


ZVM14: Using SCSI in a z/VM Environment

Eric Farman, IBM



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.
 - For a complete list of IBM Trademarks, see www.ibm.com/legal/copytrade.shtml
- 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.

Agenda

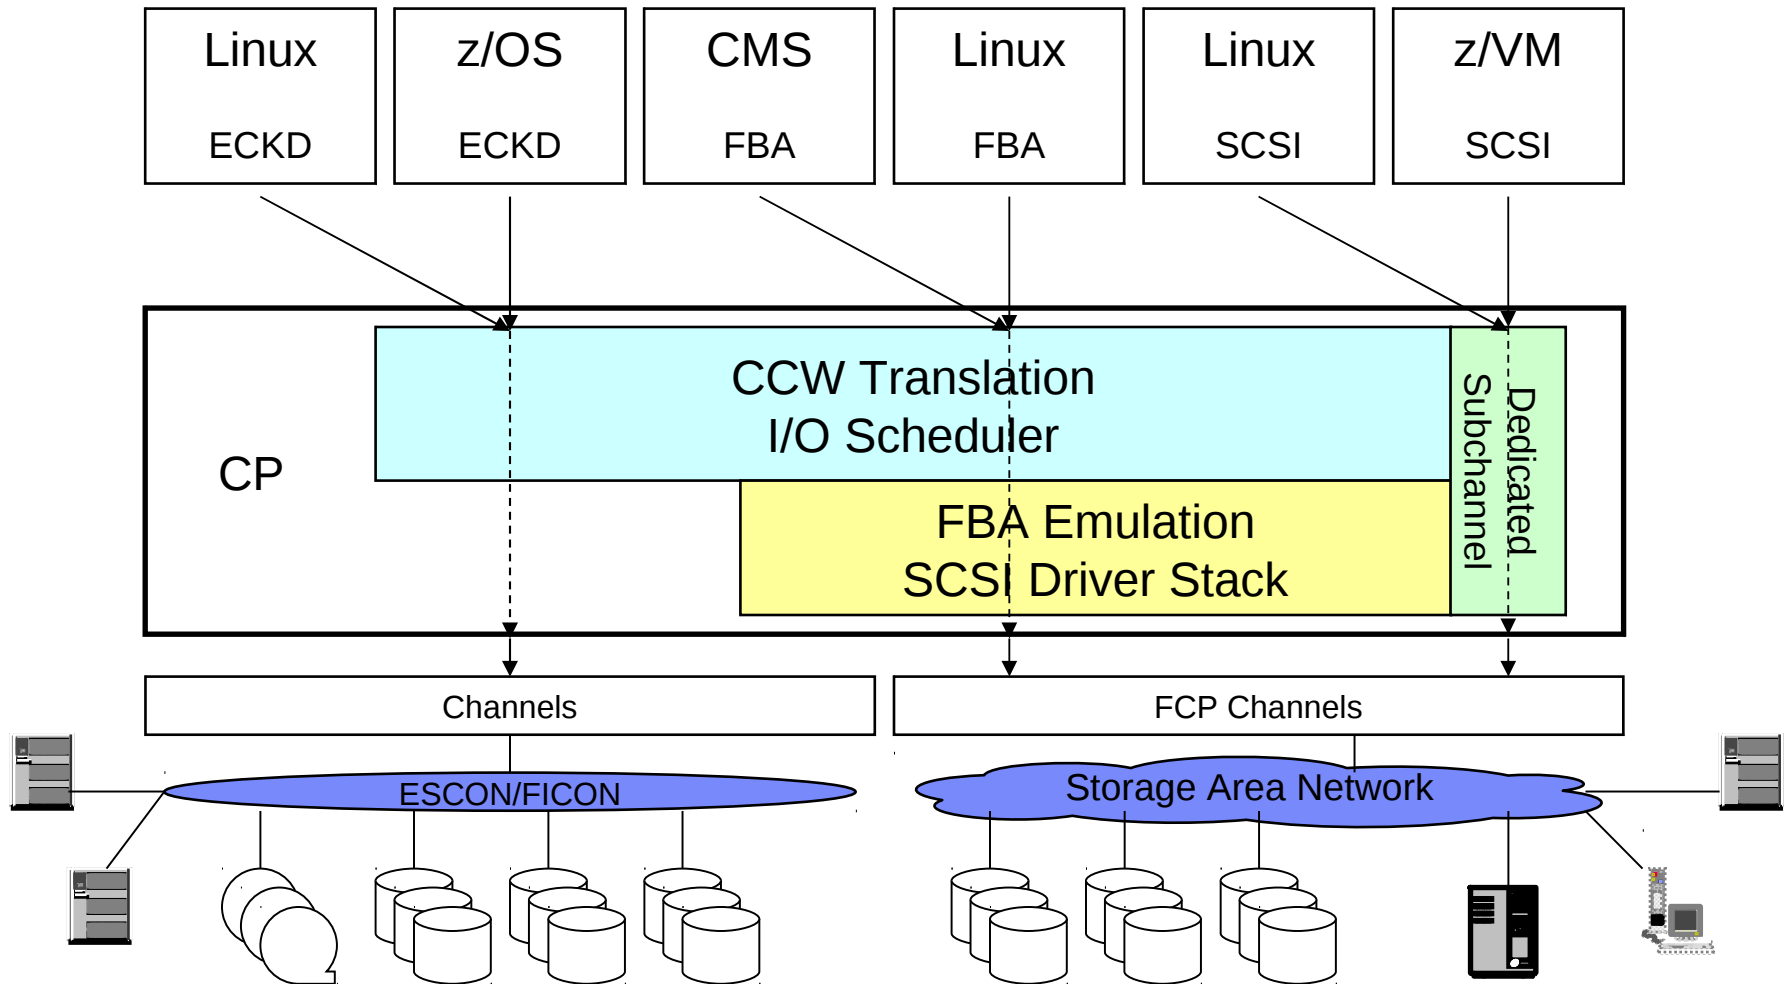
- Overview
 - System Structure
 - FBA Emulation
- Configuration Elements / IOCP
- Commands
- Configuration File Statements / HCD
- IPL
- Preferred Pathing
- Access Control
- SAN Volume Controller (SVC)
- Performance
- Reference Material

A PDF of this presentation is available at:
<http://www.vm.ibm.com/devpages/farman>

Overview

- VM provides native support for SCSI disks for paging, spooling, and other system data
- Support is provided by configuring SCSI disk LUNs to VM as 9336 FBA 512-byte block DASD
- VM guests (such as CMS, GCS, and VSE) may keep data on SCSI disk LUNs without requiring the guest software to have SCSI disk drivers
- IPL, Dump, and Service from/to SCSI disk LUNs is provided to achieve a SCSI-only VM environment
- SCSI-only as well as mixed SCSI and ECKD environments are supported
- Available on all z/VM releases beginning with Version 5 Release 1.0

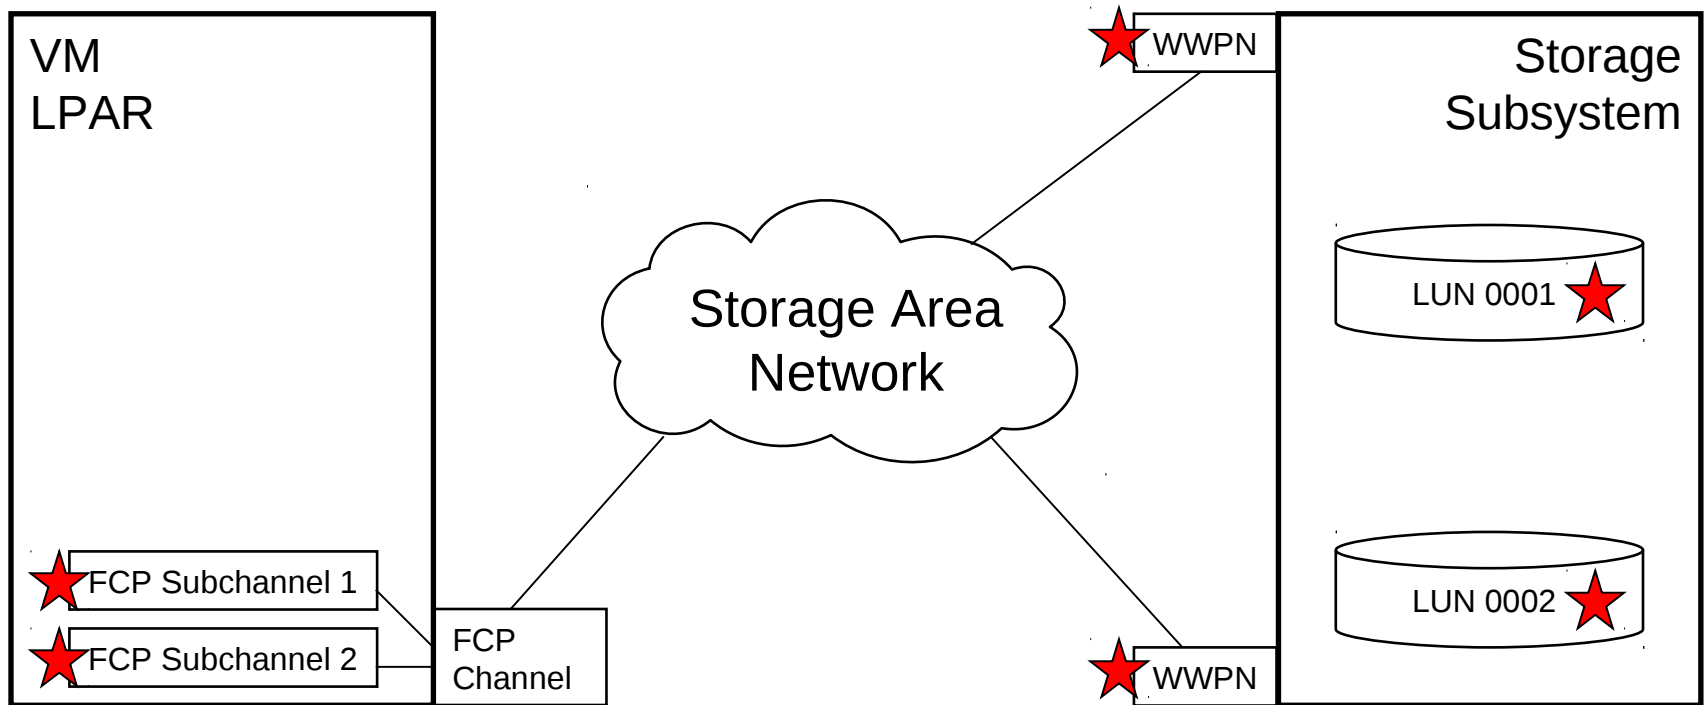
System Structure



FBA Emulation

- SCSI Disks are emulated as 9336 Model 20 FBA DASD
- FBA Emulation allows any operating system or application that supports a 9336 to utilize SCSI disks without change
- Emulated 9336 disk sizes supported are:
 - 1TB for CP with the exception that PAGE, SPOL, and DRCT allocations must remain below the 64GB (minus a 4K page) mark on a CP formatted volume since internal addressing of these slots is limited to 224 4K pages.
 - 381GB for CMS/GCS including software functions dependent on CMS functions such as DIRMAINT MOVE, COPY, ERASE and DFSMS MOVE, COPY, CHECK.
- VM officially supports the following IBM hardware as emulated 9336 DASD (other SCSI disks may work since a generic SCSI driver is provided):
 - DS8000
 - DS6000
 - ESS 750/800
 - SAN Volume Controller
 - XIV

Configuration Elements



★ = VM Configuration Element

- Generally, a shared, switched fabric is most commonly used/supported
- Point-to-Point connection (no switches) is supported on z/VM 5.3.0 and later

IOCP Configuration

```
CHPID PATH=(CSS(0), 50), SHARED, *  
      PARTITION=(LP01, LP02, LP03, LP04, LP05, LP06, LP07, LP08, LP09 *  
      , LP10, LP11, LP12, LP13, LP14, LP15), PCHID=160, TYPE=FCP  
  
CNTLUNIT CUNUMBR=5402, PATH=( (CSS(0), 50), UNIT=FCP  
  
IODEVICE ADDRESS=(5400, 128), CUNUMBR=(5402), *  
      PARTITION=( (CSS(0), LP01) ), UNIT=FCP
```


Commands

- SET EDEVICE command configures a SCSI disk LUN to the VM system as an emulated 9336 FBA DASD
 - Maximum of eight paths per device
- QUERY EDEVICE command obtains information related to SCSI disk LUNs defined to VM as emulated FBA DASD
- DELETE EDEVICE command removes a SCSI disk LUN from the VM system that has been previously defined as an emulated FBA DASD

SET EDEvice Command (Class B)

```
>>-SET--EDEvice--rdev----->
>--Type--FBA--ATTRibutes--1750--| Paths |-->
|                                |-2105-| |-ADD PATH---|
|                                |-2107-| |-DELeTe PATH-'|
|                                |-2145-|
|                                |-XIV--|
|                                '-SCSI-'
'-CLEAR-----
```

Paths:

```
<-----<
| --FCP_DEV--nnnn--WWPN--wwwwwwwwwwwwwww--LUN--1111111111111111-->
>-----|
| -PREFerred---|
| -NOTPREFerred-'
```

Notes: You can specify a maximum of 8 "paths" to the device.
PREF/NOTPREF operands for compatibility only; will be ignored.

System Configuration File Statement

```
>>-EDEVice--rdev--TYpe--FBA--ATTRibutes--.-1750-.-.-| Paths |--><
                                     |-2105-|
                                     |-2107-|
                                     |-2145-|
                                     |-XIV--|
                                     '-SCSI-'
```

Paths:

```
<-----<
|-FCP_DEV--nnnn--WWPN--wwwwwwwwwwwwwww--LUN--1111111111111111-->
>-----|
                                     |-PREferred---|
                                     '-NOTPREferred-'
```

Notes: You can specify a maximum of 8 "paths" to the device.
 PREF/NOTPREF operands for compatibility only; will be ignored.

SET EDEvice Tips

- ATTRibute keyword tells VM which SCSI driver to use
 - Choices are discrete models of IBM hardware, such as 2107 for DS8000-series, or plain “SCSI” for generic or OEM devices
- FCP_DEV is the real device number associated with an FCP subchannel that has connectivity to the LUN (as defined in IOCP)
 - Subchannel will be ATTACHed to SYSTEM when EDEV is brought online
 - At VARY ON, each FCP subchannel must be FREE, or already ATTACHed to SYSTEM
- All 16 digits of the WWPN and LUN must be specified. Otherwise, the VM command parser will add leading zeros causing the path to be invalid.
 - LUN numbers are 2-byte tiers, beginning at left-most digits
 - e.g., “LUN 1” must be specified as “0001000000000000”

SET EDEVICE Tips (Continued)

- An EDEVICE must be varied offline to CLEAR path information.
 - ADD and DELETE can be done to online EDEVICE; see next slide.
- An EDEVICE takes up a slot in VM's real device space. That is, there can't exist another real device with the same real device number.
- Once an EDEVICE is defined, it is managed on VM like a real 9336 FBA. CP commands such as VARY, ATTACH, and QUERY execute as if the emulated disk were a real FBA. This applies also to user directory and system configuration file statements.
 - Q PATHS <rdev> will not show SCSI path information. Use Q EDEV <edev> DETAILS instead.

SET EDEVICE Tips (Continued)

- All EDEVices use an amount of storage from a reserved pool that is established at IPL time
- The size of this pool is determined by the number and complexity of EDEVs in the System Configuration file
 - If none, or a few/simple devices, are defined, a “default” pool size is pre-allocated
 - For larger configurations, the pool is established to hold that configuration, plus some “growth” for new EDEVs
- Don't rely on SET EDEV commands for your entire system; use EDEVICE config statement!
 - Otherwise, devices may not be able to be defined

SET EDEvice – Dynamic path changes

- VM64743, introduced for z/VM 5.4+ in December 2009, allows paths to be added/removed from an Emulated Device while the device is online
- This permits changes to system packs, such as SYSRES or paging volumes, to be made without requiring a system IPL
- No more than eight paths can be added to a given device
- Attempts to delete the last path to a device will not be allowed
- New paths must access the same physical LUN as existing ones

Query EDEvice Command (Class B)

```

          <-----<
>>-Query-EDEvice-...-rdev-----><
      | '-rdev1-rdev2-' '-DETAILS-' |
      '-ALL-----'

```

q edev dead details

EDEV DEAD TYPE FBA ATTRIBUTES 1750

VENDOR: IBM PRODUCT: 1750500 REVISION: .514

BLOCKSIZE: 512 NUMBER OF BLOCKS: 7340032

PATHS:

FCP_DEV: 1F68 WWPN: 500507630E05FC66 LUN: 4010400900000000 **PREF**

CONNECTION TYPE: POINT_TO_POINT

FCP_DEV: 1F98 WWPN: 500507630E85FC66 LUN: 4010400900000000 **NOTPREF**

CONNECTION TYPE: SWITCHED

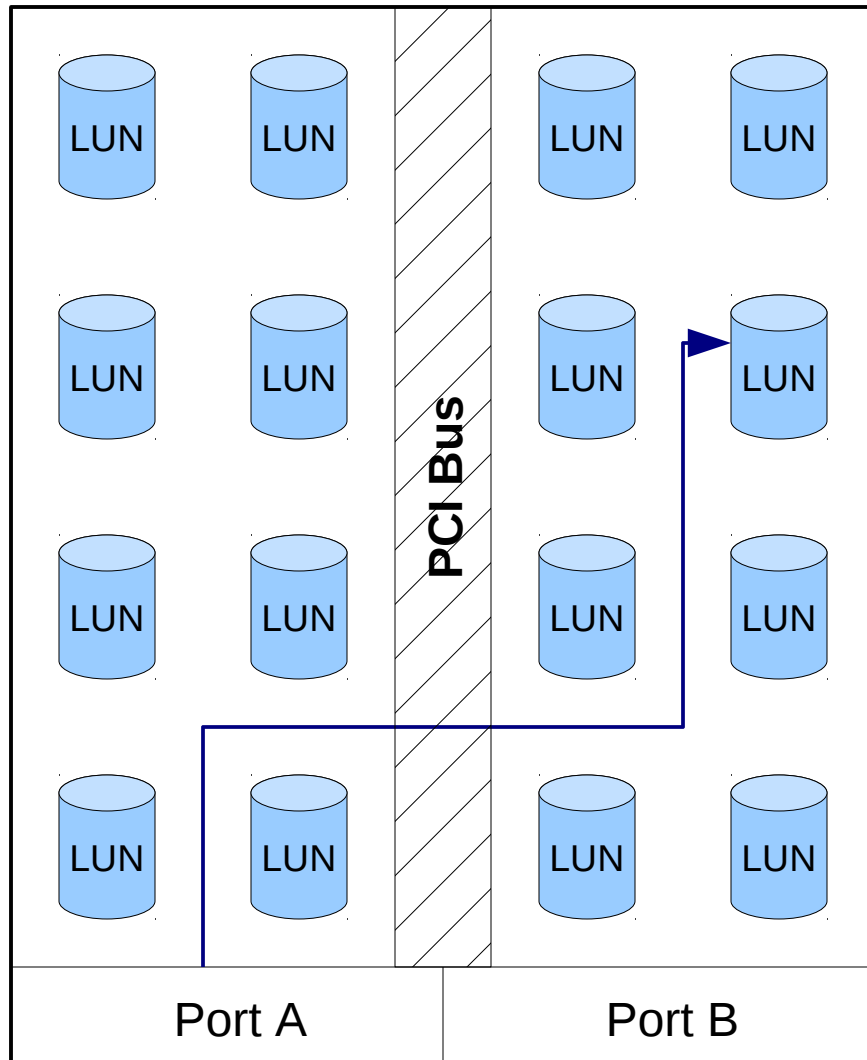
Ready; T=0.01/0.01 15:08:30

DELeTe EDEVice Command (Class B)

```
>>--DELeTe--EDEVice--.-rdev-----.-><  
      '-rdev1-rdev2-'
```

- Removes entry from real device space, unlike SET EDEV CLEAR
- An EDEVice must be offline in order to use this DELeTe command
- Note that CMS also contains a DELeTe command

Logical Path Illustration of Asymmetric Devices



- Port B is preferred path for LUNs on right side of PCI bus; Port A is non-preferred path
- Port A is preferred path for LUNs on left side of PCI bus; Port B is non-preferred

Preferred Pathing

- Some devices, such as DS6000 and SAN Volume Controller, are asymmetric either by software or hardware (or both), such that not all ports have equivalent access to the LUNs on the subsystems
 - In these hardware configurations, half the volumes need to be accessed by crossing a PCI-X bus through a given port (WWPN)
 - This bus crossing can add 30% to response time and reduce bandwidth by 50%
- Preferred versus non-preferred paths will be determined when the EDEVICE is varied online, and z/VM will use this information for its path selection algorithms
 - Prior to z/VM 5.3.0, this information needed to be specified on the SET EDEVICE command (or EDEVICE system configuration file statement)

Preferred Pathing (continued)

- z/VM uses round-robin path selection logic in its SCSI device driver to use only preferred paths and attempts to avoid non-preferred paths
 - I/O will failover (i.e., be redriven) on non-preferred paths only if no preferred paths are available or functioning
- The round-robin path selection logic is used across all paths for other devices, since the preferred vs. non-preferred anomaly is not present

IPL

- z/VM can be IPLed from a SCSI LUN using the Load panel on the SE or HMC
- Stand-Alone Program Loader (SAPL) continues to be the VM IPL interface
 - SAPL contains its own generic, stripped down SCSI stack so that it can read the VM load module from disk into memory.
- VM IPL parameter PDVOL must be specified for a SCSI IPL. It indicates the EDEVICE number of the SYSRES.
 - System stops with a wait state 6505 if not specified
 - Can be entered on the SAPL screen (PDVOL=xxxx) or defined when SAPL is set up ('burned in')

| Load - H34:EEXT2 | |
|---|--|
| CPC: | H34:EEXT2 |
| Image: | H34:EEXT2 |
| Load type | <input type="radio"/> Normal <input type="radio"/> Clear <input checked="" type="radio"/> SCSI <input type="radio"/> S |
| <input type="checkbox"/> Store status | |
| Load address | * 521E |
| Load parameter | SYSG |
| Time-out value | 60 |
| Worldwide port name | 5005076300C604D |
| Logical unit number | 5144000000000000 |
| Boot program selector | 0 |
| Boot record logical block address | 00000000000000C8 |
| Operating system specific load parameters | |

Second Level IPL

- Second-level IPL from a SCSI LUN works via support that went in originally for Linux Guest IPL from SCSI
- Use CP SET LOADDEV command to specify WWPN, LUN, Boot Program Selector and Boot Record Logical Block Address
- Use the same Boot Program Selector and Logical Block Address as first-level IPL for a CP formatted disk (i.e., 0 and 0000000000000000C8)
 - Add in start of RECOMP area for a CMS formatted disk.
- IPL virtual address of FCP subchannel to start load

Second Level IPL (example)

```
att 50aa *
```

```
FCP 50AA ATTACHED TO STEVO 50AA
```

```
Ready; T=0.01/0.01 13:16:20
```

```
q v fcp
```

```
FCP 50AA ON FCP 50AA CHPID 40 SUBCHANNEL = 000E
```

```
50AA QDIO-ELIGIBLE QIOASSIST-ELIGIBLE
```

```
Ready; T=0.01/0.01 13:16:24
```

```
set loaddev portname 50050763 00c20b8e lun 52410000 00000000 bootprog 0 br_lba  
00000000 0000000C8
```

```
Ready; T=0.01/0.01 13:16:33
```

```
q loaddev
```

```
PORTNAME 50050763 00C20B8E LUN 52410000 00000000
```

```
BOOTPROG 0 BR_LBA 00000000 0000000C8
```

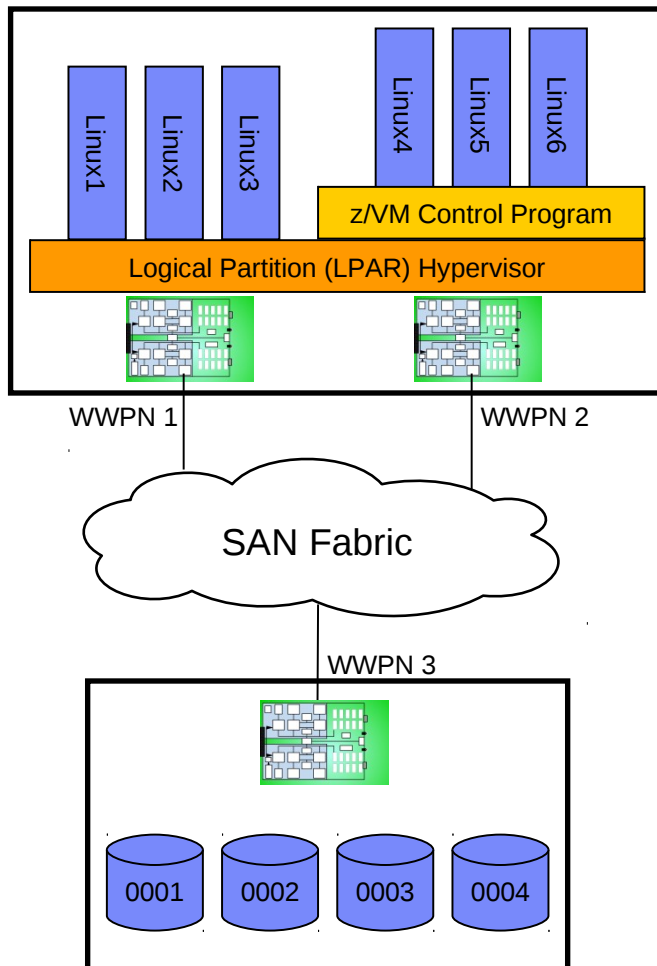
```
Ready; T=0.01/0.01 13:16:38
```

```
IPL 50aa <Hit ENTER and the IPL starts...>
```


SALIPL

- The SALIPL utility continues to set up SAPL for IPL
- SALIPL now writes to blocks 5-207 of a 9336 device (previously wrote only to blocks 5-31)
 - Affects size of RECOMP area for a CMS minidisk containing SAPL, or
 - Affects placement of allocations such as PAGE and SPOL when SAPL is put on a CP formatted volume
- SALIPL must run in a virtual machine to setup a SCSI disk for IPL SALIPL can be run against either:
 - A virtual device, such as a fullpack minidisk, on an emulated 9336
 - A virtual FCP subchannel where new SALIPL parameters WWPN and LUN designate the target SCSI disk

Access Control



- Problem: Without Access Control
 - LUN Masking and Zoning only allow access control per WWPN of processor FCP Channel Adapter
 - LUN masking (in controllers) and zoning (in switches) are industry standard methods of controlling access to target control units and LUNs
 - No individual access control/protection for Operating Systems sharing the FCP Channel Adapter
- With N_Port ID Virtualization (NPIV)
 - Hardware creates unique virtual WWPNs for each FCP subchannel on Channel Adapter
 - Can use LUN masking and zoning, based on the identity of the Virtual WWPNs associated with each FCP subchannel

N_Port ID Virtualization (NPIV)

- Virtual WWPNS are assigned to each FCP subchannel by the channel hardware and are managed at the Support Element (SE)
 - A VM guest inherits the virtual WWPNS assigned to a FCP subchannel when the subchannel is attached to the guest
 - The VM system inherits the virtual WWPNS when a FCP subchannel is attached to the system via the SET EDEVICE command
- Standard access controls such as zoning in switches and LUN masking in the controller can be used with the virtual WWPNS
- NPIV is supported on System z9 and later processors (for both Enterprise and Business classes)
- See Redbooks Paper “Introducing N_Port Identifier Virtualization for IBM System z9” (REDP-4125)

SAN Volume Controller (SVC) and Storwize V7000

- Combination Hardware/Software appliance that presents a homogeneous pool of generic SCSI devices out of heterogeneous DASD subsystems
 - Device-specific nuances are handled by SVC or V7000, not host OS
- Subsystems can be added/removed/upgraded without impacting host connections
- Earliest supported software levels:
 - SVC 4.1.0
 - Linux SLES 9 SP3
 - Linux RHEL 4 (2.6.9-42.EL.m3)
 - z/VM 5.3.0+ (with PTF for APAR VM64679)
- Refer to interoperability matrices for supported attachments
 - <http://www-03.ibm.com/systems/storage/software/virtualization/svc/index.html>
 - http://www-03.ibm.com/systems/storage/disk/storwize_v7000/interop.html

XIV

- High-end open disk storage platform, with high availability, performance, reliability and scalability
- New “XIV” ATTRibute on VM EDEVICE command/statement for these connections
- Avoid saturating an FCP subchannel (e.g., with a high memory overcommitment resulting in high paging rate) to avoid performance degradation
- Earliest supported software levels:
 - XIV 10.2.0.A (or later)
 - Linux SLES 10 SP2 (or later)
 - Linux RHEL 5.2 (or later)
 - z/VM 5.4.0+ (with PTF for APAR VM64708)
- See most current EDEVIC information at:
 - <http://www.vm.ibm.com/Storman/XIV/>

Performance Best Practices

- Bottom Line:


- Significant hypervisor path length increase with EDEVICE I/O processing compared to ECKD, resulting in higher CPU usage.

- Recommendations:

- Use dedicated FCP subchannels for Linux guests except when VM minidisks can add needed sharing capabilities, disk administrative savings, or Minidisk Cache benefits.
- Use traditional ECKD DASD for paging and spooling except when little paging/spooling activity exists, or sufficient processor resources are available to handle increased path length.
- Continue to use traditional ECKD DASD for CMS minidisks except when minidisk cache hit ratios are high, or sufficient processor resources are available to handle increased path length.
- Consider EDEVICES when no ECKD DASD is available, minidisks are required, or large disks are necessary.
- For information regarding the studies that led to these recommendations, see these writeups:
 - CP Disk I/O Performance - www.vm.ibm.com/perf/reports/zvm/html/520dasd.html
 - Linux Disk I/O Alternatives - www.vm.ibm.com/perf/reports/zvm/html/520lxd.html
 - SCSI Performance - www.vm.ibm.com/perf/reports/zvm/html/530scsi.html

Fin

Eric Farman
farman@us.ibm.com
z/VM I/O Development



Reference



HCD

- HCM GUI interface to VM's HCD support will allow an EDEV to be created with up to 8 unique paths.

Define Device <==> OS Configuration Parameters

Device Number: 8891 OS Config: CT1

Number of devices: 1 Type: VM

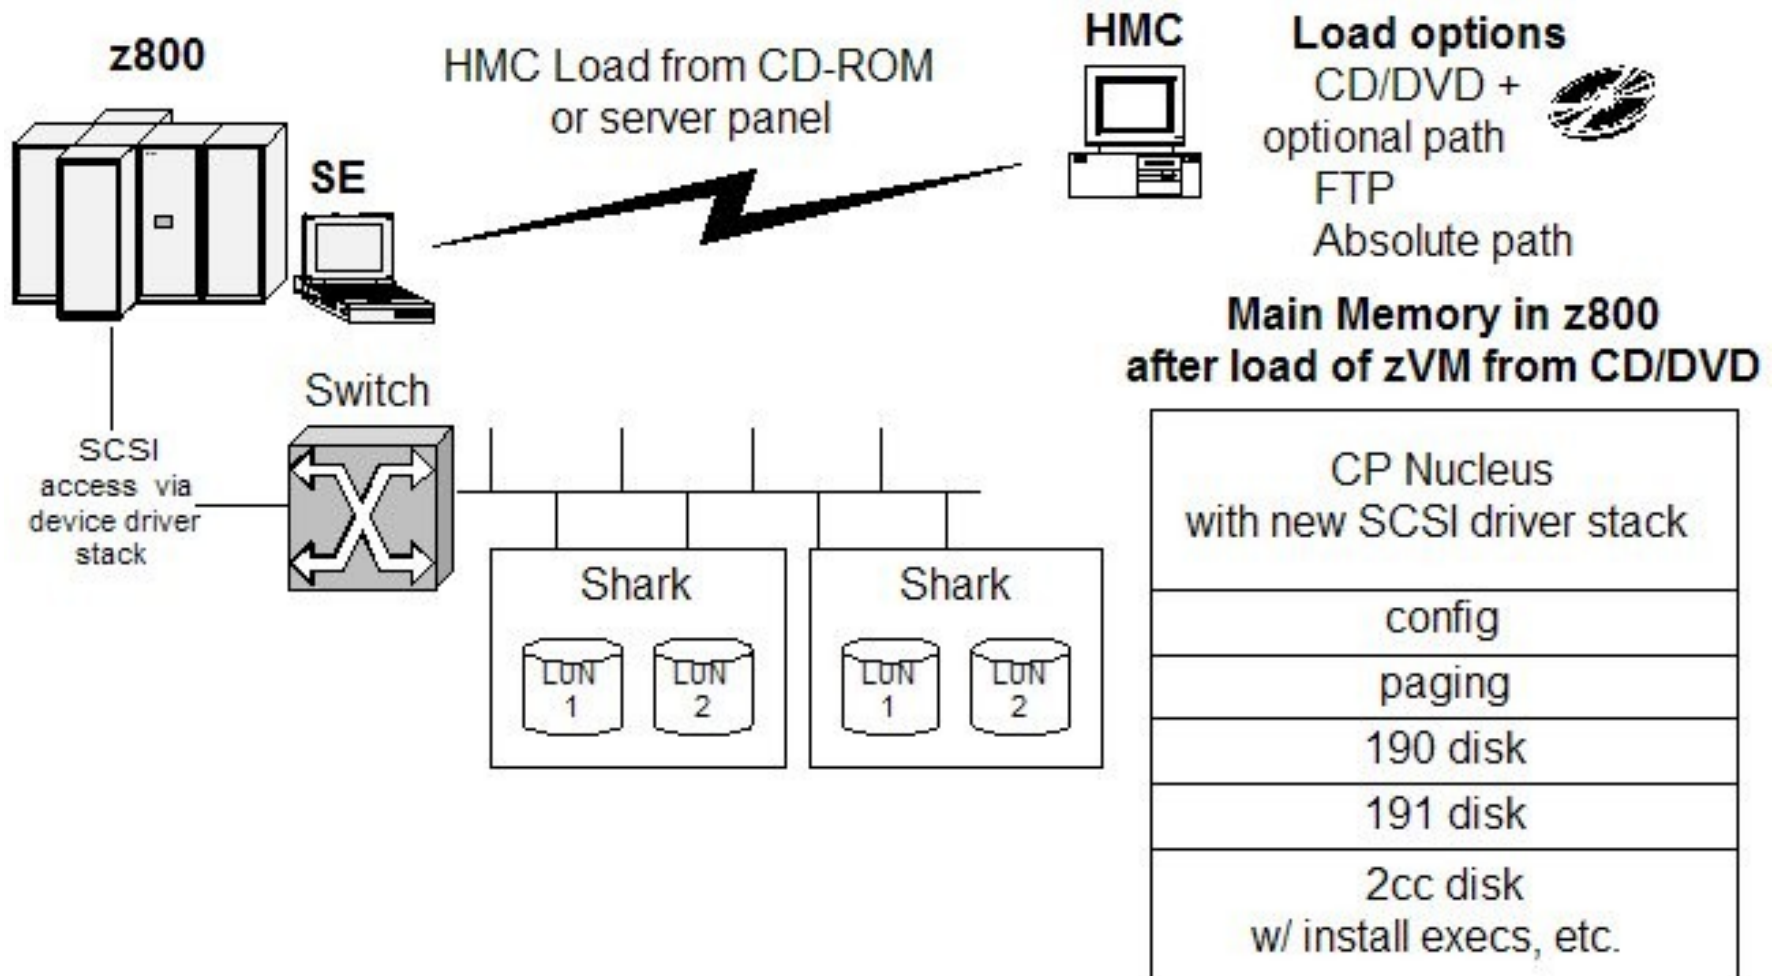
Type: FBASCSI Description: TCFT

| Parameter | Value | Req | Description | |
|-----------|------------------|-------------------------------------|-----------------------------------|------------------------------------|
| ATTR | 2105 | <input checked="" type="checkbox"/> | Name of SCSI Device attribute set | ? <input type="button" value="▲"/> |
| FCPDEV | 2A00 | <input checked="" type="checkbox"/> | Associated FCP device | ? <input type="button" value="▲"/> |
| FCPDEV -2 | 2A01 | | Associated FCP device | ? <input type="button" value="▲"/> |
| FCPDEV -3 | | | Associated FCP device | ? <input type="button" value="▲"/> |
| FCPDEV -4 | | | Associated FCP device | ? <input type="button" value="▲"/> |
| FCPDEV -5 | | | Associated FCP device | ? <input type="button" value="▲"/> |
| FCPDEV -6 | | | Associated FCP device | ? <input type="button" value="▲"/> |
| FCPDEV -7 | | | Associated FCP device | ? <input type="button" value="▲"/> |
| FCPDEV -8 | | | Associated FCP device | ? <input type="button" value="▲"/> |
| WWPN | 5005076300C204DA | <input checked="" type="checkbox"/> | World Wide Port Number | ? <input type="button" value="▲"/> |
| WWPN -2 | 5005076300CE04DA | | World Wide Port Number | ? <input type="button" value="▲"/> |
| WWPN -3 | | | World Wide Port Number | ? <input type="button" value="▲"/> |
| WWPN -4 | | | World Wide Port Number | ? <input type="button" value="▲"/> |

Install

- Install to SCSI is done via DVD from the Hardware Management Console (HMC).
- HMC software with DVD Load and Integrated 3270 console support is required. A 512MB LPAR is also required.
- Install uses a special HMC hardware interface to bring in the VM starter system. The starter system uses a RAM disk to get itself and the install procedure started. Install writes the VM system and other files out to SCSI disk using FBA emulation (see next slide).
- Second-level install to SCSI is also supported. Requires FTP transfer of install files on the first-level system. Transfer is done automatically from the install DVD by the installation utility.
- Install from DVD is also available for 3390 models 3 and 9

Install Illustration



Dump and Service

- Support has been added to take an ABEND dump (i.e., hard, soft, SNAPDUMP, VMDUMP) to a SCSI LUN. Stand-alone dump to SCSI is not supported.
- Dumps can be submitted to VM Level 2 electronically via FTP for customers without an Enterprise tape drive (e.g., 3490 or 3590)
- Dumps can be copied to tape via FBA emulation for customers with an Enterprise tape drive using existing tools such as the CMS TAPE command
- Service is obtained electronically for customers without an Enterprise tape drive
- Service is applied by existing tools without change using FBA emulation (i.e., put service files on emulated 9336 and run VMSES tools)
- Service is business as usual for customers with an Enterprise tape drive

EDEVice Performance

Positives

- Allows exploitation of minidisk cache
- Allows minidisks on SCSI to be shared
- Improves performance measurement capabilities compared to FCP subchannels dedicated to a guest
- FCP paths are faster than ESCON
- Greater flexibility in size of minidisks based on SCSI devices tending to be larger in general

Negatives

- Significant path length increase, so one needs to insure proper planning for enough processor cycles to handle the I/O rate
- Monitoring capabilities are not as strong as traditional ECKD in isolating problems at the channel, control unit, and device levels.
- FCP paths are not significantly faster than FICON (from a hardware perspective)

Monitor Records

- **MRMTRDEV (Device Configuration Data record in Monitor Domain)**
 - Changed the reserved byte at the end of the record to a new flag byte called `MTRDEV_EDEVTYPE`. This byte is a code associated with the type of SCSI device. For example, SCSI disk over FCP protocol. Also added the following information to the bottom of the record:
 - WWPN(s)
 - FCP subchannel(s)
 - LUN Address
 - Attribute Name
- **MRIODDEV (Device Activity record in I/O Domain)**
 - Added new flag byte called `IODDEV_EDEVTYPE` at the end of the record. This byte is a code associated with the type of SCSI device. Added `IODDEV_RDEVDEV` (real device number) to the end of the record. Previously, the record only contained `RDEVSID` (subchannel id) which isn't unique for SCSI devices. Adding `RDEVDEV` enables Monitor to uniquely identify the device.
- **MRIODVON (Vary On Device record in I/O Domain)**
 - Changed the reserved byte at the end of the record to a new flag byte called `IODVON_EDEVTYPE`. This byte is a code associated with the type of SCSI device. Also added WWPN, FCP subchannel, and LUN path information.

Monitor Records

- MRIODDTD (Detach Device record in I/O Domain)
 - Added IODDTD_RDEVDEV (real device number) to the end of the record so that Monitor can uniquely identify the device.
- MRIODVOF (Vary Off Device record in I/O Domain)
 - Added IODVOF_RDEVDEV (real device number) to the end of the record so that Monitor can uniquely identify the device.
- MRSTOASS - (Auxiliary Shared Storage Management record in Storage Domain)
 - Added STOASS_RDEVDEV (real device number) to the end of the record so that Monitor can uniquely identify the device.
- MRSTOATC - (Page/Spool Area of a CP Volume record in Storage Domain)
 - Added STOATC_RDEVDEV (real device number) to the end of the record so that Monitor can uniquely identify the device.
- MRMTRPAG - (Paging Configuration Data record in Monitor Domain)
 - Added MTRPAG_RDEVDEV (real device number) to the end of the record so that Monitor can uniquely identify the device.

Monitor Records

- MRSEKSEK - (Seek Data record in Seek Domain)
 - Added SEKSEK_RDEVDEV (real device number) to the end of the record so that Monitor can uniquely identify the device.
- MRIODSZI (New I/O Domain Monitor Record 24)
 - This is the SCSI Device Activity Record and will be cut in addition to the existing MRIODDEV record. The first 20 bytes of the record will contain the standard header for monitor records. The rest of the record will contain the following:
 - RDEVDEV
 - Bytes in a block (512)
 - > amount of time device is active
 - > kbytes/sec transfer rate
 - > #transfers to/from disk
 - > #blocks read from disk
 - > #blocks written to disk
 - > #seek operations
 - > I/O queue depth
 - Statistics on per path basis (same as > arrows)

HMC/SE Requirements

- The Hardware Management Console (HMC) must be communicating with the Support Element (SE). The HMC can only communicate with versions of the SE that are equal to or lower than the HMC. For example, HMC version 1.8.0 can communicate with a SE at version 1.7.3, or 1.8.0, but it can not communicate with a SE at version 1.8.2.
- The following minimum SE levels are required to install from DVD:
 - z800 - Support Element (SE) version 1.7.3, Engineering Change (EC) J11213, change level 146 or higher must be active.
 - z890 - Support Element (SE) version 1.8.2. No Licensed Internal Code changes are required.
 - z900 - Support Element (SE) version 1.7.3, Engineering Change (EC) J11213, change level 146 or higher must be active.
 - z990 - Support Element (SE) version 1.8.0, Engineering Change (EC) J12560, change level 054 or higher must be active. Or, Support Element (SE) version 1.8.2. No Licensed Internal Code changes are required.

FCP Firmware Requirements

- The following processor driver levels and FCP channel code levels are the minimum required to run z/VM with SCSI disks:
 - z800 - Driver D3G with FCP Code 0.28, MCL J11233 #014
 - z890 - Driver D55 with FCP Code 3.04, MCL J13471 #003
 - z900 - Driver D3G with FCP Code 0.28, MCL J11233 #014
 - z990 - Driver D52 with FCP Code 2.07, MCL J12951 #004 for GA2 level z990
 - z990 - Driver D55 with FCP Code 3.04, MCL J13471 #003 for GA3 level z990
 - z9 EC or BC - base code
- Driver changes require a POR. FCP code changes are concurrent; except on z800, z900, and z990-GA2 (with driver D52) where CHPID must be recycled off to on.

System Storage Firmware Requirements

- 2105 Model F
 - LIC EC Identification Number: 2.3.2.108 EC Level: SC40723
- 2105 Model 800
 - LIC EC Identification Number: 2.4.1.58 EC Level: SC40908
- 2107 (DS8000)
 - Latest microcode bundle
- 1750 (DS6000)
 - Latest microcode bundle
- 2145 (SAN Volume Controller)
 - Latest microcode bundle
- XIV
 - Minimum microcode load 10.2.0.A