

Intel Architecture Labs

Scalable Deployment of IPsec in Corporate Intranets



Intel Architecture Labs Internet
Building Blocks Initiative
Written by Prakash Iyer, Victor Lortz
and Ylian Saint-Hilaire

Legal Information and Disclaimer

Information in this document may be provided in connection with Intel products. No license, express or implied, by estoppel or otherwise, to any intellectual property rights is granted by this document. Except as provided in Intel's Terms and Conditions of Sale for such products, Intel assumes no liability whatsoever, and Intel disclaims any express or implied warranty, relating to sale and/or use of Intel products including liability or warranties relating to fitness for a particular purpose, merchantability, or infringement of any patent, copyright or other intellectual property right. Intel products are not intended for use in medical, life saving, or life sustaining applications. Intel may make changes to specifications and product descriptions at any time, without notice. The material contained herein may be used for informational purposes only.

Copyright © Intel Corporation 2000

* Third party brands and names are the property of their respective owners.

Table of Contents

Executive Summary	1
Glossary	2
A Brief Introduction to IPsec	3
Why Deploy IPsec in Intranet LANs?	3
How Network and Application Security Technologies are Deployed Today	4
IPsec Deployment Strategies	5
Workgroup-Based Deployment Model	10
Considerations for a Policy-Based IPsec Deployment	13
Impact of IPsec Deployment on Network Infrastructure	15
Recommendations and Conclusion	15
References	16
Acknowledgments	16

Executive Summary

Efforts to maintain network security are most often focused on guarding against threats to data resources from outside the enterprise. However, studies have shown that security breaches by people inside the LAN actually occur much more frequently than break-ins from the public Internet. A recent industry study found 55 percent of all network security breaches were from internal sources versus 30 percent from external sources.¹ Internal security breaches can also be very expensive. As business processes become increasingly information- and network-centric, internal network security breaches can even threaten the very survival of a business. For example, consider the potential consequences to a large mail-order business if it became known that a network eavesdropper had stolen and misused its customer credit card information database. Therefore, protecting network traffic within the LAN is a high value proposition.

One promising strategy for protecting valuable network traffic is to encrypt data communications on the LAN. Internet Protocol Security (IPsec), a network-layer security framework defined by the Internet Engineering Task Force (IETF), is rapidly gaining acceptance as an industry standard for encryption-based network security. IPsec provides end-to-end security services such as authentication, data integrity, confidentiality, and anti-replay protection for IPv4 and IPv6 data over public and private networks. Virtual private

networking (VPN) products have been early adopters of IPsec technology to securely connect remote-access users, private branch office networks, and business partner networks (extranets) over the public Internet. As the scope of internal threats gains wider recognition, however, there is growing interest in using IPsec to secure private networks as well.

IPsec is a flexible and relatively complex security framework with many possible configurations. Network administrators contemplating deployment of IPsec on their corporate LAN may need guidance to understand the full implications of many of the choices they will have to make. This paper examines the costs and benefits of deploying IPsec on a corporate Intranet, paying special attention to scalability and the trade-offs associated with different types of authentication and access control mechanisms.

In this paper, we discuss various deployment options and explain how organizations can phase in IPsec incrementally, starting with small pilot groups and a simple, role-based deployment model. If finer-grained policies or IPsec-based access control are needed, organizations can transition to a workgroup-based model by adding a trust infrastructure such as X.509 certificates or Kerberos. For many organizations, however, a simple and inexpensive role-based deployment should be sufficient.

¹ Intel paper NP1452: "IP Security: Deploying Data Protection On the Network"
http://www.intel.com/network/white_papers/ip_sec_deploy/ip_sec.pdf

Glossary

When reading this paper, it is important to have a basic understanding of the following terms:

- Encryption - obscuring the content of a message so only holders of a secret key can decipher and understand it
- Authentication - a process by which communicating peers reliably establish each other's identity
- Access control (authorization) - a process to determine which peers will be granted access to what resources. Since access control is typically identity-based, it requires prior authentication of identities.
- Trust infrastructure - a set of protocols and network services designed to support reliable authentication and access control decisions
- Security policies - a set of rules governing encryption and access control decisions
- Policy infrastructure - a set of protocols and network services designed to distribute and manage security policies across an organization

A Brief Introduction to IPsec

IPsec is a network layer (layer 3 in OSI terminology) cryptography-based security technology. It is defined native to Internet Protocol version 6 (IPv6*) and defines two types of header extensions to IPv4:

- Authentication header (AH) offers connection-less integrity, strong data origin authentication and anti-replay protection. AH is used if strong authentication of the source of data is desired. It does not encrypt IP datagrams, consequently AH provides no privacy.
- Encapsulating security payload (ESP) offers connection-less integrity, data origin authentication, confidentiality and anti-replay protection. ESP offers IP payload authentication as well as encryption.

IPsec defines two basic modes of operation using these headers:

- Transport mode typically is used in peer-to-peer communications (client-to-client or client-to-server with no intervening security gateways) such as in Intranet LANs. In this mode, AH and/or ESP headers may be applied to IP datagrams.
- Tunnel mode typically is used for remote access and site-to-site security, in general, whenever an intermediate endpoint of communication is a security gateway. In this mode, IP datagrams between two communicating peers are tunneled (encapsulated) in outer IP datagrams from an endpoint to the intermediate gateway. AH and/or ESP headers may be applied to the outer IP header.

Note that tunnel and transport modes can be combined.

AH and ESP support many industry-standard cryptographic algorithms. These include authentication codes based on MD-5 and SHA-1, and DES and 3DES for bulk encryption. AH and ESP also are extensible and can easily support newer algorithms as they become available. Note that if ESP is used, it is not necessary to also use AH (ESP includes all of the useful features of AH).

The second major component of the IPsec protocol suite is Internet Key Exchange (IKE). IKE is a protocol to securely authenticate and establish security associations (SAs), including cryptographic algorithms and their modes, keying material, IPsec modes and use of AH and ESP in these modes. The IPsec IKE protocol supports various authentication mechanisms including pre-shared keys, X.509 certificates managed by a Public Key Infrastructure (PKI) and Kerberos. The IKE protocol involves two phases to establish

SAs between two peers. Phase one establishes what is referred to as an IKE SA, a secure transaction that establishes a base set of keys. Phase two SAs essentially are triggered by network flows requiring IPsec protection and are generated using a previously established phase 1 context between the two systems.

Why Deploy IPsec in Intranet LANs?

Many industry studies have shown that a significant percentage of network traffic break-ins occur within the corporate Intranet. It's also important to note that IP is becoming ubiquitous. It runs over almost every physical medium and virtually every network protocol and application runs over IP.

IPsec offers authentication and encryption services to data at the IP layer. Therefore, IPsec can protect any type of IP traffic, regardless of higher layer protocols (e.g., TCP, FTP and HTTP) or applications (e.g., Web browser and Telnet client.) Although IPsec cannot protect data stored on servers, it can afford confidentiality to data flows between peer systems. Rendering network data opaque makes it very difficult for malicious eavesdroppers to determine session boundaries associated with a data flow and to plan and launch targeted attacks on specific data flows. Furthermore, the flexibility of key lengths and encryption algorithms and modes within IPsec allows a network administrator to configure IPsec in order to achieve an optimum balance between processor and network utilization metrics and encryption strength. Without the ability to derive contexts from network flows, it is virtually impossible to plan targeted active and passive network attacks on mission-critical data flows. Due to these characteristics, IPsec is an excellent, flexible tool to build components of a trusted network.

Another example of the need to protect Intranet network traffic originates from the world of electronic commerce. E-commerce Web sites offer good front-end protection to their customers through technologies such as X.509 certificates and secure HTTP. But confidential information such as customer accounts, credit card information, and inventory tables generally flow in the clear between the front-end Web server and backend databases in a shared Intranet. While the databases may be relatively secure, the data transfers to and from them are not. IPsec can be used to prevent employees, suppliers and others who have Intranet accounts from gaining unauthorized access to this information. Thus, IPsec is a building block to enforce information asset protection.

The benefits of deploying IPsec security solutions based on the IPsec framework can enhance several aspects of enterprise security. These could be summarized as follows:

- More secure communications within the firewall.
Most business network communication occurs between servers and client PCs over the corporate LAN. That's also where the greatest internal security threat lies. IPsec can reduce internal security risks by protecting the sensitive data of groups such as human resources or R&D.
- More secure, low-cost extranets and virtual private networks.
Businesses can save money on telephone charges and equipment by creating protected links through the Internet to branch offices, customers, vendors and other business partners. Furthermore, the additional layers of security provided by IPsec enable greater control over the network resources accessible to these externally-connected parties. With IPsec, you are not limited to an all or nothing proposition of completely trusting everyone inside the firewall and trusting no one outside of the firewall.

How Network and Application Security Technologies are Deployed Today

Security technologies are deployed in various forms in corporate LANs today. They can be broadly classified as:

- Application Layer Security
Products such as application proxies in firewalls, Web browser plug-ins, and PGP for e-mail fall into this category. Application layer mechanisms, such as Single Sign-On (SSO) and username-password authentication, enforce application layer access control.
- Session Layer Security
Protocols such as FTP and DNS are protected through extensions defined specifically for them. More generic mechanisms also are available to protect data at the session layer. Examples include Secure Sockets Layer (SSL) [TLS], which is used to protect HTTP traffic and [SOCKS] to enable authenticated firewall traversal.
- Network Infrastructure Security
To transport network data over multiple links, link layer tunneling protocols have been used. Essentially these

protocols encapsulate IP datagrams in Link Layer Protocol (LLP) specific headers. These protocols have defined proprietary or standard mechanisms to afford confidentiality to the data flows. Examples of such protocols are Point-to-Point Tunneling Protocol [PPTP] and Layer 2 Tunneling Protocol [L2TP]. Additionally, technologies such as switched 10/100 Ethernet, network intrusion detection systems (IDS) and distributed firewalls complement these protocols by performing network partitioning and local access control.

- Host intrusion detection systems
These products, which perform security-related tasks such as virus scanning, are not directly related to the network. However, they are important components in any network security solution.

IPsec: Replacement or Complementary Security?

IPsec is expected to complement application and session layer security technologies. Its main purpose within an Intranet is to protect network data flows from inappropriate use. Under certain circumstances discussed later in this paper, IPsec enables network layer access control. IPsec-enabled systems also have distributed firewall functionality, in that data flows can be processed to be forwarded in the clear, dropped (denied) or protected by IPsec. The fact that IPsec is completely configurable and manageable by policies is a key factor that enables its deployment in enterprise networks.

Eventually, a scenario is envisioned in which a corporate firewall will form the first line of defense against outside-in network attacks as well as enforce some form of access control. IPsec will be used to help protect access to the corporate Intranet as well as access within the Intranet. For example, current front-end technologies such as SSL do not protect e-commerce data on the backend, such as between Web servers and backend databases. IPsec is appropriate for use here. Application and session layer security protocols will continue to complement IPsec in providing tighter access control and authorization functions.

What IPsec Cannot Do For You

IPsec does not fully address the security needs of IPsec "multi-hop" applications. Examples include news (NNTP) servers, e-mail (SMTP) servers and relay chat services. In these cases, data are relayed over a hierarchy of servers. Servers in the public domain will continue to communicate in the clear. In the case of e-mail, for example, network data

protection between e-mail gateways and desktop clients can be enforced only on the intranet-side of a firewall or edge router. This appears to break the end-to-end paradigm for IPsec data protection. But application layer security such as PGP and S/Mime can be used to solve this class of security problems.

Furthermore, IPsec does not protect data resident on systems. Technologies such as encrypted file systems will be needed in such cases.

Essentially, IPsec is useful in building closed communities of trust; both endpoints of a conversation must agree a priori on details of a common set of IPsec policies.

IPsec Deployment Strategies

Recognizing that IPsec enhances the trustworthiness of enterprise Intranets, what are the possible strategies to deploy IPsec in such networks? A prudent strategy employs a phased deployment model for IPsec. This approach, shown in Figure 1, enables network administrators to add infrastructure components over time, while continuing to gain the benefits of an IPsec deployment. Depending on your network size and needs, different deployment models may be most cost effective and appropriate.




Figure 1: Deployment Model Alternatives

In this paper, we discuss two major deployment alternatives: role-based and workgroup-based. In the context of these major categories, we also discuss the costs and benefits of also deploying a trust infrastructure and/or a policy infrastructure.

Figure 1 depicts three alternative deployment models for IPsec. The bars in the figure are partitioned and proportioned according to various benefits provided by IPsec features.

This figure is not intended to represent precise numbers or percentages. However, it does reflect approximate benefits and costs of different IPsec deployment models. The sections below provide more detail about each of these deployment models. The key concepts to note at this point are:

- Many of the benefits of IPsec, including privacy and integrity, can be achieved with a simple, low-cost deployment strategy.
- Adding the authentication infrastructure to permit secure identification of peers adds substantially to the cost and complexity of deployment. The primary IPsec benefit that can be achieved with authentication is the capability to tightly focus access control on network resources — the workgroup model.
- Organizations interested in deploying IPsec can start gradually with the simple model and later add authentication infrastructure and tools for a scalable solution to administering workgroup policies.

The Basic Role-Based Deployment Model

In this model, shown in Figure 2, each computer (desktop client or server) is assigned a role and configured independently. The configuration typically occurs at installation and may be subsequently modified if necessary. Each role has a default initiation behavior and a default fallback behavior. The decision about which role to assign to a machine depends on the value of the content it hosts as well as its normal usage as a client or server. The model assumes that each system will exhibit the same security behavior with all peers in the Intranet (servers and/or desktops.) For multi-homed servers, it is feasible to have different roles assigned to different network interface card (NIC) interfaces. In this model, IPsec is used to provide privacy and integrity to network traffic. Remember that only a weak form of authorization (access control) is possible with this model.




Figure 2: Role-based Deployment Model

There are possibly many different possible role-based models. In the model we suggest in this paper, there are four basic roles:

1. No IPsec

When a machine is assigned to be no IPsec, IPsec is completely disabled. The machine will not initiate nor accept any secure connections of any kind. This is the default role for most systems in an unsecured network IPsec.

2. IPsec responder

An IPsec responder always initiates and accepts traffic that is not secured. However, unlike the no IPsec role, an IPsec responder will accept a secure connection if initiated by another computer. Of course, such a negotiation will succeed only if the list of proposals (encryption and/or authentication algorithms and their modes, authentication token, key sizes, etc.) from the initiator finds a match on the responder. Desktop clients may be configured as IPsec responders. In this case, communication between clients will always be in the clear, but the client will communicate securely with a server configured in IPsec always mode. This case treats IPsec as completely discretionary.

3. IPsec initiator

An IPsec initiator always will attempt to initiate IPsec on all outbound traffic flows. This also means that even if an inbound communication flow is initiated in the clear, the response data flow will cause the machine to initiate an IPsec session. However, if the secure initiation fails,

machines will revert (fallback) to communicating in the clear. Desktop clients in an organization can be configured as secure initiators, if all peer-to-peer communication needs to be secure by policy. Servers can be based on this model as well, as long as the objective is to protect network traffic within a group of systems (a logical workgroup.) Given the fallback model in this scenario, additional application level access control mechanisms may be necessary for communication with systems that are configured to not use IPsec services. This case also treats IPsec as discretionary. The difference when compared to the IPsec responder is that all outbound traffic causes IPsec to be initiated.

4. IPsec always

In an IPsec always role, a machine always will initiate and respond securely to all network data flows. If an IPsec negotiation fails in either case, the fallback mode is to deny all traffic to that peer. High value content servers in controlled logical workgroups can be configured in this mode. This is mandatory security on all network exchanges. This mode also is referred to as “lockdown.”

Note that this model does not allow a system administrator to carve out secure enclaves or workgroups on a LAN.

Also, the only cases where two machines will be unable to communicate are when a non-IPsec machine or no security machine wants to communicate to an IPsec always machine or when two machines don't have an overlapping set of security settings (proposals.)

To avoid these two situations, administrators must ensure that:

- A server is not set to IPsec always unless all clients that must access that server also are IPsec enabled, in at least IPsec responder or IPsec initiator modes.
- An organization-wide set of IPsec settings is published and used to minimize occurrences of failed key negotiation.

Role-Based Deployment With and Without Trust Infrastructure

Network administrators must make a decision about whether or not to deploy a trust (authentication) infrastructure, such as PKI for X.509v3 certificates or Kerberos, when using a role-based deployment model. In the role-based model, no attempt is made to control access. The assumption is that applications will manage their own access control. If IPsec-based access control is needed, the workgroup model offers it. By using authentication only, such as certificates, the certificate can be validated as it is issued by a trusted certificate authority (CA) and information can be logged (as in SubjectAltName) from the certificate for auditing.

The role-based model is a tradeoff between security and ease of deployment. In most cases, the role-based model is more than adequate unless you fear a man-in-the-middle attack (the bilateral authentication of the workgroup model reduces the risk of a man-in-the-middle attack).

Role-Based Model Without a Trust Infrastructure

Using a role-based model without trust infrastructure means using pre-shared keys, alternatively referred to as pass-phrases or passwords. There are two choices within this model:

“Pre-shared Key on the Wall”
(also known as “group key”)

In this case, the network administrator publishes or distributes a single well-known pre-shared key for an entire organization. The idea of a widely known password seems to imply that anyone will be able to use this password to eavesdrop on network communications. However, the pre-shared key is used only for authentication of early phases of IKE traffic. The IKE protocol generates secure session keys that are not compromised by the pre-shared key. Since only the two endpoints of the IKE negotiation know the session keys, subsequent communications protected by those keys are secure. With this model, IPsec provides only secure

communication, not access control. IPsec Application servers must employ secondary authentication and access control mechanisms to determine who can or cannot access a server or applications on that server.

Workgroup Pre-shared Keys

This model uses multiple pre-shared keys, each assigned to a different logical workgroup. Unfortunately, this model does not significantly improve network security.

To see why, consider that communication between machines belonging to different workgroups is only possible if each machine knows the password of the other workgroup. Also, knowledge of a workgroup password enables any system to access a server belonging to another workgroup.

Bottom line on role-based model with well-known pre-shared keys:

This model provides very minimal access control, but you do get the benefits of IPsec protecting your network traffic.

Role-Based Model with Trust Infrastructure

X.509 Certificates

X.509 certificates may be used as authentication tokens. Certificates are issued to users and file servers or security gateways. The certificates issued to users typically are stored on the machine and do not “follow the user” from machine to machine. Operationally they behave like system certificates, binding a public key to an IP address, fully-qualified domain name (FQDN) or email address. Note that smart cards can store certificates in a secure manner and thereby enable tying certificates to users rather than machines.

The certificate enrollment process minimally guarantees the issuance of certificates only to users or machines that are employees or assets in the organization serviced by the CA. The certificate enrollment process also distributes the CA’s public-key certificate to each client. This allows the client to “trust” the PKI.

In its simplest form, the usage model for certificates involves verification that the certificate was issued by a trusted CA. Explicit checking of the identity bound to the certificate and applying policies in that context, with or without user intervention does not occur in this model.

Points to note with the deployment of certificates:

- With certificates, IKE will verify the CA (common root of trust) during key negotiation. It also can verify the validity of the certificate and ensure that it has not been revoked, because access to a certificate revocation list (CRL) is necessary to complete the last step. This is a good step toward an identity-based policy management infrastructure. Possession of a trusted certificate guarantees membership in an organization, unless the private key bound to the certificate has been compromised. Certificate revocation and lifetime management help to reduce such risks.

One of the problems with certificate systems is that to be truly secure, either certificate lifetimes have to be very short or the IKE implementations must use CRLs or OCSP (On-line Certificate Status Protocol) to close any security holes. Another problem is that if an IPsec implementation trusts a particular CA, then it will allow anyone holding a certificate signed by that CA into the IPsec network. In other words, be certain that a trusted CA is truly trusted.

A weak form of access control—accepting only trusted certificates—can be enforced in this scenario. However, for that approach to be effective, some of the systems (for example, a group's file servers) must operate in IPsec always mode. In other modes, if the fallback mechanism is to communicate in the clear, using certificates will not even enforce minimal access control with communicating peers.

- Using certificates can enable simple network logging and auditing capabilities. The SubjectName and/or SubjectAltName from a certificate can be stored on a client and used. For IT managers evaluating the cost of PKI deployment and subsequent management, the following arguments may make the case for a phased deployment of PKI.
 - An active man-in-the-middle attack is harder to pull off. It requires physically breaking the network cable to place a computer that actively recomputes and substitutes network packets. Even when a trust infrastructure is in place, attacks are still possible from individuals within the organization with valid identities. In these cases, certificates and their associated private keys have been compromised. But IPsec, especially ESP protected traffic, makes it very difficult to know which traffic is worth attacking anyway.

- Access control to applications can also be augmented by using a layered approach, incorporating operating system passwords and application logins. Deploying an authentication infrastructure and effectively performing life-cycle management on certificates adds cost to deployment. Without effective CRL management or auto-enrollment procedures for certificates, it will be difficult to use certificates with IPsec.

Kerberos

Many companies have been using the Microsoft Windows* domain authentication as a generic access control mechanism for resources such as file systems and network printers. For these companies, Microsoft Windows authentication is the central way to control employee access to corporate Intranet resources.

Microsoft Windows 95/98/NT 4.0 security was based on NT LAN Manager (NTLM). NTLM provides one-way authentication necessary to login to the NT domain controller and gain access to network resources. Unsecured one-way authentication in a multi-access network, such as a LAN, is always subject to man-in-the-middle attacks. Therefore, one-way authentication techniques cannot by themselves be extended for use with IPsec. However, with Microsoft Windows 2000, Microsoft is switching from NTLM to Kerberos, an open standard for bilateral authentication well suited for access control to network resources. An IT manager now can choose to deploy IPsec using Kerberos as the authentication mechanism. Factors that may influence such a decision include:

- If Microsoft Windows security is widely deployed within the organization, the trust infrastructure already is set up. Using it with IPsec offers enhanced security with little extra cost.
- By reusing this trust infrastructure, employees will not have to remember another password, or carry a smart card with an X.509 certificate.
- Because it is hard to keep the user list of a trust infrastructure up-to-date, this often results in a security problem. By reusing the Microsoft Windows domain controller user list, only a single list needs to be changed as users join and depart the organization.

As with most authentication systems, unique identification of a user is an issue not addressed by Kerberos. A broad enough context needs to be established so that the right

access controls can be associated with the correct person or machine. The secondary use of an enterprise address book or domain controller helps alleviate this problem. If the organization plans a complete upgrade to Microsoft Windows 2000, considering deployment of Kerberos-based IPsec now may help ease that transition.

There are some other factors that should be considered when deploying Kerberos:

- While Kerberos is an open standard, Microsoft Windows has its own version of it. If other operating system environments are widely deployed within the organization, Kerberos may not be the best solution. Other platforms must be tested for interoperability with Kerberos in Microsoft Windows 2000.
- Unlike certificates or pre-shared keys, Kerberos is difficult to use for a remote access VPN deployment. Kerberos requires an open channel between the connecting client and the domain controller. This channel is generally not available when connecting to a VPN gateway from outside the Intranet, which requires a proxy in the firewall/VPN gateway machine. This decision is left up to organizations because it involves risks. Enterprises where “single sign-on” is a more important goal than maintaining perfect security may be willing to assume this risk.
- IPsec Kerberos is not yet widely supported by IPsec vendors.

For many organizations, the cost of deploying IPsec with Kerberos is lower than any other trust infrastructure solution. It also fits well within organizations that already have a significant percentage of Windows-based systems.

Bottom line on role-based deployment with a trust infrastructure:

Using certificates instead of pre-shared keys enables a weak form of access control. Identity certificates with systems in IPsec always mode can create small security enclaves, but at that point it may be better to consider the workgroup deployment model.

Role-Based Model With Policy Infrastructure

The basic role-based deployment model is based on a coarse-grained policy strategy. Deploying a policy management

infrastructure has a variety of benefits. The main advantage is that the distribution and updates of security policies can be managed centrally. This eases the task of network management and offers predictability to the behavior of networks with IPsec. Roles need not be defined at installation time, and subsequent updates are easier.

So while policy infrastructure is not absolutely needed with the role-based deployment model, managing a network of several 10s or 100s of nodes is made much easier.

The Need for Exception Rules

The role-based model covers communication to all systems, at least on a per NIC interface basis. As such, these rules cover all TCP/UDP ports. But this model may not be suitable for communication for all application protocols. For example, bootstrap services such as DNS, DHCP and WINS usually are in the clear or protected by protocol-specific security mechanisms. If a machine is setup as an IPsec initiator, communication with these basic service providers will incur IKE timeout penalties. Servers in IPsec always mode will fail to communicate at all. These exceptions make it desirable to install exception policies that translate to runtime exception security associations on role-based systems. Rule ordering is another factor to consider. Since most systems resolve policies on a first match basis, exception rules must precede the default behavior.

In the role-based deployment model managed through a policy infrastructure, users may be allowed to manage IPsec policies through client-side user interface. It is strongly recommended that this be an administrative option only, made available to selected users capable of understanding the implications of misconfigured policies.

Advantages of a Role-Based Deployment Model

Organizations concerned with protecting the integrity and confidentiality of traffic on their LANs, with access control as a secondary concern, should consider the role-based deployment model. Policy management infrastructure can make it easier to deploy and manage this model. It can take advantage of PKI deployment (X.509 certificates or Kerberos) to a limited extent, as described previously. It allows organizations to gradually phase-in other elements of a security infrastructure. Secure machines are free to “roam” around the networks, change name or IP address and maintain full network services. Roles can be set locally by departments to best suit their unique security requirements.

Limitations of a Role-Based Deployment Model

As mentioned earlier, this model offers very little in area of network-level access control. Implementations do not scale well if machines communicate across export control domains, in which case security proposals may have to be adjusted to disallow non-exportable encryption algorithms. A possible solution is to include a proposal list suitable for domestic and export domains.

Role-Based Deployment: An Example Scenario

Figure 3 below shows a corporate Intranet in which groups of systems have been assigned default roles.

Note that there's not much control over who can talk to whom and how in this model. In this case, fine-grained control and a bit of security are being traded for a big gain in ease of deployment. Access control must be enforced by applications using secondary authentication mechanisms.

Workgroup-Based Deployment Model

Another possible strategy to deploy IPsec is the workgroup model. In this model, each user and/or computer is a designated member of an administrated group to which specific security policies are applied. For example, a hospital may have doctors, nurses and administrative staff organized in workgroups, each of which may have specific access privileges and security requirements.

To deploy IPsec in such a situation, first build logical workgroups, each of which contain a list of all the members that belong to that group. A member of a group could be a person, computer or any entity that can be positively identified. Servers and clients are assigned to one of these workgroups. An IP address or a username is not a usable workgroup entity. A certificate, username-password pair or Kerberos token qualify as possible identifiers. Note that weak forms of authentication, such as a widely-known group key, should be avoided with the workgroup model since weak authentication undermines the basis on which policy decisions are made.

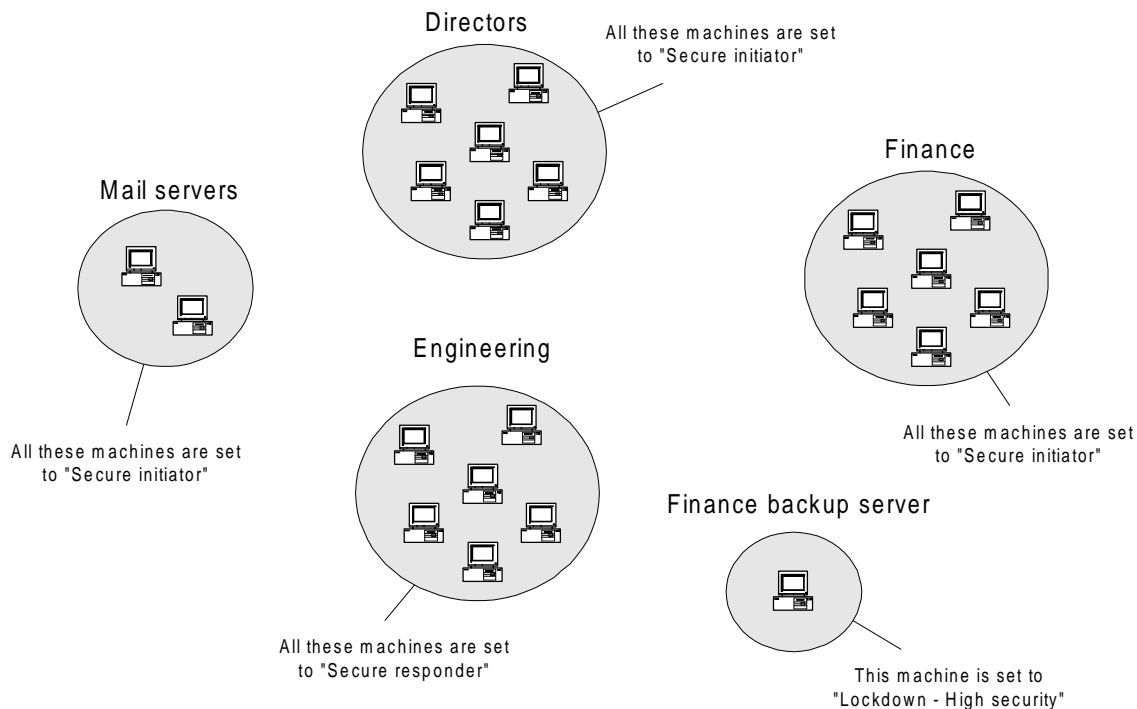


Figure 3: Example Scenario: Role-based Deployment Model

Once workgroups and workgroup members are established, policies can be specified to define secure communication within the workgroups as well as between workgroups. Policies are simply the access control rules that the IPsec implementations will enforce within the workgroup. IPsec Policies can specify that traffic must be sent in the clear, denied or secured using a set of security parameters.




Figure 4: Workgroup-based Deployment Model

In contrast to the role-based deployment model, workgroup-based deployment shown in Figure 4 is dependent on a policy-based network management infrastructure. The complex mechanics of security policy management are beyond the scope of this white paper, but this model clearly offers more tightly controlled security policy management at a potentially significant cost.

Rules for Inter- and Intra-workgroup Communication

The workgroup model can offer tightly controlled policy management. Conceptually, each client will possess an ordered list of policy objects. Each policy object will consist of a set of rules to apply to traffic flows between two endpoint lists. Each endpoint list could be an ordered list of IP addresses and/or FQDNs, subnets and address ranges. The source list identifies machines in the source workgroup; the destination list identifies machines in the destination workgroup. It is possible that the rules could degenerate to a

role. With reference to Figure 4, machines in “workgroup 1” could have rules of forms such as the following:

Policy Object 1:
From Workgroup 1 {Endpoint List $S_{11} \dots S_{1n}$ } to Workgroup 1 {Endpoint List $S_{11} \dots S_{1n}$ }
Apply ruleset X11

Policy Object 2:
From Workgroup 1 {Endpoint List $S_{11} \dots S_{1n}$ } to Workgroup 2 {Endpoint List $S_{21} \dots S_{2n}$ }
Apply ruleset X12

Policy Object 3:
From Workgroup 1 {Endpoint List $S_{11} \dots S_{1n}$ } to Workgroup 3 {Endpoint List $S_{31} \dots S_{3n}$ }
Apply ruleset X13

The complexity introduced by this model is that the rulesets must be symmetrical and synchronized across all of the workgroups; mismatches will cause IPsec communication to fail.

Workgroup-Based Deployment and Trust Infrastructure

In the workgroup model, one must carefully evaluate the role of a trust infrastructure. There are several different options to consider.

- Certificates with workgroup identifiers

1. It may be possible to deploy certificates with workgroup identifiers as IPsec approved extensions** or include such identifiers as part of the SubjectName/SubjectAltName in a certificate. In such cases, each protected server should have access to the public key certificate of each workgroup that it can service. IPsec policies on the server will be bound to the workgroup identities. All workstations in a workgroup indicate their membership by using a certificate associated with the workgroup. The certificate enrollment process will include necessary checks before certificates are issued to members in a workgroup. Typically, the certificates serve as machine identities, not user identities. However, biometric or smart card-based certificates could be used to authenticate user identities. Certificates offer a lower level of control.

This approach is not recommended unless group membership can be guaranteed to be relatively stable. Group membership should be stable because every time a user

** The IPsec Working Group has published a PKI requirements draft that describes mandatory and recommended extensions to X.509 v3 certificates for use with IPsec, RFC 2459. Any desired extensions will have to be standardized in the context of that document.

changes groups, a CRL must be issued. Also, IKE does not handle large CRLs very gracefully.

- Certificates from a trusted CA hierarchy

In this case, all that is being verified is that the certificate presented during an IKE phase one setup is from a trusted root CA—one whose public key certificate is available a priori. This essentially degenerates to the role-based model with certificates.

- Certificates with individual identities

Each certificate has a user identity in an IPsec approved extension format. Tight access control is possible through policies bound to individual certificates or lists of certificates. Each server that offers such tightly controlled access needs to store (or have secure access to another system that stores) the public key certificate of every member in an organization that it can possibly serve. Alternatively, a trusted server could be consulted during IKE negotiations to verify group membership for that identity. Clearly this places additional storage requirements on these servers or additional servers, and lookups could hurt performance.

- Integrate policy management with the trust infrastructure

Examples of this approach include the integration of Kerberos with the MS Windows 2000 domain controller or policy servers integrated with certificate authorities. During each IKE phase one negotiation, the client or server could query the policy server for authentication and policy. This has the drawback of extra network round-trips and does not scale well in mixed operating systems environments.

If an organization wants per user access controls, it must pay for this by separately configuring access controls for each user; there are no shortcuts. If they want group level access controls, they can get this at lesser cost, but they also decrease security.

Bottom line on the workgroup deployment model with a trust infrastructure:

- *For small workgroups for which network-based tight access control is desirable, deploy certificates with user identities. Users/systems outside the workgroup will not be able to access servers in these workgroups. Clients still should authenticate based on the root CA.*
 - *If group level policies are adequate, consider deploying certificates with workgroup identifiers as certificate extensions.*
-

Without PKI, a workgroup-based deployment model does not work very well.

The Need for Exception Rules

Just as in the role-based model, bootstrap services, such as DNS, DHCP and WINS may need exception rules.

Advantages of the Workgroup Model Over the Role-Based Model

In controlled environments, the workgroup model affords limited network-layer access control. This model also limits the damage an insider could cause by conducting man-in-the-middle attacks. For example, a doctor will have the credentials to conduct attacks only within his/her workgroup. With tight access control, these attacks become even more difficult.

Limitations of the Workgroup Model

The biggest consideration when deploying this model in medium to large organizations is cost, including the time and expertise needed for constant updating of the list of members in the workgroups and policies to keep them current. A slightly different approach, using a CA per workgroup, could help. Furthermore, this approach can help reduce costs, since each workgroup maintains its own set of users.

Scalability is another issue with the workgroup model. Within large organizations, it may be difficult to keep workgroups up-to-date and to build policies from every workgroup to every other workgroup in the organizations. The number of rules to link every workgroup to every other workgroup in an organization grows very quickly as the number of workgroups grows. Policy server synchronization issues need to be dealt with as well.

Workgroup-Based Deployment: An Example Scenario

In the example, the model shown in Figure 5 depicts a similar organization to the one in the role-based deployment example, using a trust infrastructure. In this scenario, each person or computer is given an identity that everyone else on the network can verify and use to apply the appropriate IPsec policy.

Workgroups may have explicit policies to communicate with other workgroups. In the absence of such rules, machines may revert to the default behavior. An example of the latter case will be applied for communication between systems in the directors' workgroup and the finance workgroup in the Figure 5.

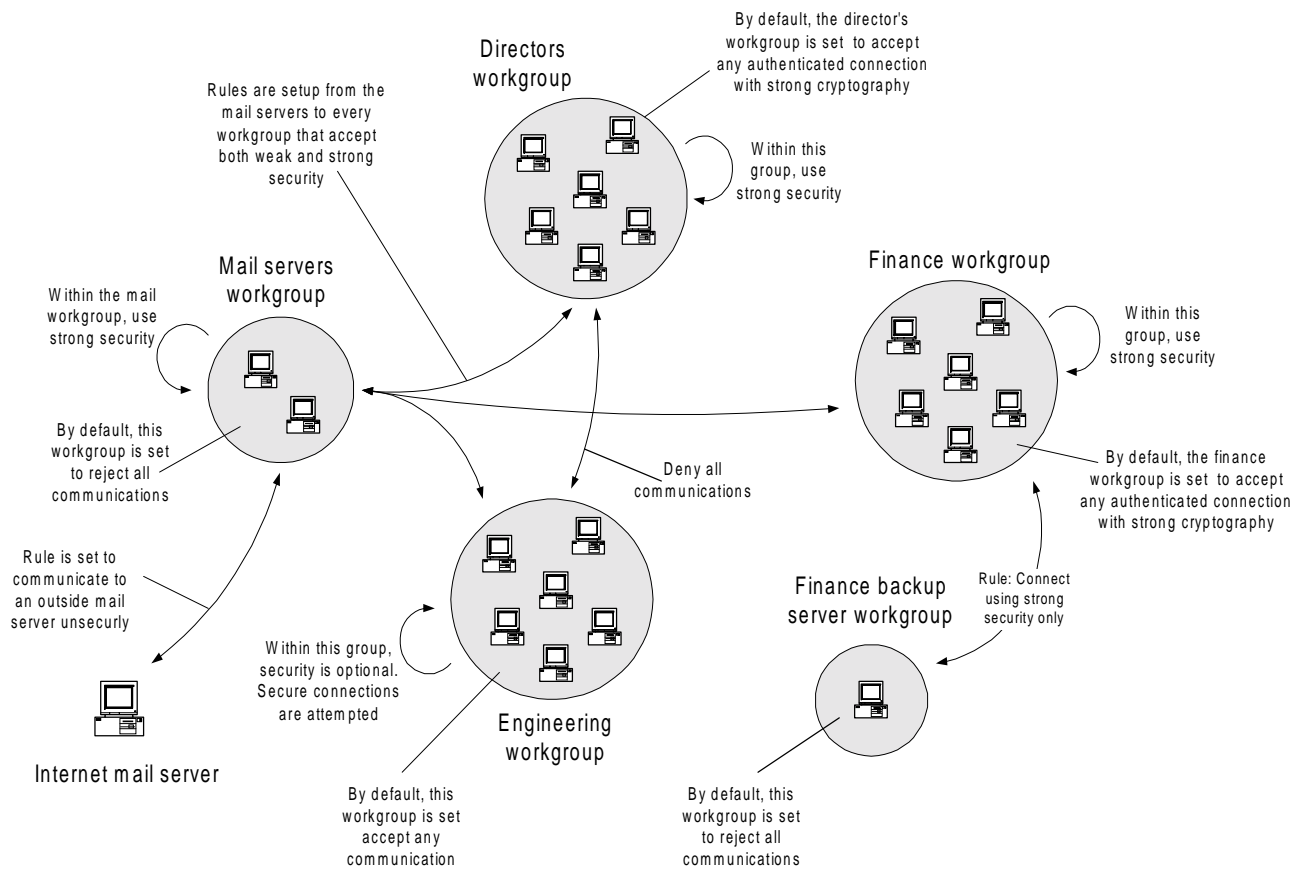


Figure 5: Communication Between Two Workgroups

This model assumes that every computer knows its own policy in relation to all others and that policy is consistent within the network. In a more complex network, policies from one machine may not exactly match policy on another. This may happen if some computers are set up in a different administration domain.

Considerations for a Policy-Based IPsec Deployment

In previous sections, there has been indirect reference to the use of a policy management infrastructure to distribute policies to IPsec enabled clients. Now let's discuss the wisdom of broad guidelines to consider when defining policies.

Security Considerations

In configuring IPsec installations, network administrators will have many choices to make: authentication and encryption algorithms, Diffie-Hellman group sizes, whether to use

perfect forward secrecy (PFS) or not, and so on. Unless they are security experts, some of these choices may be confusing to make. An important point to remember is that once all the traffic in a network is encrypted, it is difficult to detect which traffic may contain valuable information, and breaking into each IPsec connection is nearly impossible. For most organizations, IPsec, even with weak algorithms, will provide a strong deterrent to many network attacks.

The following are a few "rules of thumb" to follow in defining IPsec policies:

- Enable IPsec encryption and authentication

There is a choice to use IPsec for authentication or encryption or both. Generally it is best to make sure a policy is selected that will do both. Otherwise, many of the advantages of IPsec are lost. Actually, it never makes sense to use encryption without authentication. Combining authentication and encryption ensures data integrity and prevents an eavesdropper from reading the data. It may be sufficient to define policies based on ESP alone in the majority of cases.

- PFS is costly, so use it sparingly

When perfect forward secrecy (PFS) is activated, negotiation of session keys is very secure, but this security comes with an added computational cost. On client desktops, security computations may have a very small or no noticeable effect, but servers with a lot of secure traffic may be impacted more significantly. Network interface cards capable of offloading Diffie-Hellman exponentiation operations may be necessary.

Furthermore, if more security is needed than phase one provides, begin using a larger key size and do not use PFS. This is because PFS at most doubles the resources required for breaking the Diffie-Hellman key, and Moore's Law implies that this doubling will occur automatically in about 18 months. On the other hand, increasing the key size by 50 percent requires the adversary to spend many more years than they would have otherwise trying to break through your security.

- Choosing between DES and 3DES

There is no question that 3DES offers greater security than DES. For extremely sensitive data, 3DES is a more appropriate solution. DES also is being deprecated, which may be another reason to favor 3DES. But there are instances where DES may be preferable to 3DES. Some examples are:

- IPsec enabled servers do not have specialized hardware to accelerate 3DES.
- An export version of an IPsec product without 3DES support is being used.

- There is not much difference between MD5 and SHA1

For integrity checking, IPsec implementations generally use MD5 or SHA1. Neither of these algorithms is computation intensive. Since MD5 is known to be weaker, SHA-1 is recommended.

Canned Policies

With some IPsec deployment models, it is possible to simply establish a “canned” set of policies during installation. For example, with role-based deployment a common set of IKE phase one and phase two proposals can be specified for all peers with a “fallback clear” or “fallback deny” policy for interoperability with non-IPsec peers. The canned policy approach is simple, easy to implement and use, robust to misconfiguration, and inexpensive to deploy.

Canned policies are best suited for cases in which policy changes will be infrequent. Do keep in mind that if no

provision is made for changing canned policies after installation, this approach could lead to long-term problems. For example, changes to cryptographic export law or advances in decryption technologies could cause the initial canned policies to become inadequate. Canned policies also are problematic for deployment strategies that require policies to be dependent upon network configurations (per-subnet or per-address policies) because these configurations are likely to change frequently. Canned policies may also become problematic if they cannot be overridden locally to handle special cases. For example, end users may be unable to use IPsec to communicate with peers in external administrative domains such as at another company. Canned policies can also limit interoperability, because some IPsec implementations may not be able to implement a favorite canned policy.

In general, a successful canned policy approach will include some mechanism for updating the policies after installation. This mechanism need not be as sophisticated or expensive as a full-featured policy-based network management system, but it is important to remember that, in the end, an IPsec installation can be no more secure than the security of its policy update mechanism.

Fine-grained Versus Coarse-grained Policies

Another dimension of variability in IPsec policies is their level of granularity. This can range from a single set of IPsec proposals for all connections to individualized proposals for each port and protocol combination of each peer computer. Fine-grained policies provide more flexibility, but they can be complex to understand and expensive to administer. Furthermore, misconfiguration becomes more likely with fine-grained policies. This can lead to security holes. In general, complexity is the enemy of good security. The more complex a policy is, the harder it is to determine whether it actually implements you want it to implement. Therefore, coarse-grained policies should be preferred over fine-grained ones, but it should be possible to override these broad policies for special cases.

Choosing the Right Policy-based Network Management (PBNM) System

There are a growing number of PBNM products on the market. Rather than evaluating or making recommendations on specific products, use the same considerations you would think about to evaluate other types of network services when evaluating PBNM products. Consider whether or not the solution is cost-effective, scalable, extensible and easy to

deploy and manage. Throughout this evaluation process, keep in mind that true security is not a measure of what can be done with the product; rather, it measures what the product will prevent from occurring.

In the realm of data security, the key questions revolve around holes, not features. Since complex systems tend to have more security holes and be more difficult to analyze, deployers of PBNM systems should resist the temptation to simply choose the product with the most features. If an organization does not have in-house expertise in evaluating security products, it would be wise to hire experienced consultants to help with this process.

Impact of IPsec Deployment on Network Infrastructure

A very important issue surrounding IPsec deployment is its potential impact on other parts of the network infrastructure. Protecting network traffic with ESP makes transport protocols opaque to the network infrastructure. This can negatively affect some services already deployed in networks. Clearly, the benefits of securing network traffic via IPsec must be weighed against any possible negative impact on other network services. This is another reason we recommend a phased deployment of IPsec — phased deployment provides network administrators with an opportunity to evaluate the overall impact of IPsec on their particular network and to resolve issues prior to full-scale deployment.

Here are some examples of potential impact of IPsec on network infrastructure.

Firewalls

Firewalls filter or proxy traffic to and from the Intranet. Since ESP packets are encrypted starting from the transport header, firewalls cannot record session information within a packet. The only rule a firewall can enforce is “allow/deny all ESP traffic between X and Y.” End-to-end encryption requires distributed policy enforcement at all the machines participating in the network; it is not safe to allow encrypted end-to-end traffic through the firewall unless you can guarantee that the external node conforms to your internal policies.

Real-time Intrusion Detection

With ESP encrypting packets, it is almost impossible to detect intruders by analyzing network traffic. However, ESP also

makes it more difficult for intruders to gain any advantage or discover any weaknesses by snooping network traffic.

Quality of Service (QoS)

Resource Reservation Setup Protocol (RSVP) session identification is performed using destination address, protocol ID and an optional destination port. The destination port is necessary to distinguish multiple streams of traffic from the same machine. The destination port and the protocol ID are not exposed by IPsec, which breaks RSVP session identification. There are three possible policy driven approaches to this situation:

- IPsec can simply hide the QoS when it encapsulates.
- IPsec can propagate the QoS out of the encapsulation.
- Or IPsec can replace the QoS inside the encapsulation with some policy driven (perhaps user and/or application specific) QoS when doing the encapsulation.

RMON Probes

Remote monitoring probes gather information by snooping on transport and higher layer session or application information in IP packets. These tools break due to encryption.

Congestion Control Protocols

Congestion Control Protocols (e.g., the Berkeley snoop protocol) help reduce congestion on the network. Typically, they need to look at TCP header information, which may be hidden by IPsec.

Solutions to make ESP more “transport protocol friendly” are being looked at in the Internet Engineering Task Force. Improving client-side trusted logging and network auditing capabilities also will help.

Recommendations and Conclusion

For most organizations, a role-based strategy without a trust infrastructure is the safest and least costly solution to maintaining network security. A role-based strategy presents the least impact on the network while delivering most of the security benefits of a more complete, more costly solution. Administrators should phase-in transition to IPsec, starting out with small groups. The organizational pre-shared key and default security policy should be selected and published early so that quick adopters around the organization can all use common settings.

In the first phases, security policies should be set to fallback to clear traffic. The network should be monitored to see if traffic that should be secured is clear. Administrators should use this phase to try new tools to monitor network traffic. They should check how encryption is affecting the behavior of current monitoring tools. RMON probes and other tools will affect remote network traffic differently. Administrators should allow the first phase of security system development to be long enough for timely reporting from users about any related problems they experience with their applications.

In the next phase, local servers accessed by a small number of individuals that all support IPsec can be made to only accept secure connections. This is implemented after making sure that everyone is set up correctly, with no insecure server access.

In the last phase, Web and mail servers can be set to accept only secure connections. Take this step only after wide acceptance of IPsec throughout your organization.

References

For more information, visit the Intel Architecture Labs Internet Building Blocks Initiative on the web at <http://www.intel.com/ial/home/ibbi/>

[**SEC**] [Applied Cryptography](#) by Bruce Schneier, 1996.

[**NP1452**] Intel paper: IP Security: Deploying Data Protection On the Network http://developer.intel.com/design/security/IPsec/NP1452_wrapper.htm

[**ARCH**] RFC 2401 - Security architecture for the Internet Protocol

[**IKE**] RFC 2409 – Internet Key Exchange Protocol

[**ISAKMP**] RFC 2408 – Internet Security Association and Key Management Protocol

[**Oakley**] RFC 2412 – Oakley Key Determination Protocol

[**AH**] RFC 2402 – IP Authentication Header

[**ESP**] RFC 2406 – IP Encapsulating Security Payload

[**TLS**] www.ietf.org/html.charters/tls-charter.html

[**SOCKS**] www.ietf.org/html.charters/aft-charter.html, spiderman.socks.nec.com

[**PPTP**] www.microsoft.com/ntserver/commserv/techdetails/prodarch/understanding_pptp.asp

[**L2TP**] www.tecsec.com/overview%20of%20certification%20systems.htm

Acknowledgments

The authors wish to acknowledge Jesse Walker, John Richardson and Lynn Torrance Redlin for their comments and valuable input.

Document Code IAL_122

Copyright © Intel Corporation 2000

* Third party brands and names are the property of their respective owners.