z/VM: The Value of zSeries Virtualization Technology for Linux

BayBunch Regional User Group
San Francisco, California

September 2002

Reed A. Mullen
mullenra@us.ibm.com
zSeries Development
Endicott, New York

Trademarks

IBM @server zSeries

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

CICS*
DB2
DB2 Connect
DB2 Universal Database
e-business logo*
FICON
HiperSockets
IBM*
IBM logo
MQSeries*
Multiprise*
OS/390
RISC
S/390
S/390 Parallel Enterprise Server*
Virtual Image Facility
VM/ESA*
VSE/ESA
WebSphere
z/OS
z/VM
zSeries

* Registered trademarks of the IBM Corporation

The following are trademarks or registered trademarks of other companies.

Lotus, Notes, and Domino are trademarks or registered trademarks of Lotus Development Corporation.
Tivoli is a trademark of Tivoli Systems Inc.
Linux is a registered trademark of Linus Torvalds.
Java and all Java-related trademarks and logos are trademarks of Sun Microsystems, Inc., in the United States and other countries
UNIX is a registered trademark of The Open Group in the United States and other countries.
Microsoft, Windows and Windows NT are registered trademarks of Microsoft Corporation.

Notes:

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.

IBM considers a product “Year 2000 ready” if the product, when used in accordance with its associated documentation, is capable of correctly processing, providing and/or receiving date data within and between the 20th and 21st centuries, provided that all products (for example, hardware, software and firmware) used with the product properly exchange accurate date data with it. Any statements concerning the Year 2000 readiness of any IBM products contained in this presentation are Year 2000 Readiness Disclosures, subject to the Year 2000 Information and Readiness Disclosure Act of 1998.

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.

Topics

IBM @server zSeries

Virtualization Technology and z/VM™: the Basics

Integrated Facility for Linux (IFL) and z/VM

The Value of z/VM for Linux®

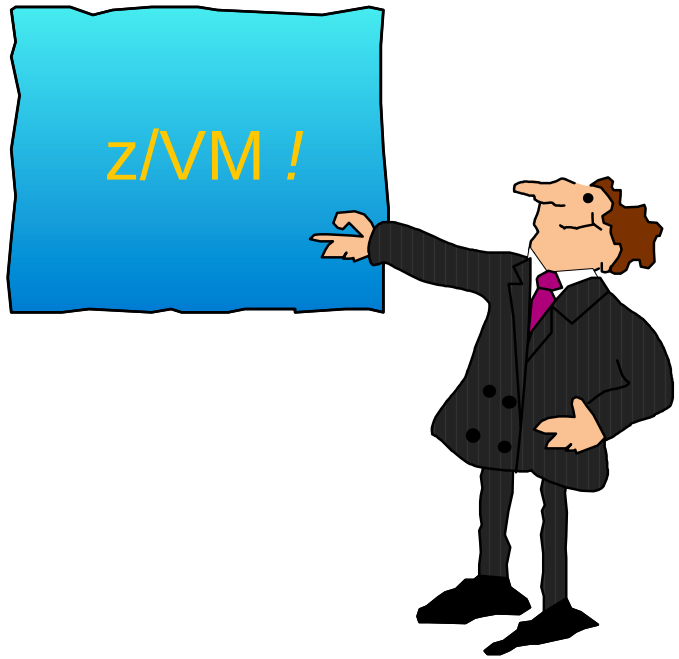

IBM Virtualization Technology Evolution

IBM @server zSeries

The Virtual Machine Concept is not new for IBM®

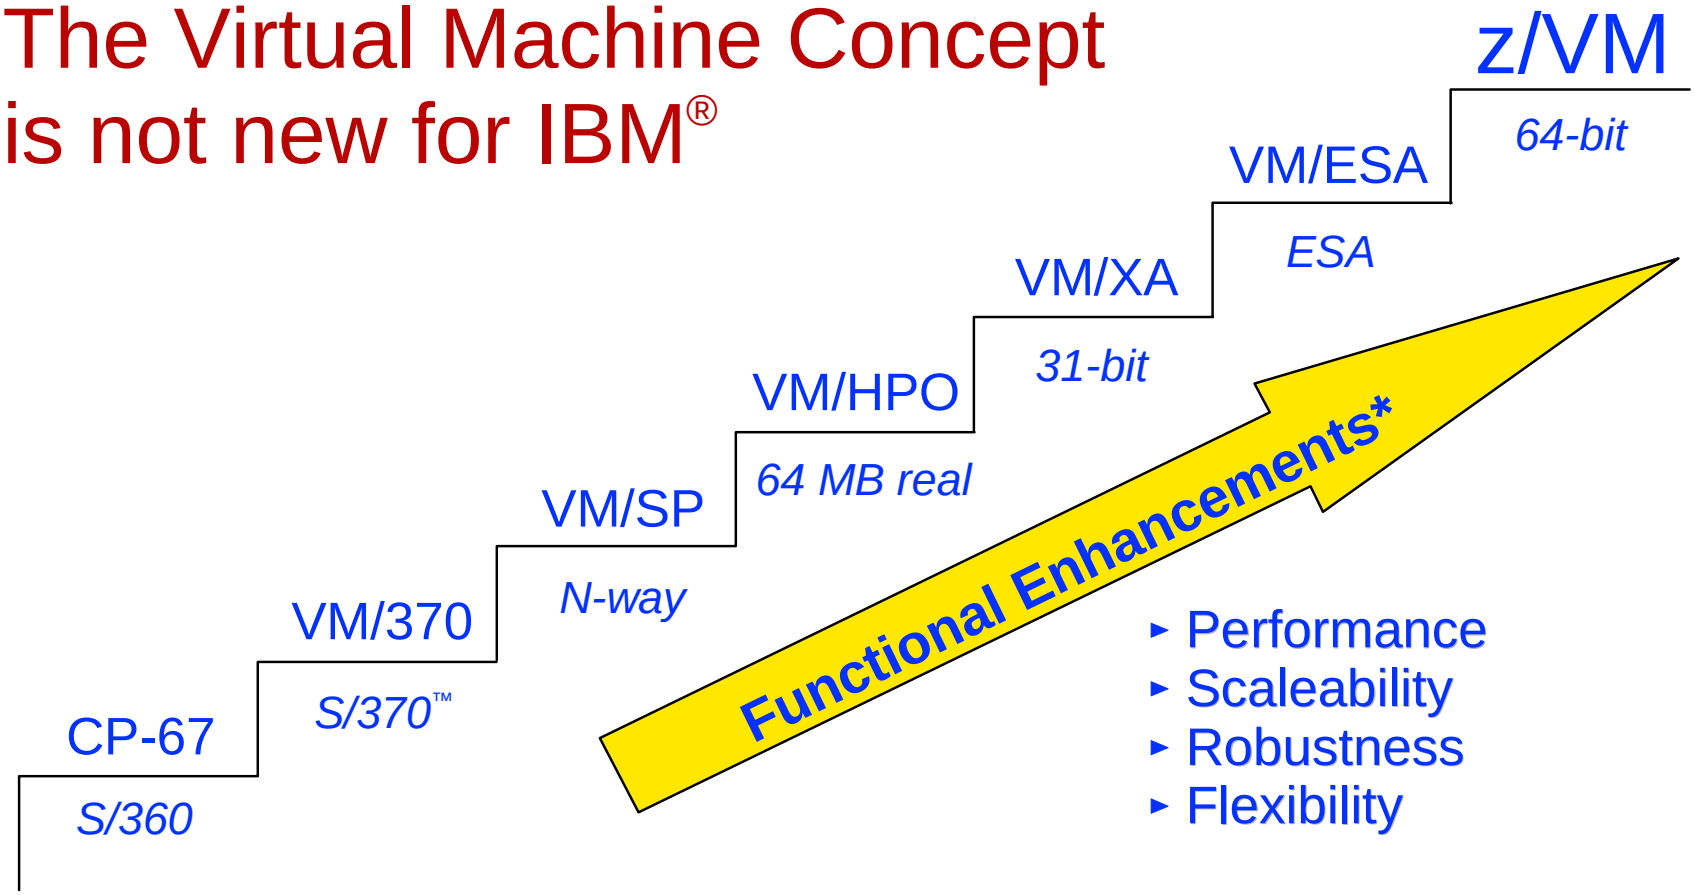

* Investments made in hardware, architecture, microcode, software

Virtual Machine Partitioning

Efficiently Exploiting the Entire Mainframe Complex

IBM @server zSeries

A *Virtual Machine* simulates the existence of a dedicated real machine, including processor functions, storage, and input/output resources.

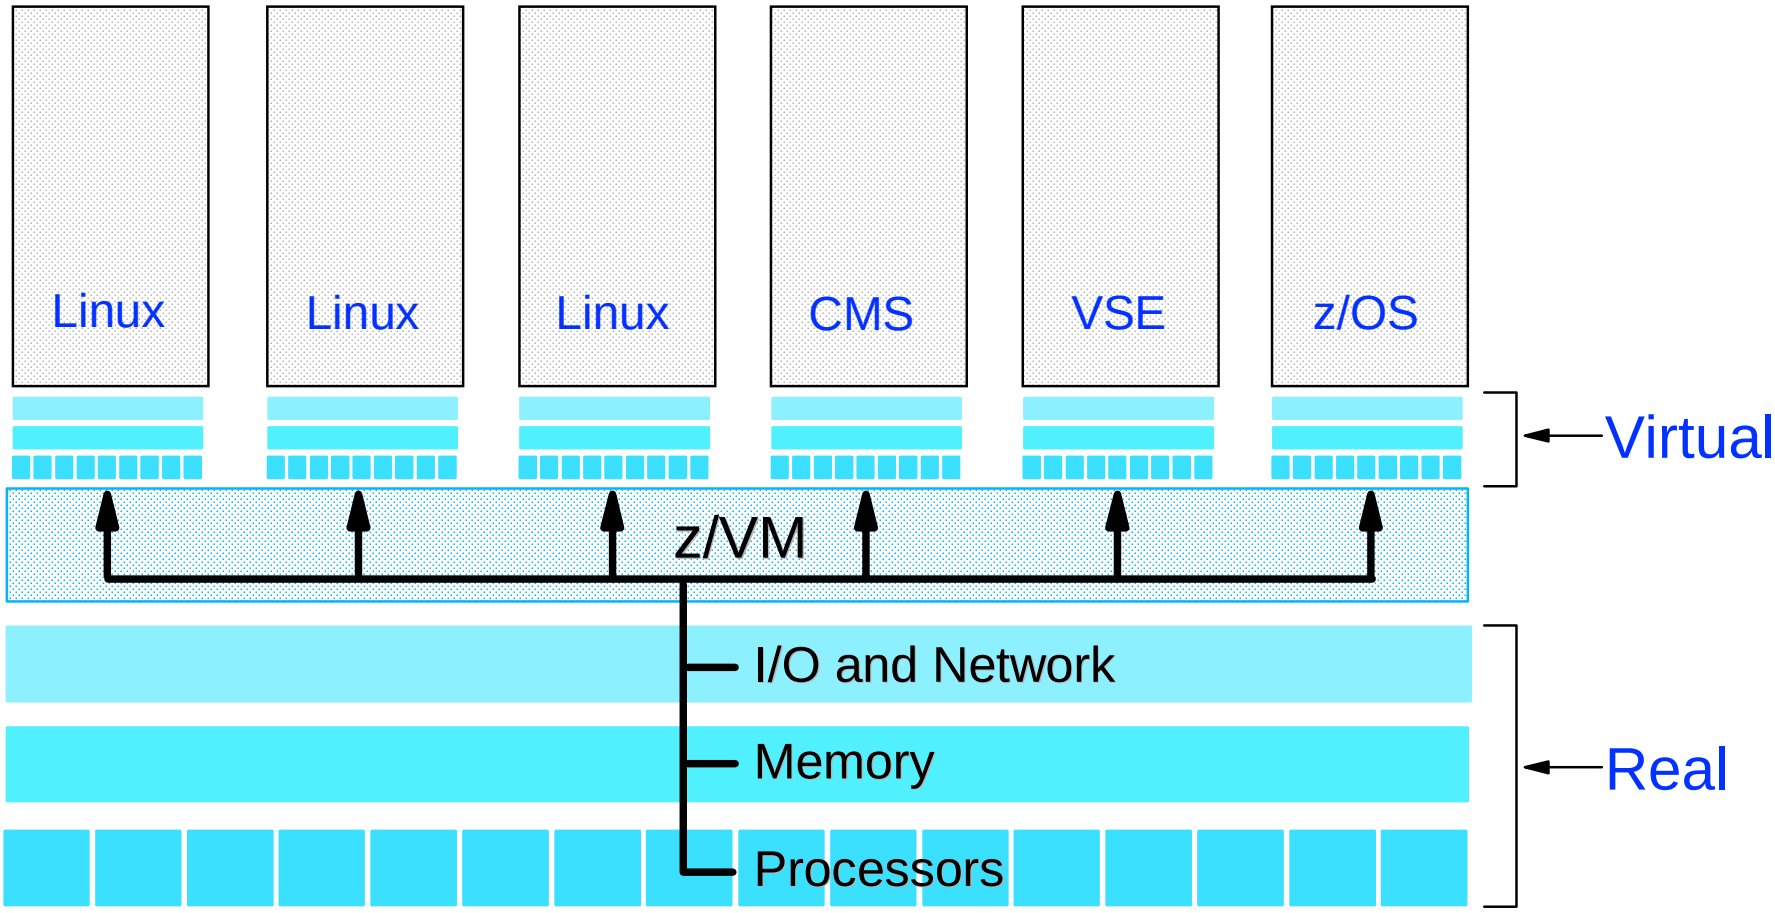

z/VM Version *4* Product Information

IBM @server zSeries

Runs on IBM G5 processor technology or better

- IBM @server zSeries™ 900 and zSeries 800 (including the z800 Model 0LF)
- IBM S/390 Parallel Enterprise Server G5 and G6
- IBM S/390® Multiprise® 3000

Runs on IFL processors as well as standard processors

IPLA software product (5739-A03)

- One-time charge (OTC) license fee, priced on a per-engine basis
- Ordered via the System Delivery Option (SDO): 5739-A04
- Serviced only via mail, fax, or e-mail under basic warranty

Optional Software Subscription and Support product (5739-SWS)

- Required to receive telephone defect support
- Entitles customers to future z/VM releases and versions
- Annual, renewable license charge

z/VM V4 IPLA-priced features (support standard or IFL engines):

- RealTime Monitor (RTM)
- Performance Reporting Facility (PRF)
- Directory Maintenance Facility (DirMaint)
- RACF for z/VM (4.3.0)
- New Performance Toolkit based on FCON (future)

S/390 Integrated Facility for Linux (IFL)
Description

IBM @server zSeries

IFLs are processors dedicated to Linux-only workloads

- Modeled after Integrated Coupling Facility (ICF) processors
- Less expensive than standard processors
- Will not support traditional mainframe operating systems
  - Only Linux and z/VM Version 4 will run on IFL engines
- Only usable in LPAR mode and cannot be "mixed" with standard processors

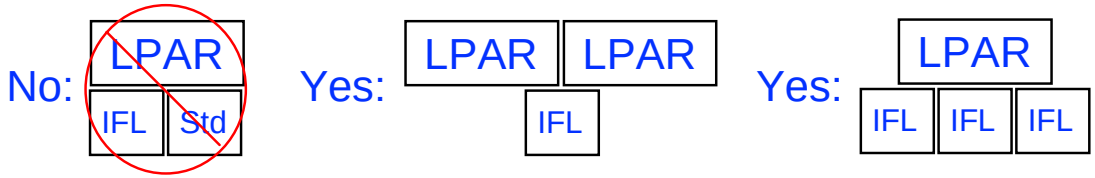

Available only with 9672 G5/G6, Multiprise 3000, zSeries servers

- One standard processor must exist before IFL processors can be added
  - Exception: the Linux-only model of the z800 is configured with IFL processors only
- Some servers don't have processors available for IFL use (e.g., z900-109)

Adding IFL processor(s) does not change a server's model designation

- No increase in fees for IBM software on standard processors
- Other software vendors have adopted this practice

Sample IFL Configuration

zSeries 900 Model 1C6 with 3 IFL Processors and z/VM V4

IBM @server zSeries

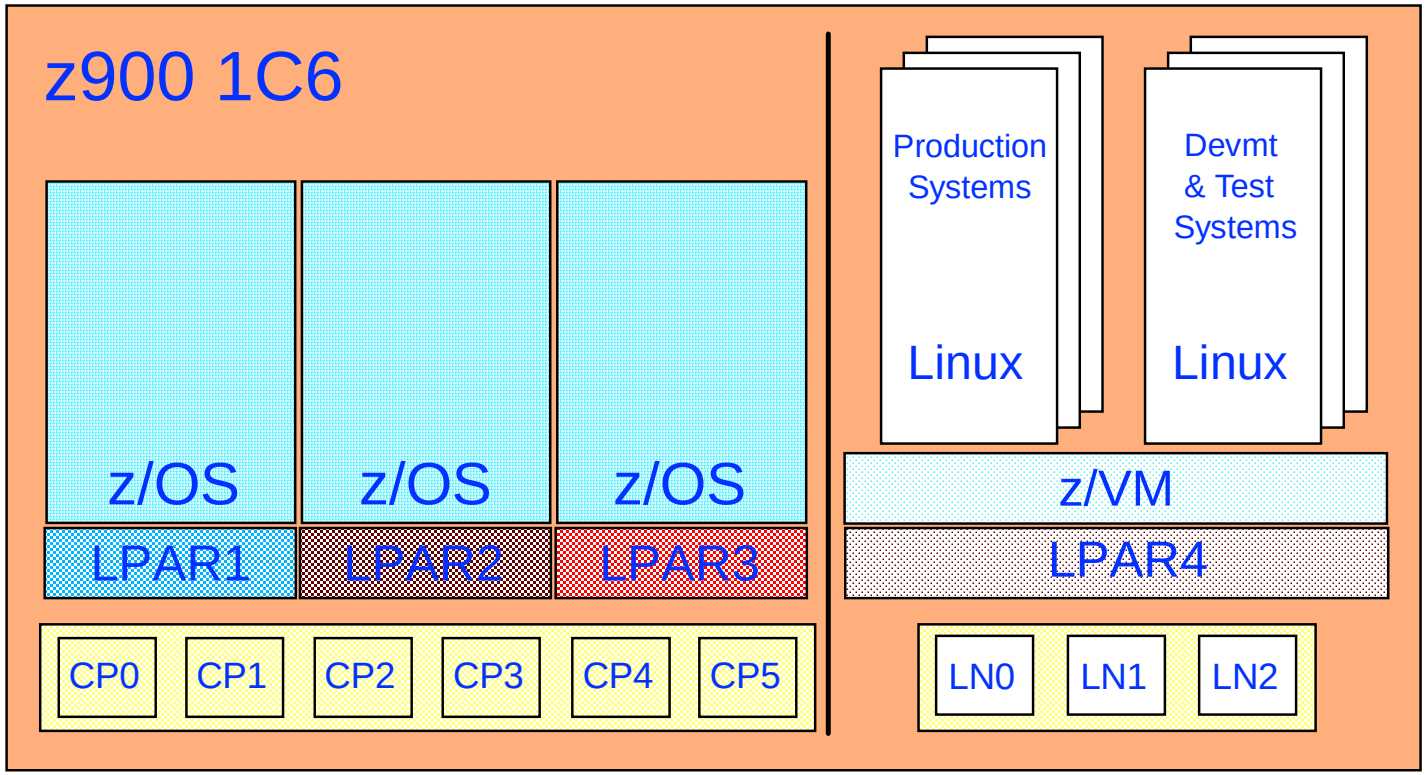

3-processor z/VM V4 charges* (U.S. prices)

Year 1	$168K	OTC plus S&S
Year 2	$33K	S&S
Year 3	$33K	S&S
3-Yr Total	$234K	

* Excludes optional priced features

Why Run Linux with z/VM?

IBM @server zSeries

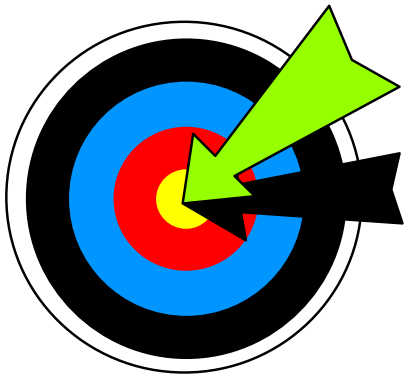

Server Consolidation

- Consolidate servers with z/VM and Linux
- "Save a bundle"
- Improve qualities of service

Speed to Market

- Deploy servers and solutions *fast*
- React quickly to challenges and opportunities
- Build innovative e-business solutions

Technology Exploitation

- Linux with z/VM offers more function than Linux alone
- Linux can exploit unique z/VM technology features

Server Consolidation with Linux on z/VM

"Server Farm in a Box"

IBM @server zSeries

Traditional Server Farm

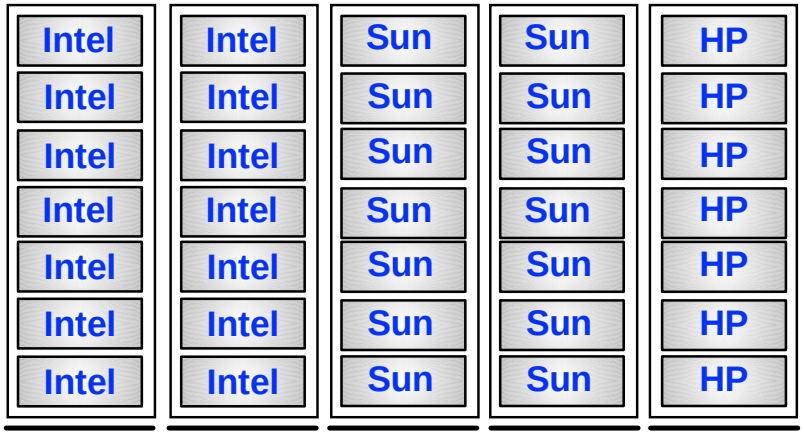

Server Farm in a Box

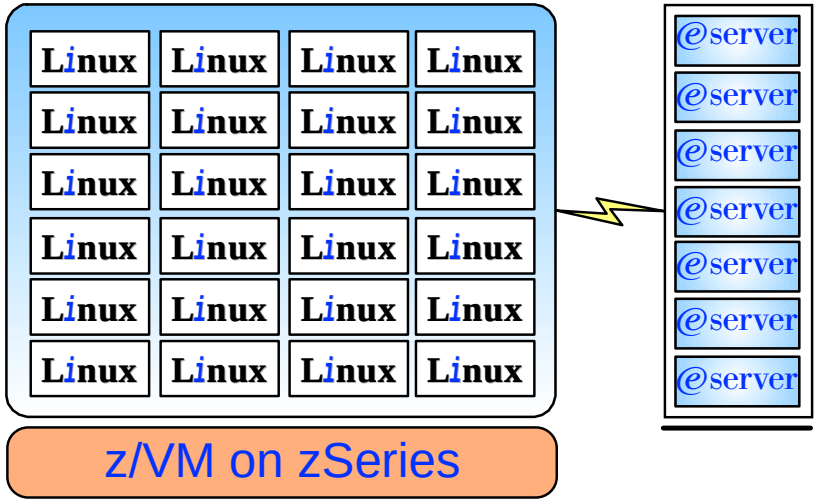

- Discrete servers consume incremental expense
  - Hardware price and maintenance
  - Floor space, power, cooling
  - Additional support staff
  - Per server (engine) software fees
- Connectivity requires kilometers of cables
- High availability ensured by spares / re-boots
- Disaster recovery rarely successfully tested

- Reduce costs without sacrificing server autonomy (one server per application)
- Virtual, high-speed, inter-server connectivity
- Exploit an architecture designed for high availability
- Mainframe qualities of service
- Proven disaster recovery services
- Connect to discrete servers as required

Server Consolidation with Linux on z/VM

Data and Application Servers on a Single @server zSeries

IBM @server zSeries

3-Tier Architecture

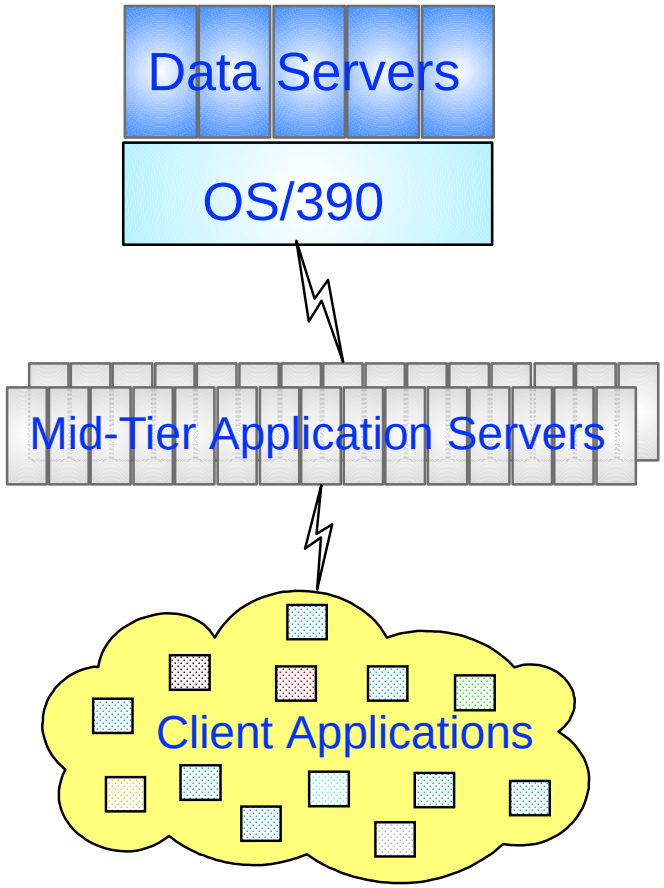

3-Tier Architecture on 2-Tiers of Hardware

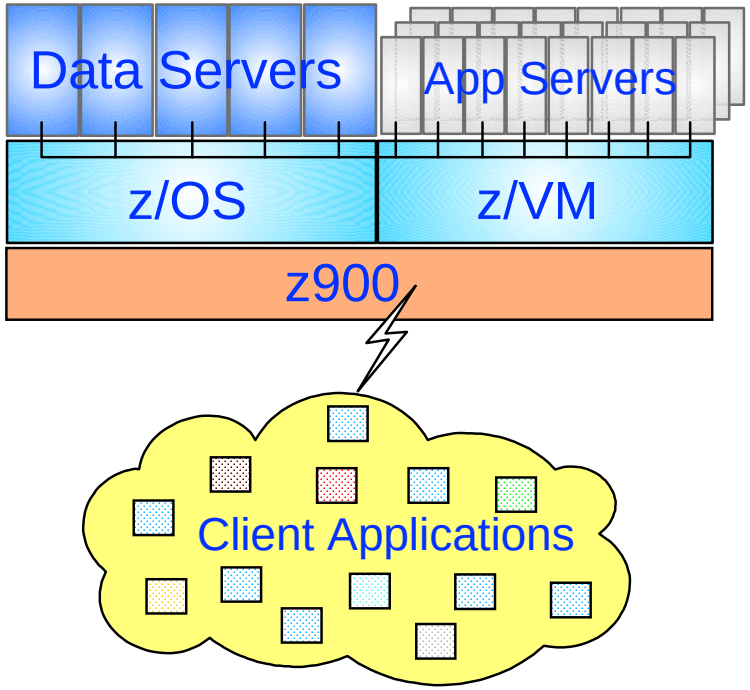

- Reduce expense (Hardware, software, floor space, energy, people)
- Exploit co-residency of application and data servers (improved performance, less complexity)
- Bring mainframe discipline to application servers
- No change to end-users

Rapid Horizontal Growth with Linux on z/VM

"Just add another server" quickly and easily with VM

IBM @server zSeries

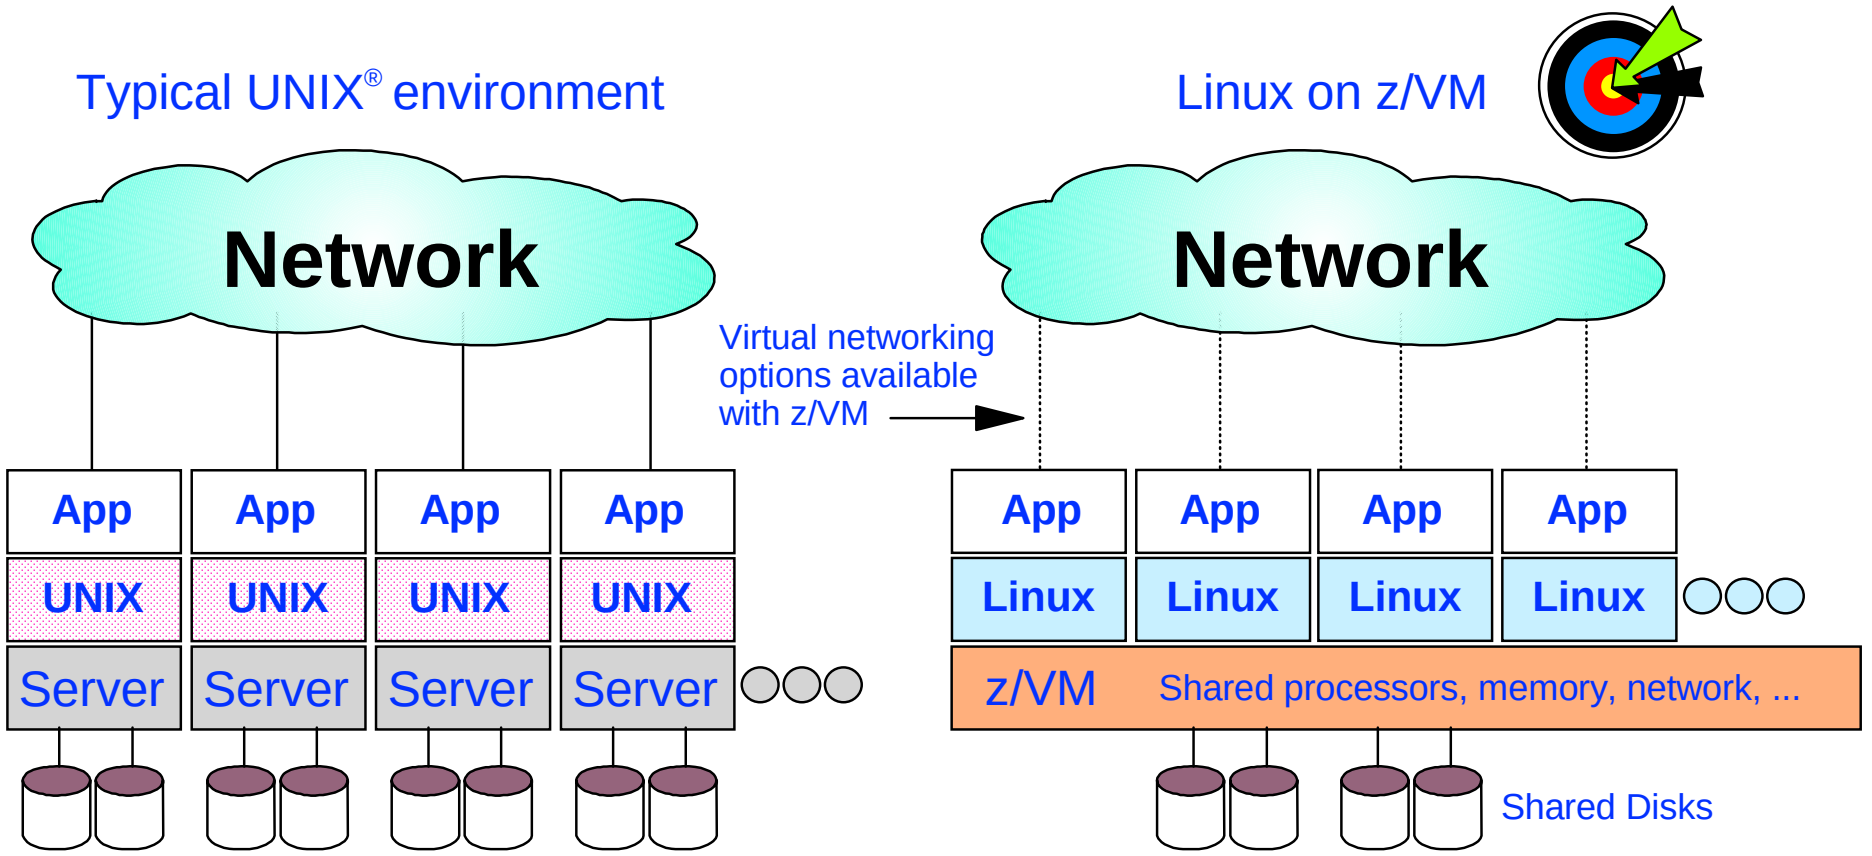

- ►Dedicated processors and disks
- ►Complex system management
- ►New servers available in days

- ►Shared resources
- ►Simplified system management
- ►New servers online in minutes

z/VM: Technology Exploitation for Linux

Performance - Scalability

IBM @server zSeries

Grow your Linux workload vertically or horizontally

- Vertical: add processor capacity, memory, and I/O to your virtual machines
- Horizontal: add more Linux virtual machines
  - Typical way of growing UNIX workloads
  - Spread the application load across multiple Linux images
  - Use z/VM technology to minimize resource duplication

Use VM's data-in-memory techniques for improved performance

- Virtual Disks in Storage (excellent swap device)
- Minidisk Cache (high-speed access to shared data)

Transparently exploit zSeries hardware when running in a virtual machine

- Cache control units, e.g., Peer-to-Peer Remote Copy (PPRC), FlashCopy
- FICON channels
- More...

Maximize throughput of a multi-Linux environment by exploiting z/VM's support for large n-way (SMP)

z/VM Technology - Large SMP Exploitation

IBM @server zSeries

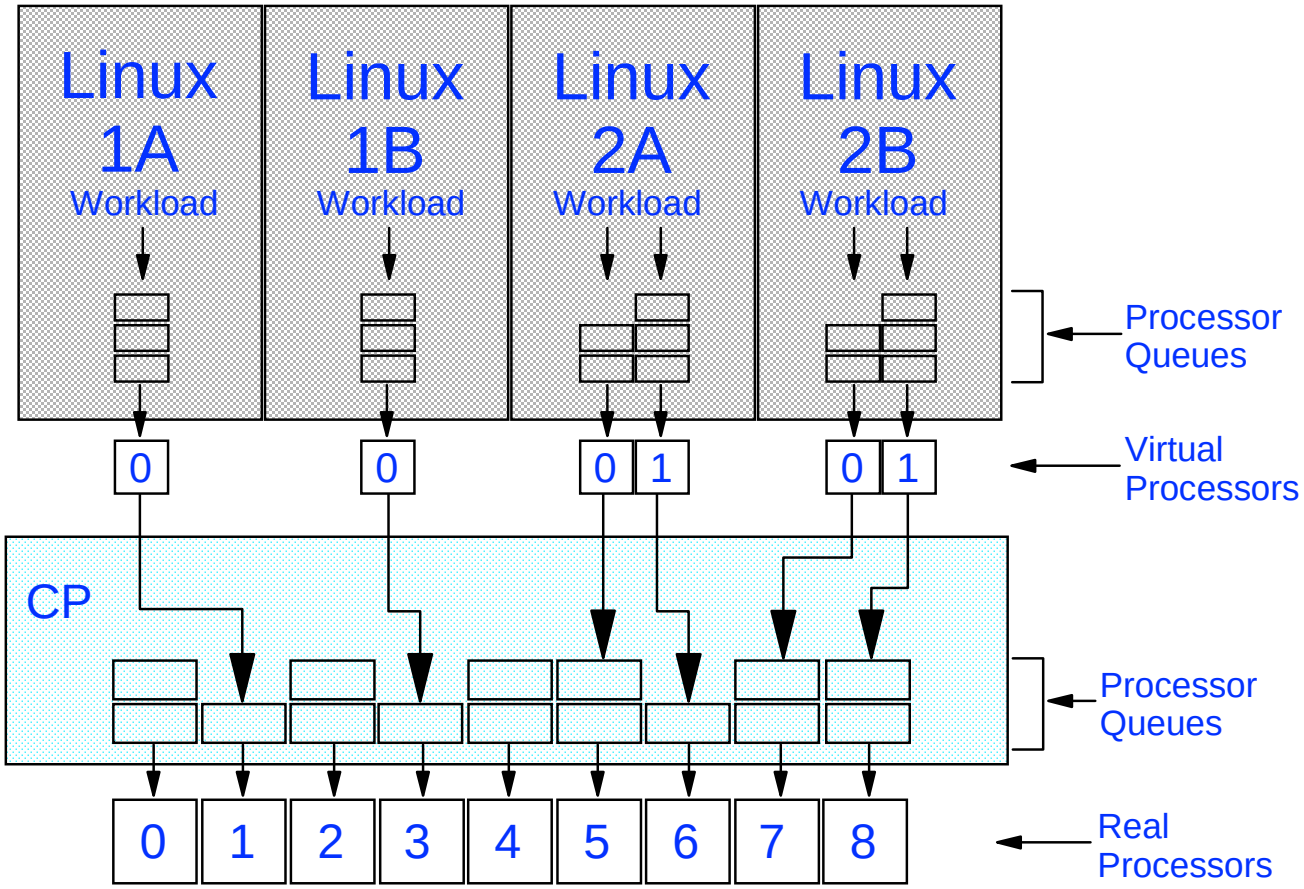

z/VM Technology - Resource Controls

IBM @server zSeries

Allocate system resources on a per-Linux-image basis

- SET SHARE command determines priority for CPU, main storage, and paging capacity
- Settings can be changed *on the fly* by command or programmed automation
- Resources are allocated to Absolute guests first, remaining resources are allocated to Relative guests

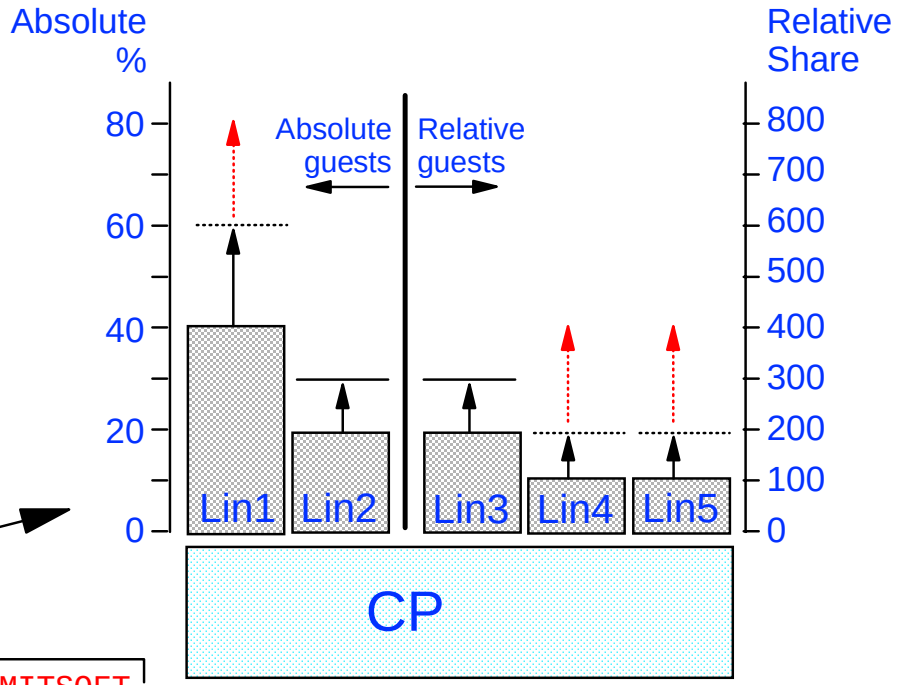


```
SET SHARE Lin1 ABSOLUTE 40% ABSOLUTE 60% LIMITSOFT
SET SHARE Lin2 ABSOLUTE 20% ABSOLUTE 30% LIMITHARD

SET SHARE Lin3 RELATIVE 200 RELATIVE 300 LIMITHARD
SET SHARE Lin4 RELATIVE 100 RELATIVE 200 LIMITSOFT
SET SHARE Lin5 RELATIVE 100 RELATIVE 200 LIMITSOFT
```

Notes:

..... = limit can be exceeded if unused capacity is available (limitsoft)

——— = limit will not be exceeded (limithard)

Virtual Networking: Point-to-Point Connections

IBM @server zSeries

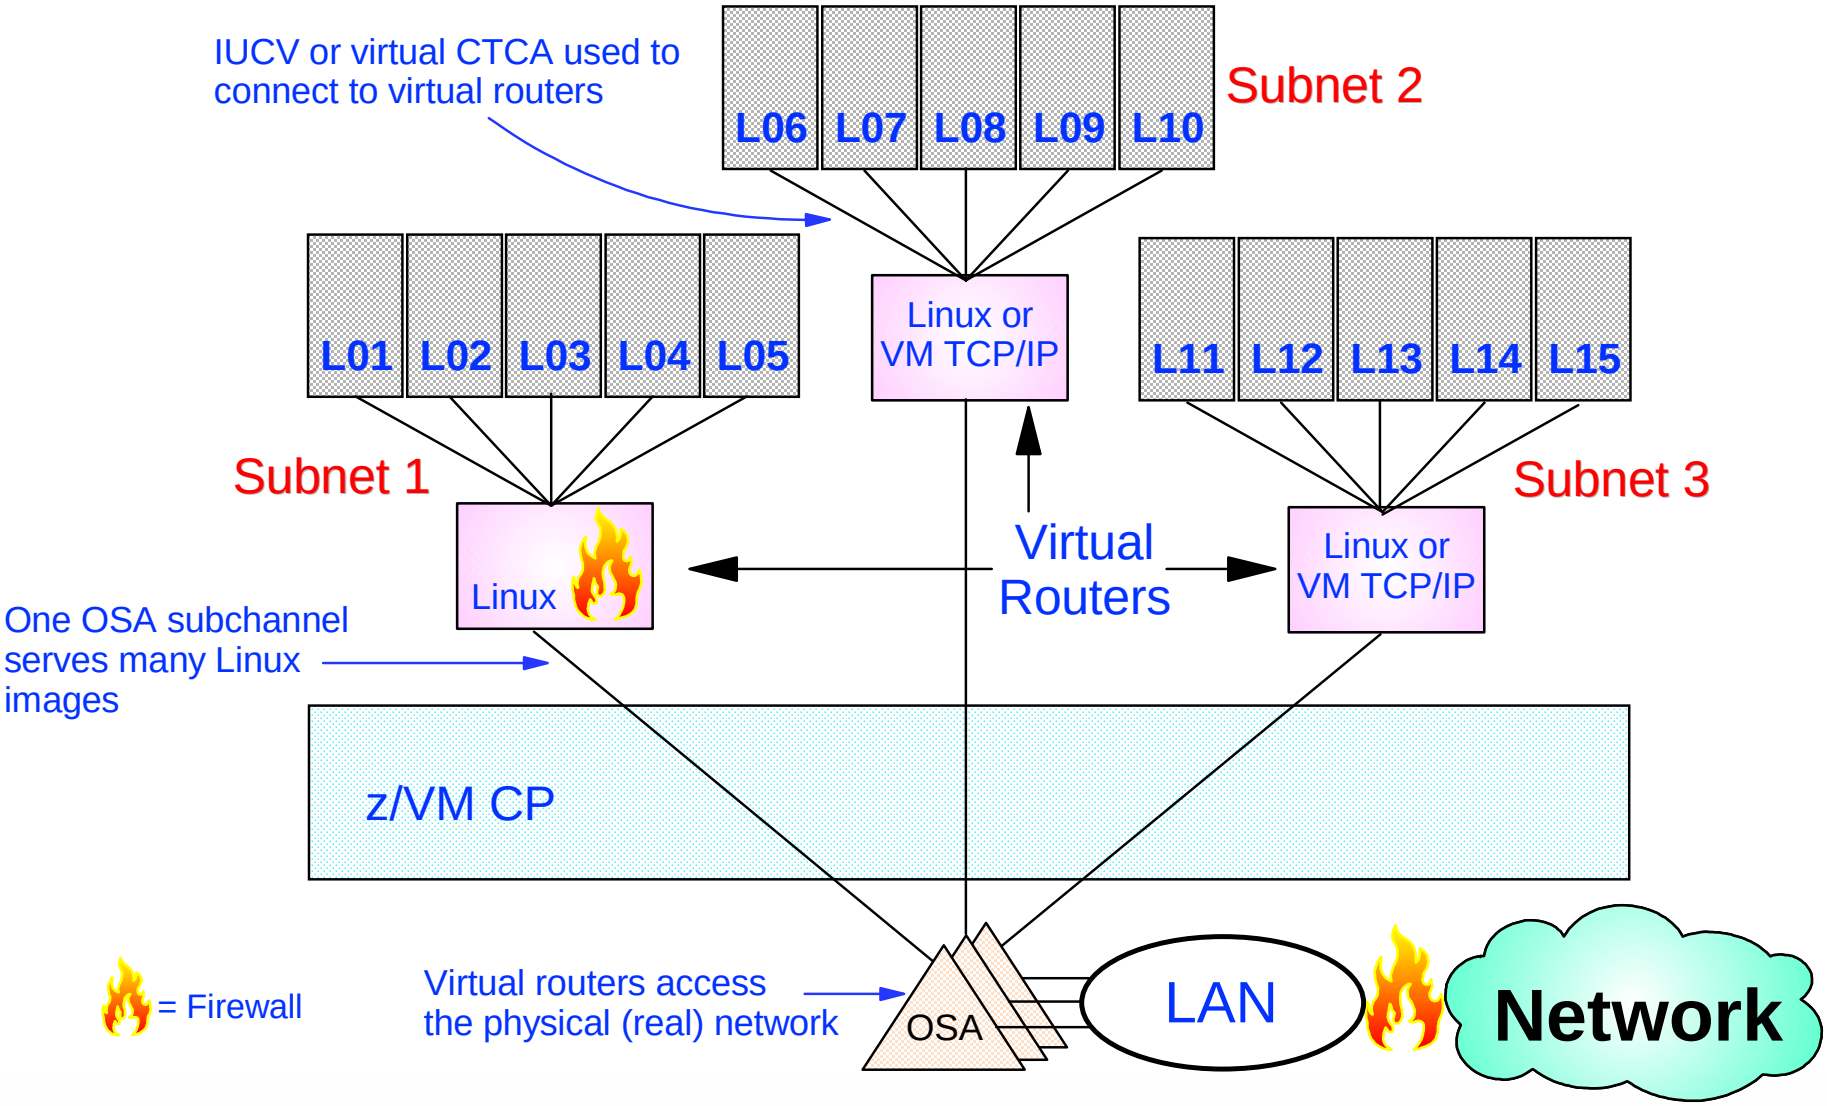

Virtual Networking: Using z/VM Guest LANs

IBM @server zSeries

One Linux guest connects to external network(s)

- Also connected to multiple Guest LANs
- Provides external routing and firewall services for guests

Other Linux guests connect to individual Guest LAN(s)

- Virtual HiperSockets and OSA Express connections supported
- Point-to-point, Multicast, and Broadcast (QDIO) supported

An ideal way to connect a server farm to z/OS

- Using real HiperSockets

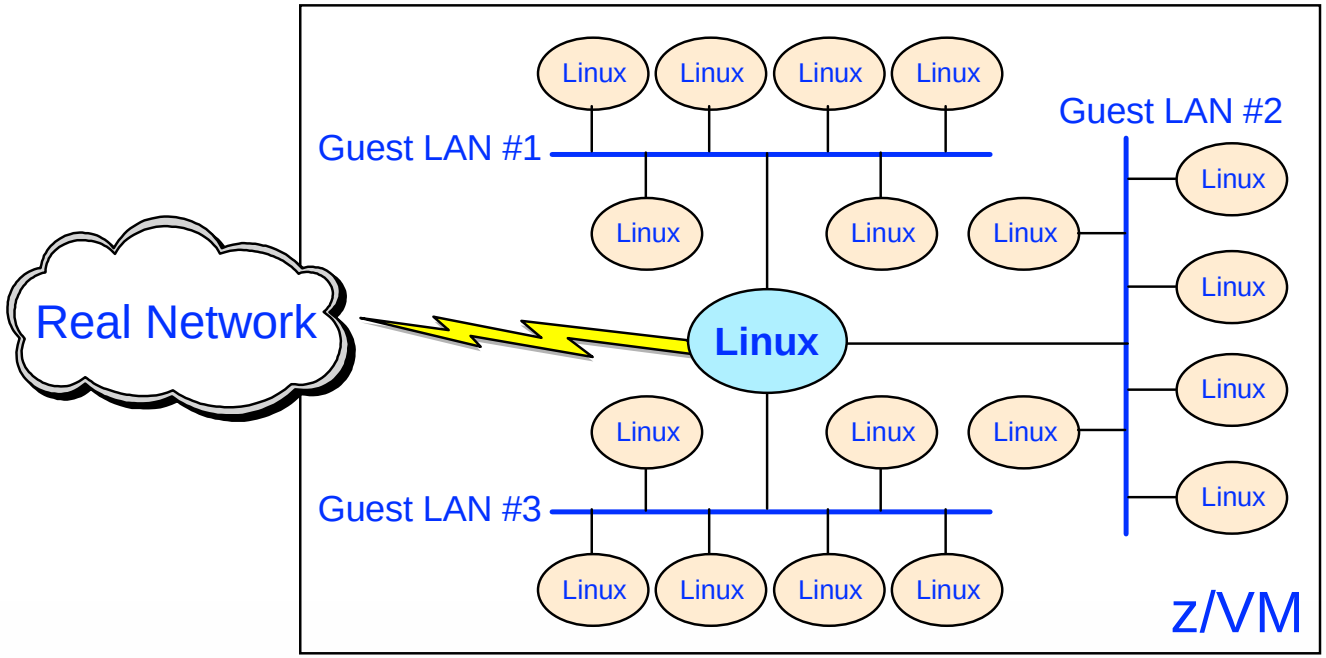

z/VM: Technology Exploitation for Linux

Productivity - Production and Test Workloads

IBM @server zSeries

Handle unexpected workload growth by adding Linux virtual machines quickly, easily, efficiently

Dynamically create Linux images for hot stand-by

- Flexible, cost-efficient high availability solution for Linux server images
- Automate or perform manually

Share data with read-only Minidisks

- Ideal for version control and centralized software upgrades
- Bonus: high-performance data access using Minidisk cache support

Use TDISKs for unexpected storage needs

- Access TDISKs using the Diagnose function of the Linux DASD driver
- Quick and easy way to provide additional temporary disk space

VM is ideal for prototyping, proof-of-concept, system testing

- Debugging is a legendary VM strength
- Execute complex test scenarios with minimal duplication of real hardware
- Give Linux systems to every developer, tester, student, etc.
- Add z/OS, OS/390, VSE, or CMS to your environment for a complete, multi-system test suite

z/VM Technology - Disk

IBM @server zSeries

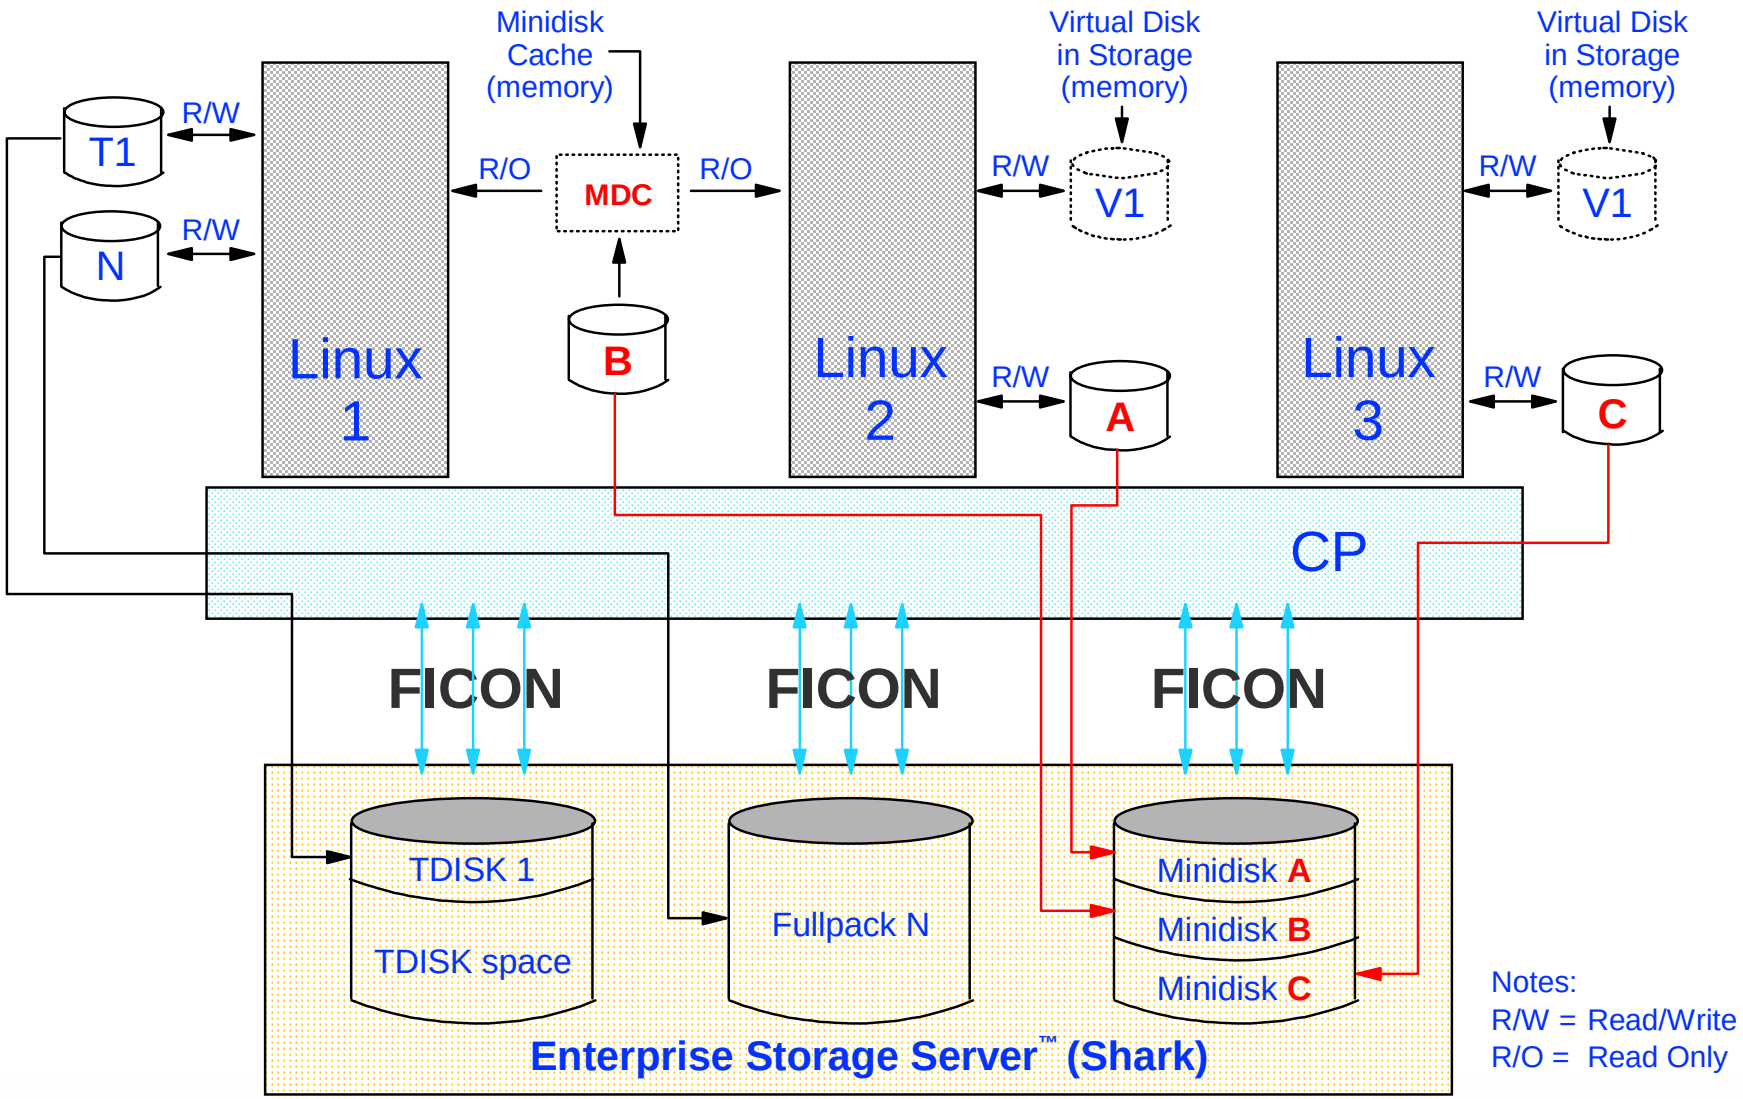

z/VM: Technology Exploitation for Linux

Systems Management - Data Backup using Shark FlashCopy

IBM @server zSeries

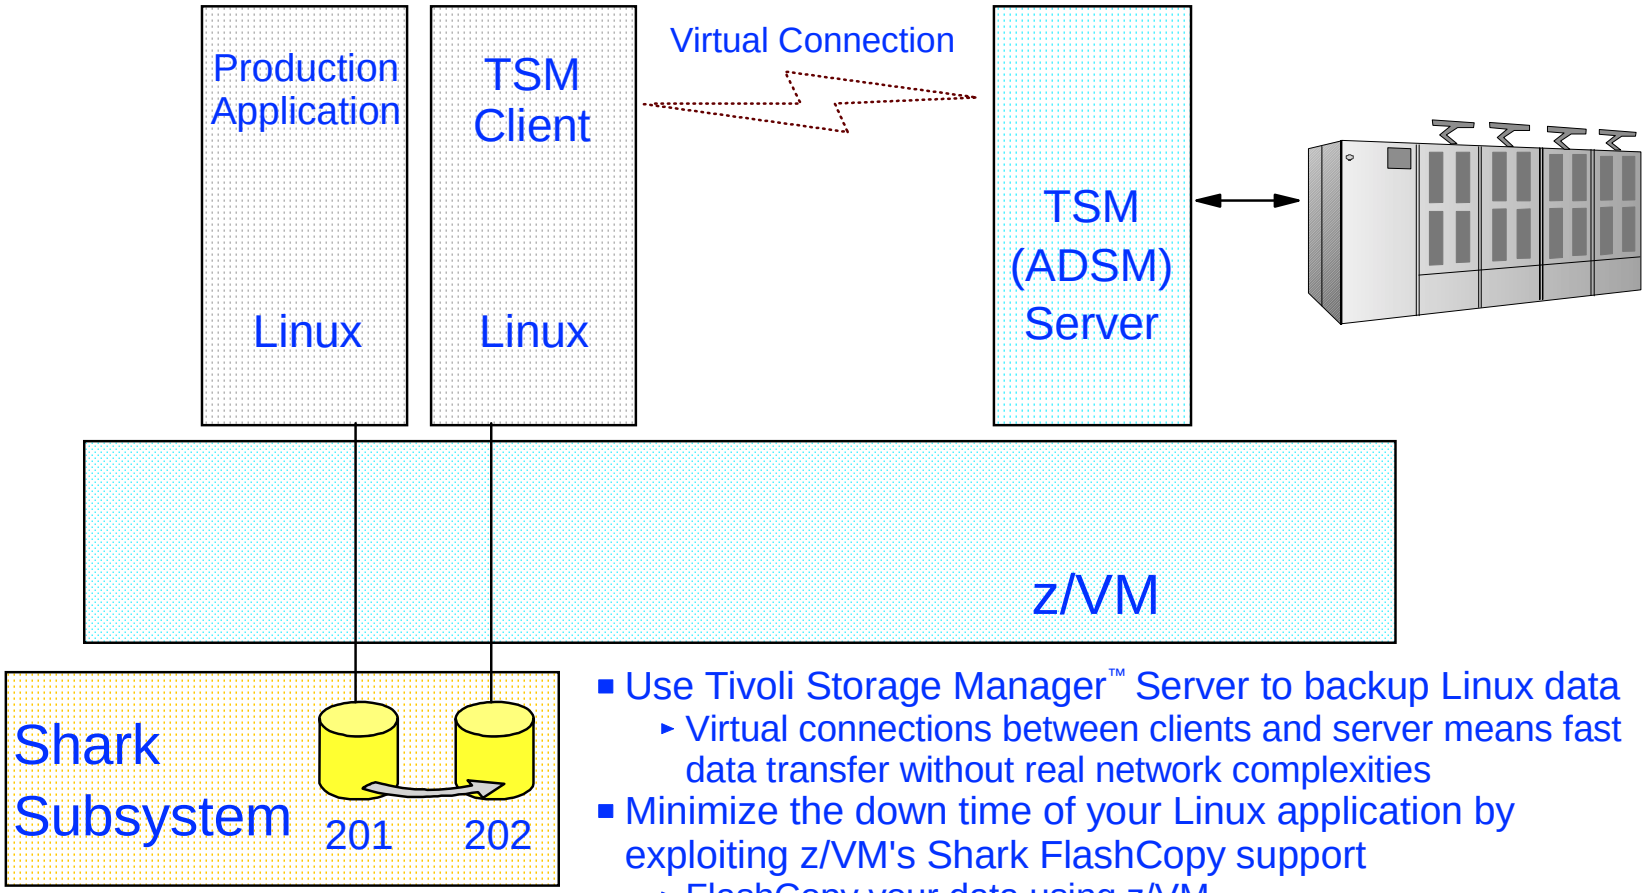

- Use Tivoli Storage Manager™ Server to backup Linux data
  - Virtual connections between clients and server means fast data transfer without real network complexities
- Minimize the down time of your Linux application by exploiting z/VM's Shark FlashCopy support
  - FlashCopy your data using z/VM
  - Run the TSM client code in a separate Linux server image
  - Your production application is back online in seconds

z/VM: Technology Exploitation for Linux

Systems Management - Command and Control

IBM @server zSeries

Systems Management functions built-into z/VM

- Performance collection and reporting for every Linux image
- Log accounting records for charge back
- Direct console output from your Linux systems to a single console
  - Using "Set Observer" or Single Console Image Facility (SCIF)
  - Automate system actions based on console output
- System Automation with CMS, Rexx, Pipelines, PROP
- Userid and password control for each Linux image
- Record and report hardware errors using EREP
- Dynamic I/O reconfiguration...add more disks without interruption

System management products are available from IBM and vendors

z/VM Technology - Systems Management

IBM @server zSeries

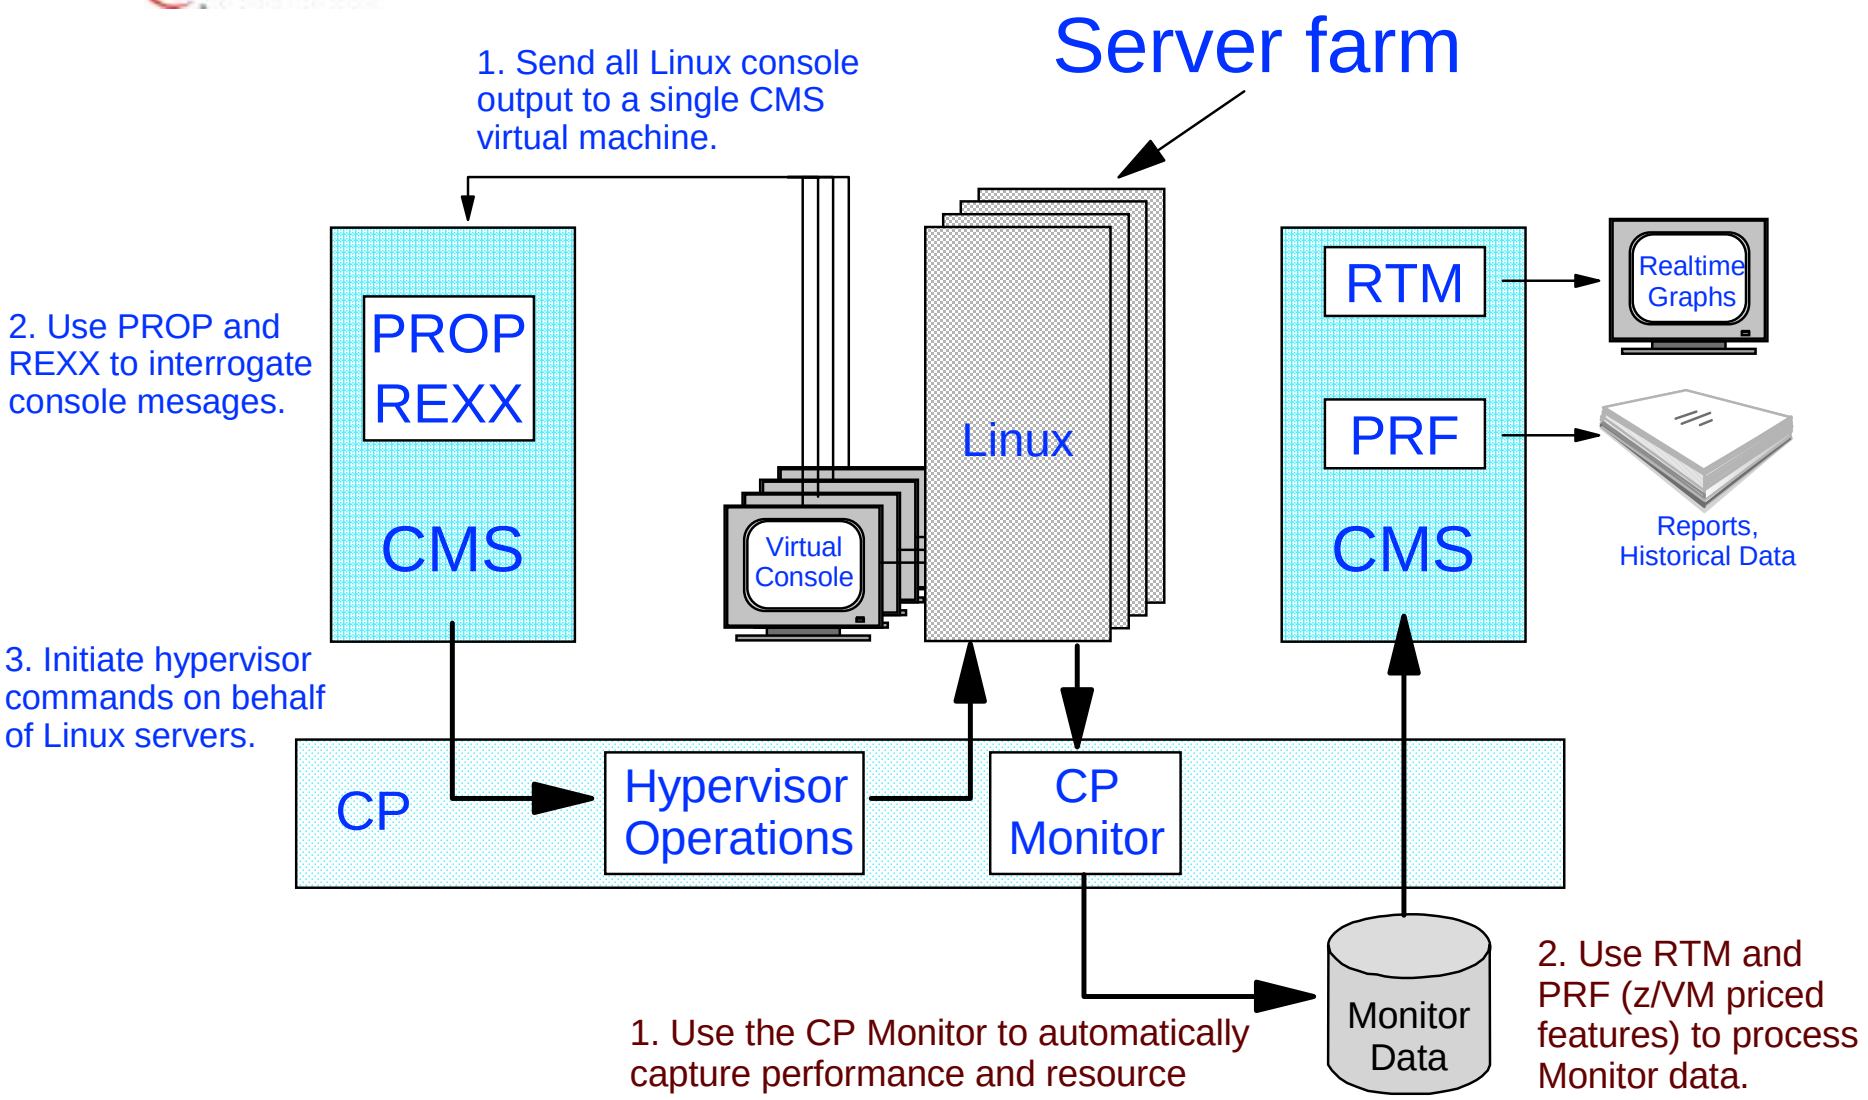

Linux and z/VM: Perfect Together
The Value of VM for Linux

IBM @server zSeries

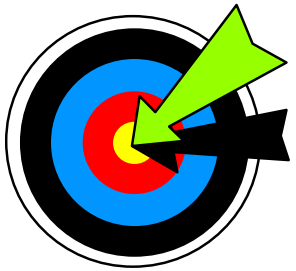

The machines are virtual -- *productivity and cost advantages*
- Ideal for prototyping, proof-of-concept, system testing
- More flexible and less disruptive than LPAR
- System resources can be virtualized, simulated, and shared
- Customers can adapt quickly with virtual system environments

Very large multi-image support -- *do more with less*
- Run tens to hundreds of images on one server (even in one LPAR)
- Ideal for consolidating multiple server images on a single processor
- Enables workload growth via server replication
- Virtual machines for every developer, tester, student, department, etc.

Architecture exploitation -- *enhanced qualities of service*
- Create a guest environment functionally richer than native or LPAR
- High-performance networking among virtual machines
- Transparent exploitation of VM's hardware support (with error recovery)
- Enhanced performance using VM's data-in-memory support
- Multi-system thruput improvements using VM's large n-way support

z/VM Version 4 Release 3

IBM @server zSeries

Announced April 2002, available May 31, 2002

New function highlights:

- Virtual machine accounting enhancements (virtual networks)
- Fibre Channel Protocol (FCP) dedicated guest support
- Automated shutdown support for guest systems
- Enhanced timer management
- Virtual Machine Resource Manager (IRD-like function)
- Shark large volume support
- Guest LAN enhancements (QDIO, Multicast, Broadcast support)
- VM TCP/IP improvements (configuration, performance, security)
- RACF/VM feature

Learn more at the IBM z/VM web site

- ibm.com/eserver/zseries/zvm

Linux and z/VM Resources

IBM @server zSeries

IBM z/VM — **ibm.com**/eserver/zseries/zvm
- Online documentation and education material
- Tools and utilities download page

Linux for zSeries — **ibm.com**/eserver/zseries/linux
- Customer stories, education offerings
- Applications (ISV support)

IBM Global Services — techsupport.services.**ibm.com**/linux/support
- Technical documents, service offerings

Linux and z/VM Redbooks
- "Linux on IBM zSeries and S/390: **Systems Management**"
  - http://publib-b.boulder.ibm.com/Redbooks.nsf/RedpieceAbstracts/sg246820.html
- "Linux on IBM zSeries and S/390: **Server Consolidation** with Linux for zSeries"
  - http://publib-b.boulder.ibm.com/Redbooks.nsf/RedpaperAbstracts/redp0222.html
- "Linux on IBM zSeries and S/390: **High Availability** for z/VM and Linux"
  - http://publib-b.boulder.ibm.com/Redbooks.nsf/RedpaperAbstracts/redp0220.html
- "Linux on IBM eServer zSeries and S/390: **Cloning** Linux Images in z/VM"
  - http://publib-b.boulder.ibm.com/Redbooks.nsf/RedpaperAbstracts/redp0301.html

z/VM Security and Integrity — **ibm.com**/servers/eserver/zseries/library/techpapers/gm130145.html

Linux for S/390 Listserver — www.marist.edu/htbin/wlvindex?linux-390
- Interact with the Linux for zSeries user community
- Scan the archive for how-to support