

PWE3 Working Group
Internet Draft
Expires: April 2007

Chris Metz
Luca Martini
Cisco Systems

Florin Balus
Jeff Sugimoto
Nortel Networks

October 4, 2006

**Pseudowire Attachment Identifiers for Aggregation and VPN
Autodiscovery
draft-ietf-pwe3-aii-aggregate-01.txt**

Status of this Memo

By submitting this Internet-Draft, each author represents that any applicable patent or other IPR claims of which he or she is aware have been or will be disclosed, and any of which he or she becomes aware will be disclosed, in accordance with [Section 6 of BCP 79](#).

This document may only be posted in an Internet-Draft.

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>

This Internet-Draft will expire on April 4, 2006.

Copyright Notice

Copyright (C) The Internet Society (2006). All Rights Reserved.

Abstract

The signaling protocols used to establish point-to-point pseudowires include type-length-value (TLV) fields that identify pseudowire endpoints called attachment individual identifiers (AII). This document defines AII structures in the form of new AII type-length-value fields that support AII aggregation for improved scalability and VPN autodiscovery. It is envisioned that this would be useful in large inter-domain virtual private wire service networks where pseudowires are established between selected local and remote PE nodes based on customer need.

Conventions used in this document

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](#) Error! Reference source not found..

Table of Contents

1. Introduction.....	2
2. Structure for the New AII Type.....	4
2.1. AII Type 1.....	4
2.2. AII Type 2.....	4
3. IANA Considerations.....	6
4. Security Considerations.....	6
5. Acknowledgments.....	6
Author's Addresses.....	7
Intellectual Property Statement.....	8
Disclaimer of Validity.....	8
Copyright Statement.....	8
Acknowledgment.....	9

[1. Introduction](#)

[RFC4447] defines the signaling mechanisms for establishing point-to-point pseudowires (PWs) between two provider edge (PE) nodes. When a PW is set up, the LDP signaling messages include a forwarding equivalence class (FEC) element containing information about the PW type and an endpoint identifier used in the selection of the PW forwarder that binds the PW to the attachment circuit at each end.

There are two types of FEC elements defined for this purpose: PWid FEC (type 128) and the Generalized ID (GID) FEC (type 129). The PWid FEC element includes a fixed-length 32 bit value called the PWid that serves as an endpoint identifier. The same PWid value must be configured on the local and remote PE prior to PW setup.

The GID FEC element includes TLV fields for attachment individual identifiers (AII) that, in conjunction with an attachment group identifier (AGI), serve as PW endpoint identifiers. The endpoint identifier on the local PE (denoted as <AGI, source AII or SAI>) is called the source attachment identifier (SAI) and the endpoint identifier on the remote PE (denoted as <AGI, target AII or TAI>) is called the target attachment identifier (TAI). The SAI and TAI can be distinct values. This is useful for applications and provisioning models where the local PE (with a particular SAI) does not know and must somehow learn (e.g. via MP-BGP auto-discovery) of remote TAI values prior to launching PW setup messages towards the remote PE.

The use of the GID FEC TLV provides the flexibility to structure (source or target) AII values to best fit particular application or provisioning model needs [L2VPN-SIG]. For example an AII structure that enables many individual AII values to be identified as a single value could significantly reduce the burden on AII distribution mechanisms (e.g. MP-BGP) and on PE memory needed to store this AII information. It should be noted that PWE3 signaling messages will always include a fully qualified AII value.

An AII that is globally unique would facilitate PW management and security in large inter-AS and inter-provider environments. Providers would not have to worry about AII value overlap during provisioning or the need for AII network address translation (NAT) boxes during signaling. Globally unique AII values could aid in troubleshooting and could be subjected to source-validity checks during AII distribution and signaling. An AII automatically derived from a provider's existing IP address space can simplify the provisioning process.

This document defines an AII structure based on [[RFC4447](#)] that:

- o Enables many discrete attachment individual identifiers to be summarized into a single AII summary value. This will enhance scalability by reducing the burden on AII distribution mechanisms and on PE memory.

- o Ensures global uniqueness if desired by the provider. This will facilitate Internet-wide PW connectivity and provide a means for providers to perform source validation on the AII distribution (e.g. MP-BGP) and signaling (e.g. LDP) channels.

This is accomplished by defining new AII types and the associated formats of the value field.

2. Structure for the New AII Type

[RFC4447] defines the format of the GID FEC TLV and the use and semantics of the attachment group identifier (AGI).

2.1. AII Type 1

AII Type 1 has been allocated by IANA for use with provisioning models requiring a fixed-length 32-bit value [L2VPN-SIG]. This value is unique on the local PE.

2.2. AII Type 2

The AII Type 2 structure permits varying levels of AII summarization to take place thus reducing the scaling burden on the aforementioned AII distribution mechanisms and PE memory. In other words it no longer becomes necessary to distribute or configure all individual AII values (which could number in the tens of thousands or more) on local PEs prior to establishing PWs to remote PEs. The details of how and where the aggregation of AII values is performed and then distributed as AII reachability information are not discussed in this document.

AII Type 2 uses a combination of a provider's globally unique identifier (Global ID), a 32-bit prefix field and an optional 4-octet attachment circuit identifier field to create globally unique AII values.

The encoding of AII Type 2 is shown in figure 1.




Figure 1 AII Type 2 TLV Structure

- o AII Type = 0x02
- o Length = length of value field in octets. The length is set to 8 if the AC ID is NULL and 12 if the AC ID is non-null.
- o Global ID = This is a 4 octet field containing a value that is unique to the provider. The global ID can contain the 2 octet or 4 octet value of the provider's Autonomous System Number (ASN). It is expected that the global ID will be derived from the globally unique ASN of the autonomous system hosting the PEs containing the actual AIIs. The presence of a global ID based on the provider's ASN ensures that the AII will be globally unique.

If the PE hosting the AIIIs is present in an autonomous system where the provider is not running BGP, chooses not to expose this information or does not wish to use the global ID, then the global ID field **MUST** be set to zero. If the global ID is derived from a 2-octet AS number, then the high-order 4 octets of this 4 octet field **MUST** be set to zero.

Please note that the use of the provider's ASN as a global ID DOES NOT have anything at all to do with the use of the ASN in protocols such as BGP.

- o Prefix = The 32-bit prefix is a value assigned by the provider or it can be automatically derived from the PE's /32 IPv4 loopback address. Note that it is not required that the 32-bit prefix have any association with the IPv4 address space used in the provider's IGP or BGP for IP reachability.
- o Attachment Circuit (AC) ID = This is a fixed length four octet field used to further refine identification of an attachment circuit on the PE. The inclusion of the AC ID is used to identify individual attachment circuits that share a common prefix.

If the AC ID is not present then the AC ID field MUST be null and the AII Length field is set to 8. If the AC ID is present then the length field is set to 12 octets.

3. IANA Considerations

This document requests that IANA allocate a value from the "Attachment Individual Identifier (AII) Type" registry defined in [\[RFC4446\]](#).

The suggested value for this AII type is 0x02.

4. Security Considerations

AII values appear in AII distribution protocols [MP-BGP-AUTO-DISC] and PW signaling protocols [\[RFC4447\]](#) and are subject to various authentication schemes (i.e. MD5) if so desired.

The use of global ID values (e.g. ASN) in the inter-provider case could enable a form of source-validation checking to ensure that the AII value (aggregated or explicit) originated from a legitimate source.

5. Acknowledgments

Thanks to Carlos Pignataro, Scott Brim, Skip Booth, George Swallow and Bruce Davie for their input into this document.

Normative References

[[RFC4447](#)], Pseudowire Setup and Maintenance Using the Label Distribution Protocol (LDP) , [RFC4447](#), April 2006

[[RFC4446](#)], "IANA Allocations for Pseudowire Edge to Edge Emulation (PWE3)", [RFC4446](#), April 2006

[L2VPN-SIG], Provisioning, Autodiscovery, and Signaling in L2VPNs , [draft-ietf-l2vpn-signaling-08.txt](#), B. Davie, et al., May 2006

[MP-BGP-AUTO-DISC], Using BGP as an Auto-Discovery Mechanism for Layer-3 and Layer-2 VPNs , Ould-Brahim, H. et al, [draft-ietf-l3vpn-bgpvpn-auto-06.txt](#), June 2005

Author's Addresses

Chris Metz
Cisco Systems, Inc.
3700 Cisco Way
San Jose, Ca. 95134
Email: chmetz@cisco.com

Luca Martini
Cisco Systems, Inc.
9155 East Nichols Avenue, Suite 400
Englewood, CO, 80112
Email: lmartini@cisco.com

Florin Balus
Nortel
3500 Carling Ave.
Ottawa, Ontario, CANADA
Email: balus@nortel.com

Jeff Sugimoto
Nortel

3500 Carling Ave.

Ottawa, Ontario, CANADA

Email: sugimoto@nortel.com

Intellectual Property Statement

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in [BCP 78](#) and [BCP 79](#).

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org

Disclaimer of Validity

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Copyright Statement

Copyright (C) The Internet Society (2006).

This document is subject to the rights, licenses and restrictions contained in [BCP 78](#), and except as set forth therein, the authors retain all their rights.

Acknowledgment

Funding for the RFC Editor function is currently provided by the Internet Society.