

October 2009 IBM

Security Zones on z/VM

Alan Altmark, IBM z/VM Security Strategist
Alan_Altmark@us.ibm.com



© 2008, 2009 IBM Corporation

This presentation is the direct result of my interactions with large IBM clients since January 2008. There are things that System z people need to understand about the network ecosystem in which they live.

In multi-tier network applications, there is usually the concept of "zones" that each contain a set of servers. Each zone is separated from others by firewalls and by access policies to ensure that there is no unwanted user access or flow of data. Learn in this presentation how to properly build virtual zones and to integrate virtual servers into your existing zones. We will also discuss using the RACF Security Server on z/VM to prevent a "red zone" server from connecting to a "green zone" network or "green zone" data.



Trademarks

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

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

* 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.



Agenda

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



System z is not delivered with a bag of pixie dust that makes it secure.




It is the detail- and process-oriented nature of mainframers that makes it secure. You have to stop and examine every part of your System z management processes and procedures to KNOW that the system is really secure. Don't get caught up in the hype. Just because Hollywood holds up the mainframe as The Holy Grail, remember that it's just a machine managed by people, neither of which is perfect.

We need to perform due-diligence to protect our companies, our clients, our families, and ourselves.

Hard work, teamwork, asking for advice, wisdom, self audit. And, yes, a bit of magic, perhaps.



Securing the Hardware



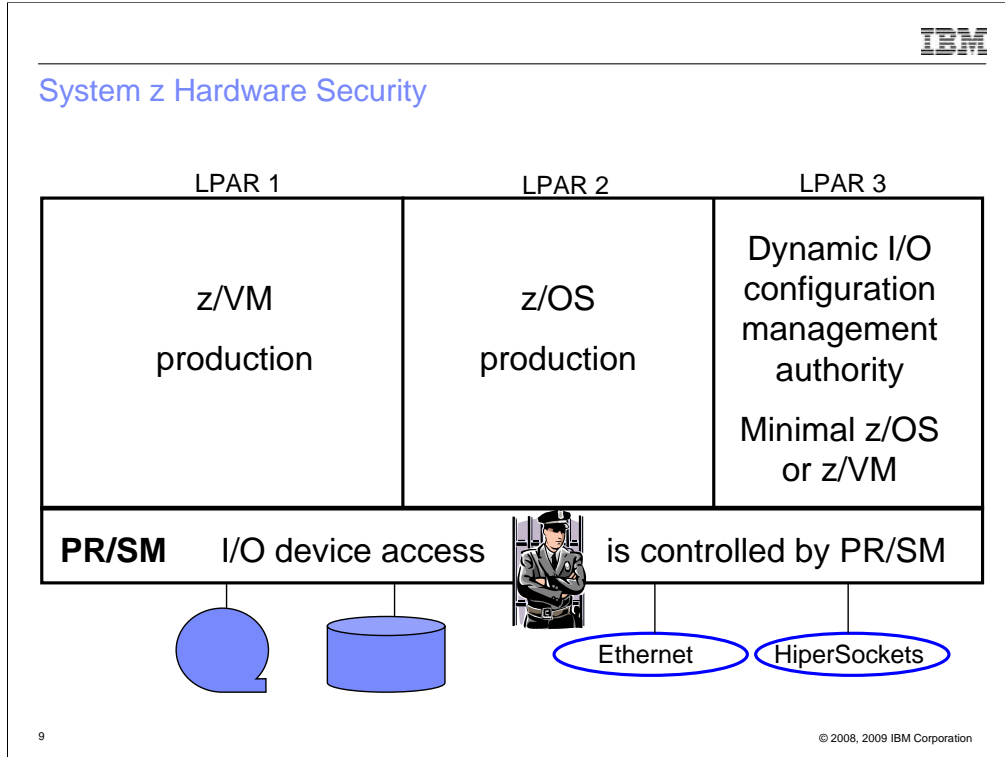
z/VM Security begins with System z security

- 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

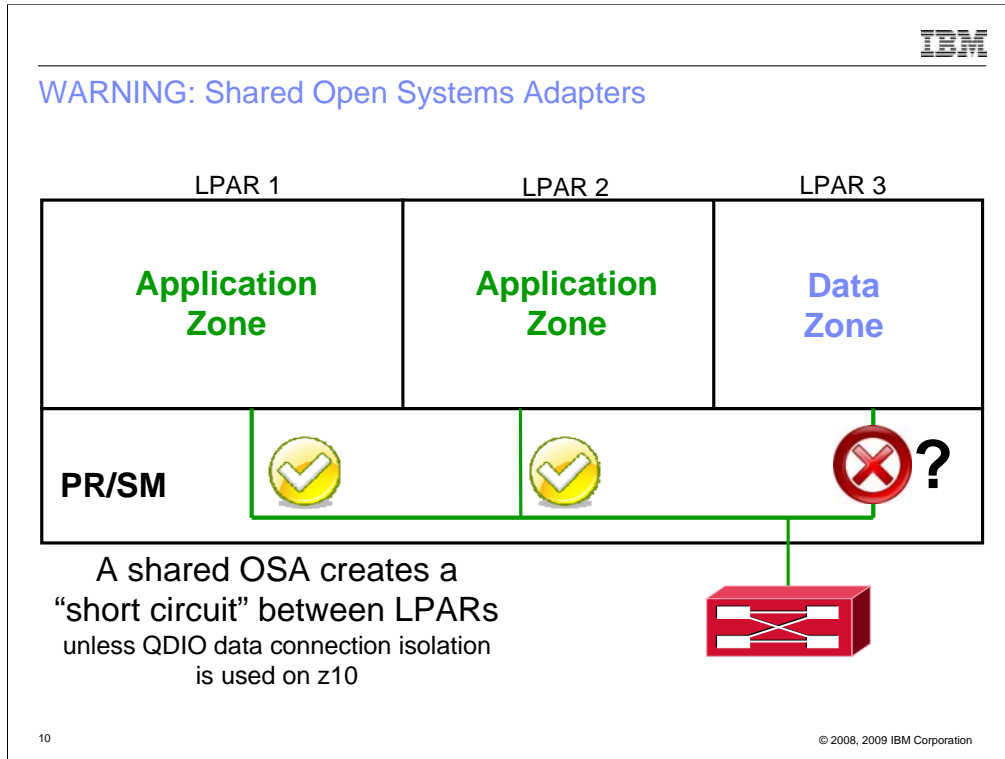
8

© 2008, 2009 IBM Corporation

Back to basics. Ultimate Power in the hands of a z/VM (or z/OS) system programmer or administrator is NOT a given. People who do not have the authority to make hardware changes should not be given the privileges to do so. Authorization to issue Dynamic I/O commands in z/VM constitutes a hardware change capability (assuming the LPAR has the privilege). Be careful.

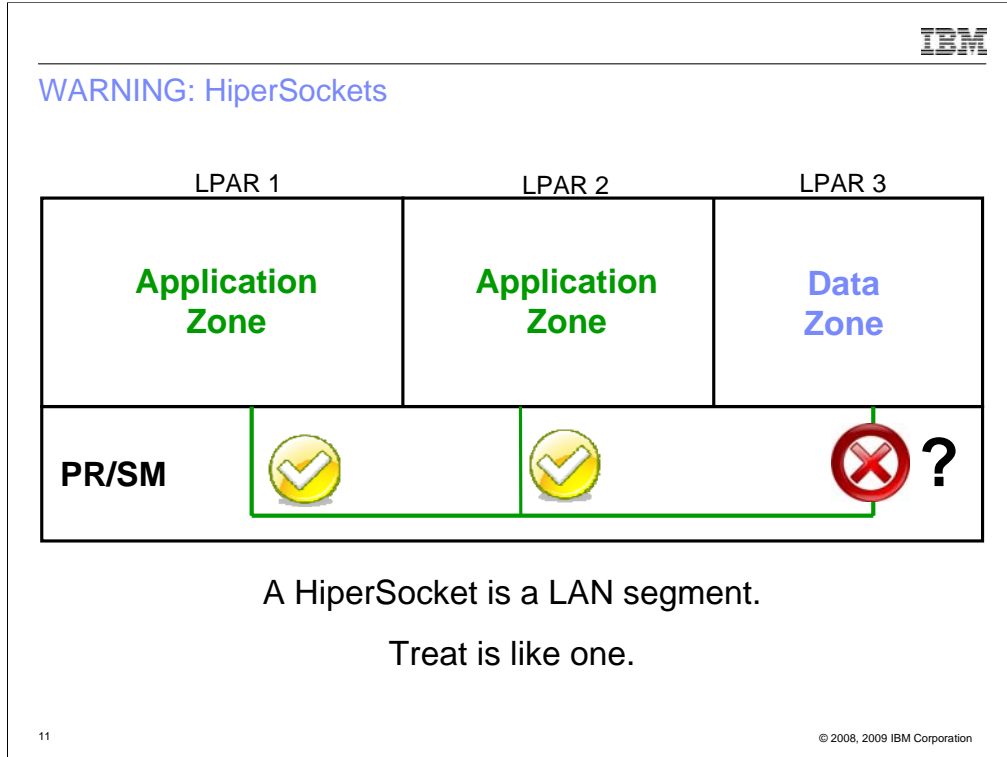


Note that the dynamic I/O partition running HCD isn't running in a production or test partition. What if the production LPARs were hacked? Would you want either of them to have control of the I/O config?



Sharing an OSA and creating a HiperSocket both create a LAN segment. Be careful about such things if it is necessary to transit a zone. You **MUST** have some sort of firewall technology.

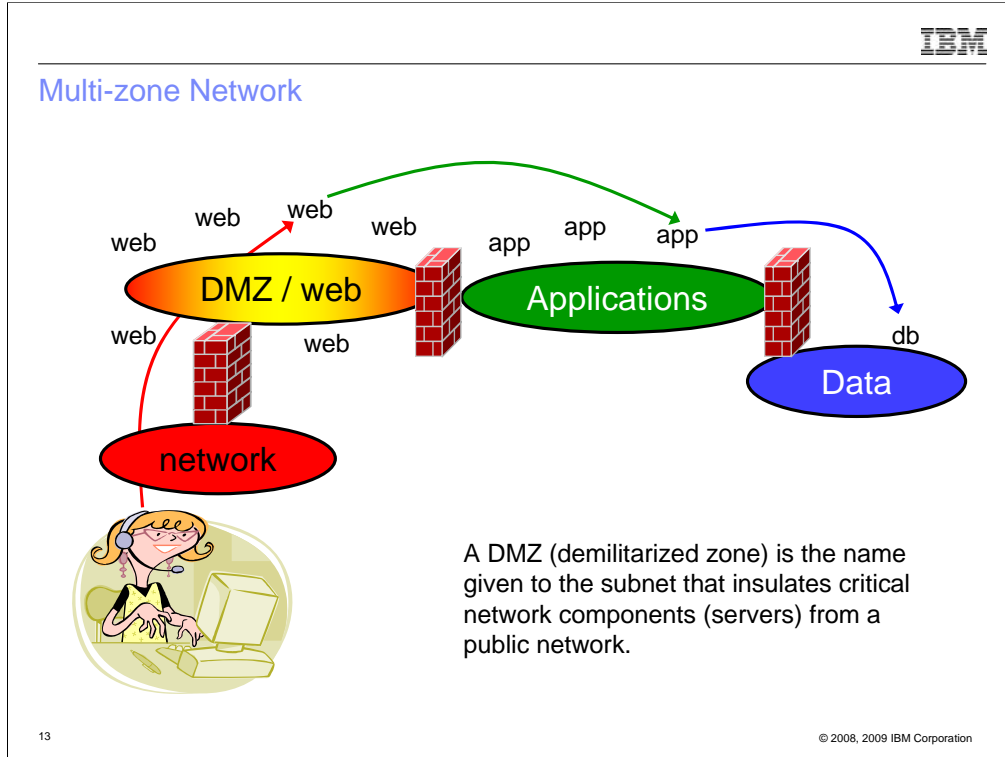
VSWITCH port isolation and QDIO data connection isolation can be used to mitigate the risk of sharing.



The only difference between a HiperSocket and a shared OSA is the lack of a built-in bridge to the Outside.



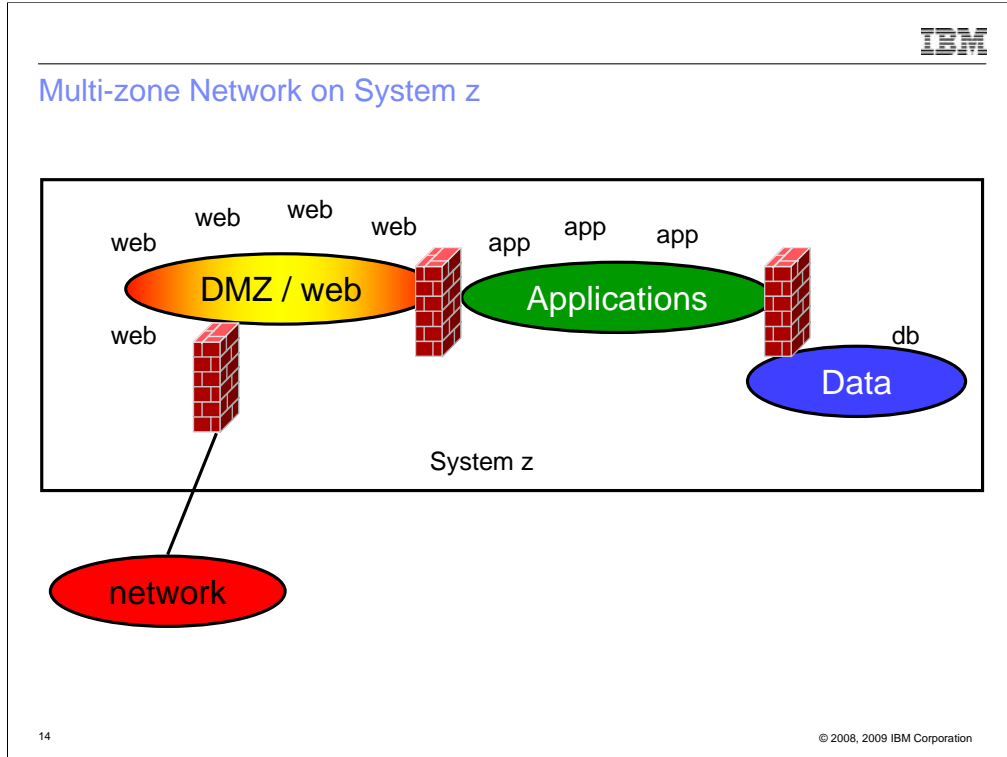
Multi-zone networks



This is a picture of a traditional 3-tier application architecture. There are 4 security zones. The definition of a zone is that it is separated by a firewall. If you get rid of a firewall, you get rid of the “higher” of the two zones.

That is, get rid of the 3rd firewall and you will lose the right to host servers in the “data” zone.

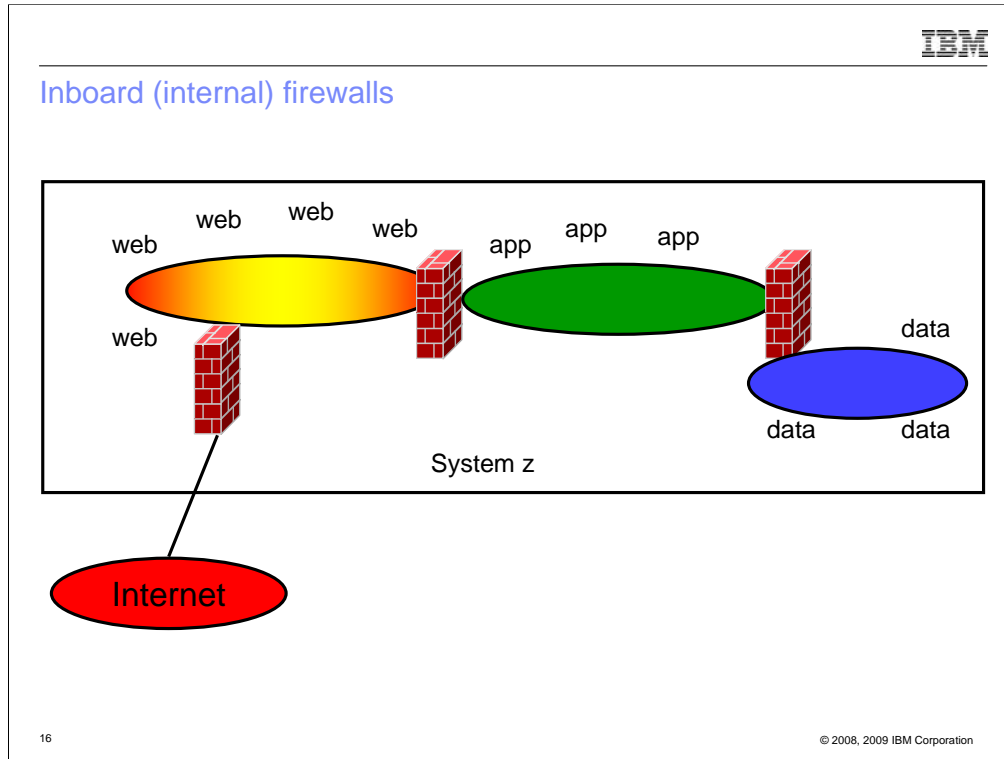
It might not make sense, but it’s true nonetheless. That’s a Best Practice for network security. (And likely required by PCI.)



This is what the System z salesman will sell. Utopia? We'll see in a few minutes.

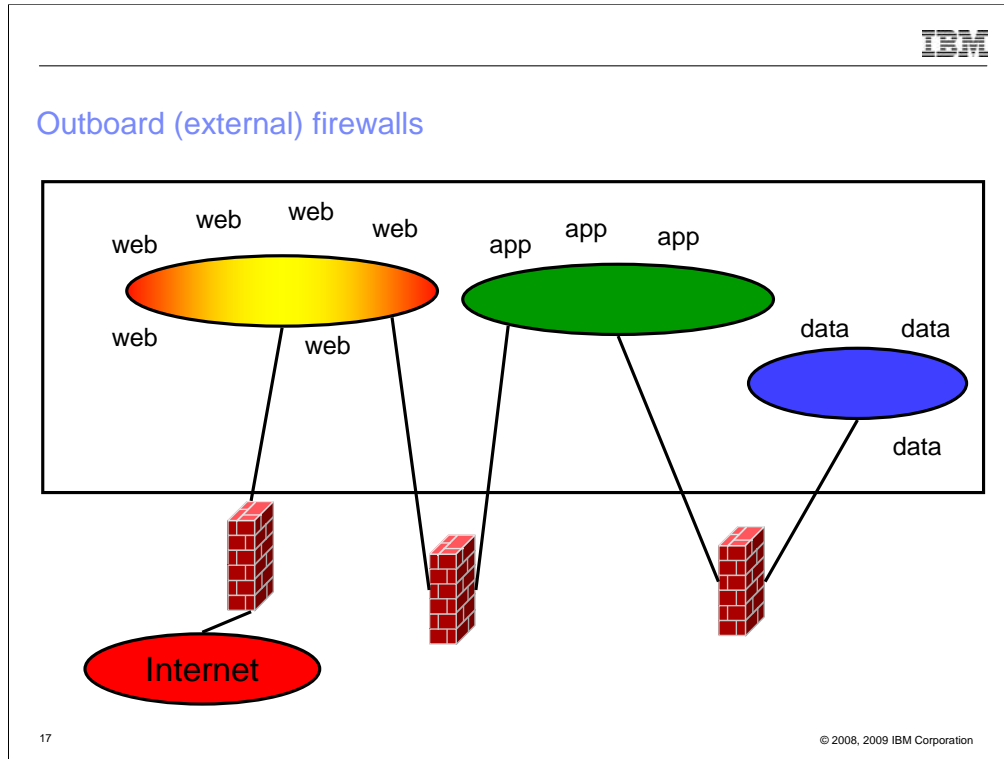


Firewalls are an important element of any network. They protect servers **and** data. Their use is mandatory in PCI-conformant designs.

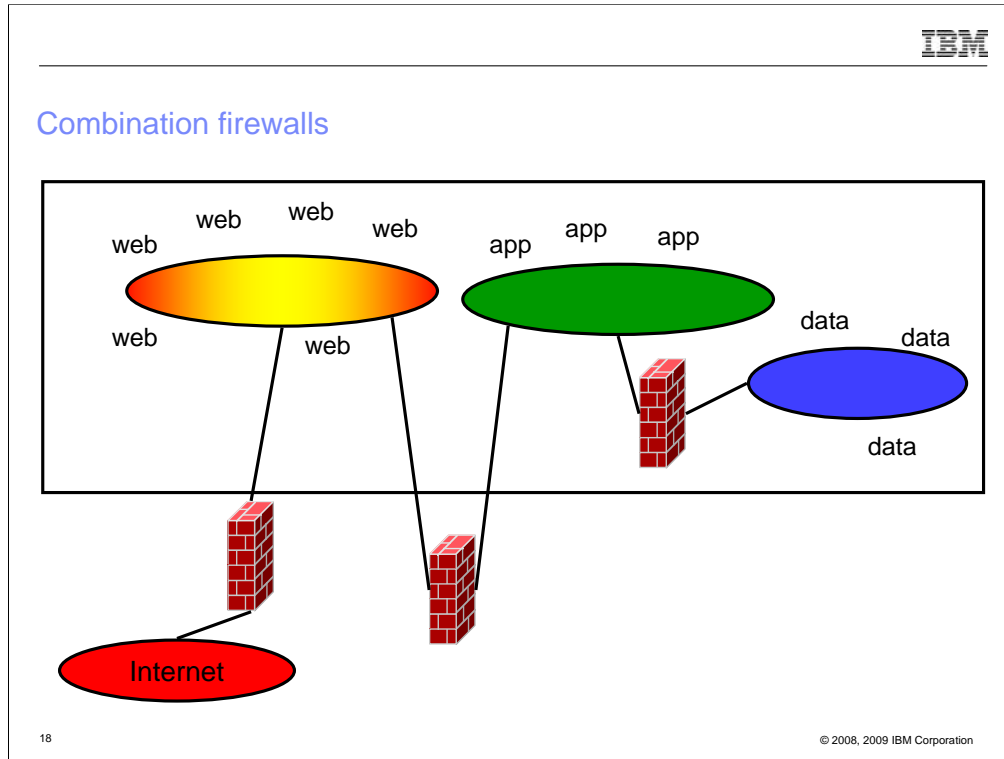


This is what the Salesman will sell, but it is NOT a given that the firewalls will be running in virtual machines. Remember that network security is not the responsibility of a z/VM systems programmer – it is the responsibility of your Network Security teams. **They** decide what firewalls are acceptable. That decision is typically based on how they manage firewalls. There is usually a piece of firewall management software that can push rules to all firewalls quickly and easily. Sometimes the firewall has built-in hot-standby capabilities. Unless you're part of the Network Security team, you won't be aware of all the issues. **DO NOT SURPRISE THEM WITH A FIREWALL TECHNOLOGY OF YOUR CHOOSING!!**

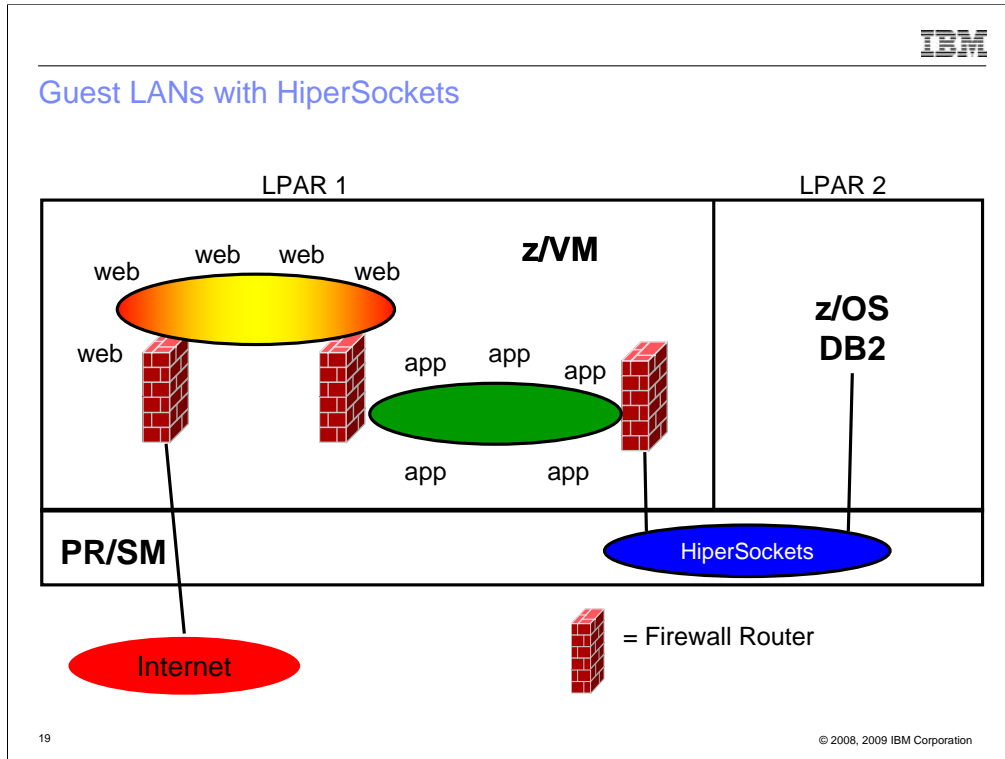
Keep an eye out for IBM Proventia Security Server for Linux on System z.



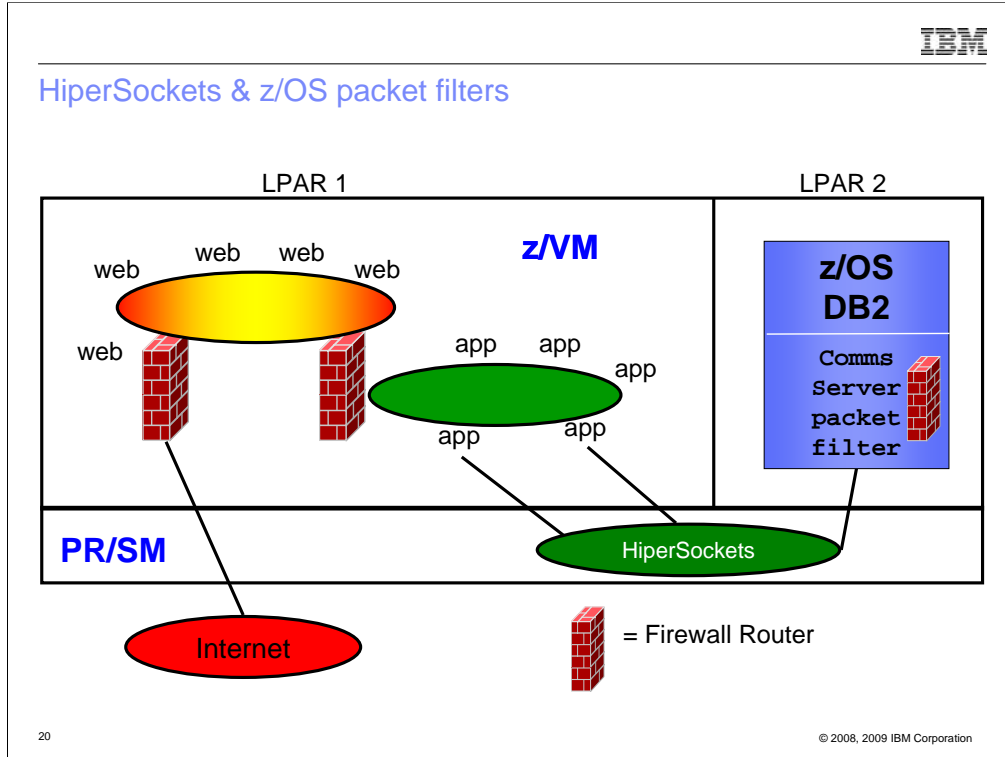
This is the easiest to build since you don't need to introduce 'alien' firewall technologies. If your Network Security team wants outboard firewall's, No Problem. We don't want them running on the mainframe anyway (esp. iptables) since they chew a lot of CPU. Yes, moving them outboard will increase latency. But as long as the transactions are "fast enough", who cares? Of course, you actually have to have a measurable standard, e.g. from a Service Level Agreement (SLA).



“The art of compromise”. This is all about risk management. Risk vs. Cost. If, working together, you and your security folks can show value to onboard firewalls, go for it. But remember that it doesn’t have to be All-or-Nothing. The workload you REALLY want are the web servers, application servers, and the databases.



One implementation.

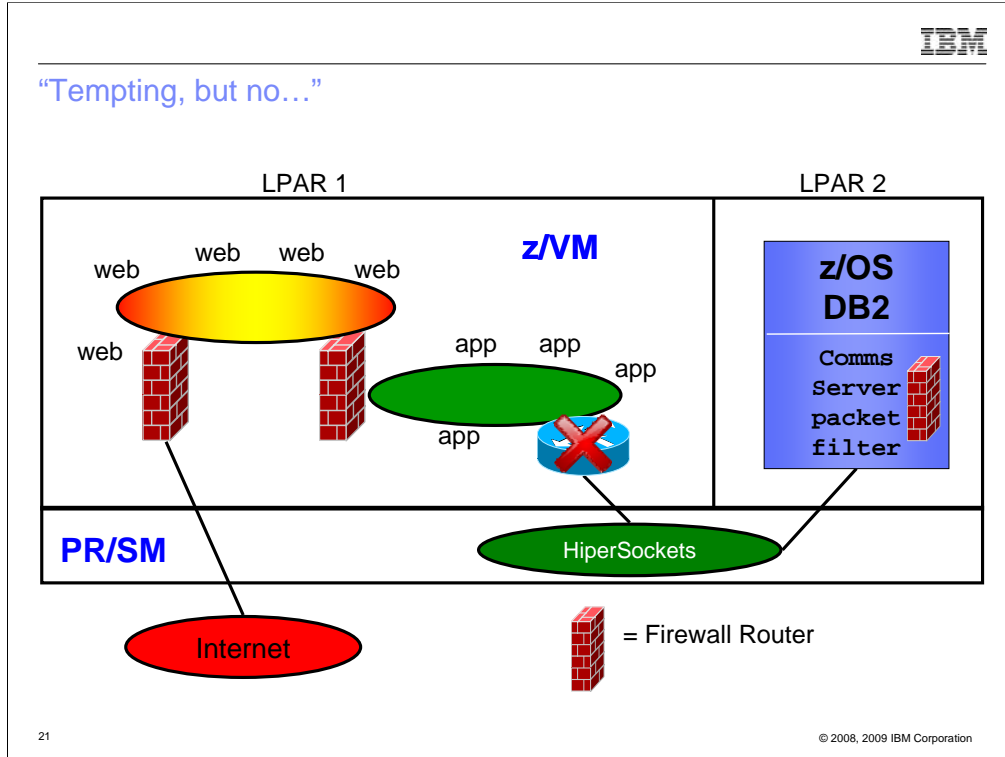


Another implementation.

Notice that the HiperSocket changed color. You can't change the color until you hit a firewall.

We're using z/OS's built-in packet filtering technology and dedicated HiperSockets (which gives us QDIO Assist!!). Of course, use of z/OS packet filters must be negotiated.

NOTE: You can use a z/OS LPAR to enter a sysplex. Could be used to avoid encryption between app and data tier.



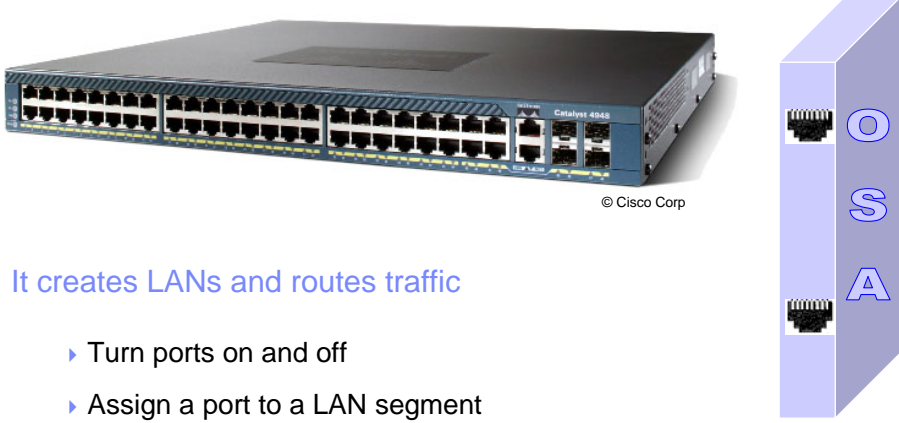
Same as previous chart, but all app->z/OS traffic is funneled through a virtual router with no firewall. BAD IDEA. It doesn't have any value add.



Virtual Switches VLANs and traffic separation

IBM

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

23

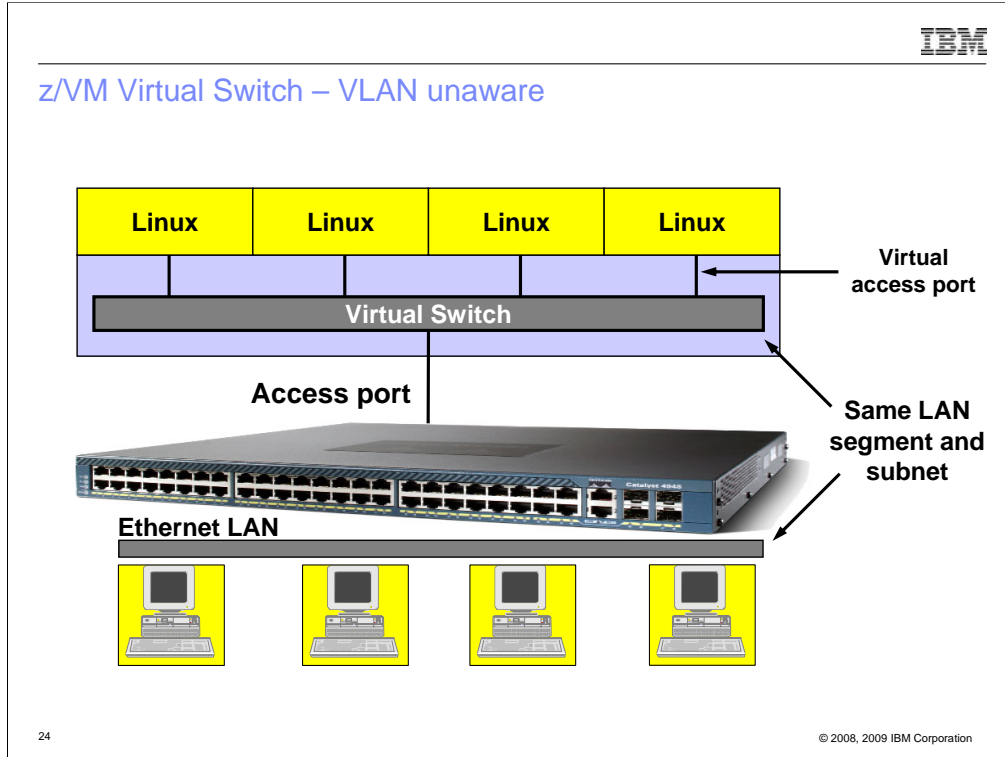
© 2008, 2009 IBM Corporation

Some folks haven't ever seen a switch. Similar to a hub (like you may have at home), but (a) better technology, (b) more function, and (c) waaaay more expensive!

“Hub” – the ports are physically connected. Each port sees all other ports. All ports the same speed. “Dumb.”


“Switch” – the ports are logically connected based on administrative settings in the switch. “Smart.”

You get what you pay for.



A VLAN Unaware VSWITCH plugs into an **access port**. It sees only a single LAN segment. Think of it as monochromatic or color blind. It neither knows nor cares about other LAN segments. This is the closest analogy to an ethernet hub.

IEEE VLANs



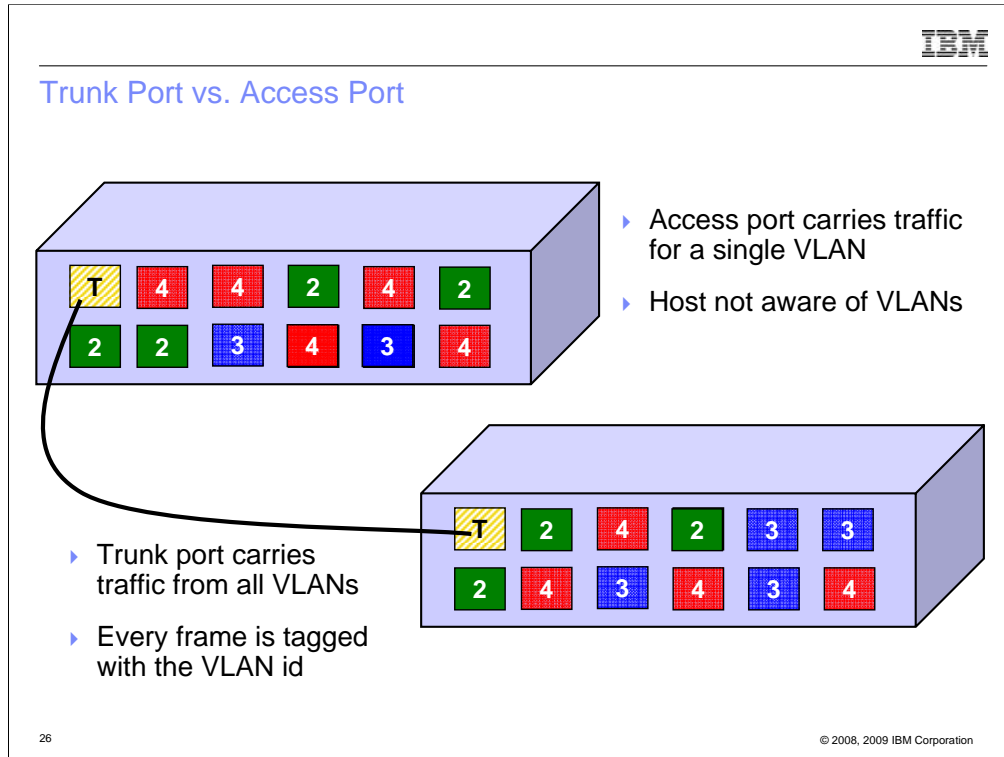
© Cisco Corp

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

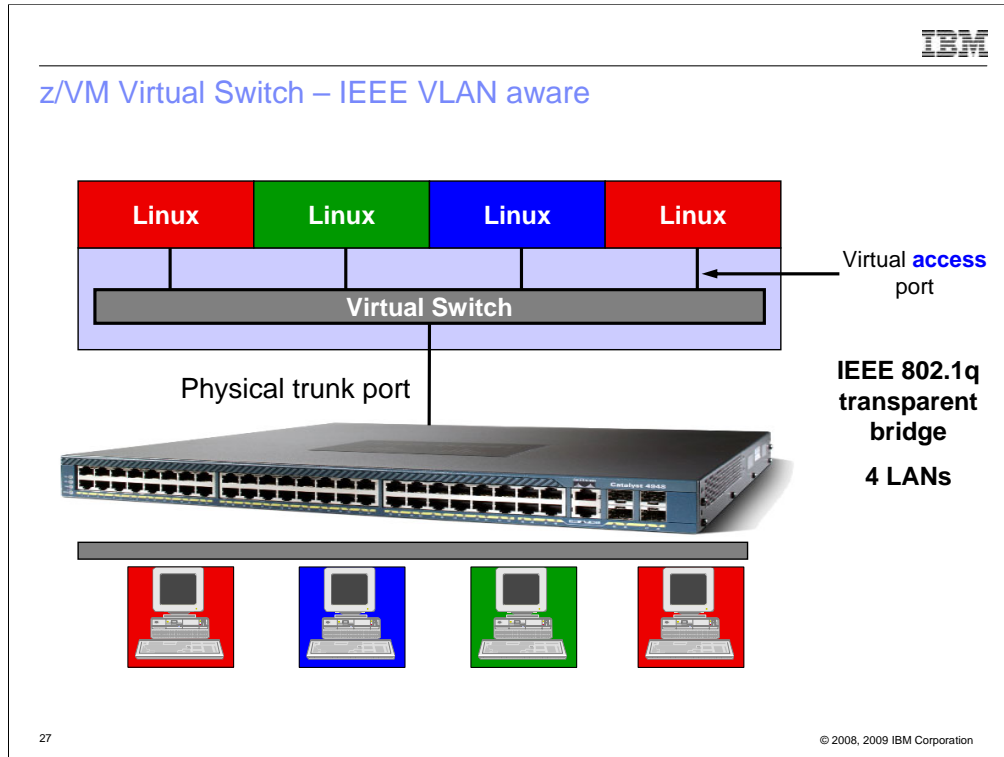
25

© 2008, 2009 IBM Corporation

Some folks haven't ever seen a switch. What do you do when it fills up? Throw it away? No, you connect it to another switch via a **TRUNK PORT**.



A **trunk port** carries ethernet frames with an extra piece of information called the *VLAN ID tag*. This *tag* tells the target switch what VLAN (LAN segment) the frame belongs to. The origin switch is responsible for adding the tag. The tag is removed before it is sent out on an **access port**. Tags are not accepted from an access port – they will be treated as a malformed frame.



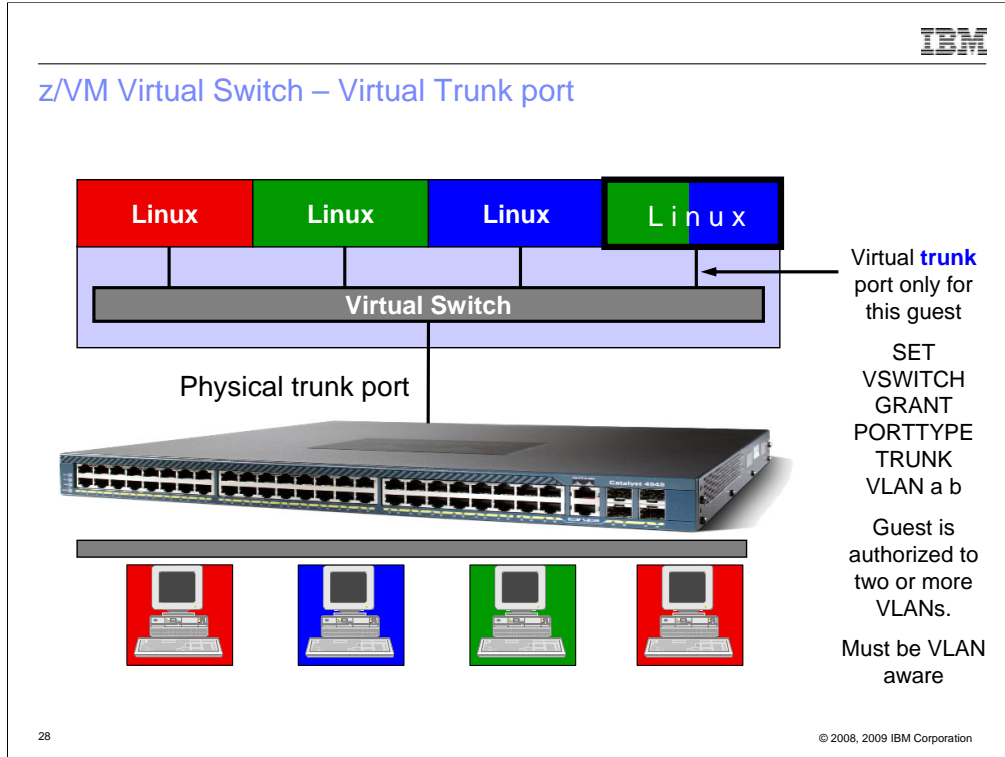
Technicolor! Each guest is placed into a VLAN.

The VSWITCH is plugged into a **trunk port** and can receive data from or send data to any VLAN which has been authorized on that port by the switch administrator. In fact, the VSWITCH is required to tag outgoing frames and remove tags from incoming frames. However, if a guest is authorized for the NATIVE VLAN id, then its frames will go untagged and the switch will assign the VLAN id. (Usually VLAN 1.) This is why it is important to know what the native VLAN id is of the switch.

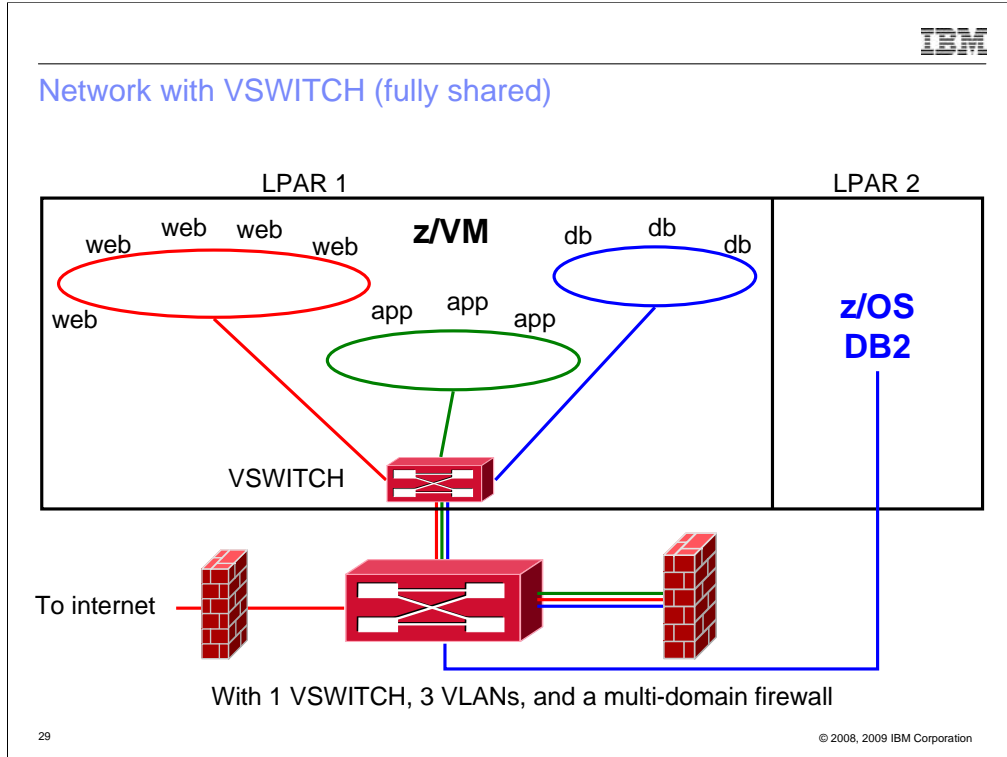
In IEEE 802.1q parlance, the native VLAN id is the default port VLAN id for a trunk port.

Remember to authorize all OSAs used by the VSWITCH with the same VLANs. If you don't, "unpredictable results are guaranteed to occur."

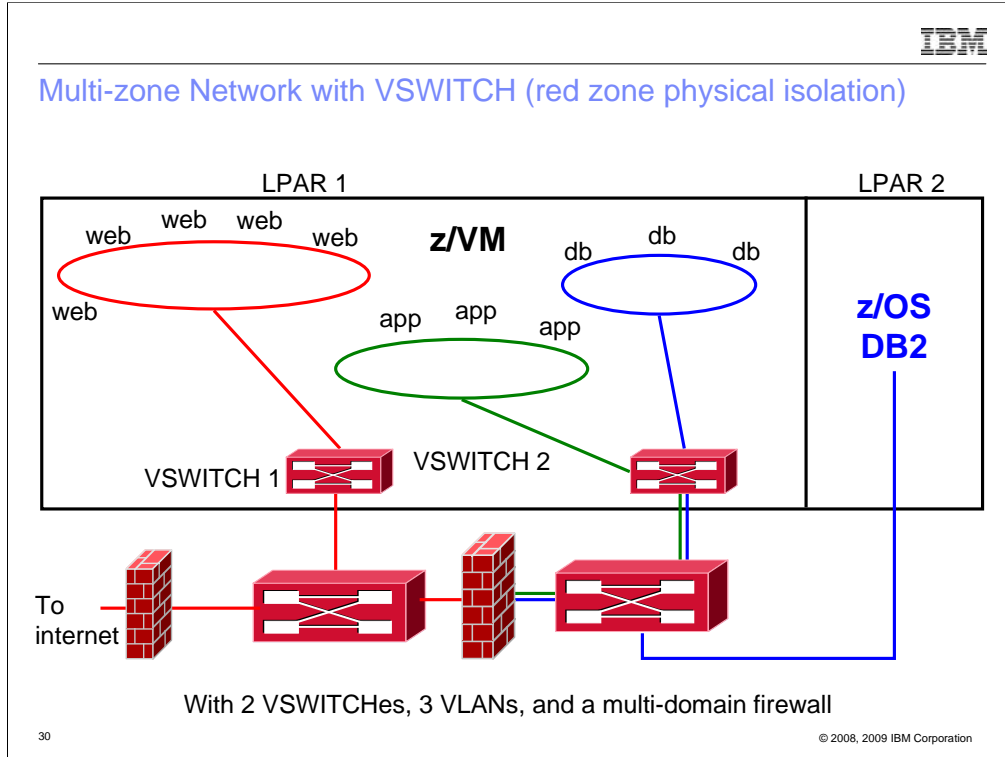
Do NOT specify PORTTYPE TRUNK on DEFINE VSWITCH!!!! That will change the default port type for each guest, having nothing to do with the OSA port. (The "VLAN" keyword on DEFINE VSWITCH is what triggers the use of VLANs.)



In this case, one of the guests has been authorized for two VLANs. It must be configured to use VLANs. If guest is not authorized for the NATIVE VLAN, untagged frames will be discarded.



One VSWITCH carrying data for three LAN segments. Only 2 OSAs are required.




Two VSWITCHes, four (4) OSA ports. The Red zone network traffic is physically isolated (layer 2) from Green and Blue. VSWITCH 1 can be VLAN unaware. VSWITCH 2 is VLAN-aware. If you want to physically isolate Green and Blue, then you need another switch and another pair of OSA ports.

With a second LPAR, you need another 6 OSA ports.



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

33 © 2008, 2009 IBM Corporation

RACF creates **objects** in **classes**. The class is the name space. Each instance of an object in a class represents another VM resource (disk, vswitch, user, ...).

On z10 and later, VSWITCH port isolation also activates QDIO data connection isolation, disabling the “short circuit” in the OSA.



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?

99 bottles of beer on the wall....



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

...99 bottles of beer! Fuhgeddaboutit. That's too hard.



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

36

© 2008, 2009 IBM Corporation

Security administrator overrides the end user and the system programmer.



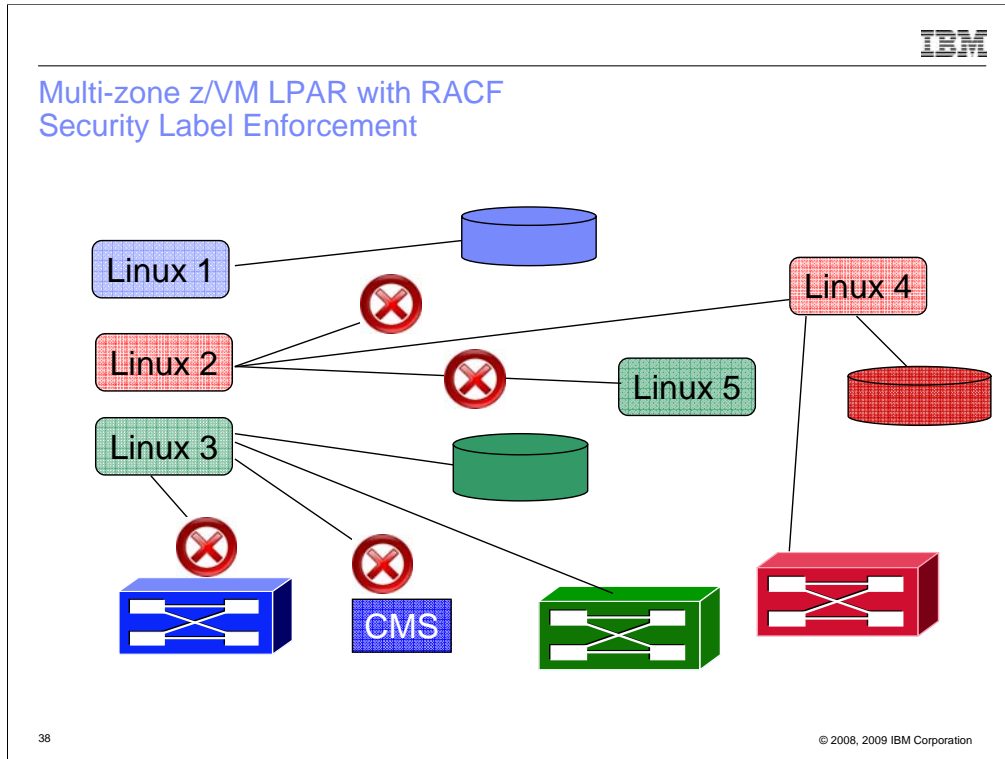
Multi-Zoning with RACF

- 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

37

© 2008, 2009 IBM Corporation

“Labeled security” is an old, well-understood concept in the industry and is part of z/VM’s Common Criteria certification. Note that security labels (aka multilevel security, MLS) are available only with RACF. None of the CA products (VM:Secure, ACF2, Top Secret) support it.



Let's virtually "color code" our guests and resources. No crossing color boundaries.



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)
```

Create 4 security labels: PUBLIC, RED, GREEN, BLUE.



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)
```

It may seem silly to have to do the PERMIT **and** the ALTUSER, but you need to assign the user a default label. (Otherwise user has no label, in which case SETROPTS controls what happens.)



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)
```

SYSNONE is pre-defined. Don't confuse it with **NONE**, **SYSHIGH**, or **SYSLOW**. The others cause certain results on a label comparison; **SYSNONE** causes the label check to be bypassed. BTW, label checks are not character string checks, but an analysis of the security level and categories that comprise the label.



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

Protect or don't protect. There is very little middle ground (by design).



Multi-Zoning with RACF

- Activate RACF protection
 - SETROPTS CLASSACT(SECLABEL VMMDISK VMLAN)
 - SETROPTS RACLIST(SECLABEL)
 - SETROPTS MACTIVE(WARNINGS)
 - If resource doesn't have a seclabel, message is issued and seclabels are ignored.
- Or
 - SETROPTS MACTIVE(FAILURES)
 - If resource doesn't have a seclabel, command fails.
 - This is more secure!

Last steps...be careful...read the book!



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!!

“Constant vigilance!”



Reference Information

- This 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>