Redbooks
ibm.com/redbooks

IBM z Systems
Performance Report on Exploiting SMT for SAP Application Servers on z13

Brenda Beane

Seewah Chan

Dr. Paul Lekkas

Veng Ly

Storage

In partnership with
IBM **Academy of Technology**

IBM

Redpaper

Introduction

With the recent availability of the IBM z13™ machine, the new processor chip offers intelligently implemented 2-way simultaneous multithreading. Simultaneous multithreading (SMT) allows two active instruction streams per core, each dynamically sharing the core's execution resources. The SMT feature is available in IBM® z13 for workloads running on the Integrated Facility for Linux (IFL) and the z Integrated Information Processor (zIIP).

SAP on IBM z Systems™ is an excellent candidate in employing zIIP processors for the SAP database server, and IFL processors for SAP application servers. Therefore, it can exploit SMT and potentially see performance gains from using it.

There are several options for running SAP application servers with a z System SAP database server. One of the more intriguing options is running the SAP application server on a z System with Linux. This option is of particular interest because it exploits both the hardware and virtualization capabilities intrinsic to z Systems, in which case, customers have the options to run these application server functions on native Linux LPARs or as Linux guests on IBM z/VM® for further virtualization and management capabilities.

z/VM provides significant virtualization support for Linux virtual servers, including exploitation of the capabilities of z Systems, non-disruptive dynamic addition of memory and processors, improved systems management, ease-of-use enhancements, performance improvements for Linux guests, and enhanced networking for guests. See *IBM z/VM V6.3 - Virtualization with Efficiency at Scale: z/VM Version 6 Release 3: Frequently Asked Questions:*

http://www.vm.ibm.com/faq/faq63.pdf

The IBM SAP on z Systems Performance Team, located in Poughkeepsie, NY, US, conducted a number of experiments to evaluate the performance effects of enabling SMT on IFLs for SAP application servers, running on Linux as z/VM guests. At the time of our study, the SMT feature was not available for Linux running as native LPARs.

We used the SAP Banking Services (SBS) Day Posting workload, which is a good representation of a customer online transaction processing (OLTP) workload. On the application server side, this workload is CPU-intensive with minimal file I/O activity. For details about this workload, see “Workload description” on page 3.

We started with baseline measurements on a zEC12 system. Then we ran on a new z13 system for core to core comparison and SMT benefit evaluation.

This IBM Redpaper™ publication documents our tests and findings. The measurements were stress tests, not certified benchmarks.

Executive summary

The z13 machine offers a vast array of functional and performance features that are well suited for the SAP business solutions on z Systems. Simultaneous multithreading (SMT) is one of the key performance features on the z13 machine. This hardware feature is available on both IFL and zIIP specialty engines.

Our measurements showed up to a 40% internal throughput rate (ITR) improvement, with SMT enabled for 16 IFL engines configured in a z/VM LPAR running SAP application servers as Linux guests as compared to zEC12.

> **ITR and ETR:** In this document, the terms internal throughput rate (ITR) and external throughput rate (ETR) are used. ETR is the transaction rate. ITR is the ETR normalized to 100% CP utilization and it gives a relative CPU time per transaction.

We ran the SAP Banking Services (SBS) Day Posting workload, which represents typical OLTP processing in an SAP banking customer environment. Both the SAP application server and the SAP database server were running in one z System CEC. The SAP application servers were configured as three Linux guests under a single z/VM LPAR with 16 IFLs. The SAP database server was in a separate IBM z/OS® LPAR with two general purpose CPs and two zIIPs. The focus of this study was the SAP application servers. Three measurements at approximately 90% CPU utilization were completed, which included the zEC12 baseline, z13 without SMT enabled, and finally z13 with SMT enabled.

For a typical SAP customer workload, the processing capacity needed for the application servers is generally several multiples of that needed by the database servers. The larger SAP customer installations can exceed hundreds of processor cores for their SAP application servers. For multi-threaded OLTP workloads such as SAP Day Posting, the more and faster cores combined with the SMT boost in the z13 machine can readily accommodate the customer's growing demand for capacity required for SAP solutions on z Systems. In addition, the cost of consolidating work on Linux under z/VM with the z13 requires less floor space, power or cooling, translating into energy and facilities cost saving.

Workload description

The SAP Banking Services Day Posting workload was used in these tests. It is an OLTP workload. We have been running this workload for many years and have much experience with it. See *IBM System zEnterprise®, System Storage, and DB2 10 for z/OS: SAP Banking Services 7.0 150 Million Accounts Measurements*

http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP101978

In this workload, a posting is a deposit or a withdrawal from a customer's account. Typical examples of a posting are a payment out of the account or a deposit into the account. This workload was developed by SAP to simulate customer environments. The workload consists of interactive "users" going through repetitive cycles of 15 dialogue steps (Table 1).

Table 1 SAP Banking Day Posting workload

Step	Operation
1	Create a total of 150 postings through five BAPI calls.
2	Create five postings.
3	Create bank statement.
4	Read postings for account.
5	Read details of postings.
6	Create five postings.
7	Create one bank statement for account.
8	Create five postings.
9	Create one bank statement for account.
10	Create five payment orders.
11	Read balances of account.
12	Create five postings.
13	Create one bank statement for account.
14	Read balances for account.
15	Read master data for account.

Test environment

Figure 1 shows an overview of the test environment with z13.

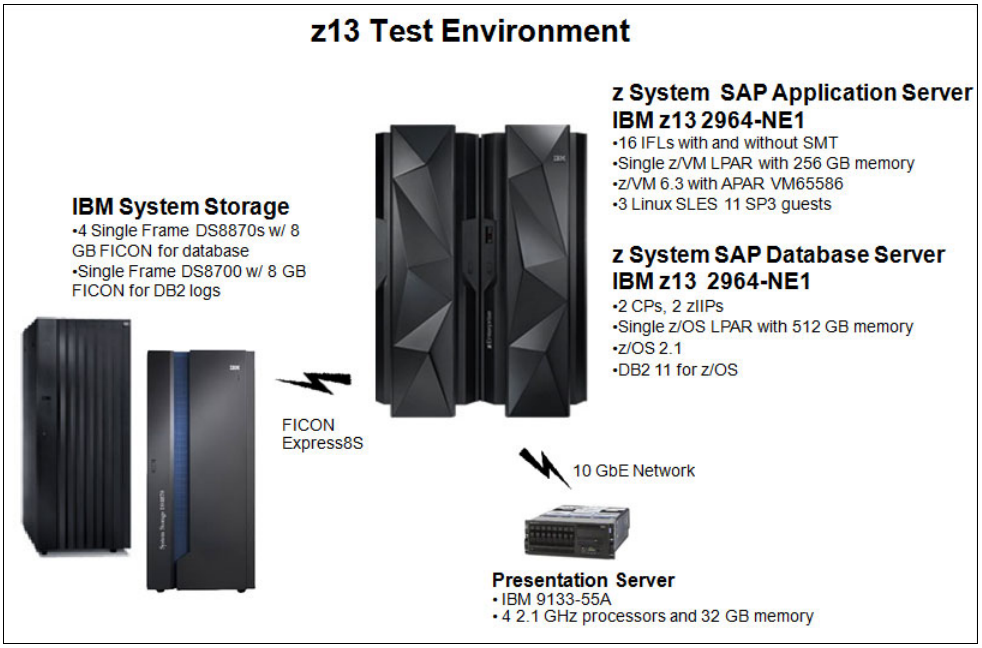

Figure 1 The z13 test environment

Figure 2 shows the baseline test environment with system zEC12. This is the same test environment as the z13 test environment except for the z System CEC.

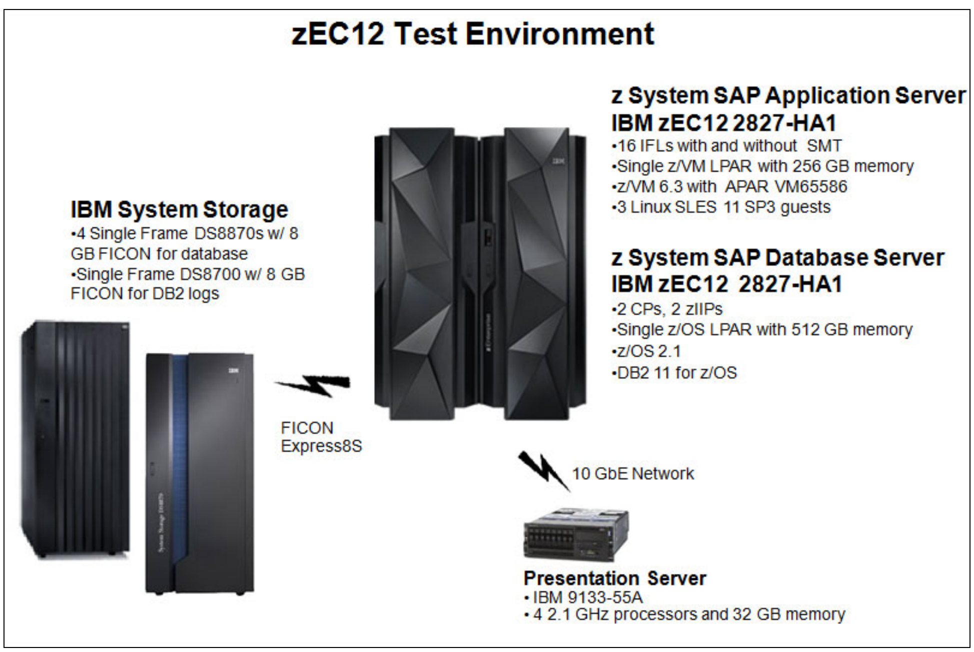

Figure 2 The zEC12 test environment

The test environment was a physical 2-tier configuration that used only one z System CEC. Both the SAP application servers and the database server resided in the same physical CEC.

Figure 3 shows our logical test configuration on z System.

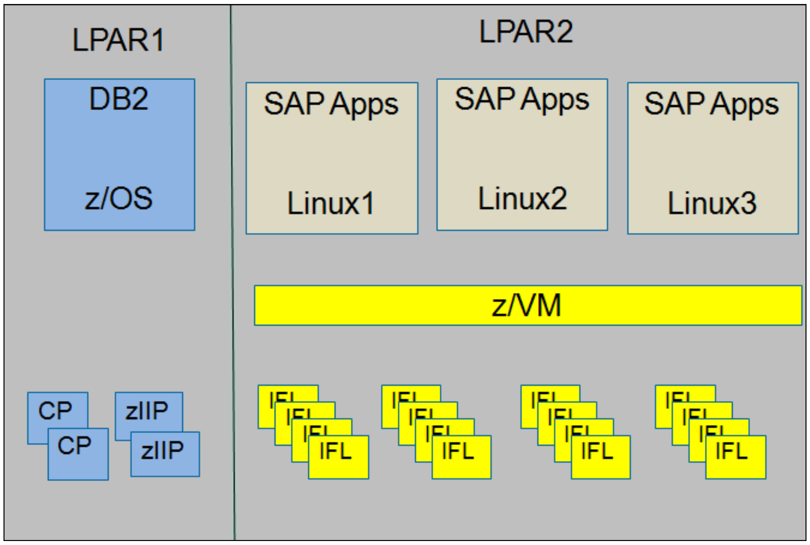

Figure 3 z System LPAR configuration

The SAP application servers resided on a z System LPAR, running as the three Linux guests under a single z/VM LPAR with 16 IFLs, as illustrated in Figure 3. The first Linux guest was configured as an SAP Message and Enqueue Server, with four logical CPs and 8 GB memory. It had one SAP instance. The second Linux guest was configured as a SAP Central and Dialog Instance server, with 16 logical CPs and 24 GB memory. It had six SAP instances. The third Linux guest was configured as a SAP dialog instance server, with 16 logical CPs and 24 GB memory. It had six SAP dialog instances.

The database server resided on another z System LPAR, on the same CEC as the LPAR where the SAP application servers resided. It was configured with two CPs and two zIIPs.

zIIP engines: zIIP engines can also exploit SMT like IFL engines can. However, for our study, we did not exploit SMT on the zIIPs. We focused on exploiting SMT on IFL engines.

Hardware

We used the following hardware.

z System SAP Database Server/ SAP Application Servers

An IBM z™ Systems 2964-NE1 (z13) or 2871-HA1(zEC12) was used depending on the test scenario. The SAP database server was in an LPAR with two dedicated CPs, two dedicated zIIPs, and 512 GB of memory. The SAP application servers were running as three Linux SLES 11 SP3 guests on a z/VM LPAR. The z/VM LPAR was configured with 16 dedicated IFLs and 256 GB memory.

Database DASD

An SBS 7.0 database with 60 million accounts was used for these tests. It resided on four single frame IBM System Storage® DS8870 (2423-961) servers, each with 16 ranks of 240 15K-RPM 300 GB hard disk drives (HDDs) configured as RAID 5. Each frame has a total of 72 TB. The unit has 256 GB regular cache, 8 GB non-volatile storage (NVS), and 16 long wave IBM FICON® Express8S attachments. The IBM DB2® subsystem, the database, and two flash copies were contained on 834 emulated 3390-mod54 volumes.

The DB2 active logs resided on a separate DASD unit from the database. The active logs were striped across four 3390-mod54 volumes on separate ranks of a single frame IBM System Storage DS8700 (2107-941) server.

Presentation Server

One IBM 9133-55A server with four 2.1 GHz processor cores and 32 GB of memory running IBM AIX® was used as the presentation server to drive the workload.

Network

A dedicated 10 Gb Ethernet network was used to connect the presentation server to the application servers, and the database server to the application servers. Four shared OSA-Express4S adapters were used to connect the SAP DB server and the SAP application servers. The Optimized Latency Mode (OLM) option of the OSA-Express4S adapters was used to improve the elapsed time of this communication. IBM HiperSockets™ is recommended for network-intensive batch workloads. We did not use them because our workload was an interactive (OLTP) type.

Software

We used the following software:

- z/OS

  z/OS release 2.1
- DB2 for z/OS

  DB2 11
- IBM DB2 Connect™

  IBM Data Server Driver for CLI that is shipped as part of DB2 Connect 10.1 FP2
- z/VM

  z/VM 6.3 with APAR VM65586
- Linux

  SLES 11 SP3
- SAP
  - SAP NetWeaver 7.1 Enhancement Package 1
  - SAP kernel level 720 EXT, 64-bit, patch number 400

Measurement results

We executed three measurements:

- zEC12 baseline
- z13 without SMT enabled
- z13 with SMT enabled

Our methodology was to hold the processor utilization consistent at approximately 90% across the measurements. This is to stress the system, and to see the effects of the z13 and SMT on throughput. The processor utilization (which is listed as “Average %CPU on application servers” in Table 2) is the percent utilization of the z/VM LPAR where all our Linux guests were running application servers. Specifically, it was from the PERFKIT Reports, see the “Understanding z/VM CPU Utilization” web page:

http://www.vm.ibm.com/perf/tips/lparinfo.html

The details of these measurements are summarized in Table 2.

Table 2 Measurement results

Run ID	S41030B1	S41107B2	S41114B2
z System	zEC12	z13	z13
SMT enabled	N/A	No	Yes
Number of z/OS LPARs	1	1	1
Number of processors per z/OS LPAR	2 CPs + 2 zIIPs	2 CPs + 2 zIIPs	2 CPs + 2 zIIPs
Real Storage Configured per z/OS LPAR	512 GB	512 GB	512 GB
z/OS Level	2.1	2.1	2.1
Number of z/VM LPARs	1	1	1
Number of processors per z/VM LPAR	16 IFL	16 IFL	16 IFL
Real Storage Configured per z/VM LPAR	256 GB	256 GB	256 GB
z/VM Level	6.3 APAR VM65586	6.3 APAR VM65586	6.3 APAR VM65586
Number of Users	1968	2240	2856
Average %CPU on application servers	92.30%	91.00%	90.50%
ETR (DS/sec)	175.34	188.29	241.69
ITR (DS/sec) on application server	189.97	206.91	267.06

Figure 4 on page 8 shows the external throughput rate (ETR), illustrating the processing capacity improvement with z13 and SMT. Our results showed that SAP application servers running as Linux guests under z/VM on z13 with SMT can process a 38% higher ETR when compared with a zEC12 machine configured with the same number of 16 IFLs. For core-to-core comparison, at the same processor utilization, without SMT enabled, the z13 can process a 7% higher ETR as compared to zEC12.

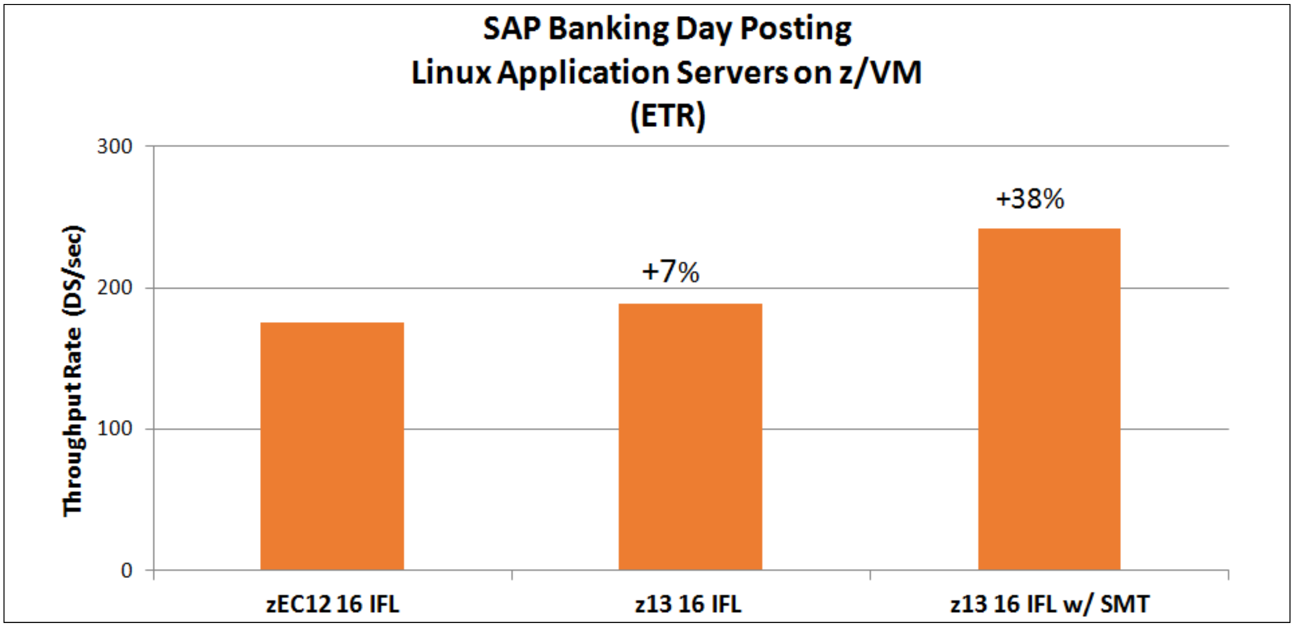

Figure 4 Processing capacity improvement with z13 and SMT

Figure 5 shows z13 and SMT effects on ITR. Again, ITR is the ETR normalized to 100% CP utilization and it gives a relative CPU time per transaction. Our results showed that SAP application servers running as Linux guests under z/VM on z13 with SMT can gain up to 40% in transactions per CPU second when compared with zEC12 machine configured with the same number of 16 IFLs. For core-to-core comparison, z13 showed a 9% ITR gain as compared to zEC12.

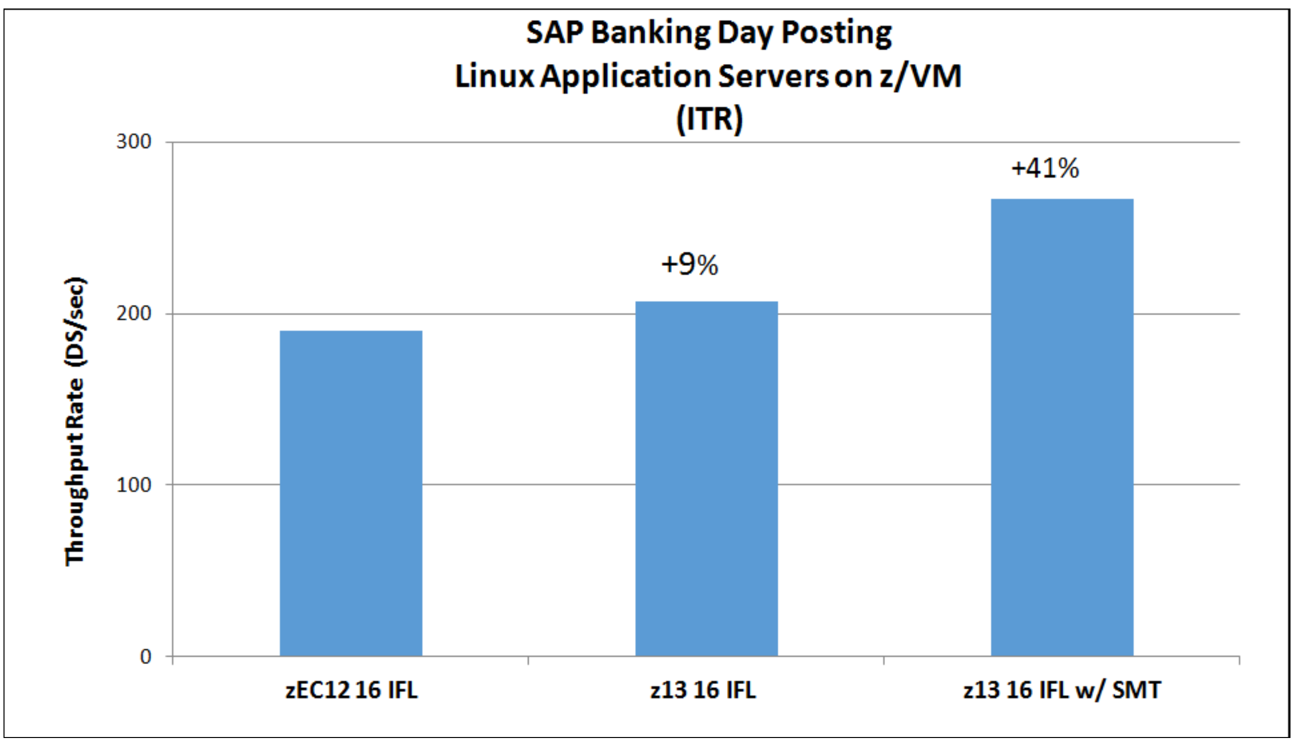

Figure 5 z13 and SMT effects on ITR

Conclusions

SMT is one of the many key features in the IBM z13™. It can provide up to a 40% ITR improvement per core as compared to zEC12, when SAP application servers are running as Linux guests under z/VM on the IBM z13.

The performance benefits of SMT can vary depending on the characteristics of the workload. Very highly computation-intensive batch workloads that require total core resource for each thread might not see significant performance benefits with SMT. The capacity and throughput gain per core depends on the overlap and interference between two threads. An overlap can result from many core resources being replicated so that each thread can make progress, or while one thread waits for cache miss, other threads can continue to run. Interference can be from some serialization points within the core. It can be threads sharing the same caches, thus cache misses can increase, or cause contention. See *z Systems Simultaneous Multithreading Revolution*, REDP-5144:

http://www.redbooks.ibm.com/abstracts/redp5144.html

The IBM z13 has more capacity per core without requiring appreciably more space, power, or cooling than the previous generation processor with single-threaded cores. It is designed to offer the capacity and processing power to improve business performance and growth demanded by explosive amounts of analytics data and rapid proliferation of mobile applications and transactions, estimated to grow to 40 trillion transactions per day by 2025. See *The Technology Economics of the Mainframe, Part 3: New Metrics and Insights for a Mobile World*, by Dr. Howard A. Rubin, January 2015

http://www.ibm.com/common/ssi/cgi-bin/ssialias?infotype=SA&subtype=WH&htmlfid=ZSLO3354USEN#loaded

References

For more information, see the following resources:

- *SAP on IBM z Systems® Reference Architecture: SAP for Banking*

  http://scn.sap.com/docs/DOC-14357
- Large Systems Performance Reference for IBM z Systems

  http://www.ibm.com/servers/resourcelink/lib03060.nsf/pages/lsprindex?OpenDocument
- *IBM DB2 11 for z/OS Technical Overview*, SG24-8180

  http://www.redbooks.ibm.com/Redbooks.nsf/RedpieceAbstracts/sg248180.html?Open
- *IBM System zEnterprise®, System Storage, and DB2 10 for z/OS: SAP Banking Services 7.0 150 Million Accounts Measurements*

  http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP101978
- *The Virtualization Cookbook for IBM z Systems Volume 1: IBM z/VM 6.3*, SG24-8147

  http://www.redbooks.ibm.com/abstracts/sg248147.html
- *The Technology Economics of the Mainframe, Part 3: New Metrics and Insights for a Mobile World*, by Dr. Howard A. Rubin, January 2015

  http://www.ibm.com/common/ssi/cgi-bin/ssialias?infotype=SA&subtype=WH&htmlfid=ZSLO3354USEN#loaded

- *IBM z/VM V6.3 - Virtualization with Efficiency at Scale: z/VM Version 6 Release 3: Frequently Asked Questions*

  http://www.vm.ibm.com/faq/faq63.pdf

- *z Systems Simultaneous Multithreading Revolution*, REDP-5144

  http://www.redbooks.ibm.com/abstracts/redp5144.html

- Understanding z/VM CPU Utilization

  http://www.vm.ibm.com/perf/tips/lparinfo.html

Authors

This paper was produced by a team of specialists from around the world working at the International Technical Support Organization (ITSO), Poughkeepsie Center.

Brenda Beane is a Senior Software Engineer at IBM in Poughkeepsie, NY. She has 33 years of experience in development and test at IBM. For the last 19 years, she has been a member of the SAP on z Systems performance test team specializing in DB2 for z/OS. Brenda holds a Bachelor of Science degree in Computer Science and Mathematics from the University of Massachusetts and a Master's degree in Computer Science from Marist College.

Seewah Chan is a Senior Software Engineer at IBM in Poughkeepsie, NY. He has 33 years of experience in development and test at IBM. For the last 19 years, he has been a member of the SAP on z Systems performance test team specializing in DB2 for z/OS. Seewah holds a Bachelor of Science degree in Physics, Mathematics, and Computer Science from the State University of New York at Albany.

Dr. Paul Lekkas is an IBM Systems Group Distinguished Engineer. He has 32 years of experience in development and test of HW and SW at IBM. For the first several years he worked as hardware developer doing chip logic and ASIC design. Since 1989, he has been member of the technical community support pre-sales teams and customers to implement complex business solutions. He joined the SAP team in 1998 and today he is the leader of the SAP performance team. Paul holds a Masters of Science in Physics and Mathematics and a PhD in Nuclear Physics.

Veng Ly is an Engineer at IBM Poughkeepsie, NY. He has 26 years of experience in development and test at IBM. For the first several years, he worked as a Hardware Engineer doing chip logic and microprocessor design for IBM Large Systems. For the last 17 years, he has been a member of the SAP on z Systems test team specializing in DB2 for z/OS. Veng holds a Bachelor of Electrical Engineering degree from the City College of the City University of New York, and a Masters of Science in Computer & Telecommunications Management from the New York University Polytechnic School of Engineering.

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

- Ernie Balloni, Andrea Fuga, Tom Litty, and Joe Stein for system programming support
- Bill Norton and Brian Wade for z/VM support

Now you can become a published author, too!

Here's an opportunity to spotlight your skills, grow your career, and become a published author—all at the same time! Join an ITSO residency project and help write a book in your area of expertise, while honing your experience using leading-edge technologies. Your efforts will help to increase product acceptance and customer satisfaction, as you expand your network of technical contacts and relationships. Residencies run from two to six weeks in length, and you can participate either in person or as a remote resident working from your home base.

Find out more about the residency program, browse the residency index, and apply online at:

ibm.com/redbooks/residencies.html

Stay connected to IBM Redbooks

- Find us on Facebook:

  http://www.facebook.com/IBMRedbooks

- Follow us on Twitter:

  http://twitter.com/ibmredbooks

- Look for us on LinkedIn:

  http://www.linkedin.com/groups?home=&gid=2130806

- Explore new Redbooks publications, residencies, and workshops with the IBM Redbooks weekly newsletter:

  https://www.redbooks.ibm.com/Redbooks.nsf/subscribe?OpenForm

- Stay current on recent Redbooks publications with RSS Feeds:

  http://www.redbooks.ibm.com/rss.html

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:
IBM Director of Licensing, IBM Corporation, North Castle Drive, Armonk, NY 10504-1785 U.S.A.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of 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.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

This document REDP-5287-00 was created or updated on October 9, 2015.

Send us your comments in one of the following ways:

- Use the online **Contact us** review Redbooks form found at:
  ibm.com/redbooks
- Send your comments in an email to:
  redbooks@us.ibm.com
- Mail your comments to:
  IBM Corporation, International Technical Support Organization
  Dept. HYTD Mail Station P099
  2455 South Road
  Poughkeepsie, NY 12601-5400 U.S.A.

Trademarks

REDP-5287-00

ISBN 0738454524

Printed in U.S.A.

Get connected