

Network Working Group
Internet-Draft
Expires: December 11, 2006

Shepherd
Farinacci
Cisco Systems
Wu
Li
Cernet
June 9, 2006

**IPv4 unicast/multicast VPNs over an IPv6 core
draft-ietf-softwire-4over6vpns-00**

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](#).

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 December 11, 2006.

Copyright Notice

Copyright (C) The Internet Society (2006).

Abstract

This document describes a method by which a Service Provider with an IPv6 backbone may provide VPNs (Virtual Private Networks) and MVPNs (Multicast Virtual Private Networks) for its IPv4 customers. The IPv6 core network need only deploy native multicast services using Protocol Independent Multicast (PIM). All additional functionality

described is Customer Edge (CE) based and there are no additional Provider (P) or Provider Edge (PE) protocols.

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](#)].

1. Introduction

Current PE based VPN solutions continue to overload functional and scaling requirements onto the PE nodes. The next logical direction for VPN expansion is to move the functionality onto the CE nodes. By doing so, we can remove the need for per-customer routetables inside any provider node. The provider network need only implement a means to control traffic distribution to only those CE nodes participating in a particular VPN instance.

This document describes a means by which an IPv6 provider network can use multicast to control traffic distribution between participating VPN CE nodes and how those CE nodes can auto discover all other VPN participating CE nodes without additional protocols nor overloaded extensions to existing protocols.

2. Requirements

- o This is a CE-managed service. That is the service provider PE and P routers only move native IPv6 packets and do not otherwise participate in the customer routing protocols.
- o The service provider infrastructure runs native multicast services as defined in [1] [[RFC2362](#)] so precise multicast replication can be performed among the VPN sites.
- o A unique IPv6 scoped multicast address is assigned to each VPN customer as defined in [2] [[RFC4291](#)]. The multicast group prefix of the VPN could be one of several possibilities: ff05, ff08, or could possibly have a new scope ID assignment. The T flag may also be 1.
- o Each participating CE of a VPN joins the VPN assigned group creating a multipoint tunnel between the VPN sites so dynamic discovery of the CE devices can occur. Broadcasting over the tunnel is realized by using the IPv6 multicast in the underlying provider network.
- o ARP protocol as in [3] [[RFC0826](#)] is used to discover the underlying tunnel endpoints. The CE nodes ARP over the tunnel for a

VPN-based next-hop - on the tunnel's subnet - and the hardware address returned is an IPv6 address internal to the provider network.

- o Participating CEs within a VPN share a common routing protocol and neighbor adjacencies through the multipoint tunnel.

3. Multicast VPNs

- o PIM runs with the Intergateway Protocol (IGP) at each customer site as well as over the multipoint tunnel through the provider network.

- o Sending PIM Hello messages are "broadcasted over the multipoint tunnel which ensures only the VPN member CE routers will get the packets.

4. Unicast VPNs

Each VPN CE member router is configured with the core IPv6 VPN multicast group address, which is effectively a VPN ID. Each CE member router joins this core IPv6 multicast group, creating a multipoint tunnel between each of the CE member routers. The VPN customer IGP runs across this multipoint tunnel, establishing neighbor adjacencies and building a complete customer routing table.

By using ARP across the multipoint tunnel to discover the next-hop of each of the CE member neighbors, the learned hardware address returned will be the core-facing IPv6 interface address of the multipoint neighbor. Unicast packets coming from one CE destined to a remote CE VPN neighbor will be unicast encapsulated with the ARP-learned IPv6 next hop of the CE VPN neighbor.

5. Packet Forwarding

5.1. Unicast

Unicast packets are forwarded at the customer site as IPv4 packets to the edge of the network following the IPv4 routed topology. The CE router will encapsulate the IPv4 packets in IPv6 and send to the hardware address learned through the multipoint tunnel across the provider network. The destination CE router will decapsulate and forward the internal IPv4 packet to the unicast destination.

5.2. Multicast

Multicast can run in any of Any Source Multicast (ASM), Source

Specific Multicast (SSM) or BiDirectional (BiDir) within each VPN. For ASM and Bidir the Rendezvous Point (RP) can be located at any of the VPN sites. For joining SSM channels, the member in the receiver site will join a (S,G) which are IPv4 addresses. The IGP routing within the VPN allows the PIM join to travel to the edge and over the multipoint tunnel. The VPN internal multicast state is setup via IPv4 PIM.

Multicast forwarding to receivers sites may be a subset of all participating VPN sites and precise replication/forwarding without unwanted traffic to non-receiver CEs may be desired. To facilitate this, the CE router(s) in the receiver sites will take the IPv4 PIM (S,G) join, after sending it over the multipoint tunnel, and the IPv6 VPN group address to build an IPv6 PIM (S,G) join where:

S is the underlying IPv6 address of the CE router at the source site.

G is a group address derived from the VPN IPv6 group address and the IPv4 (S,G) address.

The complete group address G will be:

ff18:vvvv:ssss:ssss:gggg:gggg::x where s and g are the nibbles of the IPv4 (S,G) address and vvvv is the unique 16-bit VPN ID value. The IPv6 unique VPN multicast address SHOULD comprise only the higher order bits with trailing zeros to allow for at least 64 lower bits to be used for encoding the IPv4 (S,G) address.

6. IANA Considerations

A new ARP hardware type should be specified to identify the IP address of the interface joined to the multipoint tunnel.

7. Security

The VPN member CE routers could maintain secure communications through the use of Security Architecture for the Internet Protocol as described in [4] [[RFC4301](#)].

8. Normative References

[RFC0826] Plummer, D., "Ethernet Address Resolution Protocol: Or converting network protocol addresses to 48.bit Ethernet address for transmission on Ethernet hardware", STD 37, [RFC 826](#), November 1982.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate

Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.

- [RFC2362] Estrin, D., Farinacci, D., Helmy, A., Thaler, D., Deering, S., Handley, M., and V. Jacobson, "Protocol Independent Multicast-Sparse Mode (PIM-SM): Protocol Specification", [RFC 2362](#), June 1998.
- [RFC4291] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", [RFC 4291](#), February 2006.
- [RFC4301] Kent, S. and K. Seo, "Security Architecture for the Internet Protocol", [RFC 4301](#), December 2005.

Authors' Addresses

Greg Shepherd
Cisco Systems

Email: shep@cisco.com

Dino Farinacci
Cisco Systems

Email: dino@cisco.com

Jianping Wu
Cernet

Email: jianping@cernet.edu.cn

Xing Li
Cernet

Email: xing@cernet.edu.cn

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.

