z/VSE Hardware Exploitation

Ingolf Salm
salm@de.ibm.com
Ingolf's z/VSE Blog: https://www.ibm.com/developerworks/mydeveloperworks/blogs/vse

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.

VSE Roadmap

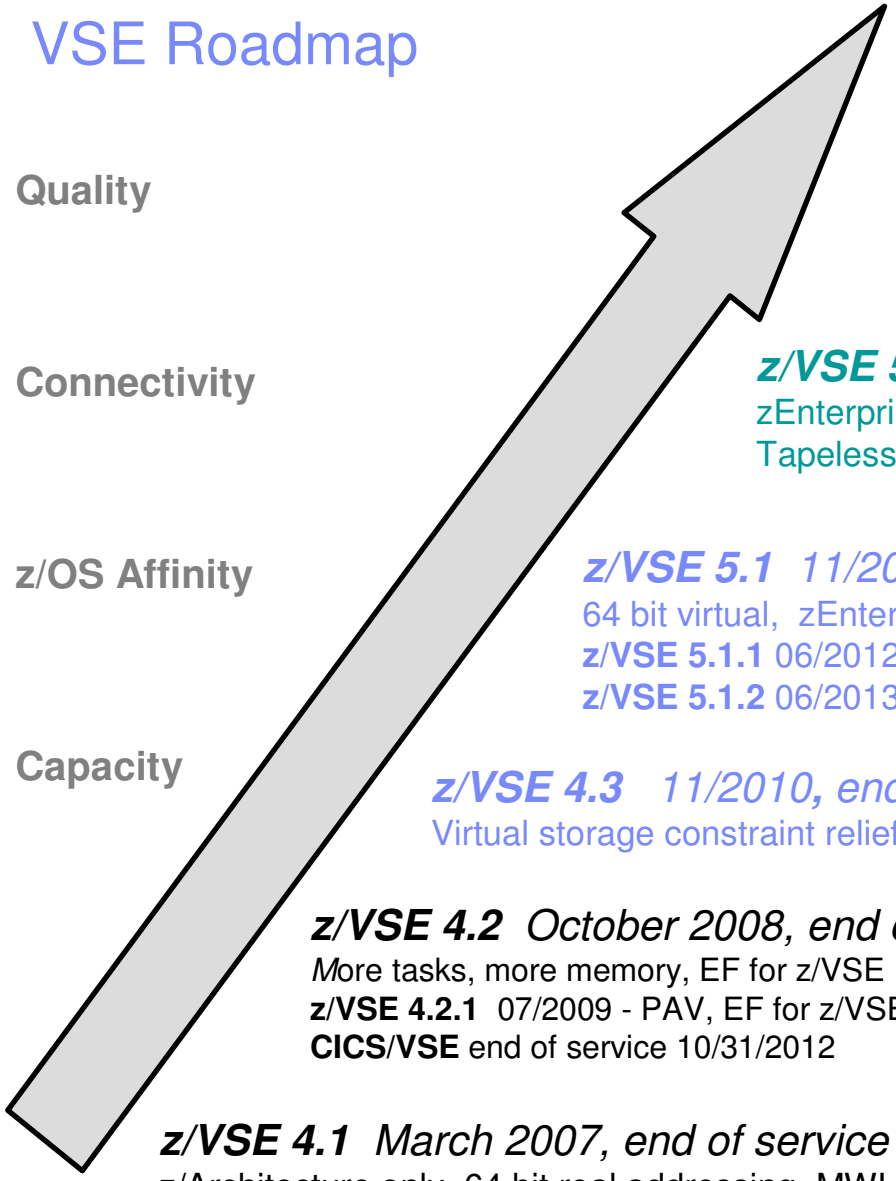

z/VSE Statement of Direction (SOD)
CICS Explorer update, Cannels & Containers,
new CICS TS for z/VSE release,
next z/VSE release: ALS to z10 or higher

z/VSE 5.2 *Ann: 04/07/2014, GA 04/25/2014*
zEnterprise exploitation, device support
Tapeless installation, networking / security enhancements

z/VSE 5.1 *11/2011*
64 bit virtual, zEnterprise exploitation, z9 or higher
z/VSE 5.1.1 06/2012: CICS Explorer, LFP in LPAR, database connector
z/VSE 5.1.2 06/2013: TS1140, 64 bit I/O, openSSL, db connector enhancements

z/VSE 4.3 *11/2010, end of service 10/31/2014*
Virtual storage constraint relief, 4 digit cuus, z/VSE 4.3.1 08/2011

z/VSE 4.2 *October 2008, end of service 10/31/2012*
*M*ore tasks, more memory, EF for z/VSE 1.1, CPU balancing, SCRT on z/VSE
z/VSE 4.2.1 07/2009 - PAV, EF for z/VSE 1.2, **z/VSE 4.2.2** 04/2010 - IPv6/VSE 05/2010
CICS/VSE end of service 10/31/2012

z/VSE 4.1 *March 2007, end of service 04/30/2011*
z/Architecture only, 64 bit real addressing, MWLC – full and sub-capacity pricing

z/VSE 5.2 – Quick Overview

- Announcement: 04/07/2014, GA: 04/25/2014
- Hardware support
  - IBM System z Enterprise support
  - Device support
    - Tape, ECKD and FCP-attached SCSI disks
- 64 bit virtual exploitation
  - Virtual disk in memory objects
- Networking enhacements
  - IPv6 support for selected z/VSE functions
- Security enhancements
  - Basic Security manager (BSM) and VSE/POWER audit enhancements
- Ease of use
  - Tapeless installation from ECKD devices
  - Stacking tape support
- Fast Service Upgrade (FSU) from z/VSE 4.3 and z/VSE 5.1

- Pricing
  - z9, z10, z196, zEC12: Midrange Workload License Charge (MWLC) pricing with sub-capacity option
  - z114, zBC12: Advanced Entry Workload License Charge (AEWLC) pricing with sub-capacity option

z/VSE 5.2 – Quick Overview ...

- Support for IBM zEnterprise EC12 and IBM zEnterprise BC12
  - Configurable Crypto Express4S feature
  - OSA-Express5S features
  - HMC based configuration for OSA-Express4 and OSA-Express5S (OSA/SF)
- Support for IBM System Storage
  - Tape support
    - Systems Managed Encryption with IBM System Storage TS1140
    - IBM System Storage TS7700 Virtualization Engine Release 3.1
  - ECKD / FCP-attached SCSI disk support
    - IBM System Storage DS8870 Release 7.2
    - Upgrade of the z/VSE support for the Parallel Access Volume (PAV) feature (ECKD)
  - FCP-attached SCSI disk support
    - IBM Storwize V7000
    - IBM Storwize V5000 Midrange Disk
    - IBM Storwize V3700 Entry Disk

IBM System z server / z/VM support

- z/VSE V5 supports IBM System z servers:

  – IBM zEnterprise EC12 (zEC12)
  – IBM zEnterprise BC12 (zBC12)
  – IBM zEnterprise 196 (z196)
  – IBM zEnterprise 114 (z114)
  – IBM System z10 (z10 EC, z10 BC)
  – IBM System z9 (z9 EC, z9 BC)

  ... and z/VSE V5 can run in an LPAR or as a z/VM guest on all supported z/VM releases
  ... in uni- or multiprocessor mode

  Please see the statement of direction in the z/VM 6.3 announcement (July 2013):

  Stabilization of z/VM V5.4 support: The IBM zEnterprise EC12 and IBM zEnterprise BC12 are planned to be the last System z servers supported by z/VM V5.4 and the last System z servers that will support z/VM V5.4 running as a guest (second level). z/VM V5.4 will continue to be supported until December 31, 2014, or until the IBM System z9 EC and IBM System z9 BC are withdrawn from support, whichever is later.

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

VSE Support for IBM System z

VSE Release	z800 / z900	z890 / z990	System z9 / z10 / z196 / z114 / zEC12 / zBC12	VSE EoS
z/VSE V5.2	No	No	Yes	tbd
z/VSE V5.1	No	No	Yes	tbd
z/VSE V4.3	Yes	Yes	Yes	10/31/2014
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

IBM zEnterprise exploitation

- 64 bit real addressing - up to 32 GB (System z), 64 bit virtual addressing – up to 90 GB
- Large page support (z10, zEnterprise)
- Dynamic add / remove of logical CPs (z10, zEnterprise)
- OSA-Express 3, OSA-Express 4, OSA-Express 5S support
- HiperSockets Completion Queue on z196, z114, zEC12, zBC12 (z/VSE 5.1.1 and higher)
- Linux Fast Path (LFP) in z/VM mode LPAR (z10, zEnterprise)
- Exploitation of the z/VSE z/VM IP Assist (zEnterprise)
- zEnterprise and zEnterprise BladeCenter Extension (zBX) support
  - Intra Ensemble Data Network (IEDN)
  - Virtual LAN support, Layer 2 support
  - IEDN communication using the z/VM VSWITCH
- 4096-bit RSA key support with configurable Crypto Express3 (z10, zEnterprise)
  …. and Crypto Express4S (zEC12, zBC12) – z/VSE V5 only
- Static power save mode supported for SCRT (z196, zEC12)

- **zEC12 / zBC12 do not support ESCON channels**

IBM zEnterprise exploitation

- Following functions are not supported in z/VM guests:
- 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
- Dynamic add of logical CPs (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

```
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                       0       16367       26978   0.606
AR 0015
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 1I40I  READY
```

Linux Fast Path (LFP)

- Routes IPv4 or IPv6 socket request to Linux on System z
  - Without using the local TCP/IP stack
- LFP on z/VM (z/VSE 4.3 or higher)
  - Uses an IUCV connection between z/VSE and Linux on System z
  - Both – z/VSE and Linux – need to be z/VM guests of the same z/VM
- Linux Fast Path using z/VSE z/VM IP Assist (VIA – z/VSE 5.1)
  - Both – z/VSE need to be a z/VM guests
- Linux Fast Path in LPAR (z/VSE 5.1 + enhancements – GA 06/15/2012)
  - LFP daemon on Linux forwards the socket request to the Linux TCP/IP stack
- LFP is transparent to IBM socket APIs
  - Supported APIs: LE/C socket API, EZA socket / EZASMI interface, …
  - Transparent to IBM applications (DB2 client, Connectors, Power PNET)
  - No standard TCP/IP applications (Telnet, FTP, …) provided
  - IPv6/VSE: TCP/IP applications can exploit LFP
- Provided with the z/VSE base product – no additional charge

LFP - z/VSE z/VM IP Assist (VIA) – z/VSE 5.1

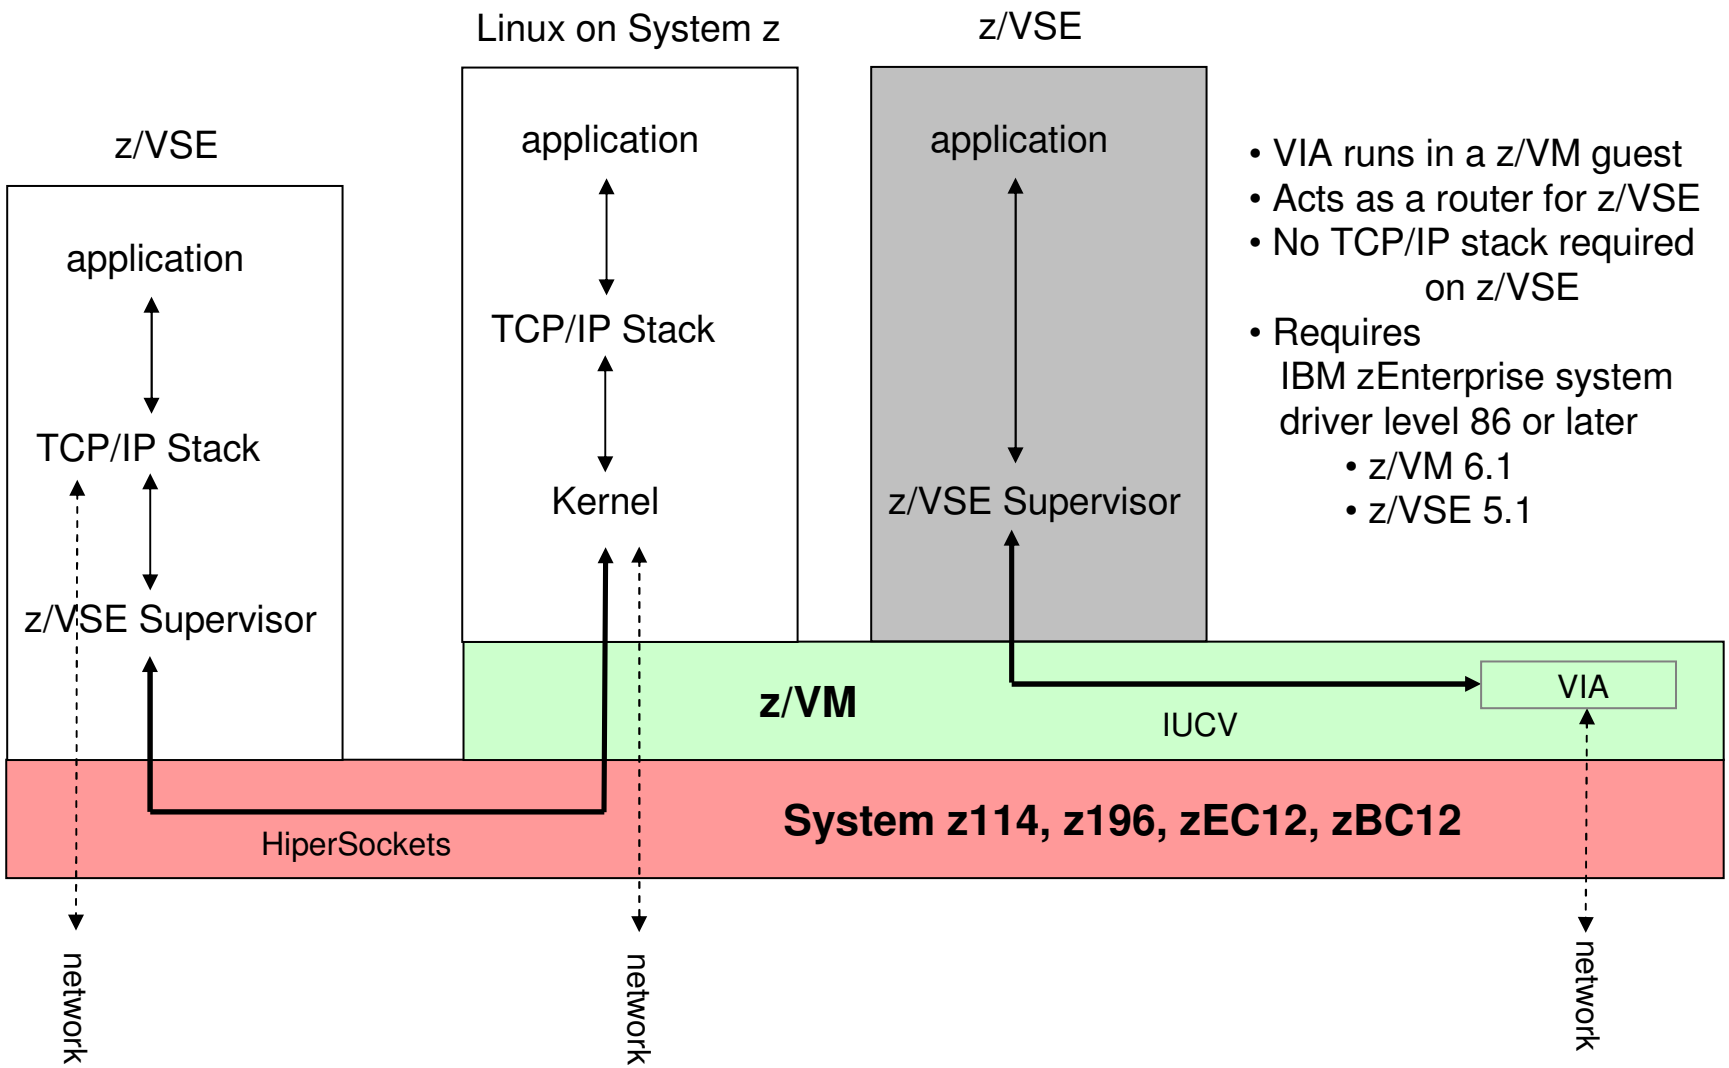

- VIA runs in a z/VM guest
- Acts as a router for z/VSE
- No TCP/IP stack required on z/VSE
- Requires IBM zEnterprise system driver level 86 or later
  - z/VM 6.1
  - z/VSE 5.1

Linux Fast Path in LPAR

- No TCP/IP stack required on z/VSE
- System requirements
  - Supported on zEnterprise
    - Exploits HiperSockets completion queue
  - Linux on System z distribution (RHEL, SLES)
  - Available with z/VSE 5.1.1 (z/VSE 5.1 + PTF)

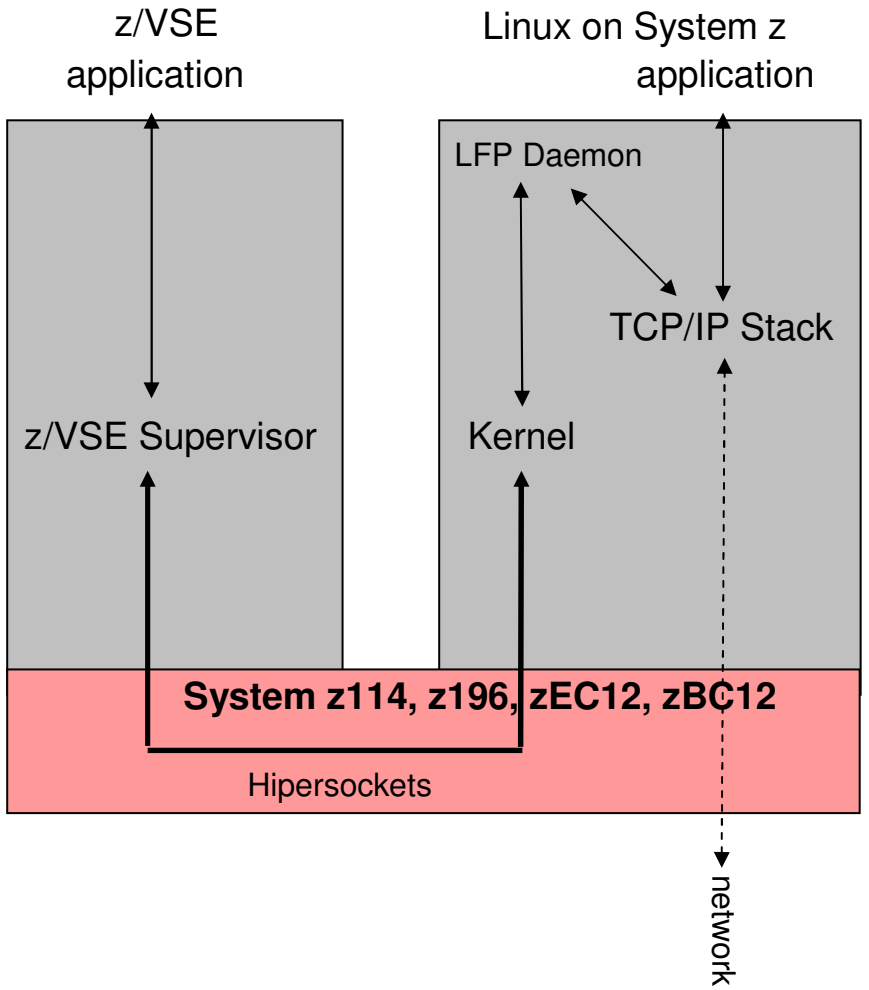

64 bit real addressing

- Processor storage support up to 32 GB
- 64 bit real addressing only, introduced with z/VSE 4.1
- z/VSE 5.1
  - Virtual address space > 2 GB
  - 64 bit virtual addressing
- Data space size remains at max. 2 GB
- 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)

64 bit virtual

- Support 64 bit virtual addressing
- 64 bit area can be used for **data only**
  - No instruction execution above the bar
- **z/OS affinity:** APIs (IARV64 services) - to manage memory objects – compatible with z/OS
  - Private memory objects for use in one address space
  - Shared memory objects to be shared among multiple address spaces
- Maximum VSIZE still limited to 90 GB
- Advantages:
  - Eases the access of large amounts of data
    - E.g. instead of using and managing data spaces
  - Reduces complexity of programs
    - Data contained in primary address space
  - Chosen design has no dependencies to existing APIs, minor impact on existing system code

64 bit virtual - address space layout

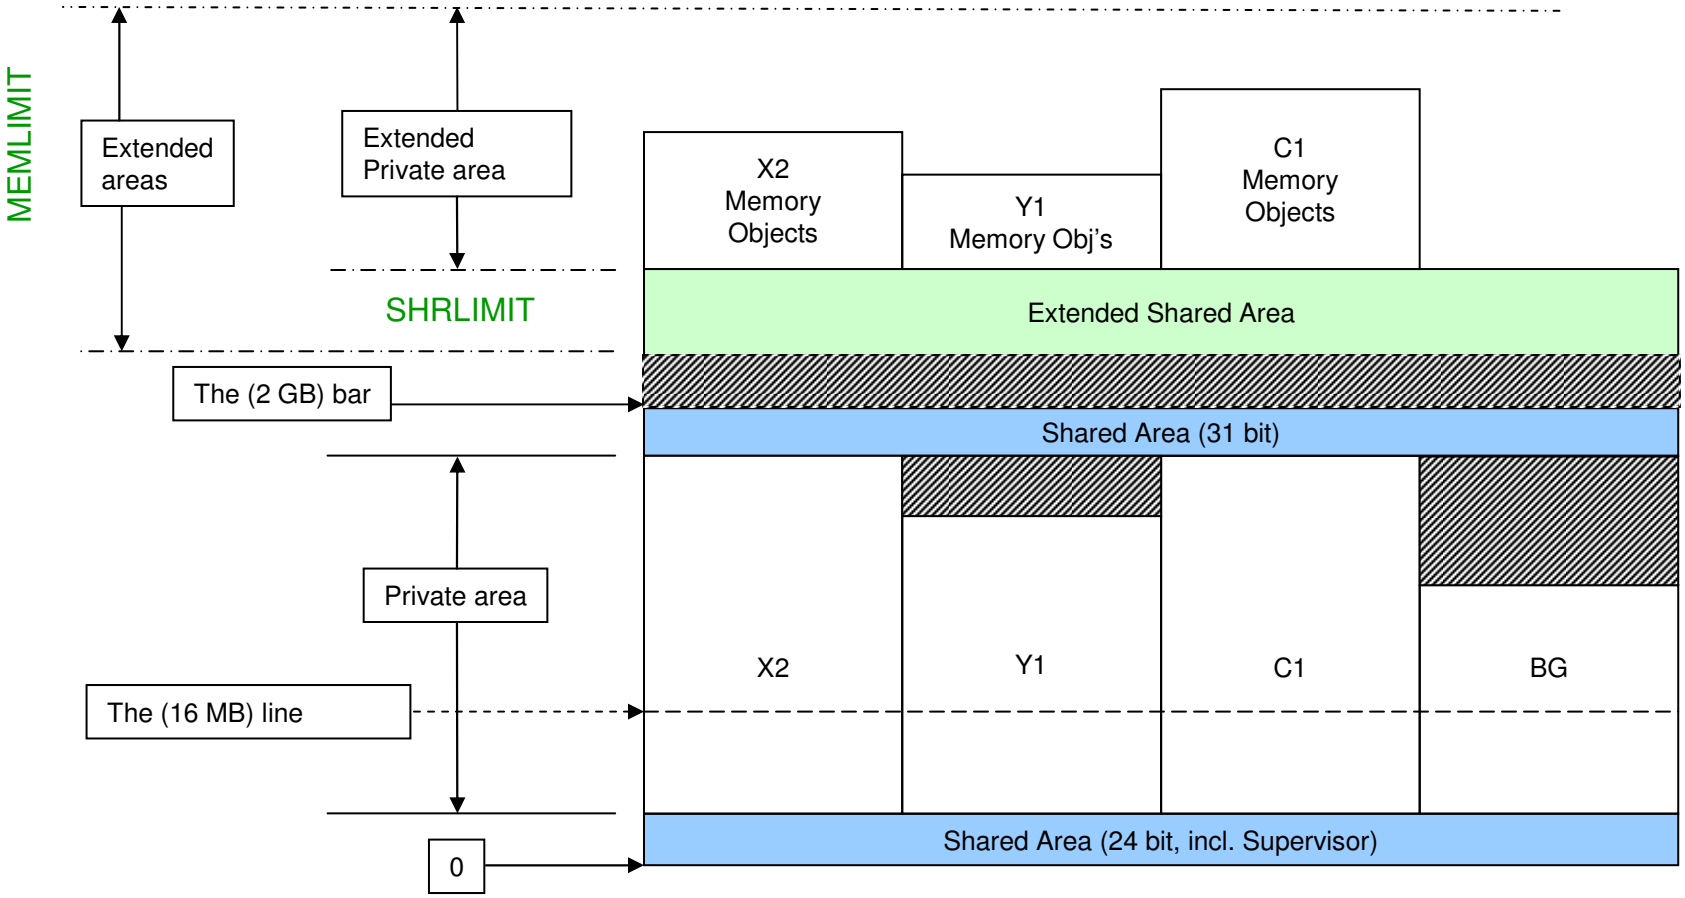

64 bit virtual I/O for applications

- Available with z/VSE 5.1 APAR DY47419
- SYSCOM bit IJBIO64E in IJBIOFL1, if 64 bit virtual I/O support available

- I/O buffers can now be created above the bar (above 2 GB)
- I/O buffers in **private memory objects** supported only
- I/O control blocks to be allocated below the bar (in 31 bit storage)

- Supported for ECKD devices

- CCB macro with a new parameter: IDAW=FORMAT2
- CCB points to a Format-0 or Format-1 CCW
- CCW with IDA-flag and data address point to a single Format-2 IDAW containing a 64 bit virtual address.

- I/O buffer will be TFIXed by I/O Supervisor, not necessary to PFIX the I/O buffer

- Not supported for
  - FBA / SCSI devices, tape devices, LIOCS

zEnterprise zEC12 / zBC12 zManager (HMC) and z/VSE

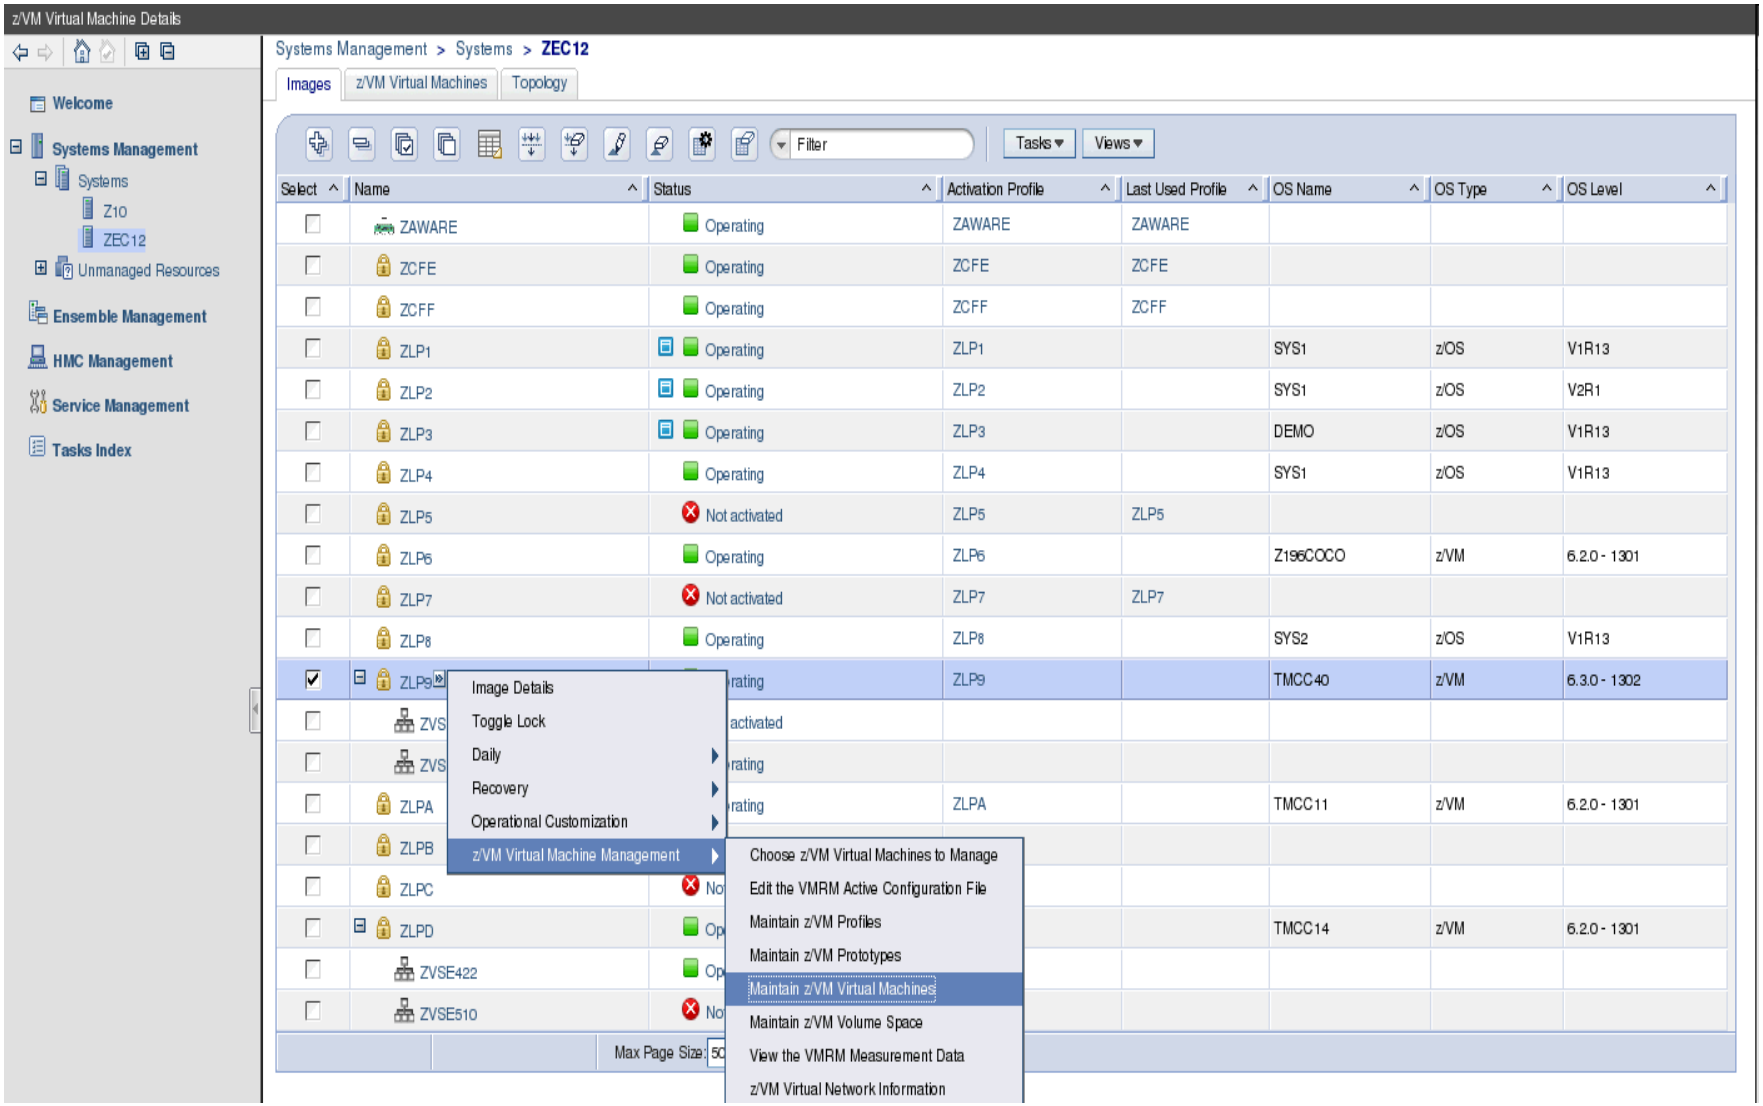

zEnterprise zManager zEC12 / zBC12 (HMC) and z/VSE

Virtual Machine Details for ZVSE510 - ZVSE510 [ZEC12:ZLP9:TMCC40]

General | Acceptable Status | **Configuration** | Virtual Network

Memory Size: 2 GB
Share Type: Relative
Share Value: 2000
Number of CPUs: 2

CPU Information

ID	Number	Status
FF11111128278000	0	Base
FF22222228278000	1	Stopped

Device Information

Address	Type
0191	DASD
019D	DASD
019E	DASD
9005	DASD
9006	DASD

Apply | Cancel | Help

Add a New z/VM Virtual Machine - ZEC12:ZLP9

Specify the virtual machine name and either select a prototype to add it with or specify the directory statements below.
If you select a prototype, any directory statement data will be ignored.
If you specify directory statements, then the initial password and initial account number may not be specified.
Click "OK" to add the virtual machine.
Click "View Prototype..." to view the selected prototype.
Click "View Profile..." to view a profile.

Virtual Machine Name: test
Virtual Machine Prototype Name: Select a prototype below
Initial Password:
Initial Password (verify):
Initial Account Number:

Select a prototype below
CMS
LINUX
ZVSE

Directory Statements:

```
USER ZVSETEST
```

View an Existing z/VM Virtual Machine - ZEC12:ZLP9

The virtual machine's directory statements are displayed below.
Click "View Profile" to view a profile.

Virtual Machine Name: ZVSE510

Directory Statements:

```
USER ZVSE510 I        3G 4G G
  ACCOUNT 3300
  CPU 0 CPUID 111111
  CPU 1 CPUID 222222
  IPL CMS PARM AUTOCR
  LOGONBY PI
  MACH ESA 2
  OPTION MAINTCCW CPUID 222222 QUICKDSP MAXCONN 25 TODENABLE
  POSIXINFO UID
*-------------------------------------------
* VSE MACHINE
*-------------------------------------------
```

Change an Existing z/VM Virtual Machine - ZEC12:ZLP9

Change the existing virtual machine's directory statements below.
Click "OK" to change the virtual machine.
Click "View Profile..." to view a profile.

Virtual Machine Name: ZVSE510

Directory Statements:

```
USER ZVSE510 I        ) 3G 4G G
  ACCOUNT 3300
  CPU 0 CPUID 111111
  CPU 1 CPUID 222222
  IPL CMS PARM AUTOCR
  LOGONBY
  MACH ESA 2
  OPTION MAINTCCW CPUID 222222 QUICKDSP MAXCONN 25 TODENABLE
  POSIXINFO UID
*-------------------------------------------
* VSE MACHINE change comment for test!!
*-------------------------------------------
```

System z Exploitation

- FICON Express8 - Higher I/O bandwidth
- Adapter interruptions (performance improvements)
  - OSA-Express3 / OSA-Express4S / OSA-Express5S (QDIO mode)
  - FICON Express8 / FICON Express8S (FCP)
- OSA-Express features
  - 10 Gigabit Ethernet, Gigabit Ethernet
  - 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 queued-I/O assist for real networking devices
  - OSA-Express adapters (CHIPID type OSD)
  - Hipersockets (CHIPID type IQD)

OSA-Express Support

- OSA-Express for high-speed communication
  - OSA-Express3 on z10, z196, z114, zEC12, zBC12
  - OSA-Express4S on z114, z196 and zEC12, zBC12
  - OSA-Express5S on zEC12, zBC12
- OSA-Express for non-QDIO environments (CHPID type OSE)
  - SNA and passthru traffic require configuration via OSA/SF
    - OSA-Express4S / OSA-Express5S on HMC
- 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
- Port specification for TCP/IP
  - OSA-Express 10 GbE features: one port per CHPID to connect to the network
  - OSA-Express 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

OSA-Express Support on zEC12 / zBC12 HMC

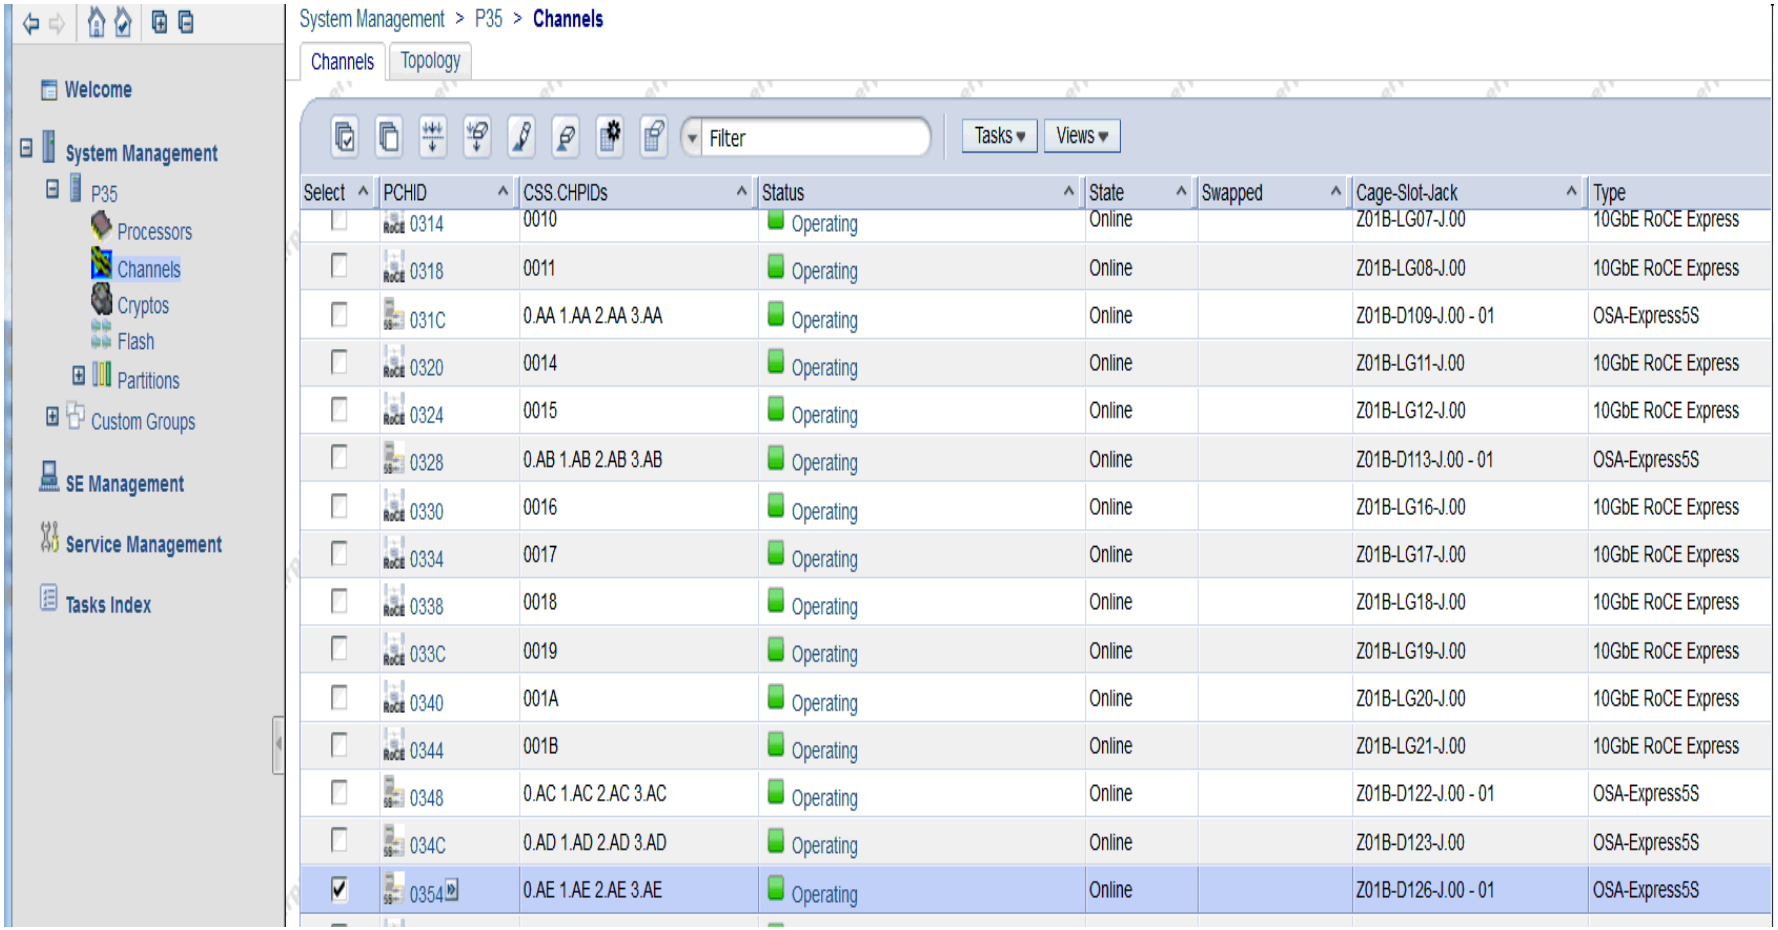

Select	PCHID	CSS.CHPIDs	Status	State	Swapped	Cage-Slot-Jack	Type
	0314	0010	Operating	Online		Z01B-LG07-J.00	10GbE RoCE Express
	0318	0011	Operating	Online		Z01B-LG08-J.00	10GbE RoCE Express
	031C	0.AA 1.AA 2.AA 3.AA	Operating	Online		Z01B-D109-J.00 - 01	OSA-Express5S
	0320	0014	Operating	Online		Z01B-LG11-J.00	10GbE RoCE Express
	0324	0015	Operating	Online		Z01B-LG12-J.00	10GbE RoCE Express
	0328	0.AB 1.AB 2.AB 3.AB	Operating	Online		Z01B-D113-J.00 - 01	OSA-Express5S
	0330	0016	Operating	Online		Z01B-LG16-J.00	10GbE RoCE Express
	0334	0017	Operating	Online		Z01B-LG17-J.00	10GbE RoCE Express
	0338	0018	Operating	Online		Z01B-LG18-J.00	10GbE RoCE Express
	033C	0019	Operating	Online		Z01B-LG19-J.00	10GbE RoCE Express
	0340	001A	Operating	Online		Z01B-LG20-J.00	10GbE RoCE Express
	0344	001B	Operating	Online		Z01B-LG21-J.00	10GbE RoCE Express
	0348	0.AC 1.AC 2.AC 3.AC	Operating	Online		Z01B-D122-J.00 - 01	OSA-Express5S
	034C	0.AD 1.AD 2.AD 3.AD	Operating	Online		Z01B-D123-J.00	OSA-Express5S
☑	0354	0.AE 1.AE 2.AE 3.AE	Operating	Online		Z01B-D126-J.00 - 01	OSA-Express5S

OSA/SF Support on zEC12 / zBC12 HMC

Advanced Facilities - PCHID0354

Channel ID: 0354
Channel type: OSE for non-QDIO
Card description: OSA-Express5S 1000BASE-T Ethernet
Select a function and click "OK".

- View code level
- Card trace/log/dump facilities...
- Card specific advanced facilities...
- Reset to defaults...

OK Cancel

Advanced Facilities - PCHID0354

Channel ID: 0354
LAN port type: 1000Base-T Ethernet
Select a function and click "OK".

- Query port status...
- View port parameters...
- Display or alter MAC address...
- Enable or disable ports...
- Run port diagnostics
- Set card mode...
- Panel configuration options...
- Manual configuration options...
- Activate configuration
- Display activate configuration errors...
- Display OAT entries...

OK Cancel

Edit OSA Address Table (OAT) Entries - PCHID0354

Channel ID:0354 LAN port type:OSE

--- Select Action --- Filter

Select	CSS	IID	Unit Address	Device Number	LPAR Name	Port Number	Session Type	IP Address	Router Indicator
◉	00	01	00, 01	0580, 0581	R35LP01	0	TCPIP	NONE	NONE
○	00	01	02, 03	0582, 0583	R35LP01	1	TCPIP	NONE	NONE
○	00	02	00, 01	0580, 0581	R35LP02	0	TCPIP	NONE	NONE
○	00	02	02, 03	0582, 0583	R35LP02	1	TCPIP	NONE	NONE
○	00	03	00, 01	0580, 0581	R35LP03	0	TCPIP	NONE	NONE
○	00	03	02, 03	0582, 0583	R35LP03	1	TCPIP	NONE	NONE
○	00	04	00, 01	0580, 0581	R35LP04	0	TCPIP	NONE	NONE
○	00	04	02, 03	0582, 0583	R35LP04	1	TCPIP	NONE	NONE
○	00	05	00, 01	0580, 0581	R35LP05	0	TCPIP	NONE	NONE
○	00	05	02, 03	0582, 0583	R35LP05	1	TCPIP	NONE	NONE
○	00	06	00, 01	0580, 0581	R35LP06	0	TCPIP	NONE	NONE
○	00	06	02, 03	0582, 0583	R35LP06	1	TCPIP	NONE	NONE

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 V5
- Virtual HiperSockets via z/VM Guest LAN support
- HiperSockets Completion Queue (z/VSE V5)

HiperSockets configurable input buffers

- Available as APAR DY47394 (z/VSE 5.1), included in z/VSE 5.2
- QDIO input queue buffers were set to 8 before
- More QDIO input buffers can improve performance
- In z/VSE you may increase the number of buffers to up to 64
- With a new configuration option you may select 8 (default), 16, 32 or 64 in the configuration file (IJBOCONF.PHASE)
- QDIO input buffers are allocated in 31 bit partition GETVIS space
- The buffers are to be PFIXed.
  - The limit for PFIX storage has to be defined with the JCL SETPFIX command
- QDIO input buffers are available for HiperSockets and OSA Express (CHPID OSD)

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 transparently
  - Encryption Facility for z/VSE (CPACF)
- Transparent for “TCP/IP” applications
  - VSE connector server, CICS Web Support, 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 / Express3 / Express4S – for asymmetric encryption
  - Encryption hardware assist for increased SSL throughput
    - Supports SSL handshaking only for applications that use the SSL crypto API
  - Crypto Express4S support (z/VSE 5.1 + PTF)
  - 2048-bit RSA key with Crypto Express2
  - 4096-bit RSA key support with configurable Crypto Express3 / Crypto Express 4S
  - Configurable Crypto Express
    - 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

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

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

Exploitation of IBM System Storage Products

- IBM System Storage TS1130 / TS1120 / TS1140 Tape Drive
- IBM System Storage TS7700 / TS7720 Virtualization Engine
  - Copy Export function of TS7700 Virtualization Engine for disaster recovery
  - Multi-Cluster Grid support of the TS7700 Virtualization Engine Series
- IBM System Storage TS3400 autoloader Tape Library
- IBM System Storage TS3500 Tape Library

- zVSE supports the channel command interface via
  - Perform Subsystem Function (PSF)
  - Perform Library Function (PLF) commands

Tape Data Encryption

- IBM TS1120 / TS1130 / TS1140 Tape Drive with encryption feature
  - Supports data encryption within the drive itself
  - Using Systems Managed Encryption with the TS1120 / TS1130 / TS1140
  - 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 cardridge.
  - 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 wirte 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

Exploitation of IBM System Storage Products …

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

- 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
- Upgrade of the z/VSE support for the Parallel Access Volume (PAV) feature (ECKD)

FlashCopy Support

- Available on DS8000, DS6000 and ESS series
- 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
- Supported by ICKDSF only
  - DS8000 Remote Mirror and Copy (RMC)
  - Peer-to Peer Remote Copy (PPRC)
    - Allows remote data replication

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

SCSI Support in z/VSE

- SCSI disks as emulated FBA disks on z/VM
  - 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
- FCP-attached SCSI disk support (IBM System Storage)
  - DS8000, DS6000 and ESS series
  - 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
  - IBM Storwize V5000 Midrange Disk
  - IBM Storwize V3700 Entry Disk

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

- 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) not 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 Support – Disk Controller Configuration

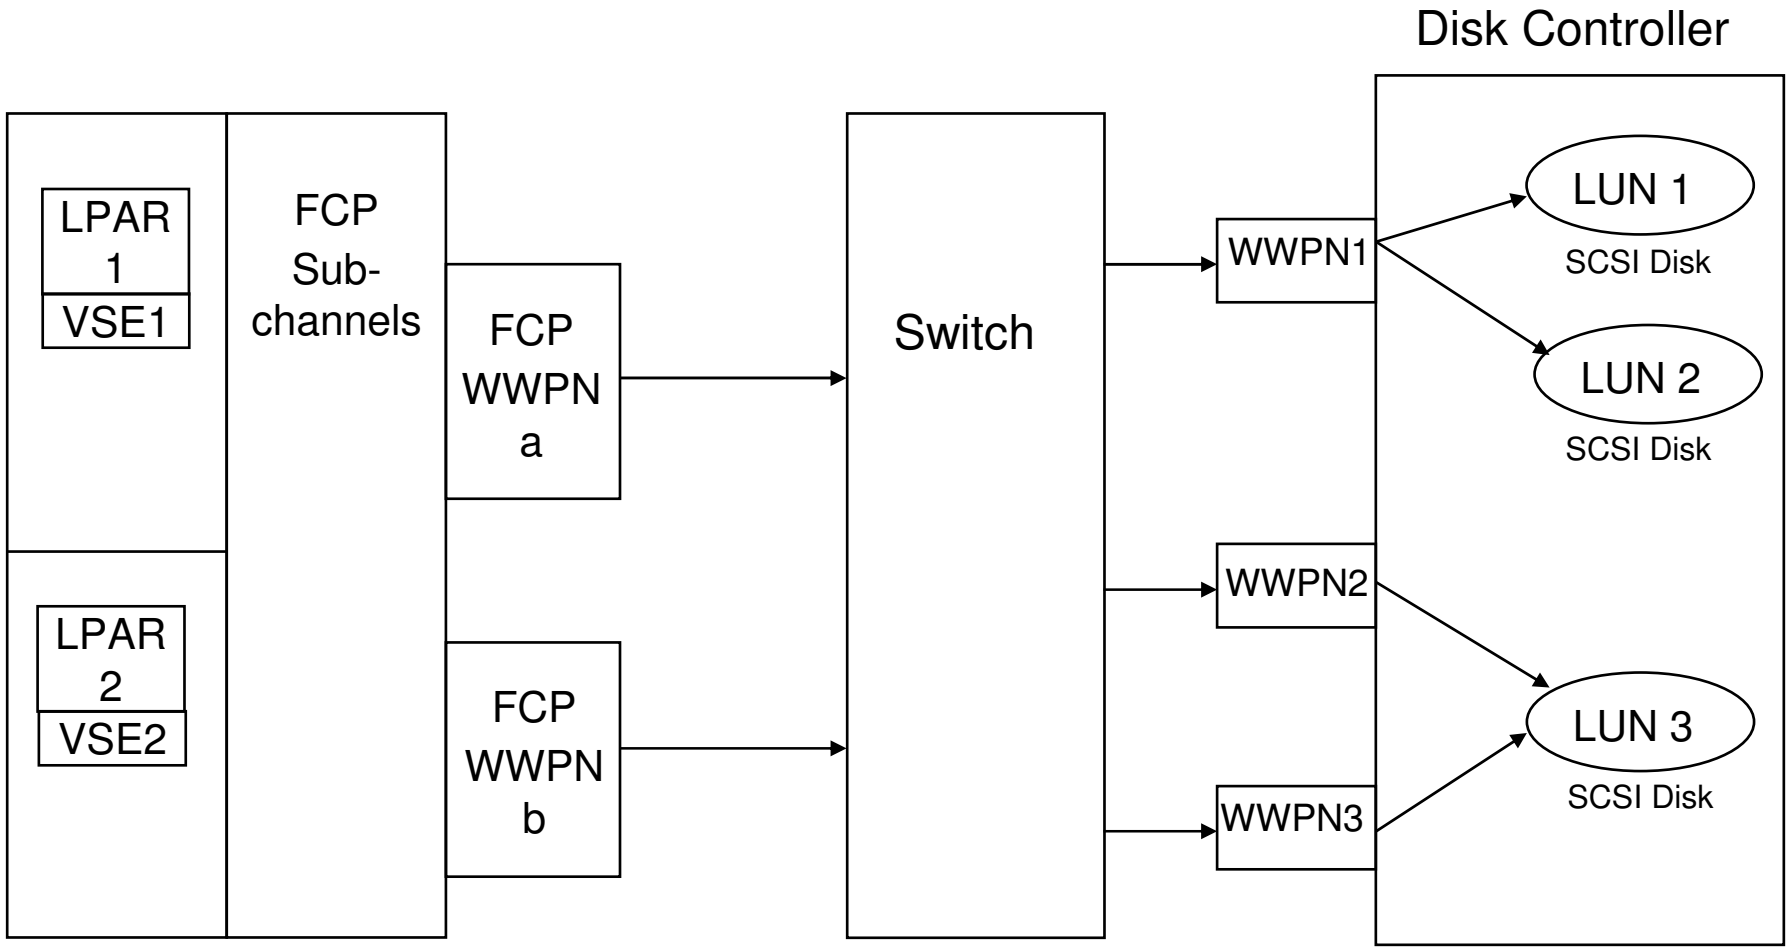

Point to point connection possible (z9 or higher possible)

More Information

- … on VSE home page: http://ibm.com/vse
- Ingolf's z/VSE blog: https://www.ibm.com/developerworks/mydeveloperworks/blogs/vse

- Hints and Tips for z/VSE 5.1:
  - http://www.ibm.com/systems/z/os/zvse/documentation/#hints

- 64 bit virtual information:
  - IBM z/VSE Extended Addressability, Version 5 Release 1
  - IBM z/VSE System Macro Reference, Version 5 Release 1

- CICS Explorer: http://www.ibm.com/software/htp/cics/explorer/

- IBM Redbooks:
  - Introduction to the New Mainframe: z/VSE Basics
    http://www.redbooks.ibm.com/abstracts/sg247436.html?Open
  - Security on IBM z/VSE – updated
  - http://www.redbooks.ibm.com/Redbooks.nsf/RedbookAbstracts/sg247691.html?Open
  - z/VSE Using DB2 on Linux for System z
  - http://www.redbooks.ibm.com/abstracts/sg247690.html?Open
  - New: Enhanced Networking on IBM z/VSE
    http://www.redbooks.ibm.com/Redbooks.nsf/RedpieceAbstracts/sg248091.html?Open

- Please contact z/VSE: https://www-03.ibm.com/systems/z/os/zvse/contact/contact.html
  or me – Ingolf Salm – salm@de.ibm.com – for any questions

Questions ?