# A Conceptual Framework for Group-Centric Secure Information Sharing

Ram Krishnan
George Mason University
rkrishna@gmu.edu

Ravi Sandhu
Univ of Texas at San Antonio
ravi.sandhu@utsa.edu

Jianwei Niu
Univ of Texas at San Antonio
niu@cs.utsa.edu

William H. Winsborough
Univ of Texas at San Antonio
wwinsborough@acm.org

## ABSTRACT

In this paper, we propose a conceptual framework for developing a family of models for Group-Centric information sharing. The traditional approach to information sharing, characterized as Dissemination-Centric in this paper, focuses on attaching attributes and policies to an object (sometimes called "sticky policies") as it is disseminated from producers to consumers in a system. In contrast, Group-Centric sharing envisions bringing the subjects and objects together in a group to facilitate sharing. The metaphor is that of a secure meeting room where participants and information come together to "share" information for some common purpose. Another metaphor is that of the subscription model where, depending on policy, joining users may or may not be authorized to access past content. We argue that in such contexts, and in accordance with different application use cases, authorizations are influenced by the temporal ordering of subject and object group membership and by the precise nature of membership operations. For instance some subjects may only get future information added to the group while others may also be able to access previously added information. We develop a lattice of models based on variations of these basic membership operations, and discuss usage scenarios to illustrate practical applications of this lattice. Two principles guide Group-Centric models. First, "share but differentiate" which promotes sharing while differentiating user authorizations depending on temporal aspect of membership. Next, "groups within groups" which advocates relationships (such as a hierarchy) between multiple groups. In this paper, we confine our attention to read accesses in a single group.

## 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


*ASIACCS'09,* March 10-12, 2009, Sydney, NSW, Australia.


## General Terms

Security

## Keywords

Information Sharing, Models, Access Control

## 1. INTRODUCTION

This paper introduces the concept of Group-Centric Secure Information Sharing (g-SIS) and provides a conceptual framework to develop a family of models. The traditional approach to information sharing, characterized as Dissemination-Centric sharing in this paper, focuses on attaching attributes and policies to an object as it is disseminated from producers to consumers in a system. These policies are sometimes described as being "sticky". As an object is disseminated further down a supply chain the policies may get modified, such modification itself being controlled by existing policies. This mode of information sharing goes back to early discussions on originator-control systems [9, 13] in the 1980's and Digital Rights Management in the 1990's and 2000's. XrML [1], ODRL [3] and XACML [2] are recent examples of policy languages developed for this purpose. Dissemination-Centric sharing describes in advance the characteristics or properties of subjects who may access the object by attaching "sticky policies" to be enforced when a subject attempts to access the object.

The vision of Group-Centric sharing differs in that it advocates bringing the subjects and objects together to facilitate sharing. The metaphor is that of a secure meeting room where participants and information come together to "share" for some common purpose. This common purpose can range from collaboration on a specific goal-oriented task (such as designing a new product) to participation in a shared activity (such as a semester long class) to subscription to a magazine (where the publisher contributes information that the participants read and possibly respond to content in associated blogs and forums). Visualize a conversation room where users may join, leave and re-join but only hear the conversation occurring during their participation period. For instance, in a Program Committee meeting Alice may be excused from the room when her paper is being discussed and may re-join the room after that portion of the discussion has concluded. In doing so, the conversation that occurred during her absence is not accessible to her. In another setting, all conversations are recorded on a white-

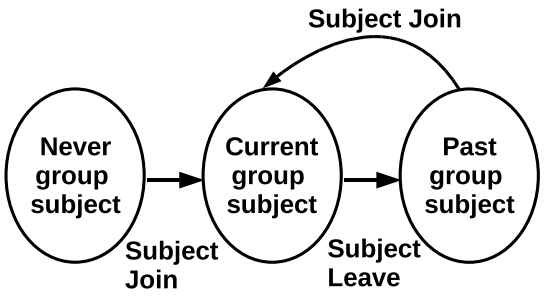


**Figure 1: Subject Membership States.**

board in the room and as Alice re-joins she is able to see what happened during her absence. Such a room may also be appropriate in a different context such as a design group wherein Alice participates as a consultant on demand.

Another metaphor is that of the subscription model where subscribers can receive content depending on when the subscription began. For example, when Alice subscribes to an online news magazine, she may be allowed to access only new content published after she paid for the subscription. In another setting (probably a higher priced subscription), she may also be allowed to access the magazine's archives.

These two metaphors illustrate two important principles in the Group-Centric approach. The first principle is "share but differentiate". As one can see, sharing is enabled by joining and adding information to group. Yet, users' access is differentiated by the time at which they join and the time at which the requested information is added to the group. The second principle is the notion of "groups within groups". That is, in a given g-SIS system, there may be any number of groups. The relationship between these groups can be of any type. One well-known structure is that of a hierarchy, where subjects at a higher level dominate those at the lower levels in terms of read access.

We envision that Dissemination-Centric and Group-Centric sharing will co-exist in a mutually supportive manner. For example, objects could be Added with "sticky" policies in a Group-Centric model. In this case, the objects may have controls imposed by both the Group-Centric model and the "sticky policies". Also, the "sticky policies" on the object could determine whether or not an object can be added to the group in the first place. It may turn out that at a theoretical level whatever Dissemination-Centric can achieve Group-Centric can also achieve and vice versa. But at a pragmatic level, we believe these are significantly different approaches to information sharing.

In this paper, we propose a conceptual framework for developing a family of models for g-SIS. We propose an abstract set of group operations: Join and Leave for subjects, Add and Remove for objects. Subjects may Join, Leave and possibly re-Join the group. This is illustarted in figure 1. Similarly, objects may be Added, Removed and re-Added to the group. Further each of these operations could be of many different types. For example, a Strict Join will only allow a joining subject to access objects added to the group after Join time. But a Liberal Join, in addition, will allow the subject to access objects added before Join time. In general, there may be any number of such variations beyond those explicitly identified in this paper.

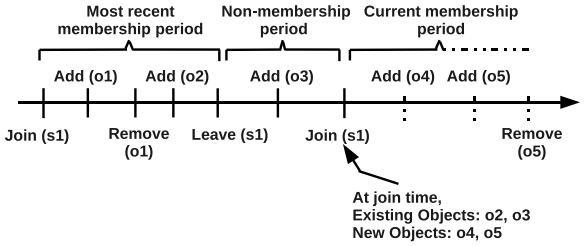


**Figure 2: Subject Operations Illustration.**

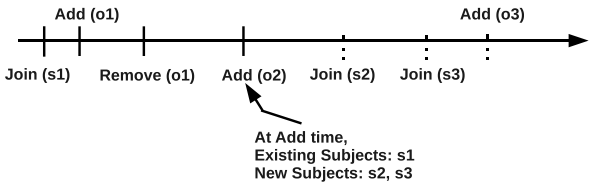


**Figure 3: Object Operations Illustration.**

Authorizations concerning who controls Join, Leave, Add and Remove are addressed as administrative tasks expressed in an administrative model. While a g-SIS administrative model is important, many approaches have been considered in the literature (see [10, 15] for example). Further, the administrative model is likely to be application dependant. Consider two different g-SIS applications, one where users need to pay to join a group and another where users are admitted based on organizational needs. The administrative model for these two applications is likely to be different. Without clear understanding of the operational semantics, an administrative model would be premature.

We believe that authorizations concerning the operational aspects that bear on group membership is a more interesting and novel problem, and this will be the focus of this paper. We leave the development of an administrative g-SIS model for future work. Furthermore, we confine our attention to correct authorization behavior with respect to *read* access in a *single group*. We have developed extensions to other forms of accesses such as *write* or *update* and *multiple groups*. Discussion of these is out of scope for this paper.

## 2. A FAMILY OF G-SIS MODELS

We now discuss a family of g-SIS models based on specific variation of subject and object operations (Join, Leave and Add, Remove respectively). The semantics of variations are based on the temporal ordering of subject and object group memberships. However, there may be any number of additional semantics beyond those identified here.

*Strict Join (SJ) Vs Liberal Join (LJ):* In SJ, the joining subject can access only those objects added after Join time. However, LJ allows the subject additionally to access objects that were added prior to the time of Join. Suppose that in figure 2 the second Join (s1) is an SJ. Then s1 can only access o4 and o5. If the Join was an LJ instead of SJ, s1 can also access o2 and o3.

*Strict Leave (SL) Vs Liberal Leave (LL):* In SL, the leaving subject loses access to all objects. In LL, the leaving subject can retain access to objects authorized prior to the time of Leave. In figure 2, on SL, s1 loses access to all group

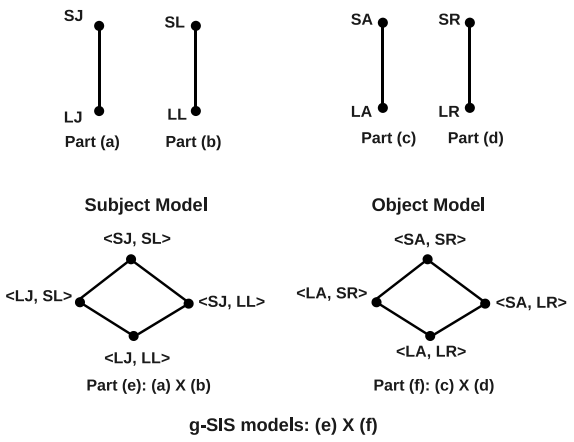


**Figure 4: A family of g-SIS models: *The Cartesian product of Subject and Object Model results in a lattice of 16 g-SIS models with fixed operation types (products are ordered pointwise).***

objects (e.g. o2) authorized during the membership period. An LL will allow s1 to retain access to o2.

*Strict Add (SA) Vs Liberal Add (LA):* In SA, the added object can be accessed only by subjects already in the group. In LA, there are no such restrictions. The added object may be accessed by subjects that join later. If Add (o2) in figure 3 is an SA, only s1 can access the object. Subjects s2 and s3, joining later, cannot access this object. But on LA current subject s1 and future subjects s2 and s3 may access o2.

*Strict Remove (SR) Vs Liberal Remove (LR):* In SR, the removed object cannot be accessed by any subject. In LR, subjects who had access to the object at the time of Remove may continue to access (of course subjects joining later are not allowed to access the removed object. In figure 3, if Remove (o1) is an SR, every group subject (including s1) loses access to o1. If Remove (o1) is an LR, s1 can continue to access o1. However s2 and s3 will not have access to o1.

## 2.1 g-SIS Lattice

Let us first consider g-SIS models where the group operations are fixed for all subjects and objects. For example, a g-SIS model may only allow Liberal operations for all subjects and objects (LJ, LL, LA, LR) in the group. That is, every subject that is admitted to the group will be given LJ. Similarly, objects will be added only with LA and so on. Another example of a fixed operations model is (SJ, SL, SA, SR) where all operations are Strict.

Thus if the type of operations are fixed for all subjects and objects, there are 16 possible models ranging from the most restrictive model allowing only Strict operations: (SJ, SL, SA, SR) to the most permissive model allowing only Liberal operations: (LJ, LL, LA, LR). This is illustrated in figure 4. Parts (a) through (d) show that the Strict operation is more restrictive than the Liberal operation. Parts (e) and (f) show the subject and object model that is obtained by the Cartesian product of subject and object operations respectively. Finally, a lattice of 16 g-SIS models can be obtained by a Cartesian product of subject and object models (parts (e) and (f)). Due to space constraints, the final lattice with 16 models is not shown. Note that if the most restrictive model (SJ, SL, SA, SR) permits a subject to access an object, the most permissive model (LJ, LL, LA, LR) should also grant access to the same object.

An authorization policy can be formally specified for each of these 16 models that specify the conditions under which a subject may access an object. For example, authorization will succeed for the (SJ, SL, SA, SR) model in any state, if the requested object was added after the subject joined the group and both the subject and object are still current group members at the time of request. This can be precisely specified using Linear Temporal Logic (LTL) [12] as follows:

Definition 2.1 (Most Restrictive Specification). *A g-SIS specification is Most Restrictive if it satisfies the following LTL formula:*

$$\Box(\text{Authz} \leftrightarrow (\neg\text{SR} \wedge \neg\text{SL})\ \mathcal{S}\ (\text{SA} \wedge (\neg\text{SL}\ \mathcal{S}\ \text{SJ})))$$

The Most Restrictive g-SIS Specification only allows Strict operations. The above formula says that a subject is authorized to access an object if and only if both are still part of the group since the object was added (indicated by $\mathcal{S}$ temporal operator). Also, when the object was added, the subject was a current member of the group. Because of SJ and SL, we only need to consider the case where an object is added after the subject joins the group since subjects cannot access objects added prior to their join time. The $\Box$ operator says that the formula should hold in every state.

A highly flexible g-SIS model could simply allow different types of operations on a case by case basis. For example, SJ for s1, LJ when s1 re-joins, LJ for s2, LL for s1, SL for s2, SJ when s2 re-joins, etc. (similarly for objects). In this case, we would have one all encompassing specification.

# 3. USAGE SCENARIOS

We now discuss two usage scenarios: a large-scale subscriptions scenario where the operations are fixed for all subjects and a small-scale collaboration scenario where the operations could be mixed.

## 3.1 Subscription Service

In general, subject operations define the semantics of most subscription models. Thus most subscription models fall into one of the four categories: (SJ, SL), (SJ, LL), (LJ, SL) and (LJ, LL). Consider a premier online news magazine ABS Corp. that offers four levels of membership:

1. Level 1; $10/year (SJ, SL): These subscribers can access news articles that are published after they started paying the subscription fee. Level 1 subscribers cannot access ABS's archives (effected by SJ). If they cancel their subscription, they lose access to all news articles.

2. Level 2; $12/year (SJ, LL): Similar to Level 1 but subscribers can retain access to news articles that they paid for even after canceling their subscription.

3. Level 3; $15/year (LJ, SL): Level 3 subscribers can access rich archives filled with post-news analysis, predictions, annotations and opinions from experts, in addition to future articles. But if they cancel their subscription, they lose access to everything including archives.

4. Level 4; $17/year (LJ, LL): Similar to Level 3, but even after canceling membership, subscribers can login and view all articles that they had access before leaving.

Object operations do not fundamentally change the subscription model's semantics. Nevertheless, they model useful scenarios. For example, if an object is added with SA, only existing subjects in the group may access. Thus SA objects model sales promotion or discounted price available only to current group members.

### 3.2 Mission Oriented Group

Consider a g-SIS model with the operation types: (LJ, SL, SA/LA, SR) where all operations are fixed except object Add. Objects can be added to the group by type SA or LA. Let us consider a simple collaboration scenario where the group is mission oriented, so many users may Join and Leave the group to contribute and receive information over time.

Consider two subjects Alice and Bob who Join the group at the same time. If Bob wants to ensure that any information he shares with Alice is not accessible to future subjects who may Join the group, he can add objects with SA. SA objects are only accessible to *existing* members at Add time. This allows current members of the group to share information privately. On the other hand, to the mission's end, information can be made available to future subjects by LA'ing objects to the group. Suppose Alice leaves the group and later Cathy joins with LJ. Cathy cannot access SA objects shared between Alice and Bob before her join time. Cathy can only access existing LA objects.

## 4. RELATED WORK AND CONCLUSION

Older approaches to Secure Information Sharing (SIS) can be classified into at least three categories. First is Discretionary Access Control (DAC) [8, 11, 7] which proposes to enforce controls on sharing information at the discretion of the "owner" of the object. Although, this is similar in objective to SIS, DAC fails to solve the problem since it does not correlate the controls on copies of information with copies of the original. The second is Mandatory Access Control (MAC) [5, 6, 7] which allows information to flow in one direction in a lattice of security labels. Copies of information made from one or more objects inherit the least upper bound of the labels from the individual objects. Thereby the copies are controlled at least as strictly as the original. Historically, one directional information flow has not been the most common requirement of SIS. The third is Originator Control or ORCON [9, 13] in which the owner of the object decides which user(s) may have access to it. The owner is the principal source of the policy to be enforced. As information flows from one container to another, the policy is also propagated. In other words, it is a "sticky policy". Recently, information sharing challenges have been considered in the context of Dynamic Coalition Problem or DCP (see [14, 4] for example). The DCP is concerned with the challenges involved when a coalition is dynamically formed, for example, in response to a crisis. Government, civilian and other commercial organizations may need to form a coalition (who may otherwise distrust each other) and share information quickly to solve the problem at hand. Our approach to information sharing is primarily different in that it focuses on authorizations involving the temporal aspect of group membership.

In this paper, we proposed a Group-Centric family of models for Secure Information Sharing. We identified a few useful variations of group operations whose semantics are temporal in nature. The framework can accommodate additional semantics that may turn out to be useful beyond those identified in the paper. In future papers we will report formalization of these models and identification of a layered set of g-SIS properties.

## 5. ACKNOWLEDGMENTS

The authors are partially supported by NSF grants IIS-0814027, IIS-0814027, CCR-0325951, CCF-0524010 and CNS -0716750, AFOSR grant FA9550-06-01-0045, THECB ARP grants 010115-0037-2007 and 010115-0037-2007 and grants from the State of Texas Emerging Technology Fund and Intel Corporation.

## 6. REFERENCES

[1] eXtensible rights Markup Language. *www.xrml.org*.
[2] OASIS eXtensible Access Control Markup Language . *www.oasis-open.org/committees/xacml/*.
[3] Open Digital Rights Lang. Initiative. *www.odrl.net*.
[4] V. Atluri and J. Warner. Automatic Enforcement of Access Control Policies Among Dynamic Coalitions. *International Conference on Distributed Computing & Internet Technology, Bhubaneswar, India, Dec*, 2004.
[5] D. Bell and L. LaPadula. Computer Security Model: Unified Exposition and Multics Interpretation. *MITRE Corp., Bedford, MA, Tech. Rep. ESD-TR-75-306, June*, 1975.
[6] D. Denning. A Lattice Model of Secure Information Flow. *Comm. of the ACM*, 19(5):236–243, 1976.
[7] DoD National Computer Security Center (DoD 5200.28-STD). *Trusted Computer System Evaluation Criteria*, December 1985.
[8] G. Graham and P. Denning. Protection-principles and practice. *Proceedings of the AFIPS Spring Joint Computer Conference*, 40:417–429, 1972.
[9] R. Graubart. On the Need for a Third Form of Access Control. *Proceedings of the 12th National Computer Security Conference*, pages 296–304, 1989.
[10] M. A. Harrison, W. L. Ruzzo, and J. D. Ullman. Protection in operating systems. *Comm. of the ACM*, pages 461–471, August 1976.
[11] B. Lampson. Protection. *ACM SIGOPS Operating Systems Review*, 8(1):18–24, 1974.
[12] Z. Manna and A. Pnueli. *The Temporal Logic of Reactive and Concurrent Systems*. Springer-Verlag, Heidelberg, Germany, 1992.
[13] C. McCollum, J. Messing, and L. Notargiacomo. Beyond the pale of MAC and DAC - defining new forms of access control. *Proc. of the IEEE Symposium on Security and Privacy*, pages 190–200, 1990.
[14] C. Phillips Jr, T. Ting, and S. Demurjian. Information sharing and security in dynamic coalitions. *Proceedings of the 7th ACM Symposium on Access Control Models and Technologies*, pages 87–96, 2002.
[15] R. Sandhu. The typed access matrix model. In *Proceedings of the IEEE Symposium on Security and Privacy*, page 122, 1992.

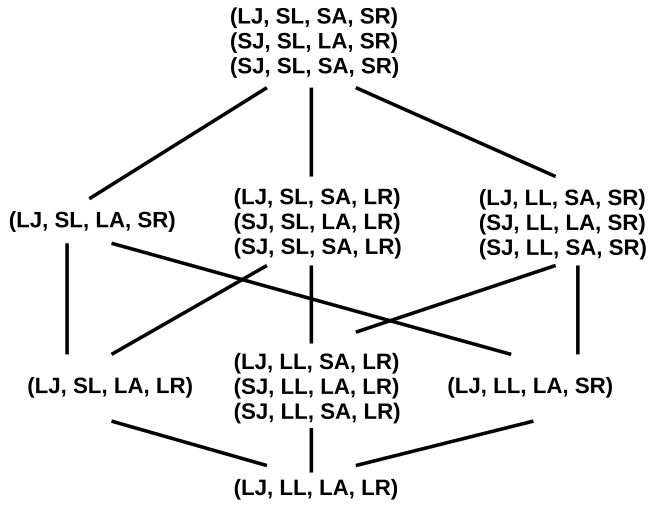


**Figure 5: Reduced lattice with 8 fixed operation specifications.**

# APPENDIX

## A. LATTICE REDUCTION

The following observations can be made about 16 possible fixed operation g-SIS specifications in figure 4:

1. The type of object add has no significance in a model allowing only SJ. This is because, with SJ, joining subjects can only access newly added objects. Thus regardless of how the object is added, a joining subject cannot access objects added prior to join time. And if the object is added after join time, the subjects can access the object regardless of how it is added.
2. Similarly, the type of subject join has no significance in a model allowing only SA.
3. LR has no significance in a model with SJ. LR allows subjects who had access at the time of remove to retain access. These are subjects who joined with SJ prior to the object add time. Subjects who joined with SJ after the removed object's add time had no access to begin with. Thus an SJ model supporting only LR is as though the model has no support for remove operation.
4. Similarly, LR has no significance in a model with SA.

Thus, based on the observations 1 and 2 above, we only have 8 unique g-SIS models with "fixed" operations as illustrated in figure 5 since:

1. (SJ, SL, SA, SR)=(SJ, SL, LA, SR)=(LJ, SL, SA, SR)
2. (SJ, SL, SA, LR)=(SJ, SL, LA, LR)=(LJ, SL, SA, LR)
3. (SJ, LL, SA, SR)=(SJ, LL, LA, SR)=(LJ, LL, SA, SR)
4. (SJ, LL, SA, LR)=(SJ, LL, LA, LR)=(LJ, LL, SA, LR)
5. (LJ, SL, LA, SR)
6. (LJ, SL, LA, LR)
7. (LJ, LL, LA, SR)
8. (LJ, LL, LA, LR)

Thus the type of Add has no significance on the authorization in these 8 specifications. In most usage scenarios, object operations do not remain fixed for all group objects. Certain objects may need to be added with SA to restrict access to existing group members while others may be added with LA when such a restriction is not required. Similarly, certain objects may be removed with SR while others with LR. For instance, in the four subscription models discussed in section 3 ((SJ, SL), (SJ, LL), (LJ, SL) and (LJ, LL)), the subject operations remain fixed for all joining and leaving subjects in their respective membership level. However, object operations may differ from one object to another. On the other hand, in many dynamic scenarios such as in emergency response or clinical work flow systems, even the subject operations may differ from one subject to another. For instance, physicians may be given an LJ so they have access to all the patient records. However, nurses are rotated in shifts and thus may be given SJ and SL. Thus the nurses get to see patient information that is pertinent during their shift period.