

DECLARATION

I, Alexa Morris, based on my personal knowledge and information, hereby declare as follows:

1. I am Executive Director of the Internet Engineering Task Force (“IETF”) and have held this position since January 1, 2008.
2. The Internet Engineering Task Force (IETF) is a large open international community of network designers, operators, vendors, and researchers concerned with the evolution of the Internet architecture and the smooth operation of the Internet. It is open to any interested individual. The goal of the IETF is to make the Internet work better. The mission of the IETF is to produce high quality, relevant technical and engineering documents that influence the way people design, use, and manage the Internet in such a way as to make the Internet work better. These documents include protocol standards, best current practices, and informational documents of various kinds. The IETF pursues this mission in adherence to several cardinal principles including Open Process - any interested person can participate in the work, know what is being decided, and make his or her voice heard on the issue. Part of this principle is our commitment to making our documents, our WG mailing lists, our attendance lists, and our meeting minutes publicly available on the Internet.
3. Among my responsibilities as IETF Executive Director, I act as the custodian of business records, including meeting minutes and proceedings, for the IETF.
4. I make this declaration based on my personal knowledge and information contained in the business records of the IETF, or confirmation with other responsible IETF personnel with such knowledge.

- 1 5. It is the regular practice of the IETF to publish minutes of meetings of IETF working
2 groups and make them available to the public on its website at www.ietf.org. The IETF
3 maintains copies of minutes of working groups in the ordinary course of its regularly
4 conducted activities.
- 5
- 6 6. A true and correct copy of the minutes of working group 2.7.8 "Network Address
7 Translators (nat)" held at IETF 43 in 1998 is attached hereto as Exhibit 1 to this
8 declaration.
- 9
- 10 7. I personally reviewed the document attached as Exhibit 1 to this declaration.
- 11
- 12 8. I hereby certify that Exhibit 1 to this declaration constitutes a record of regularly
13 conducted business activity which was (A) made at or near the time of the occurrence of
14 the matters set forth by, or from information transmitted by, a person with knowledge of
15 those matters; (B) kept in the course of the regularly conducted activity; and (C) made
16 by the regularly conducted activity as a regular practice.
- 17
- 18 9. Based on a search of IETF records, I have determined that the IETF maintained a copy
19 of Exhibit 1 to this declaration in the ordinary course of its regularly conducted
20 activities.
- 21
- 22 10. It is IETF's regular course of business to publish minutes of working group meetings
23 held during a regular IETF meeting no later than the first day of the proceedings of the
24 next IETF meeting. Based on a search of IETF records and the IETF's course of
25 conduct in publishing working group minutes, I have determined that Exhibit 1 to this
26 declaration, which constitutes minutes taken at IETF 43, was published on the IETF
27 website (www.ietf.org) no later than March 14, 1999, the first day of IETF 44. At such
28 time, such document was reasonably accessible to the public, and was disseminated or

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28

otherwise available to the extent that persons interested and ordinarily skilled in the subject matter or art exercising reasonable diligence could have located it.

Pursuant to Section 1746 of Title 28 of United States Code, I declare under penalty of perjury under the laws of the United States of America that the foregoing is true and correct and that the foregoing is based upon personal knowledge and information and is believed to be true.

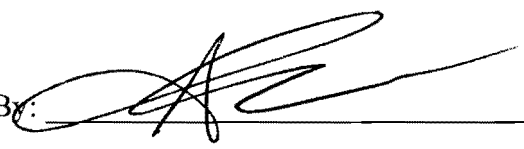
Date: August 26, 2015 By: 
Alexa Morris

EXHIBIT 1

NAT WG meeting minutes - 43rd IETF - Orlando, FL - December 9, 1998

Chairs: Matt Holdrege < matt@ascend.com >

- Pyda Srisuresh < suresh@ra.lucent.com >

Reported by: Ben Rogers < ben@ascend.com >

Edited by: Matt & Suresh

Mailing list: nat@livingston.com

To subscribe: Send e-mail to nat-request@livingston.com with "subscribe" in the body of the message.

- To unsubscribe: Send e-mail to nat-request@livingston.com with "unsubscribe" in the body of the message.

Mailing list: nat-digest@livingston.com

- (This is nat mailing list, in daily-digest format.)
- To receive the digest, you can subscribe as follows.)

To subscribe: Send e-mail to nat-digest-request@livingston.com with "subscribe" in the body of the message. To unsubscribe: Send e-mail to nat-digest-request@livingston.com with "unsubscribe" in the body of the message.

All presentations, along with WG meeting minutes will be available on-line from the following URL. <http://www.livingston.com/tech/ietf/nat>

In order to avoid confusion, the following indentation and format legend should be used as a guide to interpreting the minutes.

<Presenter> - "<presentation/Discussion Topic>"

<Any other remarks by minutes taker or editors>

Detailed Slides and/or Comments made by the presenter

Questions from the Audience:

[<Audience-Name> - <Question/Comment>]

<Response from the Presenter>

Matt Holdrege - Introduction

Two NAT Drafts sent to mailing list completed WG last call and are currently awaiting IESG review for advancing into informational RFC status.

Agenda

1. "NAT Friendly Application Design Guidelines" 20 mins.

<draft-ietf-nat-app-guide-00.txt> by Daniel Senie

2. "A Multihoming solution using NATs" 20 mins.

<draft-akkiraju-nat-multihoming-00.txt > by Praveen Akkiraju and Yakov Rekhter

3. "Security for IP Network Address Translator 20 mins.

(NAT) Domains"

<draft-ietf-nat-security-00.txt> by Pyda Srisuresh

4. "IP Network Address Translator (NAT) Protocol Issues" 20 mins.
<draft-ietf-nat-protocol-issues-01.txt> by Matt Holdrege & Pyda Srisuresh

5. "IP Network Address Translator Application 20 mins.

Programming Interface"

<draft-ietf-nat-api-00.txt> by Pyda Srisuresh

6. "IP Host Network Address (and Port) Translation" 20 mins.

<draft-ietf-nat-hnat-00.txt> by Jeffrey Lo & K. Taniguchi

Daniel Senie - "NAT Friendly Application Design Guidelines"

- <draft-ietf-nat-app-guide-00.txt>

<http://www.amarantnetworks.com/nat> - This is the link to Daniel's web page that will have the most up to date copy of applications that are NAT friendly

General Goal

- Develop new application protocols which do not require ALG assistance from NAT and Firewall implementations

- These guidelines tend to apply to both NAT and Firewall implementations as the requirements for a firewall are similar.

General Recommendations

- Use of Multiple Session Bundle

- Use single TCP or UDP port for all traffic where possible

- If possible, originate all sessions from the same endpoint (this avoids the FTP problem)

- Use DNS names, not IP addresses

- Avoid passing addressing data in payload (IP addresses, port numbers)

- End to end IPsec problematic

- Consider using TLS instead. Host NAT is a possibility for the future

- Service Location

- Avoid location via broadcast or multicast

- Subsequent Sessions

- Address bindings not guaranteed between subsequent sessions.

Reliance on the appearance of co-location can be problematic

- Operational Reliability

- TCP preferred over UDP since NAT can track TCP sessions more easily and know when sessions end.

- UDP sessions are tracked by timeouts.

ALG's can overcome this problem, but we'd rather design applications to not need this processing.

- Processing Overhead

- Each new session requires resources and processing to establish associations. Using a single session for app reduces overhead.

There is still some overhead, but this is unavoidable.

Additional Items

- Comments since the last draft
- Performance: latency and throughput can be affected by NAT. Depends on implementation.

The hardware solution can be optimized to lessen this problem

- NAT implementations presently tend to only support TCP and UDP applications (... for implementations of NAT). A device performing just NAT, without any ALG support, supports many applications as is.

Questions from the Audience:

[Eliot Lear - UDP session management. UDP may make it more difficult to maintain the mapping]

An application may maintain keep-alives to make this less of a problem.

[Someone said there are similar issues with SOCKS, and these ideas can be shared with NAT. Do we have any plan to make a utility to verify the guidelines?]

Multiple sessions can work, but are not as reliable, nor as friendly as single sessions. An analyzer might be created to diagnose traffic in a non-NAT environment.

[Can we add a recommendation to change current applications (eg. modify FTP to use passive) to avoid these problems in current protocols.]

An RFC exists on this particular issue, RFC1579 by Steve Bellovin. It doesn't reduce FTP to a single session, though, just makes the connections start from the same endpoint.

Praveen Akkiraju - "A Multihoming solution using NATs"

- <draft-akkiraju-nat-multihoming-00.txt>

Agenda

- Terminology
- Bootstrapping
- Routing configuration
- DNS configuration
- Scenario 1: internally originated connections
- Scenario 2: externally originated connections
- Discussion

Problems with Multihoming

- Creates scaling problems for the global routing infrastructure.
- Use of multiple address blocks from upstream ISPs provides a possible solution, but results in added address management complexity
- Controlling load balancing based on address assignment is complex
- Difficult to achieve symmetric flow of packets in and out of an enterprise

Terminology

- Inside Global Address (IGA): Block of address space assigned by the ISP to the enterprise.
- Inside Local Address (ILA): Address space used within the enterprise behind the NAT
- Outside Global Address (OGA): Legal address space outside the enterprise and the

NAT

- Outside Local Address (OLA): Address space in the enterprise used to translate OGA addresses

[Brian Carpenter - terminology is quite confusing and is worth rethinking, because outside and inside are backwards from what we're currently used to.]

Topology

An enterprise with 2 NAT routers interfacing with 2 different ISPs.

Routing Configurations

- NAT configured to assign IGA prefixes to internal addresses and OLA prefixes to external addresses.

- OLA prefixes must not overlap with ILA or OGA address space

- NATs EBGP peer with upstream ISPs to advertise IGA prefixes

- NATs injects OLA prefixes into the enterprise IGP routing

- NATs also IBGP peer with each other

DNS configuration

Goal: Bootstrap exchange of DNS messages across NAT

Reaching the parent zone server

- Assign an OLA prefix to the external parent server and create a translation entry in the NAT mapping the server OGA to its OLA

- Configure internal DNS server with the parent zone server's OLA

Reaching foo.com's DNS server

- Assign a IGA prefix to the internal DNS server and create the corresponding translation entry in both the NATs

- Configure DNS Server for authoritative parent zone with the IGA's from both NAT's to reach foo.com's authoritative server

(DNS message handling as in draft-ietf-nat-dns-alg-01.txt)

Initial NAT Table

Internal DNS server and Root DNS server are both mapped to provide inside/outside pairs for each.

NAT bootstrapped with foo.com and parent zone DNS server mappings.

DNS Summary

- DNS response processing in a NAT consists of 2 stages

- translation of the IP-UDP header

- translation of the reply message

(Walk-through of full packet flow for a DNS lookup, and changes to the NAT table.)

The structure of this query-response system ensures that all traffic flow passes through the same NAT.

Resolving X.foo.com (externally initiated session)

(Walk through of full packet flow for the incoming DNS-lookup)

Discussion

- Scheme places no addressing restrictions on the multihomed network except in OLA

allocation.

- Scheme is independent of the enterprises internal routing architecture.
- Changing providers doesnt require renumbering.
- Due to address allocation packet flow between a pair of hosts will always flow through the same NAT resulting in symmetry
- In case of link failure between the NAT and ISP, fallback connectivity may be provided using RFC 2260 mechanisms.
- Controlling the selection of the NAT which processes a DNS response controls the rest of the packet flow to the target host.
- Load balancing is tied to the flow of DNS packets
- Translation of the IP-UDP and DNS reply are handled in the same NAT
- Under policy control Translation of the IP-UDP and DNS reply may be handled in seperate NATs
- This allows for more flexible policy based load balancing of traffic flows.

Misc.

- Draft includes lab-tested router and DNS server configurations
- Concept tested by Ed Kern at NANOG
<http://www.academ.com/nanog/feb1998/nat.html>
- Refer draft-rfcd-info-srisuresh-05.txt for details on NAT operations and issues.

Questions from the Audience:

[Brian Carpenter - DNS server for foo.com has to be inside foo.com, and zone transfer will result in a disaster.]

Correct

[continued... Very clever zone transfer (at minimum) is needed, if this is possible at all. This is not very clear in the text.]

[Suresh - Names are apparently end-to-end unique. Is the DNS information here a duplicate effort of the DNS-ALG draft (or) Is there something particularly different in this draft?]

DNS information here is same as in DNS-ALG draft.

The information is included here for Completeness.

[continued.. Twice nat and BGP are implied as required in this draft, while they are only required for the given scenarios.]

Twice nat is not necessarily required, but BGP seems to be.

[continued.. BGP doesn't have to be on boxes connecting to two ISPs. The draft must be scoped properly. Terminology should be cleared up to be consistent with other drafts. Also, DNS based Load sharing discussed here seems orthogonal to multihoming and NAT issues.]

[Eliot Lear - We can probably get away with no BGP, but we would still need to use some dynamic routing protocol.

[Yakov - BGP is mentioned because an existing RFC 2260 for robust multihoming is based on BGP]

[Radia Perlman - Talked about ISP link going down, but not what happens when the NAT itself goes down. Can we force new DNS queries so that address caching does not occur.]

Force of 0 ttl should address this problem.

[Bob Brandt - This seems to infer that there is a 1-1 mapping between external DNS servers and internal addresses. This seems to create a scaling problem.]

Addresses are assigned for a limited time, so we can timeout the addresses quickly.

Pyda Srisuresh - "Security Model for NAT domains"

- <draft-ietf-nat-security-00.txt>

Problems with End-to-end IPsec.

Security from NAT box to external node seems to be possible

Model made to be in line with IPsec model to provide security for NAT nodes-> external nodes

Terminology

- Clear-NAT - Nat applied to Outer packet Header
- Secure-NAT - NAT applied to packets embedded within a secure tunnel
- Secure-NAT gateway device - Supports both Clear-NAT and Secure-NAT
- Terminology from IPsec RFCs + nat draft

Secure-NAT Model

- Secure Policies administered using private realm addressing
- Secure-NAT address mapping
- Security Association Parameters

Secure-NAT Gateway Outbound Packet Processing

(Flow-chart to describe NAT processing decision)

(Verify Outbound packet against outbound Secure policies. In case of a match, subject the packet to Secure-NAT prior to tunneling using Outbound IPsec SA.)

Secure-NAT Gateway Inbound IPsec Packet Processing

(Detunnel packet using IPsec SA, perform inbound Secure-NAT, then verify that it matches inbound security policies)

Slides thus far make the assumption that keys are statically configured.

IKE support to secure-NAT GW

(Large diagram describing the role of IKE-ALG in the translation of secure policies)

Secure-NAT applications

- Secure Extranet connectivity - Allows private domains to connect securely to external domains.
- Secure mobile user connectivity - allows mobile users to connect securely to enterprise routers.

There is a draft by Vipul (<draft-gupta-ipsec-remote-access-00.txt>) that describes how secure remote access is possible for a mobile user, without requiring both Home and Care-of addresses on Mobile Node.

Matt Holdrege - "NAT Protocol Issues" draft

- <draft-ietf-nat-protocol-issues-01.txt>

Title has been a point of contention (issues?)

"Limitiations" proposed on the list

"NAT Protocol Complications" is the new proposed title.

We need input from the community

We will continue to document the additional information we receive on protocols in the current format.

We are working on a better method to get more contributions.

The draft will be titled after we can put more information into this draft.

Questions from the Audience:

[Brian Carpenter - RFC on renumbering issues brought up a similar set of problems.

There was not at the time a clean way to find all the affected protocols. Y2K was solved by searching 2000 RFCs, but an appeal to the community may be better in this case.]

Diverse group in the room should be able to talk to members of other WGs to help find these types of issues.

Pyda Srisuresh - "NAT API"

- <draft-ietf-nat-api-00.txt>

Draft title is not necessarily representative of the objectives of this draft. In particular, the intent is not to mandate standardization of the NAT API.

Objective

- Identify external agents and their interface requirements with NAT
- Provide a framework for the development of one or more protocols by which agents can interface with NAT
- Identify NAT objects that could be externalized with MIB
- Provide an API framework for agents on the same device to interface with NAT.

External Agents

- Host-NAT and Host-NAPT clients need some way to obtain their addresses and routes.
- Management application needed to configure a NAT (which ALGs are available, ...)
- Backup NAT may be needed.
- ALGs

A number of agents with different requirements would need to interface with the NAT. (Hence the API name.)

NAT is doing significant resource management, and this information is useful for network management applications.

These constitute the behind-the-scene's agenda for the creation of this draft. There is no intention to standardize on an API.

NAT elements

Draft will contain more gory details than are discussed in this meeting.

- NAT Descriptor - ID, Type, Address map, etc.
- BINDing Descriptor - ID, Type, specific addresses (port) bound, Leased time, Controlling Agent ID, etc.

- Session Descriptor - ID, Controlling BIND ID, Original and Translated session parameters, Session Tag, Session Termination heuristic, etc.

External Agent descriptor

- Agent ID
- Agent Type
- Agent Call back requirements
- Agent Call-back functions
- Agent Accessibility information

Function calls provided to agents

(List of function calls that an agent could invoke)

Functions are provided to perform inquiries on bindings, sessions to register with NAT (as an ALG, ...) to set parameters within BIND or Session to free BIND or Sessions

Callback functions in agents

(List of callback function an agent could provide)

Callbacks may occur on events or packets, or simply periodically.

The ADs have advised that this WG is the right forum to discuss Host-NAT issues.

However, we need to ensure that the base documents of the WG are given priority.

Questions from the Audience:

[Eliot Lear - This seems to be required. Brian raised the issue of architectural implications of NAT. This draft addresses some of those problems, especially in terms of Host NAT.]

This doesn't provide for Host NAT, per-se.

[Brian Carpenter - Hasn't read the draft.]

Jeffrey Lo - "Host Network Address Translation"

- <draft-ietf-nat-hnat-00.txt>

This draft came out recently, so there may not have been many people who have taken a look at it. (Some indication of Distributed NAT.)

Framework

- Uses a common "control" protocol for HNAT parameter negotiation
- Uses tunneling mechanisms for routing externally addressed packets within private domain.

Motivations

- Lack of a common protocol that enables Host-NAT-client and Host-NAT-Server to interoperate.

- Hence goal is to design a common protocol that enables Host-NAT-client and Host-NAT-server to interoperate.

Protocol requirements

- Client type registration and identification
- Client ID assignment
- Enables negotiation of multiplexing fields
- Global address, port, protocol ID

- Enables negotiation of HNAT related parameters
- Max. Leased Time, NAT type
- Support negotiation of tunnel type
- Support subnet query

Considerations

- Exploitation of existing protocols?
- DHCP, ICMP, COPS, TCP/UDP/RUTS based?
- Independence
- Support negotiation of all fundamental parameters required to perform HNAT.
- Extensibility
- Flexible packet format
- Able to support vendor specific information
- New error codes
- Able to accommodate new policies.

Proposal

- Dynamic Bindings Acquisition Protocol (DBAP)
- ICMP based
- Enable assignment of various parameters

End host implementation

(Chart of relations between various pieces of client and server)

Questions from the Audience:

[Daniel Senie - Choice of ICMP versus TCP]

This was designed around a larger scale implementation.

[Nair Venugopal - Does this mean that we can use transport mode IPsec.]

Yes.