ANIMA
Internet-Draft
Intended status: Standards Track
Expires: 4 December 2026

L. Du, Ed.
X. Que
F. Deng
X. Gong
W. Wang
BUPT

2 June 2026

Autonomic SRv6 Network Fast Failover Using Bounce-back Strategy with GRASP

draft-du-anima-srv6-failover-grasp-01

Abstract

This document specifies an autonomic fast failover mechanism for SRv6 networks using a bounce-back strategy. It uses GRASP to distribute failover protection information, enabling data plane fast reroute without control plane reconvergence.

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 https://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 4 December 2026.

Copyright Notice

Table of Contents

1. Introduction

Segment Routing over IPv6 (SRv6) [RFC8986] provides a flexible source routing paradigm that enables explicit path specification for traffic engineering and service chaining. This flexibility, however, comes with a trade-off: when any node or link along an explicitly specified SRv6 path fails, the entire path is disrupted. Traditional recovery mechanisms rely on control plane reconvergence, which typically takes seconds to complete.

Existing fast protection mechanisms such as Topology-Independent Loop-Free Alternate (TI-LFA) provide local protection for IGP segments. However, as noted in [RFC9256] Section 9, TI-LFA has inherent limitations:

* SR Policies built with non-protected Adjacency SIDs do not benefit from any local protection.

* Links that span multiple IGP domains cannot benefit from TI-LFA automated local protection.

This document introduces a bounce-back strategy to address these gaps. The strategy provides dual protection across the time dimension:

* For in-flight packets at the moment of failure, the Bouncer node bounces them back to the nearest upstream Anchor node, which re-encapsulates the traffic with a pre-computed backup SRv6 segment list.

* For new packets arriving after failure detection, the Anchor node directly encapsulates them onto the backup path.

The mechanism uses GRASP [RFC8990] to autonomically distribute failover protection information during path setup. The ACP [RFC8994] and BRSKI [RFC8995] provide the security foundation.

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.

2. Terminology and Abbreviations

This document uses the terminology defined in [RFC7575] and [RFC8986].

Anchor Node:
: A node on the primary path that has a pre-computed backup path to reach the destination. When receiving bounced-back traffic, it re-encapsulates with the backup SRv6 segment list.

Bouncer Node:
: A node on the primary path that does NOT have a backup path. It can only bounce traffic upstream when detecting a downstream failure.

Bounce-back:
: The action of forwarding traffic toward the upstream direction when a downstream link failure is detected.

FPM ASA:
: Failover Path Manager Autonomic Service Agent. Manages the computation, distribution, and installation of failover protection information using GRASP.

Flow Identifier:
: A value used to uniquely identify a communication flow. This document uses the IPv6 Flow Label field for this purpose.

Primary Path:
: The main forwarding path computed for normal traffic transmission.

Backup Path:
: A pre-computed alternative path used when the primary path fails.

Cascading Bounce-back:
: When a failure occurs, traffic bounces upstream through one or more Bouncer nodes until reaching an Anchor node that can reroute traffic via a backup path.

Path Initiator:
: The entity responsible for computing paths and distributing protection information via GRASP. Typically the Ingress Edge Node or a centralized controller.

Path Responder:
: A node on the primary path that receives protection information from the Path Initiator via GRASP.

3. Failover Mechanism Overview

The bounce-back failover mechanism operates in two distinct phases: control plane setup using GRASP, and data plane failover execution.

3.1. Control Plane and Data Plane Separation

* Control Plane: Distributes failover protection information during path setup. GRASP is NOT involved in failure detection or failover execution.

* Data Plane: Performs failure detection, executes bounce-back action, and performs backup path re-encapsulation. All failover decisions are made locally based on pre-installed forwarding state.

3.2. Node Roles and Backup Path Requirements

Each node on the primary path is assigned one of two roles:

3.2.1. Anchor Node Definition

An Anchor Node (role=0) has a valid backup path and can re-encapsulate traffic. When receiving bounced-back traffic, it MUST strip the original SRv6 encapsulation, re-encapsulate with its pre-installed backup segment list, and forward toward the backup path next-hop.

3.2.2. Bouncer Node Definition

A Bouncer Node (role=1) does NOT have a valid backup path. When detecting downstream failure or receiving bounced-back traffic, it MUST forward upstream without modification.

3.2.3. Backup Path Requirements

A valid backup path MUST satisfy: Reachability, Disjointness (MUST NOT traverse any primary path link between Anchor and destination), First-hop Difference, and SRv6 Expressibility.

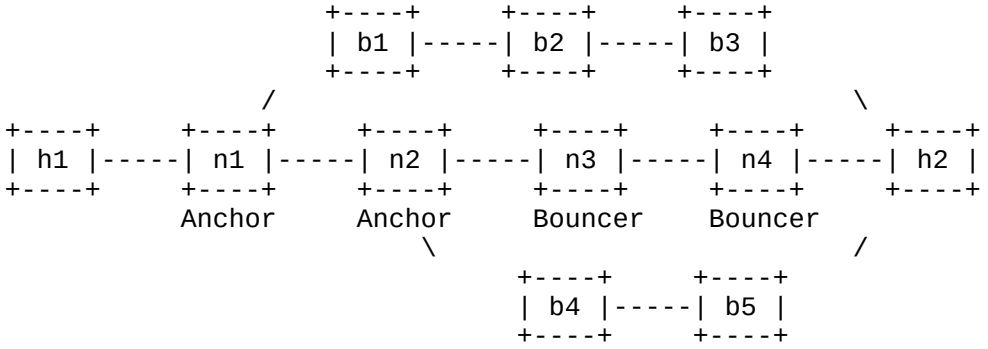

Primary Path: h1 -> n1 -> n2 -> n3 -> n4 -> h2

Node Roles:
n1: Anchor (backup: n1 -> b1 -> b2 -> b3 -> h2)
n2: Anchor (backup: n2 -> b4 -> b5 -> h2)
n3: Bouncer (no backup path)
n4: Bouncer (no backup path)

Backup Path Validity:
A valid backup MUST NOT traverse any primary path link downstream of the Anchor. For n2: backup via b4 -> b5 -> h2 is valid (avoids n2-n3, n3-n4, n4-h2 links)

Figure 1: Topology with Node Roles and Backup Paths

3.3. Before Failure: Protection Information Distribution

Before failure, protection information must be distributed to all path nodes via GRASP, including node role, upstream/downstream neighbor addresses, flow identifier, and backup segment list (for Anchors).

3.4. Failure Detection

Failure detection is performed locally using link-layer detection, BFD, or hardware port monitoring. To achieve sub-50ms failover, hardware-assisted detection SHOULD be used.

3.5. After Failure: Bounce-back and Re-encapsulation

When failure occurs: (1) Node detects downstream interface failure, (2) Immediately redirects traffic upstream, (3) Upstream nodes identify bounced traffic by arrival interface and flow-id, (4) Bouncer forwards upstream; Anchor re-encapsulates and forwards on backup path.

4. GRASP Objective for Failover Path Management

4.1. Failover Path Manager Objective

The objective name is "SRv6-Failover" conforming to [RFC8990].
Format in CDDL [RFC8610]:

```
objective = ["SRv6-Failover", objective-flags, loop-count, ?objective-value]
objective-name = "SRv6-Failover"
objective-flags = uint .bits objective-flag
loop-count = 0..255
objective-value = srv6-failover-value
```

Figure 2: GRASP Objective Format

4.2. Objective Value Definition

```
srv6-failover-value = [flow-info, primary-path-info, *node-protection-info]

flow-info = [flow-id, source-address, destination-address, lifetime]
flow-id = uint
source-address = bytes .size 16
destination-address = bytes .size 16
lifetime = uint

primary-path-info = [primary-segment-list, *anchor-backup-entry]
primary-segment-list = [*srv6-sid]
srv6-sid = bytes .size 16
anchor-backup-entry = [anchor-address, backup-segment-list]

node-protection-info = [node-address, node-role, upstream-neighbor, downstream-neighbo
r, ?backup-info]
node-role = &(anchor: 0, bouncer: 1)
backup-info = [backup-segment-list, backup-next-hop]
```

Figure 3: Flow and Path Information Structure

4.3. Objective Example

For the following topology and primary path:

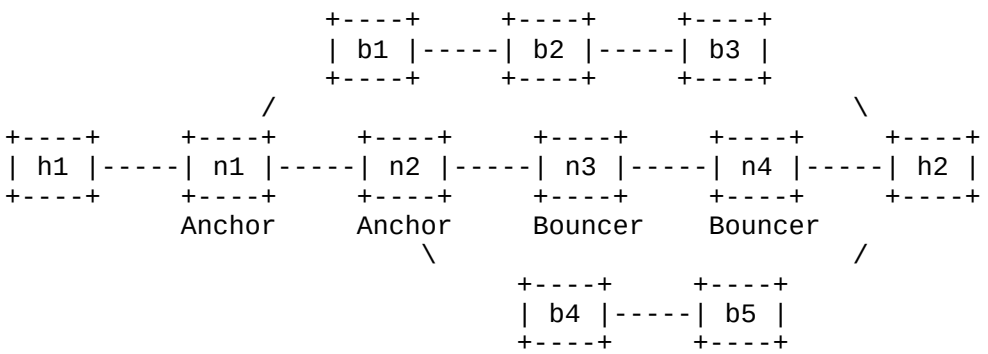

Primary Path: h1 -> n1 -> n2 -> n3 -> n4 -> h2
Flow Identifier: 0x0000a

Figure 4: Example Topology

The GRASP objective value would contain:

```
srv6-failover-value = [
  ; flow-info
  [0x0000a, h1-address, h2-address, 3600000],

  ; primary-path-info
  [
    [n1-sid, n2-sid, n3-sid, n4-sid, h2-sid],  ; primary segment list
    [n1-address, [b1-sid, b2-sid, b3-sid, h2-sid]],  ; n1's backup
    [n2-address, [b4-sid, b5-sid, h2-sid]]           ; n2's backup
  ],

  ; node-protection-info for n1 (Anchor)
  [n1-address, 0, h1-address, n2-address,
   [[b1-sid, b2-sid, b3-sid, h2-sid], b1-address]],

  ; node-protection-info for n2 (Anchor)
  [n2-address, 0, n1-address, n3-address,
   [[b4-sid, b5-sid, h2-sid], b4-address]],

  ; node-protection-info for n3 (Bouncer)
  [n3-address, 1, n2-address, n4-address],

  ; node-protection-info for n4 (Bouncer)
  [n4-address, 1, n3-address, h2-address]
]
```

Figure 5: Objective Value Example

5. Failover Procedures and Scenarios

5.1. GRASP Procedures

The FPM ASA on the Path Initiator discovers path nodes via GRASP Discovery, then sends Request messages with node-protection-info. Path Responders resolve addresses, install forwarding state, and respond with Negotiation End (ACCEPT).

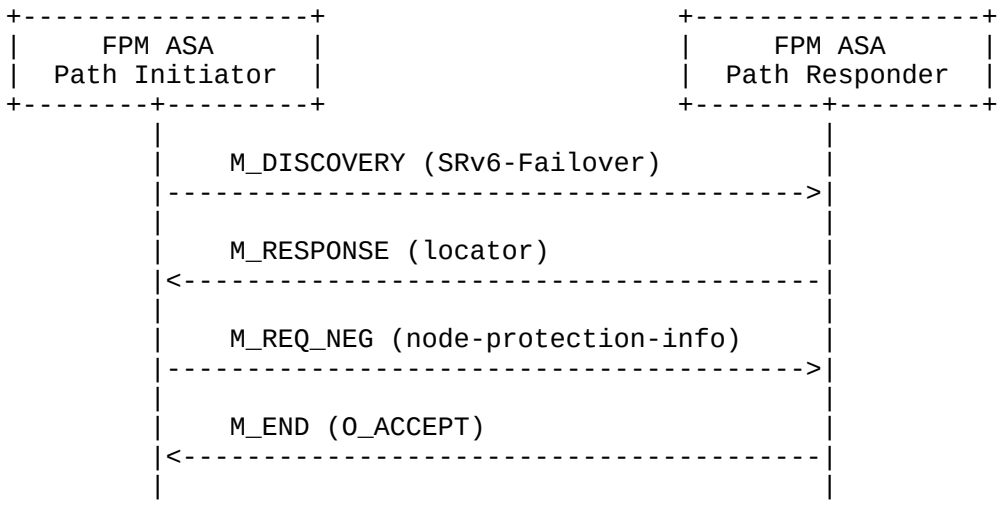

Figure 6: GRASP Negotiation Procedure

5.2. Forwarding Behavior Requirements

Nodes MUST distinguish Normal Traffic (from upstream) and Bounced Traffic (from downstream) based on arrival interface and flow identifier. Anchor nodes re-encapsulate; Bouncer nodes forward upstream.

5.3. Intra-domain Scenario

5.3.1. Topology and Configuration

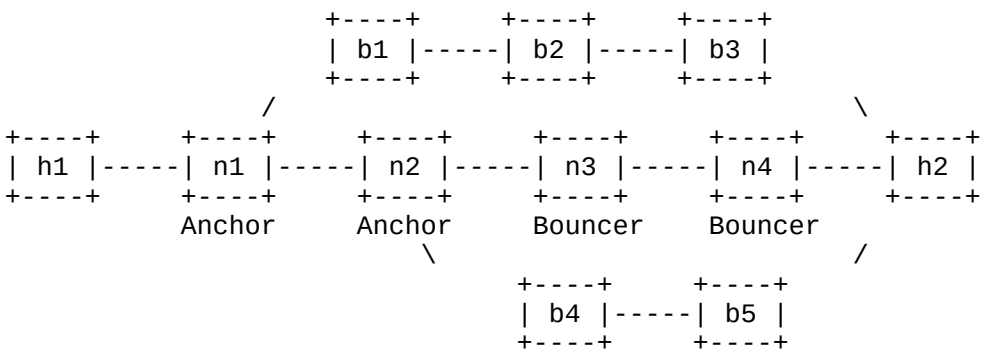

Primary Path: h1 -> n1 -> n2 -> n3 -> n4 -> h2
Flow Identifier: 0x0000a

Protection Configuration:
  n1: Anchor, upstream=h1, downstream=n2
      backup=[b1-sid, b2-sid, b3-sid, h2-sid], next-hop=b1
  n2: Anchor, upstream=n1, downstream=n3
      backup=[b4-sid, b5-sid, h2-sid], next-hop=b4
  n3: Bouncer, upstream=n2, downstream=n4
  n4: Bouncer, upstream=n3, downstream=h2

Figure 7: Intra-domain Topology and Configuration

5.3.2. Failover Execution Example

When the n3-n4 link fails:

1. n3 detects the interface toward n4 is down.

2. n3 (Bouncer) bounces in-flight traffic toward n2.

3. n2 receives traffic from n3 direction with flow-id=0x0000a, identifies it as bounced traffic.

4. n2 (Anchor) re-encapsulates traffic with backup segment list [b4-sid, b5-sid, h2-sid] and forwards toward b4.

5. Traffic reaches h2 via the backup path n2 -> b4 -> b5 -> h2.

If b4-b5 link also fails, n2 bounces traffic to n1, which re-encapsulates via its backup path n1 -> b1 -> b2 -> b3 -> h2 (cascading bounce-back).

5.4. Inter-domain Scenario

Flow identifier is preserved across domain boundaries. Each domain computes its own node roles for its portion of the path.

6. Implementation Considerations

OpenFlow implementations MAY use: in_port, eth_type, ipv6_dst, ipv6_label match fields; fast-failover groups with watch_port.

P4 implementations MAY use: ingress metadata, tables keyed by (ingress_port, flow_label, ipv6_dst), link status registers.

To achieve sub-50ms failover: use hardware-based detection, implement bounce-back in hardware, pre-install forwarding state.

7. Security Considerations

This mechanism inherits security considerations of [RFC8990] and [RFC8986].

* Authentication: All FPM ASA communication MUST occur over ACP [RFC8994]. Nodes MUST be authenticated via BRSKI [RFC8995].
* Authorization: Only authorized nodes SHOULD initiate or modify failover protection state.
* Integrity: GRASP messages MUST be integrity-protected.
* Flow Identifier Security: Ingress filtering SHOULD prevent flow identifier spoofing.
* Resource Exhaustion: Rate limiting of GRASP requests SHOULD be implemented.

8. IANA Considerations

8.1. GRASP Objective Name

IANA is requested to add "SRv6-Failover" to the "GRASP Objective Names" registry. Reference: [this document]

8.2. Node Role Registry

IANA is requested to create "SRv6-Failover Node Role" registry with initial values:

* 0: Anchor - A node with a pre-computed backup path

* 1: Bouncer - A node that can only bounce traffic upstream

Values 2-255 are reserved. New values require Standards Action.

9. References

9.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <https://www.rfc-editor.org/info/rfc2119>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <https://www.rfc-editor.org/info/rfc8174>.

[RFC8610] Birkholz, H., Vigano, C., and C. Bormann, "Concise Data Definition Language (CDDL): A Notational Convention to Express Concise Binary Object Representation (CBOR) and JSON Data Structures", RFC 8610, DOI 10.17487/RFC8610, June 2019, <https://www.rfc-editor.org/info/rfc8610>.

[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>.

[RFC8990] Bormann, C., Carpenter, B., Ed., and B. Liu, Ed., "GeneRic Autonomic Signaling Protocol (GRASP)", RFC 8990, DOI 10.17487/RFC8990, May 2021, <https://www.rfc-editor.org/info/rfc8990>.

[RFC8994] Eckert, T., Ed., Behringer, M., Ed., and S. Bjarnason, "An Autonomic Control Plane (ACP)", RFC 8994, DOI 10.17487/RFC8994, May 2021, <https://www.rfc-editor.org/info/rfc8994>.

[RFC8995] Pritikin, M., Richardson, M., Eckert, T., Behringer, M., and K. Watsen, "Bootstrapping Remote Secure Key Infrastructure (BRSKI)", RFC 8995, DOI 10.17487/RFC8995, May 2021, <https://www.rfc-editor.org/info/rfc8995>.

9.2. Informative References

[RFC7575] Behringer, M., Pritikin, M., Bjarnason, S., Clemm, A., Carpenter, B., Jiang, S., and L. Ciavaglia, "Autonomic Networking: Definitions and Design Goals", RFC 7575, DOI 10.17487/RFC7575, June 2015, <https://www.rfc-editor.org/info/rfc7575>.

[RFC9256] Filsfils, C., Talaulikar, K., Ed., Voyer, D., Bogdanov, A., and P. Mattes, "Segment Routing Policy Architecture", RFC 9256, DOI 10.17487/RFC9256, July 2022, <https://www.rfc-editor.org/info/rfc9256>.

Appendix A. Complete CDDL Definition

This appendix provides the complete CDDL definition for the SRv6-Failover GRASP objective:

```
; SRv6-Failover GRASP Objective CDDL Definition

srv6-failover-objective = [
  "SRv6-Failover",
  objective-flags,
  loop-count,
  ?srv6-failover-value
]

objective-flags = uint
loop-count = 0..255

srv6-failover-value = [
  flow-info,
  primary-path-info,
  *node-protection-info
]

flow-info = [
  flow-id: uint,
  source-address: ipv6-address,
  destination-address: ipv6-address,
  lifetime: uint
]

ipv6-address = bytes .size 16

primary-path-info = [
  primary-segment-list: [*srv6-sid],
  *anchor-backup-entry
]
```

```
srv6-sid = bytes .size 16

anchor-backup-entry = [
  anchor-address: ipv6-address,
  backup-segment-list: [*srv6-sid]
]

node-protection-info = [
  node-address: ipv6-address,
  node-role: 0..1,
  upstream-neighbor: ipv6-address,
  downstream-neighbor: ipv6-address,
  ?backup-info
]

backup-info = [
  backup-segment-list: [*srv6-sid],
  backup-next-hop: ipv6-address
]
```

Figure 8: Complete CDDL Definition

Acknowledgements

The authors thank the contributors of the ANIMA working group for their valuable feedback on autonomic networking mechanisms.

Authors' Addresses

Lintong Du (editor)
Beijing University of Posts and Telecommunications
No.10 Xitucheng Road
Beijing
Hai-Dian District, 100876
China
Email: baronmail@bupt.edu.cn

Xirong Que
Beijing University of Posts and Telecommunications
No.10 Xitucheng Road
Beijing
Hai-Dian District, 100876
China
Email: rongqx@bupt.edu.cn

Fang Deng
Beijing University of Posts and Telecommunications
No.10 Xitucheng Road
Beijing
Hai-Dian District, 100876
China
Email: dengfang@bupt.edu.cn

Xiangyang Gong
Beijing University of Posts and Telecommunications
No.10 Xitucheng Road
Beijing
Hai-Dian District, 100876
China
Email: xygong@bupt.edu.cn

Wendong Wang
Beijing University of Posts and Telecommunications
No.10 Xitucheng Road
Beijing
Hai-Dian District, 100876
China
Email: wdwang@bupt.edu.cn