

Network Working Group
INTERNET-DRAFT
Expires in: May 2008
Intended Status: Informational

Scott Poretsky
Reef Point Systems

November 2007

Considerations for Benchmarking
Link-State IGP Data Plane Route Convergence

<[draft-ietf-bmwg-igp-dataplane-conv-app-14.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

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/lid-abstracts.txt>.

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

Copyright Notice

Copyright (C) The IETF Trust (2007).

ABSTRACT

This document discusses considerations for benchmarking Interior Gateway Protocol (IGP) Route Convergence for any link-state IGP, such as Intermediate System-Intermediate System (ISIS) and Open-Shorted Path first (OSPF). A companion methodology document is to be used for benchmarking IGP convergence time through externally observable (black box) data plane measurements. A companion terminology document is to be referenced to support the benchmarking.

INTERNET-DRAFT Considerations for Benchmarking November 2007
IGP Data Plane Route Convergence

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 Plane for IGP Route Convergence Benchmarking....	4
6.	IANA Considerations.....	4
7.	Security Considerations.....	4
8.	Acknowledgements.....	5
9.	Normative References.....	5
10.	Author's Address.....	6

[1.](#) Introduction

Convergence Time is a critical performance parameter. Customers of Service Providers use convergence packet loss [[Po07t](#)] due to Interior Gateway Protocol (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 for any IGP such as Intermediate System - Intermediate System (ISIS) [[Ca90](#)] and Open-Shorted Path first (OSPF) [[Mo98](#)]. 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 [[Po07m](#)] and [[Po07t](#)]. 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

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](#) [[Br97](#)]. [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.

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 [[Vi02](#)], [[Ka02](#)], [[Fi02](#)], [[Al02](#)] and [[Al00](#)], these categories are Event Detection, Shortest Path First (SPF) Processing, IGP Advertisement, and Forwarding Information Base (FIB) Update. These have numerous components that influence the convergence time, as listed below:

Poretsky

[Page 2]

INTERNET-DRAFT

Considerations for Benchmarking November 2007
IGP Data Plane Route Convergence

-Event Detection-

- Physical Layer failure/recovery indication time
- Layer 2 failure/recovery indication time
- IGP Hello Dead Interval

-SPF Processing-

- SPF Delay Time
- SPF Hold time
- SPF Execution time

-IGP Advertisement-

- LSA/LSP Flood Packet Pacing
- LSA/LSP Retransmission Packet Pacing
- LSA/LSP Generation time

-FIB Update-

- Tree Build time
- Hardware Update time

-Increased Forwarding Delay due to Queueing

The contribution of each of these factors listed above will vary with each router vendors' architecture and IGP implementation. Routers may have a centralized forwarding architecture, in which one route table is calculated and referenced for all arriving packets, or a distributed forwarding architecture, in which the central route table is calculated and distributed to the interfaces for local look-up as packets arrive. The distributed route tables are typically maintained in hardware.

It is therefore necessary to design a convergence test that considers all of these components contributing to convergence time and is independent of the Device Under Test (DUT) architecture, The benefit of designing a test for these considerations 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 as follow:

- * administrative link removal
- * unplanned link failure
- * line card failure
- * route changes such as withdrawal, flap, next-hop change, and cost change.
- * session loss due to loss of peer or adjacency
- * link recovery
- * link insertion

When benchmarking a router it is important to measure 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 [[Po07m](#)] using the terminology provided in [[Po07t](#)].

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. IANA Considerations

This document requires no IANA considerations.

7. 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 production networks.

Poretsky

[Page 4]

INTERNET-DRAFT Considerations for Benchmarking November 2007
IGP Data Plane Route Convergence

8. Acknowledgements

Thanks to Curtis Villamizar for sharing so much of his knowledge and experience through the years. Thanks to Ron Bonica, Al Morton, David Ward, and the BMWG for their reviews and comments.

9. References

9.1 Normative References

- [Br97] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [RFC 2119](#), March 1997
- [Ca90] Callon, R., "Use of OSI IS-IS for Routing in TCP/IP and Dual Environments", [RFC 1195](#), December 1990.
- [Mo98] Moy, J., "OSPF Version 2", [RFC 2328](#), IETF, April 1998.
- [Po07m] Poretsky, S., "Benchmarking Methodology for Link-State IGP Data Plane Route Convergence", [draft-ietf-bmwg-igp-dataplane-conv-meth-14](#), work in progress, November 2007.

[Po07t] Poretsky, S., "Benchmarking Terminology for Link-State IGP Data Plane Route Convergence", [draft-ietf-bmwg-igp-dataplane-conv-term-14](#), work in progress, November 2007.

9.2 Informative References

[Al00] Alaettinoglu, C., Jacobson, V., and Yu, H., "Towards Millisecond IGP Convergence", NANOG 20, March 2000.

[Al02] Alaettinoglu, C. and Casner, S., "ISIS Routing on the Qwest Backbone: a Recipe for Subsecond ISIS Convergence", NANOG 24, March 2002.

[Fi02] Filsfils, C., "Deploying Tight-SLA Services on an Internet Backbone: ISIS Fast Convergence and Differentiated Services Design (tutorial)", NANOG 25, March 2002.

[Ka02] Katz, D., "Why are we Scared of SPF? IGP Scaling and Stability", NANOG 25, March 2002.

[Vi02] Villamizar, C., "Convergence and Restoration Techniques for ISP Interior Routing", NANOG 25, March 2002.

Poretsky

[Page 5]

INTERNET-DRAFT

Considerations for Benchmarking November 2007
IGP Data Plane Route Convergence

10. Author's Address

Scott Poretsky
Reef Point Systems
3 Federal Street
Billerica, MA 01821 USA
Phone: + 1 508 439 9008
EMail: sporetsky@reefpoint.com

Full Copyright Statement

Copyright (C) The IETF Trust (2007).

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, THE

IETF TRUST 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.

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.