

PEI Models towards Scalable, Usable and High-Assurance Information Sharing

Ram Krishnan
Laboratory for Information
Security Technology
George Mason University
Fairfax, VA, USA
rkrishna@gmu.edu

Ravi Sandhu
Institute for
Cyber-Security Research
Univ of Texas at San Antonio
San Antonio, TX, USA
ravi.sandhu@utsa.edu

Kumar Ranganathan
Intel System Research Center
Bangalore, India
kumar.ranganathan@intel.com

ABSTRACT

Secure Information Sharing (SIS) or “share but protect” is a challenging and elusive problem both because of its broad scope and complexity ranging right from conception (objective and policy) to culmination (implementation). In this paper, we consider how to solve SIS challenges with three main and conflicting objectives: *scalability*, *usability* and *high-assurance*. In the context of SIS, high-assurance requires strong controls on the client. It is widely accepted that such controls cannot be entirely software-based. In this regard, we consider solutions based on commercially emerging hardware-rooted Trusted Computing Technology. For SIS, we argue super-distribution (“protect once and access wherever authorized”) and off-line access are necessary to achieve scalability and usability. We limit super-distribution to occur within a group of Trusted Platform Module [1] or TPM-enabled machine. For simplicity, we assume all content that are distributed to be read-only. Drilling down, we discuss Policy, Enforcement and Implementation (PEI) models for SIS within a group (group-based SIS or g-SIS).

Categories and Subject Descriptors

D.4.6 [Operating Systems]: Security and Protection – Access controls; K.6.5 [Management of Computing and Information Systems]: Security and Protection – Unauthorized access

General Terms

Security

Keywords

Information Sharing, Trusted Computing, PEI, Client-side Access Control, High-Assurance

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

SACMAT'07, June 20-22, 2007, Sophia Antipolis, France.
Copyright 2007 ACM 978-1-59593-745-2/07/0006 ...\$5.00.

1. INTRODUCTION

Sharing information (objects) *while* protecting it is one of the earliest problems to be recognized in computer security, and yet remains a challenging problem to solve. Although SIS sounds like an oxymoron, its application scenarios are endless ranging from revenue centric retail Digital Rights Management (DRM) and sensitivity centric intellectual property to national security centric secret property. Classic access control models are either inherently weak or don't even address this problem domain. The Discretionary Access Control model or DAC as discussed in [4, 5, 2] is fundamentally limited in that they control access only to original objects but not to copies. If objects could be read, one can read and create a copy of this object. Objects are protected up to the point when read access is granted to a subject. From then on, the owner has no control on his object. Mandatory Access Control models or MAC (like Bell-Lapadula) as discussed in [3] address information flow but are too rigid for fine-grained access control and falling back to DAC for fine-grained access control as the Orange Book suggests [2] is pointless.

In this paper, we approach the Secure Information Sharing (SIS) problem with three main and conflicting objectives: *scalability*, *usability* and *high-assurance*. For high-assurance, we consider solutions based on commercially emerging hardware-rooted Trusted Computing Technology. Further, for SIS we argue super-distribution (“protect once and access wherever authorized”) and off-line access are necessary to achieve scalability and usability¹. We assume all content that are distributed to be read-only.

The three conflicting objectives discussed earlier forces us to consider super-distribution within a group. We define group to be a set of TPM-enabled machines that share a common property. This definition of a group is abstract and powerful enough to accommodate a wide range of real world scenarios. For instance, a group could be tightly-knit that is task oriented (E.g.: project groups) or loosely-knit with similar interests (E.g.: discussion forums) or an ecosystem (E.g.: a network of competitors, suppliers and customers working towards a common goal that benefits everyone). On the other hand, contextual properties like location, proximity, etc. could also form an ad-hoc group.

We explore policy, enforcement and implementation mod-

¹Ultimate proof of scalability, usability and high-assurance of our solutions can only come from commercial implementation and wide-spread deployment and usage. Our solutions are proposed here as steps in this direction.

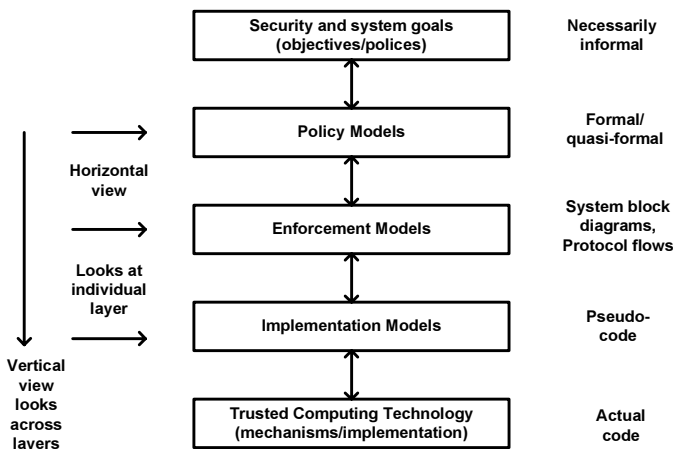


Figure 1: The PEI Models Framework.

els for the group-based secure information sharing problem. To help separate and answer *what* we want to solve from *how* we want to solve, many frameworks have been proposed in the past including the policy-mechanism separation principle. We use the more recent PEI framework [10] that suggests additional separation between enforcement and implementation models which we believe is absolutely necessary to address a complex problem such as SIS.

The remainder of this paper is organized as follows. In section 2 we give a brief overview of the PEI framework and Trusted Computing Technology. In section 3, we develop a policy framework for the g-SIS problem. In section 4, we specify policies for g-SIS under the idealized assumption that instant and preemptive revocation of group membership can take place. In section 5, we identify two approximate enforcement models to accommodate the ideal policy requirements in practical scenarios. We identify trade-offs and develop a usable model to reconcile these trade-offs. In section 6, we outline protocols for our enforcement model using Trusted Computing Technology. We conclude the paper in section 7.

2. BACKGROUND

We give a brief overview of the PEI framework and Trusted Computing Technology. The overviews are necessarily short and we direct interested readers to the references for further details.

2.1 PEI Framework

Many approaches have been proposed in the past to help solve security problems including the traditional approach of separating policy from mechanism. We use the more recent PEI framework [10] (Figure 1) to analyze our g-SIS problem and synthesize suitable solutions. At the highest level of the framework, we specify the overall goals or objectives of the security system. At this level the discussion and analysis is necessarily informal. In the Policy Models (P) level, we formally specify our objectives using an appropriate model. Many models exist including the traditional DAC and MAC, RBAC96 [9], and the more recent UCON. In the Enforcement Model (E) level, we specify system architecture to address the “how” question. In the Implementation Model (I) level, we propose concrete protocol flows and address spe-

cific issues left open in the Enforcement layer. The final layer makes system dependant decisions like technologies to use and produces actual code.

It is important to understand that in PEI the relationship between adjacent layers is many to many. A policy model could have multiple enforcement models and vice versa. Similarly, an enforcement model could have multiple implementation models and vice versa. Thus a horizontal view helps to explore various means to realize the objectives at a particular layer, and the vertical view helps to provide a specific solution.

2.2 Trusted Computing

Trusted Computing Technology is an industry standard proposed by the Trusted Computing Group (TCG) [1]. It is widely accepted that software only mechanisms cannot provide high assurance. This motivated TCG to provide a root of trust at the hardware level through the Trusted Platform Module (TPM). The technology has evolved to a great degree now and we only provide a very brief overview here. The TPM mainly offers three novel features: Trusted Storage for keys, Trusted Capabilities and Platform Configuration Registers (PCR’s). Trusted Storage for keys is provided by encrypting user’s keys with chain of keys. The root key at the top of the chain is stored within the TPM and is not accessible outside the TPM. Trusted Capabilities are capabilities exposed by the TPM that are guaranteed to be trustworthy. Users cannot modify the behavior of these capabilities. PCRs are hardware registers in the TPM that are used to store integrity metrics (hash values) of software (e.g.: boot chain). Trusted Capabilities and PCRs provide some powerful functionalities. For example, Seal is a trusted capability that encrypts and binds some data to a specific PCR value. This data can be accessed (Un-Sealed) by authorized entities only when PCR value at unseal time matches with the specified PCR value at seal time. Using this feature, one can make sure that a data blob such as a key would be available to authorized entities only when the platform (e.g.: every piece of software involved in the boot cycle up to the kernel) is in a trustworthy state (which is implied by the PCR value). There are many capabilities and features provided by TPM and abundant materials have been published by the TCG.

3. POLICY FRAMEWORK

From here on we use the word *object* to refer to information of any form (e.g. documents, voice data, etc.) that belongs to the group and inter-changeably use the words *object* and *document*. We use the word *user* to refer to any entity (machine, human, programs, etc.). We use the word *member* to refer to a user who is enrolled into a group. We now clearly state the objectives.

3.1 Objectives

We specify the following objectives for the g-SIS problem studied in this paper.

1. The objects are read-only. In the example of documents, we assume that documents can only be read and cannot be modified. If modified, the new document will not belong to the group unless explicitly added as a ‘new’ object again.

2. Objects are obtained via super-distribution. Super-distribution means that the protected (and therefore encrypted) object can be accessed anywhere (and by any-means) by an authorized user. The encryption is done in a universal manner for all users rather than customized for each individual user.
3. We want to impose policies on these objects and hence need client side access control.
4. We want to enable off-line access. That is, the member need not be connected to a server while accessing the object.
5. We assume an admin who owns the group. The admin can add and remove members from the group. We do not care how an admin is appointed. The admin may or may not be a member of the group.

We now digress briefly to compare g-SIS with a related problem –broadcast/multicast encryption. Member management in g-SIS scenario sharply differs from Secure Internet Multicast. In multicast, as users join and leave a group, remaining members go through a re-key process thereby refreshing the group key [7]. However, for secure information sharing, such a requirement is extremely un-friendly because members need not be always connected to a server to access the objects. Thus, continuing our list of objectives, we have the following.

6. When a user joins or leaves the group, remaining members should not be affected. In other words, join and leave operations should be completely oblivious to other members.
7. Objective 6 requires that current members should not be forced to be online and go through a re-key process.²
8. Secure multicast is only concerned about forward and backward secrecy [7]. Forward-secrecy means that if a user leaves a group, he should not be able to read group data that are created in the future. Backward-secrecy means that when a new user joins a group, he should not be able to read data created in the past. However, SIS is not limited to forward and backward secrecy of information. Flexible membership policies are to be enforced. When a new user joins the group, whether he can access any group documents created prior to his membership is policy-dependant. Any group document created after he joins the group are accessible. When a member leaves the group, whether he can continue to access all documents that he possesses is policy-dependant. However, he cannot access any documents exchanged in the future.
9. We use Trusted Computing Technology for high-assurance of protection of information and enforcement of policies within a group.

²Members should not be asked to re-key or contact a server to get a new key. Note that re-keying is not an efficient solution in SIS as the member needs to keep track of which document was encrypted with which key. As users join and leave a group, the remaining members will need to go through a re-key process resulting in encrypting documents with different keys along the time line. One cannot discard the old key (as done in multicast) as disseminated documents encrypted with the old key continue to persist.

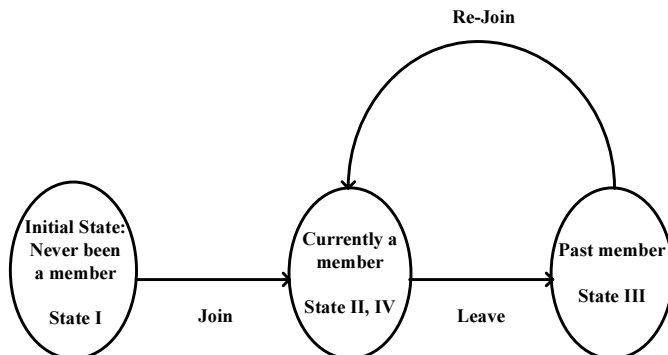


Figure 2: Various states of a subject in a group.

3.2 Member States

Figure 2 shows various states of a member in a group. We identify 4 states. In the initial state (state I), the user has never been a member of the group. In state II, he is a current member of the group. In state III, he left the group and is a past member. Note that a past member in state III can reach state II again by re-joining the group (state IV). Having classified the membership scope, we can define access policies for each state.

In State I, the access policy is straight-forward: the user has no access to any group documents.

In State II (current member), access could be allowed only to current documents (documents created after the member joined the group) or both current and past documents (documents that were created before a member joined a group). For each of these two access policy in state II, we can have rate-limited access (number of accesses allowed per unit time) or usage-limited access. These usage restrictions can be either used for refreshing membership status or for access throttling purposes.

In State III, many kinds of policies could be applied: 1. A past member may lose access to all documents. 2. A past member may access any document created during his membership time period. 3. A past member may be allowed access to the documents he accessed during his membership (the ones local to his machine). On multiple join and leave, accesses could accumulate. 4. A past member can only access the documents he accessed during his membership (the ones local to his machine). But on multiple join and leave, access is allowed only to documents acquired during his latest membership.

For State IV (member rejoin), access policies for members could get complicated. For example, when a member first left the group, he might have been denied access to all (both past and current) documents (policy state III. 1). But when he re-joins, if one allows access to all past documents (policy state II.2), in a sense it contradicts with the policy during his initial membership. These kinds of issues could get complicated over multiple re-join operations and needs to be stated with extra-care. To keep it simple, we suggest having two re-join policies to choose from: no rejoin allowed or when a past member re-joins he will join the group as a new member.

3.3 Motivating Use-cases

We now discuss a few interesting use-cases for the rich set of access policies that we have identified so far. For simplicity, we assume that no re-join is allowed in the following discussion. We consider 4 cases and corresponding use cases for each of these below.

Case 1. Suppose current members (in state II) can access only current documents and past members (in state III) lose access to all documents. We have the following use case.

Memory-less Collaboration: Many universities and corporations allow access to their content and share as long as one is within their network. Once the user leaves the network, the user loses access to the content.

Case 2. Suppose current members (in State I) can access only current documents and past members (in State III) can access documents created during his membership period. We have the following use case.

Collaborative Computing: A financial institution could recruit a software-consulting firm to provide software solutions. This forms a short-lived group. The incoming group members (from the software firm) cannot access any older documents. When they finish the project and leave the group, they can continue to have access to the documents exchanged during their membership in order to add to their profile. This is dependant on financial institution's policy.

Case 3. Suppose current members can access current documents and also the documents created before their membership period and past members lose access to all documents. We have the following use cases.

Employee management: An employee joining a company can access all the current documents. When an employee quits, he loses access to all the sensitive documents he had access to.

Supply Chain: In a supply chain situation, there are lots of partners and lots of suppliers who send quotes for a given proposal. They need to have access to the proposal and related content. But once the quote/response is submitted, their membership context for that particular or group of proposals ceases and they shouldn't have access to any of the older content that they had access to.

Case 4. Suppose current members can access current documents and also the documents created before their membership period and also past members can continue to access documents created before his leave time. We have the following use cases.

Collaborative product development: In the case of several automobile models, there are product twins –models from the same company that resemble each other, except for the division's brand name and price tag. It's less expensive for auto manufacturers to produce parts in bulk and share them than to build separate components for their various brand names. Ford and Mercury, for example, are under the same corporate umbrella, and their Taurus and Sable four-doors are among the company's twins. In such instances, there could be either a loose collaboration (e.g. shared design team, parts ordering/manufacturing but different factories) or a tight collaboration (e.g. joint manufacturing of two different models). In either case, the members from different parties join hands and share documents actively. They will need access to both old documents and current documents. Even after the collaboration period, they will need access to the old documents for further refinement and production.

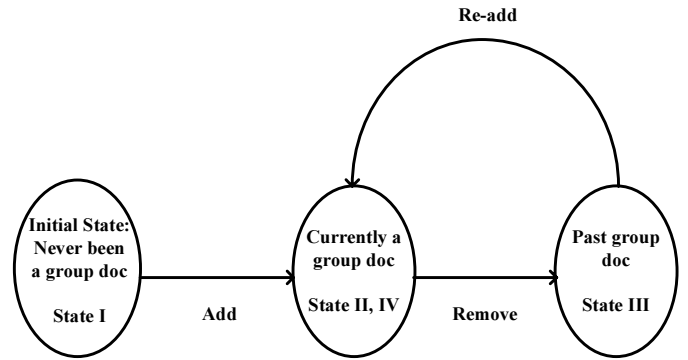


Figure 3: Various states of an object in a group.

3.4 Object States

Symmetric to states that a member goes through, we have a notion of membership of objects or documents. Figure 3 shows various states of a document in a group. We now identify policies for document membership.

For state I, members have no access to documents that are not part of the group. For state II (current document), current documents can be accessed by current members. For State III (past document) we have the following possibilities: 1. No one can access past document. 2. Any member can access. 3. Any one including non-members can access. For State IV (re-added documents), there are several possibilities: 1. Documents cannot be re-added. 2. When a document is re-added, it will be treated as a new document in the group. 3. When a document is re-added, it will be treated as an old document (its history will be preserved).

It is clear that specifying membership policies for members and documents could get complicated. But we believe we have motivated the reader to appreciate the richness of the policy framework for the g-SIS problem, and therefore the necessity for a flexible enforcement and implementation framework to accommodate these variations and switch between them as necessary.

4. IDEAL POLICIES

In the earlier section, we provided an overview of the policy framework for group membership. The policies identified there define the characteristic of a group and we call them as group-level policies or meta-policies. In addition to meta-policies, we have some key policies to consider. *Ideally*, we would want the following policies for g-SIS:

1. A Team Representative (TR, the group admin) adds and removes members and documents into/from the group.
2. Instant and preemptive revocation of a member from a group by the TR takes effect. That is if the member is removed from the group he immediately moves into state III (instant revocation) and the policy with respect to documents currently being accessed by this member are immediately and preemptively adjusted into the state III policy (preemptive revocation).
3. Instant and preemptive revocation of a document from a group by the TR takes effect, similar to the instant and preemptive revocation of a member from a group above.




Figure 4: Enforcement Architecture.

5. ENFORCEMENT MODELS

In this section, we discuss enforcement models for the specified policy framework. We propose an Ideal Enforcement Model that attempts to precisely cover our ideal policy model. We argue that in practical scenarios even the most ideal enforcement model is only an approximation of our policies. We therefore propose Approximate Enforcement Models.

5.1 Ideal Enforcement Model

Recall that our Ideal Policy Model had an important requirement: instant and preemptive revocation of members and documents from a group. Instant revocation means that members and documents should be immediately removed from the group the moment the TR decides to cancel their membership. Preemptive revocation means that should membership status of either the member or the document change during access, access should be stopped. An Ideal Enforcement Model would then be a client-server model. A member should be permitted to access documents only after verifying membership with a server. However, even such a model of “refresh on every access” may not be totally ideal. Any distributed system has an inherent latency and membership verification made at the time of read request may not be valid at the time of actual read by the member. Thus even in an ideal enforcement model the membership validity is approximated to a window of latency time. Moreover, this enforcement model would directly contradict one of our main objectives: *off-line access*. Hence in order to realize our objectives, we discuss ways to approximate our ideal policy.

5.2 Approximate Enforcement Models

Figure 4 shows our enforcement architecture. In order to facilitate various operations, we introduce a Control Center (CC). The CC acts as a single point of contact for both members and the TR. The TR would update membership status of both members and documents at the CC. The CC would then make access decisions based on current membership. Note that a group could have multiple TRs and we assume that members, TRs and the CC are configured for a group.

Steps 1 and 2 are member join operations. In step 1, the TR authorizes the Joining Member (JM) to join the group. In step 2, the CC officially enrolls the JM into the group by verifying the TR’s authorization. In step 3, a member receives approval from the TR to add a document into the group. Step 4 and step 5 differentiates ideal and approximate enforcement models. In an ideal enforcement model,

step 4 and 5 will become a single step of verifying membership with the CC before a document is read by the member. We can approximate this step based on usage count or time. A usage count based approximation would allow the member to access documents a specific number of times without having to contact the CC, thus enabling off-line access with limited usage count. A time based approximation would allow the member to access documents for a specific period of time without having to contact the CC, to provide off-line access with a time limit. Thus in approximate models, step 4 involves off-line document read by a member and step 5 would involve refreshing group membership if either the usage count or time expires. The window provided by the usage count or time is a degree of approximation of currentness of group membership. Steps 6 and 7 involve removal of members and documents from the group by the TR.

A time based approximate enforcement model would require an off-line trusted source of time. We are not aware of any TPMs that provide such a feature. This forces us to elect the usage-count based model as our approximate enforcement model.

Password based, machine based and credential based enforcements models were proposed in a previous work [10]. However, we will show later that this architecture accommodates all of these models.

Our architecture is different from Microsoft Windows Rights Management Services (RMS) in that our motivation here is document access *without* “pre-planning”. Pre-planning is planning in advance which documents one would need to access off-line in the future and “checking-out” those documents by obtaining licenses to access them from a server. Further, in RMS when a client receives a document from some channel, it needs to contact the server and obtain a use license to open the document. Our architecture incorporates the notion of group level policies which not only controls access during membership but also after members leave the group. Our implementation models as discussed in section 6 uses trusted computing based mechanisms to provide strong client-side policy enforcement.

6. IMPLEMENTATION MODELS

In this section, we give a brief overview of our implementation model based on Trusted Computing Technology [1]. As mentioned earlier, we need trust at the hardware level as provided by the TPM for the high-assurance requirements of the SIS problem. The implementation model involves a Trusted Reference Monitor (TRM) module on every group member’s machine. The TRM is a trust-worthy reference monitor that enforces group policies on the client. We also have a Trusted Viewer (TV) module that is used for viewing documents on member’s machines. For our discussion we assume that the TRM and TV are somehow provisioned on the member’s machines in a trust-worthy manner. We can use the mechanisms provided by the TPM [1] to protect the integrity of the TRM and the TV. [8] and [6] discuss practical solutions based on remote attestation using the TPM. It is beyond the scope of this paper to go over the protocols for each of the steps in figure 4. We only outline the protocol with main steps here.

In step 1, the Team Representative (TR) provides a signed credential to the Joining Member (JM) and authorizes joining the group. In step 2, the JM uses this credential to enroll into the group by contacting the Control Center (CC). The

CC verifies the credential and installs a ticket on the JM in such a manner that only the TRM can access this ticket. This is achieved using TPM functionalities like seal, unseal, etc. and integrity measurements as mentioned in section 2. The format of this ticket would be: $signCC\{\{meta - policy||K||refreshMonotonicCount||TSJoin||TSLeave||DRL\}\}$. The meta-policy specifies the group level policies that needs to be enforced. K is the group key. The refresh frequency is defined by $refreshMonotonicCount$. It is the number of times the group key can be used before a membership status refresh is forced. The TPM provides a hardware Monotonic Counter. We use a dedicated virtual monotonic counter that is extended from this hardware monotonic counter to keep track of refresh counts and prevent replay of older tickets. Every time a document is accessed, the dedicated monotonic counter is incremented by the TRM. Once the dedicated monotonic counter reaches $refreshMonotonicCount$, the ticket expires and a refresh is forced. Note that only the TRM should be authorized to update this counter. $TSJoin$ is the time-stamp of member join time as seen by the CC and $TSLeave$ is the time-stamp of leave. DRL is the Document Revocation List that lists the documents that have been removed from the group. This whole ticket is signed by the CC and installed in such a fashion that only the TRM can access it.

The TRM would only let the TV open a document. Documents are always stored on the disk in protected form and is of the format: $signTR\{doc||K\{docKey\}||hash(doc)||TSAdd\}$. doc represents the encrypted version of the document. The document is encrypted with a $docKey$, the document key. The $docKey$ is encrypted with the group key K . $TSAdd$ is the time at which the document was approved to be added to the group by the TR. $hash(doc)$ is the hash value of the document. The TR's signature ($signTR$) shows his/her approval for this document to be added to the group.

6.1 Enforcing Document Read

Here we discuss how the TRM enforces policies for document read based on the ticket and the document format we discussed earlier. The TV requests the TRM to open the document. The TRM initially verifies the integrity of TV. Then the TRM would unseal the ticket. This step would succeed only if the TRM is in the same integral state as it was when the ticket was sealed. First the TRM would check the dedicated monotonic counter with the $refreshMonotonicCount$ and make sure the ticket has not expired. It would then check that the document ($hash(doc)$) is not part of the DRL . Next the TRM would check that the meta-policy permits opening the document based on $TSAdd$ of the document and the $TSJoin$ and $TSLeave$ of the member. If these steps succeed, the TRM would decrypt the document for the TV using K and the $docKey$. Note that if the $refreshMonotonicCount$ has expired, the ticket is no longer valid and the member is forced to obtain a new ticket from the CC. The new ticket would contain updated counter value and updated time-stamps if the membership status has changed.

7. CONCLUSIONS AND FUTURE WORK

In this paper we investigated the Secure Information Sharing problem with three main objectives: scalability, usability and high-assurance. Our approach to solving this problem using the PEI framework has many advantages. The

enforcement and implementation models can be easily enhanced to accommodate interesting scenarios. For example, documents could be password protected by simply adding the password hash to the document and modifying the TRM to enforce password protection. Further passwords could be added on a per-document basis or on the whole group basis. For enforcing a password throughout the group, the password hash could be added to the ticket. The group level policies could be modified to member-level policies by constructing tickets with a different meta-policy for different members. These could be generalized to a credential based enforcement model where the required credential would be mentioned as part of the ticket. The TRM would then mandate such credentials from group members. Credentials could be role certificate, attribute certificate, etc.

A range of future work is underway. First a prototype needs to be built as a proof-of-concept for the group-based SIS problem. We need to explore generality of the policy framework to application scenarios other than document sharing. Restricting information flow across groups needs to be investigated. Support for document querying to obtain specific sections of documents instead of the whole document has many interesting usage scenarios. *Document write* is an important requirement to support true secure collaboration, and hence should be studied.

8. ACKNOWLEDGEMENTS

The authors gratefully acknowledge the collaborative and financial support of Intel Corporation in pursuing this research.

9. REFERENCES

- [1] TCG specification architecture overview. <http://www.trustedcomputinggroup.org>.
- [2] DOD trusted computer sys. eval. criteria. Dec 1985.
- [3] D. Bell and L. LaPadula. Secure computer systems: Unified exposition and multics interpretation. *Technical Report, The Mitre Corp.*, March 1975.
- [4] G. Graham and P. Denning. Protection principles and practice. *AFIPS Joint Computer Conference*, 1972.
- [5] B. Lampson. Protection. In *fifth Princeton Symposium on Information Science and Systems*, 40:437443, 1971.
- [6] J. M. McCune, S. Berger, R. Ceres, T. Jaeger, and R. Sailer. Shamon – A system for distributed mandatory access control. *22nd Annual Computer Security Applications Conference*, December 2006.
- [7] S. Rafaeli and D. Hutchison. A survey of key management for secure group communication. *ACM Computing Surveys*, pages 309–329, September 2003.
- [8] R. Sailer, X. Zhang, T. Jaeger, and L. van Doorn. Design and implementation of a TCG-based integrity measurement architecture. *13th Usenix Security Symposium*, August 2004.
- [9] R. Sandhu, E. Coyne, H. Feinstein, and C. Youman. Role-based access control models. *IEEE Computer*, 29(2):38–47, February 1996.
- [10] R. Sandhu, K. Ranganathan, and X. Zhang. Secure information sharing enabled by trusted computing and PEI models. *Proc. of ASIACCS 2006*, pages 2–12, 2006.