Managing Identity and Access in the Defence Environment

Dr. S. Zeber*
DRDC Ottawa
3701 Carling Avenue
Ottawa, ON K1A 0Z4
(613) 991-1388
Steve.Zeber@drdc-rddc.gc.ca

A. Magar*
Magar Security Architecture Inc.
Suite 310, 95 Beech Street
Ottawa, ON, K1S 3J7
(613) 233-0199
alanmagar@rogers.com

Abstract

Information in today's defence environment is managed and used by a diverse, often dynamic, population of users, with resources distributed across many separate networks. The information may be classified at different levels and may also be subject to caveat restrictions. The goal is to use the information to support operations effectively while complying with established security policies. Enforcing security policies demands that the identities and access privileges of users and administrators are managed in a trusted manner. Two innovative technologies have recently evolved that, when used collaboratively, provide this capability: Public Key Infrastructure (PKI) technology, and 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. 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.

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.

* Sponsored by Defence Research and Development Canada Ottawa.

This paper proposes a new content-based approach to information security that eliminates 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. Adding a PMI provides a more efficient mechanism to manage the access rights and privileges to information resources such as documents based on the strong authentication of users' identities.

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**

Identity management deals with the creation and modification of the electronic credentials of entities including users, processes, and systems, and the authentication of the identities of these entities when requesting access to system and information resources. In particular, the strength of the authentication mechanism is a key factor affecting the level of assurance associated with the access control mechanisms used to protect system resources.

The authentication mechanism in a CBIS environment must provide a higher level of assurance than traditional authentication mechanisms such as passwords or passphrases, and even two-factor mechanisms such as authentication using cryptographic credentials on a hardware token protected by a passphrase or a personal identification number (PIN).

3.1 *Biometric Authentication*

Biometric authentication provides the strongest authentication mechanism, in principle, because it is based on a unique personal biological characteristic. Biometric authentication can be based on techniques such as fingerprint scanning, retinal scanning, iris scanning, signature verification, voice recognition, face recognition, and hand geometry recognition. However, these techniques do have limitations [Frazee, 2001] that have impeded their widespread adoption and use, including cost, intrusiveness, and performance.

The intrusiveness of biometric authentication techniques, which involve the extraction and electronic storage of an individual's physical characteristics is considered unacceptable in the public and commercial sectors because it is seen as a violation of personal privacy, and the potential for misuse of the information is considered serious.

The performance of biometric techniques has also affected its acceptance in the public and commercial sectors. In practice, no technique is one hundred per cent accurate, and a balance must be struck between an acceptable False Acceptance Rate (FAR) and the False Rejection Rate (FRR). In a public or commercial environment any non-zero FRR can be unacceptable while even a low FAR may expose the system to an unacceptable risk.

3.2 *The Defence Environment*

The defence environment is considered well suited to the use of biometric authentication techniques since the objections of intrusiveness, performance, and to some extent, cost, do not necessarily carry the same weight in this environment. The objection of intrusiveness and violation of personal privacy has little validity because the military environment, by its very nature, requires more intrusive security measures than in the civil environment. Similarly, the performance objection is also much less significant since a non-zero FRR may be considered acceptable when balanced against the increased assurance of a lower FAR. Moreover, in a classified environment, the number of users with access to classified material is usually related inversely to the level of classification, so that the relative impact of the FRR on users decreases as the level of classification increases. Finally, since the mandates and priorities of defence departments and agencies differ from those of the civil and commercial sectors, cost may not be as significant a factor in limiting the use of biometric techniques.

Three-factor authentication mechanisms that combine a biometric technique with a password, passphrase or PIN to protect cryptographic credentials on a hardware token provide, in principle, the highest level of assurance possible. The three-factor combination provides a potentially lower FRR without raising the FAR, thereby improving the overall convenience to users while maintaining or improving the overall level of assurance.

3.3 *The Canadian 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**

Access management addresses the problem of controlling access to system resources by granting a user appropriate access rights based on the user's authenticated identity. A user's rights are associated with the roles and groups to which the user belongs, as defined by security policy. This section describes possible approaches to access management.

4.1 ***Access Control Lists***

Traditionally, a system or application controls access to an information resource or object using an Access Control List (ACL) associated with the resource. Permission to access the resource is granted if the user's rights match those required by the ACL for that resource. 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 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 or her access rights. Likewise, when a user leaves the organization, removing all of the user's access rights 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.

4.2 ***Public Key Certificates***

One possible solution to this problem is to store access rights 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 [Wilson, 2000], it has been concluded that public key certificates are the wrong mechanism to store access rights for the reasons given in the following sections.

4.2.1 ***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 rights, 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 rights. 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.

4.2.2 *Interoperability*

Public-key certificate extensions are optional. If an extension field is used for access rights, 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.

4.2.3 *Certificate Churn*

By storing access rights 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 [Wilson, 2000] have argued, with good reason, that public key certificates can be used to convey access rights 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.

4.3 *Attribute Certificates*

Once it became clear that public key certificates were the wrong mechanism to store privilege information, such as access rights, the international standards community responsible for the public key certificate format developed a similar certificate format without 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, [Grandy, 2001] 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 [Jansen and Karygiannis, 2000] 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 [Wilson, 2000] [Grandy, 2001] that could ultimately prove detrimental to their eventual widespread adoption. The following sections describe these limitations.

4.3.1 ***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.

4.3.2 ***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.

4.3.3 ***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.

4.3.4 ***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.

4.3.5 *Summary*

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 *Standard Model*

In spite of its current limitations, a review [Magar, 2001] 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.

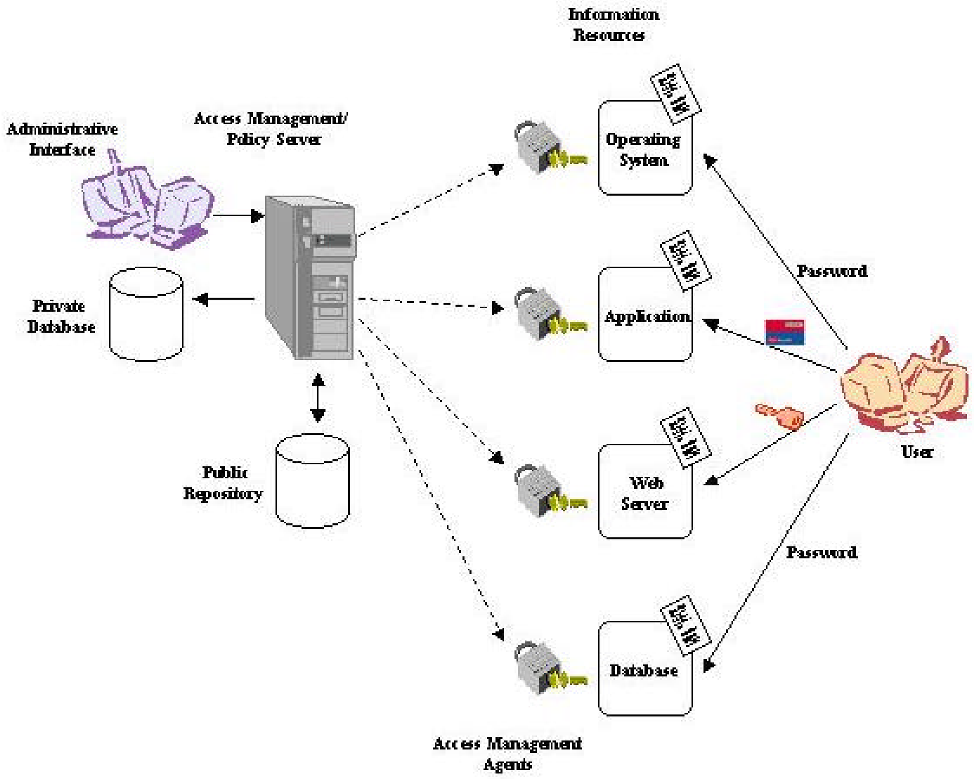

Figure 1. Standard PMI

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.

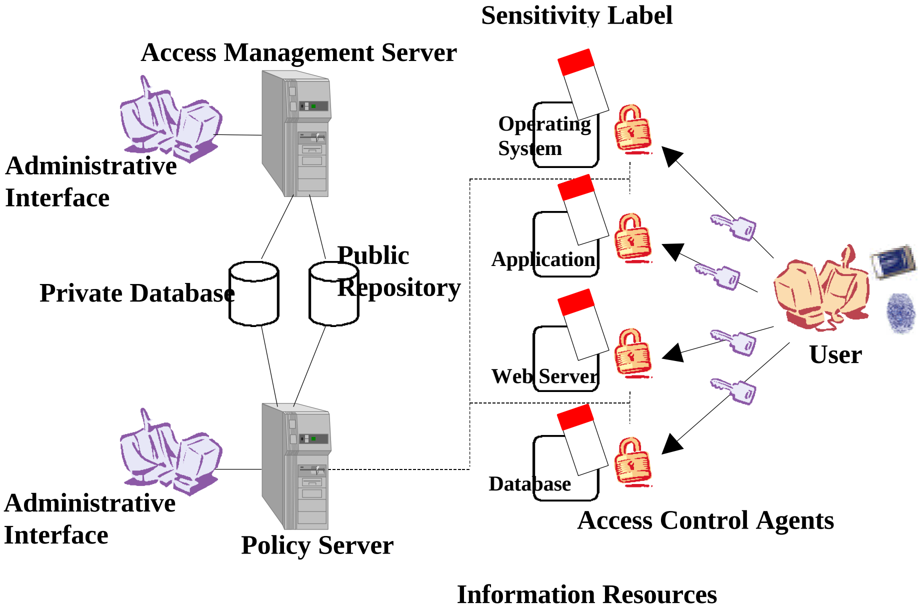

Figure 2. Enhanced PMI

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:

<u>Role-based Access Control:</u> 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.

<u>Delegated Administration:</u> 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, which 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.

<u>Rights-enablement Automation:</u> 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 approvals.

6.3 *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.

In order to facilitate the secure exchange of electronic information among NATO facilities and with member nations and allies, NATO 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*. [WP/6, 2001] *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*.[WP/7, 2001]

An electronic sensitivity label must be bound [WP/6, 2001] 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 [WP/7, 2001] capable of providing a strong binding are as follows:

Security Server – 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.

Digital Signature – 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 Decisions*

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

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 behavioral 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.

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.

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 next section).

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*

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.

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*

Although the digital signature method of binding a sensitivity label to an information resource is preferred it facilitates the detection of tampering, the security server mechanism may be easier to implement.

7.3 *Policy-Enforced Access Control*

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).

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.

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.

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**

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

[Grandy, 2001] 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.

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

[Magar, 2001] Magar, A., *Privilege Management Infrastructure*, Defence Research Establishment Ottawa, March 2001, DRDC Ottawa CR 2002-058.

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

[WP/6, 2001] *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.

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