



On-Demand Scalability for Mission-Critical Databases

IBM® DB2® pureScale™ on Intel® processor-based IBM System x® servers





Contents

Executive Summary	3
The Challenges of Scaling Mission-Critical Databases	4
Powerful, Enterprise-Class Servers.....	5
High-Performance Clustering for Near-Limitless Scalability	8
Application Transparency Helps Eliminate Hidden Costs	10
High Availability for Uninterrupted Operations	12
Servers to Match Specific Needs	14
IBM eX5 Servers: High Value for Large, Rapidly Growing Databases	14
Conclusion	15

Executive Summary

The pressures related to managing transactional databases are increasing rapidly. Business growth, the drive to consolidate databases and the need to deploy new, data-intensive technologies are fostering the massive expansion of data volumes and application workloads. At the same time, the move toward real-time computing requires faster and more reliable data access, especially when databases are used to drive customer-facing applications.

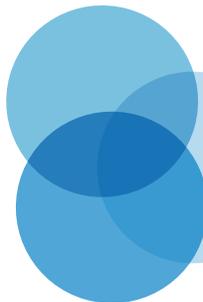
Businesses need simpler and more cost-effective strategies for expanding their database environment. IBM and Intel provide an answer with IBM® DB2® pureScale™ software and the latest generation of IBM System x® servers based on Intel® Xeon® processors. The combined solution enables customers to scale mission-critical, performance-sensitive databases simply, incrementally, and without limit using affordable, industry-standard servers.

IBM eX5 servers deliver breakthrough performance, scalability, and reliability. A fully configured IBM System x3850 X5

system provides ample capacity for heavy, data-intensive workloads, with up to 40 processing cores, 80 threads, and 3 TB of memory.¹ Servers equipped with the Intel Xeon processor E7-4800 product family can achieve up to 29 times the performance of 4-socket servers using single-core processors.² Major silicon-level enhancements help improve scalability, reliability, and data integrity.

With IBM DB2 pureScale software, multiple System x servers can be combined into a high-performance cluster that provides virtually unlimited capacity with continuous availability and complete application transparency. A new server can be integrated into the cluster with just two commands: no application changes are required, and workloads are automatically rebalanced to optimize performance—a breakthrough capability that makes it far easier to handle today's exploding data volumes. IT organizations can now take a building-block approach to database expansion by scaling large, mission-critical databases incrementally, affordably, and virtually without limit.

This white paper describes the systems and technologies that enable a powerful new strategy for expanding database environments. It provides valuable information for any IT professional looking for simple, cost-effective ways to support rapidly growing database requirements.





The Challenges of Scaling Mission-Critical Databases

We live in a time of rapid data growth. According to Gartner, IT data volumes can be expected to grow 650 percent over the next five years.³ Business growth can accelerate this increase, as can mergers, acquisitions, new application deployments, and application upgrades that increase data processing workloads. Database solutions must be able to absorb increasing data volumes and expanding workloads while sustaining performance and avoiding disruptive capacity upgrades.

Transactional databases are the lifeblood of the enterprise. Critical data must be available quickly and without interruption to keep the business running at peak efficiency. If a key database goes down or slows down, dependent applications suffer the same fate—often costing the business millions of dollars in lost revenue and dissatisfied customers. Unplanned downtime is not the only risk: in today's global, 24x7 business environment, taking down a key database for upgrades or maintenance can entail considerable cost.

Traditionally, scaling capacity for enterprise databases has been particularly challenging because of the need to host the database on a single server. Even a large multiprocessor system can take only so many processors,

memory chips, and I/O cards before reaching its full capacity. The next step would be to replace the entire system—an inherently expensive proposition.

Clustered database solutions offer an alternative approach to database scaling, one that has been available for several years. However, expanding a cluster has been a resource-intensive operation, requiring substantial time and expertise to repartition data, rebalance workloads, and optimize performance.

IBM and Intel are working to eliminate these challenges. IBM System x servers based on the Intel Xeon processor E7 platforms provide an ideal platform for data-intensive workloads in mission-critical environments. IBM DB2 pureScale software, an optional feature of DB2 for Linux*, UNIX* and Microsoft Windows*, enables these servers to be clustered together to support truly massive workloads with excellent performance, linear scalability, and high availability. Together, IBM System x servers and IBM DB2 pureScale software are fundamentally changing the economics of scaling large, mission-critical databases.

Powerful, Enterprise-Class Servers

IBM eX5 servers based on the Intel Xeon processor E7 platforms provide a robust foundation for mission-critical workloads with industry-standard servers. With on-demand scalability, flexible partitioning, exceptional memory capacity, and advanced reliability features, the eX5 servers provide powerful data center platforms that can be deployed affordably and then scaled easily to meet changing requirements.

Building on the breakthrough capabilities introduced with the Intel Xeon processor 7500 series, the Intel Xeon processor E7 platforms play an essential role in delivering the performance, scalability, and reliability to support demanding applications on large, enterprise-class systems. Servers based on

these processors can help IT organizations scale performance as workloads grow and maintain high availability—all while preserving the compelling cost advantages of Intel processor-based solutions.

Performance

The Intel Xeon processor E7 platforms offer outstanding performance for heavy, data-intensive workloads. Built on a new 32nm microarchitecture, these platforms can provide up to 10 cores, 20 threads, and 30 MB of shared cache per processor. Four advanced, high-bandwidth interconnect links allow multiple processors to be directly connected to each other to increase performance and reduce latency (see Figure 1).

Intel Microarchitecture Geared to Meet Advanced DBMS Requirements

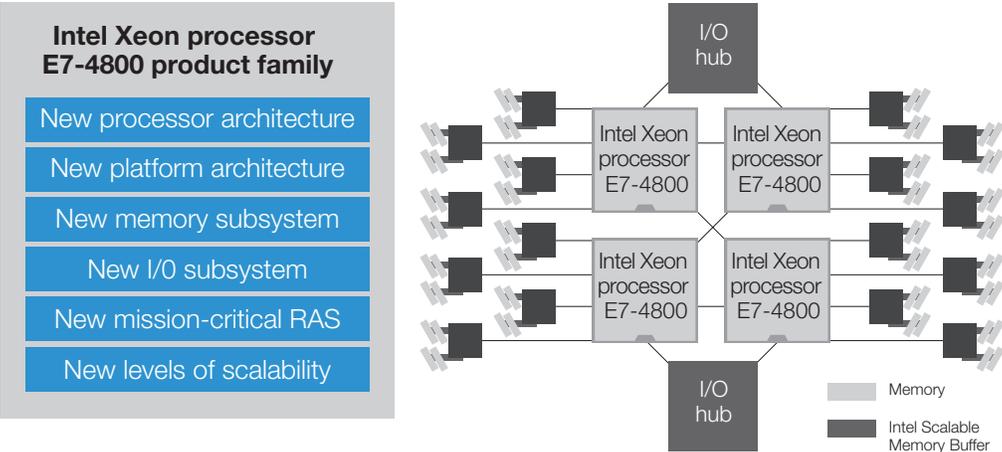


Figure 1. The Intel Xeon processor E7-4800 product family capitalizes on an Intel microarchitecture designed to provide the high-end scalability and availability needed to support large, mission-critical databases.



These new capabilities deliver tangible performance results. The Intel Xeon processor 7500 series, which preceded the Intel Xeon processor E7 platforms, delivered the largest performance leap ever for an Intel processor generation, with an average of three times higher performance across a wide range of industry benchmarks⁴ compared with the previous processor generation and up to 20 times better performance per server versus single-core servers.⁵

Intel Xeon processor E7 platforms boost performance even further. Using the Intel Xeon E7-4800 product family, organizations could accelerate database transactions by up to 40 percent compared with previous-generation 4-socket processors.⁶ As a result, organizations can generate results faster or process larger data volumes than ever before without having to increase the server footprint.

Scalability

The Intel Xeon processor E7 platforms provide the scalability required for adding or expanding workloads to accommodate new business opportunities. Intel® QuickPath Interconnects allow scaling of processors from 2 sockets to up to 8 sockets when combining two independent IBM eX5 servers to create one larger system. The Quad-Channel Integrated Memory Controller supports up to 16 memory slots per processor socket, enabling administrators to deploy up to 2 TB of memory in a 4-socket server to handle peak demands and leave

headroom for database growth. By using IBM solidDB® in-memory database software, organizations take full advantage of this large-scale memory capacity to accelerate database performance, running queries in real time and speeding the path from data to decision.

Reliability

The Intel Xeon processor E7 platforms include more than 20 features designed to deliver the reliability, security, and availability required for demanding, mission-critical workloads. For example, self-healing capabilities enable continued operation even in the event of component failures. Machine Check Architecture Recovery (MCA Recovery) technology works with the host operating system or virtual machine monitor (VMM) to enable automatic recovery from certain uncorrectable errors that would have caused previous-generation servers to crash. Intel Xeon processor E7 platforms add support for Double Data Device Correction (DDDC), which facilitates recovery from two DRAM device failures, and Partial Memory Mirroring, which enables mirroring of only critical portions of memory to improve reliability while controlling costs. Support for Advanced Encryption Standards–New Instructions (AES-NI) enhances security by significantly reducing the performance penalties usually experienced with pervasive encryption. These and other features help organizations protect data and meet service-level agreements while refocusing resources on innovation.⁷

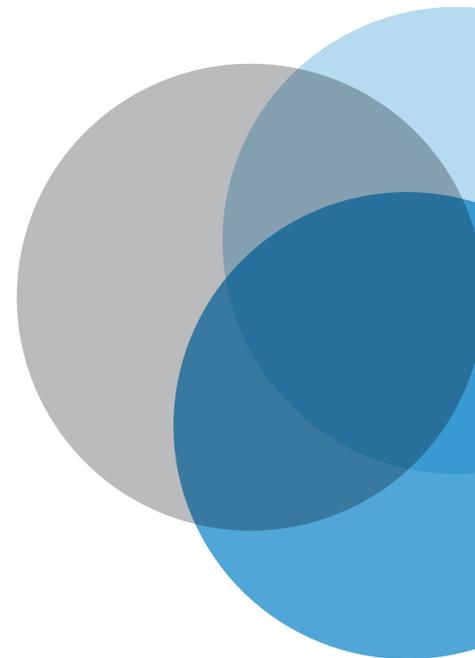
IBM eX5 servers take advantage of these processors to provide exceptional performance and scalability for mission-critical databases. This fifth generation of the IBM X-Architecture® platform represents the biggest IBM investment ever in Intel Xeon processor-based servers. Servers are available in 2- and 4-socket configurations. Organizations can scale processor and memory resources independently in each server. They can also connect a pair of 4-socket IBM System x3850 X5 servers with a high-performance link to create a single system with as many as 8 sockets, 80 cores, 160 threads, and up to 6 TB of memory using external IBM MAX5 memory expanders. A fully configured IBM eX5 system provides ample horsepower for hosting heavy, data-intensive workloads, and organizations have considerable flexibility for tailoring their servers to match their workloads.

IBM eX5 servers offer more than just scalable performance. They also deliver major enhancements to improve data integrity and system resilience. IBM supports and extends the silicon-based reliability, availability, and serviceability (RAS) features of the Intel Xeon processor E7 platforms with advanced error management, node partitioning, and automatic failover to improve utilization and uptime for

mission-critical workloads. Additional features such as Memory ProteXion™ enable the servers to exceed industry-standard reliability.

Servers based on this architecture also include features such as integrated management modules and advanced light-path diagnostics to provide sophisticated monitoring and proactive problem resolution. Redundant power supplies help eliminate a key source of potential downtime. And if issues should arise, easily accessible components help speed time to repair.

The latest version of IBM DB2 software has been tuned to take advantage of Intel Xeon processor enhancements. It also provides valuable software features such as new autonomic management tools that help to improve performance and reduce total cost. New compression capabilities in IBM DB2 software are especially valuable for large database implementations: this feature can reduce storage requirements by as much as 80 percent to help lower storage costs and improve performance. Since more data can be kept in system memory, there is less need for performance-robbing disk I/O operations. Processors have faster access to data, which improves utilization to help reduce transaction latencies and increase throughput.





High-Performance Clustering for Near-Limitless Scalability

A fully configured IBM eX5 server provides a powerful and resilient database engine. Add the DB2 pureScale feature and these servers become building blocks that can be combined without limit to support truly massive data volumes and application workloads with exceptional performance and high availability. Importantly, adding a server to the cluster is a simple process. Unlike other database clustering solutions, there is no need to repartition the database, manually rebalance workloads, or adapt applications. The technology that enables this near-limitless scalability is not altogether new. It has been at work in IBM mainframes for many years, helping those systems support

enormous workloads with the very highest levels of reliability. Continued advancements in industry-standard processor, server, and networking technologies have enabled IBM to port this technology into its Intel processor-based X-Architecture servers—something that was not feasible until these recent advances.

DB2 pureScale software employs a shared-disk architecture that gives every server in the cluster direct access to all data (see Figure 2). That architecture enables any server to process any request, which helps simplify workload balancing. It also helps ensure that if one or more servers fail, the remaining servers are not cut off from any portion of the data—enabling normal operations to continue without disruption.

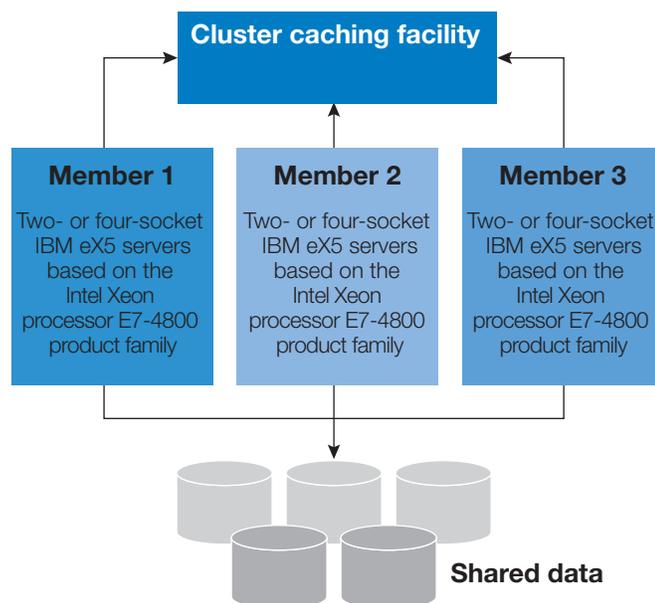


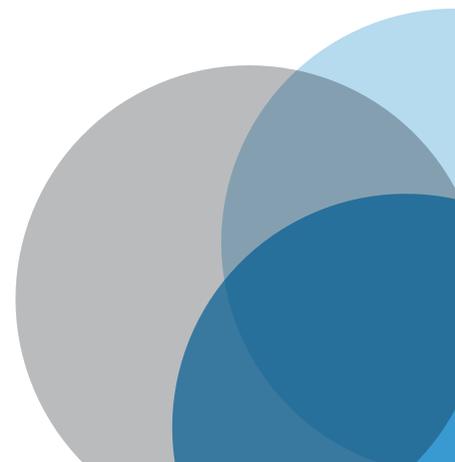
Figure 2. IBM DB2 pureScale and Intel Xeon processors enable organizations to achieve near-limitless scalability, high availability, and complete application transparency on affordable, industry-standard servers.

Every server in a DB2 pureScale cluster also has direct, point-to-point access to a centralized data caching and locking facility through a high-speed InfiniBand* network. This centralized cache is the “secret sauce” IBM has ported from its mainframe architecture and is the key to enabling low-latency database performance in a clustered server environment. When an individual server accesses data to perform a transaction, the data page is stored not only in the server’s local memory and cache subsystems but also in the centralized cache. From this repository, data and locking information can be accessed almost instantly by any other server in the cluster.

To make this architecture work, data transfers among the clustered servers and the centralized cache must be very fast—fast enough that little or no latency is introduced into database transactions. To meet this requirement, DB2 pureScale software takes advantage of Remote Direct Memory Access (RDMA). This lightweight memory access protocol allows each member in the cluster to directly access the memory of the centralized cache and vice versa. It helps dramatically reduce the communications overhead required to request and transfer data, and it enables each server to quickly access and deliver data and other critical information to the centralized cache. As a result, the centralized cache reflects the activities of each server in the cluster almost instantaneously.

Every server has high-speed access not only to all data in the shared storage system but also to all cached data and locking information. The high-speed, centralized caching and locking facility virtually eliminates the traditional performance challenges of distributed database solutions. In a traditional cluster, when a server receives a request, it has to search for cached data among other servers in the cluster. Older clustering solutions use distributed caching and locking mechanisms, leading to time-consuming searches and introducing latencies that slow performance. As the cluster grows, the searches become more complex and the latencies increase. By eliminating these complex and time-consuming operations, DB2 pureScale software enables fast performance and linear scalability even for very large clusters.

Organizations can now use affordable IBM System x servers to support data volumes and application workloads of any size. As data volumes and application workloads grow, they can simply add additional servers to the cluster to scale capacity. Organizations no longer need to replace entire servers or disrupt operations. In addition, they can build their clusters using the servers that best match their workloads and service-level requirements. Organizations gain the flexibility to optimize capability versus total cost.





Application Transparency Helps Eliminate Hidden Costs

IBM DB2 pureScale requires no application changes, database partitioning or manual workload balancing to maintain optimized performance. An application running on a single, large server can be moved onto a DB2 pureScale cluster quickly and easily. New servers can be added to the cluster at any time, all with no code changes and very little administrative overhead.

This is not the case for a traditional database cluster that relies on distributed caching and locking mechanisms. With these traditional clusters, each server in the cluster must be as close as possible to the data it needs to minimize data access latencies. Satisfying this requirement introduces a great deal of

administrative overhead. The database must be partitioned, workloads must be balanced, and applications must be made “cluster-aware” so that they can interface efficiently with this complex environment. These are time-consuming processes that require a significant level of expertise. They are also ongoing requirements because the demands of data partitioning, workload balancing, and application awareness all change as the data, applications, and hardware infrastructure evolve. Changes to applications are particularly resource-intensive, since they typically require extensive testing and validation before they can be implemented in mission-critical production systems.

Turning Exploding Data Volumes into High-Value Business Insight

The IBM Smart Analytics System on Intel processor-based IBM System x servers

IBM and Intel are focused on delivering comprehensive, high-value information solutions for companies looking to maximize the total value of their data. On-demand scalability for transactional databases with IBM DB2 pureScale software is a first step. The IBM Smart Analytics System provides the next step by delivering comparable value for business analytics. These systems provide IT organizations with a complete, fully integrated business intelligence solution that can be deployed with the ease of an appliance to achieve value in days instead of months, and then expanded easily as needs change.

The IBM Smart Analytics System includes a powerful warehouse foundation and broad analytic capabilities, all running on affordable Intel Xeon processor-based IBM System x servers. These systems integrate easily with existing data sources to provide advanced multidimensional Cubing Services, data mining, text analytics, dashboards, and reporting. The tools are powerful, easy to use, and customizable. They help ensure that users across the enterprise are always acting on the most timely, useful, and trusted information available, enabling them to make smarter, faster business decisions.

IBM DB2 pureScale software eliminates these requirements. With the software’s centralized caching and locking facility, every server is equally close to all data, so there is no need to partition the database. Any server can perform any transaction. This capability makes it relatively easy to balance workloads across the cluster—and DB2 pureScale performs this function automatically.

Just as important, the centralized caching and locking facility enables DB2 pureScale to be completely transparent to applications. No changes are required to run an application on a DB2 pureScale cluster, and no special tuning is required to optimize performance as the cluster

grows. Scaling performance and capacity is quick and simple. A new server can be integrated with just two commands, and workloads are automatically rebalanced to help optimize performance (see Figure 3). Adding and removing servers from a cluster through the cluster caching facility (CF) is easy enough that IT can even scale clusters up or down in response to seasonal or monthly workload fluctuations. Using two CF servers can help eliminate a single point of failure with duplexing. IBM offers licensing policies that can help businesses take advantage of this flexibility to achieve significant cost savings.

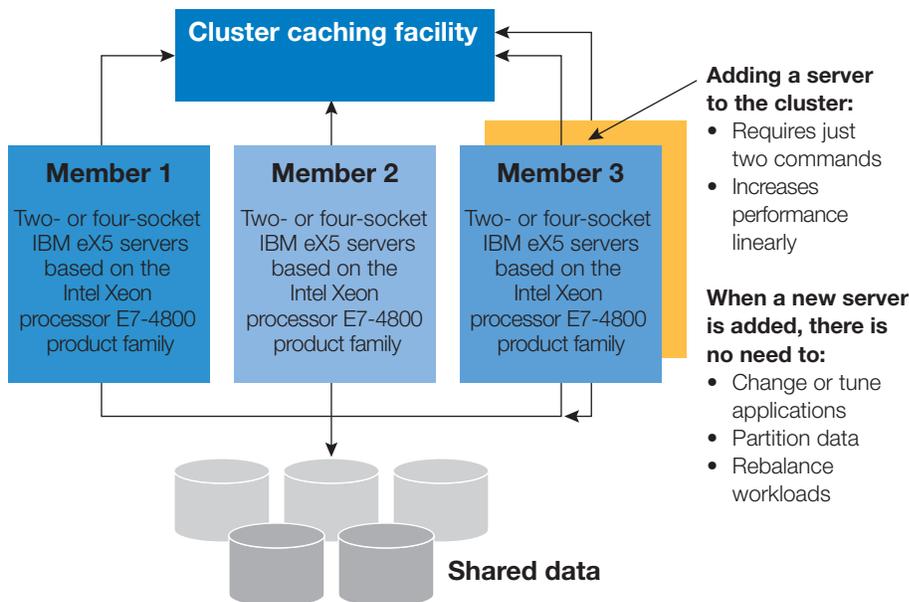


Figure 3. A server can be quickly and easily added to a DB2 pureScale cluster through the cluster caching facility.



High Availability for Uninterrupted Operations

IBM DB2 pureScale software adds yet another layer of high availability to the silicon- and system-level RAS features that are built into eX5 servers (see Figure 4). Since every member of a DB2 pureScale cluster is active and has access to all data, whether cached or stored, a failed server affects only transactions that are currently in process on that particular server. The remaining servers continue operating as usual, and they can instantly take on any new requests that would have gone to the failed system.

As in any transactional database, transactions that are interrupted when a server fails must be either recovered or rolled back. A partially completed transaction would not only corrupt the transaction itself but also compromise the integrity of the data. This is another area where the DB2 pureScale centralized caching and locking technology delivers tremendous value. With a centralized information repository, DB2 pureScale software is always aware of the status of every transaction. It also knows which data pages are “dirty”—that is, which pages contain data that has been changed by a transaction but not yet recorded to disk.

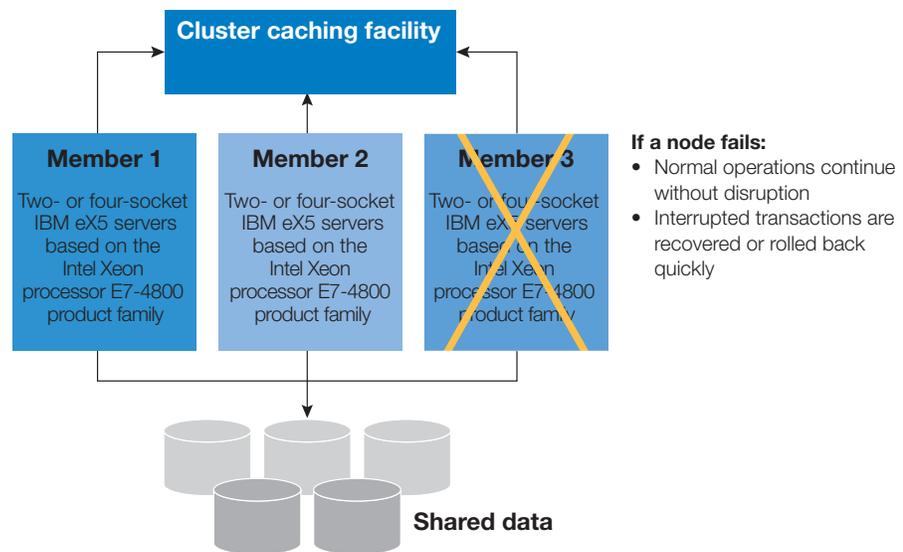


Figure 4. DB2 pureScale seamlessly maintains normal operations following a server failure by leveraging advanced RAS features.

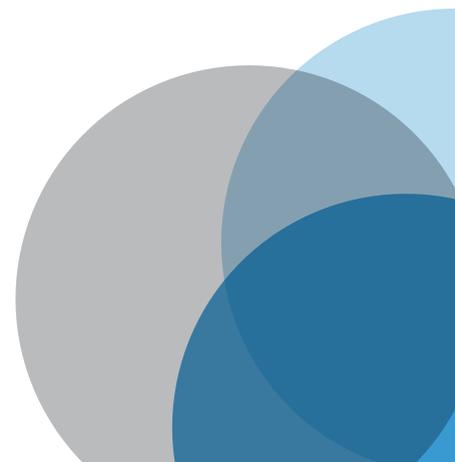
Following a server failure, therefore, DB2 pureScale software always has all the information needed to:

- **Maintain normal operations without interruption.** The server failure affects only interrupted transactions and dirty data pages directly associated with those transactions. With the distributed caching and locking architecture of a traditional cluster, there is a delay while individual servers share information to determine which data pages are dirty. Until those issues are fully resolved, no data page can be safely used because of the risk of using old data in new transactions. DB2 pureScale software eliminates this delay.
- **Recover or roll back interrupted transactions quickly.** Because of its centralized caching and locking facility, DB2 pureScale software knows which transactions need to be recovered or rolled back without accessing additional information from any other servers in the cluster. Each server also maintains recovery processes that are sitting idle but are ready in the event of a failure. As a result, there is no need for the operating system to create a process and allocate memory to it. Since no information has to be shared and recovery processes can be instantly activated, complete recovery of all interrupted transactions is exceptionally fast. Full recovery typically takes less than 20 seconds, including the time needed to detect the failure. That detection time is not always included in the database recovery specifications of other vendors.

DB2 pureScale software not only maintains normal operations without interruption but also continues to deliver optimized performance by automatically rebalancing workloads across the remaining servers in the cluster. As soon as the failed server is restored, workloads are again automatically rebalanced.

To help ensure there are no single points of failure in a cluster, the DB2 pureScale software components are always hosted on at least two physical servers. If one server fails, the other takes over instantly to maintain normal operations. With this approach, a DB2 pureScale cluster provides exceptionally high availability. Traditional single-server database implementations, in contrast, typically require an expensive, high-end hot standby server to provide comparable service-level assurances.

In addition, because DB2 pureScale software automatically rebalances workloads as servers are added or removed from the cluster, taking a server down for physical maintenance is both simple and nondisruptive. Hardware and software remain fully operational, and performance is dynamically optimized based on available hardware resources.





Servers to Match Specific Needs

IBM DB2 pureScale software can be used with any IBM System x servers in 2-, 4-, or 8-socket configurations based on the Intel Xeon processor E7-2800, E7-4800, or

E7-8800 product families, respectively.⁸ The choice will depend on workloads, expected growth rates, and service-level requirements.

IBM eX5 Servers: High Value for Large, Rapidly Growing Databases

IBM eX5 servers are the most powerful and robust System x servers. They are ideal for large to extremely large databases where response times and uptime are critical. Based on the Intel Xeon processor E7 product family, these servers provide up to 10 cores, 20 threads, and 30 MB of cache per socket, and support a number of advanced, mainframe-inspired RAS features.

Two server models are supported in IBM DB2 pureScale clusters:

- **IBM x3850 X5 servers** are ideal for the very heaviest workloads and most demanding latency requirements. Each server can be configured with up to four processors from the Intel Xeon processor E7-4800 product family or eight processors from the Intel Xeon

processor E7-8800 product family. In its 4-socket configuration, the server can accommodate up to 3 TB of memory with an optional IBM MAX5 external memory chassis. Two servers can be connected with a high-speed link to double the available resources.

- **IBM x3690 X5 servers** provide the robust capabilities of the eX5 architecture in 2-socket configurations that enable more granular scaling of an IBM pureScale cluster. Each server can be configured with two processors from the Intel Xeon processor E7-2800 product family and up to 512 GB of memory. An optional IBM MAX5 external memory chassis can be added to increase memory capacity to 1 TB.⁹

Conclusion

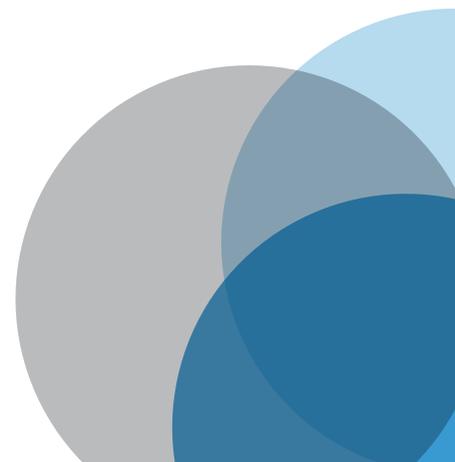
The high core counts and the large cache and memory configurations supported by IBM eX5 servers provide exceptionally high computing density and excellent computing efficiency for the demanding workloads that are common in DB2 implementations. Their advanced RAS features also help to provide enhanced data integrity and very high availability per node. As a result, IBM eX5 servers deliver excellent performance, enormous capacity, and extremely high availability in small clusters designed to minimize space, power, cooling, and administrative costs.

To learn more about the IBM eX5 Portfolio Overview, visit <http://publib-b.boulder.ibm.com/Redbooks.nsf/RedbookAbstracts/redp4650.html?OpenDocument> or the IBM Web site at ibm.com/systems/info/x86servers/ex5.

To learn more about the Intel Xeon processor E7 platform, visit intel.com/itcenter/products/xeon.

IBM eX5 servers based on the Intel Xeon processor E7 platforms can deliver the performance, scalability, and mission-critical capability for data-intensive enterprise workloads. These servers offer a powerful, high-value hardware infrastructure for IBM DB2 implementations, with affordable starting points, incremental scalability, and an array of mainframe-inspired features to support demanding requirements for high availability and data integrity.

With the IBM DB2 pureScale feature, customers can combine a large number of these powerful servers into a high-performance cluster that delivers virtually unlimited scalability with high availability and complete application transparency. By eliminating the challenges that have plagued database clustering solutions from other vendors and supporting affordable Intel Xeon processor-based servers, IBM DB2 pureScale software changes the paradigm for deploying and growing large, mission-critical databases. As increasing volumes of data flow into the enterprise, businesses can now scale their infrastructure on demand and without limit, enabling them to store and use their data far more effectively and at a much lower cost.





Software and workloads used in performance tests may have been optimized for performance only on Intel® microprocessors. Performance tests, such as SYSmark® and MobileMark®, are measured using specific computer systems, components, software, operations, and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products.

¹ An IBM System x3850 X5 server can be configured with up to 2 TB of memory without an optional external IBM MAX5 memory chassis and up to 3 TB with one.

² Estimated SPECint*_rate_base2006 performance results compare an Intel Xeon processor E7-4870 and a single-core Intel Xeon processor. For more information, visit www.intel.com/performance/server.

³ Source: Dave Cappuccio, research vice president at Gartner, in a keynote speech at Gartner IT Infrastructure, Operations & Management Summit 2009. www.gartner.com/it/content/750200/750233/iom_2009_trip_report.pdf

⁴ Average of 3x performance claim based on geometric mean of four industry-standard, common enterprise benchmarks (SPECjbb*2005, SPECint*_rate_base2006, SPECfp*_rate_base2006, and TPC Benchmark* E) comparing best published/submitted results on four-socket (4S) Intel® Xeon® processor X7560-based server platform to best published 4S Intel Xeon processor X7460-based server platform as of March 26, 2010.

⁵ Claim: “Up to 20x performance per server” Disclaimer: Intel performance comparison using SPECjbb2005* business operations per second between five-year-old single-core Intel® Xeon® processor 3.33GHz-based servers and one new Intel Xeon processor X7560-based server. Performance tests and ratings are measured using specific computer systems and/or components and reflect the approximate performance of Intel products as measured by those tests. Any difference in system hardware or software design or configuration may affect actual performance. Buyers should consult other sources of information to evaluate the performance of systems or components they are considering purchasing. For more information, visit www.intel.com/performance/server.

⁶ Up to 40% generational compute-intensive throughput claim based on SPECint*_rate_base2006 benchmark comparing Intel Xeon processor E7-4870 (30 MB cache, 2.40 GHz, 6.40 GT/s Intel QPI, formerly codenamed Westmere-EX) scoring 1,010 (includes Intel Compiler XE2011 improvements accounting for about 11% of the performance boost) to Intel Xeon processor X7560 (24 MB cache, 2.26 GHz, 6.40 GT/s Intel QPI, formerly codenamed Nehalem-EX) scoring 723 (Intel Compiler 11.1). Source: Intel SSG TR#1131. For more information, visit www.intel.com/performance/server.

⁷ Some features require support and validation from the server manufacturer and may not be provided in all server platforms. Intel® AES-NI requires a computer system with an AES-NI enabled processor, as well as non-Intel software to execute the instructions in the correct sequence. AES-NI is available on select Intel® processors. For availability, consult your reseller or system manufacturer. For more information, see <http://software.intel.com/en-us/articles/intel-advanced-encryption-standard-instructions-aes-ni/>.

⁸ The 8-socket configuration is possible when linking together two servers.

⁹ These memory capacity specifications are based on the use of 16 GB memory modules. As 32 GB memory modules become available, the potential memory capacity will double.

Copyright © IBM Corporation 2011
Copyright © 2011 Intel Corporation

Produced in the United States of America
March 2011
All rights reserved

IBM, the IBM logo, ibm.com, DB2, pureScale, solidDB, System x, and X-Architecture are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries or both. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or TM), these symbols indicate U.S. 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 “Copyright and trademark information” at ibm.com/legal/copytrade.shtml.

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

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