



Virtual Networking with z/VM Guest LANs and the z/VM Virtual Switch

Session 9125

**Alan Altmark, IBM
z/VM Development, Endicott, NY**

Note

References to IBM products, programs, or services do not imply that IBM intends to make these available in all countries in which IBM operates. Any reference to an IBM product, program, or service is not intended to state or imply that only IBM's product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe on any of the intellectual property rights of IBM may be used instead. The evaluation and verification of operation in conjunction with other products, except those expressly designed by IBM, are the responsibility of the user.

The following terms are trademarks of the International Business Machines Corporation in the United States or other countries or both:

IBM	IBM logo	eServer	zSeries
System z9	DB2	z/OS	z/VM

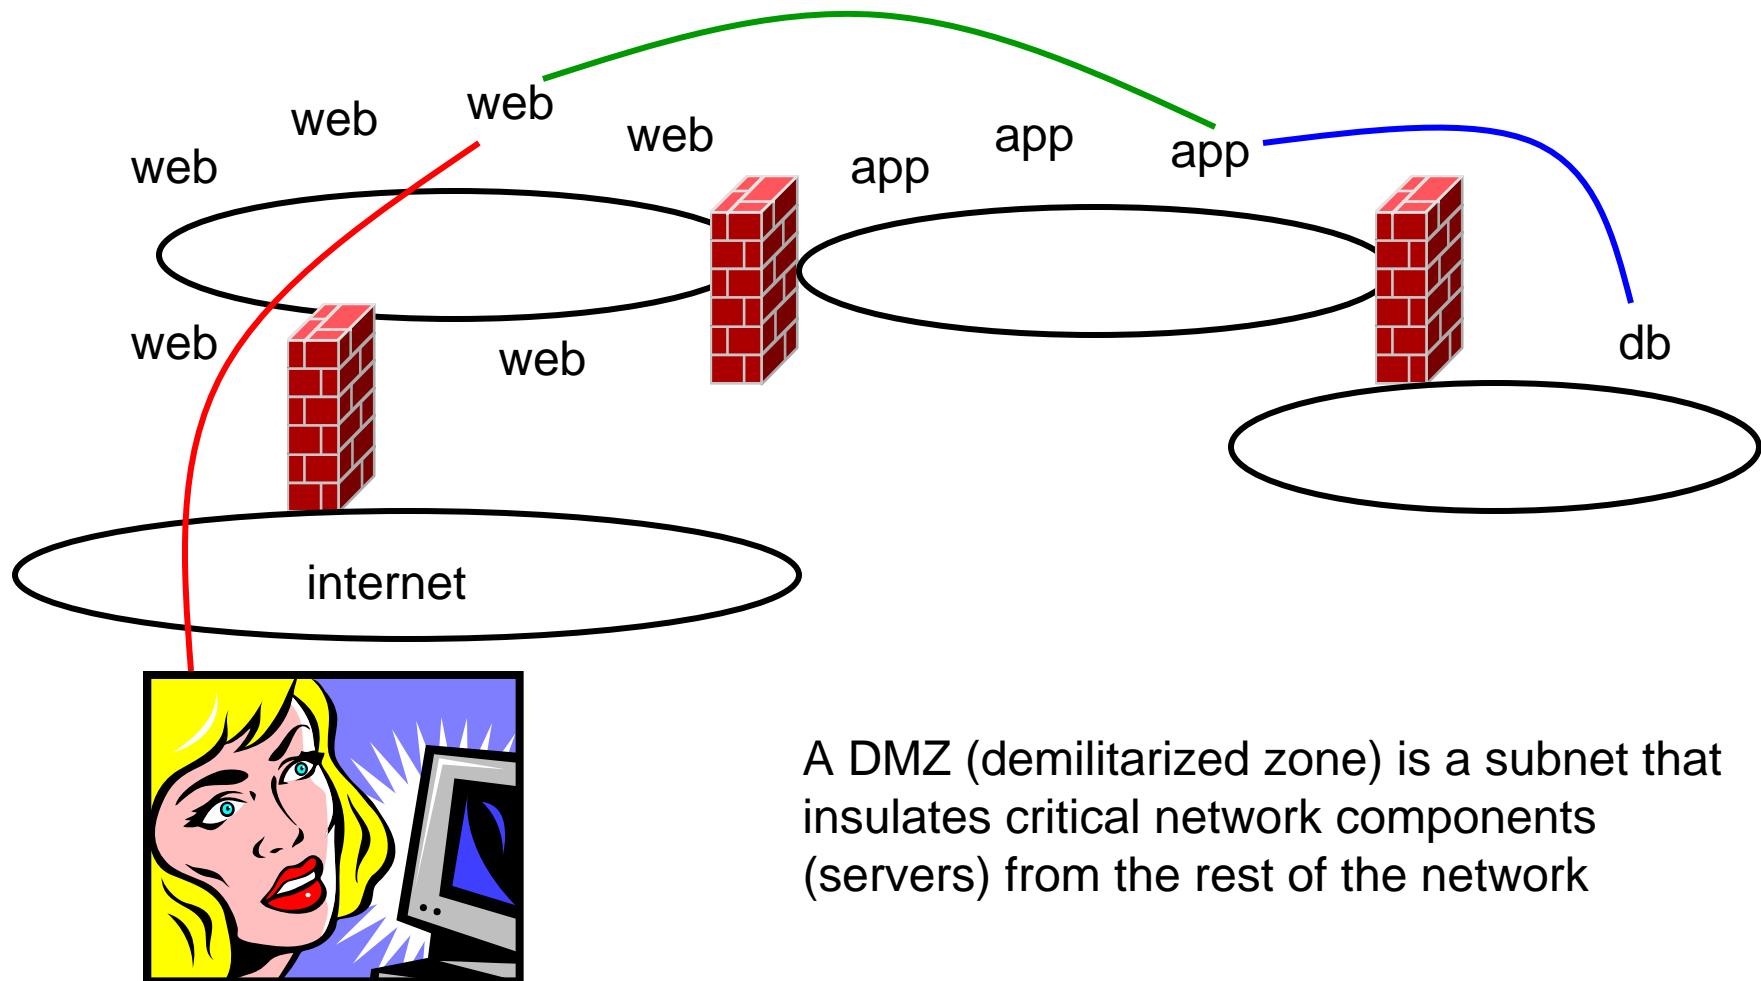
Other company, product, and service names may be trademarks or service marks of others.

© Copyright 2003, 2007 by International Business Machines Corporation

Topics


- Overview
- Guest LANs
- Virtual Network Interface Card
- Virtual Switch
- What's new in z/VM Version 5.1 and 5.2

Multi-DMZ Network

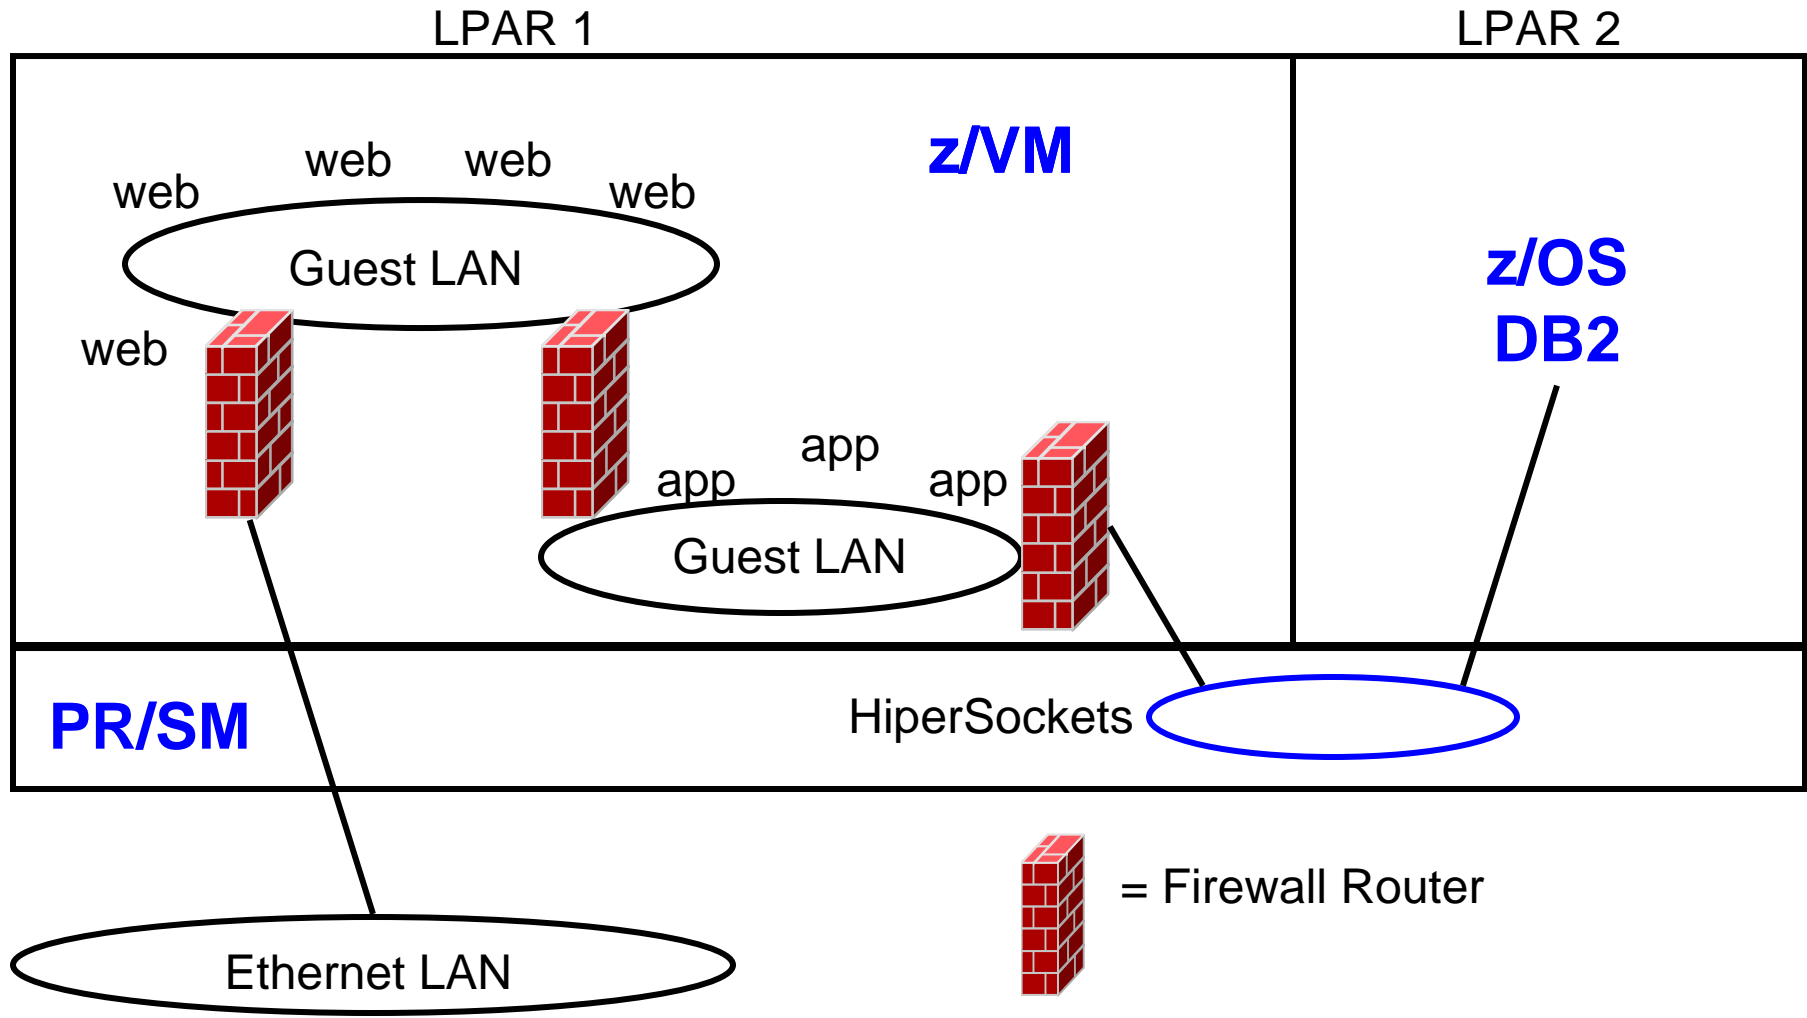


A DMZ (demilitarized zone) is a subnet that insulates critical network components (servers) from the rest of the network

Multi-DMZ Network on zSeries



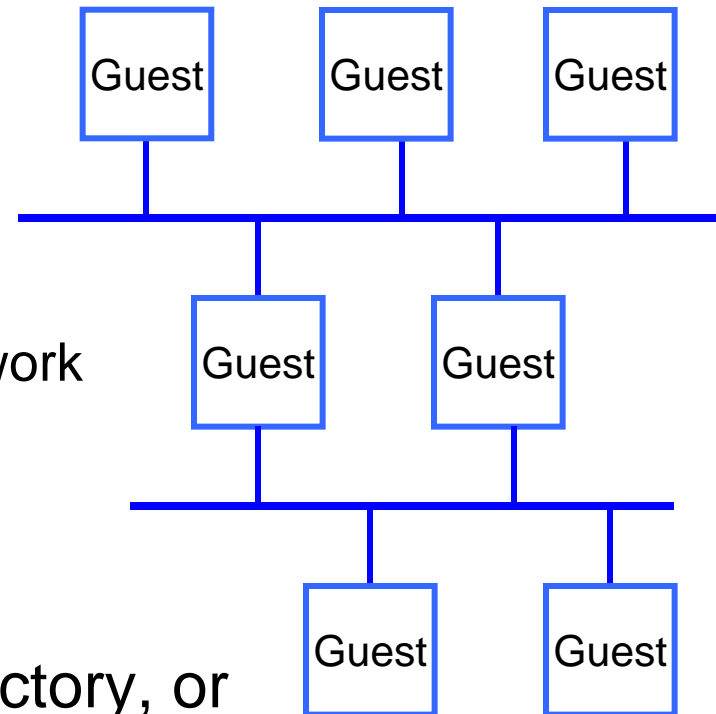
Multi-DMZ Network with Guest LANs



Guest LANs

z/VM Guest LAN

- A simulated LAN
 - ▶ Ethernet: IPv4 and IPv6
 - ▶ HiperSockets: IPv4
 - ▶ No built-in connection to outside network
- As many as you want
- Created in `SYSTEM CONFIG`, directory, or by `CP DEFINE LAN` command



Why Guest LAN instead of Dedicated Hardware ?

- Dedicated network connections may be best for some environments:
 - ▶ When intense network activity is expected and you need to guarantee network access bandwidth
 - ▶ When direct access to physical network is required

- z/VM Guest LAN may be better for other environments:
 - ▶ When network hardware is limited
 - ▶ When multiple nodes are guests in the same z/VM host image
 - ▶ When network activity must be isolated from primary network (e.g. test environments, student labs, application server access to database servers)

Primary Guest LAN Attributes

- Name & Owner
- Type
- Access list
- Maximum frame size (HiperSockets only)

- Some attributes can be changed after the LAN is defined

- There are some others not discussed here
 - ▶ Maximum number of connections
 - ▶ Accounting

LAN Name and Owner

- The LAN name is a simple 1-8 character token
- The LAN owner is a VM user ID or “SYSTEM”
- (name, owner) is unique within the system
- A Class G LAN owner can
 - ▶ modify the LAN access list
 - ▶ delete the LAN
- A Class B user can create, modify, or detach any LAN

HiperSockets or Ethernet

TYPE HIPERsockets | QDIO [IP | ETHERNET]

- HiperSockets
 - ▶ Synchronous
 - ▶ Low latency
 - ▶ Slightly smaller path length in CP (less CP time)

- QDIO
 - ▶ OSA-Express in QDIO mode
 - ▶ Asynchronous
 - ▶ Higher latency than HiperSockets
 - ▶ Higher CPU cost
 - ▶ IP = Layer 3, ETHERNET = Layer 2

Access list

■ Unrestricted

- ▶ Any user can connect (couple) to this LAN
- ▶ Hint: CP QUERY LAN can show you who is connected

■ Restricted

- ▶ Only users in the access list can connect (couple) to this LAN
- ▶ LAN owner uses CP SET LAN to GRANT or REVOKE access
- ▶ CP QUERY LAN can show you the current access list
- ▶ CP QUERY LAN can show you who is connected

■ External Security Manager

- ▶ RACF/VM support for new VMLAN objects

Maximum Frame Size (HiperSockets only)

MFS 16K | 24K | 40K | 64K

- Simulates CHPID OS=*value* specification in IOCDs for HiperSockets (TYPE=IQD) chpids
 - ▶ Does not apply to QDIO

- Largest MTU specification = (MFS - 8K)

- Hints:
 - ▶ If guest LAN is isolated, use large MFS and large MTU
 - ▶ If LAN has external gateway, use MFS 16K and match external MTU (e.g. 1492)
 - ▶ Jumbo frame (MTU 8992) gateway needs 24K MFS

Persistent vs. Transient LAN

- Persistent / Transient is inferred from other attributes
 - ▶ Any LAN owned by user “SYSTEM” is *persistent*
 - ▶ Any LAN created by SYSTEM CONFIG is *persistent*
 - ▶ All other LANs are *transient*
- A *persistent* LAN must be explicitly deleted by CP DETACH LAN
- A *transient* LAN is automatically deleted when the last user uncouples from the LAN

Setting Guest LAN defaults and limits

- Set global guest LAN attributes in the SYSTEM CONFIG file:

```
VMLAN LIMIt PERSistent INFinite|maxcount
VMLAN LIMIt TRANSient INFinite|maxcount
VMLAN ACNT|ACCOUNTing SYSTEM ON|OFF
VMLAN ACNT|ACCOUNTing USER ON|OFF
VMLAN MACPREFIX 020000-02FFFF
VMLAN MACIDRANGE SYSTEM x-y [USER a-b]
```

- VMLAN LIMIT TRANSIENT 0 prevents dynamic definition
- Use SET VMLAN to change dynamically



Virtual MAC Addresses

- Each instance of CP should have a unique VMLAN MACPREFIX
- Virtual MAC = MACPREFIX || MACID
- VMLAN MACIDRANGE
 - ▶ SYSTEM – The range of MACIDs from which CP will select a dynamically defined MAC
 - ▶ USER – The range of MACIDs reserved by CP for NICDEF. All MACIDs on NICDEFs must be in this range.
 - ▶ USER is a subset of SYSTEM

Create a Guest LAN

- DEFINE LAN in SYSTEM CONFIG

```
DEFINE LAN name [OWNERid ownerid]  
                [TYPE HIPERsockets|QDIO]  
                [MAXCONN INFinite|nnnn]  
                [MFS 16K|24K|40K|64K]  
                [ACCOUNTing ON|OFF]  
                [UNRESTRicted|RESTRicted]  
                [GRANT userlist]
```

Examples:

```
DEFINE LAN QDIO5 OWNER SYSTEM TYPE QDIO
```

- CP DEFINE LAN to create dynamically

```
DEFINE LAN NET9 OWNER SYSTEM RESTRICTED TYPE QDIO
```

Grant Guest LAN Access

- DEFINE LAN and MODIFY LAN in SYSTEM CONFIG

```
MODIFY LAN  name
             [OWNERid ownerid / OWNERID SYSTEM]
             [GRANT userid]
```

Example:

```
DEFINE LAN HIPER1 OWNER SYSTEM RESTRICTED
MODIFY LAN HIPER1 OWNER SYSTEM GRANT LINUX01
MODIFY LAN HIPER1 OWNER SYSTEM GRANT LINUX02
```


- CP SET LAN to change dynamically

```
CP SET LAN HIPER1 OWNER SYSTEM GRANT LINUX03
```

Virtual Network Interface Card

Virtual Network Interface Card (NIC)

- A simulated network adapter
 - ▶ OSA-Express QDIO
 - ▶ HiperSockets
 - ▶ Must match LAN type
- 3 or more devices per NIC
 - ▶ More than 3 to simulate port sharing on 2nd-level system or for multiple data channels
- Provides access to Guest LAN or Virtual Switch
- Created by directory or CP DEFINE NIC command



Virtual NIC - User Directory

- May be automated with USER DIRECT file:

```
NICDEF vdev [TYPE HIPERS | QDIO]
           [LAN owner name]
           [DEVICES nn]
           [CHPID xx]
           [MACID xyyyzz]           Combined with VMLAN
                                     MACPREFIX to create
                                     virtual MAC
```

Example:

```
NICDEF 1100 LAN SYSTEM SWITCH1 CHPID B1 MACID B10006
```

Virtual NIC - CP Command

- May be interactive with CP DEFINE NIC and COUPLE commands:

```
CP DEFINE NIC vdev
      [[TYPE] HIPERsockets | QDIO]
      [DEVICES devs]
      [CHPID xx]
```

```
CP COUPLE vdev [TO] owner name
```

Example:

```
CP DEFINE NIC 1200 TYPE QDIO
CP COUPLE 1200 TO SYSTEM CSC201
```

NIC CHPID parameter

CHPID xx

- Specifies the Channel Path ID number (in hex) to use for this NIC
 - ▶ Default is any available unused real CHPID number

- Needed for z/OS guest because HiperSockets are managed by CHPID number

- **This is a virtual CHPID number**


Virtual Switch

What's a 'switch' anyway?




© Cisco Corp

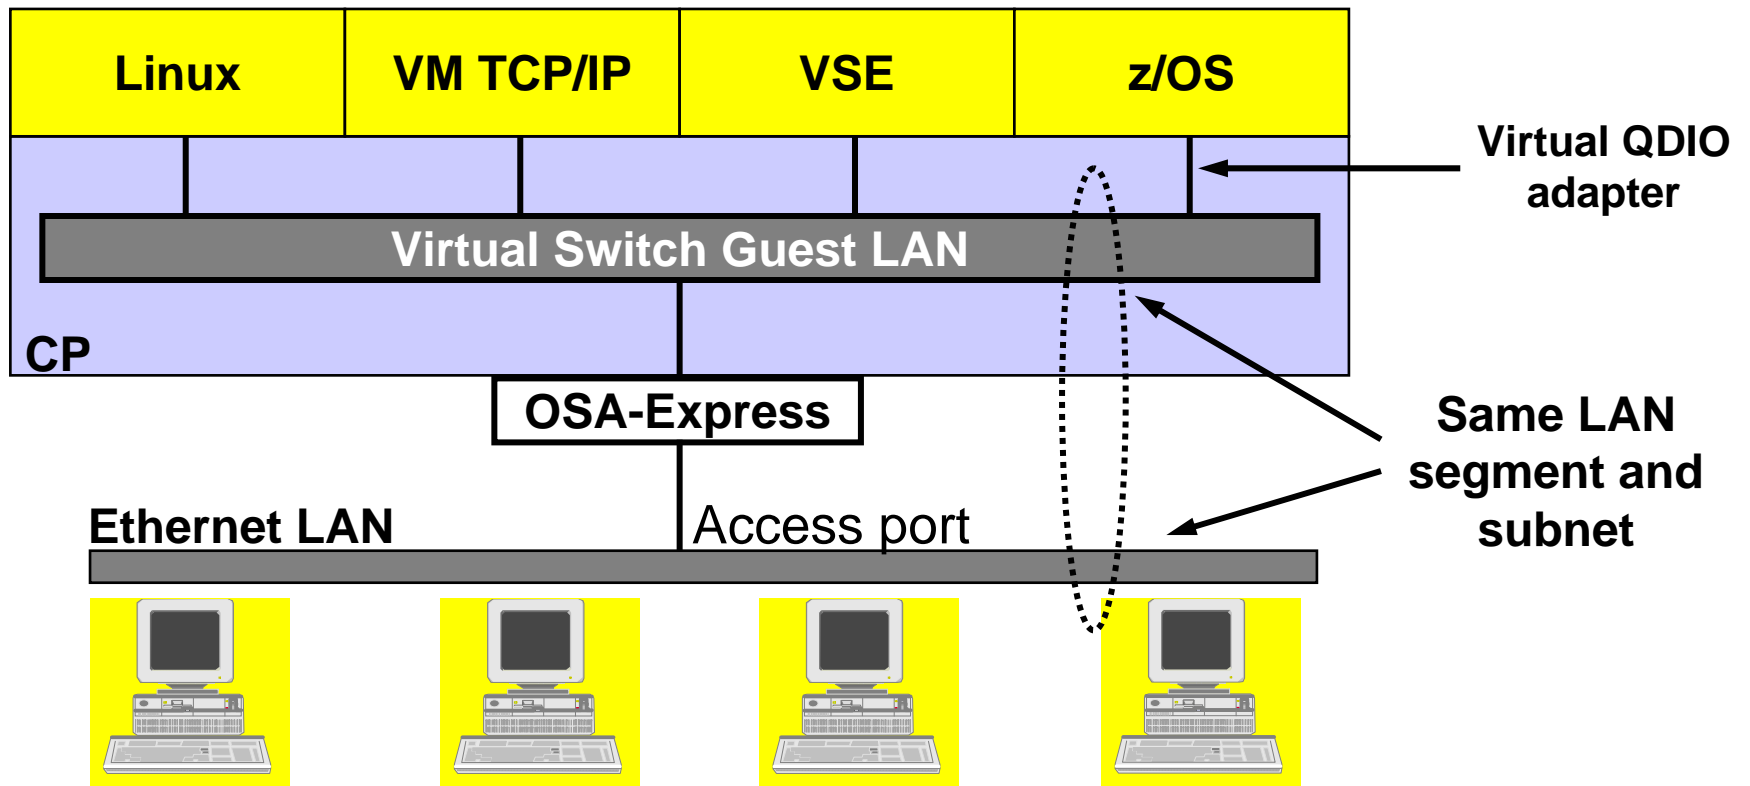
- ▶ A box that creates a LAN
- ▶ It can be remotely configured
 - ▶ E.g. Turn ports on and off
- ▶ Similar to a home router



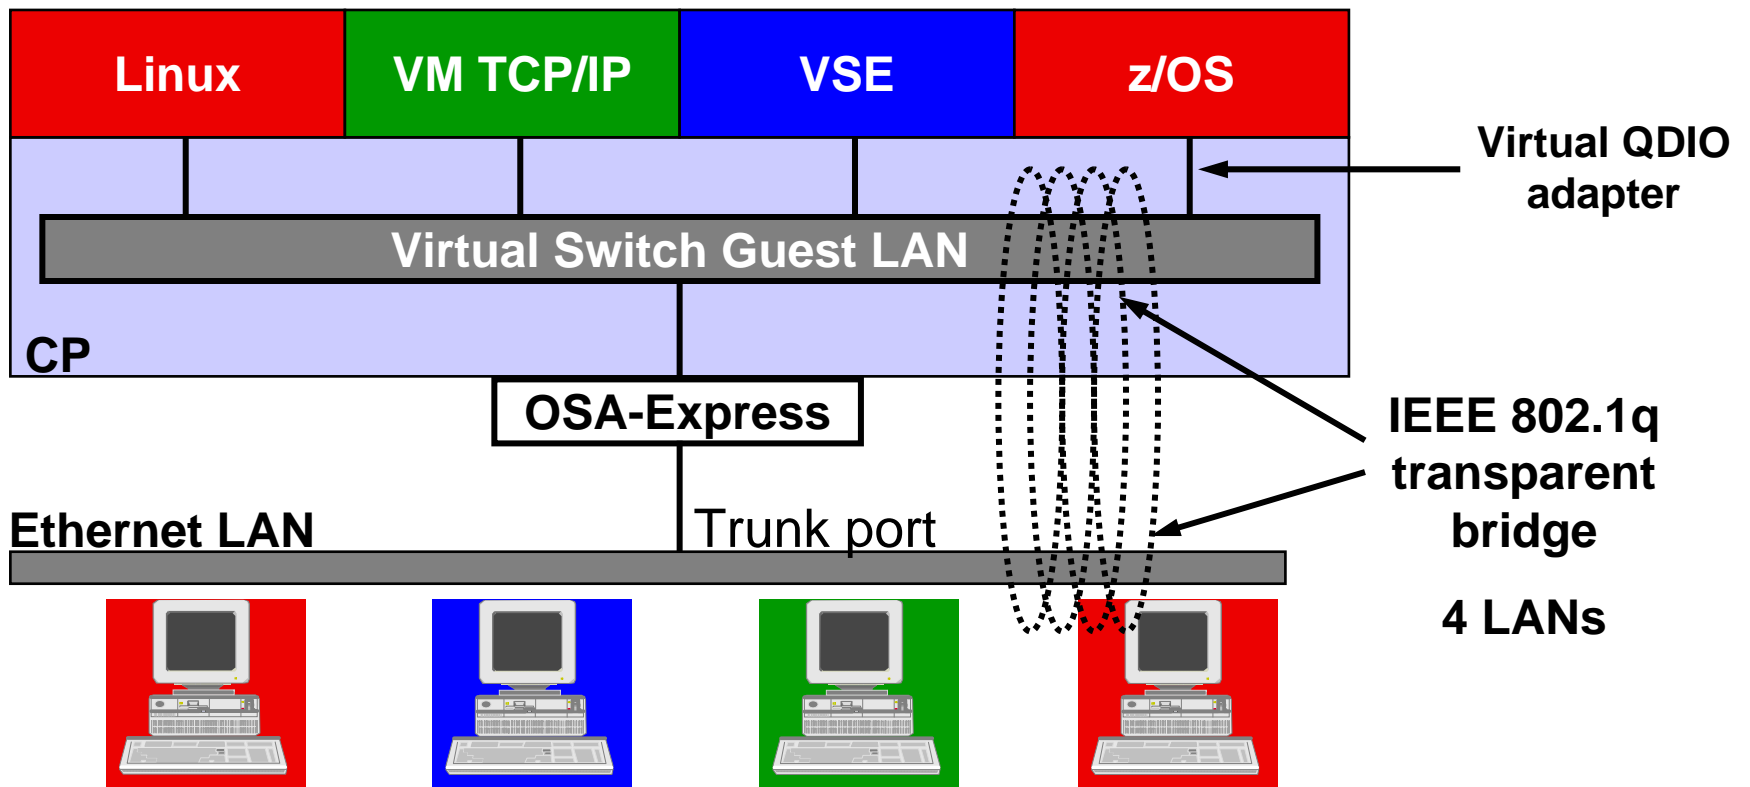
A VLAN-aware switch: An inside look




z/VM Virtual Switch – VLAN unaware




z/VM Virtual Switch – VLAN aware



Trunk Port vs. Access Port

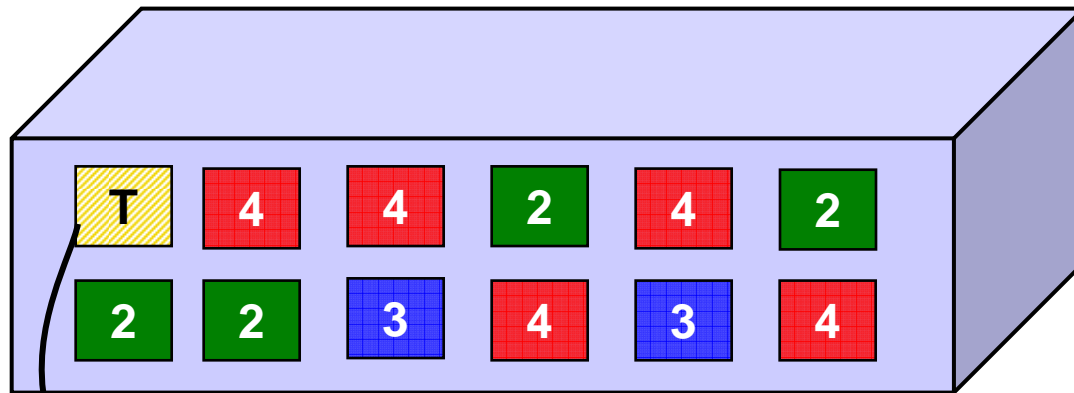


- ▶ Access port carries traffic for a single VLAN
- ▶ Host not aware of VLANs



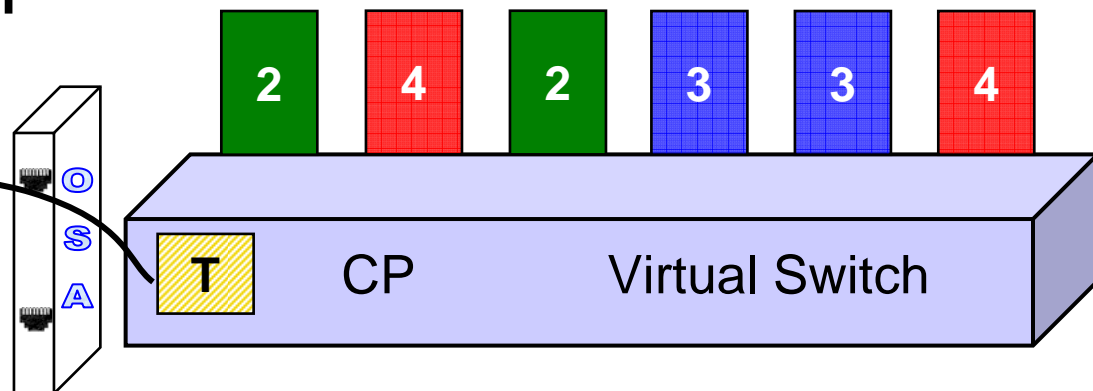
- ▶ Trunk port carries traffic from all VLANs
- ▶ Every frame is tagged with the VLAN id

Physical Switch to Virtual Switch



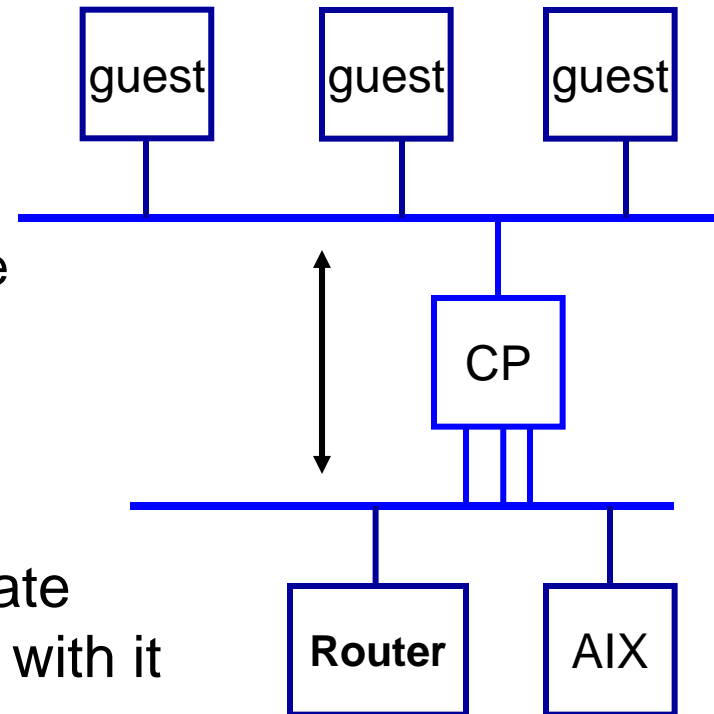
▶ Trunk port carries traffic between CP and switch

▶ Each guest can be in a different VLAN



z/VM Virtual Switch

- 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
 - ▶ User needs IUCV *VSWITCH authorization
 - ▶ User needs VSWITCH CONTROLLER statement in PROFILE TCPIP
- 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]
```

z/VM 5.3

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
```

- z/VM 4.4 supports “VLAN ANY”, but it’s removed in z/VM5.1!

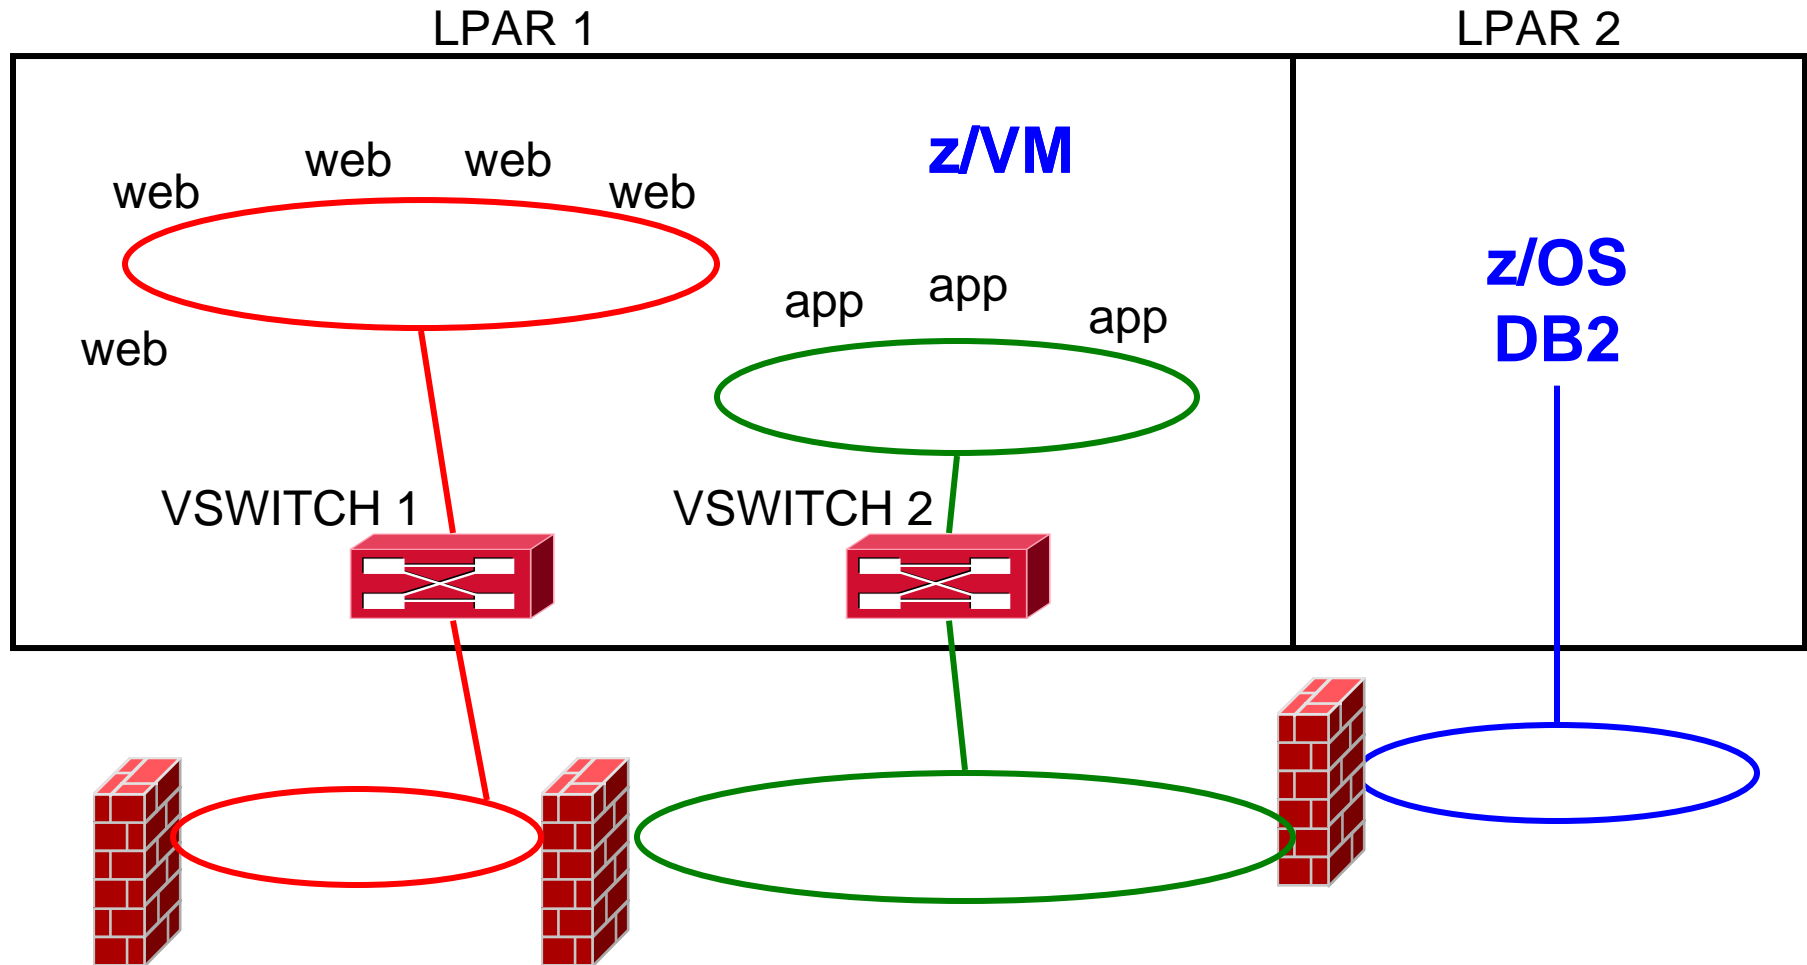
Multi-DMZ Network on zSeries - Reloaded



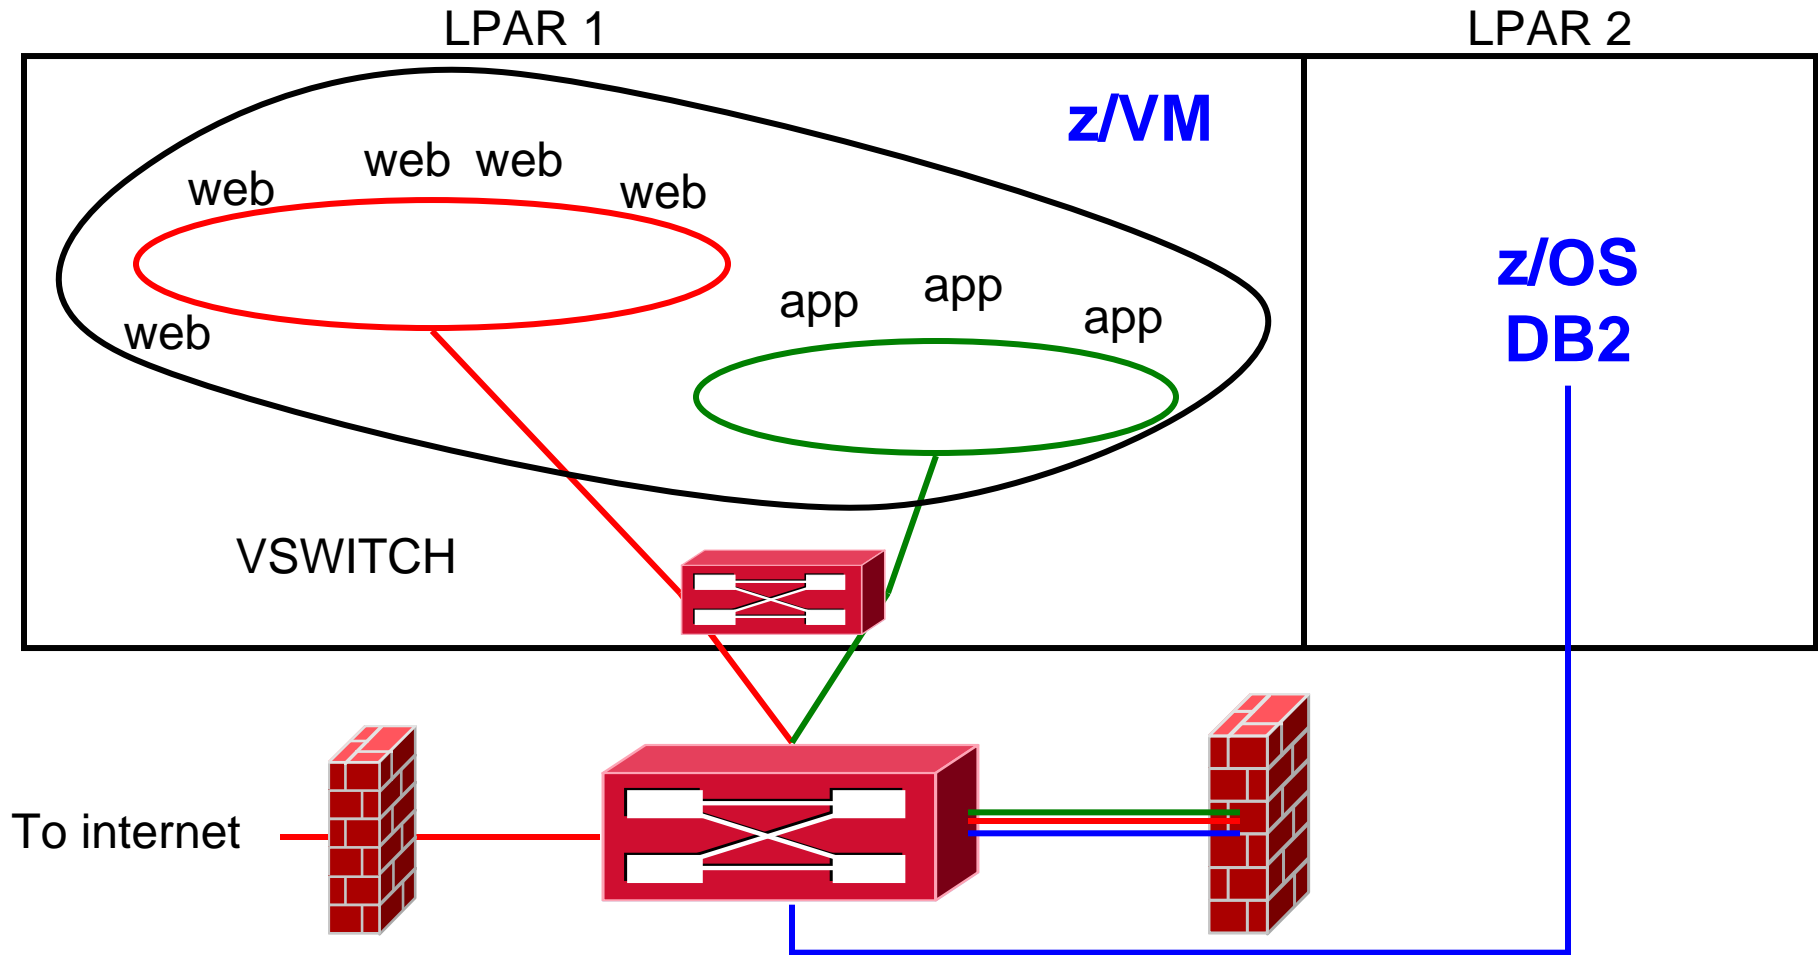
Multi-DMZ Network on zSeries with outboard firewall



Multi-DMZ Network with VSWITCH (A)



Multi-DMZ Network with VSWITCH (B)



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

What's new?

New in z/VM 5.3

- OSA link aggregation
 - ▶ IEEE 802.1ad
 - ▶ Hyperchannel
 - ▶ Channel Bonding

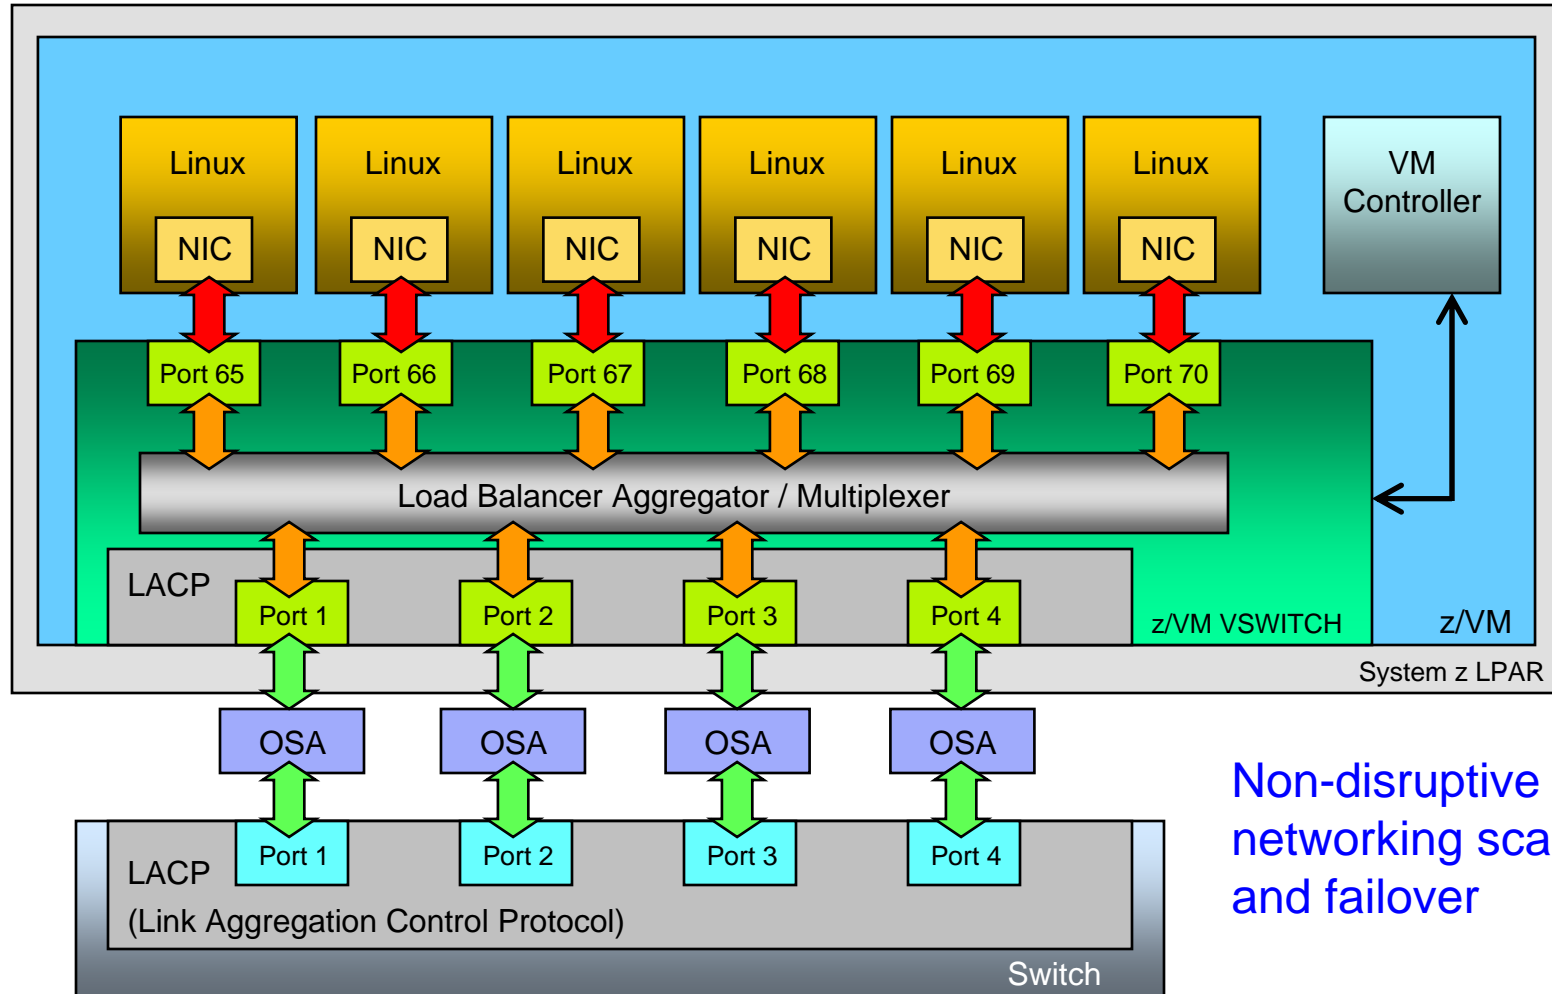
- SNMP virtual switch monitor

OSA-Express2 Link Aggregation Support

- **z/VM support for IEEE 802.3ad Link Aggregation**
- **Requires associated OSA-Express2 support announced on April 18, 2007 for IBM System z9 servers**
- **Groups available OSA-Express2 adapters for use by the z/VM Virtual Switch**
 - Up to 8 adapters can be aggregated 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
- **Enables increased scalability for virtual network I/O**
- **Includes support to recover from a failed external switch**
- **Enhances support for business continuance**

z/VM VSWITCH Link Aggregation Support


Enhanced Networking Bandwidth and Business Continuance



Non-disruptive
networking scalability
and failover

z/VM Virtual Switch SNMP Support

- **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**



New in z/VM 5.2...

■ 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

■ 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

New in z/VM 5.1...

- ESM control for all guest LANs and VSWITCHes, including VLAN ID control
 - ▶ RACF: Class VMLAN, Profile owner.laname or owner.laname.vid
 - ▶ All Guest LANs and VSwitches can be controlled

- Layer 2 (MAC) communications
 - ▶ Fulfillment of Statement of Direction
 - ▶ All types of traffic, not just IP
 - ▶ Virtual NIC MAC appears on network
 - ▶ VMLAN updates to allow specification of ranges used for automatic and static MAC address assignments


- Better VSWITCH stall detection, error reporting, and error recovery.

New in z/VM 5.1...

- IEEE 802.1q compliance changes
 - ▶ VLAN ANY is gone
 - ▶ VSWITCH can be defined as VLAN-aware (or not). Default is “not”.
 - ▶ When a NIC couples to a VLAN-aware VSWITCH, it will be assigned a PORTTYPE attribute
 - ACCESS: VLAN tags not given to or accepted from guest
 - TRUNK: VLAN tags are given to and expected from guest
 - ▶ Default PORTTYPE comes from DEFINE VSWITCH
 - Can be overridden by MODIFY VSWITCH GRANT
 - ▶ Some configurations require migration effort

Some Final Thoughts...

Guest LAN vs. Virtual Switch



- Virtual router is required
- Different subnet
- External router awareness
- Guest-managed failover
- No virtual router
- Same subnet
- Transparent bridge
- CP-managed failover

Network Configuration

- Guest LANs require a new subnet and the use of a virtual router
- A Virtual SWITCH extends the subnets you already have
- By having virtual and real configurations be the same, you can easily test network configuration before deployment with real hardware

Thanks for Listening!

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

Support Summary

z/VM V5.3	<ul style="list-style-type: none">▪ Link aggregation▪ Separation of default VLAN id from native VLAN id▪ SNMP monitor
z/VM V5.2	<ul style="list-style-type: none">▪ Virtual SPAN ports for sniffers
z/VM V5.1	<ul style="list-style-type: none">▪ Virtual trunk and access port controls▪ Removal of VLAN ANY▪ Layer 2 (MAC) frame transport▪ Improved virtual switch error detection & recovery▪ External security manager access control
z/VM V4	<ul style="list-style-type: none">▪ IPv4 Virtual Switch with IEEE VLANs▪ IPv4 HiperSocket Guest LAN▪ IPv4 and IPv6 QDIO Guest LAN

References

- 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/>

Contact Information

- By e-mail: `Alan_Altmark@us.ibm.com`
- In person: USA 607.429.3323
- On the Web: <http://ibm.com/vm/devpages/altmarka>
- Mailing lists:
`IBMTCP-L@vm.marist.edu`
`IBMVM@listserv.uark.edu`
`LINUX-390@vm.marist.edu`

<http://ibm.com/vm/techinfo/listserv.html>