

Internet Engineering Task Force (IETF)
Request for Comments: 7308
Category: Standards Track
ISSN: 2070-1721

E. Osborne
July 2014

Extended Administrative Groups in MPLS Traffic Engineering (MPLS-TE)

Abstract

MPLS Traffic Engineering (MPLS-TE) advertises 32 administrative groups (commonly referred to as "colors" or "link colors") using the Administrative Group sub-TLV. This is defined for OSPFv2 (RFC 3630), OSPFv3 (RFC 5329) and IS-IS (RFC 5305).

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.

Status of This Memo

This is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc7308>.

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.	Requirements Language	3
2.	Extended Administrative Groups Sub-TLV	3
2.1.	Packet Format	4
2.2.	Admin Group Numbering	4
2.3.	Backward Compatibility	4
2.3.1.	AG and EAG Coexistence	4
2.3.2.	Desire for Unadvertised EAG Bits	5
3.	Security Considerations	5
4.	IANA Considerations	6
5.	Acknowledgements	6
6.	References	6
6.1.	Normative References	6
6.2.	Informative References	7

1. Introduction

Do we need more than 32 bits?

The IGP extensions to support MPLS-TE (RFCs 3630 [RFC3630] and 5305 [RFC5305]) define a link TLV known as Administrative Group (AG) with a limit of 32 AGs per link. The concept of Administrative Groups comes from Section 6.2 of RFC 2702 [RFC2702], which calls them Resource Classes. RFCs 3630 [RFC3630] and 5305 [RFC5305] describe the mechanics of the TLV and use the term Administrative Groups (sometimes abbreviated herein as AGs), as does this document.

Networks have grown over time, and MPLS-TE has grown right along with them. Administrative Groups are advertised as fixed-length 32-bit bitmasks. This can be quite constraining, as it is possible to run out of values rather quickly. One such use case is #5 in Section 6.2 of RFC 2702 [RFC2702], 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 builds their network along the lines of this use case, using 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. Each metro region is given its own bit in the AG bitmask. This means that 32 bits can only (cleanly) represent 32 metro areas. It should be obvious that 32 may not be enough even for a US-based network, never mind a worldwide network.

There may be some opportunity for color reuse; that is, bit 0x8 may mean 'Seattle' or 'Prague' or '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 a Label Switched Path (LSP) that 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 Link State Advertisement (LSA) or MTU size. While an operator may one day need to go beyond these protocol-specific restrictions, allowing 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.

EAG's intended use case is within a single domain. As such, this document provides no support for signaling an EAG. It provides no analog to either the SESSION_ATTRIBUTE of C-Type 1 defined in [RFC3209] nor the LSP Attributes (LSPA) object of the Path Computation Element Communication Protocol (PCEP), defined in [RFC5440]. Since this specification provides no way of signaling an LSP's path requirements in reference to the EAG, such constraints may only be applied at the ingress.

1.1. 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].


2. Extended Administrative Groups Sub-TLV

This document defines the Extended Administrative Group (EAG) sub-TLV for both OSPF [RFC3630] and IS-IS [RFC5305]. The EAG sub-TLV is used in addition to the Administrative Groups when an operator wants to make more than 32 colors available for advertisement in a network. The EAG sub-TLV is optional. Coexistence of EAG and AG TLVs is covered in Section 2.3.1 of this document.

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 IS-IS:



The Type of the sub-TLV for OSPF is 26 and for IS-IS is 14. The Length is the size of the Extended Admin Group (EAG) value in bytes. The EAG may be of any non-zero length, but it MUST be a multiple of 4 bytes. The only limits on EAG size are those that 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 sub-TLVs are numbered 0 (least significant bit) to 31 (most significant bit). 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. If an EAG's length is more than 4 bytes, numbering for these additional bytes picks up where the previous byte left off. For example, the least significant bit in the fifth byte of an 8-byte EAG is referred to as bit 32.

2.3. Backward Compatibility

There are two questions 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 both an AG and EAG are present, a receiving node MUST use the AG as the first 32 bits (0-31) of administrative color and use the EAG for bits 32 and higher, if present.

A receiving node that notices that the AG differs from the first 32 bits of the EAG SHOULD report this mismatch to the operator.

This process allows nodes that do not support EAG to obtain some link color information from the network, while also allowing 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 may 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.

To allow for maximum interoperability, an implementation SHOULD treat desired but unadvertised EAG bits as if they were set to 0. Consider the case where a node wants to only use links where the 127th bit of an EAG is set to 1. If a link is only advertising 64 EAG bits, the setting of the 127th EAG bit is not known -- that is, it is neither explicitly 0 nor 1. The node that wants the 127th EAG bit to be 1 will not use this link when implementing the recommended behavior, as the assumption is that the unadvertised 127th bit is set to 0.

That said, each implementation makes its own choices based on necessary constraints, and there might be reasons to provide other strategies for handling this case. A strategy that deviates from the behavior this document recommends SHOULD be configurable to use the recommended behavior, in order to provide maximum interoperability.

3. Security Considerations

This extension adds no new security considerations.

4. IANA Considerations

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

For OSPF, the subregistry is the "Types for sub-TLVs of TE Link TLV (Value 2)" in the "Open Shortest Path First (OSPF) Traffic Engineering TLVs" registry.

For IS-IS, it is "Sub-TLVs for TLV 22, 141, and 222" subregistry in the "IS-IS TLV Codepoints" registry. For IS-IS, the value should be marked 'y' for Sub-TLVs 22, 141 and 222; this is identical to the allocation for the Administrative Group sub-TLV (value 3 in the same subregistry).

The assigned value from the OSPF registry is 26 and the assigned value from the IS-IS registry is 14. The sub-TLV is called "Extended Administrative Group".

5. Acknowledgements

Thanks to Santiago Alvarez, Rohit Gupta, Liem Nguyen, Tarek Saad, Robert Sawaya, Andy Malis, Les Ginsberg and Adrian Farrel for their review and comments.

6. References

6.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [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.
- [RFC5440] Vasseur, JP., Ed., and JL. Le Roux, Ed., "Path Computation Element (PCE) Communication Protocol (PCEP)", RFC 5440, March 2009.

6.2. Informative References

- [RFC2702] Awduche, D., Malcolm, J., Agogbua, J., O'Dell, M., and J. McManus, "Requirements for Traffic Engineering Over MPLS", RFC 2702, September 1999.

Author's Address

Eric Osborne

E-Mail: eric.osborne@notcom.com

