VIRTUAL ENTERPRISES AND THE ENTERPRISE SECURITY ARCHITECTURE

Tom Haigh

Secure Computing Corp.

Roseville, MN 55113

haigh@sctc.com

Abstract

The emergence of internetworked systems has given corporations and Government agencies the opportunity to share information in unprecedented fashion. This sharing can be distributed across several enterprises. In effect, actual enterprises with shared interests can form virtual enterprises. There are significant security implications in this. An enterprise must not only protect the confidentiality and integrity of its own information; it must also protect the information of the virtual enterprises to which it belongs. We suggest that this can be considered a new security paradigm, the Virtual Enterprise Security Paradigm, which has a significant impact on security policy, architecture, and mechanisms and leads to an Enterprise Security Architecture that is consistent with the trend toward client-server systems and is suitable for the emerging, internetworked environments now found in both the Government and Private Sectors. In this paper we describe the common themes and solution principles that led to this new view of computing and security. We then discuss the policy and architectural issues associated with the paradigm. We finish by drawing some tentative conclusions about the paradigm.

1. Overview

The emergence of internetworked systems has given corporations and government agencies the opportunity to share information in unprecedented fashion.Not only can this sharing be distributed geographically across a single enterprise, either a corporation or a government agency, it can also be distributed across several enterprises. In effect, actual enterprises with shared interests can form virtual enterprises. These virtual enterprises are dynamic. Some may last for years, such as the relationship between a manufacturer and a supplier or customer. Others may last only long enough to accomplish a limited set of objectives, such as a business team organized to perform on a certain project. There are significant security implications in this. An enterprise must not only protect the confidentiality and integrity of its own information, it must also protect the information of the virtual enterprises to which it belongs.

We suggest that this can be considered a new security paradigm that has significant impact on security policy, architecture, and mechanisms. In the past, emphasis has been placed on operating system policies and mechanisms. Even with the advent of networks and distributed computing, the focus has been on policy and mechanisms for securing a single enterprise. This is still the paradigm that guides the thinking of much of the DOD and Internet communities. Firewalls and guards have mechanisms to enforce a single security policy intended to protect the local enclave from the external world. With this old paradigm, cryptography is used to extend the local enclave's security perimeter to other enclaves within the same enterprise and to support the limited exchange of information between enterprises.

The Virtual Enterprise Security Paradigm (VESP) changes this perspective. It integrates several existing concepts and ideas, such as the data enclave concept [1], the notion of a firewall/application gateway [2], and the concept of multipolicy enforcing systems [3]. Application of the paradigm leads to an Enterprise Security Architecture that is consistent with the trend toward client-server systems and is suitable for the emerging, internet-

0-8186-7318-4/95 $4.00

worked environments now found in both Government and Private Sectors. This Architecture provides an enterprise with a cost-effective method for enforcing not only its own security policy but also the security policies of each virtual enterprise to which it belongs. Moreover, both the paradigm and the architecture are flexible enough to address new security concerns that are introduced as new technologies are developed.

In this paper we describe the common themes and solution principles that led to this new view of computing and security. We then discuss the policy and architectural issues associated with the paradigm. We finish by drawing some tentative conclusions about the paradigm. In particular we offer the tentative conclusion that, for the foreseeable future, perimeter protection is a necessary component of all secure internetted systems. This protection is much more complex than that provided by most firewalls available today. It involves complex filtering proxies and servers for network applications as well as the concentration of encryption services at the perimeter. Other cryptographic services, such as signatures, can be provided on individual user work-stations. This protection must be hosted on operating systems with strong, low level access control mechanisms that can be used to protect the encryption and proxies and can ensure that these mechanism are always properly applied to data that crosses the perimeter There is nothing definitive about this paper. It is intended to propose the paradigm and architecture and to present a springboard for discussion.

2. Background

Over the past two years security consultants from Secure Computing have been meeting with potential customers from the Private Sector, from Civil Government agencies, and from the DoD, particularly the Navy and the Air Force. Private Sector customers include financial institutions, healthcare providers, public utilities, and major retailers. Common themes began to emerge from these discussions, and Secure Computing has formalized a perspective and set of principles that we use in recommending solutions to the security problems of our customers. In this section we describe those themes and principles.

2.1 Themes

Users recognize the value of information and have systems in place to support the efficient processing and transmission of information. Many of these systems are essentially Command, Control, Communication, and Intelligence (C^3I) systems. Financial planners require up to the minute information on a variety of data, including information from world financial markets, events in the global business community, and even information on human and natural events that can affect financial decision-making and behavior. Utilities require a combination of near real-time and more static information to control their power generation and distribution. To meet these needs these users have developed a large installed bases of computer and communications hardware and software.

Users have the opportunity to achieve greater connectivity, more functionality, and increased openness in the future. Today this can be accomplished by gaining access to common networks, such as the Internet, and the services these networks provide. In the future the National Information Infrastructure (NII) will provide even greater opportunities. By taking advantage of these opportunities enterprises can:

- more effectively gather information required to perform their missions,
- disseminate information about the enterprise, and
- establish connectivity with expanded communities of interest outside the enterprise, thus forming virtual enterprises of one form or another.

To achieve these objectives, enterprises want to use the full complement of Internet services, from simple e-mail and ftp to more complex World-Wide-Web and distributed client-server

database applications. Moreover, these enterprises are looking at an even broader range of applications, including extensive use of Electronic Data Interchange and Electronic Commerce.

At the same time that enterprises require this increased connectivity, they are constrained by concern about security in this more open, highly connected environment. In the healthcare industry and some portions of the Civil Government sector, such as the IRS, this concern is dictated by privacy laws at both the federal and state levels. In other instances there is concern that information shared among partners in the same community of interest might be observed or even modified by agents hostile to that community of interest. The Navy and Air Force are concerned about the compromise of classified information as well as the corruption of mission critical data. Utility companies have very real concerns that terrorists of one flavor or another might target their power generation plants, particularly nuclear plants, or their distribution control systems. They also worry about preserving the confidentiality of customer and vendor proprietary information stored in their systems. In this environment the enterprise has a responsibility to address the security concerns of every stakeholder in each community of interest with which it shares data. Privacy of customers must be respected. Integrity of critical data must be preserved in the face of threats from outside the enterprise itself. In effect, the organization must enforce not one security policy, it must enforce multiple security policies.

Moreover, these policies cannot be formulated in terms of the simple access control matrices or information flow assertions that have been used in the past to state system security policies. The new policies are very application specific, constraining not only applications a given user can use to operate on various classes of data but also constraining the nature of both the output that each reader can observe and the activities that each user can initiate via these applications. These constraints are dynamic. The constraints on a given user depend more on the enterprise to which s/he belongs and the role which s/he is playing, and possibly present location than on the user's actual identity. As an enterprise changes its relationship with various organizational partners, the applicable security policies will change. The security architecture and mechanisms of the enterprise's information system must address this concern as well as several others.

As enterprises begin investigating the security issues related to increased connectivity and functionality over common networks, and as they begin to develop policies and adopt mechanisms to enforce these policies, they run into several significant roadblocks. First, the end-users understand the opportunities provided by the increased connectivity and functionality but do not adequately assess the accompanying risks. They see security measures as intrusive and as preventing them from doing their jobs as effectively as they otherwise could. Moreover, budgets are limited, and when users are presented with the choice between investing in technology that will enhance their capability to perform their missions and mechanisms that they see as limiting their capability, the choice is obvious to them. To paraphrase one of our DoD customers, "Regardless of what you people say about security, this system will go operational." In the eyes of users, the mission is paramount. Finally, enterprises have a variety of legacy systems that they cannot afford to replace over night, and their users dislike the idea of having to learn a new system in order to do what they were doing quite well with the old one.

This set of concerns:

- the large installed base,
- the desire for increased connectivity, functionality, and openness
- the limited budgets,
- the resistance of end-users, and

- the need to utilize legacy systems,

places requirements on the security solutions that enterprises adopt. At the same time that a solution supports greater connectivity and functionality across the network, it must be largely transparent to the end-users and must be compatible with the enterprise's installed base.

2.2 Principles

Secure Computing has developed the following set of principles that, when carefully applied, provide a good compromise between workable security and transparency for end-users.

- Identify the physical perimeters and defend them.
- Restrict access to data and services to any user outside the perimeter based on the user's organization and/or role.
- Use cryptography to protect and extend the perimeters.
- Utilize a strong, low level TCB access control mechanism to support the previous principles.

None of these principles are entirely new. In the rest of this section we describe the manner in which these principles can be combined and extended to address the concerns raised by the VESP. This discussion provides motivation for the Enterprise Security Architecture introduced in section 3.

Perimeter Defense

Identifying one or more physical perimeters and concentrating security mechanisms at the perimeters rather than distributing them to individual end-user systems accomplishes several things. The security mechanism can be concentrated in an Application Gateway located at the perimeter. The Application Gateway system is discussed in the next section. By applying this principle the security mechanisms are largely transparent to the end-users. They can continue to use the same hardware systems and, to a very large extent, the same applications. In particular, there is no need to modify existing application client software or to add any sort of "trusted" workstations within the security perimeter. Users will feel the constraints imposed by the security mechanisms, but these constraints will be imposed by an Application Gateway system removed from the users and their local application clients. Finally, it is more difficult for uncooperative end-users to subvert or bypass security mechanisms hosted on a remote Application Gateway rather then on their one local systems.

In the rest of this paper, we will adopt Boebert's terminology and call the system within a physical perimeter a data enclave. The term security perimeter will be used to denote the union of one or more data enclaves with a common security policy. Users within a data enclave are presumed to be trustworthy with respect to all the data and resources within the enclave. In the DOD environment, a physically protected, single-level or system high network would be a data enclave. All the users within the enclave would be cleared to the maximum sensitivity level of data within the enclave. In the Civil Government or Private Sectors, a data enclave would be a physically protected network whose users are trusted to respect the confidentiality and integrity of all the data in the enclave. This does not necessarily mean that all the users within the enclave need to or should access all the data within the enclave. Nor does it mean that the enterprise believes all the software exercised within the enclave behaves in a manner consistent with the security policy of the enterprise. Rather, it means that the enterprise believes it has augmented the controls found in COTS systems within the enclave with controls on the perimeter to achieve an acceptable level of risk, without investing in special purpose hardware or software to further control activities within the enclave. At this level of risk, the benefit of connecting the conclave to the outside exceeds the expected cost of the lost confidentiality or integrity for data or services within the enclave.

Besides preserving an enterprise's current hardware and application base, the notion of perimeter protection has the flexibility to provide cost effective solutions to new security problems. As attackers become more sophisticated in their attacks and as end users identify new applications that they require to perform their jobs, the enhanced security measures can be concentrated at the perimeter rather than distributed throughout the corporate network. Thus, the notion of perimeter protection provides for a cost-effective security solution for systems today and into the future.

It is not necessary for an enterprise to have a single security perimeter. Many enterprises naturally contain a hierarchy of nested perimeters, separating users with different organizational roles. This notion of organizational role, both within an enterprise and across enterprises, is important.

Organizational Roles

The concept of organizational role is addressed and extended by the second principle. In talking with users it became apparent that they had a desire to restrict access based on the role that an end-user played with respect to the enterprise. In the utility example, employees in accounting might be allowed limited access to the distribution system only to monitor the actual usage of individual customers. They would be allowed no access to information related to the control of power plants but would need some access to information on the cost of maintenance. Such access restrictions can be enforced by layering security perimeters within the enterprise. The power generation and control subsystem is within its own perimeter, separate from accounting, which is likely to be within a corporate headquarters perimeter. Similarly, each generating plant might be within its own data enclave nested inside the generation and control security perimeter. Accesses of end-users outside a given security perimeter can be controlled either by placing the data which they should access on an Application Gateway system or by limiting the network services they can perform across the Gateway and by filtering the data transmitted across the Gateway as a result of the service.

The notion of organizational role expands to include end-users outside the enterprise. For example, a retailer has a separate relationship with each of its wholesalers, many of whom are competitors and who also sell to competing retailers. The result is a highly connected set of relationships. Although it would be cost-effective for each wholesaler-retailer pair to exchange information electronically, there is considerable risk in doing so unless both parties can ensure the confidentiality and integrity of shared data. The parties must be assured that pricing and inventory information pertinent to that relationship is not divulged to a competitor. The retailer's orders must be placed in a well defined and controlled fashion. Both parties must also protect themselves and other partners from malicious code that might be inadvertently imported from the other party's network.

In a sense each retailer-wholesaler pair forms a virtual enterprise, and it is desirable to create an extended security perimeter that includes selected employees of both enterprises. End-users should have access to that data of both enterprises that they need to do their jobs, and both enterprises have the responsibility to protect the data shared by the other enterprise. The retailer has similar relationships with each of its customers. They share information on the customers' purchases and their credit information. Again, the confidentiality and integrity of that information is important to both parties. Thus, the retailer is part of a large number of virtual enterprises with which it shares information, and the retailer has a responsibility to protect the shared information appropriately. To share the information electronically, the retailer's computer system must extend one or more of its data enclaves to include users from other corporations. This requires the retailer to enforce a distinct information security policy for each virtual

enterprise to which it belongs, restricting access to information based on the enterprise and role of the person requesting access. In order to enforce these restrictions, there must be a mechanism for identifying the user, or at least the user's organization to the retailer. We discuss this next.

Cryptography

Cryptography plays a key role in defending and extending the security perimeter, both across data enclaves within a single enterprise, and across corporate boundaries to establish broader security perimeters. First, cryptographic Identification and Authentication (I&A) provides a means of identifying a user outside the data enclave so that access and filter decisions can be made with confidence. As noted above, in most instances it is not necessary to identify the individual user, it is only necessary to identify the enterprise associated with the user. Thus cryptographic I&A between enclaves is generally adequate. Of course, this would not be the case if it had been established that a particular enterprise had weak internal personnel or procedural controls so that unauthorized individuals from the enterprise were attempting to access data in another enterprise's enclave. When the concern is great, authorized individuals at their workstations can be provided with stronger, remote I&A or their transactions or with digital signature mechanisms that they can use to identify themselves to the remote enclave.

Second, encryption can be used as a means of preserving the confidentiality of information as it moves between two protected sites within the same community of interest or security perimeter. Again, enclave to enclave encryption is generally adequate and, in fact, preferable to encryption at end-user workstations. There are significant operational and technical problems with passing cipher text generated at the workstation outside the enclave. Operationally, the enterprise has an obligation to control the flow of information from its internal network, both to prevent the compromise of information sensitive to the enterprise and to prevent the use of the internal network as a base from which to launch attacks against other enclaves or individuals outside the enclave. That is the reason for establishing security policy and investing in enforcement mechanisms in the first place. To accomplish this, either the enterprise must prohibit end-users from transmitting ciphertext from their workstations or, as an alternative, the enterprise must require that a copy of each message encryption key be made available to the enterprise so that any encrypted message could be decrypted by the enterprise at a later date. Although there is no technical means to enforce this aspect of the security policy, the enterprise can enforce it procedurally.

The technical problem is simpler. Until there are end-user workstations with very strong security mechanisms, for both confidentiality and integrity, it is not possible to adequately protect the encryption subsystem, from compromise of the keying material, or from unauthorized modification, or from being bypassed altogether. The use of co-processor encryption, say on a PCMCIA card, can protect the confidentiality and integrity of the cryptography, thus protecting I&A and digital signature functions, but it cannot assure that data is always properly encrypted before it is released outside the enclave. Thus, the enterprise buys very little by placing encryption capability on untrusted workstations. In fact, this might actually weaken the overall security of the architecture because the proper application of the encryption is so problematic.

There are a few caveats here. In some cases workstation encryption is desirable. For instance, private personnel records must be protected within the enclave. However, it is important to recognize the limitations on the confidence that the encryption mechanism has been properly protected and applied. In the absence of a highly assured, trusted workstation, the encryption would not protect the information from a motivated, reasonably knowledgable insider. It is also the case that

some commercially available client-server application packages now provide encryption capabilities. There are questions with regard to both the protection and application of the mechanism as well as with regard to the ability of the enterprise to monitor and control the flow of information via these applications.

Finally, cryptography can be used to provide a digital signature on outgoing data. This signature can then be checked on the receiving end to verify that the data has not been corrupted in transit or prior to transmission. This is extremely important when importing data from a remote site that is trusted to maintain the integrity of stored data. Consider, for example, a digital library in which each document is signed by its author. This signature ensures that any modification to the document will be detected. This accomplishes two things. First, it provides a high degree of accountability for the contents of the document. If it contains malicious code or incorrect information, there is clear traceability to the author. Second, it protects the integrity of the author and the library. They can be confident that what the reader accesses in the library is the intended document.

TCB Access Control

The final principle, use of a low level TCB access control mechanism, is required to guarantee that the organizational security policy is properly enforced and that cryptography is applied properly. We identify this as a TCB access control mechanism to distinguish it from the application level access control discussed earlier in the paper. Application level access control is likely to be implemented in Application Gateway proxies and servers. The low level mechanism discussed here controls the accesses of processes in a single enclave to data and resources in the enclave.

The strength required for this mechanism depends on the level of concern of the enterprise. A rule of thumb is that it should be at least as strong as any other access control mechanism in the system. It should also be as strong as the cryptography it is intended to protect. Thus, a DoD site that requires B3 or A1 multilevel security and Type 1 cryptography should demand the same strength for the TCB access control mechanism. After all, if this mechanism is defeated, it will be possible for an attacker to bypass the cryptographic subsystem and gain whatever information has been protected cryptographically. On a Unix system where the protection bits are the only access control mechanism, this sort of protection could be provided by careful use of the suid and sgid features [4]. However, the strength of this mechanism is quite weak and is not likely to be satisfactory in any but the most benign environments. Certainly it is not as strong as the cryptographic protection provided by FORTEZZA or RSA and DES. A label-based, table-driven mechanism such as LOCK Type Enforcement [4] or the DTOS extensions to it [5] would be more appropriate protection for either FORTEZZA or RSA/DES.

By low level, we mean a mechanism that is applied as close to the hardware as possible. Control applied at the hardware level, in the Memory Management Unit, would be best. Control applied in the bowels of the Unix kernel would be next best. Application at any higher level in the system opens a vulnerability to attacks by code that can, in effect, tunnel below the mechanism. Once this happens, the hostile code has broad system accesses and could potentially subvert the entire set of system protection mechanisms by modifying either the code or the security database utilized by the protection mechanism. At the very least, the cryptographic subsystem could be bypassed or critical keying material could be exposed.

These four principles, perimeter defense, enterprise based access control, use of cryptography, and protection using strong, low level TCB access control, provide the motivation for the Enterprise Security Architecture. In the next section we describe the Architecture and the Application Gateway concept, which lies at the heart of the architecture.

3. Policy and Architecture

As noted earlier, viewing security from the perspective of the virtual enterprise leads to a new set of policy and enforcement issues. The related Enterprise Security Architecture is intended to address these issues and to apply the principles described in the previous section. Here we present an overview of the Architecture and the role of the Application Gateway in the Architecture. First we discuss issues related to security policy and security mechanisms.

3.1 Policy and Mechanisms

Since so many of the security concerns associated with the Virtual Enterprise Security Paradigm manifest themselves in the network applications, it is necessary to develop security mechanisms for these applications. These take the form of filtering proxies and servers that can be used in conjunction with application layer access control lists (ACLs) and strong network I&A both to sanitize information released from the enclave and to protect the enclave from data driven attacks that are mounted against the enclave from the outside or that intend to use the enclave as a base from which to launch attacks at other enclaves. In this sense the enforcement mechanisms can be viewed as a pipeline of filters with different pipelines for different applications and different virtual enterprises. This is illustrated in Figures 1 and 2.

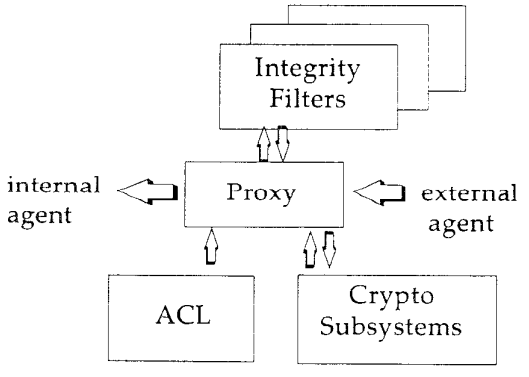

Figure 1. Proxy uses cryptographic I&A to identify external agent, consults ACL to determine if message import is authorized, decrypts all ciphertext, applies message format and source specific filters to ensure integrity of data.

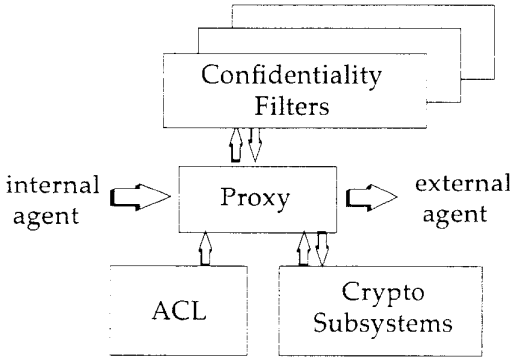

Figure 2. Proxy uses cryptographic I&A to identify internal agent (if necessary), consults ACL to determine if message export is authorized (if necessary), applies format and destination specific filters, including encryption, to ensure that data can only be observed by authorized agents.

There would be one set of filters, primarily concerned with confidentiality, to apply to data leaving the enclave. A second set of filters, concerned with integrity, would be applied to data entering the enclave. Individual filters could be reused in several different pipelines. As shown in the figures, cryptography plays a key role for both incoming and outgoing data and requests.

Below these application layer mechanisms and the cryptography, there would be a layer of supporting security mechanisms. These would include conventional access control and information flow control mechanisms. One of them would be the low level mechanism described in Section 2. This mechanism provides the support necessary to structure the pipelines of proxies and filters and to protect the cryptography. On systems where a multilevel security policy is necessary, there would also be a conventional-set-of-multilevel-security mechanisms. Mechanisms for identifying local users, auditing their activities, and labeling

output are also necessary. In many instances, these should be stronger than those found on conventional TCBs. For instance, audit capabilities should be strengthened to provide real-time detection and response. In effect, the audit system becomes the sensor in an anomaly and misuse control system. Appropriate effectors must still be developed. Similarly, stronger I&A mechanisms can be used to replace conventional password mechanisms.

Just as the security mechanisms are layered, the security policies must also be layered. The policies for the virtual enterprises are application layer policies that focus on the external behavior of the system as observed by agents on the network that use, or try to use, the applications available on the network. Conventional operating system access control and information flow control polices appear as lower layer policies that support the application layer policies. Only a small amount of work has been done to describe a theory of composition for policies at different layers to provide a unified policy for the entire system [6 & 7].

3.2 The Enterprise Security Architecture

The Architecture consists of a collection of data enclaves, each sitting behind its own Application Gateway system, and each trusting other data enclaves to greater or lesser extents, depending on what virtual enterprises the various enclaves belong to. This is illustrated in Figures 3 and 4.

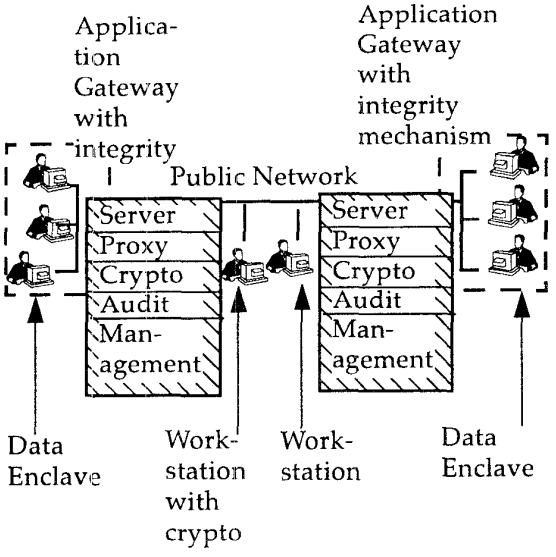

Figure 3. Crypto extends security perimeter across Public Net. I & A, ACLs and Integrity Mechanism prevent users from accessing data except in modes authorized for user and data. Proxies, Servers restrict use of authorized access modes. Audit deters would be attackers.

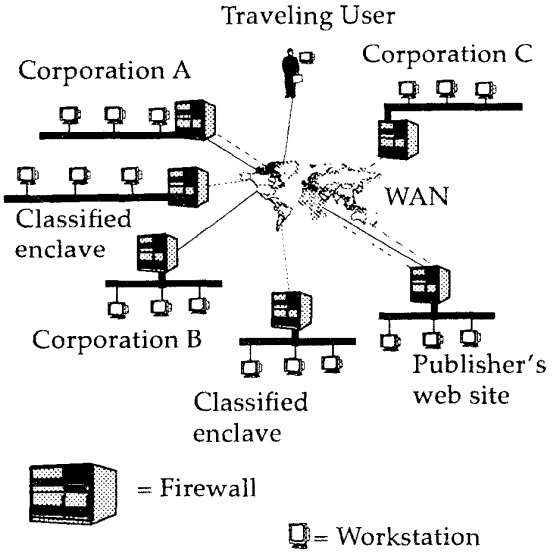

Figure 4. Crypto extends security perimeter across Public Net. I & A, ACLs and Integrity Mechanism prevent users from accessing data except in modes authorized for user and data. Proxies, Servers restrict use of authorized access modes. Audit deters would be attackers.

Each virtual enterprise is protected by its own virtual security perimeter, and each data enclave can be part of several different virtual enterprises. The function of the Application Gateway is to establish and maintain the required separation among these virtual enterprises. It does this by embodying the four principles of section 2.

The Gateway hosts a software suite that consists of servers and proxies, a cryptographic application programming interface (CAPI) to a high speed cryptographic coprocessor, and audit and management software. The underlying strong access control mechanism separates these software components from each other and protects them from attacks mounted, primarily from the outside network. It is generally the case that enterprises are willing to trust their own end-users rather than enforce stronger internal controls that would inconvenience and irritate the users. Notice that this architecture only identifies one enclave to another. This is consistent with the observation that enterprises are willing to trust their users within their own enclaves. Thus, once data is in the enclave, it is hypothetically available to all the personnel within the enclave. If individual I&A is required for certain high integrity transactions, those users could be given their own cryptographic I&A mechanism, but there would have to be strong procedural controls in place. For example, there might be designated workstations within the enclave, for which remote logins are disabled, and only those workstations would be configured to support the remote I&A required to perform the critical transactions. This could be done without requiring trusted workstations within the enclave.

Based on the cryptographically supported I&A, a foreign enclave is identified as belonging to one or more virtual enterprises to which the local enclave belongs. Based on which virtual enterprises, the agent from the foreign enclave is allowed certain access privileges in the local enclave. Generally these privileges take the form of access to certain services on the Gateway. For instance, the broadest community of interest is the general Internet community. Users from this community might not be required to provide any I&A information and might be restricted to WWW and FTP access to a limited set of files, with the services provided by servers on the Gateway. Even this quality of service might be too much in some instances. The Gateway might perform lower layer filtering and reject all connections from a list of rogue IP addresses, although this is of limited value.

When the local Gateway establishes that the foreign Gateway is part of a more restrictive community of interest to which the local enclave belongs, the local Gateway would provide a higher quality of service, possibly allowing access to one or more proxied services. In providing their services, the proxies would apply a set of filters to both incoming and outgoing data. These filters would be specific to the identified community of interest and the nature of mutual trust implied by membership in the community. For instance, regulators might be given access to browse a database inside the enclave using a limited set of client application available on the Gateway. They might also be allowed to FTP filtered version of certain files from inside the enclave to their homesites. However, these actions would be audited to provide a level of deterrence to abuse by the regulators, and the transactions would be encrypted between Gateways to preserve the confidentiality of corporate data. The ability of the regulators to modify data within the enclave would be restricted to certain database fields or to comments on files. Most likely, all modifications would have to be signed by the regulator.

Similarly, data returned to the local enclave by request of an internal user might be subjected to a set of filters dependent upon the source enclave. Data from some WWW sites might be signed for integrity with a signature that is known and trusted by the local Gateway. Such data would be passed through without further scrutiny. Data from other sites might be

restricted to certain formats and then be filtered to restrict the transmission of data driven attacks into or through the enclave. Since filtering technology is not, and never will be, perfect, this approach represents a trade between the expected cost of damage done by data driven attacks, imported into the enclave and the value to the enterprise of allowing the transaction.

4. Conclusions

As enterprises move from the closed computing environments of the past to the open, highly connected environments that are evolving today, they encounter a number of problems. Agents, users and programs, within physically separate enclaves in the enterprises naturally form communities of interest, or virtual enterprises, that span multiple physical enclaves distributed across several enterprises. Each such virtual enterprise has its own security policy that prescribes the sort of protection that the community requires for its assets. In some cases there is partial overlap among communities of interest. For example regulators have some interests in common with the enterprises they regulate, but there are other areas in which they have distinctly different interests. Enterprises must enforce the security policies of the communities to which they belong without overly constraining the activities of the end users. This is the Virtual Enterprise Security Paradigm. At the same time, financial and operational constraints preclude the wholesale replacement of existing hardware and software systems with new ones.

The Enterprise Security Architecture addresses these concerns. Key to this Architecture is the notion of an Application Gateway System that protects physically separate enclaves from the rest of the network but still allows agents within each enclave to interact with data and agents outside the enclave in a manner consistent with the security policies of all the virtual enterprises to which the enclave belongs. The Gateway hosts a variety of servers and proxies, which apply data filters and strong cryptography to enforce the security policies of the various virtual enterprises in which the enclave participates. Behind these Gateways, end-users can continue to use their existing workstations and software systems with no modifications, safe in the knowledge that their data and activities are protected from attack by the outside world.

The Virtual Enterprise Security Paradigm and the Enterprise Security Architecture have several consequences worth noting. First, Application Gateways, Guards, and Firewalls are more than temporary band-aids. They possess desirable features that make them an integral part of the Architecture. Security policy enforcement and related mechanisms can be concentrated on a small number of machines. This limits the effort involved with assuring that the security mechanisms do enforce the appropriate security policies. Rather than studying every workstation and client application suite inside the enclave, only the Gateways and their software need be examined. This increases interoperability among data enclaves and frees end-users to conduct business as usual on their workstations. It is not even necessary to modify their client software. When new applications are introduced to the enclave, it is only necessary to write the proxies, servers, and filters required for the Gateway. In many case existing filters can be reused.

A second consequence is that all encryption services, and most other cryptographic services can be concentrated on the Gateways. Again, this reduces the cost of implementing the Architecture and makes it more acceptable to the end-users. It also reduces the magnitude of the key management problem. For certain applications, it may be necessary to provide specific users, at specific workstations, a signature capability. Certainly traveling users will need cryptographic I&A capabilities on their portable systems. In effect, their portable systems become their own limited Gateways.

Finally, the conventional TCB enforcement mechanisms are not adequate for Application

Gateway operating systems. These systems need a flexible, strong, low-level access control mechanism that can be used to configure the pipelines of proxies and filters and to protect the cryptography used to enforce the security polices for the virtual enterprises.

5. References

[1] Boebert, W. E., *Building a Data Enclave*, unpublished notes, Secure Computing Corporation, 1991.

[2] Cheswick, W. R. and Steven Bellovin, *Firewalls and Internet Security*, Addison-Wesley Professional Computing Series, 1994.

[3] Hosmer, Hilary H., "The Multipolicy Paradigm", *Proceedings of the 15th National Computer Security Conference*, Baltimore, MD, October, 1992.

[4] Thomsen, Dan, "Sidewinder™: Enhanced Security for a Unix Firewall", Secure Computing Corporation, 2675 Long Lake Road, Roseville, MN 55113, March, 1995, submitted to 1995 ACSAC.

[5] Fine, Todd and Spencer E. Minear, "Assuring Distributed Trusted Mach", *Proceedings IEEE Computer Society Symposium on Research in Security and Privacy*, Oakland, CA, May, 1993.

[6] Secure Computing Corporation, Final Report, Software Crypto, CDRL A015, Secure Computing Corporation, 2675 Long Lake Road, Roseville, MN 55113, August, 1994, marked FOUO.

[7] Secure Computing Corporation, SNS Formal Security Policy Model - SNS Phase 4, CDRL AV06, Secure Computing Corporation, 2676 Long Lake Road, Roseville, MN 55113, August 6, 1995, marked FOUO.