#SHAREorg

SHARE
Technology · Connections · Results

IPv6 on z/OS – Part 1

Mike Fitzpatrick – mfitz@us.ibm.com
IBM Raleigh, NC

Thursday, March 15th, 8:00am
Session: 10830

Trademarks, notices, and disclaimers

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

- Advanced Peer-to-Peer Networking®
- AIX®
- alphaWorks®
- AnyNet®
- AS/400®
- BladeCenter®
- Candle®
- CICS®
- DataPower®
- DB2 Connect
- DB2®
- DRDA®
- e-business on demand®
- e-business (logo)
- e business(logo)®
- ESCON®
- FICON®
- GDDM®
- GDPS®
- Geographically Dispersed Parallel Sysplex
- HiperSockets
- HPR Channel Connectivity
- HyperSwap
- i5/OS (logo)
- i5/OS®
- IBM eServer
- IBM (logo)®
- IBM®
- IBM zEnterprise™ System
- IMS
- InfiniBand ®
- IP PrintWay
- IPDS
- iSeries
- LANDP®
- Language Environment®
- MQSeries®
- MVS
- NetView®
- OMEGAMON®
- Open Power
- OpenPower
- Operating System/2®
- Operating System/400®
- OS/2®
- OS/390®
- OS/400®
- Parallel Sysplex®
- POWER®
- POWER7®
- PowerVM
- PR/SM
- pSeries®
- RACF®
- Rational Suite®
- Rational®
- Redbooks
- Redbooks (logo)
- Sysplex Timer®
- System i5
- System p5
- System x®
- System z®
- System z9®
- System z10
- Tivoli (logo)®
- Tivoli®
- VTAM®
- WebSphere®
- xSeries®
- z9®
- z10 BC
- z10 EC
- zEnterprise
- zSeries®
- z/Architecture
- z/OS®
- z/VM®
- z/VSE

* All other products may be trademarks or registered trademarks of their respective companies.

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

- 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 there from.
- 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.
- InfiniBand is a trademark and service mark of the InfiniBand Trade Association.
- 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.

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.

Other IPv6 Sessions at this SHARE

10395: IPv6 Tunneling Technologies	Wednesday, March 14th: 1:30pm
10397: IPv6 Basics	Wednesday, March 14th: 8:00am
10396: IPv6 Addressing	Wednesday, March 14th: 9:30am
10398: IPv6 Design	Wednesday, March 14th: 3:00pm
10399: IPv6 Implementation	Wednesday, March 14th: 4:30pm
10400: IPv6 Planning	Wednesday, March 14th: 11:00am
10401: Transitioning to IPv6	Thursday, March 15th: 3:00pm
10413: Common IPv6 Mistakes	Friday, March 16th: 8:00am
10414: IPv6 Deep Dive	Friday, March 16th: 11:00am
10831: IPv6 on z/OS - Part 2	Thursday, March 15th: 9:30am
10836: IPv6 Configuration on z: Hands-on Lab	Thursday, March 15th: 4:30pm

Agenda

- Why is IPv6 important?
- What is IPv6?
- IPv6 penetration
- Coexistence and migration

Disclaimer: All statements regarding IBM future direction or intent, including current product plans, are subject to change or withdrawal without notice and represent goals and objectives only. All information is provided for informational purposes only, on an "as is" basis, without warranty of any kind.

Why is IPv6 important?

When do our z/OS customers believe they will need IPv6?

- The majority of z/OS customers do not know
  - Expectations are that it will be needed slightly earlier on other platforms than z/OS
- It is time to start thinking, learning, and preparing ***now*** !

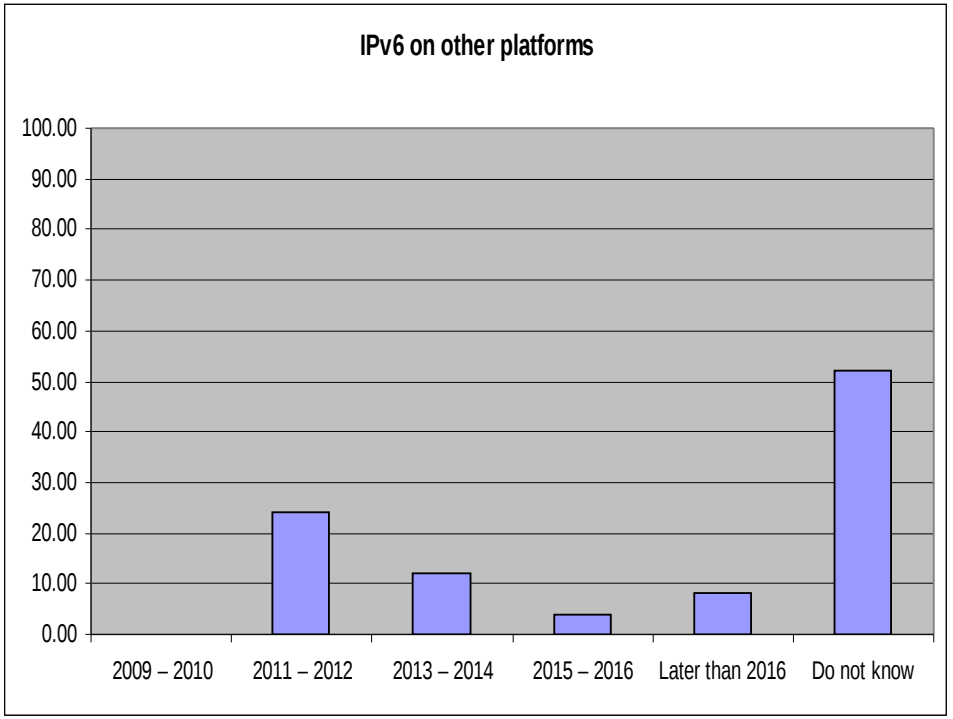

Source: Survey conducted by ENS early 2009 among a selected set of customers (39 responses to this question)

IPv4 address usage since early 1993

- ~~Projected~~ Internet Assigned Numbers Authority (IANA) Unallocated Address Pool Exhaustion
  - Feb 2011
- Projected Regional Internet Registries (RIR) Unallocated Address Pool Exhaustion
  - Apr 2011 – Apr 2014
- z/OS Communications Server continues to focus on IPv6 standards currency
  - US DoD/NIST
  - IPv6 Forum

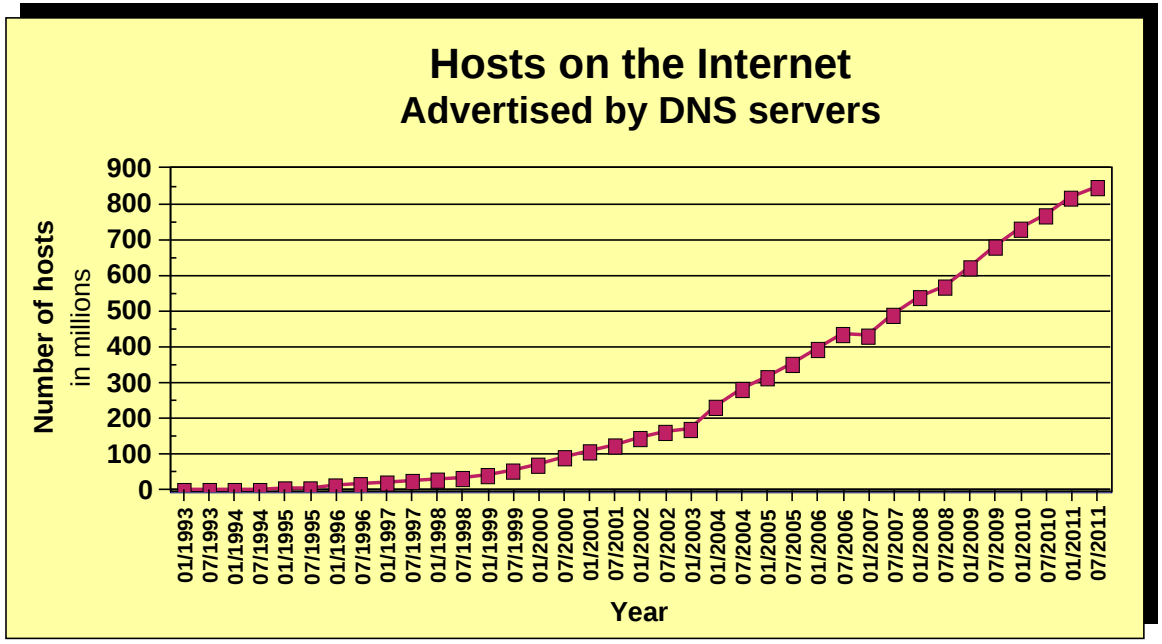

- **What is the upper practical limit (the ultimate pain threshold) for number of assigned IPv4 addresses? Some predictions said 250,000,000 (250 million), others go up to 1,000,000,000 (one billion or one milliard).**
- **Source: https://www.isc.org/solutions/survey**
- **Source: http://www.potaroo.net/tools/ipv4/index.html**
- **Source: http://penrose.uk6x.com/**

If you want to stay in business after 2011/2012, you'd better start paying attention!

Do not worry too much; the sky isn't falling – IPv4 and IPv6 will coexist for many years to come. Your applications need to be able to use both. If you write directly to the TCP/IP sockets layer, you need to start changing those applications.

When is Doomsday going to be here?

http://www.potaroo.net/tools/ipv4/index.html

IPv4 Address Report

This report generated at 30-Jan-2012 07:59 UTC.

IANA Unallocated Address Pool Exhaustion: 03-Feb-2011

Projected RIR Address Pool Exhaustion Dates:

RIR	Projected Exhaustion Date
APNIC:	19-Apr-2011
RIPENCC:	27-Jul-2012
ARIN:	19-Jul-2013
LACNIC:	29-Jan-2014
AFRINIC:	27-Oct-2014

This is no longer a future long term concern!!!!

IPv4 Address exhaustion awareness is a hot topic

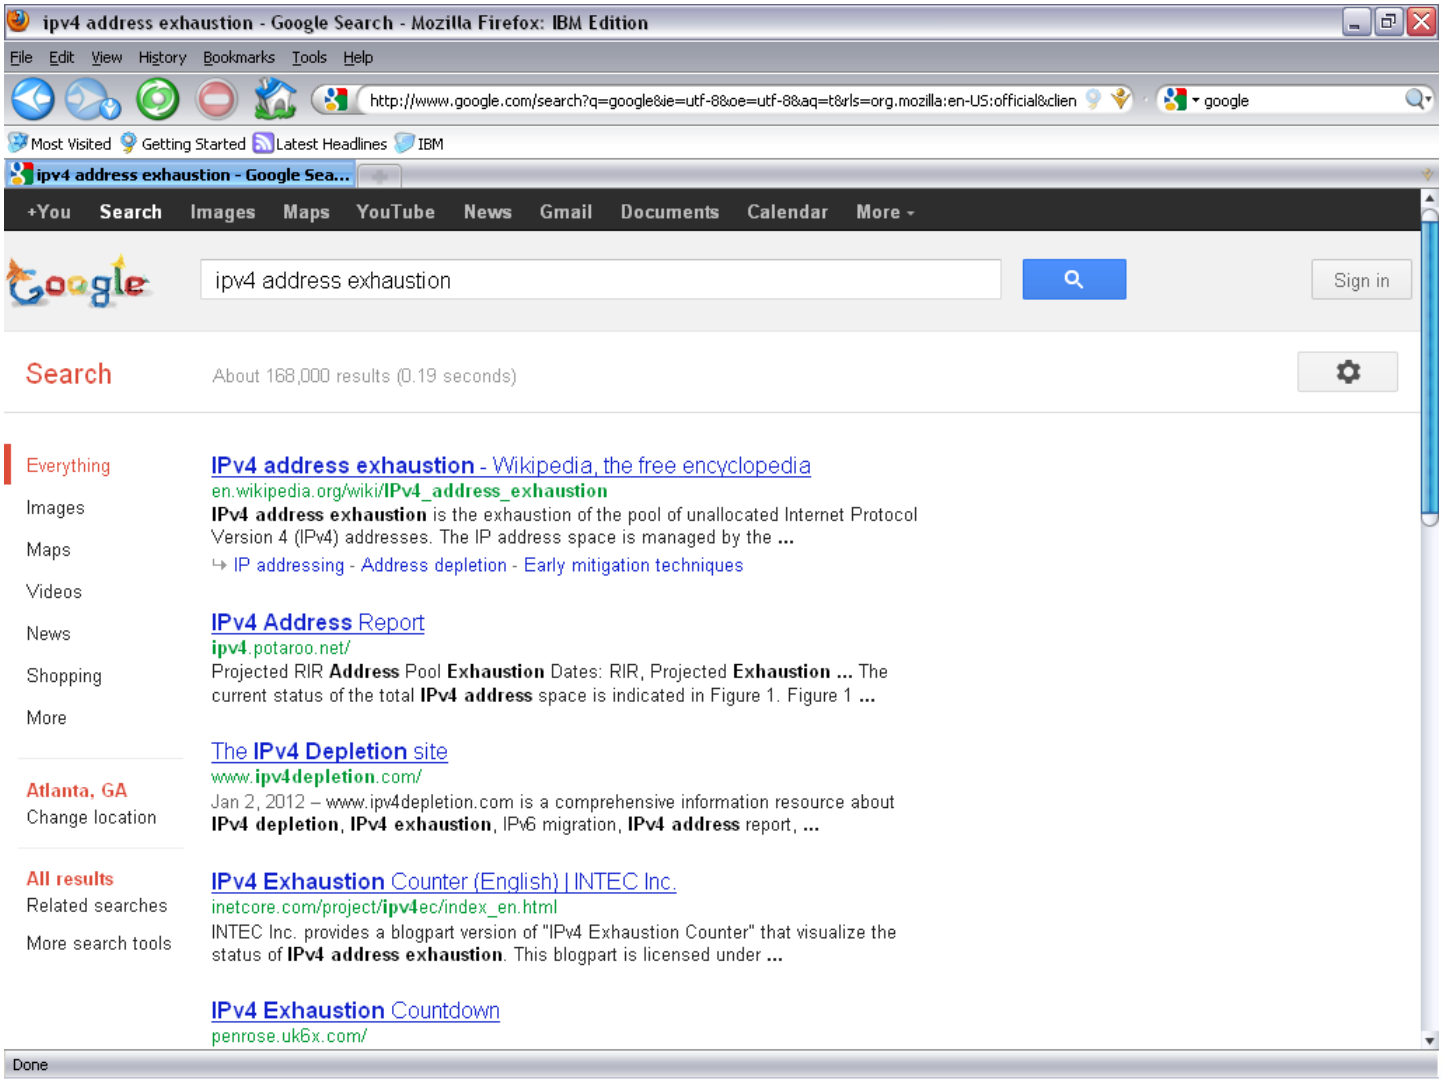

Network World – early May 2010

How the IPv4 address space is managed

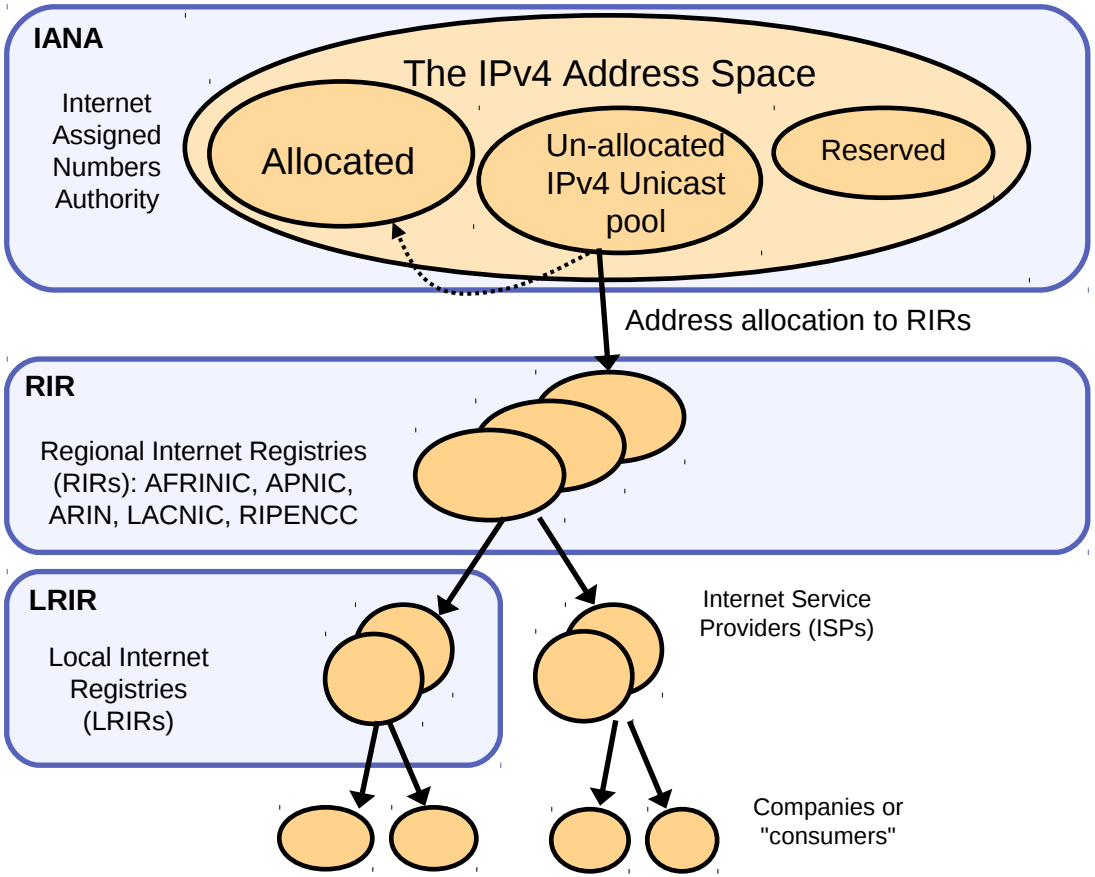

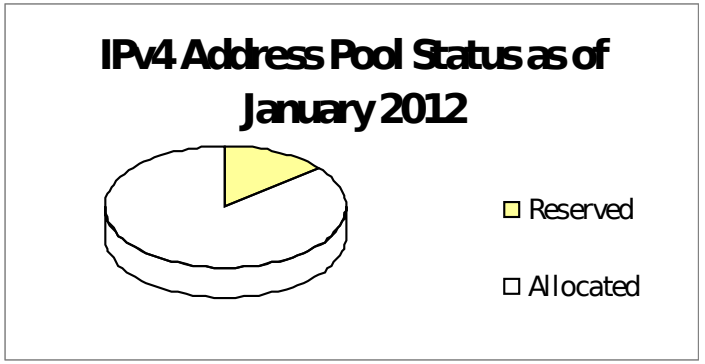

The IANA pool

of un-allocated addresses was exhausted in February 2011

Source: "IPv4 Address Report" - http://www.potaroo.net/tools/ipv4/

IPv4 address space data as of January 2012

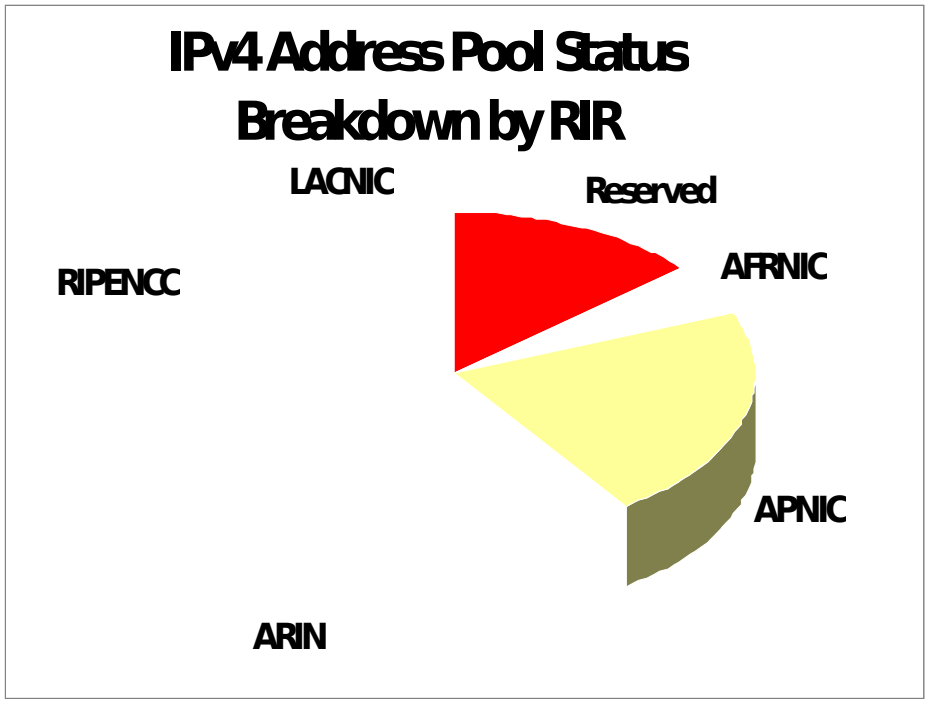

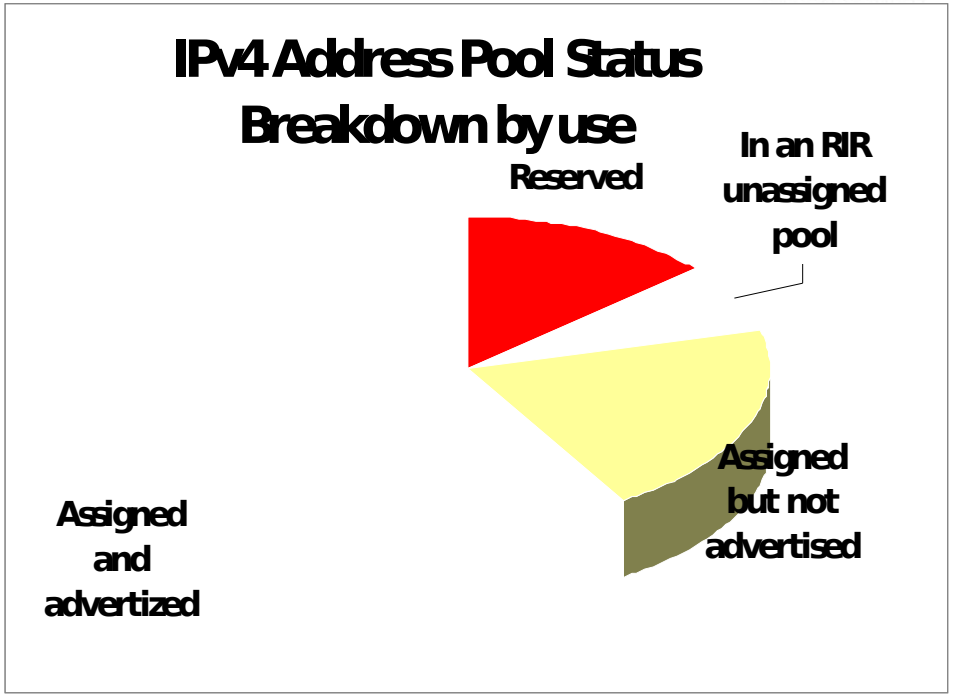

Reserved:	Reserved by the IETF
Un-allocated:	Available to be allocated to the RIRs ***(None available)***
AFRINIC:	Africa, portions of the Indian Ocean
APNIC:	Portions of Asia, portions of Oceania (includes Australia, China, India)
ARIN:	Canada, United States, islands in the Caribbean Sea and North Atlantic Ocean
RIPENCC:	Europe, the Middle East, Central Asia
LACNIC:	Latin America, portions of the Caribbean

Why IPv6? - It's really simple: IPv4 addresses are running short!

- **Forget about fancy IPv6 features as a reason for moving to IPv6**
  - (Some of them are actually good!)
- **IPv6 deployment is inevitable**
  - Literally running out of IPv4 addresses
    - IPv4 address pool already exhausted
    - To minimize disruption, IPv6 needs to be in place and in actual use soon
  - No other credible alternative to IPv6
    - Only alternative is IPv4 with significant increase in NAT
      - *Increased use of private addresses and resulting address collisions*
      - *Complete loss of globally unique addressing*
      - *Even NAT requires pools of public IPv4 addresses*
- **All major vendors have maturing IPv6 product lines**
  - Most operating systems support IPv6, with middleware and application support starting to ship as well
  - Router vendors have supported IPv6 for several years
  - Both Windows VISTA and Windows 7 were IPv6-enabled “out of the box”

The Internet - a worldwide digital utility.

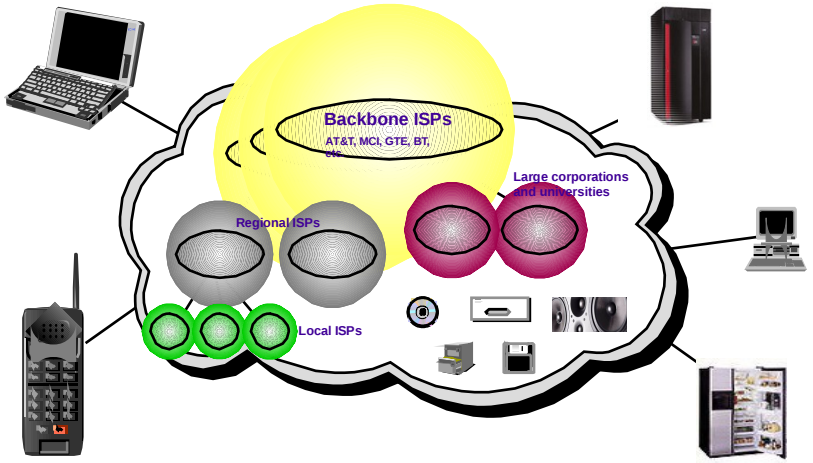

Connectivity for ***anyone*** from ***anywhere*** (car, plane, home, office) to ***anything***!

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

Current trends driving IPv6

- **Growing mobility of users**
  - Internet access from anywhere (car, airplane, 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
    - Network operations for 2008 Summer Olympics done solely on IPv6 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 China, Japan, Korea and Taiwan
  - European Commission (EC) encourages IPv6 research, education, and adoption in member countries
  - US Department of Defense - All platforms offered to DoD must meet very specific IPv6 capabilities
  - Other US government institutions through the National Institute of Standards and Technologies - NIST has also published detailed IPv6 compliance requirements

- **More and more "push" applications being deployed in the wireless market space.**
  - Clients subscribe to services that get pushed out by servers – requires public addresses for clients

- **Convergence of voice, video and data on IP**
  - Need for reliable and scalable architecture
  - “Always-on Connections"

What is IPv6?

So - what is IPv6?

- **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
  - Not backward compatible, but migration techniques defined
- **Today's IPv4 has 32 bit addresses**
  - Theoretical limit is around 4 billion addresses
  - Due to IPv4 address assignment structure and policies, the practical limit is less than 1 billion useable global addresses

IPv4 Address:
9.67.122.66

- **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 a while
  - More addresses *cannot* be retrofitted into IPv4

IPv6 Address:
2001:0DB8:4545:2::09FF:FEF7:62DC

- **Other improvements important, but to some extent secondary:**
  - Facilities for automatic configuration
  - Improved support for site renumbering
  - End to end IP security
  - Mobility with route optimization (important for wireless)
  - Miscellaneous improvements aimed at improving router performance

Important IPv6 technical features

- **IPv6 header and extensions header**
  - Streamlined IPv6 header
    - Fixed length to speed up forwarding processing in routers
  - Optional extension headers for fragmentation, security, etc.
- **Routers are no longer allowed (able) to fragment forwarded data-grams**
  - Path MTU discovery is always used
- **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
      - *3.4028 * 10**38*
    - Enough for many(!) addresses per person on the planet
  - A 64-bit subnet prefix identifies the link
  - Followed by a 64-bit Interface Identifier (IID)
- **IID may be derived from IEEE identifier (MAC address)**
  - Only leftmost 64 bits available for routing and "network addressing"
  - The rightmost 64-bits identify the host on the target link

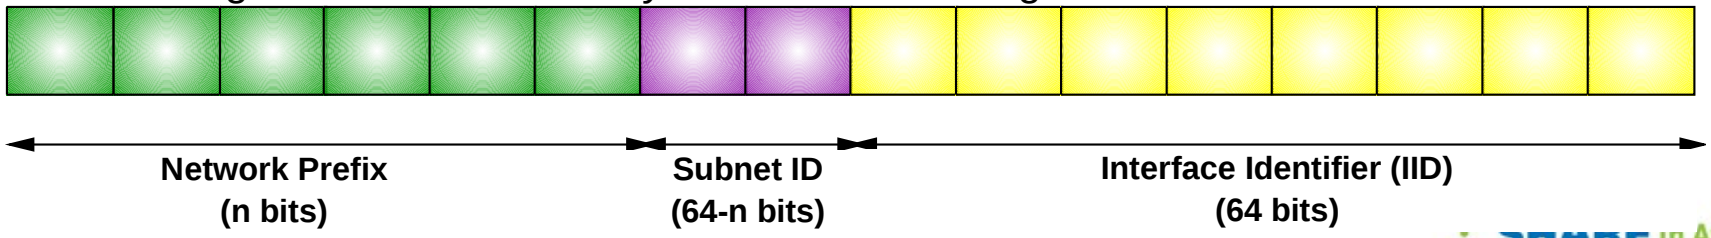

IPv6 address textual representation

- Addresses are represented as 8 segments of 4 hex digits (16 bits), separated by colons
  - **2001:0DB8:0:0:240:2BFF:FE3D:71AD**
- Two colons in a row can be used to denote one or more sets of zeroes, usually used between the prefix and the interface ID
  - **2001:0DB8::240:2BFF:FE3D:71AD**
- The prefix length can be indicated after a slash at the end
  - **2001:0DB8::240:2BFF:FE3D:71AD/64**
- A prefix alone is represented as if the interface ID bits are all zero
  - **2001:0DB8::/64**
- Obviously, this syntax may be a bit difficult for humans…..
  - Use of DNS/hostnames is no longer an optional convenience

Common IPv6 addresses and prefixes

- **::/128**
  - INADDR6_ANY (the unspecified address)
  - All zero address
- **::1/128**
  - IPv6 loopback address
- **FF00::/8**
  - Multicast addresses
- **FE80::/10**
  - Link-local addresses
- **FC00::/7**
  - Unique local addresses
- **::FFFF/96**
  - IPv4-Mapped IPv6 Address
- **2000::/3**
  - Current globally unique IPv6 address space (may change in the future)
  - In a sense, anything different from the ones above are to be considered globally unique addresses

8 bits	4 bits	4 bits	112 bits
11111111	Flags	scope	Group ID

10 bits	54 bits	64 bits
1111111010	0	Interface ID

7 bits	121 bits
1111110	Local address

80 bits	16 bits	32 bits
0	FFFF	IPv4 address

3 bits	45 bits	16 bits	64 bits
	Network prefix		
001	Network id	Subnet	Interface ID

IPv6 scoped unicast addressing

- Concept of scoped unicast addresses is part of architecture
- Link-local addresses for use on a single link
  - Primarily used for bootstrapping and infrastructure protocols such as Neighbor Discovery
  - Address = well-known link-local prefix plus node-generated IID
- Unique Local IPv6 Unicast addresses for use within a site
  - Like net 10 (not routable in the Internet backbone)
  - Site-local addresses
    - Part of early IPv6 standards -but introduced a lot of complexity
    - Has been deprecated by the IETF
- Global address prefixes are provided by ISPs

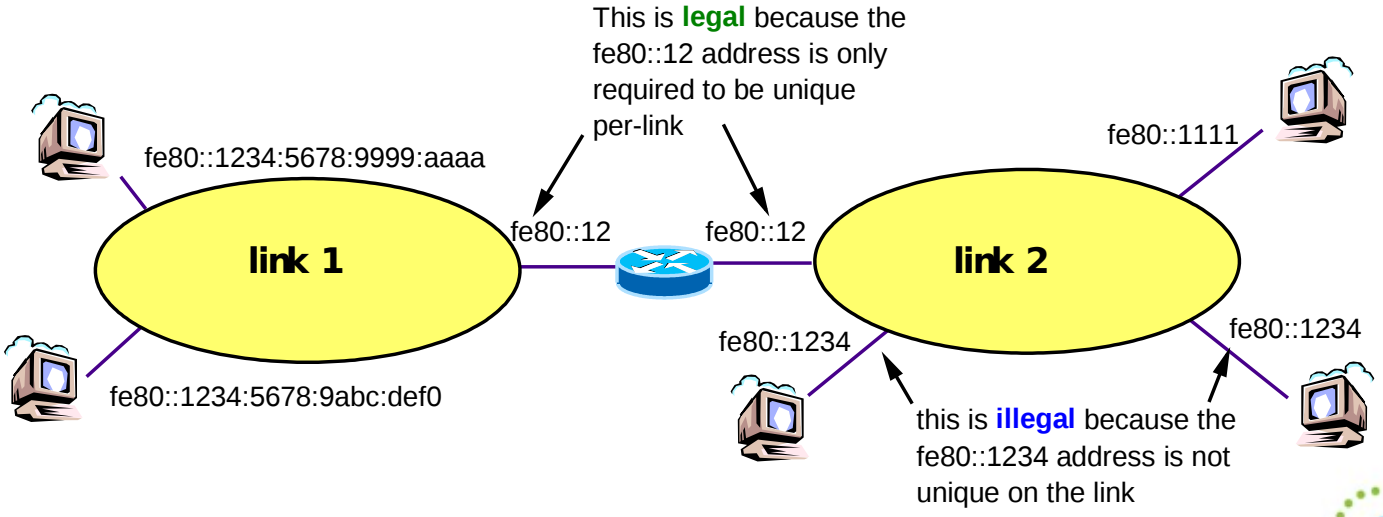

Stateless Address Auto-configuration and Neighbor Discovery

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

Router Discovery
- Router Solicitations and Router Advertisements used to find and keep track of neighboring routers
- Includes additional information for IP stack configuration

Address resolution
- Neighbor Solicitations and Neighbor Advertisements perform address resolution (i.e., ARP functions)

Neighbor Un-reachability Detection (NUD)
- Keep track of reachability of neighbors
- If path to router fails, switch to another router before TCP timeouts

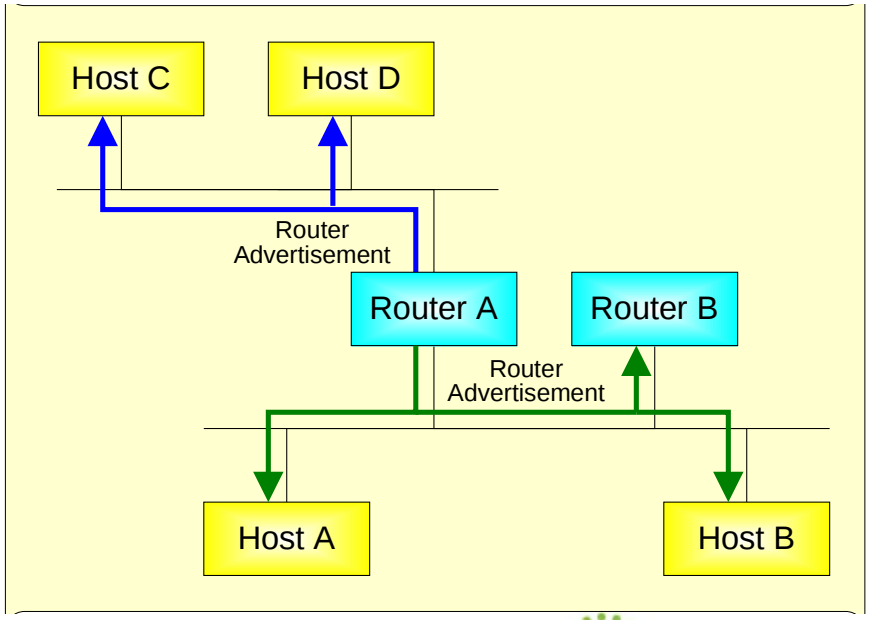

IPv6 penetration

Who is currently taking the lead on IPv6 deployment?

- **US Department of Defense (DoD)**
  - IPv6 compliance requirements detailed
  - All platforms offered to DoD must meet very specific IPv6 capabilities
- **Other US government institutions through the National Institute of Standards and Technologies**
  - NIST has also published detailed IPv6 compliance requirements
  - Generally platforms offered to any US government institution must meet these very specific IPv6 requirements
- **Worldwide, other organizations are closely looking at the IPv6 compliance tests done according to the IPv6 forum – the IPv6-Ready Phase 1 and Phase 2 logo certification (Tahi test suite)**
  - z/OS V1R10 is IPv6 certified according to the US DoD IPv6 requirements
  - z/OS V1R11 are IPv6-Ready Logo Phase 2 certified
- **Russia has begun developing similar IPv6 compliance requirements**
- **The European Union is trying to jump-start IPv6 deployment within the European Union**
  - ADVANCING THE INTERNET - Action Plan for the deployment of Internet Protocol version 6 (IPv6) in Europe (issued in 2008)
- **Japan and China have had operational IPv6 networks for a few years**
- **The mobile telephone (device) industry is moving beyond GSM into IMS (IP multimedia Subsystem)**
  - Agreed to by the industry that it has to be based on IPv6
- **The internal management network (INMN) in the zEnterprise System is IPv6 due to the facilities IPv6 offers in terms of auto configuration**
  - Cloud infrastructure solutions in general are assumed to move to IPv6 for the same reasons

IANA IPv6 Prefix allocations

- IANA assigns IPv6 prefixes to the Regional Internet Registries for both IPv4 and IPv6
- The table to the right lists the current IPv6 prefix allocation to the RIRs
  - You will typically be assigned a 48-bit prefix out of these
  - Note: they all currently start with binary 001
- The graph below shows how many Autonomous Systems (AS) currently announce IPv6 prefixes

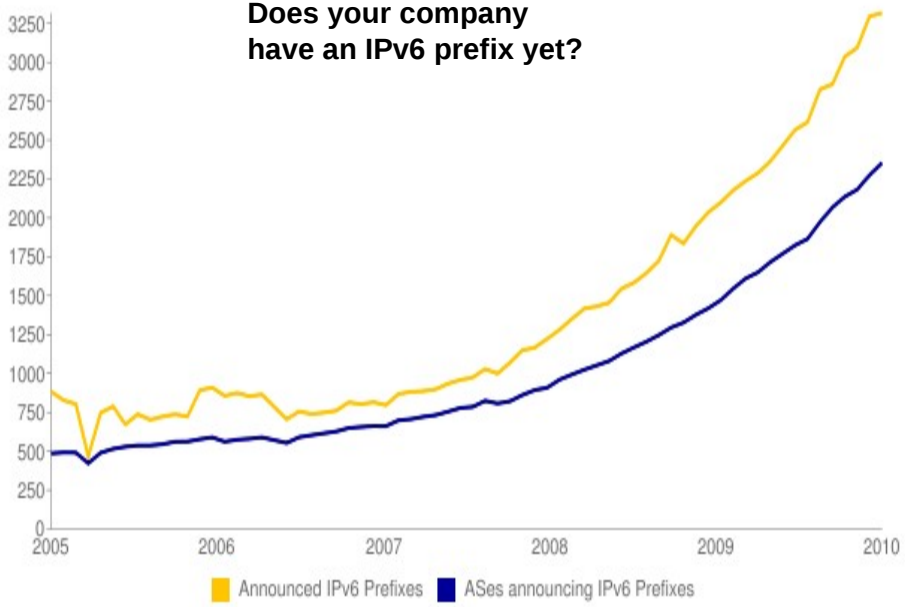

Prefix	Designation	Date	Status
2001:0000::/23	IANA	7/1/1999	ALLOCATED
2001:0200::/23	APNIC	7/1/1999	ALLOCATED
2001:0400::/23	ARIN	7/1/1999	ALLOCATED
2001:0600::/23	RIPENCC	7/1/1999	ALLOCATED
2001:0800::/23	RIPENCC	5/2/2002	ALLOCATED
2001:0A00::/23	RIPENCC	11/2/2002	ALLOCATED
2001:0C00::/23	APNIC	5/2/2002	ALLOCATED
2001:0E00::/23	APNIC	1/1/2003	ALLOCATED
2001:1200::/23	LACNIC	11/1/2002	ALLOCATED
2001:1400::/23	RIPENCC	2/1/2003	ALLOCATED
2001:1600::/23	RIPENCC	7/1/2003	ALLOCATED
2001:1800::/23	ARIN	4/1/2003	ALLOCATED
2001:1A00::/23	RIPENCC	1/1/2004	ALLOCATED
2001:1C00::/22	RIPENCC	5/4/2001	ALLOCATED
2001:2000::/20	RIPENCC	5/4/2001	ALLOCATED
2001:3000::/21	RIPENCC	5/4/2001	ALLOCATED
2001:3800::/22	RIPENCC	5/4/2001	ALLOCATED
2001:3C00::/22	IANA		RESERVED
2001:4000::/23	RIPENCC	6/11/2004	ALLOCATED
2001:4200::/23	AfriNIC	6/1/2004	ALLOCATED
2001:4400::/23	APNIC	6/11/2004	ALLOCATED
2001:4600::/23	RIPENCC	8/17/2004	ALLOCATED
2001:4800::/23	ARIN	8/24/2004	ALLOCATED
2001:4A00::/23	RIPENCC	10/15/2004	ALLOCATED
2001:4C00::/23	RIPENCC	12/17/2004	ALLOCATED
2001:5000::/20	RIPENCC	9/10/2004	ALLOCATED
2001:8000::/19	APNIC	11/30/2004	ALLOCATED
2001:A000::/20	APNIC	11/30/2004	ALLOCATED
2001:B000::/20	APNIC	3/8/2006	ALLOCATED
2002:0000::/16	6to4	2/1/2001	ALLOCATED
2003:0000::/18	RIPENCC	1/12/2005	ALLOCATED
2400:0000::/12	APNIC	10/3/2006	ALLOCATED
2600:0000::/12	ARIN	10/3/2006	ALLOCATED
2610:0000::/23	ARIN	11/17/2005	ALLOCATED
2620:0000::/23	ARIN	9/12/2006	ALLOCATED
2800:0000::/12	LACNIC	10/3/2006	ALLOCATED
2A00:0000::/12	RIPENCC	10/3/2006	ALLOCATED
2C00:0000::/12	AfriNIC	10/3/2006	ALLOCATED

IPv4 to IPv6 Internet evolution

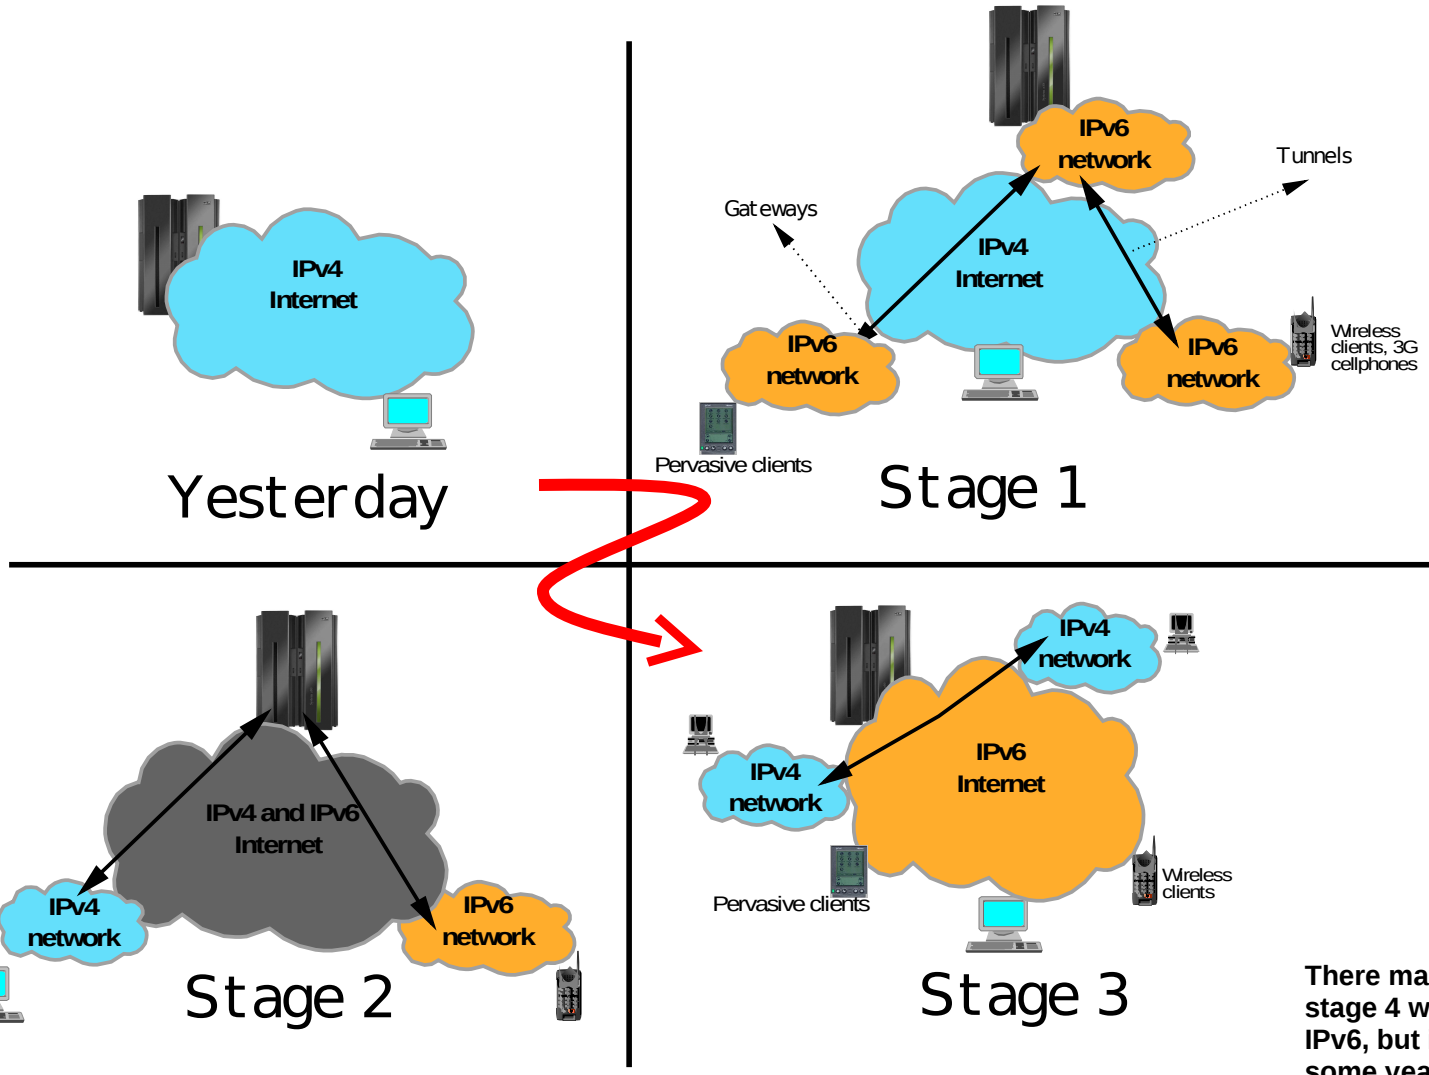

So where are we?
In stage 1, preparing to move to stage 2.

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

Coexistence and migration

General transition considerations

1

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

- Standard LAN technologies for multiplexing multiple network protocols over the same media
- Dual-mode stack (supports both IPv4 and IPv6 in one TCP/IP stack)
- Tunneling of IPv6 over IPv4

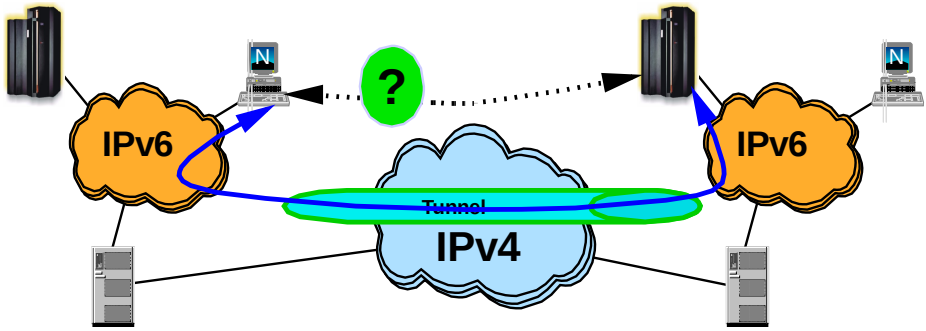

2

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

- Dual-mode stack
- Application Layer Gateways (ALG)
- Various other IPv6 transition technologies, such as ISATAP, 6to4 prefixes, SIIT, etc.

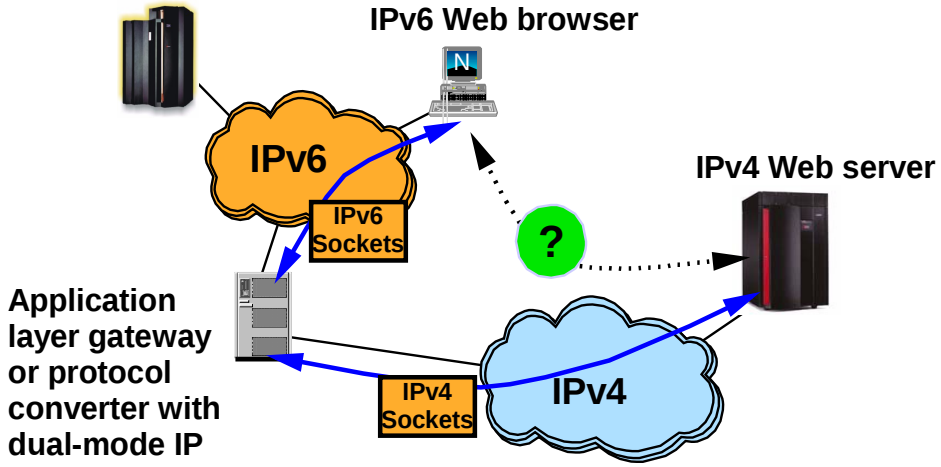

Isn't IPv6 enablement just a network engineering exercise?

- Unfortunately, no!

- A few facts:
  - The network infrastructure will have to be updated to support IPv6 network infrastructure functions, such as neighbor discovery (an auto-addressing technology), IPv6 routing tables (OSPFv3), ICMPv6, name servers with IPv4 and IPv6 addresses, DHCP servers for IPv6, etc.
    - Layer-3 routers
    - Firewalls
    - Intrusion Detection devices
    - Application layer gateways (ALGs)
    - Etc.
  - The physical media you use today can carry both IPv4 and IPv6 – so no new cabling (!)
  - A TCP/IP stack must be updated to support IPv6 – alongside with IPv4 (known as dual-mode TCP/IP stack)
  - IPv6 requires a new sockets interface, known as AF_INET6 (Addressing Family IPv6)
    - IPv4 sockets programs today use AF_INET, which is IPv4 only. An AF_INET sockets program can communicate with an IPv4 sockets partner only
    - Sockets programs that are updated to support AF_INET6 can communicate with both IPv4 and IPv6 sockets partners

- Sockets programs must be updated to talk IPv6 !!

z/OS TCP/IP is a dual-mode TCP/IP stack

- A dual-mode (or dual-stack) 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.

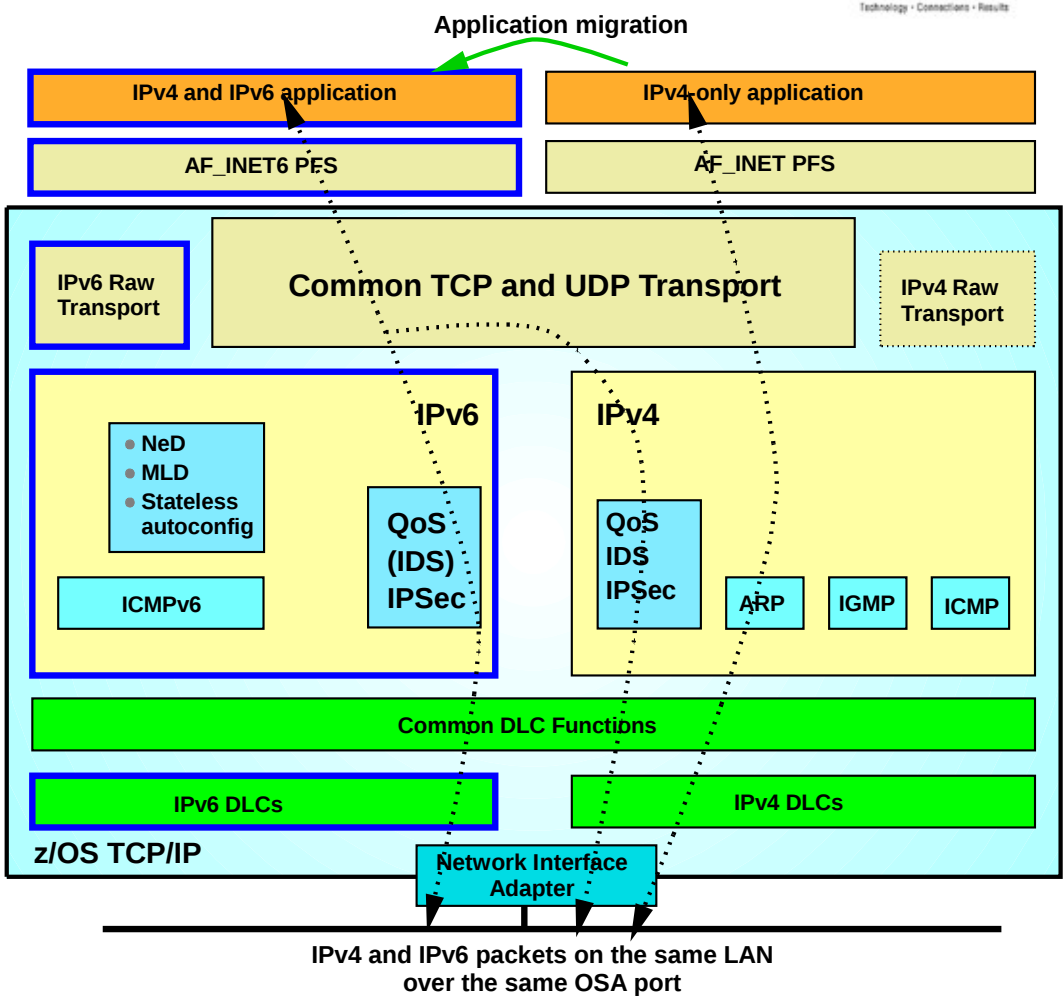

- Raw applications make the determination themselves when they choose IPv4 or IPv6 raw transport.

IPv6-enabled application on a dual mode stack

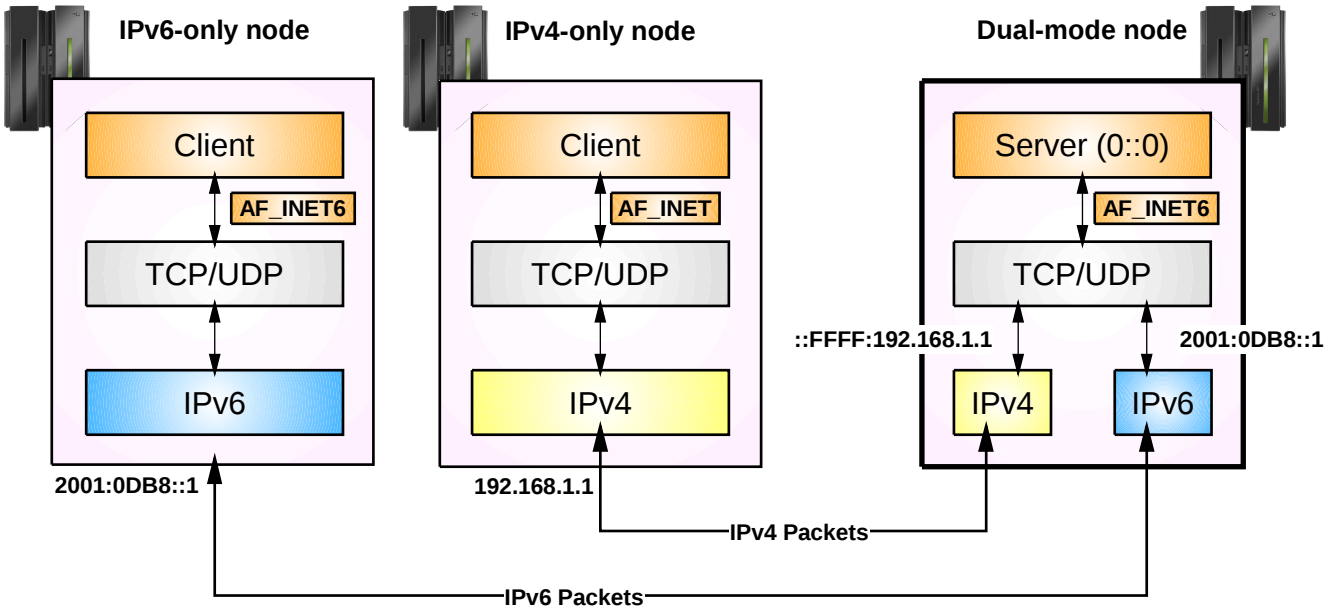

- An IPv6-enabled application can communicate with both IPv4 and IPv6 peers
  - A single socket can be used to send or receive traffic from either IPv4 or IPv6 partners
  - IPv4 packets to the IPv4 partner and IPv6 packets to the IPv6 partner
  - No changes need to be made to the partner application
- An IPv6-enabled application uses AF_INET6 sockets for both IPv4 and IPv6 partners
  - An IPv4 address is mapped to IPv6 addresses by the Transport Layer in the TCP/IP stack
  - Uses a special address format which identifies the IPv6 address as an IPv4-mapped IPv6 address
  - For example, 192.168.1.1 would be represented as ::FFFF:192.168.1.1

IPv4-only application on a dual-mode stack

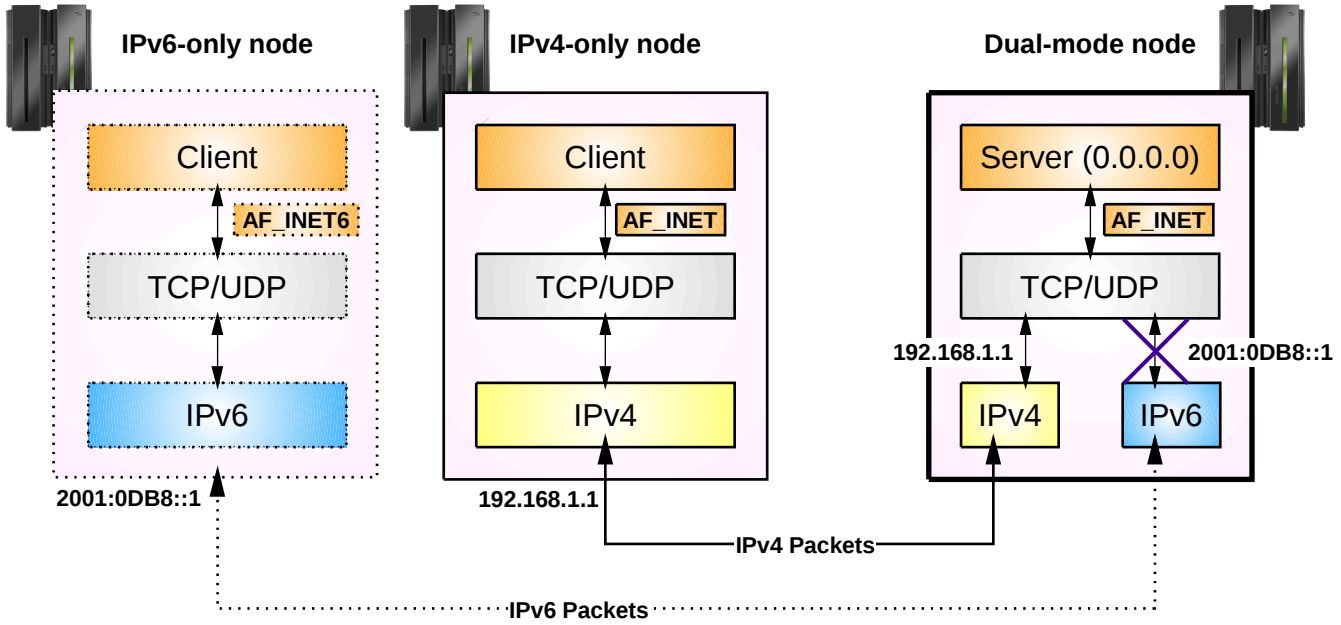

- An IPv4 application running on a dual-mode stack can communicate with an IPv4 partner.
  - The source and destination addresses will be native IPv4 addresses
  - The packet which is sent will be an IPv4 packet
- If partner is IPv6 running on an IPv6 only stack, then communication fails
  - If partner was on dual-mode stack, then it would fit in previous page discussion
  - The partner only has a native IPv6 address, not an IPv4-mapped IPv6 address
  - The native IPv6 address for the partner cannot be converted into a form the AF_INET application will understand

Accessing IPv4-only applications through an IPv6 application layer gateway (ALG)

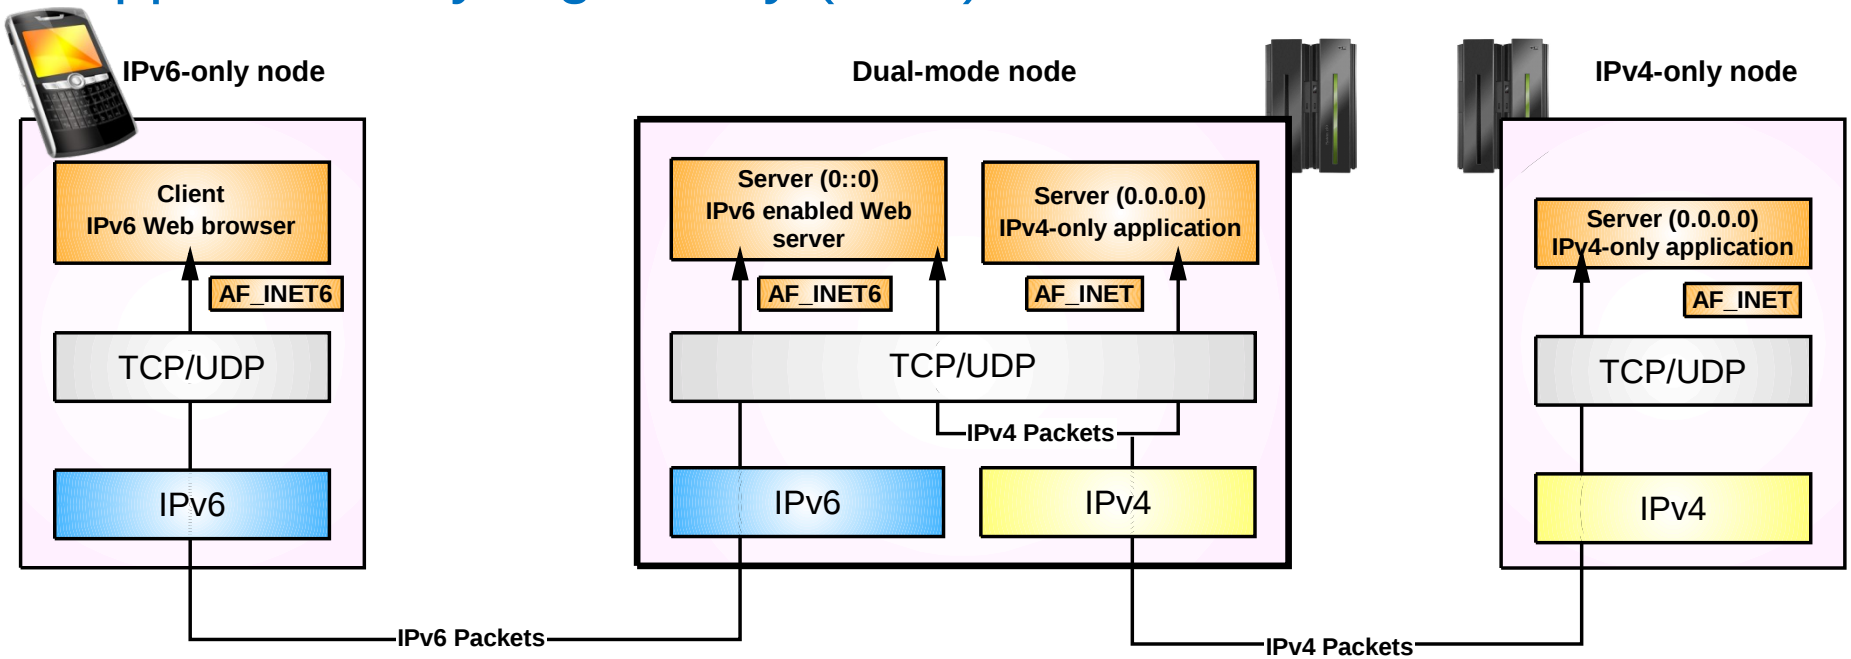

- An IPv6-only client can access IPv4-only servers via an IPv6 “proxy”
  - The IPv6 proxy communicates with the IPv6-only client using IPv6, and accesses the IPv4-only server using IPv4
  - The IPv4-only server may be on the same TCP/IP stack as the IPv6 proxy, or may reside on a different stack
  - The use of a backend IPv4-only server is, in most cases, completely transparent to the IPv6 client

For more information

URL	Content
http://www.twitter.com/IBM_Commserver	IBM z/OS Communications Server Twitter Feed
http://www.facebook.com/IBMCommserver	IBM z/OS Communications Server Facebook Page
https://www.ibm.com/developerworks/mydeveloperworks/blogs/IBMCommserver/?lang=en	IBM z/OS Communications Server Blog
http://www.ibm.com/systems/z/	IBM System z in general
http://www.ibm.com/systems/z/hardware/networking/	IBM Mainframe System z networking
http://www.ibm.com/software/network/commserver/	IBM Software Communications Server products
http://www.ibm.com/software/network/commserver/zos/	IBM z/OS Communications Server
http://www.redbooks.ibm.com	ITSO Redbooks
http://www.ibm.com/software/network/commserver/zos/support/	IBM z/OS Communications Server technical Support – including TechNotes from service
http://www.ibm.com/support/techdocs/atsmastr.nsf/Web/TechDocs	Technical support documentation from Washington Systems Center (techdocs, flashes, presentations, white papers, etc.)
http://www.rfc-editor.org/rfcsearch.html	Request For Comments (RFC)
http://www.ibm.com/systems/z/os/zos/bkserv/	IBM z/OS Internet library – PDF files of all z/OS manuals including Communications Server
http://www.ibm.com/developerworks/rfe/?PROD_ID=498	RFE Community for z/OS Communications Server
https://www.ibm.com/developerworks/rfe/execute?use_case=tutorials	RFE Community Tutorials

For pleasant reading ….