



Contents

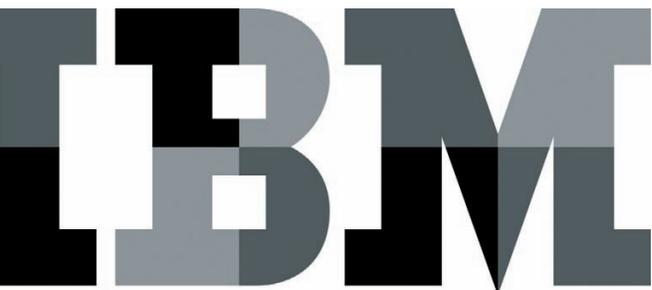
- 1 Introduction
 - 2 The solution
 - 3 Test results
 - 6 Solution component details
 - 7 Conclusion
 - 8 For more information
-

IBM FlashSystem accelerates IBM DB2 workloads

Introduction

IT systems administrators in enterprises around the globe are constantly challenged by the deluge of data they are expected to manage. How can the growing data volumes between different processing clusters, departments and sites be shared most effectively? How can the explosion of *unstructured data* be managed? As shown in Figure 1, data volumes and velocities are dramatically increasing from every direction and every source, while IT budgets and staff remain static or are slashed. The proper tools to diagnose the root cause of performance issues are lacking and, what's more, IT staff must begin deploying analytic applications in multiple environments.

High-performance, cost-effective and easy-to-deploy solutions are available from IBM. IBM® Power Systems™ provide the processing capabilities needed to scale out or scale up while powering the most complex analytics. IBM FlashSystem® storage offers performance, efficiency and reliability, as well as the virtualization tools that enable enterprises to turn the explosion of data into insight for business advantage. These IT infrastructure components provide the foundation upon which to build a next generation data platform—using IBM DB2®—that keeps transactional workflows and analytics operating at maximum efficiency.



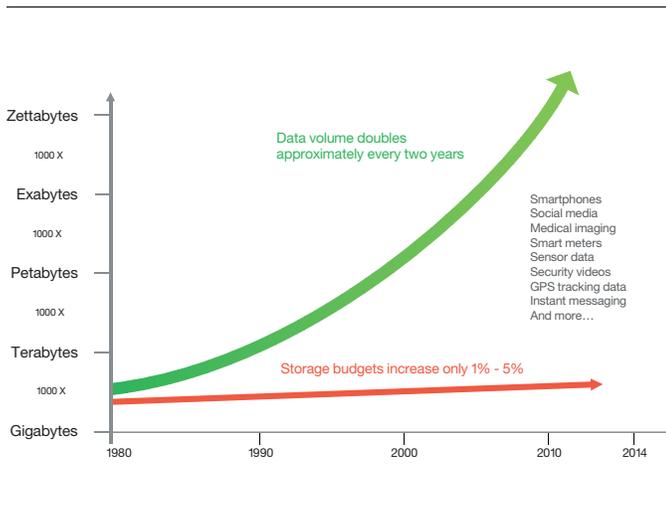


Figure 1. Storage requirements are devouring CAPEX and OPEX resources

The solution

The solution described in this paper employs four IBM products known for their ability to achieve extraordinary levels of performance and scalability:

- Power Systems servers
- IBM FlashSystem storage
- IBM AIX® operating system
- DB2 database

A set of workload tests were conducted to gauge the performance and scalability of the IBM products. The server employed in the test, the Power System S822, comprised two IBM POWER8™ processors, each with ten cores. Designed for big data, the POWER8 processor has industry-leading

performance capabilities, including 230 GB per second of memory bandwidth, 96 GB per second of inputs and outputs (I/O) bandwidth, and eight hardware threads per core. As a result, POWER8 servers have delivered industry-leading performance across a variety of domains, including SAP,¹ integer and floating point,² and application-serving.³

The storage employed in the evaluation was IBM FlashSystem. IBM FlashSystem delivers breakthrough performance compared to other solid-state disk (SSD) arrays on the market today, including achieving super-low response time through the use of IBM MicroLatency® technology and over 500,000 inputs and outputs per second (IOPS) per rack unit. The IBM FlashSystem design excels in real-world mixed workloads at high-demand rates, at high-capacity utilization and with long-term sustainable rates, well beyond just the initial burst performance that most systems can achieve.

AIX, with capabilities such as multipath I/O, asynchronous I/O and concurrent I/O, is the glue that combines IBM POWER® and IBM FlashSystem technology, so that massive amounts of data can be stored, retrieved and processed at unprecedented speed. The software that brings it all together is DB2 for Linux, UNIX and Windows. DB2 helps customers reach new levels of availability and scalability, while enabling extreme workload capacity and delivering a unique combination of performance, ease of use and cost-efficiency.

For additional information on the combination of Power Systems and IBM FlashSystem technology, see “FlashSystem: Ideal Storage for Power Systems,” available on the web at ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP102506

Workloads tested

1. *Brokerage workload* replicates the daily tasks of a brokerage firm database. This workload includes customer input such as trades, market analysis and account inquiries, and the in-turn action of executing the trades and updating the relevant account information.
2. *Transactional workload* simulates an environment in which a pool of users executes transactions against a database. The transactions could include entering and fulfilling orders, accepting payments, inquiring into the status of orders and overseeing inventory levels.
3. *Big data insights analytics workload* models a day in the life of a business intelligence (BI) application. The workload is based on a retail database with in-store, online and catalog sales of merchandise.

Configuration details

The three workloads (brokerage, transactional and big data insights analytics) were tested using the following configuration:

Table 1: *Workload testing configuration details*

Component	Details
Database	DB2 version 10.5 FP4
Server	Power System S824 (20 cores, 512 GB, eight 8 Gb Fibre Channel ports)
Operating system	AIX 7.1 TL3 SP3
Storage	IBM FlashSystem 840 (eight 4 TB flash modules, 24 TB RAID 5) <i>compared to</i> a disk storage enclosure (24 900 GB 10K SAS drives, 17.2 TB RAID 5)

Test results

Brokerage workload

The number of transactions per second (tps) was measured as well as IOPS, response time and CPU utilization. The database size needed to be approximately 20 TB, so both SAN devices used in the comparison (IBM FlashSystem and disk storage enclosure) and were sized similarly for capacity.

- IBM FlashSystem storage achieved more than 40 times better performance than the disk storage enclosure, for both transactions per second and IOPS. That translates to a 90 percent price-per-performance improvement.
- IBM FlashSystem drove over 272,000 IOPS, had faster than a 0.5 ms response time (for both read and write operations) and performed over 13,000 tps, utilizing 90 percent of CPU capacity (bottlenecked by CPUs).
- The disk storage enclosure performed only 6,000 IOPS, had a 66 ms read response time and performed only 316 tps, utilizing 2 percent of CPU capacity (bottlenecked by the disk storage enclosure).
- To reach the level of performance of one IBM FlashSystem module, a customer would need 42 disk storage enclosures, which would fill two rack units, require an order of magnitude more power and cost over five times as much.

Table 2: *Results of brokerage workload testing*

Storage	IOPS	Response time	Transactions per second	CPU utilization	Bottleneck
IBM FlashSystem storage	272,000	Less than 0.5 ms	13,000	90%	Server
Disk storage enclosure	6,000	66 ms	316	2%	Disk storage enclosure

Transactional workload

The number of transactions per minute (tpm) was measured as well as IOPS, response time and CPU utilization. The hardware configuration remained the same.

- IBM FlashSystem storage achieved 190 times better performance than the disk storage enclosure for transactions per minute and 86 times better for IOPS. That translates to a 95 percent price-per-performance improvement.
- IBM FlashSystem drove over 344K IOPS, had faster than a 0.4 ms response time (for both read and write operations) and performed over 570K tpm, utilizing 70 percent of CPU capacity.
- The disk storage enclosure performed only 4K IOPS, had a 52 ms response time, and performed only 3K tpm, utilizing 1 percent of CPU capacity (bottlenecked by the disk storage enclosure).

Table 3: *Results of transactional workload testing*

Storage	IOPS	Response time	Transactions per second	CPU utilization	Bottleneck
IBM FlashSystem storage	344,000	Less than 0.4 ms	570,000	70%	Server
Disk storage enclosure	4,000	52 ms	3,000	1%	Disk storage enclosure

As shown in Figure 2, the transactions per minute for both IBM FlashSystem and the disk storage enclosure remained constant, while the number of clients increased. Note that the number of transactions per minute was so much larger for IBM FlashSystem than it was for the disk storage enclosure that the y axis is a log scale.

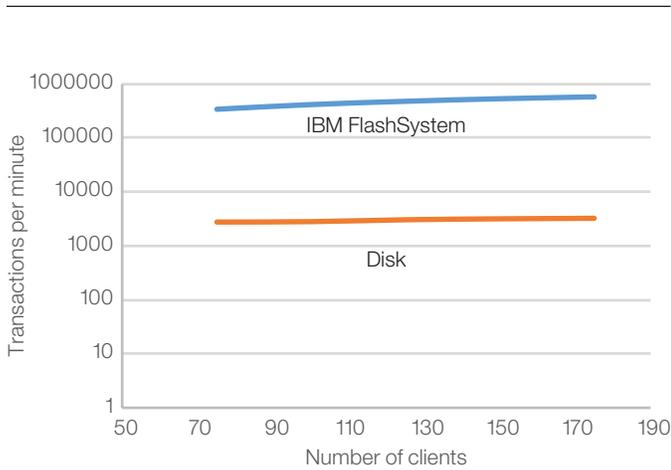


Figure 2. In tests, IBM FlashSystem performed 190 times more transactions per minute than disk

Table 4: Results of big data insights analytics workload testing

Storage	Buffer pool	Intermediate queries per hour	Complex queries per hour	CPU utilization	Bottleneck
IBM FlashSystem storage	25 GB	605	37	92%	Server
Disk storage enclosure	25 GB	85	14	24%	Disk storage enclosure

Big data insights analytics workload

The CPU utilization was measured, along with the number of intermediate and complex queries performed per hour.

- IBM FlashSystem performed 7 times better than the disk storage enclosure for intermediate queries and 2.5 times better for complex queries.
- IBM FlashSystem still had significant headroom and was bottlenecked by the server (at 92 percent CPU utilization), while the disk storage enclosure bottlenecked the query performance (at only 24 percent CPU utilization).
- IBM FlashSystem is far less sensitive to the buffer pool size and performance remains constant whether there is a small (500 MB) or average (25 GB) buffer pool size. Clients should experience easier setup and tuning with IBM FlashSystem.
- The big data insights analytics workload was used as-is, without modification or additional components, with minimal tuning. Only two parameters were changed (25 GB threshold = 10, degree = any; 500 MB threshold = 10, degree = 10).

Solution component details

IBM FlashSystem

The IBM FlashSystem family is a suite of enterprise-class all-flash storage platforms that are ideal for delivering extreme performance with low latency for database online transaction processing (OLTP) environments. IBM FlashSystem arrays can deliver one million or more read I/Os per second at hundred-microsecond latency levels, while providing dozens of terabytes of usable multidimensional Flash RAID-protected data storage in just a few rack units of space. IBM FlashSystem storage systems have enterprise-level availability and reliability with no single point of failure, multiple layers of data correction, chip redundancy and redundant hot-swap components.



Figure 3. The IBM FlashSystem storage array

Extreme performance, extremely low latency, macro-efficiency and enterprise-grade reliability make IBM FlashSystem a powerful and cost-effective tool for accelerating OLTP systems, meeting and exceeding online customer expectations and gaining competitive advantage in DB2 database environments. IBM FlashSystem performance also powers online analytics processing (OLAP) tools that help businesses better understand and more swiftly respond to customer needs and preferences.

Just as importantly, the extraordinary capabilities and capacity of IBM FlashSystem arrays enable enterprises to address multiple compute challenges in current 24x7 operational environments, while at the same time empowering growth and innovation into the future.

IBM DB2

IBM DB2 for Linux, UNIX and Windows is a next-generation data platform for transactional and analytical operations. It provides continuous availability of data to keep transactional workflows and analytics operating at maximum efficiency. And, it delivers breakthrough in-memory performance, enabling speed-of-thought analytics without the constraints of other in-memory solutions, yet with the simplicity of “load and go” setup. DB2 is optimized to deliver industry-leading performance while lowering costs and improving IT productivity. DB2 enables you to:

- Automate tasks such as memory allocation, storage management, health monitoring and business policy maintenance to improve IT productivity.
- Provide high-performing predictive and pattern analysis without moving your data into proprietary data-mining platforms.
- Ensure your database remains operational during planned or unplanned outages.
- Reduce downtime to help meet strict service-level agreements (SLAs), with no loss of data during infrastructure failures.
- Save money, lower risks and improve IT productivity by easily optimizing performance.

- Increase database administrator productivity.
- Achieve high availability and dynamic scalability for transaction processing.
- Adapt to changing business needs by allocating more resources for performance or for growth, at will.
- Reduce storage costs and increase performance with industry-leading technologies that compress tables, indexes, archive logs, temporary space, LOBs, XML and backup data.

IBM Power Systems

Designed to extract the greatest value from an enterprise's data, Power Systems deliver the foundation for organizations to bring insight to the point of impact. These first-generation systems push the physical and virtual boundaries of data center technology with innovation purpose-built to drive faster and more efficient data-centric applications that are required for today's smarter enterprise.

With new innovations, Power Systems provide the ability to:

- Gain faster insights with the POWER8 processor and smart acceleration enabled by Coherent Accelerator Processor Interface (CAPI) technologies.
- Achieve lower latency and smaller footprint with CAPI Flash.
- Move data in and out of systems more quickly with twice the memory and I/O expansion.
- Achieve greater speed and efficiency for database, OLTP and other highly multi-threaded applications, with transactional memory supported by 50 percent more cores and twice the number of simultaneous threads per core.

Power Systems are ideal for consolidation of multiple applications and infrastructure workloads in a virtualized environment, bringing together business-transaction processing with infrastructure for big data, analytics and OLTP solutions.



Figure 4. The IBM Power System S824

Conclusion

Performance testing clearly demonstrated the superiority of IBM FlashSystem storage to traditional disk storage enclosures. Whether the tests conducted looked at I/O read and write speeds, transaction volumes per second and per minute, CPU utilization or latency levels, IBM FlashSystem storage outperformed its disk counterpart.

For more information

To learn more about the IBM FlashSystem family of storage products, IBM DB2 database solutions and IBM Power Systems software subscription and support, please contact your IBM representative or IBM Business Partner, or visit the following websites:

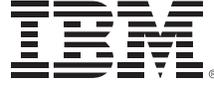
- ibm.com/systems/storage/flash
- ibm.com/software/data/db2
- ibm.com/systems/power

Additionally, IBM Global Financing can help you acquire the IT solutions that your business needs in the most cost-effective and strategic way possible. We'll partner with credit-qualified clients to customize an IT financing solution to suit your business goals, enable effective cash management, and improve your total cost of ownership. IBM Global Financing is your smartest choice to fund critical IT investments and propel your business forward. For more information, visit: ibm.com/financing

¹ IBM Power System S824 on the two-tier SAP SD standard application benchmark running SAP enhancement package 5 for the SAP ERP 6.0 application; 4 processors/24 cores/96 threads, POWER8; 3.52 GHz, 512 GB memory, 21,212 SD benchmark users, running AIX 7.1 and DB2 10.5, dialog response: 0.98 seconds, line items/hour: 2,317,330, dialog steps/hour: 6,952,000 SAPS: 115,870 database response time (dialog/update): 0.011 sec/0.019sec, CPU utilization: 99%, Certification #: 2014016. Source: www.sap.com/benchmark. (1.1) Fujitsu RX300 S8 on the two-tier SAP SD standard application benchmark running SAP enhancement package 5 for the SAP ERP 6.0 application; 2 processors/24 cores/48 threads. Intel Xeon E5-2697 processor 2.70 GHz, 256 GB memory, 10,240 SD benchmark users, running Windows Server 2012 SE and SQL Server 2012, Certification #: 2013024

² IBM Power S824 (3.5 GHz, 24 core) results. PowerEdge R920 (Intel Xeon E7-8893 v2, 3.40 GHz) results. Source: www.spec.org

³ IBM WebSphere® Application Server V8.5.5.2 and DB2 10.5 on IBM Power S824 result of 22,543.34 published on Apr 22, 2014. Oracle Weblogic Server Standard Edition Release 12.1.2 and Oracle Database 12c on Oracle Sun Server X4-2 result of 11,259.88 published on Sep 23, 2013. Source: www.spec.org



© Copyright IBM Corporation 2015

IBM Systems
Route 100
Somers, NY 10589

Produced in the United States of America
February 2015

IBM, the IBM logo, ibm.com, AIX, DB2, IBM FlashSystem, POWER, POWER8, Power Systems, and WebSphere are trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at "Copyright and trademark information" at ibm.com/legal/copytrade.shtml

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

Windows is a trademark of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

This document is current as of the initial date of publication and may be changed by IBM at any time. Not all offerings are available in every country in which IBM operates.

The performance data discussed herein is presented as derived under specific operating conditions. Actual results may vary. It is the user's responsibility to evaluate and verify the operation of any other products or programs with IBM products and programs.

THE INFORMATION IN THIS DOCUMENT IS PROVIDED "AS IS" WITHOUT ANY WARRANTY, EXPRESS OR IMPLIED, INCLUDING WITHOUT ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND ANY WARRANTY OR CONDITION OF NON-INFRINGEMENT. IBM products are warranted according to the terms and conditions of the agreements under which they are provided.

Actual available storage capacity may be reported for both uncompressed and compressed data and will vary and may be less than stated.



Please Recycle