

PacketCable™ 2.0

NAT and Firewall Traversal Technical Report

PKT-TR-NFT-C01-140314

CLOSED

Notice

This PacketCable technical report is the result of a cooperative effort undertaken at the direction of Cable Television Laboratories, Inc. for the benefit of the cable industry and its customers. You may download, copy, distribute, and reference the documents herein only for the purpose of developing products or services in accordance with such documents, and educational use. Except as granted by CableLabs® in a separate written license agreement, no license is granted to modify the documents herein (except via the Engineering Change process), or to use, copy, modify or distribute the documents for any other purpose.

This document may contain references to other documents not owned or controlled by CableLabs. Use and understanding of this document may require access to such other documents. Designing, manufacturing, distributing, using, selling, or servicing products, or providing services, based on this document may require intellectual property licenses from third parties for technology referenced in this document. To the extent this document contains or refers to documents of third parties, you agree to abide by the terms of any licenses associated with such third party documents, including open source licenses, if any.

© Cable Television Laboratories, Inc., 2006-2014

DISCLAIMER

This document is furnished on an "AS IS" basis and neither CableLabs nor its members provides any representation or warranty, express or implied, regarding the accuracy, completeness, noninfringement, or fitness for a particular purpose of this document, or any document referenced herein. Any use or reliance on the information or opinion in this document is at the risk of the user, and CableLabs and its members shall not be liable for any damage or injury incurred by any person arising out of the completeness, accuracy, or utility of any information or opinion contained in the document.

CableLabs reserves the right to revise this document for any reason including, but not limited to, changes in laws, regulations, or standards promulgated by various entities, technology advances, or changes in equipment design, manufacturing techniques, or operating procedures described, or referred to, herein.

This document is not to be construed to suggest that any affiliated company modify or change any of its products or procedures, nor does this document represent a commitment by CableLabs or any of its members to purchase any product whether or not it meets the characteristics described in the document. Unless granted in a separate written agreement from CableLabs, nothing contained herein shall be construed to confer any license or right to any intellectual property. This document is not to be construed as an endorsement of any product or company or as the adoption or promulgation of any guidelines, standards, or recommendations.

Document Status Sheet

Document Control Number: PKT-TR-NFT-C01-140314

Document Title: NAT and Firewall Traversal Technical Report

Revision History: V01 – Released 04/06/06

V02 – Released 10/13/06

V03 – Released 09/25/07

V04 – Released 11/06/07

V05 – Released 04/25/08

C01 – Closed 03/14/14

Date: March 14, 2014

Trademarks:

CableLabs® is a registered trademark of Cable Television Laboratories, Inc. Other CableLabs marks are listed at <http://www.cablelabs.com/certqual/trademarks>. All other marks are the property of their respective owners.

Abstract

Due to the rapid growth of low-cost, Network-enabled devices, high-speed data service subscribers frequently install IP routers and Ethernet hubs in their home so that they can connect multiple devices to the broadband network. Frequently, these devices contain IP Network Address and Port Translation (NA(P)T – simply referred to as NAT) and Firewall capabilities. While NATs and Firewalls provide numerous benefits to the customer, they also create numerous challenges for service providers seeking to offer seamless communication services to applications located behind the customer's NAT/Firewall device.

NATs translate IP addresses between one IP address "realm" and another. This mapping is most commonly done between a private Internet address space and a public Internet address space and is created when an outbound packet is sent (from "inside" to "outside" the NAT/Firewall device). Several challenges arise when NAT and Firewall devices are present in the communication path. For example, these mappings between realms have short lifetimes (for UDP transport) and once they expire, inbound traffic can no longer traverse the NAT. While NAT and Firewall capabilities are two distinct functions, they are often implemented together, which compounds the issues. Devices that implement NAT and also exhibit Firewall characteristics block traffic coming across the NAT (from outside to inside the NAT/Firewall device) based on certain packet filtering rules.

This Technical Report presents the challenges that NATs and Firewalls create for the service provider, outlines some technical requirements in scope of the PacketCable architecture to address these challenges, and describes how these challenges are met.

Contents

1	INTRODUCTION.....	1
1.1	PACKETCABLE OVERVIEW	1
2	REFERENCES	2
2.1	NORMATIVE REFERENCES	2
2.2	INFORMATIVE REFERENCES.....	2
2.3	REFERENCE ACQUISITION.....	3
3	TERMS AND DEFINITIONS	4
4	ABBREVIATIONS AND ACRONYMS.....	5
5	PACKETCABLE NAT REQUIREMENTS AND SCOPE.....	6
5.1	REQUIREMENTS.....	6
5.2	SCOPE	6
5.3	LIMITATIONS.....	7
6	NAT BACKGROUND	8
6.1	TYPES OF NAT AND FIREWALL DEVICES	8
6.1.1	<i>Types of NATs</i>	8
6.1.2	<i>Filtering Behavior</i>	9
6.2	NA(P)T AND FIREWALL TRAVERSAL CONSIDERATIONS	9
7	PACKETCABLE NAT ARCHITECTURE	11
7.1	RELATIONSHIP TO PACKETCABLE MULTIMEDIA	12
7.2	RELATIONSHIP TO 3GPP IMS RELEASE 7	12
7.3	RELATIONSHIP TO PACKETCABLE E-MTAS	13
7.4	PROVISIONING AND MANAGEMENT	13
7.4.1	<i>Provisioning</i>	13
7.4.2	<i>Management</i>	13
8	ARCHITECTURE DESCRIPTION.....	14
8.1	FUNCTIONAL COMPONENTS	14
8.1.1	<i>P-CSCF</i>	14
8.1.2	<i>S-CSCF</i>	14
8.1.3	<i>UE</i>	14
8.1.4	<i>STUN Servers</i>	15
8.1.5	<i>TURN Server</i>	15
8.2	PROTOCOL INTERFACES AND REFERENCE POINTS	15
8.2.1	<i>NAT Traversal for Media</i>	15
8.2.2	<i>Gm (NAT Traversal for Signaling)</i>	17
APPENDIX I	ACKNOWLEDGEMENTS	20

Figures

FIGURE 1 - TYPES OF NATS (ADDRESS MAPPING).....	8
FIGURE 2 - NAT AND FW TRAVERSAL REFERENCE POINTS	11
FIGURE 3 - ABSTRACT REFERENCE DIAGRAM	16
FIGURE 4 - NAT TRAVERSAL FOR SIP SIGNALING	17

Tables

TABLE 1 - TYPES OF NATS (ADDRESS MAPPING).....	9
TABLE 2 - TYPES OF FILTERING BEHAVIOR.....	9
TABLE 3 - PACKETCABLE NAT REFERENCE POINTS	12

1 INTRODUCTION

This technical report provides an overview of how the PacketCable architecture and associated UEs support the traversal of Network Address and Port Translation (NA(P)T – simply referred to as NAT throughout this document, see Section 6 for further details) and Firewall (FW) devices for media and signaling flows as well as for provisioning and management. To aid the reader in understanding the PacketCable NAT traversal approach and methodologies, the high-level goals and specific logical components and interfaces defined are discussed in this technical report. CableLabs has issued this technical report and associated specifications to facilitate design and field-testing leading to the manufacture and interoperability of conforming hardware and software by multiple vendors.

1.1 PacketCable Overview

PacketCable is a CableLabs specification effort designed to support the convergence of voice, video, data, and mobility technologies. There are tens of millions of cable broadband customers, and the capability of the network to provide innovative services beyond high-speed Internet access is ever-increasing. In particular, real-time communication services based on the IP protocols, such as Voice over Internet Protocol (VoIP), are rapidly evolving and consumers are embracing a wide-range of client devices and media types. It is expected that new technologies, such as Video over IP communications and the ability to display voice and video mail message notifications on a TV-set, will change the way communication and entertainment services are offered. These cutting edge technologies will present exciting new opportunities for cable operators to offer high-value services to consumers in a cost-effective manner; see [ARCH-FRM TR] for further discussion about PacketCable.

2 REFERENCES

2.1 Normative References

There are no normative references in this document.

2.2 Informative References

This Technical Report uses the following informative references.

- [ARCH-FRM TR] PacketCable Architecture Framework Technical Report, PKT-TR-ARCH-FRM-C01-140314, March 14, 2014, Cable Television Laboratories, Inc.
- [ID ICE] IETF Internet-Draft, draft-ietf-mmusic-ice-19, Interactive Connectivity Establishment (ICE): A Methodology for Network Address Translator (NAT) Traversal for Offer/Answer Protocols, October 29, 2007, work in progress.
- [ID OUTBOUND] IETF Internet-Draft, draft-ietf-sip-outbound-12, Managing Client Initiated Connections in the Session Initiation Protocol (SIP), March 21, 2008, work in progress.
- [ID STUN] IETF Internet-Draft, draft-ietf-behave-rfc3489bis-15, Session Traversal Utilities for (NAT) (STUN), February 23, 2008, work in progress.
- [ID TURN] IETF Internet-Draft, draft-ietf-behave-turn-07, Traversal Using Relays around NAT (TURN): Relay Extensions to Session Traversal Utilities for NAT (STUN), February 25, 2008, work in progress.
- [PCMM] PacketCable Multimedia Specification, PKT-SP-MM-I06-110629, June 29, 2011, Cable Television Laboratories, Inc.
- [PKT 24.229] PacketCable SIP and SDP Stage 3 Specification 3GPP TS 24.229, PKT-SP-24.229-C01-140314, March 14, 2014, Cable Television Laboratories, Inc.
- [RFC 4787] IETF RFC 4787, Network Address Translation (NAT) Behavioral Requirements for Unicast UDP, January, 2007.
- [RFC 1918] IETF RFC 1918, Address Allocation for Private Internets, 1996.
- [RFC 2327] IETF RFC 2327, SDP: Session Description Protocol, April 1998.
- [RFC 2663] IETF RFC 2663, IP Network Address Translator (NAT) Terminology and Considerations, August 1999.
- [RFC 3327] IETF RFC 3327, Session Initiation Protocol (SIP) Extension Header Field for Registering Non-Adjacent Contacts, 2002.
- [RFC 3550] IETF RFC 3550, RTP: A Transport Protocol for Real-Time Applications, July 2003.

- [RFC 3581] IETF RFC 3581, An Extension to the Session Initiation Protocol (SIP) for Symmetric Response Routing, August 2003.
- [RFC 3605] IETF RFC 3605, Real Time Control Protocol (RTCP) attribute in Session Description Protocol (SDP), October 2003.
- [TS 23.228] 3rd Generation Partnership Project; Technical Specification Group Services and System Aspects; IP Multimedia Subsystem (IMS); Stage 2 (Release 7); December 2007.

2.3 Reference Acquisition

CableLabs Specifications:

- Cable Television Laboratories, Inc., 858 Coal Creek Circle, Louisville, CO 80027; Phone +1-303-661-9100; Fax +1-303-661-9199; Internet: <http://www.cablelabs.com>
- Internet Engineering Task Force (IETF); Internet: <http://www.ietf.org/>
Note: Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time.
The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/1id-abstracts.txt>
- Third Generation Partnership Project (3GPP), Internet: <http://www.3gpp.org>

3 TERMS AND DEFINITIONS

Various terms related to NAT are defined in [RFC 2663], please refer to it for the general NAT terms and definitions not present below. In addition, this document uses the following terms:

- ALG** An Application Layer Gateway within a NAT device that attempts to sniff application signaling and modify application addresses appropriately in order to take care of changes caused by the NAT.
- NAT, NAPT** NATs perform IP address translation, typically interconnecting private and public address domains. NAPT devices also translate ports in order to save IP addresses. In this document, the term NAT also refers to NAPT devices.

4 ABBREVIATIONS AND ACRONYMS

This document uses the following abbreviations:

ALG	Application-Layer Gateway
CSCF	Call Session Control Function
CPE	Customer Premises Equipment
FW	Firewall
ICE	Interactive Connectivity Establishment
IETF	Internet Engineering Task Force
IP	Internet Protocol
NAPT	IP Network Address and Port Translation
NAT	IP Network Address Translation
PAM	PacketCable Application Manager
P-CSCF	Proxy Call Session Control Function
RTCP	Real Time Control Protocol
RTP	Real Time Protocol
S-CSCF	Serving Call Session Control Function
SDP	Session Description Protocol
SIP	Session Initiation Protocol
STUN	Simple Traversal of UDP Through NAT
TCP	Transmission Control Protocol
TURN	Traversal Using Relay NAT
UDP	User Datagram Protocol
UDPTL	UDP Transport Layer
URI	Uniform Resource Identifier
VPN	Virtual Private Network

5 PACKETCABLE NAT REQUIREMENTS AND SCOPE

The objective of this document is to provide an architecture definition for a UE to obtain access to the PacketCable network in the presence of one or more NAT device(s). In particular, it outlines a comprehensive set of mechanisms for a UE to maintain both signaling and media bindings to ensure both media and signaling traffic destined for the UE is able to traverse the NAT as well as to allow the UE to be provisioned and managed when located behind a NAT.

In addition, this architecture provides support for failover of signaling paths through a backup Proxy-CSCF (P-CSCF). While, the procedures for mid-session failover are not currently defined, guidance is provided in both [ID OUTBOUND] and [PKT 24.229] on how mid-session failover may be accomplished.

Finally, this architecture recognizes that UEs will have to interwork with non-PacketCable devices that do not support the required NAT traversal mechanisms and provide for these cases.

The following section captures the set of architecture requirements necessary to achieve the objective.

5.1 Requirements

The following list contains requirements that a general-purpose NAT Traversal solution should satisfy to support the services envisioned for PacketCable:

- Support multiple UEs (on one or more devices) behind a single NAT;
- No requirements will be imposed on the NAT devices, nor require the network to be aware of the presence of a NAT;
- Support both inbound and outbound requests to and from UEs through one or more NAT device(s);
- Maintain bindings to multiple P-CSCFs to provide reliable inbound message delivery in the face of a P-CSCF failure;
- Support the traversal of NATs between the UE and network (home NAT, visited network NAT);
- Be Application independent: the solution should not employ application-specific mechanisms that could not be used by other non-SIP based solutions. The solutions actual use may require application support;
- Avoid unnecessarily long media paths due to media pinning;
- Re-establish communications in failure situations (e.g., the NAT device re-boots and NAT bindings are lost).

5.2 Scope

The scope of the PacketCable NAT traversal solution is limited to NATs within the access network. In the case of cable access, this implies NATs that are between the UE and CMTS. Note that PacketCable E-MTAs are out of scope for this document. However, some additional requirements may be placed on PacketCable E-MTAs in order to ensure they interoperate with PacketCable UEs following the NAT traversal procedures described in this document and other specifications.

5.3 Limitations

The NAT traversal solution defined by PacketCable does not address NAT and Firewall traversal of ITU-T T.38 fax media streams over User Datagram Protocol Transport Layer (UDPTL). Rather, T.38 fax over RTP is necessary to provide NAT and Firewall traversal capability.

Operator call-back for emergency calls (such as 911) made without first registering is not currently supported.

6 NAT BACKGROUND

NATs translate addresses between one IP address realm and another. This mapping is most commonly done between a private Internet address space using addresses set aside for that purpose in [RFC 1918] and a public Internet address space. This mapping is commonly referred to as a NAT binding, as the NAT has bound together the tuple of PrivateIP:PrivatePort to PublicIP:PublicPort to allow the subsequent response packets from the external endpoint to be forwarded to the proper internal host. The term NAT in this document also refers to Network Address and Port Translation (NA(P)T) devices, which also translate port addresses in order to reduce the number of public addresses used on the public address side of the NAT (see [RFC 2663] section 4 for more details).

In addition to address translation, NAT devices also exhibit firewall characteristics. In other words, they block traffic coming across the NAT (from outside to inside the NAT/FW device) based on certain filtering rules.

6.1 Types of NAT and Firewall Devices

The following sub-sections use the definitions from the IETF BEHAVE working group as defined in [RFC 4787].

6.1.1 Types of NATs

The definitions from [RFC 4787] are included here for convenience:

Endpoint Independent Mapping: The NAT reuses the port mapping for subsequent packets sent from the same internal IP address and port to any external IP address and port.

Address Dependent Mapping: The NAT reuses the port mapping for subsequent packets sent from the same internal IP address and port to the same external IP address, regardless of the external port. If the packets are sent to a different external IP address, the mapping will be different.

Address and Port Dependent Mapping: The NAT reuses the port mapping for subsequent packets sent from the same internal IP address and port to the same external address and port. If packets are sent to a different IP address and/or port, then a different mapping will be used.

This address mapping behavior is described in Table 1 by making use of the illustration in Figure 1.




Figure 1 - Types of NATs (Address Mapping)

In Figure 1, address X:x inside the NAT is translated to address X1:x1 when communicating with Y1:y1 outside the NAT. The same address X:x translates to X2:x2 when communicating with Y2:y2.

Table 1 - Types of NATs (Address Mapping)

Type of NAT	Mapping Description
Endpoint Independent Mapping	X1:x1 always equals X2:x2 for all values of Y2:y2
Address Dependent mapping	X1:x1 equals X2:x2 only if Y1 equals Y2
Address and Port Dependent Mapping	X1:x1 equals X2:x2 only if Y1:y1 equals Y2:y2

Note that, for small NATs (e.g., residential CPE NATs), a single IP address (usually from the public space) is normally assigned as the external IP address (i.e., $X1 = X2$). However, larger NATs will assign the external IP address from a pool of available IP addresses.

6.1.2 Filtering Behavior

Filtering behavior in [RFC 4787] is described in terms of similar categories:

Endpoint Independent Filtering: sending packets from the internal side of the NAT to any external IP address is sufficient to allow any packets back to the internal endpoint.

Address Dependent Filtering: in order to receive packets from a specific external endpoint, it is necessary for the internal endpoint to send packets first to that specific external endpoint's IP address.

Address and Port Dependent Filtering: receiving packets from a specific external endpoint, it is necessary for the internal endpoint to send packets first to that external endpoint's IP address and port.

Table 2 describes this filtering behavior, which can be described in terms of the examples shown in Figure 1.

Table 2 - Types of Filtering Behavior

Type of NAT	Filtering Example
Endpoint Independent Filtering	Packets sent from X:x to Y1:y1 will enable packets from Y1:y1 or Y2:y2 to be received.
Address Dependent Filtering	Packets sent from X:x to Y1:y1 will enable packets to be received from Y1:z for any port z but will not allow packets to be received from any other IP address.
Address and Port Dependent Filtering	Packets sent from X:x to Y1:y1 will only allow packets to be sent from Y1:y1 to X:x.

6.2 NA(P)T and Firewall Traversal Considerations

While NAT devices provide a rather simple solution to IP address exhaustion, the consequence is that these devices break many existing applications. In particular, applications and real-time communication protocols such as SIP, which depend on the exchange of addressing information within the protocol itself (sometimes in SIP header lines, or more generally, inside the SDP message body of some SIP messages). If the address within the protocol is unreachable, the recipient of the message will be unable to successfully respond, resulting in failed sessions. Given the growing problem NAT devices present to communications protocols, several solutions have been used in the past and proposed for the future.

The following list provides a snapshot of some of the more prevalent approaches and their drawbacks:

- Application-Layer Gateway (ALG): One solution is for the NAT device to contain an Application-Layer Gateway, which looks inside the protocol messages and modifies them based on the NAT bindings it has created. However, this requires constant update of the ALG as protocols evolve so that the expected operation is preserved. In addition, this approach fails when the protocol includes an integrity check or is encrypted.
- Another approach is that proposed by the IETF MIDCOM group. In this approach, a signaling element directly controls the NAT device to open firewall pinholes for the media and to obtain the NAT binding information needed to update any IP addressing information (e.g., SDP) within the protocol. However, this requires NAT device support and requires that the signaling device is somehow able to determine which NAT device to control.
- Other proposed solutions include directly inserting a media relay (an additional address translator) in the path or tunneling through the NAT device using VPN technology. Each of these methods has its pros and cons, but the major disadvantage is that it forces the media along the same path as the signaling. This may not result in the most optimum media routing in certain circumstances.
- The present short-term solution defined within the IETF is to make use of the ICE methodology [ID ICE], with STUN [ID STUN] and TURN [ID TURN] for media and outbound [ID OUTBOUND] for signaling. ICE allows an endpoint to discover, advertise and find the best address for communicating using the mechanisms described in the ICE IETF Internet-Draft [ID ICE], while outbound allows the endpoint to actively manage its connectivity to the SIP network by creating and maintaining flows to its provisioned proxy(s) as described in the outbound IETF Internet-Draft [ID OUTBOUND].
- Other solutions are being examined within the IETF BEHAVE working group. However, this involves longer term modification of NAT behavior in order to solve these problems.

In addition to the problem of protocols that embed IP address information within their payload, NAT devices cause other problems such as:

- NAT devices have timeouts associated with NAT bindings and firewall pinholes. For UDP-based transports, these bindings tend to have rather short lifetimes, and if traffic is absent for a period of time, the bindings become invalid and pinholes close. To avoid this situation, mechanisms must be in place to maintain the bindings/pinholes for both media and signaling.
- The IETF RTP protocol [RFC 3550] specifies that, for UDP and similar protocols, RTP should use an even destination port number, and the corresponding RTCP stream should use the next higher (odd) destination port number. However, NAT translations will make this practice invalid, since they typically do not maintain these port relationships across the NAT.
- Another issue of concern is the routing of inbound signaling to UE devices. This signaling must be routed through the NAT over a connection for which there is an existing NAT binding.

Given the requirements provided in Section 5.1, the current IETF solution using [ID ICE](for media) and [ID OUTBOUND](for signaling) was chosen for PacketCable. Not only do these solutions satisfy the stated requirements, but they have the benefit of growing industry support.

7 PACKETCABLE NAT ARCHITECTURE

Figure 2 is a reference diagram showing key components and interfaces related to NAT/FW traversal in the PacketCable architecture. The Cable Modem is not shown since it does not contain any functions specifically related to or impacted by NAT traversal.

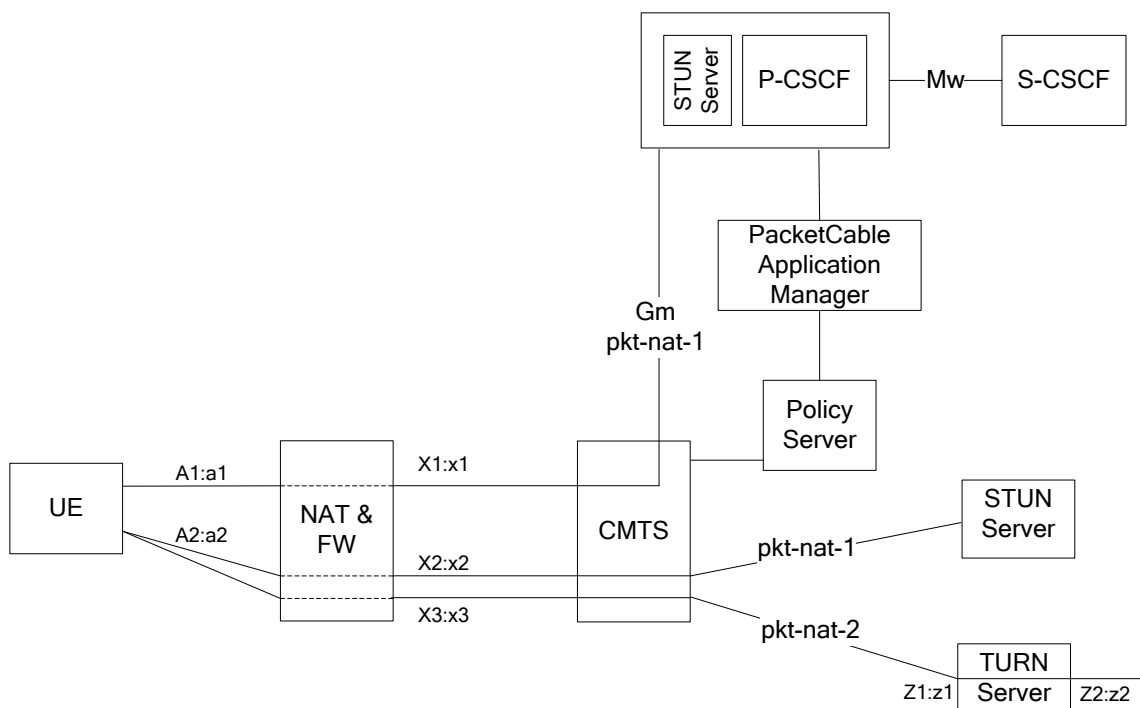


Figure 2 - NAT and FW Traversal Reference Points

The network addresses shown in the diagram are of the form "IP_address:port". Note that:

- For UE devices, the same IP address is normally used for media and signaling: A1 = A2.
- For residential CPE NAT/FW devices (e.g., deployed in residential environment), normally X1=X2=X3. For larger NATs, with large numbers of client devices, the external IP address may be selected from a pool of IP addresses.

It is understood that there are other network topologies with NAT devices that may be encountered, such as a network NAT. These other networks are currently outside the scope of PacketCable and thus are not addressed within this document. Some network topologies (such as points of interconnection) may be deemed in scope and covered in other documents.

Table 3 - PacketCable NAT Reference Points

Reference Points	PacketCable Network Elements	Reference Point Description
Gm	UE -P-CSCF	Allows the UE to communicate with the P-CSCF for registration and session control. This reference point is SIP-based and is defined in [PKT 24.229].
Mw	CSCF – CSCF	Allows the communication and forwarding of signaling messaging among CSCFs in support of registration and session control. This reference point is SIP-based and is defined in [PKT 24.229].
pkt-nat-1	UE – STUN Server UE – P-CSCF	A STUN-based interface defined by [ID STUN] and used by the UE to determine the IP address assigned to its serving NAT or to keep NAT bindings active to a P-CSCF via a keepalive mechanism.
pkt-nat-2	UE – TURN Server	A STUN-based interface defined by [ID TURN] and used by the UE to request TURN server resources for relaying media packets to/from the requesting UE.

7.1 Relationship to PacketCable Multimedia

Functional elements from the PacketCable Multimedia architecture [PCMM] are also shown in Figure 2. This includes the Application Manager (AM) as well as the Policy Server (PS). Although this document does not impact PacketCable Multimedia operation, there is a requirement for the PacketCable Application Manager (PAM) to supply appropriate packet classifier definitions for media flows to the CMTS.

The PAM builds packet classifiers for media flows using the default IP address and port as advertised in the "m=" and "c=" lines of the SDP. When a UE invokes the ICE procedure and gathers candidate addresses, it is required to use the relayed transport address and port (TURN Server assigned) as the default address in the "m=" and "c=" lines of the SDP.

When a TURN server is used, the default address advertised in the SDP as illustrated in Figure 2 is "Z2:z2". However, this is of no value in defining a packet classifier, so some unique filter must be available in the SDP for describing the flows from addresses "X3:x3" and "Z1:z1". ICE [ID ICE] and TURN [ID TURN] have provided the capability for the UE to learn and advertise X3:x3 in the SDP as part of the candidate address encoding as defined by [ID ICE].

7.2 Relationship to 3GPP IMS Release 7

PacketCable is based on Release 7 of the IP Multimedia Subsystem (IMS) as defined by the 3rd Generation Partnership Project (3GPP). 3GPP is a collaboration agreement between various standards bodies. The scope of 3GPP is to produce Technical Specifications and Technical Reports for GSM and 3rd Generation (3G) Mobile System networks.

The current IMS Release 7 architecture provides support for multiple NAT and Firewall traversal techniques; a network-managed and a UE-managed solution. The network-managed solution utilizes an application level gateway within the P-CSCF for managing the SIP signaling and a media relay (transition gateway) for media traversal. The UE-managed solution is based on the IETF procedures and capabilities described within this Technical Report. The NAT traversal techniques defined by IMS are confined to two Specifications, 3GPP TS 23.228 and PacketCable specification 24.229.

The Technical Specification 23.228 [TS 23.228] provides the requirements for supporting both NAT and Firewall traversal techniques from the architecture perspective.

The PacketCable Specification 24.229 [PKT 24.229] documents the procedures related to both signaling and media traversal of NAT and Firewall devices. The procedures define support for the Application Level Gateway for signaling and the Transition Gateway for media in addition to the IETF Outbound Internet-Draft [ID OUTBOUND] for signaling and the IETF Interactive Connectivity Establishment Internet-Draft [ID ICE] for media. Section 8 of this document provides a high-level overview of the associated interfaces and network element roles as they relate to UE-managed NAT traversal solution.

7.3 Relationship to PacketCable E-MTAs

This document does not apply to PacketCable E-MTAs. However, the proposed solution does imply a requirement on PacketCable E-MTAs to be able to accept an empty (no payload) RTP packet with a payload type of 20 as a keepalive for maintaining NAT bindings for media.

7.4 Provisioning and Management

In addition to supporting the traversal of NAT devices for signaling and media, it is also imperative that the UE be able to be provisioned and managed when behind a NAT. Provisioning refers to the processes involved in the initialization of user attributes and resources on user equipment and network components to provide services to a User. Management refers to the protocols, methodologies, and interfaces that enable monitoring, regulating, and ensuring availability of offered services in a Service Provider Network.

7.4.1 Provisioning

UE provisioning is currently out of scope and thus no procedures for how to provision a UE behind a NAT have been developed.

7.4.2 Management

UE management is currently out of scope and thus no procedures for how to manage a UE behind a NAT have been developed.

8 ARCHITECTURE DESCRIPTION

The previous section of this technical report described a set of logical network entities grouped by specific service functions (NAT), as well as a set of interfaces that support the information flows exchanged between the functional groups and network entities. This section provides a more detailed discussion of those logical elements and the associated interfaces to the PacketCable architecture. It also provides an overview of other topics related to the NAT architecture that are not documented elsewhere.

8.1 Functional Components

In this section, additional detail is provided on each of the functional elements in the PacketCable architecture and their role in NAT and Firewall traversal.

8.1.1 P-CSCF

The P-CSCF's primary role in NAT traversal is to ensure that requests and responses occur across a flow for which there is an existing NAT binding. When a registration occurs, the P-CSCF stores a flow identifier token in the SIP Path header, so that for incoming requests that contain a URI with that flow identifier token, it can identify which flow to use.

The P-CSCF also supports the rport extension [RFC 3581] to ensure that all responses to the UE, including those from mid-dialog requests, are sent to the same source IP Address and Port over which the request was received to ensure their ability to traverse the NAT.

The P-CSCF also acts as a STUN server to allow the UE to use STUN for keepalives and to check for changes in NAT bindings (e.g., to check for NAT re-boots that would result in removal of the flow).

8.1.2 S-CSCF

The S-CSCF stores the instance-id associated with the public user identity, as well as the reg-id and path header, and includes these as part of the contact information.

8.1.3 UE

The UE is responsible for managing the overall NAT discovery process and for invoking the various protocol mechanisms to implement the NAT traversal approach. Depending on the UE type (standalone, embedded, etc.), the following protocols or mechanisms are necessary:

- Outbound for signaling
- STUN Client and Server for maintaining NAT bindings (signaling and media), connectivity checks (ICE), and candidate address gathering (ICE)
- TURN client for media relay
- ICE Methodology for Media

Before the UE can receive inbound session requests (or responses to outbound session requests), it will need to invoke the procedures defined in [ID OUTBOUND] during the registration process to create a flow to its assigned P-CSCF. Once flows have been created, the UE can then initiate sessions and receive session requests through the NAT.

During the session establishment process, the UE initiates the ICE methodology to gather, advertise, and test candidate addresses.

The UE also makes use of the rport extension parameter of the Via header as defined in [RFC 3581] for symmetric response routing of SIP messages.

8.1.4 STUN Servers

STUN Servers receive STUN binding requests and provide a response containing the source IP address and Port contained in the IP header of the STUN binding request. Two STUN servers are shown in Figure 2:

- The STUN server shown as a functional component within the P-CSCF is used by the UE in order to maintain the NAT bindings for signaling. These STUN messages may also act as a keepalives, allowing the UE to determine P-CSCF availability.
- The external STUN server in the lower right-hand corner of the diagram is used as part of the ICE methodology [ID ICE] to determine one of several possible candidate media addresses using STUN [ID STUN]. For the example media stream shown with source IP address "A2:a2", the UE obtains translated address "X3:x3" via the STUN server.

8.1.5 TURN Server

In addition to the STUN servers, the architecture also contains a TURN server. This may be required if the NAT device does not use Endpoint Independent Mapping (see Section 6.1.1). When used to transfer media, the TURN server acts as a media relay. The UE sends packets from address "A2:a2" to the TURN server address "Z1:z1". The source address of these packets is first translated by the NAT to "X3:x3" and then relayed by the TURN server so that the source address becomes "Z2:z2". The TURN server also relays media packets in the opposite direction, i.e., packets sent to "Z2:z2" will be sent to "X3:x3" and then via the NAT to "A1:a2". Note that the TURN server provides Address Independent Filtering (see Section 6.1.2), thus retaining some of the filtering characteristics of a NAT, but does not maintain port restrictions, i.e., if traffic is sent to an IP address, it is allowed from that address regardless of what port from which it comes.

Note that only a single example media stream is illustrated in Figure 2. In fact there may be multiple media streams and each media stream may have an RTP stream as well as an RTCP control channel using RTCP. NAT translations and the corresponding mechanisms for communicating are relevant to both.

8.2 Protocol Interfaces and Reference Points

This technical report has identified several interfaces, or reference points, in the PacketCable NAT Traversal architecture. An overview of the various protocol interfaces is provided within this section.

8.2.1 NAT Traversal for Media

UEs communicate via a network that provides signaling components as well as STUN and TURN servers to enable NAT traversal. UE support includes the STUN and TURN protocols as well as the ICE methodology. The associated requirements are provided in the following sub-sections. The diagram in Figure 3 illustrates an abstract view of the architecture.




Figure 3 - Abstract Reference Diagram

8.2.1.1 ICE

The ICE methodology [ID ICE] consists of the following steps:

- Gathering candidate addresses for media communications.
- Advertising the candidate addresses along with the active transport address in the m/c lines of the SDP.
- Doing connectivity checks on the candidate addresses in order to select a suitable address for communications.
- Depending on the results of the connectivity checks, one of the candidate addresses may be promoted to become the active transport address.
- Maintaining the bindings for media.

If one of the endpoints does not support ICE, that endpoint will ignore any of the "a=candidate" attributes and will not provide any of these attributes. In that case, the default value in the m/c lines will be used and connectivity checks will not be done.

8.2.1.2 PKT-NAT-1

Session Traversal Utilities for (NAT) (STUN) provides a toolkit of functions. These functions allow entities behind a NAT to learn the address bindings allocated by the NAT, to keep those bindings open, and communicate with other STUN-aware devices to validate connectivity. STUN requires no changes to NATs, and works with an arbitrary number of NATs in tandem between the application entity and the public Internet.

STUN is a simple client-server protocol. A client sends a request to a server, and the server returns a response. There is a single type of request - Binding Requests, sent over UDP. Binding requests are used to determine the bindings allocated by NATs. Upon receipt of a binding request, the server examines the source IP address and port of the request, and copies them into a response that is sent back to the client.

Once the client learns the WAN address of its local NAT, it will advertise this learned address as a candidate address in the SDP for remote endpoints to try and reach through the ICE process.

8.2.1.3 PKT-NAT-2

This interface defines an extension to STUN, called TURN, that allows a client to request a relayed transport address on a TURN server. To accomplish that, this usage defines a handful of new STUN requests and indications. The Allocate request is the most fundamental component of this usage. It is used to provide the client with a transport address that is relayed through the TURN server. A transport address that relays through an intermediary is called a relayed transport address.

A TURN client first discovers the address of a TURN server based on configuration. This can be preconfigured as an IP address, domain name, or FQDN. This will allow for different TURN servers for UDP and TCP. Once a TURN server is discovered, the client sends a STUN Allocate request to the TURN server. STUN provides a mechanism for mutual authentication and integrity checks for both requests and responses, based on a shared secret. Assuming the request is authenticated and has not been tampered with, the TURN server allocates a transport address to the TURN client, called the relayed transport address, and returns it in the response to the Allocate Request. Normally, the relayed transport address will be on one of the interfaces on the TURN server itself. However, it is also allowed for the TURN server to be behind a NAT, in which case the allocated transport address may correspond to the NAT, which is then mapped to the private address of the TURN server. Proper operation of the TURN server will require it to have many bindings established in the NAT ahead of time; the means for doing so are outside the scope of this technical report.

Once the client is assigned a relayed transport address, it will advertise this address in the "c=" line of the SDP. As a result the relayed transport address will be the first-used address until the other candidate address is found to be a better path through the ICE process.

8.2.2 Gm (NAT Traversal for Signaling)

This section provides a high-level overview of the procedures defined in IETF Draft "draft-ietf-sip-outbound" [ID OUTBOUND] and roles of the various PacketCable Network elements. Note that the term flow is used in [ID OUTBOUND] and in the following sections to describe a network layer connection that uses the same IP addresses and ports (UDP or TCP) at either end of the connection.

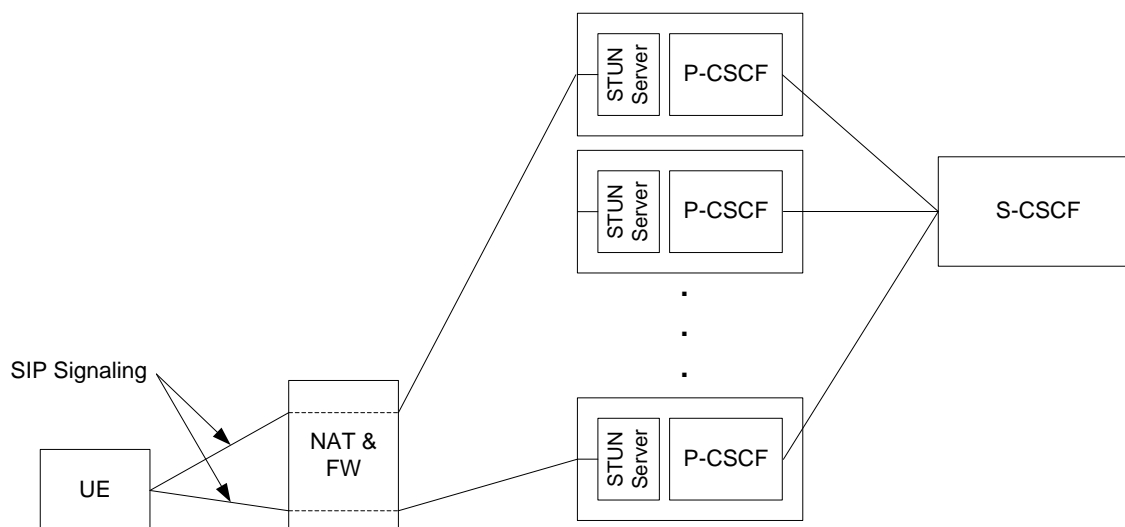


Figure 4 - NAT Traversal for SIP Signaling

As illustrated in Figure 4, a UE may be able to connect to any number of P-CSCFs. However, in order for the UE to receive incoming calls, the signaling must follow a path for which there is an existing NAT binding. Several such bindings may exist over multiple flows to Edge Proxies (e.g., for redundancy purposes). The problems of NAT traversal for SIP signaling then involve:

- Establishing an outbound connection: Setting up one or more signaling connections or flows to P-CSCF;
- Maintaining the NAT bindings and keeping FW pinholes open for those flows;
- Inbound signaling: being able to route the signaling to an appropriate P-CSCF and from there to the UE over a flow for which there is an existing NAT binding.

8.2.2.1 Establishing an Outbound Connection

SIP registration is used to set up an outbound connection and establish NAT bindings for that flow. During registration:

- The UE establishes a unique instance-id as described in [ID OUTBOUND] that remains constant over re-boots.
- The UE also uses a reg-id as described in [ID OUTBOUND] in order to identify each flow that is established with a P-CSCF. STUN is used to keep the flow (i.e., the NAT bindings and pin-holes) alive.
- The UE includes the instance-id and the reg-id when it registers. If it registers over multiple flows, then it would use the same instance-id, but a different reg-id for that different flow.
- The P-CSCF forwards the REGISTER and includes a Path header [RFC 3327] in order to establish a signaling path between P-CSCF that is terminating the specific connection/flow and the S-CSCF. The P-CSCF includes within the user portion of a loose route in the path header a unique identifier to identify the flow over which the registration occurs. The P-CSCF then maps any future requests that include that identifier to that flow.

The S-CSCF stores:

- The instance-id and reg-id as part of the contact information in addition to the time the binding was last updated.
- The Path header.

Note that multiple registrations across alternative flows (different reg-ids) allow the UE to pre-establish redundant signaling channels.

In the case where unregistered UEs are allowed to establish dialogs (e.g., emergency calls, subscribing to configuration profiles, etc.), any signaling during the life-time of that session needs to be maintained over the flow established for that session. This puts a requirement on the P-CSCF to record route and to ensure that signaling for that session occurs over that flow until the session ends.

8.2.2.2 Maintaining NAT Bindings

As indicated above, STUN is used by the UE in order to:

- Maintain the NAT bindings and keeping FW pinholes open for the signaling;
- Determine if there is a failure of the connection;
- Determine if the NAT binding has changed as a result of a NAT re-boot.

The STUN server runs on the P-CSCF on the same port that is used for signaling for that flow. The UE makes STUN requests over the flow as a keepalive mechanism for the flow as well as to determine if NAT bindings have changed as a result of a NAT re-boot.

8.2.2.3 Inbound Signaling

Note that signaling in both directions needs to be established over a flow with existing NAT bindings. In the case of UDP, this implies that SIP messages are sent and received over the same UDP port.

As a result of registrations, the S-CSCF is able to maintain contact information (including reg-ids) and Path headers in order to be able to access a given UE instance over one or more flows. Therefore, it can route an incoming call to a UE instance over a given flow; and if that fails, re-try the request over an alternative flow.

Appendix I Acknowledgements

We wish to thank the vendor participants contributing directly to this document:

Bill Foster – Cisco Systems (lead author)

Kevin Johns - CableLabs
