



Linux on zSeries and SCSI


Linux on zSeries and SCSI Implementation

Klaus Bergmann

IBM Lab Böblingen, Germany


Agenda

- **zSeries Hardware**
 - zSeries in a SAN
- **zSeries Software**
 - Linux SCSI/FCP Support
 - Multi-Pathing
- **Storage Devices**
 - Disk, Tape
- **SCSI IPL**



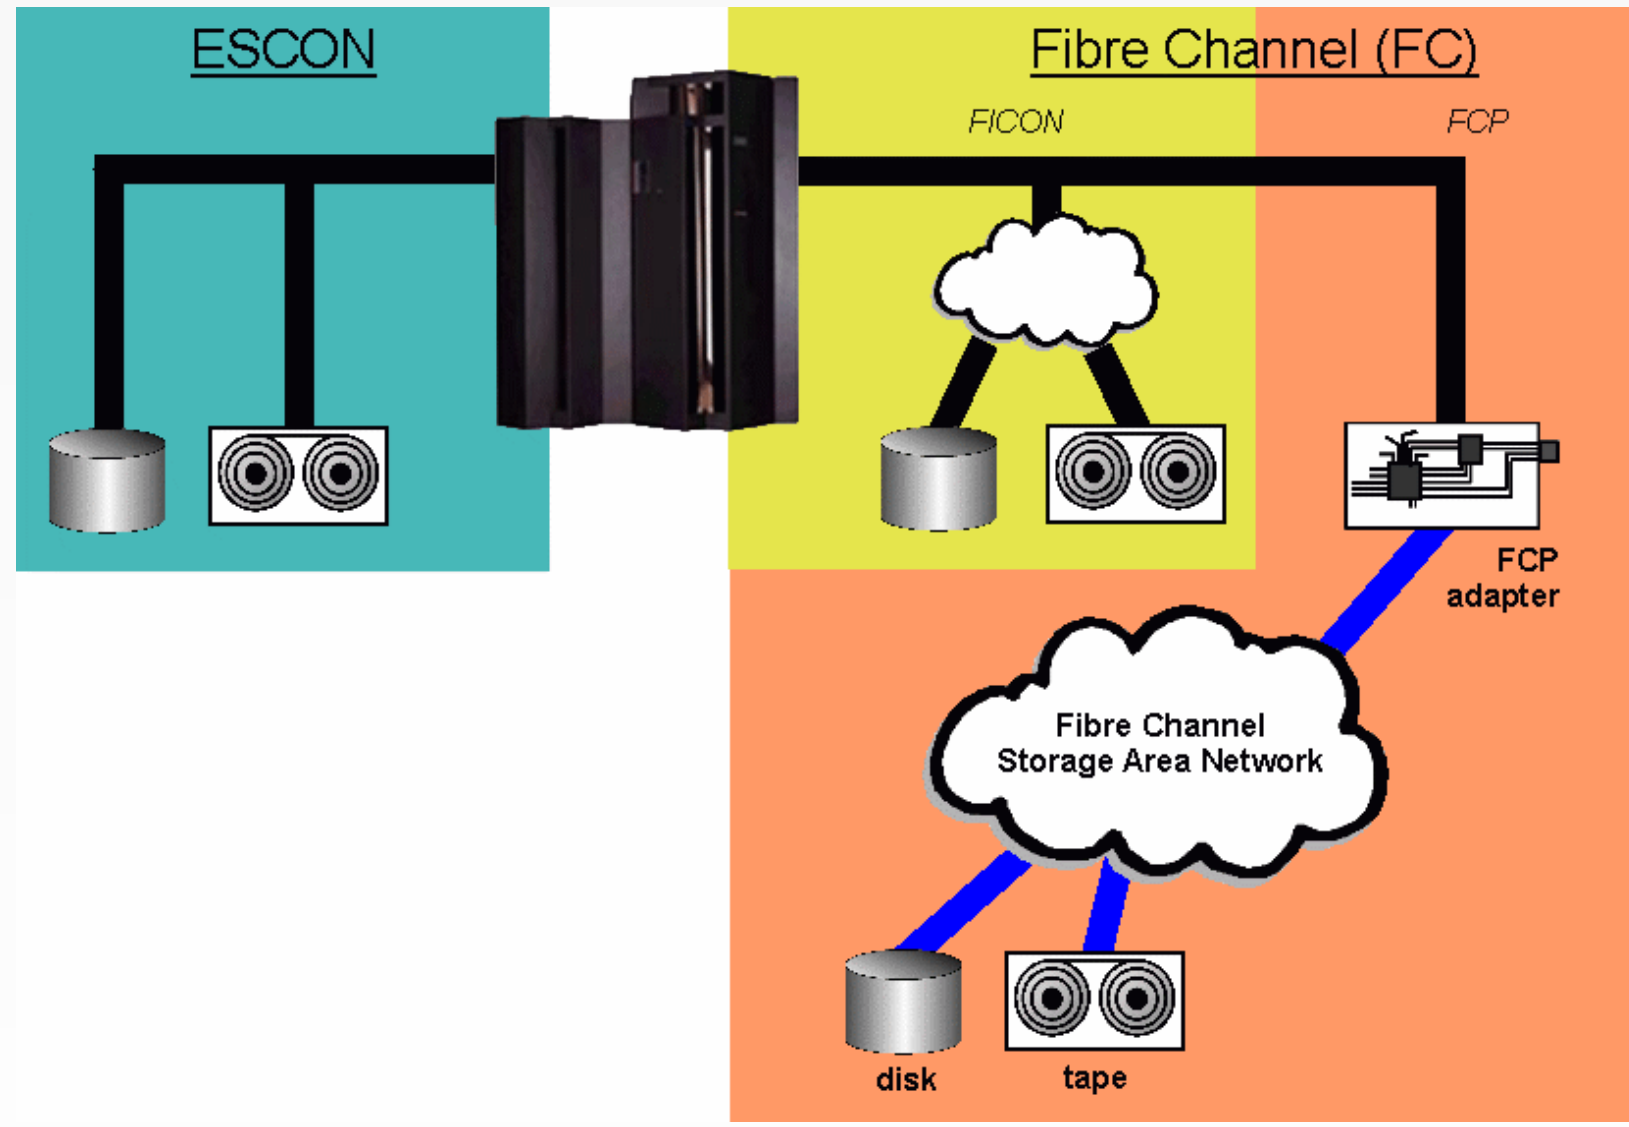


zSeries I/O – ESCON & FICON, Legacy Storage Attachments



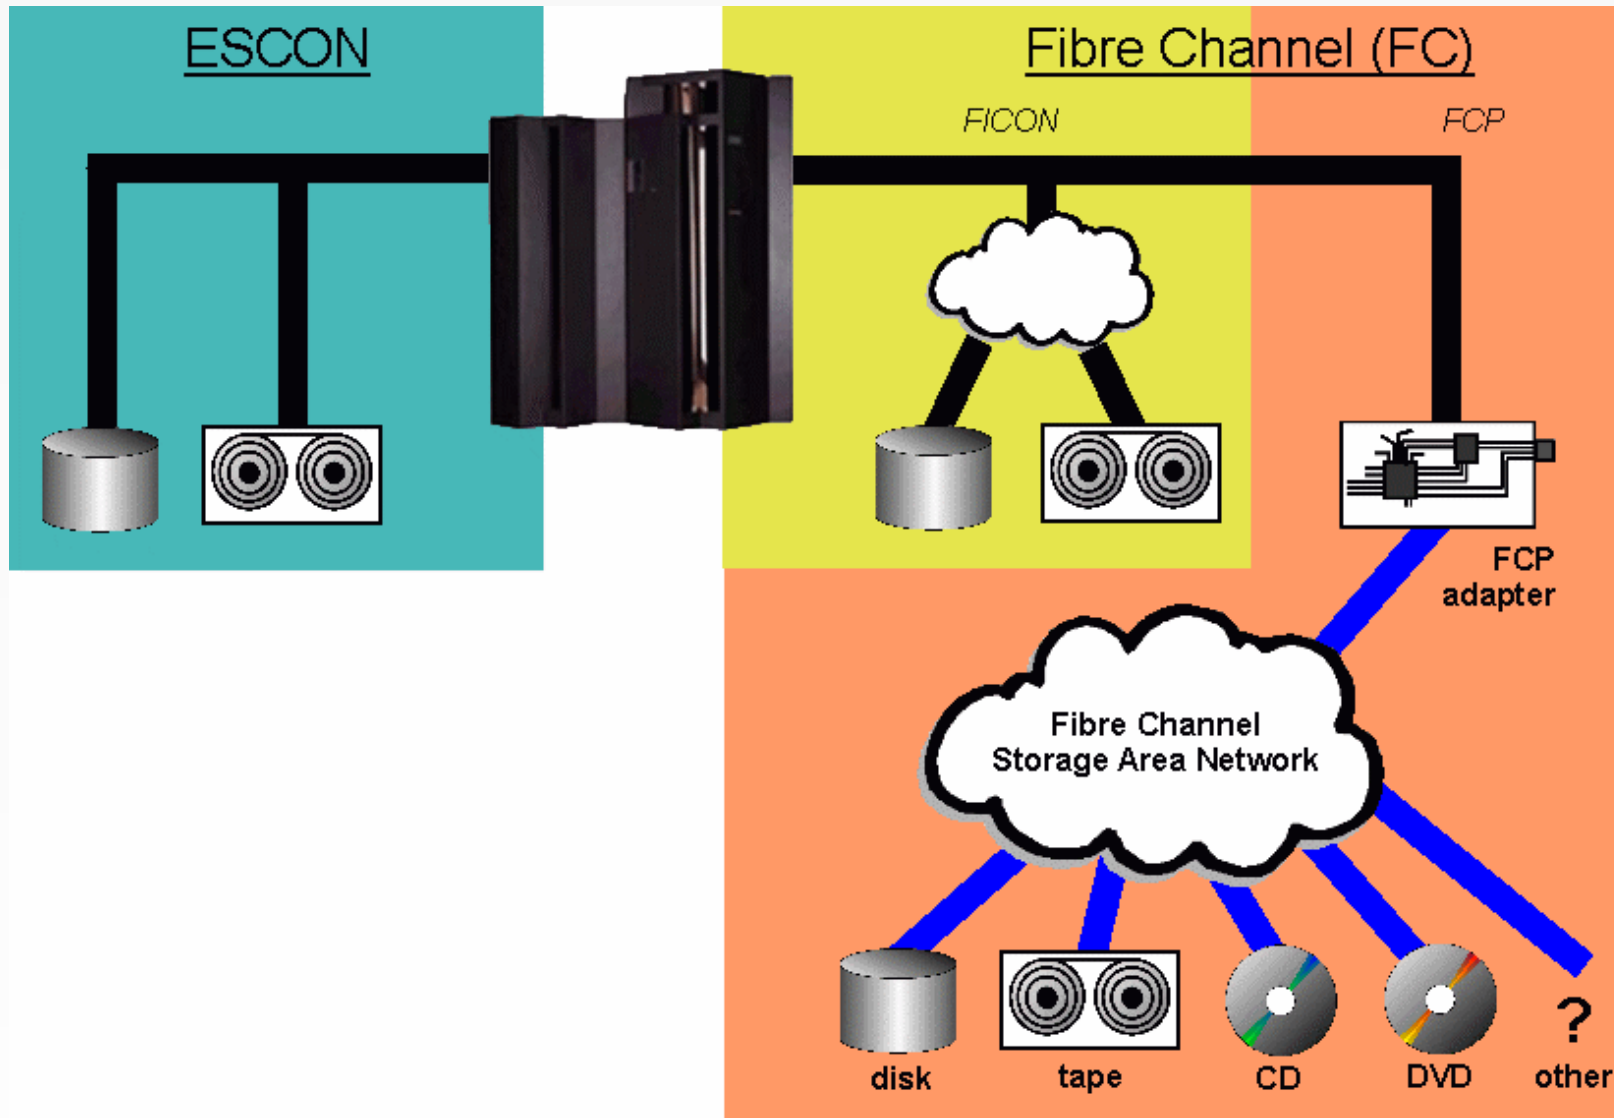


zSeries in a SAN – FCP, Connecting to the Open Storage World



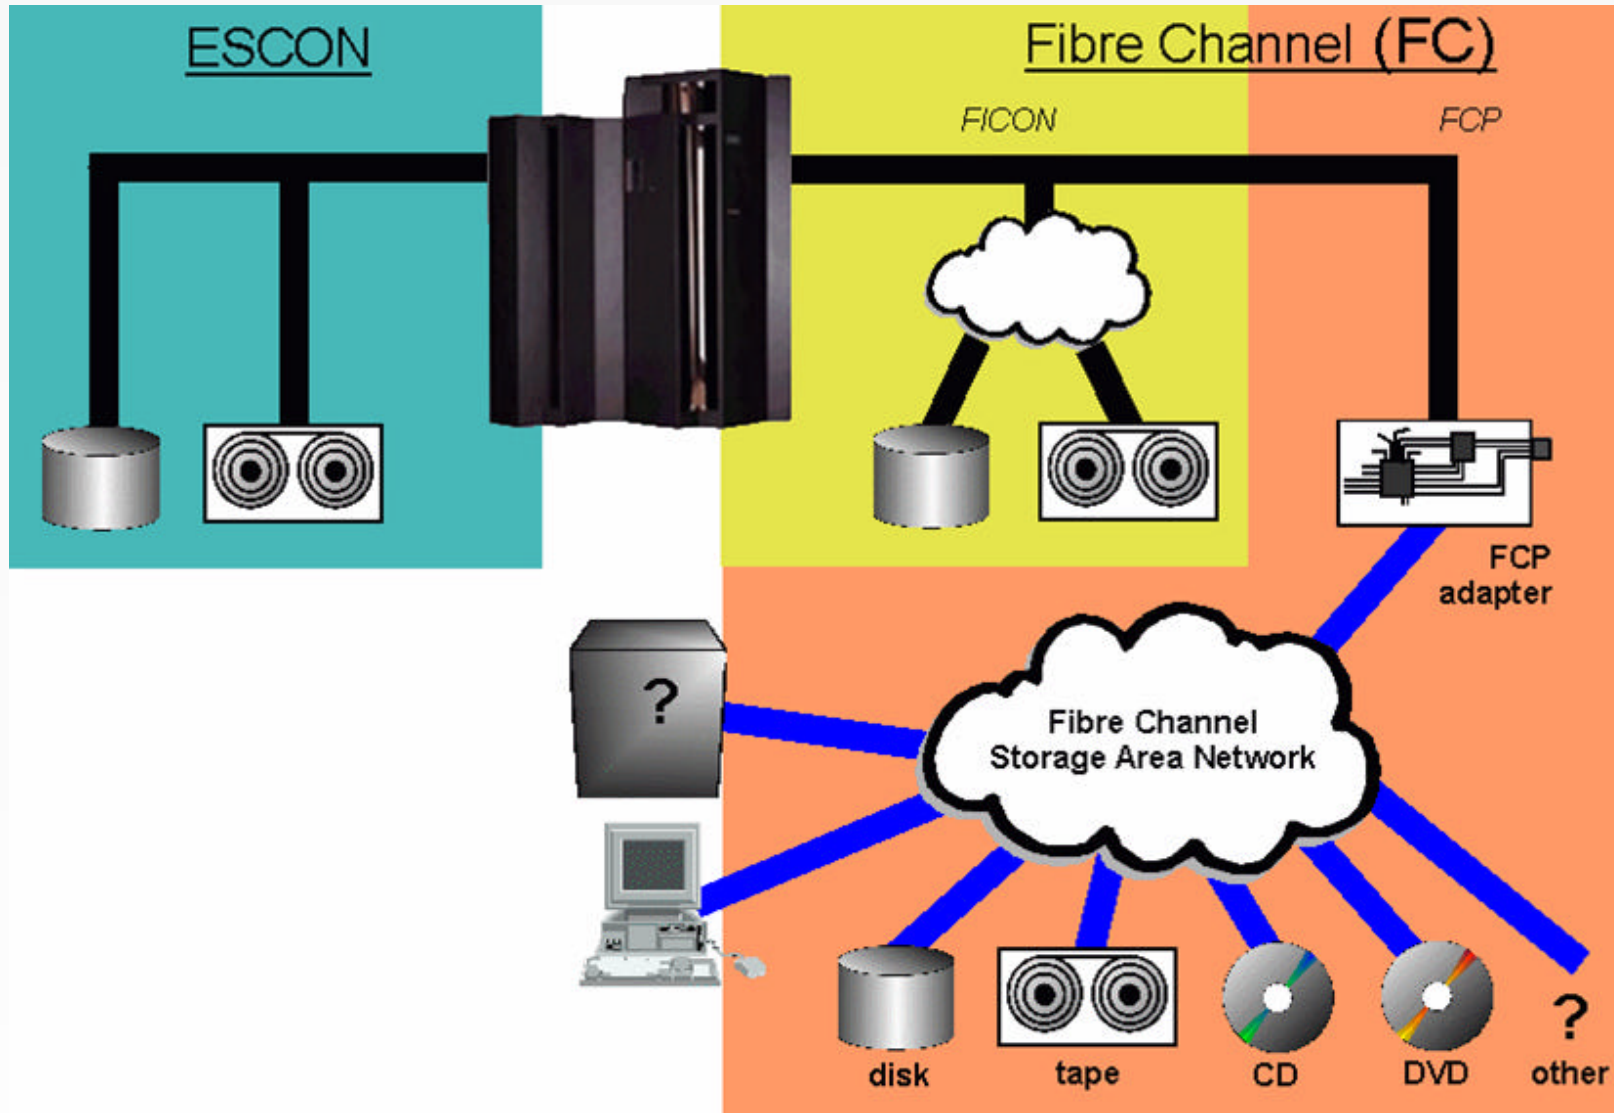


zSeries in a SAN – Additional Device Types





zSeries in a SAN – Sharing Your Storage Resources






zSeries in a SAN - Hardware Requirements

- **IBM zSeries 800, 900 or 990**
- **FICON or FICON Express adapter card**
- **Additional CHPID type FCP**
- **FC fabric switch**
- **FC attached storage devices**
- **Optional: FC<->SCSI bridge + SCSI devices**







zSeries in a SAN – Topologies



**point-to-point
not supported**



**direct-attached arbitrated loop
not supported**



**switched fabric
supported**



Fibre Channel and SCSI – Software Requirements

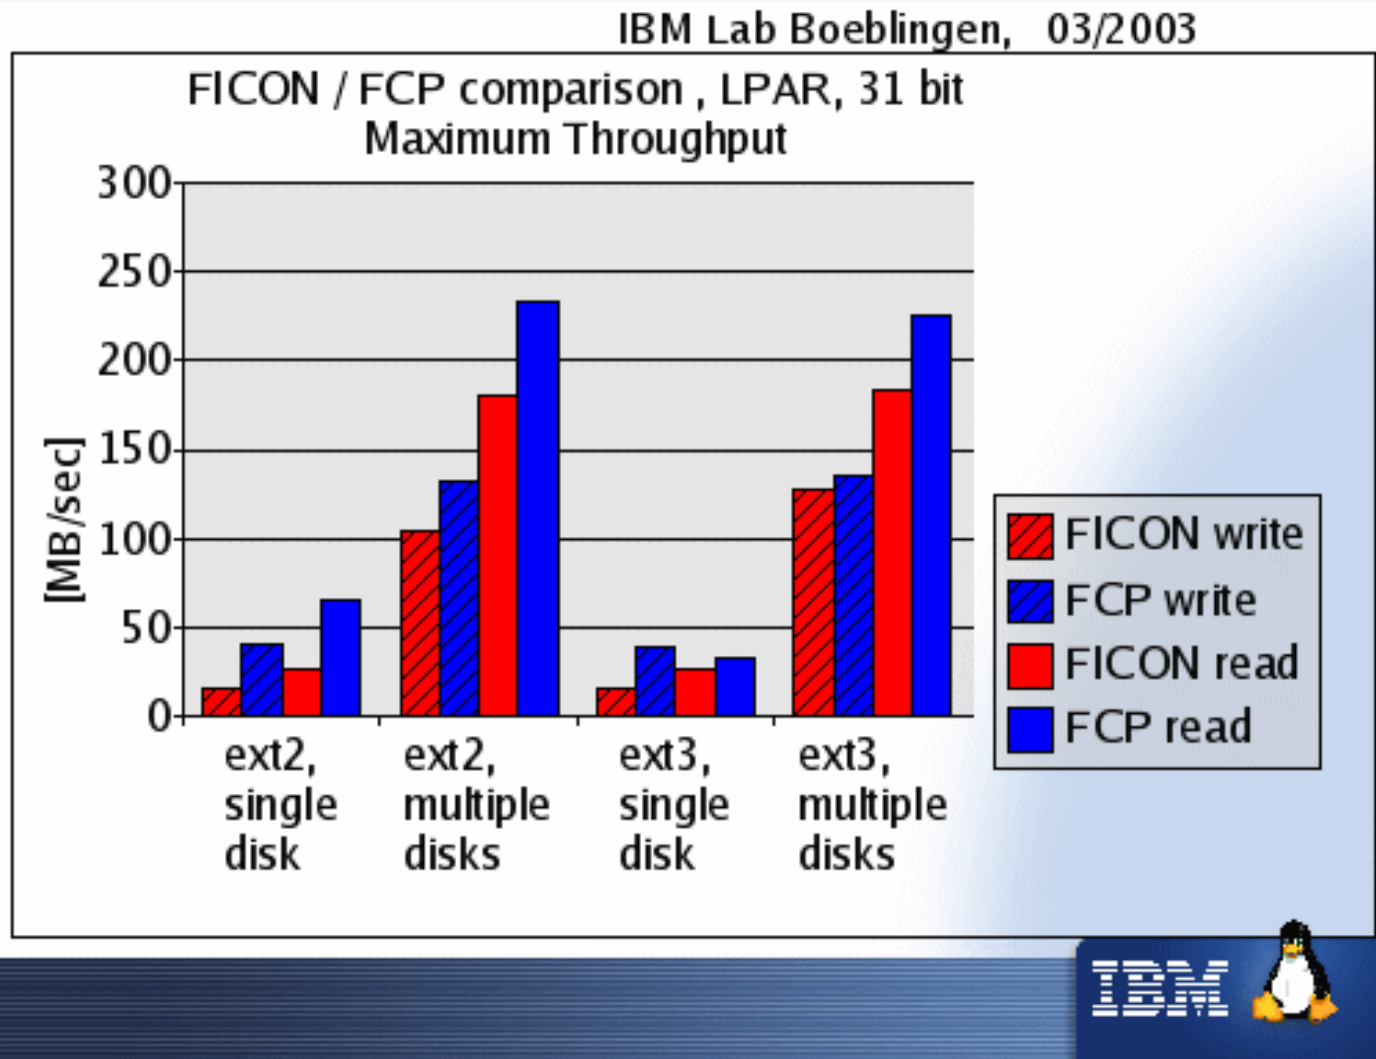
- **SUSE Linux Enterprise Server 8 (SLES8)**
 - GA November 2002
 - Currently SP3
- **Red Hat Enterprise Linux 3 (RHEL3)**
 - GA October 2003
 - Update 1 or higher required
- **z/VM 4.3**
 - GA May 2002
 - Includes FCP channel guest support for Linux
 - Currently z/VM 4.4



Disk Usage – ECKD and SCSI Comparison


	ECKD DASD	FCP/SCSI disk
Device configuration / access control	IOCDs & z/VM Operator	I/O mapping (Linux guest)
Low level disk format	dasdfmt	n/a
Partitioning	fdasd	fdisk
Create filesystem	mke2fs (or other)	mke2fs (or other)
Access filesystem	mount	mount
Access files	application	application

FICON and FCP performance comparison



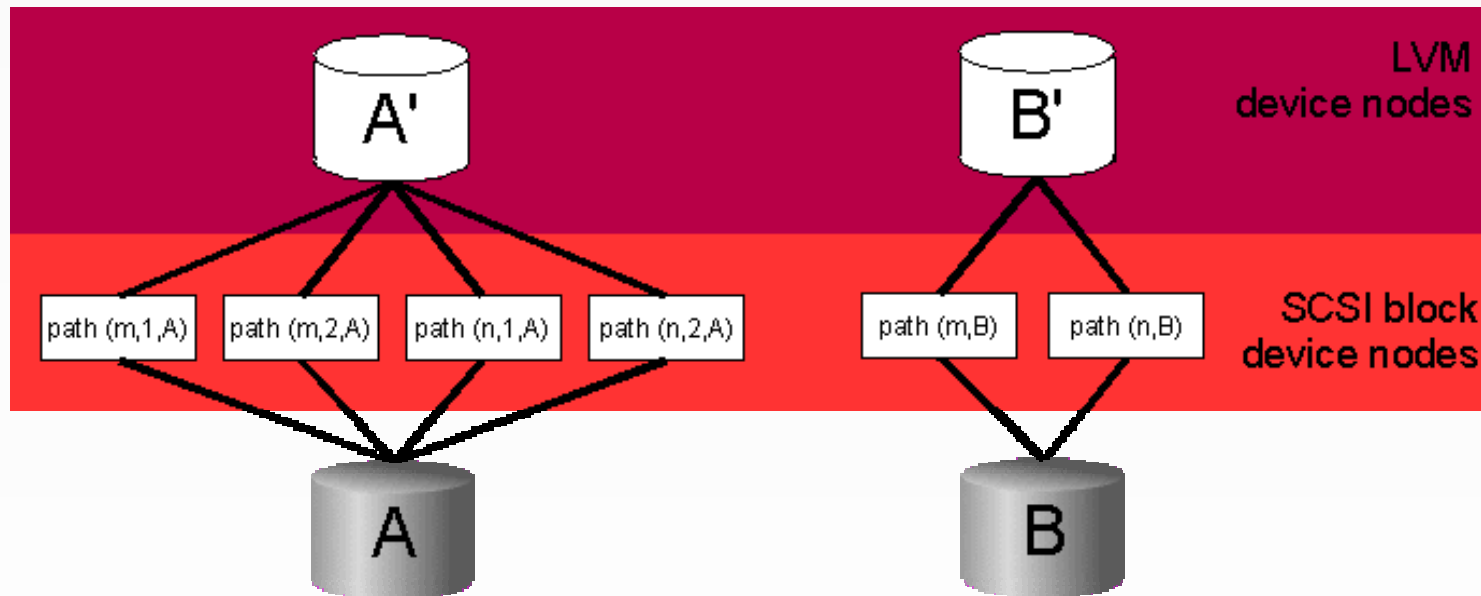



Multipathing with the Logical Volume Manager (LVM)





Multipathing with the Logical Volume Manager (LVM)





Multipathing with LVM

- **Enterprise-class multipathing solution for high-availability of disk storage**
 - Failover on path-failure
 - Failback if recovered path is detected (retries)
 - Load balancing (use of multiple paths for concurrent I/Os according to assigned priorities)
 - Designed to cover all block devices
- **Separate kernel patch**
 - Included in SLES 8






FCP/SCSI Tape Support

■ Devices:

- IBM TotalStorage Enterprise Tape System 3590
- IBM TotalStorage Enterprise Tape Drive 3592
- IBM TotalStorage Enterprise Tape Library 3494
- IBM TotalStorage UltraScalable Tape Library 3584 w/ Ultrium 2

■ **IBMtape.1.3.2 and IBMtapeutil.1.1.6 packages required**

- `/lib/modules/(Your system's kernel name)/kernel/drivers/scsi/IBMtape.o`
- `/usr/bin/IBMtapeconfig`
- `/usr/bin/IBMtaped`
- `/usr/bin/IBMtapeutil`





Device Support - Summary

- **FCP Connectivity Options supported:**
 - IBM TotalStorage Enterprise Tape System 3590
 - IBM TotalStorage Enterprise Tape Drive 3592
 - IBM TotalStorage Enterprise Tape Library 3494
 - IBM TotalStorage Enterprise Storage Server Models 800, F20, F10
 - IBM TotalStorage UltraScalable Tape Library 3584
 - CNT (INRANGE) FC/9000 64-port and 128-port models (IBM 2042)
 - McDATA Intrepid 6064 (IBM 2032)
 - McDATA Intrepid 6140 (IBM 2032)
 - McDATA Sphereon 4500 Fabric Switch (IBM 2031-224)
 - IBM TotalStorage SAN Switches 2109-F16 and S16/S08
 - McDATA ES-1000 Loop Switch (IBM 2031-L00)
- **FCP Connectivity Options currently under interoperability test:**
 - McDATA ED-5000 (IBM 2032-001)
 - IBM 2108-G07 SAN Data Gateway (parallel SCSI connectivity to non-IBM storage)



SCSI IPL & SCSI Dump

- **SCSI IPL from FCP attached SCSI disks**
- **SCSI Dump to FCP attached SCSI disk (LPAR only)**
- **Expand the world of open I/O attachments on zSeries from pure data access to allow IPL and Dump support**
- **Enhances the setup to allow Linux on zSeries to run completely on SCSI disks - incl. IPL, Data access and Dump support**





SCSI IPL & SCSI Dump – Cont.

Load

CPC: P000F12B

Image: ZFCP4

Load type: Normal Clear SCSI SCSI dump

Store status

Load address: 5C00

Load parameter:

Time-out value: 060 (60 to 600 seconds)

World wide port name: 5005076300CE93A7

Logical unit number: 5732000000000000

Boot program selector: 0

Boot record logical block address: 0000000000000000

OS specific load parameters:

OK Reset Cancel Help

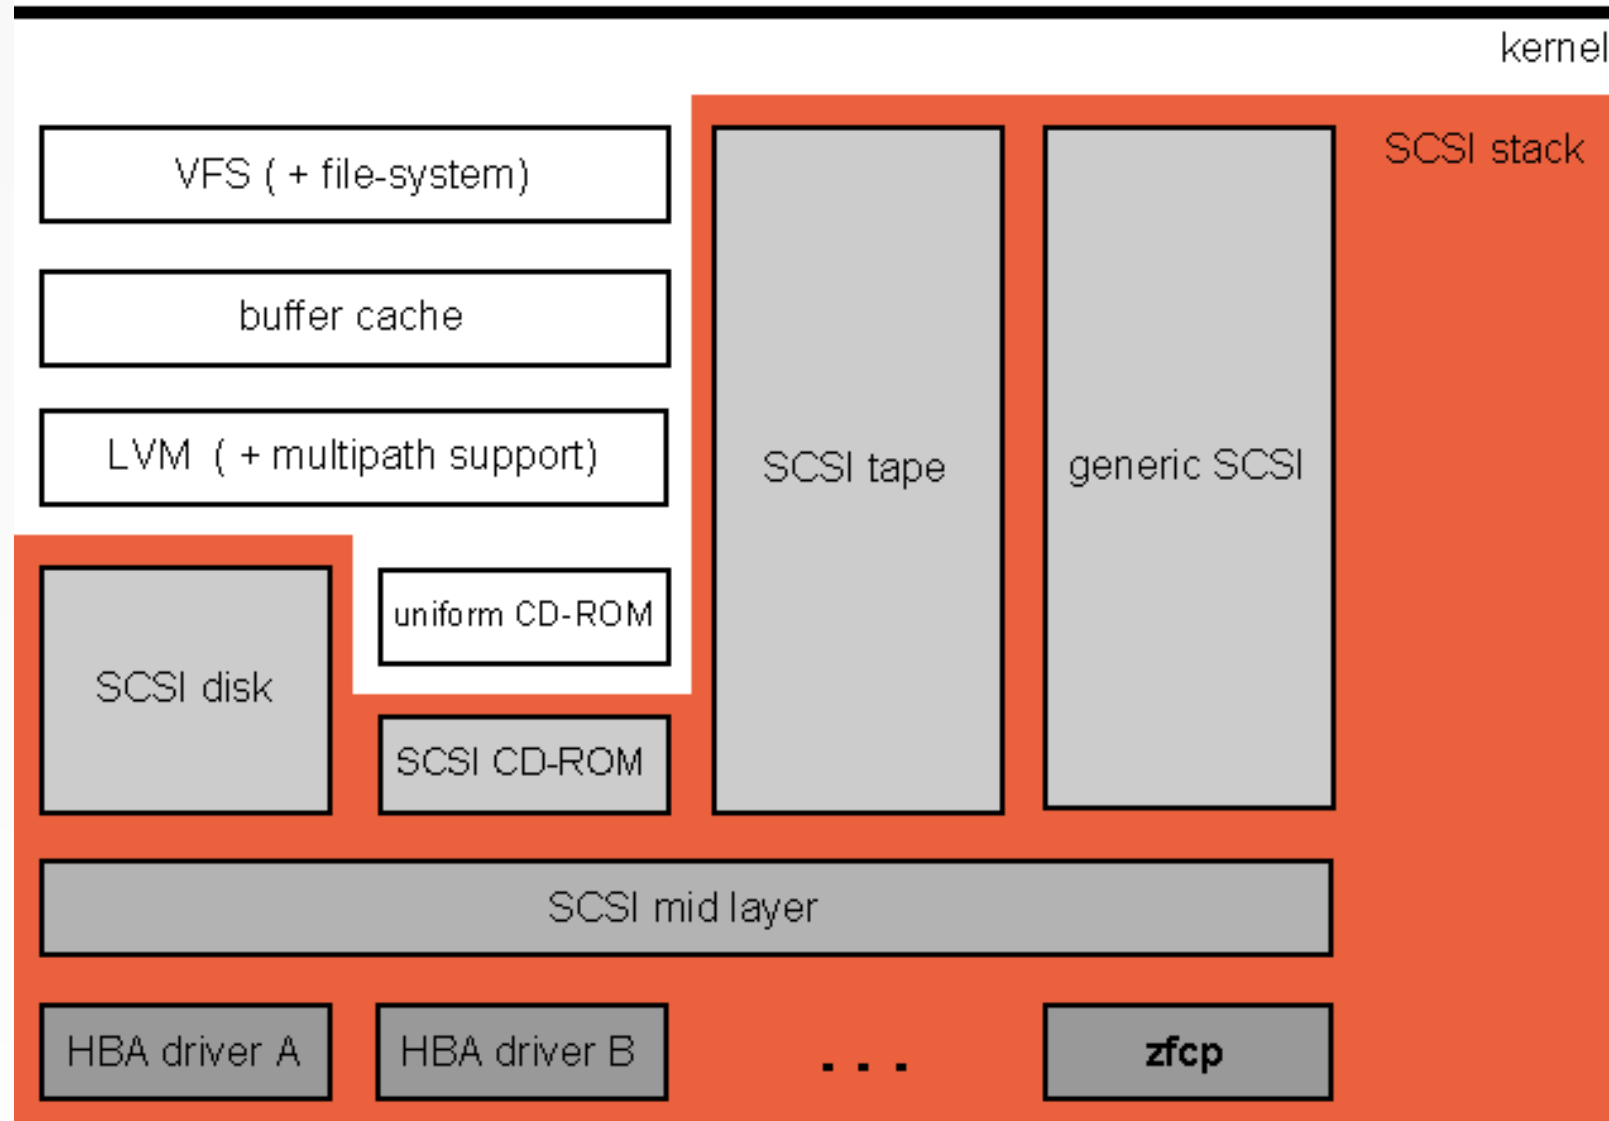
- New set of IPL parameters
- LPAR and z/VM guests supported
- Disk preparation with Linux „zipl“ tool
- Up to 31 boot configurations possible

■ Requirements

- Requires enablement by FC9904
- Requires FCP channels
- IBM zSeries server 800, 900 and 990
- z/VM 4.4 (PTF UM30989)

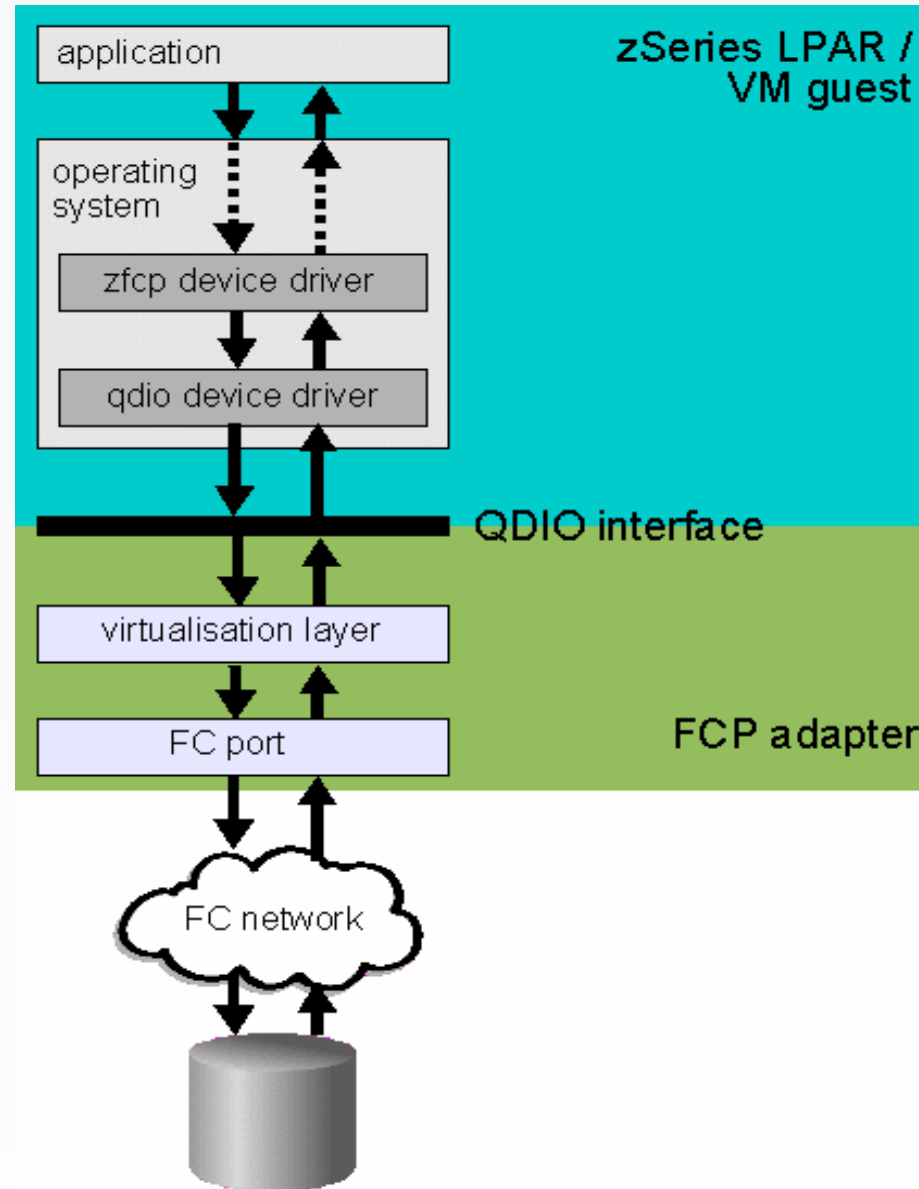


Linux SCSI Stack





Linux on zSeries SCSI Traffic



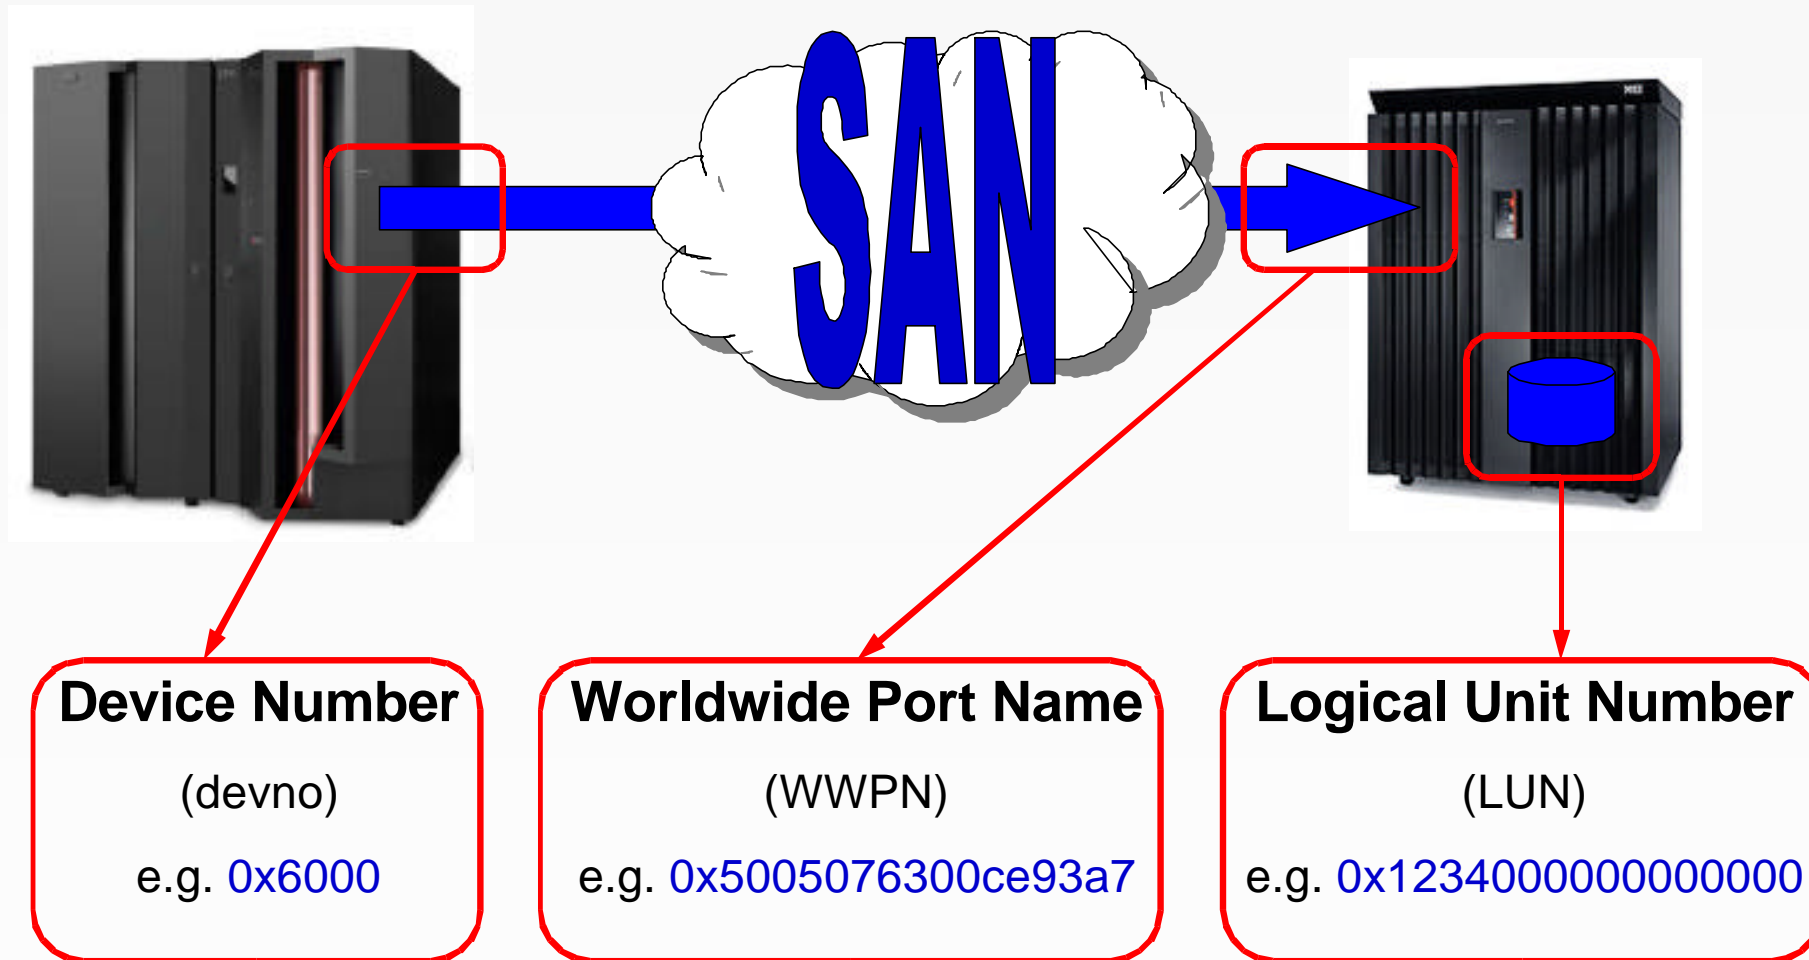


zfcplib's Task in the Linux SCSI Stack

- **zfcplib drives the zSeries FCP host bus adapter**
 - sends SCSI commands and associated data on behalf of the Linux SCSI subsystem to SCSI devices attached via a zSeries FCP adapter
 - returns replies and data from SCSI devices to the Linux SCSI subsystem
 - maintains connections through the SAN to SCSI devices attached via a zSeries FCP adapter
 - maps SAN devices to SCSI devices as seen by the Linux SCSI subsystem



SAN Addressing – How to Find the Path to an FCP Device





FCP Device Mapping – Configuration Example

```
bash# cat /proc/subchannels
```

```
Device sch.  Dev Type/Model CU  in use  PIM PAM POM CHPIDs
-----
6000  0011  1732/03  1731/03                80  80  FF  2A000000 00000000
```

```
bash# cat /proc/scsi/zfcp/map
```

```
0x6000 0x00000001:0x5005076300c38550 0x00000000:0x5249000000000000
```





FCP Device Mapping – Configuration Example

```
bash# cat /proc/subchannels
```

Device	sch.	Dev	Type/Model	CU	in use	PIM	PAM	POM	CHPIDs
6000	0011	1732/03	1731/03			80	80	FF	2A000000 00000000

device number



0x6000 0x00000001:0x5005076300c38550 0x00000000:0x5249000000000000





FCP Device Mapping – Configuration Example

```
bash# cat /proc/subchannels
```

Device	sch.	Dev	Type/Model	CU	in use	PIM	PAM	POM	CHPIDs
6000	0011	1732/03	1731/03			80	80	FF	2A000000 00000000

device number

WWPN



0x6000 0x00000001:0x5005076300c38550 0x00000000:0x5249000000000000





FCP Device Mapping – Configuration Example

```
bash# cat /proc/subchannels
```

```
Device sch.  Dev Type/Model CU  in use  PIM PAM POM CHPIDs
-----
6000  0011  1732/03  1731/03          80  80  FF  2A000000 00000000
```

device number

WWPN

FCP LUN



0x6000 0x00000001:0x5005076300c38550 0x00000000:0x5249000000000000





FCP Device Mapping – Configuration Example

```
bash# cat /proc/subchannels
```

Device	sch.	Dev	Type/Model	CU	in use	PIM	PAM	POM	CHPIDs
6000	0011	1732/03	1731/03			80	80	FF	2A000000 00000000

device number

WWPN

FCP LUN



0x6000 0x00000001:0x5005076300c38550 0x00000000:0x5249000000000000



SCSI ID assigned by user





FCP Device Mapping – Configuration Example

```
bash# cat /proc/subchannels
```

```
Device sch.  Dev Type/Model CU  in use  PIM PAM POM CHPIDs
-----
6000  0011  1732/03  1731/03          80  80  FF  2A000000 00000000
```

device number

WWPN

FCP LUN



0x6000 0x00000001:0x5005076300c38550 0x00000000:0x5249000000000000



SCSI ID assigned by user

SCSI LUN assigned by user





FCP Device Mapping – Configuration Example

```
bash# cat /proc/subchannels
```

```
Device sch.  Dev Type/Model CU  in use  PIM PAM POM CHPIDs
-----
6000  0011  1732/03  1731/03          80  80  FF  2A000000 00000000
```

device number

WWPN

FCP LUN



0x6000 0x00000001:0x5005076300c38550 0x00000000:0x5249000000000000



SCSI ID assigned by user

SCSI LUN assigned by user

```
bash# cat /proc/scsi/scsi
```

Attached devices:

```
Host: scsi0 Channel: 00 Id: 01 Lun: 00
```

```
Vendor: IBM      Model: 2105F20
```

```
Rev: 2.91
```

```
Type: Direct-Access
```

```
ANSI SCSI revision: 03
```




```
bash# cat /proc/scsi/zfcp/map
```

```
0x5800 0x00000001:0x100000e00221b89c 0x00000000:0x0005000000000000
```

```
bash# cat /proc/scsi/scsi
```

```
Attached devices:
```

```
Host: scsi0 Channel: 00 Id: 01 Lun: 00
```

```
Vendor: PLEXTOR Model: CD-R PX-W1210S Rev: 1.04
```

```
Type: CD-ROM ANSI SCSI revision: 02
```





```
bash# cat /proc/scsi/zfcp/map
```

```
0x5800 0x00000001:0x100000e00221b89c 0x00000000:0x0005000000000000  
0x5800 0x00000002:0x100000e00221b89d 0x00000000:0x0009000000000000
```

```
bash# cat /proc/scsi/scsi
```

Attached devices:

```
Host: scsi0 Channel: 00 Id: 01 Lun: 00
```

```
Vendor: PLEXTOR Model: CD-R PX-W1210S Rev: 1.04
```

```
Type: CD-ROM ANSI SCSI revision: 02
```

```
Host: scsi0 Channel: 00 Id: 02 Lun: 00
```

```
Vendor: PIONEER Model: DVD-ROM DVD-305 Rev: 1.03
```

```
Type: CD-ROM ANSI SCSI revision: 02
```



```
bash# cat /proc/scsi/zfcp/map
```

```
0x5800 0x00000001:0x100000e00221b89c 0x00000000:0x0005000000000000
```

```
0x5800 0x00000002:0x100000e00221b89d 0x00000000:0x0009000000000000
```

```
0x6000 0x00000001:0x5005076300c38550 0x00000000:0x5249000000000000
```

```
bash# cat /proc/scsi/scsi
```

Attached devices:

```
Host: scsi0 Channel: 00 Id: 01 Lun: 00
```

```
Vendor: PLEXTOR Model: CD-R PX-W1210S Rev: 1.04
```

```
Type: CD-ROM ANSI SCSI revision: 02
```

```
Host: scsi0 Channel: 00 Id: 02 Lun: 00
```

```
Vendor: PIONEER Model: DVD-ROM DVD-305 Rev: 1.03
```

```
Type: CD-ROM ANSI SCSI revision: 02
```

```
Host: scsi1 Channel: 00 Id: 01 Lun: 00
```

```
Vendor: IBM Model: 2105F20 Rev: 2.91
```

```
Type: Direct-Access ANSI SCSI revision: 03
```

FCP/SCSI Tape Support

- **IBMtape special files (created by IBMtapeconfig):**

- /dev/IBMtape0
- /dev/IBMtape0n
- /dev/IBMchanger0



- **Tape utility program (IBMtapeutil):**

```
#Mount cartridge from slot 3
```

```
IBMtapeutil -f /dev/IBMchanger0 mount 3
```

```
#backup myfile.tar to tape
```

```
IBMtapeutil -f /dev/IBMtape0 write -s myfile.tar
```

Burning CDs with Linux for zSeries

The screenshot displays the K3b - The CD Kreator application window. The main interface shows a file tree on the left with the 'linux-2.4.19.SuSE' directory selected. A 'Progress - K3b' dialog box is overlaid on the right, providing detailed information about the burning process.

Progress - K3b Information:

- Creating image in /root/.kde/share/apps/k3b/temp/image.iso
- Image successfully created in /root/.kde/share/apps/k3b/temp/image.iso
- Starting recording at 4x speed...
- in 9 seconds.
- in 8 seconds.
- in 7 seconds.
- in 6 seconds.
- in 5 seconds.
- in 4 seconds.
- in 3 seconds.

Writing data:

Writing iso data

31%

Overall progress (00:02:39) 60 of 192 MB written

65%

Buffer Status:

Writer: PLEXTOR CD-R PX-W1210S 100%

Cancel

177.54 MB 650 K3b 0.7



FCP/SCSI on Linux for zSeries - Summary

- **FCP/SCSI support for IBM zSeries**
 - New FCP channel based on FICON / FICON Express cards
 - FCP channel support in z/VM 4.3 and higher for Linux guests
 - First FCP/SCSI exploitation for zSeries in SLES8 and RHEL3
- **Integration of your zSeries into standard based Fibre Channel Storage Area Networks**
- **New device types**
- **Reduced emulation overhead in OS and ESS compared to ECKD due to native use of fixed block I/O**
- **Larger disks in comparison to ESCON/FICON**
- **Current restrictions:**
 - Only switched fabric supported
 - No LUN sharing on a single adapter -> use separate physical adapters



Useful Links

- **I/O Connectivity on IBM zSeries mainframe servers**
 - <http://www-1.ibm.com/servers/eserver/zseries/connectivity/#fcp>
- **Getting Started with zSeries Fibre Channel Protocol, IBM Redpaper**
 - <http://www.redbooks.ibm.com/redpapers/pdfs/redp0205.pdf>
- **z/VM Version 4 Release 4**
 - <http://www.vm.ibm.com/zvm440/>
- **SuSE Linux Enterprise Server 8**
 - <http://www.suse.de/de/business/products/server/sles/index.html>
- **Linux for zSeries and S/390 (June 2003 stream)**
 - http://oss.software.ibm.com/linux390/june2003_recommended.shtml
- **Linux Device Drivers and Installation Commands**
 - <http://oss.software.ibm.com/linux390/docu/lx24jun03dd02.pdf>
- **IBM TotalStorage Tape Device Drivers – Installation and User's Guide**
 - <ftp://ftp.software.ibm.com/storage/devdvr/Doc/>
- **ESS Fibre Channel Attachment White Paper**
 - <http://www.storage.ibm.com/disk/ess/support/essfcwp.pdf>



Linux on zSeries and SCSI

Questions

