Implications of Large Scale Network Address Translation (NAT)

A BROADBAND INTERNET TECHNICAL ADVISORY GROUP
TECHNICAL WORKING GROUP REPORT

A Near-Uniform Agreement Report
(100% Consensus Achieved)

Issued:
March 2012

Executive Summary

Introduction

The Internet is running out of addresses in the format in which they were originally standardized, known as IPv4, due to aspects of that format which constrain the address space to a relatively "small" number of unique addresses as compared to the burgeoning number of devices requiring those same addresses to function on the Internet. A successor address format, IPv6, has been developed to support as many devices as can conceivably be connected to the Internet for the foreseeable future. While the global transition to IPv6 is in progress, it is going to take a number of years to upgrade *all* Internet applications and services, consumer electronics devices, and networks to support IPv6. The transition period is also likely to be lengthy given that, among other things, IPv4-only equipment is still being manufactured and sold to consumers. As network operators deploy IPv6 technology into their existing IPv4 networks, IPv4 and IPv6 will thus need to co-exist until the demand for IPv4 services diminishes.

Given the amount of time it may take to migrate out of a pure IPv4, or mixed IPv4 and IPv6 network environment, to pure IPv6 service, network operators are employing a variety of techniques to extend the life of IPv4 addressing. One such technique is the use of Large Scale Network Address Translation (also known as Large Scale NAT or LSN). LSN equipment allows a large number of IPv4-enabled end devices to share a single public IPv4 address. Network Address Translation (NAT) functionality has long existed in local/private networks to help network operators manage their network addresses using private address space. NAT functionality is known to adversely impact some Internet applications; wider use of NAT as part of LSN therefore deserves careful examination.

The BITAG is interested in LSN given that IP address sharing is a key tool for extending the life of IPv4 during the transition to IPv6. LSN is likely to affect many players in the Internet ecosystem: ISPs, end users, application providers, equipment vendors, content delivery networks, and third parties such as law enforcement agencies. A broad understanding of problems that may arise has the following benefits: (1) it will help stakeholders to prepare for actions that minimize the impact on end users and applications; (2) it will inform policymakers and regulators of the motivations and trade-offs for the deployment of this technology; (3) it will accelerate the transition to IPv6; and (4) it will more generally help to reduce or preclude friction and/or conflict surrounding use of this technique among stakeholders, as some have argued that Large Scale NAT could be abused by parties for anti-competitive, discriminatory, or other non-technical purposes.

LSN Deployments and Impacts

LSN will be deployed in different ways depending on which IPv6 transition technologies are in use. These alternatives are discussed in the body of this paper. For all of these alternatives, there are a variety of technical implications of LSN for end users, ISPs, and application providers to consider.

The address sharing enabled by LSN use impacts end users in three primary ways: (1) the number of connections available per user is affected, (2) the ability to uniquely identify an end user device solely via the source IP address is lost, and (3) it becomes much more difficult to reach and maintain connectivity to end user devices. All of these impacts are present in local/private network implementation of NATs. Introduction of LSN increases the probability that users will be affected due to sharing of the port number space. The number of users affected by the limitation in port availability may also increase for the same reason. However, note that a user checking email or performing simple web browsing functions will not be affected by the LSN.

Internet Service Providers (ISPs) electing to use LSNs must balance the impacts of new network infrastructure (operational and capital costs) and of engineering this new infrastructure for scalability, resiliency, security, and capacity, as well as meeting mandates to be able to log individual customer IP address assignments, with maintaining an appropriate level of customer service. In the mobile environment, where every device must be assigned at least one IP address and where simple devices may have limited access to Internet applications, mobile operators have already implemented LSN and faced some of these challenges. However, the continued swift growth in the number of mobile customers and their likely evolution toward expecting wireless Internet service to behave in a manner comparable to wireline service presents new challenges.

LSN can have a wide variety of impacts on applications. These may relate to capacity constraints if the LSN is undersized, the handling of multiple connections to the same application server, the loss of IP-based geolocation capability, new logging requirements, and a variety of other factors.

Recommendations

BITAG has compiled the following recommendations regarding steps that can be taken to help ensure optimal user experience, balanced with efficient LSN deployments and operations:

- **Commit to rapid deployment of IPv6.** The best way to mitigate the impacts of LSN is to reach a state where IPv6 is the dominant addressing scheme. BITAG suggests that ISPs deploy IPv6, that equipment manufacturers support IPv6 in their devices, and that applications sensitive to NAT be supported via IPv6 as soon as possible.

- **Address application impacts of LSN.** BITAG suggests that vendors of LSN equipment adhere to existing requirements [Common requirements for Carrier Grade NAT (CGN)] intended to increase the likelihood that applications will function properly in the presence of LSN. BITAG also suggests that ISPs test their LSN implementations and mitigate application issues prior to deployment, and that application developers use LSN work-arounds or avoid deploying services that do not function properly in the presence of NAT or LSN.

- **Disclose LSN deployment.** To assist with end user troubleshooting, BITAG suggests that ISPs be transparent with respect to the locations and timing of LSN deployment.

- **Provide mechanisms to facilitate LSN traversal to end users.** BITAG suggests that, where feasible, ISPs and equipment vendors support mechanisms to facilitate NAT traversal, including mechanisms for the manual or automatic creation and management of port forwarding rules. Such mechanisms increase the likelihood that applications requiring inbound connections to end users can function across LSN.

- **Provide contact information.** BITAG suggests that ISPs provide a means for application providers to contact them to discuss LSN impacts and possible mitigations.

- **Consider Logging Impacts of Port Allocation.** BITAG suggests that ISPs deploying LSN consider logging and operational impacts when deciding whether to implement a deterministic or dynamic mechanism (or a hybrid of the two) for assigning ports to subscriber sessions.

- **Include Port Number When Logging Activity**. BITAG suggests that Application Providers that maintain a log of user activity include both the IPv4 address and port number in the log. This would ensure that logs accurately reflect the actions of a single ISP customer when IPv4 traffic goes across a LSN.

Table of Contents

1 About BITAG

The Broadband Internet Technical Advisory Group (BITAG) is a non-profit, multi-stakeholder organization focused on bringing together engineers and technologists in a Technical Working Group (TWG) to develop consensus on broadband network management practices and other related technical issues that can affect users' Internet experience, including the impact to and from applications, content and devices that utilize the Internet.

BITAG's mission includes: (a) educating policymakers on such technical issues; (b) addressing specific technical matters in an effort to minimize related policy disputes; and (c) serving as a sounding board for new ideas and network management practices. Specific TWG functions also may include: (i) identifying "best practices" by broadband providers and other entities; (ii) interpreting and applying "safe harbor" practices; (iii) otherwise providing technical guidance to industry and to the public; and/or (iv) issuing advisory opinions on the technical issues germane to the TWG's mission that may underlie disputes concerning broadband network management practices.

BITAG TWG reports focus primarily on technical issues. While the reports may touch on a broad range of questions associated with a particular network management practice, the reports are not intended to address or analyze in a comprehensive fashion the economic, legal, regulatory or public policy issues that the practice may raise.

2 Introduction

The explosive growth of the Internet has been driven by computers, smart phones, netbooks, tablets and the like. While bringing new innovation, it is also creating new challenges. In particular, the world is running out of IP addresses in the format that was standardized in 1981, also known as IPv4 addresses. In short, IPv4 lacks sufficient address space (4.3 billion) to keep up with the surging demand. IPv6 is the successor IP address format to IPv4 and is capable of supporting many more devices.

In February 2011, the Internet Assigned Numbers Authority (IANA) allocated the remaining last five large blocks of IPv4 address space to the Regional Internet Registries (RIRs). And to help prepare for IPv6, earlier this year, the Internet Society (ISOC) hosted the first World IPv6 Day, which involved organizations and companies from around the world offering content over IPv6 for a period of 24 hours as a sort of "test flight." The purpose was to encourage all industry stakeholders to adequately gear up for the actual transition to IPv6. Domestically, the White House Office of Management and Budget (OMB) set a June 30, 2008 deadline for Federal agencies to support IPv6 in their backbone networks and the Federal CIO has required Federal agencies to upgrade their public-facing web sites and services to support IPv6 by September 30, 2012 and to upgrade applications that communicate with the public Internet to use native IPv6 by September 30, 2014.

While the IPv6 transition should be inevitable, it is going to take a number of years to upgrade all Internet applications and services, consumer electronics (CE) devices, and small- and large-scale (enterprise and commercial) networks to support IPv6. All parties have a role in this transition. However, a key goal will be to minimize impacts to end users, who will not want to see service to the Internet disrupted, nor be required to unnecessarily replace devices that only support IPv4.

Given the time horizon to migrate to IPv6 service, providers are employing a variety of conservation techniques to extend the life of IPv4 addressing. One such technique is the deployment of Large Scale Network Address Translation (LSN, also known as Carrier Grade NAT or CGN) equipment that allows a large number of IPv4-enabled end devices to share a single IPv4 address. Since publicly routable IPv4 addresses are becoming scarce, LSN allows enterprises, Internet Service Providers (ISPs), wireless providers, and other networks that serve large numbers of devices to extend the lifespan of IPv4 connectivity and permit end users to extend the life of their existing devices. The use of LSN in a service provider network is not necessarily a solution that is desired, but one of necessity given the time horizon of IPv6 migration and IPv4 address availability.

Network Address Translation (NAT) functionality has long existed in local/private networks and is known to adversely impact some Internet applications that rely upon a unique relationship between the IP address and the end user of the application. Therefore, wide use of NATs in a service provider network presents implementation challenges that should be examined.

This document will provide background on LSN deployments, discuss the implications of LSN deployments on various elements of the Internet community, and suggest some mechanisms that may be used to help mitigate some of the challenges of LSN deployments.

3 Issue Overview

This section introduces the issues associated with IPv4 address exhaustion and the rationale for IPv4 address sharing. A brief tutorial about IPv4 and IPv6 addresses can be found at [IPv4 vs. IPv6 – What Are They, Exactly?].

3.1 IPv4 Address Exhaustion and IPv6 Transition

The transition from an IPv4 Internet to an IPv6 Internet involves a coexistence phase between the two technologies. Older equipment and software (in residential networks this includes telephones, printers, computers, set-top boxes, switches, and routers) is generally IPv4-only, while newer equipment often supports both technologies. It is unreasonable to expect all users to change all of their equipment at once. End users generally replace equipment piecemeal as issues of age or capability demand such action. Hence, network service providers must accommodate both IPv4 and IPv6 for what will likely be a lengthy

transition period. In essence, network service providers must deploy IPv6 technology into their existing IPv4 networks at a pace that allow both IPv4 and IPv6 to co-exist, ideally in a transparent manner to users, until the demand for IPv4 services diminishes to a level permitting the withdrawal of support for IPv4.

The length of this coexistence period is indeterminate. However, it is generally estimated in equipment lifetimes. Consumers own the equipment, in most instances, and therefore control when they replace devices in their homes. Often consumers tend to take an "if it works do not replace it" approach, making it highly likely that consumers will continue to use and maintain IPv4-only devices for many years to come. Additionally, given that IPv4-only devices are still manufactured and sold to consumers today, much of this equipment is at the very beginning of its lifetime.

The IETF's recommendations along these lines, notably RFC 4213, *Basic Transition Mechanisms for IPv6 Hosts and Routers,* [RFC 4213], envisioned that this parallel deployment would happen during the latter stages of IPv4 address usage, with a view to a smooth transition well in advance of the exhaustion of available IPv4 addresses. However, because of the explosive growth of Internet connections, this did not happen. Hence, by the end of 2011, the pool of available public IPv4 addresses approached depletion, access network deployments of IPv6 were just beginning, and IPv4-only consumer electronics devices were still being produced and bought by consumers. Business operations demand, therefore, that the IPv4 Internet remain viable while users discover a need to change networks and equipment, even though IPv4's basic resource – IPv4 addresses – is dwindling.

3.2 IPv4 Address Sharing and LSN Deployments

NAT technology is not new. Private networks (home, small business, and many enterprise networks) have long used NATs to manage their network addresses through private IPv4 addresses [RFC 1918] without the complexity of dealing with acquisition of public IPv4 address space for all devices and users of those private networks. Because these networks are not affected by depletion of private IPv4 addresses (i.e., they use IP address that are not exposed to or routable on the Internet), consumer electronics devices that are intended to be used inside these networks, for intra-network usages, are proving very slow to implement IPv6.

Worldwide, the deployment of Large Scale NATs (LSNs) is not a new phenomenon either. For various reasons, networks in South America, Africa, and Central Asia have used layered NATs to multiplex and expand a limited address space. Use of LSNs in this manner has resulted in service impacts, which both the operators and users have generally worked around, although in some cases it has been with some difficulty. For example, where a provider has multiple layers of NAT in their architecture, they may find it difficult or even impossible to provide an email server that is reachable both by users and by those trying to send email to these users.

Some operators in the US have also used network address translation. Most mobile networks use it to provide mobile handsets with access to the Internet. [1]

Operators providing wireline residential Internet Access Services may find that they need or will need to use LSNs to maintain their IPv4 business during the coexistence phase of IPv6 deployment (see example in Section 3.3.1). This may be as part of a dual-stack offering (IPv6 service together with IPv4 service that uses private IPv4 address space), or it may be as part of an IPv6 transition technology offering. The Dual-Stack Lite (DS-Lite) transition mechanism [RFC 6333] uses IPv6 with an IPv4 service offered "virtually" on top of the IPv6 service (this is known as "tunneling"). The IPv4 connection uses private address space, with a LSN, but without a NAT between the home network and the access network. The 6rd transition mechanism [RFC 5969] provides something of the reverse: IPv6 via a tunnel over IPv4. The IPv4 can be provided with either public address space or by using a LSN and private address space.

Other proposals for how to accomplish address sharing (other than LSN) [dIVI][dIVI-pd][4rd] are in draft form for consideration by the IETF. Because their future status is unknown, they are not addressed in this document.

[1] Sprint [New IPv6 survey released on labs.ripe.net], Verizon [Verizon Wireless], T-Mobile [Re: [v4tov6transition] draft-arkko-ipv6-transition-guidelines WGLC], China Mobile [Dual Stack Hosts Using "Bump-in-the-Host" (BIH)], China Telecom [Rapid Transition of IPv4 contents to be IPv6-accessible],

3.3 Example LSN Deployments

As noted in Section 3.2, LSN will be deployed in several different ways. This section provides greater detail on some of these.

3.3.1 Example 1: NAT444

In NAT444, the customer gateway is doing address translation (as described in Section 3.2) from the home network private IPv4 addresses to an IPv4 address provided to the customer through the broadband connection. This provider address is part of a pool that is subsequently translated by the LSN to a globally-routable IPv4 address that can be routed across the Internet.

The "444" part of NAT444 refers to the fact that a particular message will use 3 different IPv4 addresses as it is transported from an originating device to its destination. NAT444 can be used to provide a stand-alone IPv4 service, in conjunction with a native IPv6 service to provide dual-stack IPv4 and IPv6 connectivity, or in conjunction with an IPv6 tunneled transition technology, such as 6rd, to provide dual-stack IPv4 and IPv6 connectivity to the customer premises network.

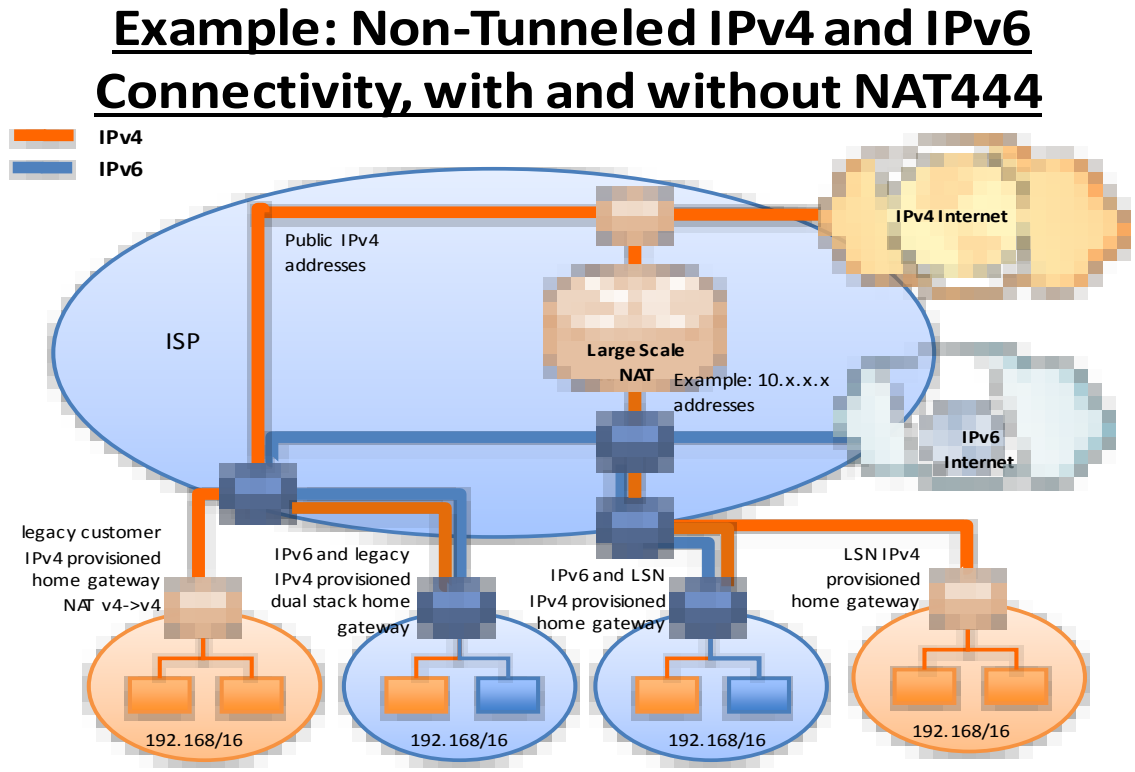

3.3.2 Example 2: Dual-Stack Lite (DS-Lite)

Dual-Stack Lite (DS-Lite) is an interesting deployment scenario because it removes the NAT from the customer gateway, so that only the LSN provides NAT functionality. DS-Lite is viewed as a transition technology that allows for continued IPv4 connectivity after IPv6 has been deployed. When DS-Lite is used, the IPv4 connection is provided as a tunnel that goes over the IPv6 connection.

Providers who deploy DS-Lite may also continue to support non-tunneled IPv4 for customers who need or want that sort of connectivity.

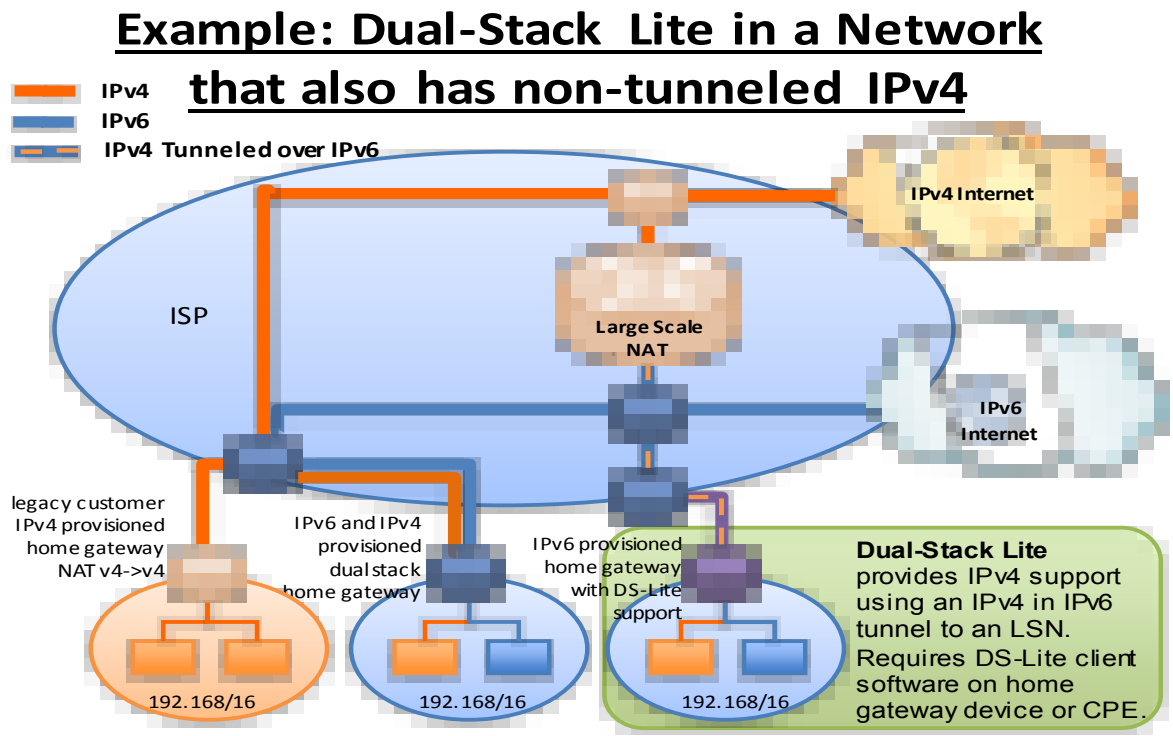

3.3.3 Example 3: Mobile Network

As dual stack is more expensive than a single IP stack to operate in a mobile network, some mobile network providers are considering giving devices either an IPv4 or an IPv6 address, but not both. Some mobile providers are considering providing those IPv4-address-only devices with the ability to access IPv6-only service by translating between the IPv4 address on the device and an IPv6 address for routing across the Internet. If the IPv4 address is a private IPv4 address, then NAT44 would also be used. Some mobile providers are considering providing NAT64 capability to those devices with only an IPv6 address (to provide access to IPv4-only service, translating between an IPv6 address on the device and a public IPv4 address for routing across the Internet). NAT64 has additional complexities that do not exist for NAT44 [RFC6052] [RFC6144] [RFC6145] [RFC6146] [RFC6147].

An additional option being considered by some mobile providers is to provide a dual stack IP address (IPv4 or IPv6, whichever the device does not have already) on demand, when the device attempts to access a site or application on the Internet that calls for that particular type of address. Since this on-demand connection does not necessarily use a NAT, it is mentioned only for completeness and will not be discussed further in this document.

Some mobile providers may provide different connections to different customers, depending on the service the customers subscribe to.

Example: NAT Options in a Mobile Network

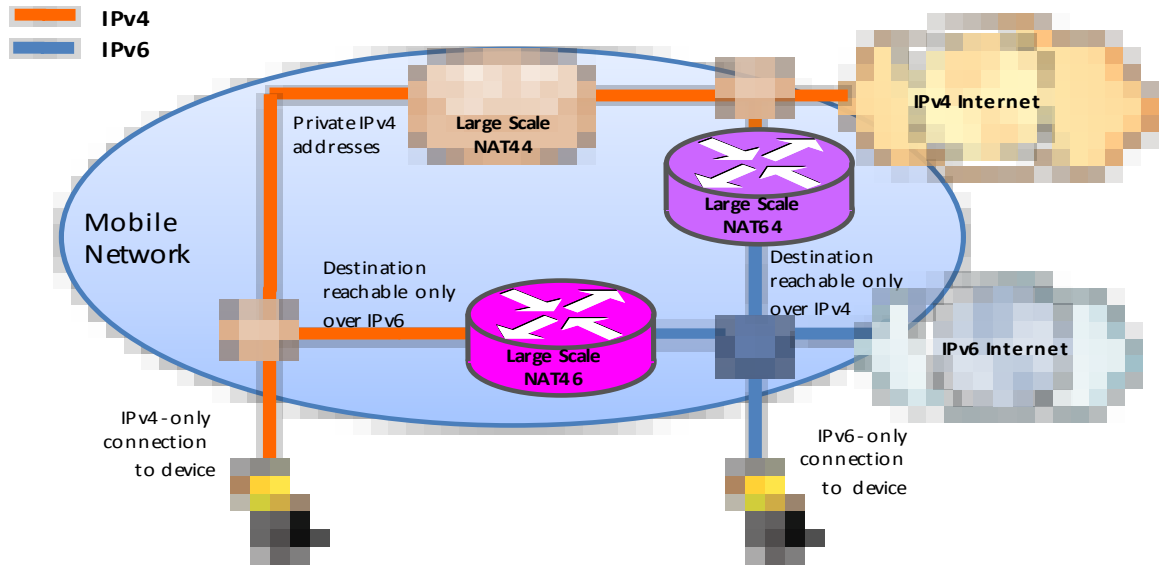

3.4 Overview of LSN Implementation Considerations

While not all Internet Service Providers (ISPs) will deploy LSNs, and some may only deploy LSNs to certain customers (e.g., new customers, customers connected through certain technologies, customers who purchase specific services, etc.), some ISPs do consider LSN as the only means to continue adding new customers and to be able to provide all customers with IPv4 addresses. LSN is a valid solution for extending the life of the IPv4 address space, but it is not a solution that an ISP will deploy without first considering all impacts and trade-offs as they relate to the ISP's business and customers.

Internet Service Providers (ISPs) must consider a variety of impacts when deploying LSNs. These include impacts of new network infrastructure (operational and capital costs) and engineering this new infrastructure for scalability, resiliency, security, and capacity, as well

as meeting mandates to be able to log individual customer IP address assignments, while maintaining a level of service that customers will accept.

Mobile Network Providers have many of the same concerns but must also deal with the constraints of efficiently operating a wireless network, as well as the fact that the number of smartphone and broadband home networking users (who expect Internet connectivity comparable to that of wireline Internet connections) is growing at a tremendous pace. Most mobile network providers are facing these challenges at the same time as they attempt to deploy new radio technologies, such as LTE, that will increase use of their networks for Internet access.

As stated earlier, IP address-sharing using LSN equipment may impact services and applications delivered to consumer electronics (CE) devices behind a LSN. For example, a recent CableLabs test [nat444-impacts] demonstrated that Netflix downloads succeed through a residential NAT but in some cases failed through multiple NATs. Thus, service providers and manufacturers of LSN equipment need to be cognizant of potential application impacts just as applications (on the Internet and in CE devices) need to be designed with the awareness that LSNs may exist. The IETF has documented a number of potential impacts of LSN deployment in RFC 6269 [RFC 6269]. It is theoretically possible to engineer LSNs to minimize adverse effects on popular/widely used higher-layer transport and application protocols. However, it is not possible to engineer for unknown, yet to be invented, applications. Hence there is a need for collaboration within the industry and for an awareness of when and to what extent LSN may be introduced.

The ability to determine accurate geo-location of users is also a concern both for many Application Providers, as well as services such as Next Generation 9-1-1 (NG9-1-1) [NENA 08-752].

4 BITAG Interest in This Issue

BITAG is interested in this issue as IP address sharing is a key tool for extending the life of IPv4 addressing during the transition to IPv6. It is likely to affect many players in the Internet ecosystem: ISPs, end-users, application providers, CPE equipment makers, content delivery networks, and third parties such as law enforcement agencies. A broad understanding of problems that may arise as a result of IP address sharing will help encourage stakeholders to take actions that minimize the impact on end-users and applications, will inform policy makers and regulators of the motivation and trade-offs for the deployment of this technology, and will accelerate the transition to IPv6. Finally, BITAG hopes to generally help reduce or preclude friction and/or conflict surrounding use of this technique among stakeholders, as some have argued that Large Scale NAT could be abused by parties for anti-competitive, discriminatory, or other non-technical purposes.

5 Impact Analysis

As described at a high level in Section 3.4, there are a variety of technical implications to consider when introducing or dealing with a LSN. This section provides further details on these implications, as they apply to the various elements of the Internet community (end users, ISPs, application providers, and other third parties). Recommendations for dealing with these implications are included in Section 6.

5.1 Implications for End Users

The address sharing enabled by LSN use impacts end users in three primary ways: (1) the number of connections available per user is affected, (2) the ability to uniquely identify an end user device solely via the source IP address is lost, and (3) it becomes much more difficult to reach and maintain connectivity to end user devices. All of these impacts are present in local/private network implementation of NATs. Introduction of LSN increases the probability that users will be affected because of the sharing of the port number space. The number of users affected by limitation in port availability may also increase for the same reason. Note that a user checking email or performing simple web browsing functions will not be affected by the LSN.

5.1.1 Port Limitations

Theoretically, a single IP address has access to a maximum of 65,535 port numbers usable to establish communication between two endpoints. In practice, routers have NAT tables which track and maintain open IP sessions for ports that are being used at a particular point in time. Consumer home routers tend to have NAT tables that are sized to allow somewhere between a few hundred to a few thousand ports to be used at one time. Most such routers are on the lower end of this range. The average consumer does not experience any limitations due to the size of the NAT table in his or her home router.

Address sharing limits the number of ports available per endpoint by sharing those ports across some number of endpoints. Depending on how many ports are available to any endpoint at a given time, application performance may be impacted. For example, some applications, such as Apple iTunes and Google Maps, make use of multiple ports for a single transaction. A single user running a single instance of an application on one personal computer may encounter no issues, while a family of four using multiple Internet connected devices concurrently could experience a situation where an open port is unavailable. Applications with no ports available for communication would not function as expected.

Accordingly, an Internet Service Provider implementing a LSN system will be faced with the challenge of how to configure port allocations per user. Two general options are available for allocating ports per public IPv4 address: dynamic port allocation and static port allocation. The use of a purely dynamic allocation model ensures the most efficient use

of a service provider's available remaining public IPv4 address as all ports are available for any customer who needs them. There are two major drawbacks to this method. First, a small number of customers are capable of utilizing all the available ports in a region affecting other users in the area. Second, dynamic port allocation generates massive logs (see Section 5.2.2), which is a challenge to implement efficiently and a costly addition for a service provider with a large customer base.

The other option for port allocation is a static model sometimes referred to as bulk port allocation. In this model, a range of predetermined ports is allocated per user by various means. The benefit of this model is that logging requirements are greatly reduced as users can be effectively determined through an IP address plus port number, and the need to log all connections for legal requirements is reduced. The disadvantage is that each user now has a set limited number of ports available for applications to use.

See Section 6.6 for recommendations around customer port assignments.

5.1.2 User-to-IP Address Association

Many applications use the source IP address used to communicate with the application as an identifier for the user. A user of such an application, who has connectivity through a LSN, will be affected if other users that share the same IPv4 address interact with the same application. For example, some moderated web forums control posting rights based on the user's source IP address. Because LSNs are likely to serve a relatively large base of customers, the odds are fairly high that an application that relies upon the source IP address to limit access by a single user will instead affect many users within the same geographic area. Similarly, an application that uses the source IP address for user authentication may set the limit for sessions from a particular source IP address to one. If the application's limit algorithm remains unchanged, only the first user of multiple parties sharing the same IP address would be granted access. Hence, application developers and application service providers need to be aware of the timing and general geographic areas impacted by LSN implementation.

See Sections 6.1, 6.2, and 6.3 for recommendations intended to help with this issue.

5.1.3 Automatically- and Manually-Created Port Forwarding Rules

Nearly every residential gateway in service today implements NAT functionality to go between public and private IPv4 addresses. NAT, by its nature, blocks all unsolicited inbound communication into the home network. This is because it does not know what device to send inbound traffic to, unless there has been recent outbound traffic using the same address and port mapping, or the NAT has a port forwarding rule that tells it what to do.

Port forwarding rules map the combination of an external address and port to a corresponding internal address and port. Because these port forwarding rules are critical for allowing devices inside the home to be reachable from the Internet, various tools for activating services through the NAT have been developed. Two mechanisms exist to enable applications to function in this environment: Automatically-created and manually-created port forwarding rules. Some devices and applications, such as game consoles, automatically create port forwarding rules in the residential gateway by using the UPnP IGD protocol [UPnP IGD] to poke holes thought the NAT barrier. In the future, devices may also use PCP [Port Control Protocol (PCP)] to create these holes. Manual creation of port forwarding rules is another option that allows a home network administrator to configure rules allowing communication through the NAT. However, having such rules in the residential gateway will no longer be sufficient when there is also a LSN present. Such rules must exist in both the residential gateway and the LSN in order for the application to function.

See Section 6.4 for recommendations on handling this concern.

5.2 Implications for ISPs

ISPs will take different approaches to IPv6 transition and IPv4 service continuity, based on a wide range of considerations, such as existing network equipment limitations, rate of IPv4 address exhaustion, economic and resource constraints. Nonetheless, it is likely that many ISPs will need to deploy some variant of LSN to continue to support connectivity to the IPv4 Internet to meet the needs of end users and their IPv4-only CE devices. Since the LSN equipment is in the service path for customer traffic, the overall design for how the LSN equipment is incorporated into the ISP network requires careful engineering. Also the operational support for the LSN must be handled appropriately to ensure that the performance delivered by the ISP network to the customer is acceptable.

The primary impact to Internet service providers is the requirement to deploy a costly new infrastructure solely for the purpose of extending IPv4 Internet addressability and reachability to users. The benefit is straightforward: the ability to continue offering IPv4 Internet service to new customer activations. The disadvantage is that it is nearly impossible to assure that all end user services using a shared IPv4 address (via a LSN) will operate in precisely the same manner compared to using an unshared IPv4 address.

See Sections 6.1 and 6.2 for recommendations to help address this impact.

5.2.1 Security, Resiliency, and Capacity

The act of placing a LSN device into the provider's network poses both security and capacity challenges. Avoiding having a single point of failure is a common network design principle for increased fault tolerance. However, insertion of a LSN potentially introduces a single point of failure and, if the LSN is undersized, creates a possible capacity bottleneck in

the network. Because the LSN is a single point through which traffic passes, it also represents a potential "attack point" for parties with malicious intent.

The presence of a LSN has both a negative and a positive impact on the ability to detect and correct user devices infected by malware. On the negative side, the presence of a LSN makes it more difficult to identify the compromised hosts participating in coordinated botnet activities, because the source IP address is no longer uniquely assigned to a single customer. On the core network side of the LSN, malware detection algorithms could determine that a botnet is active somewhere within the base of customers subtending the LSN but these algorithms would be unable to isolate the activity to a single device or household. Therefore, it would be impossible to accurately block an attack based solely on the source IP address (which is currently a frequently used remedy). As a result, malware detection will need to adapt to the presence of LSN by identifying more than just the source IP address.

On the other hand, the presence of a LSN interferes with the command and control structure typically used by botnets to conceal the main distribution infrastructure of the malware. This command and control structure typically relies on compromised devices with IP addresses directly exposed to the Internet.[2] The presence of a LSN makes these servers impossible to reach without creating a port forwarding rule in the LSN. Although the botnets will adopt the same techniques to open pinholes as used by legitimate applications, this activity is easily detected in the network and can be used to quickly identify compromised devices and remove them from the network. As all traffic must flow through the LSN, it is an ideal location for implementing malware detection and remediation algorithms.

However, if a large group of end users is supported by a single LSN, it is possible that infected hosts could impact other customers subtending the LSN and could be invisible to monitoring steps at the LSN. This situation is addressed by monitoring not only the traffic passing through the LSN but also by monitoring the traffic that does not need to pass through the LSN. Network service providers deploying LSNs should keep this design consideration in mind.

See Section 6.2 for general recommendations that will help to address this concern.

5.2.2 Logging

A shared addressing scheme can impact a service provider's ability to readily respond to legal requests from law enforcement agencies and civil litigants for logged traffic of specific end users. Currently, service providers generally fulfill such requests by maintaining a log recording customers' IPv4 addresses and the length of time the addresses are assigned to

[2] Software is installed on these devices to enable them to act as a server (e.g., DNS, HTTP, SMTP, etc) from which malware updates are distributed to compromised hosts.

those customers. This log effectively associates a residential household or business to the IPv4 address used for communication on the Internet. This assignment history may be queried in response to a warrant, court order, civil or criminal subpoena, or other request for the account information associated with a particular IP address.

Once a LSN is deployed, an IP address no longer uniquely identifies a subscriber. At a minimum, the port must also be included in order to uniquely identify the source device. Accordingly, third-party interfaces will need to change and ISPs will need to gather and store more information to accurately log activity.

The storage and processing resources required for maintaining logs of subscribers to IPv4 addresses and ports are impacted by LSN design decisions. If port numbers are predictable, subscribers can be identified by address + port using a rule when LSNs are introduced. However, if ports (or addresses) are dynamically assigned, the log must associate a subscriber with the address + port that each subscriber uses at all times. The longer that an ISP is required to maintain such a log, the greater the resources and cost required to maintain the log. There has been considerable discussion on this topic in the IETF[3].

See Section 6.6 for recommendations on this topic.

5.2.3 Using Existing Private IPv4 Address Space

Care has to be taken to ensure that the provider address pool does not conflict with the pools that are used in the home networks; otherwise the home gateway will likely fail to forward packets from the home network to the provider network. This is one of the reasons that many providers use address space other than that which is documented in the IETF's RFC 1918. In some cases, providers are using globally-routable unicast addresses (often referred to as GUA space) that have been allocated to them by a Regional Internet Registry (e.g. ARIN, APNIC, RIPE, etc.).

No specific recommendation is made to address this concern.

5.2.4 Minimizing Impacts on Users

Operators incur costs when users have negative user experiences. These include churn, support costs, etc. Therefore operators deploying LSNs have a strong incentive to minimize the adverse impacts on user experience described in Section 5.1. Many operators are

---

[3] Following is a non-exhaustive list of some of the drafts that can be found at http://tools.ietf.org/html/ and that are exemplary of the discussions that have occurred at the IETF: draft-carpenter-v6ops-icp-guidance, draft-dec-stateless-4v6, draft-donley-behave-deterministic-cgn, draft-golovinsky-cloud-services-log-format, draft-ietf-behave-lsn-requirements, draft-ietf-savi-threat-scope, draft-jpdionne-behave-cgn-mib, draft-kuarsingh-lsn-deployment, draft-operators-softwire-stateless-4v6-motivation, draft-oreirdan-mody-bot-remediation, draft-sivakumar-behave-nat-logging, draft-chown-v6ops-address-accountability.

testing LSN deployments against a variety of applications prior to deployment, so that potential problems are identified, known, and addressed – where practical – prior to LSN deployment. Some early adopters of LSNs are openly sharing their results in an effort to increase awareness of these potential impacts. An example of a report of such lab testing is at [nat444-impacts]. This practice should be encouraged.

See Section 6.2, 6.3, and 6.4 for recommendations that will help address this.

5.3 Implications for Mobile Network Providers

While there has been a rapid growth in both mobile and stationary home devices in recent years, the impact on the carriers involved is very different. A wireline broadband provider may see an increase in traffic as a result of the additional devices; however, if these devices are all behind a single home router, only a single IP address is consumed. In the mobile environment, every device must be assigned its own IP address, and in some architectures, a single device will have multiple IP addresses assigned to it for various tasks. When coupled with the rapid adoption of smart phones that are frequently communicating with the Internet over a data connection, many mobile carriers have had no other option than to implement LSN as their available IPv4 address space is depleted.

LSNs have long been used in many mobile networks for simple handheld devices that have limited ability to access the Internet and applications on the Internet (such as email and web browsing). These are used in a NAT44 configuration to translate between a private IPv4 address that is assigned to a handheld device and a public IP address that is routable over the Internet.

Even though some mobile operators have already deployed LSNs, and have effectively dealt with the challenges, new ones will arise. For example, as mobile carriers roll out 4G networks with speeds comparable to broadband, more consumers may consider replacing their existing home Internet connections with wireless connections that employ routers with a wireless modem for in-home networking. These new customers will bring with them the expectations of an Internet connection that they can use just like a wireline broadband connection. This expectation has not been faced, previously, on a wide scale.

Not surprisingly, mobile network providers who choose to deploy NAT44 or NAT64 have many of the same network challenges as those faced by ISPs. NAT64 presents additional challenges [RFC6052] [RFC6144] [RFC6145] [RFC6146] [RFC6147].

5.4 Implications for Application Providers

The IETF has documented many of the impacts that LSN has on various applications in RFC 6269 [RFC 6269]. A few of those impacts are specifically noted in this section.

5.4.1 Constraints on Real-Time Services

To the extent that LSN deployments consolidate entry points into end-user ISP networks, such deployments could introduce capacity constraints that do not currently exist. Such constraints could pose difficulties for scaling video and other real-time, Internet-based consumer services if LSN capacity is not accurately engineered.

See Section 6.1 for recommendations for dealing with this concern.

5.4.2 Performance Impacts

Early testing of HTTP video delivery by CableLabs through Large Scale NAT devices demonstrated performance issues which, if present in production deployments, could impact real-time entertainment services that provide video delivery via HTTP [nat444-impacts].

See Section 6.1 and 6.2 for recommendations for dealing with this concern.

5.4.3 Determining Client Location

Current content delivery networks largely rely on client DNS resolver and source IP address to deduce the approximate geographic location of end users. This allows the application to provide a destination address for a content server that is close to the end user (proximity generally results in better performance). Source IP address is also currently used by many content providers to determine the location of a client for content licensing purposes. Estimating general client location by source IP address is used by a number of Internet applications today. As LSN deployment becomes more widespread, these applications will need to adapt.

See Section 6.1 for recommendations for dealing with this concern.

5.4.4 Geo-Location and NG9-1-1

In NG9-1-1 [NENA 08-752], a device that cannot directly acquire its own location (through GPS or AGPS, for example) may request its location from the access network to which it is attached. In the case of a device inside a home network, this may result in the device asking the access network for the location associated with its public IPv4 address (as perceived by a VoIP service provider or Internet site). When the IPv4 address is shared by many customers, absent other steps, there will be no location that can be associated with it that is sufficiently accurate, for example, to permit dispatching emergency services to a particular street address. As LSNs are deployed, a different form of device identity (for example, IPv4 address plus port number, or a physical layer identifier) or a different method of location configuration may be necessary.

See Section 6.1 for recommendations for dealing with this concern.

5.4.5 Logging

As noted in Section 5.2.2, interfaces for requesting information that associates an IP address to a specific user will need to change in order to track and request the port as well as the IP address. This will impact content and application providers as well as third parties (not directly involved in the delivery, use, or transport of content or data) that may need to request such data of ISPs, content providers, or application providers.

See Section 6.7 for recommendations for dealing with this concern.

6 Conclusions and Recommendations

Internet Service Providers (ISPs) deploying LSN face a number of challenges and have a variety of options to choose from for addressing these challenges when determining their exact deployment architectures. There is no "right choice" that addresses every possible challenge and fully mitigates every possible impact. However, if good engineering principles are applied and an open and collaborative environment exists between network providers (including ISPs, Internet connectivity providers, and mobile network providers) and application/content providers then the deployment of LSN technology will be orderly and the majority of most end users will notice little or no difference in their Internet access service.

The recommendations in this section are intended to provide advice regarding the steps the industry can take to help ensure the best user experience (and the least breakage of applications) balanced with efficient LSN deployments and operations.

The following table summarizes which of the subsequent sub-sections contain recommendations for various entities of the industry (ISP, Application Provider, and Equipment Manufacturer).

Section	Contains Recommendations for…
6.1	ISP, Equipment Manufacturer, Application Provider
6.2	ISP, Equipment Manufacturer, Application Provider
6.3	ISP
6.4	ISP, Equipment Manufacturer
6.5	ISP
6.6	ISP
6.7	Application Provider

6.1 Commit to Rapid Deployment of IPv6

The best way to mitigate the impacts of LSN is to rapidly reach a state where IPv6 is the dominant addressing scheme. This requires efforts by device manufacturers, application developers, network providers (including ISPs, Internet connectivity providers, and mobile network providers) and end users.

ISPs can deploy IPv6 either as a native service or by using a transition technology, such as 6rd. Both of these provide end-to-end IPv6 connectivity. BITAG suggests **ISPs** (including Mobile Network Providers) deploy IPv6.

BITAG suggests that **Equipment Manufacturers** (for home network devices, mobile devices, etc.) include support for IPv6 in their devices as soon as possible and to avoid producing new devices that do not include support for IPv6.

BITAG suggests that **Application Providers** whose applications are sensitive to NAT support their applications via IPv6 as quickly as possible.

Note that some of the slowness in adoption of IPv6 has been attributed to a "chicken and egg" problem where ISPs cannot justify investing in rapid deployment of IPv6 infrastructure because there is little content available over IPv6 and many end user devices still do not support IPv6. Device manufacturers and application providers make similar claims as to why they do not support IPv6. It is important for all of these parties to work in concert to make IPv6 deployment a reality. There are considerable costs associated with deploying IPv6, for all parties involved (device manufacturers, application providers, third parties, end users, and ISPs). The best-case scenario for a successful transition is one in which all parties are equally engaged in the process.

6.2 Address application impacts of LSN

Since preliminary testing of LSN implementations shows wide variation in impact on different applications, it is important that LSN equipment vendors make their implementations as robust as possible.

The IETF has defined a set of LSN behaviors that are intended to minimize the impact on applications [Common requirements for Carrier Grade NAT (CGN)]. LSNs compliant to this recommendation do not implement any firewall functions or limit inbound connections.

BITAG suggests that **Equipment Manufacturers** who are providers of LSN equipment adhere to these IETF recommendations.

BITAG suggests that **ISPs** who intend to deploy LSN thoroughly test their LSN implementations and work with their selected vendors to resolve or mitigate issues prior to deployment.

Application Providers should be aware that LSN deployments are becoming more prevalent and should avoid deploying new services on IPv4 that will break in the presence of single or double NAT. They may wish to develop work-arounds (e.g., application support for STUN [RFC5389], relay servers, session border controllers, etc.) to make LSN as transparent to the end-user as possible.

6.3 Disclose LSN Deployment

BITAG suggests that **ISPs** be reasonably transparent with respect to the locations and timing of LSN deployment. It is useful for end users to know they are behind a LSN, as this can impact any trouble-shooting they may engage in to resolve issues, for example with other parties such as application providers.

6.4 Provide mechanisms to facilitate LSN traversal to end-users

BITAG suggests **ISPs** support mechanisms to facilitate LSN traversal where feasible. At a minimum, users need to be able to create a small set of port forwarding rules. ISPs may also consider supporting Port Configuration Protocol (PCP) [Port Control Protocol (PCP)] to allow user applications to automatically create port forwarding rules.

BITAG suggests that **Equipment Manufacturers** who implement [UPnP IGD] IGD Port Forwarding capability in CE routers support an interworking function to go from UPnP to PCP. PCP may be supported between the LSN and the CE router to automate management of port forwarding rules in the LSN.

6.5 Provide Contact Information

BITAG suggests that **ISPs** provide a means for application providers to contact them and discuss impacts caused by LSN, and consider possible mitigations.

6.6 Consider Logging Impacts of Port Allocation

Some LSN implementations randomly assign a public IPv4 address and port for each IPv4 session that emanates from a user's home network. This could cause the logging requirements in a large service provider deployment to be challenging and, in some cases, impossible to deploy. The BITAG suggests **ISPs** deploying a LSN system consider logging impacts when deciding whether to implement a deterministic port allocation model that maps a number of ports to an IP address per subscriber, a dynamic port allocation mechanism, or some hybrid approach to port allocation. The deterministic model eases the logging infrastructure in addition to enforcing an equitable distribution of LSN resources. This ensures that heavy usage customers do not lock out low usage users by utilizing all of the available ports and addresses. However, deterministic allocation may make inefficient

use of port space. New hybrid options are evolving that leverage aspects of both the dynamic and deterministic models, such as described in [deterministic-cgn].

6.7 Include Port Number When Logging Activity

BITAG suggests that **Application Providers** that maintain a log of user activity include both the IPv4 address and port number in the log. The IPv4 address alone may be insufficient to identify a unique ISP customer if the IPv4 address is shared by multiple customers via a LSN.

7 References

[UPnP IGD] UPnP Forum, "Internet Gateway Device (IGD) V 2.0", September 2, 2010, <http://upnp.org/specs/gw/igd2/>.

[RFC 1918] Rekhter, Y., "Address Allocation for Private Internets", February 1996, <http://tools.ietf.org/html/rfc1918>.

[RFC 4213] Nordmark, E., Gilligan, R., "Basic Transition Mechanisms for IPv6 Hosts and Routers." October 2005, <http://tools.ietf.org/html/rfc4213>.

[RFC5389] Rosenberg, J., et al, "Session Traversal Utilities for NAT (STUN)", October 2008, <http://tools.ietf.org/html/rfc5389>.

[RFC 5969] Townsley, M., Troan, O., "IPv6 Rapid Deployment on IPv4 Infrastructures (6rd) - Protocol Specification", August 2010, <http://tools.ietf.org/html/rfc5969>.

[RFC6052] Bao, C., et al, "IPv6 Addressing of IPv4/IPv6 Translators", October 2010, <http://tools.ietf.org/html/rfc6052>.

[RFC6144] Baker, F., et al, "Framework for IPv4/IPv6 Translation", April 2011, <http://tools.ietf.org/html/rfc6144>.

[RFC6145] Li, X. et al, "IP/ICMP Translation Algorithm", April 2011, <http://tools.ietf.org/html/rfc6145> .

[RFC6145] Bagnulo, M.. et al, "Stateful NAT64: Network Address and Protocol Translation from IPv6 Clients to IPv4 Servers", April 2011, <http://tools.ietf.org/html/rfc6146> .

[RFC6147] Bagnulo, M., et al, "DNS64: DNS Extensions for Network Address Translation from IPv6 Clients to IPv4 Servers", April 2011, <http://tools.ietf.org/html/rfc6147>.

[RFC 6269], Ford, M., et al, “Issues with IP Address Sharing”, June 2011, <http://tools.ietf.org/html/rfc6269>.

[RFC 6333], Durand, A., et al, “Dual-Stack Lite Broadband Deployments Following IPv4 Exhaustion”, August 2011, <http://tools.ietf.org/html/rfc6333>.

[Common requirements for Carrier Grade NAT (CGN)] Perreault, S., et al, “Common requirements for Carrier Grade NAT (CGN)” (work in progress), October 24, 2011, <http://tools.ietf.org/html/draft-ietf-behave-lsn-requirements-04>.

[Port Control Protocol (PCP)], Wing, D., et al, “Port Control Protocol (PCP)” (work in progress), October 31, 2011, <http://tools.ietf.org/html/draft-ietf-pcp-base-17>.

[4rd] Despres, R., et al, “IPv4 Residual Deployment across IPv6-Service networks (4rd) ISP-NAT's made optional”, March 14, 2011, <http://tools.ietf.org/html/draft-despres-intarea-4rd>.

[dIVI] Bao, C., et al, "dIVI: Dual-Stateless IPv4/IPv6 Translation", July 10, 2011, <http://tools.ietf.org/html/draft-xli-behave-divi>.

[dIVI-pd] Li, X., et al, "dIVI-pd: Dual-Stateless IPv4/IPv6 Translation with Prefix Delegation", September 23, 2011, <http://tools.ietf.org/html/draft-xli-behave-divi-pd>

[nat444-impacts] Donley, C., et al, “Assessing the Impact of Carrier-Grade NAT on Network Applications”, October 31, 2011, <http://tools.ietf.org/html/draft-donley-nat444-impacts>.

[deterministic-cgn] Donley, C., et al, “Deterministic Address Mapping to Reduce Logging in Carrier Grade NATs”, September 26, 2011, < http://tools.ietf.org/html/draft-donley-behave-deterministic-cgn>.

[IPv4 vs IPv6 – What Are They, Exactly?] Kenney, B., “IPv4 vs IPv6 – What Are They, Exactly?”, February 2011, <http://www.thetechlabs.com/tech-news/ipv4-vs-ipv6/>.

[Rapid Transition of IPv4 contents to be IPv6-accessible] Sun, Qiong, et al, "Rapid Transition of IPv4 contents to be IPv6-accessible" (individual draft), July 11, 2011, <http://tools.ietf.org/html/draft-sunq-v6ops-contents-transition>.

[Dual Stack Hosts Using "Bump-in-the-Host" (BIH)], Huang, Bill, Deng, Hui, Savolainen, Teemu, "Dual Stack Hosts Using "Bump-in-the-Host" (BIH)” (individual draft), August 10, 2011, <http://tools.ietf.org/html/draft-ietf-behave-v4v6-bih>.

[Verizon Wireless] Parker, B., "IPv6 deployment plans at Verizon Wireless", June 16, 2010 [as presented during the "Google IPv6 Implementors Conference: Mobile Networks Session"], <http://www.youtube.com/watch?v=1GlRgaFriYU>.

[NENA 08-752] NENA, "NENA Technical Requirements Document (TRD) for Location Information to Support IP-Based Emergency Services", NENA 08-752, Issue 1, December 21, 2006 < https://nena.site-ym.com/resource/collection/2851C951-69FF-40F0-A6B8-36A714CB085D/NENA_08-752-v1_Location_Information_Support_IP_Based_EmergSvcs.pdf>.

[NANOG53 weds morning session notes], NANOG, "NANOG53 weds morning session notes", October 12, 2011, <http://kestrel3.netflight.com/2011.10.12-nanog53-morning-session.txt>.

[IPv4 Address Exhaustion: A Progress Report], Huston, G., "IPv4 Address Exhaustion: A Progress Report", October 12, 2011, <http://www.nanog.org/meetings/nanog53/presentations/Wednesday/Huston.pdf>.

[2011.10.12 NANOG53 weds morning session notes email thread], Petach, M., et al, "Mailing List Archive: NANOG: users", <http://www.gossamer-threads.com/lists/nanog/users/145679>.

[Re: [v4tov6transition] draft-arkko-ipv6-transition-guidelines WGLC], Byrne, C., "v6ops@ops.ietf.org", August 25, 2010, <http://comments.gmane.org/gmane.ietf.v6ops/11305>.

[New IPv6 survey released on labs.ripe.net],George, W., "Wes E [NTK] George", April 27, 2011. <http://www.mentby.com/george-wes-e-ntk/>.

[BITAG Bylaws], "BITAG Bylaws", July 28, 2011, <http://www.bitag.org/documents/BITAG_Bylaws.pdf>.

8 Glossary of Terms

- 6rd (IPv6 Rapid Deployment on IPv4 Infrastructures): An automatic tunneling mechanism tailored to advance deployment of IPv6 to end users via a service provider's IPv4 network infrastructure. Key aspects include automatic IPv6 prefix delegation to sites, stateless operation, simple provisioning, and service that is equivalent to native IPv6 at the sites that are served by the mechanism. [RFC 5969]

- Address Sharing (also, IP address sharing): Address sharing refers to the scenario where multiple devices make use of the same public IP address, by using a NAT to go between the public IP address and private IP addresses assigned to each device.

- Application Provider: An individual or entity that is engaged substantially in the provision of products, services, caching, or solutions which are transmitted over the Internet. Definition from [BITAG Bylaws].

- Botnet: A group of compromised (infected with software that allows a hacker to have some control) computers that are connected to the Internet.

- Dual Stack: Having IPv4 and IPv6 operating at the same time in a device.

- Dual Stack – Lite (DS-Lite): Enables a broadband service provider to share IPv4 addresses among customers by combining two well-known technologies: IP in IP (IPv4-in-IPv6) and Network Address Translation (NAT). [RFC 6333]

- Geo-Location: Identifying the geographic location of a device. Information on geographic location may be general (e.g., city, state, or even country) or very specific (e.g., latitude/longitude, street address).

- Global Unicast Address (GUA): A publicly routable IP address that is associated with a single node on the Internet. Note that this node may have a NAT function that allows the GUA to be shared by multiple devices (address sharing). See also Public Address.

- Internet Connectivity Provider: An individual or entity that is engaged substantially in the provision of Internet connectivity, transmission and routing services to end users by any means including, but not limited to, digital subscriber line ("DSL"), microwave, fiber, cable, broadband over power lines ("BPL"), wireless, or satellite. Definition from [BITAG Bylaws].

- Internet Service Provider (ISP): An entity that is engaged in the provision of Internet connectivity to end users. Where the ISP is also responsible for the physical connectivity to the end user, the ISP is also an Internet Connectivity Provider. However, some ISPs resell physical connectivity that is supplied by another entity and provide only the IP layer transport and routing service.

- Layered NAT: Multiple NATs that are connected in series, so that IP traffic must go through all NATs in the sequence. A specific example of a layered NAT is a NAT444. See NAT444.

- Mobile Network Provider: A provider who uses wireless technologies to provide various services. Where a Mobile Network Provider supplies Internet connectivity, it is also an Internet Connectivity Provider and an ISP, although this may or may not be the primary use of its network.

- Native IP: IP traffic that is not tunneled over another IP layer.

- Network Address Translation (NAT): Network Address Translation is used to conserve public IP addresses and is where IP addresses are mapped from one

address realm to another, providing transparent routing to end hosts. In other words, NAT facilitates using a single public IP address to route to multiple private IP addresses. Private IP addresses are those addresses that are used within an internal network (e.g., within a home or business network) and cannot go out to the public Internet. There are many variations of network address translation. As adapted from IETF, RFC-2663: "IP Network Address Translator (NAT) Terminology and Considerations," August 1999, <http://tools.ietf.org/html/rfc2663>; see also "Assessing the Impact of NAT444 on Network Applications draft-donley-nat444-impacts-01," October 2010, <http://tools.ietf.org/html/draft-donley-nat444-impacts-01>.

- NAT44: A NAT that translates between IPv4 addresses.

- NAT444: Two layered NAT44 instances that cause IPv4 traffic to undergo two IPv4 address translations. One example where this would occur, would be in the case where a consumer's home router uses a NAT and has IPv4 service provided by the ISP via a LSN.

- NAT64: A NAT that translates between IPv6 and IPv4. This function has been defined for the case where the host is IPv6 and the accessed service is IPv4. The term is also applied to the scenario where the host is IPv4 and the accessed service is IPv6; however this case is considered by many to be problematic and is not well-defined.

- Port: Source and destination port numbers are included in TCP (Transmission Control Protocol), UDP (User Datagram Protocol), and other less common IP transport protocol headers. IP transport protocols are used in conjunction with an IP protocol to provide information necessary to transport application data between IP endpoints. Since they are independent of the actual IP header, they are also independent of IPv4 and IPv6 address numbers or headers. Port numbers are 16-bit fields, which means that they can theoretically have any value from 0 to 65,536 (2^{16}). However, many port numbers have been assigned by IANA to be used for specific applications (see http://www.iana.org/protocols). Port numbers are divided into three ranges, where numbers from 0 to 1023 are called "System Port" numbers and are strictly managed by IANA, numbers from 1024 to 49151 are "User Ports" and are available for assignment through IANA, and numbers from 49152 to 65535 are "Dynamic Ports" that cannot be assigned. Applications are free to use these dynamic port numbers at any time.

- Port Forwarding: Applications in consumer devices expect to send and receive certain application protocols on specific ports (for example, port 80 for HTTP, port 21 for FTP, port 53 for DNS). Since all devices want to use the same ports for these protocols, the NAT must translate not only the IP address but also the port number that is being used. For every port that the NAT has open on its interface to the Internet, it maintains a mapping from that port to an IP address and port on the home network interface. These mappings are stored in the NAT table and are called

Port Forwarding rules. When static mappings are created, they are often referred to as pinholes.

- Private Address: Term used to refer to addresses that cannot be routed over the Internet, especially those defined in [RFC1918], which include 10.x.x.x, 192.168.x.x, and 172.16.x.x IPv4 addresses. The term can also be used in to describe IPv6 Unique Local Addresses (ULA).

- Public Address: Term used to describe an IP address that can be routed across the Internet. See also GUA.

- Transition Technology: An IPv6 transition technology is a mechanism that allows an IPv6 connection to be established by tunneling the IPv6 over IPv4. Examples include 6rd and 6to4. An IPv4 transition technology is a mechanism that allows an IPv4 connection to be established by tunneling the IPv4 over IPv6. Transition technologies are often used where native IPv6 or IPv4 connectivity is not available. See Native IP.

- Tunneling: Running an IP connection on top of another transport protocol (usually another IP connection). This allows two disjointed networks to connect to each other across an intermediate network that knows nothing of the protocol or address scheme of the disjointed networks.

9 Document Contributors and Reviewers

Mark Clougherty, Alcatel Lucent
Charles Kalmanek, Jr., AT&T
Barbara Stark, AT&T
Alissa Cooper, Center for Democracy and Technology
Michael Fargano, CenturyLink
Fred Baker, Cisco
Tom Klieber, Comcast
Stan Barber, Cox Communications
Jeff Finkelstein, Cox Communications
Jeff Good, Disney
Vint Cerf, Google
Leslie Daigle, Internet Society
Amer Hassan, Microsoft
Ken Florance, NETFLIX
Jason Weil, Time Warner Cable, Inc.
Jeff Swinton, Verizon