SIIT-DC (Stateless IP/ICMP Translation for IPv6 Data Center Environments)

Overview

SIIT-DC defines a network deployment model to allow legacy IPv4-only networks to establish connections to and from IPv6-only networks, in other words, to allow connections between single-stack IPv4-only and IPv6-only networks. Thus, SIIT-DC facilitates the deployment of IPv6-only network infrastructures such as datacenters, servers and local area networks.
This transition mechanism is standardized in RFC7755 (2016)

SIIT-DC was especially designed:

- For deploying IPv6-only datacenters.
- As an efficient solution to public IPv4 address exhaustion.
- To simplify and avoid inefficiency in the deployment of dual-stack scenarios by considering a single-stack approach.

IPv6 - Mecanismo de Transición SIIT-DC (RFC7755)
Arquitectura

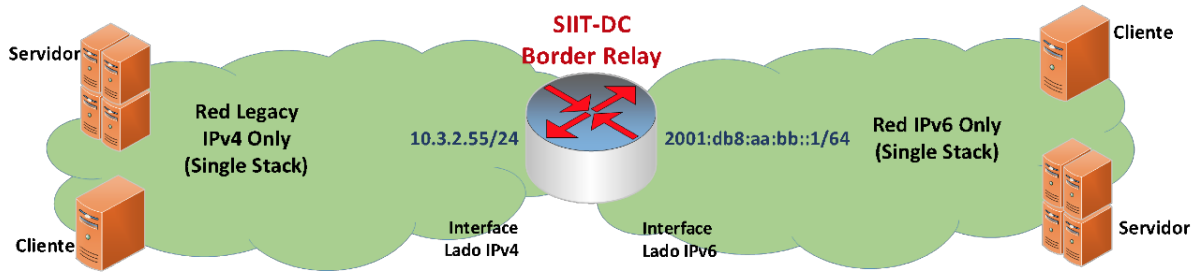

SIIT-DC BR - Tabla EAM

Prefijo IPv4	Prefijo IPv6
192.0.2.1	2001:db8:a:b::cafe
192.0.2.2/32	2001:db8:cc:dd::ee33/128
192.0.2.16/28	2001:db8:cc35:cc35::/124
192.0.2.128/26	2001:db8:cafa::/64
192.0.2.192/28	2001:db8:efff:caa0::/96
192.0.2.224/31	2001:db8:7733:7733::/127
10.0.1.0/24	64:ff9b::/96

IPv4 Suffix (32-M) ≤ IPv6 Suffix (128-Len)

While the dual-stack transition mechanism was initially widely deployed for transitioning legacy IPv4 networks to the IPv6 protocol, it is true that a dual-stack approach brings with it a series of disadvantages compared to the single-stack option, such as:

- Dual-stack deployment presents increased complexity due to the need to manage a dual network, dual routing, dual processing and dual planning.
- The dual-stack approach does not in any way help or solve the problem of IPv4 address space depletion.

Technical Characteristics

The single-stack IPv6-only approach allows datacenter operators to prepare for future scenarios, and SIIT-DC is a step in this direction. Additionally, SIIT-DC allows maintaining fully operational legacy IPv4-only networks and interconnecting them with new IPv6-only networks, all this without the need for a dual-stack operation.

With SIIT-DC, neither new IPv6-only hosts and servers nor legacy IPv4-only hosts and servers require any special adjustments or modifications, as they will operate in single-stack mode. An operator can transition to IPv6-only transport networks, which will also not require any adjustments or special support. Legacy IPv4 networks and IPv6-only networks are fully compatible with SIIT-DC.

SIIT-DC simply introduces a new functional element to the network: the SIIT-DC Border Relay (SIIT-DC BR or simply BR). The operation of the BR component is stateless; it translates IPv4 to IPv6 and vice versa using a translation algorithm and tables for mapping the translations.

One of the notable aspects of SIIT-DC is the stateless operation of the BR. The BR does not maintain any state associated with individual connections, devices or flows. The BR operates very similarly to an IPv4 or IPv6 router, but it allows IPv4 packet traffic to and from IPv6 interfaces. That said, the only limiting factor is the packet-per-second capacity of the hardware where the BR is functionally deployed (associated with network interface, CPU and RAM resources). One or multiple BR routers can be used in the operator's network for efficient traffic distribution purposes using ECMP (RFC2991). In this sense, SIIT-DC supports asymmetric routing and high availability (HA) architectures in the network when necessary.

SIIT-DC uses the translation algorithm defined in RFC6052. This algorithm allows 1:1 translations from/to IPv4 to/from IPv6. Therefore, it allows mapping an IPv4 address to an IPv6 address and also mapping a complete IPv4 address block to an IPv6 prefix. The translation algorithm always translates both the source and the destination IP addresses of the IP traffic from IPv4 into IPv6 and vice versa.

Explicit Address Mapping (EAM). In SIIT-DC translation from IPv4 to IPv6 and, by analogy, vice versa, the BR translates an IP packet's two addresses (source and destination) as specified in the EAM algorithm defined in RFC7757 (2016). The EAM algorithm is based on the use of IPv6 translation prefixes to map IPv4 address blocks. On the IPv6 side, packets forwarded to IPv4 networks must be routed to the BR's IPv6 interface using these translation prefixes. On the IPv4 side, packets forwarded to IPv6 networks must be routed to the BR's IPv4 interface, also using the corresponding translation prefixes.

The translation performed by the BR does not modify IP packet payload, it merely introduces changes in the IPv4 and IPv6 headers.

SIIT-DC is especially designed for traditional client/server applications such as HTTP, HTTPS, SSH, SNMP, NTP, among many others. It is perfectly possible to combine SIIT-DC with DNS64 to allow IPv6-only applications to initiate connections to IPv4-only servers.

Considerations regarding the choice of translation prefixes:

- IPv6 prefixes from the operator's own IPv6 address space. Define and reserve prefixes for SIIT-DC as part of the addressing plan.
- Use well-known prefix (WKP) 64:ff9b::/96. In this case, the disadvantage is that a single WKP exists.
- Using a /96 translation prefix is very useful and recommended because it allows the literal notation PrefixIPv6::X.Y.Z.W.

Advantages:

- Promotes and encourages the deployment of IPv6-only networks and datacenters.
- When deploying large datacenters with an IPv6-only (single-stack) approach, the scarcity of public IPv4 addresses is no longer an issue. In this sense, in these datacenters, IPv6-only servers that do not need to connect to the IPv4 Internet will not require IPv4 addressing or edge translation. Similarly, the use of IPv6-only optimizes network connections and addressing within the datacenters themselves.
- Fully compatible with IPv4 and IPv6 stacks. Does not require any changes in IPv4 and IPv6 hosts or servers.
- Allows performing tasks such as geolocation, logging, abuse handling, and so on.
- Even when operating in IPv6-only mode, the value-added services of an ISP (DNS, NTP, SYSLOG, RADIUS, TACACS, etc.) can serve legacy IPv4-only networks.
- SIIT-DC fully matches TCP and UDP protocols and applications based on the client-server model.
- SIIT-DC BRs are highly efficient as their operation is stateless. In addition, they allow asymmetric routing and achieving high availability.
- IPv4-only networks, servers, services and applications that are difficult to migrate and/or that cannot be upgraded to IPv6 will be able to continue operating and will even be able to connect to IPv6-only clients and servers.
- EAM tables in the BR are quick and easy to configure.
- Does not require updates or adjustments in client CPEs.

Disadvantages

- Certain issues with fragmentation in IPv4 and limited use of fragmentation in IPv6. Example: when translated, a 1500-byte IPv4-packet will result in a 1520-byte IPv6 packet.
- Issues with protocols that embed IP addresses as part of their payload, e.g., SIP and FTP. Use of these protocols will be problematic if they use NAPT44.

Translation in SIIT-DC BR based on EAM Table

IPv6 - Mecanismo de Transición SIIT-DC (RFC7755)
Ilustración del Proceso de Traducción

SIIT-DC BR - Tabla EAM

Prefijo IPv4	Prefijo IPv6
192.0.2.1	2001:db8:a:b::cafe
192.0.2.2/32	2001:db8:cc:dd::ee33/128
192.0.2.16/28	2001:db8:cc35:cc35::/124
192.0.2.128/26	2001:db8:cafa::/64
192.0.2.192/28	2001:db8:efff:caa0::/96
192.0.2.224/31	2001:db8:7733:7733::/127
10.0.1.0/24	64:ff9b::/96

IPv4 Suffix (32-M) ≤ IPv6 Suffix (128-Len)

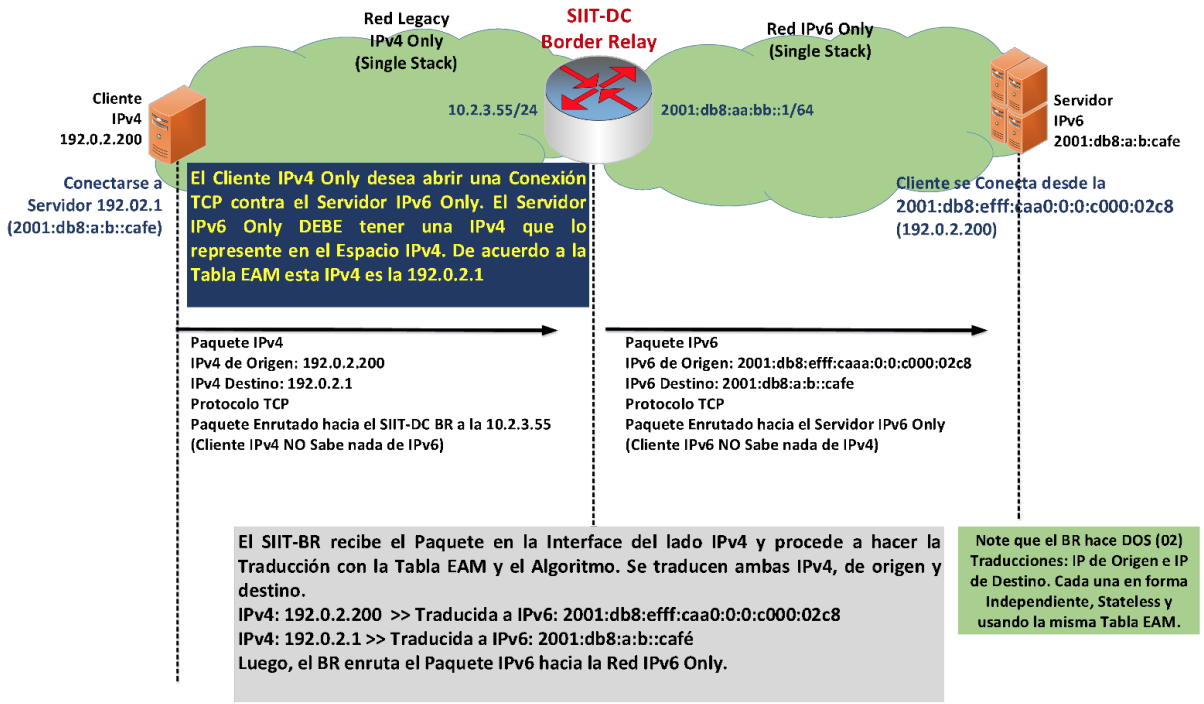

Translation in SIIT-DC BR as an IPv4aaS mechanism

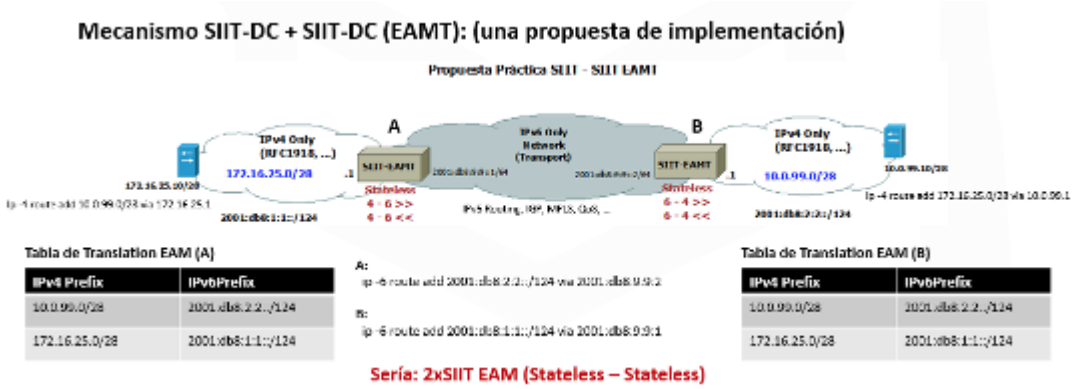

Double stateless-stateless SIIT-DC translation in two separate BRs can be applied to allow transporting IPv4 traffic over IPv6-only networks and offering IPv4aaS.