

Internet Engineering Task Force (IETF)
Request for Comments: 6412
Category: Informational
ISSN: 2070-1721

S. Poretsky
Allot Communications
B. Imhoff
F5 Networks
K. Michielsen
Cisco Systems
November 2011

Terminology for Benchmarking Link-State IGP Data-Plane Route Convergence

Abstract

This document describes the terminology for benchmarking link-state 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 IS-IS and OSPF.

Status of This Memo

This document is not an Internet Standards Track specification; it is published for informational purposes.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Not all documents approved by the IESG are a candidate for any level of Internet Standard; see Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc6412>.

Copyright Notice

Copyright (c) 2011 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must

include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

This document may contain material from IETF Documents or IETF Contributions published or made publicly available before November 10, 2008. The person(s) controlling the copyright in some of this material may not have granted the IETF Trust the right to allow modifications of such material outside the IETF Standards Process. Without obtaining an adequate license from the person(s) controlling the copyright in such materials, this document may not be modified outside the IETF Standards Process, and derivative works of it may not be created outside the IETF Standards Process, except to format it for publication as an RFC or to translate it into languages other than English.

Table of Contents

1.	Introduction and Scope	4
2.	Existing Definitions	4
3.	Term Definitions	5
3.1.	Convergence Types	5
3.1.1.	Route Convergence	5
3.1.2.	Full Convergence	5
3.2.	Instants	6
3.2.1.	Traffic Start Instant	6
3.2.2.	Convergence Event Instant	6
3.2.3.	Convergence Recovery Instant	7
3.2.4.	First Route Convergence Instant	8
3.3.	Transitions	8
3.3.1.	Convergence Event Transition	8
3.3.2.	Convergence Recovery Transition	9
3.4.	Interfaces	10
3.4.1.	Local Interface	10
3.4.2.	Remote Interface	10
3.4.3.	Preferred Egress Interface	10
3.4.4.	Next-Best Egress Interface	11
3.5.	Benchmarking Methods	11
3.5.1.	Rate-Derived Method	11
3.5.2.	Loss-Derived Method	14
3.5.3.	Route-Specific Loss-Derived Method	15
3.6.	Benchmarks	17
3.6.1.	Full Convergence Time	17
3.6.2.	First Route Convergence Time	18
3.6.3.	Route-Specific Convergence Time	18
3.6.4.	Loss-Derived Convergence Time	20
3.6.5.	Route Loss of Connectivity Period	21
3.6.6.	Loss-Derived Loss of Connectivity Period	22
3.7.	Measurement Terms	23
3.7.1.	Convergence Event	23
3.7.2.	Convergence Packet Loss	23
3.7.3.	Connectivity Packet Loss	24
3.7.4.	Packet Sampling Interval	24
3.7.5.	Sustained Convergence Validation Time	25
3.7.6.	Forwarding Delay Threshold	26
3.8.	Miscellaneous Terms	26
3.8.1.	Impaired Packet	26
4.	Security Considerations	27
5.	Acknowledgements	27
6.	Normative References	27

1. Introduction and Scope

This document is a companion to [Pollm], which contains the methodology to be used for benchmarking link-state Interior Gateway Protocol (IGP) convergence by observing the data plane. The purpose of this document is to introduce new terms required to complete execution of the Link-State IGP Data-Plane Route Convergence methodology [Pollm].

IGP convergence time is measured by observing the data plane through the Device Under Test (DUT) at the Tester. 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 Intermediate System to Intermediate System (IS-IS) [Ca90][Ho08], Open Shortest Path First (OSPF) [Mo98] [Co08], and others.

2. Existing Definitions

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

Throughput	[Br91], Section 3.17
Offered Load	[Ma98], Section 3.5.2
Forwarding Rate	[Ma98], Section 3.6.1
Device Under Test (DUT)	[Ma98], Section 3.1.1
System Under Test (SUT)	[Ma98], Section 3.1.2
Out-of-Order Packet	[Po06], Section 3.3.4
Duplicate Packet	[Po06], Section 3.3.5
Stream	[Po06], Section 3.3.2
Forwarding Delay	[Po06], Section 3.2.4
IP Packet Delay Variation (IPDV)	[De02], Section 1.2
Loss Period	[Ko02], Section 4

The keywords "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 keywords 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

3.1. Convergence Types

3.1.1. Route Convergence

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 (Section 3.4.4).

Discussion:

In general, IGP convergence does not necessarily result in a change in forwarding. But the test cases in [Pollm] are specified such that the IGP convergence results in a change of egress interface for the measurement data-plane traffic. Due to this property of the test case specifications, Route Convergence can be observed externally by the rerouting of the measurement data-plane traffic to the Next-Best Egress Interface (Section 3.4.4).

Measurement Units:

N/A

See Also:

Next-Best Egress Interface, Full Convergence

3.1.2. Full Convergence

Definition:

Route Convergence for all routes in the Forwarding Information Base (FIB).

Discussion:

In general, IGP convergence does not necessarily result in a change in forwarding. But the test cases in [Pollm] are specified such that the IGP convergence results in a change of egress interface for the measurement data-plane traffic. Due to this property of the test cases specifications, Full Convergence can be observed externally by the rerouting of the measurement data-plane traffic to the Next-Best Egress Interface (Section 3.4.4).

Measurement Units:

N/A

See Also:

Next-Best Egress Interface, Route Convergence

3.2. Instants

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 (Section 3.5.2) or the Route-Specific Loss-Derived Method (Section 3.5.3) to benchmark IGP convergence time, and the applied Convergence Event (Section 3.7.1) does not cause instantaneous traffic loss for all routes at the Convergence Event Instant (Section 3.2.2), 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:

seconds (and fractions), reported with resolution sufficient to distinguish between different instants

See Also:

Loss-Derived Method, Route-Specific Loss-Derived Method, Convergence Event, Convergence Event Instant

3.2.2. Convergence Event Instant

Definition:

The time instant that a Convergence Event (Section 3.7.1) occurs.

Discussion:

If the Convergence Event (Section 3.7.1) causes instantaneous traffic loss on the Preferred Egress Interface (Section 3.4.3), the Convergence Event Instant is observable from the data plane as the instant that no more packets are received on the Preferred Egress Interface.

The Tester SHOULD collect a timestamp on the Convergence Event Instant if the Convergence Event does not cause instantaneous traffic loss on the Preferred Egress Interface (Section 3.4.3).

Measurement Units:

seconds (and fractions), reported with resolution sufficient to distinguish between different instants

See Also:

Convergence Event, Preferred Egress Interface

3.2.3. Convergence Recovery Instant**Definition:**

The time instant that Full Convergence (Section 3.1.2) has completed.

Discussion:

The Full Convergence completed state MUST be maintained for an interval of duration equal to the Sustained Convergence Validation Time (Section 3.7.5) 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 (Section 3.4.4) without impairments.

Measurement Units:

seconds (and fractions), reported with resolution sufficient to distinguish between different instants

See Also:

Sustained Convergence Validation Time, Full Convergence, Next-Best Egress Interface

3.2.4. First Route Convergence Instant

Definition:

The time instant the first route entry completes Route Convergence (Section 3.1.1)

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 that is not an Impaired Packet (Section 3.8.1) is received from the Next-Best Egress Interface (Section 3.4.4) or, for the test cases with Equal Cost Multi-Path (ECMP) or Parallel Links, the instant that the Forwarding Rate on the Next-Best Egress Interface (Section 3.4.4) starts to increase.

Measurement Units:

seconds (and fractions), reported with resolution sufficient to distinguish between different instants

See Also:

Route Convergence, Impaired Packet, Next-Best Egress Interface

3.3. Transitions

3.3.1. Convergence Event Transition

Definition:

A time interval following a Convergence Event (Section 3.7.1) in which the Forwarding Rate on the Preferred Egress Interface (Section 3.4.3) gradually reduces to zero.

Discussion:

The Forwarding Rate during a Convergence Event Transition may or may not decrease linearly.

The Forwarding Rate observed on the DUT egress interface(s) may or may not decrease to zero.

The Offered Load, the number of routes, and the Packet Sampling Interval (Section 3.7.4) influence the observations of the Convergence Event Transition using the Rate-Derived Method (Section 3.5.1).

Measurement Units:

seconds (and fractions)

See Also:

Convergence Event, Preferred Egress Interface, Packet Sampling Interval, Rate-Derived Method

3.3.2. Convergence Recovery Transition

Definition:

A time interval following the First Route Convergence Instant (Section 3.4.4) in which the Forwarding Rate on the DUT egress interface(s) gradually increases to equal to the Offered Load.

Discussion:

The Forwarding Rate observed during a Convergence Recovery Transition may or may not increase linearly.

The Offered Load, the number of routes, and the Packet Sampling Interval (Section 3.7.4) influence the observations of the Convergence Recovery Transition using the Rate-Derived Method (Section 3.5.1).

Measurement Units:

seconds (and fractions)

See Also:

First Route Convergence Instant, Packet Sampling Interval, Rate-Derived Method

3.4. Interfaces

3.4.1. Local Interface

Definition:

An interface on the DUT.

Discussion:

A failure of a Local Interface indicates that the failure occurred directly on the DUT.

Measurement Units:

N/A

See Also:

Remote Interface

3.4.2. Remote Interface

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

See Also:

Local Interface

3.4.3. 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 (Section 3.7.1).

Measurement Units:

N/A

See Also:

Convergence Event, Next-Best Egress Interface

3.4.4. Next-Best Egress Interface

Definition:

The outbound interface or set of outbound interfaces in an Equal Cost Multipath (ECMP) set or parallel link set of the Device Under Test (DUT) for traffic routed to the second-best next-hop.

Discussion:

The Next-Best Egress Interface becomes the egress interface after a Convergence Event (Section 3.4.4).

For the test cases in [Pollm] using test topologies with an ECMP set or parallel link set, the term Preferred Egress Interface refers to all members of the link set.

Measurement Units:

N/A

See Also:

Convergence Event, Preferred Egress Interface

3.5. Benchmarking Methods

3.5.1. Rate-Derived Method

Definition:

The method to calculate convergence time benchmarks from observing the Forwarding Rate each Packet Sampling Interval (Section 3.7.4).

Discussion:

Figure 1 shows an example of the Forwarding Rate change in time during convergence as observed when using the Rate-Derived Method.




Figure 1: Rate-Derived Convergence Graph

To enable collecting statistics of Out-of-Order Packets per flow (see [Th00], Section 3), the Offered Load SHOULD consist of multiple Streams [Po06], and each Stream SHOULD consist of a single flow. If sending multiple Streams, the measured traffic 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 of the Packet Sampling Interval SHOULD be as small as possible, but the presence of IP Packet Delay Variation (IPDV) [De02] may require that a larger Packet Sampling Interval be used.

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)

where minDelay and maxDelay indicate (respectively) the minimum and 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 Impaired Packet count.

To measure convergence time benchmarks for Convergence Events (Section 3.7.1) 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 (Section 3.2.2), 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. Therefore, the Rate-Derived Method SHOULD NOT be used to benchmark Route Loss of Connectivity Period (Section 3.6.5).

Measurement Units:

N/A

See Also:

Packet Sampling Interval, Convergence Event, Convergence Event Instant, Next-Best Egress Interface, Route Loss of Connectivity Period

3.5.2. Loss-Derived Method

Definition:

The method to calculate the Loss-Derived Convergence Time (Section 3.6.4) and Loss-Derived Loss of Connectivity Period (Section 3.6.6) benchmarks from the amount of Impaired Packets (Section 3.8.1).

Discussion:

To enable collecting statistics of Out-of-Order Packets per flow (see [Th00], Section 3), the Offered Load SHOULD consist of multiple Streams [Po06], and each Stream SHOULD consist of a single flow. If sending multiple Streams, the measured traffic 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 the 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 for Convergence Events that cause instantaneous traffic loss for all routes at the Convergence Event Instant, the Tester SHOULD observe the Impaired Packet count on all DUT egress interfaces (see Connectivity Packet Loss (Section 3.7.3)).

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 Impaired Packet count separately on the Next-Best Egress Interface (see Convergence Packet Loss (Section 3.7.2)).

Since Loss-Derived Method does not distinguish between traffic destinations and the Impaired Packet 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 (Section 3.6.4) and Loss-Derived Loss of Connectivity Period (Section 3.6.6).

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:

N/A

See Also:

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

3.5.3. Route-Specific Loss-Derived Method

Definition:

The method to calculate the Route-Specific Convergence Time (Section 3.6.3) benchmark from the amount of Impaired Packets (Section 3.8.1) 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]. Each Stream has a single destination address matching a different route entry, for all routes or a statistically representative subset of all routes. Each Stream SHOULD consist of a single flow (see [Th00], Section 3). Convergence Packet Loss is measured for each Stream separately.

Route-Specific Loss-Derived Method SHOULD always be combined with the Rate-Derived Method in order to observe Full Convergence completion. The total amount of Convergence Packet Loss (Section 3.7.2) for each Stream is collected after Full Convergence completion.

Route-Specific Loss-Derived Method is the RECOMMENDED method to measure convergence time benchmarks.

To measure convergence time and loss of connectivity benchmarks for Convergence Events that cause instantaneous traffic loss for all routes at the Convergence Event Instant, the Tester SHOULD observe Impaired Packet count on all DUT egress interfaces (see Connectivity Packet Loss (Section 3.7.3)).

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 (see Convergence Packet Loss (Section 3.7.2)).

Since Route-Specific Loss-Derived Method uses traffic streams to individual routes, it observes Impaired Packet count 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:

N/A

See Also:

Route-Specific Convergence Time, Route Loss of Connectivity Period, Connectivity Packet Loss, Convergence Packet Loss

3.6. Benchmarks

3.6.1. Full Convergence Time

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.

$$\text{Full Convergence Time} = \text{Convergence Recovery Instant} - \text{Convergence Event Instant}$$

Equation 2

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

For the test cases described in [Pollm], it is expected that Full Convergence Time equals the maximum Route-Specific Convergence Time when benchmarking all routes in the 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 (and fractions)

See Also:

Full Convergence, Rate-Derived Method, Route-Specific Loss-Derived Method, Convergence Event Instant, Convergence Recovery Instant

3.6.2. First Route Convergence Time

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 test cases described in [Pollm], it is expected that First Route Convergence Time equals the minimum Route-Specific Convergence Time when benchmarking all routes in the 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 (and fractions)

See Also:

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

3.6.3. Route-Specific Convergence Time

Definition:

The amount of time it takes for Route Convergence to be completed for a specific route, as calculated from the amount of Impaired Packets (Section 3.8.1) 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 Impaired Packet count 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-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 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 the route entry or entries. 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 [Pollm].

Measurement Units:

seconds (and fractions)

See Also:

Route-Specific Loss-Derived Method, Convergence Event, Convergence Event Instant, Convergence Packet Loss, Connectivity Packet Loss, Route Convergence

3.6.4. Loss-Derived Convergence Time

Definition:

The average Route Convergence time for all routes in the Forwarding Information Base (FIB), as calculated from the amount of Impaired Packets (Section 3.8.1) 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 (Section 3.7.3) should be observed. Connectivity Packet Loss is the combined Impaired Packet count 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.

$$\text{Loss-Derived Convergence Time} = \frac{\text{Connectivity Packet Loss}}{\text{Offered Load}}$$

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 (Section 3.7.2) separately on the Next-Best Egress Interface. When benchmarking Loss-Derived Convergence Time, Convergence Packet Loss is measured and Equation 7 is applied.

$$\text{Loss-Derived Convergence Time} = \frac{\text{Convergence Packet Loss}}{\text{Offered Load}} - (\text{Convergence Event Instant} - \text{Traffic Start Instant})$$

Equation 7

Measurement Units:

seconds (and fractions)

See Also:

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

3.6.5. Route Loss of Connectivity Period**Definition:**

The time duration of packet impairments 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 [Pollm].

For the test cases described in [Pollm], the Route Loss of Connectivity Period is expected to be a single Loss Period [Ko02].

When benchmarking the 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 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 (and fractions)

See Also:

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

3.6.6. Loss-Derived Loss of Connectivity Period**Definition:**

The average time duration of packet impairments 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 [Pollm].

For the test cases described in [Pollm], each route's Route Loss of Connectivity Period is expected to be a single Loss Period [Ko02].

When benchmarking the 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.

$$\text{Loss-Derived Loss of Connectivity Period} = \frac{\text{Connectivity Packet Loss for all routes}}{\text{Offered Load}}$$

Equation 9

The Loss-Derived Loss of Connectivity Period SHOULD be measured using the Loss-Derived Method.

Measurement Units:

seconds (and fractions)

See Also:

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

3.7. Measurement Terms

3.7.1. Convergence Event

Definition:

The occurrence of an event in the network that will result in a change in the egress interface of the DUT for routed packets.

Discussion:

All test cases in [Pollm] are defined such that a Convergence Event results in a change of egress interface of the DUT. Local or remote triggers that cause a route calculation that does not result in a change in forwarding are not considered.

Measurement Units:

N/A

See Also:

Convergence Event Instant

3.7.2. Convergence Packet Loss

Definition:

The number of Impaired Packets (Section 3.8.1) as observed on the Next-Best Egress Interface of the DUT during convergence.

Discussion:

An Impaired Packet is considered as a lost packet.

Measurement Units:

number of packets

See Also:

Connectivity Packet Loss

3.7.3. Connectivity Packet Loss**Definition:**

The number of Impaired Packets observed on all DUT egress interfaces during convergence.

Discussion:

An Impaired Packet is considered as a lost packet. Connectivity Packet Loss is equal to Convergence Packet Loss if the Convergence Event causes instantaneous traffic loss for all egress interfaces of the DUT except for the Next-Best Egress Interface.

Measurement Units:

number of packets

See Also:

Convergence Packet Loss

3.7.4. Packet Sampling Interval**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] 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 (and fractions)

See Also:

Rate-Derived Method

3.7.5. Sustained Convergence Validation Time

Definition:

The amount of time for which the completion of Full Convergence is maintained without additional Impaired Packets being observed.

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 Benchmarking Methodology Working Group (BMWG) selected 5 seconds based upon [Br99], 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 (and fractions)

See Also:

Full Convergence, Convergence Recovery Instant

3.7.6. Forwarding Delay Threshold

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 packets with a too large Forwarding Delay to be considered lost, as is required for some applications (e.g. voice, video, etc.). The Forwarding Delay Threshold is a parameter of the methodology, and it MUST be reported. [Br99] recommends waiting 2 seconds for residual frames to arrive.

Measurement Units:

seconds (and fractions)

See Also:

Convergence Packet Loss, Connectivity Packet Loss

3.8. Miscellaneous Terms

3.8.1. Impaired Packet

Definition:

A packet that experienced at least one of the following impairments: loss, excessive Forwarding Delay, corruption, duplication, reordering.

Discussion:

A lost packet, a packet with a Forwarding Delay exceeding the Forwarding Delay Threshold, a corrupted packet, a Duplicate Packet [Po06], and an Out-of-Order Packet [Po06] are Impaired Packets.

Packet ordering is observed for each individual flow (see [Th00], Section 3) of the Offered Load.

Measurement Units:

N/A

See Also:

Forwarding Delay Threshold

4. Security Considerations

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. Acknowledgements

Thanks to Sue Hares, Al Morton, Kevin Dubray, Ron Bonica, David Ward, Peter De Vriendt, Anuj Dewagan, Adrian Farrel, Stewart Bryant, Francis Dupont, and the Benchmarking Methodology Working Group for their contributions to this work.

6. Normative References

- [Br91] Bradner, S., "Benchmarking terminology for network interconnection devices", RFC 1242, July 1991.
- [Br97] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [Br99] Bradner, S. and J. McQuaid, "Benchmarking Methodology for Network Interconnect Devices", RFC 2544, March 1999.
- [Ca90] Callon, R., "Use of OSI IS-IS for routing in TCP/IP and dual environments", RFC 1195, December 1990.

- [Co08] Coltun, R., Ferguson, D., Moy, J., and A. Lindem, "OSPF for IPv6", RFC 5340, July 2008.
- [De02] Demichelis, C. and P. Chimento, "IP Packet Delay Variation Metric for IP Performance Metrics (IPPM)", RFC 3393, November 2002.
- [Ho08] Hopps, C., "Routing IPv6 with IS-IS", RFC 5308, October 2008.
- [Ko02] Koodli, R. and R. Ravikanth, "One-way Loss Pattern Sample Metrics", RFC 3357, August 2002.
- [Ma98] Mandeville, R., "Benchmarking Terminology for LAN Switching Devices", RFC 2285, February 1998.
- [Mo98] Moy, J., "OSPF Version 2", STD 54, RFC 2328, April 1998.
- [Po06] Poretsky, S., Perser, J., Erramilli, S., and S. Khurana, "Terminology for Benchmarking Network-layer Traffic Control Mechanisms", RFC 4689, October 2006.
- [Pollm] Poretsky, S., Imhoff, B., and K. Michielsen, "Benchmarking Methodology for Link-State IGP Data-Plane Route Convergence", RFC 6413, November 2011.
- [Th00] Thaler, D. and C. Hopps, "Multipath Issues in Unicast and Multicast Next-Hop Selection", RFC 2991, November 2000.

Authors' Addresses

Scott Poretsky
Allot Communications
300 TradeCenter
Woburn, MA 01801
USA

Phone: + 1 508 309 2179
EMail: sporetsky@allot.com

Brent Imhoff
F5 Networks
401 Elliott Avenue West
Seattle, WA 98119
USA

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

Kris Michielsen
Cisco Systems
6A De Kleetlaan
Diegem, BRABANT 1831
Belgium

EMail: kmichiel@cisco.com

