



Session V24

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

Tracy Adams

IBM
SYSTEM z9 AND zSERIES EXPO
October 9 - 13, 2006

Orlando, FL

© IBM Corporation 2006

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, 2006 by International Business Machines Corporation

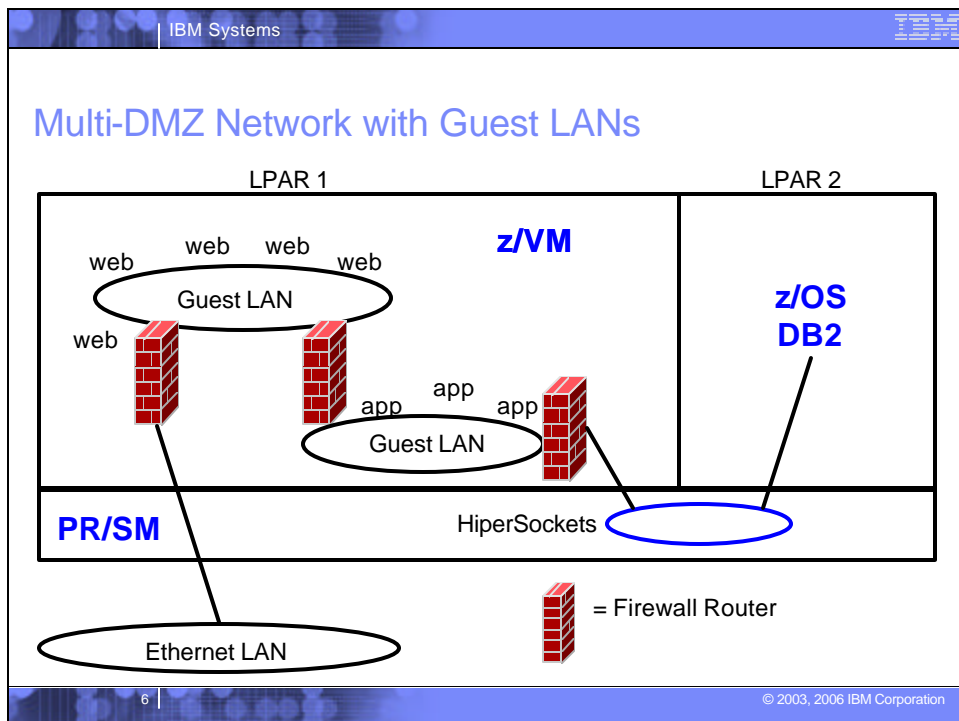
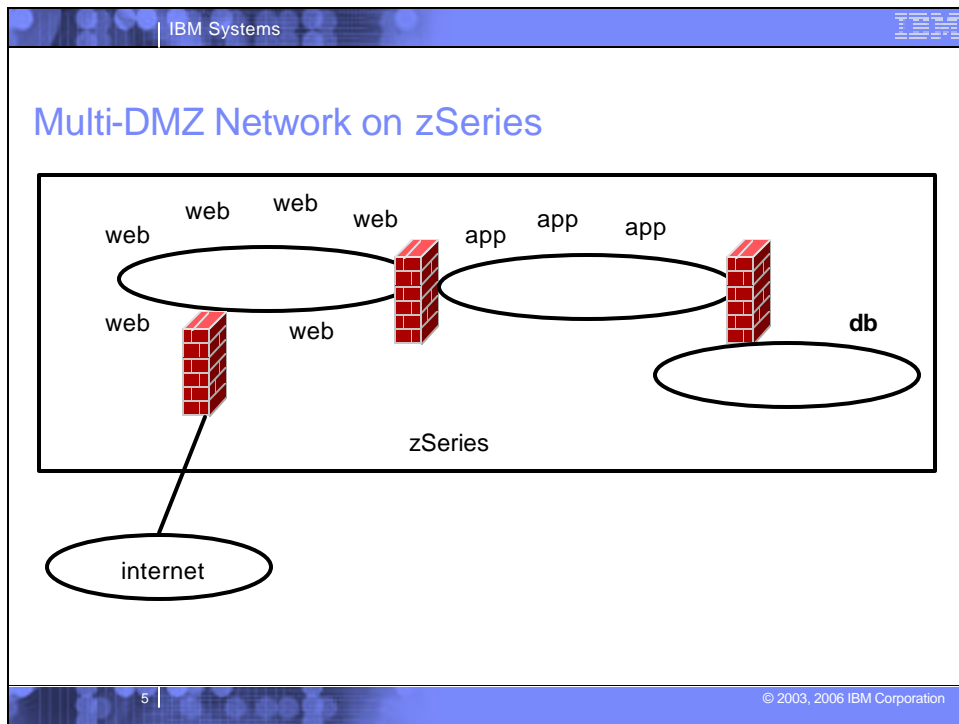
Topics

- 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





IBM Systems

z/VM Guest LAN

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

```
graph TD; G1[Guest] --- H1[ ]; G2[Guest] --- H1; G3[Guest] --- H1; H1 --- H2[ ]; G4[Guest] --- H2; G5[Guest] --- H2;
```

8 | © 2003, 2006 IBM Corporation

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

 z/VM 5.1

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 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 VM 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]
```

z/VM 5.1

- Maxcount* of 0 prevents dynamic definition
- 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]
            [DEVICES devs]
            [LAN owner name]
            [CHPID xx]
            [MACID xxyyzz]
```

z/VM 5.1

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
- 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?

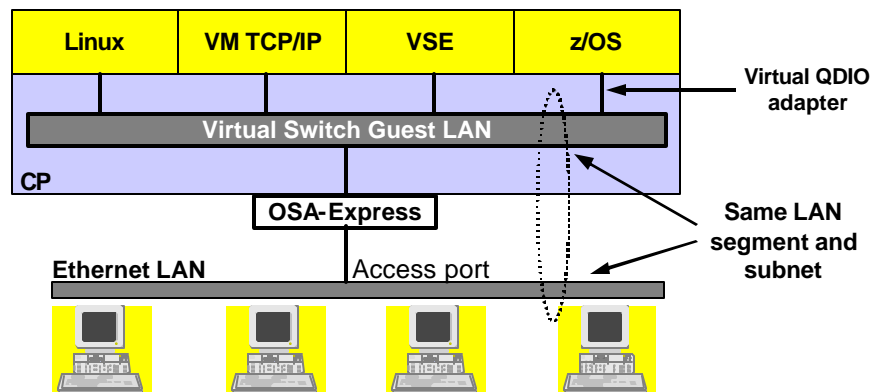


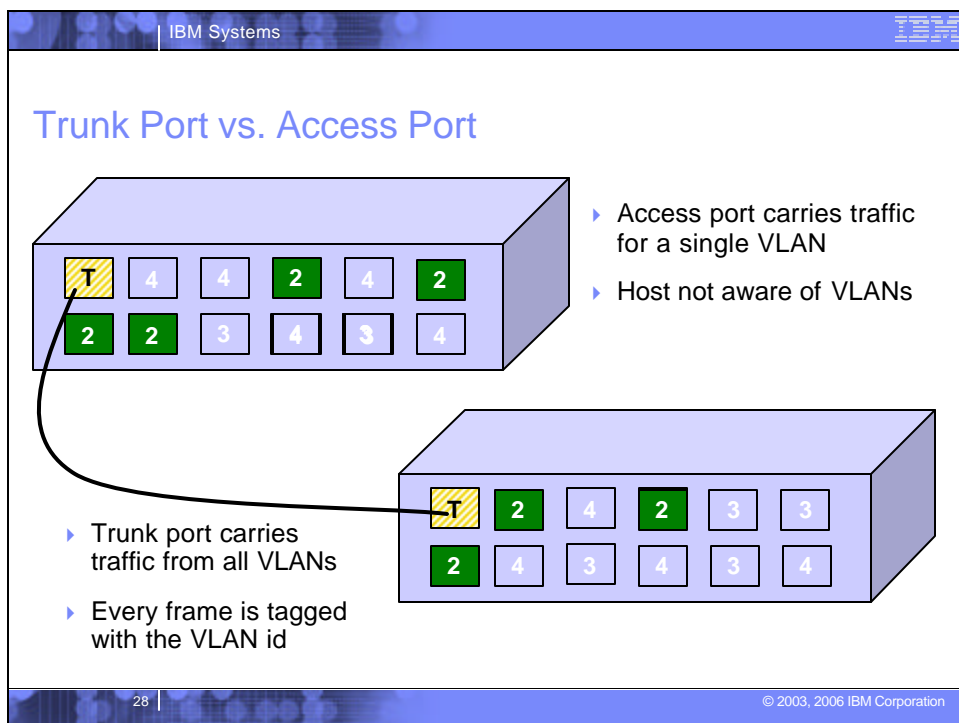
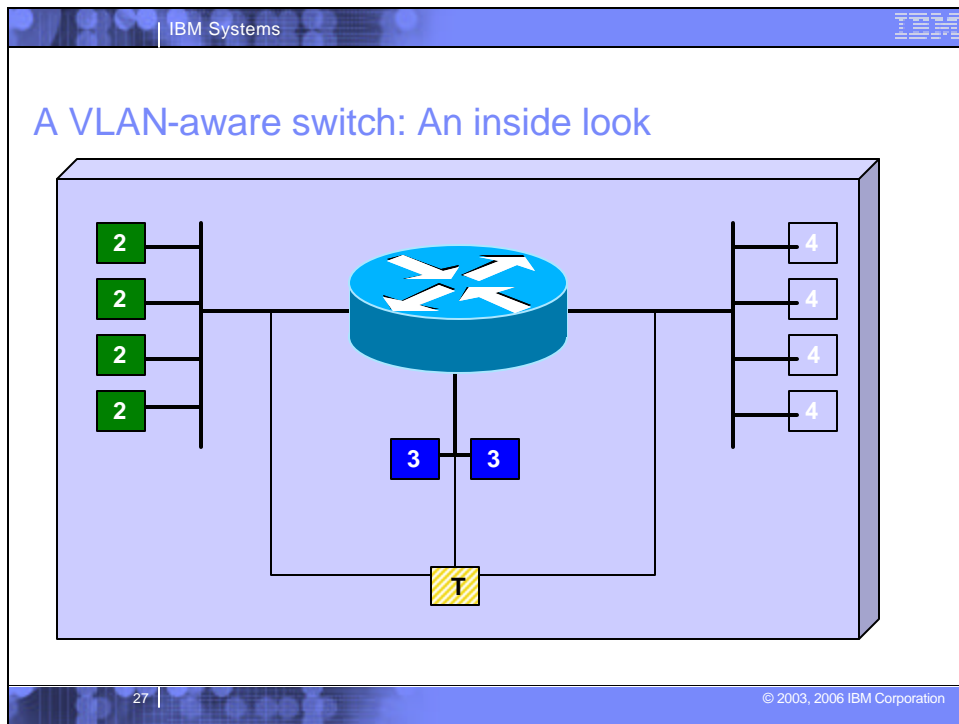
© Cisco Corp

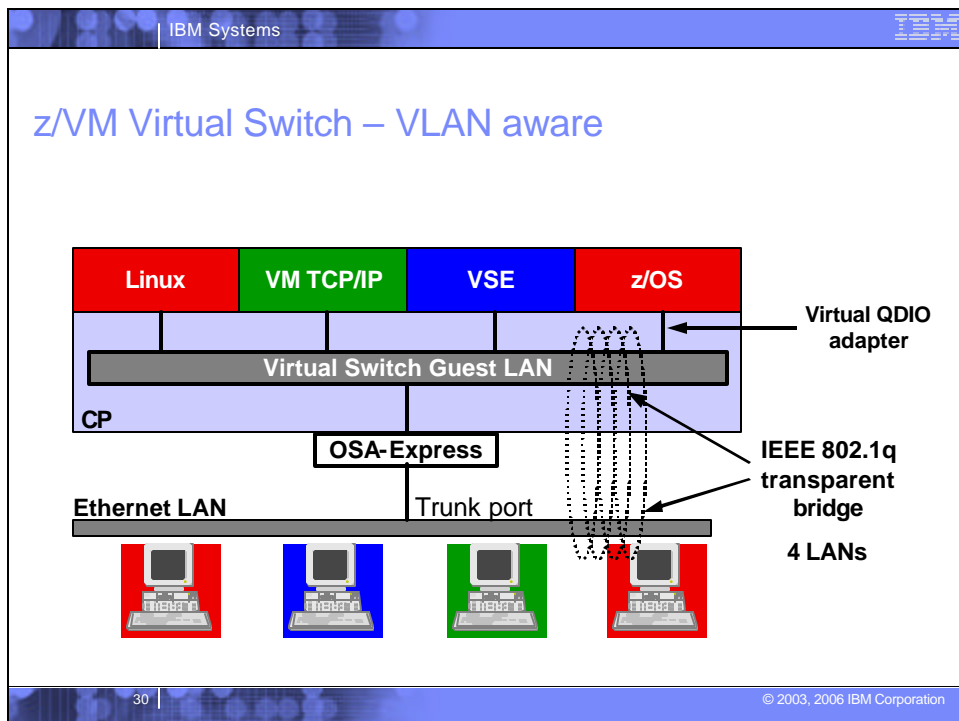
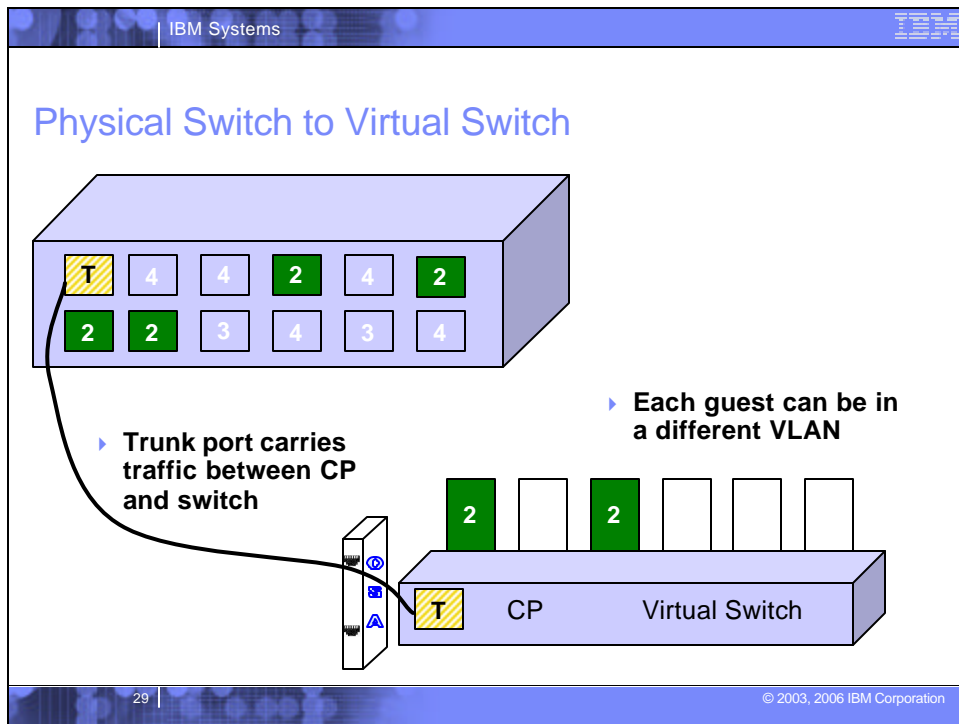
- ▶ A box that creates a LAN
- ▶ It can be remotely configured
 - ▶ E.g. Turn ports on and off
- ▶ Similar to a home router



z/VM Virtual Switch – VLAN unaware








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]
    [IP IPTIMEOUT 5 NONROUTER | ETHERNET]
    [VLAN UNAWARE | VLAN native_vid]
    [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:

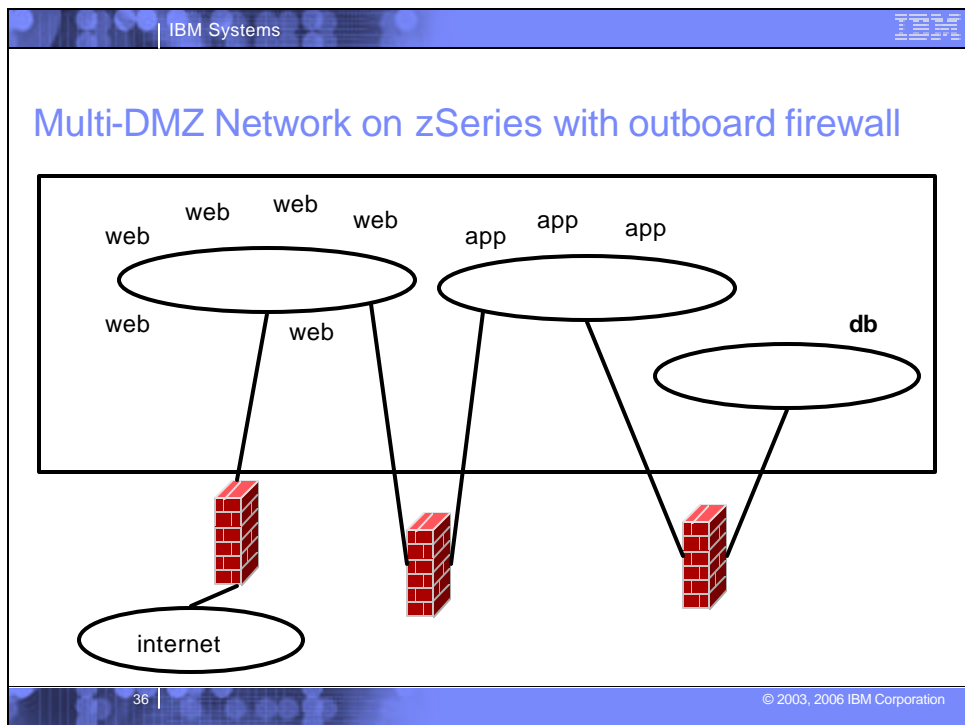
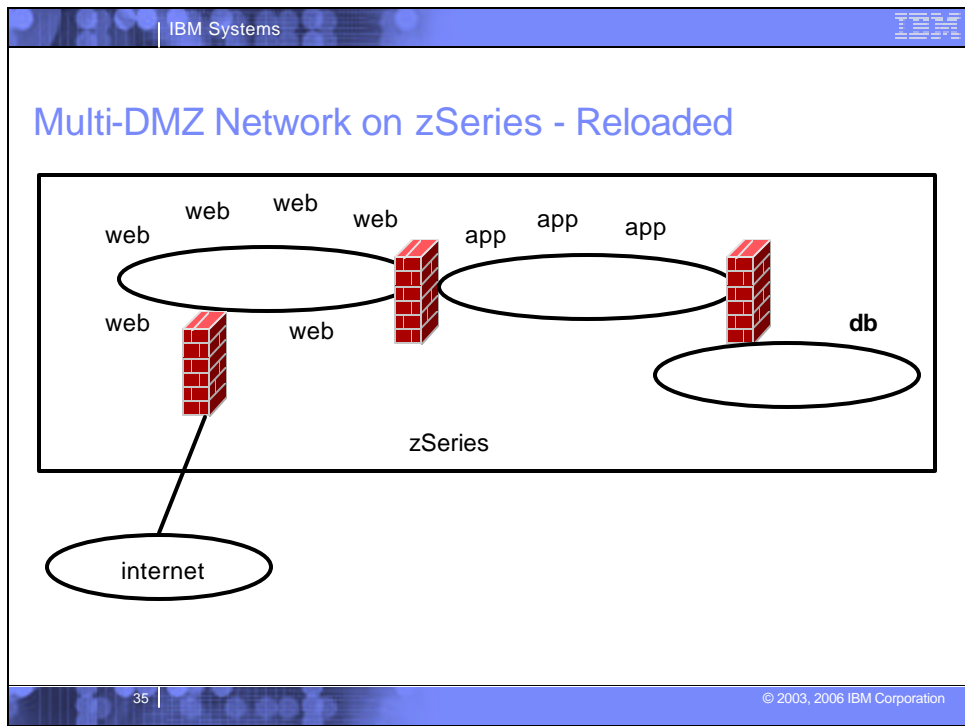
```

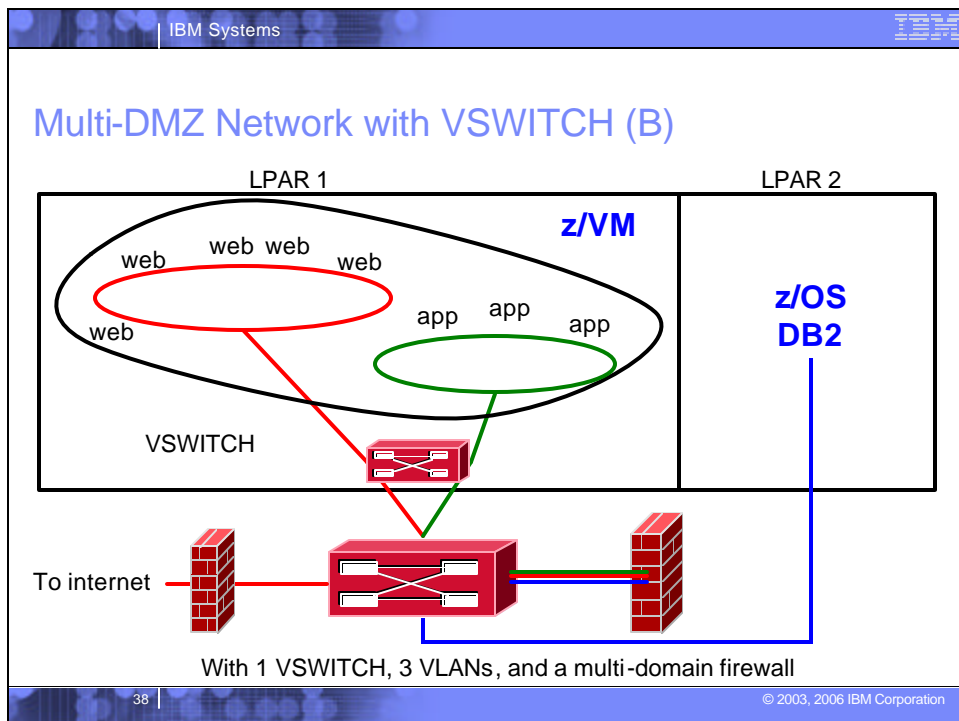
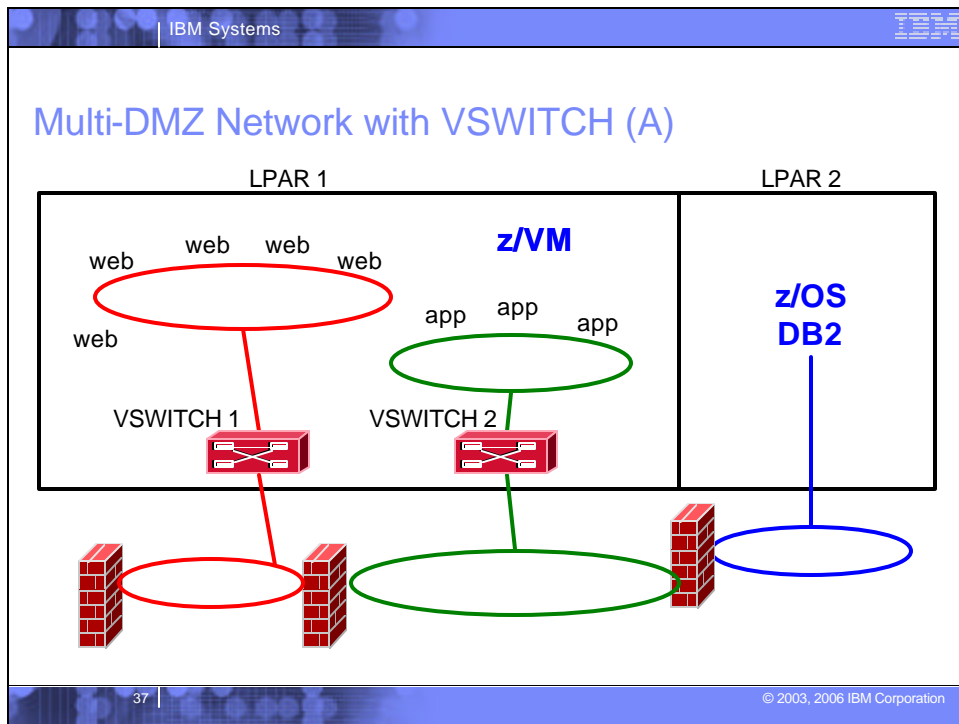
MODIFY VSWITCH SWITCH12 GRANT LNX01 VLAN 3 7 105
CP SET VSWITCH SWITCH12 GRANT LNX02 PORTTYPE TRUNK
                                           VLAN 4-20 22-29
  
```


```

CP SET VSWITCH SWITCH12 GRANT LNX02 PRO
  
```

- z/VM 4.4 supports "VLAN ANY", but it's removed in z/VM5.1!







New in z/VM 5.1...

- 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
- 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

IBM Systems

IBM

Systems

IBM

Systems

IBM

Systems

Some Final Thoughts...

43

© 2003, 2006 IBM Corporation

IBM Systems

IBM

Systems

IBM

Systems

Guest LAN vs. Virtual Switch

<ul style="list-style-type: none"> ▪ Virtual router is required ▪ Different subnet ▪ External router awareness ▪ Guest-managed failover 	<ul style="list-style-type: none"> ▪ No virtual router ▪ Same subnet ▪ Transparent bridge ▪ CP-managed failover
---	---

44

© 2003, 2006 IBM Corporation

Network Configuration

- In general, configure a Guest LAN network like any other network
 - Subnet routing
- Use the VSWITCH whenever possible
 - Exploit IEEE VLAN if you can
- By having virtual and real configurations be the same, you can easily test network configuration before deployment with real hardware

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 V5.2	<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 controlsRemoval of VLAN ANYLayer 2 (MAC) frame transportImproved virtual switch error detection & recoveryExternal security manager access control
z/VM V4	<ul style="list-style-type: none">IPv4 Virtual Switch with IEEE VLANsIPv4 HiperSocket Guest LANIPv4 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/>
 - <http://www.vm.ibm.com/virtualnetwork/>

IBM Systems

49 | © 2003, 2006 IBM Corporation

Contact Information

- By e-mail: bolinda@us.ibm.com
- In person: USA 607.429.5469
- Mailing lists: IBMVM@listserv.uark.edu
LINUX-390@vm.marist.edu

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

IBM Systems

50 | © 2003, 2006 IBM Corporation

Thanks for Listening!