

- draft version -
 Behavior and Classification of NAT devices and
 implications for NAT Traversal
 - draft version -

Andreas Müller, Andreas Klenk and Georg Carle
 Chair for Network Architectures and Services
 Technische Universität München
 {mueller, klenk, carle}@net.in.tum.de
 http://www.net.in.tum.de

Please note that this is only a draft version of the paper. The full paper was published in the IEEE Network Special Issue on Implications and Control of Middleboxes in the Internet, October 2008.

Abstract—For a long time, traditional client-server communication was the predominant communication paradigm of the Internet. Network Address Translation devices (NAT) emerged to help with the limited availability of IP addresses and were designed with the hypothesis of asymmetric connection establishment in mind. But with the growing success of peer-to-peer applications, this assumption is no longer true. Consequently NAT Traversal became a field of intensive research and standardization for enabling efficient operation of new services. This article provides a comprehensive overview of Network Address Translation and introduces established NAT Traversal techniques. A new categorization of applications into four NAT Traversal Service Categories helps to determine applicable techniques for NAT Traversal. The Interactive Connectivity Establishment (ICE) framework is categorized and a new framework is introduced that addresses scenarios that are not supported by ICE. Current results from a field test on NAT behavior and the success ratio of NAT Traversal techniques support the feasibility of this classification.

I. INTRODUCTION

When the Internet Protocol (IP) was designed, its growth to today's size was not imaginable. Therefore, it was reasonable to use a fixed 32 bit field to identify a host based on its IP address. This limited address range makes it impossible to assign globally unique IPv4 addresses to the growing number of networked devices. Furthermore, requesting an IP address for every newly added device results in an unacceptable administration overhead. The authors in [1] propose to assign a number of public IP addresses to a designated border router instead of configuring certain hosts with globally routeable addresses. The border router is then responsible for translating IP addresses between the private and the public domains, allowing as many simultaneous connections as public IP addresses were assigned. This allows a host within the local

network to access the Internet even though it has a private IP address. This technique became known as Network Address Translation (NAT). Since the translation of addresses breaks the end-to-end connectivity model of the Internet Protocol, newly developed services following the peer-to-peer (P2P) paradigm such as file sharing, instant messaging and VoIP applications suffer from the existence of NAT. Thus, NAT Traversal is an important problem today. And even in the future, after a possible success of IPv6, companies and home users might still deploy NAT devices in order to hide their topologies from the ISPs. There are two possible approaches to the problem. One direction within the IETF Behave Working Group [2] is to cope with existing NAT implementations and to establish standards for the detection of NAT behavior and for NAT Traversal. On the other hand, the IETF also standardizes behavioral properties for NATs to work in conjunction with IETF protocols (e.g. DCCP, ICMP, SCTP). Enterprise class NATs are among the first to incorporate new features introduced through standardization. However, the large scale deployment of residential gateways with NAT functionality prohibits the change of the NAT and requires the use of protocols that work with existing NATs. This is also the focus of this paper where we treat NATs as black boxes rather than trying to change them.

II. NAT BEHAVIOR

Today, a Network Address Translation device (NAT) is usually used to share a single public IP address among a number of private end systems. The NAT maintains a table, listing all connections between the public and the private domain. For every connection attempt (e.g. a TCP SYN packet) coming from an internal host, the NAT creates a new entry in the list. In NAT terminology this entry is called a binding [3]. Each entry contains the source IP address and the source port. The NAT replaces the source IP address with its public IP address. The source port is replaced using one of the strategies further explained later in this section.

While the concept of NAT has been published as early as 1994 [1], no common approach for NAT has emerged. Current

TABLE I
NAT BEHAVIOR CATEGORIES AND POSSIBLE NAT PROPERTIES

Classification	NAT Property
Port Binding	Port Preservation No Port Preservation Port Overloading Port Multiplexing
NAT Binding	Endpoint Independent Address (Port) Dependent Connection Dependent
Endpoint Filtering	Independent Address Restricted Address and Port Restricted

NAT implementations not only differ from vendor to vendor, but also from model to model, which leads to compatibility issues. If an application works with one particular NAT, this does not imply that it always works in a NATed environment. It therefore is very important to understand and classify existing NAT implementations in order to design applications that can work in combination with current NATs. The classification in this article is mainly derived from STUN [4], while the address binding and mapping behavior follows the terminology used in RFC 4787 [5]. This section only covers topics necessary for the understanding of this article. A detailed discussion and further information (including test results) is given in [6] (for TCP) and [5] (for UDP).

Binding covers "context based packet translation" [7], which describes the strategy the NAT uses to assign a public transport address (combination of IP address and port) to a new state in the NAT. Filtering, or packet discard, shows how the NAT handles (or discards) packets trying to use an existing mapping. Table I shows the different categories and their possible properties. Port Binding describes the strategy a NAT uses for the assignment. With *Port Preservation* the NAT assigns an external port to a new connection, it attempts to preserve the local port number, if possible. *Port Overloading* is problematic and happens rarely. A new connection takes over the binding and the old connection is dropped. *Port Multiplexing* is a very common strategy where ports are demultiplexed based on the destination transport address. Incoming packets can now carry the same destination port and are distinguished by the source transport address.

NAT Binding deals with the reuse of existing bindings. That is, if an internal host closes a connection and establishes a new one from the same source port, NAT Binding describes the assignment strategy for the new connection. As shown in Table I, the NAT Binding is organized into three categories. With *Endpoint Independent* the external port is only dependent on the source transport address of the connection. As long as a host establishes a connection from the same source IP address and port, the mapping does not change. The assignment is dependent on the internal and the external transport address with the *Address (Port) Dependent* strategy. As long as consecutive connections from the same source to the same destination are established, the mapping does not change. As soon as we use

a different destination, the NAT changes the external port. With a *Connection Dependent* binding, the NAT assigns a new port to every connection. We distinguish between NATs that increase the new port number by a specific (and well predictable) delta and NATs that assign random port numbers to the new mappings.

Endpoint Filtering describes how existing mappings can be used by external hosts and how a NAT handles incoming connection attempts not being part of a response. *Independent Filtering* allows inbound connections independent of the source transport address of the packet. As long as the destination transport address of a packet matches an existing state, the packet is forwarded. With *Address Restricted Filtering*, the NAT only forwards packets coming from the same host (matching IP address) the initial packet was sent to. *Address and Port Restricted Filtering* also compares the source port of the inbound packet in addition to Address Restricted Filtering.

III. NAT TRAVERSAL PROBLEM

In order to work properly, the NAT must have access to the protocol headers at Layers 3 and 4 (in case of a NAPT). Additionally, for every incoming packet it is required that the NAT already has a state listed in its table. Otherwise it will not be able to find the related internal host the packet belongs to. According to RFC 3027 [8], the NAT Traversal problem can be separated into three categories that are presented in this section. In addition to the three problems, we identified *Unsupported Protocols* as a new category.

The first problem occurs if a protocol uses *Realm Specific IP Addresses* in its payload. That is, if an application layer protocol such as the Session Initiation Protocol (SIP) uses a transport address from the private realm within its payload signaling where it expects a response. Since regular NATs do not operate above Layer 4, application layer protocols typically fail in such scenarios. A possible solution is the use of an Application Layer Gateway (ALG) that extends the functionality of a NAT for specific protocols. However, an ALG only supports the application layer protocols that are specifically implemented, and may fail when encryption is used.

The second category is *Peer-to-Peer Applications*. The traditional Internet consists of servers located in the public realm and clients that actively establish connections to these servers. This structure is well suited for NATs, because for every connection attempt (e.g. a TCP SYN) coming from an internal client, the NAT can add a mapping to its table. But unlike client-server applications, a P2P connection can be initiated by any of the peers regardless of their location. However, if a peer in the private realm tries to act as a traditional server (e.g. listening for a connection on a socket), the NAT is unaware of incoming connections and drops all packets. A solution could be that the peer located in the private domain always establishes the connection. But what if two peers, both behind a NAT, want to establish a connection to each other? Even if the security policy would allow the connection, it cannot be established.

The third category is a combination of the first two. *Bundled Session Applications*, such as FTP or SIP/SDP, carry realm




Fig. 1. NAT Traversal Service Categories for Applications: a) RNT; b) GSP; c) SPPS; d) SSP

specific IP addresses in their payload in order to establish an additional session. The first session is usually referred to as the control session, while the newly created session is called the data session. The problem here is not only the realm specific IP addresses, but the fact that the data session is often established from the public Internet towards the private host, a direction the NAT does not permit (e.g. active FTP).

Unsupported Protocols are typically newly developed transport protocols such as the Stream Control Transmission Protocol (SCTP) or the Datagram Congestion Control Protocol (DCCP) that cause problems with NATs even if an internal host initiates the connection establishment. This is because current NATs simply do not have built-in support for these protocols. The Unsupported Protocols also cover protocols that cannot work with NATs because their Layer 3 or Layer 4 header is not available for translation. This happens when using encryption protocols such as IPSec.

IV. NAT TRAVERSAL SERVICE CATEGORIES

Instead of classifying the NAT behavior (see classification in STUN [4]), we defined four NAT Traversal Service Categories, each making different assumptions about the purpose of the connection establishment and the infrastructure that is available. Our categorization emphasizes that the applicability of many NAT Traversal techniques depends on the support of a combination of requester, the responder, globally reachable infrastructure nodes and the role of the application. On the one hand, server applications set up a socket and wait for connections (that also applies to P2P applications). On the other hand, client applications such as VoIP clients actively initiate a connection and wait for an answer on a different port (Bundled Session Applications). Other applications only work across NATs if both ends participate in the connection establishment (Unsupported Protocols). Thus, we differentiate between supporting a service and supporting a client. In this article, the client is called the requester, because it actively initiates a connection.

The behavior of the NAT is important, because it allows or prohibits certain NAT Traversal techniques within one Service Category. If only one end implements NAT Traversal support (e.g. by running a standalone framework or by built-in NAT Traversal functionality), NAT Traversal techniques that rely on a collaboration of both ends (e.g. ICE), are not applicable.

Our first category *Requester side NAT Traversal* (RNT) covers scenarios where only the requester side supports NAT Traversal (e.g. the application or the NAT itself). RNT helps applications that actively participate in the connection establishment and still suffer from the existence of NATs. Typical

examples are applications having problems with realm specific IP Addresses in their payload. This applies to protocols using in-band signaling on the application layer, which is related to Bundled Session Applications with asymmetric connection establishment (e.g. VoIP using SIP/SDP).

The second category, *Global Service Provisioning* (GSP), assumes that the host providing the service implements NAT Traversal support, helping to make a service globally accessible. This is done by creating and maintaining a NAT mapping which then accepts multiple connections from previously unknown clients (see Figure 1). This is the main difference to RNT, which only creates a NAT mapping for one particular session (e.g. one call in case of VoIP).

The last two categories assume support at both ends, the service and the requester. On the one side, NAT Traversal is needed in order to make a service behind a NAT globally accessible, while on the other side the support at the requester allows the use of sophisticated techniques through coordinated action. Thus, *Service Provisioning using Pre-Signaling* (SPPS) extends the GSP category by the assumption that both hosts have interoperable frameworks (e.g. ICE [9], NUTSS [10], NATBlaster [11] or NatTrav [12]) running. This allows to select from all available NAT Traversal solutions, which leads to a high success rate of NAT Traversal. In Figure 1, the two hosts use a Rendezvous Point to agree on a NAT Traversal technique. After creating the mapping in step 2, the service may be accessible by any host depending on the chosen NAT Traversal technique and the filtering strategy of the NAT. SPPS supports all kinds of services where a one-to-one connection is sufficient and pre-signaling is available.

The last category, *Secure Service Provisioning* (SSP), is an extension of SPPS and addresses scenarios that require authorization of the remote party before initiating the NAT Traversal process. The hereby established channel must only be accessible by the authorized remote party. This requires additional functionality that enforces this policy and only allows authorized users to access the service. The policy enforcement can be done at the NAT itself, at a data relay or at a firewall. Table II depicts all four Service Categories with popular NAT Traversal techniques and shows the implications for automated NAT Traversal and necessary signaling. First we distinguish between the service and the requester. "Support at the service" means that, for example, a framework needs to be deployed at the same host providing the service. The same applies to the requester. "RP" means that a Rendezvous Point is needed for relaying data back and forth. "Signaling messages" means that some sort of signaling protocol is used for NAT Traversal. Again, we differentiate between signaling

TABLE II
SERVICE CATEGORIES AND THEIR IMPLICATIONS FOR AUTOMATED NAT TRAVERSAL - RP DENOTES RENDEZVOUS POINT

Service Category	NAT Traversal Techniques	Requires support at				Signaling messages				Stream independent
		Service	Requester	RP	NAT	Service	Requester	RP	STUN	
RNT	NAT with ALG UPnP (for bundled session applications)		X		X X		X			X
GSP	UPnP (port forwarding)	X			X	X				X
	Hole Punching - independent filtering	X				X			X	X
	open data relay (e.g. RSIP)	X		X		X		X		X
SPPS	Hole Punching - independent binding	X	X			X	X	X	X	
	UPnP	X	X		X	X	X	X		
	closed/open data relay (e.g. TURN, Skype)	X	X	X		X	X	X		
	Tunneling (e.g. over UDP)	X	X			X	X			X
SSP	Hole Punching - restricted filtering	X	X			X	X	X	X	
	NSIS NATFW NSLP	X	X		X	X	X			
	closed data relay (e.g. TURN)	X	X	X		X	X	X		
	Tunneling (e.g. over secure channel)	X	X			X	X			X

at the service and signaling at the requester. A Rendezvous Point for signaling messages is needed in case of pre-signaling. Finally, "stream independent" describes the requirement for consecutive connections. For example, a port forwarding entry has to be created only once, while Hole Punching [13] requires to send a new Hole Punching packet for every new stream (with restricted filtering).

Altogether, Table II shows the main differences of our Service Categories. RNT deals with Bundled Session Applications that wait on a port after initiating a session (e.g. via a SIP INVITE). GSP only needs support of the service and aims to make a service globally reachable for multiple clients. SPPS and SSP combine these categories and require support at both ends. The requester initiates pre-signaling in order to exchange information about a global endpoint. The service then creates a mapping in the NAT which can be used by the client.

V. APPLICABILITY OF NAT TRAVERSAL TECHNIQUES FOR NAT TRAVERSAL SERVICE CATEGORIES

There are many different techniques for solving the NAT Traversal problem in specific scenarios, but none of them provides a solution that works well with all NATs, applications and network topologies. A former article already explains many of the available protocols for NAT Traversal [14] in general. This section describes the applicability of existing techniques from the applications point of view.

Requester side NAT Traversal (RNT) is required for protocols using in-band signaling (Bundled Session Applications). One common approach is therefore to integrate RNT into these applications (e.g. the VoIP client), in order to establish port bindings on the fly. One possibility is the integration of a UPnP client. Another option is to use Application Layer Gateways (ALG) that are integrated in the NAT, interpreting in-band signaling and establishing mappings accordingly. ALGs are no general solution, because the NAT must implement the necessary logic for each protocol and end to end security prohibits the interpretation of the signaling by the NAT.

Global-Service Provisioning (GSP) depends on NAT Traversal techniques that allow unrestricted access to a public

endpoint. A control protocol can be used to directly establish a port forwarding entry in the mapping tables of the NAT, for instance, with *Universal Plug and Play* (UPnP) [15]. Port forwarding entries created by UPnP are easy to maintain and work independently from NAT behavior. However, UPnP only works if the NAT is in the local network on the path to the other endpoint. Thus, nested NATs are not allowed and path changes break the connectivity.

Hole-Punching is an alternative if UPnP is not applicable and works for NATs with an independent filtering strategy. The mapping has to be refreshed periodically, for instance, by sending keep-alive packets. For NATs other than Full-Cone, Hole-Punching for GSP cannot be used since the source port of the request is unknown in advance.

Service-Provisioning using Pre-Signaling (SSPS) makes no assumption about the accessibility of a created mapping, thus all possible techniques are applicable. Different to GSP, Hole-Punching for SPPS works as long as port prediction is possible. For NATs implementing restricted filtering, pre-signaling helps to create the appropriate mapping since the 5-tuple of the connection is exchanged. Pre-signaling also allows the establishment of an UDP tunnel, allowing the encapsulation of unsupported protocols. SPPS can also use UPnP to establish port forwarding entries for one session.

Secure-Service Provisioning (SSP) is an extension to SPPS that only allows authorized hosts to allocate and to use a mapping. Protocols that authorize requests and assume control over the middlebox, such as *MIDCOM* [16] or the *NAT/Firewall NSIS Signaling Layer Protocol* [17] qualify for SSP. The advantage of NSIS is that it is able to discover and configure multiple middleboxes along the data path, thus supporting complex scenarios with nested NATs and multipath routing. However, if one NAT on the path does not support the protocol, NSIS fails. Using NSIS and MIDCOM for SSP requires restrictive rules that only allow authorized clients to use the mapping, for instance, by opening pinholes for IP 5-tuples. UPnP is not useful for SSP, because it forwards inbound packets without considering the source transport address. Hole-Punching can only be used with SSP if the NAT implements a restricted filtering strategy. All cases discussed

above rely on additional measures to prohibit IP spoofing. The use of secure tunnels impedes IP spoofing and allows secure NAT Traversal, even for unsupported protocols (e.g. IPSec, SCTP, DCCP). SSP can also be achieved by using TURN with authentication, authorization and secure communication (e.g. via TLS).

The *Interactive Connectivity Establishment* (ICE) [9] is under standardization by the IETF and strives to combine several techniques into a framework flexible enough to work with all network topologies. Since ICE requires both peers to have an ICE implementation running, it can be seen as a technique for SPPS or SSP, depending on the accessibility and the security policies of the public endpoint.

The same is true for solutions such as *Traversal Using Relays around NAT* (TURN) [18]. TURN is a promising candidate for SPPS, because it provides a relay with a public transport address allowing the exchange of data packets between a TURN client and a public host.

A. Why Unilateral Solutions Exist

Despite the great flexibility of SPPS and SSP, both categories involve a number of assumptions that are not always satisfied. The most important one is the need for both ends (and sometimes also the infrastructure), to support compatible versions of the NAT Traversal framework. It remains to be seen if the future will bring a sufficiently big deployment of one framework to rely on for arbitrary applications. The chances are better within homogeneous problem domains, like telecommunication, where such frameworks can be integrated with the applications and be distributed in large numbers. For instance, the adoption of ICE is mainly happening within the VoIP/SIP community and focusing on VoIP specific use cases. These drawbacks are the reason why RNT and GSP as unilateral solutions for the NAT Traversal problems exist. It is easier to enhance an infrastructure under one responsibility, than to rely on a solution that requires a global deployment. However, unilateral solutions are limited to the middleboxes in the given domain. They fail to provide solutions to scenarios with nested NATs and depend on the network topology.

VI. COALESCING UNILATERAL AND COOPERATIVE APPROACHES FOR NAT TRAVERSAL

When investigating existing NAT Traversal techniques, we determined that none of them can be used in all scenarios. For example, UPnP only supports globally accessible endpoints, while ICE requires both hosts to run the framework. In [19] we proposed ANTS, a new framework which aims towards providing an Advanced NAT-Traversal Service supporting all four Service Categories. The concept of ANTS is based on the idea of reusing previously obtained knowledge about the topology of the network and the capability of the NAT. A small component of ANTS, the NAT tester, is responsible for gathering this information and will be presented (together with some test results) in the next chapter.

If a user decides that a particular application should be reachable from the public Internet, he registers it at a Session Manager keeping track of all applications requesting NAT

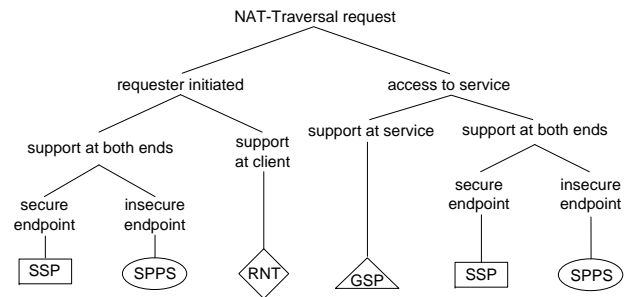


Fig. 2. Decision Tree for ANTS

Traversal support. With the Session Manager, ANTS is able to provide GSP and RNT directly. Whenever an application is added and associated with GSP or RNT, the Session Manager calls the NAT Traversal logic and asks to allocate an appropriate mapping in the NAT. This also requires ANTS to have sufficient knowledge about the applicability of the integrated techniques regarding the Service Categories. For example, UPnP cannot be used for SSP since it violates the idea of an endpoint that is only accessible by authenticated hosts.

Figure 2 shows a decision tree that ANTS uses in order to establish a mapping in the NAT. First, we distinguish between requester initiated NAT Traversal on the one hand and the access to a service on the other hand. We then need to know which ends actually implement ANTS. If both hosts have the framework running, pre-signaling is possible which leads to a wide choice of techniques dependent on the security considerations of the mapping. If only one end supports ANTS, only techniques belonging to GSP or RNT are applicable.

Despite some unsolved issues such as the question how to connect legacy applications to ANTS (e.g. by using a library or a TUN-based approach), the idea of a knowledge based framework seems to be the right answer. Thus once implemented, ANTS will be able to help many existing services by integrating several techniques and making its choice based on knowledge about the NAT and the requirements of the application.

VII. FIELD TEST ON NAT TRAVERSAL

To prove that existing techniques can be adapted to our Service Categories, we implemented a NAT Tester which acts as a cornerstone for our new framework. This section presents the results of a field test investigating 185 NATs in the wild. For a detailed description including all results, please see our website: <http://gex.cs.uni-tuebingen.de>.

The first test queries a public STUN server in order to determine the type of the NAT. Afterwards, the NAT Tester performs the following connection tests and tries to establish a connection to the host behind the NAT: UPnP, Hole Punching and connecting to a data relay (each for both protocols, UDP and TCP).

We then adapted the test results to our work and evaluated the success rates of the individual techniques regarding to our defined Service Categories. Table III shows the categories

TABLE III
RESULTS OF THE FIELD TEST: SUCCESS RATES OF NAT TRAVERSAL
TECHNIQUES DEPENDING ON SERVICE CATEGORIES

S. Cat.	Prot.	Condition	Suc. Rate
RNT	UDP	(UPnP or HP-UDP)	90.27 %
	TCP	(UPnP or HP-TCP)	77.84 %
GSP	UDP	(Full Cone and HP-UDP)	27.03 %
	TCP	(Full Cone and HP-TCP)	17.30 %
	UDP	(UPnP or (Full Cone and HP-UDP))	50.27 %
	TCP	(UPnP or (Full Cone and HP-TCP))	44.32 %
SSP	UDP	(HP-UDP)	88.65 %
	TCP	(HP-TCP)	71.35 %
	TCP	(HP-TCP or HP-UDP)	94.59 %
	UDP	(UPnP or HP-UDP)	90.27 %
	TCP	(UPnP or HP-TCP)	77.84 %
SPPS	UDP	(Restricted NAT and HP-UDP)	48.65 %
	TCP	(Restricted NAT and HP-TCP)	38.38 %

and the conditions that have to be met according to the considerations made above. For example, GSP requires the use of UPnP or Hole Punching support in combination with a Full Cone NAT in order to make a service globally accessible. Therefore, 50.27% of our tested NATs supported a direct connection for UDP and category GSP (44.32% for TCP). In all other cases (the remaining percentages) an external relay has to be used to provide GSP.

For SPPS, which makes no security assumptions, we divided our results into two categories. First we determined the success rates without considering UPnP. With 88.65% of all NATs we were able to establish a direct connection to the host behind the NAT (71.35% for TCP). This rate increased slightly (for TCP to 77.84%) when UPnP was an option. The highest success rate for TCP NAT Traversal (95.14%) was discovered when we also allowed the tunneling of TCP packets through UDP.

SSP only allows authorized hosts to create and to use a mapping. Therefore, a suitable technique for SSP is Hole Punching in combination with a NAT implementing a restricted filtering strategy. This was supported by 48.65% for TCP and 38.38% for UDP.

The success rate for RNT depends on the effort that is made for the specific protocol. For example, if we assume that we can inspect each signaling packet on the Application Layer thoroughly, we could adopt the results from SPPS to RNT. If we would only modify the packets in a way that the internal port is reachable by any client, the success rate of GSP would apply to RNT. Finally, we did not measure the effect of NATs with integrated ALGs in this field test.

VIII. CONCLUSION

With the increasing popularity of peer-to-peer communication, the NAT Traversal problem has become more urgent than ever before. Existing solutions have the drawback of only supporting certain types of NATs and can not be seen as a general solution to the problem. When analyzing the NAT Traversal problem more thoroughly, we discovered that the question who supports the NAT Traversal framework determines which NAT

Traversal techniques are applicable. Therefore, we identified four NAT Traversal Service Categories that differentiate between support by service, client and infrastructure and listed applicable NAT Traversal techniques for each category. Our findings from a field test showed that there are a number of prospective NAT Traversal techniques that enable connectivity for each NAT Traversal Service Category. We emphasized how to build upon this categorization to develop a knowledge based NAT Traversal framework. Future frameworks that aspire to support the typical connectivity scenarios of nowadays applications should support all four service categories.

REFERENCES

- [1] K. Egevang and P. Francis, "The IP Network Address Translator (NAT)," RFC 1631, Internet Engineering Task Force, May 1994.
- [2] IETF, "Behavior Engineering for Hindrance Avoidance (behave)," <http://www.ietf.org>.
- [3] P. Srisuresh and M. Holdrege, "IP Network Address Translator (NAT) Terminology and Considerations," RFC 2663, Internet Engineering Task Force, August 1999.
- [4] J. Rosenberg, J. Weinberger, C. Huitema, and R. Mahy, "STUN: Simple Traversal of User Datagram Protocol (UDP) Through Network Address Translators (NATs)," RFC 3489, Internet Engineering Task Force, March 2003.
- [5] E. F. Audet and C. Jennings, "NAT Behavioral Requirements for Unicast UDP," RFC 4787, Internet Engineering Task Force, January 2007.
- [6] S. Guha and P. Francis, "Characterization and Measurement of TCP Traversal through NATs and Firewalls," In Proceedings of ACM Internet Measurement Conference (IMC), Berkeley, CA, October 2005.
- [7] G. Huston, "Anatomy: A Look Inside Network Address Translators," *The Internet Protocol Journal*, vol. 7, pp. 2–32, 2004.
- [8] M. Holdrege and P. Srisuresh, "Protocol Complications with the IP Network Address Translator," RFC 3027, Internet Engineering Task Force, January 2001.
- [9] J. Rosenberg, "Interactive Connectivity Establishment (ICE): A Protocol for Network Address Translator (NAT) Traversal for Offer/Answer Protocols," Internet-Draft - work in progress, Internet Engineering Task Force, October 2007.
- [10] P. Francis, S. Guha, and Y. Takeda, "NUTSS: A SIP-based Approach to UDP and TCP Network Connectivity," Cornell University, Panasonic Communications, Tech. Rep., 2004.
- [11] A. Biggadike, D. Ferullo, G. Wilson, and A. Perrig, "NATBLASTER: Establishing TCP connections between hosts behind NATs," in *ACM SIGCOMM Asia Workshop, Beijing, China*, 2005.
- [12] J. Eppinger, "TCP Connections for P2P Applications - A Software Approach to Solving the NAT Problem," Carnegie Mellon University, Pittsburgh, PA, Tech. Rep., 2005.
- [13] B. Ford, P. Srisuresh, and D. Kegel, "Peer-to-Peer Communication Across Network Address Translation," Massachusetts Institute of Technology, Tech. Rep., 2005.
- [14] H. Khlifi, J. Gregoire, and J. Phillips, "VoIP and NAT/Firewalls: Issues, Traversal Techniques, and a Real-World Solution," *IEEE Communications Magazine*, July 2006.
- [15] U. Forum, "Internet gateway device (IGD) standardized device control protocol," November 2001.
- [16] P. Srisuresh, J. Kuthan, J. Rosenberg, A. Molitor, and A. Rayhan., "Middlebox communication architecture and framework," RFC 3303, Internet Engineering Task Force, August 2002.
- [17] M. Stiernerling, H. Tschofenig, C. Aoun, and E. Davies, "NAT/Firewall NSIS Signaling Layer Protocol (NSLP)," Internet-Draft, Internet Engineering Task Force, February 2008.
- [18] J. Rosenberg, R. Mahy, and P. Matthews, "Traversal Using Relays around NAT (TURN)," Internet Draft - work in progress, Internet Engineering Task Force, June 2008.
- [19] A. Müller, A. Klenk, and G. Carle, "On the Applicability of knowledge-based NAT-Traversal for future Home Networks," In Proceedings of IFIP Networking 2008, Springer, Singapore, May 2008.