

Network Working Group	Hamid Ould-Brahim (Nortel Networks)
Internet Draft	Don Fedyk (Nortel Networks)
Expiration Date: June 2009	Yakov Rekhter (Juniper Networks)
Intended Status: Proposed Standard	

BGP Traffic Engineering Attribute

[draft-ietf-softwire-bgp-te-attribute-04.txt](#)

Status of this Memo

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

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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".

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/1id-abstracts.txt>

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

Copyright (c) 2008 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.

Abstract

This document defines a new BGP attribute, Traffic Engineering attribute, that enables BGP to carry Traffic Engineering information.

The scope and applicability of this attribute currently excludes its use for non-VPN reachability information.

1. Specification of Requirements

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. Introduction

In certain cases (e.g., L1VPN [[RFC5195](#)]) it may be useful to augment VPN reachability information carried in BGP with the Traffic Engineering information.

This document defines a new BGP attribute, Traffic Engineering attribute, that enables BGP [[RFC4271](#)] to carry Traffic Engineering information.

[Section 4 of \[RFC5195\]](#) describes one possible usage of this attribute.

The scope and applicability of this attribute currently excludes its use for non-VPN reachability information.

Procedures for modifying the Traffic Engineering attribute, when re-advertising a route that carries such attribute are outside the scope of this document.

3. Traffic Engineering Attribute

Traffic Engineering attribute is an optional non-transitive BGP attribute.

The information carried in this attribute is identical to what is carried in the Interface Switching Capability Descriptor, as specified in [[RFC4203](#)], [[RFC5307](#)].

The attribute contains one or more of the following:

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			
+--+			
Switching Cap	Encoding		Reserved
+--+			
	Max LSP Bandwidth at priority 0		

```

+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               Max LSP Bandwidth at priority 1                               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               Max LSP Bandwidth at priority 2                               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               Max LSP Bandwidth at priority 3                               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               Max LSP Bandwidth at priority 4                               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               Max LSP Bandwidth at priority 5                               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               Max LSP Bandwidth at priority 6                               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               Max LSP Bandwidth at priority 7                               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|          Switching Capability-specific information          |
|          (variable)          |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

The Switching Capability (Switching Cap) field contains one of the values specified in [Section 3.1.1 of \[RFC3471\]](#).

The Encoding field contains one of the values specified in [Section 3.1.1 of \[RFC3471\]](#).

The Reserved field SHOULD be set to 0 on transmit and MUST be ignored on receive.

Maximum LSP Bandwidth is encoded as a list of eight 4 octet fields in the IEEE floating point format [\[IEEE\]](#), with priority 0 first and priority 7 last. The units are bytes (not bits!) per second.

The content of the Switching Capability specific information field depends on the value of the Switching Capability field.

When the Switching Capability field is PSC-1, PSC-2, PSC-3, or PSC-4, the Switching Capability specific information field includes Minimum LSP Bandwidth and Interface MTU.

```


      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
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               Minimum LSP Bandwidth                               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               Interface MTU                               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

The Minimum LSP Bandwidth is encoded in a 4 octet field in the IEEE floating point format. The units are bytes (not bits!) per second. The Interface MTU is encoded as a 2 octet integer.

When the Switching Capability field is L2SC, there is no Switching Capability specific information field present.

When the Switching Capability field is TDM, the Switching Capability specific information field includes Minimum LSP Bandwidth and an indication of whether the interface supports Standard or Arbitrary SONET/SDH.



The Minimum LSP Bandwidth is encoded in a 4 octet field in the IEEE floating point format. The units are bytes (not bits!) per second. The indication of whether the interface supports Standard or Arbitrary SONET/SDH is encoded as 1 octet. The value of this octet is 0 if the interface supports Standard SONET/SDH, and 1 if the interface supports Arbitrary SONET/SDH.

When the Switching Capability field is LSC, there is no Switching Capability specific information field present.

4. Implication on aggregation

Routes that carry the Traffic Engineering Attribute have additional semantics that could affect traffic forwarding behavior. Therefore, such routes SHALL NOT be aggregated unless they share identical Traffic Engineering Attributes.

Constructing the Traffic Engineering Attribute when aggregating routes with identical Traffic Engineering attributes follows the procedure of [RFC4201].

5. Implication on scalability

The use of the Traffic Engineering Attribute does not increase the number of routes, but may increase the number of BGP Update messages required to distribute the routes depending on whether these routes share the same BGP Traffic Engineering attribute or not (see below).

When the routes differ in other than the Traffic Engineering Attribute (e.g., differ in the set of Route Targets, and/or NEXT_HOP), use of Traffic Engineering Attribute has no impact on the number of BGP Update messages required to carry the routes. There is also no impact when routes share all other attribute information and have an aggregated or identical Traffic Engineering Attribute. When

routes share all other attribute information and have different Traffic Engineering Attributes, routes must be distributed in per-route BGP Update messages rather than a single message.

6. IANA Considerations

This document defines a new BGP attribute. This attribute is optional and non-transitive.

7. Security Considerations

This extension to BGP does not change the underlying security issues currently inherent in BGP. BGP security considerations are discussed in [RFC 4271](#)

8. Acknowledgements

The authors would like to thank John Scudder and Jeffrey Haas for their review and comments.

9. Normative References

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

[RFC4201] Kompella, K., Rekhter, Y., Berger, L., "Link Bundling in MPLS Traffic Engineering (TE)", [RFC 4201](#), October 2005<P>

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

[RFC3471] Berger, L., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Functional Description", [RFC 3471](#), January 2003.

[IEEE] IEEE, "IEEE Standard for Binary Floating-Point Arithmetic", Standard 754-1985, 1985 (ISBN 1-5593-7653-8).

10. Non-Normative References

[RFC4203] Kompella, K., Rekhter, Y., "OSPF Extensions in Support of

Generalized Multi-Protocol Label Switching (GMPLS)", [RFC4203](#), October 2005

[RFC5307] Kompella, K., Rekhter, Y., "Intermediate System to Intermediate System (IS-IS) Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)", [RFC5307](#), October 2005

[RFC5195] Ould-Brahim, H., Fedyk, D., Rekhter, Y., "BGP-Based Auto-Discovery for Layer-1 VPNs", [RFC5195](#), June 2008

11. Author Information

Hamid Ould-Brahim
Nortel Networks
Email: hbrahim@nortel.com

Don Fedyk
Nortel Networks
Email: dwfedyk@nortel.com

Yakov Rekhter
Juniper Networks, Inc.
email: yakov@juniper.com