IPv6 Operations
Internet-Draft
Intended status: Standards Track
Expires: March 19, 2015

T. Anderson
Redpill Linpro

September 15, 2014

SIIT-DC: Dual Translation Mode
draft-anderson-v6ops-siit-dc-2xlat-00

Abstract

This document describes an extension of the SIIT-DC [I-D.anderson-v6ops-siit-dc] architecture, which allows applications that are incompatible with IPv6, SIIT-DC and/or Network Address Translation in general to operate correctly in an SIIT-DC environment. This is accomplished by introducing a new component called a SIIT-DC Host Agent, which reverses the translations made by an SIIT-DC Gateway. The application is thus provided with seemingly native IPv4 connectivity.

The reader is expected to be familiar with the SIIT-DC architecture described in [I-D.anderson-v6ops-siit-dc].

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 March 19, 2015.

Copyright Notice

Table of Contents

1. Introduction

SIIT-DC [I-D.anderson-v6ops-siit-dc] describes an architecture where IPv4-only users can access IPv6-only services through a stateless translation gateway. However, this only works for applications that are compatible with Network Address Translation (NAT), due to the fact that the SIIT-DC Gateway will rewrite the addresses in the IP header as part of the translation process. SIIT-DC will also fail to work correctly for applications that make use of legacy IPv4-only socket calls.

This document remedies this problem by defining an extension to SIIT-DC. Translations performed by the SIIT-DC Gateway will also be done in reverse by an SIIT-DC Host Agent running on the server. The resulting IPv4 packets are then passed to the application. This way, the application will be able to use legacy IPv4-only socket calls and /or include references to its own IPv4 address in the application payload, while maintaining correct operation.

The approach is heavily inspired by and very similar to 464XLAT [RFC6877]. The SIIT-DC Host Agent described in this document is almost identical to the CLAT component in 464XLAT, except for the

fact that it will be located on a server, rather than on the customer-side node. Furthermore, an SIIT-DC Host Agent uses statically configured public IP addresses, whereas a 464XLAT CLAT uses a dynamic IPv6 address and a private IPv4 address. The SIIT-DC Gateway described in [I-D.anderson-v6ops-siit-dc] is used instead of the PLAT described by 464XLAT.

2. Terminology

This document makes use of the following terms:

IPv4 Service Address A public IPv4 address with which IPv4-only clients will communicate. This communication will be translated to IPv6 by the SIIT-DC Gateway.

IPv6 Service Address A public IPv6 address assigned to a server or application in the IPv6 network. IPv6-only and dual stacked clients communicates with this address directly without invoking SIIT-DC. IPv4-only clients also communicate with this address through the SIIT-DC Gateway and via an IPv4 Service Address.

SIIT-DC Host Agent A logical function very similar to an SIIT-DC Gateway that resides on a server and provides virtual IPv4 connectivity to applications, by performing [I-D.anderson-v6ops-siit-dc] translation on packets passing through it. See Section 3.

SIIT-DC Gateway A device or a logical function that translates between IPv4 and IPv6 in accordance with [I-D.anderson-v6ops-siit-dc].

Static Address Mapping A bi-directional mapping between an IPv4 Service Address and an IPv6 Service Address configured in the SIIT-DC Gateway. When translating between IPv4 and IPv6, the SIIT-DC Gateway changes the adddress fields in the translated packet’s IP header according to any matching Static Address Mapping.

Translation Prefix An IPv6 prefix into which the entire IPv4 address space is mapped. This prefix is routed to the SIIT-DC Gateway’s IPv6 interface. It is either an Network-Specific Prefix or a Well-Known Prefix as specified in [RFC6052]. When translating between IPv4 and IPv6, the SIIT-DC Gateway prepends or strips the Translation Prefix from the address fields in the translated packet’s IP header, unless a Static Address Mapping exists for the IP address in question.

3. SIIT-DC Host Agent Specification

The SIIT-DC Host Agent runs on the servers hosting application which do not work correctly with the SIIT-DC architecture as specified by [I-D.anderson-v6ops-siit-dc]. Its task is the performing the exact same packet translation as the SIIT-DC Gateway, only in reverse. It therefore shares the same implementation requirements as the SIIT-DC Gateway defined in Section 4 of [I-D.anderson-v6ops-siit-dc], with one exception: The SIIT-DC Host Agent is not required to support configuring an arbitrary number of Static Address Mappings, but it must support at least one.

The SIIT-DC Host Agent must be configured with a Static Address Mapping that corresponds exactly with the same mapping found on the SIIT-DC Gateway. The IPv4 address of the Static Address Mapping (i.e., the IPv4 Service Address) must be configured on a virtual network interface which applications running on the server can bind to, and the server is expected to install a default IPv4 route poining to this virtual IPv4 interface. The IPv6 address of the Static Address Mapping must be a secondary address that is routed to the server by the IPv6 network. The server must forward all packets it receives destined for this IPv6 address to the SIIT-DC Host Agent.

4. Architectural Overview

The following figure shows how an application (that is presumably incompatible with standard SIIT-DC) is being made available to the IPv4 Internet on the IPv4 address 192.0.2.4. The application will be able to know that this is its local address and thus be able to provide correct references to it in application payload.

The figure also shows how the same application is available over IPv6 on its IPv6 Service Address 2001:db8:12:34::3. This is included in order to illustrate how native IPv6 connectivity is not impacted by the SIIT-DC Host Agent, and also to illustrate how the address assigned to the SIIT-DC Host Agent (2001:db8:12:34::4) is separate from the primary IPv6 address of the server. It is however important to note that the application in question does not have to be dual-stack capable at all. IPv4-only applications would also be able to operate behind a SIIT-DC Host Agent in the exact same manner.

Note that the figure below could be considered a more detailed view of Customer A’s FTP server from the example topology figure in Appendix A of I-D.anderson-v6ops-siit-dc [I-D.anderson-v6ops-siit-dc]. Both figures intentionally use the exact same example IP addresses and prefixes.

SIIT-DC Host Agent Architecture

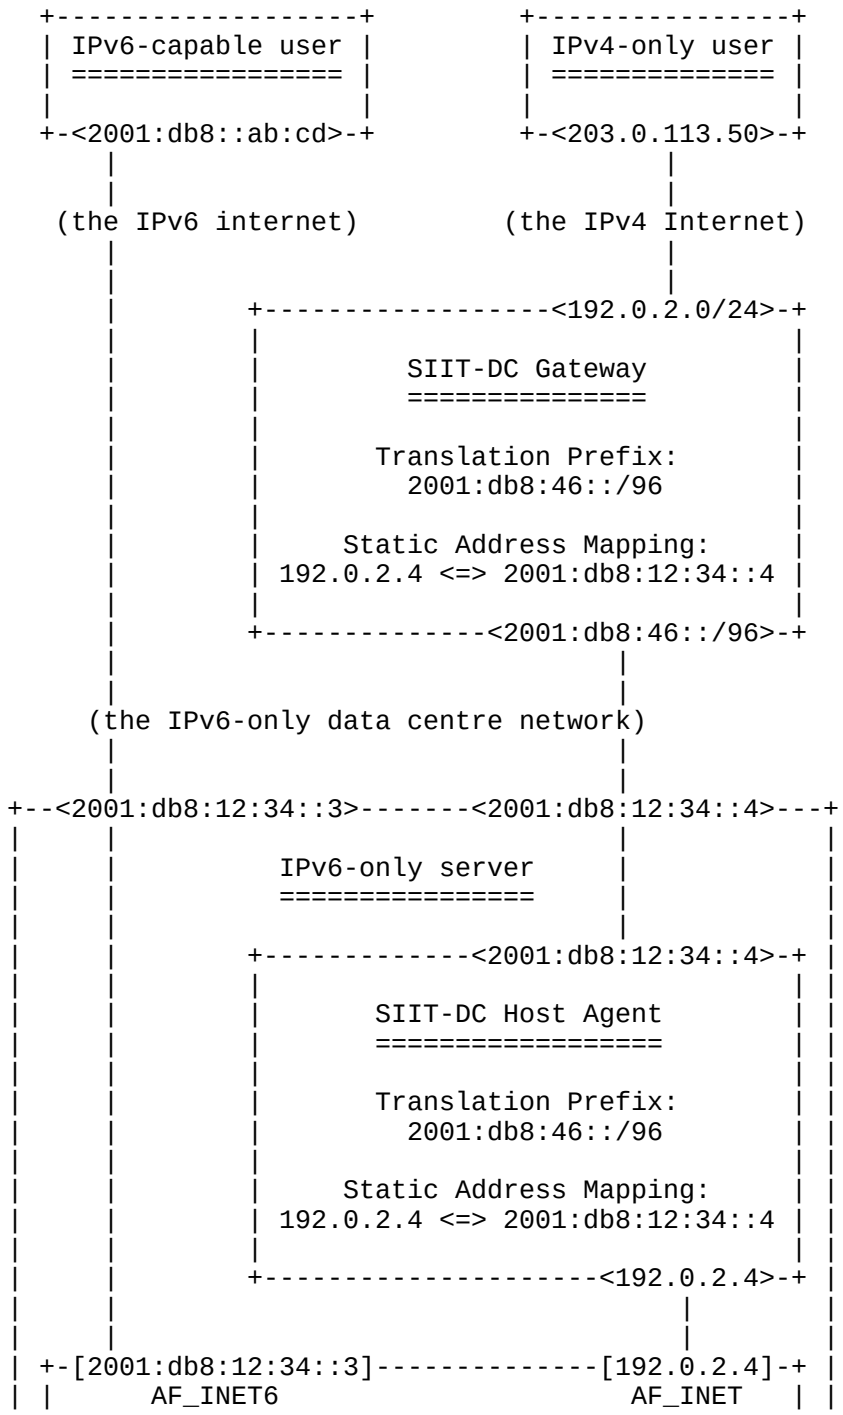
IPv6-capable user
<2001:db8::ab:cd>
IPv4-only user
<203.0.113.50>
(the IPv6 internet)
(the IPv4 Internet)
<192.0.2.0/24>
SIIT-DC Gateway
Translation Prefix:
2001:db8:46::/96
Static Address Mapping:
192.0.2.4 <=> 2001:db8:12:34::4
<2001:db8:46::/96>
(the IPv6-only data centre network)
<2001:db8:12:34::3>
<2001:db8:12:34::4>
IPv6-only server
<2001:db8:12:34::4>
SIIT-DC Host Agent
Translation Prefix:
2001:db8:46::/96
Static Address Mapping:
192.0.2.4 <=> 2001:db8:12:34::4
<192.0.2.4>
[2001:db8:12:34::3]
[192.0.2.4]
AF_INET6
AF_INET

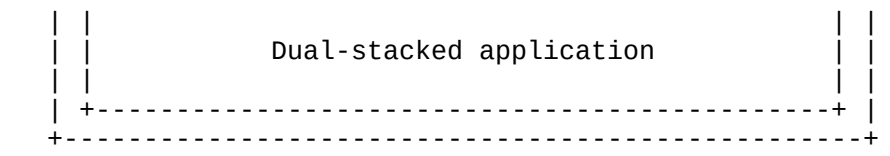

Figure 1

5. Deployment Considerations

5.1. IPv6 Path MTU

The IPv6 Path MTU between the SIIT-DC Host Agent and the SIIT-DC Gateway will typically be larger than the default value defined in Section 4 of [RFC6145] (1280), as it will typically contained within a single administrative domain. Therefore, it is recommended that the IPv6 Path MTU configured in the SIIT-DC Host Agent is raised accordingly. It is RECOMMENDED that the SIIT-DC Host Agent and the SIIT-DC Gateway use identical configured IPv6 Path MTU values.

5.2. IPv4 MTU

In order to avoid fragmentation, it is RECOMMENDED that the virtual IPv4 interface is configured with an MTU value identical to the configured IPv6 Path MTU - 20. This ensures that the application may do its part in avoiding IP-level fragmentation from occurring, e.g., by segmenting/fragmenting outbound packets at the application layer, and advertising the maximum size its peer may use for inbound packets (e.g., through the use of the TCP MSS option).

6. Acknowledgements

The author would like to especially thank the authors of 464XLAT [RFC6877]: Masataka Mawatari, Masanobu Kawashima, and Cameron Byrne. The architecture described by this document is merely an adaptation of their work to a data centre environment, and could not have happened without them.

The author would like also to thank the following individuals for their contributions, suggestions, corrections, and criticisms: Fred Baker, Tobias Brox, [YOUR NAME GOES HERE].

7. 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 [RFC2119].

8. IANA Considerations

This draft makes no request of the IANA. The RFC Editor may remove this section prior to publication.

9. Security Considerations

This section discusses security considerations specific to the use of a SIIT-DC Host Agent. See the Security Considerations in I-D .anderson-v6ops-siit-dc [I-D.anderson-v6ops-siit-dc] for additional security considerations applicable to the SIIT-DC architecture in general.

9.1. Address Spoofing

If the SIIT-DC Host Agent receives an IPv4 packet from the application from a different source address than the one it has a Static Address Mapping for, the both the source and destination addresses will be rewritten according to [RFC6052]. After undergoing the reverse translation in the SIIT-DC Gateway, the resulting IPv4 packet routed to the IPv4 network will have a spoofed IPv4 source address. The SIIT-DC Host Agent should therefore ensure that ingress filtering (cf. BCP38 [RFC2827]) is used on the SIIT-DC Host Agent’s IPv4 interface, so that such packets are immediately discarded.

If the SIIT-DC Host Agent receives an IPv6 packet with both the source and destination address equal to the one it has a Static Address Mapping for, the resulting packet would appear to the application as locally generated, as both the source address and the destination address will be the same address as the one configured on the virtual IPv4 interface. This could trick the application into thinking this packet came from a trusted source, and give elevated privileges accordingly. To prevent this, the SIIT-DC Host Agent should discard any received IPv6 packets that have a source address that is equal either to either the IPv4 (after undergoing [RFC6052] translation) or the IPv6 address in the Static Address Mapping.

10. References

10.1. Normative References

[I-D.anderson-v6ops-siit-dc]
Anderson, T., "SIIT-DC: Stateless IP/ICMP Translation in IPv6 Data Centre Environments", draft-anderson-v6ops-siit-dc-00 (work in progress), September 2014.

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

10.2. Informative References

[RFC2827] Ferguson, P. and D. Senie, "Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing", BCP 38, RFC 2827, May 2000.

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

[RFC6145] Li, X., Bao, C., and F. Baker, "IP/ICMP Translation Algorithm", RFC 6145, April 2011.

[RFC6877] Mawatari, M., Kawashima, M., and C. Byrne, "464XLAT: Combination of Stateful and Stateless Translation", RFC 6877, April 2013.

Author’s Address

Tore Anderson
Redpill Linpro
Vitaminveien 1A
0485 Oslo
NORWAY

Phone: +47 959 31 212
Email: tore@redpill-linpro.com