Ilya Krutov

Workload Optimization with IBM X6 Servers

The increasing demand for cloud computing and business analytical workloads by enterprises to meet business needs drives innovation to find new ways to build informational systems. Clients are looking for cost-optimized fit-for-purpose IT solutions that manage large amounts of data, easily scale performance, and provide reliable real-time access to actionable information.

Built on decades of innovation, IBM® introduces its sixth generation of Enterprise X-Architecture® technology, IBM X6 servers. IBM X6 servers are designed to be *fast*, *agile*, and *resilient*:

- Fast application performance means immediate access to actionable information.
- Agile system design helps to reduce acquisition costs and provide the ability to host multiple generations of technology in a single server.
- Resilient platforms maximize application uptime and promote easy integration in virtual environments.

IBM X6 servers continue to lead the way as the shift toward mission-critical scalable databases, business analytics, virtualization, enterprise applications, and cloud applications accelerates.

This IBM Redpaper™ document describes IBM X6 technology innovations that help address challenges that clients are facing in the enterprise environments, by delivering workload-optimized systems for their applications. The following topics are covered:

This paper is intended for IT professionals who want to understand and learn more about workload optimization with IBM X6 technologies.

Introduction to the IBM X6 portfolio

The IBM X6 product portfolio represents the sixth generation of servers that are built upon Enterprise X-Architecture. Enterprise X-Architecture is the culmination of generations of IBM technology and innovation that is derived from our experience in high-end enterprise servers. Now, with X6, IBM scalable systems can be expanded on demand and configured by using a building block approach that optimizes system design for your workload requirements.

These servers scale to more CPU cores, memory, and I/O than previous systems, enabling them to handle greater workloads than the systems that they supersede. Power efficiency and server density are optimized, making them affordable to own and operate.

IBM X6 systems allow your enterprise to grow in processing, input/output (I/O), and memory dimensions. Therefore, you can provision what you need now and expand the system to meet future requirements. System redundancy and availability technologies are more advanced than those previously available in the x86 systems.

Fast application performance

The x3850 X6 and x3950 X6 servers deliver fast application performance thanks to an innovative scalable design and new storage technology that is designed to optimize overall solution performance. The X6 servers are the first servers designed and optimized for IBM eXFlash memory-channel storage. With eXFlash memory-channel storage, they can deliver up to 12.8 TB of ultra-low latency flash storage. With the new Intel Xeon processor E7-4800 v2 and E7-8800 v2 product families, the x3850 X6 and x3950 X6 servers can deliver up to 6.0 TB or 12 TB of memory and 60 or 120 cores of processing power, respectively. Armed with these capabilities, you can host essential mission-critical applications, implement large virtual machines, or run sizeable in-memory databases without compromises in performance, capacity, or scalability.

These enterprise-class servers use unique IBM eXFlash memory-channel storage to deliver an exceptional level of performance and value. eXFlash memory-channel storage offers significantly lower write latency - as low as 5 µs. As more eXFlash dual inline memory modules (DIMMs) are added, input/output operations per second (IOPS) increases, yet latency does not increase.

IBM eXFlash DIMMs feature exclusive features called WriteNow and RAID 1 to substantially get higher performance and enterprise-class reliability.

With IBM eXFlash memory-channel storage, you gain consistent performance even if you are running mixed workloads. The on-memory bus design alleviates potential I/O contention. In addition, databases from 200 GB to 12.8 TB have deterministic response times and consistent performance.

The eXFlash memory-channel storage supports several times more virtual machines per server without degradation of service and is ideal for large databases and highly virtualized systems. They represent a highly scalable form factor that provides greater performance and granular capacity growth. eXFlash DIMMs provide the following advantages:

- Use universal DIMM slots, which are ideal for all server types including IBM Flex System™
- Are interoperable with standard DDR3 RDIMMs
- Support 200 GB to 12.8 TB of memory-channel storage
- Deliver performance that scales in a linear fashion with additional modules, keeping latency consistently low

eXFlash memory-channel storage contributes to lower licensing costs by enabling a higher consolidation ratio. It also helps reduce storage costs. Using internal eXFlash storage reduces or eliminates the need for storage area network (SAN) or network-attached storage (NAS) storage.

IBM FlashCache Storage Accelerator is advanced intelligent caching software that enables IBM eXFlash memory-channel storage and hard disk drive storage to work together transparently to maximize performance and to minimize cost.

IBM X6 servers support more DIMMs per channel and higher memory speeds compared to standard Intel specifications, allowing support for bigger workloads with consistent memory performance.

Agile design characteristics

Change is inevitable, and managing change is a must to achieve or to maintain market leadership. Changes in IT infrastructure typically drive complexity and cost. Managing an evolving technology, divergent customer needs, and fluctuating costs requires an agile approach to platform design. Having flexible systems to create fit-for-purpose solutions is essential.

The unique, adaptive modular rack design of the X6 family delivers agility that enables you to design a solution that meets your needs. At the same time, you can realize infrastructure cost savings by hosting multiple generations of technology in a single platform, without compromising performance or capacity.

The X6 platforms provide the following capabilities:

- You can configure the server to fit the unique requirements of your applications and workloads. You can add, modify, or upgrade the X6 platforms easily with selectable modular Book components from each of the types of the X6 Books, one for each of the major subsystems (that is, storage, compute, and I/O).
- You can scale capacity and performance from 4-sockets to 8-sockets to deliver twice the performance for growing applications without creating IT sprawl.
- You can use IBM Fast Setup software for automated provisioning of a cluster of servers and realize time-to-value in minutes rather than days.
- You can realize agile system design that is ready to host multiple generations of technology in a single server.

Resilient enterprise platforms

The growth of new applications has elevated database processing and business analytics to the top of the list of crucial x86 workloads for enterprise businesses. These environments demand continuous uptime to rapidly achieve the most valuable result: massive amounts of business-critical data. The enterprise platforms that host these workloads must deliver data at a high velocity and with continuous availability.

Through differentiated X6 self-healing technology, the X6 servers maximize uptime by proactively identifying potential failures and transparently taking the necessary corrective actions.

The X6 servers include the following unique features:

- Advanced Page Retire proactively protects applications from corrupted pages in memory, which is crucial for scaling memory to terabytes.
- Advanced Processor Recovery allows the system to automatically switch access and control of networking, management, and storage in the event of a processor 1 failure, which provides higher availability and productivity.
- Upward Integration Modules for standard hypervisors enable the creation and management of policies to maintain high availability of virtual machines and concurrent updating of the system firmware, with no impact on application performance or availability.
- The x3850 X6 and x3950 X6 modular design reduces service time by enabling quick and easy replacement of failed components.

These built-in technologies drive the outstanding system availability and uninterrupted application performance that are needed to host business-critical applications.

IBM X6 systems overview

The IBM X6 rack portfolio consists of the following flagship servers of the IBM x86 server family:

- IBM System x3850 X6
- IBM System x3950 X6

The x3850 X6 server is a 4U rack-optimized server scalable to four sockets, and the x3950 X6 server is a 8U rack-optimized server scalable to eight sockets. These systems are designed for maximum usage, reliability, and performance for compute-intensive and memory-intensive workloads.

Figure 1 shows the IBM System x3850 X6 server.

Figure 1 IBM System x3850 X6 server

The x3850 X6 server has the following key characteristics:

- Up to four Intel Xeon processor E7-4800 v2 or E7-8800 v2 product family processors
- Up to 96 DIMM slots (24 DIMM slots per processor) for up to 6 TB of memory (using 64 GB DIMMs)
- Up to 1600 MHz DDR3 memory speeds and up to 2667 MHz SMI2 link speeds
- Up to 12.8 TB of IBM eXFlash memory-channel storage
- Up to eight 2.5-inch hot-swap drives or up to sixteen 1.8-inch hot-swap solid-state drives (SSDs)
- Support for 12 Gbps SAS connectivity for the internal storage
- Mezzanine LOM slot for the integrated NIC functionality (choice of dual-port 10 GbE or quad-port 1 GbE adapters)
- Up to 11 PCIe 3.0 I/O slots
- Internal USB port for the embedded hypervisor

Unlike eX5 servers, the x3950 X6 server employs a single-server design with a single backplane, without any external connectors and cables.

Figure 2 shows the IBM System x3950 X6 server.

Figure 2 IBM System x3950 X6 server

The x3950 X6 server has the following key characteristics:

- Up to eight Intel Xeon processor E7-8800 v2 product family processors
- Up to 192 DIMM slots (24 DIMM slots per processor) for up to 12 TB of memory (using 64 GB DIMMs)

- Up to 1600 MHz DDR3 memory speeds and up to 2667 MHz SMI2 link speeds
- Up to 12.8 TB of IBM eXFlash memory-channel storage
- Up to 16 2.5-inch hot-swap drives or up to 30 two 1.8-inch hot-swap SSDs
- Support for 12 Gbps SAS connectivity for the internal storage
- Two mezzanine LOM slots for the integrated NIC functionality (choice of dual-port 10 GbE or quad-port 1 GbE adapters)
- Up to 22 PCIe 3.0 I/O slots
- Two internal USB ports for the embedded hypervisors

The supported operating systems for both x3850 X6 server and x3950 X6 server include Microsoft Windows Server 2012 R2, 2012, and 2008 R2, Red Hat Enterprise Linux 6, SUSE Linux Enterprise Server 11, and VMware vSphere (ESXi) 5.1 and 5.5.

IBM X6 technology overview

New IBM X6 technologies offer significant improvements in resiliency, scalability, and performance to meet the client's growing needs and expectations. IBM X6's superior resiliency gives clients more fault tolerance and better systems management capability to save hours of downtime from unscheduled or scheduled incidents. Seamless scalability delivers the lowest total cost of ownership through unique "pay-as-you-grow" capability with support for multiple technology generations and "fit-for-purpose" design. Breakthrough performance with the latest Intel Xeon processor E7 v2 family and flash storage technologies provides lower latencies, faster response times and smarter data management in real time.

Performance and scalability

The current models of the X6 systems use the Intel Xeon processor E7-4800 v2/8800 v2 product families. The Intel Xeon processors that are used in the X6 systems are follow-ons to the Intel Xeon processor E7-4800/8800 product family. New processors feature the Intel microarchitecture (formerly code-named *IvyBridge-EX*) and the 22 nm manufacturing process that provide higher core count, larger cache sizes, higher core frequencies, and higher memory speeds. In addition, these processors support more memory with up to 24 DIMMs per processor and faster low-latency I/O with integrated PCIe 3.0 controllers.

Table 1 summarizes key X6 scalability improvements compared to the eX5 generation.

Table 1 X6 scalability improvements

Feature	8-way eX5 system	8-way eX6 system	Improvement ratio
Processor			
Max. cores	80	120	1.5x
Max. core frequency (at max. core count)	2.4 GHz	2.8 GHz	1.17x
Max. cache size	30 MB	37.5 MB	1.25x
Memory			
Max. DDR3 memory speed	1066 MHz	1600 MHz	1.5x
Max. SMI link speed	1066 MHz	2667 MHz	2.5x

Feature	8-way eX5 system	8-way eX6 system	Improvement ratio
Max. memory bandwidth (per memory controller)	8.5 GBps	21.3 GBps	2.5x
Max. memory capacity	6 TB	12 TB	2x
Storage			
Max. SAS speed	6 Gbps	12 Gbps	2x
Max. flash storage capacity	12.8 TB[a]	25.6 TB[b]	2x
I/O			
Max. I/O slots	16	24	1.5x
Max. I/O bandwidth	72 GBps	252 GBps	3.5x

a. 32x 400 GB SSDs
b. 32x 400 GB SSDs + 32x 400 GB eXFlash DIMMs

Increased number of cores and core frequency together with larger cache size increase the CPU processing capabilities of the X6 systems by almost two times.

Another major improvement is in memory capacity and speed. IBM X6 servers support twice more memory compared to their IBM predecessors (or three times more compared to the previous generation of other industry-standard 4+-way x86 servers), and they also offer more than twice higher memory throughput.

IBM X6 servers feature 12 Gbps SAS internal storage connectivity. Twelve Gb SAS doubles the data transfer rate compared to 6 Gb SAS solutions to fully unlock the potential of the PCIe 3.0 interface and to maximize performance for storage I/O-intensive applications.

The X6 family of servers supports the latest generation of PCI Express (PCIe) protocol, Version 3.0. PCIe 3.0 is the evolution of the PCI Express I/O standard that brings doubled bandwidth over PCIe 2.0. In addition, X6 servers offer better I/O expansion capabilities with more PCIe slots and 3.5 times higher aggregated I/O throughput.

For both traditional and virtualized environments, X6 servers can help lower the overall total cost of ownership by increasing performance density and providing savings in the following areas:

- Server hardware: fewer servers to acquire, deploy, and manage
- Software: fewer sockets or cores require lower software licensing costs
- Operations: fewer servers consume less power and generate less heat
- Management: fewer servers occupy less space and require less effort to deploy and support

IBM Flash internal storage

IBM Flash storage offerings can help to close the disk I/O performance gap by combining extreme IOPS performance with low latency. IBM X6 servers use flash storage offerings to achieve fastest response time for analytical workloads, transactional databases, and virtualized environments.

The X6 servers include the following IBM Flash storage technologies:

- IBM eXFlash memory-channel storage (eXFlash DIMMs): Innovative flash storage that leverages the memory channels to provide direct connectivity to the processors, therefore delivering the lowest latencies and highest IOPS density per GB among IBM flash storage offerings.
- IBM eXFlash SSDs: Innovative high-density design of the drive cages and the performance-optimized storage controllers with the reliable high-speed SSD technology.
- IBM High IOPS MLC Adapters: Utilize the latest enterprise-level solid-state storage technologies in a standard PCIe form factor and include sophisticated advanced features to optimize flash storage and help deliver consistently high levels of performance, endurance, and reliability.
- 2.5-inch Enterprise SSDs: Designed to be flexible across a wide variety of enterprise workloads in delivering outstanding performance, reliability, and endurance at an affordable cost.
- IBM FlashCache Storage Accelerator: All-in-one flash-caching product that leverages the speed, management, capacity and breadth of the IBM qualified solid-state storage and integrates them into a high-speed server-side caching service that seamlessly accelerates the most important data with little or minimal IT overhead in both physical and virtual servers.

Table 2 compares features of the IBM Flash storage devices.

Table 2 IBM Flash storage devices

Feature	eXFlash DIMMs	eXFlash SSDs	PCIe SSD adapters	2.5-inch SSDs
Form factor	LP DIMM	1.8-inch drive	PCIe adapter	2.5-inch drive
Interface	DDR3 1333 MHz	6 Gbps SATA	PCIe 2.0 x8	6 Gbps SAS or SATA
Capacity	400 GB	Up to 400 GB	Up to 2.4 TB	Up to 1.6 TB
Max. random read IOPS	Up to 125,000	Up to 75,000	Up to 285,000	Up to 100,000
Write latency	As low as 5 µs	65 µs	15 µs	Less than 100 µs
Hot-swap capabilities	No	Yes	No	Yes
RAID support	Driver software mirroring with pairs	Yes	Chip-level redundancy	Yes

As a general consideration, 2.5-inch SSDs can be an entry-level solution that is optimized for commodity servers with a conventional hard disk drive (HDD) tray, with moderate storage IOPS performance requirements.

Both PCIe SSD adapters and eXFlash SSDs are optimized for the storage I/O-intensive enterprise workloads. PCIe adapters offer significantly lower write latency and the eXFlash SSDs offer better IOPS density and use the convenience of traditional hot-swap drives with a hardware RAID protection. At the same time, PCIe SSD adapters do not support hot-swap capabilities, and they can use operating system's software RAID capabilities to offer data protection, if required.

The highest IOPS density and the lowest latency are provided by the eXFlash DIMMs that are highly optimized for both IOPS and latency, and they can support the most demanding IOPS-intensive and latency-sensitive applications, including financial services, high-performance databases, big data analytics, and virtualization and cloud computing.

Modular design

IBM X6 servers offer “pay-as-you-grow” scalability with modular “bookshelf” design. This design allows the system to be scaled from two to eight processors. The x3850 X6’s base is a 4U chassis, and the x3950 X6’s base is a 8U chassis. All modular components are interconnected through a passive midplane (4U or 8U), and there are no external scalability connectors and cables.

The modular component that can be installed in a chassis is called *Book*. There are several types of books available:

- Compute Books

  The Compute Book contains one processor and 24 DIMM slots. It is accessible from the front of the server.

- Storage Books

  The Storage Book contains standard 2.5-inch or IBM eXFlash 1.8-inch hot swap drive bays. It also provides front USB and video ports, and it has two PCIe slots for internal storage adapters. It is accessible from the front of the server.

- I/O Books

  The I/O Book is a container that provides PCIe expansion capabilities. I/O Books are accessible from the rear of the server.

  There are three types of I/O Books:

  - Primary I/O Book. This Book provides core I/O connectivity, including the eXLOM unique slot for an onboard network, three PCIe slots, Integrated Management Module II, and rear ports (USB, video, serial, and management).
  - Full-length I/O Book. This hot-swap Book provides three optional full-length PCIe slots, and two of them are capable of hosting Graphics Processing Unit (GPU) adapters up to 300 W.
  - Half-length I/O Book. This hot-swap Book provides three optional half-length PCIe slots.

Such a modular design prepares the infrastructure for future technology upgrades, minimizing the impact on system availability and upgrade complexity and costs.

X6 is designed to replace only the Compute Books at each technology refresh, therefore reducing capital expenses and eliminating operational overhead, such as the re-configuration of the networking and storage infrastructure.

Resiliency

One of the key X6 design points is resiliency. The term *resiliency* describes the system’s ability not only to continue its operations in case of a component failure but to employ the overall strategy to avoid planned or unplanned downtime.

IBM X6 servers offer superior resiliency features that are an essential part of a mission-critical system, for example:

- Extensive memory protection with IBM Chipkill and Redundant Bit Steering (RBS).

  The combination of IBM Chipkill and RBS (also known as Double Device Data Correction or DDDC) provides robust memory protection that sustains to two sequential memory DRAM chip failures without affecting the overall system performance.

- Advanced Processor Recovery with redundant processor-to-primary I/O interconnect links enable self-recovery from a processor failure.

  If the primary processor (the one used for booting the operating system) fails, an X6 system can use a second processor to boot the OS, because the system still has access to the integrated I/O devices through the redundant links between the processors and primary I/O devices.

- IBM Advanced Page Retire coordinates memory recovery features by balancing the goals of maximum uptime and minimum repair actions.

  Advanced Page Retire uses memory hardware recovery features, followed by software recovery features, to optimize recovery results for both newer and older operating systems and hypervisors. When recovery features are exhausted, firmware issues a PFA. Memory that has failed completely is held offline during reboots until repaired.

- Predictive Failure Analysis (PFA).

  PFA allows the server to monitor the status of critical subsystems and to notify the system administrator when components appear to be degrading. In most cases, replacements of failing parts can be performed as part of planned maintenance activity. As a result, unscheduled outages can be prevented and your system continues to run.

- Light path diagnostics LCD display.

  Light path diagnostics allows systems engineers and administrators to easily and quickly diagnose hardware problems on the IBM System x® servers. Hardware failures that in the past might have taken hours to locate and diagnose can be detected in seconds, avoiding or reducing downtime.

- Intel Xeon processor E7-4800 v2/8800 v2 product family RAS features (Intel RunSure).

  The Intel Xeon processor E7 family of processors offers self-healing and failover capabilities for their interconnect links (SMI and QPI).

- Machine Check Architecture (MCA) recovery.

  MCA recovery is a set of technologies that enable the handling of system errors that otherwise require that the operating system be halted, for example:

  - Non-execution path recovery: Ability to work with hardware and software to recognize and isolate unrecoverable errors that were detected in system memory.
  - Execution path recovery: Ability to work with hardware and software to recognize and isolate unrecoverable errors that were delivered to the execution engine (core).
  - MCA I/O: Ability to report uncorrectable (both fatal and non-fatal) errors to the software layers for further handling, such as determining the root cause of failure or preventive maintenance.
  - Enhanced MCA (eMCA) Gen 1: Provides enhanced error log information to the operating system, hypervisor, or application that can be used to provide better diagnostic and predictive failure analysis for the system. This enables higher levels of uptime and reduced service costs.

- IBM Upward Integration Modules provide IT administrators with the ability to integrate the management features of the System x servers with the third-party virtualization and systems management tools.

  Upward Integration Modules provide the discovery, configuration, event management, monitoring, and automation functions that are needed to reduce cost and complexity through server consolidation and simplified management.
- Lid-less fixed chassis design.

  There is no need to pull the chassis in or out of the rack because all chassis components can be accessed either from the front or from the rear. This design provides faster maintenance by simplifying system servicing.
- PCIe adapter hot-replace and failover.
- Redundant hot-swap components.

For more information about X6 technologies, see *IBM X6 Servers: Technical Overview,* REDP-5059.

IBM X6 workloads

IBM X6 servers are designed to support mission-critical, scalable workloads, such as databases, business analytics, virtualization, and enterprise applications. This section describes typical challenges that clients might face in their environments, how IBM X6 systems address these issues, and the benefits of the IBM solution.

Mission-critical workloads

Core business services and applications deployed in modern data centers must be available 24x7. Because of their 24x7 availability requirements, we call them *mission-critical applications*. If such an application stops because a certain event occurred and the mission-critical service is disrupted, this downtime can be costly. This situation is especially important in the case of large-scale server deployments where, statistically, component failures might occur often.

Typical examples of mission-critical applications include core databases, business intelligence, corporate applications, such as enterprise resource planning (ERP), and consolidated virtualized environments that host multiple business-critical application and infrastructure services. These environments are the ones where RAS is important; better RAS means fewer outages.

There are many ways to define reliability, availability, and serviceability (RAS). A useful definition of RAS for server hardware is shown:

- *Reliability*: How infrequently a defect or fault is seen in a server.
- *Availability*: How infrequently the functionality of a system or application is negatively affected by a fault or defect.
- *Serviceability*: How well faults and their impacts are communicated to users and service personnel and how efficiently and nondisruptively they are repaired.

Using this definition, reliability in hardware is about how often a hardware fault requires a system to be serviced. The less frequent the failures, the greater the reliability. Availability is how infrequently such a failure affects the operation of the system or application.

For high levels of availability, the correct system operation must not be affected adversely by hardware faults. A highly available system design ensures that most hardware failures do not result in an application outage. Serviceability concerns identifying what fails and ensuring an efficient repair.

The ultimate design goal for IBM X6 servers as a part of the Enterprise X-Architecture strategy is to prevent hardware faults from causing an outage, and to minimize the number of system reboots and touches. The part selection for reliability, redundancy, recovery, self-healing techniques, and degraded operational modes is used in a fault-resilient RAS strategy to avoid application outages (see Figure 3).

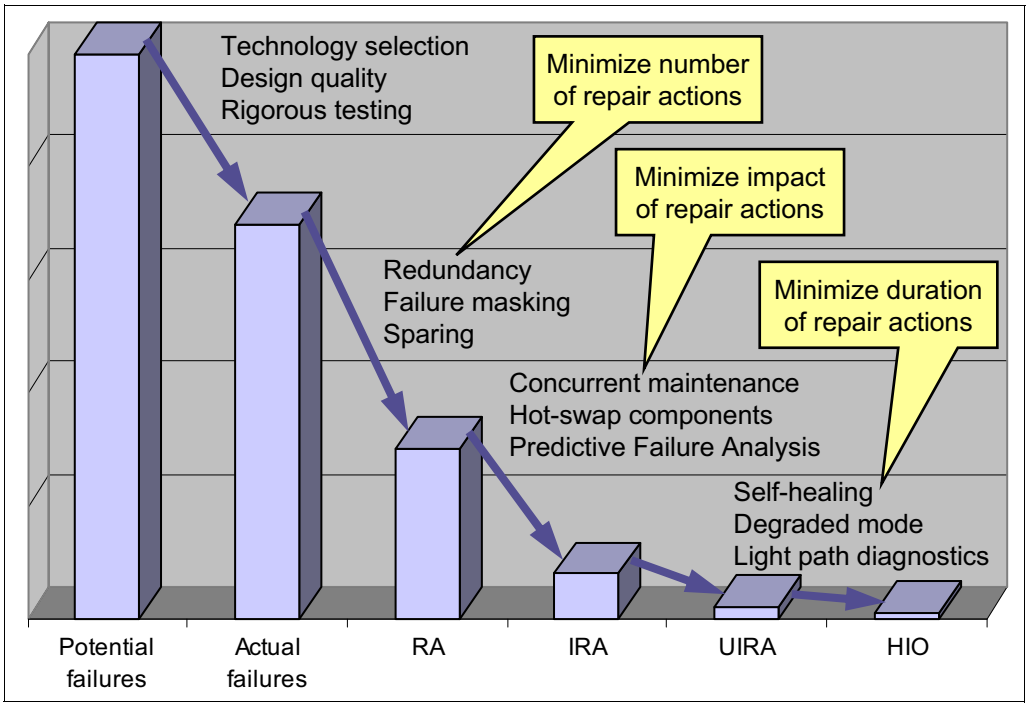

Figure 3 IBM Enterprise X-Architecture RAS strategy

The intent of 24x7 availability is to reduce the impact of these failures on system operations. IBM traditionally classifies hardware failures in multiple ways:

- Repair Action (RA): RAs are related to the industry standard definition of mean time between failure (MTBF). An RA is any hardware event that requires service on a system. Repair actions include incidents that affect system availability and incidents that are concurrently repaired.
- Interrupt Repair Action (IRA): An IRA is a hardware event that requires a scheduled system outage to repair.
- Unscheduled Incident Repair Action (UIRA): A UIRA is a hardware event that causes a system to be rebooted in full or degraded mode. The system experiences an unscheduled outage. The restart might include some level of capability loss, but the remaining resources are made available for productive work.
- High Impact Outage (HIO): An HIO is a hardware failure that triggers a system crash that is not recoverable by immediate reboot. This failure is usually caused by the failure of a component that is critical to system operation and is, in some sense, a measure of system single points of failure. HIOs result in the most significant availability impact on the system because repairs cannot be effected without a service call.

The core principles guiding the IBM engineering design of the mission-critical servers are reflected in the advanced RAS architecture. The goals of an X6 server design are listed:

- Achieve a highly reliable design through the extensive use of highly reliable components and rigorous compatibility testing.
- Employ a server architecture that allows the system to recover from intermittent errors in the components and fail over to redundant components when necessary.
- Develop server hardware that can detect and report on failures and impending failures.
- Create server hardware that is self-healing, that is, it automatically initiates actions to effect error correction, repair, or component replacement.

Standard approaches to server RAS focus on errors that can be handled at the hardware level. In such a case, any uncorrectable error causes a system outage. In this respect, another design element of the IBM X6 RAS strategy is to enable software layers (operating system, hypervisors, and applications) to participate in the recovery process. The idea is to make the software that runs on a system both hardware fault-aware (both soft and hard errors) and to use its capabilities to prevent a service failure or to restore service operations. This combination of hardware and software self-recovery techniques is part of the advanced RAS features that increase the availability of services that must be 24x7.

Figure 4 shows the different operational states of a system that follows the X6 RAS strategy.

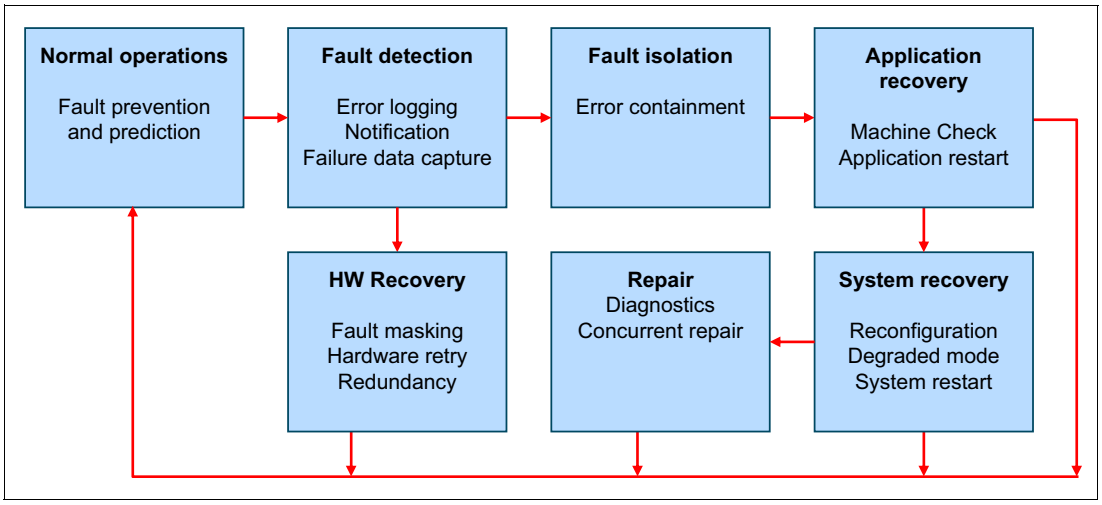

Figure 4 IBM X6 system RAS operations

IBM technologies that are used to achieve the highest levels of resiliency and fault-tolerance, along with ease of serviceability, can help achieve these benefits:

- Higher availability by reducing the number of unscheduled incident repair actions due integrated self-healing capabilities
- Faster maintenance with no or little downtime due to simplified servicing with modular “bookshelf” design with easily accessible components
- Simplified ongoing management and maintenance with IBM advanced systems management tools and upward integration with third-party systems management solutions to leverage IBM-unique monitoring and reporting capabilities to perform policy-based workload management and intelligent system software upgrades

In scalable infrastructures, Enterprise X-Architecture servers can potentially lower the overall number of repair actions by three times[1], compared to the traditional dual-socket rack server-based environments. Among them, the number of scheduled repair actions is reduced by four times, the number of unscheduled repair actions recovered through self-healing by using degraded mode is cut by five times, and the number of unexpected outages that require manual intervention to repair drops by more than six times.

The IBM X6 RAS design ensures that the clients can operate their mission-critical enterprise workloads on IBM X6 servers with a high level of confidence.

Figure 5 illustrates an example of how an application can use multiple levels of X6 recovery features to handle system errors to prevent itself from being terminated in the case of a system error.

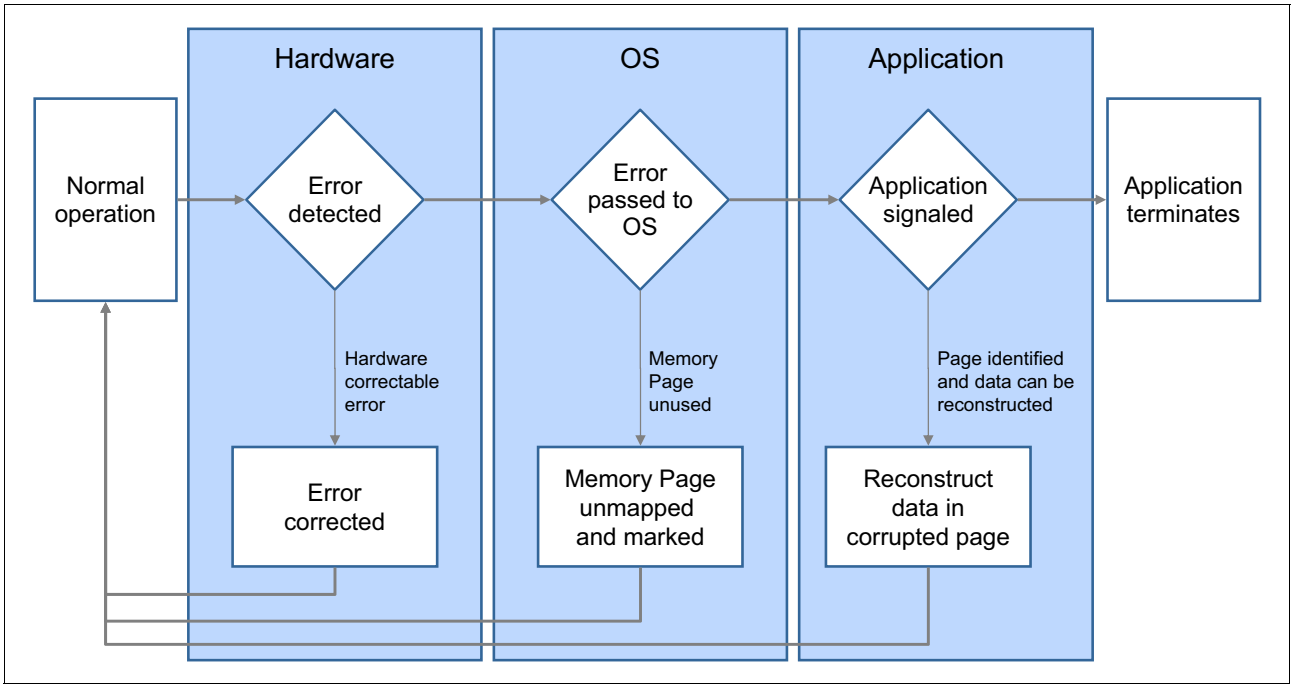

Figure 5 Machine Check Architecture recovery

If a recoverable memory error is detected, the hardware corrects it in a manner that is transparent to the operating system and applications. If a memory error is encountered that cannot be corrected by the hardware, the processor sends an MCA recovery signal to the operating system. An operating system that supports MCA recovery now determines whether the affected memory page is in use by an application. If unused, it unmaps the memory page and marks it as in error. If the page is used by an application, the OS usually holds that application or stops all processing and halts the system.

With the MCA-aware application, the operating system can signal the error situation to the application, giving it the chance to try to repair the effects of the memory error. Using MCA recovery, the application can take an appropriate action at the level of its own data structures to ensure a smooth return to normal operation, avoiding a time-consuming restart or loss of information.

[1] *Reliability, Availability, and Serviceability Features of the IBM eX5 Portfolio*, REDP-4864, at this website: http://www.redbooks.ibm.com/abstracts/redp4864.html

Transactional (OLTP) databases

Online transaction processing (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 X6 offerings address these challenges by providing low latency, extreme performance, and efficient transaction management.

For most OLTP systems, the processor, memory, and I/O subsystem in a server are well balanced and are not considered performance bottlenecks. The major source of performance issues in OLTP environments is typically related to the storage I/O activity. The speed of traditional hard disk drive (HDD)-based storage systems does not match the processing capabilities of the servers. As a result, often a situation occurs where a powerful processor sits idle, waiting for the storage I/O requests to complete, negatively affecting user and business productivity. The negative impact on productivity extends the time to achieve a return on investment (ROI) and increases the overall total cost of ownership (TCO).

This performance imbalance became even sharper with the introduction of the IBM X6 servers featuring Intel Xeon processor E7-4800 v2/8800 v2 product families. Therefore, storage IOPS performance and latency become strategic considerations for business. It is critical to ensure that the response time goals are met and that performance optimization is realized for other system resources (processor and memory).

In general, clients might experience the following challenges in OLTP environments:

- Failure to meet user expectations and service levels because of slow application response time
- Decreased user and business productivity
- Application and data availability concerns (slow batch processing, long backup windows, and hardware failure rates)
- Increased storage performance and capacity requirements
- Scalability constraints because of data center space, power, and cooling limits
- Increasing TCO:
  - Rising data center power and cooling costs
  - Increasing software licensing fees
  - Rising server, network, and storage infrastructure management and support costs
- Longer lead time to ROI because of inefficient utilization of the existing resources

The OLTP workload optimization goal for the IBM X6 systems is to address these issues in the most cost-efficient way in the following areas:

- In-memory database
- Flash storage as a cache
- Flash storage as a primary data store

In-memory database

The main memory is the fastest storage type that can hold a significant amount of data. Data in main memory can be accessed more than a hundred thousand times faster than data on a spinning hard disk, and even flash technology storage is about a thousand times slower than main memory.

Main memory is connected directly to the processors through a high-speed bus, where hard disks are connected through a chain of buses (QPI, PCIe, and SAN) and controllers (I/O hub, RAID controller or SAN adapter, and storage controller).

Compared with keeping data on disk, keeping the data in main memory can dramatically improve database performance just by the advantage in access time. However, there is one potential drawback.

In database technology, atomicity, consistency, isolation, and durability (ACID) is a set of requirements that guarantees that database transactions are processed reliably:

- A transaction must be atomic. That is, if part of a transaction fails, the entire transaction has to fail and leave the database state unchanged.
- The consistency of a database must be preserved by the transactions that it performs.
- Isolation ensures that no transaction interferes with another transaction.
- Durability means that after a transaction is committed, it will remain committed.

Although the first three requirements are not affected by the in-memory concept, durability is a requirement that cannot be met by storing data in main memory alone. Main memory is volatile storage. That is, it loses its content when it is out of electrical power. To make data persistent, it must reside on non-volatile storage, such as hard disk drives, solid-state drives (SSDs), or flash devices. Therefore, some sort of permanent storage is still needed.

IBM X6 offers workload-optimized models that are specifically designed for the SAP HANA in-memory database.

Figure 6 illustrates the SAP HANA scale-out cluster solution, showing a 3-node configuration as an example.

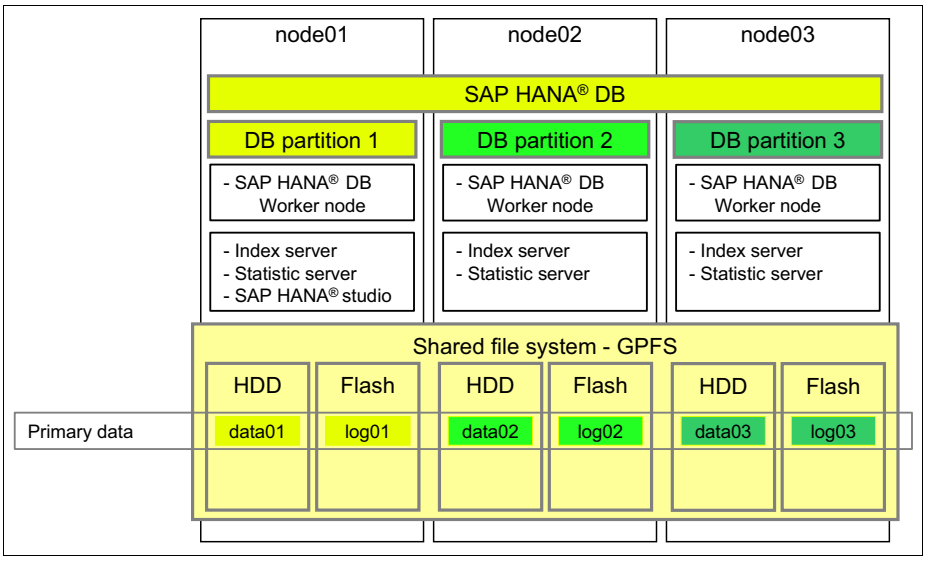

Figure 6 A 3-node clustered solution without failover capabilities

This scale-out solution consists of the following components:

- The solution consists of a homogeneous cluster of X6 building blocks, which are interconnected with two separate 10 Gb Ethernet networks (not shown in Figure 6): one for the SAP HANA application and one for the General Parallel File System (GPFS™) file system communication.

- The SAP HANA database is split into partitions, forming a single instance of the SAP HANA database.
- Each node of the cluster holds its own savepoints and database logs on the local storage devices of the server.
- The GPFS file system spans all nodes of the cluster, making the data of each node available to all other nodes of the cluster.

In addition, this scenario can be expanded with failover capabilities by adding a standby node.

For more information about SAP HANA on X6 servers, see *SAP HANA on IBM X6 Systems: Reference Architecture,* SG24-8214.

Flash storage as a cache

When X6 flash storage is utilized for caching, IBM FlashCache Storage Accelerator is used. IBM FlashCache Storage Accelerator is a suite of caching software and tools that are designed to significantly increase your server and storage performance and efficiency by helping eliminate I/O bottlenecks, and keep the most active data closer to the application. This in turn helps you gain greater system utilization, lower software and hardware costs, and save power, cooling, and floor space.

Used in conjunction with IBM flash storage offerings, IBM FlashCache Storage Accelerator can provide an efficient, cost-effective, and easy-to-implement solution for virtual and physical environments.

IBM FlashCache Storage Accelerator offers these key features:

- Transforms IBM Flash storage (IBM SSDs, High IOPS SSD Adapters, and eXFlash) into a transparent acceleration device to cache frequently accessed data off of any storage system, whether SAN or DAS.
- Caches data within servers, close to the application workloads, delivering extremely low latency.
- Works in physical (non-virtualized) and virtual environments.
- Transparently and dynamically rebalances I/O.
- Includes disk, file, volume, and virtual machine (VM)-specific caching.
- Is supported on a wide variety of operating systems, including VMware, Microsoft Windows Server, and Linux.
- Supports both native operating system high availability clustering (such as Microsoft Failover Clustering and Linux high availability clusters) and advanced workload availability and rebalancing for virtualized environments (VMware vMotion, HA, FT, and Distributed Resource Scheduling or DRS).

An OLTP solution with eXFlash and FlashCache Storage Accelerator typically consists of the following components:

- IBM System x database servers (typically, IBM X6 systems) that run data management software such as IBM DB2®, Microsoft SQL Server, or Oracle Database. These database servers also host the following products:
  - IBM FlashCache Storage Accelerator Direct Edition for Windows
  - IBM eXFlash (DIMMs or SSDs)
- External shared storage system that hosts the entire data set
- Storage area network (SAN) that is used to provide connectivity across database servers and storage systems

The OLTP database acceleration scenario is shown in Figure 7.

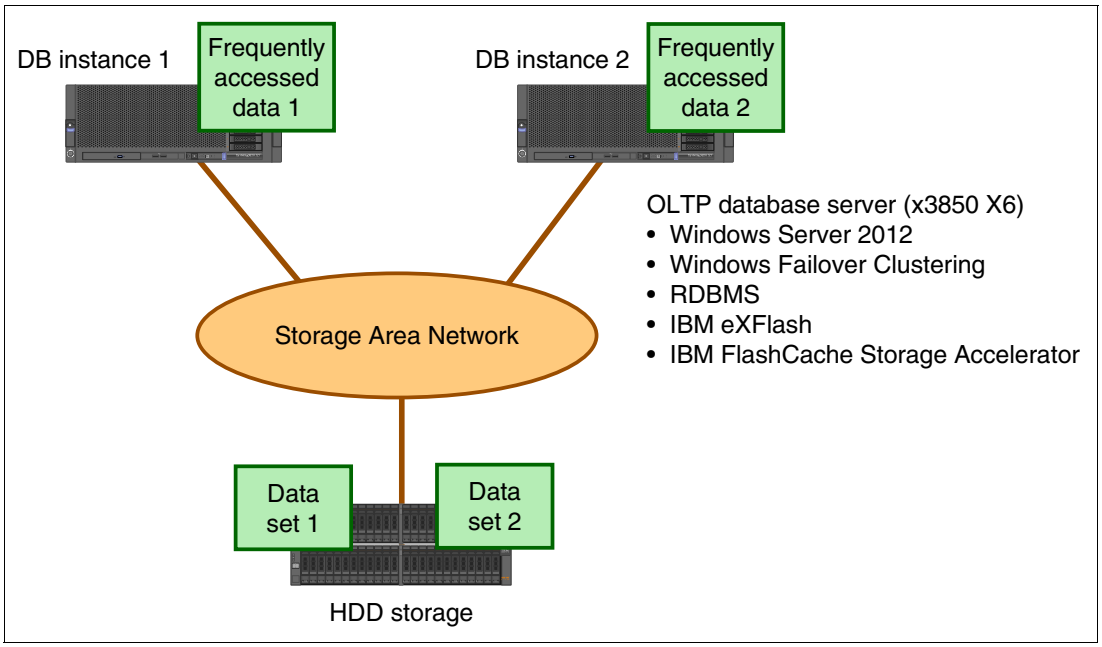

Figure 7 OLTP database acceleration

In this scenario, database server systems are connected to the external shared storage via Fibre Channel SAN. Host systems run the Microsoft Windows Server 2012 operating system. They support multipathing to allow redundant storage connections through SAN, including dual-port storage system interface connections, dual-port host bus adapters (HBAs) on the host systems, and redundant SAN switched fabric.

Database high availability is achieved by using the Windows Failover Clustering service, and there are two database instances that are running in the cluster. The entire data sets for the database instances (Data set 1 and Data set 2) are hosted on the external SAN-attached storage, and server-side caching is implemented with the IBM FlashCache Storage Accelerator and IBM eXFlash.

The most frequently used data from the data sets on the SAN are dynamically and transparently cached on the server's local eXFlash storage to increase performance and lower response time.

For more information, see *The Benefits of IBM FlashCache Storage Accelerator in Enterprise Solutions,* REDP-5080.

Flash memory as a primary data store

In the primary data store scenario, the entire data set is placed on eXFlash internal storage in the server. There is no external storage, and the eXFlash storage is configured with RAID protection to provide data redundancy. The database can be deployed on a single node or on multiple nodes by using the database partitioning feature.

The partitioning feature of many databases (for example, IBM DB2) can help to split the workload among several nodes, increasing overall performance, availability, and capacity.

In addition, to ensure that the high availability requirements are met in case of a node failure, several techniques can be utilized depending on the database vendor.

The following techniques are included:

- Log shipping
- Replication
- Database mirroring

If, for some reason, the entire database cannot be placed on IBM eXFlash, consider putting at least part of the data there. Look at the following areas:

- Log files
- Temporary table space
- Frequently accessed tables
- Table partitions
- Indexes

IBM X6 benefits for OLTP workloads

IBM X6 in-memory computing and flash storage offerings can help to deliver the following benefits in the OLTP environments:

- Dramatically boosting the performance of existing applications and lowering cost per IOPS ratio without a need for re-architecture
- Increasing user productivity with better response times, improving business efficiency
- Increasing 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
- Increasing performance and capacity while decreasing power, cooling, and space requirements
- Reducing TCO:
  - Reducing energy costs because of lower power and cooling requirements
  - Reducing the number of systems, devices, and components that are required to build the solution by increasing the usage of available resources
  - Reducing software license fees because fewer systems or processors are required
  - Reducing management and support costs because there are fewer components to deploy and support
  - Faster ROI because of better resource usage

The IBM X6 servers help offer the following expected performance increase:

- Up to two times more performance due to Intel Xeon processor E7-4800/8800 v2 product families and support for larger memory capacity.
- Up to ten times less I/O latency with IBM eXFlash DIMMs compared to traditional solid-state drives.
- eXFlash DIMMs provide up to 40% better read IOPS performance, and they also improve write IOPS performance by almost 60%.

Business analytics

Data warehouses are commonly used with online analytical processing (OLAP) workloads in decision support systems, such as financial analysis. Unlike OLTP, where transactions are typically relatively simple and deal with small amounts of data, OLAP queries are more complex and process larger volumes of data.

OLAP databases are normally separated from OLTP databases and tend to consolidate historical and reference information from multiple sources. Queries are submitted to OLAP databases to analyze consolidated data from different points of view to make better business decisions in a timely manner.

For OLAP workloads, a fast response time is critical to ensure that strategic business decisions can be made quickly in dynamic market conditions. Delays can significantly increase business and financial risks. Usually, decision making is stalled or delayed because of a lack of accurate, real-time operational data for analytics, which means missed opportunities for the following reasons:

- Inability to gain insight into a business
- Inability to predict business outcomes
- Explosion of volume, variety, and velocity of information

IBM X6 systems can help to make businesses more agile and analytics-driven by providing up-to-the-minute analytics based on real-time data, and not yesterday's news.

The delays come primarily from batch data loads and performance issues due to handling heavy complex queries that use I/O resources. A common performance bottleneck in OLAP environments is the I/O that is required for reading massive amounts of data (frequently referred to as *big data*) from storage for processing in the OLAP database server. The server ability to process this data is usually a nonfactor because the servers typically have significant amounts of RAM and processing power, parallelizing tasks across the computing resources of the servers.

In general, clients might experience the following challenges in OLAP environments:

- Slow query execution and response times, which delay business decision making
- Dramatic growth in data, which requires deeper analysis

As with OLTP workloads, in-memory databases or flash storage are used for workload optimization (see "Transactional (OLTP) databases" on page 15).

A sample OLAP solution with eXFlash as a primary data storage (shown in Figure 8) consists of the following components:

- IBM System x database servers (x3850 X6 systems) that run data management software. These database servers also host IBM eXFlash storage for the partitioned data sets.
- A private network, such as 10 Gb Ethernet or QDR/FDR InfiniBand, that is used to provide high-speed connectivity across database servers in a cluster.

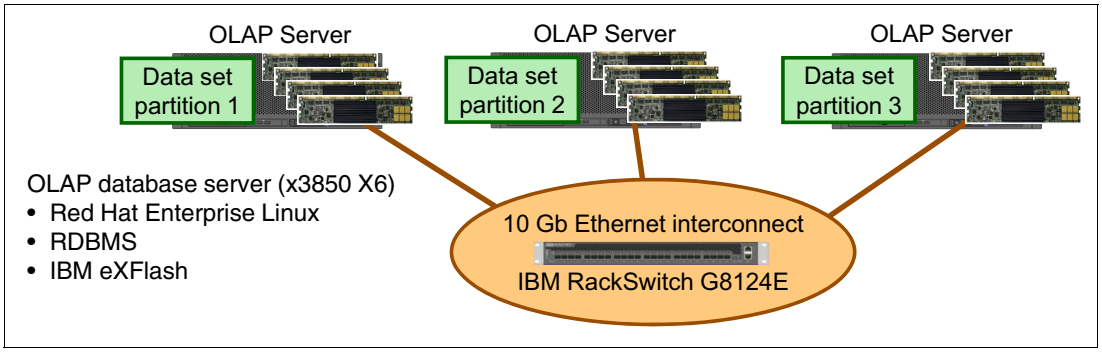

Figure 8 OLAP database acceleration

Server hosts, or nodes, are interconnected with the isolated high-speed network, such as 10 Gb Ethernet with IBM RackSwitch™ G8124E, that is used for the inter-node data exchange.

Each node runs a copy of the OLAP database application, and the analyzed data set is partitioned and distributed across the storage systems. Depending on the database management software that is used and its architecture, each node might have access to only a certain portion of data. OLAP queries are distributed and processed across the nodes.

This solution can scale easily by adding more similarly configured nodes. In such a case, storage capacity and I/O bandwidth are incremented linearly with the increasing number of nodes, which can help to eliminate storage I/O bottlenecks in OLAP workloads.

IBM X6 benefits for OLAP workloads

IBM X6 systems can help to address OLAP challenges in the following ways:

- Dramatically boosting the performance of OLAP workloads with distributed scale-out architecture, providing almost linear and virtually unlimited performance and capacity scalability
- Significantly improving response time for better and timely decision making

For more information, see *The Benefits of IBM eXFlash Memory-Channel Storage in Enterprise Solutions,* REDP-5089.

Virtualization and cloud computing

While many users realize benefits on savings from deploying server virtualization solutions, there are still some questions and concerns that must be addressed to move the adoption of virtualization to the next level, including applications that were not able to be virtualized before due to existing storage I/O constraints.

One of the key requirements of any virtualized environment is high availability for virtual machines (VM) in addition to easy VM migration. However, most high availability technologies require shared access to storage from multiple physical servers to enable the migration of VMs across hosts.

Traditional data center architectures rely on expensive SAN-based storage systems for shared storage access. However, SAN storage can become a performance bottleneck when deployed with clusters of VM-dense servers, and is expensive to scale. The limited IOPS available for a SAN storage array have to be shared with a large number of VMs across multiple physical servers. Additionally, there is no easy way to segregate I/O traffic on a per-VM basis to allocate IOPS to the specific VMs that need them.

However, the overall architecture still needs to maintain shared storage access across physical servers to enable VM mobility. While it is possible to use local flash in physical servers as static cache, this approach ties VMs to physical servers and breaks VM mobility. Furthermore, in a mixed virtualized environment, we can observe different kinds of workloads and storage access patterns, even on the same physical server.

IBM X6 servers with eXFlash storage, combined with IBM FlashCache Storage Accelerator, can help solve many challenges of a virtualized environment by transparently and dynamically adopting to the changing workload patterns while keeping essential VM mobility and availability features and increasing storage IOPS performance.

In virtual environments, the FlashCache Storage Accelerator is deployed as a part of the hypervisor (VMware ESXi), as shown in Figure 9. Optionally, for more granular cache management, FlashCache Storage Accelerator elements can be installed within the guest operating system, in addition to the hypervisor installation.

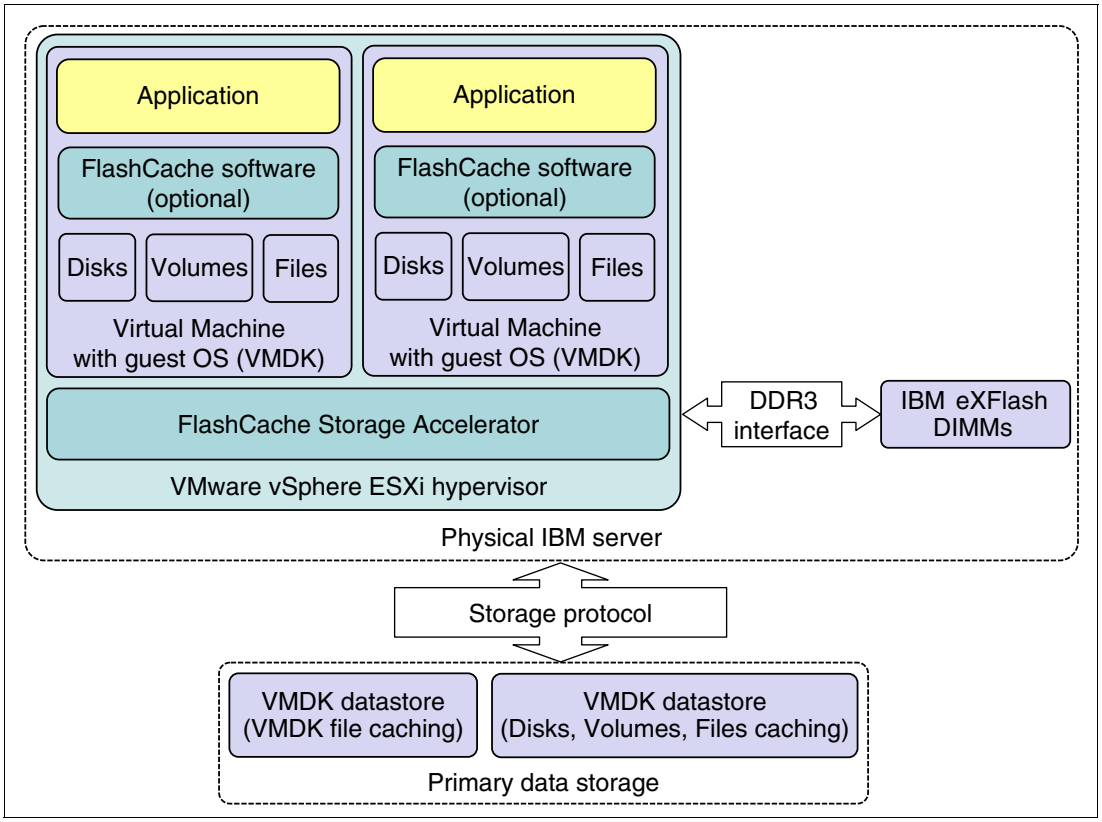

Figure 9 FlashCache Storage Accelerator in virtual environments

The following caching levels are supported in the virtual environments when the FlashCache Storage Accelerator components are installed on the VMware ESXi hypervisor and within the guest operating system (OS):

- Disk caching
- Volume caching
- File caching
- VMDK file caching

If the FlashCache Storage Accelerator is installed only as a part of the hypervisor, only VMDK file caching is available.

Performance consideration: As a general rule, better performance results can be achieved when the FlashCache Storage Accelerator is deployed within the guest OS.

The FlashCache Storage Accelerator supports storage-specific VMware multi-pathing software to take full advantage of redundant SAN storage connections.

The FlashCache Storage Accelerator transparently supports VMware high availability clustering (VMware HA), live migration (vMotion), and dynamic resource reallocation (Distributed Resource Scheduling or DRS).

IBM X6 benefits for virtualization solutions

IBM X6 systems can help to provide the following benefits in virtualized solutions:

- Ability to virtualize data-intensive enterprise applications that were unable to be virtualized before due to storage I/O or memory constraints
- Higher reliability and availability of services due to fewer number of components used to build the solution
- Lower acquisition costs due to fewer number of systems and components
- Shorten ROI time frame and decrease overall TCO with the efficient utilization of server resources and lower software, power, cooling, and management costs

For more information, see *The Benefits of IBM FlashCache Storage Accelerator in Enterprise Solutions,* REDP-5080.

Enterprise applications

Enterprise applications, such as ERP or customer relationship management (CRM), represent a mixed workload where both transaction processing and a certain level of real-time reporting exist. In a 2-tier implementation, both database server and application modules reside on the same server. The key performance metric is response time, such as with OLTP and OLAP workloads.

IBM X6 offerings provide low latency, extreme performance, and efficient transaction management to accommodate mixed workload requirements.

IBM X6 offerings are the ideal platform for clients to implement mission-critical applications, such as SAP, and who are looking for reliability, manageability, and scalability with the flexibility to run Windows or Linux. This new family of products provides a cost-efficient, scalable platform with superior performance and unmatched reliability and resiliency. The IBM X6 family is perfectly suited for enterprise solutions because it offers the following unmatched advantages:

- Scalability

  IBM X6 offers unparalleled scalability inside the 4-way (x3850 X6) and 8-way (x3950 X6) servers. Using 15 cores per processor, the top configuration can reach 120 cores and 12 TB of memory.

- Performance

  IBM X6 servers are purpose-built for high-end enterprise workloads. They offer an increase in the total available memory for the solutions with high memory demands (such as SAP databases and virtualization scenarios), and they deliver greater memory bandwidth for enhanced performance.

  IBM eXFlash memory-channel storage, with up to 32 eXFlash DIMMs per system in combination with a high-speed memory controller, deliver extreme IOPS performance and consistently low latency to replace the limited performance capabilities of traditional hard disk drives.

- Reliability

  It is imperative for a critical server environment to have built-in redundancy on a maximum number of components to ensure that the servers are always available to the users with no downtime. With unique advanced RAS (reliability, availability, and serviceability) features such as Machine Check Architecture recovery, Chipkill, redundant bit steering, memory rank sparing, and memory mirroring, the X6 family is the highly optimized platform to meet this requirement for high-end solutions (such as SAP).

- Cost-effective deployment

  Because the IBM X6 servers offer extremely high scalability and high performance, the overall hardware sizing requirement is reduced for any application. Moreover, the simplification of deployment and maintenance allows SAP customers to focus on their implementation and better manage their infrastructure costs.

- Virtualization and provisioning

  Virtualization technologies in conjunction with the right provisioning tools can help clients manage compute resources in a more efficient manner. This helps increase the overall use of assets, while providing additional capacity on demand. It improves collaboration across spheres of any organization, helping the organization to focus on growing their business. The IBM X6 servers help reduce license costs for virtualization by providing up to 1.5 TB of memory per Computer Book.

- Manageability

  IBM helps clients improve their IT infrastructure management by reducing the complexity of IT. Through autonomic computing technologies, IBM helps clients to better use their critical skills, allowing IT to do more with less. This helps position them to respond dynamically to unpredictable workloads and better manage the proliferation of systems.

- Availability

  IBM created a hot-swap server infrastructure that allows clients to add and replace hard disk drives (SAS/SATA), power supplies, fans, and optional I/O Books dynamically. This capability increases the availability of the server and reduces the cost of maintenance of the server and systems running inside the box.

  Advanced self-healing architecture and resiliency features provide automated recovery from certain failures, helping minimize unplanned downtime.

- I/O throughput and latency

  The X6 server family offers integrated PCIe 3.0 I/O connectivity, which helps achieve lower-access latency and higher throughput capabilities of up to 31.5 GBps per processor.

For more information, see *SAP Business Suite on X6 Servers: Reference Architecture,* REDP-5073.

Summary

The IBM X6 family of scalable rack servers consists of the IBM System x3850 X6 server, a 4U four-socket server, and the x3950 X6 server, an 8U eight-socket server. Leveraging the proven technologies of the previous generations of Enterprise X-Architecture, these new servers introduce the following benefits:

- New levels of fault tolerance and resiliency with advanced RAS features implemented in hardware and software
- Agility and scalability with a fit-for-purpose modular “bookshelf” design that is ready to support multiple technology upgrades
- Significant improvements in response time with ultralow latency, stretched memory speeds that exceed Intel specifications, and innovative flash memory-channel storage offerings

IBM X6 servers continue to lead the way as the shift toward mission-critical scalable databases, business analytics, virtualization, enterprise applications, and cloud applications accelerates.

IBM X6 systems using the Intel Xeon processor E7 v2 family deliver an extensive and robust set of integrated advanced RAS features that prevent hardware faults from causing an outage. Part selection for reliability, redundancy, recovery, and self-healing techniques and degraded operational modes are used in a RAS strategy to avoid application outages. Using this strategy, IBM X6 systems can help increase application availability and help reduce downtime by enabling 24x7 mission-critical capabilities to run your core business services.

Related publications

The following IBM Redbooks® publications provide additional information about the topic in this document. Note that some publications referenced in this list might be available in softcopy only.

- *IBM X6 Servers: Technology Overview*, REDP-5059
- *IBM System x3850 X6 and x3950 X6 Product Guide,* TIPS-1084
- *IBM System x3850 X6 and x3950 X6 Implementation Guide*, SG24-8208
- *The Benefits of IBM FlashCache Storage Accelerator in Enterprise Solutions,* REDP-5080
- *SAP Business Suite on X6 Servers: Reference Architecture,* REDP-5073
- *SAP HANA on IBM X6 Systems: Reference Architecture,* SG24-8214
- *The Benefits of IBM eXFlash Memory-Channel Storage in Enterprise Solutions,* REDP-5089
- *Reliability, Availability, and Serviceability Features of the IBM eX5 Portfolio*, REDP-4864

Author

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 150 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:

From International Technical Support Organization, Raleigh Center:

- Kevin Barnes
- Tamikia Barrow
- David Bennin
- Ella Buslovich
- Mary Comianos
- Cheryl Gera
- Chris Rayns
- David Watts
- Debbie Willmschen

From IBM:

- Randy Kolvick
- Randall Lundin

- Iliyas Pathan
- Doug Petteway

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 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.

This document REDP-5058-00 was created or updated on February 14, 2014.

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