Network Working Group
Internet-Draft
Expires: August 30, 2005

P. Bose, Ed.
Lockheed Martin
R. Bonica
Juniper Networks
R. White
Cisco Systems

February 26, 2005

Secure Layer 3 Virtual Private Networks
draft-bose-secure-l3vpn-00

Status of this Memo

Copyright Notice

Abstract

This document presents an framework for secure layer 3 Virtual Private Networks (VPNs) for customer networks that exchange encrypted information through a service provider network. It also describes the necessary interactions between the various functions (e.g. routing and encryption) at the VPN boundary to construct the secure

VPN.

Table of Contents

1. Introduction

VPNs are typically constructed in one of two ways. One is that a service provider offers the VPN, and provides an underlying circuit network, often MPLS, that connects the underlying endpoints as defined in a contract. These are referred to as "Provider-Provisioned VPNs" [RFC3809]. The other, generally referred to as "customer-provisioned", is that the edge routers themselves provide tunnels over an underlying network using one of a variety of types of IP tunnel technologies such as loose source routes as specified in DVMRP [RFC1075], IP/IP [RFC2003], IPsec/ESP [RFC2401][RFC2406], L2TP [RFC2661], GRE [RFC2784] and others.

In this context, a "Secure VPN" is an example of an IPsec or IPsec-like VPN, and is therefore "customer-provisioned". Such networks have in the past been a complex collection of tunnels, a star or a multi-star network in which a large set of client sites maintain static or semi-static tunnels with a much smaller set of service sites. This document presents a Layer 3 VPN architecture for such networks by taking into account the underlying network architecture and requirements of a secure VPN where plaintext (red) enclaves connect to remote plaintext enclaves through a ciphertext (black) network.

1.1. Key Requirements for Secure VPNs in the GIG

Secure VPNs presents a unique set of requirements which include:

- Scalability of VPN sites up to a large number of routes and prefixes: Each secure enclave may have hundreds or thousands of reachable destinations, and there may be thousands of secure enclaves. A solution proven to be able to map large number of destinations to some indicator of the enclave those destinations belong to, ideally a ciphertext (black) next hop address representing the secure enclave. Two types of systems have these scaling characteristics in today's global Internet, name resolution systems, and interdomain routing systems.

- Mobility of customer and service provider networks: Secure enclaves will tend to be mobile in most environments, and reachable destinations will tend to be mobile between secure enclaves themselves. This implies any solution must be able to adjust to changes in enclave or reachable destinations within an enclave very quickly, within the application requirements for the network as a whole. In most large scale internetworks, the times acceptable for reacting to changes in reachability or connectivity is on the order of seconds to tens of seconds, rather than minutes to hours. Interdomain routing systems deployed in the global

Internet can meet these convergence time criteria.

- Secure routing and discovery of hosts, routers and gateways: Destination reachability, in some environments, cannot become known within the general routing table of all devices. This implies not only authentication and authorization of reachability information must be provided, but also that confidientality of reachability information. Current interdomain routing protocols can be used over secure point-to-point links, providing this type of security. Further, authorization to advertise a specific set of destinations must also be authorizable, in some fashion. Interdomain routing protocols allow the validation of authorization to advertise reachability, and this is an area of current research and work for interdomain routing protocols.

- Minimal information transfer at the VPN boundary: Minimal amounts of information may be passed between secure enclaves and the internetwork through which these secure enclaves communicate. Primarily, for this work, this includes reachability information, as described above; reachability information cannot be passed from the secure enclave into the insecure internetwork. Thus, there must be some way to associate a public address with a set of destinations within a secure enclave, without adding any information to the insecure internetwork. Current interdomain routing allows this attachment through the use of indirection; the next hop towards a particular destination does not need to be the transmitter of the routing information itself.

- Reachability policies are required: Because of the nature of the routing information, the system used to advertise and discover reachability information within secure enclaves must be able to attach policy about where and how that routing information may be shared to the routing information itself. Such policies are also required to provide inter-enclave routing with information about the best possible entry point into the secure enclave to reach any specific destination within the enclave, as well as other preferences. Current interdomain routing within the global Internet has similar policy requirements, and thus the current interdomain routing protocol supports such policy mechanisms.

- Dynamic on-demand discovery of reachability In some environments, it may be more feasible to request routes on demand, rather than learning all possible routes through a connection to a server or other device. Current interdomain routing does not support this mode of operation, but could be easily modified to provide this type of service.

2. Functional Architecture of a Secure L3 VPN

A simple representation of a secure VPN is shown in Figure 1.

```
       Plaintext(PT)|         Ciphertext(CT)          | Plaintext(PT)
          (Red)      |            (Black)              |    (Red)
                     |                                 |
             +-------+-------+                 +-------+-------+
             |       |       |                 |       |       |
             |  +----+---------------------------------+----+  |
             |  +----+---------------------------------+----+  |
             |    VPN|Router |    Encrypted    |    VPN|Router |
             +-------+-------+     Tunnel      +-------+-------+
                     |                                 |
                     |                                 |
                     |                                 |
                     |                                 |
                     |                                 |
```

Figure 1: Simple VPN

Each customer network has a plaintext (red) enclave which participates in a virtual private network with other red enclaves. This is accomplished via an encrypted tunnel which originates and terminates in the red enclave and traverses through the intermediate ciphertext (black) service provider network. The VPN Router (depicted here as a single logical node) performs the following functions:

o Intra-Domain Routing

o Encryption/Decryption

o Transfer of necessary control plane information between plaintext and ciphertext domains (e.g. IP addresses)

o Inter-Domain Routing

o Data Forwarding

Each of these functions may be performed by the same or different physical entity in the network. A description of the various physical implementation alternatives of these functions is beyond the scope of this document.

The simple view presented in Figure 1 can be expanded in terms of a typical layer 3 framework as described in [RFC4110] and depicted in Figure 2 below.

```
       Plaintext(PT) |         Ciphertext(CT)          | Plaintext(PT)
          (Red)      |            (Black)              |    (Red)
                     |                                 |
           +-----+   +   +-----+   +-----+   +-----+   +  +-----+
      IGP  | CE  |   |   | CE  |   | PE  |   | CE  |   |  | CE  |IGP
     <==== | ================================================   |===>
    Static |  +------+---------------------------------+------+ |Static
           |     |   |   |     |   |     |   |     |   |  |     |
           +-----+   +   +-----+   +-----+   +-----+   +  +-----+
                     |IGP                           IGP|
                     |<==>     <==>  BGP  <==>     <==>|
                     |Static                     Static|
                     |                                 |
                     |              IPsec              |
                  <=========================================>
                     |                                 |
                     |     BGP(For PT to CT Nexthop)   |
              <=================================================>
```

Figure 2: Secure Layer 3 VPN Architecture

2.1. Proposed L3 VPN Architecture

The proposed Layer 3 VPN architecture as described in Figure 2 overlays a BGP VPN over IPsec tunnels. It is assumed for this description that the two first hop routers at the VPN boundary form the customer edge router which may or may not be the same device. The red CE router learns and summarizes red prefixes using an IGP in the red enclave such as OSPF. These prefixes are imported into the inter-domain routing function at the red CE router and advertised to a remote red enclave via BGP. BGP information exchange between remote enclaves flows along with other data in the security associations (e.g. IPsec). Any given red CE router learns the black address of the remote enclave initially via configuration, a routing service such as a BGP Route Reflector or a red-to-black address translation scheme.

Data is forwarded into the appropriate security association as per the virtual forwarding table constructed at the VPN boundary by BGP. The black domain may use static, IGP or BGP routing as appropriate. The figure above presents one possible routing model for the black domain. The information exchange between the red and black domains in this model is restricted to the knowledge in the red domain of the black address to red address mapping of the remote enclave CE router.

The proposal is further explained by describing the different functional components at the VPN boundary and the roles played by them in the VPN as shown in Figure 3

```
                                PLAINTEXT               | CIPHERTEXT
                                  (Red)                 |  (Black)
                                                        |
                            +---------------+           |
                PT Prefixes | Intra-Domain  |           |
                 R1,R2,R3   |   Routing     |           |
                            +---------------|           |
                                    |                   |
                                    |                   |
                               Route|Redistribute       |
                                    |                   |
                                    |                   |
   +-----------+   CT-PT    +---------------+  CT IP    |
   | BGP Route |<---------->| Inter-Domain  |<--------+ |
   | Reflector |  Routes    |   Routing     |  Addr   | |
   +-----------+            +---------------+  =      | |
      R1->B1.H1                     |          B1.H1  | |
      R2->B1.H1                 FIB | Table           | |
      R3->B1.H1             +---------------+     +-----+------+
                            |  Forwarding   |-----| Encryption |
                            +---------------+     +-----+------+
                                                        |CT IP Addr
                                                        |
```

Figure 3: Functional Components at Secure VPN Boundary

2.2. Functional components at the VPN boundary

The role of the various functions at the VPN boundary as depicted in Figure 3 are:

- Intra-Domain Routing: This function advertises, withdraws, and summarizes the routes in the red enclave. It maintains an internal routing table for the red domain. This routing table is provided to the Inter-Domain Routing function for advertisement to remote enclaves.

- Encryption: This function encrypts and decrypts all data traffic into IPsec security associations with peer enclaves. It establishes these security associations with a peer encryptor using the IP address provided by the Forwarding function. The encryption function has an IP address on the black side (CT IP Addr) and an IP address on the red side (PT IP Addr), which is used for forwarding data in the respective domains.

- Inter-Domain Routing: This function imports the internal routing table from the Intra-Domain Routing function. It also advertises the CT IP Addr of the encryption as the next-hop for the internal

red routes of the enclave. It provides this information to its peer functions in the remote enclaves through a protocol such as BGP. It also provides this information to a Route Reflector which can be queried by remote enclaves to learn the CT IP Addr of the encryption function for an internal red route. This function also manages the information transfer of control plane information between the red and black domains. In secure VPNs, the CT IP Addr of the encryption function is accessed by the Inter-Domain Routing function via manual configuration, SNMP, a simple routing protocol such as RIP etc.

- Route Reflector: The route reflector function maintains the inter-domain routing tables as provided the Inter-Domain Routing function. It converses with other route reflectors to exchange this routing information. A red enclave Inter-Domain Routing function can query the route reflector to learn the next-hop information for remote enclaves. Routes learnt from the Route Reflector by the Inter-Domain routing function can be used to populate the Security Policy Database of the encryption function via manual configuration, SNMP etc.

- Forwarding: This function maintains the forwarding table which is used to forward data to remote enclaves. It receives this forwarding table from the Inter-Domain Routing function. The IP addresses maintained in this forwarding table are used by the encryption function to set up security associations to remote enclaves.

These functions and the described interactions construct the proposed virtual private network. Unicast routing between peer enclaves is conducted via routing exchanges inside the security associations by the peer Inter-Domain routing functions or by querying the neighborhood route reflector.

2.3. Hierarchical Distribution of Reachability Information

In many cases, secure enclaves will not have the correct information to directly connect and build BGP peering sessions; the correct set of public internetwork addresses to connect to will be difficult to discover on a large scale. To resolve this problem, a set of route servers, or route reflectors [RFC2796], may be maintained for each secure enclave attached to the insecure internetwork. The BGP speaker contained within the secure enclave would have a manually configured list of these route reflectors. Some form of anycast addressing may be applicable to this situation, though this has not been thoroughly investigated.

In some situations, it may be difficult to scale a single group of

route reflectors or route servers large enough to support the number of sites within an interconnected secure enclave. Hierarchical route reflectors, or hierarchical eBGP route servers, may be used to scale in these situations. This is common practice in the public Internet.

3. IANA Considerations

This document makes no request of the IANA.

Note to RFC Editor: in the process assigning numbers and building IANA registries prior to publication, this section will have served its purpose. It may therefore be removed upon publication as an RFC.

4. Security Considerations

One security concern addressed in this proposal is the transfer of CT IP address information to the plaintext side. In the current proposal this is achieved via an authorized manual/automated network management function on the plaintext side which queries the encryption function. Alternatively a simple routing protocol like RIP is used on the plaintext enclave only to provide this information.

The security policy database for the encryption function, denoted in the proposal as the forwarding table is also populated with CT IP address nexthop information for remote enclaves. This information is again configured via an an authorized manual/automated network management function on the plaintext side. Only ciphertext IP addresses fronting remote VPN enclaves are therefore used within the plaintext enclave for the purposes of routing. Since this information is only available on the plaintext enclave, which due to the VPN characteristics of these networks a more secure environment, misuse of this information within the plaintext enclave is expected to be minimal.

5. Contributors

- Fred Baker (fred@cisco.com)
- Eric Fleischman (eric.fleischman@boeing.com)
- Julie Tarr (tarr@itd.nrl.navy.mil)
- Tony De Simone (antonio.desimone@jhuapl.edu)

6. Acknowledgements

Initial introductory text is borrowed from [I-D.baker-nested-vpn-routing] .

7. References

7.1. Normative References

[I-D.baker-nested-vpn-routing]
Baker, F., "Routing across Nested VPNs", draft-baker-nested-vpn-routing-01 (work in progress), July 2005.

[RFC3809] Nagarajan, A., "Generic Requirements for Provider Provisioned Virtual Private Networks (PPVPN)", RFC 3809, June 2004.

7.2. Informative References

[RFC1075] Waitzman, D., Partridge, C., and S. Deering, "Distance Vector Multicast Routing Protocol", RFC 1075, November 1988.

[RFC2003] Perkins, C., "IP Encapsulation within IP", RFC 2003, October 1996.

[RFC2401] Kent, S. and R. Atkinson, "Security Architecture for the Internet Protocol", RFC 2401, November 1998.

[RFC2406] Kent, S. and R. Atkinson, "IP Encapsulating Security Payload (ESP)", RFC 2406, November 1998.

[RFC2661] Townsley, W., Valencia, A., Rubens, A., Pall, G., Zorn, G., and B. Palter, "Layer Two Tunneling Protocol "L2TP"", RFC 2661, August 1999.

[RFC2784] Farinacci, D., Li, T., Hanks, S., Meyer, D., and P. Traina, "Generic Routing Encapsulation (GRE)", RFC 2784, March 2000.

[RFC4110] Callon, R. and M. Suzuki, "A Framework for Layer 3 Provider-Provisioned Virtual Private Networks (PPVPNs)", RFC 4110, July 2005.

Authors' Addresses

Pratik Bose (editor)
Lockheed Martin
700 North Frederick Ave
Gaithersburg, Maryland 20876
USA

Phone: +1-240-462-9083
Fax: +1-301-428-5415
Email: pratik.bose@lmco.com

Ron Bonica
Juniper Networks
Virginia
USA

Phone:
Fax:
Email: rbonica@juniper.net

Russ White
Cisco Systems
North Carolina
USA

Phone:
Fax:
Email: riw@cisco.com

Full Copyright Statement

Intellectual Property

Acknowledgment

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