



**SHARE**  
Technology • Connections • Results

Session  
5519  
9226

# Linux and IBM Storage

Tony Pearson  
IBM Storage Systems Group  
SSG Software Strategy and Architecture  
tpearson@us.ibm.com

March 5, 2002

# Trademarks



The following are trademarks or registered trademarks of the International Business Machines Corporation:

AIX, AS/400, AS/400e, CICS, DB2, DB2 Universal Database, e-business (logo), Enterprise Storage Server, the eServer logo, ESCON, FlashCopy, IBM, Intellistation, iSeries, Magstar, Modular Storage Server, MQSeries, Netfinity, NUMA-Q, OS/390, OS/400, Parallel Sysplex, pSeries, RS/6000, S/390, SANergy, Seascape, Sequent, Sequent (logo), SP, SP2, SSA, StorWatch, Thinkpad, Tivoli, Tivoli Storage Manager, Ultrastar, WebSphere, xSeries, zSeries.

Lotus, Lotus Notes and Lotus Domino are trademarks of Lotus Development Corporation.

Microsoft, Windows, Windows NT and the Windows logo are registered trademarks of Microsoft Corporation. Intel and Pentium are registered trademarks of Intel Corporation. UNIX is a registered trademark licensed exclusively through the OPEN group. LINUX is a registered trademark of Linus Torvalds. Java and all Java-based trademarks and logos are trademarks of Sun Microsystems, Inc.

The "National Medal of Technology" is a federally registered trademark and service mark of the United States Department of Commerce.

Red Hat, the Red Hat "Shadow Man" logo, RPM and all Red Hat-related logos are trademarks or registered trademarks of Red Hat, Inc. Caldera Systems, and the C-logo, are trademarks or registered trademarks of Caldera Systems, Inc. Turbolinux and "lightning bolt" logo are registered trademarks of Turbolinux, Inc. SuSE, and SuSE "lizard" logo, are trademarks of SuSE, Inc.

Other company, product and service names may be trademarks or service marks of others.

# 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.

# Agenda



- Linux Overview
- eServer Platforms
- Storage Connectivity Options
- IBM Storage Offerings
  - ▶ Disk Storage
  - ▶ Tape Storage
- Storage Management
- Resources

# Open e-business platform

## Presentation/Access

Portable & Interchangeable

## Data/Documents

Portable & Interchangeable

## Business

Logic

Portable

&


Interchangeable

Development

& Operating

Environment

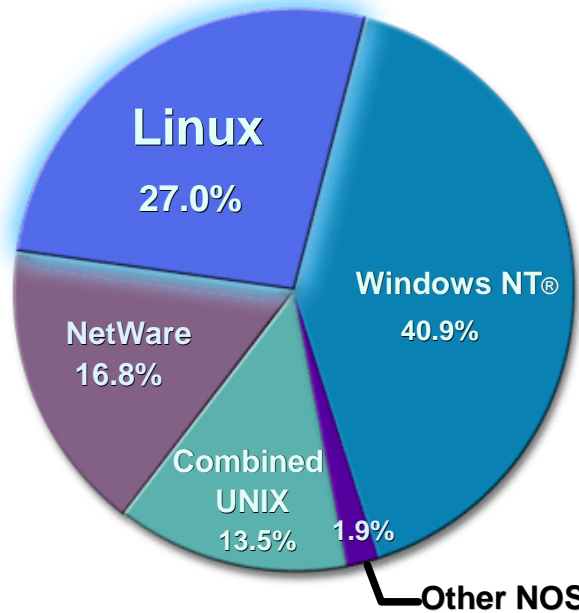
Open & Portable



***An Integrated Standards and Open Source Model***

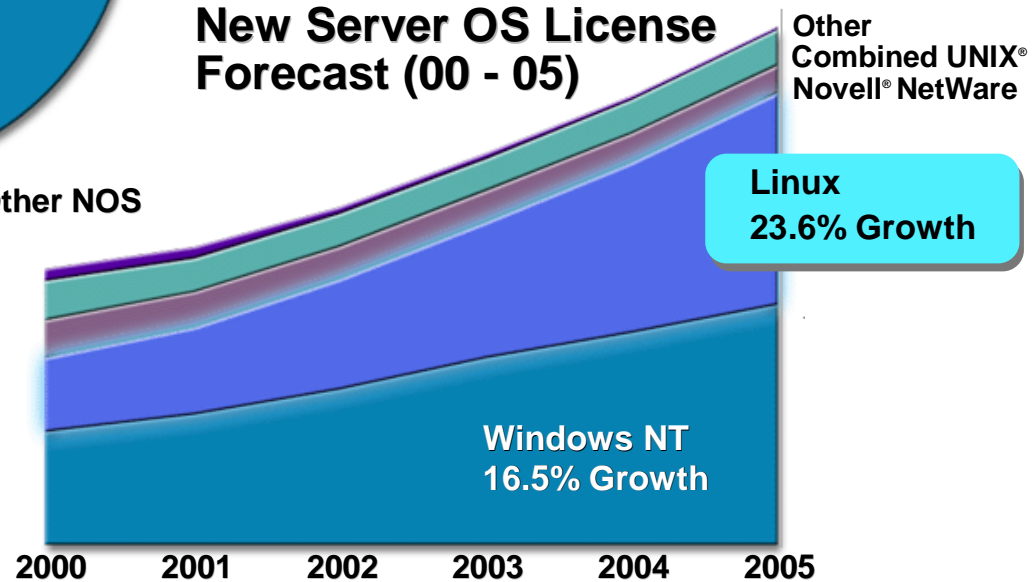
# Linux Momentum Building

## 2000 New Server OS Shipments



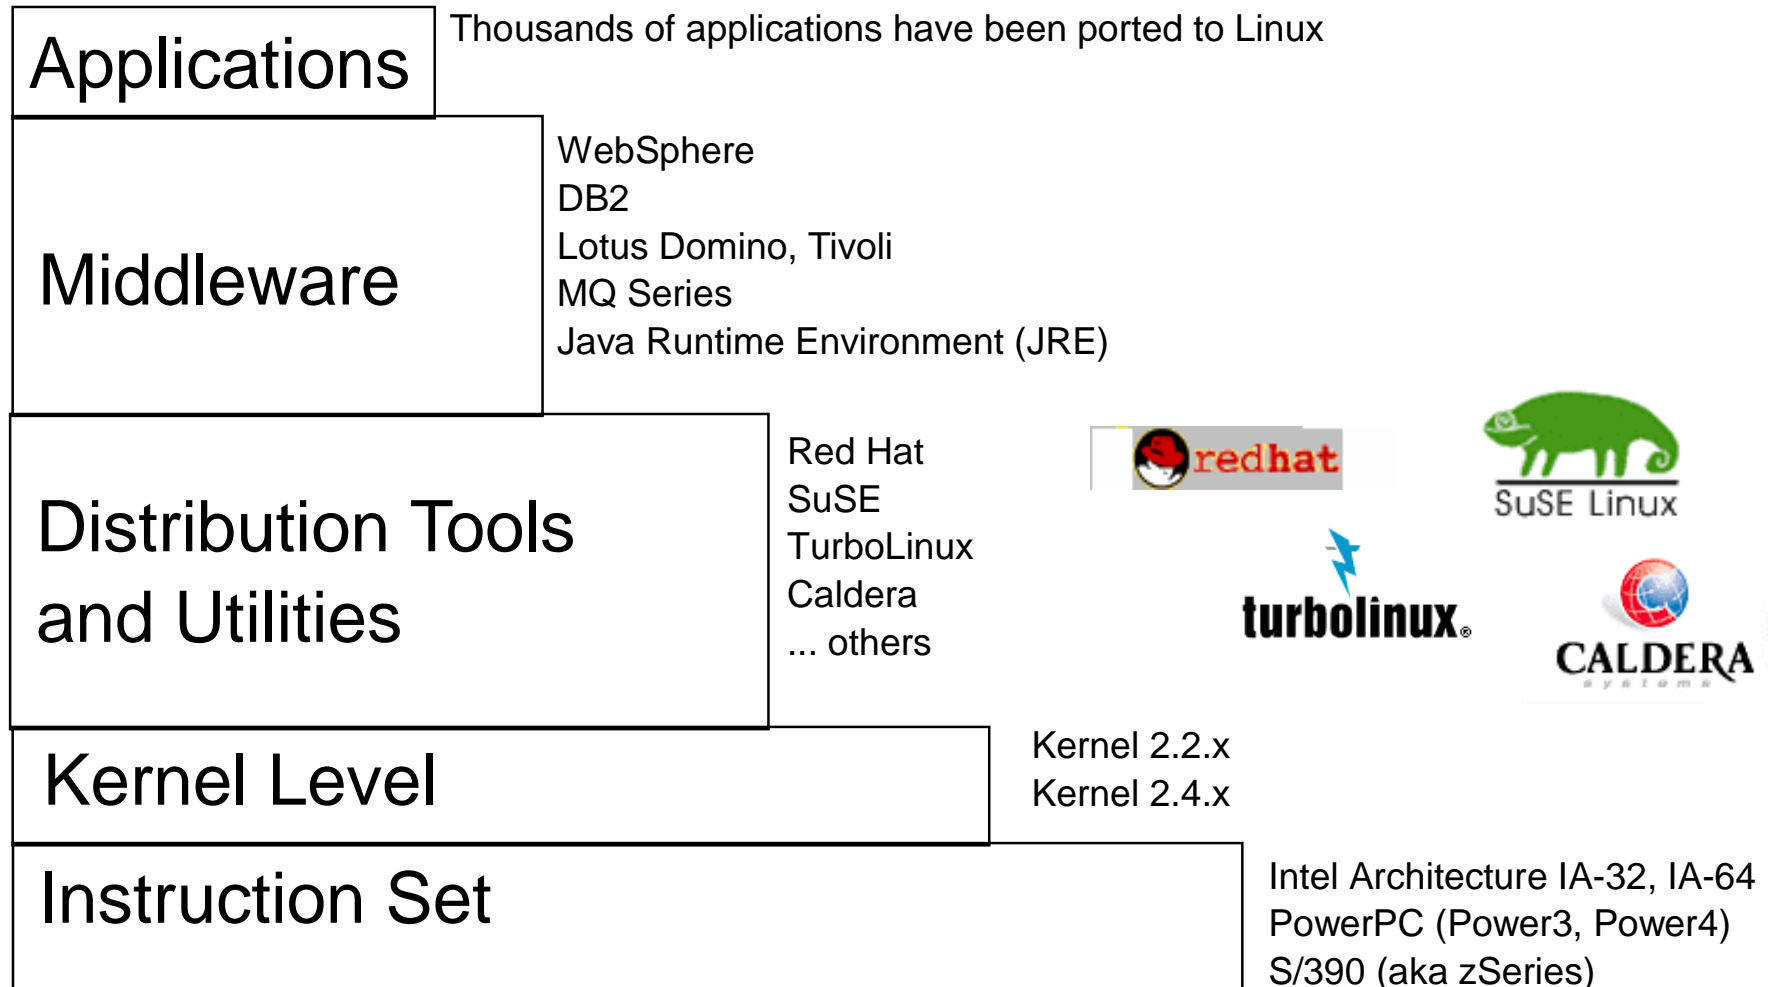
- 1999 - 2000 Linux shipments grew 24% YTY
- 2000: #2 volume OS
- Linux shipments projected to have highest growth (99 - 05)

## New Server OS License Forecast (00 - 05)



Source: IDC, 2001

# Linux Terminology explained



# IBM eServer xSeries



## Linux available for xSeries:

### ✓ *IBM Netfinity servers*


- Rack-optimized - More computing power per square foot
- X-architecture brings enterprise-class scalability, reliability & security to Intel-based servers
- Certified under the IBM ServerProven Program for *Caldera, Red Hat, SuSE and TurboLinux*
- Preloads for Caldera, Red Hat and TurboLinux

### ✓ *IBM Thinkpads*

- All 4 distributions supported

### ✓ *IBM Intellistations*

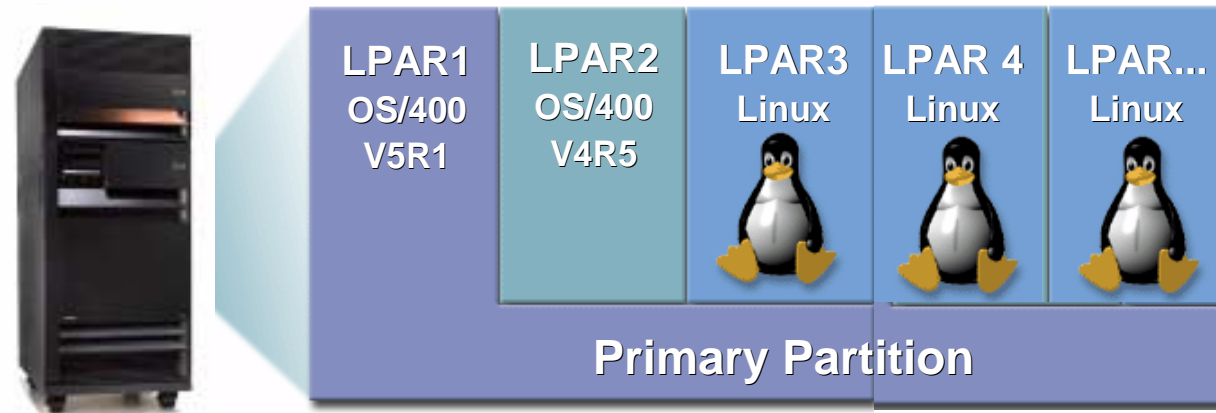
- All 4 distributions supported





# Linux on eServer iSeries

## The Point of Coexistence - Where Linux Complements Integrated e-business Solutions



### iSeries support of Linux provides:

- Linux runs in iSeries logical partition up to 31 Linux servers on one iSeries
- Flexible allocation of process, memory and I/O resources to applications
- iSeries scalability, availability and ease-of-use for Linux applications

# IBM eServer pSeries



Linux available for pSeries:

✓ **IBM RS/6000 and pSeries servers**

- Industry-leading clustering
- Rugged, high-density, rack-mounted systems
- RS/6000 43P-150, B50 "Pizzazz", F50
- pSeries 640, 680 and 690 "Regatta"

✓ **Linux for PowerPC Distributions**

- SuSE
- Red Hat
- TurboLinux (under development)


✓ **AIX Affinity for Linux Strategy**

- AIX Toolbox for Linux for AIX 4.3.3 and 5L
- "Linux GNU" tools, utilities, RPM for AIX
- Integrated Linux APIs into AIX lib.c


✓ **Partitioning Support**

- pSeries 690 allows up to 16 partitions

RPM = Red Hat Package  
Manager



# IBM eServer zSeries



Linux available today for:

✓ ***S/390 and zSeries family***

- World's most scalable server
- z800 and z900 models supported
- Bullet-proof reliability
- Dynamic workload management
- DB2, WebSphere, MQSeries, Java, Tivoli

✓ ***Linux Distributions***

- Linux for S/390 for 32-bit architecture
- Linux for zSeries for 64-bit architecture } Linux on S/390
- Linux distributions from: SuSE, TurboLinux, Red Hat

✓ ***Variety of Operational Modes***

- Native
- Logical Partition (LPAR)
  - Integrated Facility for Linux (IFL)
- Guests under VM, z/VM
  - z/VM 4.2 can run on IFL engines



# Storage Connectivity Options



- Positioning
  - ▶ Storage Area Network (SAN)
  - ▶ Block-oriented protocol (iSCSI)
  - ▶ File-oriented protocols (NAS)
  
- Storage Protocol comparison
  
- Connectivity Options

# SAN, iSCSI and NAS positioning

## SAN


- Block I/O optimization
- High Performance & Scalability
- Fibre Channel protocols
- Longer design and installation time
- Typically more expensive
- Requires significant IT support

## iSCSI

- Block I/O optimization
- Better Performance & Scalability than file-based solutions
- Ethernet connections - TCP/IP protocols
- Self-contained and quick to install
- Less expensive than SAN storage

## NAS

- File I/O optimization
- Performance & Scalability Considerations
- Requires little IT support
- Ethernet connections - TCP/IP protocols
- Self-contained and quick to install
- Typically less expensive

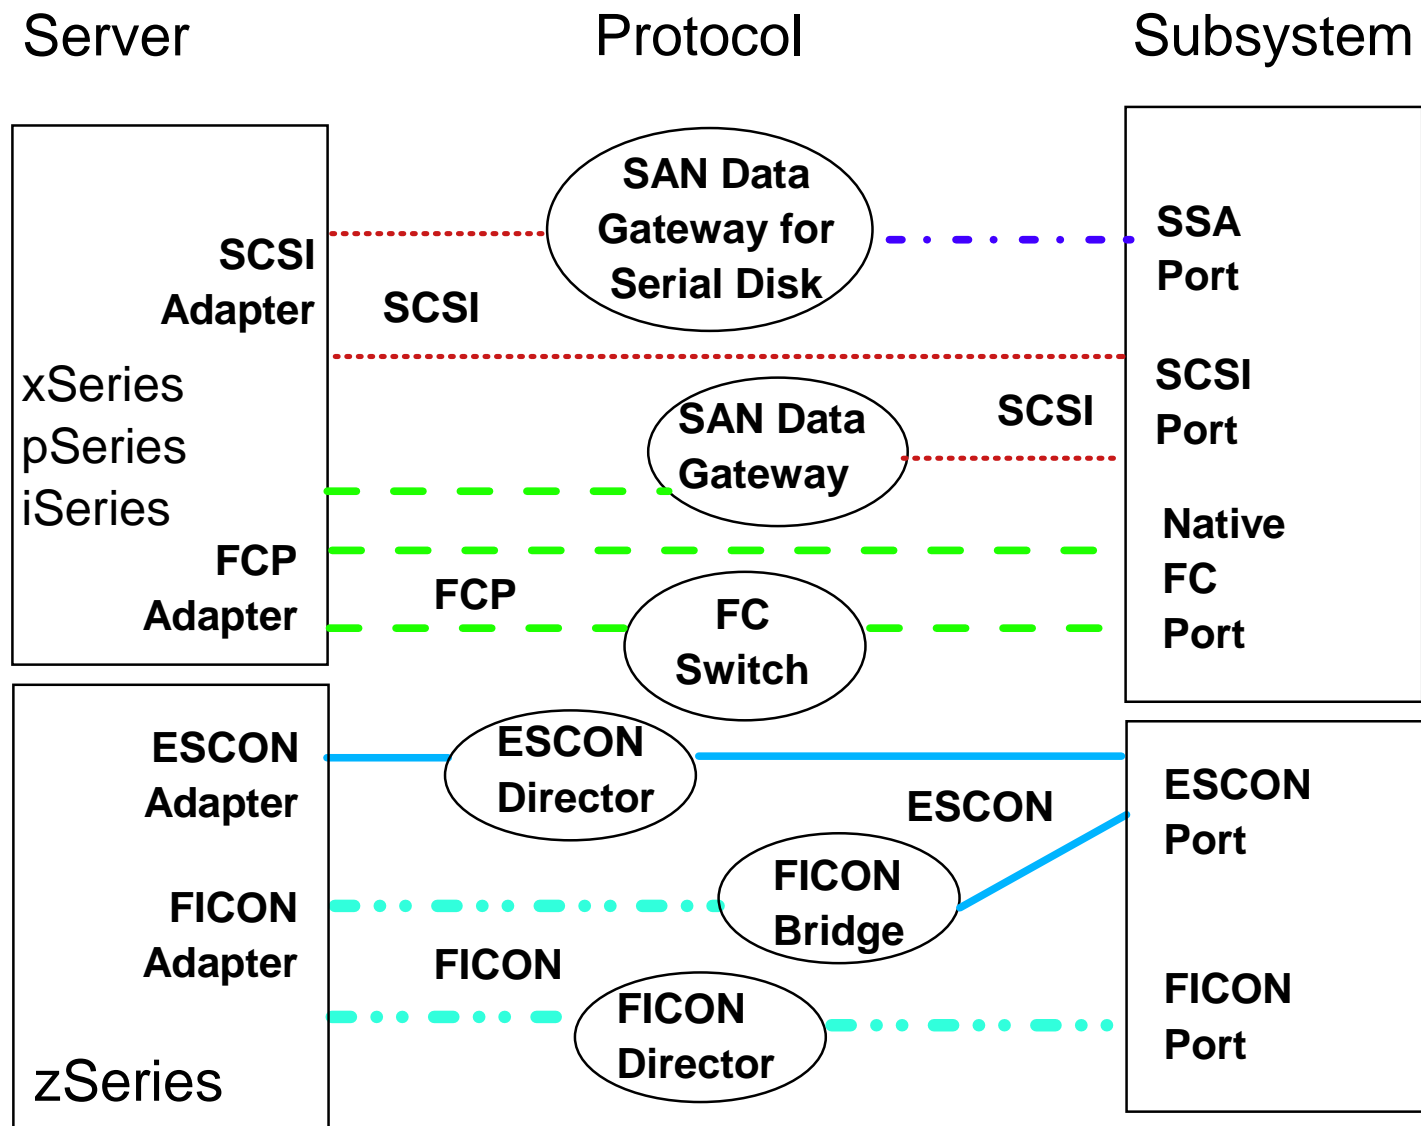


# Storage Protocols

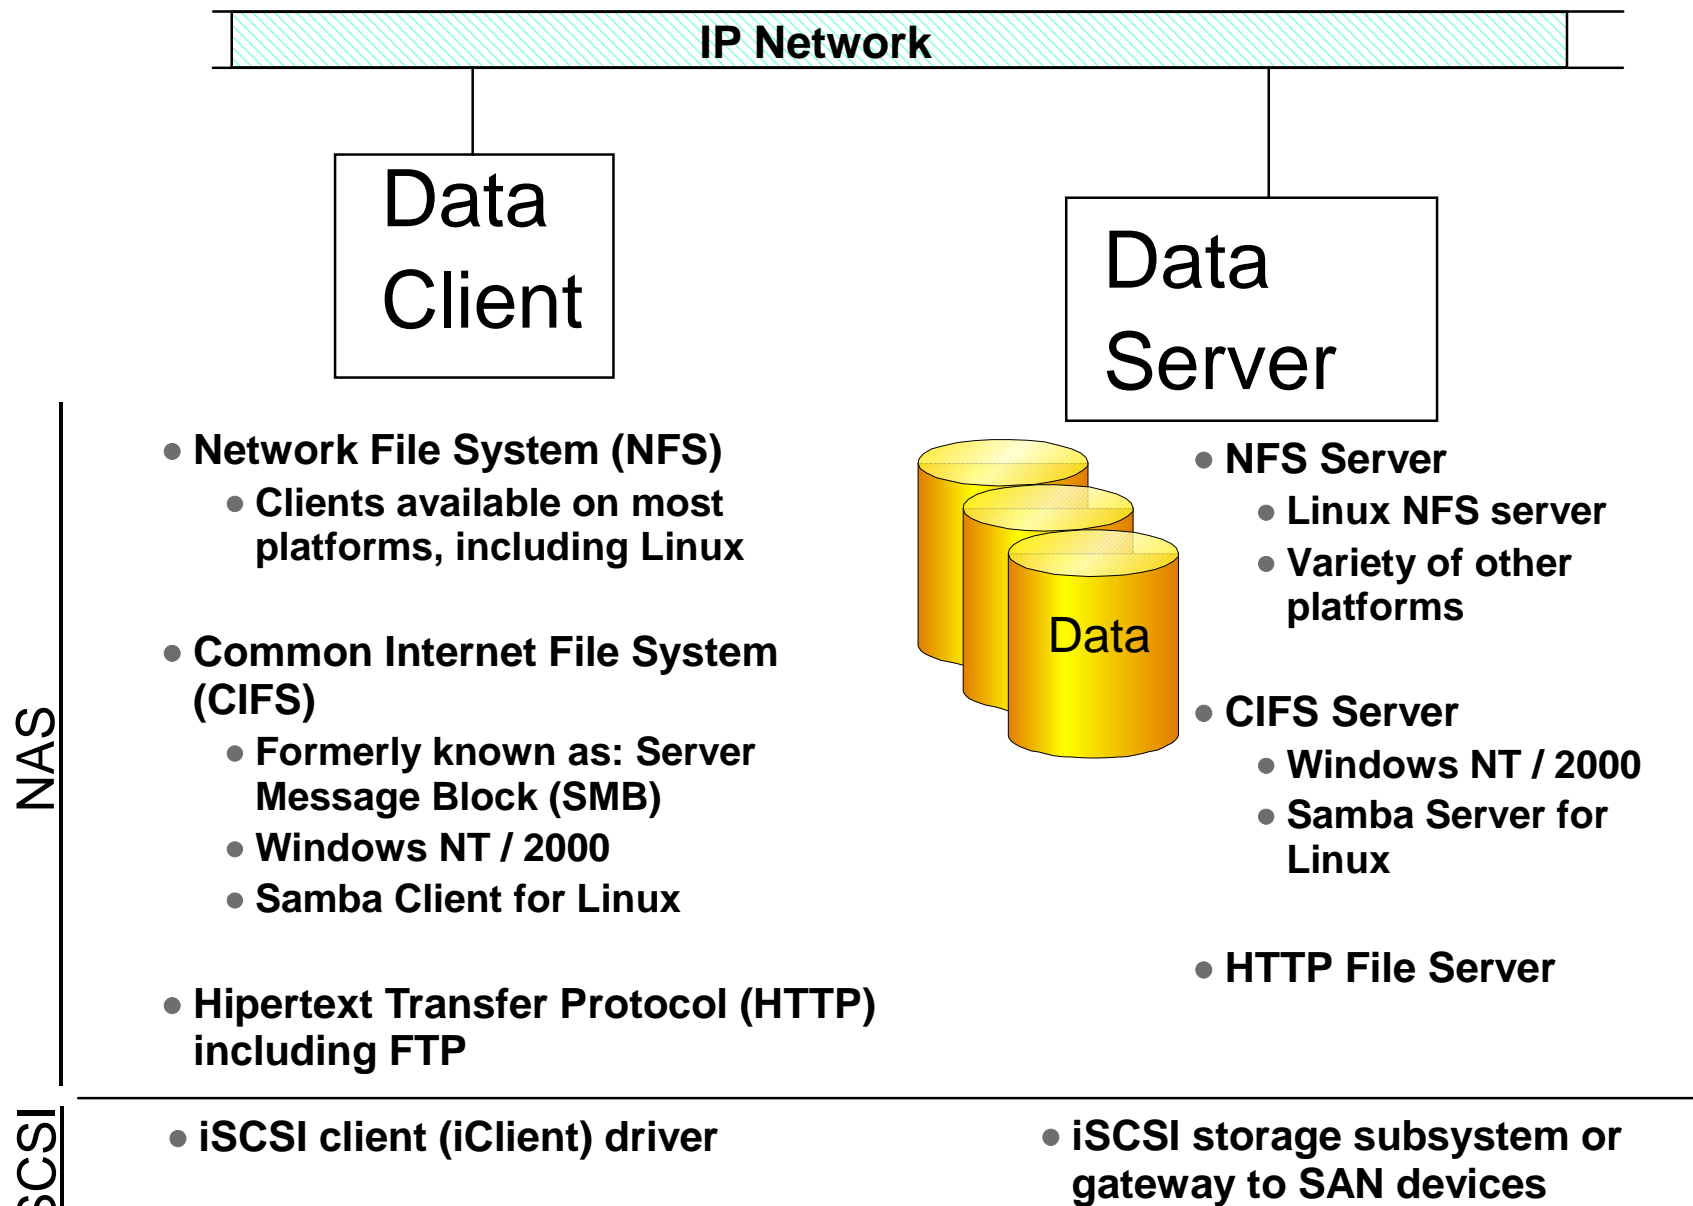


- Small Computer System Interface (SCSI)
  - ▶ Industry standard parallel interface
  - ▶ Ultra, Ultra 2, Ultra 3, and SPI-4
- Serial Storage Architecture (SSA)
  - ▶ Industry standard serial interface
- Fibre Channel Protocol (FCP) SCSI over Fibre Channel
  - ▶ Point-to-Point
  - ▶ Arbitrated Loop
  - ▶ Switched Fabric
- Enterprise Storage Connection (ESCON)
  - ▶ Enhanced Count-Key-Data (ECKD) for S/390, zSeries
- Fibre Connection (FICON)
  - ▶ ESCON protocol over Fibre Channel

# Storage Network Connectivity



# IP Network Connectivity





# IBM Storage Offerings



- Questions for customers to ask a storage vendor
  
- IBM Disk Storage
  - ▶ Expandable Storage Plus
  - ▶ Fibre Array Storage Technology (FAStT)
  - ▶ Enterprise Storage Server (ESS)
  - ▶ IBM IP Storage 200i
  
- IBM Tape Storage
  - ▶ Linear Tape-Open (LTO) Ultrium Drives and Libraries
  - ▶ Magstar Tape Drives and Libraries



# Questions to Ask a Storage Vendor



- What host attachment ports are available on your storage subsystem?
  - ▶ **SCSI, FC-AL, Native (switched) FC, ESCON/FICON**
- Which host bus adapters have you tested with?
  - ▶ **There are a variety for each platform**
- Which distributions of Linux have these drivers been tested with?
  - ▶ **TurboLinux, SuSE, Red Hat, Caldera, etc.**
- Which levels of Linux kernels are supported by these drivers?
  - ▶ **Newer 2.4.x level has better Error Recovery Process (ERP) than the 2.2.x level**
- Will your device work with my existing server, adapter, driver, distribution and kernel level?
  - ▶ **This information will help evaluate devices**

# 2104 Expandable Storage Plus (EXP Plus)



## Host Attachments available:

- ✓ SCSI-2, Ultra2 SCSI, and Ultra3 SCSI

## Powerful Features:

- ✓ Rack-mounted drawer or Tower
- ✓ Multiple RAID levels supported when used v with Ultra3 SCSI RAID attachments
- ✓ Scalability:  
Up to 509GB per tower or drawer
- ✓ Redundant, Hot-swappable components



## Supported Platforms:

- ✓ AIX 4.2.1 and above
- ✓ Linux applications through AIX 5L

# Fibre Array Storage Technology (FAStT)



## Host Attachments available:

- ✓ Fibre Channel
  - Point-to-Point, Arbitrated Loop (FC-AL)
  - Switched Fabric

## Powerful Features:

- ✓ Rack-mounted
- ✓ Multiple RAID levels supported
- ✓ Scalability / FC Support
  - FAStT200: 36GB to 4.3TB / 1 Gb
  - FAStT500: 36GB to 16TB / 1 Gb
  - FAStT700: 36GB to 16TB / 2 Gb
- ✓ Redundant, Hot-swappable components

## Supported Platforms:

- ✓ FAStT Storage Manager v7.1 and 8.0 on Linux-Intel Red Hat 7.1
- ✓ Windows, HP-UX, Sun Solaris, NetWare, Linux-Intel Red Hat 7.1 (2.4.x kernel)
- ✓ FAStT500 also supports AIX
- ✓ Clustering for MC/Service Guard, Veritas, and MSCS



FAStT200

FAStT EXP 500



FAStT500

FAStT700

# Enterprise Storage Server (ESS)



## Host Attachments available:

- ✓ UltraSCSI
- ✓ FC-AL and Switched Fibre Channel
- ✓ ESCON and FICON

## Powerful Features:

- ✓ JBOD and RAID5 levels supported
- ✓ Subsystem Device Driver (SDD)
- ✓ Peer-to-Peer Remote Copy (PPRC)
- ✓ FlashCopy
- ✓ Scalable from 420GB to 22TB useable capacity
- ✓ Full fault tolerance and redundancy

## Command Line Interface:

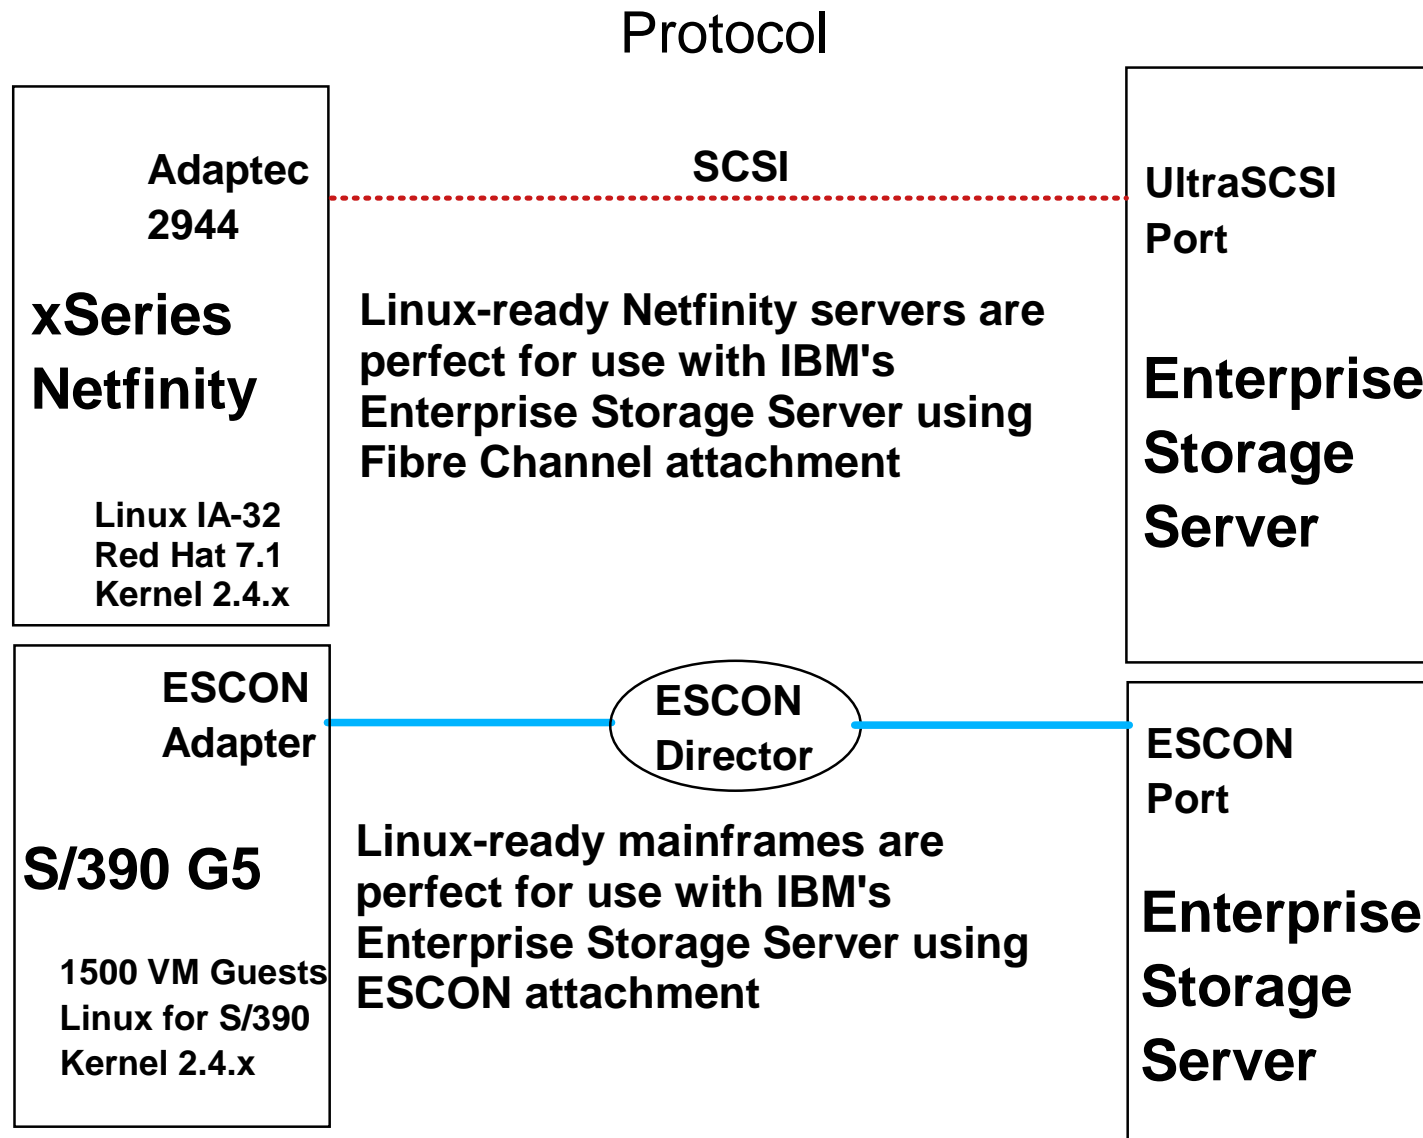
- ✓ Requires Java 1.1.8
- ✓ Available for AIX, Windows, Sun Solaris, HP-UX, and Linux!

## Linux Support:

- ✓ zSeries: SuSE, TurboLinux, Red Hat (2.4.x kernel), ESCON
- ✓ xSeries: Linux-Intel Red Hat 7.1, SuSe 7.2, FC on F10/F20 models
- ✓ xSeries: E10/E20 models or SCSI attachment - RPQ



# Example Configurations



# iSCSI - Block-oriented protocol

New standard jointly developed between IBM and Cisco

## ■ IBM IP Storage 200i

- ▶ Scalable pooled storage appliance: 108GB - 3.5TB
- ▶ High availability redundancies
- ▶ Optimized for block I/O database access
- ▶ Easy integration into existing network
- ▶ Supports Windows NT/2000
- ▶ Supports Red Hat 6.2 Linux (2.2.x)
- ▶ Linux on S/390 being tested



**IP Storage 200i**

## ■ Cisco SN 5420 Storage Router

- ▶ Extend SAN to IP infrastructure
- ▶ Designed for block I/O
- ▶ Use existing IP network skills
- ▶ Supports Windows NT, Sun, Linux (2.2.x)



**Cisco SN 5420 Storage Router**

# Linear Tape-Open - Ultrium

**New standard in tape jointly developed between IBM, Hewlett-Packard and Seagate**

## **Ultrium Tape Drive 3580:**

- 100 GB native capacity cartridge
- 15 MB/sec native data rate
- Connectivity: UltraSCSI and FC

## **Ultrium Tape Autoloader 3581:**

- Seven cartridge autoloader (700GB native capacity)
- Connectivity: UltraSCSI and FC

## **Ultrium Scalable Tape Library 3583:**

- 1-6 drives, 18-72 cartridges
- 7.2 TB native capacity
- FC connectivity via internal SAN Data Gateway module

## **UltraScalable Tape Library 3584:**

- 1-72 drives, up to 2,481 cartridges
- 248 TB native capacity
- DLT frame option, 1-12 drives, up to 360 cartridges
- Native FC LTO drives



3580 Tape Drive



3581 Tape Autoloader



Scalable Tape Library



UltraScalable Tape Library

Linux-Intel support

- ✓ Red Hat 7.1  
(kernel 2.4.x)



# Magstar Tape Drive

## 3590 B11/B1A tape drive

- 10 or 20 GB native capacity cartridge
- 9 MB/sec native data rate
- Connectivity
  - ✓ Ultra SCSI with dual data ports
  - ✓ FCP via San Data Gateway
  - ✓ ESCON

## 3590 E11/E1A tape drives

- 20 or 40 GB native capacity cartridge
- 14 MB/sec native data rate
- Investment Protection
  - ✓ Field upgradable from B1X Models
  - ✓ Media Reuse
- Connectivity
  - ✓ Ultra SCSI with dual data ports
  - ✓ FCP via San Data Gateway to SCSI drives
  - ✓ Native Fibre Channel
  - ✓ ESCON



3590 Rack-Mount



3494 Tape Library

**Linux-Intel support Red Hat 7.1  
(kernel 2.4.x) for Ultra SCSI**

**LinuxPPC for iSeries LPAR  
Red Hat 7.1 SCSI or ESCON**

**Linux on S/390 support  
ESCON emulating 3490E**

# Storage Management



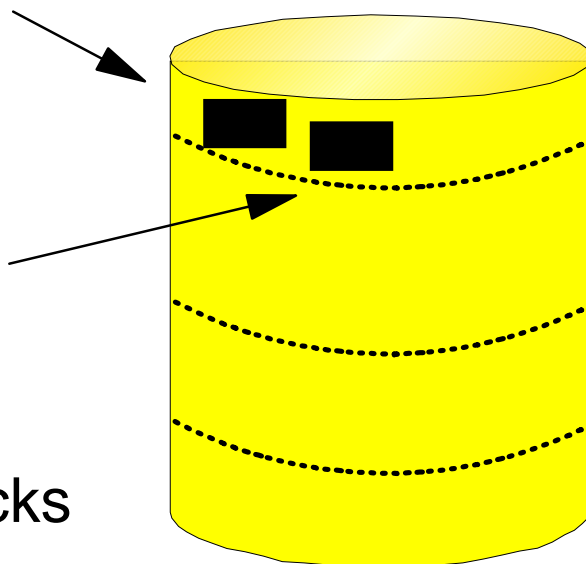
- Storage Management can provide backup, archive and space management facilities for your data
  
- Linux on S/390
  - Compatible Disk Layout (CDL)
- Tivoli Storage Manager
  - ▶ Client and Server platforms supported
  - ▶ Client Requirements for Linux IA-32
  - ▶ Network Storage Manager
- Tivoli SANergy

# 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 2.4.5


# Linux on S/390 Compatible Disk Layout

## Linux view:

/dev/dasd/019b/part1  
mounted as /programs →

/dev/dasd/019b/part2  
mounted as /data →

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



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

## z/OS view:

Label: VOL1VOL1APP037  
VTOC: 12 DSCB's


LINUX.VAPP037.PART0001.NATIVE

LINUX.VAPP037.PART0002.NATIVE

LINUX.VAPP037.PART0003.SWAP

Kernel 2.4.5

# Linux on S/390 Compatible Disk Layout




Linux volumes can be varied online to OS/390 and z/OS

## DFSMSdss Support

- COPY FULL
- DUMP
  - FULL
  - physical DATASET
- RESTORE
  - FULL
  - physical DATASET
- PRINT
- COPYDUMP

Kernel 2.4.5

# Linux on S/390 Compatible Disk Layout



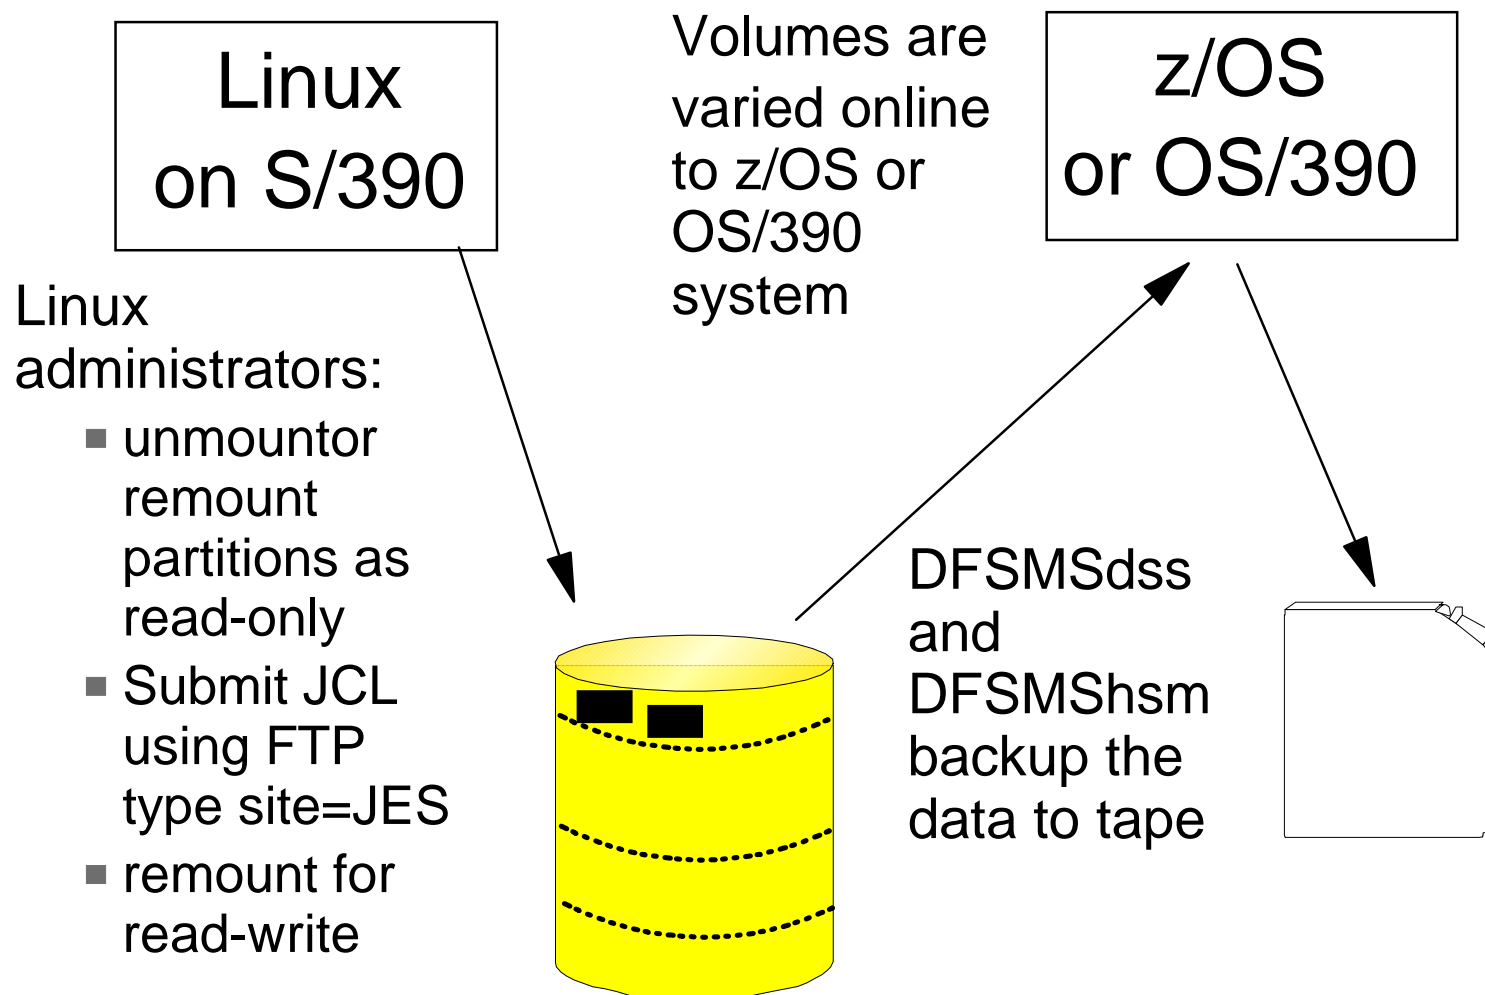
Linux volumes can be varied  
online to OS/390 and z/OS

## DFSMSHsm Support

- ADDVOL PRIMARY  
NOAUTORECALL  
NOAUTOBACKUP  
NOAUTOMIGRATE
- Automatic DUMP
- BACKVOL DUMP
- RECOVER FROMDUMP

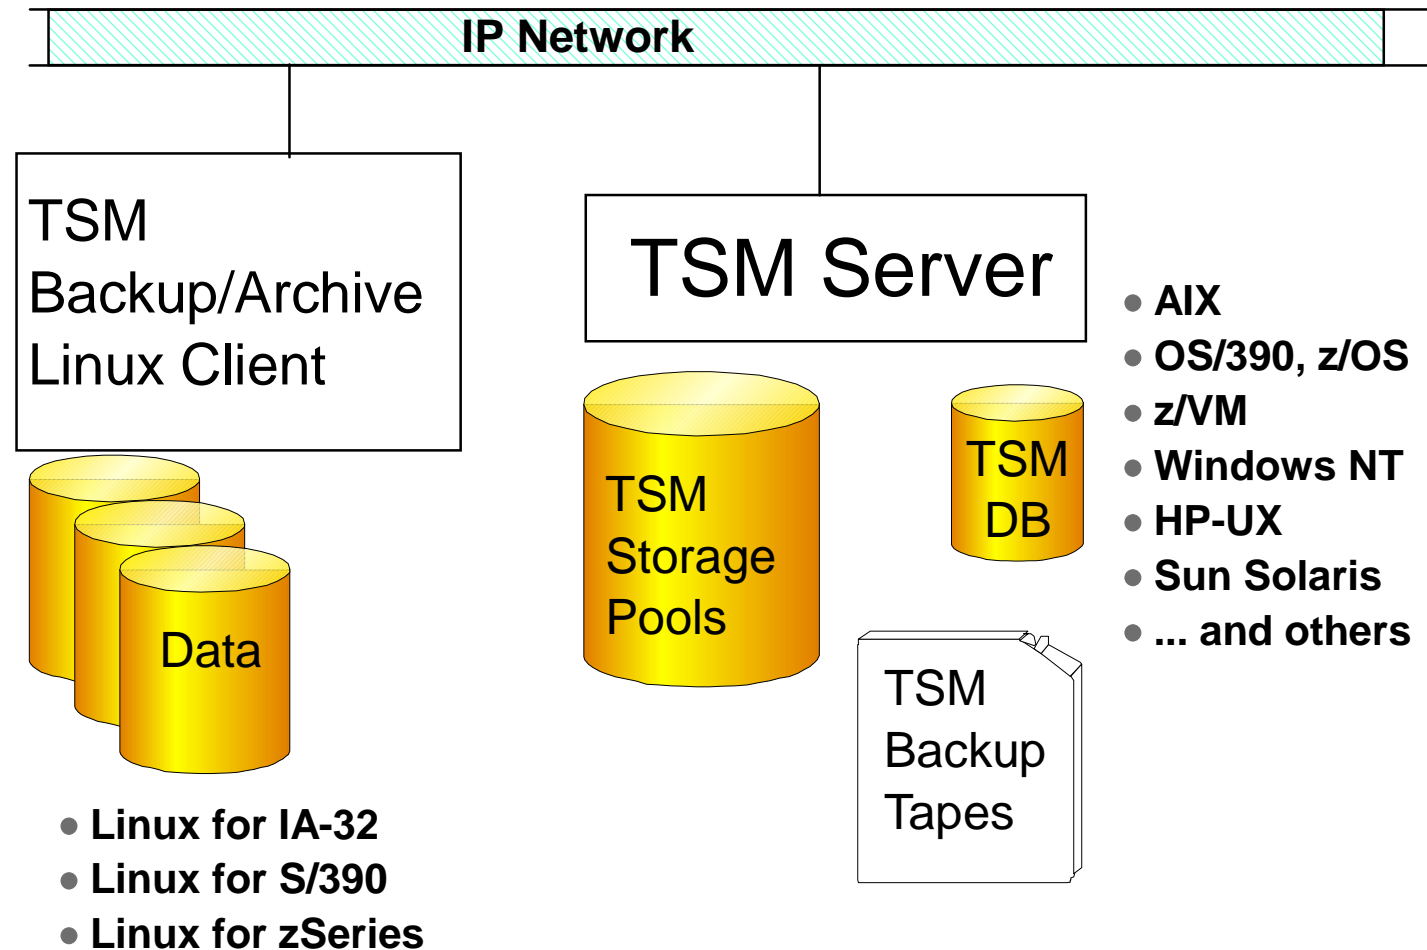
Kernel 2.4.5

# Linux on S/390 Compatible Disk Layout



Kernel 2.4.5

# Tivoli Storage Manager (TSM)





# Tivoli Storage Manager: Linux-Intel



## ■ IA-32 TSM features

- ▶ Backup/Archive client
- ▶ Admin client
- ▶ Web client
- ▶ API

## ■ Requirements

- ▶ Linux kernel 2.2.x or higher
- ▶ RPM package
- ▶ Certified distributions known to provide the base requirements:
  - Red Hat 6.1, SUSE 6.3
  - TurboLinux 6.0, Caldera 2.3



IA-32 = Intel Architecture 32-bit

# 3466 Network Storage Manager



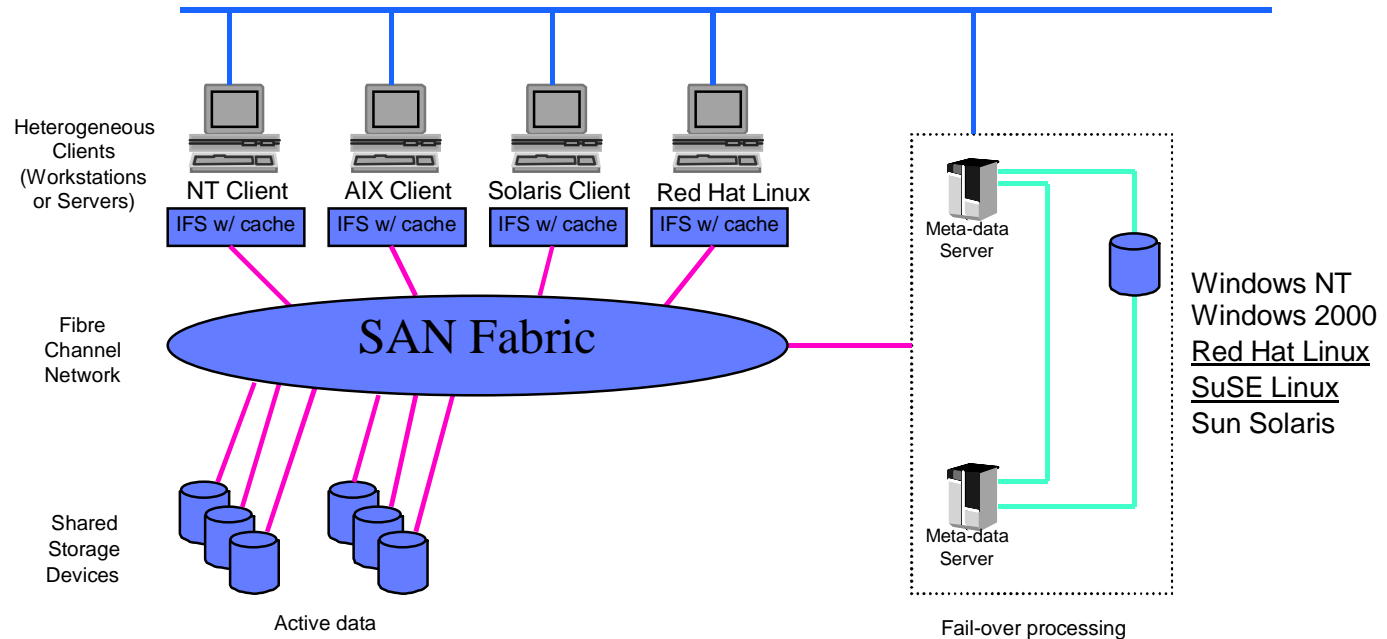
**A sophisticated Data Protection solution**

**Integrated hardware and software to provide a network/data protector**

- **Lights-out operation**
- **Centralizes Backup/Recovery**
- **Enterprise-wide Administration, Monitoring and Reporting**
  - ✓ **Tivoli Storage Manager queries**
  - ✓ **Tivoli Storage Manager accounting records**
- **Attaches to tape library**



# Tivoli SANergy File Sharing



- File sharing capability of a LAN with the speed of a SAN
- Metadata server owns and formats the data
  - ▶ Metadata platforms: NY, W2K, Sun, Red Hat Linux-Intel
- SAN-based file sharing among heterogeneous servers
  - ▶ NT, W2K, Sun, AIX, MacOS, IRIX, Red Hat Linux-Intel, TRU64, DG/UX
- Improved performance vs. NAS filers

# Resources...



- [www.ibm.com/linux](http://www.ibm.com/linux)
  - ▶ Links to IBM Linux & open source sites
- [www.ibm.com/storage](http://www.ibm.com/storage)
  - ▶ IBM Storage Hardware and Software
  - ▶ Specifications
  - ▶ System Requirements
  - ▶ Host attachments
- [www.ibm.com/redbooks](http://www.ibm.com/redbooks)
  - ▶ IBM online publications and guide books
  - ▶ Implementing Linux with IBM Disk Storage, SG24-6261
  - ▶ Implementing IBM LTO in Linux and Windows Environments, SG24-6268
- [www.tivoli.com](http://www.tivoli.com)
  - ▶ Overview of Tivoli software product suite

HOWTO's for  
DFSMsds and  
DFSMShsm!

# 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**