Getting Started With IPv6

By Nalini Elkins

Is IPv6 hype or reality? How soon should we start to think about migration? We've heard about the shortage of IPv4 addresses, which will supposedly force us to use IPv6. How ready are the platforms we use daily for IPv6? If we need to migrate, can we really do it? How ready is the Internet for IPv6?

Research shows that the Internet appears to have at least some native IPv6 routing and backbone infrastructure. However, as far as commercial IPv6 services, it's scarce; few Internet Service Providers (ISPs) offer native IPv6 services in the U.S.

Federal agencies, particularly the Department of Defense (DoD) and National Institute of Standards and Technology (NIST), have led the adoption of IPv6 in the U.S. Supposedly, the network backbone for federal agencies must be IPv6-compliant by 2008. Actual implementation status among federal agencies varies widely.

The U.S. federal government has defined what Request for Comments (RFCs) must be implemented in various hardware and software equipment, a test plan, and the results of the test plan. There are two sets of IPv6 profiles for the U.S. government. The Joint Interoperability Test Command (JITC) produces the set of requirements for the DoD; NIST produces the set of requirements for the remainder of the U.S. government. The JITC has published V1 of the profiles, and is updating to V1.1. The NIST profiles are in draft state, recently completing a request for comments from the technical community.

You can access the profiles at these sites:

- **JITC:** http://jitc.fhu.disa.mil/adv_ip/register/register.html
- **NIST:** www.antd.nist.gov/usgv6-v1-comments.html.

Business community adoption in the U.S. is mostly in preliminary stages. Many are waiting for a business driver to justify the migration costs. Experts differ as to when we'll run out of IPv4 addresses; the speculation ranges from five years to 17 years. Nevertheless, implementations, lab activity, and research into IPv6 have started at some major corporations. Even if IPv6 won't become an issue until seven or eight years from now, there are so many transition issues that it's wise to become familiar with the new protocols, features, and software and hardware requirements. It may take large companies several years to completely convert to IPv6.

This article examines the state of Internet infrastructure for native IPv6 and use of IPv6 in today's workhorse systems: Windows XP and the IBM mainframe operating system, z/OS.

IPv6 Peer Exchanges

The first requirement for IPv6 over the Internet is to have peer exchanges capable of transporting native IPv6. Internet peer or public exchange facilities provide a common location where multiple service providers can meet and exchange customer traffic. These public exchange services let customers peer with each other without the need for dedicated links. Many such peering facilities are available for IPv4, including the familiar MAE West and MAE East peer exchanges.

What's available for IPv6? Such exchange points using native IPv6 are the first step to connectivity and acceptance of IPv6 as a routine production service. Current native IPv6 peer exchange points include:

- www.6iix.net/ (NY and LA)
- www.6tap.net/ (Chicago; see also www.6ren.net/)
- www.ny6ix.net/ (NY)
- www.nttmcl.com/html/ComIPv6.html
- www.paix.net/
- www.ams-ix.net/home.html (Amsterdam)
- www.uk6x.com/ (London; see also www.labs.bt.com/projects/ipv6.htm)
- www.inxs.de/ipv6.shtml (Munich)
- www.wide.ad.jp/nspixp6/ (Tokyo).

Let's track the exchange points used by a large ISP, Nippon Telephone and Telegraph (NTT), which is quite active in the IPv6 world (see Figure 1). In fact,

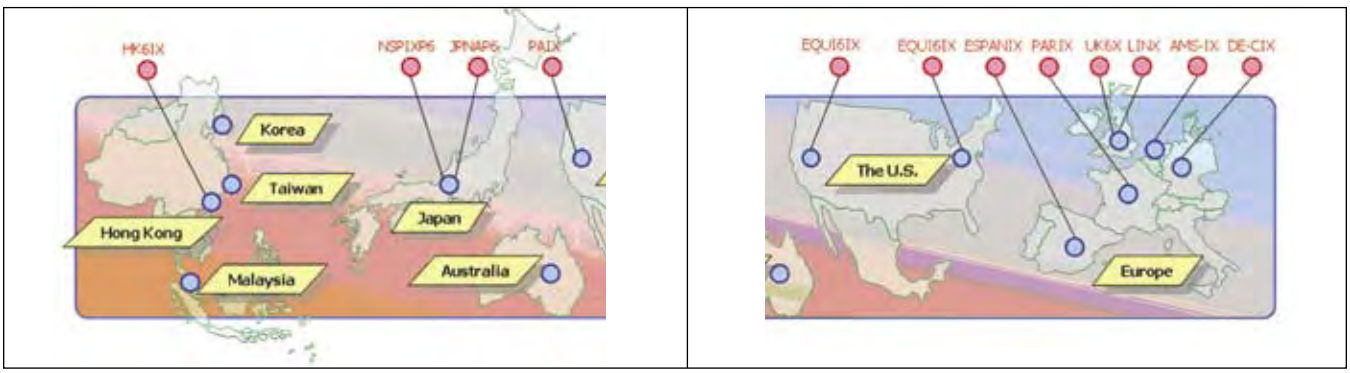

Figure 1: NTT IPv6 Exchange Points

they're one of the few ISPs to offer commercial IPv6 addresses in the U.S.

Peer Exchange Members

Some exchanges provide a list of their participants. This is illuminating, as it reveals which ISPs have IPv6 capability and have exploited it to the extent of registering to peer traffic. Figure 2 shows some of the ISPs that have registered with the Big APE in New York City. Google appears in both the Big APE and the INXS (Munich) peer exchanges. INXS in Munich is a widely used peering exchange in Europe run by cable and wireless interests (see Figure 3).

IPv6 on Windows XP

It's easy to enable IPv6 on Windows XP. From a DOS window, just type: ipv6 install. This is what you'll see:

```
C:\ipv6 install
Installing…
Succeeded.
```

To ping an IPv6 address from Windows, key in the code shown in Figure 4. To do a trace route to an IPv6 address from Windows, type:

```
tracert6 [ipaddress]
```

Figure 5 is a trace route to the Japan IPv6 portal; here, we're going out over IPv6 addresses all the way to Japan!

IPv6 on z/OS

On z/OS, to bring up a simple system with IPv6 enabled and taking all the defaults, BPXPRMxx and the TCP profile must be changed. There'll be two TCP stacks running IPv6 and IPv4. Currently, dual stack is the only supported mode. Of course, you may run IPv4 only, but you can't run only IPv6.

BPXPRMxx Changes

Your BPXPRMxx will look similar to Figure 6. To add support for IPv6, add the socket definitions for IPv6 to BPXPRMxx. You may change MAXSOCKETS to a number you choose; the other parameters must be as specified in Figure 7.

TCP Profile Changes

The changes to the TCP profile are in the IPconfig6 macro. You currently have an IPv4 macro called "IPCONFIG." This controls the configuration for IP version 4. The IPCONFIG macro has a number of optional parameters. The one in your TCP profile may look something like this:

```
IPConfig IGNOREREDIRECT …
```

To add support for IPv6, add the following line to the TCP profile:

```
IPCONFIG6
```

This is the set of configuration parameters that control IPv6. If we code only the statement "IPCONFIG6," this means we are taking all the defaults for the IPv6 parameters.

IPv6 Start-Up on z/OS

Now, re-IPL or restart TCP/IP. Figure 8 shows the start-up messages for TCP/IP on the system console. The last two messages in Figure 8 indicate that TCP start-up was successfully completed. However, the messages don't show if the IPv6 stack was successfully started or the configuration defaults taken. To see that, you must enter the NETSTAT UP command in Figure 9. If IPv6 isn't enabled, you'll see the message in Figure 10.

Now we can be sure IPv6 is really enabled. You can check the defaults taken with the IPCONFIG6 command by doing NETSTAT CONFIG (see Figure 11). You'll see separate definitions for IPv4 and IPv6. Notice some new parameters for IPv6:

- Ignore Router Hop Limit (don't take a different limit received in Router Advertisement)

About BIG APE

The Big Apple Peering Exchange [BIG APE] (Formerly NY6IX) is an IP-based Gigabit-Ethernet Exchange (IX) located at tel^{x} 60 Hudson Street, 9th Floor in New York City for the purpose of exchanging Internet Protocol (IP) traffic

Current Participants

Org	ASN	IP	Contact
Stealth Communications	8002	198.32.238.11 2001:0458:0026:2::80	Peering@Stealth.net
Quest	209	198:32.238.12 2001:0458:0026:2::300	Peering@mail.quest.net
ISCf.rootservers.netServers	27319	198:32.238.12 2001:0458:0026:2::600	Peering@isc.org
WV Fiber	19151	198:32.238.14	Peering@wvfiber.com
PPL Telcom	32327	198:32.238.16	Peering@ppltelcomip.com
Progress Telecom	19962	198:32.238.17	noc@progresstelecom.com
ericom	5466	198:32.238.18	peering@ericom.net
Net2Phone	7270	198:32.238.19	peering@Net2Phone.com
Teleglobe	6453	198:32.238.20 2001:0458:0026::380	Peering@teleglobe.com
Lexent	22270	198:32.238.21	dhaluza@lexent.com
Google	15169	198:32.238.22	Peering@google.com
PAETEC Communications	15270	198:32.238.23	bgp-admin@paetec.com
Spiritlink	6175	2001:0458:0026:2::100	ipv6-support@spirit.net
France Telecom	5511	2001:0458:0026:2::400	peering@opentransit.net
IIJ	2497	2001:0458:0026:2::200	V6-peering@iij.ad.jp
DISN-LES-V6	5994	2001:0458:0026:2::180	info@v6.les.disa.mil
KDDI	2516	2001:0458:0026:2::280	ipv6@kiddi.com
UPC Distribution Services	6830	2001:0458:0026:2::480	peering@aorta.net
Hurricane Electric	6939	2001:0458:0026:2::500	peering@he.net
Easynet	4589	2001:0458:0026:2::580	peering@easynet.net
Cable&Wireless	1273	2001:0458:0026:2::680	ipv6@cw.net

Figure 2: Big APE Exchange Members

AS	Name	Status IPv4	IPv4Multicast	IPv6
42	WoodyNet	A		
174	CCD Cogent Communications Deutschland GmbH	A		
286	KPN Eurorings	A	I	A
1273	Cable & Wireless Telecommunication Services GmbH	A	A	A
3209	Arcor AG & Co	A		
3291	PSINet Europe Bv.	A		
3320	T-Systems (Deutsche Telekom)	A		A
3856	PCH	A		
4589	Easynet GmbH	A		A
5539	SpaceNet GmbH	A	A	A
6805	Telefonica Deutschland GmbH	A		
8220	COLT Internet (UK)	A		
8222	netplace Telematic GmbH	A		
8481	Transnet GmbH	A		A
8560	1 & 1 Internet AG	A		I
8767	M-Net TeleKomuniKations GmbH	A	P	A
8928	Interoute Communications Ltd.	P		P
9132	Boardnet Mediascape Communications AG	A		
12306	Plus. line AG	A		
12337	noris network AG	A		A
12339	Scan Plus GmbH	A		A
13132	Cbyerways Gmbh	A		
13184	Hansenet Gmbh	A		
13237	Lambdanet Communications	A		
13246	InterNetWire Communications GmbH	A		
15169	Google Inc	P	P	P

Figure 3: INXS Peer Exchange Members

```
ping6 [ipaddress]

Pinging www.kame.net [2001:200:0:8002:203:47ff:fea5:3085]
with 32 bytes of data:

Reply from 2001:200:0:8002:203:47ff:fea5:3085: time=224ms
Reply from 2001:200:0:8002:203:47ff:fea5:3085: time=223ms
Reply from 2001:200:0:8002:203:47ff:fea5:3085: time=226ms
Reply from 2001:200:0:8002:203:47ff:fea5:3085: time=223ms

Ping statistics for 2001:200:0:8002:203:47ff:fea5:3085:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
    Minimum = 223ms, Maximum = 226ms, Average = 224ms
```

Figure 4: Ping to Japan IPv6 Portal

```
Tracing route to www.kame.net [2001:200:0:8002:203:47ff:fea5:3085]
over a maximum of 30 hops:
  1    <1 ms    <1 ms    <1 ms  2001:4840:ffff:c012:214:bfff:feba:45f9
  2    74 ms    71 ms    67 ms  ipv601-smv.research.EARTHLINK.NET [::209.179.5.34]
  3    67 ms   102 ms    72 ms  2001:4840:ffff::1
  4    69 ms    77 ms    70 ms  earthlink-gw.customer.klax.occaid.net [2001:4830:ea:3::2]
  5    70 ms    70 ms    71 ms  2.fe-0-0.cr1.lax3.us.occaid.net [2001:4830:ea:3::1]
  6    70 ms    72 ms    70 ms  8.ge-0-0.cr1.lax1.us.occaid.net [2001:4830:ff:11ab::1]
  7    80 ms    79 ms    79 ms  so-1-0-0.cr1.sjc2.us.occaid.net [2001:4830:ff:1201::1]
  8    80 ms    97 ms    89 ms  38.fe0-0.cr1.sfo2.us.occaid.net [2001:4830:ff:12ea::1]
  9    89 ms    83 ms    82 ms  sf-guest.r2.sfo2.isc.org [2001:4f8:4:b:290:6900:b30:541f]
 10   242 ms   226 ms   225 ms  isc.200paul.wide.ad.jp [2001:4f8:4:236::10]
 11   224 ms   223 ms   231 ms  hitachi1.otemachi.wide.ad.jp [2001:200:0:4401::3]
 12   241 ms   231 ms   225 ms  2001:200:0:1c04:230:13ff:feae:5b
 13   231 ms   230 ms   227 ms  2001:200:0:4800::7800:1
 14   224 ms   224 ms   224 ms  orange.kame.net [2001:200:0:8002:203:47ff:fea5:3085]

Trace complete.
```

Figure 5: Trace Route to the Japan IPv6 Portal

```
NETWORK  DOMAINNAME(AF_INET)
         DOMAINNUMBER(2)
         MAXSOCKETS(30000)
         TYPE(INET)
         INADDRANYPORT(5555)
         INADDRANYCOUNT(1000)

NETWORK  DOMAINNAME(AF_UNIX)
         DOMAINNUMBER(1)
         MAXSOCKETS(2000)
         TYPE(UDS)
```

Figure 6: BPXPRMxx

```
NETWORK  DOMAINNAME(AF_INET6)
         DOMAINNUMBER(19)
         MAXSOCKETS(3000)
         TYPE(INET)
```

Figure 7: Changes to BPXPRMxx for IPv6

- ICMP Error Limit (how many ICMP messages should be allowed in a burst).

Fun With the Loopback Interface

But what will this allow us to do? We haven't defined any interfaces (see Figure 12). Loopback is the only active interface. To define interfaces, we must have an interface capable of supporting IPv6 and a router attached. But we can still do some interesting things with a loopback interface. Let's try a ping to the loopback interface ::1. Enter the following:

```
Ping ::1
```

The response will be something like:

```
CS V1R7: Pinging host ::1
Ping #1 response took 0.002 seconds
```

IPv6 Packets

What does an IPv6 packet really look like? How different will it be from an IPv4 packet? If we take a trace of the packet, we'll see the ICMPv6 packet as shown in Figure 13. This packet seems familiar in many ways. The ICMPv4 PING is the same as the ICMPv6 PING. The differences are in the IPv6 and IPv4 main headers.

IPv4 vs. IPv6 Header

You'll see all the changes in the header. For the IPv4 header:

- Source and destination addresses are 32 bits (4 bytes) in length.
- Fragmentation is supported at originating hosts and intermediate routers.
- IP header includes a checksum.
- IP header includes options.
- No identification of payload for Quality of Service (QoS) handling by routers is present in the IPv4 header.

```
STC00027 00000090  EZZ0300I OPENED PROFILE FILE DD:PROFILE
STC00027 00000090  EZZ0309I PROFILE PROCESSING BEGINNING FOR DD:PROFILE
         00000090  BPXI004I OMVS INITIALIZATION COMPLETE
STC00027 00000090  EZZ0316I PROFILE PROCESSING COMPLETE FOR FILE DD:PROFILE
STC00027 00000090  EZZ0641I IP FORWARDING NOFWDMULTIPATH SUPPORT IS ENABLED
STC00027 00000090  EZZ0335I ICMP WILL NOT IGNORE REDIRECTS
STC00027 00000090  EZZ0337I CLAWUSEDOUBLENOP IS CLEARED
STC00027 00000090  EZZ0345I STOPONCLAWERROR IS DISABLED
STC00027 00000090  EZZ0338I TCP PORTS 1 THRU 1023 ARE RESERVED
STC00027 00000090  EZZ0338I UDP PORTS 1 THRU 1023 ARE RESERVED
STC00027 00000090  EZZ0336I A LIMIT ON INCOMING UDP DATAGRAM QUEUE SET
STC00027 00000090  EZZ4202I OPENEDITION-TCP/IP CONNECTION ESTABLISHED FOR TCPIP
STC00027 00000090  EZZ4314I INITIALIZATION COMPLETE FOR DEVICE LCS1, LINK ETH1
STC00027 00000090  EZB6473I TCP/IP STACK FUNCTIONS INITIALIZATION COMPLETE.
STC00027 00000090  EZAIN11I ALL TCPIP SERVICES FOR PROC TCPIP ARE AVAILABLE.
```

Figure 8: The Start-Up Messages for TCP/IP on the System Console

```
                          NETSTAT UP - Display Time TCP/IP Started
                                   Address:192.168.1.231
MVS TCP/IP NETSTAT CS V1R6                TCPIP NAME: TCPIP         22:35:42
TCPIP STARTED AT 20:16:49 ON 01/24/2006 WITH IPV6 ENABLED
```

Figure 9: Netstat UP With IPv6 Enabled

```
MVS TCP/IP NETSTAT CS V1R6           TCPIP NAME: TCPIP               22:24:20
TCPIP STARTED AT 20:22:43 ON 02/23/2006 WITH IPV6 DISABLED
```

Figure 10: Netstat UP With IPv6 Disabled

```
MVS TCP/IP NETSTAT CS V1R6     TCPIP NAME: TCPIP                  22:41:40
IP CONFIGURATION TABLE:
FORWARDING:       YES       TIMETOLIVE:         00064              RSTIMEOUT:00060
FIREWALL:         NO
ARPTIMEOUT:       00300     MAXRSMSIZE:         65535     FORMAT:   LONG
IGREDIRECT:       NO        SYSPLXROUT:         NO        DOUBLENOP:NO
STOPCLAWER:       NO        SOURCEVIPA:         NO        VARSUBNET:NO
MULTIPATH:        NO        PATHMTUDSC:         NO        DEVRTRYDUR:0000000090
DYNAMICXCF:       NO
IQDIOROUTE:       NO
TCPSTACKSRCVIPA: NO

IPV6 CONFIGURATION TABLE:
FORWARDING:         YES     HOPLIMIT:           00255     IGREDIRECT:         NO
SOURCEVIPA:         NO      MULTIPATH:          NO        ICMPERRLIM:         00003
IGRTRHOPLIMIT: NO
DYNAMICXCF:       NO
TCPSTACKSRCVIPA: NO
```

Figure 11: Netstat Config for IPv4 and IPv6

```
767 ADCD      PACKET    00000004 12:44:29.187575 Packet Trace
From Interface     : LOOPBACK6           Device: Loopback6        Full=304
Tod Clock          : 2007/05/22 12:44:29.187569                   Intfx: 3
Sequence #         : 0                   Flags: Pkt Home Ping
IpHeader: Version  : 6                   Header Length: 40
Class:             : 00                  Flow: 000000
Payload Length     : 264
Hops               : 255                 Protocol: ICMPV6
Source             : ::1
Destination        : ::1
ICMPv6
Type/Code          : ECHO Request        CheckSum: BFFF FFFF
Id                 : 0031                Seq: 1
Time               : 2006/01/22 20:44:29.186326
Echo Data          : 256
000000 43D3EEAD 0002D7D6 08090A0B 0C0D0E0F
|C...............................|

More data…

IP Header           : 40
60000000 01083AFF 00000000 00000000  00000000 00000001 00000000 0000000
00000000 00000001

Protocol Header     : 8
000000 8000BFFF 00310001

Data                : 256     Data Length: 256
000000 43D3EEAD 0002D7D6 08090A0B 0C0D0E0F |.L....PO........ C...............|
```

Figure 13: The ICMPv6 Packet

For the IPv6 header:

- Source and destination addresses are 128 bits (16 bytes) in length.
- Fragmentation isn't supported at routers; it's supported only at the originating host.
- IP header doesn't include a checksum.
- All optional data is moved to IPv6 extension headers.
- IPSec support is required in a full IPv6 implementation.
- Payload identification for QoS handling by routers is included in the IPv6 header using the flow label field.

Summary

There's considerable support in various platforms (Windows, z/OS, Unix, Cisco) for IPv6. The Internet appears to have at least some native IPv6 routing and backbone infrastructure. It's likely that IPv6 is inevitable. The sooner you start to learn, the more comfortable you'll be with all the new protocol structures.

Inside Products is leading a discussion forum for private sector businesses interested in IPv6 migration. It will be organized under the auspices of the California IPv6 Task Force (CAv6TF), which is part of the National IPv6 Task Force (NAv6TF). Contact ipv6@insidethestack.com for more information.

About the Author

Nalini Elkins, the CEO and founder of Inside Products, Inc., is a recognized leader in the field of computer performance measurement and analysis. In addition to being an experienced software product designer, developer and planner, she is a formidable businesswoman. She has been the founder or co-founder of two start-ups in the high-tech arena. She has a strong computer networking background, but specializes in network performance analysis, measurement, monitoring, tuning, and troubleshooting of large enterprise computing networks, including TCP/IP and SNA.
Email: nalini_elkins@inside-products.com

```
MVS TCP/IP NETSTAT CS V1R6
HOME ADDRESS LIST:
LINKNAME:        ETH1
  ADDRESS:       192.168.1.231
    FLAGS:       PRIMARY
LINKNAME:        LOOPBACK
  ADDRESS:       127.0.0.1
    FLAGS:
INTFNAME:        LOOPBACK6
  ADDRESS:       ::1
    TYPE:        LOOPBACK
    FLAGS:
```

Figure 12: Netstat Home With IPv6 Loopback