Behavior Engineering for Hindrance Avoidance
Internet-Draft
Intended status: Informational
Expires: July 10, 2011

R. Penno
Juniper Networks
T. Saxena
Cisco Systems
M. Boucadair
France Telecom
S. Sivakumar
Cisco Systems
January 6, 2011

Analysis of 64 Translation
draft-penno-behave-64-analysis-06

Abstract

Due to specific problems, NAT-PT was deprecated by the IETF as a mechanism to perform IPv6-IPv4 translation. Since then, new efforts have been undertaken within IETF to standardize alternative mechanisms to perform IPv6-IPv4 translation. This document evaluates how the new translation mechanisms avoid the problems that caused the IETF to deprecate NAT-PT.

Requirements Language

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 RFC 2119 [RFC2119].

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at http://datatracker.ietf.org/drafts/current/.

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. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on July 10, 2011.

Copyright Notice

Table of Contents

1. Introduction 3
1.1. Definition 3
1.2. Context 3
1.3. Scope 4
2. Analysis of 64 Translation Against Concerns of RFC4966 4
2.1. Problems Not Addressed by 64 4
2.2. Problems Addressed by 64 7
3. Conclusions 9
4. IANA Considerations 11
5. Security Considerations 11
6. Acknowledgements 11
7. References 12
7.1. Normative References 12
7.2. Informative References 13
Authors' Addresses 14

1. Introduction

1.1. Definition

This document uses 64 proposal (or 64 for short) to refer to the mechanisms defined in the following documents:

- Stateful NAT64: Network Address and Protocol Translation from IPv6 Clients to IPv4 Servers [I-D.ietf-behave-v6v4-xlate-stateful]
- DNS64: DNS extensions for Network Address Translation from IPv6 Clients to IPv4 Servers [I-D.ietf-behave-dns64]
- IPv6 Addressing of IPv4/IPv6 Translators [RFC6052]
- Framework for IPv4/IPv6 Translation [I-D.ietf-behave-v6v4-framework]

1.2. Context

The current 64 proposal is widely seen as the next step in the evolution of interconnection techniques enabling communications between IPv6-only and IPv4-only networks. One of the building blocks of this proposal is decoupling the DNS functionality from the protocol translation itself.

This approach is pragmatic in the sense that there is no dependency on DNS implementation for the successful NAT handling. As long as there is a function (e.g., DNS64 [I-D.ietf-behave-dns64] or other means) that can construct an IPv6-Embedded IPv4 address with a pre-configured IPv6 prefix, an IPv4 address and a suffix (refer to [RFC6052]), NAT64 will work just fine.

To understand the 64 proposal, we must keep in mind that the focus of this proposal is on the deployment and not the implementation details. As long as a NAT64 implementation conforms to the expected behaviour, as desired in the deployment scenario, the details are not very important as mentioned in this excerpt from [I-D.ietf-behave-v6v4-xlate-stateful]:

> "A NAT64 MAY perform the steps in a different order, or MAY perform different steps, but the externally visible outcome MUST be the same as the one described in this document."

1.3. Scope

This document provides an analysis of how the proposed set of documents that specify stateful IPv6-only to IPv4-only translation and replace NAT-PT [RFC2766] address the issues raised in [RFC4966].

As a reminder, it is worth mentioning the 64 proposal analysis is limited in the sense that hosts from IPv6 networks can initiate a communication to IPv4 network/Internet, but not vice-versa. This corresponds to Scenario 1 and Scenario 5 described in [I-D.ietf-behave-v6v4-framework]. Hence, the scenario of servers moving to IPv6 while clients remaining IPv4 remains unaddressed. Of course, IPv6 to IPv4 communications can also be supported if static bindings are configured on the stateful NAT64.

The 64 proposal, just like any other technique under development, has some positives and some drawbacks. The ups and downs of the proposal must be clearly understood while going forward with its future development.

The scope of this document does not include stateless translation.

2. Analysis of 64 Translation Against Concerns of RFC4966

Of the set of problems pointed out in [RFC4966], the 64 proposal addresses some of them, whereas leaves others unaddressed.

Some issues mentioned in [RFC4966] were solved by [RFC4787], [RFC5382] and [RFC5508]. At the time when NAT-PT was published these recommendations were not in place but they are orthogonal to the translation algorithm per se, therefore they could be implemented with NAT-PT. On the other hand, NAT64 explicitly mentions that these recommendations need to be followed and thus should be seen as a complete specification.

It is also worth pointing out that the scope of the 64 proposal is reduced when compared to NAT-PT. Following is a point by point analysis of the problems.

2.1. Problems Not Addressed by 64

Problems discussed in [RFC4966], which are not addressed by the 64 proposal:

1. Disruption of all protocols that embed IP addresses (and/or ports) in packet payloads or apply integrity mechanisms using IP addresses (and ports).

Analysis: In the case of FTP [RFC0959] this problem is addressed by the use of a FTP64 ALG [I-D.ietf-behave-ftp64] which is a workaround solution. In the case of SIP [RFC3261], no specific issue is induced by 64; the same techniques for NAT traversal can be used when a NAT64 is involved in the path (e.g., ICE [RFC5245], Hosted NAT Traversal [RFC5853], embedded SIP ALGs, etc.). The functioning of other protocols is left unaddressed. Note that the traversal of NAT64 by application embedding IP address literal is not specific to NAT64 but generic to all NAT-based solutions.

2. Inability to redirect traffic for protocols that lack de-multiplexing capabilities or are not built on top of specific transport-layer protocols for transport address translations.

   Analysis: This issue is not specific to 64 but to all NAT-based solutions.

3. Loss of information due to incompatible semantics between IPv4 and IPv6 versions of headers and protocols.

   Analysis: This issue is not specific to 64 but due to the design of IPv4 and IPv6.

4. Need for additional state and/or packet reconstruction in dealing with packet fragmentation. Otherwise, implement no support for fragments.

   Analysis: This issue is not specific to 64 but to all NAT-based solutions. [I-D.ietf-behave-v6v4-xlate-stateful] specifies how to handle fragmentation; appropriate recommendations to avoid fragmentation-related DoS attacks are proposed (e.g., limit resources to be dedicated to out of order fragments).

5. Interaction with SCTP [RFC4960] and multihoming.

   Analysis: SCTP is out of scope of 64. Only TCP and UDP transport protocols are within the scope of 64.

6. Need for the NAT64-capable device to act as proxy for correspondent node when IPv6 node is mobile, with consequent restrictions on mobility.

   Analysis: This is not specific to NAT64 but to all NAT flavors. Refer to [I-D.haddad-mext-nat64-mobility-harmful] for an early analysis on mobility complications encountered

when NAT64 is involved.

7. Inability to handle multicast traffic.

   Analysis: This problem is not addressed by the current 64 specifications.

8. Scalability concerns together with introduction of a single point of failure and a security attack nexus.

   Analysis: This is not specific to NAT64 but to all stateful NAT flavors.

9. Creation of a DoS (Denial of Service) threat relating to exhaustion of memory and address/port pool resources on the translator.

   Analysis: This specific DoS concern on Page 6 of [RFC4966] is under a DNS-ALG heading in that document, and refers to NAT-PT's creation of NAT mapping state when a DNS query occurred. With the new IPv6-IPv4 translation mechanisms, DNS queries do not create any mapping state. Thus, this concern is fully eliminated with the new IPv6-IPv4 translation mechanisms.

10. Restricted validity of translated DNS records: a translated record may be forwarded to an application that cannot use it.

    Analysis: If a node on the IPv4 side forwards the address of the other endpoint to a node which cannot reach the NAT box or is not covered under the endpoint-independent constraint of NAT, then the new node will not be able to initiate a successful session.

    Actually, this is not a limitation of 64 (or even NAT-PT) but a deployment context where shared IPv4 addresses managed by the NAT64 are not globally reachable. The same limitation can be encountered when referrals (even without any NAT in the path) include reachability information with limited reachability scope (See [I-D.carpenter-behave-referral-object] for more discussion about scope-related issues).

11. Unless UDP encapsulation is used for IPsec [RFC3948], traffic using IPsec AH (Authentication Header), in transport and tunnel mode, and IPsec ESP (Encapsulating Security Payload), in transport mode, is unable to be carried through NAT-PT without terminating the security associations on the NAT-PT, due to their usage of cryptographic integrity protection.

Analysis: This is not specific to NAT64 but to all NAT flavours.

12. Address selection issues when either the internal or external hosts implement both IPv4 and IPv6.

    Analysis: This is out of scope of 64 since Scenarios 1 and 5 of [I-D.ietf-behave-v6v4-framework] assume IPv6-only hosts.

    Therefore this issue is not resolved and mitigation techniques outside the 64 need to be used. These techniques may allow to offload NAT64 resources and prefer native communications which do not involve address family translation. Avoiding NAT devices in the path is encouraged for mobile nodes in order to save power consumption due to keepalive messages which are required to maintain NAT states ("always-on" services). An in-depth discussion can be found in [I-D.wing-behave-dns64-config].

2.2. Problems Addressed by 64

Problems, identified in [RFC4966], which are adequately addressed by the 64 proposal:

1. Constraints on network topology (as it relates to DNS-ALG; see Section 3.1 of [RFC4966]).

   Analysis: This issue has mitigated severity as the DNS is separated from the NAT functionality. Nevertheless, a minimal coordination may be required to ensure that the NAT64 to be crossed (the one to which the IPv4-Converted IPv6 address returned to a requesting host) must be in the path and has also sufficient resources to handle received traffic.

2. Inappropriate translation of responses to A queries from IPv6 nodes.

   Analysis: DNS64 [I-D.ietf-behave-dns64] does not resolve A queries.

3. Address selection issues and resource consumption in a DNS-ALG with multi-addressed nodes.

   Analysis: Since the DNS-ALG is not there and communications initiated from the IPv4 side are not supported, there is no need to maintain temporary states in anticipation of connections.

4. Limitations on DNS security capabilities when using a DNS-ALG.

   Analysis: A DNSSEC validating stub resolver behind a DNS64 in server mode is not supported. Therefore if a host wants to do its own DNSSEC validation, and it wants to use a NAT64, the host has to also perform its own DNS64 synthesis. Refer to Section 3 of[I-D.ietf-behave-dns64] for more details.

5. Creation of a DoS (Denial of Service) threat relating to exhaustion of memory and address/port pool resources on the translator.

   Analysis: This specific DoS concern on Page 6 of [RFC4966] is under a DNS-ALG heading in that document, and refers to NAT-PT's creation of NAT mapping state when a DNS query occurred. With the new IPv6-IPv4 translation mechanisms, DNS queries do not create any mapping state in the NAT64. Thus, this concern is fully eliminated in 64.

6. Requirement for applications to use keepalive mechanisms to workaround connectivity issues caused by premature timeout for session table and BIB entries.

   Analysis: Since NAT64 follows some of the [RFC4787], [RFC5382] and [RFC5508] requirements, there is a high lower bound for the lifetime of sessions. In NAT-PT this was unknown and applications needed to assume the worst case. For instance, in NAT64, the lifetime for a TCP session is approximately 2 hours, so not much keep-alive signalling overhead is needed.

   Application clients (e.g., VPN clients) are not aware of the timer configured in the NAT device. For unmanaged services, a conservative approach would be adopted by applications which issue frequent keepalive messages to be sure that an active mapping is still be maintained by any involved NAT64 device even if the NAT64 complies with TCP/UDP/ICMP BEHAVE WG specifications.

   Note that keepalive messages may be issued by applications to ensure that an active entry is maintained by a firewall, with or without a NAT in the path, which is located in the boundaries of a local domain.

7. Lack of address mapping persistence: Some applications require address retention between sessions. The user traffic will be disrupted if a different mapping is used. The use of the DNS-ALG to create address mappings with limited lifetimes means that applications must start using the address shortly after the

mapping is created, as well as keep it alive once they start using it.

Analysis: In the context of 64, the external IPv4 address (representing the IPv6 host in the IPv4 network) is assigned by the NAT64 machinery and not the DNS64 function. Address persistence can be therefore easily ensured by the NAT64 function (which complies with BEHAVE NAT recommendations). Address persistence should be guaranteed for both dynamic and static bindings.

In the IPv6 side of the NAT64, the same IPv6 address is used to represent an IPv4 host; no issue about address persistence is raised in IPv6 network.

3. Conclusions

The above analysis of the solutions provided by the 64 proposal shows that the majority of the problems that are not directly related to the decoupling of NAT and DNS remain unaddressed. Some of these problems are not specific to 64 but are generic to all NAT-based solutions.

This points to several shortcomings of 64 proposal which must be addressed if the future network deployments have to move reliably towards 64 as a solution to IPv6-IPv4 interconnection.

Some of the issues, as pointed out in [RFC4966], have possible solutions. However these solutions will require significant updates to the 64 proposal, increasing its complexity.

The following table summarizes the conclusions based on the analysis of 64 proposal.

Issue	NAT-PT Specific	Exists in NAT64	DNS ALG Specific	Generic NAT	Can be solved?
Protocols embedding addresses	No	Yes	No	Yes	Yes
Protocols without demux capability	No	Yes	No	Yes	No
Binding state decay	No	Yes	No	Yes	No
Loss of information	No	Yes	No	No	No
Fragmentation	No	No	No	Yes	Yes
SCTP and Multihoming interaction	No	Yes	No	Yes	Yes
Proxy corresp node for MIPv6	Yes	Yes	No	Yes	??
Multicast	No	Yes	No	Yes	Yes
Topology constraints with DNS-ALG	Yes	No	Yes	No	Yes
Scale and Single point of failure	No	Yes	No	Yes	Yes
Lack of address persistence	No	Yes	No	Yes	No
DoS attacks	No	Yes	No	Yes	Yes

Address selection issues with Dual stack hosts	Yes	No	Yes	No	Yes
Non-global validity of Translated RR records	Yes	No	Yes	Yes	Yes
Incorrect translation of A responses	Yes	No	Yes	No	Yes
DNS-ALG and Multi-addressed nodes	No	Yes	No	Yes	Yes
DNSSEC limitations	No	Yes	No	Yes	Yes

Table 1: Summary of NAT64 analysis

4. IANA Considerations

This document makes no request of IANA.

Note to RFC Editor: this section may be removed on publication as an RFC.

5. Security Considerations

This document does not specify any new protocol or architecture. It only analyses how BEHAVE WG 64 documents mitigate concerns raised in [RFC4966] and which ones are still unaddressed.

6. Acknowledgements

Many thanks to Marcelo Bagnulo for his comments.

7. References

7.1. Normative References

[I-D.ietf-behave-dns64]
Bagnulo, M., Sullivan, A., Matthews, P., and I. Beijnum, "DNS64: DNS extensions for Network Address Translation from IPv6 Clients to IPv4 Servers", draft-ietf-behave-dns64-11 (work in progress), October 2010.

[I-D.ietf-behave-v6v4-xlate-stateful]
Bagnulo, M., Matthews, P., and I. Beijnum, "Stateful NAT64: Network Address and Protocol Translation from IPv6 Clients to IPv4 Servers", draft-ietf-behave-v6v4-xlate-stateful-12 (work in progress), July 2010.

[RFC0959] Postel, J. and J. Reynolds, "File Transfer Protocol", STD 9, RFC 959, October 1985.

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

[RFC4787] Audet, F. and C. Jennings, "Network Address Translation (NAT) Behavioral Requirements for Unicast UDP", BCP 127, RFC 4787, January 2007.

[RFC4960] Stewart, R., "Stream Control Transmission Protocol", RFC 4960, September 2007.

[RFC4966] Aoun, C. and E. Davies, "Reasons to Move the Network Address Translator - Protocol Translator (NAT-PT) to Historic Status", RFC 4966, July 2007.

[RFC5382] Guha, S., Biswas, K., Ford, B., Sivakumar, S., and P. Srisuresh, "NAT Behavioral Requirements for TCP", BCP 142, RFC 5382, October 2008.

[RFC5508] Srisuresh, P., Ford, B., Sivakumar, S., and S. Guha, "NAT Behavioral Requirements for ICMP", BCP 148, RFC 5508, April 2009.

[RFC6052] Bao, C., Huitema, C., Bagnulo, M., Boucadair, M., and X. Li, "IPv6 Addressing of IPv4/IPv6 Translators", RFC 6052, October 2010.

7.2. Informative References

[I-D.carpenter-behave-referral-object]
Carpenter, B., Boucadair, M., Halpern, J., Jiang, S., and K. Moore, "A Generic Referral Object for Internet Entities", draft-carpenter-behave-referral-object-01 (work in progress), October 2009.

[I-D.haddad-mext-nat64-mobility-harmful]
Haddad, W. and C. Perkins, "A Note on NAT64 Interaction with Mobile IPv6", draft-haddad-mext-nat64-mobility-harmful-01 (work in progress), April 2010.

[I-D.ietf-behave-ftp64]
Beijnum, I., "An FTP ALG for IPv6-to-IPv4 translation", draft-ietf-behave-ftp64-06 (work in progress), November 2010.

[I-D.ietf-behave-v6v4-framework]
Baker, F., Li, X., Bao, C., and K. Yin, "Framework for IPv4/IPv6 Translation", draft-ietf-behave-v6v4-framework-10 (work in progress), August 2010.

[I-D.wing-behave-dns64-config]
Wing, D., "DNS64 Resolvers and Dual-Stack Hosts", draft-wing-behave-dns64-config-02 (work in progress), February 2010.

[RFC2766] Tsirtsis, G. and P. Srisuresh, "Network Address Translation - Protocol Translation (NAT-PT)", RFC 2766, February 2000.

[RFC3261] Rosenberg, J., Schulzrinne, H., Camarillo, G., Johnston, A., Peterson, J., Sparks, R., Handley, M., and E. Schooler, "SIP: Session Initiation Protocol", RFC 3261, June 2002.

[RFC3948] Huttunen, A., Swander, B., Volpe, V., DiBurro, L., and M. Stenberg, "UDP Encapsulation of IPsec ESP Packets", RFC 3948, January 2005.

[RFC5245] Rosenberg, J., "Interactive Connectivity Establishment (ICE): A Protocol for Network Address Translator (NAT) Traversal for Offer/Answer Protocols", RFC 5245, April 2010.

[RFC5853] Hautakorpi, J., Camarillo, G., Penfield, R., Hawrylyshen, A., and M. Bhatia, "Requirements from Session Initiation Protocol (SIP) Session Border Control (SBC) Deployments", RFC 5853, April 2010.

Authors' Addresses

Reinaldo Penno
Juniper Networks
1194 N Mathilda Avenue
Sunnyvale, California 94089
USA

Email: rpenno@juniper.net

Tarun Saxena
Cisco Systems

Email: tasaxena@cisco.com

Mohamed Boucadair
France Telecom
Rennes 35000
France

Email: mohamed.boucadair@orange-ftgroup.com

Senthil Sivakumar
Cisco Systems
7100-8 Kit Creek Road
Research Triangle Park, North Carolina 27709
USA

Email: ssenthil@cisco.com