Diameter Maintenance and Extensions
Internet Draft

F. Brockners
S. Bhandari
S. Vikram
P. Mishra
Cisco Systems

Intended status: Informational
Expires: December 5, 2009

July 3, 2009

Review of NAT Control Protocols
draft-brockners-nat-control-protocols-review-00.txt

Status of this Memo

Copyright Notice

Abstract

This document reviews NAT control capabilities of a set of protocols and evaluates their applicability to per endpoint control of a Large Scale NAT device.

Table of Contents

1. Introduction

With the foreseeable depletion of available IPv4 addresses from the IANA pool, service providers are starting to consider network designs which no longer assign unique global IPv4 addresses to their subscribers. One of the approaches considered, is the deployment of a provider-operated large scale NAT device between the end-users and the Internet. Nishitani et al. [I-D.nishitani-cgn] call this NAT device a "Large Scale NAT (LSN)".

LSNs will be inserted into the existing subscriber access and aggregation networks which typically provide for per-endpoint service management and control as well as per-endpoint accounting. Per-endpoint rules include those which relate to service offerings of the SP (e.g. access bandwidth, time or volume based access restrictions) as well as rules which follow legal regulations of the "National Regulation Authorities (NRA)". The introduction of a LSN impacts the per-endpoint service offerings as well as the regulatory requirements and gives rise to new control requirements within the service provider network: Service providers need to dynamically manage the behavior of the LSN on a per-endpoint basis.

This document reviews a set of protocols and protocol frameworks that offer capabilities for NAT control. Based on the review, the document assesses the applicability of the different protocols to per-endpoint management of a LSN.

2. Terminology

Conventions used in this document

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119].

Abbreviations are used in this document:

Endpoint: Device representing a user, subscriber, server, or similar that requires the establishment of NAT-Bindings to communicate with other endpoints.

AAA: Authentication, Authorization, Accounting

AFT: Address Family Translation.

LSN: Large Scale NAT device.

NAT: Network Address Translation. Additional NAT related terminology definitions are found in [RFC2663].

NAT-Binding or Binding: Association of two IP-address/port pairs (with one IP-address typically being private and the other one public) to facilitate NAT.

NAT-Manager, NAT-Device: A protocol for NAT control is assumed to run between a "NAT-Manager" and a "NAT-Device". The NAT-Device is a transport network element which performs the actual NAT operations, whereas the NAT-Manager is a control entity, which controls the way the NAT-Device performs network address translation.

NAS: Network Access Server.

SP: Service Provider.

3. Protocol capabilities for NAT Control

This section briefly outlines the set of capabilities used for evaluating NAT control protocols and puts those into perspective to the NAT control capabilities expected for controlling a "Large Scale NAT" device. The review includes control capabilities for IPv4 to IPv4 NAT as well as AFT between IPv6 and IPv4.

3.1. Outline of control capabilities

Control of a NAT device can either be performed using manual configuration (e.g. using a command line or web interface to the NAT device) or through the use of a protocol. Deployment dependent, manual configuration can also be combined with a control protocol. An example for this case is, NAT-control through a web-portal where the user manually inputs parameters which are then transmitted to a NAT device using a control protocol.

This document reviews the different protocols for NAT control using the following list of capabilities:

(C1) Endpoint awareness: An endpoint aware protocol associates NAT control with an endpoint. Endpoints can be identified through a "Global Endpoint ID": The global endpoint ID will allow for common identification of an endpoint on a LSN as well as other endpoints - or subscriber-aware devices such as a Network Access Server (NAS) or an AAA system. Endpoints are identified through a single or a set of classifiers such as IP address, VLAN identifier, or interface identifier which uniquely identify the traffic associated with a particular global endpoint.

(C2) Configure NAT-Binding Limits: Define/restrict the maximum number of NAT-bindings on a per-endpoint basis. This enables service providers to offer differentiated services based on the number of bindings and hence optimize the consumption of IP-address/port-ranges. Per-endpoint NAT-binding limits also allow for protection against denial of service attacks.

(C3) Configure full NAT-Binding: Request the allocation of a pre-defined NAT-binding. Both the internal as well as the external IP-address/port pair are specified within the request. Some deployment cases, such as access to a web-server within a user's home network with IP-address and port, could benefit from statically configured bindings.

(C4) Configure half NAT-Binding: Request the allocation of an external IP-address for a given internal IP-address and report the allocated external IP-address back to the requestor. In some deployment scenarios, the application requires immediate knowledge of the allocated binding for a given internal IP-address but does not control the allocation of the external IP-address (e.g. SIP-proxy server deployments).

(C5) Configure Address pools: Define the external address-pool(s) and port ranges to be used for allocating an external IP-address. External address-pools can either be pre-defined on the LSN, or specified within a request. If pre-defined address-pools are used, a request would just include a reference (e.g. name) to an already defined address pool on LSN. Otherwise, the request will contain a description of the IP-address pool(s) to be used (e.g. list of IP-subnets).

(C6) Accounting/Reporting: Report established bindings for a particular user. Accounting can either be done on a per-binding level (referred to as (C6a: Accounting - per NAT-binding)) or on a per IP-address/port-range level (referred to as (C6b: Accounting - per range)). The later assumes that the LSN reserves ranges of port to endpoints (i.e. a block of ports which corresponds to the maximum number of NAT-bindings allowed). With per port-range accounting, accounting records only indicate that some number of NAT-bindings is allocated to a particular endpoint. Details on the particular NAT-binding are omitted.
Apart from statistical and charging purposes, binding reporting is also required for legal reasons. Most National Regulatory Authorities (NRA) require that service providers provide the identity of a user upon request. The service provider needs to be able to correlate a tuple (public IP-address, port, time) to a particular user or endpoint.

(C7) NAT-Binding Information query: Report details and statistics of bindings for a single endpoint or a set of endpoints through an external interface which integrates with the overall per-endpoint management suite. A flexible information query can be used to retrieve information about a binding that was established through the control protocol as well as information about bindings which were allocated by the NAT device autonomously for a particular endpoint.

(C8) Support address latching control: Latching describes the process of a NAT device correlating incoming data to a requested NAT-binding and as a result ignoring the remote IP address and port received in a NAT-binding request and replacing it with the incoming ("learned") addresses of the data. Details on address latching control can be found in [ITU-H.248.37]. Hosted NAT traversal solutions leverage support for address latching.

(C9) Support Address Family Translation (AFT): Network address and protocol translation (i.e. translation between IPv4 and IPv6 addresses).

(C10) Support Twice-NAT: RFC 2663 defines Twice-NAT as a "variation of NAT in that both the source and destination addresses are modified by NAT as a datagram crosses address realms. This is in contrast to Traditional-NAT and Bi-Directional NAT, where only one of the addresses (either source or destination) is translate". See also [RFC2663], section 4.3. Firewalls or session border controllers typically require a translation of source and destination address.

(C11) Support Soft-State Configs: Soft-state is a temporary state governed by periodic expiration. A NAT device supporting soft-state configuration via a dedicated NAT management protocol often leverages periodic messages to refresh the state.

(C12) Connection initiation: The protocol association between the NAT-Manager and the NAT-Device can either be initiated by the NAT-Device or the NAT-Manager. Request initiation from the NAT-Device is typically referred to as "pull" mode, whereas request initiation from the NAT-Manager is referred to as "push" mode.

(C13) Transport specific bindings: Some NAT control protocols are agnostic to the transport protocol used (e.g. they allow to refer to a protocol port number, but don't allow to differentiate between UDP or TCP), while other allow for a specific description of the layer 4 transport protocol; in which case the NAT-devices could support translation tables per transport protocol.

3.2. Control Capabilities for a Large Scale NAT

A control protocol for Large Scale NAT does not necessarily need to incorporate all the protocol capabilities which are summarized in section 3.1. Capabilities C1 to C7 as well as C13 are expected to be typically supported by a LSN. C9 will be a requirement for those LSN deployments which include AFT between IPv4 and IPv6. Twice-NAT (C10) and Support for address latching (C8) are not necessarily required, as these capabilities refer to specific deployment environments such as that of a session border controller or firewall.

Due to the fact that a LSN will be deployed within the SP domain, only protocols which support a deployment within the SP domain apply to a LSN. With this said, it should be noted that NAT-control for the user can be combined with NAT control protocols for the SP through the use of an appropriate protocol proxy. For example UPnP IGD, NAT-PmP, or NAT-PmP relay, initiated at the endpoint, can be used in

conjunction with SP focused NAT control protocols to provide a solution for NAT control.

4. Review of protocols for NAT control

The protocol review will not only factor in the different capabilities a protocol offers for NAT control, but also how these capabilities are realized. Where applicable, the document notes how a specific capability can be achieved using the semantics that the protocol offers, and whether a capability can be offered efficiently (e.g. with regards to the number of message exchanges used).

4.1. Protocols for control by the SP

4.1.1. MIDCOM

MIDCOM is not a protocol in its own right, but specifies a set of protocol semantics (see [RFC5189]) for middlebox communication. Middleboxes in the MIDCOM sense are intermediate devices (such as firewalls or NAT-gateways) which require application intelligence (see [RFC3303], [RFC3304]). MIDCOM defines the protocol semantics in terms of transactions. Those can be request transactions, or asynchronous transactions - used either for configuration or monitoring. The protocol semantics of MIDCOM can be implemented by any protocol which supports the required transaction types and procedures laid out. [RFC4097] evaluates a set of protocols (i.e. SNMP, RSIP, Megaco, Diameter, COPS) for MIDCOM compliance. [RFC5190] specifies a MIDCOM implementation using SNMP.

Capability	Supported
C1: Endpoint awareness	(Note1)
C2: Configure NAT-Binding Limits	No
C3: Configure full NAT-Binding	Yes
C4: Configure half NAT-Binding	Yes
C5: Configure Address pools	Yes (Note2)
C6a: Accounting - per NAT-binding	No (Note3)
C6b: Accounting - per range	No (Note3)
C7: NAT-Binding Information query	No (Note3)
C8: Support address latching	No
C9: Support AFT	Yes
C10: Support Twice-NAT	Yes
C11: Support Soft-State Configs	Yes
C12: Connection initiation	Push
C13: Transport specific bindings	Yes
Base protocol	n/a

```
+-------------------------------------------------+
```

Note 1: MIDCOM agents establish a session with middleboxes to facilitate configuration and monitoring of the middlebox. Both agents as well as middleboxes can maintain multiple sessions. While MIDCOM assumes only a single session between any pair of agent and middlebox (see [RFC5189], section 2.2), there are no explicit mechanisms put in place which would disallow multiple sessions between an agent and a middlebox. If an agent is endpoint aware, the session concept in MIDCOM could be leveraged for per-endpoint control of the middlebox.

Note 2: Address pool configuration is supported through MIDCOM policy rules.

Note 3: MIDCOM supports asynchronous communication from the middlebox to the agent, though this is limited to notifications about session termination, as well as notification on policy or group events. The scope of MIDCOM's transactions is at policy rule level. Hence MIDCOM offers no semantics for procedures which are specific to transient objects (i.e. NAT bindings allocated by the middlebox without explicit specification of the binding in a policy rule) controlled by a policy rule.

4.1.2. SIMCO

SIMCO [RFC4540] is a protocol following the MIDCOM framework and protocol semantics. Consequently, the NAT control capabilities of SIMCO resemble those of MIDCOM. SIMCO is not an internet standard, but a protocol specified by NEC.

```
+-------------------------------------------------+
| Capability                      | Supported     |
+-------------------------------------------------+
|C1: Endpoint awareness           | No (Note 1)   |
|C2: Configure NAT-Binding Limits |     No        |
|C3: Configure full NAT-Binding   |     Yes       |
|C4: Configure half NAT-Binding   |     Yes       |
|C5: Configure Address pools      |     Yes       |
|C6a: Accounting - per NAT-binding|     No        |
|C6b: Accounting - per range      |     No        |
|C7: NAT-Binding Information query|     No        |
|C8: Support address latching     |     No        |
|C9: Support AFT                  |     Yes       |
|C10: Support Twice-NAT           |     Yes       |
```

C11: Support Soft-State Configs	Yes
C12: Connection initiation	Push
C13: Transport specific bindings	Yes
Base protocol	SIMCO

Note 1: MIDCOM, while not specifically designed for it, could support multiple sessions between an agent and a middlebox. With this, the session could be used to represent an endpoint. This approach can only be used if the protocol used to implement the MIDCOM protocol semantics incorporates a native session concept (such as for example DIAMETER). SIMCO also leverages a session model to associate agents and middleboxes, but does not supply a session model of its own. SIMCO messages do not incorporate session identification information, hence the session context would need to be supplied by the underlying transport protocol, e.g. TCP. This renders the approach of using a SIMCO session per endpoint as non feasible.

4.1.3. ETSI Ia

ETSI Ia is a H.248 based protocol is defined in [ETSI-ES283018]. It is defined for use between the Service Policy Decision Function (SPDF) and the Border Gateway Function (BGF) within the ETSI TISPAN NGN architecture. Ia is a H.248-based control protocol which includes NAT control capabilities. In addition, it offers gate control (i.e. packets filtering depending on "IP address / port"), resource allocation and bandwidth reservation, usage metering as well as traffic policing capabilities.

ETSI Ia offers the following capabilities for NAT control:

Capability	Supported
C1: Endpoint awareness	No
C2: Configure NAT-Binding Limits	No
C3: Configure full NAT-Binding	Yes
C4: Configure half NAT-Binding	Yes
C5: Configure Address pools	No
C6a: Accounting - per NAT-binding	No
C6b: Accounting - per range	No
C7: NAT-Binding Information query	No
C8: Support address latching	Yes
C9: Support AFT	Yes
C10: Support Twice-NAT	Yes
C11: Support Soft-State Configs	No
C12: Connection initiation	Push

C13: Transport specific bindings	No
Base protocol	H.248

4.1.4. ETSI Gq'

ETSI Gq' is a Diameter application defined in [ETSI-TS183017].

The Gq' interface is the interface between the Application Function (AF) and the Service Policy Decision Function (SPDF) and is used for session based policy set-up information exchange between the SPDF and the AF.

Capability	Supported
C1: Endpoint awareness	Yes (Note1)
C2: Configure NAT-Binding Limits	No
C3: Configure full NAT-Binding	Yes
C4: Configure half NAT-Binding	Yes
C5: Configure Address pools	No
C6a: Accounting - per NAT-binding	No
C6b: Accounting - per range	No
C7: NAT-Binding Information query	No
C8: Support address latching	Yes
C9: Support AFT	Yes
C10: Support Twice-NAT	No
C11: Support Soft-State Configs	Yes
C12: Connection initiation	Push
C13: Transport specific bindings	No
Base protocol	Diameter

Note 1: Gq' leverages the session concept of the Diameter base specification [RFC3588]. Within the typical Gq' deployment context, the session-id refers to a resource reservation initiated by an application function. Given that the Session-Id is globally unique and is meant to uniquely identify a user session without reference to any other information it can be leveraged to uniquely identify an endpoint.

4.1.5. ITU Rs

ITU Rs is a Diameter application defined in [ITU-T-Q.3301.1]. Rs runs between service control entities (SCE) and the Policy Decision Physical Entity (PD-PE) in the Resource and Admission Control

Function block. Similar to Gq' the protocol includes capabilities for NAT control. In addition it can be used to request and commit transport resources, to retrieve address mapping information that can be used to modify application signaling, and to receive reports on transport resource usage for charging.

Capability	Supported
C1: Endpoint awareness	Yes (Note1)
C2: Configure NAT-Binding Limits	No
C3: Configure full NAT-Binding	Yes
C4: Configure half NAT-Binding	Yes
C5: Configure Address pools	No
C6a: Accounting - per NAT-binding	No
C6b: Accounting - per range	No
C7: NAT-Binding Information query	No
C8: Support address latching	Yes
C9: Support AFT	Yes
C10: Support Twice-NAT	No
C11: Support Soft-State Configs	Yes
C12: Connection initiation	Push
C13: Transport specific bindings	No
Base protocol	Diameter

Note 1: The considerations related to the session-id made as part of the Gq' protocol discussion apply to Rs as well.

4.1.6. ITU Rw

ITU Rw is a Diameter application defined in [ITU-T-Q.3303.3].

Rw is defined between policy decision and enforcement points for QoS resource control, gate control, and NAPT/NAT traversal control.

Capability	Supported
C1: Endpoint awareness	Yes
C2: Configure NAT-Binding Limits	No
C3: Configure full NAT-Binding	Yes
C4: Configure half NAT-Binding	Yes
C5: Configure Address pools	No
C6a: Accounting - per NAT-binding	No
C6b: Accounting - per range	No

C7: NAT-Binding Information query	No
C8: Support address latching	Yes
C9: Support AFT	Yes
C10: Support Twice-NAT	No
C11: Support Soft-State Configs	Yes
C12: Connection initiation	Push/Pull
C13: Transport specific bindings	No
Base protocol	Diameter

4.2. Protocols for control by the End-User

4.2.1. UPnP IGD

UPnP IGD (see [UPnP-IGD]) was designed to be used primarily for home NAT gateways. UPnP IGD provides methods for following:

- Learning public IP address
- Enumerating existing port mappings
- Adding and removing port mappings
- Assigning lease times to mappings

Capability	Supported
C1: Endpoint awareness	No
C2: Configure NAT-Binding Limits	No
C3: Configure full NAT-Binding	No
C4: Configure half NAT-Binding	Yes(Note1)
C5: Configure Address pools	No
C6a: Accounting - per NAT-binding	No
C6b: Accounting - per range	No
C7: NAT-Binding Information query	No
C8: Support address latching	No
C9: Support AFT	No
C10: Support Twice-NAT	No
C11: Support Soft-State Configs	Yes
C12: Connection initiation	Push(Note2)
C13: Transport specific bindings	Yes
Base protocol	XML based appln on UDP

Note 1: UPnP IGD allows querying for public IP that will be used and port mappings created.

Note 2: UPnP IGD is for home subscribers to request/query for NAT binding information with NAT device.

4.2.2. Bonjour NAT-PMP

Bonjour NAT (see [I-D.cheshire-nat-pmp]) port mapping protocol was defined for automating the process of creating Network NAT port mappings. It includes a method for retrieving the external IP address of a NAT gateway, thus allowing a client to make this external IP address and port number known to peers that may wish to communicate with it.

This protocol is designed for small home networks, with a single logical link (subnet) where the client's default gateway is also the NAT translator for that network. For more complicated networks where the NAT translator is some device other than the client's default gateway, this protocol is not appropriate.

Capability	Supported
C1: Endpoint awareness	No
C2: Configure NAT-Binding Limits	No
C3: Configure full NAT-Binding	No
C4: Configure half NAT-Binding	Yes(Note1)
C5: Configure Address pools	No
C6a: Accounting - per NAT-binding	No
C6b: Accounting - per range	No
C7: NAT-Binding Information query	No
C8: Support address latching	No
C9: Support AFT	No
C10: Support Twice-NAT	No
C11: Support Soft-State Configs	Yes
C12: Connection initiation	Push(Note2)
C13: Transport specific bindings	Yes
Base protocol	New appln over UDP

Note 1: NAT-PmP allows querying for external IP that would be used without influencing it and allows requesting for actual port number to be mapped.

Note 2: NAT-PmP is for home subscribers to request/query for NAT binding information with NAT device.

4.3. Protocols for control by the End-User or SP

4.3.1. NAT-PMP relay

Applicability of NAT-PMP with Service Provider Deployments of Network Address Translation is studied in [I-D.woodyatt-spnatpmp-appl].

NAT-PMP relay extends Bonjour NAT-PMP to service provider deployments. The primary requestor for port mapping would be the subscriber end device. This specification describes a way to relay the NAT-PmP messages from subscriber home gateway to service provider NAT gateway. Though it adds a response code for port mapping request to indicate "Per-subscriber resource limit would be exceeded" it is unspecified as to how a subscriber is identified and this decision made.

Capability	Supported
C1: Endpoint awareness	No
C2: Configure NAT-Binding Limits	No
C3: Configure full NAT-Binding	No
C4: Configure half NAT-Binding	Yes(Note1)
C5: Configure Address pools	No
C6a: Accounting - per NAT-binding	No
C6b: Accounting - per range	No
C7: NAT-Binding Information query	No
C8: Support address latching	No
C9: Support AFT	No
C10: Support Twice-NAT	No
C11: Support Soft-State Configs	Yes
C12: Connection initiation	Push
C13: Transport specific bindings	Yes
Base protocol	NAT-PmP

Note 1: NAT-PmP allows querying for external IP that would be used

without influencing it and allows requesting for actual port number to be mapped.

4.3.2. NSLP

NSLP (NAT/Firewall NSIS Signaling Layer Protocol) is specified in [I-D.ietf-nsis-nslp-natfw].

The NATFW NSLP is designed to request the dynamic configuration of NATs and/or firewalls along the data path. Dynamic configuration includes enabling data flows to traverse these devices without being obstructed, as well as blocking of particular data flows at inbound firewalls.

Capability	Supported
C1: Endpoint awareness	Yes(Note1)
C2: Configure NAT-Binding Limits	No
C3: Configure full NAT-Binding	No
C4: Configure half NAT-Binding	Yes
C5: Configure Address pools	No
C6a: Accounting - per NAT-binding	No
C6b: Accounting - per range	No
C7: NAT-Binding Information query	No
C8: Support address latching	No
C9: Support AFT	No
C10: Support Twice-NAT	No
C11: Support Soft-State Configs	Yes
C12: Connection initiation	Push(Note2)
C13: Transport specific bindings	Yes
Base protocol	GIST(Note3)

Note 1: NATFW NSLP has signaling session concept. Signaling session is initiated by Data sender and created at every hop in the data path. Each NSIS node that supports NATFW NSLP can authenticate and authorize the signaling session before applying the policy requested. Thus there will be multiple NSLP sessions per endpoint based on every unique IP 5 tuple.

Note 2: NSLP protocol defines dynamic configuration of NAT device by signaling messages initiated by Data sender towards Data receiver. NAT device is not capable of querying for NAT policy when it encounters data packets for which no policy is installed.

Note 3: General Internet Signaling Transport(GIST) - the implementation of the NTLP) defined in [I-D.ietf-nsis-ntlp].

4.3.3. NAT with explicit control (NAT-XC)

Details on NAT-XC can be found in [I-D.moore-nat-xc]. Although focused on IPv4/IPv6 NAT, NAT-XC could be applied to IPv4/IPv4 NAT as well. NAT-XC is designed to allow applications to be explicitly aware of, and control, their address bindings. By closely associating the NAT-Device (i.e. "translator") with the application it provides a predictable programming environment for applications. Consequently the typical operational model of the NAT-Device does not resemble the model of a default router, but the application clients will send packets directly to the NAT-Device - making the assumed operational model similar to for example that of a session border controller. The NAT-XC binding definition is also reflected by the operational model: A binding in the sense of NAT-XC is a 9-tuple consisting of transport protocol, local and remote address and port of the NAT-device(s), as well as the local and remote endpoint addresses.

Note that NAT-XC was at draft state at the time this document was written. Several protocol details weren't fully defined.

Capability	Supported
C1: Endpoint awareness	(Note1)
C2: Configure NAT-Binding Limits	No (Note2)
C3: Configure full NAT-Binding	Yes
C4: Configure half NAT-Binding	No (Note2)
C5: Configure Address pools	No (Note2)
C6a: Accounting - per NAT-binding	Yes (Note3)
C6b: Accounting - per range	No
C7: NAT-Binding Information query	Yes (Note4)
C8: Support address latching	No
C9: Support AFT	Yes
C10: Support Twice-NAT	Yes
C11: Support Soft-State Configs	Yes
C12: Connection initiation	push
C13: Transport specific bindings	Yes
Base protocol	NAT-XC

Note 1: Not fully applicable. Due to the tight integration of NAT-Device and NAT-Manager a concept of NAT-endpoint awareness at the NAT-device is not really applicable.

Note 2: Not applicable due to the operational model of NAT-XC, which assumes explicit control of the bindings created by the application.

Note 3: "Binding Notification Messages" of NAT-XC (see [I-D.moore-nat-xc], section 3.7.6) could be leveraged to supply per NAT-binding accounting.

Note 4: NAT-XC supplies a "Get Binding List" operation (see [I-D.moore-nat-xc], section 3.7.5) which allows a NAT-Manager to retrieve a list of all existing bindings.

Note that the typical operational model for NAT-XC does not match the deployment context of a LSN, because NAT-XC assumes a close integration with the application and requires the application to control every binding on the LSN. This contradicts the default model of a LSN which autonomously creates bindings. NAT-XC includes autonomous operation of the a NAT-device as a corner case. Section 2.3 of the draft [I-D.moore-nat-xc] mentions a model where NAT-Manager and NAT-Device would be implemented on the same physical device. In this case, the NAT-Manager leverages traffic analysis and heuristics to control the NAT-Device. In case NAT-Manager and NAT-Device are co-located on the same physical device, remote management of the NAT-Device is naturally excluded.

5. Security Considerations

This document studies the generic capabilities of several protocols as they apply to the control of a NAT device. Please refer to the specifications of the individual protocols for detailed information on the security solutions they provide.

6. IANA Considerations

This document has no actions for IANA.

7. References

7.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC2663] Srisuresh, P. and M. Holdrege, "IP Network Address Translator (NAT) Terminology and Considerations", RFC 2663, August 1999.

[RFC3303] Srisuresh, P., Kuthan, J., Rosenberg, J., Molitor, A., and A. Rayhan, "Middlebox communication architecture and framework", RFC 3303, August 2002.

[RFC3304] Swale, R., Mart, P., Sijben, P., Brim, S., and M. Shore, "Middlebox Communications (midcom) Protocol Requirements", RFC 3304, August 2002.

[RFC3588] Calhoun, P., Loughney, J., Guttman, E., Zorn, G., and J. Arkko, "Diameter Base Protocol", RFC 3588, September 2003.[RFC4540] Stiemerling, M., Quittek, J., and C. Cadar, "NEC's Simple Middlebox Configuration (SIMCO) Protocol Version 3.0", RFC 4540, May 2006.

[RFC4097] Barnes M., "Middlebox Communications (MIDCOM) Protocol Evaluation", RFC 4097, June 2005.

[RFC5189] Stiemerling, M., Quittek, J., and T. Taylor, "Middlebox Communications (MIDCOM) Protocol Semantics", RFC 5189, March 2008.

[RFC5190] Quittek, J., Stiemerling, M., Srisuresh, P., "Definitions of Managed Objects for Middlebox Communication", RFC 5190, March 2008.

[I-D.ietf-nsis-nslp-natfw] M. Stiemerling,H. Tschofenig, C. Aoun and E. Davies, "NAT/Firewall NSIS Signaling Layer Protocol (NSLP)", draft-ietf-nsis-nslp-natfw-20.txt(work in progress), November 2008

[I-D.ietf-nsis-ntlp] Schulzrinne, H. and R. Hancock, "GIST: General Internet Signalling Transport", draft-ietf-nsis-ntlp-20 (work in progress), June 2009.

[I-D.moore-nat-xc] Moore, K., "IPv4/v6 NAT With Explicit Control (NAT-XC)", draft-moore-nat-xc-02(work in progress), March 2009.

[I-D.woodyatt-spnatpmp-appl] Woddyatt, J., "Applicability of NAT-PMP with Service Provider Deployments of Network Address Translation", draft-woodyatt-spnatpmp-appl-01 (work in progress), November 2008.

[I-D.cheshire-nat-pmp] Cheshir, S., Krochmal, M., Sekar, K., "NAT Port Mapping Protocol (NAT-PMP)", draft-cheshire-nat-pmp (work in progress), April 2008.

[I-D.nishitani-cgn] Nishitani, T. and S. Miyakawa, "Common Functions of Large Scale NAT (LSN) ", draft-nishitani-cgn-02 (work in progress), May 2009.

[ITU-H.248.37] "Gateway control protocol: IP NAPT traversal package", ITU-T H.248.37, September 2005.

[ITU-T-Q.3303.3] "Resource control protocol no. 3 (rcp3): Protocol at the Rw interface between the Policy Decision Physical Entity (PD-PE) and the Policy Enforcement Physical Entity (PE-PE): Diameter", ITU-T Recommendation Q.3303.3, 2008.

[ITU-T-Q.3301.1] "Resource control protocol No. 1 - Protocol at the Rs interface between service control entities and the policy decision physical entity", ITU-T Q.3301.1, March 2007.

[ETSI-TS183017] Telecommunications and Internet converged Services and Protocols for Advanced Networking (TISPAN);Resource and Admission Control: DIAMETER protocol for session based policy set-up information exchange between the Application Function (AF) and the Service Policy Decision Function (SPDF);Protocol specification; ETSI TS 183017, Version 2.3.1, September 2008.

[ETSI-ES283018] Telecommunications and Internet converged Services and Protocols for Advanced Networking (TISPAN);Resource and Admission Control: H.248 Profile for controlling Border Gateway Functions (BGF) in the Resource and Admission Control Subsystem (RACS);Protocol specification, ETSI ES 283 018, Version 2.3.1, June 2008.

[UPnP-IGD] UPnP Forum, "Universal Plug and Play (UPnP) Internet Gateway Device (IGD)", November 2001, <http://www.upnp.org/standardizeddcps/igd.asp>

8. Author's Addresses

Frank Brockners
Cisco Systems, Inc.
Hansaallee 249
40549 Duesseldorf
Germany
Email: fbrockne@cisco.com

Shwetha Bhandari
Cisco Systems, Inc.
Cessna Business Park,
Sarjapura Marathalli Outer Ring Road
Bangalore, KARNATAKA 560 087
India
Email: shwethab@cisco.com

Shashank Vikram
Cisco Systems, Inc.
Cessna Business Park,
Sarjapura Marathalli Outer Ring Road
Bangalore, KARNATAKA 560 087
India
Email: svikram@cisco.com

Pallavi Mishra
Cisco Systems, Inc.
Cessna Business Park,
Sarjapura Marathalli Outer Ring Road
Bangalore, KARNATAKA 560 087
India
Email: palmishr@cisco.com