z/VM Security and Integrity

Alan Altmark, z/VM Architecture and Design

Alan_Altmark@us.ibm.com

Trademarks

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

DB2	System z9
DS8000*	System z10
Enterprise Storage Server*	z9*
IBM*	z10
IBM eServer	z/OS*
IBM logo*	z/VM
System Storage*	z/VSE
System z	zSeries*

* Registered trademarks of IBM Corporation

The following are trademarks or registered trademarks of other companies.

Intel is a trademark of Intel Corporation in the United States, other countries, or both.

Java and all Java-related trademarks and logos are trademarks of Sun Microsystems, Inc., in the United States and other countries

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

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.

Red Hat, the Red Hat "Shadow Man" logo, and all Red Hat-based trademarks and logos are trademarks or registered trademarks of Red Hat, Inc., in the United States and other countries.

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

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.

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.

Prices subject to change without notice. Contact your IBM representative or Business Partner for the most current pricing in your geography.

Virtualization security risks being overlooked, Gartner warns

Gartner raises warning on virtualization and security.

Companies in a rush to deploy virtualization technologies for server consolidation efforts could wind up overlooking many security issues and exposing themselves to risks, warns research firm Gartner.

“Virtualization, as with any emerging technology, will be the target of new security threats,” said Neil MacDonald, a vice president at Gartner, in a published statement.

Network World
April 6, 2007

Integrity

What is system integrity?

1. The ability of the hypervisor (CP) to operate without interference or harm, intentional or not, from the guest virtual machines
2. The inability of a virtual machine to circumvent system security features and access controls
3. The ability of the hypervisor to protect virtual machines from each other

System Integrity

- But how is that actually done?
- Answer: Interpretive Execution Facility

Interpretive Execution Facility

- Start Interpretive Execution (SIE) instruction runs a virtual machine
  - Registers, PSW (Program Status Word), memory
  - Interception conditions (a.k.a. "SIE break")
    - Time slice expires
    - Unassisted I/O
    - Instructions that require CP's help
      - e.g. Set Clock
  - Certain program interrupts
- SIE runs until interception condition raised

Hardware Access to Virtual Memory

- SIE uses CP-maintained dynamic address translation tables to convert virtual addresses to real addresses

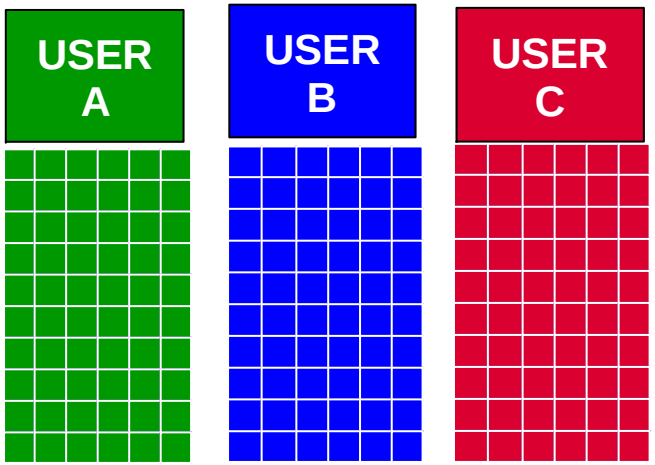

- CP provides page, segment, and region tables to SIE
- Page table entries are 'invalid' until initialized by CP

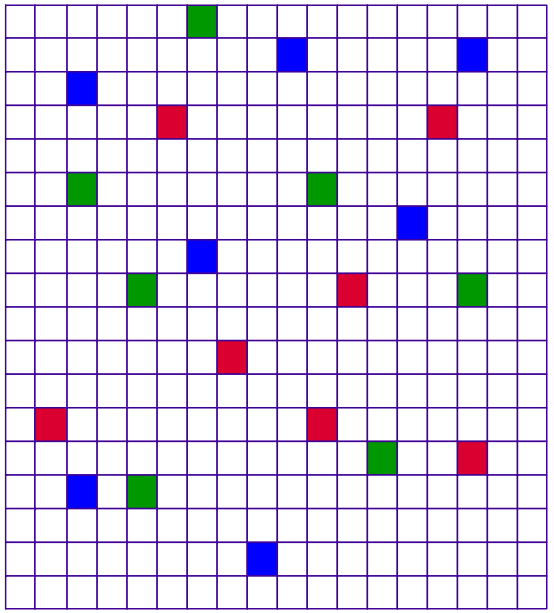

CP Real Memory

Interpretive Execution Facility

- The only virtualization technology on the market the provides not one, but two levels of hardware support for virtualization.
- The need exists for both “hard” virtualization (partitioning) and “soft” virtualization (z/VM)

What is a logical partition?

- A virtual machine created by LPAR hypervisor using PR/SM technology

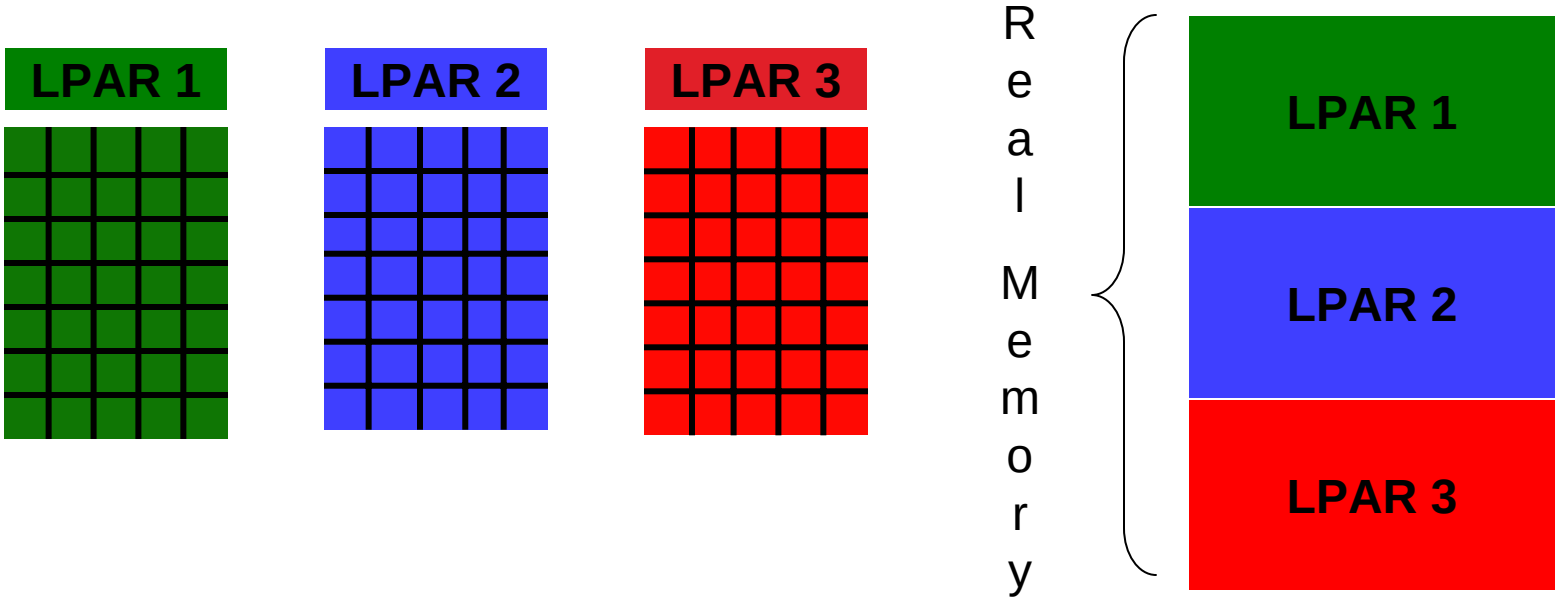

Virtual I/O

- SIE break – CP examines I/O request
  - Translates CCW virtual addresses to real addresses
  - Pins user pages in memory
  - Looks for harmful operations
  - Alters minidisk cylinder locations, if required
  - Inserts device limits whenever possible
    - DEFINE EXTENT for minidisks

DEFINE EXTENT

- A virtual machine has access to a “minidisk”
- CP translates virtual disk location (0-99) to an actual location
- DEFINE EXTENT I/O command forces control unit to confine I/O to the actual disk extent

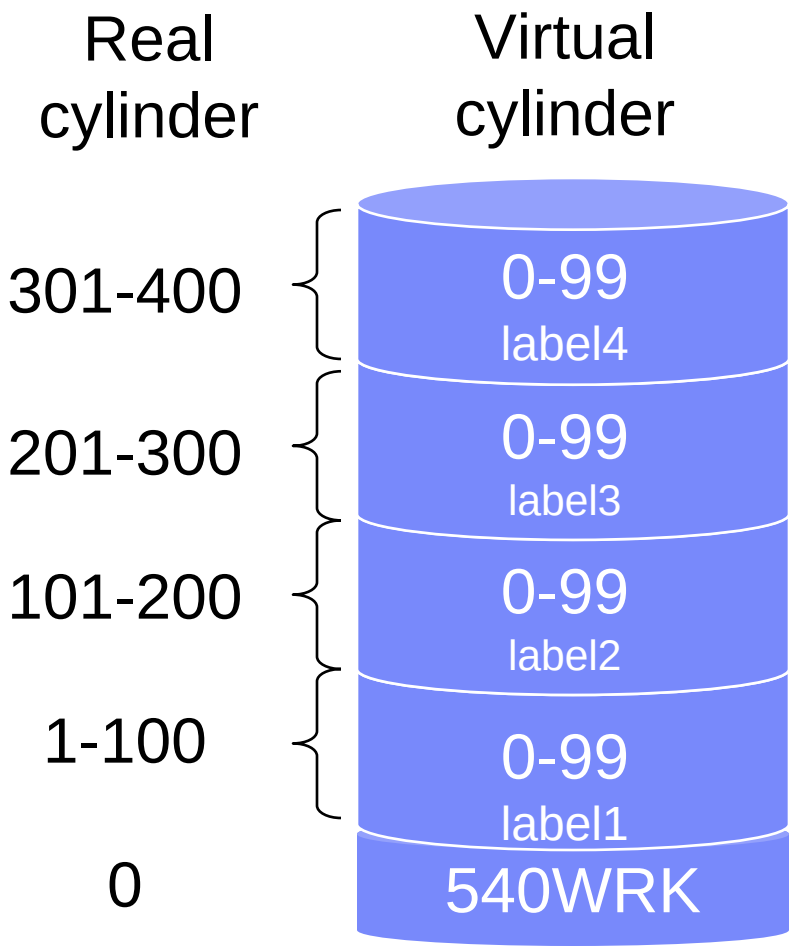

I/O Hardware Assist

- Interpretive Execution Facility handles I/O request
  - No SIE break, so no involvement of CP
  - CP and hardware share address tables
- Dedicated QDIO devices only
  - OSA and Fibre Channel (FCP)

Security vs. Integrity

- Security is only meaningful in the presence of system integrity!
  - Integrity prevents bypass of security controls
  - Audit trail confirms conformance

Security

What is System Security?

- Authentication
- Authorization
- Audit

An integrated set of system functions that control access to a system and its resources, and that provides a record of those accesses.

What is System Security?

- Authentication
- Authorization
- Audit

Authentication

- Three forms of identification
  - What you have (key)
  - What you know (password)
  - Who you are (fingerprint)
- Combinations may be used
  - Two-factor authentication (“2FA”)

Authentication

- z/VM uses a password or phrase to establish your identity
  - Logon
  - FTP
  - Rexec
  - NFS
  - …

What is System Security?

- Authentication
- Authorization
- Audit

Authorization

- Ensures that a user has access only to system resources specifically permitted or within scope of responsibility
  - Must be authenticated first!
- Applies to commands, interfaces, and data

Authorization

- Primary authorization mechanism is privilege class
  - Specified in USER DIRECT
  - Assigned to every command or DIAGNOSE instruction
- Class G is a General User
- Class A to F have special privileges

Privilege Class

- System administrators have class A, B, C, D, and/or E
  - Potential to bypass system integrity and security controls
- Do not assign privilege class B just to get the MSGNOH command
- Customer can alter assignment and contents of privilege classes

Privilege Class

- Excess privilege is the root of all Evil
  - Do not give privilege to untrusted virtual machines
- Alternatives
  - COMMAND statement in the directory
  - Alter privilege class of specific commands
  - Use automation
- “Less than Class G” – see the Internet

Directory COMMAND

- User directory entry can contain CP commands:

```
USER ALAN mypass G
COMMAND SET SHARE …
COMMAND VARY ON …
COMMAND ATTACH …
```

- Command is executed regardless of user privilege class
- Actual user assigned privilege class takes effect when guest begins running

Network with VSWITCH (with VLANs)

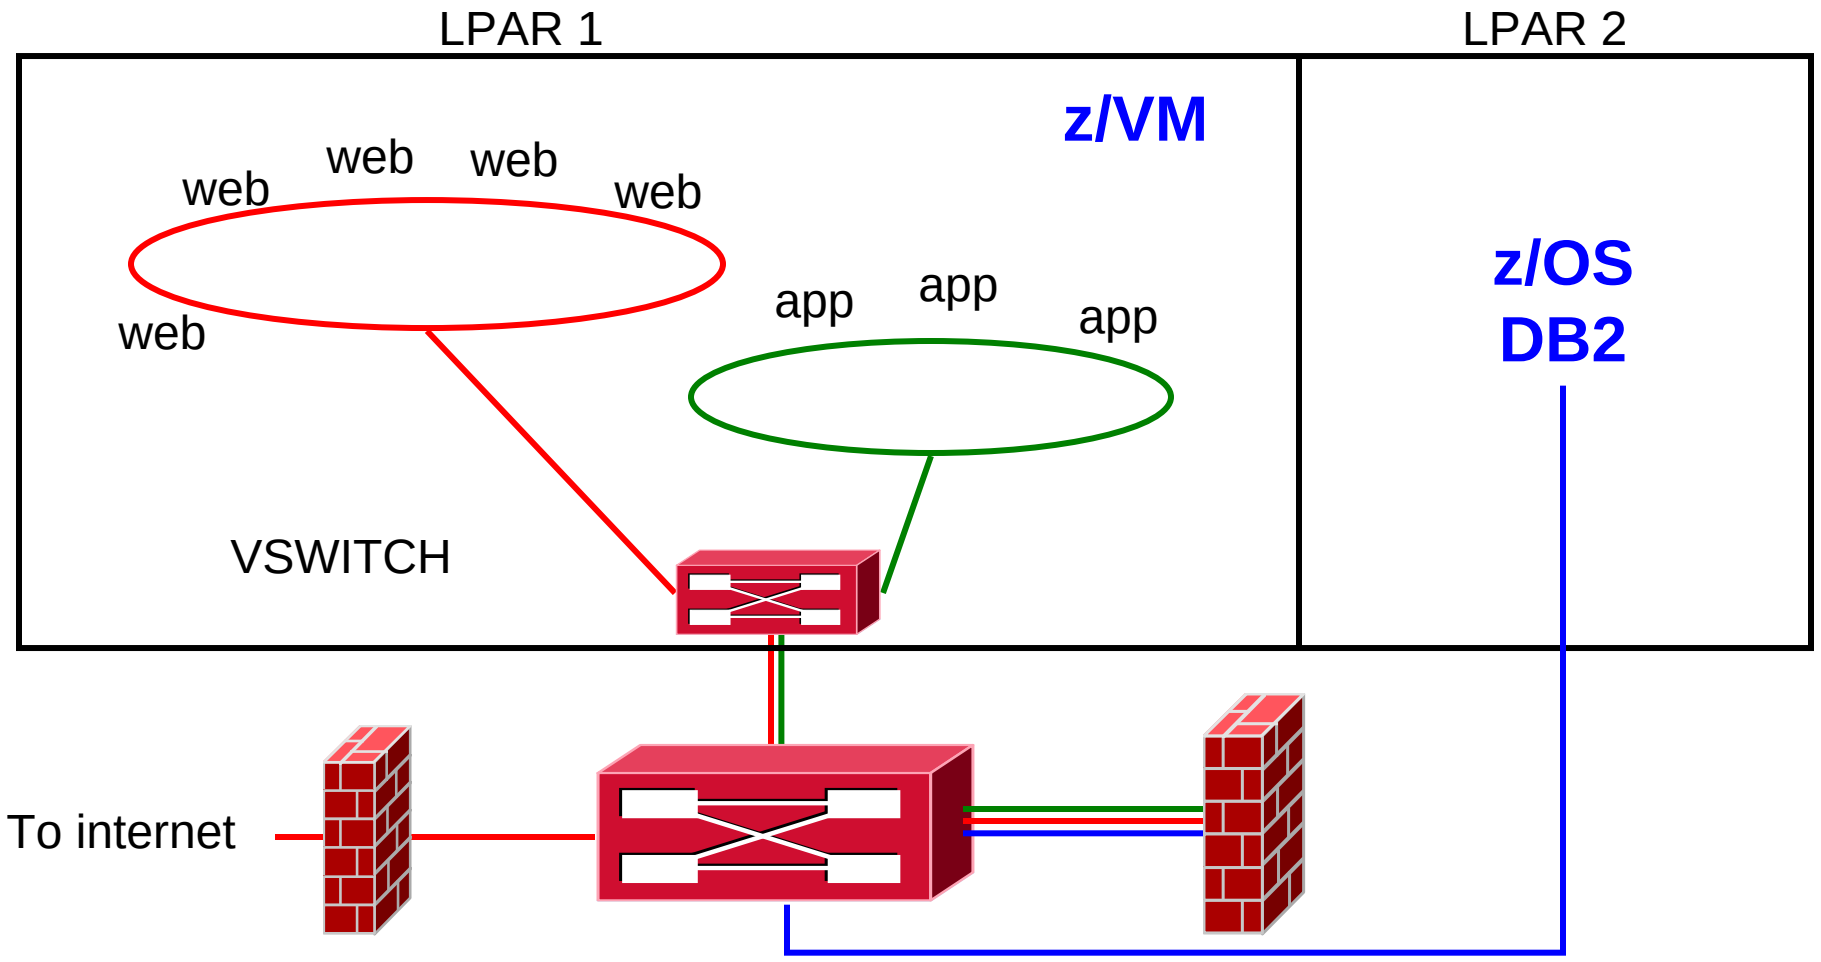

With 1 VSWITCH, 3 VLANs, and a multi-domain firewall

z/VM Virtual Switch – VLAN assignment

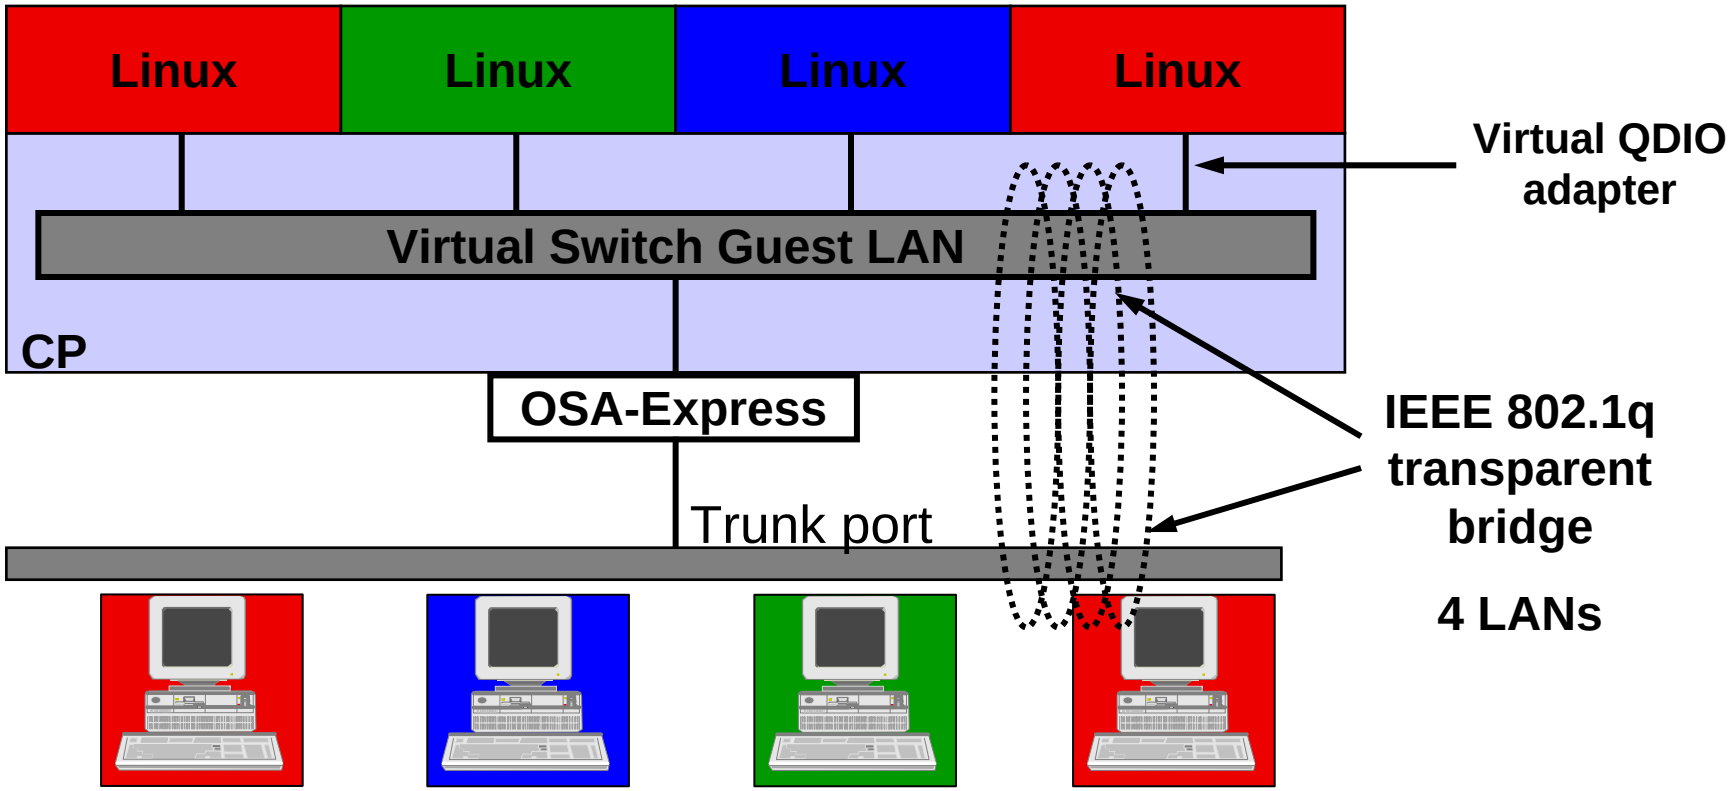

Virtual Switch

- You must authorize access
  - SET VSWITCH GRANT, or
  - VMLAN class in the RACF Security Server
  - Including VLAN ID
- Sniffer mode requires additional authorization
- Port isolation
  - Guests in same VLAN cannot talk to each other

Mandatory Access Controls

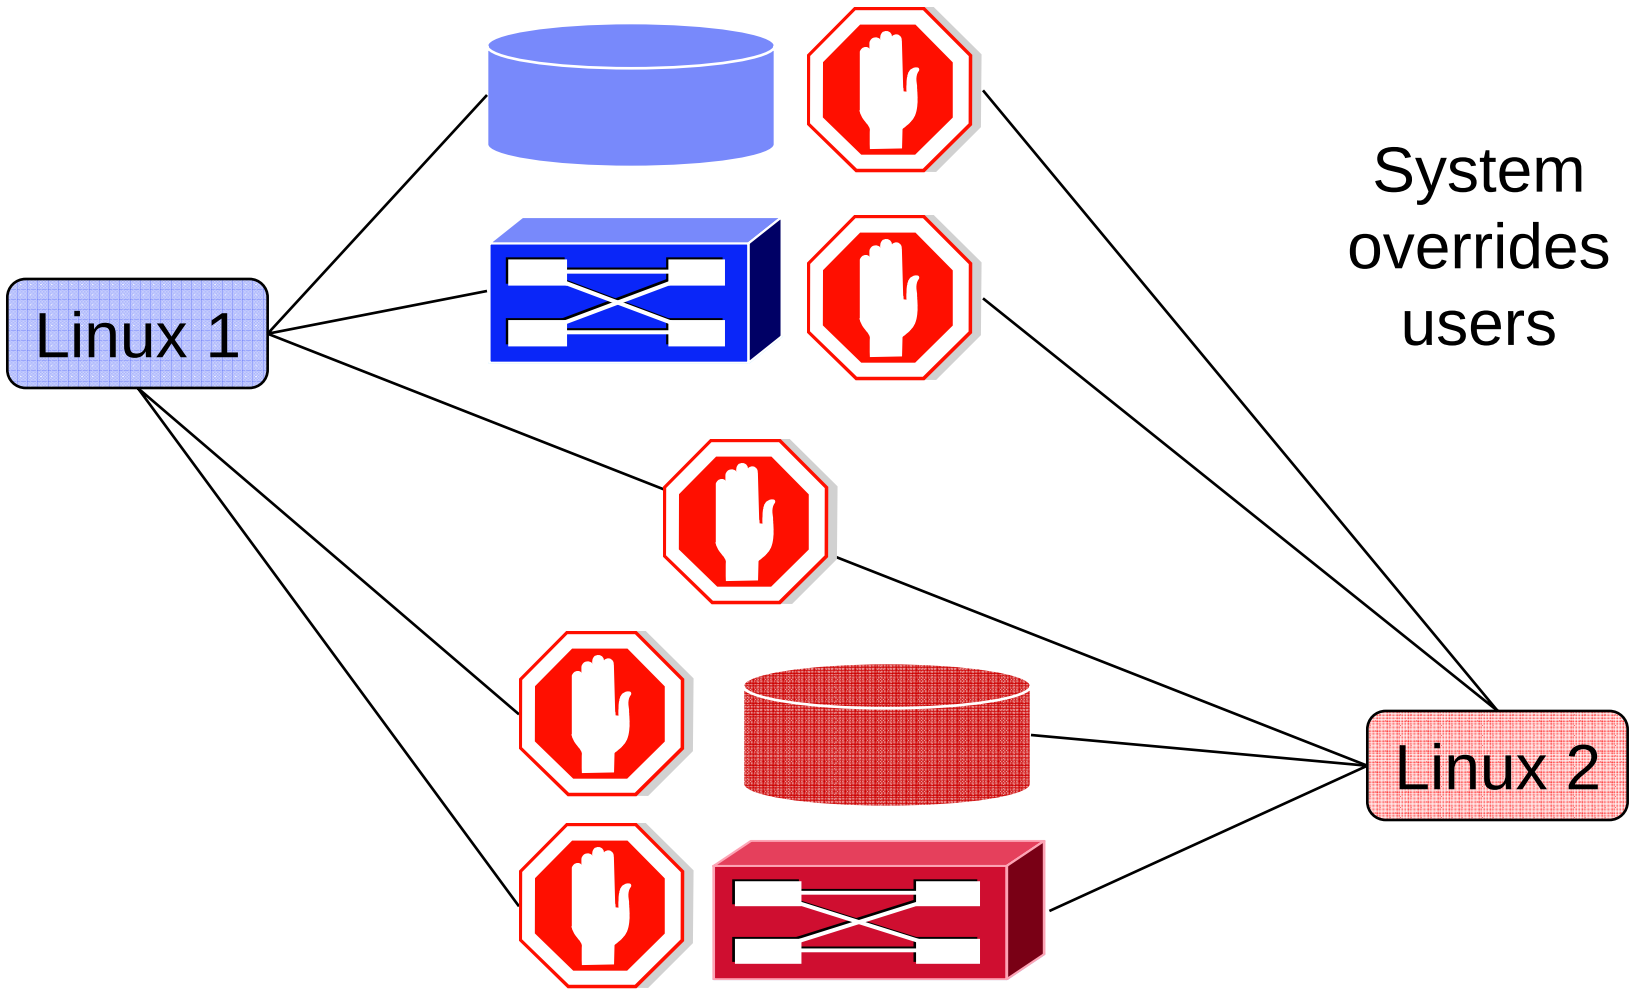

Mandatory Access Controls

- Mandatory access controls override discretionary controls
  - Users are assigned to one or more named projects
  - Minidisks, guest LANs, VSWITCHes, and VLAN IDs all represent data in those same projects
  - Users can only access data in their assigned projects
  - Overrides user- or admin-given permissions

What is System Security?

- Authentication
- Authorization
- Audit

Audit

- Knowing what security-relevant events have occurred
  - Successes (access, not function)
  - Failures (access, not function)
  - Who
  - What
  - When
  - Where

Audit

- The audit trail is management's assurance that the system is being operated according to policy
- It is the most important data asset
  - How do you know that your business data has not had unauthorized out-of-band updates?
- External Security Manager (add-on)
  - Full record of any command or system interface

External Security Manager

- Enhances auditing, authentication, and access controls
- Encrypt user passwords
- Use Access Control List for minidisks instead of minidisk password

ESM Security Controls

- Mixed-case passwords and long phrases
- Virtual Switches and Guest LANs
- VLANs
- Minidisks
- Shared memory
- Shared virtual machines
- Spool files
- Terminals (restricted login)
- Mandatory access controls
  - Multiple security zones (projects)
- Certain commands (e.g. STORE HOST)
- Control Program interfaces
- Full audit: interface, command, virtual machine

Command authorization security flow

- Directory privilege
  - Privilege class
  - Option
- Additional ESM privilege check
- Audit

Security Highlights

- RACF Security Server
- LDAP
- Transparent SSL/TLS
- DIRMAINT
- Virtual Switch

z/VM RACF Security Server

- Pre-installed optional feature of z/VM
- Trusted brand
  - Shared heritage with flagship z/OS version
  - In business since 1976
  - On z/VM since 1986
- Competitive arena
  - Priced separately

LDAP server and utilities

- Enables remote hosts or applications to securely authenticate users against the RACF database on z/VM
  - E.g. Linux PAM
- Enables central management of remote host passwords on z/VM
- Remote audit via LDAP extended operation
- CMS client utilities
  - ldapadd, ldapsrch, ldapmdfy, ldapmrdn, ldapdlet

SSL Server

- Provides transparent SSL/TLS support for client and server applications
  - Any server
- Some applications can negotiate a connection from clear-text to encrypted:
  - Telnet
  - FTP
  - SMTP

SSL Server

- SSL services provided by System SSL
  - Same as z/OS System SSL
  - Exploits CPACF integrated cryptographic function
  - No exploitation of Cryptographic Coprocessors (cards)

SSL Server

- Certificate management via gskeyman
  - Create user certficiates in response to a request
  - Create intermediate CAs and trusted CAs
  - Certificate export, import, renewal
  - Menu driven (linemode, so automation is possible)

DIRMAINT

- Interfaces with RACF
  - User add and delete
  - Password changes (user or administrator)
  - Mindisk create / change / delete
  - Optional, allowing Separation of Duties
- Exits available to override or extend

IBM Commitment

- Continued investment
  - Built on 40+ years of previous investment
  - CP/67
  - Common Criteria (ISO)
- Prompt response to incidents reported to the IBM Support Center

IBM Commitment

- No public disclosure of IBM System z vulnerabilities
  - May disclose to individuals or groups that have demonstrated to IBM a legitimate need to know
- Commitment published in z/VM General Information manual

Common Criteria

- Common Criteria ensures
  - A set of meaningful security functions
    - Access control
    - Audit
  - Extensive testing of those functions
  - Effective processes
  - Good documentation
  - Developed by US National Security Agency
- Assurance levels 1 through 7
  - Evaluation by accredited firms
  - Certification by government agencies
  - CommonCriteriaPortal.org

Common Criteria

- Controlled Access Protection Profile (CAPP)
  - Discretionary access controls
  - “I choose to give you access”
  - User- or administrator-controlled access
- Labeled Security Protection Profile (LSPP)
  - Mandatory access controls (MAC)
  - System overrides user
  - Security clearances and compartmentalization enforced

Customer Commitments

- Define and deploy a security policy
- Examine audit trails periodically
- Apply recommended service

Common Criteria

- z/VM compliance
  - Includes CP, TCP/IP stack with telnet, and RACF
  - First evaluation: z/VM 5.1, October 2005, EAL 3+
  - Second evaluation: **z/VM 5.3**, **August 2008**, EAL 4+
- z/VM 5.4 will **not** be certified.
  - "Designed to meet the requirement"

…but z/VM Security begins with System z security

- Protect the HMC
  - Don’t share user IDs
  - …but don’t be afraid to connect it to your internal network
  - Limit span of control as appropriate
- Protect the I/O configuration
  - Create a separate LPAR that is authorized to modify the I/O config
  - Give partitions access only to devices they require

System z Hardware Security

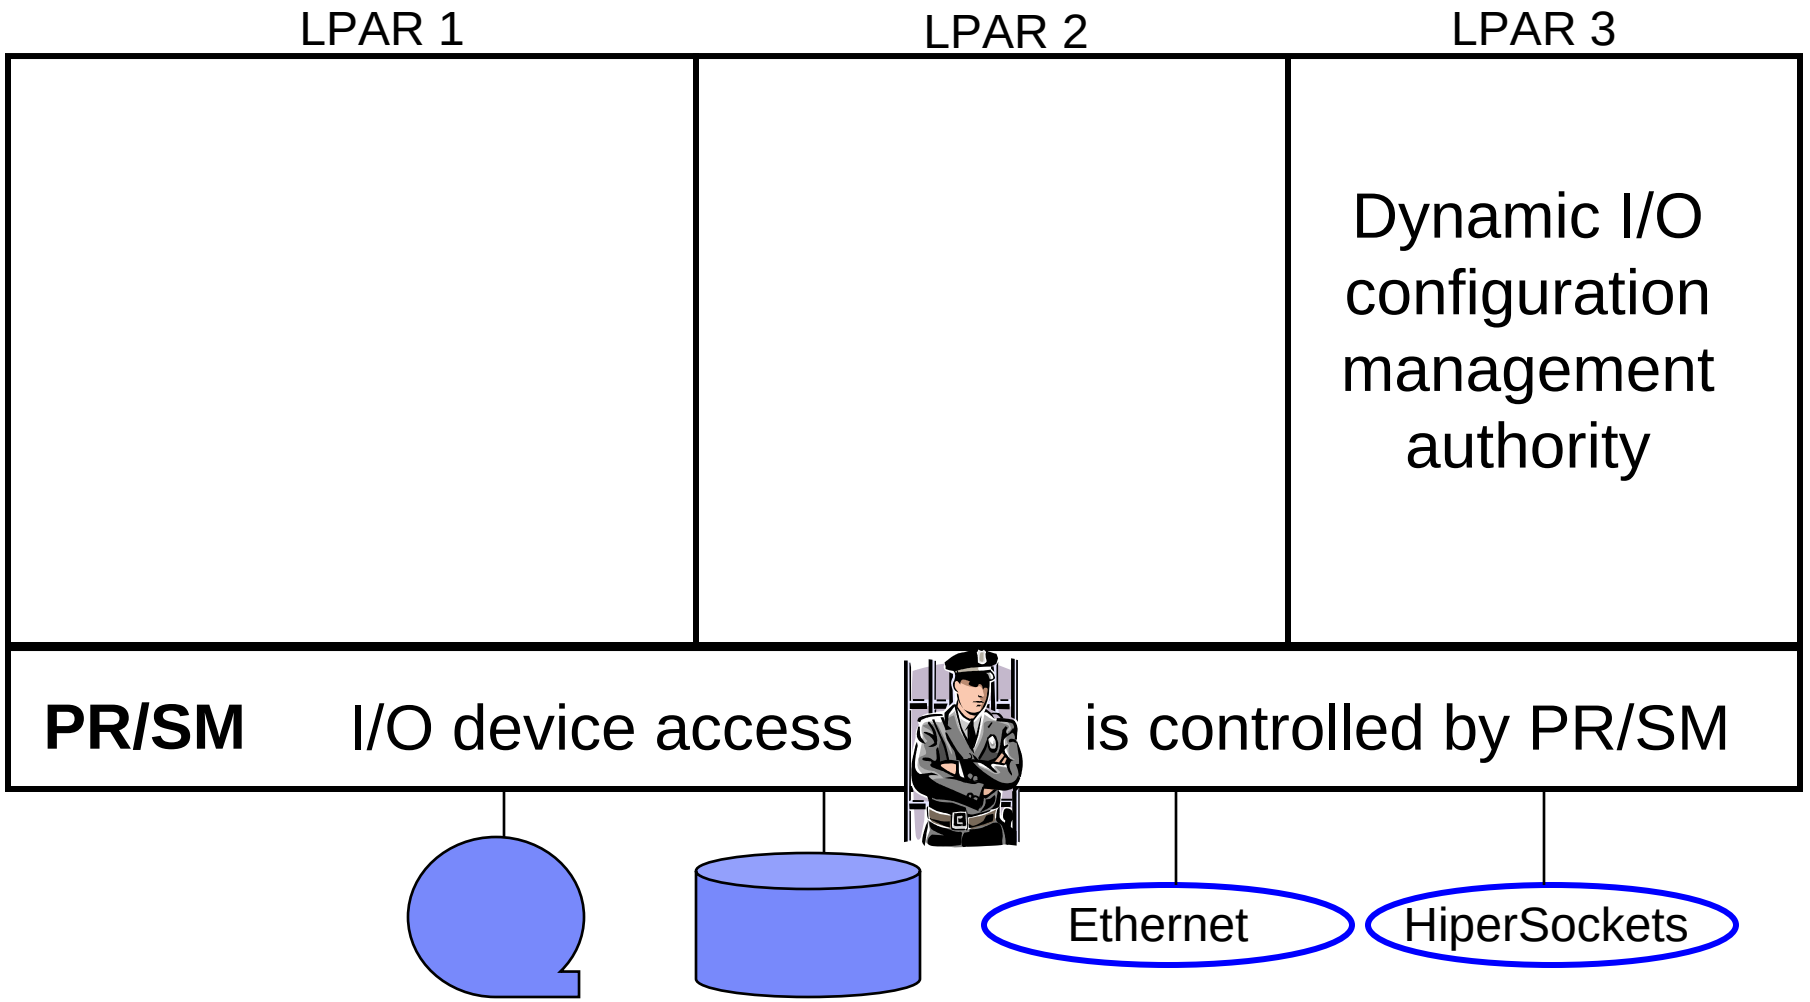

Shared Open Systems Adapters

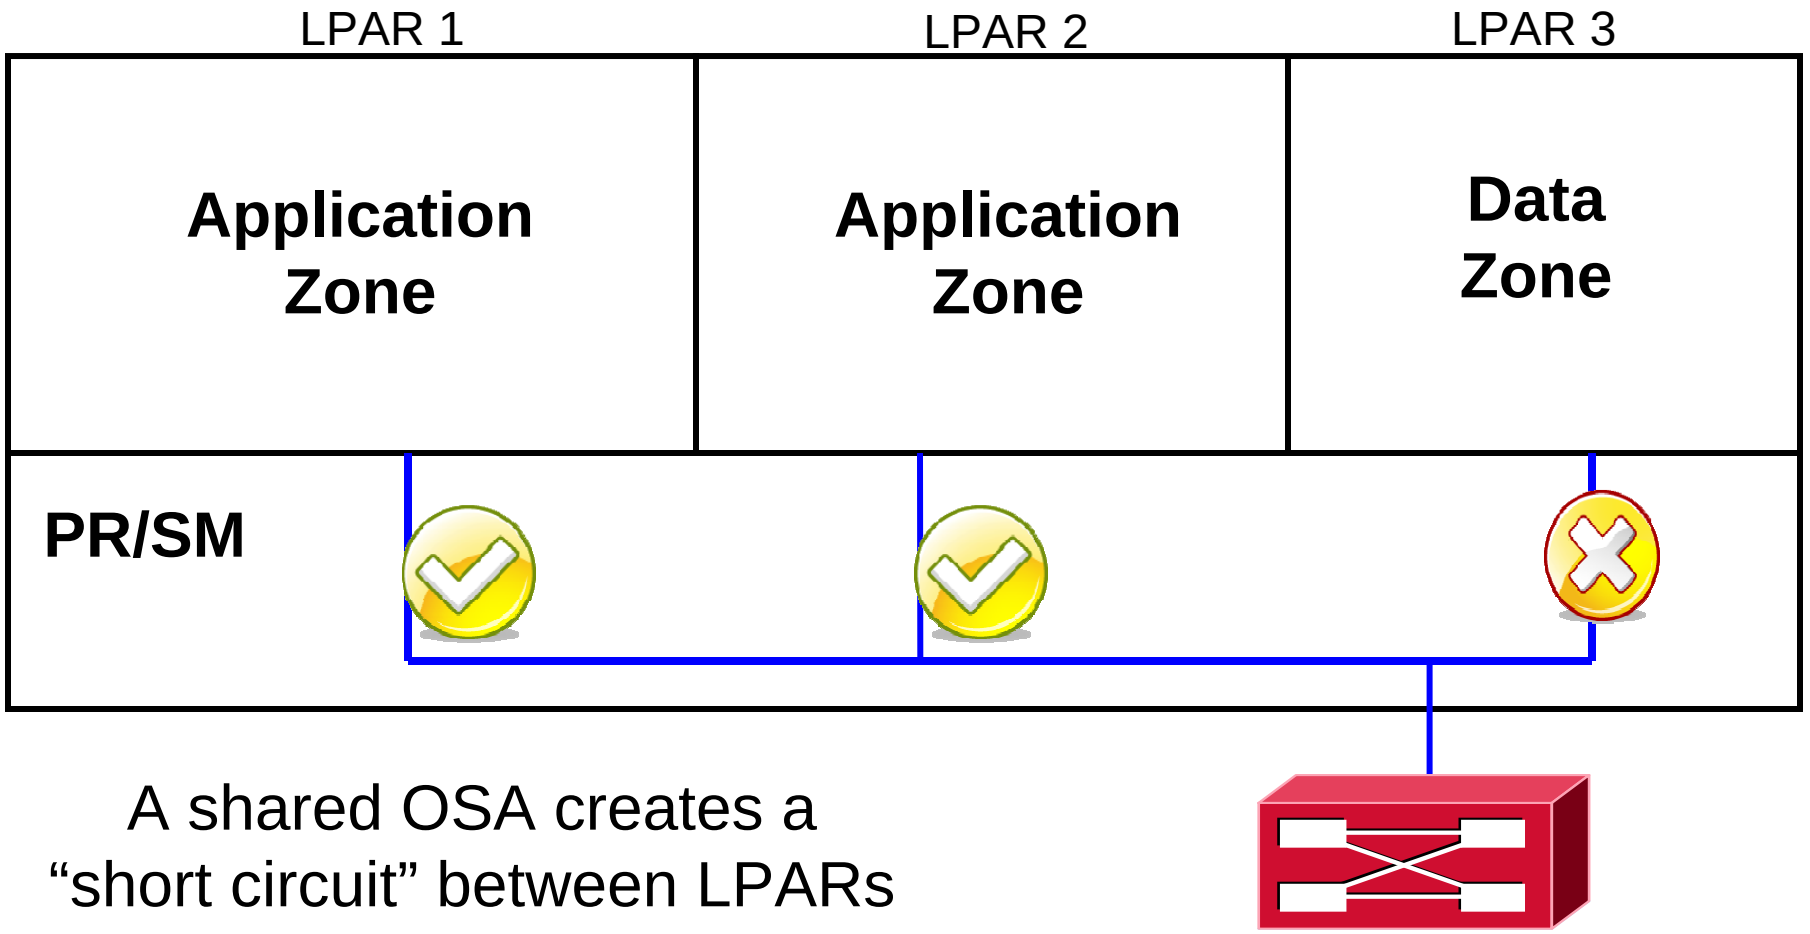

A shared OSA creates a "short circuit" between LPARs

Summary

- z/VM was designed to host virtual machines
- System z hardware provides facilities used by z/VM to ensure the integrity of the system is maintained
- Backed by 40+ years of practical experience in maintaining virtual machines
- IBM commitment
- Customer-defined security policy

Summary

- An external security manager such as RACF Security Server is recommended
  - Privileged command audit trail
  - Encrypted passwords
  - ACLs for minidisks instead of passwords
  - Finer grain of control
- A full discussion of z/VM security and integrity features can be found in publication GM13-0145-01 (April 2005)
  - Link at http://www.VM.ibm.com/security

Reference Information

- z/VM Security resources
  - http://www.VM.ibm.com/security
- IBM Redbook “Security on z/VM”
  - http://www.redbooks.ibm.com/abstracts/sg247471.html?Open
- System z Security
  - http://www.ibm.com/systems/z/advantages/security/
- z/VM Home Page
  - http://www.vm.ibm.com