e-business

Communications Server z/OS V1R5 Technical Update

The journey towards the next generation Internet - IPv6

z/OS V1R5

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

Topics

e-business

- Brief recap of why IPv6 is of interest
- IPv6 basics review
- IPv4/IPv6 coexistance and migration overview
- Basics of IPv6 on z/OS
- z/OS V1R5
  - Sendmail 8.12 IPv6 support (*)
  - CICS sockets IPv6 enabled (*)
  - IPv6-enabled another batch of applications
    - SNTPD, SyslogD, TFTPD, DCAS, remote execution commands and servers (*)
  - IPv6-enabled SNMP environment
  - New IPv6 interface support
    - XCF - dynamic and static
    - IUTSAMEH
    - MPCPTP6
  - Policy agent IPv6 enabled - IPv6 QoS policy support (*)
- z/OS V1R6
  - Full sysplex support for IPv6
    - `DVIPAs, Sysplex Distributor, SourceVIPA

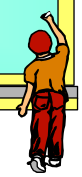

(*) - non-IPv6 enhancements included in this section

z/OS V1R5 has been certified with the IPv6 Ready logo

IPv6 Ready Logo Program by IPv6 Forum - Microsoft Internet Explorer

http://www.ipv6ready.org/logo_db/logo_search2.php?logoid_number=01-000156&btm=S

Item	Content
Logo ID	01-000156
Vendor Name	IBM Corporation
Country Name	US
Product Name (Original)	z/OS
Product version (Original)	V1R5
Product Description (Original)	Highly secure scalable high-performance enterprise operating system
Product Name (Update)	
Product version (Update)	
Product Description (Update)	
Product Category	Host
Applied date	20031217
Application ID	US-20031217-000136
Current Status	Approved
Certificated Date	20040326

e-business

Why do we need a new Internet protocol?

IBM

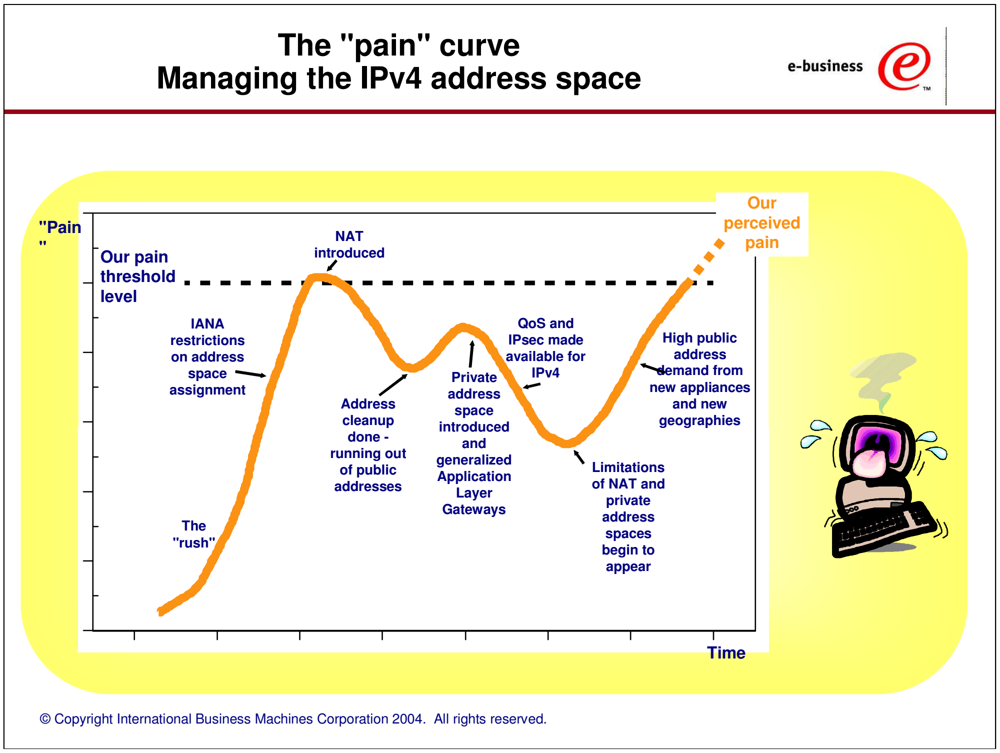

Shortage of IPv4 addresses has led to extensive use of dynamic addresses and private (not globally reachable) addresses
Requires Network Address Translators (NATs) at Intranet/Internet boundaries
NATs delay, but do not obviate need for IPv6
Pain of NATs depends on one's perspective

NATs are a barrier to continued Internet scaling
NATs break protocols (H.323, FTP, IPsec, DRDA, etc.) which rely on globally unique addresses

NATs have operational and administrative scaling problems

Always-on devices need permanent, global addresses (NATs prevent this)
Barrier to deployment of new types of

Visible IPv4 hosts on the Internet through the last 10 years

e-business

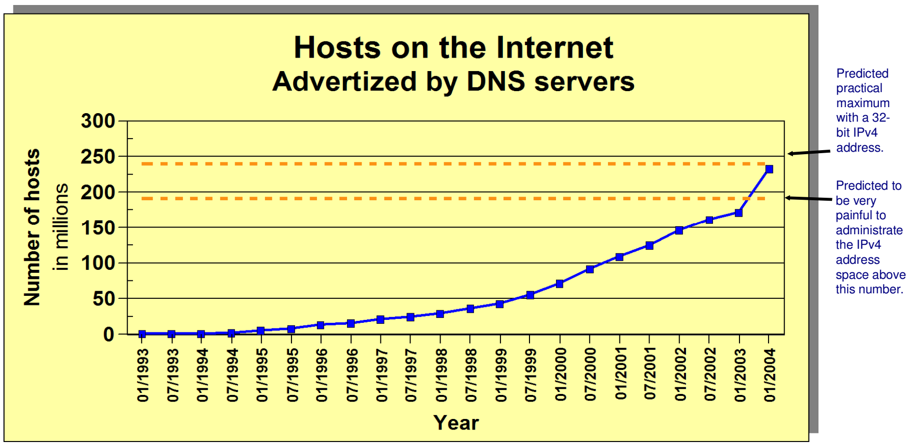

Source: http://www.isc.org/ds and "*An update on the H ratio*", RFC3194

These figures show 'visible' hosts - hosts behind firewalls are not visible and not counted in these figures.

The actual number of hosts that have access to information on the Internet is probably between 50 and 100 million.

Secure IP connectivity for anyone from anywhere to anything!

e-business

- **Growing mobility of users**
  - Internet access from anywhere (car, home, office)
  - Multiple addresses per person
  - Pervasive Computing
- **Continued rapid growth of the Internet**
  - China plans to roll out ~1 billion Internet nodes, starting with a 320 million student educational network
  - Asia/Pacific, and to a lesser extent Europe, missed out on the early IPv4 address allocations
- **Government support**
  - Wide-scale IPv6 promotion underway in Japan, Korea, and Taiwan
  - European Commission (EC) encourages IPv6 research, education, and adoption in member countries
  - Government agencies beginning to mandate IPv6 capable technology
- **Convergence of voice, video and data on IP**
  - Need for reliable and scalable architecture
  - "Always-on Connections"
- **New application opportunities**
  - Potentially unlimited number of IP nodes (such as vehicles, devices, components, and individual parts)
- **Security becomes more and more important**
  - Various optional security features have been patched on top of IPv4
  - IPv6 has security features defined as part of the base protocol

The Internet - a worldwide digital utility.

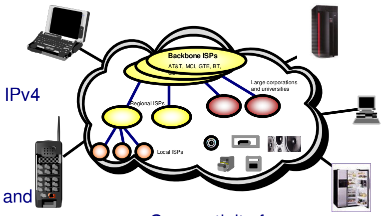

Connectivity for ***anyone*** from ***anywhere*** (car, plane, home, office) to ***anything***!

IPv6 promises true end-to-end connectivity for peer-based collaborative solutions.

The Third Generation Partnership Project (3GPP), which is responsible for the standardization of the third-generation mobile networks, has designated the session initiation protocol (SIP) as the call control protocol and Internet protocol version 6 (IPv6) as the only network protocol for 3G IP-based wireless networks.

Couldn't we just add more Network Address Translating (NAT) firewalls to deal with the limited number of IPv4 addresses?

e-business

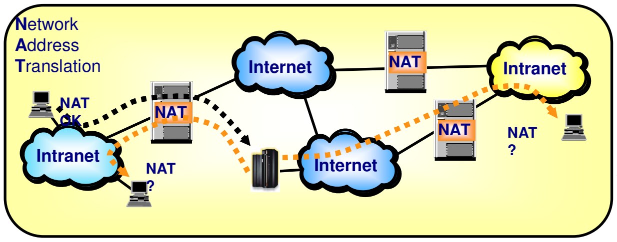

NATs work best in small end-sites, client-only

- All connections originate from clients (outbound only)
- Only a subset of clients need Internet access at any point in time
- A pool of public addresses matching the number of clients who need concurrent Internet access need to be available even when NAT is used
- Little configuration/administration is needed
- Limited applicability

When clients are servers (inbound connections):

- Static NATing (manual configuration on NAT device)
- If most/all clients are servers, NAT multiplexing premise fails

- Shortage of IPv4 addresses has led to extensive use of private (not globally reachable) addresses
  - Requires Network Address Translators (NATs) or application layer gateways at intranet/Internet boundaries
  - Every NAT node between a private and a public network needs a pool of public IP addresses - adding more NAT nodes requires more public IP addresses
- NATs are a pain to design around and are generally a severe barrier to continued Internet scaling
  - NATs break protocols (FTP, IPSec, DRDA, EE, etc.) that rely on globally unique addresses
  - NATs are very often sensitive to application data being encrypted (SSL/TLS, IPSec, Kerberos)
  - NATs have operational and administrative scaling problems
  - Always-on devices need permanent, global addresses (NATs prevent this)
  - Barrier to deployment of new types of applications (true peer-to-peer)
  - Convergence of voice, video, and data over IP
- IPv6 alleviates these problems and removes barriers to continued Internet expansion

Why has deployment been slow so far?

e-business

- Economic slowdown has slowed growth and spending
  - •Network infrastructure vendors are not introducing new products quickly
  - •Service providers are not upgrading and expanding networks
- IPv6 upgrades to network infrastructure are expensive
  - •IPv6 routing performance requires hardware upgrades
  - •New technology requires staff training
  - •New code/additional complexity will cause added support burdens
  - •No current revenue stream to justify the costs

- Major technology markets are comfortable with IPv4
  - •US and Europe have (relatively) many IPv4 addresses
  - •Address shortages have been mitigated by the use of NAT
- Benefits of IPv6 are not widely understood or not compelling
  - •Desire that it solves more problems (e.g., multihoming)
- Need critical mass of IPv6 peers for tangible benefits
  - •Chicken and egg problem; limited incentive for legacy IPv4 sites
  - •Deployments of new devices and associated new infrastructure do not have these constraints
  - •ISPs will not move until pressured to do so by customers
- Potential for rapid adoption when critical mass is reached
  - •Applications + Middleware + Infrastructure (OS, routers)
  - •A few big customers will show the way

e-business

Brief technical introduction to what IPv6 is

IBM

What is IPv6?

e-business

- **IPv6 is an evolution of the current version of IP, which is known as IPv4**
  - Work on new IETF standard started in early 90's under the name IPng (IP next generation)
  - Not backward compatible, but migration techniques defined
- **Today's IPv4 has 32-bit addresses**
  - Theoretical limit is 4,294,967,295 addresses
  - Practical limit is significantly less - predictions range from around 250,000,000 to 1,000,000,000
- **IPv6 provides almost unlimited number of addresses**
  - IPv6 addresses are 128 bits
  - No practical limit on global addressability
  - Enough address space to meet all imaginable needs for the whole world and for generations to come
  - More addresses *cannot* be retrofitted into IPv4
- **Other impartant improvements:**
  - Facilities for automatic configuration
  - Improved support for site renumbering
  - End to end IP security
  - Mobility with route optimization (important for wireless)
  - Miscellaneous minor improvements

IPv4 Address:
9.67.122.66

IPv6 Address:
2001:0DB8:4545:2::09ff:fef7:62dc

An added advantage of IPv6:

- You don't have to try and remember those IP addresses any longer - IPv6 addresses are plain impossible to remember!
- A DNS infrastructure is an absolute requirement in an IPv6 environment.

Expanded routing and addressing

e-business

Expanded size of IP address space

- Address space increased to 128 bits
  - Provides 340,282,366,920,938,463,463,374,607,431,768,211,456 addresses

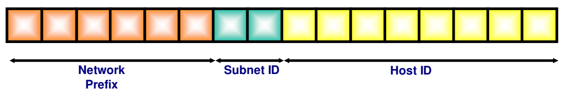

Prefix-bits designate address scope:

- **link-local scope**
- **site-local scope**
- **global scope**

- As important as the expanded address space is the use of hierarchical address formats:
  - Allocation architecture
    - Network portion of address allocated by ISPs
    - Subnet portion of address is allocated by customer
    - Host Identifier is derived from the MAC address of the interface adapter
  - Facilitates efficient routing architectures
    - IPv6 uses CIDR (Classless InterDomain Routing), first introduced in IPv4
  - IPv6 hierarchical routing likely only viable method for keeping the size of the backbone router tables under control
    - Even with hierarchical routing, the current IPv4 Internet backbone maintains 90,000 or more routes

IPv6 address types and scopes

e-business

Unicast:

- Assigned to one interface. Packets destined for a unicast address are sent to only one node.
- Can be link-local scope, site-local scope or global scope

Multicast:

- Provides a means for a source to communicate with a group

Anycast:

- Allows the source to communicate with the closest member of a group

The first few bits of the address (format prefix) identify the address type and scope:

Address type	Binary prefix	IPv6 notation
Unspecified	00...0 (128 bits	::/128
Loopback	00...1 (128 bits)	::1/128
Multicast	11111111	FF00::/8
Link-local unicast	1111111010	FE80::/10
~~Site-local unicast~~	~~1111111011~~	~~FEC0::/10~~
Global unicast	(everything else	

Site-local addresses have been deprecated by the IETF, but they may still be part of some implementations.

Multicast addresses (groups) also have scopes

IPv6 header format

e-business

- Changes to the IP header
  - Header is fixed length
    - Optional headers daisy-chained
  - Address space quadrupled to 16 bytes
  - Fragmentation fields moved out of base header
  - Header checksum eliminated
    - Performed by the link layer
  - No hop-by-hop fragmentation
    - Path MTU discovery determines MTU size
  - Flow label added
  - Some fields renamed
    - Time to Live -> Hop Limit
    - Protocol -> Next Header
    - Type of Service -> Traffic Class

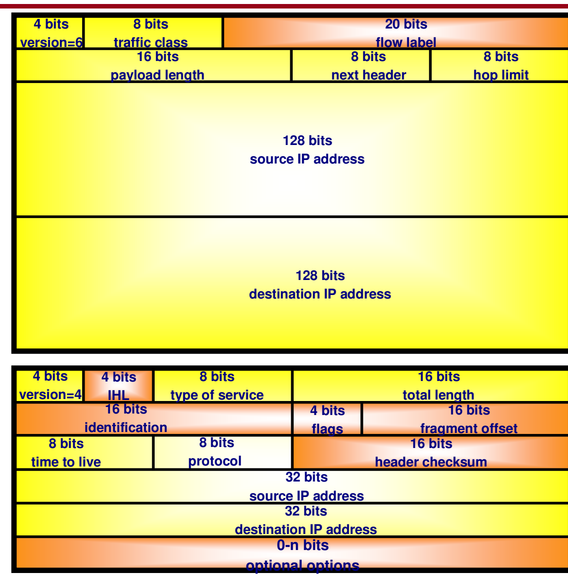

Stateless Address Autoconfiguration

e-business

- Address Configuration without separate DHCP server
  - Router *is* the server, advertising key address configuration information
- Address formed by combining routing prefix with Interface ID (IID)
- Link-local address configured when an interface is enabled
  - Allows immediate communication with devices on the local link
  - Primarily used for bootstrapping and discovery
  - Well-known prefix combined with locally-generated 64-bit IID
- Other addresses configured via Routing Advertisements
  - RA advertises 64-bit prefixes (e.g., on-link, form an address)
  - Public (e.g., server) addresses formed from IID

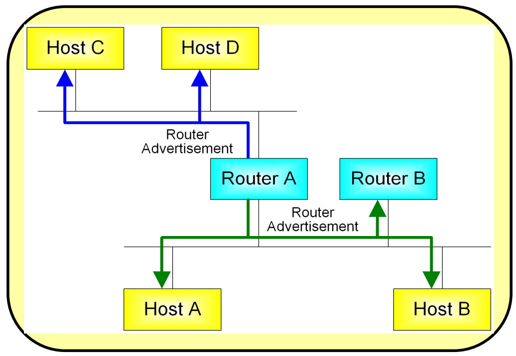

In this example, Router cisco4b is originating two ***unique*** *Router Advertisements, one onto Link A and one onto Link B.*

The Router Advertisement sent on Link A will inform all nodes on that link that:

cisco4b exists on Link A as a router

cisco4b either is or is not to be used as a default router on Link A

optionally, what prefixes exist on Link A and how they are to be used
what internet parameters (MTU and hop-limit) are configured for Link A

Similarly, the Router Advertisement on Link B will inform all nodes on that link:

cisco4b exists on Link B as a router

cisco4b either is or is not to be used as a default router on Link B

optionally, what prefixes exist on Link B and how they are to be used

what internet parameters (MTU and hop-limit) are configured for Link B

As a side note, this diagram demonstrates the ability of one OSA-E adapter being shared between an IPv4-only stack on one LPAR, and an IPv6-enabled stack on another LPAR.

Mobility With Route Optimization

e-business

- IPv6 includes enhanced support for mobile clients
  - All hosts include support for communicating directly with mobile clients, without having to send packets through an intermediate proxy
  - Avoids triangular routing problems found in IPv4
- Basic processing is as follows:
  - As the mobile client roams, it notifies its Home Agent on its current location by sending its Care of Address **(A)**.
  - When a server wishes to communicate with the mobile client and does not already know the current Care of Address, it sends a packet to the mobile client's home address **(B)**.
  - The Home Agent intercepts the packet and forwards it to the mobile client at its current Care of Address **(C)**.
  - The mobile client sends a response directly to the server, and includes its Care of Address in the packet **(D)**.
  - Subsequently, the client and server send packets directly to one another, without having to send packets through the Home Agent **(D)**, **(E)**.

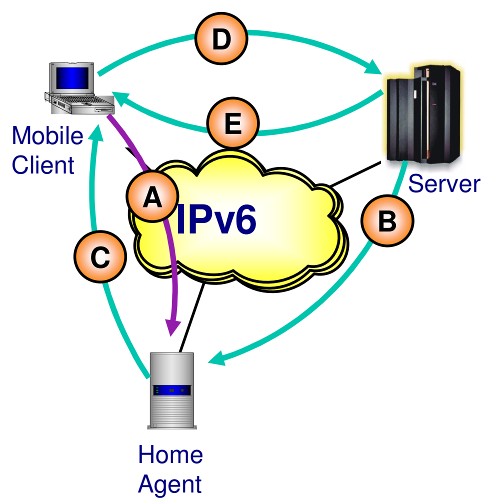

One of the major improvements found in IPv6 is the ehanced support for mobile clients
Mobile IPv6 avoids a major shortcoming found in IPv4 mobility support, namely that IPv4 requires all packets sent or received by an IPv4 mobile node to traverse a proxy server located at the home site of the IPv4 mobile node

In Mobile IPv6, the mobile client and a server are able to communicate directly with one another instead of through a proxy server
This not only improves performance, but dramatically improves the scalability of Mobile IPv6 support by placing a much lighter strain on the mobile IPv6 proxy nodes
And, with the expected number of mobile devices in an IPv6 enviornment, this is a necessary benefit for mobility support to work properly

e-business

IPv4 / IPv6 coexistence and migration

IPv4 to IPv6 Internet evolution - it won't happen overnight!

e-business

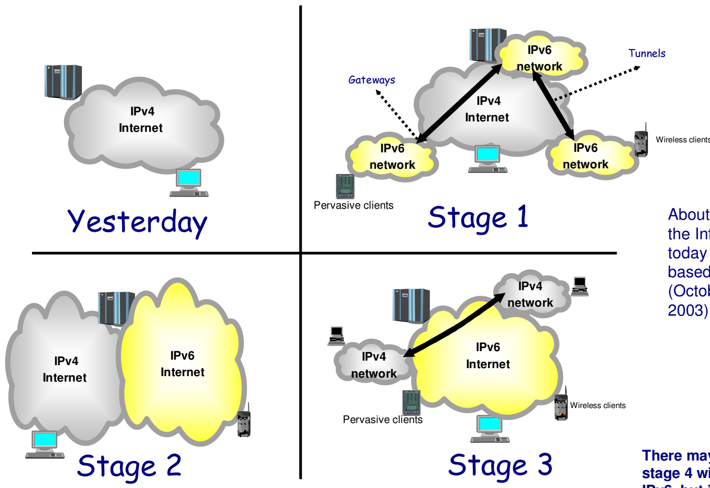

About 2% of the Internet is today IPv6-based (October 2003)

There may be a stage 4 with only IPv6, but it will take some years to get there.

General IPv4 - IPv6 migration issues

e-business

1 ***How do we share the physical network so that both IPv4 and IPv6 can be transported over one and the same physical network?***

- ?Dual-stack
- ?Tunneling of IPv6 over IPv4

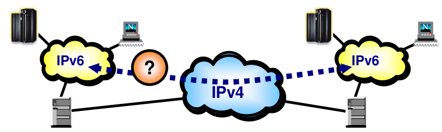

2 ***How do applications that have not yet been enhanced to support IPv6 communicate with applications that have been enhanced to support IPv6?***

- ?Dual-stack
- ?Application Layer Gateways (ALG)
- ?Network Address Translation - Protocol Translation (NAT-PT)
- ?Bump-In-the-Stack (BIS) or Bump-In-the-API (BIA)

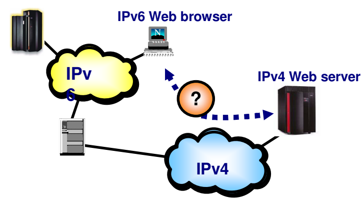

IPv6 over IPv4 network tunneling overview

e-business

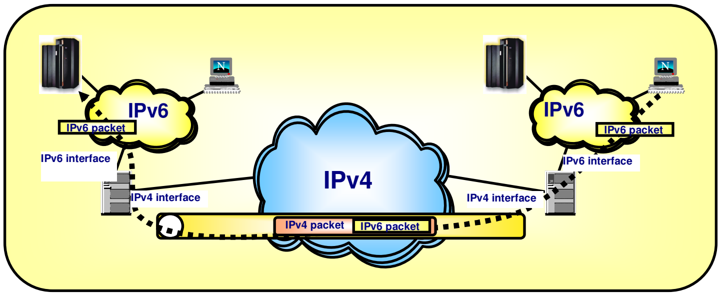

- ➢Tunneling: encapsulating an IPv6 packet in an IPv4 packet and send the IPv4 packet to the other tunnel endpoint IPv4 address.
- ➢Requires applications on both endpoints to use AF_INET6 sockets
- ➢Tunnels endpoints can be in hosts or routers
  - •The tunnel endpoint may be an intermediate node, the final endpoint, or a mixture of the two
- ➢The tunnel endpoint placement depends on connectivity needs
  - •Placing endpoints in routers allows entire sites to be connected over an IPv4 network
  - •Placing endpoints in hosts allows access to remote IPv6 networks without requiring updates to the routing infrastructure

Generalized dual-mode TCP/IP structure

e-business

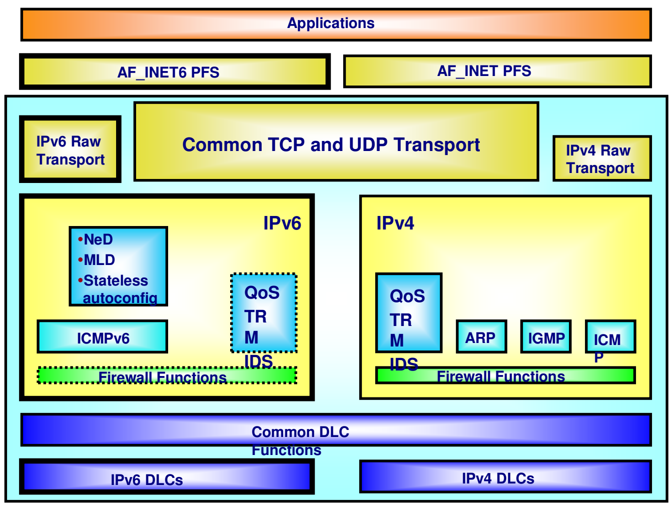

A dual-mode (or dual-stack) TCP/IP implementation supports both IPv4 and IPv6 interfaces - and both AF_INET and AF_INET6 applications.

The dual-mode TCP/IP implementation is a key technology for IPv4 and IPv6 coexistance in an internet.

z/OS uses the dual-mode implementation - a single TCP/IP address space that handles both IPv4 and IPv6.

Communication between IPv6 nodes and IPv4 nodes or applications

- Tools which enable communication between IPv6 nodes and IPv4 nodes or applications typically involve some form of translation
- This translation can be performed at the IP, transport, or application layer
  - At the IP layer, Simple IP/ICMP Translator (SIIT) may be used
    - Network Address Translator-Protocol Translator NAT-PT is built on top of SIIT
  - At the transport layer, SOCKS has been updated to allow IPv6/IPv4 relaying
    - The TCP or UDP connections are terminated at the boundary of the IPv6 domain and relayed to the IPv4 domain
  - At the application layer, proxies (sometimes referred to as Application Layer Gateways or ALGs) can be run on dual mode stacks

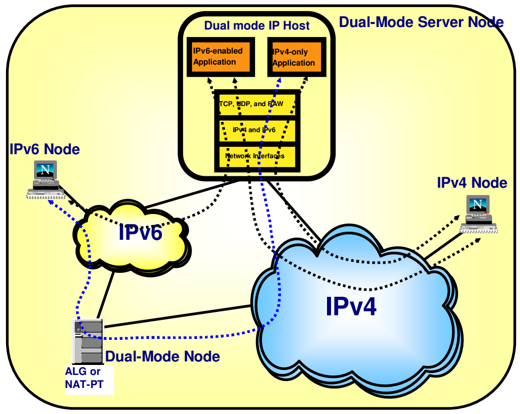

e-business

IPv6 on z/OS specifics

IBM

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

e-business

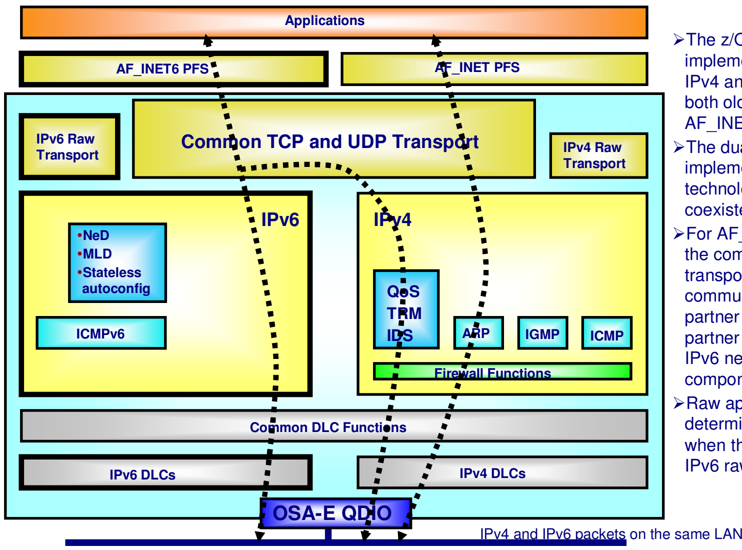

- The z/OS dual-mode TCP/IP implementation supports both IPv4 and IPv6 interfaces - and both old AF_INET and new AF_INET6 applications.
- The dual-mode TCP/IP implementation is a key technology for IPv4 and IPv6 coexistence in an internet.
- For AF_INET6 applications, the common TCP or UDP transport 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

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

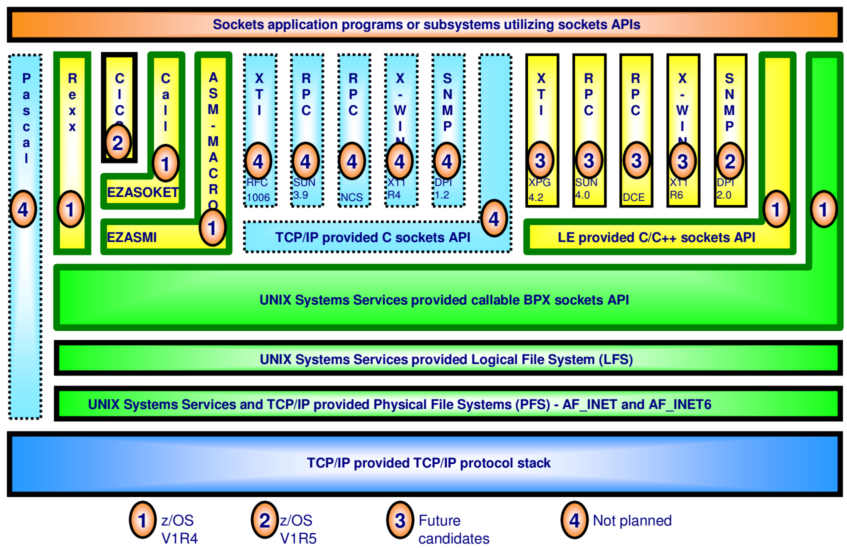

IPv6 network interface support

e-business

IPv6 network interfaces supported by CS in z/OS V1R5:

- ✓IPv6 Loopback interface
- ✓OSA-Express QDIO interface (zSeries hardware)
  - Gigabit Ethernet
  - Fast Ethernet
- ✓IUTSAMEHOST to other stacks in same LPAR
- ✓XCF to other stacks in same Sysplex
- ✓ESCON/FICON (MPCPTP) to another z/OS image (not to any known Channel-attached Routers)

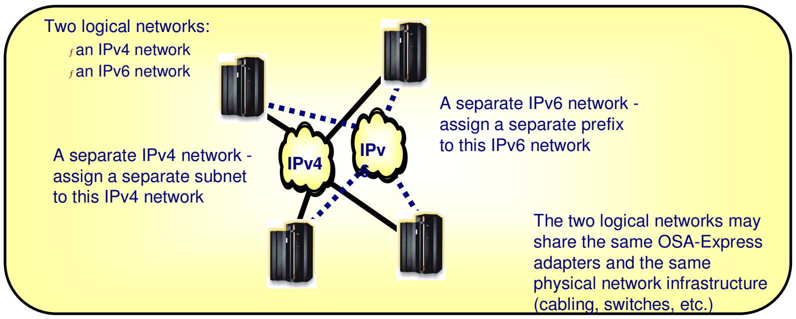

IPv6 Sysplex support added in z/OS V1R6

- IPv6 Dynamic VIPA support
  - Up to 1024 IPv6 Dynamic VIPA addressed per stack
- IPv6 support for stack managed Dynamic VIPA addresses (VIPADEFINE/VIPABACKUP)
- IPv6 support for application-specific Dynamic VIPA addresses (VIPARANGE)
- IPv6 support for distributed Dynamic VIPA addresses and distribution of IPv6 workload by Sysplex Distributor (VIPADISTRIBUTE):
  - WLM-based distribution
  - Round-robin distribution
  - Server affinity
  - Passive mode FTP support
  - Fast connection reset support
- IPv6 support for Sysplex sockets
- IPv6 support for source VIPA address use:
  - Interface-based selection of source VIPA
  - Sysplex-wide source VIPA addresses
  - Job-specific source VIPA
- SNMP MIB support for IPv6 dynamic VIPA addresses

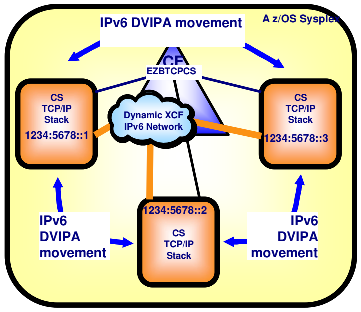

Dynamic routing and network management

e-business

➢ **OMPROUTE has in z/OS V1R5 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
- OSPF for IPv6 added in z/OS V1R6

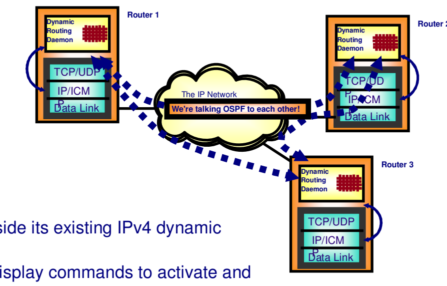

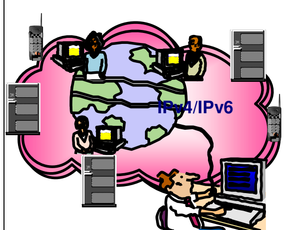

➢ **Network management SNMP support**
- Support SNMP agent (OSNMPD)
- IPv6 MIB support - new RFC drafts have been published that define IP version neutral objects
  - RFC2011 (IP and ICMP)
  - RFC2012 (TCP)
  - RFC2096 (IP routes)
  - RFC2233 (Interfaces) - this one is not version neutral

➢ **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

How do we enable IPv6 support on z/OS and what are the consequences?

e-business

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.

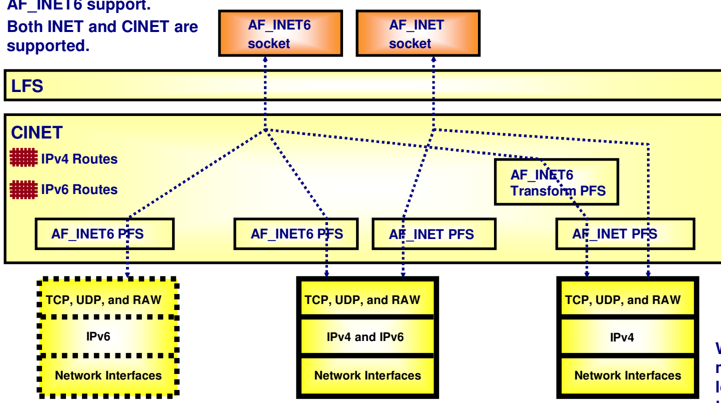

IPv6-only TCP/IP Stack

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

e-business

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

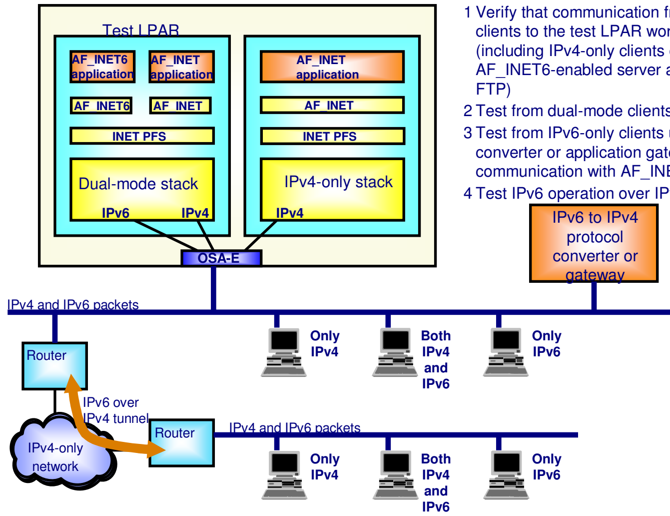

Steps for moving to an IPv6 environment

e-business

1 **Network access**

- A LAN can carry both IPv4 and IPv6 packets over the same media
- An OSA-E port can be used for both IPv4 and IPv6
- Update TCP/IP Profile to include the INTERFACE statement(s) for any IPv6 interfaces
- For LPAR-LPAR communication for IPv6, this must be accomplished via intermediate LAN (both LPARs using QDIO to a shared LAN)

2 **IPv6 address selection**

- Obtain an address block from your ISP, or use one of your IPv4 addresses to create a 6to4 prefix
- For test purposes, site-local IPv6 addresses is sufficient, but avoid using them in production
- IPv6 addresses can be assigned to the IPv6 Interfaces and static VIPAs
- Addresses can be manually configured on the INTERFACE statement in the TCP/IP Profile or autoconfigured using Neighbor Discovery Stateless Autoconfiguration (VIPA addresses must be manually configured)

3 **DNS setup**

- DNS BIND 9 Name Server can be used for both IPv4 and IPv6 resources
- Continue to use the existing host name for IPv4 connectivity to avoid possible disruption in network connectivity and IPv4-only applications on an IPv6-enabled stack
- Create a new host name to be used for IPv6 and IPv4 connectivity
- Optionally, a third host name which may be used only for IPv6 can be configured
- If using stateless autoconfiguration to define IPv6 addresses, static VIPA addresses should be stored in DNS since the autoconfigured addresses will change over time and no Dynamic DNS support is available on z/OS

Steps for moving to an IPv6 environment

e-business

4 **INET or Common INET**

- Both are supported for IPv6, but INET is simplier
- Running IPv4 and dual-mode stacks under CINET is not recommended - run dual-mode stacks in a separate LPAR from IPv4 only stacks
- AF_INET6 NETWORK statement must be coded in BPXPRMxx before starting IPv6-enabled stacks

5 **Selection and placement of IPv6 to IPv4 protocol converter or application gateway**

- z/OS does not implement any functions that will allow IPv6-only nodes to communicate with z/OS-resident AF_INET applications, so an outboard protocol converter or application-layer gateway component may be needed
- This component will only be needed if the test configuration includes IPv6-only platforms
- Various technologies are being made available by various vendors; SOCKS64 seems the simplest technology right now

6 **Connectivity to non-local IPv6 locations**

- Tunneling may be needed between a router connected to the LAN that z/OS is connected to, and a router at another location where IPv6 test equipment is located

The Journey to IPv6 for z/OS Communications Server

e-business

IPv6 deployment phases

–The first phase (z/OS V1R4)

- Stack support for IPv6 base functions - (APIs, Protocol layers)
- Resolver
- High speed attach (OSA Express QDIO))
- Service tools (Trace, Dump, etc.)
- Configuration and netstat, ping, traceroute, SMF
- Static Routing
- FTP, otelnetd,unix rexec, unix rshd/rexecd

–The second phase (z/OS 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 w/ OMPROUTE (only RIPng)

–The third phase (z/OS V1R6)

- Sysplex Exploitation (Dynamic VIPA, Sysplex Distributor functions)
- Dynamic Routing Protocol w/ OMPROUTE (OSPFv3)
- Additonal Network Management MIBs

–After z/OS V1R6

- Integrated IPSec
- HiperSockets DLC
- Advanced Socket APIs
- Extended Stats MIB, OSPFv3 MIB
- Intrusion Detection Services
- IPv6 mobility support

Objective is to have IPv6 production ready on the platform when you need it!

The Internet - a worldwide digital utility.

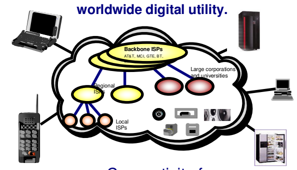

Connectivity for ***anyone*** from ***anywhere*** (car, plane, home, office) to ***anything***!

Trademarks, Copyrights, and Disclaimers

e-business