Internet Engineering Task Force (IETF)
Request for Comments: 5814
Category: Standards Track
ISSN: 2070-1721
March 2010

W. Sun, Ed.
SJTU
G. Zhang, Ed.
CATR

Label Switched Path (LSP) Dynamic Provisioning Performance Metrics in Generalized MPLS Networks

Abstract

Generalized Multi-Protocol Label Switching (GMPLS) is one of the most promising candidate technologies for a future data transmission network. GMPLS has been developed to control and operate different kinds of network elements, such as conventional routers, switches, Dense Wavelength Division Multiplexing (DWDM) systems, Add-Drop Multiplexers (ADMs), photonic cross-connects (PXCs), optical cross-connects (OXCs), etc. These physically diverse devices differ drastically from one another in dynamic provisioning ability. At the same time, the need for dynamically provisioned connections is increasing because optical networks are being deployed in metro areas. As different applications have varied requirements in the provisioning performance of optical networks, it is imperative to define standardized metrics and procedures such that the performance of networks and application needs can be mapped to each other.

This document provides a series of performance metrics to evaluate the dynamic Label Switched Path (LSP) provisioning performance in GMPLS networks, specifically the dynamic LSP setup/release performance. These metrics can be used to characterize the features of GMPLS networks in LSP dynamic provisioning.

Status of This Memo

Table of Contents

1. Introduction

Generalized Multi-Protocol Label Switching (GMPLS) is one of the most promising control plane solutions for future transport and service network. GMPLS has been developed to control and operate different kinds of network elements, such as conventional routers, switches, Dense Wavelength Division Multiplexing (DWDM) systems, Add-Drop Multiplexors (ADMs), photonic cross-connects (PXCs), optical cross-connects (OXCs), etc. These physically diverse devices differ drastically from one another in dynamic provisioning ability.

The introduction of a control plane into optical circuit switching networks provides the basis for automating the provisioning of connections and drastically reduces connection provision delay. As more and more services and applications are seeking to use GMPLS-controlled networks as their underlying transport network, and increasingly in a dynamic way, the need is growing for measuring and characterizing the performance of LSP provisioning in GMPLS networks, such that requirement from applications and the provisioning capability of the network can be mapped to each other.

This document defines performance metrics and methodologies that can be used to characterize the dynamic LSP provisioning performance of GMPLS networks, more specifically, performance of the signaling protocol. The metrics defined in this document can be used to characterize the average performance of GMPLS implementations.

2. Conventions Used in This Document

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 [RFC2119].

3. Overview of Performance Metrics

In this memo, to characterize the dynamic LSP provisioning performance of a GMPLS network, we define three performance metrics: unidirectional LSP setup delay, bidirectional LSP setup delay, and LSP graceful release delay. The latency of the LSP setup/release signal is conceptually similar to the Round-trip Delay in IP networks. This enables us to refer to the structures and notions introduced and discussed in the IP Performance Metrics (IPPM) Framework documents, [RFC2330] [RFC2679] [RFC2681]. The reader is assumed to be familiar with the notions in those documents.

Note that data-path-related metrics, for example, the time between the reception of a Resv message by the ingress node and when the forward data path becomes operational, are defined in another document [CCAMP-DPM]. It is desirable that both measurements are performed to complement each other.

4. A Singleton Definition for Single Unidirectional LSP Setup Delay

This section defines a metric for single unidirectional Label Switched Path setup delay across a GMPLS network.

4.1. Motivation

Single unidirectional Label Switched Path setup delay is useful for several reasons:

- o Single LSP setup delay is an important metric that characterizes the provisioning performance of a GMPLS network. Longer LSP setup delay will most likely incur higher overhead for the requesting application, especially when the LSP duration itself is comparable to the LSP setup delay.

- o The minimum value of this metric provides an indication of the delay that will likely be experienced when the LSP traverses the shortest route at the lightest load in the control plane. As the delay itself consists of several components, such as link propagation delay and nodal processing delay, this metric also reflects the status of the control plane. For example, for LSPs traversing the same route, longer setup delays may suggest congestion in the control channel or high control element load. For this reason, this metric is useful for testing and diagnostic purposes.

- o The observed variance in a sample of LSP setup delay metric values variance may serve as an early indicator on the feasibility of support of applications that have stringent setup delay requirements.

The measurement of single unidirectional LSP setup delay instead of bidirectional LSP setup delay is motivated by the following factors:

- o Some applications may use only unidirectional LSPs rather than bidirectional ones. For example, content delivery services with multicasting may use only unidirectional LSPs.

4.2. Metric Name

Single unidirectional LSP setup delay

4.3. Metric Parameters

- o ID0, the ingress Label Switching Router (LSR) ID
- o ID1, the egress LSR ID
- o T, a time when the setup is attempted

4.4. Metric Units

The value of single unidirectional LSP setup delay is either a real number of milliseconds or undefined.

4.5. Definition

The single unidirectional LSP setup delay from ingress node ID0 to egress node ID1 [RFC3945] at T is dT means that ingress node ID0 sends the first bit of a Path message packet to egress node ID1 at wire-time T, and that ingress node ID0 received the last bit of responding Resv message packet from egress node ID1 at wire-time T+dT.

The single unidirectional LSP setup delay from ingress node ID0 to egress node ID1 at T is undefined means that ingress node ID0 sends the first bit of Path message packet to egress node ID1 at wire-time T and that ingress node ID0 does not receive the corresponding Resv message within a reasonable period of time.

The undefined value of this metric indicates an event of Single Unidirectional LSP Setup Failure and would be used to report a count or a percentage of Single Unidirectional LSP Setup failures. See Section 14.5 for definitions of LSP setup/release failures.

4.6. Discussion

The following issues are likely to come up in practice:

- o The accuracy of unidirectional LSP setup delay at time T depends on the clock resolution in the ingress node; but synchronization between the ingress node and egress node is not required since unidirectional LSP setup uses two-way signaling.
- o A given methodology will have to include a way to determine whether a latency value is infinite or whether it is merely very large. Simple upper bounds MAY be used, but GMPLS networks may accommodate many kinds of devices. For example, some photonic cross-connects (PXCs) have to move micro mirrors. This physical motion may take several milliseconds, but the common electronic

switches can finish the nodal processing within several microseconds. So the unidirectional LSP setup delay varies drastically from one network to another. In practice, the upper bound SHOULD be chosen carefully.

- If the ingress node sends out the Path message to set up an LSP, but never receives the corresponding Resv message, the unidirectional LSP setup delay MUST be set to undefined.

- If the ingress node sends out the Path message to set up an LSP but receives a PathErr message, the unidirectional LSP setup delay MUST be set to undefined. There are many possible reasons for this case; for example, the Path message has invalid parameters or the network does not have enough resources to set up the requested LSP, etc.

4.7. Methodologies

Generally, the methodology would proceed as follows:

- Make sure that the network has enough resources to set up the requested LSP.

- At the ingress node, form the Path message according to the LSP requirements. A timestamp (T1) may be stored locally on the ingress node when the Path message packet is sent towards the egress node.

- If the corresponding Resv message arrives within a reasonable period of time, take the timestamp (T2) as soon as possible upon receipt of the message. By subtracting the two timestamps, an estimate of unidirectional LSP setup delay (T2-T1) can be computed.

- If the corresponding Resv message fails to arrive within a reasonable period of time, the unidirectional LSP setup delay is deemed to be undefined. Note that the "reasonable" threshold is a parameter of the methodology.

- If the corresponding response is a PathErr message, the unidirectional LSP setup delay is deemed to be undefined.

4.8. Metric Reporting

The metric result (either a real number or undefined) MUST be reported together with the selected upper bound. The route that the LSP traverses MUST also be reported. The route MAY be collected via use of the record route object, see [RFC3209], or via the management plane. The collection of routes via the management plane is out of scope of this document.

5. A Singleton Definition for Multiple Unidirectional LSPs Setup Delay

This section defines a metric for multiple unidirectional Label Switched Paths setup delay across a GMPLS network.

5.1. Motivation

Multiple unidirectional Label Switched Paths setup delay is useful for several reasons:

- o Carriers may require that a large number of LSPs be set up during a short time period. This request may arise, e.g., as a consequence to interruptions on established LSPs or other network failures.

- o The time needed to set up a large number of LSPs during a short time period cannot be deduced from single LSP setup delay.

5.2. Metric Name

Multiple unidirectional LSPs setup delay

5.3. Metric Parameters

- o ID0, the ingress LSR ID

- o ID1, the egress LSR ID

- o Lambda_m, a rate in reciprocal milliseconds

- o X, the number of LSPs to set up

- o T, a time when the first setup is attempted

5.4. Metric Units

The value of multiple unidirectional LSPs setup delay is either a real number of milliseconds or undefined

5.5. Definition

Given Lambda_m and X, the multiple unidirectional LSPs setup delay from the ingress node to the egress node [RFC3945] at T is dT means:

- o ingress node ID0 sends the first bit of the first Path message packet to egress node ID1 at wire-time T;
- o all subsequent (X-1) Path messages are sent according to the specified Poisson process with arrival rate Lambda_m;
- o ingress node ID0 receives all corresponding Resv message packets from egress node ID1; and
- o ingress node ID0 receives the last Resv message packet at wire-time T+dT.

If the multiple unidirectional LSPs setup delay at T is "undefined", this means that:

- o ingress node ID0 sends all the Path messages toward egress node ID1,
- o the first bit of the first Path message packet is sent at wire-time T, and
- o ingress node ID0 does not receive one or more of the corresponding Resv messages within a reasonable period of time.

The undefined value of this metric indicates an event of Multiple Unidirectional LSP Setup Failure and would be used to report a count or a percentage of Multiple Unidirectional LSP Setup failures. See Section 14.5 for definitions of LSP setup/release failures.

5.6. Discussion

The following issues are likely to come up in practice:

- o The accuracy of multiple unidirectional LSPs setup delay at time T depends on the clock resolution in the ingress node; but synchronization between the ingress node and egress node is not required since unidirectional LSP setup uses two-way signaling.
- o A given methodology will have to include a way to determine whether a latency value is infinite or whether it is merely very large. Simple upper bounds MAY be used, but GMPLS networks may accommodate many kinds of devices. For example, some photonic cross-connects (PXCs) have to move micro mirrors. This physical

motion may take several milliseconds, but electronic switches can finish the nodal processing within several microseconds. So the multiple unidirectional LSP setup delay varies drastically from one network to another. In practice, the upper bound SHOULD be chosen carefully.

- o If the ingress node sends out the multiple Path messages to set up the LSPs, but never receives one or more of the corresponding Resv messages, multiple unidirectional LSP setup delay MUST be set to undefined.

- o If the ingress node sends out the Path messages to set up the LSPs but receives one or more PathErr messages, multiple unidirectional LSPs setup delay MUST be set to undefined. There are many possible reasons for this case. For example, one of the Path messages has invalid parameters or the network does not have enough resources to set up the requested LSPs, etc.

- o The arrival rate of the Poisson process Lambda_m SHOULD be chosen carefully such that on the one hand, the control plane is not overburdened. On the other hand, the arrival rate is large enough to meet the requirements of applications or services.

- o It is important that all the LSPs MUST traverse the same route. If there are multiple routes between the ingress node ID0 and egress node ID1, Explicit Route Objects (EROs), or an alternate method, e.g., static configuration, MUST be used to ensure that all LSPs traverse the same route.

5.7. Methodologies

Generally, the methodology would proceed as follows:

- o Make sure that the network has enough resources to set up the requested LSPs.

- o At the ingress node, form the Path messages according to the LSPs' requirements.

- o At the ingress node, select the time for each of the Path messages according to the specified Poisson process.

- o At the ingress node, send out the Path messages according to the selected time.

- o Store a timestamp (T1) locally on the ingress node when the first Path message packet is sent towards the egress node.

- o If all of the corresponding Resv messages arrive within a reasonable period of time, take the final timestamp (T2) as soon as possible upon the receipt of all the messages. By subtracting the two timestamps, an estimate of multiple unidirectional LSPs setup delay (T2-T1) can be computed.

- o If one or more of the corresponding Resv messages fail to arrive within a reasonable period of time, the multiple unidirectional LSPs setup delay is deemed to be undefined. Note that the "reasonable" threshold is a parameter of the methodology.

- o If one or more of the corresponding responses are PathErr messages, the multiple unidirectional LSPs setup delay is deemed to be undefined.

5.8. Metric Reporting

The metric result (either a real number or undefined) MUST be reported together with the selected upper bound. The route that the LSPs traverse MUST also be reported. The route MAY be collected via use of the record route object, see [RFC3209], or via the management plane. The collection of routes via the management plane is out of scope of this document.

6. A Singleton Definition for Single Bidirectional LSP Setup Delay

GMPLS allows establishment of bidirectional symmetric LSPs (not of asymmetric LSPs). This section defines a metric for single bidirectional LSP setup delay across a GMPLS network.

6.1. Motivation

Single bidirectional Label Switched Path setup delay is useful for several reasons:

- o LSP setup delay is an important metric that characterizes the provisioning performance of a GMPLS network. Longer LSP setup delay will incur higher overhead for the requesting application, especially when the LSP duration is comparable to the LSP setup delay. Thus, measuring the setup delay is important for application scheduling.

- o The minimum value of this metric provides an indication of the delay that will likely be experienced when the LSP traverses the shortest route at the lightest load in the control plane. As the delay itself consists of several components, such as link propagation delay and nodal processing delay, this metric also reflects the status of the control plane. For example, for LSPs traversing the same route, longer setup delays may suggest congestion in the control channel or high control element load. For this reason, this metric is useful for testing and diagnostic purposes.

- o LSP setup delay variance has a different impact on applications. Erratic variation in LSP setup delay makes it difficult to support applications that have stringent setup delay requirement.

The measurement of single bidirectional LSP setup delay instead of unidirectional LSP setup delay is motivated by the following factors:

- o Bidirectional LSPs are seen as a requirement for many GMPLS networks. Its provisioning performance is important to applications that generate bidirectional traffic.

6.2. Metric Name

Single bidirectional LSP setup delay

6.3. Metric Parameters

- o ID0, the ingress LSR ID
- o ID1, the egress LSR ID
- o T, a time when the setup is attempted

6.4. Metric Units

The value of single bidirectional LSP setup delay is either a real number of milliseconds or undefined

6.5. Definition

For a real number dT, the single bidirectional LSP setup delay from ingress node ID0 to egress node ID1 at T is dT means that ingress node ID0 sends out the first bit of a Path message including an Upstream Label [RFC3473] heading for egress node ID1 at wire-time T, egress node ID1 receives that packet, then immediately sends a Resv message packet back to ingress node ID0, and that ingress node ID0 receives the last bit of the Resv message packet at wire-time T+dT.

If the single bidirectional LSP setup delay from ingress node ID0 to egress node ID1 at T is "undefined", this means that ingress node ID0 sends the first bit of a Path message to egress node ID1 at wire-time T and that ingress node ID0 does not receive that response packet within a reasonable period of time.

The undefined value of this metric indicates an event of Single Bidirectional LSP Setup Failure and would be used to report a count or a percentage of Single Bidirectional LSP Setup failures. See Section 14.5 for definitions of LSP setup/release failures.

6.6. Discussion

The following issues are likely to come up in practice:

- The accuracy of single bidirectional LSP setup delay depends on the clock resolution in the ingress node; but synchronization between the ingress node and egress node is not required since single bidirectional LSP setup uses two-way signaling.

- A given methodology will have to include a way to determine whether a latency value is infinite or whether it is merely very large. Simple upper bounds MAY be used, but GMPLS networks may accommodate many kinds of devices. For example, some photonic cross-connects (PXCs) have to move micro mirrors. This physical motion may take several milliseconds, but electronic switches can finish the nodal processing within several microseconds. So the bidirectional LSP setup delay varies drastically from one network to another. In the process of bidirectional LSP setup, if the downstream node overrides the label suggested by the upstream node, the setup delay may also increase. Thus, in practice, the upper bound SHOULD be chosen carefully.

- If the ingress node sends out the Path message to set up the LSP, but never receives the corresponding Resv message, single bidirectional LSP setup delay MUST be set to undefined.

- If the ingress node sends out the Path message to set up the LSP, but receives a PathErr message, single bidirectional LSP setup delay MUST be set to undefined. There are many possible reasons for this case. For example, the Path message has invalid parameters or the network does not have enough resources to set up the requested LSP.

6.7. Methodologies

Generally, the methodology would proceed as follows:

- Make sure that the network has enough resources to set up the requested LSP.

- o At the ingress node, form the Path message (including the Upstream Label or suggested label) according to the LSP requirements. A timestamp (T1) may be stored locally on the ingress node when the Path message packet is sent towards the egress node.

- o If the corresponding Resv message arrives within a reasonable period of time, take the final timestamp (T2) as soon as possible upon the receipt of the message. By subtracting the two timestamps, an estimate of bidirectional LSP setup delay (T2-T1) can be computed.

- o If the corresponding Resv message fails to arrive within a reasonable period of time, the single bidirectional LSP setup delay is deemed to be undefined. Note that the "reasonable" threshold is a parameter of the methodology.

- o If the corresponding response is a PathErr message, the single bidirectional LSP setup delay is deemed to be undefined.

6.8. Metric Reporting

The metric result (either a real number or undefined) MUST be reported together with the selected upper bound. The route that the LSP traverses MUST also be reported. The route MAY be collected via use of the record route object, see [RFC3209], or via the management plane. The collection of routes via the management plane is out of scope of this document.

7. A Singleton Definition for Multiple Bidirectional LSPs Setup Delay

This section defines a metric for multiple bidirectional LSPs setup delay across a GMPLS network.

7.1. Motivation

Multiple bidirectional LSPs setup delay is useful for several reasons:

- o Upon traffic interruption caused by network failure or network upgrade, carriers may require a large number of LSPs be set up during a short time period.

- o The time needed to set up a large number of LSPs during a short time period cannot be deduced by single LSP setup delay.

7.2. Metric Name

Multiple bidirectional LSPs setup delay

7.3. Metric Parameters

- o ID0, the ingress LSR ID
- o ID1, the egress LSR ID
- o Lambda_m, a rate in reciprocal milliseconds
- o X, the number of LSPs to set up
- o T, a time when the first setup is attempted

7.4. Metric Units

The value of multiple bidirectional LSPs setup delay is either a real number of milliseconds or undefined

7.5. Definition

Given Lambda_m and X, for a real number dT, the multiple bidirectional LSPs setup delay from ingress node to egress node at T is dT, means that:

- o Ingress node ID0 sends the first bit of the first Path message heading for egress node ID1 at wire-time T;
- o All subsequent (X-1) Path messages are sent according to the specified Poisson process with arrival rate Lambda_m;
- o Ingress node ID1 receives all corresponding Resv message packets from egress node ID1; and
- o Ingress node ID0 receives the last Resv message packet at wire-time T+dT.

If the multiple bidirectional LSPs setup delay from ingress node to egress node at T is "undefined", this means that the ingress node sends all the Path messages to the egress node and that the ingress node fails to receive one or more of the response Resv messages within a reasonable period of time.

The undefined value of this metric indicates an event of Multiple Bidirectional LSP Setup Failure and would be used to report a count or a percentage of Multiple Bidirectional LSP Setup failures. See Section 14.5 for definitions of LSP setup/release failures.

7.6. Discussion

The following issues are likely to come up in practice:

- o The accuracy of multiple bidirectional LSPs setup delay depends on the clock resolution in the ingress node; but synchronization between the ingress node and egress node is not required since bidirectional LSP setup uses two-way signaling.

- o A given methodology will have to include a way to determine whether a latency value is infinite or whether it is merely very large. Simple upper bounds MAY be used, but GMPLS networks may accommodate many kinds of devices. For example, some photonic cross-connects (PXCs) have to move micro mirrors. This physical motion may take several milliseconds, but electronic switches can finish the nodal process within several microseconds. So the multiple bidirectional LSPs setup delay varies drastically from a network to another. In the process of multiple bidirectional LSPs setup, if the downstream node overrides the label suggested by the upstream node, the setup delay may also increase. Thus, in practice, the upper bound SHOULD be chosen carefully.

- o If the ingress node sends out the Path messages to set up the LSPs, but never receives all the corresponding Resv messages, the multiple bidirectional LSPs setup delay MUST be set to undefined.

- o If the ingress node sends out the Path messages to set up the LSPs, but receives one or more responding PathErr messages, the multiple bidirectional LSPs setup delay MUST be set to undefined. There are many possible reasons for this case. For example, one or more of the Path messages have invalid parameters or the network does not have enough resources to set up the requested LSPs.

- o The arrival rate of the Poisson process Lambda_m SHOULD be carefully chosen such that on the one hand, the control plane is not overburdened. On the other hand, the arrival rate is large enough to meet the requirements of applications or services.

- o It is important that all the LSPs MUST traverse the same route. If there are multiple routes between the ingress node ID0 and egress node ID1, EROs, or an alternate method, e.g., static configuration, MUST be used to ensure that all LSPs traverse the same route.

7.7. Methodologies

Generally, the methodology would proceed as follows:

- Make sure that the network has enough resources to set up the requested LSPs.

- At the ingress node, form the Path messages (including the Upstream Label or suggested label) according to the LSPs' requirements.

- At the ingress node, select the time for each of the Path messages according to the specified Poisson process.

- At the ingress node, send out the Path messages according to the selected time.

- Store a timestamp (T1) locally in the ingress node when the first Path message packet is sent towards the egress node.

- If all of the corresponding Resv messages arrive within a reasonable period of time, take the final timestamp (T2) as soon as possible upon the receipt of all the messages. By subtracting the two timestamps, an estimate of multiple bidirectional LSPs setup delay (T2-T1) can be computed.

- If one or more of the corresponding Resv messages fail to arrive within a reasonable period of time, the multiple bidirectional LSPs setup delay is deemed to be undefined. Note that the "reasonable" threshold is a parameter of the methodology.

- If one or more of the corresponding responses are PathErr messages, the multiple bidirectional LSPs setup delay is deemed to be undefined.

7.8. Metric Reporting

The metric result (either a real number or undefined) MUST be reported together with the selected upper bound. The route that the LSPs traverse MUST also be reported. The route MAY be collected via use of the record route object, see [RFC3209], or via the management plane. The collection of routes via the management plane is out of scope of this document.

8. A Singleton Definition for LSP Graceful Release Delay

There are two different kinds of LSP release mechanisms in GMPLS networks: graceful release and forceful release. This document does not take forceful LSP release procedure into account.

8.1. Motivation

LSP graceful release delay is useful for several reasons:

- o The LSP graceful release delay is part of the total cost of dynamic LSP provisioning. For some short duration applications, the LSP release time cannot be ignored.

- o The LSP graceful release procedure is more preferred in a GMPLS controlled network, particularly the optical networks. Since it doesn’t trigger restoration/protection, it is "alarm-free connection deletion" in [RFC4208].

8.2. Metric Name

LSP graceful release delay

8.3. Metric Parameters

- o ID0, the ingress LSR ID

- o ID1, the egress LSR ID

- o T, a time when the release is attempted

8.4. Metric Units

The value of LSP graceful release delay is either a real number of milliseconds or undefined

8.5. Definition

There are two different LSP graceful release procedures: one is initiated by the ingress node, and another is initiated by the egress node. The two procedures are depicted in [RFC3473]. We define the graceful LSP release delay for these two procedures separately.

For a real number dT, if the LSP graceful release delay from ingress node ID0 to egress node ID1 at T is dT, this means that ingress node ID0 sends the first bit of a Path message including an Admin Status Object with the Reflect (R) and Delete (D) bits set to the egress node at wire-time T, that egress node ID1 receives that packet, then immediately sends a Resv message including an Admin Status Object with the Delete (D) bit set back to the ingress node. Ingress node ID0 sends the PathTear message downstream to remove the LSP, and egress node ID1 receives the last bit of PathTear packet at wire-time T+dT.

Also, as an option, upon receipt of the Path message including an Admin Status Object with the Reflect (R) and Delete (D) bits set, egress node ID1 may respond with a PathErr message with the Path_State_Removed flag set.

The LSP graceful release delay from ingress node ID0 to egress node ID1 at T is undefined means that ingress node ID0 sends the first bit of Path message to egress node ID1 at wire-time T and that (either the egress node does not receive the Path packet, the egress node does not send a corresponding Resv message packet in response, or the ingress node does not receive that Resv packet, and) egress node ID1 does not receive the PathTear message within a reasonable period of time.

If the LSP graceful release delay from egress node ID1 to ingress node ID0 at T is dT, this means that egress node ID1 sends the first bit of a Resv message including an Admin Status Object with the Reflect (R) and Delete (D) bits set to the ingress node at wire-time T. Ingress node ID0 sends a PathTear message downstream to remove the LSP, and egress node ID1 receives the last bit of PathTear packet at wire-time T+dT.

If the LSP graceful release delay from egress node ID1 to ingress node ID0 at T is "undefined", this means that egress node ID1 sends the first bit of Resv message including an Admin Status Object with the Reflect (R) and Delete (D) bits set to the ingress node ID0 at wire-time T and that (either the ingress node does not receive the Resv packet or the ingress node does not send PathTear message packet in response, and) egress node ID1 does not receive the PathTear message within a reasonable period of time.

The undefined value of this metric indicates an event of LSP Graceful Release Failure and would be used to report a count or a percentage of LSP Graceful Release failures. See Section 14.5 for definitions of LSP setup/release failures.

8.6. Discussion

The following issues are likely to come up in practice:

- In the first (second) circumstance, the accuracy of LSP graceful release delay at time T depends on the clock resolution in the ingress (egress) node. In the first circumstance, synchronization between the ingress node and egress node is required, but it is not in the second circumstance.

- A given methodology has to include a way to determine whether a latency value is infinite or whether it is merely very large. Simple upper bounds MAY be used, but the upper bound SHOULD be chosen carefully in practice.

- In the first circumstance, if the ingress node sends out Path message including an Admin Status Object with the Reflect (R) and Delete (D) bits set to initiate LSP graceful release, but the egress node never receives the corresponding PathTear message, LSP graceful release delay MUST be set to undefined.

- In the second circumstance, if the egress node sends out the Resv message including an Admin Status Object with the Reflect (R) and Delete (D) bits set to initiate LSP graceful release, but never receives the corresponding PathTear message, LSP graceful release delay MUST be set to undefined.

8.7. Methodologies

In the first circumstance, the methodology may proceed as follows:

- Make sure the LSP to be deleted is set up;

- At the ingress node, form the Path message including an Admin Status Object with the Reflect (R) and Delete (D) bits set. A timestamp (T1) may be stored locally on the ingress node when the Path message packet is sent towards the egress node.

- Upon receiving the Path message including an Admin Status Object with the Reflect (R) and Delete (D) bits set, the egress node sends a Resv message including an Admin Status Object with the Delete (D) and Reflect (R) bits set. Alternatively, the egress node sends a PathErr message with the Path_State_Removed flag set upstream.

- When the ingress node receives the Resv message or the PathErr message, it sends a PathTear message to remove the LSP.

- o The egress node takes a timestamp (T2) once it receives the last bit of the PathTear message. The LSP graceful release delay is then (T2-T1).

- o If the ingress node sends the Path message downstream, but the egress node fails to receive the PathTear message within a reasonable period of time, the LSP graceful release delay is deemed to be undefined. Note that the "reasonable" threshold is a parameter of the methodology.

In the second circumstance, the methodology would proceed as follows:

- o Make sure the LSP to be deleted is set up;

- o On the egress node, form the Resv message including an Admin Status Object with the Reflect (R) and Delete (D) bits set. A timestamp may be stored locally on the egress node when the Resv message packet is sent towards the ingress node.

- o Upon receiving the Admin Status Object with the Reflect (R) and Delete (D) bits set in the Resv message, the ingress node sends a PathTear message downstream to remove the LSP.

- o The egress node takes a timestamp (T2) once it receives the last bit of the PathTear message. The LSP graceful release delay is then (T2-T1).

- o If the egress node sends the Resv message upstream, but it fails to receive the PathTear message within a reasonable period of time, the LSP graceful release delay is deemed to be undefined. Note that the "reasonable" threshold is a parameter of the methodology.

8.8. Metric Reporting

The metric result (either a real number or undefined) MUST be reported together with the selected upper bound and the procedure used (e.g., either from the ingress node to the egress node or from the egress node to the ingress node; see Section 8.5 for more details). The route that the LSP traverses MUST also be reported. The route MAY be collected via use of the record route object, see [RFC3209], or via the management plane. The collection of routes via the management plane is out of scope of this document.

9. A Definition for Samples of Single Unidirectional LSP Setup Delay

In Section 4, we defined the singleton metric of single unidirectional LSP setup delay. Now we define how to get one particular sample of single unidirectional LSP setup delay. Sampling means to take a number of distinct instances of a skeleton metric under a given set of parameters. As in [RFC2330], we use Poisson sampling as an example.

9.1. Metric Name

Single unidirectional LSP setup delay sample

9.2. Metric Parameters

- o ID0, the ingress LSR ID
- o ID1, the egress LSR ID
- o T0, a time
- o Tf, a time
- o Lambda, a rate in the reciprocal milliseconds
- o Th, LSP holding time
- o Td, the maximum waiting time for successful setup

9.3. Metric Units

A sequence of pairs; the elements of each pair are:

- o T, a time when setup is attempted
- o dT, either a real number of milliseconds or undefined

9.4. Definition

Given T0, Tf, and Lambda, compute a pseudo-random Poisson process beginning at or before T0, with average arrival rate Lambda, and ending at or after Tf. Those time values greater than or equal to T0 and less than or equal to Tf are then selected. At each of the times in this process, we obtain the value of unidirectional LSP setup delay sample. The value of the sample is the sequence made up of the resulting <time, LSP setup delay> pairs. If there are no such pairs, the sequence is of length zero and the sample is said to be empty.

9.5. Discussion

The parameter Lambda should be carefully chosen. If the rate is too high, too frequent LSP setup/release procedure will result in high overhead in the control plane. In turn, the high overhead will increase unidirectional LSP setup delay. On the other hand, if the rate is too low, the sample might not completely reflect the dynamic provisioning performance of the GMPLS network. The appropriate Lambda value depends on the given network.

The parameters Td should be carefully chosen. Different switching technologies may vary significantly in performing a cross-connect operation. At the same time, the time needed in setting up an LSP under different traffic may also vary significantly.

In the case of active measurement, the parameters Th should be carefully chosen. The combination of Lambda and Th reflects the load of the network. The selection of Th should take into account that the network has sufficient resources to perform subsequent tests. The value of Th MAY be constant during one sampling process for simplicity considerations.

Note that for online or passive measurements, the arrival rate and LSP holding time are determined by actual traffic; hence, in this case, Lambda and Th are not input parameters.

It is important that, in obtaining a sample, all the LSPs MUST traverse the same route. If there are multiple routes between the ingress node ID0 and egress node ID1, EROs, or an alternate method, e.g., static configuration, MUST be used to ensure that all LSPs traverse the same route.

9.6. Methodologies

- o Select the times using the specified Poisson arrival process,
- o Set up the LSP as the methodology for the singleton unidirectional LSP setup delay, and obtain the value of unidirectional LSP setup delay, and
- o Release the LSP after Th, and wait for the next Poisson arrival event.

Note: it is possible that before the previous LSP release procedure completes, the next Poisson arrival event arrives and the LSP setup procedure is initiated. If there is resource contention between the two LSPs, the LSP setup may fail. Ways to avoid such contention are outside the scope of this document.

9.7. Typical Testing Cases

9.7.1. With No LSP in the Network

9.7.1.1. Motivation

Single unidirectional LSP setup delay with no LSP in the network is important because this reflects the inherent delay of a Resource Reservation Protocol - Traffic Engineering (RSVP-TE) implementation. The minimum value provides an indication of the delay that will likely be experienced when an LSP traverses the shortest route with the lightest load in the control plane.

9.7.1.2. Methodologies

Make sure that there is no LSP in the network and proceed with the methodologies described in Section 9.6

9.7.2. With a Number of LSPs in the Network

9.7.2.1. Motivation

Single unidirectional LSP setup delay with a number of LSPs in the network is important because it reflects the performance of an operational network with considerable load. This delay may vary significantly as the number of existing LSPs vary. It can be used as a scalability metric of an RSVP-TE implementation.

9.7.2.2. Methodologies

Set up the required number of LSPs, and wait until the network reaches a stable state; then, proceed with the methodologies described in Section 9.6.

9.8. Metric Reporting

The metric results including both real and undefined values MUST be reported together with the total number of values. The context under which the sample is obtained, including the selected parameters, the route traversed by the LSPs, and the testing case used, MUST also be reported.

10. A Definition for Samples of Multiple Unidirectional LSPs Setup Delay

In Section 5, we defined the singleton metric of multiple unidirectional LSPs setup delay. Now we define how to get one particular sample of multiple unidirectional LSPs setup delay.

Sampling means to take a number of distinct instances of a skeleton metric under a given set of parameters. As in [RFC2330], we use Poisson sampling as an example.

10.1. Metric Name

Multiple unidirectional LSPs setup delay sample

10.2. Metric Parameters

- ID0, the ingress LSR ID
- ID1, the egress LSR ID
- T0, a time
- Tf, a time
- Lambda_m, a rate in the reciprocal milliseconds
- Lambda, a rate in the reciprocal milliseconds
- X, the number of LSPs to set up
- Th, LSP holding time
- Td, the maximum waiting time for successful multiple unidirectional LSPs setup

10.3. Metric Units

A sequence of pairs; the elements of each pair are:

- T, a time when the first setup is attempted
- dT, either a real number of milliseconds or undefined

10.4. Definition

Given T0, Tf, and Lambda, compute a pseudo-random Poisson process beginning at or before T0, with an average arrival rate Lambda and ending at or after Tf. Those time values greater than or equal to T0 and less than or equal to Tf are then selected. At each of the times in this process, we obtain the value of multiple unidirectional LSP setup delay sample. The value of the sample is the sequence made up of the resulting <time, setup delay> pairs. If there are no such pairs, the sequence is of length zero and the sample is said to be empty.

10.5. Discussion

The parameter Lambda is used as an arrival rate of "batch unidirectional LSPs setup" operation. It regulates the interval in between each batch operation. The parameter Lambda_m is used within each batch operation, as described in Section 5

The parameters Lambda and Lambda_m should be carefully chosen. If the rate is too high, overly frequent LSP setup/release procedure will result in high overhead in the control plane. In turn, the high overhead will increase unidirectional LSP setup delay. On the other hand, if the rate is too low, the sample might not completely reflect the dynamic provisioning performance of the GMPLS network. The appropriate Lambda and Lambda_m value depends on the given network.

The parameters Td should be carefully chosen. Different switching technologies may vary significantly in performing a cross-connect operation. At the same time, the time needed in setting up an LSP under different traffic may also vary significantly.

It is important that, in obtaining a sample, all the LSPs MUST traverse the same route. If there are multiple routes between the ingress node ID0 and egress node ID1, EROs, or an alternate method, e.g., static configuration, MUST be used to ensure that all LSPs traverse the same route.

10.6. Methodologies

- o Select the times using the specified Poisson arrival process,
- o Set up the LSP as the methodology for the singleton multiple unidirectional LSPs setup delay, and obtain the value of multiple unidirectional LSPs setup delay, and
- o Release the LSP after Th, and wait for the next Poisson arrival event.

Note: it is possible that before the previous LSP release procedure completes, the next Poisson arrival event arrives and the LSP setup procedure is initiated. If there is resource contention between the two LSPs, the LSP setup may fail. Ways to avoid such contention are outside the scope of this document.

10.7. Typical Testing Cases

10.7.1. With No LSP in the Network

10.7.1.1. Motivation

Multiple unidirectional LSPs setup delay with no LSP in the network is important because this reflects the inherent delay of an RSVP-TE implementation. The minimum value provides an indication of the delay that will likely be experienced when LSPs traverse the shortest route with the lightest load in the control plane.

10.7.1.2. Methodologies

Make sure that there is no LSP in the network and proceed with the methodologies described in Section 10.6.

10.7.2. With a Number of LSPs in the Network

10.7.2.1. Motivation

Multiple unidirectional LSPs setup delay with a number of LSPs in the network is important because it reflects the performance of an operational network with considerable load. This delay can vary significantly as the number of existing LSPs vary. It can be used as a scalability metric of an RSVP-TE implementation.

10.7.2.2. Methodologies

Set up the required number of LSPs, and wait until the network reaches a stable state; then, proceed with the methodologies described in Section 10.6.

10.8. Metric Reporting

The metric results including both real and undefined values MUST be reported together with the total number of values. The context under which the sample is obtained, including the selected parameters, the route traversed by the LSPs, and the testing case used, MUST also be reported.

11. A Definition for Samples of Single Bidirectional LSP Setup Delay

In Section 6, we defined the singleton metric of single bidirectional LSP setup delay. Now we define how to get one particular sample of single bidirectional LSP setup delay. Sampling means to take a number of distinct instances of a skeleton metric under a given set of parameters. As in [RFC2330], we use Poisson sampling as an example.

11.1. Metric Name

Single bidirectional LSP setup delay sample with no LSP in the network

11.2. Metric Parameters

- o ID0, the ingress LSR ID
- o ID1, the egress LSR ID
- o T0, a time
- o Tf, a time
- o Lambda, a rate in the reciprocal milliseconds
- o Th, LSP holding time
- o Td, the maximum waiting time for successful setup

11.3. Metric Units

A sequence of pairs; the elements of each pair are:

- o T, a time when setup is attempted
- o dT, either a real number of milliseconds or undefined

11.4. Definition

Given T0, Tf, and Lambda, compute a pseudo-random Poisson process beginning at or before T0, with an average arrival rate Lambda, and ending at or after Tf. Those time values greater than or equal to T0 and less than or equal to Tf are then selected. At each of the times in this process, we obtain the value of bidirectional LSP setup delay sample. The value of the sample is the sequence made up of the resulting <time, LSP setup delay> pairs. If there are no such pairs, the sequence is of length zero and the sample is said to be empty.

11.5. Discussion

The parameters Lambda should be carefully chosen. If the rate is too high, overly frequent LSP setup/release procedure will result in high overhead in the control plane. In turn, the high overhead will increase bidirectional LSP setup delay. On the other hand, if the rate is too low, the sample might not completely reflect the dynamic provisioning performance of the GMPLS network. The appropriate Lambda value depends on the given network.

The parameters Td should be carefully chosen. Different switching technologies may vary significantly in performing a cross-connect operation. At the same time, the time needed to set up an LSP under different traffic may also vary significantly.

In the case of active measurement, the parameters Th should be carefully chosen. The combination of Lambda and Th reflects the load of the network. The selection of Th SHOULD take into account that the network has sufficient resources to perform subsequent tests. The value of Th MAY be constant during one sampling process for simplicity considerations.

Note that for online or passive measurements, the arrival rate and the LSP holding time are determined by actual traffic; hence, in this case, Lambda and Th are not input parameters.

It is important that, in obtaining a sample, all the LSPs MUST traverse the same route. If there are multiple routes between the ingress node ID0 and egress node ID1, EROs, or an alternate method, e.g., static configuration, MUST be used to ensure that all LSPs traverse the same route.

11.6. Methodologies

- o Select the times using the specified Poisson arrival process,
- o Set up the LSP as the methodology for the singleton bidirectional LSP setup delay, and obtain the value of bidirectional LSP setup delay, and
- o Release the LSP after Th, and wait for the next Poisson arrival event.

Note: it is possible that before the previous LSP release procedure completes, the next Poisson arrival event arrives and the LSP setup procedure is initiated. If there is resource contention between the two LSPs, the LSP setup may fail. Ways to avoid such contention are outside the scope of this document.

11.7. Typical Testing Cases

11.7.1. With No LSP in the Network

11.7.1.1. Motivation

Single bidirectional LSP setup delay with no LSP in the network is important because this reflects the inherent delay of an RSVP-TE implementation. The minimum value provides an indication of the delay that will likely be experienced when an LSP traverses the shortest route with the lightest load in the control plane.

11.7.1.2. Methodologies

Make sure that there is no LSP in the network and proceed with the methodologies described in Section 11.6.

11.7.2. With a Number of LSPs in the Network

11.7.2.1. Motivation

Single bidirectional LSP setup delay with a number of LSPs in the network is important because it reflects the performance of an operational network with considerable load. This delay can vary significantly as the number of existing LSPs varies. It can be used as a scalability metric of an RSVP-TE implementation.

11.7.2.2. Methodologies

Set up the required number of LSPs and wait until the network reaches a stable state; then, proceed with the methodologies described in Section 11.6.

11.8. Metric Reporting

The metric results including both real and undefined values MUST be reported together with the total number of values. The context under which the sample is obtained, including the selected parameters, the route traversed by the LSPs, and the testing case used, MUST also be reported.

12. A Definition for Samples of Multiple Bidirectional LSPs Setup Delay

In Section 7, we defined the singleton metric of multiple bidirectional LSPs setup delay. Now we define how to get one particular sample of multiple bidirectional LSP setup delay.

Sampling means to take a number of distinct instances of a skeleton metric under a given set of parameters. As in [RFC2330], we use Poisson sampling as an example.

12.1. Metric Name

Multiple bidirectional LSPs setup delay sample

12.2. Metric Parameters

- o ID0, the ingress LSR ID
- o ID1, the egress LSR ID
- o T0, a time
- o Tf, a time
- o Lambda_m, a rate in the reciprocal milliseconds
- o Lambda, a rate in the reciprocal milliseconds
- o X, the number of LSPs to set up
- o Th, LSP holding time
- o Td, the maximum waiting time for successful multiple unidirectional LSPs setup

12.3. Metric Units

A sequence of pairs; the elements of each pair are:

- o T, a time when the first setup is attempted
- o dT, either a real number of milliseconds or undefined

12.4. Definition

Given T0, Tf, and Lambda, compute a pseudo-random Poisson process beginning at or before T0, with an average arrival rate Lambda and ending at or after Tf. Those time values greater than or equal to T0 and less than or equal to Tf are then selected. At each of the times in this process, we obtain the value of multiple unidirectional LSP setup delay sample. The value of the sample is the sequence made up of the resulting <time, setup delay> pairs. If there are no such pairs, the sequence is of length zero and the sample is said to be empty.

12.5. Discussion

The parameter Lambda is used as an arrival rate of "batch bidirectional LSPs setup" operation. It regulates the interval in between each batch operation. The parameter Lambda_m is used within each batch operation, as described in Section 7.

The parameters Lambda and Lambda_m should be carefully chosen. If the rate is too high, overly frequent LSP setup/release procedure will result in high overhead in the control plane. In turn, the high overhead will increase unidirectional LSP setup delay. On the other hand, if the rate is too low, the sample might not completely reflect the dynamic provisioning performance of the GMPLS network. The appropriate Lambda and Lambda_m values depend on the given network.

The parameters Td should be carefully chosen. Different switching technologies may vary significantly in performing a cross-connect operation. At the same time, the time needed to set up an LSP under different traffic may also vary significantly.

It is important that, in obtaining a sample, all the LSPs MUST traverse the same route. If there are multiple routes between the ingress node ID0 and egress node ID1, EROs, or an alternate method, e.g., static configuration, MUST be used to ensure that all LSPs traverse the same route.

12.6. Methodologies

- o Select the times using the specified Poisson arrival process,

- o Set up the LSP as the methodology for the singleton multiple bidirectional LSPs setup delay, and obtain the value of multiple unidirectional LSPs setup delay, and

- o Release the LSP after Th, and wait for the next Poisson arrival event.

Note: it is possible that before the previous LSP release procedure completes, the next Poisson arrival event arrives and the LSP setup procedure is initiated. If there is resource contention between the two LSPs, the LSP setup may fail. Ways to avoid such contention are outside the scope of this document.

12.7. Typical Testing Cases

12.7.1. With No LSP in the Network

12.7.1.1. Motivation

Multiple bidirectional LSPs setup delay with no LSP in the network is important because this reflects the inherent delay of an RSVP-TE implementation. The minimum value provides an indication of the delay that will likely be experienced when an LSPs traverse the shortest route with the lightest load in the control plane.

12.7.1.2. Methodologies

Make sure that there is no LSP in the network and proceed with the methodologies described in Section 10.6.

12.7.2. With a Number of LSPs in the Network

12.7.2.1. Motivation

Multiple bidirectional LSPs setup delay with a number of LSPs in the network is important because it reflects the performance of an operational network with considerable load. This delay may vary significantly as the number of existing LSPs vary. It may be used as a scalability metric of an RSVP-TE implementation.

12.7.2.2. Methodologies

Set up the required number of LSPs, and wait until the network reaches a stable state; then, proceed with the methodologies described in Section 12.6.

12.8. Metric Reporting

The metric results including both real and undefined values MUST be reported together with the total number of values. The context under which the sample is obtained, including the selected parameters, the route traversed by the LSPs, and the testing case used, MUST also be reported.

13. A Definition for Samples of LSP Graceful Release Delay

In Section 8, we defined the singleton metric of LSP graceful release delay. Now we define how to get one particular sample of LSP graceful release delay. We also use Poisson sampling as an example.

13.1. Metric Name

LSP graceful release delay sample

13.2. Metric Parameters

- o ID0, the ingress LSR ID
- o ID1, the egress LSR ID
- o T0, a time
- o Tf, a time
- o Lambda, a rate in reciprocal milliseconds
- o Td, the maximum waiting time for successful LSP release

13.3. Metric Units

A sequence of pairs; the elements of each pair are:

- o T, a time, and
- o dT, either a real number of milliseconds or undefined

13.4. Definition

Given T0, Tf, and Lambda, we compute a pseudo-random Poisson process beginning at or before T0, with an average arrival rate Lambda and ending at or after Tf. Those time values greater than or equal to T0 and less than or equal to Tf are then selected. At each of the times in this process, we obtain the value of LSP graceful release delay sample. The value of the sample is the sequence made up of the resulting <time, LSP graceful delay> pairs. If there are no such pairs, the sequence is of length zero and the sample is said to be empty.

13.5. Discussion

The parameter Lambda should be carefully chosen. If the rate is too large, overly frequent LSP setup/release procedure will result in high overhead in the control plane. In turn, the high overhead will increase unidirectional LSP setup delay. On the other hand, if the rate is too small, the sample might not completely reflect the dynamic provisioning performance of the GMPLS network. The appropriate Lambda value depends on the given network.

It is important that, in obtaining a sample, all the LSPs MUST traverse the same route. If there are multiple routes between the ingress node ID0 and egress node ID1, EROs, or an alternate method, e.g., static configuration, MUST be used to ensure that all LSPs traverse the same route.

13.6. Methodologies

Generally, the methodology would proceed as follows:

- o Set up the LSP to be deleted
- o Select the times using the specified Poisson arrival process,
- o Release the LSP as the methodology for the singleton LSP graceful release delay, and obtain the value of LSP graceful release delay, and
- o Set up the LSP, and restart the Poisson arrival process, wait for the next Poisson arrival event.

13.7. Metric Reporting

The metric results including both real and undefined values MUST be reported together with the total number of values. The context under which the sample is obtained, including the selected parameters, and the route traversed by the LSPs MUST also be reported.

14. Some Statistics Definitions for Metrics to Report

Given the samples of the performance metric, we now offer several statistics of these samples to report. From these statistics, we can draw some useful conclusions of a GMPLS network. The value of these metrics is either a real number of milliseconds or undefined. In the following discussion, we only consider the finite values.

14.1. The Minimum of Metric

The minimum of the metric is the minimum of all the dT values in the sample. In computing this, undefined values SHOULD be treated as infinitely large. Note that this means that the minimum could thus be undefined if all the dT values are undefined. In addition, the metric minimum SHOULD be set to undefined if the sample is empty.

14.2. The Median of Metric

Metric median is the median of the dT values in the given sample. In computing the median, the undefined values MUST NOT be included.

14.3. The Maximum of Metric

The maximum of the metric is the maximum of all the dT values in the sample. In computing this, undefined values MUST NOT be included. Note that this means that measurements that exceed the upper bound are not reported in this statistic. This is an important consideration when evaluating the maximum when the number of undefined measurements is non-zero.

14.4. The Percentile of Metric

The "empirical distribution function" (EDF) of a set of scalar measurements is a function F(x), which, for any x, gives the fractional proportion of the total measurements that were <= x.

Given a percentage X, the X-th percentile of the metric means the smallest value of x for which F(x) >= X. In computing the percentile, undefined values MUST NOT be included.

See [RFC2330] for further details.

14.5. Failure Statistics of Metric

In the process of LSP setup/release, it may fail due to various reasons. For example, setup/release may fail when the control plane is overburdened or when there is resource shortage in one of the intermediate nodes. Since the setup/release failure may have significant impact on network operation, it is worthwhile to report each failure cases, so that appropriate operations can be performed to check the possible implementation, configuration or other deficiencies.

Five types of failure events are defined in previous sections:

- o Single Unidirectional LSP Setup Failure
- o Multiple Unidirectional LSP Setup Failure
- o Single Bidirectional LSP Setup Failure
- o Multiple Bidirectional LSP Setup Failure
- o LSP Graceful Release Failure

Given the samples of the performance metric, we now offer two statistics of failure events of these samples to report.

14.5.1. Failure Count

Failure Count is defined as the number of the undefined value of the corresponding performance metric (failure events) in a sample. The value of Failure Count is an integer.

14.5.2. Failure Ratio

Failure Ratio is the percentage of the number of failure events to the total number of requests in a sample. The calculation for Failure Ratio is defined as follows:

X type failure ratio = Number of X type failure events/(Number of valid X type metric values + Number of X type failure events) * 100%.

15. Discussion

It is worthwhile to point out that:

- o The unidirectional/bidirectional LSP setup delay is one ingress-egress round-trip time plus processing time. But in this document, unidirectional/bidirectional LSP setup delay has not taken the processing time in the end nodes (ingress and/or egress) into account. The timestamp T2 is taken after the endpoint node receives it. Actually, the last node has to take some time to process local procedures. Similarly, in the LSP graceful release delay, the memo has not considered the processing time in the end node.

- o This document assumes that the correct procedures for installing the data plane are followed as described in [RFC3209], [RFC3471], and [RFC3473]. That is, by the time the egress receives and processes a Path message, it is safe for the egress to transmit data on the reverse path, and by the time the ingress receives and processes a Resv message it is safe for the ingress to transmit data on the forward path. See [CCAMP-SWITCH] for detailed explanations. This document does not include any verification that the implementations of the control plane software are conformant, although such tests MAY be constructed with the use of suitable signal generation test equipment. In [CCAMP-DPM], we defined a series of metrics to do such verifications. However, it is RECOMMENDED that both the measurements defined in this document and the measurements defined in [CCAMP-DPM] are performed to complement each other.

- o Note that, in implementing the tests described in this document, a tester should be sure to measure the time taken for the control plane messages including the processing of those messages by the nodes under test.

- o Bidirectional LSPs may be set up using three-way signaling, where the initiating node will send a ResvConf message downstream upon receiving the Resv message. The ResvConf message is used to notify the terminate node that it can transfer data upstream. Actually, both directions should be ready to transfer data when the Resv message is received by the initiating node. Therefore, the bidirectional LSP setup delay defined in this document does not take the confirmation procedure into account.

16. Security Considerations

Samples of the metrics can be obtained in either active or passive manners.

In active measurement, ingress nodes inject probing messages into the control plane. Since the measurement endpoints must be conformant to signaling specifications and behave as normal signaling endpoints, it will not incur other security issues than normal LSP provisioning. However, the measurement parameters must be carefully selected so that the measurements inject trivial amounts of additional traffic into the networks they measure. If they inject "too much" traffic, they can skew the results of the measurement, and, in extreme cases, cause congestion and denial of service.

When samples of the metrics are collected in a passive manner, e.g., by monitoring the operations on real-life LSPs, the implementation of the monitoring and reporting mechanism must be careful so that they will not be used to attack the control plane. A typical implementation may use the Management Information Base (MIB) to collect/store the metrics and access to the MIB is limited to the Network Management Systems (NMSs). In this case, passive monitoring will not incur other security issues than implementing the MIBs and NMSs. If an implementation chooses to expose the performance data to other applications, then it must take into account the possible security issues it may face. For example, when exposing the performance data through Simple Network Management Protocol (SNMP), certain authentication methods should be used to ensure that the entity maintaining the performance data are not subject to unauthorized readings and modifications. Rate limiting on the performance query may also be desirable to reduce the risk that the entity maintaining the performance data are overwhelmed by too many query requests. It is RECOMMENDED that implementers consider the security features as provided by the SNMPv3 framework (see [RFC3410], section 8), including full support for the SNMPv3 cryptographic mechanisms (for authentication and privacy).

Additionally, the security considerations pertaining to the original RSVP protocol [RFC2205] and its TE extensions [RFC3209] also remain relevant.

17. Acknowledgments

We wish to thank Dan Li, Fang Liu (Christine), Zafar Ali, Monique Morrow, Adrian Farrel, Deborah Brungard, Lou Berger, Thomas D. Nadeau for their comments and help. Lou Berger and Adrian Farrel have made text contributions to this document.

We wish to thank experts from IPPM and BMWG -- Reinhard Schrage, Al Morton, and Henk Uijterwaal -- for reviewing this document. Reinhard Schrage has made text contributions to this document.

This document contains ideas as well as text that have appeared in existing IETF documents. The authors wish to thank G. Almes, S. Kalidindi, and M. Zekauskas.

We also wish to thank Weisheng Hu, Yaohui Jin, and Wei Guo in the state key laboratory of advanced optical communication systems and networks for the valuable comments. We also wish to thank the support from National Natural Science Foundation of China (NSFC) and 863 program of China.

18. References

18.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC2205] Braden, B., Zhang, L., Berson, S., Herzog, S., and S. Jamin, "Resource ReSerVation Protocol (RSVP) -- Version 1 Functional Specification", RFC 2205, September 1997.

[RFC2679] Almes, G., Kalidindi, S., and M. Zekauskas, "A One-way Delay Metric for IPPM", RFC 2679, September 1999.

[RFC2681] Almes, G., Kalidindi, S., and M. Zekauskas, "A Round-trip Delay Metric for IPPM", RFC 2681, September 1999.

[RFC3209] Awduche, D., Berger, L., Gan, D., Li, T., Srinivasan, V., and G. Swallow, "RSVP-TE: Extensions to RSVP for LSP Tunnels", RFC 3209, December 2001.

[RFC3471] Berger, L., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Functional Description", RFC 3471, January 2003.

[RFC3473] Berger, L., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Extensions", RFC 3473, January 2003.

[RFC3945] Mannie, E., "Generalized Multi-Protocol Label Switching (GMPLS) Architecture", RFC 3945, October 2004.

[RFC4208] Swallow, G., Drake, J., Ishimatsu, H., and Y. Rekhter, "Generalized Multiprotocol Label Switching (GMPLS) User-Network Interface (UNI): Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Support for the Overlay Model", RFC 4208, October 2005.

18.2. Informative References

[CCAMP-DPM] Sun, W., Zhang, G., Gao, J., Xie, G., Papneja, R., Gu, B., Wei, X., Otani, T., and R. Jing, "Label Switched Path (LSP) Data Path Delay Metric in Generalized MPLS/ MPLS-TE Networks", Work in Progress, December 2009.

[CCAMP-SWITCH] Shiomoto, K. and A. Farrel, "Advice on When It is Safe to Start Sending Data on Label Switched Paths Established Using RSVP-TE", Work in Progress, October 2009.

[RFC2330] Paxson, V., Almes, G., Mahdavi, J., and M. Mathis, "Framework for IP Performance Metrics", RFC 2330, May 1998.

[RFC3410] Case, J., Mundy, R., Partain, D., and B. Stewart, "Introduction and Applicability Statements for Internet-Standard Management Framework", RFC 3410, December 2002.

Appendix A. Authors' Addresses

Jianhua Gao
Huawei Technologies Co., LTD.
China

Phone: +86 755 28973237
EMail: gjhhit@huawei.com

Guowu Xie
University of California, Riverside
900 University Ave.
Riverside, CA 92521
USA

Phone: +1 951 237 8825
EMail: xieg@cs.ucr.edu

Rajiv Papneja
Isocore
12359 Sunrise Valley Drive, STE 100
Reston, VA 20190
USA

Phone: +1 703 860 9273
EMail: rpapneja@isocore.com

Bin Gu
IXIA
Oriental Kenzo Plaza 8M, 48 Dongzhimen Wai Street, Dongcheng District
Beijing 200240
China

Phone: +86 13611590766
EMail: BGu@ixiacom.com

Xueqin Wei
Fiberhome Telecommunication Technology Co., Ltd.
Wuhan
China

Phone: +86 13871127882
EMail: xqwei@fiberhome.com.cn

Tomohiro Otani
KDDI R&D Laboratories, Inc.
2-1-15 Ohara Kamifukuoka Saitama
356-8502
Japan

Phone: +81-49-278-7357
EMail: otani@kddilabs.jp

Ruiquan Jing
China Telecom Beijing Research Institute
118 Xizhimenwai Avenue
Beijing 100035
China

Phone: +86-10-58552000
EMail: jingrq@ctbri.com.cn

Editors' Addresses

Weiqiang Sun (editor)
Shanghai Jiao Tong University
800 Dongchuan Road
Shanghai 200240
China

Phone: +86 21 3420 5359
EMail: sunwq@mit.edu

Guoying Zhang (editor)
China Academy of Telecommunication Research, MIIT, China.
No.52 Hua Yuan Bei Lu,Haidian District
Beijing 100083
China

Phone: +86 1062300103
EMail: zhangguoying@mail.ritt.com.cn