



IBM System z

Technical University 2011

System z Hardware Exploitation in z/VSE

zDG07

Ingolf Salm
salm@de.ibm.com

Trademarks

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

IBM*
IBM Logo*

* Registered trademarks of IBM Corporation

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.

INFINIBAND, InfiniBand Trade Association and the INFINIBAND design marks are trademarks and/or service marks of the INFINIBAND Trade Association.

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.

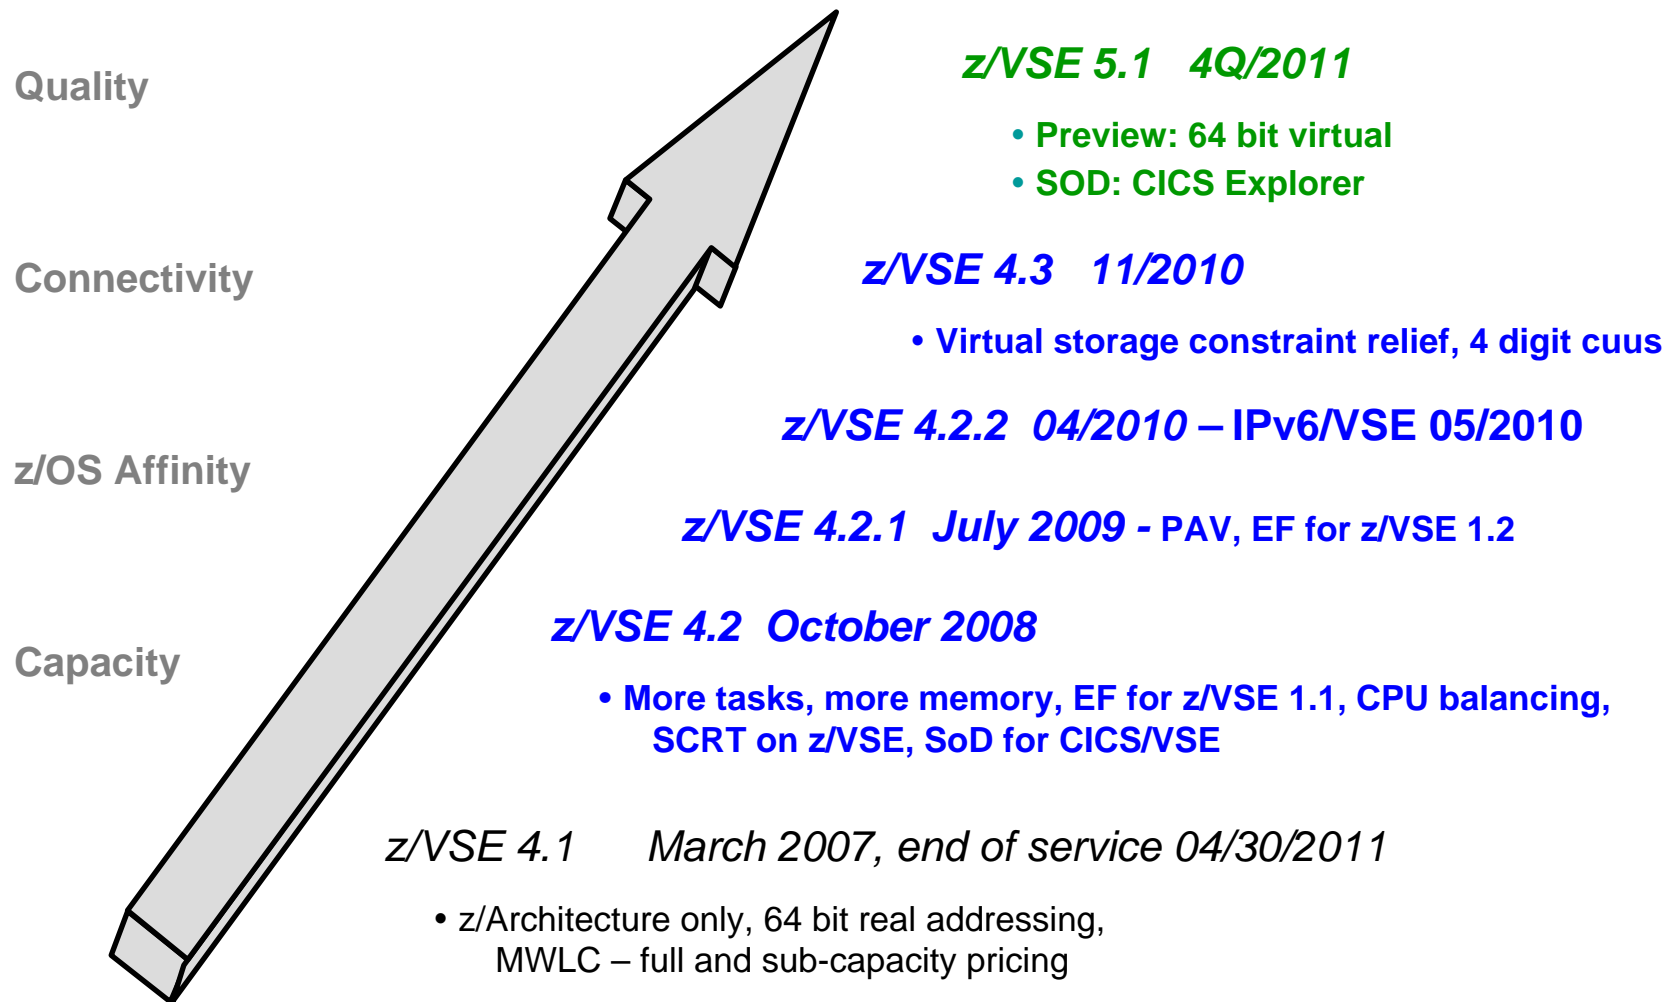
Information about non-IBM products is obtained from the manufacturers of those products or their published announcements. IBM has not tested those products and cannot confirm the performance, compatibility, or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products 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.

Agenda

- Roadmap
- VSE strategy
- z/VSE 4.3
- Processor support
- Device support

VSE Roadmap



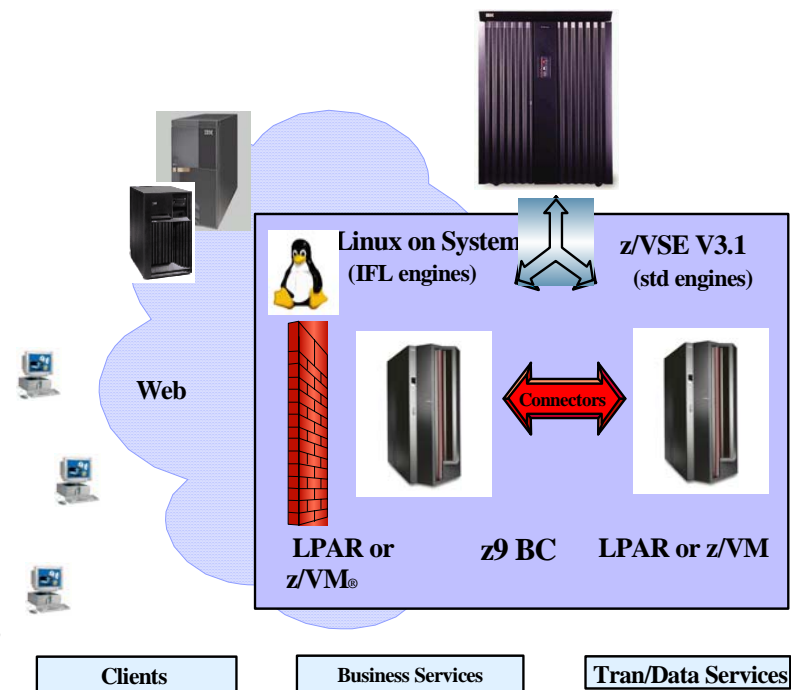
VSE Strategy

- Helps **Protect** your existing investments in core VSE programs, data, equipment, IT skills, *plus* business processes, end user training, etc.
 - modernize, i.e. extend VSE resources to Web
 - exploit **IBM servers, storage** and software

- **Integrate** VSE with the rest of your IT based on open and industry standards
 - IBM middleware
 - VSE connectors and web services

- **Extend** with Linux on System z
 - infrastructure consolidation/simplification
 - add new infrastructure and/or line-of-business applications

Why Not Think Inside the Box?



z/VSE 4.3

- Preview: 10/2009, planned GA: 4Q/2010
- z/VSE 4.3.1 GA: 08/12/2011

- Virtual storage constraint relief for 24 bit (CICS) programs

- 4 digit device addresses (cuus)

- IBM System z10 exploitation
 - Dynamically add of CPUs
 - Large (1 megabyte) page support
 - FICON Express 8

- IBM System Storage support
 - DS8000 Remote Mirror and Copy (RMC) support through ICKDSF

- Basic Security Manager (BSM) will allow to protect MQ resources

- Monitoring agent based on SNMP (Simple Network Management Protocol)

- Midrange Workload License Charges (MWLC) with sub-capacity mode
 - Sub-Capacity Reporting Tool (SCRT) – available with z/VSE 4.1 and later (z9 / z10 only)

- FSU from z/VSE 4.1 and 4.2

z/VSE 4.3

- Preview: 10/2009, Announced: 10/2010, GA: 11/26/2010
- z/VSE 4.3.1 GA: 07/

- IBM System z10 / z196 exploitation
 - Dynamically add of CPUs
 - Large (1 megabyte) page support
 - Static power saving mode for SCRT (z196)

- Virtual storage constraint relief for 24 bit (CICS) programs

- 4 digit device addresses (CUUs)

- Basic Security Manager (BSM) will allow to protect MQ resources

- Monitoring agent based on SNMP (Simple Network Management Protocol)

- Linux Fast Path

- Midrange Workload License Charges (MWLC) with sub-capacity mode
 - Sub-Capacity Reporting Tool (SCRT) – available with z/VSE 4.1 and later (z9 / z10 / z196 / z114 only)

- FSU from z/VSE 4.1 and 4.2

Supported System z Environments

- z/VSE 4.2 and 4.3 support the following platforms
 - IBM e-server zSeries processors (z800, z900, z890, z990)

- z/VSE 4.2, 4.3 and 5.1 run on the following platforms
 - IBM System z9 (z9 BC, z9 EC)
 - IBM System z10 (z10 BC, z10 EC)
 - IBM System zEnterprise (z114, z196)

and support

- uni- and multiprocessors
- Basic mode (z800, z900 only), as z/VM guest or in LPAR
- z/VSE 4.2, 4.3 and 5.1 (will) run under all supported z/VM releases.



VSE Support for System z

VSE Release	z800 / z900	z890 / z990	System z9 / z10 / z196 / z114	VSE EoS
z/VSE V5.1 (GA 4Q/2011)	No	No	Yes	tbd
z/VSE V4.3	Yes	Yes	Yes	tbd
z/VSE V4.2	Yes	Yes	Yes	10/31/2012
z/VSE V4.1	Yes	Yes	Yes	04/30/2011
z/VSE V3.1	Yes	Yes	Yes	07/31/2009
VSE/ESA V2.7	Yes	Yes	Yes	02/28/2007
VSE/ESA V2.6	Yes	Yes	Yes	03/2006
VSE/ESA V2.5	Yes	No	No	12/2003
VSE/ESA V2.4	Yes	No	No	06/2002
VSE/ESA V2.3	No	No	No	12/2001

64 bit real Addressing

- Processor storage support up to 32 GB
- 64 bit real addressing only, introduced with z/VSE 4.1
- Virtual address/data space size remains at max. 2 GB
- 64 bit virtual addressing not supported
- 64 bit addressing mode not supported for applications or ISVs
- Implementation transparent to user applications
- Performance: 64 bit real can reduce / avoid paging
- Many z/VSE environments can run without a page dataset (NOPDS option)

IBM zEnterprise exploitation

- Large page (1 megabyte page) support for data spaces (z10, zEnterprise)
 - Better exploitation of large processor storage, may improve performance
 - No configuration options required
 - Transparent to applications
 - Not supported in z/VM guests

- Dynamic add of logical CPUs (z10, zEnterprise)
 - Ability to dynamically add logical central processors (CPs) without preplanning
 - Logical processor add from HMC/SE
 - Allows adding CPs to LPAR without re-IPL of the z/VSE system
 - Capacity of the z/VSE V4.3 system may be in-/decreased dependent on workload needs
 - New SYSDEF TD parameters (STARTSBY / STOPSBY) to manage the additional CPs
 - Not supported in z/VM guests

```

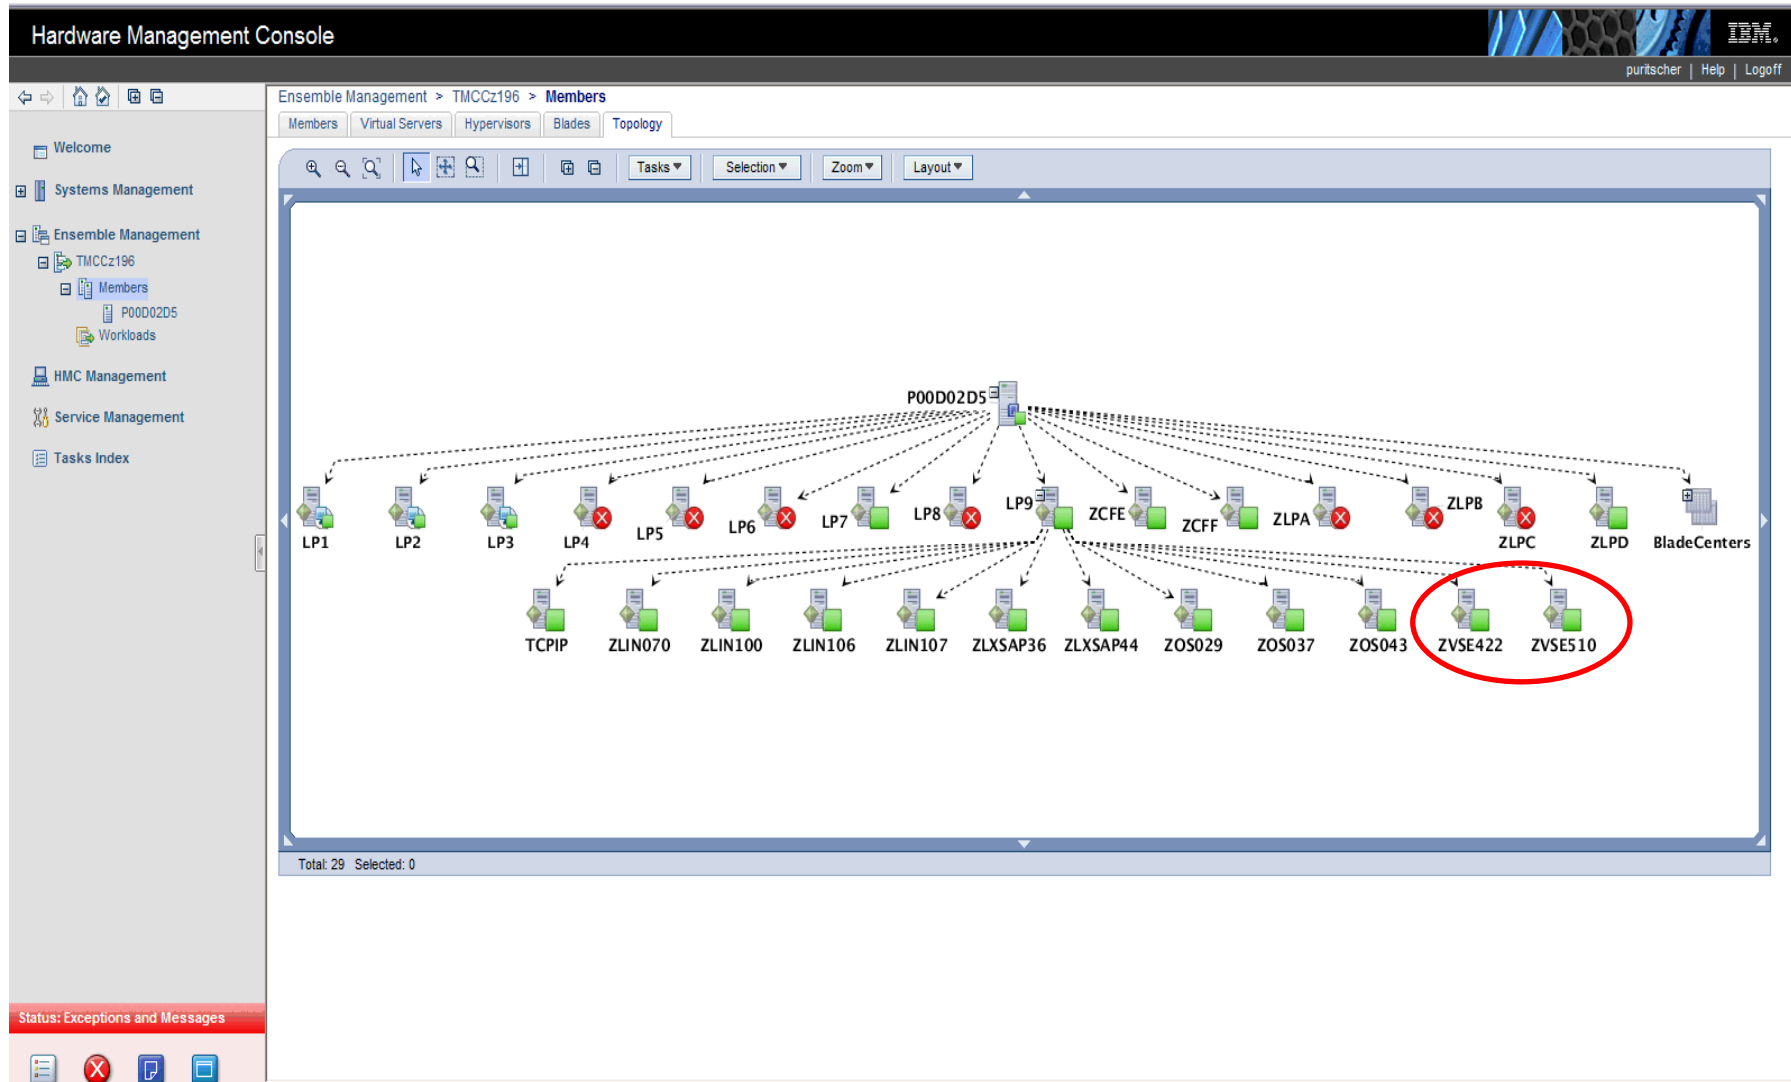
query td
AR 0015 CPU STATUS SPIN_TIME NP_TIME TOTAL_TIME NP/TOT
AR 0015 00 ACTIVE 0 16367 26978 0.606
AR 0015 01 INACTIVE
AR 0015 02 INACTIVE
AR 0015 03 STANDBY
AR 0015
AR 0015 TOTAL -----
AR 0015 0 16367 26978 0.606
AR 0015 NP/TOT: 0.606 SPIN/(SPIN+TOT): 0.000
AR 0015 OVERALL UTILIZATION: 0% NP UTILIZATION: 0%
AR 0015
AR 0015 CPU BALANCING: NOT ACTIVATED
AR 0015
AR 0015 ELAPSED TIME SINCE LAST RESET: 4026069
AR 0015 11401 READY

```

IBM zEnterprise exploitation ...

- Linux Fast Path (LFP) in z/VM mode LPAR (z10, zEnterprise)
- 4096-bit RSA key support with configurable Crypto Express3 (z10, zEnterprise)
- z114 - Statement of general direction: Hipersockets Completion Queue (zEnterprise)
- zEnterprise and zEnterprise BladeCenter Extension (zBX) support
 - “native” Intra Ensemble Data Network (IEDN) – z/VSE 5.1
 - IEDN communication using the z/VM VSWITCH – z/VSE V4 and 5.1
- Static power save mode supported for SCRT (z196 only)

zEnterprise zManager (HMC) and z/VSE



zEnterprise zManager (HMC) and z/VSE

Hardware Management Console

Ensemble Management > TMCCz196 > Members

Members | Virtual Servers | Hypervisors | Blades | Topology

Select	Name	Member	Status	Processors	Memory (MB)	Type	Auto Start	Shutdown Timeout
<input type="checkbox"/>	P00D02D5		Operating			PR/SM		
<input type="checkbox"/>	LP9	P00D02D5	Operating			z/VM	--	300
<input type="checkbox"/>	TCP/IP	P00D02D5	Operating	1	128	z/VM		
<input type="checkbox"/>	ZLIN070	P00D02D5	Operating	2	3,072	z/VM		
<input type="checkbox"/>	ZLIN100	P00D02D5	Operating	2	1,024	z/VM		
<input type="checkbox"/>	ZLIN106	P00D02D5	Operating	4	6,144	z/VM		
<input type="checkbox"/>	ZLIN107	P00D02D5	Operating	1	1,024	z/VM		
<input type="checkbox"/>	ZLXSAP36	P00D02D5	Operating	2	6,144	z/VM		
<input type="checkbox"/>	ZLXSAP44	P00D02D5	Operating	2	6,144	z/VM		
<input type="checkbox"/>	ZOS029	P00D02D5	Operating	4	2,048	z/VM		
<input type="checkbox"/>	ZOS037	P00D02D5	Operating	4	16,384	z/VM		
<input type="checkbox"/>	ZOS043	P00D02D5	Operating	4	16,384	z/VM		
<input checked="" type="checkbox"/>	ZVSE422	P00D02D5	Operating	2	1,024	z/VM		
<input type="checkbox"/>	ZVSE510	P00D02D5	Operating	3	2,048	z/VM		
<input type="checkbox"/>	ZLPA	P00D02D5	Operating			z/VM	--	300

Virtual Server Details: View virtual server details. - Click to launch

Tasks: ZVSE422

- Virtual Server Details
- Toggle Lock
- Daily
 - Activate
 - Deactivate
 - Grouping
- Configuration
 - Delete Virtual Server
 - New Virtual Server Based On
- Monitor
 - Monitor System Events

Status: Exceptions and Messages

zEnterprise zManager (HMC) and z/VSE

Virtual Server Details - ZVSE422 [P00D02D5:LP9:TMCC40]

Name | Status | **Processors** | Memory | Network | Storage | Options | Workloads | Performance

Initial virtual processors: * 2

Maximum virtual processors: * 2

Share limit: None

Initial share mode: Relative

Initial relative shares: * 2000

OK | Apply | Cancel | Help


System z HiperSockets

- “network in the box”,
TCP/IP based communication at near memory speed within one system
 - System z Logical Partitions (LPARs)
 - z/VM guests (via virtual guest LAN)
 - z/VM guests and LPARs

- z/VSE may communicate with
 - Linux on System z
 - z/OS
 - z/VM
 - z/VSE V4 or z/VSE 5.1

- Virtual HiperSockets via z/VM Guest LAN support

HiperSockets Example



System z HiperSockets ...

- z114 enhancement
 - Statement of general direction (see z114 hardware announcement – http://www-01.ibm.com/common/ssi/rep_ca/6/897/ENUS111-136/ENUS111-136.PDF):

HiperSockets Completion Queue:

IBM plans to support transferring HiperSockets messages asynchronously, in addition to the current synchronous manner on z196 and z114. This could be especially helpful in burst situations. The Completion Queue function is designed to allow HiperSockets to transfer data synchronously if possible and asynchronously if necessary, thus combining ultra-low latency with more tolerance for traffic peaks. HiperSockets Completion Queue is planned to be supported in the z/VM and z/VSE environments.

See IBM Hardware Announcement 111-136, July 12, 2011

All statements regarding IBM's plans, directions, and intent are subject to change or withdrawal without notice.

System z Exploitation

- FICON Express8 - Higher I/O bandwidth
- Adapter interruptions (performance improvements)
 - OSA Express2 / OSA Express3 (QDIO mode), FICON Express8 (FCP)
- Open Systems Adapter (OSA Express3) features
 - OSA Express3 10 Gigabit Ethernet (2 ports), Gigabit Ethernet (4 ports)
 - OSA Express3 1000BASE-T Ethernet (4 modes of operation)
 - ICC (Integrated Console Controller)
 - QDIO (Queued Direct I/O) for TCP/IP traffic
 - Non-QDIO for TCP/IP and SNA traffic
 - OSN (Open System Adapter for NCP) works with IBM Communication Controller for Linux on System z
- z/VM queue-I/O assist for real networking devices
 - OSA Express adapters (CHIPID type OSD)
 - Hipersockets (CHIPID type IQD)

OSA Express3 Support

- OSA Express3 – Open Systems Adapter card for z10, z196, z114
 - Successor of OSA Express2

- z/VSE supports the Gigabit Ethernet (GbE) and 10 Gigabit Ethernet (10 GbE) features
 - To be configured in IOCDS as CHPID type OSD (other CHPID types not supported)
 - Exploited by TCP/IP via DEFINE LINK,TYPE=OSAX command
 - OSA Express3 10 GbE (2 ports), GbE (4 ports)

- Port specification for TCP/IP
 - OSA Express2 and OSA Express3 10 GbE features:
 - only one port per CHPID to connect to the network
 - OSA Express3 GbE: two ports per CHPID – port 0 and port 1
 - To use port 0, no port specification is necessary
 - To use port 1, the port needs to be specified, e.g.:
`DEFINE LINK,TYPE=OSAX,DEV=D00,DATAPATH=D02,OSAPORT=1`

System z hardware cryptographic support

- Enhances Internet security
- Encryption support via crypto cards or on the processor itself (CPACF)
- Cryptographic assists
 - Exploited by the SSL support of TCP/IP for VSE/ESA's transparently
- Transparent for "TCP/IP for VSE/ESA" applications
 - VSE connector server, CWS, VSE/Power PNET
- No definition necessary

System z hardware cryptographic support

- CPACF – for symmetric encryption
 - AES for 128-bit keys (z9 EC, z9 BC), AES for 256 keys (z10 EC or higher)

- Crypto Express2 / Crypto Express3 (crypto card) – for asymmetric encryption
 - Encryption hardware assist for increased SSL throughput
 - Supports SSL handshaking only for applications that use the SSL crypto API

 - 2048-bit RSA key with Crypto Express2
 - 4096-bit RSA key support with configurable Crypto Express3 (z/VSE 4.3 or higher)

 - Configurable Crypto Express2
 - Dynamically configurable in coprocessor or accelerator mode

 - Dynamic change of cryptographic processors
 - Add/remove cryptographic processor of z10 LPAR or higher

 - AP (adjunct processor)-queue adapter-interruption facility
 - May accelerate the SSL throughput
 - Available on Crypto Express2 or Crypto Express3 feature

Signal Quiesce (Signal Shutdown) Support

- If e.g. an IML or IPL is performed via the HMC / SE or z/VM SIGNAL SHUTDOWN , a signal-quiesce event is generated.
- Need to be enabled via IPL SYS QUIESCE=YES | NO
- If QUIESCE=YES a message is generated:

0W01D DO YOU WANT TO CONTINUE SYSTEM SHUTDOWN (WILL BE FORCED AFTER TIMEOUT)?
REPLY 'YES' TO ENTER HARD WAIT STATE OR 'NO'

- If the operator reply is **yes**,
 - The system will enter the disabled wait state
- If the operator reply is **no** or does not reply, the system will wait for a predefined time interval
 - Console automation can initiate a controlled system shutdown
- z/VSE does not provide controlled shutdown processing

Exploitation of IBM System Storage Products

- IBM System Storage TS1130 / TS1120 Tape Drive
- IBM System Storage TS7700 / TS7720 Virtualization Engine
- IBM System Storage TS3400 autoloader Tape Library
- IBM System Storage TS3500 Tape Library
- IBM TS7680 ProtecTIER Deduplication Gateway for System z
 - Disk-only virtual tape solution
- zVSE supports the S/390 channel command interface via
 - Perform Subsystem Function (PSF)
 - Perform Library Function (PLF) commands

Exploitation of IBM System Storage Products ... Large Volume (64K cylinder) support

- IBM System Storage DS8000/DS6000 64K cylinder support
 - Supported by BAM and VSE/VSAM
 - Allows consolidation of smaller disks volumes

- VSAM supports more than 1,500 clusters per catalog

- VSAM FAT-BIG DASD support
 - **Small DASD** (normal): smaller than 64k tracks per volume
 - 3390 in LISTCAT
 - Large DASD with two subtypes:
 - **Big DASD**: more than 64k tracks per volume
BIG-3390 in LISTCAT
Support of up to 10017 cylinders
 - **Fat DASD**: up to 64k cylinders
FAT-3390 in LISTCAT
New type of volume

Parallel Access Volume (PAV)

- Optional licensed feature of DS8000, DS6000, ESS series
- Enables z/VSE to simultaneous process multiple I/O operations to the same volume
 - Can provide enhanced throughput
 - Can help to consolidate small volumes to large volumes
- Multiple logical addresses to the same physical device
 - = Base and alias volumes for concurrent processing of I/O operations
 - Configuration in DASD, IOCDs and z/VSE
 - Base device: physical device to be added during IPL
 - Alias device(s) are associated to the base device.
 - z/VSE supports up to 7 alias devices
- Multiple z/VSE jobs can transfer data to or from the same physical volume in parallel
- All z/VSE references to I/O devices (e.g. in JCL) relate to the base device
- In z/VSE PAV processing can be dynamically activated or deactivated via the AR/JCL command `SYSDEF PAV=START` or `STOP`
- Max. 1023 I/O devices can be added, if PAV to be activated
- Delivered via PTF (included in z/VSE 4.2.1)

Parallel Access Volume (PAV) ...

- Multiple z/VSE jobs can transfer data to or from the same physical volume in parallel
 - A VSE task can only process one I/O at a time

- PAV processing can be dynamically
 - Activated via the AR/JCL command `SYSDEF PAV=START`
 - PAV devices will be sensed during START processing

sysdef system,pav=Start

```
AR 0028 1K01I ALIAS DEVICE 778 FOR BASE 777 ESTABLISHED
AR 0028 1K01I ALIAS DEVICE 77E FOR BASE 777 ESTABLISHED
AR 0028 1K01I ALIAS DEVICE 779 FOR BASE 777 ESTABLISHED
AR 0028 1K01I ALIAS DEVICE 77A FOR BASE 777 ESTABLISHED
AR 0028 1K01I ALIAS DEVICE 77B FOR BASE 777 ESTABLISHED
AR 0015 1I40I READY
```
 - Deactivated via the AR/JCL command `SYSDEF PAV=STOP`
 - All I/Os on alias devices need to complete, delays possible.

- `QUERY SYSTEM` and `SIR AR` commands show, if PAV is active
 - E.g. `QUERY SYSTEM: AR 0015 PARALLEL ACCESS VOLUME (PAV): ACTIVE / INACTIVE`

Parallel Access Volume (PAV) ...

- VOLUME AR command indicates a base volume via a “*B” in the CODE column:

VOLUME

```
AR 0015 CUU CODE DEV.-TYP VOLID  USAGE  SHARED STATUS CAPACITY
....
....
AR 0015 261 6E    2107-900  VIS001 UNUSED          20 CYL
AR 0015 262 6E    2107-900  VIS002 UNUSED          20 CYL
AR 0015 777 6E*B 2105-000  FRA740 UNUSED        10017 CYL
AR 0015 1I40I READY
```

- VOLUME AR command, DETAIL parameter

VOLUME 777,DETAIL

```
AR 0015 CUU CODE DEV.-TYP VOLID  USAGE  SHARED STATUS CAPACITY
AR 0015 777 6E*B 2105-000  FRA740 UNUSED        10017 CYL
AR 0015          BASE TO 778,77E,779,77A,77B
AR 0015 1I40I READY
```

Parallel Access Volume (PAV) ...

- Alias device cuus can not be added at IPL
 - ADD statement with an alias cuu will be ignored.
- 3 digit cuu limitation apply to base and alias devices
- Planning information:
 - Alias devices make use of copy blocks, channel queue entries and other I/O resources



Parallel Access Volume (PAV) ...

- Subsystem Monitoring Facility (SMF) shows I/O distribution to base and alias devices
 - Documentation available on ibm.com/vse
 - Hints and Tips for z/VSE
 - PAV White Paper (available with PAV PTF)

SIR SMF,VSE,777

AR 0015 TIMINGS FOR 777 BASED ON	2368 I/O INSTRUCTION				
AR 0015 BASE 777	0 I/O INSTRUCTION				
AR 0015 ALIAS 778	0 I/O INSTRUCTION				
AR 0015 ALIAS 77E	208 I/O INSTRUCTION				
AR 0015 ALIAS 779	310 I/O INSTRUCTION				
AR 0015 ALIAS 77A	832 I/O INSTRUCTION				
AR 0015 ALIAS 77B	1018 I/O INSTRUCTION				
AR 0015					
AR 0015 QUEUED	PENDING	CONNECT	DISCONN	DEV.BUSY	TOTAL
AR 0015 msec/SSCH	msec/SSCH	msec/SSCH	msec/SSCH	msec/SSCH	msec/SSCH
AR 0015 0.000	0.000	0.812	0.000	0.000	0.812
AR 0015 1I40I	READY				

z/VSE 4.3: 4 digit CUUs

- Ease of use and infrastructure simplification
 - In mixed environments running z/VSE together with z/VM, Linux on system z or z/OS
 - Removes the requirement for a z/VSE specific IOCDS configuration
 - Provides more flexibility

- 4 digit CUUs transparent to applications and most system programs
 - Implemented via mapping to 3 digit CUUs during IPL
 - z/VSE will only use 3 digit CUUs after IPL complete
 - Exception: z/VM DIAG instruction use 4 digit CUUs

FlashCopy Support

- Available on DS8000. DS6000 and ESS
- Source and copied data almost available immediately
- NOCOPY option
 - Direct copy to backup device
- Dataset Copy
 - Source and target volumes may have different sizes
 - Should not be used for VSAM files
- Elimination of Logical Subsystems
 - Source and target volume can span LSS
- Multiple relationship FlashCopy
 - Up to 12 volumes from one source in a single FlashCopy operation

FlashCopy Support

- IBM System Storage DS8000 FlashCopy SE (Space Efficient)
 - Allocates storage on target volume only “as-needed”, if copied tracks from source volume

- FlashCopy Consistency Group
 - Allows to create a consistent point-in-time copy across multiple volumes

- DS8000 Remote Mirror and Copy (RMC) support through ICKDSF

- Peer-to Peer Remote Copy (PPRC)
 - Allows remote data replication
 - Supported via ICKDSF

- z/VSE does not support:
 - Incremental FlashCopy
 - Persistent FlashCopy relationship
 - Inband Commands over Remote Mirror link

SCSI Support in z/VSE

- SCSI disks as emulated FBA disks on z/VM V5.2 or higher
 - z/VSE supports a max. size of 2 GB
- Direct attached SCSI disks
 - z/VSE supports up to 24 GB (VSAM: 16 GB)
 - z/VSE supports SCSI disk devices only
 - Impact on applications
 - Transparent to all VSE applications and subsystems,
 - Reasons for transparency:
 - z/VSE's SCSI implementation is based on FBA support
 - applications can not exploit SCSI commands directly
 - FBA to SCSI emulation on low level I/O interface
- SAN Volume Controller (SVC)
 - To access FCP-SCSI disks in DS8000, DS6000, DS4000 and ESS series as well as disk subsystems from other manufacturers supported by SVC
- IBM XIV Storage System
- IBM Storwize V7000 Midrange Disk System

SCSI Support in z/VSE

- Access SCSI devices through Fibre Channel Protocol (FCP)
- z/VSE's SCSI support includes:
 - SCSI for system and data device (SCSI only system)
 - Multipathing for fail-over
- SCSI support transparent to existing (I/O) APIs
- SCSI disk devices utilize fixed block sectors
 - Block size restricted to 512 bytes,
even if the SCSI device can be configured with larger block sizes
- FSU from SCSI to SCSI device only

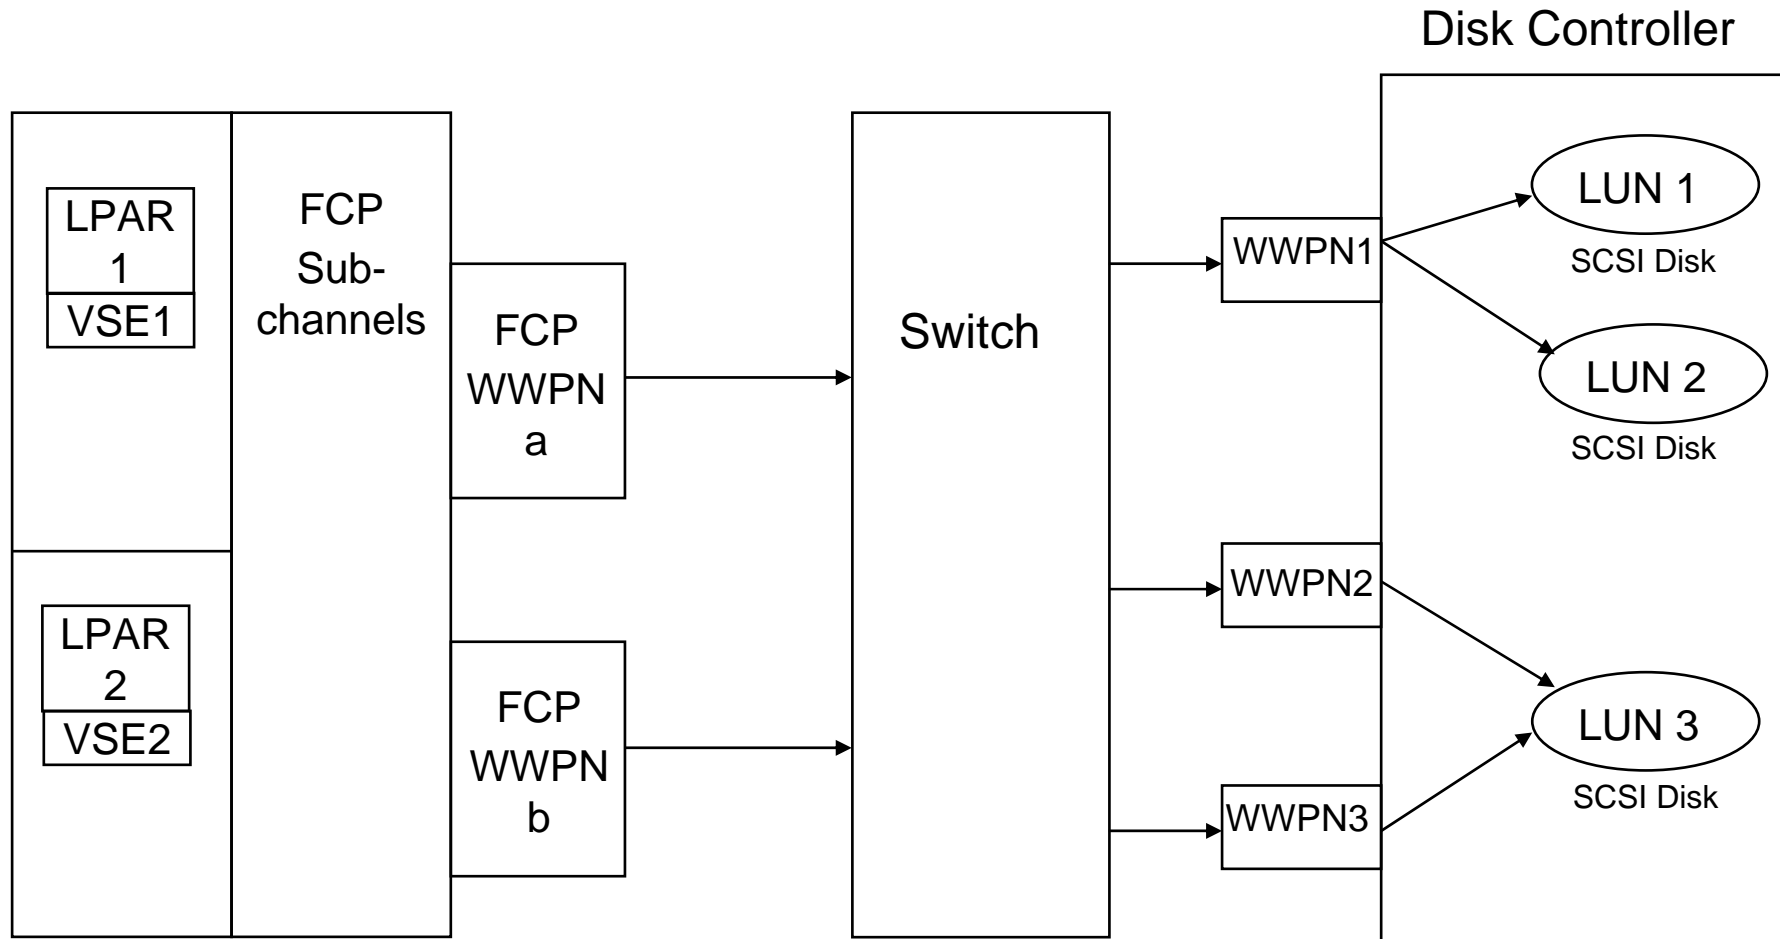
SCSI Support - Configuration

- New IPL / JCL commands and dialog to define and query a SCSI device
- Required steps to get a SCSI device known to z/VSE
 - Device configuration
 - Switch configuration
 - In case of point to point connections (System z9 or higher) no longer necessary
 - FCP Adapter to be configured in IOCDS (CHIPID type FCP)
 - FCP adapter and SCSI disk to be defined in VSE via
 - IPL ADD commands to define FCP and FBA device
 - IPL DEF or JCL SYSDEF command to define connection to LUN

SCSI Configuration in z/VSE (Example)


- Define FCP Devices, SCSI Disks and Connection Paths to z/VSE
 - FCP Devices
 - ADD C00,FCP (ADD C00:C0F,FCP)
 - ADD D00,FCP (ADD D00:D0F,FCP)
 - FBA Devices:
 - ADD 700:701,FBA
 - Note: these devices must not exist in the IOCP or under VM
 - Define a Connection Path (IPL)
 - DEF SCSI,FBA=700,FCP=C00,WWPN=5005076300CA9A76,LUN=5600
 - DEF SCSI,FBA=701,FCP=C00,WWPN=5005076300CA9A76,LUN=5601
 - Only one FCP cuu required to access the LUNs
 - Define a Connection Path (after IPL)
 - SYSDEF SCSI,FBA=702,FCP=C00,WWPN=5005076300CA9A76,LUN=5602
 - Note: The FBA and FCP devices added during IPL.
 - IUI Dialogs are available to configure SCSI Devices

SCSI Support – Disk Controller Configuration



Point to point connection
with z9/z10 possible

SCSI Support – SVC Configuration



SCSI Support

- N_Port ID Virtualization (NPIV) for (CHPID type) FCP channels (z9 or higher)
 - Multiple virtual FCP channels can be defined each with its own unique Fibre Channel port name and FC N_Port ID
 - NPIV allows sharing the Lock file on SCSI between multiple z/VSE systems using the same physical FCP adapter (CHPID)
 - DEF SCSI,FBA=600,FCP=C00,WWPN=5005076300CA9A76,LUN=5750 (VSE1)
 - DEF SCSI,FBA=600,FCP=C01,WWPN=5005076300CA9A76,LUN=5750 (VSE2)

600 is the lock file disk. With NPIV, C00 and C01 can be on same FCP CHPID
 - To use NPIV, the Fibre Channel switch must support NPIV
 - Without NPIV,
 - Each FCP channel(device) has the portname of the FCP CHPID
 - Each z/VSE needs its own physical FCP adapter to access the lock file
- FCP point-to-point attachments (z9 or higher)
 - FCP feature can directly attach to storage devices. No switch required.

Data Encryption

- IBM TS1120 / TS1130 Tape Drive with encryption feature
 - Supports data encryption within the drive itself
 - Using Systems Managed Encryption with the TS1120 / TS1130
 - z/VSE support will require the **Encryption Key Manager** component running on another operating system other than z/VSE using an out-of-band connection.
 - Generation and communication of encryption keys for tape drive
 - TCP/IP connection between EKM and the tape controller
 - Data encryption is transparent to z/VSE applications
 - Data encryption
 - Data will be encrypted and compressed, when specified
 - Default: encryption disabled
 - Encryption re-keying support to encrypt data key of encrypted tape cartridge

Data Encryption ...

- Encryption Key Manager (EKM)
 - EKM is a Java application, used to generate and protect AES keys
 - On request EKM generates AES (256 bit) data keys and protects those keys
 - Key encryption key label (KEKL) identifies the encryption keys
 - The KEKL or the hash value of the public key can be stored on the cartridge.

 - You may download EKM from the internet

- In z/VSE jobs must have an ASSGN statement and KEKL statement to access or write encrypted data

- ASSGN statement
 - ASSGN SYSnnn, cuu, mode
 - cuu = device address
 - mode =
 - 03 encryption write mode
 - 0B encryption and IDRC write mode
 - 23 encryption and unbuffered (compression) write mode
 - 2B encryption and IDRC and unbuffered write mode

- KEKL statement
 - // KEKL UNIT=cuu, KEKL1=key_label_1, KEM={L|H}
 - KEM = key encoding mechanism
 - L = label, H = public key hash

Data Encryption ...

- Write encryption data example
 - // JOB ENCRYPT
 - // ASSGN SYS005,480,3
 - // KEKL UNIT=480, KEKL1=,HUSKEK1',KEM1=L
 - // EXEC LIBR
 - BACKUP LIB=PRD2 TAPE=SYS005
 - /*
 - /&

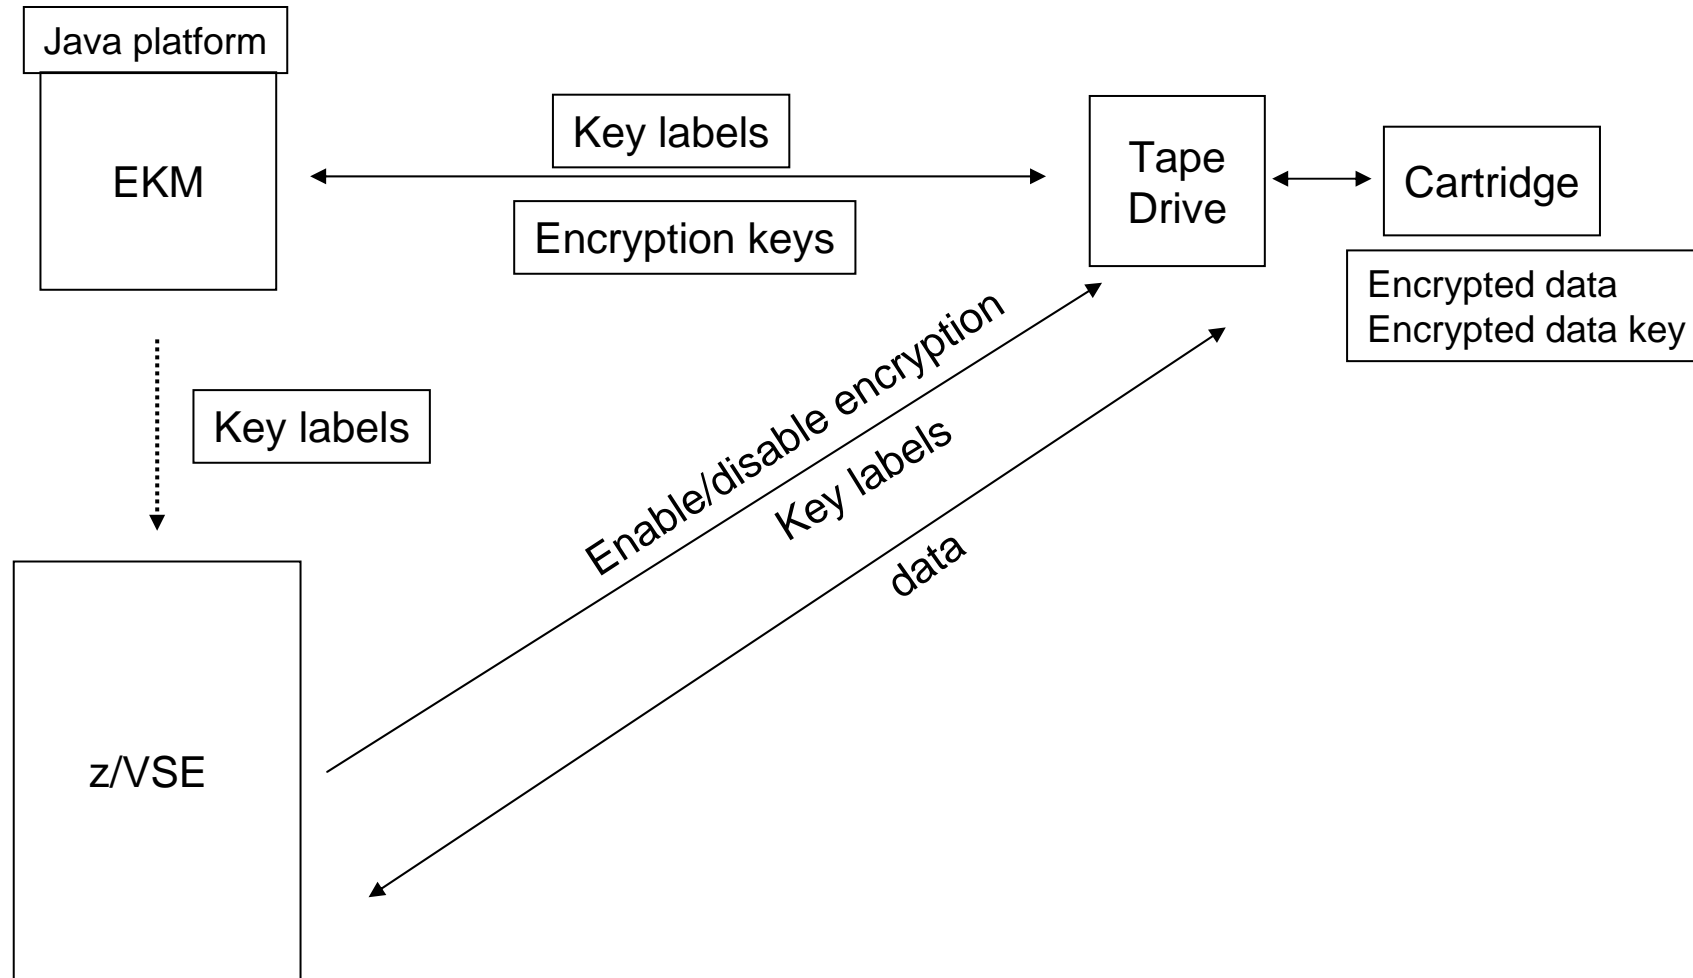
- Read encrypted data
 - No need to specify the ASSGN mode or KEKL
 - The control unit recognizes the encrypted tape and tries a key exchange with the EKM and the KEKL saved in cartridge memory.

Data Encryption ...

- Steps to encryption
 1. Load cartridge
 2. EKM to tape drive: specify encryption, provide key labels
 3. Tape drive requests data key from EKM
 4. EKM generates key and encrypts with public and session keys
 5. EKM to tape drive: Encrypted keys transmitted
 6. Tape drive writes encrypted data and stores encrypted data key on cartridge

- Implementation in z/VSE
 - VSE JCL enhancements
 - For encryption setting (via ASSGN)
 - Key Encryption Key Label (KEKL) may be specified
 - I/O Supervisor
 - retrieves encryption information, activates encryption and transfers KEKL

Data Encryption ...



More Information

... on VSE home page: <http://ibm.com/vse>

- Hints and Tips for z/VSE:

<ftp://public.dhe.ibm.com/eserver/zseries/zos/vse/pdf3/zvse42/hintamm2.pdf>

- z/VSE Planning, SC33-8301-03