

Network Working Group
Request for Comments: 4378
Category: Informational

D. Allan, Ed.
Nortel Networks
T. Nadeau, Ed.
Cisco Systems, Inc.
February 2006

A Framework for Multi-Protocol Label Switching (MPLS)
Operations and Management (OAM)

Status of This Memo

This memo provides information for the Internet community. It does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (2006).

Abstract

This document is a framework for how data plane protocols can be applied to operations and maintenance procedures for Multi-Protocol Label Switching (MPLS). The document is structured to outline how Operations and Management (OAM) functionality can be used to assist in fault, configuration, accounting, performance, and security management, commonly known by the acronym FCAPS.

Table of Contents

1. Introduction	2
2. Terminology	2
3. Fault Management	2
3.1. Fault Detection	2
3.2. Diagnosis	6
3.3. Availability	7
4. Configuration Management	7
5. Accounting	7
6. Performance Management	7
7. Security Management	8
8. Security Considerations	9
9. Acknowledgements	9
10. Normative References	9

1. Introduction

This memo outlines in broader terms how data plane protocols can assist in meeting the Operations and Management (OAM) requirements outlined in [RFC4377] and [Y1710] and can apply to the management functions of fault, configuration, accounting, performance, and security (commonly known as FCAPS) for MPLS networks, as defined in [RFC3031]. The approach of the document is to outline functionality, the potential mechanisms to provide the function, and the required applicability of data plane OAM functions. Included in the discussion are security issues specific to use of tools within a provider domain and use for inter-provider Label Switched Paths (LSPs).

2. Terminology

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].

OAM	Operations and Management
FCAPS	Fault management, Configuration management, Administration management, Performance management, and Security management
FEC	Forwarding Equivalence Class
ILM	Incoming Label Map
NHLFE	Next Hop Label Forwarding Entry
MIB	Management Information Base
LSR	Label Switching Router
RTT	Round Trip Time

3. Fault Management

3.1. Fault Detection

Fault detection encompasses the identification of all data plane failures between the ingress and egress of an LSP. This section will enumerate common failure scenarios and explain how one might (or might not) detect the situation.

3.1.1. Enumeration and Detection of Types of Data Plane Faults

Lower-layer faults:

Lower-layer faults are those in the physical or virtual link that impact the transport of MPLS labeled packets between adjacent LSRs at the specific level of interest. Some physical links (such as SONET/SDH) may have link-layer OAM functionality and detect and notify the LSR of link-layer faults directly. Some physical links (such as Ethernet) may not have this capability and require MPLS or IP layer heartbeats to detect failures. However, once detected, reaction to these fault notifications is often the same as those described in the first case.

Node failures:

Node failures are those that impact the forwarding capability of a node component, including its entire set of links. This can be due to component failure, power outage, or reset of the control processor in an LSR employing a distributed architecture, etc.

MPLS LSP mis-forwarding:

Mis-forwarding occurs when there is a loss of synchronization between the data and the control planes in one or more nodes. This can occur due to hardware failure, software failure, or configuration problems.

It will manifest itself in one of two forms:

- packets belonging to a particular LSP are cross-connected into an NHLFE for which there is no corresponding ILM at the next downstream LSR. This can occur in cases where the NHLFE entry is corrupted. Therefore, the packet arrives at the next LSR with a top label value for which the LSR has no corresponding forwarding information, and is typically dropped. This is a No Incoming Label Map (No ILM) condition and can be detected directly by the downstream LSR that receives the incorrectly labeled packet.
- packets belonging to a particular LSP are cross-connected into an incorrect NHLFE entry for which there is a corresponding ILM at the next downstream LSR, but is associated with a different LSP. This may be detected by the following:
 - o some or all of the misdirected traffic is not routable at the egress node, or

- o OAM probing is able to detect the fault by detecting the inconsistency between the data path and the control plane state.

Discontinuities in the MPLS Encapsulation

The forwarding path of the FEC carried by an LSP may transit nodes or links for which MPLS is not configured. This may result in a number of behaviors that are undesirable and not easily detected.

- if exposed, payload is not routable at the LSR, resulting in silent discard, OR
- the exposed MPLS label was not offered by the LSR, which may result in either silent discard or mis-forwarding.

Alternately, the payload may be routable and packets successfully delivered but may bypass associated MPLS instrumentation and tools.

MTU problems

MTU problems occur when client traffic cannot be fragmented by intermediate LSRs and is dropped somewhere along the path of the LSP. MTU problems should appear as a discrepancy in the traffic count between the set of ingress LSRs and the egress LSRs for an FEC and will appear in the corresponding MPLS MIB performance tables in the transit LSRs as discarded packets.

TTL Mishandling

The implementation of TTL handling is inconsistent at penultimate hop LSRs. Tools that rely on consistent TTL processing may produce inconsistent results in any given network.

Congestion

Congestion occurs when the offered load on any interface exceeds the link capacity for sufficient time that the interface buffering is exhausted. Congestion problems will appear as a discrepancy in the traffic count between the set of ingress LSRs and the egress LSRs for an FEC and will appear in the MPLS MIB performance tables in the transit LSRs as discarded packets.

Mis-ordering

Mis-ordering of LSP traffic occurs when incorrect or inappropriate load sharing is implemented within an MPLS network. Load sharing typically takes place when multiple equal-cost paths exist between the ingress and egress of an LSP. In these cases, traffic is split among these equal-cost paths using a variety of algorithms. One such algorithm relies on splitting traffic between each path on a per-packet basis. When this is done, it is possible for some packets along the path to be delayed due to congestion or slower links, which may result in packets being received out of order at the egress. Detection and remedy of this situation may be left up to client applications that use the LSPs. For instance, TCP is capable of re-ordering packets belonging to a specific flow (although this may result in re-transmission of some of the mis-ordered packets).

Detection of mis-ordering can also be determined by sending probe traffic along the path and verifying that all probe traffic is indeed received in the order it was transmitted. This will only detect truly pathological problems as mis-ordering typically is an insufficiently predictable and repeatable problem.

LSRs do not normally implement mechanisms to detect mis-ordering of flows.

Payload Corruption

Payload corruption may occur and may be undetected by LSRs. Such errors are typically detected by client payload integrity mechanisms.

3.1.2. Timeliness

The design of Service Level Agreements (SLAs) and management support systems requires that ample headroom be allotted in terms of their processing capabilities in order to process and handle all necessary fault conditions within the bounds stipulated in the SLA. This includes planning for event handling using a time budget that takes into account the over-all SLA and the time required to address any defects that arise. However, it is possible that some fault conditions may surpass this budget due to their catastrophic nature (e.g., fibre cut) or due to incorrect planning of the time processing budget.




Figure 1: Fault Correction Budget

In figure 1, we represent the overall fault correction time budget by the maximum time as specified in an SLA for the service in question. This time is then divided into two subsections, the first encompassing the total time required to detect a fault and notify an operator (or optionally automatically correct the defect). This section may have an explicit maximum time to detect defects arising from either the application or a need to do alarm management (i.e., suppression), and this will be reflected in the frequency of OAM execution. The second section indicates the time required to notify the operational systems used to diagnose, isolate, and correct the defect (if they cannot be corrected automatically).

3.2. Diagnosis

3.2.1. Characterization

Characterization is defined as determining the forwarding path of a packet (which may not be necessarily known). Characterization may be performed on a working path through the network. For example, this is done to determine equal-cost multi-paths (ECMP), the MTU of a path, or simply to know the path occupied by a specific FEC. Characterization will be able to leverage mechanisms used for isolation.

3.2.2. Isolation

Isolation of a fault can occur in two forms. In the first case, the local failure is detected, and the node where the failure occurred is capable of issuing an alarm for such an event. The node should attempt to withdraw the defective resources and/or rectify the situation prior to raising an alarm. Active data plane OAM mechanisms may also detect the failure conditions remotely and issue their own alarms if the situation is not rectified quickly enough.

In the second case, the fault has not been detected locally. In this case, the local node cannot raise an alarm, nor can it be expected to

rectify the situation. In this case, the failure may be detected remotely via data plane OAM. This mechanism should also be able to determine the location of the fault, perhaps on the basis of limited information such as a customer complaint. This mechanism may also be able to automatically remove the defective resources from the network and restore service, but should at least provide a network operator with enough information by which they can perform this operation. Given that detection of faults is desired to happen as quickly as possible, tools which possess the ability to incrementally test LSP health should be used to uncover faults.

3.3. Availability

Availability is the measure of the percentage of time that a service is operating within a specification, often specified by an SLA.

MPLS has several forwarding modes (depending on the control plane used). As such, more than one model may be defined and more than one measurement technique may be required.

4. Configuration Management

Data plane OAM can assist in configuration management by providing the ability to verify the configuration of an LSP or of applications utilizing that LSP. This would be an ad-hoc data plane probe that should verify path integrity (a complete path exists) and that the path function is synchronized with the control plane. As part of the payload, the probe would carry relevant control plane information that the receiver would be able to compare with the local-control plane configuration.

5. Accounting

The requirements for accounting in MPLS networks, as specified in [RFC4377], do not place any requirements on data plane OAM.

6. Performance Management

Performance management permits the information transfer characteristics of LSPs to be measured, perhaps in order to be compared against an SLA. This falls into two categories: latency (where jitter is considered a variation in latency) and information loss.

Latency can be measured in two ways: one is to have precisely synchronized clocks at the ingress and egress such that time-stamps in PDUs flowing from the ingress to the egress can be compared. The other is to use an exchange of PING type PDUs that gives a round trip

time (RTT) measurement, and an estimate of the one-way latency that can be inferred with some loss of precision. Use of load spreading techniques, such as ECMP, mean that any individual RTT measurement is only representative of the typical RTT for an FEC.

To measure information loss, a common practice is to periodically read ingress and egress counters (i.e., MIB module counters). This information may also be used for offline correlation. Another common practice is to send explicit probe traffic that traverses the data plane path in question. This probe traffic can also be used to measure jitter and delay.

7. Security Management

Providing a secure OAM environment is required if MPLS specific network mechanisms are to be used successfully. To this end, operators have a number of options when deploying network mechanisms including simply filtering OAM messages at the edge of the MPLS network. Malicious users should not be able to use non-MPLS interfaces to insert MPLS-specific OAM transactions. Provider initiated OAM transactions should be able to be blocked from leaking outside the MPLS cloud.

Finally, if a provider does wish to allow OAM messages to flow into (or through) their networks, for example, in a multi-provider deployment, authentication and authorization are required to prevent malicious and/or unauthorized access. Also, given that MPLS networks often run IP simultaneously, similar requirements apply to any native IP OAM network mechanisms in use. Therefore, authentication and authorization for OAM technologies is something that **MUST** be considered when designing network mechanisms that satisfy the framework presented in this document.

OAM messaging can address some existing security concerns with the MPLS architecture. That is, through rigorous defect handling, operator's can offer their customers a greater degree of integrity protection that their traffic will not be incorrectly delivered (for example, by being able to detect leaking LSP traffic from a VPN).

Support for inter-provider data plane OAM messaging introduces a number of security concerns as, by definition, portions of LSPs will not be within a single provider's network the provider has no control over who may inject traffic into the LSP, which can be exploited for denial of service attacks. OAM PDUs are not explicitly identified in the MPLS header and therefore are not typically inspected by transit LSRs. This creates opportunity for malicious or poorly behaved users to disrupt network operations.

Attempts to introduce filtering on target LSP OAM flows may be problematic if flows are not visible to intermediate LSRs. However, it may be possible to interdict flows on the return path between providers (as faithfulness to the forwarding path is to a return path requirement) to mitigate aspects of this vulnerability.

OAM tools may permit unauthorized or malicious users to extract significant amounts of information about network configuration. This would be especially true of IP based tools as, in many network configurations, MPLS does not typically extend to untrusted hosts, but IP does. For example, TTL hiding at ingress and egress LSRs will prevent external users from using TTL-based mechanisms to probe an operator's network. This suggests that tools used for problem diagnosis or which, by design, are capable of extracting significant amounts of information will require authentication and authorization of the originator. This may impact the scalability of such tools when employed for monitoring instead of diagnosis.

8. Security Considerations

This document describes a framework for MPLS Operations and Management. Although this document discusses and addresses some security concerns in Section 7, it does not introduce any new security concerns.

9. Acknowledgements

The editors would like to thank Monique Morrow from Cisco Systems and Harmen van Der Linde from AT&T for their valuable review comments on this document.

10. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC3031] Rosen, E., Viswanathan, A., and R. Callon, "Multiprotocol Label Switching Architecture", RFC 3031, January 2001.
- [RFC4377] Nadeau, T., Morrow, M., Swallow, G., Allan, D., and S. Matsushima, "Operations and Management (OAM) Requirements for Multi-Protocol Label Switched (MPLS) Networks", RFC 4377, February 2006.
- [Y1710] ITU-T Recommendation Y.1710(2002), "Requirements for OAM Functionality for MPLS Networks".

Authors' Addresses

David Allan
Nortel Networks
3500 Carling Ave.
Ottawa, Ontario, CANADA

Phone: +1-613-763-6362
EMail: dallan@nortel.com

Thomas D. Nadeau
Cisco Systems
300 Beaver Brook Drive
Boxborough, MA 01824

Phone: +1-978-936-1470
EMail: tnadeau@cisco.com

Full Copyright Statement

Copyright (C) The Internet Society (2006).

This document is subject to the rights, licenses and restrictions contained in BCP 78, and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in BCP 78 and BCP 79.

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

Acknowledgement

Funding for the RFC Editor function is provided by the IETF Administrative Support Activity (IASA).

