VSWITCH
für komplexe z/VM Umgebungen
Session VM05

Dr. Manfred Gnirss
Senior IT Specialst
Technical Marketing Competence Center Europe,
R&D Support Support Centers Boeblingen
IBM Deutschland Research & Development GmbH
gnirss@de.ibm.com

Trademarks

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

IBM*
IBM logo*
System Storage*
System z*
System z9*
System z10*
z9*
z10
z/OS*
z/VM*

* 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.

Acknowledgement

My very best thanks belong to

Alan Altmark

for his input to this presentation

Agenda

- Introduction
- A multi-zone network
- Securing System z hardware
- Firewalls
- VLANs and traffic separation
- Enforcing the rules




The Myth of Mainframe Security

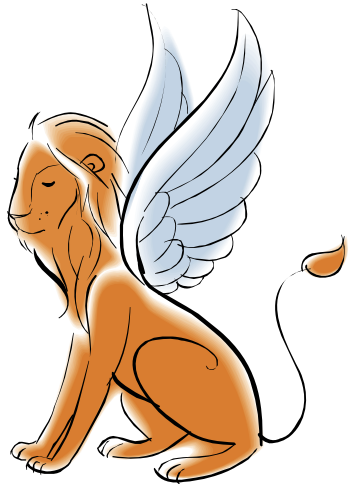


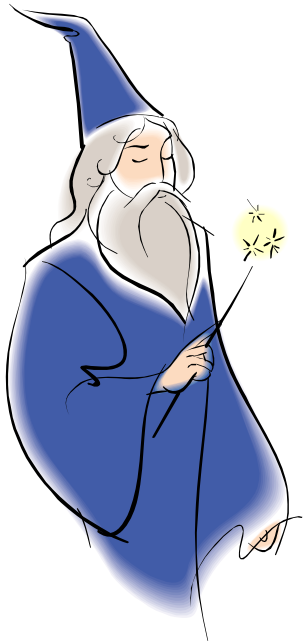


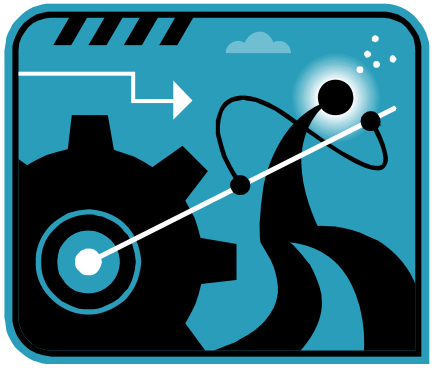

The Reality of Mainframe Security

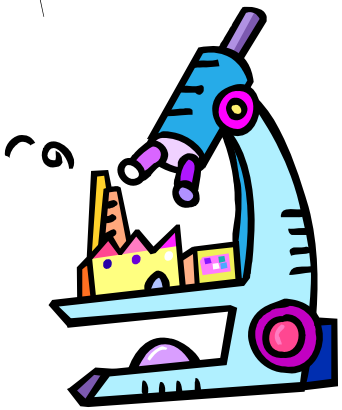

Multi-zone networks

Multi-zone Network

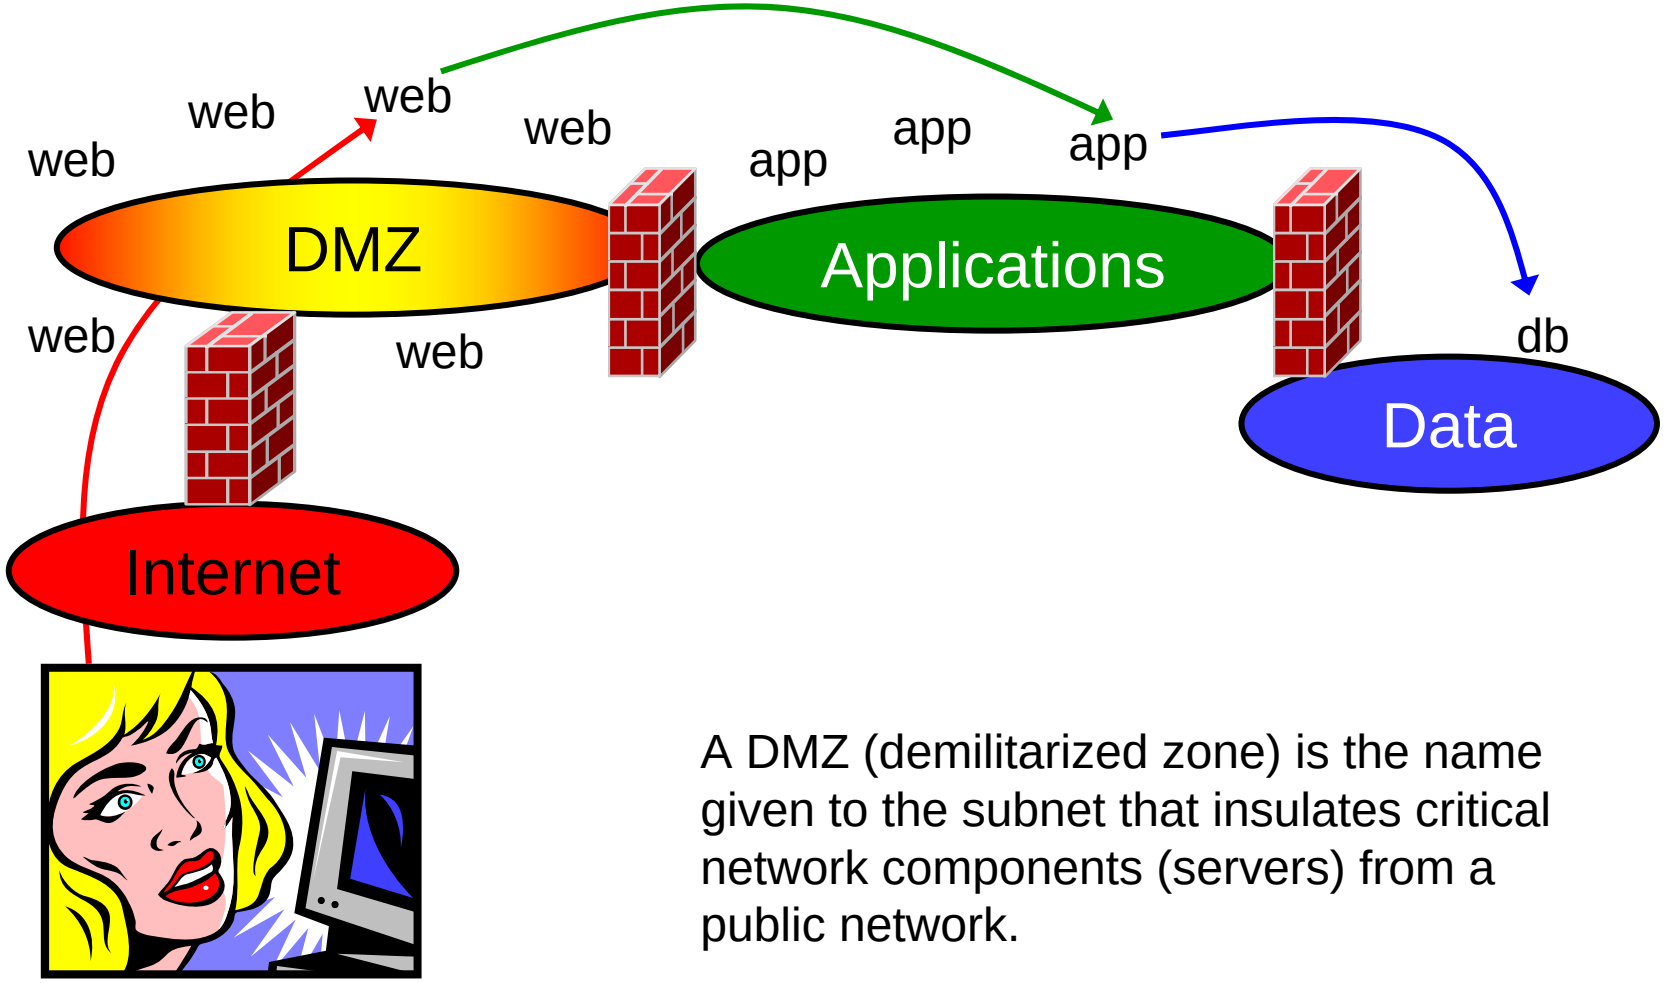

A DMZ (demilitarized zone) is the name given to the subnet that insulates critical network components (servers) from a public network.

Multi-zone Network on System z

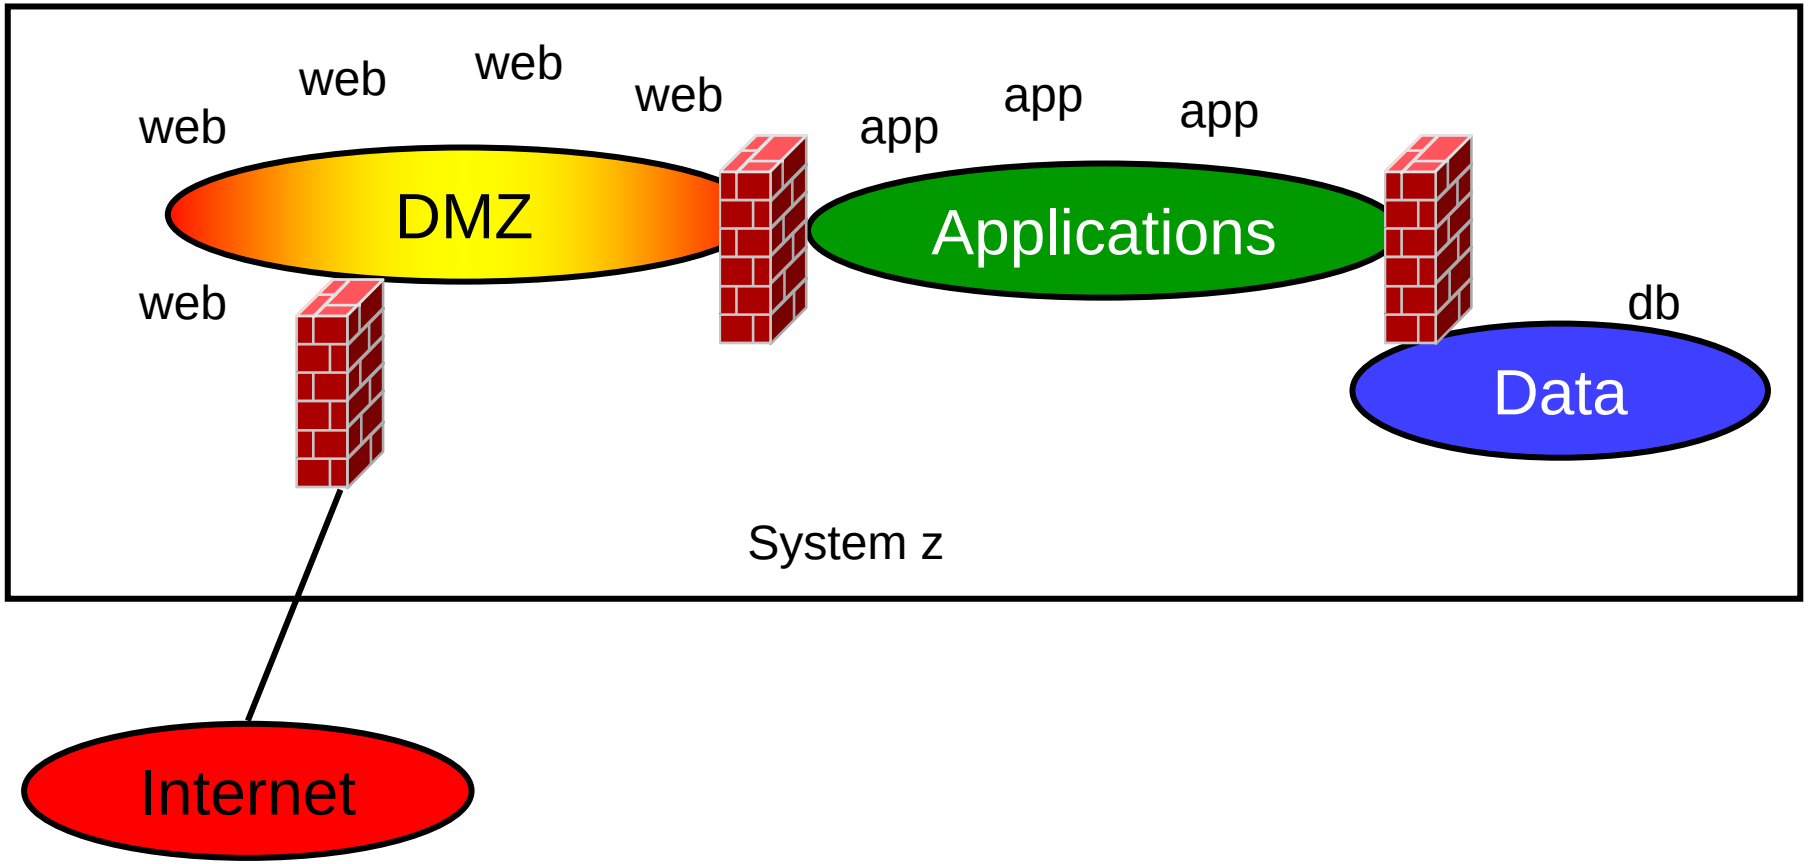

Securing the Hardware

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

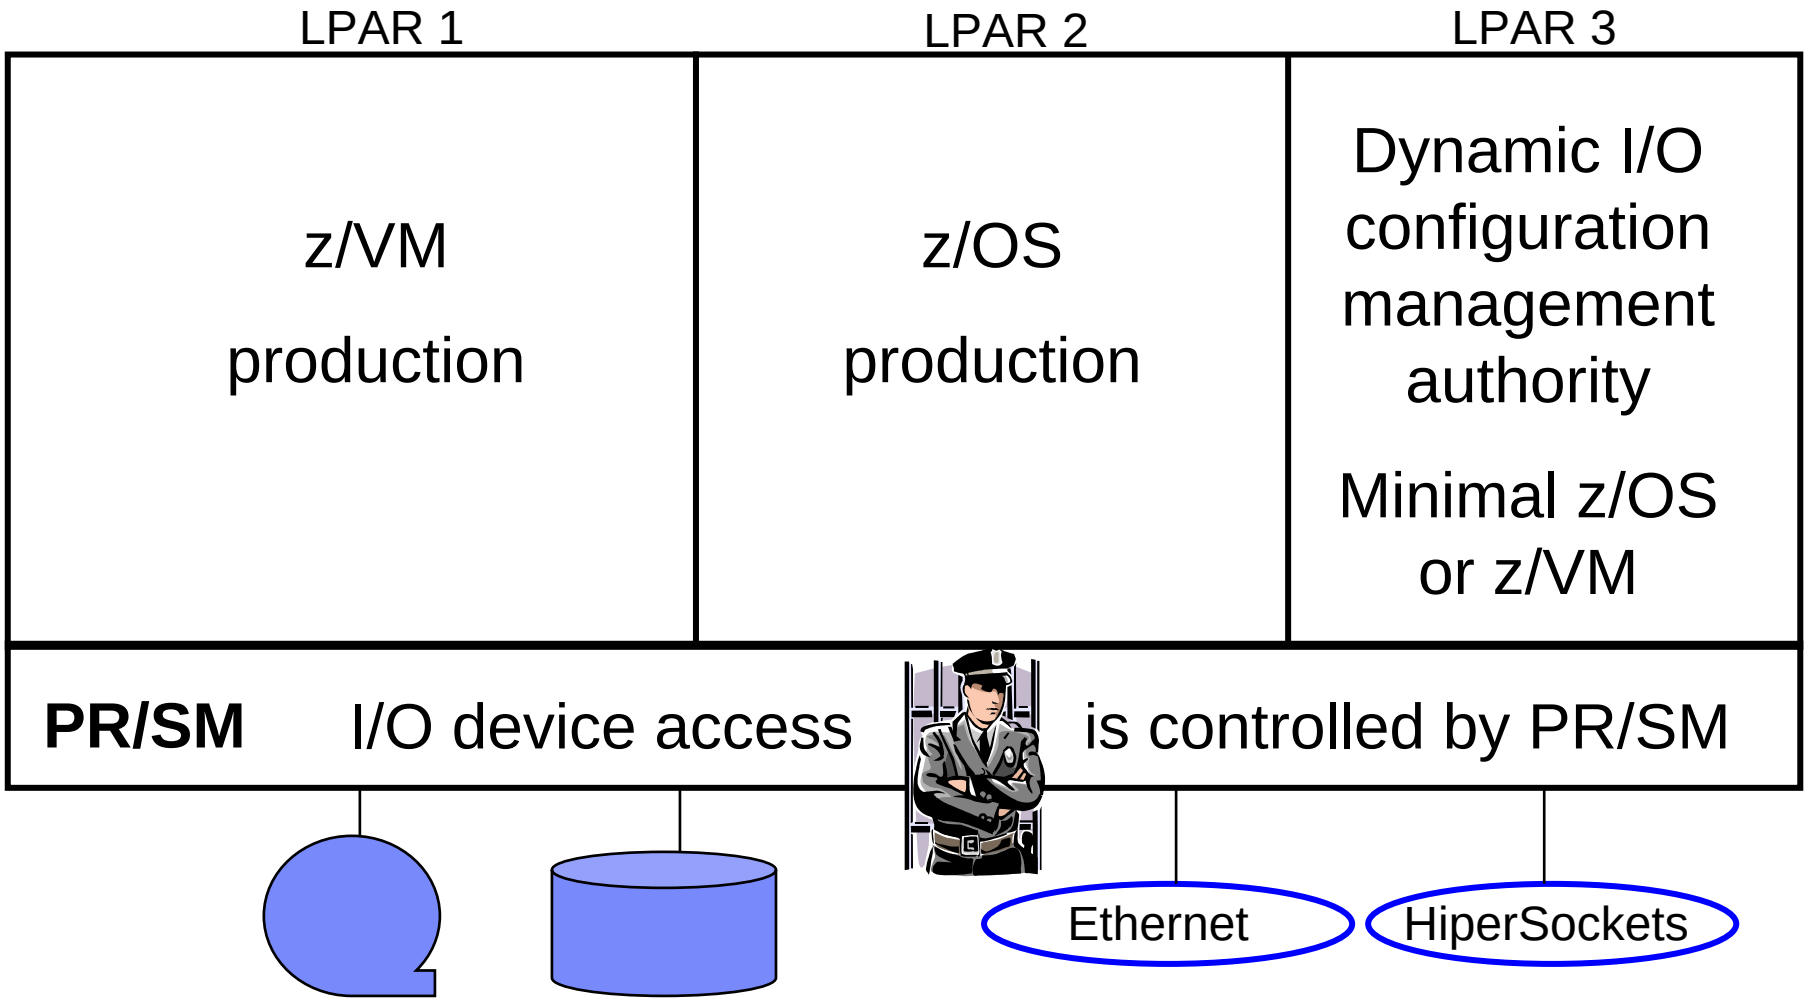

WARNING: Shared Open Systems Adapters

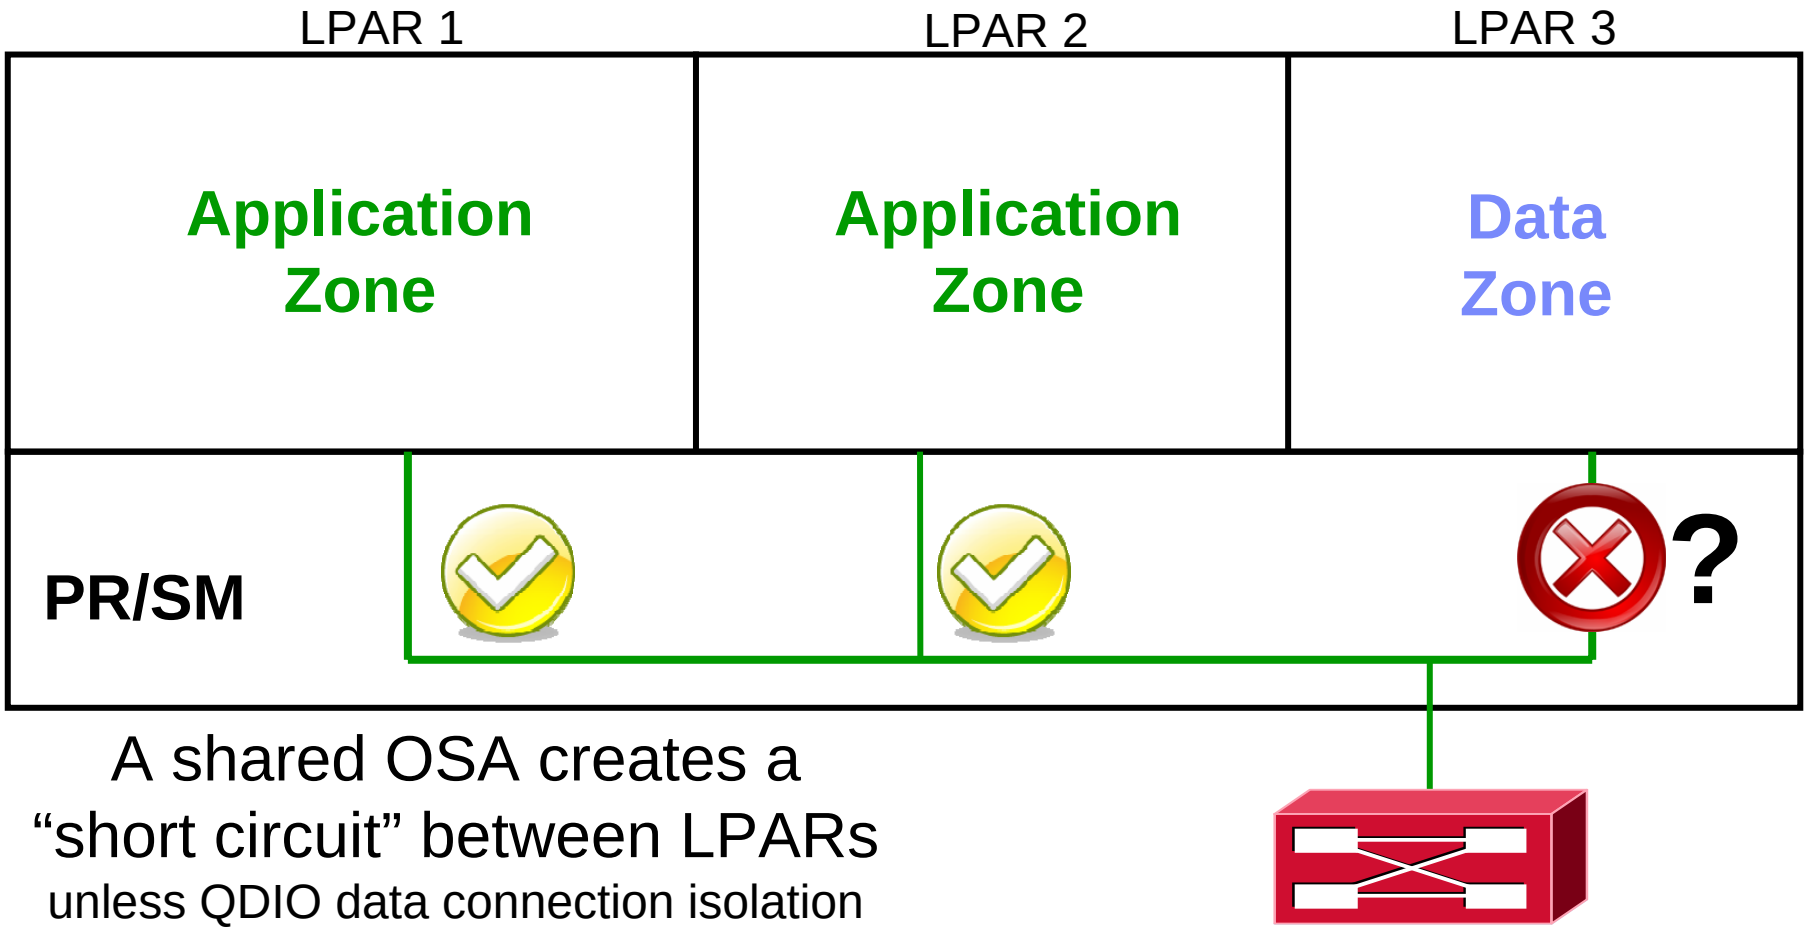

A shared OSA creates a
"short circuit" between LPARs
unless QDIO data connection isolation
is used on z10

WARNING: HiperSockets

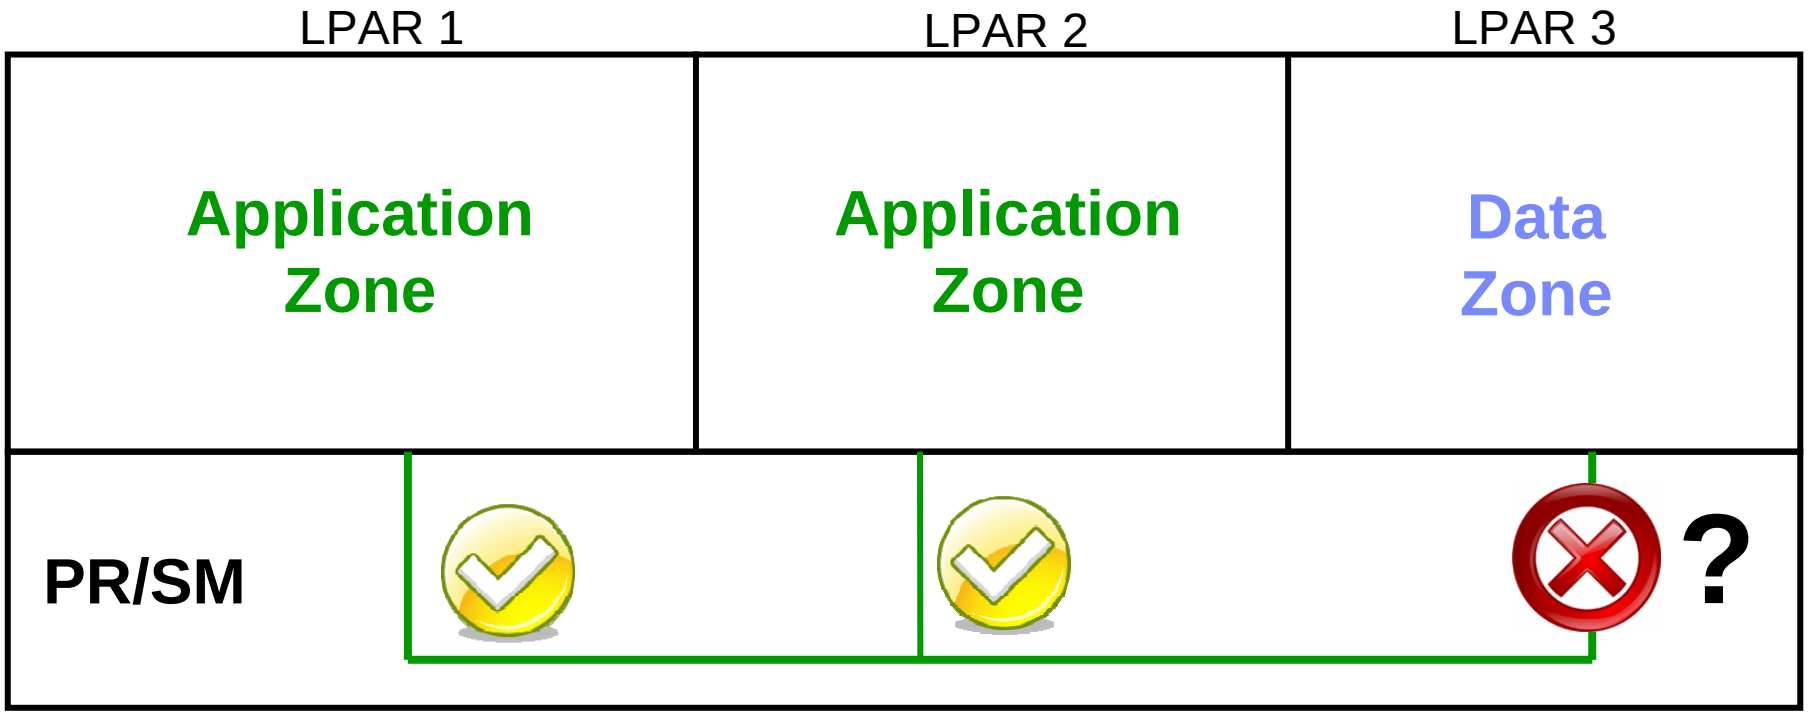

A HiperSocket is a LAN segment.

Treat is like one.

Firewalls

“Where, oh, where has my firewall gone?”

Inboard (internal) firewalls

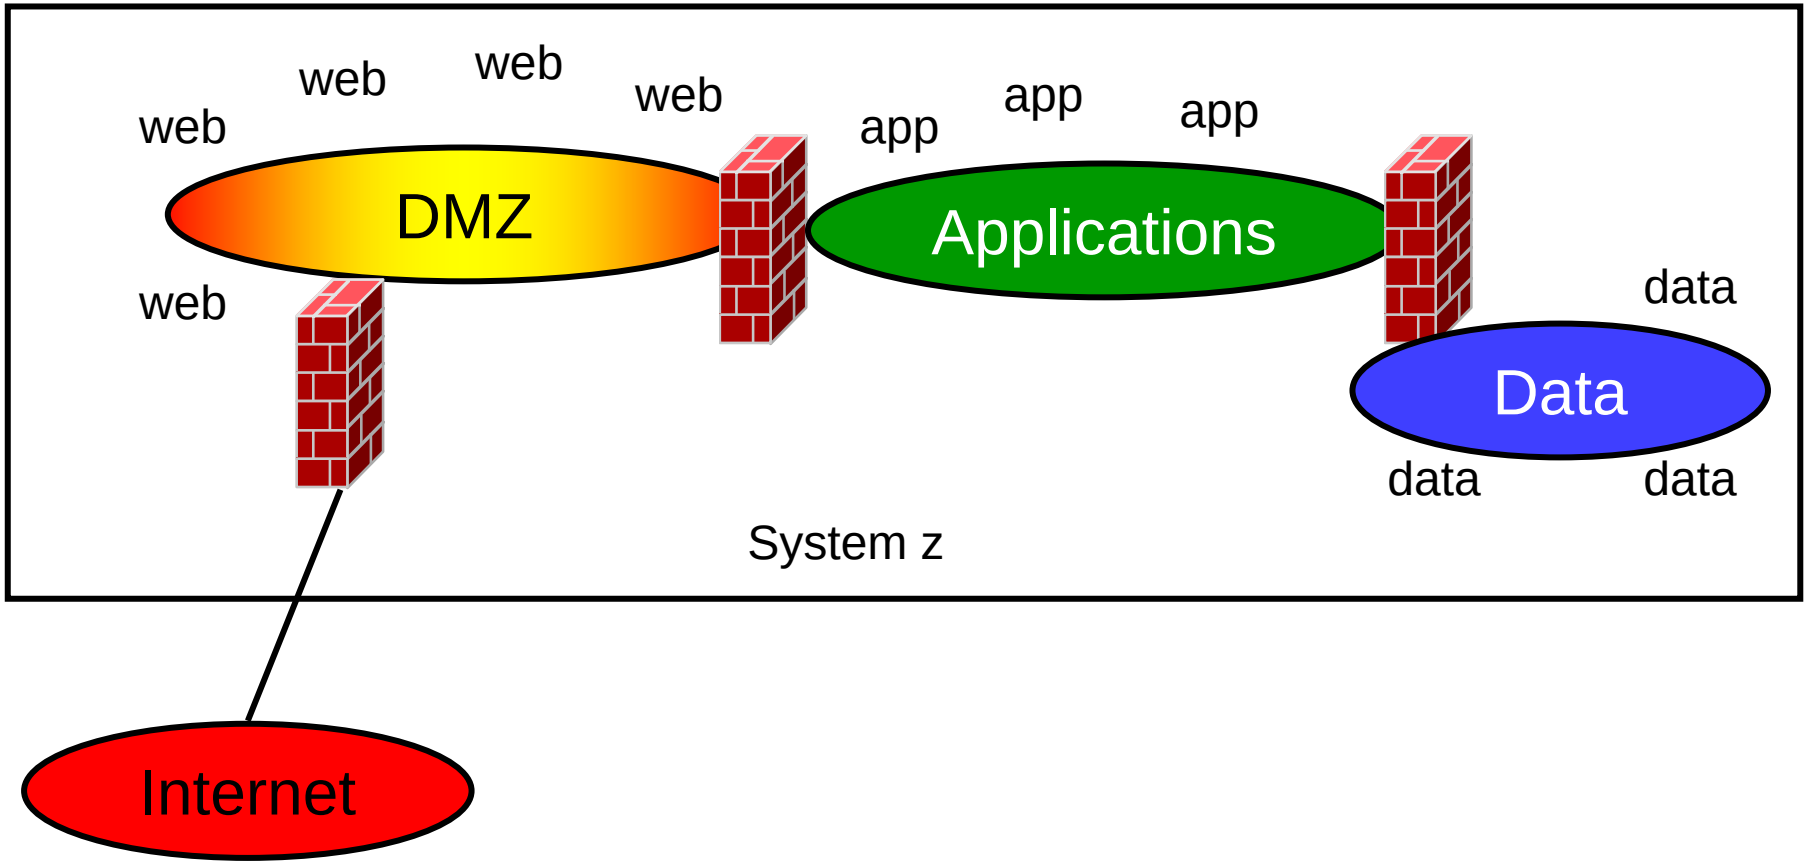

Outboard (external) firewalls

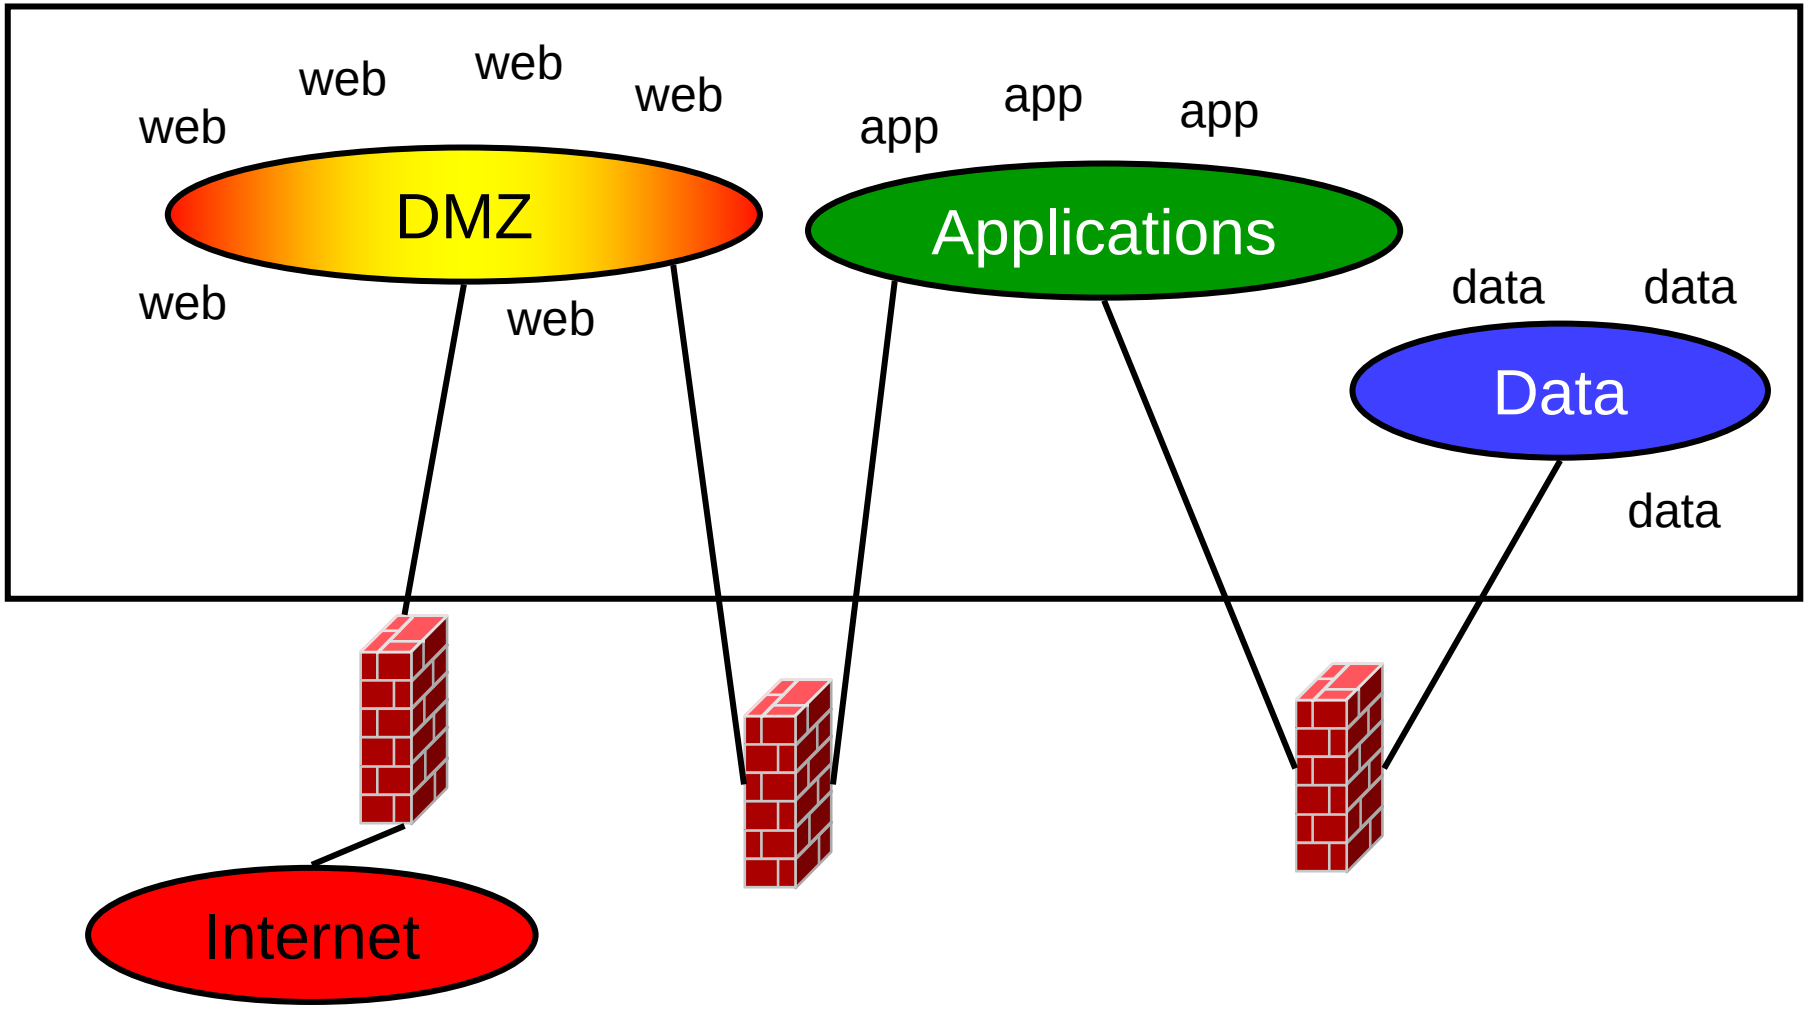

Combination firewalls

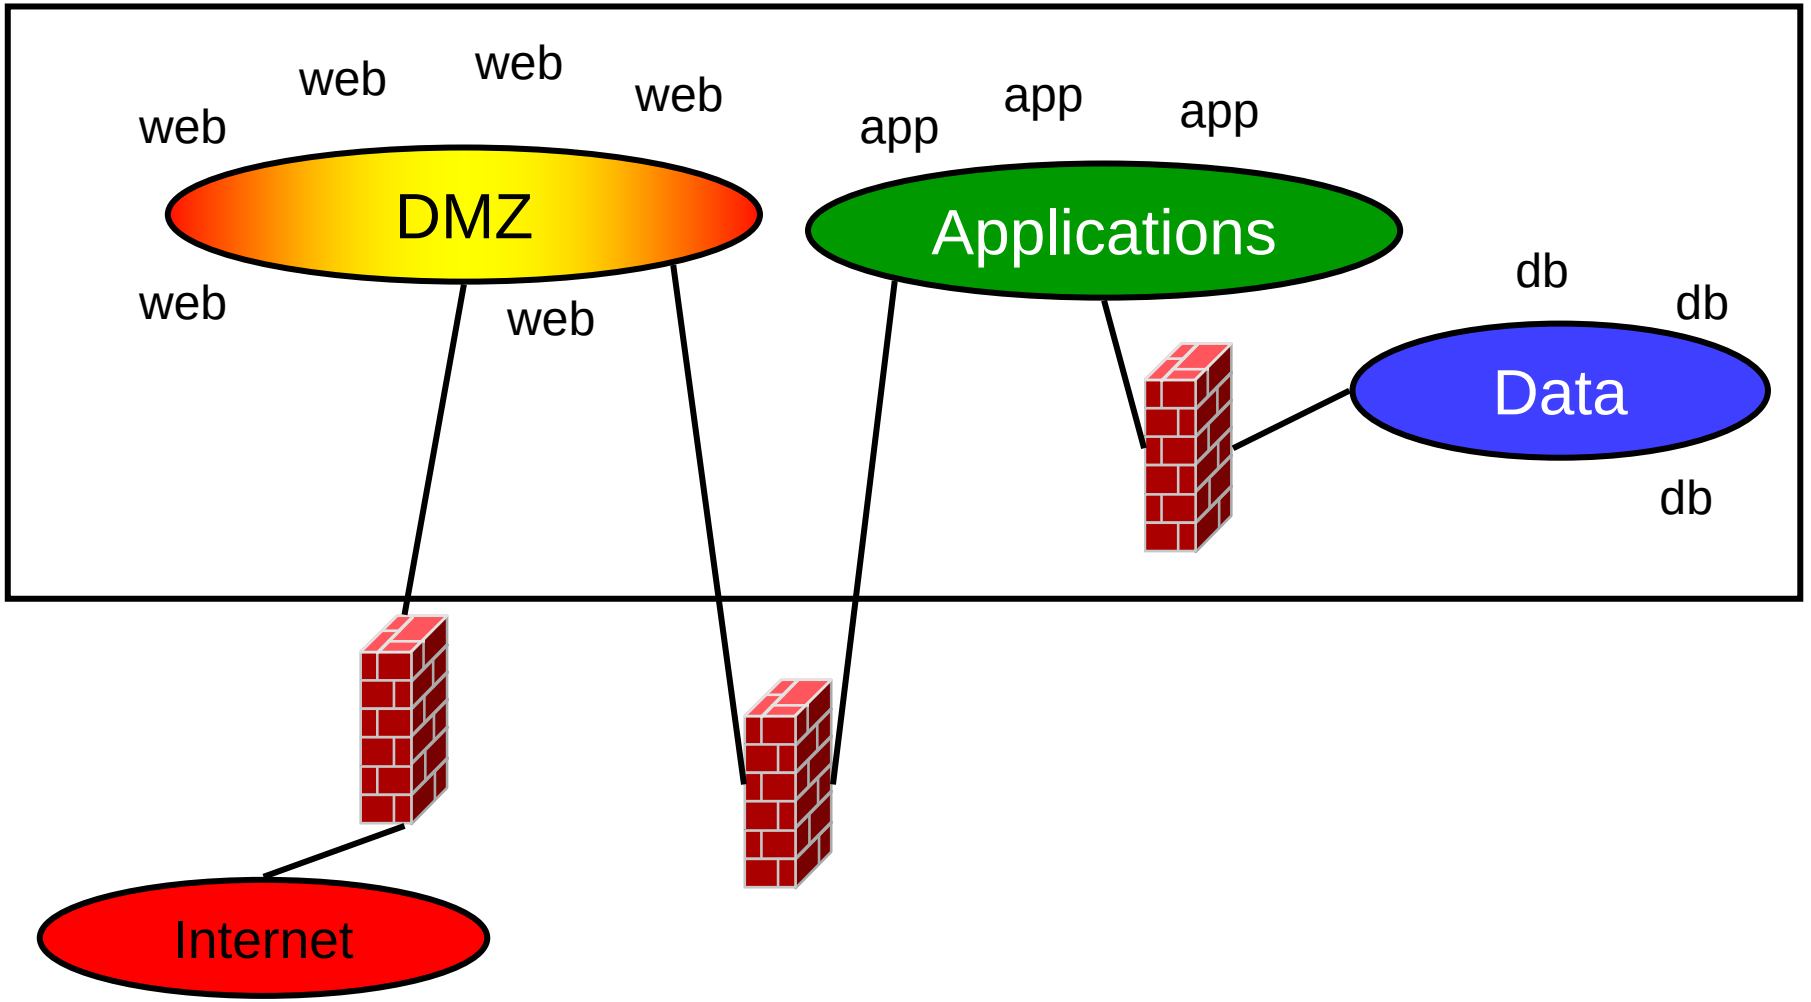

Guest LANs with HiperSockets

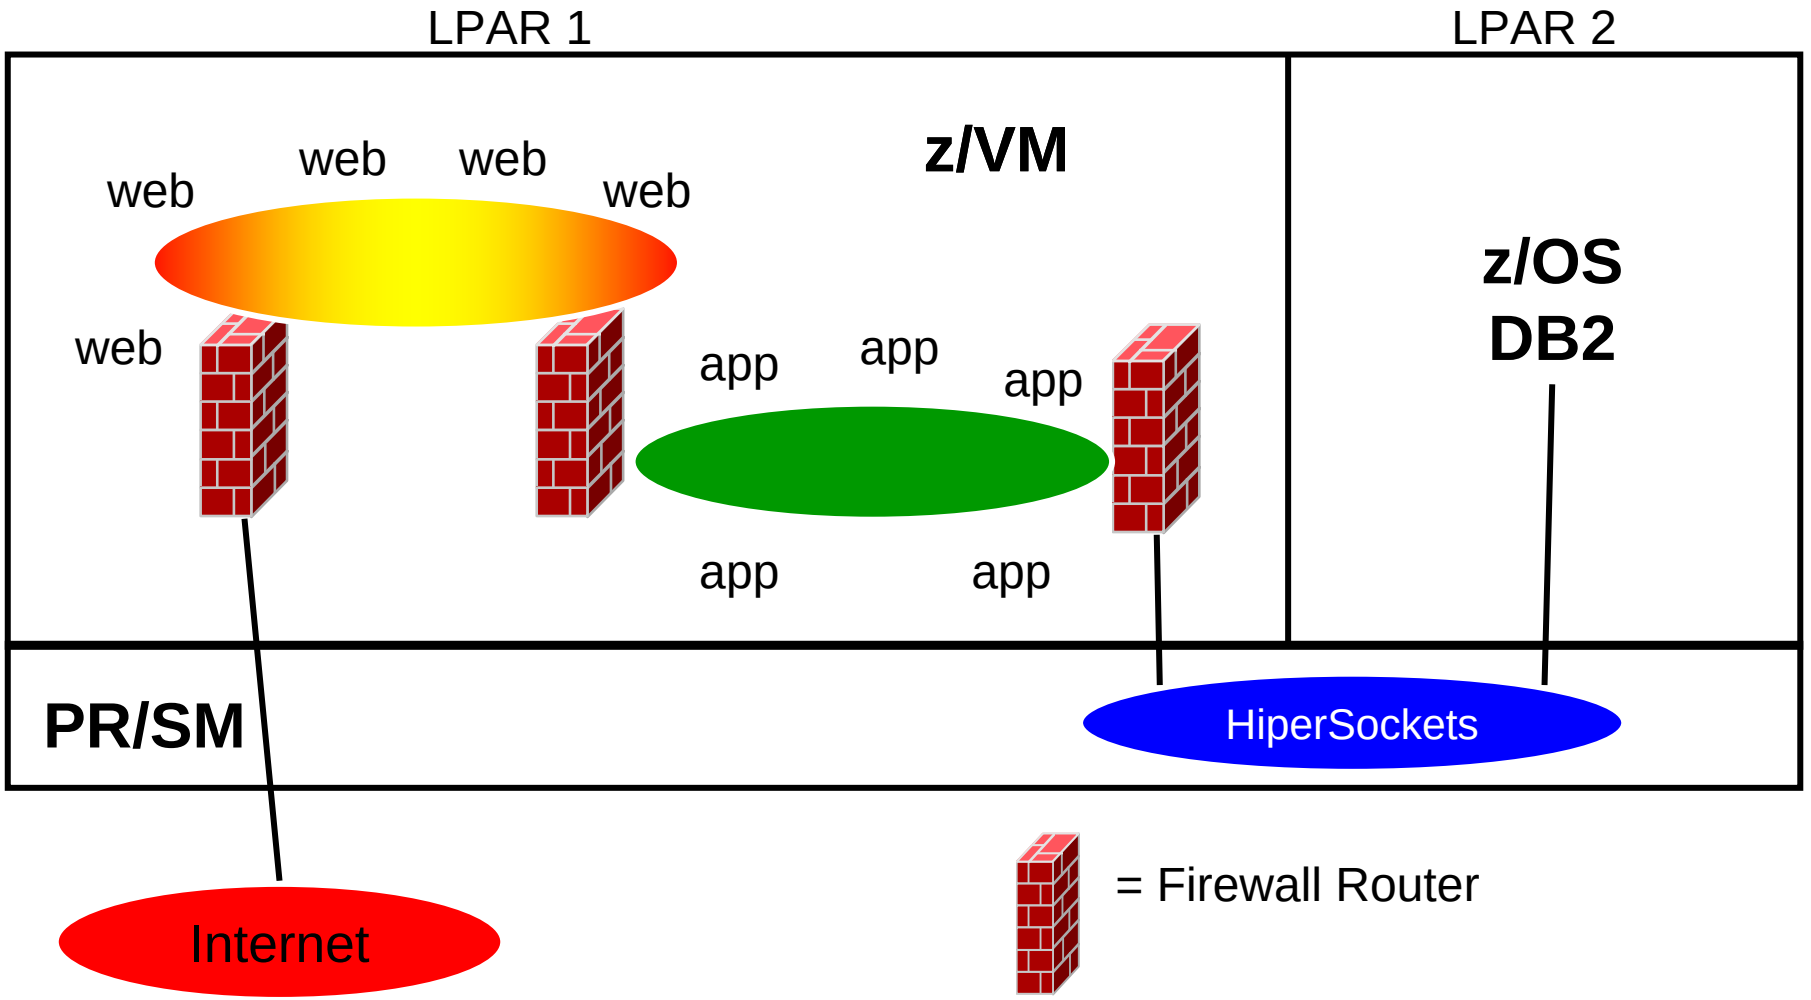

HiperSockets & z/OS packet filters

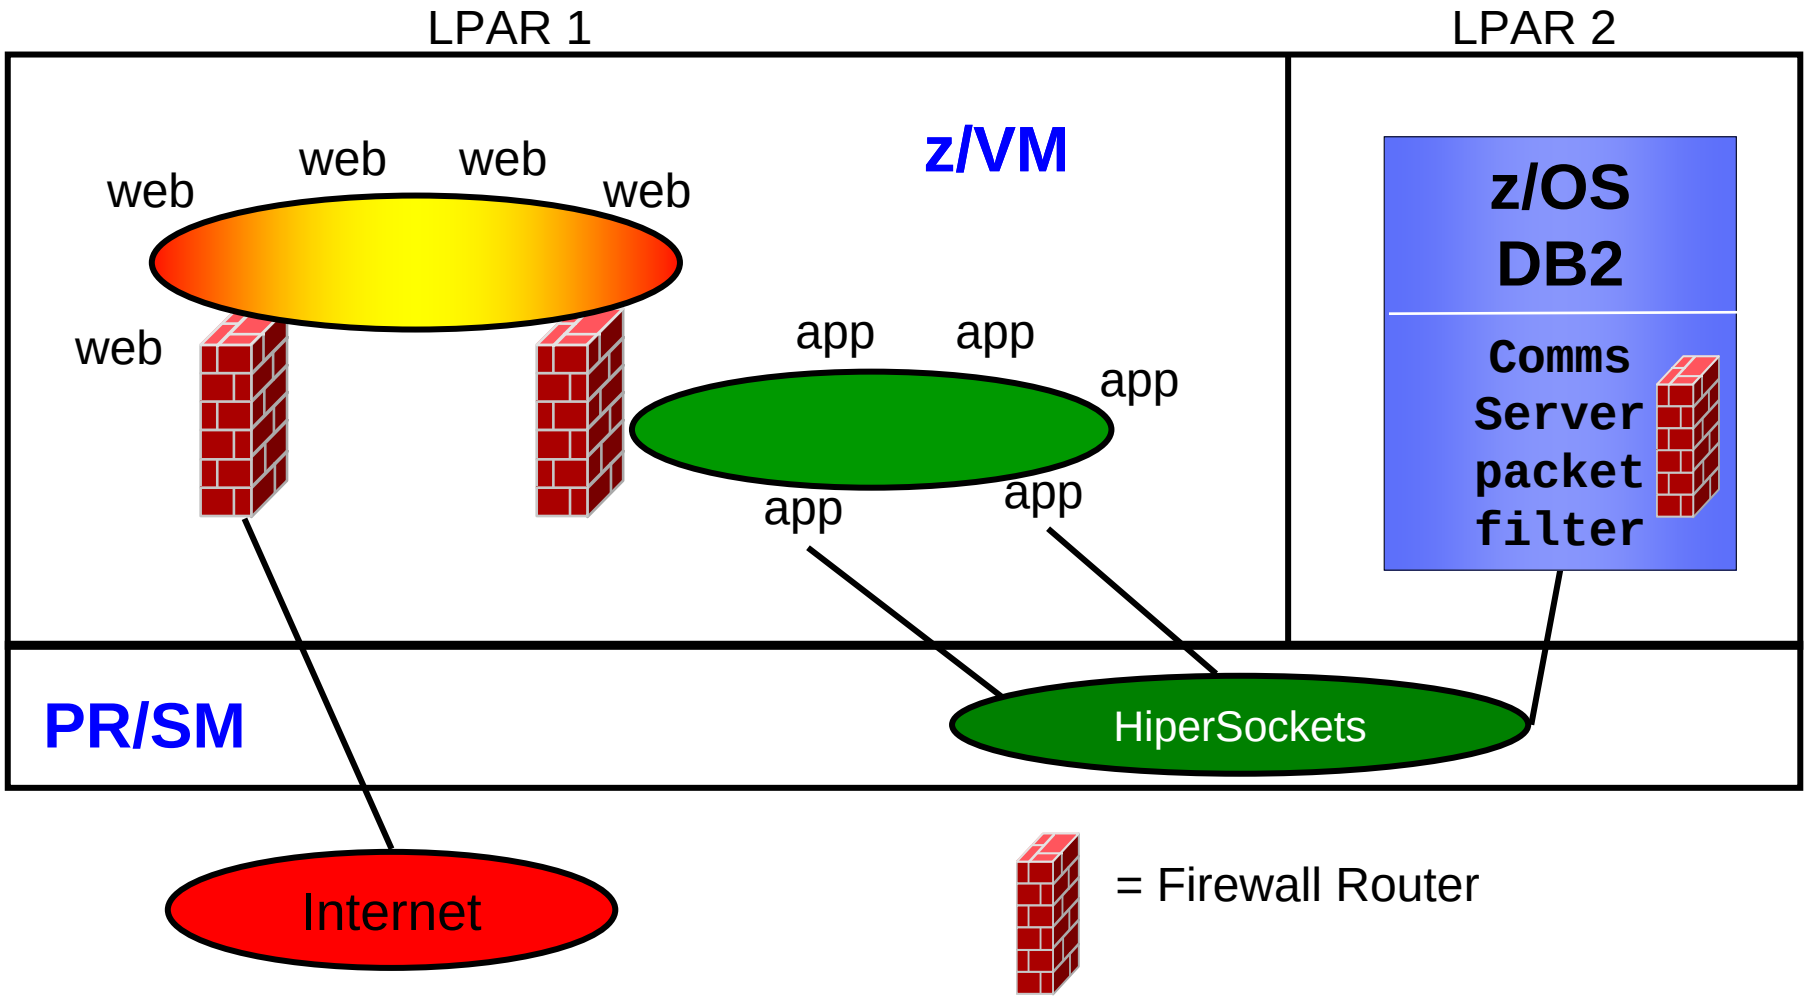

“Tempting, but no…”

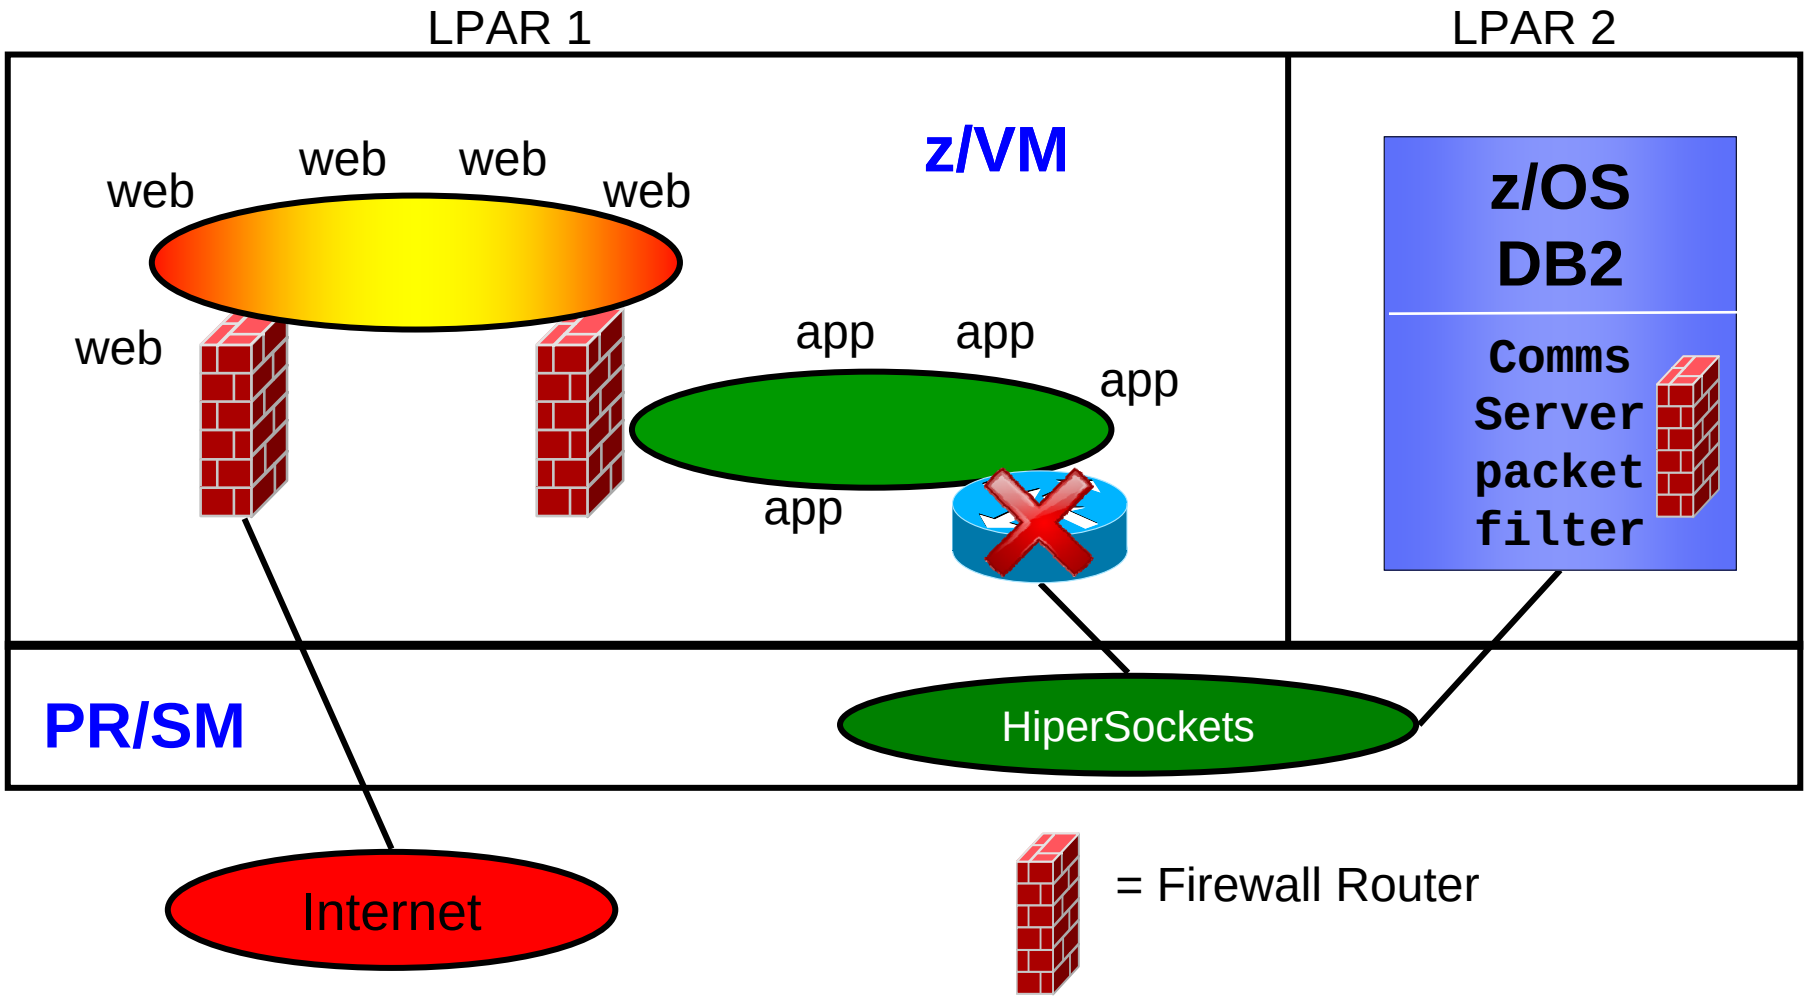

Virtual Switches
VLANs and traffic separation

What's a 'switch' anyway?

© Cisco Corp

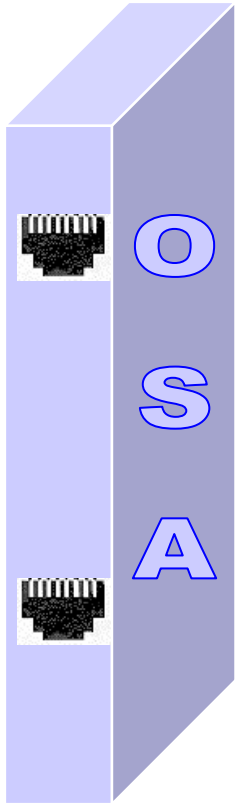

It creates LANs and routes traffic

- Turn ports on and off
- Assign a port to a LAN segment
- Provides LAN sniffer ports

z/VM Virtual Switch – VLAN unaware

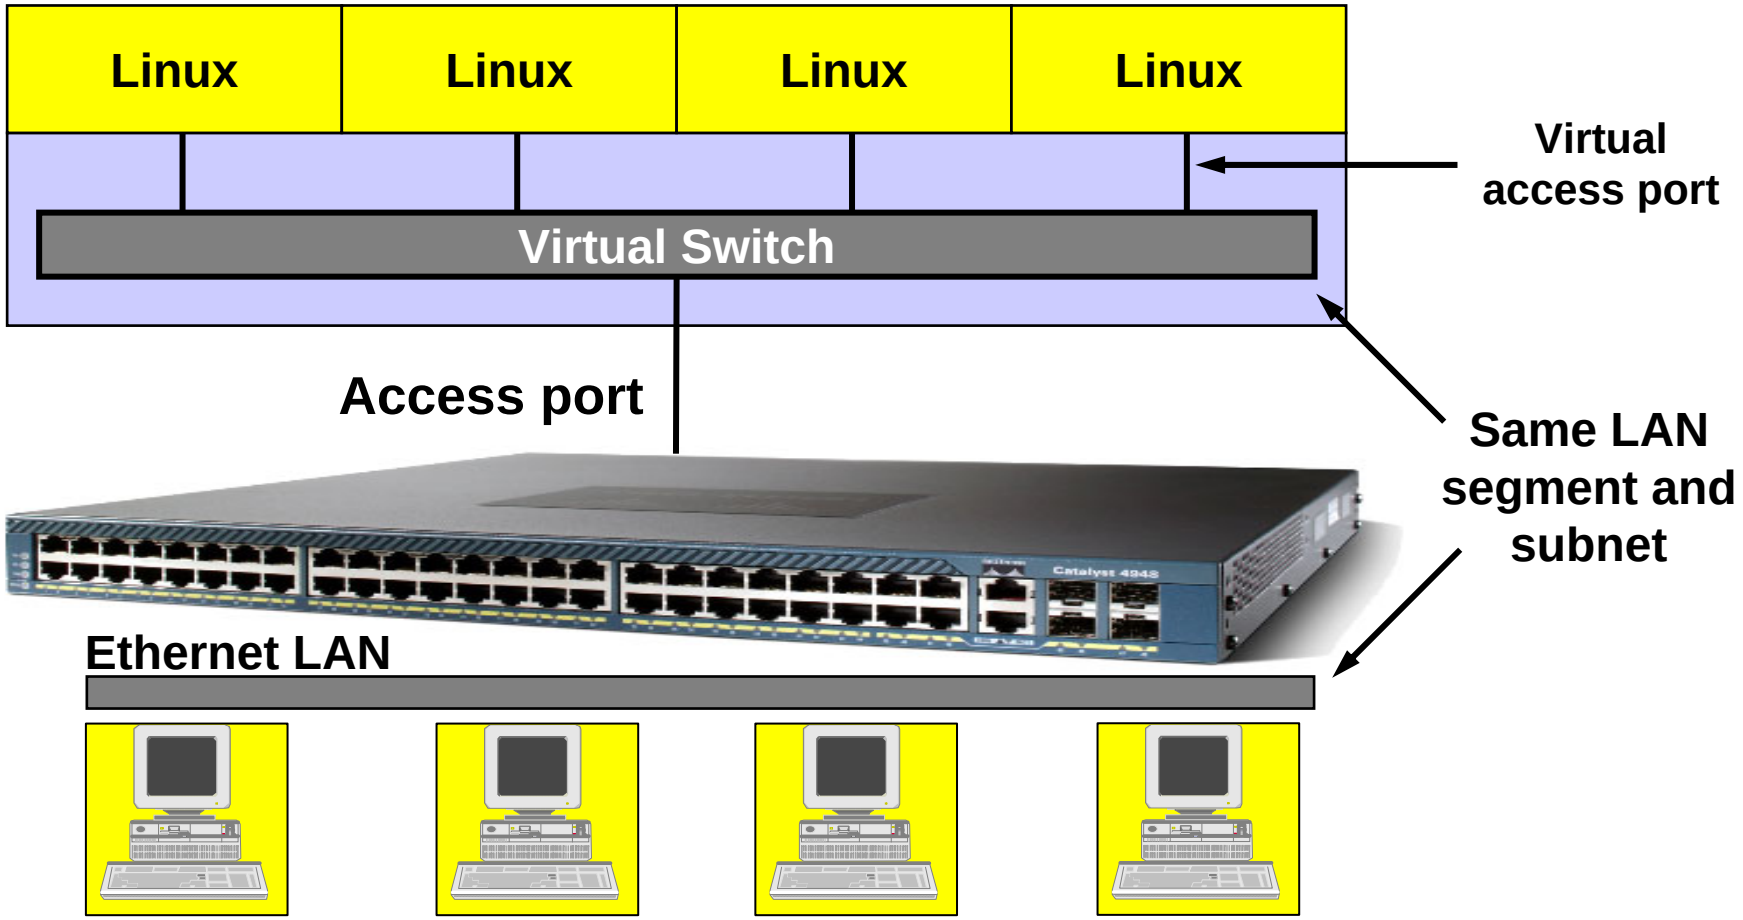

IEEE VLANs

© Cisco Corp

- If you run out of ports, you don't throw it away, you daisy chain ("trunk") it to another switch.

Trunk Port vs. Access Port

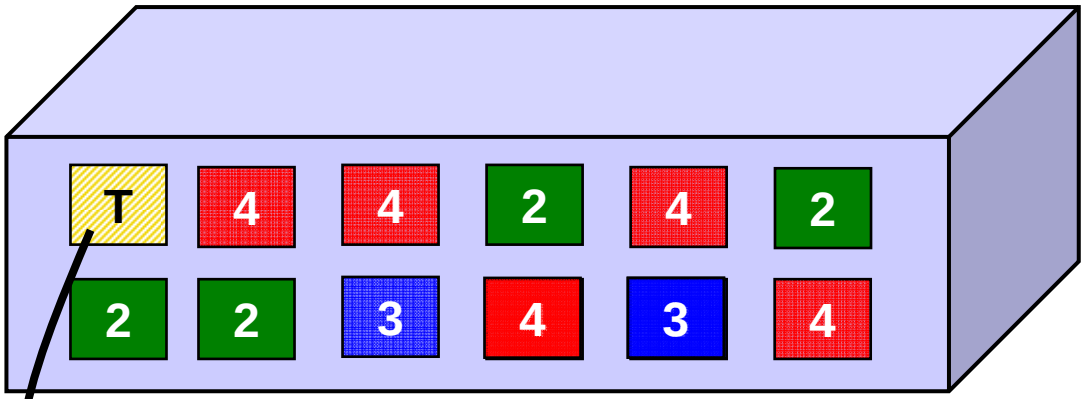

- Access port carries traffic for a single VLAN
- Host not aware of VLANs

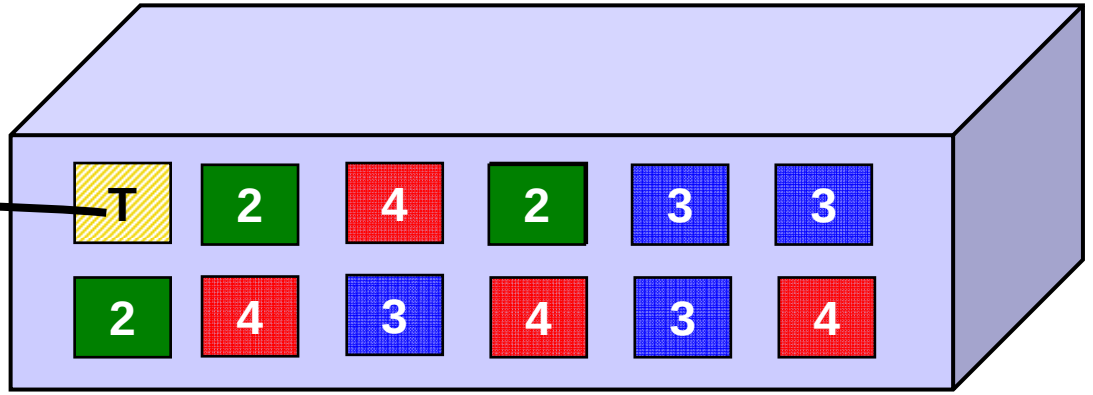

- Trunk port carries traffic from all VLANs
- Every frame is tagged with the VLAN id

Physical Switch to Virtual Switch

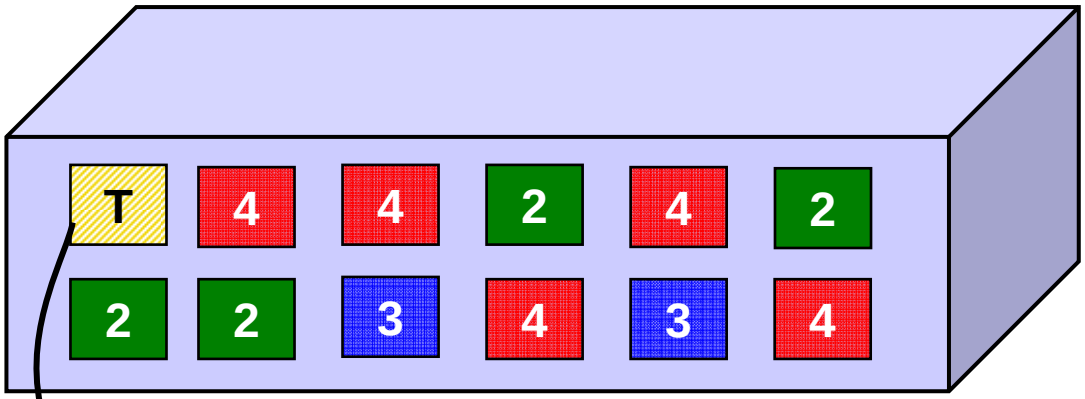

- **Trunk port carries traffic between CP and switch**
- **Each guest can be in a different VLAN**

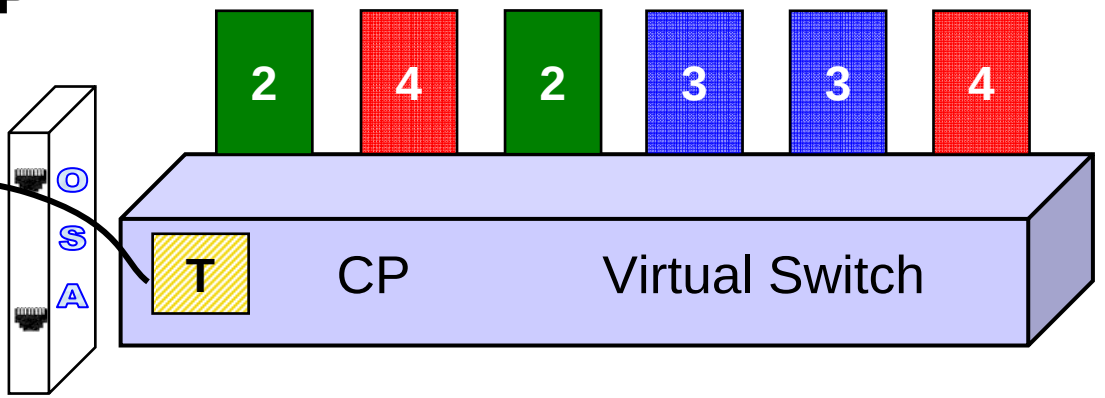

z/VM Virtual Switch – IEEE VLAN aware

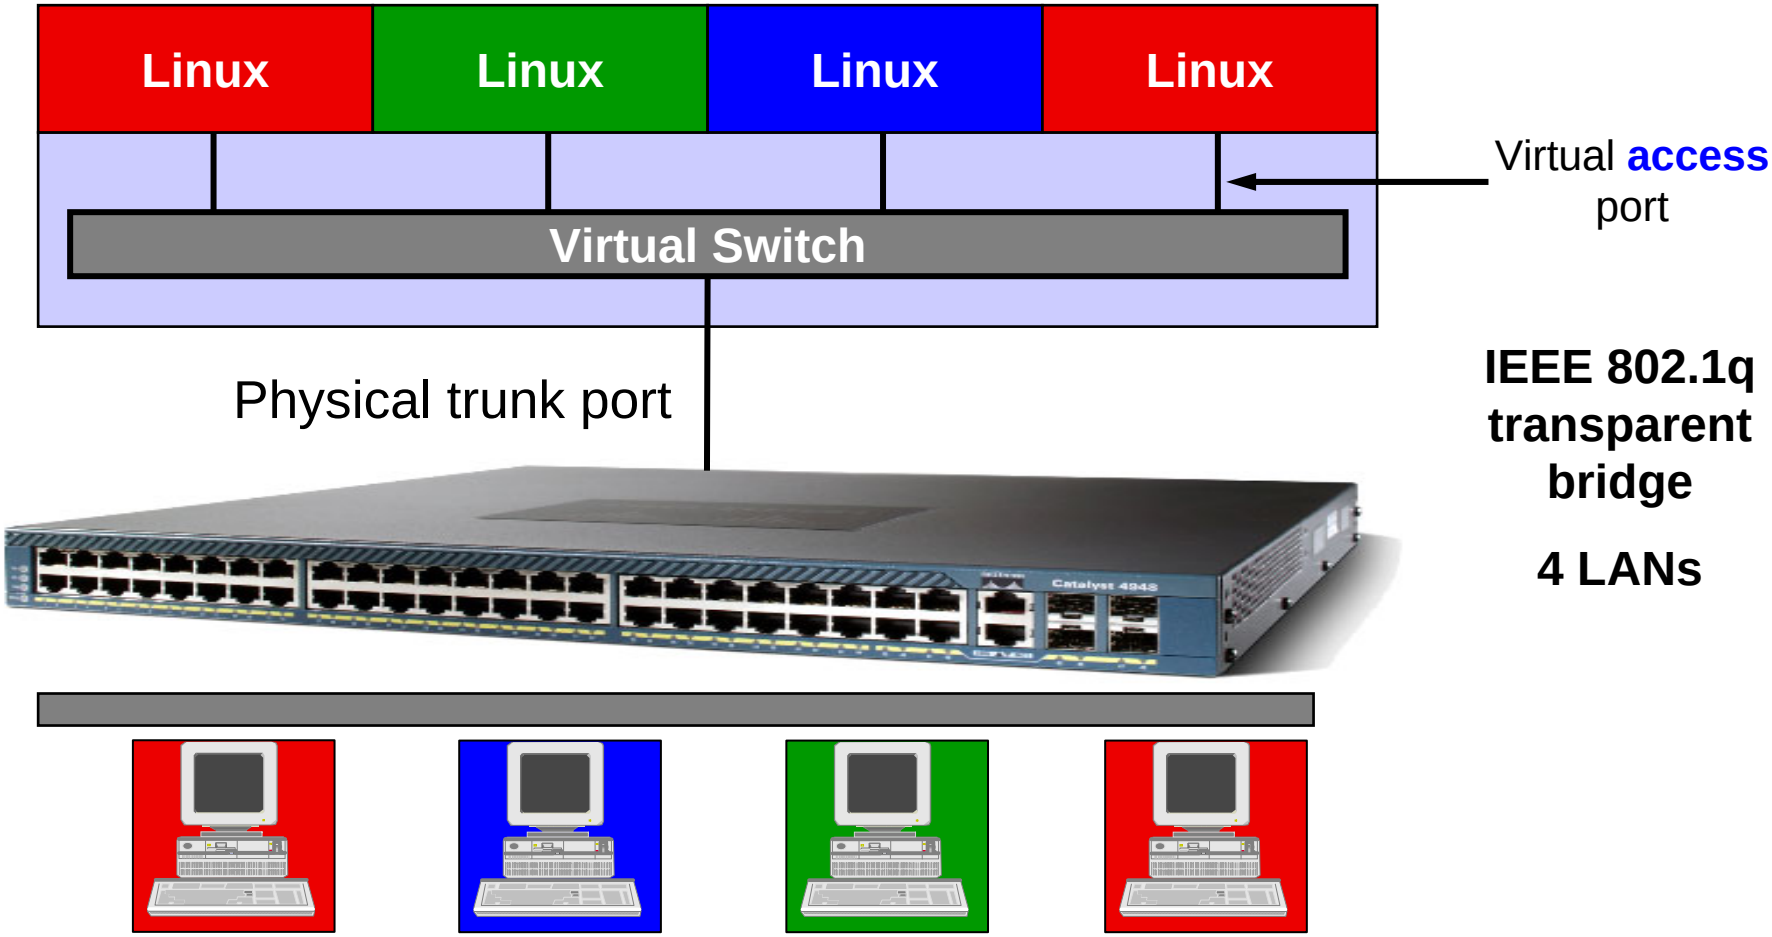

z/VM Virtual Switch – Virtual Trunk port

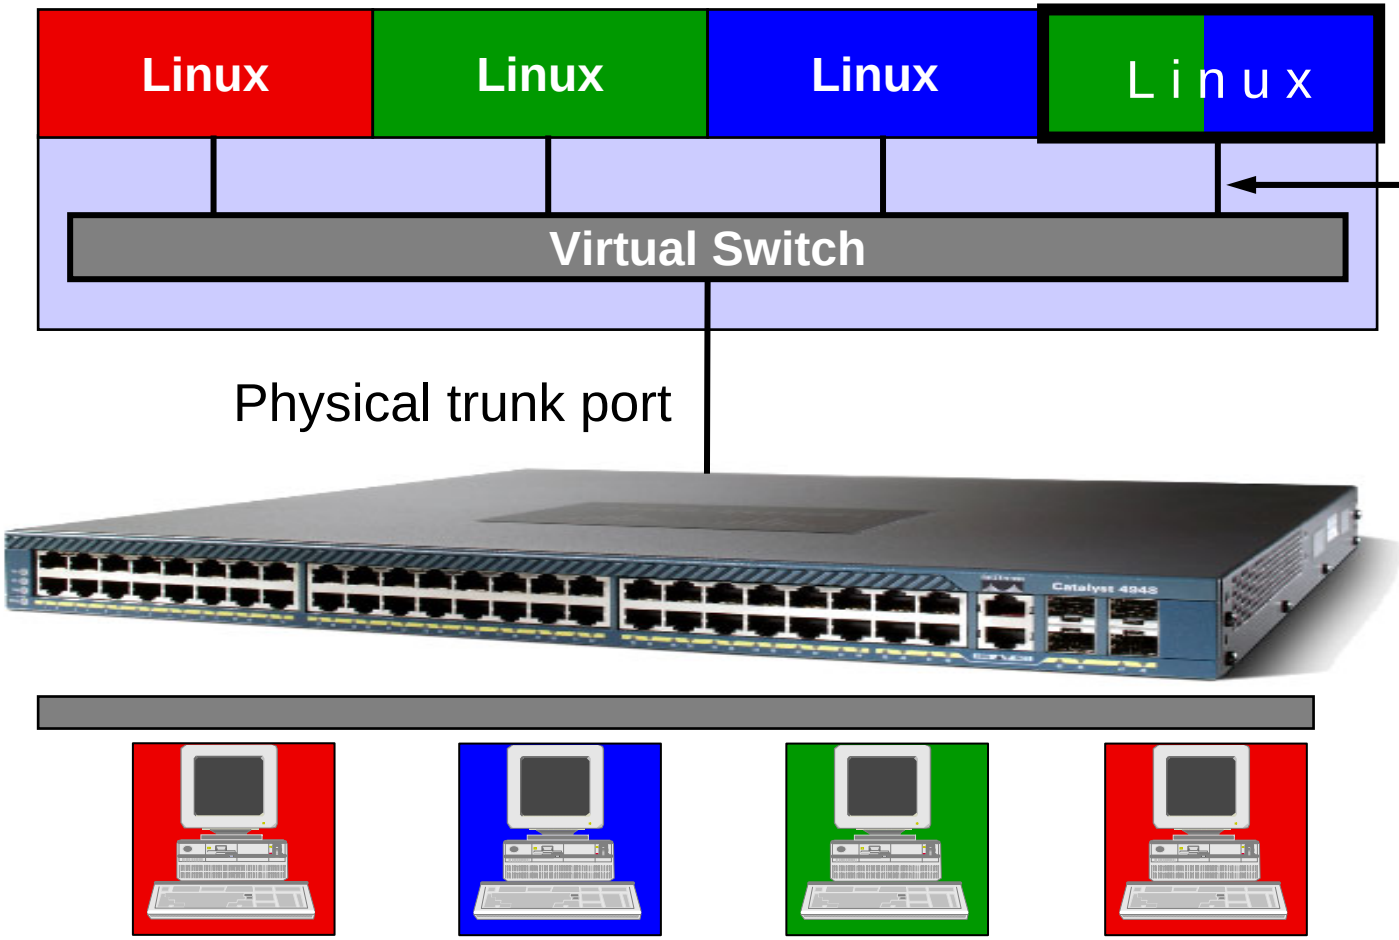

Virtual **trunk** port only for this guest

SET VSWITCH GRANT PORTTYPE TRUNK VLAN a b

Guest is authorized to two or more VLANs.

Must be VLAN aware

Network with VSWITCH (fully shared)

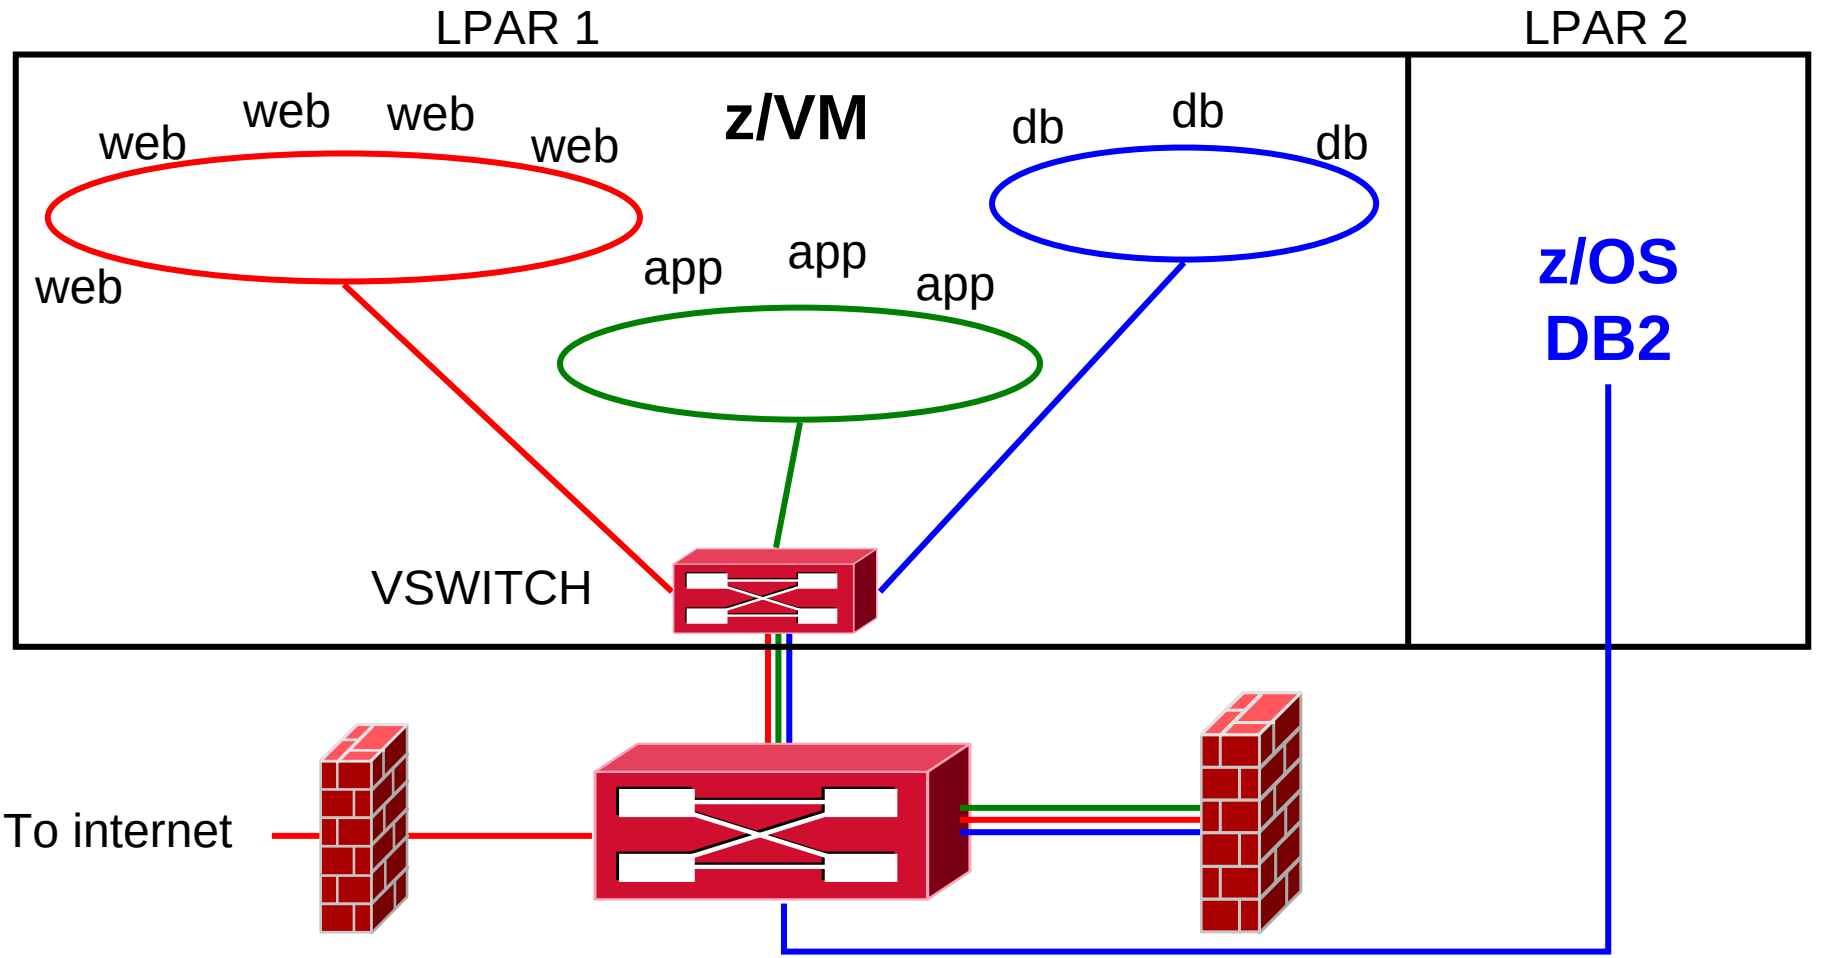

With 1 VSWITCH, 3 VLANs, and a multi-domain firewall

Multi-zone Network with VSWITCH (red zone physical isolation)

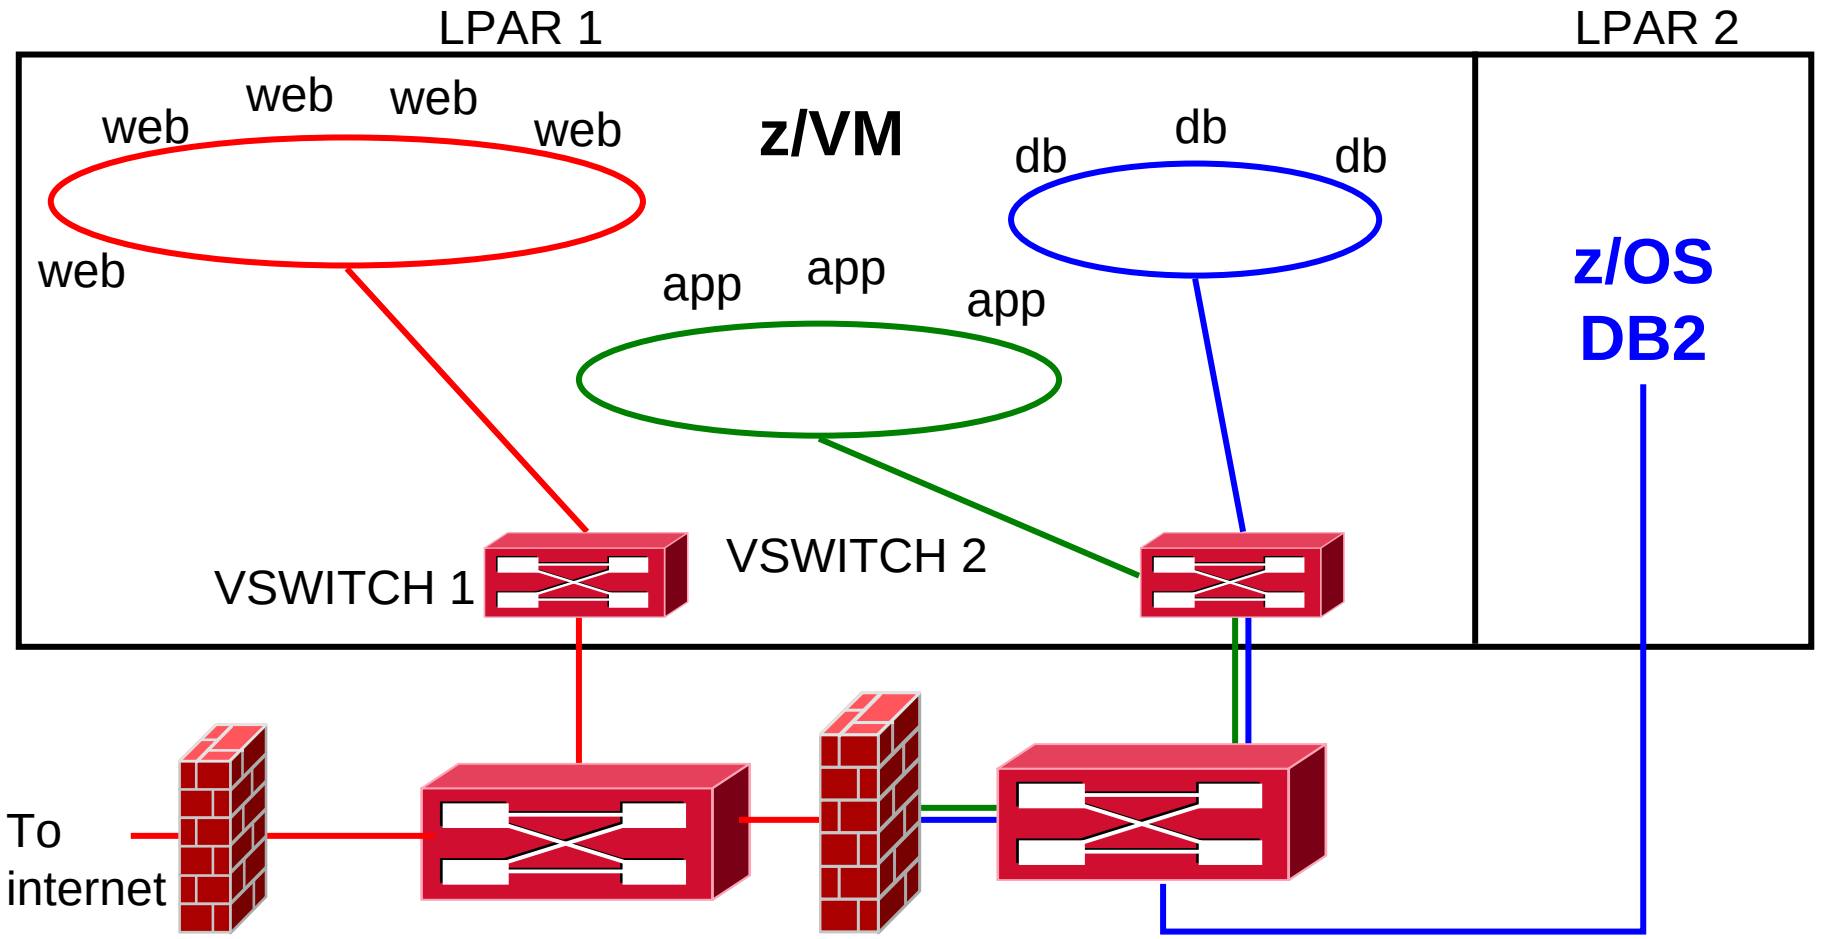

With 2 VSWITCHes, 3 VLANs, and a multi-domain firewall

z/VM Virtual Switch

- A special-purpose Guest LAN
  - Ethernet IPv4
  - Built-in IEEE 802.1q bridge to outside network
  - IEEE VLAN capable
- Each Virtual Switch has up to 3 separate OSA-Express connections associated with it
- Created in SYSTEM CONFIG or by CP DEFINE VSWITCH command

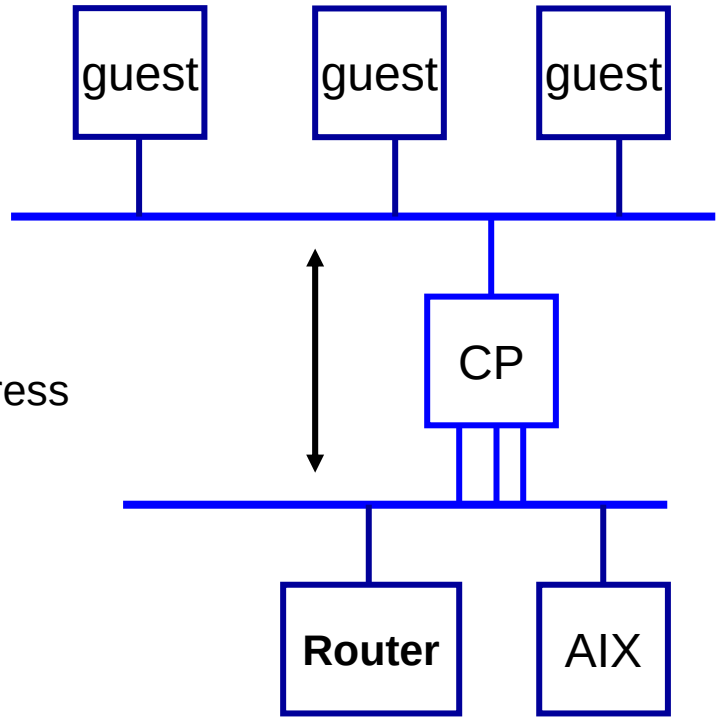

Virtual Switch Attributes

- Name
- Associated OSAs (maximum 3)
- A controlling virtual machine (minimal VM TCP/IP stack server)
  - Controller not involved in data transfer
  - Do not ATTACH or DEDICATE
  - DTCVSW1 and DTCVSW2
- Similar to Guest LAN
  - Owner SYSTEM
  - Type QDIO
  - Persistent
  - Restricted

Create a Virtual Switch

- SYSTEM CONFIG or CP command:

```
DEFINE VSWITCH name
              [RDEV NONE | cuu [cuu [cuu]] ]
              [CONNECT | DISCONNECT]
              [CONTROLLER * | userid]
              [NONROUTER | PRIROUTER]

              [VLAN UNAWARE | VLAN default_vid]
              [NATIVE native_vid]
              [GROUP group_name]

              [PORTTYPE ACCESS | PORTTYPE TRUNK]

Example:

DEFINE VSWITCH SWITCH12 RDEV 1E00 1F04 CONNECT
```

Change the Virtual Switch access list

- Specify after DEFINE VSWITCH statement in SYSTEM CONFIG to add users to access list

```
MODIFY VSWITCH name GRANT  userid
SET                        [VLAN vid1 vid2 vid3 vid4]
                           [PORTTYPE ACCESS | TRUNK]
                           [PROmiscuous | NOPROmiscuous]

SET    VSWITCH name REVOKE userid

Examples:
MODIFY VSWITCH SWITCH12 GRANT LNX01 VLAN 3 7 105
CP SET VSWITCH SWITCH12 GRANT LNX02 PORTTYPE TRUNK
                                    VLAN 4 20-22 29 302


CP SET VSWITCH SWITCH12 GRANT LNX02 PROMISCUOUS
```

IEEE 802.3ad Link Aggregation

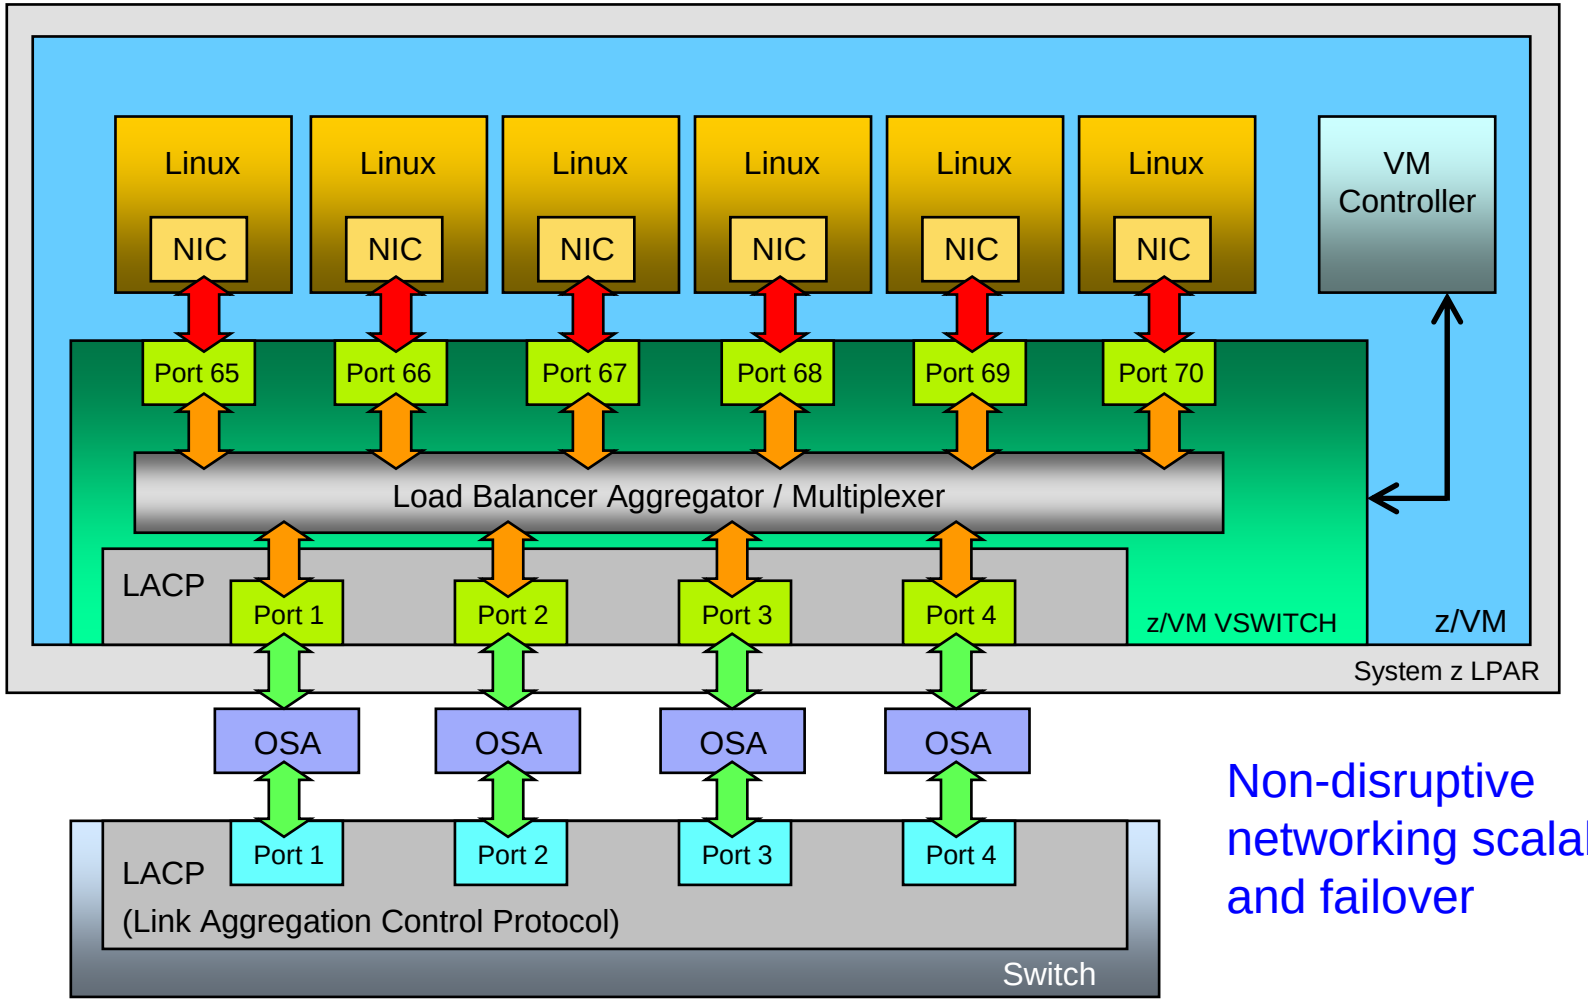

Non-disruptive networking scalability and failover

IEEE 802.3ad Link Aggregation

- System z9 and z10
- Groups available OSA-Express2 ports for use by the z/VM Virtual Switch
  - Up to 8 ports per virtual switch
  - Increases Virtual Switch bandwidth and provides near seamless failover in the event of a failed controller, link or switch
  - Only supported for Layer 2 switches
- Includes support to recover from a failed external switch

IEEE 802.3ad Link Aggregation

- Define an OSA port group with SET PORT GROUP
- DEFINE VSWITCH … GROUP
- No sharing of OSA!

z/VM Virtual Switch SNMP MIB

- **Helps enhance virtual network management with additional support for Simple Network Management Protocol (SNMP)**
- **Provides an SNMP subagent that will return Bridge MIB (Management Information Base) data for the z/VM Virtual Switch**
  - MIB data returned is defined by RFC 1493
  - The subagent acquires the information using a Control Program Diagnose interface (Diagnose x‘26C’)
- **Helps network administrators:**
  - Manage virtual network performance
  - Find and solve virtual network problems
  - Plan virtual network growth
- **Support also provides a programming interface to obtain information about virtual networks**

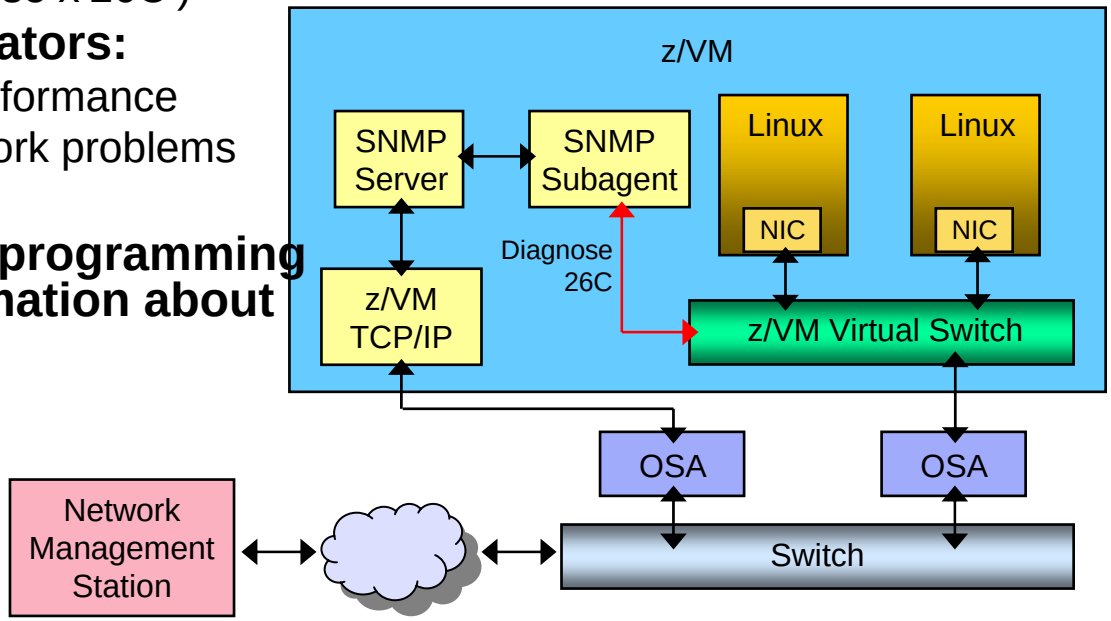

z/VM Virtual Switch – VLAN aware – No OSA

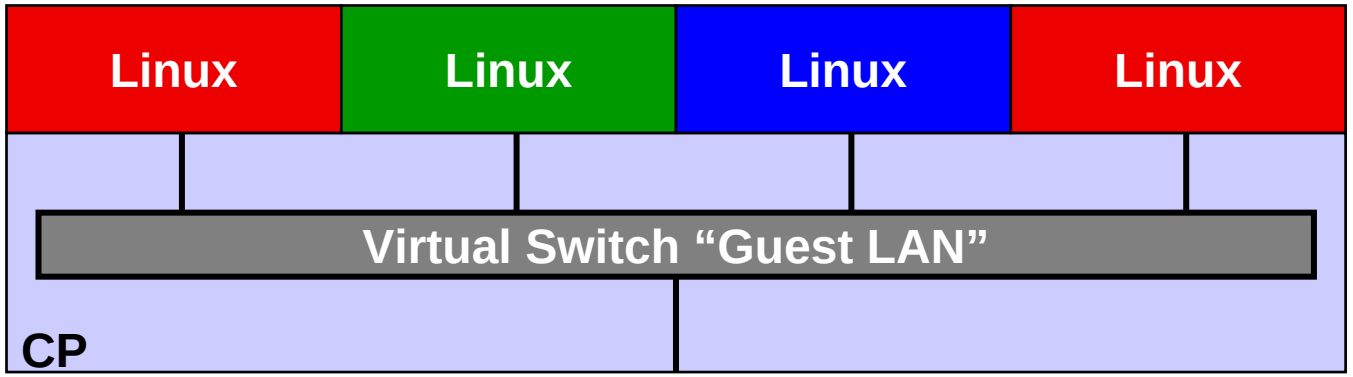

Use this instead of a Guest LAN!

(More controls)

Security controls

- ESM control for all guest LANs and VSWITCHes, including VLAN ID control
  - RACF: Class VMLAN, Profile owner.lanname or owner.lanname.vid
  - All Guest LANs and VSwitches can be controlled
- Support for LAN Sniffers
  - CP command or device driver control (“promiscuous mode”)
    - SET VSWITCH GRANT, SET LAN GRANT, SET NIC
  - External security manager
    - RACF/VM CONTROL access to VMLAN profile
  - Guest receives copies of all frames sent or received

Additional Features

- Pre-defined VSWITCH controllers
  - DTCVSW1 and DTCVSW2
  - Same as shown in Getting Started with Linux
    - Add them to AUTOLOG1
    - Remove “VSWITCH CONTROLLER ON” from PROFILE TCPIP in your production stacks

Enforcing the Rules with RACF

Virtual Switch

- Access controlled by VMLAN class in RACF
  - SYSTEM.*name* or SYSTEM.*name.vlanid*
  - *owner.name* (for Guest LANs)
- PERMIT SYSTEM.VSW01 CLASS(VMLAN) ID(ALAN) ACCESS(UPDATE)
  - Sniffer mode requires CONTROL access
- Port isolation
  - SET VSWITCH *name* ISOLATE
  - Guests cannot talk to each other
  - System z10 OSA QDIO data connection isolation: No cross-talk on shared OSA to/from the VSWITCH

Turn off backchannel communications

- No user-defined Guest LANs
  - VMLAN LIMIT TRANSIENT 0
- No virtual CTC
  - MODIFY COMMAND DEFINE IBMCLASS G PRIVCLASS M
- No IUCV
  - Use explicit IUCV authorization in the directory, not IUCV ALLOW or IUCV ANY
- No secondary consoles
  - MODIFY COMMAND SET SECUSER IBMCLASS G PRIV M

- But what else might there be?

Turn off backchannel communication

- VMCF
  - MODIFY DIAGNOSE DIAG068 IBMCLASS G PRIVCLASS M
- ESA/XC mode address space sharing (ADRSPACE PERMIT)
- DCSS
- And we can add new interfaces in an APAR

- Google “less than class g” by Rob van der Heij

- Too hard for some folks

- Consider RACF Mandatory Access Controls instead

- AppArmor and SELinux provide the same capabilities for Linux

Multi-Zoning with RACF

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

Multi-Zoning with RACF

- A **Security Label** combines the concepts of
  - Security clearance (secret, top secret, eyes only)
  - Information zones
- Information zones apply to any place data may exist
  - disks, networks, and other users
- Security clearance
  - Ensures servers cannot see extra-sensitive data in their information zone
  - Prevents copying of data to medium that is readable by servers with lower security clearance (“No write down”)
  - Not prevalent since there is no equivalent in distributed networking solutions
- Label “dominance” is established based on intersection of zones and security clearance
  - Not just a simple string comparison

Multi-zone z/VM LPAR with RACF Security Label Enforcement

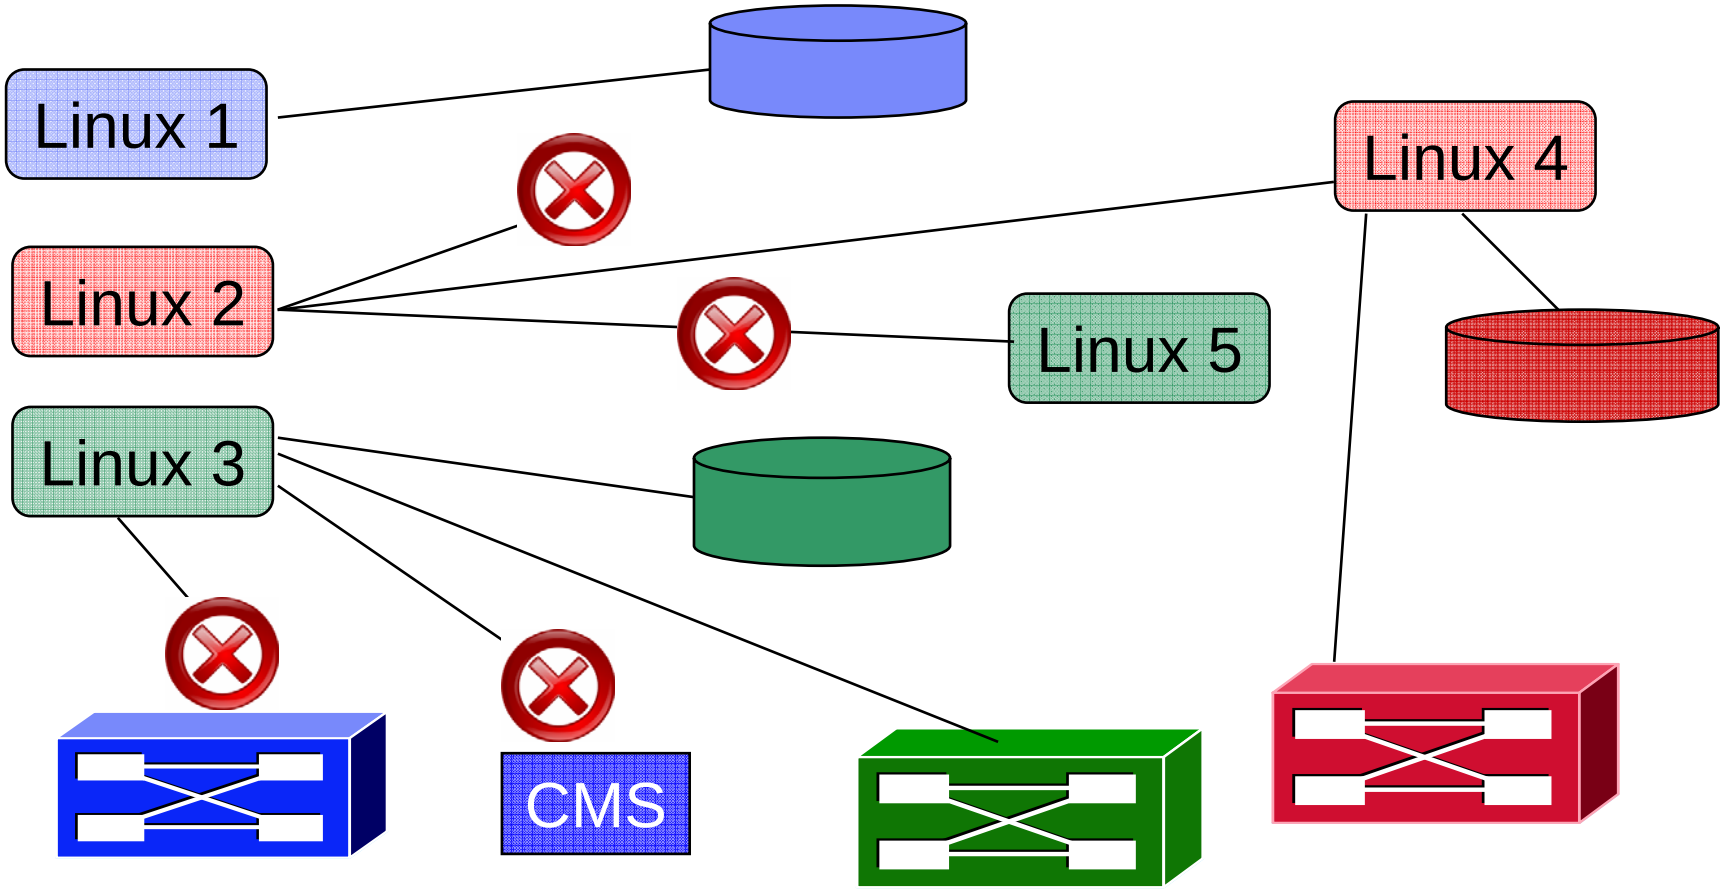

Multi-Zoning with RACF

Create security levels and data partitions

```
RDEFINE SECDATA SECLEVEL ADDMEM(DEFAULT/100)

RDEFINE SECDATA CATEGORY ADDMEM(INTERNET DMZ APPS DATA COMMON)

RDEFINE SECLABEL PUBLIC SECLEVEL(DEFAULT)ADDCATEGORY(COMMON)
    UACC(NONE)

RDEFINE SECLABEL RED SECLEVEL(DEFAULT)ADDCATEGORY(DMZ COMMON)
    UACC(NONE)

RDEFINE SECLABEL GREEN SECLEVEL(DEFAULT) ADDCATEGORY(APPS COMMON)
    UACC(NONE)

RDEFINE SECLABEL BLUE SECLEVEL(DEFAULT) ADDCATEGORY(DATA COMMON)
    UACC(NONE)
```

Multi-Zoning with RACF

Assign virtual machines their SECLABELs

```
PERMIT RED CLASS(SECLABEL) ID(LXHTTP01) ACCESS(READ)
ALTUSER LXHTTP01 SECLABEL(RED)

PERMIT GREEN CLASS(SECLABEL) ID(LXWAS001) ACCESS(READ)
ALTUSER LXWAS001 SECLABEL(GREEN)
```

Multi-Zoning with RACF

- But sometimes a server serves the Greater Good, providing services to all users
- Exempt server from label checking
- Assign system servers label SYSNONE

```
PERMIT SYSNONE CLASS(SECLABEL) ID(TCPIP) ACCESS(READ)

ALTUSER TCPIP SECLABEL(SYSNONE)
```

Multi-Zoning with RACF

- Assign labels to resources
  - VMMDISK – Minidisk
  - VMLAN – Guest LANs and Virtual Switches

  - `RALTER VMMDISK LXHTTP01.201 SECLABEL(RED)`

  - `RALTER VMLAN SYSTEM.NET1 SECLABEL(RED)`

  - `RALTER VMLAN SYSTEM.NET2.0307 SECLABEL(GREEN)`
  - `RALTER VMLAN SYSTEM.NET2.0410 SECLABEL(BLUE)`

- If you intend to activate TERMINAL or VMSEGMT classes, those resources all need SECLABELs

Multi-Zoning with RACF

- Activate RACF protection
  - SETROPTS CLASSACT(SECLABEL VMMDISK VMLAN)
  - SETROPTS RACLIST(SECLABEL)
  - SETROPTS MLACTIVE(WARNINGS)
    - If resource doesn't have a seclabel, message is issued and seclabels are ignored.

  Or

  - SETROPTS MLACTIVE(FAILURES)
    If resource doesn't have a seclabel, command fails.
    This is more secure!

Summary

- Check network design with network architect

- Don’t whine about firewalls
- Optimize with host-resident firewalls later

- Protect the hardware
- Protect your data
- Protect your servers
- Protect your company

- Protect yourself!!

Appendix

Reference Information

- Security zones on z/VM presentation
  - http://www.VM.ibm.com/devpages/altmarka/present.html
- z/VM Security resources
  - http://www.VM.ibm.com/security
- z/VM Secure Configuration Guide
  - http://publibz.boulder.ibm.com/epubs/pdf/hcss0b30.pdf
- System z Security
  - http://www.ibm.com/systems/z/advantages/security/
- z/VM Home Page
  - http://www.VM.ibm.com

Reference Information . . .

- Publications:
  - z/VM CP Planning and Administration
  - z/VM CP Command and Utility Reference
  - z/VM TCP/IP Planning and Customization
  - z/VM Connectivity Planning, Administration and Operation

- Links:
  - http://www.ibm.com/servers/eserver/zseries/os/linux/
  - http://www.linuxvm.org/

Built-in Diagnostics

- **CP QUERY VMLAN**
  - to get global VM LAN information (e.g. limits)
  - to find out what service has been applied
- **CP QUERY LAN ACTIVE**
  - to find out which users are coupled
  - to find out which IP addresses are active
- **CP QUERY NIC DETAILS**
  - to find out if your adapter is coupled
  - to find out if your adapter is initialized
  - to find out if your IP addresses have been registered
  - to find out how many bytes/packets sent/received