IBM Technical Brief

IBM z Systems®: Performance Report on Exploiting Large Memory for DB2 Buffer Pools with SAP®

Version 3.0

February 27, 2015

Authors:

Brenda Beane
Seewah Chan
Dr. Paul Lekkas

Document Owner:

Brenda Beane
IBM SAP on System z Performance
Version: 3.00
Date: February 27, 2015
Filename: Large_Memory_withSAP.pdf

Table of Contents

Tables

Figures

Disclaimers

Trademarks

Acknowledgements

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

- Akiko Hoshikawa for her DB2 analysis support
- Tom Litty for his z/OS system programming support
- Dave Surman for his Parallel Sysplex analysis support
- Steve Turnbaugh for his DB2 development support

Feedback

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

Change History

Version 1.0 is the original document.

Version 2.0 is the original document with the following updates.

- Added results for measurements with the SAP Night Balancing workload

Version 3.0 is version 2 of the document with the following updates.

- Added results for measurements with the SAP Day Posting workload on the new IBM z13 processor

1.0 Introduction

It is reasonable to assume that using more memory will improve overall system performance. One obvious way to make use of more memory is to allocate more memory to DB2 buffer pools. Reducing I/O by caching more data in buffer pools should reduce response time, increase throughput, and provide CPU savings.

The IBM zEnterprise 196 and EC12 systems can support up to 3 TB of real memory per server and 1 TB per LPAR. The IBM z13 system, announced in January 2015, supports up to 10 TB of real memory per server, more than three times that of the two previous generations of systems, and IBM has announced plans to support up to 4 TB of real memory per LPAR in the future. Currently, DB2 10 and 11 for z/OS theoretically allow up to 1 TB of memory to be used for all buffer pools in a given member. Special pricing for memory on the z13 system makes using larger amounts of memory more attractive and affordable.

The IBM SAP on System z Performance Team, located in Poughkeepsie, NY, conducted a number of experiments to evaluate the performance effects of using large amounts of memory for DB2 buffer pools. We systematically increased the sizes of DB2 buffer pools and measured the effects on system performance.

We used the SAP Banking Services (SBS) Day Posting and Night Balancing workloads. Day Posting is an online transaction processing (OLTP) workload, which is memory intensive, accesses a large number of tables, and exhibits random I/O behavior. Night Balancing is a batch processing workload, which is insert intensive and creates large volumes of logs. And, data processing is done sequentially in a clustering order and the data is processed only once. Both of these workloads are good representations of customer workloads. See section 3.0 *Workload* Descriptions for details of these workloads.

We ran with DB2 11 for z/OS in both single system and data sharing environments. In data sharing, the group buffer pools on the coupling facility (CF) provide an extra layer of caching. We explored the effects of adding memory to both the local buffer pools and the group buffer pools.

We also experimented with reducing the number of buffer pools used. Although isolating or separating objects into their own buffer pools can provide essential monitoring capabilities and performance optimizations, it can also produce a large number of buffer pools over time which can be labor intensive and time consuming to maintain and tune. We conducted a set of experiments to evaluate the performance impact of reducing the number of buffer pools with hopes of simplifying buffer pool management without adversely affecting performance.

With the introduction of the IBM z13 system, we repeated a couple of the measurements that we did on the zEC12 on a z13 to validate that the results would be similar.

This paper documents our tests and findings. It is not intended to be a guide to tuning DB2 buffer pools. See references [1,2] on page 59 for DB2 buffer pool tuning guidelines. The measurements that were done were stress tests, not certified benchmarks.

2.0 Executive Summary

Significant performance improvements were seen when more memory was used for larger DB2 buffer pools while running the SAP Banking Services (SBS) Day Posting workload on an IBM zEnterprise EC12 (zEC12). Improvements were seen when running in single system and data sharing environments. Similar performance improvements were also seen while running the SAP Day Posting workload on an IBM z13 system.

Our measurements showed reductions in response time of up to 70%, increases in transaction rates of up to 37%, and savings in CPU time per transaction of up to 25%. These performance improvements were a result of up to a 97% reduction in synchronous I/O. Note that in this document, the terms External Throughput Rate (ETR) and Internal Throughput Rate (ITR) are used. ETR is the transaction rate. ITR is the ETR normalized to 100% CP utilization and it gives a relative CPU time per transaction. See reference [4] on page 59 for more of an explanation of these terms.

Our measurements with the SBS Night Balancing workload did not show much, if any, performance improvements with additional memory for the DB2 buffer pools. This is mostly due to the workload being sequential in nature with little re-reference of the data in the buffer pools. Although adding memory to the DB2 buffer pools reduced the amount of synchronous I/O, it was not enough to make an impact on performance.

We found that the performance improvements that can be achieved by increasing the sizes of DB2 buffer pools are very much configuration and workload dependent. Reducing the amount of synchronous I/O is key to achieving overall system performance improvements. Many of our clients should see measurable performance benefits by adding memory to their DB2 buffer pools, however, it is important to note that performance benefits will vary and may not be seen in all environments.

We saw significant performance improvements with the SAP Day Posting workload since its data access pattern is highly random and it is a “pure DB2 workload”. Only DB2 is running on the z System, the application processing is done on another machine. Our measurements were run in a physical 3-tier environment where each of the three tiers (database server, application server, and presentation server) were on separate machines. Our SAP database server was an IBM zEC12 or z13, depending on the test scenario, with 12 general purpose processors.

Tuning buffer pools to minimize I/O is a very worthwhile task that can improve overall system performance. An essential part of buffer pool tuning can be to use additional memory to increase the size of buffer pools. In data sharing environments, buffer pool tuning involves tuning both the local and group buffer pools.

To help simplify buffer pool management and tuning, we experimented with using a smaller number of buffer pools. We found that there was essentially no impact on ITR performance when using a smaller number of buffer pools with a total buffer pool size of up to 320 GB. This is great news for clients who want to simplify their buffer pool strategy. It also supports our out-of-the-box installation strategy. See reference [6] on page 59 for the initial recommended DB2 buffer pool settings for SAP installations.

Late in 2014, DB2 11 for z/OS introduced support for buffer pool simulation which allows you to simulate a larger buffer pool size while running your regular workloads to see how the larger size would affect the number of synchronous reads. This support is delivered in APAR PI22091. Good candidates for buffer pool simulation are buffer pools with significant number of synchronous reads and that contain pages that are likely to be referenced again. DB2 will report an estimate of the number of sync reads that could be avoided if the larger buffer pool is used. For more information about DB2 buffer pool simulation, see reference [8] on page 59.

Based on the number of avoidable sync read I/O reported from the buffer pool simulation, you can estimate CPU time savings. In our measurements on zEC12 and z13, we saw that each avoided synchronous read I/O saves between 20 and 40 microseconds of CPU time. See reference [9] on page 59 for how to translate the I/O savings into percent CPU time saved.

In our measurements, our database was on an IBM System Storage DS8870 server with only Hard Disk Drives (HDD). The DS8870 server has a Solid State Drive (SSD) option. This feature may provide better disk response time for workloads with high database I/O, which can improve overall transaction response time. However, merely replacing HDDs with SSDs doesn't reduce the number of synchronous I/Os. Therefore, the CPU savings that we observed in this study using HDDs should also be seen with SSDs when synchronous I/Os are reduced by adding memory to DB2 buffer pools. See reference [7] on page 59 for the results of a study comparing the performance of HDDs and SSDs using SAP Banking workloads.

3.0 Workload Descriptions

Two different workloads were used in these tests. One was the SAP Banking Services Day Posting workload and the other was the SAP Banking Services Night Balancing workload. Day Posting is an online transaction processing (OLTP) workload. Night Balancing is a batch processing workload. We have been running these workloads for many years. We have quite a bit of experience with them. See reference [3] on page 59.

3.1 SAP Day Posting

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

3.2 SAP Night Balancing

This workload performs account balancing on all accounts in the database. Account balancing is done periodically for each account to calculate charges (fees) and interest as well as posting these. It is usually a series of batch jobs executed at night once a month or quarter depending on the institution and the type of account. Within each batch job, the data is processed for each account and the same account is not re-referenced. Usually account balancing is done for all accounts as soon as possible after the period close such as month-end.

This workload was also developed by SAP. All the accounts have identical attributes and characteristics. Before doing the account balance, each account has 20 items posted on 20 workdays. The key metric of throughput is accounts balanced per unit time.

4.0 Test Environment

Since the focus of these tests is on the database server, a physical 3-tier environment where each of the three layers resided on separate machines was used. The SAP Database Server was on an IBM zEC12 or a z13, depending on the test scenario, running z/OS. The SAP application servers were IBM Power7 Blade Servers running AIX. The presentation server was an IBM Power5 p55A running AIX.

The following figure provides an overview of our test environment.

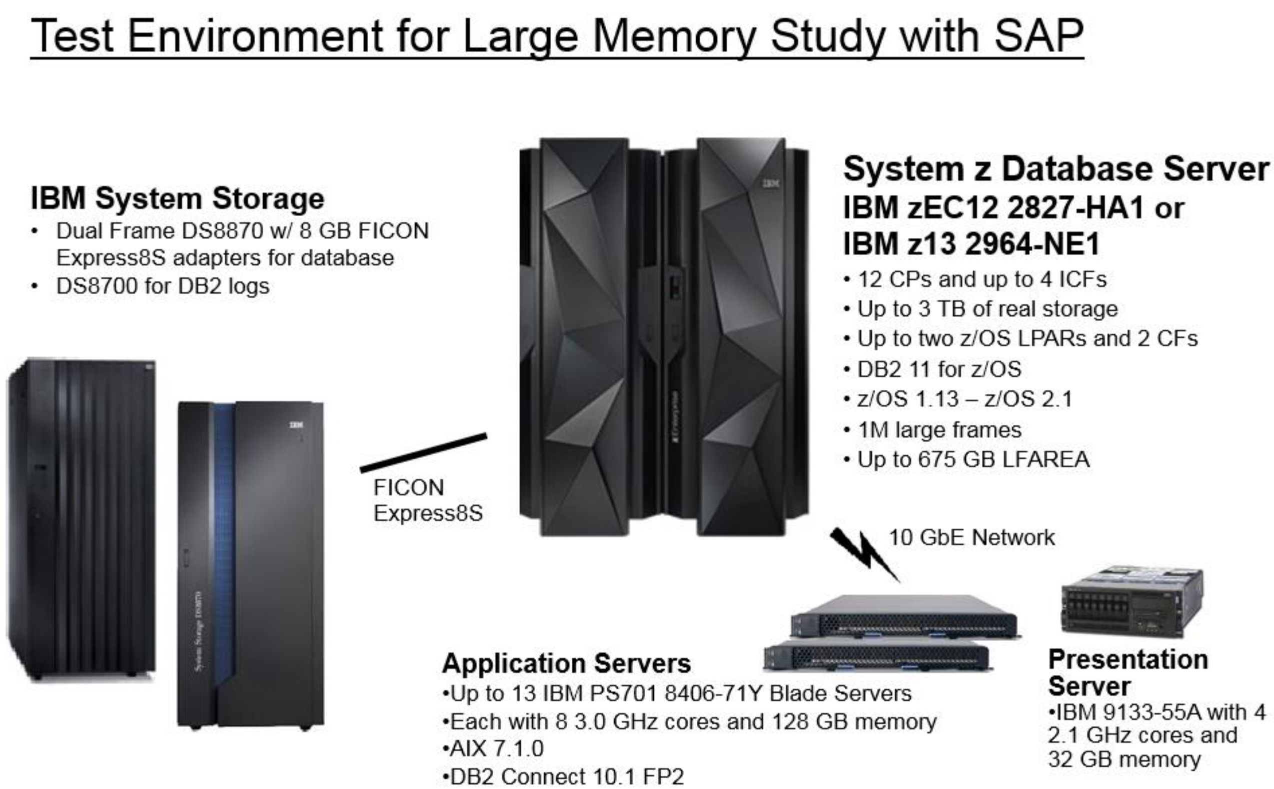

Figure 1: Test Environment – Large Memory Study with SAP

4.1 Hardware

System z Database Server

An IBM zEnterprise Class 2827-HA1 (zEC12) or an IBM z Systems 2964-NE1 (z13), each with 12 dedicated CPs and up to 4 internal coupling facility (ICF) processors configured online, was used for these tests. These systems had 3 TB of real storage. Varying amounts of real storage was configured online depending on the test scenario.

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, up to 675 GB of real storage was set aside for 1M fixed large frames. Using 1M fixed large frame support for DB2 buffer pools is recommended and is identified by IBM as a “best practice” for performance and CPU utilization.

Database DASD

A SBS 7.0 database with 60M accounts was used for these tests. It resided on a dual frame IBM System Storage DS8870 (2107-961) server with 60 ranks of 15K-RPM 146 GB Hard Disk Drives (HDDs) configured as RAID 5. The effective total capacity was 896 3390-mod54 volumes, about 46 TB. The unit has 256 GB regular cache, 8 GB non-volatile storage (NVS), and 16 long wave FICON Express8S attachments. The DB2 subsystem, the database, and two flash copies were contained on 834 emulated 3390-mod54 volumes.

The DB2 active logs resided on a separate DASD unit from the database. The active logs 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.

Even though in this document we are highlighting the benefits of reducing I/O, it is still important to maintain a good DASD I/O subsystem. By adding memory to DB2 buffer pools, DB2 synchronous reads can be reduced, but there still will be DB2 asynchronous I/O, including prefetch and deferred writes, as well as the critical DB2 synchronous logging I/O.

Application Servers

Up to 25 IBM PS701 8406-71Y Blade Servers running AIX were used for the SAP application servers. Each of these Power7 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. Up to four SAP dialog instances were on each of the remaining blades. Four instances per blade were used for the Day Posting workload. Two instances per blade were used for the Night Balancing workload. The actual number of SAP application servers used depended on the configuration, single system or data sharing.

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 1.13 for Single System measurements with Day Posting
z/OS release 2.1 for all the other measurements

DB2 for z/OS

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

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

4.3 DB2 Configuration

Buffer pools

- All buffer pools were page fixed (PGFIX=YES). This means that the page frames that hold the buffer pool will be fixed in real storage and can't be stolen by z/OS. This is good for buffer pools that have I/O activity, provided that there is enough real storage available. This is highly recommended for improved performance.
- 1M fixed large frame support was used for the buffer pools that were page fixed. Using 1 MB frames helps the z/OS Real Storage Manager (RSM) efficiently manage the overall real storage. This is highly recommended for improved performance.
- The sizes of the buffer pools varied depending on the test scenario.
- All objects used GBPCACHE CHANGED. This means that updated (or changed) pages are written to the group buffer pool when there is inter-DB2 R/W interest on the object.

Buffer Pool Assignments

Two different buffer pool configurations were used in this study. One is called "isolated buffer pools" and the other is called "consolidated buffer pools".

The **"isolated buffer pool"** configuration was developed over many years of running the SAP Banking Services Day Posting and Night Balancing workloads. This configuration has 43 buffer pools defined. Key objects used by the banking applications are separated (or "isolated") into different buffer pools. This allows for easy monitoring of these tables and indexes. Objects with similar access patterns are also separated and grouped into specific buffer pools.

The following table shows how objects are assigned to buffer pools in the "isolated buffer pool" configuration. In addition to the buffer pools listed in the table, BP1 is also used. It contains 4K sort work files, however, there is no sorting in the Day Posting or Night Balancing workloads.

Buffer Pool	Objects	Buffer Pool	Objects
BP0	4K DB2 catalog tables and indexes	BP32	BCA_RCN_SUMS_IN
BP2	4K tables	BP33	Indexes for tables in BP32
BP3	4K indexes	BP34	BCA_RCN_SUMS_OUT
BP4	BCA_GL_PAYMITEM	BP35	Indexes for tables in BP34
BP5	Indexes for tables in BP4	BP37	BCA_CNSP_ACCT indexes
BP6	BCA92, BKK92_POSTINGS, BKK92_SUMS	BP40	LOBs
BP7	Indexes for tables in BP6	BP42	BUT000, BUT020, BUT021, BUT021RF
BP8	BCA_ACCTBAL, BCA92_RESTART	BP43	Indexes for tables in BP42
BP9	Indexes for tables in BP7	BP44	BANK_PP_APPLLOCK, ADR6, ADRC, ADRCT, ADRP, BANK_RECMN_RECR, BCA_CN_RELATION, BCA_INFOITEM, BCA_PRENOTE
BP10	BCA_COUNTER, BCA_CONTRACT	BP45	Indexes for tables in BP44
BP11	Indexes for tables in BP11	BP8K0	8K DB2 catalog tables BCA_PAYMITEM BCA_PAYMITEM_ENQ BCA_PO_IT
BP20	BCA_CN_EVENT	BP8K1	8K tables BCA_PAYMITEM~S01 index
BP21	Indexes for tables in BP20	BP8K2	BCA_PAYMITEM~S02 index BCA_PAYMITEM~S03 index BCA_PAYMITEM~0 index BCA_PAYMITEM~S04 index
BP22	BCA_BANO_DUE, BCA96	BP16K0	16K DB2 catalog tables
BP23	Indexes for tables in BP22	BP16K1	16K tables
BP24	BCA_CN_PER_ACBAL	BP32K	32K DB2 catalog tables DFKKCODCLUST
BP25	Indexes for tables in BP24	BP32K1	32K tables BCA_CN_LINK~0 index
BP26	BALDAT	BP32K2	BCA_CN_LINK
BP27	Indexes for tables in BP26	BP32K3	BCA_CN_LINK~S02 index BCA_CN_LINK~S01 index
BP28	BCA_TRANSFIG	BP32K4	DFKKCOH DFKKCOHARC DFKKCOHI
BP29	Indexes for tables in BP28	BP32K5	DFKKCODCLUST indexes DFKKCOH indexes DFKKCOHARC indexes DFKKCOH indexes

Table 2: Buffer Pool Assignments - Isolated Buffer Pools

The **"consolidated buffer pool"** configuration uses a much smaller number of buffer pools. Objects have been "consolidated" into 13 buffer pools. This significantly simplifies buffer pool management.

The following table shows how objects are assigned to buffer pools in the "consolidated buffer pool" configuration. In addition to the buffer pools listed in the table, BP1 is also used. It contains 4K sort work files, however, there is no sorting in the Day Posting or Night Balancing workloads.

Buffer Pool	**Objects**
BP0	4K DB2 catalog tables / indexes
BP2	4K tables
BP3	4K indexes
BP40	LOBs
BP8K0	8K DB2 catalog tables
BP8K1	8K tables, including BCA_PAYMITEM
BP8K2	8K indexes
BP16K0	16K DB2 catalog tables
BP16K1	16K tables
BP32K	32K DB2 catalog tables
BP32K1	32K tables
BP32K2	32K indexes

Table 3: Buffer Pool Assignments - Consolidated Buffer Pools

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 zEC12 or z13, which was used as the SAP Database Server.

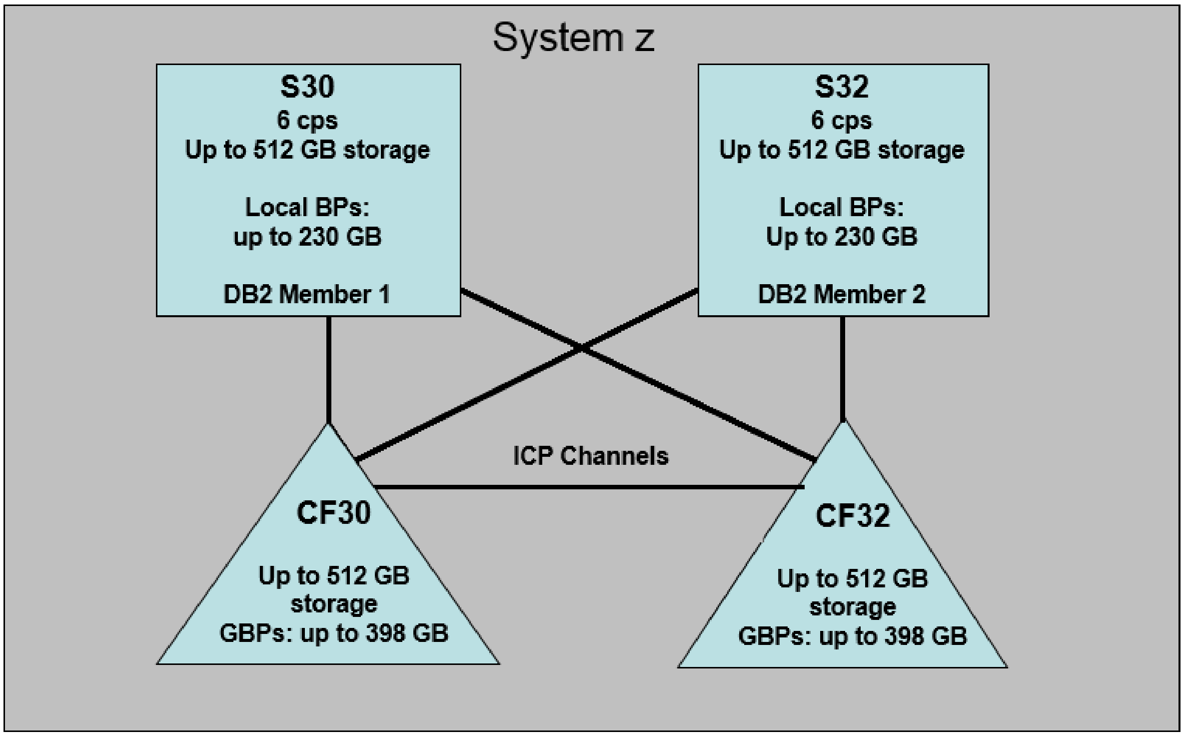

Figure 2: System z LPAR Configuration for DB2 2-way Data Sharing

The SAP Database Server had two z/OS LPARs, named S30 and S32. Each of these LPARs were configured with 6 dedicated CPs and up to 512 GB of real storage. Member 1 of the DB2 data sharing group was on S30 and member 2 of the data sharing group was on S32. Each DB2 member had local buffer pools defined with up to 230 GB of storage.

There also were two LPARs for Coupling Facilities, named CF30 and CF32. Each of these LPARs were configured with 2 dedicated ICF processors and up to 512 GB of real storage. Group buffer pools were defined with up to 398 GB of storage. Group buffer pool duplexing was disabled for most of the measurements. However, having two CFs configured allows for easy enabling of GBP duplexing.

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. Keep in mind that this is a test environment so virtual links are used for simplicity.

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

4.5 Database Configuration

The SBS 7.0 database had the following attributes / customizations :

- Unicode
- Compression for tables, not indexes
- 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

5.0 Single System Results and Analysis – Day Posting

We executed two different sets of buffer pool scaling measurements in a single system environment with the Day Posting workload. Both sets were identical except for the buffer pool configuration used. One set used the “isolated buffer pool” configuration and the other used the “consolidated buffer pool” configuration. See section *4.3 DB2 Configuration* for an explanation of these configurations.

These single system measurements were run on an IBM zEC12 with 12 general purpose processors. The amount of storage that was configured online varied from 256 GB to 1 TB depending on the amount of storage used for the DB2 buffer pools in the measurement. See section *4.0 Test Environment* for more details of the test environment.

5.1 Buffer Pool Scaling with Isolated Buffer Pools

In this series of measurements, we systematically increased the total amount of memory used for the DB2 buffer pools by pretty much doubling it for each measurement point. We started with about 160 GB of memory for the buffer pools in the first measurement. We then increased it to about 320 GB for the second measurement, and then to about 638 GB for the third measurement.

In these measurements, we used the “isolated buffer pool” configuration. It had 43 buffer pools defined. See section *4.3 DB2 Configuration* for an explanation of this configuration and see Table 2 for how objects were assigned to each of these buffer pools.

As we increased the total amount of memory used for the buffer pools, we had to decide how to allocate the additional memory to the individual pools. There were 33 buffer pools used when running the Day Posting workload. We looked at the BP hit ratio and the amount of synchronous reads of these individual pools. While there still was I/O that could be eliminated for a given pool, we increased the size of that buffer pool.

See Table 5 for the sizes of the individual buffer pools used in these measurements.

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

Run id	S30801B1	S30802B1	S30906B2
Real Storage Configured per z/OS LPAR	256 GB	512 GB	1024 GB
Buffer Pool Storage per DB2 Member	**161 GB**	**320 GB**	**638 GB**
Database Server	zEC12	zEC12	zEC12
Number of z/OS LPARs	1	1	1
Number of CPs per z/OS LPAR	12	12	12
z/OS Level	1.13	1.13	1.13
LFAREA (for 1M frames)	180 GB	380 GB	675 GB
Number of Users	10560	10560	10560
ETR (DS/sec)	708.77	819.35	975.71
Average %CPU on z/OS	71.46%	72.92%	78.85%
ITR (DS/sec)	991.84	1123.63	1237.43
Database (DB) Request Time in secs	0.695	0.428	0.209
Total DASD Rate (I/O per sec)	70,450	45,407	27,644
Total DB2 Synchronous Reads per sec	38.4K	11.7K	939.34
DB2 Synchronous Reads per sec – TOT4K	29.8K	5456.72	801.88
DB2 Synchronous Reads per sec – TOT8K	2284.04	1724.34	135.57
DB2 Synchronous Reads per sec – TOT32K	6256.33	4478.41	1.89
Total Buffer Pool Hit Ratio	90.30%	95.47%	99.84%
Buffer Pool Hit Ratio – TOT4K	91.84%	98.40%	99.80%
Buffer Pool Hit Ratio – TOT8K	85.96%	87.62%	99.87%
Buffer Pool Hit Ratio – TOT32K	86.42%	91.52%	100.00%

Table 4: Measurement Results - Buffer Pool Scaling with Isolated BPs

The following table shows how the storage was allocated among the buffer pools for each of these measurements. See Table 2 for how objects were assigned to each of these buffer pools.

Run id	S30801B1		S30802B1		S30906B2	
	Pages (K)	GB	Pages (K)	GB	Pages (K)	GB
BP0	10	0.038	10	0.038	10	0.038
BP2	1,500	5.722	3,000	11.444	5,000	19.073
BP3	1,500	5.722	3,000	11.444	4,000	15.259
BP4	120	0.458	150	0.572	500	1.907
BP5	6,000	22.888	12,000	45.776	12,000	45.776
BP7	2.5	0.010	2.5	0.010	2.5	0.010
BP8	400	1.526	800	3.052	1,000	3.815
BP9	400	1.526	800	3.052	800	3.052
BP10	4,500	17.166	9,000	34.332	9,000	34.332
BP11	4,000	15.259	10,000	38.147	10,000	38.147
BP20	1,500	5.722	3,000	11.444	3,000	11.444
BP21	768	2.930	1,000	3.815	1,000	3.815
BP24	150	0.572	300	1.144	700	2.670
BP25	150	0.572	300	1.144	700	2.670
BP28	768	2.930	1,800	6.866	10,000	38.147
BP29	4,000	15.259	9,000	34.332	9,000	34.332
BP32	2	0.008	2	0.008	4	0.015
BP33	1	0.004	1	0.004	1	0.004
BP34	1	0.004	1	0.004	1	0.004
BP35	1	0.004	1	0.004	1	0.004
BP37	75	0.286	100	0.381	300	1.144
BP42	200	0.763	300	1.144	800	3.052
BP43	180	0.687	210	0.801	500	1.907
BP44	200	0.763	400	1.526	400	1.526
BP45	150	0.572	300	1.144	300	1.144
BP8K0	1,536	11.719	2,036	15.533	6,000	45.776
BP8K1	1,536	11.719	2,036	15.533	6,000	45.776
BP8K2	1,000	7.629	2,000	15.259	14,000	106.811
BP32K	30	0.916	50	1.526	80	2.441
BP32K2	600	18.311	1,300	39.673	4,700	143.433
BP32K3	220	6.714	540	16.479	540	16.479
BP32K4	30	0.916	40	1.221	60	1.831
BP32K5	64	1.953	100	3.052	400	12.207
Total Size		161.46		320.10		638.24

Table 5: Buffer Pool Sizes for Buffer Pool Scaling with Isolated BPs

5.2 Buffer Pool Scaling with Consolidated Buffer Pools

This series of measurements is the same as those documented in section *5.1 Buffer Pool Scaling with Isolated Buffer Pools*, except we used the "consolidated buffer pool" configuration. We consolidated the number of buffer pools used when running the Day Posting workload from 33 to 9. See section *4.3 DB2 Configuration* for an explanation of this configuration and see Table 3 for how objects were assigned to each of these buffer pools.

Using fewer buffer pools simplified the management and tuning of the buffer pools tremendously. It was much easier to look at the BP hit ratio and the amount of synchronous reads of the individual pools to decide where to allocate additional storage.

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

Run id	S31114B1	S31115B1	S31118B1
Real Storage Configured per z/OS LPAR	256 GB	512 GB	1024 GB
Buffer Pool Storage per DB2 Member	**161 GB**	**320 GB**	**638 GB**
Database Server	zEC12	zEC12	zEC12
Number of z/OS LPARs	1	1	1
Number of CPs per z/OS LPAR	12	12	12
z/OS Level	1.13	1.13	1.13
LFAREA (for 1M frames)	180 GB	380 GB	675 GB
Number of Users	10560	10560	10560
ETR (DS/sec)	715.63	818.91	940.27
Average %CPU on z/OS	72.19%	73.04%	78.67%
ITR (DS/sec)	991.31	1121.18	1195.20
Database (DB) Request Time in secs	0.707	0.420	0.253
Total DASD Rate (I/O per sec)	66,591	39,725	26,544
Total DB2 Synchronous Reads per sec	34.0K	6619.57	1119.53
DB2 Synchronous Reads per sec – TOT4K	25.7K	695.41	779.90
DB2 Synchronous Reads per sec – TOT8K	2636.25	1930.35	134.58
DB2 Synchronous Reads per sec – TOT32K	5654.49	3993.81	205.05
Total Buffer Pool Hit Ratio	89.52%	96.27%	99.80%
Buffer Pool Hit Ratio – TOT4K	91.30%	99.80%	99.80%
Buffer Pool Hit Ratio – TOT8K	85.18%	87.18%	99.88%
Buffer Pool Hit Ratio – TOT32K	86.30%	91.16%	99.61%

Table 6: Measurement Results - Buffer Pool Scaling with Consolidated BPs

The following table shows how the storage was allocated among the buffer pools for each of these measurements. See Table 3 for how objects were assigned to each of these buffer pools.

Run id	S31114B1		S31115B1		S31118B1	
	Pages (K)	GB	Pages (K)	GB	Pages (K)	GB
BP0	10	0.038	10	0.038	10	0.038
BP2	9,341	35.633	18,755	71.545	30,405	115.986
BP3	17,179	65.533	36,716	140.060	38,606	147.270
BP8K0	100	0.763	10	0.076	10	0.076
BP8K1	2,517	19.203	2,972	22.675	8,900	67.902
BP8K2	1,545	11.787	3,090	23.575	17,090	130.386
BP32K	30	0.916	50	1.526	80	2.441
BP32K1	630	19.226	1,340	40.894	4,760	145.264
BP32K2	287	8.759	643	19.623	943	28.778
Total Size		161.94		320.09		638.29

Table 7: Buffer Pool Sizes for Buffer Pool Scaling with Consolidated BPs

5.3 Analysis – Single System

Our main objective for these buffer pool scaling measurements was to show the effects of adding memory to the DB2 buffer pools on the overall system performance. It was not our objective to perfectly optimize each measurement. However, as we scaled the overall buffer pool memory, we did add memory to the individual pools that could benefit from the additional memory. For the most part, we looked at the BP hit ratio and the amount of synchronous reads of individual pools to decide whether or not to increase the size of a given pool.

We found that as the total buffer pool memory was scaled from 161 GB to 638 GB using the isolated buffer pool configuration, there was a 97% reduction in DB2 synchronous reads and the BP Hit ratio improved from 90% to close to 100%. There was a 25% improvement in ITR, a 38% improvement in ETR, and a 70% drop in Database (DB) Request time.

The DB Request Time is the time it takes the application server to get a request back from the database server. It consists of the network time and the database processing time, including database I/O and logging I/O. It does not include the application server processing time.

The following table summarizes the effects on key performance indicators as memory is added to DB2 buffer pools using an “isolated buffer pool” configuration while running the SAP Day Posting workload.

Run id	BP Size	ITR DS/sec	ITR Delta	%CPU	ETR DS/sec	ETR Delta	DB Req Time (secs)	DB Req Time Delta	Sync Reads per sec	Sync Reads Delta
S30801B1	**161 GB**	991.84	n/a	71.46%	708.77	n/a	0.695	n/a	38.4K	n/a
S30802B1	**320 GB**	1123.63	13%	72.92%	819.35	16%	0.428	-38%	11.7K	-70%
S30906B2	**638 GB**	1237.43	25%	78.85%	975.71	38%	0.209	-70%	0.9K	-97%

Table 8: Performance Improvements with Isolated Buffer Pools

Another objective of these measurements was to show the effects of reducing the number of DB2 buffer pools used on the overall system performance. Although isolating or separating objects into their own buffer pools can provide essential monitoring capabilities and performance optimizations, it can also produce a large number of buffer pools over time which can be labor intensive and time consuming to maintain and tune.

By consolidating the number of buffer pools from 33 down to 9, we found that there was essentially no impact on ITR performance with a total buffer pool size of up to 320 GB. This is great news for clients who want to simplify their buffer pool strategy. We saw that there was a small cost in ITR performance (less than 3%) with a total buffer pool size of 638 GB.

The following table summarizes the effects on key performance indicators as memory is added to DB2 buffer pools using a “consolidated buffer pool” configuration while running the SAP Day Posting workload. These results can be compared to those in Table 8 to see the differences in using a consolidated versus isolated buffer pool strategy.

Run id	BP Size	ITR DS/sec	ITR Delta	%CPU	ETR DS/sec	ETR Delta	DB Req Time (secs)	DB Req Time Delta	Sync Reads per sec	Sync Reads Delta
S31114B1	**161 GB**	991.31	n/a	72.19%	715.63	n/a	0.707	n/a	34.0K	n/a
S31115B1	**320 GB**	1121.18	13%	73.04%	818.91	14%	0.420	-41%	6.6K	-81%
S31118B1	**638 GB**	1195.20	21%	78.67%	940.27	31%	0.253	-64%	1.1K	-97%

Table 9: Performance Improvements with Consolidated Buffer Pools

We found that DB2 11 for z/OS manages large buffer pools very efficiently, up to 140 GB in a single buffer pool. When we ran with a total buffer pool size of 320 GB using our consolidated buffer pool configuration, our largest single 4K buffer pool was 140 GB, our largest single 8K buffer pool was 24 GB, and our largest single 32K buffer pool was 41 GB.

When we ran with a total buffer pool size of 638 GB using our consolidated buffer pool configuration, our largest single 4K, 8K, and 32K buffer pools were 147 GB, 130 GB, and 145 GB, respectively. Using our isolated buffer pool configuration, our largest single 4K, 8K, and 32K buffer pools were 46 GB, 61 GB, and 143 GB, respectively.

The following graph shows the dramatic reduction in total DB2 synchronous reads as the total size of the buffer pools is increased when running the SAP Day Posting workload with both the isolated buffer pool and consolidated buffer pool configurations. Even though our buffer pool hit ratio with a total buffer pool size of 160 GB was close to 90%, there were still about 35K or so synchronous reads per second. By doubling the amount of memory for the buffer pools, this was cut to 11K or less per second and we achieved a buffer pool hit ratio of about 95%. As we further increased the size of the buffer pools, the synchronous reads continued to drop, but at a slower rate as we approached close to a 100% buffer pool hit ratio.

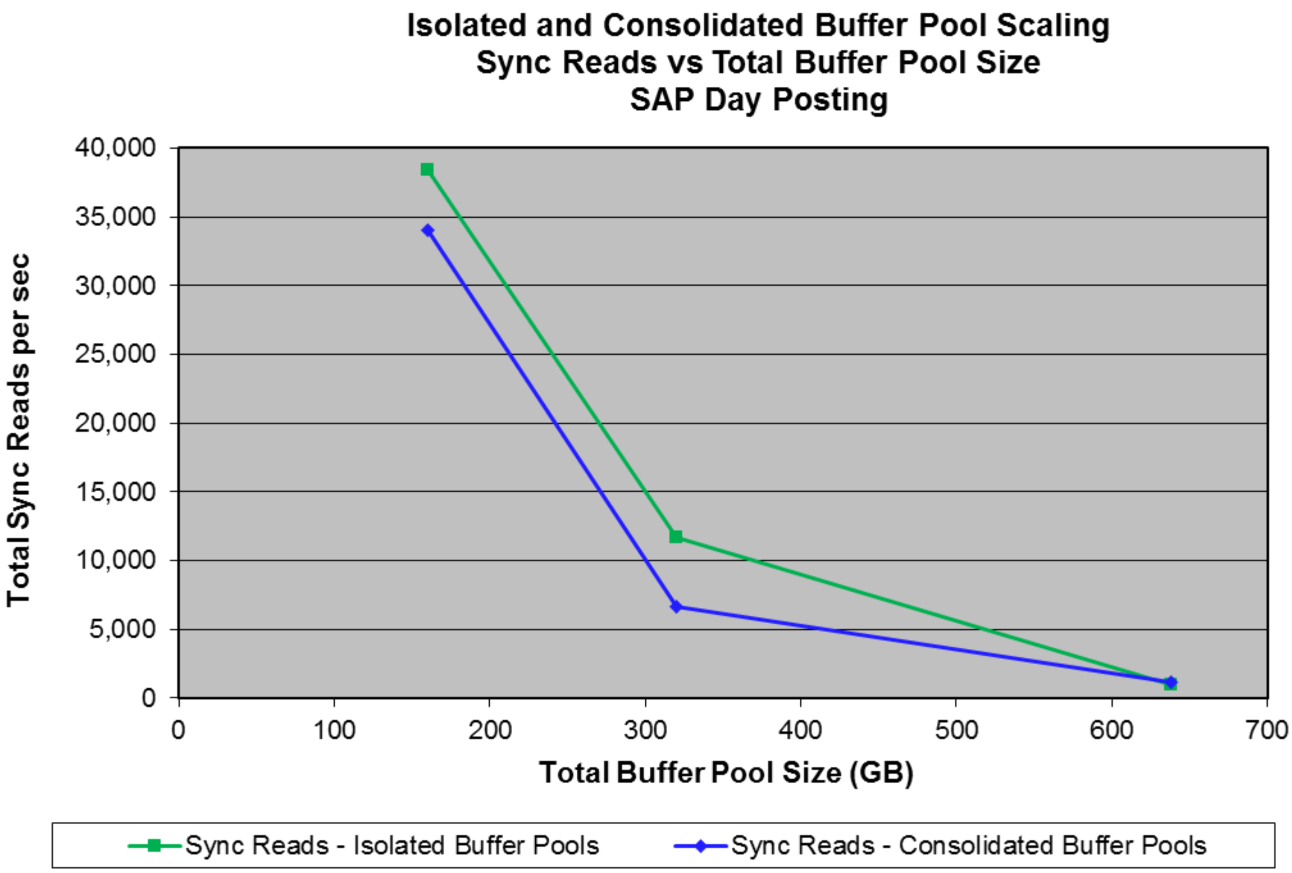

Figure 3: Buffer Pool Scaling Effects on Synchronous Reads

The following graph shows the significant reduction in DB request time as the total size of the buffer pools is increased when running the SAP Day Posting workload with both the isolated buffer pool and consolidated buffer pool configurations. With larger buffer pools, DB2 finds the data it needs in the buffer pools more often instead of having to read it from disk. Less I/O activity and fewer context switches due to the reduction in I/O activity improves the DB request time.

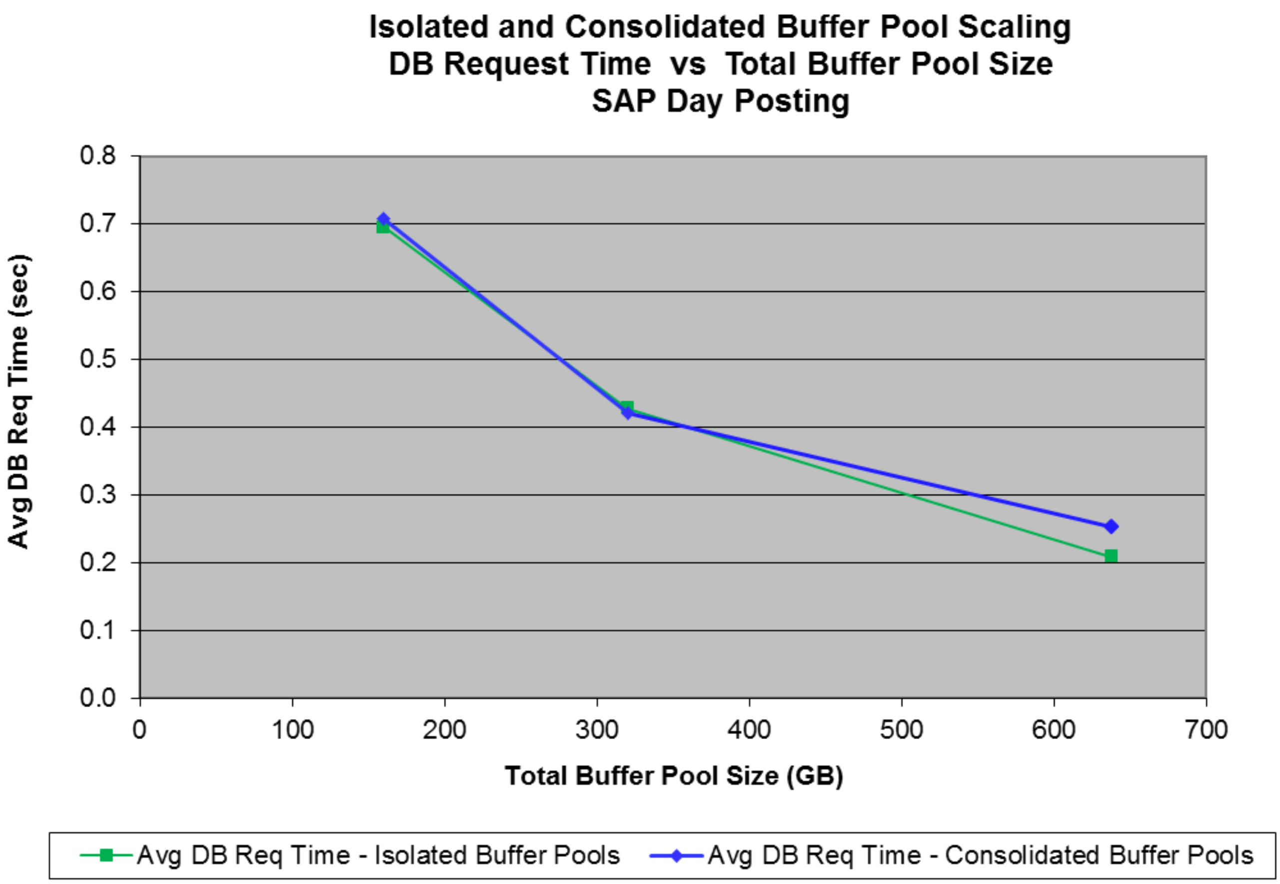

Figure 4: Buffer Pool Scaling Effects on DB Request Time

The following graph shows the significant increase in ETR as the total size of the buffer pools is increased when running the SAP Day Posting workload with both the isolated buffer pool and consolidated buffer pool configurations.

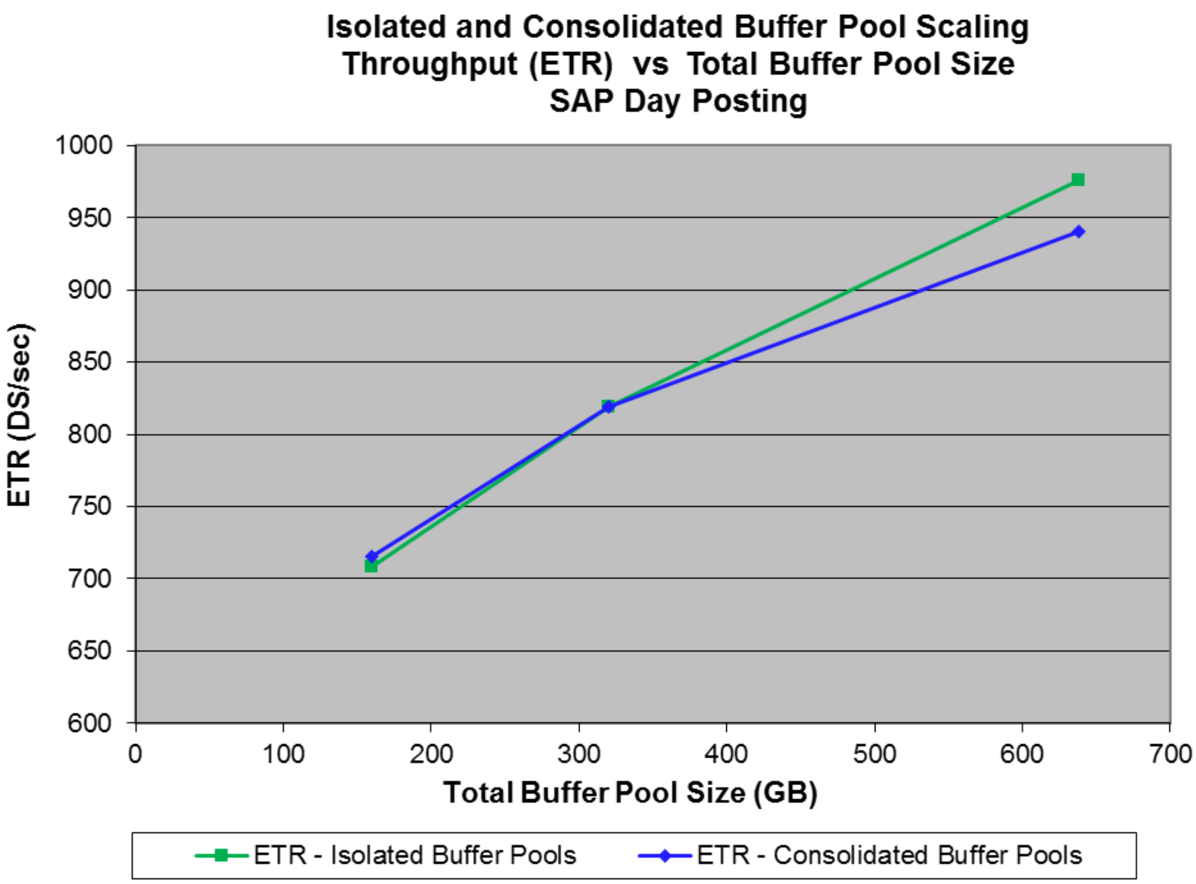

Figure 5: Buffer Pool Scaling Effects on ETR

The following graph shows the significant increase in ITR as the total size of the buffer pools is increased when running the SAP Day Posting workload with both the isolated buffer pool and consolidated buffer pool configurations.

There are several contributors to this improvement in ITR, including the reduction in the synchronous database I/O and the reduction in DB request time. The reduced sync I/O decreases DB2 lock/latch contention and context switches which improve the ITR.

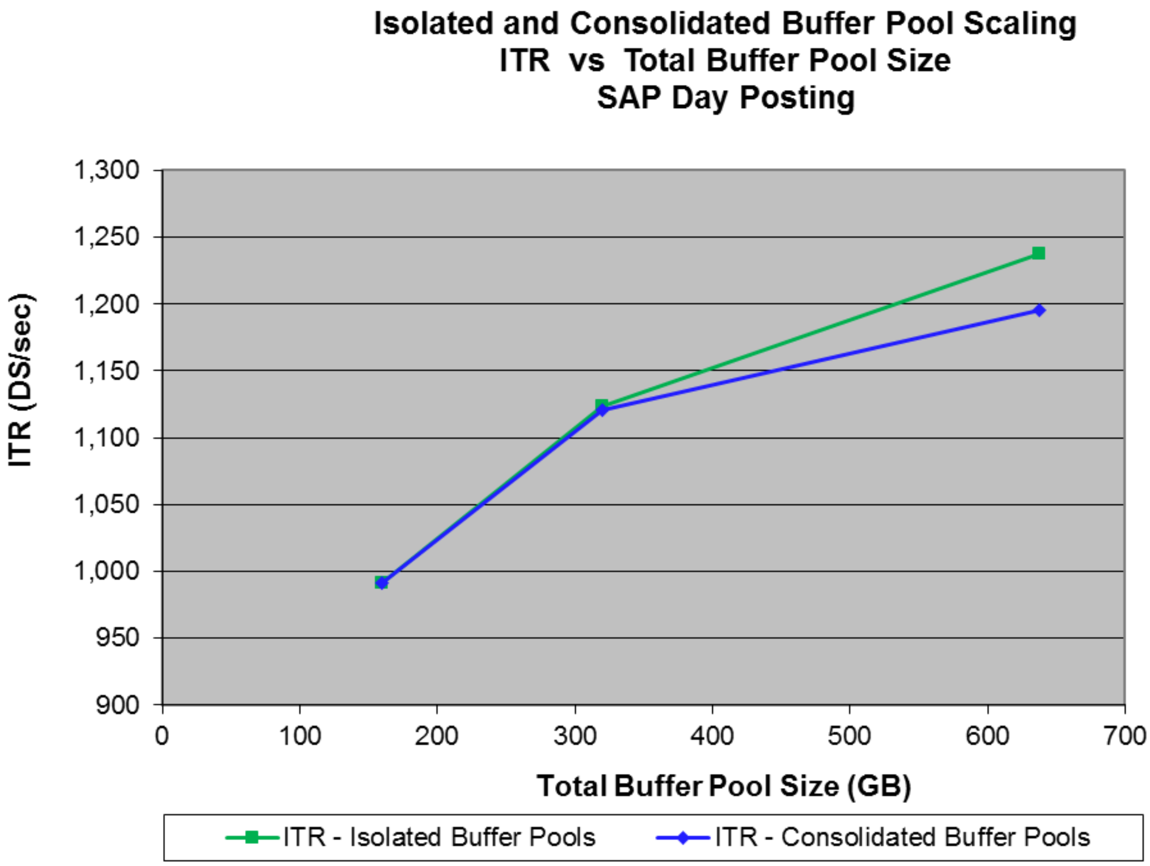

Figure 6: Buffer Pool Scaling Effects on ITR

The performance benefits that we saw with the SAP Day Posting workload from adding memory to the DB2 buffer pools in a single system environment may appear to be extreme. However, keep in mind, that in our 3-tier environment the database serving portion of the Day Posting workload running on the zEC12 is a “pure DB2 workload”. The SAP application server is running outboard. Also, the application data access pattern in this workload is highly random.

However, our measurements do illustrate that performance improvements can be achieved depending on the configuration and workload by adding memory to DB2 buffer pools.

If your buffer pools are already well-tuned with minimal synchronous I/O and good buffer pool hit ratios then adding memory to your buffer pools may not provide any significant performance benefits. The data access patterns of a particular workload will also influence the results. Eliminating random I/O, which results from random data access, is more likely to reduce DB2 synchronous reads than sequential I/O.

6.0 Data Sharing Results and Analysis – Day Posting

Adding more memory to DB2 buffer pools is a little more complex in a data sharing environment since there are local buffer pools in each DB2 member and group buffer pools in the Coupling Facility. In order to understand the effects of adding memory to DB2 buffer pools on overall system performance in a data sharing environment, we executed a number of measurements. First, we added memory to just the local buffer pools. Next, we added memory to just the group buffer pools. Then, we added memory to both the local buffer pools and the group buffer pools. We learned a lot by using this staged approach.

For these measurements, we ran 2-way data sharing on an IBM zEC12 with two z/OS LPARs, each with 6 general purpose processors and up to 512 GB of storage. We used the "consolidated buffer pool" configuration. See section *4.0 Test Environment* for more details of the test environment.

6.1 Local Buffer Pool Scaling

In these measurements, we systematically increased the total amount of memory used for the local buffer pools. We started with about 29 GB of memory for the local buffer pools in each DB2 member. We then increased it to 58 GB, 115 GB, and finally to 230 GB per member. We held the total group buffer pool size constant at 64 GB.

Each group buffer pool is divided into two parts, directory entries and data entries. Directory entries track the location and status of all data and index pages that occupy a spot in the local buffer pools in the data sharing group. Data entries are blocks where the data pages are cached. A group buffer pool should be at least large enough to avoid directory entry reclaims. A reclaim is basically a steal of a currently in-use directory entry to be used to register a new page. Directory entry reclaims result in invalidation of clean pages cached in local buffer pools. This may increase database I/O activity.

To avoid directory entry reclaims, there should be enough directory entries to register all the different pages that could be cached in the group buffer pool and in the associated local buffer pools at any one time. One way to increase directory entries is to make the group buffer pool CF cache structure larger. The Coupling Facility Usage Summary report within the CF Activity Report within RMF shows directory reclaims (DIR REC) and cross invalidations due to directory reclaims (DIR REC XI's).

For this series of measurements, we sized our group buffer pools to accommodate our largest local buffer pool configuration of 230 GB. The total amount of memory for the group buffer pools needed to be 64 GB to avoid directory reclaims. Once the group buffer pools were tuned for this highest point, we used this same group buffer pool configuration for the lower points.

Our objective for these local buffer pool scaling measurements was to show the effects of adding memory to the DB2 local buffer pools on the overall system performance. It was not our objective to perfectly optimize each measurement. However, as we scaled the overall local buffer pool memory, we did add memory to the individual local pools that could benefit from the additional memory. Similar to what we did in our single system measurements, we looked at the BP hit ratio and the amount of synchronous reads of individual pools to decide whether or not to increase the size of a given pool.

See Table 11 for the sizes of the individual local and group buffer pools used in these measurements.

The details of these measurements are summarized in the following table. Note that the "Total DASD Rate" and the "Total DB2 Synchronous Reads per sec" are the sum from each LPAR. Keep in mind that twice the amount of memory is needed for group buffer pools for redundancy. In a data sharing environment, there must be storage available in a second CF for the GBP cache structures to be rebuilt in a recovery scenario.

Run id	S40303B2	S40303B1	S40227B1	S40209B1
Real Storage Configured per z/OS LPAR	64 GB	128 GB	256 GB	512 GB
Local Buffer Pool Storage per DB2 Member	**29 GB**	**58 GB**	**115 GB**	**230 GB**
Group Buffer Pool Storage	64 GB	64 GB	64 GB	64 GB
Database Server	zEC12	zEC12	zEC12	zEC12
Number of z/OS LPARs	2	2	2	2
Number of CPs per z/OS LPAR	6	6	6	6
z/OS Level	2.1	2.1	2.1	2.1
LFAREA (for 1M frames)	40 GB	80 GB	160 GB	320 GB
Number of Users	7680	7680	7680	7680
ETR (DS/sec)	581.69	594.30	588.51	642.26
Average %CPU on z/OS	84.37%	84.47%	81.20%	81.80%
ITR (DS/sec)	689.45	703.56	724.81	785.21
Database (DB) Request Time in secs	0.584	0.575	0.542	0.364
Total DASD Rate (I/O per sec)	70,029	60,386	49,698	29,716
Average %CF Utilization	28.9%	28.7%	27.6%	28.2%
Total DB2 Synchronous Reads per sec	51.7K	43.2K	32.8K	11.4K
DB2 Synchronous Reads per sec – TOT4K	39.9K	33.0K	25.0K	4629.83
DB2 Synchronous Reads per sec – TOT8K	3428.04	2641.48	2129.07	1773.82
DB2 Synchronous Reads per sec – TOT32K	8454.81	7579.18	5723.67	5008.93
Total Buffer Pool Hit Ratio	82.28%	85.22%	88.17%	94.80%
Buffer Pool Hit Ratio – TOT4K	83.03%	86.47%	89.55%	98.26%
Buffer Pool Hit Ratio – TOT8K	83.17%	84.44%	85.94%	87.32%
Buffer Pool Hit Ratio – TOT32K	75.96%	78.75%	83.61%	87.09%
Total Group Buffer Pool Hit Ratio	37%	33%	30%	53%
Group Buffer Pool Hit Ratio – TOT4K	30%	26%	26%	59%
Group Buffer Pool Hit Ratio – TOT8K	48%	42%	28%	27%
Group Buffer Pool Hit Ratio – TOT32K	82%	81%	71%	48%

Table 10: Measurement Results for Local Buffer Pool Scaling

The "Total Group Buffer Pool Hit Ratio" indicates the effectiveness of finding a referenced page in the group buffer pool given that it was not found in the local buffer pool. It is calculated from data in the DB2 Performance Expert Statistics Report.

Total Group Buffer Pool Hit Ratio =
(SYN.READ(XI)-DATA RETURNED + SYN.READ(NF)-DATA RETURNED) /
(SYN.READ(XI)-DATA RETURNED + SYN.READ(XI)-NO DATA RETURN +
SYN.READ(NF)-DATA RETURNED + SYN.READ(NF)-NO DATA RETURN)

The following table shows how the total amount of storage for the local buffer pools was allocated among the individual local buffer pools for each of these measurements. It also shows how the storage for the group buffer pools was allocated among the individual group buffer pools. The same group buffer pool configuration was used in all these measurements. See Table 3 for how objects were assigned to each of these buffer pools.

	Local Buffer Pool Sizes per DB2 Member								GBP Sizes
Run id	S40303B2		S40303B1		S40227B1		S40209B1		
	Pages (K)	GB	Pages (K)	GB	Pages (K)	GB	Pages (K)	GB	
BP0	10	0.04	10	0.04	10	0.04	10	0.04	977M
BP2	2,000	7.63	4,000	15.26	8,000	30.52	16,000	61.04	14G
BP3	3,000	11.44	6,000	22.89	12,000	45.78	24,000	91.55	23G
BP8K0	50	0.38	50	0.38	50	0.38	50	0.38	196M
BP8K1	480	3.66	960	7.32	1,920	14.65	3,840	29.30	10G
BP8K2	300	2.29	600	4.58	1,200	9.16	2,400	18.31	13G
BP32K	10	0.31	10	0.31	10	0.31	10	0.31	196M
BP32K1	60	1.83	120	3.66	240	7.32	480	14.65	977M
BP32K2	60	1.83	120	3.66	240	7.32	480	14.65	977M
Total Size		29.41		58.10		115.47		230.22	64G

Table 11: LBP and GBP Sizes for Local Buffer Pool Scaling

We found that as the total local buffer pool memory was scaled from 29 GB to 230 GB, there was a 78% reduction in DB2 synchronous reads and the BP hit ratio improved from 82% to close to 95%. There was a 14% improvement in ITR, a 10% improvement in ETR, and a 38% drop in DB request time. This is the same behavior that we saw in our single system buffer pool scaling measurements. However, keep in mind that in data sharing, memory needs to be added to the local buffer pools in each data sharing member.

The following table shows the effects on key performance indicators as memory is added to just the local buffer pools while running the SAP Day Posting workload.

Run id	LBP Size per member ---- GBP Size	ITR DS/sec	ITR Delta	%CPU	ETR DS/sec	ETR Delta	DB Req Time (secs)	DB Req Time Delta	Sync Reads per sec	Sync Reads Delta
S40303B2	**29 GB** 64 GB	689.45	n/a	84.37%	581.69	n/a	0.584	n/a	51.7K	n/a
S40303B1	**58 GB** 64 GB	703.56	2%	84.47%	594.30	2%	0.575	-2%	43.2K	-16%
S40227B1	**115 GB** 64 GB	724.81	5%	81.20%	588.51	1%	0.542	-7%	32.8K	-37%
S40209B1	**230 GB** 64 GB	785.21	14%	81.80%	642.26	10%	0.364	-38%	11.4K	-78%

Table 12: Performance Improvements with Local Buffer Pool Scaling

The following graph shows the dramatic reduction in total DB2 synchronous reads as the total size of the local buffer pools is increased when running the SAP Day Posting workload. With a total local buffer pool size of 29 GB, the buffer pool hit ratio was about 82% with 51.7K synchronous reads per second. By doubling the amount of memory for the local buffer pools, this was cut to 43.2K per second and we achieved a buffer pool hit ratio of 85%. As we further increased the size of the local buffer pools to 230 GB, the synchronous reads continued to steadily drop. The buffer pool hit ratio in our final measurement was close to 95%.

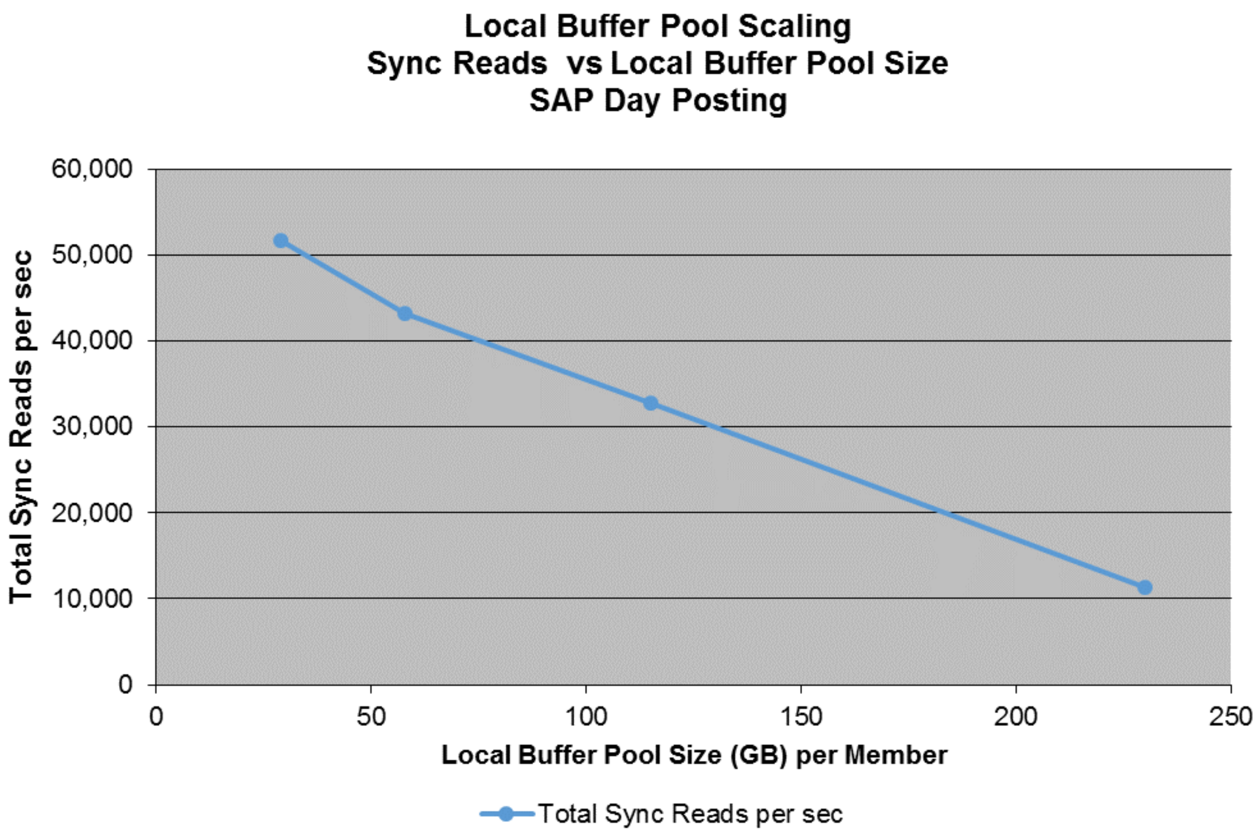

Figure 7: Local Buffer Pool Scaling Effects on Synchronous Reads

The following graph shows the reduction in DB request time as the total size of the local buffer pools is increased when running the SAP Day Posting workload. With larger local buffer pools, DB2 finds the data it needs in the local buffer pools more often instead of having to read it from the group buffer pool or disk. Less I/O activity and fewer context switches due to the reduction in I/O activity improves the DB request time.

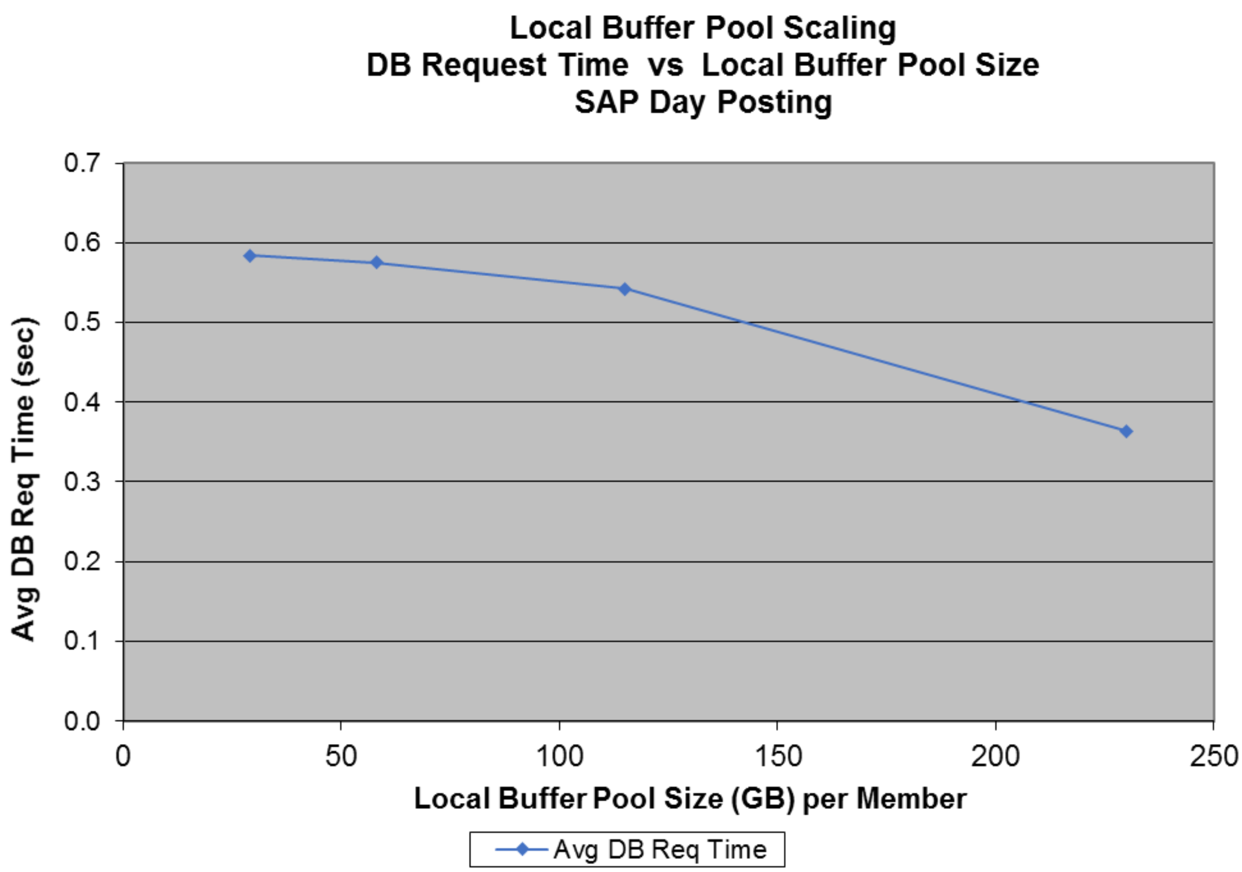

Figure 8: Local Buffer Pool Scaling Effects on DB Request Time

The following graph shows the significant increase in ITR as the total size of the local buffer pools is increased when running the SAP Day Posting workload. There are several contributors to this improvement in ITR, including the reduction in the synchronous database I/O and the reduction in DB request time. The reduced sync I/O decreases DB2 lock/latch contention and context switches which improve the ITR.

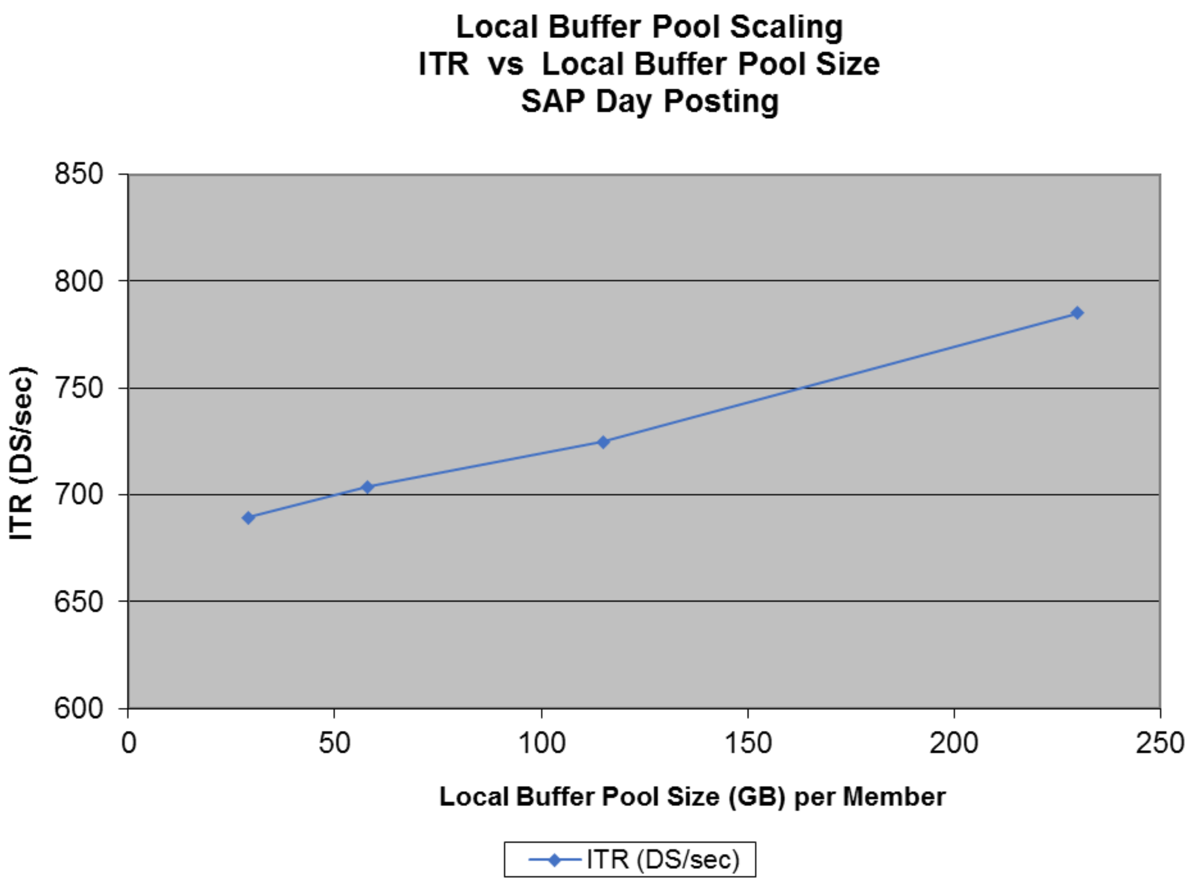

Figure 9: Local Buffer Pool Scaling Effects on ITR

See section *6.4 Analysis – Data Sharing* for the overall analysis and findings of the study of adding memory to DB2 buffer pools to improve system performance in a data sharing environment.

6.2 Group Buffer Pool Scaling

In these measurements, we systematically increased the total amount of memory used for the group buffer pools. We started with 52 GB of memory for the group buffer pools. We then increased it to 103 GB, to 204 GB, and then to 398 GB. We held the total local buffer pool size constant at 58 GB.

In our base measurement in this series with a total local buffer pool size of 58 GB and a group buffer pool size of 52 GB, the group buffer pool was sized or tuned appropriately to not have any directory entry reclaims. What we are looking at in these measurements is to see if adding more memory beyond what is needed to avoid directory entry reclaims provides a performance benefit.

The details of these measurements are summarized in the following table. Note that the "Total DASD Rate" and the "Total DB2 Synchronous Reads per sec" are the sum from each LPAR. Keep in mind that twice the amount of memory is needed as group buffer pools are increased in size for redundancy. In a data sharing environment, there must be storage available in a second CF for the GBP cache structures to be rebuilt in a recovery scenario.

Run id	S40201B1	S40218B1	S40211B1	S40225B1
Real Storage Configured per z/OS LPAR	128 GB	128 GB	128 GB	128 GB
Local Buffer Pool Storage per DB2 Member	58 GB	58 GB	58 GB	58 GB
Group Buffer Pool Storage	**52 GB**	**103 GB**	**204 GB**	**398 GB**
Database Server	zEC12	zEC12	zEC12	zEC12
Number of z/OS LPARs	2	2	2	2
Number of CPs per z/OS LPAR	6	6	6	6
z/OS Level	2.1	2.1	2.1	2.1
LFAREA (for 1M frames)	80 GB	80 GB	80 GB	80 GB
Number of Users	7680	7680	7680	7680
ETR (DS/sec)	586.01	581.53	590.24	601.46
Average %CPU on z/OS	83.33%	80.94%	79.34%	79.86%
ITR (DS/sec)	703.28	718.51	743.98	753.05
Database (DB) Request Time in secs	0.577	0.526	0.455	0.411
Total DASD Rate (I/O per sec)	61,398	57,496	50,592	42,450
Average %CF Utilization	28.2%	27.6%	27.2%	28.0%
Total DB2 Synchronous Reads per sec	44.0K	40.3K	34.2K	26.5K
DB2 Synchronous Reads per sec – TOT4K	32.9K	30.7K	24.5K	16.6K
DB2 Synchronous Reads per sec – TOT8K	3593.53	2487.31	2505.34	2480.62
DB2 Synchronous Reads per sec – TOT32K	7506.47	7086.72	7148.11	7382.96
Total Buffer Pool Hit Ratio	84.78%	85.79%	87.69%	90.48%
Buffer Pool Hit Ratio – TOT4K	86.29%	87.07%	89.62%	93.21%
Buffer Pool Hit Ratio – TOT8K	82.85%	84.66%	85.38%	86.88%
Buffer Pool Hit Ratio – TOT32K	78.89%	79.89%	79.95%	80.24%
Total Group Buffer Pool Hit Ratio	29%	38%	57%	82%
Group Buffer Pool Hit Ratio – TOT4K	25%	32%	56%	87%
Group Buffer Pool Hit Ratio – TOT8K	20%	44%	45%	47%
Group Buffer Pool Hit Ratio – TOT32K	81%	92%	96%	96%

Table 13: Measurement Results for Group Buffer Pool Scaling

The "Total Group Buffer Pool Hit Ratio" indicates the effectiveness of finding a referenced page in the group buffer pool given that it was not found in the local buffer pool. It is calculated from data in the DB2 Performance Expert Statistics Report.

Total Group Buffer Pool Hit Ratio =
(SYN.READ(XI)-DATA RETURNED + SYN.READ(NF)-DATA RETURNED) /
(SYN.READ(XI)-DATA RETURNED + SYN.READ(XI)-NO DATA RETURN +
SYN.READ(NF)-DATA RETURNED + SYN.READ(NF)-NO DATA RETURN)

The following table shows how the total amount of storage for the group buffer pools was allocated among the individual group buffer pools for each of these measurements. It also shows how the storage for the local buffer pools was allocated among the individual local buffer pools. The same local buffer pool configuration was used in all these measurements. See Table 3 for how objects were assigned to each of these buffer pools.

	Group Buffer Pool Sizes				LBP Sizes per DB2 Member	
Run id	S40201B1	S40218B1	S40211B1	S40225B1	Pages (K)	GB
BP0	977M	977M	977M	977M	10	0.04
BP2	12G	25G	50G	99G	4,000	15.26
BP3	19G	38G	76G	153G	6,000	22.89
BP8K0	196M	196M	196M	196M	50	0.38
BP8K1	10G	19G	38G	76G	960	7.32
BP8K2	8G	15G	31G	61G	600	4.58
BP32K	196M	196M	196M	196M	10	0.31
BP32K1	977M	2G	4G	4G	120	3.66
BP32K2	977M	2G	4G	4G	120	3.66
Total Size	52G	103G	204G	398 GB		58.10

Table 14: LBP and GBP Sizes for Group Buffer Pool Scaling

We found that adding more memory beyond what is needed to avoid directory entry reclaims provides a performance benefit. As the total group buffer pool memory was scaled from 52 GB to 398 GB, there was a 40% reduction in DB2 synchronous reads and the BP hit ratio improved from 85% to over 90% and the group buffer pool hit ratio improved from 29% to 82% There was a 7% improvement in ITR, a 3% improvement in ETR, and a 29% drop in DB request time.

The group buffer pools act as a second layer of cache. However, the amount of GBP dependent data in a workload influences the performance benefits of adding memory to the group buffer pools when GBPCACHE CHANGED is used.

GBPCACHE is a DB2 parameter that can be specified on the DB2 object level via DDL or on the group buffer pool level as a parameter on the ALTER GROUPBUFFERPOOL command. When GBPCACHE CHANGED is specified for an object, only updated (or changed) pages are written to the group buffer pool when there is inter-DB2 R/W interest on the object. When there is no inter-DB2 R/W interest, the group buffer pool is not used. Inter-DB2 R/W interest exists when more than one member in the data sharing group has the object open and at least one member has it open for update. Data is GBP dependent if there is inter-DB2 R/W interest in it.

The SAP Day Posting workload in our environment has about 58% GBP dependent data. Since we are running with GBPCACHE CHANGED, only the GBP dependent data has potential for being cached in the group buffer pool. The amount of GBP dependent data in client workloads varies widely, 10% to 100%.

Another factor that influenced the effectiveness of the group buffer pool cache in our measurements with the SAP Day Posting workload is the use of the DB2 member cluster feature. This feature significantly improves insert performance in data sharing environments. When a table space is defined with the member cluster option, DB2 manages space for inserts on a member-by-member basis instead of using one centralized space map. Each DB2 member has its own designated pages to use so there are a very small number of distinct pages that are cached in the group buffer pool.

The DB2 member cluster feature is highly recommended for specific SAP banking tables when running in a data sharing environment.

In our measurements, even though we were using a consolidated buffer configuration where all the 8K tables were assigned to BP8K1, there was basically only one table with activity in this buffer pool. This table was BCA_PAYMITEM which is member clustered and is accessed as heavy INSERT and random SELECT. Because of this, adding memory to this specific group buffer pool did not provide a performance benefit.

The following table shows the effects on key performance indicators as memory is added to the group buffer pools while running the SAP Day Posting workload.

Run id	LBP Size per member ---- GBP Size	ITR DS/sec	ITR Delta	%CPU	ETR DS/sec	ETR Delta	DB Req Time (secs)	DB Req Time Delta	Sync Reads per sec	Sync Reads Delta
S40201B1	58 GB **52 GB**	703.28	n/a	83.33%	586.01	n/a	0.577	n/a	44.0K	n/a
S40218B1	58 GB **103 GB**	718.51	2%	80.94%	581.53	<1%	0.526	-9%	40.3K	8%
S40211B1	58 GB **204 GB**	743.98	6%	79.34%	590.24	<1%	0.455	-21%	34.2K	-22%
S40225B1	58 GB **398 GB**	753.05	7%	79.86%	601.46	3%	0.411	-29%	26.5K	-40%

Table 15: Performance Improvements with Group Buffer Pool Scaling

The following graph shows the reduction in total DB2 synchronous reads as the total size of the group buffer pools is increased when running the SAP Day Posting workload. With a total group buffer pool size of 52 GB, there were 44K synchronous reads per second. By doubling the amount of memory for the group buffer pools, this was cut to 40.3K. As we further increased the size of the group buffer pools to 398 GB, the synchronous reads continued to drop to 26.5K per second.

DB2 synchronous reads are avoided when the data is found in the local buffer pool or, since we are running with GBPCACHE CHANGED, when GBP dependent data is found in the group buffer pool.

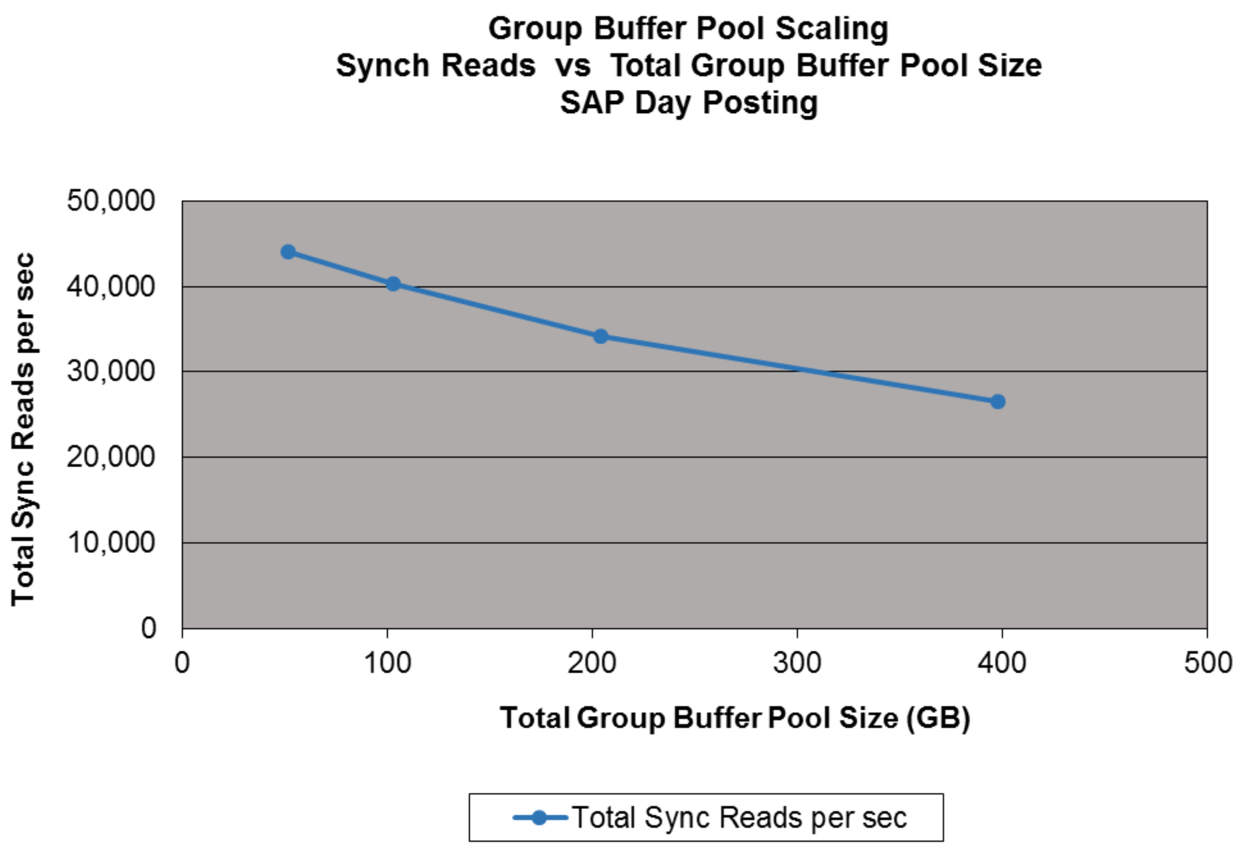

Figure 10: Group Buffer Pool Scaling Effects on Synchronous Reads

The following graph shows the reduction in DB request time as the total size of the group buffer pools is increased when running the SAP Day Posting workload. With larger group buffer pools, DB2 finds some of the GBP dependent data it needs in the group buffer pools more often instead of having to read it from disk, therefore, the DB request time improves.

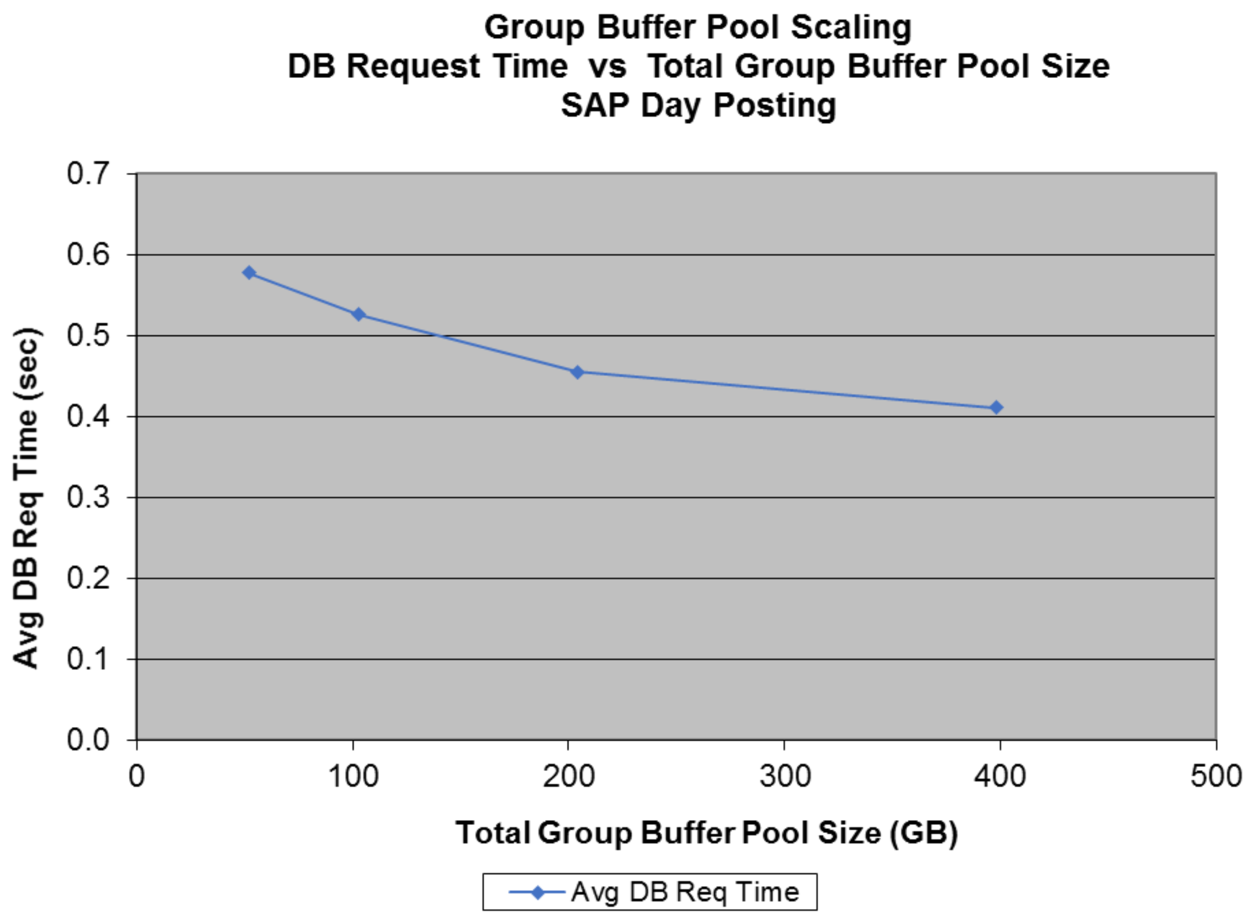

Figure 11: Group Buffer Pool Effects on DB Request Time

The following graph shows the increase in ITR as the total size of the group buffer pools is increased when running the SAP Day Posting workload. There are several contributors to this improvement in ITR, including the reduction in the synchronous database I/O and the reduction in DB request time. The reduced sync I/O decreases DB2 lock/latch contention and context switches which improve the ITR.

There is a point of diminishing return in the ITR improvement as memory is added solely to the group buffer pools. The significant contributor to this is the amount of GBP dependent data in our workload for these measurements. Since we are running with GBPCACHE CHANGED and our workload has only a little more than half of its data as GBP dependent, there is only so much data eligible to be cached in the group buffer pools.

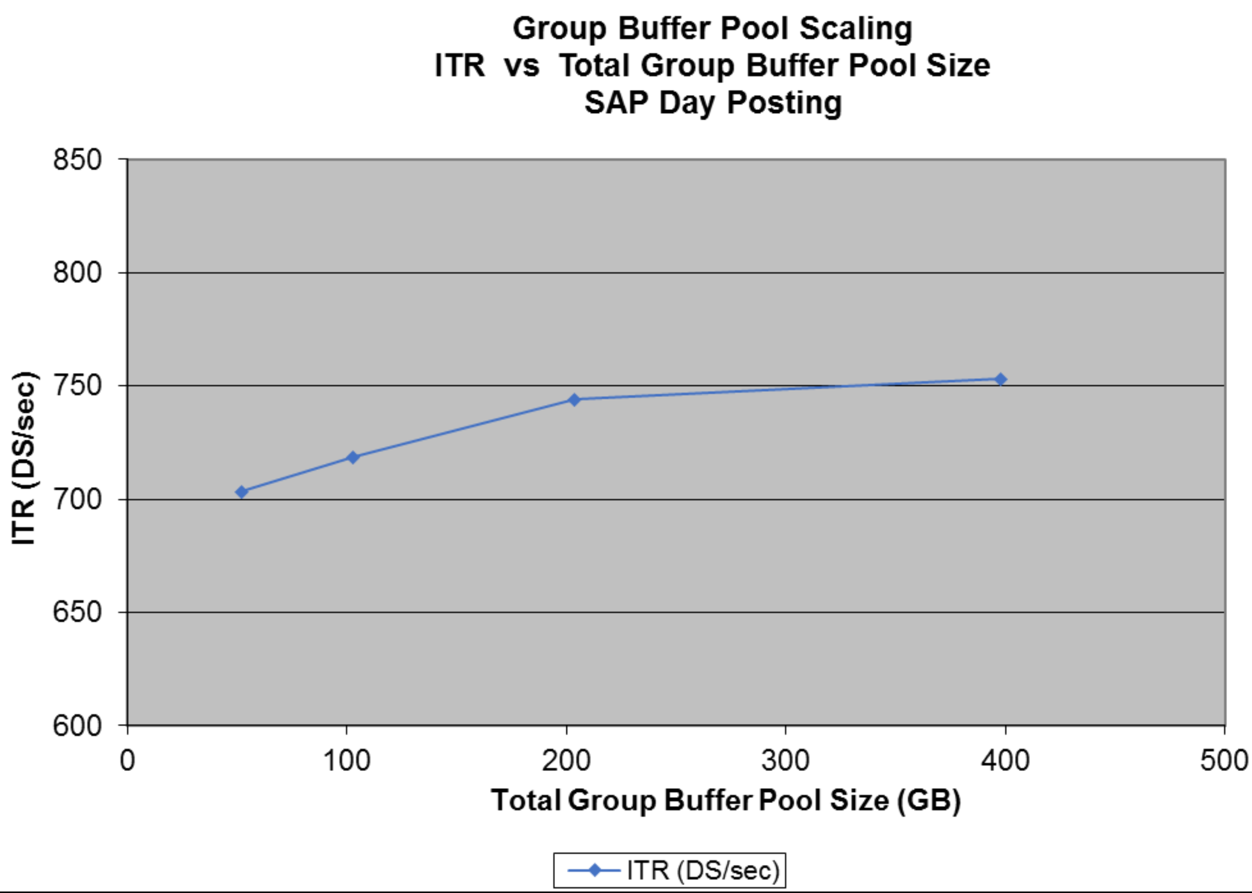

Figure 12: Group Buffer Pool Scaling Effects on ITR

See section *6.4 Analysis – Data Sharing* for the overall analysis and findings of the study of adding memory to DB2 buffer pools to improve system performance in a data sharing environment.

6.3 Large Local Buffer Pools and Large Group Buffer Pools

In this measurement, we added memory to both the local buffer pools and the group buffer pools. We took the best results from the local buffer pool scaling and group buffer pool scaling measurements to decide on the total amount of memory to use for both the local and group buffer pools.

The choice of 230 GB for the local buffer pools was obvious. With this amount of memory for the local pools, we achieved a buffer pool hit ratio of close to 95%. And, comparing this with using 29 GB for the local buffer pools, there was a 78% reduction in DB2 synchronous reads and a 14% improvement in ITR. See Table 12 to see how adding memory to the local buffer pools affected the key performance indicators when running the SAP Day Posting workload.

We chose to use 204 GB of memory for the group buffer pools since this gave us the "most bang for our buck". Comparing this with using 52 GB for the group buffer pools, there was a 22% reduction in DB2 synchronous reads and a 6% improvement in ITR. Adding another 194 GB to the group buffer pools, only bought us another 1% improvement in ITR. See Table 15 to see how adding memory to the group buffer pools affected the key performance indicators when running the SAP Day Posting workload.

The details of this measurement with a local buffer pool size of 230 GB and a group buffer pool size of 204 GB is summarized in the following table.

Run id	S40409B1
Real Storage Configured per z/OS LPAR	512 GB
Local Buffer Pool Storage per DB2 Member	**230 GB**
Group Buffer Pool Storage	**204 GB**
Database Server	zEC12
Number of z/OS LPARs	2
Number of CPs per z/OS LPAR	6
z/OS Level	2.1
LFAREA (for 1M frames)	320 GB
Number of Users	7680
ETR (DS/sec)	638.97
Average %CPU on z/OS	78.30%
ITR (DS/sec)	816.11
Database (DB) Request Time in secs	0.382
Total DASD Rate (I/O per sec)	24,822
Average %CF Utilization	25.7%
Total DB2 Synchronous Reads per sec	7873.57
DB2 Synchronous Reads per sec – TOT4K	2535.02
DB2 Synchronous Reads per sec – TOT8K	1695.13
DB2 Synchronous Reads per sec – TOT32K	3643.41
Total Buffer Pool Hit Ratio	95.75%
Buffer Pool Hit Ratio – TOT4K	99.02%
Buffer Pool Hit Ratio – TOT8K	87.85%
Buffer Pool Hit Ratio – TOT32K	90.06%
Total Group Buffer Pool Hit Ratio	73%
Group Buffer Pool Hit Ratio – TOT4K	77%
Group Buffer Pool Hit Ratio – TOT8K	32%
Group Buffer Pool Hit Ratio – TOT32K	96%

Table 16: Measurement Results for Large LBPs and Large GBPs

Table 11 (column for run id S40209B1) shows how the 230 GB of memory was allocated among the individual local buffer pools for this measurement. Table 14 (column for run id S40211B1) shows how the 204 GB of memory was allocated among the individual group buffer pools for this measurement.

6.4 Analysis – Data Sharing

In section *6.1 Local Buffer Pool Scaling*, the benefits of adding memory to the local buffer pools for the SAP Day Posting workload were discussed. In section *6.2 Group Buffer Pool Scaling*, the benefits of adding memory to the group buffer pools were discussed. In both scenarios, we saw that adding memory to these buffer pools reduced the DB2 synchronous reads and the database request time which, in turn, improved ITR.

In this section, we will look at the effects on overall system performance of adding memory to both the local and the group buffer pools. See Run id S40409B1 in the table below. As expected, this gave us the best performance improvements. However, we found that how additional memory is allocated between the local and group buffer pools to get the most bang for your buck is very much configuration and workload dependent.

The following table shows the effects on key performance indicators as memory is added in varying amounts to the local and/or group buffer pools while running the SAP Day Posting workload. We included in this table our base run with a relatively small amount of memory allocated to both the local buffer pools and the group buffer pools, our "best" run from our group buffer pool scaling measurements, our "best" run from our local buffer pool scaling measurements, and finally our run with a relatively large amount of memory allocated to both the local and group buffer pools.

Run id	LBP Size per member ---- GBP Size	ITR DS/sec	ITR Delta	%CPU	ETR DS/sec	ETR Delta	DB Req Time (secs)	DB Req Time Delta	Sync Reads per sec	Sync Reads Delta
S40303B2	**29 GB** **64 GB**	689.45	n/a	84.37%	581.69	n/a	0.584	n/a	51.7K	n/a
S40211B1	**58 GB** **204 GB**	743.98	8%	79.34%	590.24	1%	0.455	-22%	34.2K	-34%
S40209B1	**230 GB** **64 GB**	785.21	14%	81.80%	642.26	10%	0.364	-38%	11.4K	-78%
S40409B1	**230 GB** **204 GB**	816.11	18%	78.30%	638.97	10%	0.382	-35%	7.9K	-85%

Table 17: Performance Improvements with Local and Group Buffer Pool Scaling

Looking at the measurements in the above table, we see a steady reduction in DB2 synchronous reads and a steady increase in ITR as we added memory to the local and/or group buffer pools.

Using 230 GB of memory for our local buffer pools and 204 GB for our group buffer pools, compared to using 29 GB and 64 GB, respectively, there was an 85% reduction in sync reads and a 35% reduction in DB request time. There was also a reduction in CF traffic. Using large local buffer pools may avoid accessing the CF as long as the pages in the local buffer pools have not been cross-invalidated. The CF utilization decreased from 29% to 26%. All these things contributed to the 18% improvement in ITR.

There are some other interesting comparisons to highlight.

Comparing run id S40409B1 (230 GB / 204 GB) with run id S40209B1 (230 GB / 64 GB) where the only difference was an additional 140 GB of memory allocated to the group buffer pools (280 GB is needed for redundancy), there was a 31% reduction in synchronous reads and a 4% improvement in ITR.

Comparing run id S40409B1 (230 GB / 204 GB) with run id S40211B1 (58 GB / 204 GB) where the only difference was an additional 172 GB of memory allocated to the local buffer pools per member for a total of 344 GB since we are running 2-way data sharing, there was a 77% reduction in synchronous reads and a 10% improvement in ITR.

From these measurements and from our other local buffer pool and group buffer pool scaling measurements, we know that adding memory to the local buffer pools is more beneficial than adding it to the group buffer pools for our configuration and workload. Our SAP Day Posting workload with key banking tables member clustered has a large amount (about 42%) of non-GBP dependent data. This makes the scaling of the local buffer pools very effective. We also ran in a 2-way data sharing environment making it very reasonable to scale our local buffer pools.

The following chart plots the ITRs for the measurements in Table 17 and it shows the percent improvement in ITR with varying local and group buffer pool sizes.

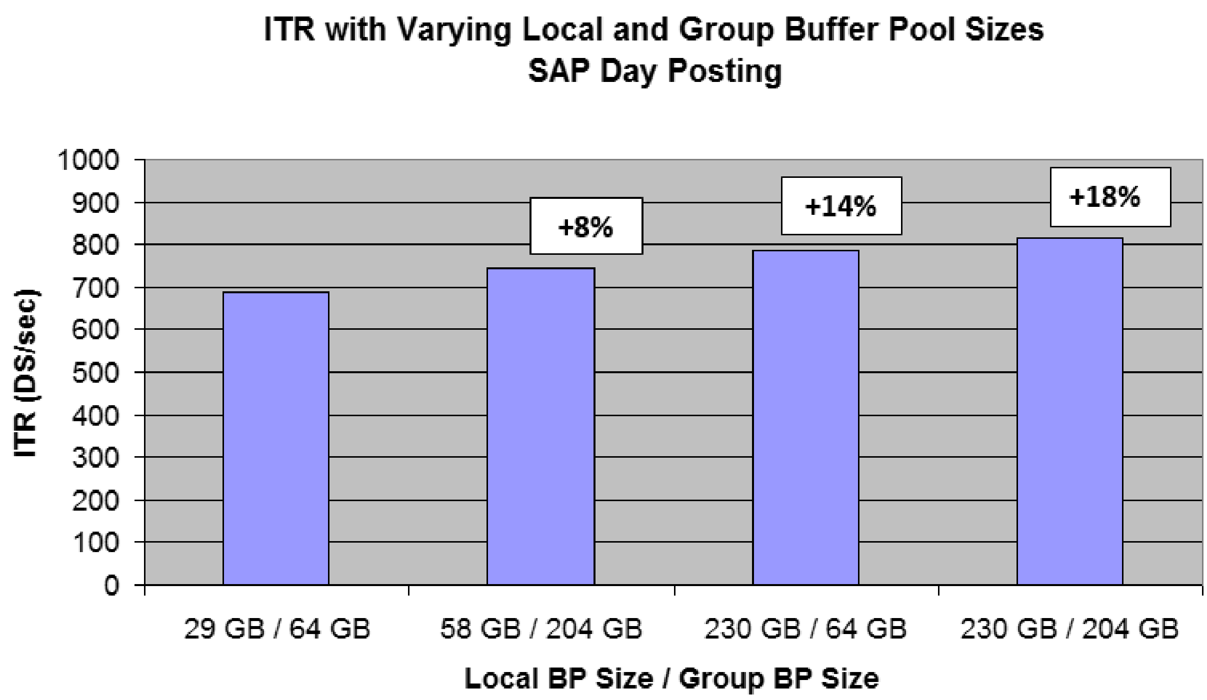

Figure 13: Local and Group Buffer Pool Scaling Effects on ITR

7.0 Single System Results and Analysis – Night Balancing

We repeated the same set of single system measurements using the "consolidated buffer pool" configuration that we did with the SAP Day Posting workload using the SAP Night Balancing workload. The only difference was the z/OS level. For the SAP Night Balancing measurements, we used z/OS 2.1. See section *4.3 DB2 Configuration* for an explanation of the "consolidated buffer pool" configuration.

These single system measurements were run on an IBM zEC12 with 12 general purpose processors. The amount of storage that was configured online varied from 256 GB to 1 TB depending on the amount of storage used for the DB2 buffer pools in the measurement. See section *4.0 Test Environment* for more details of the test environment.

The details of these measurements are summarized in the following table. See Table 7 on page 22 for how the storage was allocated among the individual buffer pools for these measurements. It was allocated the same way as it was for the single system Day Posting measurements.

Run id	S40919B2	S40918B1	S40919B1
Real Storage Configured per z/OS LPAR	256 GB	512 GB	1024 GB
Buffer Pool Storage per DB2 Member	**161 GB**	**320 GB**	**638 GB**
Database Server	zEC12	zEC12	zEC12
Number of z/OS LPARs	1	1	1
Number of CPs per z/OS LPAR	12	12	12
z/OS Level	2.1	2.1	2.1
LFAREA (for 1M frames)	180 GB	380 GB	675 GB
Number of Accounts	60M	60M	60M
Number of Batch Jobs	240	240	240
Average %CPU on z/OS	58.30%	57.50%	57.00%
Maximum Elapsed Time (in secs)	13,116	13,242	13,388
ETR (accounts per hour)	16.47M	16.31M	16.13M
Total DASD Rate (I/O per sec)	13,381	12,285	11,987
Total DB2 Synchronous Reads per sec	4008.89	3147.27	2785.63
DB2 Synchronous Reads per sec – TOT4K	2903.72	2049.77	1711.17
DB2 Synchronous Reads per sec – TOT8K	1017.47	1010.26	988.98
DB2 Synchronous Reads per sec – TOT32K	87.69	87.24	85.48
Total DB2 Dynamic Prefetch Reads per sec	814.50	810.98	794.82
DB2 Dynamic Prefetch Reads per sec – TOT4K	519.20	516.96	507.11
DB2 Dynamic Prefetch Reads per sec – TOT8K	163.21	162.50	159.05
DB2 Dynamic Prefetch Reads per sec – TOT32K	132.09	131.52	128.65
Total DB2 Pages Read via Dyn Prefetch per sec	8985.92	8945.69	8750.29
DB2 Pages Read via Dyn Prefetch per sec – TOT4K	7635.11	7600.53	7435.68
DB2 Pages Read via Dyn Prefetch per sec – TOT8K	996.98	992.86	970.14
DB2 Pages Read via Dyn Prefetch per sec – TOT32K	353.82	352.29	344.46
Total Buffer Pool Hit Ratio	94.30%	94.67%	94.82%
Buffer Pool Hit Ratio – TOT4K	93.66%	94.17%	94.36%
Buffer Pool Hit Ratio – TOT8K	96.66%	96.67%	96.67%
Buffer Pool Hit Ratio – TOT32K	72.55%	71.17%	71.17%

Table 18: Measurement Results - Buffer Pool Scaling with Night Balancing - Single System

The SAP Night Balancing workload is a batch processing workload. In these measurements, we processed 60 Million accounts using 240 batch jobs running in parallel. The maximum elapsed time for these jobs was in the neighborhood of just over 3.5 hours for all these measurements. All the jobs for a given measurement completed in about the same elapsed time due to the application implementation. The throughput (ETR) is the number of accounts processed per hour. Maximum elapsed time and throughput are the key metrics for the Night Balancing workload.

We found that as the total buffer pool memory was scaled from 161 GB to 638 GB, there was a 31% reduction in DB2 synchronous reads. However, the raw number of synchronous reads avoided with the larger buffer pools was only at most 1223 per second. Since there were approximately 230,000 getpages per second, this had very little impact on the total BP hit ratio, which was very good to start with. The processing of I/Os is very small compared to the processing of SQL.

We also found that there was very little reduction in sequential I/O as the buffer pools were scaled. This can be seen from the “DB2 Dynamic Prefetch Reads”, which is the number of prefetch I/Os triggered, and the “DB2 Pages Read per Dyn Prefetch”, which is the number of pages brought into the buffer pool by the prefetch I/Os. The SAP Night Balancing workload is sequential in nature with little re-reference of the data brought into the buffer pools. Therefore, adding memory to the DB2 buffer pools doesn’t improve the sequential hit ratio and it doesn’t improve the performance of this workload.

There was a small reduction in CP utilization, but no visible improvement in elapsed time or throughput. Given the moderate number of sync reads and the database residing on a DS8870 that does fast I/O, the I/O wait time was not a significant factor in the elapsed time.

The following table summarizes the effects on key performance indicators as memory is added to DB2 buffer pools while running the SAP Night Balancing workload.

Run id	BP Size	%CPU	%CPU Delta	Max Elapsed Time (secs)	Max Elapsed Time Delta	ETR (acct/hr)	ETR Delta	Sync Reads per sec	Sync Reads Delta
S40919B2	**161 GB**	58.30%	n/a	13,116	n/a	16.47M	n/a	4009	n/a
S40918B1	**320 GB**	57.50%	-1.4%	13,242	1.0%	16.31M	-1.0%	3147	-21%
S40919B1	**638 GB**	57.00%	-2.2%	13,388	2.1%	16.13M	-2.0%	2786	-31%

Table 19: Performance Summary - Night Balancing – Single System

8.0 Data Sharing Results and Analysis – Night Balancing

We repeated two of the 2-way data sharing measurements that we did with the SAP Day Posting workload using the SAP Night Balancing workload. One was the 29 GB LBP / 64 GB GBP measurement and the other was the 230 GB LBP / 204 GB GBP measurement. These are the measurements in Table 17 on page 42 using the smallest and largest amounts of storage.

The same buffer pool assignments and sizes were used for the Day Posting and Night Balancing measurements. However, the actual storage allocated to the group buffer pools in the Night Balancing measurements is less than what was allocated with Day Posting, because fewer buffer pools have group buffer pool dependent data. GBP8K2 (8K indexes) and GBP32K2 (32K indexes) do not get allocated by Night Balancing.

For these measurements, we ran 2-way data sharing on an IBM zEC12 with two z/OS LPARs, each with 6 general purpose processors and up to 512 GB of storage. We used the "consolidated buffer pool" configuration. See section *4.0 Test Environment* for more details of the test environment.

The details of these measurements are summarized in the following table. Note that the "Total DASD Rate", "Total DB2 Synchronous Reads per sec", "Total DB2 Dynamic Prefetch Reads per sec", and the "Total DB2 Pages Read via Dyn Prefetch per sec" are the sum from each LPAR.

Run id	S41001B1	S41002B1
Real Storage Configured per z/OS LPAR	64 GB	512 GB
Local Buffer Pool Storage per DB2 Member	**29 GB**	**230 GB**
Group Buffer Pool Storage	**49 GB**	**169 GB**
Database Server	zEC12	zEC12
Number of z/OS LPARs	2	2
Number of CPs per z/OS LPAR	6	6
z/OS Level	2.1	2.1
LFAREA (for 1M frames)	40 GB	320 GB
Number of Accounts	60M	60M
Number of Batch Jobs	240	240
Average %CPU on z/OS	57.55%	57.15%
Maximum Elapsed Time (in secs)	13,217	13,198
ETR (accounts per hour)	16.34M	16.37M
Total DASD Rate (I/O per sec)	18,022	15,172
Average %CF Utilization	1.1%	1.1%
Total DB2 Synchronous Reads per sec	6063.32	3512.55
DB2 Synchronous Reads per sec – TOT4K	4965.54	2402.61
DB2 Synchronous Reads per sec – TOT8K	1010.70	1021.27
DB2 Synchronous Reads per sec – TOT32K	87.07	88.67
Total DB2 Dynamic Prefetch Reads per sec	808.46	821.22
DB2 Dynamic Prefetch Reads per sec – TOT4K	515.67	523.98
DB2 Dynamic Prefetch Reads per sec – TOT8K	161.83	164.32
DB2 Dynamic Prefetch Reads per sec – TOT32K	130.97	132.94
Total DB2 Pages Read via Dyn Prefetch per sec	8914.26	9055.76
DB2 Pages Read via Dyn Prefetch Read per sec – TOT4K	7575.07	7697.49
DB2 Pages Read via Dyn Prefetch Read per sec – TOT8K	988.39	1002.11
DB2 Pages Read via Dyn Prefetch Read per sec – TOT32K	350.78	356.13
Total Buffer Pool Hit Ratio	93.38%	94.52%
Buffer Pool Hit Ratio – TOT4K	92.38%	93.95%
Buffer Pool Hit Ratio – TOT8K	96.66%	96.67%
Buffer Pool Hit Ratio – TOT32K	72.76%	72.66%

Table 20: Measurement Results - Buffer Pool Scaling with Night Balancing - Data Sharing

The following table shows how the total amount of storage for the local buffer pools was allocated among the individual local and group buffer pools for each of these measurements.

	S41001B1			S41002B1		
Run id	LBP Sizes per Member		GBP Sizes	LBP Sizes per Member		GBP Sizes
	Pages (K)	GB		Pages (K)	GB	
BP0	10	0.04	977M	10	0.04	977M
BP2	2,000	7.63	14G	16,000	61.04	50G
BP3	3,000	11.44	23G	24,000	91.55	76G
BP8K0	50	0.38		50	0.38	
BP8K1	480	3.66	10G	3,840	29.30	38G
BP8K2	300	2.29		2,400	18.31	
BP32K	10	0.31	196M	10	0.31	196M
BP32K1	60	1.83	977M	480	14.65	4G
BP32K2	60	1.83		480	14.65	
Total Size		29.41	49G		230.22	169G

Table 21: LBP and GBP Sizes for Night Balancing

These measurements show that using larger local and group buffer pools doesn't provide any significant overall system performance improvement when running the Night Balancing workload. We already have a good total buffer pool hit ratio with 29 GB of storage for local BPs and 49 GB for GBPs. So, even though there was a 42% reduction in the DB2 synchronous reads, the magnitude of the sync reads is only about 6000 per second in the base run. And, there was no reduction in sequential I/O as the buffer pools were scaled. This can be seen from the "DB2 Dynamic Prefetch Reads", which is the number of prefetch I/Os triggered, and the "DB2 Pages Read per Dyn Prefetch", which is the number of pages brought into the buffer pool by the prefetch I/Os.

As mentioned in section *7.0 Single System Results and Analysis – Night Balancing*, the SAP Night Balancing workload is sequential in nature with little re-reference of the data brought into the buffer pools. Therefore, adding memory to the DB2 buffer pools doesn't improve the sequential hit ratio and it doesn't improve the performance of this workload.

The SAP Night Balancing workload in our environment results in only 14% of pages being GBP dependent. There is minimal group buffer pool interest because of the way the application "assigns" accounts to SAP instances for processing and how we partition the key banking table spaces. Under normal circumstances, each SAP instance is assigned to a specific application server and each application server is assigned to a specific data sharing member. Each data sharing member has affinity to its set of adjacent accounts. The very low CF utilization also illustrates the minimal group buffer pool interest.

Since we are running with GBPCACHE CHANGED, only the GBP dependent data has the potential for being cached in the group buffer pool. Therefore, adding memory to the group buffer pools doesn't provide much benefit. Also, as recommended, we us the DB2 member cluster feature. See section *6.2 Group Buffer Pool Scaling* for more information on GBPCACHE and the effects of using DB2 member clustered tables.

The "Total Group Buffer Pool Hit Ratio" indicates the effectiveness of finding a referenced page in the group buffer pool given that it was not found in the local buffer pool. It is calculated from data in the DB2 Performance Expert Statistics Report.

Total Group Buffer Pool Hit Ratio =
(SYN.READ(XI)-DATA RETURNED + SYN.READ(NF)-DATA RETURNED) /
(SYN.READ(XI)-DATA RETURNED + SYN.READ(XI)-NO DATA RETURN +
SYN.READ(NF)-DATA RETURNED + SYN.READ(NF)-NO DATA RETURN)

Since there is minimal group pool interest and little traffic to the GBP with the Night Balancing workload, the "Total Group Buffer Pool Hit Ratio" can be misleading, and therefore, we did not document it for these measurements. Both of these measurements have very few GBP misses (eg. NO DATA RETURN) that result in synchronous I/O.

The following table summarizes the effects on key performance indicators as memory is added to DB2 local and group buffer pools while running the SAP Night Balancing workload.

Run id	LBP Size per member --- GBP Size	%CPU	%CPU Delta	Max Elapsed Time (secs)	Max Elapsed Time Delta	ETR (acct/hr)	ETR Delta	Sync Reads per sec	Sync Reads Delta
S41001B1	**29 GB** **49 GB**	57.55%	n/a	13,217	n/a	16.34M	n/a	6063	n/a
S41002B1	**230 GB** **169 GB**	57.15%	-0.7%	13,198	-0.1%	16.37M	0.1%	3513	-42%

Table 22: Performance Summary - Night Balancing – Data Sharing

9.0 z13 Measurement Results and Analysis

We repeated two of the single system measurements that we did on an IBM zEC12 with the SAP Day Posting workload using the “consolidated buffer pool” configuration on an IBM z13 machine. We wanted to confirm that the results on z13 would be similar to those that we saw on zEC12. We chose to repeat the tests that used the “smallest” (161 GB) and “largest” (638 GB) amounts of memory for the buffer pools. Due to the limited availability of test time on the z13 machine, we had to limit the scope of our tests.

These single system measurements were run on an IBM z13 with 12 general purpose processors with z/OS 2.1. The amount of storage configured online ranged from 256 GB to 1 TB depending on the test. The “consolidated buffer pool” configuration was used. See section *4.0 Test Environment* on page 10 for more details on the test environment. The only change in the test environment between the zEC12 and z13 measurements is the System z processor for the SAP Database Server. See section *4.3 DB2 Configuration* for an explanation of the buffer pool configuration. Table 7 on page 22 shows how the storage was allocated among the buffer pools for these measurements.

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

Run id	S41103B1	S41102B1
Real Storage Configured per z/OS LPAR	256 GB	1024 GB
Buffer Pool Storage per DB2 Member	**161 GB**	**638 GB**
Database Server	z13	z13
Number of z/OS LPARs	1	1
Number of CPs per z/OS LPAR	12	12
z/OS Level	2.1	2.1
LFAREA (for 1M frames)	180 GB	675 GB
Number of Users	10560	10560
ETR (DS/sec)	709.50	939.40
Average %CPU on z/OS	68.25%	72.63%
ITR (DS/sec)	1039.56	1293.40
Database (DB) Request Time in secs	0.729	0.245
Total DASD Rate (I/O per sec)	66,026	26,807
Total DB2 Synchronous Reads per sec	34.4K	1.0K
DB2 Synchronous Reads per sec – TOT4K	25.9K	759.87
DB2 Synchronous Reads per sec – TOT8K	2630.96	126.13
DB2 Synchronous Reads per sec – TOT32K	5912.08	125.02
Total Buffer Pool Hit Ratio	89.65%	99.82%
Buffer Pool Hit Ratio – TOT4K	91.27%	99.80%
Buffer Pool Hit Ratio – TOT8K	85.59%	99.88%
Buffer Pool Hit Ratio – TOT32K	87.06%	99.79%

Table 23: z13 Measurement Results with Day Posting - Single System

We found the results on z13 from buffer pool scaling to be very good, very similar to what we saw on zEC12. As the total buffer pool memory was scaled from 161 GB to 638 GB, there was a 97% reduction in DB2 synchronous reads and the BP Hit Ratio improved from close to 90% to close to 100%. There was a 24% improvement in ITR, a 32% improvement in ETR, and a 66% drop in Database (DB) Request time.

The following table summarizes the effects on key performance indicators as memory is added to DB2 buffer pools using the "consolidated buffer pool" configuration while running the SAP Day Posting workload on z13.

Run id	BP Size	ITR DS/sec	ITR Delta	%CPU	ETR DS/sec	ETR Delta	DB Req Time (secs)	DB Req Time Delta	Sync Reads per sec	Sync Reads Delta
S41103B1	**161 GB**	1039.56	n/a	68.25%	709.50	n/a	0.729	n/a	34.4K	n/a
S41102B1	**638 GB**	1293.40	24%	72.63%	939.40	32%	0.245	-66%	1.0K	-97%

Table 24: Performance Improvements with Buffer Pool Scaling on z13

The following four charts show the results of these measurements pictorially. Bar charts are used since we only have two formal measurement points. However, based on an informal measurement that we did at 320 GB BP size on z13 that is not documented here, the z13 results show sloping curves similar to the ones from our zEC12 measurements. The analysis of the zEC12 measurements discussed in section *5.3 Analysis – Single System* applies to these measurements on z13.

The following chart shows the dramatic reduction in DB2 synchronous reads as the DB2 buffer pools are scaled from 161 GB to 638 GB when running the Day Posting workload on z13. The sync reads dropped from 34.4K per sec to 1K per sec, a 97% decrease.

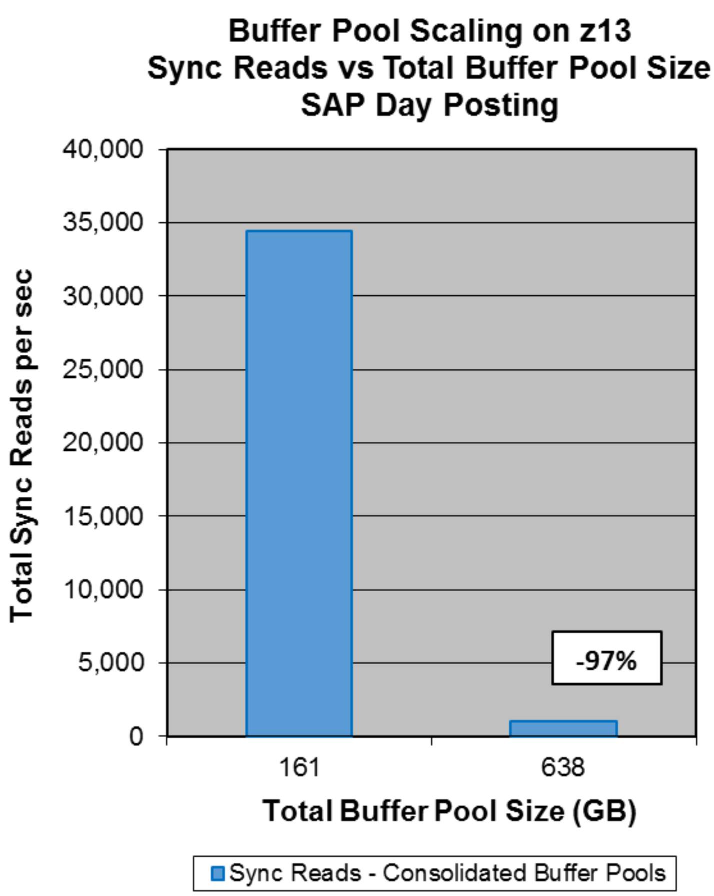

Figure 14: Buffer Pool Scaling Effects on Synchronous Reads on z13

The following chart shows the significant reduction in DB request time as the DB2 buffer pools are scaled from 161 GB to 638 GB when running the Day Posting workload on z13. The DB request time dropped from 0.729 seconds to 0.245 seconds, a reduction of 66%.

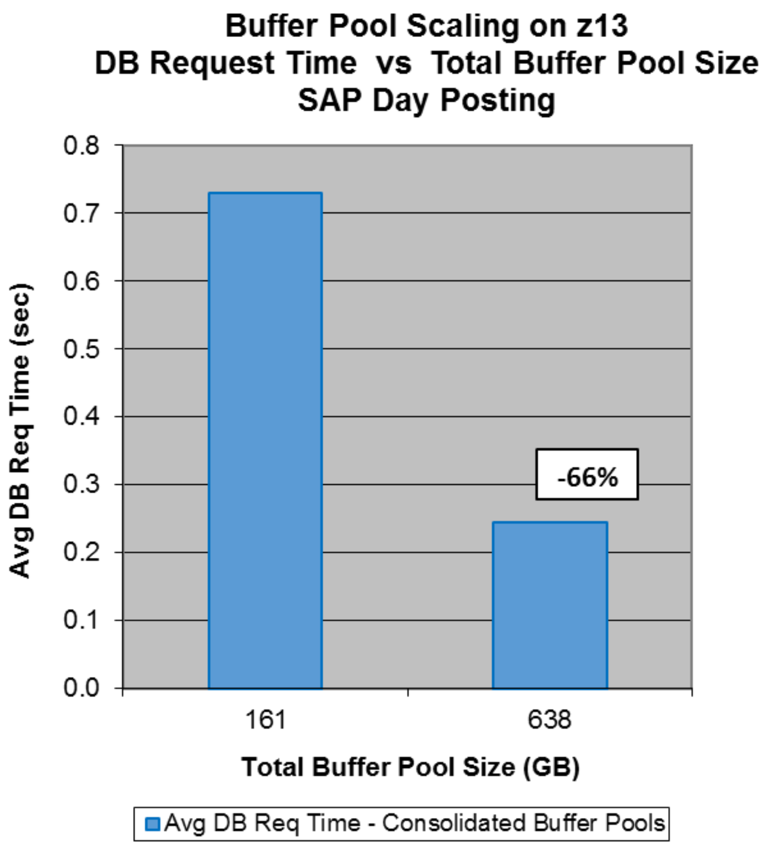

Figure 15: Buffer Pool Scaling Effects on DB Request Time on z13

The following chart shows the significant improvement in thoughput (ETR) as the buffer pools are scaled from 161 GB to 638 GB when running the Day Posting workload on z13. The throughput increased from 709.50 DS/sec to 939.40 DS/sec, an improvement of 32%.

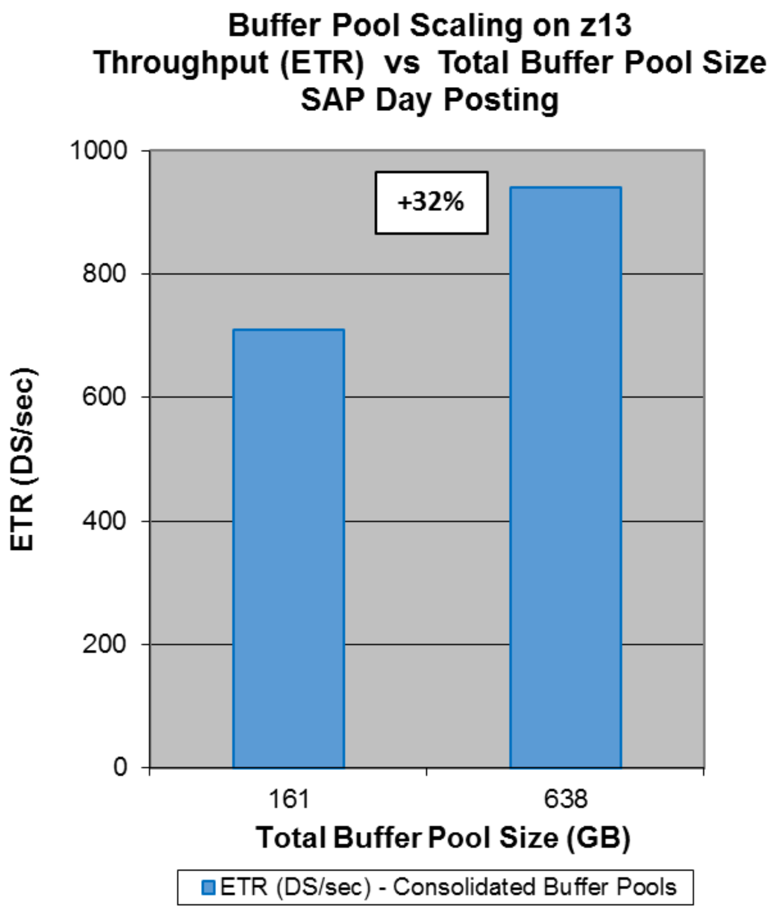

Figure 16: Buffer Pool Scaling Effects on ETR on z13

The following chart shows the significant improvement in ITR as the buffer pools are scaled from 161 GB to 638 GB when running the Day Posting workload on z13. The ITR increased from 1039.56 DS/sec to 1293.40 DS/sec, an improvement of 24%.

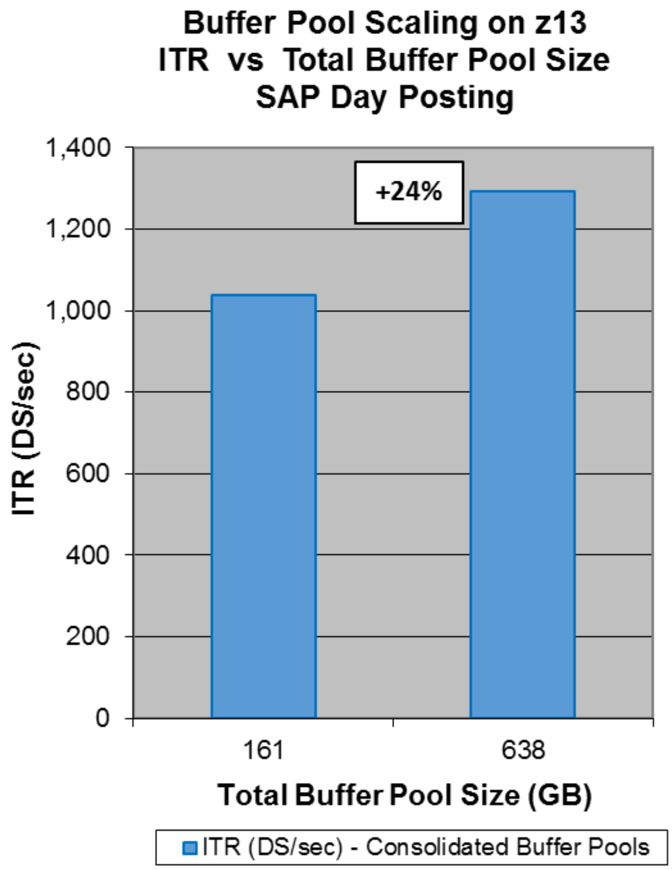

Figure 17: Buffer Pool Scaling Effects on ITR on z13

The significant performance benefits seen on zEC12 from adding memory to the DB2 buffer pools extend to z13. The following chart clearly illustrates this. It shows the buffer pool scaling ITR ratios (ITRR) for zEC12 and z13. Using the 161 GB BP size as a base point of 1 on both zEC12 and z13, it shows the ITR improvement gained by scaling the DB2 buffer pools from 161 GB to 638 GB is very similar on both machines. The ITR ratio on zEC12 is 1.21 or a 21% improvement. On z13, the ITR ratio is 1.24 or a 24% improvement. With more real memory available on z13 and with attractive pricing of memory, it behooves clients to explore opportunites to make use of more memory.

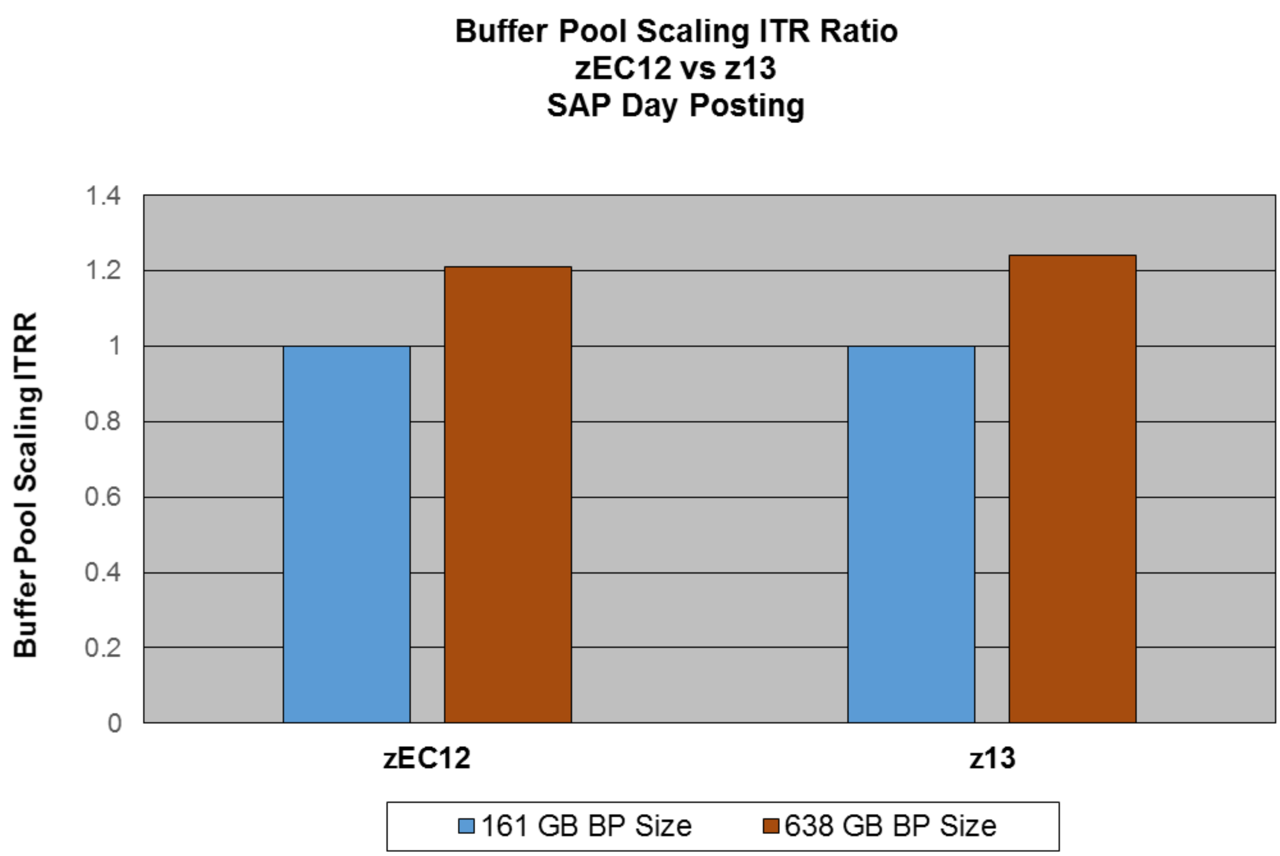

Figure 18: Buffer Pool Scaling ITRR - zEC12 and z13

10.0 Considerations

The performance benefits from adding memory to buffer pools and the allocation of the additional memory between local buffer pools and group buffer pools is very much configuration and workload dependent. Reducing DB2 synchronous read I/O is key to achieving performance improvements.

Keep in mind the following:

- When scaling the size of the local buffer pools, the amount of memory needed can multiply quickly depending on the number of data sharing members in the data sharing group.
- When adding memory to the local buffer pools, be sure to tune the group buffer pools appropriately to avoid directory entry reclaims.
- When adding memory to the group buffer pools, twice the amount of memory being added to the group buffer pools is needed for redundancy. In a data sharing environment, there must be storage available in a second CF for the GBP cache structures to be rebuilt in a recovery scenario.

Although it is highly recommended to use GBP duplexing, all the measurements detailed in this report did not use it. Our test methodology was to use a staged approach introducing factors one at a time to see how each affected the results. However, we did do two measurements with GBP duplexing enabled, one used 52 GB of memory for the group buffer pools and the other used 204 GB. We saw similar performance improvements with and without GBP duplexing when memory was added to the group buffer pools.

All the measurements detailed in this report were run in a 2-way data sharing environment. However, we also executed a number of measurements in a 4-way data sharing environment. We found that the performance benefits of adding memory to the buffer pools that we saw in 2-way data sharing extended to 4-way data sharing.

The performance benefits that we saw with the SAP Day Posting workload in a data sharing environment from adding memory to the local and group buffer pools may appear to be high. However, keep in mind, that in our 3-tier environment the database serving portion of the Day Posting workload running on the zEC12 is a "pure DB2 workload". The SAP application server is running outboard. Also, the application data access pattern in this workload is highly random.

However, our measurements do illustrate that performance improvements can be achieved depending on the configuration and workload, although perhaps on a more modest scale, by adding memory to local and group buffer pools.

There are several factors that should be considered when adding memory to buffer pools.

- Consider your current buffer pool tuning. If your buffer pools are already well-tuned with minimal synchronous I/O, good buffer pool hit ratios, and little group buffer pool directory reclaim activity then adding memory to your buffer pools will not provide any significant performance benefits.
- Consider your DB2 configuration. If you are running data sharing then consider the number of DB2 members in your data sharing group. This will affect the total amount of memory needed when increasing the size of local buffer pools. The total amount of memory needed for the group buffer pools is twice the defined size for redundancy. In a data sharing environment, there must be storage available in a second CF for the GBP cache structures to be rebuilt in a recovery scenario.
- Consider your workload characteristics. The data access patterns of a particular workload will influence the results of adding memory to buffer pools. Eliminating random I/O, which results from random data access, is more likely to reduce DB2 synchronous reads than sequential I/O. In data sharing, the update intensity of the workload, the amount of GBP dependent data, and the use of the

DB2 member cluster feature may influence where additional memory should be added, to the local or the group buffer pools.

- Consider using the DB2 buffer pool simulation introduced in DB2 11 APAR PI22091 to identify buffer pools that could benefit from additional memory. Good candidates for buffer pool simulation are buffer pools with a significant number of synchronous reads and that contain pages that are likely to be referenced again.

For guidance on buffer pool tuning, see references [1,2] on page 59. For more information about DB2 buffer pool simulation, see reference [8] on page 59.

11.0 Conclusions

Significant performance improvements were seen in both single system and data sharing environments when more memory was used for larger DB2 buffer pools while running the SAP Banking Services (SBS) Day Posting workload on an IBM zEC12 or a z13.

The performance improvements that can be achieved by increasing the sizes of DB2 buffer pools are very much configuration and workload dependent. Reducing the amount of synchronous I/O is key to achieving overall system performance improvements. Many of our clients should see measurable performance benefits by adding memory to their DB2 buffer pools, however, it is important to note that performance benefits will vary and may not be seen in all environments.

Below are some of our other findings and observations that are worth re-iterating.

- The cost of consolidating or reducing the number of buffer pools, while keeping the total amount of memory used for them constant, was little to none in a single system environment.
- Adding memory to the DB2 buffer pools to improve performance is more complex in data sharing than single system since you need to decide how to allocate the additional storage between the local and group buffer pools. The decision on how to spread the additional memory is configuration and workload dependent.
- If you add memory to local buffer pools in a data sharing environment then you may need to tune the group buffer pools to avoid directory entry reclaims.
- In general, for most workloads, adding some memory to both the local and group buffer pools will be the best approach.
- Even though in this document we are highlighting the benefits of reducing I/O, it is still important to maintain a good DASD I/O subsystem. By adding memory to DB2 buffer pools, DB2 synchronous reads can be reduced, but there still will be DB2 asynchronous I/O, including prefetch and deferred writes, as well as the critical DB2 synchronous logging I/O.
- It is highly recommended to “page fix” local buffer pools where appropriate. When a buffer pool is page fixed, the page frames that hold the buffer pool will be fixed in real storage and can’t be stolen by z/OS. Fixed page frames also allow more efficient CF activity. This is good for buffer pools that have I/O activity, provided that there is enough real storage available.
- It is highly recommended to use 1M fixed large frame support for the buffer pools that are page fixed. Using 1 MB frames helps the z/OS Real Storage Manager (RSM) efficiently manage the overall real storage.

12.0 References

1. IBM Corp. 2014. *DB2 11 for z/OS Buffer Pool Monitoring and Tuning Redpaper* REDP-5092-00 http://www.redbooks.ibm.com/redpieces/abstracts/redp5092.html?Open
2. IBM Corp. 2013. *DB2 for z/OS Best Practices – DB2 for z/OS Data Sharing Tuning Group Buffer Pools*; John Campbell; https://www.ibm.com/developerworks/community/wikis/home?lang=en#!/wiki/W7c1e3103820b_4d9e_add6_b5ade3280dcb/page/DB2%20for%20zOS%20Data%20Sharing%20Tuning%20Group%20Buffer%20Pools
3. 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
4. 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
5. IBM Corp. 2013. *DB2 11 for z/OS Technical Overview* http://www.redbooks.ibm.com/Redbooks.nsf/RedpieceAbstracts/sg248180.html?Open
6. *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)
7. IBM Corp. 2009. *IBM System z and System Storage DS8800: Accelerating the SAP Deposits Management Workload With Solid State Drives* http://www-01.ibm.com/support/docview.wss?uid=tss1wp101442
8. IBM Corp. 2014. *DB2 11 for z/OS Managing Performance*, SC19-4060-06
9. IBM Corp. 2014. *System z: Advantages of Configuring More Memory for DB2 Buffer Pools* http://www-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP102464