

Network Working Group
Internet Draft
Intended status: Informational
Expires: September 3, 2022

Yisong Liu
W. Cheng
China Mobile
C. Lin
New H3C Technologies
X. Geng
Huawei Technologies
Y. Liu
ZTE
March 3, 2022

Considerations for Protection of SRv6 Networks
draft-liu-rtgwg-srv6-protection-considerations-01

Abstract

This document describes the considerations for protection of SRv6 network.

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), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/lid-abstracts.txt>

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>

This Internet-Draft will expire on September 3, 2022.

Copyright Notice

Copyright (c) 2022 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..... | 3 |
| 1.1. | Requirements Language..... | 3 |
| 1.2. | Terminology..... | 3 |
| 2. | Forwarding over SRv6 Network..... | 3 |
| 2.1. | SRv6 BE Path..... | 3 |
| 2.2. | SRv6 TE Path..... | 4 |
| 3. | Protection Mechanisms..... | 5 |
| 3.1. | Path Protection..... | 5 |
| 3.1.1. | Local Protection Mechanisms..... | 5 |
| 3.1.2. | Liveness Check For Local Protection..... | 6 |
| 3.1.3. | Micro-Loop Avoidance..... | 6 |
| 3.1.4. | End-to-End Protection Mechanisms..... | 6 |
| 3.1.5. | Liveness Check For End-to-End Protection..... | 7 |
| 3.2. | Service Protection..... | 8 |
| 3.2.1. | Local Repair..... | 8 |
| 3.2.2. | Ingress Node Switchover..... | 8 |
| 4. | Implementation Recommendations..... | 9 |
| 4.1. | SRv6 BE..... | 11 |
| 4.2. | SRv6 TE..... | 12 |
| 5. | Security Considerations..... | 15 |
| 6. | IANA Considerations..... | 15 |
| 7. | Contributors..... | 15 |
| 8. | References..... | 16 |
| 8.1. | Normative References..... | 16 |
| 8.2. | Informative References..... | 17 |
| | Authors' Addresses..... | 18 |

1. Introduction

Segment Routing [[RFC8402](#)] instantiated on the IPv6 dataplane (SRv6) provides network programming capability to create interoperable overlays with underlay optimization [[RFC8986](#)].

This document describes the common failure scenarios and protection mechanisms in SRv6 networks. Then implementation recommendations for protection of SRv6 networks are proposed.

1.1. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [BCP 14](#) [[RFC2119](#)] [[RFC8174](#)] when, and only when, they appear in all capitals, as shown here.

1.2. Terminology

BE: Best Effort

TE: Traffic Engineering

G-SRv6: Generalized SRv6 Network Programming

2. Forwarding over SRv6 Network


Segment Routing [[RFC8402](#)] leverages the source routing paradigm. Segment Routing instantiated on the IPv6 dataplane is referred to as SRv6. SRv6 provides network programming capability to create interoperable overlays with underlay optimization [[RFC8986](#)].

In an SRv6 network, the ingress node encapsulates a received packet in an outer IPv6 header, followed by an optional Segment Routing Header (SRH) [[RFC8754](#)], which instructs the SRv6 network to forward the packet via a specific path to the egress node. The forwarding path is either an SRv6 BE path or an SRv6 TE path.

2.1. SRv6 BE Path

In the SRv6 BE path, the ingress PE encapsulates the payload in an outer IPv6 header where the destination address is the SRv6 Service

SID provided by the egress PE. The underlay P nodes between the PEs only need to perform plain IPv6 shortest path forwarding.




Ingress PE ---> P nodes ---> Egress PE

Figure 1: Forwarding over SRv6 BE

2.2. SRv6 TE Path

In the SRv6 TE path, the ingress PE steers the traffic flow into an SR Policy [[I-D.ietf-spring-segment-routing-policy](#)], and encapsulates the payload packet in an outer IPv6 header with the Segment Routing Header (SRH) carrying the segment list of the SR policy. The underlay P nodes whose SRv6 SID's are part of the SRH segment list are called endpoint nodes. They will be involved in the forwarding path and execute the function associated with the SID.




Ingress PE ---> P nodes ---> Egress PE

Figure 2: Forwarding over SRv6 TE

If Compressed Segment List encoding is enabled in the SRv6 network [[I-D.ietf-spring-srv6-srh-compression](#)], the segment list in the SRH will be encoded in the compressed way. The compressed SRv6 Segment-List encoding can optimize the packet header length by avoiding the repetition of the Locator-Block and trailing bits with each individual SID.

The G-SRv6 mechanism will be used as an example for the encoding of SRv6 TE path in this document. Figure 3 shows the encapsulation of packet using the G-SRv6 mechanism.



Ingress PE ---> P nodes ---> Egress PE

Figure 3: Forwarding over G-SRv6 Encoded TE

[3. Protection Mechanisms](#)

3.1. Path Protection

[3.1.1. Local Protection Mechanisms](#)

Local protection is performed by the node adjacent to the failed component using fast-reroute techniques [[RFC5286](#)] [[RFC5714](#)]. The common method of local repair is to provide a repair path for the destination avoiding the failed component.

[I-D.ietf-rtgwg-segment-routing-ti-lfa] describes the Topology Independent Loop-free Alternate Fast Re-route technology (TI-LFA) using Segment Routing, which is able to provide a loop free backup path irrespective of the topologies used in the network. For each destination in the network, TI-LFA pre-installs a backup forwarding entry for each protected destination ready to be activated upon

detection of the failure of a link used to reach the destination.

In SRv6 dataplane, the TI-LFA repair path is encoded as an SRv6 SID list, and encapsulated in the SRH along with an outer IPv6 header. If Compressed Segment List encoding is enabled, the repair node should check the G-SRv6 capability of nodes along the repair path and try to use G-SIDS to encode the repair path, which will help to optimize the packet header length.

[3.1.2.](#) Liveness Check For Local Protection

In order to perceive the failures of links and neighbors, a node should monitor the liveness of its adjacent components.

[RFC5880] and [[RFC7880](#)] provide widely used mechanisms for liveness check, called Bidirectional Forwarding Detection (BFD) and Seamless Bidirectional Forwarding Detection (S-BFD).

BFD can be associated with the interface state to detect the failure of directly-connected links. Two adjacent nodes may establish BFD or S-BFD sessions between each other, and send BFD control packets to monitor the liveness of each other. In another way, a node may send BFD echo packets to all the neighbors, and they will reflect the packets back, without establishing BFD sessions.

Other OAM methods, such as Ping, TWAMP or STAMP, may also be used for liveness check for local protection, which will not be enumerated here in detail.

[3.1.3.](#) Micro-Loop Avoidance

When a component fails or comes back up, the topology is changed. The routing convergence happens in each node at different times and during a different lapse of time. These transient routing inconsistencies may cause micro-loops.

[I-D.bashandy-rtgwg-segment-routing-uloop] provides a mechanism leveraging segment routing to ensure loop-freeness during the IGP reconvergence process, which relies on the temporary use of SR policies ensuring loop-freeness over the post-convergence paths from the converging node to the destination.

In SRv6 dataplane, the loop-free post-convergence path is encoded as

an SRv6 SID list, and encapsulated in the SRH along with an outer IPv6 header. If Compressed Segment List encoding is enabled, the converging node should check the G-SRv6 capability of nodes along the post-convergence path and try to use G-SIDs to encode the path.

[3.1.4.](#) End-to-End Protection Mechanisms

End-to-end protection lets the ingress PE node be in charge of the failure recovery. The ingress node should steer the flow from the failed path into another alive path.

Liu, et al.

Expires September, 2022

[Page 6]

Internet-Draft

SRv6 Protection Considerations

March 2022

In the case of SRv6 TE path, the SR Policy itself allows for multiple candidate paths, of which at any point in time there is a single active candidate path that is provisioned in the forwarding plane and used for traffic steering [I-D.ietf-spring-segment-routing-policy]. The candidate path with highest preference is selected as the primary path, and the candidate path with second highest preference can be selected as the hot-standby backup. When the primary candidate path fails, switchover to the backup candidate path can be triggered by fast re-route mechanism.

If all the candidate paths fail, the ingress node may use SRv6 BE path for best-effort forwarding.

[3.1.5.](#) Liveness Check For End-to-End Protection

It is essential that the ingress PE node should check the end-to-end liveness of paths, including primary path and backup path. So that the ingress PE node can perceive the path failure and then trigger the switchover.

In the case of SRv6 TE path, BFD or S-BFD can be used to monitor the liveness of SR Policy at the level of segment list. If all the BFD sessions associated with segment lists in a candidate path are down, the candidate path is deemed to be failed. If all the candidate paths is failed, the SR Policy is deemed to be failed.

Moreover, If the SRv6 TE path is strict (every hop along the path appearing in the SID list), the reverse path of the BFD packets should be the same with the forward path. Otherwise, the failure in the reverse path may cause the misjudgement of the liveness of SR Policy. To achieve the consistence of forward path and reverse path,

the egress node should be instructed to use specific path to send packets back to the ingress node.

Other OAM methods, such as Ping, TWAMP or STAMP, may also be used for liveness check for end-to-end protection, which will not be enumerated here in detail.

Local protection and end-to-end protection may both be used in the same SRv6 network. Since the speed of failure detection for local protection is faster than end-to-end protection, local protection usually performs the local repair in advance, which allows the path to remain alive. In this case, the ingress node will not perceive the failure and does not need to trigger end-to-end protection.

3.2. Service Protection

If the failure occurs on the egress PE node, the service provided by that PE is not accessible anymore. TI-LFA or the hot-standby backup candidate path of SR Policy will not work under this circumstance. To provide protection, the packet should be forwarded to another backup Egress PE node of the same service, if it exists.

[3.2.1. Local Repair](#)

In the case of egress PE node failure, the local repair node should forward packet to another Egress PE node.

[I-D.ietf-rtgwg-srv6-egress-protection] provides a method to use Mirror SID for egress protection. The Mirror SID is configured on the backup egress PE to protect the primary egress PE, and it will be used by the repair node to encode the segment list of repair path.

[3.2.2. Ingress Node Switchover](#)

If there are multiple egress PE nodes, the ingress PE node receives all their advertisements of the same service, and builds paths for each of them respectively. The ingress PE node may use Fast Reroute (FRR) for these different paths. When the primary egress PE node fails, the ingress node steers the flow to the path belonging to another egress PE node for protection.

BFD can be used to monitor the liveness of the service SID, locator or interface address of the egress PE node. If the BFD session is down, the egress PE node is deemed to be unreachable.

Service protection and path protection may both be used in the same SRv6 network. Among the different paths to the same egress PE node and the paths to different egress PE nodes, one is selected as the primary path and others are used as backup. The priorities of multiple backup paths may be decided by the egress-node-first strategy or the TE-first strategy.

By the Egress-node-first strategy, paths to the primary egress PE nodes are prioritized. For example, if a failure occurs on the primary path, the ingress PE node will select another path still leading to the primary egress PE nodes. Unless all the paths to the primary egress PE node are failed, the ingress PE node would use the path to the backup egress PE node.

By the TE-first strategy, SRv6 TE paths to any egress PE node have higher priorities than SRv6 BE paths. For example, if a failure occurs on the primary path and there is no other alive SRv6 TE paths to the primary egress PE node, the ingress node will select an SRv6 TE path to the backup egress PE node, rather than an SRv6 BE path still leading to the primary egress PE node.

[4.](#) Implementation Recommendations

This section will introduce the implementation recommendations of protection for SRv6 BE and SRv6 TE scenarios in SRv6 network:

Figure 5 is used as a reference topology in this section. PE1 and PE3 are primary PE nodes for VPN service access. PE2 and PE4 are used as backup. The prefix of CE2, along with VPN service SID, is advertised by BGP routes from PE3 and PE4 to PE1 and PE2. The VPN traffic is from CE1 to CE2.




Figure 5: Reference Topology

The link metrics are configured as follows:

- o Metrics of PE1-P2, PE2-P1, P1-P4, P2-P3, P3-P6, P4-P5, P5-P8, P6-P7, P7-PE4, and P8-PE3 links are 11.
- o Metrics of all other links are 5.
- o Link metrics are bidirectional.

All P and PE nodes are capable of G-SRv6 compression. The SRv6 SIDs are configured as follows:

| | Locator | End.DT |
|-----|------------------|-------------------|
| PE1 | 2001:DB8:A1::/48 | 2001:DB8:A1:100:: |
| PE2 | 2001:DB8:A2::/48 | 2001:DB8:A2:100:: |
| PE3 | 2001:DB8:A3::/48 | 2001:DB8:A3:100:: |
| PE4 | 2001:DB8:A4::/48 | 2001:DB8:A4:100:: |
| P1 | 2001:DB8:B1::/48 | - |
| P2 | 2001:DB8:B2::/48 | - |
| P3 | 2001:DB8:B3::/48 | - |
| P4 | 2001:DB8:B4::/48 | - |
| P5 | 2001:DB8:B5::/48 | - |
| P6 | 2001:DB8:B6::/48 | - |
| P7 | 2001:DB8:B7::/48 | - |
| P8 | 2001:DB8:B8::/48 | - |

| | End | End with COC |
|-----|-----------------|-----------------|
| PE1 | 2001:DB8:A1:1:: | 2001:DB8:A1:2:: |
| PE2 | 2001:DB8:A2:1:: | 2001:DB8:A2:2:: |
| PE3 | 2001:DB8:A3:1:: | 2001:DB8:A3:2:: |
| PE4 | 2001:DB8:A4:1:: | 2001:DB8:A4:2:: |
| P1 | 2001:DB8:B1:1:: | 2001:DB8:B1:2:: |
| P2 | 2001:DB8:B2:1:: | 2001:DB8:B2:2:: |
| P3 | 2001:DB8:B3:1:: | 2001:DB8:B3:2:: |
| P4 | 2001:DB8:B4:1:: | 2001:DB8:B4:2:: |

P1 2001:DB8:B5:1:: 2001:DB8:B5:2::
P2 2001:DB8:B6:1:: 2001:DB8:B6:2::
P3 2001:DB8:B7:1:: 2001:DB8:B7:2::
P4 2001:DB8:B8:1:: 2001:DB8:B8:2::

The SR Policies on PE1 are configured as follows:

SR Policy 1 (Strict Path to PE3)

Candidate Path 1

Preference: 20

Segment List: 2001:DB8:B1:2::, 2001:DB8:B3:2::, 2001:DB8:B5:2::,
2001:DB8:B7:2::, 2001:DB8:A3:1::

Candidate Path 2

Preference: 10

Segment List: 2001:DB8:B2:2::, 2001:DB8:B4:2::, 2001:DB8:B6:2::,
2001:DB8:B8:2::,2001:DB8:A3:1::

SR Policy 2 (Loose Path to PE4)

Candidate Path 1

Preference: 20

Segment List: 2001:DB8:B4:2::, 2001:DB8:B8:2::,2001:DB8:A4:1::

4.1. SRv6 BE

In this scenario, SRv6 BE paths are used to steer the VPN service.
The deployments of protection are as follows:

- o All nodes enable TI-LFA for local protection.

- o All nodes enable BFD for links and neighbors.
- o Ingress PE node enables FRR of SRv6 BE path to backup egress PE node for service protection.
- o Ingress PE node enables BFD for locator of egress PE node to monitor the liveness of SRv6 BE path.

PE1 installs the SRv6 BE path to PE3 with destination address 2001:DB8:A3:100:: as the primary next-hop for the VPN flow. Meanwhile, PE1 also installs the SRv6 BE path to PE4 with destination address 2001:DB8:A4:100:: as the backup next-hop.


PE1 enables BFD for locator 2001:DB8:A3::/48 and 2001:DB8:A4::/48 to monitor the liveness of SRv6 BE paths.


TI-LFA is enabled on all nodes. Take P5 for example. The shortest path from P5 to PE3 is via neighbor P7. In order to provide local protection for P7 node failure, P5 computes and installs the repair path P5->P6->P8->PE3, using [2001:DB8:B6:2::, B8:2, A3:1] as the G-SRv6 SID list.

All nodes use BFD to monitor the liveness of links and adjacent nodes.

Under normal circumstances, PE1 encapsulates the VPN payload in an outer IPv6 header where the destination address is 2001:DB8:A3:100::.

Assume that a failure occurs on P7. The fail-timer of BFD echo from P5 to P7 expires, so P5 perceives the failure. When P5 forwards the VPN packet, the TI-LFA repair path is used. Then, P5 encapsulates the packet in an outer IPv6 Header with SRH carrying a compressed segment-list of [2001:DB8:B6:2::, B8:2, A3:1], as shown in the following figure. The packet is forwarded in the repair path P5->P6->P8->PE3 according to the outer IPv6 Header and SRH. So the failure is repaired by local protection.





Assume that a failure occurs on PE3. TI-LFA does not work and the packets along the SRv6 BE path are dropped. Then the BFD session from PE1 to locator 2001:DB8:A3::/48 is down, so PE1 triggers the switchover to the SRv6 BE path to PE4 and encapsulates the VPN payload in an outer IPv6 header where the destination address is 2001:DB8:A4:100::. After that, the VPN traffic from CE1 to CE2 is recovered.

4.2. SRv6 TE

In this scenario, the SRv6 TE strict path with G-SRv6 compression is used to steer the VPN traffic flows to the primary egress node PE3, and the SRv6 TE loose path with G-SRv6 compression is used for the backup egress node PE4.

The deployments of protection are as follows:

- o In the SR Policy of SRv6 TE strict path, disjoint backup candidate path is used as hot standby for end-to-end protection.
- o Ingress PE node uses SRv6 BE paths as backup for end-to-end protection of SRv6 TE paths.
- o Ingress PE node enables BFD for SR Policy. In the case of SRv6 TE strict path, the reverse path of BFD packet keeps consistent with forward path.
- o Ingress PE node enables BFD for locator of egress PE node to monitor the liveness of SRv6 BE path.
- o Ingress PE node enables FRR of paths to backup egress PE node for service protection.

- o All nodes enable TI-LFA for local protection. All nodes enable BFD for links and neighbors.

PE1 installs SR Policy 1, which is the SRv6 TE strict path to PE3, as the primary next-hop for the VPN flow. SR Policy 1 has two disjoint candidate paths. The candidate path with higher preference is selected as the primary candidate path, and the candidate path with lower preference is selected as hot standby backup.

Meanwhile, the SRv6 BE path to PE3, the SRv6 TE loose path to PE4 (SR Policy 2), and the SRv6 BE path to PE4 are also installed as backup next-hops. The priorities of multiple backup paths may be decided by either of the egress-node-first strategy or the TE-first strategy.

Egress-node-first strategy:

- o primary: SRv6 TE path to primary egress node PE3 (SR Policy 1)
- o backup(1st priority): SRv6 BE path to primary egress node PE3
- o backup(2nd priority): SRv6 TE path to backup egress node PE4 (SR Policy 2)
- o backup(3rd priority): SRv6 BE path to backup egress node PE4

TE-first strategy:

- o primary: SRv6 TE path to primary egress node PE3 (SR Policy 1)

- o backup(1st priority): SRv6 TE path to backup egress node PE4 (SR Policy 2)
- o backup(2nd priority): SRv6 BE path to primary egress node PE3
- o backup(3rd priority): SRv6 BE path to backup egress node PE4

Egress-node-first strategy is used as an example below.

PE1 enables BFD for SR Policy 1 and SR Policy 2 to monitor the liveness of SRv6 TE paths. For SR Policy 1 which is the strict path, the forward and reverse paths of BFD packet should be the same. For

example, the primary path of SR Policy 1 is PE1->P1->P3->P5->P7->PE3, so the reverse path should be PE3->P7->P5->P3->P1->PE1. A segment list of such reverse path is installed on PE3, and the BSID is 2001:DB8:A3:200. PE1 may send BFD echo packet with the segment list of SR Policy 1 along with the BSID of reverse path, which is [2001:DB8:B1:2::, B3:2, B5:2, B7:2, A3:1, 2001:DB8:A3:200]. When the BFD echo packet is forwarded along the strict path to PE3, PE3 will add an outer IPv6 header with SRH carrying the segment list of [2001:DB8:B7:2::, B5:2, B3:2, B1:2, A1:1], which instructs the packet to be forwarded along the same strict path back to PE1.

PE1 enables BFD for locator 2001:DB8:A3::/48 and 2001:DB8:A4::/48 to monitor the liveness of SRv6 BE paths.

TI-LFA is enabled on all nodes. BFD are used to monitor the liveness of links and adjacent nodes.

Under normal circumstances, PE1 encapsulates the VPN payload in an outer IPv6 header with SRH carrying the segment list of primary candidate path of SR Policy 1 along with the VPN SID advertised by PE3. Using G-SRv6 compression, the segment list will be encoded as [2001:DB8:B1:2::, B3:2, B5:2, B7:2, A3:1, 2001:DB8:A3:100::].

Assume that a failure occurs on P3. The packets are dropped since the failed P3 is on the path. The BFD session of the segment list in the primary candidate path of SR Policy 1 is down, so PE1 triggers the switchover to the backup candidate path of SR Policy 1. Then PE1 encapsulates the VPN payload in an outer IPv6 header with SRH carrying the segment list of [2001:DB8:B2:2::, B4:2, B6:2, B8:2, A3:1, 2001:DB8:A3:100::].

Before the recovery of P3, assume that P6 also fails. The BFD session of the segment list in the backup candidate path of SR Policy 1 is also down. Then PE1 triggers the switchover to the 1st

priority backup next-hop which is the SRv6 BE path to PE3. PE1 encapsulates the VPN payload in an outer IPv6 header where the destination address is 2001:DB8:A3:100::.

Assume that a failure occurs on PE3. Both the BFD sessions of SR Policy 1 and locator 2001:DB8:A3::/48 are down, which means the primary next-hop and the 1st priority backup next-hop are down. So PE1 triggers the switchover to the 2nd priority backup next-hop, which is the SRv6 TE loose path to PE4. Then PE1 encapsulates the

VPN payload in an outer IPv6 header with SRH carrying the segment list of [2001:DB8:B4:2::, B8:2, A4:1, 2001:DB8:A4:100::].

Before the recovery of PE3, assume that a failure occurs on P6. The fail-timer of BFD echo from P4 to P6 expires, so P4 perceives the failure. When P4 forwards the VPN packet, the TI-LFA repair path is used. Then, P4 encapsulates the packet in an outer IPv6 Header with SRH carrying a compressed segment-list of [2001:DB8:B3:2::, B5:2, A8:1]. The packet is forwarded in the repair path P4->P3->P5->P8 according to the outer IPv6 Header and SRH. So the failure is repaired by local protection.

Before the recovery of PE3, assume that a failure occurs on P8. When P6 forwards the VPN packet to destination address 2001:DB8:B8:2:: which is one of the segments in the segment list of SRH, the TI-LFA on P6 does not work, since the failed node P8 is the destination. So the packets are dropped. The BFD session of SR Policy 2 is down, and PE1 triggers the switchover to the 3rd priority backup next-hop which is the SRv6 BE path to PE4. Then PE1 encapsulates the VPN payload in an outer IPv6 header where the destination address is 2001:DB8:A4:100::. If the routing convergence is not completed at the moment, P6 will use TI-LFA repair path P6->P5->P7->PE4 to forward the packet. After the routing convergence is done, P nodes will forward the packet along new shortest path excluding P8.

[5. Security Considerations](#)

TBD.

[6. IANA Considerations](#)

This document has no IANA actions.

[7. Contributors](#)

In addition to the authors listed on the front page, the following co-authors have also contributed to this document:

Liu, et al.

Expires September, 2022

[Page 15]

Internet-Draft

SRv6 Protection Considerations

March 2022

Mengxiao Chen

H3C

Email: chen.mengxiao@h3c.com

[8. References](#)

8.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in [RFC 2119](#) Key Words", [BCP 14](#), [RFC 8174](#), May 2017
- [RFC8754] Filsfils, C., Ed., Dukes, D., Ed., Previdi, S., Leddy, J., Matsushima, S., and D. Voyer, "IPv6 Segment Routing Header (SRH)", [RFC 8754](#), DOI 10.17487/RFC8754, March 2020, <<https://www.rfc-editor.org/info/rfc8754>>.
- [I-D.ietf-spring-segment-routing-policy] Filsfils, C., Talaulikar, K., Voyer, D., Bogdanov, A., and P. Mattes, "Segment Routing Policy Architecture", [draft-ietf-spring-segment-routing-policy-18](#) (work in progress), February 2022.
- [I-D.ietf-spring-srv6-srh-compression] Cheng, W., Filsfils, C., Li, Z., Decraene, B., Cai, D., Clad, F., Zadok, S., Guichard, J., Aihua, L., Raszuk, R. and C. Li, "Compressed SRv6 Segment List Encoding in SRH", [draft-ietf-spring-srv6-srh-compression-00](#) (work in progress), February 2022.
- [I-D.ietf-rtgwg-segment-routing-ti-lfa] Litkowski, S., Bashandy, A., Filsfils, C., Francois, P., Decraene, B., and D. Voyer, "Topology Independent Fast Reroute using Segment Routing", [draft-ietf-rtgwg-segment-routing-ti-lfa-08](#) (work in progress), January 2022.
- [RFC5880] Katz, D. and D. Ward, "Bidirectional Forwarding Detection (BFD)", [RFC 5880](#), DOI 10.17487/RFC5880, June 2010, <<https://www.rfc-editor.org/info/rfc5880>>.
- [RFC7880] Pignataro, C., Ward, D., Akiya, N., Bhatia, M., and S. Pallagatti, "Seamless Bidirectional Forwarding Detection (S-BFD)", [RFC 7880](#), DOI 10.17487/RFC7880, July 2016, <<https://www.rfc-editor.org/info/rfc7880>>.

8.2. Informative References

- [RFC8402] Filsfils, C., Ed., Previdi, S., Ed., Ginsberg, L.,

Decraene, B., Litkowski, S., and R. Shakir, "Segment Routing Architecture", [RFC 8402](#), DOI 10.17487/RFC8402, July 2018, <<https://www.rfc-editor.org/info/rfc8402>>.

[RFC8986] Filsfils, C., Ed., Camarillo, P., Ed., Leddy, J., Voyer, D., Matsushima, S., and Z. Li, "Segment Routing over IPv6 (SRv6) Network Programming", [RFC 8986](#), DOI 10.17487/RFC8986, February 2021, <<https://www.rfc-editor.org/info/rfc8986>>.

[RFC5286] Atlas, A., Ed. and A. Zinin, Ed., "Basic Specification for IP Fast Reroute: Loop-Free Alternates", [RFC 5286](#), DOI 10.17487/RFC5286, September 2008, <<https://www.rfc-editor.org/info/rfc5286>>.

[RFC5714] Shand, M. and S. Bryant, "IP Fast Reroute Framework", [RFC 5714](#), DOI 10.17487/RFC5714, January 2010, <<https://www.rfc-editor.org/info/rfc5714>>.

[I-D.bashandy-rtgwg-segment-routing-uloop] Bashandy, A., Filsfils, C., Litkowski, S., Decraene, B., Francois, P. and P., Psenak, "Loop avoidance using Segment Routing", [draft-bashandy-rtgwg-segment-routing-uloop-12](#) (work in progress), December 2021.

[I-D.ietf-rtgwg-srv6-egress-protection] Hu, Z., Chen, H., Chen, H., Wu, P., Toy, M., Cao, C., He, T., Liu, L., and X. Liu, "SRv6 Path Egress Protection", Work in Progress, Internet-Draft, [draft-ietf-rtgwg-srv6-egress-protection-04](#), 17 October 2021, <<https://www.ietf.org/archive/id/draft-ietf-rtgwg-srv6-egress-protection-04.txt>>.

Yisong Liu
China Mobile
China

Email: liuyisong@chinamobile.com

Weiqiang Cheng
China Mobile
China

Email: chengweiqiang@chinamobile.com

Changwang Lin
New H3C Technologies
China

Email: linchangwang.04414@h3c.com

Xuesong Geng
Huawei Technologies
China

Email: gengxuesong@huawei.com

Yao Liu
ZTE Corp.
China

Email: liu.yao71@zte.com.cn