

Technical Guide to Mapping of Address and Port (MAP)

Last Updated: January 2014

Executive Summary

Mapping of Address and Port (MAP) is an IPv6 transition and interworking technology that leverages the natural aggregation capability of the address and port space to allow IPv4 addresses to be translated or encapsulated in IPv6, without requiring a stateful translator on the service provider network. This provides a production-quality, deployable IPv6 transition mechanism, which allows service providers to share scarce IPv4 address resources, while deploying an IPv6-only provider network. Two flavors of MAP exist today—MAP-T and MAP-E -which use translation and encapsulation respectively. The key benefit of MAP is its stateless implementation on the SP router, which allows it to scale proportionally to the number of aggregated prefixes and traffic volume instead of by the number of end user connections or states. On the Cisco ASR9000, MAP is processed inline on the 2nd generation Typhoon line cards, which allows for line-rate performance.

MAP-T was introduced on the Cisco ASR9000 in 4.3.0; MAP-E was added in 4.3.1.

Goals

This whitepaper aims to provide the reader with a technical overview of *Mapping of Address and Port (MAP)*, along with its relevance and benefits as an IPv6 transition and interworking technology. Details of the port mapping algorithm, as well as both variants of the MAP implementation (MAP-E and MAP-T), will be described. This document will enable a solid understanding of MAP technology and the ability to position it effectively with service providers in regard to its benefits, relative to other alternatives (e.g., Dual-Stack Lite).

Solution Overview

Mapping of Address and Port (MAP) addresses the problem space faced by service providers in phased migrations from IPv4 to IPv6. In these migrations, there may be situations where residual IPv4 deployments in different parts of the network are required. This is in spite of IPv6 networks having become commonplace and deployed in most parts of the network. This could be the case where an IPv6-only solution cannot be offered to end users due to various limitations, which may include specific protocols not working across address family translation techniques such as NAT64, or in the case whereby the subscriber end stations, which are beyond the administrative control of the service provider, do not support IPv6.

Mapping of Address and Port (MAP) exploits the natural aggregation model, which has allowed IP networks to scale to the size of the Internet today. This aggregation model was made possible through the use of Classless Inter-Domain Routing (CIDR) and route summarization across the Internet backbone. Via an organized allocation of the address space represented by both IP addresses and L4 ports, MAP extends the CIDR concept into a combined 48-bit address space (32 bits IPv4 + 16 bits TCP/UDP).

MAP leverages an IPv6-address-to-IPv4-address mapping, as well as port mapping algorithms that use specific bits in the IPv6 address space to represent both IPv4 addresses and L4 ports. In MAP-T, the original L3 and L4 information is available to regular IPv6 data plane devices via the IPv6 header. In comparison, MAP-E encapsulates the entire original IPv4 datagram with an IPv6 header with the IPv6 source and destination corresponding to the tunnel endpoints, so that no IPv4 header information is directly exposed to the IPv6 data plane.

NAT44 is a component of the MAP solution, but the NAT44 in MAP differs from traditional NAT44 or NAT444 deployments in that instead of assigning a public IPv4 address range to each CPE for translation (in the case of NAT), or a single public IPv4 address for translation (in the case of PAT) to each CPE, it extends the granularity beyond a single public IPv4 address, by being able to assign a port range to each of the CPEs sharing the same IPv4 public address. This unique address and port range combination is then translated into the IPv6 address space when transitioning into the IPv6 domain using the MAP CPE. The MAP algorithm still retains the ability to assign the full IPv4 address or an IPv4 prefix to the MAP CPE.

In this way, there are stateless mappings between the IPv4 and IPv6 address without requiring a large stateful translator in the network. More information on the details of the MAP algorithm is covered in the later section of this document.

The following are some of the key benefits of implementing Mapping of Address and Port (MAP):

- **Natural Aggregation**—The MAP algorithm for the IPv4 address and port range assignment ensures that the natural IPv6 aggregation logic can be used to direct packets to their destination CPE. The network design remains “clean” and does not contain extraneous per-subscriber or per-connection state.
- **No Translation Logging**—Translation Logging is required in today’s stateful NAT44 or NAT64 implementations. Removing the need for logging results in reduced CAPEX and OPEX for the service provider, as they do not need to invest in an infrastructure to collect and analyze translation logs, which are often required in order to comply with law enforcement requirements. It also removes the complexity and costs associated with sizing the logging solution for peak loads, since the service provider needs to ensure that all transaction records are recorded at all times.

-
- **Flexible Topology**—The use of stateless translation allows traffic to flow asymmetrically in and out of the MAP domain, without requiring the CGN device to see traffic in both directions to generate translation state. This implies more flexible deployment options for the customer such that the network ingress and egress points may be different.
 - **Inline Processing**—On the Cisco ASR9000, MAP is processed in the line card forwarding hardware on the 2nd generation (Typhoon) line cards. This allows for high performance line-rate processing on the line card, reducing the costs per subscriber associated with deploying the solution. The solution then scales with the volume of subscriber traffic, and is not constrained by the number of sessions or session setup rate like NAT44 or DS-Lite.
 - **Improved Resilience**—Since no state is created on the translator, it does not automatically become a chokepoint in the event of a DDoS attack, and generally there are fewer available to use to deny service. Every time a mapping is done, there is an automatic reverse check to make sure that a spoofed source or destination is not translated.

Alternative Technologies

There are various alternative technologies that were intended to address the same, or overlapping, problem space addressed by MAP. The following sections provide a brief overview of the various technologies, and its brief comparison with MAP.

Dual-Stack Lite (DS-Lite) started off as *draft-durand-softwire-dual-stack-lite* and was later ratified as RFC6333. The idea behind Dual Stack Lite (DS-Lite) was conceived prior to Mapping of Address and Port (MAP) and employs the use of IPv4-over-IPv6 tunnels between CPEs and a massive stateful tunnel concentrator, known as the Address-Family Transition Router (AFTR).

DS-Lite allows the service provider to deploy an IPv6-only infrastructure in the aggregation network, and the carriage of IPv4 traffic across the IPv6 infrastructure through the use of tunneling. IPv4-over-IPv6 Tunnels are created between the CPE (known as the Basic Bridging Broadband Element or B4) and the Address Family Transition Router (AFTR), which resides in the service provider network. The B4 provides bridging and encapsulation of the subscriber IPv4 traffic onto an IPv6 tunnel, which terminates at the AFTR. The AFTR performs de-encapsulation of the tunneled IPv4 traffic, and subsequently stateful NAT44 prior to sending it towards the intended destination. The B4 endpoint address is used to differentiate between traffic originating from different subscribers, which may use the same private or internal source address range.

The key advantages, offered by DS-Lite, lies in its ability to automate the tunneling of IPv4 traffic to the CGN NAT44 device (which is also possible with MAP). Because the B4 does not perform NAT44, the solution does not require any additional IPv4 address space, and it also does away with the requirement of double NAT as in the case for a traditional NAT44 or NAT444 solution.

One of the most significant deployment considerations for DS-Lite is in the number of stateful translations supported on the AFTR, as is the case of any stateful NAT44 solution. This is the case even when traffic is being passed within the service provider, as all IPv4 traffic needs to be tunneled from the B4 to the AFTR and back to the other subscriber's B4. This directly impacts the number of subscribers that may be terminated on an AFTR, and multiple AFTRs may need to be deployed in the service provider network, in order to provide the required scale. In the context of the IOS-XR based CGv6 solutions, this would translate to requiring another service blade in order to scale up the solution. In contrast, MAP is processed inline by the line card forwarding processors on the 2nd Generation (Typhoon) line cards, and provides for line-rate performance, which scales to the volume of traffic required rather than the amount of state information.

The other consideration that needs to be taken when a stateful translation technology, such as DS-Lite is being used, is in logging. Proper sizing for both the peak amount of logging information, as well as storage for a period of time, is required in order to comply with law enforcement requirements. This is also another area that has been addressed by MAP—being stateless, it does not require extensive logging of the subscriber's translation details, and this translates to simplified operations and a reduction in both CAPEX and OPEX required to deploy the solution. Stateful solutions attempt to limit the amount of logging, by introducing Bulk Port Allocation (BPA), which pre-allocates groups of ports to each subscriber, and keeps them assigned in an effort to lower the amount of logging. Setting the parameters for BPA (allocation block size) requires a tradeoff between address utilization efficiency (best with small blocks) and logging reduction (best with large blocks), so a compromise must always be made.

Lightweight 4over6 (draft-ietf-softwire-lw4over6) is functionally similar to DS-Lite. However, it attempts to relocate the centralized NAT functionality from the AFTR to the B4 in order to reduce the amount of state information that needs to be kept on the AFTR from a per-flow basis down to a per-subscriber basis. The required NAT functionality is now performed at the Lightweight B4 (lwB4). However, this means that the lwB4 now needs to be provisioned with the public IPv4 address and port set it is allowed to utilize. The Lightweight AFTR (lwAFTR) also needs to be made aware of the binding between the IPv6 address of each subscriber and IPv4 address and port set allocated, so that it can perform ingress filtering upstream, and encapsulation downstream. This means that the information that is maintained by the provisioning system, and that maintained by the lwAFTR, must always be synchronized.


Similar to DS-Lite, it does not address certain topologies, such as those requiring directly meshed connectivity between subscribers, but without packets traversing the AFTR. Overall, Lightweight 4over6 has the same deployment considerations as DS-Lite, with the exception of requiring less state information to be kept on the AFTR.

More information on Lightweight 4over6 is available from draft: [draft-ietf-softwire-lw4over6](#).

Compared to DS-lite, Lightweight 4over6 does get rid of a large amount of state (per-session state used to create the translations). However, it still retains a smaller amount of state that is bound to the subscriber—the per-subscriber state used to derive the IPv4 address and port set allocated. This state amounts to some kind of function that, given the CPE, returns the IPv4 address and port set for this CPE—and the state is the table-driven output of the function.

MAP Terminology

Figure 1. MAP Network Topology



A MAP network comprises Customer Edge (CE) and Border Relay (BR) routers, along with a set of parameters and rules, which define how they work together to deliver the MAP solution. The CEs provide IPv4 and IPv6 service to their users, while the path between the CE and BR is IPv6-only.

The following sections provide additional details on MAP, including Domain Parameters, Mapping Rules, and packet forwarding.

MAP Domain Parameters

All nodes in a MAP Domain must be provisioned with a set of parameters, which are used to implement the MAP functions. These parameters may be configured via the CLI (usually on the BR), or via protocols such as DHCP (most common for the CE nodes). Collectively, these values are called the Basic Mapping Rule (BMR). The BMR is used for both provisioning and forwarding. The provisioned parameters are:

- SP's IPv6 Prefix and its length
- SP's IPv4 Prefix and its length
- Border Relay IPv6 address
- Length of the Embedded Address (EA) bits—how many IPv6 bits are used to represent IPv4
- PSID Offset (optional—explained later in this document)

Mapping Rules

Mapping rules define the forwarding behavior for a MAP domain. Together, they make up the Mapping Rule Table (MRT). The MRT effectively serves as a routing table for the BR and CE to make forwarding decisions based on 48-bit longest match. There are three types of mapping rules:


- Basic Mapping Rule (BMR)—The BMR uses the parameters set during provisioning to forward packets to and from the BR within the MAP domain (hub and spoke model).
- Default Mapping Rule (DMR)—The DMR is used for destinations outside the MAP domain. In the Map Rule Table, it is 0.0.0.0::0.0/0 (48 bits for all addresses and all ports) and points to the BR.
- Forwarding Mapping Rules (FMR)—FMRs are optional additional mapping rules allowing direct CE to CE communication when using a mesh model.

Mapping Algorithm

A core component of the MAP solution is the algorithm that enables the CPEs to determine their IPv4/L4 address space from an IPv6 address assigned via the MAP provisioning. This is the key to stateless operation. It is important to draw a distinction between the addressing used to assign address space (NAT44 pool) to a CE and the addressing on data packets. This should become clear with further examples. The IPv4 address and ports for the CE are encoded in the MAP IPv6 address it is assigned. The MAP IPv6 address contains the following fields:


- *SP IPv6 prefix.*
- *EA bits*—Define the unique address space available to the CE (i.e., the NAT pool). Some of the EA bits are used to complete the IPv4 address and some are used to set the range of L4 ports that may be used by the CE.
 - IPv4 address bits are combined with the IPv4 prefix to set the IPv4 address of the CE. The number of EA bits used for the IPv4 address is determined by the *IPv4 Prefix Length*. Specifically, the number of EA bits used for the IPv4 address is 32 minus the *IPv4 Prefix Length*.
 - Port Set ID (PSID) is the remainder of the EA bits. It defines the L4 ports that may be used by the CE. All the L4 ports used by the CE will have the same PSID bits. Note that the PSID is usually offset within the IPv4 L4 port range. The remaining bits are used to generate the ports in the NAT pool available to the CE. The number of shared bits determines the extent to which the IPv4 address is divided, and is called the sharing ratio. Specifically $2^{(\# \text{ of PSID bits})} = \text{sharing ratio}$.
- Subnet ID—bits for CE to use for subnetting.
- Interface ID—same as regular IPv6, defines hosts.
- PSID Offset (optional)—defines which bits of the L4 address space are used for the PSID. This allows for bits both before and after the PSID to be used for the addresses, which results in noncontiguous port ranges, thus allowing for better security by making the range less predictable. By default, there are 6 offset bits.

Figure 2. EA bits mapping



In summary, IPv4 NAT pool assigned to the CE is shown in the following diagram:

Figure 3. Full Address Format



Mapping Rule Example

The following example shows a complete example of determining the valid IPv4 NAT pool from the provisioned parameters and the assigned MAP IPv6 address.

Provisioned Parameters

Table 1. Provisioned Parameters

Field	Value
SP IPv6 Prefix	2001:db8:9500::
SP IPv6 Mask Length	/40
IPv4 Prefix	198.51.100.190
IPv4 Mask Length	/24
Number of EA Bits	16
IPv6 address assigned to the CE	2001:db8:95BE:EF00::
PSID Offset	4

Derived Values

The EA bits are the first 16 bits (per the provisioned value) after the SP IPv6 prefix—0x**BEEF**. There are 8 bits remaining in the IPv4 address, so the IPv4 address uses the first 8 EA bits—0x**BE** which is 190. Therefore, the IPv4 address used for the NAT pool is **198.51.100.190**. The PSID is the remaining EA bits. Since EA bits length is provisioned as 16, there are 8 (16 minus 8) bits used for the PSID. Therefore, the PSID is 0x**EF**. The PSID offset is 4, so bits 4-11 of the IPv4 L4 port range will be fixed for this CE, and bits 0-3 and 12-15 will be allocated to the NAT pool. There is a rule prohibiting all of the EA bits from being set to zero. This rule prevents the use of

reserved ports (e.g., HTTP, FTP...). Therefore, there are 15 blocks of 16 L4 ports (total of 240) available for this CE. The ranges are: 7920-7935, 12016-12031...65264-65279.

The table below shows some of the ports in the NAT pool for the CE.

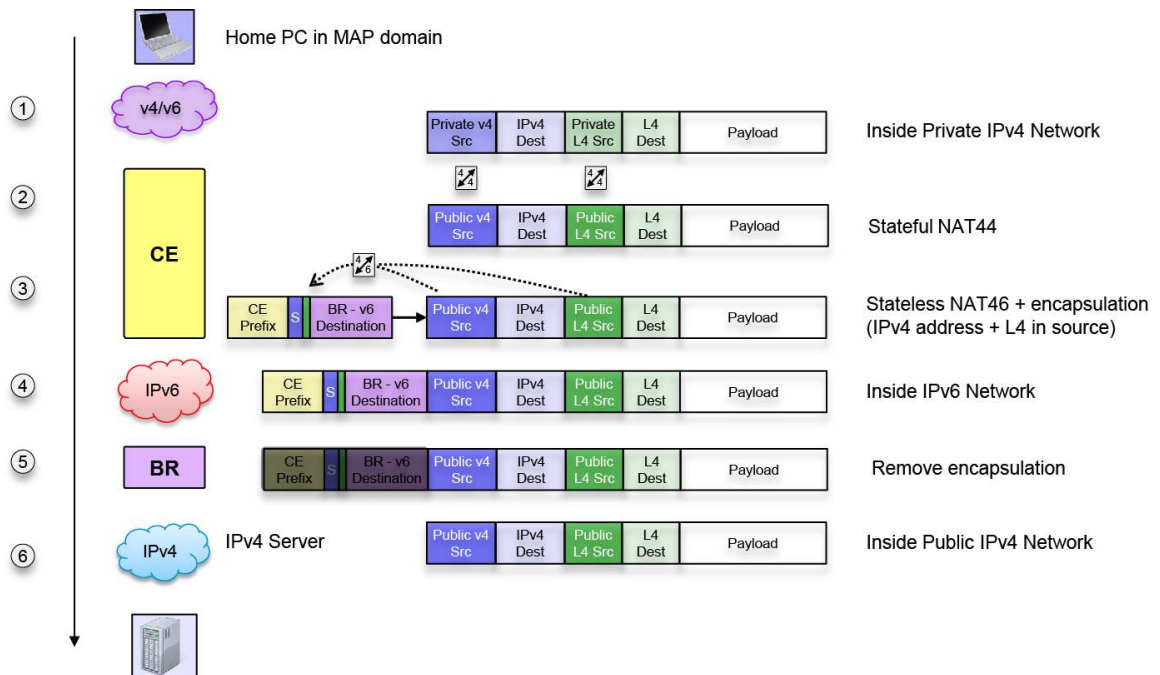
Table 2. Address/Port Assignment Example

First 4 bits	PSID	Last 4 bits	L4 port hex	L4 port
0x1	0xEF	0x0	0x1EF0	7920
0x1	0xEF	0x1	0x1EF1	7921
0x1	0xEF	0x2	0x1EF2	7922
0x2	0xEF	0x0	0x2EF0	12016
0x2	0xEF	0x1	0x2EF1	12017
0x2	0xEF	0x2	0x2EF2	12018
0xF	0xEF	0x0	0xFEFF	65279

Packet Forwarding—MAP-E

For traffic leaving the MAP domain, the source is the CE and the destination is the BR. For traffic returning from outside the MAP domain, the source is the BR and the destination is the CE. For forwarding, these paths can be viewed as point-to-point tunnels. An example of the CE to BR path for traffic leaving the MAP domain is shown below.

Figure 4. MAP-E Packet Path



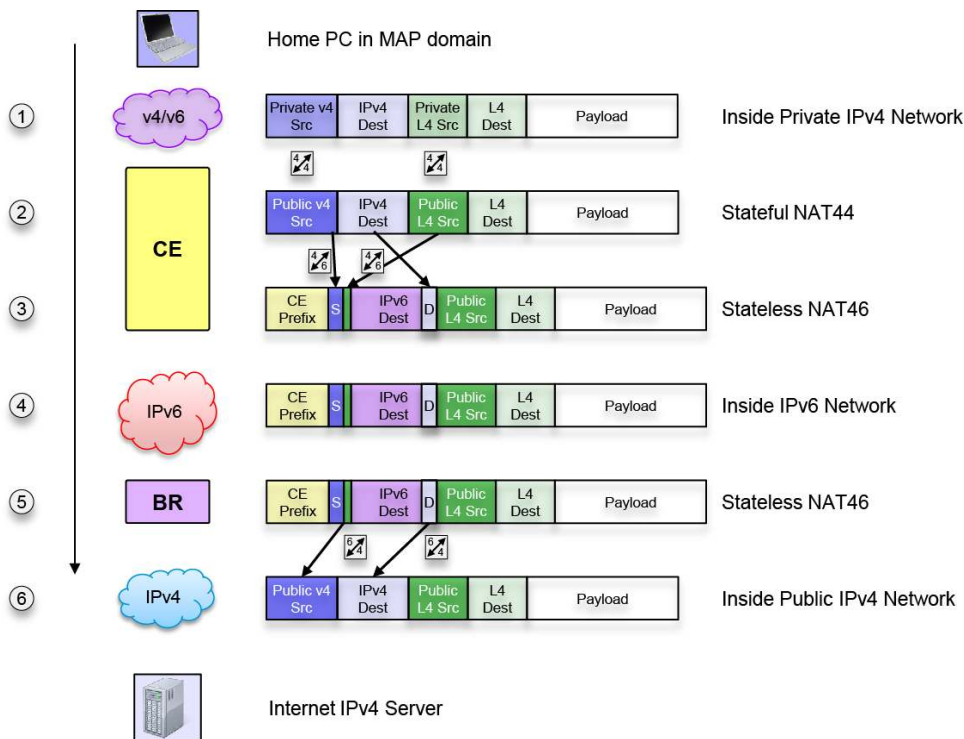
The following steps occur in forwarding a packet from an IPv4 client in a MAP-E domain, to an IPv4 server in the Internet.

1. The packet is sent with a private source IP address and port and the destination IP returned from DNS.
2. At the CE, the packet will undergo a stateful NAT44 translation of its source address and port. The public IP address and port are chosen from the NAT pool that was created via the MAP rules and the DHCP-assigned MAP IPv6 address.
3. The IPv6 header will be constructed as follows: The IPv6 source address is a combination of the CE's IPv6 address and the client's translated IPv4 address and port. The destination IPv6 address is the BR.
4. The IPv6 packet is transmitted over the IPv6 network towards the BR.
5. The BR removes the IPv6 header and forwards the IPv4 packet toward the destination.
6. IPv4 forwarding happens normally until reaching the destination.

Packet Forwarding—MAP-T

There are four packet forwarding operations in MAP-T: CE v4/v6, CE v6/v4, BR v4/v6, and BR v6/v4. As an example, the packet from a home network PC to an Internet IPv4 server utilizes the CE v4/v6 and BR v6/v4 operations as shown below.

Figure 5. MAP-T Packet Path



1. The IPv4 packet is sent normally from the IPv4 client.
2. The MAP-T CE performs the Stateful NAT44 translation into the port-restricted IP/L4 range. The IPv4/L4 destination is then used to select the appropriate MAP rule (which is used to determine the NAT64 operation) from the MRT.
3. The MAP-T CE performs a Stateless NAT46 operation to encode the IPv4 source, source L4 port, and IPv4 destination in the IPv6 header. The type of rule followed (e.g., DMR or FMR) will determine the parameters for the encoding. The details of this encoding are outside the scope of this paper, but can be found in the references.
4. The IPv6 packet on the wire has no IPv4 header since the IPv4 addresses are encoded in the IPv6 header.
5. The MAP-T BR performs a lookup on the **source** IPv6 address to find the appropriate MAP domain/CE, and a lookup on the **destination** address to find the longest match in the Mapping Rule Table. The BR then derives the IPv4 addresses using the chosen mapping rule, and performs the Stateless NAT64 operation. In this case, the packet is being forwarded outside of the MAP domain, so the DMR is used.
6. The packet is then forwarded in the public IPv4 network to the destination.

IOS XR MAP Configuration

The router configuration for MAP-E and MAP-T is similar. The difference will be highlighted in blue in the example below.

Table 3. Table 3- MAP Configuration

Configuration	Notes
<code>service cgn demo</code>	Create the CGN instance
<code>service-type map-t 1</code>	Configure the MAP-T instance (other CGN types may be active on the "demo" service). Set to <code>map-e</code> for encapsulation mode.
<code>cpe-domain ipv4 prefix 192.0.2.0</code>	Specify IPv4 Prefix and Length
<code>cpe-domain ipv6 prefix 2001:db8:b001::/48</code>	Specify IPv6 Prefix and Length
<code>external-domain ipv6 prefix 2001:db8:b001:ffff::/64</code>	
<code>sharing-ratio 8</code>	Sets length of EA bits (together with v4 prefix length)
<code>consecutive-ports 4</code>	Sets PSID offset
<code>address-family ipv4</code> <code>interface ServiceApp1</code>	Virtual tunnel to IPv4 side of the MAP function (implemented on LC and ISM/VSM)
<code>address-family ipv6</code> <code>interface ServiceApp2</code>	Virtual tunnel to IPv6 side of the MAP function (implemented on LC and ISM/VSM)
<code>router static</code> <code>address-family ipv4 unicast</code> <code>192.0.2.0/28 ServiceApp2 192.0.2.41</code> <code>address-family ipv6 unicast</code> <code>2001:db8:b001:ffff::/64 ServiceApp1</code>	Route traffic to the ServiceApps IPv4 prefix to IPv4 side IPv6 domain prefix to IPv6 side
<code>interface ServiceApp1</code> <code>ipv6 address 2001:db8::1/64</code> <code>service cgn demo service-type map-t</code> <code>interface ServiceApp2</code> <code>ipv4 address 192.0.2.42 255.255.255.0</code> <code>service cgn demo service-type map-t</code>	Connect ServiceApps to the CGN service

MAP-T vs. MAP-E Considerations

MAP-T and MAP-E use the same algorithm for port and address translation. However, MAP-T does not encapsulate the original IPv4 header with an additional IPv6 header like MAP-E, but instead performs the conversion of the original IPv4 address and port into an IPv6 address and port combination. This results in less overhead and simpler deep packet inspection for MAP-T. There are some other (sometimes subtle) differences between the two variants:

- Reduced overhead for MAP-T, due to not carrying the IPv4 header.
- QoS classification—Classification is done on encapsulation endpoints or a tunnel-aware data plane for MAP-E. For MAP-T, it leverages the IPv6 data plane features directly.
- IP options are fully supported in MAP-E, but not MAP-T; there isn't room for them in the IPv6 header.
- IPv6 does not have all the IPv4 ICMP codes. MAP-E preserves them. MAP-T does not, although some heuristics can be used to try to convey the semantics of the message without a one-to-one translation.
- Fragmentation of IPv6 packets will result in some IPv6 packets not having the IPv4 header, and break anycast since only a single BR can reassemble the packet. MAP-E solves this problem by requiring that the fragmentation occur in the IPv4 stack. MAP-T preserves the IPv4 addresses in all packets and enables the fragmentation to occur normally at the endpoint.
- Deep packet inspection is more challenging with MAP-E due to the extra header.

Deployment Considerations

There are a many decisions that must be made in planning a MAP deployment. Obviously using MAP-T vs. MAP-E is one, but there are others that are equally, if not more, important. These have been discussed in draft-ietf-softwire-map-deployment (<http://tools.ietf.org/html/draft-ietf-softwire-map-deployment>) which is outlined below. Reading this is highly recommended for further detail on deployment considerations, but also to reinforce the material presented in this paper.

- Dividing the network into MAP domains
- Available address space for IPv4 and IPv6 and appropriate mask lengths for CE assignments
- Sharing ratio—number of supported CEs and ports per CE. A network may have multiple sharing ratios (PSID length) by using multiple MAP domains.
- Mesh (allows CE to CE path), hub and spoke (all traffic via BR), or hybrid (some CEs have FMRs). Mesh is recommended if there is high CE-to-CE traffic and a smaller number (10) of MAP rules in use within the domain.
- Assignment mechanism for the MAP IPv6 address—when to use DHCP and when use static configuration
- BR placement—Placing the BR close to the CEs minimizes the path for CE to CE communication in the hub-and-spoke model, but also requires more of the IPv4 network support, which impacts operations.
- Whether or not to use anycast (via advertising a shared address in the DMR). Anycast provides redundancy in the case of BR failures.
- Setting MTUs to avoid or minimize fragmentation.

-
- Logging of address mapping for legal or other requirements. Although per-session logging is not required due to the static allocations, MAP rule changes must be recorded so that the mappings at any given time are known.
 - As MAP requires port translation, only TCP, UDP, and ICMP are supported. Ensure that other types of traffic are accommodated appropriately.

Cisco ASR 9000 Notes

On the Cisco ASR9000 specifically, inline MAP processing (MAP-E/T) requires the use of 2nd generation (Typhoon) line cards. This results in a line rate performance of 30 Gbps per Typhoon NPU. However, at least one service engine (initially the A9K-ISM-100) will need to be present in the chassis to perform the control plane and special packet processing (e.g., ICMP and fragments). Similarly, IPv4 UDP packets with checksum values of zero are also punted to the ISM for processing.

Summary

Mapping of Address and Port (MAP) is an elegant IPv6 transition and internetworking technology, which leverages on the use of the L3/L4 address, and port space mappings, to allow IPv4 addresses to be translated to IPv6 addresses without the use of a stateful translator. Its key benefits include scaling linearly with traffic volume, and not requiring the use of a logging solution. In addition, it provides the service provider with a flexible network topology, and provides for improved resiliency in the face of DDoS attacks, which no longer have a stateful resource to exhaust.

References

The following documents provide additional reading materials on Mapping of Address and Port (MAP).

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

http://www.cisco.com/en/US/solutions/collateral/ns341/ns525/ns1017/solution_overview_c11-726499.html

Mapping of Address and Port-Encapsulation Mode

<http://tools.ietf.org/html/draft-ietf-softwire-map>

Mapping of Address and Port-Translation Mode

<http://tools.ietf.org/html/draft-ietf-softwire-map-t>

Mapping of Address and Port (MAP)—Deployment Considerations

<http://tools.ietf.org/html/draft-ietf-softwire-map-deployment>

MAP Addressing Simulation Tool (*These are links to download Cisco developed software tools*)

<http://6lab.cisco.com/map/MAP.php>

<https://itunes.apple.com/us/app/cisco-map-calculator/id561121079?mt=8> (iTunes)

https://play.google.com/store/apps/details?id=map.calculator&hl=en_GB (Android)



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)

Printed in USA

C11-729800-00 01/14