# z/VSE SCSI Support and Migration Options

Monika Zimmermann
Email: Monika.Zimmermann@de.ibm.com


z/VSE 5.2

http://www.ibm.com/zVSE
http://twitter.com/IBMzVSE

**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 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.

# Why SCSI?

- SCSI (Small Computer System Interface) disks are widely used in ‘open‘ systems
- SCSI controllers supporting the FCP (Fibre Channel Protocol) can be attached to IBM System z
- z/VSE SCSI support was introduced with z/VSE 3.1 (GA March 2005)
  - z/VSE manages the relationship with SCSI hardware
- Offers more storage choices
- Hybrid systems are part of z/VSE‘s ‘PIE‘ strategy
  - z/VSE SCSI support fits well in hybrid environments
  - SCSI controllers can be shared by z/VSE, Linux, or z/VM
- Disk controllers qualified for use with z/VSE SCSI:
  - IBM System Storage DS8000 Series
  - IBM Storwize V7000, V5000, V3700, V3500
  - IBM SAN Volume Controller (SVC)
  - IBM XIV Storage System
  - DS6000 (EOM), some models of IBM Total Storage ESS (EOM)

# What is supported

- SCSI disks can be used as system and/or data disks
- z/VSE supports SCSI only systems (DOSRES and SYSWK1 on SCSI)
  - Installation on SCSI
- z/VSE can use SCSI disks in an LPAR
  - z/VSE‘s SCSI support is used to access FCP-attached SCSI disks
- z/VSE can use SCSI disks in a z/VM guest environment
  - Using z/VSE‘s SCSI support or
  - Using z/VM‘s SCSI support via emulated FBA disks
    - For z/VSE these are real FBA disks
    - Maximum size is 2GB (real FBA size supported by z/VSE)
- High availability through Multi-Pathing (failover)
- DASD sharing
- N_Port ID Virtualization (NPIV)

# Integrating SCSI in z/VSE

- z/VSE‘s SCSI implementation is transparent to application programs
- SCSI disks are seen in z/VSE as FBA disks
  - Both SCSI and FBA (Fixed Block Architecture) disks have an underlying (512-bytes) block structure
  - Few configuration commands are needed to define and work with SCSI disks
  - Once configured, z/VSE application and system programs see SCSI disks as FBA disks
  - SCSI (FBA) disks are accessed by system and user programs using FBA interfaces
  - FBA I/O channel commands are internally translated into SCSI commands
- User applications, vendor and system programs will run unchanged
  - Provided they are device-independent or use FBA channel programs
- SCSI disk size when using z/VSE‘s SCSI support
  - 8MB up to 24GB, VSE/VSAM can use first 16GB of a SCSI disk
- z/VSE applications can not use SCSI commands

# SCSI Configuration

- Hardware
  - IBM System z server
    - FCP adapter (FICON Express card configured as CHPID type FCP)
    - FCP adapter passes SCSI commands to the SCSI disk
    - FCP adapter is connected to the switch or controller
  - FCP-capable switch (switched network)
    - Offers greatest flexibility and reliability
    - Required by SCSI only disk systems (Storwize, SVC, XIV)
  - Alternative: Point-to-Point Connection (no switch)
    - FCP adapter is directly attached to the disk controller
    - Disadvantage: Less flexibility
  - Qualified disk controller - define your SCSI disks
    - Define the LUN (Logical Unit Number), gives e.g. LUN= 4020402600000000
- Software
  - z/VSE running in an LPAR or as z/VM guest

# System z Configuration

- FICON Express adapter configured as CHPID type FCP in the IOCDS
- In the IOCDS for each CHPID the IODEVICEs are configured
  - Example: CHPID 9D, IODEVICE FA0,FA1, ...
  - IODEVICE is the FCP device used by z/VSE to access the FCP adapter
- Associated with the FCP CHPID is a unique physical WWPN (World Wide Port Name)
  - same for all IODEVICEs on the CHPID
- WWPN is used within the Storage Area Network (SAN) to grant access to SCSI disks
  - when the physical WWPN is used: all IODEVICEs are allowed to access a SCSI disk
- N_Port ID Virtualization (NPIV) allows to define virtual WWPNs
- Recommendation: Always use NPIV
  - Each IODEVICE has its own virtual WWPN – grant access on a IODEVICE base
  - Provides better access control - prevents unauthorized access of LUNs (SCSI disk)
  - Requires System z9 or later

# Steps to configure SCSI in z/VSE

- All I/O devices used by z/VSE must be ADDed during IPL
  - ADD FA0,FCP Add the FCP device (FA0) defined in the IOCDS.
    It carries the SCSI commands
  - ADD FB0,FCP Add the FCP device (FB0) as defined in the IOCDS.
    Needed for multi-pathing only
  - ADD 700:702,FBA Add (virtual) FBA devices under which the SCSI devices are known to z/VSE.
    FBA device (cuu) must not be defined in the IOCDS or z/VM guest.
- Connect the FBA device with the SCSI device – define the connection path
  - (IPL) DEF SCSI command required for system disks (DOSRES,SYSWK1,DLF,DPD..)
    - DEF SCSI,FBA=700,FCP=FA0,WWPN=500507630A08C066,LUN=4020406700000000
    - DEF SCSI,FBA=701,FCP=FA0,WWPN=500507630A08C066,LUN=4020406800000000
  - (JCL/AR) SYSDEF SCSI command can be used for data disks - recommended
    - SYSDEF SCSI, FBA=702,FCP=FA0,WWPN=500507630A08C066,LUN=4020406900000000
- Now the SCSI disks can be used

# IPL z/VSE from SCSI

- IPL the FCP device (not the SCSI device)
- Connection path must be defined prior to IPL

Initiating IPL of z/VSE (when running as z/VM guest)

- SET LOADDEV PORT 50050763 00C295A5 LUN 40204067 00000000
- IPL FA0 (IPL fcp_device_number)

Initiating IPL of z/VSE (when running in an LPAR)

In the Load Panel select SCSI and specify:

| | |
|---|---|
| World Wide Port Name | 5005076300C295A5 |
| Logical Unit Number | 4020406700000000 |
| LOAD Address | FA0 |

# Delete a SCSI device

- Delete a SCSI device
  - (AR) OFFLINE cuu command
    - OFFLINE 700 (FBA cuu)
    - Rejected if I/O is ongoing
  - SYSDEF SCSI,DELETE,FBA=
    - Either delete one connection path
    - Delete the SCSI device (all connection pathes)

# High-Availability through Multi-Pathing

- More than one connection path to the same SCSI device is defined
- Protects the z/VSE system against potential outages of
  - FCP adapter and/or disk controller port and thus enhances availability
- I/O against the SCSI device is repeated on the alternate path in case of failure
  - Failure not presented to the application program
- Failing connection path is automatically recovered by z/VSE
- Multi-Pathing is not used for workload balancing
- Example:
  - DEF SCSI,FBA=700,FCP=FA0,WWPN=500507630A08C066,LUN=40204067
  - DEF SCSI,FBA=700,FCP=FB0,WWPN=500507630A03C066,LUN=40204067

# z/VSE SCSI Configuration using a Disk Controller - Example

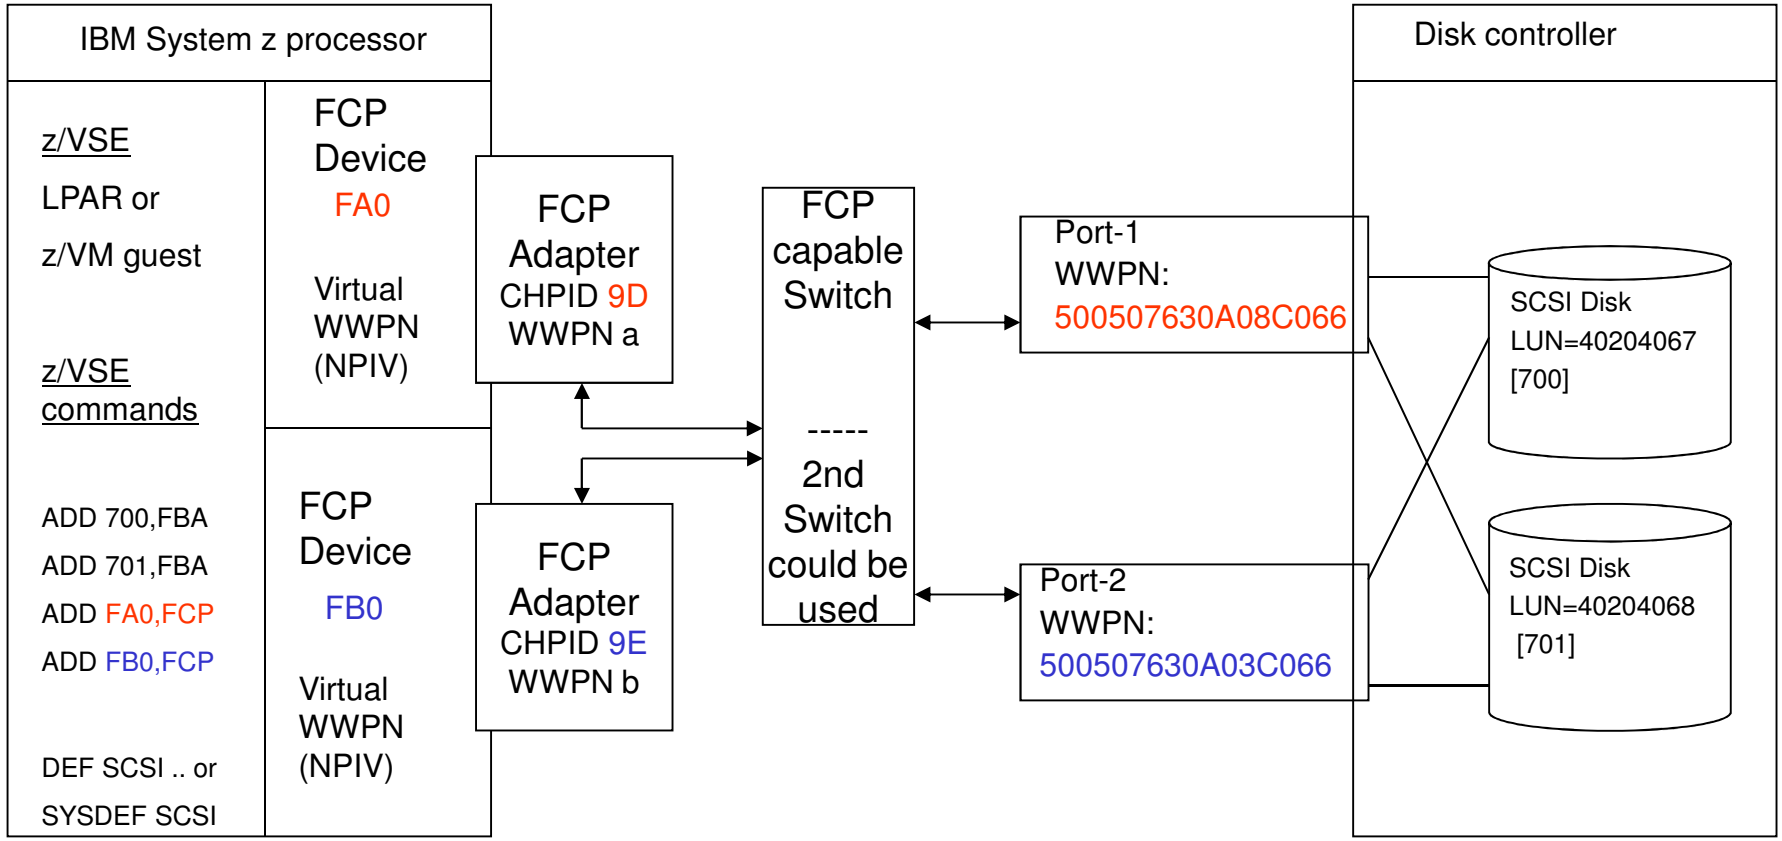


**DEF SCSI,FBA=700,FCP=FA0,WWPN=500507630A08C066,LUN=40204067**

**DEF SCSI,FBA=700,FCP=FB0,WWPN=500507630A03C066,LUN=40204067**

**DEF SCSI,FBA=701,FCP=FA0,WWPN=500507630A08C066,LUN=40204068**

**DEF SCSI,FBA=701,FCP=FB0,WWPN=500507630A03C066,LUN=40204068**

# Switch Configuration

- When SVC is used
  - Connect the FCP adapter, the SVC and the disk controller with the switch
  - FCP and SVC must be in the same zone
  - SVC and storage controller must be in the same zone
- When disk controller is used
  - Connect the FCP adapter and the disk controller with the switch
  - FCP and disk controller must be in the same zone

# Controller Configuration

- IBM Disk Controller
  - Define the LUN (Logical Unit Number), gives e.g. 40204067, 40204068
  - Configure which ports of the controller (port WWPNs) are allowed to access the LUNs
  - Configure which FCP adapters (FCP WWPNs) are allowed to access the LUNs
  - Use NPIV instead of physical WWPN of FCP adapter:
    - Each FCP IODEVICE has its own virtual WWPN – grant access on IODEVICE base
    - Provides better access control – prevents unauthorized access of LUNs (SCSI disks)
    - Access restricted to an FCP IODEVICE (and not the whole FCP CHPID)
- SVC (similar for Storwize)
  - Define disk storage in the disk controller, managed by the SVC
  - Use the SVC 'Create VDISKs' (to create a LUN)
  - Define the FCP adapter as host in the SVC (use the WWPN of the FCP adapter)
    - Since FCPs are not automatically detected use 'Additional Ports' to type in the WWPN
  - Use 'Map VDISKs to a Host' to assign the VDISK to the host (FCP adapter)
    - This gives the SCSI LUN IDs 1,2,3,...
    - These LUN IDs relate to LUN 0001, 0002, 0003 .. in z/VSE

The terminology for the configuration may change

# z/VSE SCSI Configuration using a SAN Volume Controller - Example

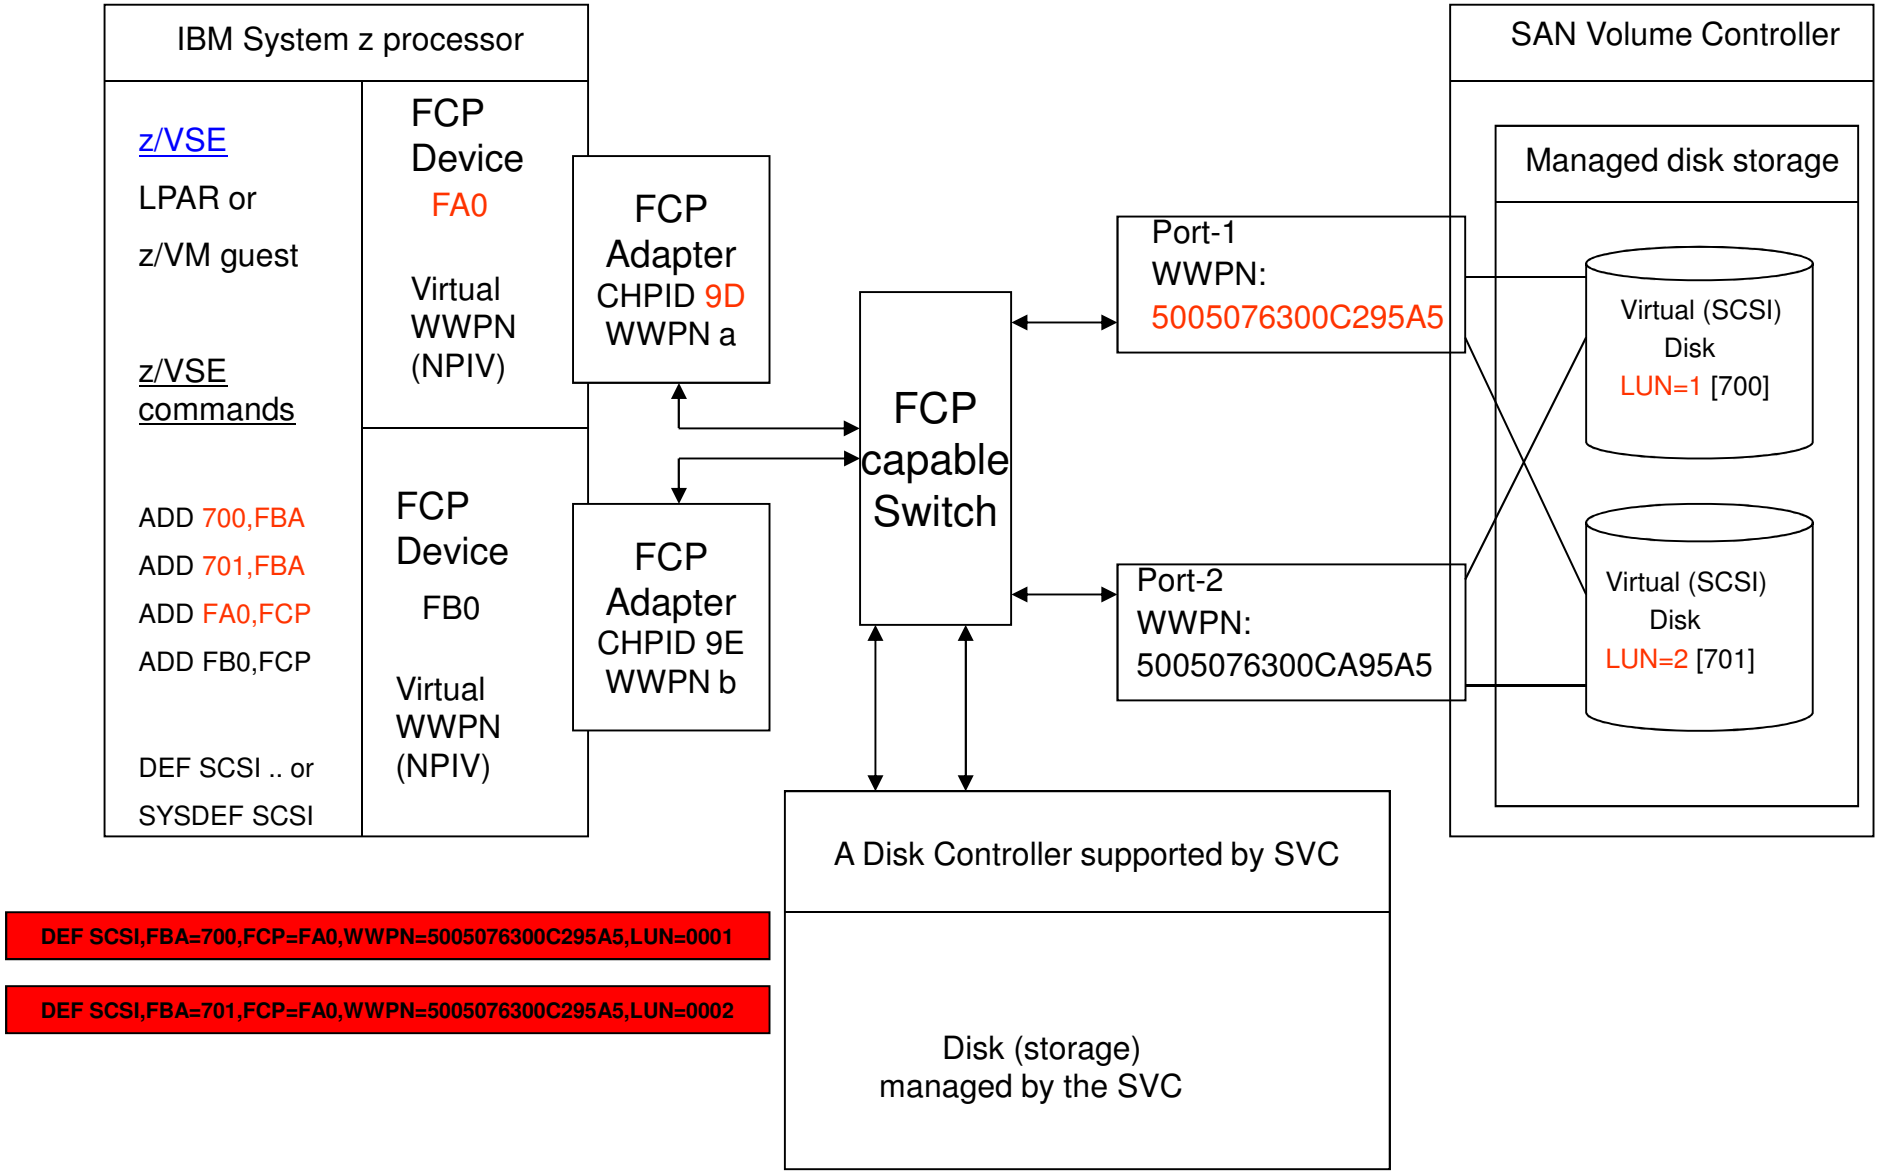

# General Recommendations for z/VSE systems

- Run your system with NOPDS
  - z/VSE supports up to 32GB processor storage
  - Exploit processor storage of IBM zEnterprise servers
  - Results in improved performance since no page I/O is done
- Shared Systems – Lock file on SCSI
  - Place the lock file (DLF statement) on a separate disk
  - Do not specifiy a multi-path connection to the lock file disk

# Useful Commands

## QUERY SCSI

```
AR 0015 FBA-CUU  FCP-CUU  WORLDWIDE PORTNAME  LOGICAL UNIT NUMBER  PORT-STATUS

AR 0015    700          FA0     5005076300C295A5          00010000000000000          -
AR 0015    700MP        FB0     5005076300CA95A5          00010000000000000          -

AR 0015    701          FA0     5005076300C295A5          00020000000000000          -
AR 0015    701MP        FB0     5005076300CA95A5          00020000000000000          -
```

## VOLUME 700

```
AR 0015 CUU  CODE DEV.-TYP   VOLID     USAGE    SHARED  STATUS  CAPACITY
AR 0015  700   90      1732-FCP   DOSRES   USED                              4194304 BLK
```

## STATUS (FCP) cuu

```
AR 0015 SCHIB DEV  INT-PARM ISC FLG LP PNO LPU PI MBI  PO PA CHPID0-3

AR 0015 000E  0FA0  00004650   3     83    80  00     80   80  0000 FF 80  9D000000

...

REQUEST IS STARTED           DEVICE IS BUSY
```

## STATUS (FBA) cuu

```
AR 0015 DEVICE IS AN FCP-CONTROLLED SCSI DEVICE

AR 0015   PUB=00004608  PUBX=000B0510  PUB2=000ADE82  POWN=00006520

AR 0015  VCTE=000AD9D0  POWNX=002B6DF4
```

# Migration from ECKD to SCSI

- Initial Installation
  - z/VSE system on SCSI requires initial installation
  - Fast Service Upgrade (FSU) from ECKD to SCSI is not supported
- Device Structure
  - (E)CKD devices are cylinder / track oriented
  - FBA and SCSI devices are block oriented
- IPL Procedure
  - At least one ADD statement for FCP device
  - Each SCSI disk requires an ADD statement for an FBA device
  - DEF SCSI for each system disk
  - SYSDEF SCSI for each data disk
- SYS BUFSIZE value – copy blocks
  - FBA channel programs might have different need of copy blocks compared to ECKD
- IPL/JCL Statements and VSE/VSAM definitions
  - File, catalog, and space definitions, extent statement
    - Replace cylinder / track specification by block specification
  - Calculation:
    - One (3390) track (cylinder) is about 112 (1680) blocks on SCSI
    - A library VSE library block is 1024 bytes, requires 2 blocks on SCSI

# Data Migration from ECKD to SCSI

- Use the utilities provided by the components (DB2, DL/1, VSAM,.. )
- Adapt your JCL definitions (e.g. replace cylinder/track by block specification)
- Adapt your VSAM definitions
  - replace cylinder/track by block / record specification
- VSE Libraries
  - Librarian Backup / Restore function
- Sequential (SAM) files
  - DITTO copy function
- Migration of VSAM data
  - VSAM LISTCAT function shows size of currently used clusters
  - VSAM Backup / Restore: defines clusters in target catalog and moves data
  - VSAM Export / Import: defines one cluster and moves data
  - VSAM Repro: only moves data, cluster has to be defined
- Migration of VSE Power Queues
  - POFFLOAD command

# Data Migration from ECKD to SCSI cont.

- Migration of DB2 data
  - Increase copy blocks prior to initialize / extend DB2 database on target system
    - SYS BUFSIZE =12000
  - Decrease afterwards
  - Define VSAM catalogs and clusters
  - DBSU UNLOAD / RELOAD
  - DB2 application programs: no changes required.

- Migration of DL/I data
  - Define VSAM catalog and clusters
  - New DBD generation with adapted DATASET statement
    - Specify FBA in the DEVICE parameter
    - Adapt BLOCK parameter to CISIZE depending on KSDS or ESDS HD database
  - CISIZE BLOCK parameter
  - New ACB generation (DLZUACB0 utility) with DMB=YES on BUILD control statement
  - DL/I Image Copy and DL/I Recovery utilities
  - DL/I Unload and DL/I Reload utilites
  - DL/I application programs: no changes required

# Applications

- Vendor programs
  - Contact your vendor if updates are required when using FBA / SCSI disks
- Device independent programs
  - No change required
- Device dependent programs
  - Change channel programs and/or language specifications
  - Use FBA interfaces instead of (E)CKD interfaces
- COBOL – RPG – PL/1 Programs
  - COBOL/VSE fully compliant with ANSI-85 COBOL, compiled with COBOL/VSE should run unchanged
  - 'Older' COBOL programs with device specific language elements might need migration
  - RPG programs with device specific language elements might need migration
  - PL/1 programs using regional datasets need to be migrated
  - Details can be found in language documentation.

# What is not supported

- Standalone dump on SCSI disk
- z/VSE Flashcopy for SCSI disks
  - Copy services offered by the disk controller can be used
- Attachment of non-disk SCSI devices (for example SCSI tapes)
- Concurrent microcode upgrade for FCP-attached SCSI disks
- z/VSE installation disk on SCSI (for tapeless initial installation)

# Performance

- Test jobs with heavy I/O load showed improved elapsed times compared with ECKD disks
- CPU utilization may increase for jobs using SCSI disks:
  - FBA to SCSI command translation
- Check your CPU utilization before migrating to SCSI
- Results may vary depending on environment

# Summary

- Transparent to applications
- Easy to setup
- Offers more storage choices
- Fits well in hybrid environments
- May reduce elapsed times of your jobs

# Documentation

- z/VSE Administration publication
- z/VSE Planning publication
- z/VSE SCSI Support and Migration Options (Whitepaper)

# Thank You


Please forward your questions or remarks to
zvse@de.ibm.com
Monika.Zimmermann@de.ibm.com

# z/VSE Live Virtual Classes

z/VSE @ http://www.ibm.com/zvse/education/

LINUX + z/VM + z/VSE @ http://www.vm.ibm.com/education/lvc/

Read about upcoming LVCs on @ http://twitter.com/IBMzVSE

Join the LVC distribution list by sending a short mail to alina.glodowski@de.ibm.com