

Network Working Group

Internet Draft

Document: [draft-luciani-ppvnp-vpn-discovery-00.txt](#)

James Luciani
Crescent Networks

Matt Squire
Hatteras Networks

Marty Borden
Atrica

Cedell Alexander
Olen Stokes
Extreme Networks

Pierre Lin
Yipes

Juha Heinanen
Telia

Loa Andersson
Utfors

Ryan Brooks
Time Warner Telecom

September 2001

Using DNS for VPN Discovery

Status of this Memo

This document is an Internet-Draft and is in full conformance with all provisions of [Section 10 of RFC2026](#) [1].

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/lid-abstracts.txt>

The list of Internet-Draft Shadow Directories can be accessed at

<http://www.ietf.org/shadow.html>.

Abstract

Virtual private networks are becoming a common service offered by service providers. There are many technologies over which to implement a VPN service from IPsec to GRE to MPLS. One common requirement of VPN methodologies is the need to discover all of the sites, or at least all the provider equipment associated with the sites, that are in the VPN. DNS provides a simple and commonly available means for site discovery that is independent

of any signaling protocol.

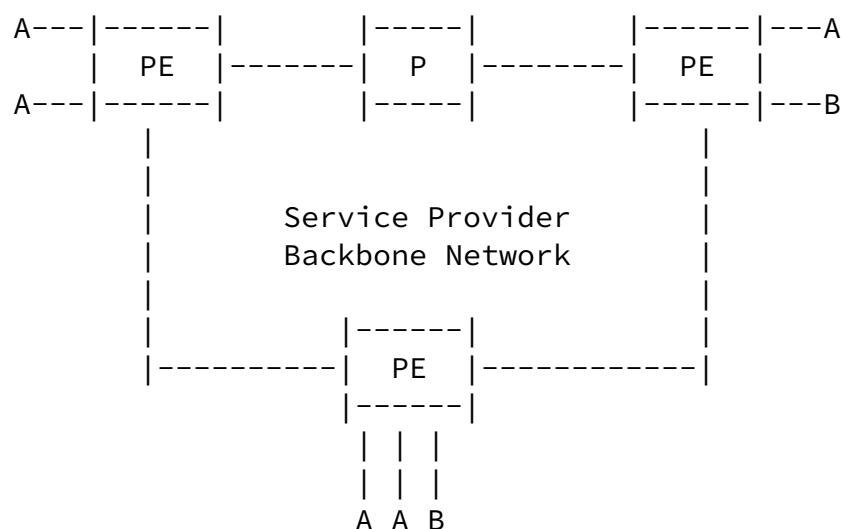
1 Introduction

[Page 1]

[draft-luciani-ppvpn-VPN-discovery-00.txt](#) September 2001

Virtual networking services are being offered by more and more Service providers. There are many flavors of VPNs available in the market today, depending on the customer requirements and provider abilities. There is a variety of data encapsulation protocols used to transport data between customer sites. VPNs may be offered as Layer 2 or Layer 3 services. VPNs may be based on an overlay model or a virtual router model. Other variations are possible.

A VPN consists of a set of customer sites interconnected by one or more provider networks, providing the semblance of private connectivity between the sites over either a private or public backbone network. The following terminology will be used throughout. Customer edge equipment (CE) is located at each customer site and potentially operated independently from the service provider equipment, and may be operated by the customer. The CE equipment is connected to provider edge equipment (PE) that sits at the boundary of the provider network. The PE equipment surrounds a core provider equipment (P). This is depicted in Figure 1.



A: Company A Virtual Network

B: Company B Virtual Network

Figure 1. Virtual Network Model

Each PE supporting a particular VPN must be in a multi-access network with all other PEs supporting that same VPN.

There is a concept called Pseudo Wires (PW), where L2 information is carried in a point to point fashion transparently across a network [pwe3]. Further the concept of a Transparent LAN Service [TLS] also exists wherein an Ethernet virtual 802.1d bridge is simulated for a given set of users. It delivers a layer 2 broadcast domain that is fully capable of learning and forwarding on Ethernet MAC addresses that is closed to a given set of users.

[Page 2]

[draft-luciani-ppvnp-VPN-discovery-00.txt](#) September 2001

All of these services would benefit from a common set of mechanisms and functions in order to promote interoperability and co-existence. In this document the term VPN is used to cover both L2 and L3 VPNs as well as the PWs and TLS.

Certain base functions are common to many of the technologies used to build VPNs. Two of those functions are Discovery and Signaling:

- * Discovery. An optional but incredibly beneficial function is if PE device involved in a particular VPN can discover the other PE equipment in that VPN.
- * Signaling. In many cases, a signaling protocol is required between PE equipment so that particular data flows can be identified and correlated with the VPN.

These functions can be and are implemented in many and various ways.

To date, proposals for discovery have focused on piggy-backing VPN information on BGP and IGP routing protocols. This method has the unfortunate effect of increasing the size of routing tables within the set of affected provider domains, even for those devices that are not involved in the VPN. This increase in size may be quite significant in some cases.

There are other disadvantages to linking discovery and signaling to each other, and to an existing routing protocol. Routing changes or recalculations could interfere with the discovery and

signaling functions of VPNs. When PE equipment is connected with explicitly routed LSPs, for example, such interference is completely unwarranted. Likewise, VPN changes (adding or deleting VPN support) have an impact on routing as this information must be propagated across the network via the routing protocol.

Signaling between PE equipment is required to identify which tunnels are used for which VPNs, to correlate two unidirectional tunnels together to form a bi-directional virtual link, and to give some indication on how to mux/demux traffic from multiple VPNs onto a single tunnel. VPN signaling enhancements have been proposed for LDP, RSVP-TE, CR-LDP, BGP, and OSPF.

[1.1](#) 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](#).

[Page 3]

[draft-luciani-ppvnp-VPN-discovery-00.txt](#) September 2001

[2](#) Using DNS for Discovery

It is highly desirable to use hierarchical identifiers to identify VPN specific information. Hierarchical identifiers permit each organization to guide the use of its own identifier space.

URLs satisfy a hierarchical naming scheme, and they have an advantage over numeric schemes in that the identifier can often be interpreted by the user (e.g., bobsVpn.serviceProvider.net versus <AS 10, ID 64>).

DNS provides mechanisms to resolve a DNS name into a set of IP addresses. Normally, these addresses are interpreted as an 'anycast' identifier; i.e., any of these services can be used to provide connectivity to the named service. When using DNS for VPN resolution, **all** of the addresses are used and are taken to identify the set of PE equipment that supports the named VPN.

Thus, when a PE 10.1.1.1 resolves bobsVpn.serviceProvider.net into {10.1.1.1, 10.2.1.1, 10.5.1.1}, that PE has the IP addresses of the other PEs serving customer sites in bobsVpn.serviceProvider.net. It can then initiate signaling to these other addresses in order to establish the bi-directional tunnel for data transfers.

3 Interactions with Signaling

Current signaling proposals use some variation of a VPN identifier to indicate the VPN that will be used on that specific data channel. These VPN identifiers are of fixed length and potentially with some semantic interpretation.

Although DNS discovery can be used without modifications to signaling, configuration is reduced if the identifiers used in signaling matches the identifier used in discovery. Without signaling enhancements, the VPN DNS name must be mapped to the VPN identifier via manual configuration. Note that this is still preferable to no discovery at all as using DNS names still provides a mechanism to add and delete customer sites to particular VPNs. The <vpn name, vpnid> mapping issue might also be resolved by including the VPNID (route distinguisher) in the name, e.g., 64.10.vpn.isp.fi.

Some guidelines when using DNS with an explicit signaling protocol are:

1. Resolvers SHOULD refresh VPN DNS names resolved for VPN purposes before their TTL expires.
2. A PE MUST use the its address as identified in the A record of the DNS entry as its source address when signaling other PE equipment in the VPN.

[Page 4]

3. If a PE receives a signaled request from a PE not currently in the set of PE addresses associated with a VPN, the PE SHOULD re-request the DNS information for the VPN DNS name. If the requestor source IP address is still not in the list of A records, the request SHOULD be rejected. If the requestor source IP address is in the list of A records, the request SHOULD be accepted.
4. When a refresh or new query results in any A records to which the local PE is not currently connected to for this VPN, and which is not one of its own IP addresses, the local

- PE equipment SHOULD initiate signaling to those newly discovered addresses.
5. When a signaled request to a PE device that was listed in the A records for a VPN DNS name is rejected by the destination, the request should be retried using exponential backoff.

As a potential modification to the above approach, it might be preferable to design a new resource record type which is explicitly designed to apply to the semantics of DNS based VPN discovery. This is for further study.

As a general note on doing name resolution across providers, concerns about ownership of the namespace should be somewhat allayed by the fact that even if a VPN spans multiple providers and ASes, it tends to be "owned" by one of them.

4 Examples

[MARTINI] provide mechanisms for forming a point-to-point L2 VPN between two sites. In the proposal, each side must be configured with the address of the other endpoint of the tunnel, a VC ID, and a group ID. The VC ID and group ID have no semantics, they are used simply to identify the two unidirectional components of a logical bi-directional link. The group ID has additional function in the wildcard removals of associations, but that function is not applicable to this discussion. At the PE equipment, a particular VPN (VLAN, DLCI, etc.) is associated with the tunnel definition (endpoint, VC ID, group ID) via configuration.

Unfortunately, the VC ID and group ID are not hierarchical, and thus if when crossing administrative boundaries its conceivable that matching numbers are not available in all domains. Thus the short flat VC ID space is very limited. Coordination among the domains managing the edge devices is required.

When generalizing [[MARTINI](#)] to a full mesh topology, the problem of configuring the peers becomes more problematic as each peer must be configured with the address of every other. Additionally, the configuration of more PEs must be correlated in the group and VC ID.

It would be simpler if the PEs could simply be configured with the VPN DNS name, the associated VPN (VLAN, DLCI, etc.), and the group ID. The peers could then be discovered via DNS resolution, and the VPN DNS name could be used in signaling (instead of the VC ID) to determine the data channels for this VPN.

Although this section discussed DNS based discovery based on the [\[MARTINI\]](#) techniques, the discovery mechanism is generally applicable to any environment where LDP, CR-LDP, or RSVP-TE is used for signaling (rather than routing protocols as discussed in earlier sections).

[5](#) Security Considerations

A Virtual private network, by its very nature implements a policy, as agreed upon between the customer and provider, that hides information about the customer's VPN from others. It is reasonable to assume that this policy would require restricting anyone other than the customer from deriving that customer's sites using mechanisms of the provider. An even more important security requirement is that a customer not be able to manipulate the site information about another customer. (It is possible that the provider may wish to allow customers to manipulate their own site information, although this would likely be done through an indirect method.)

DNS itself provides many security extensions that could be used to protect the identity of the PE equipment in a particular VPN. DNS also provides a dynamic update ability that could conceivably be used to provide PE equipment with the ability to register itself with the DNS server upon configuration into a particular VPN. These possibilities are recognized but not investigated within this draft.

[6](#) Issues

We list some issues that are not resolved or discussed in this draft that will be considered in future revisions of the draft.

- o Solutions for security requirements must be given.
- o Since DNS provides required information only upon demand, a mechanism must be developed so that PEs will update their endpoint information when necessary. This might be as simple as a timeout and query again type of mechanism.
- o Specific recommendations for TLV formats for LDP, RSVP-TE, and CR-LDP and inter-working with current Martini proposal.

[7](#) References

[draft-luciani-ppvnp-VPN-discovery-00.txt](#) September 2001

[RFC2026] Bradner, S., "The Internet Standards Process -- Revision 3", [BCP 9](#), [RFC 2026](#), October 1996.

[MARTINI] Martini, Luca, et al., "Transport of Layer 2 Frames Over MPLS", [draft-martini-l2circuit-trans-mpls-05.txt](#), Work in Progress.

[TLS] Lasserre, Marc, et al., "Transparent VLAN Services over MPLS", [draft-lasserre-tls-mpls-00.txt](#), Work in Progress.

[PWE3] Xiao, XiPeng, et al., "Requirements for Pseudo-Wre Emulation Edge-to-Edge (PWE3)", [draft-ietf-pwe3-requirements-01.txt](#), Work in Progress.

[8](#) Acknowledgments

[9](#) Author's Addresses

Matt Squire
Hatteras Networks
639 Davis Drive
Research Triangle Park, NC 27709
Email: msquire@hatterasnetworks.com

James V. Luciani
Crescent Networks
900 Chelmsford
Lowell, MA 01851
Email: jluciani@crescentnetworks.com

Marty Borden
Atrica, Inc.
Email: mborden@acm.org

Cedell Alexander
Olen Stokes
Extreme Networks
Research Triangle Park, NC 27709
Email: calexander@extremenetworks.com, ostokes@extremenetworks.com

Pierre Lin
Yipes Communications, Inc.

Juha Heinanen
Telia Finland
Email: jh@telia.fi

Loa Andersson
Utfors Research, Architecture and Future Lab (URAX)

[Page 7]

[draft-luciani-ppvnpn-VPN-discovery-00.txt](#)

September 2001

Utfors AB
Rörsundavägen 12
Box 525, 169 29 Solna, Sweden
Office: +46 8 5270 2000
Email: loa.andersson@utfors.se

Ryan K. Brooks
Time Warner Telecom, Inc.
3235 Intertech Drive
Brookfield, WI 53045
Email: ryan@twtelecom.net

[Page 8]