

IPv6 Local Network Protection with Cisco IOS Routers and Switches

Last updated: May 2008

The primary purpose of Network Address Translation (NAT) is to extend the available IPv4 address space by allowing multiple devices on a private network to share a single public IP address. This can introduce complications and cause issues with a number of applications, yet it is perceived that using NAT can provide a number of useful management and security benefits to a network. IPv6 was designed with a significantly larger address space than IPv4. As a result, NAT is not needed with IPv6 and no IPv6 NAT standard has been developed. The IETF informational document “Local Network Protection (LNP) for IPv6” (RFC 4864) describes how IPv6 can provide management and security benefits that are equal to or better than those that can be achieved with IPv4 NAT, without the need for address translation. This document outlines how to implement Local Network Protection using Cisco IOS Software.

A Practical Guide to Achieving “NAT-like” Enterprise Network Security with IPv6 Summary

Due to historical and market reasons, home and Enterprise networks are often connected to the IPv4 Internet through a Network Address Translation (NAT) middleware device. The motivation for the use of NAT devices arose from a need to extend the available address space due to the ongoing depletion of available IPv4 addresses. Initially, NAT was seen as a quick and short term solution for the IPv4 address depletion problem. The long-term solution was the development of a new version of the Internet Protocol which would provide a much larger address space. Efforts to develop a new version of IP resulted in the creation of IPv6, which increases the address size to 128 bits from the 32 bit addresses used with IPv4. In addition to an enlarged address space, IPv6 provided other enhanced capabilities such as auto-configuration, improved multicast, removal of broadcast packets, address scoping, and device mobility.

IPv6 was designed with the intention of making NAT unnecessary. Therefore, no IPv6 NAT standard has been written describing how to handle translation of specific IPv6 features, such as option headers. Many perceived benefits obtained by using an address translator have been widely publicized, while, even though they have been outlined in several IETF RFCs, the negative impacts of NAT on applications are not as widely known. It is sometimes thought that by moving from IPv4 to IPv6, you will lose some of the added functionality provided by NAT. This document, based on RFC 4864, describes some of the uses of an IPv4 address translator and demonstrates how to achieve the same goals using native IPv6 on Cisco IOS Routers and Switches.

As far as security and privacy are concerned, this document considers how to mitigate a number of threats, but should not be considered exhaustive. Some threats are obviously external, such as a hacker or a worm-infected outside machine trying to penetrate and attack the local network. Some threats are local, such as a disgruntled employee disrupting business operations or the unintentional negligence of a user downloading malware, which then proceeds to attack the

network from within. Other threats may be inherent in the device hardware ("embedded"), such as the firmware in a domestic appliance "phoning home" to its manufacturer without the user's consent.

This document will also consider the view that NAT can be used to fulfill the goals of a security policy. On the one hand, NAT does satisfy some policy goals, such as topology hiding, but at the same time it defeats others, such as the ability to produce an end-to-end audit trail at the network level. That being said, there are artifacts of NAT devices that do provide some value:

1. The need to establish state before anything gets through from the outside to the inside.
2. The expiration of state to stop receiving any packets when finished with a flow.
3. The ability for nodes to appear to be attached at the edge of the network.
4. The ability to have addresses that are not publicly routed.

This paper will address the above aspects and will demonstrate how they can be achieved using Cisco IOS Routers and Switches.

Technology Analysis between IPv4 NAT and IPv6

The following table summarizes the perceived benefits from using address translation with IPv4 and how to obtain the same benefit when using IPv6 without address translation.

Table 1. RFC4864—Market perceived benefits of IPv4 NAT

Function	IPv4	IPv6
Simple Gateway	DHCP: Single address upstream DHCP: Limited number of individual devices downstream	DHCP-PD: Arbitrary length customer prefix upstream SLAAC via RA downstream
Simple Security	Filtering side effect due to lack of translation state	Explicit Context Based Access Control (Reflexive ACL)
Local Usage Tracking	NAT state table	Address uniqueness
End System Privacy	NAT transforms device ID bits in the address	Temporary use privacy addresses
Topology Hiding	NAT transforms subnet bits in the address	Untraceable addresses using IGP host routes /or MIPv6 tunnels for stationary
Addressing Autonomy	RFC 1918	RFC 3177 and ULA
Global Address Pool Conservation	RFC 1918	340,282,366,920,938,463,463,374,607,431,768,211,456 (3.4*10 ³⁸) addresses
Renumbering and Multihoming	Address translation at the border	Preferred lifetime per prefix and multiple addresses per interface

Simple Gateway

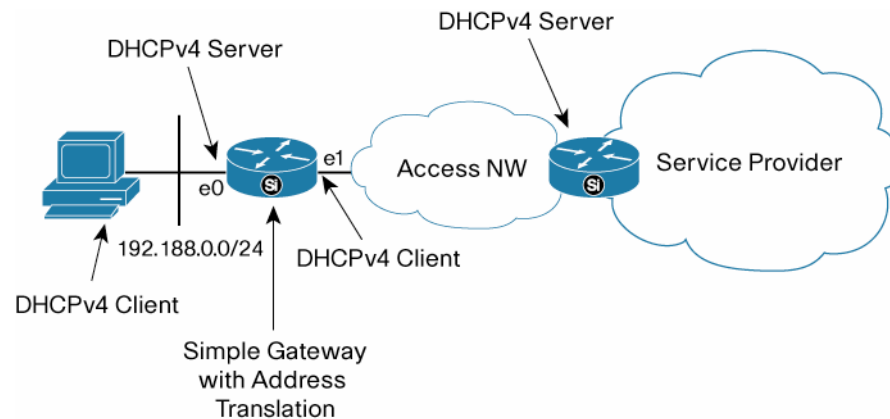
Wide-scale deployments have shown that NAT technology can provide simple gateway functionality in an IPv4 network. This functionality allows for example, the usage of private IPv4 addresses (RFC1918) on the local network in combination with a global IPv4 address assigned on the outside of the end-user edge router. Often, this outside address is a dynamically learned Internet IPv4 address.

When a user wants to connect from the local network to the Internet, the 'Simple Gateway' device will intercept the packet and may translate the original private IPv4 addresses, into unique IPv4 addresses without user intervention. The NAT component of the 'Simple Gateway' solution complements the dynamical behavior of the outside address assignment (ie: DHCPv4) in such a

way that there is no need for additional human interaction, when connecting a local network to the Internet or to external networks.

Below is an example IOS configuration for a router acting as a gateway to connect a local network to the Internet. The router is running NAT and DHCP. On the outbound interface, possibly connected to ie: a cable modem, the DHCP-client is used to learn an IPv4 address from the Service Provider. The router is configured to be a DHCP server to the local network. The IPv4 DHCP server will provide the local network IPv4 addresses within the range 192.168.0.0/24. The NAT functionality will translate these addresses (the configured range is 192.168.0.0/24) to use the dynamically learned IPv4 address on Ethernet1.

Figure 1. IPv4 Addresses Configured Dynamically by DHCPv4



```

!
ip dhcp excluded-address 192.168.0.254
!
ip dhcp pool DHCPPOOL
  import all
  network 192.168.0.0 255.255.255.0
  default-router 192.168.0.254
  dns-server 10.0.0.21 10.0.0.22
!
no ip dhcp-client network-discovery
!
interface Ethernet0
  ip address 192.168.0.254 255.255.255.0
  ip nat inside
!
interface Ethernet1
  ip address dhcp
  ip nat outside

```

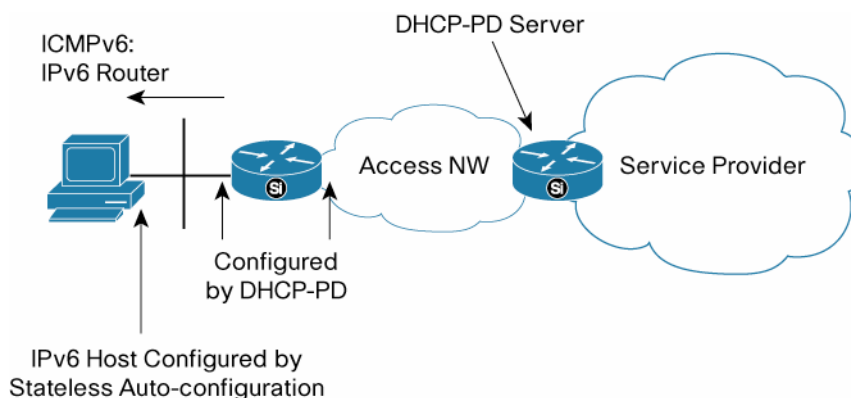
```

ip nat inside source route-map NAT_MAP interface ethernet1 overload
access-list 100 permit ip 192.168.0.0 0.0.0.255 any
route-map NAT_MAP permit 10
  match ip address 100
!

```

While there is no address translation functionality defined for IPv6, the goal of a simple gateway to allow automatic addressing without human intervention can also be achieved when using DHCPv6 with Prefix delegation extensions. DHCP exists for both IPv4 (RFC2131) and IPv6 (RFC 3315). These protocols typically have a DHCP-client receive an IP address from a DHCP-server. In addition to this functionality, the DHCP technology has been extended in IPv6 to not only exchange IPv6 addresses, but to also exchange a complete range of IPv6 addresses. This technology is standardized as 'RFC3633: IPv6 Prefix Options for Dynamic Host Configuration Protocol (DHCP) version 6'. By using IPv6 Prefix Delegation (PD) it is possible for a router to receive from its Service Provider an IPv6 prefix and use this learned prefix on its local network interfaces, so that all devices within the local network can receive a unique IPv6 address.

Figure 2. IPv6 Addresses Configured Dynamically by DHCP-PD



To achieve a similar goal with IPv6 as shown in the example above for IPv4, the following configuration can be used:

```

!
ipv6 cef
!
ipv6 dhcp pool foo
  dns-server 2001:4::1
  domain-name cisco.com

interface Ethernet0
  ipv6 address autoconfig
  ipv6 dhcp client pd PREFIX1
!
interface Ethernet1

```

```
ipv6 address PREFIX1 ::1:0:0:0:1/64
ipv6 nd other-config-flag
ipv6 dhcp server foo
!
ipv6 route ::/0 Ethernet0
!
```

With this IPv6 configuration, the hosts within the local network will receive an IPv6 address from the router. The router learned this address from its connection to the Service Provider. The learned prefix will typically be globally unique, and hence, there is no need for address translation. The solution works by using a combination of basic Stateless Address Auto-configuration (SLAAC) and Stateless DHCPv6.

1. The router has the DHCP-PD feature enabled on Ethernet0 and makes a PD request to the upstream ISP router. The requesting router will add a DHCPv6 Unique Identifier (DUID) inside this request.
2. The upstream ISP router will see this request and can delegate a prefix to the requesting router. In this example, we assume that the IPv6 prefix is 48 bits long. (We will assume this prefix to be 2001:DB8:1::/48). The upstream will answer the request from the requesting router and will delegate the IPv6 prefix.
3. The requesting router received the 2001:DB8:1::/48 prefix. Locally, the Cisco IOS Router will abstract this 48 bit IPv6 prefix as 'PREFIX1' (this is a user defined name).
4. On Ethernet1, this abstracted name 'PREFIX1' is used to provide a globally unique IPv6 address on the interface. The final address in this example assigned to Ethernet1 is 2001:DB8:1:1:0:0:0:1 with a prefix length of 64 bits.
5. Next, the router will start automatically announcing this prefix in its periodic IPv6 Router Advertisement packets. Router Advertisements are standard IPv6 packets that inform all connected IPv6 hosts that there is a router connected to the network. In addition to this function, there is extra information contained inside these packets that allows the connected hosts to get an IPv6 address assigned, as described in 'RFC2462—IPv6 Stateless Address Autoconfiguration.'
6. The RA packets also contain flags (M and O bits in the RA packet) to indicate whether the connected hosts should use DHCPv6 for requesting a managed IPv6 address and/or need to request additional information through DHCPv6 (for example DNS server, domain name, etc.). The 'O' bit is set by 'ipv6 nd other-config-flag' command under the router interface connecting to the IPv6 hosts. This will instruct the connected IPv6 hosts to request through DHCPv6, additional information (DNS and domain name) without requesting a managed IPv6 address. The IPv6 address used by the IPv6 hosts in the example is configured with the help of Stateless Address Auto-Configuration (SLAAC). An alternative choice would be the use of a DHCPv6 server to assign managed IPv6 addresses to the hosts. This is a user decision. To set the 'O' flag, the command 'ipv6 nd managed-config-flag' can be used. This flag will instruct the connected IPv6 hosts to use DHCPv6 to obtain an IPv6 address, instead of using SLAAC technology (Note that Cisco CNR has IPv6 DHCP-server support).
7. In the example, the hosts receive the router announcements with the 'O' bit set, but with the 'M' bit cleared and they will use the SLAAC mechanism to automatically create an IPv6 address and will use DHCPv6 to obtain DNS and domain name.

The above IPv6 scenario will provide 'simple gateway' functionality without the need to manipulate the IP header fields due to address translation technology. The dynamic behavior of the address assignment through DHCPv6-PD is in such a way that there is no need for additional human interaction when connecting the local network to the Internet, or to external networks.

Detailed information about Cisco DHCPv6 can be found at:

http://www.cisco.com/en/US/products/ps6553/products_white_paper09186a00801e199d.shtml

Another use case of IPv4 NAT's simple gateway functionality is where the local network is connected to an upstream Service Provider and receives exactly 'one' IPv4 address from this provider. On the local network side of the edge router, private IPv4 addresses are used. The majority of home-networks use this design and they tend to work sufficiently for Internet connectivity. Using this design will cause the original IPv4 address of the user to be translated each time by the edge router. If a user would like to make a server or a service available to the Internet, then the user will need to configure IPv4 NAT port mapping on the edge router. If, for example, the user would like to make a HTTP server available, then he or she will have to map tcp port number 80 to an assigned internal private IPv4 address. This mapping can only be done once for each external UDP or TCP port number. If the user would like to add a second HTTP server, he or she will have to choose a different external UDP or TCP port to make the server available to the Internet, because port 80 cannot be mapped a second time. One can conclude that the more servers or services are offered to the Internet with this network design, the more complex the configuration becomes. In addition, there is need for a configuration change for each additional server or service added to the local network.

Using IPv6 and simple gateway functionality will remove the incremental configuration complexity and also the need to reconfigure the edge router with each added server or service. With IPv6 all devices attached to the local network will have globally reachable IPv6 address, assuming that the local IPv6 network received from its upstream service provider a global IPv6 address range. Each connected device can use SLAAC (Stateless Address Auto-Configuration) to create a globally unique IPv6 address, while DHCPv6 can be used to assign DNS and other details to the device. The device is reachable from the internet via its unique global IPv6 address. There is no need for port mapping or incremental configuration for each added server or service.

Simple Security

An IPv4 address translator is a stateful device. When a user on a private local network wants to connect through a NAT to a device on the outside of this network or to the Internet, then the NAT creates a translation slot. This translation slot is in essence a memory structure, which maps the packet's source and/or destination IPv4 address into the IPv4 addresses seen in the packet on the outside of the network. NAT devices typically also have Application Level Gateway functionality and keep state about translations that are needed within the payload of the IPv4 packet. This translation slot will allow the return traffic to be translated back to the inside addresses of the network. Return traffic without a translation slot can not be translated and will typically be dropped at the address translator.

This functionality leads to the belief that the act of translating address bits within the header does not provide security in itself. The perceived security from NAT does not come from the act of translating the address itself, but from the lack of a pre-established or permanent mapping state. On a Cisco router, these translation slots can be viewed by using the 'show ip nat translations' command:


```
NAT_Router# show ip nat translations
```

```
Pro  Inside global      Inside local      Outside local     Outside global
udp  10.69.233.209:1220  192.168.1.95:1220  172.69.2.132:53  172.69.2.132:53
tcp  10.69.233.209:11012 192.168.1.89:11012 172.69.1.220:23  172.69.1.220:23
tcp  10.69.233.209:1067  192.168.1.95:1067  172.69.1.161:23  172.69.1.161:23
```

The ability to filter based on a pre-established or permanent mapping is a function typically provided by a firewall, but this pre-established state does not protect against abuse of the translation slot. Any hacker can send traffic through the NAT device by using the established translation slots. In addition to the pre-established state filtering, a firewall will also perform packet inspection for technologies such as ftp, http, and DNS. This functionality provides a much more valuable security infrastructure, but requires slightly more operational investment.

Both forms of simple security mentioned above are also available in IPv6. Cisco IOS and the PIX both provide extensive IPv6 firewall support. An example Cisco IOS Firewall configuration can be found below.

Figure 3.

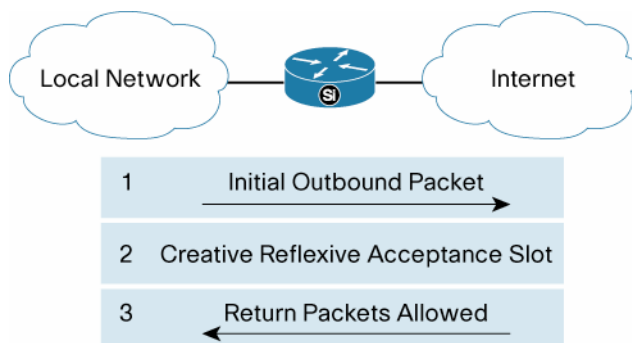


```
interface ethernet0
  ipv6 address 2001:DB8:100:F::1/64
  ipv6 inspect FW-TRAF in
!
interface serial 0
  ipv6 address 2001:DB8:E::1/64
  ipv6 traffic-filter TRAF in
!
ipv6 access-list TRAF Permit udpany host 2001:DB8:E::20 eq53
ipv6 access-list TRAF deny ipv6 any any
!
ipv6 inspect name FW-TRAF tcp
ipv6 inspect name FW-TRAF udp
ipv6 inspect name FW-TRAF ftp
!
```

When the network user is only interested in providing filtering based on a pre-established state, then he or she may choose to use a Cisco IOS reflexive Access-List (ACL). Reflexive access-lists

exist for both IPv4 and IPv6. A reflexive ACL is an ACL that allows a router to control traffic flows based on session initiation.

Figure 4.



The principle of a reflexive access-list is demonstrated in the figure above. It provides the same capability of filtering based on a pre-established mapping that can be done by using an address translator.

A “normal” ACL has no sense of state. Packets are evaluated individually, so there is no concept of a collection of packets or sessions. In an HTTP connection, for example, packets flow both ways. The HTTP client initiates the session by sending packets to the HTTP server, and then the server sends packets to the HTTP client. With simple ACLs, the router must have an ACL to permit each flow.

With a reflexive ACL, we’ll see that the “return traffic” ACL can be written on-the-fly and closed after the session has concluded. For example, examine the case where you have traffic entering the router from FastEthernet0/0, headed out FastEthernet0/1, destined for an HTTP server on port 80 on the Internet. With a reflexive ACL applied to FastEthernet0/1, you can allow that traffic out to the Internet, and dynamically create a rule allowing the return traffic from the HTTP server back to the requesting node. The inbound FastEthernet0/1 rule created would be temporary, and as specific as the outbound FastEthernet0/1 rule. To make the same traffic flow above without using Reflexive ACL, you would need two static ACLs: one on FE0/1 for the outbound traffic and one on FE0/1 for the inbound traffic. It’s better to dynamically and temporarily create the FE0/1 inbound rule.

```

!
ip reflexive-list timeout 120
!
ipv6 access-list MY-REFL-OUT-LIST
  permit tcp any any eq 80 reflect REFLECT-OUT
ipv6 access-list MY-REFL-IN-LIST
  evaluate REFLECT-OUT
!

```

```

int f0/1

  ipv6 traffic-filter MY-REFL-OUT-LIST out
  ipv6 traffic-filter MY-REFL-IN-LIST in
!

```

Additional information and configuration examples for reflexive access-lists can be found at: http://www.cisco.com/en/US/docs/ios/12_2t/ipv6/v6_tffw_ps6350_TSD_Products_Configuration_Guide_Chapter.html

Local Usage Tracking

Another perceived benefit from using NAT is the ability for the local network administrator to track user and application traffic. The idea is that by using the NAT translation table, the administrator will be able to determine which Internet locations are being accessed by users on the private network. While some rudimentary user tracking information can be obtained by capturing the NAT translation table, the transient nature of the table makes obtaining viable user tracking data more difficult than it initially appears. Depending on the NAT configuration used, tracking users with the translation table may not even be possible.

When ordinary NAT is configured each global address for the Internet can only be used by one local address on the private network at a time. Only a single translation is created for each private address, even if traffic is being sent to multiple outside destinations.

```

router# show ip nat translations

Pro Inside global    Inside local    Outside local    Outside global
--- 172.16.6.14      10.10.10.4     ---             ---
--- 172.16.11.70    10.10.50.4     ---             ---

```

As a result no outside address information is contained within the NAT translation table and detailed user tracking is not possible. The administrator can discover which internal addresses are connecting to the outside network, but will be unable to determine where the traffic is going by using NAT.

When NAT overloading, or Port Address Translation (PAT), is configured each global address can be used by multiple local addresses at the same time. This is accomplished by using unique source port numbers for each translation.

```

router# show ip nat translations

Pro Inside global      Inside local      Outside local      Outside global
udp 172.16.233.209:1220 192.168.1.95:1220 172.16.2.132:53    172.16.2.132:53
tcp 172.16.233.209:11012 192.168.1.89:11012 172.16.1.220:23    172.16.1.220:23
tcp 172.16.233.209:1067 192.168.1.95:1067 172.16.1.161:23    172.16.1.161:23

```


A separate translation entry is created for each source/destination pair. This allows the administrator to discover what outside addresses the users are contacting. However, due to the transient nature of the address translations, the NAT translation table is continually changing. In order to capture a complete set of user tracking information, the administrator will need to obtain snapshots of the translation table at frequent intervals.

The preferred way to perform user/application tracking is to use a facility, such as Cisco IOS NetFlow, which was designed with that purpose in mind. By capturing packet flow data over time

using Cisco IOS NetFlow, the administrator can obtain detailed information on what external locations are being contacted by each internal address. NAT can still create issues when using NetFlow to perform user/application tracking. When dynamic NAT is used a single global address will likely represent different local addresses at different times. In order to determine which user the global address represents at any given time, the administrator will need a time-correlated list of the address and port mappings created by NAT.

The greatly increased address space available with IPv6 renders the use of address translation unnecessary. All addresses used for communication will be unique between local devices and to devices outside the local network. Each data traffic stream can be uniquely identified thanks to the globally unique IPv6 addresses for both the source and destination. This helps to resolve some of the issues with using NetFlow to track users created by IPv4 NAT.

Figure 5.



Since IPv6 does not use NAT, the “free” rudimentary user tracking provided by the translation table is not available. Therefore, an outside user tracking facility must be used. In the following example, Cisco IOS NetFlow is used for this purpose. In the network below, hosts on the private network are given IPv6 addresses using the 2001:DB8::/64 prefix.

NetFlow is configured on the gateway device to capture data for outbound IPv6 flows on the public Internet facing interface.

```
ipv6 flow-export version 9
ipv6 flow-export source Loopback0
ipv6 flow-export destination 10.0.0.50 9991
!
interface FastEthernet1/0
  description Connection to Internet
  ipv6 address 2001:DB8:10::100/64
  ipv6 flow egress
```

When IPv6 traffic leaves the private network via the FastEthernet1/0 interface NetFlow will capture information about each of the traffic flows. By using the "show ipv6 flow cache" command the administrator can view information about the current IPv6 traffic flows.

```
Router#sh ipv6 flow cache
IP packet size distribution (3635 total packets):
  1-32  64  96 128 160 192 224 256 288 320 352 384 416 448 480
  .000 .004 .002 .993 .000 .000 .000 .000 .000 .000 .000 .000 .000 .000 .000
.000
  512 544 576 1024 1536 2048 2560 3072 3584 4096 4608
  .000 .000 .000 .000 .000 .000 .000 .000 .000 .000 .000
```

```

IP Flow Switching Cache, 475168 bytes
  11 active, 4085 inactive, 38 added
  10817 ager polls, 0 flow alloc failures
  Active flows timeout in 10 minutes
  Inactive flows timeout in 600 seconds
IP Sub Flow Cache, 33992 bytes
  0 active, 1024 inactive, 0 added, 0 added to flow
  0 alloc failures, 0 force free
  1 chunk, 1 chunk added
SrcAddress      InpIf  DstAddress      OutIf  Prot SrcPrt DstPrt Packets
FE80::A...01:F710 Local  FE80::A...1:F900 Fa1/0  0x3A 0x0000 0x8800 2
FE80::A...01:F710 Local  FE80::A...1:F800 Fa1/0  0x3A 0x0000 0x8800 2
2001:DB8:10::100 Local  FE80::A...1:F800 Fa1/0  0x3A 0x0000 0x8800 2
2001:DB8:10::100 Local  FE80::A...1:F900 Fa1/0  0x3A 0x0000 0x8800 2
2001:DB8::1      Fa0/0  2001:DB8:E::1    Fa1/0  0x3A 0x0000 0x8000 750
2001:DB8::2      Fa0/0  2001:DB8:E::1    Fa1/0  0x3A 0x0000 0x8000 675
2001:DB8::1      Fa0/0  2001:DB8:E::2    Fa1/0  0x3A 0x0000 0x8000 925
2001:DB8::2      Fa0/0  2001:DB8:E::2    Fa1/0  0x3A 0x0000 0x8000 1200
FE80::A...01:F710 Local  FE80::A...1:F900 Fa1/0  0x3A 0x0000 0x8700 2
FE80::A...01:F710 Local  FE80::A...1:F800 Fa1/0  0x3A 0x0000 0x8700 2
FE80::A...01:F710 Local  FF02::1          Fa1/0  0x3A 0x0000 0x8600 2

```

The output from this command shows traffic flows from the internal addresses 2001:DB8::1 and 2001:DB8::2 to the external addresses 2001:DB8:E::1 and 2001:DB8:E::2.

The “show ipv6 flow cache” command only displays the current contents of the NetFlow cache for IPv6 traffic. By combining the NetFlow capabilities of Cisco IOS Software on the gateway device with an external tool to gather and analyze the data captured by the router, an administrator can obtain a complete picture of what external addresses are being contacted and which internal addresses are contacting them.

NetFlow will provide details about the IPv6 addresses, but since each user or device may have multiple addresses it will still be up to the administrator to match the IPv6 address with its user. In some cases, such as when short-lifetime privacy addresses are used, additional information will be needed in order to match an IPv6 address with a particular user.

For more information:

Implementing NetFlow for IPv6:

http://cisco.com/en/US/docs/ios/12_2t/ipv6/v6ntflw_ps6350_TSD_Products_Configuration_Guide_Chapter.html

Cisco NetFlow Monitoring Applications:

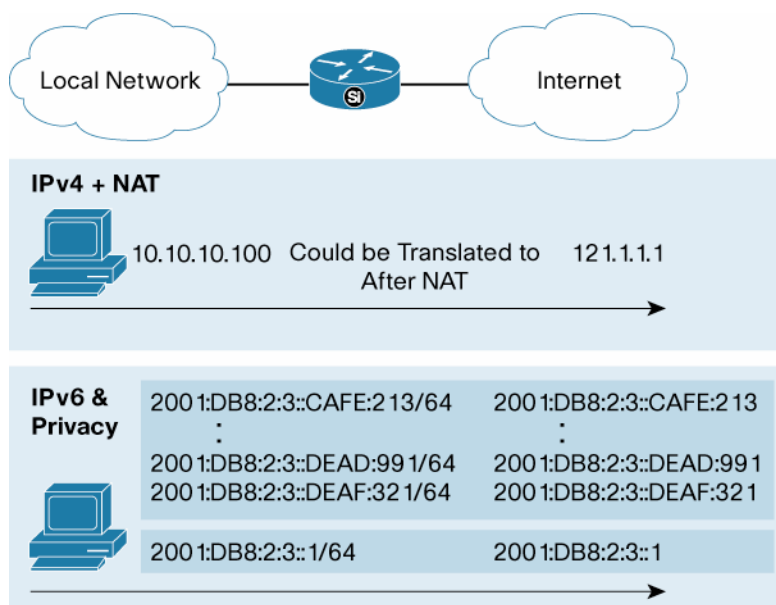
http://www.cisco.com/en/US/prod/iosswrel/ps6537/ps6555/ps6601/networking_solutions_products_genericcontent0900aec805ff720.html

NetFlow on Cisco.com: <http://www.cisco.com/go/netflow>

End System Privacy

When a user on the local network is communicating through a NAT device to non-local devices, there is a masking of the users source IPv4 address. This makes the user anonymous on the non-local network because there is no direct correlation between the IPv4 address used and the user's identity due to the address translation.

Figure 6. End System Privacy



Concern about the ability to map an IP address to a specific user by untrusted networks was a privacy concern for IPv6 and resulted in the standardization of a new type of address: IPv6 privacy addresses (RFC3041). Privacy addresses allow a network user to generate a new unique IPv6 address periodically or for each session. Following the IPv6 leading practices design rules, a Local Area Network is normally allocated with a 64 bit prefix and a 64 bit host-ID. This allows a diversity of 2^{64} possible IPv6 addresses per LAN and will make it virtually impossible for entities on the outside to map a user to a single IPv6 address. At this moment most operating systems for client PC's have support for IPv6 rfc3041 privacy addresses and can be configured to use them.

Note that usage of privacy addresses within Enterprise networks has consequences for the operational support (troubleshooting, perimeter security, etc.) of the network as indicated in section 4 of RFC3041 (<http://www.tools.ietf.org/html/rfc3041-section-4>).

Topology Hiding

Topology hiding provides the ability for hosts to appear attached to the Enterprise network border router. This provides a mechanism to hide the internal structure of the network towards the outside world. To some IT managers, this is seen as a useful artifact from address translation done at the Enterprise edge border router. The value perception of NAT's topology hiding can be analyzed. It does indeed make it harder for an outside instance to guess the addressing and routing structure

behind an address translation gateway, however NAT does not mask the egress point of the network, and this egress point may become an easy attack vector. Most well designed Enterprise networks do not have a need for security by obscurity, as it provides a false perception of a secure infrastructure. A secure Enterprise network tends to be much less concerned about hiding the location of an internal device, in comparison with the routability from the outside to the inside and with direct communication of the hosts itself. The act of address translation or topology hiding does not provide security in itself; for example, consider a configuration with static NAT translation and all inbound ports translating to a single machine. In such a scenario, the security risk for that machine is identical to the case with no NAT device in the communication path. As result there is no specific security value in the topology hiding or address translation function. The perceived security of NAT comes from the lack of pre-established or permanent mapping state. Dynamically establishing state in response to internal requests reduces the threat of unexpected external connections to internal devices. This role with explicit management and context control is what a real firewall does best and provides a much more secure environment than an address translator .

Topology hiding is a result due to decoupling the location of the device within the local network from the IPv6 address seen on the outside world. As documented in the previous section, the incremental security this ability provides is marginal in most networks. For IPv6, the ability of topology hiding is possible, however it must be enabled on the network as a network service. The decoupling between the external address and the location of an IPv6 device can be achieved using two methods:

1. Injecting host route entries

Using host routes will mask the visibility about the location of the hosts within the local network towards external networks. With IPv6 the prefix will be a 64 bit length for host LANs or segments, and an external user may be able to guess the network topology by looking at the addresses used by the devices on the local network. Injecting host routes for the devices on a LAN segment will remove the correlation potential between the IPv6 device and its attachment to the LAN segment or interface. These IPv6 host routes can, on a Cisco IOS router, be created by static IPv6 routing entries. This method may be a solution for small networks, however, for large scale Enterprise networks, this is not a greatly scalable solution due to the large amount of IGP route entries and the increased network convergence time.

2. Using Mobile IPv6 (MIPv6)

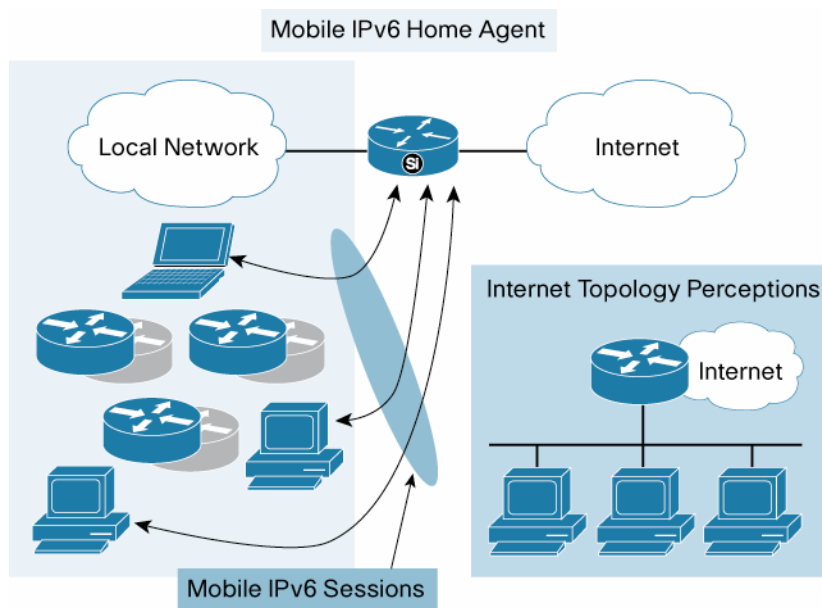
The second method to achieve topology hiding is by masking the local network addresses, by using Mobile IPv6 technology. The method exists out of using a Mobile IPv6 Home-Agent instead of NAT middleware at the edge of the network, while the local devices must run a mobile IPv6 client protocol stack (note that not all operating systems have a client Mobile IPv6 protocol stack). There is no need for the devices in the Internet to have Mobile IPv6 technology support.

On the Home-Agent (HA) there is a LAN attached which will have all the devices connected for which topology hiding is required. This LAN segment will typically be addressed with a 64 bit IPv6 prefix. This IPv6 address will be the devices IPv6 Home-Address. If there are many hosts that need to enjoy topology hiding services, then they need to be distributed amongst a set of Home-agents for scalability reasons. If any host on the external side of the Home-Agent wants to speak with any of the local devices, then the external device will need to communicate with the IPv6 Home Address of the local device.

Obviously, a typical Enterprise network does not consist out of one single massive LAN segment, but will probably be built with a hierarchical topology to increase its scalability, redundancy and performance. When using IPv6 this basic network design principle will remain unchanged. Each of these devices will be placed on their final location and will be assigned with a hierarchical IPv6 address. Within Mobile IPv6 this address is called the Care-of-

Address (CoA). This address is a real IPv6 address and might be used to communicate towards the Internet. However, if this IPv6 address is used, then the device is not using IPv6 topology hiding. If this local IPv6 device wants to start using topology hiding, then its Mobile IPv6 protocol stack needs to be enabled and configured with its IPv6 Home-Address (The IPv6 address it would have if the device would be connected to the large massive LAN segment of the Home-Agent) and the IPv6 address of its Home-Agent. These pieces of information will allow the creation of a session between the Home-Agent and the local IPv6 device. This session can be best visualized as a dynamically created IPv6 tunnel. If the local IPv6 device now wants to communicate to a device on the Internet then it will send a packet through this dynamically created tunnel to the Home-Agent. The source address of the tunneled IPv6 packet is the devices IPv6 Home-Address, while the destination address is the IPv6 address of the remote device on the Internet. Keep in mind that the tunnel source IPv6 address is the Mobile IPv6 Care-of-Address. The packet will look just like any other IPv6 packet once the Home-Agent de-capsulated the tunneled packet (it stripped the IPv6 tunnel addresses) and forwarded the packet onwards to the device on the Internet. The remote IPv6 device on the Internet has no awareness at all about the packets Care-of-Address as that address was stripped by the Mobile IPv6 Home-Agent. For the return traffic, a similar flow can be seen. The device on the Internet will send the packet towards the IPv6 Home-Address of the local IPv6 device. The Home-Agent will intercept this packet and will push it into the dynamically created tunnel towards the local IPv6 device's Care-of-Address. For the Home-Agent to do this, he needs to maintain stateful table to correlate which Home-Address matches which Care-of-Address. The Internet device does not need to have awareness and neither use Mobile IPv6.

Figure 7. Topology Hiding with IPv6



The big difference between using NAT as middleware or by using Mobile IPv6 Home-Agent is that NAT breaks the peer-to-peer, communication, while with MIPv6 that is not the case. The connection between the local device and the external device is established between the external device's IPv6 address and the home-address of the local device, and 'not' of its Care-of-Address.

Cisco IOS Software starts with Mobile IPv6 Home-Agent capabilities since version Release 12.3(14)T and Release 12.4.

Mobile IPv6 configuration information can be found at:


http://www.cisco.com/en/US/docs/ios/12_2t/ipv6/SA_mobv6_ps6350_TSD_Products_Configuration_Guide_Chapter.html

Addressing Autonomy

The usage of an IPv4 address translator gives users autonomy to select an address range for the local network. Under IPv4 address space policy restrictions, each subnet must be optimized, so one has to look periodically into the number of hosts on a segment and the subnet size allocated to the segment and rebalance. When using a NAT it is possible to almost randomly select any IPv4 address space and use it for the numbering of the local network and reduce the operational overhead. Most commonly the prescribed 'IPv4 private addresses' defined by rfc1928 are used for this purpose. This provides the user to design the local network with greater flexibility, disregarding the increasing complexity to get and maintain global unique IPv4 address space.

IPv6 provides for autonomy in local use addresses through Unique Local Addresses (ULAs). The goal of these addresses is to provide address independency from any external influences. In essence, they are similar as the IPv4 'private' addresses, but with the major difference, that these addresses are with very high probability globally unique, which provides additional benefits (ie: during the merger of two private networks). The usage of ULA addresses on the local network does not preclude usage of globally unique IPv6 addresses for communication with devices and services on external networks. IPv6 is designed from the base assumption that each interface will have multiple IPv6 addresses assigned per interface, and IPv6 will select automatically a local address for communication within the local network, and globally unique addresses for external communication.

Figure 8.



For IPv6 however ...

there is no problem of address Autonomy:

Large Address space per site or user (/48)

RFC3177 describes the allocation of IPv6 address space
 Typical site will get /48 (this provides 16 bits for subnets =
 65536 networks per site (even for your home-network))

Unique Local Addresses

RFC4193

Provides Unique private address space for internal independent usage

The operational overhead to periodically assess the number of network users and size of subnets is eliminated when using IPv6. The need to ask for more address space will become far less likely, due to the increased size of the subnets. Each IPv6 subnet will typically allow 2^{64} hosts to be connected. The current IPv6 address allocation policy at the Internet Registries prescribe for a default allocation of either a /56 or /48 to home and Enterprise networks. These allocations allow

256 or 65536, respectively IPv6 subnets, each capable of supporting up to 2^{64} unique IPv6 hosts.

Global Address Pool Conservation

Address Translation was initially specified to provide a solution for the ongoing IPv6 address depletion. This is a function that has been achieved very well. However, address translation slowed the IPv4 address depletion down, but it did not stop it. There are studies on address depletion available on the Internet which anticipates a hard stop for available IPv4 addresses around 2009.


Some reference information with respect to IPv4 address depletion:

<http://www.potaroo.net/tools/ipv4/>

<http://www.networkworld.com/community/?q=node/14969&neth=051007dailynews1&>

<http://www.tndh.net/~tony/ietf/IPv4-pool-combined-view-070314.pdf>

Figure 9.



- IPv4
 - 32 bits
 - $\approx 4,200,000,000$ possible addressable nodes
- IPv6
 - 128 bits: 4 times the size in bits
 - $\approx 3,4 \times 10^{38}$ possible addressable nodes
 - $\approx 340,282,366,920,938,463,374,607,432,768,211,456$
 - $\approx 10^{30}$ addresses per person on the planet

Conclusion: No need for Global Address Pool Conservation in IPv6 due to a legacy protocol limitation

In contrast with IPv4 there is currently a gigantic address pool available when using IPv6. IPv6 has an 128 bit address space, while IPv4 has a theoretical address space of 32 bits. This means that the available IPv6 address pool is 2^{96} times the full 32 bit IPv4 address space. This IPv6 address space will be available throughout the next decennia for all next generation services, applications and networks.

Renumbering and Multihoming

Allowing a network to be 'multihomed' and 'renumbered' are quite different functions. However, these are argued together as reasons for using NAT, because making a network multihomed is often a transitional state required as part of network renumbering, and NAT interacts with both in the same way.

Enterprise networks are typically designed from a performance, security and resiliency perspective. It is highly desirable for these networks to be connected with two or more connections to more than one Service Provider, and to allow the Enterprise the flexibility to change Service Provider at will. Unfortunately, IPv4 was not designed to facilitate this maneuver very well. When a NAT is used however, only this component needs to deal with the multihoming and renumbering issues.

Multihoming of an Enterprise network is on first glance similar between IPv4 and IPv6 and can in theory be achieved in identical configuration, however, due to the enlarged address space there is an increased concern about scalability of the Internet. This resulted into an address allocation policy which is more 'controlled' and leaves less room for 'hole poking' as done with IPv4. IPv6 was designed to allow sites and hosts to run with several simultaneous allocated prefixes, and with several simultaneous ISPs. Hosts and Cisco IOS routers and switches support the IPv6 address selection mechanism (RFC 3484), so that all devices will behave consistently when several addresses are simultaneously valid. The fundamental difficulty that IPv4 has in regard to multiple addresses therefore does not apply to IPv6. IPv6 sites can and do run today with multiple ISPs active, and the processes for adding, removing, and renumbering active prefixes.

The options for multihoming are next to the IETF efforts to re-engineer these concepts currently the same as with IPv4:

- Multi-home to a single provider
- Provider Independent (PI) addresses
- As in IPv4, poke holes in the aggregation
- Multihoming with tunnels (RFC3178)

To aid the renumbering capabilities of IPv6, there are various configurable and tunable IPv6 parameters. A first step is the capability to abstract the received IPv6 prefix from Service Providers to a –name-, known as a 'general-prefix'. This name can now be used instead of the real IPv6 prefix. When there is need for renumbering, the configuration impact per Cisco device is dramatically reduced, because only the mapping between the assigned IPv6 address and the –name- needs to be reconfigured. When using DHCP-PD this can even be automated because IOS allows dynamic mapping between the IPv6 prefix, which is delegated to the Cisco IOS edge router. In this situation, there is no need for human intervention to renumber the router. An example IOS configuration can be found in the 'simple gateway' section. While a 'general-prefix' is a good solution for some address related tasks, it is currently not a full solution for all address related policies, like for example access-lists, neither is there availability for central management within an Enterprise network. Ongoing technological innovation will however not restrict further development to make reduce the complexity of network renumbering.

```
!
ipv6 general-prefix MY-PREFIX 2001:DB8:20::/48
!
i interface FastEthernet0/1
  ipv6 address MY-PREFIX 0:0:0:1::/64 eui-64
!
```

Additionally IPv6 has a variety of supported mechanisms to influence the addressing of hosts through the Neighbor Discovery (ND) protocol. The capabilities of Cisco IOS with respect to ND are detailed at:


Implementing IPv6 Addressing and Basic Connectivity:

http://www.cisco.com/en/US/docs/ios/12_2t/ipv6/v6addres_ps6350_TSD_Products_Configuration_Guide_Chapter.html

IPv6 host renumbering can be done through DHCPv6, just like an IPv4 host can be renumbered with DHCPv4. However, when using IPv4, the host will typically drop all existing sessions when the

host receives a new IPv4 address. With IPv6, it is possible to have a smooth transition with minimal impact on the hosts being renumbered, by manipulating valid and preferred timers attached to each IPv6 address.

Figure 10.



- **Tentative:** The address is in the process of being verified as unique.
- **Preferred:** A node can send and receive unicast traffic to and from a preferred address.
- **Deprecated:** The address is still valid, but using it for new communication is discouraged.
- **Invalid:** The address can no longer send unicast traffic to or receive it from a node. An address enters this state after the valid lifetime expires.

IPv6 host renumbering is possible by manipulating via Cisco IOS, the ND parameters to allow host renumbering in a non-disruptive make-before-brake procedure. Assume an example where initially all hosts are addressed with 'PREFIX1' and all need to be renumbered to 'PREFIX2'. This procedure can be simplified in following steps:

1. Configure the new 'PREFIX2' additionally on the router

```
interface ethernet x/y
  ipv6 nd prefix PREFIX1::/64 2592000 604800
  ipv6 nd prefix PREFIX2::/64 2592000 604800
```

2. Tune prefix timers within the ND messages

```
interface ethernet x/y
  ipv6 nd prefix PREFIX1::/64 2592000 0 - -> The prefix is valid for
2592000 seconds, but is 0 seconds preferred for creation of new sessions
  ipv6 nd prefix PREFIX2::/64 2592000 604800- -> The prefix is valid for
2592000 seconds, and is 604800 seconds preferred for creation of new
sessions
```

3. Goal: The hosts will create new IP sessions with 'PREFIX2', while existing sessions remain using 'PREFIX1'
4. After some considerable time (1 hour/day/week) tune the address timing parameters again to make 'PREFIX1' invalid

```
interface ethernet x/y
  ipv6 nd prefix PREFIX1::/64 0 0 - -> The prefix is valid for
0 seconds, and is 0 seconds preferred for creation of new sessions
  ipv6 nd prefix PREFIX2::/64 2592000 604800 -> The prefix is valid for
2592000 seconds, and is 604800 seconds preferred for creation of new
sessions
```

5. Goal: Drop all sessions on all hosts which were still using 'PREFIX1'
6. Finally the 'PREFIX1' address can be removed and the hosts are renumbered

```
interface ethernet x/y  
  
ipv6 nd prefix PREFIX2::/64 2592000 604800
```

However, multihoming and renumbering remain technically challenging even with IPv6 with regards to session continuity across multihoming events or interactions with ingress filtering. With the usage of the 'IPv6 general-prefix' it becomes less challenging than with IPv4.

Example Usage Cases

Medium/Large Private Networks

The majority of private Enterprise, academic, research, or government networks fall into this category. Many of these networks have one or more exit points to the Internet. Many of these networks have one or more exit points to the Internet, and enjoy the availability of a Network Operation Center (NOC). This type of network tends to have sufficient means to obtain independent address space, however there may be several reasons why NAT is deployed:

- From the ISP perspective, there is no need to import the end-user address space to facilitate summarization
- The network design allows a larger available address space by using private RFC1918 IPv4 addresses
- Privacy for the hosts and topology of the internal network
- Reduce the overhead in changing towards a new ISP, as only the NAT device may be impacted

These motivations to use NAT can be addressed by the enhanced features of IPv6 and Cisco IOS routers and switches. Security is for IPv6 as important as for IPv4. It is highly recommended to secure the perimeter of the IPv6 network with firewall technology. Both Cisco PIX as Cisco IOS Firewall have IPv6 protocol support.

The IPv6 address allocation policies allow an ISP to assign a sub-range of its assigned IPv6 address space to end user-networks. This will, if desired, allow the ISP to summarize all customer networks based on logical, regional or functional segmentation.

The IPv6 available address space and the prescribed IPv6 prefix allocation allows for an unparalleled flexibility in creating a network address plan. There is no need for usage of private addresses for that purpose. Under current RIR policies Enterprise networks will be granted with either a /56 or a /48 network prefix. This allows either 256 or respectively 65.536 different /64 subnets. For a majority of the Enterprise networks, this will be sufficient to generate a high quality address plan.

To allow a user of the Enterprise network to mask its identity, the RFC3041 privacy extensions can be enabled and used. However, while most Enterprise networks operating software support RFC3041 functionality, it may break their ability for internal tracking of host sessions to internal resources, and current usage is not frequently seen on Enterprise networks. The RFC3041 extensions allow a random time-limited diversity to the host part of the users IPv6 address. This will make it very hard for an external entity to gain awareness of all IPv6 addresses allocated to each, and will avoid an external element that can track and collect sent and received information from a certain user with a known IPv6 address.

For some networks it is desirable to mask in addition to its users identity, the topology of the internal infrastructure. Accomplishing this goal will need:

1. Injecting static IPv6 host routes for those hosts that need decoupling between the IPv6 address and the location. The assignment of the IPv6 addresses to the hosts can be achieved through a DHCPv6 server. Cisco CNR6.2 has supports for IPv6 address assignment, while a Cisco IOS router has DHCPv6 relay agent functionality available. Next, the injected routes need to be carried through the network by an IGP. This will make these topology hidden hosts reachable for both local and external users. This technology does provide dynamic topology masking, but there is a scalability limitation, as an IGP is generally not designed to carry many thousands of IPv6 prefixes.
2. Alternatively, Mobile IPv6 could be used. This will allow hosts to appear attached to the Enterprise border router when observing the Enterprise network from an external location. An additional traffic filter needs to be implemented at the border router/firewall to filter out leaking attachment information. While Cisco has Mobile IPv6 Home-Agent support available on Cisco IOS Software, there is a consideration for a Mobile IPv6 stack support on some client operating systems.

To minimize the impact of changing between ISP's, Unique Local Addresses (ULA) can be used. These addresses provide the Enterprise network with permanent addresses, similar to IPv4 RFC1918 private addresses. The IPv6 ULA addresses can be used to address internal servers and services, and the service would remain available even during global address renumbering.

Small Office—Home Office

These types of networks are also known as SOHO (Small Office/Home Office) networks. These networks mainly exist out of one router and one network egress point. Usually, there is no Network Operation Center (NOC) and they are connected to the Internet or external network via broadband access or dial-up connectivity. In the IPv4 situation, these networks tend to often use NAT as the cheapest available solution for address management and network connectivity.

When deploying IPv6 Local Network Protection (LNP) for these types of networks, it is important to acknowledge the requirement for simplicity and auto-configuration, while not reducing the stateful filter capabilities, which Address Translation provided for IPv4. A Cisco IOS router has stateful IPv6 firewall capabilities, or when 'just' NAT-like stateful packet filtering is required, a simple reflexive access-list can be utilized, however this minimal filtering is overshadowed with the security capabilities of the Cisco IOS Firewall.

There are two possible approaches with respect to IPv6 prefix address assignment for these SOHO networks:

1. ISP uses DHCP-PD to assign dynamically an IPv6 address range to the SOHO network
2. Statically assigned IPv6 address range by the ISP

For the DHCPv6-PD solution, a dynamic address allocation approach is chosen. By means of the enhanced DHCPv6 protocol, it is possible to have the ISP push down an IPv6 prefix range automatically towards the small private network, and populate all interfaces in that small private network dynamically. This reduces the burden for administrative overhead because everything happens automatically.

For the static configuration, the mechanisms used could be the same as the medium/large Enterprises. Typically, the need for masking the topology will not be of high priority for these users, and the usage of IPv6 privacy extensions on the hosts (RFC3041) could be sufficient.

For both alternatives, the ISP has the unrestricted capability for summarization of its RIR-allocated IPv6 prefix, while the small private network administrator has all flexibility in using the received IPv6 prefix to its advantage, because it will be of sufficient size to allow all the local nodes to have a public address and full range of ports available whenever necessary.

Conclusion

When using IPv6 LNP functionality, similar functionality with respect to simple gateway functionality, address autonomy, advanced security, usage tracking, renumbering automation, end system privacy and topology hiding are all available, without a requirement or need for address translation. In addition to these aspects, IPv6 provides additional features and capabilities, due to the protocol enhancements and the lack of Address Translation middleware.

IPv6 restored the capability for universal any-to-any connectivity. The dramatic growth of the IPv6 address space enables all devices to be able to get a global IPv6 Internet address. All devices having such global IPv6 address have the potential to communicate between each other with Address translation middleware. This provides a fertile environment for technology and application innovations (peer-to-peer, instant messaging, Internet telephony, etc.).

The merging of two private IPv4 networks can cause severe complications, especially if the networks are using overlapping private rfc1918 compliant addresses. When using IPv6, this not the case anymore. IPv6 uses a concept of almost globally unique IPv6 addresses (Unique Local Addresses—RFC4193) which can be used on a site in similar fashion as the IPv4 private addresses. Due to the high probability that these local IPv6 addresses are globally unique, the complication caused by overlapping addresses is not experienced.

The revised auto-configuration capabilities will make it possible to add an almost unlimited amount of devices to a LAN with almost zero requirement for human interaction. With IPv4, this would require a stateful DHCP protocol stack on the added devices and a powerful DHCP server to keep track of stateful assigned IPv4 addresses.

Multicast services with IPv4 are severely restricted due to the NAT middleware and the limited amount of available multicast groups. IPv6 adds an almost unrestricted amount of multicast groups and introduces the concept of multicast scope.

Due to the restored aspect that both source and destination IPv6 addresses are visible in the IPv6 packet and to both communicating peers, is it possible to keep track of all flows and to use IPsec directly between the peers. An IPv4 firewall with NAT functionality can deny or allow IPv4 packets. A firewall has for IPv6, an undoubted importance and can deny flows and packets based on its security policies; however it does not interfere with the traffic being allowed, and allows end-to-end auditing or IP level identification.

IPv6 has enhanced mobility capabilities, and does not use a Mobile IP foreign-agent, unlike IPv4. Mobile IPv6 does, in addition, support direct communication between two nodes, if they both have mobile IPv6 protocol stack, while with Mobile IPv4, all traffic is forced to flow over the Mobile IP Home-Agent.



Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV
Amsterdam, The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

CCDE, CCENT, Cisco Eos, Cisco Lumin, Cisco StadiumVision, the Cisco logo, DCE, and Welcome to the Human Network are trademarks; Changing the Way We Work, Live, Play, and Learn is a service mark; and Access Registrar, Aironet, AsyncOS, Bringing the Meeting To You, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, CCVP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Collaboration Without Limitation, EtherFast, EtherSwitch, Event Center, Fast Step, Follow Me Browsing, FormShare, GigaDrive, HomeLink, Internet Quotient, IOS, iPhone, iQ Expertise, the iQ logo, iQ Net Readiness Scorecard, iQuick Study, IronPort, the IronPort logo, LightStream, Linksys, MediaTone, MeetingPlace, MGX, Networkers, Networking Academy, Network Registrar, PCNow, PIX, PowerPanels, ProConnect, ScriptShare, SenderBase, SMARTnet, Spectrum Expert, StackWise, The Fastest Way to Increase Your Internet Quotient, TransPath, WebEx, and the WebEx logo are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or Website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0804R)