



S/MIME
Secure Messaging Architecture
Overview

Version 1.0

Preface

Objective of this document

The objective of this document is to describe the S/MIME Secure Messaging Architecture which defines the components of a secure messaging system based on the S/MIME set of standards, their relationships to each other and to the environment in which they operate.

Messaging professionals who know and understand the content of this architecture should be capable of designing and/or deploying a secure messaging system based on the S/MIME set of standards.

Scope

This architecture is constrained to secure e-mail systems that use cryptographic approaches based on the IETF S/MIME and related standards. It includes both desktop-to-desktop and gateway-to-gateway modes of operation.

This document offers a high level description of the S/MIME Secure Messaging Architecture. References are included to full definitions of individual components, upon which The Open Group's S/MIME Secure Messaging Certification program is based.

Target audience

This document is intended for anyone involved in the planning, design or deployment of a secure messaging system based on the S/MIME set of standards and in particular for organizations and individuals planning to achieve certification under The Open Group's S/MIME Secure Messaging (SSM) certification program.

"This architecture assumes expertise in information security principles as defined in the S/MIME Secure Messaging Certification Core Syllabus".

Future Directions

Future versions of this document are intended to increase it's scope to address approaches to secure messaging that are not based on the use of the S/MIME set of standards, for example:

- Other cryptographic approaches
- The use of secure Web mail
- Securing other forms of messaging

Acknowledgements

The Open Group would like to recognise the contributions of the following people to the development of this document.

Wen Fang, The Boeing Company
Stephan Wappler, Noventum Consulting

Table of Contents

Preface	2
<i>Objective of this document</i>	2
<i>Scope</i>	2
<i>Target audience</i>	2
<i>Future Directions</i>	2
<i>Acknowledgements</i>	2
Table of Contents	3
1. Introduction	3
1.1 <i>Context</i>	3
1.2 <i>Vendor Specific Solutions</i>	3
1.3 <i>Complexity and Risks</i>	3
1.4 <i>An Architecture for Secure E-Mail</i>	3
1.5 <i>Relationship to TOGAF</i>	3
2. Scope of the S/MIME Secure Messaging Architecture	3
2.1 <i>Perimeter Architecture</i>	3
2.2 <i>Internal Architecture</i>	3
3. Components of a Secure Messaging Architecture	3
3.1 <i>E-mail System</i>	3
3.1.1 <i>"To-End" Solutions</i>	3
3.1.2 <i>"To-Site" Solutions</i>	3
3.2 <i>Public Key Infrastructure (PKI)</i>	3
3.2.1 <i>Registration Authority</i>	3
3.2.2 <i>Certification Authority</i>	3
3.2.3 <i>Directory Services</i>	3
3.2.4 <i>Validation Services</i>	3
4. "To-End" Desktop Solution	3
4.1 <i>Applicability</i>	3
4.2 <i>Components</i>	3
4.3 <i>Solution Architecture</i>	3
4.4 <i>Operation</i>	3
4.5 <i>Implementation Considerations</i>	3
5. "To-Site" Gateway Solution	3
5.1 <i>Proxy Certificates</i>	3

5.2	<i>Domain Certificates</i>	3
5.3	<i>Internal Architecture Implications</i>	3
6.	Certification Authority Service Models	3
6.1	<i>Self-sign CA Model</i>	3
6.2	<i>Purchase CA Model</i>	3
6.2.1	Purchase certificate(s) from a full service PKI vendor	3
6.2.2	Purchase full service PKI solution	3
6.2.3	Purchase self-service PKI solution	3
7.	Directory Services	3
7.1	<i>Manual Key</i>	3
7.2	<i>Search and Discovery using LDAP</i>	3
7.3	<i>Publication of Certificates and CRL's</i>	3
7.4	<i>Enabling Access to Directories</i>	3
7.5	<i>LDAP Proxy System</i>	3
7.5.1	LDAP Proxy as External Directory Presence	3
7.5.2	LDAP Proxy as Internal LDAP Service Provider	3
7.5.3	With the LDAP Protocol through the corporate firewall	3
7.5.4	LDAP proxy as single point of presence in organizations	3
8.	Validation Services	3
8.1	<i>Certificate Revocation Lists (CRL)</i>	3
8.2	<i>Online Certificate Status Protocol (OCSP)</i>	3
9.	Bridge Infrastructure Solution	3
9.1	<i>Bridge-CA Solutions</i>	3
9.1.1	Signed List	3
9.1.2	Cross Certification	3
9.2	<i>Bridge-Directory Solution</i>	3
9.3	<i>Bridge-Validation Solution</i>	3
	Appendix 1 : Abbreviations	3
	Appendix B : References	3

1. Introduction

1.1 Context

Increasingly, organizations have a requirement for secure electronic communication with their external business associates. This need is being driven by a combination of regulation and business concerns.

Individual standards exist today to define individual components of a secure communications infrastructure such as

- S/MIME for encoding/decoding e-mail messages,
- PKI for certificate trust hierarchy and key escrow, and
- LDAP to provide interoperable electronic directory systems between companies.

Little has been done to date to integrate the existing protocols and e-mail components into a total architecture for the e-mail user community, and to address the organizational aspects of the required solution.

1.2 Vendor Specific Solutions

Many software vendors are addressing this problem with proprietary protocols and components. These vendor specific solutions are frequently difficult to integrate and/or interoperate with the organization's existing e-mail system. If an organization deploys one vendor specific product for e-mail encryption, without an architectural approach to ensure interoperability, then all of that organization's external partners will be forced to use the same product to exchange e-mail securely, which would in practical terms mean that organizations would need to support multiple e-mail systems for communication with different partners, which rapidly multiplies the cost of the encryption solution.

It is not acceptable for any large organization to base its secure e-mail infrastructure on a solution that is only available from one vendor.

Another type of proprietary solution – that of the privately owned, sole authority third party key broker, has similar drawbacks. It requires the use of a single third party that every organization must trust to broker key information and ensure that the information is shared only with the appropriate partners. This solution places the whole architecture of the key management infrastructure in the hands of the third party providing the broker service.

1.3 Complexity and Risks

Much effort has been focussed on the "desktop-to-desktop encryption" with the "person-to-person key exchange" solution. While this approach provides a high degree of security, it brings with it a number of disadvantages:

- This solution does not scale well because each user must manually and individually exchange certificates with every person to whom they wish to send or receive encrypted e-mail. The

management of keys, including revocation becomes a complex and expensive activity, which may be beyond the capabilities of small organizations.

- Certificates are widely available on the Internet, which creates the possibility that employees may obtain their own keys directly from the Internet, rather than from an organization owned or approved certificate supplier, in which case, the organizations may not be able to decrypt and will thus lose all the data if the employee loses his or her private key or if the employee leaves the organization.
- Desktop-to-desktop encrypted messages can only be scanned for viruses or other forms of unacceptable contents at the desktop. Enterprise level filtering has no access to the content of the message.

1.4 *An Architecture for Secure E-Mail*

The goal of this document is to describe an architecture and standard methodology for exchanging encrypted e-mail between organizations. The architecture described

- is based on existing Internet standards;
- is vendor neutral;
- and can be integrated into existing e-mail systems.

The architecture provides both the internal and perimeter architecture views, and optional configurations, which address the differing needs of different types and sizes of organization.

Two distinct modes of secure e-mail communication are described

- In the "To End" Desktop solution, e-mail is delivered encrypted to the end-user for decryption at the desktop.
- In the "To Site" Gateway solution, e-mail is delivered to a gateway system for decryption and onward transmission to the end-user.

1.5 *Relationship to TOGAF*

The S/MIME Secure Messaging Architecture can be viewed as a Common Systems Architecture within the overall continuum of architectures defined within The Open Group Architecture Framework.

As such it defines a self-contained secure e-mail "building block" that can be integrated into a more complete enterprise IT architecture.

It contains within it a number of specific Architecture Patterns for different applications of secure e-mail.

2. Scope of the S/MIME Secure Messaging Architecture

The scope of the S/MIME Secure Messaging Architecture is limited to the configuration of the perimeter services and the assertion that certain industry standard protocols and conventions be used. The intention is to make the architecture as widely adoptable as possible by organizations that have already deployed messaging and directory and validation infrastructures.

2.1 *Perimeter Architecture*

All perimeter systems shall be required to adhere to the following:

- **Directory:** Each organization must provide a publicly accessible LDAP directory to support lookup of e-mail addresses and user certificates
- **Certificates:** All certificates must be issued using X.509 v3 CA Services. Certificates may be self-signed or commercially purchased. Certificates must use RSA algorithm with a minimum 1024 bit key length
- **Directory Protocols:** All perimeter directories must support LDAP v3.
- **Validation Protocols:** All perimeter services must support CRL's and may optionally also support OCSP.
- **E-mail:** S/MIME compliant e-mail client or gateway must be capable of requesting certificates from directories. Additionally the clients or gateways must be able to validate certificates via CRL's or via OCSP.

In restricted closed-community deployments of "to-Site" solutions, alternative forms of key exchange (e.g. taking certificates from a digitally signed message from another gateway system) may be practical. In such deployments, the perimeter system is not required to support directory or validation protocols.

2.2 *Internal Architecture*

The organization e-mail architecture must be S/MIME compliant to the point at which messages are first decrypted. The specific internal architecture for mail handling between that system and the originator or recipient is out of scope of this architecture.

For "to-End" solutions, the messaging clients must support the request of certificates from an LDAP directory and certificate validation. For "to-Site" solutions, the messaging gateway must either support the request of certificates from an LDAP directory and certificate validation, or must provide a mechanism for manual key exchange.

3. Components of a Secure Messaging Architecture

3.1 *E-mail System*

3.1.1 "To-End" Solutions

In a "To-End" solution, e-mail is delivered encrypted to the end-user for decryption at the desktop.

Based on the accepted certificate practices, the recipient's e-mail encryption certificate will be published into the LDAP directory when issued. The certificate will be removed from the directory system when become invalid for any reason. The validity of a certificate can be checked using either Certificate Revocation Lists (CRL) or Online Certificate Status Protocol (OCSP).

The e-mail originator will request the recipient's certificate from an LDAP directory with their e-mail client. After receiving the recipient's certificate, the sender validates it with either CRL or OCSP at their local desktop. If the validation of the certificate is successful, then the user can encode and sign the e-mail with their local e-mail client. If the sender is unable to validate the recipient's message, then the decision on whether to proceed and encrypt the message depends on the recipient's PKI practice and/or the sender's local policy.

On receipt of an encrypted and/or signed e-mail the recipient decrypts the e-mail and validates the signature with their local e-mail client.

This is characterized as "To-End" solution since all the procedures and trust are done on either the originator's or recipient's desktop.

3.1.2 "To-Site" Solutions

In a To-Site solution, e-mail is delivered to a secure e-mail gateway system inside the recipient's intranet for decryption and onward transmission to the end-user

The e-mail encryption certificate for the recipient domain may be obtained from the recipient's LDAP directory, or may be obtained by an alternative means such as taking the certificates from a digitally signed message from the recipient's gateway system.

The message can be sent to the originator's gateway either encrypted or as clear text, depending on the company policy. The secure e-mail gateway system will then forward the message to the intended recipient(s) encrypted with the appropriate recipient. The secure e-mail gateway system may validate the certificate with CRL or OCSP before encrypting and forwarding the e-mail message.

The originator of the e-mail does not need to be aware of the encryption process and does not need access to the recipient gateway system's certificate. If the originator's local policy requires encryption of the message during transmission to the outbound gateway system, the originator will need to store the outbound gateway system's public key in the e-mail client, in which case the gateway will decrypt the message before encrypting with the recipient's gateway certificate for onward transmission to the recipient(s).

This type of operation is characterized as the "To-Site" solution since all the procedures are done on the site gateway system and not on the e-mail client.

3.2 Public Key Infrastructure (PKI)

The PKI provides the X.509 V.3 certificates required to enable a secure exchange of digital signed and/or encrypted e-mails in open networks.

PKI components:

- Certificate Authority (CA)
- Registry Authority (RA) or Local Registry Authorities (LRA)
- Directory Service
- Online Validation Service

3.2.1 Registration Authority

The Registration Authority (RA) is responsible for recording and verifying all information the CA needs to issue a certificate. The RA can be

- an online application for accepting certificate request,
- a set of procedures,
- a person in charge of validating the identity of the person that requesting the certificate,
- or any combinations of them.

The Registration Authority has two main functions:

- Verify and record the identity of the requestors
- Interact with the CA to issue and deliver the requested certificates

3.2.2 Certification Authority

The Certificate Authority (CA) is the entity responsible for issuing and administering the digital certificates. The CA acts as the trusted agent in the PKI.

A CA performs the following main functions:

- Generate public/private key pair, or accept A user generated key pair
- Package public keys in digital certificate
- Sign certificate
- Publish certificate in designated Directory
- Issues certificate revocation lists (CRLs)

The foundation upon which a PKI is built is trust - in other words the user community must trust the CA to distribute, revoke, and manage keys and certificates in such a way as to prevent any security breaches. As long as users trust the CA and its business processes, they can trust certificates that the CA issues.

The certificate issued from the CA server is a signed electronic document. Based on the accepted digital signature validation process, a user or application can determine if a certificate has been altered. The CA's signature in a certificate ensures that any changes to its contents will be detected. Such certificates can be distributed publicly, and users retrieving a public key from a certificate can be assured of the validity that the key:

- Belongs to the entity specified in the certificate
- Can be used safely in the manner for which the CA certified it

Users need to be able to determine the degree of assurance or trust that can be placed in the authenticity and integrity of the public keys contained in certificates the CA issues. The information upon which such determinations can be made is documented in the Certificate Policy and the Certification Practice Statement of the CA.

A CA has the following tasks:

- Generate the certificate based on a public key. Typically a Trust Center generates the pair of keys on a smart card or a USB token.
- Guarantees the uniqueness of the pair of keys and links the certificate to a particular user
- Manages published certificates
- Is part of cross certification with other CAs

3.2.3 Directory Services

The current PKI best practice is integrating CA and Directory in the certificate publishing process. The Directory system needs to support the industry accepted directory access protocol, LDAP V.3. The directory service has two main functions:

- Publish certificates
- Publish a Certificate Revocation List

3.2.4 Validation Services

The function of the online validation service is to determine whether a certificate is valid via the standard Online Certificate Status Protocol (OCSP). This is an alternative to using Certificate Revocation Lists. The online validation service is also referred to as the Responder.

4. "To-End" Desktop Solution

4.1 *Applicability*

There are 2 main reasons for choosing an "end-To-End" e-mail encryption solution:

- PKI Certificate Policy (CP) – requires that the user private key for decryption can only be held and accessed by the holder of the certificate. This policy requirement makes it almost impossible to implement a "To-Site" solution if you need to send encrypted e-mail to a recipient and only the recipient, not the gateway system, can decrypt the e-mail since the recipient is the only person who has the decryption private key.
- Company internal security policy – when a company in its normal course of business has the requirement to encrypt and protect data within its internal network.

4.2 *Components*

The To-End solution involves the following components:

- E-mail client application that support S/MIME standards for encryption/decryption of e-mail messages, and LDAP protocol for search and discovery of the recipient's encryption certificates
- E-mail server that can deliver S/MIME encoded e-mail messages
- PKI and CA systems to issue encryption certificates, and publish them to the directory
- Directory system allowing automated search and discovery of e-mail encryption certificates
- Optional LDAP Proxy to facilitate publishing certificates externally and also search and discovery of recipient certificates

4.3 Solution Architecture

Figure 4.1 shows the structure of the End-To-End e-mail encryption architecture:

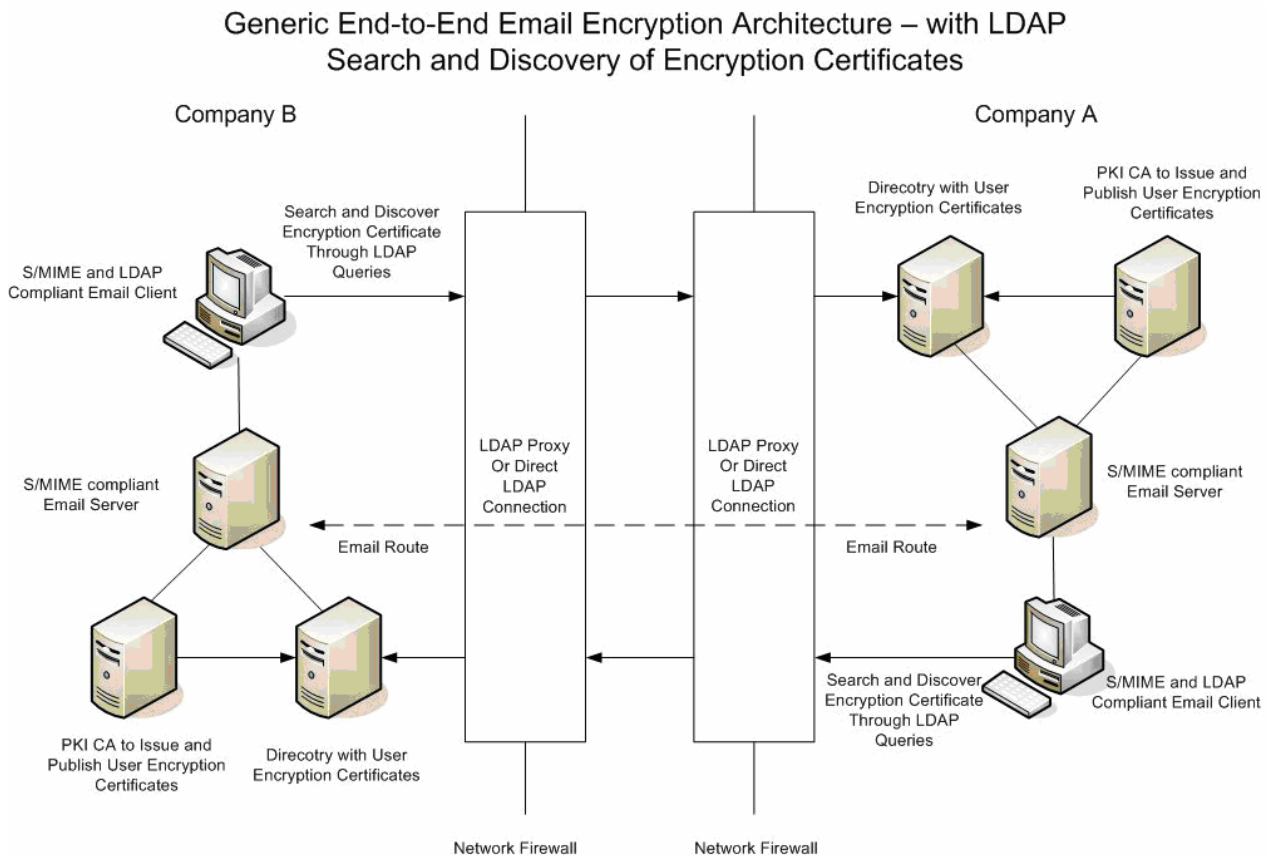


Figure 4.1

4.4 Operation

The overall end-To-End e-mail encryption process is as follow:

- The sender and recipient both have the required components – PKI/X.509 Certificate, S/MIME compliant e-mail server, S/MIME and LDAP compliant e-mail client, and LDAP Proxy or server.
- The sender will search for the recipient's encryption certificate through the LDAP proxy
- The LDAP Proxy will parse the request for e-mail address and e-mail domain name
- The LDAP Proxy will try to match the e-mail domain name in its configuration entries

- The LDAP Proxy will make a query to the designated LDAP Proxy or server if a match is found
- When a successful query result is received, the LDAP Proxy will check the certificate to make sure its capable for encryption before forward back to the sender that requesting the certificate
- The sender e-mail client will need to validate the certificates locally for:
 - PKI Trust Relationship
 - Validate the recipient's certificate with CRL and/or OCSP
- The sender will then encode the message with recipient's public key to create S/MIME compliant message body, and add digital signature with sender's public key
- The message is sent as a regular SMTP message through the normal message delivery routes
- The recipient will decrypt the message with his private key, and validate any digital signature for:
 - PKI Trust
 - Validate the sender's certificate with CRL and/or OCSP

4.5 *Implementation Considerations*

The to-End or "end-to-end" architecture can provide the highest level of security, and can support trust models between individuals. The precise level of trust supported will depend on the specific certificate policies in place at each end.

The "to end" architecture tends to have the highest installation and operational costs of the models described due to the need to manage certificates for each end user, as well as due to the need for end-user training in certificate management, client issues, and general security.

5. "To-Site" Gateway Solution

The major benefit of the To-Site Gateway solution is that all incoming and outgoing encrypted e-mail are managed at the central gateway.

At this central point the e-mail messages may be scanned for viruses, Trojan horses and for critical and dangerous content of the e-mail.

There are two different to-Site Gateway architectures:

- Proxy Certificates
- Domain Certificates

5.1 *Proxy Certificates*

One mechanism to enable scanning at a central point for viruses, Trojan horses and for critical and dangerous content of e-mails is to store the private key of the company users at the central secure gateway environment.

The private keys of all the users associated with this gateway are stored in a secure environment at the gateway. If the gateway receives an encrypted incoming e-mail, then the gateway attempts to locate the private keys associated with the recipient address.

If the private key associated with the public key used to encrypt the message is stored at the gateway, then the gateway decrypts the message allowing for it to be scanned for viruses and Trojan horses and filtered for unacceptable content.

If the e-mail was digital signed, the gateway validates the signature of the e-mail and information about the signature status is added to the original e-mail.

If following decryption and subsequent filtering the message is accepted, then, based on local policy, the Gateway

- forwards the original encrypted e-mail to the intended recipient or
- forwards the decrypted e-mail to the intended recipient or
- encrypts the e-mail with the public key of an representative of the original intended recipient or
- forwards the decrypted e-mail to a representative of the original intended recipient.

If the message cannot be decrypted, or is unacceptable for any other reason. the Gateway forwards the e-mail to an administrator account.

Depending on local policy, outgoing e-mails may also be handled at the gateway:

- If the user encrypts and/or signs outgoing e-mails at his desktop, the gateway can either forward such e-mails to the internet for authroized users or block them.

- The gateway may remove the user signature and add a signature at the gateway. The gateway can generate a "personal" signature for an outgoing e-mail, using the private key associated with the sender stored at the gateway.
- If the gateway has a valid public key for the recipient or if the gateway is able to retrieve a valid public key from a directory then the gateway is able to encrypt the outbound e-mail.

For selected recipients of outgoing e-mails a rule may be created that e-mail has to be encrypted or the sender may add special commands to the gateway defining how it has to handle this e-mail (with or without a signature; or e-mail transfer is only allowed if the e-mail is encrypted).

This solution is a to-Site solution. The private key of the user is under the control of the gateway and the gateway administrators. It is not necessary for the user to interact with the gateway.

This architecture, often referred to as "proxy certificates" (as the gateway effectively proxies for the end user) emulates an end-to-end architecture (in that each end user has an individual certificate) while attempting to provide centralized management. It is a hybrid solution that has some distinct advantages over a "true" end-to-end solution in reduced costs through centralized certificate management and lowered user training. It also is fully interoperable with existing S/MIME client software at external sites.

The trust model here is really site-to-site, or domain-level, as encryption and signatures are applied only at the gateway, and are not under the end-user's control. A risk of this architecture is that, since it masquerades as end-to-end, external users might believe that they are operating in an end-to-end trust environment, with the corresponding implications on confidentiality and signature semantics. This risk must be recognized and dealt with administratively when this model is deployed.

5.2 Domain Certificates

A second mechanism to enable scanning at a central point for viruses, Trojan horses and for critical and dangerous content of e-mails is to store the private key of the company users at the central secure gateway environment and to encrypt and sign e-mails or decrypt and validate signatures of e-mails is to use domain certificates.

A gateway is associated with one or more domains or sub-domains. Messages between gateways are handled based on their domains, not on the basis of individual certificates. Gateway certificates are called "Domain Certificates", and have the same format as S/MIME Version 3 certificates [CERT31].

The sender of an e-mail encrypts the e-mail with the public key contained in the domain certificate of the receiver prior to transfer over the internet to the intended recipient user domain.

The received encrypted e-mail is decrypted at the domain gateway using the related private key of the domain to decrypt the e-mail allowing for it to be scanned for viruses and Trojan horses and filtered for unacceptable content. If the e-mail was digital signed than the gateway validates the signature of the e-mail. Information about the signature status is added to the original e-mail.

If following decryption and subsequent filtering the message is accepted, then, based on local policy, the Gateway

- forwards the decrypted e-mail to the intended recipient or
- forwards the decrypted e-mail to a representative of the original intended recipient.

If the message cannot be decrypted, or is unacceptable for any other reason. then the Gateway forwards the e-mail to an administrator account.

Depending on local policy, outgoing e-mails may also be handled at the gateway:

- If the user encrypts and/or signs outgoing e-mails at his desktop, the gateway can either forward such e-mails to the internet for authroized users or block them.
- The gateway may remove the user signature and add a signature at the gateway. The gateway can generate a "personal" signature for an outgoing e-mail, using the private key associated with the sender stored at the gateway.
- If the gateway has a valid public key for the recipient or if the gateway is able to retrieve a valid public key from a directory then the gateway is able to encrypt the outbound e-mail.

For selected recipients of outgoing e-mails a rule may be created that e-mail has to be encrypted or the sender may add special commands to the gateway defining how it has to handle this e-mail (with or without a signature; or e-mail transfer is only allowed if the e-mail is encrypted).

This solution is a to-Site solution. The only key required is a domain key. and this key is under control of the gateway and the gateway administrators. It is not necessary for the user to interact with the gateway.

This architecture is a "pure" site-to-site or domain solution. It explicitly implements a domain-level trust model and does not masquerade as anything else. Unfortunately most current e-mail clients do not support domain-level S/MIME, so it is applicable only for a gateway-to-gateway implementation, unlike proxy certificates, which can be used in a gateway-to-end-user mode as well.

Advantages of the domain-to-domain model include the reduced costs associated with centralization of certificate management and lowered end-user training costs. In addition, this model reduces the number of certificates required for operation, often by several orders of magnitude. This means that in certain environments and communities the model can be deployed without requiring the use of automated certificate exchange processes and on-line certificate validation. These capabilities are desirable, and necessary as larger communities of domain-level encryption are deployed.

5.3 Internal Architecture Implications

In either "to-site solution" confidentiality of e-mail is assured only to the gateway, i.e. to the site perimeter. Additional security is desirable within an organization, e.g to protect against the possibility of "sniffing" of a corporate network (always possible, but much more common in a wireless-enabled network). Some e-mail products include encryption within the organization. In the absence of this link encryption of server-to-server and client-to-server communications can be highly effective. These are implemented using SSL or TLS between e-mail components.

TLS can also be used in a server-to-server environment outside the organization. While TLS is a low-cost and low-impact encryption solution, it does not provide for encryption of the data at rest (i.e. on servers and relays), so cannot be recommended as a strong security model.

6. Certification Authority Service Models

The PKI infrastructure provides a logical level framework of trust for companies. The CA and X.509 certificate provide a physical link for validating the PKI trust. The secure e-mail, either to-user or To-Site, can open operate with trusted entities.

There are the basic three models of CA services:

- Self-sign CA Model: The organization creates its own infrastructure for the issuance and management of certificates
- Purchase CA Model: The organization purchases certificates from a trusted commercial vendor
- Bridging Model [addressed in detail in chapter 9]

6.1 *Self-sign CA Model*

There are many reasons why an organization might choose this model and setup its own self-sign CA servers.

- For example, self-sign CA service has a cost benefit over purchase CA for organizations with a requirements for a large number of certificates
- The self-sign CA service carries a stronger corporate identity than use of a purchase CA model.

A company and its business associates must determine whether a self-sign CA model is appropriate to their environment.

In this model, the Root CA server is signed by itself, and cannot be traced to any known or trusted PKI vendor. The root CA server can then sign for subordinate CA servers for the issuing of certificates to users for specific purposes, such as e-mail encryption, digital signature, identity management etc. Having at least 2 levels of CA servers is recommended to protect the Root CA server.

The disadvantage of this CA model relates to the issue of certificate trust. The root and subordinate CA certificates of most commercial PKI vendors are already installed in most operating systems and Internet browser packages. Users who purchase certificate from these vendors do not need to worry about trust issues. Users with self-sign certificates will need to manage the PKI trust with their external business associates.

6.2 *Purchase CA Model*

In this CA model, the certificates are issued by a PKI vendor (also know as Trust Center). There are several different service offerings from PKI vendors within this model, depending on the extent to which the purchaser assumes management responsibility:

- Purchase individual certificate(s) from a full service PKI vendor
- Purchase a full service PKI solution, managed by the PKI vendor

- Purchase self-service PKI solution, managed by the purchaser

6.2.1 Purchase certificate(s) from a full service PKI vendor

In this case, the user purchases one or more certificate(s) individually from the PKI vendor. The PKI vendor provides all of the necessary interfaces for certificate management (request, issue, download, archive, and revoke). The PKI vendor also provides the certificate validation services for CRL checking and optionally OSCP validation, and publishes the certificate(s) to Internet accessible directory systems.

6.2.2 Purchase full service PKI solution

In the previous option, the user certificates are issued from the PKI vendor's general CA servers. This option differs from the previous one in the CA server that issues the certificates. In this option, a company can purchase a dedicated subordinate CA server to issue certificates only to its authorized persons.

The PKI vendor manages all of the certificate management tasks (request, issuing, archive, revoke etc.). The PKI vendor provides the CRL and OSCP validation services, and publishes the certificate(s) to directory systems.

6.2.3 Purchase self-service PKI solution

This option is similar to the full service managed PKI solution except that the certificate management tasks are done by the company's authorized personnel. The PKI vendor will provide the needed self-service interface for the company to manage the certificate management.

7. Directory Services

Companies and organizations which deal with secure e-mails have to address two different problems:

- They must provide their communication partners with the X.509 encryption certificates of their own employees in a form that they can access. This involves making information about the directory services of their partners available to e-mail clients and provision of internet access via internal networks.
- They are faced with the task of enabling their users to access the directories of their business partners to retrieve the encryption certificates of their business partners' users.

This situation becomes more complex when mobile devices and mobile users need to be able to access directory information from different locations inside and outside of the company.

7.1 Manual Key

The first e-mail encryption used manual encryption key exchange, partially due to the lack of Internet standards for automated encryption key delivery. Technically manual key exchange works. It is not, however, scalable for enterprise-wide deployment for a large user base.

For each individual with whom you need to exchange encrypted e-mail, it is necessary to send out your encryption certificate and receive back the other person's encryption certificate via digitally signed e-mail messages, a process known as "manual key exchange". For an enterprise with 10,000 e-mail users, if each one of them needs to exchange encrypted e-mail with 10 other external business associates, there will be 100,000 (10,000x10) certificates to be sent out and another 100,000 to be received back and processed. This is not an automated process and is likely to generate a significant technical support load.

Based on the current PKI Certificate Policy, e-mail encryption certificates expire and need to be renewed once a year. This means that a lengthy and error-prone manual-key exchange would have to be repeated annually.

7.2 Search and Discovery using LDAP

An alternative to the manual key exchange method is to search and discover encryption certificates through Light-weight Directory Access Protocol (LDAP) queries. The LDAP protocol has been established as an industry standard for accessing electronic directory systems. The LDAP protocol is supported by almost all major e-mail software vendors. Using LDAP, e-mail users can simply request the recipient's encryption certificate through LDAP queries from the designated electronic directory system.

7.3 Publication of Certificates and CRL's

Every organization that uses certificates for e-mail encryption needs to publish their certificates on the public internet either directly themselves or by using a service provider to enable their communication partners to retrieve certificates for an encrypted communication without the need for any prior communication from sender to recipient. This problem would be solved by a public directory service on the internet.

Additionally, organizations need to publish information about the validity of their issued certificates, such as revoked but not yet expired certificates, i.e. invalid certificates. Such information is required by the originator before using the recipient certificate to encrypt the e-mail.

Revocation information can be published in a Certificate Revocation List (CRL). The certificate itself should define the location from which this information may be downloaded (e.g. Web server or Directory service).

7.4 *Enabling Access to Directories*

There are problems associated with enabling business partners to access directories from their user e-mail clients or secure mail gateways:

- It is necessary to open the firewall for the LDAP-Protocol and;
- It is necessary to manage external access to information contained in the directories
- It is necessary to enable internal access to information contained in directories of business partners

These problems may be addressed using an LDAP Proxy system

7.5 *LDAP Proxy System*

Most companies will not allow external access to their complete directory systems. Most organizations directories contain information that should not be made externally available. Maintaining a separate extract of the internal directory for access by external partners introduces duplication and the need for synchronization.

The LDAP Proxy approach been developed to address this problem.

The LDAP Proxy is designed to deploy in the company network De-Militarized Zone (DMZ). It has the dual functionalities:

- An external directory presence to share certificates
- An internal LDAP service interface for finding external certificates associated with business partners

7.5.1 *LDAP Proxy as External Directory Presence*

When the LDAP Proxy is deployed in the DMZ, it can connect to the internal directory system to share user certificates with external business associates, with the following security functions:

- LDAP Proxy only accepts queries including a valid e-mail address
- LDAP Proxy does not respond to queries with partial or wild-card queries
- LDAP Proxy only passes the following 3 attributes back in a query result:
 - E-mail address (which has to be identical to the e-mail address in the original query)
 - Common Name

- Encryption certificate
- Since a Distinguished Name (DN) name must be returned in any successful query result, LDAP Proxy re-writes the DN to hide internal directory schema structure

7.5.2 LDAP Proxy as Internal LDAP Service Provider

The LDAP Proxy can be deployed to be the Internal LDAP Service Provider to enable internal e-mail users to find external certificates. The LDAP Proxy is designed to store information from multiple external directories in its configuration, grouped by the e-mail domain name.

When an LDAP query for an external certificate is received by the LDAP Proxy, it parses the query to find the e-mail address for which the certificate is requested. The LDAP Proxy will further parse the address for the e-mail domain and attempt to locate the entries for that domain in its configuration file.

If there is a match of the e-mail domain entry, the LDAP Proxy enables the LDAP query for certificate. If a successful search result is received, the LDAP Proxy examines the certificate for the following attributes:

- "E-mail" or "Subject Alternative Name" contains the e-mail address in the original request
- Check the expiration date of the certificate to make sure its not expired
- Check the key usage for "key encipherment" for the e-mail encryption certificate

Only when the certificate(s) passes all the checks, is it returned to the internal e-mail user for encryption.

The LDAP Proxy accepts multiple entries in its configuration for the same e-mail domain name to provide directory redundancies. In this case, the LDAP Proxy will perform sequential searches from the first configuration entry to the next.

The LDAP Proxy has the ability to create proxy chaining. The local proxy forwards any un-resolved queries to a higher level proxy server. This chain of proxies can facilitate the search and discovery of certificate for domains not known to the local proxy.

7.5.3 With the LDAP Protocol through the corporate firewall

The LDAP protocol is a problem for many companies.

- Many companies are unwilling to open corporate firewalls for the LDAP protocol and thus permit insecure access to the Internet from the internal network.
- Not all operators of a directory service use the standard port 389, but instead use their own ports for different reasons. This means that other ports have to be opened on the firewalls of their partner.

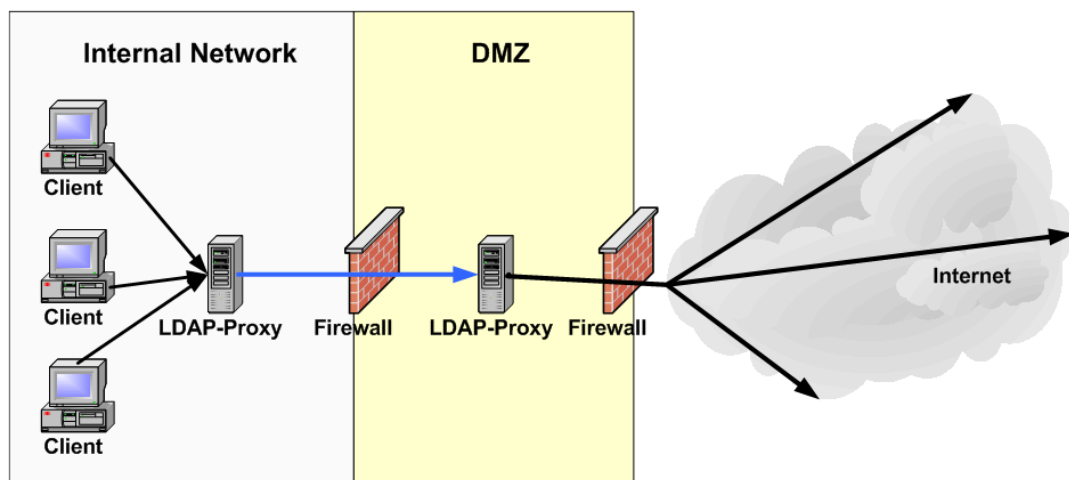


Figure 7.1 : Cascading LDAP Proxy Servers

These problems can be resolved by cascading LDAP proxy servers (see Figure 7.1).

- The internal clients access the first proxy server which is installed in the internal network using the standard port.
- This passes the query onto the second LDAP proxy server installed in the DMZ using a port which is known only between these two servers and which is opened only at the internal firewall. Thus, the servers cannot be accessed from both inside the company through the internal firewall and the Internet through the external firewall using the same port.
- All partner directory services, including the resulting port problems, are managed within the DMZ, in other words on the second LDAP proxy server.
- For this purpose, only the external firewall needs to be opened for the private ports of the partner directory services.

7.5.4 LDAP proxy as single point of presence in organizations

The LDAP proxy allows a single point of administration for an internal network (see figure 7.2). All e-mail clients or secure mail gateways need only be configured with one directory service entry. All partner directory services are configured and administered at the proxy. The clients or gateways send their queries to the proxy where they are checked (as in 7.5.2) and then sent to the corresponding external directory services.




Figure 7.2 : LDAP Proxy as a single point of presence

With this solution, it is very simple for system administrators to add a new partner directory service or to configure a new company e-mail client.

8. Validation Services

8.1 *Certificate Revocation Lists (CRL)*

The validation of certificates can be done by either the e-mail client or the gateway itself. The client/gateway needs the Certificate Revocation List (CRL) of the CA which issued the certificate which has to be validated. The client requests the CRL at the CRL distribution point (CDP) which can be found in the CDP attribute of the user certificate. Normally, the CDP is given as an LDAP or HTTP URL.

A CRL includes the serial number, the time and the reason for the revocation of the certificate [RFC3280]. The CLR is digitally signed by the issuing CA to ensure the integrity of its contents.

8.2 *Online Certificate Status Protocol (OCSP)*

An alternative way to validate a certificate is the use of the Online Certificate Status Protocol (OCSP) [RFC2560]. The Client can check the status of a certificate, online, by using an OCSP Responder.

The OCSP Responder can send three status results back to the requesting client:

- Valid
- Not valid
- Status unknown

The Response is digitally signed by the Responder to ensure integrity.

In principal, OCSP Responders can be chained, the OCSP Responder can forward the request of the client to a further OCSP Responder, effectively acting as a bridge. In this case the client needs the certificate of the bridge OCSP to be able to verify the response from the bridge OCSP responder.

9. Bridge Infrastructure Solution

The objectives of a Bridge Infrastructure are to link CAs into a "virtual infrastructure" which allows one CA to deliver security certificates and associated information from other CAs to which it is linked by the Bridge Infrastructure. It enables secure and trustworthy communication between organizations who are using different CAs.

The main tasks of the bridge infrastructure are:

- Creation of trust relationship between organizations,
- Access on participants and/or end entity-certificates,
- Validation of Certificates and
- Creation of a contractual framework.

The bridge infrastructure establishes an organizational and contractual framework including:

- Assessment of security level for the participating CAs by examination of
 - Certificate Policies
 - Certification Practice Statements (or disclosed part of it)
 - Definition of Policy Mappings
- Contractual documents
 - Relying Party Agreement
 - Audits through the Bridge Infrastructure

9.1 Bridge-CA Solutions

The trustworthiness of a certificate depends on being able to link that certificate to a trusted Root certificate or subordinated CA certificate. The exchange of such certificates is a complex task.

Root or subordinated CA certificates of many Certificate Service Providers are now already pre-loaded into the certificate storage of commercial software, including Internet browser software, from different software manufacturers. If a user purchases a new device or installs such a software system, at the same time he has also acquired these Root or subordinated CA certificates and has installed these in the certificate storage.

However, the rigorous acceptance criteria for being stored in these list of base certificates can not be fulfilled by every Certification Service Provider and the same is true for organizations which operate their own CA. The Root certificates of these company CAs must be installed later. The installation and the required verification of such certificates is currently a very time-consuming process.

Manual exchange of the Root or subordinated CA certificates using an out-of-band root certificate delivery method or a cross-certification is not scaleable. The exchange via a Bridge-CA infrastructure currently provides the only practical alternative.

Where necessary, the creation of trust relationship between organizations will be done by the Bridge-CA:

- Exchange of Root- and Sub-CA Certificates or
- Cross-Certification of Sub-CA Certificates

9.1.1 Signed List

A prospective PKI provider or a Certification Service Provider who would like to participate in the Bridge-CA applies for admission to participate. A representative of the Bridge-CA checks the technical (e.g., interoperability) and organizational conditions for the PKI integration into the Bridge-CA. After all tests are concluded and the participants' contracts are signed, new participants take their root certificates to a RA where they must prove their identity. The root or subordinated CA certificate is registered and stored in the Bridge-CA directory in one of two ways:

Signed list:

- The certificate is stored in the centrally administrated list of the CA certificates of the other qualified participants
- The Bridge-CA signs the list and guarantees the trustworthiness of the list
- The signed list is offered to the participants (e.g., by e-mail or downloaded)

Bridge directory:

- The certificate is stored in the centrally administrated directory of CA certificates of the other qualified participants
- The Bridge-CA guarantees the trustworthiness of the stored certificates through inclusion the entry in the directory
- Every participant has access to the directory and can download the certificates of the other participants

This approach guarantees that interoperability exists between the PKIs. Participants of the Bridge-CA can individually decide whether they recognize all or only some of the supplied CA certificates as reliable and which CA certificates they import into their certificate storage.

Expired or revoked certificates are removed immediately from the directory or the signed list. For the distribution of the revocation information the CRL-Distribution Points (CDPs) of connected PKI's or the OCSP Responder Service of the Bridge-CA can be used. The OCSP Responder of the Bridge-CA also verifies the validity of the Root and subordinated CA certificates at the CRL-Distribution Points of the Certification Service Provider and PKIs.

9.1.2 Cross Certification

The trust between the member PKIs can be achieved by cross-certifying the corresponding member CAs with the bridge CA. The generation of cross-certificates is the basis for the secure relationship between the different PKI environments.

9.2 Bridge-Directory Solution

The exchange of certificates can be achieved by using directory technologies. The Bridge-Infrastructure can operate a Directory Proxy Service which connects the directories of the participating Certification Service Provider and the private CA-operators into a single virtual Directory.

The advantages of virtual directories are:

- Central administration point for the access data for all connected Certification Service Provider and PKIs directories
- Standardized access to the virtual Directory
- A client can search for information in several domains from a single entry point.
- Transformation of the requested data from the connected directories to the format required for transmission to the inquiring Client
- The Clients are able to access directories of different standards and formats, e.g.
 - Active Directory Service (ADS of Microsoft)
 - Novell Directory Service (NDS of Novell)
 - LDAP Directories of the different versions (LDAPv1, LDAPv2, LDAPv3)
 - X.500 Directories
 -
- Real time LDAP schema and attribute translation to eliminate directory incompatibility problem
- Single point of configuration for all participating directories
- Testing the request's legitimacy
- No redundant data - changes to entries in the directories are immediately active

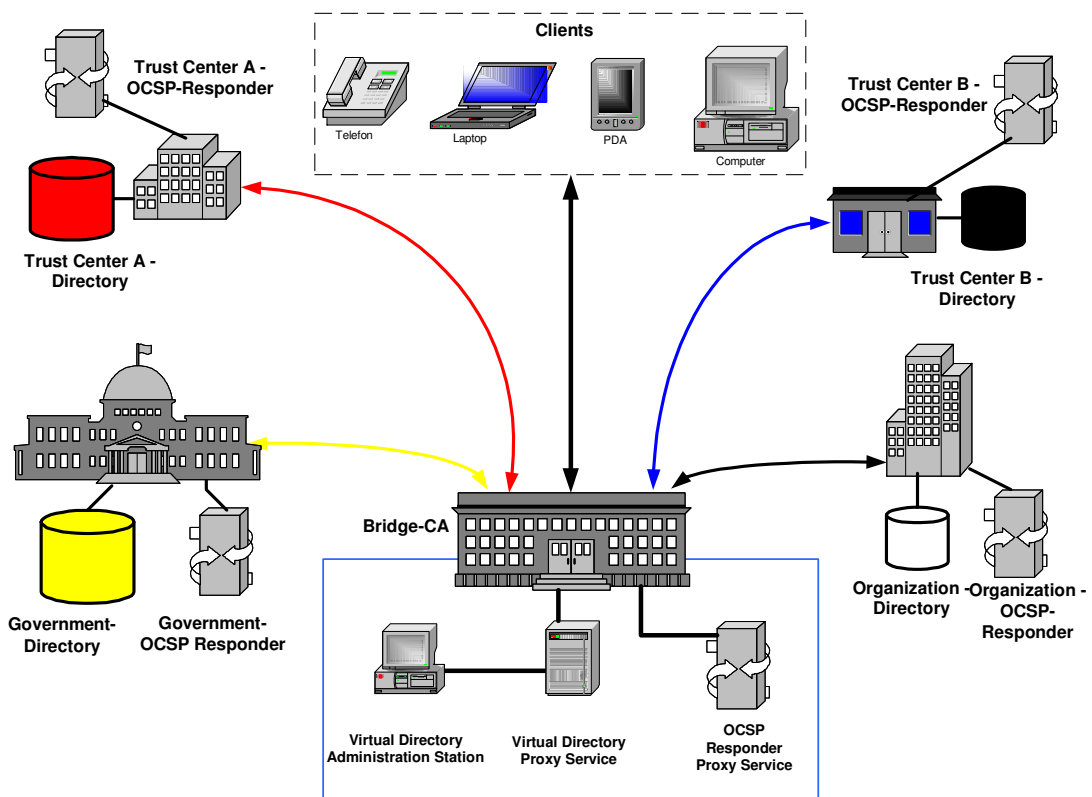


Figure 9.1 – Bridge CA Infrastructure

A precondition for this architecture is that the directories of the connected Certification Service Provider or company CAs can be accessed via the public network.

9.3 Bridge-Validation Solution

Validation of certificates is another major problem. The Bridge Infrastructure can operate an OCSP Responder Proxy Service which provides on virtual OCSP Responder which connects to the responders of the participating Certification Service Providers and the private CA-operators.

The advantages of such a Virtual OCSP Responder Service are:

- Central administration point for access to the OCSP Responders of all connected Certification Service Providers and private CAs
- Standardized access to the virtual OCSP Responder
- A single entry point through which clients can search for information in several domains
- Transformation of the data retrieved from the connected responders to the format required for transmission to the inquiring client
- Changes in the entries in the OCSP Responder Data Sources become effective immediately, rather than waiting for a replication phase.
- CRLs may be used as a basis for revocation information.

- Single point of configuration for all participating validation services

A precondition for this architecture is that the validation information from the connected Certification Service Providers or company CAs can be accessed via the public network.

Appendix 1 : Abbreviations

ADS	Active Directory Service
CA	Certification Authority
CDP	CRL Distribution Point
CRL	Certification Revocation List
CSP	Certification Service Provider
HTTP	Hyper Text Transfer Protocol
LDAP	Lightweight Directory Access Protocol
LRA	Local Registration Authority
NDS	Novell Directory Service
OCSP	Online Validation Status Protocol
PKI	Public Key Infrastructure
RA	Registration Authority
S/MIME	Secure MIME
TOGAF	The Open Group Architecture Framework
URL	Uniform Resource Locator

Appendix B : References

[CERT31]	S/MIME Version 3.1 Certificate Handling, Internet Draft, draft-ietf-smime-rfc2632bis 324
[MSG31]	S/MIME Version 3.1 Message Specification, Internet Draft, draft-ietf-smime-rfc2633bis
[Reed00]	A. Reed: Implementing Directory Services, In: McGraw-Hill (2000)
[RFC2251]	M. Wahl, T. Howes, S. Kille: RFC 2251 - Lightweight Directory Access Protocol (v3), RFC2251 (December 1997)
[RFC2560]	X.509 Internet Public Key Infrastructure - Online Certificate Status Protocol – OCSP (1999)
[RFC3280]	Internet X.509 Public Key Infrastructure - Certificate and Certificate Revocation List (CRL) Profile (2002)
[X.509]	ITU-T Recommendation X.509 (1997) ISO/IEC 9594-8:1998
[TOGAF]	The Open Group Architecture Framework : Version 8.1 Enterprise Edition. ISBN 1931624569 or http://www.opengroup.org/architecture/togaf8-doc/arch/
