

Network Working Group
Internet Draft
Expiration Date: September 2004
File name: [draft-ietf-idr-bgp-gr-survey-01.txt](#)

John G. Scudder
Cisco Systems
Editor
March 2004

BGP Graceful Restart - Implementation Survey
draft-ietf-idr-bgp-gr-survey-01.txt

Status of this Memo

This document is an Internet-Draft and is subject to all provisions of [Section 10 of RFC2026](#).

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/1id-abstracts.html>

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

Abstract

This document provides a survey of BGP-4 Graceful Restart implementations.

1. Survey Summary

This document provides a survey of BGP-4 Graceful Restart [\[1\]](#) implementations. After a brief summary, each response is listed. The editor makes no claim as to the accuracy of the information provided.

The following organizations reported having implementations of Graceful Restart: Cisco Systems, IP Infusion, Juniper Networks, Laurel Networks, Redback Networks, Riverstone Networks and Tenor Networks. Andrew Partan responded regarding independent interoperability testing.

Respondents reported having tested interoperability between the following:

	Cisco	Juniper	Laurel	Redback	Riverstone
Cisco	-	X	X	X	X
Juniper	X	-	X	X	X

In addition to responses from implementors, Andrew Partan <asp@partan.com> reports that he has tested interoperation between Cisco, Juniper and Redback.

2. Survey Forms

2.1. Cisco Systems, Inc.

Person filling out this form: Ruchi Kapoor <ruchi@cisco.com>

Does your Graceful Restart implementation do the following as defined in [draft-ietf-idr-restart-05.txt](#)?

Send the End-of-RIB marker upon completion of initial routing update according to [section 4](#)?

Yes.

Exchange the Graceful Restart capability according to [section 5](#)?

Yes.

Restart BGP sessions normally (i.e., a full restart not a Graceful Restart) when terminated with a NOTIFICATION message according to [section 6](#)?

Yes.

Function as a Restarting Speaker according to [section 6.1](#)?

Yes.

Function as a Receiving Speaker according to [section 6.2](#)?

Yes.

List other implementations that you have tested for Graceful Restart interoperability.

JunOS.

2.2. IP Infusion, Inc.

Person filling out this form: Kunihiro Ishiguro
<kunihiro@ipinfusion.com>, Wei Cao <wei@ipinfusion.com>

Does your Graceful Restart implementation do the following as defined in [draft-ietf-idr-restart-05.txt](#)?

Send the End-of-RIB marker upon completion of initial routing update according to [section 4](#)?

Yes.

Exchange the Graceful Restart capability according to [section 5](#)?

Yes, except the "only be a receiver" case.

Restart BGP sessions normally (i.e., a full restart not a Graceful Restart) when terminated with a NOTIFICATION message according to [section 6](#)?

Yes.

Function as a Restarting Speaker according to [section 6.1](#)?

Yes. And we support both planned and unplanned restart.

Function as a Receiving Speaker according to [section 6.2](#)?

Yes.

List other implementations that you have tested for Graceful Restart interoperability.

None.

2.3. Juniper Networks

Person filling out this form: Chaitanya Kodeboyina
<ck@juniper.net>

Does your Graceful Restart implementation do the following as defined in [draft-ietf-idr-restart-05.txt](#)?

Graceful restart functionality is not on by default and one has to configure 'graceful-restart' for BGP and other protocols to implement this functionality. So all answers below are conditional on this knob being configured. (There is one knob for all protocols.)

Send the End-of-RIB marker upon completion of initial routing update according to [section 4](#)?

Yes.

Exchange the Graceful Restart capability according to [section 5](#)?

Yes.

Restart BGP sessions normally (i.e., a full restart not a Graceful Restart) when terminated with a NOTIFICATION message according to [section 6](#)?

Yes.

Function as a Restarting Speaker according to [section 6.1](#)?

Yes, except that we don't have a capability today (needs a knob at most) to restrict graceful-restart to planned restarts only. But this is an internal matter which has no bearing on interop and I don't see why this paragraph should be part of the draft:

If one wants to apply graceful restart only when the restart is planned (as opposed to both planned and unplanned restart), then one way to accomplish this would be to set the Forwarding State bit to 1 after a planned restart, and to 0 in all other cases. Other approaches to accomplish this are outside the scope of this document.

Function as a Receiving Speaker according to [section 6.2](#)?

Yes.

List other implementations that you have tested for Graceful Restart interoperability.

None.

Expires September 2004

[Page 4]

2.4. Laurel Networks, Inc.

Person filling out this form: Ardas Cilingiroglu
<ardas@laurelnetworks.com>

Does your Graceful Restart implementation do the following as defined in [draft-ietf-idr-restart-05.txt](#)?

Send the End-of-RIB marker upon completion of initial routing update according to [section 4](#)?

Yes. End-of-RIB is generated upon completion of the initial update messages if graceful-restart is configured, even when it's a normal bgp (re)start.

Exchange the Graceful Restart capability according to [section 5](#)?

Yes.

Restart BGP sessions normally (i.e., a full restart not a Graceful Restart) when terminated with a NOTIFICATION message according to [section 6](#)?

Yes.

Function as a Restarting Speaker according to [section 6.1](#)?

Yes. We support both planned and unplanned restarts. Also, implemented a configurable upper bound to defer route selection.

Function as a Receiving Speaker according to [section 6.2](#)?

Yes. Implemented a configurable upper bound to retain stale Rib-In routes.

List other implementations that you have tested for Graceful Restart interoperability.

Cisco and Juniper.

2.5. Redback Networks, Inc.

Person filling out this form: Jenny Yuan <jenny@redback.com>

Does your Graceful Restart implementation do the following as defined in [draft-ietf-idr-restart-05.txt](#)?

Send the End-of-RIB marker upon completion of initial routing update according to [section 4](#)?

Yes.

Exchange the Graceful Restart capability according to [section 5](#)?

Yes.

Restart BGP sessions normally (i.e., a full restart not a Graceful Restart) when terminated with a NOTIFICATION message according to [section 6](#)?

Yes.

Function as a Restarting Speaker according to [section 6.1](#)?

Yes.

Function as a Receiving Speaker according to [section 6.2](#)?

Yes.

List other implementations that you have tested for Graceful Restart interoperability.

Cisco.

2.6. Riverstone Networks

Person filling out this form: Greg Hankins
<ghankins@riverstonenet.com>

Does your Graceful Restart implementation do the following as defined in [draft-ietf-idr-restart-05.txt](#)?

Send the End-of-RIB marker upon completion of initial routing update according to [section 4](#)?

Yes.

Exchange the Graceful Restart capability according to [section 5](#)?

Yes.

Restart BGP sessions normally (i.e., a full restart not a Graceful Restart) when terminated with a NOTIFICATION message according to [section 6](#)?

Yes.

Function as a Restarting Speaker according to [section 6.1](#)?

Yes.

Function as a Receiving Speaker according to [section 6.2](#)?

Yes.

List other implementations that you have tested for Graceful Restart interoperability.

Cisco, Juniper.

Other information:

I defined a 'Resync-Time' to cover the following:

a) Restarter

This timer is started after establishment, and sets an upper limit on time the restarter will wait for EOR markers before sending routes.

b) Helper

This timer is started at the same point, and is the upper limit on the time taken by the restarter to refresh the 'stale' routes.

2.7. Tenor Networks

Person filling out this form: Jim Tsillas
<jtsillas@tenornetworks.com>

Does your Graceful Restart implementation do the following as defined in [draft-ietf-idr-restart-05.txt](#)?

Send the End-of-RIB marker upon completion of initial routing update according to [section 4](#)?

Yes

Exchange the Graceful Restart capability according to [section 5](#)?

Yes

Restart BGP sessions normally (i.e., a full restart not a Graceful Restart) when terminated with a NOTIFICATION message according to [section 6](#)?

Yes

Function as a Restarting Speaker according to [section 6.1](#)?

Yes

Function as a Receiving Speaker according to [section 6.2](#)?

Yes

List other implementations that you have tested for Graceful Restart interoperability.

(None given.)

3. References

- [1] Sangli, S., Y. Rekhter, R. Fernando, J. Scudder and E. Chen,
"Graceful Restart Mechanism for BGP," Work in Progress ([draft-ietf-idr-restart-08.txt](#)), September 2003.

4. Security Considerations

This document does not address any security issues.

5. IANA Considerations

No parameters are defined in this document.

6. Author's Address

John G. Scudder
Cisco Systems, Inc.
100 S. Main Suite 200
Ann Arbor, MI 48104
Email: jgs@cisco.com

7. Full Copyright Statement

Copyright (C) The Internet Society (2004). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assigns.

This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING

TASK FORCE DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.