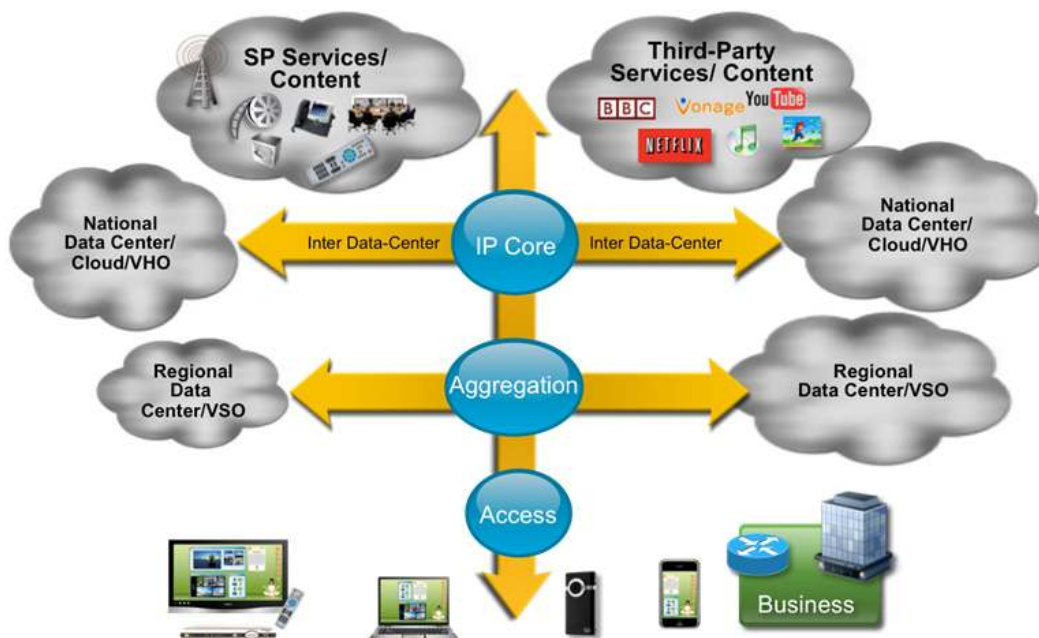


Carrier-Grade IPv6: Mapping Address and Port-Translation Technical Brief

Introduction

IPv6 deployment is today's new normal. The proliferation of IPv6 in service provider networks and allocation to end users is a reality now, but it is not always as simple as just flipping a switch. Many dependencies and considerations are involved as operators choose their path to IPv6. Each place in the network (Figure 1) has different characteristics that dictate how easy or hard it will be to turn on IPv6 and eventually turn off IPv4.

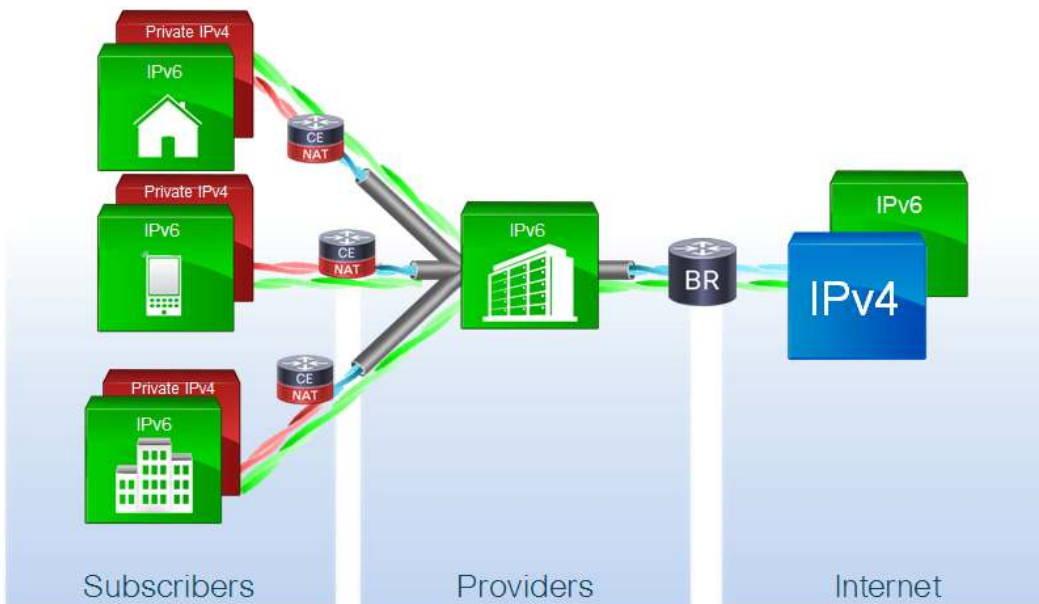
Figure 1. Example of a Service Provider Network



Core networks largely evolved to be dual-stack-enabled in almost all operator networks. This means that either native IPv6 traffic originating from end users or from the content owners, or IPv4 traffic, can traverse the core network without requiring translation or tunneling. Moving to the edge of the operator network, the situation starts to differ depending on the equipment deployed and how far the service provider has progressed on the path to IPv6 enablement. In large part, the stage operators find themselves in depends on what is enabled on their access equipment and that of their end users. In cases where a single stack is employed at any point in the access or at the end-user stage, either a translation such as Network Address Translation (NAT) or a tunneling mechanism such as IPv6 rapid deployment (6rd) would be used.

After several years of IPv6 deployment, service providers are starting to transition to IPv6-only single-stack architectures. However, existing IPv4 communications persist from the end user to the content providers, and must operate across that single-stack IPv6 domain. Figure 2 depicts a scenario of this type.

Figure 2. Example of an IPv6 Single-Stack Provider Network



The use of an IPv4–IPv6 translation-based solution is an optimal way to address these requirements, particularly in combination with stateless translation or encapsulation techniques that seek to minimize complexities. Mapping Address and Port (MAP), is such a solution, and one that builds on existing stateless IPv4 and IPv6 address translation techniques specified in the [RFC6052](#), [RFC6144](#), and [RFC6145](#) standards. MAP has two variants: MAP encapsulation (MAP-E) and MAP translation (MAP-T). The difference between the two options is evident in their names. MAP-E uses IPv6 to encapsulate and decapsulate IPv4 traffic, whereas MAP-T uses NAT64 to translate IPv4 to IPv6 and back. This document describes MAP-T and recent interoperability and proof-of-concept testing done in relation to MAP-T. While MAP-T is covered here, MAP-E performs similarly and provides the same level of benefits by enlisting a stateless algorithm. MAP-E will also be supported on the Cisco platforms mentioned in this document.

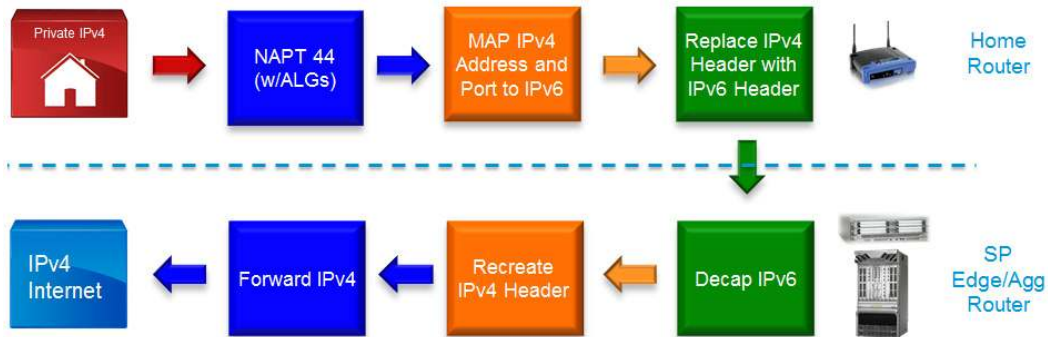
Mapping Address and Ports-Translation

The MAP-T solution presents an operator with the prospect of a full transition of a domain to IPv6-only in a manner that:

- Retains the ability for IPv4 end-hosts to communicate across the IPv6 domain with other IPv4 hosts while also permitting both individual IPv4 address assignment and address sharing
- Allows communication between IPv4-only end hosts, as well as any IPv6-enabled end hosts, to native IPv6-only servers in the domain that are using an IPv4-mapped IPv6 address
- Does not require the operation of a stateful IPv4 overlay network, nor the introduction of nonnative-IPv6 network device or server functionality
- Allows the use of IPv6 native network operations, including the ability to classify IP traffic, as well as to perform IP traffic routing optimization policies such as routing optimization, based on peering policies for Internet IPv4 destinations outside of the domain
- Extends stateless IPv4–IPv6 translation with algorithmic address and port mapping

Figure 3 gives a high-level illustration of how MAP-T works

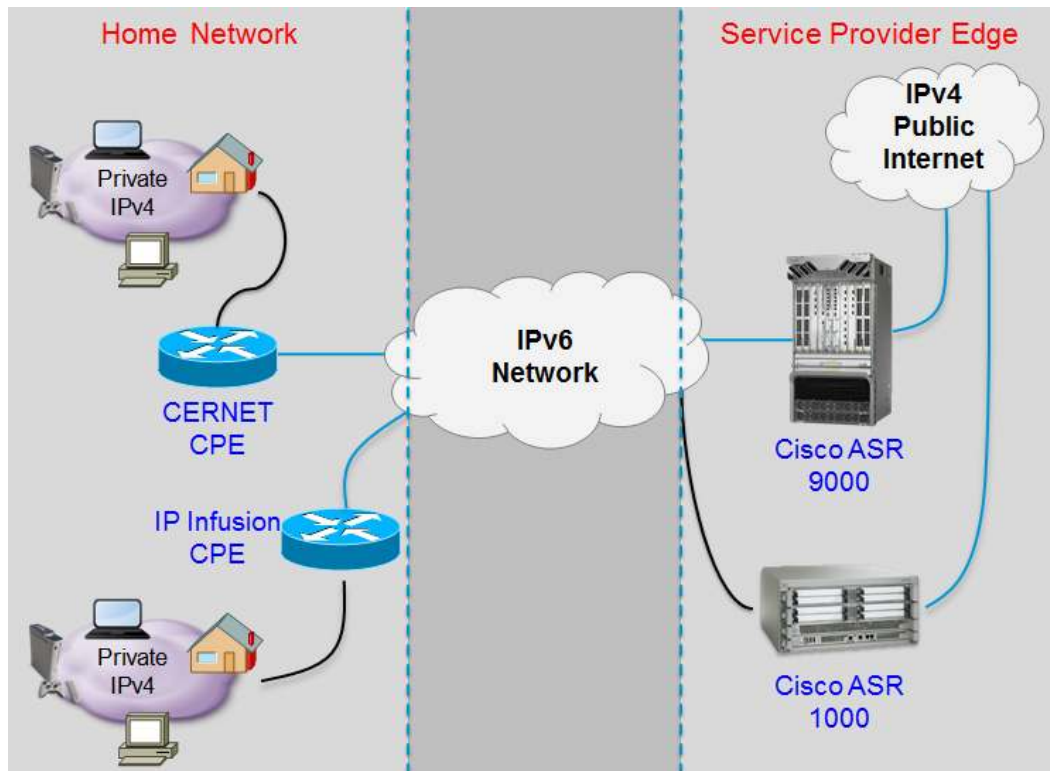
Figure 3. High-Level Packet Flow for MAP-T



Proof of Concept

Recently, Cisco conducted a series of interoperability and proof-of-concept tests using MAP-T. MAP-T requires participation by both the home router and the service provider edge and aggregation routers to form a complete end-to-end solution. Cisco collaborated with IP Infusion, Linksys, and CERNET to provide two sample home routers to pair with the Cisco ASR 9000 and the Cisco ASR 1000 Aggregation Services Routers serving as the service provider edge routers. Figure 4 gives a simple representation of the network topology that was used as part of the proof-of-concept testing.

Figure 4. Simplified Interoperability Testing Topology



This topology was used to validate the MAP-T technology and its stateless nature in terms of performance and ability to support today's commonly used applications. More specifically, the testing simulated an environment that included the customer premises equipment (CPE) of an average household, where the end user accessed commonly used applications. Table 1 summarizes the applications that were tested and the results.

Table 1. Summary of Tested Applications and Results

Application	Result	Notes
Standard web browsing	Successful	Several different sites were tested with no issues.
YouTube	Successful	YouTube and other video content were accessed with no issues.
Skype	Successful	Calls were placed and received with no issues.
FTP	Successful	Transfer of files was successfully tested.
FaceTime	Successful	Calls were placed and received with no issues.

After establishing the viability of the MAP-T solution in terms of application success, the next factor that was verified was the scalability of the solution. Generally speaking, when evaluating translation technologies, the two variables that are monitored are the total number of translations and the amount of throughput that can be done at the translation device. In the proof of concept for MAP-T, both of these variables were evaluated.

The number of translations that can be performed using MAP-T is almost unlimited due to its stateless nature. This is the advantage of stateless technologies: no state needs to be tracked or maintained at the translation device, eliminating the potential limitation of a specific number of translations. Avoiding this potential limitation not only results in a simpler implementation, but also reduces cost. Keeping track of states can be expensive in terms of operating and capital expenses. The operating expense is due to the cost of managing the state tables and the possible need to manipulate the tables. Capital expenses are based on the requirement and resulting cost of storage devices for the state tables as well as the increased cost of the translation devices. By removing the requirement to track, maintain, and store translations, the solution becomes simpler and the cost of the solution is greatly reduced.

The second variable to be considered is the throughput that translation devices can support. The amount of throughput relates directly to the number of devices and users that the translation device can support and thus affects the total cost of the solution. In a perfect world, the edge aggregation device in a service provider network would be able to implement any and all features with no effect on the performance capabilities of the aggregation device. Of course, this is not realistic as some features directly affect the performance of the aggregation device. This is especially true when dealing with features that typically are delivered on service cards, which is true for most IPv4–IPv6 translation and transition technologies. In those instances, performance capabilities are limited by the throughput of the service card. However, in the case of MAP-T on the Cisco ASR 9000 and ASR 1000, the functionality required on the aggregation device can be delivered as part of the normal packet processing pipeline. This is an especially important advantage for the ASR 9000, where the MAP packet processing is distributed to the individual line cards. For both routers, performance was validated and is summarized in Table 2. (Tests were performed with the Cisco ASR 9010 Router, using thirty-six 10-Gigabit Ethernet cards at line rate with bidirectional traffic.)

Table 2. Platform Performance Testing Results Summary

Platform	10GE Result	Platform Result
Cisco ASR 1000	9.81 Gbps (full duplex)	19.62 Gbps
Cisco ASR 9000	Line-rate (full duplex)	5.76 Tbps

Summary

MAP-T is an emerging solution for transporting private IPv4 addressing over IPv6 networks, which brings many advantages over alternative options. The advantages are primarily rooted in the stateless nature of the protocol, which reduces operating and capital expenses across the network. The Cisco ASR 9000 and ASR 1000 Series routers implement a solution that has been demonstrated with customer premises equipment based on IP Infusion, Linksys, and CERNET, confirming the advantages of the MAP-T solution.



Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: www.cisco.com/go/trademarks. Third party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)