



Towards Better Key Exchange Performance in IPSec-Based VPNs

Yongqing Han

Cisco Systems, Inc., 275 East Tasman Drive, San Jose, CA 95134, USA, frankhan@cisco.com

Dorina C. Petriu

Carleton University, Systems & Computer Eng., Ottawa, ON, Canada, K1S 5B6, petriu@sce.carleton.ca

George Yee

National Research Council Canada, 1200 Montreal Road, M-50, Ottawa, Canada, K1A 0R6, george.yee@nrc-cnrc.gc.ca

ABSTRACT

Virtual Private Networks (VPNs) provide an inexpensive and scalable solution for the transfer of sensitive data through an unsecured network by creating a "tunnel" from sender to receiver. One of the most popular protocols for creating VPNs is the IPSec protocol suite, where secure key negotiation and exchange must be done first, before any encryption of data can take place. This article examines the latest VPN technologies focusing on one of the factors that have an effect on VPN performance and scalability, namely security key management. A new aggregation key exchange approach compatible to current technologies is proposed for improving the key exchange performance in large VPN systems. The new approach represents a trade-off between performance and security. A simulation model based on the Network Simulator (ns) was developed for this new approach. Simulation experiments for various scenarios were conducted and their results were compared to the traditional key exchange scheme.

INTRODUCTION

The Internet, as a global public network, provides an ideal backbone for data communication due to its low-cost and ubiquitous access. Many companies and agencies are using Virtual Private Networks (VPN) to build secure networks over the Internet for their private use. A VPN is formally defined as a class of services using a shared network to emulate the characteristics of a private network, expressed in terms of requirements for performance, reliability, security and quality of service [1].

VPN solutions are designed to deal with these requirements. Using special tunneling protocols and effective encryption procedures, data integrity and privacy is achieved in point-to-point connections. IPSec is a popular protocol for building secure VPNs over the Internet that assure data integrity, authentication, and privacy [7]. IPSec uses a variety of protocol exchanges and encapsulations at tunnel endpoints to authenticate and encrypt user data packets forwarded across the public Internet [12].

Although VPN technologies over the Internet are very promising, many specialists have serious concerns with the scalability and security of these techniques [1]. The scalability implications affect the following aspects: memory used to maintain per-VPN or per-site state, processing power, and management load.

This article focuses only on one of the factors that affect the VPN performance and scalability, namely security key management. Although the latency of encryption and decryption may be felt by the end user as the most expensive operations related to a secure communication, the management of a very large number of security keys (as is bound to happen when the number of VPNs and sites per VPN increase) will put a high load on the network nodes responsible for it. While encryption/decryption are performed at the user's end, key management is concentrated in some network nodes, adding to their load and raising the potential for node bottleneck. Moreover, for a high-security level it is recommended that the keys be replaced very often (the extreme case being a new key for every message exchanged). Considering, for

example, that the process of negotiating a new key takes in IPSec from three to six messages [7], frequent key replacement may have a strong impact on the traffic levels in the network.

A new Aggregation Approach for key exchange, compatible with current technologies, is proposed in the paper for improving the key exchange performance in large VPN systems. The new approach represents a trade-off between performance and security. A simulation model based on the Network Simulator (ns) [5] was developed for the proposed approach. Simulation experiments for various scenarios were conducted and their results were compared to the traditional key exchange scheme.

This paper is organized as follows: Section 2 presents the most common IP-based VPN models, Section 3 describes the proposed aggregation approach, Section 4 analyzes the security implications, Section 5 discusses the simulation model and its results, and Section 6 presents the conclusions.

VPN MODELS

A typical VPN consists of a number of geographically dispersed customer sites, which are attached to Customer Edge (CE) devices and communicate with each other via a shared public network. Each CE is directly connected to a Provider Edge (PE) device. In terms of size, a reasonable estimate for the number of PE in a public network is 50, where each PE can support on average 500 CEs [1]. The general strategy used in today's VPN models is to concentrate the VPN intelligence at the edges of the core network, leaving the core network elements unaffected [1]. The VPN deployed today are classified as *customer premises equipment-based (CE-to-CE-based)* and *network-based VPN* (also known as *PE-to-PE-based*) [1], as illustrated in Figure 1.a and 1.b, respectively.

In CE-to-CE VPN, all the VPN routing and tunnel setup are implemented and maintained by the customer equipment (a.k.a. VPN gateway). The provider has no knowledge of a customer's VPN routing or addressing scheme, and sees only normal IP packets traveling through the shared network [12]. The drawback of this VPN model is the expensive and heavy management load on the customer gateway. It gives relative poor performance on key negotiation and exchange, especially for large networks. However, it gives strong protection of inter-site traffic through the Internet [1].

With PE-to-PE VPN, customer routers need not implement VPN specific functions like tunneling. Customer sites are connected to a Provider's Edge (PE) device through a CE. All the secure tunnels are established between PEs. A PE will be responsible for information exchange and encryption/decryption. A significant disadvantage is the unprotected link between CE and PE.

In this paper we assume that the IPSec protocol suite is used for realizing secure communication in different VPN models. IPSec is designed to provide interoperable, high quality, cryptographically-based security for IPv4 and IPv6. It offers data authentication, protection

against replay (integrity), and algorithms for encryption/ decryption (confidentiality). These services are provided at the IP layer, to protect the IP and upper layer protocols. IPsec is comprised of three basic protocols: IKE, AH and ESP. We are interested especially in the Internet Key Exchange (IKE) protocol, which allows communicating parties to negotiate methods of secure communication.

AGGREGATION APPROACH FOR KEY EXCHANGE

A new approach to key management, named Aggregation Approach for Key Exchange, is proposed in this section for the network-based IP VPN. First of all, we assume that asymmetric encryption is used. With this scheme, each user can generate a public/private key pair [4][8]. The sender and the recipient exchange only the public keys over the public Internet, and each holds its own private key. The sender uses the peer’s public key to encrypt the data and forward the packet to the recipient; the recipient uses his own private key to decrypt the data when he receives the packet.

In network-based IP VPNs, the provider edge (PE) router is the entity able to manage all key negotiations on behalf of the users associated with it. As illustrated in Figure 1.c, a group of users of VPN1 are located at different sites and are connected to PE1 through the gateways CE1, CE2, and CE3, and to PE2 through CE5. When any of the users in CE1, CE2, and CE3 initiates a VPN communication to users in CE5, the two provider edge devices, PE1 and PE2, will negotiate, generate the public/private key pairs, and exchange the public keys and other information, such as the encryption and authentication algorithm, for this communication. This is done by using the IKE protocol, which has two phases: in phase one a secure channel is set up between two PE peers, and in phase two the public/private key pairs are negotiated, generated, and the public keys exchanged [7].

In the proposed Aggregation Approach, two PE peers will establish an encrypted tunnel between themselves by executing IKE phase one, and then will execute one or more phase two in order to exchange the keys and assign them to the connected gateways. The following alternatives are possible:

- a) *Strong aggregation/lower security level:* use the same public/private key pair for all VPNs connected to a PE. For example, in Figure 1.c, PE1 shares the same key pair with CE1, CE2, CE3 and CE4.
- b) *Medium aggregation/better security level:* a key pair is shared only among the CEs representing the same VPN. The PE peers will execute a new IKE phase two for each VPN connected to them. For example, in Figure 1.c, PE1 will share a key pair with CE1, CE2, CE3, and another key pair with CE4. It is also possible to define security groups (divide the VPN users at a site in different

groups) and have an IKE phase two key exchange for each pair of groups. The security implications of these alternatives are discussed in the next section.

In the proposed Aggregation Approach, the security-related responsibilities are divided between PE and CE. While the PEs have the responsibility to setup secure channels and to negotiate and exchange security keys and other necessary information, the CEs are responsible for the actual encryption and decryption of data, by using the keys assigned by the corresponding PE. In this way, the proposed approach reduces the management load by maintaining fewer secure channels and by managing fewer security keys, without compromising too much the end-to-end protection of user data.

Assume that a public network has $p=50$ PEs, and each PE supports on average $c=500$ CEs [1]. In the “traditional” key management approach, the number of CE-to-CE secure channels that need to be created is $cp(cp-1)/2$ (almost 315 million), and the number of public keys need to be exchanged is $cp(cp-1)$ (almost 625 million). On the other hand, in the proposed Aggregation Approach the number of PE-to-PE tunnels is significantly reduced to only $p(p-1)/2 = 1225$ according to the previous estimate, and the number of keys exchanged to $p(p-1)=2450$.

However, the advantage of exchanging less information over the public network comes at a price in terms of security capabilities, as discussed in the next section.

SECURITY ANALYSIS


We identify here some of the security weaknesses of the Aggregation Approach for PE-to-PE VPN, discuss how these weaknesses can be overcome, and give some additional ways of improving security.

Since the link between CE and PE is unprotected, it is possible that, when the public/private key pair is transmitted by the PE to the CE, the keys may be compromised on this link. This will threaten the entire set of traffic flows within the secure channel between the respective PEs.

We propose the following way around this: a) prior to key transmission, each CE attached to a PE requests a certificate from the PE and the CE-PE pair then sets up a secure channel between themselves using SSL [3], b) the PE then encrypts the public and private keys using the SSL session key for transmission to each CE. Note that the PE will use a different session key corresponding to each CE. Also, the public key transmitted by the PE to the CE in the certificate for SSL setup is the PE’s public key for encrypting data meant for the PE itself, with the certificate supplied by some central CA (Certification Authority). In this way, a PE can securely transmit its generated public/private key pair to each CE that is required to share them.

For a given PE with CEs that represent different VPNs, users of these different VPNs may be able to eavesdrop on one another depending on how the PE generated public/private key pair is shared among the CEs. There are 3 cases: a) the key pair is shared among CEs representing different VPNs, b) a key pair is shared only among CEs representing the same VPN, as a new key pair is generated for each different VPN (this can be done with new second phases of IKE), and c) combinations of a) and b), i.e. some different VPNs share the same key pair, others have their own distinct key pair. Case b) is similar to the CE-to-CE VPN model and gives the same level of security, since users of the same VPN are in a trusted environment and can use the same key pair. However, case b) still gives a performance enhancement over the CE-to-CE type since only 1 key pair is generated for different CEs of the same VPN as opposed to a new key-pair for each CE as in the CE-to-CE type of VPN. Case a) and to some extent, case c), are the problem cases for they allow users of different VPNs to eavesdrop on one another and this should definitely not be allowed as it goes against the very idea of having a VPN in the first place. To this we have no clean solution. It truly is a security-performance trade-off: we can have more security as in case b) but at the cost of performance (more key pairs have to be generated and more key exchanges need to take place). We can, however, lessen the seriousness of this weakness, by postulating that when the same key pair is shared among CEs representing different VPNs, the users of these different VPNs all belong to a trust community and they can share the same key pair. This can happen in real life, for example, where users from different

Figure 1. VPN Models



enterprises belong to a trust community for the purpose of collaborating on a common project.

Another security weakness is that an attacker can compromise the CE, the PE, or the link between the CE and PE (man in the middle attack) and redirect the packets to other destinations. This type of attack is not particular to the Aggregation Approach for PE-to-PE VPNs, but Internet routers and links are susceptible as well. One defense against these attacks is to use a form of onion routing [10] and encrypt the address of the receiver in the inner most layer of an onion, which is then covered with multiple layers of encryption where each layer only defines the next hop in a hop-by-hop routing scheme. Each node that receives the onion would decrypt its corresponding layer to discover the next hop in the route. The last node in the route decrypts the receiver's address. Under such a scheme, an attacker would find it next to impossible to redirect packets. Of course, the use of such a scheme incurs overhead and we arrive once more at the security-performance tradeoff.

To complete this security analysis, we point out two additional ways to improve security for our Aggregation Approach PE-to-PE VPN. The first way is to allow for groups of users to be grouped into communities of trust within which they can share the same key pair. A grouping would be done for various purposes and needs, such as for the purpose of collaboration. The PE would generate a different key pair for use by each group. This offers flexibility in that a group could be very large, where security needs are low, to very small, where security needs are high. We could even have groups of one person each for even greater security. A second way to improve security is to allow users to set up IPSec secure channels to their communication partners through the PE-PE secure tunnel. This can be done quite easily as IPSec comes with MS Windows 2000, for example. In this way, there is a high security double layer of encryption traversing the PE-PE secure tunnel. Similar multi-level encryption schemes can be defined, but again at the cost of performance.

SIMULATION MODEL

A simulation model for the traditional and aggregation approach (case a) for key exchange has been designed and implemented by using the Network Simulator (ns) [5]. The simulation analysis took into account the effect on key management performance of three key experimental factors: network size, node connectivity degree and lifetime of a secure channel.

The network topology used in the simulation experiments follows the "transit-stub" model and was generated with the tool GT-ITM (Georgia Tech Inter-network Topology Model), embedded within the ns package [2]. The parameters of a topology are the network size, and the connectivity degree (defined as the average probability of connection between two nodes). Six random network topologies were generated, characterized by three sizes (52, 108 and 150 nodes) and two connectivity degrees (low at 2.68 and high at 3.69). The lifetime of a secure channel was measured by the number of single communication sessions

Fig. 2. Average Key Exchange Time Vs. Network Size for different key exchange approaches and connectivity degrees

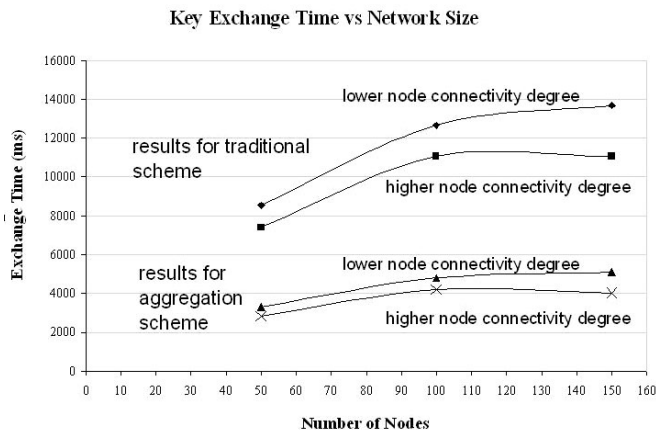
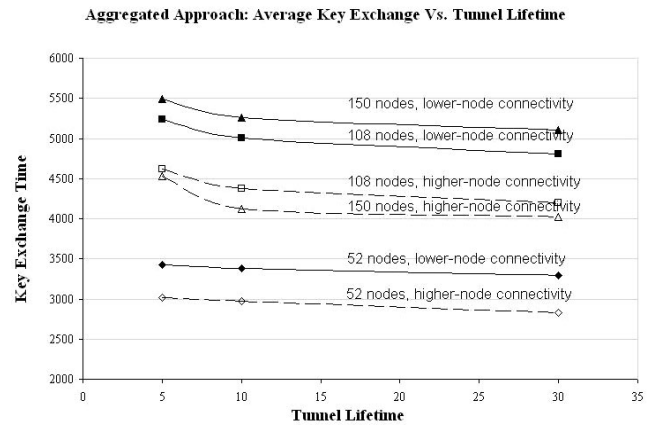


Fig. 3. Aggregated Approach: the effect of Tunnel lifetime on the Average Key Exchange Time



aggregated into the same channel, rather than in time units. After a certain number of sessions have been aggregated, the channel is considered no longer secure, and it is torn down and re-established if needed. The lifetime values used in the simulation experiments were 5, 10 and 30. The traffic generator chosen for the simulation experiments is attached randomly to the nodes in the "stub" domain, and generates traffic according to an exponential "on" and "off" distribution: packets are sent at a fixed rate in every "on" period and no packets are sent during the "off" period. The numerical values used for the traffic generator were 500 ms for the mean duration of the "on" and "off" periods, and 100Kbps arrival rate in the "on" period. The simulation results were obtained with confidence intervals of 4% of the mean or better at 98% confidence level.

Figure 2 gives the average key exchange time for both the traditional and aggregation approaches, for different network sizes and connectivity degrees. In these cases, the keys were considered valid until the communication session was completed. The following can be concluded:

- The average key exchange completion time is reduced significantly when the aggregation approach is applied. In the cases shown in the figure, the improvement is of approximately 60%.
- For the same network size, raising the connectivity degree improves the key exchange and negotiation time. This effect is stronger in larger networks.

Figure 3 show the effect of tunnel lifetime on key exchange performance for different network sizes. As expected, the overhead of executing the key negotiation and exchange is larger when the tunnel lifetime is shorter. Also, higher connectivity degrees lead to lower key exchange times. From the figure it can be seen that the frequency of changes in tunnel lifetime has a higher impact in larger networks.

CONCLUSIONS

This paper studied the trade-off between performance and security for the key negotiation and exchange in IPSec-based VPN networks. A new Aggregation Approach for Key Exchange was proposed in order to improve the key exchange performance in large VPN systems. The security implication of the new approach were discussed. A simulation model based on the Network Simulator (ns) was developed, and simulation experiments for various scenarios were conducted and analyzed.

This is a performance and security trade-off solution targeted at improving the performance of secure key exchange. Depending on the customer's request, different levels of performance and security may be engineered into the VPN design as needed.

REFERENCES

J. D. Clercq, O. Paridaens, "Scalability Implications of Virtual Private Networks", IEEE Communications Magazine, May 2002.

- K. Calvert, E. Zegura, "GT Internetwork Topology Models (GT-ITM)" Georgia Institute of Technology, 1996, <http://www.cc.gatech.edu/fac/Ellen.Zegura/gt-itm/gt-itm.tar.gz>.
- W. Chou, "Inside SSL: The Secure Sockets Layer Protocol", IT Professional, Vol.4, Issue 4, Jul/Aug2002 pp. 47 -52.
- W. Diffie, M. Hellman, "New Directions in Cryptography" IEEE Transactions on Information Theory, Vol. IT-22, No. 6, November 1976.
- K. Fall and V. Kannan, 2000, "ns Notes and Documentation" <http://www.isi.edu/nsnam/ns/doc/index.html>
- B. Gleeson, A. Lin, J. Heinanen, G. Armitage, and A. Malis, "A Framework for IP-Based Virtual Private Network", RFC 2764, Feb. 2000.
- D. Harkins, and D. Carrel, "The Internet Key Exchange (IKE)", RFC 2409, Nov 1998.
- M. E. Hellman. "An Overview of Public Key Cryptography", IEEE Communications Magazine, November 1978 - Vol 16, Number 6.
- S. Kent, R. Atkinson, "Security Architecture for the Internet Protocol", RFC2401, 1998.
- L. Korba, R. Song, and G. Yee, "Anonymous Communications for Mobile Agents", Proceedings, 4th International Workshop on Mobile Agents for Telecommunication Applications (MATA 2002), Barcelona, Spain, October 2002, pp. 171-181.
- Dave Kosiur, "Building and Managing Virtual Private Networks", John Wiley & Sons, New York, 1998.
- Metz, C., "The latest in Virtual Private Networks: Part I", IEEE Internet Computing, pp. 87-91, Jan/Feb. 2003.
- R. Perlman, and Charlie Kaufman, "Analysis of the IPsec Key Exchange Standard", Tenth IEEE International Workshops on Enabling Technologies: Infrastructure for Collaborative Enterprises, pp.150-156, Massachusetts, 2001.
- R. Perlman, and C. Kaufman, "Key Exchange in IPsec: Analysis of IKE", IEEE Internet Computing, Nov/Dec 2000.
- E. Rosen, J. De Clercq, O. Paridaens, Y. T'Joens, C. Sargor, "Use of PE-PE IPsec in RFC2547 VPNs", IETF, Network Working Group, Internet Draft, August 2002.

0 more pages are available in the full version of this document, which may be purchased using the "Add to Cart" button on the publisher's webpage:

www.igi-global.com/proceeding-paper/towards-better-key-exchange-performance/32435

Related Content

Implementing Enterprise Resource Planning

Kijpokin Kasemsap (2015). *Encyclopedia of Information Science and Technology, Third Edition* (pp. 798-807).

www.irma-international.org/chapter/implementing-enterprise-resource-planning/112473

Virtual Communities of Inquiry

Pellas Nikolaos (2015). *Encyclopedia of Information Science and Technology, Third Edition* (pp. 7564-7573).

www.irma-international.org/chapter/virtual-communities-of-inquiry/112458

Theoretical Constructs and Relationships in Information Systems Research

Brent Furneaux and Michael Wade (2009). *Handbook of Research on Contemporary Theoretical Models in Information Systems* (pp. 1-17).

www.irma-international.org/chapter/theoretical-constructs-relationships-information-systems/35821

Method to Reduce Complexity and Response Time in a Web Search

María R. Romagnano, Silvana V. Aciar and Martín G. Marchetta (2015). *International Journal of Information Technologies and Systems Approach* (pp. 32-46).

www.irma-international.org/article/method-to-reduce-complexity-and-response-time-in-a-web-search/128826

The Future of High-Performance Computing (HPC)

Herbert Cornelius (2018). *Encyclopedia of Information Science and Technology, Fourth Edition* (pp. 4004-4017).

www.irma-international.org/chapter/the-future-of-high-performance-computing-hpc/184108