


Metropolitan VM Users Association
April 20th, 2010



TCP/IP for z/VM Update

Brian W. Hugenbruch, CISSP
zVM Development Team, IBM
Endicott, NY, USA



Trademarks

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

Not all common law marks used by IBM are listed on this page. Failure of a mark to appear does not mean that IBM does not use the mark nor does it mean that the product is not actively marketed or is not significant within its relevant market.

Those trademarks followed by ® are registered trademarks of IBM in the United States; all others are trademarks or common law marks of IBM in the United States.

For a complete list of IBM Trademarks, see www.ibm.com/legal/copytrade.shtml:

*, AS/400®, e business (logo)®, DBE, ESCO, eServer, FICON, IBM®, IBM (logo)®, iSeries®, MVS, OS/390®, pSeries®, RS/6000®, S/30, VM/ESA®, VSE/ESA, WebSphere®, xSeries®, z/OS®, zSeries®, z/VM®, System i, System i5, System p, System p5, System x, System z, System z9®, BladeCenter®

The following are trademarks or registered trademarks of other companies.

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.

Cell Broadband Engine is a trademark of Sony Computer Entertainment, Inc. in the United States, other countries, or both and is used under license therefrom.

Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

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

Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

ITIL is a registered trademark, and a registered community trademark of the Office of Government Commerce, and is registered in the U.S. Patent and Trademark Office.

IT Infrastructure Library is a registered trademark of the Central Computer and Telecommunications Agency, which is now part of the Office of Government Commerce.

* 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

- General IPv6 Support
- Level 530 Enhancements
- Level 540 Enhancements
- Level 610 Enhancements
- Service Strategy

IPv6 support currently in z/VM

- CP support for IPv6
 - QDIO and HIPERSOCKETS Guest LANs support IPv6
 - Layer 2 VSWITCH supports IPv6

- TCP/IP support for IPv6
 - HiperSockets (QDIOIP) and OSA-Express (QDIOETHERNET) devices
 - Dynamic routing with MPRROUTE
 - Static routing and IPv6 Router Advertisements
 - IFCONFIG, IPWIZARD, NETSTAT, PING, TRACERTE, and **TELNET (new in level 540)**
 - Failover and Virtual IP address (VIPA) support

Function: IPv6

- Moving toward support for IPv6 networks
 - Address constraint relief
 - Auto-configuration
 - Other improvements

- Support for IPv6 networks connected through OSA Express (QDIO) adapter
 - Static routing
 - Router Advertisements
 - TRACERTE**, **PING**, and **IFCONFIG** support
 - IPv6 sockets through Language Environment and OpenExtensions Callable Services

Function: IPv6 ...

- v4 and v6 networks treated separately
 - Separate **HOME** lists, filters (**BLOCK** statement) address translation tables, static routing tables (**GATEWAY** statement), **PORT** lists

 - No routing between networks

 - New **DEVICE OSD** statement options
 - **IPv6PriRouter**
 - **IPv6SecRouter**
 - **IPv6NonRouter**

 - New **LINK QDIOEthernet** statement options
 - **EnableIPv6**
 - **DupAddrXmits**

Function: IPv6 ...

- New *RouterAdv* statement
 - Defines characteristics of router advertisements for a link

- New *RouterAdvPrefix* statement
 - Defines address prefix to be used for link router advertisements and associated on-link determination, autonomous, and lifetime characteristics

- New *AssortedParms* statement options
 - ***IgnoreIPv6Redirect***
 - ***EqualCostIPv6MultiPath****

Function: IPv6 ...

- New *NCBPoolSize* statement
 - Defines size of IPv6 Neighbor Control Block pool

- New *ICMPErrLimit* statement
 - Define maximum rate per second of IPv6 ICMP error packets transmitted on a link

- New *Neighbor* and *DelNeighbor* functions of *NETSTAT*
 - Display/delete neighbor cache entries

- *NETSTAT DEVLINKS* reports
 - Maximum frame size (Hipersockets links)
 - MTU size
 - IPv6 status
 - Multicast addresses

TCP/IP Level 530 New Function

- LDAP Server and Client
- IP Takeover (IPv4 and IPv6)
- Delete Device and Link
- SSL upgrade and TLS support
- SNMP for Virtual Switches
- MPRROUTE V1R8
- RouteD and BootP discontinued

530 Function: LDAP

- Lightweight Directory Access Protocol (RFC 2251)
- Solves a problem: the ability to have RACF be a central repository for your z/VM and Linux passwords
- Standard way for a client to retrieve data stored in a Directory Information Tree (DIT)
- z/OS 1.8 IBM Tivoli Directory Server (ITDS)

530 Function: LDAP Server and Client

- LDAP Server provides:
 - Multiple concurrent database instances (referred to as backends)
 - Interoperability with LDAP V2 or V3 protocol-capable clients
 - LDAP Version 2 and Version 3 protocol support
 - Native authentication using Challenge-Response Authentication Method (CRAM-MD5), DIGEST-MD5
 - Authentication, and Simple (unencrypted) authentication
 - Root DSE information master/slave and peer-to-peer replication

530 Function: LDAP Server and Client

- LDAP Server also provides:
 - The ability to refer clients to additional directory servers
 - The capability to create an alias entry in the directory to point to another entry in the directory
 - Access controls on directory information
 - Change logging
 - Schema publication and update
 - SSL communication (SSL V3 and TLS V1)
 - Client and server authentication using SSL/TLS*

530 Function: LDAP Server and Client

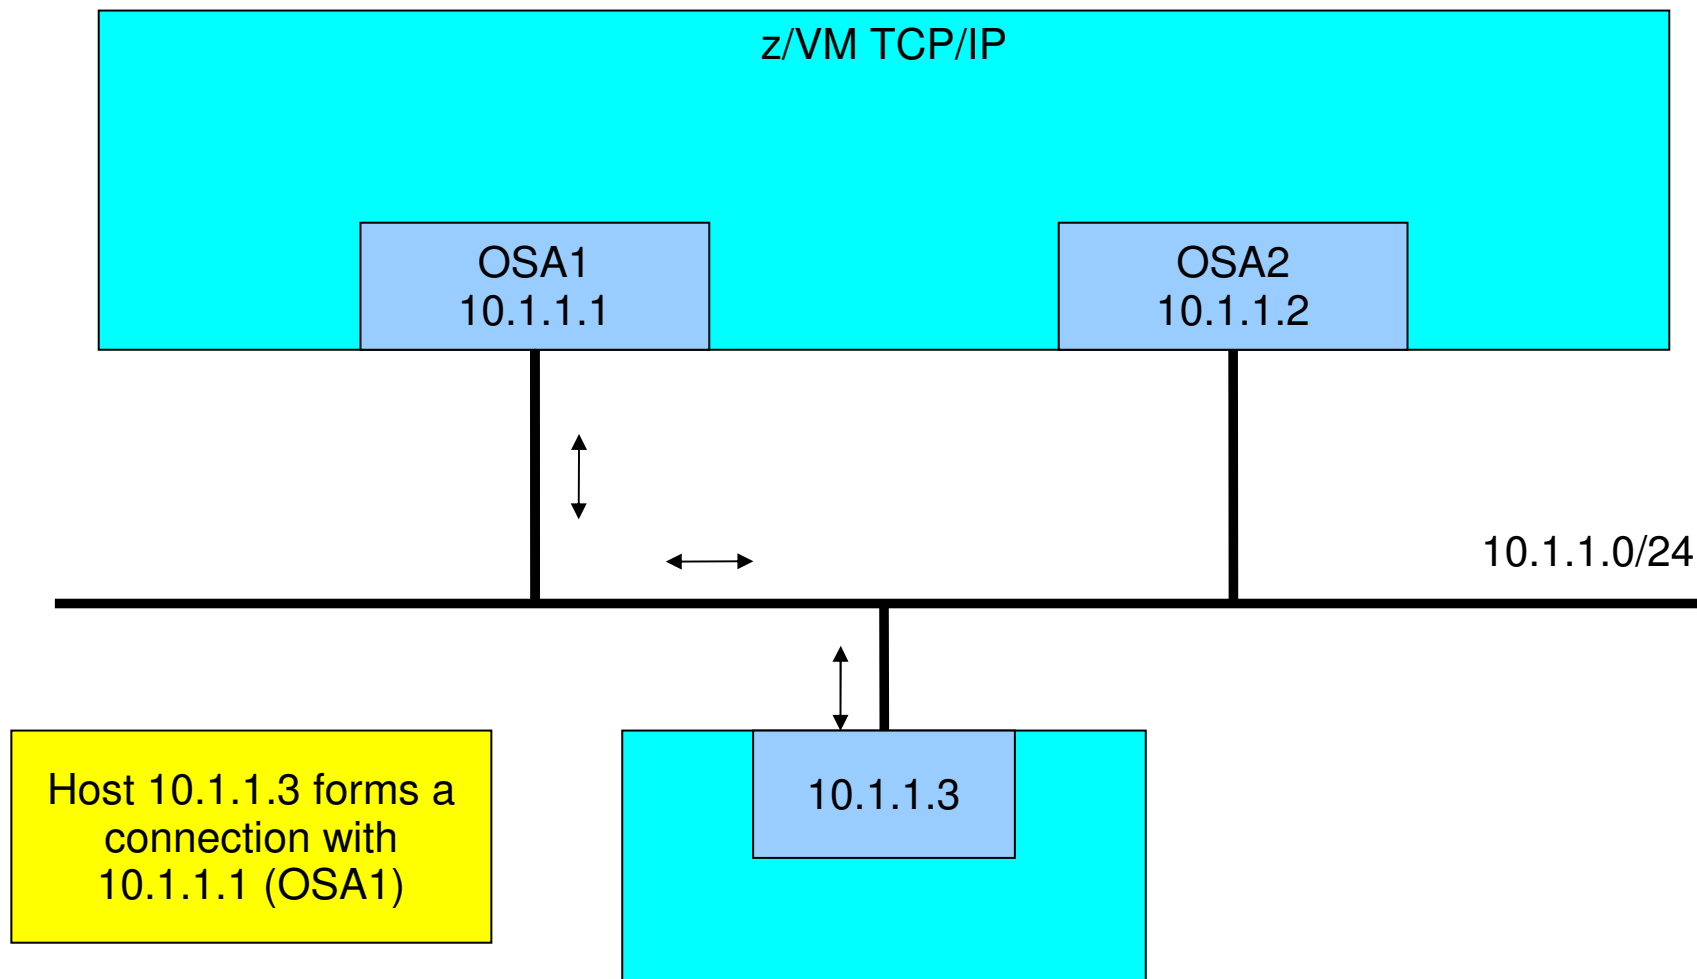
- LDAP client utilities provides a way to add, modify, search, and delete entries in any server that accepts LDAP protocol requests.

Interface High Availability – IP Takeover

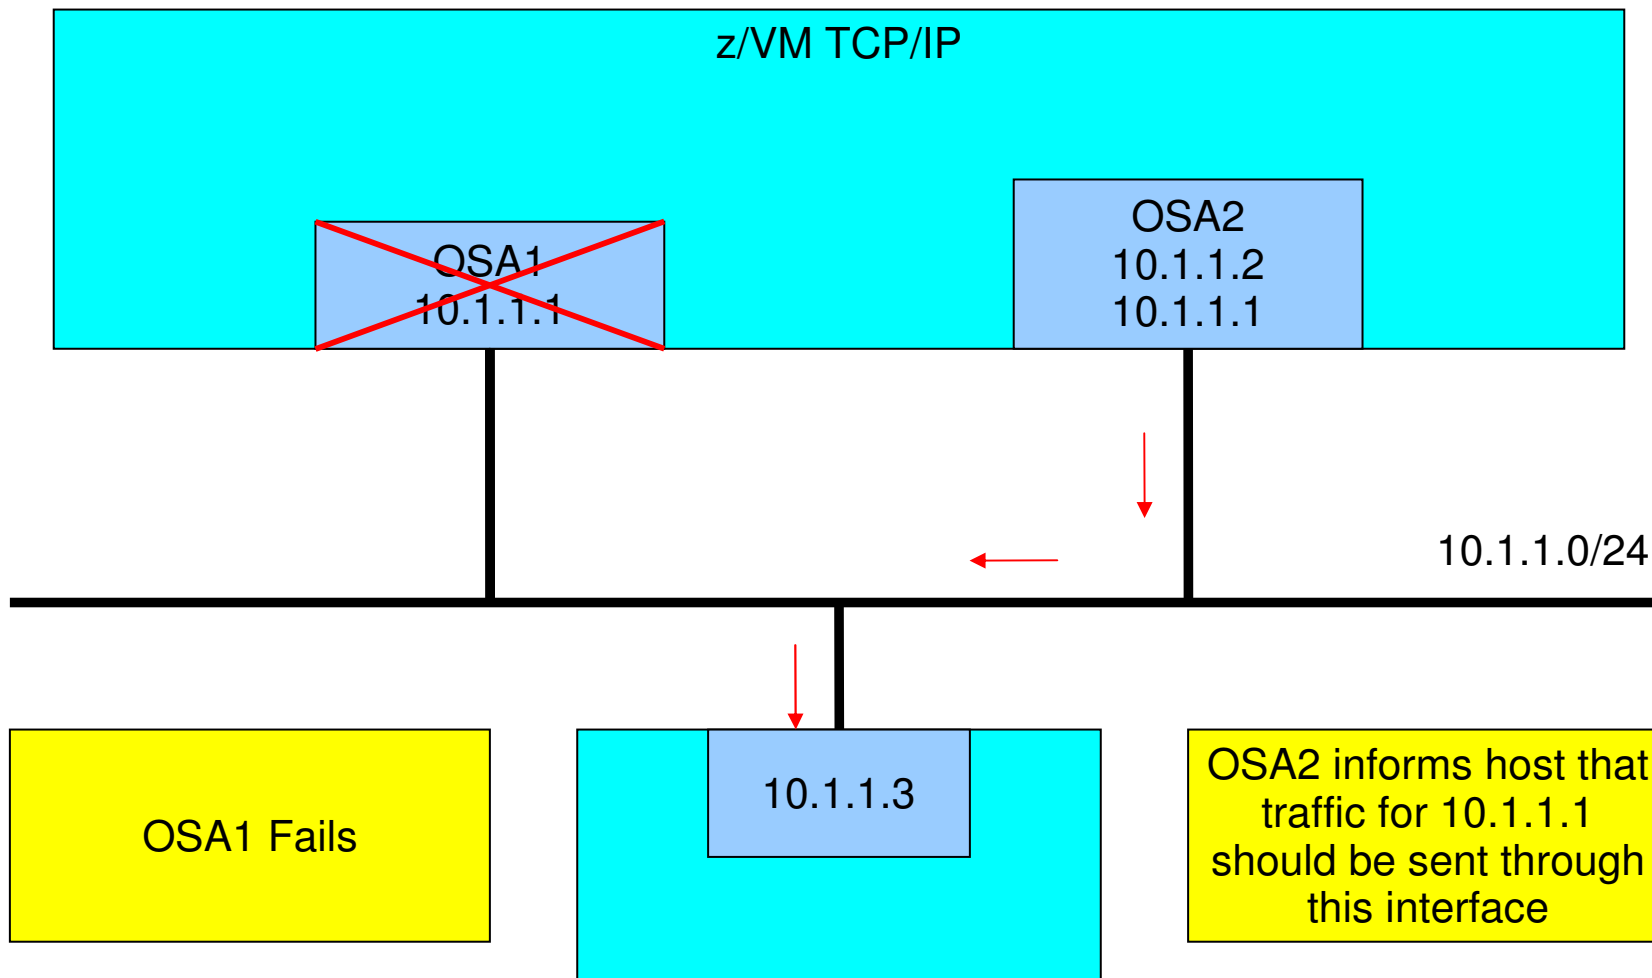
- IP takeover is supported to minimize the impact of an hardware interface failure
 - QDIO ethernet and LCS ethernet devices only

- No special parameters or options necessary
 - If the TCP/IP stack determines two interfaces are on the same network, IP takeover will be enabled for those interfaces
 - For IPv4, determination is based on the IP addresses and subnet masks of the interfaces
 - Subnet masks may be defined on the HOME statement, the GATEWAY statement, or in the MPROUTE CONFIG file


IP Takeover Details



IP Takeover Details (cont.)



IP Takeover Details (cont.)



530 Function: Delete Device and Link

- Device and Link statements can now be dynamically removed from the z/VM TCP/IP stack.

- New -Remove option for IFCONFIG

`IFCONFIG -REMOVE`

- New SIOCDINTERFACE subcommand for REXX and C

530 Function: SSL upgrade

- Support for

- Novell(R) SUSE(R) Linux Enterprise Server (SLES) 9 Service Pack 3 (64-bit)
- Novell SUSE Linux Enterprise Server (SLES) 9 Service Pack 3 (31-bit)
- Red Hat Enterprise Linux(R) (RHEL) AS 4 Update 4 (64-bit)
- Red Hat Enterprise Linux (RHEL) AS 4 Update 4 (31-bit)

530 Function: TLS Support

- Secure Sockets Layer/Transport Layer Security (SSL/TLS)
 - FTP
 - Telnet
 - SMTP

- Data Transmission can start in clear text and be converted to secure text at a later time.

530 Function: SNMP for Virtual Switches

- Management IP address for Virtual Switch
 - New HOME statement
- Generic SNMP Subagent
- Bridge MIBS for Virtual Switch reporting

530 Function: MPROUTE

- MPROUTE support upgraded to V1R8

RouteD and BootP support discontinued in 530

- MPROUTE and DHCP are available and recommended to provide the services formally performed by RouteD and BootP.

TCP/IP Level 540 New Function

- Binder Upgrade
- LDAP Upgrade
- LE Upgrade
- Support withdrawn for SNA LU.0 and X25
- Change to Device Autorestart
- IPv6 for Telnet
- TCPIP OSD Layer 2
- Path MTU discovery
- OSA Multi-Port

540 Upgrades

- The z/VM Binder has been upgraded:
 - From z/OS 1.6 to z/OS 1.9

- LDAP is has been upgraded:
 - From z/OS 1.8 to z/OS 1.10

- LE has been upgraded:
 - From z/OS 1.4 to z/OS 1.9

Support Withdrawn in 540

- SNA LU.0 and X.25 device driver support has been withdrawn from the product as of level 540.

- Support for the following has been withdrawn in the Level 540 release
 - Network Database (NDB) system
 - Trivial File Transfer Protocol (TFTP)
 - X25 (including X25IBI server) interface
 - SNALINK server

- Statement of Direction issued:
 - IBM TCPIP intends to withdraw support in a future z/VM release for the Kerberos Authentication System.

540 Function: Device Autorestart

- Autorestart is now the default on the OSD device and HiperSockets device statement

```
                .-AUTORestart---.  
>>--DEVICE--device_name--...--+-----+-----+-----+-----><  
                '-NOAUTORestart-'
```

540 Function: IPv6 support for Telnet

- The z/VM Telnet Client and Server have been updated to support IPv6 addresses. The Pascal APIs that Telnet uses have also been updated.

Telnet Command

- TELNET command will recognize and accept IPv6 addresses when entered directly on the command line
 - `TELNET 50C6:A2C2::14:5E00:9B7:127E`
- If a hostname is entered, IPv6 addresses will be considered in the resolution of the name by default. To specify the type of address returned, use the new ADDRTYPE option
 - `TELNET gdlect2 (addrtype ipv6`
- TELNET command does not support SECURE option for IPv6 (SSL support for IPv4 only)
- The SECURETELNETCLIENT statement provides the default client security value to use when neither the SECURE nor NOSECURE option is specified on the TELNET command. If YES is specified on this statement, IPv6 connections will fail.

IPv6 support for Telnet (Continued)

- New IPv6-capable Pascal APIs are provided for TCP connections.
 - Tcp6Open – initiates a TCP/IP version 6 connection. Control returns immediately
 - Tcp6WaitOpen – initiates a TCP/IP version 6 connection. Control will not be returned until the connection is established.
 - Tcp6Status – obtains the current status of an IP version 6 connection

IPv6 support for Telnet (Continued)

- Telnet server (internal to the TCP/IP server)
 - TELNET server supports IPv6
 - Will be enabled for IPv6 connections whenever the TCP/IP stack is enabled for IPV6 and there is at least one port reserved for INTCLIEN on the PORT statement
 - Same ports used for IPv4 and IPv6
 - PORT statement allows IPv6 addresses for INTCLIEN
 - No configuration changes needed
 - No support for secure IPv6 connections (SSL support for IPv4 only)

- The TN3270E statement now accepts IPv6 connections.

IPv6 support for Telnet (Continued)


- Exits – parameter lists have been updated
 - Telnet Session Connection Exit
 - Telnet Printer Management Exit

OSD Layer 2 support

What is Layer 2 Transport?

- Layer 2 or ETHERNET transport handles communication in the Data Link Layer
- All communication is based on the MAC address instead of the IP address.
- Data is transported and delivered in Ethernet frames.
- Before communicating, an IP address must be resolved to a MAC address. This is done using Address Resolution Protocol, or ARP.

OSI Model



OSD Layer 2 Support

- In z/VM 530 and earlier:
 - The TCP/IP stack can only operate OSA-Express OSD devices in IP mode.
 - The Virtual Switch both ETHERNET and IP transports.
 - IP transport is for IPv4 only.
 - ETHERNET supports IPv6 and non-IP based protocols

OSD Layer 2 Support

- The constraints soon became apparent ...
 - Cannot connect a z/VM TCP/IP stack to a z/VM guest LAN or VSWITCH operating in ETHERNET mode (layer 2).
 - Having Linux on one Virtual Switch and the TCP/IP stack on another or a dedicated OSA-Express connection increases the complexity of the LAN configuration.
 - The lack of support in the native VM TCP/IP stack was viewed as an inhibitor for some customers considering a move to a Layer 2 VSWITCH or new deployment on a Virtual Switch.

OSD Layer 2 support

- New Support in 540

- z/VM TCP/IP can now deploy an Ethernet connection on a real or virtual OSA Express adapter.
- Supported on both OSA-Express2, or OSA-Express3 feature when configured in Queued Direct I/O (QDIO) mode (CHPID type OSD).

OSD Layer 2 Support

- This new support allows connecting a z/VM TCP/IP stack to
 - an IPv4 or IPv6 Layer 2 Virtual Switch (VSWITCH).
 - a Layer 2 guest LAN
 - or (via a real OSA-Express) to a physical LAN segment in Layer 2 mode.

- A z/VM TCP/IP stack connected to a VSWITCH operating in Layer 2 mode is now supported in Link Aggregation configurations.
 - VSWITCH Link Aggregation is designed to provide both increased bandwidth and more continuous network connectivity

What's needed for Layer 2 support?

- MAC address
- ARP support
- IPv6

MAC Address

- For an Ethernet connection, the stack now deploys a unique Media Access Control (MAC) address, defined locally by CP, instead of sharing the OSA MAC address.
- When a device is started CP will return (via DIAGNOSE x'26C' subcode x'30') a locally defined MAC for a real OSA adapter or virtual NIC.

Mac Address

- A MAC address is assigned by concatenating the VMLAN MACPREFIX (as defined by the System Configuration file) with the MACID found first in the following priority list:
 1. The MAC ID from a SET NIC MACID command issued for the device,
 2. The MAC ID specified on a NICDEF directory control statement defining the device, which must be in the USER subset of the VMLAN MACIDRANGE SYSTEM range
 3. An available MACID from the SYSTEM section of the VMLAN MACIDRANGE

ARP Support

- Since packets will be sent based on MAC addresses, native ARP (Address Resolution Protocol) support will be provided for OSD devices in Layer 2 mode. TCP/IP will maintain its own ARP cache.
 - Supports both Ethernet (DIX) and 802.3 frames

- When a device is started, gratuitous ARPs are sent out to notify other systems of our MAC address.
 - Send 2 Ethernet, 1 802.3, then 2 Ethernet

ARP Support

- When sending to an unknown MAC address ARP requests are sent to determine MAC address
 - Send both Ethernet and 802.3 formats: 5 Ethernet, 5 802.3, 5 Ethernet
 - Accept both Ethernet and 802.3 responses
 - Prefer Ethernet responses

IPv6 Layer 2

- For IPv6 enabled interfaces, an Interface Identifier and a link-local address for the layer 2 link will be created (see RFC 2464: “Transmission of IPv6 Packets over Ethernet Networks”).

- The Interface Identifier is based on the MAC address.
 - MAC address : 02-6F-5C-60-00-42
 - Interface Identifier : 00-6F-5C-FF-FE-60-00-42

- The IPv6 link-local address for an Ethernet interface is formed by appending the Interface Identifier, as defined above, to the prefix FE80::/64.
 - link-local address : FE80::006F:5CFF:FE60:0042

Configuring Layer 2

Several commands were updated for Layer 2 configuration:

1. IFCONFIG

- A new Interface Operand was added to the IFCONFIG command. The new operand will be IP or ETHERNET.
 - To define an interface to an Ethernet Open Systems Adapter (OSA-2) at address 3000–3003 that will operate in layer 2 mode, enter the following:

```
ifconfig eth1 10.11.12.26 eth 3000 ethernet mtu 1500
```

Configuring Layer 2 (continued)

2. IPWIZARD

- The panels have been updated to allow users to choose the transport type for QDIO interfaces.

3. LINK STATEMENT

- Transport type (IP or ETHERNET) was added to the QDIOETHERNET LINK statement for OSD devices
 - **LINK DEV45A0 QDIOETHERNET DEV45A0 ETHERNET**

Configuring Layer 2 (continued)

- Configuring the MAC address
 - CP will automatically generate a unique local MAC address based on the MAC settings
 - Settings are configured via VMLAN system configuration statement.
 - MACPREFIX – default is 02-00-00
 - MACIDRANGE

Configuring Layer 2 MAC Address (continued)

–If you need to have specific or persistent MAC addresses, then you can configure the MACID:

- Set the MACID via SET NIC command

```
CP SET NIC 1E35 MACID USER 50FF00
```

MAC address 02-6F-5C-50-FF-00 has been assigned to device 1E35

- NICDEF directory control statement

```
NICDEF 1E35 TYPE QDIO LAN SYSTEM PEGL2 MACID 50FF00
```

Verifying Layer 2

- NETSTAT DEVLINKS output will indicate the transport type and the local MAC address

```

Device DEV45A0                Type: OSD                Status: Ready
Queue size: 0                 CPU: 0                   Address: 45A0            Port name:
UNASSIGNED
  Link DEV45A0                Type: QDIOETHERNET      Port number: 0
  Transport Type: Ethernet MAC: 02-6F-5C-60-00-54
  Speed: 1000000000
  BytesIn: 2000                BytesOut: 1826
  Forwarding: Enabled          MTU: 8992                IPv6: Enabled
  IPv4 Path MTU Discovery: Disabled
  IPv4 VIPA ARP
  IPv6 VIPA ND
  Multicast Group              Members
  -----
  224.0.0.1                    1
  FF02::1:FF60:54              2
  FF02::1                       1

```


Verifying Layer 2 (continued)

- IFCONFIG output will display the transport type and the local MAC address

```
ifconfig dev45a0
```

```
DEV45A0 inet addr: 11.11.11.45 mask: 255.255.255.0
```

```
UP MTU: 8992
```

```
vdev: 45A0 type: QDIO ETHERNET portname: UNASSIGNED
```

```
transport type: ETHERNET MAC address: 02-6F-5C-60-00-54
```

```
ipv6: ENABLED
```

```
LAN owner: SYSTEM name: PEG L2
```

```
cpu: 0 forwarding: ENABLED ipv4 path MTU discovery: DISABLED
```

```
RX bytes: 2000 TX bytes: 1826
```

```
IPv6 Addresses:
```

```
50C0:C2C1:0:0:6F:5CFF:FE60:54
```

```
FE80:0:0:0:6F:5CFF:FE60:54
```

540 Function: Path MTU Discovery

- The MTU may be configured in one (or more) of three places
 - Link statement in the TCP/IP configuration file
 - Gateway statement in the TCP/IP configuration file
 - INTERFACE, RIP_INTERFACE, or OSPF_INTERFACE statement in the MPRoute configuration file

- The LINK statement is the preferred method of configuring MTU values

What is Path MTU Discovery?

- **A process through which the TCP/IP server determines the optimal MTU value to use when sending data over a particular path**
- **The MTU to be used should be the smallest MTU of all the networks traversed on the path**
 - For IPv4, this improves network performance by causing packet fragmentation to occur at the source rather than at intermediate routers
 - For IPv6, fragmentation can only be performed at the source node. Nodes that do not implement Path MTU Discovery should only send packets at the IPv6 minimum MTU (1280). This support allows the stack to send larger packets

How Does Path MTU Discovery Work?

- **The TCP/IP server sends packets using the initial MTU value**
 - With the “Don't Fragment” bit set in the IP header for IPv4
- **Intermediate routers that cannot accommodate the packet size respond with an ICMP error message which (hopefully) includes the MTU of the network that could not handle the packet**
- **The TCP/IP server updates the MTU value for that destination and resends the packet**
 - For non-host routes, a host route will be created for the destination with the new MTU value
- **The process repeats as necessary**

How is Path MTU Discovery Configured?

- **PATHMTUAGE** statement determines how many minutes Path MTU Discovery information is kept before expiring
- **IPv4**
 - Path MTU Discovery is disabled on all interfaces by default
 - PATHMTU option on the ASSORTEDPARMS statement
 - When specified, this enables Path MTU Discovery on all interfaces
 - PATHMTU/NOPATHMTU options on the LINK statement
 - These options can be used to override the ASSORTEDPARMS statement for a particular interface
- **IPv6**
 - Path MTU Discovery is always enabled for IPv6

OSA Multi-port Support

- New hardware features with the OSA-Express3 10 GbE and OSA-Express3 GbE SX or LX
- Multiple ports are now available on the hardware and can be specified in the TCP/IP stack configuration.
- Virtual Switch also exploits multi-port support.

OSA Multi-Port Support (continued)

- OSD Device statement updated to specify port number
- Format:
 - **PORTNUMBER pp**
 - pp is the port number
 - If a port number is not specified, the default is port 0.
- Example:
 - **DEVICE DEV1300 OSD 1300 PORTNUMBER 1**

OSA Multi-Port Support (continued)

- **NETSTAT DEVLINKS** updated with Port Number
–replaced Net Number for OSD devices
- **Example output:**

```

VM TCP/IP Netstat Level 540    TCP/IP Server Name: TCPIP0D

Device DEV0D00                Type: OSD                Status: Ready
Queue size: 0  CPU: 0  Address: 9C09  Port name: UNASSIGNED
IPv4 Router Type: NonRouter  Arp Query Support: Yes
Link TONETA                    Type: QDIOETHERNET  Port number: 1
  Transport Type: IP
  Speed: 1000000000
  BytesIn: 261  BytesOut: 529
  Forwarding: Enabled  MTU: 8992  IPv6: Disabled
  IPv4 Path MTU Discovery: Disabled
  Broadcast Capability: Yes
  Multicast Capability: Yes
  IPv4 VIPA ARP
  Multicast Group              Members
  -----
  224.0.0.9                    1
  224.0.0.1                    1

```

OSA Multi-Port Support (continued)

- Other externals updated with port number:
 - IFCONFIG
 - TCP/IP messages

TCP/IP Level 610 New Function

- CMS-based SSL Server
- QDIO Data Connection Isolation

CMS-based SSL Server

- Move from a Linux based solution to a CMS based solution

- First delivered for z/VM 540 post-GA through APAR PK65850 and associated service

CMS-based SSL Server

- Encryption and decryption means replaced
 - Port of z/OS System SSL v1.10
 - No FIPS 140-2 support

- Supports SSL 3.0 and TLS 1.0, along with new cipher suites (AES)

- Stores certificates in a BFS directory
 - .kdb files compatible with z/OS certificate databases from System SSL.
 - New GSKADMIN virtual machine preconfigured with appropriate permissions and authorities

CMS-based SSL Server

- Console-based tracing of server activity

- Certificate management now handled by GSKKYMANT
 - Import certificates and keys, manage extensions

- Can be your own Certificate Authority

QDIO Data Connection Isolation

- Port isolation security for the OSA-Express
- Restrict communication inside a VSWITCH, and between other partitions sharing that OSA-Express
- Helpful for designing virtual networks with strict traffic-separation policies

QDIO Data Connection Isolation

- OSA-Express2, OSA-Express3
 - z10 BC and z10 EC
 - Exclusive to CHPID type OSD
-
- Also available for z/VM V5.3 and V5.4 with the PTFs for APARs VM64463 and PK67610.

Recommended Service Strategy

- Apply the latest RSU

- Visit the TCP/IP for z/VM Feature home page for late-breaking service news

Summary

- TCP/IP for VM is alive and well
- Level 530 delivered major advances
- Levels 540 and 610 contain exciting advancements
- We still have more to do
 - Anticipate where most z/VM TCP/IP customers are going next
 - Your requirements are important to us

Contact Information

- Presenter: Brian W. Hugenbruch
 - E-mail: bwhugen@us.ibm.com
 - Phone: USA 607.429.3660

- Web sites:
 - <http://www.vm.ibm.com/networking/> -- zVM Virtual Networking
 - <http://www.vm.ibm.com/networking/ipv6/> -- IPv6 support in z/VM
 - <http://www.vm.ibm.com/related/tcpip/> -- zVM TCPIP
 - <http://ibm.com/vm/techinfo/listserv.html> -- List of Mailing Lists

- Feature Page
 - <http://www.rfc-editor.org/rfc.html>
 - <http://www.redbooks.ibm.com/>

- Mailing Lists:
 - IBMTCP-L@VM.MARIST.EDU
 - IBMVM@LISTSERV.UARK.EDU
 - LINUX-390@VM.MARIST.EDU

- Special Thanks to: Tracy Adams, Miguel Delapaz, Alan Altmark