2G1305 Internetworking/Internetteknik
Spring 2005, Period 4

Module 10: IPv6

Lecture notes of G. Q. Maguire Jr.

KTH Information and
Communication Technology

For use in conjunction with *TCP/IP Protocol Suite*, by Behrouz A. Forouzan, 3rd Edition, McGraw-Hill.

For this lecture: Chapter 27

Last modified: 2005.05.14:17:34

Lecture 6: Outline

- IPv6

Internet Protocol Version 6 (IPv6)

- Successor of current IPv4
- Internet needs to change IP in order to continue growth
- Defines a transition from IPv4 to IPv6

Specified by RFC 2460: Internet Protocol, Version 6 (IPv6) Specification, December 1998[78].

Growth

- Currently IPv4 serves a market doubling every ~12 months
- In addition, new and very large markets are developing rapidly:
  - Nomadic Computing
  - Networked Entertainment
  - Device Control

Nomadic Computing

- Wireless computers
  - supporting multimedia
  - replacing pagers, cellular telephones, …
- IPv6 includes support for mobility
  - low overhead (?)
  - auto configuration
  - mobility

Networked Entertainment

Your TV will be an Internet Host!
[consider the network attached Personal Video Recorders (PVR), such as TiVo's DVR, SONICblue's ReplayTV, Sony's SVR-2000, Philips' PTR, …)]

- 500 channels of television
- large scale routing and addressing
- auto-configuration
- requires support for real-time data

SonicBlues's ReplayTV 4000 a networked Digital Video Recorder (DVR) {i.e., coder/decoder + very big disk) that takes advantage of your broadband Internet connection - enables you to capture and transfer videos.

Providing "narrowcast" content via broadband ⇒ all the time is "primetime".

Device Control

- Control everyday devices for
  - lightning, heating and cooling, motors, ...
  - new street light controllers already have IP addresses!
  - electrical outlets with addresses
  - networked vehicles (within the vehicle[1], between vehicles, and vehicles to infrastructure)[2]
- Market size is enormous
- Solution must be
  - simple, robust, easy to use
  - very low cost
  - potential power savings by (remote) network management based control may be quite large

There is already a networked: Toaster, a Coke machine, … .

1. On-Board Diagnostic systems (OBD-II), see slide 8 [79]
2. See InternetCAR, slide 4 (showing a Yokohama City bus) [79]

IPv6 features

- Expanded Addressing Capabilities
  - 128 bit address length
  - supports more levels of hierarchy
  - improved multicast routing by using a scope field
  - new cluster addresses to identify topological regions
- Header Format Simplification
  - some IPv4 fields have been dropped, some made optional
  - header is easier to compute
- Improved Support for Extensions and Options
  - more efficient for forwarding of packets
  - less stringent limits to length of options
  - greater flexibility for introduction of future options
- Flow Labeling Capability
  - labeling of packets belonging to a particular “flow”
  - allows special handling of, e.g., real-time, packets
- Authentication and Privacy Capabilities
  - Extensions to support authentication, data integrity, and (optional) data confidentiality

IPv6 header format

version 4 bits	Class 8 bits	flow label 20-bits		
"payload" length (in octets) 16 bit			next header 8 bits	hop limit 8 bits
Source Address 128 bits				
Destination Address 128 bits				

IPv6 header (total length = 40 bytes)

IPv6: 6 fields + 2 addresses

versus

IPv4: 10 fixed fields + 2 addresses + options

Demultiplexing

Initially, it was assumed that by keeping the version field the same that IPv4 and IPv6 could be mixed over the same links with the same link drivers.

However, now IPv6 will be demultiplexed at the link layer:
hence, IPv6 been assigned the Ethernet type 0x86DD (instead of IPv4's 0x8000)

Simplifications

IPv6 builds on 20 years of internetworking experience - which lead to the following simplifications and benefits:

Simplification	Benefits
Use fixed format headers	Use extension headers instead, thus no need for a header length field, simpler to process
Eliminate header checksum	Eliminate need for recomputation of checksum at each hop (relies on link layer or higher layers to check the integrity of what is delivered)
Avoid hop-by-hop segmentation	No segmentation, thus you must do **Path MTU discovery** or only send small packets (1996: 536 octets, 1997: proposed 1500 octets) (for observed PMTUs see [81]) • This is because we should have units of **control** based on the units of **transmitted** data.
Eliminate Type of Service (ToS) field	Instead use (labeled) flows

Quality-of-Service Capabilities

- for packet streams
- Flow characterized by flow id + source address + destination address
- unique **random** flow id for each source

CLASS (8 bits)	FLOW ID (20 bits)

- Class field

D (1 bit)	**Network-wide priority (3 bits)**	**Reserved (4 bits)**
Delay sensitive	Encodes the priority of traffic, can be used to provide "Differentiated services"	Researchers would like to use two of these bits for congestion avoidance control: ♦ one bit which could be set by routers to indicate that congestion was experienced; ♦ the other bit could be used by the source to mark that it is "ready to adapt".

- Flow ID - indicates packets which should all be handled the same way.

The original specified in *RFC 1809: Using the Flow Label Field in IPv6*

Subsequently updated - see Chapter 6 of Huitema, **2nd edition**; this change occurred because of Steve McCanne's SigComm'96 paper [83]. **Note that chapter 27 in Forouzan is incorrect!**

Payload length

Payload length is the length of the data carried after the header.

As the length field is 16 bits ⇒ maximum packet size of 64 kilobytes; but there is a provision for "jumbograms" [via the Hop-by-Hop option header with option type 194]. See RFC 2675 [84].

IPv4 Protocol type ⇒ IPv6 Next Header type

Tells how to interpret the next header which follows, it is either the payload type or the type of the next header. [Payload types use the IPv4 protocol type values]

Decimal	Keyword	Header type
0	HBH	Hop-by-hop options
2	ICMP	IPv6 ICMP
3	GGP	Gateway-to-Gateway Protocol
5	ST	Stream
6	TCP	Transmission Control Protocol
17	UDP	User Datagram Protocol
43	RH	IPv6 Routing Header
44	FH	IPv6 Fragmentation Header
45	IDRP	Inter-domain Routing Protocol
51	AH	Authentication Header
52	ESP	Encrypted Security Payload
59	Null	No next Header (IPv6)
60		IPv6 Destination Options Header
88	IGRP	IGRP
89	OSPF	Open Shortest Path First
255		Reserved

Extension headers

- Each header is a multiple of 8 octets long
- order (after IPv6 header):
  - Hop-by-hop option,
  - Destination options header (1)
  - Routing header,
  - Fragment header,
  - Authentication header,
  - Encapsulating security payload header,
  - Destination options header (2)
  - Followed by the upper layer header (e.g., TCP, UDP, …)

So a TCP packet looks like:

IPv6 header Next Header = **TCP**	TCP header + data

If we wanted to explicitly route the above packet, we simply add a routing header:

IPv6 header Next Header = **Routing**	Routing header Next Header = **TCP**	TCP header + data

Addressing

- 128 bits long
- three types: unicast, multicast, anycast

Unicast	identifies exactly one interface
Multicast	identifies a group of interfaces; a packet sent to a multicast address will be delivered to all members of the group
Anycast	delivered to the nearest member of the group

- 2^{96} times more addresses than IPv4 are available !!!

IPv6 addresses per m^2

Earth: 511,263,971,197,990 m^2
$\Rightarrow$ 665,570,793,348,866,943,898,599 / m^2

- pessimistic estimate with hierarchies: ~1,564 addresses / m^2
- optimistic: 3,911,873,538,269,506,102 /m^2

Writing an IPv6 address

The 128 bit IPv6 address is written as eight 16 bit integers using hexadecimal digits.

The integers are separated by colons, for example:
2001:0DB8:7654:3210:FEDC:BA98:7654:3210

A number of abbreviations are allowed:

- leading zeros in integers can be suppressed
- a single set of consecutive 16 bit integers with the value null, can be replaced by double colon, i.e.,
  2001:DB8:0:0:0:0:7654:3210 becomes 2001:DB8::7654:3210
- When an IPv4 address is turned into an IPv6 address we prepend 96 bits of zero; but we can write it as:
  ::10.0.0.1 - hence combining dotted-decimal and IPv6 forms
- Prefixes can be denoted in the same manner as for IPv4, i.e., CIDR:
  2001:DB8::/32 - for a 32 bit long prefix

Address Allocation [94] and [99]

Binary prefix	Hex. prefix	Fraction of address space	Assignment
0000 0000	::/8	1/256	Reserved
0000 0001	100::/8	1/256	Unassigned
0000 001	200::/7	1/128	Network Service Access Point (NSAP) Allocation-RFC 1888
0000 01	400::/6	1/64	Unassigned (first half was formerly Novell’s IPX)
0000 1	800::/5	1/32	Unassigned
0001	1000::/4	1/16	Unassigned
001	**2000::/3**	**1/8**	**Global Unicast - RFC 2374 see RFC 3587**
010	4000::/3	1/8	Unassigned (formerly provider based unicast addresses)
011	6000::/3	1/8	Unassigned
100	8000::/3	1/8	Unassigned (formerly Geographic-based Unicast Addresses)
101	A000::/3	1/8	Unassigned
110	C000::/3	1/8	Unassigned
1110	E000::/4	1/16	Unassigned
1111 0	F000::/5	1/32	Unassigned
1111 10	F800::/6	1/64	Unassigned
1111 110	FC00::/7	1/128	Unassigned
1111 1110 0	FE00::/9	1/512	Unassigned
1111 1110 10	**FE80::/10**	**1/1024**	**Link Local Use Addresses**
1111 1110 11	**FEC0::/10**	**1/1024**	**Reserved for IANA (was Site Local Use Addresses)**
1111 1111	**FF00::/8**	**1/256**	**Multicast Addresses**

Global Unicast Addresses

RFC 2374 defined an IPv6 address allocation structure that which featured Top Level Aggregator (TLAs) and Next Level Aggregator (NLAs) - this has been replaced (see RFC 3587[95]) by a coordinated allocation policy defined by the Regional Internet Registries (RIRs) [96]

The Subnet Local Aggregator (SLAs) of RFC 2374 ⇒ now called the “subnet ID”

001 (3 bits)	global routing prefix (45 bits)	subnet ID (16 bits)	Interface ID (64 bits)

Thus the Regional Internet Registries are allocating addresses from 2000:/3

For a table of IPv6 unicast assignment see

http://www.iana.org/assignments/ipv6-unicast-address-assignments

For an analysis of use from the point of view of RIPE see [101]

Interface ID

Must be unique to the link, but there are some advantages of making it more globally unique.

Hence, most will be based on the IEEE EUI-64 format, but with the "u" (unique) bit inverted.

- The "u" bit is the 7th most significant bit of a 64 bit EUI.
- The inversion was necessary because 0:0:0:0 is a valid EUI, but this would collide with one of the IPv6 special addresses.
- u=1, when the address comes from a valid EUI, and is 0 otherwise.

To go from a 48 bit IEEE 802, you insert 0xFFFE in between the 3rd and 4th octets of an IEEE 802 address, i.e., 123456789abc becomes 123456FFFE789abc.

Special Address Formats

Unspecified address

“::” == “0:0:0:0:0:0:0:0 - can only be used as a source address by a station which does not yet have an address

Loop-back address

0:0:0:0:0:0:0:1 - used to send an IPv6 datagram to yourself

IPv4-based address

prefix the 32 bit IPv4 address with 96 zero bits

Site local addresses

Site local address can not be routed on the global internet, but they can be used by sites that are not connected to the internet or for communication within the site.

1111111011 (10 bits)	0 (38 bits)	Subnet (16 bits)	Interface ID (64 bits)

Link local addresses

Link local address are simply unique to a given link - they can be used by stations that have not yet been assigned a provider-based address.

1111111010 (10 bits)	0 (54 bits)	Interface ID (64 bits)

Multicast Addresses

1111 1111	Flags (4 bit)	Scope (4 bit)	112 bit - group id
	xxxT		

T == Transient

T = 0	well-known permanent - assigned by the IANA
T = 1	non-permanent

Scope	
0	reserved
1	node local scope
2	link local scope
3, 4	unassigned
5	site local scope
6, 7	unassigned
8	organization local scope
9, A, B, C, D	unassigned
E	global scope
F	reserved

Permanently assigned groups

For example, group 0x43 has been assigned to the Network Time Protocol (NTP), hence:

FF01::43	represents all NTP servers on the same node as the sender
FF02::43	represents all NTP servers on the same link as the sender
FF05::43	represents all NTP servers on the same site as the sender
FF08::43	represents all NTP servers within the same organization as the sender
FF0E::43	represents all NTP servers in the Internet

IANA has assigned a whole series of group identifiers, including:

FF0X:0:0:0:0:0:0	Reserved multicast address - this can not be used within any scope
FF01:0:0:0:0:0:1	All Nodes on this node address
FF02:0:0:0:0:0:1	All Nodes on this link address
FF01:0:0:0:0:0:2	All Routers on this node address
FF02:0:0:0:0:0:2	All Router address on this link

FF02:0:0:0:0:0:3	unassigned
FF02:0:0:0:0:1:1	Link Name
FF02:0:0:0:0:1:2	All DHCP agents on this link
FF02:0:0:0:0:1:3	All DHCP servers on this link
FF02:0:0:0:0:1:4	All DHCP relays on this link
FF05:0:0:0:0:1:2	All DHCP agents at this site
FF05:0:0:0:0:1:3	All DHCP servers at this site
FF05:0:0:0:0:1:4	All DHCP relays at this site
FF0X:0:0:0:0:2:7FFE	Session Announcement Protocol (SAP) v1 Announcements

Multimedia conferences:

FF0X:0:0:0:0:2:8000 .. FF0X:0:0:0:0:2:FFFF	multimedia conferences

X=2 -- this link; X=5 -- this site

Use SAP to announce the conference - repeatedly until the end of the conference.

Prefix for IPv6 documentation

The IPv6 unicast address prefix `2001:DB8::/32` is reserved for use in examples for books, RFCs, … [85].

Note that this is a **non-routable** range - to help avoid problems!

Anycast

Sending a packet to a generic address to get a specific service from the “nearest” instance. This puts the burden of determining which instance to deliver it to on the routing system.

Requires defining a router entry for each anycast address.

Subnet Anycast Address:

Subnet prefix (n bits)	0 (128-n bits)

Thus the host ID of zero is treated as the subnet.

IPv6 Routing

- all standard routing protocols
- routing extensions
  - Provider Selection
  - Host Mobility (route to current location)
  - Auto-Readdressing (route to new address)

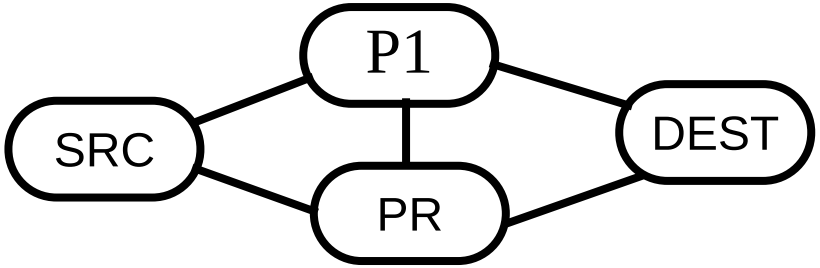

Figure 92: IPv6 Routing Option: provider specifies: SRC, PR, P1, Dest
reply: Dest, PR, P1, SRC

Routing header

Next Header (8 bits)	Header Ext Length (8 bits)	Routing Type=0 (8 bits)	Segments Left (8 bits)
reserved (32 bits)			
address[1] (128 bits)			
address[2]			
...			
address[n]			

Next Header identifies the next header in the chain of headers.

Header Ext. Length. - number of 64 bit words (not including the first 64 bits).

Routing type=0, is the generic routing header which all IPv6 implementations must support.

Number of Segments is the number of segments left in the list (between 0 and 23).

Fragment header

Next Header (8 bits)	Reserved (8 bits)	Fragment offset (13 bits)	RESERVED (2 bits)	M (1 bit)
Identification				

Fragment offset - in units of 64 bit words, the field is the most significant 13 bits of a 16 bit words.

M == More fragment bit, set in all but the last fragment

Identification - a 32 bit number

Destination Options header

Next Header (8 bits)	Header Ext. Length (8 bits)	
	Options	

Each options field is encoded as:

Option Type (8 bits)	Option Data Length (8 bits)	Option Data (n octets)

The option type:

Action (2 bits)	C (1 bit)	Number (6 bits)

Action tells what action must be taken if the processing nodes does **not** recognize the option.

Bits	**Action**
00	Skip over this option
01	Discard packet silently (i.e., without sending an ICMP report)
10	Discard packet and send an ICMP report - even if destination is multicast
11	Discard packet and send an ICMP report - only if destination is **not** multicast

C == change en route bit -- indicates that this option may be changed by intermediate relays on the way to the destination

Currently only two options are defined:

Pad1 == a null byte - for use in padding to a 64-bit boundary; note it does not have a null option length field after it - as it is the whole field

PadN - the length field says how many null bytes are needed to file to a 64-bit boundary.

Hop-by-Hop Options header

Same basic format as Destination option header, but the hop-by-hop header will be processed at each hop along the way.

Next Header (8 bits)	Header Ext. Length (8 bits)	
Options		

Each options field is encoded as:

Option Type (8 bits)	Option Data Length (8 bits)	Option Data (n octets)

Currently three options are defined: Pad1, PadN, and

- Jumbo payload option (option type =194) - the option Data Length is 4 and is followed by a 32 bit Jumbo Payload Length value.

See RFC 2113: Router Alert Option [86].

Security

- Header Authentication with signatures
  - Must have support for Message Digest 5 (MD5) algorithm [88]
- RFC 1810 [89] examines MD5 performance
- Packet Encapsulation with e.g., DES

For more information see Chapter 5 of *IPv6*, 2nd edition, by Christian Huitema.

IPSEC IPv6 implementation

The US Naval Research Lab (NRL) IPv6/IPsec Software Distribution

- a reference implementation of IPv6 and IP Security for the 4.4BSD-Lite networking software.
- Freely distributable (subject to U.S. export controls) and usable for commercial and non-commercial purposes (you must adhere to the NRL and UC Berkeley license terms) see also: *http://web.mit.edu/network/isakmp*
  - DOD ISAKMP Distribution
  - Cisco's ISAKMP Distribution
  - NRL's IPv6 + IPSEC Alpha 7.1 Distribution (Dec '98)
  - Portland State University's Mobile IP with IPSEC for FreeBSD 2.2.1.

See also the list of IPv6 implementations at:

http://playground.sun.com/pub/ipng/html/ipng-implementations.html

IPv6 ICMP [90]

Type (8 bits)	Code (8 bits)	Checksum (16 bits)
Message Body		

Currently defined ICMP Types

Type	Purpose
1	Destination Unreachable
2	Packet too big
3	Time exceeded
4	Parameter problem
128	Echo Request
129	Echo Reply
130	Group Membership Query
131	Group Membership Report
132	Group Membership Reduction
133	Router Solicitation
134	Router Advertisement
135	Neighbor Solicitation
136	Neighbor Advertisement
137	Redirect

IPv6 ICMP Error Messages

Type: 1, 2, 3, or 4:

Type	Code	Checksum	
Parameter			
As much of invoking packet as will fit - without the overall ICMP packet exceeding 576 octets			

For type 1 the code reveals the reason for discarding the datagram

IPv6 ICMP Echo Request/Reply (PING)

Type: Echo Request = 128, Echo Reply = 129

<table>
<tr><td>Type</td><td>Code</td><td colspan="2">Checksum</td></tr>
<tr><td colspan="2">Identifier</td><td colspan="2">Sequence number</td></tr>
<tr><td colspan="4">Data</td></tr>
</table>

IPv6 ICMP and groups

Three group membership messages (type 130, 131, and 132):

Type	Code	Checksum
Maximum Response Delay		Unused
Multicast Address		

The Group Membership Reduction is used when a node leaves group.

Reports are always sent to the same group address that is reported.

Maximum response delay is the time in milliseconds that the responding report messages can be delayed. Responding stations are supposed to spread their responses uniformly over this range of delays (to prevent everyone from responding at once).

Summary of IPv6 ICMP

- incorporates IPv4's ARP (via neighbor solicitation and advertisement) and IGMP (via group membership messages)
- RARP is **dropped** since BOOTP provides the same functionality
- **dropped** IPv4's Source Quench
- added Packet Too Big message to simplify learning MTU size

For more information about ICMP see: RFC 2463 [90]

DNS and IPv6

A new record type "AAAA" which contains a 128 bit address.

Just as for the "in-addr.arpa" domain used for converting Ipv4 addresses into names, IPv6 defines an "ipv6.int" domain:
thus the address 2001:0DB8:1:2:3:4:567:89ab is represented as:
b.a.9.8.7.6.5.0.4.0.0.0.3.0.0.0.2.0.0.0.1.0.0.0.8.b.d.0.1.0.0.2.IP6.INT

For further information see RFC 3596 [91].

IPv6 Transition Mechanisms

- Incremental update and deployment
  - first step: dual stack hosts and routers
  - Encapsulation of IPv6 in IPv4 packets ⇒ tunnels
- Minimal upgrade dependencies (must first upgrade DNS)
- Easy addressing (upgraded routers can use IPv4 address)
- FreeBit Co., Ltd.'s Feel6 - secure IPv6 over IPv4[82], see slide 12 [79] and *http://start.feel6.jp/*
- See also [92]

Why IPv6?

- solves Internet scaling problem
  - “eliminates” the problem of running out of addresses
  - allows route aggregation - which allows the size of the routing tables in the backbone routers to decrease
- flexible transition (interworks with IPv4)
- meets the needs of new markets
- new functionality
- real-time flows
- provider selection
- host mobility
- end-to-end security
- auto-configuration - chapter 4, “Plug and Play” in *IPv6*, 2nd edition, by Christian Huitema - this a **very major** advantage of IPv6. See also [98]

IPv6 networks

6Bone - *http://www.6bone.net/* a testbed for deployment of IPv6

- Note the phase out of the 3FFE::/16 prefix
- prefix will be returned to the unassigned address pool on 6 June 2006 [93].

vBNS - *http://www.vbns.net*

6NET *http://www.6net.org/* - project co-funded by European Commission

Euro6IX: European IPv6 Internet Exchanges Backbone

http://www.euro6ix.org/main/index.php - project co-funded by European Commission

For some issues concerning IPv6 deployment see [100]

RIR assignments of IPv6 addresses

Total number of allocated IPv6 prefixes per RIR on 13/05/2005 [102]

RIR	Size in /48s	Count
APNIC	416,579,590	376
ARIN	9,699,375	195
LACNIC	1,048,577	17
RIPE NCC	1,091,723,264	635
Total:	1,519,050,806	1,223

Further information

See:

http://www.ipv6.org/

http://www.ipv6forum.com/

Measurements of dual stack IPv6 implementations: *http://mawi.wide.ad.jp/mawi/dualstack/*

See also: [80] and [81].

Summary

This lecture we have discussed:

- IPv6

References

[78] S. Deering and R. Hinden, “Internet Protocol, Version 6 (IPv6) Specification”, IETF RFC 2460, December 1998 *http://www.ietf.org/rfc/rfc2460.txt*

[79] Jun Murai, “WIDE report”, 5th CAIDA-WIDE Workshop, Information Sciences Institute, Marina del Rey, CA, 15 March 2005

http://www.caida.org/projects/wide/0503/slides/murai.pdf

[80] Kenjiro Cho, “Measuring IPv6 Network Quality” (part 1), 5th CAIDA-WIDE Workshop, Information Sciences Institute, Marina del Rey, CA, 15 March 2005 *http://www.caida.org/projects/wide/0503/slides/kenjiro-1.pdf*

[81] Kenjiro Cho, “Measuting IPv6 Network Quality” (part 2), *Internet Iniative Japan (IIJ)* / WIDE, 5th CAIDA-WIDE Workshop, Information Sciences Institute, Marina del Rey, CA, 15 March 2005

http://www.caida.org/projects/wide/0503/slides/kenjiro-2.pdf

[82] “Trying Out for Yourself: Smooth use of IPv6 from IPv4 by Feel6 Farm”,

IPv6Style, NTT Communications, 7 March 2003

http://www.ipv6style.jp/en/tryout/20030307/index.shtml

[83] S. McCanne, V. Jacobson, M. and Vetterli, “Receiver-driven Layered Multicast”, ACM SIGCOMM, August 1996, Stanford, CA, pp. 117-130.

ftp://ftp.ee.lbl.gov/papers/mccanne-sigcomm96.ps.gz

[84] D. Borman, S. Deering, and R. Hinden, “IPv6 Jumbograms”, IETF RFC 2675 August 1999 *http://www.ietf.org/rfc/rfc2675.txt*

[85] G. Huston, A. Lord, and P. Smith, “IPv6 Address Prefix Reserved for Documentation”, IETF RFC 3849, July 2004 *http://www.ietf.org/rfc/rfc3849.txt*

[86] Dave Katz, “IPv6 Router Alert Option”, IETF RFC 2113, February 1997

http://www.ietf.org/rfc/rfc2113.txt

[87] R. Gilligan,S. Thomson, J. Bound, and W. Stevens, “Basic Socket Interface Extensions for IPv6”, IETF RFC 2133, April 1997

http://www.ietf.org/rfc/rfc2133.txt

[88] R. Rivest, “The MD5 Message-Digest Algorithm”, IETF RFC 1321, April 1992 *http://www.ietf.org/rfc/rfc1321.txt*

[89] J. Touch, “Report on MD5 Performance”, IETF RFC 1810, June 1995 *http://www.ietf.org/rfc/rfc1810.txt*

[90] A. Conta and S. Deering, “Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification”, IETF RFC 2463, December 1998 *http://www.ietf.org/rfc/rfc2463.txt*

[91] S. Thomson, C. Huitema, V. Ksinant, and M. Souissi, “DNS Extensions to Support IP Version 6”, IETF RFC 3596, October 2003

[92] C. Huitema, R. Austein, S. Satapati, and R. van der Pol, “Unmanaged Networks IPv6 Transition Scenarios”, IETF RFC 3750 , April 2004 *http://www.ietf.org/rfc/rfc3750.txt*

[93] R. Fink and R. Hinden, “6bone (IPv6 Testing Address Allocation) Phaseout”, IETF RFC 3701, March 2004 *http://www.ietf.org/rfc/rfc3701.txt*

[94] R. Hinden and S. Deering, “Internet Protocol Version 6 (IPv6) Addressing Architecture”, IETF RFC 3513, April 2003 *http://www.ietf.org/rfc/rfc3513.txt*

[95] R. Hinden, S. Deering, and E. Nordmark, “IPv6 Global Unicast Address Format”, IETF RFC 3587, August 2003 *http://www.ietf.org/rfc/rfc3587.txt*

[96] APNIC, ARIN, and RIPE NCC, "IPv6 Address Allocation and Assignment Policy", Document ID: ripe-267, January 22, 2003 *http://www.ripe.net/ripe/docs/ipv6policy.html*

[97] IAB, “IAB/IESG Recommendations on IPv6 Address Allocations to Sites”, IETF RFC 3177, September 2001 *http://www.ietf.org/rfc/rfc3177.txt*

[98] S. Thomson and T. Narten, “IPv6 Stateless Address Autoconfiguration”, IETF RFC 2462, December 1998

[99] Toshiyuki Hosaka, “IPv6 Address Allocation and Policy: PART1 IPv6 Address Basics”, Tech Tutorials, IPv6Style, NTT Communications, 18 November 2004 *http://www.ipv6style.jp/en/tech/20041117/index.shtml*

[100]Cisco, “IPv6 Deployment Strategies”, 23 December 2002 ,
http://www.cisco.com/univercd/cc/td/doc/cisintwk/intsolns/ipv6_sol/ipv6dswp.htm

[101]Gert Döring, “Impressions:An overview of the global IPv6 routing table”, RIPE 50, Stockholm, SE, 3 May 2005, *http://www.space.net/~gert/RIPE/R50-v6-table.pdf*

[102] RIPE, “Total number of allocated IPv6 prefixes per RIR on 13/05/2005”, web page accessed 2005.05.14 *http://www.ripe.net/rs/ipv6/stats/index.html*

[103]J. Rajahalme, A. Conta, B. Carpenter, and S. Deering. “IPv6 Flow Label Specification”, IETF RFC 3697, March 2004 *http://www.ietf.org/rfc/rfc3697.txt*