

Workgroup: LSR Working Group
Internet-Draft:
draft-basavaraj-lsr-ospf-gr-enhancements-08
Published: 26 November 2023
Intended Status: Standards Track
Expires: 29 May 2024
Authors: S. Boutros A. Dubey V. Basavaraj
 VMware VMware VMware
 A. Lindem
 LabN Consulting, L.L.C.

OSPF Graceful Restart Enhancements

Abstract

This document describes enhancements to the OSPF graceful restart procedures to improve routing convergence in some OSPF network deployments. This document updates RFC 3623 and RFC 5187.

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 29 May 2024.

Copyright Notice

Copyright (c) 2023 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 (<https://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 Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

- [1. Introduction](#)
- [2. Conventions used in this document](#)
- [3. Graceful Restart Enhancements](#)
 - [3.1. Stub-Link Network Scenarios](#)
 - [3.2. Multiple Failure Scenarios](#)
- [4. Security Considerations](#)
- [5. IANA Considerations](#)
- [6. Acknowledgements](#)
- [7. References](#)
 - [7.1. Normative References](#)
 - [7.2. Informative References](#)
- [Authors' Addresses](#)

1. Introduction

This document describes the enhancements to the current OSPF [RFC2328] and OSPFv3 [RFC5340] graceful restart procedures as respectively defined in OSPF Graceful Restart [RFC3623] and OSPFv3 Graceful Restart [RFC5187] to improve routing convergence in certain OSPF network deployment scenarios. The goal is for both the restarting OSPF node and the helper OSPF node to terminate the OSPF graceful restart procedure faster and not wait for the grace period expiry in those network scenarios and hence improve the overall OSPF network convergence.

2. Conventions used in this document

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.

3. Graceful Restart Enhancements

3.1. Stub-Link Network Scenarios



Figure 1: OSPF topology with Graceful Restart

As graphically depicted in figure 1, Router-2 is an area border router (ABR) with OSPF links in 2 areas. Furthermore, Router-2 has formed full adjacencies only in Area 0. In Area 1, Router-2 has an OSPF link enabled but Router2 could not form an adjacency either because Router-3 is down or Router-3 does not have OSPF enabled. Hence, Router-2 will only have a stub link in Area 1.

On restart, the ABR router Router-2, having only a stub link in the Area 1, will never receive its pre-restart LSA in this area and will never form an adjacency and will have to wait for the grace period expiry leading to slower OSPF routing convergence.

For this scenario, if no OSPF control packets are received within the dead interval on a link in Area 1 as per the above network scenario, Router-2 MUST mark the link as stub and MUST not wait for the grace period to form an adjacency on this link to successfully Exit GR.

3.2. Multiple Failure Scenarios

In scenarios where more than one router is restarting at the same time in the same OSPF area and StrictLSAChecking is disabled, restarting OSPF routers will end up waiting the entire grace interval to exit GR. If the restarting routers receive a Grace Link State Advertisement (LSA) from another router in a given area after restart, and the helper routers receive grace LSAs from more than one router, this will indicate that there have been multiple failures. Therefore, the helper and restarting routers MUST terminate GR and avoid any unnecessary delay in OSPF routing convergence.

4. Security Considerations

The security considerations in OSPF Graceful Restart [RFC3623] and OSPFv3 Graceful Restart [RFC5187] are applicable. This document does not introduce any additional security considerations.

5. IANA Considerations

This specification doesn't require any IANA assignments.

6. Acknowledgements

TBD

7. References

7.1. 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>>.
- [RFC3623] Moy, J., Pillay-Esnault, P., and A. Lindem, "Graceful OSPF Restart", RFC 3623, DOI 10.17487/RFC3623, November 2003, <<https://www.rfc-editor.org/info/rfc3623>>.
- [RFC5187] Pillay-Esnault, P. and A. Lindem, "OSPFv3 Graceful Restart", RFC 5187, DOI 10.17487/RFC5187, June 2008, <<https://www.rfc-editor.org/info/rfc5187>>.

[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>>.

7.2. Informative References

[RFC2328] Moy, J., "OSPF Version 2", STD 54, RFC 2328, DOI 10.17487/RFC2328, April 1998, <<https://www.rfc-editor.org/info/rfc2328>>.

[RFC5340] Coltun, R., Ferguson, D., Moy, J., and A. Lindem, "OSPF for IPv6", RFC 5340, DOI 10.17487/RFC5340, July 2008, <<https://www.rfc-editor.org/info/rfc5340>>.

Authors' Addresses

Sami Boutros
VMware

Email: sboutros@vmware.com

Ankur Dubey
VMware

Email: adubey@vmware.com

Vijayalaxmi Basavaraj
VMware

Email: vbasavaraj@vmware.com

Acee Lindem
LabN Consulting, L.L.C.
301 Midenhall Way
Cary, NC 27513
United States of America

Email: acee.ietf@gmail.com