



| IBM System z

IBM z/VSE V4.1 and V4.2 Performance Considerations



Trademarks

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

CICS*	IBM*	Virtual Image
DB2*	IBM logo*	Facility
DB2 Connect	IMS	VM/ESA*
DB2 Universal	Intelligent	VSE/ESA
Database	Miner	VisualAge*
e-business logo*	Multiprise*	VTAM*
Enterprise Storage	MQSeries*	WebSphere*
Server	OS/390*	xSeries
HiperSockets	S/390*	z/Architecture
	SNAP/SHOT	z/VM
	*	z/VSE
		zSeries

* Registered trademarks of IBM Corporation

The following are trademarks or registered trademarks of other companies.

LINUX is a registered trademark of Linus Torvalds

Tivoli is a trademark of Tivoli Systems Inc.

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

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

Microsoft, Windows and Windows NT are registered trademarks of Microsoft Corporation.

SET and Secure Electronic Transaction are trademarks owned by SET Secure Electronic Transaction LLC.

Intel is a registered trademark of Intel Corporation.

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.

Contents

§ Hardware support

§ z/VSE V4.1 Considerations

- z/Architecture
- 64-bit real addressing
- Paging

§ z/VSE V 4.2 Considerations

- More tasks
- PAV Performance Considerations
- CPU Balancing

§ Sizing a system for z/VSE

§ Miscellaneous Considerations

z/VSE V4.1 and V4.2 Hardware support

§ z/VSE V4.1 and V4.2 runs on the following machines

- IBM System z10 BC or z10 EC
- IBM System z9 BC or z9 EC (formerly z9-109)
- IBM zSeries: z800, z900, z990, z890
- z/VM V5.2 (or later) is a prerequisite for running z/VSE V4 under VM.

§ z/VSE 3.1 and VSE/ESA 2.7 runs on the following machines

- IBM System z10 BC or z10 EC
- IBM System z9 BC or z9 EC (z9-109)
- IBM zSeries: z800, z900, z990, z890
- 9672 Parallel Enterprise Server (G5/G6)
- Multiprice 3000 (7060)
- equivalent emulators (Flex-ES)

Supported VSE Releases

VSE Release	Available	End of Marketing	End of Service
z/VSE 4.2	10/17/2008		
z/VSE 4.1	03/16/2007	10/17/2008	04/30/2010
z/VSE 3.1	03/04/2005	05/31/2008	07/31/2009
VSE/ESA 2.7	03/14/2003	09/30/2005	02/28/2007 (out of service)
VSE/ESA 2.6	12/14/2001	03/14/2003	03/31/2006 (out of service)
VSE/ESA 2.5	09/29/2000	12/14/2001	12/31/2003 (out of service)
VSE/ESA 2.4	06/25/1999	09/29/2000	06/30/2002 (out of service)
VSE/ESA 2.3	07/12/1997	06/30/2000	12/31/2001 (out of service)

Running z/VSE V4 under z/VM

§ z/VM V5.2 (or later) is a prerequisite for running z/VSE V4.2 or V4.2under z/VM

- If you IPL z/VSE V4.1 or V4.2 in a guest system of z/VM version 4 or z/VM 5.1, you may experience severe performance problems
- Because of that the following message is issued during IPL:
 - 0J86I WARNING: VM RELEASE NOT SUPPORTED BY VSE 4.1
– Z/VM 5.2 OR LATER REQUIRED
- If you receive this message, you must urgently upgrade your VM system to z/VM 5.2 or a later release.

§ Note: It is not required to run z/VSE under z/VM, you can also run z/VSE in an LPAR

VSE Server Support

IBM Server	z/VSE 4.2	z/VSE 4.1	z/VSE 3.1	VSE/ESA 2.7 (*)	VSE/ESA 2.6 (*)	VSE/ESA 2.5 (*)	VSE/ESA 2.4/2.3 (*)
IBM System z10 BC/EC	Yes	Yes (PTF required)	Yes (PTF required)	Yes (PTF required)	Yes (PTF required)	Yes (PTF required)	No
IBM System z9 BC/EC (z9-109)	Yes	Yes	Yes	Yes	Yes (PTF required)	Yes (PTF required)	No
zSeries 890, 990	Yes	Yes	Yes	Yes	Yes (PTF required)	Yes (PTF required)	No
zSeries 800, 900	Yes	Yes	Yes	Yes	Yes	Yes	Yes
S/390 Parallel Enterprise Server G5/G6	No	No	Yes	Yes	Yes	Yes	Yes
S/390 Multiprise 3000	No	No	Yes	Yes	Yes	Yes	Yes
S/390 Parallel Enterprise Server G3/G4	No	No	No	No	Yes	Yes	Yes
S/390 Multiprise 2000	No	No	No	No	Yes	Yes	Yes
S/390 Integrated Server	No	No	No	No	Yes	Yes	Yes
S/390 Parallel Enterprise Server G2 / G1 (out of Service)	No	No	No	No	Yes	Yes	Yes
ES/9000 – 9221, 9121, 9021 (out of Service)	No	No	No	No	Yes	Yes	Yes
P/390 and R/390 (out of Service)	No	No	No	No	Yes	Yes	Yes

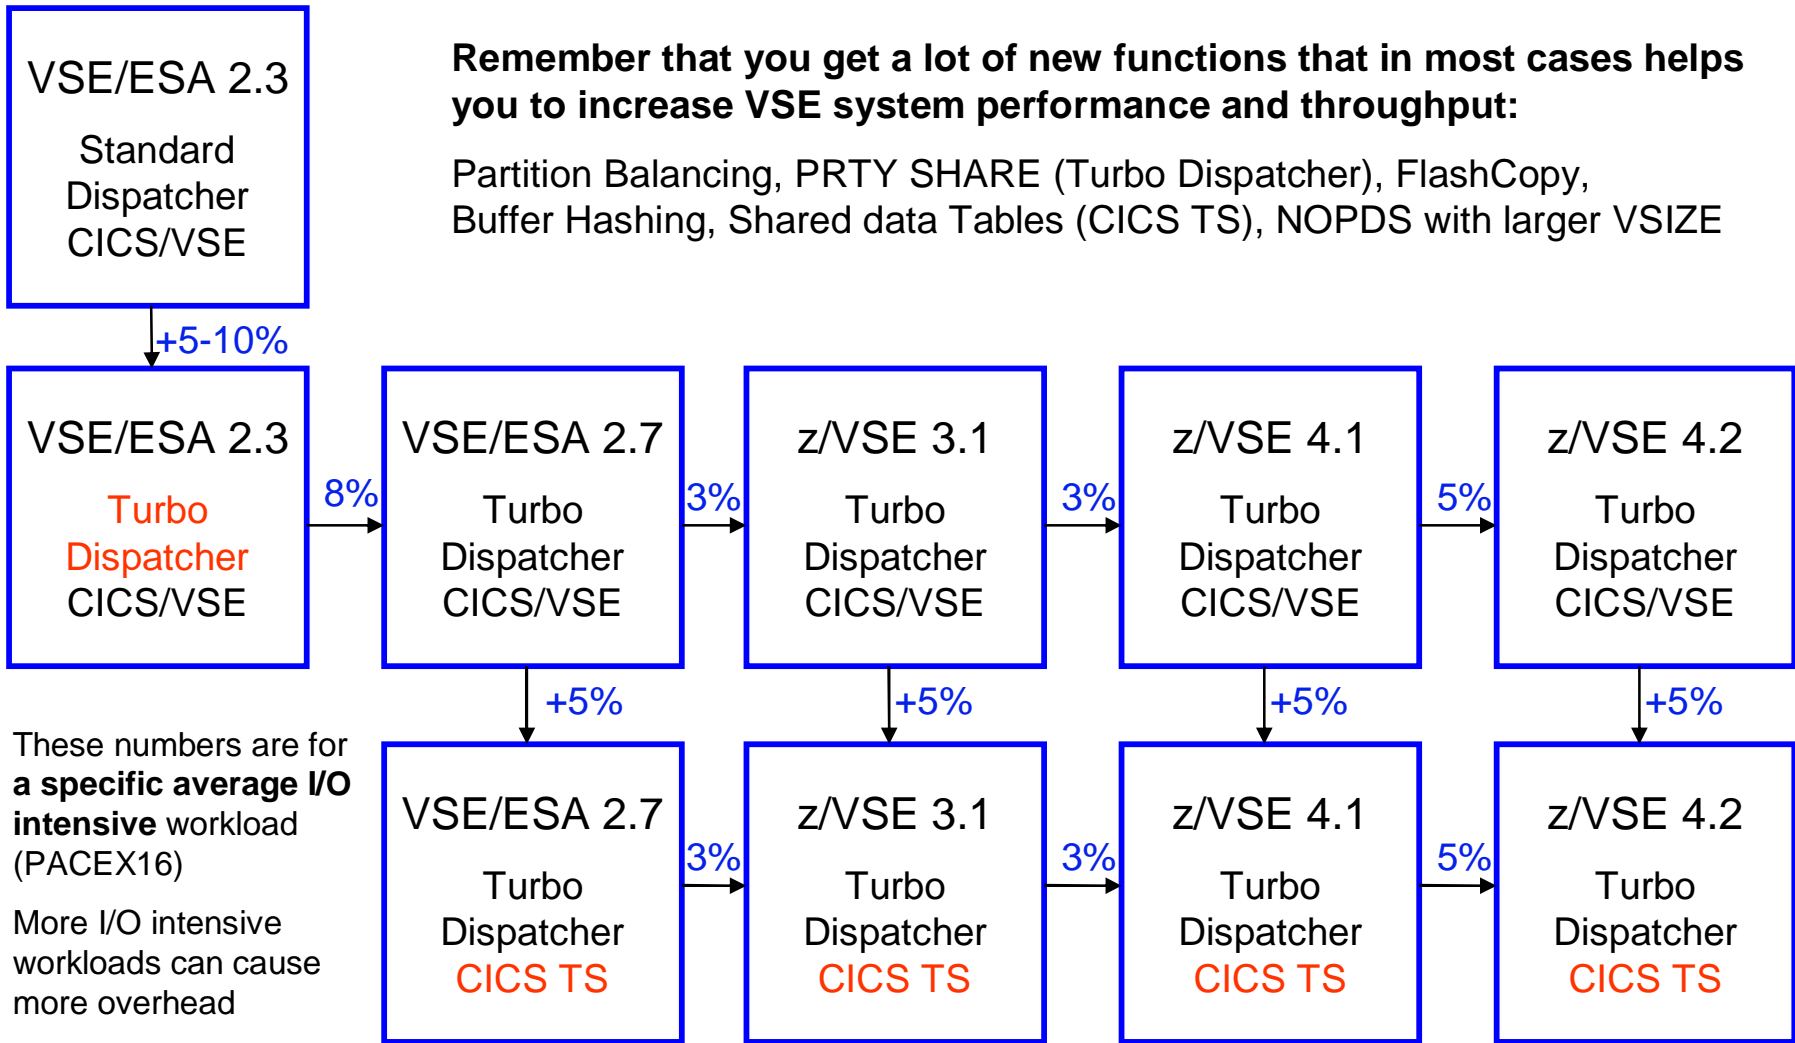
(*) Note: Although VSE/ESA 2.7 or earlier releases technically run on selected servers, these releases are Out-of-Service anyway.

VSE Hardware Support

VSE Release	HiperSockets	OSA Express (QDIO mode)	Hardware Crypto
z/VSE 4.2	Yes	Yes	Yes (PCICA, CEX2C, CEX2A, CPACF)
z/VSE 4.1	Yes	Yes	Yes (PCICA, CEX2C, CEX2A, CPACF)
z/VSE 3.1	Yes	Yes	Yes (PCICA, CEX2C, CEX2A, CPACF)
VSE/ESA 2.7	Yes	Yes	Yes (PCICA, CPACF)
VSE/ESA 2.6	No	Yes	No
VSE/ESA 2.5 or earlier	No	No	No

Crypto Card	z800	z900	z890	z990	z9 BC/EC	z10 BC/EC
PCICA	No	Yes	Yes	Yes	No	No
CEX2C	No	No	Yes	Yes	Yes	Yes
CPACF	No	No	Yes	Yes	Yes	Yes
CEX2A	No	No	No	No	Yes	Yes

Overhead Deltas for VSE Releases



z/VSE V4.1 and V4.2 - z/Architecture mode

§ z/VSE V4.1 and V4.2:

- Supports z/Architecture-capable (64-bit) processors.
- Executes in z/Architecture mode only.
- Supports 64-bit real addressing for selected system functions.
- Supports processor storage up to 8 GB (z/VSE 4.1) / 32GB (z/VSE 4.2).
- The storage beyond 2 GB is managed exclusively by the z/VSE operating system.

§ z/VSE V4.1 and V4.2 does not support:


- 64-bit virtual addressing. The size of a virtual address or data space remains restricted to 2 GB.
- For user applications, 64-bit addressing and operations that use 64-bit registers.

§ 64-bit real addressing is transparent to your user applications providing you use IBM-supplied standard interfaces.

§ Customers with especially large z/VSE environments might benefit from lower paging rates.

§ Many z/VSE environments might be able to run without a page data set (using the NOPDS option).

What does 64 Bit 'real addressing' mean ?



What does 64 Bit 'real addressing' mean ?

§ VSE/ESA V2.x and z/VSE 3.1

- z/VSE 3.1 or below can address only 2 GB of processor storage
 - 31 bit (real) addressing
- Page data set required if VSIZE (+VIO) > ~ 2GB
- No page data set required as long 2 GB processor storage is sufficient

§ z/VSE V4.1 and V4.2

- z/VSE V4.1 can address up to 8 GB of processor storage, z/VSE V4.2 up to 32 GB
 - 64 bit (real) addressing
- No page data set required
 - If processor storage >= VSIZE (+VIO)
- Virtual address spaces and data spaces are still limited to 2 GB
 - 31 bit (virtual) addressing
 - No changes required to applications

z/VSE V4 – Example with 3 GB of real storage

```

map real
AR 0015 AREA          R-SIZE   R-ADDR   PFIX(BELOW)      PFIX(ABOVE)
AR 0015                ACTUAL   LIMIT      ACTUAL   LIMIT
AR 0015 SUP           52K      0
AR 0015 SYS-24        592K      14136K
AR 0015 BG V          0K        0K         0K       0K
AR 0015 F1 V          148K      400K       0K       1400K
AR 0015 F2 V          32K       144K       0K       0K
AR 0015 F3 V          88K       424K       0K       0K
AR 0015 F4 V          0K        0K         0K       0K
AR 0015 F5 V          0K        0K         0K       0K
AR 0015 F6 V          0K        0K         0K       0K
AR 0015 F7 V          200K      400K       1044K    2100K
AR 0015 F8 V          0K        0K         0K       0K
AR 0015 F9 V          0K        0K         0K       0K
AR 0015 FA V          0K        0K         0K       0K
AR 0015 FB V          0K        0K         0K       0K
AR 0015 SYS-31        7404K     2052264K
AR 0015 DYN-PA        0K        0K         0K       0K
AR 0015 AVAIL         64K
AR 0015 SYSTEM        25068K
AR 0015 TOTAL         3145728K
AR 0015
AR 0015 AVAILABLE FOR SETPFIX: 13544K      2044860K
AR 0015
AR 0015 1I40I READY

```

Exploiting 64 Bit real storage

- § **Even on VSE/ESA or z/VSE 3.1 the VSIZE could ‘theoretically’ be up to 90G**
- § **‘Practically’ you are limited by**
 - Page dataset size and number of extents
 - Page I/O rate
 - Too heavy page I/O rates makes VSE almost unusable
- § **With z/VSE V4 you can run with VSIZE+VIO up to approximately 8GB (z/VSE V4.1) / 32 GB (z/VSE 4.2) without a page dataset (NOPDS)**
 - If enough processor storage is available
 - No time consuming page I/Os
- § **You can have more large partitions in your system**

Paging considerations

§ **'Paging' is another word for Page Manager activities**

- Assigning real pages to virtual pages
- Writing pages to the page data set (page-out I/O)
- Reading pages from the page data set (page-in I/O)

§ **Even with 8GB real storage and NOPDS paging happens**

- Assigning real pages to virtual pages
- Moving pages from <2G to >2GB (above the bar) and vice versa

§ **'Paging' as such is not bad**

- But page I/Os are 'bad' (dependent on the page I/O rate)

Paging considerations - PMRMON

§ **New SIR command: SIR PMRMON**

- Displays reports from the 'Page Manager Monitor'
- Shows number of page faults, page I/Os, page exchanges (31 <-> 64), ...

§ **Usage:**

- SIR PMRMON=ON (resets counters)
- // run your workload
- SIR PMRMON
- SIR MPRMON=OFF

§ **Output example, see next foil**

§ **The output displays mainly internal counters that are for evaluation by IBM support persons**

Paging considerations - PMRMON

```

AR 0015                PAGE MANAGER MONITORING REPORT
AR 0015                (BASED ON A 0000:01:04.801 INTERVAL)
AR 0015 IPFQ 31-BIT    =      517126    IPFQ 64-BIT    =      115185
AR 0015  PSQ 31-BIT    =       1500     PSQ 64-BIT    =       527
AR 0015 PF EXCH TOTAL  =       4673     PF EXCH 31->64 =         0
AR 0015 PF EXCH 64->31 =       4673     PGFLT TOTAL    =      25825
AR 0015 PGFLT PMGR     =         0      PGFLT USER    =      25825
AR 0015 PGFLT IMM PO 31 =         0      PGFLT IMM PO 64 =         0
AR 0015 SELCT ON PSQ 31 =         0      SELCT ON PSQ 64 =         0
AR 0015 SELC R=1 MAX 31 =         0      SELC R=1 MAX 64 =         0
AR 0015 RECLAIMS       =         0      NPSQ LOW      =         0
AR 0015 PGOUT I/O TOTAL =         0      PGIN I/O TOTAL =         0
AR 0015 PGOUT I/O UNC. =         0      PGOUT I/O PRE. =         0
AR 0015 LRA PGM CHECK  =       3904     TFIX 64-BIT FR =       768
AR 0015 1I40I  READY

```

Description of values see next foil

Paging considerations - PMRMON

IPFQ 31-BIT	Number of unused page frames below 2G
PSQ 31-BIT	Number of page frames below 2G in page selection queue
IPFQ 64-BIT	Number of unused page frames above 2G
PSQ 64-BIT	Number of page frames above 2G in page selection queue
PF EXCH TOTAL	Total number of times a page frame exchange is requested
PF EXCH 31->64	Number of times a page frame is exchanged from below 2G to above 2G (e.g. instead of page-out I/O)
PF EXCH 64->31	Number of times a page frame is exchanged from above 2G to below 2G (e.g. TFIX)
PGFLT TOTAL	Total number of page faults
PGFLT PMGR	Number of page faults handled by page manager task
PGFLT USER	Number of page faults handled immediately without activating page manager task
PGFLT IMM PO 31	Number of page faults forcing an immediate page-out of a page frame below 2G
PGFLT IMM PO 64	Number of page faults forcing an immediate page-out of a page frame above 2G
SELCT ON PSQ 31	Number of page selection events on page frames below 2G
SELCT ON PSQ 64	Number of page selection events on page frames above 2G
SELC R=1 MAX 31	Maximum number of cycles on a page selection event for page frames below 2G
SELC R=1 MAX 64	Maximum number of cycles on a page selection event for page frames above 2G
RECLAIMS	Number of times page frames are reclaimed from the page-out queue
NPSQ LOW	Number of times the available page frames are below a critical minimum
PGOUT I/O TOTAL	Total number of page-out I/O
PGIN I/O TOTAL	Total number of page-in I/O
PGOUT I/O UNC.	Total number of unconditional page-out I/O
PGOUT I/O PRE.	Total number of pre-page-out I/O
LRA PGM CHECK	Number of special operation program exceptions on LRA
TFIX 64-BIT FR	Number of TFIX requests with page frames above 2G involved

z/VSE V4.2 Performance Considerations

§ More tasks

- z/VSE V4.2 now supports up to 512 tasks

§ PAV Performance Considerations

- PAV = Parallel Access Volumes

§ CPU Balancing

- Automatically start and stops CPUs according to the current workload

z/VSE V4.2 – More tasks

§ z/VSE 4.2 supports up to 512 concurrent tasks

- the previous limit was 255 concurrent tasks).
- Additional task IDs will be in the range X'0100' to X'0200'.
- A QUERY SYSTEM command has been provided to display information on the current task usage

§ Please note:


- The maximum **number of tasks per partition remains 32**.
- The default maximum remains 255. If you wish to increase the maximum number of tasks, you must use the SYSDEF command to increase it to any number up to 512.

§ Benefits:

- The workload resulting from CICS and batch programs can now be more efficiently distributed
- Previously, some customers might have been reluctant to perform migration from CICS/VSE to the CICS Transaction Server for VSE/ESA because of its more intensive usage of VSE tasks (compared to CICS/VSE). This inhibitor has now been removed.
- Using more tasks have **no measurable performance impact** (neither negative nor positive)

z/VSE V4.2 – More tasks

- § The **SYSDEF SYSTEM** command is used to activate the new tasks support system-wide.



§ **NTASKS=nnn|MAX**

- nnn specifies the total number of tasks (old and new) that can be allocated to the system and application programs. This is possible only during BG ASI processing. The value of nnn can be between 255 and 512.


§ **TASKS=ANY|OLD**

- This is the system-wide default mode for subtask attaching, it can be run at any time after IPL:
 - If TASKS=ANY is specified, new or old tasks can be attached as subtasks. However, primarily new tasks will be attached.
 - OLD specifies that only old tasks will be attached as subtasks.

§ Use **QUERY SYSTEM** to display what has been specified with **SYSDEF SYSTEM**

z/VSE V4.2 – More tasks

§ The **TASKS** specification in the **SYSDEF SYSTEM** command **can be overwritten** with the **TASKS** parameter in the **JCL EXEC** statement:



- When the SYSDEF SYSTEM command is used with TASKS=OLD, application programs which want to make use of the new tasks, must be started with parameter TASKS=ANY in the EXEC statement.
- When the SYSDEF SYSTEM command is used with TASKS=ANY, but an application program needs to be restricted to old tasks, the parameter TASKS=OLD has to be specified in the EXEC statement.

z/VSE V4.2 – Parallel Access Volumes (PAV)


§ **Parallel Access Volume (PAV) is an optional licensed feature on the IBM System Storage DS6000/DS8000 series**

§ **Parallel access volumes:**


- Are managed by creating **multiple addresses (aliases)** for a single **logical device**.
- Allow your z/VSE system to **access volumes in parallel**.
- Enable **more than one I/O operation** to be processed for a single logical device.
- Depending upon workload, might represent a significant performance improvement by the storage unit over traditional I/O processing.
- Can significantly **reduce device queue delays** to busy devices (all I/O operations are normally processed from a single queue).
- Allow your system to access a single volume from a single z/VSE host with **multiple concurrent requests**.



Base



Alias X



Alias Y

z/VSE V4.2 – Parallel Access Volumes (PAV)

- § **To gain a performance benefit from PAV, it is required that:**
 - the workload can **utilize multiple paths** (subchannels) to the same volume
 - A single job that **starts an I/O and waits** until its completion before it starts another I/O **does not benefit** from PAV by itself
 - However, **running multiple such jobs concurrently** in multiple partitions can utilize parallel access to the same volume.
- § **Besides the simple ‘Start I/O and wait’ behavior, more sophisticated programs (like CICS) may be capable to **start multiple I/Os and wait for its completion asynchronously****
 - This behavior also is eligible for benefiting from PAV, if the I/Os are for the same volume
- § **Jobs that use multiple subtasks that perform I/O operations may benefit from PAV**
- § **Running **multiple jobs concurrently** in separate partitions and access the same volume are also good candidates for PAV**

z/VSE V4.2 – Parallel Access Volumes (PAV)

§ Candidates for PAV volumes

- In general, all volumes that contain files, catalogs of data in general that are **used from multiple jobs** can benefit from PAV
 - Without PAV, I/Os might get queued and therefore delayed in case of parallel access to the same volume
 - PAV allows parallel access to these files from different jobs without disturbing each other.
- VSAM Catalogs
- Shared VSAM Clusters
- VSE Libraries

§ Besides volumes that contain files that are shared between multiple jobs, also volumes that contain **heavily used system files** can benefit from PAV

- Although there might not be any parallel access to those system files as such, there may be **parallel access to other files residing on the same volume**
 - Without PAV, I/Os might get queued and therefore delayed in case of parallel access to the same volume
 - PAV allows parallel access to these volumes without being disturbed by access to the system files.
- POWER Spool files
- Workfiles
- Log files

z/VSE V4.2 – Parallel Access Volumes (PAV)

§ Besides the general PAV setup, **no changes** in application programs or JCL are required

- PAV is enabled in z/VSE via the SYSDEF SYSTEM command:

```
SYSD E F  S Y S T E M  , P A V =  S T A R T  S T O P
```

§ z/VM support for PAV (since z/VM 5.2)

- Dedicated PAV devices
- PAV Mindisks.
- For more information see here:
<http://www.vm.ibm.com/storman/pav/pav2.html>

z/VSE V4.2 – CPU Balancing

§ Using **CPU Balancing** can reduce the multiprocessing overhead

- When CPU Balancing is active, the Turbo Dispatcher only selects CPUs that are required for the current workload
- CPU Balancing might also reduce the overhead or performance-degradation for I/O-intensive workloads that are running in one partition only (the remaining partitions are idle) when multiple CPUs are active


§ With **CPU Balancing** active, Turbo Dispatcher inspects the CPU utilization periodically

- When the the utilization falls below a certain threshold, CPUs are dynamically stopped
- When the the utilization increases over the threshold, CPUs are dynamically started

§ This results in only as many CPUs being active that are utilized by the current workload

z/VSE V4.2 – CPU Balancing

§ CPU Balancing is activated using the SYSDEF TD command



§ **INT=nn** - nn defines the interval in seconds, after which the CPUs utilization is to be inspected.

- When **0** is specified, CPU balancing will be deactivated
- The default value is **0**.
- The value must be in the range of: **nn=0..99**

§ **THR=nn** - nn defines the threshold value in percent

- An additional CPU is activated, when the CPU utilization is larger than nn
- The default value is **50**
- The value must be in the range of: **nn=10..99**


§ You can use the QUERY TD command to obtain information about the current status of CPU balancing

z/VM 5.4 Considerations

- § **z/VM V5.4 removes one of the consequences of stopped virtual processors for virtual processor prioritization**
- § **Prior to z/VM V5.4, a guest's CPU share was distributed equally among its virtual processors**
 - by dividing its share value by the number of non-dedicated virtual processors in the configuration
 - **regardless of whether the virtual processors were in a stopped or started state**
- § **z/VM V5.4 performs share redistribution whenever a virtual processor is started or stopped**
 - and **no longer includes stopped virtual processors** in the calculation of how much share to distribute to each virtual processor
- § **The share redistribution also continues to occur in z/VM V5.4, as in previous releases**
 - at the time a SET SHARE command is issued
 - or when the number of defined virtual processors is altered.


z/VM 5.4 Considerations

<= z/VM 5.3



A guest's CPU share is distributed equally among its virtual processors by dividing its share value by the number of processors, regardless of whether the virtual processors were in a stopped or started state.

z/VM 5.4



z/VM V5.4 performs share redistribution whenever a virtual processor is started or stopped and **no longer includes stopped virtual processors** in the calculation of how much share to distribute to each virtual processor.

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

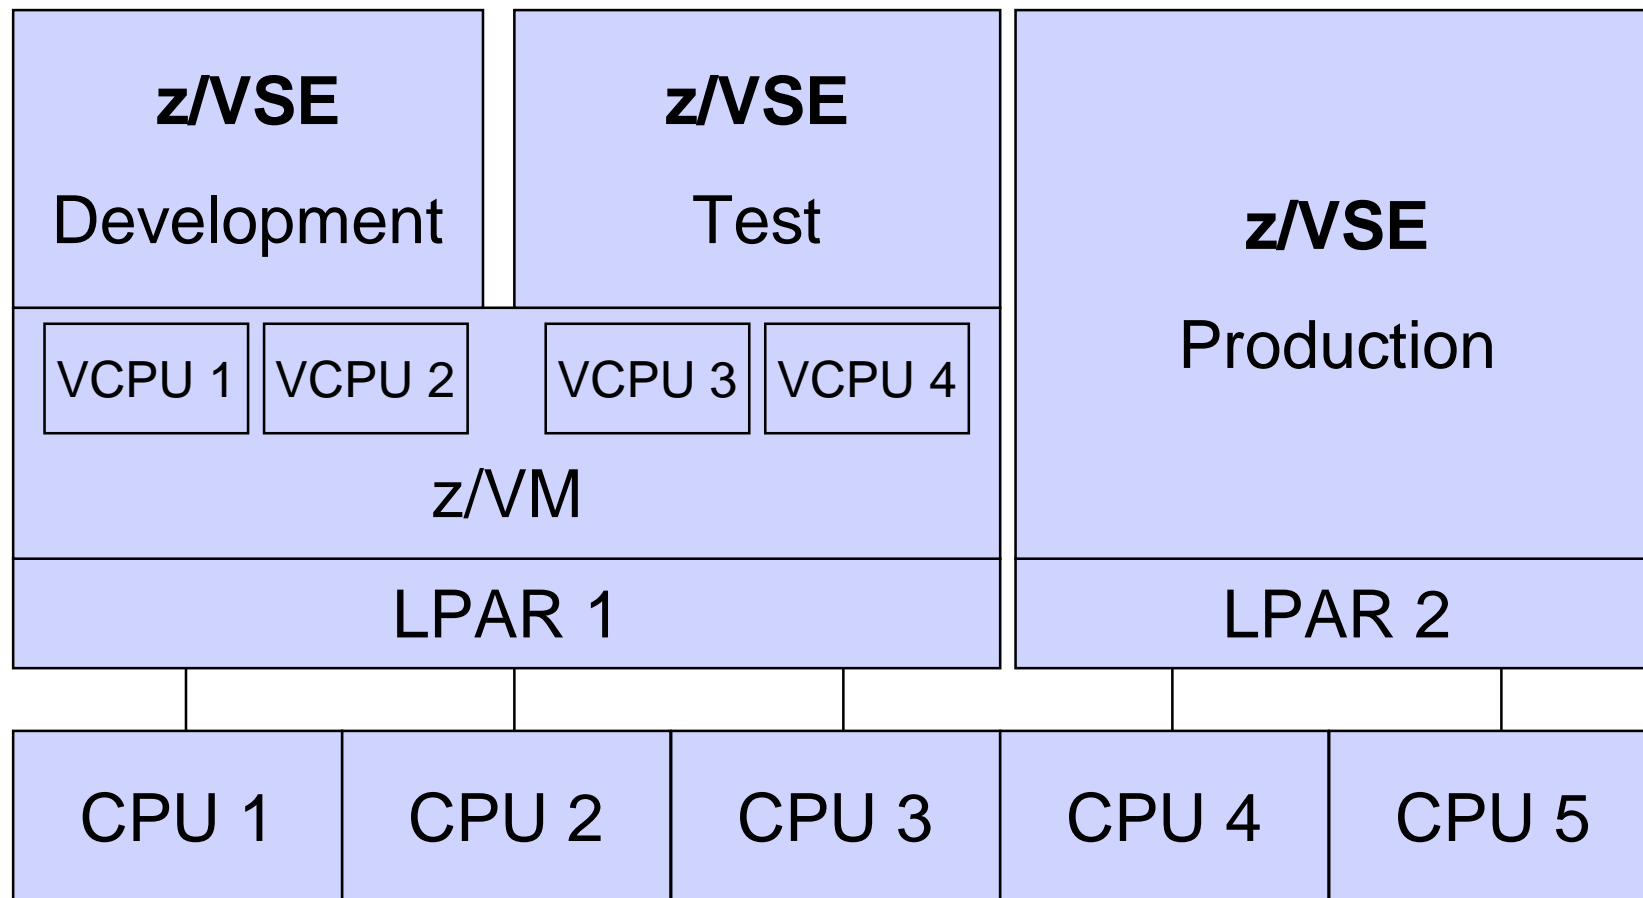
§ Additional CPUs can be useful when multiple LPARs or z/VM Guests are used

- 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 (see page 24) instead with the CB-L workload

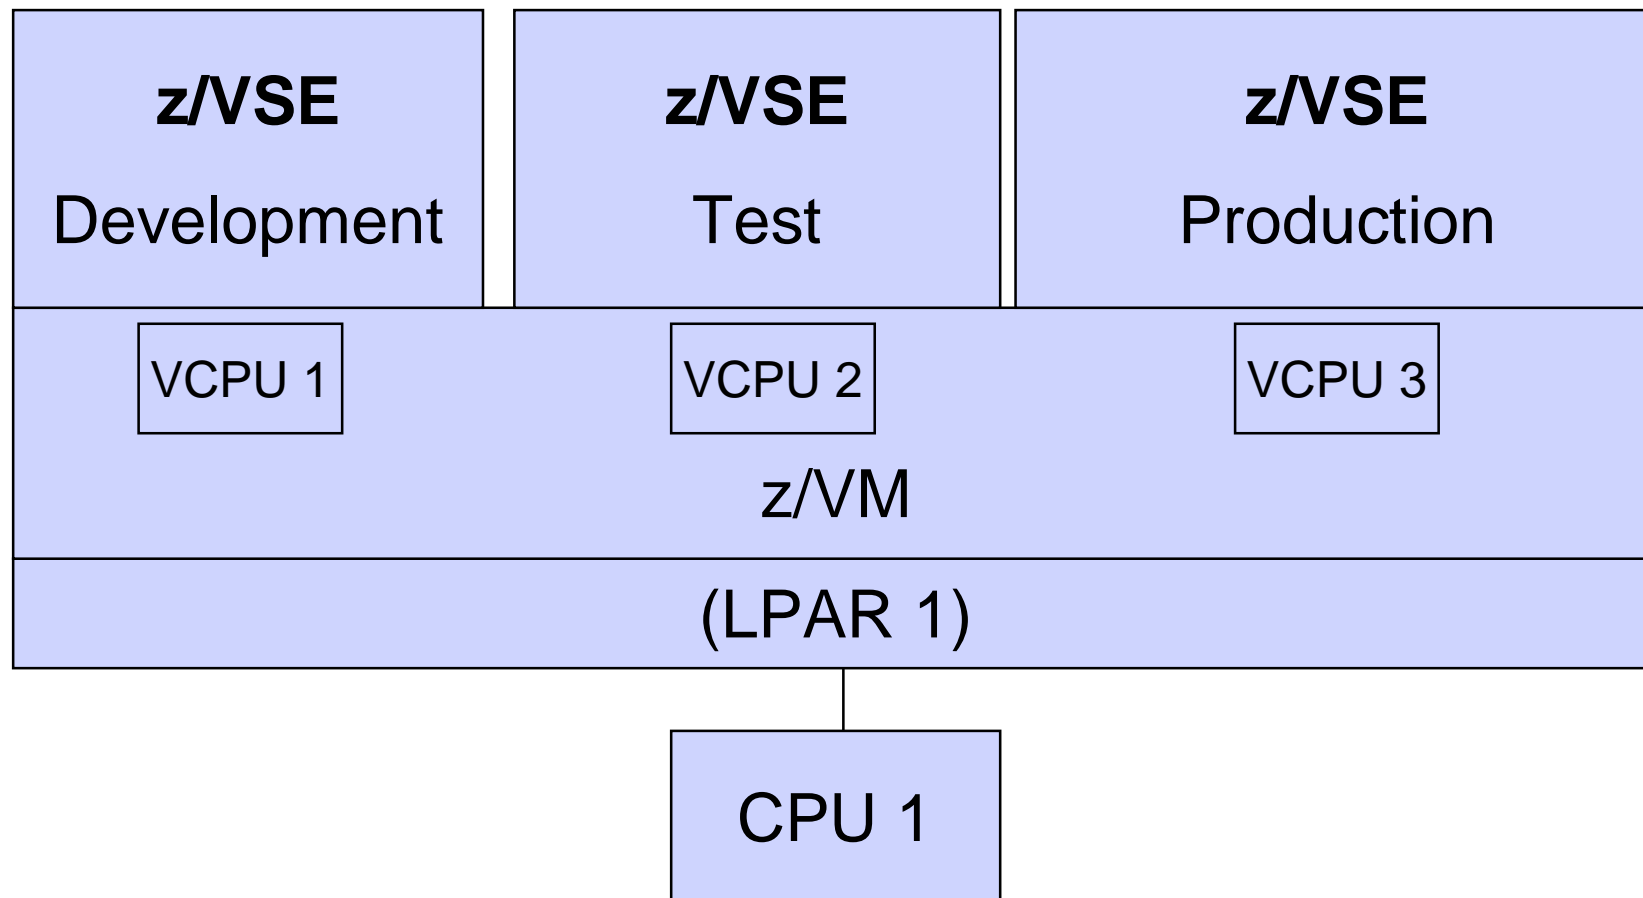
Sizing a system for z/VSE – Example 1



§ Although the CEC has 5 CPUs, each VSE only sees 2 CPUs

§ Test & Development systems have no influence on production system

Sizing a system for z/VSE – Example 2



§ z/VM does a very good job in sharing a single CPU across multiple guests systems

§ Use SET SHARE (ABSOLUTE/RELATIVE) command in z/VM to prioritize production systems

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 z9/z10 and eServer zSeries processors running various workload environments**

§ **Capacity results are based on IBM's LSPR data supporting all IBM System z9/z10 and eServer zSeries processors**

- Large System Performance Reference:
<http://www.ibm.com/servers/eserver/zseries/lspr/>

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

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**
- § **Download**
 - <http://www.ibm.com/servers/eserver/zseries/zvse/downloads/tools.html>
 - 'As is', no official support, e-mail to zvse@de.ibm.com

VSE CPU Monitor Tool


- § CPUMON **periodically** issues a TDSERV FUNC=TDINFO macro to get performance relevant data.
- § The data provided by the macro is the same as command **QUERY TD** shows.
- § The data from each measurement interval is printed to SYSLST in a comma separated format.
- § Later on this data can be imported into a spreadsheet (EXCEL)
- § CPUMON runs in a VSE partition (dynamic or static).
- § CPUMON is started using:

```
// EXEC DTRIATTN,PARM='SYSDEF TD,RESETCNT`  
/*  
// EXEC CPUMON,PARM='nn`   nn = interval in seconds  
/*
```

- § The tool can be stopped by entering the following command:

```
MSG xx,DATA=EXIT           xx = partition id
```

VSE CPU Monitor Tool




IBM System z9 BC (1-4 CPUs)




Note: Do not use MIPS to do any kind of capacity planning, use the zPCR tool instead !

IBM System z9 EC subcapacity models (1-8 CPUs)




Note: Do not use MIPS to do any kind of capacity planning, use the zPCR tool instead !

IBM System z9 EC (9-54 CPUs)




Note: Do not use MIPS to do any kind of capacity planning, use the zPCR tool instead !

IBM System z10 BC (1-5 CPUs)




Note: Do not use MIPS to do any kind of capacity planning, use the zPCR tool instead !

IBM System z10 EC (Models 400, 500, 600)



Note: Do not use MIPS to do any kind of capacity planning, use the zPCR tool instead !

IBM System z10 EC (Models 700)



Note: Do not use MIPS to do any kind of capacity planning, use the zPCR tool instead !

Midrange Workload License Charge (MWLC)

- § **MWLC is a new monthly license charge price metric on the IBM System z9 servers**
- § **It applies to z/VSE V4 and 12 key VSE-related middleware programs**
 - such as CICS TS for VSE, ACF/VTAM for VSE, and DB2 Server for VSE.
- § **MWLC is only available on z9 EC and z9 BC servers with z/VSE V4.**
- § **It is NOT a performance topic**
 - Just for pricing
- § **Capacity Measurement Tool**
 - Measures used MSUs (Millions of Service Units) per image (z/VM guest or LPAR)
 - Measurement interval = 30 minutes
 - Calculates 4 hour rolling average
 - Not to be used for performance tuning !
- § **For more details see**
 - **IBM System z Software Pricing:**
<http://www.ibm.com/servers/eserver/zseries/swprice/>
 - **IBM's MSU ratings for System z:**
<http://www.ibm.com/servers/eserver/zseries/library/swpriceinfo/hardware.html>

z890, z990, z9 and z10 Considerations

- § **The z890, z990, z9 and z10 are LPAR-only machines**
 - No basic mode any more
 - Even if you run just one VSE system, it now runs in an LPAR
 - Running z/VSE systems under z/VM means
 - running z/VSE in z/VM in an LPAR
 - No I/O Assist in LPARs
 - Only available if z/VM runs in basic mode, but no basic mode available on z890, z990, z9 and z10

z/VM V5 considerations

§ z/VM V5 no longer supports V=R and V=F guests

§ z/VM V5 no longer support I/O Assist

- If you currently run with preferred guests, you will need to estimate and plan for a likely increase in processor requirements as those preferred guests become V=V guests as part of the migration.
- Refer to Preferred Guest Migration Considerations at <http://www.vm.ibm.com/perf/tips/z890.html> for assistance and background information

§ How to size the impact (on your current system)

- **Loss of I/O Assist:** Run your workload with CP SET IOASSIST OFF and measure the increase
- **Loss of V=R/F:** Run your workload with V=V and use the CP Monitor to watch for increased CPU consumption

§ How to tune

- **Dedicated processors:** CP SET SHARE ABSOLUTE
- **Dedicated memory:** CP SET RESERVED
- **I/O Assist:** Use minidisks, turn minidisk caching on (MDC)

§ Note: z/VM V5.2 (or later) is a prerequisite for running z/VSE V4 under z/VM

Performance Tips

- § **A partition can only exploit 1 CPU at a time**
 - 2 CPUs do not have any benefit for a single CICS partition
 - Use as many partitions as required for selected n-way
- § **Use/define only as many CPUs as really needed**
 - additional CPUs create more overhead, but no benefit
- § **Partitions setup**
 - Set up more batch and/or (independent) CICS partitions
 - Split CICS production partitions into multiple partitions (AOR, TOR, FOR)
- § **Try to exploit Turbo Dispatcher functions**
 - Priority settings
 - Partition balancing
 - Partition balancing groups

Performance Tips (2)

- § **1 CPU** must be able to handle **all non-parallel workload**
- § **Non-parallel code limits the n-Way exploitation**
 - QUERY TD: $NP/TOT = NPS$ (non parallel share)
 - Measure NPS before migration
 - **max CPUs = 0.8 - 0.9 / NPS**

NPS	#CPUs	NPS	#CPUs
0.20	4.0-4.5 (4)	0.45	1.8-2.0 (2)
0.25	3.2-3.6 (3)	0.50	1.6-1.8 (2)
0.30	2.7-3.0 (3)	0.55	1.5-1.6 (2)
0.35	2.3-2.6 (2)	0.60	1.3-1.5 (1)
0.40	2.0-2.2 (2)	0.65	1.2-1.4 (1)

Performance Tips (3)

- § **Non-parallel code limits the maximum MP exploitation**
- § **System code (Key 0) increases non-parallel share**
 - Vendor code can have significant impact
- § **Overhead increases when NP code limits throughput**
- § **Data In Memory (DIM) reduces non-parallel code**
 - less system calls (I/Os)
 - may increase throughput
 - CICS Shared Data Tables
 - Large/many VSAM Buffers (with buffer hashing)
 - Virtual Disks
- § **Change VSE/POWER startup to WORKUNIT=PA**
- § **Switch tracing/DEBUG off for production**

zSeries Remarks – Split cache

- § **Prior to zSeries there is one cache for data and instructions**
- § **zSeries has split data and instruction cache**
- § **Performance implications:**
 - If **program variables** and **code that updates** these program variables are **in the same cache line** (256 byte)
 - Update of program variable invalidates instruction cache
 - Performance decrease if update is done in a loop
 - See APAR PQ66981 for FORTRAN compiler

zSeries Remarks – Split cache - example

Killer example:

```
*   prepare length
BCTR  R2,0   ADJUST FOR SS-INSTR.
STC   R2,*+5
MVC   RECEIVER(*-*),SENDER
```

STC instruction modifies the next instruction to set the length.

Better code:

```
*   prepare length
BCTR  R2,0   ADJUST FOR SS-INSTR.
EX    R2,MVC01
...
MVC01 MVC    RECEIVER(*-*),SENDER
```

Use EXECUTE instruction instead.

zSeries Performance: Processor Design Considerations:

<http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/FLASH10208>

zSeries Remarks – Split cache - example

Not causing a problem:

```

LA      R1,PHASNAME      POINT AT PHASE NAME
CDDELETE (1)
+*     SUPERVISOR - CDDELETE - 5686-032-06
+      CNOP  0,4
+      BAL   15,*+8
+      DC    A(B'00010010')
+      L     15,0(,15)
+      SVC   65           ISSUE SVC FOR CDDELETE
+      DS    0H

```

CDDELETE uses an inline flag byte,
but does not modify it

Can cause a problem:

```

WTO TEXT=DATA
+      CNOP  0,
+      BAL   1,IHB0003A  BRANCH AROUND MESSAGE
+      DC    AL2(8)      TEXT LENGTH
+      DC    B'0000000000010000'  MCSFLAGS
+      DC    AL4(0)      MESSAGE TEXT ADDR
+      ...
+IHB0003A DS    0H
+      LR    14,1        FIRST BYTE OF PARM LIST
+      SR    15,15       CLEAR REGISTER 15
+      AH    15,0(1,0)  ADD LENGTH OF TEXT + 4
+      AR    14,15       FIRST BYTE AFTER TEXT
+      LA    15,DATA     LOAD TEXT VALUE
+      ST    15,4(0,1)  STORE ADDR INTO PLIST
+*     SUPERVISOR - SIMSVC - 5686-032
+      ...
+      SVC   35           ISSUE SVC 35
@GE00016 DS    0H

```

WTO uses an inline parameter list,
but modifies the parameter list

Note: WTO can be coded with an external
parameter list: WTO ...,MF=(E,addr)

Possible performance issues with PPRC

§ Issue may occur if

- PPRC is used
- VSE runs in native or in LPAR
- Not all devices that are defined in IOCP are also defined in VSE ADD statements

§ In case there is an PPRC state change, interrupts are sent to all LPARs where the related device are defined in IOCP.

- If the device is defined in VSE ADD, no problem occurs: VSE will process the interrupt correctly.
- If the device is NOT defined in VSE ADD, the interrupt is ignored by VSE and the interrupt is resent very quickly to that LPAR
 - Results in very high channel activity (up to 100%)

§ Solution:

- Define ALL devices in VSE ADD that are defined in IOCP

VSE/POWER POFFLOAD Performance Issues

- § **Caused by incompatibility between VSE/POWER tape format and new tape drives**
- § **3490F empties cache for FSF used by POFFLOAD LOAD**
 - Install **DY46164/DY46245** for VSE/ESA 2.7/2.6
- § **3590 synchronizes cache with tape for each WTM**
 - Install microcode **FC0520** on A60 controller + VSE/AF **APAR DY45817** + AR command **TAPE WTM=NOSYNC**
 - Unfortunately controller A50 is too small to install FC0520

New POWER POFFLOAD Append function

§ New APPEND option (new with z/VSE 4.1):

- POFFLOAD {BACKUP|SAVE|PICKUP},
 queue,tapeaddr,density,class,APPEND...
- WAVV requirement WAVV200433
- Requires 3592 tape unit

§ Great performance benefit when using 3592 tape unit:

- Spool time dramatically reduced when using 3592 tape unit (spool time cut half)
- No wasting of time for skipping existing spool entries on tape when APPENDING using "Space End of Data" command

Tape unit	Search for End-Of-Data to begin APPENDING new Offload Entries	Time to Spool 206 new entries at End-of-data
3592	About 1 second	102 seconds
2490E	307 seconds (not supported)	204 seconds

Documentation

§ z/VSE homepage:

- <http://www.ibm.com/servers/eserver/zseries/zvse/>

§ VSE Performance:

- <http://www.ibm.com/servers/eserver/zseries/zvse/documentation/performance.html>

§ z/VM homepage:

- <http://www.ibm.com/vm>

§ z/VM Performance:

- <http://www.vm.ibm.com/perf/>

§ z/VM Preferred Guest Migration Considerations

- <http://www.vm.ibm.com/perf/tips/z890.html>

§ IBM System z Software Pricing

- <http://www.ibm.com/servers/eserver/zseries/swprice/>

§ IBM's MSU ratings for IBM System z

- <http://www.ibm.com/servers/eserver/zseries/library/swpriceinfo/hardware.html>