

Extended Administrative Groups in MPLS-TE
draft-osborne-mpls-extended-admin-groups-03

Abstract

This document provides additional administrative groups (sometimes referred to as "link colors") to the IGP extensions for MPLS-TE.

Requirements Language

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 [RFC 2119](#) [[RFC2119](#)].

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on July 13, 2014.

Copyright Notice

Copyright (c) 2014 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must

include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1.	Introduction	2
1.1.	Do we need more than 32 bits?	2
2.	Extended Administrative Groups sub-TLV	3
2.1.	Packet Format	3
2.2.	Admin group numbering	4
2.3.	Backward compatability	4
2.3.1.	AG and EAG coexistence	4
2.3.2.	Desire for unadvertised EAG bits	4
3.	Signaling Extended Administrative Groups in RSVP	5
4.	Security Considerations	5
5.	IANA Considerations	5
6.	Acknowledgements	5
7.	Normative References	6
	Author's Address	6

[1.](#) Introduction

MPLS-TE advertises 32 administrative groups (commonly referred to as "colors" or "link colors") using the Administrative Group sub-TLV of the Link TLV. This is defined for OSPFv2 [[RFC3630](#)], OSPFv3 [[RFC5329](#)] and ISIS [[RFC5305](#)].

This document adds a sub-TLV to the IGP TE extensions, "Extended Administrative Group". This sub-TLV provides for additional administrative groups (link colors) beyond the current limit of 32.

[1.1.](#) Do we need more than 32 bits?

The IGP extensions to support MPLS-TE (RFCs 3630 and 5305) define a link TLV known as Administrative Group (AG) with a limit of 32 AGs per link. This property comes from [section 6.2 of RFC 2702](#) [[RFC2702](#)]. RFCs 3630 and 5305 describe the mechanics of the TLV; the actual definition of the field comes from [RFC 2702](#).

Networks have grown over time, and MPLS-TE has grown right along with them. Implementing network-wide policies such as the ones listed in [RFC 2702 section 6.2](#) with only 32 bits gives the operator only five bits per policy with two bits left over. This can be quite constraining - AGs are a bit mask, so five bits does not mean 32 possible values, it means 5. Running a country-wide or worldwide MPLS-TE network with only five possible values for each case is clearly too constraining.

Even if an operator wishes to use AGs to implement only a single policy it is possible to run out of bit values. One such use case is #5, using AGs to constrain traffic within specific topological regions of the network. A large network may well have far more than 32 geographic regions. One particular operator uses AGs to flag network regions down to the metro scale, e.g. Seattle, San Francisco, Dallas, Chicago, St. Louis, etc. MPLS-TE tunnels are then specified with affinities to include or exclude specific metro regions in their path calculation. It is clear that 32 may not be enough even for a US-based network, nevermind a worldwide network.

There may be some opportunity for color reuse; that is, bit 0x8 may mean 'Seattle' and 'Prague' and 'Singapore' depending on the geography in which it is used. In practice, coordinating this reuse is fraught with peril and the reuse effectively becomes the limiting factor in MPLS-TE deployment. With this example it is not possible to build an LSP which avoids Seattle while including Prague, as it is the same AG value.

This document provides Extended Administrative Groups (EAGs). The number of EAGs has no fixed limit, it is constrained only by protocol-specific restrictions such as LSA or MTU size. While an operator may one day need to go beyond these protocol-specific restrictions, allow for an arbitrary number of EAGs should easily provide the operator with hundreds or thousands of bit values, thus no longer making the number of AGs an impediment to network growth.

2. Extended Administrative Groups sub-TLV

The Extended Administrative Groups sub-TLV is used in addition to the Administrative Groups when a node wishes to advertise more than 32 colors for a link. The EAG sub-TLV is optional.

This document uses the term 'colors' as a shorthand to refer to particular bits with an AG or EAG. The examples in this document use 'red' to represent the least significant bit in the AG (red == 0x1), 'blue' to represent the second bit (blue == 0x2). To say that a link has a given color or that the specified color is set on the link is to say that the corresponding bit or bits in the link's AG are set to 1.

2.1. Packet Format

The format of the Extended Administrative Groups sub-TLV is the same for both OSPF and ISIS:


```

      0              1              2              3
    0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|                               Extended Admin Group                               |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|                               .....                               |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|                               Extended Admin Group                               |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+

```

The Type of the sub-TLV for OSPF and ISIS is TBD. The Length is the size of the Extended Admin Group (EAG) value in bytes. The EAG may be of any length, but MUST be a multiple of 4 bytes. The only limits on EAG size are those which are imposed by protocol-specific or media-specific constraints (e.g. max packet length).

[2.2. Admin group numbering](#)

By convention, the existing Administrative Group TLVs are numbered 0 (LSB) to 31 (MSB). The EAG values are a superset of AG. That is, bits 0-31 in the EAG have the same meaning and MUST have the same values as an AG flooded for the same link.

[2.3. Backward compatability](#)

There are two things to consider for backward compatibility with existing AG implementations - how do AG and EAG coexist, and what happens if a node has matching criteria for unadvertised EAG bits?

[2.3.1. AG and EAG coexistence](#)

If a node advertises EAG it MAY also advertise AG. If a node advertises both AG and EAG then the first 32 bits of the EAG MUST be identical to the advertised AG. If the AG and EAG advertised for a link differ, the EAG MUST take priority. This allows nodes which do not support EAG to obtain some link color information from the network, but also allow for an eventual migration away from AG.

[2.3.2. Desire for unadvertised EAG bits](#)

The existing AG sub-TLV is optional; thus a node may be configured with a preference to include red or exclude blue, and be faced with a link that is not advertising a value for either blue or red. What does an implementation do in this case? It shouldn't assume that red is set, but it is also arguably incorrect to assume that red is NOT set, as a bit must first exist before it can be set to 0.

Practically speaking this has not been an issue for deployments, as many implementations always advertise the AG bits, often with a default value of 0x00000000. However, this issue may be of more concern once EAGs are added to the network. EAGs may exist on some nodes but not others, and the EAG length may be longer for some links than for others.

Each implementation is free to choose its own method for handling this question. However, to encourage maximum interoperability an implementation SHOULD treat specified but unadvertised EAG bits as if they are set to 0. A node MAY provide other (configurable) strategies for handling this case.

3. Signaling Extended Administrative Groups in RSVP

RSVP provides the ability to signal link affinity via the SESSION_ATTRIBUTE object with C-Type 1 in[RFC3209]. At first glance it seems useful to extend RSVP to provide a session attribute which can signal extended affinities. As it turns out, there are several non-trivial things to tackle were one to provide such an extension. In addition, an informal survey of the field, both MPLS-TE implementors and network operators, suggests that the ability to signal affinity bits in a SESSION_ATTRIBUTE object is not widely deployed today. It is thus likely that signaling EAG in a SESSION_ATTRIBUTE would see virtually no deployment. As this work would be both non-trivial and aimed at a solution unlikely to be deployed, it is not addressed in this document.

This document does not preclude solving this problem in the future should it be necessary.

4. Security Considerations

This extension adds no new security considerations.

5. IANA Considerations

This document requests a sub-TLV allocation in both OSPF and ISIS.

6. Acknowledgements

Thanks to Santiago Alvarez, Rohit Gupta, Liem Nguyen, Tarek Saad, and Robert Sawaya for their review and comments.

7. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC2702] Awduche, D., Malcolm, J., Agogbua, J., O'Dell, M., and J. McManus, "Requirements for Traffic Engineering Over MPLS", [RFC 2702](#), 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.
- [RFC3630] Katz, D., Kompella, K., and D. Yeung, "Traffic Engineering (TE) Extensions to OSPF Version 2", [RFC 3630](#), September 2003.
- [RFC5305] Li, T. and H. Smit, "IS-IS Extensions for Traffic Engineering", [RFC 5305](#), October 2008.
- [RFC5329] Ishiguro, K., Manral, V., Davey, A., and A. Lindem, "Traffic Engineering Extensions to OSPF Version 3", [RFC 5329](#), September 2008.

Author's Address

Eric Osborne
Cisco Systems

Email: eosborne@cisco.com

