SHARE – Orlando – August 2011

Session 9564

z/VM Security and Integrity: How it Works

Alan Altmark
IBM Lab Services z/VM and Linux Consultant
Alan_Altmark@us.ibm.com

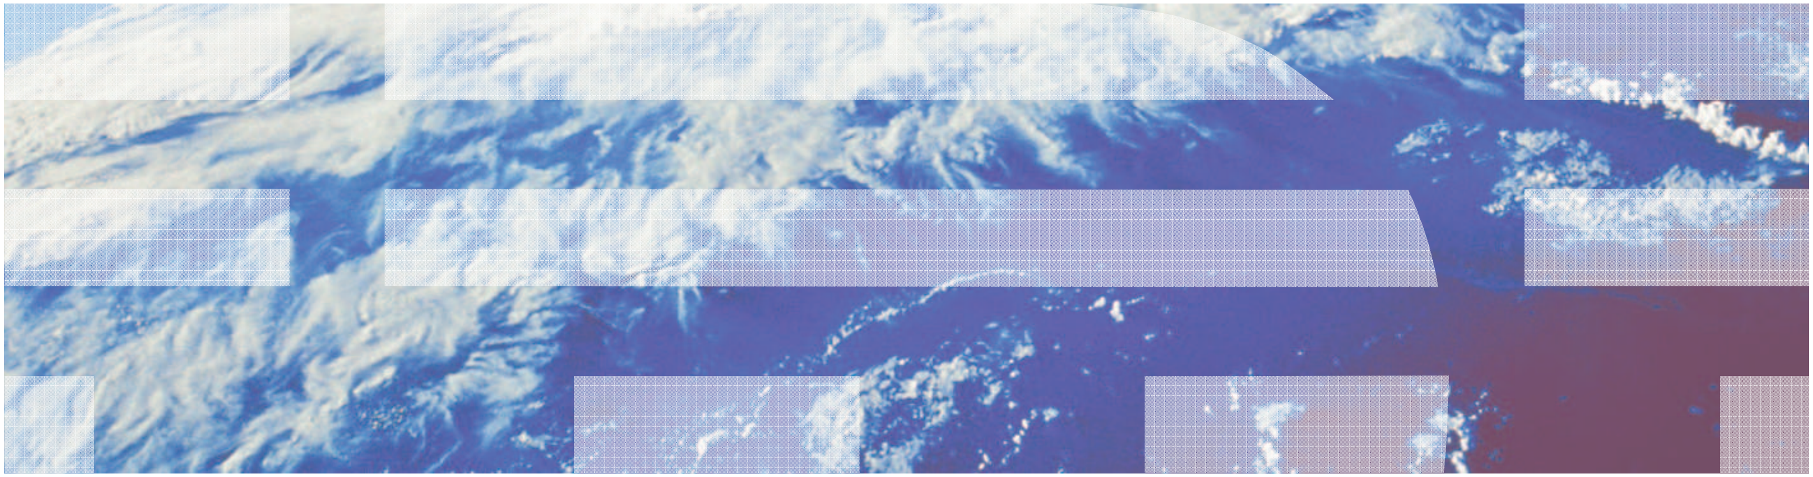

Trademarks

The following are trademarks of International Business Machines Corp., registered in many jurisdictions worldwide.

IBM
System z
z/OS
IBM logo
zEnterprise
z/VM
ibm.com
BladeCenter
System z10
z/Architecture
S/390

Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at “Copyright and trademark information” at www.ibm.com/legal/copytrade.shtml.

Java is a trademark of Sun Microsystems, Inc. in the United States, other countries, or both.

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 and Windows are registered trademarks of Microsoft Corporation in the United States, other countries, or both.

Notes:

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.

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 hypervisor assist
  - Certain program interrupts
- Runs until interception condition raised
- Uses hypervisor-maintained address mapping to convert guest addresses to real addresses
  - Region, Segment, Page
  - Zone offset

Interpretive Execution Facility

- Q. What is a virtual machine?
- A. An execution environment that conforms to the rules specified in the System z Principles of Operation
  - What the virtual machine sees as “real” is a virtual reality created by the underlying *hypervisor*
  - Fixed or Variable memory mapping, with or without overcommittment

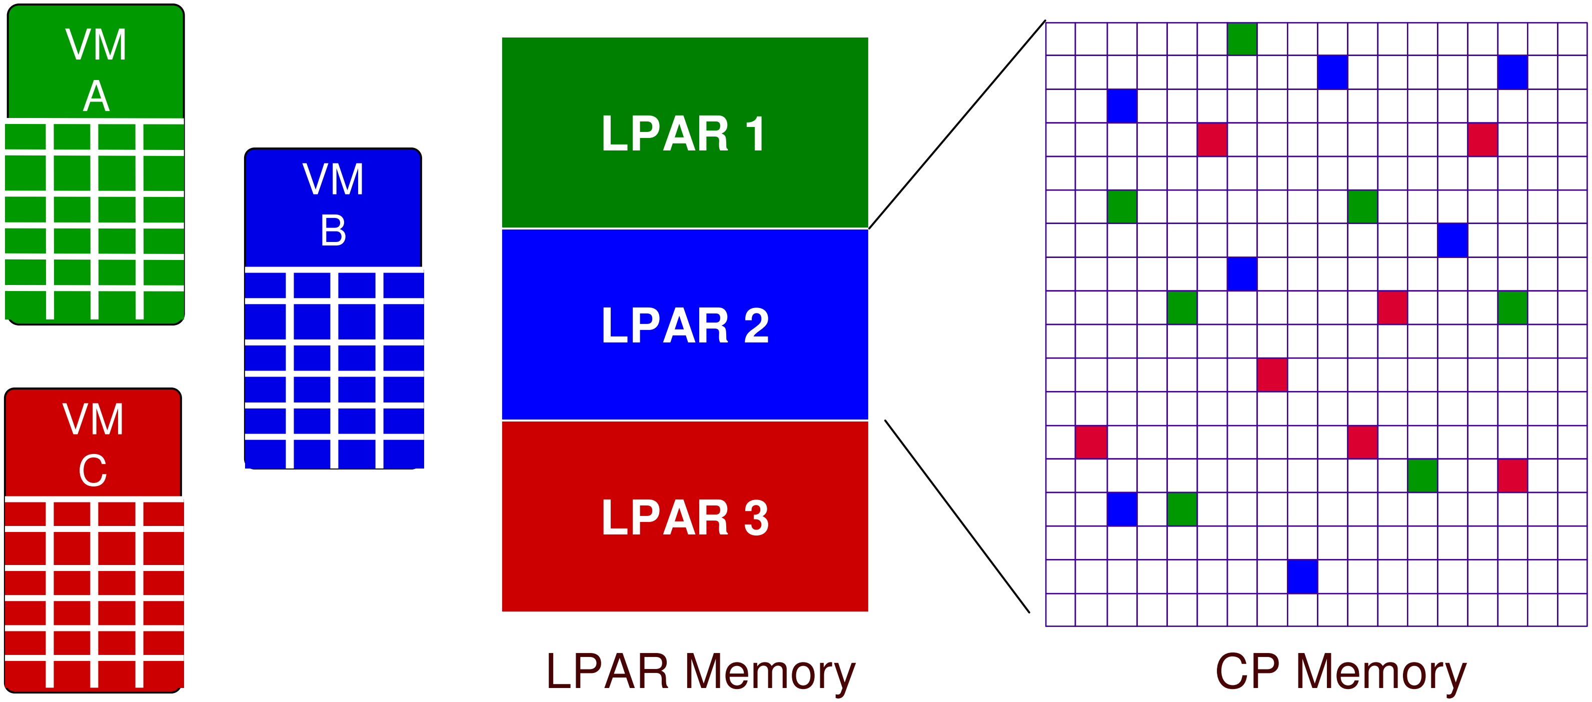

Interpretive Execution Facility

- 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” partitioning (LPAR) and “soft” partitioning (z/VM)

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

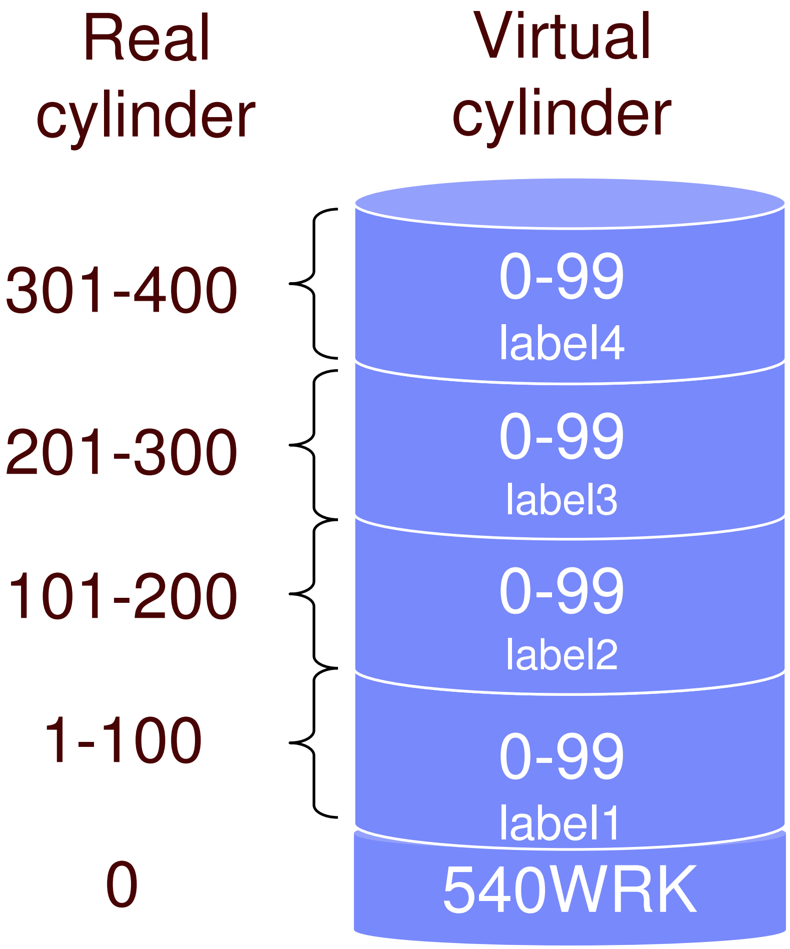

I/O Hardware Assist

- 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 provide a record of those accesses.

What is System Security?

- Authentication
- Authorization
- Audit

Reliably identify the people and processes that access your system and its resources

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
  - …
- z/VM does not provide two-factor authentication

Passwords

- They are stored in clear text in USER DIRECT
- They are obfuscated in the object directory
- External Security Managers such as RACF provide for secure, encrypted passwords
- Password phrases require External Security Manager
  - 100 characters
  - Mixed case
  - Special characters
  - Blanks

What is System Security?

- Authentication
- Authorization
- Audit

Ensure that a user has access only to system resources specifically permitted and within scope of responsibility

Applies to commands, interfaces, and data

Command / Function authorization security flow

- Directory privilege
  - Privilege class
  - Option

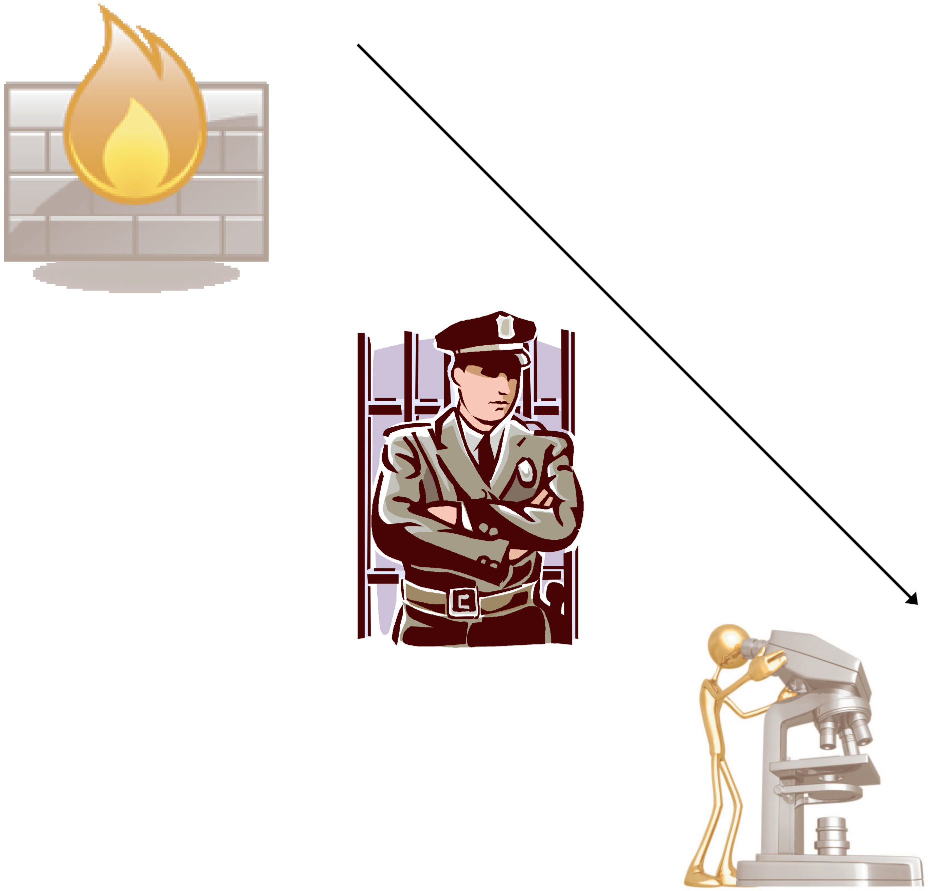

- Additional ESM privilege check

- Audit

Privilege Class

- First line of protection is the **privilege class**
  - A to Z and 1 to 9
    - IBM uses A through H
  - Each class identifies a defined set of commands and DIAGNOSE functions
  - Class “any” functions can be used by any virtual machine without regard to the virtual machine’s privilege class

- Specified in USER DIRECT

- Most virtual machines have class G (“general”)

- Trusted virtual machines have class A, B, C, D, and/or E
  - Potential to bypass system integrity and security controls
  - Give only to system administrators and trusted servers

Matrix Management of Privilege

Command or function	**A**	**B**	**C**	**D**	**E**	**F**	**G**	**Any**
QUERY TIME							X	
QUERY NAMES							X	
SHUTDOWN	X							
ATTACH		X						
QUERY RDR ALL				X				
SAVESEG (create a shared memory object)					X			
LOGOFF							X	
LOGON								X
DISPLAY HOST (hypervisor) memory			X					
STORE HOST memory			X+					
DIAG 0x64 - Access a shared memory object							X	
DIAG 0x08 – Issue a hypervisor command								X
…								

+ Can be restricted to specific virtual machines via ESM

Privilege Class

- Excess privilege is the root of all Evil
  - DO NOT give extra privilege to untrusted virtual machines!
- DO use the COMMAND statement in the directory
  - Runs with all privileges
- DO use automation
- DO change the privilege classes assigned to commands and DIAGNOSE instructions
  - `MODIFY CMD SHUTDOWN PRIVCLASS S`
  - `MODIFY CMD QUERY SUBCMD NAMES IBMCLASS G PRIVCLASS Z`
  - `MODIFY DIAGNOSE 8 PRIVCLASSES ABCDE`
    - Bad idea!

z/VM native resource access controls

- Virtual Switches and Guest LANs
  - SET VSWITCH NET9 GRANT ALAN
  - Associate with a VLAN: VLAN 10
  - Allow sniffing: PROMISCUOUS
  - Disallow sharing: ISOLATION DROP
- Minidisks
  - Read, write, multi-write passwords
  - LINK in USER DIRECT automatically grants access
- Shared memory
  - DEFSYS/DEFSEG with RSTD option
  - NAMESAVE statement in USER DIRECT grants access
- Shared virtual machines
  - LOGONBY statement (maximum 8 users)

z/VM native resource access controls

- Special passwords
  - NOLOG: User cannot logon or be authenticated
  - NOPASS: Password not required
  - LBYONLY: Accessible only via LOGON BY
  - AUTOONLY: User can only be XAUTOLOGed; no authentication possible

Network: VLAN-aware Virtual Switch

Each guest authorized to one or more VLANs

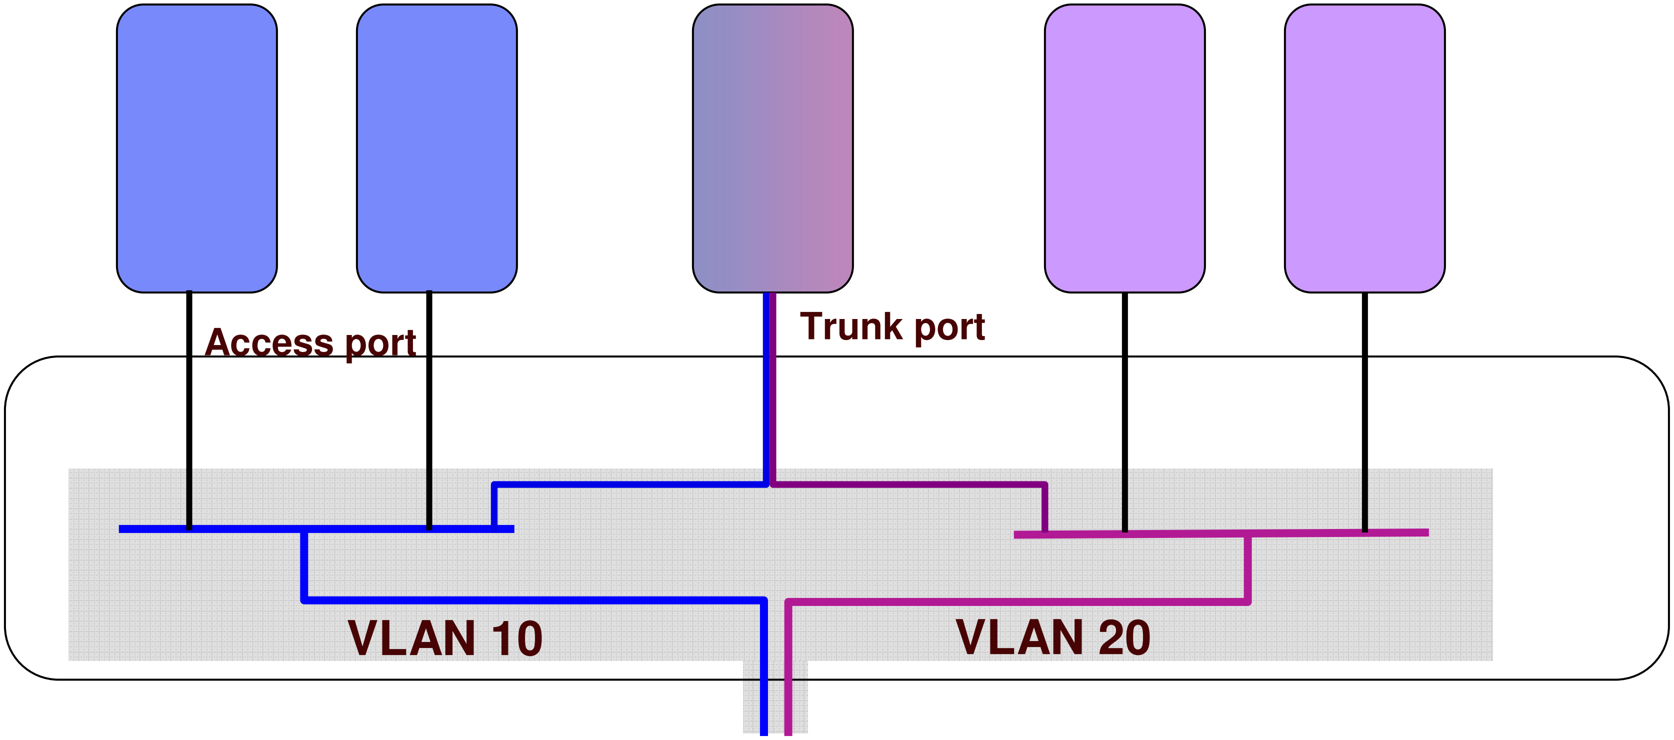

OSA plugged into
[real] trunk port

External Security Manager resource access controls

- Virtual Switches and Guest LANs
- Minidisks
  - ACLs
- Shared memory
- Shared virtual machines
  - LOGON BY
- Spool files
- Terminals (restricted login)
- Mandatory access controls
  - Multiple security zones (projects)
- Further restrictions on certain commands and functions (e.g. STORE HOST)

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

Mandatory Access Controls

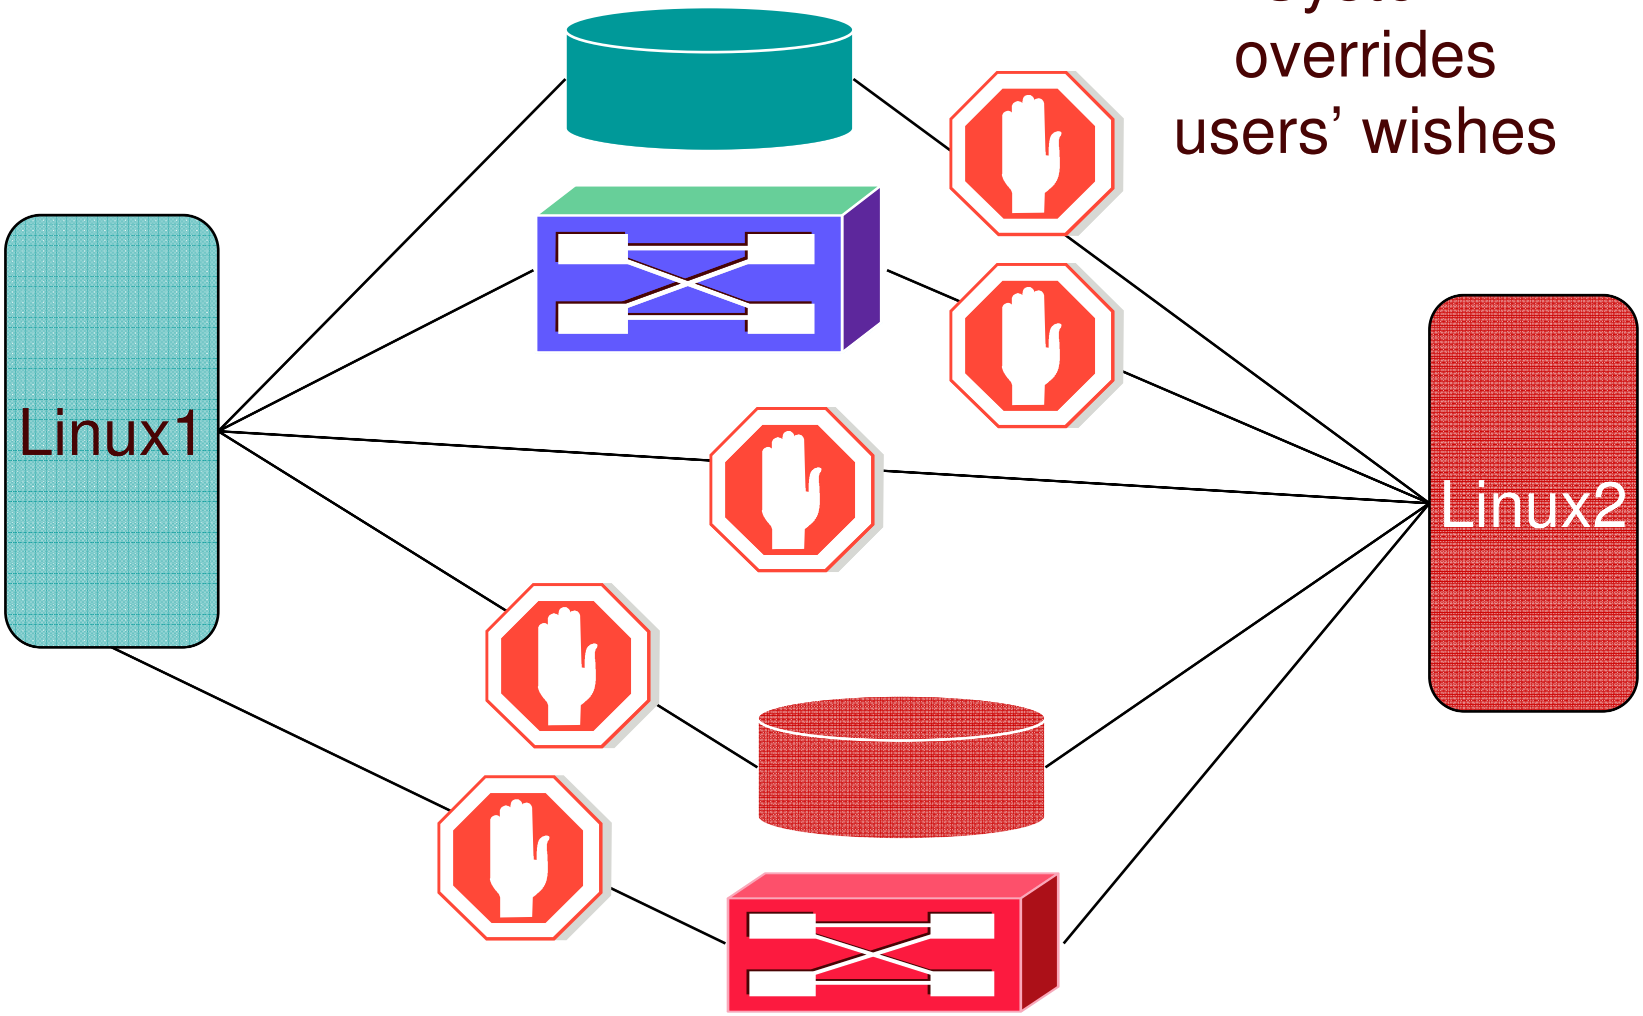

System overrides users' wishes

What is System Security?

- Authentication
- Authorization
- Audit

Knowing what security-relevant events have occurred

- Who has entered the system
- Who was denied access to the system
- Resources accessed
- Resources denied
- Functions used
- Where
- When

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
  - Full record of any command or system interface
  - Includes reporting tools

- CP records some activity in the accounting data
  - Logon, logoff, link to disk, dedicated device usage, APPCVM CONNECT, SET PRIVCLASS, virtual network traffic
  - Minimal controls over what is collected

Security Components

- RACF Security Server
- LDAP
- SSL/TLS
  - Telnet
  - FTP
  - SMTP
  - Transparent SSL/TLS also available
    - E.g. for a simple web server
- DIRMAINT

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

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 was **not** certified.
  - “Designed to meet the requirement”
- z/VM 6.1 certification is underway!!
  - Operating System Protection Profile with labeled security extensions, EAL 4+

Customer Commitments

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

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 more than 40 years of practical experience in maintaining virtual machines
- IBM’s commitment remains strong
- Customer-defined security policy is key to success

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

- Alan Altmark
  - Alan_Altmark@us.ibm.com
- z/VM Secure Configuration Guide
- 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