

A NETWORK MANAGEMENT VIEWPOINT ON SECURITY IN E-SERVICES

Raouf Boutaba
Brent Ishibashi
Basem Shihada

*School of Computer Science
University of Waterloo**

{rboutaba,bkishiba,bshihada}@bbcr.uwaterloo.ca

Abstract With the advent and the rapid growth of the Internet, e-services have proliferated. Indeed, e-commerce activities have played a vital role in expanding current business transactions to much higher levels by allowing a larger number of potential customers and companies to interact in a shorter time with lower costs. E-services include business information, processes, resources, and applications, which are supported through the Internet.

As the popularity of e-services have grown, so has the need for effective security. All aspects of the e-service must be secured, using a variety of security mechanisms, objects, and functions. In order to maintain a secure system as a whole, security components must be managed. Therefore, the implementation of secure e-services cannot take place without full support from network management.

Network management monitors and controls the network in order to ensure that it is providing its services efficiently. It also shapes the network's evolution through integrating new technology and supporting new services. There are five widely accepted network management functional areas: fault, configuration, accounting, performance, and security management. Security management involves several services including access control; authentication; confidentiality; integrity; non-repudiation; availability; and accountability.

This paper will highlight essential and common network management architectures and protocols in constructing a complete view of how network management enables security for e-services.

Keywords: network management, security management, security services, e-services

*This work has been funded by the Natural Sciences and Engineering Research Council of Canada.

1. INTRODUCTION

A chain is only as strong as its weakest link. So too is the security of an e-service. Added to this, providing security for e-services is of great importance, even more so than for many other distributed systems. Therefore, great care must be taken to create a secure infrastructure on which truly secure e-services can be offered. At the base of this infrastructure is the network itself, with network management working to provide the security mechanisms and services on which a truly secure e-service can be built.

The importance of proper security infrastructure becomes even more important when we consider the type of networks these systems are intended for. If we were dealing with a closed private corporate LAN, security issues might be considerably simpler. Providing a service to all users of the network (or even some subset of the users) would be fairly straightforward, as corporate policy and simple controls could restrict who could go where and who could do what. Additionally, traffic starts, finishes, and always remains within controlled network. Maintaining a secure system in this scenario, while still not trivial, is hardly insurmountable.

However, we of course wish to provide e-services on large, open networks, especially the Internet. Here, we have considerably less control over the network as a whole; in fact, we have practically none. Here the network is made up of a large number of separate domains, under the control of many different organizations or individuals. Some of these parties may be trusted, but most of course are not. As each of these parties may deal with their section of the network as they choose, traffic flowing through these sections will be subject to the security of each section. Hence, there are potentially a very large number of weak links within the system.

Clearly, we cannot hope to secure the entire system. Instead we must provide mechanisms and services to allow us to secure our own networks and systems, as well as securely deliver our e-services to the customer, regardless of whether that customer is internal or external to our own network and organization. Network management (NM) aims to provide this infrastructure.

Network management by itself will not create a secure environment for e-services to be offered within. Only an entirely trusted, closed system could offer such an environment. Clearly, such an environment is impossible with the networks we wish to offer our services upon. Instead, network management can only aim to provide services in order for systems to adequately protect themselves.

2. ORGANIZATION OF THIS PAPER

Throughout this paper, we aim to describe the role network management plays in securing e-services. This will involve building up a picture of e-services, and how network management, and specifically security management, fit into that picture. Included throughout will be examples and descriptions of current technologies that are being used and developed.

In the next section we will state our goals in providing secure e-services, including our description of an e-service. These are the motivations behind the rest of the paper. In section 4, we will proceed with a description of security, including security policies, objectives, and risks. Section 5 will focus on the security services needed to fulfil our security objectives. Within this section, we also look at the use of cryptosystems as an important part of providing security. As an example of a security service, we briefly describe the authentication service provided by Kerberos.

Section 6 turns to security management, describing how these services can be provided and supported. Several technologies are discussed, including public key infrastructures and policy-based networking. Section 7 continues on to tie security management into the field of network management as a whole. Within this section we also focus on the necessity of securing the management infrastructure itself. Section 8 provides a look at assuring security within a service, as well as creating customer trust. Finally, section 9 provides a conclusion, tying the pieces together and connecting management with the e-services.

3. SECURE E-SERVICES

As a first step in our discussion, what exactly are we trying to accomplish? What exactly are we aiming for in delivering secure e-services? Even before that, what do we mean by an e-service?

3.1. E-SERVICES

Remembering that we are looking at this from a network management point-of-view, we will define an e-service as the provision of some electronic asset by a provider to a customer. Some e-services that have been suggested include e-Commerce, e-Government, e-Health, and e-Education. To date, e-commerce has been the quickest to catch on, with many businesses offering online ordering systems, banks offering online access to banking services, as well as new online electronic payment systems. In fact, any website is an e-service, delivering information to




Figure 1 A provider offers e-services to a customer using the Internet

a customer. Usually, this is provided openly, and therefore there are few security requirements. However, many websites sell access to their material, increasing the need for security.

In all these services, we have a provider who has an asset, whether it is some piece of information or some capability to perform some service. The provider wishes to offer this asset to a customer, or number of customers. The delivery of this service will be electronically, via some data network connecting the providers system or systems and the customers system(s). We will consider our general model of the data network to be the Internet. This idea of an e-service is depicted in figure 1.

3.2. MODELLING THE E-SERVICE

Our e-service model is made up of three basic parts: the service providers system(s), the customers systems(s), and the connecting network. This is a distributed system that can be viewed as a client-server architecture. In fact, provider and customer systems may be (and likely are) made up of a number of connected machines, however from a network management viewpoint, the interactions between these individual machines can be handled in the same way (or possibly in a simpler way) as the client-server interactions.

Additional security mechanisms may be needed for group communications if the server side is composed of a group of replicated servers, for fault tolerance or performance reasons. Similarly, certain e-service transactions may require the inclusion of a third party (or fourth, etc.), but these too would be similarly modelled.




Figure 2 The client-server model made up of application, system, and network layers. Data is exchanged between the application layers using the system and network layers.

Within the client-server model, we have a client application on the customers system contacting the server application on the providers system. Each application runs on top of their respective systems hardware and operating system. In turn, the systems reside on the network. Data is kept and used by the applications using mechanisms provided by the system. Data messages are exchanged between the client and server applications, by using the facilities provided by their systems. The systems in turn pass the messages to the other system via the network. At the other end, the message is passed back up through the layers to the destination application. This layered model is depicted in figure 2. It is possible to further refine this model into further sub-layers, such as the layering of protocols for providing system access to the network (OSI or TCP/IP protocol architectures), however this level of abstraction is sufficient for our purposes.

3.3. SECURING THE E-SERVICE

Securing an e-service (like other distributed systems) focuses primarily on two tasks: protecting communications via a secure channel between any communicating parties[1], and controlling access to the systems and resources involved in the service[2]. The secure channel protects the confidentiality, integrity, and authenticity of the messages it carries. Access control verifies that parties are authorized to have access to a resource, and bars unauthorized users from the system. We will elaborate further in the next section on security.

Due to the layered nature of providing an e-service, security must be considered at all levels of the model. Obviously the e-service applications themselves must be concerned with security. However, because the applications run on top of the operating systems, the systems themselves must provide secure services to the application. A security breach within the system could jeopardize the security of the application. Similarly, security issues within the network can put the systems and the applications at risk.

We should clarify that this does not mean that the entire network must be secure. It has already stated that this is impossible in an environment such as the Internet. However, the network must provide sufficient security and security services to allow the systems and applications to secure themselves to the extent required.

This raises a question - what extent is required? This will be dependent on the service being provided, and the potential risk involved, however we have two basic goals. From a network management point-of-view, our goal is to protect, or allow the service provider to protect, all the providers assets from inappropriate use, theft, or damage. These dangers may come from sources external to the e-service (elsewhere in the network), or from either within the providers or the customers systems. As a secondary goal, we should protect and allow the provider and customer to protect the customers assets as well.

Why should network management concern itself with the provider first, and the customer second? First, it is likely that the providers assets and systems will be far more extensive and susceptible to attack. Second, the providers systems will likely be located within a domain where their organization has at least partial control over the management of the network. Third, it is in the providers best interest to offer its services securely for both itself and the customer, as security risks to the customer will decrease the service providers reputation and utility.

4. SECURITY

Although we have broadly defined our goals, a better understanding of security issues is essential. In our goals we stated that we wished to protect all assets from inappropriate use, theft, or damage, from sources internal, and external to the e-service. We will now elaborate on this.

4.1. SECURITY OBJECTIVES, RISKS AND ATTACKS

Our overall goal of security can be broken down into a list of security objectives. Different objectives relate to different aspects of the overall

system. Some of these objectives deal with data and messages. Confidentiality, data integrity, and availability can be considered primary goals pertaining to data. Authentication and non-repudiation focus on security of communication between entities. Finally, access control, audit trails, and security alarms relate to the security architecture itself. All of these objectives are summarized in the first part of table 1[3].

By achieving these objectives, we eliminate security risks. These risks can be similarly categorized, as shown in the second part of table 1 [4]. There is of course a strong similarity between the two lists. In many cases, such as disclosure and confidentiality, the risk (disclosure) directly contravenes the objective (confidentiality). However, some risks may contravene multiple objectives, or conversely, one objective may serve a role in eliminating several risks. For example, access control obviously work to prevent unauthorized access, however authentication may also play a role.

Real world security attacks can be described in terms of these risks and objectives. For example, a data tampering attack would be a modification risk, contravening the data integrity objective. Sniffing or snooping of data packets would be a disclosure risk, and contravene confidentiality. Again, there is not a one-to-one correspondence, as a specific attack may be a combination of risks, and/or contravene any number of objectives.

4.2. SECURITY POLICY

Some security is provided at each of the layers of our model. For example, the operating system on both the provider and customer systems provides some mechanisms for restricting access to their respective file systems. However, if security was provided solely on a layer-by-layer basis, it would be extremely difficult, if not impossible, to achieve a completely secure system. Certain security issues may be best addressed at a particular layer, however many require mechanisms that coordinate across the layers in order to adequately protect the overall system. In addition, building all kinds of security mechanisms into an e-service does not make sense unless it is known how those mechanisms are to be used and against what. This requires that we know about the security policy that is to be enforced.

In order to meet the security objectives for the entire e-service, an organizational security policy should be created. A security policy defines the security issues that an organization faces, and identifies strategies that can be used in order to achieve the organizations goals (the goals we have just defined). Creating a security policy can be broken down into

<i>Security Objective</i>	<i>Description</i>
Confidentiality	Protect information held or communicated within the e-service from unauthorized access or eavesdropping.
Data integrity	Prevent information (held or communicated) from being changed or lost.
Availability	Ensure that the service is available at all times that it is needed.
Authentication	Ensure the identity of communication partners, and ensure the authenticate the origin and integrity of messages.
Non-repudiation	Provide proof of origin and proof of delivery of messages.
Access control	Limit who or what is allowed access to services and resources based on authorizations.
Audit trail	Provide evidence of who did what, and when.
Security alarm	Minimize risks by detecting actual or potential security failures.
<i>Security Risk</i>	<i>Description</i>
Disclosure	Release of information from within or about the service, to an unauthorized party.
Unauthorized access	Improper access to services and resources, by a party who does not have those privileges.
Modification	Changing of information, whether within a system, or as the modification of a communication message.
Misuse	Use of a resource or service for reasons other than their intended use.
Abuse	Legitimate users may make abusive use of the service or other related resources. ie. Use of more than their share of resources.
Fraud	Misrepresentation of identity or intention in using the provided service.
Repudiation	Denial by a legitimate user that they have made use of a service or resource, including denial the of sending or receiving a message.
Denial of Service	Inability for legitimate users to properly use the provided service.

Table 1 Description of Security Objectives and Risks

<i>Layer</i>	<i>Possible Assets</i>
E-service Application	Customer records Financial transactions The e-Service itself
System	System access Files Physical hardware
Network	Device settings Bandwidth Physical hardware

Table 2 Description of Security Objectives and Risks

three steps: identification of assets, threat analysis, and threat elimination.

4.2.1 Identification of Assets. First, all assets must be identified. This includes everything of value within the e-system, which must be protected. Proceeding layer-by-layer, we can identify assets to be protected within the applications, systems, and network layers. Table 4.2.1 lists a few examples of assets that must be protected. In addition to identifying each asset, a value for that asset should also be assigned.

4.2.2 Threat Analysis. Once the assets have been determined, all potential threats to those assets should be identified. To aid in finding solutions to all possible threats, policy-makers should investigate not only how each asset might be attacked, but also where the threat might come from, what intent the attacker has, and, in the case of mobile assets such as transmitted data, where the asset might be attacked. These points are summarized in table 4.2.2.

We mentioned earlier, in our description of e-services, that an attack may originate internally or externally to the e-system. If we approach this from the providers point-of-view, the most common concern is of course an external attack, where an outside attacker (elsewhere on the network) poses some form of threat to the provider, customer, or intermediate network. It is also possible for a security threat to originate from either the customer (or someone posing as a legitimate customer), or even from within the service provider, including threats such as an equipment failure.

<i>Questions</i>	<i>Description</i>	<i>Possible Answers</i>
Source	Identify who might attack the asset.	Outside attacker Malicious customer Within organization
Location (source)	Identify where the threat comes from.	Internal threat External threat
Location (asset)	Where is the asset vulnerable?	Provider systems On Network Customer systems
Intent	Is the attack caused intentionally?	Malicious Accidental
Type of attack	How does the attack affect the system?	Active Passive
Risks	What risks are the asset in danger to?	All security risks
Objectives	Which objectives need to be protected?	All security objectives

Table 3 Identifying and Classifying Security Threats to an Asset

The intent of the threat source may also need to be considered. A customer may accidentally or carelessly pose a security risk. However, someone posing as a customer could maliciously pose a similar or greater risk. The security solutions required to close security holes may differ depending on the intent of potential threat sources. Also, the response or punishment towards a malicious attack will likely be far greater for a malicious transgression than an accidental one.

Different assets may be susceptible to different types of attacks. Attacks may be either active or passive. In an active attack, data or communications may be altered or tampered with, whereas in a passive attack, attackers only observe. Although active attacks may be more difficult to perform, they can also do much more damage. In an open network like the Internet, preventing passive attacks may be impossible, however methods can be used to limit the amount of information such an attack can yield.

Finally, we must determine what types of security risks each asset is vulnerable to, and which security objectives these threats endanger. By

identifying these risks and objectives, a better understanding of each potential threat will be gained. The better the entire situation is understood, the more easily solutions for eliminating or reducing those threats can be found.

4.2.3 Elimination of Threats. Once all the threats to each of the assets have been identified, the security policy can then turn to finding solutions for eliminating those threats. First, it should be acknowledged that not all threats can be eliminated. For example, no protection has been found that is effective in preventing a denial-of-service attack. In such a case, a certain amount of risk must be accepted. However, some security measures may minimize the risk involved, rather than accepting the full risk.

Additionally, there may be times when protecting an asset is not practical. With a value associated to each asset, there may be times where the cost to provide additional security for an asset outweighs the asset's value to the organization. In such a case, there is no reason to provide this additional protection. It would cost more to secure the asset than it would cost the organization if the attack occurred successfully.

There may be a need for some of the solutions specified within the security policy to not in fact be focused on electronic aspects of the organization and service. Some may be physical mechanisms, such as the protection of hardware (locks and doors in the real world). Others might be organizational policies, such as a policy on prosecution or punishment. However, for issues directly relating to the electronic aspects of the e-service, security services should be utilized to minimize or eliminate all preventable risks.

5. SECURITY SERVICES

Creating a security policy is the first step to providing secure e-services. In creating the policy, decisions must be made on what assets need protection, and what security services will be used in order to protect them. Next, we must take a look at what types of services may be provided.

Consider the securing of a simple e-commerce service. The provider has a web server handling incoming transaction requests. The customer is using a web browser. A number of services must be used in order to make the transaction a secure one. In order to ensure that each party is who they say they are, some authentication and certification services will be required. To prevent the transaction data, or other data such as passwords from being overheard, a confidentiality service will be used. A key management service will provide and facilitate the exchange of

<i>Service</i>	<i>Description</i>
Authentication	Ensure that parties involved in the service are who they say they are.
Key management	Securely provide authentic cryptographic keys.
Integrity	Prevent modification of data by unauthorized users.
Availability	Protect against malicious service disruption, misuse and abuse.
Accountability	Hold users accountable for their actions, including billing users for service usage.

Table 4 A list of possible Security Services

cryptographic keys for the confidentiality service. Using these services, a secure channel can be created between the customer and provider systems, and the transaction can be completed.

The services provided to a network and e-service will depend on the management policies for the system. They might include services for authentication, key management, authorization, access control, confidentiality, or others. A brief description of these services is described in table 4.

5.1. CONFIDENTIALITY SERVICES

Confidentiality services are essential to most other security services within the network. Without the ability to maintain secrecy regarding the content of messages transmitted on the network, many other security mechanisms are rendered useless. For example, if an authentication system transmits a password from client to server without any protection, the interception of that password easily defeats the security of the system.

All data within the e-service may need protection. This includes messages being passed between communicating parties, data stored within the system, or even data about the service itself. Protecting the content of messages being sent between the customer and provider from prying eyes is obviously needed, however it may also be important (and more difficult) to hide the fact that a communication even occurred. We will first briefly introduce the use of cryptosystems to protect the content of messages, followed by a description of how the IETFs IP Security Protocol (IPSec) protects against traffic analysis.

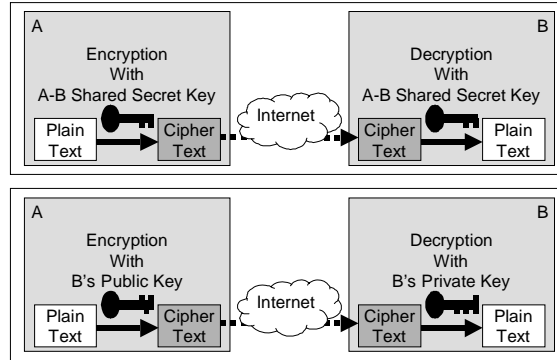


Figure 3 The use of secret (top) and public (bottom) key cryptosystems.

5.1.1 Cryptosystems. The primary mechanism used to protect messages is, of course, data encryption. Cryptographic systems are used to encrypt the data prior to placing it on the insecure network. At the receiver the data is decrypted, and assuming only the encryption and decryption keys are secure, only the sender and receiver have access to the usable data.

Two types of cryptosystems are in use. The first is a secret key or symmetric cryptosystem. Here, the same key is used for both encryption and decryption. As the same key is used to both encrypt and decrypt the message, it is essential that the key remains private - only the proper communicating parties should have knowledge of it. This also means that a key should only be used for a short length of time, in order to avoid data analysis attacks and minimize compromised data in the case of a stolen key. A large number of secret key algorithms exist, such as Data Encryption Standard (DES)[5] or Blowfish[6].

The second type is public key or asymmetric cryptosystems. In public key cryptography (PKC), two different keys are used, one each for encryption and decryption. The keys are the inverse of each other, so one key can be used to encrypt a message, and the other used to decrypt it. The public key is mathematically derived from the private key, however it is extremely difficult (computationally hard) to find the private key from the public key. Therefore the owner of the key must keep the private key secret, however everyone can know about the public key. RSA (Rivest-Shamir-Adleman)[7] is the most widely deployed public key cryptosystem. Secret and public key cryptosystems are illustrated in figure 3.

Unfortunately, PKC is fairly slow compared to secret key systems, as the keys used must be much longer, and the algorithm more complex. Therefore, secret key cryptography is much more efficient for protecting large amounts of data. A combined system is often used - secret key encryption is used to protect the data, using session keys (the session key is a short-lived key that is changed regularly,) while public key cryptography is used to facilitate the establishment of the session keys. Public key cryptography is also important in digital signatures and certification, which will be discussed later.

5.1.2 IPSec. The Internet Protocol Security Protocol (IPSec) [8], is an IETF standard for providing confidentiality, integrity, and authenticity to packets carried over an IP-based network. IPSec provides network layer data encryption, with data placed into several new packet formats. One of the features that IPSec allows is the ability for packets to tunnel. In tunnelling, a packet is encrypted (including all headers) and encapsulated in a new packet. This packet is then sent to its destination (not necessarily the same as the original destination). This node removes the header, decrypts the original packet and forwards it on to the destination.

Using IPSec, a provider organization can use a device (such as a router) to act as an IPSec proxy. The proxy encrypts and encapsulates the packet, then forwards it to an IPSec-enabled router at the other end. That device then forwards the packet to the proper destination. Although this does not fully hide the fact that traffic has travelled from through the network, the actual end system addresses have been hidden, and the traffic cannot be identified from other traffic flowing between those devices.

5.2. KERBEROS NETWORK AUTHENTICATION SERVICE

As an example of providing a security service, we will present the widely used Kerberos network authentication service. The Kerberos system was originally developed at MIT in the late 1970s. A wide number of free and commercial distributions of the latest version, v5, are available.

Kerberos[9] helps clients create a secure channel, by providing authentication and certain key management services. We will focus on the authentication service, which is based on secret keys. Consider a user trying to log in to a server from their workstation. The user will enter their login identity, which will be sent to the Kerberos authentication




Figure 4 Authenticating user A to contact server B in the Kerberos Network Authentication System. Messages in white are transmitted in clear text, while messages in light gray are encrypted. Note that tickets are themselves already encrypted with a key shared between the creator of the ticket and the destination.

server (AS). The AS returns a session key that the user can use to contact the ticket granting service (TGS), along with a ticket to be given to the TGS. The ticket is encrypted in a secret key shared between the AS and TGS, while the entire message is encrypted in a secret key between the AS and the user.

The user must decrypt this message in order to access the secret key and ticket that are needed in order to contact the TGS. The key shared between the user and authentication service is based on the users password - with the proper password the user can properly decrypt the message. The user then provides the ticket, as well as the identity of the server they wish to access and a timestamp, to the TGS. This ticket proves to the ticket granting service that the user has been properly authenticated. The TGS then returns a secret key to be used between the user and the server (encrypted of course), as well as a ticket to be provided to the server containing the users identity and a copy of the key. This ticket is encrypted in a secret key shared between the TGS and the server. This exchange is illustrated in figure 4.

Several other common authentication services exist. Diameter[10] is the current IETF authentication protocol that evolved from RADIUS [11] (Remote Authentication Dial-In User Service), although RADIUS is also still in use. These protocols are the basis for IETFs AAA (Authentication, Authorization, Accounting) working group. TACACS (Terminal Access Controller Access Control System)[12] is older but less popular than RADIUS, due in large part to the fact the latest version, TACACS+, is proprietary to Cisco Systems.

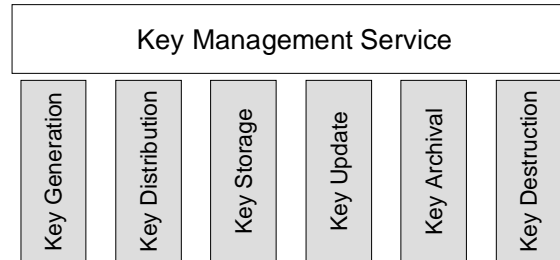


Figure 5 Services are composed out of a number of individual functions. Here, the functions of a key management service are shown.

5.3. KEY MANAGEMENT SERVICES

Clearly, cryptographic keys play an important role in providing security. It is critical that any key used is handled in a secure fashion, throughout the keys lifetime - in fact, even after the key is no longer needed. Management of those keys is therefore crucial to the security of the cryptosystems. We have already stated that a public key system can be used to exchange session keys. The session keys can be randomly generated and exchanged using this system (and destroyed after use). However, the public-private keys used in PKC must be managed, in order to ensure that they are properly handled at all times.

To complete our look at security services, we should show that services can be broken down into many individual functions. Each service in fact performs a number of different tasks, each provided by some sub-function of the service. Some component functions for a key management service are shown in figure 5. They include functions for creating, maintaining, destroying, and storing keys. Providing these functions is the responsibility of security management.

6. SECURITY MANAGEMENT

Lets return to the idea of layering. The e-service application performs some security tasks to deal with certain security issues. It might perform actions such as end-to-end encryption, to protect messages to be passed between systems. It can also deal directly with ensuring that all transactions are performed correctly.

The application must rely on the system to protect it. The operating system protects the application, as well as other processes within its execution space from unauthorized interference. It also protects the




Figure 6 Security management provides security functions to all of the layers: the e-service application, the systems, and the network.

application from other processes. The files system adds mechanisms to protect data stored by the application. Access to all resources is carefully controlled.

The system in turn relies on the network to provide communication services. The network itself, and its resources must be protected from all types of security threats. In an open network such as the Internet, even maintaining service availability can be problematic.

The application, system and network form a layered structure. E-service security is built on top of system security, which is in turn built on top of network security. Again, the question arises - how do we tie it all together in order to build a complete, secure e-service?

Security management provides functions to be utilized by all three layers, as shown in figure 6. By providing the necessary infrastructure, security management allows the creation and maintenance of a sufficient set of security services in order to effectively enforce our security policies. This infrastructure consists of the mechanisms and data structures required to create, operate, and maintain the security services. On their own, the management functions will not provide security, however they provide the basis for implementing the required security services. A number of important security management functions are summarized in table 5. We will elaborate on several of these functions.

6.1. MANAGEMENT OF SECURITY SERVICES

In order to provide the requisite service set to the system, we need to be able to manage the security services. Security management must provide mechanisms to support each implemented service throughout the services lifecycle. A security plan must be in place to follow an individual service from the development stage, through implementation,

<i>Function</i>	<i>Description</i>
Services and Mechanisms	Functions to manage the offered set of security services and mechanisms throughout each service's lifetime.
Keys	Create, deliver, store, and destroy cryptographic keys.
User registration and information	Provide infrastructure for creating, changing, and storing user-related information.
Access control information	Manage mechanisms for restricting access to resources by unauthorized parties.
Security audit trails	Maintain reporting and recording facilities to be able to retrace security-critical actions within the system.
Distribution of Security information	Allow security-related information to be delivered safely throughout the system, wherever it is needed.
Security Event Reporting	Provide mechanisms for reporting current security statuses or triggering alarms to the required management stations.

Table 5 Functions Security management must provide and manage.

maintenance, and finally decommission. The security and effectiveness of a service could be compromised if a security lapse is allowed to occur at any step within a services lifespan. For example, even after the removal of a service from a system, if encryption keys from that system are leaked or can somehow be reconstructed, this could compromise the confidentiality of previously encrypted messages and data.

Maintenance of the security services applies to both the individual services and to the set as a whole. During the update or replacement of a particular service, it is critical that the service be changed in such a way as to minimize the security implications. A brief interruption of a particular service may create a security hole on its own, or may impact on other services, creating a hole. A mechanism might be provided for performing updates with the system online, resulting in a smooth and secure switchover, however alternatively (as is often the case,) the solution in many systems may simply be to take the system offline, perform the change, and then restart the system. While this compromises the availability of the service for a short time, it may be viewed as the most secure, and least difficult solution.

6.2. MANAGEMENT OF KEYS

Earlier we showed the importance of keys and how a key management service was made up of a number of component functions. We will now reprise this topic by looking at the need for these components, and by introducing a public key infrastructure (PKI). Within this section, we will also discuss the use of certificates and digital signatures.

With public key cryptosystems, the owner of a key should control the private key. The public key should be available to anyone who wishes to communicate with the owner. This requires that there must be some method for creating the public-private key pair, a method for allowing the owner to have the private key, as well as a method for making the public key available. One option is to simply allow the owner to create the key pair and give the public key to anyone who wants it. The problem with this is in identifying the owner. How can someone receiving the public key be assured that a) the owner is who he says he is, and b) the key really is the owners?

If there is no way to tie a key to an owner, then anybody can provide a public key claiming they are someone else (or that the key is someone elses). In order to provide a guarantee that keys are properly paired with their owners, certificates are used. Before we can understand how a certificate works, we first must discuss digital signatures.




Figure 7 Creating and checking a message signed by a digital signature.

An interesting property arises out of the fact that the public and private keys are inverses of each other. While normally the public key is used for encryption and the private key for decryption, the reverse also works. If a message is encrypted using the private key, anyone holding the public key will be able to decrypt the message. If the identity of the public key is assured, then successful decryption of the message with the public key assures that the owner of the key sent the message. For a digital signature, often a hash of the message (using an algorithm such as MD5[13] or SHA[14],) is encrypted and sent, rather than the encrypting the entire message. If the hash can be successfully decrypted and matched to a recomputed hash of the received message, the message is assured of being unaltered and sent by the owner of the key. This process is shown in figure 7.

Certification relies on digital signatures. A certificate is a public key matched with an identity. In order to ensure that it cannot be forged, the certificate is signed by a certification authority (CA). The CA must be trusted by anyone using the system, otherwise the certificate cannot be trusted. All users in system must also have the CA's public key. They can therefore verify that the key and identity are correctly matched. This certificate can now be freely published, as it is protected by a digital signature.

A public key infrastructure is a set of hardware, software, people, policies and procedures needed to manage public-private key pairs. A list of PKI functions is presented in table 7. This list is of course very similar to the list presented while discussing the composition of services (figure 5).

PKIX is the IETF standard for PKIs[15]. It is based on the X.509 version 3 certificates[16], the most widely used certificate standard. The PKIX architecture is made up of certification authorities (CAs), organisational registration authorities (ORAs), and repositories. The CA is responsible for issuing the PK certificates, only after the ORA has ap-

Registration	Key expiry	Cross certification
Initialisation	Key compromise	Revocation
Certification	Key generation	Certificate Distribution
Key-pair recovery	Key update	Revocation Notice Publication

Table 6 Functions provided by a public key infrastructure

proved and vouched for the certificate. After the certificate is issued, it is stored and made available by the repository. Additionally, the repository handles revocation lists, controlling certificates that have been revoked.

6.3. ACCESS CONTROL INFORMATION

In order to control access within the e-service, methods must be in place to record which users are allowed access to which assets, and to what degree. There are several ways this information can be provided, and security management must provide the mechanisms to do so.

Access control matrices store information regarding the actions a subject (represented by a matrix row) is allowed to perform on a particular object (represented by a column). For scalability reasons (as the matrix might be very large and sparse), an access control list (ACL) might be used instead. An ACL is kept by each object, each list containing only the subjects who have some permission on that object.

Another approach involves giving the list to the user (subject) rather than the object. In this case, the user gets a list of capabilities. When the subject wishes to make use of an object, it presents a capability along with its request. The resource checks the capability and, if the appropriate action is permitted, performs the request. These capabilities are usually stored in a privilege attribute certificate, which is timestamped, authenticated, and possibly encrypted to ensure security.

6.4. DISTRIBUTING SECURITY INFORMATION

A number of mechanisms are used for distributing security information to the required points within the system. We have already discussed mechanisms relating to key management and how keys and certificates can be safely distributed. Other types of security information must also be distributed. Another example is the delivery of security settings to devices throughout the network. Security information is just one type

of management-related information that needs to be distributed. Often it can be distributed to devices in the same way as all other settings.

One mechanism popularly used in network management is the concept of policy-based networking (PBN). Network policies are sets of rules that are used to administer or manage resources within the network. Policies provide a way for network administrators to deal with situations across multiple devices within the network in a consistent manner. The IETF Distributed Management Task Force (DMTF) provides a policy framework to deliver and enforce these policies[17].

The policy framework consists four elements. A policy management tool is used to define policies to be used. These policies are stored within a policy repository. The policies themselves are delivered and enforced at policy enforcement points (PEPs). The PEPs take actions upon devices within the network to manage their behaviour. The actual policies are retrieved from the repository by the policy decision points (PDPs). The PDP acts as a policy server, by interpreting the policies and then choosing the appropriate policies to be delivered to each PEP for enforcement.

The Common Open Policy Server protocol can be used for policy provisioning (COPS-PR)[18]. COPS-PR has been created for use as a management protocol within policy-based networks. It allows all types of policies to be delivered and maintained, including security policies. Other management protocols (especially SNMP) have also been used in PBNs, however COPS is specifically oriented for use with policies. We will return to SNMP however, in section 7.1.

6.5. AUDITING

In some cases, despite the best efforts, security problems may occur. Audit trails create a record of activities, by logging certain events as they occur. By receiving reports on user registration, login attempts (especially failed attempts), access to critical resources, and other suspicious, unusual, or particularly dangerous activity, security management points can log these events. In the case of a security breach, the actions of the intruder can then be retraced and the appropriate response can be made. Auditing is closely related to both security management and accounting management, another area of network management.

7. SECURITY OF NETWORK MANAGEMENT

Security Management is one of the five OSI management functions. Known as FCAPS, the five areas are fault, configuration, accounting,

performance, and security management. Together, the five areas make up network security. The five areas are summarized in table 7. It is important to note that the five areas are not discrete - there is, of course, a strong interdependence between the different areas. For example, consider the crippling effect faults or poor configuration can have on the performance of the network.

<i>Management Area</i>	<i>Description</i>
Fault	Detection, recovery, and documentation of network anomalies and faults.
Configuration	Recording and maintenance of network configurations and updates to ensure normal network operations.
Accounting	Handle user management, administration functions, and billing.
Performance	Provide reliable and quality network service, including QoS provisioning and regulation of performance parameters.
Security	Provide protection against all security threats to resources, services, and data.

Table 7 The FCAPS functional areas of network management.

The other functional areas of management support services similar to security management. In fact, like mentioned, there is no distinction as the different areas are interdependent. However, we must consider the ramifications of all network management functions on the security of the system as a whole. All services, and the delivery of those services must be handled in a secure way.

A look at configuration management easily shows this effect. Configuration management functions are provided to allow network devices and settings to be maintained or modified. Poor configuration of a network device (consider an improperly configured firewall) can potentially create a large security risk. However, an inadequately secured configuration management function that allows inappropriate access to configuration settings could be even more dangerous. Therefore, properly securing all management functions is critical. Any security flaw within the management services can be extremely dangerous.

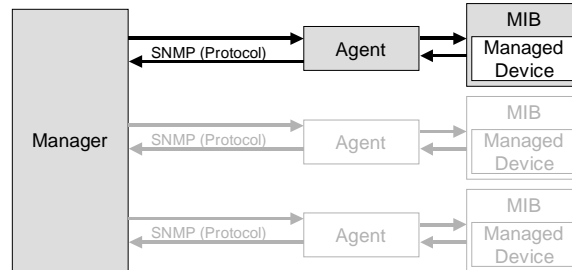


Figure 8 The SNMP manager-agent architecture for network management.

7.1. SIMPLE NETWORK MANAGEMENT PROTOCOL

To illustrate the need to secure network management, we will trace the evolution of the Simple Network Management Protocol (SNMP) from its first version to the current version 3. SNMP defines a management architecture for managing devices within a network. The basic SNMP manager-agent architecture is shown in figure 8. Additionally, SNMP also defines a protocol for the communication between managers and agents[19].

In the SNMP architecture, an agent is attached to each managed device. It may reside within the device, or where such capabilities do not exist, it may reside on a proxy device. The agent handles all management communication with the managed device, through the use of the Management Information Base, or MIB. The SNMP agent monitors and controls the device by reading and setting respectively, the settings in the MIB. The manager agent acts as a central control point within the network. The manager send messages to the agents on each managed device, asking for current information, or instructing a setting to be changed.

SNMP v1 suffers from a generally weak security concept. First, data transmission within the protocol itself is performed via clear text (no encryption). Obviously, this violates a number of security objectives. Second, it relies on connectionless UDP communication to transmit SNMP messages, meaning messages can easily be lost. Despite this, SNMP became, and remains, the dominant management architecture for data communication networks. Primarily, this is due to the fact that SNMP is designed to be simple. This gained it early adoption within the Internet.

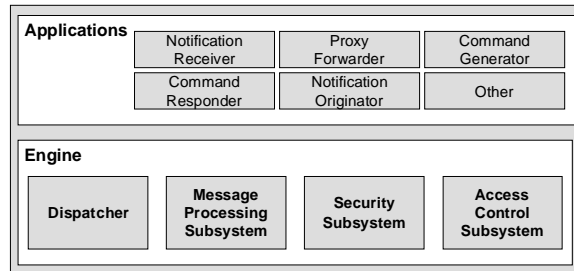


Figure 9 The SNMP v3 Entity structure.

With the security concerns in mind, development of SNMP v2 began. Several versions were designed, however version 2p contained the most significant set of security features. In v2p, messages were encrypted using DES. Also, packet sources and content were authenticated using MD5. Other features included weakly synchronized time stamping of management packets, to prevent attackers from replaying or reordering packets, as well as the addition of security levels. SNMP v2 also adopted a distributed manager architecture, by allowing a hierarchy of manager-agents to be formed.

Unfortunately from a security viewpoint, none of the SNMP v2 variants were widely accepted. The primary reason for this was the fact that SNMP v2 was not backwards compatible with v1. The messages used in the two versions differed, and v1 agents could not handle the v2 messages. This lack of interoperability prevented SNMP v2s use, as v1 was already widely implemented. SNMP v2p has now been classified as historic[20].

Version 3 addressed both the security concerns and the compatibility issue. Intended to address the security and administration deficiencies of the previous versions, SNMP v3 created a new v3 Entity, added the use of User-based Security Models (USMs), as well as adding the View-based Access Control Model (VACM). The modular design of the new entity is shown in figure 9. The dispatcher allows multiple versions of SNMP messages to be accepted within the engine, by forwarding different packets to different modules. Therefore, within an SNMP v3 entity, there may be several message processing subsystems to handle the different protocol versions. This allows the co-existence of SNMP versions within the network.

The entity also added a dedicated security subsystem, along with an access control subsystem. The access control module allows the entity

to make use of authorization services. The security module takes care of authentication and privacy concerns, as well as allowing for the use of multiple security models.

The User-based Security Model[22] offers SNMP v3 message-level security protection. First, it protects from modification of information, masquerade attacks, and message stream modification, via the use of MD5 and SHA (Secure Hash Algorithm). Second, it protects messages from disclosure by using DES for encryption, although this is considered an optional component. The chosen protocols were deemed acceptably secure, however the model allows for changes to be made if they are deemed necessary.

The VACM[23] provides an access control facility. The VACM allows agents to be configured to allow different access rights to different managers. A particular manager may be allowed full access to the agents full MIB, or a restricted view of only selected fields within the MIB. Additionally, this view may be restricted to read-only access. The access control policies must be pre-configured. Combined with the integrity and authentication protection offered by the USM, the VACM ensures that only the appropriate parties have access to the MIB, and therefore to the device.

8. SECURITY ASSURANCES AND TRUST

One more security aspect needs to be discussed. After all the services and mechanisms are in place, and the e-service can be securely offered, how can we ensure that it is in fact as secure as we think? Also, in a distributed system, that spans across multiple organizational domains, how can we agree that the required security tasks have been performed? For example, how do customers know that security measures are in place to protect their confidential information?

There are several aspects to this. First, both (or all) parties involved must be able to decide on the level of security required—what must be done to ensure the transaction is secure. Second, the provider must be able to verify that the mechanisms are working properly, and that the overall security is effective. Third, the provider must be able to convince the customer that their security claims fulfil the customer's security requirements.

8.1. SECURITY SERVICE LEVEL AGREEMENTS

One way to deal with these issues is the use of Service Level Agreements (SLAs). An SLA is a formal agreement between provider and

customer that contains both parties negotiated QoS requirements and responsibilities. Security demands are one aspect of QoS, so a security SLA could include what details of what mechanisms must be in place, and what procedures should be followed[24]. The SLA is then a legal contract, forcing both parties to fulfil the terms of that contract.

Different levels of security may be provided. In some situations it may not be critical to provide maximum security for the e-service. For example, e-mail is not generally required to be encrypted, however in some cases the involved parties may decide that more confidentiality is required. Depending on the level required, a simple encryption technique could be used, such as PGP (Pretty Good Protection), which can be cracked if enough time and resources are used. Or, a stronger encryption algorithm as discussed previously could offer a completely safe solution.

Both parties must negotiate and agree on what level is required for their particular service contract. Once it is in place, both parties must be assured that the security level described is in fact being performed. From the provider's point-of-view, methods must exist to verify the methods and test the overall security. Methods for quantifying security are required. From the customer's perspective, they must be assured that the provider's security claims are in fact valid.

8.2. TRUST

This raises the issue of trust. In traditional business, trust is based in large part on reputation and the permanence of brick-and-mortar. However, in the geographically distributed and dynamic environment of the Internet, these assurances are of little value. We have already described the use of certificates as one mechanism for establishing a form of trust, at least in terms of identity. Even with certificates though, the Certificate Authority must be a trusted party, or must be certified by a higher level CA who is trusted.

One method to increase trust in a system is the increased use of standardization. When both parties are using standardized procedures and services, they can at least know what they are dealing with. Although exposing a flaw in a standardized package may compromise a large number of systems, it may also aid in closing such holes.

Another method is through testing. Security mechanisms have often been tested in a rather loose manner. In order for the customer to trust the provider's security, it must have confidence in both the company's security claims and their testing methods. A solution to this is the use of third-party trusted testing organizations, paired with standardized testing procedures. An example of this would be the United States' Na-




Figure 10 The completed picture of network management providing services to the different layers of the e-service.

tional Information Assurance Program (NIAP, which includes research and development involving the creation of both Security Requirement Profiles (SRPs) and Common Criteria for Information Technology Security Evaluation (CCs)[25]. A NIAP-certified testing facility can then issue a certificate to an organization or service as an indication of its trustworthiness.

9. CONCLUSION

Throughout this paper we have tried to create a picture of how network management relates to providing security in e-services. This picture is now complete, as seen in figure 10. In this model, the e-service application exists on top of the system, which in turn exists on the network. The application relies on system services, and both application and system rely on network services - the application making use of those services through the system.

We have shown how security issues exist across all of these levels, with many issues common to all of them. A full set of security services must be offered to all layers in order to secure the system as a whole. This is the role network management plays, by providing and supporting these network-wide services. The service set is managed by the network management architecture that provides the necessary infrastructure in order to develop, implement, and maintain each of the individual services.

Finally, we stepped back to look at security management as one of the five functional sections of network management. The importance of providing all management functions in a secure way was stressed, depicted in the diagram as security offering additional protection to each of the other four areas. As the areas are so strongly interdependent, the diagram generalizes from an offered set of security services, to a set of

management services. Network management provides this set of services to each of the components of the e-service.

We briefly touched on the idea establishing security assurances and trust. This is an issue of critical importance, especially to an e-service. It is also an area in its infancy, with much work to be done in developing the system and methodologies.

Clearly, security must be of great concern to the provider of an e-service. It is the providers assets, systems, reputation, and bottom line that are ultimately at risk. As it is those providers who in large part will purchase the technological solutions to secure their e-services, it is up to the network management field to continue to develop and improve those technologies to protect them. That's *our* bottom line at stake.

References

- [1] V.L. Voydock and S.T. Kent, "Security mechanisms in high-level network protocols," *ACM Comp. Surv.*, 1983, 15(2), pp. 35-71.
- [2] A.S. Tanenbaum and M. van Steen, *Distributed Systems: Principles and Paradigms*, Prentice-Hall, Upper Saddle River, N.J., 2002.
- [3] A. Langsford and J.D. Moffett, *Distributed Systems Management*, Addison-Wesley, Wokingham, England, 1993.
- [4] P.A. Janson, "Security for Management and Management of Security," *Network and Distributed Systems Management*, M. Sloman, ed., Addison-Wesley, Wokingham, England, 1984, IP Security Document Roadmap, pp. 403-430.
- [5] National Bureau of Standards, "Data Encryption Standard", FIPS PUB 46, January 1977.
- [6] B. Schneier, "Description of a New Variable-Length Key, 64-Bit Block Cipher (Blowfish)," *Fast Software Encryption, Cambridge Security Workshop Proceedings (December 1993)*, Springer-Verlag, 1994, pp.191-204.
- [7] R.L.Rivest, A.Shamir, L.M.Adleman, "A Method for Obtaining Digital Signatures and Public-Key Cryptosystems," *Communications of the ACM*, February 1978.

- [8] R. Thayer et al., “IP Security Document Roadmap,” IETF RFC 2411, November 1998.
- [9] B.C. Neuman and T. Ts'o, “Kerberos: An Authentication Service for Computer Networks,” *IEEE Communications*, 32(9):33-38. September 1994.
- [10] P. R. Calhoun et al., Diameter Base Protocol, IETF Internet Draft, draft-ietf-aaa-diameter-12.txt, July 2002. [Work in Progress]
- [11] C. Rigney et al., “Remote Authentication Dial In User Service (RADIUS),” IETF RFC 2865, June 2000.
- [12] C. Finseth, “An Access Control Protocol, Sometimes Called TACACS,” IETF RFC 1492, July 1993.
- [13] R. Rivest, “The MD5 Message Digest Algorithm,” RFC 1321, MIT Laboratory for Computer Science (April 1992).
- [14] National Institute of Standards and Technology, “Secure Hash Standard,” FIPS PUB 180-1, April 1995.
- [15] C. Adams and S. Farrell, “Internet X.509 Public Key Infrastructure Certificate Management Protocols,” IETF RFC 2510, March 1999.
- [16] ITU-T Recommendation X.509 (1997 E): Information Technology—Open Systems Interconnection—The Directory: Authentication Framework, June 1997.
- [17] A. Westerinen et al., “Terminology for Policy Based Management,” IETF RFC 3198, November 2001.
- [18] K. Chan et al., “COPS Usage for Policy Provisioning (COPS-PR),” IETF RFC 3084, March 2001.
- [19] J.D. Case et al., “Simple Network Management Protocol,” IETF Standard 0015, May 1990.
- [20] J. Galvin and K. McCloghrie, “Security Protocols for Version 2 of the Simple Network Management Protocol (SNMPv2),” IETF RFC 1446 (Historic), April 1993.
- [21] D. Harrington et al., “An Architecture for Describing SNMP Management Frameworks,” IETF RFC 2571, April 1999.

- [22] U. Blumenthal and B. Wijnen, “User-based Security Model (USM) for version 3 of the Simple Network Management Protocol (SNMPv3) ,” IETF RFC 2574, April 1999.
- [23] B. Wijnen et al., “View-based Access Control Model (VACM) for the Simple Network Management Protocol (SNMP),” IETF RFC 2575, April 1999.
- [24] R.R. Henning, “Security Service Level Agreements: Quantifiable Security for the Enterprise?” *ACM 1999 New Security Paradigm Workshop*, Ontario, Canada, 2000.
- [25] P.J. Brusil et al., “Emerging Security Testing, Evaluation and Validation: The Key to Enhancing Consumer Trust in Security-Enhanced Products,” in *Handbook of Communication Technologies: The Next Decade*, CRC Press, to be published. <http://niap.nist.gov/article.html>