



What's New in z/OS Communications Server?

Enterprise Networking Solutions, Raleigh

Sam Reynolds - samr@us.ibm.com

Mac Devine - wdevine@us.ibm.com

Alfred B Christensen - alfredch@us.ibm.com

© Copyright International Business Machines Corporation 2004. All rights reserved.



The logo for IBM z/OS Communications Server, featuring a red '@' symbol followed by the word 'server' in a black, lowercase, sans-serif font.

Agenda

z/OS V1R5 Communications Server


- **Overview**
- **Sysplex**
- **Applications**
- **Security**
- **Other TCP/IP Enhancements**
- **Enterprise Extender and SNA**
- **IPv6 Support**

z/OS Communications Server Disclaimer

- Plans for the z/OS Communications server are subject to change prior to general availability
- Information provided in this presentation may not reflect what is actually shipped for z/OS Communications Server
- This presentation includes an early overview of future z/OS Communications Server enhancements



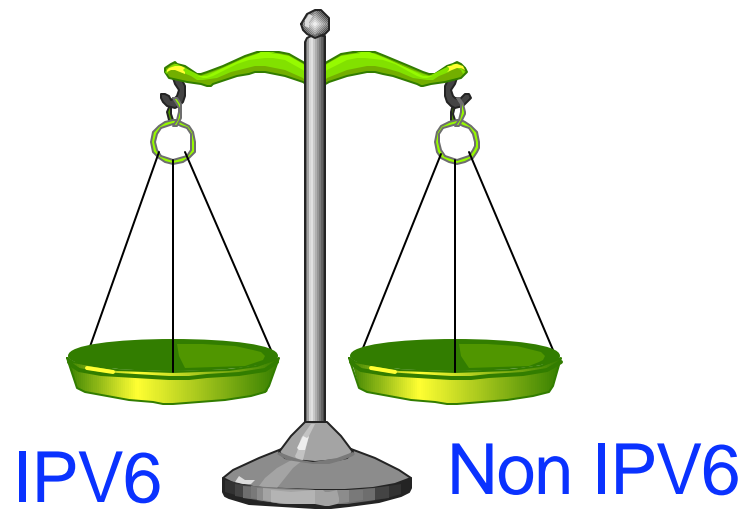
Note: Plans are subject to change!



z/OS V1R5 Communications Server

z/OS V1R5 continues the effort started in z/OS V1R4 to provide IPV6 on z/OS and provides a lot of Non-IPV6 function also!

★ IPV6 support is technology driven in preparation for future



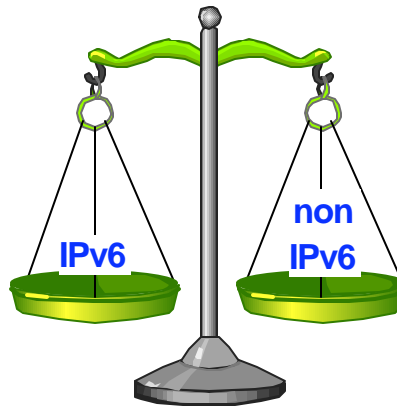
V1R5 - Continue to provide a balance of IPv6 and non-IPv6 functions

IPv6 focus in V1R5:

- ✓ Applications
- ✓ APIs
- ✓ DLCs
- ✓ Network management
- ✓ Policy
- ✓ Routing
- ✓ Security
- ✓ Enterprise Extender

non-IPv6 focus in V1R5:

- ✓ TN3270 server enhancements
- ✓ FTP enhancements
- ✓ Sysplex enhancements
- ✓ Network management
- ✓ Policy
- ✓ Security
- ✓ Msys
- ✓ SNA enhancements
- ✓ Enhancements for new hardware
- ✓ Routing enhancements
- ✓ Mail



Sysplex Enhancements

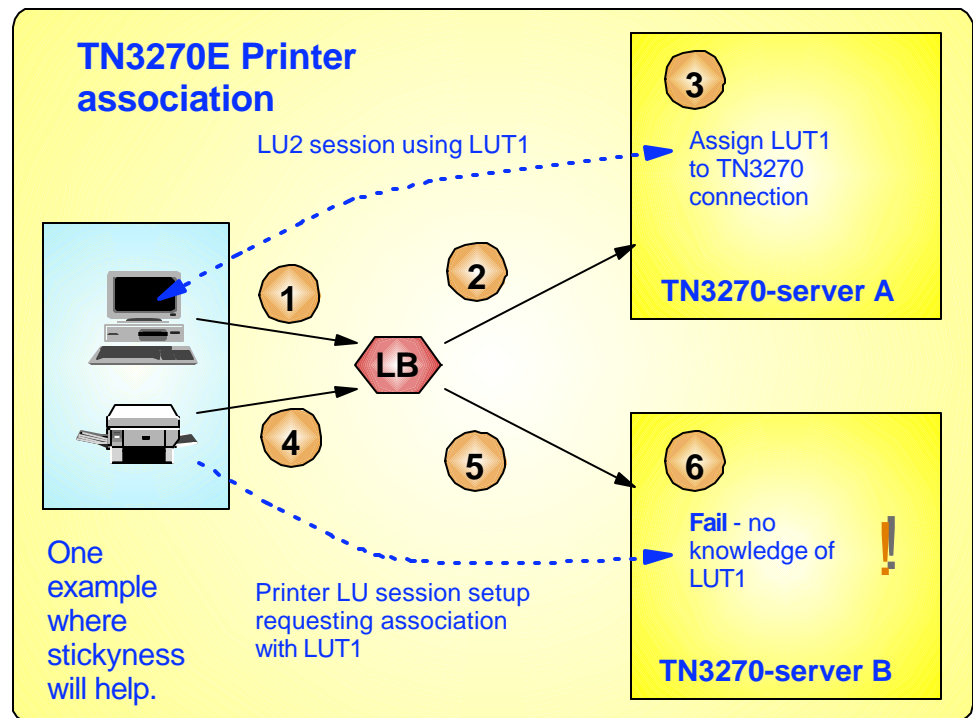
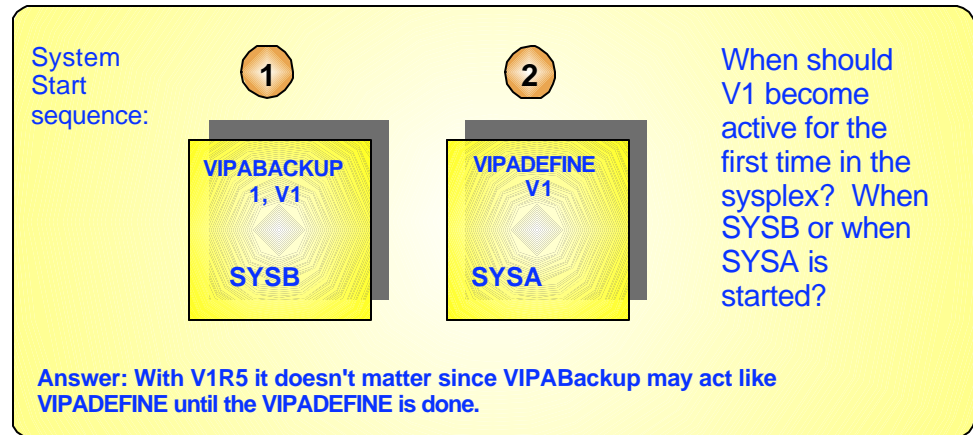
Copyright International Business Machines Corporation 2004. All rights reserved.



Sysplex enhancements in z/OS V1R5

Sysplex Enhancements

- ▶ Increase ports on VIPADISTRIBUTE from 4 to 64 (PTFed back to z/OS V1R2 - APAR PQ65205)
- ▶ Dynamic port definition for VIPADISTRIBUTE dynamic VIPA when server binds to dynamic VIPA
- ▶ Increase limit of DVIPAs per stack from 256 to 1024
- ▶ Support DVIPA activation based on VIPABACKUP before VIPADEFINE ever processed
- ▶ Sysplex Distributor affinity
 - Configurable timer-based stickyness per source IP address, server DVIPA and port
- ▶ New round-robin distribution method in Sysplex Distributor (PTFed back to z/OS V1R4 - APAR PQ76866)
 - Alternative to WLM-based distribution
 - Useful where availability is more important than capacity




Application Enhancements

Copyright International Business Machines Corporation 2004. All rights reserved.

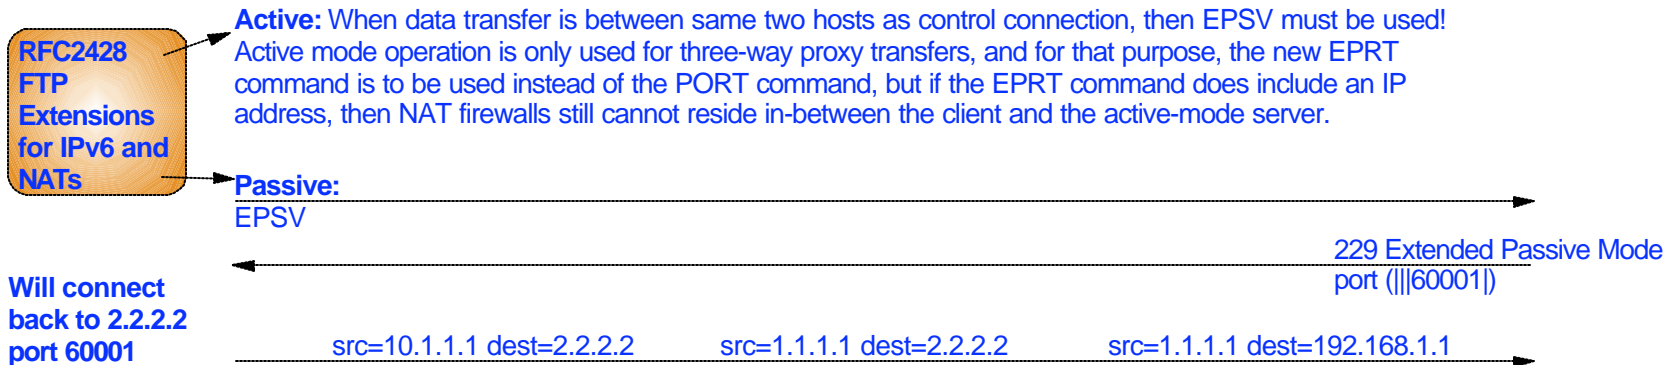
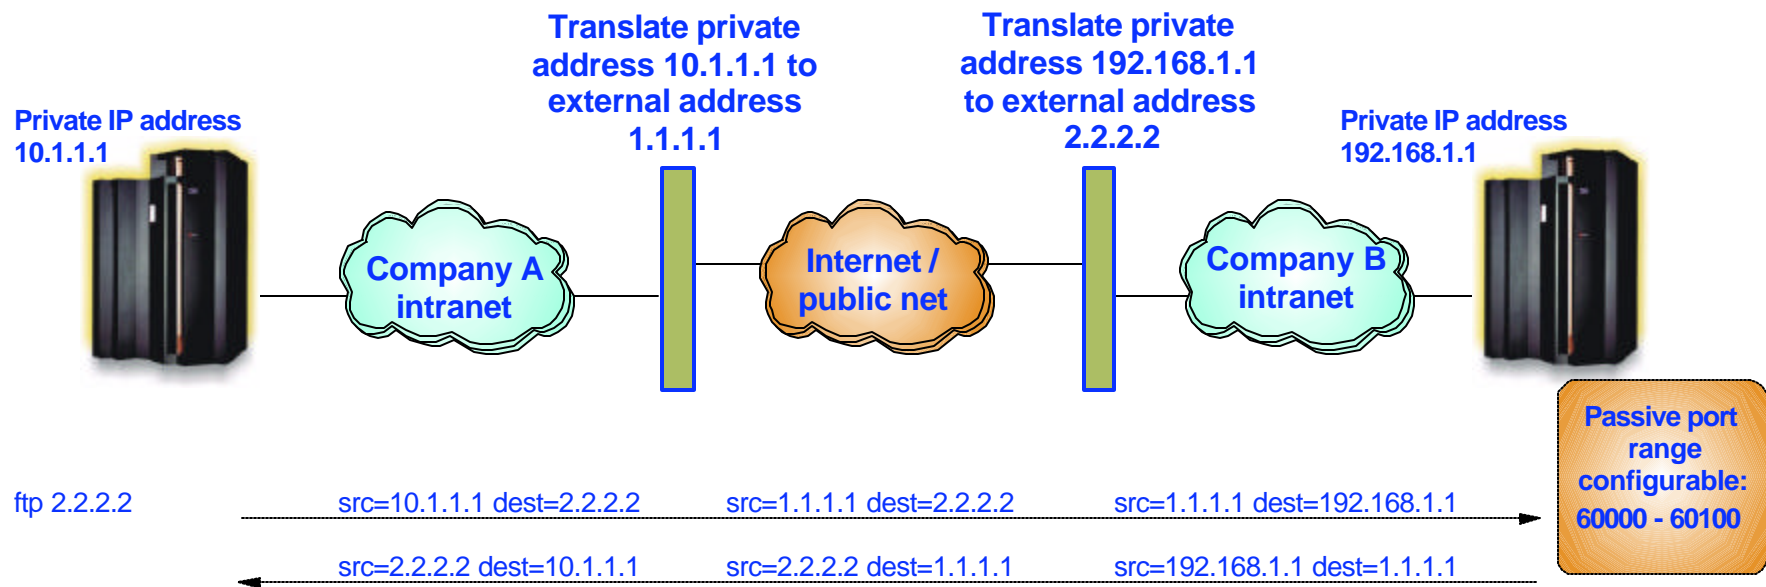


FTP enhancements in z/OS V1R5

- Support allocation of target PDS(E) like existing source PDS(E)
- Enable user to specify if a PDS or a PDSE should be created
- Allow configuration control of passive data port range on FTP server
- Support use of extended passive (EPSV) and active mode (EPRT) for NAT firewall relief
- Provide consistent error codes from FTP client
- Enable MVS syslog message when batch FTP client operation fails (for automation purposes)
- Allow SSL/TLS login to FTP server without password when client authentication is used
- Deliver FTP server load module as RMODE=ANY



Secure FTP and Network Address Translation (NAT) firewalls in z/OS V1R5



Client EPSV support and server passive data port range configuration
PTF'ed back to z/OS V1R4 - APAR PQ80281


TN3270 server enhancements in z/OS V1R5

- **Improve Tn3270/e connection recovery**
 - ▶ End user need not know LUNAME to recover session

- **Provide SNMP network management data for tn3270/e**
 - ▶ Support in new Enterprise-specific TN3270 MIB and new TN3270 subagent
 - ▶ Focus on connection performance data
 - ▶ Provides end-to-end stats over Tcp/Ip and SNA for connections
 - ▶ Data accessible via MIB or D TCPIP, TELNET, CONN, CONN=connid detail command

- **Enhance usability with the following items:**
 - ▶ Support range of ipaddress on IPGroup in profile
 - ▶ Better compatibility between type ahead-capable clients and applications
 - UNLOCKKEYBOARD statement in PROFILE.TCPIP

- **Performance enhancement for definite Rsp. sessions**
 - ▶ Server turns off delay acks for clients that have negotiated it. Client can do many small sends without waiting for acks from server.



Security

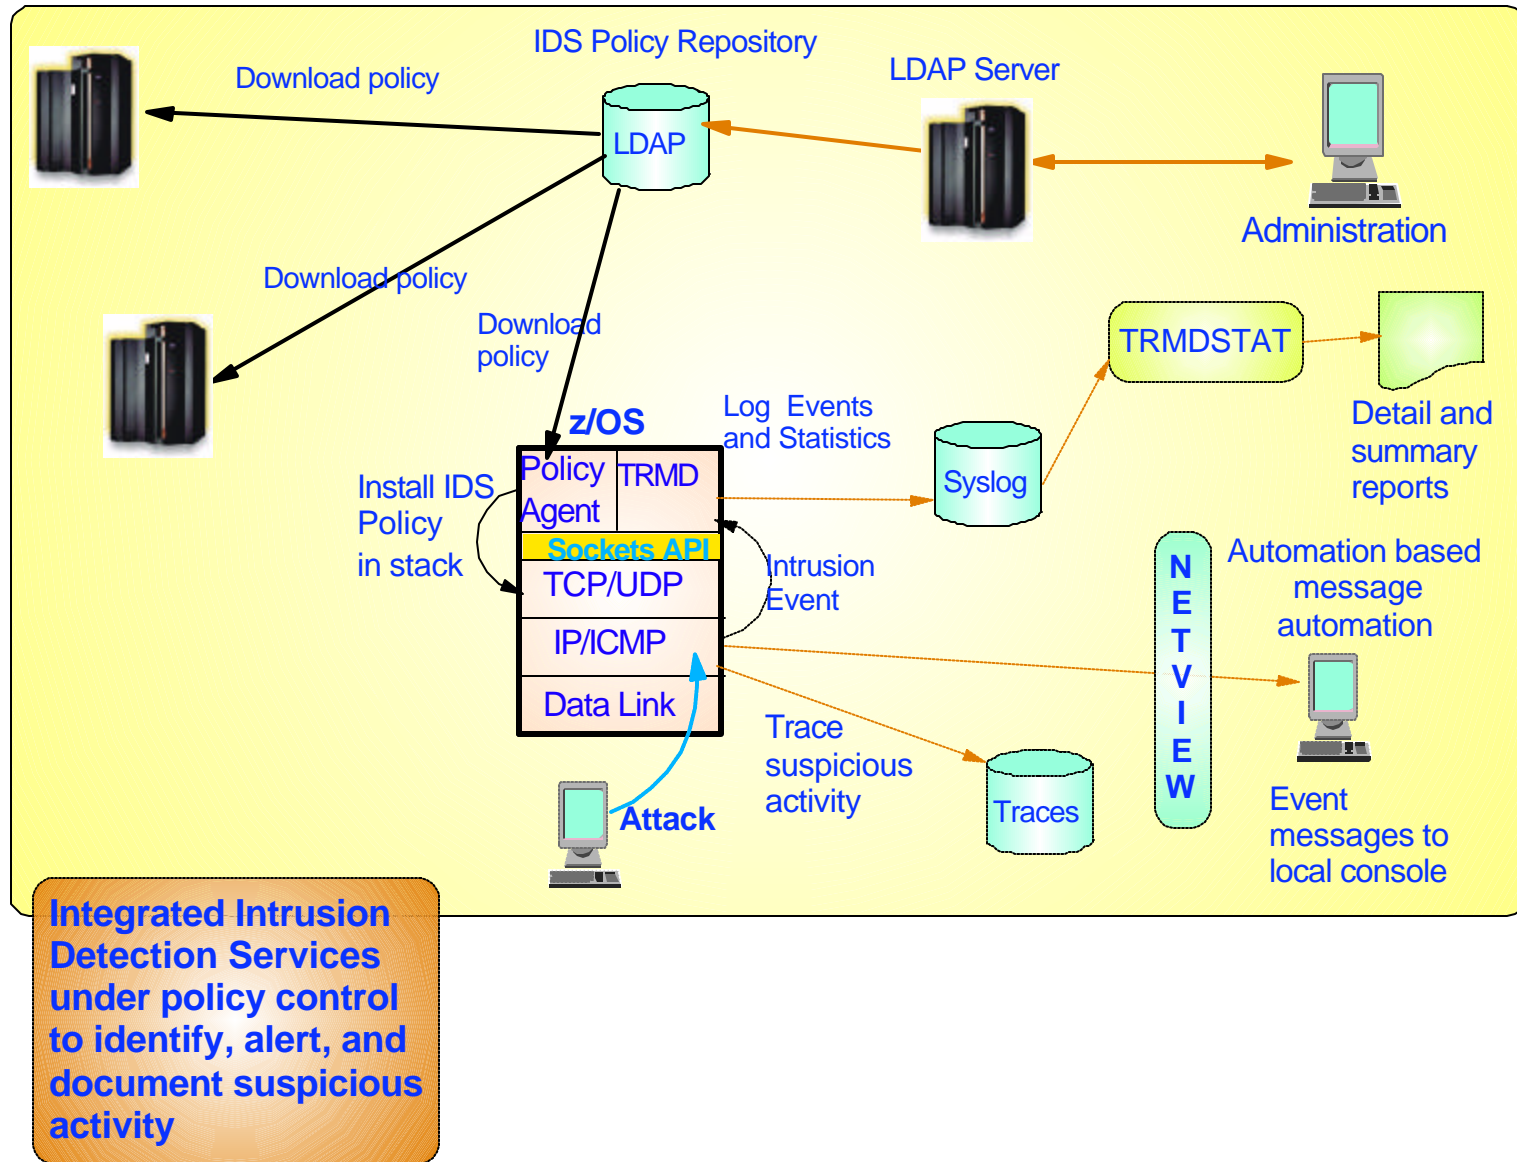
Copyright International Business Machines Corporation 2004. All rights reserved.



Intrusion Detection Services enhancements

Intrusion Detection Services

- ▶ Added new Policy-based interface flood detection
 - Detects floods of "garbage" data on a per-interface basis (ex:)
 - ★ Unsupported protocol
 - ★ Examine packet discards



Multi-Level security enhancements

➤ Governments require control over information accessed through computer networks

- ▶ Applications/users and networks are classified and assigned a level of security
 - ▶ Classifications are done in RACF with SECLABEL
- ▶ MLS for Communication's Server TCPIP stack is based upon the existing NETACCESS function
 - ▶ IP addresses, both INBOUND and OUTBOUND are defined to be in a security zone
 - ▶ Security zone is defined to RACF using a SERVAUTH class profile
 - ▶ User may be RACF PERMITTED to access the SERVAUTH profile (ie. security zone)
 - ▶ In an MLS environment, user and SERVAUTH profile (zone) are assigned a SECLABEL

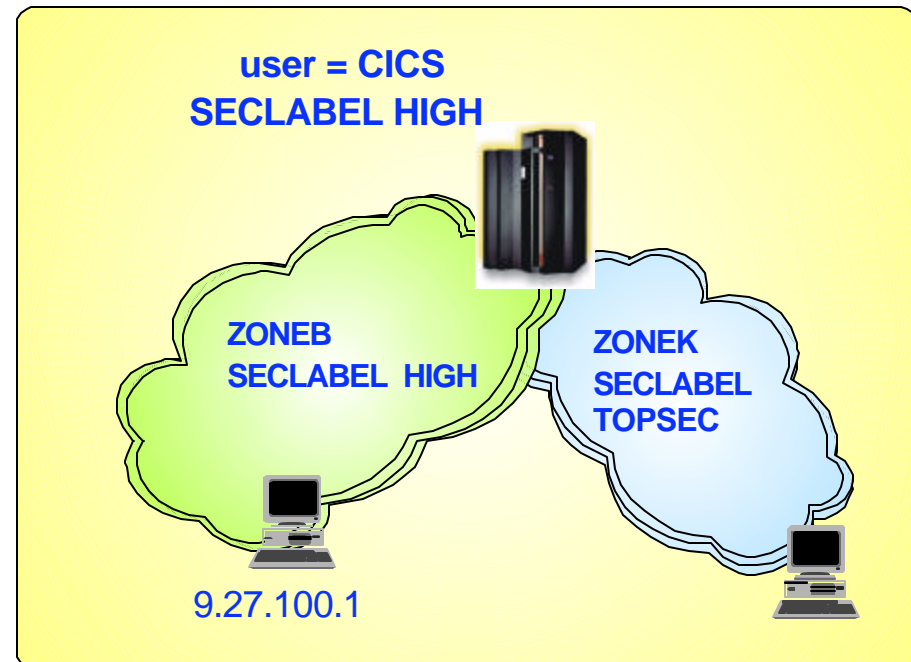
Key: For a user to communicate with a partner in a security zone, SECLABELS must be equivalent (user = zone)

TCPIP Profile:

```
NETACCESS INBOUND OUTBOUND  
9.27.100.0/24 ZONEB  
ENDNETACCESS
```

In RACF:

```
RDEFINE SERVAUTH ZONEB  
PERMIT CICS ZONEB
```



MLS notes

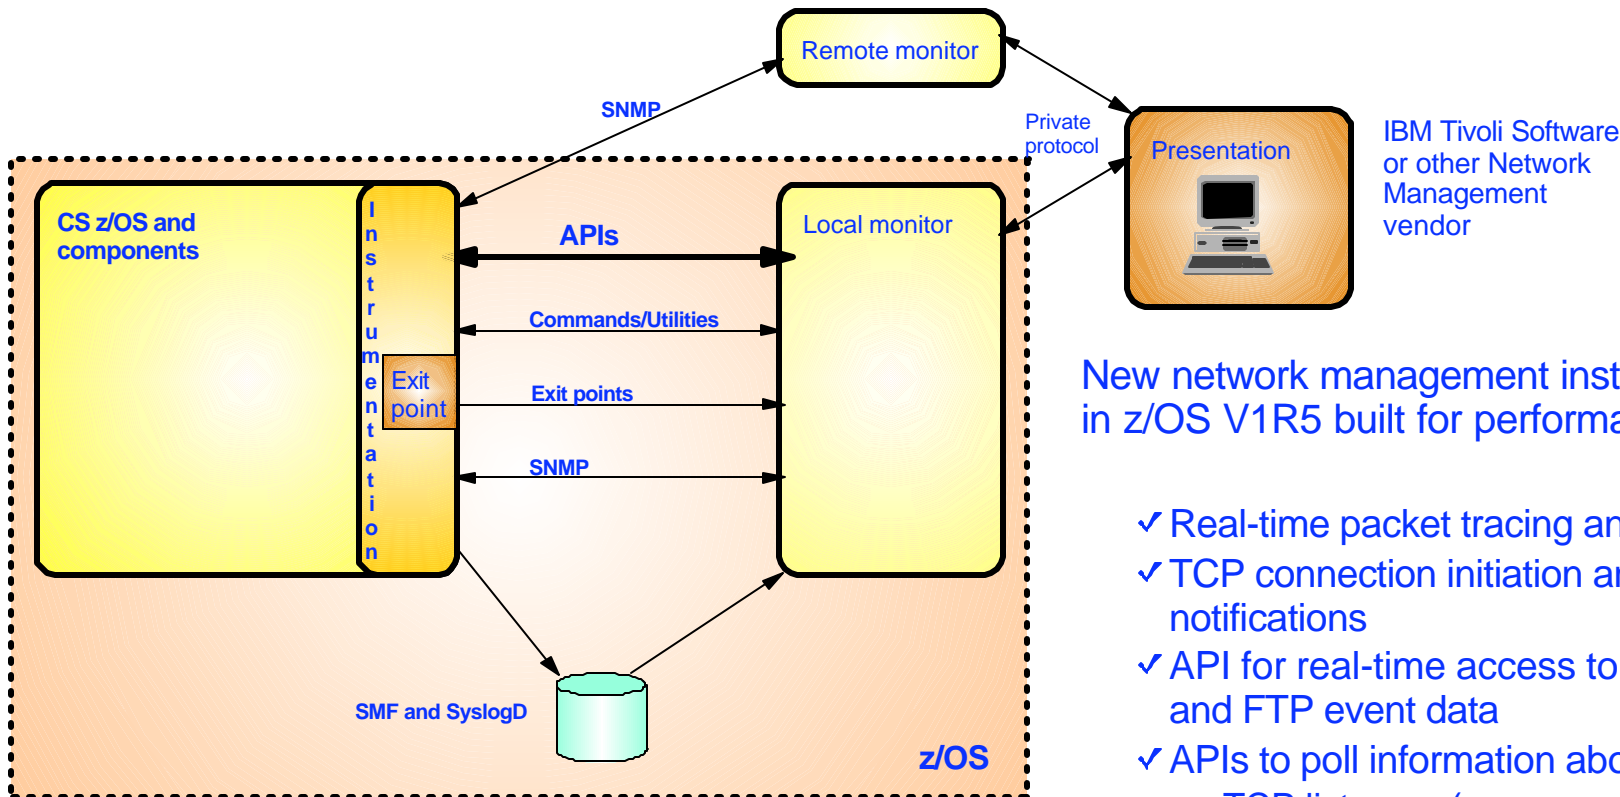
- For z/OS users communicating in an MLS environment across a sysplex (IUTSAMEHOST and XCF links), network packets carry the SECLABEL of the sender and are checked for SECLABEL equivalence at the destination host
- Packets sent /received outside of a sysplex are assume to have equivalent SECLABELS if NETACCESS checks are successful
- NETACCESS and MLS for Applications
 - ▶ FTP Poe - RACF Port of Entry (PoE) checks are done to determines if client/user is permitted to logon. Applies to MLS and nonMLS environments.
 - Determines if the client's NETACCESS zone is permitted to the zone of the logon userid.
 - For MLS, client and logon userid must have equivalent SECLABELs.
 - ▶ Tn3270 server Poe - determines if the client's NETACCESS zone has an equivalent SECLABEL to the Tn3270 LUname selected
 - LUGROUP gets a SECLABEL
 - ▶ Tn3270 can also determine if a client is NETACCESS permitted to a Tn3270 port
 - NonMLS
 - Port gets a security zone assigned in Tn3270 profile

Other TCP/IP Enhancements

Copyright International Business Machines Corporation 2004. All rights reserved.



Network Management Instrumentation Architecture



New network management instrumentation APIs in z/OS V1R5 built for performance:

- ✓ Real-time packet tracing and formatting
- ✓ TCP connection initiation and termination notifications
- ✓ API for real-time access to TN3270 server and FTP event data
- ✓ APIs to poll information about currently active
 - TCP listeners (server processes)
 - TCP connections (detailed information about individual connections)
 - UDP endpoints
 - CS storage usage
- ✓ API to receive and poll for Enterprise Extender management data

Tivoli's IBM Tivoli Monitor / Network Performance (ITM/NP) use these new APIs

The APIs and their documentation will be available as part of CS z/OS for use by network management vendors and customer network management applications.

Other TCPIP enhancements

- **OMPROUTE will be changed in this release to:**
 - ▶ Improved trace performance by using CTRACE
 - ▶ Multipath enhanced to support 16 routes to same destination
 - ▶ Support advertisement for more than 254 interfaces (no upper limits)
 - ▶ Some display clean-up to show N/A where an option doesn't apply
 - ▶ Ignore undefined interfaces (don't have to define everything to OSPF)

- **Enhance MVS system symbol resolution**
 - ▶ In TCPIP.DATA
 - ▶ No longer requires use of EZACFSM1 preprocessor utility
 - ▶ In Resolver setup file

- **Increase Max. Number of allowed sockets from 2000 to 65535**
 - ▶ Macro API (EZASMI), Sockets Extended (callable EZASOKET), CICS Sockets (callable EZASOKET), and IMS Sockets

- **Improve performance for asynchronous stream socket receive operations**
 - ▶ LE C/C++ API
 - ▶ z/OS UNIX Assembler Callable Services Sockets API

- **MSys support for complete configuration of the FTP client and server**

Full VLAN support by z/OS and other OSA enhancements for z990

➤ Provide Full VLAN support for OSA-Express

- z/OS V1R2 supported only VLAN priority-tagged frames
 - ▶ supported via Policy Agent configuration statement as parameter on SetSubnetPrioTosMask
- Extend VLAN in z/OS V1R5 to support non-null VLAN ID (Full VLAN) with OSA Express
 - ▶ Supported for OSA Express QDIO Gigabit Ethernet and Fast Ethernet with IBM Server zSeries z900 (OSA driver level 3G), and z990
 - ▶ VLAN ID is configured in TCPIP profile on the LINK (IPv4) or Interface (IPv6) Statement
 - ▶ Supports only one VLAN ID for IPv4 interface and one VLAN ID for IPv6 (IDs can be different) per OSA port

➤ OSA Express QDIO performance improvements for z990

- Offload checksum to Improve performance for IPv4 packets using QDIO mode
- Provide control of storage utilization inbound per Tcpiplink/interface
- Provide some controls over the inbound performance per link/interface - latency vs. CPU cost

➤ Broadcast support for Hipersockets Interface (IPv4) for z990

- ▶ IEEE standard 802.1Q defines operation of Virtual Bridged LANs
- ▶ Defines additional information in the MAC frame (tag carries VLAN-ID)
- ▶ Each port on the VLAN-aware bridge has a VLAN-ID
- ▶ VLAN-aware bridges can be configured to:
 - ✓ Admit only VLAN-tagged frames (i.e.. have both priority and non-null VLAN ID)
 - ✓ Admit all frames

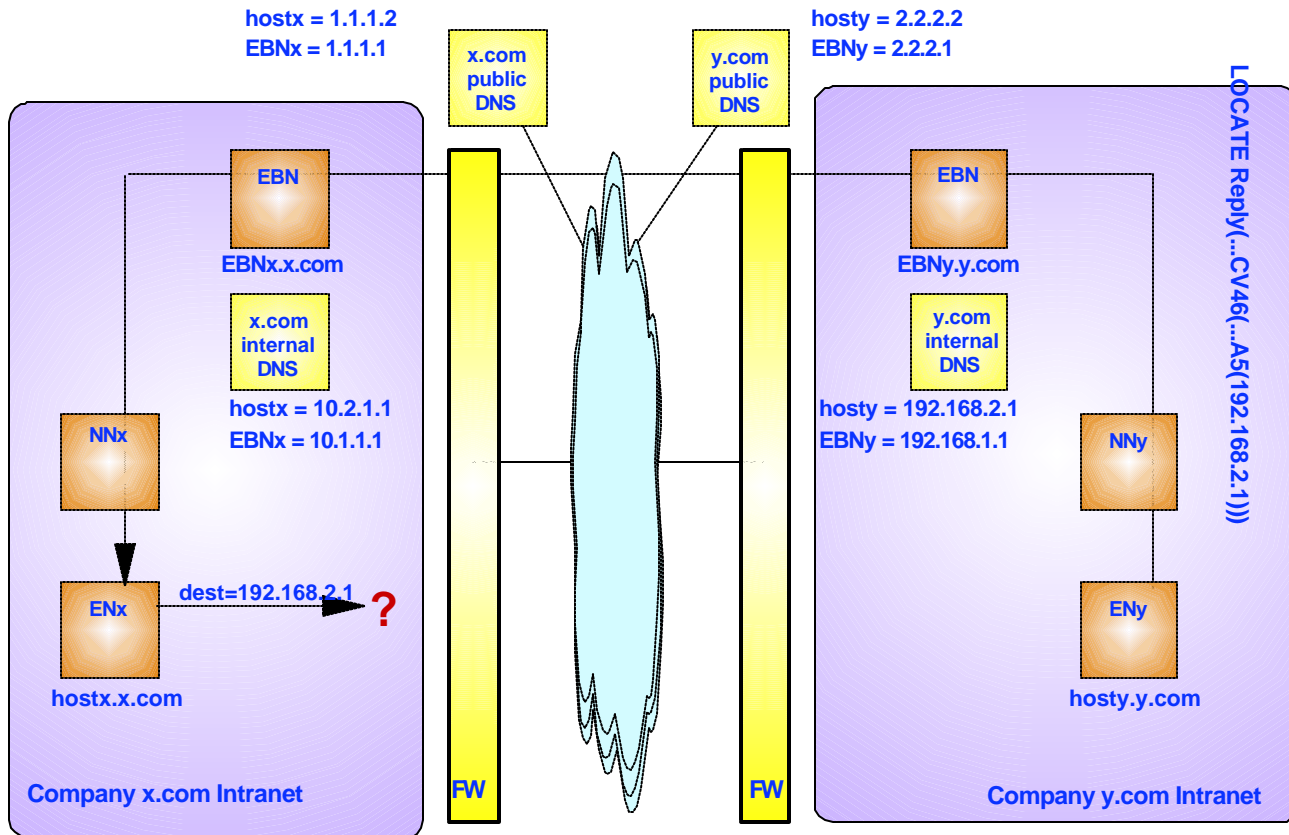
★ Allows packets to be isolated to selected network interfaces and physical networks and adapters to be shared for multiple purposes, such as various security zones or customers..

Enterprise Extender and SNA Enhancements

Copyright International Business Machines Corporation 2004. All rights reserved.



NAT Compatibility with EE VRN

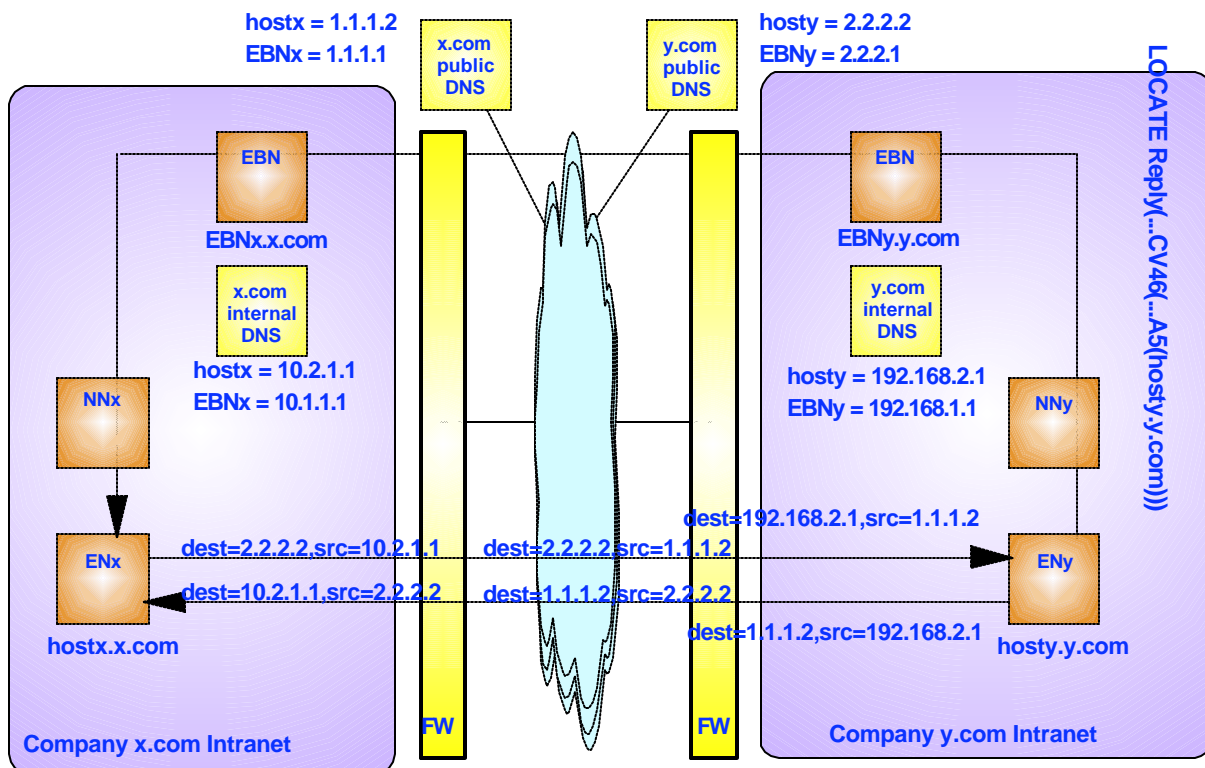


- The control vectors representing an EE connection network carry the IP address of the node that owns the vectors.
- Applications that carry IP addresses in their data stream are fundamentally incompatible with Network Address Translation (NAT).

Intranet #	Public #
10.1.1.1	1.1.1.1
10.2.1.1	1.1.1.2

Public #	Intranet #
2.2.2.1	192.168.1.1
2.2.2.2	192.168.2.1

Enhanced HOSTNAME Support




Intranet #	Public #
10.1.1.1	1.1.1.1
10.2.1.1	1.1.1.2

Public #	Intranet #
2.2.2.1	192.168.1.1
2.2.2.2	192.168.2.1

- EE architecture has been updated to allow the EE connection network control vectors to carry the hostname corresponding to the EE VIPA.
- IPv4 connection network control vectors will continue to carry IP address as well as hostname to ensure compatibility with downlevel nodes.
- Administrative requirement of coordinating NAT tables and public DNS entries is a known administrative procedure to installations that use NAT.
- There will be recommended maximum fully-qualified hostname lengths due to limited space in the route selection control vector.

Multiple VRN/VIPA Support



- V1R5 will allow the specification of multiple local and/or multiple global EE connection networks.
 - In the diagram above, Node B defines 2 local VRNs (both IPv4) and 2 global VRNs (one IPv4 and one IPv6)
- EE will allow multiple (static) VIPAs, defined on a GROUP basis in the EE XCA major node.
 - All EE VIPAs must still belong to a single TCP/IP stack

Enterprise Extender Model PUs

- Currently, unless the Configuration Services XID exit is used, EE does not allow the user to code a model for non-connection- network dynamic PUs created on the dial-in side.
 - If DYNPU=YES is specified, EE creates a "vanilla" dynamic PU.
- A new model type, DYNTYPE=EE, is added to the PU statement in the Model Major Node.
- An EE model PU allows for customizing dynamic non-connection network PUs in various manners:
 - Coding the DISCNT operand so that a disconnect delay time can be specified.
 - Specifying the DWINOP, REDIAL, and REDDELAY operands to specify whether to drive redial attempts, and how often and how many times to attempt redial, when INOPs occur for the connections using these dynamic PUs.
 - Overriding the default TG characteristics that in the past were used for these dynamic PUs.

```
MODEL A1A VBUILD TYPE=MODEL
*
EEMODEL PU DYNTYPE=EE,
CAPACITY=100K,
COSTTIME=5,
CPCP=YES,
DISCNT=NO,
DWINOP=YES,
REDIAL=30,
REDDelay=60
```


D RTPS Enhancement

- In V1R5, the DISPLAY RTPS command is enhanced to:
 - Allow the display of all RTP pipes using a particular link for the first hop. The link can be specified by either:
 - ALSNAME
 - TG number & partner CP name
 - Allow the display of all RTP pipes going to a particular Net ID (either as the destination or as the first hop). This is done by specifying *netid.** where a CP name is specified.

```
d net,rtps,firsttg=21,firstcp=sscp2a
```

```
IST097I DISPLAY ACCEPTED
```

```
IST350I DISPLAY TYPE = RTPS
```

IST1695I	PU NAME	CP NAME	COS NAME	SWITCH	CONGEST	SESSIONS
IST1696I	CNR00005	NETA.SSCP2A	#INTER	NO	NO	1
IST1696I	CNR00004	NETA.SSCP2A	#BATCH	NO	NO	1

```
IST1454I 2 RTP(S) DISPLAYED
```

```
IST314I END
```

```
d net,rtps,cpname=neta.*
```

```
IST097I DISPLAY ACCEPTED
```

```
IST350I DISPLAY TYPE = RTPS
```

IST1695I	PU NAME	CP NAME	COS NAME	SWITCH	CONGEST	SESSIONS
IST1696I	CNR00005	NETA.SSCP2A	#INTER	NO	NO	1
IST1696I	CNR00004	NETA.SSCP2A	#BATCH	NO	NO	1
IST1696I	CNR00001	NETA.SSCP1A	CPSVCMG	NO	NO	1

```
IST1454I 3 RTP(S) DISPLAYED
```

```
IST314I END
```

Other EE/SNA enhancements

- When dial failures occur across a connection network, the virtual routing node state is usually not affected, so the APPN topology and routing component continues to select the VRN if it is the preferred path. V1R5 adds a more specific failure message to indicate that a dial across a VRN has failed, providing an easier trigger for customer notification and automation:

```
IST1903I FAILURE OVER VRN NETA.SSCPVN TO CP NETA.SSCP2A
```

- V1R5 adds new messages for accounting and problem determination capability for DLUR-based sessions
 - Messages issued when a DLUR-served PU session is established or terminated with VTAM:
 - New message to include DLUR name on ACTPU and ACTLU failures
- Support for multiple concurrent APINGs
- Remove restriction that the backup NNS must be in same sysplex as the End Node for generic resource to work (allows the End Node to be searched for the resource)
- HPR Resequencing Optimization
 - Improves inbound processing of out of order segmented packets
- Modify IOBUF pool expansion limit without recycling VTAM

Other EE/SNA Enhancements...

Serviceability/Problem Determination

- Usability improvements to CSDUMP command (dump for sense code or message)
 - ▶ Display settings of CSDUMP command and allow for the deletion of triggers
- APPN Trace enhancement
 - ▶ New SUBTRACE option TGVC will provide TG Vectors in appropriate trace records
- Track CSM Buffers internally through components
- Enable dumping for VTAM Inoperative conditions with new Modify INOPCode command
 - ▶ Works in conjunction with existing Modify INOPDump command
- Allow for extended information to be displayed for search failures that cause sessions not to be established: DSIRFMSG start option

Performance/Storage

- Multipath Channel HPDT Packing (nonXCF) examines packet sizes for optimal outbound transmission
 - ▶ Packing turned on/off at TRLE. Requires some workload examination for tuning.

New/Enhanced Functionality

- Sift parameters from the Group on model major node for dynamic switched LUs
 - ▶ Allows for definition reduction
- SWNORDER and DLRORDER
 - ▶ Enhanced to allow greater control over PU selection during connection processing. On start option or in XCA or NCP major nodes.
- New search option, SSCPORD, provides granular control over searches for resources
 - ▶ defined on start option or in ADJSSCP table

EE Performance Improvement

- OSA Microcode Level: 0326
- z/OS Communications Server Maintenance:

APAR	Purpose	PTF	Notes
OW53393	ARB Enhancements	UW94491	V1R2 only, base in V1R4
OW56896	LAN Idle	UA00067	V1R2 only, base in V1R4
OW52291	EE Packing & QDIOSTG Option	UA00131 - V1R2 UA00132 - V1R4	
OW53978	EE Outbound Data Ordering	UA00131 - V1R2 UA00132 - V1R4	Coreq: PQ69398
PQ69398	Fast UDP Outbound Ordering	UQ73923	V1R2 only, base in V1R4
OW57459	HPR Resequencing	UA00131	V1R2 only, base in V1R4
OA02213	Send SRB Optimization	UA01999 - V1R2 UA02000 - V1R4	(2)



1. Monitor the EE Info APAR (II12223) for other announcements
2. Will prereq other VTAM APARs in table
 - Recommendation: Apply OA02213 & PQ69398; Upgrade OSA microcode to 0326

EE Performance Improvement...

- Relative to our base system:

(Includes OSA 0326 microcode, plus OW53393)

- ▶ RR CPU reduced up to 5%
- ▶ Stream CPU reduced 16-24% with an 8K RU
 - Reduced about 5% with 1920 byte RU
 - Flat for 16K RU
- ▶ RR Throughput relatively flat
- ▶ Stream Throughput up 14-16% (1920 byte and 8K)
 - Up 31% for 16K RU



- Relative to previous V1R2 measurements:

(pre-0320 OSA microcode, no maintenance)

- ▶ Stream Throughput up 94.9%



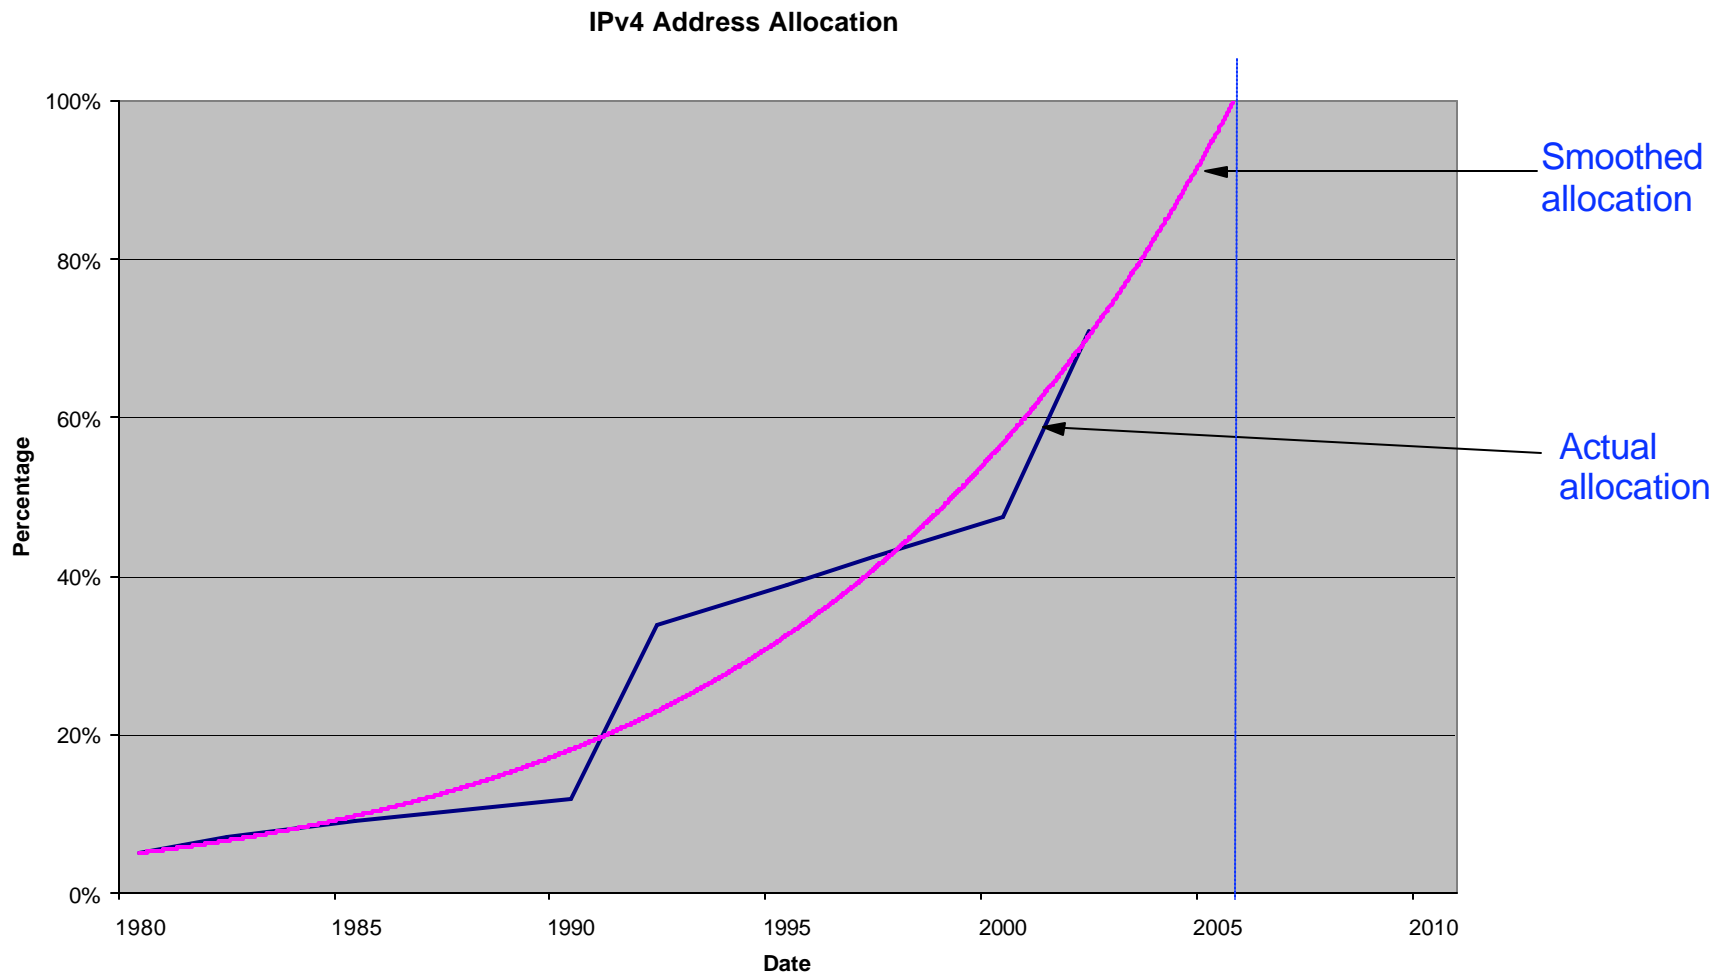
IPv6 on z/OS Communications Server

Copyright International Business Machines Corporation 2004. All rights reserved.




We have IPv4 addresses enough - or do we? Latest IPv4 address space usage overview

The chart shows IPv4 address allocation over time. The "blue" line is the actual allocation, the "purple" line is the smoothed allocation. This shows the IPv4 address pool being depleted in the next 5 years or so.



Visible IPv4 hosts on the Internet through the last 10 years



Source: <http://www.isc.org/ds/WWW-200301/index.html> and "An update on the H ratio", RFC3174

z/OS CS V1R4 and V1R5 IPv6 contents overview

Basic IPv6 Support

- ✓ IPv6 dual-mode stack, including Neighbor Discovery, MLD, etc.
- ✓ DNS BIND 9.2
- ✓ Static VIPA
- ✓ Service Tools

Network Management

- ▶ IPv6 mibs
- ▶ osnmp
- ▶ osnmpd & assoc'd functions
- ▶ MVS TCPIP subagent
- ▶ SMF119 records
- ▶ netstat

DLCs

- ✓ OSA-E QDIO
 - ▶ XCF
 - ▶ Samehost
 - ▶ ESCON/FICON (MPCPTP)

Security

- ▶ NetAccess

Policy

- ▶ OoS stack
- ▶ Pagent
- ▶ Pasearch
- ▶ SLAPM Mib #2

APIs

- ✓ Basic Socket APIs
 - UNIX (LE) C/C++
 - USS Callable
 - TCP/IP Macro API
 - TCP/IP Call Instruction API
 - TCP/IP Rexx Sockets
- ✓ Advanced Socket APIs
 - UNIX (LE) C/C++
 - USS Callable
- ▶ CICS Sockets
 - CICS C
 - EZASOKET

Routing

- ✓ Static Routing
- ▶ RIPng with OMPRoute

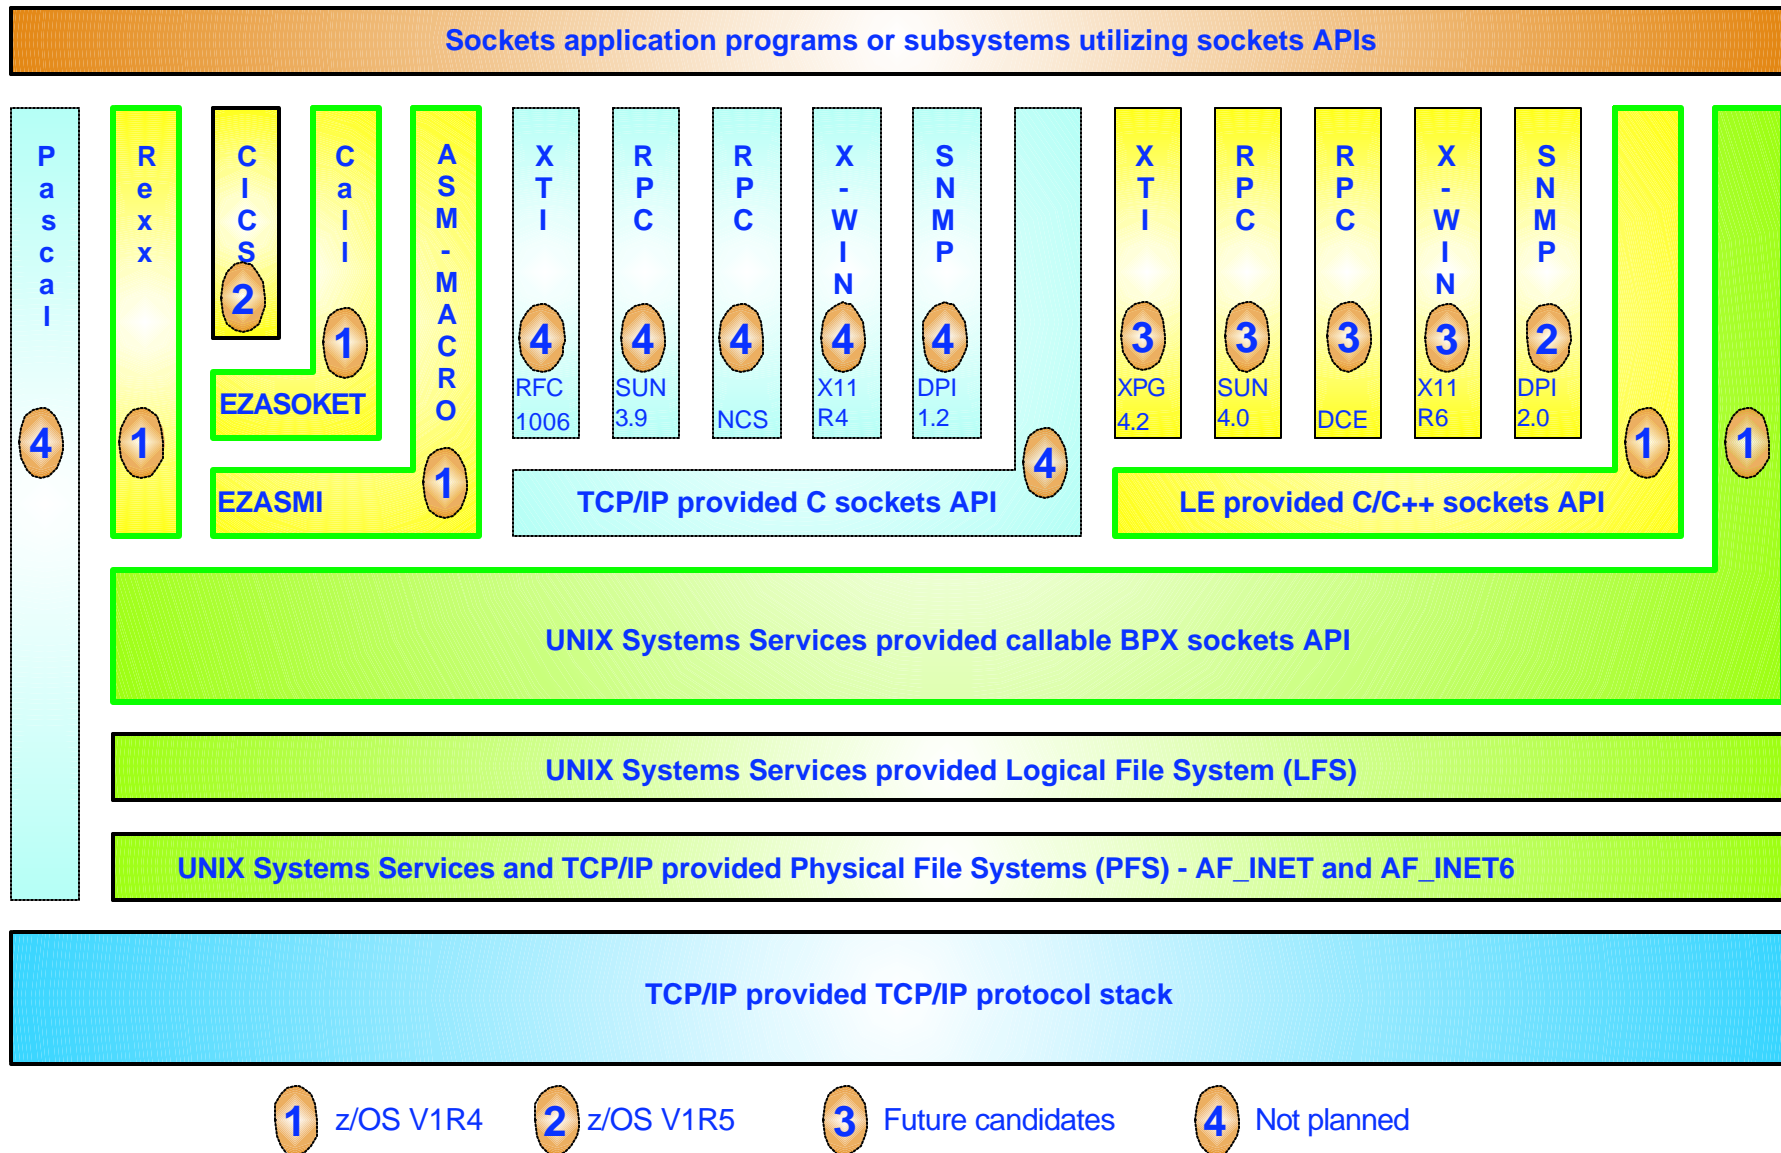
Applications

- ✓ inetd, ftpd, ftp, telnetd, USS rshd, USS rexec, USS rexecd, Ping, Tracert, Netstat
- ▶ TN3270 server
- ▶ Enterprise Extender
- ▶ syslogd
- ▶ dcas
- ▶ sntp
- ▶ tfptd
- ▶ MVS rexecd / rshd
- ▶ TSO rsh /rexec client
- ▶ sendmail
- ▶ UNIX rsh clientT

Legend:

- ✓ Provided in V1R4
- ▶ Provided in V1R5

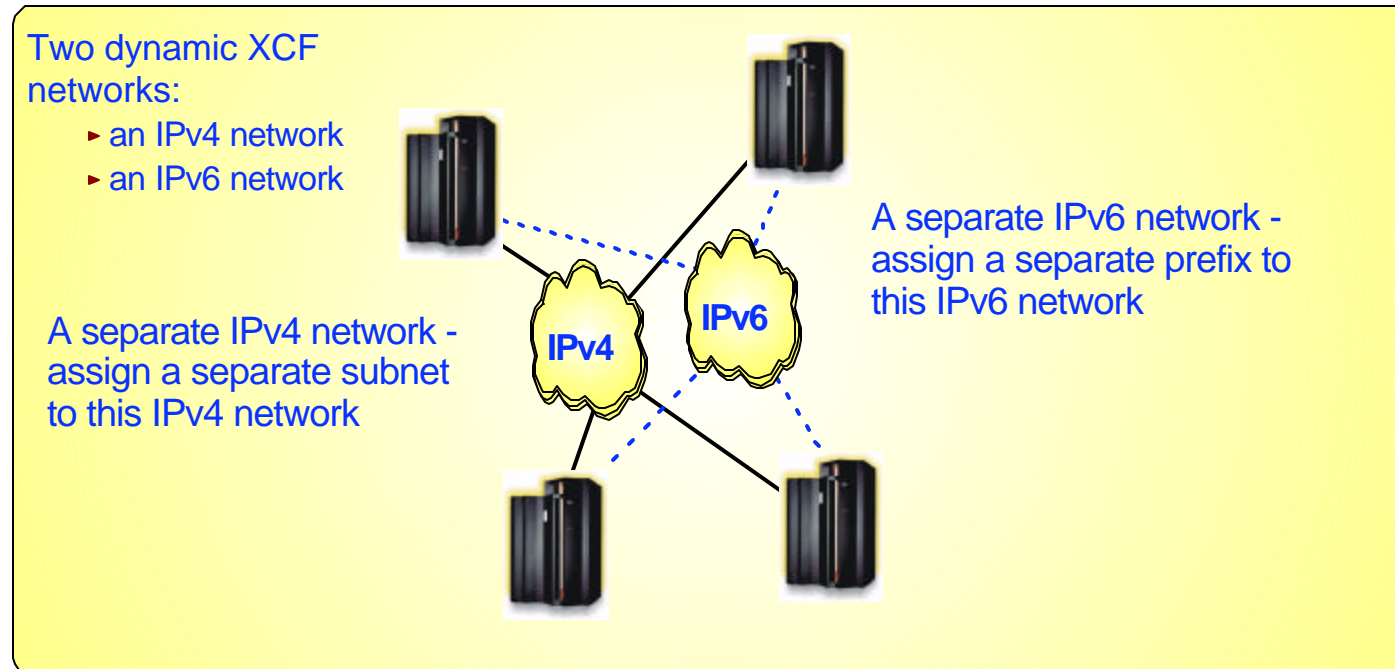
IPv6: Sockets-related API AF_INET6 enablement overview, status, and plans



IPv6: New interface support

IPv6 MPCPTP-type Interface support - Extends connectivity options for IPV6

- ▶ IUTSAMEHOST to other stacks in same LPAR
- ▶ XCF to other stacks in same Sysplex
 - ▶ Both static and dynamic XCF, but not at the sametime! Must either have static xcf or dynamic.
- ▶ ESCON/FICON to another z/OS image (not to any known Channel-attached Routers)



IPv6: Another batch of AF_INET6 enabled applications

IPv6-enable another batch of applications

- tftpd (trivial file transfer server)
- syslogd
- dcas (digital certificate access server)
- sntpd (simple network time protocol server)
- TN3270 Server: provide connectivity to SNA applications from IPv6 TN3270 clients
- Enterprise Extender: allow SNA applications to communicate over IPv6 networks
- sendmail 8.12.5 (new port of sendmail picks up IPv6 enablement too)
- MVS rshd/rexecd server
 - the UNIX rexecd and rshd servers were IPv6-enabled in z/OS V1R4
- TSO rsh/rexec clients
 - provide updated version that can be used in all z/OS environments (batch, TSO, REXX, etc.)
- New UNIX rsh client that is IPv6-enabled from start
 - the UNIX rexec client was IPv6-enabled in z/OS V1R4

Changes to support partners with V6 addresses:

- ◆ Changes to socket api calls and structures
- ◆ Configuration file changes

IPv6-enable the CICS listener, sockets infrastructure, and sockets libraries

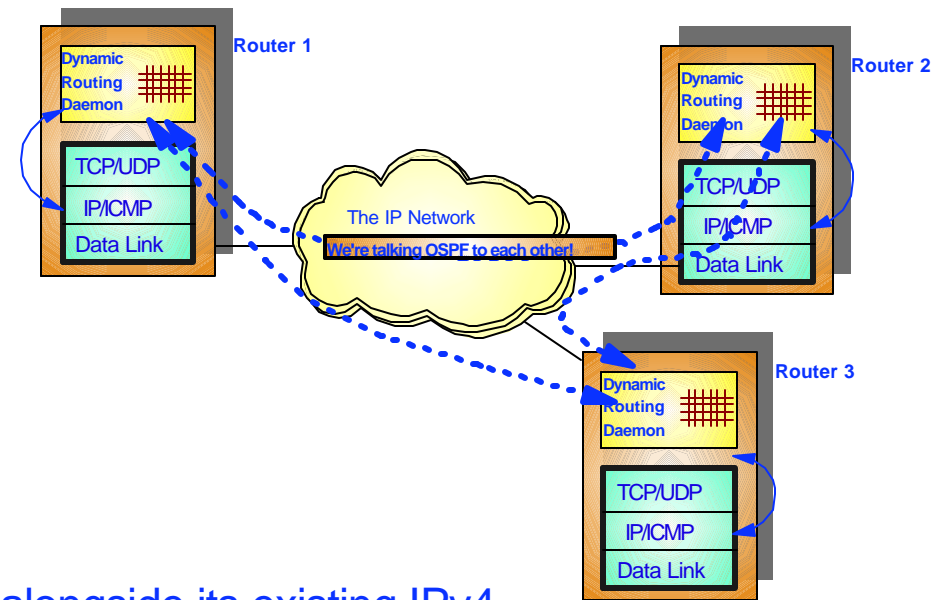
IPv6: Add RIPng to OMPRoute and other OMPROUTE enhancements

➤ OMPROUTE has been enhanced to support IPv6 and RIPng

- Like IPv4, there is support for the routing protocol, plus support for basic IPv6 routing concepts
 - ▶ generic interfaces
 - ▶ static routes
 - ▶ direct routes
 - ▶ prefix and router advertisement routes
- This new support has been added to OMPROUTE alongside its existing IPv4 dynamic routing support
- You use new sets of IPv6 configuration statements and display commands to activate and monitor this new support

➤ OMPROUTE has also been changed in this release to:

- Improved trace performance by using CTRACE
- Multipath enhanced to support 16 routes to same destination
- Support advertisement for more than 254 interfaces (no upper limits)
- Some display clean-up to show N/A where an option doesn't apply
- Ignore undefined interfaces (don't have to define everything to OSPF)

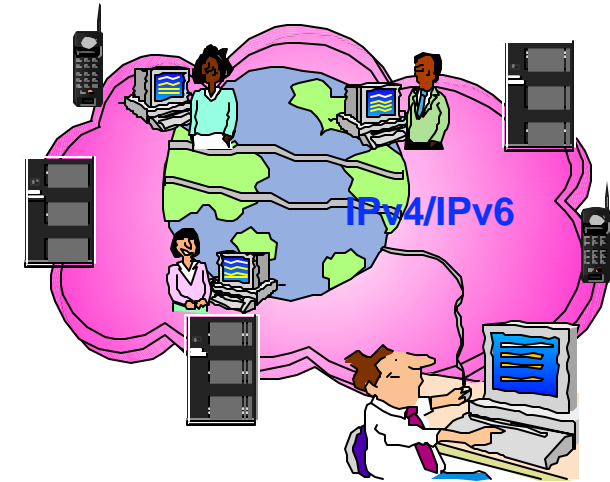


IPv6: Network management support for IPv6

➤ Network management SNMP support

- Support SNMP agent (OSNMPD)
- DPI 2.0 enabled for AF_INET6
- Support MVS TCPIP subagent
- osnmp command
- The trap forwarder daemon enabled for AF_INET6
- IPv6 MIB support - new RFC drafts have been published that define IP version neutral objects
 - RFC2011 (IP and ICMP)
 - RFC2012 (TCP)
 - RFC2096 (IP routes)
 - RFC2233 (Interfaces) - this one is not version neutral

Must be able to monitor an IPv6 network



➤ SMF119 records support

- The redesign in z/OS V1R2 did factor in IPv6 addresses, so most subtypes are already in z/OS V1R4 supporting IPv6 addresses
- Some changes needed to selected records to capture additional IPv6-related data, such as interface records and statistics records

➤ Non IPv6-specific enhancements to SMF119 records:

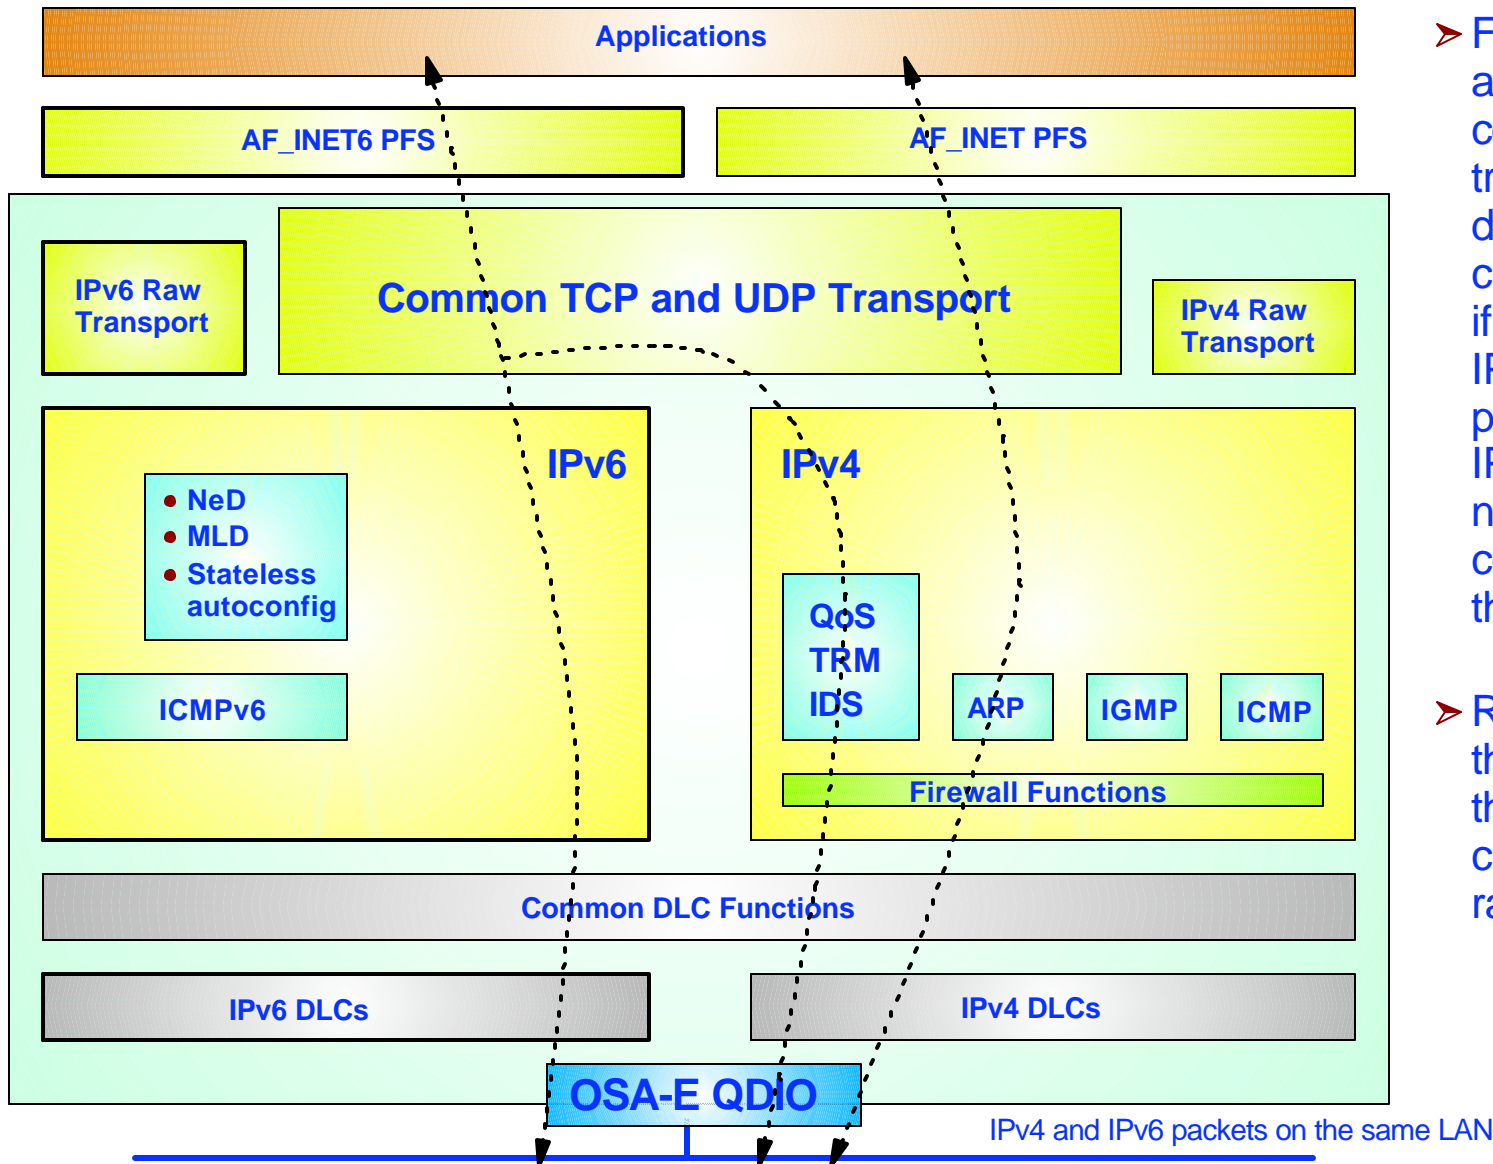
- Security information has been added to the FTP server and client file transfer completion records
- TCP connection termination records now, for TN3270 server connections, include TN3270-specific information: LU name, SNA Application name, protocol, etc.

➤ And some netstat goodies:

- add support for a hostname filter for connection-type reports
- add existing IP address filter to BYETINFO report
- add interface/link filter to DEVLINKS reports
- add interface statistics sections to DEVLINKS report

At this time, the UDP MIB draft was very unstable and too risky to implement - will follow in a later release when RFCs have stabilized.

Flow between LAN and applications on an IPv6-enabled z/OS LPAR



➤ For AF_INET6 applications, the common TCP or UDP transport layer determines per communication partner if the partner is an IPv4 or an IPv6 partner - and chooses IPv4 or IPv6 networking layer component based on that.

➤ Raw applications make the determination themselves when they choose IPv4 or IPv6 raw transport.

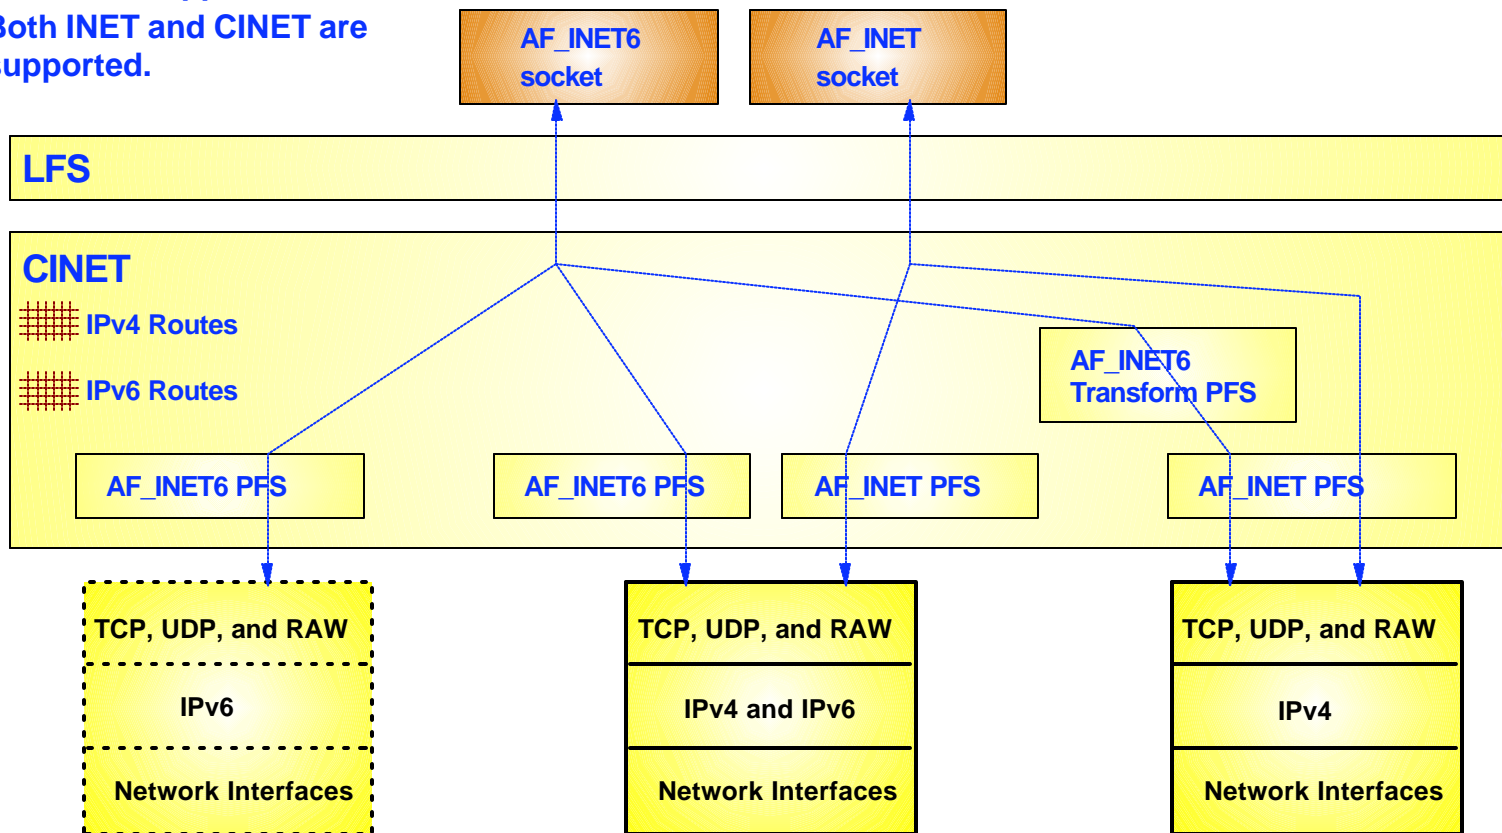
OSA-E for IPv6 requires zSeries hardware

How do we enable IPv6 support and what are the consequences?

IPv6 is enabled at an LPAR level via an option in BPXPRMxx to enable AF_INET6 support.

Both INET and CINET are supported.

When IPv6 is enabled, a z/OS V1R4 TCP/IP stack will always have an IPv6 Loopback interface. You can define real IPv6 interfaces in addition to the loopback interface.



IPv6-only TCP/IP Stack

This will not be the case on z/OS for the foreseeable future! An AF_INET6 stack is required to also support AF_INET!

Dual Mode TCP/IP Stack

A z/OS V1R4 TCP/IP stack will always come up as dual-mode if AF_INET6 is enabled in BPXPRMxx

IPv4-only TCP/IP Stack

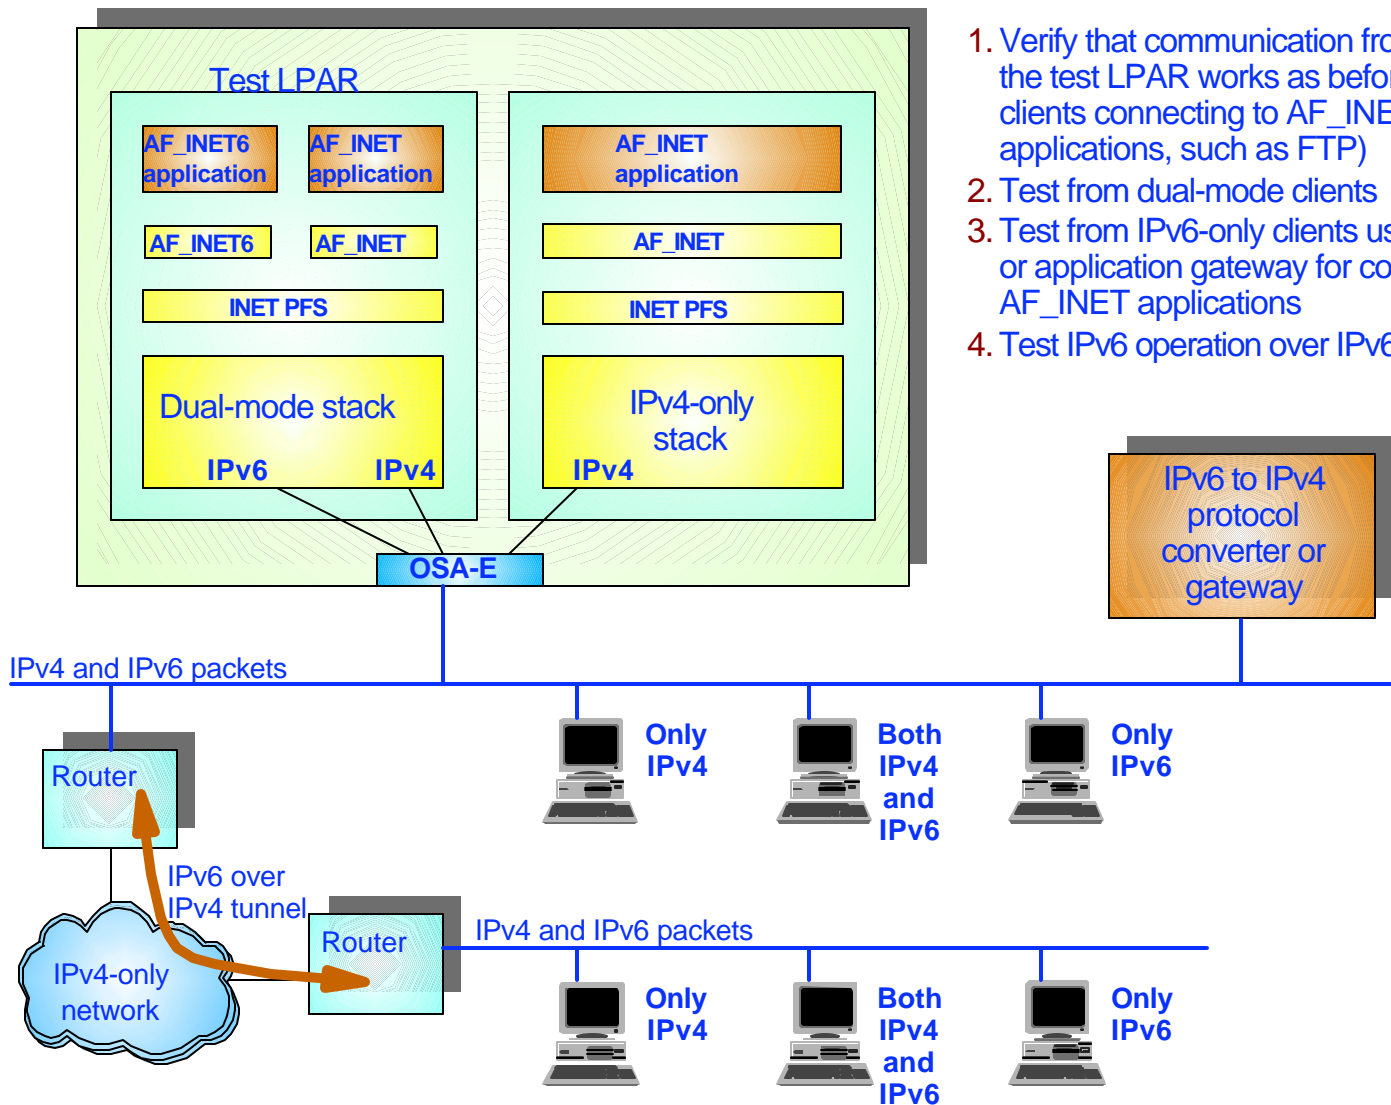
(such as AnyNet or an OEM TCP/IP stack)

- ▶ Existing AF_INET sockets programs will continue to work as they always did - no difference in behavior or support.
- ▶ AF_INET6 enabled sockets programs will be able to communicate with IPv4 partners (just as before they were changed to support IPv6), but in addition to that they will also be able to communicate with IPv6 partners.

When IPv6 is enabled, most netstat reports will look different because of the potential for long IPv6 addresses.

Make sure you have modified any netstat screen-scraping REXX programs you might have developed in the past!

Start testing IPv6 on z/OS



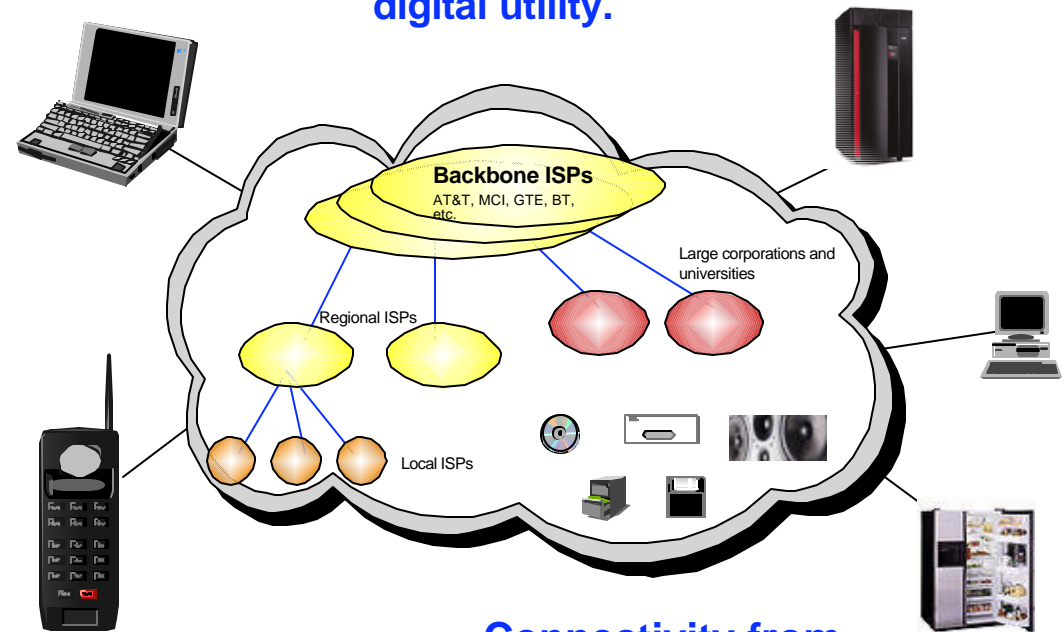
1. Verify that communication from IPv4-only clients to the test LPAR works as before (including IPv4-only clients connecting to AF_INET6-enabled server applications, such as FTP)
2. Test from dual-mode clients
3. Test from IPv6-only clients using protocol converter or application gateway for communication with AF_INET applications
4. Test IPv6 operation over IPv6 tunnel

The Journey to IPv6 for z/OS Comm Server

IPv6 Development Phases

- ▶ The first phase (V1R4)
 - ▶ Stack support for IPv6 base functions
 - ▶ Resolver
 - ▶ High speed attach (QDIO DLCs)
 - ▶ Service tools (Trace, Dump, etc.)
 - ▶ Configuration and netstat
 - ▶ Static Routing
 - ▶ FTP, otelnetd, unix rexec, unix rshd/rexecd
- ▶ The second phase (V1R5)
 - ▶ Network Management
 - Applications and DPI
 - Version neutral Tcp/Ip Standard MIBs
 - Additional SMF records
 - ▶ Applications/Clients/APIs
 - Tn3270 server, CICS sockets, sendmail, ntp, dcas, rxserve, rsh client
 - ▶ Enterprise Extender
 - ▶ Point to Point - type DLCS
 - ▶ Dynamic Routing Protocol (RIPng)
- ▶ To follow:
 - ▶ Sysplex Exploitation
 - ▶ Dynamic Routing Protocol (OSPFv3)
 - ▶ Complete Network Management MIBs
 - ▶ HiperSockets
 - ▶ Security (IPSec, IDS, SWSA etc.)
 - ▶ Remaining applications and DLCs
 - ▶ Additional IPV6 technologies

The Internet - a worldwide digital utility.



Connectivity from anywhere (car, home, office) to anything!

Note: Plans are subject to change!

For More Information....

URL	Content
http://www.ibm.com/servers/eserver/zseries	IBM eServer zSeries Mainframe Servers
http://www.ibm.com/servers/eserver/zseries/networking	Networking: IBM zSeries Servers
http://www.ibm.com/servers/eserver/zseries/networking/technology.html	IBM Enterprise Servers: Networking Technologies
http://www.ibm.com/software/network	Networking & communications software
http://www.ibm.com/software/network/commserver	Communications Server
http://www.ibm.com/software/network/commserver/library	Communications Server white papers, product documentations, etc.
http://www.redbooks.ibm.com	ITSO redbooks
http://www.ibm.com/software/network/commserver/support	Communications Server technical Support
http://www.ibm.com/support/techdocs/	Advanced technical support (flashes, presentations, white papers, etc.)
http://www.rfc-editor.org/rfcsearch.html	Request For Comments (RFC)