IBM Technical Brief

IBM System z®: SAP® Bank Analyzer 7.0 Tests

Authors:

Seewah Chan
Paul Lekkas
Veng K. Ly
Georg Mayer
Howard E. Poole
Michael R. Sheets

Document Owner:

Michael R. Sheets
IBM Americas Solutions Advanced Technical Skills

Version: 1.00
Date: December 6, 2010

Filename: System_z_SAP_BA.pdf

TABLE OF CONTENTS

TABLES

FIGURES

1 Introduction

IBM participates in several partnership programs with SAP AG. As part of this partnership, IBM and SAP work together to run different SAP applications. This benefits SAP because it broadens the exposure of their applications and personnel to different environments and gives them the opportunity to try new features and functions and improve their products. IBM products and personnel similarly benefit and it is part of IBM's continuous testing of its products.

This paper describes several tests we performed with IBM System z, DB2 for z/OS and SAP's Bank Analyzer (BA). BA is an integrated family of financial sub-products intended to provide bank management a consistent and prompt view of the institution's risk and return. It uses data from the institution's operations. BA supports several industry standards such as Integrated Finance and Risk Architecture (IFRA), International Accounting Standards (IAS), Basel II, Risk Adjusted Performance Measurement, and Sarbanes-Oxley [1]. In particular, we ran SAP's new Bank Analyzer 7.0 (BA7). We chose this combination because of previous, and continuing, experience with SAP's core banking applications [2, 3] as well as non-SAP banking and financial services experiences with customers.

The focus of our efforts and this paper were on the interactions of the application with the IBM supplied infrastructure. Infrastructure examples are servers, storage subsystem, operating systems, relational database and associated items. We did not focus on BA7 from a functional level.

These tests were not performance tests or formal benchmarks. While we did report several performance related metrics, these tests were more in the nature of function tests. There was very limited exposure of the application to the test environment. Sometimes there was just a single run. In addition, there were relatively low business volumes attempted. As a result, there was very limited tuning, investigation of bottlenecks, and investigation of how to scale up the volumes. Further, we ran on relatively old systems with old features and functions. For example, the DB Server's System z9 was two generations down level from today's zEnterprise System. This resulted in relatively slow adapters for networks and disks and no capability to use features like zHPF [4, 5]. We did run relatively recent versions of z/OS, DB2 for z/OS.

BA7 was new with several functions, features, and extensions added as the tests progressed. The project enabled SAP to enhance this application to run more efficiently with DB2 for z/OS.

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 rigorous performance measurements. The workloads changed and there was little, if any, chance to tune the workloads because of schedules. 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

SAP, R/3, mySAP, mySAP.com, xApps, xApp, 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.

AnyNet, CICS, DB2, DFSMSdfp, DFSMShsm, DFSMSrmm, DFSORT, DS6000, DS8000, ESCON, FICON, GDPS, HiperSockets, HyperSwap, IBM, IBM eServer, IBM logo, IMS, InfoPrint, Language Environment, Multiprise, NetView, OMEGAMON, Open Class, OS/390, Parallel Sysplex, PrintWay, RACF, RMF, S/390, Sysplex Timer, System Storage, System z, System z9, System z10, SystemPac, Tivoli, VSE/ESA, VTAM, WebSphere, z/Architecture, z/OS, z/VM, z/VSE, and zSeries are trademarks or registered trademarks of the International Business Machines Corporation in the United States and other countries.

Java and all Java-related trademarks and logos are trademarks of Sun Microsystems, Inc., in the United States and other countries

Linux is a trademark of Linus Torvalds in the United States and other countries.

Unicode is a trademark of Unicode, Inc.

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

Other products may be trademarks or registered trademarks of their respective 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 of IBM, and Michael Reuter, Corina Farcas, Hendrik Rusch, and Nabil Nouri, all of SAP AG.

6 Workload and Test Background

We tested four scenarios.

- Integrated Financial and Management Accounting (IFMA)
  - for loans
  - for accounts
- Accounting for Financial Instruments with aggregation (AFI) for accounts
- Imported Sub-ledger Documents (ISD) for loans

Each of these scenarios, whose results will be discussed in more detail in subsequent sections, consisted of several steps. In general, the final test for each scenario went through all the steps – but not necessarily in the order shown below.

As mentioned in Section 1, "Introduction", on page 4, we focused on application interactions with the infrastructure. We did not run SAP Standard Application Benchmarks (SAB). The Bank Analyzer has many more possible scenarios and process steps we did not test. These four scenarios were selected by SAP as being reasonably common and having broad test coverage. The tests were more in the nature of function or integration tests. Further, we did not tune these tests as one might for a benchmark or performance test. For example, because of the modest hardware environment, we did no virtual storage tuning.

Keeping in mind our focus on the infrastructure, and not the functional level, here are some short summaries of the functions. More details about the process steps tested in the scenarios can be found in Section 8, "Test Results" on page 12.

6.1 Integrated Financial and Management Accounting (IFMA)

Financial accounting provides information needed by external stakeholders. Management accounting provides similar information for the bank's management. This scenario runs them both together and provides a standard view of the financial position of the bank. There are separate scenarios for loans and accounts. One million of each were processed for each scenario.

The loan scenario has a term of five years, to be repaid in 60 monthly payments for each loan. The test simulates end of month processing consisting of two business transactions per loans. The transactions are principal repayment and interest payment. In total, there are two million business transactions processed for one million loans. These transactions are posted, evaluated, extracted to the general ledger, and balanced.

The current account scenario has three business transactions (withdrawal, deposit, and fee) per account, for end of month processing. Thus, in total, there are three million business transactions for one million accounts. These transactions are aggregated, posted, evaluated, extracted to general ledger and balanced. For accounts, there is a fee at month end.

Most of the IFMA process steps are AFI related (see Section 8.1, "Integrated Financial and Management Accounting (IFMA) Results", on page 12) – but without aggregation.

6.2 Accounting for Financial Instruments with Aggregation (AFI)

This scenario produces sub-ledger postings by applying valuation methods based on local Generally Accepted Accounting Principles (GAAP) accounting standards. These sub-ledger postings are aggregated to the granularity of the general ledger. This scenario does this by looking at ten million current accounts.

As mentioned in Section 6.1, "Integrated Financial and Management Accounting (IFMA)", above, there are AFI processes in the IFMA scenarios. This scenario differs because of its use of account aggregation for AFI. Loans AFI with aggregation is determined from IFMA for loans.

6.3 Imported Sub-Ledger Documents (ISD)

Postings for two million loans and their accrued interest are imported, general ledger accounts are derived, and posted in the sub-ledger. Finally, the sub-ledger documents are aggregated and transferred to the general ledger. Each loan generates one financial transaction. Each of these transactions has one document of two line items – one for principal and the other for interest.

7 Configurations

7.1 Hardware Environment

System z DB Server: Tests were performed on a single z9 Enterprise Class Model S38 with a total of 128 GB installed. The runs utilized one dedicated LPAR for z/OS DB2 9 with 16 processors and 64 GB. This represents about 10,500 SAPS using Business Suite 7 and Unicode.

Storage: IBM System Storage Server DS8700 Model 2421-941 with 128 HDDs and 128 GB cache. The IBM System Storage Server was attached to the z9 by eight long wave FICON Express4 connections. The capacity was about 26 TB of available storage capacity for database, logs and FlashCopy sets.

The BA 7.0 database used 128 3390 Mod 54 volumes - about 6.5 TB. It had 32 active logs of 4GB size each, striped across four volumes on four different ranks. The SAP database size had about two TB of allocated space, which contained approximately one TB of active pages in table spaces and one TB of active pages in index spaces.

Application Servers: We used a total of five application servers. One for the Central Instance, the others were regular application Servers. They were System p 550 POWER6 8204-E8A servers each with eight 5.0 GHz processor cores using SMT and 128 GB of memory. These five servers represent about 85,000 SAPS using BS7 and Unicode.

Network: Gigabit Ethernet networks were used for all connections. Each of the application servers was connected via a Gigabit Ethernet switch to the z9 via OSA-Express2 adapters. There were two cards each with one of their ports attached to the DB2 LPAR.

Below is a conceptual view of the configuration.

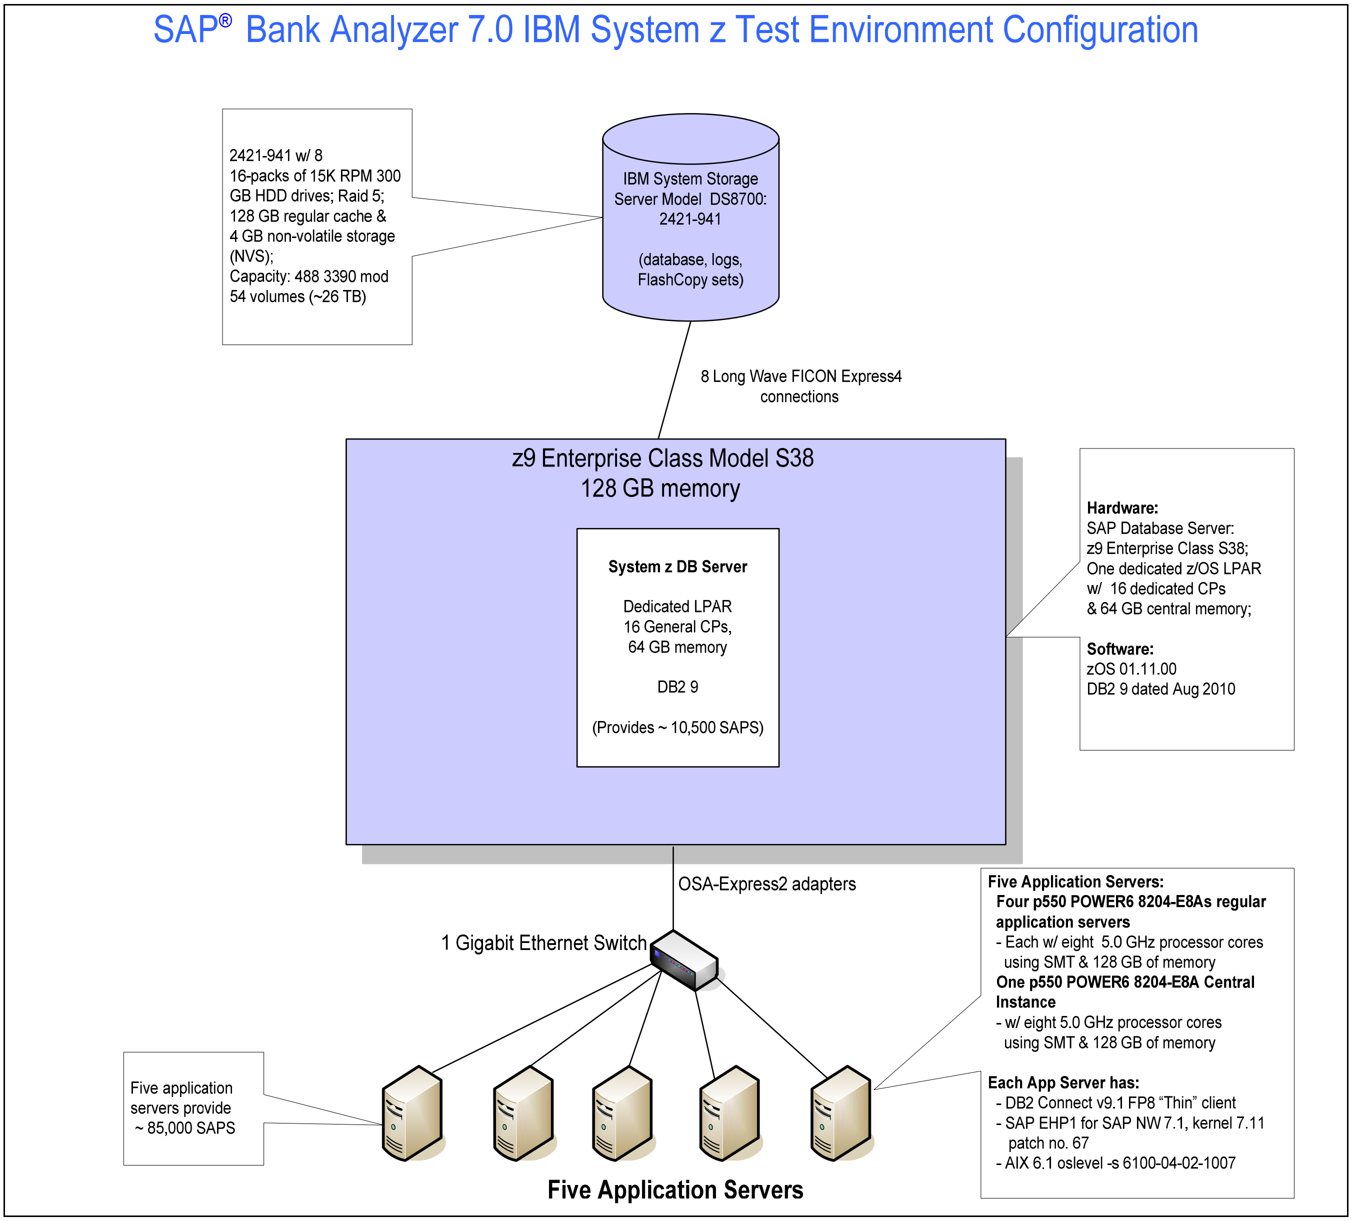

Figure 1: Conceptual View of BA7 Test Environment Configuration

7.2 Software Environment

z/OS

z/OS release 01.11.00 (R1.11)
DB2 9 dated August 2010

The bufferpool settings are listed below. There was no attempt made to tune the bufferpools.

BP name	PGFIX	VP Pages (K)	VPSEQT [%]	VPPSEQT [%]	VPXPSEQT [%]	DWQT [%]	VDWQT [%]	VDWQT [pages]	Page steal
BP0	YES	3,000	50	50	0	50	10	0	LRU
BP1	NO	120	99	50	0	50	10	0	LRU
BP2	NO	120	50	50	0	30	5	0	LRU
BP3	NO	400	40	50	0	30	5	0	LRU
BP4	NO	1,200	99	50	0	50	10	0	LRU
BP5	NO	120	99	50	0	50	10	0	LRU
BP6	YES	400	99	50	0	50	10	0	LRU
BP7	YES	400	99	50	0	50	10	0	LRU
BP8	YES	400	99	50	0	50	10	0	LRU
BP9	YES	400	99	50	0	50	10	0	LRU
BP10	YES	800	99	50	0	50	10	0	LRU
BP11	YES	800	99	50	0	50	10	0	LRU
BP40	NO	10	50	50	0	50	10	0	LRU
BP8K0	YES	800	50	50	0	30	5	0	LRU
BP8K1	NO	2	10	50	0	70	50	0	LRU
BP16K0	NO	9	50	50	0	30	5	0	LRU
BP16K1	NO	3	50	50	0	30	5	0	LRU
BP32K	NO	15	50	50	0	30	5	0	LRU
BP32K1	NO	15	50	50	0	30	5	0	LRU
BP32K2	NO	15	50	50	0	30	5	0	LRU
BP32K3	NO	15	99	50	0	50	5	0	LRU
BP32K4	NO	150	99	50	0	50	5	0	LRU
BP32K5	YES	30	50	50	0	30	5	0	LRU

Table 1: DB2 Bufferpool Settings

IBM DB2 Connect "Thin client" side: Driver for ODBC, CLI, JDBC and SQLJ - Version 9.1 FP8

AIX

AIX 6.1
oslevel -s
6100-04-02-1007

SAP Application Levels

SAP EHP1 for SAP NetWeaver 7.1, kernel release 711 patch no 67

Software Component	Release	Level	Highest Support Package	Short Description of Software Component
SAP_ABA	711	4	SAPKA71104	Cross-Application Component
SAP_BASIS	711	4	SAPKB71104	SAP Basis Component
PI_BASIS	711	4	SAPK-71104INPIBASIS	Basis Plug-In
ST-PI	2008_1_710	1	SAPKITLRE1	SAP Solution Tools Plug-In
SAP_BW	711	4	SAPKW71104	SAP Business Warehouse
FINBASIS	700	9	SAPK-70009INFINBASIS	Financial Basis
SEM-BW	700	9	SAPK-70009INSEMBW	SEM-BW 700: Add-On Installation
BI_CONT	711	0	-	Business Intelligence Content
FSAPPL	300	4	SAPK-30004INFSAPPL	SAP Banking Services
ST-A/PI	01L_BCO710	0	-	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 7.0 requires Unicode, and the System z has hardware data compression. However, because of the function test nature of this effort, data compression was used on only a few tables. If these tests were executed as a rigorous measurement, we would have made much more extensive use of data compression.

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. In general, these runs were not tuned as one might with a real benchmark. Each of these scenarios consists of many steps which may not individually be considered as complex. However, each step has its unique characteristics. Combining the sheer number of steps, and their uniqueness, makes the BA scenarios a very complex workload.

The tested scenarios have a heavy I/O profile with skewed hot spots across the various steps. The I/O profile is discussed more in Section 9.1, "I/O Rates" on page 17. Further, it is extremely tedious to tune and balance the I/O manually. Newer alternatives to help both the DB2 and I/O issues are discussed in Section 10, "Conclusions" on page 19.

Listed below are the test results selected as being the most useful, given the time constraints of this effort. Each set of tests will be discussed more in Section 9, "Analysis", on page 17.

For each process step of each scenario, we report run times, utilizations of the Application Servers, utilization of the DB Server, DB I/O rates, DB I/O rates normalized for processor utilization, and database growth. These last two metrics are discussed in more detail in Section 9, "Analysis", on page 17. The "Average" utilization metrics listed in the following tables are for sustained busy periods during the run. The time duration considered in this average corresponds to the runtime when the 96 jobs were running in parallel.

8.1 Integrated Financial and Management Accounting (IFMA) Results

As mentioned earlier, tests were run for both loans and accounts.

8.1.1 IFMA for Loans Results

Each step processed one million loans. Below are listed the scenario test results. The shaded steps are management accounting related. The rest are AFI related.

Process step	Runtime (sec)	Avg. App CPU	Avg. DB CPU	Avg. DB IO/Sec.	Normalized I/O Rate	DB Growth (GB)
1. Post external business transactions	3,417 (3,200 parallel)	52%	32%	15,223	47,572	10.77
2. Update secondary business transactions	6,356 (5,960 parallel)	96%	26%	14,752	56,738	21.40
3. Key date valuation for all financial positions	4,595 (3,840 parallel)	95%	28%	19,924	71,157	5.78
4. Update costing for period	2,727 (2,510 parallel)	92%	32%	9,504	29,700	0.66
5. Calculation for all Financial Positions for Period (PICC)	12,266 (7,320 parallel)	90%	28%	11,350	40,536	34.41
6. GL Connector	3,521 (3,511 parallel)	1%	25%	1,700	6,800	0.66
7. Balance processing step 1 (SV)	2,926 (2,680 parallel)	81%	46%	28,121	61,133	5.49
8. Balance processing step 2	3,076	85%	56%	25,769	46,016	16.96

Process step	Runtime (sec)	Avg. App CPU	Avg. DB CPU	Avg. DB IO/Sec.	Normalized I/O Rate	DB Growth (GB)
(ATP)	(3,050 parallel)					

Table 3: IFMA for Loans Test Results

8.1.2 IFMA for Accounts Results

Each step processed one million accounts. Below are listed the scenario test results. The shaded steps are management accounting related. The rest are AFI related.

Process step	Runtime (sec)	Avg. App CPU	Avg. DB CPU	Avg. DB IO/Sec.	Normalized I/O Rate	DB Growth (GB)
1. Post external business transactions	11,455 (10,138 parallel)	24%	18%	11,018	61,211	18.11
2. Update secondary business transactions	1,135 (918 parallel)	94%	21%	5,671	27,005	0.26
3. Key date valuation for all financial positions	11,907 (3,976 parallel)	82%	19%	6,952	36,589	5.91
4. Financial Positions for Period (PICC) event processing	116 (66 parallel)	5%	9%	1,411	15,678	0
5. Update costing for period	1,994 1,740 parallel)	81%	42%	9,776	23,276	1.56
6. Calculation for all PICC	13,609 (11,844 parallel)	81%	24%	9,251	38,546	52.51
7. GL Connector	3,349 (3,333 parallel)	1%	9%	3,110	34,556	0.02
8. Balance processing step 1 (SV)	60,984 (2,145 parallel)	83%	43%	18,891	43,933	4.90
9. Balance processing step 2 (ATP)	14,492 (2,857 parallel)	76%	50%	15,266	30,532	12.30

Table 4: IFMA for Accounts Test Results

8.2 Accounting for Financial Instruments with Aggregation (AFI) for Accounts Results

Each step processed 10 million current accounts. Below are listed the scenario test results. No utilization data were collected for the last step (number 9) because of its short runtime.

Process step	Runtime (Sec.)	Avg. App CPU	Avg. DB CPU	Avg. DB IO/Sec.	Normalized I/O Rate	DB Growth (GB)
1. Generate Positions	9,958 (9,281 parallel)	31%	40%	41,009	101,230	23.67
2. Aggregate Financial Transactions	2,333 (1,772 parallel)	17%	41%	29,999	73,168	6.23
3. Prepare Business Transactions for Aggregation	5,760 (5,565 parallel)	15%	55%	65,854	119,735	1.40
4. Process Reassignments for Business Transactions	758 (733 parallel)	32%	54%	16,957	31,402	1.30
5. Aggregate Business Transactions	7,561 (7,542 parallel)	2%	11%	9,946	90,418	7.01
6. Aggregate Current Accrual Results	633 (576 parallel)	70%	52%	12,374	23,360	0.15
7. Process Reassignments for Accrual results	487 (206 parallel)	42%	37%	24,026	64,935	0.46
8. Aggregate Retroactive	2,843	50%	72%	38,419	53,360	1.64

Process step	Runtime (Sec.)	Avg. App CPU	Avg. DB CPU	Avg. DB IO/Sec.	Normalized I/O Rate	DB Growth (GB)
Accrual Results	(1,023 parallel)					
9. Derive Application Events for source data aggregation	37					

Table 5: AFI for Accounts Test Results

In contrast to the other test scenarios, all the steps for AFI were done in a single day in the sequence shown above. As a result, we can show graphically the CPU utilization for one of the Application Servers (sapf66), the CPU utilization of the DB Server, and the DB Server I/O rates for each step. The numbers in the DB Server utilization graph correspond to the process step (i.e., the data row).

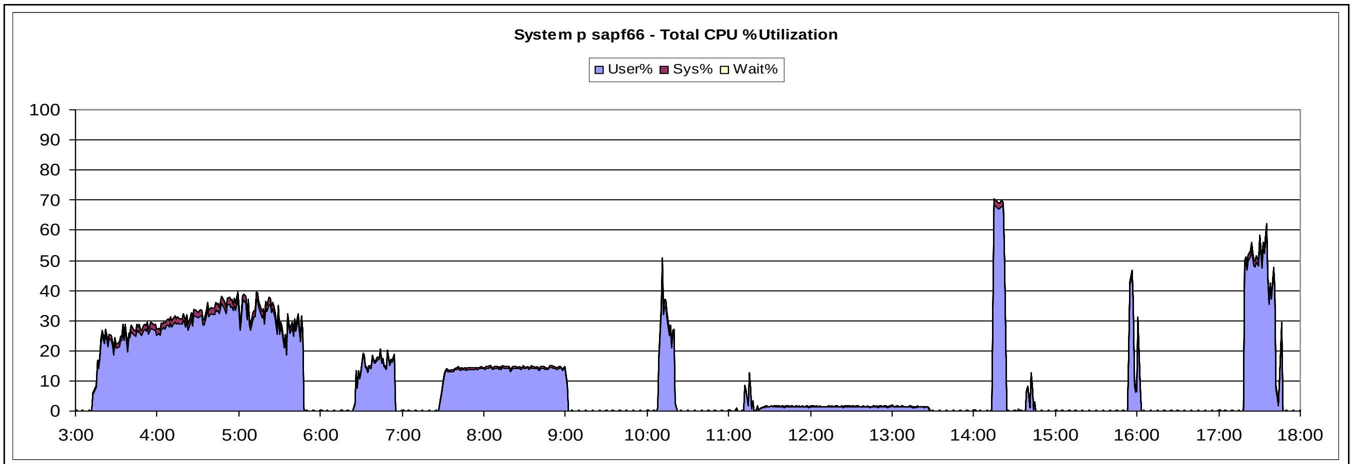

Figure 2: Application Server System sapf66 CPU Utilization for AFI

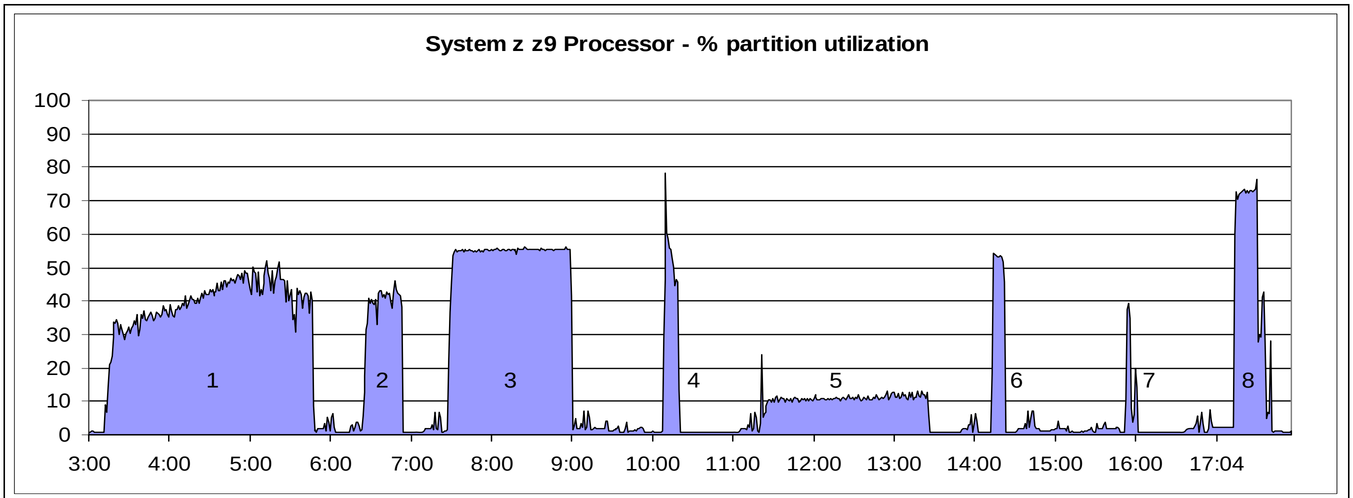

Figure 3: DB Server System Utilization for AFI

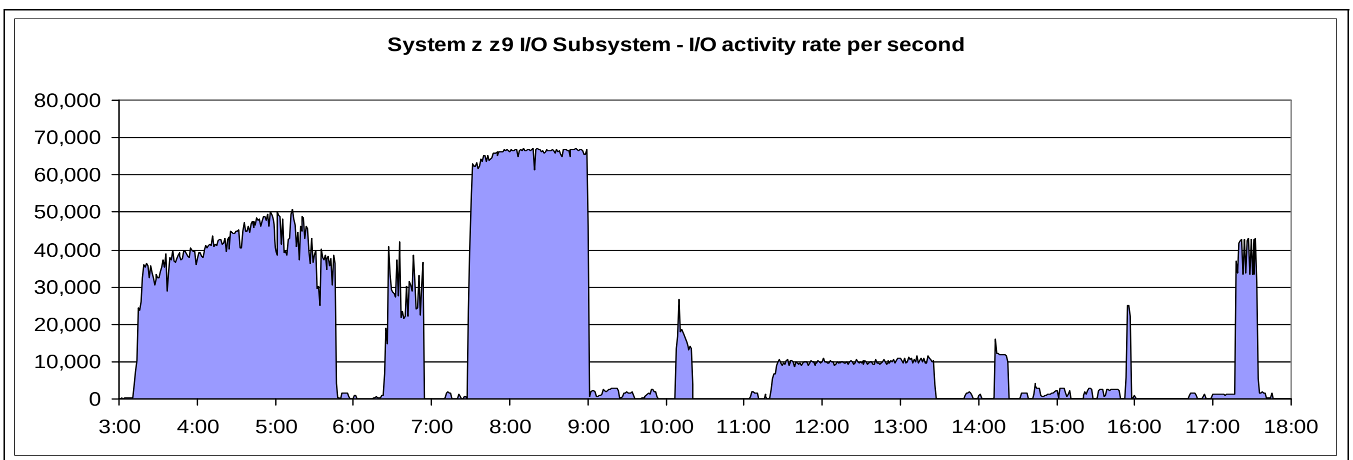

Figure 4: DB Server I/O rates for AFI

8.3 Imported Sub-Ledger Documents (ISD) for Loans Results

There was one step run in this test and it processed two million loans. Below are listed the scenario test results.

Process step	Runtime (Sec.)	Avg. App CPU	Avg. DB CPU	Avg. DB IO/Sec.	Normalized I/O Rate	DB Growth (GB)
1. ISD sub-ledger import	2,411 (2,390 parallel)	24%	91%	8,378	9,207	8.42

Table 6: ISD for Loans Test Results

9 Analysis

Given that these tests were essentially function tests, there were two characteristics about these BA7 tests that especially stood out. These were the I/O rates and the database growth. Each is discussed below.

9.1 I/O Rates

The I/O rates of the different process steps varied considerably. This is to be expected when looking at any application that involves many steps. Several of the process steps achieved extremely high I/O rates. These rates are not instantaneous peaks. The rates shown in Table 3, Table 4, Table 5, and Table 6 are averages over the entire duration. Some of the I/O intensive steps had long durations.

Another metric to look at is an I/O rate normalized for DB Server utilization – a normalized I/O rate. This metric is similar to the Large Systems Performance Report's [6] Internal Throughput Rate metric and can facilitate a comparison of different steps and workloads' I/O intensity. Further, by also adjusting for SAP DB Server Internal Throughput Rate Ratios, we can compare to other workloads run on different System z systems. The maximum normalized I/O rate for any step on this system was about 120,000 I/O per second. Our previous experience with SAP's core banking [3] had a maximum normalized rate of about 41,000 on this system. Even a specially modified version of this workload to stress an SSD storage subsystem [4] had only about 52,000 normalized I/Os per second. In contrast, for less stressful production systems we have a rough rule of thumb for this DB server of about 20,000 normalized I/O's per second. Even SAP's more conservative (e.g., it does not take compression into account) independently derived rule of thumb predicts about 47,000 I/O's per second.

While we did not instrument all the process steps, of those which we did (i.e., IFMA for Loans), we could see the vast majority of the I/O operations were random. Conversely, very few of these I/O operations were sequential. Other BA7 tests on other platforms have also seen this. The preponderance of random I/O's is a result of the data base design of BA7.

Of the SQL statements we looked at, they all seemed to be simple to medium in complexity. We saw no star schema.

As shown earlier, some of the buffer pools were rather large. The total RAM memory requirement of all the buffer pools is about 47 GB of the LPAR's allocation of 64 GB. Regardless, we did not do much tuning of them. However, given the randomness of the I/Os and the skews of the I/O rates from step to step, buffer pool tuning might not be very effective.

The normalized I/O rates are shown graphically below for the three most significant scenarios.

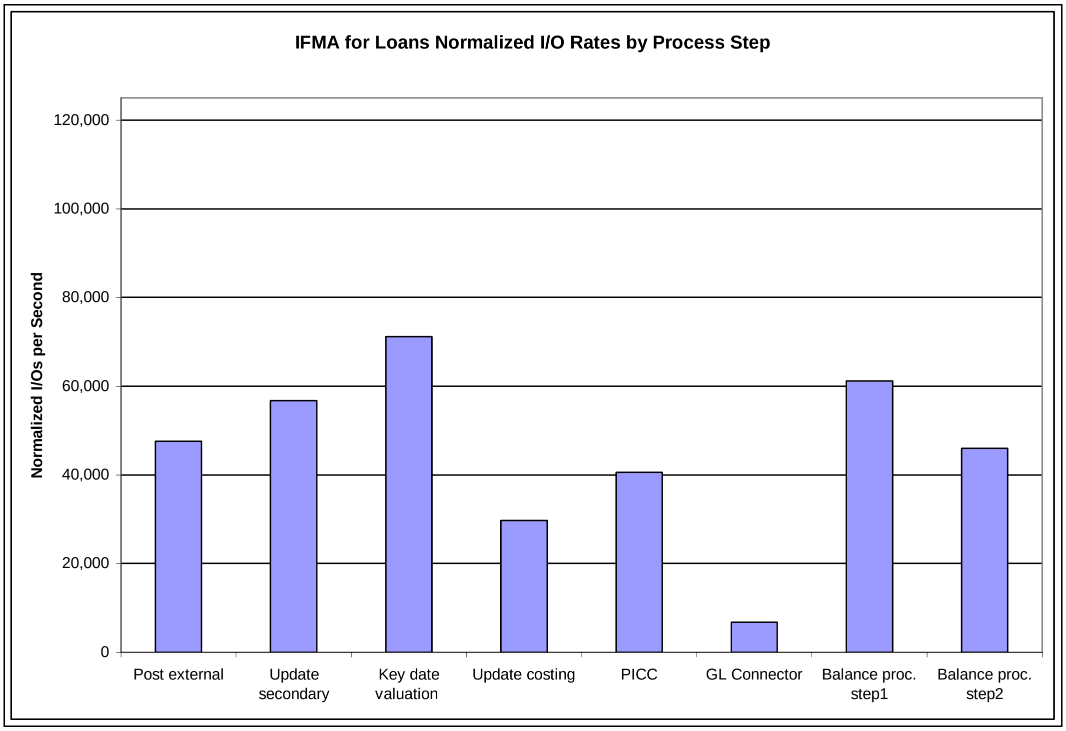

Figure 5: IFMA for Loans Normalized I/O Rates by Process Step

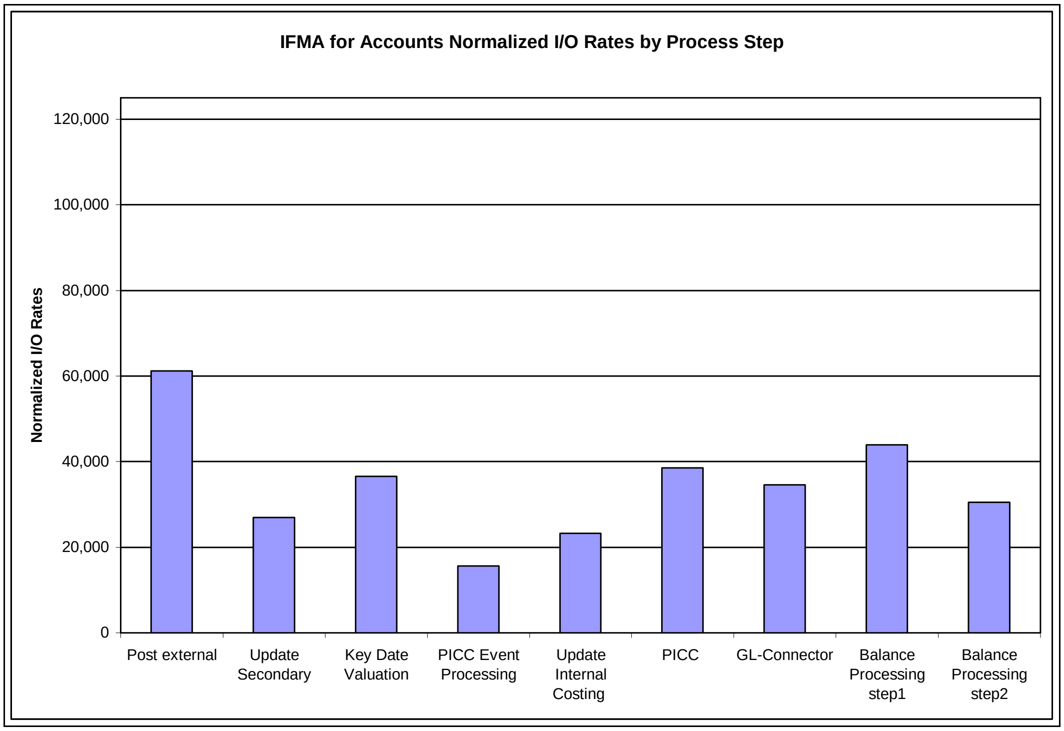

Figure 6: IFMA for Accounts Normalized I/O rates by Process

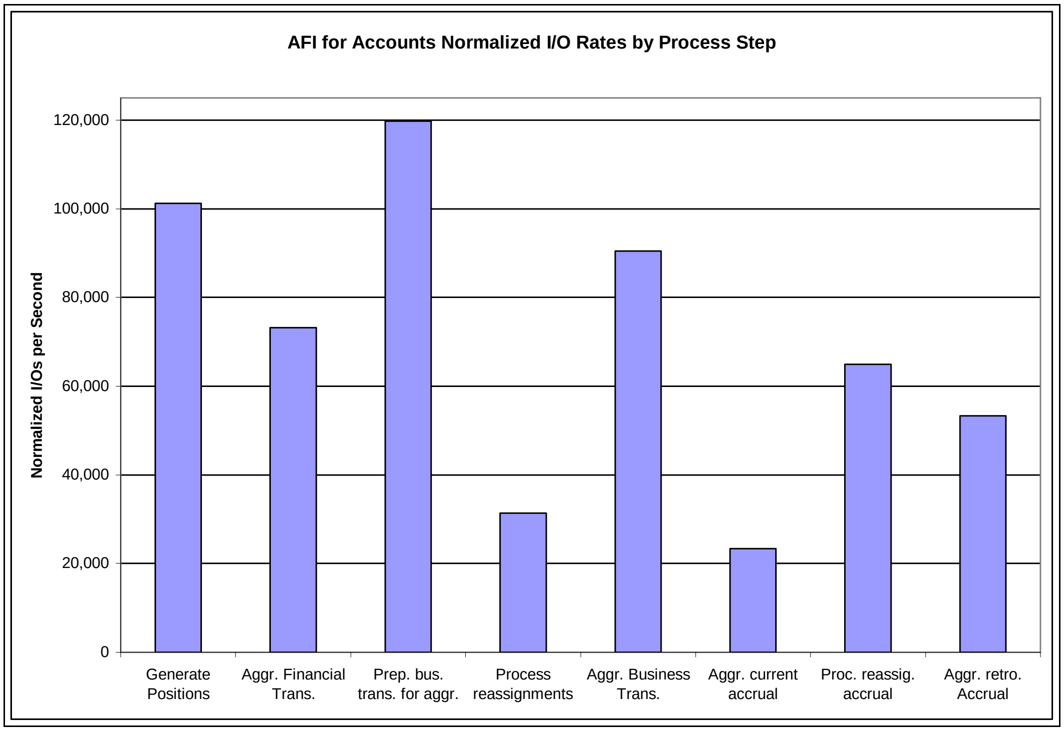

Figure 7: AFI for Accounts Normalized I/O Rates by Process Step

9.2 Data Base Growth

As with the I/O rates, some of the process steps caused significant database growth due to insertion of data. This metric is also listed in Table 3, Table 4, Table 5, and Table 6. In production environments, high database growth must be planned for – not only to have available space for this growth, but also to implement procedures for data archiving to keep the database size under control. As an example, the process step with the greatest growth was 52.51 GB in "Calculation for all PICC" in IFMA for Accounts. This single step could account for a 2.7% growth of our two TB of SAP DB.

Often database growth implies the occurrence of significant SQL INSERT processing by DB2. We found several cases of this in these tests. In fact, some of these cases were INSERTs to tables that had many secondary indexes. With DB2 9 this is problematic, resulting in processing being blocked while complex and lengthy disk operations were going on. This can limit application scalability.

10 Conclusions

As mentioned in Section 1, "Introduction" on page 4, these tests were not performance tests or formal benchmarks. These tests were more in the nature of function tests. Because of schedules, there was limited exposure of the application to the test environment. In addition, there were relatively low business volumes attempted. As a result, there was limited tuning or investigation of how to scale up the volumes.

All of that notwithstanding, overall BA7 behaved like an INSERT intensive On-Line Transaction Processing (OLTP) workload with significant random I/O.

There are several infrastructure product feature and function enhancements now available that could significantly improve the performance of these tests. The first is the zEnterprise [7]. It's z196 shows significantly better SAP DB Server performance than the z9 used for these tests. As well, the z196 supports significantly faster I/O capabilities such as OSA-Express3 for network operations and FICON Express8 and the associated High Performance FICON for System z (zHPF) for file operations. These improved I/O capabilities are an excellent match for the BA7 workload. Outside the zEnterprise, there are several new Solid State Disk (SSD) options now available for IBM's System Storage DS8000 systems. In the past [3], we have found SSD's unique performance characteristics to be a very good match for random I/O. An excellent example would be IBM's DS8800 with SSD and the Easy Tier self-tuning feature [8].

The recently delivered DB2 10 for z/OS [9] has several features and functions specifically implemented to improve SAP environments. Our BA7 tests could have benefited from DB2 10 too. For example, DB2 10's new parallel index update feature should help BA7's heavy INSERT onto tables with lots of secondary indexes. DB2 10's improved optimizer is expected to also help. Certainly, the DB2 10 virtual storage constraint relief will help all large SAP on z systems – including large BA7 systems.

11 References

[1] SAP AG 2010. SAP Bank Analyzer (FS-BA)
http://help.sap.com/saphelp_banking60/helpdata/en/f9/cdc23b95c8cb6fe10000000a114084/frameset.htm.

[2] IBM Corp. 2009. IBM System z® *and System Storage DS8000: Accelerating the SAP® Deposits Management Workload With Solid State Disks*
http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP101442 or msheets@us.ibm.com

[3] IBM Corp. 2008. *SAP® Transaction Banking: IBM System z® DB Server Using DB2 V9.1 for z/OS Large Database Measurements*
http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP101243 or msheets@us.ibm.com

[4] IBM Corp. 2008. US Announcement Letter 108-869: *IBM System Storage DS8000 series (Machine type 2107) delivers new functional capabilities (zHPF and RMZ resync) for System z environments*
http://www.ibm.com/common/ssi/cgi-bin/ssialias?infotype=AN&subtype=CA&htmlfid=897/ENUS108-869&appname=USN

[5] IBM Corp. 2009. *IBM System z10 I/O and High Performance FICON for System z Channel Performance* ftp://ftp.software.ibm.com/common/ssi/sa/wh/n/zsw03058usen/ZSW03058USEN.PDF

[6] IBM Corp. 2010. *Large Systems Performance Reference*
https://www.ibm.com/servers/resourcelink/lib03060.nsf/pages/lsprindexpdf/$file/SC28118714_20100714.pdf

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

[8] IBM Corp. 2010. *New IBM System Storage DS8800*
http://www.ibm.com/systems/storage/news/center/disk/enterprise/

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

