IPsecME Working Group
Internet-Draft
Intended status: Informational
Expires: September 6, 2012

S. Hanna
Juniper

March 5, 2012

Point to Point VPNs Problem Statement
draft-ietf-ipsecme-p2p-vpn-problem-00

Abstract

This document describes the problem of enabling a large number of systems to communicate directly using IPsec to protect the traffic between them. Manual configuration of all possible tunnels is too cumbersome in such cases, so an automated method is needed.

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at http://datatracker.ietf.org/drafts/current/.

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."

This Internet-Draft will expire on September 6, 2012.

Table of Contents

1. Introduction

IPsec [RFC4301] is used in several different cases, including tunnel-mode site-to-site VPNs and Remote Access VPNs. Host to host communication employing transport mode also exists, but is far less commonly deployed.

The subject of this document is the problem presented by large scale deployments of IPsec. These may be a large collection of VPN gateways connecting various sites, a large number of remote endpoints connecting to a number of gateways or to each other, or a mix of the two. The gateways and endpoints may belong to a single administrative domain or several domains with a trust relationship.

Section 4.4 of RFC 4301 describes the major IPsec databases needed for IPsec processing. It requires an extensive configuration for each tunnel, so manually configuring a system of many gateways and endpoints becomes infeasible and inflexible.

The difficulty is that all the configuration mentioned in RFC 4301 is not superfluous. IKE implementations need to know the identity and credentials of all possible peer systems, as well as the addresses of hosts and/or networks behind them. A simplified mechanism for dynamically establishing point-to-point tunnels is needed. Section 2 contains several use cases that motivate this effort.

1.1. Terminology

Endpoint - A host that implements IPsec for its own traffic but does not act as a gateway.

Gateway - A network device that implements IPsec to protect traffic flowing through the device.

Point-to-Point - Direct communication between two parties without active participation (e.g. encryption or decryption) by any other

parties.

Security Association (SA) - Defined in [RFC4301].

1.2. 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 [RFC2119].

2. Use Cases

This section presents the key use cases for large-scale point-to-point VPN.

In all of these use cases, the participants (endpoints and gateways) may be from a single organization or from multiple organizations with an established trust relationship. When multiple organizations are involved, products from multiple vendors are employed so open standards are needed to provide interoperability. Establishing communications between participants with no established trust relationship is out of scope for this effort.

2.1. Endpoint-to-Endpoint P2P VPN Use Case

Two endpoints wish to communicate securely via a direct, point-to-point SA.

The need for secure endpoint to endpoint communications is often driven by a need to employ high-bandwidth, low latency local connectivity instead of using slow, expensive links to remote gateways. For example, two users in close proximity may wish to place a direct, secure video or voice call without needing to send the call through remote gateways, which would add latency to the call, consume precious remote bandwidth, and increase overall costs.

2.2. Gateway-to-Gateway P2P VPN Use Case

Two gateways suddenly need to exchange a lot of data.

For example, a mobile worker from one government agency may sit down in a shared remote office and start up his VOIP or video phone software. He should rapidly get an efficient, secure, low latency connection to his voice mail system and to anyone that he might call. This user, his voice mail system, and other people that he calls will probably be operating behind gateways but those gateways may have little advance warning of the need to establish secure connectivity between them.

2.3. Endpoint-to-Gateway P2P VPN Use Case

An endpoint wants to connect directly to the most efficient gateway for accessing a particular service.

For example, a mobile user roaming on the Internet may need to open a remote desktop connection to a virtual machine hosted on a particular server or to a service provided by a variety of servers distributed around the globe. The user should be able to establish a connection

directly to the gateway closest to the service desired. If multiple gateways can suffice, load balancing and failover across gateways may be useful.

3. Inadequacy of Existing Solutions

Several solutions exist for the problems described above. However, none of these solutions is adequate, as described here.

3.1. Exhaustive Configuration

One simple solution is to configure all gateways and endpoints in advance with all the information needed to determine which gateway or endpoint is optimal and to establish an SA with that gateway or endpoint. However, this solution does not scale in a large network with hundreds of thousands of gateways and endpoints, especially when multiple organizations are involved and things are rapidly changing (e.g. mobile endpoints). A more dynamic system for securely and scalably establishing SAs between gateways is needed.

3.2. Star Topology

The most common way to address this problem today is to use what has been termed a "star topology". In this case one or a few gateways are defined as "core gateways", while the rest of the systems (whether endpoints or gateways) are defined as "satellites". The satellites never connect to other satellites. They only open tunnels with the core gateways.

For a large number of gateways in one administrative domain, one gateway may be defined as the core, and the rest of the gateways and remote access clients connect only to that gateway. If the packet destination is behind another gateway, then the core gateway will re-encrypt the traffic, and send it through the other tunnel. If we have two collections of gateways under two administrative domains, then each domain has its own core, and the administrators only need to define an IPsec tunnel between the two cores. This tunnel is often referred to as a "trunk".

One problem with stars and trunks is that it creates a high load on the core gateways as well as on the trunk connection. This load is both in processing power and in network bandwidth. A single packet in the trunk scenario can be encrypted and decrypted three times. It would be much preferable if these gateways and clients could initiate tunnels between them, bypassing the core gateways. Additionally, the path bandwidth to these core gateways may be lower than that of the path between the satellites. For example, two remote access users may be in the same building with high-speed wifi (for example, at an IETF meeting). Channeling their conversation through the core gateways of their respective employers seems extremely wasteful, as well as having lower bandwidth.

The challenge is how to build large scale, fully meshed IPsec protected networks that can dynamically change with minimum administrative overhead.

3.3. Proprietary Approaches

Several vendors offer proprietary solutions to these problems. However, these solutions offer no interoperability between equipment

from one vendor and another. This means that they are generally restricted to use within one organization. Multiple organizations cannot be expected to all choose the same equipment vendor.

4. Requirements

This section will be completed when the use cases are agreed upon.

5. Security Considerations

The solution to the problems presented in this draft may involve dynamic updates to databases defined by RFC 4301, such as the Security Policy Database (SPD) or the Peer Authorization Database (PAD).

RFC 4301 is silent about the way these databases are populated, and it is implied that these databases are static and pre-configured by a human. Allowing dynamic updates to these databases must be thought out carefully, because it allows the protocol to alter the security policy that the IPsec endpoints implement.

One obvious attack to watch out for is stealing traffic to a particular site. The IP address for www.example.com is 192.0.2.10. If we add an entry to an IPsec endpoint's SPD that says that traffic to 192.0.2.10 is protected through peer Gw-Mallory, then this allows Gw-Mallory to either pretend to be www.example.com or to proxy and read all traffic to that site. Updates to this database requires a clear trust model.

More to be added.

6. IANA Considerations

No actions are required from IANA for this informational document.

7. Acknowledgements

Many people have contributed to the development of this problem statement and many more will probably do so before we are done with it. While we cannot thank all contributors, some have played an especially prominent role. Yoav Nir, Jorge Coronel Mendoza, Chris Ulliott, and John Veizades wrote the document upon which this draft was based. Geoffrey Huang, Suresh Melam, Praveen Sathyanarayan, Andreas Steffen, and Brian Weis provided essential input.

8. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC4301] Kent, S. and K. Seo, "Security Architecture for the Internet Protocol", RFC 4301, December 2005.

Author's Address

Steve Hanna
Juniper Networks, Inc.
1194 N. Mathilda Ave.
Sunnyvale, CA 94089
USA

Email: shanna@juniper.net