

IBM Technical Brief

**IBM zEnterprise System[®] and System Storage Server DS8870:
SAP[®] Bank Analyzer 8.0 AFI Current Accounts Tests**

Authors:

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

Document Owner:

**Veng K. Ly
SAP on System z Performance
IBM Poughkeepsie**

**Version: 1.00
Date: October 8, 2013**

Filename: System_z_SAP_BA8.CA.pdf

TABLE OF CONTENTS

| | | |
|------------|--|-----------|
| 1 | INTRODUCTION..... | 4 |
| 2 | REMINDEES | 5 |
| 3 | TRADEMARKS..... | 6 |
| 4 | FEEDBACK | 6 |
| 5 | ACKNOWLEDGEMENTS | 7 |
| 6 | WORKLOAD AND TEST BACKGROUND..... | 8 |
| 6.1 | General Tuning..... | 9 |
| 6.2 | Application Server Tuning | 9 |
| 6.3 | Database Tuning..... | 9 |
| 7 | CONFIGURATIONS | 13 |
| 7.1 | Hardware Environment..... | 13 |
| 7.2 | Software Environment | 15 |
| 8 | TEST RESULTS | 16 |
| 8.1 | AFI Current Account Results..... | 16 |
| 8.1.1 | BA8 AFI 1.5 Million Accounts – 32 Jobs with 8 CP Results..... | 17 |
| 8.1.2 | BA8 AFI 1.5 Million Accounts – 64 Jobs with 8 CP Results..... | 17 |
| 8.1.3 | BA8 AFI 1.5 Million Accounts – 96 Jobs with 8 CP Results..... | 18 |
| 8.1.4 | BA8 AFI 1.5 Million Accounts – 128 Jobs with 8 CP Results..... | 18 |
| 8.1.5 | BA8 AFI 1.5 Million Accounts – 128 Jobs with 6 CP Results..... | 19 |
| 8.1.6 | BA8 AFI 1.5 Million Accounts – 128 jobs with 3 CP + 3 zIIP Results | 19 |
| 8.2 | Detailed Data for 3 CP + 3 zIIP..... | 20 |
| 9 | ANALYSIS | 25 |
| 10 | CONCLUSIONS | 26 |
| 11 | REFERENCES..... | 27 |

FIGURES

| | |
|--|----|
| Figure 1: Conceptual View of BA8 Test Environment Configuration | 14 |
| Figure 2: 1.5 Million Accounts Test Results Summary | 16 |
| Figure 3: 3 CP + 3 zIIP PEBT DB Server Processor Utilization and I/O rates | 21 |
| Figure 4: 3 CP + 3 zIIP USBT DB Server Processor Utilization and I/O rates | 22 |
| Figure 5: 3 CP + 3 zIIP KDV DB Server Processor Utilization and I/O rates | 23 |
| Figure 6: 6 GL DB Server Processor Utilization and I/O rates | 24 |

TABLES

| | |
|--|----|
| Table 1: DB2 Buffer Pool Settings..... | 10 |
| Table 2: SAP Application Levels | 15 |
| Table 3: BA8 AFI 1.5 Million Accounts – 32 Jobs with 8 CP Results | 17 |
| Table 4: BA8 AFI 1.5 Million Accounts – 64 Jobs with 8 CP Results | 17 |
| Table 5: BA8 AFI 1.5 Million Accounts -- 96 Jobs with 8 CP Results | 18 |
| Table 6: BA8 AFI 1.5 Million Accounts – 128 jobs with 8 CP results | 18 |
| Table 7: BA8 AFI 1.5 Million Accounts – 128 Jobs with 6 CP Results | 19 |
| Table 8: BA8 AFI 1.5 Million Accounts – 128 jobs with 3CP + 3 zIIP Results | 19 |
| Table 9: 3 CP + 3 zIIP Parent Batch Detailed Data..... | 20 |
| Table 10: 3 CP + 3 zIIP Child Batch Detailed Data..... | 20 |
| Table 11: 3 CP + 3 zIIP Application Server Detailed Data | 20 |
| Table 12: 3 CP + 3 zIIP DB Server Detailed Data | 20 |

1 Introduction

This paper documents some of the tests that we did with the AFI Current Account scenario of 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 8.0 AFI Loans [1] and Bank Analyzer 7.0 (BA7) [2]. 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 [3]. 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* [4].

Our test environment was DB2 10 for z/OS [5, 6, 7, 8] and z/OS 1.13. We ran on the latest zEnterprise System zEC12 [9] and System Storage Server DS8870 [10,11,12]. 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](http://www.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 vengly@us.ibm.com

5 Acknowledgements

The authors would like to recognize the contributions of:

- Andrea Fuga (IBM System z programmer in Poughkeepsie)
- Tom Litty (IBM System z programmer in Poughkeepsie)
- Allen Wright (IBM Disk Products Program Management in Tucson)
- Ralf Knoebl (SAP Project Manager, Product Management Analytical Banking)
- Knut Willems (SAP Performance Group)
- Corina Farcas (SAP Global Banking, Active Global Support)

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 accounts processing to construct a workload. This was a different workload than we had for our earlier BA8 tests [1], which focused on various types of loan scenarios. For these tests, we focused on AFI current accounts scenario, without source data layer aggregation (SDA). For scenarios with very large volumes of accounts, SDA can improve elapsed time, and it should be considered.

We used generated test data, with every account having a separate business partner (a 1:1 relationship). On an average working day, for each account, one business transaction was loaded and posted.

By default, there was no daily key date valuation. Most time-critical processing was for month end processing. At the end of the month, each account was processed with four business transactions (withdrawal, deposit, interest rate on deposit, and fee). For each account, one accrual was loaded into Results Data Layer (RDL).

The processing was done via SAP batch as a chain of process steps. Typically, this chain has four logical steps:

- **Post External Business Transactions (PEBT):** External business transactions are generally from a feeder system. However, in our case, the test data was created within our BA system. The data is read and converted from the Source Data Layer (SDL) format to the specified accounting format, before transferring to the specified accounting system. Examples of the different accounting systems/formats are the International Accounting Standard (IAS) or German Accounting Standard Handelsgesetzbuch (HGB). Then the data is posted initially to a position clearing account in financial position management.
- **Update Secondary Business Transactions (USBT):** The position clearing account is cleared when the secondary business transactions are updated.
- **Key Date Validation (KDV):** The valuation of a position for financial instruments is based on a key date, and goes through multiple calculation steps. Possible step types include amortization, fair value calculation, and distribution of valuation reserves across time, and deferred taxes. Customers can define their own types of key date valuations. An example could be changing a tax rate.
- **General Ledger Connector Document Preparation (GL):** To create general ledger documents, required data are extracted from the results data layer. This function checks the consistency of the general ledger documents and saves them in the persistence layer.

As a result, for these tests, our key performance indicator (KPI) metric was the total elapsed time of PEBT, USBT, KDV, and GL. 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”.

6.1 General Tuning

General tuning was similar to our previous measurements on Bank Analyzer 8.0 AFI loan scenario since we used almost the same test environment.

- NetWeaver 7.1 EHP1 to exploit 64-bit addressability with AIX 7.1.
- Optimized Latency Mode (OLM) option of the OSA-Express4S adapter to improve the elapsed time for application server to DB server communication.
- HyperPAV for the DS8870 to reduce disk I/O queuing.
- High Performance FICON for System z (zHPF) w/ multi-track support to improve efficiency of I/O resources.
- DB2 10 and System z hardware and software features including Index I/O Parallelism for Insert, Safe Query Optimization, Large 1 MB Page Support.
- DB2 striped logs.
- MAXKEEPD=64K to minimize the number of prepares.

6.2 Application Server Tuning

- Adjusted the number of SAP instances to eight per Power 780 to increase the SAP memory pools to accommodate more parallel jobs.
- Adjusted the number of parallel batch jobs based on application server CPU consumption of the workload and the available processor capacity. This had a profound effect on the utilization of both the application server and the DB server.
- Installed AIX 7.1 to take advantage of improvements in AIX as well as its exploitation of 16 MB pages.

6.3 Database Tuning

- Identified “hot” tables using SAP transaction st04, and spread out these tables/indexes on DASD.
- Isolated these tables/indexes to unique DB2 buffer pools.
- Used larger DB2 buffer pools.
- Used member cluster for heavily inserted tables.
- 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 | 12,000 | 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 |
| 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

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 /IBA/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 | Strucname | 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 | |
| /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_OIXCH701 | T | 701 | SBA | S_KFCBD2 | | BP12 | Yes |
| /1BA/HM_OIXCH701~0 | I | | | | | BP13 | |
| /1BA/HM_OIXCH70150 | I | | | | | BP13 | |
| | | | | | | | |
| /1BA/HM_OIXCV701 | T | 701 | SBA | S_KFCBD2 | | BP12 | No |
| /1BA/HM_OIXCV701~0 | I | | | | | BP13 | |
| /1BA/HM_OIXCV70150 | I | | | | | BP13 | |
| /1BA/HM_OIXCV701VG | I | | | | | BP13 | |
| | | | | | | | |
| /1BA/HM_OIXC_701 | T | 701 | SBA | S_KFCBD2 | | BP12 | No |
| /1BA/HM_OIXC_701Z1 | I | | | | | BP13 | |
| /1BA/HM_OIXC_701~0 | I | | | | | BP13 | |
| | | | | | | | |
| /1BA/HM_1J3IH701 | T | 701 | SBA | S_FPCBD2 | | BP14 | Yes |
| /1BA/HM_1J3IH701~0 | I | | | | | BP15 | |



| Table or Index Name | Table or Index | Client | Area | Object Name | Strucname | BP | Member Cluster |
|---------------------|----------------|--------|------|-------------|--------------------------|------|----------------|
| /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 | |
| /1SGS/4F0AD302 | T | | | | _R0_SBA_____SF_ _____ | BP4 | Yes |
| /1SGS/4F0AD302~0 | I | | | | | BP5 | |
| /1SGS/4F0AD302~DEF | I | | | | | BP5 | |

7 Configurations

7.1 *Hardware Environment*

System z DB Server: Tests were performed on a single zEC12. The runs utilized one dedicated LPAR for z/OS DB2 10 with up to 8 processors and 235 GB of real storage.

Storage: One dual frame IBM System Storage Server DS8870 Model 2421-961/96E with 768 HDDs and 256 GB cache. The IBM System Storage Servers were attached to the zEC12 by sixteen long wave FICON Express8S connections. The capacity was about 164 TB of available storage capacity for database, logs, and FlashCopy sets.

The BA 8.0 database used 103 3390 Mod 54 volumes - about 5.6 TB. It had 32 active logs of 4GB size each, striped across four volumes on four different ranks.

Application Servers: One IBM Power 780 system with 64 3.9 GHz cores utilizing SMT and 512 GB memory. It had an AIX LPAR with eight SAP instances.

Network: A dedicated 10 Gb Ethernet network was used for all connections. The application server was connected via two 10 Gb Ethernet adapters through a 10 Gb Ethernet switch to the zEC12 via two OSA-Express4S adapters.

Below is a conceptual view of the test hardware and software configuration.

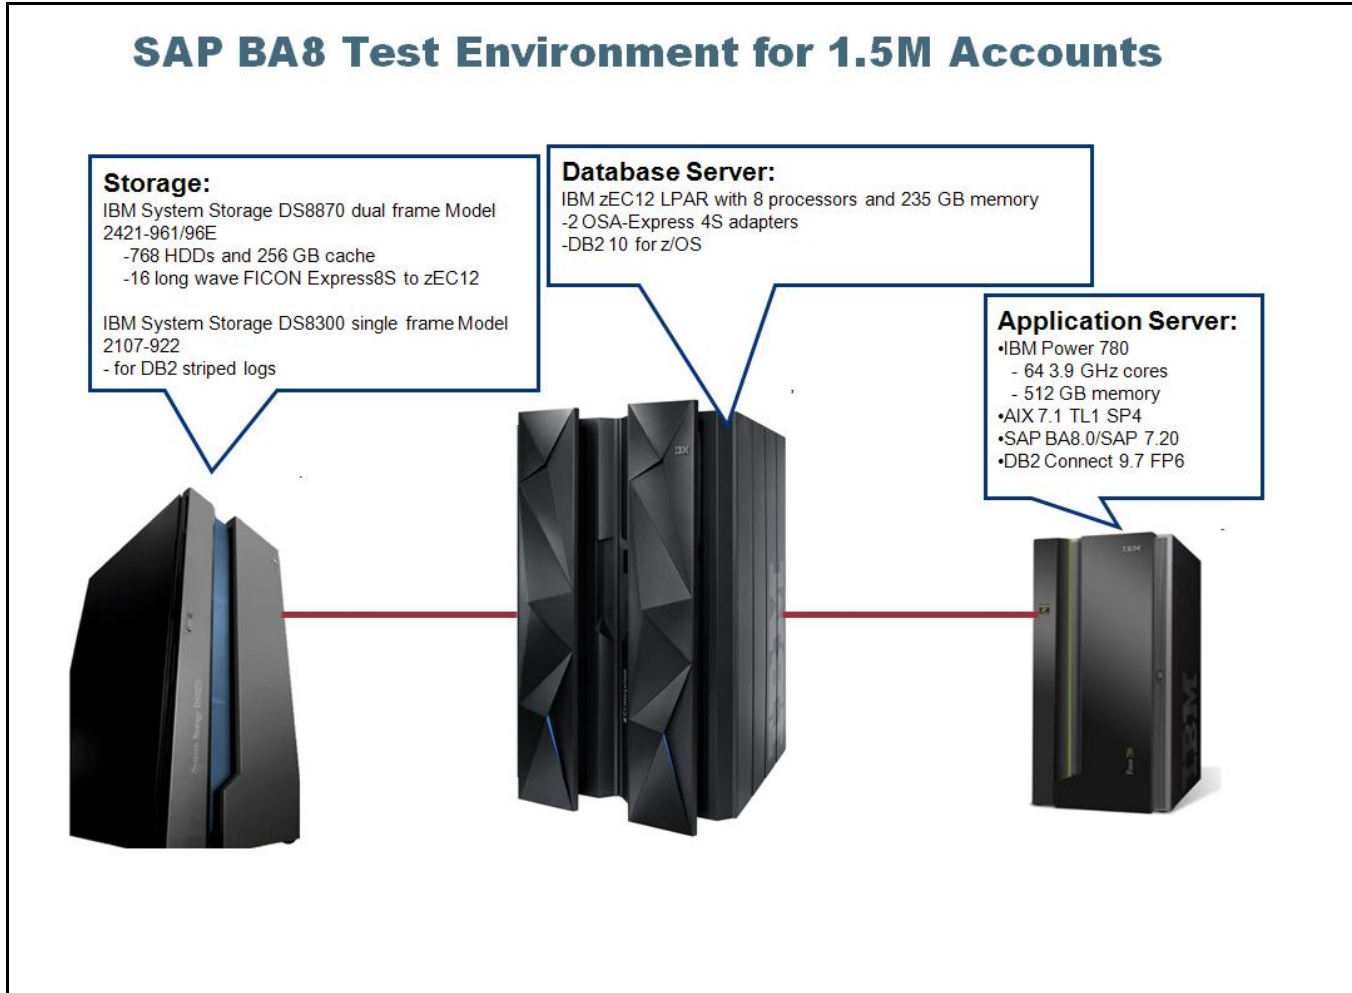


Figure 1: Conceptual View of BA8 Test Environment Configuration

7.2 Software Environment

z/OS

z/OS release 01.13.00 (R1.13)

DB2 10 dated January 2012

AIX

AIX 7.1 TL1 SP4

SAP Kernel Information

SAP Kernel : 720

Kernel Patch number : 300

DBSL Patch No. : 221

ICU Version : 4.0.1 Unicode Version 5.1

libsap16 Version : 1.0025 Jul 6 2012 19:22:32

IBM DB2 Connect “Thin client” side: Driver for ODBC, CLI, JDBC and SQLJ - Version 9.7 FP6

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 | 0010 | SAPKA71107 | Cross-Application Component |
| SAP_BASIS | 711 | 0010 | SAPKB71107 | SAP Basis Component |
| PI_BASIS | 711 | 0010 | SAPK-71107INPIBASIS | Basis Plug-In |
| ST-PI | 2008_1_710 | 0007 | SAPKITLRE4 | SAP Solution Tools Plug-In |
| SAP_BW | 711 | 0010 | SAPKW71107 | SAP Business Warehouse |
| FINBASIS | 700 | 0019 | SAPK-70016INFINBASIS | Financial Basis |
| SEM-BW | 700 | 0019 | SAPK-70009INSEMBW | SEM-BW 700: Add-On Installation |
| BI_CONT | 711 | 0002 | SAPK-71102INBICONT | Business Intelligence Content |
| FSAPPL | 400 | 0004 | 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.

8.1 AFI Current Account Results

Subsequent sections show each measurement as well as some detailed measurement data. The following figure shows the summary of results of six measurement points, scaling from 32 to 128 parallel jobs, and the demonstration of utilizing zIIP, and optimizing the DB server configuration.

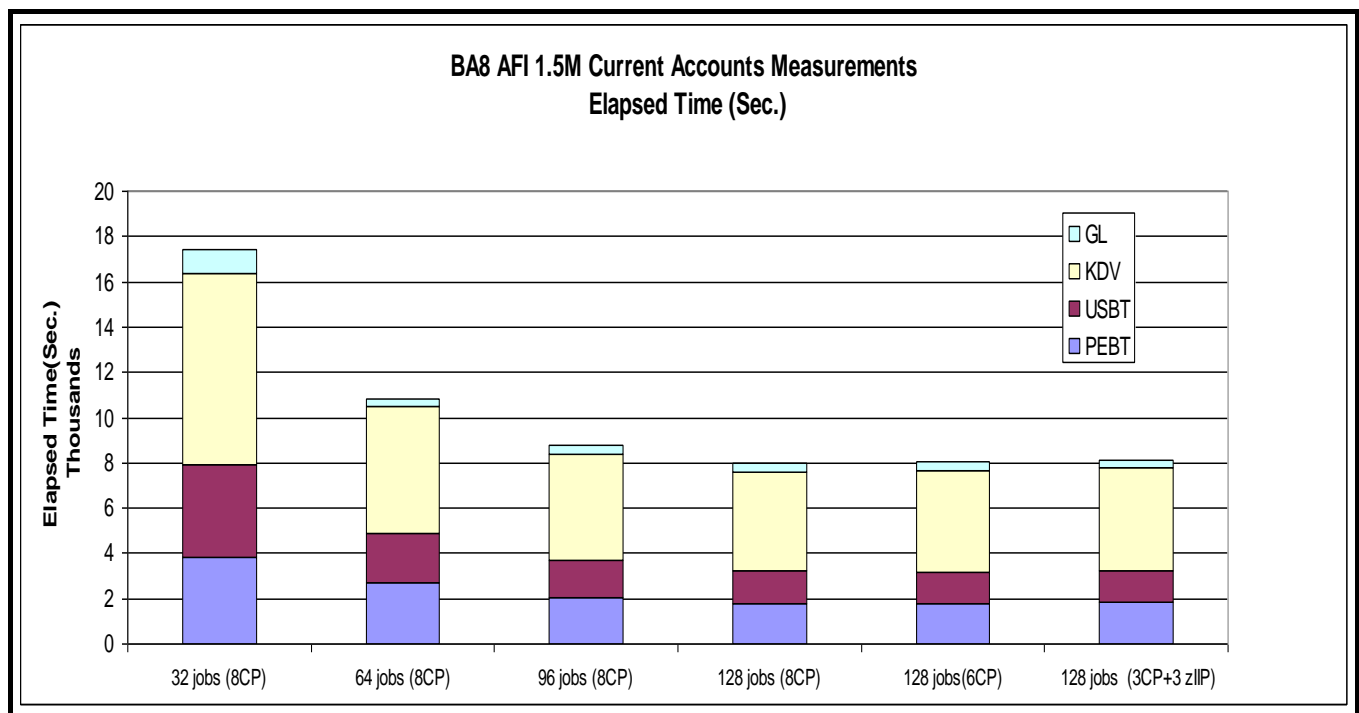


Figure 2: 1.5 Million Accounts Test Results Summary

8.1.1 BA8 AFI 1.5 Million Accounts – 32 Jobs with 8 CP Results

We started the first sequence of measurements with a zEC12 SAP DB server having 8 general processors (CPs) and 32 parallel batch jobs.

| Runid | Process Step | # of Batch Jobs | Elapsed Time (sec) | DB Server | App. Server | DB Server Util. | App. Server Util. | DB I/O/Sec. (Peak 5M Avg) |
|----------|--------------------------|-----------------|--------------------|-----------|--------------|-----------------|-------------------|---------------------------|
| S30503A1 | Post External (PEBT) | 32 | 3,801 | 2827-708 | p780-64 core | 18% | 22% | 14,414 |
| S30503A2 | Update Secondary (USBT) | 32 | 4,088 | 2827-708 | p780-64 core | 16% | 19% | 13,198 |
| S30503A3 | Key Date Valuation (KDV) | 32 | 8,465 | 2827-708 | p780-64 core | 12% | 21% | 12,100 |
| S30503A4 | GL Connector (prep) | 32 | 1,047 | 2827-708 | p780-64 core | 26% | 1% | 1,507 |
| | TOTAL (sec) | | 17,401 | | | | | |
| | TOTAL (hour) | | 4.8 | | | | | |

Table 3: BA8 AFI 1.5 Million Accounts – 32 Jobs with 8 CP Results

8.1.2 BA8 AFI 1.5 Million Accounts – 64 Jobs with 8 CP Results

This measurement is almost the same as the previous one, except we increased the number of parallel jobs from 32 to 64, and the maximum number of parallel jobs for GL step was limited at 48 by SAP code.

| Runid | Process Step | # of Batch Jobs | Elapsed Time (sec) | DB Server | App. Server | DB Server Util. | App. Server Util. | DB I/O/Sec. (Peak 5M Avg) |
|----------|--------------------------|-----------------|--------------------|-----------|--------------|-----------------|-------------------|---------------------------|
| S30510A1 | Post External (PEBT) | 64 | 2,690 | 2827-708 | p780-64 core | 26% | 33% | 21,715 |
| S30510A2 | Update Secondary (USBT) | 64 | 2,193 | 2827-708 | p780-64 core | 29% | 35% | 24,297 |
| S30510A3 | Key Date Valuation (KDV) | 64 | 5,616 | 2827-708 | p780-64 core | 22% | 40% | 22,148 |
| S30606A1 | GL Connector (prep) | 48 | 347 | 2827-708 | p780-64 core | 49% | 1% | 1,892 |
| | TOTAL (sec) | | 10,846 | | | | | |
| | TOTAL (hour) | | 3.0 | | | | | |

Table 4: BA8 AFI 1.5 Million Accounts – 64 Jobs with 8 CP Results

8.1.3 BA8 AFI 1.5 Million Accounts – 96 Jobs with 8 CP Results

In this measurement, to get better total elapsed time, we increased the number of parallel jobs again from 64 to 96.

| Runid | Process Step | # of Batch Jobs | Elapsed Time (sec) | DB Server | App. Server | DB Server Util. | App. Server Util. | DB I/O/Sec. (Peak 5M Avg) |
|----------|--------------------------|-----------------|--------------------|-----------|--------------|-----------------|-------------------|---------------------------|
| S30513A1 | Post External (PEBT) | 96 | 2,038 | 2827-708 | p780-64 core | 37% | 46% | 25,066 |
| S30513A2 | Update Secondary (USBT) | 96 | 1,633 | 2827-708 | p780-64 core | 43% | 50% | 34,891 |
| S30513A3 | Key Date Valuation (KDV) | 96 | 4,729 | 2827-708 | p780-64 core | 33% | 59% | 30,479 |
| S30606A1 | GL Connector (prep) | 48 | 347 | 2827-708 | p780-64 core | 49% | 1% | 1,892 |
| | TOTAL (sec) | | 8,747 | | | | | |
| | TOTAL (hour) | | 2.4 | | | | | |

Table 5: BA8 AFI 1.5 Million Accounts -- 96 Jobs with 8 CP Results

8.1.4 BA8 AFI 1.5 Million Accounts – 128 Jobs with 8 CP Results

In this measurement, again, we increased the number of parallel jobs from 96 to 128, to get better elapsed time.

| Runid | Process Step | # of Batch Jobs | Elapsed Time (sec) | DB Server | App. Server | DB Server Util. | App. Server Util. | DB I/O/Sec. (Peak 5M Avg) |
|----------|--------------------------|-----------------|--------------------|-----------|--------------|-----------------|-------------------|---------------------------|
| S30514A1 | Post External (PEBT) | 128 | 1,802 | 2827-708 | p780-64 core | 42% | 56% | 26,381 |
| S30514A2 | Update Secondary (USBT) | 128 | 1,451 | 2827-708 | p780-64 core | 52% | 62% | 37,649 |
| S30514A3 | Key Date Valuation (KDV) | 128 | 4,361 | 2827-708 | p780-64 core | 40% | 71% | 36,831 |
| S30606A1 | GL Connector (prep) | 48 | 347 | 2827-708 | p780-64 core | 49% | 1% | 1,892 |
| | TOTAL (sec) | | 7,961 | | | | | |
| | TOTAL (hour) | | 2.2 | | | | | |

Table 6: BA8 AFI 1.5 Million Accounts – 128 jobs with 8 CP results

8.1.5 BA8 AFI 1.5 Million Accounts – 128 Jobs with 6 CP Results

In this measurement, we stayed at 128 parallel jobs, and we optimized the zEC12 SAP DB server configuration by changing from 8 to 6 general processors (CPs), while still maintaining the same total elapsed time from the previous measurement.

| Runid | Process Step | # of Batch Jobs | Elapsed Time (sec) | DB Server | App. Server | DB Server Util. | App. Server Util. | DB I/O/Sec. (Peak 5M Avg) |
|----------|--------------------------|-----------------|--------------------|-----------|--------------|-----------------|-------------------|---------------------------|
| S30529A1 | Post External (PEBT) | 128 | 1,793 | 2827-706 | p780-64 core | 53% | 56% | 25,959 |
| S30529A2 | Update Secondary (USBT) | 128 | 1,361 | 2827-706 | p780-64 core | 70% | 64% | 40,314 |
| S30529A3 | Key Date Valuation (KDV) | 128 | 4,518 | 2827-706 | p780-64 core | 55% | 70% | 36,730 |
| S30605A1 | GL Connector (prep) | 48 | 393 | 2827-706 | p780-64 core | 59% | 1% | 1,475 |
| | TOTAL (sec) | | 8,065 | | | | | |
| | TOTAL (hour) | | 2.2 | | | | | |

Table 7: BA8 AFI 1.5 Million Accounts – 128 Jobs with 6 CP Results

8.1.6 BA8 AFI 1.5 Million Accounts – 128 jobs with 3 CP + 3 zIIP Results

In this measurement, we changed the zEC12 SAP DB server configuration to utilize zIIP engines by having three general processors (CPs) and three zIIPs, while still maintaining the same total elapsed time.

| Runid | Process Step | # of Batch Jobs | Elapsed Time (sec) | DB Server | App. Server | Avg. DB Server Util. | App. Server Util. | DB I/O/Sec. (Peak 5M Avg) |
|----------|--------------------------|-----------------|--------------------|------------------|--------------|----------------------|-------------------|---------------------------|
| S30603A1 | Post External (PEBT) | 128 | 1,831 | 2827-703+3 zIIPs | p780-64 core | 54% | 56% | 21,755 |
| S30603A2 | Update Secondary (USBT) | 128 | 1,432 | 2827-703+3 zIIPs | p780-64 core | 68% | 61% | 36,047 |
| S30603A3 | Key Date Valuation (KDV) | 128 | 4,503 | 2827-703+3 zIIPs | p780-64 core | 54% | 69% | 37,464 |
| S30605A3 | GL Connector (prep) | 48 | 332 | 2827-703+3 zIIPs | p780-64 core | 64% | 1% | 756 |
| | TOTAL (sec) | | 8,098 | | | | | |
| | TOTAL (hour) | | 2.2 | | | | | |

Table 8: BA8 AFI 1.5 Million Accounts – 128 jobs with 3CP + 3 zIIP Results

8.2 Detailed Data for 3 CP + 3 zIIP

Listed below are some more detailed measurement results for 3 CP + 3 zIIP measurement. The first two tables (Tables 9 and 10) show metrics from SAP.

| Parent Batch: Process Step | Elapsed Time | App. Server CPU Time | DBREQ Time | RFC Time | ENQ Time |
|----------------------------|--------------|----------------------|------------|----------|----------|
| PEBT | 1,831 | 25 | 128 | 76 | 33 |
| USBT | 1,432 | 42 | 176 | 53 | 23 |
| KDV | 4,503 | 28 | 2,624 | 75 | 34 |
| GL | 332 | 2 | 48 | 8 | 1 |
| Total (sec) | 8,098 | | | | |
| Total (hr) | 2.2 | | | | |

Table 9: 3 CP + 3 zIIP Parent Batch Detailed Data

Note that the child batch elapsed times (Table 10) should be close to the parent batch elapsed time (Table 9). The KDV step has largest difference. Please refer to page 23 for an explanation.

| Child Batch: Process Step | # of Child Batches | Elapsed Time | App. Server CPU Time | DBREQ Time | RFC Time | ENQ Time |
|---------------------------|--------------------|--------------|----------------------|------------|----------|----------|
| PEBT | 128 | 1,695 | 406 | 722 | 0 | 5 |
| USBT | 128 | 1,257 | 351 | 549 | 0 | 5 |
| KDV | 128 | 1,930 | 652 | 609 | 0 | 3 |
| GL | 48 | 313 | 0.1 | 313 | 0 | 0 |

Table 10: 3 CP + 3 zIIP Child Batch Detailed Data

This table shows data from the application server. The data comes from VMSTAT and SAP statistics.

| Process Step | # of Child Batches | Avg. CPU Util. | 30 Sec. Peak CPU Util. | 30 Sec. Peak Active Memory (GB) | Peak IOWAIT |
|--------------|--------------------|----------------|------------------------|---------------------------------|-------------|
| PEBT | 128 | 56% | 58% | 107.2 | 0% |
| USBT | 128 | 61% | 63% | 109.6 | 0% |
| KDV | 128 | 69% | 74% | 107.9 | 0% |
| GL | 48 | 1% | 1% | 108.8 | 0% |

Table 11: 3 CP + 3 zIIP Application Server Detailed Data

This table shows data from the DB server, which comes from RMF and Tivoli Omegamon XE for DB2.

| Process Step | CPU Util. | zIIP Util. | BPOOL Size | BPOOL Hit Ratio | 5 Min. Peak I/O Rate (/sec) | AVG Reads (/sec) | AVG Prefetch Reads (/sec) | AVG Async. Writes (/sec) | 5 Min. Peak Log Writes (/sec) | zHPF |
|--------------|-----------|------------|------------|-----------------|-----------------------------|------------------|---------------------------|--------------------------|-------------------------------|------|
| PEBT | 54% | 55% | 127 GB | 96.50% | 21,755 | 10,016 | 1,009 | 8,728 | 1,793 | 99% |
| USBT | 56% | 80% | 127 GB | 70.40% | 36,047 | 19,955 | 12,706 | 346 | 680 | 99% |
| KDV | 47% | 61% | 127 GB | 97.92% | 37,464 | 8,586 | 4825 | 2424 | 866 | 99% |
| GL | 64% | 65% | 127 GB | 99.52% | 756 | 539 | 123 | 1 | 10 | 99% |

Table 12: 3 CP + 3 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 collected zIIP data over time for the first three 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.

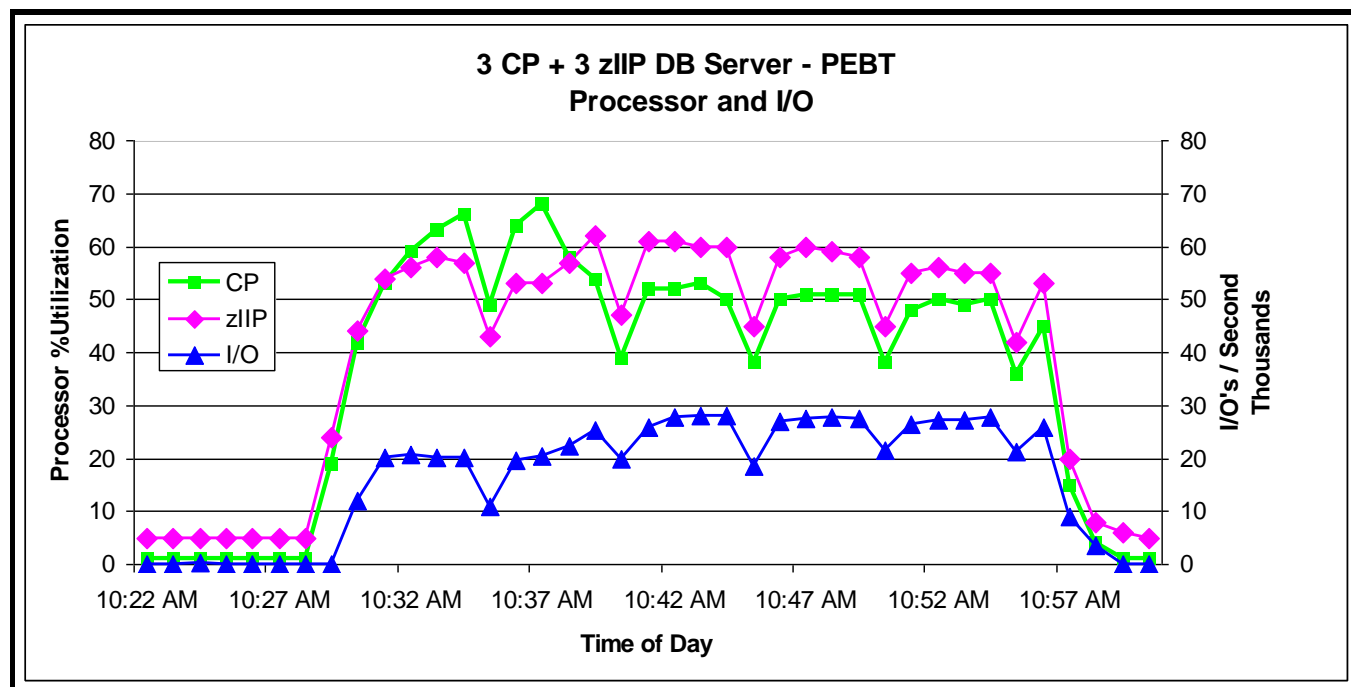


Figure 3: 3 CP + 3 zIIP PEBT DB Server Processor Utilization and I/O rates

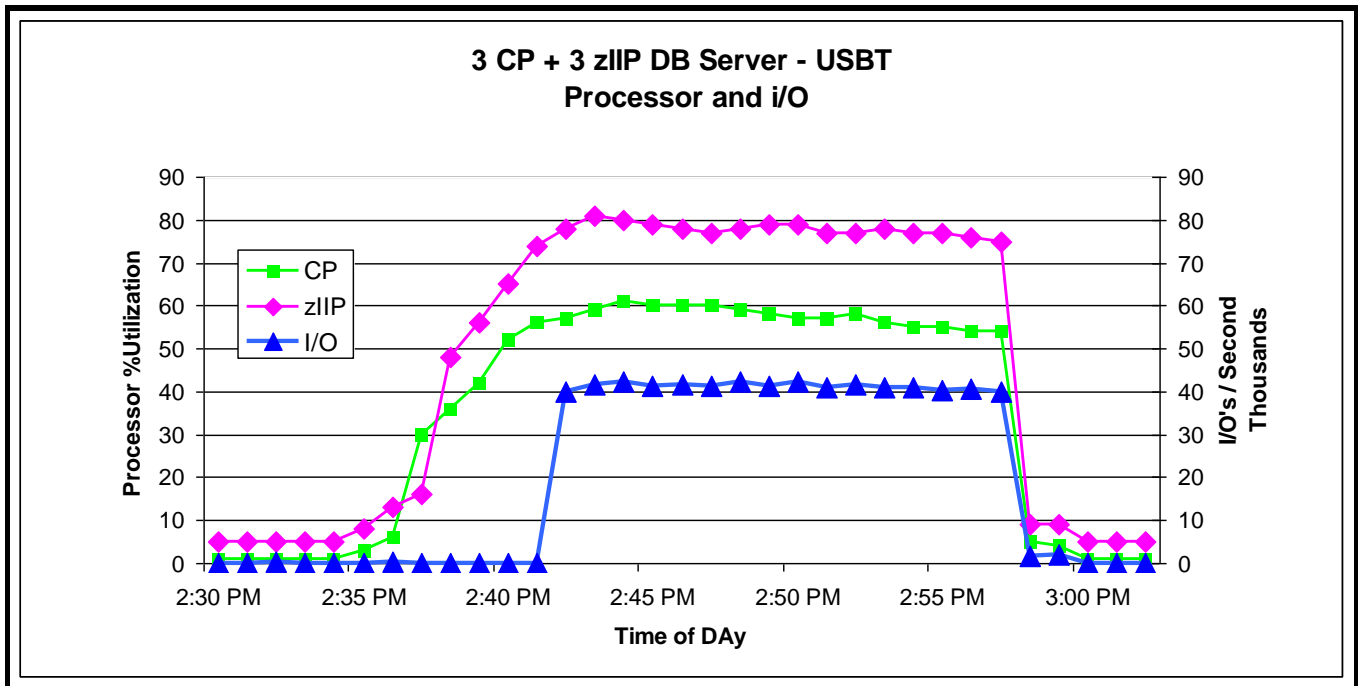


Figure 4: 3 CP + 3 zIIP USBT DB Server Processor Utilization and I/O rates

The KDV step has the longest elapsed time of all the steps [Figure 5]. As a result, the x-axis is not to the same scale as other figures. Note that the first part of the graph shows high utilization, and the second part shows low utilization. In the first part, there were 128 batch jobs dispatched to process the KDV step in parallel. After all 128 jobs completed, one batch job (parent task) remained to perform the final processing of KDV step. This is why the parent elapsed time is so much more than the child batch time. SAP has planned an enhancement to improve this in the future.

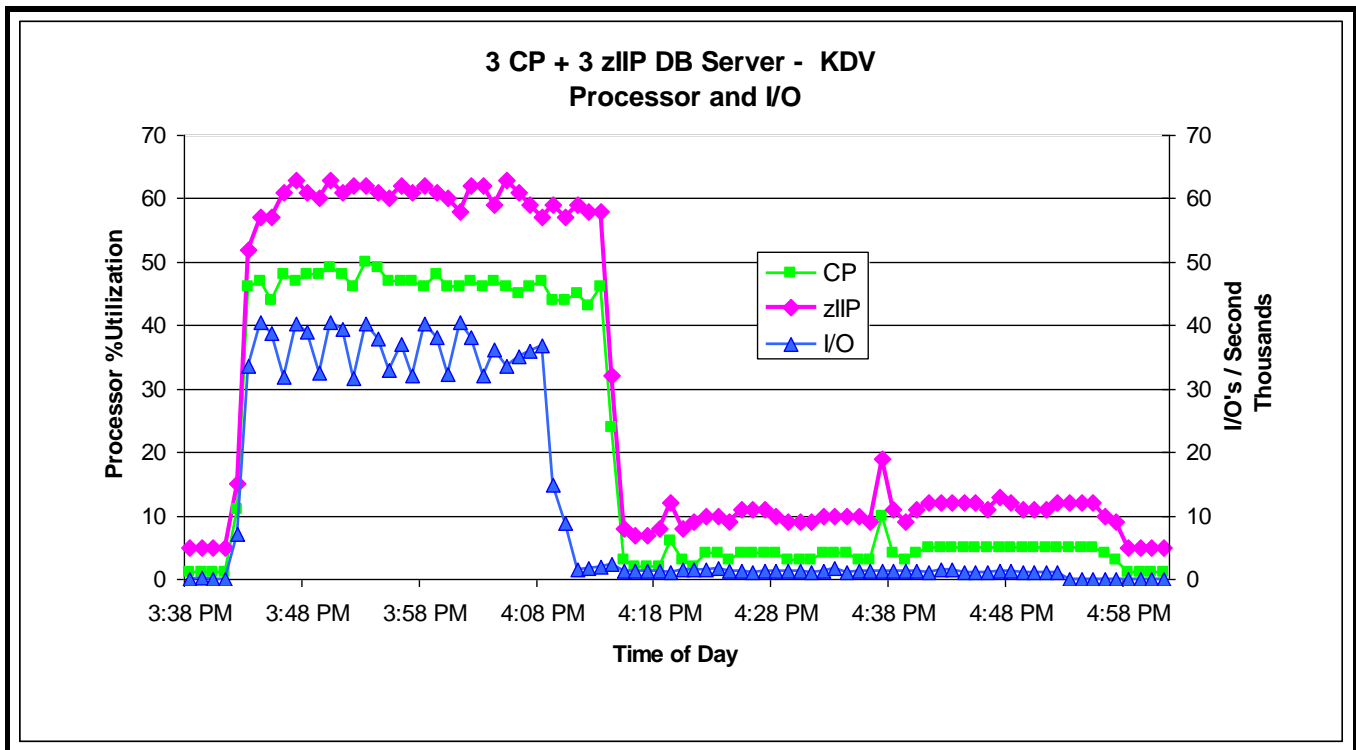


Figure 5: 3 CP + 3 zIIP KDV DB Server Processor Utilization and I/O rates

Shown below is a similar graph for GL. However, since we did not collect the zIIP information in one minute intervals for this step, the graph shown below is from the 6 CP run. As shown in Table 12, the zIIP utilization was slightly higher than the general CPs. We would expect a 3 CP + 3 zIIP graph to have a similar plot.

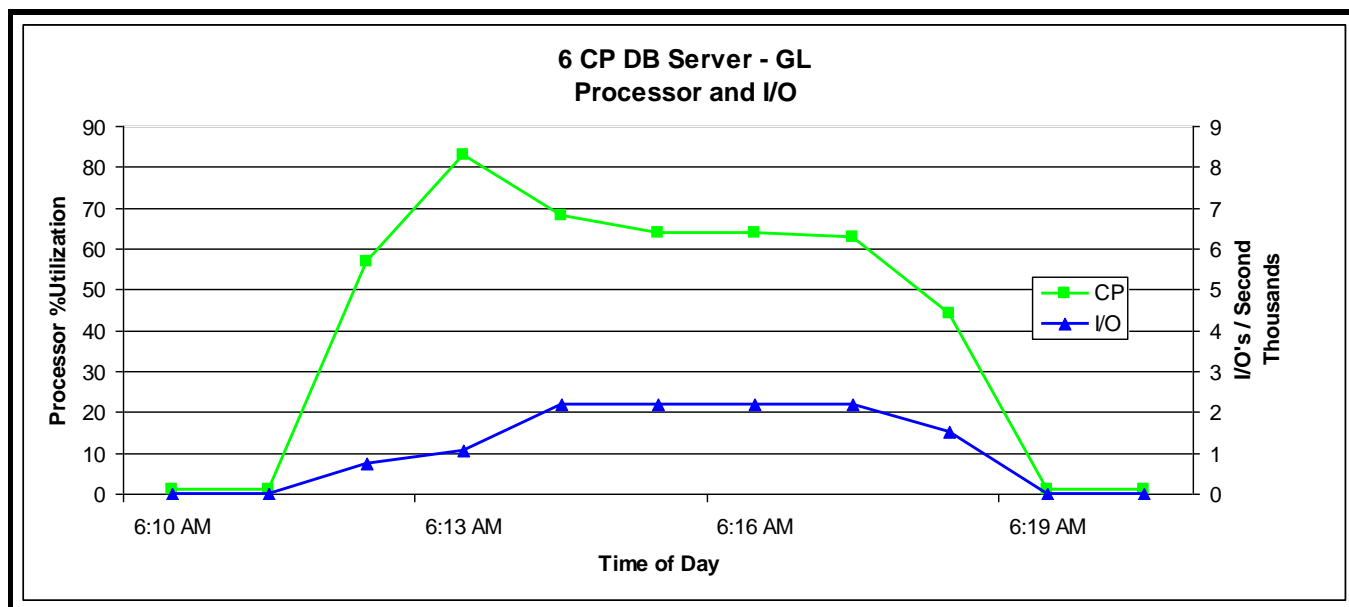


Figure 6: 6 GL DB Server Processor Utilization and I/O rates

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”. Discussed below are several analysis aspects.

Scaling by Varying Number of Parallel Jobs

Total elapsed time to process 1.5 million accounts was reduced from 4.8 hours to 2.2 hours, more than half, by increasing the number of parallel jobs from 32 to 128 jobs.

Database Server

The first four measurement runs used 8 CPs for the DB Server. To optimize the DB2 server configuration, we reduced the number of processors to 6 CPs. Then we did one measurement with 3 CPs + 3 zIIPs, in an attempt to test a more economical DB server configuration, and demonstrate zIIPs. Using zIIPs can reduce cost.

zIIP Exploitation

Analysis, based on data shown in “Table 12: 3 CP + 3 zIIP DB Server Detailed Data” on page 20, indicates that the percent redirect to zIIPs for the entire LPAR ranged, depending on the step, from 55% to 80%. 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., HSM).

Application Server

The application server processing in the AFI account scenario was not as intensive as in the AFI loan scenario. We observed this by determining the CPU cores for each step normalized for utilization and elapsed time. When looking at all four steps combined, the application server cores were 91% and the DB server cores were only 9% of the total used CPU cores; as compared to AFI loan scenario, 96% and 4%, respectively. Essentially, the AFI current account scenario has a core ratio of DB to application server typical as most other SAP application workloads.

The AFI loan scenario is much more complex because of the huge amount of different conditions (ie., full disbursement, partial disbursements, interest calculation, amortization, charges, and discounts). In this test scenario with the AFI current accounts, we only have a limited set of conditions (ie., deposit, withdrawal, and charges). In addition, in the AFI loan scenario, a cash flow has to be calculated for the whole fixed interest rate period. In the AFI current account scenario, we do not have a cash flow generation.

I/O Profile

The BA8 AFI current account scenario has significant I/O rates. One step, Update Secondary (USBT), had a peak 5 minute I/O rate of 40,314 I/Os per second as compared to 24,900 I/Os per second for 15 minutes in the AFI loan scenario. When normalized for utilization, this is more than twice our rough rule of thumb peak I/O rate for this DB server. This means that customers considering BA should have a particularly robust I/O configuration for the DB server. An example would involve large amounts of main memory, high-speed FICON adapters, and a particularly strong storage subsystem.

One factor that can significantly improve I/O performance and FICON adapter efficiency is zHPF. DB2 in an SAP environment is typically a good exploiter of zHPF. Further, FICON Express8S is particularly efficient processing zHPF requests. The BA8 AFI account scenario continues that tradition. In our measurements, we saw 99+% exploitation in all the steps. A second factor in improving I/O performance is to have a storage subsystem that can deliver good response time. In our measurements, we saw each step's peak I/O rate interval had an excellent average DASD response times - 0.4 to 1.5 milliseconds.

10 Conclusions

We can draw three main conclusions from these BA measurements:

- Application server resources are key to overall elapsed time performance.
- BA has high I/O rates.
- Utilizing zIIP can be cost effective.

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.

11 References

- [1] IBM Corp. 2013. *IBM Enterprise System: SAP Bank Analyzer 8.0 AFI Loan Tests*
<http://www-03.ibm.com/support/techdocs/atmastr.nsf/WebIndex/WP102206> or msheets@us.ibm.com
- [2] IBM Corp. 2010. *IBM System z: SAP Bank Analyzer 7.0 Tests*
<http://www.ibm.com/support/techdocs/atmastr.nsf/WebIndex/WP101812> or msheets@us.ibm.com
- [3] SAP AG 2012. SAP Documentation: Bank Analyzer (FS-BA)
http://help.sap.com/saphelp_banking80/helpdata/en/4e/1217c4e1905707e10000000a42189e/frameset.htm
- [4] 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>
- [5] IBM Corp. 2010. *DB2 10 for z/OS* <http://www.ibm.com/software/data/db2/zos/>
- [6] IBM Corp. 2010. *DB2 10 for z/OS Technical Overview*
<http://www.redbooks.ibm.com/redbooks/pdfs/sg247892.pdf>
- [7] 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>
- [8] IBM Corp. 2011. *DB2 10 for z/OS with SAP on IBM System z Performance*
<http://www.ibm.com/support/techdocs/atmastr.nsf/WebIndex/WP101845>
or msheets@us.ibm.com
- [9] IBM Corp. 2010. *IBM zEnterprise System* <http://www.ibm.com/systems/z/hardware/zenterprise/>
- [10] IBM Corp. 2010. *New IBM System Storage DS8800*
<http://www.ibm.com/systems/storage/news/center/disk/enterprise/>
- [11] 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/atmastr.nsf/WebIndex/WP101442> or msheets@us.ibm.com
- [12] 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>
- [13] 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>

[14] 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/atmastr.nsf/WebIndex/WP101243> or msheets@us.ibm.com

[15] IBM Corp. 2010. *Large Systems Performance Reference*

[https://www.ibm.com/servers/resourceLink/lib03060.nsf/pages/lspindexpdf/\\$file/SC28118714_20100714.pdf](https://www.ibm.com/servers/resourceLink/lib03060.nsf/pages/lspindexpdf/$file/SC28118714_20100714.pdf)