IBM Technical Brief

IBM Enterprise System®: SAP® Bank Analyzer 8.0 AFI Loan Tests

Authors:

Brenda W. Beane
Seewah Chan
John W. Langer
Paul Lekkas
Veng K. Ly
Michael R. Sheets

Document Owner:

Michael R. Sheets
IBM Americas Solutions Advanced Technical Skills

Version: 1.01
Date: January 31, 2013

Filename: System_z_SAP_BA8.pdf

TABLE OF CONTENTS

FIGURES

TABLES

1 Introduction

This paper documents some of the tests we did with SAP's most recent release of Bank Analyzer (BA) - Release 8.0 (BA8). This is a continuation of and builds on our previous experiences with SAP's Bank Analyzer 7.0 (BA7) [1]. It is assumed the reader is familiar with that work. We will not go into much depth on BA's functionality. However, SAP provides some of this in [2]. IBM and SAP have worked together for several years on many aspects of SAP Banking solutions for large enterprises. One example is the *SAP for Banking on System z Reference Architecture* [3].

In addition to up-leveling the BA application, we also up-leveled to DB2 10 for z/OS [4, 5, 6, 7], z/OS, the zEnterprise Server [8], the System Storage [9], as well as changing the workload. All this will be discussed in more detail in later sections.

These tests were not formal benchmarks.

2 Reminders

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

The results shown are based on specific workloads run in a somewhat controlled environment. However, these tests were not intended to be definitive performance measurements. The actual throughput that any user will experience will vary considerably from these results. Therefore, no assurance can be given that an individual user will achieve throughput equivalent to the performance stated here.

All customer examples cited or described in this presentation are presented as illustrations of the manner in which some customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics will vary depending on individual customer configurations and conditions.

This publication was produced in the United States. IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the product or services available in your area.

Information about non-IBM products is obtained from the manufacturers of those products or their published announcements. IBM has not tested those products and cannot confirm the performance, compatibility, or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

3 Trademarks

IBM, the IBM logo, and ibm.com® are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate US registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the web at “Copyright and trademark information” at www.ibm.com/legal/us/en/copytrade.shtml.

SAP, SAP R/3, SAP ERP, mySAP, mySAP.com, xApps, xApp, ABAP, BAPI, SAP NetWeaver and all SAP product and service names mentioned herein are trademarks or registered trademarks of SAP AG in Germany and in several other countries all over the world.

Adobe is a registered trademark of Adobe Systems Incorporated in the United States, and/or other countries.

InfiniBand is a registered trademark of the InfiniBand Trade Association.

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

Windows and Microsoft are registered trademarks of Microsoft Corporation in the United States, other countries, or both.

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

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

Unicode is a trademark of Unicode, Inc.

Other product and service names might be trademarks of IBM or other companies.

4 Feedback

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

5 Acknowledgements

The authors would like to recognize the contributions of Andrea Fuga, John Mullin, Georg Mayer, Wen-Tzer T. Chen, and Damon Bull (all with IBM), the IBM System z Benchmark Center in Poughkeepsie, and Michael Reuter, Corina Farcas, Hendrik Rusch, Ralf Knoebl, and Knut Willems (all with SAP AG).

6 Workload and Test Background

SAP's BA offers an extremely wide set of features and functions. For this set of tests, we used a narrow set of BA functions oriented towards loans processing to construct a workload. This was a different workload than we had for our earlier BA7 tests. For our BA8 tests, we used SAP's Accounting for Financial Instruments (AFI) with Financial Calculator (FCL) cash flow generation functions.

Previously with BA7, we used BA's Financial Mathematics (FIMA) cash flow engine. Since then, SAP BA8 made the Financial Calculator (FCL) available as an alternative cash flow engine. FCL is more flexible but also requires more CPU resources. We estimate that FCL consumed about twice the application server CPU resources as Financial Mathematics (FIMA).

We used a mix of loan types.

- 70% were installment loans with two year fixed interest rate loans.
- 30% were mortgage loans with 10-year variable interest loans.

The processing was done via SAP batch as a chain of process steps. Typically, this chain has five logical steps - four are required and one is optional.

- Post External Business Transactions (PEBT): This process reads initialization and operational business transactions that are relevant for accounting from the Source Data Layer (SDL). The external business transactions are converted from the SDL format to the Accounting format and transferred to Accounting. In Accounting, the system posts the accounting documents for the specified accounting system. The system posts items from the operational business transactions that are to be processed in financial position management to clearing figures in financial position management. External business transactions are posted initially to a position clearing account. The position clearing account is cleared later in the Update of Secondary Business Transactions.
- Results Data Layer (RDL) Pre-aggregation: This step is optional but recommended because it could improve performance of the accounting process by creating aggregated results for an aggregation level and a filter. Delta aggregation is also possible for merging pre-aggregated data and newer data to provide up-to-date aggregated results. However, in the case of loans, we did not see any performance improvement with RDL pre-aggregation.
- Update Secondary Business Transactions (USBT): This process is used to generate internal business transactions, such as determining the exchange rate obtained for a sale. Updating secondary business transactions triggers position clearing.
- Key Date Valuation (KDV): This process values financial positions on a key date. The key date valuation comprises multiple calculation steps. These steps include the calculation of amortization, fair value, deferred tax, and the distribution of revaluation reserves across time. Users can define their own key date valuation types. For example, users can do this to reflect a change in tax rates. During the key date valuation process, the system calls different functions (managers) such as the foreign currency valuation. Users can display the accounting documents that are generated by the key date valuation in the Results Data Layer.
- General Ledger Connector Document Prep (GL): This process is used to create general ledger documents, in which postings from sub-ledger documents are aggregated. To create general ledger documents, users extract the required data from the Results Data Layer (RDL). The system checks the consistency of the general ledger documents and saves them in the persistence layer.

While we did run the optional RDL pre-aggregation steps in all cases, we did not report it in this paper because it can be run outside the critical batch window. As mentioned above, in the case of loans, we did not see any performance improvement with RDL. As a result, for these tests, our key performance indicator (KPI) metric was the total elapsed time of PEBT, USBT, KDV, and GL. All but one of the tests shown here were done with 1.5 million loans. All these options were chosen because they were felt to be representative of what real customers might use.

The measurements we did were stress tests – not SAP certified benchmarks. We did invest some time in doing typical tuning, as discussed below. While a major goal of these measurements was to get as much throughput as possible within schedule constraints, we did not resort to "benchmark specials". For example, we used the same buffer pool settings for all the steps. Below is an overview of the different types of tuning we performed.

6.1 General Tuning

- We used NetWeaver 7.1 EHP1 to exploit 64-bit addressability with AIX 7.1 and its better memory management.
- We used the Optimized Latency Mode (OLM) option of the OSA-Express4S adapter to improve the elapsed time for application server to DB server communication.
- We used HyperPAV for the DS8800 to reduce disk I/O queuing.
- We used High Performance FICON for System z (zHPF) w/ multi-track support to improve efficiency of I/O resources.
- We used DB2 10 and System z hardware and software features including Index I/O Parallelism for Insert, Safe Query Optimization, and large 1 MB Page Support.
- We used DB2 striped logs.
- We used MAXKEEPD=64K to minimize the number of prepares.

6.2 Application Server Tuning

- We adjusted the number of SAP instances to eight per Power 780 to increase the SAP memory pools to accommodate more parallel jobs.
- We adjusted the number of parallel batch jobs based on application server CPU consumption of the workload, the number of loans being processed, and the available processor capacity.
- We installed AIX 7.1 to take advantage of improvements in AIX as well as its exploitation of 16 MB pages.

6.3 Database Tuning

- We identified "hot" tables using SAP transaction st04.
- We spread out these tables/indexes on DASD.
- We isolated these tables/indexes to unique DB2 buffer pools.
- We used larger DB2 buffer pools.
- We used the DB2 member cluster feature for heavily inserted tables.
- We improved SQL access paths with up-to-date RUNSTATS.

The buffer pool settings were:

BP name	PGFIX	VP Pages (K)	VPSEQT [%]	VPPSEQT [%]	VPXPSEQT [%]	DWQT [%]	VDWQT [%]	VDWQT [pages]	Page steal
BP0	YES	100	80	50	0	50	5	0	LRU
BP1	NO	120	100	50	0	50	5	0	LRU
BP2	YES	4,000	80	50	0	50	5	0	LRU
BP3	YES	4,000	50	50	0	50	5	0	LRU
BP4	YES	400	80	50	0	50	5	0	LRU
BP5	YES	1,200	50	50	0	50	5	0	LRU
BP7	YES	4,000	50	50	0	50	5	0	LRU
BP8	YES	400	80	50	0	50	5	0	LRU
BP9	YES	1,200	50	50	0	50	5	0	LRU
BP10	YES	400	80	50	0	50	5	0	LRU
BP11	YES	400	50	50	0	50	5	0	LRU
BP12	YES	400	80	50	0	50	5	0	LRU
BP13	YES	400	50	50	0	50	5	0	LRU
BP14	YES	400	80	50	0	50	5	0	LRU
BP15	YES	400	50	50	0	50	5	0	LRU
BP40	YES	10	80	50	0	50	5	0	LRU
BP8K0	YES	800	80	50	0	50	5	0	LRU
BP8K1	YES	800	80	50	0	50	5	0	LRU
BP8K2	YES	1,600	80	50	0	50	5	0	LRU
BP16K0	YES	10	80	50	0	50	5	0	LRU
BP16K1	YES	10	80	50	0	30	5	0	LRU
BP32K	YES	15	80	50	0	50	5	0	LRU
BP32K1	YES	15	80	50	0	50	5	0	LRU
BP32K3	NO	15	80	50	0	50	10	0	LRU

Table 1: DB2 Buffer Pool Settings

To help improve performance, these BA8 measurements had more than twice the total defined total buffer pool size compared to our previous BA7 measurements.

The hot tables are listed below. In general, these are the largest and most heavily inserted tables. These tables are Reorder Row Format (RRF), Partition by Growth (PBG), and compressed. Many tables in BA are "generated" tables so the table names are unique to each SAP system. RDL table names begin with /1BA/HM_. The character to the left of the client number (e.g., 701 in our case) describes Header tables (H), Version tables (V), Buffer tables (B), and Data tables (_).

Table or Index Name	Table or Index	Client	Area	Object Name	Structure Name	BP	Member Cluster
/1BA/HM_G3MG_701	T	701	SBA	SF		BP8K2	yes
/1BA/HM_G3MG_701~0	I					BP7	
/1BA/HM_G3MG_701UI	I					BP7	
/1BA/HM_G3MG_701FP	I					BP7	
/1BA/HM_G3MG_701UH	I					BP7	
/1BA/HM_G3MG_701BK	I					BP7	
/1BA/HM_G3MG_701UA	I					BP7	

Table or Index Name	Table or Index	Client	Area	Object Name	Structure Name	BP	Member Cluster
/1BA/HM_G3MG_701ZG	I					BP7	
/1BA/HM_G3MGV701	T	701	SBA	SF		BP4	no
/1BA/HM_G3MGV701~0	I					BP5	
/1BA/HM_G3MGV701VG	I					BP5	
/1BA/HM_G3MGH701	T	701	SBA	SF		BP4	no
/1BA/HM_G3MGH701~0	I					BP5	
/1BA/HM_G3MGH70150	I					BP5	
/1BA/HM_NYC4_701	T	701	SBA	S_FA_FPB1D		BP8	no
/1BA/HM_NYC4_701~0	I					BP9	
/1BA/HM_NYC4_701LF	I					BP9	
/1BA/HM_NYC4_701Z1	I					BP9	
/1BA/HM_NYC4B701	T	701	SBA	S_FA_FPB1D		BP8	no
/1BA/HM_NYC4B701~0	I					BP9	
/1BA/HM_NYC4B701Z2	I					BP9	
/1BA/HM_NYC4H701	T	701	SBA	S_FA_FPB1D		BP8	no
/1BA/HM_NYC4H701~0	I					BP9	
/1BA/HM_NYC4H70150	I					BP9	
/1BA/HM_NYC4V701	T	701	SBA	S_FA_FPB1D		BP8	no
/1BA/HM_NYC4V701~0	I					BP9	
/1BA/HM_NYC4V70150	I					BP9	
/1BA/HM_NYC4V701VG	I					BP9	
/1BA/HM_0IXCH701	T	701	SBA	S_KFCBD2		BP12	yes
/1BA/HM_0IXCH701~0	I					BP13	
/1BA/HM_0IXCH70150	I					BP13	
/1BA/HM_0IXCV701	T	701	SBA	S_KFCBD2		BP12	no
/1BA/HM_0IXCV701~0	I					BP13	
/1BA/HM_0IXCV70150	I					BP13	
/1BA/HM_0IXCV701VG	I					BP13	
/1BA/HM_0IXC_701	T	701	SBA	S_KFCBD2		BP12	no
/1BA/HM_0IXC_701Z1	I					BP13	
/1BA/HM_0IXC_701~0	I					BP13	
/1BA/HM_1J3IH701	T	701	SBA	S_FPCBD2		BP14	yes
/1BA/HM_1J3IH701~0	I					BP15	
/1BA/HM_1J3IH70150	I					BP15	
/1BA/HM_1J3IV701	T	701	SBA	S_FPCBD2		BP14	no
/1BA/HM_1J3IV70150	I					BP15	
/1BA/HM_1J3IV701VG	I					BP15	
/1BA/HM_1J3IV701~0	I					BP15	
/1BA/HM_1J3I_701	T	701	SBA	S_FPCBD2		BP14	no
/1BA/HM_1J3I_701Z1	I					BP15	
/1BA/HM_1J3I_701~0	I					BP15	

Table or Index Name	Table or Index	Client	Area	Object Name	Structure Name	BP	Member Cluster
/1SGS/4F0AD302	T				_R0_SBA________SF_ ________	BP4	yes
/1SGS/4F0AD302~0	I					BP5	
/1SGS/4F0AD302~DEF	I					BP5	
/1SGS/4F0AE202	T				_R0_SBA________S_FA_FPB1D	BP10	no
/1SGS/4F0AE202~0	I					BP11	
/1SGS/4F0AE202~DEF	I					BP11	
/1SGS/4F0AE502	T				_R0_SBA________S_FPCBD2__	BP10	yes
/1SGS/4F0AE502~0	I					BP11	
/1SGS/4F0AE502~DEF	I					BP11	
/1SGS/4F0AE902	T				_R0_SBA________S_KFCBD2__	BP10	yes
/1SGS/4F0AE902~0	I					BP11	
/1SGS/4F0AE902~DEF	I					BP11	
DB2DB02TS_HIS	T					BP14	yes
DB2DB02TS_HIS~0	I					BP15	
DB2DB02TS_HIS~TS	I					BP15	

7 Configurations

7.1 Hardware Environment

System z DB Server: Tests were performed on a single z196 Model M15 with a total of 192 GB installed. The runs utilized one dedicated LPAR for z/OS DB2 10 with up to 15 processors and 192 GB.

Storage: Two dual frame IBM System Storage Server DS8800 Model 2421-951/95E with 576 HDDs and 384 GB cache each. The IBM System Storage Servers were attached to the z196 by sixteen long wave FICON Express8 connections. The capacity was about 164 TB of available storage capacity for database, logs, and FlashCopy sets.

The BA 8.0 database used 161 3390 Mod 54 volumes - about 8.7 TB. It had 32 active logs of 4GB size each, striped across four volumes on four different ranks.

Application Servers: We used up to three application servers. One for the Enqueue server, the others were regular application servers. The Enqueue server was run on an IBM Power 740 with 16 3.5 GHz processor cores utilizing SMT and 256 GB memory. The remaining application server functions were run on up to two IBM Power 780 systems each with 64 3.9 GHz cores utilizing SMT and 512 GB memory. Each of these had one AIX LPAR with eight SAP instances.

Network: A dedicated 10 Gb Ethernet network was used for all connections. Each of the application servers was connected via four 10 Gb Ethernet adapters through a 10 Gb Ethernet switch to the z196 via four OSA-Express4S adapters.

Below is a conceptual view of the configuration.

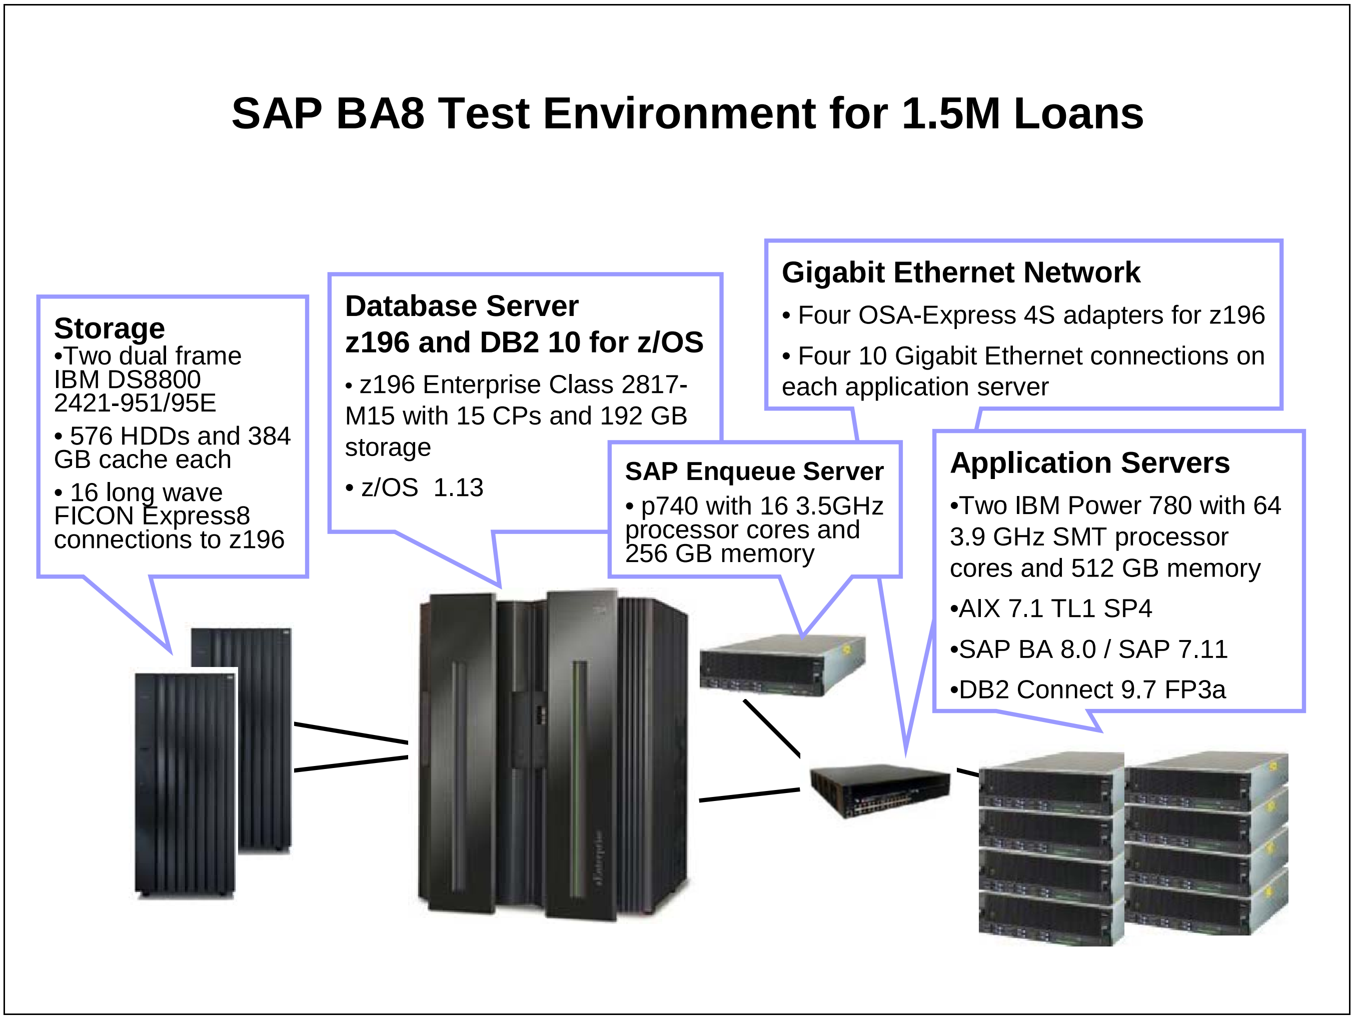

Figure 1: Conceptual View of BA8 Test Environment Configuration

7.2 Software Environment

<u>z/OS</u>
z/OS release 01.13.00 (R1.13)
DB2 10 dated January 2012

<u>AIX</u>
AIX 7.1 TL1 SP4

<u>SAP Kernel Information</u>

SAP Kernel : 711
Kernel Patch number : 135
DBSL Patch No. : 157
ICU Version : 3.4 Unicode Version 4.1
libsapu16 Version : 1.0025 May 16 2011 20:12:01
IBM DB2 Connect "Thin client" side: Driver for ODBC, CLI, JDBC and SQLJ - Version 9.7 FP3a

SAP Application Levels
SAP EHP1 for SAP NetWeaver 7.1

Software Component	Release	Level	Highest Support Package	Short Description of Software Component
SAP_ABA	711	0007	SAPKA71107	Cross-Application Component
SAP_BASIS	711	0007	SAPKB71107	SAP Basis Component
PI_BASIS	711	0007	SAPK-71107INPIBASIS	Basis Plug-In
ST-PI	2008_1_710	0004	SAPKITLRE4	SAP Solution Tools Plug-In
SAP_BW	711	0007	SAPKW71107	SAP Business Warehouse
FINBASIS	700	0016	SAPK-70016INFINBASIS	Financial Basis
SEM-BW	700	9	SAPK-70009INSEMBW	SEM-BW 700: Add-On Installation
BI_CONT	711	0002	SAPK-71102INBICONT	Business Intelligence Content
FSAPPL	400	0001	SAPK-40001INFSAPPL	SAP Banking Services
ST-A/PI	01N_710BCO	0001	SAPKITAB8H	Servicetools for other App./Netweaver200

Table 2: SAP Application Levels

As can be seen above, BA consists of SAP's core banking, BI Content, and BW. BA requires Unicode, and the System z has hardware data compression. As a result, hardware data compression was used extensively.

SAP CSS Notes Applied
1704783: fix memory leak in detail log
1673941: memory leak when reading aggregated RDL results
1689312: KDV redundant selects on flow table
1670331: filter variants to reduce logging
1237311: fix full table scan when processing transactions (index on FPO characteristics)

8 Test Results

During the course of this effort many runs were performed. Some were to get familiar with the environment and the workload. Some were for debugging. It is beyond the scope of this paper to show them all. Listed below is a sub-set of data from what we believe is a consistent set of measurements. With the exception of the first measurement with 200,000 loans (see 18.1.1, "BA8 AFI 200K Mixed Loans FCL – FCL Fix + DS8800 Results" on page 15), all the other measurements ran the same 1.5M loans through the four process steps.

8.1 AFI / FCL Mixed Loans Results

The progression of this set of measurements is intended to show different factors and their effect on performance. The KPI for these measurements was the total elapsed time for the four process steps. With the exception of the first and the last two measurements, each measurement is built on the previous measurements.

A summary of the 1.5M loan measurements' overall elapsed times is shown in "Figure 2: BA8 AFI 1.5 Million Mixed Loans FCL Test Results Summary" below. The following table outlines the changes in the different runs and the environment (i.e., DB server or application server) affected by each change.

Measurement Run	DB	App
FCL Fix + DS8800	Y	Y
2x p780		Y
Member Cluster	Y	
AIX 7.1 + 16M Pages		Y
KDV 224 Jobs		Y
6 CP + 6 zIIP	Y	
5 CP + 5 zIIP	Y	

Table 3: Measurement Runs and Environment Affected

Subsequent sections show each measurement as well as some detailed measurement data on the last measurement.

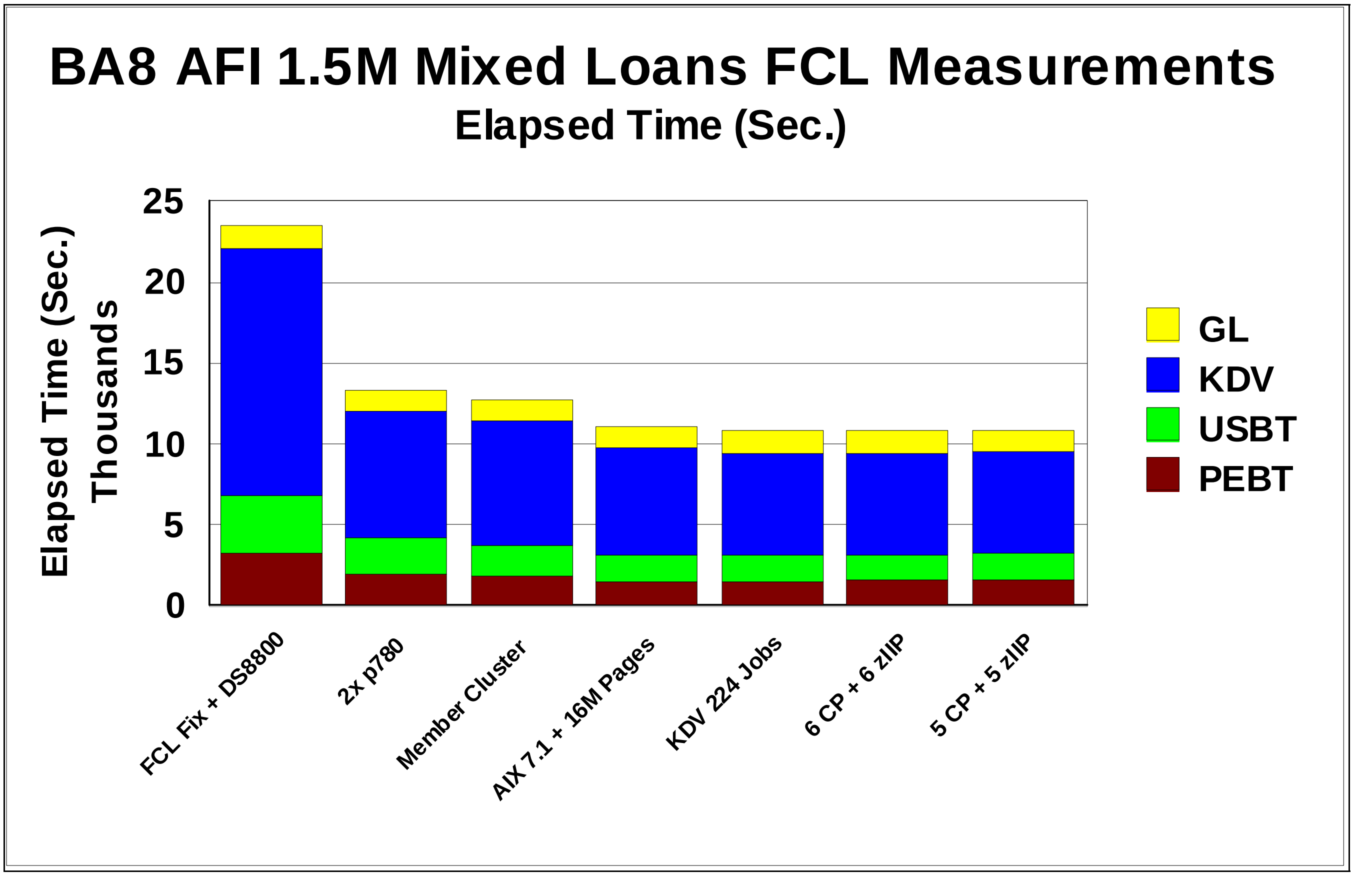

Figure 2: BA8 AFI 1.5 Million Mixed Loans FCL Test Results Summary

8.1.1 BA8 AFI 200K Mixed Loans FCL – FCL Fix + DS8800 Results

In the preparation for these tests, we were able to get a fix from SAP to improve the overall FCL performance. In addition, we obtained two DS8800 storage subsystems (see "Figure 1: Conceptual View of BA8 Test Environment Configuration" on page 11). This was a preliminary "short" run with only 200,000 loans.

Runid	Process Step	# of Batch Jobs	Elapsed Time (sec)	DB Server	App. Server	DB Server Util.	Apps. Server Util.	DB I/O/Sec. (Peak 15M Avg)	Comment
S2061 5A2	Post External (PEBT)	96	588	2817-715	p780-64 core	14%	58%	3,240	FCL fix DS8800
S2061 5A4	Update Secondary (USBT)	96	481	2817-715	p780-64 core	16%	63%	1,723	FCL fix DS8800
S2061 5A5	Key Date Valuation (KDV)	96	1845	2817-715	p780-64 core	4%	76%	961	FCL fix DS8800
S2061 5A6	GL Connector (prep)	48	19	2817-715	p780-64 core	15%	1%	24	FCL fix DS8800
	TOTAL (sec)		2933						
	TOTAL (hour)		**0.8**						

Table 4: BA8 AFI 200K Mixed Loans FCL – FCL Fix + DS8800 Results

8.1.2 BA8 AFI 1.5M Mixed Loans FCL – FCL Fix + DS8800 Results

This is basically the same experiment as the previous measurement – but with a full 1.5 million loans. This was effectively the start of this measurement sequence.

Runid	Process Step	# of Batch Jobs	Elapsed Time (sec)	DB Server	App. Server	DB Server Util.	Apps. Server Util.	DB I/O/Sec. (Peak 15M Avg)	Comment
S206-21A1	Post External (PEBT)	96	3,307	2817-715	p780-64 core	21%	69%	13,493	FCL fix DS8800
S206-23A1	Update Secondary (USBT)	96	3,518	2817-715	p780-64 core	21%	63%	9,858	FCL fix DS8800
S206-24A1	Key Date Valuation (KDV)	96	15,271	2817-715	p780-64 core	4%	76%	3,953	FCL fix DS8800
S206-24A2	GL Connector (prep)	48	1,391	2817-715	p780-64 core	12%	1%	11,053	FCL fix DS8800
	TOTAL (sec)		23,487						
	TOTAL (hour)		**6.5**						

Table 5: BA8 AFI 1.5M Mixed Loans FCL – FCL Fix + DS8800 Results

8.1.3 BA8 AFI 1.5M Mixed Loans FCL – 2x p780 Results

In this measurement, we added a second p780 application server.

Runid	Process Step	# of Batch Jobs	Elapsed Time (sec)	DB Server	App. Server	DB Server Util.	Apps. Server Util.	DB I/O/Sec. (Peak 15M Avg)	Comment
S207-10A1	Post External (PEBT)	128	1,997	2817-715	2x p780-64 core	36%	53%	17,692	2x p780
S207-11A2	Update Secondary (USBT)	128	2,219	2817-715	2x p780-64 core	32%	47%	12,607	2x p780
S207-11A3	Key Date Valuation (KDV)	128	7,770	2817-715	2x p780-64 core	8%	59%	7,293	2x p780
S207-11A4	GL Connector (prep)	48	1,371	2817-715	2x p780-64 core	13%	1%	9,995	2x p780
	TOTAL (sec)		13,357						
	TOTAL (hour)		**3.7**						

Table 6: BA8 AFI 1.5M Mixed Loans FCL – 2x p780 Results

8.1.4 BA8 AFI 1.5M Mixed Loans FCL – Member Cluster Results

In this measurement, we used a DB2 option called member cluster on heavily inserted DB2 tables to reduce contention on the space map pages.

Runid	Process Step	# of Batch Jobs	Elapsed Time (sec)	DB Server	App. Server	DB Server Util.	Apps. Server Util.	DB I/O/Sec. (Peak 15M Avg)	Comment
S207-28A1	Post External (PEBT)	192	1,874	2817-715	2x p780-64 core	41%	65%	21,933	member cluster
S207-30A2	Update Secondary (USBT)	192	1,813	2817-715	2x p780-64 core	44%	62%	19,289	member cluster
S207-30A3	Key Date Valuation (KDV)	192	7,695	2817-715	2x p780-64 core	8%	75%	7,041	member cluster
S207-30A4	GL Connector (prep)	48	1,319	2817-715	2x p780-64 core	13%	1%	9,544	member cluster
	TOTAL (sec)		12,701						
	TOTAL (hour)		**3.5**						

Table 7: BA8 AFI 1.5M Mixed Loans FCL – Member Cluster Results

8.1.5 BA8 AFI 1.5M Mixed Loans FCL – AIX 7.1 + 16M Page Results

In this measurement, we added to the application servers an up-level AIX kernel (7.1 TL1 SP4) as well as using 16 MB pages. Each application server allocated 8192 pages of 16 MB pages - 128 GB per machine. Each job used about 800 MB.

Runid	Process Step	# of Batch Jobs	Elapsed Time (sec)	DB Server	App. Server	DB Server Util.	Apps. Server Util.	DB I/O/Sec. (Peak 15M Avg)	Comment
S208-10A1	Post External (PEBT)	192	1,514	2817-715	2x p780-64 core	50%	59%	22,901	AIX 7.1 TL1 SP4 16M pages
S208-10A2	Update Secondary (USBT)	192	1,621	2817-715	2x p780-64 core	53%	58%	22,590	AIX 7.1 TL1 SP4 16M pages
S208-10A3	Key Date Valuation (KDV)	192	6,588	2817-715	2x p780-64 core	10%	75%	9,258	AIX 7.1 TL1 SP4 16M pages
S208-09A4	GL Connector (prep)	48	1,373	2817-715	2x p780-64 core	13%	1%	8,557	AIX 7.1 TL1 SP4 16M pages
	TOTAL (sec)		11,096						
	TOTAL (hour)		**3.1**						

Table 8: BA8 AFI 1.5M Mixed Loans FCL – AIX 7.1 + 16M Page Results

8.1.6 BA8 AFI 1.5M Mixed Loans FCL – KDV 224 Results

In this measurement, we changed the KDV step's number of batch jobs from 192 to 224.

Runid	Process Step	# of Batch Jobs	Elapsed Time (sec)	DB Server	App. Server	DB Server Util.	Apps. Server Util.	DB I/O/Sec. (Peak 15M Avg)	Comment
S208-10A1	Post External (PEBT)	192	1,514	2817-715	2x p780-64 core	50%	59%	22,901	KDV 224
S208-10A2	Update Secondary (USBT)	192	1,621	2817-715	2x p780-64 core	53%	58%	22,590	KDV 224
S208-11A3	Key Date Valuation (KDV)	224	6,297	2817-715	2x p780-64 core	11%	83%	9,888	KDV 224
S208-09A4	GL Connector (prep)	48	1,373	2817-715	2x p780-64 core	13%	1%	8,557	KDV 224
	TOTAL (sec)		10,805						
	TOTAL (hour)		**3.0**						

Table 9: BA8 AFI 1.5M Mixed Loans FCL – KDV 224 Results

8.1.7 BA8 AFI 1.5M Mixed Loans FCL – 6 CP + 6 zIIP Results

In this measurement, we started with the KDV 224 measurement in Section 8.1.6, "BA8 AFI 1.5M Mixed Loans FCL – KDV 224 Results" above but we changed the z196 SAP DB server to utilize six CP and six zIIPs.

Runid	Process Step	# of Batch Jobs	Elapsed Time (sec)	DB Server	App. Server	DB Server Util.	Apps. Server Util.	DB I/O/Sec. (Peak 15M Avg)	Comment
S208-14A2	Post External (PEBT)	192	1,603	2817-706	2x p780-64 core	62%	55%	21,352	6 CP + 6 zIIP
S208-15A1	Update Secondary (USBT)	192	1,525	2817-706	2x p780-64 core	64%	61%	17,655	6 CP + 6 zIIP
S208-16A1	Key Date Valuation (KDV)	224	6,339	2817-706	2x p780-64 core	15%	83%	9,597	6 CP + 6 zIIP
S208-15A3	GL Connector (prep)	48	1,360	2817-706	2x p780-64 core	19%	1%	10,674	6 CP + 6 zIIP
	TOTAL (sec)		10,827						
	TOTAL (hour)		**3.0**						

Table 10: BA8 AFI 1.5M Mixed Loans FCL – 6 CP + 6 zIIP Results

8.1.8 BA8 AFI 1.5M Mixed Loans FCL – 5 CP + 5 zIIP Results

In this measurement, we started with the KDV 224 measurement in Section 8.1.6, "BA8 AFI 1.5M Mixed Loans FCL – KDV 224 Results" above but we changed the z196 SAP DB server to utilize five CP and five zIIPs.

Runid	Process Step	# of Batch Jobs	Elapsed Time (sec)	DB Server	App. Server	DB Server Util.	Apps. Server Util.	DB I/O/Sec. (Peak 15M Avg)	Comment
S208-21A1	Post External (PEBT)	192	1,619	2817-705	2x p780-64 core	72%	54%	21,871	5 CP + 5 zIIP
S208-21A2	Update Secondary (USBT)	192	1,602	2817-705	2x p780-64 core	78%	58%	24,900	5 CP + 5 zIIP
S208-21A3	Key Date Valuation (KDV)	224	6,287	2817-705	2x p780-64 core	19%	83%	9,714	5 CP + 5 zIIP
S208-21A4	GL Connector (prep)	48	1,358	2817-705	2x p780-64 core	22%	1%	11,423	5 CP + 5 zIIP
	TOTAL (sec)		10,866						
	TOTAL (hour)		**3.0**						

Table 11: BA8 AFI 1.5M Mixed Loans FCL – 5 CP + 5 zIIP Results

8.2 Detailed Data for 5 CP + 5 zIIP

Listed here below are some more detailed measurement results for the measurement shown in 8.1.8, "BA8 AFI 1.5M Mixed Loans FCL – 5 CP + 5 zIIP Results" on page 18. Tables 12 and 13 show data from SAP.

Parent Batch: Process Step	Elapsed Time	App. Server CPU Time	DBREQ Time	RFC Time	ENQ Time	Avg. DB Bytes Added per Loan
PEBT	1,619	30	91	47	20	20,053
USBT	1,602	53	210	51	32	9,570
KDV	6,287	119	987	325	72	14,437
GL	1,358	6	106	1	3	8
Total (sec)	10,866					
Total (hr)	**3.0**					

Table 12: 5 CP + 5 zIIP Parent Batch Detailed Data

Here is an example of one of the batch jobs for each step. The rest of the batch jobs are similar.

Child Batch: Process Step	# of Child Batches	Elapsed Time	App. Server CPU Time	DBREQ Time	RFC Time	ENQ Time
PEBT	192	1,533	540	528	0	30
USBT	192	1,415	518	364	0	116
KDV	224	5,787	2,721	285	0	12
GL	48	1,352	0.3	1,352	0	0

Table 13: 5 CP + 5 zIIP Example of One Child Batch Detailed Data

This table shows data from both application servers. The data comes from VMSTAT.

Process Step	# of Child Batches	Avg. CPU Util.	30 Sec. Peak CPU Util.	30 Sec. Peak Active Memory (GB)	Peak IOWAIT
PEBT	192	54%	66%	113.9	0%
USBT	192	58%	61%	116.0	0%
KDV	224	83%	84%	247.4	0%
GL	48	1%	1%	191.3	0%

Table 14: 5 CP + 5 zIIP Application Server Detailed Data

This data shows data from the DB server. The data comes from RMF and DB2PE.

Process Step	CPU Util.	zIIP Util.	BPOOL Size	BPOOL Hit Ratio	15 Min. Peak I/O Rate (/sec)	15 Min. Peak Sync. Reads (/sec)	15 Min. Peak Prefetch Reads (/sec)	15 Min. Peak Async. Writes (/sec)	15 Min. Peak Log Writes (/sec)	zHPF
PEBT	69%	75%	84 GB	99.51%	21,871	4,144	100	8,753	6,020	83%
USBT	72%	83%	84 GB	94.42%	24,900	7,121	5,654	5,986	3,098	85%
KDV	15%	22%	84 GB	84.23%	9,714	3,987	2,911	1,465	388	89%
GL	19%	25%	84 GB	94.52%	11,423	8,423	877	1	9	98%

Table 15: 5 CP + 5 zIIP DB Server Detailed Data

Below are some graphs showing the processor utilizations and I/O rates over time of the different steps. We only plotted zIIP data for the last two steps. The y-axis scales are the same for all steps. However, this is not the case for the x-axis as it is a function of the elapsed time, which varied considerably depending on the step.

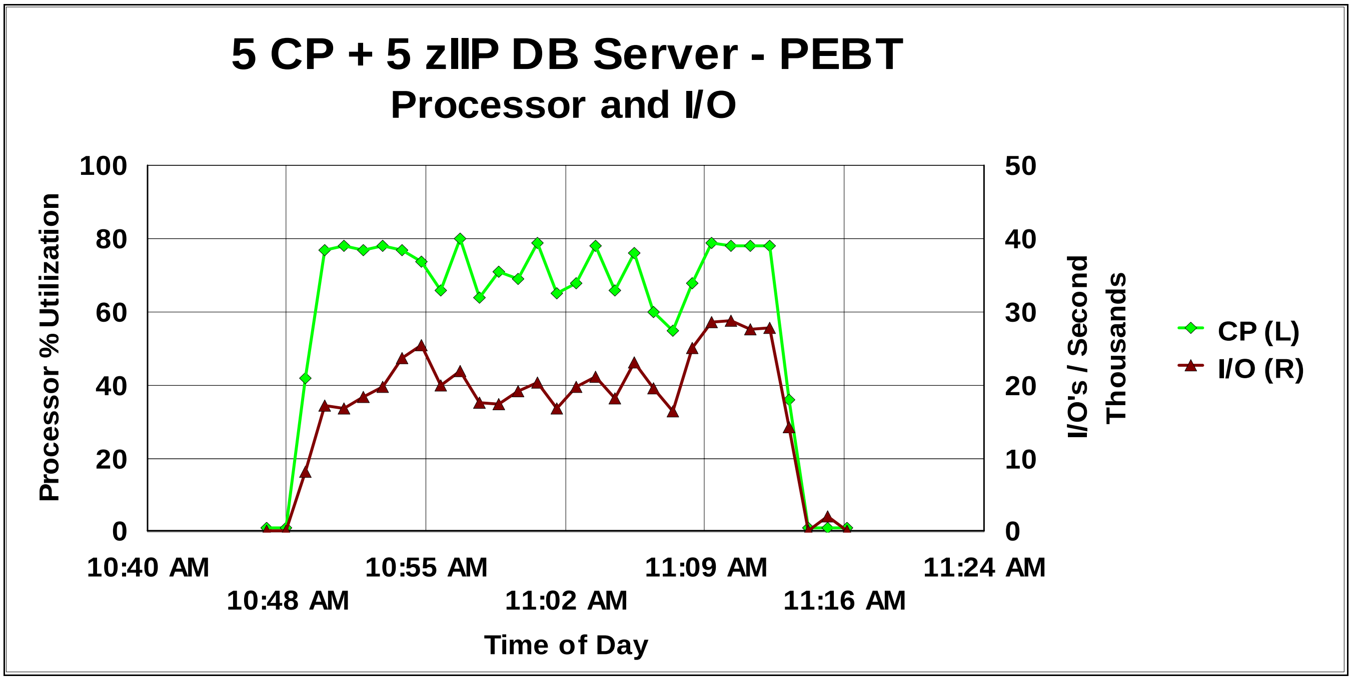

Figure 3: 5 CP + 5 zIIP DB Server Processor Utilization and I/O Rates – PEBT

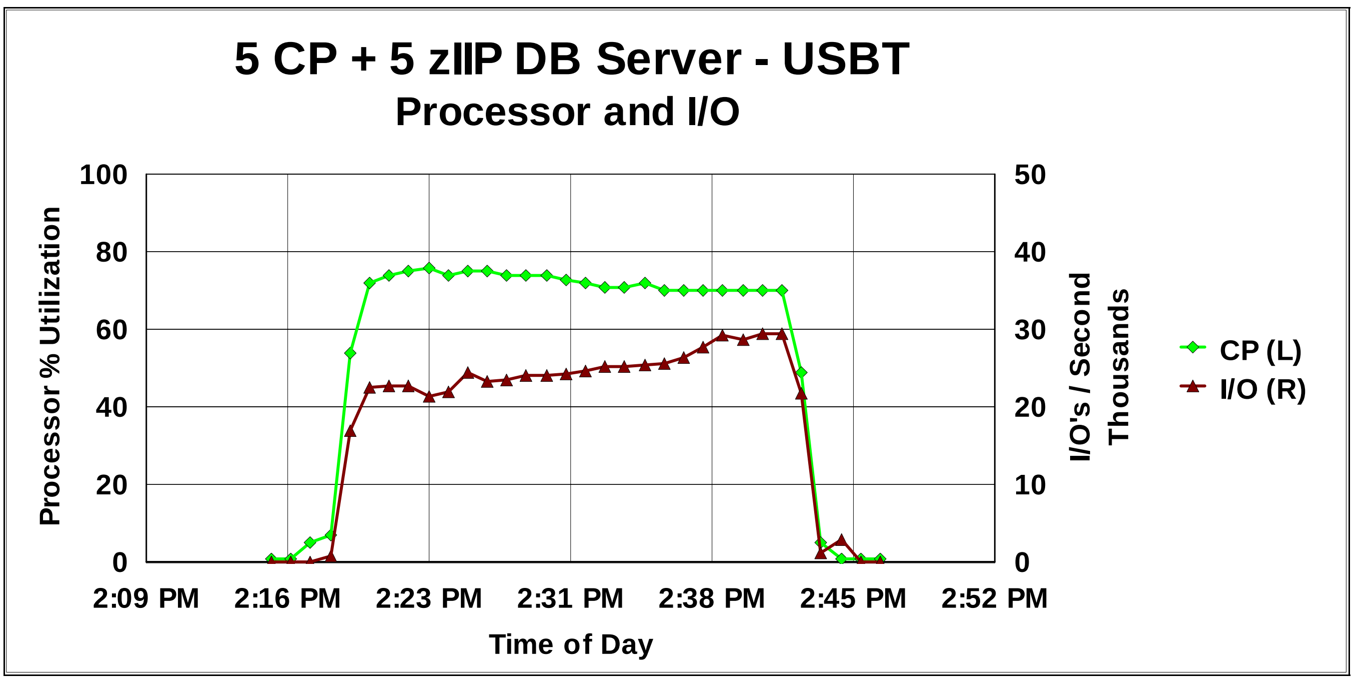

Figure 4: 5 CP + 5 zIIP DB Server Processor Utilization and I/O Rates - USBT

KDV had the longest elapsed time for any step. As a result, the x-axis is not to the same scale as other figures.

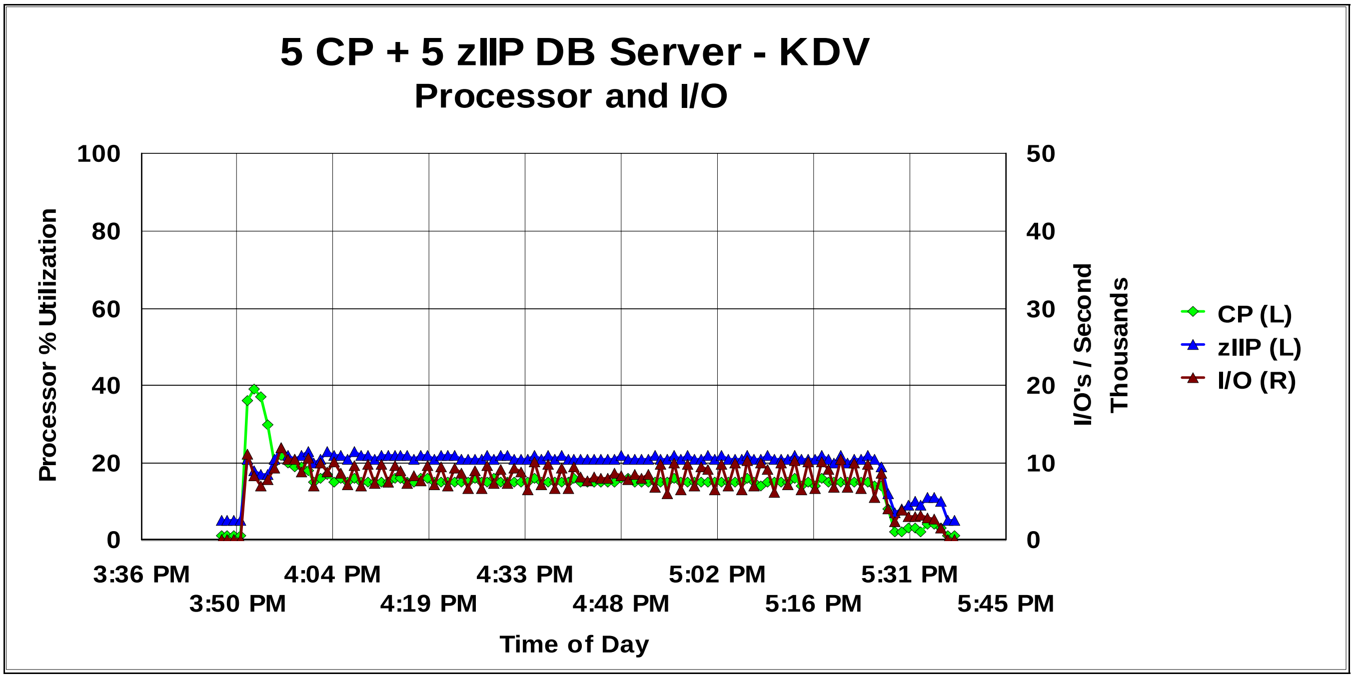

Figure 5: 5 CP + 5 zIIP DB Server Processor Utilization and I/O Rates - KDV

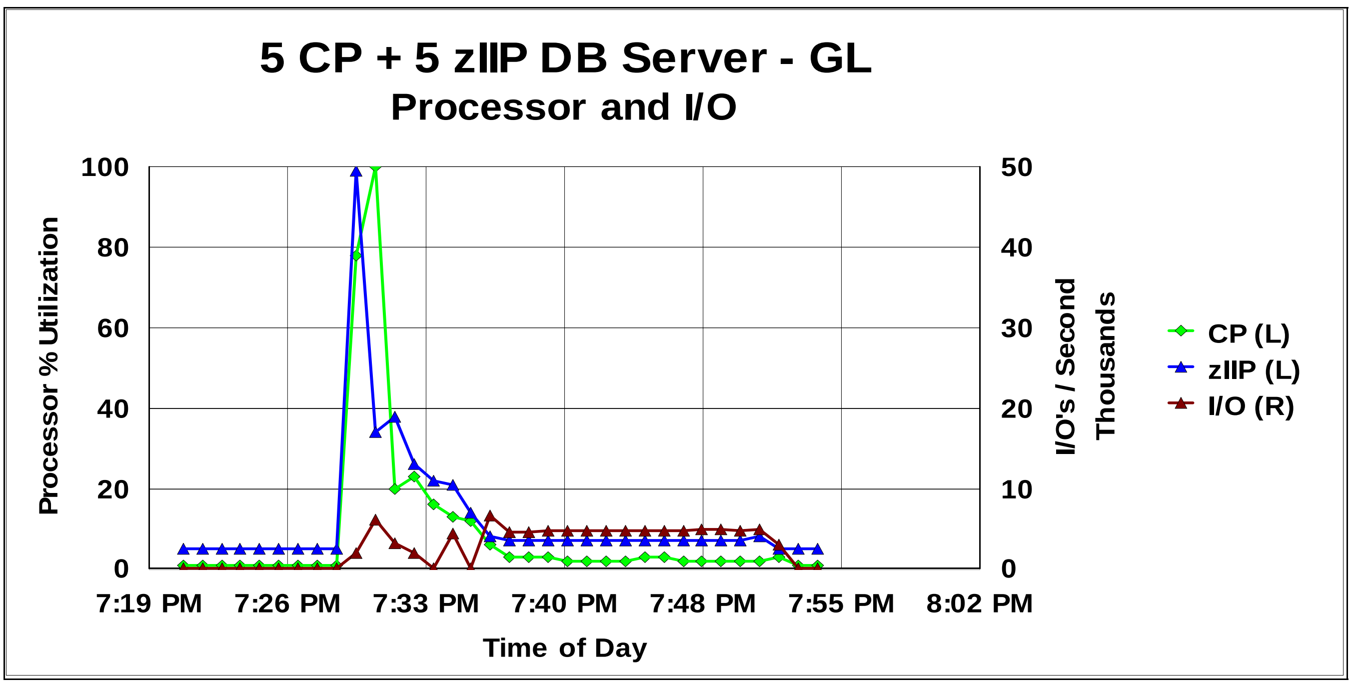

Figure 6: 5 CP + 5 zIIP DB Server Processor Utilization and I/O Rates - GL

9 Analysis

Once again, while a major goal of these measurements was to get as much throughput as possible in the limited time available, we did not resort to "benchmark specials". There were several significant differences between these measurements and the measurements we had done with BA7. Three main differences were:

- We measured a significantly different set of functions and workload.
- We significantly upgraded the robustness of the IBM infrastructure hardware and software.
- We used a new and enhanced BA application.

Discussed below are several analysis aspects.

Scaling From 200,000 to 1.5 Million Loans

If we compare the KPI elapsed time of the 200,000 loan run to its closest equivalent 1.5 M loan run, we see the elapsed time was 8.01 times the 200K run. This corresponds to 7.5 times as many loans processed as well as the databases' active data size. That is, the longer run took only 7% more time to process each loan. This indicates not only good scalability for this sequence of steps, but DB2's efficiency at managing a significantly larger database.

Test Evolution

In terms of the overall performance for these measurements, the primary factor was the application server. This is discussed in more detail later. This fact influenced our efforts – but not to the exclusion of everything else. We had a combination of application server and database server tuning experiments. This is illustrated by the items shown in "Table 3: Measurement Runs and Environment Affected" and "Figure 2: BA8 AFI 1.5 Million Mixed Loans FCL Test Results Summary" both on page 13. Here is a short summary of each measurement and its effectiveness.

- FCL Fix and DS8800: We started out with a robust test environment. As mentioned earlier, the FCL fix from SAP about doubled the application server performance. Our previous BA7 experience showed that BA could exhibit very intensive I/O behavior. Therefore, for these measurements we made sure we had two DS8800 high performance storage subsystems. See below for more discussion about the I/O profiles.
- 2x p780: Adding a second 64-core application server was the single biggest improvement in these measurements. All indications are that a third application server could have helped even more. See the application server discussion below.
- Member Cluster: Utilizing this DB2 feature provided some DB server improvement.
- AIX 7.1 and 16MB pages: The combination of improvements in AIX and its exploitation of 16 megabyte pages provided the second best improvement to the measurements. Section 8.1.5, "BA8 AFI 1.5M Mixed Loans FCL – AIX 7.1 + 16M Page Results" on page 17 describes how much memory was allocated for 16M pages.
- KDV 224 Jobs: Increasing the numbers of parallel jobs in this step from 192 to 224 was the final performance improvement. Because of schedules, we were not able to increase the other steps.
- 6 CP plus 6 zIIPs: Previous runs all used 15 CPs (i.e., a z196-715) for the DB server. In an attempt to test a smaller and more economical DB configuration, and demonstrate zIIPs, this measurement used a total of 12 cores (i.e., a z196-706 plus 6 zIIPs). Reducing the DB server cores had very little effect on the elapsed time.

- 5 CP plus 5 zIIPs: In our final attempt to test a smaller and even more economical DB configuration, and demonstrate zIIPs, this measurement used a total of 10 cores (i.e., a z196-705 plus 5 zIIPs). This measurement also had very little effect on the elapsed time. This measurement was chosen to show the most measurement data (see Section 8.2, "Detailed Data for 5 CP + 5 zIIP" on page 19).

Application Server[1]

The application server processing, especially for the KDV step, dominated this particular workload. For example, simply doubling the application servers (see Section 8.1.3, "BA8 AFI 1.5M Mixed Loans FCL – 2x p780 Results" on page 16 and "Figure 2: BA8 AFI 1.5 Million Mixed Loans FCL Test Results Summary" on page 14) reduced the elapsed time KPI by 43%.

Another way to look at this is to determine the CPU cores for each step normalized for utilization and elapsed time. When looking at all four steps combined, the application server cores were 96%, and the DB server cores only 4% of the total used CPU cores. Most other SAP workloads have a substantially higher DB server percentage.

A third indication of the application server dominance, especially with respect to the KDV step, is to look at some of the DB server detailed data in Section 8.2, "Detailed Data for 5 CP + 5 zIIP" on page 19 and following pages. In "Figure 5: 5 CP + 5 zIIP DB Server Processor Utilization and I/O Rates" on page 21, we can see that the DB server CPU cores are generally less than 23% utilized over this step's long elapsed time period. The DB server is waiting for application server work to finish. This is the reason we were able to reduce the DB server CPU cores from 15 to 10 with virtually no effect on elapsed time.

I/O Profile

Our previous tests with BA7 resulted in some extremely high I/O rates. With BA8, we had a much more powerful system that had significantly lower I/O rates. For example, with BA8, our peak 15 minute average rate was 24,900 I/Os per second. Earlier with BA7, we had an average of 65,854 I/Os per second **for more than 90 minutes**. Several factors contribute to this reduction in I/O rates with BA8.

A major contributor to this decrease was that we had more than twice the buffer pool space and, as a result, were able to do much better buffer pool tuning. Another contributor is that BA8 had a significant amount of efficient list pre-fetch I/O operations where more data is moved per operation. In contrast, BA7's I/O operations were dominated by random I/Os that each moved less data. We suspect that the list pre-fetch exploitation in these measurements was due to intrinsic workload differences as well as an improved DB2 optimizer.

These improvements in I/O rates not withstanding, BA8 still has significant I/O rates. Each step had a peak 15 minute I/O rate that, when normalized for utilization, exceeded our pre-install sizing rough rule of thumb peak I/O rate for this DB server. The step with the highest normalized rate was about twice our rough rule of thumb. Despite these still significant I/O rates, each step had excellent DASD response times. Each step's peak I/O rate interval had average DASD response times of 0.2 to 0.4 milliseconds.

[1] All of this analysis ignored the ENQ server – which had minimal utilization.

These excellent DASD response times had several contributing factors. One is the use of high-speed FICON Express8 adapters. Second, DB2 in an SAP environment is typically a good exploiter of zHPF – which improves the efficiency of I/O operations. BA8 continues that tradition. In our measurements, we saw 83% to 98% zHPF exploitation depending on the step. Third, list pre-fetch I/O operations, coupled with the recent support of multi-track operations in zHPF and the DS8800 system storage, further improved the efficiency of zHPF. Finally, the DS8800 system storage itself is designed to deliver good response time.

Customers considering BA should also consider all these factors to make sure they have a particularly robust I/O configuration for the DB server.

zIIP Exploitation

Analysis based on data shown in "Table 15: 5 CP + 5 zIIP DB Server Detailed Data" on page 19, indicates that the percent redirect to zIIPs for the entire LPAR ranged, depending on the step, from 52% to 59%. Contributing to this redirect is the frequency of list pre-fetch. This is probably a little higher than a typical production environment. This is because we did not run some of the typical functions customers use for the care and feeding of a real production database (e.g., Data Facility Hierarchical Storage Manager - DFHSM).

DB Growth

Another observation from our previous work with BA7 was significant DB growth running BA. BA8 continues to show this workload characteristic. For the four required steps we measured, we saw 43 k bytes of growth per processed loan. That amounts to 16% growth of the whole database. If we include the optional RDL pre-aggregation step, it was a total of 61 kB or 23% growth.

10 Conclusions

We saw three main workload characteristics in these BA measurements:

- Application server resources are key to overall elapsed time performance.
- BA has high I/O rates compared to our experiences with other SAP workloads.
- BA has high database growth compared to our experiences with other SAP workloads.

Should even more throughput be needed, there are further options. For example, adding more application servers to the current configuration could help. Similarly, DB scalability and its near continuous availability could be enhanced with a data sharing Parallel Sysplex.

As mentioned earlier, SAP's BA offers an extremely wide set of features and functions. For this set of tests, we used a narrow set of BA functions. As a result, we expect that there can be significant variances in behavior amongst BA implementations.

11 References

[1] IBM Corp. 2010. *IBM System z: SAP Bank Analyzer 7.0 Tests*
http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP101812 or msheets@us.ibm.com

[2] SAP AG 2012. SAP Documentation: Bank Analyzer (FS-BA)
http://help.sap.com/saphelp_banking80/helpdata/en/4e/1217c4e1905707e10000000a42189e/frameset.htm

[3] SAP AG 2011. *SAP for Banking on System z Reference Architecture*
http://www.sdn.sap.com/irj/sdn/db2?rid=/library/uuid/a00e4718-314f-2b10-19a6-a76f257addaf

[4] IBM Corp. 2010. *DB2 10 for z/OS* http://www.ibm.com/software/data/db2/zos/

[5] IBM Corp. 2010. *DB2 10 for z/OS Technical Overview*
http://www.redbooks.ibm.com/redbooks/pdfs/sg247892.pdf

[6] SAP AG 2010. *SAP on DB2 10 for z/OS – Being More Productive, Reducing Costs and Improving Performance* http://www.sdn.sap.com/irj/scn/index?rid=/library/uuid/005c6b33-aaf0-2d10-fcbb-b42e89ac5791

[7] IBM Corp. 2011. *DB2 10 for z/OS with SAP on IBM System z Performance*
http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP101845
or msheets@us.ibm.com

[8] IBM Corp. 2010. *IBM zEnterprise System* http://www.ibm.com/systems/z/hardware/zenterprise/

[9] IBM Corp. 2012. IBM Redbooks *DS8800 Performance Monitoring and Tuning*
http://www.redbooks.ibm.com/abstracts/sg248013.html