Network Working Group
Internet Draft
Expires December 2000

Cristallo, De Clercq
Alcatel

June 2002

BGP Tunnel Attribute
<draft-cristallo-bgp-tunnel-attr-00.txt>

Status of this Memo

2. Abstract

This document proposes the Tunnel Attribute, which may be used by a given BGP Speaker to indicate whether it expects to receive all, some, or none of its traffic through a tunnel and the types of tunnel that may be used.

3. Introduction

Several IETF proposals require the establishement of tunnels ([NGTRANS-BGP], [PPVPN]). There is currently no mechanism for the automatic configuration of these tunnels while the drafts [NGTRANS-BGP] and [IPsec-2547], for example, require the use of a BGP attribute to achieve a similar behavior. The above mentionned proposals being all based on BGP, the BGP protocol appears as a key component for the automatic establishement and configuration of these tunnels. The aim of this draft is to provide a BGP-based mechanism for the automatic configuration and establishment of tunnels. This draft proposes a new BGP attribute, the Tunnel Attribute, that will be used by BGP speakers to specify the types of tunnel that may be used for the transmission of traffic towards certain destination.

4. Tunnel Attribute

The TUNNEL Attribute is an optional non-transitive attribute which conveys tunneling-related information associated to the routes described in the NLRI field or in the MP_REACH_NLRI field [BGP-MP] of the BGP Update message. The Type Code of the Tunnel Attribute is TBD_IANA and is subject to the IANA considerations section of this draft.

The TUNNEL Attribute contains one or several Tunnel Type values, each encoded on one octet. The Tunnel Type specifies the type of tunnel that may be used by the BGP peer that receives the route, in order to transfer traffic toward the set of destinations specified in the NLRI field or in the MP_REACH_NLRI field of the Update message. When several Tunnel Type values are specified, the receiver of the message chooses one of these values for the transmission of the traffic. The following types have been identified so far:

- 0x00: No Tunnel Type specified
- 0x01: IP in IP
- 0x02: GRE
- 0x03: IPSec
- 0x04: MPLS
- 0x05: L2TP

Tunnel Type 0x00 SHOULD not be included in a list with other types different from 0x00. If a list contains type 0x00 and other types, the receiver SHOULD ignore the other types.

5. Operation

A given BGP speaker may use the Tunnel Attribute to indicate whether it expects to receive all, some, or none of its traffic through a tunnel and to specify the corresponding encapsulation that may be used for the transmission of traffic towards the set of destinations. The Tunnel Attribute May contain one or several Tunnel Types, each encoded on one octet.

A BGP speaker receiving a route that does not have the Tunnel Attribute MAY add this attribute to the Update Message when propagating it to its peers. In this case, tunnels will only be established between itself and its peers.

A BGP speaker receiving a route with the Tunnel Attribute which does not recognize the attribute MUST NOT install the route (Normal processing of an unrecognized non-transitive attribte [BGP4]). The reason is that the originator of the route wants to receive its traffic only via one of the specified types of tunnel.

A BGP speaker receiving a route with a recognized Tunnel Attribute may delete, modify or leave it unchanged when propagating the Update

Message to its peers.

In addition, a BGP speaker that recognizes the Tunnel Attribute in a route received from a peer may choose any particular value from the list of Tunnel Types specified by the peer, and establish a tunnel of

that type for transmission of traffic toward these destinations. As explained in section 6 hereafter, thanks to the negociation of capabilities [BGP-CAP], the list of tunnel types specified by the peer contains only types that are supported by both peers. The tunnel endpoint MUST be the BGP NEXT_HOP specified in the Update message. Several cases are possible:

- Tunnel Type is 0x00: The originator of the route wants to receive its traffic through a tunnel, but does not want to specify the type, or just means that it supports any type of Tunnel. Actually, in this case, the choice of the value of the Tunnel Type is made by the receiver of the BGP Update message with the Tunnel Attribute. The receiver will pick one of the values that are supported by both speakers, as negociated at connection setup (see Section 6).

- The list of Tunnel Types contains one or several types different from 0x00: The receiver chooses one of the specified values, establishes a tunnel of that type and uses the corresponding encapsulation for transmission of traffic towards the specified set of destinations.

If a BGP speaker receives an Update message without the Tunnel Attribute, the BGP speaker should send the traffic to the announced destinations according to a configured or default encapsulation.

6. Use of Capabilities Advertisement with BGP-4

A BGP speaker that uses the Tunnel Attribute SHOULD use the Capabilities Advertisement procedures, as defined in [BGP-CAP], so that it might be able to determine if it can use such an attribute with a particular peer.

The use and meaning of these fields are as follows:

- The Capability Code is a one octet field that unambiguously identifies individual capabilities.

- The Capability Length is a one octet field that contains the length of the Capability Value field in octets.

- The Capability Value field is a variable length field that contains one or several Tunnel Type values, as defined in section 4 of this draft, each encoded on one octet. These are the tunnel types that the peer support a-priori. The use and meaning of this field is different from the Tunnel Attribute. The later specifies for which destinations a particular type of tunnel should be used, while the Capabilities Attribute specifies a set of Tunnel Types that the

peer recognizes and supports.

If a BGP speaker that supports the Tunnel Attribute determines that its peer doesn't, the speaker may send a NOTIFICATION message to the peer, and terminate peering. The Error Subcode in the message is set

to Unsupported Capability.

If a BGP speaker that supports the Tunnel Attribute determines that its peer also supports it, the speaker will check the Tunnel Types specified in the Value field of the Capabilities Attribute. Only the Tunnel Types that are supported by both speaker could be used later on when the connection will be established. In case no Tunnel Type is supported by both peers, the Tunnel Attribute should never be exchanged between the peers.

Thanks to this negociation phase, the list of Tunnel Types specified by a BGP speaker will only contain values that are also supported by the peer.

7. Examples

- Example 1: Consider the simple example below. Site A and Site D are from the same corporation and make use of private addresses, while Sites B and E make use of public addresses. For this reason, traffic between site A and site D must be sent in a Tunnel, a GRE Tunnel for example. R3 sends to R4 routing information for Site A and Site B. R3 adds a Tunnel Attribute with value 0x02 in the Update Message sent to R4 and relative to Site A. R3 will not add the Tunnel Attribute in the message relative to Site B. R4 MUST encapsulate traffic toward Site A in an GRE Tunnel and SHOULD forward traffic to Site B in a configured or default encapsulation. The behavior for the traffic in the reverse direction is analogu-ous.

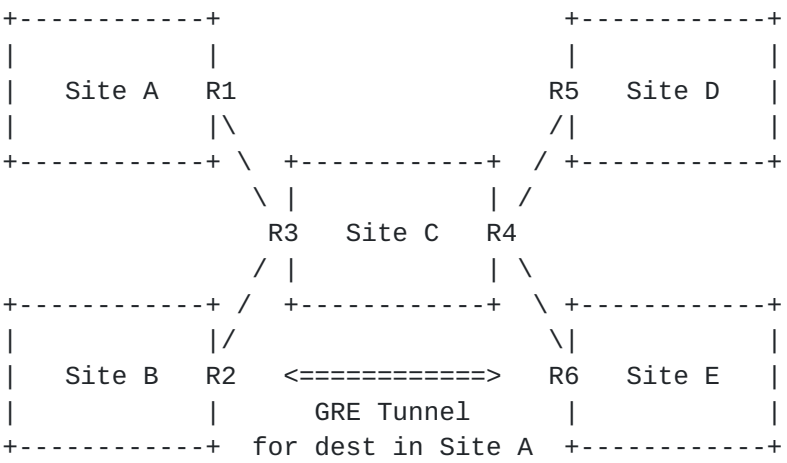

- Example 2: In [NGTRANS-BGP], BGP speakers announce IPv6 routes to each other with BGP-MP [BGP-MP], over an IPv4 network. The IPv6 traffic needs to be encapsulated first before traversing the IPv4 network. Normally the used encapsulation must be configured at both

BGP speakers. The Tunnel Attribute proposed in this draft allows for signalling and thus less configuration effort.

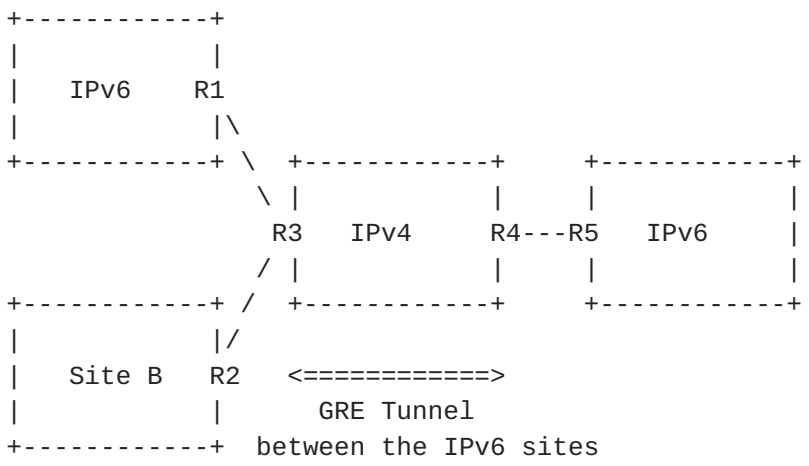

- Example 3: In [IPsec-2547], BGP peering PE routers that normally use MPLS to forward VPN traffic to each-other, require the possiblity to signal to each other that certain traffic should be tunneled through IPSec instead of through MPLS.

8. IANA Considerations

The Type Code of the TUNNEL Attribute (Section 4) will be assigned by IANA.

9. Security Considerations

This additional BGP4 attribute specification does not change the underlying security issues inherent in the existing BGP-4 protocol specification [RFC2385].

10. Acknowledgements

The authors would like to thanks the authors of the [NGTRANS-BGP] and [IPsec-2547] drafts, who first suggested the use of a BGP Tunnel type attribute.

11. References

[BGP4] "A Border Gateway Protocol 4 (BGP-4)", RFC 1771, Y. Rekhter and T. Li, March 1995.

[BGP-MP] "Multiprotocol Extensions for BGP-4", RFC 2858, T. Bates, R. Chandra, D. Katz and Y. Rekhter, June 2000.

[NGTRANS-BGP] "Connecting IPv6 Islands across IPv4 Clouds with BGP", draft-ietf-ngtrans-bgp-tunnel-04.txt, Work In Progress, De J.

Clercq, G. Gastaud, T. Nguyen, D. Ooms, S. Prevost and F. Le Faucheur, January, 2002

[IPsec-2547]
"Use of PE-PE IPsec in RFC2547 VPNs", draft-ietf-ppvpn-ipsec-

2547-01.txt, Work In Progress, E. Rosen, J. De Clercq, O. Paridaens, Y. T'Joens and C. Sargor, February 2002

[PPVPN] "Using BGP as an Auto-Discovery Mechanism for Network-based VPNs", draft-ietf-ppvpn-bgpvpn-auto-02.txt, Work In Progress, Hamid Ould-Brahim, Bryan Gleeson, Peter Ashwood-Smith, Eric C. Rosen, Yakov Rekhter, Luyuan Fang, Jeremy De Clercq, Riad Hartani, Tissa Senevirathne, January 2002

[RFC2385] "Protection of BGP sessions via the TCP MD5 Signature Option", RFC 2385, A. Heffernan, August 1998.

[BGP-CAP] "Capabilities Advertisement with BGP-4", RFC 2842, Chandra, R., Scudder, J., May 2000.

11. Authors Addresses

Cristallo Geoffrey
Alcatel
Fr. Wellesplein 1, 2018 Antwerp, Belgium
E-mail: geoffrey.cristallo@alcatel.be

Jeremy De Clercq
Alcatel
Fr. Wellesplein 1, 2018 Antwerpen, Belgium.
E-mail: Jeremy.De_Clercq@alcatel.be