IBM Technical Brief

IBM zEnterprise System®: DB2 11 for z/OS with SAP® Performance Report

Version 1.0

December 16, 2013

Authors:

Brenda Beane
Seewah Chan
Akiko Hoshikawa
Dr. Paul Lekkas
Michael Sheets

Document Owner:

Brenda Beane
IBM SAP on System z Performance
Version: 1.00
Date: December 16, 2013
Filename: DB211andSAP_performance.pdf

Contents

Tables

Figures

Disclaimers

Neither this document nor any part of it may be copied or reproduced in any form or by any means or translated into another language, without the prior consent of the IBM Corporation. IBM makes no warranties or representations with respect to the content hereof and specifically disclaims any implied warranties of merchantability or fitness of any particular purpose. IBM assumes no responsibility for any errors that may appear in this document. The information contained in this document is subject to change without any notice. IBM reserves the right to make any such changes without obligation to notify any person of such revision or changes. IBM makes no commitment to keep the information contained herein up to date.

Performance data and customer experiences for IBM and non-IBM products and services contained in this document were derived under specific operating and environmental conditions. The actual results obtained by any party implementing any such products or services will depend on a large number of factors specific to such party's operating environment and may vary significantly. IBM makes no representation that these results can be expected or obtained in any implementation of any such products or services.

The information in this document is provided "as-is" without any warranty, either express or implied. IBM expressly disclaims any warranties of merchantability, fitness for a particular purpose or infringement.

Trademarks

IBM, the IBM logo, zSeries and z/OS are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries or both. AnyNet, CICS, DB2, DFSMSdfp, DFSMShsm, DFSMSrmm, DFSORT, DS6000, DS8000, ESCON, FICON, GDPS, HiperSockets, HyperSwap, IBM, IBM eServer, IBM logo, IMS, InfoPrint, Language Environment, Multiprise, NetView, OMEGAMON, Open Class, OS/390, Parallel Sysplex, PrintWay, RACF, RMF, S/390, Sysplex Timer, System Storage, System z, System z10, System z10, SystemPac, Tivoli, VSE/ESA, VTAM, WebSphere, z/Architecture, z/OS, z/VM, z/VSE, and zSeries are trademarks or registered trademarks of the International Business Machines Corporation in the United States and other countries.

SAP, the SAP logo, and all other SAP product and service names mentioned herein are trademarks or registered trademarks of SAP AG in Germany and in several other countries all over the world.

Unicode is a trademark of Unicode, Inc.

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

Acknowledgements

The authors of this document would like to acknowledge the following individuals for their contributions to this project:

- Judi Bank (System z Performance)
- Andrea Fuga (z/OS System Programming)
- Tom Litty (z/OS System Programming)
- Veng Ly (IBM SAP on System z Performance)
- Johannes Schuetzner (IBM SAP on DB2 for z/OS Development)
- Steve Turnbaugh (DB2 for z/OS Development)

We worked very closely with the DB2 for z/OS team. We applaud them for the performance and scalability improvements that they delivered in DB2 11.

Feedback

Please send comments or suggestions for changes to beane@us.ibm.com

1.0 Introduction

The performance of DB2 for z/OS is critical to the SAP solution on System z. As a result, IBM invests heavily in it. With DB2 11 for z/OS, large investments were made in both development and test. Over forty new features were put in DB2 11 specifically for SAP. Changes were made in DB2 11 to improve performance and to enhance continuous availability, scalability, and security for business critical information. These investments have paid off with DB2 11 for z/OS in the SAP environment.

DB2 11 for z/OS delivers "out-of-the-box" DB2 CPU savings of up to 10% for complex online transaction processing (OLTP) and up to 10% for update-intensive batch workloads, when compared to DB2 10 running the same workloads. Even greater savings can be seen with the SAP Banking Services (SBS) Day Posting workload. Complex queries may enjoy up to 40% CPU savings. Utility performance improvements have also been delivered.

DB2 11 for z/OS delivers scalability improvements. Support for larger relative byte address (RBA) and logical record sequence number (LRSN) is available. A larger dynamic statement cache is supported. Reductions in DB2 latch class contention in buffer management have been delivered.

DB2 11 for z/OS delivers efficient partition scaling. Partitioned table spaces can have a larger number of partitions without adversely affecting performance.

DB2 11 for z/OS delivers more efficient data sharing performance and scalability. Enhancements have been made to group buffer pool (GBP) processing to take advantage of the GBP write around function. Castout processing was also enhanced to avoid GBP full situations and performance impacts. More efficient index split processing and full LRSN spin avoidance have also been implemented.

DB2 11 for z/OS delivers greater efficiency and CPU savings by making more work eligible for zIIP processing, especially in DB2 system address space operations. IBM System z has offered Integrated Information Processors (zIIPs) since the z9 generation of machines. These processors are designed to help free-up general computing capacity and lower the overall cost of computing. The overall performance of a system using a mix of general purpose CPs and zIIPs is equivalent or better than a system using all general purpose CPs. See A.0 Appendix: DB2 11 for z/OS zIIP Eligible Processing on page 39 for details of the zIIP updates in DB2 11.

DB2 11 for z/OS delivers many more benefits and value. See reference [1] on page 38 for more information.

SAP AG completed certification of DB2 11 for z/OS at the same time that IBM made DB2 11 for z/OS generally available. This is the earliest SAP certification of any DB2 for z/OS release. See SAP Note 1850403 (Release of DB2 11 for SAP Components).

The IBM SAP on System z Performance Team, located in Poughkeepsie, NY, worked very closely with the DB2 for z/OS Development and Performance teams for over two and a half years throughout the development of DB2 11 for z/OS. We evaluated the performance of early DB2 11 drivers, prototypes, and tuning recommendations. We collected and analyzed performance data. We identified and helped debug performance issues, as well as functional problems. We drove performance improvements into DB2 11 for z/OS. And, finally, we executed a set of measurements comparing DB2 11 with DB2 10.

We focused on running DB2 11 for z/OS with SAP on an IBM Enterprise Class machine, a zEC12 . Specifically, we ran the SAP Banking Services (SBS) Day Posting workload, which is a good representation of a customer OLTP workload. See section 3.0 Workload Description on page 10 for details of this workload.

We used various configurations for these tests to show how DB2 11 performed in different environments and how it compared to DB2 10. We focused on areas that are important to our customers running SAP on System z. We ran in both DB2 single system and data sharing environments. We ran with different database sizes and different number of partitions in the database. We ran with all general purpose processors (CPs) and a mix of CPs and zIIPs.

This paper is a culmination of all this work. It is a joint effort between the IBM SAP on System z Performance Team and the IBM DB2 for z/OS Performance Team. In it, we will discuss the results of key measurements that were done and our experiences with DB2 11. The measurement tests that were done were stress tests, not certified benchmarks.

2.0 Summary of Results

Our measurement results show that DB2 11 for z/OS delivers significant performance and scalability improvements compared to DB2 10.

For the tests documented in this paper, we measured the performance of DB2 11 in terms of Internal Throughput Rate (ITR). See reference [6] on page 38 for an explanation of this term.

In a DB2 single system environment, we saw performance improvements in the range of 7% to 18% with DB2 11 depending on the configuration when running the SAP Banking Services (SBS) Day Posting workload.

Using a SBS 6.0 database with 5 million accounts and 128 partitions, we compared DB2 11 with DB2 10 on a zEC12 with a single z/OS logical partition (LPAR) with 2 CPs. There was a 7% improvement in ITR with DB2 11.

Using a SBS 7.0 database with 60 million accounts and 512 partitions, we compared DB2 11 with DB2 10 on three different configurations of a single z/OS LPAR on a zEC12. The results were impressive. With 2 CPs, there was a 15% improvement in ITR with DB2 11. With 4 CPs, there was a 16% improvement. And, with 8 CPs, there was an 18% improvement in ITR with DB2 11. The results show that DB2 11 delivers efficient CP scaling. As the DB2 system got larger, the performance improvement of DB2 11 grew.

The measurement with the SBS 6.0 database can not be compared to the measurements with the SBS 7.0 database. The differences in the SAP levels (SBS 6.0 vs SBS 7.0), the database size (5M accounts vs 60M accounts), and the number of database partitions (128 vs 512) make it like different workloads were used, even though both sets of measurements used the SAP Day Posting workload. However, we do know that the number of partitions in the databases used in these measurements affected the degree of the performance improvement seen between DB2 11 and DB2 10.

We saw significant performance improvements with partition scaling in DB2 11. We tested the efficiency of partition scaling with both DB2 10 and DB2 11. We ran the SAP Day Posting workload on a zEC12 with a single z/OS LPAR with 3 CPs using a database with 512 partitions and one with 1024 partitions. There was a 3.7% degradation in ITR going from 512 to 1024 partitions with DB2 10. There was less than a 1% degradation in ITR going from 512 partitions to 1024 partitions with DB2 11. This is a significant improvement in scalability with DB2 11 for transactions accessing one or a few partitions out of a large number of partitions.

Therefore, the performance delta between DB2 10 and DB2 11 can be larger when there are a larger number of partitions. Table spaces are typically partitioned for large tables to improve performance and ease maintenance. As the data in databases grow, it may be beneficial to use a larger number of partitions.

In a DB2 2-way data sharing environment, **we saw up to an 11% improvement in performance with DB2 11** when running the SAP Banking Services Day Posting workload.

Using a SBS 7.0 database with 60M accounts and 512 partitions, we ran 2-way data sharing on a zEC12 with two z/OS LPARs, each with 6 general purposes CPs. There was a 9% improvement in ITR with DB2 11 compared to DB2 10.

We observed a larger performance improvement with DB2 11 in 2-way data sharing when we used zIIPs. Running on a zEC12 with two z/OS LPARs, each with 3 general purpose CPs and 3 zIIPs, there was an 11% improvement in ITR with DB2 11.

The performance improvements that we saw with DB2 11 can be attributed to several things.

- A big contributor was the efficient partition scaling.

- Another significant contributor was the enhancements to column processing. DB2 uses customized instructions for repeated operations (xProcs). The new column procedures introduced in DB2 11 are created at prepare and used to pass the output columns to the application. Since a typical SAP application has large numbers of columns in the select list, this optimization is effective especially in SAP applications.

- Also contributing to the performance improvements was the moving of the log buffers to 64-bit common storage, eliminating the need for cross address space operations between the DBM1 and MSTR address spaces for logging.

- The exploitation of the synchronous TCP/IP receive model by Distributed Data Facility (DDF) also provided CPU savings.

- Enhancements in group buffer pool and castout processing, more efficient index split processing, and full LRSN spin avoidance contributed to the performance improvements in data sharing.

DB2 11 continues its deep integration with System z and this provides value to our customers. DB2 11 makes more work eligible for zIIP processing. Using zIIPs is not only cost efficient, but also provides equivalent or better performance. In most of our measurements, we saw improved performance when using a mix of general purpose CPs and zIIPs versus just using general purpose CPs. We also saw more processing being redirected to zIIPs with DB2 11, including work from the DB2 MSTR address space.

We found that the following scalability improvements in DB2 11 had or could have a positive effect on performance.

- The extended Relative Byte Address (RBA) and Log Record Sequence Number (LRSN) addressing support to 1 yottabyte (10^{24}) is important since some customers are reaching the limit of the 6-byte RBA format and manual recovery action is needed to reset it to zero. In addition to this, using the larger 10-byte RBA / LRSN provides full LRSN spin avoidance in data sharing which may improve performance. This function is available in DB2 11 New Function Mode (NFM). We implemented this in our 60M account database.

  We did not see the performance benefit of the full LRSN spin avoidance delivered in DB2 11 in our data sharing measurement since the SAP Day Posting workload has a lot of random access which is typical for OLTP. However, benefits of this should be seen in data sharing with batch workloads where data is updated or deleted through sequential access.

- Reductions in latch class contention benefits performance. We saw this throughout our measurements, especially for latch class 14 which is related to the buffer manager.

- A larger local dynamic statement cache can greatly improve performance. The size of the local statement cache is controlled by the MAXKEEPD DB2 system parameter. The limit was 64K in DB2

10. In DB2 11, the limit was increased to 200K. We did not need to use a MAXKEEPD value greater than 64K in our measurements since we ran a single, homogeneous workload and our local statement cache hit ratio was already 100%. Exploiting larger MAXKEEPD values for larger configurations and/or mixed workloads is highly recommended to improve performance and to reduce response time if you are seeing less than a 100% local statement cache hit ratio. See reference [7] on page 38 for more details.

- A group buffer pool castout threshold (CLASST) of less than 1% of the buffer pool capacity can improve performance. The CLASST parameter on the ALTER GROUPBUFFERPOOL can be set to an absolute number of pages, as well as to a percentage. This can be useful for very large group buffer pools to more evenly spread out asynchronous writes. We did not need to make use of this feature in our measurements, but it can benefit large scale systems. It would have been beneficial in the SBS 7.0 150 Million Account Measurements that our team did with DB2 10 for z/OS. See reference [4] on page 38.

Our tests with DB2 11 for z/OS and SAP show that DB2 11 has delivered significant performance and scalability improvements that will benefit our customers.

3.0 Workload Description

The SAP Banking Services Day Posting workload was used in these tests. It is an online transaction processing (OLTP) workload. We have been running this workload for many years. We have quite a bit of experience with it. See references [4, 5] on page 38.

In this workload, a posting is a deposit or a withdrawal from a customer's account. Typical examples of a posting are a payment out of the account or a deposit into the account. This workload was developed by SAP to simulate customer environments. The workload consists of interactive "users" going through repetitive cycles of 15 dialogue steps.

Step	Operation
1	Create a total of 150 postings via five BAPI calls
2	Create five postings
3	Create bank statement
4	Read postings for account
5	Read details of postings
6	Create five postings
7	Create one bank statement for account
8	Create five postings
9	Create one bank statement for account
10	Create five payment orders
11	Read balances of account
12	Create five postings
13	Create one bank statement for account
14	Read balances for account
15	Read master data for account

Table 1: SAP Banking Day Posting Workload

4.0 Test Environment

Our tests were done with various configurations. Since the main focus is on the database server, all test scenarios were done using a physical 3-tier environment where each of the three layers resided on separate machines. The SAP Database Server was on an IBM zEnterprise system (zEC12) running z/OS. The SAP application servers were IBM p7 Blade Servers running AIX. The presentation server was an IBM POWER5 p55A running AIX.

The following figure provides an overview of our test environment.

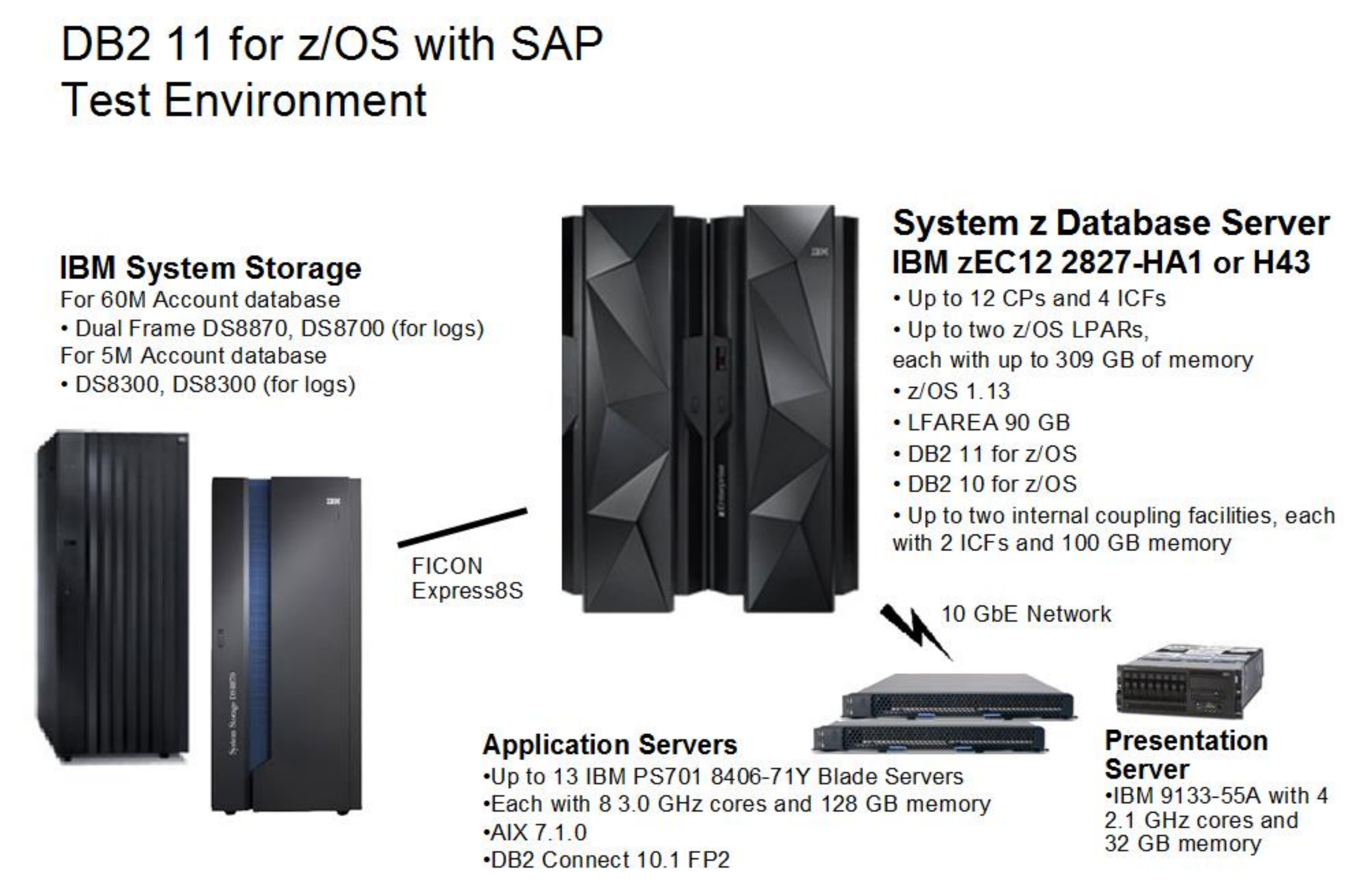

Figure 1: Test Environment - DB2 11 for z/OS with SAP

4.1 Hardware

System z Database Server

An IBM zEnterprise Class 2827-H43 (zEC12) with up to 8 dedicated CPs configured online was used for our DB2 single system measurements. This system had up to 309 GB of real storage configured on one z/OS LPAR.

An IBM zEnterprise Class 2827-HA1 (zEC12) with 12 dedicated CPs and 4 internal coupling facility (ICF) processors configured online was used for our DB2 2-way data sharing measurements. This system had 670 GB of real storage configured across two z/OS LPARs and two Coupling Facilities. For more information on this configuration, see page 16.

The specific hardware configuration used for the SAP Database Server in each test is documented in the details describing the individual tests.

1M fixed large frame support was used for all the measurements. This is a feature available on System z hardware starting with the z10 machine. It requires z/OS 1.10 or higher. It allows some storage to be set aside to be used for 1 MB pages. Starting with DB2 10, DB2 uses large frames, if available, for buffer pools that are page fixed. Storage is allocated for large frames by setting the LFAREA parameter in IEASYSxx. For these tests, 90 GB of real storage was set aside in each z/OS LPAR for 1M fixed large frames. Using 1M fixed large frame support for DB2 buffer pools is recommended since it provides a performance benefit.

zEC12 and z/OS 2.1 introduced support for 2G large frames. This support was also ported back to z/OS 1.13 and made available as a web deliverable. DB2 11 supports both 1M and 2G large frames for buffer pools that are page fixed. We did not use 2G large frame support.

Database DASD

Two different SAP Banking databases were used in these tests. There was a SBS 6.0 database with 5 million accounts and there was a SBS 7.0 database with 60 million accounts. All the tests except one used the SBS 7.0 database with 60 million accounts. Both of these databases contain unicode data and all the database tables were compressed. The indexes were not compressed.

The SBS 6.0 database with 5M accounts resided on a single frame IBM TotalStorage Turbo Model DS8300 (2107-932) server with 15K RPM disks configured as RAID 5 (Redundant Array of Independent Disks). The effective total capacity was five hundred twenty 3390-mod54 volumes – about 27 TB. The unit has 128 GB regular cache and 4 GB non-volatile storage (NVS). It was attached to the zEC12 via eight long wave FICON Express8S attachments, but they had an effective rate of only 4 Gbps due to older/slower host adapters on the DASD and switch. The DB2 subsystem, the database, and one flash copy were contained on 158 of these volumes.

The SBS 7.0 database with 60M accounts resided on a dual frame IBM System Storage DS8870 (2107-961) server with 15K RPM disks configured as RAID 5. The effective total capacity was eight hundred ninety six 3390-mod54 volumes – about 46 TB. The unit has 256 GB regular cache, 8 GB non-volatile storage (NVS), and sixteen long wave FICON Express8S attachments. The DB2 subsystem, the database, and two flash copies were contained on 834 of these volumes.

The DB2 active logs associated with each of these databases resided on separate DASD units from the application databases. The active logs for the 5M account database were striped across four 3390-mod54 volumes on separate ranks of a single frame IBM System Storage DS8300 (2107-922) server.

The active logs for the 60M account database were striped across four 3390-mod54 volumes on separate ranks of a single frame IBM System Storage DS8700 (2107-941) server.

HyperPAV support provided on the DS8000 family of DASD was used to reduce disk I/O queueing.

High Performance FICON for System z (zHPF) with multi-track support was used to improve the efficiency of I/O resources.

Application Servers

Up to 13 IBM PS701 8406-71Y Blade Servers running AIX were used for the SAP application servers. Each of these p7 blades had eight 3.0 GHz processor cores and 128 GB of memory. The SAP central instance resided on one of these blades, along with a stand-alone SAP enqueue server. Four SAP dialog instances were on each of the remaining blades. The actual number of SAP application servers used varied depending on the specific test.

Presentation Server

One IBM 9133-55A server with four 2.1 GHz processor cores and 32 GB of memory running AIX was used as the presentation server to drive the workload.

Network

A dedicated 10 Gb Ethernet network was used to connect the presentation server, the applications servers, and the database server. The application servers were connected via 10 Gb Ethernet adapters through a 10 Gb Ethernet switch to the zEC12 via two OSA-Express4S adapters. The Optimized Latency Mode (OLM) option of the OSA-Express4S adapters was used to improve the elapsed time of this communication.

4.2 Software

z/OS

z/OS release 01.13

Note: This is the minimum z/OS level supported by DB2 11.

DB2 for z/OS

DB2 10 (June 2012)
DB2 11

DB2 Connect

IBM Data Server Driver for CLI that is shipped as part of DB2 Connect 10.1 FP2

AIX

AIX 7.1.0

SAP

For the system used with the 5M account database:

- SAP NetWeaver 7.1
- SAP kernel level 720 EXT, 64-bit, patch number 114

For the system used with the 60M account database:

- SAP NetWeaver 7.1 Enhancement Package 1
- SAP kernel level 720 EXT, 64-bit, patch number 400

4.3 DB2 Configuration

Buffer pools

- About 8 GB of storage was allocated for the buffer pools used with the 5M account database.
- About 60 GB of storage was allocated for the buffer pools used with the 60M account database.
- Most of the buffer pools were paged fixed (PGFIX=YES).
- 1M fixed large frame support was used for the buffer pools that were paged fixed.

Buffer pool allocation was changed in DB2 11 to allocate storage for the entire buffer pool upon first reference. This differs from DB2 10 where storage for the buffer pool was allocated as needed up to the maximum size specified.

DB2 11 provides a parameter FRAMESIZE on the ALTER BUFFERPOOL command to specify the frame size to be used. Possible values are 4K, 1M, and 2G. If you are migrating from DB2 10 to DB2 11 and you used 1M pages in DB2 10 then 1M pages will be the default in DB2 11.

DB2 Logging

DB2 striped logs were used.

MAXKEEPD

MAXKEEPD=64K was used. This is the limit in DB2 10. DB2 11 supports a MAXKEEPD value of up to 200K. We did not need to use a larger MAXKEEPD value in our measurements since our local statement cache hit ratio was 100%. Exploiting larger MAXKEEPD values is highly recommended to increase the local statement cache hit rate. See reference [7] on page 38 for more details.

Large RBA (or 10-byte RBA format) at DB2 11 NFM

The table spaces / indexes in the 60M account database used by the SAP Day Posting workload were converted to the 10-byte RBA format. This is done using the REORG utility. It is recommended to set the DB2 system parameters UTILITY_OBJECT_CONVERSION and OBJECT_CREATE_FORMAT to EXTENDED so objects will be converted to the 10-byte format during a utility (such as REORG or LOAD) or when an object is created. To take full advantage of extended log RBA, both the BSDS and user objects need to be converted.

Rebind of SAP Application at DB2 11 NFM

The SAP application was rebound in DB2 11 NFM with the DB2 system parameter APPLCOMPAT set to V11R1.

Multi-row INSERT

The IBM Data Server Driver for CLI can automatically convert chains of SQL INSERT statements into multi-row INSERT (MRI) statements that insert up to 100 records at once into a DB2 table. This capability was used for both DB2 10 and DB2 11.

4.4 2-way Data Sharing Environment

The following figure describes the 2-way data sharing environment used in our measurements. It shows the LPARs that were configured on the IBM zEnterprise Class 2827-HA1, which was used as the SAP Database Server.

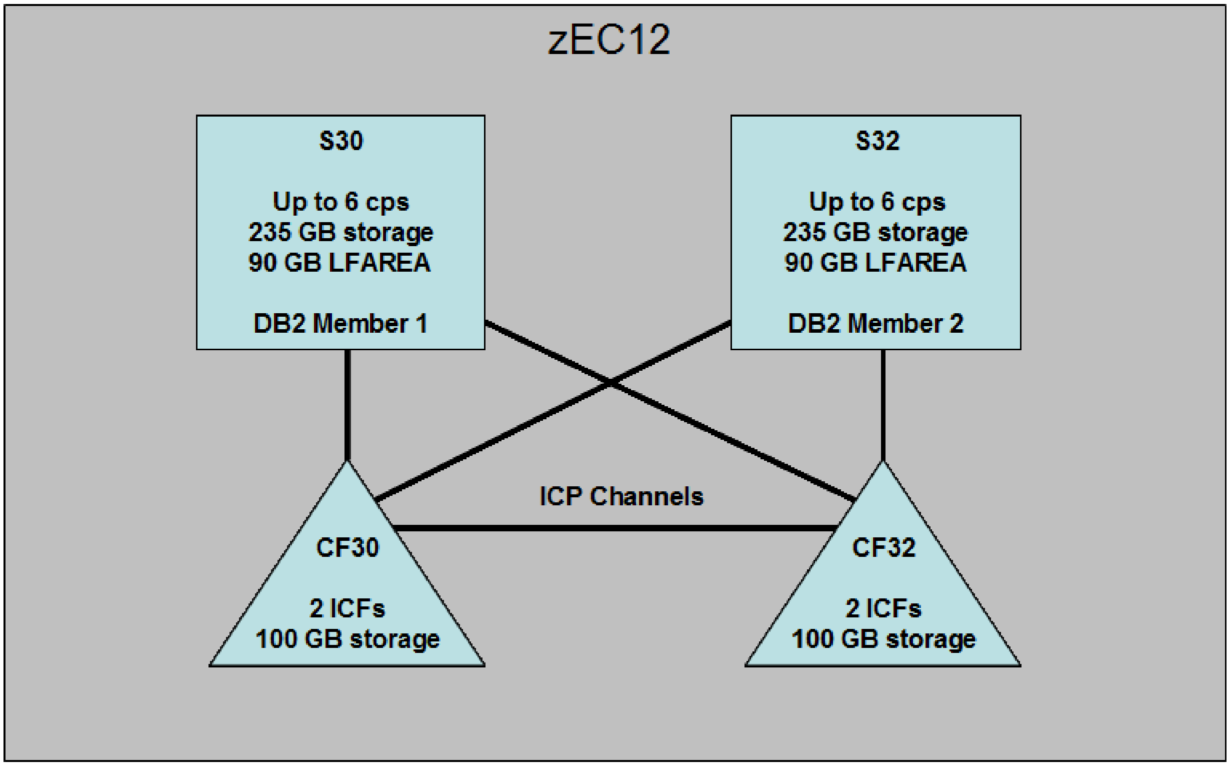

Figure 2: zEC12 LPAR Configuration for DB2 2-way Data Sharing

The SAP Database Server had two z/OS LPARs, named S30 and S32. Each of these LPARs were configured with up to 6 dedicated CPs and 235 GB of real storage. On each of these images, 90 GB of the storage was set aside for 1M fixed large frames by using the z/OS LFAREA parameter. Member 1 of the DB2 data sharing group was on S30 and member 2 of the data sharing group was on S32.

There also were two LPARs for Coupling Facilities, named CF30 and CF32. Each of these LPARs were configured with 2 dedicated ICF processors and 100 GB of real storage.

The Coupling Facility LPARs were connected to the z/OS LPARs using Internal Coupling Peer (ICP) channels. There were three ICP channels each between S30 and CF30 and S30 and CF32. Likewise, for S32 and CF32 and S32 and CF30. There were two ICP channels between CF30 and CF32.

Group buffer pool duplexing was enabled. CF32 contained the primary group buffer pool structures. CF30 contained the secondary group buffer pool structures and the Shared Communication Area (SCA) and lock structures. CF30 had 53 GB of storage allocated. CF32 had 46 GB of storage allocated.

The Global Resource Serialization (GRS) Star configuration was used in the parallel sysplex environment.

4.5 Database Configuration

The SBS 6.0 database with 5M accounts had the following attributes / customizations:

- Unicode
- Compression for tables, not indexes
- Reordered Row Format (RRF)
- Partition-by-Growth (PBG) universal table spaces and classic partitioned table spaces according to SAP Note 496904 (Performance notes database parameters FS-AM for DB2 for z/OS)
- MEMBER CLUSTER option was used for heavily inserted tables according to SAP Note 496904
- VARGRAPHIC data type
- Key banking tables customized with 128 partitions

The SBS 7.0 database with 60M accounts had the following attributes / customizations :

- Unicode
- Compression for tables, not indexes
- Basic Row Format (BRF)
- Partition-by-Growth (PBG) and Partition by Range (PBR) universal table spaces according to SAP Note 496904 (Performance notes database parameters FS-AM for DB2 for z/OS)
- MEMBER CLUSTER option was used for heavily inserted tables according to SAP Note 496904
- VARGRAPHIC data type
- Key banking tables customized with 512 partitions

5.0 Measurements and Results

We executed a number of tests with DB2 11 using the SAP Day Posting workload. We used various configurations to show how DB2 11 performed in different environments. We focused on areas that are important to our customers running SAP on System z. We executed comparable tests with DB2 10.

We completed the following measurements.

- DB2 Partition Scaling
- DB2 Single System – 5M Accounts (SBS 6.0)
- DB2 11 Migration and 10-byte RBA Conversion
- DB2 Single System – 60M Accounts (SBS 7.0)
- DB2 2-way Data Sharing – 60M Accounts (SBS 7.0)
- zIIP Exploitation

The results of these measurements are documented in this section. We use LSPR terminology which is explained in reference [6] on page 38 to describe our results.

We tested with DB2 single system and 2-way data sharing environments since these provide the most insights into the performance of DB2 11. We are not suggesting that these are the ideal configurations for our customers running SAP. For guidelines on how best to implement and run SAP for Banking solutions on System z, see reference [9] on page 38.

5.1 DB2 Partition Scaling

Our first set of measurements was defined to demonstrate the improvements in partition scaling that were delivered in DB2 11. This set of measurements included running both DB2 10 and DB2 11 with a database with 512 partitions and the same database with 1024 partitions.

We used the SBS 7.0 60M account database for these tests. This database initially had about 13 key banking tables used in the SAP Day Posting workload defined as Partition By Range (PBR) table spaces with 1024 partitions. We had two sets of flash copies (backups) defined for this database so we could maintain two copies of the database, one with 1024 partitions and one with 512 partitions. We created the copy with 512 partitions by unloading the data from these 13 tables with 1024 partitions, dropping/recreating these table spaces to define 512 partitions, concatenating the unloaded data from every two partitions into single flat files, reloading the data back into 512 partitions, and running RUNSTATS. Both of these flash copies were maintained at DB2 10 New Function Mode (NFM). To do a measurement at DB2 11, we first migrated from DB2 10 to DB2 11 Conversion Mode (CM).

With DB2 10, we measured a 3.7% degradation in ITR going from 512 partitions to 1024 partitions. This degradation was because DB2 has been assuming that the application would eventually access all the partitions defined. This assumption is not always correct, especially in online transaction processing. DB2 11 is now optimized to achieve linear scalability by accessing a few partitions on a table space with a large number of partitions defined. With DB2 11, we measured less than a 1% degradation in ITR going from 512 partitions to 1024 partitions.

The details of these measurements are summarized in the following table.

DB2 level	**DB2 10**		**DB2 11**	
Run id	S30220B1	S30219B1	S30221B1	S30225B1
# of Database Partitions	512	1024	512	1024
Database Server	zEC12	zEC12	zEC12	zEC12
# of z/OS LPARs	1	1	1	1
# of CPs per z/OS LPAR	3	3	3	3
Real Storage	309 GB	309 GB	309 GB	309 GB
LFAREA (for 1M frames)	90 GB	90 GB	90 GB	90 GB
# of Users	2500	2500	2500	2500
# of DB2 Threads	329	329	329	329
%CPU on z/OS	69.33%	73.33%	62.56%	64.41%
ETR (DS/sec)	202.13	204.99	200.99	205.46
ITR (DS/sec)	291.55	279.54	321.28	318.99
Response Time (secs)	2.368	2.196	2.438	2.168

Table 2: DB2 Partition Scaling Measurement Results

The following figure shows the effect of the number of database partitions on the ITR for both DB2 10 and DB2 11. It also includes trend lines for the ITR as the number of partitions decrease toward 128. The DB2 partition scaling with DB2 11 is impressive.

Table spaces are typically partitioned for large tables to improve performance and ease maintenance. As the data in databases grow, it may be beneficial to use a larger number of partitions. With DB2 11, this can be done without being concerned about performance.

This figure also shows how the number of database partitions can affect the performance delta between DB2 10 and DB2 11. Greater performance improvements may be seen with DB2 11 with databases that have larger number of partitions.

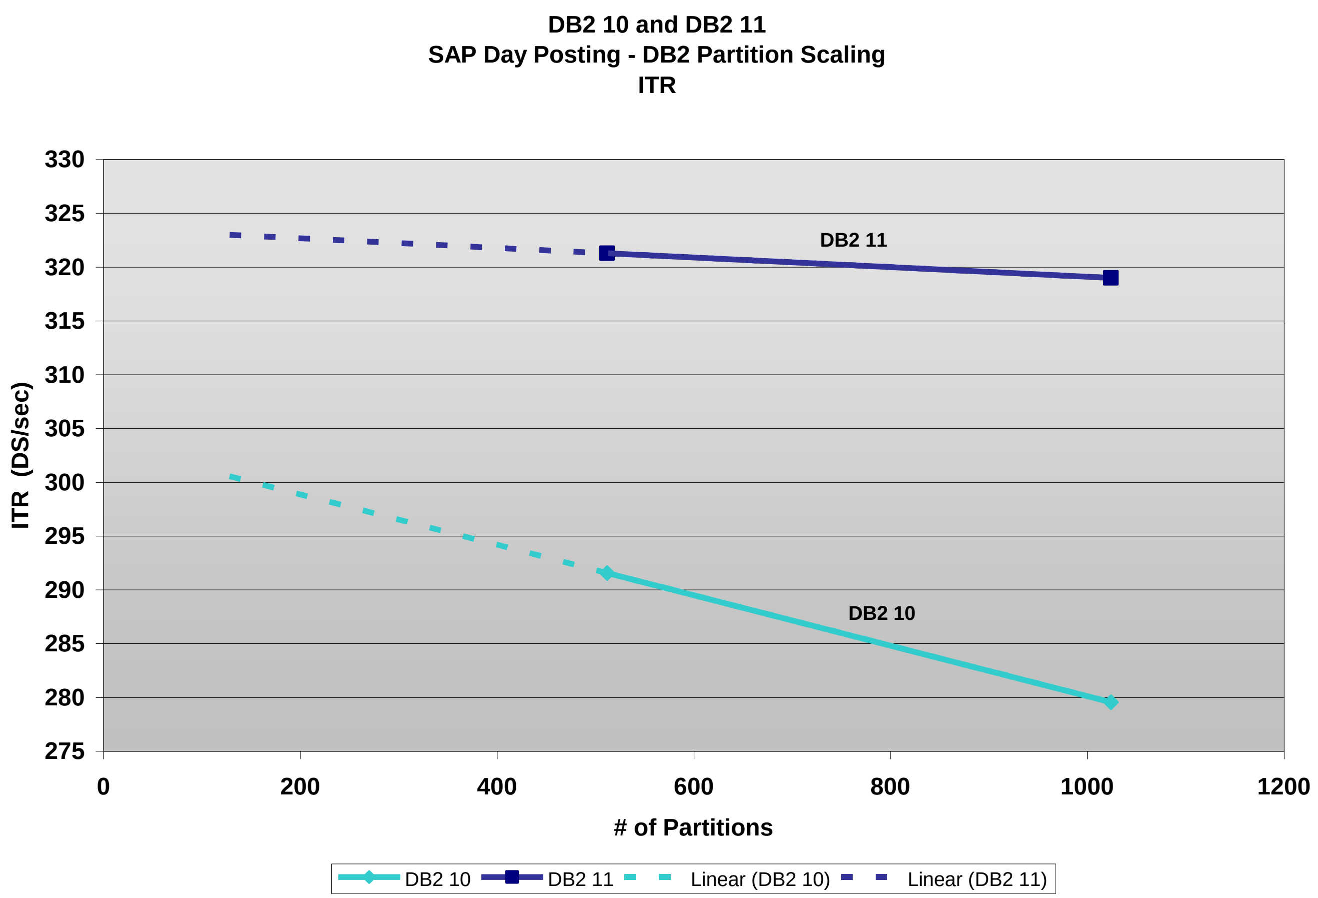

Figure 3: DB2 10 and DB2 11 - Partition Scaling - ITR

5.2 DB2 Single System – 5M Accounts (SBS 6.0)

In these measurements, we ran the SAP Day Posting workload with the SBS 6.0 database with 5M accounts. We ran with DB2 10 and DB2 11 CM on a zEC12 with a single z/OS LPAR with 2 CPs. There was a 7% improvement in ITR with DB2 11. The SBS 6.0 5M account database had the key banking tables defined with 128 partitions. Systems with databases with a similar number of partitions to this may see performance improvements more in this range. See section 5.1 DB2 Partition Scaling on page 19 for more information on this.

The details of these measurements are summarized in the following table.

DB2 level	**DB2 10**	**DB2 11**
Run id	S30221U1	S30516U1
# of Database Partitions	128	128
Database Server	zEC12	zEC12
# of z/OS LPARs	1	1
# of CPs per z/OS LPAR	2	2
Real Storage	309 GB	309 GB
LFAREA (for 1M frames)	90 GB	90 GB
# of Users	2000	2000
# of DB2 Threads	169	169
%CPU on z/OS	84.11%	79.17%
ETR (DS/sec)	186.98	188.55
ITR (DS/sec)	222.30	238.16
Response Time (secs)	0.610	0.521

Table 3: DB2 Single System 5M Accounts (SBS 6.0) Measurement Results

The following figure shows the ITR results from running the SAP Day Posting workload with the SBS 6.0 5M account database with DB2 10 and DB2 11. It shows the 7% improvement in ITR with DB2 11.

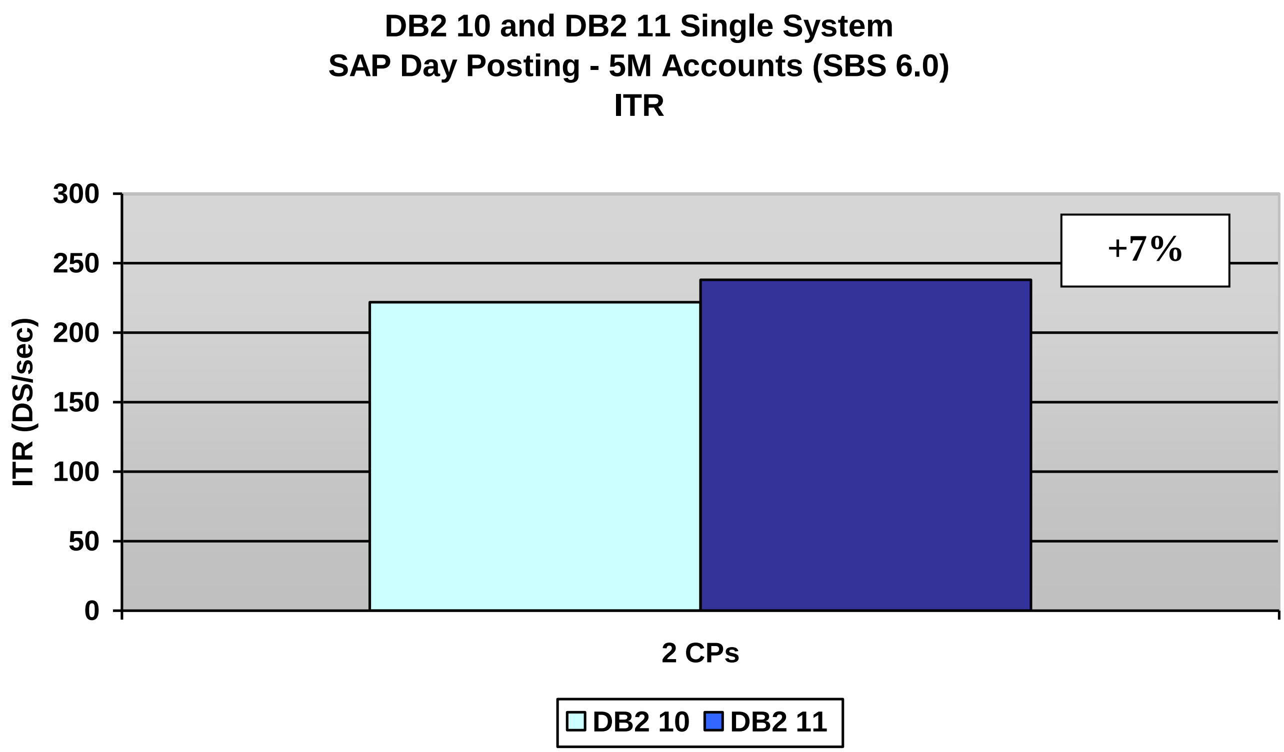

Figure 4: DB2 10 and DB2 11 - 5M Accounts - ITR

5.3 DB2 11 Migration and 10-byte RBA Conversion

This set of measurements was defined to quantify the performance of DB2 11 during the various states of the migration to DB2 11 NFM and the conversion to the 10-byte RBA / LRSN format. For these measurements, we ran the SAP Day Posting workload with the SBS 7.0 database with 60M accounts in a DB2 2-way data sharing environment.

First, we measured the performance of DB2 11 in Conversion Mode (CM) with the APPLCOMPAT zparm set to V10R1. The 6-byte RBA format is the only option in Conversion Mode.

Second, we measured the performance of DB2 11 New Function Mode (NFM). To get to this state, we ran the jobs to migrate from CM to Enabling-New-Function-Mode (ENFM) and from ENFM to NFM. We set the APPLCOMPAT zparm to V11R1 and we rebound the SAP application. We were still using the 6-byte RBA format in this measurement.

Third, we measured the performance of DB2 11 in NFM with the BSDSs for both DB2 members and the DB2 catalog converted to the 10-byte RBA format. This enables logging with the larger RBA / LRSN format. DB2 provides jobs to do these conversions.

Lastly, we measured the performance of DB2 11 in NFM with the table spaces / indexes used by the SAP Day Posting workload converted to the 10-byte RBA format, along with the BSDSs and DB2 catalog. The REORG utility was used to convert the table spaces / indexes to the 10-byte RBA format. The zparms OBJECT_CREATE_FORMAT and UTILITY_OBJECT_CONVERSION should be set to EXTENDED so options do not need to be specified on the REORG statement.

For more information on migrating SAP systems to DB2 11 for z/OS, see reference [2] on page 38.

In the early days of the development of DB2 11, we were concerned that the DB2 internal processing for converting the RBA between the 6-byte and 10-byte format caused degradation in overall system performance. The conversion between the RBA formats is required every time the RBA is accessed since the DB2 11 code uses the 10-byte RBA format, but the RBAs stored in the DB2 catalog, BSDSs, and user databases are in the 6-byte format until they are explicitly converted to 10-bytes.

To compensate for any degradation in overall system performance due to these required RBA conversions at runtime, the DB2 team moved the log buffers into 64-bit common storage. This improved performance by removing cross address space operations between DBM1 and MSTR for logging.

The use of the 10-byte RBA / LRSN is optional, but highly recommended. Using the larger RBA / LRSN not only protects against reaching the limit of the 6-byte RBA, but it also provides LRSN spin avoidance in data sharing environments. The 6-byte LRSN value has precision to only 16 microseconds. This can cause LRSN “spinning” while waiting for a unique value.

DB2 has continued to eliminate the need for unique LRSNs to avoid LRSN spins. Starting in DB2 9, a unique LRSN is no longer necessary for updating different data or index pages for insert or update. In DB2 10, a non-unique LRSN is allowed for inserting into the same pages. However, there is still a need for unique LRSNs for sequential update or delete and some insert processing, and, therefore, there still is potential for spins. Batch updates, inserts, or deletes in a data sharing environment may see significant improvement by avoiding LRSN spins once the BSDS and the objects are converted to use the 10-byte extended LRSN.

In the end, when running the SAP Day posting workload in a 2-way data sharing environment, we found that the performance remained constant throughout the various phases of migration to DB2 11 that we tested. We did not see a benefit from LRSN spin avoidance due to the random access nature of our OLTP workload.

There was some concern that using the 10-byte RBA format would increase the space usage of the database since the RBA is stored on every page in all the table spaces and indexes. So, we compared the space usage by looking at the difference in the number of NPAGES / NLEAF in the key banking tables / indexes between DB2 10 with 6-byte RBAs and DB2 11 NFM with 10-byte RBAs. NPAGES is the number of distinct pages with active rows in the partition or table space. NLEAF is the number of active leaf pages in the index. We did not find a significant increase in space usage. There was less than a half percent increase in the total number of NPAGES / NLEAF in the banking tables / indexes that we converted to use the 10-byte RBA.

However, our measurements did show a 22% increase in the log bytes written when using the extended RBA. Once the BSDS is converted to the 10-byte format, DB2 records the extended RBA (10-byte format) in the log.

The details of these measurements are summarized in the following table.

DB2 level	**DB2 11 CM**	**DB2 11 NFM**	**DB2 11 NFM BSDS and Catalog w/ large RBA**	**DB2 11 NFM Banking tables w/ large RBA**
Run id	S30607B2	S30607B1	S30601B1	S30614B1
# of Database Partitions	512	512	512	512
Database Server	zEC12	zEC12	zEC12	zEC12
# of z/OS LPARs	2	2	2	2
# of CPs per z/OS LPAR	5	5	5	5
Real Storage	235 GB	235 GB	235 GB	235 GB
LFAREA (for 1M frames)	90 GB	90 GB	90 GB	90 GB
# of Users	5120	5120	5120	5120
# of DB2 Threads	523	523	523	523
%CPU on z/OS	69.06%	68.88%	69.34%	69.65%
ETR (DS/sec)	441.20	441.25	443.29	444.40
ITR (DS/sec)	638.86	640.61	639.30	638.05
Response Time (secs)	1.598	1.599	1.540	1.518
Log bytes written per commit	10,457	10,461	12,697	12,799

Table 4: DB2 11 Migration and 10-byte RBA Conversion Measurement Results

The following figure shows the ITR results of the above measurements. It clearly shows the consistent performance across the runs.

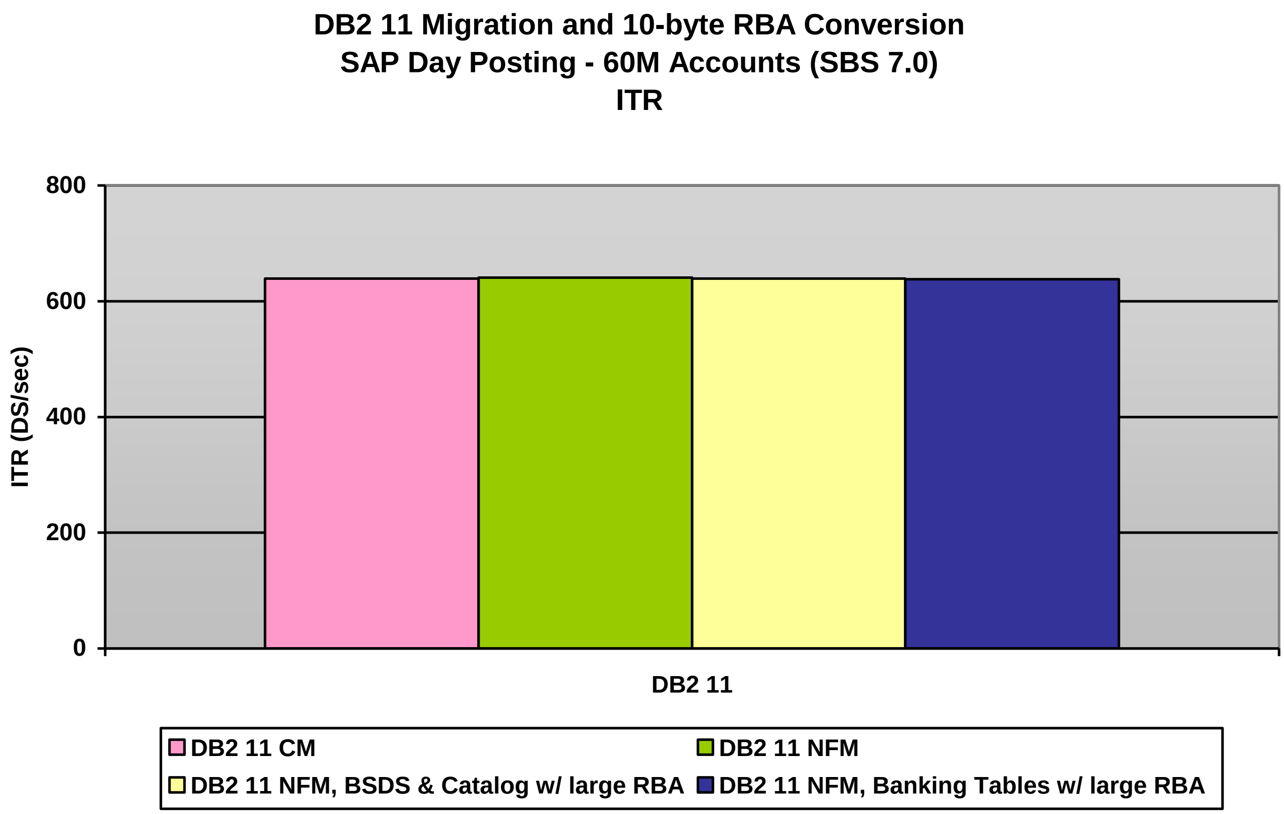

Figure 5: DB2 11 Migration and 10-byte RBA Conversion - ITR

5.4 DB2 Single System – 60M Accounts (SBS 7.0)

In these measurements, we ran the SAP Day Posting workload with the SBS 7.0 database with 60M accounts. We ran with DB2 10 and DB2 11. With DB2 11, we were in NFM using the large 10-byte RBA format. We ran a series of measurements on a zEC12 with a single z/OS LPAR with 2 CPs, 4 CPs, and 8 CPs.

The results we saw were impressive. We measured a 15% improvement in ITR with DB2 11 on the zEC12 with 2 CPs. There was a 16% improvement in ITR with 4 CPs. And, there was an 18% improvement in ITR with 8 CPs. These results show that DB2 11 delivers efficient CP scaling. As the DB2 system got larger, the performance improvement of DB2 11 compared to DB2 10 grew.

A significant contributor to the performance improvements seen with DB2 11 in these measurements was the large number of partitions in this database. The SBS 7.0 60M account database had the key banking tables defined with 512 partitions. Systems with databases with a similar number of partitions to this may see performance improvements more in this range. See section 5.1 DB2 Partition Scaling on page 19 for more information on this.

Another significant contributor to the performance improvements seen in these measurements was the enhancements made in DB2 11 in the area of column processing. DB2 uses customized instructions for repeated operations (xProcs). The new column procedures introduced in DB2 11 are created at prepare and used to pass the output columns to the application. Since a typical SAP application has large numbers of columns in the select list, this optimization is effective especially in SAP applications.

Also contributing to the performance improvements was the moving of the log buffers to 64-bit common storage, eliminating the need for cross address space operations between the DBM1 and MSTR address spaces for logging.

In addition, the DDF Sync Receive mode change contributed to reducing the DB2 and TCP/IP CPU time. This was possible through a new Communications Server Termination Notification added by PM80004. Once the communications server provides a new Connection Termination Notification, DDF is changed to a synchronous receive interface except when the thread is inactive. This provides CPU saving by avoiding suspend/resume operations.

These measurements with the SBS 7.0 60M account database can not be compared with the one done with the SBS 6.0 5M account database. The differences in the SAP levels (SBS 6.0 vs SBS 7.0), the database size (5M accounts vs 60M accounts), and the number of database partitions (128 vs 512) make it like different workloads were used, even though both sets of measurements used the SAP Day Posting workload.

The details of these DB2 single system measurements with the SBS 7.0 60M account database are summarized in the following table.

DB2 level	DB2 10			DB2 11		
Run id	S30910B2	S30819B2	S30819B1	S30910B1	S30911B1	S30813B1
# of Database Partitions	512	512	512	512	512	512
Database Server	zEC12	zEC12	zEC12	zEC12	zEC12	zEC12
# of z/OS LPARs	1	1	1	1	1	1
# of CPs per z/OS LPAR	2	4	8	2	4	8
Real Storage	235 GB	235 GB	235 GB	235 GB	235 GB	235 GB
LFAREA (for 1M frames)	90 GB	90 GB	90 GB	90 GB	90 GB	90 GB
# of Users	1792	3360	6000	1792	3360	6000
# of DB2 Threads	266	458	778	266	458	778
%CPU on z/OS	83.61%	84.27%	83.90%	74.21%	72.47%	74.12%
ETR (DS/sec)	164.77	294.63	509.44	168.39	294.97	531.08
ITR (DS/sec)	197.07	349.63	607.20	226.91	407.02	716.51
Response Time (secs)	0.876	1.404	1.778	0.642	1.391	1.298

Table 5: DB2 Single System 60M Accounts (SBS 7.0) Measurement Results

The following figure shows the ITR results from running the SAP Day Posting workload with the SBS 7.0 60M account database on a zEC12 with a single z/OS LPAR with three different configurations with both DB2 10 and DB2 11.

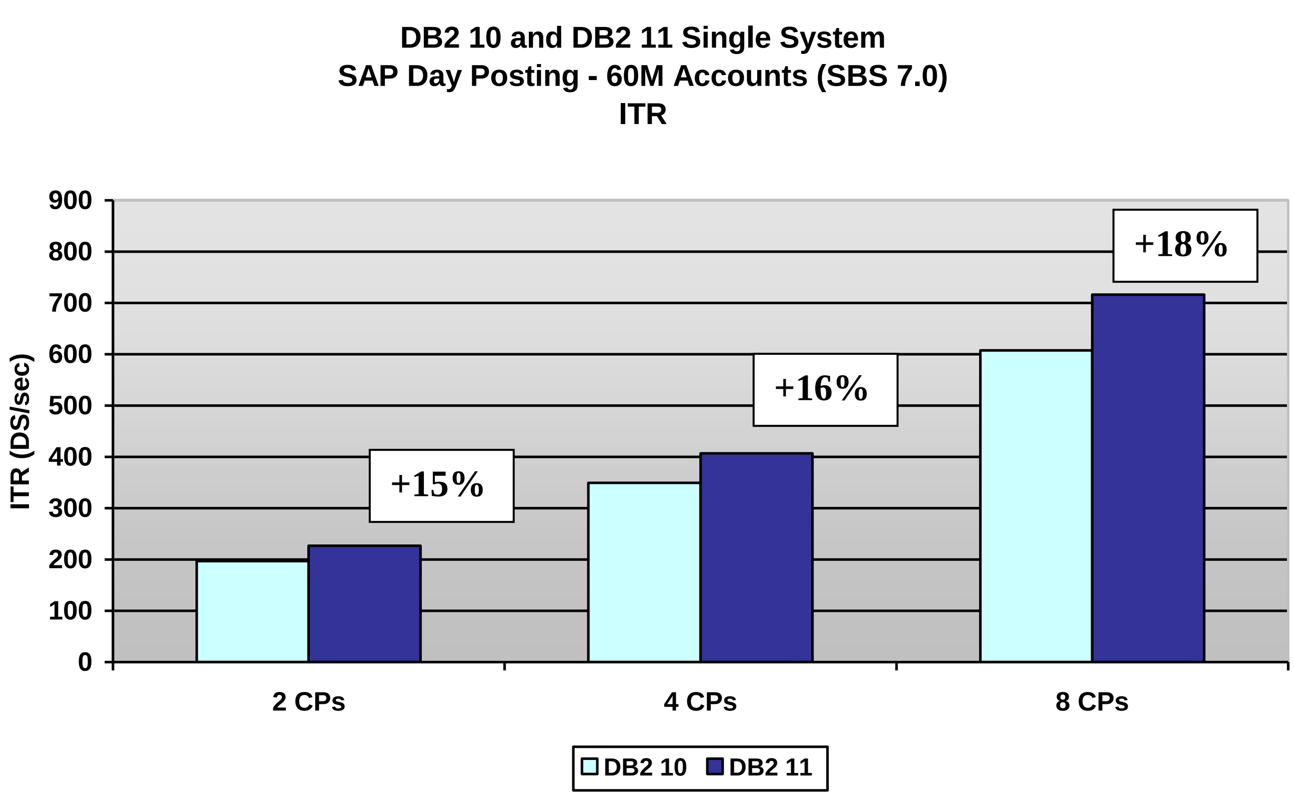

Figure 6: DB2 10 and DB2 11 Single System - 60M Accounts - ITR

We also looked at the virtual storage and real storage usage of DB2 11 compared to DB2 10 using the largest configuration in this set of single system measurements.

DB2 10 provided the significant relief in virtual storage usage below the 2GB bar in DBM1. However, DB2 11 continues the trend to reduce the usage of below the bar storage, such as by moving xProc storage above the bar. There was a 48% reduction in DBM1 below the bar virtual storage usage with DB2 11 compared to DB2 10 and equivalent real storage usage. There was no difference in virtual storage usage in the DIST address space between DB2 10 and 11.

The following table shows the storage usage details.

DB2 level	**DB2 10**	**DB2 11**	Delta
Run id	S30819B1	S30813B1	
Database Server	zEC12	zEC12	
# of z/OS LPARs	1	1	
# of CPs per z/OS LPAR	8	8	
DBM1 Below 2G (MB)	224	115	-48.66%
Number of user threads	778	778	0.00%
DIST Below 2GB (MB)	183	184	0.55%
Buffer Pool Allocation (MB)	77344	77426	0.11%
Real Storage In Use (MB)	81686.77	81734.28	0.06%
31 BIT In Use (MB)	373.63	239.74	-35.83%
64 BIT In Use (MB)	81313.15	81494.53	0.22%
Real Storage other than buffer pools (MB)	4343	4308	-0.79%

Table 6: DB2 Single System Virtual and Real Storage Usage

5.5 DB2 2-way Data Sharing – 60M Accounts (SBS 7.0)

In these measurements, we ran the SAP Day Posting workload with the SBS 7.0 database with 60M accounts in a DB2 2way data sharing environment. We ran with DB2 10 and DB2 11. With DB2 11, we were in NFM using the large 10-byte RBA / LRSN. We ran these measurements on a zEC12 with two z/OS LPARs, each with 6 CPs. See page 16 for more information on this environment. There was a 9% improvement in ITR with DB2 11.

We did not see the performance benefit of the full LRSN spin avoidance delivered in DB2 11 in this data sharing measurement due to the random access nature of the Day Posting workload. However, the benefits of this should be seen in batch workloads with more sequential access.

We saw similar trends in virtual storage and real storage usage in 2-way data sharing as we saw in single system. There was a 31% reduction in virtual storage usage in DB2 11 compared to DB2 10. There was a slight increase in real storage usage for “other than buffer pools” with DB2 11.

The details of these measurements are summarized in the following table.

DB2 level	**DB2 10**	**DB2 11**	**Delta**
Run id	S30717B1	S30718B1	
# of Database Partitions	512	512	
Database Server	zEC12	zEC12	
# of z/OS LPARs	2	2	
# of CPs per z/OS LPAR	6	6	
Real Storage	235 GB	235 GB	
LFAREA (for 1M frames)	90 GB	90 GB	
# of Users	6400	6400	
# of DB2 Threads	651	651	
Avg %CPU on z/OS	78.82%	71.04%	
ETR (DS/sec)	534.63	525.90	
ITR (DS/sec)	678.29	738.88	9%
Response Time (secs)	1.958	2.192	
DBM1 Below 2G (MB) - Avg	138.5	96.0	-31%
DIST below 2GB (MB) - Avg	77.0	78.5	2%
Buffer Pool Allocation (MB) - Avg	77364.5	77427.0	0%
Real Storage In Use (MB) – Avg	81702.81	81967.43	0%
31 BIT In Use (MB) - Avg	295.63	230.42	-22%
64 BIT In Use (MB) - Avg	81407.19	81737.01	0%
Real Storage other than buffer pools (MB) – Avg	4338	4540	5%

Table 7: DB2 2-way Data Sharing 60M Accounts (SBS 7.0) Measurement Results

The following figure shows the ITR results from running the SAP Day Posting workload with the SBS 7.0 60M account database with DB2 10 and DB2 11 in a 2way data sharing environment.

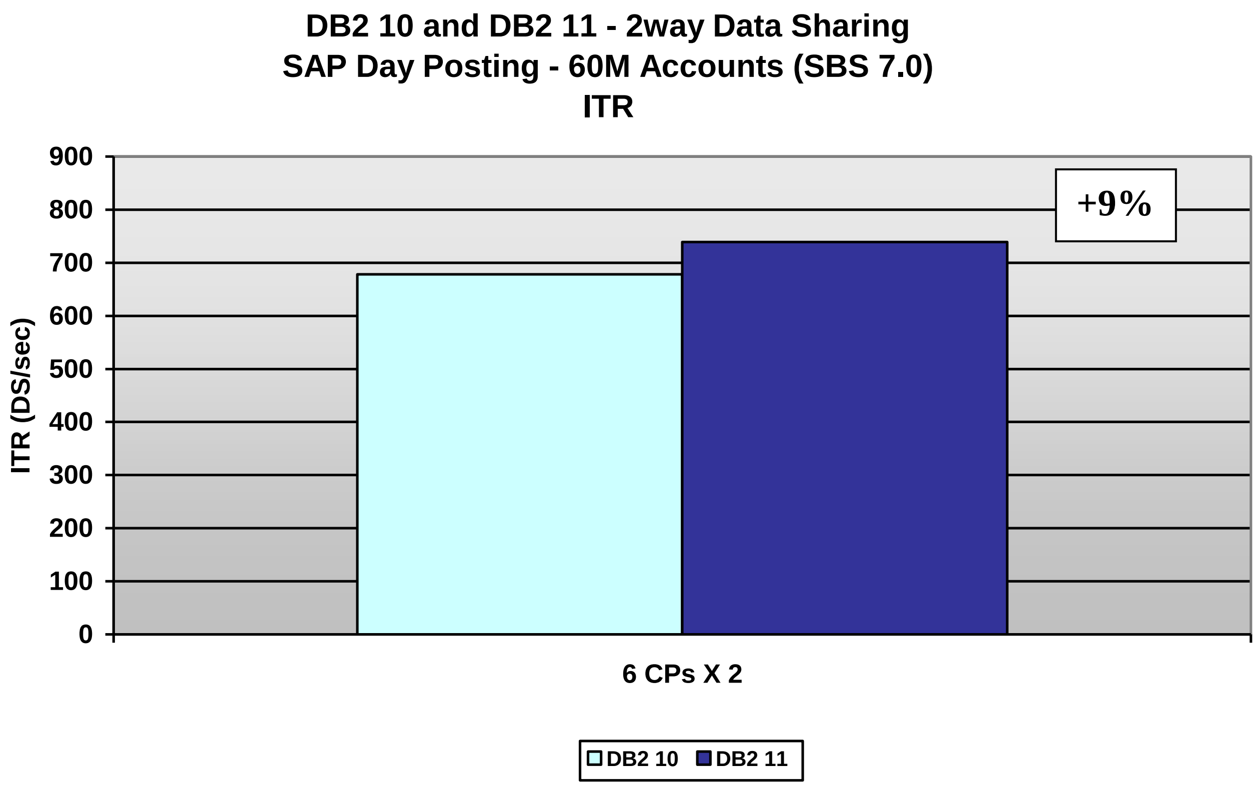

Figure 7: DB2 10 and DB2 11 2way Data Sharing - 60M Accounts - ITR

5.6 zIIP Exploitation

The IBM System z Integrated Information Processor (zIIP) is a special purpose processor and is available on all IBM zEnterprise servers. It 's is designed to help free-up general computing capacity and lower overall total cost of computing for select data and transaction processing workloads. The general types of DB2 work eligible for zIIP exploitation in DB2 10 include a portion the following:

- Remote DRDA access via TCP/IP
- Parallel query operations
- XML parsing in DB2
- Certain DB2 utilities processing
- Certain system level tasks

See A.0 Appendix: DB2 11 for z/OS zIIP Eligible Processing on page 39 for the “word” from DB2 on the additional processing in DB2 11 that will be authorized to execute on a zIIP.

The SAP Day Posting workload accesses DB2 for z/OS via DRDA over TCP/IP connections. Therefore, it is an eligible zIIP exploiter on the zEC12 server. To assess its zIIP exploitation potential in DB2 11 compared to DB2 10, we ran a set of measurements in both single system and 2-way data sharing environments.

5.6.1 zIIP Exploitation – Single System

In a DB2 single system environment, we ran the SAP Day Posting workload with the SBS 7.0 60M account database with DB2 10 and DB2 11 on a zEC12. One set of runs was on a single z/OS LPAR with 8 general purpose CPs. The other set of runs was on a single z/OS LPAR with 4 general purpose CPs and 4 zIIPs.

Comparing the DB2 10 and DB2 11 measurements done with 4 general purpose CPs and 4 zIIPs, we saw a reduction in the CP and the zIIP utilization with DB2 11. We also saw that more work is redirected to zIIPs in DB2 11, specifically 53% compared to 50%. The DB2 MSTR address space shows more zIIP redirection in DB2 11 mainly due to asynchronous log write processing.

Comparing the measurements that use a mix of general purpose CPs and zIIPs with those that use just general purpose CPs, we saw that using a mix provided equivalent or better performance. This is true for both DB2 10 and DB2 11.

The details of these measurements are summarized in the following table.

DB2 level	**DB2 10**		**DB2 11**	
Run id	S30819B1	S30912B2	S30813B1	S30913B1
Data Sharing Members	1	1	1	1
# of Database Partitions	512	512	512	512
Database Server	zEC12	zEC12	zEC12	zEC12
# of z/OS LPARs	1	1	1	1
# of CPs per z/OS LPAR	8	4	8	4
# of zIIPs per z/OS LPAR	0	4	0	4
# of Users	6000	6000	6000	6000
# of DB2 Threads	778	778	778	778
CP util %	83.90%	80.56%	74.12%	68.70%
zIIP util %	n/a	83.27%	n/a	78.42%
zIIP redirect %	n/a	50.83%	n/a	53.30%
DDF zIIP redirect %	n/a	55%	n/a	58%
DBM1 zIIP redirect %	n/a	92%	n/a	98%
MSTR zIIP redirect %	n/a	0%	n/a	93%
Avg %CPU on z/OS	83.90%	81.92%	74.12%	73.56%
ETR (DS/sec)	509.44	515.76	531.08	527.58
ITR (DS/sec)	607.20	629.63	716.51	717.21
Response Time (secs)	1.778	1.633	1.298	1.373

Table 8: DB2 Single System zIIP Measurement Results

The next two figures show results from running the SAP Day Posting workload with DB2 10 and DB2 11 on a single z/OS LPAR with 4 general purpose CPs and 4 zIIPs. The first figure shows the utilization of the general purpose CPs and zIIPs. The second figure shows the amount of zIIP redirects.

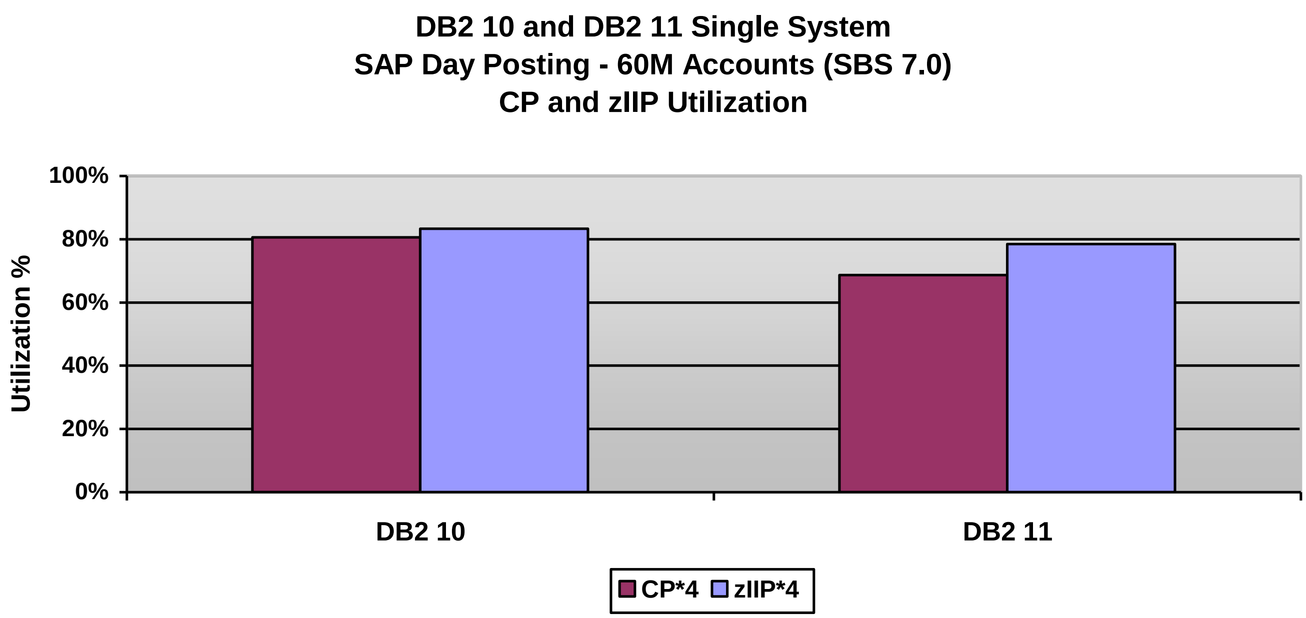

Figure 8: DB2 Single System – zIIP Measurements – CP and zIIP Utilization

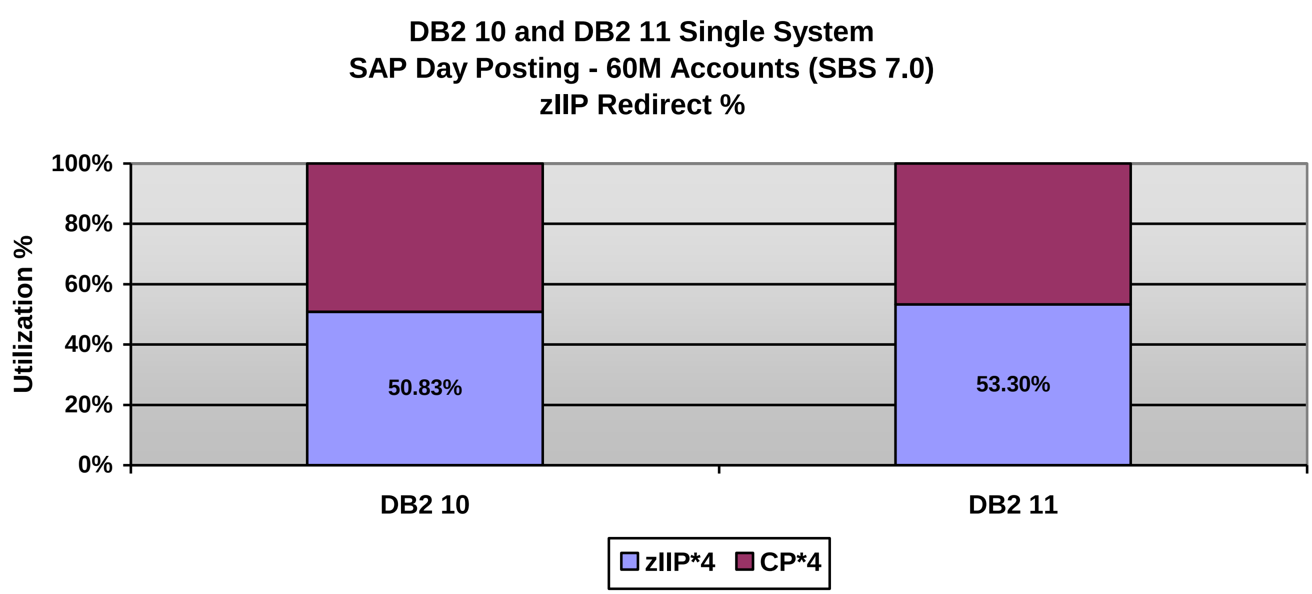

Figure 9: DB2 Single System – zIIP Measurements – zIIP Redirect

The following figure shows the ITR results from running the SAP Day Posting workload using just general purpose CPs and a mix of general purpose CPs and zIIPs with DB2 10 and DB2 11 in a single system environment. It shows that using a mix of general purpose CPs and zIIPs provided equivalent or better performance.

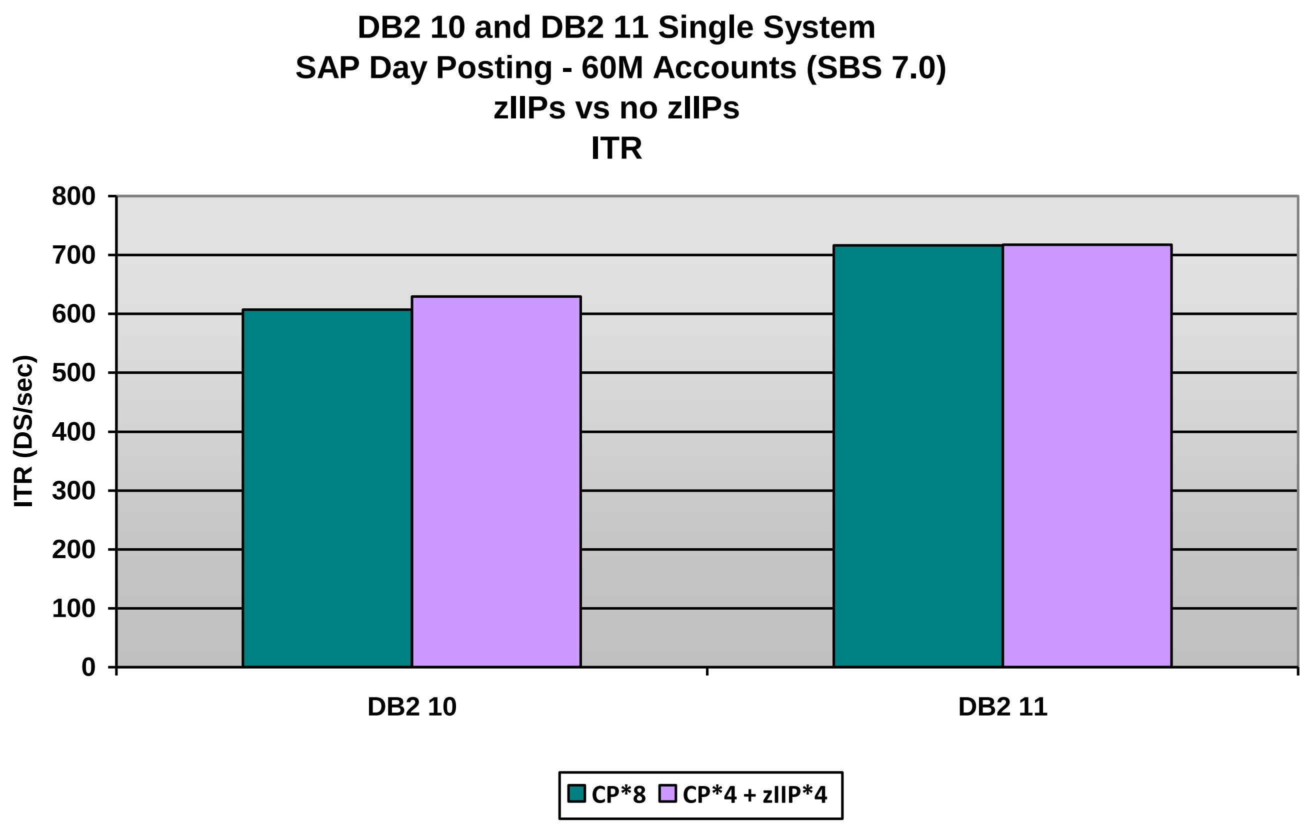

Figure 10: DB2 Single System – zIIPs vs no zIIPs - ITR

5.6.2 zIIP Exploitation – 2-way Data Sharing

In a 2-way data sharing environment, we ran the SAP Day Posting workload with the SBS 7.0 60M account database with DB2 10 and DB2 11 on a zEC12. One set of runs was on two z/OS LPARs, each with 6 general purpose CPs. The other set of runs was on two z/OS LPARs, each with 3 general purpose CPs and 3 zIIPs.

Comparing the DB2 10 and DB2 11 measurements done with a mix of general purpose CPs and zIIPs, we saw a reduction in the CP and the zIIP utilization with DB2 11. We also saw that more work is redirected to zIIPs in DB2 11, specifically 48% compared to 52%. The DB2 MSTR address space shows more zIIP redirection in DB2 11 mainly due to asynchronous log write processing.

Comparing the measurements that use a mix of general purpose CPs and zIIPs with those that use just general purpose CPs, we saw that using a mix provided better performance in DB2 10 and DB2 11.

The details of these measurements are summarized in the following table.

DB2 level	**DB2 10**		**DB2 11**	
Run id	S30717B1	S30809B1	S30718B1	S30808B1
Number of Data Sharing Members	2	2	2	2
# of Partitions	512	512	512	512
Database Server	zEC12	zEC12	zEC12	zEC12
# of z/OS LPARs	2	2	2	2
# of CPs per z/OS LPAR	6	3	6	3
# of zIIPs per z/OS LPAR	0	3	0	3
# of Users	6400	6400	6400	6400
# of DB2 Threads	651	651	651	651
CP util %	78.82%	78.75%	71.04%	62.84%
zIIP util %	n/a	72.50%	n/a	68.37%
zIIP redirect %	n/a	47.93%	n/a	52.11%
DDF zIIP redirect %	n/a	55%	n/a	58%
DBM1 zIIP redirect %	n/a	30%	n/a	83%
MSTR zIIP redirect %	n/a	0%	n/a	95%
Avg %CPU on z/OS	78.82%	75.63%	71.04%	65.61%
ETR (DS/sec)	534.63	538.45	525.90	519.63
ITR (DS/sec)	678.29	712.00	738.88	792.06
Response Time (secs)	1.958	1.887	2.192	2.311

Table 9: DB2 2-way Data Sharing zIIP Measurement Results

The next two figures show results from running the SAP Day Posting workload with DB2 10 and DB2 11 on two z/OS LPARs, each with 3 general purpose CPs and 3 zIIPs. The first figure shows the utilization of the general purpose CPs and zIIPs. The second figure shows the amount of zIIP redirects.

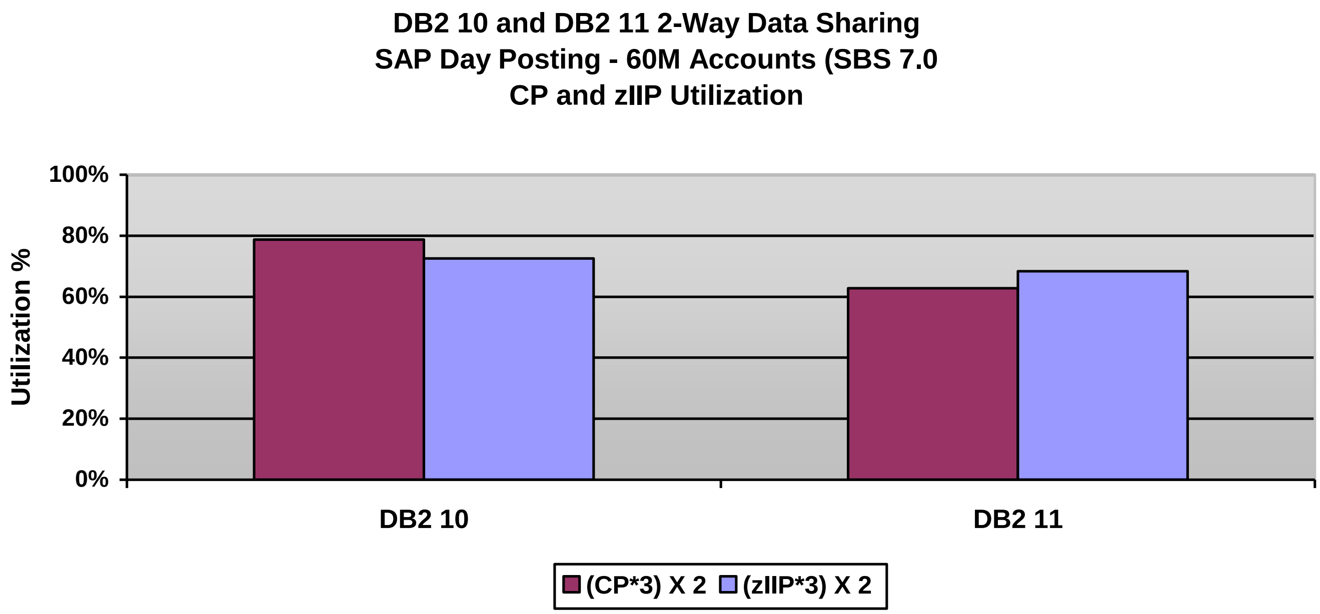

Figure 11: DB2 2-way Data Sharing – zIIP Measurements – CP and zIIP Utilization

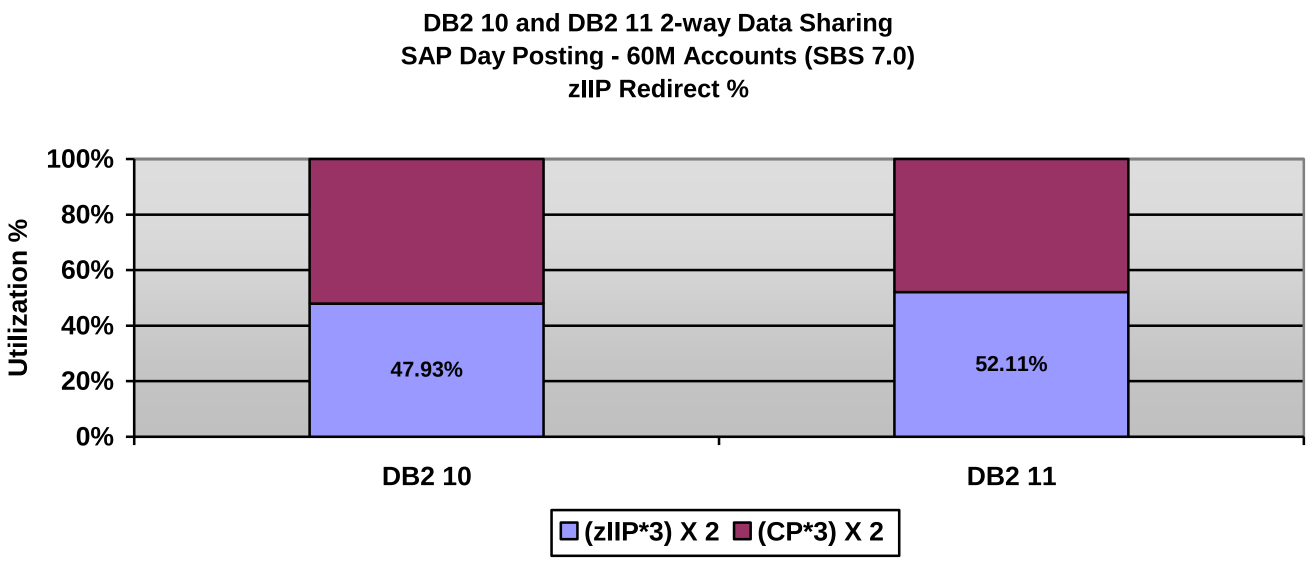

Figure 12: DB2 2-way Data Sharing – zIIP Measurements - zIIP Redirect

The following figure shows the ITR results from running the SAP Day Posting workload using just general purpose CPs and a mix of general purpose CPs and zIIPs with DB2 10 and DB2 11 in a 2-way data sharing environment. It shows that using a mix of general purpose CPs and zIIPs provided better performance.

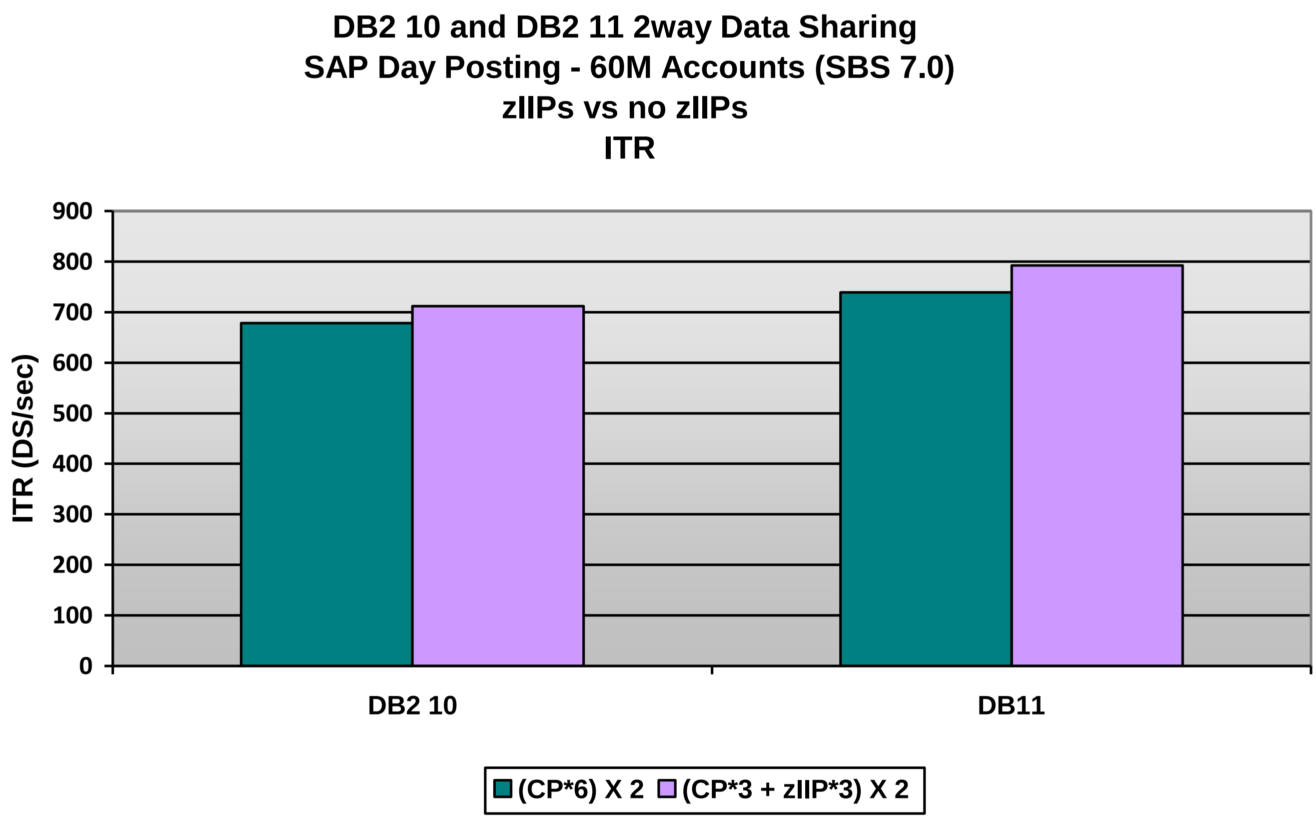

Figure 13: DB2 2-way Data Sharing – zIIPs vs no zIIPs - ITR

6.0 Conclusions

DB2 11 for z/OS delivers significant performance and scalability improvements.

- We saw performance improvements in the range of 7% to 18% in a DB2 single system environment.
- We saw up to an 11% improvement in performance in a 2-way data sharing environment.
- Optimizations to column processing are particularly beneficial to SAP applications which fetch large numbers of columns.
- The exploitation of the synchronous TCP/IP receive model by DDF benefits SAP applications.
- Using 1M fixed large frame support for buffer pools that are page fixed is still recommended.
- DB2 11 continues the trend to reduce virtual storage usage below the bar. Real storage usage remains relatively flat.
- DB2 11 delivers more efficient partition scaling. A larger number of partitions can be used without worrying about adverse performance effects.
- DB2 11 provides extended RBA and LRSN addressing support to 1 yottabyte. Converting from the 6-byte RBA to the 10-byte extended RBA is highly recommended for scalability and performance.
- Latch class contention was reduced, especially for latch class 14, relating to the buffer manager.
- DB2 11 supports a larger local dynamic statement cache. The new limit for the MAXKEEPD zparm is 200K.
- The group buffer pool castout threshold (CLASST) can be set to an absolute number of pages, as well as to a percentage of the buffer pool size.
- More work is eligible to be executed on zIIPs. Using zIIPs reduces costs and provides equal or better performance.

7.0 References

1. IBM Corp. 2013. *IBM United States Software Announcement 213-376, IBM DB2 11 for z/OS: The database for data and analytics* http://www-01.ibm.com/common/ssi/rep_ca/6/897/ENUS213-376/ENUS213-376.PDF
2. *SAP Best-Practices Document: Migrating SAP Systems to DB2 11 for z/OS* https://websmp108.sap-ag.de/~sapidb/011000358700001057822013E (requires SAP user id)
3. *SAP DBA Guide: DB2 11 for z/OS), Database Administration Guide: SAP on IBM DB2 for z/OS* https://websmp108.sap-ag.de/~sapidb/011000358700001053572013E (requires SAP user id)
4. IBM Corp. 2012. *IBM System zEnterprise, System Storage, and DB2 10 for z/OS: SAP Banking Services 7.0 150 Million Accounts Measurements* http://www-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP101978
5. IBM Corp. 2010. *SAP® Transaction Banking: Case Study Understanding Data Sharing Effects as a Function of Long Distance Coupling Links IBM System z®* http://www-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP101692
6. IBM Corp. 2013. *Large Systems Performance Reference*, Document Number SC28-1187-17 https://www-304.ibm.com/servers/resourcelink/lib03060.nsf/pages/lsprindex/$file/SC28118717.pdf
7. IBM Corp. 2011. *DB2 10 for z/OS with SAP on IBM System z Performance Report* http://www-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP101845
8. IBM Corp. 2013. *DB2 11 for z/OS Technical Overview* http://www.redbooks.ibm.com/Redbooks.nsf/RedpieceAbstracts/sg248180.html?Open
9. SAP AG 2011. *SAP Banking on System z Reference Architecture* http://www.sdn.sap.com/irj/sdn/db2?rid=/library/uuid/a00e4718-314f-2b10-19a6-a76f257addaf

A.0 Appendix: DB2 11 for z/OS zIIP Eligible Processing

The following is [additional] DB2 11 for z/OS processing that will be authorized to execute on a zIIP:[1]

Asynchronous processing that is executed under enclave SRBs (Service Request Blocks) and that will be "charged" for CPU consumption purposes to an IBM DB2 address space (rather than to user applications), with the exception of P-lock negotiation processing.
Such zIIP eligible processing includes:
1) Cleanup of pseudo deleted index entries as part of DB2 system task cleanup
2) Cleanup of XML multi-version documents (available in DB2 10 for z/OS via APAR PM 72526)
3) Log write and log read

The DB2 base LOAD, REORG and REBUILD INDEX utility processing of inline statistics collection that DB2 directs to be executed under enclave SRBs (Service Request Blocks)[2]

The DB2 base processing of RUNSTATS utility Column Group Distribution statistics collection that DB2 directs to be executed under enclave SRBs (Service Request Blocks)[2]

The DB2 base LOAD utility index management processing when running LOAD REPLACE that DB2 directs to be executed under enclave SRBs (Service Request Blocks)[2]

[1] This information provides only general descriptions of the types and portions of workloads that are eligible for execution on Specialty Engines (e.g, zIIPs, zAAPs, and IFLs) ("SEs"). IBM authorizes customers to use IBM SE only to execute the processing of Eligible Workloads of specific Programs expressly authorized by IBM as specified in the “Authorized Use Table for IBM Machines” provided at www.ibm.com/systems/support/machine_warranties/machine_code/aut.html (“AUT"). No other workload processing is authorized for execution on an SE. IBM offers SE at a lower price than General Processors/Central Processors because customers are authorized to use SEs only to process certain types and/or amounts of workloads as specified by IBM in the AUT.

[2] DB2 does not direct all such base utility processing to be executed under enclave SRBs.