

z/VSE Performance Update

Ingo Franzki, IBM



Trademarks

The following are trademarks of the International Business Machines Corporation in the United States, other countries, or both.

Not all common law marks used by IBM are listed on this page. Failure of a mark to appear does not mean that IBM does not use the mark nor does it mean that the product is not actively marketed or is not significant within its relevant market.

Those trademarks followed by ® are registered trademarks of IBM in the United States; all others are trademarks or common law marks of IBM in the United States.

For a complete list of IBM Trademarks, see www.ibm.com/legal/copytrade.shtml:

*, AS/400®, e business(logo)®, DBE, ESCO, eServer, FICON, IBM®, IBM (logo)®, iSeries®, MVS, OS/390®, pSeries®, RS/6000®, S/30, VM/ESA®, VSE/ESA, WebSphere®, xSeries®, z/OS®, zSeries®, z/VM®, System i, System i5, System p, System p5, System x, System z, System z9®, BladeCenter®

The following are trademarks or registered trademarks of other companies.

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries. Cell Broadband Engine is a trademark of Sony Computer Entertainment, Inc. in the United States, other countries, or both and is used under license therefrom.

Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

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

Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

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

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

ITIL is a registered trademark, and a registered community trademark of the Office of Government Commerce, and is registered in the U.S. Patent and Trademark Office.

IT Infrastructure Library is a registered trademark of the Central Computer and Telecommunications Agency, which is now part of the Office of Government Commerce.

* All other products may be trademarks or registered trademarks of their respective companies.

Notes:

Performance is in Internal Throughput Rate (ITR) ratio based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput improvements equivalent to the performance ratios stated here.

IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply.

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.

All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

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.

Prices subject to change without notice. Contact your IBM representative or Business Partner for the most current pricing in your geography.

Notice Regarding Specialty Engines (e.g., zIIPs, zAAPs and IFLs):

- § Any information contained in this document regarding Specialty Engines ("SEs") and SE eligible workloads provides only general descriptions of the types and portions of workloads that are eligible for execution on Specialty Engines (e.g., zIIPs, zAAPs, and IFLs). IBM authorizes customers to use IBM SE only to execute the processing of Eligible Workloads of specific Programs expressly authorized by IBM as specified in the "Authorized Use Table for IBM Machines" provided at http://www.ibm.com/systems/support/machine_warranties/machine_code/aut.html ("AUT").
- § No other workload processing is authorized for execution on an SE.
- § IBM offers SEs at a lower price than General Processors/Central Processors because customers are authorized to use SEs only to process certain types and/or amounts of workloads as specified by IBM in the AUT.

Disclaimer

The information contained in this document has not been submitted to any formal IBM test and is distributed on an "AS IS" basis without any warranty either express or implied. The use of this information or the implementation of any of these techniques is a customer responsibility and depends on the customer's ability to evaluate and integrate them into the operational environment. While each item may have been reviewed by IBM for accuracy in a specific situation, there is no guarantee that the same or similar results will be obtained elsewhere. Customers attempting to adapt these techniques to their own environments do so at their own risk.

In this document, any references made to an IBM licensed program are not intended to state or imply that only IBM's licensed program may be used; any functionally equivalent program may be used instead.

Any performance data contained in this document was determined in a controlled environment and, therefore, the results which may be obtained in other operating environments may vary significantly. Users of this document should verify the applicable data for their specific environments.

It is possible that this material may contain reference to, or information about, IBM products (machines and programs), programming, or services that are not announced in your country. Such references or information must not be construed to mean that IBM intends to announce such IBM products, programming or services in your country.

Agenda

§ General Considerations

§ z/VSE V4.3 Performance Considerations

- GETVIS constraint relief
- Dynamic start/stop of CPUs
- Queue-I/O Assist (QIOASSIST)
- Crypto Express3 and AP queue interrupt support


§ z/VSE V5.1 Performance Considerations

- 64 bit virtual
- Standalone Dump

§ Sizing a System for z/VSE

§ z/VM and Linux considerations

§ Performance Measurement Tools



Supported VSE Releases

VSE Release	Available	End of Marketing	End of Service
z/VSE 5.1	11/25/2011		
z/VSE 4.3	11/26/2010	06/25/2012	05/31/2014
z/VSE 4.2	10/17/2008	11/26/2010	10/31/2012
z/VSE 4.1	03/16/2007	10/17/2008	04/30/2011 (out of service)
z/VSE 3.1	03/04/2005	05/31/2008	05/31/2009 (out of service)
VSE/ESA 2.7	03/14/2003	09/29/2005	08/2007 (out of service)
VSE/ESA 2.6	12/14/2001	06/30/2000	06/2006 (out of service)
VSE/ESA 2.5	09/29/2000	11/26/2010	03/2003 (out of service)
VSE/ESA 2.4	06/25/1999	06/30/2000	02/2002 (out of service)
VSE/ESA 2.3	07/12/1997	06/30/2000	12/31/2001 (out of service)

Attention:
 z/VSE V4.2, CICS/VSE V2.3, DL/I V1.10 and DL/I V1.11 was withdrawn from service effective at 10/31/2012.
 z/VSE V4.2 was the last release that offered CICS/VSE V2.3 and DL/I V1.10.
 CICS/VSE V2.3 and DL/I V1.10 are not included in any current or future version or release of z/VSE.

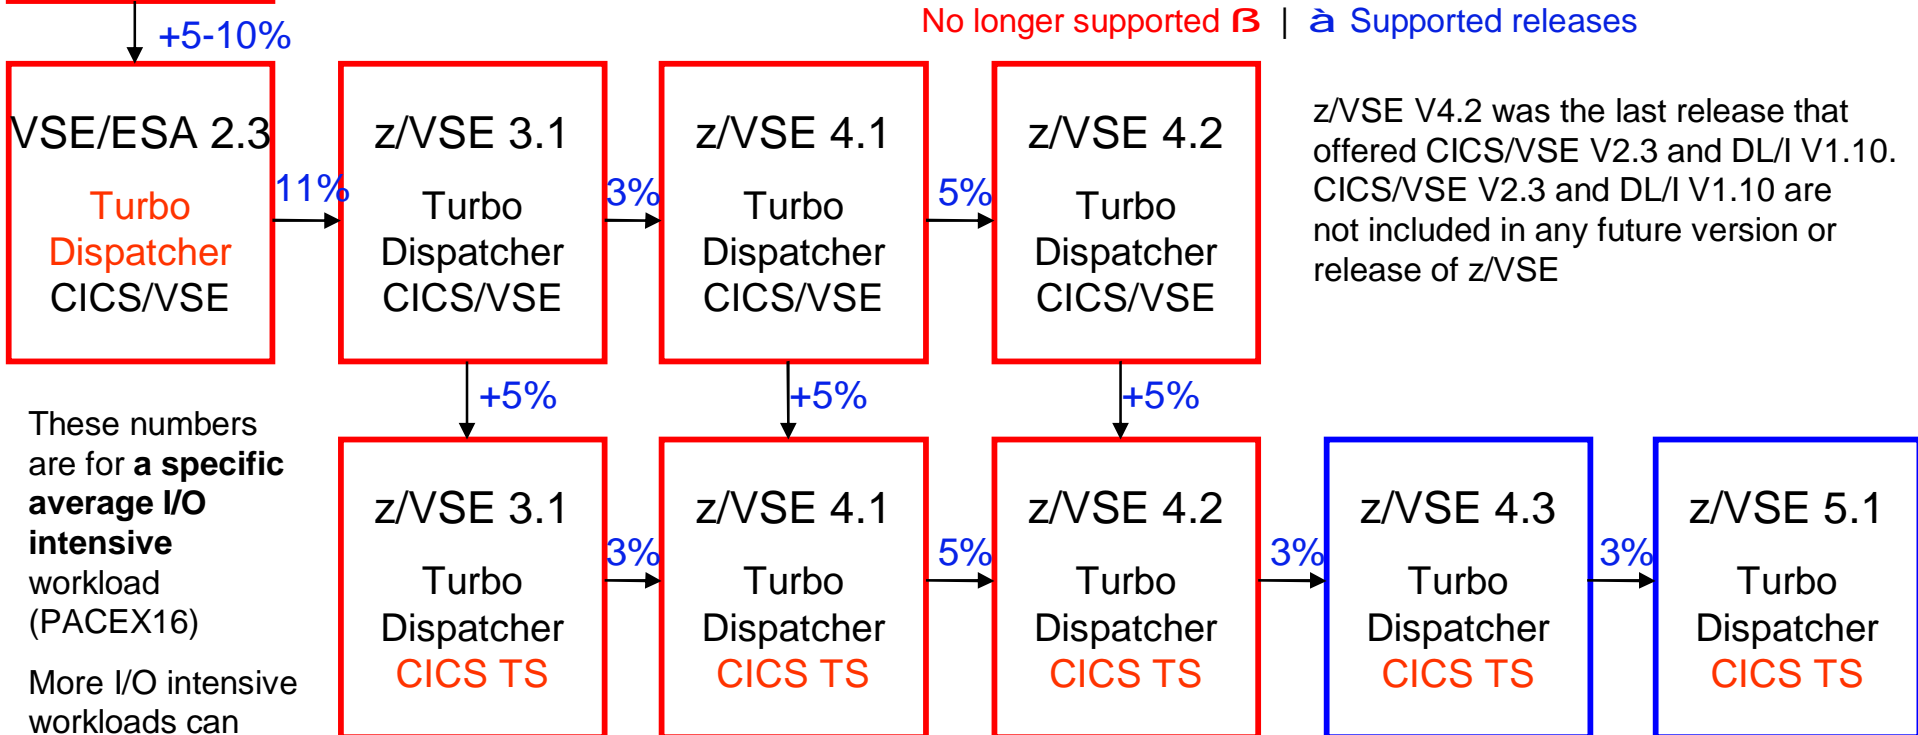
<http://www.ibm.com/systems/z/os/zvse/about/status.html>

Overhead Deltas for VSE Releases

VSE/ESA 2.3
Standard
Dispatcher
CICS/VSE

Remember that you get a lot of new functions that in most cases helps you to increase VSE system performance and throughput:

Partition Balancing, PRTY SHARE (Turbo Dispatcher), FlashCopy, Buffer Hashing, Shared data Tables (CICS TS), NOPDS with larger VSIZE



These numbers are for a **specific average I/O intensive workload (PACEX16)**
More I/O intensive workloads can cause more overhead

You can also use the zSoftCap tool to determine release migration overhead:

<http://www-03.ibm.com/support/techdocs/atmastr.nsf/WebIndex/PRS268>



New: zSoftCap Tool

You can use the [zSoftCap tool](http://www-03.ibm.com/support/techdocs/atmastr.nsf/WebIndex/PRS268) to determine release migration overhead:
<http://www-03.ibm.com/support/techdocs/atmastr.nsf/WebIndex/PRS268>



zSoftCap is a PC-based productivity tool designed to assess the effect on capacity for IBM System z processors, when migrating to more current releases of the operating system or major subsystems. zSoftCap assumes that hardware remains constant while software releases change.

The screenshot shows the 'Software Migration Scenario' window. It includes a menu bar (File, Help), a toolbar with icons for NEW, HTML, a person, and a question mark, and the version 'zSoftCap V4.3'. The main area is titled 'Define Current and Future Software Environment'. It shows 'Processor Family' as 'z196/700' and 'GP CPs' as '2', with 'Processor Model' '2817-702'. Below this are two tabs: 'z/OS' (selected) and 'z/VSE'. The interface is split into 'Current' (blue background) and 'Future' (green background) sections. In the 'Current' section, 'VSE' is 'z/VSE-4.1', 'Dispatcher' is 'Turbo', and 'CICS' is 'CICS/TS'. In the 'Future' section, 'VSE' is 'z/VSE-5.1', 'Dispatcher' is 'Turbo', and 'CICS' is 'CICS/TS'. At the bottom of the main area, 'Current Utilization' is shown as '60%' and 'Utilization After Upgrades' is '66.6%', with a 'Benefit(+)' of '-11.0%' (shown in red). A disclaimer at the bottom states: 'IBM does not guarantee the results from this tool. This information is provided "as is", without warranty, expressed or implied. You are responsible for the results obtained from your use of this tool.'

z/VSE V4.3 – GETVIS constraint relief

§ I/O Constraint Relief:

- The z/VSE I/O supervisor routines will run in AMODE(31)
- Depending on the IODEV statement, the control blocks will be allocated either in 24-bit area or in 31-bit area:
 - Specifying `IODEV=1023` will result in an allocation of the control blocks *below the line* (24-bit area)
 - Specifying `IODEV=1024` will result in an allocation of the control blocks *above the line* (31-bit area)

à Note: in either way the z/VSE limit of 1024 devices does apply and is as well maintained !

– *Example:*

```
...  
BG 0000 $$A$SUPI, VSIZE=264M, VIO=512K, VPOOL=64K, LOG, IODEV=1024  
...
```

- The z/VSE 4.3/5.1 system is shipped with IODEV=1024
- If you FSU to z/VSE 4.3/5.1, the value remains IODEV=1023

z/VSE V4.3 – Dynamic Starting/Stopping of CPUs

§ **z/VSE allows you to start CPUs that were not online at IPL**

§ **These can be either:**

- CPUs that were in a “standby” state at IPL
- CPUs that were added to the z/VSE LPAR profile after IPL

§ Note that z/VSE uses the term CPU to refer to Central Processors (CPs).

§ **Using the HMC/SE “Logical Processor Add” task, you can add CPUs to an LPAR after IPL. This feature is available on IBM z10 and z196 platforms.**

- When CPUs are added to an LPAR, z/VSE automatically updates its CPU configuration.

§ **You can then use the:**

- SYSDEF TD,STARTSBY=cpuaddr command to set CPUs that are in a “standby” state to an “online” state and start these CPUs
- SYSDEF TD,STOPSBY=cpuaddr command to change the CPU state from “online” to “standby”.

§ **These functions allow you to exploit z10 and z196 technology and update the CPU configuration depending on workload needs**

§ Note: “Standby” CPUs do not consume any CPU share of their LPAR

z/VSE V4.3 – Dynamic Starting/Stopping of CPUs

Customize Image Profiles: H05:H05LP56 : H05LP56 : Processor

Group Name: <Not Assigned>

Logical Processor Assignments

Dedicated processors

Select	Processor Type	Initial	Reserved
<input checked="" type="checkbox"/>	Central processors (CPs)	2	1
<input type="checkbox"/>	zSeries application assist processors (zAAPs)	0	0
<input type="checkbox"/>	System z integrated information processors (zIIPs)	0	0

Not Dedicated Processor Details

Initial processing weight: 10 (1 to 999) Initial capping

Enable workload manager

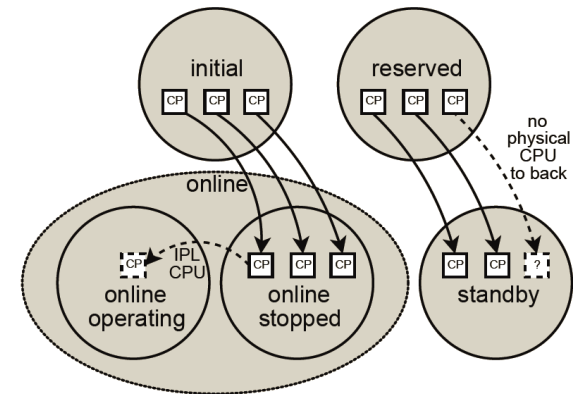
Minimum processing weight: 0

Maximum processing weight: 0

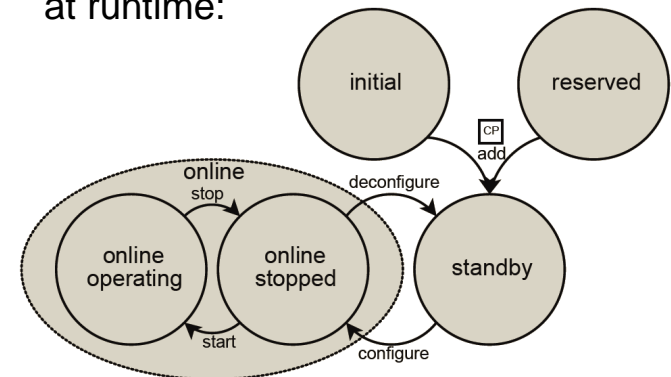
Logical Processor Add: H05LP56 (Active) - H05:H05LP56

CP Type	Number of Initial CPs	Number of Reserved CPs	Capping	Dedicated	Initial Weight	Minimum Weight	Maximum Weight
GP	1	2	<input type="checkbox"/>	<input type="checkbox"/>	10		
zAAP	0	0	<input type="checkbox"/>	<input type="checkbox"/>			
zIIP	0	0	<input type="checkbox"/>	<input type="checkbox"/>			

LPAR CPU state transitions during LPAR activation and IPL:



LPAR CPU state transitions at runtime:



z/VSE V4.3 – Dynamic Starting/Stopping of CPUs

```
query td
```

```
AR 0015 CPU STATUS SPIN_TIME NP_TIME TOTAL_TIME NP/TOT  
AR 0015 00 ACTIVE 0 43 23070 0.001
```

```
[...]
```

```
AR 0030 0W03I 00002 STANDBY CPUS HAVE BEEN ADDED TO THE CONFIGURATION
```

```
query td
```

```
AR 0015 CPU STATUS SPIN_TIME NP_TIME TOTAL_TIME NP/TOT  
AR 0015 00 ACTIVE 0 65 23095 0.002  
AR 0015 01 STANDBY  
AR 0015 02 STANDBY
```

```
[...]
```

```
sysdef td,startsby=01
```

```
AR 0015 1YH7I NUMBER OF CPU(S) - ACTIVE: 1 - QUIESCED: 0 - INACTIVE: 0 -  
STANDBY: 1
```

```
AR 0015 1I40I READY
```

z/VSE V4.3 – Queue-I/O Assist (QIOASSIST)

§ The z/VM function **queue-I/O assist** (QIOASSIST) provides performance improvements for V=V guests using real adapters and real networking devices that use the Queued Direct I/O (QDIO) facility

§ z/VSE exploits the queue-I/O assist function for:

- OSA Express devices (CHPID type OSD)
- HiperSockets devices (CHPID type IQD)

§ To use the queue-I/O assist (QIOASSIST) function in z/VSE, use the z/VM CP command:

- SET QIOASSIST ON

§ After you have enabled the queue-I/O assist function, each z/VSE DEFINE LINK,TYPE=OSAX command will then **automatically exploit** the queue-I/O assist function

§ If you do not wish to use the queue-I/O assist (QIOASSIST) function in z/VSE, disable it using this z/VM CP command:

- SET QIOASSIST OFF

z/VSE V4.3 – Crypto Express 3 and AP queue interrupt support

§ **Support for AP-interrupts is a new function of IBM System z10 and IBM zEnterprise**

§ **A hardware interrupt is issued when a response is ready for de-queuing from a card.**

- Removes the need for the formerly used polling mechanism
- User can switch between polling and interrupts (default: polling)
- Using interrupts **increase throughput** for certain workloads without increasing CPU load

§ **Only available in LPAR mode, not available under z/VM!**

§ **Supported cards are:**

- Crypto Express 2
- Crypto Express 3

§ **The VSE crypto device driver provides new commands:**

- **APEAI**, enable AP interrupts for all APs
- **APDAI**, disable AP interrupts for all APs



z/VSE V5.1 – Enhancements

§ 64 bit virtual addressing

- Support for (64-bit) Memory Objects in 64-bit Address Spaces

§ Exploitation of IBM System Storage Tape Technology

- TS7700 Copy Export Feature
- TS7700 Multi-Cluster Grids

§ Exploitation of IBM System Storage Disk Technology (supported by z/VSE V4.2 or later)

- IBM Storwize V7000
- IBM XIV Storage System

§ Security/Encryption Enhancements

- Improved Crypto Commands
- New AP Crypto Command
- 4096-Bit Key Lengths

§ e-business Connectors Enhancements

§ System Management Enhancements

- Enhancements of SNMP Trap client
- GDPS Support



z/VSE V5.1 – Enhancements

§ VSE/POWER Enhancements

- Support for the TKN Attribute

§ TCP/IP and Networking Enhancements

- Layer 2 Support
- Virtual LAN (VLAN) Support
- Support for Intra-Ensemble Data Network (IEDN) using OSA Express for zBX
- z/VSE z/VM IP Assist (z/VSE VIA)
- IPv6/VSE
- Support for IPv6 in Fast Path to Linux on System z

§ Dump Handling Enhancements

- Support for Memory Objects
- Support for 64-Bit Address Space

§ Available via PTFs (see announcement letter)

- CICS Explorer
- Database connector
- 64-bit I/O



z/VSE V5.1 – 64 bit Virtual Performance Considerations

à **64 bit virtual does not make the system run faster or slower (*)**

(*) Besides the general release overhead

§ But it allows an application to keep more data in memory

– Memory Objects above the 2 GB bar

§ Examples:

– Extended caching for database type of applications

– Large TCP windows support of IPv6/VSE product

à Exploitation requires changes in application or vendor code

§ You may have to increase the VSIZE of your system to utilize 64 bit memory objects

– Check the real storage, to avoid paging



Standalone Dumps of large systems

§ Dumping huge amounts of data may take a considerable long time

- Large VSIZE
- Huge partitions (up to 2GB)
- 64 bit Memory Objects (above or below the 2 GB bar)



§ Be prepared for taking a Stand-Alone Dump in case of a hardware

- Check the SADUMP options
 - Dump only what is really needed
- Have an updated Stand-Alone Disk/Tape available
 - The creation of the standalone dump program on disk can take a long time to complete
 - à Every data record on IJSYSDU is initialized
 - à This job can run in background and should not affect other workload
- Have enough space available

§ In general, it is recommended to use standalone dump on disk

- Only if the size of the system is > 50GB, then use standalone dump on tape
- Advantage of disk:
 - It is faster
 - IJSYSDU file can be easily transferred via FTP

Sizing a system for z/VSE

§ Sizing a system for z/VSE is **different** from sizing a system for z/OS

- Although z/VSE supports multiprocessing, z/VSE does not scale as good as z/OS does
 - Do not use more than 3 active processors per z/VSE LPAR or z/VM Guest



§ In general, **a faster single CPU is better** than multiple smaller CPUs

- One partition can only exploit the power of one CPU
 - The largest partition (e.g. CICS) must fit into one single CPU
- Dependent on nonparallel share (NPS) value

§ Additional CPUs can be useful when multiple LPARs or z/VM Guests are used

- Define **only up to 3 CPUs** per LPAR or z/VM Guest, even if more than 3 CPUs are available on the CEC

§ **Do not use MIPS tables** for capacity planning purposes

- Use **zPCR Tool** instead with the z/VSE workloads **Batch, Online or Mixed**
- Use free of charge Capacity Planning Services from IBM

Sizing a system for z/VSE

**The fastest
uni-processor
is (almost always *)
the best processor**

(*) from a single VSE-image point o view

IBM Processor Capacity Reference for zSeries (zPCR)

- § The zPCR tool was released for customer use on October 25, 2005
 - <http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/PRS1381>
 - ‘As is’, no official support, e-mail to zpcr@us.ibm.com

§ PC-based productivity tool under Windows

§ It is designed to provide capacity planning insight for IBM System z processors running various workload environments



§ Capacity results are based on IBM's LSPR data supporting all IBM System z processors

- Large System Performance Reference:
<http://www.ibm.com/systems/z/advantages/management/lspr/>

§ For VSE use z/VSE workloads **Batch**, **Online** or **Mixed**

z/VSE CPU Monitor Tool

§ Intended to help customers to **measure the CPU utilization** of their VSE system **over a period of time**.

§ When you plan for a processor upgrade it is very important to know the CPU utilization of your VSE system over a day or a week.

– Helps you to estimate the size of the new processor.

§ The VSE CPU Monitor Tool is not intended to replace any existing monitoring product provided by partners.

§ It provides only very **basic monitoring** capabilities on **an overall VSE system level**.

§ No details about CPU usage of certain applications are provided

§ New version available (XML Output) for z/VSE Capacity Planning

§ **Download**


– <http://www.ibm.com/systems/z/os/zvse/downloads/tools.html>

– ‘As is’, no official support, e-mail to zvse@de.ibm.com



z/VSE CPU Monitor Tool

Example CPU utilization chart



z/VSE Capacity Planning Offering

§ **A new z/VSE Capacity Planning Offering is now available**

- Available for Business Partners
- and Customers

§ **Performance data collection is based on a new version of the CPUMON Tool**

§ **Analysis is done using zCP3000**

§ **Contact techline@us.ibm.com**
and ask for z/VSE Capacity
Planning Support

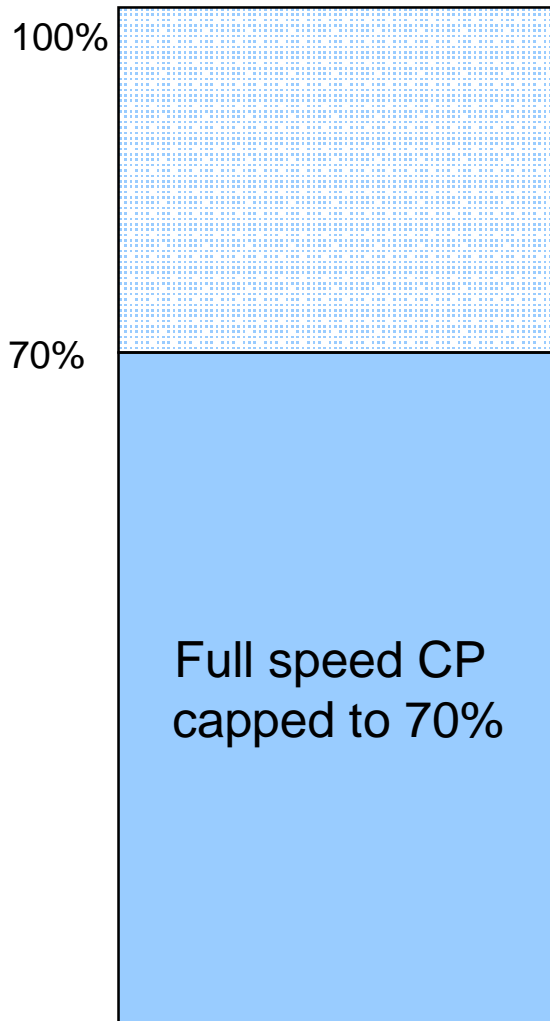



Capping versus Capacity Settings


Attention: Do not use Capping to simulate Capacity Settings !

§With **Capping**, the processor runs on its full speed, until the capping stops the guest from getting dispatched by the LPAR hipervisor or z/VM (timeslicing)

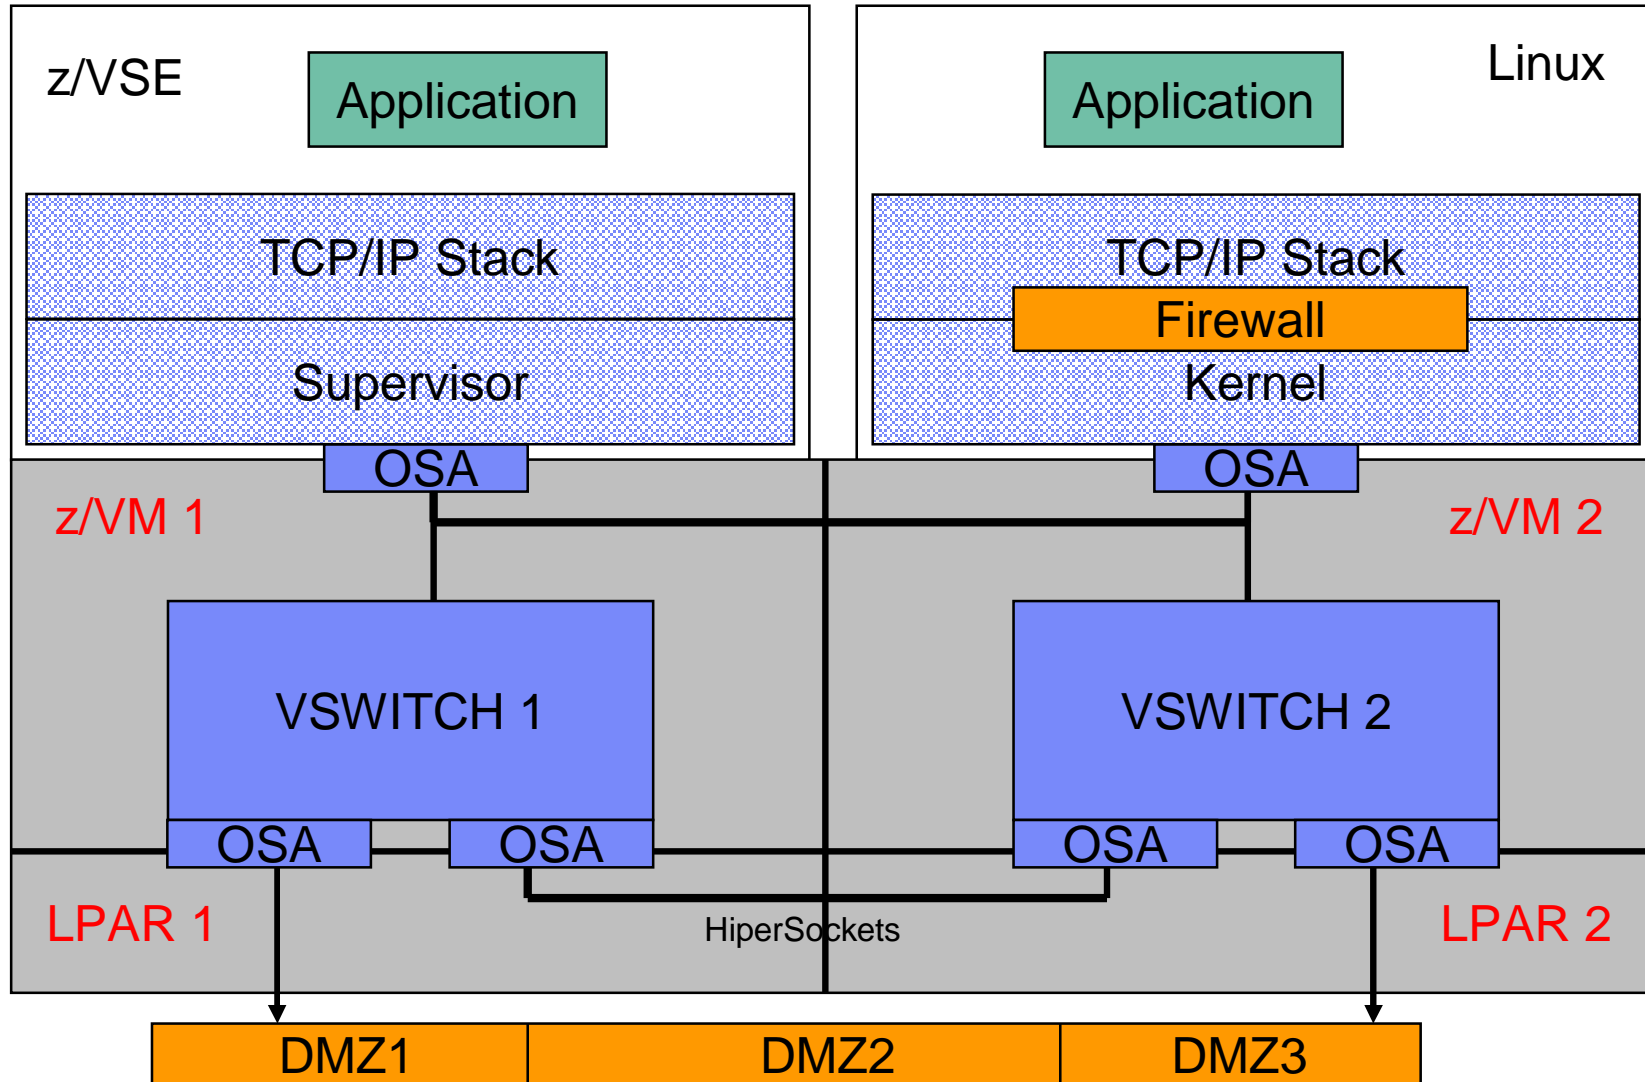
§With a **Capacity Setting**, the processor runs on a slower speed (and all related tasks as well, like Hipersockets memory copy, Hipervisor processing, etc)




Capping is NOT equivalent to Capacity Settings !



TCP/IP Tuning: A simple picture might not be that simple in reality



Shared OSA Adapter versus HiperSockets

To connect a z/VSE system with a Linux on System z you have 2 options:

1. Using a shared OSA Adapter

- § All traffic is passed through the OSA Adapter
- § The OSA Adapter has **its own processor**
 - § Processing occurs asynchronous
 - § Processing in OSA Adapter does not affect host processors

2. Using HiperSockets

- § Direct memory copy from one LPAR/Guest to the other
- § Memory copy is **handled by the host processors**
 - § Processing occur synchronous
 - § Consider mixed speed processors (full speed IFLs and throttled CPs)
 - à Memory copy performed by throttled CP is slower than memory copy performed by full speed IFL



TCP/IP Tuning: Performance tuning for HiperSockets

§ When using HiperSockets to communicate between z/VSE and Linux, you may run into a “**Target Buffer Full**” condition

- This happens when z/VSE sends faster/more than Linux can receive
- Per default Linux has 16 inbound buffers (64K per buffer = 1M per link)
- To **increase the number of buffers on Linux**, use QETH option “**buffer_count=128**”
 - Use YAST to configure, or sysconfig scripts
 - Maximum of 128 buffers require 8MB of storage per link

§ When TCP/IP for VSE encounters this situation (**BUSY**), it waits 500 msec until it retries to send the packet

- Any additional packets to be sent are queued up
- Problem can become dramatic, if more than 16 packets are queued up to be sent after BUSY situation
 - The resend will immediately flood the Linux buffers again, leading to the next BUSY situation, and so on....

§ You can check via **QUERY STATS, LINKID=xxxx [,RESET]** if you have ever run into the **BUSY situation** (RESET resets the counters)

```
C1 0065 0004: IPL615I  Busy mode.....0   B see here
C1 0065 0004: IPL615I  Busy mode, longest.....0
```

§ You can configure a shorter **BUSY** wait time via **DEFINE LINK** command

- **BUSY=nnn** (shortest possible wait time is 100 msec)

Best practice activities for monitoring and tuning

Any performance improvement process includes the following fundamental steps:

1. Define the performance objectives.
2. Establish performance indicators for the major constraints in the system.
3. Develop and execute a performance monitoring plan.
4. Continually analyze monitoring results to determine which resources require tuning.
5. Make one adjustment at a time.

If an application or a set of applications does not run as fast as expected, you must become familiar with the affected systems and components.

- A distributed environment, such as z/VSE connected to a DB2 database server running on Linux on System z, requires a **look at the entire picture** rather than at each site or component separately.

The data and requests flow through several different software layers and services on its way from the z/VSE application to the server on Linux on System z and back.

- Each layer adds a certain amount of overhead and processing time.
- To learn why such an application does not perform as expected, you might need to **look at each layer**, which includes the z/VSE application itself, the middleware running on z/VSE, the TCP/IP stack on both systems, the network, and middleware running on Linux on System z.
- In addition to the layers that are directly involved, you also must look at both operating systems, z/VSE and Linux on System z.

Best practice activities for monitoring and tuning

§ Remember the law of diminishing returns

- The greatest performance benefits usually come from your initial efforts.

§ Do not tune just for the sake of tuning

- Tune to relieve identified constraints.
- If you tune resources that are not the primary cause of performance problems, you can make subsequent tuning work more difficult.

§ Consider the whole system

- You cannot tune one parameter or resource in isolation.
- Before you make an adjustment, consider how the change will affect the system as a whole.
- Performance tuning requires trade-offs among various system resources.
 - For example, you might increase buffer pool sizes to achieve improved I/O performance, but larger buffer pools require more memory, and that might degrade other aspects of performance.

§ Change one parameter at a time

- Do not change more than one factor at a time.
- Even if you are sure that all the changes will be beneficial, you will have no way to assess the contribution of each change.



Best practice activities for monitoring and tuning



§ Measure and configure by levels

- Tune one of the following levels of your system at a time:
 - Hardware
 - Operating system
 - Middleware components
 - Application programs

§ Check for both hardware and software problems

- Some performance problems can be corrected by applying service to your hardware, your software, or both.
- Do not spend excessive time monitoring and tuning your system before applying service to the hardware or software.

§ Understand the problem before you upgrade your hardware

- Even if it seems as though additional storage or processor power might immediately improve performance, take the time to understand where the bottlenecks are.
- You might spend money on additional disk storage, only to find that you do not have the processing power or the channels to exploit it.

§ Put fallback procedures in place before you start tuning

- If tuning efforts result in unexpected performance degradation, reverse the changes that are made before you attempt an alternative approach. Save your original settings so that you can easily undo changes that you do not want to keep.

Best practice activities for monitoring and tuning

1. Collect data:

- § The data collection step involves the usage of various performance monitor tools and optionally tracing tools on all components that are affected
- § Part of this step proves that the monitoring and tracing tools work as expected and produce the kind of output that is needed for further analysis.

2. Analyze the data:

- § The data analysis step uses the data that was collected in the previous step and analyzes it.
- § Depending on the performance problem, concentrate on unusual behavior that catches your attention such as delays, high CPU utilization, long wait times, high water marks, and so on.

3. Tune the systems:

- § Based on the results from analyzing the data, you tune the different parts of the environment.
- § The goal is to change settings that influence the behavior when analyzing the data but in a positive way.

4. Run the test case:

- § Re-run the test case to check whether the tuning has helped.
- § Don't forget to activate the performance monitor and tracing tools
- § Make sure that the preconditions of the test case are the same on all test runs

z/VSE monitoring tools

§ **System Activity Dialogs (SYS fast path 361 and 362)**

- Displays real-time performance information about the System, CPU, partitions and I/O

§ **QUERY TD command**

- Displays information about CPU usage on the console

§ **SIR SMF command**

- Displays I/O related performance information on the console

§ **Job Accounting Exit (SKJOBACC in ICCF library 59)**

- Prints performance related information (CPU, I/O) to SYSLST after each job step

§ **MAP and GETVIS commands**

- Displays memory related information on the console

§ **z/VSE CPUMON Tool**

- Monitors overall system CPU usage and performance counters

§ **CICS Statistics**

- Prints CICS statistics


§ **CICS built-in tools like CEMT INQUIRE**

- Displays information about CICS resources

§ **A z/VSE performance monitor product for batch and CICS**

- Like CA Explore, ASG TMON, etc.

§ **z/VSE SNMP Agent**



Linux and z/VM monitoring tools

§ **sysstat utilities**

- A collection of performance monitoring tools for Linux

§ **iostat utility**

- Monitors disk utilization

§ **top utility**

- Prints system data for each process. It shows an overview of the currently running system processes.

§ **oprofile utility**

- profiles all running code on Linux systems, providing a variety of statistics

§ **IBM Tivoli OMEGAMON XE on z/VM and Linux**


- Displays data that is collected from multiple systems in one, flexible interface called the Tivoli Enterprise Portal

§ **Performance Toolkit for z/VM**

- Provides capabilities to monitor and report performance data for a z/VM system

§ **3rd Party Performance Monitoring Tools**

- e.g. Tools from Velocity Software, Inc



Questions ?



THANK YOU