

Network Working Group
INTERNET-DRAFT
Expires in: June 2006

Scott Poretsky
Reef Point Systems

Brent Imhoff
Juniper Networks

January 2006

Terminology for Benchmarking IGP Data Plane Route Convergence

<[draft-ietf-bmwg-igp-dataplane-conv-term-09.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/1id-abstracts.txt>.

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

Copyright Notice

Copyright (C) The Internet Society (2006).

ABSTRACT

This document describes the terminology for benchmarking IGP Route Convergence as described in Applicability document [1] and Methodology document [2]. The methodology and terminology are to be used for benchmarking Convergence Time 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 [2].

Table of Contents

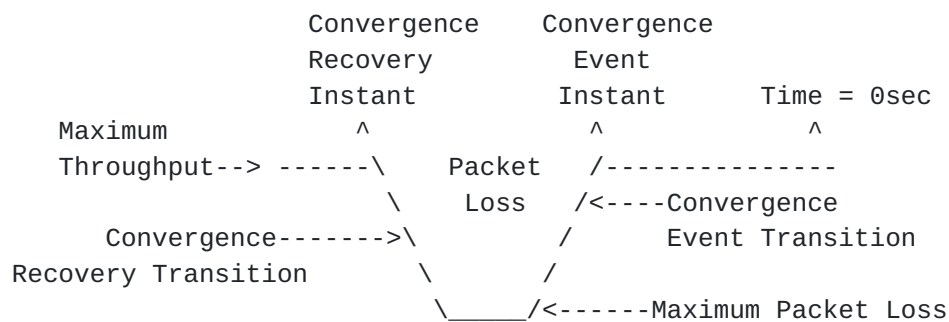
1.	Introduction	2
2.	Existing definitions	3
3.	Term definitions.....	3
3.1	Convergence Event.....	3
3.2	Route Convergence.....	4
3.3	Network Convergence.....	4
3.4	Full Convergence.....	5
3.5	Convergence Packet Loss.....	5
3.6	Convergence Event Instant.....	6
3.7	Convergence Recovery Instant.....	6
3.8	Rate-Derived Convergence Time.....	7
3.9	Convergence Event Transition.....	7
3.10	Convergence Recovery Transition.....	8
3.11	Loss-Derived Convergence Time.....	8
3.12	Sustained Forwarding Convergence Time.....	9
3.13	Restoration Convergence Time.....	9
3.14	Packet Sampling Interval.....	10
3.15	Local Interface.....	11
3.16	Neighbor Interface.....	11
3.17	Remote Interface.....	11
3.18	Preferred Egress Interface.....	12
3.19	Next-Best Egress Interface.....	12
3.20	Stale Forwarding.....	13
3.21	Nested Convergence Events.....	13
4.	IANA Considerations.....	13
5.	Security Considerations.....	14
6.	Normative References.....	14
7.	Author's Address.....	14

[1. Introduction](#)

This draft describes the terminology for benchmarking IGP Route Convergence. The motivation and applicability for this benchmarking is provided in [\[1\]](#). The methodology to be used for this benchmarking is described in [\[2\]](#). The methodology and terminology to be used for benchmarking Route Convergence can be applied to any link-state IGP such as ISIS [\[3\]](#) and OSPF [\[4\]](#). The data plane is measured to obtain black-box (externally observable) convergence benchmarking metrics. The purpose of this document is to introduce new terms required to complete execution of the IGP Route Convergence Methodology [\[2\]](#). These terms apply to IPv4 and IPv6 traffic as well as IPv4 and IPv6 IGPs.

An example of Route Convergence as observed and measured from the data plane is shown in Figure 1. The graph in Figure 1 shows Throughput versus Time. Time 0 on the X-axis is on the far right of the graph. The components of the graph and metrics are

defined in the Term Definitions section.



X-axis = Time

Y-axis = Throughput

Figure 1. Convergence Graph

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](#). [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. The term Throughput is defined in [RFC 2544](#).

3. Term Definitions

3.1 Convergence Event

Definition:

The occurrence of a planned or unplanned action in the network that results in a change in the egress interface of the Device Under Test (DUT) for routed packets.

Discussion:

Convergence Events include link loss, routing protocol session loss, router failure, configuration change, and better next-hop learned via a routing protocol.

Measurement Units:

N/A

Issues:

None

See Also:

Convergence Packet Loss

Convergence Event Instant

3.2 Route Convergence

Definition:

Recovery from a Convergence Event indicated by the DUT
Throughput equal to the offered load.

Discussion:

Route Convergence is the action of all components of the router being updated with the most recent route change(s) including the Routing Information Base (RIB) and Forwarding Information Base (FIB), along with software and hardware tables. Route Convergence can be observed externally by the rerouting of data Traffic to a new egress interface.

Measurement Units:

N/A

Issues:

None

See Also:

Network Convergence
Full Convergence
Convergence Event

3.3 Network Convergence

Definition:

The completion of updating of all routing tables, including the FIB, in all routers throughout the network.

Discussion:

Network Convergence is bounded by the sum of Route Convergence for all routers in the network. Network Convergence can be determined by recovery of the Throughput to equal the offered load, with no Stale Forwarding, and no blenders[5][6].

Measurement Units:

N/A

Issues:

None

See Also:

Route Convergence
Stale Forwarding

3.4 Full Convergence

Definition:

Route Convergence for an entire FIB.

Discussion:

When benchmarking convergence, it is useful to measure the time to converge an entire FIB. For example, a Convergence Event can be produced for an OSPF table of 5000 routes so that the time to converge routes 1 through 5000 is measured. Full Convergence is externally observable from the data plane when the Throughput of the data plane traffic on the Next-Best Egress Interface equals the offered load.

Measurement Units:

N/A

Issues: None

See Also:

Network Convergence

Route Convergence

Convergence Event

3.5 Convergence Packet Loss

Definition:

The amount of packet loss produced by a Convergence Event until Route Convergence occurs.

Discussion:

Packet loss can be observed as a reduction of forwarded traffic from the maximum Throughput. Convergence Packet Loss includes packets that were lost and packets that were delayed due to buffering. The maximum Convergence Packet Loss observed in a Packet Sampling Interval may or may not reach 100% during Route Convergence (see Figure 1).

Measurement Units:

number of packets

Issues: None

See Also:

Route Convergence

Convergence Event

Rate-Derived Convergence Time

Loss-Derived Convergence Time

Packet Sampling Interval

3.6 Convergence Event Instant

Definition:

The time instant that a Convergence Event becomes observable in the data plane.

Discussion:

Convergence Event Instant is observable from the data plane as the precise time that the device under test begins to exhibit packet loss.

Measurement Units:

hh:mm:ss:nnn, where 'nnn' is milliseconds

Issues:

None

See Also:

Convergence Event
Convergence Packet Loss
Convergence Recovery Instant

3.7 Convergence Recovery Instant

Definition:

The time instant that Full Convergence is measured and then maintained for an interval of duration equal to the Sustained Forwarding Convergence Time

Discussion:

Convergence Recovery Instant is measurable from the data plane as the precise time that the device under test achieves Full Convergence.

Measurement Units:

hh:mm:ss:uuu

Issues:

None

See Also:

Sustained Forwarding Convergence Time
Convergence Packet Loss
Convergence Event Instant

3.8 Rate-Derived Convergence Time

Definition:

The amount of time for Convergence Packet Loss to persist upon occurrence of a Convergence Event until occurrence of Route Convergence.

Rate-Derived Convergence Time can be measured as the time difference from the Convergence Event Instant to the Convergence Recovery Instant, as shown with Equation 1.

$$\text{(eq 1) Rate-Derived Convergence Time} = \text{Convergence Recovery Instant} - \text{Convergence Event Instant.}$$

Discussion:

Rate-Derived Convergence Time should be measured at the maximum Throughput. Failure to achieve Full Convergence results in a Rate-Derived Convergence Time benchmark of infinity.

Measurement Units:

seconds/milliseconds

Issues:

None

See Also:

Convergence Packet Loss
Convergence Recovery Instant
Convergence Event Instant
Full Convergence

3.9 Convergence Event Transition

Definition:

The characteristic of a router in which Throughput gradually reduces to zero after a Convergence Event.

Discussion:

The Convergence Event Transition is best observed for Full Convergence. The Convergence Event Transition may not be linear.

Measurement Units:

seconds/milliseconds

Issues:

None

See Also:

Convergence Event
Rate-Derived Convergence Time

Convergence Packet Loss
Convergence Recovery Transition

Poretsky, Imhoff

[Page 7]

3.10 Convergence Recovery Transition

Definition:

The characteristic of a router in which Throughput gradually increases to equal the offered load.

Discussion:

The Convergence Recovery Transition is best observed for Full Convergence. The Convergence Event Transition may not be linear.

Measurement Units:

seconds/milliseconds

Issues:

None

See Also:

Full Convergence
Rate-Derived Convergence Time
Convergence Packet Loss
Convergence Event Transition

3.11 Loss-Derived Convergence Time

Definition:

The amount of time it takes for Route Convergence to be achieved as calculated from the Convergence Packet Loss. Loss-Derived Convergence Time can be calculated from Convergence Packet Loss that occurs due to a Convergence Event and Route Convergence.as shown with Equation 2.

(eq 2) Loss-Derived Convergence Time =

Convergence Packets Loss / Offered Load

NOTE: Units for this measurement are
packets / packets/second = seconds

Discussion:

Loss-Derived Convergence Time gives a better than actual result when converging many routes simultaneously. Rate-Derived Convergence Time takes the Convergence Recovery Transition into account, but Loss-Derived Convergence Time ignores the Route Convergence Recovery Transition because it is obtained from the measured Convergence Packet Loss. Ideally, the Convergence Event Transition and Convergence Recovery Transition are instantaneous so that the Rate-Derived Convergence Time = Loss-Derived Convergence Time. However, router implementations are less than ideal.

For these reasons the preferred reporting benchmark for IGP
Route Convergence is the Rate-Derived Convergence Time.

Guidelines for reporting Loss-Derived Convergence Time are provided in [2].

Measurement Units:
seconds/milliseconds

Issues:
None

See Also:
Route Convergence
Convergence Packet Loss
Rate-Derived Convergence Time
Convergence Event Transition
Convergence Recovery Transition

3.12 Sustained Forwarding Convergence Time

Definition:
The amount of time for which Full Convergence is maintained without additional packet loss.

Discussion:
The purpose of the Sustained Forwarding Convergence Time is to produce Convergence benchmarks protected against fluctuation in Throughput after Full Convergence is observed. The Sustained Forwarding Convergence Time to be used is calculated as shown in Equation 3.

(eq 3)
$$\text{Sustained Forwarding Convergence Time} = 5 \text{ packets/Offered Load}$$

units are packets/pps = sec

for which at least one packet per route in the FIB for all routes in the FIB MUST be offered to the DUT per second.

Measurement Units:
seconds or milliseconds

Issues: None

See Also:
Full Convergence
Convergence Recovery Instant

3.13 Restoration Convergence Time

Definition:
The amount of time for the router under test to restore

traffic to the original outbound port after recovery from
a Convergence Event.

Discussion:

Restoration Convergence Time is the amount of time for routes to converge to the original outbound port. This is achieved by recovering from the Convergence Event, such as restoring the failed link. Restoration Convergence Time is measured using the Rate-Derived Convergence Time calculation technique, as provided in Equation 1. It is possible to have the Restoration Convergence Time differ from the Rate-Derived Convergence Time.

Measurement Units:

seconds or milliseconds

Issues:

None

See Also:

Convergence Event

Rate-Derived Convergence Time

3.14 Packet Sampling Interval**Definition:**

The interval at which the tester (test equipment) polls to make measurements for arriving packet flows.

Discussion:

Metrics measured at the Packet Sampling Interval may include Throughput and Convergence Packet Loss.

Measurement Units:

seconds or milliseconds

Issues:

Packet Sampling Interval can influence the Convergence Graph. This is particularly true when implementations achieve Full Convergence in less than 1 second. The Convergence Event Transition and Convergence Recovery Transition can become exaggerated when the Packet Sampling Interval is too long. This will produce a larger than actual Rate-Derived Convergence Time. The recommended value for configuration of the Packet Sampling Interval is provided in [\[2\]](#).

See Also:

Convergence Packet Loss

Convergence Event Transition

Convergence Recovery Transition

3.15 Local Interface

Definition:

An interface on the DUT.

Discussion:

None

Measurement Units:

N/A

Issues:

None

See Also:

Neighbor Interface

Remote Interface

3.16 Neighbor Interface

Definition:

The interface on the neighbor router or tester that is directly linked to the DUT's Local Interface.

Discussion:

None

Measurement Units:

N/A

Issues:

None

See Also:

Local Interface

Remote Interface

3.17 Remote Interface

Definition:

An interface on a neighboring router that is not directly connected to any interface on the DUT.

Discussion:

None

Measurement Units:

N/A

Issues:
None

Poretsky, Imhoff

[Page 11]

See Also:

Local Interface

Neighbor Interface

3.18 Preferred Egress Interface

Definition:

The outbound interface from the DUT for traffic routed to the preferred next-hop.

Discussion:

The Preferred Egress Interface is the egress interface prior to a Convergence Event.

Measurement Units:

N/A

Issues:

None

See Also:

Next-Best Egress Interface

3.19 Next-Best Egress Interface

Definition:

The outbound interface from the DUT for traffic routed to the second-best next-hop. It is the same media type and link speed as the Preferred Egress Interface

Discussion:

The Next-Best Egress Interface becomes the egress interface after a Convergence Event.

Measurement Units:

N/A

Issues:

None

See Also:

Preferred Egress Interface

3.20 Stale Forwarding

Definition:

Forwarding of traffic to route entries that no longer exist or to route entries with next-hops that are no longer preferred.

Discussion:

Stale Forwarding can be caused by a Convergence Event and is also known as a "black-hole" since it may produce packet loss. Stale Forwarding exists until Network Convergence is achieved.

Measurement Units:

N/A

Issues:

None

See Also:

Network Convergence

3.21 Nested Convergence Events

Definition:

The occurrence of Convergence Event while the route table is converging from a prior Convergence Event.

Discussion:

The Convergence Events for a Nested Convergence Events MUST occur with different neighbors. A common observation from a Nested Convergence Event will be the withdrawal of routes from one neighbor while the routes of another neighbor are being installed.

Measurement Units:

N/A

Issues:

None

See Also:

Convergence Event

[4. IANA Considerations](#)

This document requires no IANA considerations.

5. Security Considerations

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

6. References

6.1 Normative References

- [1] Poretsky, S., "Benchmarking Applicability for IGP Data Plane Route Convergence", [draft-ietf-bmwg-igp-dataplane-conv-app-09](#), work in progress, January 2006.
- [2] Poretsky, S., "Benchmarking Methodology for IGP Data Plane Route Convergence", [draft-ietf-bmwg-igp-dataplane-conv-meth-09](#), work in progress, January 2006.
- [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.

6.2 Informative References

- [5] S. Casner, C. Alaettinoglu, and C. Kuan, "A Fine-Grained View of High Performance Networking", NANOG 22, June 2001.
- [6] L. Ciavattone, A. Morton, and G. Ramachandran, "Standardized Active Measurements on a Tier 1 IP Backbone", IEEE Communications Magazine, pp90-97, May 2003.

7. Author's Address

Scott Poretsky
Reef Point Systems
8 New England Executive Park
Burlington, MA 01803
USA

Phone: + 1 508 439 9008
EMail: sporetsky@reefpoint.com

Brent Imhoff
Juniper Networks
1194 North Mathilda Ave
Sunnyvale, CA 94089
USA

Phone: + 1 314 378 2571
EMail: bimhoff@planetspork.com

Full Copyright Statement

Copyright (C) The Internet Society (2006).

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.

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.