CHAPTER 13

Cisco and IPv6[1]

Cisco Systems, the premier IP vendor, is committed to the evolution of the Internet and of intranets and considers the next generation IP to be a key component of their growth. Cisco has taken a leadership role in the definition and implementation of the IPv6 protocols within the IETF and within its IOS™ software. Recognizing the magnitude of the migration involved, Cisco also is implementing techniques (discussed later in this chapter) that facilitate the transition from IPv4 to IPv6. Its current IOS™ implementation is in Beta, and Cisco expects to ship its comprehensive IOS IPv6 support near the end of 1998.

[1]This chapter was written with the help of Martin McNealis, IOSTM product line manager at Cisco Systems, Inc. Without his help, this chapter would not have been possible. The author wants to thank Martin for his contribution, his advice, and his friendship.

IOS runs on Cisco routers and is a very powerful router and switch operating system supporting more than 15,000 features and various protocols.

IPv6 will be one of the protocols supported by IOS, and it will be fully integrated into the operating system.

Leveraging its unparalleled experience in building the world's largest network including, of course, the Internet, Cisco has developed optimum layer 3 switching techniques such as *Cisco Express Forwarding* and *Tag Switching*, which will encompass support for IPv6.

Cisco's Express Forwarding (CEF) technology is a scalable, distributed, layer 3 switching solution designed to meet the future performance requirements of the Internet and Enterprise networks. CEF represents the ultimate advance in Cisco IOS switching capabilities, which include *NetFlow™ Switching* and *Distributed Switching*. CEF is also a key component of Cisco's *Tag Switching* architecture.

The position of Cisco—as premium IP vendor—is not to force the users to migrate to IPv6 but to enable users to decide the right moment to migrate based upon their unique network condition. For many customers, the transition to IPv6 is a decision that they won't need to make for several years. Cisco has already developed extensions to IPv4, incorporating in IPv4 many of the advantages of IPv6. For example:

- *Classless Inter-Domain Routing* (CIDR) and *Network Address Translation* (NAT) provide an effective means of resolving the current limitations of IP address assignment.
- *Virtual Private Networks* (VPNs) made with IPv4 tunnels are an effective solution for Enterprise networks and when integrated with NAT mitigate the lack of IPv4 address space.
- IPSec available in IPv4 addresses the security concerns of network managers.
- DHCP servers and relays address the need for user mobility and for plug-and-play configuration.
- *Resource Reservation Protocol* (RSVP) and *Weighted Fair Queuing* (WFQ) are among the options available for defining quality of service on existing IP networks.

In particular, NAT [1] supports the connectivity in the presence of nonunique addresses. The NAT technology enables each organization connected to the Internet to reuse the same block of addresses (for example,

the addresses defined in [2]), while requiring only a small number (relative to the total number of addresses used by the organization) of globally unique addresses for external connectivity.

Cisco recognises that continued growth of the Internet and demand for IP addressing will be fueled for example by the *Voice over IP* (VoIP), the new on-line devices such as *Personal Digital Assistants* (PDAs), hybrid mobile phones, and set-top boxes, all of which are becoming Network-aware and IP manageable and as such IPv6 provides a clear path to such expansion.

Of course, there are also some caveat and inefficiencies introduced by IPv6: while the regular and simple structure of the IPv6 header will simplify the streamline processing of packets without options, the larger header size will no longer make possible to fully contain a TCP ACK response in a single ATM cell (as in IPv4)—introducing a substantial overhead.

Another important advantage of IPv6 is the provider-based addressing, that will introduce an efficient aggregation hierarchy with the related benefits (there is a clear analogy with telephony network). With the current proposal of Top-Level Aggregator, Next-Level Aggregator, Site-Level Aggregator, etc., it is possible that the Internet core router would carry only 8,000 prefixes on the Internet backbone.

Cisco's strategy is to minimize the transition pain and leverage existing proven technology, like translation. The most likely deployment scenario will see the Enterprise first with Cisco routers performing translation for the backbone Internet until a major ISP seeks first-mover advantage.

Going forward, Cisco understands that both IPv4/NAT and IPv6 will coexist for a long period of time and, therefore, it is ready to support both of them in an integrated way in IOS.

Cisco maintains an official IPv6 web server at the following address: **`http://www.cisco.com/IPv6`**.

IPv6 in IOS™

At the time of writing, Cisco has a Beta version of IOS, which includes the IPv6 support. Information presented here is not based on the final implementation, and therefore, users are invited to read official Cisco manuals before configuring the router.

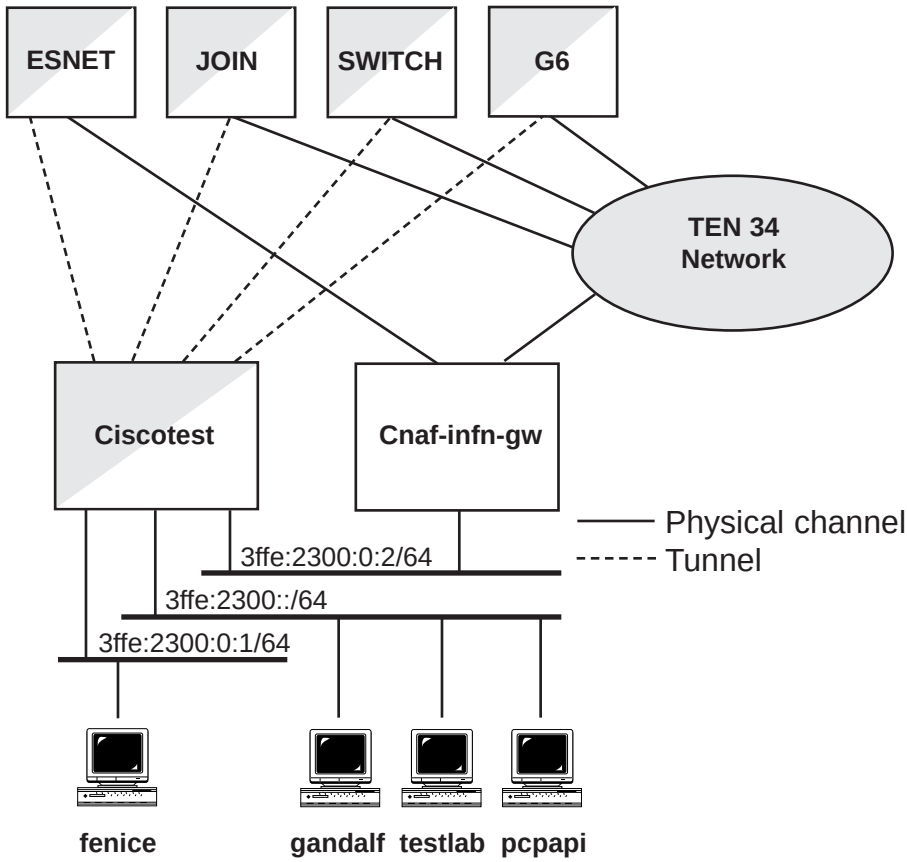

Figure 13-1
The CNAF/INFN 6bone node.

The explanation is based on an example courtesy of the backbone node of 6bone of the Italian research network (GARR), which is run in Bologna by CNAF/INFN[2]. Figure 13-1 depicts the architecture of the 6bone node present at CNAF/INFN (for details, see **`http://www.cnaf.infn.it`**).

Routers colored white and gray run both IPv4 and IPv6.

The following description is related to the node "CISCOTEST," a Cisco 7505 router, running an appropriate version of IOS.

Before going on with a description of the configuration, it is important to understand the IPv6 addressing plan of 6bone at the time of this writing.

[2] The author is in debt to the people of CNAF/INFN for their help and in particular to Antonia Ghiselli, Cristina Vistoli, and Luca dell'Agnello, who provided all the valuable information.

The CNAF/INFN asked of 6bone a *pseudo Top Level Aggregation Identifier* (pTLA) for GARR. The word "pseudo" means that this TLA will only be used during the testing phase of 6bone. 6bone is seen from IANA as a TLA, and IANA has assigned to 6bone the TLA-ID 0x1fe on 13 bits (see Figure 13-2). Adding the Aggregatable Address Format Prefix equal to 001 on 3 bits, we can derive the 6bone prefix **`3ffe::/16`** on 16 bits.

The first 8 bits of the *Next Level Aggregation* (NLA) identify all the IPv6 networks of GARR and have been set by 6bone equal to 0x23. Therefore, the IPv6 prefix of GARR is **`3ffe:2300::/24`**. GARR has assigned to CNAF/INFN the remaining 24 bits of the NLA equal to zero, and therefore the CNAF/INFN prefix is **`3ffe:2300::/48`**.

The router CISCOTEST has three Ethernet interfaces that run IPv6. A different IPv6 subnet using a different value in the SLA-ID field is associated to each Ethernet network. The three subnet prefixes are **`3ffe:2300::/64`**, **`3ffe:2300:0:1::/64`** , and **`3ffe:2300:0:2::/64`**.

Figure 13-3 lists the significant sections of the configuration file of the router CISCOTEST. The ellipsis indicates an omission of material not relevant to IPv6 configuration.

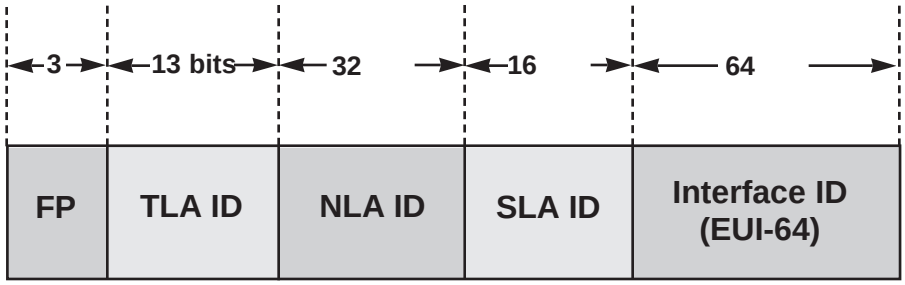

Figure 13-2 6bone Aggregatable Address

```
!
...
hostname ciscotest
!
...
ipv6 unicast-routing
ipv6 bgp redistribute connected
ipv6 bgp neighbor 3FFE:700:20:2::9 remote-as 293
ipv6 bgp neighbor 3FFE:302:11:2:0:2:0:51 remote-as 1717
ipv6 bgp neighbor 3FFE:2000:0:1::61 remote-as 559
ipv6 bgp neighbor 3FFE:401::2C0:33FF:FE02:14 remote-as 1275
ipv6 bgp network 3FFE:2300::0/24 summary
```

Figure 13-3 CISCOTEST configuration file.

continues

Figure 13-3
Continued.

```
!
interface Tunnel100
 description tunnel BGP4+ ---> ESNET
 no ip address
 ipv6 address 3FFE:700:20:2::A/126
 tunnel source FastEthernet0/0/0
 tunnel destination 198.128.2.27
 tunnel mode ipv6ip
!
interface Tunnel101
 description tunnel BGP4+ ---> IMAG
 no ip address
 ipv6 address 3FFE:302:11:2:0:2:0:52/124
 tunnel source FastEthernet0/0/0
 tunnel destination 129.88.26.7
 tunnel mode ipv6ip
!
interface Tunnel102
 description tunnel BGP4+ ---> SWITCH
 no ip address
 ipv6 address 3FFE:2000:0:1::62/124
 tunnel source FastEthernet0/0/0
 tunnel destination 130.59.15.6
 tunnel mode ipv6ip
!
interface Tunnel103
 description tunnel BGP4+ ---> JOIN
 no ip address
 ipv6 enable
 ipv6 address 3FFE:2300:0:FFFF::9/126
tunnel source FastEthernet0/0/0
 tunnel destination 128.176.191.66
 tunnel mode ipv6ip
!
interface Tunnel200
 description static tunnel ---> UNIBO
 no ip address
 ipv6 address 3FFE:2300:0:FFFF::5/126
 tunnel source FastEthernet0/0/0
 tunnel destination 137.204.198.2
 tunnel mode ipv6ip
!
interface Tunnel201
 description static tunnel ---> DEMOKRITOS
 no ip address
 ipv6 address 3FFE:2300:0:FFFF::D/126
 tunnel source FastEthernet0/0/0
 tunnel destination 192.108.114.29
 tunnel mode ipv6ip
!
...
interface FastEthernet0/0/0
 ip address 131.154.3.58 255.255.255.0
```

```
 ipv6 address 5F15:4100:839A:300:0:2E0:14C5:6B60/80
 ipv6 address 3FFE:2300::0/64 eui-64
!
interface Ethernet3/0
 ip address 131.154.100.1 255.255.255.0
 ipv6 address 3FFE:2300:0:2::0/64 eui-64
!
...
router bgp 137
!
...
ipv6 route 3FFE:401::2C0:33FF:FE02:14/128 Tunnel103
ipv6 route 3FFE:2300:31::0/48 Tunnel200
ipv6 route 3FFE:23FF::0/32 Tunnel201
!
```

IPv6 commands

Let's examine the most relevant commands of Figure 13-3 and also the output of some show commands.

show ipv6 route

```
show ipv6 route
```

This command displays the IPv6 routing table.

Figure 13-4
Output of show IPv6 route.

```
ciscotest>show ipv6 route
*** This output has been cut to fit into one page ***
IPv6 Routing Table - 120 entries
Codes: C - Connected, L - Local, S - Static, R - RIP, B - BGP
Timers: Uptime/Expires

B 3FFE:301:DEC2::0/48 [20/6]
  via FE80::C3F:5B96:B, Tunnel101, 00:01:27/never
B 3FFE:301:DEC0::0/44 [20/4]
  via FE80::C3F:5B96:B, Tunnel101, 02:01:30/never
L 3FFE:302:11:2:0:2:0:52/128 [0/0]
  via 3FFE:302:11:2:0:2:0:52, Tunnel101, 06:24:17/never
C 3FFE:302:11:2:0:2:0:50/124 [0/0]
  via 3FFE:302:11:2:0:2:0:52, Tunnel101, 06:24:17/never

...
```

continues

Figure 13-4
Continued.

```
L 3FFE:2300:0:2:2E0:14FF:FEC5:6B60/128 [0/0]
  via 3FFE:2300:0:2:2E0:14FF:FEC5:6B60, Ethernet3/0,
      06:24:30/never
C 3FFE:2300:0:2::0/64 [0/0]
  via 3FFE:2300:0:2:2E0:14FF:FEC5:6B60, Ethernet3/0,
      06:24:30/never
L 3FFE:2300:0:FFFF::5/128 [0/0]
  via 3FFE:2300:0:FFFF::5, Tunnel200, 06:24:19/never
C 3FFE:2300:0:FFFF::4/126 [0/0]
  via 3FFE:2300:0:FFFF::5, Tunnel200, 06:24:19/never
L 3FFE:2300:0:FFFF::9/128 [0/0]
  via 3FFE:2300:0:FFFF::9, Tunnel103, 06:24:19/never
C 3FFE:2300:0:FFFF::8/126 [0/0]
  via 3FFE:2300:0:FFFF::9, Tunnel103, 06:24:19/never
L 3FFE:2300:0:FFFF::D/128 [0/0]
  via 3FFE:2300:0:FFFF::D, Tunnel201, 06:24:19/never
C 3FFE:2300:0:FFFF::C/126 [0/0]
  via 3FFE:2300:0:FFFF::D, Tunnel201, 06:24:19/never
S 3FFE:2300:31::0/48 [1/0]
  via 0::0, Tunnel200, 06:24:19/never
S 3FFE:23FF::0/32 [1/0]
  via 0::0, Tunnel201, 06:24:19/never
L FE80::0/64 [0/0]
  via 0::0, Null0, 06:24:38/never
                    ciscotest>
```

show ipv6 tunnel

```
show ipv6 tunnel
```

This command displays, for each tunnel running IPv6, the tunnel unit number, the name of the dynamic routing protocol in use, the time of the last input, the number of input packets, and the description string.

Figure 13-5
Output of show IPv6 tunnel.

```
ciscotest>show ipv6 tunnel
Tun Route   LastInp   Packets Description
100   -   00:00:00     12356 tunnel BGP4+ —-> ESNET
101   -   00:00:00      6992 tunnel BGP4+ —-> IMAG
102   -   00:00:01      5841 tunnel BGP4+ —-> SWITCH
103   -   03:55:00         9 tunnel BGP4+ —-> JOIN
200   -      never         0 tunnel statico —-> UNIBO
201   -      never         0 tunnel statico —-> DEMOKRITOS
```

show ipv6 neighbors

```
show ipv6 neighbors [<ipv6addr> | <interface>]
```

This command displays neighbor adjacency entries from the IPv6 *Neighbor Discovery* (ND) table (see Section 6.5). It includes the state of the adjacency entry, its lifetime, and the associated MAC and IPv6 addresses.

Figure 13-6
Output of show IPv6 neighbors.

```
ciscotest>show ipv6 neighbors
IPv6 Address                            Age MAC Address    State
      Interface
3FFE:2300::2A0:24FF:FE99:DA7           24 00a0.2499.0da7 REACH
      FastEthernet0/0/0
FE80::2A0:24FF:FE99:DA7                 24 00a0.2499.0da7 REACH
      FastEthernet0/0/0
```

show ipv6 interface

```
show ipv6 interface [<interface>]
```

This command displays IPv6 interface related parameters and addresses.

Figure 13-7
Output of show IPv6 interface.

```
ciscotest>  show ipv6 int FastEthernet0/0/0
FastEthernet0/0/0 is up, line protocol is up
  IPv6 is enabled, link-local address is
      FE80::2E0:14FF:FEC5:6B00
  Global unicast address(es):
    5F15:4100:839A:300:0:2E0:14C5:6B60, subnet is
      5F15:4100:839A:300::0/80
    3FFE:2300::2E0:14FF:FEC5:6B00, subnet is 3FFE:2300::0/64
  Joined group address(es):
    FF02::1
    FF02::2
    FF02::1:FFC5:6B00
    FF02::1:FFC5:6B60
    FF02::1:14C5:6B60
    FF02::1:FEC5:6B00
  MTU is 1500 bytes
  ICMP error messages limited to one every 500 milliseconds
  ND advertised reachable time is 0 milliseconds
  ND advertised retransmit interval is 0 milliseconds
  ND router advertisements are sent every 200 seconds
  ND router advertisements live for 1800 seconds
  Hosts use stateless autoconfig for addresses.
                    ciscotest>
```

show ipv6 traffic

```
show ipv6 traffic
```

This command displays IPv6 related traffic statistics.

traceroute ipv6

```
traceroute ipv6 <destination>
```

This command traces the route for IPv6 packets between the node where the command is entered and the destination address.

ping ipv6

```
ping ipv6 <destination>
```

This command sends ICMPv6 echo request packets (see Sections 5.6.1 and 5.6.2) to **<destination>**, i.e., to an IPv6 host name or address.

ipv6 unicast-routing

```
ipv6 unicast-routing
```

This command enables the routing of IPv6 unicast packets. The default setting is disabled.

interface tunnel

```
interface tunnel
```

Tunneling provides a way to encapsulate arbitrary packets inside another protocol (see Section 12.2). It is implemented as a virtual interface to provide a simple configuration.

In the preceding example it is used to create an IPv6 tunnel over IPv4. The IPv4 end-points are specified with the commands:

- tunnel source **<interface>**
- tunnel destination **<IPv4 address>**

Because tunnels are point-to-point links, a separate tunnel is configured for each link.

The command **no ip address** specifies that there is no IPv4 address associated to this tunnel, while the command **ipv6 address <IPv6 address>** assigns an IPv6 address to the tunnel interface. Finally, the command **tunnel mode ipv6ip** configures a *static* tunnel interface (a "configured tunnel" according to RFC 1933 [1]). This interface can be used like any other interface (static routes can point to it or a dynamic routing protocol can run over it).

ipv6 address

```
[no] ipv6 address <ipv6addr>[/<prefix-length>]
```

This command enables IPv6 and configures an IPv6 address on the interface. Optionally, a prefix length may be specified. In this case the router will autoconfigure the remaining bits.

ipv6 address ... eui-64

```
[no] ipv6 address <ipv6prefix>/<prefix-length> eui-64
```

This command is used to enable IPv6 and to autoconfigure an IPv6 address on an interface using the EUI-64 style "Interface ID" (see section 4.10). If the **<prefix-length>** specified is greater than 64, the prefix bits will have precedence over the EUI-64 ID.

ipv6 unnumbered

```
[no] ipv6 unnumbered <interface>
```

It is also possible to enable and to configure an interface without requiring a global IPv6 address. The **<interface>** parameter must specify the name of an interface that does have a global IPv6 address. This command is used to reduce address administration for a network administrator.

ipv6 route

```
[no] ipv6 route <prefix> {<next-hop> | <interface>} [<distance>]
```

This command configures a static IPv6 route. **<prefix>** specifies the IPv6 prefix for which the route is created. **<next-hop>** is the host name or IPv6 address of the next-hop to reach the destination prefix. **<interface>** can be used in place of **<next-hop>** for point-to-point interfaces like serial links or tunnels. The default value for **<distance>** is 1, which gives static routes precedence over any other type of route with the exception of directly connected routes.

ipv6 mtu

```
[no] ipv6 mtu <bytes>
```

This command configures the *Maximum Transmission Unit* (MTU) for IPv6 packets on an interface. The default value is the link MTU. If a non-default value is configured, an MTU option will be included in Router Advertisements (see Section 5.6.5).

ipv6 hop-limit

```
ipv6 hop-limit <value>
```

This command configures the router to use **<value>** as the IPv6 Hop Limit value used in Router Advertisements (see Section 5.6.5) and in all IPv6 packets generated within the router. The default value is 255.

ipv6 auto-tunnel

```
[no] ipv6 auto-tunnel
```

This command configures IPv6 in IPv4 automatic tunneling (see RFC 1933 [1]). Automatic tunneling is performed when a destination address in an IPv6 packet contains an IPv4 compatible IPv6 address (see section 4.7.8).

RIP Protocol

The Cisco implementation of IPv6 supports RIPv6 (see section 7.5.1). RIP routing is started whenever RIP is enabled on at least one interface. It is also possible to redistribute static routes over RIP.

BGP4+

During the standardization process of IPv6, it was decided to adopt IDRPv2 as Exterior Routing Protocol (see Section 7.5.2). This new protocol, derived from OSI, has been implemented by some companies, but it does not seem to gain acceptance among users. Cisco's decision to implement a generalized BGP rather than IDRP was based upon the fact that the Service Provider community preferred to leverage a time-proven/deployed protocol with integrated support for IPv4 and IPv6 rather than run another protocol in the ships-in-the-night mode. This was a very realistic, pragmatic approach to deployment which Cisco wholly endorsed with the support of BGP4+ or more formally "Multiprotocol Extensions for BGP-4" [2]. BGP4+ defines extensions to BGP-4 to enable it to carry routing information for multiple Network Layer protocols (e.g., IPv6, IPX, etc...). The extensions are backward compatible—a router that supports the extensions can interoperate with a router that doesn't support the extensions.

To configure BGP4+ it is therefore necessary first to configure and start the IPv4 BGP with the classical command:

```
router bgp <as-number>
```

The definition of IPv6 neighbors and parameters is however done in a different section of the configuration file. The principal commands used are described in the following sections.

ipv6 bgp redistribute connected

```
[no] ipv6 bgp redistribute connected
```

This command configures the redistribution of routing information learned on directly connected networks into bgp.

ipv6 bgp redistribute static

```
[no] ipv6 bgp redistribute static
```

This command configures the redistribution of static routes into bgp.

ipv6 bgp redistribute rip

```
[no] ipv6 bgp redistribute rip
```

This command configures the redistribution of routes learned via rip process into bgp.

ipv6 bgp neighbor

```
ipv6 bgp neighbor <IPv6 address> remote-as <as-num>
no ipv6 bgp neighbor remote-as
```

This command defines a BGP neighbor. External neighbors must be directly connected. Neighbors must be specified by global addresses.

ipv6 bgp network

```
[no] ipv6 bgp network <prefix>
```

This command originates a BGP route for each route found on the IPv6 routing table that matches with the given prefix.

NAT

Chapter 12 presents the migration from IPv4 to IPv6 and explains that NAT between IPv4 and IPv6 is not a mandatory feature according to IETF.

Cisco decided to provide the NAT feature related to both protocols and address translation between IPv6 and IPv4 in IOS from the beginning. This is an important value-added feature that will greatly simplify the introduction of IPv6 in Enterprise Networks.

NAT devices would enable the interconnection of hosts that have IPv6-only addresses (hosts that do not have IPv4-compatible addresses) with hosts that have IPv4-only addresses. If assigning globally unique IPv4 addresses would become impossible (due to the exhaustion of the IPv4 address space) before a sufficient number of the Internet hosts would transition to IPv6, then NAT devices would allow continuing (and completing) the transition, even in the absence of the globally unique IPv4 addresses.

Cisco IPv6 NAT is designed to allow an IPv6 network to access and be accessed by the IPv4 Internet.

CONCLUSIONS

With the design decision made in the implementation of IPv6, Cisco confirms to be the leading company in IP routing. The Internet is today mostly powered by Cisco routers and so are many Intranets. The Cisco implementation of IPv6 will greatly simplify the migration phase from IPv4 to IPv6 and the unavoidable coexistence of IPv4 and IPv6 nodes. From the end of 1998 IPv6 will be a standard feature of Cisco's strategic IOS-based routing and switching platforms.

REFERENCES

[1] P. Francis, K. Egevang, *The IP Network Address Translator (NAT),* RFC 1631, May 1994.

[2] Y. Rekhter, B. Moskowitz, D. Karrenbergde, G. Groot, E. Lear, *Address Allocation for Private Internets*, RFC 1918, February 1996.

[3] R. Gilligan, E. Nordmar, *RFC 1933: Transition Mechanisms for IPv6 Hosts and Routers*, April 1996.

[4] T. Bates, R. Chandra, D. Katz, Y. Rekhter, *Multiprotocol Extensions for BGP-4,* <draft-bates-bgp4-multiprotocol-03.txt>, July 1997.