Medical Informatics and the Internet in Medicine

ISSN: 1463-9238 (Print) (Online) Journal homepage: http://www.tandfonline.com/loi/imif19

An integrated architecture for deploying a virtual private medical network over the Web

S. Gritzalis, D. Gritzalis, C. Moulinos, J. Iliadis

An integrated architecture for deploying a virtual private medical network over the Web

S. GRITZALIS[†][‡]*, D. GRITZALIS[§], C. MOULINOS[§][¶], J. ILIADIS[†]

[†] Department of Information and Communication Systems, University of the Aegean Research Unit, 30 Voulgaroktonou St., Athens, GR-11472, Greece
[‡] Department of Informatics, Technological Educational Institute of Athens, Ag. Spiridonos St., Athens GR-12210, Greece
[§] Department of Informatics, Athens University of Economics and Business, 76 Patission St., Athens GR-10434, Greece
[¶] Greek Data Protection Authority, 8 Omirou St., Athens GR-10564, Greece

(*Received January 2001*)

Abstract. In this paper we describe a pilot architecture aiming at protecting Web-based medical applications through the development of a virtual private medical network. The basic technology, which is utilized by this integrated architecture, is the Trusted Third Party (TTP). In specific, a TTP is used to generate, distribute, and revoke digital certificates to/from medical practitioners and healthcare organizations wishing to communicate in a secure way. Digital certificates and digital signatures are, in particular, used to provide peer and data origin authentication and access control functionalities. We also propose a logical Public Key Infrastructure (PKI) architecture, which is robust, scalable, and based on standards. This architecture aims at supporting large-scale healthcare applications. It supports openness, scalability, flexibility and extensibility, and can be integrated with existing TTP schemes and infrastructures offering transparency and adequate security. Finally, it is demonstrated that the proposed architecture enjoys all desirable usability characteristics, and meets the set of criteria, which constitutes an applicable framework for the development of trusted medical services over the Web.

Keywords: Security; Privacy; Trusted Third Party (TTP); Public Key Infrastructure (PKI); Virtual private medical network (VPMN).

1. Introduction

The Internet and the World Wide Web ('Web') offer many advantages to organizations and to commercial and government enterprises; however, many are reluctant to use the Internet due to the security risks entailed. Since the Internet is far from being secure, large-scale web-based security infrastructures have been developed, capable of meeting the essential security requirements.

The healthcare sector is an application area that has a lot to gain from the development of a web-based security infrastructure [1,2]. The main objectives of the activities that are currently in place in this area, are:

- to increase information availability,
- to reduce costs,
- to increase the efficiency of healthcare practice,

* Author for correspondence; e-mail: sgritz@aegean.gr

Medical Informatics & The Internet in Medicine ISSN 1463-9238 print/ISSN 1464-5238 online © 2001 Taylor & Francis Ltd
http://www.tandf.co.uk/journals
DOI: 10.1080/14639230110035480

- to improve patient privacy capabilities, and
- to support new applications (tele-diagnostics, surgeries using robots, etc.).

To meet these objectives, a set of essential security requirements needs to be defined. The major requirements are confidentiality, integrity, and non-repudiation, which can be enforced by means such as ciphers, digital signatures, authentication protocols, and access control lists.

The security of a secure web-based healthcare infrastructure is mainly based on two characteristics: (a) the authenticity of the data carried, and (b) the actions performed by the participating entities. The first is achieved through digital signatures, while the second through the use of attributes (i.e. roles, access rights, authorizations, etc.).

The most well-known carrier for both the above two characteristics is the ***public key***. A ***certificate*** binds a public key value to a set of information that fully identifies the entity, known as the subject of the certificate. In order for public keys to be trusted from a vast user community, a reliable entity that is trusted by all users should verify the authentic link between an entity and its public key. These entities are known as ***Trusted Third Parties*** (*TTPs*).

A TTP is an impartial organization delivering business confidence, through commercial and technical security features, to an electronic transaction. It supplies technically and legally reliable means of carrying out, facilitating, producing independent evidence about and/or arbitrating on an electronic transaction. A TTP consists of several distributed ***Registration Authorities*** (*RAs*), and a ***Certification Authority*** (*CA*). RAs handle identity verification material, so that a certificate can be issued to a user, and issue certificate requests, on behalf of the user. The CAs can issue and revoke a certificate. The degree to which a user can trust the assurance offered by a certificate depends on several factors (e.g. the practices followed by the RA in authenticating the subject, the CA's operating policy, the subject obligations, warranties, limitations on liability, etc.). A large-scale deployment of the public key cryptography requires multiple TTPs. A ***Public Key Infrastructure*** (*PKI*) consists of one or several TTPs, and refers to an infrastructure that is used to issue and revoke public keys and public keys certificates.

In this paper, we present a TTP framework capable of providing the essential security services. We also propose a logical PKI architecture, which is robust, scalable, and based on standards. The architecture aims at supporting large-scale healthcare applications.

The paper is organized as follows: In section 2 the security requirements for medical applications are presented, while in section 3 a TTP framework which guarantees the provision of robust security services is described. In section 4, an integrated PKI architecture for deploying trusted medical services, involving entities that have registered in different TTPs, is presented, while section 5 includes the results of the work.

2. Security requirements for medical applications

2.1. *Web security issues*

The main requirements of a secure web-enabled application are [3]:

- security of the Web server and the data residing on it,

- security of the information that is transmitted between the Web server and the user, and
- security of the computer used by the user.

Moreover, we have to address a number of additional issues, such as:

- the identification and authentication between communicating parties,
- the patient privacy,
- the information integrity, and
- the logging and auditing of information about the transaction.

2.2. *Security threats*

The Web has not been designed with security in mind. Therefore, any application using the Web and the Internet as navigation and communication tools is vulnerable to specific threats. These threats may jeopardize the functionality of a medical application operating over the Internet. The most important threats that have to be dealt with, in order to establish a secure medical system over the Internet, are presented below.

1. *Monitoring of communication lines*: By monitoring communication lines wiretappers may gain unauthorized access to medical data, thereby violating a patient privacy.
2. *Shared key guessing*: If shared keys are used in order to encrypt communicated medical data one may attempt and succeed in guessing those keys. Knowledge of the keys can lead to the disclosure of a patient medical data.
3. *Shared key stealing*: If the shared keys used for encrypting communicated medical data are transmitted in cleartext, or if the protocol used for the exchange of these keys is not robust, then a third party may steal these keys and gain access to a patient medical data.
4. *Unauthorized modification of information in transit*: Medical records may be modified on their course to their recipient. Modification may be performed in such a way that the receiving entity will not be aware of it.
5. *Forged network addresses*: A healthcare organization considers received medical data as valid if they are sent by another healthcare organization. In this case an unauthorized party may transmit medical data to the first organization that will be accepted as valid, by forging the proper network address.
6. *Masquerade*: A malicious user may masquerade the identity of a web site to that of a valid medical site.
7. *Password stealing*: If the passwords are transmitted in cleartext form are used to authenticate medical personnel, a third party may steal these passwords and impersonate authorized medical personnel.
8. *Unauthorized access*: Unauthorized access from invalid users of a medical system may cause the storage of false, corrupted, or modified data, resulting in the false diagnosis of a patient.
9. *Repudiation of origin*: A malicious user may forge his/her authentication credentials to gain unauthorized access to a server holding medical data. Moreover, a valid user may maliciously modify medical data and repudiate his/her actions at a later stage.
10. *Private key stealing*: If a malicious user steals the private key of a valid user of a medical organization, he/she can impersonate him/her. This user may

proceed in digitally signing false or illegally modified medical records thereby validating them.

11. *Private key compromise*: If another user compromises the private key of a valid medical user, the later can make use of that private key to digitally sign false or illegally modified patient records or other medical data, thereby validating them.

2.3. *User requirements*

The need for a medical infrastructure capable of exchanging information leads to the use of public networks. The Internet can undertake that role because it provides a worldwide communication infrastructure, which is available to the global medical community at a low cost. Moreover, the Web can serve as a transport mechanism and navigation tool for audio-visual medical data stored and communicated between geographically distant medical organizations and healthcare professionals. These data can take virtually any form, from plain text to x-rays and other medical examinations requiring the storage of visual or audio components. Such an infrastructure must provide the means for communication, exchange of information and co-operation, regardless of the underlying computing equipment.

According to the Council of Europe (CoE) Recommendation on the Protection of Medical Data [4], '...Appropriate technical and organizational measures shall be taken to protect Personal Data processed in accordance with this recommendation against accidental or illegal destructure, accidental loss as well as against unauthorized access, alteration, communication or any other form of processing. Such measures shall ensure an appropriate level of security taking into account, on the one hand, of the technical state-of-the-art and, on the other hand, of the sensitive nature of medical data and the evaluation of potential risks'.

Taking into consideration this recommendation, a medical network has to be designed and implemented with security in mind. This imposes a set of security requirements which have to be met by a telemedical application; especially in the case where it is deployed in wide scale and uses public networks. It should be stressed that security should not act as a restrictive factor towards the normal operation of healthcare organizations and professionals in such a network. This is possible because most of the security mechanisms are put in force transparently to the end-user (i.e. the user is not aware of the fact that a security service is provided), so effective and robust security mechanisms should be viewed as an enabler for these operations. On the other hand, security and performance are orthogonal attributes; therefore an upgrade of security may eventually lead to some downgrade of performance.

The development of a virtual private medical network (VPMN) should meet the following security requirements.

- *Confidentiality of medical data*: Medical data should be disclosed to authorized personnel, only.
- *Integrity of medical data*: Mechanisms should be put in place in order to prevent medical data from being modified by unauthorized parties.
- *Transparency of security mechanisms*: Security mechanisms should not produce too much overhead in the normal operation of healthcare organizations and professionals. The security mechanisms should operate as transparently as possible.

- *Provision of interoperability mechanisms*: A security infrastructure, providing the above mentioned security mechanisms to the medical community, must also provide interoperability.
- *Outsourcing the operation/maintenance of the security infrastructure*: Medical staff are often not supported by security experts. For this reason, the medical community should not be expected to maintain or operate the security infrastructure by itself. This infrastructure may be operated and/or maintained by third parties with the essential security expertise. These parties should be responsible for the deployment, operation, and maintenance of the security infrastructure as a whole.

3. Deploying a framework for the provision of security services

3.1. *Trusted Third Party services*

The services provided by a TTP will be described in this section. These services can provide the medical community with the means for implementing a VPMN [5].

3.1.1. Electronic registration. A healthcare professional wishing to communicate with other healthcare professionals or medical organizations should register with a TTP. At first, he/she submits his/her registration request to the RA. The latter will verify, via out-of-band mechanisms, the identification data included in the registration request. If it is valid, it will forward the completed registration form to CA, which will proceed with the issuance of a certificate for that entity. Finally, the CA will generate the certificate and communicate it to the healthcare professional that requested it. When the healthcare professional confirms that it has received and installed a valid certificate, the CA will store that certificate in the Directory.

3.1.2. Initialization. This function covers the session between the requesting entity and the CA. It allows the requester and the CA to technically synchronize their cryptographic environments (i.e. both parties interchange information about the selected cryptographic algorithms, cryptographic protocols, certification and key exchange protocols in order to establish a secure communication channel), and comprises the Secure Session Layer (SSL) handshake [6]. Essentially, the CA public certificate is sent to the client, in order for it to be able to send, in encrypted form, the symmetric key and the algorithm selected for that session. Furthermore, this function entails the appropriate actions for the entity's authentication process. When the authentication process succeeds, the authorization authority invites a CA to issue the certificate.

3.1.3. Key personalization, generation, and repository. Key personalization is the process of associating a key-pair with the registered name of a healthcare professional. It is possible that the key-pair is created in a transparent way for the healthcare professional. However, according to Directive 1999/93/EC [7], the procedure is described as follows: A user asks a TTP to generate a key-pair; the CA generates the key-pair and sends to the healthcare professional his/her secret key, via out of band methods; a user may decide that he/she wishes the TTP to keep backup of his/her secret key for key recovery purposes.

3.1.4. Naming. The naming of the healthcare professionals and the medical organizations is performed in accordance with the X.500 specifications [8], in order to provide the means to identify them uniquely, without depending on the identification methods used by the various medical organizations.

3.1.5. Certificates: Structure, Generation, Distribution, Storage, and Retrieval. The certificates of the healthcare professionals and the medical organizations should be digitally signed using the private key of the CA, in order to achieve integrity (i.e. non-modification) of the message and authenticity of the CA. The latter sends the issued certificates to the Directory, and keeps a backup copy at a local repository. The healthcare professional is notified that he/she may download his/her digitally signed certificate from the Web server of the TTP or the Directory; he/she can even ask to have the certificate sent by secure e-mail using e.g. S/MIME (Secure Multipurpose Internet Mail Extensions) technology [9]. It is useful for the CAs to perform management functions on the certificates it generates. In order to provide these services (e.g. notifying a user when a certificate is about to expire or revoking certificates) the CA will use the local repository to store and retrieve the certificates it generates. Expired or revoked certificates should be removed from the Directory.

3.1.6. Auditing. To provide additional assurance of the trusted nature of TTP and to provide information to agency personnel conducting internal audits, the actions of each CA and RA should be audited. In general, internal auditing procedures are performed for internal purposes of the TTP itself (either by its own resources, or by consulting experts) and are essential for the operation and progress of the healthcare organization. External auditing is conducted for the needs of external organizations or bodies, which set regulations for the TTP and PKI operation and/or supervise the compliance with them. These bodies hold the responsibility for checking whether the TTP and PKI are implementing the appropriate actions, in order to achieve conformance with existing standards [7]. Audit records and audit trails are generated for events such as user registration, certificate request and receipt, compromised key reports, etc.

3.1.7. Certificate Directory Management. According to ISO and CCITT [8], the Directory acts as a distributed repository of identification and authentication data, such as the certificates of healthcare professionals. Servers consult the Directory to retrieve the latest version of the Certificate Revocation Lists (CRLs) and the identification data for a user that is attempting to access them. Such a repository may be implemented in many ways (e.g. using ftp servers, mail servers, Web servers or X.500 and Lightweight Directory Access Protocol (LDAP) [10] schemes, etc.) An efficient way to implement a Directory for the storage and retrieval of identification and authentication data of healthcare professionals and medical organizations would be an X.500-based Directory, with a LDAP front-end. LDAP is a streamlined and simplified version of the Open Systems Interconnection (OSI) Directory Access Protocol (DAP) [8] that is used to access X.500 Directory services. This scheme provides the administrators of the VPMN with the ability to store and retrieve medical data, certificates and CRLs. The communication with the Directory can be performed through LDAP protocol over SSL [6]. The de facto standard SSL protocol is an effective security protocol for the Internet and the Web. The structure and operation of the Directory can be independent of the Certification and Trust

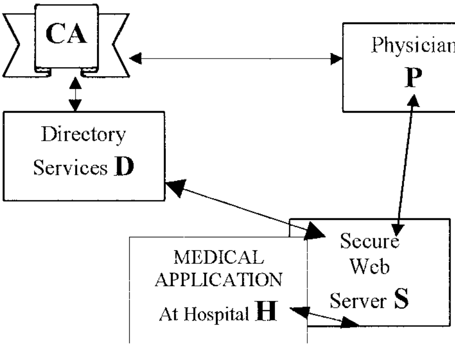

Figure 1. TTP infrastructure.

scheme (e.g. hierarchy, cross-certification). The design of the internal structure of the Directory depends on the medical organizations involved in a VPMN, their internal organizational schemes, and on the independent healthcare professionals communicating within such a network.

3.1.8 CRLs: Structure, Generation and Maintenance, Distribution, Storage, and Retrieval. Certificate Revocation Lists (CRLs) are lists that contain the certificates that have been revoked [11]. They include information such as the CRL issuer's identifier, the serial numbers of the revoked certificates, and the date each certificate that was revoked. The CRL is digitally signed with the private key of the CA that issued the specific CRL. The CRL is published in the Directory so everyone can access it. The CA has to ensure that the information within the CRL is as current as possible.

3.1.9. Time stamping services. In many cases, time and data stamps must be affixed to the documentation, in order to denote when the documentation was received or sent. Certification information is always time-stamped. If the documentation is generated by and sent via electronic means, then the data and time-stamp must also be generated and affixed to the document electronically. Moreover, medical data may be required to be timestamped too. If this is the case, the private keys of the entities, which communicate such data, can be used to timestamp the medical data before transmitting them to another entity.

3.1.10. Integrity of the root public and private keys. The integrity of the root private key is guaranteed by the rigid security measures put in place in the TTP. Physical security and IT security mechanisms should be deployed in the TTP organization to ensure that the root private key is secure. Furthermore, the integrity of the root public key and root public certificate, when communicated to entities outside the TTP, is ensured by secure communication protocols.

3.2. *Technical infrastructure*

The core modules of a TTP security scheme are as follows (figure 1):

- *Directory services.* They act as repositories of registration, identification, and authentication information of the entities participating in the Security

Architecture (e.g. healthcare professionals, servers of medical organizations). The LDAP protocol could be preferred for the implementation of these Directory Services.

- *Certificate servers.* They provide the X.509v3 certificates [12] and thus validate the digital signatures of the healthcare professionals and the servers of the medical organizations involved in the VPMN. SSL can be used for the establishment of secure communications between these entities and S/MIME (Secure Multipurpose Internet Mail Extensions) [9] can be used to establish secure communication between healthcare professionals. The produced certificates and CRL should be stored on a local RDBMS (Relational DataBase Management System) and in the aforementioned Directory, in order to be accessed by the medical entities for verification purposes. The Certification Services should comply with the relevant standards (e.g. X.509v3, SSLv3, LDAP and PKCS).
- *Secure Web servers.* They operate as platforms for the storage and retrieval of medical data, and for the execution of the web-enabled medical applications. These SSL-enabled Web Servers can either host an entire medical application, or provide a Web front-end for a standalone medical application, operating at a local level. Additionally, these Web Servers can be used to store medical data, or as a front-end for accessing medical data maintained in a local database.

The specific scheme that is proposed for the implementation of a VPMN is based on open specifications and standards. Therefore, interoperability with other secure medical schemes, based on similar, open architectures, can be achieved at low cost.

3.3. *Confronting with security threats*

The functions of a VPMN render it capable of confronting successfully with the threats presented in section 2.2. In this section, the way these threats can be averted to, will be discussed in some detail.

1. *Monitoring of communication lines*: The medical data communicated between a healthcare professional and a medical organization or another healthcare professional are encrypted with the use of shared session keys. These keys are randomly generated in the beginning of every communication session and used for the encryption of that session only.
2. *Shared key guessing*: Tackling this threat involves the use of substantially large keys and cryptologically secure random number generators. After the end of each communication session, the keys are discarded. Random number generation algorithms ensure that the same-shared session key will not be used again in another communication session.
3. *Shared key stealing*: The software used by the healthcare professional and medical organizations (Web browsers, Web servers) encrypts the randomly generated session keys before communicating them. Encryption of the shared session keys is performed by the use of asymmetric algorithms, which use the public keys of the healthcare professional and the medical organization involved in their respective certificates.
4. *Unauthorized modification of information in transit*: The integrity of medical data communicated over the VPMN is supported by the use of Message

Authentication Codes and secure hashing algorithms [13, 14]. These algorithms can use the private key of the entity that transmits data; therefore the receiving entity can verify the integrity of data received.

5. *Forged network addresses*: Existing protocols [15] aim to avoid the forging of a network address. Until these protocols are tested in depth, one can use the X.509v3 certificate of the communicating entity in order to establish the origin of communication.
6. *Masquerade*: Verification of the identity of communicating end-entities can be performed by the exchange of X.509v3 certificates. The validity of these certificates can be verified against the Directory maintained by the TTP that has issued the certificate.
7. *Password stealing*: The use of certificates and shared session keys for authenticating healthcare professionals and medical organizations, as well as for encrypting the communicated medical data, limits the use of passwords to a minimum. However, if passwords are used they are transmitted encrypted, using the shared communication session keys.
8. *Unauthorized access*. Access to the medical resources of an organization is controlled independently by that organization. The means, which the latter can use in order to authenticate healthcare professionals and grant them the appropriate access rights to the resources, are the certificates owned by these entities.
9. *Repudiation of origin*: Any entity communicating with another can make use of the authentication credentials presented to verify the identity of the communicating entity. These credentials should be logged for future reference [16].
10. *Private key stealing*: The private keys of both the healthcare professionals and of those corresponding to servers, hosting medical data and maintained by medical organizations should be kept encrypted while kept in the storage medium. The optimum solution for the protection of the private keys is to use tamper-resistant smart cards in order to store them.
11. *Private key compromise*: A malicious user holding a compromised private key may impersonate the entity that has the private key. It is imperative to inform the TTP after a private key is compromised, in order to add the respective certificate in the CRL. The CRL is digitally signed by the TTP; therefore anyone can verify which keys have been compromised.

3.4. *Operational, organizational and legal issues*

The infrastructure required by a healthcare professional, in order to access medical applications securely, is merely a Web browser, an Internet connection, and a registration to the CA. Furthermore, the Web Server of a medical organization need only support SSL and it has to be given a certificate from a TTP like the one described in this paper. Therefore, the upgrade in existing equipment is minor, compared to the advantages offered by the provision of a VPMN, able to operate on a global scale.

The internal structure of the Directory, as well as the Certification scheme used, are issues that are of interest to the implementers of a VPMN. Deciding upon these issues depends mainly on the structure and the number of medical organizations and healthcare professionals involved.

The relevant legal framework and the healthcare codes of practice reflect on specific requirements for the development of secure web-enabled medical application. However, in order to exploit the security services offered by a VPMN, implementing them would not be enough. The medical personnel have to be eager to familiarize them with using security services. The friendliness of the proposed framework is achieved through the transparent way the proposed solution meets the security requirements. The end-users need only to acquire a certificate from the TTP, and use the certificates presented to them by other healthcare professionals or medical organizations, that should they wish to identify the later. Moreover, the medical organizations need not make radical changes in their access control methods, at their Web Servers and medical databases. Access control remains the responsibility of individual medical organizations and it is up to them to decide on the way they will implement it. The usage of certificates in order to identify and authenticate entities and then grant them with their respective rights is a procedure that the medical organizations should implement.

The proposed security scheme is characterized by scalability. Certificates can be granted to any healthcare professional and medical organization requesting them. The deployed Directory can be managed to provide new user groups and organizational units in order to store the certificates and identification information of new entities wishing to join the VPMN.

Protection of the medical data is governed by the European Convention on Human Rights [17], and by the Recommendation on the Protection of Medical Data [4]. The proposed framework for developing a VPMN makes extensive use of TTP infrastructures and digital signature schemes. The legal recognition of the digital signature concept is emerging in most of the European Union (EU) Member States [7], as well as in other countries (e.g. USA, Canada, etc.). It is expected that the legal recognition process will be completed in the next couple of years. In the case of the EU, TTPs operation should comply with the requirements set forth by the EU Data Protection Directive. Finally, a security policy should be developed and put in force for the healthcare organizations, in the context of the upcoming ISO 17799 [18].

3.5. *Issues remain to be solved*

TTP supplies reliable means for carrying out, facilitating, and producing independent evidence about and/or arbitrating on an electronic transaction. Its services are provided and underwritten by technical, legal, financial and/or structural means. In a global Public Key Infrastructure it is unlikely that all users will be connected to a unique TTP. The cases where the involvement of more than one TTP is necessary in a transaction occur very often. Such a situation will be faced:

- when the transacting parties do not belong to the same geographical or national domain, or to the same healthcare sector,
- when a user requests a service that his/her home TTP does not support and requires the communication with another TTP that provides such a service,
- when multiple concurrent users require the provision of a service. Although high-speed connections are available, vast concurrent transactions are not efficiently served in practice by TTP products, and/or
- when mobile users using different entry points and acting under different, and in many cases incompatible, underlying technologies, compatibility problems may be faced.

Due to the above reasons, a web of TTPs may be established. This set of TTPs is connected through chains of trust (usually called certificate paths), in order to provide a web of trust, called a Public Key Infrastructure (PKI). Furthermore, different users belonging to different TTPs should enjoy the same interface when they request a service from the PKI. As a result, an integrated PKI architecture is needed. This architecture is outlined in the next section.

4. A PKI architecture for deploying trusted medical services

4.1. *The reference model*

In this section, a unified, extensible, scaleable, robust and flexible PKI-based architecture will be described. This architecture is based on standards and is useful across different healthcare application domains. The architecture is addressed at two levels of abstraction: (a) the reference model and (b) the functional architecture. We call it the *KEYSTONE* architecture, and it includes users, TTPs, and other elements (e.g. application programs, managers, etc.). An overview of the abstract model is given in the following figure 2.

The organizational entities are ***users*** and *TTPs*. Inside TTPs, activities must be performed in order for a TTP to be able to provide trust services meeting specific healthcare user requirements. These activities can be clustered in ***roles***. These roles can be defined as integrated actions performing specific and well-defined tasks, aiming at providing trust services in open distributed systems.

A user may use trust services offered by one TTP operating in his/her domain. TTPs in different domains should be able to interact with other TTPs. One or many users may exist in every security domain or healthcare sector. One or more roles may also exist within every TTP.

The users can be healthcare employees, patients, or application programs. There is a specific interface for every communication channel of the proposed architecture:

- TTP-User interface (communication between the TTP and the user).
- User-TTP interface (communication between user and the TTP).
- TTP-TTP interface (communication among different TTPs). This interface is important, in order to allow TTPs to provide services to one another, supporting the operation of a scalable PKI.

The proposed architecture evolves unlimited number of TTPs communicating with each other using the TTP-TTP interface. Technology incompatibilities are hidden within the implementation of each interface. As a result, technology changes will not influence the overall architecture.

4.2. *The functional architecture*

The *KEYSTONE* functional architecture describes the TTP information system, in terms of functional units interacting across clearly defined interfaces. The list of functional units is presented along with the description of individual units, their interfaces, and the overall picture of information processing in the TTP information system.

4.2.1. Functional architecture elements. The functional architecture is made up of a number of ***functional units***, each one performing a specific task within the TTP. The functional unit definition specifies the functionality provided without being

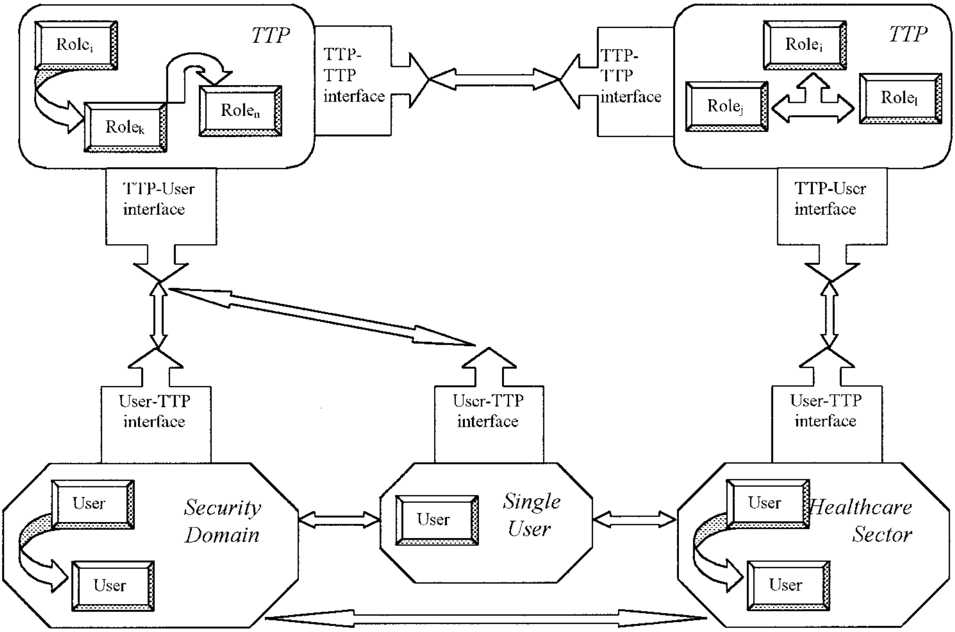

Figure 2. A reference model for PKI.

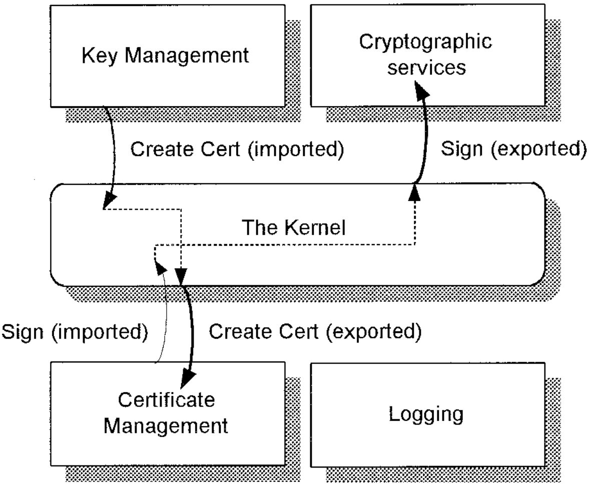

Figure 3. Kernel as a bus for abstract primitives.

tied to any particular technology. The functional units provide services to one another by means of abstract primitives.

In order to simplify the management of the functional units, all importing and exporting of abstract primitives is from and to the ***kernel*** (figure 3). The use of a central kernel allows the enhancement of new functional units, and thus facilitating the provision of new services by the TTP. By forbidding any exchange of abstract primitives, except via the kernel, the functional architecture divorces each functional unit from the internal details of the other. This makes the TTPs easier to develop and maintain, and also facilitates distribution of TTP functionality, if required.

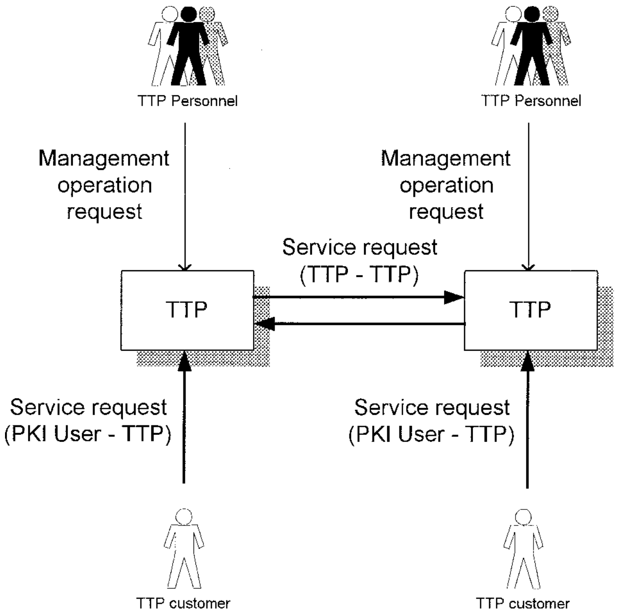
TTP Personnel
TTP Personnel
Management operation request
Management operation request
Service request (TTP - TTP)
TTP
TTP
Service request (PKI User - TTP)
Service request (PKI User - TTP)
TTP customer
TTP customer

Figure 4. TTP and its users.

Such architecture is suited to implementation using an object management technology, such as CORBA [19, 20].

Since the functional unit is not tied to any particular technology, any technology that provides the functions defined for the functional unit may be used. In the event of a new technology appearing, such as a new encryption algorithm or a message digest technique, this may be added to the TTP without any changes to any functional unit other than the one concerned. Thus, the functional unit acts as a gateway between a particular technology and a set of functions, which are required by the TTP.

4.2.2. TTP and its users. In order to design the *KEYSTONE* functional architecture, the relationship between the TTP information system and its users has to be defined. This relationship is schematically depicted in figure 4. Each TTP information system deals with three different kinds of external entities: ***TTP personnel***, ***TTP customer***, and ***other TTPs***. There are two fundamental types of interaction:

- ***service provision***, which takes place between a TTP and the TTP customers or between two TTPs. It is an interaction where the TTP offers the requested service to a TTP customer or to another TTP for a certain reward
- ***management operation***, which takes place between a TTP and the TTP personnel. Initiated by the TTP personnel, a management operation modifies a certain aspect of TTP behaviour.

Service provision or management operation is initiated by a ***service request*** or a ***management operation request*** respectively; these are issued by external entities. The purpose of a service request is similar to a paper based order form. It facilitates service provision and has to carry information present in usual order forms:

- service description,
- service delivery method description,
- payment method description,
- invoice and receipt delivery method description,
- date and time, and
- requester's identification and digital signature.

The format of the service request must be standardized to facilitate interoperability. The rules must be defined to convert service requests from different formats into the *KEYSTONE* service request format. The rest of the commands are given to the TTP information system in the form of management operation requests.

4.2.3. Functional architecture overview. From the information processing point of view, the *KEYSTONE* TTP functions fall into six groups:

1. *Managerial functions* – all decisions that can be taken by TTP personnel (e.g. application form validation, operations manual specification, etc.)
2. *Management access* – all technical means of providing TTP personnel with controlled access to the TTP configuration (processing of management operation requests).
3. *Customer access* – all technical means of providing TTP customers and other TTPs with controlled access t the trusted services (processing of service requests).
4. *Trusted services* – all technical functions implementing trusted services (e.g. certification path validation, time stamp generation, etc.)
5. *Localized supporting services* – technical functions locally implementing basic mechanisms required by the rest of the TTP (e.g. encryption, archiving, database management, etc.).
6. *Infrastructure supporting services* – technical functions providing access to distributed network services essential for TTP functioning (e.g. directory service access, access to Visa/MasterCard secure electronic transaction infrastructure, etc.).

The inter-relations between the functional groups are schematically depicted in figure 5.

4.2.4. Analysis of functional units. Service requests are first processed by the customer access functions that perform customer authentication, access rights control, and payment. Following this, the appropriate functional unit performs the requested service. Similarly, management operation requests are first processed by the management access functions that perform request originator authentication and access rights control. Following this, the corresponding functional unit performs the requested management operation. Supporting services, such as cryptographic computations, database management, and logging can be used in every step of the user request processing.

The *KEYSTONE* functional architecture splits the functional groups of figure 5 into interrelated functional units, and specifies interfaces between them. The logical view of the *KEYSTONE* functional architecture (the kernel is not shown) is presented in figure 6. The Customer Access group is divided into the ***Secure customer access***, ***Customer access rights control***, ***payment***, and ***Customers' accounts*** functional units.

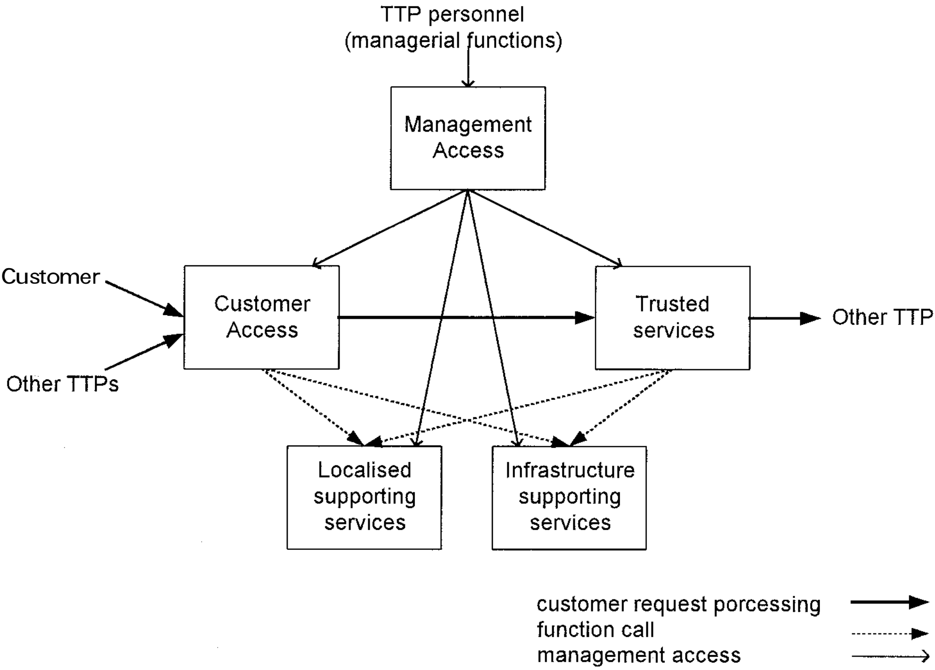

Figure 5. Functional groups and their interrelation.

- The *Secure customer access* functional unit is responsible for establishing the secure communication between the TTP and the TTP customers over an insecure network. It provides service request *authentication*, data exchange *integrity* and *confidentiality*.
- The *Customer access rights control* functional unit verifies that the customer has the right to use the requested type of service. The decision is made on the basis of the access control information stored in the customer's account.
- The *Payment* functional unit is responsible for checking whether the customer has to pay for the requested service or not and if payment is required for charging the customer's account or for initiating an electronic payment transaction.
- The *Customer's accounts* functional unit maintains the database of customers' accounts. A customer's account holds the entire information about the customer.

The Secure Customer Access, the Customer Access Rights Control, and the Payment functional unit sequentially process the service request. Then, the kernel to the appropriate functional unit dispatches the request.

The Management Access group is divided into three functional units: ***Secure management access***, ***Management access rights control***, and ***Management accounts***.

- The *Secure management access* functional unit is responsible for establishing the secure communication between the TTP and the TTP personnel over an insecure network. It provides management operation request *authentication*, data exchange *integrity* and *confidentiality*.

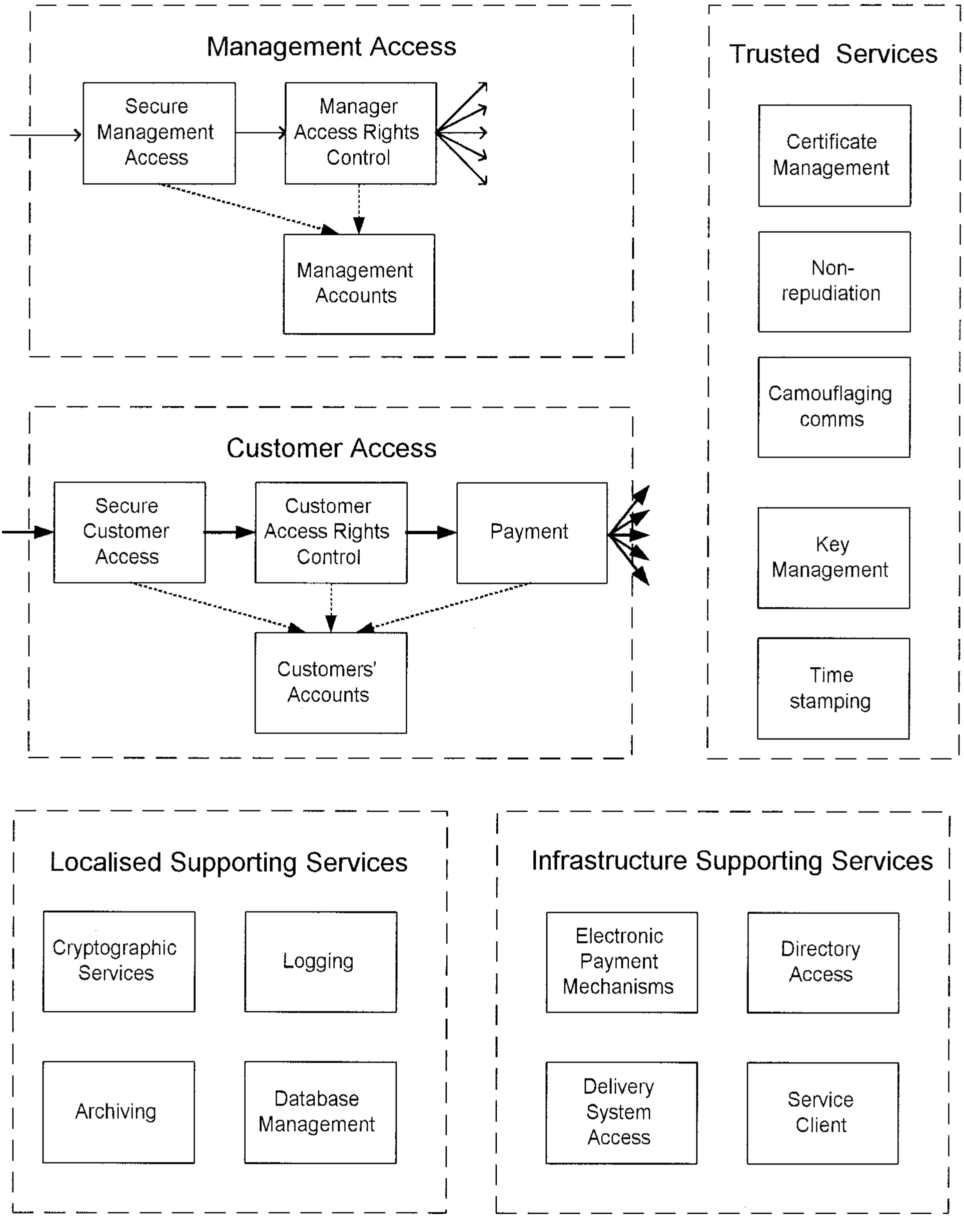

Figure 6. *KEYSTONE* functional architecture (the kernel is not shown).

- The *Management access rights control* functional unit verifies that the particular member of TTP personnel has the right to initiate the requested management operation. The decision is made on the basis of the access control information stored in his/her management account.
- The *Management accounts* functional unit maintains the database that holds information about the members of the TTP personnel necessary to allow them to perform their managerial functions. Each record stores identification information, access rights, etc.

The Trusted Services functional group is divided into five functional units, each one implementing a particular type of trusted service:

- The *Certificate management* functional unit supports the certificate management service. It performs certificate generation, distribution, storage and retrieval, and revocation.
- The *Key management* functional unit supports the key management service. It performs functions such as key generation, personalization, distribution of keys, key storage, retrieval, recovery, etc.
- The *Non-repudiation* functional unit supports the non-repudiation service. It performs functions such as generation of records about events, storage of these records, and presenting these records for dispute resolution.
- The *Time-stamping* functional unit supports the time-stamping service. It performs retrieval of the time/date data for the time-stamp, link of time-stamps to a message, verification of the validity of the time-stamp certificate, maintenance of a database of time-stamp certificates, maintenance of a log of time-stamping authority activity, etc.
- The *Camouflaging communications* functional unit supports the camouflaging communications service. It takes the message submitted for camouflaged transmission and passes it through the network of camouflaging TTPs to its destination. Onion routing and data flow concealing are used to provide camouflaging.

This service-per-unit approach used in the Trusted Services group simplifies addition and removal of support for trusted services by actual TTPs.

The remaining two functional groups (Localized supporting services and Infrastructure supporting services) provide general-purpose functionality used in trusted services, and access related functional units.

The Localized Supporting Services group is split into four functional units: ***Cryptographic services***, ***Logging***, ***Archiving***, and ***Database management***.

- The *Cryptographic services* functional unit provides various functions that perform cryptographic computations on a block of data.
- The *Logging* functional unit provides functions for creating and subsequent analysis of various logs of events.
- The *Archiving* functional unit provides access to archiving facilities. Its major functions are to save a named block of data in the archive, and to restore it from the archive.
- The *Database management* functional unit provides functionality necessary for creating and maintaining various databases.

The *Infrastructure supporting services* functional group is divided into four functional units: ***Electronic payment access***, ***Directory service access***, ***Delivery system access***, and ***Service client***:

- The *Electronic payment mechanisms* functional unit supports various on-line payment mechanisms. Its major function is to accept payments for services from TTP clients and other TTPs. It has to maintain databases of certificates and be registered with various on-line payment systems such as Visa/MasterCard SET, MONDEX, etc.
- The *Directory access* functional unit provides access to various Directory services and on-line databases. Its functions are to retrieve a known object

from a Directory or an on-line database, to search a Directory or an on-line database for specific objects, and to provide access to the record(s) about the TTP in the Directory service(s) and/or on-line database(s).

- The *Delivery system access* functional unit provides access to various delivery services such as e-mail, electronic file transfer, fax, and postal mail. These services are used for key distribution, camouflaged message delivery, etc. The major function of this functional unit is to transmit a block of data to a specified destination using a specified means of transport.
- The *Service client* functional unit provides other functional units (primarily trusted services) with a mechanism to request trusted services from other TTPs. This is essential for camouflaged communications, disclosing a newly generated key to the governmental key, etc.

4.2.5. Technology evaluation. In order to ensure TTP to TTP interoperation, standards must be adopted at the TTP operation level, as well at the TTP to TTP interconnection level. Standard status gives a technology additional competing advantage on the market. Such technologies are considered to be particularly promising candidates. Several PKI related international standards have been reviewed in [21]. The applicability of PKI standards for different *KEYSTONE* PKI services is summarized in table 1.

A review of PKI related standards demonstrated that many PKI services are covered with standards of some form. There are international standards dedicated to encrypted communications, digital signatures, certificates, non-repudiation services, key management, TTP security assurance and TTP management. Other PKI services or their elements are discussed as parts of large framework standards or as supporting services for other PKI services.

Every service referred to table 1 corresponds to one or more functional units through its supporting functions. In order to implement these functional units, available technologies have been evaluated. The rest of this section evaluates available technologies in each functional area in the *KEYSTONE* functional architecture and highlights the most promising candidates in each functional area.

We understand *functional area* as a group of functional units that are based on the same technologies. The *KEYSTONE* functional architecture consists of fifteen functional areas:

- *System interconnect* (the kernel)
- *Secure customer and management access* (Secure customer access and Secure management access, and Service client functional units)
- *Access rights and payment control* (Customer access rights control, Payment, Management access rights control functional units)
- *Certificates* (Certificate management functional unit)
- *Key management* (Key management functional unit)
- *Time-stamping* (time-stamping functional unit)
- *Non-repudiation* (Non-repudiation functional unit)
- *Camouflaging communications* (Camouflaging communications functional unit)
- *Cryptographic services* (Cryptographic services functional unit)
- *Database management* (Database management functional unit)
- *Archiving* (Archiving functional unit)

Table 1. Standards and *KEYSTONE* PKI services.

	KEYSTONE PKI services														
Standards and specifications	Registration	Digital sign	Encryption	Time stamping	Non-repudiation	Key-management	Certificates	Inf Repository	Directory	Camouflaging	Authorisation	Audit	Assurance	Customer	TTP to TTP
ISO 9594|ITU-T X.500 (The Directory)	+	+				+	+	+	+		+			+	+
ISO 9796 (Digital signature giving message recovery)		+	+												
ISO 14888 (Digital signatures with appendix)		+													
ISO 11770 (Key management)						+	+								
ISO 9798 (Entity authentication)		+		+							+				
ISO 13888 (Non repudiation framework)					+			+							
ISO 13335 (Management of IT security)QN											+	+			
ISO 15416 (Management of TTPs)												+	+		+
ISO 15408 (Evaluation criteria for IT security)												+	+		
ISO TR-13569 (Information security guidelines in banking)								+				+	+		
ISO 9735 (EDIFACT)			+												
ENV 12388 (Algorithm for digital signature services)		+													
PKIX						+	+								+
SPKI						+	+				+				+
LDAP	+								+						
WHOIS++	+								+						
PEM			+			+	+								+
S/MIME			+												
SSL			+												
IPSEC			+												
DNSSEC	+						+	+	+						
GSS-API		+	+			+									
Microsoft CryptoAPI		+	+			+									
GCS-API		+	+			+									
CORBA Security Services		+	+		+	+					+	+			
RSA PKCS	+	+			+	+									
SET			+		+										

+ means that the standard describes the mechanisms and/or guidelines that implement the service on their own or in co-operation with other mechanisms and/or guidelines.

Table 2. *KEYSTONE* technology profile.

Functional area	Most promising candidate technologies
System interconnection	● CORBA
Secure management and customer access	● WWW+SSL, ● SSH+custom application ● Secure email (S/MIME, PGP) ● Postal mail
Access rights and payment control	● Policy Maker or a proprietary system
Certificates	● X.509 ● SPKI
Key management	● ISO 8732 ● ISO 11770 ● PKCS
Time stamping	● U.S. patent 5,136,647, ● Annex to ISO 13888-3
Non-repudiation	● ISO 13888 ● CORBA Non-repudiation service
Camouflaging communications	● Onion routing ● Traffic padding
Cryptographic services	● RSA Cryptoki ● Microsoft CryptoAPI ● Open Group's GCS-API
Database management systems	● Medium-class or high-end DBMS supporting SQL
Archiving	● Unix tar ● IBM ADSM
Logging	● Unix logging ● CORBA Audit service
Electronic payment mechanisms	● SET ● ECash ● Mondex
Directory access	● X.500/LDAP ● Z.39.50
Delivery	● Secure e-mail (S/MIME, PGP) ● Postal mail

- *Logging* (Logging functional unit)
- *Electronic payment mechanisms* (Electronic payment mechanisms functional unit)
- *Directory access* (Directory access functional unit)
- *Delivery* (Delivery system access functional unit)

There are two more functional units that are based on services provided by other functional units and do not require additional technology, the Customers' accounts and the Management accounts functional units.

- *A precise specification* of service provision or data processing, which defines data formats, algorithms, protocols, etc. There may be a software or hardware implementation of the specification, which can be used as a building block for a TTP.
- *A widely accepted scenario/approach* for service provision or data processing. In the absence of a precise specification, this can be used in as a guideline for designing a good proprietary solution.

The technology profile presented in table 2 lists the most promising candidate technologies on a per-functional area basis. Of course, the above technology profile is not fixed. When a new suitable technology appears, it can be added to it.

4.3. *The* KEYSTONE *abstract PKI architecture*

The main characteristics of the proposed *KEYSTONE* architecture are the following:

1. *Openness*: Open data networks should be employed as the carriage for the *KEYSTONE* PKI instead of closed proprietary Value Added Networks (VANs). Despite its security risks, the Internet is a good candidate for the *KEYSTONE* PKI. VANs are costly and base security on their closed nature rather than on their robust security mechanisms. In addition, there is significant progress in solving Internet security problems (e.g. IPv6). Furthermore, the proposed infrastructure is based on well-established standards that have been implemented over the Internet (i.e. HTTP, SSL, S/MIME, LDAP, X.509, X500 etc.). As a result it can integrate other non-compliant PKIs due to the use of protocol gateways that should be established at the intersection points of the PKI and the proprietary PKIs.
2. *Scalability*: An evolution from small–medium scale organizational networks to large scale internets is witnessed. The proposed infrastructure has been designed in such a way that it provides the same as the information and computing resources grow and become distributed.
3. *Flexibility-extensibility*: As new technologies are adopted and old ones become obsolete, a PKI must be capable of easily merging any changing to the corporate infrastructure. The TTP internal structure has been designed with Object Orientation principles in mind. As a result, modularity and logical independence between functional units has been achieved.
4. *Integration* with *existing* information and technological *infrastructure*: The architecture has been designed to coexist with pre-established technological solutions. The CORBA (its security enhanced version) kernel guarantees that legacy infrastructure (software code as well as legacy hardware) can interoperate with the *KEYSTONE* PKI.
5. *Transparency*. The proposed model enables users that use specific software or hardware platforms to transparently interact with users that use different and non-compliant with their equipment. Platform independence is achieved through the use of CORBA at the kernel level. CORBA offers mechanisms that platform dependent code can wrap and exported through higher-level common interfaces.
6. *Security* of reserved information. The proposed architecture has adopted the relevant state-of-the-art services, functions, mechanisms and standards.

The proposed abstract *KEYSTONE* PKI is depicted in figure 7. The basic communications that take place in such a model are the following:

1. *User to TTP communication*: Each user is equipped with the Customer access and Directory access functional units. Using the Internet, the user is connected to the Customer access or the Directory access functional units at the TTP's side. Process of a communication request and the relationships between the involved functional units is described in Figure 5 and 6.
2. KEYSTONE–*PKI–member–TTP to* KEYSTONE–*PKI–member–TTP communication*: The communication between two different *KEYSTONE* TTPs takes place through the Service Client functional unit, which belongs to the Infrastructure Supporting Services group. Since this communication is based on well-established standards, no further analysis is needed at this

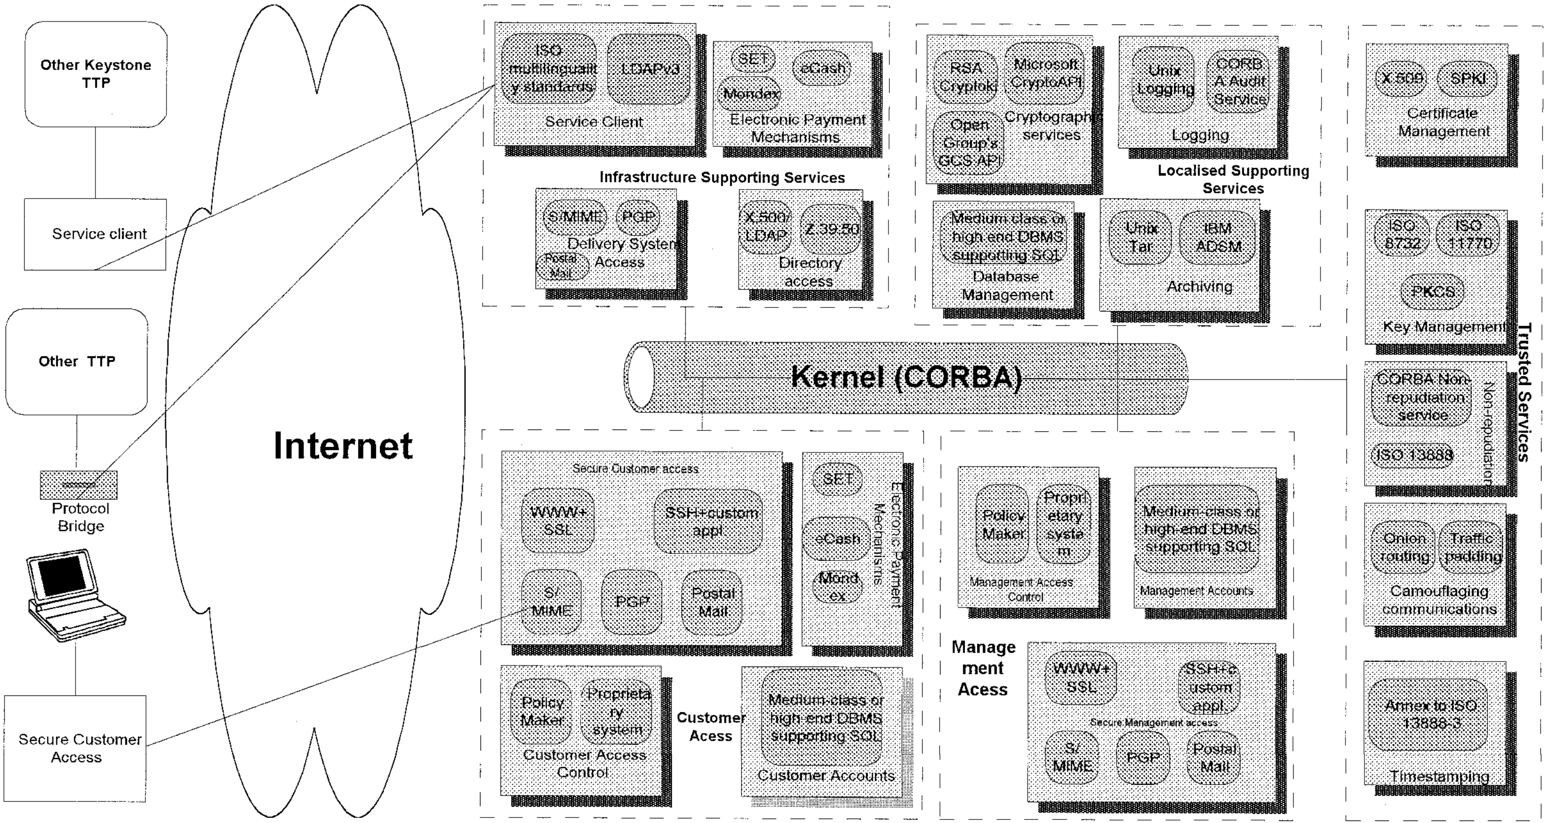

Figure 6. Unified *KEYSTONE* abstract Public Key Infrastructure architecture.

point. The same holds true for non-*KEYSTONE* PKI member TTP (or PKI) that implements the same (or compliant with those) standards that the *KEYSTONE* PKI utilizes.

3. *Other–TTP to* KEYSTONE*–PKI–member–TTP communication.* In that case protocol and standard converters or bridges should be utilized in order to translate transmitted information from *KEYSTONE* based format to other PKI's based format and vice versa. These bridges should be the intersection points between the *KEYSTONE* PKI and the foreign PKI. The bridges need not necessarily reside in sites that belong to *KEYSTONE* infrastructure.

5. Concluding remarks

A schema supporting a robust security framework for telemedical applications operating over the Web has been described in the previous sections. The schema is based on a Trusted Third Party architecture, under which Certification Authorities store the public key certificates of hospitals and medical practitioners. Digital signatures are used to provide peer and data origin authentication, and, in combination with access control lists, to provide access control.

The deployed infrastructure is based on off-the-shelf available clients and servers, and provides functions for electronic registration of participants, session initialization, user authentication, key generation and personalization, certificate generation, distribution, storage and retrieval, certificate revocation lists, and auditing.

Furthermore, the *KEYSTONE* Public Key Infrastructure architecture has been described. It includes a set of services that could be offered, as well as a set of functions implementing these services; an abstract reference model describing the operation of a PKI in terms of roles and actions; a functional specification comprising functional units and a communication kernel; a set of technologies and relevant standards implementing the defined functional units.

This integrated architecture fulfills the desirable characteristics and meets the criteria that are essential for a PKI to constitute a successful framework for the development of inter-domain and international telemedical trusted services. The object methodology followed during the architecture evolution ensures that this architecture can be adjusted to future technological variations. Moreover, the adoption of standards suitable for the TCP/IP protocol stack guarantees that the proposed architecture may constitute a good vehicle for deploying secure telemedical services, taking advantage of the Internet as the information highway backbone.

Acknowledgements

This research work was supported in part by the European Commission under DGXIII/ETS-I project 20820, and DGXIII/ETS-II project 23187. The authors would like to thank their project partners for many helpful comments.

References

1 The SEISMED Consortium (Eds.), 1995, *Data Security for Health Care*, *Vol. I*, *II*, *III*, (Amsterdam: IOS Press).

2 Barber, B., *et al.* (Eds.), 1999, *Standardisation in Health Care Security* (Amsterdam: IOS Press).

3 Garfinkel, S., and Spafford, E., 1997, *Practical Unix and Internet Security* (Sebastopol: O'Reilly & Associates, Inc.).

4 Council of Europe Recommendation R(97)5, 1997, *On the Protection of Medical Data* (Strasbourg: CoE).

5 Gritzalis, S., Iliadis, J., Gritzalis, D. Spinellis, D. and Katsikas, S., 1999, Developing secure Web-based medical applications. *Medical Informatics*, **24**, 75–90.

6 Freier, A., Karlton, P. and Kocher, P., 1996, available at http://home.netscape.com/newsref/std/SSL.html.

7 European Parliament and the Council, 1999, Directive 1999/93/EC on a Community Framework for Electronic Signatures, *Official Journal of the European Communities*, 19.1.2000, **L13/12**, EN.

8 CCITT, 1988, Recommendations X.500-X.521, Data Communication Networks Directory (Geneva: CCITT).

9 RSA Data Security Inc., 1995, S/MIME Implementation Guide, Interoperability Profile, Ver.1. (Massachusetts: RSA Inc.).

10 Yeong, W., Howes, T. and Kille, S., 1995, LDAP Lightweight Directory Access Protocol, University of Michigan, ISODE Consortium, Request For Comments RFC 1777.

11 Iliadis, J., Spinellis, D., Katsikas, S., Gritzalis, D. and Preneel, B., 2000, Evaluating Certificate Status Information Mechanisms, *Proceedings of the 7th ACM Conference on Computer and Communication Security CCS'2000*, November 2000 (New York: ACM Press) pp. 1–8.

12 CCITT Blue Book, 1988, Recommendations X.509 and ISO 9594–8, Information Processing Systems – OSI – The Directory Authentication Framework (Geneva: CCITT).

13 Kaliski, B., 1992, The MD2 Message-Digest Algorithm, Request For Comments RFC 1319.

14 Rivest, R. and Dusse, S., 1992, The MD5 Message-Digest Algorithm, Request For Comments RFC 1321.

15 Atkinson, R. J., 1995, Security Architecture for the Internet Protocol, Request For Comments RFC 1825.

16 Schneier, B. and Kelsey, J., 1998, Cryptographic Support for Secure Logs on Untrusted Machines, *Proceedings of the 7th USENIX Security Symposium*, (Berkeley: USENIX) pp. 53–62.

17 Council of Europe, 1981, Convention for the protection of individuals with regard to automatic processing of personal data, Convention No. 108 (Strasbourg: CoE).

18 ISO 17799, 2001, Code of Practice for Information Security Management (work in progress), available at www.securityauditor.net/ISO17799/.

19 OMG, 1995, CORBA: The Common Object Request Broker Architecture: Architecture and Specification.

20 OMG, 1997, CORBA: The Common Object Request Broker Architecture: Services Specification.

21 Moulinos, K. and Gritzalis, D., 2000, Cryptographic Libraries as a Means to Support Privacy-Enhanced Information Systems, *Proceedings of the 7th ACM Conference on Computer and Communication Security CCS'2000*, *Workshop on Security and Privacy in Electronic Commerce*, November 2000 (New York: ACM Press).