

IPsecME Working Group
Internet-Draft
Intended status: Informational
Expires: January 11, 2013

S. Hanna
Juniper
V. Manral
HP
July 10, 2012

Auto Discovery VPN Problem Statement and Requirements
draft-ietf-ipsecme-p2p-vpn-problem-02

Abstract

This document describes the problem of enabling a large number of systems to communicate directly using IPsec to protect the traffic between them. It then expands on the requirements, for such a solution.

Manual configuration of all possible tunnels is too cumbersome in many such cases. In other cases the IP address of end points change or the end points may be behind NAT gateways, making static configuration impossible. The Auto Discovery VPN solution is chartered to address these requirements.

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 January 11, 2013.

Copyright Notice

Copyright (c) 2012 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of

publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	3
1.1. Terminology	3
1.2. Conventions Used in This Document	4
2. Use Cases	5
2.1. Endpoint-to-Endpoint P2P VPN Use Case	5
2.2. Gateway-to-Gateway AD VPN Use Case	5
2.3. Endpoint-to-Gateway AD VPN Use Case	6
3. Inadequacy of Existing Solutions	7
3.1. Exhaustive Configuration	7
3.2. Star Topology	7
3.3. Proprietary Approaches	8
4. Requirements	9
4.1. Gateway and End Point Requirements	9
5. Security Considerations	10
6. IANA Considerations	11
7. Acknowledgements	12
8. Normative References	13
Authors' Addresses	14

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 and the requirements on a solution to address the problem. 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 device 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.

Hub - The central point in a star topology, generally implemented in a gateway

Spoke - The edge devices in a star topology, implemented in endpoints or gateways

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. Such a usecase also enables connectivity when both endpoints are behind NAT gateways. Such usecase should allow for seamless connectivity even as Endpoints roam, in being or away from gateways.

In a hub and spoke topology when two end-points communicate, they must use a mechanism for authentication, such that they do not expose them to impersonation by the other spoke endpoint.

2.2. Gateway-to-Gateway AD VPN Use Case

A typical Enterprise traffic model is Hub and Spoke, with the Gateways connecting to each other using IPsec tunnels.

However for the voice and other rich media traffic that occupies a lot of bandwidth and the traffic tromboning to the Hub can create traffic bottlenecks on the Hub and can lead to a increase cost. It is for this purpose Spoke-to-Spoke tunnels are dynamically created and torn-down.

The Spoke Gateways can themselves come up and down, getting different IP addresses in the process, making th static configuration impossible.

Also for the reasons of cost and manual error reduction, it is desired there be minimal or even no configuration on the Hub as a new Spoke Router is added or removed.

In a hub and spoke topology when two spoke gateways communicate, they must use a mechanism for authentication, such that they do not expose them to impersonation by the other gateways spoke.

2.3. Endpoint-to-Gateway AD VPN Use Case

An endpoint should be able to use the most efficient gateway as it roams in the internet.

A mobile user roaming on the Internet may connect to a gateway, which because of roaming is no longer the most efficient gateway to use (reasons could be cost/ efficiency/ latency or some other factor). The mobile user should be able to discover and then connect to the current most efficient gateway without having to reinitiate the connection.

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). Such a solution is also limited by the smallest endpoint/ gateway, as the same exhaustive configuration is to be applied on all endpoints/ gateways. A more dynamic, secure and scalable system for 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 "Hub gateways", while the rest of the systems (whether endpoints or gateways) are defined as "spokes". The spokes never connect to other spokes. They only open tunnels with the core gateways. Also 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.

This solution however does not work when the spokes, get dynamic IP address which the "core gateways" cannot be configured with. It is also desired that there is minimal to no configuration on the Hub as the number of spokes increases and new spokes are added and deleted randomly.

Another 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 to build a large scale, 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, and it is harder to move off such solutions as the features are not standardized. Besides multiple organizations cannot be expected to all choose the same equipment vendor.

4. Requirements

This section is currently being updated and hence under flux.

4.1. Gateway and End Point Requirements

1. For any network topology (whether Hub-and-Spoke or Full Mesh) Gateways/ end points MUST allow for minimal configuration changes when a new Gateway or end-point is added, removed or changed. The solution should allow for such configuration on a global basis.
2. Gateways/ end-points MUST allow IPsec Tunnels to be setup without any configuration changes, even as peer addresses gets updated every time the device comes up.
3. Gateways MUST allow tunnel binding, such that applications like Routing using the tunnels can work seamlessly without any updates to the higher level application configuration i.e. OSPF configuration.
4. In a Hub-and-Spoke topology, Spoke Gateways/ en-points MUST allow for direct communication with other Spoke Gateways/ end-points, using authentication that does not expose them to other Gateway Spoke.
5. Gateways SHOULD allow for easy handoff of sessions in case end-points are roaming and cross policy boundaries.
6. Gateways SHOULD allow for easy handoff of a session to another gateway, to optimize latency, bandwidth or other factor, based on policy.
7. Gateways/ End-points MUST be able to work, behind NAT boxes.

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, Yaron Scheffer, 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.

Authors' Addresses

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

Email: shanna@juniper.net

Vishwas Manral
Hewlett-Packard Co.
19111 Pruneridge Ave.
Cupertino, CA 95113
USA

Email: vishwas.manral@hp.com