

# Improving I/O Performance Using Solid-State Drive Technology



## Applies to:

IBM Storage, SAP on DB2 for z/OS, SAP on IBM i

For more information, visit the [Performance homepage](#).

## Summary

Solid-State Drive (SSD) technology is an important new technology that is being introduced into IBM System Storage solutions to significantly improve performance and reduce energy consumption. Measurements have shown that combining IBM System Storage with SSD, high-speed fiber connectivity, and DB2 can provide more options for improved system responsiveness and throughput for SAP applications.

**Author:** Dr. Heiko Gerwens, SAP AG, Germany  
Rainer Pfister, SAP AG; Germany  
Wolfgang Reichert, IBM Deutschland Research & Development GmbH  
Mike Sheets, IBM USA, Advanced Technical Support  
Johannes Schützner, IBM Deutschland Research & Development GmbH  
Steve Tlusty, IBM USA, IBM i ERP Development

**Company:** SAP AG

**Created on:** 20 July 2009

## Author Bio



Heiko Gerwens is the leading performance expert for SAP financial service applications. He joined SAP in 1997 and has worked in the performance and data management team. During this time he has been involved in the development of many applications including customer projects that implement these applications. He supports customers, partners, consulting agents, and the field organization that deals with all aspects of performance, scalability, and sizing. He can be reached at [Heiko.Gerwens@sap.com](mailto:Heiko.Gerwens@sap.com).



Rainer Pfister is a development architect in the SAP database platform organization. He studied computer science at Mannheim University of Applied Sciences and joined SAP in 1987. Since then he has worked on various SAP kernel and communication development projects. Rainer is responsible for System z porting projects. In his role as advocate of these projects, he is the central contact for cross-platform-related issues of third-party products and acquired companies. He can be reached at [rainer.pfister@sap.com](mailto:rainer.pfister@sap.com).



Wolfgang Reichert is an IBM senior technical staff member. Since 1995 he has been working with SAP in supporting SAP business solutions on IBM systems. As lead architect he is responsible for SAP-related development across all IBM platforms. Wolfgang studied physics at the University of Duisburg, Germany. He is a member of the IBM System z Technology Center for SAP Applications. He can be reached at [w\\_reichert@de.ibm.com](mailto:w_reichert@de.ibm.com)



Mike Sheets is an executive information technology specialist based in Poughkeepsie, NY, USA. Since 1996 he has been working with SAP on the System z project and its customers. Prior to that, he worked in the Poughkeepsie development laboratory on IBM's operating systems and hardware products – in particular, their performance. He studied chemistry at the University of Iowa and computer science at the University of Rhode Island. He can be reached at [msheets@us.ibm.com](mailto:msheets@us.ibm.com).



Johannes Schuetzner divides his time between the IBM Boeblingen Lab and SAP Labs in Walldorf and St. Leon-Rot, and works on all aspects of the use of IBM DB2 for z/OS for SAP applications. Johannes studied computer science at the University of Stuttgart and the University of Connecticut at Storrs. He belongs to the IBM System z Technology Center for SAP Applications and is a member of the IBM Technical Expert Council. He can be reached at [schuetzner@de.ibm.com](mailto:schuetzner@de.ibm.com).



Steve J. Tlusty is a member of the SAP on System i Team, as part of i5/OS Development in Rochester. He has eight years of experience with SAP Business Intelligence. Before joining the SAP Team, Steve worked in IBM Global Services developing business intelligence solutions for retail and manufacturing customers. He can be reached at [tlusty@us.ibm.com](mailto:tlusty@us.ibm.com).

## Table of Contents

SAP's Typical Access Pattern to Storage.....	4
Solid-State Drive Technology .....	5
What Are SSDs? .....	5
IBM Storage Solution .....	5
Measurement and Recommendations.....	6
1. SAP for Banking on System z .....	6
Background .....	6
Environmentals: Floor Space, Electrical Power, and Drives Comparisons .....	8
Performance.....	9
2. SAP NetWeaver Business Warehouse Performance on IBM i .....	10
Background .....	10
Performance.....	10
Using SSDs to Improve the Database Performance .....	12
The End-to-End Perspective .....	12
How to Find the Sweet Spot? .....	12
Related Content.....	13
Copyright .....	14

## SAP's Typical Access Pattern to Storage

SAP provides a broad range of solutions for large enterprises based on the SAP Business Suite, including SAP CRM and SAP ERP, or based on industry solutions for different areas such as financial and public services or manufacturing.

A key factor in the success of running many of these business-critical applications is that the underlying infrastructure is able to support the processing of very large data volumes or a high number of transactions with satisfactory performance. Particularly in very large systems with several terabytes of application data, the storage system is one of the main infrastructure components.

Many of these applications have a pattern for how they access the storage system that is quite typical of any OLTP-like workload. This pattern mainly involves random synchronous read I/O operations at very high frequency. The storage system must be able to support this requirement with good response times. This type of access pattern poses a challenge for current storage technology based on hard disk drives and can be solved only by using a higher number of hard disk drives than is actually needed for the amount of data in question.

A typical example of this behavior can be found in the industry solution for SAP for Banking. SAP Deposits Management is a flexible banking solution for all products concerned with balances. It is designed for banks that have a large number of accounts or a high volume of transactions. The application contains extensive functions for checking all accounts, including savings accounts, as well as functions for time deposits and card management. Due to the extremely high number of transactions that have to be processed every day with this application, a test workload has been created to simulate the typical daily volumes of a retail bank. This includes the quantity of payment transaction operations received from external payment transaction systems, single real-time payment transactions received from ATM machines, real-time account statements, as well as the volume of account information read from the system. For more information, refer to *SAP Benchmark: Transaction Banking* [1]. In the following chapters, we will demonstrate the advantages of SSDs for this type of workload.

In addition to OLTP like workloads, analytical applications such as SAP NetWeaver Business Intelligence are challenging for storage subsystems. Huge data volumes of 2-digit TByte have to be accessed and processed with good performance. To be able to compare different technologies on a cross-platform basis, a Standard Application Benchmark, the [SAP NetWeaver Business Intelligence mixed-load benchmark](#), has been created. This workload has also been used to demonstrate the benefits of SSDs.

## Solid-State Drive Technology

Hard disk drive technology has improved considerably over the years, sustaining more I/O operations per second and increased throughput. However, there are inherent physical limitations to the response times (access time and latency) that can be achieved with this technology. These limitations and the continuing demand for better response times and more throughput have created the need for a different kind of storage technology that avoids seeks and rotational delays. This new technology, known as Solid-State Drives, can provide access times measured in microseconds rather than milliseconds.

### What Are SSDs?

Solid-State Drives (SSDs) are data storage devices that use solid-state or flash memory to store persistent data. SSDs are similar to traditional hard disk drives (HDDs) in that data is retained when the drive is powered off. Data can also be accessed randomly and the life of the drive is limited by the number of write/erase operations that can be performed. To extend the lifetime of the drive and to ensure its integrity, SSDs have built-in dynamic and static wear-leveling and bad-block mapping algorithms. This is further complemented by features such as over-provisioning, error detection code, and error correction algorithms to ensure data reliability.

The technology used in SSDs is based on semiconductor material and contains no moving parts for retrieving data. This provides a number of advantages over HDDs. Firstly, it almost entirely eliminates the risk of mechanical failures since it can tolerate extreme shocks, higher altitude, vibration, and extreme temperature. However, SSDs can still fail and must be RAID-protected like HDDs. Secondly, SSDs provide improved I/O transaction-based performance with random I/O up to 100 times faster than HDDs on the device level. Finally, SSDs only consume power when reading or writing data. Therefore, they consume less power, and produce less heat and noise.

At present, SSDs are more expensive than HDDs. Industry experts expect that the gap in price will most likely narrow over time and SSDs might in the future come in line with the price of spinning disks. As the price of SSDs starts to go down, the percentage of SSD backend storage for new storage systems will increase.

### IBM Storage Solution

In the first half of 2009, IBM expanded its support of SSDs across all storage and server lines. This support allows SSDs to coexist with HDDs in any IBM System Storage™ DS8000® Turbo series or internal storage configuration. In practice, a database might be spread over both disk technologies as selected tables and indices can be placed on SSDs while others reside on HDDs. Disk technologies like synchronous mirroring (Metro Mirror), FlashCopy or z/OS DFSMSHsm copy pools are fully compatible with SSDs. Data can be transparently moved from one storage device to another.

Earlier, in October 2008, IBM announced a complementary feature called High Performance FICON for System z (zHPF). zHPF exploits a new channel protocol specially designed for more efficient I/O operations. The maximum number of I/Os on a single 4KB channel with a System z10 server was raised to 31,000 I/Os per second. DB2 primarily benefits from synchronous I/Os, certain log I/Os, and certain deferred write operations.

## Measurement and Recommendations

### 1. SAP for Banking on System z

#### Background

The SAP on System z performance team in Poughkeepsie (USA) has been working with SAP's core banking application and SAP AG for several years, as part of its ongoing testing and measuring activities. The main goal here was to gain experience with the application as customers use it and as it evolves. All the measurements shown in this document used the most current Banking Services for SAP 6.0 and, in particular, its Deposits Management functions. The workload tested was a slightly modified SAP TRBK (Transaction Banking) Standard Application Benchmark [1]. However, the measurements we did were stress tests, not SAP certified benchmarks. Furthermore, no tuning was performed on these measurements - as one might for a benchmark.

There were several measurements taken of the SAP Deposits Management application, comparing storage subsystems with these technologies, using an IBM System Storage DS8300 with both SSDs and HDDs. Our team wanted to see the possible effects in a customer-like application environment. Since the use of different files varies, the team also wanted to experiment with placing different files on the different technologies. A balanced approach was used by not optimizing just one performance metric at the expense of many others.

Two distinct workloads were measured: a modified day processing or posting scenario and a night processing scenario.

In the day processing scenario, the benchmark program simulates typical daily volumes of a retail bank, including the quantity of payment transactions received from external payment transaction systems, single real-time payment transactions received from ATM machines, real-time account statements, as well as the volume of account information read from the system. These transactions can occur anytime, day or night, especially as Internet banking is used more widely. The majority of postings take place during business days. There are high file I/O rates consisting of mostly random synchronous read I/Os. The modification increases the number of payment transactions relative to other transactions in order to increase the I/O relative to the CPU workload.

In the night processing scenario, an account balancing job is performed periodically for each account to calculate charges (fees) and interest and also to post these. The I/O characteristics of this workload are quite different from posting transactions. In general, the workload is more like a batch update workload. As such, it has a sequential access pattern with a heavy transfer rate to the DB2 log.

The following graphic shows the standard storage configuration used as a base line in comparison with the SSD configuration that yields best results.

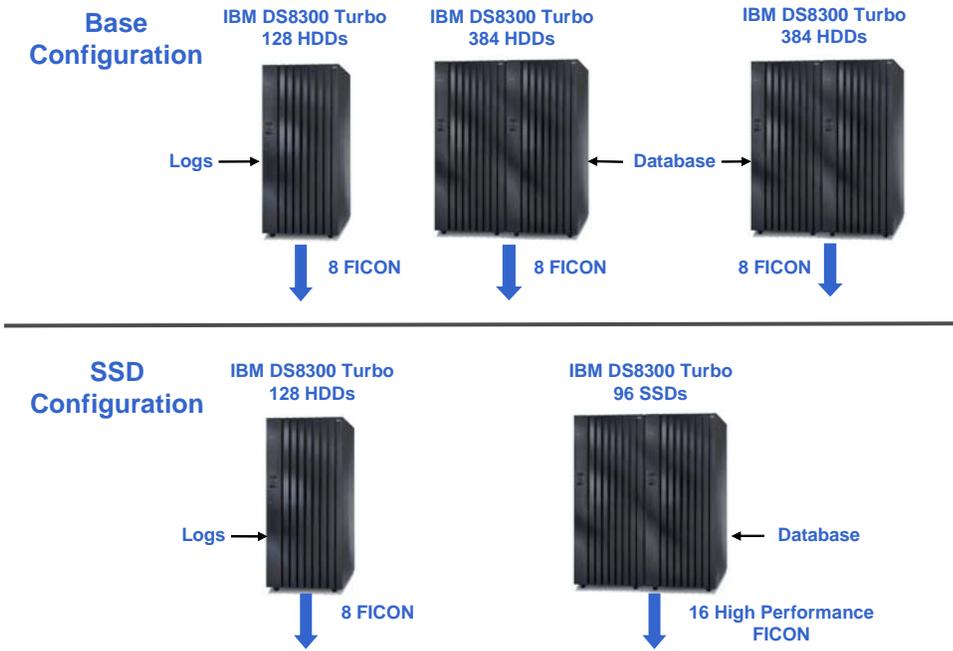
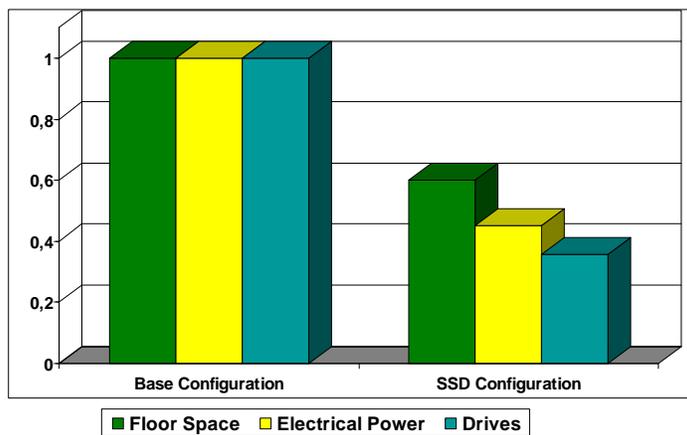


Figure 1: Conceptual View of HDD/SSD Hardware Configuration

## Environmentals: Floor Space, Electrical Power, and Drives Comparisons

There are several ways to compare the overall characteristics of the configurations. In Figure 2 below, we have compared floor space, electrical power consumption, and the number of drives. The base case is three IBM System Storage servers as shown in *Figure 1* above. We can see that the SSD configuration has significantly less floor space, electrical power consumption, and drives than the base case, which demonstrates some of the benefits of including SSD. We also performed some runs where we put both the logs and the database on a single IBM System Storage server using just SSDs. In this case, the reductions for floor space, electrical power, and the number of drives are even greater.

### SAP for Banking – Day Posting with SSD Floor Space, Electrical Power and Disk Drives



**Figure 2: Storage Configuration: Floor Space, Electrical Power Consumption, and Drives**

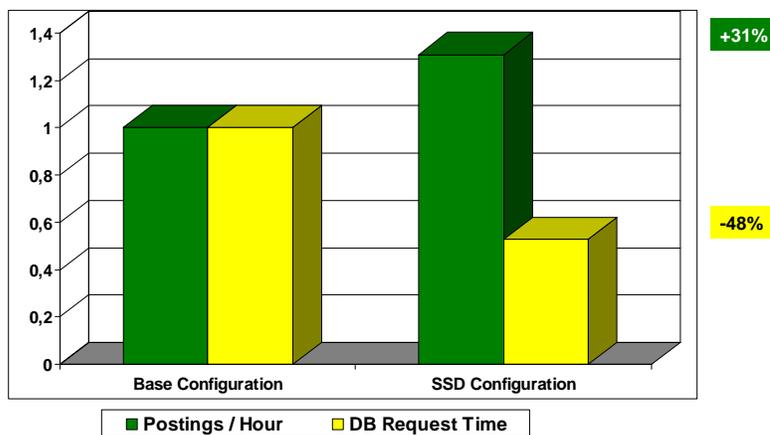
## Performance

The best postings throughput was 31% more than the base. This was achieved by implementing an SSD for the database and the logs on a separate IBM System Storage server with HDDs using zHPF. In all cases, using SSDs improved both the response time and the I/O rate compared to HDD runs.

A comparison of the different runs shows that, as the configuration evolved, throughput increased even though the user think time remained the same. At the same time, each of the two different responsiveness metrics improved substantially.

Using the SSD for just the database and keeping the logs on a separate IBM System Storage server with HDDs further improved these metrics. This is because the logs benefited from having dedicated channels. It is a simplistic example of the benefits that can occur when you segregate the I/O and place the disparate activities on different technologies. The results are plotted in terms of ratio to the base run in *Figure 3*.

### SAP for Banking – Day Posting with SSD Performance



**Figure 3: Posting Workload – Throughput and DB Request Time**

As mentioned earlier, the balancing workload is more like a batch update workload. Its unique I/O resource characteristics are writes and volume of data written to the DB2 logs. As a result, we did not expect much advantage for SSD or zHPF. The measurement results confirmed this.

More details about these measurements can be found in the paper *IBM System z and System Storage DS8000: Accelerating the SAP Deposits Management Workload With Solid State Drives* [2].

## 2. SAP NetWeaver Business Warehouse Performance on IBM i

### Background

The SAP on IBM i performance team in Rochester (USA) has been working with SAP's core business intelligence application for several years as part of our ongoing testing and measuring work. Our aim is to test database and operating system functions and to ensure proper application performance with each release of IBM i. The two sets of measurements performed used the SAP BI Mixed Load (BI-MXL) Standard Application Benchmark and a modified BW workload. The measurements taken in the SAP BI-MXL Standard Application Benchmark were certified and run in a highly-tuned environment. The measurements in the modified BW workload were run as stress tests in an environment with very little tuning.

The SAP BI-MXL Standard Application Benchmark demonstrates the Business Intelligence capabilities of SAP NetWeaver by simulating multiple users querying the database with delta data loads running in parallel. The multiple users query 10 Sales and Distribution InfoCubes that contain an initial total of 300 million records. The InfoCubes are extended with 3 separate delta data loads during the high load activity period of the benchmark run.

The modified BW workload runs only the user phase of the SAP BI-MXL Standard Application Benchmark with a few modifications to stress test the disk subsystem. To achieve a high amount of disk activity, the benchmark application was changed - along with some modifications - to the normal database settings in a benchmark environment. In order to prevent the simulated users from synchronizing on the same dialog steps, the think time assigned between dialog steps was reduced to 0 and, because the number of users was smaller than in a normal benchmark environment, the users were assigned to specific InfoCubes rather than being given a random assignment. For this workload, the database was prevented from reusing previous cached results by turning the CACHE\_RESULTS parameter to \*NONE in the query options file. Also, the database was not tuned since the workload only used the indices that are provided during installation. No aggregates, materialized query tables, new indices, or other performance enhancements were made available during the modified BW workload.

Both of these workloads were run several times on an IBM Power 550 system using different configurations of internal solid-state and traditional drives. The SAP BI-MXL certifications were performed using either all SSDs or all HDDs to demonstrate how the SSDs provide the same performance with fewer drives. The modified BW workload tested three configurations to demonstrate the potential effects of SSDs in a customer-like environment. The configurations tested were all SSDs, all HDDs, and a combination of SSDs and HDDs. The modified BW workload also demonstrates the use of the ASP Balancer to assist with the placement of frequently accessed, or hot, data on the solid-state drives.

### Performance

In April 2009, two SAP BI-MXL Standard Application Benchmarks were certified, thus demonstrating that the same throughput level can be achieved with 77 percent fewer drives in the solid-state configuration than in the traditional hard disk drive configuration. Both runs achieved over 90,000 dialog steps per hour. The HDD configuration [6] required 96 drives while the SSD configuration [7] only needed 22 drives.

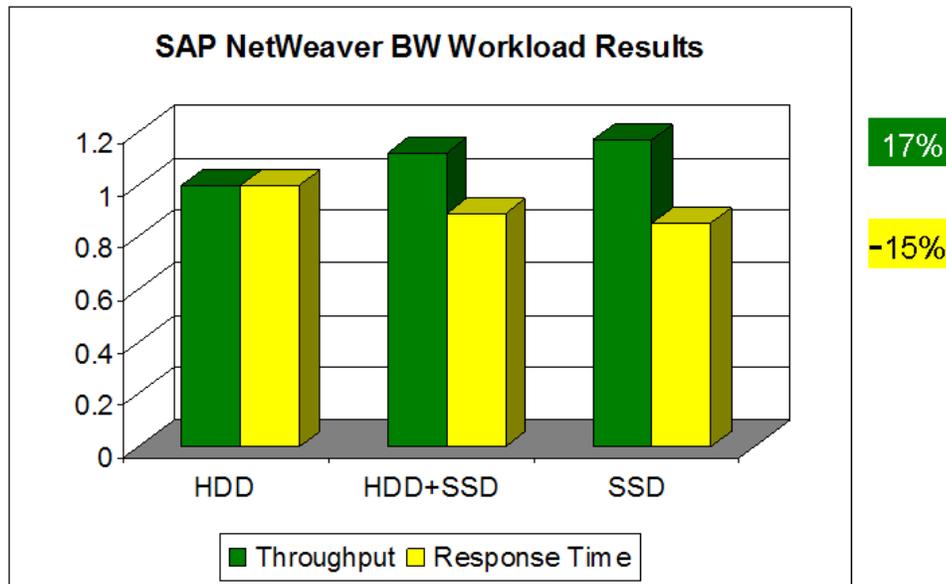
	Throughput (DS / Hour)	Number of Drives
HDD	90,492	96
SDD	90,635	22

77% fewer drives

**Figure 4: SAP BI-MXL Standard Application Benchmark Results**

The modified BW workload demonstrated that the best performance was achieved with a complete SSD configuration using 16 SSDs. There was a 17 percent improvement in throughput over the traditional drive configuration using 96 HDDs. A 12 percent improvement over the HDD configuration was achieved by adding 4

SSD drives to the HDD configuration. Improvements in both response time and I/O rates were seen in all runs involving SSDs.



**Figure 5: BW Workload – Throughput and Response Time**

The mixed environment of 96 HDDs and 4 SDDs used the ASP Balancer tool to identify the hot data and move the identified hot data to the previously empty SDDs. This is only one way that data can be moved to SSDs on IBM i systems. Other methods are described in the paper *Performance Value of Solid State Drives using IBM i* [8].

## Using SSDs to Improve the Database Performance

### The End-to-End Perspective

The best and most obvious benefit of SSDs in both the DS8000 and internal storage options is the improvement for DB2 synchronous read I/Os, as frequently seen with SAP online transactions and analytical queries.

Disk drives are just one aspect of the I/O subsystem. The entire software and hardware stack from the database to the storage device has been optimized to overcome the relatively slow disk I/O rate. Whenever possible, data is buffered and cached to avoid synchronous disk access. Sophisticated DB2 algorithms determine which data might be requested next and then pre-fetch it in such a way that the data is available in the memory whenever the application needs it. Moreover, DS8000 provides a large cache that allows I/O requests to be satisfied without accessing the HDD drives if the requested data is in the cache.

DB2 buffer pools and cache in the storage server continue to be important. Enterprise systems often experience average DB2 synchronous I/O times of two milliseconds or less, and this does not include the zero I/O time for DB2 buffer hits. So, it is simplistic to argue that the busiest data sets belong on expensive fast-performing devices, such as SSD. However, data sets with low access rates do not belong on SSD either. Rather, databases with a large working set and high access rates have the most to gain from SSD performance. But large databases require large disk capacity. Consequently, you cannot always expect to purchase just a small amount of SSD storage and expect to realize a large performance benefit. The benefits will vary based on your particular scenario and the ability to place the appropriate data on the SSDs.

Another benefit of SSD is that it can help to facilitate database administration. As the relevance of the clustering sequence of tables or indexes diminishes, you have the opportunity to perform database management tasks – such as table reorganizations - less frequently in certain cases.

### How to Find the Sweet Spot?

IBM delivers tools and capabilities to optimize the use of SSDs and thereby maximize their value. This can help you to identify tables and indexes that are the best candidates for SSD. For example, new z/OS SAS-based tooling is available for downloading from the IBM Web site [5]. Its purpose is to provide a pragmatic plan for managing the transition to SSD.

When you have found the tables and indexes that benefit most from SSD, you can assign them selectively to SSDs. With DB2 for Linux, UNIX, and Windows, you can create a new tablespace and assign file or device containers that reside on SSD. With DB2 for z/OS, you can take advantage of DFSMS policy-based storage management that allows you to control the allocation of the data sets that contain DB2 tables, indexes, or partitions. With DB2 for IBM i, you can assign objects to SSDs by creating a library in an auxiliary storage pool (ASP) that contains SSDs, assigning a media preference to database tables and indices, or by using the ASP Balancer.

IBM and SAP continues to evaluate SSD usage through customer feedback and in-lab testing with a view to optimizing the technology and identifying further product improvements.

## Related Content

- [1] *SAP Benchmark: Transaction Banking* at <http://www.sap.com/benchmark>
- [2] *IBM System z and System Storage DS8000: Accelerating the SAP Deposits Management Workload With Solid State Drives*, white paper available at:  
<http://www.ibm.com/support/techdocs/atmastr.nsf/WebIndex/WP101442>
- [3] *DS8000: Introducing Solid State Drives*, IBM Redpaper available at  
<http://www.redbooks.ibm.com/redpieces/abstracts/redp4522.html>
- [4] *Ready to Access DB2 for z/OS Data on Solid-State Drives*, IBM Redbook available at  
<http://www.redbooks.ibm.com/abstracts/redp4537.html>
- [5] *FLASHDA – z/OS SAS®-based tooling for Solid State Drives*, available at  
<http://www.ibm.com/systems/z/os/zos/downloads/flashda.html>
- [6] The following benchmark results fully comply with the SAP Benchmark Council regulations and have been audited and certified by SAP AG. SAP Standard Application Benchmarks for SAP BW Releases available prior to Release 3.0 are not comparable with benchmarks for SAP BW Release 3.0 or later. Following benchmark is an SAP BI Mixed Load (BI-MXL) Standard Application Benchmark running on SAP NetWeaver 7.0 on IBM i 6.1 using DB2 for i 6.1 and was performed by IBM in Rochester, MN, USA on March 31, 2009. Certification 2009010. Certified April 29, 2009. Configuration: Central server: IBM Power 550, 2 processors / 4 cores / 8 threads, POWER6, 5.0 GHz, 128 KB L1 cache and 4 MB L2 cache per core, 32 MB L3 cache per processor, 64 GB main memory. Achieved throughput of 90,492 query navigation steps per hour.  
<http://www.sap.com/solutions/benchmark/bimxl.epx>
- [7] The following benchmark results fully comply with the SAP Benchmark Council regulations and have been audited and certified by SAP AG. SAP Standard Application Benchmarks for SAP BW Releases available prior to Release 3.0 are not comparable with benchmarks for SAP BW Release 3.0 or later. Following benchmark is an SAP BI Mixed Load (BI-MXL) Standard Application Benchmark running on SAP NetWeaver 7.0 on IBM i 6.1 using DB2 for i 6.1 and was performed by IBM in Rochester, MN, USA on March 19, 2009. Certification 2009011. Certified April 29, 2009. Configuration: Central server: IBM Power 550, 2 processors / 4 cores / 8 threads, POWER6, 5.0 GHz, 128 KB L1 cache and 4 MB L2 cache per core, 32 MB L3 cache per processor, 64 GB main memory. Achieved throughput of 90,635 query navigation steps per hour.  
<http://www.sap.com/solutions/benchmark/bimxl.epx>
- [8] *Performance Value of Solid State Drives using IBM i*, white paper available at  
<http://www.ibm.com/systems/i/advantages/perfmgmt/resource.html>

## Copyright

© Copyright 2009 SAP AG. All rights reserved.

No part of this publication may be reproduced or transmitted in any form or for any purpose without the express permission of SAP AG. The information contained herein may be changed without prior notice.

Some software products marketed by SAP AG and its distributors contain proprietary software components of other software vendors.

Microsoft, Windows, Excel, Outlook, and PowerPoint are registered trademarks of Microsoft Corporation.

IBM, DB2, DB2 Universal Database, System i, System i5, System p, System p5, System x, System z, System z10, System z9, z10, z9, iSeries, pSeries, xSeries, zSeries, eServer, z/VM, z/OS, i5/OS, S/390, OS/390, OS/400, AS/400, S/390 Parallel Enterprise Server, PowerVM, Power Architecture, POWER6+, POWER6, POWER5+, POWER5, POWER, OpenPower, PowerPC, BatchPipes, BladeCenter, System Storage, GPFS, HACMP, RETAIN, DB2 Connect, RACF, Redbooks, OS/2, Parallel Sysplex, MVS/ESA, AIX, Intelligent Miner, WebSphere, Netfinity, Tivoli and Informix are trademarks or registered trademarks of IBM Corporation.

Linux is the registered trademark of Linus Torvalds in the U.S. and other countries.

Adobe, the Adobe logo, Acrobat, PostScript, and Reader are either trademarks or registered trademarks of Adobe Systems Incorporated in the United States and/or other countries.

Oracle is a registered trademark of Oracle Corporation.

UNIX, X/Open, OSF/1, and Motif are registered trademarks of the Open Group.

Citrix, ICA, Program Neighborhood, MetaFrame, WinFrame, VideoFrame, and MultiWin are trademarks or registered trademarks of Citrix Systems, Inc.

HTML, XML, XHTML and W3C are trademarks or registered trademarks of W3C®, World Wide Web Consortium, Massachusetts Institute of Technology.

Java is a registered trademark of Sun Microsystems, Inc.

JavaScript is a registered trademark of Sun Microsystems, Inc., used under license for technology invented and implemented by Netscape.

SAP, R/3, SAP NetWeaver, Duet, PartnerEdge, ByDesign, SAP Business ByDesign, and other SAP products and services mentioned herein as well as their respective logos are trademarks or registered trademarks of SAP AG in Germany and other countries.

Business Objects and the Business Objects logo, BusinessObjects, Crystal Reports, Crystal Decisions, Web Intelligence, Xcelsius, and other Business Objects products and services mentioned herein as well as their respective logos are trademarks or registered trademarks of Business Objects S.A. in the United States and in other countries. Business Objects is an SAP company.

All other product and service names mentioned are the trademarks of their respective companies. Data contained in this document serves informational purposes only. National product specifications may vary.

These materials are subject to change without notice. These materials are provided by SAP AG and its affiliated companies ("SAP Group") for informational purposes only, without representation or warranty of any kind, and SAP Group shall not be liable for errors or omissions with respect to the materials. The only warranties for SAP Group products and services are those that are set forth in the express warranty statements accompanying such products and services, if any. Nothing herein should be construed as constituting an additional warranty.