

2012

**IBM System z Technical University**

Enabling the infrastructure for smarter computing

# **SCSI over FCP for Linux on System z**

**zLG09**

Susanne Wintenberger



## 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](http://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.

\* All other products may be trademarks or registered trademarks of their respective companies.

IT Infrastructure Library is a registered trademark of the Central Computer and Telecommunications Agency, which is now part of the Office of Government Commerce.

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

- Introduction to FCP on System z
- FCP with Linux on System z
- IPL over FCP
- SCSI dump
- Multipathing
- N-Port Id Virtualization (NPIV)
- Troubleshooting

## FCP in a Nutshell



- Storage Area Networks (SANs) are specialized networks dedicated to the transport of mass storage data
- Today the most common SAN technology used is Fibre Channel Protocol (FCP)
- With this technology the SCSI protocol is used to address and transfer raw data between the servers and the storage device
- Each server is equipped with a least one adapter which provides the physical connection to the SAN
- For System z any supported FCP adapter, such as FICON Express, FICON Express2, FICON Express4 or FICON Express8 can be used for this purpose.
- The Fibre Channel (FC) standard was developed by the National Committee of Information Technology Standards (NCITS)

## Why FCP?

- Performance advantages
  - FCP is much faster than FICON
  - Reason 1: asynchronous I/O
  - Reason 2: no ECKD emulation overhead
- No disk size restrictions
- Up to 15 partitions (16 minor numbers per device)
- SCSI disks do not waste disk space (no low-level formatting)
- System z integration in existing FC SANs
- Use of existing FICON infrastructure
  - FICON Express adapter cards
  - FC switches / Cabling
  - Storage subsystems
- Dynamic configuration
  - Adding of new storage subsystems possible without IOCDS change
- Does NOT require more CPU than FICON

# SAN topologies and System z


point-to-point




direct attached  
arbitrated loop



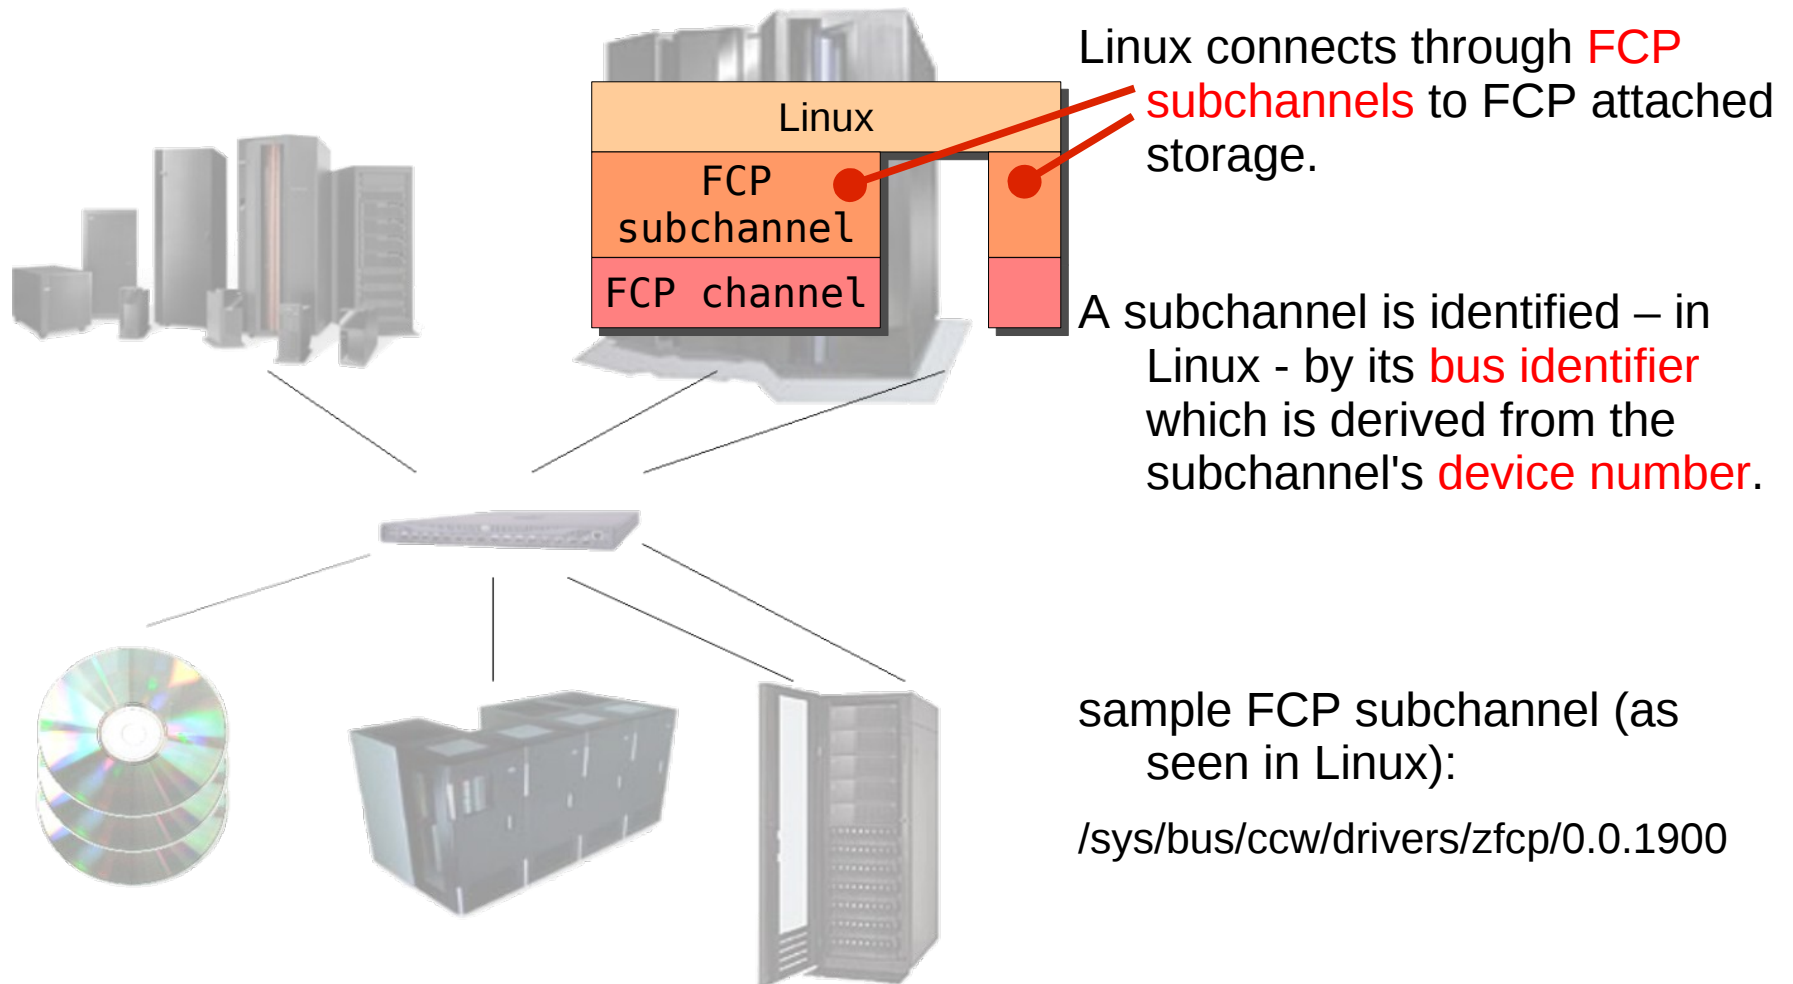
**not supported**




switched fabric



## FCP channel and subchannel



# World Wide Port Names (WWPN)



Storage devices and servers attach through Fibre Channel ports (called N\_Ports).


An N\_Port is identified by its **World-Wide Port Name (WWPN)**.

For redundancy, servers or storage may attach through several N\_Ports.

sample WWPN:  
0x5005076303000104



## Logical Unit Number (LUNs)

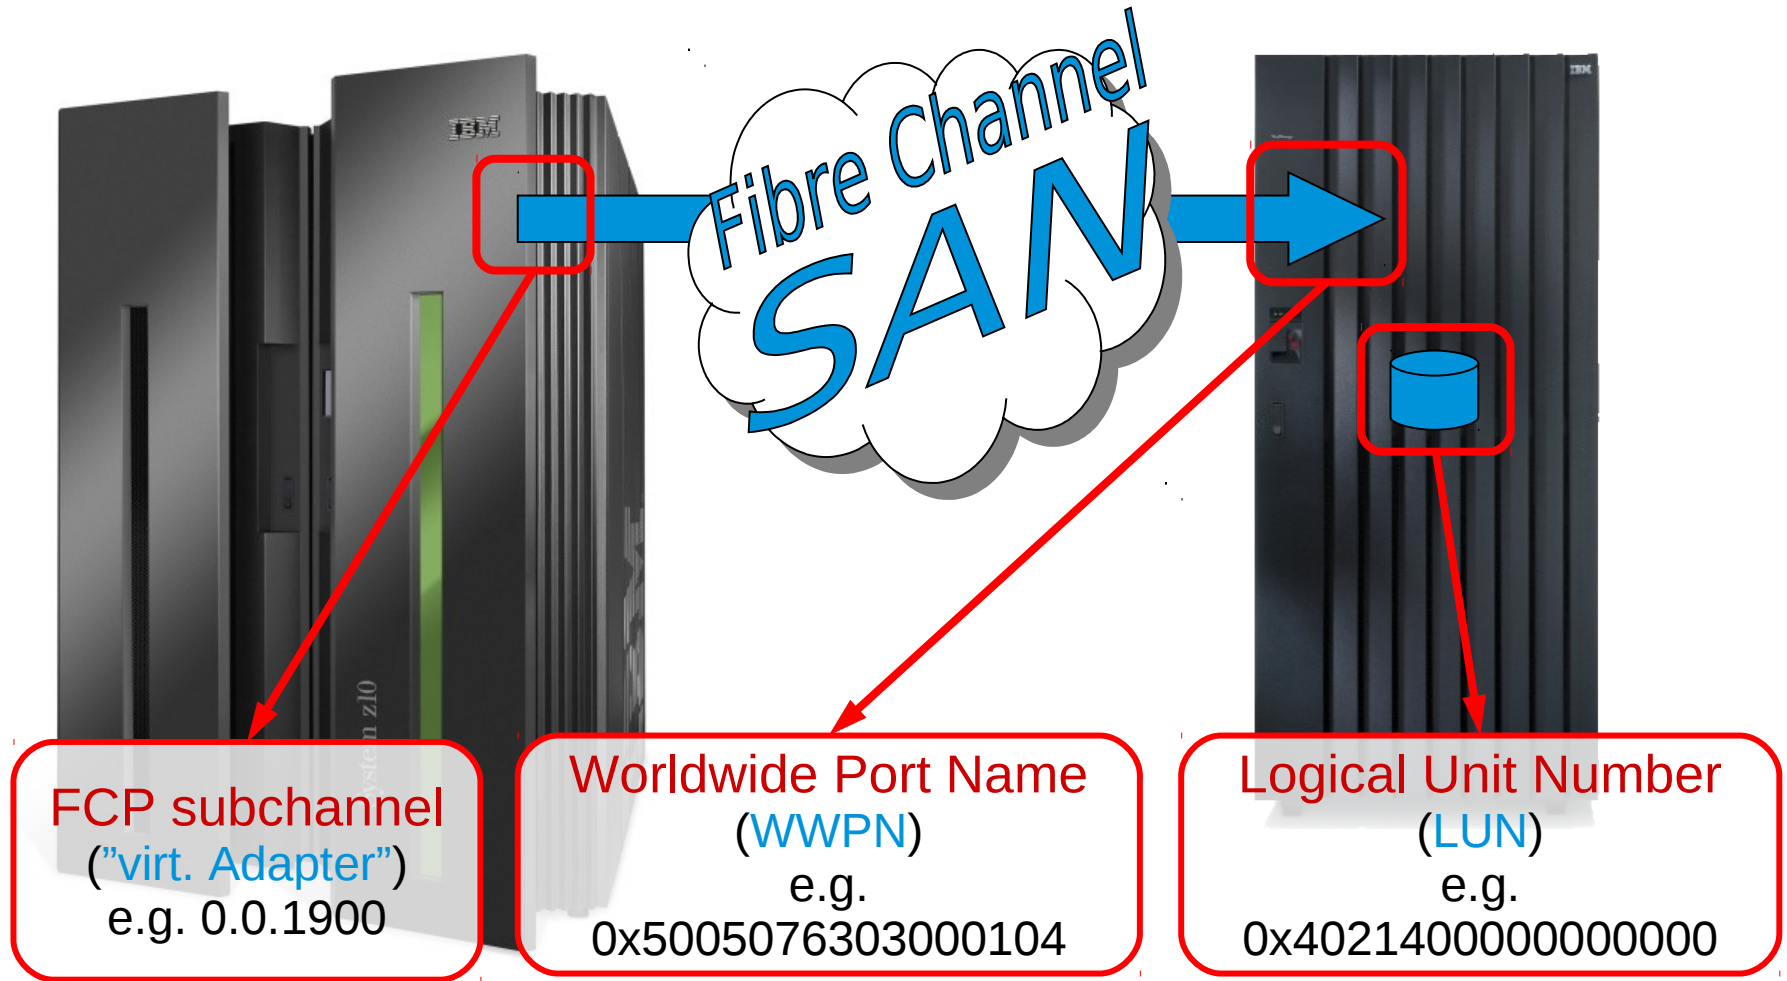


Storage devices usually comprise many logical units (volumes, tape drives, ...).

A logical unit is identified by its **Fibre Channel Protocol Logical Unit Number (FCP LUN)**.

sample FCP LUN:  
0x40**21**40**00**00000000  
Beware of LUN translation!

# Navigating in a SAN



## SCSI compared to Channel I/O

### ▪ SCSI / FCP

- Channel subsystem device represents the virtual adapter to the Fibre Channel SAN
- FCP adapter is defined in System z I/O configuration
- Ports and LUNs attachment handled in Operating Systems
- Multipathing handled in Operating System
- No disk size restrictions for SCSI disks
- Additional configuration outside System z necessary
  - Zoning in the SAN fabric
  - LUN masking on the storage server

### ▪ Channel I/O

- Channel subsystem device represents disk volume (ECKD)
- Device defined in System z I/O configuration
- Ports attachment handled in System z I/O configuration
- Multipathing handled in System z firmware
- Disk size restrictions to Mod 54 / Mod 224
- Switch configuration via System z I/O configuration

## ZFCP getting started

- To access a Fibre Channel-attached SCSI device follow these configuration steps:
  - 1) Configure a Fibre Channel host adapter within the mainframe (I/O Definition File).
  - 2) Configure **zoning** for the Fibre Channel host adapter to gain access to desired target ports within a SAN.
    - Segmentation of a switched fabric is achieved through zoning. It can be used to partition off certain portions of the switched fabric, allowing only the members of a zone to communicate with that zone.
  - 3) Configure **LUN masking** for the Fibre Channel host adapter at the target device to gain access to desired LUNs.
    - A LUN represents a portion of a controller, such as a disk device. With the use of LUNs, a controller can be logically divided into independent partitions. Access to these LUNs can be restricted to distinctive WWPNs as part of the controller configuration
  - 4) **In Linux, configure target ports and LUNs of the SCSI device at the target port for use of zfc.**
- Note: If the Fibre Channel host adapter is directly attached to a target device (point-to-point connection), step 2 is not needed.

## Hardware: Define FCP adapter in IOCDs

```

CHPID PATH=(CSS(0,1,2,3),51),SHARED, *
NOTPART=( (CSS(1), (TRX1), (=) ), (CSS(3), (TRX2, T29CFA), (=) ) ) *
PCHID=1C3, TYPE=FCP

CNTLUNIT CUNUMBR=3D00, *
PATH=( (CSS(0), 51), (CSS(1), 51), (CSS(2), 51), (CSS(3), 51)), *
UNIT=FCP

IODEVICE ADDRESS=(3D00, 001), CUNUMBR=(3D00), UNIT=FCP

IODEVICE ADDRESS=(3D01, 007), CUNUMBR=(3D00), *
PARTITION=( (CSS(0), T29LP11, T29LP12, T29LP13, T29LP14, T29LP*
15), (CSS(1), T29LP26, T29LP27, T29LP29, T29LP30), (CSS(2), T29*
LP41, T29LP42, T29LP43, T29LP44, T29LP45), (CSS(3), T29LP56, T2*
9LP57, T29LP58, T29LP59, T29LP60)), UNIT=FCP

IODEVICE ADDRESS=(3D08, 056), CUNUMBR=(3D00), *
PARTITION=( (CSS(0), T29LP15), (CSS(1), T29LP30), (CSS(2), T29*
LP45), (CSS(3), T29LP60)), UNIT=FCP

```

## Manual LUN configuration with s390-tools

### Show available FCP devices:

```
# lscss -t 1732/03,1732/04
```

Device	Subchan.	DevType	CU Type	Use	PIM	PAM	POM	CHPIDs
0.0.1900	0.0.0010	1732/03	1731/03		80	80	ff	3b000000 00000000
0.0.1940	0.0.0011	1732/03	1731/03		80	80	ff	46000000 00000000

### Enable FCP devices:

```
# chccwdev -e 0.0.1900,0.0.1940
```

```
Setting device 0.0.1900 online
```

```
Done
```

```
Setting device 0.0.1940 online
```

```
Done
```

### Show online FCP devices and corresponding SCSI hosts:

```
# lszfcp
```

```
0.0.1900 host0
```

```
0.0.1940 host1
```

## Manual LUN configuration with s390-tools (cont'd)

### **Attach ports, only required for older distributions (SLES10, RHEL5, ...):**

```
# echo 0x5005076303000104 > /sys/bus/ccw/drivers/zfcp/0.0.1900/port_add
# echo 0x5005076303100104 > /sys/bus/ccw/drivers/zfcp/0.0.1940/port_add
```

### **Show available storage ports:**

```
# lszfcp -P
0.0.1900/0x5005076303000104 rport-0:0-8
0.0.1940/0x5005076303100104 rport-1:0-0
```

### **Attach LUNs:**

```
# echo 0x4021400000000000 >
/sys/bus/ccw/drivers/zfcp/0.0.1900/0x5005076303000104/unit_add
# echo 0x4021400000000000 >
/sys/bus/ccw/drivers/zfcp/0.0.1940/0x5005076303100104/unit_add
```

### **List available LUNs:**

```
# lszfcp -D
0.0.1900/0x5005076303000104/0x4021400000000000 0:0:8:1074872353
0.0.1940/0x5005076303100104/0x4021400000000000 1:0:0:1074872353
```

## Linux Tools: lsscsi and lslns

**lsscsi: Show SCSI devices and block device nodes:**

```
# lsscsi
[0:0:8:1074872353]disk      IBM          2107900      .280  /dev/sda
[1:0:0:1074872353]disk      IBM          2107900      .280  /dev/sdb
```

**lslns: Which LUNs are available on the storage system for use?**

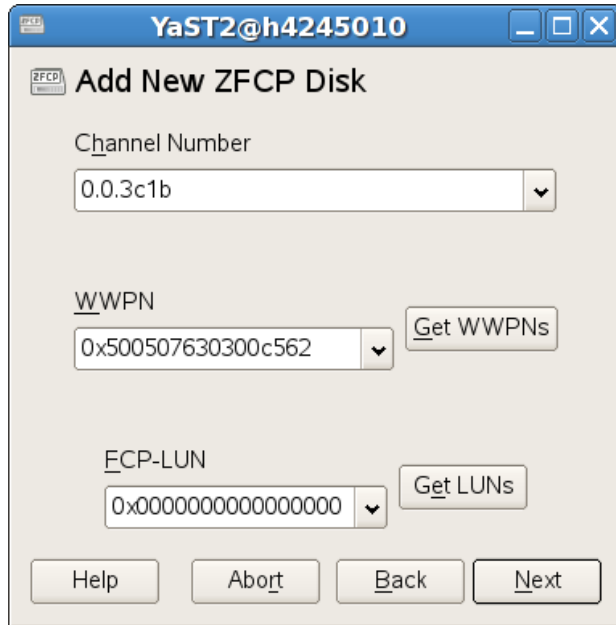
```
# lslns -c 0.0.1900 -p 0x5005076303000104
Scanning for LUNs on adapter 0.0.1900
  at port 0x5005076303000104:
    0x4021400000000000
    0x4021400100000000
    0x4021400200000000
    0x4021400300000000
    0x4021400400000000
    0x4021400500000000
    0x4021400600000000
    0x4021400700000000
    0x4021400800000000
    0x4021400900000000
```



## zfcf configuration files

- Configuring LUNs manually through sysfs is not persistent, will be lost after IPL
- Use configuration mechanism provided by distribution for persistent LUN configuration
- See documentation from distributions for details
- zfcf configuration file in RHEL
  - /etc/zfcf.conf
- zfcf configuration files in SLES 10
  - /etc/sysconfig/hardware/hwcfg-zfcf-bus-ccw-\*
- zfcf configuration files in SLES 11
  - /etc/udev/rules.d/51-zfcf\*

## SLES: GUI Setup



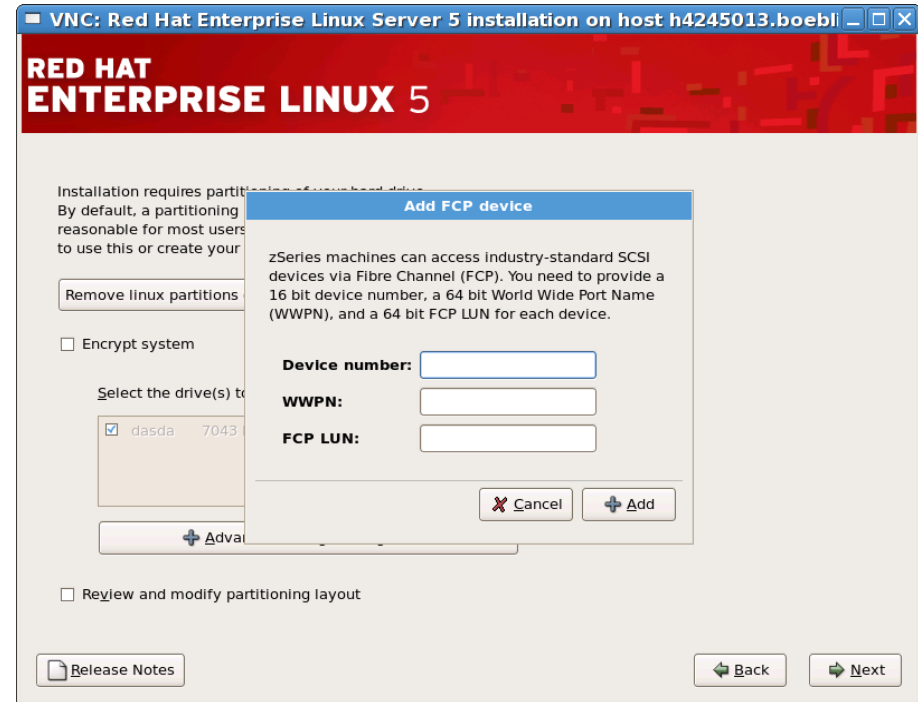
- zfcplib dialog in YaST simplifies setup of SAN attached devices
- Auto detects available FCP subchannels, WWPNs, and LUNs
- copy&paste WWPNs and FCP\_LUNs from configuration file obtained from SAN management tools or administrator

- alternatively on command line

- SLES 10: `/etc/sysconfig/hardware/hwcfg-zfcplib-bus-ccw-0.0.*`
- SLES 11: `zfcplib_{host|disk}_configure → /etc/udev/rules.d/51-zfcplib-0.0.*.rules`

## RHEL: GUI Setup

- Ignore subsequent complaints in case of DASD-less system.
- GUI only available during installation. Define FCP devices in */etc/zfcp.conf* for permanent addition.



```
# cat /etc/zfcp.conf
0.0.1900 0x5005076303000104 0x4021400000000000
# cat /etc/modprobe.conf (RHEL5)
alias scsi_hostadapter zfcp
# cat /etc/modprobe.d/local.conf (RHEL6)
alias scsi_hostadapter zfcp
# /sbin/zfcpconf.sh
```

## zfc: SCSI disk usage

```
# fdisk /dev/sda
```

```
Command (m for help): p
```

```
Disk /dev/sda: 5368 MB, 5368709120 bytes  
166 heads, 62 sectors/track, 1018 cylinders  
Units = cylinders of 10292 * 512 = 5269504 bytes
```

Device	Boot	Start	End	Blocks	Id	System
/dev/sda1		1	1018	5238597	83	Linux

```
# mke2fs -j /dev/sda1
```

## SCSI IPL

- The traditional initial program load (IPL) process relies on accessing a device using System z channel attachment
- For IPL from a FCP-attached device, this is not possible
- SCSI IPL expands the set of IPL'able devices
  - SCSI disks as Linux boot file system possible
- New set of IPL parameters
- Requires to address the SCSI disk
  - FCP adapter id
  - Remote port
  - LUN
- LPAR and z/VM guests supported
- SCSI (IPL) with z/VM
  - z/VM Version 4.4 (PTF UM30989) or newer
  - z/VM Version 5.3 (current version)

## SCSI IPL example LPAR

Load - H05:H05LP26	
CPC:	H05:H05LP26
Image:	H05:H05LP26
Load type	<input type="radio"/> Normal <input type="radio"/> Clear <input checked="" type="radio"/> SCSI <input type="radio"/> SCSI dump
<input type="checkbox"/> Store status	
Load address	* 5900
Load parameter	
Time-out value	60 <input type="button" value="↑"/> <input type="button" value="↓"/> 60 to 600 seconds
Worldwide port name	50050763030BC562
Logical unit number	4011400B00000000C
Boot program selector	0
Boot record logical block address	0
Operating system specific load parameters	
<input type="button" value="OK"/> <input type="button" value="Reset"/> <input type="button" value="Cancel"/> <input type="button" value="Help"/>	

## SCSI IPL example z/VM

```
set loaddev port 50050763 03000104 lun 40214000 00000000
Ready; T=0.01/0.01 22:11:01
```

In hexadecimal format with a blank separating the first 8 from the final 8 digits

WWPN

LUN

```
query loaddev
```

```
PORTNAME 50050763 03000104 LUN 40214000 00000000 BOOTPROG 0
BR_LBA 00000000 00000000
Ready; T=0.01/0.01 22:11:06
```

is the device number of the FCP subchannel that provides access to the SCSI boot disk.

```
i 1900
00: HCPLDI2816I Acquiring the machine loader from the processor controller.
00: HCPLDI2817I Load completed from the processor controller.
00: HCPLDI2817I Now starting the machine loader.
00: MLOEVL012I: Machine loader up and running (version 0.18).
00: MLOPDM003I: Machine loader finished, moving data to final storage location.
Linux version 2.6.16-18.x.20060403-s390xdefault (wirbser@t2944002) (gcc version
4.1.0) #1 SMP PREEMPT Mon Apr 3 09:56:54 CEST 2006
We are running under VM (64 bit mode)
Detected 4 CPU's
Boot cpu address 0
Built 1 zonelists
Kernel command line: dasd=e960-e962 root=/dev/sda1 ro noinitrd zfcp.device=0.0.1900,
0x5005076303000104,0x4021400000000000
```

## SCSI dump

- Dump memory of one LPAR to disk for problem analysis
- Similar to VMDUMP and dump to DASD
- SCSI dump supported for LPARs and as of z/VM 5.4
- Preparation summary:
  - large SCSI disk (system memory + 11 MB)
  - fdisk /dev/sda
  - mke2fs /dev/sda1
  - mount /dev/sda1 /mnt
  - zipl -D /dev/sda1 -t /mnt
  - umount /mnt



## SCSI dump from HMC

- Select CPC image for LPAR to dump
- Goto Load panel
- Issue SCSI dump
  - FCP device ID
  - WWPN
  - LUN

The screenshot shows the 'Load' panel in the HMC interface. The panel is titled 'Load' and contains the following configuration options:

- CPC: T63
- Image: T63LP22
- Load type:  Normal  Clear  SCSI  SCSI dump
- Store status
- Load address: \*4B49
- Load parameter: [ ]
- Time-out value: 60 [ ] 60 to 600 seconds
- Worldwide port name: 5005076305194786
- Logical unit number: 40FB400300000000
- Boot program selector: 0
- Boot record logical block address: 0
- Operating system specific load parameters: [ ]

At the bottom of the panel are four buttons: OK, Reset, Cancel, and Help.


## SCSI dump under z/VM

- SCSI dump from z/VM is supported as of z/VM 5.4
- Issue SCSI dump

```
#cp cpu all stop
#cp cpu 0 store status
#cp set dumpdev portname 47120763 00ce93a7 lun 40104020 00000000
bootprog 0
#cp ip1 4b49 dump
```

- To access the dump, mount the dump partition

## FCP Multipathing



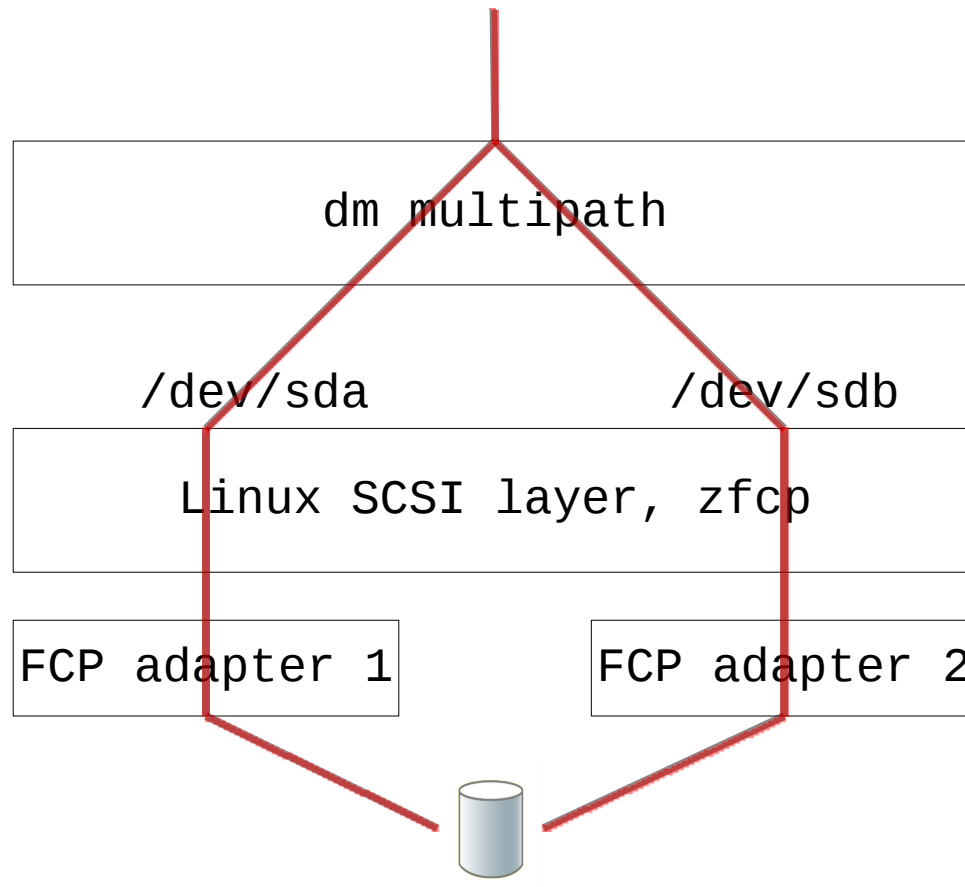
2 paths to disk through independent FCP adapters and independent controllers.

## Multipathing for disks

- Use multiple paths from operating system to storage
- Why multipathing ?
  - Avoid single points of failure
  - Performance: I/O requests can be spread across multiple paths,
  - Usually a requirement for SAN connected storage servers
    - e.g. when one storage controller is in maintenance mode I/O continues to run through second controller
- Multipathing in Linux
  - Implemented in Linux in multipath-tools package, together with the device-mapper in the Linux kernel
  - The default configuration is already included, configuration changes only for special requirements
  - SCSI device (“LUN”) in Linux represents one path to the disk volume on the storage server
- Multipath devices are block devices in Linux, can be used e.g.
  - Partitions
  - Directly for a filesystem
  - Logical Volume Manager (LVM)

## Multipathing for disks – linux device mapper

The device mapper creates one block device for the LUN /dev/mapper/xxx  
 /dev/mapper/36005076303ffc56200000000000010cc



unique WWID

(World-Wide Identifier) from storage server identifies volume



## zfcpx setup for multipathing

- zfcpx and SCSI report each path as device
- multipathing happens on higher layer
- avoid shared components in different paths

```
# lsscsi
[0:0:0:0]    disk    IBM      2107900    2.27    /dev/sda
[1:0:1:0]    disk    IBM      2107900    2.27    /dev/sdb
# lszfcp -D
0.0.3c00/0x500507630313c562/0x401040cc00000000 0:0:0:0
0.0.3d00/0x500507630303c562/0x401040cc00000000 1:0:1:0
```

different adapters and  
different ports to avoid  
single points of failures

usually same  
FCP LUN (check  
on storage server)

## Multipathing setup and configuration

- Config file */etc/multipath.conf*: Default settings are usually good, but can be overwritten
- Ensure devices are not blacklisted and multipathd is running

```
# cat /etc/multipath.conf
blacklist {
    devnode "^(ram|raw|loop|fd|md|dm-|sr|scd|st)[0-9]*"
    devnode "^hd[a-z][[0-9]*]"
    devnode "^cciss!c[0-9]d[0-9]*[p[0-9]*]"
    devnode "^dasd[a-z]+[0-9]*"
}
```

- Multipath devices are created automatically when SCSI LUNs are attached

WWID for  
volume

pathgroup

```
# multipath -ll
36005076303ffc562000000000000010cc dm-0 IBM,2107900
[size=5.0G][features=1 queue_if_no_path][hw_handler=0]
\_ round-robin 0 [prio=2][active]
\_ 1:0:0:0 sdb 8:16 [active][ready]
\_ 0:0:0:0 sda 8:0 [active][ready]
...
```

## Multipathing setup and configuration (cont'd)

- Device to work with: `/dev/mapper/36005076303ffc562000000000000010cc`

```
# mkfs.ext3 /dev/mapper/36005076303ffc562000000000000010cc
...
# mount /dev/mapper/36005076303ffc562000000000000010cc /mnt
```

- Example: Change `rr_min_io` to 10 - (number of IO to route to a path before switching to the next in the same path group)
- See *multipath.conf* manpage and documentation from distributions for details

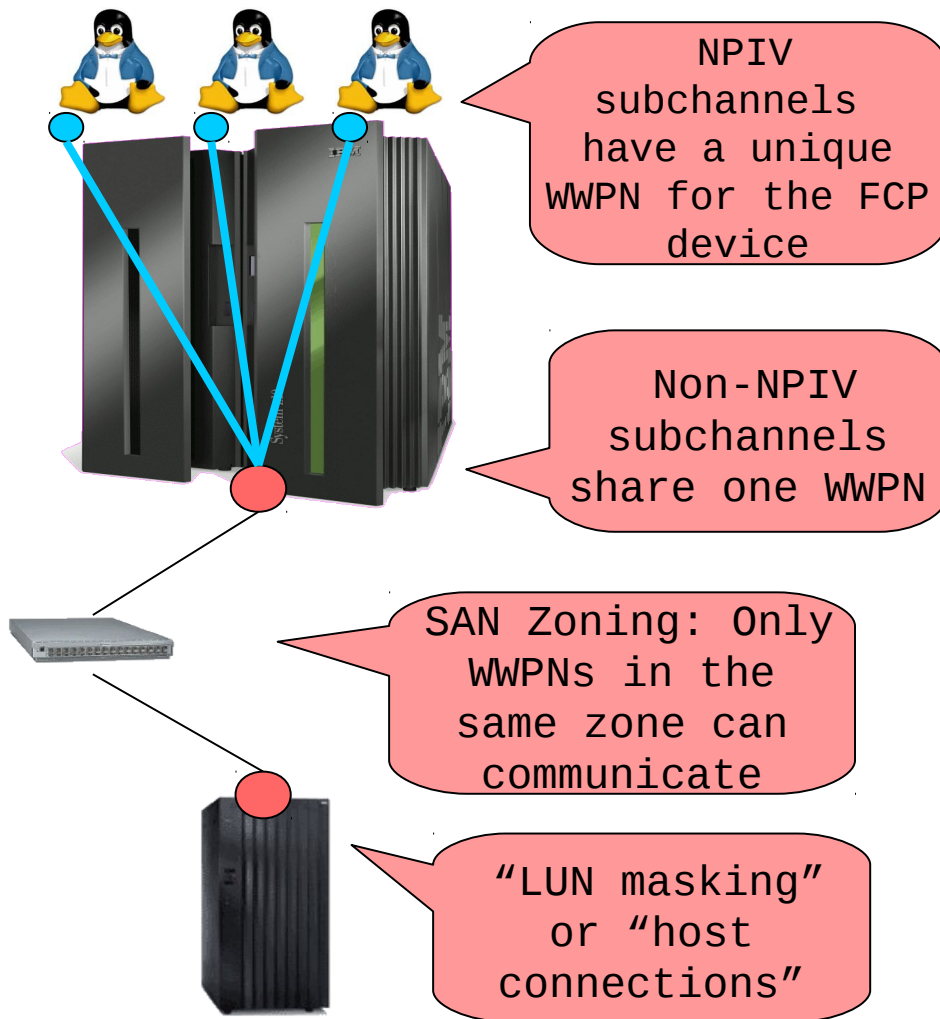
```
# cat /etc/multipath.conf
multipaths {
    multipath {
        wwid "36005076303ffc562000000000000010cc"
        rr_min_io "10"
        path_selector "round-robin 0"
    }
}
```



## Root filesystem on SCSI multipath volume

- Possibility to put root filesystem on SCSI multipath volume
- Older zipl versions do not support writing on multipath device
- Workaround: Use additional single path volume for /boot, use this volume for IPL
- Newer Linux distributions (e.g. SLES 11 SP1, RHEL 6.2) support root filesystem on multipath device without workaround with s390-tools 1.8.3 or newer
- <http://www.ibm.com/developerworks/linux/linux390/s390-tools-1.8.3.html>
  - zipl: Add support for device mapper devices.
  - zipl now allows installation of and booting from a boot record on logical devices, i.e. devices managed by device mapper (or similar packages), e.g. multipath devices.

## N-Port Id Virtualization (NPIV)



- NPIV: Each virtual adapter has its own WWPN in the SAN. This is the foundation for restricting access to storage.
- SAN zoning: Only WWPNs in same zone can communicate
- "LUN masking" or "host connections": Each disk volume is only available for a specific list of WWPNs
- With NPIV: Restricted configuration possible where disk volumes are defined for each operating system
- Without NPIV: First operating system to access a disk volume can use it.

## NPIV requirements



- NPIV is available on System z9, System z10 and z/Enterprise
  - FICON Express 2 adapter running with MCL003 on EC J99658
- z/VM
  - z/VM 5.2 or 5.3
  - z/VM 5.1 with the PTF for APAR VM63744
- Linux Distribution
  - Currently SLES10, SLES11, RHEL5, RHEL6 (LPAR mode or z/VM)
- NPIV-Capable Switch
  - only required for switch adjacent to System z
  - Mostly firmware upgrades possible (e.g. McData, Brocade)

## NPIV

```
# lszfc -a | grep port_name  
    permanent_port_name = "0xc05076ffe5005611"  
    port_name           = "0xc05076ffe5005350"
```

- “port\_name” is the WWPN used by the FCP subchannel
- “permanent\_port\_name” is the WWPN assigned to the FCP channel
- Compare to find out if NPIV is in use:
  - If both are the same, the FCP subchannel does NOT use NPIV
  - If they differ, the FCP subchannel uses NPIV
- Worldwide portname prediction tool
- <http://www.ibm.com/servers/resourceink/>

## Troubleshooting

- Check kernel messages that are possibly related to SCSI on Linux on System z:
  - scsi (common SCSI code)
  - sd (SCSI disk)
  - rport (common SCSI code FC remote port messages)
  - qdio (communication between Linux and FCP Channel)
  - zfcplib driver kernel messages
  - See “Kernel Messages” book on <https://www.ibm.com/developerworks/linux/linux390/>
  - “device-mapper: multipath”
- Other syslog messages
  - Multipathd
- zfcplib driver traces available in `/sys/kernel/debug/s390dbf/`
- Collect data with `dbginfo.sh` when reporting a problem to capture configuration, messages and traces

## Troubleshooting: scsi\_logging\_level

```
# scsi_logging_level -g  
Current scsi logging level:  
dev.scsi.logging_level = 0  
SCSI_LOG_ERROR=0  
SCSI_LOG_TIMEOUT=0  
SCSI_LOG_SCAN=0  
SCSI_LOG_MLQUEUE=0  
SCSI_LOG_MLCOMPLETE=0  
SCSI_LOG_LLQUEUE=0  
SCSI_LOG_LLCOMPLETE=0  
SCSI_LOG_HLQUEUE=0  
SCSI_LOG_HLCOMPLETE=0  
SCSI_LOG_IOCTL=0
```

- More SCSI output in kernel messages
- Higher levels can create lots of messages and slow down system

```
# scsi_logging_level -s -a 1  
New scsi logging level:  
dev.scsi.logging_level = 153391689  
SCSI_LOG_ERROR=1  
SCSI_LOG_TIMEOUT=1  
SCSI_LOG_SCAN=1  
SCSI_LOG_MLQUEUE=1  
SCSI_LOG_MLCOMPLETE=1  
SCSI_LOG_LLQUEUE=1  
SCSI_LOG_LLCOMPLETE=1  
SCSI_LOG_HLQUEUE=1  
SCSI_LOG_HLCOMPLETE=1  
SCSI_LOG_IOCTL=1
```

## Troubleshooting: zfcplib

- Query Fibre Channel nameserver about ports available for my system:

```
# zfcplib -n
Local Port List:
    0x500507630313c562 / 0x656000 [N_Port] proto = SCSI-FCP   FICON
    0x50050764012241e4 / 0x656100 [N_Port] proto = SCSI-FCP
    0x5005076401221b97 / 0x656400 [N_Port] proto = SCSI-FCP
```

- Query SAN topology, requires FC management server access:

```
# zfcplib
Interconnect Element Name      0x100000051e4f7c00
Interconnect Element Domain ID 005
Interconnect Element Type      Switch
Interconnect Element Ports     256
  ICE Port 000  Online
    Attached Port [WWPN/ID] 0x50050763030b0562 / 0x650000 [N_Port]
  ICE Port 001  Online
    Attached Port [WWPN/ID] 0x50050764012241e5 / 0x650100 [N_Port]
  ICE Port 002  Online
    Attached Port [WWPN/ID] 0x5005076303008562 / 0x650200 [N_Port]
  ICE Port 003  Offline
...

```

## Troubleshooting: zfcplib

- Check if remote port responds (requires FC management service access):

```
# zfcplib ping 0x5005076303104562
Sending PNG from BUS_ID=0.0.3c00 speed=8 GBit/s
  echo received from WPN (0x5005076303104562) tok=0 time=1.905 ms
  echo received from WPN (0x5005076303104562) tok=1 time=2.447 ms
  echo received from WPN (0x5005076303104562) tok=2 time=2.394 ms

----- ping statistics -----
min/avg/max = 1.905/2.249/2.447 ms
-----
```

- zfcplib\_show and zfcplib\_ping are part of the zfcplib-hbaapi 2.1 package:

<http://www.ibm.com/developerworks/linux/linux390/zfcplib-hbaapi-2.1.html>



## Summary of FCP

- available for zSeries and System z
- based on existing Fibre Channel infrastructure
- runs on all available z/VM and RHEL/SLES versions
- integrates System z into standard SANs
- connects to switched fabric or point-to-point
- multipathing for SCSI disks is a must
- gives you new storage device choices
- usually performs better than FICON
- buys you flexibility at the cost of complexity
- tooling available, receiving better integration

## More Information

I/O Connectivity on IBM zSeries mainframe servers

[www.ibm.com/systems/z/connectivity/](http://www.ibm.com/systems/z/connectivity/)

Linux on zSeries: Fibre Channel Protocol Implementation Guide

[www.redbooks.ibm.com/redpapers/pdfs/redp0205.pdf](http://www.redbooks.ibm.com/redpapers/pdfs/redp0205.pdf)

How to use FC-attached SCSI devices with Linux on System z

[download.boulder.ibm.com/ibmdl/pub/software/dw/linux390/docu/l26cts00.pdf](http://download.boulder.ibm.com/ibmdl/pub/software/dw/linux390/docu/l26cts00.pdf)

Linux for IBM System z Device Drivers Book and other documentation

[http://www.ibm.com/developerworks/linux/linux390/documentation\\_dev.html](http://www.ibm.com/developerworks/linux/linux390/documentation_dev.html)

SLES 11 SP2: Storage Administration Guide

[http://www.novell.com/documentation/sles11/stor\\_admin/?page=/documentation/sles11/stor\\_admin/data/bookinfo.html](http://www.novell.com/documentation/sles11/stor_admin/?page=/documentation/sles11/stor_admin/data/bookinfo.html)

Red Hat Enterprise Linux 6: Storage Administration Guide

[http://docs.redhat.com/docs/en-US/Red\\_Hat\\_Enterprise\\_Linux/6/html/Storage\\_Adminis](http://docs.redhat.com/docs/en-US/Red_Hat_Enterprise_Linux/6/html/Storage_Adminis)

# Questions?



***Susanne Wintenberger***

*Certified IT Specialist  
Linux on System z*

*Schönaicher Strasse 220  
71032 Böblingen, Germany*

*Phone +49 (0)7031-16-3514  
swinten@de.ibm.com*