



IBM Technical Brief

**IBM z Systems®:
SAP® Insurance Analyzer**

Authors:

**Seewah Chan
Paul DeLessio
Paul Lekkas**

Document Owner:

**Paul D. DeLessio
SAP on z Systems Performance
IBM Poughkeepsie**

**Version: 1.0
Date: February 20, 2017**

Filename: System_z_SAP_IA-1.0.pdf

1 Table of Contents

1	Introduction	4
2	Reminders	4
3	Trademarks	5
4	Feedback	6
5	Project Team	6
6	Acknowledgements	6
7	IA Business Processes	7
8	Workload and Test Background	9
8.1	KDV 1M, 9M and 18M policy scaling	9
8.2	KDV 1 versus 2 Application Server comparison using 1M policies	10
9	Configurations	11
9.1	Hardware Environment	11
9.2	Software Environment	12
10	Test Results and Analysis.....	14
10.1	KDV 1M, 9M and 18M policy scaling results	15
10.1.1	Scaling Results - Elapsed Time	16
10.1.2	Scaling Results - CPU Utilization.....	18
10.1.3	Scaling Results - DB Growth Rate in GB	19
10.2	KDV 1 versus 2 Application Server comparison measurements	22
10.2.1	KDV 1 vs 2 Application Server - Elapsed Time.....	22
10.2.2	KDV 1 vs 2 Application Server - DB Server CPU Utilization	23
10.2.3	KDV 1 vs 2 Application Server - Application Server Average CPU Utilization	24
11	Summary	25
12	References	26

Tables

Table 1: DB2 Bufferpool Settings	12
Table 2: SAP Application Levels	13
Table 3: KDV 1M, 9M and 18M policy scaling results	15
Table 4: Scaling Results - DB Growth Rate in GB.....	19
Table 5: Scaling Results - DB Growth Rate in GB - Normalized to 1M Objects	21
Table 6: KDV 1 vs 2 Application Server - Elapsed Time	22

Figures

Figure 1: Key IA Business Processes	7
Figure 2: Conceptual View of Insurance Analyzer Hardware Landscape	12
Figure 3: KDV for Feb. - Elapse Time in Hours	16
Figure 4: KDV for Feb. - Elapse Time in Hours - Normalized to 1M Objects.....	17
Figure 5: KDV for Feb. - DB Server CPU Utilization	18
Figure 6: KDV for Feb. - Application Server CPU Utilization	19
Figure 7: KDV for Feb. - DB Growth Rate in GB.....	20
Figure 8: KDV for Feb. - DB Growth Rate in GB - Normalized to 1M Objects.....	21
Figure 9: KDV for Feb. 1 vs 2 Application Server - Elapsed Time in hours	22
Figure 10: KDV for Feb. 1 vs 2 Application Server - DB Server CPU Utilization	23
Figure 11: KDV for Feb. 1 vs 2 Application Server - Average Application Server CPU Utilization.....	24

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. MSG Global/Nexontis, as a joint developer of Insurance Analyzer (IA) with SAP, is also an integral partner in this project.

This paper describes tests we performed with IBM z Systems, DB2 for z/OS and SAP's Insurance Analyzer (IA). IA is an integrated family of financial software intended to help insurance providers comply with the new International Financial Reporting Standards (IFRS 4) accounting principles and Solvency II risk guidelines. The standards are required of insurance companies operating out of the European Union (EU), or international insurance companies intending to do business in the EU.

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

Our tests were not performance tests or formal benchmarks. While we did report several performance related metrics, our tests were more in the nature of function tests. The exposure of the application was to our available on-hand test environment hardware with existing features and functions. For example, we ran in a zEC12 LPAR, with dedicated processors, versus running on a dedicated z13 server. Despite some limitations with time and the environment, it was adequate to demonstrate the scalability of SAP Insurance Analyzer on IBM Systems.

IA was new with several functions, features, and extensions added as the tests progressed. The project enabled SAP to enhance the IA application to run more efficiently with the recent versions of z/OS and 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 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

© IBM Corporation 1994-2017. All rights reserved. References in this document to IBM products or services do not imply that IBM intends to make them available in every country.

AIX, DB2, DFSMSdftp, DFSMSHsm, DFSMSrmm, DS8800, FICON, FICON Express, IBM, IBM eServer, IBM logo, Language Environment, OSA, OSA Express, OS/390, Power7, RACF, RMF, S/390, System Storage, System z, VTAM, z Systems, zEC12, z/Architecture, z/OS, and zSeries are trademarks or registered trademarks of the International Business Machines Corporation in the United States and other countries.

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

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.

Other products may be trademarks or registered trademarks of their respective companies.

Information is provided "AS IS" without warranty of any kind.

Information concerning non-IBM products was obtained from a supplier of these products, published announcement material, or other publicly available sources and does not constitute an endorsement of such products by IBM. Sources for non-IBM list prices and performance numbers are taken from publicly available information, including vendor announcements and vendor worldwide homepages. IBM has not tested these products and cannot confirm the accuracy of performance, capability, or any

other claims related to non-IBM products. Questions on the capability of non-IBM products should be addressed to the supplier of those products.

4 Feedback

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

5 Project Team

Role /Area responsible	Person responsible	Company
Executive Sponsor	Armin Fischer	SAP AG
Executive Sponsor	Louise Cooke	Nexontis Consulting
Executive Sponsor	Paul Lekkas	IBM Germany
Project Management	Justine Angeles	Nexontis Consulting
Technology – Database DB2	Seewah Chan	IBM Poughkeepsie
Test Data Support	Jonathan Chitarra	Nexontis Consulting
Project Lead IBM	Paul DeLessio	IBM Poughkeepsie
Application Expert	Klaus Mitulla	Nexontis Consulting
Technical Advisor - SAP / IA	Nikolai Sauerwald	SAP AG
Project Steering and Management	Klaus Scholten	IBM Germany

6 Acknowledgements

The authors would like to thank the following colleagues for their contributions and support:

- Andrea Fuga, IBM Poughkeepsie, for System Programming support.
- Thomas Litty, IBM Poughkeepsie, for System Programming support.
- Veng Ly, of IBM Poughkeepsie, for test support and coverage.
- Rose Manz, IBM Poughkeepsie, for test support and coverage.
- Johannes Schuetzner, IBM Boeblingen, for DB2 support.

7 IA Business Processes

Insurance companies need to follow complicated accounting principles and risk guidelines such as the International Financial Reporting Standards (IFRS 4) accounting principles and Solvency II risk guidelines. SAP Insurance Analyzer is an integrated family of financial software which enables insurance companies to address the new regulations.

The SAP Insurance Analyzer solution offers the advantages of integration of finance and risk, subledger approach for assets and liabilities, integration of SAP general ledger and other reporting systems, multiledger capabilities, versioning and drill-down, and a full audit trail. In this study, we have narrowed the scope of tests to the IFRS 4 relevant SAP Accounting for Insurance Contracts solution process.

SAP Insurance Analyzer is a “sequential” batch workload that involves a set of process steps. Listed below are the key process steps with brief descriptions:

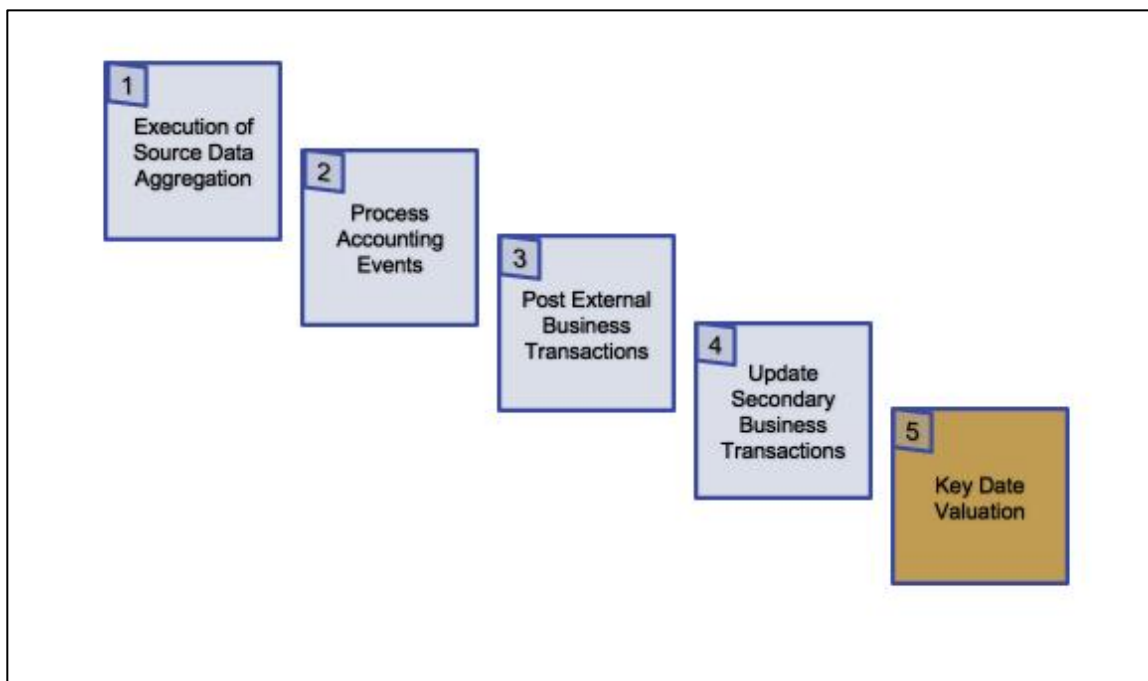


Figure 1: Key IA Business Processes

1. **Execution of Source Data Aggregation:** Generates aggregated objects from single objects in the Source Data Layer (SDL) and in the Results Data Layer (RDL) that together form an abstract aggregation object. This reduces the volume of input data for subsequent processes, such as accounting event processing.
2. **Processing of Accounting Events:** Generates a worklist for the accounting processes. This worklist is based on application events that the system derives from operational events in the Source Data Layer (SDL) and from analytical results in the Results Data Layer (RDL).

3. **Post External Business Transactions (PEBT):** The PEBT process step performs the posting of external business transactions. 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 an Accounting format, before transferring to the specified accounting system. The step then posts the external business transactions to the RDL.
4. **Update Secondary Business Transactions (USBT):** This process performs IFRS 4 Phase II valuations and stores the results in the RDL.
5. **Key Date Validation (KDV):** This process is used to determine valuation of a position for financial instruments based on a key date, and goes through multiple calculation steps.

8 Workload and Test Background

The SAP Insurance Analyzer (IA) process is a sequential batch workload that involves a set of process steps. Our measurements primarily focused on the Key Date Valuation (KDV) step, since it is the critical "Month End Process". Other processes are normally run in advance, on a weekly or daily basis, with lower volumes and less time-sensitive batch windows. In our efforts, we did not perform the Execution of Source Data Aggregation and Processing of Accounting Events. Other SAP IA process steps, such as Post External Business Transactions (PEBT) and Update Secondary Business Transactions (USBT) were executed, with very limited tuning in our efforts and were not the focus of measurements.

Each of our test scenarios, whose results will be discussed in more detail in subsequent sections, consisted of several steps. In general, the final measurement for each scenario went through a set of account batch process steps, but with only measurements of the KDV batch process step. This is because the KDV step has limited monthly period ending batch windows, where other steps would normally be iterative updates throughout a month, and not limited to a time-sensitive batch window.

Measurements focused on February, because the first month, coincidentally January in this case, has a one-time contract initialization overhead and includes extra data preparation costs. This effect is a test scenario cost, not an annual January activity cost, and why January was not used in testing.

As mentioned in section 1, "[Introduction](#)", we focused on the application interactions with the infrastructure. We did not run standard SAP application benchmarks. SAP Insurance Analyzer has many more possible scenarios and process steps we did not test. The test scenarios were selected by SAP, IBM and Nexontis stakeholders as being reasonable to evaluate scalability. 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 little virtual storage tuning.

Using the KDV month end process, we executed two comparison test scenarios.

- Key Date Valuation (KDV) for February 1 Million (M), 9M and 18M policy scaling measurements.
- Key Date Valuation (KDV) for February, at 1M policies, with 1 vs 2 application servers, elapsed time comparison.

With a focus on infrastructure, here are some short summaries of the test scenarios. More details can be found in section 10, "[Test Results and Analysis](#)".

8.1 KDV 1M, 9M and 18M policy scaling

At the beginning of this project, we had targeted the high-end scaling measurement of 18M policies which was higher than the largest profiled SAP Insurance customer scenario at 17M policies. We also had defined two lower scaling measurement points of 1M and 9M policies.

The three scalability measurement points at 1M, 9M and 18M policies equate to 3M, 27M and 54M insurance coverages respectively. For all three scaling points, we kept the hardware configuration and

the load the same with 120 parallel jobs. The Application and DB servers processed the same pace for 1M, 9M and 18M. Therefore, the CPU% numbers stay consistent, but the elapsed times showed a near linear 20X growth from 1M to 18M policy scenarios.

For a definition of policy extents and calculations behind 1M, 9M and 18M measurement points see section 10, “[Test Results and Analysis](#)”.

For the KDV 1M, 9M and 18M policy scaling measurement related detail see section 10.1, “[KDV 1M, 9M and 18M policy scaling results](#)”.

8.2 KDV 1 versus 2 Application Server comparison using 1M policies

This scenario simply was to show the effect of adding a second application server to the KDV 1M policy scenario. The definition of policy extents and calculations remain the same for both measurements as for the 1M policy object test above, and detailed in section 10, “[Test Results and Analysis](#)”. In fact, the single application server run here is the same 1M measurement in the scaling measurements. The goal of this test was to determine what elapsed time reductions could be achieved, and what the effects would be to the DB Server and Application Servers.

For the KDV 1 versus 2 Application server comparison using 1M policies measurement related detail, see section 10.2, “[KDV 1 versus 2 Application Server comparison measurements](#)”.

9 Configurations

9.1 Hardware Environment

System z DB Server: Tests were performed on a single zEC12 2827-716. The runs utilized one dedicated LPAR for z/OS zDB2 11 with 16 processors and 200 GB of real storage, this LPAR represented 36,700 SAPS.

Storage: IBM System Storage Server DS8800 Model 2421-951/95E, with a total capacity of 39 TB 15K rpm HDD storage space and 384 GB cache. The IBM System Storage Server was attached to the zEC12 via two FICON Express8S LX connections. The usable capacity was for database, logs and FlashCopy sets.

The SAP IA 1.0 SP5 system used a zDB2 11 database with (100) to (248) 3390 Mod 54 volumes - about 5 to 12TB of allocated space, depending on the test scenario size. Approximately half of all active pages were in table spaces and half in index spaces.

Application Servers: For the primary measurements one dialog application server plus a standalone enqueue server were used and consisted of the following:

- (1) IBM Power7 780 (9179-MHD) dialog application server with (64) 4.42 GHz SMT4 processor cores and 1 TB of memory. This was a 178,770 SAPS rated server.
- (1) IBM Power7 750 (8233-E8B) standalone enqueue server with (32) 3.6 GHz SMT4 processor cores and 128 GB of memory. This 66,220 SAPS server was oversized for the effort.

With a secondary (2) application server comparison effort, an identical IBM Power7 780 (9179-MHD) dialog application server was added, for a combined 357,540 dialog application server SAPS.

Network: 10 Gigabit Ethernet networks were used for all connections. Each of the application servers was connected via a 10 Gigabit Ethernet switch to the zEC12 via OSA-Express4S 10 Gbps Ethernet SR adapter.

Below is a conceptual view of the hardware landscape.

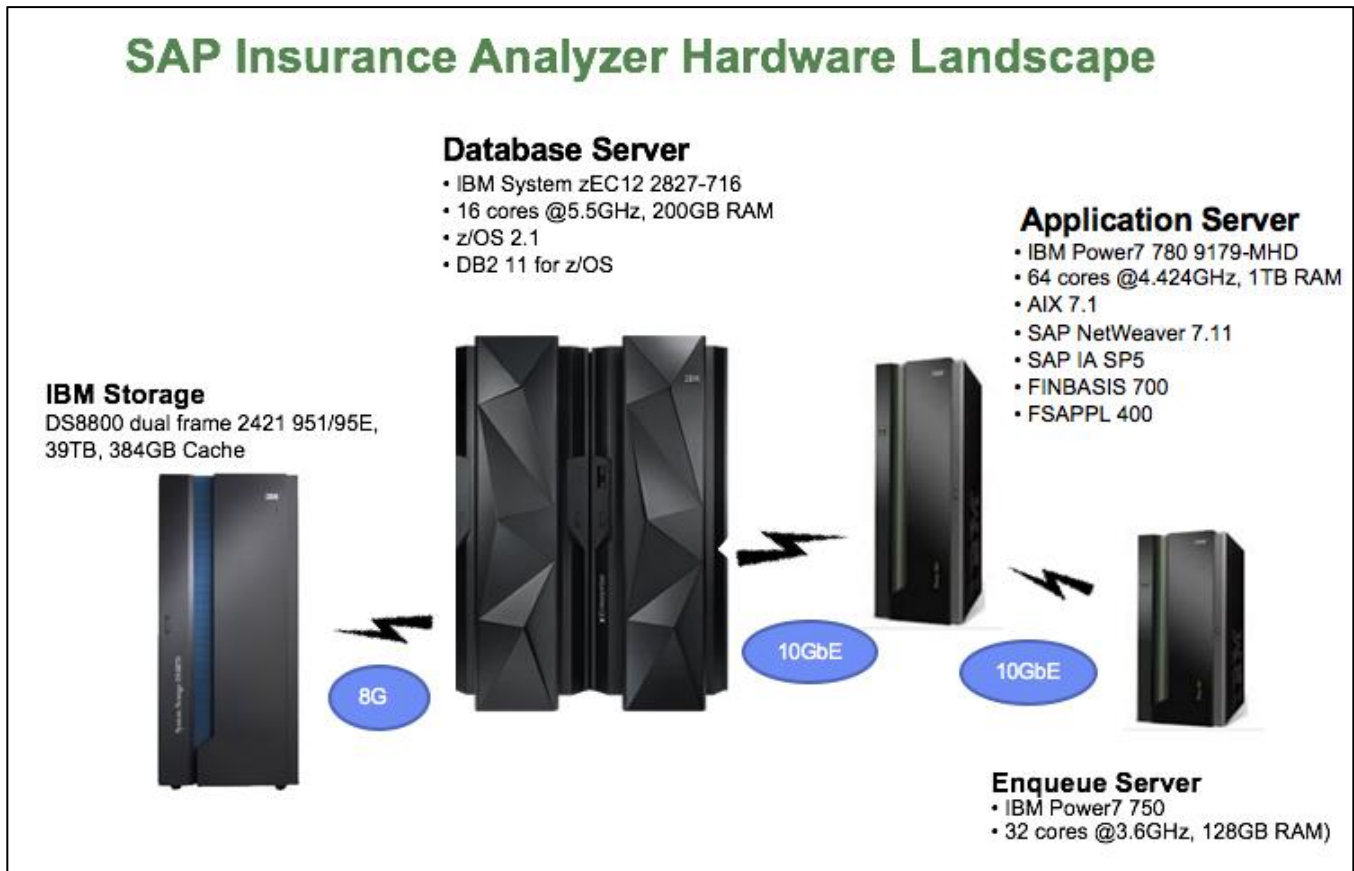


Figure 2: Conceptual View of Insurance Analyzer Hardware Landscape

9.2 Software Environment

z/OS

z/OS release 02.01.00 (R2.01)

DB2 11 dated September 2014

Total allocated bufferpool size was 70.4GB. There was little attempt made to tune the bufferpools. The settings are listed below.

BP name	PGFIX	VP Pages (K)	BP Size (MB)	VPSEQT [%]	DWQT [%]	VDWQT [%]	Page steal
BP0	YES	4,000	16,000	50	50	10	LRU
BP2	YES	4,000	16,000	50	50	10	LRU
BP3	YES	4,000	16,000	50	50	10	LRU
BP8K0	YES	10	80	50	50	10	LRU
BP8K1	YES	2,000	16,000	50	50	10	LRU
BP32K	YES	100	3,200	50	50	10	LRU
BP32K1	YES	100	3,200	50	50	10	LRU

Table 1: DB2 Bufferpool Settings

IBM DB2 Connect “Thin client” side: Driver for ODBC, CLI, JDBC and SQLJ - Version 10.5 FP2 special 31569.

AIX

AIX 7.1

oslevel -s

7100-03-05-1524

SAP Application Levels

SAP EHP1 for SAP NetWeaver 7.11, kernel release 721 patch no 525

Software Component	Release	Level	Highest Support Package	Short Description of Software Component
SAP_ABA	711	13	SAPKA71113	Cross-Application Component
SAP_BASIS	711	13	SAPKB71113	SAP Basis Component
PI_BASIS	711	13	SAPK-71113INPIBASIS	Basis Plug-In
ST-PI	2008_1_710	8	SAPKITLRE8	SAP Solution Tools Plug-In
SAP_BW	711	13	SAPKW71113	SAP Business Warehouse
FINBASIS	700	22	SAPK-70022INFINBASIS	Financial Basis
SEM-BW	700	22	SAPK-70022INSEMBW	SEM-BW 700: Add-On Installation
BI_CONT	711	2	SAPK-71102INBICONT	Business Intelligence Content
FSAPPL	400	11	SAPK-40011INFSAPPL	SAP Banking Services
NXC	100_400	5	SAP-14005INNXC	Insurance Analyzer
ST-A/PI	01R_710	2	SAPKITAB8P	Service tools for other App./Netweaver200

Table 2: SAP Application Levels

As can be seen above, IA consists of SAP’s core insurance analyzer NXC component, and banking BI Content and BW. IA requires Unicode, and z Systems has hardware data compression, which was used to compress tablespaces. With the function test nature of the IA measurement effort, numerous transport corrections were created and applied by the team became official software component transport updates.

10 Test Results and Analysis

During 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 formal benchmark. Each of the scenarios consisted of a set of steps with unique characteristics, along with the measured KDV step. Combining the sheer number of steps, and their uniqueness, makes the IA KDV scenario a complex workload to measure.

Below are some definitions and meanings behind test results.

A policy is defined by the following extents:

- Each set of data represents an insurance policy portfolio.
- Each policy contains 2 to 4 underlying coverages, with an average of 3 per policy.
- All coverages are issued by a single Legal Entity (LE).
- Monthly premium transactions - 24 premiums Business Transactions (BT) per coverage per year.
Note: 24 premiums BT = (1 premium due and 1 premium paid per month) for 12 months.
- Average of 2 "claims" BTs every 60 coverages.
- "Claims" dates randomly spread throughout the year.
- Cash flow best estimates are provided for the portfolio.

Calculation information with respect to KDV for 1M, 9M and 18M policy scaling:

To estimate the number of calculations per coverage, use the formula:

- Calculations for 1 coverage = 5 flow types * 2 posting key figures * 3 calculation methods.

Flow types (5):

- Insurance other cash inflow
- Insurance expense and other cash flows
- Insurance discretionary benefit
- Insurance guaranteed benefit / claim
- Insurance premium

Posting key figures (2):

(Debit, Credit)

Sample calculation methods (3):

(Accrual, Deferral and Component Valuation)

To substitute the coverage calculation formula within the scenarios:

- 1M Policies or 3M Coverages: $(3M * 5 * 2 * 3) = 90M$ Calculations
- 9M Policies or 27M Coverages: $(27M * 5 * 2 * 3) = 810M$ Calculations
- 18M Policies or 54M Coverages: $(54M * 5 * 2 * 3) = 1,620M$ Calculations

SAP Insurance Analyzer, like SAP Bank Analyzer, has a heavy I/O profile with skewed hot spots across various steps. Due to the limited focus on the calculation heavy SAP IA KDV process, it was not the intent of this paper to tune or balance SAP IA I/O. Listed below are the test results selected as being the most useful, given the time and environment constraints of this effort.

For each process step of each scenario, we report Number of Policies processed, Number of batch threads, run times, run times normalized per 1M policy objects processed, Application Server utilization, DB Server utilization, DB growth and DB growth normalized per 1M policies processed. These last two metrics are discussed in more detail see section 10.1.3, “*Scaling Results - DB Growth Rate in GB*”. The utilization metrics listed in the following tables are average sustained busy periods during the run.

10.1 KDV 1M, 9M and 18M policy scaling results

The run time corresponds to the duration of time the run took when the number of parallel batch threads is set to 120, while using one P7 780 application server to drive workload, in our measurement environment.

Number of Policies	Number batch of threads	Runtime (hours)	Runtime normalized per 1M (hours)	DB Server Utilization	AppServer Utilization	Total DB Growth (GB)	DB Growth Normalized per 1M (GB)
1M	120	2.3	2.3	15%	83%	116	116
9M	120	23.8	2.6	15%	78%	1,013	112
18M	120	45.5	2.5	14%	79%	2,059	115

Table 3: KDV 1M, 9M and 18M policy scaling results



10.1.1 Scaling Results - Elapsed Time

Shown graphically the KDV measurements for 1M, 9M and 18M policies exhibit a near linear 20X elapsed time growth. The small deviation from linearity is in part due to memory limitation of our test environment DB Server LPAR and in part due to the high database I/O activity, especially for the 9M and 18M cases. More physical memory on the DB Server and DASD with newer technology and larger cache would benefit these larger DB size scenarios.

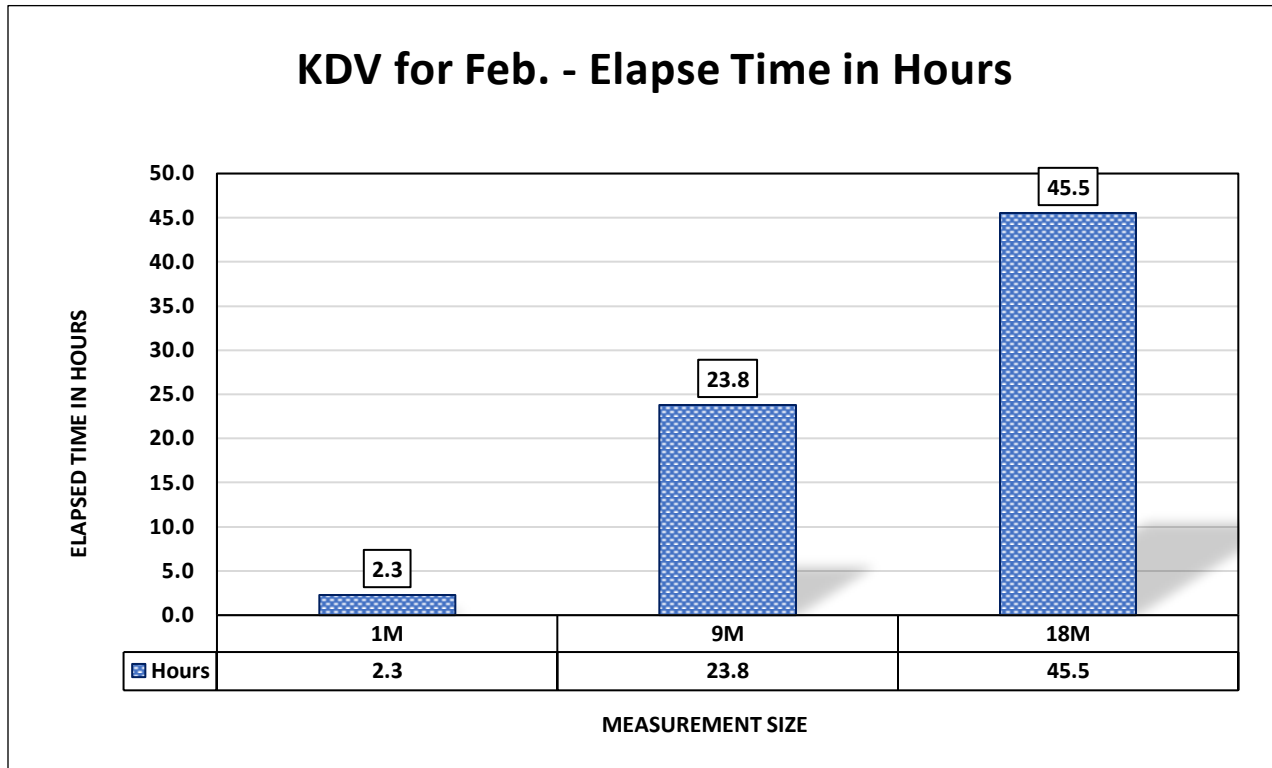


Figure 3: KDV for Feb. - Elapse Time in Hours

Taking our result data and normalizing the measurement per 1M objects shows only small deviations in per 1M policy object elapsed time processing.




Figure 4: KDV for Feb. - Elapse Time in Hours - Normalized to 1M Objects

10.1.2 Scaling Results - CPU Utilization

The Insurance Analyzer workload is dominated by the application server utilization. Because we kept the load constant with 120 parallel batch jobs, both the DB and Application Server utilization remained consistent across and throughout the 1M, 9M and 18M measurement points, while the elapsed time had the near linear 20X increase. The zEC12 DB Server is a non-SMT processor and had low utilizations. The single Application Server is a Power7 780 and had comparatively high utilization for the effort. The CPU Utilization on the application server was captured via the AIX "vmstat" command. The utilization numbers and used SAPS do not have a linear correlation. For actual sizing details, contact your IBM vendor. The figures below show KDV for February - DB Server Utilization, followed by the KDV for February - Application Server CPU Utilization.




Figure 5: KDV for Feb. - DB Server CPU Utilization

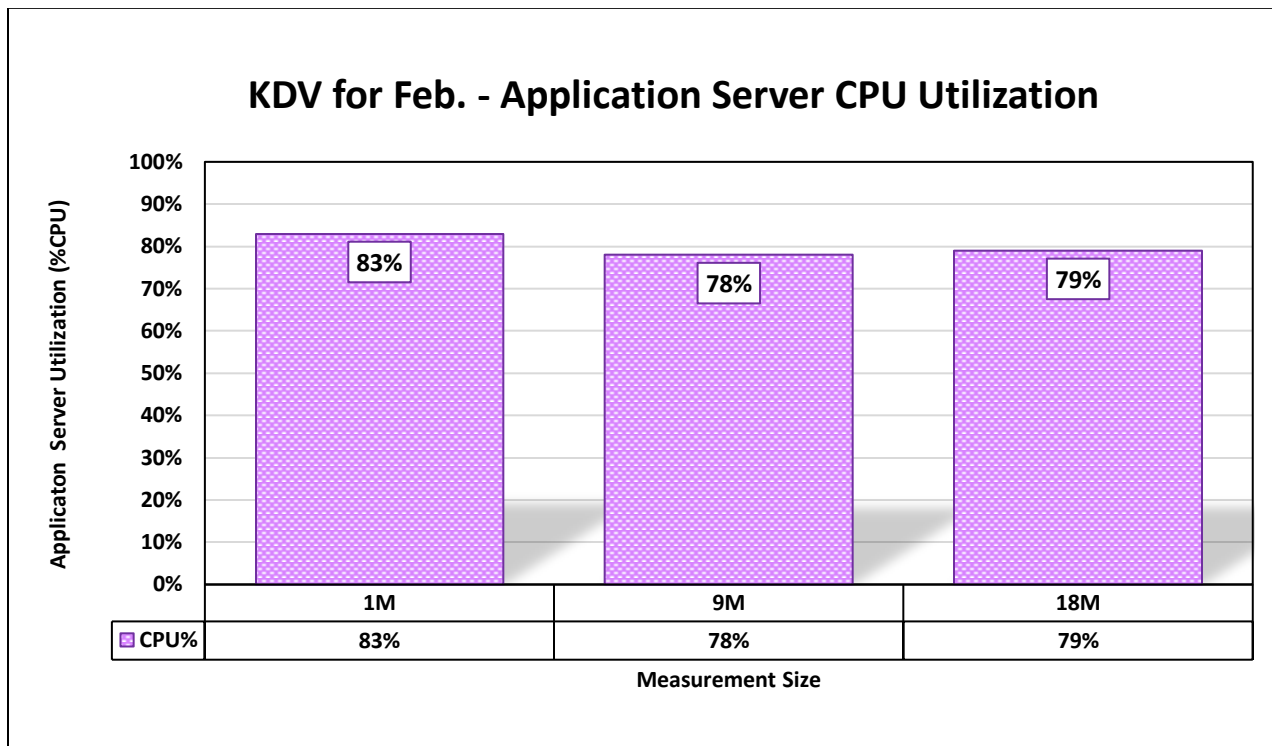


Figure 6: KDV for Feb. - Application Server CPU Utilization

10.1.3 Scaling Results - DB Growth Rate in GB

The database growth rate was captured via SAP's DB02 transaction, with before and after KDV process step comparisons. The growth in the database, whether tablespace or index is linear across the 1M, 9M and 18M policy measurements. Note that tablespace compression was done via z Systems hardware compression and indexes were not compressed.

Number of Policies	Tablespace Growth (GB)	Index Growth (GB)	Total DB Growth (GB)	Tablespace Growth Normalized per 1M (GB)	Index Growth Normalized per 1M (GB)	DB Growth Normalized per 1M (GB)
1M	53	63	116	53	63	116
9M	452	561	1,013	50	62	112
18M	932	1,127	2,059	52	63	115

Table 4: Scaling Results - DB Growth Rate in GB




Figure 7: KDV for Feb. - DB Growth Rate in GB

Normalizing the Database Growth Rate per 1M policies again shows the linear growth in both index and tablespace allocations, across the 1M, 9M and 18M policy measurements.

Number of Policies	Tablespace Growth (GB)	Index Growth (GB)	Total DB Growth (GB)	Tablespace Growth Normalized per 1M (GB)	Index Growth Normalized per 1M (GB)	DB Growth Normalized per 1M (GB)
1M	53	63	116	53	63	116
9M	452	561	1,013	50	62	112
18M	932	1,127	2,059	52	63	115

Table 5: Scaling Results - DB Growth Rate in GB - Normalized to 1M Objects




Figure 8: KDV for Feb. - DB Growth Rate in GB - Normalized to 1M Objects

10.2 KDV 1 versus 2 Application Server comparison measurements

One activity that was done to reduce elapsed time was the addition of an identical second Power7 780 application server. See the “[Hardware Landscape](#)” picture for the application server hardware and software specifications, then double the IBM Power7 780 in the picture.

10.2.1 KDV 1 vs 2 Application Server - Elapsed Time

The results of the (2) application server measurement showed a linear reduction of the runtime, cutting the elapsed time approximately in half. The application server and DB server activity was approximately doubled during the run, or about what would be expected by doubling the number of batch threads from 120 to 240 parallel threads. (2 Application Servers * 120 threads each).

Number of Policies	Number batch of threads	Runtime (hours)	DB Server Utilization	AppServer #1 Utilization	AppServer #2 Utilization
1M	120	2.3	15%	83%	na
1M	240	1.2	31%	82%	81%

Table 6: KDV 1 vs 2 Application Server - Elapsed Time

A graphical representation of these measurements shows the linear elapsed time reduction when adding the second Power7 780 application server. The 9M and 18M scenarios were not executed with 2 application servers, but for 1M we achieved a throughput that would be high enough to process 18M in less than 24 hours.

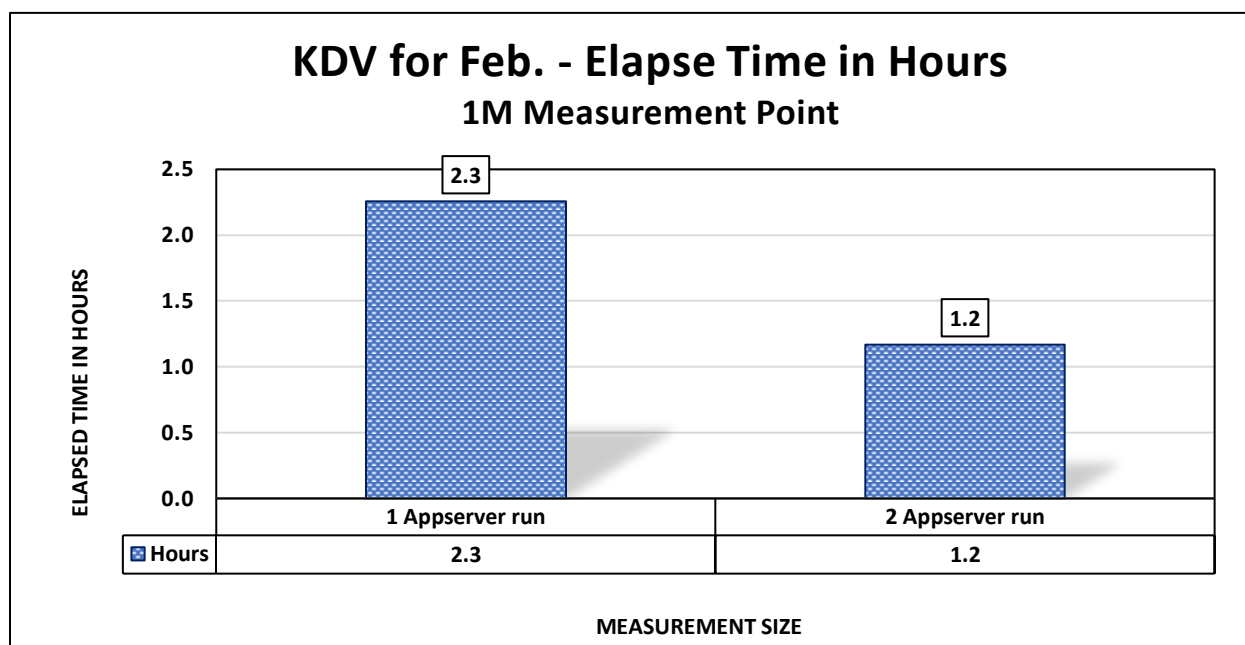


Figure 9: KDV for Feb. 1 vs 2 Application Server - Elapsed Time in hours

10.2.2 KDV 1 vs 2 Application Server - DB Server CPU Utilization

Below is a graphical representation of both the CPU utilization of both the DB Server and Application Servers. While elapsed time was cut in half, work handled by the DB Server doubled, as reflected by the increased DB Server utilization.

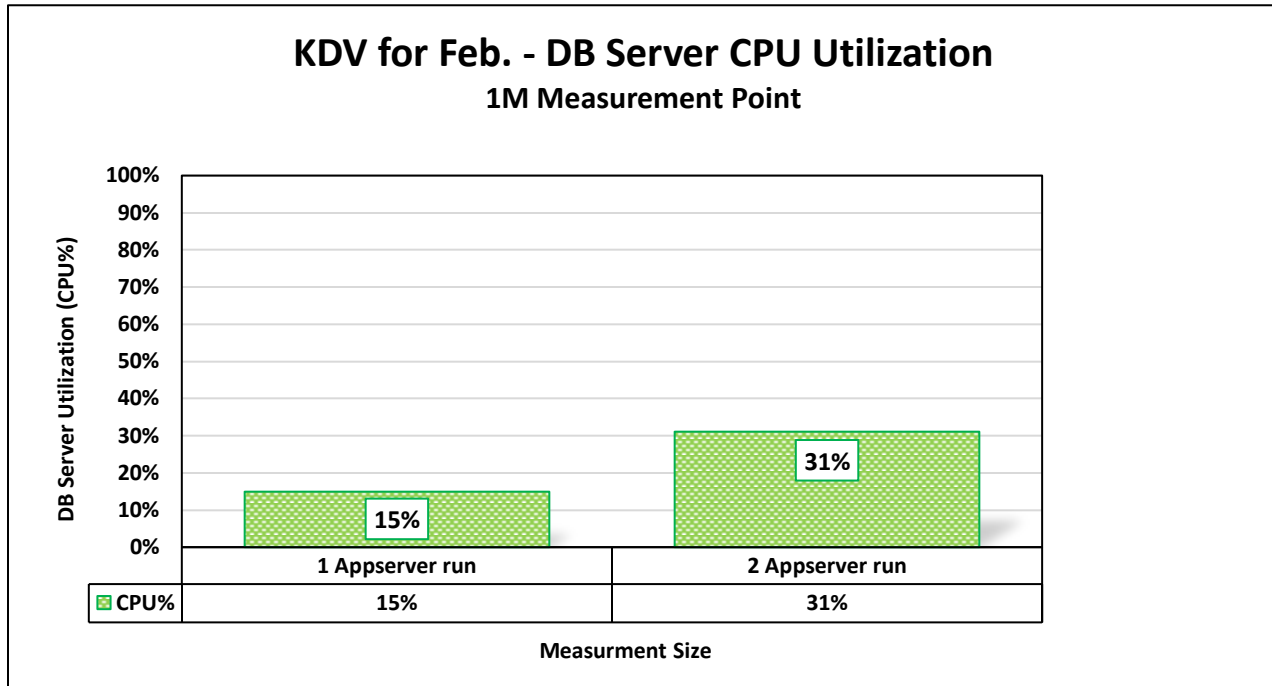


Figure 10: KDV for Feb. 1 vs 2 Application Server - DB Server CPU Utilization

10.2.3 KDV 1 vs 2 Application Server - Application Server Average CPU Utilization

The average CPU utilization across (2) application servers was close to that of (1) application server. Note the utilization shown here and SAPS do not have a linear correlation. For actual sizing details, contact your IBM Vendor.

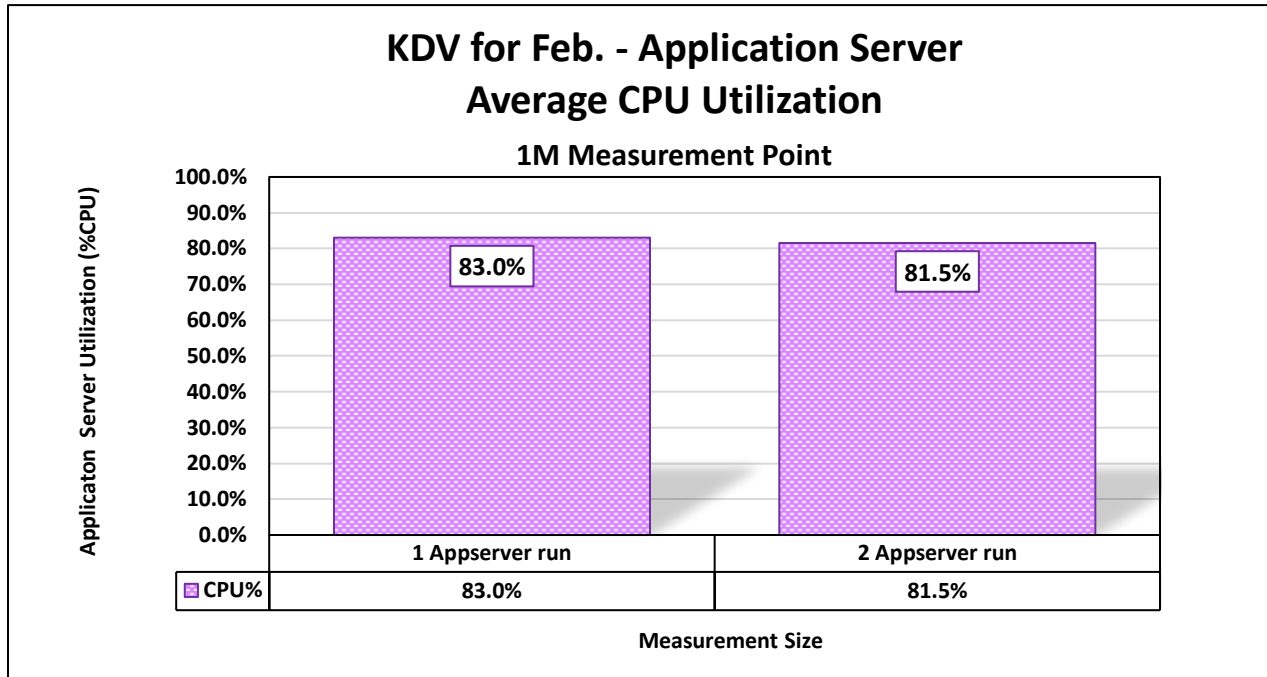


Figure 11: KDV for Feb. 1 vs 2 Application Server - Average Application Server CPU Utilization

11 Summary

The Insurance Industry is evolving. There is stiffer competition for customers using new technologies and there are new regulations meant to establish guidelines for insurers. IFRS 4 accounting principles and Solvency II risk guidelines are such regulations. With DB2 on IBM z Systems at its core, SAP Insurance Analyzer is a strong and robust solution to address these needs.

Within the SAP Insurance Analyzer solution there are many "sequential" batch process steps. Some have time-sensitive batch execution window requirements. One such process with a critical end of month time-sensitive window is the IFRS 4 accounting process called Key Date Valuation (KDV). This paper focused specifically on the KDV step running on the z Systems platform. We scaled KDV process tests to 18M policies, which is higher than currently known SAP customer environments. Additionally, we measured the impact of doubling application server capacity, successfully cutting elapsed time in half. Though the 18M scenario was not executed with 2 application servers, but for the 1M scenario, we achieved a throughput that would be high enough to process 18M in less than 24 hours.

SAP has partnered with IBM z Systems because IBM z Systems is the industry leading platform for Enterprise computing. IBM z Systems can sustain the stringent business continuity of 24x7 application availability and provide strict security for vital customer and sensitive data. IBM z Systems can accommodate varying system loads and database growth with its robust workload management and highly scalable architecture. IBM z Systems has all the underlying IT infrastructure criteria required by the insurance industry.

This study has demonstrated key strengths of DB2 on z Systems - strength such as database scalability when increasing from 1M to 18M policies and strength such as 24x7 operations as demonstrated by the long running 18M measurement. All of this is evidence of DB2 on z Systems to be the ideal database management system for the SAP Insurance Analyzer solution.

Insurers need to deliver superior service by truly connecting with customers, while delivering their product to market faster and getting the right information to the right customers at the right time. They need an underlying IT infrastructure that is particularly strong in "Business Continuity", "Security" and "Performance & Scalability". They need the robustness of DB2 on z Systems to be the database management system for the SAP Insurance Analyzer solution to best meet new regulatory guidelines.

IBM z Systems with its performance, scalability and integrated security, is the ideal platform for SAP Insurance Analyzer to satisfy the insurance industry needs.

12 References

- [1] IBM Corp. 2016. *IBM zEnterprise System* <http://www.ibm.com/systems/z/hardware/zenterprise/>
- [2] IBM Corp. 2016. *SAP on IBM z System Reference Architecture - SAP for Insurance* http://www.sap.com/documents/2015/07/ee933188-5b7c-0010-82c7-eda71af511fa.html?cm_mc_uid=93595253815714818426407&cm_mc_sid_50200000=1482350855
- [3] IBM Corp. 2011. *IBM System z: SAP for Insurance* <http://www-03.ibm.com/support/techdocs/atmastr.nsf/WebIndex/WP101890>
- [4] IBM Corp. 2014. *IBM zEnterprise System: DB2 10 for z/OS Large Objects with SAP Insurance Solutions* <http://www-03.ibm.com/support/techdocs/atmastr.nsf/WebIndex/TD106262>
- [5] IBM Corp. 2016. *SAP on IBM z Systems Wiki > Reference Architectures* [https://www.ibm.com/developerworks/community/wikis/home?lang=en-!/wiki/W52acd0d6cb4d_4532_ac8c_2a5a8b0529b9/page/Reference Architectures](https://www.ibm.com/developerworks/community/wikis/home?lang=en-!/wiki/W52acd0d6cb4d_4532_ac8c_2a5a8b0529b9/page/Reference%20Architectures)
- [6] IBM Corp. 2016. *DB2 11 for z/OS* <http://www-01.ibm.com/software/data/db2/zos/family/db211/index.html>
- [7] IBM Corp. 2013. *IBM zEnterprise System: DB2 11 for z/OS with SAP Performance Report* <http://www-03.ibm.com/support/techdocs/atmastr.nsf/WebIndex/WP102394>
- [8] IBM Corp. 2013. *IBM zEnterprise System and System Storage Server DS8870: SAP Bank Analyzer 8.0 AFI Current Accounts Tests* <http://www-03.ibm.com/support/techdocs/atmastr.nsf/WebIndex/WP102354>
- [9] IBM Corp. 2014. *IBM zEnterprise System: SAP Bank Analyzer 8.0 AFI Loan Tests* <http://www-03.ibm.com/support/techdocs/atmastr.nsf/WebIndex/WP102206>