Paper ID# 901011

MULTI-DOMAIN TRUST MANAGEMENT IN VARIABLE-THREAT ENVIRONMENTS – A USER-CENTRIC MODEL

Mike Burmester
Florida State University
Tallahassee, FL

Prasanta Das, Martin Edwards
The MITRE Corporation
McLean, VA

Alec Yasinsac
University of South Alabama
Mobile, AL

ABSTRACT

Trust Management (TM) systems are trust infrastructures that support authorization for security-critical actions in decentralized environments. In this paper we present a user-centric view to address trust management as it impacts the unanticipated user and/or user behavior for multi-domain applications. This protection can be tuned to deal with users who may be responsible for an elevated threat level, and builds upon a resource-centric architecture.

Our model is suitable for variable-threat environments and allows for temporary adjustments of trust levels. The expectation is to enable a Trust Management Agent to determine appropriateness of the unanticipated user or behavior, and reverse restrictions without compromising actions that took place during such periods —we term this, rollback-access. We argue that a rollback-access capability is an essential feature for security-critical applications, and is appropriate for today's military and intelligence community coalitions as they execute their particular missions in the Global War on Terrorism.

1. INTRODUCTION

To support Coalition Information Sharing (CIS) network capabilities, there is a need to establish a Trust Management (TM) infrastructure among members of the coalition. This infrastructure may be ad hoc and dependent on validation of the identity of a member by others within a trust community. To meet this need requires the U.S. to establish trust relationships with coalition partners.

This paper will investigate how a capability, that we term rollback access (RA), can be used manage trust for increased or decreased functionality across the Global Information Grid in support of multinational information sharing in a net-centric environment from a user-centric view. Our goal is to establish dynamic and flexible trust mechanisms for complex coalition environments that address user-centric threats. This work extends and complements earlier work which considered the threat model from a resource-centric view [7].

The approach is a fundamental change in the nature of trust. In most existing TM models, trust is static: users are either trusted to access a resource or untrusted, depending on their clearance level and the security level of the resource. Coalition trust environments have inherent complexities such as dealing effectively with conflicting trust information that cannot be accurately captured in fixed value trust models. A variable trust valued model that reflects the dynamics of the local (or global) network as impacted by (perceived or actual) adversarial actions better reflects the myriad of subtleties that characterize modern coalitions and consortiums (e.g., trust, mistrust, malicious hosts, insider attacks, passive adversaries, sleeper cells, etc.)

This paper presents an approach to develop a working prototype of a dynamic TM system with RA functionality for distributed systems that enables user-centric trust mechanisms for accessing system resources by unanticipated authorized users while denying access to unauthorized users. This can be combined with the model in [7] to support both user-centric and resource-centric trusted services.

Background. There is extensive work in the literature on modeling access control and TM systems. Access control can be discretionary, mandatory or role-based [2, 3, 10, 13]. Recent work focuses on decentralized TM systems (see e.g., [5]), and on flexible TM systems that are appropriate for open network applications (e.g., [1]). Several papers address implementation issues (e.g., [12, 6]). Role-based systems such as RT [11] combine the flexibility of role-based access control (RBAC) with the strength of TM.

Our contribution. In today's environment of asymmetric warfare and homeland security, the formation of coalition partnerships among governmental and non-governmental organizations within the U.S. as well as U.S. collaboration with international partners to share information is essential. The premise of such information sharing among ad hoc domains is that it is on a need-to-share basis for security critical missions. Information sharing is based on trust policies that determine the internal or external trust value sets which enable the two sides to establish trust so they can interact.

The trust relations among ad hoc coalition partners

introduce additional security uncertainty based on changing external and internal threat levels. Since the coalition networks may contain multiple domains the ability to share information across these domains is paramount. The information trustworthiness will depend on the security implementations within each domain. The dynamic trust model is planned to be able to adjudicate between these varying domains and broker the appropriate access based on the perceived and calculated threat levels. Therefore, a static TM model based on a fixed set of criteria is less applicable and prone to security vulnerability. We present a dynamic TM model that addresses access requests by unanticipated users or unexpected user behavior in a flexible way, while at the same time dealing with potential access abuses (e.g., identity theft, impersonation attacks, spoofing attacks, etc). Anomalous or unanticipated behavior by an authorized user is checked by a rollback access (RA) service that may restrict the user's access until their mission partners establish (validate) the trust level of the user for the particular mission.

User-centric vs resource-centric trust management. For security critical applications in variable threat environments access actions should be managed so as to protect both users and resources. In the earlier [7] model, access to resources may be denied if this is deemed to be a threat to the system when the prevailing threat level is elevated, even though such access would otherwise be granted. The user-centric model focuses on users and addresses potential threats to the system posed by unanticipated user ID accesses, impersonation and spoofing attacks. It provides a rollback access functionality to manage suspended actions as well as support an access enabling functionality for service flexibility (for a forgetful Alice, Section 3.1). The protection is independent of resource protection, and can be tuned to deal with users who may be responsible for the elevated threat level. The model assigns a trust value to the relation between a user ID and the actual user accessing the service. For $vanilla$ trust, access is restricted to basic services, whereas for $high$ trust all services normally available to the user can be accessed. On the basis of trust assessment a user is provided vanilla access initially and as the trust level is validated by the trust agent the user's access increases to the higher trust levels. If the trust agent is unable to validate the user ID with the actual user, the access will remain unchanged or vanilla.

2. SCENARIOS

To motivate our methodology and the rollback access functionality we consider two scenarios. The first scenario stresses the need for a dynamic, open-door functionality, while the second underlines the intricacies of managing rollback-access.

2.1. Scenario A

The U.S. President through the Department of Homeland Security (DHS) and the Office of the Director of National Intelligence (DNI) has established a policy where counter-terrorism information is to be shared to the greatest extent possible. The point of the policy is to ensure all participants in the national counterterrorism effort are provided the most accurate and current information available.

Operational environment. Key to the implementation of this policy is the reality that establishing a single network that all the responsible agencies within federal, state, local and tribal organizations is too expensive and will take too long. Therefore the implementation guidance stresses the need for all data/information producers and owners to instantiate "open door" capabilities to their networks and data stores. The policy does not call for wholesale exposure of data and information, but for open visibility of data and information to those needing the information to execute their day-to-day mission.

Trust Model. The intelligence community is ready to embrace this policy. It is, however, looking for reassurance that each agency ensures that the information provided is only made available to those with an established role on a need-to-share basis (e.g., information needed to execute their mission). As an agency system is monitoring the use of its data sources the system is alerted by the Trust Management Agent (TMA) that a user ID is making data requests to a data source that are very different from those its assigned role causes it to normally make. Automatically, the TMA rolls back the user ID's access to $vanilla$, which provides minimal information. The concern is that a bad actor has entered the enclave by spoofing the rightful owner of the user ID. The TMA analyzes the aberrant behavior based on the established network threat level and validates the role and "normal" accesses associated with the user ID. If the TMA validates the user ID's role was expanded due to new mission requirements, then the TMA identifies the change of the role to the System Administrator for final (human-in-the-loop) confirmation. The System Administrator confirms that the user ID's roles were expanded to require access to the identified data source due to the transfer of personnel and an inability to replace the individual for an extended period. Based on this confirmation, the TMA rolls the user ID's accesses back to the level appropriate for the validated roles. If however, the TMA determines the user ID's role was not changed the TMA would further reduce the system's access for the user ID to affect not only the area of new requests, but also areas formerly open to the user ID. The System

Administrator would be alerted to allow for the appropriate security actions to be taken. The monitoring of this trust by the TMA must be dynamic and able to respond to the changing needs of tactical missions.

2.2. Scenario B

A coalition of 12 national militaries (e.g., U.S., Germany, Belgium, France, U.K. etc), governmental agencies (e.g., DOS, DOE, CIA) and non-governmental organizations (NGOs) (e.g., Médecins Sans Frontières, American Red Cross, UNICEF, Red Crescent) are involved in a stabilization and humanitarian relief effort in Orange Land, a sub-Sahara nation, in the midst of inter-tribal conflict and a three year drought. The Orange Land government is generally pro-west, but there are at least two factions within the government that have ties to terrorist organizations through their rhetoric and tribal affiliations.

Operational environment. A tactical wide area network was established to support the coordination and cooperation in all facets of coalition operations. As such, each national military and governmental agency as well as the NGOs is on the network with common access based on attributes associated with the group. The military consistently presents information on insurgent locations and dangerous areas (e.g., improvised explosive device (IED) locations) to allow non-military group use of the information for safety and planning. Additionally, the military provides time-lines for general operations that will go force-on-force with insurgents to ensure the non-military efforts are not caught up in these operations, which could result in civilian casualties.

Key to the level of information sharing provided is the trust established between the organizations that information would be available to each group, but groups would not share between themselves the information. Thus each user has its own "information container" which prevents cross-talk, but allows for coordinated approaches to resolving issues. Over the last months this trust relationship has allowed the military to successfully eliminate a number of insurgent strongholds and clearly map the IEDs planted. Most of the IEDs were destroyed, but some are still in areas too "hot" to get into, but the military is planning operations to solve that.

A trust problem. As operations continue the TMA detects efforts by a recently added user ID to post information about locations of IEDs as well as possible "bad guy" locations into the general information stores location. The TMA is unable to determine the association of this user ID with other members of the coalition.

Rollback access capability. The TMA allows the new user ID to post the information into a "safe zone" where it is visible but has caveats as to its accuracy. The TMA then queries other members of the coalition to determine if the new user ID is recognized and should have an ability to post this type of information. The TMA is doing the identity validation through automated means such as using voting buttons in the request. The community responds that the new user ID is from an NGO newly assigned to the coalition and deployed into a different geographical location than other members. With this confirmation of the user ID's identity the TMA rolls back the user ID's accesses to the data stores such that they are the same as other members of the coalition, moves the information provided into the normal data stores environment, and removes the caveats. These actions allow the new user ID access to the data store so the information can be updated directly without going through the "safe zone" process. Safe zones are implemented by using information containers.

Information containers. Managing information containers in dynamic networks is a major challenge, particularly if we allow for granularity. There should be a clear separation between the different instantiations of information compartments: e.g., it should not be possible to write to earlier instantiations unless/until the trust level justifies this.

3. OUR APPROACH

3.1. Rollback access: a state of suspension

Our approach is based on, and extends the human-centric TM model in [8] with rollback access (RA). This model exploits the effectiveness that humans have in understanding their roles in their peer communities to support access services when traditional mechanisms fail.

Suppose for example that a user, Alice, has forgotten her password or pin. RA will still allow her access to some basic vanilla services, for short periods. The system will then contact her peer-community and if sufficient trust is mustered, Alice will get full access. An important enabling feature of RA is that if vanilla-Alice logs out before successful authentication is accomplished, the session manifestation can be maintained in a suspended state, neither committed nor discarded. If the questionable session is later authenticated, all manifestations can be triggered and the system state updated as though the actions were taken at the time they were initiated by vanilla-Alice. Conversely, if an impersonation attempt is recognized, RA restore mode can revert the system to its original state, essentially rolling back all changes that vanilla-Alice performed. RA is effectively a user-centric escrow recovery mechanism.

In our approach we shall use such a mechanism not only

for a forgetful Alice, or an Alice who exhibits unanticipated behavior (e.g., an Alice who was inactive for some time, an impersonator-Alice, a hacker-Alice, etc), but also for a tactical mission Alice who needs the support of her partners to access resources supporting her mission.

3.2. Our model

We build our model on a TM system that provides adequate flexibility: e.g., a TM based on credentials [4, 9], or roles [11]. For our purpose it is sufficient that the trust will support our additional functionality. TM systems provide a unified approach in specifying and interpreting security policies, credentials and relationships. Their functionality is to *authorize* actions of entities (individual users or processes).

Let TM^{auth} be the authorization functionality specified by the TM system. We say that TM^{auth} *realizes* the TM system. In our model, the functionality TM^{auth} is restricted by the trust level of the entity $U \in \mathcal{U}$ that invokes it. If ϕ is the trust level of U, the restricted functionality for U is denoted by $TM^{auth}_{U,\phi}$. Trust levels for entities can be affected by the local (domain) or global state of the system, and may be linear or non-linear. For simplicity we consider a linear trust level structure $(\Phi, \succeq)$ with two distinguished values: $vanilla$ and $high$. A third $mission$ value may also be used for access to data needed to support a specific mission.

We denote the set of TM systems that Φ induces on TM by,

$$\mathrm{TM}_\Phi = \{\mathrm{TM}_{U,\phi}\}_{U\in\mathcal{U},\phi\in\Phi},$$

and call it, a multi-domain TM system with rollback access. TM_Φ is realized by the functionalities: $TM^{auth}_{U,\phi}$, where $U \in \mathcal{U}$ is an entity and $\phi \in \Phi$ the trust level of that entity.

There is a natural dominance relation "$\succeq_{auth}$" between the TM systems $TM^{auth}_{U,\phi}$ in which: $TM^{auth}_{U,\phi_1} \succeq_{auth} TM^{auth}_{U,\phi_2}$, if every action authorized by TM^{auth}_{U,ϕ_2} is also authorized by TM^{auth}_{U,ϕ_1}. In this model the relation "$\succeq$" between trust levels ϕ of the entity U is proportional to the TM dominance:

$$\phi_1 \succeq \phi_2 \quad \Rightarrow \quad TM^{auth}_{U,\phi_1} \succeq_{auth} TM^{auth}_{U,\phi_2}\,. \qquad (1)$$

Consequently by raising the trust level of an entity, authorization is extended until eventually it is fully restored. Conversely by lowering the trust level, authorization is restricted until eventually it is reduced to $vanilla$.

3.3. The rollback access service

Assign to each user U a *trust level* $\phi \in \Phi$ and to each resource an *access trust threshold* $\psi \in \Phi$. The trust level ϕ is determined dynamically by several factors that relate to the perceived threat U may pose to the system (based on the relation between the user ID and the actual user accessing the system, as well as unanticipated user behavior) at a given point in time and location (domain), and managed by a Trust Management Agent (TMA). The access trust thresholds ψ for the network resources are determined independently and relate to the nature of the resources.

Resources with threshold access level ψ can only be accessed by users whose current trust level ϕ is at least as large as ψ. For example if the current trust level of a user's ID is $\phi = vanilla$ than that user can only access resources for which $\psi = vanilla$, and which the TM system (the MAC and DAC) will authorize.

When the trust level ϕ of user U is lowered to ϕ^-, rollback access (RA) is triggered and the functionality TM^{auth}_{U,ϕ^-} is invoked: actions that are executed while the trust level was ϕ, and which are not authorized by the new functionality, get suspended (RA: suspend mode) and a record of their partially executed state is temporarily stored (for later retrieval).

To capture this we introduce the concept of an *information container* (IC). For each user U and trust level ϕ we define the container $IC_{U,\phi}$. $IC_{U,\phi}$ is a (logical) memory block in which are stored records of partially executed actions that get unauthorized when the trust level of U is lowered to the next level below ϕ. In particular, when the execution of an access action is suspended because the trust level of U is lowered from ϕ to ϕ^-, a record of its suspended state is stored in $IC_{U,\phi}$. In general, several actions may be suspended when the trust level is lowered involving intermediary containers IC_{U,ϕ^*}, $\phi \succ \phi^* \succeq \phi^-$.

If the trust level is later raised to ϕ, then the TMA will rollback the access to those records of suspended executions of access actions in $IC_{U,\phi}$ that get authorized by the new functionality. We describe these two actions in more detail below. Initially, $IC_{U,\phi} \leftarrow \emptyset$ for all users $U \in \mathcal{U}$ and trust levels $\phi \in \Phi$.

Rollback access, suspend mode: $\phi \to \phi^-$

1. Put in the information container IC_{U,ϕ^*}, $\phi \succ \phi^* \succeq \phi^-$, a record of every suspended access action α requiring trust level ϕ^*.
2. Invoke the functionality $TM^{auth}_{\phi^-}$.
3. Every object β produced while the trust level is ϕ^- is assigned the trust value ϕ^- (in addition to the classification of the underlying TM system).

Rollback access, restore mode: $\phi \to \phi^+$

1. All records in IC_{U,ϕ^*}: $\phi^+ \succeq \phi^* \succeq \phi$, authorized by the new functionality $TM^{auth}_{\phi^+}$ get restored: they get labeled as objects with trust value ϕ^*, and removed from IC_{ϕ^*}.
2. Invoke the functionality $TM^{auth}_{\phi^+}$.

3. Every object β produced while the trust level is ϕ^+ is assigned the trust value ϕ^+.

3.4. Architecture

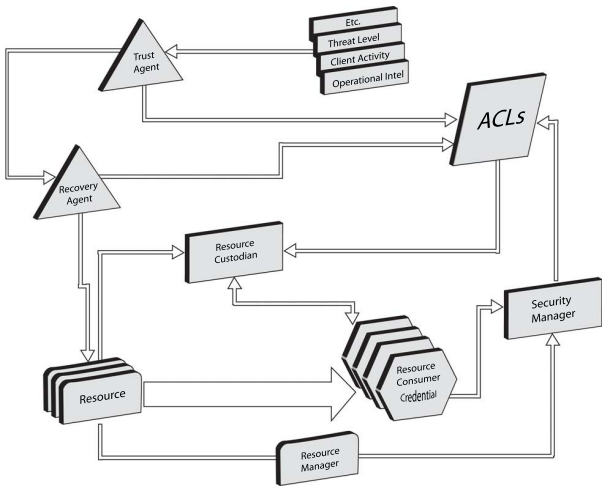

Figure 1: An architecture for variable threat TM system with rollback access functionality.

Variable threat TM systems recognize that the trust relation between a user ID and the expected behavior of the user may be weak, and is not binary. In the extreme case when the user cannot substantiate adequately this relation (e.g., forgot the password, lost the ID card, failed the biometric scan, exhibits erratic behavior, is an impersonator, an insider, etc) the relation is reduced to $vanilla$. If later the relation gets partially substantiated the trust level is raised allowing access to more services. Conversely when the relation is weakened, services get suspended. The model is flexible and adaptive, and allows for an access action in progress to be suspended. This procedure is managed by the RA functionality. To minimize the loss of service, partly executed actions are stored in information containers. RA in restore mode will make these available when the trust relation is restored.

The architecture of a variable threat TM system with rollback access functionality consists of:

- A Trust Management Agent (TMA), a Security Manager and a Resource Manager
- A Resource Custodian and a Recovery Agent
- Environmental inputs (e.g., client activities, threat level, operational intelligence, etc.)
- Access Control Lists (ACLs)
- Resource Owners, Resources, Resource Consumers.

Based on the environmental conditions, and the user behavior, the TMA provides input to the ACL system in terms of the user trust level $\phi \in \Phi$. The Security Manager coordinates with the Resource Manager to properly represent desired resource security properties in terms of entity credentials and the ACL system. The Resource Custodian makes access decisions based on the credentials presented, the ACL system and the access trust threshold ψ of the resource, to provide access to the client.

Operationally, the TMA of a variable threat TM system executes similarly to the traditional Bell-LaPadula model. When the trust level of a user ID is low, or lowered due to unanticipated behavior and/or other environmental conditions, the TMA modifies the corresponding ACLs according to the new restrictions that come into play for each of the resources and the resource consumers. The Resource Custodian then starts to rollback access (RA: suspend mode) to those resources whose threshold level is breached. Rollback may include documents and data objects that the resource consumer was an author of, so that they are not able to modify and possibly even read the object at this low trust level. In those instances where operations must continue a new time stamped version of the data object may be established which allows modification within the new trust level.

During this time, when the trust level is reduced, the Recovery Agent is assessing the status established by the TMA to determine whether conditions have changed such that resource consumers can have their access rollback (RA: restore mode) to its original openness. As the environment returns to "normal operations" and the user behavior is ascertained, the user trust level is raised and the TMA returns to the resource consumer the visibility and modification rights previously enjoyed. Additionally, the Recoverey Agent will assess whether the information introduced on the new time stamped version of the data object is valid and acceptable to be consumed by the earlier data object. Thus work done while the user trust level was diminished is preserved but also adjudicated prior to wholesale acceptance as valid.

4. A WORKING PROTOTYPE

A proof of concept prototype will be developed to embrace the TM functionality described herein. The prototype uses a variable threat TM system with a linear trust level structure: $\Phi = (high \succ \phi_a \succ \phi_b \succ \cdots \succ vanilla)$. We only discuss the additional RA functionality.

4.1. Rollback access

Assign to each user U a variable trust level $\phi \in \Phi$ (depending on user behavior), and to every object α produced by U the access trust threshold ϕ. Any access operation β on α

requires a trust level of at least ϕ. The value of ϕ reflects the vulnerability of β to external/internal threats.

1. If the trust level of U is reduced below ϕ while β is executed, this action is suspended: an object $suspend(\beta)$ is generated and assigned the access trust threshold ϕ.
2. If (later) the trust level of U is raised back to ϕ, then $suspend(\beta)$ becomes available and the execution of β can be completed, provided this is authorized by the TM system (e.g., by its owner, or anybody assigned access by the owner). Objects produced while the trust level of U was below ϕ get reassigned the access trust threshold ϕ.

4.2. Compatibility

The RA capability will support the security of the underlying TM system (which is based on controlling information flows—the simple security property [2]) because by (1) we have: $\phi^+ \succeq \phi \ \Rightarrow \ TM^{auth}_{\phi^+} \succeq_{auth} TM^{auth}_{\phi}$.

4.3. Example

Suppose the trust level of Alice is $high$, the trust threshold of access β to resource α is $high$, and that Alice has TM authorization for β. Then Alice has authorized β-access to resource α. If the trust level of Alice is (later) reduced to $vanilla$ while β is executed, then this action is suspended: an object $suspend(\beta)$ is generated with access trust threshold $high$. Now Alice cannot access α, nor its partly executed state (for example, if she was writing a report regarding insurgent activities in Orange Land this report is suspended), even if the TM functionality allows it: her trust level $vanilla$ overrules this. However she may be able to continue with a new report, using $vanilla$ material.

For Bob, access β is not TM-authorized (he doesn't have discretionary access). He cannot execute β even if his trust level is $high$.

Remark 1. According to the definitions of suspend and restore in Section 3.3, objects created while the trust level of Alice was $vanilla$ are assigned an access trust threshold $vanilla$. This may cause sensitive information to leak to an imposter. To prevent this Alice may assign a $high$ access trust threshold to such objects—however she will not be able to access them until her trust level is restored.

4.4. Access control in variable threat TM systems

There are three levels at which an access action β has to be authorized:

- the *discretionary* level,
- the *mandatory* level, and
- the (user) *trust* level.

The first two define the functionality of the TM system. The last defines the extended functionality proposed in this paper. Access based on the trust level ϕ is temporal and locational,[1] and is determined by the relation between the expected behavior of the user and the actual behavior of the user's ID, and the access trust threshold $\phi(\beta)$. We refer to this authorization as, *trust-level* (tl)-authorization. We have:

Simple trust-level (stl) property:

- *If $\phi \succeq \phi(\beta)$, then the action β is* tl-authorized *and the suspended state of any incomplete β-instantiation is* tl-restored.[2]
- *If $\phi(\beta) \succ \phi$, then the action β is not tl-authorized, and any incomplete β-instantiations that are not already tl-suspended, get suspended and assigned the access trust threshold $\phi(\beta)$.*

The stl-property is a counterpart of the ss-(simple security) property of the Bell-LaPadula model [2]. In our case it is used to protect objects from potential hackers and/or insiders in variable-threat environments. As in [2] it will protect objects (information containers) rather than contents (the information itself). In Bell-LaPadula, a *-property is used to protect information flows. Our model assumes a secure TM infrastructure, and in particular the Bell-LaPadula security requirements: consequently it inherits this level of security.

The easiest way to show this is through an illustration. Suppose that Alice in U.S. has write access to an object α, with $\phi(\alpha) = mission$, that was generated by Bob in Orange Land, who cannot complete it because of a sudden change in his trust level (it is reported that he has been kidnapped and his trust level is reduced to $vanilla$). Suppose that Alice completed the task in U.S. and produced the object γ. Then by the TM-functionality requirements, the objects α, γ have the same security level, and by our requirements in Section 4.1 have the same access trust threshold $\phi(\alpha) = \phi(\gamma) = mission$. This prevents "illegal" information flows.

5. THE WAY AHEAD

There are several areas in which research on variable threat TM systems with rollback access shows promise. Below we highlight three such areas:

[1] Domains are not necessarily geographical.

[2] Full $TM^{auth}_{U,\phi}$ authorization requires TM authorization.

1. *The user trust level structure.* In this paper we have focused on a global, linear structure. Local structures that address issues such as, user trust levels in Orange Land being different from those in the U.S., capture more fully the scenarios described in the Introduction. Observe that if an action is suspended in Orange Land because of a reduction of the trust in the user executing it, there is no requirement to suspend the action for all users in the global enterprise. Thus while the user in Orange Land is restricted from continuing the action because his/her access was rolled back to a more restrictive vanilla access, a user in the U.S. may be able to continue the activity and complete the task. The reason this is possible is the access of the U.S. based user remains higher, or at least less restrictive than the Orange Land user's vanilla access, because of the user's location. This brings into play the aspect of locality as part of the trust management equation similar to what was described for the resource centric view [7].

2. *The impact of trust level dominance on the functionality of TM systems.* We have not discussed how this works, other than require that it is proportional to the access threshold: in particular an increased user trust level will support additional functionality (Section 3.2, end). In general, when modeling access to resources with a high threshold trust level one may want to distinguish the commander in chief from a field worker. So the relation between the trust levels of users and the threshold levels of resources need not be smooth. For example, we may use a model for which the trust value for the commander in chief is always $high$. Alternatively trust values may be linked to clearance levels.

3. *Extending the trust model to allow for a user-centric functionality.* By the nature of the effort being exploratory, we anticipate demonstrating the feasibility of the approach through developing a prototype and initiating a set of indicators of dynamic trust levels. Through the process of development, trust indicators will be formalized and attributed with greater granularity.

6. CONCLUSION

Access control and trust management are the basic components of a trusted information system. In this paper we propose a new access control mechanism that supports a more flexible approach to trust management. This mechanism is triggered by the trust the system has in the relation between the expected behavior of a user and the actual behavior of the user ID (a measure of unanticipated behavior). If this trust falls below a certain level, then access to system resources is reduced to $vanilla$. Later, when the trust gets established (e.g., other users with $high$ trust levels confirm the identity of the user), these are restored thus providing a rollback access functionality.

References

[1] L. Bauer, M.A. Schneider, and E.W. Felten. A General and Flexible Access-Control System for the Web. *Proc. 11th USENIX Security Symp.*, 2002.

[2] D. Elliott Bell and Leonard J. LaPadula. Secure Computer Systems: Mathematical Foundations. TR #2547, MITRE Corp. 1973.

[3] K.J. Biba. Integrity Considerations for Secure Computer Systems. TR #3153, MITRE Corp. 1977.

[4] M. Blaze, J. Feigenbaum, J. Ioannidis, and A. Keromytis. KeyNote Trust Management system, v.2. ftp://ftp.isi.edu/in-notes/rfc2704.txt.

[5] M. Blaze, J. Feigenbaum, and A. Keromytis. The role of trust management in distributed systems security. *Secure Internet Programming*, pp. 185–210, 1999.

[6] M. Blaze, J. Ioannidis, and A. Keromytis. Experience with the KeyNote Trust Management System: Applications and Future Directions. *iTrust, LNCS #2692*, 2003.

[7] M. Burmester, P. Das, M. Edwards, and A. Yasinsac. Multi-Domain Trust Management in Variable Threat Environments Using Rollback-Access. *MILCOM 2008*.

[8] M. Burmester, B. de Medeiros, and A. Yasinsac. Community-Centric Vanilla-Rollback Access. *Security Protocols Workshop, LNCS #4631*, 2005.

[9] D.E. Clarke, J.E. Elien, C.M. Ellison, M. Fredette, A. Morcos, and R.L. Rivest. Certificate Chain Discovery in SPKI/SDSI. *Journal of Computer Security*, 9(4):285–322, 2001.

[10] M.A. Harrison, W.L. Ruzzo, and J. D. Ullman. Protection in Operating Systems. *Communications of ACM*, 19(8):461–471, 1992.

[11] N. Li, J.C. Mitchell, and W.H. Winsborough. Design of a Role-Based Trust-Management Framework. *IEEE Symposium on Security and Privacy*, 2002.

[12] N. Li, W.H. Winsborough, and J.C. Mitchell. Distributed credential chain discovery in trust management. *Journal of Computer Security*, 11(1):35–86, 2003.

[13] R.S. Sandhu, E.J. Coyne, H.L. Feinstein, and C.E. Youman. Role-Based Access Control Models. *IEEE Computer*, 29(2):38–47, 1996.