacity workloads.

For both traditional and virtualized environments, x3850 X5 servers, MAX5 expansion units, and Windows Server 2012 family, can help to lower overall total cost of ownership by increasing VM density and providing savings on the following items:

- ▶ Server hardware: fewer servers to acquire, deploy, and manage
- ▶ Software: fewer sockets or cores requires lower software licensing costs
- ▶ Operations: fewer servers consumes less power and generates less heat
- ▶ Management: fewer servers occupies less space and requires less efforts to deploy and support

VMware vSphere 5.1

VMware vSphere 5.1 offerings, paired with eX5 server, can help you with the following tasks and goals:

- ▶ Consolidate operations, reduce or eliminate redundancy in infrastructure and personnel, and increase overall systems control.
- ▶ Improve service levels, decrease both planned and unplanned downtime, and improve employee productivity.
- ▶ Respond to new business opportunities quickly and process more information in real-time.
- ▶ Build a virtualization solution that will allow you to adjust and transform your business by establishing a solid foundation for growth.

Table 3 summarizes key enhancements in scalability features of the VMware vSphere family of hypervisors.

Table 4 VMware vSphere family scalability features

Feature	vSphere 4.1	vSphere 5.1	Improvement ratio
Host			
Logical processors	160	160	-
Physical memory	1 TB	2 TB	2x
vCPU per host	512	2,048	4x
Virtual machine (VM)			
vCPU per VM	8	64	8x
Memory per VM	255 GB	1 TB	4x
Active VMs per host	320	512	1.6x

The key scalability improvement is doubling the amount of supported memory per physical host. For memory-demanding virtual machines, that means that you can potentially deploy twice more virtual machines x3850 X5 or x3690 X5. In addition, you can configure VMs with more memory and higher number of vCPUs, allowing to consolidate higher capacity workloads.

Also, you can potentially deploy almost two times more virtual machines on a single x3850 X5 or x3690 X5 if vCPU was the limiting factor.

For virtualized environments, x3850 X5 servers together with VMware vSphere 5.1, can help to lower overall total cost of ownership by increasing VM density and providing savings on the following items:

- ▶ Server hardware: fewer servers to acquire, deploy, and manage
- ▶ Software: fewer sockets or cores requires lower software licensing costs
- ▶ Operations: fewer servers consumes less power and generates less heat
- ▶ Management: fewer servers occupies less space and requires less efforts to deploy and support

IBM eXFlash

IBM eXFlash technology is a server-based high performance internal storage solution, which is based on SSDs and performance-optimized disk controllers (both RAID and non-RAID).

Single eXFlash unit accommodates up to eight hot-swap SSDs, and can be connected to up to two performance-optimized controllers. eXFlash is supported on IBM System x3690 X5, x3850 X5, and x3950 X5 Intel Xeon processor E7 family-based servers, and also Intel Xeon processor E5 family-based x3650 M4 and x3750 M4 servers.

Figure 2 shows an eXFlash unit, with the status lights assembly on the left side.




Figure 2 IBM eXFlash unit

IBM eXFlash requires the following components:

- ▶ IBM eXFlash hot-swap SAS SSD backplane
- ▶ 1.8-inch solid-state drives (SSDs)
- ▶ IBM disk controllers

Each eXFlash unit can accommodate up to eight 1.8-inch hot-swap front-accessible SSDs, and occupies four 2.5-inch SAS hard disk drive bays. You can install the following number of eXFlash units:

- ▶ The x3850 X5 can have up to sixteen 1.8-inch SSDs with up to two eXFlash units.
- ▶ The x3690 X5 can have up to twenty four 1.8-inch SSDs with up to three eXFlash units.
- ▶ The x3650 M4 and x3750 M4 can have up to thirty two 1.8-inch SSDs with up to four eXFlash units.

With the introduction of Enterprise Value and next generation Enterprise SSDs, the IBM eXFlash increases storage I/O performance and capacity by almost two times, with the choice to optimize the cost per IOPS ratio depending on the workload type. These new SSD offerings significantly expand the capabilities of the eX5 servers in supporting resource-demanding enterprise workloads.

Enterprise and Enterprise Value SSDs

Two types of SSDs can be used with IBM eXFlash: *Enterprise SSDs* and *Enterprise Value SSDs*. While both Enterprise SSDs and Enterprise Value SSDs typically have similar read IOPS characteristics, the key difference between them is write IOPS performance and their endurance (or life expectancy).

SSDs have a huge but finite number of program/erase (P/E) cycles, which affects how long they can perform write operations and thus their life expectancy. Enterprise SSDs have significantly better endurance but higher cost/IOPS ratio compared to Enterprise Value SSDs. SSD write endurance is typically measured by the number of P/E cycles that the drive can incur over its lifetime, listed as Total Bytes Written (TBW) in the device specification.

Because of such behavior by Enterprise Value solid-state drives, careful planning must be done to use them only in read-intensive environments to ensure that the TBW of the drive will not be exceeded prior to the required life expectancy.

Therefore, for optimal cost per IOPS ratio, use Enterprise SSDs for write-intensive workloads and use Enterprise Value SSDs for read-intensive workloads, as shown in Table 5.

Table 5 Application workload by SSD type

SSD type → ↓ Application type	Enterprise	Enterprise Value	Good for eXFlash
OLTP Database	Yes		Yes
Data warehouse		Yes	Yes
E-mail server	Yes		Yes
Medical imaging		Yes	Yes
Video on demand		Yes	Yes
Web/Internet		Yes	Yes
Web 2.0	Yes		Yes

Typical IBM eXFlash usage scenarios include the following examples:

- ▶ High-speed read cache in a local or SAN-based storage environment
- ▶ Temporary local storage space for mid-tier applications and databases
- ▶ Main (Tier 0) local data storage in a single server environments or in a distributed scale-out environment with local-only storage or mixed local and SAN-based storage

Table 6 summarizes key specifications of the new Enterprise and Enterprise Value SSDs compared to the first and second generation (legacy) eXFlash SSDs.

Table 6 Enterprise and Enterprise Value SSD specifications

Specification	IBM Enterprise SSD (legacy)	IBM Enterprise SSD (next generation)	IBM Enterprise Value SSD
Maximum capacity	200 GB	400 GB	512 GB
Interface	3 Gb SATA	6 Gb SATA	6 Gb SATA
IOPS read ^a	30,000	60,000	50,000
IOPS write ^a	20,000	40,000	7,500
Sequential read rate ^b	250 MBps	520 MBps	350 MBps
Sequential write rate ^b	250 MBps	500 MBps	140 MBps

a. 4 KB block transfers

b. 128 KB block transfers

New SSDs in the eXFlash offer the following advantages compared to legacy SSDs:

- ▶ Twice more random and sequential I/O performance for read-intensive and write-intensive workloads with IBM Enterprise SSDs
- ▶ Almost 50% more random and sequential I/O performance for read-intensive workloads with IBM Enterprise Value SSDs
- ▶ Twice more capacity
- ▶ Twice more interface bandwidth
- ▶ Optimized cost per IOPS ratio with the choice of Enterprise and Enterprise Value SSDs

IBM eXFlash benefits

In theory, the OLTP I/O performance of the single eXFlash unit combined with the IBM SATA Enterprise SSDs configured in a RAID 5 array is equivalent to the storage system consisting of more than 1000 traditional spinning HDDs. Besides HDDs themselves, building such a massive I/O-intensive high-performance storage system requires external deployment with many additional infrastructure components including host bus adapters (HBAs), switches, storage controllers, disk expansion enclosures, and cables. Consequently, this leads to more capital expenses, floor space, electrical power requirements, and operation and support costs. Because eXFlash is based on internal server storage, it does not require all those components discussed previously, and helps to eliminate additional expenses and environmental requirements.

In summary, IBM eXFlash solution can provide the following benefits:

- ▶ Significantly lower implementation cost of high performance I/O-intensive storage systems with much better cost per IOPS ratio
- ▶ Significantly higher performance of I/O-intensive applications like databases with up to ten times less response time
- ▶ Significant savings in power and cooling with high performance per watt ratio
- ▶ Significant savings in floor space with extreme performance per U ratio
- ▶ Simplified management and maintenance with internal server-based configurations (no external power and information infrastructure needed)

IBM eXFlash is optimized for a heavy mix of random read and write operations, such as transaction processing, data mining, business intelligence and decision support, and other random I/O-intensive applications. In addition to its superior performance, eXFlash offers superior uptime with three times the reliability of mechanical disk drives. SSDs have no moving parts to fail. They also use Enterprise Wear-Leveling to extend their use even longer.

In environments where RAID protection is required (that is, eXFlash is used as a master data storage), use a RAID controller with Performance Accelerator key enabled to ensure the peak IOPS can be reached.

The main advantage of ServeRAID M5014, M5015, and M5110 with Performance Key controllers for SSDs is a Cut Through I/O (CTIO) feature enabled. CTIO optimizes highly random read and write I/O operations for small data blocks to support the high IOPS capabilities of SSD drives and to ensure the fastest response time to the application. For example, by enabling CTIO on a RAID controller with SSDs up to two times more IOPS can be achieved compared to the controller with CTIO feature disabled.

Note: A single eXFlash unit requires a dedicated controller (or two controllers). When used with eXFlash, these controllers cannot be connected to the HDD backplanes.

In a non-RAID environment where eXFlash can be used as a high-speed read cache, use the IBM 6 Gb Performance Optimized HBA to ensure maximum random I/O read performance is achieved. Only one 6 Gb SSD HBA is supported per single SSD backplane.

It is possible to mix RAID and non-RAID environments, however, the maximum number of disk controllers that can be used with all SSD backplanes in a single system is four.

Workload-optimized models

The IBM System x3950 X5 models offer the same core features as the x3850 X5, and they are workload-optimized to deliver performance and rapid deployment of database applications, virtualized environments or SAP In-Memory Appliance (SAP HANA), depending on your business requirements.

The most recent additions to the workload optimized portfolio are as follows:

- ▶ Database-optimized models featuring IBM eXFlash with next-generation Enterprise SSDs installed
- ▶ SAP HANA models featuring second generation of High IOPS PCIe SSD adapters

Database-optimized models

In February 2013, IBM announced new IBM System x3950 X5 Workload Optimized System for Database models. These models enhance the server line by providing new levels of performance, and price for performance. These refreshed models use the capabilities of the next generation IBM Enterprise SSDs in the IBM eXFlash to double the performance of local storage I/O.

Table 7 summarizes key specifications for the newly announced database-optimized models.

Table 7 Workload-optimized systems for databases

Machine type model	Processor	Memory	Drive controller	Drives
7143-D3x	4x E7-4860 10C 2.26 GHz 24 MB 1066 MHz	32x 4 GB	2x 6 Gb Performance Optimized HBA	16x 200 GB SSD
7143-D4x	4x E7-4860 10C 2.26 GHz 24 MB 1066 MHz	32x 4 GB	4x ServeRAID M5015 with Performance Accelerator feature enabled	16x 200 GB SSD

SAP HANA optimized models

In October of 2012, IBM announced a new set of IBM System x workload-optimized models for SAP HANA, updating some of the components with newer generation versions. These next-generation components increase the performance and capacity of flash-based I/O subsystem. Table 8 shows all building blocks announced in 2012 and their features.

Table 8 IBM System x workload-optimized models for SAP HANA, 2012 models

Building block	Server (MTM)	CPUs	Main memory	Log storage	Data storage	Preload
XS	x3690 X5 (7147-HAx ^a)	2x Intel Xeon E7-2870	128 GB DDR3 (8x 16 GB)	10x 200 GB 1.8" MLC SSD (combined log and data)		Yes
S	x3690 X5 (7147-HBx)	2x Intel Xeon E7-2870	256 GB DDR3 (16x 16 GB)	10x 200 GB 1.8" MLC SSD (combined log and data)		Yes
S+	x3950 X5 (7143-HAx)	2x Intel Xeon E7-8870	256 GB DDR3 (16x 16 GB)	1.2 TB High IOPS adapter	8x 900 GB 10 K SAS HDD	Yes
M	x3950 X5 (7143-HBx)	4x Intel Xeon E7-8870	512 GB DDR3 (32x 16 GB)	1.2 TB High IOPS adapter	8x 900 GB 10 K SAS HDD	Yes
L Option	x3950 X5 (7143-HCx)	4x Intel Xeon E7-8870	512 GB DDR3 (32x 16 GB)	1.2 TB High IOPS adapter	8x 900 GB 10 K SAS HDD	No

a. The letter "x" equals the country-specific letter (for example, EMEA MTM is 7147-HAG, and the US MTM is 7147-HAU). Contact your IBM representative for regional part numbers.

All models (except for 7143-HCx) are preloaded with SUSE Linux Enterprise Server for SAP Applications (SLES for SAP) 11 SP1, IBM GPFS™, and the SAP HANA software stack. Licenses and maintenance fees (for three years) for SLES for SAP and GPFS are included. The licenses for the SAP software components must be acquired separately from SAP.

The L-Option building blocks (7143-HCx) are intended as an extension to an M building block (7143-HBx). When building an L-Size SAP HANA system, one M building block has to be combined with an L-Option building block, leveraging eX5 scalability. Both systems then act as one single eight-socket, 1 TB server. Therefore, the L-Option building blocks do not require a software preload, L-Option building block includes however the required additional software licenses for GPFS and SLES for SAP.

The building blocks are configured to match the SAP HANA sizing requirements. The main memory sizes match the number of CPUs, to give the correct balance between processing power and data volume. Also, the storage devices in the systems provide the storage capacity required to match the amount of main memory.

All systems include storage for both the data volume and the log volume. Savepoints are stored on a RAID protected array of 10 K SAS hard drives, optimized for data throughput. The SAP HANA database logs are stored on flash technology storage devices:

- ▶ RAID-protected, hot swap eXFlash SSD drives on the models based on IBM System x3690 X5
- ▶ Flash-based High IOPS PCIe adapters for the models based on IBM System x3950 X5

These flash technology storage devices are optimized for high IOPS performance and low latency to provide the SAP HANA database with a log storage that allows the highest possible performance. Because a transaction in the SAP HANA database can only return after the corresponding log entry is written to the log storage, high IOPS performance and low latency are key to database performance.

The building blocks based on the IBM System x3690 X5 include combined data and log storage on an array of RAID-protected, hot-swap eXFlash SSD drives. Optimized for throughput, high IOPS performance, and low latency, these building blocks give extra flexibility when dealing with large amounts of log data, savepoint data, or backup data.

The IBM High IOPS PCIe Adapters combine high IOPS performance with low latency. As an example, with 512 KB block random reads, the IBM 1.2TB High IOPS MLC Mono Adapter can deliver 143,000 IOPS, compared with 420 IOPS for a 15 K RPM 146 GB disk drive. The read access latency is about 68 microseconds, which is one hundredth of the latency of a 15 K RPM 146 GB disk drive (about 5 ms or 5000 microseconds). The write access latency is even less, with about 15 microseconds.

16 Gb Fibre Channel connectivity

With the virtualization trends in data centers, enterprise workloads demand higher I/O bandwidth to match the capabilities of IBM eX5 server's multi-core processors and increased amounts of memory, allowing the higher number of virtual machines (VMs) to be hosted on a single physical system. Higher I/O bandwidth, including storage I/O, can help to achieve better server utilization and higher VM per server ratio.

Data warehouses and business analytics are additional examples of the workload that requires higher storage I/O bandwidth to allow faster data processing, making strategic business decisions in a timely manner.

With support for the 16 Gb Fibre Channel adapters, x3850 X5 can help to increase the utilization and enhance I/O scaling capabilities in virtualized environments by increasing the speed of storage fabric and the number of available I/O slots. Combined with the reliable high-speed solid-state drive technology and storage tiering, 16 Gb Fibre Channel fabric can help to significantly decrease storage I/O response time to match the processing power of the

x3850 X5 server. In general, 16 Gb Fibre Channel means twice higher bandwidth compared to the 8 Gb Fibre Channel.

The use of 16 Gb Fibre Channel connectivity can help to achieve the following benefits:

- ▶ Higher performance:
 - Higher IOPS
 - Higher throughput
 - Lower latency
- ▶ Infrastructure simplification:
 - Simplified deployment and management
 - Higher VM density and lower number of physical systems
- ▶ Improved TCO:
 - Reduced acquisition costs
 - Reduced power and cooling costs
 - Reduced support and maintenance costs

Performance

The maximum theoretical throughput of 16 Gb Fibre Channel is 1.6 GBps, while 8 Gb Fibre Channel has a throughput of 0.8 GBps. This fact means that throughput-intensive applications such as OLAP databases can achieve better response time.

Also, 16 Gb Fibre Channel can potentially be beneficial for IOPS-intensive applications:

- ▶ The 16 Gb infrastructure can deliver up to 200,000 IOPS per port using 8 KB I/O blocks.
- ▶ The 8 Gb can deliver only 100,000 IOPS per port.

Infrastructure simplification

Because the performance of a single 16 Gb Fibre Channel port is equivalent to the performance of two 8 Gb Fibre Channel ports, you can reduce the number of ports, cables, and SFP+ modules while keeping the same performance levels.

Also, if 8 Gb Fibre Channel is a limiting factor for overall server performance, you can potentially support heavier workloads on a same system by moving to 16 Gb, such as increasing VM density for I/O intensive workloads while lowering overall number of physical systems.

Improved TCO

With the lower number of systems and components used to build the infrastructure, you can potentially reduce acquisition costs and operational, management, and support costs because of simplified infrastructure.

Two possible deployment scenarios where eX5 servers with 16 Gb Fibre Channel connectivity can provide certain benefits are as follows:

- ▶ Data warehouses
- ▶ Mixed virtualized workloads

IBM flash storage offerings

IBM eX5 servers can leverage I/O performance of the recently announced IBM FlashSystem™ flash storage, balancing the computing power with the storage I/O capabilities.

IBM FlashSystem delivers extraordinary business value with extreme performance, efficiency, and reliability through easy-to-deploy all-flash systems. The recent acquisition by IBM of Texas Memory Systems (TMS) extends IBM leadership in flash optimized storage infrastructures. By using patented Variable Stripe RAID technology and 2D Flash RAID, IBM flash storage capability enhances system resiliency without sacrificing performance or usable capacity. For a truly self-optimized and automated storage infrastructure, you can combine FlashSystem with IBM System Storage® SAN Volume Controller virtualization and IBM System Storage Easy Tier®, intelligent data placement software. The result is further optimization of workloads, improved performance and service levels, and faster response to changing business requirements.

IBM FlashSystem transforms the data center environment and enhances performance and resource consolidation to gain the most from business processes and critical applications. Examples of such processes and applications include online transaction processing (OLTP), business intelligence (BI), online analytical processing (OLAP), virtual desktop infrastructures, high-performance computing, and content delivery solutions (such as cloud storage and video on demand).

IBM FlashSystem storage systems deliver over 500,000 read IOPS with less than 100 microseconds latency, while they provide up to 20 TB of usable data storage just in 1U of rack space. IBM flash storage systems have enterprise-level availability and reliability with no single point of failure, multiple layers of data correction, chip redundancy, and redundant hot swap components.

IBM eX5 servers can leverage these unique performance, capacity, density, availability, and efficiency features of the IBM flash storage to achieve the following benefits:

- ▶ Dramatically boost application performance and lower cost per IOPS ratio.
- ▶ Increase user productivity with better response times, improving business efficiency.
- ▶ Increase data availability by using advanced system-level high availability and reliability technologies, reducing the number of solution components and shortening batch processing and backup times.
- ▶ Increase storage performance and capacity while decreasing power, cooling, and space requirements.
- ▶ Reduce TCO:
 - Reduce energy costs because of lower power and cooling requirements.
 - Reduce the number of systems, devices, and components that are required to build the solution by increasing usage of available resources.
 - Reduce software license fees because fewer systems or processors are required.
 - Reduce management and support costs because of fewer components to deploy and support.
 - Offer faster ROI because of better server resource usage.

As one of the examples, consider the following deployment scenario. OLTP workloads are characterized by small, interactive transactions that generally require subsecond response times. The key performance indicator (KPI) of the transactional system is latency, because

the user expects to receive the requested product information or to place an order quickly. Inability to meet these user expectations leads to customer dissatisfaction and revenue loss. IBM flash storage addresses these challenges by providing low latency, extreme performance, and efficient transaction management.

The planning of information infrastructure includes choosing the most cost-effective way to fulfill the application requirements for storage access with respect to speed, capacity, and availability. To describe these requirements, and to establish the framework for the deployment of the storage infrastructure, the storage tiering approach was established.

Each storage tier defines a set of characteristics to meet the application requirements. There are four tiers, each with performance, availability, capacity, and access pattern characteristics for the data residing on that tier. Knowing your data access requirements will help you to place data on the appropriate storage tier, thereby ensuring that your storage infrastructure is capable of running your workloads in a cost-efficient manner.

Sometimes identifying data sets that will be used most frequently to place them onto an appropriate storage tier is difficult. In this case, automated storage tiering can help to identify most frequently used sets of data (frequently referred to as hot data). It provides continuous I/O monitoring and dynamic placement onto a performance optimized storage (focus on cost per IOPS). Automated storage tiering keeps infrequently used data (known as cold data) on a capacity optimized storage (focus on cost per GB).

IBM flash storage systems can be combined with System Storage SAN Volume Controller and its integrated IBM System Storage Easy Tier feature to provide automated storage tiering and, therefore, dynamically optimize storage performance as illustrated in Figure 3.

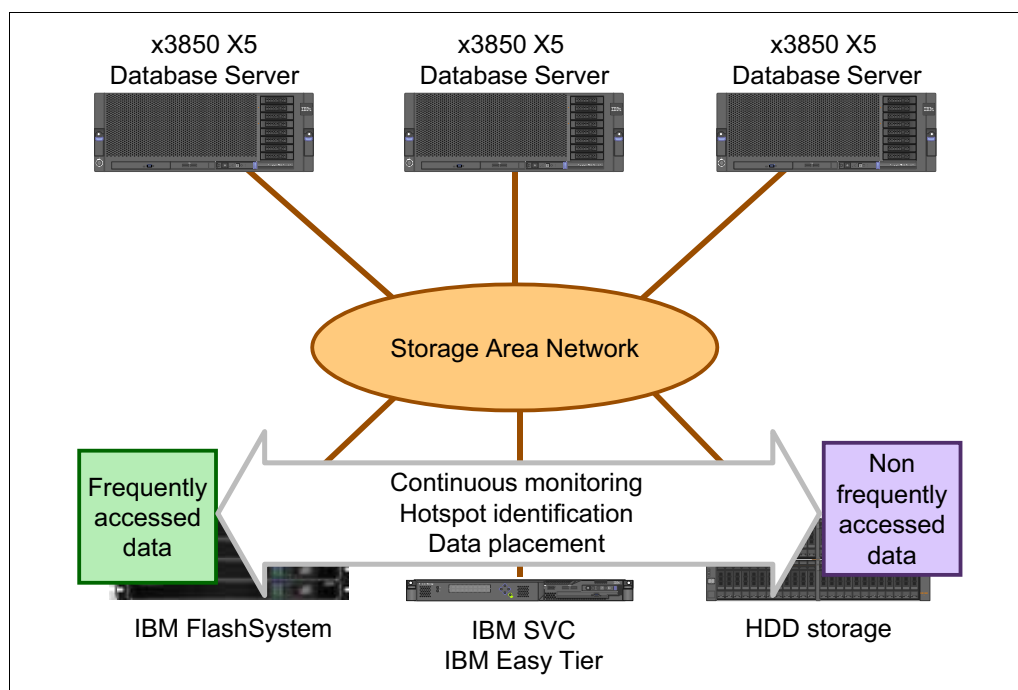


Figure 3 IBM flash storage systems with SAN Volume Controller and IBM Easy Tier

IBM SAN Volume Controller provides storage virtualization through pooled management of diverse storage environments by using a simple, common interface that remains consistent regardless of storage type. SAN Volume Controller hides physical changes within the storage infrastructure to help improve availability and to streamline the task of provisioning, making dynamic management of storage assets far faster and more efficient than manual processes.

To further simplify management, SAN Volume Controller integrates with IBM Tivoli® Storage Productivity Center to create a single view of physical and virtual storage resources and their relationships with physical and virtual servers.

IBM System Storage Easy Tier is a function that automatically and nondisruptively moves frequently accessed data from hard disk drives to solid-state storage, placing such data in a faster tier of storage. System Storage Easy Tier eliminates manual intervention when assigning highly active data on volumes to faster responding storage. In this dynamically tiered environment, data movement is seamless to the host application regardless of the storage tier in which the data resides.

Optionally, SAN Volume Controller supports the IBM Real-time Compression™ feature, which performs inline data compression without performance impact, helping to increase usable storage space up to five times. Unlike other approaches to compression, IBM Real-time Compression can be used with active primary data, such as production databases and email applications, which dramatically expands the range of candidate data that can benefit from compression. IBM Real-time Compression operates immediately as data is written to disk so that no space is wasted by storing uncompressed data that is awaiting post-processing.

Conclusion

This paper discussed key technology enhancements that expand capabilities of the IBM eX5 servers in mission critical database, enterprise application, and virtualized environments.

Through ongoing technology innovations and improvements, IBM eX5 servers continue to provide leadership capabilities in scalability and performance, reliability, availability, and serviceability, and manageability and operations:

- ▶ IBM x3850 X5 delivered world record TPC-E benchmark result, and best in class 4-socket TPC-E performance result.
- ▶ IBM x3850 X5 in a 8-socket configuration can potentially host twice more virtual machines, and 4-socket with MAX5 configuration can host up to 50% more VMs, without a need to buy additional servers and software licenses.
- ▶ With Enterprise Value and next generation Enterprise SSDs, the IBM eXFlash increases storage I/O performance and capacity by almost two times, with the choice to optimize the cost per IOPS ratio depending on the workload type.
- ▶ IBM flash storage systems deliver up to 500,000 read IOPS with less than 100 microsecond latency, while providing up to 10 TB of usable data storage in only 1U of rack space. IBM eX5 servers can leverage high performance capabilities of the IBM flash storage in OLTP and OLAP solutions.
- ▶ With support for the 16 Gb Fibre Channel adapters, eX5 server can help to increase resource utilization and enhance I/O scaling capabilities in virtualized environments by increasing the speed of storage fabric and the number of available I/O slots. Combined with the reliable high-speed solid-state drive technology and storage tiering, 16 Gb Fibre Channel fabric can help to significantly decrease storage I/O response time to match the processing power of the eX5 servers.

Related publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this paper.

- ▶ *IBM eX5 Portfolio Overview: IBM System x3850 X5, x3950 X5, x3690 X5, and BladeCenter HX5*, REDP-4650
<http://www.redbooks.ibm.com/abstracts/redp4650.html>
- ▶ *Reliability, Availability, and Serviceability Features of the IBM eX5 Portfolio*, REDP-4864
<http://www.redbooks.ibm.com/abstracts/redp4864.html>
- ▶ *In-memory Computing with SAP HANA on IBM eX5 Systems*, SG24-8086
<http://www.redbooks.ibm.com/abstracts/sg248086.html>
- ▶ *IBM SATA 1.8-inch and 2.5-inch MLC Enterprise SSDs for IBM System x*, TIPS0908
<http://www.redbooks.ibm.com/abstracts/tips0908.html>
- ▶ *IBM SATA 1.8-inch and 2.5-inch MLC Enterprise Value SSDs*, TIPS0879
<http://www.redbooks.ibm.com/abstracts/tips0879.html>
- ▶ *IBM FlashSystem in OLTP Database Environments*, TIPS0973
<http://www.redbooks.ibm.com/abstracts/tips0973.html>
- ▶ IBM System x performance benchmarks
<http://www.ibm.com/systems/x/resources/benchmarks/results>

Authors

This paper was produced by a team of specialists from around the world working at the International Technical Support Organization (ITSO), Raleigh Center.

Ilya Krutov is a Project Leader at the ITSO Center in Raleigh and has been with IBM since 1998. Before he joined the ITSO, Ilya served in IBM as a Run Rate Team Leader, Portfolio Manager, Brand Manager, Technical Sales Specialist, and Certified Instructor. Ilya has expertise in IBM System x, BladeCenter, and Flex System products; server operating systems; and networking solutions. He has authored over 120 books, papers, Product Guides, and Solution Guides. He has a Bachelor's degree in computer engineering from the Moscow Engineering and Physics Institute.

Thanks to the following people for their contributions to this project:

Kevin Barnes, Tamikia Barrow, Mary Comianos, Shari Deiana, Cheryl Gera, Linda Robinson, David Watts

International Technical Support Organization, Raleigh Center

Randall Lundin
IBM

Now you can become a published author, too!

Here's an opportunity to spotlight your skills, grow your career, and become a published author—all at the same time! Join an ITSO residency project and help write a book in your area of expertise, while honing your experience using leading-edge technologies. Your efforts will help to increase product acceptance and customer satisfaction, as you expand your network of technical contacts and relationships. Residencies run from two to six weeks in length, and you can participate either in person or as a remote resident working from your home base.

Find out more about the residency program, browse the residency index, and apply online at:

ibm.com/redbooks/residencies.html

Stay connected to IBM Redbooks

- ▶ Find us on Facebook:
<http://www.facebook.com/IBMRedbooks>
- ▶ Follow us on Twitter:
<http://twitter.com/ibmredbooks>
- ▶ Look for us on LinkedIn:
<http://www.linkedin.com/groups?home=&gid=2130806>
- ▶ Explore new IBM Redbooks® publications, residencies, and workshops with the IBM Redbooks weekly newsletter:
<https://www.redbooks.ibm.com/Redbooks.nsf/subscribe?OpenForm>
- ▶ Stay current on recent Redbooks publications with RSS Feeds:
<http://www.redbooks.ibm.com/rss.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 grant 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 websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites 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.

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.

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.

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.

© Copyright International Business Machines Corporation 2013. All rights reserved.

Note to U.S. Government Users Restricted Rights -- Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

This document REDP-5002-00 was created or updated on May 19, 2013.

Send us your comments in one of the following ways:

- ▶ Use the online **Contact us** review Redbooks form found at:
ibm.com/redbooks
- ▶ Send your comments in an email to:
redbooks@us.ibm.com
- ▶ Mail your comments to:
IBM Corporation, International Technical Support Organization
Dept. HYTD Mail Station P099
2455 South Road
Poughkeepsie, NY 12601-5400 U.S.A.



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:

BladeCenter®
Easy Tier®
FlashSystem™
GPFS™
IBM®

IBM FlashSystem™
Real-time Compression™
Redbooks®
Redpaper™
Redbooks (logo) ®

System Storage®
System x®
Tivoli®
X-Architecture®

The following terms are trademarks of other companies:

Texas Memory Systems, and the Texas Memory Systems logo are trademarks or registered trademarks of Texas Memory Systems, an IBM Company.

Intel, Intel Xeon, Intel logo, Intel Inside logo, and Intel Centrino logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

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