Network Working Group
Internet Draft
Expires: April 2009
Intended Status: Informational

S. Poretsky
Allot Communications

Brent Imhoff
Juniper Networks

October 15, 2008

**Benchmarking Methodology for
Link-State IGP Data Plane Route Convergence**

<draft-ietf-bmwg-igp-dataplane-conv-meth-16.txt>

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

Table of Contents

1. Introduction

This document describes the methodology for benchmarking Interior Gateway Protocol (IGP) Route Convergence. The applicability of this testing is described in [Po07a] and the new terminology that it introduces is defined in [Po07t]. Service Providers use IGP Convergence time as a key metric of router design and architecture. Customers of Service Providers observe convergence time by packet loss, so IGP Route Convergence is considered a Direct Measure of Quality (DMOQ). The test cases in this document are black-box tests that emulate the network events that cause route convergence, as described in [Po07a]. The black-box test designs benchmark the data plane and account for all of the factors contributing to convergence time, as discussed in [Po07a]. The methodology (and terminology) for benchmarking route convergence can be applied to any link-state IGP such as ISIS [Ca90] and OSPF [Mo98] and others. These methodologies apply to IPv4 and IPv6 traffic and IGPs.

2. Existing definitions

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14, RFC 2119 [Br97]. RFC 2119 defines the use of these key words to help make the intent of standards track documents as clear as possible. While this document uses these keywords, this document is not a standards track document.

This document adopts the definition format in Section 2 of RFC 1242 [Br91]. This document uses much of the terminology defined in [Po07t]. This document uses existing terminology defined in other BMWG work. Examples include, but are not limited to:

Throughput	[Ref.[Br91], section 3.17]
Device Under Test (DUT)	[Ref.[Ma98], section 3.1.1]
System Under Test (SUT)	[Ref.[Ma98], section 3.1.2]
Out-of-order Packet	[Ref.[Po06], section 3.3.2]
Duplicate Packet	[Ref.[Po06], section 3.3.3]
Packet Loss	[Ref.[Po07t], Section 3.5]

3. Test Setup

3.1 Test Topologies

Figure 1 shows the test topology to measure IGP Route Convergence due to local Convergence Events such as Link Failure, Layer 2 Session Failure, IGP Adjacency Failure, Route Withdrawal, and route cost change. These test cases discussed in section 4 provide route convergence times that include the Event Detection time, SPF Processing time, and FIB Update time. These times are measured by observing packet loss in the data plane at the Tester.

Figure 2 shows the test topology to measure IGP Route Convergence time due to remote changes in the network topology. These times are measured by observing packet loss in the data plane at the Tester. In this topology the three routers are considered a System Under Test (SUT). A Remote Interface [Po07t] failure on router R2 MUST result in convergence of traffic to router R3. NOTE: All routers in the SUT must be the same model and identically configured.

Figure 3 shows the test topology to measure IGP Route Convergence time with members of an Equal Cost Multipath (ECMP) Set. These times are measured by observing packet loss in the data plane at the Tester. In this topology, the DUT is configured with each Egress interface as a member of an ECMP set and the Tester emulates multiple next-hop routers (emulates one router for each member).

```
     ---------       Ingress Interface         ---------
     |       |<--------------------------------|       |
     |       |                                 |       |
     |       |    Preferred Egress Interface   |       |
     |  DUT  |-------------------------------->| Tester|
     |       |                                 |       |
     |       |~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~>|       |
     |       |    Next-Best Egress Interface   |       |
```

```
  ---------                                 ---------
```

Figure 1. Test Topology 1: IGP Convergence Test Topology for Local Changes

```
                -----                       ---------
                |   | Preferred             |       |
        -----   |R2 |---------------------->|       |
        |   |-->|   | Egress Interface      |       |
        |   |   -----                       |       |
        |R1 |                               |Tester |
        |   |   -----                       |       |
        |   |-->|   |   Next-Best           |       |
        -----   |R3 |~~~~~~~~~~~~~~~~~~~~~~>|       |
          ^     |   |   Egress Interface    |       |
          |     -----                       ---------
          |                                     |
          |-------------------------------------
                      Ingress Interface
```

      Figure 2. Test Topology 2: IGP Convergence Test Topology
                for Convergence Due to Remote Changes

```
        ---------       Ingress Interface         ---------
        |       |<--------------------------------|       |
        |       |                                 |       |
        |       |     ECMP Set Interface 1        |       |
        |  DUT  |-------------------------------->| Tester|
        |       |                .                |       |
        |       |                .                |       |
        |       |                .                |       |
        |       |~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~>|       |
        |       |     ECMP Set Interface N        |       |
        ---------                                 ---------
```

      Figure 3. Test Topology 3: IGP Convergence Test Topology
                for ECMP Convergence

```
        ---------       Ingress Interface         ---------
        |       |<--------------------------------|       |
        |       |                                 |       |
        |       |     Parallel Link Interface 1   |       |
        |  DUT  |-------------------------------->| Tester|
        |       |                .                |       |
        |       |                .                |       |
        |       |                .                |       |
        |       |~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~>|       |
        |       |     Parallel Link Interface N   |       |
        ---------                                 ---------
```

      Figure 4. Test Topology 4: IGP Convergence Test Topology
                for Parallel Link Convergence

Figure 4 shows the test topology to measure IGP Route Convergence time with members of a Parallel Link. These times are measured by observing packet loss in the data plane at the Tester. In this topology, the DUT is configured with each Egress interface as a member of a Parallel Link and the Tester emulates the single next-hop router.

3.2 Test Considerations

3.2.1 IGP Selection

The test cases described in section 4 MAY be used for link-state IGPs, such as ISIS or OSPF. The Route Convergence test methodology is identical. The IGP adjacencies are established on the Preferred Egress Interface and Next-Best Egress Interface.

3.2.2 Routing Protocol Configuration

The obtained results for IGP Route Convergence may vary if other routing protocols are enabled and routes learned via those protocols are installed. IGP convergence times MUST be benchmarked without routes installed from other protocols.

3.2.3 IGP Route Scaling

The number of IGP routes will impact the measured IGP Route Convergence. To obtain results similar to those that would be observed in an operational network, it is RECOMMENDED that the number of installed routes and nodes closely approximates that of the network (e.g. thousands of routes with tens of nodes). The number of areas (for OSPF) and levels (for ISIS) can impact the benchmark results.

3.2.4 Timers

There are some timers that will impact the measured IGP Convergence time. Benchmarking metrics may be measured at any fixed values for these timers. It is RECOMMENDED that the following timers be configured to the minimum values listed:

Timer	Recommended Value
Link Failure Indication Delay	<10milliseconds
IGP Hello Timer	1 second
IGP Dead-Interval	3 seconds
LSA Generation Delay	0
LSA Flood Packet Pacing	0
LSA Retransmission Packet Pacing	0
SPF Delay	0

3.2.5 Interface Types

All test cases in this methodology document may be executed with any interface type. All interfaces MUST be the same media and Throughput [Br91][Br99] for each test case. The type of media may dictate which test cases may be executed. This is because each interface type has a unique mechanism for detecting link failures and the speed at which that mechanism operates will influence the measure results. Media and protocols MUST be configured for minimum failure detection delay to minimize the contribution to the measured Convergence time. For example, configure SONET with the minimum carrier-loss-delay. All interfaces SHOULD be configured as point-to-point.

3.2.6 Packet Sampling Interval

The Packet Sampling Interval [Po07t] value is the fastest measurable Rate-Derived Convergence Time [Po07t]. The RECOMMENDED value for the Packet Sampling Interval is 10 milliseconds. Rate-Derived Convergence Time is the preferred benchmark for IGP Route Convergence. This benchmark must always be reported when the Packet Sampling Interval is set <= 10 milliseconds on the test equipment. If the test equipment does not permit the Packet Sampling Interval to be set as low as 10 milliseconds, then both the Rate-Derived Convergence Time and Loss-Derived Convergence Time [Po07t] MUST be reported.

3.2.7 Offered Load

The offered load MUST be the Throughput of the device as defined in [Br91] and benchmarked in [Br99] at a fixed packet size. At least one packet per route in the FIB for all routes in the FIB MUST be offered to the DUT within the Packet Sampling interval. Packet size is measured in bytes and includes the IP header and payload. The packet size is selectable and MUST be recorded. The Forwarding Rate [Ma98] MUST be measured at the Preferred Egress Interface and the Next-Best Egress Interface. The duration of offered load MUST be greater than the convergence time. The destination addresses for the offered load MUST be distributed such that all routes are matched and each route is offered an equal share of the total Offered Load. This requirement for the Offered Load to be distributed to match all destinations in the route table creates separate flows that are offered to the DUT. The capability of the Tester to measure packet loss for each individual flow (identified by the destination address matching a route entry) and the scale for the number of individual flows for which it can measure packet loss should be considered when benchmarking Route-Specific Convergence [Po07t].

3.2.8 Selection of Convergence Time Benchmark Metrics

The methodologies in the section 4 test cases MAY be applied to benchmark Full Convergence and Route-Specific Convergence with benchmarking metrics First Route Convergence Time, Loss-Derived Convergence Time, Rate-Derived Convergence Time, Reversion Convergence Time, and Route-Specific Convergence Times [Po07t].

When benchmarking Full Convergence the Rate-Derived Convergence Time benchmarking metric MAY be measured. When benchmarking Route-Specific Convergence the Route-Specific Convergence Time benchmarking metric MUST be measured and Full Convergence MAY be obtained from max(Route-Specific Convergence Time). The First Route Convergence Time benchmarking metric MAY be measured when benchmarking either Full Convergence or Route-Specific Convergence.

3.2.9 Tester Capabilities

It is RECOMMENDED that the Tester used to execute each test case have the following capabilities:

1. Ability to insert a timestamp in each data packet's IP payload.
2. An internal time clock to control timestamping, time measurements, and time calculations.
3. Ability to distinguish traffic load received on the Preferred and Next-Best interfaces.
4. Ability to disable or tune specific Layer-2 and Layer-3 protocol functions on any interface(s).

3.3 Reporting Format

For each test case, it is recommended that the reporting table below is completed and all time values SHOULD be reported with resolution as specified in [Po07t].

Parameter	Units
Test Case	test case number
Test Topology	(1, 2, 3, or 4)
IGP	(ISIS, OSPF, other)
Interface Type	(GigE, POS, ATM, other)
Packet Size offered to DUT	bytes
IGP Routes advertised to DUT	number of IGP routes
Nodes in emulated network	number of nodes
Packet Sampling Interval on Tester	milliseconds
IGP Timer Values configured on DUT:	
Interface Failure Indication Delay	seconds
IGP Hello Timer	seconds
IGP Dead-Interval	seconds
LSA Generation Delay	seconds
LSA Flood Packet Pacing	seconds
LSA Retransmission Packet Pacing	seconds
SPF Delay	seconds
Forwarding Metrics	
Total Packets Offered to DUT	number of Packets
Total Packets Routed by DUT	number of Packets
Convergence Packet Loss	number of Packets
Out-of-Order Packets	number of Packets
Duplicate Packets	number of Packets
Convergence Benchmarks	
Full Convergence	
First Route Convergence Time	seconds
Rate-Derived Convergence Time	seconds
Loss-Derived Convergence Time	seconds
Route-Specific Convergence	
Number of Routes Measured	number of flows
Route-Specific Convergence Time[n]	array of seconds
Minimum R-S Convergence Time	seconds
Maximum R-S Convergence Time	seconds
Median R-S Convergence Time	seconds
Average R-S Convergence Time	seconds
Reversion	
Reversion Convergence Time	seconds
First Route Convergence Time	seconds
Route-Specific Convergence	
Number of Routes Measured	number of flows
Route-Specific Convergence Time[n]	array of seconds
Minimum R-S Convergence Time	seconds

Maximum R-S Convergence Time	seconds
Median R-S Convergence Time	seconds
Average R-S Convergence Time	seconds

4. Test Cases

It is RECOMMENDED that all applicable test cases be executed for best characterization of the DUT. The test cases follow a generic procedure tailored to the specific DUT configuration and Convergence Event[Po07t]. This generic procedure is as follows:

1. Establish DUT configuration and install routes.
2. Send offered load with traffic traversing Preferred Egress Interface [Po07t].
3. Introduce Convergence Event to force traffic to Next-Best Egress Interface [Po07t].
4. Measure First Route Convergence Time.
5. Measure Full Convergence from Loss-Derived Convergence Time, Rate-Derived Convergence Time, and optionally the Route-Specific Convergence Times.
6. Wait the Sustained Convergence Validation Time to ensure there no residual packet loss.
7. Recover from Convergence Event.
8. Measure Reversion Convergence Time, and optionally the First Route Convergence Time and Route-Specific Convergence Times.

4.1 Convergence Due to Local Interface Failure

Objective

To obtain the IGP Route Convergence due to a local link failure event at the DUT's Local Interface.

Procedure

1. Advertise matching IGP routes from Tester to DUT on Preferred Egress Interface [Po07t] and Next-Best Egress Interface [Po07t] using the topology shown in Figure 1. Set the cost of the routes so that the Preferred Egress Interface is the preferred next-hop.
2. Send offered load at measured Throughput with fixed packet size to destinations matching all IGP routes from Tester to DUT on Ingress Interface [Po07t].
3. Verify traffic is routed over Preferred Egress Interface.
4. Remove link on DUT's Preferred Egress Interface. This is the Convergence Event [Po07t] that produces the Convergence Event Instant [Po07t].
5. Measure First Route Convergence Time [Po07t] as DUT detects the link down event and begins to converge IGP routes and traffic over the Next-Best Egress Interface.
6. Measure Rate-Derived Convergence Time [Po07t] as DUT detects the link down event and converges all IGP routes and traffic over the Next-Best Egress Interface. Optionally, Route-Specific Convergence Times [Po07t] MAY be measured.
7. Stop offered load. Wait 30 seconds for queues to drain. Restart offered load.
8. Restore link on DUT's Preferred Egress Interface.

9. Measure Reversion Convergence Time [Po07t], and optionally measure First Route Convergence Time [Po07t] and Route-Specific Convergence Times [Po07t], as DUT detects the link up event and converges all IGP routes and traffic back to the Preferred Egress Interface.

Results
The measured IGP Convergence time is influenced by the Local link failure indication, SPF delay, SPF Hold time, SPF Execution Time, Tree Build Time, and Hardware Update Time [Po07a].

4.2 Convergence Due to Remote Interface Failure

Objective
To obtain the IGP Route Convergence due to a Remote Interface Failure event.

Procedure
1. Advertise matching IGP routes from Tester to SUT on Preferred Egress Interface [Po07t] and Next-Best Egress Interface [Po07t] using the topology shown in Figure 2. Set the cost of the routes so that the Preferred Egress Interface is the preferred next-hop.
2. Send offered load at measured Throughput with fixed packet size to destinations matching all IGP routes from Tester to SUT on Ingress Interface [Po07t].
3. Verify traffic is routed over Preferred Egress Interface.
4. Remove link on Tester's Neighbor Interface [Po07t] connected to SUT's Preferred Egress Interface. This is the Convergence Event [Po07t] that produces the Convergence Event Instant [Po07t].
5. Measure First Route Convergence Time [Po07t] as SUT detects the link down event and begins to converge IGP routes and traffic over the Next-Best Egress Interface.
6. Measure Rate-Derived Convergence Time [Po07t] as SUT detects the link down event and converges all IGP routes and traffic over the Next-Best Egress Interface. Optionally, Route-Specific Convergence Times [Po07t] MAY be measured.
7. Stop offered load. Wait 30 seconds for queues to drain. Restart offered load.
8. Restore link on Tester's Neighbor Interface connected to DUT's Preferred Egress Interface.
9. Measure Reversion Convergence Time [Po07t], and optionally measure First Route Convergence Time [Po07t] and Route-Specific Convergence Times [Po07t], as DUT detects the link up event and converges all IGP routes and traffic back to the Preferred Egress Interface.

Results
The measured IGP Convergence time is influenced by the link failure indication, LSA/LSP Flood Packet Pacing, LSA/LSP Retransmission Packet Pacing, LSA/LSP Generation time, SPF delay, SPF Hold time, SPF Execution Time, Tree Build Time, and Hardware Update Time [Po07a]. This test case may produce Stale Forwarding [Po07t] due to microloops which may increase the measured convergence times.

4.3 Convergence Due to Local Adminstrative Shutdown

Objective

To obtain the IGP Route Convergence due to a administrative shutdown at the DUT's Local Interface.

Procedure

1. Advertise matching IGP routes from Tester to DUT on Preferred Egress Interface [Po07t] and Next-Best Egress Interface [Po07t] using the topology shown in Figure 1. Set the cost of the routes so that the Preferred Egress Interface is the preferred next-hop.
2. Send offered load at measured Throughput with fixed packet size to destinations matching all IGP routes from Tester to DUT on Ingress Interface [Po07t].
3. Verify traffic is routed over Preferred Egress Interface.
4. Perform adminstrative shutdown on the DUT's Preferred Egress Interface. This is the Convergence Event [Po07t] that produces the Convergence Event Instant [Po07t].
5. Measure First Route Convergence Time [Po07t] as DUT detects the link down event and begins to converge IGP routes and traffic over the Next-Best Egress Interface.
6. Measure Rate-Derived Convergence Time [Po07t] as DUT converges all IGP routes and traffic over the Next-Best Egress Interface. Optionally, Route-Specific Convergence Times [Po07t] MAY be measured.
7. Stop offered load. Wait 30 seconds for queues to drain. Restart offered load.
8. Restore Preferred Egress Interface by administratively enabling the interface.
9. Measure Reversion Convergence Time [Po07t], and optionally measure First Route Convergence Time [Po07t] and Route-Specific Convergence Times [Po07t], as DUT detects the link up event and converges all IGP routes and traffic back to the Preferred Egress Interface.

Results

The measured IGP Convergence time is influenced by SPF delay, SPF Hold time, SPF Execution Time, Tree Build Time, and Hardware Update Time [Po07a].

4.4 Convergence Due to Layer 2 Session Loss

Objective

To obtain the IGP Route Convergence due to a local Layer 2 loss.

Procedure

1. Advertise matching IGP routes from Tester to DUT on Preferred Egress Interface [Po07t] and Next-Best Egress Interface [Po07t] using the topology shown in Figure 1. Set the cost of

the routes so that the IGP routes along the Preferred Egress Interface is the preferred next-hop.

2. Send offered load at measured Throughput with fixed packet size to destinations matching all IGP routes from Tester to DUT on Ingress Interface [Po07t].

3. Verify traffic is routed over Preferred Egress Interface.
4. Tester removes Layer 2 session from DUT's Preferred Egress Interface [Po07t]. It is RECOMMENDED that this be achieved with messaging, but the method MAY vary with the Layer 2 protocol. This is the Convergence Event [Po07t] that produces the Convergence Event Instant [Po07t].
5. Measure First Route Convergence Time [Po07t] as DUT detects the Layer 2 session down event and begins to converge IGP routes and traffic over the Next-Best Egress Interface.
6. Measure Rate-Derived Convergence Time [Po07t] as DUT detects the Layer 2 session down event and converges all IGP routes and traffic over the Next-Best Egress Interface. Optionally, Route-Specific Convergence Times [Po07t] MAY be measured.
7. Stop offered load. Wait 30 seconds for queues to drain. Restart offered load.
8. Restore Layer 2 session on DUT's Preferred Egress Interface.
9. Measure Reversion Convergence Time [Po07t], and optionally measure First Route Convergence Time [Po07t] and Route-Specific Convergence Times [Po07t], as DUT detects the session up event and converges all IGP routes and traffic over the Preferred Egress Interface.

Results

The measured IGP Convergence time is influenced by the Layer 2 failure indication, SPF delay, SPF Hold time, SPF Execution Time, Tree Build Time, and Hardware Update Time [Po07a].

4.5 Convergence Due to Loss of IGP Adjacency

Objective

To obtain the IGP Route Convergence due to loss of the IGP Adjacency.

Procedure

1. Advertise matching IGP routes from Tester to DUT on Preferred Egress Interface [Po07t] and Next-Best Egress Interface [Po07t] using the topology shown in Figure 1. Set the cost of the routes so that the Preferred Egress Interface is the preferred next-hop.
2. Send offered load at measured Throughput with fixed packet size to destinations matching all IGP routes from Tester to DUT on Ingress Interface [Po07t].
3. Verify traffic is routed over Preferred Egress Interface.
4. Remove IGP adjacency from Tester's Neighbor Interface [Po07t] connected to Preferred Egress Interface. The Layer 2 session MUST be maintained. This is the Convergence Event [Po07t] that produces the Convergence Event Instant [Po07t].
5. Measure First Route Convergence Time [Po07t] as DUT detects the loss of IGP adjacency and begins to converge IGP routes and

traffic over the Next-Best Egress Interface.

6. Measure Rate-Derived Convergence Time [Po07t] as DUT detects the IGP session failure event and converges all IGP routes and traffic over the Next-Best Egress Interface. Optionally, Route-Specific Convergence Times [Po07t] MAY be measured.

7. Stop offered load. Wait 30 seconds for queues to drain. Restart offered load.
8. Restore IGP session on DUT's Preferred Egress Interface.
9. Measure Reversion Convergence Time [Po07t], and optionally measure First Route Convergence Time [Po07t] and Route-Specific Convergence Times [Po07t], as DUT detects the session recovery event and converges all IGP routes and traffic over the Preferred Egress Interface.

Results
The measured IGP Convergence time is influenced by the IGP Hello Interval, IGP Dead Interval, SPF delay, SPF Hold time, SPF Execution Time, Tree Build Time, and Hardware Update Time [Po07a].

4.6 Convergence Due to Route Withdrawal

Objective
To obtain the IGP Route Convergence due to Route Withdrawal.

Procedure
1. Advertise matching IGP routes from Tester to DUT on Preferred Egress Interface [Po07t] and Next-Best Egress Interface [Po07t] using the topology shown in Figure 1. Set the cost of the routes so that the Preferred Egress Interface is the preferred next-hop. It is RECOMMENDED that the IGP routes be IGP external routes for which the Tester would be emulating a preferred and a next-best Autonomous System Border Router (ASBR).
2. Send offered load at measured Throughput with fixed packet size to destinations matching all IGP routes from Tester to DUT on Ingress Interface [Po07t].
3. Verify traffic is routed over Preferred Egress Interface.
4. Tester withdraws all IGP routes from DUT's Local Interface on Preferred Egress Interface. The Tester records the time it sends the withdrawal message(s). This MAY be achieved with inclusion of a timestamp in the traffic payload. This is the Convergence Event [Po07t] that produces the Convergence Event Instant [Po07t].
5. Measure First Route Convergence Time [Po07t] as DUT detects the route withdrawal event and begins to converge IGP routes and traffic over the Next-Best Egress Interface. This is measured from the time that the Tester sent the withdrawal message(s).
6. Measure Rate-Derived Convergence Time [Po07t] as DUT withdraws routes and converges all IGP routes and traffic over the Next-Best Egress Interface. Optionally, Route-Specific Convergence Times [Po07t] MAY be measured.
7. Stop offered load. Wait 30 seconds for queues to drain. Restart offered load.
8. Re-advertise IGP routes to DUT's Preferred Egress Interface.

9. Measure Reversion Convergence Time [Po07t], and optionally measure First Route Convergence Time [Po07t] and Route-Specific Convergence Times [Po07t], as DUT converges all IGP routes and traffic over the Preferred Egress Interface.

Results
The measured IGP Convergence time is the SPF Processing and FIB Update time as influenced by the SPF or route calculation delay, Hold time, Execution Time, and Hardware Update Time [Po07a].

4.7 Convergence Due to Cost Change
Objective
To obtain the IGP Route Convergence due to route cost change.

Procedure
1. Advertise matching IGP routes from Tester to DUT on Preferred Egress Interface [Po07t] and Next-Best Egress Interface [Po07t] using the topology shown in Figure 1. Set the cost of the routes so that the Preferred Egress Interface is the preferred next-hop.
2. Send offered load at measured Throughput with fixed packet size to destinations matching all IGP routes from Tester to DUT on Ingress Interface [Po07t].
3. Verify traffic is routed over Preferred Egress Interface.
4. Tester increases cost for all IGP routes at DUT's Preferred Egress Interface so that the Next-Best Egress Interface has lower cost and becomes preferred path. This is the Convergence Event [Po07t] that produces the Convergence Event Instant [Po07t].
5. Measure First Route Convergence Time [Po07t] as DUT detects the cost change event and begins to converge IGP routes and traffic over the Next-Best Egress Interface.
6. Measure Rate-Derived Convergence Time [Po07t] as DUT detects the cost change event and converges all IGP routes and traffic over the Next-Best Egress Interface. Optionally, Route-Specific Convergence Times [Po07t] MAY be measured.
7. Stop offered load. Wait 30 seconds for queues to drain. Restart offered load.
8. Re-advertise IGP routes to DUT's Preferred Egress Interface with original lower cost metric.
9. Measure Reversion Convergence Time [Po07t], and optionally measure First Route Convergence Time [Po07t] and Route-Specific Convergence Times [Po07t], as DUT converges all IGP routes and traffic over the Preferred Egress Interface.

Results
There should be no measured packet loss for this case.

4.8 Convergence Due to ECMP Member Interface Failure

Objective
To obtain the IGP Route Convergence due to a local link failure event of an ECMP Member.

Procedure
1. Configure ECMP Set as shown in Figure 3.
2. Advertise matching IGP routes from Tester to DUT on each ECMP member.
3. Send offered load at measured Throughput with fixed packet size to destinations matching all IGP routes from Tester to DUT on Ingress Interface [Po07t].
4. Verify traffic is routed over all members of ECMP Set.
5. Remove link on Tester's Neighbor Interface [Po07t] connected to one of the DUT's ECMP member interfaces. This is the Convergence Event [Po07t] that produces the Convergence Event Instant [Po07t].
6. Measure First Route Convergence Time [Po07t] as DUT detects the link down event and begins to converge IGP routes and traffic over the other ECMP members.
7. Measure Rate-Derived Convergence Time [Po07t] as DUT detects the link down event and converges all IGP routes and traffic over the other ECMP members. At the same time measure Out-of-Order Packets [Po06] and Duplicate Packets [Po06]. Optionally, Route-Specific Convergence Times [Po07t] MAY be measured.
8. Stop offered load. Wait 30 seconds for queues to drain. Restart offered load.
9. Restore link on Tester's Neighbor Interface connected to DUT's ECMP member interface.
10. Measure Reversion Convergence Time [Po07t], and optionally measure First Route Convergence Time [Po07t] and Route-Specific Convergence Times [Po07t], as DUT detects the link up event and converges IGP routes and some distribution of traffic over the restored ECMP member.

Results
The measured IGP Convergence time is influenced by Local link failure indication, Tree Build Time, and Hardware Update Time [Po07a].

4.9 Convergence Due to ECMP Member Remote Interface Failure

Objective
To obtain the IGP Route Convergence due to a remote interface failure event for an ECMP Member.

Procedure
1. Configure ECMP Set as shown in Figure 2 in which the links from R1 to R2 and R1 to R3 are members of an ECMP Set.
2. Advertise matching IGP routes from Tester to SUT to balance traffic to each ECMP member.
3. Send offered load at measured Throughput with fixed packet size to destinations matching all IGP routes from Tester to SUT on Ingress Interface [Po07t].
4. Verify traffic is routed over all members of ECMP Set.
5. Remove link on Tester's Neighbor Interface to R2 or R3. This is the Convergence Event [Po07t] that produces the Convergence Event Instant [Po07t].
6. Measure First Route Convergence Time [Po07t] as SUT detects the link down event and begins to converge IGP routes and traffic over the other ECMP members.
7. Measure Rate-Derived Convergence Time [Po07t] as SUT detects the link down event and converges all IGP routes and traffic over the other ECMP members. At the same time measure Out-of-Order Packets [Po06] and Duplicate Packets [Po06]. Optionally, Route-Specific Convergence Times [Po07t] MAY be measured.
8. Stop offered load. Wait 30 seconds for queues to drain. Restart offered load.
9. Restore link on Tester's Neighbor Interface to R2 or R3.
10. Measure Reversion Convergence Time [Po07t], and optionally measure First Route Convergence Time [Po07t] and Route-Specific Convergence Times [Po07t], as SUT detects the link up event and converges IGP routes and some distribution of traffic over the restored ECMP member.

Results
The measured IGP Convergence time is influenced by Local link failure indication, Tree Build Time, and Hardware Update Time [Po07a].

4.10 Convergence Due to Parallel Link Interface Failure

Objective
To obtain the IGP Route Convergence due to a local link failure event for a Member of a Parallel Link. The links can be used for data Load Balancing

Procedure

1. Configure Parallel Link as shown in Figure 4.
2. Advertise matching IGP routes from Tester to DUT on each Parallel Link member.

3. Send offered load at measured Throughput with fixed packet size to destinations matching all IGP routes from Tester to DUT on Ingress Interface [Po07t].
4. Verify traffic is routed over all members of Parallel Link.
5. Remove link on Tester's Neighbor Interface [Po07t] connected to one of the DUT's Parallel Link member interfaces. This is the Convergence Event [Po07t] that produces the Convergence Event Instant [Po07t].
6. Measure First Route Convergence Time [Po07t] as DUT detects the link down event and begins to converge IGP routes and traffic over the other Parallel Link members.
7. Measure Rate-Derived Convergence Time [Po07t] as DUT detects the link down event and converges all IGP routes and traffic over the other Parallel Link members. At the same time measure Out-of-Order Packets [Po06] and Duplicate Packets [Po06]. Optionally, Route-Specific Convergence Times [Po07t] MAY be measured.
8. Stop offered load. Wait 30 seconds for queues to drain. Restart offered load.
9. Restore link on Tester's Neighbor Interface connected to DUT's Parallel Link member interface.
10. Measure Reversion Convergence Time [Po07t], and optionally measure First Route Convergence Time [Po07t] and Route-Specific Convergence Times [Po07t], as DUT detects the link up event and converges IGP routes and some distribution of traffic over the restored Parallel Link member.

Results
The measured IGP Convergence time is influenced by the Local link failure indication, Tree Build Time, and Hardware Update Time [Po07a].

5. IANA Considerations

This document requires no IANA considerations.

6. Security Considerations

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

7. Acknowledgements

Thanks to Sue Hares, Al Morton, Kevin Dubray, Ron Bonica, David Ward, Kris Michielsen and the BMWG for their contributions to this work.

8. References

8.1 Normative References

[Br91] Bradner, S., "Benchmarking Terminology for Network Interconnection Devices", RFC 1242, IETF, March 1991.

[Br97] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", RFC 2119, March 1997

[Br99] Bradner, S. and McQuaid, J., "Benchmarking Methodology for Network Interconnect Devices", RFC 2544, IETF, March 1999.

[Ca90] Callon, R., "Use of OSI IS-IS for Routing in TCP/IP and Dual Environments", RFC 1195, IETF, December 1990.

[Ma98] Mandeville, R., "Benchmarking Terminology for LAN Switching Devices", RFC 2285, February 1998.

[Mo98] Moy, J., "OSPF Version 2", RFC 2328, IETF, April 1998.

[Po06] Poretsky, S., et al., "Terminology for Benchmarking Network-layer Traffic Control Mechanisms", RFC 4689, November 2006.

[Po07a] Poretsky, S., "Considerations for Benchmarking Link-State IGP Convergence", draft-ietf-bmwg-igp-dataplane-conv-app-16, work in progress, October 2008.

[Po07t] Poretsky, S., Imhoff, B., "Benchmarking Terminology for Link-State IGP Convergence", draft-ietf-bmwg-igp-dataplane-conv-term-16, work in progress, October 2008.

8.2 Informative References

None

9. Author's Address

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

Acknowledgement

Funding for the RFC Editor function is currently provided by the Internet Society.