

Jim Guichard
Robert Hanzl
Dan Tappan
Scott Wainner
Cisco Systems, Inc

Vic Locicero
INFONET Services Corporation

IETF Internet Draft

Expires: November, 2003

Document: [draft-guichard-ce-ce-ipsec-00.txt](#)

May, 2003

CE-CE IPSec within an [RFC-2547](#) Network

Status of this Memo

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

This document describes a reference architecture that may be used to tightly integrate CE-CE [IPSec] encryption with the any-to-any connectivity model of [\[RFC2547\]](#). Using this mechanism, a Service Provider is able to provide an IP VPN service with data encryption between customer edge routers, but without the need of direct routing protocol exchange, or IP-based tunnels such as provided by [GRE].

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 [\[RFC-2119\]](#).

[1.1](#) Terminology

Several terms used within this document are defined as follows:

"Security Gateway" - A router that is a member of a [\[RFC2547\]](#) VPN and serves as a termination point for [IKE] and [IPSec] Security Associations

"Security Gateway Identity" - An IPv4 address representing the identity of a router serving as an security gateway for the establishment of [IKE] and [IPSec] Security Associations

"Trusted Subnet" - A range of IP addresses represented as a network and mask that a security gateway protects

"Security Policy" - A set of policies and attributes used to protect information as described in [IKE]

[2](#) Introduction

[RFC2547] provides an attractive service architecture that is able to build an any-to-any data path between VPN sites. However, this model does not provide any inherent data encryption services; therefore, customers that wish to encrypt their traffic must do so before it enters the [\[RFC2547\]](#) network. This is typically achieved by enabling [IPSec] encryption and running [IPSec] tunnels between CE routers that belong to the "encrypted" VPN.

The deployment of [IPSec] tunnel meshes is analogous to the "overlay" model used in Frame-relay or ATM networks. One of the key success drivers for [\[RFC2547\]](#) is its ability to avoid such topologies, and provide any-to-any connectivity while eliminating pre-configuration of a mesh of CE-to-CE circuits. It is desirable to align the CE-to-CE protection methodologies with the any-to-any connection model provided by [\[RFC2547\]](#) so that the customer experience is seen as the

"best of both worlds". This enhanced model supports CE-to-CE data protection while eliminating the requirement for pre-established IP-based tunnels or routing adjacencies between [IPSec] security gateways.

Some Service Providers have provided some integration of [[RFC2547](#)] and [IPSec] by terminating [IPSec] tunnels into a Virtual Routing & Forwarding Instance (VRF). This solution, although network-based, does not extend [IPSec] between security gateways in different

Guichard et. al

2

customer sites, and is used more for extending the reach of the Service Provider so that remote customer locations are able to access the VPN.

This document provides a mechanism that is able to maintain the any-to-any connectivity nature of [[RFC2547](#)], but also enables the dynamic establishment of the CE-CE [IPSec] topology. The [IPSec] security associations established can be thought of as "Security & Forwarding Associations" in the sense that they are used to exchange encrypted data packets between the CE routers; however there is no requirement that they be used to exchange routing information. Thus, the routing scalability property of [[RFC2547](#)] is preserved.

[3](#) Service Provider Infrastructure Reference Model

A Service Provider may deploy an [[RFC2547](#)] service using a number of backbone tunneling techniques such as those described in [[RFC2547](#)], [MPLS-in-IP], or [PE-PE-IPsec]. [[RFC2547](#)] uses a hierarchical routing model that provides scalable distribution of route forwarding attributes. CE-CE [IPsec] encryption, as described within this document, relies upon the IP address partitioning and route forwarding state created by the [[RFC2547](#)] infrastructure and it can be deployed independently of the backbone tunneling technique chosen. The CE-CE [IPSec] topology requires a point-to-point relationship between CE's for data protection; however, the routing plane associated with the CE-CE topology leverages the [[RFC2547](#)] hierarchical routing model.

[4](#) Coupling of CE Security Policy and PE Routing Planes

Each CE router ascertains through configuration, or other means outside the scope of this document, whether it is used as a [IPsec]

security gateway. Once this information is discovered, the CE router MUST advertise it's "Security Gateway Identity" used for [IKE] and [IPSec] peer end-point termination to the PE router using [BGP-4]. The identity must be associated with each 'trusted subnet' represented as a prefix that the CE router protects. The identity will typically be an IPv4 address where the [IKE] and [IPSec] authentication and encryption services will be established. The PE router MUST then use [MP-BGP] to advertise the trusted subnet prefixes and the associated identity information to other PE routers.

A PE router that receives this information via [MP-BGP] MUST be able to (a) identify which VPN the prefix and security gateway identity end-point is associated with, and (b) advertise that information to any security gateway CE routers that belong to the VPN. Identification of which VPN the update belongs to is determined by

the "route-target" extended-community attribute as described in [\[RFC2547\]](#).

5 Distribution of [IPSec] Security Gateway End-points

A CE router MAY send encrypted and non-encrypted traffic toward the PE router for delivery to other members of its VPN. A CE router that belongs to an "encrypted" VPN needs to be able to build a Security Association (SA) with any remote CE router that also belongs to the same VPN, and is a member of the encryption service.

When traffic that needs to be encrypted is sent from a CE router that belongs to an "encrypted" VPN, the CE router MUST be able to establish a Security Association (SA) with the remote CE router through which the destination of the incoming packet is reachable. To achieve this aim, the CE router needs to discover the remote peer's trusted subnet prefix and the associated security gateway identity of the peer, and then build the [IKE] and [IPSec] security association.

Discovery of the end-point addresses is achieved through direct [BGP-4] exchange with the PE router. If [BGP-4] is not established with the customer site, then a different discovery protocol is necessary and is outside the scope of this draft.

Many [\[RFC2547\]](#) deployments use [BGP-4] on the PE-CE links. Typically these sessions only carry standard [BGP-4] attributes. In order to

use the mechanisms described within this document, the CE router MUST support the capabilities as specified in [EXTCOM].

When a CE router advertises routes from an "encrypted" VPN into the backbone it MUST attach a new BGP extended-community attribute, hereby referred to as the "Security Gateway Identity" attribute, to all trusted subnet prefixes for which encryption is desired. A PE router that receives such an update MUST export those trusted subnet prefixes along with the "Security Gateway Identity" attribute.

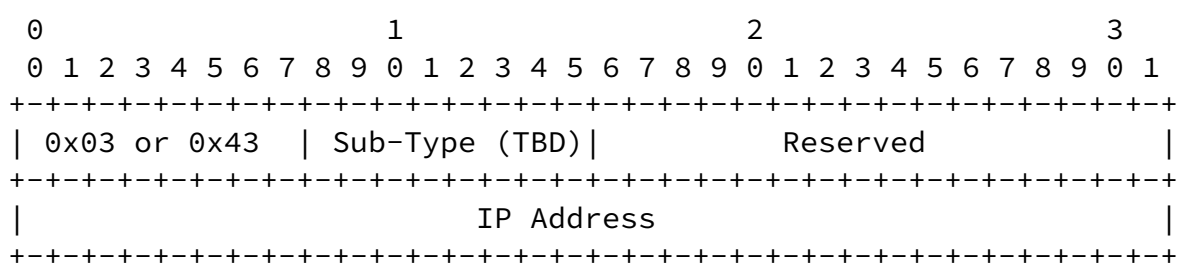
A PE router that receives the update MUST advertise the trusted subnet prefixes and security gateway identities to any relevant CE routers that are (a) members of the "encrypted" VPN, and (b) are running [BGP-4] with the PE router.

6 BGP "Security Gateway Identity" Extended-community Attribute

The Extended Community Attribute is a transitive optional BGP attribute, with type code 16, as specified in [EXTCOM]. The solution described within this document proposes to use the Opaque Extended Community, as specified in [section 6.4](#) of [EXTCOM]. The format of this community is as follows:

Guichard et. al

4



The value of the high-order octet of this extended type is either 0x03 or 0x43. The low-order octet of this extended type carries the sub-type with value (TBD) and indicates that this is a "Security Gateway Identity".

The value field consists of:

Reserved - 2 octets

Reserved for future use and should be set to 0x0000

IP Address - 4 octets

Security Gateway IP address

7 CE-router IPsec requirements

A CE router that wishes to belong to an "encrypted" VPN, and use the mechanisms described within this document, MUST conform to the procedures described in [IPsec]. This means that the CE router MUST provide a Security Policy Database (SPD), as described in [section 4.4.1](#) of [IPsec], which is used to determine the disposition of all IP traffic inbound or outbound from the router. Each entry within the database specifies whether traffic matching the policy should bypass IPsec processing, be discarded, or be subject to IPsec processing.

A CE router MUST provide the ability to specify "Selectors", as described in [section 4.4.2](#) of [IPsec]. These are a set of IP and upper layer protocol field values that are used by the Security Policy Database (SPD) to map traffic to a Security Association (SA).

The Security Policy Database and Selector attributes MUST be populated with the "Security Gateway Identity" and the associated "trusted subnet" prefixes. The population of the SPD from [BGP-4] may be an automated process with the appropriate [BGP-4] controls provided by the CE.

Each CE router MUST enable the creation of security associations in the Security Association Database (SAD), as described in [section 4.4.3](#) of [IPsec], that contains parameters derived by traffic

matching the [BGP-4] injected Selectors in the SPD. This database is used to determine what [IPsec] services are offered to IP packets.

7.1 CE-CE Security Association Setup using IKE

A CE router MUST support an automated Security Association/Key management protocol for the purpose of establishing and maintaining Security Associations between two [IPsec] peer end-points. One example of such a protocol is [IKE].

There are several options available to the CE routers with respect to IPsec tunnel setup and encryption of traffic:

CE-CE authentication and/or encryption of selective packets based on traffic flow initiated establishment of security associations

CE-CE authentication and/or encryption of selective packets based on pre-established [IPSec] security associations

Each of these options is described in the following sections. Regardless of which option is used, on receipt of traffic that is matched to an SPD policy that requires [IPSec] processing, a CE router MUST check whether a Security Association (SA) already exists with the [IPSec] Security Gateway address. If an SA already exists, then the CE router can encrypt the traffic and forward it toward the PE router. If no SA exists, the CE router MUST use [IKE] or similar protocol to establish the SA with the security gateway identity.

[7.1.1](#) CE-CE encryption of selective packets based on traffic flow

CE-CE encryption may be driven by traffic flow and a CE router MAY choose to selectively encrypt packets based on a 'Selector' match. On receipt of a packet that is matched by the CE router's SPD for encryption, the CE router MUST be able to establish an SA with the remote CE router through which the destination is reachable.

As the CE router is running [BGP-4] with the PE router, it can dynamically build the 'Selector' criteria based on receipt of routing updates that carry the "Security Gateway Identity" attribute. Using this information, the CE router is able to identify which routes are associated with a remote site, and also which of these routes need encryption. For the routes that need encryption, the CE is able to determine the "Security Gateway Identity" associated with those routes.

The CE MAY dynamically establish [IPSec] SA's between the CE and PE routers. These [IPSec] tunnels may be used to protect the [BGP-4] exchange of 'Trusted Subnets' and 'Security Gateway Identities'

between the PE and CE. Alternatively, the CE and PE routers MAY use [BGP-MD5] on the [BGP-4] session to authenticate the prefixes and the associated "Security Gateway Identity".

[7.1.2](#) CE-CE Encryption with pre-established IPSec Security

Associations

A CE router MAY choose to pre-establish [IPSec] tunnels between CE routers. [IPSec] SA's may be established automatically upon population of the SPD that occurs upon receipt of a 'Trusted Subnet' prefix with a valid "Security Gateway Identity". The CE router MUST encrypt all traffic destined to a route via the established [IPSec] security association.

The CE MAY have pre-established [IPSec] SA's between the CE and PE routers. These [IPSec] tunnels may be used to protect the [BGP-4] exchange of 'Trusted Subnets' and 'Security Gateway Identities' between the PE and CE. Alternatively, the CE and PE routers MAY use [BGP-MD5] on the [BGP-4] Session to authenticate the prefixes and the associated "Security Gateway Identity".

8 References

[RFC2547], Rosen, E. et al., "BGP/MPLS VPNs", [draft-ietf-ppvpn-rfc2547bis-03](#), October, 2002.

[GRE], Li, T. et al, "Generic Routing Encapsulation (GRE)", [RFC 1701](#), October, 1994.

[IPSec], Kent and Atkinson, "Security Architecture for the Internet Protocol", [RFC 2401](#), November 1998.

[MPLS-in-IP], Rosen E. et al., "Encapsulating MPLS in IP or GRE", [draft-ietf-mpls-in-ip-or-gre-00](#), January, 2003.

[PE-PE-IPsec], Rosen E. et al., "Use of PE-PE IPsec in [RFC2547](#) VPNs", [draft-ietf-ppvpn-ipsec-2547-03](#), February, 2003.

[MP-BGP], Rekhter, Y. et al., "Multiprotocol Extensions for BGP-4", [RFC 2858](#), June, 2000.

[BGP-4], Rekhter, Y. et al., "A Border Gateway Protocol 4 (BGP-4)", [RFC 1771](#), March, 1995.

[EXTCOM], Tappan, D. et al., "BGP Extended Communities Attribute", [draft-ietf-idr-bgp-ext-communities-05](#), May, 2002.

[IKE], Harkins, D. et al., "The Internet Key Exchange (IKE)", [RFC 2409](#), November, 1998.

[BGP-MD5], Heffernan, A., "Protection of BGP Sessions via the TCP MD5 Signature Option", [RFC 2385](#), August 1998.

[9](#) Authors' Address

Jim Guichard
Cisco Systems, Inc.
1414 Massachusetts Avenue
Boxborough, MA, 01719
Email : jguichar@cisco.com

Robert Hanzl
Cisco Systems, Inc.
1414 Massachusetts Avenue
Boxborough, MA, 01719
Email : rhanzl@cisco.com

Dan Tappan
Cisco Systems, Inc.
1414 Massachusetts Avenue
Boxborough, MA, 01719
Email : tappan@cisco.com

Scott Wainner
Cisco Systems, Inc.
13600 Dulles Technology Drive
Herndon
Virginia, 20171
Email : swainner@cisco.com

Vic Locicero
INFONET Services Corporation
2160 E. Grand Ave.
El Segundo, CA 90245
Email : vic_locicero@infonet.com

