Independent Submission
Request for Comments: 5684
Category: Informational
ISSN: 2070-1721

P. Srisuresh
EMC Corporation
B. Ford
Yale University

February 2010

Unintended Consequences of NAT Deployments with Overlapping Address Space

Abstract

This document identifies two deployment scenarios that have arisen from the unconventional network topologies formed using Network Address Translator (NAT) devices. First, the simplicity of administering networks through the combination of NAT and DHCP has increasingly lead to the deployment of multi-level inter-connected private networks involving overlapping private IP address spaces. Second, the proliferation of private networks in enterprises, hotels and conferences, and the wide-spread use of Virtual Private Networks (VPNs) to access an enterprise intranet from remote locations has increasingly lead to overlapping private IP address space between remote and corporate networks. This document does not dismiss these unconventional scenarios as invalid, but recognizes them as real and offers recommendations to help ensure these deployments can function without a meltdown.

Status of This Memo

Copyright

Table of Contents

1. Introduction and Scope 3
2. Terminology and Conventions Used 4
3. Multi-Level NAT Network Topologies 4
   3.1. Operational Details of the Multi-Level NAT Network 6
        3.1.1. Client/Server Communication 7
        3.1.2. Peer-to-Peer Communication 7
   3.2. Anomalies of the Multi-Level NAT Network 8
        3.2.1. Plug-and-Play NAT Devices 10
        3.2.2. Unconventional Addressing on NAT Devices 11
        3.2.3. Multi-Level NAT Translations 12
        3.2.4. Mistaken End Host Identity 13
4. Remote Access VPN Network Topologies 14
   4.1. Operational Details of the Remote Access VPN Network 17
   4.2. Anomalies of the Remote Access VPNs 18
        4.2.1. Remote Router and DHCP Server Address Conflict 18
        4.2.2. Simultaneous Connectivity Conflict 20
        4.2.3. VIP Address Conflict 21
        4.2.4. Mistaken End Host Identity 22
5. Summary of Recommendations 22
6. Security Considerations 24
7. Acknowledgements 24
8. References 25
   8.1. Normative References 25
   8.2. Informative References 25

1. Introduction and Scope

The Internet was originally designed to use a single, global 32-bit IP address space to uniquely identify hosts on the network, allowing applications on one host to address and initiate communications with applications on any other host regardless of the respective host's topological locations or administrative domains. For a variety of pragmatic reasons, however, the Internet has gradually drifted away from strict conformance to this ideal of a single flat global address space, and towards a hierarchy of smaller "private" address spaces [RFC1918] clustered around a large central "public" address space. The most important pragmatic causes of this unintended evolution of the Internet's architecture appear to be the following.

1. Depletion of the 32-bit IPv4 address space due to the exploding total number of hosts on the Internet. Although IPv6 promises to solve this problem, the uptake of IPv6 has in practice been slower than expected.

2. Perceived Security and Privacy: Traditional NAT devices provide a filtering function that permits session flows to cross the NAT in just one direction, from private hosts to public network hosts. This filtering function is widely perceived as a security benefit. In addition, the NAT's translation of a host's original IP addresses and port number in a private network into an unrelated, external IP address and port number is perceived by some as a privacy benefit.

3. Ease-of-Use: NAT vendors often combine the NAT function with a DHCP server function in the same device, which creates a compelling, effectively "plug-and-play" method of setting up small Internet-attached personal networks that is often much easier in practice for unsophisticated consumers than configuring an IP subnet. The many popular and inexpensive consumer NAT devices on the market are usually configured "out of the box" to obtain a single "public" IP address from an ISP or "upstream" network via DHCP ([DHCP]), and the NAT device in turn acts as both a DHCP server and default router for any "downstream" hosts (and even other NATs) that the user plugs into it. Consumer NATs in this way effectively create and manage private home networks automatically without requiring any knowledge of network protocols or management on the part of the user. Auto-configuration of private hosts makes NAT devices a compelling solution in this common scenario.

[NAT-PROT] identifies various complications with application protocols due to NAT devices. This document acts as an adjunct to [NAT-PROT]. The scope of the document is restricted to the two

scenarios identified in sections 3 and 4, arising out of unconventional NAT deployment and private address space overlap. Even though the scenarios appear unconventional, they are not uncommon to find. For each scenario, the document describes the seeming anomalies and offers recommendations on how best to make the topologies work.

Section 2 describes the terminology and conventions used in the document. Section 3 describes the problem of private address space overlap in a multi-level NAT topology, the anomalies with the topology, and recommendations to address the anomalies. Section 4 describes the problem of private address space overlap with remote access Virtual Private Network (VPN) connections, the anomalies with the topology, and recommendations to address the anomalies. Section 5 describes the security considerations in these scenarios.

2. Terminology and Conventions Used

In this document, the IP addresses 192.0.2.1, 192.0.2.64, 192.0.2.128, and 192.0.2.254 are used as example public IP addresses [RFC5735]. Although these addresses are all from the same /24 network, this is a limitation of the example addresses available in [RFC5735]. In practice, these addresses would be on different networks.

Readers are urged to refer to [NAT-TERM] for information on NAT taxonomy and terminology. Unless prefixed with a NAT type or explicitly stated otherwise, the term NAT, used throughout this document, refers to Traditional NAT [NAT-TRAD]. Traditional NAT has two variations, namely, Basic NAT and Network Address Port Translator (NAPT). Of these, NAPT is by far the most commonly deployed NAT device. NAPT allows multiple private hosts to share a single public IP address simultaneously.

3. Multi-Level NAT Network Topologies

Due to the pragmatic considerations discussed in the previous section and perhaps others, NATs are increasingly, and often unintentionally, used to create hierarchically interconnected clusters of private networks as illustrated in figure 1 below. The creation of multi-level hierarchies is often unintentional, since each level of NAT is typically deployed by a separate administrative entity such as an ISP, a corporation, or a home user.

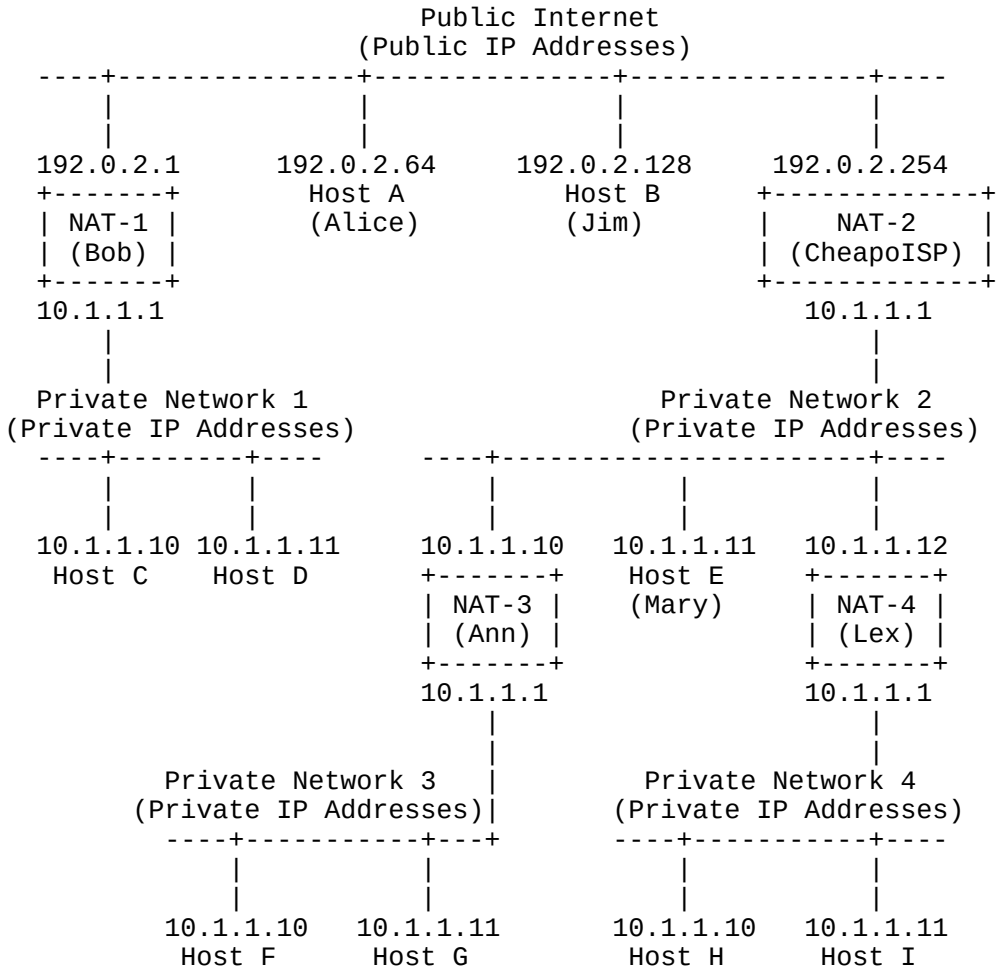

Figure 1. Multi-Level NAT Topology with Overlapping Address Space

In the above scenario, Bob, Alice, Jim, and CheapoISP have each obtained a "genuine", globally routable IP address from an upstream service provider. Alice and Jim have chosen to attach only a single machine at each of these public IP addresses, preserving the originally intended architecture of the Internet and making their hosts, A and B, globally addressable throughout the Internet. Bob, in contrast, has purchased and attached a typical consumer NAT box. Bob's NAT obtains its external IP address (192.0.2.1) from Bob's ISP via DHCP, and automatically creates a private 10.1.1.x network for Bob's hosts C and D, acting as the DHCP server and default router for this private network. Bob probably does not even know anything about IP addresses; he merely knows that plugging the NAT into the Internet

as instructed by the ISP, and then plugging his hosts into the NAT as the NAT's manual indicates, seems to work and gives all of his hosts access to Internet.

CheapoISP, an inexpensive service provider, has allocated only one or a few globally routable IP addresses, and uses NAT to share these public IP addresses among its many customers. Such an arrangement is becoming increasingly common, especially in rapidly developing countries where the exploding number of Internet-attached hosts greatly outstrips the ability of ISPs to obtain globally unique IP addresses for them. CheapoISP has chosen the popular 10.1.1.x address for its private network, since this is one of the three well-known private IP address blocks allocated in [RFC1918] specifically for this purpose.

Of the three incentives listed in section 1 for NAT deployment, the last two still apply even to customers of ISPs that use NAT, resulting in multi-level NAT topologies as illustrated in the right side of the above diagram. Even three-level NAT topologies are known to exist. CheapoISP's customers Ann, Mary, and Lex have each obtained a single IP address on CheapoISP's network (Private Network 2), via DHCP. Mary attaches only a single host at this point, but Ann and Lex each independently purchase and deploy consumer NATs in the same way that Bob did above. As it turns out, these consumer NATs also happen to use 10.1.1.x addresses for the private networks they create, since these are the configuration defaults hard-coded into the NATs by their vendors. Ann and Lex probably know nothing about IP addresses, and in particular they are probably unaware that the IP address spaces of their own private networks overlap not only with each other but also with the private IP address space used by their immediately upstream network.

Nevertheless, despite this direct overlap, all of the "multi-level NATed hosts" -- F, G, H, and I in this case -- all nominally function and are able to initiate connections to any public server on the public Internet that has a globally routable IP address. Connections made from these hosts to the main Internet are merely translated twice: once by the consumer NAT (NAT-3 or NAT-44) into the IP address space of CheapoISP's Private Network 2 and then again by CheapoISP's NAT-2 into the public Internet's global IP address space.

3.1. Operational Details of the Multi-Level NAT Network

In the "de facto" Internet address architecture that has resulted from the above pragmatic and economic incentives, only the nodes on the public Internet have globally unique IP addresses assigned by the official IP address registries. IP addresses on different private networks are typically managed independently -- either manually by

the administrator of the private network itself, or automatically by the NAT through which the private network is connected to its "upstream" service provider.

By convention, nodes on private networks are usually assigned IP addresses in one of the private address space ranges specifically allocated to this purpose in RFC 1918, ensuring that private IP addresses are easily distinguishable and do not conflict with the public IP addresses officially assigned to globally routable Internet hosts. However, when plug-and-play NATs are used to create hierarchically interconnected clusters of private networks, a given private IP address can be and often is reused across many different private networks. In figure 1 above, for example, private networks 1, 2, 3, and 4 all have a node with IP address 10.1.1.10.

3.1.1. Client/Server Communication

When a host on a private network initiates a client/server-style communication session with a server on the public Internet, via the server's public IP address, the NAT intercepts the packets comprising that session (usually as a consequence of being the default router for the private network), and modifies the packets' IP and TCP/UDP headers so as to make the session appear externally as if it were initiated by the NAT itself.

For example, if host C above initiates a connection to host A at IP address 192.0.2.64, NAT-1 modifies the packets comprising the session so as to appear on the public Internet as if the session originated from NAT-1. Similarly, if host F on private network 3 initiates a connection to host A, NAT-3 modifies the outgoing packet so the packet appears on private network 2 as if it had originated from NAT-3 at IP address 10.1.1.10. When the modified packet traverses NAT-2 on private network 2, NAT-2 further modifies the outgoing packet so as to appear on the public Internet as if it had originated from NAT-2 at public IP address 192.0.2.254. The NATs in effect serve as proxies that give their private "downstream" client nodes a temporary presence on "upstream" networks to support individual communication sessions.

In summary, all hosts on the private networks 1, 2, 3, and 4 in figure 1 above are able to establish a client/server-style communication sessions with servers on the public Internet.

3.1.2. Peer-to-Peer Communication

While this network organization functions in practice for client/server-style communication, when the client is behind one or more levels of NAT and the server is on the public Internet, the lack

of globally routable addresses for hosts on private networks makes direct peer-to-peer communication between those hosts difficult. For example, two private hosts F and H on the network shown above might "meet" and learn of each other through a well-known server on the public Internet, such as host A, and desire to establish direct communication between G and H without requiring A to forward each packet. If G and H merely learn each other's (private) IP addresses from a registry kept by A, their attempts to connect to each other will fail because G and H reside on different private networks. Worse, if their connection attempts are not properly authenticated, they may appear to succeed but end up talking to the wrong host. For example, G may end up talking to host F, the host on private network 3 that happens to have the same private IP address as host H. Host H might similarly end up unintentionally connecting to host I.

In summary, peer-to-peer communication between hosts on disjoint private networks 1, 2, 3, and 4 in figure 1 above is a challenge without the assistance of a well-known server on the public Internet. However, with assistance from a node in the public Internet, all hosts on the private networks 1, 2, 3, and 4 in figure 1 above are able to establish a peer-to-peer-style communication session amongst themselves as well as with hosts on the public Internet.

3.2. Anomalies of the Multi-Level NAT Network

Even though conventional wisdom would suggest that the network described above is seriously broken, in practice it still works in many ways. We break up figure 1 into two sub-figures to better illustrate the anomalies. Figure 1.1 is the left half of figure 1 and reflects the conventional single NAT deployment that is widely prevalent in many last-mile locations. The deployment in figure 1.1 is popularly viewed as a pragmatic solution to work around the depletion of IPv4 address space and is not considered broken. Figure 1.2 is the right half of figure-1 and is representative of the anomalies we are about to discuss.

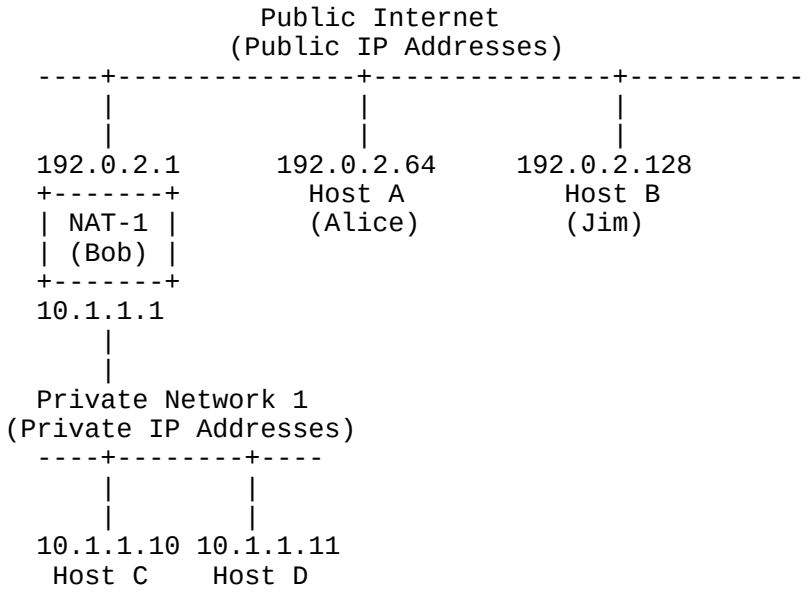

Figure 1.1. Conventional Single-level NAT Network topology

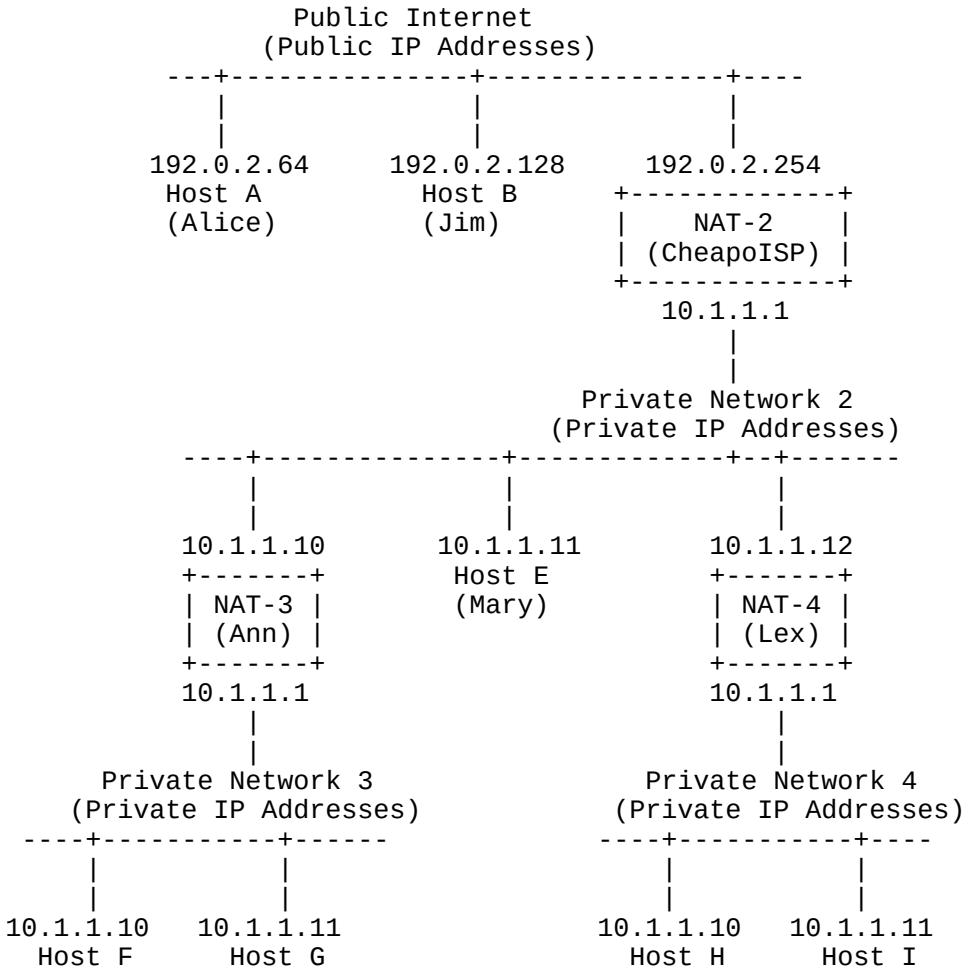

Figure 1.2. Unconventional Multi-Level NAT Network Topology

3.2.1. Plug-and-Play NAT Devices

Consumer NAT devices are predominantly plug-and-play NAT devices, and assume minimal user intervention during device setup. The plug-and-play NAT devices provide DHCP service on one interface and NAT function on another interface. Vendors of the consumer NAT devices make assumptions about how their consumers configure and hook up their PCs to the device. When consumers do not adhere to the vendor assumptions, the consumers can end up with a broken network.

A plug-and-play NAT device provides DHCP service on the LAN attached to the private interface, and assumes that all private hosts at the consumer site have DHCP client enabled and are connected to the single LAN. Consumers need to be aware that all private hosts must be on a single LAN, with no router in between.

A plug-and-play NAT device also assumes that there is no other NAT device or DHCP server device on the same LAN at the customer premises. When there are multiple plug-and-play NAT devices on the same LAN, each NAT device will offer DHCP service on the same LAN, and may even be from the same private address pool. This could result in multiple end nodes on the same LAN ending up with identical IP addresses and breaking network connectivity.

As it turns out, most consumer deployments have a single LAN where there they deploy a plug-and-play NAT device and the concerns raised above have not been an issue in reality.

3.2.2. Unconventional Addressing on NAT Devices

Let us consider the unconventional addressing with NAT-3 and NAT-4. NAT-3 and NAT-4 are apparently multi-homed on the same subnet through both their interfaces. NAT-3 is on subnet 10.1.1/24 through its external interface facing NAT-2, as well as through its private interface facing clients host F and host G. Likewise, NAT-4 also has two interfaces on the same subnet 10.1.1/24.

In a traditional network, when a node has multiple interfaces with IP addresses on the same subnet, it is natural to assume that all interfaces with addresses on the same subnet must be on a single connected LAN (bridged LAN or a single physical LAN). Clearly, that is not the case here. Even though both NAT-3 and NAT-4 have two interfaces on the same subnet 10.1.1/24, the NAT devices view the two interfaces as being on two disjoint subnets and routing realms. The plug-and-play NAT devices are really not multi-homed on the same subnet as in a traditional sense.

In a traditional network, both NAT-3 and NAT-4 in figure 1.2 should be incapable of communicating reliably as a transport endpoint with other nodes on their adjacent networks (e.g., private networks 2 and 3 in the case of NAT-3 and private Networks 2 and 4 in the case of NAT-4). This is because applications on either of the NAT devices cannot know to differentiate packets from hosts on either of the subnets bearing the same IP address. If NAT-3 attempts to resolve the IP address of a neighboring host in the conventional manner by broadcasting an Address Resolution Protocol (ARP) request on all of its physical interfaces bearing the same subnet, it may get a different response on each of its physical interfaces.

Even though both NAT-3 and NAT-4 have hosts bearing the same IP address on the adjacent networks, the NAT devices do communicate effectively as endpoints. Many of the plug-and-play NAT devices offer a limited number of services on them. For example, many of the NAT devices respond to pings from hosts on either of the interfaces. Even though a NAT device is often not actively managed, many of the NAT devices are equipped to be managed from the private interface. This unconventional communication with NAT devices is achievable because many of the NAT devices conform to REQ-7 of [BEH-UDP] and view the two interfaces as being on two disjoint routing domains and distinguish between sessions initiated from hosts on either interface (private or public).

3.2.3. Multi-Level NAT Translations

Use of a single NAT to connect private hosts to the public Internet as in figure 1.1 is a fairly common practice. Many consumer NATs are deployed this way. However, use of multi-level NAT translations as in figure 1.2 is not a common practice and is not well understood.

Let us consider the conventional single NAT translation in figure 1.1. Because the public and private IP address ranges are numerically disjoint, nodes on private networks can make use of both public and private IP addresses when initiating network communication sessions. Nodes on a private network can use private IP addresses to refer to other nodes on the same private network, and public IP addresses to refer to nodes on the public Internet. For example, host C in figure 1.1 is on private network 1 and can directly address hosts A, B, and D using their assigned IP addresses. This is in spite of the fact that hosts A and B are on the public Internet and host D alone is on the private network.

Next, let us consider the unconventional multi-level NAT topology in figure 1.2. In this scenario, private hosts are able to connect to hosts on the public Internet. But, private hosts are not able to connect with all other private hosts. For example, host F in figure 1.2 can directly address hosts A, B, and G using their assigned IP addresses, but F has no way to address any of the other hosts in the diagram. Host F in particular cannot address host E by its assigned IP address, even though host E is located on the immediately "upstream" private network through which F is connected to the Internet. Host E has the same IP address as host G. Yet, this addressing is "legitimate" in the NAT world because the two hosts are on different private networks.

It would seem that the topology in figure 1.2 with multiple NAT translations is broken because private hosts are not able to address each other directly. However, the network is not broken. Nodes on

any private network have no direct method of addressing nodes on other private networks. The private networks 1, 2, 3, and 4 are all disjoint. Hosts on private network 1 are unable to directly address nodes on private networks 2, 3, or 4 and vice versa. Multiple NAT translations were not the cause of this.

As described in sections 3.1.1 and 3.1.2, client-server and peer-to-peer communication can and should be possible even with multi-level NAT topology deployment. A host on any private network must be able to communicate with any other host, no matter to which private network the host is attached or where the private network is located. Host F should be able to communicate with host E and carry out both client-server communication and peer-to-peer communication, and vice versa. Host F and host E form a hairpin session through NAT-2 to communicate with each other. Each host uses the public endpoint assigned by the Internet-facing NAT (NAT-2) to address its peer.

When the deployed NAT devices conform to the hairpin translation requirements in [BEH-UDP], [BEH-TCP], and [BEH-ICMP], peer nodes are able to connect even in this type of multi-level NAT topologies.

3.2.4. Mistaken End Host Identity

Mistaken end host identity can result in accidental malfunction in some cases of multi-level NAT deployments. Consider the scenario in figure 1.3. Figure 1.3 depicts two levels of NATs between an end-user in private network 3 and the public Internet.

Suppose CheapoISP assigns 10.1.1.11 to its DNS resolver, which it advertises through DHCP to NAT-3, the gateway for Ann’s home. NAT-3 in turn advertises 10.1.1.11 as the DNS resolver to host F (10.1.1.10) and host G (10.1.1.11) on private network 3. However, when host F sends a DNS query to 10.1.1.11, it will be delivered locally to host G on private network 3 rather than CheapoISP’s DNS resolver. This is clearly a case of mistaken identity due to CheapoISP advertising a server that could potentially overlap with its customers’ IP addresses.

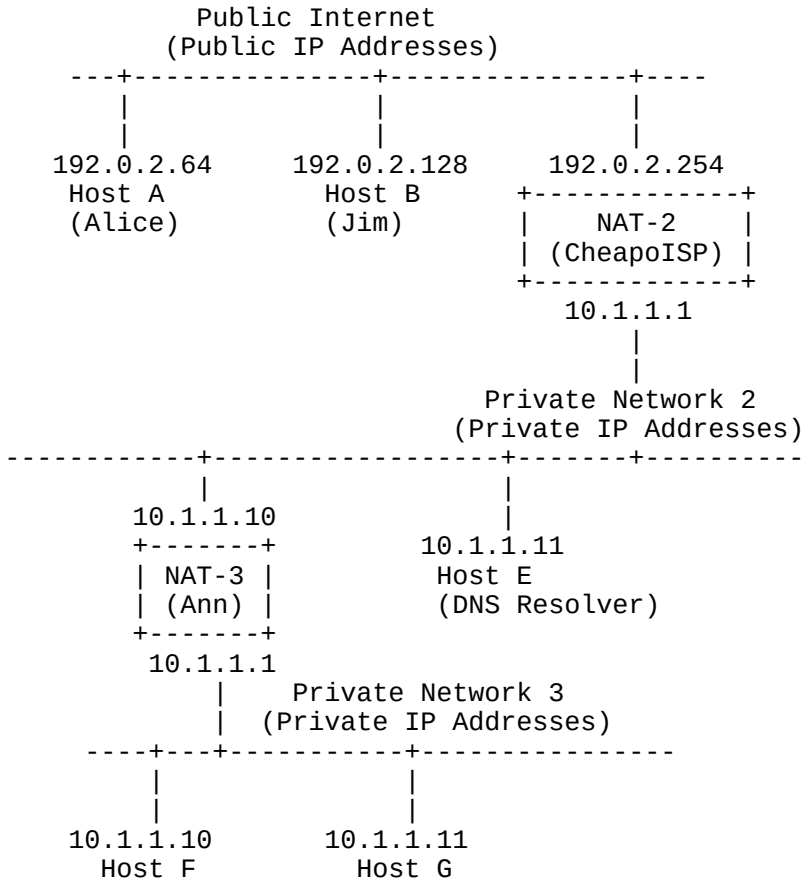

Figure 1.3. Mistaken Server Identity in Multi-Level NAT Topology

Recommendation-1: ISPs, using NAT devices to provide connectivity to customers, should assign non-overlapping addresses to servers advertised to customers. One way to do this would be to assign global addresses to advertised servers.

4. Remote Access VPN Network Topologies

Enterprises use remote access VPN to allow secure access to employees working outside the enterprise premises. While outside the enterprise premises, an employee may be located in his/her home office, hotel, conference, or a partner's office. In all cases, it is desirable for the employee at the remote site to have unhindered access to his/her corporate network and the applications running on

the corporate network. While doing so, the employee should not jeopardize the integrity and confidentiality of the corporate network and the applications running on the network.

IPsec, Layer 2 Tunneling Protocol (L2TP), and Secure Socket Layer (SSL) are some of the well-known secure VPN technologies used by the remote access vendors. Besides authenticating employees for granting access, remote access VPN servers often enforce different forms of security (e.g., IPsec, SSL) to protect the integrity and confidentiality of the run-time traffic between the VPN client and the VPN server.

Many enterprises deploy their internal networks using private address space as defined in RFC 1918 and use NAT devices to connect to the public Internet. Further, many of the applications in the corporate network refer to information (such as URLs) and services using private addresses in the corporate network. Applications such as the Network File Systems (NFS) rely on simple source-IP-address-based filtering to restrict access to corporate users. These are some reasons why the remote access VPN servers are configured with a block of IP addresses from the corporate private network to assign to remote access clients. VPN clients use the virtual IP (VIP) address assigned to them (by the corporate VPN server) to access applications inside the corporate network.

Consider the remote access VPN scenario in figure 2 below.

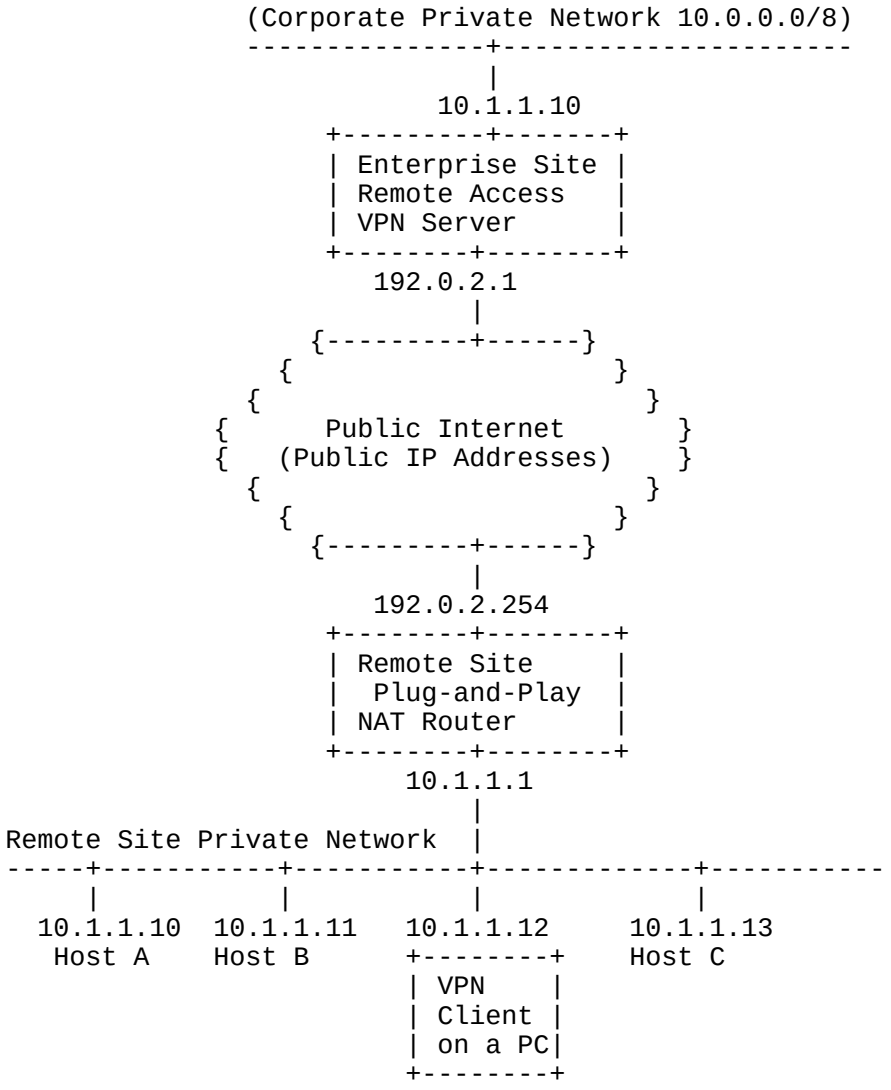

Figure 2. Remote Access VPN with Overlapping Address Space

In the above scenario, say an employee of the corporation is at a remote location and attempts to access the corporate network using the VPN client, the corporate network is laid out using the address pool of 10.0.0.0/8 as defined in RFC 1918, and the VPN server is configured with an address block of 10.1.1.0/24 to assign virtual IP addresses to remote access VPN clients. Now, say the employee at the remote site is attached to a network on the remote site that also happens to be using a network based on the RFC 1918 address space and

coincidentally overlaps the corporate network. In this scenario, it is conventionally problematic for the VPN client to connect to the server(s) and other hosts at the enterprise.

Nevertheless, despite the direct address overlap, the remote access VPN connection between the VPN client at the remote site and the VPN server at the enterprise should remain connected and should be made to work. That is, the NAT device at the remote site should not obstruct the VPN connection traversing it. Additionally, the remote user should be able to connect to any host at the enterprise through the VPN from the remote desktop.

The following subsections describe the operational details of the VPN, anomalies with the address overlap, and recommendations on how best to address the situation.

4.1. Operational Details of Remote Access VPN Network

As mentioned earlier, in the "de facto" Internet address architecture, only the nodes on the public Internet have globally unique IP addresses assigned by the official IP address registries. Many of the networks in the edges use private IP addresses from RFC 1918 and use NAT devices to connect their private networks to the public Internet. Many enterprises adapted the approach of using private IP addresses internally. Employees within the enterprise’s intranet private network are "trusted" and may connect to any of the internal hosts with minimal administrative or policy enforcement overhead. When an employee leaves the enterprise premises, remote access VPN provides the same level of intranet connectivity to the remote user.

The objective of this section is to provide an overview of the operational details of a remote access VPN application so the reader has an appreciation for the problem of remote address space overlap. This is not a tutorial or specification of remote access VPN products, per se.

When an employee at a remote site launches his/her remote access VPN client, the VPN server at the corporate premises demands that the VPN client authenticate itself. When the authentication succeeds, the VPN server assigns a Virtual IP (VIP) address to the client for connecting with the corporate intranet. From this point onwards, while the VPN is active, outgoing IP packets directed to the hosts in the corporate intranet are tunneled through the VPN, in that the VPN server serves as the next-hop and the VPN connection as the next-hop link for these packets. Within the corporate intranet, the

outbound IP packets appear as arriving from the VIP address. So, IP packets from the corporate hosts to the remote user are sent to the remote user's VIP address and the IP packets are tunneled inbound to the remote user's PC through the VPN tunnel.

This works well so long as the subnets in the corporate network do not conflict with subnets at the remote site where the remote user's PC is located. However, when the corporate network is built using RFC 1918 private address space and the remote location where the VPN client is launched is also using an overlapping network from RFC 1918 address space, there can be addressing conflicts. The remote user's PC will have a conflict in accessing nodes on the corporate site and nodes at the remote site bearing the same IP address simultaneously. Consequently, the VPN client may be unable to have full access to the employee's corporate network and the local network at the remote site simultaneously.

In spite of address overlap, remote access VPN clients should be able to successfully establish connections with intranet hosts in the enterprise.

4.2. Anomalies of the Remote Access VPNs

Even though conventional wisdom would suggest that the remote access VPN scenario with overlapping address space would be seriously broken, in practice it still works in many ways. Let us look at some anomalies where there might be a problem and identify solutions through which the remote access VPN application could be made to work even under the problem situations.

4.2.1. Remote Router and DHCP Server Address Conflict

Routing and DHCP service are bootstrap services essential for a remote host to establish a VPN connection. Without DHCP lease, the remote host cannot communicate over the IP network. Without a router to connect to the Internet, the remote host is unable to access past the local subnet to connect to the VPN server at the enterprise. It is essential that neither of these bootstrap services be tampered with at the remote host in order for the VPN connection to stay operational. Typically, a plug-and-play NAT device at the remote site provides both routing and DHCP services from the same IP address.

When there is address overlap between hosts at the corporate intranet and hosts at the remote site, the remote VPN user is often unaware of the address conflict. Address overlap could potentially cause the remote user to lose connectivity to the enterprise entirely or lose connectivity to an arbitrary block of hosts at the enterprise.

Consider, for example, a scenario where the IP address of the DHCP server at the remote site matched the IP address of a host at the enterprise network. When the remote user's PC is ready to renew the lease of the locally assigned IP address, the remote user's VPN client would incorrectly identify the IP packet as being addressed to an enterprise host and tunnel the DHCP renewal packet over the VPN to the remote VPN server. The DHCP renewal requests simply do not reach the DHCP server at the remote site. As a result, the remote PC would eventually lose the lease on the IP address and the VPN connection to the enterprise would be broken.

Consider another scenario where the IP address of the remote user's router overlapped with the IP address of a host in the enterprise network. If the remote user's PC were to send a ping or some type of periodic keep-alive packets to the router (say, to test the liveness of the router), the packets would be intercepted by the VPN client and simply redirected to the VPN tunnel. This type of unintended redirection has the twin effect of hijacking critical packets addressed to the router as well as the host in the enterprise network (bearing the same IP address as the remote router) being bombarded with unintended keep-alive packets. Loss of connectivity to the router can result in the VPN connection being broken.

Clearly, it is not desirable to route traffic directed to the local router or DHCP server to be redirected to the corporate intranet. A VPN client on a remote PC should be configured such that IP packets whose target IP address matches any of the following are disallowed to be redirected over the VPN:

a) IP address of the VPN client's next-hop router, used to access the VPN server.

b) IP address of the DHCP server, providing address lease on the remote host network interface.

Recommendation-2: A VPN client on a remote PC should be configured such that IP packets whose target IP address matches *any* of (a) or (b) are disallowed to be redirected over the VPN:

a) IP address of the VPN client's next-hop router, used to access the VPN server.

b) IP address of the DHCP server, providing address lease on the remote host network interface.

4.2.2. Simultaneous Connectivity Conflict

Ideally speaking, it is not desirable for the corporate intranet to conflict with any of the hosts at the remote site. As a general practice, if simultaneous communication with end hosts at the remote location is important, it is advisable to disallow access to any corporate network resource that overlaps the client's subnet at the remote site. By doing this, the remote user is able to connect to all local hosts simultaneously while the VPN connection is active.

Some VPN clients allow the remote PC to access the corporate network over VPN and all other subnets directly without routing through the VPN. Such a configuration is termed as "Split VPN" configuration. "Split VPN" configuration allows the remote user to run applications requiring communication with hosts at the remote site or the public Internet, as well as hosts at the corporate intranet, unless there is address overlap with the remote subnet. Applications needing access to the hosts at the remote site or the public Internet do not traverse the VPN, and hence are likely to have better performance when compared to traversing the VPN. This can be quite valuable for latency-sensitive applications such as Voice over IP (VoIP) and interactive gaming. If there is no overriding security concern to directly accessing hosts at the remote site or the public Internet, the VPN client on remote PC should be configured in "Split VPN" mode.

If simultaneous connectivity to hosts at the remote site is not important, the VPN client may be configured to direct all communication traffic from the remote user to the VPN. Such a configuration is termed as "Non-Split VPN" configuration. "Non-Split VPN" configuration ensures that all communication from the remote user's PC traverses the VPN link and is routed through the VPN server, with the exception of traffic directed to the router and DHCP server at the remote site. No other communication takes place with hosts at the remote site. Applications needing access to the public Internet also traverse the VPN. If the goal is to maximize the security and reliability of connectivity to the corporate network, the VPN client on remote PC should be configured in "Non-Split VPN" mode. "Non-Split VPN" configuration will minimize the likelihood of access loss to corporate hosts.

Recommendation-3: A VPN client on a remote PC should be configured in "Non-Split VPN" mode if the deployment goal is (a), or in "Split VPN" mode if the deployment goal is (b):

a) If the goal is to maximize the security and reliability of connectivity to the corporate network, the VPN client on the remote PC should be configured in "Non-Split VPN" mode. "Non-Split VPN" mode ensures that the VPN client directs all traffic

from the remote user to the VPN server (at the corporate site), with the exception of traffic directed to the router and DHCP server at the remote site.

b) If there is no overriding security concern to directly accessing hosts at the remote site or the public Internet, the VPN client on the remote PC should be configured in "Split VPN" mode. "Split VPN" mode ensures that only the corporate traffic is directed over the VPN. All other traffic does not have the overhead of traversing the VPN.

4.2.3. VIP Address Conflict

When the VIP address assigned to the VPN client at the remote site is in direct conflict with the IP address of the existing network interface, the VPN client might be unable to establish the VPN connection.

Consider a scenario where the VIP address assigned by the VPN server directly matched the IP address of the networking interface at the remote site. When the VPN client on the remote host attempts to set the VIP address on a virtual adapter (specific to the remote access application), the VIP address configuration will simply fail due to conflict with the IP address of the existing network interface. The configuration failure in turn can result in the remote access VPN tunnel not being established.

Clearly, it is not advisable to have the VIP address overlap the IP address of the remote user’s existing network interface. As a general rule, it is not advisable for the VIP address to overlap any IP address in the remote user’s local subnet, as the VPN client on the remote PC might be forced to respond to ARP requests on the remote site and the VPN client might not process the handling of ARP requests gracefully.

Some VPN vendors offer provisions to detect conflict of VIP addresses with remote site address space and switch between two or more address pools with different subnets so the VIP address assigned is not in conflict with the address space at remote site. Enterprises deploying VPNs that support this type of vendor provisioning are advised to configure the VPN server with a minimum of two distinct IP address pools. However, this is not universally the case.

Alternately, enterprises may deploy two or more VPN servers with different address pools. By doing so, a VPN client that detects conflict of a VIP address with the subnet at the remote site will have the ability to switch to an alternate VPN server that will not conflict.

Recommendation-4: Enterprises deploying remote access VPN solutions are advised to adapt a strategy of (a) or (b) to avoid VIP address conflict with the subnet at the remote site.

a) If the VPN server being deployed has been provisioned to configure two or more address pools, configure the VPN server with a minimum of two distinct IP address pools.

b) Deploy two or more VPN servers with distinct IP address pools. By doing so, a VPN client that detects conflicts of VIP addresses with the subnet at the remote site will have the ability to switch to an alternate VPN server that will not conflict.

4.2.4. Mistaken End Host Identity

When "Split VPN" is configured on the VPN client on a remote PC, there can be a potential security threat due to mistaken identity. Say, a certain service (e.g., SMTP mail service) is configured on exactly the same IP address on both the corporate site and the remote site. The user could unknowingly be using the service on the remote site, thereby violating the integrity and confidentiality of the contents relating to that application. Potentially, remote user mail messages could be hijacked by the ISP’s mail server.

Enterprises deploying remote access VPN servers should allocate global IP addresses for the critical servers the remote VPN clients typically need to access. By doing this, even if most of the private corporate network uses RFC 1918 address space, this will ensure that the remote VPN clients can always access the critical servers regardless of the private address space used at the remote attachment point. This is akin to Recommendation-1 provided in conjunction with multi-level NAT deployments.

Recommendation-5: When "Split VPN" is configured on a VPN client of a remote PC, enterprises deploying remote access VPN servers are advised to assign global IP addresses for the critical servers the remote VPN clients are likely to access.

5. Summary of Recommendations

NAT vendors are advised to refer to the BEHAVE protocol documents ([BEH-UDP], [BEH-TCP], and [BEH-ICMP]) for a comprehensive list of conformance requirements for NAT devices.

The following is a summary of recommendations to support the unconventional NAT topologies identified in this document. The recommendations are deployment-specific and addressed to the personnel responsible for the deployments. These personnel include ISP administrators and enterprise IT administrators.

Recommendation-1: ISPs, using NAT devices to provide connectivity to customers, should assign non-overlapping addresses to servers advertised to customers. One way to do this would be to assign global addresses to advertised servers.

Recommendation-2: A VPN client on a remote PC should be configured such that IP packets whose target IP address matches *any* of (a) or (b) are disallowed to be redirected over the VPN:

a) IP address of the VPN client’s next-hop router, used to access the VPN server.

b) IP address of the DHCP server, providing address lease on the remote host network interface.

Recommendation-3: A VPN client on a remote PC should be configured in "Non-Split VPN" mode if the deployment goal is (a), or in "Split VPN" mode if the deployment goal is (b):

a) If the goal is to maximize the security and reliability of connectivity to the corporate network, the VPN client on the remote PC should be configured in "Non-Split VPN" mode. "Non-Split VPN" mode ensures that the VPN client directs all traffic from the remote user to the VPN server (at the corporate site), with the exception of traffic directed to the router and DHCP server at the remote site.

b) If there is no overriding security concern to directly accessing hosts at the remote site or the public Internet, the VPN client on the remote PC should be configured in "Split VPN" mode. "Split VPN" mode ensures that only the corporate traffic is directed over the VPN. All other traffic does not have the overhead of traversing the VPN.

Recommendation-4: Enterprises deploying remote access VPN solutions are advised to adapt a strategy of (a) or (b) to avoid VIP address conflict with the subnet at the remote site.

a) If the VPN server being deployed has been provisioned to configure two or more address pools, configure the VPN server with a minimum of two distinct IP address pools.

b) Deploy two or more VPN servers with distinct IP address pools. By doing so, a VPN client that detects conflicts of VIP addresses with the subnet at the remote site will have the ability to switch to an alternate VPN server that will not conflict.

Recommendation-5: When "Split VPN" is configured on a VPN client of a remote PC, enterprises deploying remote access VPN servers are advised to assign global IP addresses for the critical servers the remote VPN clients are likely to access.

6. Security Considerations

This document does not inherently create new security issues. Security issues known to DHCP servers and NAT devices are applicable, but not within the scope of this document. Likewise, security issues specific to remote access VPN devices are also applicable to the remote access VPN topology, but not within the scope of this document. The security issues reviewed here only those relevant to the topologies described in sections 2 and 3, specifically as they apply to private address space overlap in the topologies described.

Mistaken end host identity is a security concern present in both topologies discussed. Mistaken end host identity, described in sections 2.2.4 and 3.2.4 for each of the topologies reviewed, essentially points the possibility of application services being hijacked by the wrong application server (e.g., Mail server). Security violation due to mistaken end host identity arises principally due to critical servers being assigned RFC 1918 private addresses. The recommendation suggested for both scenarios is to assign globally unique public IP addresses for the critical servers.

It is also recommended in section 2.1.2 that applications adapt end-to-end authentication and not depend on source IP address for authentication. Doing this will thwart connection hijacking and denial-of-service attacks.

7. Acknowledgements

The authors wish to thank Dan Wing for reviewing the document in detail and making helpful suggestions in reorganizing the document format. The authors also wish to thank Ralph Droms for helping with rewording the text and Recommendation-1 in section 3.2.4 and Cullen Jennings for helping with rewording the text and Recommendation-3 in section 4.2.2.

8. References

8.1. Normative References

[BEH-ICMP] Srisuresh, P., Ford, B., Sivakumar, S., and S. Guha, "NAT Behavioral Requirements for ICMP", BCP 148, RFC 5508, April 2009.

[BEH-TCP] Guha, S., Ed., Biswas, K., Ford, B., Sivakumar, S., and P. Srisuresh, "NAT Behavioral Requirements for TCP", BCP 142, RFC 5382, October 2008.

[BEH-UDP] Audet, F., Ed., and C. Jennings, "Network Address Translation (NAT) Behavioral Requirements for Unicast UDP", BCP 127, RFC 4787, January 2007.

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

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

[RFC1918] Rekhter, Y., Moskowitz, B., Karrenberg, D., de Groot, G., and E. Lear, "Address Allocation for Private Internets", BCP 5, RFC 1918, February 1996.

8.2. Informative References

[DHCP] Droms, R., "Dynamic Host Configuration Protocol", RFC 2131, March 1997.

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

[RFC5735] Cotton, M. and L. Vegoda, "Special Use IPv4 Addresses", BCP 153, RFC 5735, January 2010.

Authors' Addresses

Pyda Srisuresh
EMC Corporation
1161 San Antonio Rd.
Mountain View, CA 94043
U.S.A.

Phone: +1 408 836 4773
EMail: srisuresh@yahoo.com

Bryan Ford
Department of Computer Science
Yale University
51 Prospect St.
New Haven, CT 06511

Phone: +1-203-432-1055
EMail: bryan.ford@yale.edu