Defence Research and Development Canada
Recherche et développement pour la défense Canada

DEFENCE R&D DÉFENSE

Managing Identity and Access in the Defence Environment

Position Paper submitted to ADM(IM)

S. Zeber and A. Magar

Defence R&D Canada - Ottawa

TECHNICAL MEMORANDUM
DRDC Ottawa TM 2002-056
April 2002

Canada

Managing Identity and Access in the Defence Environment

Position Paper submitted to ADM(IM)

S. Zeber
DRDC Ottawa

A. Magar
Magar Security Architecture Inc.

DEFENCE R&D CANADA - OTTAWA

Technical Memorandum

DRDC OTTAWA TM 2002-056

April 2002

Abstract

Information in the defence environment is managed across many separate networks and a variety of system resources by a diverse, often dynamic, population of users. The information is distributed across different classification levels and information at a particular classification level may be subject to further caveat separation restrictions. It is both a requirement and a challenge in this environment to ensure that the information and the system resources are used and managed to support operations effectively, but in compliance with established security policies. Enforcing security policies in this environment requires the capability to manage the identities and access privileges of users and administrators in a trusted manner. Two innovative technologies have recently evolved that, when used collaboratively, provide this capability in support of security policy enforcement. One is Public Key Infrastructure (PKI) technology, and the other is Privilege Management Infrastructure (PMI) technology. This paper presents the results of initial studies undertaken to determine how these two technologies can be combined in a content-based information security model to enable the enforcement of trusted multi-caveat separation and, eventually, multi-level security for this environment. The results indicate that existing commercial-off-the-shelf PKI and PMI products do not meet current defence security policy requirements. The paper proposes enhancements to address these deficiencies and proposes a practical proof-of-concept demonstration to refine the model further. The resulting model should be easily adaptable to any government or corporate environment with similar or less rigorous security requirements.

Résumé

Dans l’environnement de la Défense, l’information est gérée, souvent de façon dynamique, par un ensemble hétérogène d’utilisateurs dans nombre de réseaux distincts et dans une variété de ressources de systèmes. L’information est communiquée à des utilisateurs de divers niveaux de classification et peut être soumise à d’autres restrictions plus rigoureuses de diffusion. Dans un tel environnement, l’utilisation et la gestion de l’information et des ressources du système de façon à assurer le soutien efficace des opérations, conformément aux politiques de sécurité en vigueur, constituent à la fois une nécessité et un défi. Pour assurer le respect de ces politiques de sécurité, il importe de gérer de façon sûre l’identité et les droits d’accès des utilisateurs et des administrateurs. Deux technologies novatrices récemment mises au point, l’Infrastructure à clés publiques (PKI ou ICP) et l’Infrastructure de gestion des droits (PMI ou IGD), permettent d’assurer un tel respect des politiques de sécurité lorsqu’elles sont utilisées conjointement. Le présent document énonce les conclusions d’études préliminaires menées dans le but de déterminer comment ces deux technologies peuvent être combinées en un modèle de sécurité de l’information basé sur le contenu qui permette l’application de plusieurs restrictions valides et, par la suite, assure une sécurité multiniveaux dans cet environnement. Les résultats obtenus indiquent que les produits PKI et PMI commerciaux existants ne satisfont pas aux exigences de l'actuelle politique de sécurité de la Défense. Le document propose certaines améliorations destinées à corriger ces lacunes ainsi qu'une démonstration pratique de validation de principe destinée à affiner davantage le modèle. Le modèle ainsi obtenu devrait pouvoir être adapté sans difficulté à tout environnement gouvernemental ou corporatif dont les exigences de sécurité sont semblables ou moins rigoureuses.

Executive summary

The traditional approach for managing access to information resources in the defence environment relies on the physical separation of networks for different classification levels and caveats. This approach inhibits the information sharing that is often required for effective operational support. Additionally, managing the information resources and access privileges is complex, inefficient and error-prone. Duplication of information in the different network domains creates synchronization difficulties which can result in unintentional security vulnerabilities.

This paper proposes a new paradigm for managing access to information resources in a single network based on the use cryptographic separation to enforce security policy, including classification levels and caveat separation. This approach, which is known as content-based information security (CBIS) is also the subject of an Advanced Concept Technology Demonstration jointly sponsored by the U.S. Joint Forces Command and the SPAWAR Systems Centre San Diego.

Enforcing security policies in a CBIS environment depends critically on the capability to authenticate identities and manage the access privileges of users and administrators in a trusted manner. Two innovative technologies have recently evolved that, when used collaboratively, can provide these capabilities. One is Public Key Infrastructure (PKI) technology, and the other is Privilege Management Infrastructure (PMI) technology.

This paper synthesizes the results of initial studies undertaken to determine how commercial implementations of these two technologies can be combined to provide a CBIS environment capable of enforcing caveat separation and, eventually, multi-level security in a trusted manner. The results indicate that existing commercial PKI and PMI products do not meet current security policy requirements due to deficiencies in, or the lack of, authentication, enhanced access control, and sensitivity labelling mechanisms. However, the paper concludes that commercial products with suitable enhancements to address these deficiencies may provide an acceptable solution and, therefore, proposes a proof-of-concept demonstration to refine the elements of potential solution further. The resulting solution should be easily adaptable to any government or corporate environment with similar or less rigorous security requirements.

Zeber, S., Magar, A.. 2002. Managing Identity and Access in the Defence Environment. DRDC Ottawa TM 2002-056. Defence R&D Canada – Ottawa.

Sommaire

Jusqu'à maintenant, la gestion de l'accès aux ressources d'information dans l'environnement de la Défense était basée sur la séparation physique des réseaux en fonction des divers niveaux de classification et de restrictions. Une telle approche a pour effet d'empêcher le partage d'information souvent essentiel à la prestation d'un soutien opérationnel efficace. En outre, la gestion des ressources d'information et des privilèges d'accès est un processus complexe, inefficace et sujet à l'erreur. La duplication de l'information dans les domaines des différents réseaux pose des problèmes de synchronisation qui peuvent donner naissance à des vulnérabilités accidentelles dans le domaine de la sécurité.

Le présent document propose un nouveau modèle de gestion des accès aux ressources d'information dans un seul réseau, basé sur la vérification cryptographique pour assurer le respect de la politique de sécurité, et comportant plusieurs niveaux de classification et de restrictions de diffusion. Cette approche, connue sous le nom de 'sécurité de l'information en fonction du contenu (CBIS)' fait également l'objet d'une démonstration de concepts technologiques évolués, parrainée conjointement par le U.S. Joint Forces Command et le SPAWAR Systems Centre de San Diego.

Le respect des politiques de sécurité dans un environnement CBIS dépend essentiellement de la capacité à authentifier l'identité et à gérer de façon sûre les privilèges d'accès des utilisateurs et des administrateurs. Deux technologies novatrices récemment mises au point, l'Infrastructure à clés publiques (PKI ou ICP) et l'Infrastructure de gestion des droits (PMI ou IGD) permettent d'assurer un tel respect des politiques de sécurité, lorsqu'elles sont utilisées conjointement.

Le présent document contient un résumé des conclusions d'études préliminaires entreprises afin de déterminer comment les versions commerciales de ces deux technologies pourraient être fusionnées de façon à créer un environnement CBIS qui puisse appliquer, de façon sûre, des restrictions de diffusion et, éventuellement, un environnement de sécurité multiniveaux. Les résultats de ces travaux révèlent que les produits commerciaux PKI et PMI existants ne satisfont pas aux exigences de la politique de sécurité actuelle en raison des carences ou de l'absence de mécanismes d'authentification, de contrôle d'accès renforcé et d'attribution de labels de sensibilité. Le document conclut cependant que ces lacunes pourraient être comblées par certaines améliorations apportées aux produits commerciaux, et propose de procéder à une démonstration de validation de principe qui permettrait de cerner avec plus de précision les avenues potentielles de solution. Le modèle ainsi obtenu devrait pouvoir être adapté sans difficulté à tout environnement gouvernemental ou corporatif dont les exigences de sécurité sont semblables ou moins rigoureuses.

Zeber, S., Magar, A. 2002. Managing Identity and Access in the Defence Environment. DRDC Ottawa TM 2002-056. R & D pour la défense Canada - Ottawa

Table of contents

List of figures

This page intentionally left blank.

1. Introduction

The challenge for information management in today's defence environment is to optimize information sharing in support of operational requirements across heterogeneous systems and multiple networks while simultaneously protecting the information resources in accordance with security policy requirements. The traditional approach relies on networks replicated at different classification levels, further partitioned according to caveat requirements The user population is diverse and highly dynamic, including military personnel, civilian defence personnel, contractors, and possibly, allies. User authentication and access privileges are managed separately for each network based on identity authentication, security clearances and role authorizations. This environment results in inefficient duplication and inhibits information sharing particularly between classification levels. Furthermore, difficulties in synchronizing changes to user security attributes across all networks, including deleting all accounts and privileges when a user leaves the organization, provide opportunities to compromise security.

This paper proposes a new content-based approach to information security that eliminates many of the inefficiencies resulting from duplication and facilitates the sharing of information while rigorously enforcing security policy requirements. The proposed model integrates standard commercial off-the-shelf (COTS) public key infrastructure (PKI) and privilege management infrastructure (PMI) technologies with a number of enhancements to meet defence security policy requirements. The combination of PKI and PMI technologies provides a robust basis for security by leveraging their individual strengths, while the enhancements address deficiencies identified in standard COTS implementations of these technologies.

This proposal is consistent with and supports the evolving information technology security strategy in the Department of National Defence (DND) where DND is deploying a PKI in all security domains. The PKI provides strong authentication of the individual user within the formal chain of command but does not provide an efficient method for granting access rights and privileges to groups or communities of interest as is often done by functional authorities. For example, a PKI can encrypt a document for a pre-defined group of users but the document must be decrypted and reencrypted for the entire group whenever the group changes. This is particularly inefficient and may become impractical when the group is large or changes frequently. Adding a PMI provides a more efficient mechanism to manage the access rights and privileges of the group to information resources such as documents based on the strong authentication of the users in the group.

2. Content-Based Information Security

Content-based information security (CBIS) seeks to protect information based on the encryption of its content at the point of origin and not based on the classification of the network in which is used. In this environment, encryption is the mechanism used to enforce the security policy requirement for the separation of objects at different classification levels. The information is protected by encryption both on servers and when in transit across a network. The CBIS approach, whose goal is to improve the capability for information sharing in a multi-level secure coalition environment, adopts a strategy which relies on the trusted labelling of information, strong authentication of users, and authorization management based on matching the information labels and the user's security attributes. CBIS is currently the subject of an Advanced Concept Technology Demonstration jointly sponsored by the U.S. Joint Forces Command and the SPAWAR Systems Centre San Diego.

3. Managing Identity

A user's identity is represented in a system or network by a data structure, usually called an account, that holds the electronic credentials of the user, and other information such as roles, groups and security attributes. Identity management deals with the creation, modification and deletion of these electronic accounts and the authentication of an individual requesting the use of an account to access information resources.

Authentication is the process that verifies a user's identity to a system, and may also verify a system identity to another system, using a trusted channel. In the case of an individual user, by verifying the user's electronic credentials, the authentication process establishes a binding between the individual and the electronic identity of that individual in the system. It is through this binding that the individual can be held accountable for all actions attributed to the corresponding electronic identity. Therefore the strength of the authentication process is a key element in the overall security of the system. The following discusses modern authentication techniques and proposes a combination of techniques suitable for a CBIS environment.

3.1 Traditional Authentication

A user's electronic credentials traditionally consist of an alphanumeric string known as a username or userID and a second alphanumeric string that may be known as a password, a pass phrase, or a Personal Identification Number (PIN). During account creation each username or userID is assigned a unique password, passphrase or PIN. In the traditional authentication process the user enters his username or userID and the corresponding password, pass phrase, or PIN. The authentication succeeds if the password, pass phrase or PIN entered matches that stored with the corresponding identity. Clearly, the integrity of this process is based on the premise that the password, passphrase or PIN is known only to the user to whom it was originally assigned. Once the user is successfully authenticated, all further access to information resources is granted based on this authenticated identity. The effectiveness of the access controls, then, depends on the strength of, and trust in, the authentication mechanism.

PINs, passwords and pass phrases are considered a weak form of authentication [1] in that they can be purposely shared, guessed, stolen or attacked using trojan horses, viruses, keyboard sniffers, shoulder surfing, eavesdropping and social engineering. Increasing the complexity of the password can reduce the effectiveness of some of these attacks. However, increasing the complexity can also be counter-productive since the user may either write it down to avoid forgetting it, which compromises the security, or he may forget it, thereby preventing legitimate access, and increasing the administrative burden by requiring a password change. To complicate matters further, users are typically required to remember a password for each system and application to which they have access.

3.2 Cryptographic Authentication

Cryptographic credentials provide a much stronger authentication mechanism. These credentials, which are unique to a particular individual, can be used in a challenge-response scheme for authentication. Cryptographic credentials can be issued and managed in a seamless and transparent manner using PKI technology. While PKI solves most of the problems with respect to managing and authenticating identity, it does have its own potential weaknesses.

The trust in PKI authentication is based on two main factors. One is the rigour with which the identity of the individual to whom the cryptographic keys are issued is verified. This is known as user registration or enrolment. The other is the strength of the mechanism used to ensure that the private key used in the authentication process is accessible only to the individual to whom it was issued. Therefore, once the keys have been issued, protection of the private key is critical to the integrity of the authentication process.

In the simplest case, private keys, may be stored in a computer system, either locally on a permanent or removable disk, or on a remote server, such as in a public repository, to facilitate roaming. In this case the keys can be encrypted using a password, pass phrase, or PIN to construct the encryption key. Such protection schemes, based on “something you know”, have the same vulnerability to attacks as the traditional password-based authentication schemes. Using cryptographic authentication shifts the vulnerability from the authentication process to the protection of the private key. An attacker will now direct his attack to attempt to obtain the private key rather than to break any cryptographically protected communication.

3.3 Hardware Tokens

The weaknesses of a password-only protection scheme can be mitigated through the use of a password-protected hardware token. This scheme combines protection based on “something you know” with “something you have”, increasing the level of security. For the highest level of assurance, only those tokens offering a tamper resistant, mobile platform on which the private key is generated, stored, and which performs all cryptographic functions should be considered. An attacker wishing to gain access to a user’s private key must now steal the hardware token in addition to the password. Knowledge of the password by itself is useless without possession of the token. This adds an additional level of security to the protection of the private key. Furthermore, the token provides a mobile platform to transport the private key securely, which allows users to access information resources from any system on the network rather than limiting them to a single workstation.

The “smart card” is by far the most common of these types of hardware tokens. However, not all smart cards are equal and even the better ones may not be suitable for the defence environment. For example, many smart cards do not perform certain cryptographic operations including random number generation, hashing and symmetric cryptography. Not only is it extremely difficult to generate the random numbers used to generate the cryptographic keys in a self-contained, sealed environment but, due to the low bandwidth interface, both bulk hashing and symmetric cryptography are extremely slow. Furthermore, a number of theoretical attacks on smart cards have been devised (protocol, micro probing, side-channel and fault analysis attacks). While many of these have been found to be impractical, some have proven successful when undertaken by a determined, well-equipped attacker.

A better, but significantly more expensive option for this environment would be a Federal Information Processing Standard (FIPS) 140-1 Level 2 or higher, validated Personal Computer (PC) Token (sometimes referred to as a Personal Computer Memory Card International Association (PCMCIA) Token). This type of token has most of the advantages of a smart card in addition to a faster processor capable of performing all of the cryptographic operations required. Rather than requiring a smart card reader, a PC Token requires a Type II PCMCIA reader that comes as a standard item with most notebook computers and some desktop computers.

3.4 Biometric Authentication

A protection mechanism based on “something you are”, known as biometric authentication, promises to provide an even stronger mechanism than that provided by cryptographic credentials, and may be used to protect such credentials. Biometric techniques include: fingerprint scanning, retinal scanning, iris scanning, signature verification, voice recognition, face recognition, and hand geometry recognition. These techniques have long been a desirable alternative to passwords (something you know) as they do not involve items that can be forgotten, guessed, stolen or easily attacked. Unfortunately, the desirable properties of biometric techniques have always been offset by a number of limitations [2] that have impeded widespread deployment of these techniques. These limitations include cost, intrusiveness, and performance.

The intrusiveness of biometric authentication methods has considerably reduced their acceptability in public and commercial sectors. Many consider the extraction and electronic storage of an individual’s physical characteristics to be offensive and unacceptable because it violates personal privacy. There is also concern over the potential for the misuse of the information.

The historical performance of biometric techniques has also affected its acceptance in the public and commercial sectors. The performance is assessed on two factors: the False Acceptance Rate (FAR) and the False Rejection Rate (FRR). The FAR is the rate at which unauthorized individuals are mistakenly authorized. The FRR is the rate at which valid users are mistakenly rejected. In a public or commercial environment where the goal is to provide optimum service to customers with a minimum of inconvenience, such as in the banking Automated Teller Machine (ATM) network, it is considered unacceptable to reject even a single valid user, even at the cost of accepting a relatively high number of invalid users. Since commercial enterprises must avoid inconveniencing their customers they require an effective FRR of zero, which may be unrealistic. They must also calculate the cost and effect of a possibly high FAR.

It can be argued that the objections of intrusiveness, performance and, to some extent, cost, do not apply to the defence environment in the same manner, so this environment is considered more suited to the use of biometric authentication techniques. First, the intrusiveness of the technology is considered to be less of an impediment in the military environment because military personnel are fully aware of the need for security and are more accustomed to intrusive security measures than are civilians. Second, when security must be balanced against convenience, the military has historically been willing to sacrifice a certain level of convenience in the interest of maintaining a desired level of security. Since a high FAR would be far more compromising to security than the inconvenience of false rejections, a higher FRR would be more acceptable in the military environment. As well, in a classified environment, the number of users with access to classified material is inversely related to the sensitivity of the classified material. As the sensitivity of the material increases the number of individuals with access to this material decreases and so does relative impact of the FRR. Finally, since the military do not have the same priorities as commercial enterprises, cost may not be as significant a factor in limiting the use of biometric techniques.

Biometric techniques may be combined with passwords to obtain acceptance of a lower FFR, thereby improving the overall convenience to the users. The lower FRR with biometric authentication may be acceptable when combined with other authentication techniques because of the additional security provided by the other two-factor technique (passwords and hardware tokens). This is in accordance with the principle of “defence in depth” which holds that the weaknesses of one safeguard should be balanced by the strengths of another. In this case, passwords and password or PIN-protected hardware tokens are both more susceptible to theft while tokens protected with biometric techniques are less so.

3.5 The Defence Environment

PKI-based authentication using PIN-protected smart cards for private key storage is being implemented in the designated domain of DND and is planned for the classified domain. PIN protected hardware tokens provide a low to medium level of assurance suitable for many applications and environments, such as the designated domain. For classified environments requiring a higher level of assurance, the hardware tokens should be protected with a good biometric authentication mechanism When hardware tokens are used in conjunction with biometric protection, the user registration process can combine enrolment in the PKI, the issuing of hardware tokens, and enrolment for biometric authentication. Furthermore, the user registration process should always require the user to appear in person with appropriate credentials to provide the necessary level of assurance.

4. Managing Access

Once a user’s identity has been authenticated successfully, the system must determine what access rights and privileges the user has. These rights and privileges are associated with roles and groups to which the user belongs, as defined by security policy. Access to information resources is usually controlled in each system or application using Access Control Lists (ACLs) associated with the information objects. Access to a resource is granted if the user’s rights and privileges match those required by the ACL for that object.

Often, because the ACLs are managed independently for the various systems and applications, it can be difficult, if not impossible, in a large organization to determine what a particular user's rights and privileges entitle him or her to do across all systems and applications in the organization. Thus when a new user joins the organization it can require a major effort to establish all of his access privileges. Likewise, when a user leaves the organization, removing all of the user’s access privileges from all systems and applications can also be a significant task. Unless the management of the ACLs is coordinated across the organization, the ACLs for different systems and applications can quickly lose synchronization and become obsolete, and accounts that should have been removed may remain active, opening security vulnerabilities. This problem is compounded when dealing with separate networks at different classification levels.

4.1 Public Key Certificates

One possible solution to this problem is to store privilege as well as identity information in public key certificates. This is attractive because the certificates are cryptographically protected by a digital signature. Although discussion of this solution is still ongoing [3], it has been concluded that public key certificates are the wrong mechanism to store privilege information for the reasons given in the following sections.

4.1.1 Jurisdiction

The entity responsible for issuing public-key certificates is generally not the same entity responsible for authorizing access to information resources. Public-key certificates are issued by a central, trusted authority which has a formal relationship with the individual. Authority to grant access privileges, on the other hand, is often delegated throughout the organization to the working level, where local management is familiar with the user’s requirements, and can respond quickly to changes in requirements for access privileges. Assigning responsibility for both functions to one role or department increases the probability of a security compromise. Separating these functions makes it more difficult for malicious individuals to compromise security.

4.1.2 Interoperability

Public-key certificate extensions are optional. If an extension field is used for access privileges, then applications must be designed or enabled to understand how to interpret this field. Applications without this "intelligence" will not be able to interpret the extension field and will not be interoperable with those that do if the field is critical to the application. Therefore if extension fields are used to store authorization information, interoperability becomes problematic.

4.1.3 Certificate Churn

By storing access privileges in a public-key certificate one drastically reduces the lifetime of that certificate since this information changes much more frequently than does authentication information. Authorization information may change with a change in job function, such as a promotion or a change in responsibility. Any change to the data in a certificate requires the old certificate to be revoked and a new certificate to be issued. Frequent changes lead to the phenomenon of certificate churn. Not only does this increase the size of Certificate Revocation Lists (CRLs) substantially, it also increases the administrative costs associated with revoking and reissuing public-key certificates. A number of security practitioners [3] have argued, with good reason, that public key certificates can be used to convey privilege in environments where the burden of proof of identity during the registration process is not so onerous. In these environments it would be relatively easy to re-issue public key certificates. However, in a defence environment using high assurance public key certificates, each certificate reissue would require the individual to appear before the local registration authority and present appropriate credentials as proof of identity.

4.2 Attribute Certificates

Once it became clear that public key certificates were the wrong mechanism to store access privileges, the international standards community responsible for the public key certificate format developed a similar certificate format without a the public key, for the express purpose of storing privilege information. Like public key certificates, however, these attribute certificates require an infrastructure to manage the certificates throughout their lifecycle. This infrastructure is commonly referred to as a PMI, and it is for this reason that many information security practitioners equate a PMI with attribute certificates.

The idea of the attribute certificate is to store privilege information in a certificate structure similar to that of a public key certificate but one that does not contain cryptographic key material. While the concept of attribute certificates was embraced within the international standards community, it has not been widely implemented. Attribute certificates are currently used in a small number of information security products ranging from web-based authorization solutions to Virtual Private Networks

(VPNs). In each case they are used to convey privilege information internally within the product rather than between products as one would expect from a true infrastructure product.

Other more ambitious uses of attribute certificates have also been proposed. For example, [4] proposes the use of attribute certificates to facilitate the electronic procurement process for the Canadian Forces. Role Specification Certificates (RSCs) would contain the privileges associated with a particular role while the Role Assignment Certificate (RAC) would assign individuals to a particular role. This design enables roles to be altered without affecting the assignment of roles. There is also a proposal [5] to use attribute certificates as a form of passport that would enable mobile agents to execute code on a given system based on the contents of the attribute certificate.

Unfortunately, while attribute certificates are an interesting manner in which to convey privilege, they suffer from a number of limitations [3, 4] that could ultimately prove detrimental to their eventual widespread adoption. The following sections describe these limitations.

4.2.1 Complexity

Deploying a PKI is a complex, expensive undertaking, which has significantly delayed its widespread adoption. A PMI for attribute certificates has much of the same complexity as a PKI, but there are viable alternatives.

4.2.2 Dependency

Attribute certificates are cryptographically protected from alteration by a digital signature. This is highly beneficial in that it allows attribute certificates to be posted to a public directory or transmitted over a network without fear of modification. Unfortunately, it also creates an extremely restrictive dependency that limits the deployment of this technology to those environments with an established PKI. The alternative is for organizations to attempt to deploy the two technologies together or in quick succession, thereby drastically increasing the complexity of the deployment.

4.2.3 Interoperability

As interoperability testing of PKI products from different vendors has shown, compliance to standards is no guarantee of interoperability between products. In the case of attribute certificates, interoperability problems are exacerbated by the use of attribute certificate extensions which can be designated “critical”, leading to certificate rejection if critical extensions are not recognized by other implementations.

4.2.4 Performance

Since attribute certificates are digitally signed they require PKI services. Verification of an attribute certificate involves validating the corresponding digital signature, which in turn requires the verification of at least one public key certificate. In the case of a large attribute certificate-based PMI, privilege will be delegated through a number of levels. Validating these delegation paths can place significant performance demands on an organization's information systems, resulting in performance degradation that may not be fully understood until an attribute certificate-based PMI has been widely deployed throughout the organization. The performance implications for an open environment may be even more severe.

Although attribute certificates provide a theoretically attractive mechanism to convey privilege, the practical limitations just discussed and the relatively immature state of the technology mitigate against the wide-spread implementation of attribute certificate-based PMI at this time. In fact, there are currently no large-scale implementations of PMIs that use public key or attribute certificates. Attribute certificate technology may eventually achieve the maturity and interoperability required for widespread deployment, however, this is a long term prospect. For the present, attribute certificates can provide only a partial solution, rather than a complete solution.

5. The Standard PMI Solution

5.1 The Standard Model

In spite of its current limitations, a review [6] of commercial PMI offerings provides a valid conceptual description of a standard PMI as *an enterprise-wide authorization management system capable of providing controlled access by communities of users to diverse information resources located on disparate computer systems according to a unified security policy. It is also centrally managed with delegated, de-centralized administration.* The essential elements of the standard PMI include a central Access Management Policy Server with an administrative interface, a private database, a public repository, and distributed Access Management Agents that control access to resources locally in accordance with the security policy defined in the central policy server. This concept of a standard PMI is illustrated in Figure 1. A Windows 2000 domain is an example of a single-vendor implementation of such a PMI.

In most environments the access control capabilities provided by the applications and systems, and managed locally, are sufficient. This provides distributed, locally managed access control. The standard PMI can provide centrally managed access control in accordance with an organization-wide security policy through the use of distributed Access Management Agents. These Agents, co-located with the various systems and applications across the organization, communicate with a central Policy Server to distribute and synchronize centrally-defined user, group, and privilege information. Access Management Agents have limited functionality as they merely configure the local access control mechanisms. They do not enhance them. Adding, removing, and modifying a user's privileges for each system and application in the organization is done once at the central Policy Server and distributed via the Access Management Agents. Furthermore, the Access Management Agents can be setup to reconcile the differences between what has been defined centrally and what exists in the local system. Thus, unauthorized privileges, dormant accounts, etc., can all be automatically deleted, thereby improving the overall security posture of the organization.

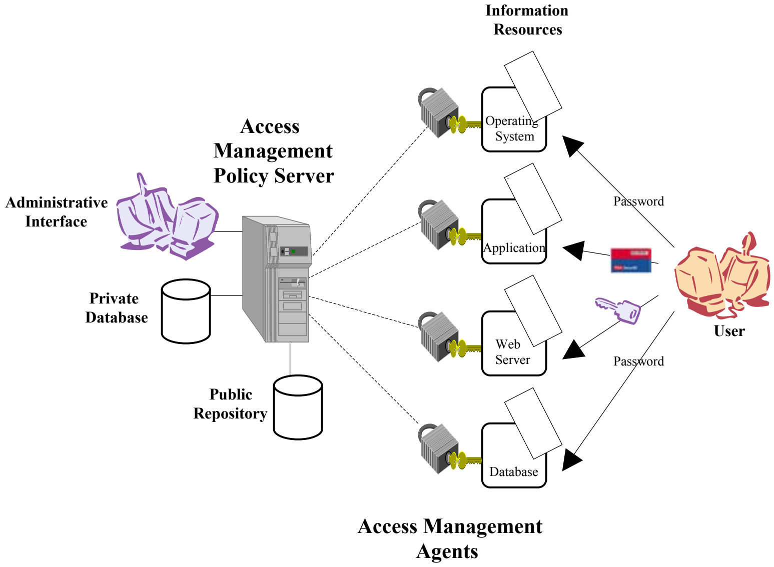

Figure 1. *A Standard PMI*

5.2 Deficiencies of the Standard Model

While the standard PMI provides a basic capability for managing access privileges, it currently lacks a number of capabilities that would be required to implement a CBIS environment, particularly for a classified domain. These deficiencies are described in the following sections.

5.2.1 Strong Authentication

Most systems and applications do not currently support PKI-based authentication using hardware tokens and biometrics. Since the protection afforded by access controls depends on the strength of the authentication mechanism used, a PMI in a CBIS environment must include PKI-based authentication using hardware tokens and biometrics.

5.2.2 Enhanced Access Control

The standard PMI defines the security policy centrally but relies on the native access control capabilities of the systems and applications for enforcement. For example, a user attempting to access a file on a system running the Solaris operating system would be permitted or denied access by the Solaris native access control capabilities even though the access control lists governing such access have been established by a centrally defined access policy. However, the native access control capabilities of these systems and applications may not meet the security policy requirements for Caveat Separation and a Multi-Level Secure mode of operation. Therefore, enhanced access control capabilities will be required.

5.2.3 Sensitivity Labelling

In a defence environment information is classified according to its sensitivity and is managed in accordance with security policy directives applicable to this level of sensitivity. The CBIS approach relies on a trusted labelling mechanism to be able to protect and control access to an information resource in accordance with its classification, caveats, and the security attributes of users, as required by security policy.

6. An Enhanced PMI Solution

6.1 Proposed Model

A standard PMI, with enhancements to implement a CBIS environment is proposed as a suitable model for managing identity and access in a defence environment that requires caveat separation and multi-level secure operation. The proposed model, illustrated in Figure 2, is referred to as an enhanced PMI.

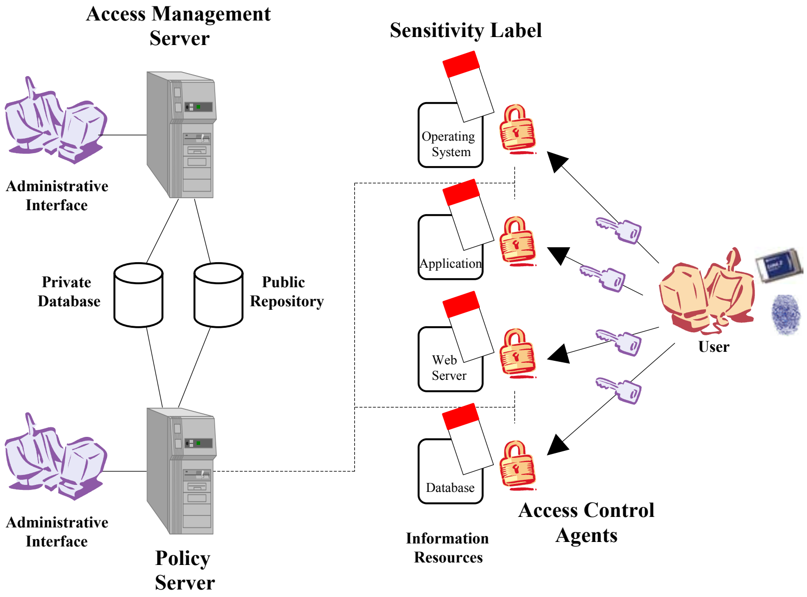

Figure 2. *An Enhanced PMI*

The enhancements to the standard PMI are the following:

- Access Rights Management implemented via distributed Access Control Agents,
- Electronic Sensitivity Labelling, and

- Policy-Enforced Access Control.

The following sections describe the properties of an enhanced PMI in more detail.

6.2 Access Rights Management

Access Rights Management deals with a user's access rights and privileges to information resources. These can depend on such factors as security clearance, role, the classification of the information resource, and caveat restrictions, which, in turn, are typically based on factors such as nationality, rank, and role. Access Rights Management involves the creation and management of identities, the assignment and management of corresponding rights and privileges, and the population of this information in a consolidated data store (either a private database or a public directory). The Policy-Enforced Access Control component uses the information in the data store to determine whether or not to grant a particular user access to a particular information resource. Access Rights Management in this proposed model has the following characteristics:

6.2.1 Role-Based Access Control

Instead of managing privileges for individual users, users are assigned to groups or roles according to their functions, and the privileges are managed for the groups and roles. This is referred to as Role-Based Access Control (RBAC). For example, a group INT may be defined for the intelligence community, and a Stabilization Force (SFOR) group may be defined to include all personnel assigned to the NATO operation in Bosnia.

6.2.2 Delegated Administration

It is generally not effective for a single individual or group to administer all identities and access privileges in a defence environment. While the scale of the task alone may prohibit this, it also violates the "separation of duties" principle of security. This principle requires that responsibility for sensitive functions be divided amongst multiple personnel so that no one individual can compromise the function. Delegated administration solves this problem by delegating the management of users and groups to a variety of individuals throughout the organization. This has the benefit of providing separation of duties while at the same time giving control to the appropriate authority. For example, the administration of the SECRET and TOP SECRET security clearance groups might be assigned to the Special Investigation Unit (SIU) of the military police. Likewise, the command responsible for supporting the NATO operation in Bosnia could be responsible for administering the SFOR group. If access to resources is based on the aggregate of group and role memberships, then no one individual would be able to compromise the entire system.

6.2.3 Rights Enablement Automation

When a new user is added to the system an identity is created and assigned to one or more groups or roles to enable access to the information resources required to accomplish his job. Likewise, when a user changes jobs or leaves the defence organization entirely, his identity must be removed from the respective groups and roles. While access can be, and likely will be, terminated immediately by revoking the departing individual's public key credentials, it is prudent to maintain access rights as accurately as possible. Since the administration of the various groups and roles may have been distributed throughout the organization, carrying out these processes is no easy task. Rights Enablement Automation facilitates these processes by providing a workflow capability that would automatically contact the appropriate authorities for the necessary approval. For example, if a new user were to be added to a classified domain, rights enablement automation may be configured to send requests to the SIU, the Canadian NATO and the Signals Intelligence (SIGINT) administrators to add this particular user to their respective groups. By corroborating this request with their own information the administrator would decide whether or not to approve the request. Provided that the requests were approved, the user would have the appropriate entitlements required for him to access, for example, the SECRET NATO signals intelligence information resources required to fulfill his job.

6.3 Electronic Sensitivity Labelling

As noted previously, information in a defence environment is classified according to its sensitivity, which is related to the potential consequences of the information being compromised. For paper-based information this marking appears as a character string label on each page and possibly also for each paragraph of text. Labelled material must be handled in accordance with the dictates of security policy as it applies for that label.

NATO, in order to facilitate the secure exchange of electronic information among NATO facilities and with member nations and allies, is developing guidance for the labelling and handling of electronic information such as electronic messages and electronic documents. In this environment the label *is a piece of electronic data that has been encoded to represent the same sensitivities as in the paper environment*. [7] *Attaching a label to electronic information ... promotes originator awareness of the requirement for correct and consistent marking, facilitates automated access and release control, enables the use of multi-level security systems, and removes the need to thoroughly examine electronic information in order to determine its sensitivity.*[8]

An electronic sensitivity label must be bound [7] to the information resource in such a way that it cannot be removed or altered by an unauthorized person. This binding between the label and the information resource must be at least as strong as the security provided by other components of the PMI. A weak binding would be susceptible to an attack which would allow an attacker to change the classification or the release control in order to gain access to the information. Two alternative implementations [8] capable of providing a strong binding are as follows:

- <u>Security Server</u> – A Security Server could store electronic sensitivity labels for each information resource in such a way that when the information resource is accessed the corresponding label would automatically be processed as well.

- <u>Digital Signature</u> – A digital signature could bind the electronic sensitivity label to the information resource. Any modification of either the information resource or the label would invalidate the binding.

Detailed specifications of these implementations are beyond the scope of this paper.

6.4 Policy-Enforced Access Control

The Policy-Enforced Access Control component embodies the process which grants or denies a request for access to a resource. This process includes an *access control decision function* and an *access control enforcement function*. The access control enforcement function grants access to an information resource if and only if the access control decision function approves the access.

6.4.1 Access Control Decision

The Policy-Enforced Access Control server (or Policy server) provides the access control decision function as well as the properties described in the following sections.

6.4.1.1 Policy Control

The Policy server enables an organization to define security policies centrally while enforcing them consistently throughout the organization. The security policy or rules defined centrally can be as simple or complex as required. Complex policies can make access decisions based on dynamic information, take behavioural patterns into consideration and even react to access attempts in various ways. It is critical that the security policy governing access to a particular information resource cannot be circumvented merely by copying the resource to a new system. The security policy must migrate with the resource so that it is consistently protected regardless of where it is located within the organization.

6.4.1.2 Monitoring

The Policy server allows an organization to monitor all accesses to information resources and to store this information in a protected audit log. The system can be configured to record access attempts without enforcing the security policy and to indicate whether an access would have been allowed or denied had the policy been enforced. This capability allows an organization to test their security policies prior to actually enforcing them.

6.4.1.3 Reaction

The Policy server can detect and react to security policy and access violations. For example, if a particular information resource is accessed, either successfully or unsuccessfully, an e-mail notifying the owner of the resource can be sent automatically.

6.4.2 Access Control Enforcement

In an enhanced PMI, distributed Access Control Agents provide the access control enforcement function. When an authenticated user attempts to access a protected information resource, an Access Control Agent blocks the attempt and sends the identity of the user and the sensitivity label of the resource to the Policy Server. The Policy Server evaluates the user's access request against the defined security policy and returns an "access approved" or "access denied" response to the Access Control Agent which enforces the decision. If access is approved it allows access to the information resource. If access is denied it blocks the access attempt and sends an appropriate message to the user. Whereas Access Management Agents in a standard PMI rely on the native access control capability of the local system or application, Access Control Agents in an enhanced PMI supplement the native access control capability of the system or application. As a result, they are more complex than Access Management Agents. Access Control Agents can also be used to provide strong authentication using public key credentials (see section 7).

7. An Integrated PKI/PMI Solution

The previous discussions on authentication and access management suggest that the best improvement in organization-wide security may result from integrating the PKI and enhanced PMI technologies. This integrated approach leverages the strengths of both technologies to provide a level of security suitable for the defence environment. The integrated PKI/PMI solution has the following advantages:

7.1 Access Rights Management

7.1.1 Administration

The user registration process benefits from integration. Registering users and groups in the PKI and then repeating the process for a PMI is inefficient because it duplicates effort. Significant savings can be achieved by combining the two registration processes within a single administrative role. A number of COTS products have a common interface or an Access Management Agent capable of adding a user to a PKI once that user has been created in the PMI. However, consolidating the administrative roles for these two processes may not be desirable in organizations where the responsibility for these two roles resides in different organizational units. Furthermore, unless the consolidation is managed properly , combining two sensitive functions within a single role can actually increase the probability of a security compromise.

7.1.2 Rights Enablement Automation

Rights Enablement Automation benefits from integration. The digital signature provided by the PKI can be used to provide enhanced security and non-repudiation for the Access Rights Management workflow capability. Requests for additional entitlements and role/group membership can be digitally signed by human resources or the user himself. These requests could then be automatically forwarded to the appropriate authorities who would in turn digitally sign the request. Provided that the digital signatures were valid the user would receive the requested entitlements.

7.2 Electronic Sensitivity Labelling

7.2.1 Binding

The digital signature method of binding a sensitivity label to an information resource can be used. This is preferred because the cryptographic binding facilitates the detection of any tampering.

7.3 Policy-Enforced Access Control

7.3.1 Certificate-based Authentication

In an enhanced PMI, Access Control Agents can support certificate-based user authentication. An Access Control Agent typically extracts the pertinent user information from the certificate and passes it to the Policy Server to complete the authentication process. This information is likely to include the identity of the user, but it can be expanded to include other information stored in the certificate extensions. An integral component of certificate-based authentication is CRL checking. This can include basic CRL checking or advanced CRL checking including support for CRL distribution points and the Online Certificate Status Protocol (OCSP).

7.3.2 Component Integration

Communications between Access Control Agents, and the Policy-Enforced Access Control Server must be protected with confidentiality, integrity and mutual authentication security services. This can best be accomplished by issuing public-key credentials to the Policy Server and to each Agent. This requires either an enterprise-wide PKI or PKI functionality built into the PMI. An additional benefit to this approach is the capability to revoke a compromised agent instantly.

7.3.3 Common Repository

Non-sensitive or protected (using digital signatures or encryption) access privileges and identity information should be stored in a public repository. Likewise, sensitive or unprotected access and identity information should be stored in a private database. These common data stores facilitate administration, enhance interoperability and improve the auditability of identity and access throughout an organization.

7.3.4 Deployment

As with any infrastructure product, deploying a PKI is a labour-intensive, time-consuming process that can take months or even years depending on the size of the organization and the distribution of users. An enhanced PMI can facilitate the deployment of a PKI by providing authenticated access to protected information resources using basic (user name and password) and certificate-based authentication. When the PKI deployment is complete, access to information resources can be limited to certificate-based authentication across the whole organization.

8. Conclusions

The challenge in the current defence environment is to optimize information sharing in support of operational requirements across heterogeneous systems and multiple networks while simultaneously protecting the information resources in accordance with security policy requirements. The traditional approach which achieves security domain separation by network separation results in inefficient duplication and inhibits necessary information sharing between domains. This paper has proposed a new content based approach to information security that provides security domain separation through cryptography. This approach eliminates many of the inefficiencies resulting from duplication and facilitates the sharing of information while rigorously enforcing security policy requirements.

The proposed model is based on integrating standard COTS PKI and PMI technologies, with a number of enhancements to meet military requirements. These enhancements include PKI authentication using hardware tokens and biometric techniques, an electronic sensitivity labelling capability, access rights management using access control agents, and policy-enforced access control. The combination of PKI and PMI technologies provides a robust basis for security by leveraging their individual strengths, while the enhancements seek to address deficiencies identified in standard COTS implementations of the technologies. In the Classified domain, this approach would provide a more flexible infrastructure to support new coalition connectivity tasks. In the Designated domain, it could support e-Commerce and specific communities of interest such as hospital medical staff.

Before this proposed model can be adopted for operational use, however, it will be necessary to validate its assumptions and test its integrity in a practical proof-of-concept laboratory demonstration. Such a demonstration will provide a practical evaluation of the model, its various components and its implementation strengths and weaknesses that cannot be predicted by a theoretical analysis.

It must also be emphasized that once validated, the model should be readily easily adaptable to other environments including both the government and the private sector.

9. References

1. *The Challenge of User Authentication*, White Paper, Ankari[1], 2001.

2. Frazee, S., *Biometrics … Why Bother?*, SANS Institute, June 29, 2001.

3. Wilson, S., *Some Limitations of Attribute Certificates*, beTRUSTed, Cryptographic Centre of Excellence (CCE) Journal, Issue 3, 2000.

4. Grandy, C., *Using a Privilege Management Infrastructure to Support Business Processes Within the Department of National Defence and the Canadian Forces*, Master of Engineering Thesis, Royal Military College, April 2001.

5. Jansen, W., and Karygiannis, T., *Privilege Management of Mobile Agents*, National Information System Security Conference, October 2000.

6. Magar, A., *Privilege Management Infrastructure*, Defence Research Establishment Ottawa, March 30, 2001.

7. *Infosec Technical Directive for Labelling of NATO Information in Electronic Format*, Version 2.0, AC/322(SC/4-AHWG/6)WP/6, September 20, 2001, NATO UNCLASSIFIED.

8. *Electronic Labelling of NATO Information*, Version 2.0, AC/322(SC/4-AHWG/6)WP/7, September 20, 2001, NATO UNCLASSIFIED.

[1] The American Biometric Company Ltd, a privately-owned Ottawa company conducting business as Ankari, was acquired by ActivCard of Freemont, California in November 2001.

List of symbols/abbreviations/acronyms/initialisms

ACL	Access Control List
ATM	Automated Teller Machine
CBIS	Content-based Information Security
COTS	Commercial off-the-shelf
CRL	Certificate Revocation List
DND	Department of National Defence
FAR	False Acceptance Rate
FRR	False Rejection Rate
FIPS	Federal Information Processing Standard
NATO	North Atlantic Treaty Organization
OCSP	Online Certificate Status Protocol
PC	Personal Computer
PCMCIA	Personal Computer Memory Card International Association
PIN	Personal Identification Number
PKI	Public Key Infrastructure
PMI	Privilege Management Infrastructure
RAC	Role Assignment Certificate
RBAC	Role-Based Access Control
RSC	Role Specification Certificate
SFOR	Stabilization Force

SIGINT	Signals Intelligence
SIU	Special Investigation Unit
VPN	Virtual Private Network

DOCUMENT CONTROL DATA

(Security classification of title, body of abstract and indexing annotation must be entered when the overall document is classified)

1. ORIGINATOR (the name and address of the organization preparing the document. Organizations for whom the document was prepared, e.g. Establishment sponsoring a contractor's report, or tasking agency, are entered in section 8.)

DRDC Ottawa
Ottawa Ontario
K1A 0Z4

2. SECURITY CLASSIFICATION (overall security classification of the document, including special warning terms if applicable)

UNCLASSIFIED

3. TITLE (the complete document title as indicated on the title page. Its classification should be indicated by the appropriate abbreviation (S,C or U) in parentheses after the title.)

Managing Identity and Acces in the Defence Environment (U)

4. AUTHORS (Last name, first name, middle initial)

Zeber, Dr. S., Magar, A.

5. DATE OF PUBLICATION (month and year of publication of document)

April 2002

6a. NO. OF PAGES (total containing information. Include Annexes, Appendices, etc.)

24

6b. NO. OF REFS (total cited in document)

8

7. DESCRIPTIVE NOTES (the category of the document, e.g. technical report, technical note or memorandum. If appropriate, enter the type of report, e.g. interim, progress, summary, annual or final. Give the inclusive dates when a specific reporting period is covered.)

Technical Memorandum

8. SPONSORING ACTIVITY (the name of the department project office or laboratory sponsoring the research and development. Include the address.)

IO section, DRDC Ottawa
3701 Carling Avenue
Ottawa K1A 0Z4

9a. PROJECT OR GRANT NO. (if appropriate, the applicable research and development project or grant number under which the document was written. Please specify whether project or grant)

5BF27

9b. CONTRACT NO. (if appropriate, the applicable number under which the document was written)

10a. ORIGINATOR'S DOCUMENT NUMBER (the official document number by which the document is identified by the originating activity. This number must be unique to this document.)

DRDC Ottawa TM 2002-056

10b. OTHER DOCUMENT NOS. (Any other numbers which may be assigned this document either by the originator or by the sponsor)

11. DOCUMENT AVAILABILITY (any limitations on further dissemination of the document, other than those imposed by security classification)

(X) Unlimited distribution
(   ) Distribution limited to defence departments and defence contractors; further distribution only as approved
(   ) Distribution limited to defence departments and Canadian defence contractors; further distribution only as approved
(   ) Distribution limited to government departments and agencies; further distribution only as approved
(   ) Distribution limited to defence departments; further distribution only as approved
(   ) Other (please specify): □□□□□□□□□□

12. DOCUMENT ANNOUNCEMENT (any limitation to the bibliographic announcement of this document. This will normally correspond to the Document Availability (11). However, where further distribution (beyond the audience specified in 11) is possible, a wider announcement audience may be selected.)

Unlimited

13. ABSTRACT (a brief and factual summary of the document. It may also appear elsewhere in the body of the document itself. It is highly desirable that the abstract of classified documents be unclassified. Each paragraph of the abstract shall begin with an indication of the security classification of the information in the paragraph (unless the document itself is unclassified) represented as (S), (C), or (U). It is not necessary to include here abstracts in both official languages unless the text is bilingual).

Information in the defence environment is managed across many separate networks and a variety of system resources by a diverse, often dynamic, population of users. The information is distributed across different classification levels and information at a particular classification level may be subject to further caveat separation restrictions. It is both a requirement and a challenge in this environment to ensure that the information and the system resources are used and managed to support operations effectively, but in compliance with established security policies. Enforcing security policies in this environment requires the capability to manage the identities and access privileges of users and administrators in a trusted manner. Two innovative technologies have recently evolved that, when used collaboratively, provide this capability in support of security policy enforcement. One is Public Key Infrastructure (PKI) technology, and the other is Privilege Management Infrastructure (PMI) technology. This paper presents the results of initial studies undertaken to determine how these two technologies can be combined in a content-based information security model to enable the enforcement of trusted multi-caveat separation and, eventually, multi-level security for this environment. The results indicate that existing commercial-off-the-shelf PKI and PMI products do not meet current defence security policy requirements. The paper proposes enhancements to address these deficiencies and proposes a practical proof-of-concept demonstration to refine the model further. The resulting model should be easily adaptable to any government or corporate environment with similar or less rigorous security requirements.

14. KEYWORDS, DESCRIPTORS or IDENTIFIERS (technically meaningful terms or short phrases that characterize a document and could be helpful in cataloguing the document. They should be selected so that no security classification is required. Identifiers such as equipment model designation, trade name, military project code name, geographic location may also be included. If possible keywords should be selected from a published thesaurus. e.g. Thesaurus of Engineering and Scientific Terms (TEST) and that thesaurus-identified. If it is not possible to select indexing terms which are Unclassified, the classification of each should be indicated as with the title.)

Access management, access control, access rights, information security, multi-level security, public key infrastructure, PKI, X.509, certificates, privilege management infrastructure, PMI, security policy, provisioning, authentication, biometrics, caveat separation, smart card, sensitivity labelling, content based information security