


# Part 2: Linux & Outlook

Concepts & Implementation  
***FCP for zSeries***

GSE Frühjahrstagung 2005, Berlin, 18.-20.5. 2005  
Martin Peschke, Linux for zSeries Development  
[mpeschke@de.ibm.com](mailto:mpeschke@de.ibm.com)

# Part 2 - Agenda



- Kernel 2.6 Concepts**
- Recent Linux Distributions**
- Multipathing**
- Supported Devices**
- Examples of Use**
- Summary & Outlook**

# Kernel 2.6

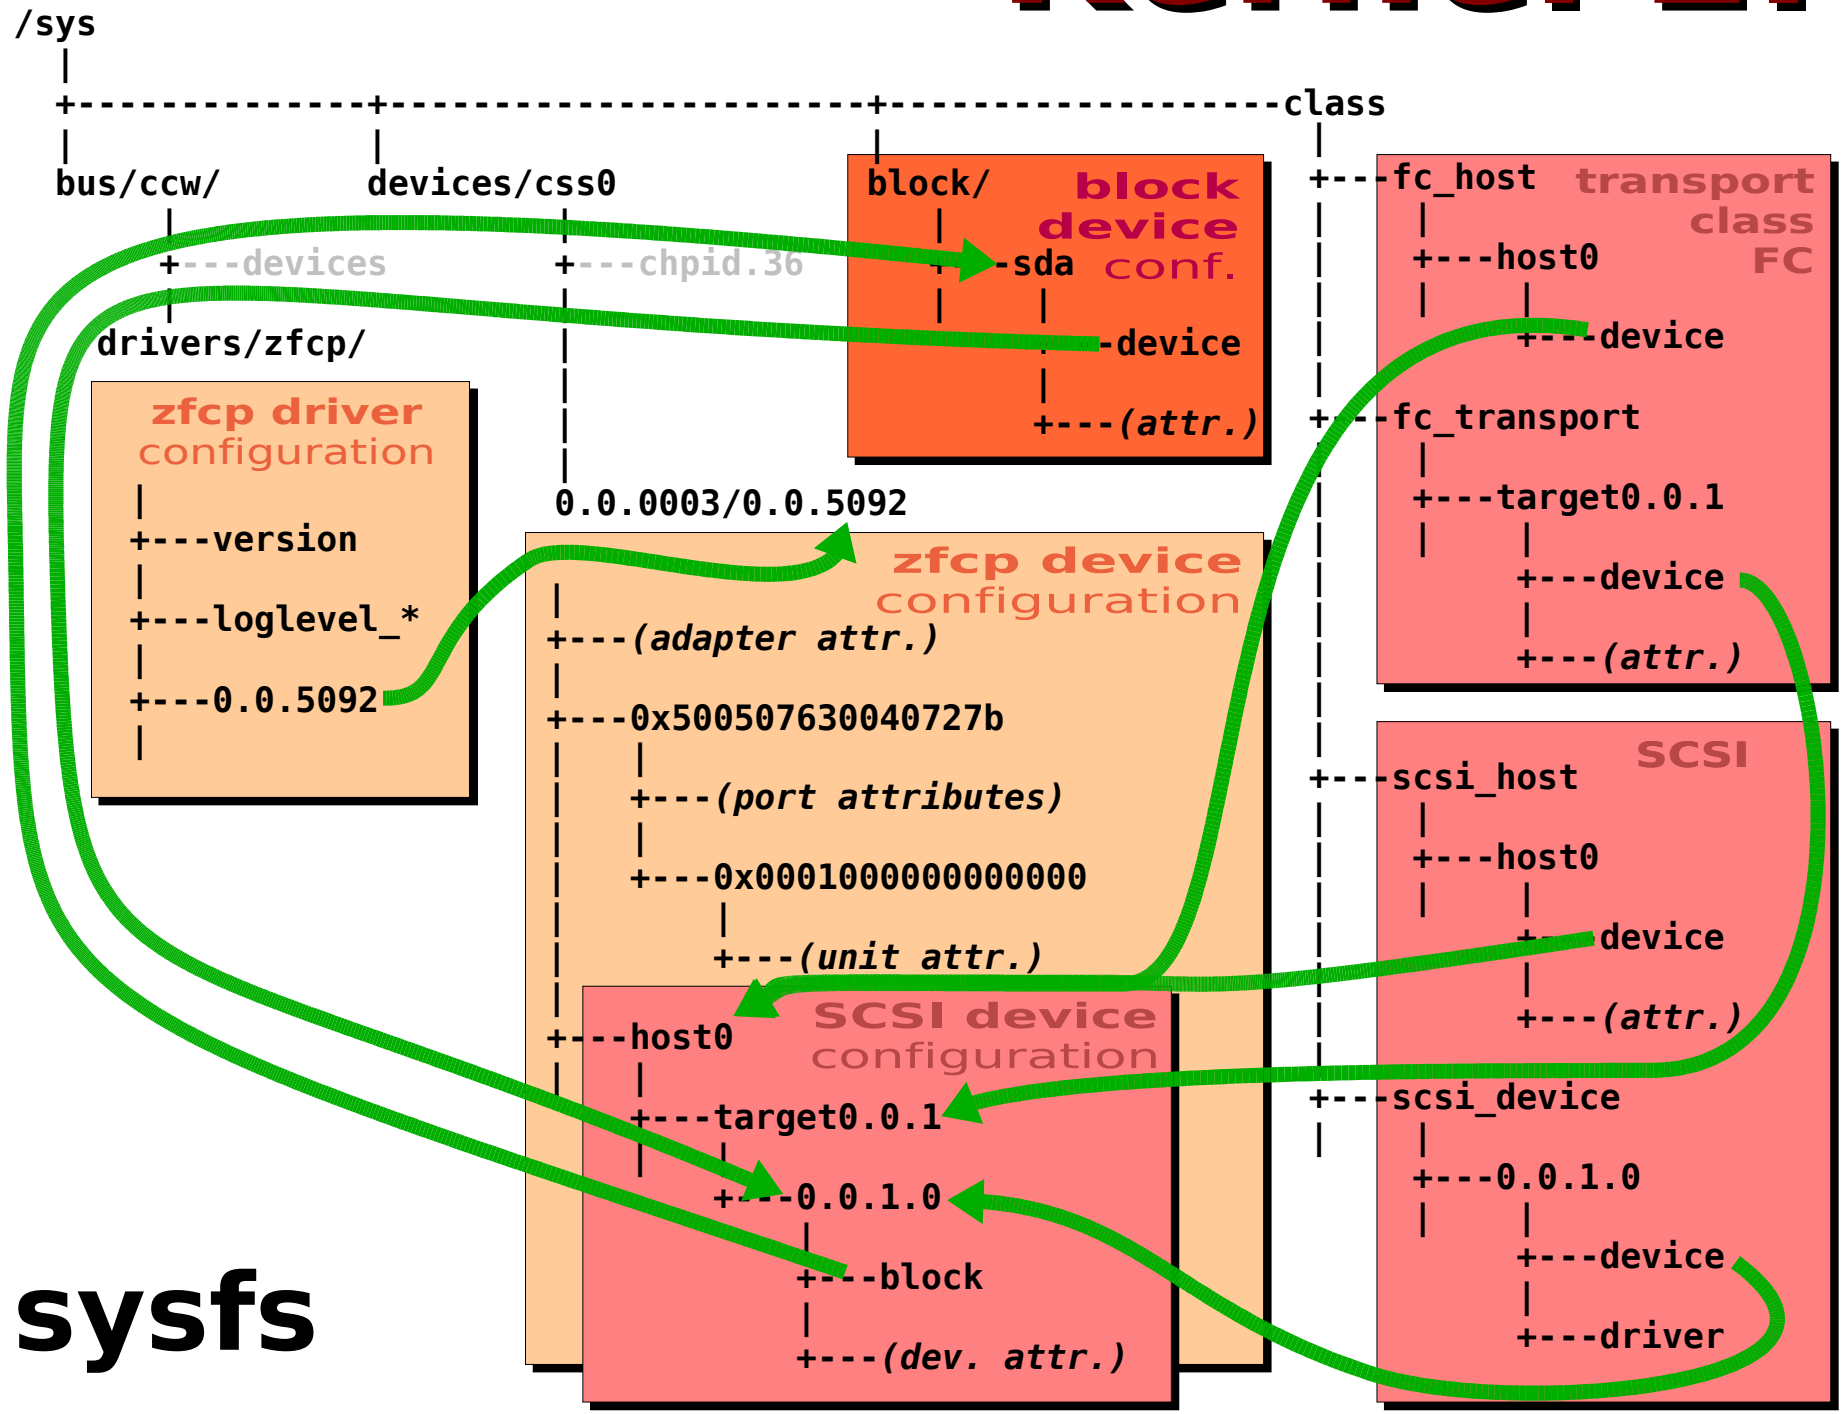


photo:  
seeds

# Kernel 2.6

## zfcplib Setup

- change to the zfcplib directory in sysfs**
- pick detected FCP subchannel you are going to use**
- set FCP subchannel online**
- configure an FCP device by adding a target port's WWPN**
- pick the newly added port**
- configure logical units (e.g. volumes of disk storage) by adding their FCP\_LUNs**

```
cd /sys/bus/ccw/drivers/zfcplib
```

```
cd 0.0.50d4
```

```
echo 1 > online
```

```
echo 0x5005076300c20b8e > port_add
```

```
cd 0x5005076300c20b8e
```

```
echo 0x5256000000000000 > unit_add
```

```
echo 0x525e000000000000 > unit_add
```

```
echo 0x525f000000000000 > unit_add
```

photo:  
gardening tools

# Kernel 2.6


## Tools

```
[root@tel15v39 root]# lsscsi -d
[0:0:1:0]    disk      IBM      2105F20      .487  /dev/sda[8:0]
[0:0:1:1]    disk      IBM      2105F20      .487  /dev/sdb[8:16]
[0:0:1:2]    disk      IBM      2105F20      .487  /dev/sdc[8:32]

[root@tel15v39 root]# lsscsi -l
[0:0:1:0]    disk      IBM      2105F20      .487  /dev/sda
             state=running queue_depth=32 scsi_level=4 type=0 device_blocked=0
[0:0:1:1]    disk      IBM      2105F20      .487  /dev/sdb
             state=running queue_depth=32 scsi_level=4 type=0 device_blocked=0
[0:0:1:2]    disk      IBM      2105F20      .487  /dev/sdc
             state=running queue_depth=32 scsi_level=4 type=0 device_blocked=0

[root@tel15v39 root]# scsi_id -s /block/sda
1IBM      2105      25614735
[root@tel15v39 root]# scsi_id -s /block/sdb
1IBM      2105      25E14735
[root@tel15v39 root]# scsi_id -s /block/sdc
1IBM      2105      25F14735
```

# SLES9

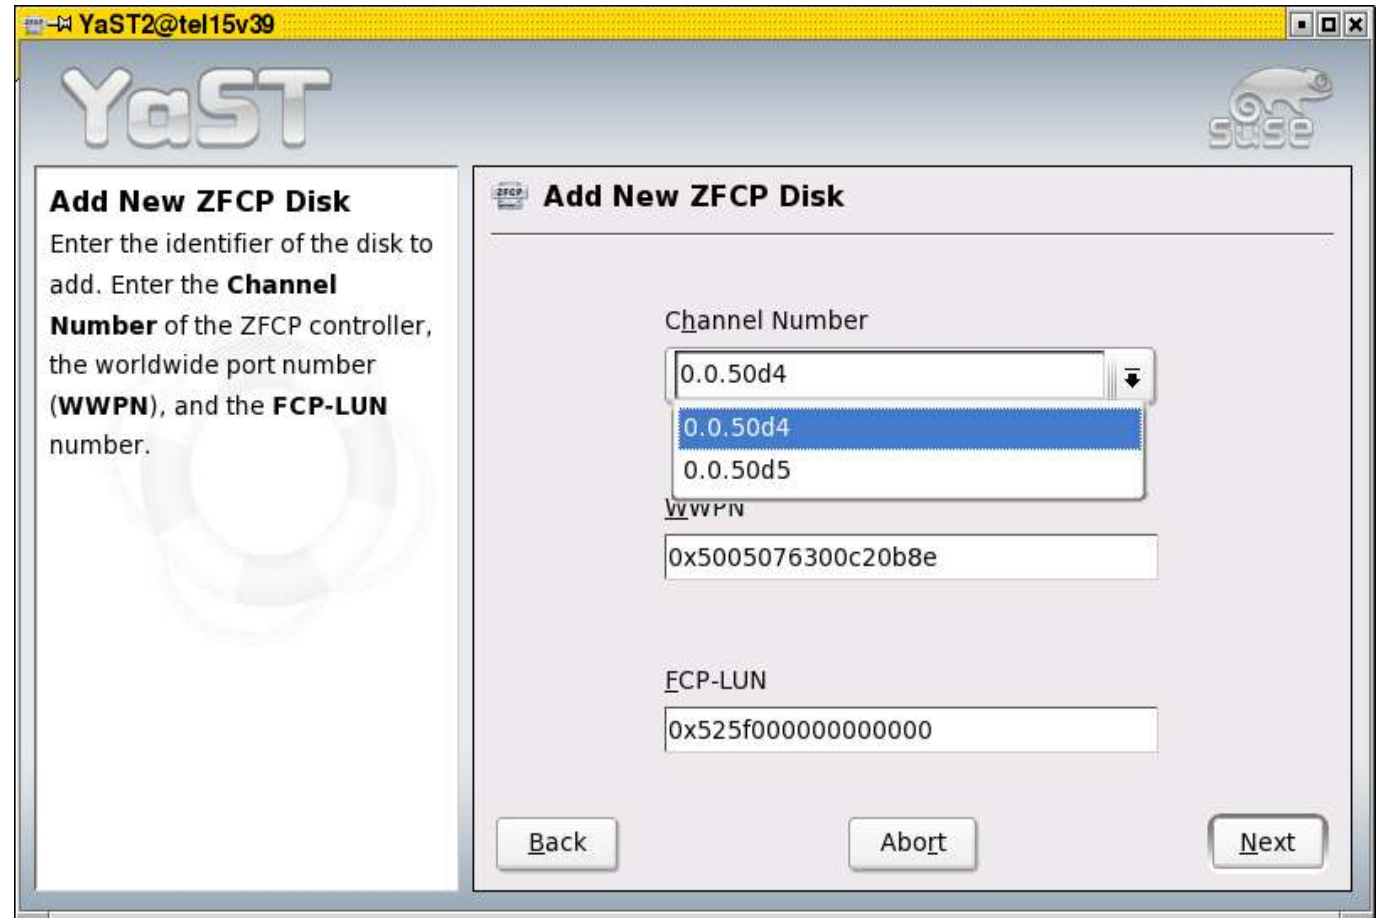


→ **new zfcplib dialog in YaST simplifies setup of SAN attached devices**

change photo:  
green pepper

**autodetects available FCP subchannels**

**copy'n paste WWPNs and FCP\_LUNs from configuration file obtained from SAN management tools or administrator**



## Persistent Device Names

photo:  
green pepper

```
tel15v39:~ # insserv boot.udev

tel15v39:~ # ls -Al /dev/disk/*
/dev/disk/by-id:
1IBM_2105_25F14735 -> ../../sda
1IBM_2105_25F14735p1 -> ../../sda1
1IBM_2105_25F14735p2 -> ../../sda2
/dev/disk/by-path:
ccw-0.0.50d4-zfcp-0x5005076300c20b8e:0x525f000000000000 -> ../../sda
ccw-0.0.50d4-zfcp-0x5005076300c20b8e:0x525f000000000000p1 -> ../../sda1
ccw-0.0.50d4-zfcp-0x5005076300c20b8e:0x525f000000000000p2 -> ../../sda2

tel15v39:~ # head -n 2 /etc/fstab
/dev/disk/by-id/1IBM_2105_25F14735p2 /          ext3    acl,user_xattr  1 1
/dev/disk/by-id/1IBM_2105_25F14735p1 swap       swap    pri=42           0 0

tel15v39:~ # mkinitrd

tel15v39:~ # zipl
```

see *SLES9 Documentation*  
also

<http://www.novell.com/documentation/sles9/>

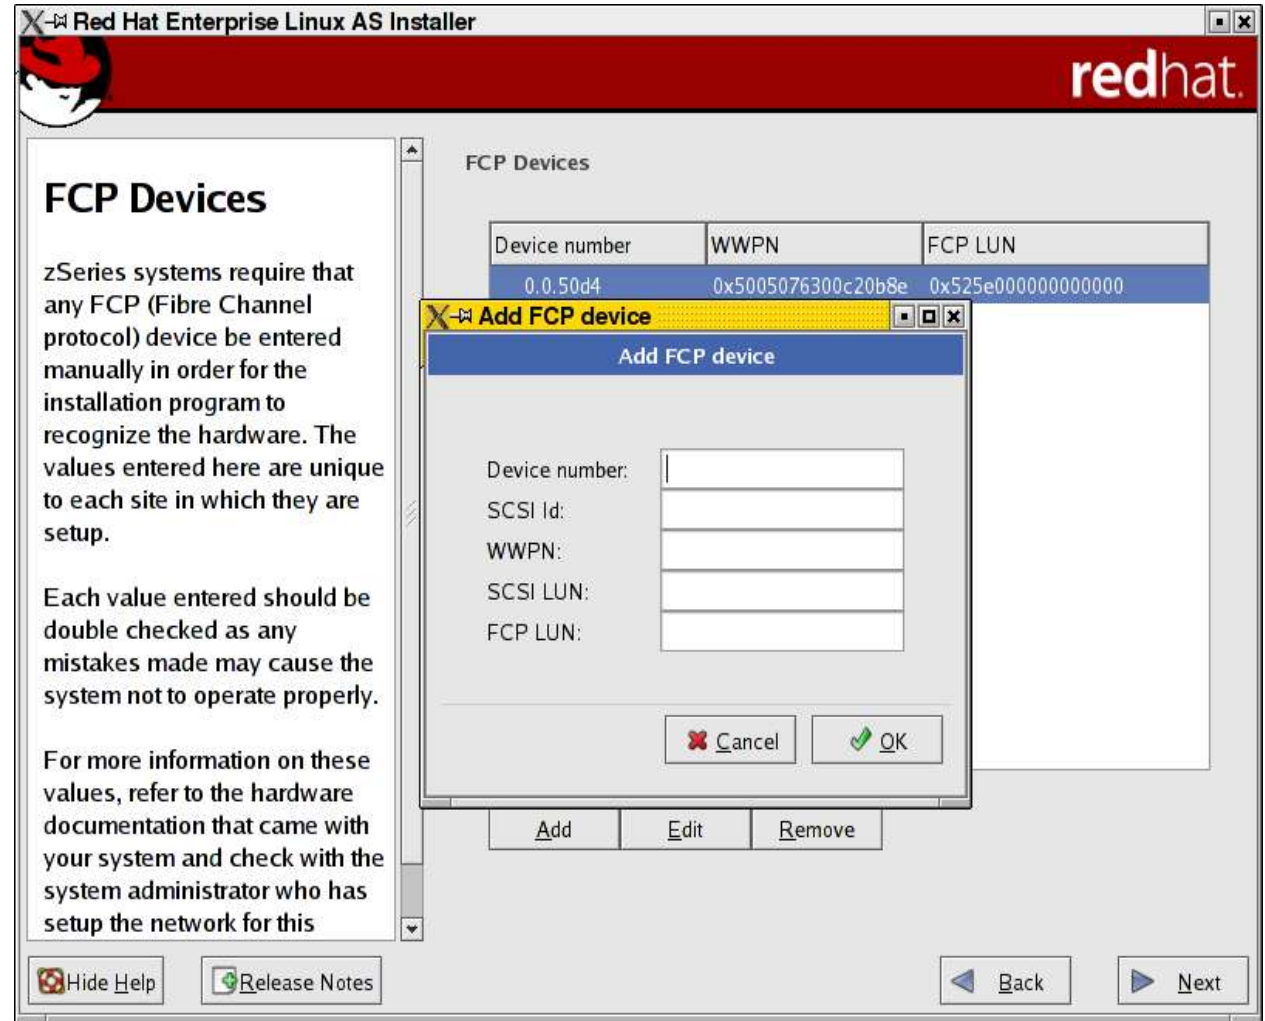
photo:  
tomato

**Attention:**

**SCSI ID and SCSI LUN shouldn't be needed. Be careful with assignments, though.  
SCSI ID: 1, 2, 3, ...  
SCSI LUN: 0, 1, 2, ...**

**Ignore subsequent complaints in case of DASD-less system.**

**GUI only available during installation. Use sysfs for addition of more FCP devices. Define FCP devices in /etc/zfcp.conf for permanent addition.**





# Multipathing

## EVMS

- **all-purpose storage management tool which is among many other things capable of setting up a multipath disk access**
- **attention: make sure to learn the EVMS terminology and read available documentation before setting out on a long quest - otherwise frustration is guaranteed**

**see also** *Enterprise Volume Management System*  
[evms.sourceforge.net](http://evms.sourceforge.net)      [www.tldp.org/LDP/EVMSUG/](http://www.tldp.org/LDP/EVMSUG/)

**see also** *The Linux Multipath Implementation*  
<http://christophe.varoqui.fr/multipath.html>

photo:

"Waldenbuch" - "hcubnedlaW"

# Multipathing

## md

**configure paths of  
FCP devices**

**create md multipath  
array with all paths**

**query UUID of md  
multipath array**

**permanently add  
array by UUID for  
more reliability**

**only use md device!**

```
[root@tel15v39 device]# cat /etc/zfc.conf
0.0.50d4 0x01 0x5005076300c20b8e 0x0 0x525e000000000000
0.0.50d4 0x01 0x5005076300c20b8e 0x1 0x526f000000000000
0.0.50d5 0x01 0x5005076300cc0b8e 0x0 0x526f000000000000
```

```
[root@tel15v39 ~]# mdadm -C /dev/md0 \
> --level=multipath \
> --raid-devices=2 /dev/sdb /dev/sdc
mdadm: array /dev/md0 started.
```

```
[root@tel15v39 ~]# mdadm --examine --scan /dev/sd*
ARRAY /dev/md0 level=multipath num-devices=2
        UUID=6b083d48:2c1e2d54:9afaa3c0:f68cd8f3
```

```
[root@tel15v39 ~]# cat /etc/mdadm.conf
DEVICE /dev/sd*
ARRAY /dev/md0 level=multipath num-devices=2
        UUID=6b083d48:2c1e2d54:9afaa3c0:f68cd8f3
```

```
[root@tel15v39 ~]# mke2fs -j /dev/md0
[root@tel15v39 ~]# mount -t ext3 /dev/md0 \
> /mnt/md-multipath-disk/
```

see  
also

*Linux on zSeries: Fibre Channel Protocol Implementation Guide*

<http://www.redbooks.ibm.com/redbooks/pdfs/sg246344.pdf>



# Supported Devices

	SLES8	RHEL3	SLES9	RHEL4
Tape System 3590, Drive 3592, Library 3494	supported	unsupported	supported	open
Tape Library 3584 (Ultrium 2 and 3 drives, 3592 drives)	supported	unsupported	supported	open
Tape Libraries 3581 U2, 3582, 3583 (Ultrium 2 and 3 drives)	supported	unsupported	supported	open
ESS 800, 750, F20, F10	supported	unsupported	in test	open
DS4300 (FASTT600)	supported	unsupported	unsupported	open
DS4400 (FASTT700)	supported	unsupported	unsupported	open
DS4500 (FASTT900)	supported	unsupported	unsupported	open
DS6000	supported	unsupported	unsupported	open
DS8000	supported	unsupported	in test	open
non-IBM devices	supported	unsupported	unsupported	open
color-coding:	supported	aiming for support	unsupported	open

see also *FCP for zSeries Connectivity* <http://www.ibm.com/servers/eserver/zseries/connectivity/#fcp>

see also *IBM TotalStorage® (refer to interoperability matrices)* <http://www.storage.ibm.com/>

# Backup Solution

## IBMtape and TSM



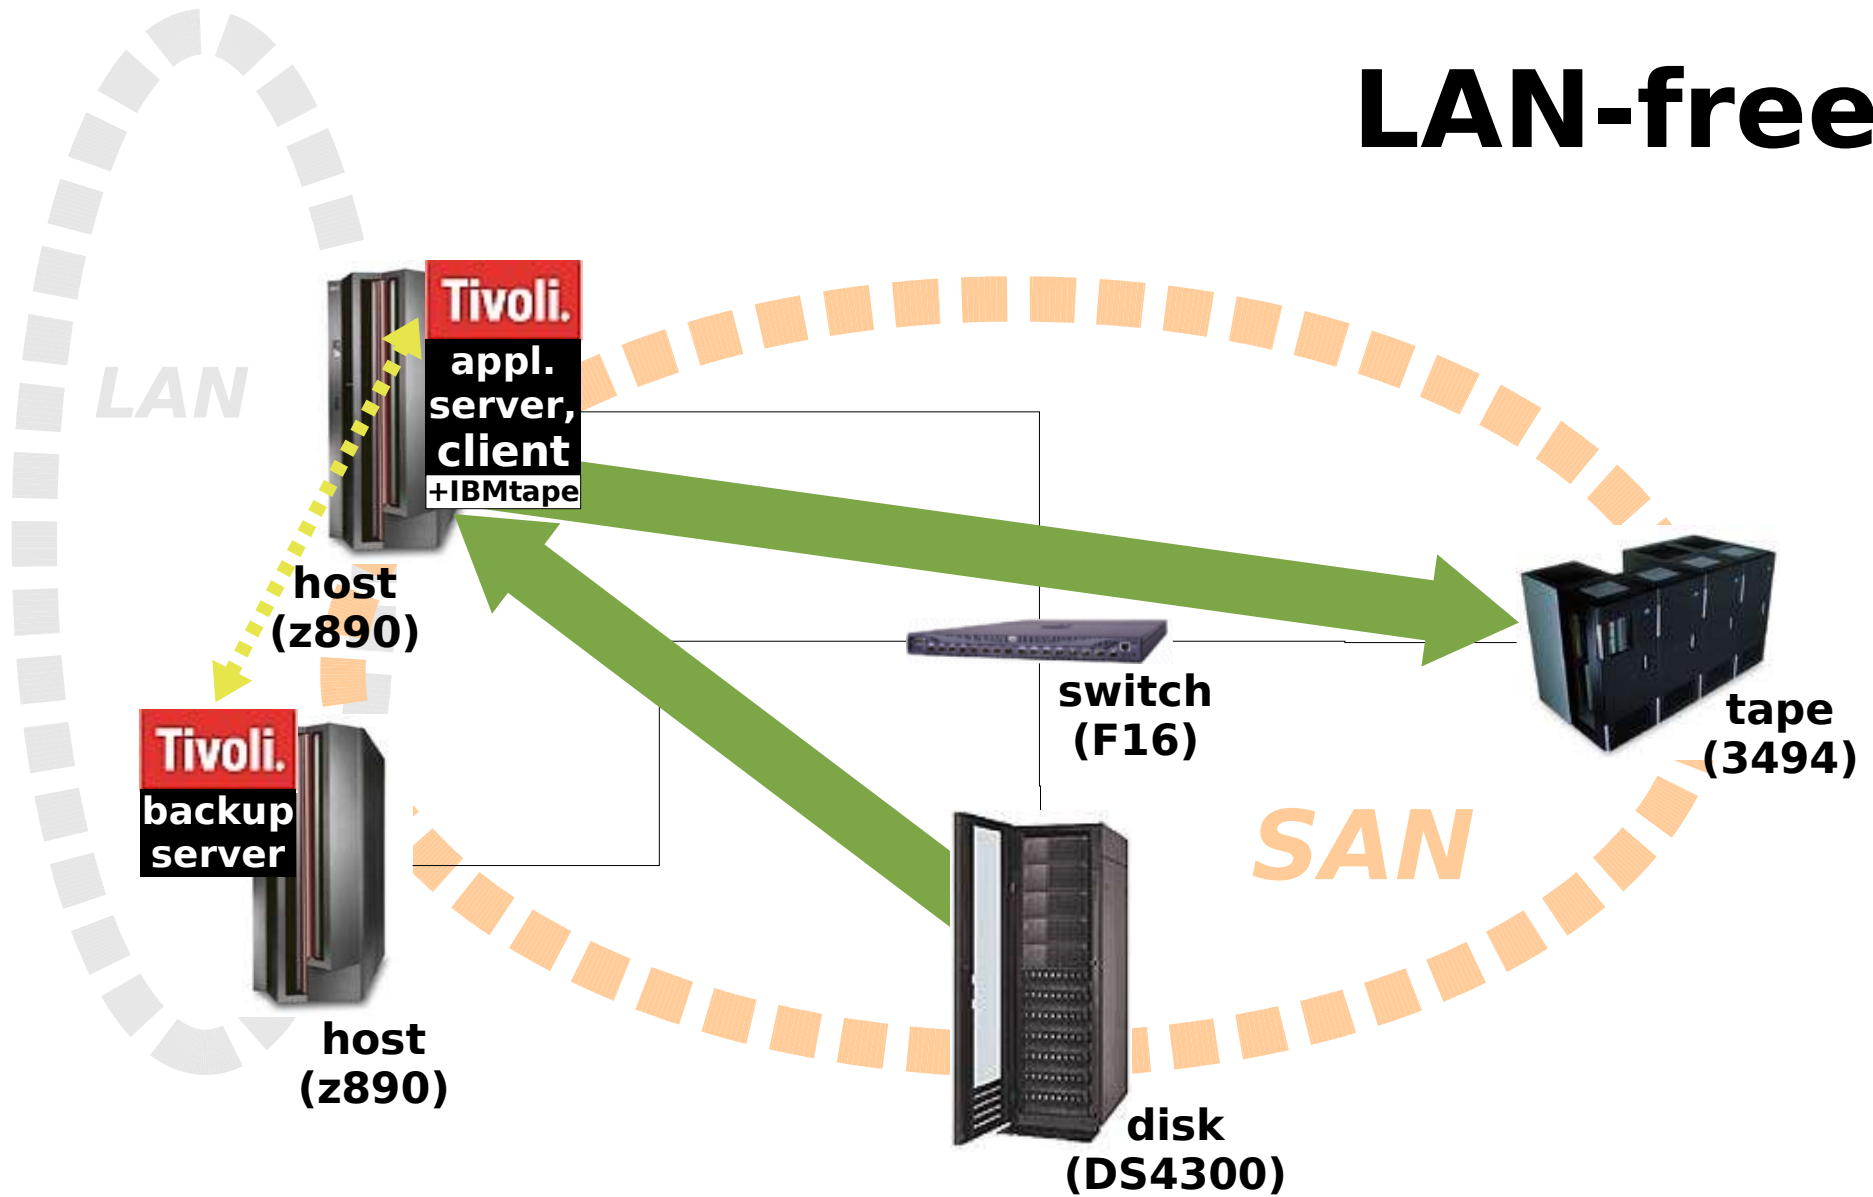
- “stand-alone” Linux backup solution without the need for help from z/OS
- TSM supports many SCSI tape devices, including OEM devices (IBM devices only supported by zSeries FCP so far)
- TSM facilitates modern SAN and SCSI features and provides efficient modes for backup and restore (LAN-free, server-free)
- both TSM client and TSM server are available for Linux on zSeries

**see also** *TSM server and client download* <ftp://service.boulder.ibm.com/storage/tivoli-storage-management/maintenance/>

**see also** *IBMtape download* <ftp://ftp.software.ibm.com/storage/devdvr/Linux/>

# Backup Solution

LAN-free



see also *IBM Tivoli Storage Manager (RedBooks Technote)* <http://publib-b.boulder.ibm.com/Redbooks.nsf/RedbookAbstracts/tips0118.html>

# More Good Uses

## Include Anything Requiring High Throughput

- as storage for databases: particularly, databases taking advantage of the 2.6 kernel's asynchronous and direct I/O capabilities should see noticable performance gains with FCP
- as fast swap device for systems that highly utilize swap
- as shared storage accessed through Samba



# Summary

photo:  
"in a nutshell"

	<b>z/VM guests</b>	<b>z/VM CP</b>	<b>z/VSE</b>	<b>SLES</b>	<b>RHEL</b>
<b>SCSI over FC</b>	4.3	5.1	3.1	8	4 (recommended)
<b>SCSI IPL</b>	4.4	5.1	3.1	8	4 (recommended)
<b>direct installation from SCSI media</b>	4.4	no	no	no	no
<b>direct installation onto SCSI media</b>	4.4	5.1	3.1	9	4
<b>SCSI device types accessible</b>		disk	disk	any	any
<b>FBA emulation for SCSI disk</b>	optional, by z/VM	required, by z/VM	required, by z/VSE or z/VM	optional, by z/VM (native SCSI recommended)	optional, by z/VM (native SCSI recommended)
<b>multipathing</b>	failover, failback, load balancing	failover, failback, load balancing	failover, failback (FBA emul. by z/VSE)	failover, failback, load balancing (not for root device)	failover, failback (not for root device)
<b>persistent device naming</b>	yes	yes	yes	yes	UUID of md as a circumvention

# Outlook - Linux

## Point-to-Point

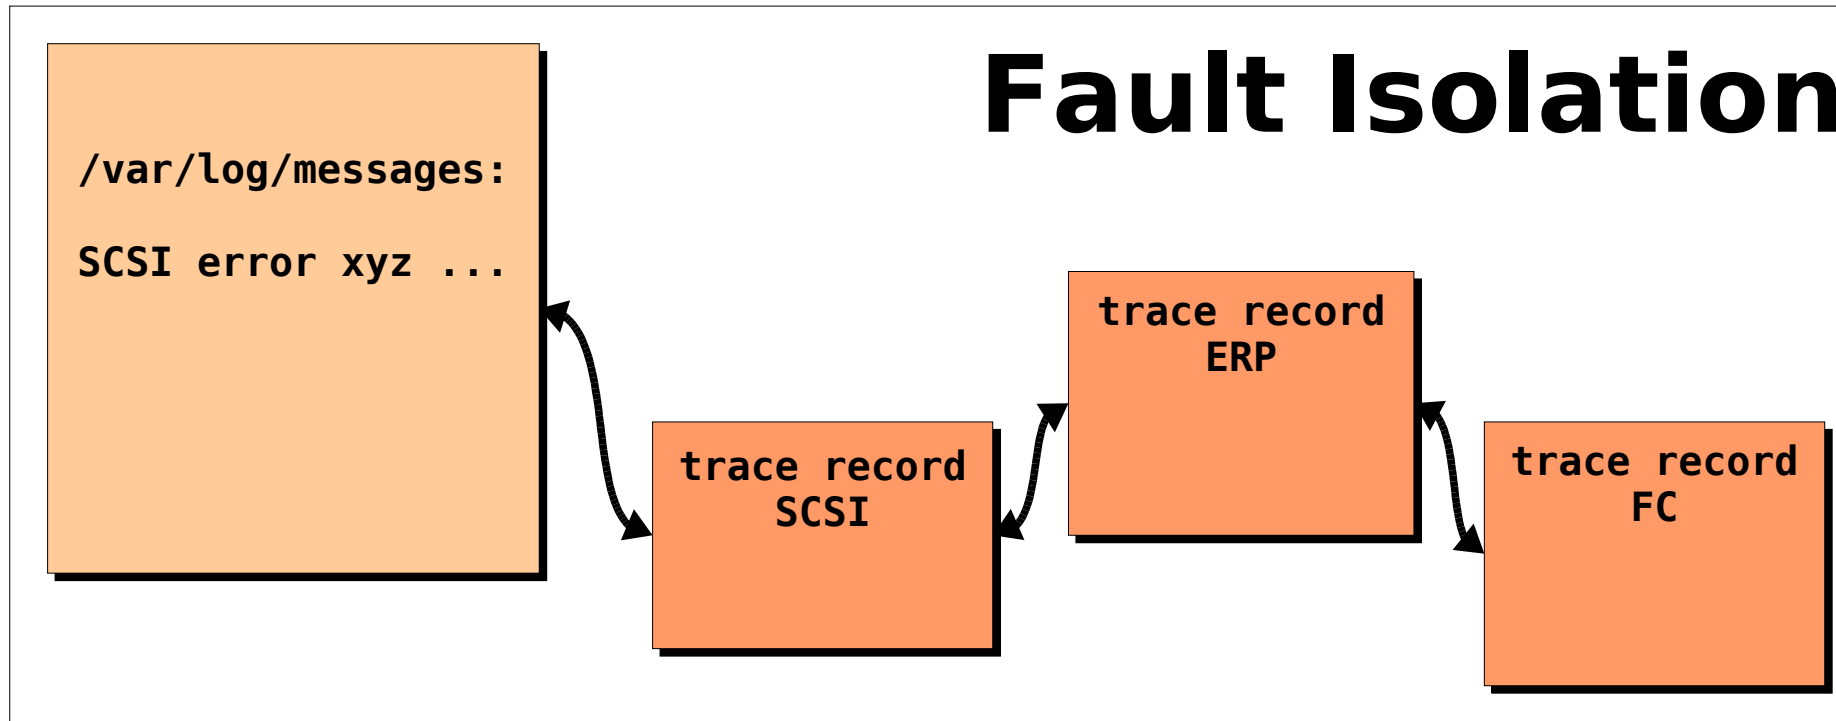
→ will allow for inexpensive configurations more easily setup for the purpose of evaluating FCP for zSeries with Linux at customer sites





# Outlook - Linux

## Fault Isolation



→ improving serviceability by restructuring the Linux for zSeries FCP device driver's error reporting and compiling a sort of zfcplib troubleshooting guide



# Outlook - Standards

## More Security

**see** *End-To-End  
also Data Protection*

<http://www.t10.org/drafts.htm#sbc2>  
<http://www.t10.org/drafts.htm#spc3>

**see** *High Integrity Fabric  
also*

<ftp://ftp.t11.org/t11/pub/fc/fs-2/04-045v3.pdf>

## & More Speed

**see** *4 Gigabit link speed  
also*

<ftp://ftp.t11.org/t11/pub/fc/pi-2/04-217v3.pdf>





The following are trademarks of the International Business Machines Corporation in the United States and/or other countries:

**Enterprise Storage Server, IBM\*, IBM logo\*, IBM eServer, z/VM, zSeries**

**\*Registered trademarks of IBM Corporation**

**Linux is a registered trademark of Linus Torvalds.**

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

# Trademarks