

Network Working Group
INTERNET-DRAFT
Expires in: July 2004

Scott Poretsky
Quarry Technologies

January 2004

Benchmarking Applicability for IGP Data Plane Route Convergence

<[draft-ietf-bmwg-igp-dataplane-conv-app-02.txt](#)>

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

This draft describes the applicability of IGP Route Convergence benchmarking methodology [[1](#)] and IGP Route Convergence benchmarking terminology [[2](#)]. The methodology and terminology is to be used for benchmarking route convergence and can be applied to any link-state IGP such as ISIS [[3](#)] and OSPF [[4](#)]. The data plane is measured to obtain the convergence benchmarking metrics described in [[1](#)].

Table of Contents

1. Introduction	2
2. Existing definitions	2
3. Factors for IGP Route Convergence Time	2
4. Network Events that Cause Route Convergence	3
5. Use of Data Traffic for IGP Route Convergence Benchmarking	3

[6.](#) Security Considerations.....[4](#)
[7.](#) Acknowledgements.....[4](#)
[8.](#) References.....[4](#)

9. Author's Address.....	5
10. Full Copyright Statement.....	5

1. Introduction

IGP Convergence is a critical performance parameter. Customers of Service Providers use packet loss due to IGP Convergence as a key metric of their network service quality. Service Providers use IGP Convergence time as a key metric of router design and architecture. Fast network convergence can be optimally achieved through deployment of fast converging routers. The fundamental basis by which network users and operators benchmark convergence is packet loss, which is an externally observable event having direct impact on their application performance.

IGP Route Convergence is a Direct Measure of Quality (DMOQ) when benchmarking the data plane. For this reason it is important to develop a standard router benchmarking methodology and terminology for measuring IGP convergence that uses the data plane as described in [\[1\]](#) and [\[2\]](#). This document describes all of the factors that influence a convergence measurement and how a purely black box test can be designed to account for all of these factors. This enables accurate benchmarking and evaluation for route convergence time.

2. Existing definitions

For the sake of clarity and continuity this RFC adopts the template for definitions set out in [Section 2 of RFC 1242](#). Definitions are indexed and grouped together in sections for ease of reference.

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 [RFC 2119](#).

3. Factors for IGP Route Convergence Time

There are four major categories of factors contributing to the measured Router IGP Convergence Time. As discussed in [\[5\]](#), [\[6\]](#), [\[7\]](#), [\[8\]](#) and [\[9\]](#), these categories are Event Detection, SPF Processing, IGP Advertisement, and FIB Update. These have numerous components that influence the convergence time. These are listed as follow:

-Event Detection-

- SONET failure indication time
- PPP failure indication time
- IGP Hello Dead Interval

-SPF Processing-

- SPF Delay Time

SPF Hold time
SPF Execution time

Poretsky

[Page 2]

- IGP Advertisement-

- LSA/LSP Flood Packet Pacing

- LSA/LSP Retransmission Packet Pacing

- LSA/LSP Generation time

- FIB Update-

- Tree Build time

- Hardware Update time

The contribution of each of these factors listed above will vary with each router vendors' architecture and IGP implementation. It is therefore necessary to design a convergence test that considers all of these components, not just one or a few of these components. The additional benefit of designing a test for all components is that it enables black-box testing in which knowledge of the routers' internal implementations is not required. It is then possible to make valid use of the convergence benchmarking metrics when comparing routers from different vendors.

4. Network Events that Cause Convergence

There are different types of network events that can cause IGP convergence. These network events are administrative link removal, unplanned link failure, line card failure, and route changes such as withdrawal, flap, next-hop change, and cost change. When benchmarking a router it is important to measure the convergence time for local and remote occurrence of these network events. The convergence time measured will vary whether the network event occurred locally or remotely due to varying combinations of factors listed in the previous sections. This behavior makes it possible to design purely black-box tests that isolate measurements for each of the components of convergence time.

5. Use of Data Plane for IGP Route Convergence Benchmarking

Customers of service providers use packet loss as the metric to calculate convergence time. Packet loss is an externally observable event having direct impact on customers' application performance. For this reason it is important to develop a standard router benchmarking methodology and terminology that is a Direct Measure of Quality (DMOQ) for measuring IGP convergence. Such a methodology uses the data plane as described in [\[1\]](#) and [\[2\]](#).

An additional benefit of using packet loss for calculation of IGP Route Convergence time is that it enables black-box tests to be designed. Data traffic can be offered to the device under test (DUT), an emulated network event can be forced to occur, and packet loss can be externally measured to calculate the convergence time. Knowledge of the DUT architecture and IGP

implementation is not required. There is no need to rely on the DUT to produce the test results. There is no need to build intrusive test harnesses for the DUT.

Use of data traffic and measurement of packet loss on the data plane also enables Route Convergence methodology test cases that consider the time for the Route Controller to update the FIB on the forwarding engine of the hardware. A router is not fully converged until all components are updated and traffic is rerouted to the correct egress interface. As long as there is packet loss, routes have not converged. It is possible to send diverse traffic flows to destinations matching every route in the FIB so that the time it takes for the router to converge an entire route table can be benchmarked.

6. Security Considerations

Documents of this type do not directly effect 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. Acknowledgements

Thanks to Curtis Villamizar for sharing so much of his knowledge and experience through the years. Also, special thanks to the many Network Engineers and Network Architects at the Service Providers who are always eager to discuss Route Convergence.

8. References

- [1] Poretsky, S., "Benchmarking Methodology for IGP Data Plane Route Convergence", [draft-ietf-bmwg-igp-dataplane-conv-meth-01](#), work in progress, October 2004.
- [2] Poretsky, S., "Benchmarking Terminology for IGP Data Plane Route Convergence", [draft-ietf-bmwg-igp-dataplane-conv-term-01](#), work in progress, October 2004.
- [3] Callon, R., "Use of OSI IS-IS for Routing in TCP/IP and Dual Environments", [RFC 1195](#), December 1990.
- [4] Moy, J., "OSPF Version 2", [RFC 2328](#), IETF, April 1998.
- [5] Villamizar, C., "Convergence and Restoration Techniques for ISP Interior Routing", NANOG 25, October 2002.
- [6] Katz, D., "Why are we Scared of SPF? IGP Scaling and Stability", NANOG 25, October 2002.
- [7] Filsfils, C., "Deploying Tight-SLA Services on an Internet Backbone: ISIS Fast Convergence and Differentiated Services Design (tutorial)", NANOG 25, October 2002.

- [8] Alaettinoglu, C. and Casner, S., "ISIS Routing on the Qwest Backbone: a Recipe for Subsecond ISIS Convergence", NANOG 24, October 2002.
- [9] Alaettinoglu, C., Jacobson, V., and Yu, H., "Towards Millisecond IGP Convergence", NANOG 20, October 2000.

9. Author's Address

Scott Poretsky
Quarry Technologies
8 New England Executive Park
Burlington, MA 01803
USA

Phone: + 1 781 395 5090
EMail: sporetsky@quarrytech.com

10. Full Copyright Statement

Copyright (C) The Internet Society (1998). 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.