

Spring
Internet-Draft
Intended status: Informational
Expires: October 15, 2017

J. Brzozowski
J. Leddy
Comcast
C. Filsfils
R. Maglione, Ed.
M. Townsley
Cisco Systems
April 13, 2017

IPv6 SPRING Use Cases
draft-ietf-spring-ipv6-use-cases-10

Abstract

The objective of this document is to illustrate some use cases that need to be taken into account by the Source Packet Routing in Networking (SPRING) architecture in the context of an IPv6 environment.

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 October 15, 2017.

Copyright Notice

Copyright (c) 2017 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must

include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1.	Introduction	2
2.	IPv6 SPRING use cases	4
2.1.	SPRING in the Home Network	4
2.2.	SPRING in the Access Network	5
2.3.	SPRING in the Data Center	6
2.4.	SPRING in the Content Delivery Networks	6
2.5.	SPRING in the Core networks	7
3.	Contributors	8
4.	Acknowledgements	9
5.	IANA Considerations	9
6.	Security Considerations	9
7.	References	9
7.1.	Informative References	9
7.2.	Normative References	10
	Authors' Addresses	10

[1.](#) Introduction

Source Packet Routing in Networking (SPRING) architecture leverages the source routing paradigm. An ingress node steers a packet through a controlled set of instructions, called segments, by prepending the packet with SPRING header. The SPRING architecture is described in [[I-D.ietf-spring-segment-routing](#)].

In today's networks, source routing is typically accomplished by encapsulating IP packets in MPLS LSPs that are signaled via RSVP-TE. Therefore, there are scenarios where it may be possible to run IPv6 on top of MPLS, and as such, the MPLS Segment Routing architecture described in [[I-D.ietf-spring-segment-routing-mpls](#)] could be leveraged to provide spring capabilities in an IPv6/MPLS environment.

However, there are other cases and/or specific network segments (such as for example the Home Network, the Data Center, etc.) where MPLS may not be available or deployable for lack of support on network elements or for an operator's design choice. In such scenarios a non-MPLS based solution would be preferred by the network operators of such infrastructures.

In addition there are cases where the operators could have made the design choice to disable IPv4, for ease of management and scale (return to single-stack) or due to an address constraint, for example because they do not possess enough IPv4 addresses resources to number

all the endpoints and other network elements on which they desire to run MPLS.

In such scenario the support for MPLS operations on an IPv6-only network would be required. However today's IPv6-only networks are not fully capable of supporting MPLS. There is ongoing work in the MPLS Working Group, described in [[RFC7439](#)] to identify gaps that must be addressed in order to allow MPLS-related protocols and applications to be used with IPv6-only networks. This is an another example of scenario where a solution relying on IPv6 without requiring the use of MPLS could represent a valid option to solve the problem and meet operators' requirements.

It is important to clarify that today, it is possible to run IPv6 on top of an IPv4 MPLS network by using the mechanism called 6PE, described in [[RFC4798](#)]. However this approach does not fulfill the requirement of removing the need of IPv4 addresses in the network, as requested in the above use case.

In summary there is a class of use cases that motivates an IPv6 data plane. This document identifies some fundamental scenarios that, when recognized in conjunction, strongly indicate an IPv6 data plane:

1. There is a need or desire to impose source-routing semantics within an application or at the edge of a network (for example, a CPE or home gateway)
2. There is a strict lack of an MPLS dataplane in a portion of the end to end path
3. There is a need or desire to remove routing state from any node other than the source, such that the source is the only node that knows and will know the path a packet will take, a priori
4. There is a need to connect millions of addressable segment endpoints, thus high routing scalability is a requirement. IPv6

addresses are inherently summarizable: a very large operator could scale by summarizing IPv6 subnets at various internal boundaries. This is very simple and is a basic property of IP routing. MPLS node segments are not summarizable. To reach the same scale, an operator would need to introduce additional complexity, such as mechanisms known with the industry term Seamless MPLS [[I-D.ietf-mpls-seamless-mpls](#)].

In any environment with requirements such as those listed above, an IPv6 data plane provides a powerful combination of capabilities for a network operator to realize benefits in explicit routing, protection and restoration, high routing scalability, traffic engineering,

service chaining, service differentiation and application flexibility via programmability.

[2.](#) IPv6 SPRING use cases

This section will describe some scenarios where MPLS may not be present and it will highlight the need for the spring architecture to take them into account.

The use cases described in the section do not constitute an exhaustive list of all the possible scenarios; this section only includes some of the most common envisioned deployment models for IPv6 Segment Routing. In addition to the use cases described in this document the spring architecture should be able to be applied to all the use cases described in [[RFC7855](#)] for the spring MPLS data plane, when an IPv6 data plane is present.

[2.1.](#) SPRING in the Home Network

An IPv6-enabled home network provides ample globally routed IP addresses for all devices in the home. An IPv6 home network with multiple egress points and associated provider-assigned prefixes will, in turn, provide multiple IPv6 addresses to hosts. A homenet performing Source and Destination Routing ([[I-D.ietf-rtgwg-enterprise-pa-multihoming](#)]) will ensure that packets exit the home at the appropriate egress based on the associated delegated prefix for that link.

A spring enabled home provides the ability to steer traffic into a

specific path from end-hosts in the home, or from a customer edge router in the home. If the selection of the source routed path is enabled at the customer edge router, that router is responsible for classifying traffic and steering it into the correct path. If hosts in the home have explicit source selection rules, classification can be based on source address or associated network egress point, avoiding the need for DPI-based implicit classification techniques. If the traffic is steered into a specific path by the host itself, it is important to know which networks can interpret the spring header. This information can be provided as part of host configuration as a property of the configured IP address.

The ability to steer traffic to an appropriate egress or utilize a specific type of media (e.g., low-power, WIFI, wired, femto-cell, bluetooth, MOCA, HomePlug, etc.) within the home itself are obvious cases which may be of interest to an application running within a home network.

Steering to a specific egress point may be useful for a number of reasons, including:

- o Regulatory
- o Performance of a particular service associated with a particular link
- o Cost imposed due to data-caps or per-byte charges
- o Home vs. work traffic in homes with one or more teleworkers, etc.
- o Specific services provided by one ISP vs. another

Information included in the spring header, whether imposed by the end-host itself, a customer edge router, or within the access network of the ISP, may be of use at the far ends of the data communication as well. For example, an application running on an end-host with application-support in a data center can utilize the spring header as a channel to include information that affects its treatment within the data center itself, allowing for application-level steering and load-balancing without relying upon implicit application

classification techniques at the data-center edge. Further, as more and more application traffic is encrypted, the ability to extract (and include in the spring header) just enough information to enable the network and data center to load-balance and steer traffic appropriately becomes more and more important.

[2.2.](#) SPRING in the Access Network

Access networks deliver a variety of types of traffic from the service provider's network to the home environment and from the home towards the service provider's network.

For bandwidth management or related purposes, the service provider may want to associate certain types of traffic to specific physical or logical downstream capacity pipes.

This mapping is not the same thing as classification and scheduling. In the Cable access network, each of these pipes are represented at the DOCSIS [[DOCSIS](#)] layer as different service flows, which are better identified as differing data links. As such, creating this separation allows an operator to differentiate between different types of content and perform a variety of differing functions on these pipes, such as byte capping, regulatory compliance functions, and billing.

In a cable operator's environment, these downstream pipes could be a specific QAM [[QAM](#)], a DOCSIS [[DOCSIS](#)] service flow or a service group.

Similarly, the operator may want to map traffic from the home sent towards the service provider's network to specific upstream capacity pipes. Information carried in a packet's spring header could provide the target pipe for this specific packet. The access device would not need to know specific details about the packet to perform this mapping; instead the access device would only need to know the interpretation of the spring header and how to map it to the target pipe.

[2.3.](#) SPRING in the Data Center

Some Data Center operators are transitioning their Data Center infrastructure from IPv4 to native IPv6 only, in order to cope with IPv4 address depletion and to achieve larger scale. In such environment, source routing (through Segment Routing IPv6) can be used to steer traffic across specific paths through the network. The specific path may also include a given function one or more nodes in the path are requested to perform.

In addition one of the fundamental requirements for Data Center architecture is to provide scalable, isolated tenant networks. In such scenario Segment Routing can be used to identify specific nodes, tenants, and functions and to build a construct to steer the traffic across that specific path.

[2.4.](#) SPRING in the Content Delivery Networks

The rise of online video applications and new, video-capable IP devices has led to an explosion of video traffic traversing network operator infrastructures. In the drive to reduce the capital and operational impact of the massive influx of online video traffic, as well as to extend traditional TV services to new devices and screens, network operators are increasingly turning to Content Delivery Networks (CDNs).

Several studies showed the benefits of connecting caches in a hierarchical structure following the hierarchical nature of the Internet. In a cache hierarchy one cache establishes peering relationships with its neighbor caches. There are two types of relationship: parent and sibling. A parent cache is essentially one level up in a cache hierarchy. A sibling cache is on the same level. Multiple levels of hierarchy are commonly used in order to build efficient caches architecture.

In an environment, where each single cache system can be uniquely identified by its own IPv6 address, a list containing a sequence of the caches in a hierarchy can be built. At each node (cache) in the list, the presence of the requested content is checked. If the requested content is found at the cache (cache hits scenario) the sequence ends, even if there are more nodes in the list; otherwise next element in the list (next node/cache) is examined.

[2.5.](#) SPRING in the Core networks

MPLS is a well-known technology widely deployed in many IP core networks. However there are some operators that do not run MPLS everywhere in their core network today, thus moving forward they would prefer to have an IPv6 native infrastructure for the core network.

While the overall amount of traffic offered to the network continues to grow and considering that multiple types of traffic with different characteristics and requirements are quickly converging over single network architecture, the network operators are starting to face new challenges.

Some operators are looking at the possibility to setup an explicit path based on the IPv6 source address for specific types of traffic in order to efficiently use their network infrastructure. In case of IPv6 some operators are currently assigning or plan to assign IPv6 prefix(es) to their IPv6 customers based on regions/geography, thus the subscriber's IPv6 prefix could be used to identify the region where the customer is located. In such environment the IPv6 source address could be used by the Edge nodes of the network to steer traffic and forward it through a specific path other than the optimal path.

The need to setup a source-based path, going through some specific middle/intermediate points in the network may be related to different requirements:

- o The operator may want to be able to use some high bandwidth links for specific type of traffic (like video) avoiding the need for over-dimensioning all the links of the network;
- o The operator may want to be able to setup a specific path for delay sensitive applications;
- o The operator may have the need to be able to select one (or multiple) specific exit point(s) at peering points when different peering points are available;

- o The operator may have the need to be able to setup a source based

path for specific services in order to be able to reach some servers hosted in some facilities not always reachable through the optimal path;

- o The operator may have the need to be able to provision guaranteed disjoint paths (so-called dual-plane network) for diversity purposes

All these scenarios would require a form of traffic engineering capabilities in IP core networks not running MPLS and not willing to run it.

3. Contributors

Many people contributed to this document. The authors of this document would like to thank and recognize them and their contributions. These contributors provided invaluable concepts and content for this document's creation.

Ida Leung
Rogers Communications
8200 Dixie Road
Brampton, ON L6T 0C1
CANADA

Email: Ida.Leung@rci.rogers.com

Stefano Previdi
Cisco Systems
Via Del Serafico, 200
Rome 00142
Italy

Email: sprevidi@cisco.com

Christian Martin
Cisco Systems

Email: martincj@cisco.com

[4.](#) Acknowledgements

The authors would like to thank Brian Field, Robert Raszuk, Wes George, Eric Vyncke, Fred Baker, John G. Scudder and Yakov Rekhter for their valuable comments and inputs to this document.

[5.](#) IANA Considerations

This document does not require any action from IANA.

[6.](#) Security Considerations

This document presents use cases to be considered by the spring architecture and potential IPv6 extensions. As such, it does not introduce any security considerations. However, there are a number of security concerns with source routing at the IP layer [[RFC5095](#)]. It is expected that any solution that addresses these use cases to also address any security concerns.

[7.](#) References

[7.1.](#) Informative References

[DOCSIS] "DOCSIS Specifications Page",
<<http://www.cablelabs.com/news/new-generation-of-docsis-technology/>>.

[I-D.ietf-mpls-seamless-mpls]
Leymann, N., Decraene, B., Filsfils, C., Konstantynowicz, M., and D. Steinberg, "Seamless MPLS Architecture", [draft-ietf-mpls-seamless-mpls-07](#) (work in progress), June 2014.

[I-D.ietf-rtgwg-enterprise-pa-multihoming]
Baker, F., Bowers, C., and J. Linkova, "Enterprise Multihoming using Provider-Assigned Addresses without Network Prefix Translation: Requirements and Solution", [draft-ietf-rtgwg-enterprise-pa-multihoming-00](#) (work in progress), March 2017.

[I-D.ietf-spring-segment-routing]
Filsfils, C., Previdi, S., Decraene, B., Litkowski, S., and R. Shakir, "Segment Routing Architecture", [draft-ietf-spring-segment-routing-11](#) (work in progress), February 2017.

[I-D.ietf-spring-segment-routing-mpls]

Filsfils, C., Previdi, S., Bashandy, A., Decraene, B., Litkowski, S., and R. Shakir, "Segment Routing with MPLS data plane", [draft-ietf-spring-segment-routing-mpls-08](#) (work in progress), March 2017.

[QAM] "QAM specification", <ITU-T Recommendation J.83 Annex B (J.83b)>.

[RFC4798] De Clercq, J., Ooms, D., Prevost, S., and F. Le Faucheur, "Connecting IPv6 Islands over IPv4 MPLS Using IPv6 Provider Edge Routers (6PE)", [RFC 4798](#), DOI 10.17487/RFC4798, February 2007, <<http://www.rfc-editor.org/info/rfc4798>>.

[RFC5095] Abley, J., Savola, P., and G. Neville-Neil, "Deprecation of Type 0 Routing Headers in IPv6", [RFC 5095](#), DOI 10.17487/RFC5095, December 2007, <<http://www.rfc-editor.org/info/rfc5095>>.

[RFC7439] George, W., Ed. and C. Pignataro, Ed., "Gap Analysis for Operating IPv6-Only MPLS Networks", [RFC 7439](#), DOI 10.17487/RFC7439, January 2015, <<http://www.rfc-editor.org/info/rfc7439>>.

[7.2.](#) Normative References

[RFC7855] Previdi, S., Ed., Filsfils, C., Ed., Decraene, B., Litkowski, S., Horneffer, M., and R. Shakir, "Source Packet Routing in Networking (SPRING) Problem Statement and Requirements", [RFC 7855](#), DOI 10.17487/RFC7855, May 2016, <<http://www.rfc-editor.org/info/rfc7855>>.

Authors' Addresses

John Brzozowski
Comcast

Email: john_brzozowski@cable.comcast.com

John Leddy
Comcast

Email: John_Leddy@cable.comcast.com

Brzozowski, et al.

Expires October 15, 2017

[Page 10]

Internet-Draft

IPv6 Segment Routing Use Cases

April 2017

Clarence Filsfils
Cisco Systems
Brussels
BE

Email: cfilsfil@cisco.com

Roberta Maglione (editor)
Cisco Systems
Via Torri Bianche 8
Vimercate 20871
Italy

Email: robmgl@cisco.com

Mark Townsley
Cisco Systems

Email: townsley@cisco.com

