

Internet Engineering Task Force (IETF)
Request for Comments: 6414
Category: Informational
ISSN: 2070-1721

S. Poretsky
Allot Communications
R. Papneja
Huawei
J. Karthik
S. Vapiwala
Cisco Systems
November 2011

Benchmarking Terminology for Protection Performance

Abstract

This document provides common terminology and metrics for benchmarking the performance of sub-IP layer protection mechanisms. The performance benchmarks are measured at the IP layer; protection may be provided at the sub-IP layer. The benchmarks and terminology can be applied in methodology documents for different sub-IP layer protection mechanisms such as Automatic Protection Switching (APS), Virtual Router Redundancy Protocol (VRRP), Stateful High Availability (HA), and Multiprotocol Label Switching Fast Reroute (MPLS-FRR).

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/rfc6414>.

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	4
1.1. Scope	4
1.2. General Model	5
2. Existing Definitions	8
3. Test Considerations	9
3.1. Paths	9
3.1.1. Path	9
3.1.2. Working Path	10
3.1.3. Primary Path	10
3.1.4. Protected Primary Path	11
3.1.5. Backup Path	11
3.1.6. Standby Backup Path	12
3.1.7. Dynamic Backup Path	12
3.1.8. Disjoint Paths	13
3.1.9. Point of Local Repair (PLR)	13
3.1.10. Shared Risk Link Group (SRLG)	14
3.2. Protection	14
3.2.1. Link Protection	14
3.2.2. Node Protection	15

3.2.3. Path Protection	15
3.2.4. Backup Span	16
3.2.5. Local Link Protection	16
3.2.6. Redundant Node Protection	17
3.2.7. State Control Interface	17
3.2.8. Protected Interface	18
3.3. Protection Switching	18
3.3.1. Protection-Switching System	18
3.3.2. Failover Event	19
3.3.3. Failure Detection	19
3.3.4. Failover	20
3.3.5. Restoration	20
3.3.6. Reversion	21
3.4. Nodes	22
3.4.1. Protection-Switching Node	22
3.4.2. Non-Protection-Switching Node	22
3.4.3. Headend Node	23
3.4.4. Backup Node	23
3.4.5. Merge Node	24
3.4.6. Primary Node	24
3.4.7. Standby Node	25
3.5. Benchmarks	26
3.5.1. Failover Packet Loss	26
3.5.2. Reversion Packet Loss	26
3.5.3. Failover Time	27
3.5.4. Reversion Time	27
3.5.5. Additive Backup Delay	28
3.6. Failover Time Calculation Methods	28
3.6.1. Time-Based Loss Method (TBLM)	29
3.6.2. Packet-Loss-Based Method (PLBM)	29
3.6.3. Timestamp-Based Method (TBM)	30
4. Security Considerations	31
5. References	32
5.1. Normative References	32
5.2. Informative References	32
6. Acknowledgments	32

1. Introduction

The IP network layer provides route convergence to protect data traffic against planned and unplanned failures in the Internet. Fast convergence times are critical to maintain reliable network connectivity and performance. Convergence Events [6] are recognized at the IP Layer so that Route Convergence [6] occurs. Technologies that function at sub-IP layers can be enabled to provide further protection of IP traffic by providing the failure recovery at the sub-IP layers so that the outage is not observed at the IP layer. Such sub-IP protection technologies include, but are not limited to, High Availability (HA) stateful failover, Virtual Router Redundancy Protocol (VRRP) [8], Automatic Link Protection (APS) for SONET/SDH, Resilient Packet Ring (RPR) for Ethernet, and Fast Reroute for Multiprotocol Label Switching (MPLS-FRR) [9].

1.1. Scope

Benchmarking terminology was defined for IP-layer convergence in [6]. Different terminology and methodologies specific to benchmarking sub-IP layer protection mechanisms are required. The metrics for benchmarking the performance of sub-IP protection mechanisms are measured at the IP layer, so that the results are always measured in reference to IP and independent of the specific protection mechanism being used. The purpose of this document is to provide a single terminology for benchmarking sub-IP protection mechanisms.

A common terminology for sub-IP layer protection mechanism benchmarking enables different implementations of a protection mechanism to be benchmarked and evaluated. In addition, implementations of different protection mechanisms can be benchmarked and evaluated. It is intended that there can exist unique methodology documents for each sub-IP protection mechanism based upon this common terminology document. The terminology can be applied to methodologies that benchmark sub-IP protection mechanism performance with a single stream of traffic or multiple streams of traffic. The traffic flow may be unidirectional or bidirectional as to be indicated in the methodology.

1.2. General Model

The sequence of events to benchmark the performance of sub-IP protection mechanisms is as follows:

1. Failover Event - Primary Path fails
2. Failure Detection - Failover Event is detected
3. Failover - Backup Path becomes the Working Path due to Failover Event
4. Restoration - Primary Path recovers from a Failover Event
5. Reversion (optional) - Primary Path becomes the Working Path

These terms are further defined in this document.

Figures 1 through 5 show models that MAY be used when benchmarking sub-IP protection mechanisms, which MUST use a Protection-Switching System that consists of a minimum of two Protection-Switching Nodes, an Ingress Node known as the Headend Node and an Egress Node known as the Merge Node. The Protection-Switching System MUST include either a Primary Path and Backup Path, as shown in Figures 1 through 4, or a Primary Node and Standby Node, as shown in Figure 5. A Protection-Switching System may provide link protection, node protection, path protection, local link protection, and high availability, as shown in Figures 1 through 5, respectively. A Failover Event occurs along the Primary Path or at the Primary Node. The Working Path is the Primary Path prior to the Failover Event and the Backup Path after the Failover Event. A Tester is set outside the two paths or nodes as it sends and receives IP traffic along the Working Path. The tester MUST record the IP packet sequence numbers, departure time, and arrival time so that the metrics of Failover Time, Additive Latency, Packet Reordering, Duplicate Packets, and Reversion Time can be measured. The Tester may be a single device or a test system. If Reversion is supported, then the Working Path is the Primary Path after Restoration (Failure Recovery) of the Primary Path.

Link Protection, as shown in Figure 1, provides protection when a Failover Event occurs on the link between two nodes along the Primary Path. Node Protection, as shown in Figure 2, provides protection when a Failover Event occurs at a Node along the Primary Path. Path Protection, as shown in Figure 3, provides protection for link or node failures for multiple hops along the Primary Path. Local Link Protection, as shown in Figure 4, provides sub-IP protection of a link between two nodes, without a Backup Node. An example of such a sub-IP protection mechanism is SONET APS. High Availability Protection, as shown in Figure 5, provides protection of a Primary Node with a redundant Standby Node. State Control is provided between the Primary and Standby Nodes. Failure of the Primary Node

is detected at the sub-IP layer to force traffic to switch to the Standby Node, which has state maintained for zero or minimal packet loss.

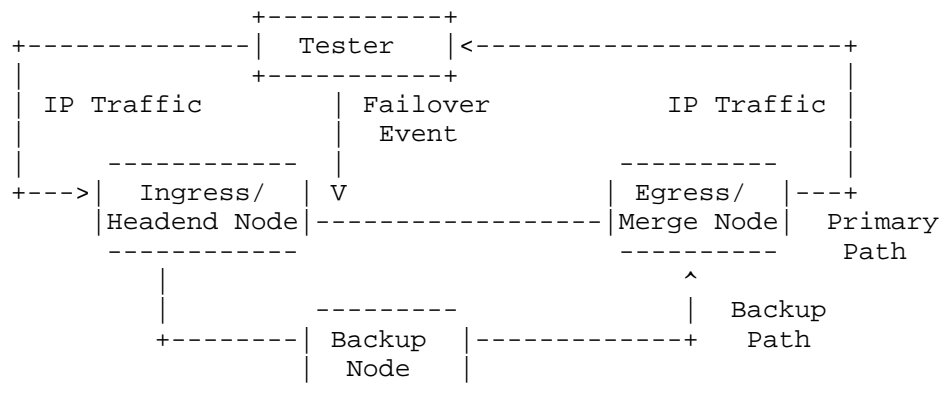


Figure 1. System Under Test (SUT) for Sub-IP Link Protection

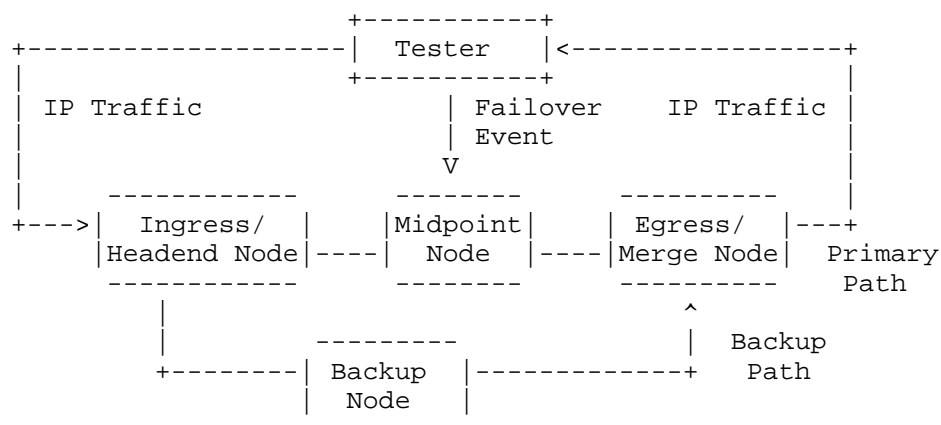


Figure 2. System Under Test (SUT) for Sub-IP Node Protection

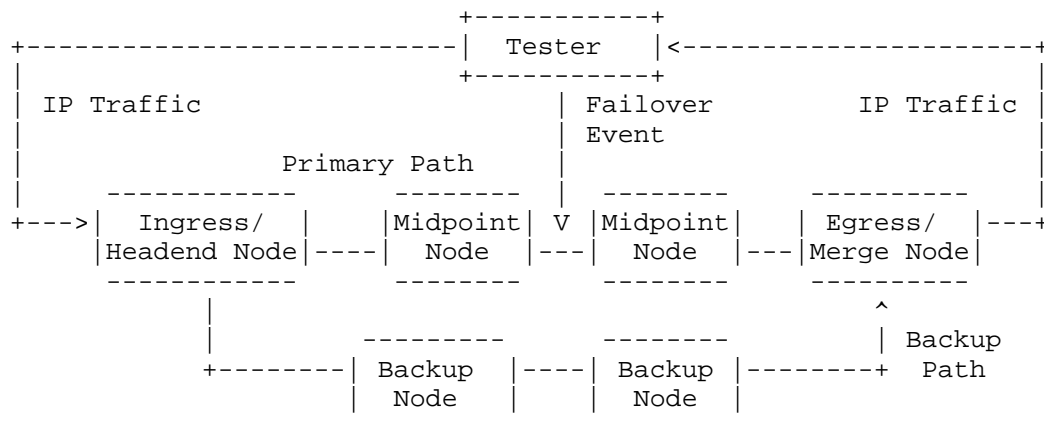


Figure 3. System Under Test (SUT) for Sub-IP Path Protection




Figure 4. System Under Test (SUT) for Sub-IP Local Link Protection

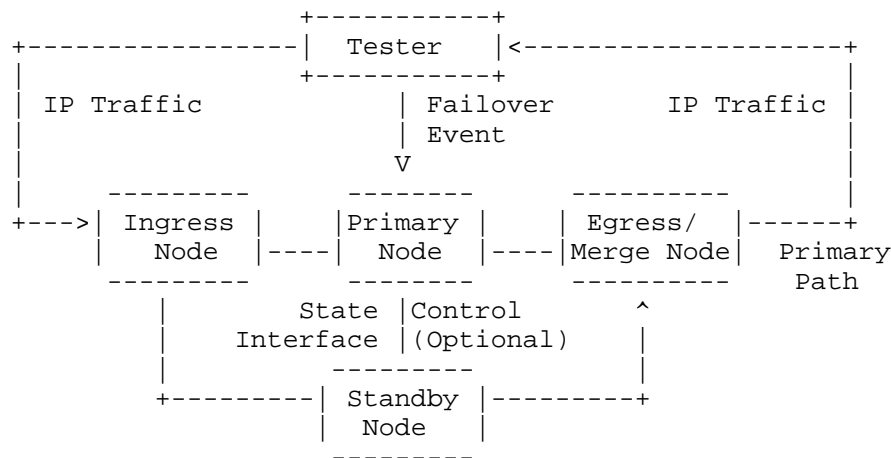


Figure 5. System Under Test (SUT)
for Sub-IP Redundant Node Protection

Some protection-switching technologies may use a series of steps that differ from the general model. The specific differences SHOULD be highlighted in each technology-specific methodology. Note that some protection-switching technologies are endowed with the ability to re-optimize the working path after a node or link failure.

2. Existing Definitions

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

Latency	[2], Section 3.8
Frame Loss Rate	[2], Section 3.6
Throughput	[2], Section 3.17
Device Under Test (DUT)	[3], Section 3.1.1
System Under Test (SUT)	[3], Section 3.1.2
Offered Load	[3], Section 3.5.2
Out-of-order Packet	[4], Section 3.3.4
Duplicate Packet	[4], Section 3.3.5
Forwarding Delay	[4], Section 3.2.4
Jitter	[4], Section 3.2.5
Packet Loss	[6], Section 3.5
Packet Reordering	[7], Section 3.3

This document has the following frequently used acronyms:

DUT	Device Under Test
SUT	System Under Test

This document adopts the definition format in Section 2 of RFC 1242 [2]. Terms defined in this document are capitalized when used within this document.

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 [5]. 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. Test Considerations

3.1. Paths

3.1.1. Path

Definition:

A unidirectional sequence of nodes $\langle R_1, \dots, R_n \rangle$ and links $\langle L_{12}, \dots, L_{(n-1)n} \rangle$ with the following properties:

- a. R_1 is the ingress node and forwards IP packets, which input into DUT/SUT, to R_2 as sub-IP frames over link L_{12} .
- b. R_i is a node which forwards data frames to $R_{(i+1)}$ over Link $L_{i(i+1)}$ for all i , $1 < i < n-1$, based on information in the sub-IP layer.
- c. R_n is the egress node, and it outputs sub-IP frames from DUT/SUT as IP packets. $L_{(n-1)n}$ is the link between the $R_{(n-1)}$ and R_n .

Discussion:

The path is defined in the sub-IP layer in this document, unlike an IP path in RFC 2026 [1]. One path may be regarded as being equivalent to one IP link between two IP nodes, i.e., R_1 and R_n . The two IP nodes may have multiple paths for protection. A packet will travel on only one path between the nodes. Packets belonging to a microflow [10] will traverse one or more paths. The path is unidirectional. For example, the link between R_1 and R_2 in the direction from R_1 to R_2 is L_{12} . For traffic flowing in the reverse direction from R_2 to R_1 , the link is L_{21} . Example paths are the SONET/SDH path and the label switched path for MPLS.

Measurement Units:

n/a

Issues:

"A bidirectional path", which transmits traffic in both directions along the same nodes, consists of two unidirectional paths. Therefore, the two unidirectional paths belonging to "one bidirectional path" will be treated independently when benchmarking for "a bidirectional path".

See Also:

Working Path
Primary Path
Backup Path

3.1.2. Working Path

Definition:

The path that the DUT/SUT is currently using to forward packets.

Discussion:

A Primary Path is the Working Path before occurrence of a Failover Event. A Backup Path shall become the Working Path after a Failover Event.

Measurement Units:

n/a

Issues:

None.

See Also:

Path
Primary Path
Backup Path

3.1.3. Primary Path

Definition:

The preferred point-to-point path for forwarding traffic between two or more nodes.

Discussion:

The Primary Path is the Path that traffic traverses prior to a Failover Event.

Measurement Units:

n/a

Issues:

None.

See Also:

Path
Failover Event

3.1.4. Protected Primary Path

Definition:

A Primary Path that is protected with a Backup Path.

Discussion:

A Protected Primary Path must include at least one Protection-Switching Node.

Measurement Units:

n/a

Issues:

None.

See Also:

Path
Primary Path

3.1.5. Backup Path

Definition:

A path that exists to carry data traffic only if a Failover Event occurs on a Primary Path.

Discussion:

The Backup Path shall become the Working Path upon a Failover Event. A Path may have one or more Backup Paths. A Backup Path may protect one or more Primary Paths. There are various types of Backup Paths:

- a. dedicated recovery Backup Path (1+1) or (1:1), which has 100% redundancy for a specific ordinary path
- b. shared Backup Path (1:N), which is dedicated to the protection for more than one specific Primary Path
- c. associated shared Backup Path (M:N) for which a specific set of Backup Paths protects a specific set of more than one Primary Path

A Backup Path may be signaled or un signaled. The Backup Path must be created prior to the Failover Event. The Backup Path generally originates at the point of local repair (PLR) and terminates at a node along a primary path.

Measurement Units:

n/a

Issues:

None.

See Also:

Path
Working Path
Primary Path

3.1.6. Standby Backup Path

Definition:

A Backup Path that is established prior to a Failover Event to protect a Primary Path.

Discussion:

The Standby Backup Path and Dynamic Backup Path provide protection, but are established at different times.

Measurement Units:

n/a

Issues:

None.

See Also:

Backup Path
Primary Path
Failover Event

3.1.7. Dynamic Backup Path

Definition:

A Backup Path that is established upon occurrence of a Failover Event.

Discussion:

The Standby Backup Path and Dynamic Backup Path provide protection, but are established at different times.

Measurement Units:

n/a

Issues:

None.

See Also:

Backup Path

Standby Backup Path

Failover Event

3.1.8. Disjoint Paths

Definition:

A pair of paths that do not share a common link or nodes.

Discussion:

Two paths are disjoint if they do not share a common node or link other than the ingress and egress.

Measurement Units:

n/a

Issues:

None.

See Also:

Path

Primary Path

SRLG

3.1.9. Point of Local Repair (PLR)

Definition:

A node capable of Failover along the Primary Path that is also the ingress node for the Backup Path to protect another node or link.

Discussion:

Any node along the Primary Path from the ingress node to the penultimate node may be a PLR. The PLR may use a single Backup Path for protecting one or more Primary Paths. There can be multiple PLRs along a Primary Path. The PLR must be an ingress to a Backup Path. The PLR can be any node along the Primary Path except the egress node of the Primary Path. The PLR may simultaneously be a Headend Node when it is serving the role as ingress to the Primary Path and the Backup Path. If the PLR is also the Headend Node, then the Backup Path is a Disjoint Path from the ingress to the Merge Node.

Measurement Units:

n/a

Issues:

None.

See Also:

Primary Path
Backup Path
Failover

3.1.10. Shared Risk Link Group (SRLG)

Definition:

SRLG is a set of links that share the same risk (physical or logical) within a network.

Discussion:

SRLG is considered the set of links to be avoided when the primary and secondary paths are considered disjoint. The SRLG will fail as a group if the shared resource (physical or anything abstract such as software version) fails.

Measurement Units:

n/a

Issues:

None.

See Also:

Path Primary Path

3.2. Protection

3.2.1. Link Protection

Definition:

A Backup Path that is signaled to at least one Backup Node to protect for failure of interfaces and links along a Primary Path.

Discussion:

Link Protection may or may not protect the entire Primary Path.
Link Protection is shown in Figure 1.

Measurement Units:

n/a

Issues:
None.

See Also:
Primary Path Backup Path

3.2.2. Node Protection

Definition:
A Backup Path that is signaled to at least one Backup Node to protect for failure of interfaces, links, and nodes along a Primary Path.

Discussion:
Node Protection may or may not protect the entire Primary Path. Node Protection also provides Link Protection. Node Protection is shown in Figure 2.

Measurement Units:
n/a

Issues:
None.

See Also:
Link Protection

3.2.3. Path Protection

Definition:
A Backup Path that is signaled to at least one Backup Node to provide protection along the entire Primary Path.

Discussion:
Path Protection provides Node Protection and Link Protection for every node and link along the Primary Path. A Backup Path providing Path Protection may have the same ingress node as the Primary Path. Path Protection is shown in Figure 3.

Measurement Units:
n/a

Issues:
None.

See Also:

- Primary Path
- Backup Path
- Node Protection
- Link Protection

3.2.4. Backup Span

Definition:

The number of hops used by a Backup Path.

Discussion:

The Backup Span is an integer obtained by counting the number of nodes along the Backup Path.

Measurement Units:

number of nodes

Issues:

None.

See Also:

- Primary Path
- Backup Path

3.2.5. Local Link Protection

Definition:

A Backup Path that is a redundant path between two nodes and does not use a Backup Node.

Discussion:

Local Link Protection must be provided as a Backup Path between two nodes along the Primary Path without the use of a Backup Node. Local Link Protection is provided by Protection-Switching Systems such as SONET APS. Local Link Protection is shown in Figure 4.

Measurement Units:

n/a

Issues:

None.

See Also:

- Backup Path
- Backup Node

3.2.6. Redundant Node Protection

Definition:

A Protection-Switching System with a Primary Node protected by a Standby Node along the Primary Path.

Discussion:

Redundant Node Protection is provided by Protection-Switching Systems such as VRRP and HA. The protection mechanisms occur at sub-IP layers to switch traffic from a Primary Node to Backup Node upon a Failover Event at the Primary Node. Traffic continues to traverse the Primary Path through the Standby Node. The failover may be stateful, in which the state information may be exchanged in-band or over an out-of-band State Control Interface. The Standby Node may be active or passive. Redundant Node Protection is shown in Figure 5.

Measurement Units:

n/a

Issues:

None.

See Also:

Primary Path
Primary Node
Standby Node

3.2.7. State Control Interface

Definition:

An out-of-band control interface used to exchange state information between the Primary Node and Standby Node.

Discussion:

The State Control Interface may be used for Redundant Node Protection. The State Control Interface should be out-of-band. It is possible to have Redundant Node Protection in which there is no state control or state control is provided in-band. The State Control Interface between the Primary and Standby Node may be one or more hops.

Measurement Units:

n/a

Issues:

None.

See Also:

- Primary Node
- Standby Node

3.2.8. Protected Interface

Definition:

An interface along the Primary Path that is protected by a Backup Path.

Discussion:

A Protected Interface is an interface protected by a Protection-Switching System that provides Link Protection, Node Protection, Path Protection, Local Link Protection, and Redundant Node Protection.

Measurement Units:

n/a

Issues:

None.

See Also:

- Primary Path
- Backup Path

3.3. Protection Switching

3.3.1. Protection-Switching System

Definition:

A DUT/SUT that is capable of Failure Detection and Failover from a Primary Path to a Backup Path or Standby Node when a Failover Event occurs.

Discussion:

The Protection-Switching System must include either a Primary Path and Backup Path, as shown in Figures 1 through 4, or a Primary Node and Standby Node, as shown in Figure 5. The Backup Path may be a Standby Backup Path or a Dynamic Backup Path. The Protection-Switching System includes the mechanisms for both Failure Detection and Failover.

Measurement Units:

n/a

Issues:

None.

See Also:

Primary Path Backup Path Failover

3.3.2. Failover Event

Definition:

The occurrence of a planned or unplanned action in the network that results in a change in the Path that data traffic traverses.

Discussion:

Failover Events include, but are not limited to, link failure and router failure. Routing changes are considered Convergence Events [6] and are not Failover Events. This restricts Failover Events to sub-IP layers. Failover may be at the PLR or at the ingress. If the failover is at the ingress, it is generally on a disjoint path from the ingress to egress.

Failover Events may result from failures such as link failure or router failure. The change in path after Failover may have a Backup Span of one or more nodes. Failover Events are distinguished from routing changes and Convergence Events [6] by the detection of the failure and subsequent protection switching at a sub-IP layer. Failover occurs at a PLR or Primary Node.

Measurement Units:

n/a

Issues:

None.

See Also:

Path

Failure Detection

Disjoint Path

3.3.3. Failure Detection

Definition:

The process to identify at a sub-IP layer a Failover Event at a Primary Node or along the Primary Path.

Discussion:

Failure Detection occurs at the Primary Node or ingress node of the Primary Path. Failure Detection occurs via a sub-IP mechanism such as detection of a link down event or timeout for receipt of a control packet. A failure may be completely isolated. A failure

may affect a set of links that share a single SRLG (e.g., port with many sub-interfaces). A failure may affect multiple links that are not part of the SRLG.

Measurement Units:

n/a

Issues:

None.

See Also:

Primary Path

3.3.4. Failover

Definition:

The process to switch data traffic from the protected Primary Path to the Backup Path upon Failure Detection of a Failover Event.

Discussion:

Failover to a Backup Path provides Link Protection, Node Protection, or Path Protection. Failover is complete when Packet Loss [6], Out-of-order Packets [4], and Duplicate Packets [4] are no longer observed. Forwarding Delay [4] may continue to be observed.

Measurement Units:

n/a

Issues:

None.

See Also:

Primary Path Backup Path Failover Event

3.3.5. Restoration

Definition:

The state of failover recovery in which the Primary Path has recovered from a Failover Event, but is not yet forwarding packets because the Backup Path remains the Working Path.

Discussion:

Restoration must occur while the Backup Path is the Working Path. The Backup Path is maintained as the Working Path during Restoration. Restoration produces a Primary Path that is

recovered from failure, but is not yet forwarding traffic. Traffic is still being forwarded by the Backup Path functioning as the Working Path.

Measurement Units:

n/a

Issues:

None.

See Also:

Primary Path
Failover Event
Failure Recovery
Working Path
Backup Path

3.3.6. Reversion

Definition:

The state of failover recovery in which the Primary Path has become the Working Path so that it is forwarding packets.

Discussion:

Protection-Switching Systems may or may not support Reversion. Reversion, if supported, must occur after Restoration. Packet forwarding on the Primary Path resulting from Reversion may occur either fully or partially over the Primary Path. A potential problem with Reversion is the discontinuity in end-to-end delay when the Forwarding Delays [4] along the Primary Path and Backup Path are different, possibly causing Out-of-order Packets [4], Duplicate Packets [4], and increased Jitter [4].

Measurement Units:

n/a

Issues:

None.

See Also:

Protection-Switching System
Working Path
Primary Path

3.4. Nodes

3.4.1. Protection-Switching Node

Definition:

A node that is capable of participating in a Protection Switching System.

Discussion:

The Protection-Switching Node may be an ingress or egress for a Primary Path or Backup Path, such as used for MPLS Fast Reroute configurations. The Protection-Switching Node may provide Redundant Node Protection as a Primary Node in a Redundant chassis configuration with a Standby Node, such as used for VRRP and HA configurations.

Measurement Units:

n/a

Issues:

None.

See Also:

Protection-Switching System

3.4.2. Non-Protection-Switching Node

Definition:

A node that is not capable of participating in a Protection Switching System, but may exist along the Primary Path or Backup Path.

Discussion:

None.

Measurement Units:

n/a

Issues:

None.

See Also:

Protection-Switching System
Primary Path
Backup Path

3.4.3. Headend Node

Definition:

The ingress node of the Primary Path.

Discussion:

The Headend Node may also be a PLR when it is serving in the dual role as the ingress to the Backup Path.

Measurement Units:

n/a

Issues:

None.

See Also:

Primary Path
PLR
Failover

3.4.4. Backup Node

Definition:

A node along the Backup Path.

Discussion:

The Backup Node can be any node along the Backup Path. There may be one or more Backup Nodes along the Backup Path. A Backup Node may be the ingress, midpoint, or egress of the Backup Path. If the Backup Path has only one Backup Node, then that Backup Node is the ingress and egress of the Backup Path.

Measurement Units:

n/a

Issues:

None.

See Also:

Backup Path

3.4.5. Merge Node

Definition:

A node along the Primary Path where Backup Path terminates.

Discussion:

The Merge Node can be any node along the Primary Path except the ingress node of the Primary Path. There can be multiple Merge Nodes along a Primary Path. A Merge Node can be the egress node for a single Backup Path or multiple Backup Paths. The Merge Node must be the egress to the Backup Path. The Merge Node may also be the egress of the Primary Path or Point of Local Repair (PLR).

Measurement Units:

n/a

Issues:

None.

See Also:

Primary Path
Backup Path
PLR
Failover

3.4.6. Primary Node

Definition:

A node along the Primary Path that is capable of Failover to a redundant Standby Node.

Discussion:

The Primary Node may be used for Protection-Switching Systems that provide Redundant Node Protection, such as VRRP and HA.

Measurement Units:

n/a

Issues:

None.

See Also:

Protection-Switching System Redundant Node Protection Standby Node

3.4.7. Standby Node

Definition:

A redundant node to a Primary Node; it forwards traffic along the Primary Path upon Failure Detection of the Primary Node.

Discussion:

The Standby Node must be used for Protection-Switching Systems that provide Redundant Node Protection, such as VRRP and HA. The Standby Node must provide protection along the same Primary Path. If the failover is to a Disjoint Path, then it is a Backup Node. The Standby Node may be configured for 1:1 or N:1 protection.

The communication between the Primary Node and Standby Node may be in-band or across an out-of-band State Control Interface. The Standby Node may be geographically dispersed from the Primary Node. When geographically dispersed, the number of hops of separation may increase failover time.

The Standby Node may be passive or active. The Passive Standby Node is not offered traffic and does not forward traffic until Failure Detection of the Primary Node. Upon Failure Detection of the Primary Node, traffic offered to the Primary Node is instead offered to the Passive Standby Node. The Active Standby Node is offered traffic and forwards traffic along the Primary Path while the Primary Node is also active. Upon Failure Detection of the Primary Node, traffic offered to the Primary Node is switched to the Active Standby Node.

Measurement Units:

n/a

Issues:

None.

See Also:

Primary Node
State Control Interface

3.5. Benchmarks

3.5.1. Failover Packet Loss

Definition:

The amount of packet loss produced by a Failover Event until Failover completes, where the measurement begins when the last unimpaired packet is received by the Tester on the Protected Primary Path and ends when the first unimpaired packet is received by the Tester on the Backup Path.

Discussion:

Packet loss can be observed as a reduction of forwarded traffic from the maximum forwarding rate. Failover Packet Loss includes packets that were lost, reordered, or delayed. Failover Packet Loss may reach 100% of the offered load.

Measurement Units:

Number of Packets

Issues:

None.

See Also:

Failover Event
Failover

3.5.2. Reversion Packet Loss

Definition:

The amount of packet loss produced by Reversion, where the measurement begins when the last unimpaired packet is received by the Tester on the Backup Path and ends when the first unimpaired packet is received by the Tester on the Protected Primary Path.

Discussion:

Packet loss can be observed as a reduction of forwarded traffic from the maximum forwarding rate. Reversion Packet Loss includes packets that were lost, reordered, or delayed. Reversion Packet Loss may reach 100% of the offered load.

Measurement Units:

Number of Packets

Issues:

None.

See Also:
Reversion

3.5.3. Failover Time

Definition:

The amount of time it takes for Failover to successfully complete.

Discussion:

Failover Time can be calculated using the Time-Based Loss Method (TBLM), Packet-Loss-Based Method (PLBM), or Timestamp-Based Method (TBM). It is RECOMMENDED that the TBM is used.

Measurement Units:
milliseconds

Issues:
None.

See Also:
Failover
Failover Time
Time-Based Loss Method (TBLM)
Packet-Loss-Based Method (PLBM)
Timestamp-Based Method (TBM)

3.5.4. Reversion Time

Definition:

The amount of time it takes for Reversion to complete so that the Primary Path is restored as the Working Path.

Discussion:

Reversion Time can be calculated using the Time-Based Loss Method (TBLM), Packet-Loss-Based Method (PLBM), or Timestamp-Based Method (TBM). It is RECOMMENDED that the TBM is used.

Measurement Units:
milliseconds

Issues:
None.

See Also:
Reversion
Primary Path
Working Path
Reversion Packet Loss

Time-Based Loss Method (TBLM)
Packet-Loss-Based Method (PLBM)
Timestamp-Based Method (TBM)

3.5.5. Additive Backup Delay

Definition:

The amount of increased Forwarding Delay [4] resulting from data traffic traversing the Backup Path instead of the Primary Path.

Discussion:

Additive Backup Delay is calculated using Equation 1 as shown below:

(Equation 1)

Additive Backup Delay =
Forwarding Delay(Backup Path) -
Forwarding Delay(Primary Path)

Measurement Units:

milliseconds

Issues:

Additive Backup Latency may be a negative result. This is theoretically possible but could be indicative of a sub-optimum network configuration.

See Also:

Primary Path
Backup Path
Primary Path Latency
Backup Path Latency

3.6. Failover Time Calculation Methods

The following Methods may be assessed on a per-flow basis using at least 16 flows spread over the routing table (using more flows is better). Otherwise, the impact of a prefix-dependency in the implementation of a particular protection technology could be missed. However, the test designer must be aware of the number of packets per second sent to each prefix, as this establishes sampling of the path and the time resolution for measurement of Failover time on a per-flow basis.

3.6.1. Time-Based Loss Method (TBLM)

Definition:

The method to calculate Failover Time (or Reversion Time) using a time scale on the Tester to measure the interval of Failover Packet Loss.

Discussion:

The Tester must provide statistics that show the duration of failure on a time scale based on occurrence of packet loss on a time scale. This is indicated by the duration of non-zero packet loss. The TBLM includes failure detection time and time for data traffic to begin traversing the Backup Path. Failover Time and Reversion Time are calculated using the TBLM as shown in Equation 2:

(Equation 2)

(Equation 2a)

$$\text{TBLM Failover Time} = \text{Time(Failover)} - \text{Time(Failover Event)}$$

(Equation 2b)

$$\text{TBLM Reversion Time} = \text{Time(Reversion)} - \text{Time(Restoration)}$$

Where

Time(Failover) = Time on the tester at the receipt of the first unimpaired packet at egress node after the backup path became the working path

Time(Failover Event) = Time on the tester at the receipt of the last unimpaired packet at egress node on the primary path before failure

Measurement Units:

milliseconds

Issues:

None.

See Also:

Failover
Packet-Loss-Based Method

3.6.2. Packet-Loss-Based Method (PLBM)

Definition:

The method used to calculate Failover Time (or Reversion Time) from the amount of Failover Packet Loss.

Discussion:

PLBM includes failure detection time and time for data traffic to begin traversing the Backup Path. Failover Time can be calculated using PLBM from the amount of Failover Packet Loss as shown below in Equation 3. Note: If traffic is sent to more than 1 destination, PLBM gives the average loss over the measured destinations.

(Equation 3)

(Equation 3a)

PLBM Failover Time =
(Number of packets lost / Offered Load rate) * 1000)

(Equation 3b)

PLBM Restoration Time =
(Number of packets lost / Offered Load rate) * 1000)

Units are packets/(packets/second) = seconds

Measurement Units:
milliseconds

Issues:
None.

See Also:
Failover Time-Based Loss Method

3.6.3. Timestamp-Based Method (TBM)

Definition:

The method to calculate Failover Time (or Reversion Time) using a time scale to quantify the interval between unimpaired packets arriving in the test stream.

Discussion:

The purpose of this method is to quantify the duration of failure or reversion on a time scale based on the observation of unimpaired packets. The TBM is calculated from Equation 2 with the values obtained from the timestamp in the packet payload, rather than from the Tester clock (which are used with the TBLM).

Unimpaired packets are normal packets that are not lost, reordered, or duplicated. A reordered packet is defined in Section 3.3 of [7]. A duplicate packet is defined in Section 3.3.5 of [4]. Unimpaired packets may be detected by checking a

sequence number in the payload, where the sequence number equals the next expected number for an unimpaired packet. A sequence gap or sequence reversal indicates impaired packets.

For calculating Failover Time, the TBM includes failure detection time and time for data traffic to begin traversing the Backup Path. For calculating Reversion Time, the TBM includes Reversion Time and time for data traffic to begin traversing the Primary Path.

Measurement Units:
 milliseconds

Issues:
 None.

See Also:
 Failover
 Failover Time
 Reversion
 Reversion Time

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. References

5.1. Normative References

- [1] Bradner, S., "The Internet Standards Process -- Revision 3", BCP 9, RFC 2026, October 1996.
- [2] Bradner, S., "Benchmarking Terminology for Network Interconnection Devices", RFC 1242, July 1991.
- [3] Mandeville, R., "Benchmarking Terminology for LAN Switching Devices", RFC 2285, February 1998.
- [4] Poretsky, S., Perser, J., Erramilli, S., and S. Khurana, "Terminology for Benchmarking Network-layer Traffic Control Mechanisms", RFC 4689, October 2006.
- [5] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [6] Poretsky, S., Imhoff, B., and K. Michielsen, "Terminology for Benchmarking Link-State IGP Data Plane Route Convergence", RFC 6412, November 2011.
- [7] Morton, A., Ciavattone, L., Ramachandran, G., Shalunov, S., and J. Perser, "Packet Reordering Metrics", RFC 4737, November 2006.
- [8] Nadas, S., Ed., "Virtual Router Redundancy Protocol (VRRP) Version 3 for IPv4 and IPv6", RFC 5798, March 2010.

5.2. Informative References

- [9] Pan, P., Ed., Swallow, G., Ed., and A. Atlas, Ed., "Fast Reroute Extensions to RSVP-TE for LSP Tunnels", RFC 4090, May 2005.
- [10] Nichols, K., Blake, S., Baker, F., and D. Black, "Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers", RFC 2474, December 1998.

6. Acknowledgments

We would like thank the BMWG and particularly Al Morton and Curtis Villamizar for their reviews, comments, and contributions to this work.

Authors' Addresses

Scott Poretsky
Allot Communications
300 TradeCenter
Woburn, MA 01801
USA
Phone: + 1 508 309 2179
EMail: sporetsky@allot.com

Rajiv Papneja
Huawei Technologies
2330 Central Expressway
Santa Clara, CA 95050
USA
Phone: +1 571 926 8593
EMail: rajiv.papneja@huawei.com

Jay Karthik
Cisco Systems
300 Beaver Brook Road
Boxborough, MA 01719
USA
Phone: +1 978 936 0533
EMail: jkarthik@cisco.com

Samir Vapiwala
Cisco System
300 Beaver Brook Road
Boxborough, MA 01719
USA
Phone: +1 978 936 1484
EMail: svapiwal@cisco.com

