Network Working Group   G. Van de Velde
Internet-Draft   T. Hain
Expires: July 28, 2005   R. Droms
                                                           Cisco Systems
                                                            B. Carpenter
                                                         IBM Corporation
                                                                E. Klein
                                                             TTI Telecom
                                                        January 24, 2005

IPv6 Network Architecture Protection
<draft-vandevelde-v6ops-nap-01.txt>

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 does not support NAT by design and this document shows how Network Architecture Protection (NAP) using IPv6 can provide the same or more benefits without the need for NAT.

Table of Contents

1. Introduction

Although there are many perceived benefits to Network Address Translation (NAT), its primary benefit of "amplifying" available address space is not needed in IPv6. The serious disadvantages of ambiguous "private" address space and of Network Address Translation (NAT) [2][6] have been well documented [5][7]. However, given its wide market acceptance NAT undoubtedly has some perceived benefits. Indeed, in an Internet model based on universal any-to-any connectivity, product marketing departments have driven a perception that some connectivity and security concerns can only be solved by using a NAT device or by using logically separated LAN address spaces. This document describes the market-perceived reasons to utilize a NAT device in an IPv4 environment and shows how these needs can be met and even exceeded with IPv6. The use of the facilities from IPv6 described in this document avoids the negative impacts of translation and may be described as Network Architecture Protection (NAP).

As far as security and privacy is concerned, this document considers how to mitigate a number of threats. Some are obviously external, such as having a hacker trying to penetrate your network, or having a worm infected machine outside your network trying to attack it. 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 any device on the LAN. Some may be embedded such as having some firmware in a domestic appliance "call home" to its manufacturer without the user's consent.

This document describes several techniques that may be combined on an IPv6 site to protect the integrity of its network architecture. These techniques, known collectively as NAP, retain the concept of a well defined boundary between "inside" and "outside" the private network, and allow firewalling, topology hiding, and privacy and will achieve these goals without address translation.

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

Function	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 for stationary systems see section 4.4
Addressing Autonomy	RFC 1918 see section 2.5	RFC 3177 & ULA see section 4.5
Global Address Pool Conservation	RFC 1918 see section 2.6	340,282,366,920,938,463,463,374,607,431,768,211,456 (3.4*10^38) addresses see section 4.6
Renumbering and Multi-homing	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 NAP can provide each of them. It

concludes with a IPv6 NAP case study and a gap analysis of work that remains to be done for a complete NAP solution.

2. Perceived benefits of NAT and its impact on IPv4

This section provides visibility into the generally perceived benefits of the use of IPv4 NAT. The goal of this description is not to analyze these benefits or discuss the accuracy of the perception (detailed discussions in [5]) , 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 internal network

A NAT device can connect a private network with any kind of address (ambiguous [RFC 1918] or global registered address) towards the Internet. The address space of the private network can be built from globally unique addresses, from ambiguous address space or from both simultaneously. Without specific configuration from public to private, the NAT device enables access between the client side of an application in the private network with the server side in the public Internet.

Wide-scale deployments have shown that using NAT to attach a private IPv4 network to the Internet is simple and practical for the non-technical end user. Frequently a simple user interface is sufficient for configuring both device and application access rights.

Additionally, thanks to successful marketing campaigns it is perceived by end users that their equipment is protected from the bad elements and attackers on the public IPv4 Internet.

2.2 Simple security due to stateful filter implementation

It is frequently believed that a NAT device puts in an extra barrier to keep the private network protected from evil outside influences due to the session-oriented character of NAT technology. Since a NAT typically keeps state only for individual sessions, attackers, worms, etc. cannot exploit this state to attack a host in general and on any port. This benefit may be partially real; however, experienced hackers are well aware of NAT devices and are very familiar with private address space, and have devised methods of attack (such as trojan horses) that readily penetrate NAT boundaries. The secure feeling is in vain.

Address translation 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 address translation function. The perceived security 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.

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 a simple router. In situations where 2 or more devices need to host the same application this complexity shifts from difficult to impossible.

2.3 User/Application tracking

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 could in theory be done by logging the network address translation details of the private and the public addresses of the NAT devices state database.

The 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 an 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

The ability of NAT to provide internet access by the use of a single (or few) global IPv4 routable addresses to a large community of users offers a simple mechanism to hide the internal topology of a network. In this scenario the large community will be reflected in the internet by a single (or few) IPv4 address(es).

The use of NAT then results in a user behind a NAT gateway actually appearing on the Internet as a user inside the NAT box itself; i.e., the IPv4 address that appears on the Internet is only sufficient to identify the NAT. 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 host computers by means of their IP addresses. At the same time a NAT creates a smaller pool of addresses for a much more focused point of attack.

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 enterprises prefer to hide as much as possible of their internal network topology from outsiders. Mostly this is achieved by blocking "traceroute" etc., but NAT of course entirely hides the internal subnet topology, which some network managers believe is a useful precaution to mitigate scanning attacks. Scanning for IPv6 can be much harder in comparison with IPv4 as described in [18]

Once a list of available devices and IP addresses 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. These open ports could be used for initiating attacks on an end system.

2.5 Independent control of addressing in a private network

Many private IPv4 networks take benefit from using 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 clean and hierarchical addressing structure in the network.

Another benefit is due to the usage of independent addresses on majority of the network infrastructure there is an increased ability to change provider with less operational difficulties.

2.6 Global address pool conservation

Due to the ongoing depletion of the IPv4 address range, the remaining pool of unallocated IPv4 addresses is below 30%. While mathematical models based on historical IPv4 prefix consumption periodically attempt to predict the future exhaustion date of the IPv4 address pool, a direct result of this continuous resource consumption is that

the administrative overhead for acquiring globally unique IPv4 addresses will continue increasing in direct response 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. In turn this has restricted the number of and types of applications that can be deployed by these ISPs and their customers. 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

The elements of multihoming and renumbering are quite different. However, multihoming is often a transitional state for renumbering, and NAT interacts with both in the same way.

For enterprise networks, it is highly desirable 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 CIDR prefix [1] 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, 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 allow more time to renumber the network components. The longterm solution is a single network without usage of NAT to avoid the ongoing operational complexity of overlapping addresses.

This solution may be sufficient for some networks; however when the merging networks were already using address translation it will create huge problems due to admistrative difficulties of the merged address space.

3. Description of the IPv6 tools

This section describes several features that can be used to provide the protection features associated with IPv4 NAT.

3.1 Privacy addresses (RFC 3041)

There are situations where it is desirable to prevent device

profiling, such as by contacted web sites, so IPv6 privacy addresses were defined to provide that capability. IPv6 addresses consist of a routing prefix subnet-id part (SID) and an interface identifier part (IID). For interfaces that contain embedded IEEE Link Identifiers the interface identifier is typically derived from it, though this practice facilitates tracking and profiling of a device as it moves around the Internet. RFC 3041 describes an extension to IPv6 stateless address autoconfiguration for interfaces [8]. Use of the privacy address extension causes nodes to generate global-scope addresses from interface identifiers that change over time, even in cases where the interface contains an embedded IEEE link identifier. 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 side 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 is expected to be from end systems running stateless autoconfiguration, there is nothing that prevents a DHCP server from running the RFC 3041 algorithm for any new IEEE identifier it hears, then remembering that for future queries. This would allow using them in DNS for registered services since the assumption of a 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, 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 3.4. Additional considerations are discussed in [17]. Providing privacy for a subnet ID will require different technology.

3.2 Unique Local Addresses

Local network and application services stability during periods of intermittent connectivity between one or more providers requires address management autonomy. Such autonomy in a single routing prefix environment would lead to massive expansion of the global routing tables, so IPv6 provides for simultaneous use of multiple prefixes. The Unique Local Address prefix (ULA) [16] has been set aside for use in local communications. The ULA address prefix for any network is routable over a locally defined collection of routers.

These prefixes are NOT to be routed on the public global Internet as that would have a serious negative impact on global routing.

ULAs have the following characteristics:
o Globally unique prefix
   * Allows networks to be combined or privately interconnected without creating any address conflicts or requiring renumbering of interfaces using these prefixes
   * If accidentally leaked outside of a network via routing or DNS, there is no conflict with any other addresses
o ISP independent and can be used for communications inside of a network without having any permanent or intermittent Internet connectivity
o Well known prefix to allow for easy filtering at network boundaries
o In practice, applications may treat these addresses like global scoped addresses

3.3 DHCPv6 prefix delegation

The Prefix Delegation (DHCP-PD) options [11] provide a mechanism for automated delegation of IPv6 prefixes using the Dynamic Host Configuration Protocol (DHCP) [10]. This mechanism (DHCP-PD) is intended for delegating a long-lived prefix from a delegating router to a requesting router, 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

These should be globally routable IPv6 addresses which can be randomly and independently assigned to IPv6 devices.

The random assignment has as purpose to confuse the outside world on the structure of the local network. However for the local network there is a correlation between the location of the device and the untraceable IPv6 address. This correlation could be done by generating IPv6 host route entries or by utilizing an aggregation device like a Mobile IPv6 Home Gateway.

The main goal of untraceable IPv6 addresses is to create an apparently unpredictable network infrastructure as seen from external networks to protect the local infrastructure from malicious outside influences or from mapping 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 reachable on very different parts of the local network instead of belonging to the same subnet next to each

other.

4. Using IPv6 technology to provide the market perceived benefits of NAT

The facilities in IPv6 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 has the role of managing both packet Routing and local address management. A basic IPv6 router could have a default configuration to advertize inside the site a locally generated random ULA prefix, independently from the state of any external connectivity. This would allow local nodes to communicate amongst themselves prior to establishing a global connection. If the network happened to concatenate with another local network, this is highly unlikely to result in address collisions. With external connectivity the simple gateway could also use DHCP-PD to acquire a routing prefix from the service provider for use when connecting to the global Internet. End node connections involving other nodes on the global Internet will always use the global IPv6 addresses [9] derived from this prefix delegation. In the very simple case there is no explicit routing protocol and a single default route is used out to the global Internet. A slightly more complex case might involve local 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 a default route would continue to be consistent with the connectivity.

4.2 IPv6 and Simple security

The vulnerability of an IPv6 host is similar as for an IPv4 host directly connected towards the Internet, and firewall and IDS systems are recommended. However, with IPv6, the following protections are available without the use of NAT:

1. Short lifetimes on privacy extension suffixes reduce the attack profile since the node will not respond to the address once the address is no longer valid.
2. IPsec is a mandatory service for IPv6 implementations. IPsec functions to prevent session hijacking, prevent content tampering, and optionally masks the packet contents. While IPsec might be available in IPv4 implementations, deployment in NAT environments either breaks the protocol or requires complex helper services with limited functionality or efficiency.

3. The size of the typical subnet ::/64 will make a network ping sweep and resulting port-scan virtually impossible due to the amount of possible combinations available. This goes from the assumption that the attacker has no access to a local connection. If an attacker has local access then he could use ND [4] and ping6 to ff02::1 to detect local neighbors. (Of course, a locally connected attacker has many scanning options with IPv4 as well.) It is recommended for site administrators to take [18] into consideration to achieve the expected goal.

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 pose some security holes while many people assume their network is secure due to the usage of NAT. This is directly the opposite of what IPv6 security best-practices are trying to achieve.

An example of a potential set of firewall rules could be:

```
        Source_A:       IPv6 Home broadband user
                        located on the internal network
        Destination_B:  IPv6 HTTP server
                        located on the external network

        On the edge broadband router a security rule could be:

        Internal network -> external network:

       Actions:
            Allow all traffic
            Create reflective session state (true) for the session

        External network -> internal network

       Actions:
         If the session had reflective 'true'-state,
              then allow the inbound traffic
         If the session had reflective 'false' state,
              then drop the traffic
```

This simple rule would create similar protection and security holes the typical IPv4 NAT device will offer and may for example be enabled by default on all broadband edge-routers,with that difference that the security caveats will be documented, and may hence be removed with the next revision of the rule. The goal is that at every iteration, the IPv6 internet will become more secure for the oblivious users.

Assuming the network administrator is aware of [18] the increased size of the IPv6 address will make topology probing much harder, and almost impossible for IPv6 devices. What one does when topology probing is to get an idea of the available hosts inside an enterprise. This mostly starts with a ping-sweep. This is an automated procedure of sending Internet Control Message Protocol (ICMP) echo requests (also known as PINGs) to a range of IP addresses and recording replies. This can enable an attacker to map the network. Since the IPv6 subnets are 64 bits worth of address space, this means that an attacker has to send out a simply unrealistic number of pings to map the network, and virus/worm propagation will be thwarted in the process. At full rate 40Gbps (400 times the typical 100Mbps LAN, and 13,000 times the typical DSL/Cable access link) it takes over 5000 years to scan a single 64 bit space.

4.3 User/Application tracking

IPv6 enables the collection of information about data flows. Due to the fact that 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. 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.

4.4 Privacy and topology hiding using IPv6

Partial host privacy is achieved in IPv6 using pseudo-random privacy addresses (RFC 3041) which are generated as required, so that a session can use an address that is valid only for a limited time. Exactly like IPv4 NAT, this only allows such a session to be traced back to the subnet that originates it, but not immediately to the actual host.

If a network manager wishes to conceal the internal IPv6 topology, and the majority of its host computer addresses, a good option will be to run all internal traffic using ULA since such packets can by definition never exit the site. For hosts that do in fact need to generate external traffic, by using multiple IPv6 addresses (ULAs and one or more global addresses), it will be possible to hide and mask some or all of the internal network. As discussed above, there are multiple parts to the IPv6 address, and different techniques to manage privacy for each.

When a network manager also wishes to conceal the internal IPv6 topology, by using explicit host routes it is possible to locate nodes on one segment while they appear externally to be on another.

An alternative method to hide the internal topology would be to use Mobile IPv6 internally without route optimization where the public facing addresses are consolidated on an edge Home Agent (HA), then use MIPv6 in bidirectional tunnel mode between the HA and topology masked node using the ULA as a COA This truly masks the internal topology as all nodes with global access appear to share a common subnet. There is no reason that rack mounted devices shouldn't be considered mobile nodes to mask the internal topology. It looks equivalent to running a VPN to a central server, however it does not involve any encryption or significant overhead.

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 a NAT is not required (or even defined) between the ULA and the public Internet. Nodes that need access to the public Internet may have a ULA for local use, and will have a global use address because the global use IPv6 address space is not a scarce resource like the global use IPv4 space. While global IPv6 allocation policy is managed through the Regional Internet Registries, it is expected that they will continue with derivatives of RFC 3177 for the foreseeable future.

When using IPv6, the need to ask for more address space will become far less likely due to the increased size of the subnets. These subnets typically allow 2^64 hosts per subnet and an enterprise will typically receive a /48 which allows segmentation into at least 2^16 different subnets.

The ongoing subnet size maintenance may become simpler when IPv6 technology is utilised. If IPv4 address space is optimised one has periodically to look into the number of hosts on a segment and the subnet size allocated to the segment; an enterprise today may have a mix of /28 - /23 size subnets for example, and may shrink/grow these as their network user base/etc changes. In v6, it's all /64.

4.6   Global address pool conservation

IPv6 provides sufficient space to completely avoid the need for overlapping address space, 340,282,366,920,938,463,463,374,607,431,768,211,456 (3.4*10^38) total possible addresses. 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 which will lead to the IPv4 practice 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 IPv4 constricting scenarios of multiple devices sharing a scarce global address.

4.7 Multihoming and renumbering

Multihoming and renumbering remain technically challenging with IPv6 (see the Gap Analysis below). However, IPv6 was designed to allow sites and hosts to run with several simultaneous CIDR-like prefixes and thus with several simultaneous ISPs. An address selection mechanism [12] is specified so that hosts will behave consistently when several addresses are simultaneously valid. The fundamental difficulty that IPv4 has in this regard therefore does not apply to IPv6. IPv6 sites can and do run today with multiple ISPs active, and the processes for adding and removing active prefixes at a site have been documented [15] and [19].

The IPv6 address space allocated by the ISP will be dependent upon the connecting Service provider. This may result in a renumbering effort if the network changes from Service Provider. When changing ISPs or ISPs readjusting their addressing pool, DHCP-PD [13] can be used as the 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

It is possible to divide the type of networks in different categories. This can be done on various criteria. The criteria used within this document are based on the number of components or connections. For each of these category of networks we can use IPv6 Network Architecture 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

Under this category fall the majority of private enterprise networks. Many of these networks have one or more exit points to the Internet. Though these organizations have sufficient resources to acquire addressing independence 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 address 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 one avoids the need for network renumbering. Finally, the customer can provide privacy about 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. A single /48 alloaction provides an enterprise network with 65536 different /64 prefixes.

The summarization benefit for the ISP is happening based on the IPv6 allocation rules. This means that the ISP will provide the enterprise with an IPv6 address-range (typically a one or multiple range(s) of '/48') from its RIR assigned IPv6 address-space. The goal of this allocation mechanism is to decrease the total amount of entries in the internet routing table.

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 known IPv6 address. Privacy extensions add a random 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 host on the inside network. The usage of IPv6 privacy extensions does not mask the internal network structure of an enterprise network.

If there is need to mask the internal structure towards the external IPv6 internet, then some form of 'Untraceable' addresses may be used. These addresses will be derived from a local pool, and may be assigned to those hosts for which topology masking is required or

which want to reach the IPv6 Internet or other external networks. The technology to assign these addresses to the hosts could be based on DHCPv6. To complement the 'Untraceable' addresses it is needed to have at least awareness of the IPv6 address location when routing an IPv6 packet through the internal network. This could be achieved by 'route-injection' in the 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. 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 trade-off that masking is time and subnet based.

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. Dynamically allocated 'Untraceable' addresses may also facilitate and simplify the connectivity to the outside networks during renumbering, because the existing IPv6 central address pool could be swapped for the newly allocated IPv6 address pool.

The use of permanent ULA addresses on a site provides the benefit that even if an enterprise would change its ISP, the renumbering is only affecting 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 which 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 through economic pressure are typically forced today to use NAT. In most cases the received global IPv4 prefix is not fixed over time and is too long 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 makes it quite hard for the equipment to become a world wide Internet server (i.e. 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 a server, only one server of any given protocol type is possible per address, and then only when the address from the ISP is public.

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

- DHCPv6 Prefix-Delegation
- 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 towards the single user connection because there are enough global addresses available at essentially the same cost. Also if the single device

wants to mask its identity to the called party or its attack profile over a short time window it will need to enable IPv6 privacy extensions, which in turn 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 which 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 which 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 makes it quite hard for the equipment to become a world wide internet server (i.e. HTTP, FTP, etc.) due to the stateful operation of the NAT equipment.

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

In 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 towards 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 IPv4 access and transport services. They tend to have three separate IPv4 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 network which addresses their backbone and access switches, and other hardware, this is separate for both 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 a and b and sometimes c) or can be from a pool of RFC 1918 addresses used with 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 NAP is utilized in these three domains then for the first category it will be possible to use the same solutions as described in chapter 5.1. The second domain of the ISP/carrier is the Operations network. This environment tends to be a closed environment, and consequently intra- communication can be done based on ULA addresses. This would give a stable configuration with respect to a local IPv6 address plan. Using these local scope addresses would also prevent from being accessed from the external network. 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 global 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 3GPP attached devices. When looking in chapter 2.3 of RFC3314 'Recommendations for IPv6 in 3GPP Standards September 2002' [9] it is found that the IPv6 WG recommends that one or more /64 prefixes should be assigned to each primary 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 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 of NAP. None of these items are show-stoppers for immediate usage of NAP in roles where there are no current gaps.

6.1 Completion of work on ULAs

As noted above, a new form of Unique Local IPv6 Unicast Addresses (ULAs) is being standardized by the IETF. Experience to date has shown that most network managers want to gain some operational familiarity with IPv6 in their local environment before exposing their network to the live global Internet. Since these addresses allow autonomy for local deployment of IPv6 in private networks, this work should be completed as soon as possible. In addition to autonomy the routing limitation of ULA addresses protects nodes that are only for local use from global exposure.

6.2 Subnet topology masking

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

As an alternative some work in Mobile IP to define a policy message where a mobile node would learn from the home agent that it should not even try to inform its correspondent about route optimization (and thereby expose its real location) would allow a border home agent using internal tunneling to the logically mobile node (potentially rack mounted) to completely mask all internal topology while avoiding the strain from a large number of host routes in the IGP. This work should be pursued in the IETF.

6.3 Minimal traceability of privacy addresses

Privacy addresses (RFC 3041) 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 with topology masking is required. This work remains to be done, and should be pursued by the IETF.

6.4 Renumbering procedure

Documentation of site renumbering procedures [15] should be completed. It should also be noticed that ULAs will help here too, since a change of ISP prefix will only affect hosts that need an externally routeable address as well as a ULA.

6.5 Site multihoming

This complex problem has never been well solved for IPv4, which is exactly why NAT has been used as a partial solution. For IPv6, after several years' work, the relevant IETF WG is finally converging on an architectural approach intended to reconcile enterprise and ISP perspectives. Again, ULAs will help since they will provide stable addressing for internal communications that are not affected by multihoming.

6.6 Untraceable addresses

The details of the untraceable addresses, along with any associated mechanisms such as route injection, must be worked out and specified.

7. IANA Considerations

This document requests no action by IANA

8. Security Considerations

While issues which are potentially security related are discussed throughout the document, the approaches herein do not introduce any new security concerns. Product marketing departments have widely sold IPv4 NAT as a security tool, though the misleading nature of those claims has been previously documented in RFC 2663 [3] and RFC 2993 [5].

This document defines IPv6 approaches which collectively achieve the goals of the network manager without the negative impact on applications or security that are inherent in a NAT approach. To the degree that these techniques improve a network manager's ability to explicitly know about or control access, and thereby manage the overall attack exposure of local resources, they act to improve local network security. In particular the explicit nature of a content aware firewall in NAP will be a vast security improvement over the NAT artifact where lack of translation state has been widely sold as a form of protection.

9. 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 Network Architecture 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, Network Architecture 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 NAP.

10. 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 draft. Editorial comments and contributions have been received from: Fred Templin, Chao Luo, Pekka Savola, Tim Chown, Jeroen Massar,

Salman Asadullah, Patrick Grossetete and other members of the v6ops WG.

11. References

11.1 Normative References

11.2 Informative References

[1] Fuller, V., Li, T., Yu, J. and K. Varadhan, "Classless Inter-Domain Routing (CIDR): an Address Assignment and Aggregation Strategy", RFC 1519, September 1993.

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

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

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

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

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

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

[8] Narten, T. and R. Draves, "Privacy Extensions for Stateless Address Autoconfiguration in IPv6", RFC 3041, January 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] Bush, R., Durand, A., Fink, B., Gudmundsson, O. and T. Hain, "Representing Internet Protocol version 6 (IPv6) Addresses in the Domain Name System (DNS)", RFC 3363, August 2002.

[12] Draves, R., "Default Address Selection for Internet Protocol

version 6 (IPv6)", RFC 3484, February 2003.

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

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

[15] Baker, F., Lear, E. and R. Droms, "Procedures for Renumbering an IPv6 Network without a Flag Day", Internet-Draft draft-ietf-v6ops-renumbering-procedure-03, November 2004.

[16] Hinden, R. and B. Haberman, "Unique Local IPv6 Unicast Addresses", Internet-Draft draft-ietf-ipv6-unique-local-addr-08, November 2004.

[17] Dupont, F. and P. Savola, "RFC 3041 Considered Harmful (draft-dupont-ipv6-   rfc3041harmful-05.txt)", June 2004.

[18] Chown, T., "IPv6 Implications for TCP/UDP Port Scanning (chown-v6ops- port-scanning-implications-01.txt)", July 2004.

[19] Chown, T., Tompson, M. and A. Ford, "Things to think about when Renumbering an IPv6 network (draft-chown-v6ops-renumber-thinkabout-00)", October 2004.

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 Corporation
Sauemerstrasse 4
Rueschlikon,   8803
Switzerland

Email: brc@zurich.ibm.com

Eric Klein
TTI Telecom
Petach Tikvah,
Israel

Phone: +972 3 926-9130
Email: erick@tti-telecom.com

Appendix A. Additional benefits due to Native IPv6 and universal unique addressing

The users of native IPv6 technology and global 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. Whereas 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, 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 on the local link only 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 'Rendez-Vous Point' (or RP) [14] directly in the IPv6 multicast address when using ASM multicast. this is not possible with 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 encryption and authentication, once a simple key/trust management model is developed, while the usage of some other less secure mechanisms is avoided (i.e. md5 password hash for neighbor authentication).
- On a local network, any user will have more security awareness. This awareness will motivate the usage of simple firewall applications/devices to be inserted on the border between the external network and the local (or home network) as there is no Address Translater and hance no false safety perception.
- 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 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 would be using different address-ranges on some parts of the network infrastructure due to the legitimate usage of RFC 1918 private addressing. This potential overlap in address space may complicate a merge 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..) which need to be accessed by the 2nd 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.

A.7 Community of interest

Although some Internet-enabled devices will function as fully-fledged Internet hosts, it is believed that many will be operated in a highly restricted manner functioning largely or entirely within a Community of Interest. By Community of Interest we mean a collection of hosts that are logically part of a group reflecting their ownership or function. Typically, members of a Community of Interest need to communicate within the community but should not be generally accessible on the Internet. They want the benefits of the connectivity provided by the Internet, but do not want to be exposed to the rest of the world. This functionality will be available through the usage of NAP and native IPv6 dataflows, without any stateful device in the middle. It will also allow to build virtual organization networks on the fly, which is very difficult to do in IPv4+NAT scenarios.

Appendix B. Revision history

B.1 Changes from *-nap-00 to *-nap-01

- o Document introduction has been revised and overview table added
- o Comments and suggestions from nap-00 draft have been included.
- o Initial section of -00 draft 2.6 and 4.6 have been aggregated into a new æcase studyÆ section 5.
- o The list of additional IPv6 benefits has been been placed into appendix.
- o new security considerations section added.
- o GAP analysis revised.
- o Section 2.6 and 4.6 have been included.

Intellectual Property Statement

Disclaimer of Validity

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Copyright Statement

Acknowledgment

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