Communications Server z/OS IPv6

IPv6 agenda

- **Introduction**
- **Advanced sockets API for IPv6 update**
- **Maintain 2 IPv6 routers in default list**

Introduction

Visible IPv4 hosts on the Internet through the last 12 years

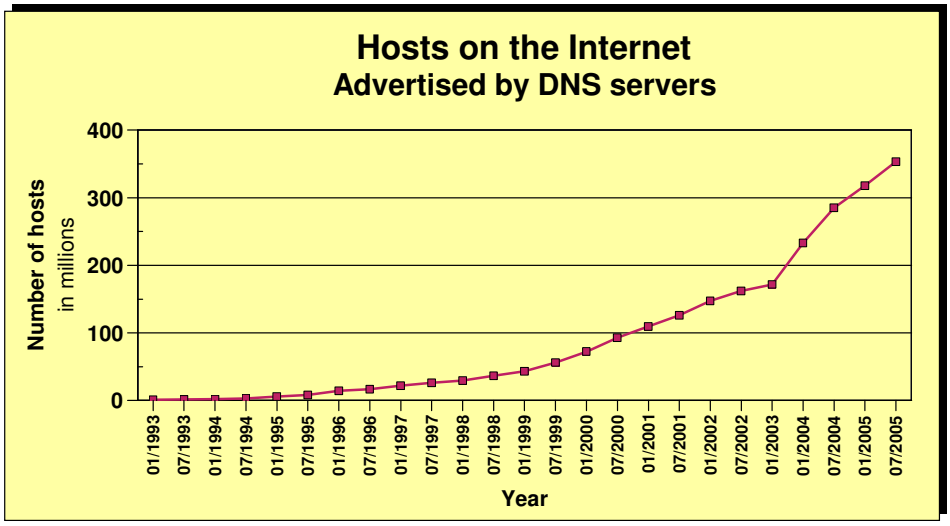

- **What is the upper practical limit (the ultimate pain threshold) for number of assigned IPv4 addresses? Some predictions say 250,000,000 (250 million), others go up to 1,000,000,000 (one billion or one milliard).**
- **Source: http://www.isc.org/index.pl?/ops/ds**

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.

Source is at http://www.isc.org/index.pl?/ops/ds

Both z/OS V1R5 and V1R6 have been certified with the IPv6 Ready logo

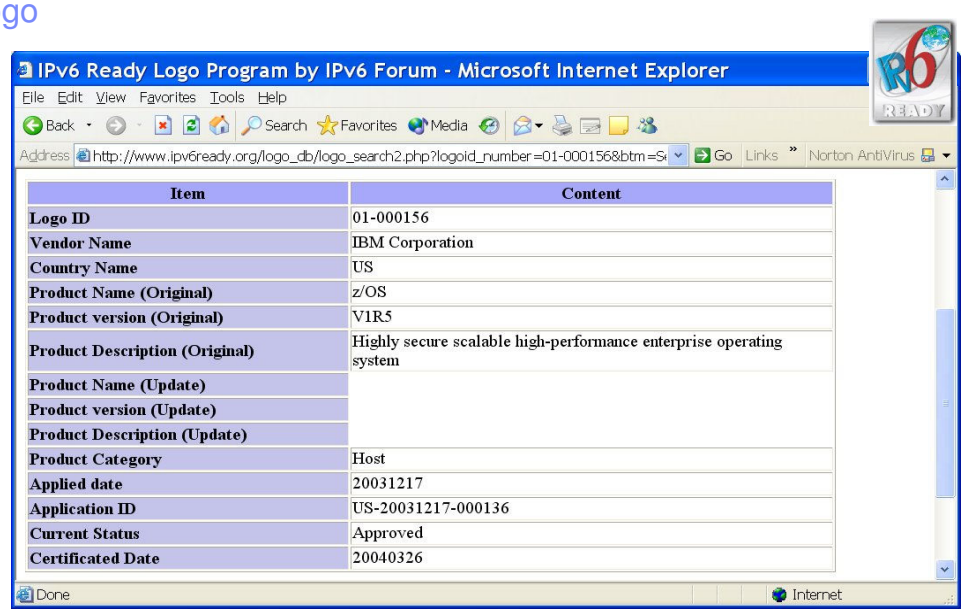

Item	Content
Logo ID	01-000156
Vendor Name	IBM Corporation
Country Name	US
Product Name (Original)	z/OS
Product version (Original)	V1R5
Product Description (Original)	Highly secure scalable high-performance enterprise operating system
Product Name (Update)	
Product version (Update)	
Product Description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

CS z/OS V1R7 is in the process of being certified too.

The Journey to IPv6 for z/OS Communications Server

- **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
- **The fourth phase (z/OS V1R7)**
  - SNMP UDP standard MIB (RFC2013) and IBM MVS TCP/IP Enterprise-specific MIB for UDP
  - Advanced Socket API support - RFC3542
  - IPv6 Two Default Routers - required for IPv6 compliance
  - HiperSockets DLC (z9-109)
- **After z/OS V1R7**
  - Integrated IPSec
  - Complete Advanced Socket APIs
  - Extended Stats MIB, OSPFv3 M
  - 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.

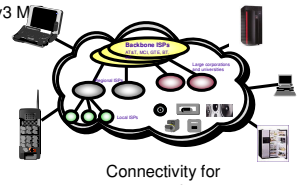

Connectivity for ***anyone*** from ***anywhere*** (car, home, office) to ***anything***!

Advanced sockets API for IPv6 update

Advanced sockets API for IPv6

- **Lets applications modify and receive information about packets**
  - Control and modify outbound packet information, such as
    - First hop address and routing headers
    - Hop options and destination options
    - Traffic class
    - Packet fragmentation
    - MTU discovery
  - Receive inbound packet information such as
    - Arriving interface
    - Destination IP address
    - Hop limit
    - Source routing
    - IPv6 options (routing headers, destination options, etc.) set by the sender
- **Defined first in RFC 2292, then in RFC 3542**
  - Defined for use by 'advanced' IPv6 applications
    - For example, ping, traceroute, and routing daemons
    - Geared more towards applications using RAW sockets
  - Separate from the Basic IPv6 Socket APIs in RFC 3493

Existing IPv6 advanced socket API support as of z/OS V1R4

- **z/OS V1R4 introduced partial support of the IPv6 advanced socket API**
- **The V1R4 implementation was based on draft RFC 2292.**
- **The following options were supported by the V1R4 implementation**

Level	Option Name	Get	Set	Data type	Data path	Transports supported
IPPROTO_IPV6	IPV6_USE_MIN_MTU	Y	Y	int	Outbound	TCP, UDP, RAW
	IPV6_PKTINFO	Y	Y	in6_pktinfo	Outbound	UDP, RAW
	IPV6_RECVHOPLIMIT	Y	Y	int	Inbound	TCP, UDP, RAW
	IPV6_RECVPKTINFO	Y	Y	in6_pktinfo	Inbound	UDP, RAW
	IPV6_CHECKSUM	Y	Y	int	Outbound	RAW
IPPROTO_ICMPV6	ICMP6_FILTER	Y	Y	icmp6_filter	Inbound	RAW

- **Outbound options could be set**
  - 'Sticky' using setsockopt - all packets on the socket used the option
  - Or 'per packet' using sendmsg ancillary data - affects only that packet
- **Inbound packet information is received on recvmsg() as ancillary data.**
- **The options could be used on z/OS UNIX callable services and using LE C/C++ APIs**

New advanced socket API options for IPv6

- **RFC 3542 is now implemented for UDP and RAW sockets and partially supported for TCP sockets.**
  - TCP only supports the very basic set of the IPv6 Advanced socket APIs
  - The more advanced APIs are geared towards UDP and RAW sockets
- **Means to provide RACF authentication to allow/disallow users and applications from using the APIs are provided.**
  - The RACF authentication is granular enough to specify access restrictions for each option of the API.
- **The options are supported only for z/OS UNIX callable services and LE C/C++ APIs**

Level	Option Name	Data Path	Transports Supported
IPPROTO_IPV6	IPV6_HOPOPTS	Outbound	UDP, RAW
	IPV6_RECVHOPOPTS	Inbound	UDP, RAW
	IPV6_RTHDR	Outbound	UDP, RAW
	IPV6_RECVRTHDR	Inbound	UDP, RAW
	IPV6_RTHDRDSTOPTS	Outbound	UDP, RAW
	IPV6_DSTOPTS	Outbound	UDP, RAW
	IPV6_RECVDSTOPTS	Inbound	UDP, RAW
	IPV6_RECVTCLASS	Inbound	TCP, UDP, RAW
	IPV6_TCLASS	Outbound	TCP, UDP, RAW
	IPV6_NEXTHOP	Outbound	UDP, RAW
	IPV6_RECVPATHMTU	Outbound	UDP, RAW
	IPV6_PATHMTU	Outbound	UDP, RAW
	IPV6_DONTFRAG	Outbound	UDP, RAW

These are the new advanced socket API options for IPv6 that are provided in z/OS V1R7.

RACF protection of the advanced IPv6 socket options

- **Access to the socket options is allowed under only three conditions**
  - Application is APF authorized - or -
  - User executing the application has super user authority - or -
  - Option resource name defined and the application has at least READ access to the resource.

- **The resource names are:**

API option	RACF Resource Name
IPV6_NEXTHOP	EZB.SOCKOPT.sysname.tcpname.IPV6_NEXTHOP
IPV6_TCLASS	EZB.SOCKOPT.sysname.tcpname.IPV6_TCLASS
IPV6_RTHDR	EZB.SOCKOPT.sysname.tcpname.IPV6_RTHDR
IPV6_HOPOPTS	EZB.SOCKOPT.sysname.tcpname.IPV6_HOPOPTS
IPV6_DSTOPTS	EZB.SOCKOPT.sysname.tcpname.IPV6_DSTOPTS
IPV6_RTHDRDSTOPTS	EZB.SOCKOPT.sysname.tcpname.IPV6_RTHDRDSTOPTS
IPV6_HOPLIMIT	EZB.SOCKOPT.sysname.tcpname.IPV6_HOPLIMIT
IPV6_PKTINFO	EZB.SOCKOPT.sysname.tcpname.IPV6_PKTINFO

- **Migration concerns are only for a multilevel security (MLS) environment**
  - IPV6_PKTINFO changed authorization in an MLS environment
  - To use the options in an MLS environment, the resource name must be defined and the application must have at least READ access to the resource.

Maintain 2 IPv6 routers in default list

Maintain 2 IPv6 routers in default list

- **IPv6 standards require a minimum of 2 default routers**
  - Required for the IPv6-Ready logo certification
- **In certain situations, z/OS CS does not meet this requirement**
  - If default routes are being removed from the stack routing table by OMPROUTE due to lost network connectvitity, the number of default IPv6 routers may go to zero
- **When the last default route is deleted from the routing table**
  - Add the default routers back to the routing table
- **No new configuration options and no migration concerns**

Trademarks, Copyrights, and Disclaimers

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

IBM	CICS	IMS	MQSeries	Tivoli
IBM(logo)	Cloudscape	Informix	OS/390	WebSphere
e(logo)business	DB2	iSeries	OS/400	xSeries
AIX	DB2 Universal Database	Lotus	pSeries	zSeries

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 registered trademarks of Microsoft Corporation in the United States, other countries, or both.

Intel, ActionMedia, LANDesk, MMX, Pentium and ProShare are trademarks of Intel Corporation in the United States, other countries, or both.

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

Linux is a registered trademark of Linus Torvalds.

Other company, product and service names may be trademarks or service marks of others.

Product data has been reviewed for accuracy as of the date of initial publication. Product data is subject to change without notice. This document could include technical inaccuracies or typographical errors. IBM may make improvements and/or changes in the product(s) and/or program(s) described herein at any time without notice. Any statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only. References in this document to IBM products, programs, or services does not imply that IBM intends to make such products, programs or services available in all countries in which IBM operates or does business. Any reference to an IBM Program Product in this document is not intended to state or imply that only that program product may be used. Any functionally equivalent program, that does not infringe IBM's intellectual property rights, may be used instead.

Information is provided "AS IS" without warranty of any kind. THE INFORMATION PROVIDED IN THIS DOCUMENT IS DISTRIBUTED "AS IS" WITHOUT ANY WARRANTY, EITHER EXPRESS OR IMPLIED. IBM EXPRESSLY DISCLAIMS ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NONINFRINGEMENT. IBM shall have no responsibility to update this information. IBM products are warranted, if at all, according to the terms and conditions of the agreements (e.g., IBM Customer Agreement, Statement of Limited Warranty, International Program License Agreement, etc.) under which they are provided. Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products in connection with this publication and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. IBM makes no representations or warranties, express or implied, regarding non-IBM products and services.

The provision of the information contained herein is not intended to, and does not, grant any right or license under any IBM patents or copyrights. Inquiries regarding patent or copyright licenses should be made, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. All customer examples described are presented as illustrations of how those customers have used IBM products and the results they may have achieved. The actual throughput or performance 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 or performance improvements equivalent to the ratios stated here.