Network Working Group
Request for Comments: 4864
Category: Informational

G. Van de Velde
T. Hain
R. Droms
Cisco Systems
B. Carpenter
IBM
E. Klein
Tel Aviv University

May 2007

Local Network Protection for IPv6

Status of This Memo

Copyright Notice

Abstract

Although there are many perceived benefits to Network Address Translation (NAT), its primary benefit of "amplifying" available address space is not needed in IPv6. In addition to NAT's many serious disadvantages, there is a perception that other benefits exist, such as a variety of management and security attributes that could be useful for an Internet Protocol site. IPv6 was designed with the intention of making NAT unnecessary, and this document shows how Local Network Protection (LNP) using IPv6 can provide the same or more benefits without the need for address translation.

Table of Contents

1. Introduction

There have been periodic claims that IPv6 will require a Network Address Translation (NAT), because network administrators use NAT to meet a variety of needs when using IPv4 and those needs will also have to be met when using IPv6. Although there are many perceived benefits to NAT, its primary benefit of "amplifying" available address space is not needed in IPv6. The serious disadvantages and impact on applications by ambiguous address space and Network Address Translation [1] [5] have been well documented [4] [6], so there will not be much additional discussion here. However, given its wide deployment NAT undoubtedly has some perceived benefits, though the bulk of those using it have not evaluated the technical trade-offs. Indeed, it is often claimed that some connectivity and security concerns can only be solved by using a NAT device, without any mention of the negative impacts on applications. This is amplified through the widespread sharing of vendor best practice documents and sample configurations that do not differentiate the translation function of address expansion from the state function of limiting connectivity.

This document describes the uses of a NAT device in an IPv4 environment that are regularly cited as ’solutions’ for perceived problems. It then shows how the goals of the network manager can be met in an IPv6 network without using the header modification feature of NAT. It should be noted that this document is ’informational’, as it discusses approaches that will work to accomplish the goals of the network manager. It is specifically not a Best Current Practice (BCP) that is recommending any one approach or a manual on how to configure a network.

As far as security and privacy are concerned, this document considers how to mitigate a number of threats. Some are obviously external, such as having a hacker or a worm-infected machine outside trying to penetrate and attack the local network. Some are local, such as a disgruntled employee disrupting business operations or the unintentional negligence of a user downloading some malware, which then proceeds to attack from within. Some may be inherent in the device hardware ("embedded"), such as having some firmware in a domestic appliance "call home" to its manufacturer without the user’s consent.

Another consideration discussed is 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; at the same time it defeats others, such as the ability to produce an end-to-end audit trail at network level. That said, there are artifacts of NAT devices that do provide some value.

1. The need to establish state before anything gets through from outside to inside solves one set of problems.

2. The expiration of state to stop receiving any packets when finished with a flow solves a set of problems.

3. The ability for nodes to appear to be attached at the edge of the network solves a set of problems.

4. The ability to have addresses that are not publicly routed solves yet another set (mostly changes where the state is and scale requirements for the first one).

This document describes several techniques that may be combined in an IPv6 deployment to protect the integrity of its network architecture. It will focus on the 'how to accomplish a goal' perspective, leaving most of the 'why that goal is useful' perspective for other documents. These techniques, known collectively as Local Network Protection (LNP), retain the concept of a well-defined boundary between "inside" and "outside" the private network, while allowing firewalling, topology hiding, and privacy. LNP will achieve these security goals without address translation while regaining the ability for arbitrary any-to-any connectivity.

IPv6 Local Network Protection can be summarized in the following table. It presents the marketed benefits of IPv4+NAT with a cross-reference of how those are delivered in both the IPv4 and IPv6 environments.

Goal	IPv4	IPv6
Simple Gateway as default router and address pool manager	DHCP - single address upstream DHCP - limited number of individual devices downstream see Section 2.1	DHCP-PD - arbitrary length customer prefix upstream SLAAC via RA downstream see Section 4.1
Simple Security	Filtering side effect due to lack of translation state see Section 2.2	Explicit Context Based Access Control (Reflexive ACL) see Section 4.2
Local Usage Tracking	NAT state table see Section 2.3	Address uniqueness see Section 4.3
End-System Privacy	NAT transforms device ID bits in the address see Section 2.4	Temporary use privacy addresses see Section 4.4
Topology Hiding	NAT transforms subnet bits in the address see Section 2.4	Untraceable addresses using IGP host routes /or MIPv6 tunnels see Section 4.4
Addressing Autonomy	RFC 1918 see Section 2.5	RFC 3177 & 4193 see Section 4.5
Global Address Pool Conservation	RFC 1918 << 2^48 application end points topology restricted see Section 2.6	17*10^18 subnets 3.4*10^38 addresses full port list / addr unrestricted topology see Section 4.6
Renumbering and Multihoming	Address translation at border see Section 2.7	Preferred lifetime per prefix & multiple addresses per interface see Section 4.7

This document first identifies the perceived benefits of NAT in more detail, and then shows how IPv6 LNP can provide each of them. It concludes with an IPv6 LNP case study and a gap analysis of standards work that remains to be done for an optimal LNP solution.

2. Perceived Benefits of NAT and Its Impact on IPv4

This section provides insight into the generally perceived benefits of the use of IPv4 NAT. The goal of this description is not to analyze these benefits or the accuracy of the perception (detailed discussions in [4]), but to describe the deployment requirements and set a context for the later descriptions of the IPv6 approaches for dealing with those requirements.

2.1. Simple Gateway between Internet and Private Network

A NAT device can connect a private network with addresses allocated from any part of the space (ambiguous [1]or global registered and unregistered addresses) towards the Internet, though extra effort is needed when the same range exists on both sides of the NAT. The address space of the private network can be built from globally unique addresses, from ambiguous address space, or from both simultaneously. In the simple case of private use addresses, without needing specific configuration the NAT device enables access between the client side of a distributed client-server application in the private network and the server side located in the public Internet.

Wide-scale deployments have shown that using NAT to act as a simple gateway attaching a private IPv4 network to the Internet is simple and practical for the non-technical end user. Frequently, a simple user interface or even a default configuration is sufficient for configuring both device and application access rights.

This simplicity comes at a price, as the resulting topology puts restrictions on applications. The NAT simplicity works well when the applications are limited to a client/server model with the server deployed on the public side of the NAT. For peer-to-peer, multi-party, or servers deployed on the private side of the NAT, helper technologies must also be deployed. These helper technologies are frequently complex to develop and manage, creating a hidden cost to this ’simple gateway’.

2.2. Simple Security Due to Stateful Filter Implementation

It is frequently believed that through its session-oriented operation, NAT puts in an extra barrier to keep the private network protected from outside influences. Since a NAT device typically keeps state only for individual sessions, attackers, worms, etc.,

cannot exploit this state to attack a specific host on any other port. However, in the port overload case of Network Address Port Translation (NAPT) attacking all active ports will impact a potentially wide number of hosts. This benefit may be partially real; however, experienced hackers are well aware of NAT devices and are very familiar with private address space, and they have devised methods of attack (such as trojan horses) that readily penetrate NAT boundaries. While the stateful filtering offered by NAT offers a measure of protection against a variety of straightforward network attacks, it does not protect against all attacks despite being presented as a one-size-fits-all answer.

The act of translating address bits within the header does not provide security in itself. For example, consider a configuration with static NAT 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 any connection to the public address will be delivered to the mapped target.

The perceived security of NAT comes from the lack of pre-established or permanent mapping state. This is often used as a 'better than nothing' level of protection because it doesn't require complex management to filter out unwanted traffic. Dynamically establishing state in response to internal requests reduces the threat of unexpected external connections to internal devices, and this level of protection would also be available from a basic firewall. (A basic firewall, supporting clients accessing public side servers, would improve on that level of protection by avoiding the problem of state persisting as different clients use the same private side address over time.) This role, often marketed as a firewall, is really an arbitrary artifact, while a real firewall often offers explicit and more comprehensive management controls.

In some cases, NAT operators (including domestic users) may be obliged to configure quite complex port mapping rules to allow external access to local applications such as a multi-player game or web servers. In this case, the NAT actually adds management complexity compared to the simple router discussed in Section 2.1. In situations where two or more devices need to host the same application or otherwise use the same public port, this complexity shifts from difficult to impossible.

2.3. User/Application Tracking

One usage of NAT is for the local network administrator to track user and application traffic. Although NATs create temporary state for active sessions, in general they provide limited capabilities for the

administrator of the NAT to gather information about who in the private network is requesting access to which Internet location. This is done by periodically logging the network address translation details of the private and the public addresses from the NAT device's state database.

The subsequent checking of this database is not always a simple task, especially if Port Address Translation is used. It also has an unstated assumption that the administrative instance has a mapping between a private IPv4-address and a network element or user at all times, or the administrator has a time-correlated list of the address/port mappings.

2.4. Privacy and Topology Hiding

One goal of 'topology hiding' is to prevent external entities from making a correlation between the topological location of devices on the local network. The ability of NAT to provide Internet access to a large community of users by the use of a single (or a few) globally routable IPv4 address(es) offers a simple mechanism to hide the internal topology of a network. In this scenario, the large community will be represented in the Internet by a single (or a few) IPv4 address(es).

By using NAT, a system appears to the Internet as if it originated inside the NAT box itself; i.e., the IPv4 address that appears on the Internet is only sufficient to identify the NAT so all internal nodes appear to exist at the demarcation edge. When concealed behind a NAT, it is impossible to tell from the outside which member of a family, which customer of an Internet cafe, or which employee of a company generated or received a particular packet. Thus, although NATs do nothing to provide application level privacy, they do prevent the external tracking and profiling of individual systems by means of their IP addresses, usually known as 'device profiling'.

There is a similarity with privacy based on application level proxies. When using an application level gateway for browsing the web for example, the 'privacy' of a web user can be provided by masking the true identity of the original web user towards the outside world (although the details of what is -- or is not -- logged at the NAT/proxy will be different).

Some network managers prefer to hide as much as possible of their internal network topology from outsiders as a useful precaution to mitigate scanning attacks. Mostly, this is achieved by blocking "traceroute", etc., though NAT entirely hides the internal subnet topology. Scanning is a particular concern in IPv4 networks because the subnet size is small enough that once the topology is known, it

is easy to find all the hosts, then start scanning them for vulnerable ports. Once a list of available devices has been mapped, a port-scan on these IP addresses can be performed. Scanning works by tracking which ports do not receive unreachable errors from either the firewall or host. With the list of open ports, an attacker can optimize the time needed for a successful attack by correlating it with known vulnerabilities to reduce the number of attempts. For example, FTP usually runs on port 21, and HTTP usually runs on port 80. Any vulnerable open ports could be used for access to an end-system to command it to start initiating attacks on others.

2.5. Independent Control of Addressing in a Private Network

Many private IPv4 networks make use of the address space defined in RFC 1918 to enlarge the available addressing space for their private network, and at the same time reduce their need for globally routable addresses. This type of local control of address resources allows a sufficiently large pool for a clean and hierarchical addressing structure in the local network.

Another benefit is the ability to change providers with minimal operational difficulty due to the usage of independent addresses on a majority of the network infrastructure. Changing the addresses on the public side of the NAT avoids the administrative challenge of changing every device in the network.

Section 2.7 describes some disadvantages that appear if independent networks using ambiguous addresses [1] have to be merged.

2.6. Global Address Pool Conservation

While the widespread use of IPv4+NAT has reduced the potential consumption rate, the ongoing depletion of the IPv4 address range has already taken the remaining pool of unallocated IPv4 addresses well below 20%. While mathematical models based on historical IPv4 prefix consumption periodically attempt to predict the future exhaustion date of the IPv4 address pool, a possible result of this continuous resource consumption is that the administrative overhead for acquiring globally unique IPv4 addresses will at some point increase noticeably due to tightening allocation policies.

In response to the increasing administrative overhead, many Internet Service Providers (ISPs) have already resorted to the ambiguous addresses defined in RFC 1918 behind a NAT for the various services they provide as well as connections for their end customers. This happens even though the private use address space is strictly limited in size. Some deployments have already outgrown that space and have begun cascading NAT to continue expanding, though this practice

eventually breaks down over routing ambiguity. Additionally, while we are unlikely to know the full extent of the practice (because it is hidden behind a NAT), service providers have been known to announce previously unallocated public space to their customers (to avoid the problems associated with the same address space appearing on both sides), only to find that once that space was formally allocated and being publicly announced, their customers couldn't reach the registered networks.

The number of and types of applications that can be deployed by these ISPs and their customers are restricted by the ability to overload the port range on the public side of the most public NAT in the path. The limit of this approach is something substantially less than 2^48 possible active *application* endpoints (approximately [2^32 minus 2^29] * [2* 2^16 minus well-known port space]), as distinct from addressable devices each with its own application endpoint range. Those who advocate layering of NAT frequently forget to mention that there are topology restrictions placed on the applications. Forced into this limiting situation, such customers can rightly claim that despite the optimistic predictions of mathematical models, the global pool of IPv4 addresses is effectively already exhausted.

2.7. Multihoming and Renumbering with NAT

Allowing a network to be multihomed and renumbering a network 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.

For enterprise networks, it is highly desirable to provide resiliency and load-balancing to be connected to more than one Internet Service Provider (ISP) and to be able to change ISPs at will. This means that a site must be able to operate under more than one Classless Inter-Domain Routing (CIDR) prefix [18] and/or readily change its CIDR prefix. Unfortunately, IPv4 was not designed to facilitate either of these maneuvers. However, if a site is connected to its ISPs via NAT boxes, only those boxes need to deal with multihoming and renumbering issues.

Similarly, if two enterprise IPv4 networks need to be merged and RFC 1918 addresses are used, there is a high probability of address overlaps. In those situations, it may well be that installing a NAT box between them will avoid the need to renumber one or both. For any enterprise, this can be a short-term financial saving and allows more time to renumber the network components. The long-term solution is a single network without usage of NAT to avoid the ongoing operational complexity of overlapping addresses.

The addition of an extra NAT as a solution may be sufficient for some networks; however, when the merging networks were already using address translation it will create major problems due to administrative difficulties of overlapping address spaces in the merged networks.

3. Description of the IPv6 Tools

This section describes several features that can be used as part of the LNP solution to replace the protection features associated with IPv4 NAT.

The reader must clearly distinguish between features of IPv6 that were fully defined when this document was drafted and those that were potential features that still required more work to define them. The latter are summarized later in the 'Gap Analysis' section of this document. However, we do not distinguish in this document between fully defined features of IPv6 and those that were already widely implemented at the time of writing.

3.1. Privacy Addresses (RFC 3041)

There are situations where it is desirable to prevent device profiling, for example, by web sites that are accessed from the device as it moves around the Internet. IPv6 privacy addresses were defined to provide that capability. IPv6 addresses consist of a routing prefix, a subnet-id (SID) part, and an interface identifier (IID) part. As originally defined, IPv6 stateless address auto-configuration (SLAAC) will typically embed the IEEE Link Identifier of the interface as the IID part, though this practice facilitates tracking and profiling of a device through the consistent IID. RFC 3041 [7] describes an extension to SLAAC to enhance device privacy. Use of the privacy address extension causes nodes to generate global-scope addresses from interface identifiers that change over time, consistent with system administrator policy. Changing the interface identifier (thus the global-scope addresses generated from it) over time makes it more difficult for eavesdroppers and other information collectors to identify when addresses used in different transactions actually correspond to the same node. A relatively short valid lifetime for the privacy address also has the effect of reducing the attack profile of a device, as it is not directly attackable once it stops answering at the temporary use address.

While the primary implementation and source of randomized RFC 3041 addresses are expected to be from end-systems running stateless auto-configuration, there is nothing that prevents a Dynamic Host Configuration Protocol (DHCP) server from running the RFC 3041 algorithm for any new IEEE identifier it hears in a request, then

remembering that for future queries. This would allow using them in DNS for registered services since the assumption of a DHCP server-based deployment would be a persistent value that minimizes DNS churn. A DHCP-based deployment would also allow for local policy to periodically change the entire collection of end-device addresses while maintaining some degree of central knowledge and control over which addresses should be in use at any point in time.

Randomizing the IID, as defined in RFC 3041, is effectively a sparse allocation technique that only precludes tracking of the lower 64 bits of the IPv6 address. Masking of the subnet ID will require additional approaches as discussed below in Section 3.4. Additional considerations are discussed in [19].

3.2. Unique Local Addresses

Achieving the goal of autonomy, that many perceive as a value of NAT, is required for local network and application services stability during periods of intermittent connectivity or moving between one or more providers. Such autonomy in a single routing prefix environment would lead to massive expansion of the global routing tables (as seen in IPv4), so IPv6 provides for simultaneous use of multiple prefixes. The Unique Local Address (ULA) prefix [17] has been set aside for use in local communications. The ULA prefix for any network is routable over a locally defined collection of routers. These prefixes are not intended to be routed on the public global Internet as large-scale inter-domain distribution of routes for ULA prefixes would have a negative impact on global route aggregation.

ULAs have the following characteristics:

- For all practical purposes, a globally unique prefix
  - allows networks to be combined or privately interconnected without creating address conflicts or requiring renumbering of interfaces using these prefixes, and
  - if accidentally leaked outside of a network via routing or DNS, is highly unlikely that there will be a conflict with any other addresses.
- They are ISP independent and can be used for communications inside of a network without having any permanent or only intermittent Internet connectivity.

- They have a well-known prefix to allow for easy filtering at network boundaries preventing leakage of routes and packets that should remain local.

- In practice, applications may treat these addresses like global-scope addresses, but address selection algorithms may need to distinguish between ULAs and ordinary global-scope unicast addresses to ensure stability. The policy table defined in [11] is one way to bias this selection, by giving higher preference to FC00::/7 over 2001::/3. Mixing the two kinds of addresses may lead to undeliverable packets during times of instability, but that mixing is not likely to happen when the rules of RFC 3484 are followed.

- ULAs have no intrinsic security properties. However, they have the useful property that their routing scope is limited by default within an administrative boundary. Their usage is suggested at several points in this document, as a matter of administrative convenience.

3.3. DHCPv6 Prefix Delegation

One of the functions of a simple gateway is managing the local use address range. The Prefix Delegation (DHCP-PD) options [12] provide a mechanism for automated delegation of IPv6 prefixes using the DHCP [10]. This mechanism (DHCP-PD) is intended for delegating a long-lived prefix from a delegating router (possibly incorporating a DHCPv6 server) to a requesting router, possibly across an administrative boundary, where the delegating router does not require knowledge about the topology of the links in the network to which the prefixes will be assigned.

3.4. Untraceable IPv6 Addresses

The main goal of untraceable IPv6 addresses is to create an apparently amorphous network infrastructure, as seen from external networks, to protect the local infrastructure from malicious outside influences and from mapping of any correlation between the network activities of multiple devices from external networks. When using untraceable IPv6 addresses, it could be that two apparently sequential addresses are allocated to devices on very different parts of the local network instead of belonging to devices adjacent to each other on the same subnet.

Since IPv6 addresses will not be in short supply even within a single /64 (or shorter) prefix, it is possible to generate them effectively at random when untraceability is required. They will be globally routable IPv6 addresses under the site’s prefix, which can be

randomly and independently assigned to IPv6 devices. The random assignment is intended to mislead the outside world about the structure of the local network. In particular, the subnet structure may be invisible in the address. Thus, a flat routing mechanism will be needed within the site. The local routers need to maintain a correlation between the topological location of the device and the untraceable IPv6 address. For smaller deployments, this correlation could be done by generating IPv6 host route entries, or for larger ones by utilizing an indirection device such as a Mobile IPv6 Home Agent. Additional details are in Section 4.7.

4. Using IPv6 Technology to Provide the Market Perceived Benefits of NAT

The facilities in IPv6 described in Section 3 can be used to provide the protection perceived to be associated with IPv4 NAT. This section gives some examples of how IPv6 can be used securely.

4.1. Simple Gateway between Internet and Internal Network

As a simple gateway, the device manages both packet routing and local address management. A basic IPv6 router should have a default configuration to advertise inside the site a locally generated random ULA prefix, independently from the state of any external connectivity. This would allow local nodes in a topology more complex than a single link to communicate amongst themselves independent of the state of a global connection. If the network happened to concatenate with another local network, the randomness in ULA creation is highly unlikely to result in address collisions.

With external connectivity, the simple gateway should use DHCP-PD to acquire a routing prefix from the service provider for use when connecting to the global Internet. End-system connections involving other nodes on the global Internet that follow the policy table in RFC 3484 will always use the global IPv6 addresses derived from this prefix delegation. It should be noted that the address selection policy table should be configured to prefer the ULA prefix range over the DHCP-PD prefix range when the goal is to keep local communications stable during periods of transient external connectivity.

In the very simple case, there is no explicit routing protocol on either side of the gateway, and a single default route is used internally pointing out to the global Internet. A slightly more complex case might involve local internal routing protocols, but with the entire local network sharing a common global prefix there would

still not be a need for an external routing protocol as the service provider could install a route for the prefix delegated via DHCP-PD pointing toward the connecting link.

4.2. IPv6 and Simple Security

The vulnerability of an IPv6 host directly connected to the Internet is similar to that of an IPv4 host. The use of firewalls and Intrusion Detection Systems (IDSs) is recommended for those that want boundary protection in addition to host defenses. A proxy may be used for certain applications, but with the caveat that the end-to-end transparency is broken. However, with IPv6, the following protections are available without the use of NAT while maintaining end-to-end reachability:

1. Short lifetimes on privacy extension suffixes reduce the attack profile since the node will not respond to the address once its lifetime becomes invalid.

2. IP security (IPsec) is often cited as the reason for improved security because it is a mandatory service for IPv6 implementations. Broader availability does not by itself improve security because its use is still regulated by the availability of a key infrastructure. IPsec functions to authenticate the correspondent, prevent session hijacking, prevent content tampering, and optionally mask the packet contents. While IPsec is commonly available in some IPv4 implementations and with extensions can support NAT traversals, NAT support has limitations and does not work in all situations. The use of IPsec with NATs requires an additional UDP encapsulation and keepalive overhead [13]. In the IPv4/NAT environment, the usage of IPsec has been largely limited to edge-to-edge Virtual Private Network (VPN) deployments. The potential for end-to-end IPsec use is significantly enhanced when NAT is removed from the network, as connections can be initiated from either end. It should be noted that encrypted IPsec traffic will bypass content-aware firewalls, which is presumed to be acceptable for parties with whom the site has established a security association.

3. The size of the address space of a typical subnet (64 bits of IID) will make a complete subnet ping sweep usually significantly harder and more expensive than for IPv4 [20]. Reducing the security threat of port scans on identified nodes requires sparse distribution within the subnet to minimize the probability of scans finding adjacent nodes. This scanning protection will be nullified if IIDs are configured in any structured groupings within the IID space. Provided that IIDs are essentially randomly distributed across the available space, address

> scanning-based attacks will effectively fail. This protection exists if the attacker has no direct access to the specific subnet and therefore is trying to scan it remotely. If an attacker has local access, then he could use Neighbor Discovery (ND) [3] and ping6 to the link-scope multicast ff02::1 to detect the IEEE-based address of local neighbors, then apply the global prefix to those to simplify its search (of course, a locally connected attacker has many scanning options with IPv4 as well).

Assuming the network administrator is aware of [20] the increased size of the IPv6 address will make topology probing much harder, and almost impossible for IPv6 devices. The intention of topology probing is to identify a selection of the available hosts inside an enterprise. This mostly starts with a ping sweep. Since the IPv6 subnets are 64 bits worth of address space, this means that an attacker has to simply send out an unrealistic number of pings to map the network, and virus/worm propagation will be thwarted in the process. At full-rate full-duplex 40 Gbps (400 times the typical 100 Mbps LAN, and 13,000 times the typical DSL/cable access link), it takes over 5,000 years to scan the entirety of a single 64-bit subnet.

IPv4 NAT was not developed as a security mechanism. Despite marketing messages to the contrary, it is not a security mechanism, and hence it will offer some security holes while many people assume their network is secure due to the usage of NAT. IPv6 security best practices will avoid this kind of illusory security, but can only address the same threats if correctly configured firewalls and IDSs are used at the perimeter.

> It must be noted that even a firewall doesn’t fully secure a network. Many attacks come from inside or are at a layer higher than the firewall can protect against. In the final analysis, every system has to be responsible for its own security, and every process running on a system has to be robust in the face of challenges like stack overflows, etc. What a firewall does is prevent a network administration from having to carry unauthorized traffic, and in so doing reduce the probability of certain kinds of attacks across the protected boundary.

To implement simple security for IPv6 in, for example, a DSL or cable modem-connected home network, the broadband gateway/router should be equipped with stateful firewall capabilities. These should provide a default configuration where incoming traffic is limited to return traffic resulting from outgoing packets (sometimes known as reflective session state). There should also be an easy interface that allows users to create inbound ’pinholes’ for specific purposes such as online gaming.

Administrators and the designers of configuration interfaces for simple IPv6 firewalls need to provide a means of documenting the security caveats that arise from a given set of configuration rules so that users (who are normally oblivious to such things) can be made aware of the risks. As rules are improved iteratively, the goal will be to make use of the IPv6 Internet more secure without increasing the perceived complexity for users who just want to accomplish a task.

4.3. User/Application Tracking

IPv6 enables the collection of information about data flows. Because all addresses used for Internet and intra-/inter-site communication are unique, it is possible for an enterprise or ISP to get very detailed information on any communication exchange between two or more devices. Unless privacy addresses [7] are in use, this enhances the capability of data-flow tracking for security audits compared with IPv4 NAT, because in IPv6 a flow between a sender and receiver will always be uniquely identified due to the unique IPv6 source and destination addresses.

At the same time, this tracking is per address. In environments where the goal is tracking back to the user, additional external information will be necessary correlating a user with an address. In the case of short-lifetime privacy address usage, this external information will need to be based on more stable information such as the layer 2 media address.

4.4. Privacy and Topology Hiding Using IPv6

Partial host privacy is achieved in IPv6 using RFC 3041 pseudo-random privacy addresses [7] which are generated as required, so that a session can use an address that is valid only for a limited time. This only allows such a session to be traced back to the subnet that originates it, but not immediately to the actual host, where IPv4 NAT is only traceable to the most public NAT interface.

Due to the large IPv6 address space available, there is plenty of freedom to randomize subnet allocations. By doing this, it is possible to reduce the correlation between a subnet and its location. When doing both subnet and IID randomization, a casual snooper won’t be able to deduce much about the network’s topology. The obtaining of a single address will tell the snooper very little about other addresses. This is different from IPv4 where address space limitations cause this not to be true. In most usage cases, this concept should be sufficient for address privacy and topology hiding, with the cost being a more complex internal routing configuration.

As discussed in Section 3.1, there are multiple parts to the IPv6 address, and different techniques to manage privacy for each which may be combined to protect the entire address. In the case where a network administrator wishes to fully isolate the internal IPv6 topology, and the majority of its internal use addresses, one option is to run all internal traffic using Unique Local Addresses (ULAs). By definition, this prefix block is not to be advertised in the public routing system, so without a routing path external traffic will never reach the site. For the set of hosts that do in fact need to interact externally, by using multiple IPv6 prefixes (ULAs and one or more global addresses) all of the internal nodes that do not need external connectivity, and the internally used addresses of those that do, will be masked from the outside. The policy table defined in [11] provides a mechanism to bias the selection process when multiple prefixes are in use such that the ULA would be preferred when the correspondent is also local.

There are other scenarios for the extreme situation when a network manager also wishes to fully conceal the internal IPv6 topology. In these cases, the goal in replacing the IPv4 NAT approach is to make all of the topology hidden nodes appear from the outside to logically exist at the edge of the network, just as they would when behind a NAT. This figure shows the relationship between the logical subnets and the topology masking router discussed in the bullet points that follow.

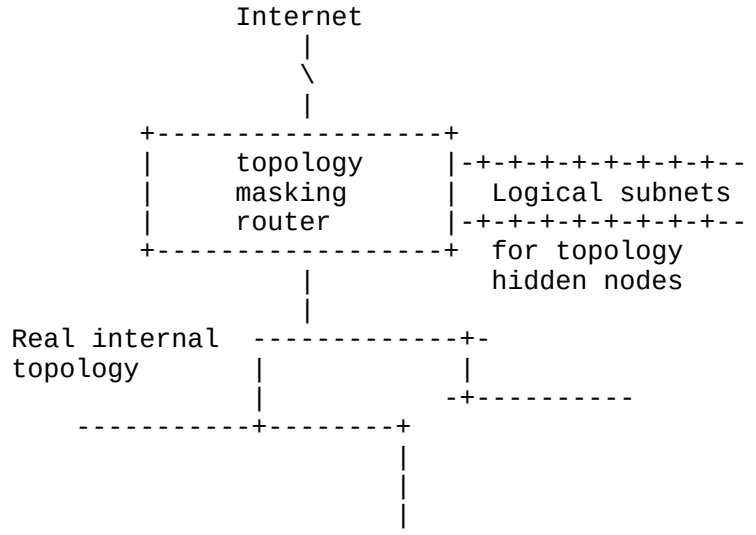

- One approach uses explicit host routes in the Interior Gateway Protocol (IGP) to remove the external correlation between physical topology attachment point and end-to-end IPv6 address. In the figure above the hosts would be allocated prefixes from one or more logical subnets, and would inject host routes into the IGP to internally identify their real attachment point. This solution does however show severe scalability issues and requires hosts to securely participate in the IGP, as well as have the firewall block all external to internal traceroutes for the logical subnet. The specific limitations are dependent on the IGP protocol, the physical topology, and the stability of the system. In any case, the approach should be limited to uses with substantially fewer than the maximum number of routes that the IGP can support (generally between 5,000 and 50,000 total entries including subnet routes). Hosts should also listen to the IGP for duplicate use before finalizing an interface address assignment as the duplicate address detection will only check for use on the attached segment, not the logical subnet.

- Another technical approach to fully hide the internal topology is use of a tunneling mechanism. Mobile IPv6 without route optimization is one approach for using an automated tunnel, as it always starts in tunnel mode via the Home Agent (HA). In this deployment model, the application perceived addresses of the nodes are routed via the edge HA acting as the topology masking router (above). This indirection method truly masks the internal topology, as from outside the local network all nodes with global access appear to share the prefix of one or more logical subnets attached to the HA rather than their real attachment point. Note that in this usage context, the HA is replacing the NAT function at the edge of the network, so concerns about additional latency for routing through a tunnel to the HA do not apply because it is effectively on the same path that the NAT traffic would have taken. Duplicate address detection is handled as a normal process of the HA binding update. While turning off all binding updates with the correspondent node would appear to be necessary to prevent leakage of topology information, that approach would also force all internal traffic using the home address to route via the HA tunnel, which may be undesirable. A more efficient method would be to allow internal route optimizations while dropping outbound binding update messages at the firewall. Another approach for the internal traffic would be to use the policy table of RFC 3484 to bias a ULA prefix as preferred internally, leaving the logical subnet Home Address external for use. The downside to a Mobile IPv6-based solution is that it requires a Home Agent in the network and the configuration of a security association with the HA for each hidden node, and it consumes some amount of bandwidth for tunnel overhead.

- Another method (where the layer 2 topology allows) uses a virtual LAN approach to logically attach the devices to one or more subnets on the edge router. This approach leads the end nodes to believe they actually share a common segment. The downside of this approach is that all internal traffic would be directed over suboptimal paths via the edge router, as well as the complexity of managing a distributed logical LAN.

One issue to be aware of is that subnet scope multicast will not work for the logical hidden subnets, except in the VLAN case. While a limited scope multicast to a collection of nodes that are arbitrarily scattered makes no technical sense, care should be exercised to avoid deploying applications that expect limited scope multicast in conjunction with topology hiding.

Another issue that this document will not define is the mechanism for a topology hidden node to learn its logical subnet. While manual configuration would clearly be sufficient, DHCP could be used for address assignment, with the recipient node discovering it is in a hidden mode when the attached subnet prefix doesn't match the one assigned.

4.5. Independent Control of Addressing in a Private Network

IPv6 provides for autonomy in local use addresses through ULAs. At the same time, IPv6 simplifies simultaneous use of multiple addresses per interface so that an IPv6 NAT is not required between the ULA and the public Internet because nodes that need access to the public Internet will have a global use address as well. When using IPv6, the need to ask for more address space will become far less likely due to the increased size of the subnets, along with an allocation policy that recognizes that table fragmentation is also an important consideration. While global IPv6 allocation policy is managed through the Regional Internet Registries (RIRs), it is expected that they will continue with derivatives of [8] for the foreseeable future so the number of subnet prefixes available to an organization should not be a limitation that would create an artificial demand for NAT.

Ongoing subnet address maintenance may become simpler when IPv6 technology is utilized. 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. For example, an enterprise today may have a mix of IPv4 /28 - /23 size subnets, and may shrink/grow these as its network user base changes. For IPv6, all subnets have /64 prefixes, which will reduce the operational and configuration overhead.

4.6. Global Address Pool Conservation

IPv6 provides sufficient space to completely avoid the need for overlapping address space. Since allocations in IPv6 are based on subnets rather than hosts, a reasonable way to look at the pool is that there are about 17*10^18 unique subnet values where sparse allocation practice within those provides for new opportunities such as SEcure Neighbor Discovery (SEND) [15]. As previously discussed, the serious disadvantages of ambiguous address space have been well documented, and with sufficient space there is no need to continue the increasingly aggressive conservation practices that are necessary with IPv4. While IPv6 allocation policies and ISP business practice will continue to evolve, the recommendations in RFC 3177 are based on the technical potential of the vast IPv6 address space. That document demonstrates that there is no resource limitation that will require the adoption of the IPv4 workaround of ambiguous space behind a NAT. As an example of the direct contrast, many expansion-oriented IPv6 deployment scenarios result in multiple IPv6 addresses per device, as opposed to the constriction of IPv4 scenarios where multiple devices are forced to share a scarce global address through a NAT.

4.7. Multihoming and Renumbering

IPv6 was designed to allow sites and hosts to run with several simultaneous CIDR-allocated prefixes, and thus with several simultaneous ISPs. An address selection mechanism [11] is specified so that hosts 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 at a site have been documented in [16] and [21]. However, multihoming and renumbering remain technically challenging even with IPv6 with regards to session continuity across multihoming events or interactions with ingress filtering (see the Gap Analysis below).

The IPv6 address space allocated by the ISP will be dependent upon the connecting service provider. This will likely result in a renumbering effort when the network changes between service providers. When changing ISPs or ISPs readjust their addressing pool, DHCP-PD [12] can be used as an almost zero-touch external mechanism for prefix change in conjunction with a ULA prefix for internal connection stability. With appropriate management of the lifetime values and overlap of the external prefixes, a smooth make-before-break transition is possible as existing communications will continue on the old prefix as long as it remains valid, while any new communications will use the new prefix.

5. Case Studies

In presenting these case studies, we have chosen to consider categories of networks divided first according to their function either as carrier/ISP networks or end user (such as enterprise) networks with the latter category broken down according to the number of connected end hosts. For each category of networks, we can use IPv6 Local Network Protection to achieve a secure and flexible infrastructure, which provides an enhanced network functionality in comparison with the usage of address translation.

- Medium/Large Private Networks (typically >10 connections)
- Small Private Networks (typically 1 to 10 connections)
- Single User Connection (typically 1 connection)
- ISP/Carrier Customer Networks

5.1. 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. Though these organizations have sufficient resources to acquire addressing independence when using IPv4, there are several reasons why they might choose to use NAT in such a network. For the ISP, there is no need to import the IPv4 address range from the remote end-customer, which facilitates IPv4 route summarization. The customer can use a larger IPv4 address range (probably with less administrative overhead) by the use of RFC 1918 and NAT. The customer also reduces the overhead in changing to a new ISP, because the addresses assigned to devices behind the NAT do not need to be changed when the customer is assigned a different address by a new ISP. By using address translation in IPv4, one avoids the expensive process of network renumbering. Finally, the customer can provide privacy for its hosts and the topology of its internal network if the internal addresses are mapped through NAT.

It is expected that there will be enough IPv6 addresses available for all networks and appliances for the foreseeable future. The basic IPv6 address range an ISP allocates for a private network is large enough (currently /48) for most of the medium and large enterprises, while for the very large private enterprise networks address ranges can be concatenated. The goal of this assignment mechanism is to decrease the total amount of entries in the public Internet routing table. A single /48 allocation provides an enterprise network with 65,536 different /64 subnet prefixes.

To mask the identity of a user on a network of this type, the usage of IPv6 privacy extensions may be advised. This technique is useful when an external element wants to track and collect all information sent and received by a certain host with a known IPv6 address. Privacy extensions add a random time-limited factor to the host part of an IPv6 address and will make it very hard for an external element to keep correlating the IPv6 address to a specific host on the inside network. The usage of IPv6 privacy extensions does not mask the internal network structure of an enterprise network.

When there is a need to mask the internal structure towards the external IPv6 Internet, then some form of 'untraceable' addresses may be used. These addresses will appear to exist at the external edge of the network, and may be assigned to those hosts for which topology masking is required or that want to reach the IPv6 Internet or other external networks. The technology to assign these addresses to the hosts could be based on DHCPv6 or static configuration. To complement the 'Untraceable' addresses, it is necessary to have at least awareness of the IPv6 address location when routing an IPv6 packet through the internal network. This could be achieved by 'host based route-injection' in the local network infrastructure. This route-injection could be done based on /128 host-routes to each device that wants to connect to the Internet using an untraceable address. This will provide the most dynamic masking, but will have a scalability limitation, as an IGP is typically not designed to carry many thousands of IPv6 prefixes. A large enterprise may have thousands of hosts willing to connect to the Internet.

An alternative for larger deployments is to leverage the tunneling aspect of MIPv6 even for non-mobile devices. With the logical subnet being allocated as attached to the edge Home Agent, the real attachment and internal topology are masked from the outside. Dropping outbound binding updates at the firewall is also necessary to avoid leaking the attachment information.

Less flexible masking could be to have time-based IPv6 prefixes per link or subnet. This may reduce the amount of route entries in the IGP by a significant factor, but has as a trade-off that masking is time and subnet based, which will complicate auditing systems. The dynamic allocation of 'Untraceable' addresses can also limit the IPv6 access between local and external hosts to those local hosts being authorized for this capability.

The use of permanent ULA addresses on a site provides the benefit that even if an enterprise changes its ISP, the renumbering will only affect those devices that have a wish to connect beyond the site. Internal servers and services would not change their allocated IPv6 ULA address, and the service would remain available even during global address renumbering.

5.2. Small Private Networks

Also known as SOHO (Small Office/Home Office) networks, this category describes those networks that have few routers in the topology and usually have a single network egress point. Typically, these networks:

- are connected via either a dial-up connection or broadband access,
- don't have dedicated Network Operation Center (NOC), and
- today, typically use NAT as the cheapest available solution for connectivity and address management

In most cases, the received global IPv4 prefix is not fixed over time and is too long (very often a /32 giving just a single address) to provide every node in the private network with a unique, globally usable address. Fixing either of those issues typically adds an administrative overhead for address management to the user. This category may even be limited to receiving ambiguous IPv4 addresses from the service provider based on RFC 1918. An ISP will typically pass along the higher administration cost attached to larger address blocks, or IPv4 prefixes that are static over time, due to the larger public address pool each of those requires.

As a direct response to explicit charges per public address, most of this category has deployed NAPT (port demultiplexing NAT) to minimize the number of addresses in use. Unfortunately, this also limits the Internet capability of the equipment to being mainly a receiver of Internet data (client), and it makes it quite hard for the equipment to become a worldwide Internet server (HTTP, FTP, etc.) due to the stateful operation of the NAT equipment. Even when there is sufficient technical knowledge to manage the NAT to enable external access to a server, only one server can be mapped per protocol/port number per address, and then only when the address from the ISP is publicly routed. When there is an upstream NAT providing private address space to the ISP side of the private NAT, additional negotiation with the ISP will be necessary to provide an inbound mapping, if that is even possible.

When deploying IPv6 LNP in this environment, there are two approaches possible with respect to IPv6 addressing.

- o DHCPv6 Prefix-Delegation (PD)
- o ISP provides a static IPv6 address range

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 for 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 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.

While a full prefix is expected to be the primary deployment model, there may be cases where the ISP provides a single IPv6 address for use on a single piece of equipment (PC, PDA, etc.). This is expected to be rare, though, because in the IPv6 world the assumption is that there is an unrestricted availability of a large amount of globally routable and unique address space. If scarcity was the motivation with IPv4 to provide RFC 1918 addresses, in this environment the ISP will not be motivated to allocate private addresses to the single user connection because there are enough global addresses available at essentially the same cost. Also, it will be likely that the single device wants to mask its identity to the called party or its attack profile over a shorter time than the life of the ISP attachment, so it will need to enable IPv6 privacy extensions. In turn, this leads to the need for a minimum allocation of a /64 prefix rather than a single address.

5.3. Single User Connection

This group identifies the users that are connected via a single IPv4 address and use a single piece of equipment (PC, PDA, etc.). This user may get an ambiguous IPv4 address (frequently imposed by the ISP) from the service provider that is based on RFC 1918. If

ambiguous addressing is utilized, the service provider will execute NAT on the allocated IPv4 address for global Internet connectivity. This also limits the Internet capability of the equipment to being mainly a receiver of Internet data, and it makes it quite hard for the equipment to become a worldwide Internet server (HTTP, FTP, etc.) due to the stateful operation of the NAT equipment.

When using IPv6 LNP, this group will identify the users that are connected via a single IPv6 address and use a single piece of equipment (PC, PDA, etc.).

In the IPv6 world, the assumption is that there is unrestricted availability of a large amount of globally routable and unique IPv6 addresses. The ISP will not be motivated to allocate private addresses to the single user connection because he has enough global addresses available, if scarcity was the motivation with IPv4 to provide RFC 1918 addresses. If the single user wants to mask his identity, he may choose to enable IPv6 privacy extensions.

5.4. ISP/Carrier Customer Networks

This group refers to the actual service providers that are providing the IP access and transport services. They tend to have three separate IP domains that they support:

- For the first, they fall into the medium/large private networks category (above) for their own internal networks, LANs, etc.
- The second is the Operations address domain, which addresses their backbone and access switches, and other hardware. This address domain is separate from the other address domains for engineering reasons as well as simplicity in managing the security of the backbone.
- The third is the IP addresses (single or blocks) that they assign to customers. These can be registered addresses (usually given to category 5.1 and 5.2 and sometimes 5.3) or can be from a pool of RFC 1918 addresses used with IPv4 NAT for single user connections. Therefore they can actually have two different NAT domains that are not connected (internal LAN and single user customers).

When IPv6 LNP is utilized in these three domains, then for the first category it will be possible to use the same solutions as described in Section 5.1. The second domain of the ISP/carrier is the Operations network. This environment tends to be a closed environment, and consequently communication can be done based on ULAs. However, in this environment, stable IPv6 Provider Independent addresses can be used. This would give a solid and scalable

configuration with respect to a local IPv6 address plan. By the usage of proper network edge filters, outside access to the closed environment can be avoided. The third is the IPv6 addresses that ISP/carrier network assign to customers. These will typically be assigned with prefix lengths terminating on nibble boundaries to be consistent with the DNS PTR records. As scarcity of IPv6 addresses is not a concern, it will be possible for the ISP to provide globally routable IPv6 prefixes without a requirement for address translation. An ISP may for commercial reasons still decide to restrict the capabilities of the end users by other means like traffic and/or route filtering, etc.

If the carrier network is a mobile provider, then IPv6 is encouraged in comparison with the combination of IPv4+NAT for Third Generation Partnership Project (3GPP)-attached devices. In Section 2.3 of RFC 3314, 'Recommendations for IPv6 in 3GPP Standards' [9], it is found that the IPv6 WG recommends that one or more /64 prefixes should be assigned to each primary Protocol Data Packet (PDP) context. This will allow sufficient address space for a 3GPP-attached node to allocate privacy addresses and/or route to a multi-link subnet, and it will discourage the use of NAT within 3GPP-attached devices.

6. IPv6 Gap Analysis

Like IPv4 and any major standards effort, IPv6 standardization work continues as deployments are ongoing. This section discusses several topics for which additional standardization, or documentation of best practice, is required to fully realize the benefits or provide optimizations when deploying LNP. From a standardization perspective, there is no obstacle to immediate deployment of the LNP approach in many scenarios, though product implementations may lag behind the standardization efforts. That said, the list below identifies additional work that should be undertaken to cover the missing scenarios.

6.1. Simple Security

Firewall traversal by dynamic pinhole management requires further study. Several partial solutions exist including Interactive Connectivity Establishment (ICE) [23], and Universal Plug and Play (UPNP) [24]. Alternative approaches are looking to define service provider mediated pinhole management, where things like voice call signaling could dynamically establish pinholes based on predefined authentication rules. The basic security provided by a stateful firewall will require some degree of default configuration and automation to mask the technical complexity from a consumer who merely wants a secure environment with working applications. There is no reason a stateful IPv6 firewall product cannot be shipped with

default protection that is equal to or better than that offered by today's IPv4/NAT products.

6.2. Subnet Topology Masking

There really is no functional standards gap here as a centrally assigned pool of addresses in combination with host routes in the IGP is an effective way to mask topology for smaller deployments. If necessary, a best practice document could be developed describing the interaction between DHCP and various IGPs that would in effect define Untraceable Addresses.

As an alternative for larger deployments, there is no gap in the HA tunneling approach when firewalls are configured to block outbound binding update messages. A border Home Agent using internal tunneling to the logical mobile (potentially rack mounted) node can completely mask all internal topology, while avoiding the strain from a large number of host routes in the IGP. Some optimization work could be done in Mobile IP to define a policy message where a mobile node would learn from the Home Agent that it should not try to inform its correspondent about route optimization and thereby expose its real location. This optimization, which reduces the load on the firewall, would result in less optimal internal traffic routing as that would also transit the HA unless ULAs were used internally. Trade-offs for this optimization work should be investigated in the IETF.

6.3. Minimal Traceability of Privacy Addresses

Privacy addresses [7] may certainly be used to limit the traceability of external traffic flows back to specific hosts, but lacking a topology masking component above they would still reveal the subnet address bits. For complete privacy, a best practice document describing the combination of privacy addresses and topology masking may be required. This work remains to be done and should be pursued by the IETF.

6.4. Site Multihoming

This complex problem has never been completely solved for IPv4, which is exactly why NAT has been used as a partial solution. For IPv6, after several years of work, the IETF has converged on an architectural approach intended with service restoration as initial aim [22]. When this document was drafted, the IETF was actively defining the details of this approach to the multihoming problem. The approach appears to be most suitable for small and medium sites, though it will conflict with existing firewall state procedures. At this time, there are also active discussions in the address

registries investigating the possibility of assigning provider-independent address space. Their challenge is finding a reasonable metric for limiting the number of organizations that would qualify for a global routing entry. Additional work appears to be necessary to satisfy the entire range of requirements.

7. Security Considerations

While issues that are potentially security related are discussed throughout the document, the approaches herein do not introduce any new security concerns. IPv4 NAT has been widely sold as a security tool, and suppliers have been implementing address translation functionality in their firewalls, though the true impact of NATs on security has been previously documented in [2] and [4].

This document defines IPv6 approaches that collectively achieve the goals of the network manager without the negative impact on applications or security that are inherent in a NAT approach. While Section 6 identifies additional optimization work, to the degree that these techniques improve a network manager's ability to explicitly audit or control access, and thereby manage the overall attack exposure of local resources, they act to improve local network security.

8. Conclusion

This document has described a number of techniques that may be combined on an IPv6 site to protect the integrity of its network architecture. These techniques, known collectively as Local Network Protection, retain the concept of a well-defined boundary between "inside" and "outside" the private network and allow firewalling, topology hiding, and privacy. However, because they preserve address transparency where it is needed, they achieve these goals without the disadvantage of address translation. Thus, Local Network Protection in IPv6 can provide the benefits of IPv4 Network Address Translation without the corresponding disadvantages.

The document has also identified a few ongoing IETF work items that are needed to realize 100% of the benefits of LNP.

9. Acknowledgements

Christian Huitema has contributed during the initial round table to discuss the scope and goal of the document, while the European Union IST 6NET project acted as a catalyst for the work documented in this note. Editorial comments and contributions have been received from: Fred Templin, Chao Luo, Pekka Savola, Tim Chown, Jeroen Massar, Salman Asadullah, Patrick Grossetete, Fred Baker, Jim Bound, Mark

Smith, Alain Durand, John Spence, Christian Huitema, Mark Smith, Elwyn Davies, Daniel Senie, Soininen Jonne, Kurt Erik Lindqvist, Cullen Jennings, and other members of the v6ops WG and IESG.

10. Informative References

[1] Rekhter, Y., Moskowitz, R., Karrenberg, D., Groot, G., and E. Lear, "Address Allocation for Private Internets", BCP 5, RFC 1918, February 1996.

[2] Srisuresh, P. and M. Holdrege, "IP Network Address Translator (NAT) Terminology and Considerations", RFC 2663, August 1999.

[3] Narten, T., Nordmark, E., and W. Simpson, "Neighbor Discovery for IP Version 6 (IPv6)", RFC 2461, December 1998.

[4] Hain, T., "Architectural Implications of NAT", RFC 2993, November 2000.

[5] Srisuresh, P. and K. Egevang, "Traditional IP Network Address Translator (Traditional NAT)", RFC 3022, January 2001.

[6] Holdrege, M. and P. Srisuresh, "Protocol Complications with the IP Network Address Translator", RFC 3027, January 2001.

[7] Narten, T. and R. Draves, "Privacy Extensions for Stateless Address Autoconfiguration in IPv6", RFC 3041, January 2001.

[8] IAB and IESG, "IAB/IESG Recommendations on IPv6 Address Allocations to Sites", RFC 3177, September 2001.

[9] Wasserman, M., "Recommendations for IPv6 in Third Generation Partnership Project (3GPP) Standards", RFC 3314, September 2002.

[10] Droms, R., Bound, J., Volz, B., Lemon, T., Perkins, C., and M. Carney, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 3315, July 2003.

[11] Draves, R., "Default Address Selection for Internet Protocol version 6 (IPv6)", RFC 3484, February 2003.

[12] Troan, O. and R. Droms, "IPv6 Prefix Options for Dynamic Host Configuration Protocol (DHCP) version 6", RFC 3633, December 2003.

[13] Huttunen, A., Swander, B., Volpe, V., DiBurro, L., and M. Stenberg, "UDP Encapsulation of IPsec ESP Packets", RFC 3948, January 2005.

[14] Savola, P. and B. Haberman, "Embedding the Rendezvous Point (RP) Address in an IPv6 Multicast Address", RFC 3956, November 2004.

[15] Arkko, J., Kempf, J., Zill, B., and P. Nikander, "SEcure Neighbor Discovery (SEND)", RFC 3971, March 2005.

[16] Baker, F., Lear, E., and R. Droms, "Procedures for Renumbering an IPv6 Network without a Flag Day", RFC 4192, September 2005.

[17] Hinden, R. and B. Haberman, "Unique Local IPv6 Unicast Addresses", RFC 4193, October 2005.

[18] Fuller, V. and T. Li, "Classless Inter-domain Routing (CIDR): The Internet Address Assignment and Aggregation Plan", BCP 122, RFC 4632, August 2006.

[19] Dupont, F. and P. Savola, "RFC 3041 Considered Harmful", Work in Progress, June 2004.

[20] Chown, T., "IPv6 Implications for TCP/UDP Port Scanning", Work in Progress, October 2005.

[21] Chown, T., Tompson, M., Ford, A., and S. Venaas, "Things to think about when Renumbering an IPv6 network", Work in Progress, September 2006.

[22] Huston, G., "Architectural Commentary on Site Multi-homing using a Level 3 Shim", Work in Progress, July 2005.

[23] Rosenberg, J., "Interactive Connectivity Establishment (ICE): A Methodology for Network Address Translator (NAT) Traversal for Offer/Answer Protocols", Work in Progress, October 2006.

[24] "Universal Plug and Play Web Site", July 2005, <http://www.upnp.org/>.

Appendix A. Additional Benefits Due to Native IPv6 and Universal Unique Addressing

The users of native IPv6 technology and globally unique IPv6 addresses have the potential to make use of the enhanced IPv6 capabilities, in addition to the benefits offered by the IPv4 technology.

A.1. Universal Any-to-Any Connectivity

One of the original design points of the Internet was any-to-any connectivity. The dramatic growth of Internet-connected systems coupled with the limited address space of the IPv4 protocol spawned address conservation techniques. NAT was introduced as a tool to reduce demand on the limited IPv4 address pool, but the side effect of the NAT technology was to remove the any-to-any connectivity capability. By removing the need for address conservation (and therefore NAT), IPv6 returns the any-to-any connectivity model and removes the limitations on application developers. With the freedom to innovate unconstrained by NAT traversal efforts, developers will be able to focus on new advanced network services (i.e., peer-to-peer applications, IPv6-embedded IPsec communication between two communicating devices, instant messaging, Internet telephony, etc.) rather than focusing on discovering and traversing the increasingly complex NAT environment.

It will also allow application and service developers to rethink the security model involved with any-to-any connectivity, as the current edge firewall solution in IPv4 may not be sufficient for any-to-any service models.

A.2. Auto-Configuration

IPv6 offers a scalable approach to minimizing human interaction and device configuration. IPv4 implementations require touching each end system to indicate the use of DHCP vs. a static address and management of a server with the pool size large enough for the potential number of connected devices. Alternatively, IPv6 uses an indication from the router to instruct the end systems to use DHCP or the stateless auto-configuration approach supporting a virtually limitless number of devices on the subnet. This minimizes the number of systems that require human interaction as well as improves consistency between all the systems on a subnet. In the case that there is no router to provide this indication, an address for use only on the local link will be derived from the interface media layer address.

A.3. Native Multicast Services

Multicast services in IPv4 were severely restricted by the limited address space available to use for group assignments and an implicit locally defined range for group membership. IPv6 multicast corrects this situation by embedding explicit scope indications as well as expanding to 4 billion groups per scope. In the source-specific multicast case, this is further expanded to 4 billion groups per scope per subnet by embedding the 64 bits of subnet identifier into the multicast address.

IPv6 allows also for innovative usage of the IPv6 address length and makes it possible to embed the multicast Rendezvous Point (RP) [14] directly in the IPv6 multicast address when using Any-Source Multicast (ASM). This is not possible with the limited size of the IPv4 address. This approach also simplifies the multicast model considerably, making it easier to understand and deploy.

A.4. Increased Security Protection

The security protection offered by native IPv6 technology is more advanced than IPv4 technology. There are various transport mechanisms enhanced to allow a network to operate more securely with less performance impact:

- IPv6 has the IPsec technology directly embedded into the IPv6 protocol. This allows for simpler peer-to-peer authentication and encryption, once a simple key/trust management model is developed, while the usage of some other less secure mechanisms is avoided (e.g., MD5 password hash for neighbor authentication).

- While a firewall is specifically designed to disallow applications based on local policy, it does not interfere with those that are allowed. This is a security improvement over NAT, where the work-arounds to enable applications allowed by local policy are effectively architected man-in-the-middle attacks on the packets, which precludes end-to-end auditing or IP level identification.

- All flows on the Internet will be better traceable due to a unique and globally routable source and destination IPv6 address. This may facilitate an easier methodology for back-tracing Denial of Service (DoS) attacks and avoid illegal access to network resources by simpler traffic filtering.

- The usage of private address space in IPv6 is now provided by Unique Local Addresses, which will avoid conflict situations when merging networks and securing the internal communication on a local network infrastructure due to simpler traffic filtering policy.

- The technology to enable source-routing on a network infrastructure has been enhanced to allow this feature to function, without impacting the processing power of intermediate network devices. The only devices impacted with the source-routing will be the source and destination node and the intermediate source-routed nodes. This impact behavior is different if IPv4 is used, because then all intermediate devices would have had to look into the source route header.

A.5. Mobility

Anytime, anywhere, universal access requires MIPv6 services in support of mobile nodes. While a Home Agent is required for initial connection establishment in either protocol version, IPv6 mobile nodes are able to optimize the path between them using the MIPv6 option header, while IPv4 mobile nodes are required to triangle route all packets. In general terms, this will minimize the network resources used and maximize the quality of the communication.

A.6. Merging Networks

When two IPv4 networks want to merge, it is not guaranteed that both networks are using different address ranges on some parts of the network infrastructure due to the usage of RFC 1918 private addressing. This potential overlap in address space may complicate a merging of two and more networks dramatically due to the additional IPv4 renumbering effort, i.e., when the first network has a service running (NTP, DNS, DHCP, HTTP, etc.) that needs to be accessed by the second merging network. Similar address conflicts can happen when two network devices from these merging networks want to communicate.

With the usage of IPv6, the addressing overlap will not exist because of the existence of the Unique Local Address usage for private and local addressing.

Authors' Addresses

Gunter Van de Velde
Cisco Systems
De Kleetlaan 6a
Diegem 1831
Belgium

Phone: +32 2704 5473
EMail: gunter@cisco.com

Tony Hain
Cisco Systems
500 108th Ave. NE
Bellevue, Wa.
USA

EMail: alh-ietf@tndh.net

Ralph Droms
Cisco Systems
1414 Massachusetts Avenue
Boxborough, MA 01719
USA

EMail: rdroms@cisco.com

Brian Carpenter
IBM
8 Chemin de Blandonnet
1214 Vernier,
CH

EMail: brc@zurich.ibm.com

Eric Klein
Tel Aviv University
Tel Aviv,
Israel

EMail: ericlklein.ipv6@gmail.com

Full Copyright Statement

Intellectual Property

Acknowledgement

Funding for the RFC Editor function is currently provided by the Internet Society.