

**Extended Administrative Groups in MPLS-TE
draft-ietf-mpls-extended-admin-group-05**

Abstract

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.

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 October 23, 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
2.	Extended Administrative Groups sub-TLV	3
2.1.	Packet Format	4
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	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
	Author's Address	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 a fixed-length 32-bit bitmask. This can be quite constraining, as it is possible to run out of vaues 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,

Osborne

Expires October 23, 2014

[Page 2]

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, nevermind 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 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.

EAG's intended use case is within a single domain. As such, this document provides no support for signaling EAG. It provides no analog to either the SESSION_ATTRIBUTE of C-Type 1 defined in [\[RFC3209\]](#), nor the LSPA object of the Path Computation Element Communication Protocol (PCEP) defined in [\[RFC5440\]](#).

2. Extended Administrative Groups sub-TLV

This document defines a sub-TLV of the Link TLV for both OSPF [\[RFC3630\]](#) and ISIS [\[RFC5305\]](#) called the Extended Administrative Groups (EAG) sub-TLV. The EAG 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. 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.

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 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 allow for maximum interoperability an implementation **SHOULD** treat desired but unadvertised EAG bits as if they are 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, clearly the 127th EAG bit is not defined - that is, it is neither explicitly 0 nor 1. The node which wants the 127th EAG bit to be 1 **MUST NOT** use this link when implementing the recommended behavior, as the assumption is that an unadvertised bit is set to 0.

A node **MAY** provide other strategies for handling this case. A strategy which deviates from the recommended behavior in this document **SHOULD** be configurable, in order to provide maximum interoperability.

3. Security Considerations

This extension adds no new security considerations.

4. IANA Considerations

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

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

For ISIS, it is "Sub-TLVs for TLV 22, 141, and 222" 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 both registries the first free value should be assigned. As of this writing, that's 26 in the OSPF registry and 14 in the IS-IS registry. The Sub-TLV should be 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. and JL. Le Roux, "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

Email: eric.osborne@notcom.com