Network Working Group
Internet Draft
Category: Standards Track
Expiration Date: January 2006

George Swallow
Cisco Systems, Inc.

Stewart Bryant
Cisco Systems, Inc.

Loa Andersson
Acreo

July 2005

Avoiding Equal Cost Multipath Treatment in MPLS Networks

draft-ietf-mpls-ecmp-bcp-01.txt

Status of this Memo

Abstract

This document describes the Equal Cost Multipath (ECMP) behavior of currently deployed MPLS networks and makes best practice recommendations for anyone defining an application to run over an MPLS network and wishes to avoid such treatment.

Contents

1. Introduction

This document describes the Equal Cost Multipath (ECMP) behavior of currently deployed MPLS networks and makes best practice recommendations for anyone defining an application to run over an MPLS network and wishes to avoid such treatment. While disabling ECMP behavior is an option open to most operators, few (if any) have chosen to do so. Thus ECMP behavior is a reality that must be reckoned with.

1.1. Terminology

ECMP	Equal Cost Multipath
FEC	Forwarding Equivalence Class
IP ECMP	A forwarding behavior in which the selection of the next-hop between equal cost routes is based on the header(s) of an IP packet
Label ECMP	A forwarding behavior in which the selection of the next-hop between equal cost routes is based on the label stack of an MPLS packet
LSP	Label Switched Path
LSR	Label Switching Router

2. Current EMCP Practices

The MPLS label stack and Forwarding Equivalence Classes are defined in [RFC3031]. The MPLS label stack does not carry a Protocol Identifier. Instead the payload of an MPLS packet is identified by the Forwarding Equivalence Class (FEC) of the bottom most label. Thus it is not possible to know the payload type if one does not know the label binding for the bottom most label. Since an LSR which is processing a label stack need only know the binding for the label(s) it must process, it is very often the case that LSRs along an LSP are unable to determine the payload type of the carried contents.

As a means of potentially reducing delay and congestion, IP networks have taken advantage of multiple paths through a network by splitting traffic flows across those paths. The general name for this practice is Equal Cost Multipath or ECMP. In general this is done by hashing on various fields on the IP or contained headers. In practice, within a network core, the hashing in based mainly or exclusively on the IP source and destination addresses. The reason for splitting

aggregated flows in this manner is to minimize the re-ordering of packets belonging to individual flows contained within the aggregated flow. Within this document we use the term IP ECMP for this type of forwarding algorithm.

In the early days of MPLS, the payload was almost exclusively IP. Even today the overwhelming majority of carried traffic remains IP. Providers of MPLS equipment sought to continue this IP ECMP behavior. As shown above, it is not possible to know whether the payload of an MPLS packet is IP at every place where IP ECMP needs to be performed. Thus vendors have taken the liberty of guessing what the payload is. By inspecting the first nibble beyond the label stack, it can be inferred that a packet is not IPv4 or IPv6 if the value of the nibble (where the IP version number would be found) is not 0x4 or 0x6 respectively. Most deployed LSRs will treat a packet whose first nibble is equal to 0x4 as if the payload were IPv4 for purposes of IP ECMP.

A consequence of this is that any application which defines a FEC which does not take measures to prevent the values 0x4 and 0x6 from occurring in the first nibble of the payload may be subject to IP ECMP and thus having their flows take multiple paths and arriving with considerable jitter and possibly out of order. While none of this is in violation of the basic service offering of IP, it is detrimental to the performance of various classes of applications. It also complicates the measurement, monitoring and tracing of those flows.

New MPLS payload types are emerging such as those specified by the IETF PWE3 and AVT working groups. These payloads are not IP and, if specified without constraint might be mistaken for IP.

It must also be noted that LSRs which correctly identify a payload as not being IP, may still need to load-share this traffic across multiple equal-cost paths. In this case a LABEL ECMP algorithm is employed, where a hash is computed on all or part(s) of the label stack. Any reserved label, no matter where it is located in the stack, may be included in the computation for load balancing. Modification of the label stack between packets of a single flow could result in re-ordering that flow. That is, were an explicit null or a router-alert label to be added to a packet, that packet could take a different path through the network.

Note that for some applications, being mistaken for IPv4 may not be detrimental. The trivial case where the payload behind the top label is a packet belonging to an MPLS IPv4 VPN. Here the real payload is IP and most (if not all) deployed equipment will locate the end of the label stack and correctly perform IP ECMP.

A less obvious case is when the packets of a given flow happen to have constant values in the fields upon which IP ECMP would be performed. For example if an ethernet frame immediately follows the label and the LSR does not do ECMP on IPv6, then either the first nibble will be 0x4 or it will be something else. If the nibble is not 0x4 then no IP ECMP is performed, but Label ECMP may be performed. If it is 0x4, then the constant values of the MAC addresses overlay the fields that would be occupied by the source and destination addresses of an IP header.

3. Recommendations for Avoiding ECMP Treatment

The field in the figure below tagged "Application Label" is a label of the FEC Type used/defined by the application. It is the bottom most label in the label stack. As such its FEC Type defines the payload which follows. Anyone defining an application to be transported over MPLS is free to define new FEC Types and the format of the payload which will be carried.

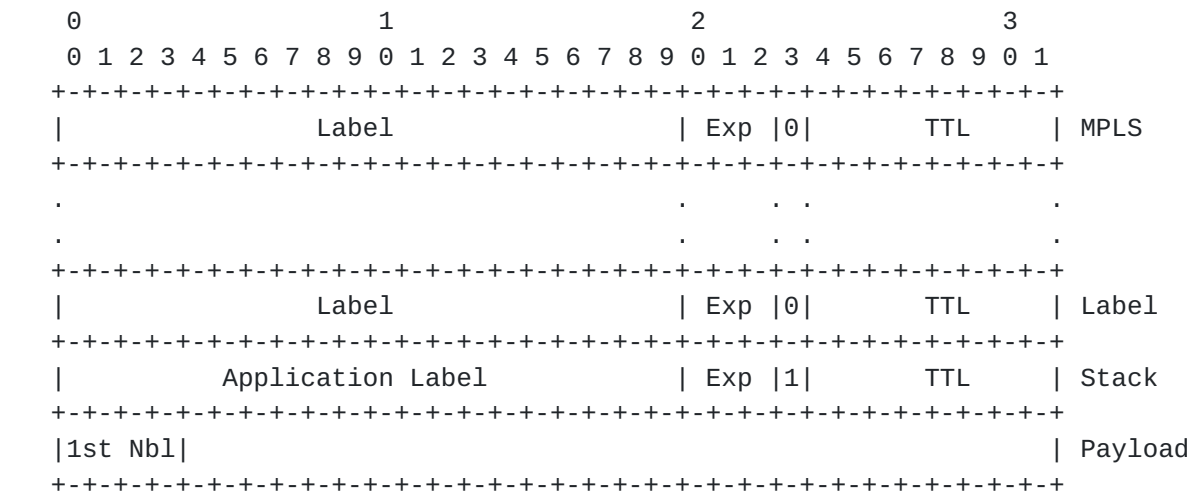

In order to avoid IP ECMP treatment it is necessary that an application take precautions to not be mistaken as IP by deployed equipment that snoops on the presumed location of the IP Version field. Thus, at a minimum, the chosen format must disallow the values 0x4 and 0x6 in the first nibble of their payload.

It is strongly recommended, however, that applications restrict the first nibble values to 0x0 and 0x1. This will ensure that that their traffic flows will not be affected if some future routing equipment does similar snooping on some future version of IP.

4. Security Considerations

This memo documents current practices. As such it creates no new security considerations.

5. References

5.1. Normative References

[RFC3031] Rosen, E. et al., "Multiprotocol Label Switching Architecture", January 2001.

6. Authors' Addresses

Loa Andersson
Acreo

Email: loa@pi.se

Stewart Bryant
Cisco Systems
250, Longwater,
Green Park,
Reading, RG2 6GB, UK

Email: stbryant@cisco.com

George Swallow
Cisco Systems, Inc.
1414 Massachusetts Ave
Boxborough, MA 01719

Email: swallow@cisco.com

Expiration Date

January 2006

Disclaimer of Validity