

Network Working Group
INTERNET-DRAFT
Expires in: January 2006

Scott Poretsky
Reef Point Systems

Shankar Rao
Qwest Communications

July 2005

**Methodology for Benchmarking Accelerated Stress
with Operational EBGP Instabilities**
<[draft-ietf-bmwg-acc-bench-meth-ebgp-00.txt](#)>

Intellectual Property Rights (IPR) statement:

By submitting this Internet-Draft, each author represents that any applicable patent or other IPR claims of which he or she is aware have been or will be disclosed, and any of which he or she becomes aware will be disclosed, in accordance with [Section 6 of BCP 79](#).

Status of this Memo

This document is an Internet-Draft and is in full conformance with 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/ietf/1id-abstracts.txt>

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

ABSTRACT

Routers in an operational network are simultaneously configured with multiple protocols and security policies while forwarding traffic and being managed. To accurately benchmark a router for deployment it is necessary that the router be tested in these simultaneous operational conditions, which is known as Stress Testing. This document provides the Methodology for performing Stress Benchmarking of networking devices when subjected to instability with eBGP-4. Descriptions of Test Topology, Benchmarks and Reporting Format are provided in addition to procedures for conducting various test cases. This

methodology is based upon the accelerated stress methodology guidelines [\[6\]](#) and is to be used with the companion terminology document [\[4\]](#).

Poretsky and Rao
1]

[Page

Table of Contents

1.	Introduction	2
2.	Existing definitions	2
3.	Test Setup.....	2
4.	Test Cases.....	3
4.1	Failed Primary EBGP Peer.....	3
4.2	Establish New EBGP Peer.....	3
4.3	BGP Route Explosion.....	4
4.4	BGP Policy Configuration.....	4
4.5	Persistent BGP Flapping.....	5
4.6	BGP Route Flap Dampening.....	6
4.7	Nested Convergence Events.....	6
5.	IANA Considerations.....	7
6.	Security Considerations.....	7
7.	References.....	7
8.	Author's Address.....	8

[1.](#) Introduction

Routers in an operational network are simultaneously configured with multiple protocols and security policies while forwarding traffic and being managed. To accurately benchmark a router for deployment it is necessary that the router be tested in these simultaneous operational conditions, which is known as Stress Testing. This document provides methodologies based upon the accelerated stress methodology guidelines [[6](#)] and is to be used with the companion terminology document [[4](#)]. This document provides the methodologies for performing Stress Benchmarking of networking devices when subjected to instability with eBGP-4. Test cases are provided for such conditions as a failed EBGP peer, establishing a new EBGP peer, BGP Route Explosion, BGP Policy Configuration, Persistent BGP Flapping, Route Flapping with BGP Dampening enabled, Nested Convergence Events, and secure BGP. Descriptions of Test Topology, Benchmarks and Reporting Format are provided in addition to the procedures to be used for conducting various test cases.

[2.](#) Existing definitions

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [BCP 14](#), [RFC 2119](#) [[7](#)]. [RFC 2119](#) defines the use of these key words to help make the intent of standards track documents as clear as possible. While this document uses these keywords, this document is not a standards track document.

Terms specific to Accelerated Stress Benchmarking are defined in [[4](#)].

[3.](#) Test Setup

Test Setup, Test Topologies, Considerations, and Reporting Format
MUST be as described in [\[6\]](#).

Poretsky and Rao
2]

[Page

4. Test Cases

4.1 Failed Primary EBGP Peer

Objective

The purpose of this test is to benchmark the performance of the DUT during stress conditions when losing an EBGP Peer from which most FIB routes have been learned.

Procedure

1. Report Configuration Set
2. Begin Startup Conditions with the DUT
3. Establish Configuration Sets with the DUT
4. Report benchmarks (for stability)
5. Apply Instability Conditions
6. Remove link to EBGP peer with most FIB routes. This SHOULD be achieved by losing physical layer connectivity with a local fiber pull. Loss of the peering session SHOULD cause the DUT to withdraw 100,000 or greater routes.
7. Report benchmarks (for instability)
8. Stop applying all Instability Conditions
9. Report benchmarks (for recovery)
10. Optional - Change Configuration Set and/or Instability Conditions for next iteration

Results

It is expected that there will be significant packet loss until the DUT converges from the lost EBGP link. Other DUT operation should be stable without session loss or sustained packet loss. Recovery time should not be infinite.

4.2 Establish New EBGP Peer

Objective

The purpose of this test is to benchmark the performance of the DUT during stress conditions when establishing a new EBGP Peer from which routes are learned.

Procedure

1. Report Configuration Set
2. Begin Startup Conditions with the DUT
3. Establish Configuration Sets with the DUT
4. Report benchmarks (for stability)
5. Apply Instability Conditions
6. Configure a new EBGP peering session at DUT and peering router. Physical and Data Link Layer connectivity SHOULD already exist to perform this step. Establishment of the peering session MUST result in the DUT learning X routes from the BGP peer and advertising X routes to the BGP peer.

NOTE 1: Number X of BGP Routes MUST be recorded.

NOTE 2: X routes MUST be installed in the FIB and MUST be

verified installed in the FIB by sending data to each NLRI.

7. Report benchmarks (for instability)

Poretsky and Rao
3]

[Page

8. Stop applying all Instability Conditions
9. Report benchmarks (for recovery)
10. Optional - Change Configuration Set and/or Instability Conditions for next iteration

Results

It is expected that there will be zero packet loss as the DUT learns the new routes. Other DUT operation should be stable without session loss or sustained packet loss.

4.3 BGP Route Explosion

Objective

The purpose of this test is to benchmark the performance of the DUT during stress conditions when there is BGP Route Explosion experienced in the network.

Procedure

1. Report Configuration Set
2. Begin Startup Conditions with the DUT
3. Establish Configuration Sets with the DUT
4. Report benchmarks (for stability)
5. Apply Instability Conditions
6. Advertise X BGP routes to the DUT from a single EBGP neighbor.

NOTE 1: Number X of BGP Routes MUST be recorded.

NOTE 2: X routes MUST be installed in the FIB and MUST be verified installed in the FIB by sending data to each NLRI.

7. Report benchmarks (for instability)
8. Stop applying all Instability Conditions, including BGP route advertisement.
9. Report benchmarks (for recovery)
10. Optional - Change Configuration Set and/or Instability Conditions for next iteration

Results

It is expected that there will be no additional packet loss from the advertisement of routes from the BGP neighbor. Other DUT operation should be stable without session loss.

4.4 BGP Policy Configuration

Objective

The purpose of this test is to benchmark the performance of the DUT during stress conditions when there is continuous reconfiguration of BGP Policy at the DUT.

Procedure

1. Report Configuration Set
2. Begin Startup Conditions with the DUT

Poretsky and Rao

[Page

4]

3. Establish Configuration Sets with the DUT
4. Report benchmarks (for stability)
5. Apply Instability Conditions
6. Configure BGP Policy on the DUT for each established neighbor.
The BGP Policy SHOULD filter X% of the routes learned from that neighbor.

NOTE 1: Note that the specific policy configuration to achieve the filtering may be device specific.

NOTE 2: Number X% of BGP Policies MUST be recorded.

NOTE 3: The policies MUST be applied so that routes in the FIB are impacted.

7. Every 30 minutes remove the BGP Policy configuration and then configure it again so that it is reapplied.
8. Report benchmarks (for instability)
9. Stop applying all Instability Conditions, including Policy changes
10. Report benchmarks (for recovery)
11. Optional - Change Configuration Set and/or Instability Conditions for next iteration

Results

It is expected that there will be no packet loss resulting from the continuous configuration and removal of BGP Policy for BGP neighbors. Other DUT operation should be stable without session loss.

4.5 Persistent BGP Flapping

Objective

The purpose of this test is to benchmark the performance of the DUT during stress conditions when flapping BGP Peering sessions for an infinite period.

Procedure

1. Report Configuration Set
2. Begin Startup Conditions with the DUT
3. Establish Configuration Sets with the DUT
4. Report benchmarks (for stability)
5. Apply Instability Conditions
6. Repeatedly flap an IBGP and an EBGp peering session.
This SHOULD be achieved by losing physical layer connectivity via a local interface shutdown/no shutdown every 180 seconds with a delay of 10 seconds between the shut and no shut.
The loss of the EBGp peering session MUST cause the DUT to withdraw 100,000 routes that are re-learned when the session re-establishes. Route Flap Dampening SHOULD NOT be enabled.
7. Report benchmarks (for instability)
8. Stop applying all Instability Conditions, including flapping
9. Report benchmarks (for recovery)

10. Optional - Change Configuration Set and/or Instability
Conditions for next iteration

Poretsky and Rao
5]

[Page

Results

It is expected that there will be significant packet loss from repeated Convergence Events. Other DUT operation should be stable without session loss. Recovery time should not be infinite.

4.6 BGP Route Flap Dampening

Objective

The purpose of this test is to benchmark the performance of the DUT during stress conditions when flapping BGP Peering sessions for an infinite period and route flap dampening is enabled.

Procedure

1. Report Configuration Set
2. Configure Route Flap Dampening on the DUT with DEFAULT parameter values.
3. Begin Startup Conditions with the DUT
4. Establish Configuration Sets with the DUT
5. Report benchmarks (for stability)
6. Apply Instability Conditions
7. Repeatedly flap an IBGP and an EBGP peering session.
This SHOULD be achieved by losing physical layer connectivity via a local interface shutdown/no shutdown every 180 seconds with a delay of 10 seconds between the shut and no shut.
The loss of the EBGP peering session MUST cause the DUT to withdraw 100,000 or greater routes that are re-learned when the session re-establishes.
8. Report benchmarks (for instability)
9. Stop applying all Instability Conditions
10. Report benchmarks (for recovery)
11. Optional - Change Route Flap Dampening parameter values
12. Optional - Change Configuration Set and/or Instability Conditions for next iteration

Results

It is expected that there will be significant packet loss from repeated Convergence Events and flap dampening. Other DUT operation should be stable without session loss. Recovery time should not be infinite.

4.7 Nested Convergence Events [\[5\]](#)

Objective

The purpose of this test is to benchmark the performance of the DUT during stress conditions when flapping BGP Peering sessions causes Nested Convergence Events.

Procedure

1. Report Configuration Set

2. Begin Startup Conditions with the DUT
3. Establish Configuration Sets with the DUT
4. Report benchmarks (for stability)

Poretsky and Rao

[Page

6]

5. Apply Instability Conditions
6. Repeatedly flap an IBGP and an EBGP peering session.
This SHOULD be achieved by losing physical layer connectivity via a local interface shutdown/no shutdown every 10 seconds with a delay of 1 second between the shut and no shut.
The loss of the EBGP peering session MUST cause the DUT to withdraw 100,000 or greater routes that are re-learned when the session re-establishes. Route Flap Dampening SHOULD NOT be enabled.
7. Report benchmarks (for instability)
8. Stop applying all Instability Conditions, including flapping
9. Report benchmarks (for recovery)
10. Optional - Change Configuration Set and/or Instability Conditions for next iteration

Results

It is expected that there will be significant packet loss from Nested Convergence Events. New Other DUT operation should be stable without session loss. Recovery time should not be infinite.

5. IANA Considerations

This document requires no IANA considerations.

6. Security Considerations

Documents of this type do not directly affect the security of the Internet or of corporate networks as long as benchmarking is not performed on devices or systems connected to operating networks.

7. References

7.1 Normative References

- [1] Bradner, S., Editor, "Benchmarking Terminology for Network Interconnection Devices", [RFC 1242](#), July 1991.
- [2] Mandeville, R., "Benchmarking Terminology for LAN Switching Devices", [RFC 2285](#), June 1998.
- [3] Bradner, S. and McQuaid, J., "Benchmarking Methodology for Network Interconnect Devices", [RFC 2544](#), March 1999.
- [4] Poretsky, S. and Rao, S., "Terminology for Accelerated Stress Benchmarking", [draft-ietf-bmwg-acc-bench-term-05](#), work in progress, July 2005.
- [5] Poretsky, S., "Benchmarking Terminology for IGP Data Plane Route Convergence", [draft-ietf-bmwg-igp-dataplane-conv-term-05](#), work in progress, July 2005.
- [6] Poretsky, S. and Rao, S., "Methodology Guidelines for Accelerated Stress Benchmarking", [draft-ietf-bmwg-acc-bench-meth-03](#), work in progress, July 2005.

[7] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [RFC 2119](#), March 1997.

Poretsky and Rao

[Page

7]

7.2 Informative References

- [[RFC3871](#)] [RFC 3871](#) "Operational Security Requirements for Large Internet Service Provider (ISP) IP Network Infrastructure. G. Jones, Ed.. IETF, September 2004.
- [NANOG25] "Core Router Evaluation for Higher Availability", Scott Poretsky, NANOG 25, June 8, 2002, Toronto, CA.
- [IEEECQR] "Router Stress Testing to Validate Readiness for Network Deployment", Scott Poretsky, IEEE CQR 2003.
- [CONVMETH] Poretsky, S., "Benchmarking Methodology for IGP Data Plane Route Convergence", [draft-ietf-bmwg-igp-dataplane-conv-meth-05](#), work in progress, July 2005.
- [NANOG30] Poretsky, S. and Imhoff, B., "BGP Testing: Why be so Negative?", NANOG 30, Miami, Feb 2004.

8. Author's Address

Scott Poretsky
Reef Point Systems
8 New England Executive Park
Burlington, MA 01803
USA
Phone: + 1 781 395 5090
EMail: sporetsky@reefpoint.com

Shankar Rao
1801 California Street
8th Floor
Qwest Communications
Denver, CO 80202
USA
Phone: + 1 303 437 6643
Email: shankar.rao@qwest.com

Full Copyright Statement

Copyright (C) The Internet Society (2005).

This document is subject to the rights, licenses and restrictions contained in [BCP 78](#), and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM 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.

Poretsky and Rao
8]

[Page

INTERNET-DRAFT Methodology for Accelerated Stress Benchmarking July 2005 Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in [BCP 78](#) and [BCP 79](#).

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

Acknowledgement

Funding for the RFC Editor function is currently provided by the Internet Society.

