

z/VSE SCSI Support and Migration Options

Monika Zimmermann
z/VSE Development
IBM Labor Böblingen, Germany

Table of Contents

Terminology.....	4
Trademarks.....	5
z/VSE SCSI Support.....	6
How SCSI disks are integrated in z/VSE.....	7
Internals.....	8
SCSI Configuration.....	9
Hardware Configuration.....	9
Steps to Configure SCSI in z/VSE.....	10
High-Availability Through Multi-Pathing.....	11
Point-to-Point Connections.....	12
N-Port ID Virtualization (NPIV).....	12
IPL from SCSI.....	13
Initiating an IPL of z/VSE in an LPAR.....	14
Initiating an IPL of z/VSE in a z/VM Guest.....	14
z/VSE SCSI Support - Summary.....	14
What the z/VSE SCSI provides.....	14
What the z/VSE SCSI not provides.....	15
Concurrent micro code upgrade.....	15
z/VSE running as z/VM guest.....	16
Using emulated FBA disks.....	16
Migration of z/VSE from ECKD to SCSI.....	17
Device structure.....	17
Initial Installation Required.....	17
Configuration Changes.....	17
IPL Procedure.....	17
Copy Block Consideration - IPL statement SYS BUFSIZE.....	17
JCL/ IPL Statements and VSE/VSAM definitions.....	17
How to calculate the number of required blocks.....	18
Data Migration.....	18
Migration of z/VSE Libraries.....	18
Migration of Sequential (SAM) files.....	18
Migration of VSAM data.....	18
Migration of DB2 data.....	19
Migration of DL/I data.....	20
Migration of VSE POWER Queues.....	21
Migration of the DTSSFILE (ICCF library).....	21
Vendor programs.....	22
Device independent applications.....	22
Device dependent applications.....	22
COBOL – PL/1 – RPG programs.....	22
Device Characteristics.....	23

z/VSE SCSI Support

Performance and CPU Usage.....	23
Conclusion.....	23

Terminology

The source system is the z/VSE system on ECKD.

The target system is the z/VSE system on SCSI.

A SCSI disk is an FCP-attached SCSI disk.

World Wide Port Name (WWPN)

Logical Unit Number (LUN)

LUN means a SCSI disk.

Trademarks

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

IBM System z z/VSE z/VM DB2 DS8000 DS6000 XIV Storwize

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

Other company, product, or service names, may be the trademarks or service marks of others.

z/VSE SCSI Support

Small Computer System Interface (SCSI) disks are widely used in “open” systems. SCSI disks that support the Fibre Channel Protocol (FCP) can be attached to IBM System z Servers, as well.

Starting with z/VSE 3.1 (GA 2005), z/VSE provides native SCSI support for FCP-attached SCSI disks. The z/VSE SCSI support manages the relationship with the SCSI hardware.

In this document, the term SCSI disk always means an FCP-attached SCSI disk.

z/VSE 3.1 SCSI support was qualified for use with IBM System Storage DS8000 series, DS6000 series and certain ESS models.

Starting with z/VSE 4.2, IBM XIV Storage System, IBM San Volume Controller, and IBM Storwize V7000 followed.

These are the disk controllers that are qualified for use with z/VSE SCSI:

- IBM System Storage DS8000 Series
- IBM XIV Storage System
- IBM Storwize V7000, V5000, V3700, V3500
 - z/VSE supports FCP-attached SCSI disks in members of the Storwize family, provided they can be attached through Fibre Channel (FC) switches to z/VSE.
 - z/VSE does not support Storwize file storage.
- IBM San Volume Controller (SVC)
- IBM System Storage DS6000 Series
- IBM TotalStorage Enterprise Storage Server Model F20, 800, and 800 Turbo

DS8000 series, DS6000 series, and ESS disk storage systems offer to define both (E)CKD and SCSI devices.

XIV, Storwize, and SVC are SCSI only storage systems.

This document summarizes the z/VSE SCSI support. Details can be found in the manuals

- z/VSE Planning
- z/VSE Installation
- z/VSE Administration

The migration chapter in this document summarizes the migration options from z/VSE on (E)CKD to z/VSE on SCSI. For data migration please see also the appropriate product documentation.

Supported z/VSE Releases:

The z/VSE SCSI support described in this document is available with any supported z/VSE release.

IBM System z Processor Support

If z/VSE is running on a System z9 or later, all features mentioned in this document are available.

IBM FICON Express feature

The term FICON Express feature is used for the family of FICON Express features that is FICON Express, FICON Express2, and higher.

A FICON Express feature can have one or more ports to connect to the Storage Area Network (SAN). Each port corresponds to a CHPID (Channel Path Identifier) in the input/output configuration data set (IOCDS).

The ports on a FICON Express feature can be configured with different CHPID types.

If a port is to be used to access FCP-attached SCSI disks it has to be configured as CHPID type FCP. For a port configured as CHPID type FCP, we use the term FCP adapter.

A CHPID can be shared among LPARs.

The I/O devices to be used with a CHPID are defined in the IOCDS with the IODEVICE statement. This I/O device is also referred as subchannel.

An I/O device on a CHPID FCP corresponds to a cuu of type FCP in z/VSE.

How SCSI disks are integrated in z/VSE

SCSI disks use fixed block 512 bytes sectors and therefore SCSI I/O commands are block oriented.

The same is true for Fixed Block Architecture (FBA) disks. FBA disks are organized in fixed block sectors of 512 bytes. A record may consist of multiple blocks and is addressed by its relative block number. I/O programming is done using FBA channel programs.

Since both SCSI and FBA devices have an underlying block structure z/VSE SCSI was implemented using the existing z/VSE FBA support.

Few configuration commands are needed to define and work with SCSI disks. Once configured, SCSI disks are seen to the operator, administrator and programmer as FBA disks and can be used with existing FBA interfaces.

FBA I/O channel commands are translated internally into SCSI I/O commands. This is done at low level I/O interfaces and therefore transparent to user and system and vendor programs.

User, system, and vendor programs will run unchanged provided they are device-independent or use FBA channel programs.

Internals

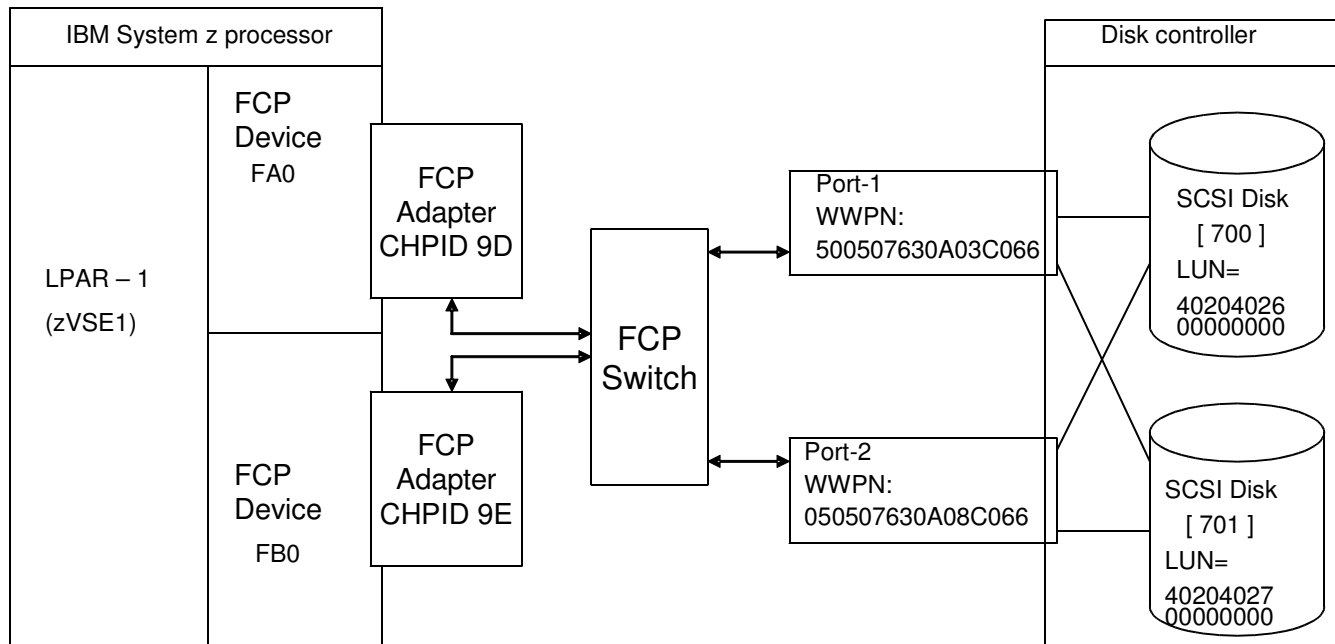
The existing FBA support made it relatively easy to provide z/VSE support for SCSI. Transparency for user, vendor, and z/VSE system programs is ensured by the z/VSE implementation on low-level I/O interfaces and the z/VSE SCSI device driver. Whenever I/O occurs against an FBA (SCSI) disk, the SCSI device driver intercepts the I/O, translates the FBA channel program (FBA I/O command) into a Fibre Channel SCSI I/O command and passes it to the FCP adapter, which then transfers the SCSI command to the SCSI disk. The FCP adapter presents the interrupt with the completion of the SCSI command to z/VSE. The SCSI device driver intercepts the interrupt and translates it into an FBA I/O interrupt.

SCSI Configuration

A typical SCSI environment consists of an FCP adapter, an FCP-capable switch and a disk controller. This switched network offers the greatest flexibility and reliability.

The following picture shows an example of a switched network. This picture is used to describe the hardware configuration and the steps needed to access a SCSI disk in z/VSE.

Switched Network:



First, you must configure the hardware. This includes the IOCDs, the switch, and the disk controller. Then you must configure software.

The hardware and software configuration in the next paragraphs is described based on the above picture.

Hardware Configuration

- You need a FICON Express adapter configured in the IOCDs as CHPID (channel path id) type FCP, that is an FCP adapter (example: CHPID 9D). This FCP adapter is connected to the switch. If you want to implement Multi-pathing you need a second FCP adapter (example: CHPID 9E). The FCP adapter 9D has the FCP device FA0, the FCP adapter 9E the FCP device FB0 configured in the IOCDs, using the IODEVICE statement.
- You need a switch that is FCP-capable.

- You need a qualified disk controller.
The disk controller in the example has two ports; both are connected to the switch. Port 1 has the WWPN (World Wide Port Name) 500507630A03C066 and port 2 has the WWPN 500507630A08C066. You need the second port if you want to implement Multi-pathing.
Use the disk controller's dialogs to configure SCSI disks and grant access. Two SCSI disks are defined with the Logical Unit Number (LUN) 4020402600000000 and 4020402700000000. A LUN represents a SCSI disk. The LUNs have been configured so that they can be accessed via Port 1 and Port 2 of the disk controller and both FCP adapters (WWPN of the FCP adapter).
The representation of a LUN depends on the disk controller. Above is the representation of a DS8000.
The cuu addresses 700 and 701 show the FBA addresses used in z/VSE to map the SCSI disk.

Now that the hardware is configured, you need to configure SCSI in z/VSE.

Steps to Configure SCSI in z/VSE

- You need an ADD statement for the FBA device under which the SCSI disk is known in z/VSE.
- You need an ADD statement for the FCP device that carries the SCSI commands.
- Then you can either use the IPL DEF SCSI command or the AR/JCL SYSDEF SCSI command to connect the FBA device with the SCSI device. In case the SCSI device is not needed during IPL, it is recommended to use the SYSDEF SCSI command to establish the connection path. This gives you more flexibility than using the IPL command.

ADD FA0,FCP ADD FB0,FCP	Add the FCP devices FA0 Add the FCP devices FB0. This is only needed for Multi-pathing.
ADD 700,FBA	Add the FBA device 700 that's to reflect the SCSI disk (LUN) 4020402600000000.
ADD 701,FBA	Add the FBA device 701 that's to reflect the SCSI disk (LUN) 4020402700000000. Any cuu can be used that isn't backed by a real device. When z/VSE is running in LPAR mode the cuu must not be defined in the IOCDS. When z/VSE is running as a z/VM guest the cuu must not be defined in z/VM.

Define the connection path to the SCSI disks using the sample configuration above: The connection path relates the FBA device (700) to the SCSI disk 4020402600000000, specifies the FCP device (FA0) that carries the SCSI commands and which port on the disk controller side is to be taken (500507630A03C066).

For a system device needed during IPL, for example DOSRES, and SYSWK1 use

```
DEF SCSI,FBA=700,FCP=FA0,WWPN= 500507630A03C066,LUN=4020402600000000
DEF SCSI,FBA=701,FCP=FA0,WWPN= 500507630A03C066,LUN=4020402700000000
```

For a non-system device use

```
SYSDEF SCSI,FBA=700,FCP=FA0,WWPN= 500507630A03C066,LUN=4020402600000000
SYSDEF SCSI,FBA=701,FCP=FA0,WWPN= 500507630A03C066,LUN=4020402700000000
```

High-Availability Through Multi-Pathing

z/VSE allows the definition of additional connection paths to the same SCSI disk. This is called "Multi-pathing". If you want an even more reliable connection to protect your z/VSE system against a potential outage of the FCP adapter or disk controller port, you can use Multi-pathing. If an outage occurs, it isn't presented to the application. Instead the I/O command is automatically repeated on the alternate path. The failing connection path is

automatically recovered by z/VSE. Multi-pathing can enhance availability, but isn't designed to improve performance. It is not used for workload balancing.

Define a second connection path using the above example:

```
DEF SCSI,FBA=700,FCP=FB0,WWPN= 500507630A08C066,LUN=4020402600000000
DEF SCSI,FBA=701,FCP=FB0,WWPN= 500507630A08C066,LUN=4020402700000000
```

```
SYSDEF SCSI,FBA=700,FCP=FB0,WWPN=500507630A08C066,LUN=4020402600000000
SYSDEF SCSI,FBA=701,FCP=FB0,WWPN=500507630A08C066,LUN=4020402700000000
```

Point-to-Point Connections

It is also possible to choose a point-to-point connection where the FCP adapter directly attaches to the disk controller without going through a switch. This may allow a lower-cost configuration, since it eliminates the need for a switch.

Point-to-Point connections require a DS8000 or DS6000 series disk controller.

XIV, Storwize, and SVC always require a switch.

N-Port ID Virtualization (NPIV)

NPIV is an industry standard technology that provides the capability to assign multiple unique world wide port names (WWPNs) to one FCP adapter.

NPIV is supported on an IBM System z9 or later processor. The NPIV feature can be activated for a selected LPAR or for all LPARs sharing an FCP adapter.

Since an FCP adapter (CHPID) can be shared by LPARs, it can be shared by operating systems running in different LPARs and/or z/VM guests.

When the NPIV feature is not used:

Each FCP adapter (CHPID) is assigned a permanent 64-bit WWPN during configuration of the System z server. All subchannels (cuus) on this CHPID have the same permanent WWPN.

Access rights in the SAN and in the disk controller are granted using the WWPN of the FCP adapter. That is each operating system, that uses a subchannel (cuu) on this FCP adapter (CHPID) is identified to the network and in the disk controller by this permanent WWPN.

This means all operating system images sharing this FCP adapter have the same access rights:

- Zone membership for all images sharing the FCP adapter
- Logical Unit Number (LUN) access rights for all images sharing the FCP adapter

When the NPIV feature is used:

The Service Element (SE) creates a own WWPN for each subchannel of the FCP adapter (CHPID).

Each operating system is identified to the network and the disk controller by the NPIV generated WWPN(s) of the subchannel(s) it is using.

Advantage:

- Access rights
NPIV allows to grant access right for LPARs and or z/VM guests on a subchannel base and prevents unauthorized access of LUNs (SCSI disks).
- DASD sharing:
With NPIV you can share the same SCSI disk (LUN) between different z/VSE systems using subchannels on the same FCP adapter.
Without NPIV, z/VSE system requires its own FCP adapter to share SCSI disks (hardware requirement).
- Cost:
Sharing of an FCP adapter between operating systems can save cost for additional adapters.
- Isolation:
Operating systems using the same FCP adapter cannot interfere with each other.

Recommendation:

Always use NPIV.

For details how to setup NPIV, refer to your IBM System z documentation, for example the Support Element Guide. The concepts of NPIV are explained in various publications, one is Redbook REDP-4125-00.

Note:

- Please be aware, that on an IBM System z9, if you replace your FICON Express feature, you will get new WWPNs, also for NPIV. This requires to change the network and disk controller configuration with the new WWPNs.
- Starting with IBM System z10, the same WWPN is generated if the same CHPID is used for the replaced FICON Express feature. This is true also for NPIV.

IPL from SCSI

You perform an IPL of the FCP device and not of the SCSI disk. Therefore, prior to the IPL, you must define the connection path.

Initiating an IPL of z/VSE in an LPAR

On the Load Panel select 'SCSI'. Then enter these values:

Load Address:

Physical address of FCP device (for example FA0)

World Wide Port Name:

Port of the disk controller (for example 500507630A03C066)

Logical Unit Number:

The LUN of the SCSI disk from which z/VSE is to be IPLed
(for example 4020402600000000)

Initiating an IPL of z/VSE in a z/VM Guest

Use the CP SET LOADDEV command to specify the connection path to the SCSI disk. Then IPL the FCP device.

Example:

```
SET LOADDEV PORTNAME 50050763 0A03C066 LUN 40204026 00000000  
IPL FA0
```

z/VSE SCSI Support - Summary

What the z/VSE SCSI provides

- You can use SCSI disks both in an LPAR and a z/VM guest environment.
- Installation on SCSI
- You can install z/VSE on SCSI disks and build a complete z/VSE system using SCSI disks only.
- This includes DASD sharing with the lock file on SCSI. For details on the usage of shared SCSI disks, see the z/VSE Administration manual.
- You can mix and match by, for example, having your z/VSE system on ECKD and using SCSI disks in addition for selected user data.
- SCSI disk size
- z/VSE supports SCSI disks from 8MB to 24GB. The first 4MB of a SCSI disk are internally used by z/VSE and are not shown to the user.
- VSE/VSAM files must be placed in the first 16GB of a SCSI disk.
- Multi-Pathing, that is multiple paths to the same SCSI disk for high availability.
- Point-to-Point connections (DS8000, DS6000 only)
- N_Port ID Virtualization (NPIV) – requires IBM System z9 or later

What the z/VSE SCSI not provides

- Standalone dump on a SCSI disk, that is if you have a SCSI only system you need a physical tape to take a standalone dump.
- Flashcopy for SCSI disks.
- However, copy services offered by the disk controller can be used.
- Attachment of non-disk SCSI devices (for example SCSI tapes)
- Applications using SCSI I/O commands directly.

Concurrent micro code upgrade

z/VSE does not support concurrent micro code upgrade for FCP-attached SCSI disks.

z/VSE running as z/VM guest

When running z/VSE as a z/VM guest, you can choose between the

- the z/VSE SCSI support using FCP-attached SCSI disks
- emulated FBA disks using the z/VM SCSI support

Using emulated FBA disks

An emulated FBA disk is defined to z/VM using the CP SET EDEVICE command. These emulated FBA disks can be attached to virtual machines, or can hold user minidisks. For details how to define and use emulated FBA disks, see the z/VM CP Planning and Administration manual.

Once the emulated FBA disk is defined, it is managed on z/VM like a real 9336 FBA disk; that is, z/VSE views this emulated FBA disk as a real FBA disk and not as a SCSI disk. The z/VSE SCSI support is not involved. z/VSE supports emulated FBA disks with a maximum size of 2 GB (the size of a real FBA disk).

Migration of z/VSE from ECKD to SCSI

This chapter describes the migration of a z/VSE system on (E)CKD to a z/VSE system on SCSI. This migration consists of two steps:

- Initial Installation of z/VSE
- Data Migration

Device structure

(E)CKD devices are cylinder / track oriented. FBA and SCSI devices are block oriented.

Initial Installation Required

Migration of a z/VSE system on ECKD to a z/VSE system on SCSI, that is DOSRES and SYSWK1 on SCSI, requires an initial installation of z/VSE.

Configuration Changes

IPL Procedure

- At least one ADD statement for an FCP device is needed.
- Each SCSI disk requires an ADD statement for an FBA device.
- Each SCSI disk that is required during IPL (system disk) requires a DEF SCSI statement.
- Each additional SCSI requires either a SYSDEF SCSI statement or an IPL DEF SCSI statement. It is recommended to use the SYSDEF SCSI statement. Due to storage constraints about 100 SCSI disks can be defined using DEF SCSI.

Copy Block Consideration - IPL statement SYS BUFSIZE

FBA channel programs might have a different need of copy blocks than (E)CKD channel programs. Depending on the application it might be necessary to adapt the SYS BUFSIZE value.

The copy block usage can be monitored using the AR command SIR.

JCL/ IPL Statements and VSE/VSAM definitions

Cylinder and track specifications for file, catalog, and space definitions must be replaced by block specifications.

How to calculate the number of required blocks

- One track of a 3390 type ECKD disk is about 112 blocks on SCSI; one cylinder is about 1680 blocks.
- A library block of a VSE library is 1024 bytes in length and requires two blocks on SCSI.

Data Migration

Migration of z/VSE Libraries

You can use the Librarian Backup / Restore function to migrate your libraries. A library block of a VSE library is 1024 bytes in length and requires two blocks on SCSI.

Migration of Sequential (SAM) files

- In the JCL EXTENT statements replace track definition by block definition.
- You can use the DITTO copy function to migrate your SAM files.

Migration of VSAM data

The VSE/VSAM LISTCAT function can be used to show the size of currently used clusters. You can take these values to redefine your clusters on the target SCSI disks.

When defining your clusters on the target system it is recommended to use records instead of blocks to be device independent.

Note, that VSE/VSAM might re-calculate and overwrite the CISIZE specification, especially for KSDS clusters.

There are several ways to migrate VSAM clusters:

VSE/VSAM Backup / Restore Function:

This function allows to migrate several clusters all at once. It does not only move data but also defines the clusters in the target catalog.

You can Backup/Restore to and from tape or use a sequential disk file.

If the Restore function fails you should use instead the VSE/VSAM Repro function to migrate the data. A failure of the Restore function is normally indicated by message

- IDC31337I CANNOT RESTORE <dataset name> WITH SPECIFIED MODIFICATIONS

For example Restore of KSDS datasets might fail when they are defined with:

KEY length > 56, CISIZE less than 4096

KEY length > 36, CISIZE less than 2048

KEY length > 6, CISIZE less than 1024

In this case, the Restore will calculate a CISIZE different from the one of the Backup which causes a Restore failure.

VSE/VSAM Export/Import Function:

This function allows to migrate one cluster at a time. Same as Backup/Restore it does not only move data but also defines the cluster in the target catalog.

You can use the Export/Import function to and from a tape or use a sequential disk file.

VSE/VSAM Repro Function:

This function allows to migrate one cluster at a time. It only moves the data. It requires that the cluster is defined in the target catalog.

You can either Repro to and from tape or use Repro from cluster to cluster when you are sharing catalogs between the two systems.

Note:

- Always use Repro for migrating to SCSI devices from clusters to which the IMBED, REPLICATE, or RECOVERABLE attributes currently apply. These attributes are no longer supported.
- Always use Repro for migrating to SCSI devices from clusters which have the SPANNED option specified.

For details refer to the VSE/VSAM documentation.

Migration of DB2 data

- **Increase Copy Blocks – IPL statement SYS BUFSIZE**

When initializing and/or extending the DB2 database on the target system, DB2 programs might have an increased need of z/VSE copy blocks. This depends on the size of the DB2 database and/or extension.

When the z/VSE system runs out of copy blocks the DB2 job to initialize the database might be canceled.

The recommendation is to increase the number of copy blocks temporarily to initialize and/or extend the DB2 database.

Specify SYS BUFSIZE=12000 either in your IPL procedure or use IPL with the LOADPARM to change the BUFSIZE during IPL.

Once you have initialized your DB2 database you can decrease the SYS BUFSIZE value again.

- **Data migration**

Define the VSAM catalog and clusters on the target system. Replace the cylinder / track specification by a block specification.

When defining your clusters on the target system it is recommended to use records instead of blocks to be device independent.

You can use the DBSU UNLOAD / RELOAD to migrate your data.

DB2 application programs:

No changes are required.

Migration of DL/I data

You can use the VSE/VSAM LISTCAT function to show the VSAM usage of the DL/I database. Translate these values in blocks for your target catalog and clusters.

If you are using records for your cluster definition you are device independent.

For the CISIZE definition use the value of the existing DL/I database. Please note that VSE/VSAM might re-calculate and overwrite your CISIZE specification. This can be checked with the LISTCAT function once the catalog and clusters on the target system are created. If CISIZE has changed, DL/I generation has to be adapted. This is discussed in more detail later on.

There are two utilities to migrate an existing DL/I database:

- The DL/I Image Copy utility to backup the data and the DL/I Recovery utility to restore the data.
- The DL/I Unload utility to backup the data and the DL/I Reload utility to restore the data.

Both utilities support tape and disk to migrate the data.

Usually DL/I Image Copy / Recovery is faster. However, the advantage of DL/I Unload / Reload is that the database is reorganized which might provide faster data access afterwards.

Before restoring the data (either utility):

- Define the VSE/VSAM catalog and clusters on the target system.
- Provide a new DBD generation with an adapted DATASET statement:
 - specify FBA in the DEVICE parameter and
 - optionally adapt the SCAN parameter with a block specification
- Provide a new ACB generation (DLZUACB0 utility) with parameter DMB=YES on the BUILD control statement.

CISIZE has changed on the target system:

The following changes are required and/or should be considered if CISIZE on the target system is not identical with CISIZE on the source system. Otherwise restoring of the data might fail, indicated by message DLZ025I for example.

- **If you have a KSDS index database:**

- In the DATASET statement for the DBD generation you have to specify the BLOCK parameter with BLOCK = xx, where xx = CISIZE / record length. The record length can be seen in the DBD listing of the source database.

- **If you have an ESDS HD (HDAM or HIDAM) database:**

- In the DATASET statement for the DBD generation you have to specify the BLOCK parameter with BLOCK = CISIZE, with a CISIZE as in the the target system.
- SCAN specifies a range in which DL/I searches free space to insert a segment. Since CISIZE has changed the number of blocks that are scanned is affected. You may adapt the SCAN parameter to optimize the search.
- You may consider to adapt the FRSPC parameter because CISIZE affects the storage that is left free.

- **If you have an ESDS HDAM database**

To achieve an optimized distribution of the data within the database:

- If you are using ACCESS=HDAM on your DBD statement then you may adapt the operands of the RMNAME parameter.
- If you are using ACCESS=HD on your DBD statement then you may adapt the parameters CIANPT, PRIMCI and RILIM on the DBD statement and RMRTN on the ACCESS statement.

DL/I application programs:
No changes are required.

Migration of VSE POWER Queues

To migrate VSE POWER queues use the POFFLOAD command. You can either use a real or a virtual tape.

Migration of the DTSFILE (ICCF library)

IUI (Interactive User Interface) offers a dialog to backup the ICCF library. You can use:

- 3 Operations
- 7 Backup/Restore
- 2 Backup/Restore Library Objects
- 4 Backup ICCF Library on Tape
- (5 Restore ICCF Library from Tape)

You can either use a real tape or a virtual tape.

Vendor programs

You should check with your vendor, if updates are required when using FBA disks.

Device independent applications

No change is required.

Device dependent applications

Device dependent applications have to change their channel programs and/or language specifications to use FBA interfaces instead of (E)CKD interfaces.

COBOL – PL/1 – RPG programs

These are only a few hints. Dependent on the language or compiler version used, the according documentation needs to be reviewed.

COBOL/VSE programs that are fully compliant with ANSI-85 COBOL, and have been compiled with COBOL/VSE should run unchanged (except for JCL or VSAM definitions).

'Older' COBOL programs that use device specific language elements (for example BLOCK CONTAINS 0) need to be reviewed. Migrated might be needed to be device independent.

RPG programs that use device specific language elements need to be reviewed. Migrated might be needed to be device independent. pg pgpg

Migrated PL/1 programs from DOS/PL1 may still use regional datasets. These datasets are only supported on (E)CKD devices by PL/1. So any applications using these datasets will need to be changed to VSAM instead or to another non-regional type SAM dataset.

Device Characteristics

The z/VSE SCSI support emulates a SCSI disk as FBA 9336 Model 20.

A Read Device Characteristics and Sense ID command returns 777 (x309) blocks per cylinder.

The GETVCE macro uses a 2-bytes field to hold the number of cylinders that is SCSI disks can be up xFFFF cylinders that is about 24 GB.

Note:

The above commands use cylinder terminology even if it is an FBA disk.

Performance and CPU Usage

Our test jobs with heavy I/O load showed improved elapsed times compared with ECKD disks. Because FBA commands are internally translated into SCSI I/O commands, CPU usage increases for jobs using SCSI disks.

In one example, our standard z/VSE batch performance workload was used with an average I/O intensive share. This workload is a job mixture with VSAM and DL/I access, COBOL, and FORTRAN compiles and SORT programs.

CPU usage increased about 13 percent.

So before migrating to a SCSI system, you should check your CPU usage if there is enough capacity for increased CPU usage.

However, every environment is unique and results may vary.

Conclusion

The z/VSE SCSI support offers more storage choices, and fits well into mixed environments where the customer uses SCSI for z/VM and Linux on System z, and may provide better elapsed times for heavy I/O workloads compared to ECKD.