L07

Bill Worthington, bworthin@us.ibm.com
ATS Business Solutions

Linux on zSeries and DASD

VSE Technical Conference

...with credit to Tony Pearson, tpearson@us.ibm.com

Trademarks

The following are trademarks or registered trademarks of the IBM Corporation in the United States, other countries, or both:

- AIX®
- DFSMS/MVS®
- Enterprise Storage Server®
- ESCON®
- @server®
- FlashCopy®
- FICON®
- iSeries™
- IBM®
- Magstar®
- OS/400®
- pSeries™
- TotalStorage®
- xSeries®
- VSE/ESA
- zSeries®
- z/OS®
- z/VM®

Other company product and service names may be trademarks or registered trademarks of their respective owners.

© Copyright, IBM Corportion, 2003

Agenda

- Introduction to Linux on zSeries
- SAN and IP Storage connectivity options
- Volume and File-level management
- Resources

What is Linux?

- A UNIX-like operating system that was designed to provide PC users a very low-cost operating system comparable to traditional, more expensive UNIX systems
- Linux's kernel developed by Linus Torvalds at the University of Helsinki
- Makes use of components developed for the GNU project
- Portable Open Systems Interface (POSIX) compliant
- Supports Intel, PowerPC, Sparc, Alpha, iSeries, zSeries, et al.

Source: whatis.techtarget.com

Linux is based on Layers

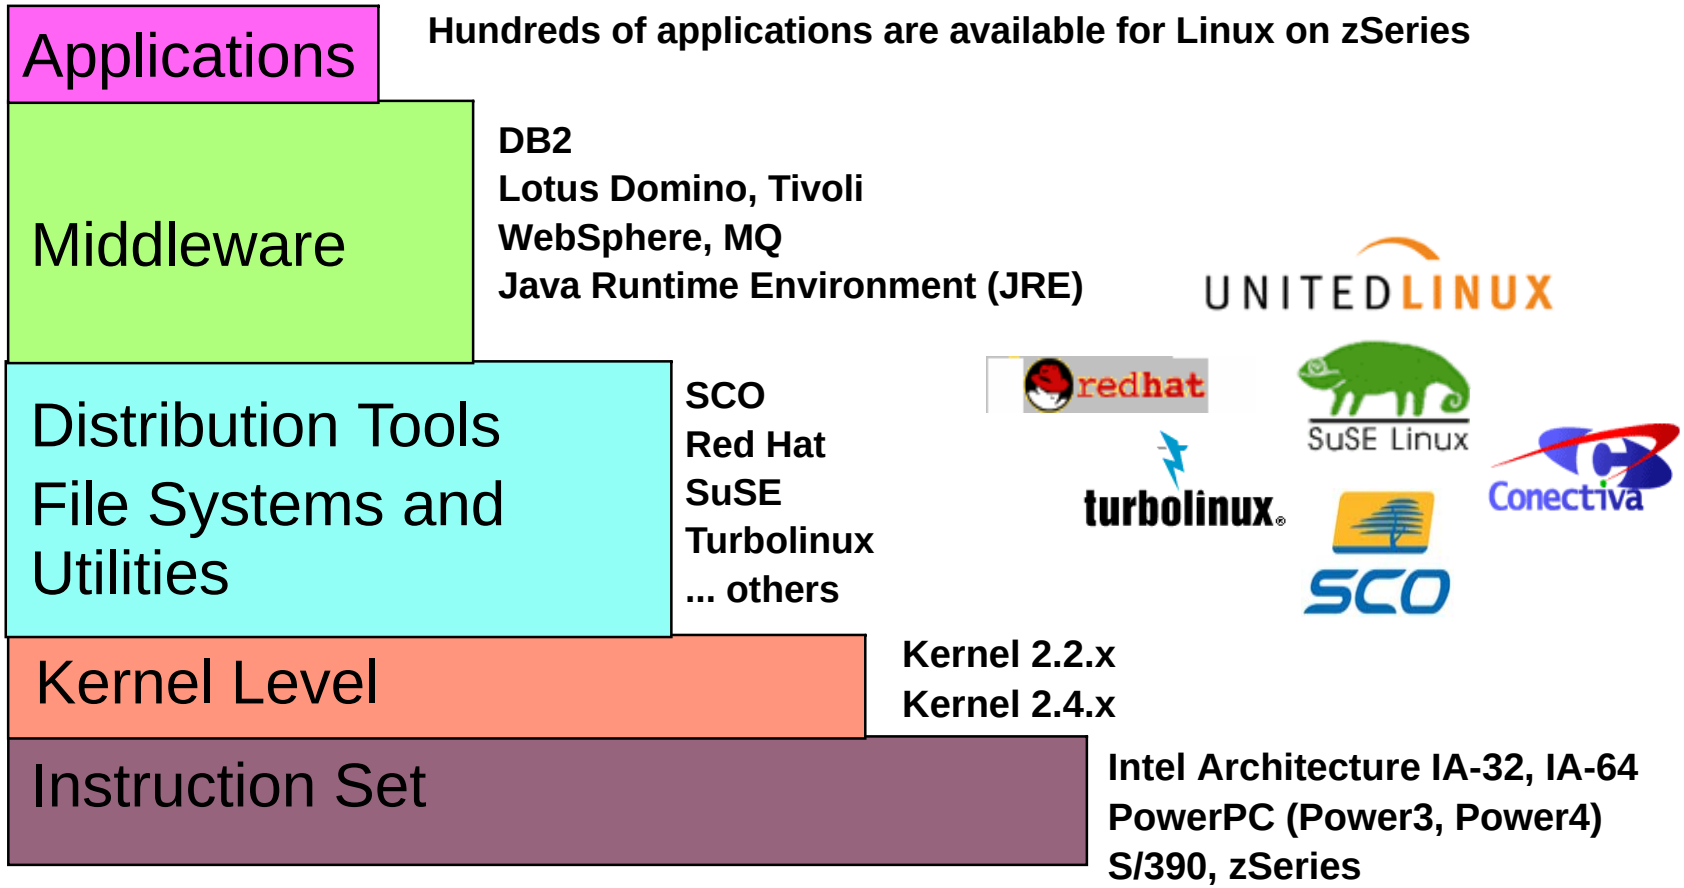

© Copyright, IBM Corportion, 2003

Linux is the Unifying Operating System

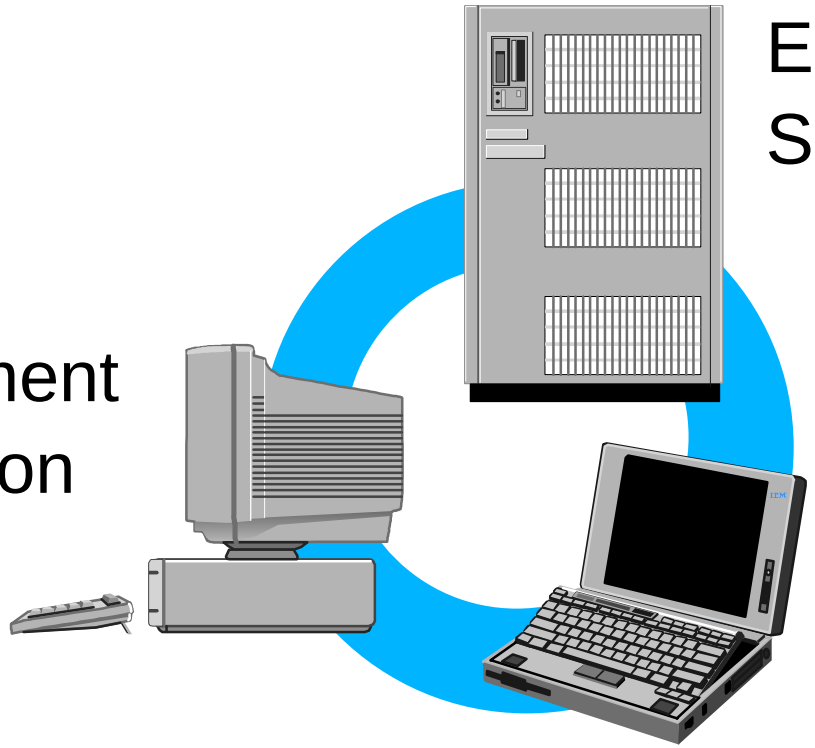

Enterprise Server

Development Workstation

Embedded Devices

© Copyright, IBM Corportion, 2003

Linux Distributions

IBM Linux Distribution Partners:

← all @server supported → ← xSeries only →

Enterprise Linux:

Red Hat
Enterprise Linux (RHEL) AS, ES, WS

UnitedLinux
SuSE Linux Enterprise Server (SLES 8)
+
Turbolinux, SCO, Conectiva

Linux 2.4.x kernel provides new features for Storage

- Dynamically adding/removing disk without reboot
- Enhanced Error Reporting
- Unified page and buffer caching
- Aggressive read-ahead for I/O, new RawIO driver
- Distributed interrupts (hardware and software)
- The VFS layer for all file systems from 31-bit to 44-bit, supporting files and filesystems up to 16 TB in size
- NFS v3 protocol compliance
- IPv4 compliance
- Journaling file systems
- Logical Volume Manager (LVM)
- SCSI command tag queuing support

IBM TotalStorage Linux Support Strategy

- New products support current distributions at GA:
  - Enterprise Levels:
    - Red Hat Enterprise Linux (RHEL)
    - SuSE Linux Enterprise Server (SLES)
  - Exceptions:
    - Red Hat Professional for xSeries only on FAStT
    - Interim support of Red Hat 7.1 and 7.2 on zSeries
- Support new distribution release or kernel patch within 30 days:
  - New Linux distribution every 12 to 18 months
  - New patch kernels about every 3 months
- Interoperability of storage products supported

IBM eServer zSeries

Linux available today for:

- S/390 and zSeries family
  - World's most scalable server
  - z800, z900 and z990 models supported
  - S/390 G5, G6, Multiprise 3000
  - Bulletproof reliability
  - Dynamic workload management
  - DB2, WebSphere, MQ, Java, Tivoli
- Linux Distributions
  - **Linux for S/390** for 32-bit architecture
  - **Linux for zSeries** for 64-bit architecture
  - Linux distributions from: Red Hat, SuSE, and Turbolin
- Variety of Operational Modes
  - Logical Partition (LPAR)
    - Integrated Facility for Linux (IFL) engines
  - Guests under VM/ESA, z/VM
    - z/VM v4 can run on IFL engines

General and IFL engines

To run z/OS or VSE guests under z/VM, you must run General purpose engines

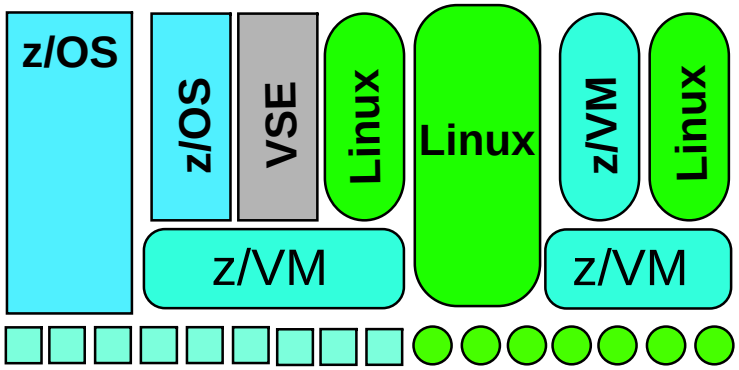

Linux can run in its own LPAR, or as a guest under z/VM, with General or IFL engines

LPARs cannot mix engines

General purpose engines can run z/OS, VSE, TPF, z/VM and Linux

Integrated Facility for Linux (IFL) engines can run Linux and z/VM 4.2 and above

What is a Storage Area Network?

Storage Area Network (SAN)

- A network dedicated to storage transfer, separating servers from storage subsystems
- Servers connect with "Host Bus Adapters" (HBA)
- First SAN was ESCON on mainframe environments, introduced over 10 years ago
- SAN now used primarily to refer to Fibre Channel Protocol (FCP)
- ESCON/FICON use of Channel Command Words (CCW) now referred to as Channel-Attach or **CON

Channel-Attached storage

- zSeries can have ESCON and
- FICON host bus adapters
  - Hundreds of ESCON channels
  - Up to 32 FICON on z800
  - Up to 96 FICON on z900
  - Up to 120 FICON on z990
- FICON bridge allows FICON cards
- on zSeries to attach to legacy ESCON disk and tape devices
- Enhanced Count-Key-Data (ECKD)
- used for disk layout (cylinders/tracks)

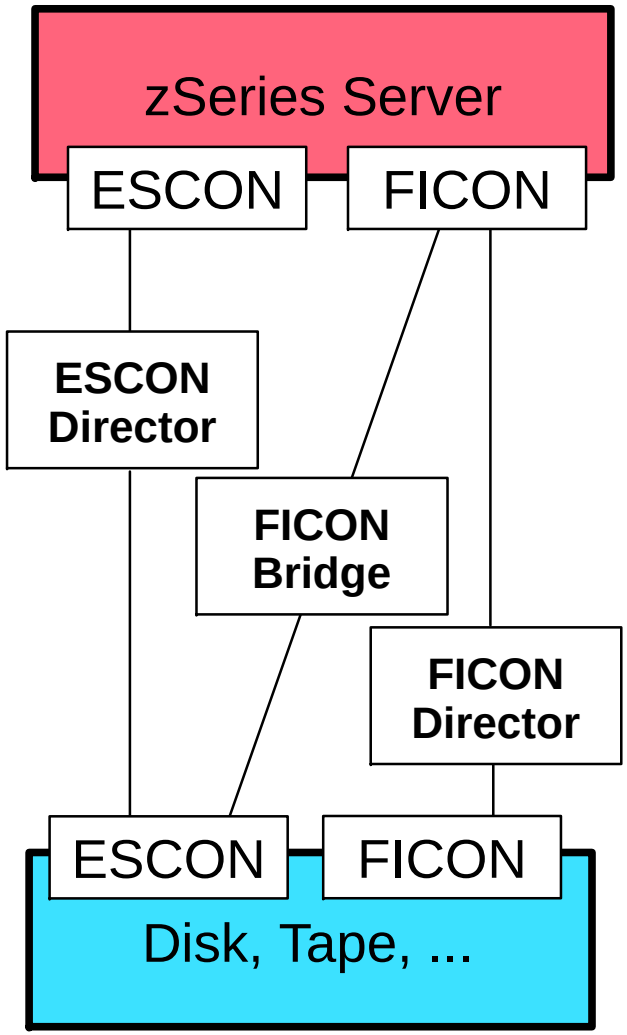

Mapping a File System to Disk

- File Systems map a directory structure, file attributes and file data blocks onto fixed block architecture (FBA)
- Linux supports a variety of File Systems:
  - Ext2 - the most common, most stable
  - Ext3 - similar, but adds journaling to improve recoverability
  - XFS - ported from Silicon Graphics' IRIX
  - ReiserFS - based on binary-tree, similar to DBMS mapping
  - JFS - IBM's Journaling File System, similar to ReiserFS
  - GPFS - IBM's General Parallel File System - clustering
  - NTFS - Microsoft Windows NT file system
  - UDF/ISO 9660 - The file system of CDR and CDRW

Disk Layouts for Linux

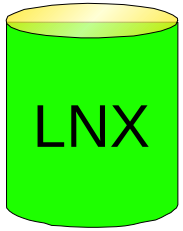

ECKD tracks formatted into fixed blocks (one 3390 track = twelve 4K blocks) to emulate Fixed Block Architecture (FBA)

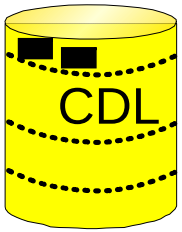

Compatible Disk Layout (CDL)
Same as above, except Track 0 is volume label, Track 1 is VTOC, rest are data blocks.
Can be carved up to three partitions.

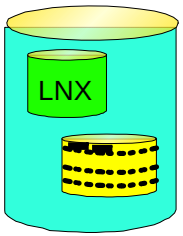

z/VM allows volume to be carved into multiple minidisks, each minidisk can be LNX or CDL layout

Linux on S/390 Compatible Disk Layout

dasdfmt utility

Track 0
Volume Label

Track 1
Volume Table
of Contents
(VTOC)

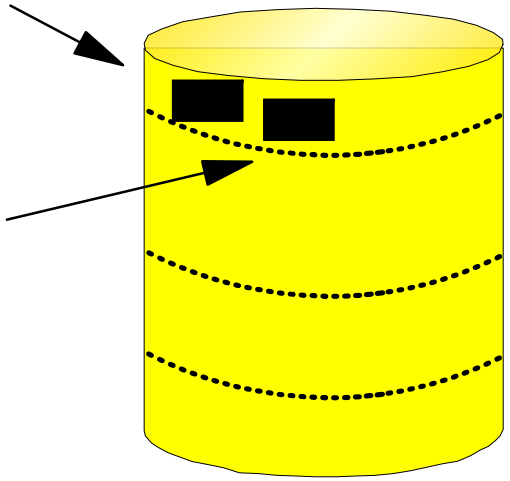

Remaining Tracks
fixed block size
(512, 1024, 2048 or 4096)
4096 is most efficient

fdasd utility

Tracks 2-n
can be carved
up into 1-3
native or swap
partitions

Each partition
described in
VTOC as data
set

Kernel 2.4.x

Linux on S/390 Compatible Disk Layout

Linux view:

z/OS, z/VM, VSE view

/dev/dasd/019b/part1
mounted as / (root)

/dev/dasd/019b/part2
mounted as /usr

/dev/dasd/019b/part3
mounted as swap space

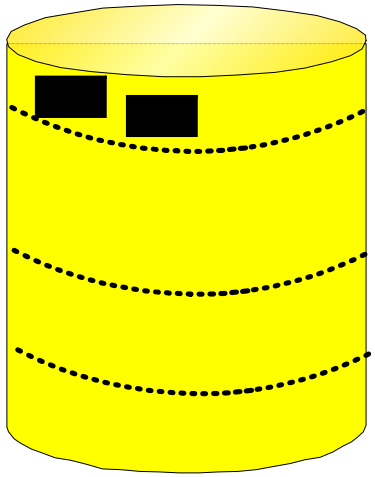

device=019B
volser=APP037
unit=3390
blocksize=4096

Label:
VOL1VOL1APP037
VTOC: 12 DSCB's

LINUX.VAPP037.PART0001.NATIVE

LINUX.VAPP037.PART0002.NATIVE

LINUX.VAPP037.PART0003.SWAP

Kernel 2.4.x

SAN Storage Protocols

- Small Computer System Interface (SCSI)
  - Industry standard parallel interface
  - Ultra, Ultra 2, Ultra 3, and SPI-4
- Serial Storage Architecture (SSA)
  - Industry standard serial interface
- Fibre Channel Protocol (FCP) SCSI over Fibre Channel
  - Point-to-Point
  - Arbitrated Loop (FC-AL)
  - Switched Fabric (FC-SW)

The evolution of Storage Protocol Standards

	Parallel	Serial	Fibre
SCSI commands	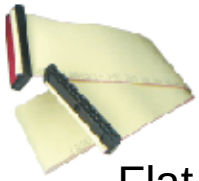 Flat Ribbon Cable	SSA: Serial Storage Architecture 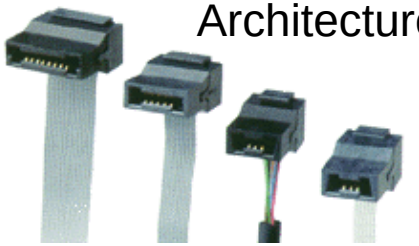	FCP: Fibre Channel Protocol 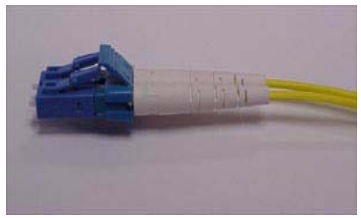 LC Duplex Connector
CCW commands	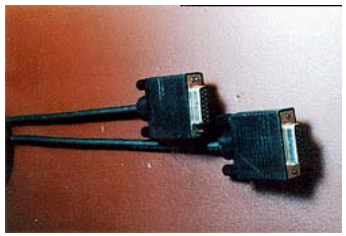 OEMI "Bus & Tag"	 ESCON: Enterprise System Connection	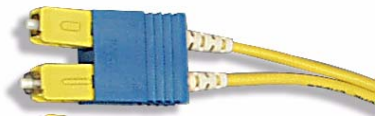 SC Duplex Connector FICON: Fiber Connection

Why support FCP on zSeries?

- Improve Performance
  - Avoid FBA-to-ECKD translation overhead
- Migration for Server Consolidation
  - Replace Windows, Sun, and Linux-Intel servers with zSeries
  - Keep data on existing storage devices
  - Exploit SAN technologies
- Potential new devices for zSeries
  - FAStT, LTO Ultrium, CD/DVD-ROM, CD/DVD-Writers, ...

FCP-attachment on zSeries

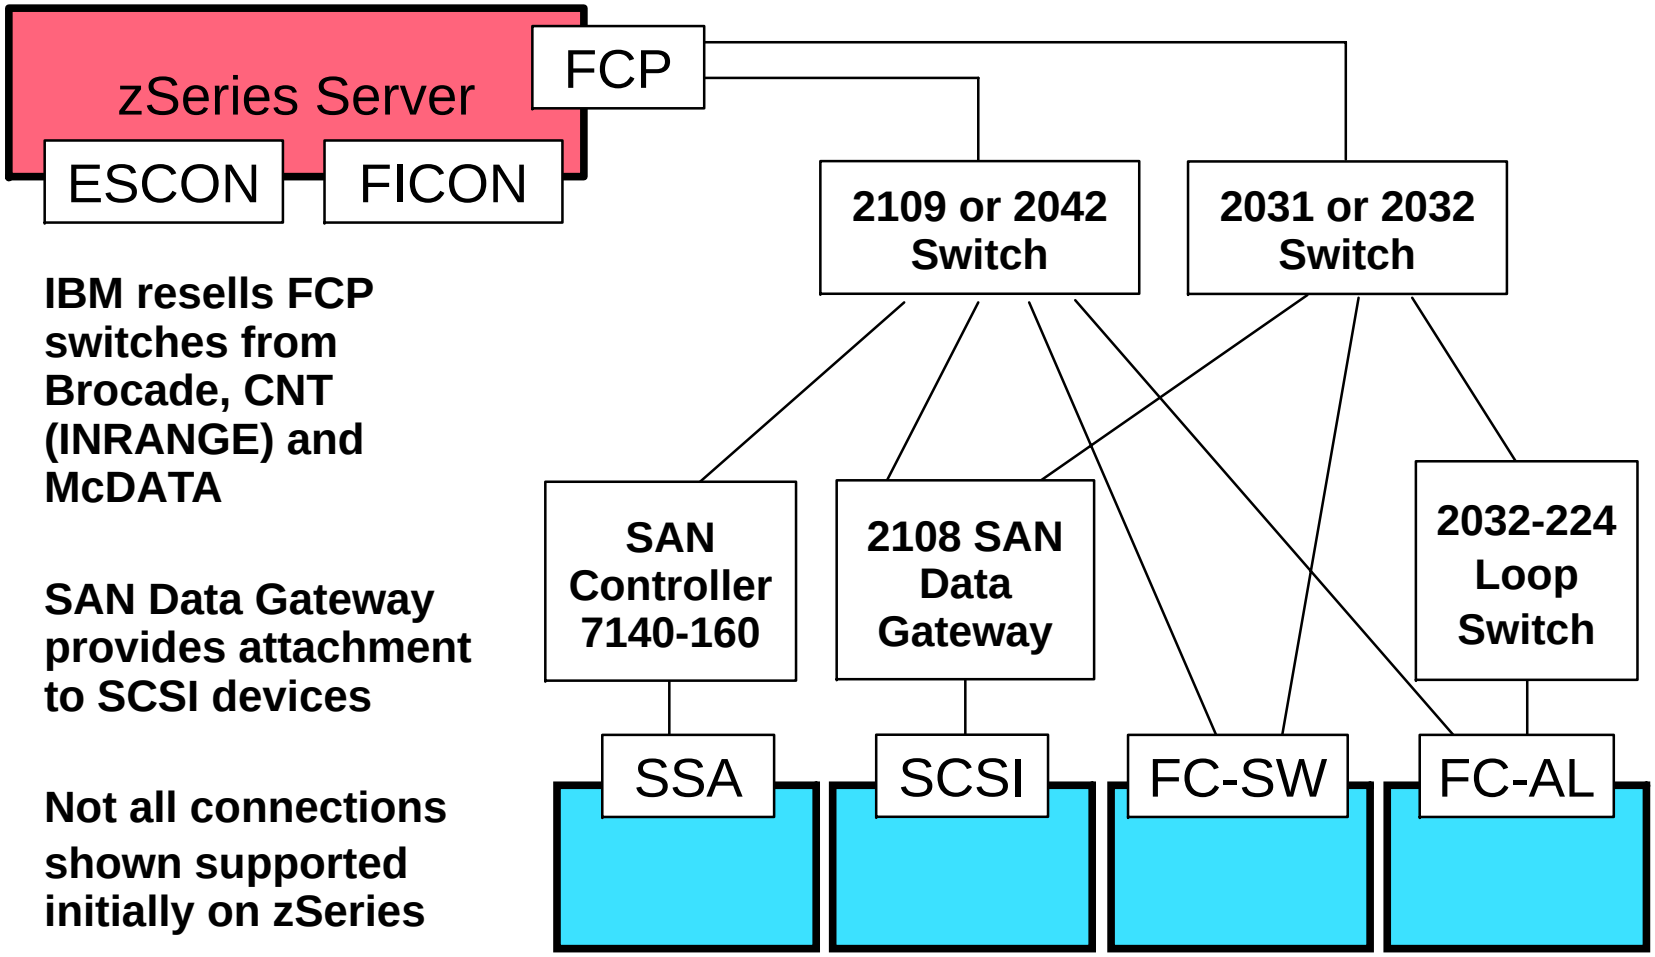

IBM resells FCP switches from Brocade, CNT (INRANGE) and McDATA

SAN Data Gateway provides attachment to SCSI devices

Not all connections shown supported initially on zSeries

FCP/FICON Intermix, FICON Cascading

You can now have both FCP and FICON ports on the same director

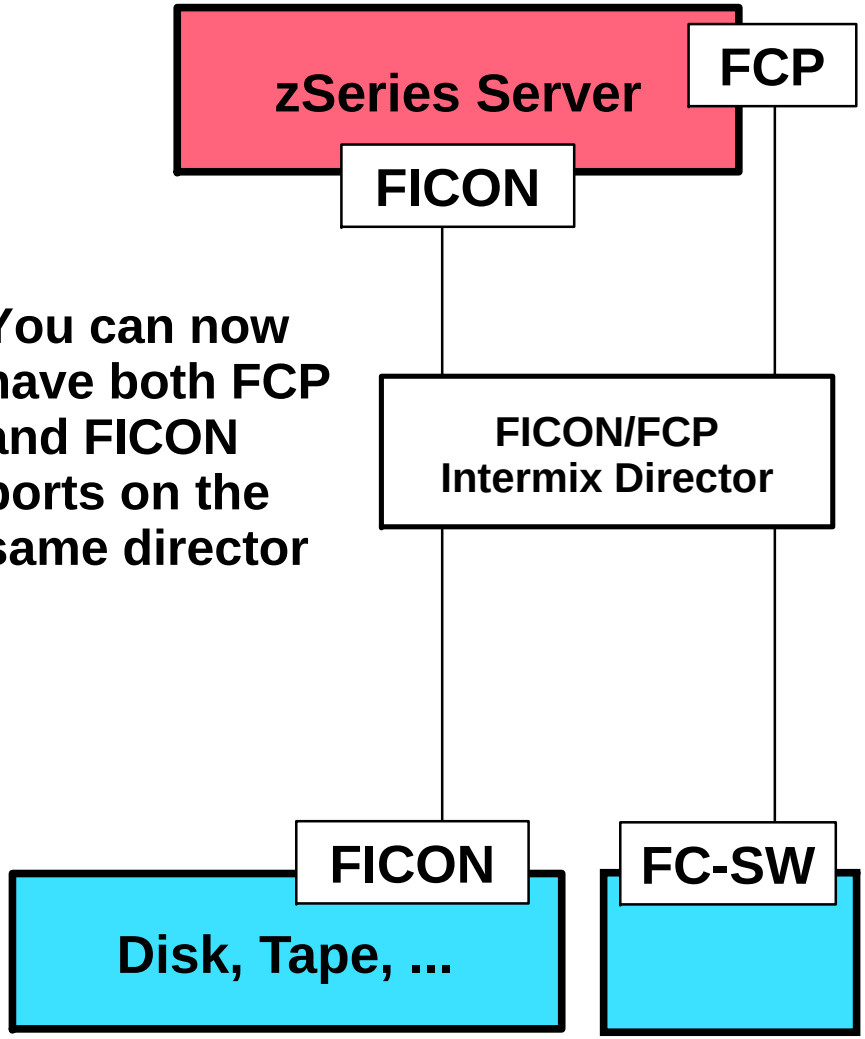

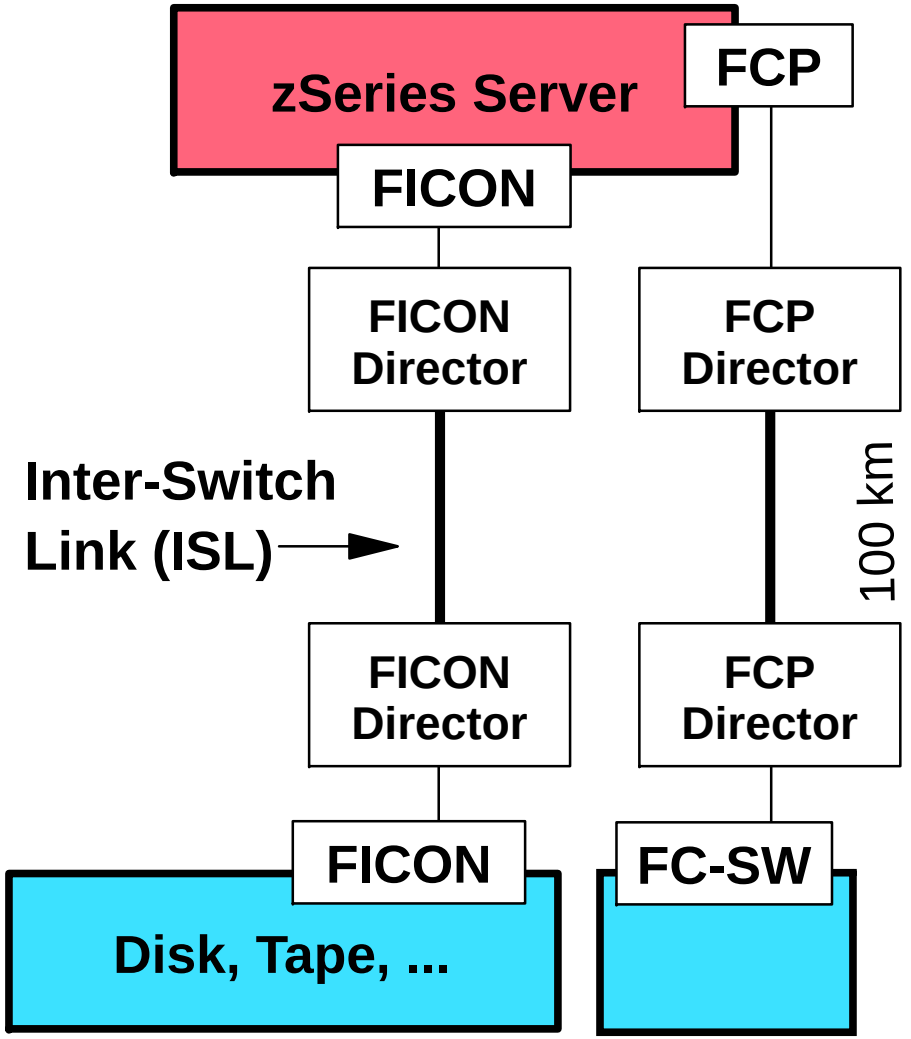

FCP on zSeries - Requirements

- IBM eServer zSeries (z800, z900 or z990)
- FICON/Express adapters, running in FCP mode (special microcode, limited availability)
  - FICON 2315, 2318 for z900
  - FICON Express 2319, 2320 for z800 and z900
- (FCP enabling) zSeries HW µcode (GA3)
- FC switch or director (no Direct Attach support provided)
- LPAR, or z/VM 4.3 to enable FCP support for Linux guests
- Linux 2.4 (including developerworks patches)
  - QDIO and zFCP open source drivers
  - SCSI stack + tools

FCP on zSeries - Restrictions

- Up to 240 Linux images per FCP card
  - All images represent the same WWN to the FC Switch
  - No LUN sharing over single FCP card
  - LUN sharing between Linux images requires multiple cards
  - Do not use for server hosting, trusted environments only
- Limit to number of FCP cards per zSeries
  - See FICON limit - a FICON or FICON Express can be either FICON or FCP mode
- SCSI request limited to 2152 KB (538 pages @4 KB)
- Cannot boot Linux yet from FCP-attach disk (see z/VM 4.4)
  - Must use ESCON/FICON disk to boot Linux
- z/VM, z/OS and VSE cannot access FCP storage devices directly

What is an IP network?

Internet Protocol (IP)

- A multipurpose, multi-protocol network for storage, messaging, configuration, control and inter-server communication
- Servers connect using Network Interface Card (NIC)
  - OSA Express supports Gigabit Ethernet
- Can support a variety of distances
  - Local Area Network (LAN), Wide Area Network (WAN)
- HiperSockets and z/VM Guest LANs allow IP communication between LPAR/guests inside server

© Copyright, IBM Corportion, 2003

IP Network Connectivity

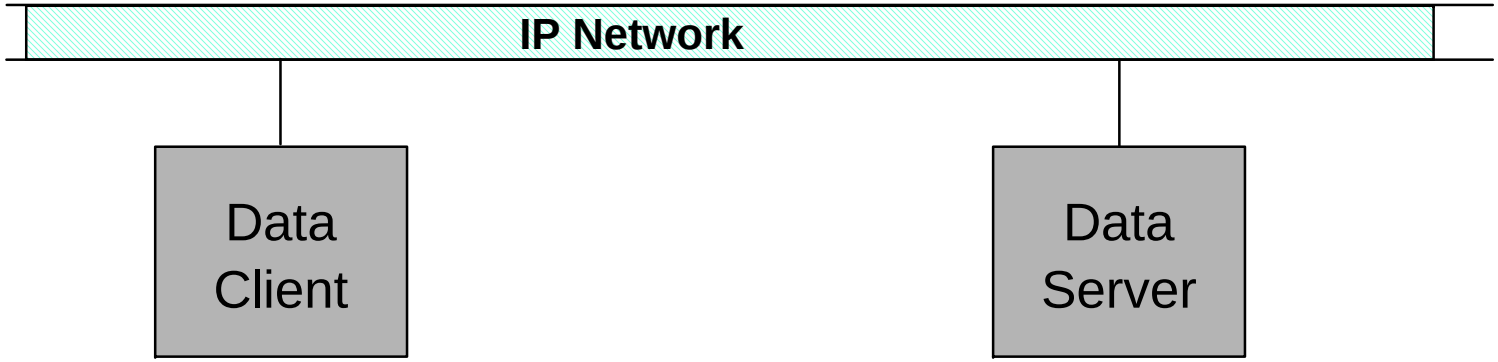

- **Network File System (NFS)**
  - **Clients available on most platforms, including Linux**
- **Common Internet File System (CIFS)**
  - **Formerly known as: Server Message Block (SMB)**
  - **Windows NT / 2000**
  - **Samba Client for Linux (smbfs)**
- **Hypertext Transfer Protocol (HTTP) including FTP**

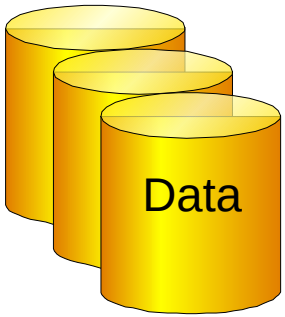

- **NFS Server**
  - **Linux NFS server**
  - **Variety of other platforms**
- **CIFS Server**
  - **Windows NT / 2000**
  - **Samba Server for Linux**
- **HTTP File Server**
  - **Apache, etc.**

IP Storage

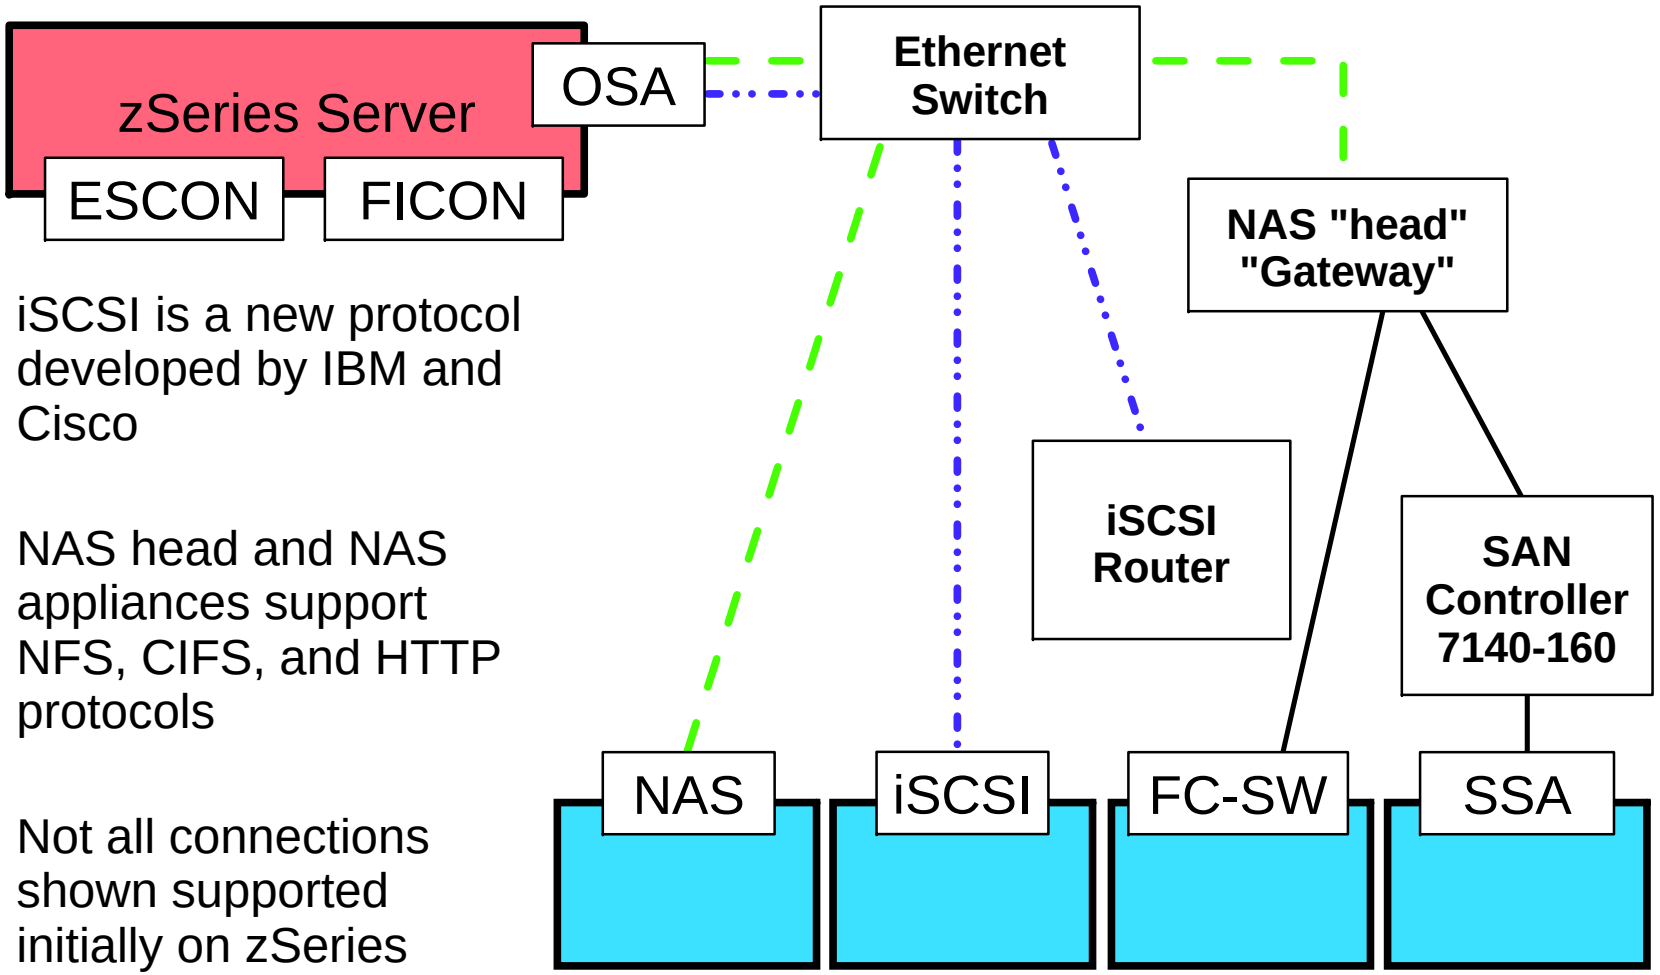

- iSCSI is a new protocol developed by IBM and Cisco
- NAS head and NAS appliances support NFS, CIFS, and HTTP protocols
- Not all connections shown supported initially on zSeries

© Copyright, IBM Corportion, 2003

Linux on zSeries as NAS "head"

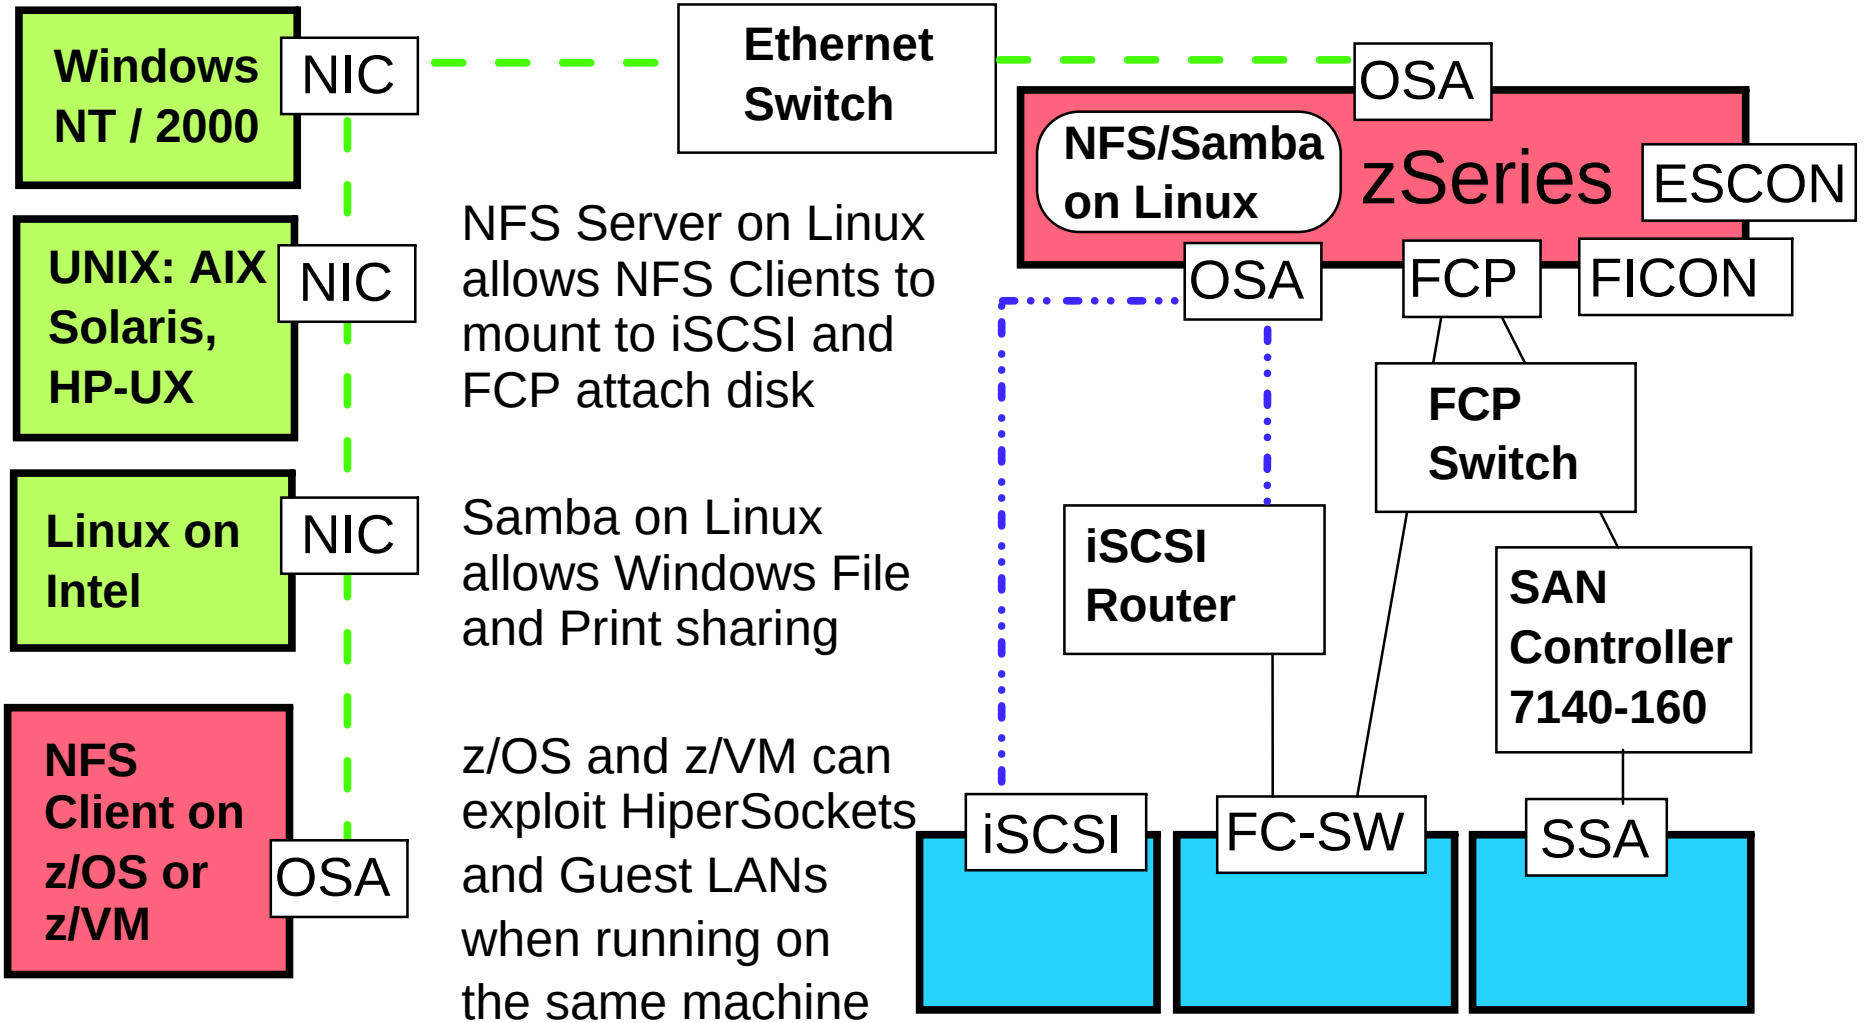

Example: Using Linux as a Gateway for FCP

- Linux partitions on zSeries and iSeries can act as a data gateway to access FCP-based storage solutions
- z/OS and z/VM access data mounted via UNIX System Services
- **NB:** This is just an example, and not any announcement of any future product or service

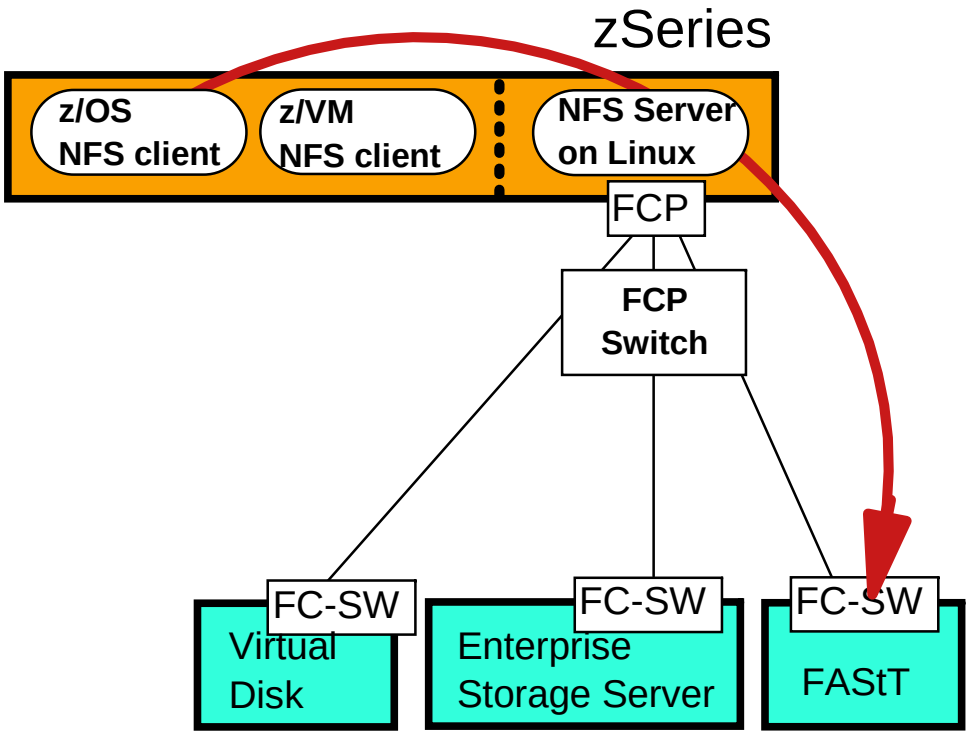

Linux on S/390 Compatible Disk Layout

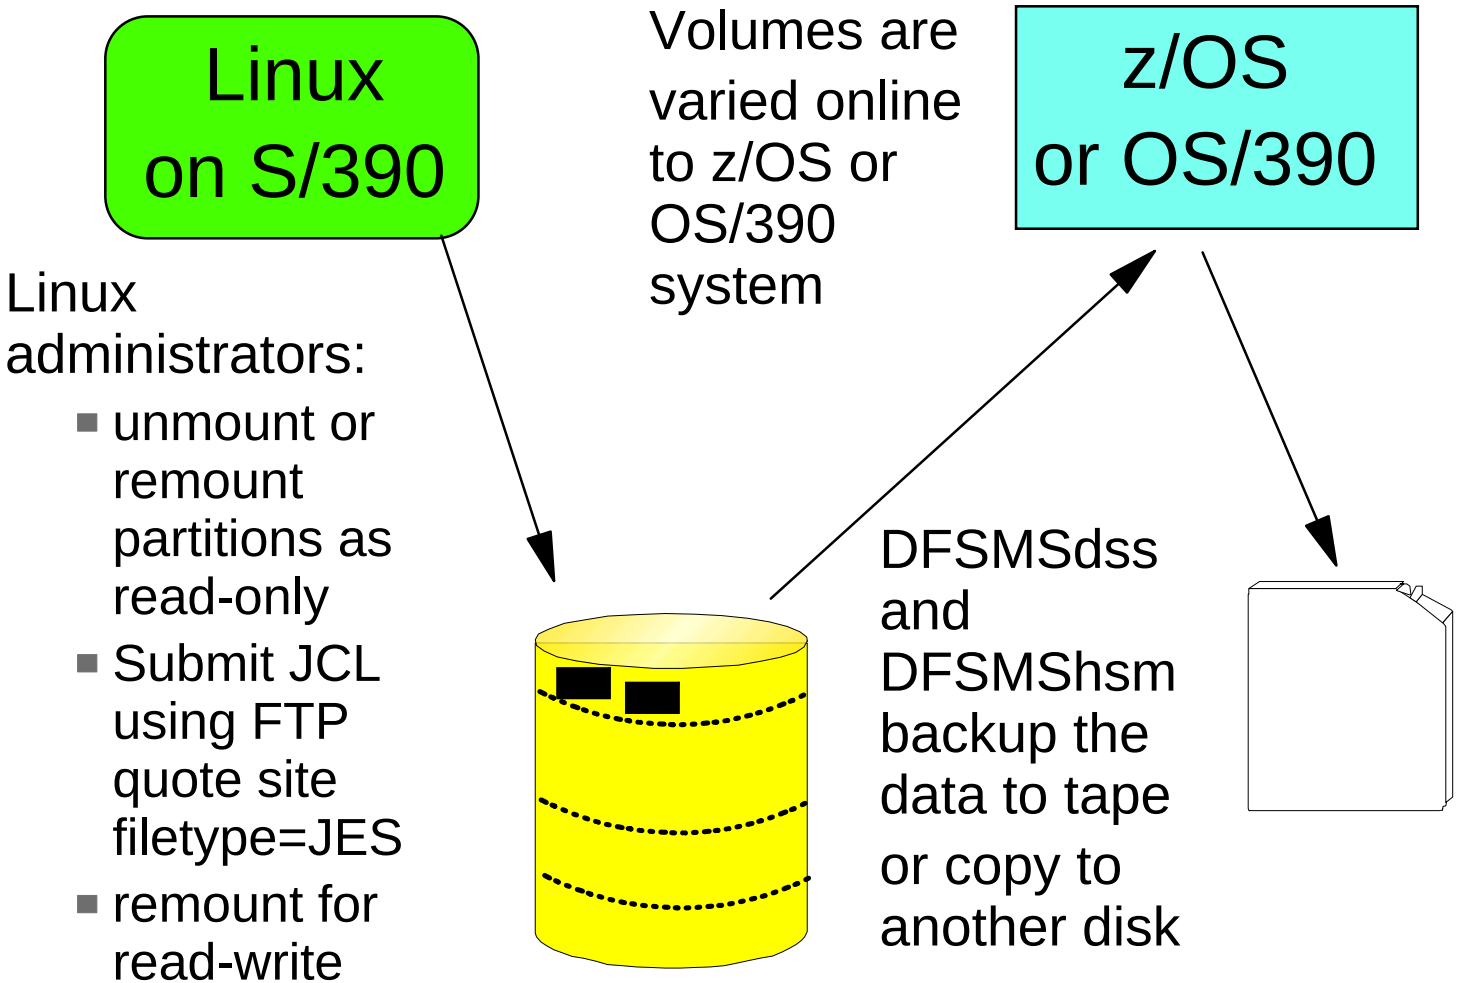

Kernel 2.4.x

What DFSMSdss can do

- Operations
  - Copy - Disk to Disk
  - Dump - Disk to Tape, or Disk to Disk
  - Restore - Tape to Disk
    - Stand-alone Restore can recover a Linux image directly from a DFSMSdss dump tape
  - CopyDump - Copy Dump Tape to another Tape
- Options
  - Full - Process entire volume
  - Dataset - Process individual partitions

What DFSMShsm can do

- Invokes DFSMSdss to perform Full volume dumps
- Output of DFSMShsm DUMP identical to DFSMSdss
- Simplifies processing:
  - Can create up to 5 output copies
  - Maintains inventory of up to 100 generations (which volumes were dumped on which days and are located on which tapes)
  - Automatic Expiration and Version Roll-Off

DFSMSdss vs. DFSMShsm

DFSMSdss

- Copy Linux disk to another disk
  - Uses Concurrent Copy, FlashCopy or SnapShot if available
  - Can be useful in cloning images
- Copy, Dump and Restore individual partitions

DFSMShsm

- Full Volume Dump only
- Maintains inventory of up to 100 generations
- Format identical to DFSMSdss dump tape
- Can be restored by DFSMSdss Stand-alone Restore

Peer-to-Peer Remote Copy (PPRC)

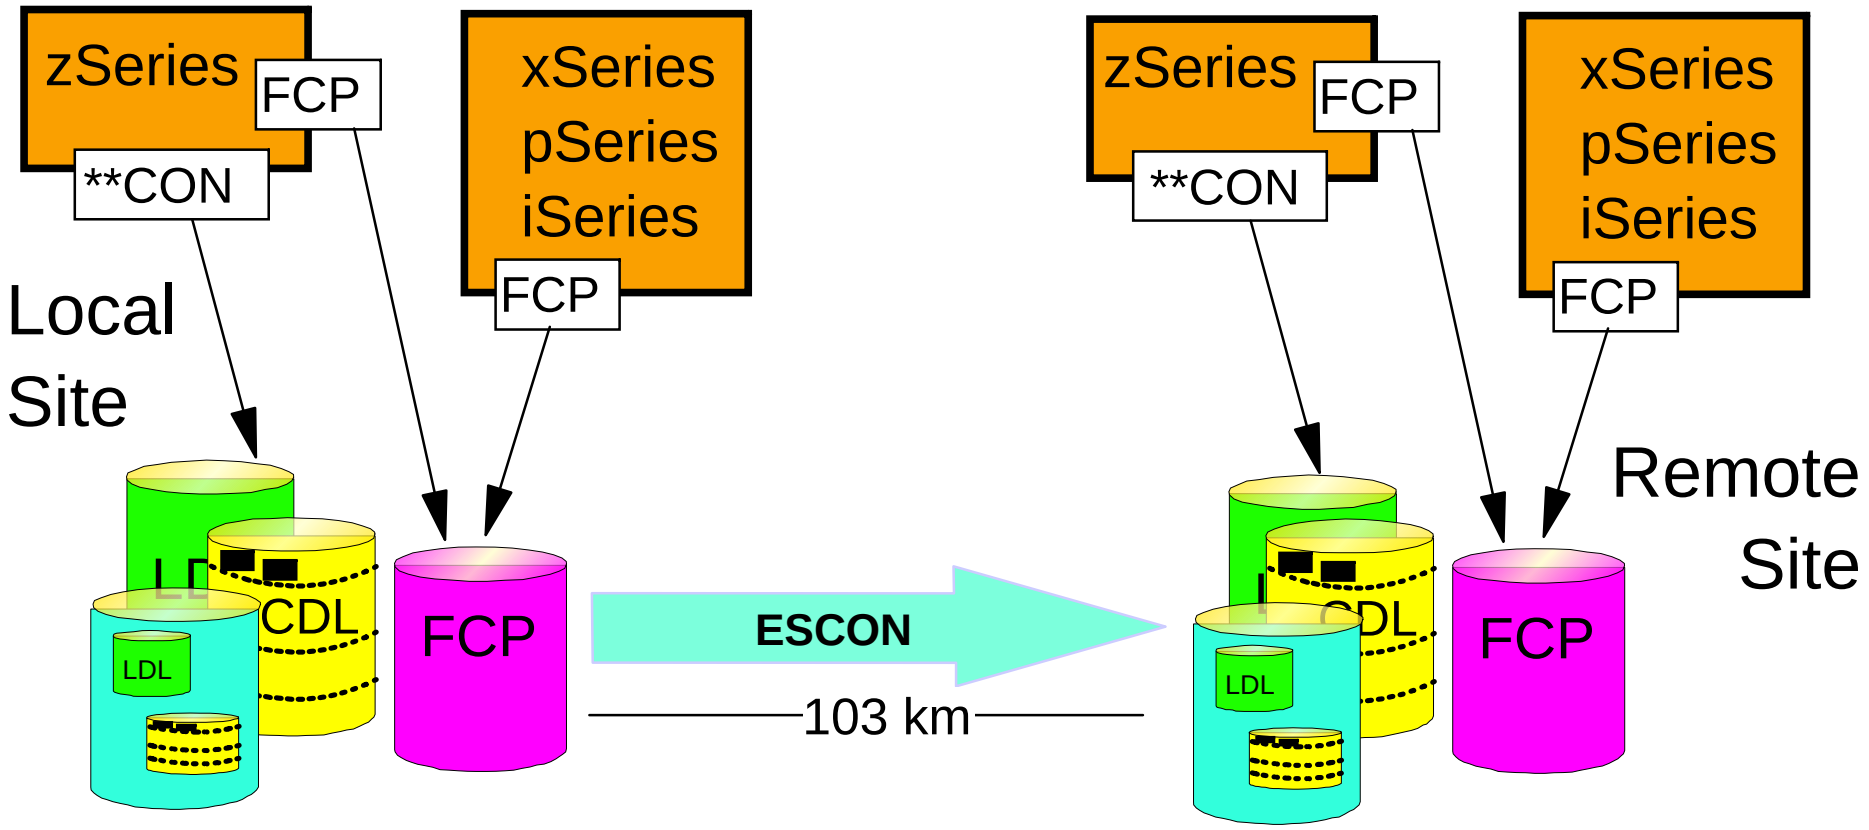

Using the PPRC feature of the ESS disk, you can mirror the volumes up to 103 km away

Peer-to-Peer Remote Copy Extended Distance (PPRC-XD)

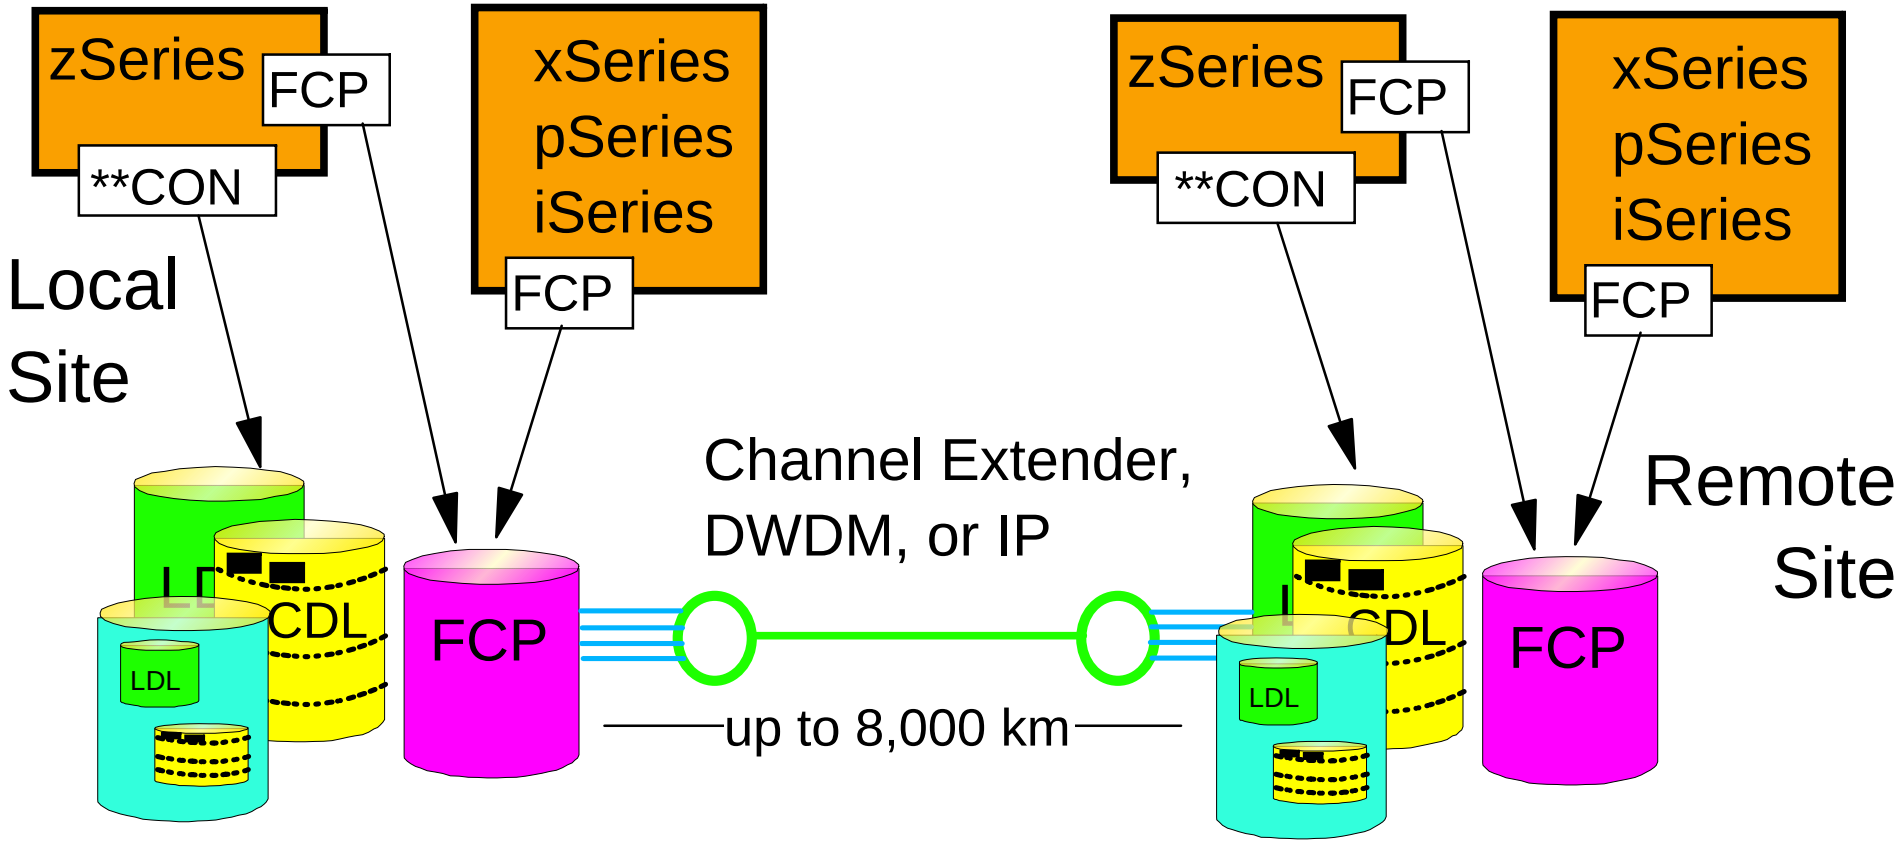

PPRC-XD feature of the ESS disk, you can mirror the volumes asynchronously (continental distances)

Extended Remote Copy (XRC)

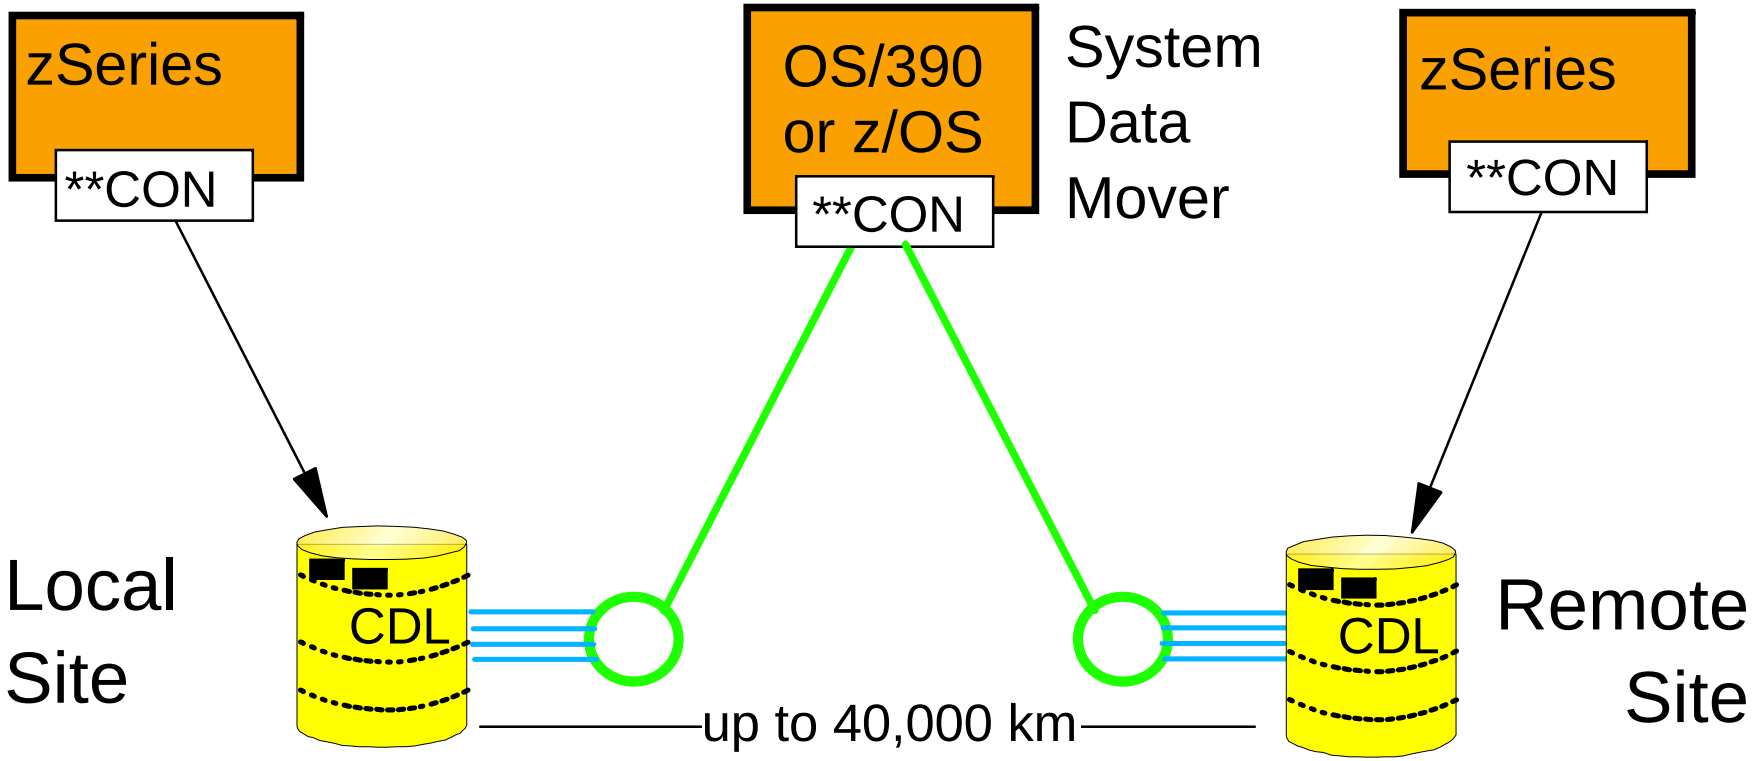

XRC is designed for large-scale remote mirroring of data, from anywhere on the planet, to anywhere on the planet

DDR under z/VM CMS

- DDR from your Linux Guest
  - IPL CMS
  - Attach Tape
  - Run DDR
- DDR from another guest
  - Requires Password and LNKNOPAS authority
  - Attach disks and tape
  - Run DDR

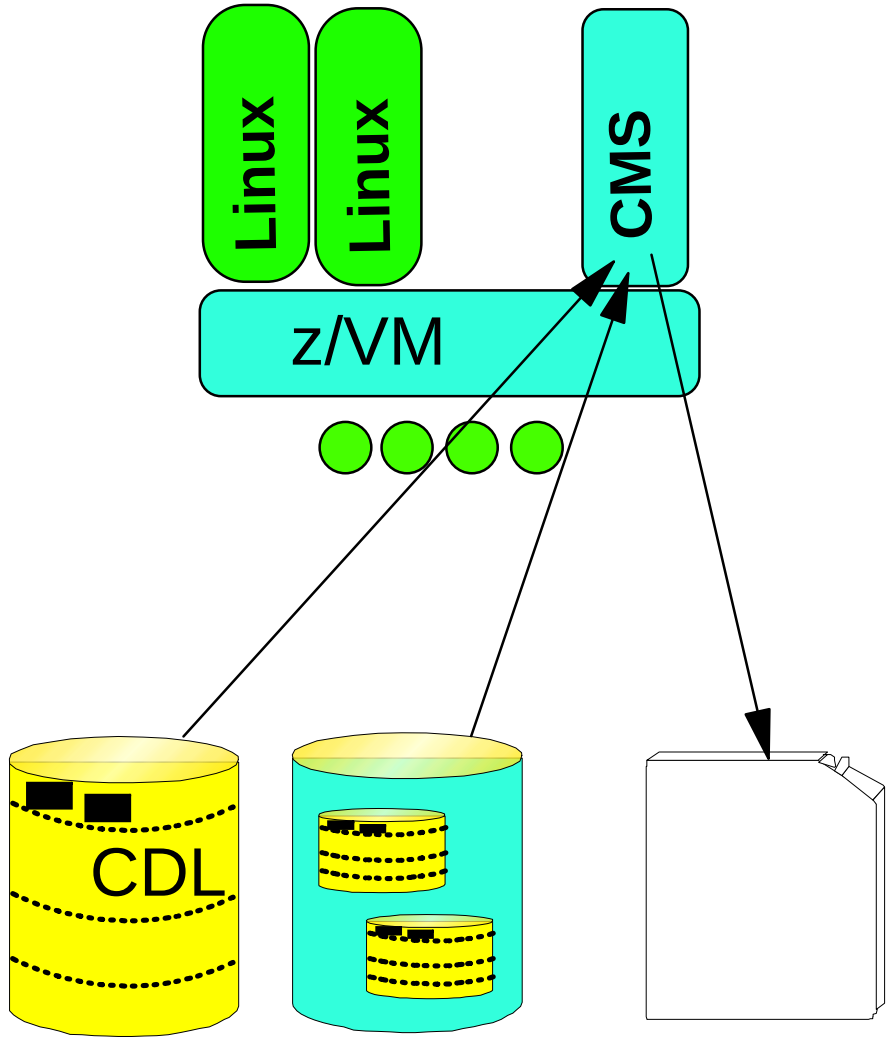

Tivoli Storage Manager (TSM)

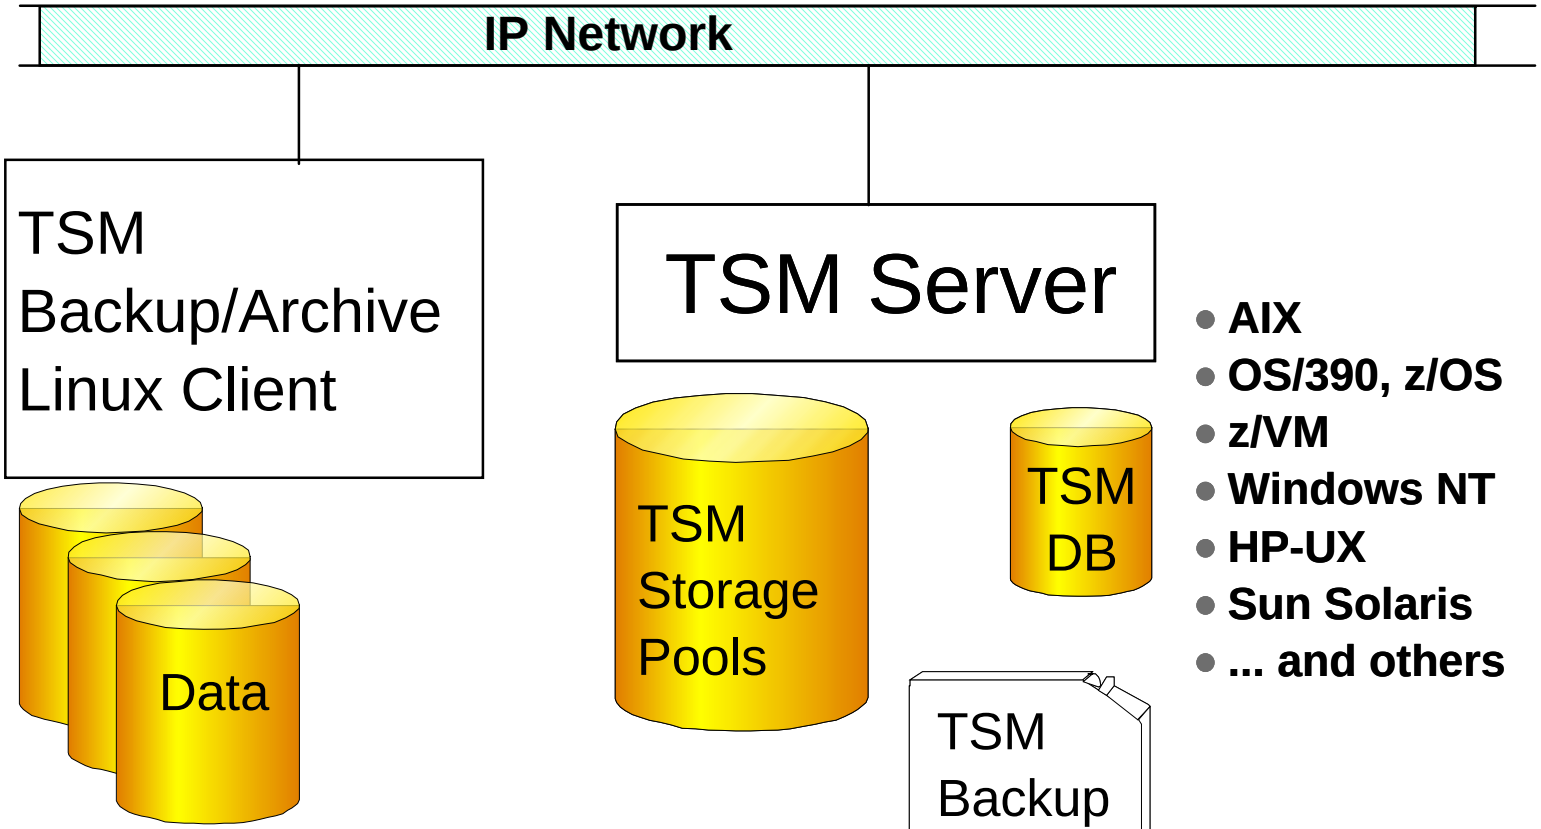

Three ways to run TSM for Linux on S/390

TSM Client on Linux

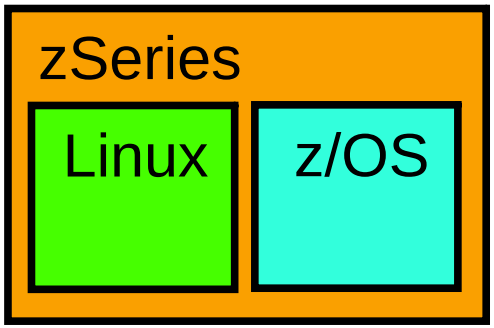

TSM v5.1
Server on z/OS

TCP/IP over
HiperSockets,
CTC, or LAN

TSM Client on Linux

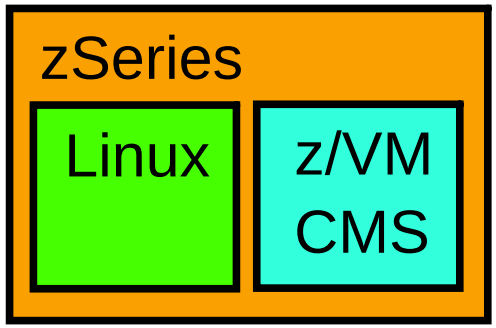

TSM v3.1
Server on z/VM
(CMS)

TCP/IP over
HiperSockets,
IUCV, Guest LAN,
or LAN

TSM Client on Linux

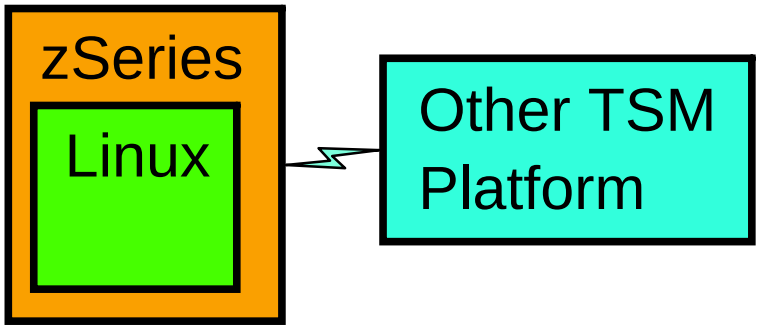

TSM Server on
external server

TCP/IP over LAN

Tivoli Support for Linux

- Linux for zSeries TSM Server:
  - Latest functionality, similar to TSM Server on Linux for Intel
  - Will support FCP-attach Tape only
    - FCP-attachment requires zSeries (z800, z900, z990)
    - Will not support Multiprise 3000, G5, G6
  - Will support both channel-attached and FCP-attached disk

TSM Server on Linux

- Linux for zSeries TSM Server:
  - Latest functionality, similar to TSM Server on Linux for Intel
  - Will support FCP-attach Tape only
    - FCP-attachment requires zSeries (z800, z900, z990)
    - Will not support Multiprise 3000, G5, G6
  - Will support both Channel-attach and FCP-attach disk

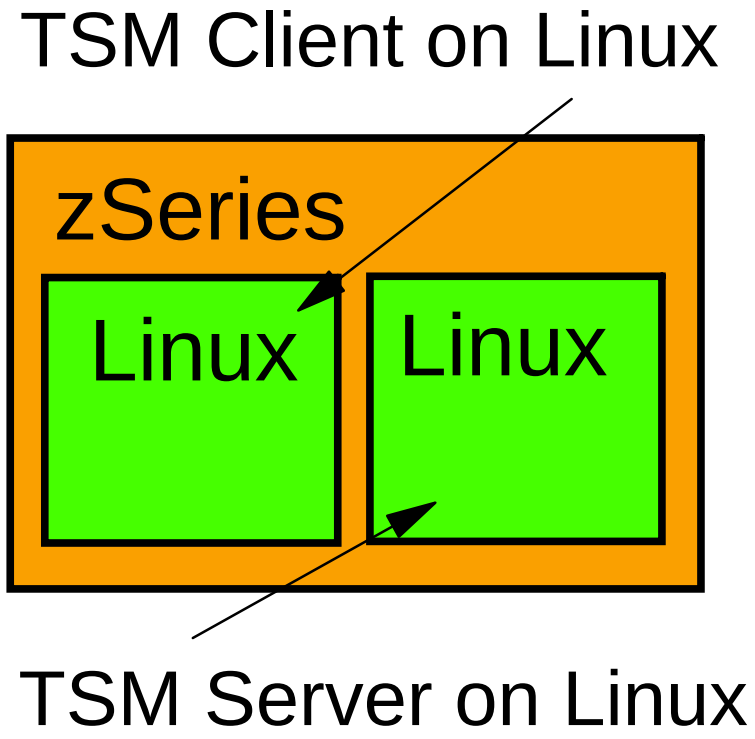

Summary

- IBM is committed to support LINUX across its server platforms
- IBM leads the industry in storage networking based on open, industry standards
- IBM delivers world-class disk and tape storage hardware, storage management software, and integrated solutions

Resources

HOWTO's for DFSMSdss and DFSMShsm!

- Linux
  - http://www.ibm.com/linux
  - Links to IBM Linux & open source sites
- Storage
  - http://www.ibm.com/totalstorage
  - IBM Storage Hardware and Software
  - Specifications
  - System Requirements
  - Host attachments
- Redbooks
  - http://www.ibm.com/redbooks
  - IBM online publications and guide books
  - *Linux with zSeries and ESS: Essentials*, SG24-7025
  - *Implementing Linux with IBM Disk Storage*, SG24-6261
- Tivoli
  - http://www.tivoli.com
  - Overview of IBM's Tivoli software product suite

Disclaimers

Both Linux and Storage are rapidly changing environments.

This information is presented "as is" without any warranty of any kind. Customers are responsible for determining the suitability to their respective environments.

IBM was awarded the 2000 National Medal of Technology for 40 years of innovation in data storage technology. IBM may have patents or pending patent applications covering subject matter in this presentation. The furnishing of this presentation does not give you any license to these patents.

Only a representative subset of the IBM offerings are presented here. Products not mentioned should not be interpreted as a lack or withdrawal of support of those products.

Information concerning non-IBM products was obtained from a supplier of these products, published announcement material, or other publicly available sources. Questions on the capability of non-IBM products should be addressed to the supplier of those products.

Some information in this presentation addresses anticipated future capabilities. Such information is not intended as a definitive statement of a commitment to specific levels of performance, function or delivery schedules with respect to any future products. Such commitments are only made in official IBM product announcements. All statements regarding IBM future direction and intent are subject to change or withdraw without notice, and represent goals and objectives only.

The information is presented here to communicate IBM's current investment and development activities as a good faith effort to help with our customers' future planning. Contact your local IBM business contact for details on specific products, programs or services.