

Inter-Domain Routing
Internet-Draft
Intended status: Standards Track
Expires: August 20, 2015

P. Sarkar, Ed.
H. Gredler
Juniper Networks, Inc.
S. Litkowski
Orange
February 16, 2015

Advertising Per-node Admin Tags in BGP Link-State Advertisements
draft-psarkar-idr-bgp-ls-node-admin-tag-extension-00

Abstract

This document describes the protocol extensions to collect per-node administrative tags advertised in IGP Link State advertisements and disseminate the same in BGP Link-State advertisement protocol, to facilitate inter-AS TE applications that may need the same per-node administrative tags to associate a subset of network devices spanning across more than one AS with a specific functionality.

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 August 20, 2015.

Copyright Notice

Copyright (c) 2015 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. Per-Node Administrative Tag	4
3. BGP-LS Extensions for Per-Node Administrative Tags	4
3.1. Node Admin Tag TLV	5
4. Elements of Procedure	7
5. Applications	8
6. IANA Considerations	9
7. Manageability Considerations	9
7.1. Operational Considerations	9
7.1.1. Operations	9
8. TLV/Sub-TLV Code Points Summary	9
9. Security Considerations	9
10. Acknowledgements	9
11. References	9
11.1. Normative References	9
11.2. Informative References	10
11.3. URIs	10
Authors' Addresses	11

1. Introduction

Advertising Per-node Administrative Tags in Link State protocols like IS-IS [I-D.ietf-isis-node-admin-tag] and OSPF [I-D.ietf-ospf-node-admin-tag] allows adding an optional operational capability, that allows tagging and grouping of the nodes in a IGP domain. This, among other applications, allows simple management and easy control over route and path selection, based on local configured policies. However per-node administrative tags advertised in IGP advertisements let network operators associate nodes within a single AS (if not a single area). This limits the use of such per-node administrative tags and applications that need to associate a subset of network devices spanning across multiple AS with a specific functionality cannot use them.

To address the need for applications that require visibility into LSDB across IGP areas, or even across ASes, the BGP-LS address-

family/sub-address-family have been defined that allows BGP to carry LSDB information. The BGP Network Layer Reachability Information (NLRI) encoding format for BGP-LS and a new BGP Path Attribute called BGP-LS attribute are defined in [I-D.ietf-idr-ls-distribution]. The identifying key of each LSDB object, namely a node, a link or a prefix, is encoded in the NLRI and the properties of the object are encoded in the BGP-LS attribute. Figure 1 describes a typical deployment scenario. In each IGP area, one or more nodes are configured with BGP-LS. These BGP speakers form an IBGP mesh by connecting to one or more route-reflectors. This way, all BGP speakers - specifically the route-reflectors - obtain LSDB information from all IGP areas (and from other ASes from EBGP peers). An external component connects to the route-reflector to obtain this information (perhaps moderated by a policy regarding what information is sent to the external component, and what information isn't).

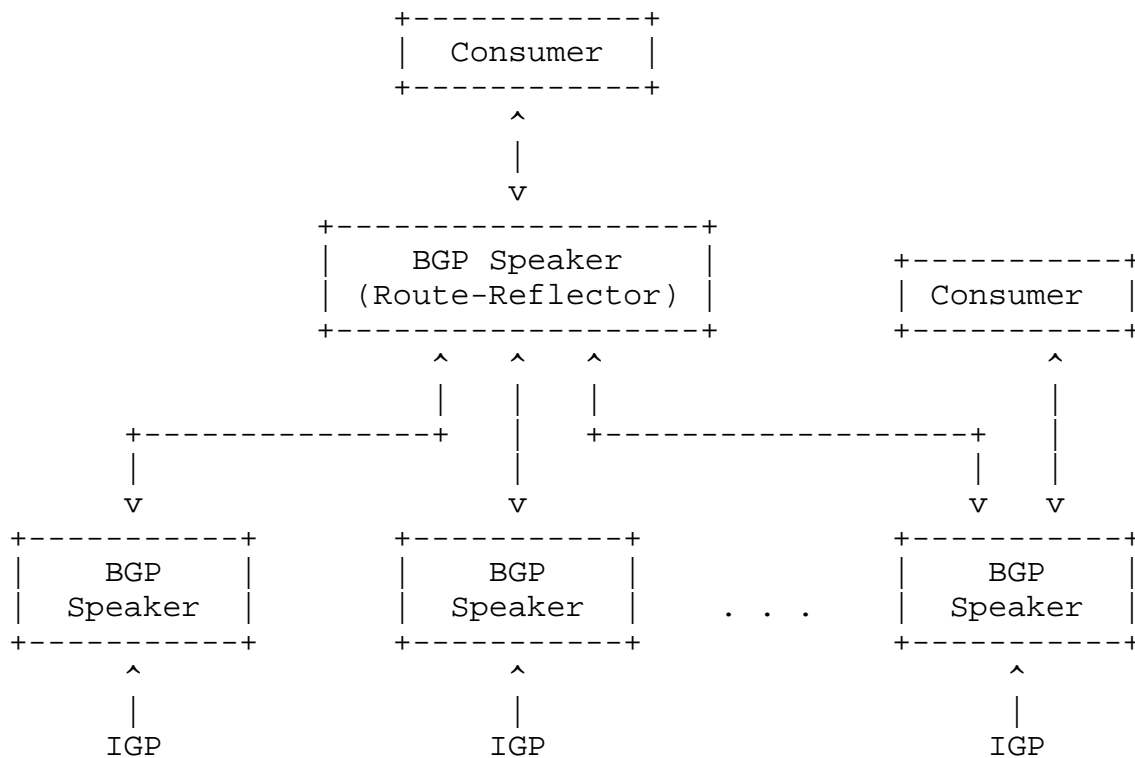


Figure 1: Link State info collection

For the purpose of advertising per-node administrative tags within BGP Link-State advertisements, a new Node Attribute TLV to be carried in the corresponding BGP-LS Node NLRI is proposed. For more details on the Node Attribute TLVs please refer to section 3.3.1 in [I-D.ietf-idr-ls-distribution]

2. Per-Node Administrative Tag

An administrative Tag is a 32-bit integer value that can be used to identify a group of nodes in the entire routing domain. The new sub-TLV specifies one or more administrative tag values. A BGP Link-State speaker that also participates in the IGP link state advertisements exchange may learn one or more per-node administrative tags advertised by another router in the same IGP domain. Such BGP-LS speaker shall encode the same set of per-node administrative tags in the corresponding Node Attribute TLV representing the network device that originated the per-node administrative tags.

The per-node administrative tags advertised in IGP link state advertisements will have either per-area (or levels in IS-IS) scope or 'global' scope. Operator may choose to a set of per-node administrative tags across areas (or levels in IS-IS) and another advertise set of per-node administrative tags within the specific area (or level). But evidently two areas within the same AS or two different may use the same per-node administrative tag for different purposes. In such case applications will need to distinguish between the per-area (or level) scoped administrative tags originated from a specific node against those originated from the same node with 'global' scope.

A BGP-LS router in a given AS while copying the per-node administrative tags learnt from IGP link-state advertisements, MUST also copy the scope associated with the per-node administrative tags. Refer to Section 3.1 for how to encode the associated scope of a per-node administrative tags as well.

To be able to distinguish between the significance of a per-area (or level) administrative tag learnt in one area, from that advertised in another area, or another AS, any applications receiving such a BGP-LS advertisements MUST consider the scope associated with each per-node administrative tag with 'per-area (or per-level) along with the area (or level in IS-IS) associated with corresponding IGP link state advertisement and the AS number associated with the originating node. The area (or level) associated with corresponding IGP link state advertisement and the AS number associated with the originating node can be derived from appropriate node attributes (already defined in BGP-LS [I-D.ietf-idr-ls-distribution]) attached with the corresponding Node NLRI.

3. BGP-LS Extensions for Per-Node Administrative Tags

The BGP-LS NLRI can be a node NLRI, a link NLRI or a prefix NLRI. The corresponding BGP-LS attribute is a node attribute, a link attribute or a prefix attribute. BGP-LS

[I-D.ietf-idr-ls-distribution] defines the TLVs that map link-state information to BGP-LS NLRI and BGP-LS attribute. This document adds an new Node Attribute TLV called 'Node Admin Tag TLV' to encode per-node administrative tags information.


[I-D.ietf-isis-node-admin-tag] defines the 'Per-node Admin Tag' sub-TLV in the Router Capability TLV (type 242) in IS-IS Link State PDUs to encode per-node administrative tags. Similarly [I-D.ietf-isis-node-admin-tag] defines the 'Per-node Administrative Tag' TLV in OSPF Router Information LSAs to encode per-node administrative tags in OSPF Link State update packets. The per-node administrative tags TLVs learnt from the IGP link state advertisements of a specific node will all be inserted in a new Node Admin Tag TLV and added to the corresponding Node are mapped to the corresponding BGP-LS Node NLRI. Per-node administrative tags from IGP advertisements are mapped to the corresponding Node Admin Tag TLV in the following way.

TLV Code Point	Description	Length	IS-IS TLV /sub-TLV	OSPF LSA/TLV
TBD	Node Admin Tag TLV	Variable	242/TBD [1]	RI-LSA/TBD [2]

Table 1: Node Admin Tag TLV Mapping from IGP

3.1. Node Admin Tag TLV

The new Node Administrative Tag TLV, like other BGP-LS Node Attribute TLVs, is formatted as Type/Length/Value (TLV) triplets. Figure 2 below shows the format of the new TLV.

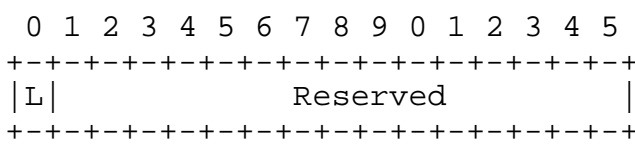


Type : A 2-octet field specifying code-point of the new sub-TLV type. Code-point: TBA (suggested 1040)

Length: A 2-octet field that indicates the length of the value portion in octets and will be a multiple of 4 octets dependent on the number of tags advertised.

Value: A 2-octet 'Flags' field, followed by a sequence of multiple 4 octets defining the administrative tags.

Flags: A 2-octet field that carries flags associated with all the administrative flags encoded in this TLV. Following is the format of this field.



The following bit flags are defined:

L bit : If the L bit is set (1), it signifies that all administrative flags encoded in this TLV has per-area(or level in IS-IS) scope, and should not be mixed with ones with same value but with 'global' scope (L bit reset to 0).

Figure 2: BGP Link-State Node Administrative Tag TLV

This new type of 'Node Admin Tag' TLVs can ONLY be added to the Node Attribute associated with the Node NLRI that originates the corresponding per-node administrative tags in IGP domain.

All the per-node administrative tags with 'per-area' (or per-level) scope, originated by a single node in IGP domain SHALL be re-originated in a single 'Node Admin Tag' TLV and inserted in the Node NLRI generated for the same node. Similarly, all the per-node administrative tags with 'global' scope originated by the same node in IGP domain SHALL be re-originated in another 'Node Admin Tag' TLV and inserted in the same Node NLRI generated for the originating node. Multiple instances of a TLV may be generated by the BGP-LS router for a given node in the IGP domain. This MAY happen if the original node's link state advertisement carries more than 16383 per-node administrative groups and a single TLV does not provide sufficient space. As such multiple occurrence of the 'Node Admin Tag' TLVs under a single BGP LS NLRI is cumulative.

While copying per-node administrative tags from IGP link-state advertisements to corresponding BGP-LS advertisements, the said BGP-LS speaker MAY run all the per-node administrative flags through a locally configured policy that selects which ones should be exported and which ones not. And then the per-node administrative tag is copied to the BGP-LS advertisement if it is permitted to do so by the said policy.

4. Elements of Procedure

Meaning of the Per-node administrative tags is generally opaque to BGP Link-State protocol. Router advertising the per-node administrative tag (or tags) may be configured to do so without knowing (or even explicitly supporting) functionality implied by the tag.

Interpretation of tag values is specific to the administrative domain of a particular network operator. The meaning of a per-node administrative tag is defined by the network local policy and is controlled via the configuration. However multiple administrative domain owners may agree on a common meaning implied by a administrative tag for mutual benefit.

The semantics of the tag order has no meaning. There is no implied meaning to the ordering of the tags that indicates a certain operation or set of operations that need to be performed based on the ordering.

Each tag SHOULD be treated as an independent identifier that MAY be used in policy to perform a policy action. Per-node administrative

tags carried by the Node Admin Tag TLV SHOULD be used to indicate a independent characteristics of the node in IGP domain that originated it. The TLV SHOULD be considered as an unordered list. Whilst policies may be implemented based on the presence of multiple tags (e.g., if tag A AND tag B are present), they MUST NOT be reliant upon the order of the tags (i.e., all policies should be considered commutative operations, such that tag A preceding or following tag B does not change their outcome).

For more details on guidance on usage of per-node administrative tags please refer to section 4 [3] in [I-D.ietf-isis-node-admin-tag].

5. Applications

[I-D.ietf-isis-node-admin-tag] and [I-D.ietf-ospf-node-admin-tag] present some applications of node administrative tags.

The policy-based Explicit routing use case can be extended to inter-area or inter-AS scenarios where an end to end path needs to avoid or include nodes that have particular properties. Following are some examples.

1. Geopolitical routing : preventing traffic from country A to country B to cross country C. In this case, we may use node administrative tags to encode geographical information (country). Path computation will be required to take into account node administrative tag to permit avoidance of nodes belonging to country C.
2. Legacy node avoidance : in some specific cases, it is interesting for service-provider to force some traffic to avoid legacy nodes in the network. For example, legacy nodes may not be carrier class (no high availability), and service provider wants to ensure that critical traffic only uses nodes that are providing high availability.

In case of inter-AS Traffic-Engineering applications, different ASes SHOULD share their admin tag policies. They MAY also need to agree upon some common tagging policy for specific applications.

For more details on some possible applications with per-node administrative tags please refer to section 5 [4] in [I-D.ietf-isis-node-admin-tag].

6. IANA Considerations

This document requests assigning code-points from the registry for BGP-LS attribute TLVs based on table Table 2.

7. Manageability Considerations

This section is structured as recommended in [RFC5706].

7.1. Operational Considerations

7.1.1. Operations

Existing BGP and BGP-LS operational procedures apply. No new operation procedures are defined in this document.

8. TLV/Sub-TLV Code Points Summary

This section contains the global table of all TLVs/Sub-TLVs defined in this document.

TLV Code Point	Description	Length
1040	Node Admin Tag	variable

Table 2: Summary Table of TLV/Sub-TLV Codepoints

9. Security Considerations

Procedures and protocol extensions defined in this document do not affect the BGP security model. See the 'Security Considerations' section of [RFC4271] for a discussion of BGP security. Also refer to [RFC4272] and [RFC6952] for analysis of security issues for BGP.

10. Acknowledgements

TBD.

11. References

11.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC4271] Rekhter, Y., Li, T., and S. Hares, "A Border Gateway Protocol 4 (BGP-4)", RFC 4271, January 2006.

11.2. Informative References

[I-D.ietf-idr-ls-distribution]

Gredler, H., Medved, J., Previdi, S., Farrel, A., and S. Ray, "North-Bound Distribution of Link-State and TE Information using BGP", draft-ietf-idr-ls-distribution-07 (work in progress), November 2014.

[I-D.ietf-isis-node-admin-tag]

Sarkar, P., Gredler, H., Hegde, S., Litkowski, S., Decraene, B., Li, Z., Aries, E., Rodriguez, R., and H. Raghuvver, "Advertising Per-node Admin Tags in IS-IS", draft-ietf-isis-node-admin-tag-00 (work in progress), December 2014.

[I-D.ietf-ospf-node-admin-tag]

Hegde, S., Raghuvver, H., Gredler, H., Shakir, R., Smirnov, A., Li, Z., and B. Decraene, "Advertising per-node administrative tags in OSPF", draft-ietf-ospf-node-admin-tag-00 (work in progress), October 2014.

[RFC4272] Murphy, S., "BGP Security Vulnerabilities Analysis", RFC 4272, January 2006.

[RFC5706] Harrington, D., "Guidelines for Considering Operations and Management of New Protocols and Protocol Extensions", RFC 5706, November 2009.

[RFC6952] Jethanandani, M., Patel, K., and L. Zheng, "Analysis of BGP, LDP, PCEP, and MSDP Issues According to the Keying and Authentication for Routing Protocols (KARP) Design Guide", RFC 6952, May 2013.

11.3. URIs

[1] <http://tools.ietf.org/html/draft-ietf-isis-node-admin-tag-00#section-3.1>

[2] <http://tools.ietf.org/html/draft-ietf-ospf-node-admin-tag-00#section-4.1>

[3] <http://tools.ietf.org/html/draft-ietf-isis-node-admin-tag-00#section-4>

[4] <http://tools.ietf.org/html/draft-ietf-isis-node-admin-tag-00#section-5>

Authors' Addresses

Pushpasis Sarkar (editor)
Juniper Networks, Inc.
Electra, Exora Business Park
Bangalore, KA 560103
India

Email: psarkar@juniper.net

Hannes Gredler
Juniper Networks, Inc.
1194 N. Mathilda Ave.
Sunnyvale, CA 94089
US

Email: hannes@juniper.net

Stephane Litkowski
Orange

Email: stephane.litkowski@orange.com