Network Working Group	S. Poretsky
Internet-Draft	Allot Communications
Intended status: Informational	B. Imhoff
Expires: November 11, 2010	Juniper Networks
	K. Michielsen
	Cisco Systems
	May 10, 2010

Terminology for Benchmarking Link-State IGP Data Plane Route Convergence
draft-ietf-bmwg-igp-dataplane-conv-term-21

Abstract

This document describes the terminology for benchmarking Interior Gateway Protocol (IGP) Route Convergence. The terminology is to be used for benchmarking IGP convergence time through externally observable (black box) data plane measurements. The terminology can be applied to any link-state IGP, such as ISIS and OSPF.

Status of this Memo

Copyright Notice

Table of Contents

---

1. Introduction and Scope

TOC

This draft describes the terminology for benchmarking Link-State Interior Gateway Protocol (IGP) Convergence. The motivation and applicability for this benchmarking is provided in [Po09a] (Poretsky, S., “Considerations for Benchmarking Link-State IGP Data Plane Route Convergence,” March 2009.). The methodology to be used for this benchmarking is described in [Po10m] (Poretsky, S., Imhoff, B., and K. Michielsen, “Benchmarking Methodology for Link-State IGP Data Plane Route Convergence,” March 2010.). The purpose of this document is to introduce new terms required to complete execution of the IGP Route Methodology [Po10m] (Poretsky, S., Imhoff, B., and K. Michielsen, “Benchmarking Methodology for Link-State IGP Data Plane Route Convergence,” March 2010.).
IGP convergence time is measured on the data plane at the Tester by observing packet loss through the DUT. The methodology and terminology to be used for benchmarking IGP Convergence can be applied to IPv4 and IPv6 traffic and link-state IGPs such as ISIS [Ca90] (Callon, R., “Use of OSI IS-IS for routing in TCP/IP and dual environments,” December 1990.)[Ho08] (Hopps, C., “Routing IPv6 with IS-IS,” October 2008.), OSPF [Mo98] (Moy, J., “OSPF Version 2,” April 1998.) [Co08] (Coltun, R., Ferguson, D., Moy, J., and A. Lindem, “OSPF for IPv6,” July 2008.), and others.

TOC

2. Existing Definitions

This document uses existing terminology defined in other BMWG work. Examples include, but are not limited to:

Frame Loss Rate	[Ref.[Br91] (Bradner, S., "Benchmarking terminology for network interconnection devices," July 1991.), section 3.6]
Throughput	[Ref.[Br91] (Bradner, S., "Benchmarking terminology for network interconnection devices," July 1991.), section 3.17]
Offered Load	[Ref.[Ma98] (Mandeville, R., "Benchmarking Terminology for LAN Switching Devices," February 1998.), section 3.5.2]
Forwarding Rate	[Ref.[Ma98] (Mandeville, R., "Benchmarking Terminology for LAN Switching Devices," February 1998.), section 3.6.1]
Device Under Test (DUT)	[Ref.[Ma98] (Mandeville, R., "Benchmarking Terminology for LAN Switching Devices," February 1998.), section 3.1.1]
System Under Test (SUT)	[Ref.[Ma98] (Mandeville, R., "Benchmarking Terminology for LAN Switching Devices," February 1998.), section 3.1.2]
Out-of-Order Packet	[Ref.[Po06] (Poretsky, S., Perser, J., Erramilli, S., and S. Khurana, "Terminology for Benchmarking Network-layer Traffic Control Mechanisms," October 2006.), section 3.3.4]
Duplicate Packet	[Ref.[Po06] (Poretsky, S., Perser, J., Erramilli, S., and S. Khurana, "Terminology for Benchmarking Network-layer Traffic Control Mechanisms," October 2006.), section 3.3.5]
Packet Reordering	[Ref.[Mo06] (Morton, A., Ciavattone, L., Ramachandran, G., Shalunov, S., and J. Perser, "Packet Reordering Metrics," November 2006.), section 3.3]
Stream	[Ref.[Po06] (Poretsky, S., Perser, J., Erramilli, S., and S. Khurana, "Terminology for Benchmarking Network-layer Traffic Control Mechanisms," October 2006.), section 3.3.2]
Forwarding Delay	[Ref.[Po06] (Poretsky, S., Perser, J., Erramilli, S., and S. Khurana, "Terminology for Benchmarking Network-layer Traffic Control Mechanisms," October 2006.), section 3.2.4]
IP Packet Delay Variation (IPDV)	[Ref.[De02] (Demichelis, C. and P. Chimento, "IP Packet Delay Variation Metric for IP Performance Metrics (IPPM)," November 2002.), section 1.2]

Loss Period	[Ref.[Ko02] (Koodli, R. and R. Ravikanth, “One-way Loss Pattern Sample Metrics,” August 2002.), section 4]

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] (Bradner, S., “Key words for use in RFCs to Indicate Requirement Levels,” March 1997.). 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. Term Definitions

TOC

3.1. Convergence Types

TOC

3.1.1. Route Convergence

TOC

Definition:
The process of updating all components of the router, including the Routing Information Base (RIB) and Forwarding Information Base (FIB), along with software and hardware tables, with the most recent route change(s) such that forwarding for a route entry is successful on the Next-Best Egress Interface.
Discussion:
Route Convergence MUST occur after a Convergence Event. Route Convergence can be observed externally by the rerouting of data traffic for a destination matching a route entry to the Next-best Egress Interface. Completion of Route Convergence may or may not be sustained over time.
Measurement Units: N/A
Issues: None
See Also:
Network Convergence, Full Convergence, Convergence Event

TOC

3.1.2. Full Convergence

Definition:
Route Convergence for all routes in the FIB.
Discussion:
Full Convergence MUST occur after a Convergence Event. Full Convergence can be observed externally by the rerouting of data traffic to destinations matching all route entries to the Next-best Egress Interface. Completion of Full Convergence is externally observable from the data plane when the Forwarding Rate of the data plane traffic on the Next-Best Egress Interface equals the Offered Load.
Completion of Full Convergence may or may not be sustained over time.
Measurement Units: N/A
Issues: None
See Also:
Network Convergence, Route Convergence, Convergence Event, Full Convergence Time, Convergence Recovery Instant

---

TOC

3.1.3. Network Convergence

Definition:
Full Convergence in all routers throughout the network.
Discussion:
Network Convergence includes all Route Convergence operations for all routers in the network following a Convergence Event.
Completion of Network Convergence can be observed by recovery of the network Forwarding Rate to equal the Offered Load, with no Stale Forwarding, and no Blenders [Ca01] (Casner, S., Alaettinoglu, C., and C. Kuan, “A Fine-Grained View of High Performance Networking,” June 2001.)[Ci03] (Ciavattone, L., Morton, A., and G. Ramachandran, “Standardized Active Measurements on a Tier 1 IP Backbone,” May 2003.).
Completion of Network Convergence may or may not be sustained over time.
Measurement Units: N/A
Issues: None
See Also:
Route Convergence, Full Convergence, Stale Forwarding

TOC

3.2. Instants

TOC

3.2.1. Traffic Start Instant

Definition:
The time instant the Tester sends out the first data packet to the DUT.
Discussion:
If using the Loss-Derived Method or the Route-Specific Loss-Derived Method to benchmark IGP convergence time, and the applied Convergence Event does not cause instantaneous traffic loss for all routes at the Convergence Event Instant then the Tester SHOULD collect a timestamp on the Traffic Start Instant in order to measure the period of time between the Traffic Start Instant and Convergence Event Instant.
Measurement Units:
hh:mm:ss:nnn:uuu, where 'nnn' is milliseconds and 'uuu' is microseconds.
Issues: None
See Also:
Convergence Event Instant, Route-Specific Convergence Time, Loss-Derived Convergence Time.

3.2.2. Convergence Event Instant

TOC

Definition:
The time instant that a Convergence Event occurs.
Discussion:
If the Convergence Event causes instantaneous traffic loss on the Preferred Egress Interface, the Convergence Event Instant is observable from the data plane as the instant that the DUT begins to exhibit packet loss.
The Tester SHOULD collect a timestamp on the Convergence Event Instant if it is not observable from the data plane.
Measurement Units:
hh:mm:ss:nnn:uuu, where 'nnn' is milliseconds and 'uuu' is microseconds.
Issues: None
See Also: Convergence Event

3.2.3. Convergence Recovery Instant

TOC

Definition:
The time instant that Full Convergence has completed.
Discussion:
The Full Convergence completed state MUST be maintained for an interval of duration equal to the Sustained Convergence Validation Time in order to validate the Convergence Recovery Instant.

The Convergence Recovery Instant is observable from the data plane as the instant the DUT forwards traffic to all destinations over the Next-Best Egress Interface.
Measurement Units:
hh:mm:ss:nnn:uuu, where 'nnn' is milliseconds and 'uuu' is microseconds.
Issues: None
See Also:
Sustained Convergence Validation Time, Full Convergence

3.2.4. First Route Convergence Instant

TOC

Definition:
The time instant the first route entry completes Route Convergence following a Convergence Event
Discussion:
Any route may be the first to complete Route Convergence. The First Route Convergence Instant is observable from the data plane as the instant that the first packet is received from the Next-Best Egress Interface.
Measurement Units:
hh:mm:ss:nnn:uuu, where 'nnn' is milliseconds and 'uuu' is microseconds.
Issues: None
See Also: Route Convergence

3.3. Transitions

TOC

3.3.1. Convergence Event Transition

TOC

Definition:
A time interval following a Convergence Event in which Forwarding Rate on the Preferred Egress Interface gradually reduces to zero.
Discussion:
The Forwarding Rate during a Convergence Event Transition may not decrease linearly.
The Forwarding Rate observed on all DUT egress interfaces may or may not decrease to zero.
The Offered Load, the number of routes, and the Packet Sampling Interval influence the observations of the Convergence Event Transition

using the Rate-Derived Method. This is further discussed with the term "Rate-Derived Method".
Measurement Units: seconds
Issues: None
See Also:
Convergence Event, Rate-Derived Method

3.3.2. Convergence Recovery Transition

TOC

Definition:
A time interval following the First Route Convergence Instant in which Forwarding Rate on the Next-Best Egress Interface gradually increases to equal the Offered Load.
Discussion:
The Forwarding Rate observed during a Convergence Recovery Transition may not increase linearly.
The Offered Load, the number of routes, and the Packet Sampling Interval influence the observations of the Convergence Recovery Transition using the Rate-Derived Method. This is further discussed with the term "Rate-Derived Method".
Measurement Units: seconds
Issues: None
See Also:
Full Convergence,First Route Convergence Instant, Rate-Derived Method

---

3.4. Interfaces

TOC

3.4.1. Local Interface

TOC

Definition:
An interface on the DUT.
Discussion:
A failure of the Local Interface indicates that the failure occurred directly on the DUT.
Measurement Units: N/A
Issues: None
See Also: Remote Interface

3.4.2. Remote Interface

TOC

Definition:
An interface on a neighboring router that is not directly connected to any interface on the DUT.
Discussion:
A failure of a Remote Interface indicates that the failure occurred on a neighbor router's interface that is not directly connected to the DUT.
Measurement Units: N/A
Issues: None
See Also: Local Interface

3.4.3. Preferred Egress Interface

TOC

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

3.4.4. Next-Best Egress Interface

TOC

Definition:
The outbound interface from the DUT for traffic routed to the second-best next-hop.
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

3.5. Benchmarking Methods

TOC

3.5.1. Rate-Derived Method

TOC

Definition:
The method to calculate convergence time benchmarks from observing Forwarding Rate each Packet Sampling Interval.
Discussion:
Figure 1 (Rate-Derived Convergence Graph) shows an example of the Forwarding Rate change in time during convergence as observed when using the Rate-Derived Method.

```
         ^         Traffic                      Convergence
    Fwd  |         Start                        Recovery
    Rate |         Instant                      Instant
         | Offered  ^                             ^
         | Load --> ----------\                   /-----------
         |                     \                 /<--- Convergence
         |                      \     Packet    /      Recovery
         |       Convergence --->\     Loss    /       Transition
         |       Event            \           /
         |       Transition        \---------/ <-- Max Packet Loss
         |
         +---------------------------------------------------------->
                         ^                   ^                 time
                    Convergence         First Route
                    Event Instant       Convergence Instant
```

Figure 1: Rate-Derived Convergence Graph

The Offered Load SHOULD consist of a single Stream [Po06] (Poretsky, S., Perser, J., Erramilli, S., and S. Khurana, “Terminology for Benchmarking Network-layer Traffic Control Mechanisms,” October 2006.). If sending multiple Streams, the measured traffic rate statistics for all Streams MUST be added together.
The destination addresses for the Offered Load MUST be distributed such that all routes or a statistically representative subset of all routes are matched and each of these routes is offered an equal share of the Offered Load. It is RECOMMENDED to send traffic to all routes, but a statistically representative subset of all routes can be used if required.
At least one packet per route for all routes matched in the Offered Load MUST be offered to the DUT within each Packet Sampling Interval. For maximum accuracy the value for the Packet Sampling Interval SHOULD be as small as possible, but the presence of IP Packet Delay Variation (IPDV) [De02] (Demichelis, C. and P. Chimento, “IP Packet Delay

Variation Metric for IP Performance Metrics (IPPM)," November 2002.)
may enforce using a larger Packet Sampling Interval.

The Offered Load, IPDV, the number of routes, and the Packet Sampling Interval influence the observations for the Rate-Derived Method. It may be difficult to identify the different convergence time instants in the Rate-Derived Convergence Graph. For example, it is possible that a Convergence Event causes the Forwarding Rate to drop to zero, while this may not be observed in the Forwarding Rate measurements if the Packet Sampling Interval is too large.

IPDV causes fluctuations in the number of received packets during each Packet Sampling Interval. To account for the presence of IPDV in determining if a convergence instant has been reached, Forwarding Delay SHOULD be observed during each Packet Sampling Interval. The minimum and maximum number of packets expected in a Packet Sampling Interval in presence of IPDV can be calculated with Equation 1.

```
 number of packets expected in a Packet Sampling Interval
   in presence of IP Packet Delay Variation
     = expected number of packets without IP Packet Delay Variation
       +/-( (maxDelay - minDelay) * Offered Load)
 with minDelay and maxDelay the minimum resp. maximum Forwarding Delay
   of packets received during the Packet Sampling Interval

                              Equation 1
```

To determine if a convergence instant has been reached the number of packets received in a Packet Sampling Interval is compared with the range of expected number of packets calculated in Equation 1.

If packets are going over multiple ECMP members and one or more of the members has failed then the number of received packets during each Packet Sampling Interval may vary, even excluding presence of IPDV. To prevent fluctuation of the number of received packets during each Packet Sampling Interval for this reason, the Packet Sampling Interval duration SHOULD be a whole multiple of the time between two consecutive packets sent to the same destination.

Metrics measured at the Packet Sampling Interval MUST include Forwarding Rate and packet loss.

Rate-Derived Method is a RECOMMENDED method to measure convergence time benchmarks.

To measure convergence time benchmarks for Convergence Events that do not cause instantaneous traffic loss for all routes at the Convergence Event Instant, the Tester SHOULD collect a timestamp of the Convergence Event Instant and the Tester SHOULD observe Forwarding Rate separately on the Next-Best Egress Interface.

Since the Rate-Derived Method does not distinguish between individual traffic destinations, it SHOULD NOT be used for any route specific measurements. Therefor Rate-Derived Method SHOULD NOT be used to benchmark Route Loss of Connectivity Period.

Measurement Units: N/A

Issues: None

See Also:
Packet Sampling Interval, Convergence Event, Convergence Event Instant, Full Convergence

---

3.5.2. Loss-Derived Method

TOC

Definition:
The method to calculate the Loss-Derived Convergence Time and Loss-Derived Loss of Connectivity Period benchmarks from the amount of packet loss.
Discussion:
The Offered Load SHOULD consist of a single Stream [Po06] (Poretsky, S., Perser, J., Erramilli, S., and S. Khurana, “Terminology for Benchmarking Network-layer Traffic Control Mechanisms,” October 2006.). If sending multiple Streams, the measured traffic rate statistics for all Streams MUST be added together.
The destination addresses for the Offered Load MUST be distributed such that all routes or a statistically representative subset of all routes are matched and each of these routes is offered an equal share of the Offered Load. It is RECOMMENDED to send traffic to all routes, but a statistically representative subset of all routes can be used if required.
Loss-Derived Method SHOULD always be combined with Rate-Derived Method in order to observe Full Convergence completion. The total amount of Convergence Packet Loss is collected after Full Convergence completion.
To measure convergence time and loss of connectivity benchmarks, the Tester SHOULD in general observe packet loss on all DUT egress interfaces (Connectivity Packet Loss).
To measure convergence time benchmarks for Convergence Events that do not cause instantaneous traffic loss for all routes at the Convergence Event Instant, the Tester SHOULD collect timestamps of the Start Traffic Instant and of the Convergence Event Instant, and the Tester SHOULD observe packet loss separately on the Next-Best Egress Interface (Convergence Packet Loss).
Since Loss-Derived Method does not distinguish between traffic destinations and the packet loss statistics are only collected after Full Convergence completion, this method can only be used to measure average values over all routes. For these reasons Loss-Derived Method can only be used to benchmark Loss-Derived Convergence Time and Loss-Derived Loss of Connectivity Period.
Note that the Loss-Derived Method measures an average over all routes, including the routes that may not be impacted by the Convergence Event, such as routes via non-impacted members of ECMP or parallel links.
Measurement Units: seconds
Issues: None
See Also:

Loss-Derived Convergence Time, Loss-Derived Loss of Connectivity Period, Convergence Packet Loss

---

3.5.3. Route-Specific Loss-Derived Method

TOC

Definition:
The method to calculate the Route-Specific Convergence Time benchmark from the amount of packet loss during convergence for a specific route entry.
Discussion:
To benchmark Route-Specific Convergence Time, the Tester provides an Offered Load that consists of multiple Streams [Po06] (Poretsky, S., Perser, J., Erramilli, S., and S. Khurana, "Terminology for Benchmarking Network-layer Traffic Control Mechanisms," October 2006.). Each Stream has a single destination address matching a different route entry, for all routes or a statistically representative subset of all routes. Convergence Packet Loss is measured for each Stream separately. Route-Specific Loss-Derived Method SHOULD always be combined with Rate-Derived Method in order to observe Full Convergence completion. The total amount of Convergence Packet Loss for each Stream is collected after Full Convergence completion.
Route-Specific Loss-Derived Method is a RECOMMENDED method to measure convergence time benchmarks.
To measure convergence time and loss of connectivity benchmarks, the Tester SHOULD in general observe packet loss on all DUT egress interfaces (Connectivity Packet Loss).
To measure convergence time benchmarks for Convergence Events that do not cause instantaneous traffic loss for all routes at the Convergence Event Instant, the Tester SHOULD collect timestamps of the Start Traffic Instant and of the Convergence Event Instant, and the Tester SHOULD observe packet loss separately on the Next-Best Egress Interface (Convergence Packet Loss).
Since Route-Specific Loss-Derived Method uses traffic streams to individual routes, it measures packet loss as it would be experienced by a network user. For this reason Route-Specific Loss-Derived Method is RECOMMENDED to measure Route-Specific Convergence Time benchmarks and Route Loss of Connectivity Period benchmarks.
Measurement Units: seconds
Issues: None
See Also:
Route-Specific Convergence Time, Route Loss of Connectivity Period, Convergence Packet Loss

---

TOC

3.6. Benchmarks

3.6.1. Full Convergence Time

TOC

Definition:
The time duration of the period between the Convergence Event Instant and the Convergence Recovery Instant as observed using the Rate-Derived Method.
Discussion:
Using the Rate-Derived Method, Full Convergence Time can be calculated as the time difference between the Convergence Event Instant and the Convergence Recovery Instant, as shown in Equation 2.

```
Full Convergence Time =
    Convergence Recovery Instant - Convergence Event Instant
```

Equation 2

The Convergence Event Instant can be derived from the Forwarding Rate observation or from a timestamp collected by the Tester.
For the testcases described in [Po10m] (Poretsky, S., Imhoff, B., and K. Michielsen, “Benchmarking Methodology for Link-State IGP Data Plane Route Convergence,” March 2010.), it is expected that Full Convergence Time equals the maximum Route-Specific Convergence Time when benchmarking all routes in FIB using the Route-Specific Loss-Derived Method.
It is not possible to measure Full Convergence Time using the Loss-Derived Method.
Measurement Units: seconds
Issues: None
See Also:
Full Convergence, Rate-Derived Method, Route-Specific Loss-Derived Method

---

3.6.2. First Route Convergence Time

TOC

Definition:
The duration of the period between the Convergence Event Instant and the First Route Convergence Instant as observed using the Rate-Derived Method.
Discussion:
Using the Rate-Derived Method, First Route Convergence Time can be calculated as the time difference between the Convergence Event Instant and the First Route Convergence Instant, as shown with Equation 3.

```
First Route Convergence Time =
    First Route Convergence Instant - Convergence Event Instant
```

Equation 3

The Convergence Event Instant can be derived from the Forwarding Rate observation or from a timestamp collected by the Tester.

For the testcases described in [Po10m] (Poretsky, S., Imhoff, B., and K. Michielsen, “Benchmarking Methodology for Link-State IGP Data Plane Route Convergence,” March 2010.), it is expected that First Route Convergence Time equals the minimum Route-Specific Convergence Time when benchmarking all routes in FIB using the Route-Specific Loss-Derived Method.

It is not possible to measure First Route Convergence Time using the Loss-Derived Method.

Measurement Units: seconds

Issues: None

See Also:

Rate-Derived Method, Route-Specific Loss-Derived Method, First Route Convergence Instant

---

3.6.3. Route-Specific Convergence Time

TOC

Definition:

The amount of time it takes for Route Convergence to be completed for a specific route, as calculated from the amount of packet loss during convergence for a single route entry.

Discussion:

Route-Specific Convergence Time can only be measured using the Route-Specific Loss-Derived Method.

If the applied Convergence Event causes instantaneous traffic loss for all routes at the Convergence Event Instant, Connectivity Packet Loss should be observed. Connectivity Packet Loss is the combined packet loss observed on Preferred Egress Interface and Next-Best Egress Interface. When benchmarking Route-Specific Convergence Time, Connectivity Packet Loss is measured and Equation 4 is applied for each measured route. The calculation is equal to Equation 8 in Section 3.6.5 (Route Loss of Connectivity Period).

```
Route-Specific Convergence Time =
   Connectivity Packet Loss for specific route/Offered Load per route
```

Equation 4

If the applied Convergence Event does not cause instantaneous traffic loss for all routes at the Convergence Event Instant, then the Tester SHOULD collect timestamps of the Traffic Start Instant and of the Convergence Event Instant, and the Tester SHOULD observe Convergence

Packet Loss separately on the Next-Best Egress Interface. When benchmarking Route-Specific Convergence Time, Convergence Packet Loss is measured and Equation 5 is applied for each measured route.

```
 Route-Specific Convergence Time =
     Convergence Packet Loss for specific route/Offered Load per route
     - (Convergence Event Instant - Traffic Start Instant)
```

Equation 5

The Convergence Event Instant and Traffic Start Instant SHOULD be collected by the Tester.

The Route-Specific Convergence Time benchmarks enable minimum, maximum, average, and median convergence time measurements to be reported by comparing the results for the different route entries. It also enables benchmarking of convergence time when configuring a priority value for route entry(ies). Since multiple Route-Specific Convergence Times can be measured it is possible to have an array of results. The format for reporting Route-Specific Convergence Time is provided in [Po10m] (Poretsky, S., Imhoff, B., and K. Michielsen, “Benchmarking Methodology for Link-State IGP Data Plane Route Convergence,” March 2010.).

Measurement Units: seconds

Issues: None

See Also:

Convergence Event, Convergence Packet Loss, Connectivity Packet Loss, Route Convergence

---

3.6.4. Loss-Derived Convergence Time

TOC

Definition:

The average Route Convergence time for all routes in FIB, as calculated from the amount of packet loss during convergence.

Discussion:

Loss-Derived Convergence Time is measured using the Loss-Derived Method.

If the applied Convergence Event causes instantaneous traffic loss for all routes at the Convergence Event Instant, Connectivity Packet Loss should be observed. Connectivity Packet Loss is the combined packet loss observed on Preferred Egress Interface and Next-Best Egress Interface. When benchmarking Loss-Derived Convergence Time, Connectivity Packet Loss is measured and Equation 6 is applied.

```
               Loss-Derived Convergence Time =
                   Connectivity Packet Loss/Offered Load
```

Equation 6

If the applied Convergence Event does not cause instantaneous traffic loss for all routes at the Convergence Event Instant, then the Tester

SHOULD collect timestamps of the Start Traffic Instant and of the Convergence Event Instant and the Tester SHOULD observe Convergence Packet Loss separately on the Next-Best Egress Interface. When benchmarking Loss-Derived Convergence Time, Convergence Packet Loss is measured and Equation 7 is applied.

```
Loss-Derived Convergence Time =
    Convergence Packet Loss/Offered Load
    - (Convergence Event Instant - Traffic Start Instant)
```

Equation 7

The Convergence Event Instant and Traffic Start Instant SHOULD be collected by the Tester.
Measurement Units: seconds
Issues: None
See Also:
Convergence Packet Loss, Connectivity Packet Loss, Route Convergence

---

3.6.5. Route Loss of Connectivity Period

TOC

Definition:
The time duration of traffic loss for a specific route entry following a Convergence Event until Full Convergence completion, as observed using the Route-Specific Loss-Derived Method.
Discussion:
In general the Route Loss of Connectivity Period is not equal to the Route-Specific Convergence Time. If the DUT continues to forward traffic to the Preferred Egress Interface after the Convergence Event is applied then the Route Loss of Connectivity Period will be smaller than the Route-Specific Convergence Time. This is also specifically the case after reversing a failure event.
The Route Loss of Connectivity Period may be equal to the Route-Specific Convergence Time if, as a characteristic of the Convergence Event, traffic for all routes starts dropping instantaneously on the Convergence Event Instant. See discussion in [Po10m] (Poretsky, S., Imhoff, B., and K. Michielsen, “Benchmarking Methodology for Link-State IGP Data Plane Route Convergence,” March 2010.).
For the testcases described in [Po10m] (Poretsky, S., Imhoff, B., and K. Michielsen, “Benchmarking Methodology for Link-State IGP Data Plane Route Convergence,” March 2010.) the Route Loss of Connectivity Period is expected to be a single Loss Period [Ko02] (Koodli, R. and R. Ravikanth, “One-way Loss Pattern Sample Metrics,” August 2002.).
When benchmarking Route Loss of Connectivity Period, Connectivity Packet Loss is measured for each route and Equation 8 is applied for each measured route entry. The calculation is equal to Equation 4 in Section 3.6.3 (Route-Specific Convergence Time).

Route Loss of Connectivity Period =
   Connectivity Packet Loss for specific route/Offered Load per route

Equation 8

Route Loss of Connectivity Period SHOULD be measured using Route-Specific Loss-Derived Method.

Measurement Units: seconds

Issues: None

See Also:

Route-Specific Convergence Time, Route-Specific Loss-Derived Method, Connectivity Packet Loss

---

3.6.6. Loss-Derived Loss of Connectivity Period

TOC

Definition:

The average time duration of traffic loss for all routes following a Convergence Event until Full Convergence completion, as observed using the Loss-Derived Method.

Discussion:

In general the Loss-Derived Loss of Connectivity Period is not equal to the Loss-Derived Convergence Time. If the DUT continues to forward traffic to the Preferred Egress Interface after the Convergence Event is applied then the Loss-Derived Loss of Connectivity Period will be smaller than the Loss-Derived Convergence Time. This is also specifically the case after reversing a failure event.

The Loss-Derived Loss of Connectivity Period may be equal to the Loss-Derived Convergence Time if, as a characteristic of the Convergence Event, traffic for all routes starts dropping instantaneously on the Convergence Event Instant. See discussion in [Po10m] (Poretsky, S., Imhoff, B., and K. Michielsen, “Benchmarking Methodology for Link-State IGP Data Plane Route Convergence,” March 2010.).

For the testcases described in [Po10m] (Poretsky, S., Imhoff, B., and K. Michielsen, “Benchmarking Methodology for Link-State IGP Data Plane Route Convergence,” March 2010.) each route's Route Loss of Connectivity Period is expected to be a single Loss Period [Ko02] (Koodli, R. and R. Ravikanth, “One-way Loss Pattern Sample Metrics,” August 2002.).

When benchmarking Loss-Derived Loss of Connectivity Period, Connectivity Packet Loss is measured for all routes and Equation 9 is applied. The calculation is equal to Equation 6 in Section 3.6.4 (Loss-Derived Convergence Time).

Loss-Derived Loss of Connectivity Period =
   Connectivity Packet Loss for all routes/Offered Load

Equation 9

Loss-Derived Loss of Connectivity Period SHOULD be measured using Loss-Derived Method.
Measurement Units: seconds
Issues: None
See Also:
Loss-Derived Convergence Time, Loss-Derived Method, Connectivity Packet Loss

---

3.7. Measurement Terms

TOC

3.7.1. Convergence Event

TOC

Definition:
The occurrence of a planned or unplanned event in the network that will result in a change in the egress interface of the Device Under Test (DUT) for routed packets.
Discussion:
Convergence Events include but are not limited to 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 Event Instant

3.7.2. Packet Loss

TOC

Definition:
The number of packets that should have been forwarded by a DUT under a constant Offered Load that were not forwarded due to lack of resources.
Discussion:
Packet Loss is a modified version of the term "Frame Loss Rate" as defined in [Br91] (Bradner, S., “Benchmarking terminology for network interconnection devices,” July 1991.). The term "Frame Loss" is intended for Ethernet Frames while "Packet Loss" is intended for IP packets.
Measurement units: Number of offered packets that are not forwarded.
Issues: None
See Also: Convergence Packet Loss

3.7.3. Convergence Packet Loss

TOC

Definition:
The number of packets lost due to a Convergence Event until Full Convergence completes, as observed on the Next-Best Egress Interface.
Discussion:
Convergence Packet Loss is observed on the Next-Best Egress Interface. It only needs to be observed for Convergence Events that do not cause instantaneous traffic loss at Convergence Event Instant.
Convergence Packet Loss includes packets that were lost and packets that were delayed due to buffering. The maximum acceptable Forwarding Delay (Forwarding Delay Threshold) is a parameter of the methodology, if it is applied it MUST be reported.
Measurement Units: number of packets
Issues: None
See Also:
Packet Loss, Full Convergence, Convergence Event, Connectivity Packet Loss

---

3.7.4. Connectivity Packet Loss

TOC

Definition:
The number of packets lost due to a Convergence Event until Full Convergence completes.
Discussion:
Connectivity Packet Loss is observed on all DUT egress interfaces.
Connectivity Packet Loss includes packets that were lost and packets that were delayed due to buffering. The maximum acceptable Forwarding Delay (Forwarding Delay Threshold) is a parameter of the methodology, if it is applied it MUST be reported.
Measurement Units: number of packets
Issues: None
See Also:
Packet Loss, Route Loss of Connectivity Period, Convergence Event, Convergence Packet Loss

---

3.7.5. Packet Sampling Interval

TOC

Definition:
The interval at which the Tester (test equipment) polls to make measurements for arriving packets.
Discussion:

At least one packet per route for all routes matched in the Offered Load MUST be offered to the DUT within the Packet Sampling Interval. Metrics measured at the Packet Sampling Interval MUST include Forwarding Rate and received packets.

Packet Sampling Interval can influence the convergence graph as observed with the Rate-Derived Method. This is particularly true when implementations complete Full Convergence in less time than the Packet Sampling Interval. The Convergence Event Instant and First Route Convergence Instant may not be easily identifiable and the Rate-Derived Method may produce a larger than actual convergence time.

Using a small Packet Sampling Interval in the presence of IPDV [De02] (Demichelis, C. and P. Chimento, "IP Packet Delay Variation Metric for IP Performance Metrics (IPPM)," November 2002.) may cause fluctuations of the Forwarding Rate observation and can prevent correct observation of the different convergence time instants.

The value of the Packet Sampling Interval only contributes to the measurement accuracy of the Rate-Derived Method. For maximum accuracy the value for the Packet Sampling Interval SHOULD be as small as possible, but the presence of IPDV may enforce using a larger Packet Sampling Interval.

Measurement Units: seconds

Issues: None

See Also: Rate-Derived Method

3.7.6. Sustained Convergence Validation Time

TOC

Definition:

The amount of time for which the completion of Full Convergence is maintained without additional packet loss.

Discussion:

The purpose of the Sustained Convergence Validation Time is to produce convergence benchmarks protected against fluctuation in Forwarding Rate after the completion of Full Convergence is observed. The RECOMMENDED Sustained Convergence Validation Time to be used is the time to send 5 consecutive packets to each destination with a minimum of 5 seconds. The BMWG selected 5 seconds based upon [Br99] (Bradner, S. and J. McQuaid, "Benchmarking Methodology for Network Interconnect Devices," March 1999.) which recommends waiting 2 seconds for residual frames to arrive (this is the Forwarding Delay Threshold for the last packet sent) and 5 seconds for DUT restabilization.

Measurement Units: seconds

Issues: None

See Also:

Full Convergence, Convergence Recovery Instant

3.7.7. Forwarding Delay Threshold

TOC

Definition:
The maximum waiting time threshold used to distinguish between packets with very long delay and lost packets that will never arrive.
Discussion:
Applying a Forwarding Delay Threshold allows to consider packets with a too large Forwarding Delay as being lost, as is required for some applications (e.g. voice, video, etc.). The Forwarding Delay Threshold is a parameter of the methodology, if it is applied it MUST be reported.
Measurement Units: seconds
Issues: None
See Also:
Convergence Packet Loss, Connectivity Packet Loss

3.8. Miscellaneous Terms

TOC

3.8.1. Stale Forwarding

TOC

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 can manifest as a "black-hole" or microloop that produces packet loss, or out-of-order packets, or delayed packets. Stale Forwarding can exist until Network Convergence is completed.
Measurement Units: N/A
Issues: None
See Also: Network Convergence

3.8.2. Nested Convergence Event

TOC

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

The Convergence Events for a Nested Convergence Event MUST occur with different neighbors. A possible 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. Security Considerations

TOC

Benchmarking activities as described in this memo are limited to technology characterization using controlled stimuli in a laboratory environment, with dedicated address space and the constraints specified in the sections above.
The benchmarking network topology will be an independent test setup and MUST NOT be connected to devices that may forward the test traffic into a production network, or misroute traffic to the test management network.
Further, benchmarking is performed on a "black-box" basis, relying solely on measurements observable external to the DUT/SUT.
Special capabilities SHOULD NOT exist in the DUT/SUT specifically for benchmarking purposes. Any implications for network security arising from the DUT/SUT SHOULD be identical in the lab and in production networks.

---

5. IANA Considerations

TOC

This document requires no IANA considerations.

---

6. Acknowledgements

TOC

Thanks to Sue Hares, Al Morton, Kevin Dubray, Ron Bonica, David Ward, Peter De Vriendt, Anuj Dewagan and the BMWG for their contributions to this work.

---

7. References

TOC

7.1. Normative References

TOC

[Br91]	Bradner, S., "Benchmarking terminology for network interconnection devices," RFC 1242, July 1991 (TXT).
[Br97]	Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels," BCP 14, RFC 2119, March 1997 (TXT, HTML, XML).
[Br99]	Bradner, S. and J. McQuaid, "Benchmarking Methodology for Network Interconnect Devices," RFC 2544, March 1999 (TXT).
[Ca90]	Callon, R., "Use of OSI IS-IS for routing in TCP/IP and dual environments," RFC 1195, December 1990 (TXT, PS).
[Co08]	Coltun, R., Ferguson, D., Moy, J., and A. Lindem, "OSPF for IPv6," RFC 5340, July 2008 (TXT).
[De02]	Demichelis, C. and P. Chimento, "IP Packet Delay Variation Metric for IP Performance Metrics (IPPM)," RFC 3393, November 2002 (TXT).
[Ho08]	Hopps, C., "Routing IPv6 with IS-IS," RFC 5308, October 2008 (TXT).
[Ko02]	Koodli, R. and R. Ravikanth, "One-way Loss Pattern Sample Metrics," RFC 3357, August 2002 (TXT).
[Ma98]	Mandeville, R., "Benchmarking Terminology for LAN Switching Devices," RFC 2285, February 1998 (TXT, HTML, XML).
[Mo06]	Morton, A., Ciavattone, L., Ramachandran, G., Shalunov, S., and J. Perser, "Packet Reordering Metrics," RFC 4737, November 2006 (TXT).
[Mo98]	Moy, J., "OSPF Version 2," STD 54, RFC 2328, April 1998 (TXT, HTML, XML).
[Po06]	Poretsky, S., Perser, J., Erramilli, S., and S. Khurana, "Terminology for Benchmarking Network-layer Traffic Control Mechanisms," RFC 4689, October 2006 (TXT).
[Po09a]	Poretsky, S., "Considerations for Benchmarking Link-State IGP Data Plane Route Convergence," draft-ietf-bmwg-igp-dataplane-conv-app-17 (work in progress), March 2009 (TXT).
[Po10m]	Poretsky, S., Imhoff, B., and K. Michielsen, "Benchmarking Methodology for Link-State IGP Data Plane Route Convergence," draft-ietf-bmwg-igp-dataplane-conv-meth-20 (work in progress), March 2010 (TXT).

7.2. Informative References

TOC

[Ca01]	Casner, S., Alaettinoglu, C., and C. Kuan, "A Fine-Grained View of High Performance Networking," NANOG 22, June 2001.
[Ci03]	Ciavattone, L., Morton, A., and G. Ramachandran, "Standardized Active Measurements on a Tier 1 IP Backbone," IEEE Communications Magazine p90-97, May 2003.

Authors' Addresses

TOC

	Scott Poretsky
	Allot Communications
	67 South Bedford Street, Suite 400
	Burlington, MA 01803
	USA
Phone:	+ 1 508 309 2179
Email:	sporetsky@allot.com
	Brent Imhoff
	Juniper Networks
	1194 North Mathilda Ave
	Sunnyvale, CA 94089
	USA
Phone:	+ 1 314 378 2571
Email:	bimhoff@planetspork.com
	Kris Michielsen
	Cisco Systems
	6A De Kleetlaan
	Diegem, BRABANT 1831
	Belgium
Email:	kmichiel@cisco.com