Journal of Information Assurance and Security
ISSN 1554-1010 Volume 6 (2011) pp. 322-330

A Global Virtual Machine Attribute Access Control Policy for Auditing Federated Digital Identities within a Compute Cloud

Sean Thorpe[1], Indrajit Ray[2], Indrakshi Ray[2], Tyrone Grandison[3], Abbie Barbir[4]

[1]Faculty of Engineering and Computing
University of Technology,
Kingston, Jamaica
sthorpe@utech.edu.jm

[2]Colorado State University
Fort Collins, USA
(indrajit, iray)@cs.colostate.edu

[3]IBM Almaden Research
Silicon Valley, USA,
tyroneg@us.ibm.com

[4]Bank of America
abbie.barbir@bankofamerica.com

Abstract: **An audit framework is necessary to enable policy makers across vertical markets (healthcare, mobile payment, digital content) to set the security and privacy bar for identity providers, identity brokers and relying parties. For sure, across all vertical markets, the sharing of identity requires a baseline of best practices for security, and privacy as it facilitates customer adoption of cloud identity services. In this paper the authors focus on representing the policy concerns required in administering system properties of an interoperable virtual cloud computing network. This need is particularly important in order for evaluation of attribute consistency against the portable virtual machine (VM) identity.**

We study these concerns from the perspective of the security administrator for whose purpose we have named the framework the "Global Virtual Machine Attribute Policy Auditor" (GVMAPA). Audit analysis of the virtual machine (VM) data is unfortunately still absent from the literature on utility clouds. From the perspective of the global system policy administrator, the inability to forecast the behavior of interoperable cloud resources in a collaborative and distributed environment is a growing concern for both academic and industrial communities. As the size of the virtual compute cloud communities grows the need to articulate specific policy attribute constraints for the VM digital identity is a critical requirement not only for law enforcement but the society in general. The motivation for this paper is based on the authors own work in [1] [2].

The paper is divided into five (5) sections -: Introduction and Motivation, Background and Related Work, Scenario Overview, Concepts and Definitions, Scenario Evaluation and Discussion and Conclusion and Future work.

Keywords**:** Cloud, Audit, Attribute, Identities, Interoperable**.**

I. Introduction and Motivation

By capping liability risk through certification, an identity audit framework would make it commercially easier for large Internet consumers, commercial banks and online payment systems to participate as identity providers in high assurance transactions such as health care, eGov services and all new types of compute cloud services. In essence, this is not too different from the VISA model, where a consortium of financial institutions establishes the network blueprint for online payment, as well as the necessary security controls.

If the same is to hold true for the virtualization communities the necessary security policy controls and their enforcements will have to be clearly defined. The data services that run within a virtual compute cloud are not a tangible entity. And hence provisions of measurable audits for such platforms are not only difficult but literally non existent. The focus of the authors work in [1] formally specifies the attribute parameters for the virtual compute cloud stack. The compute cloud stack by definition uses Software as a Service, Platform as a Service, and Infrastructure as a service as three(3) distinct classes. These service classes can be considered meta attributes. Refining the definition shows that the specific attribute parameters must include a physical machine with a set of prescribed location parameters, running from a mapped physical data centre for which the physical co-located machines are abstracted as logical Internet protocol and machine addresses.

It becomes clear that in order to provide VM digital identity access and authentication policy control across distributed and interoperable hybrid clouds -: an attribute based audit policy is required to enable such flexibility.

For the security administrator a VM attribute binary matrix theorizes that there may be an assigned one to one or many to one relationship between a co-located physical data centre and its VMs. As one expands the definition to look at a virtual compute cloud network, the VM binary matrix resizes to include its incremental values for each newly added system property. Different models of access control however have been recommended however fail to capture fine grained property constraints of the dynamic VM environment. Generally the models fall into the general categories of Mandatory access (exclusive rights control),

and Discretionary Access (shared access control). Some of these models include the well known Biba, Bell –Lapudala, Clark-Wilson, Chinese Wall, Task based, Role Based and Attribute based access controls.

These suggested policy models however are context dependent. For a distributed and multi-centric environment like a compute cloud, the issues of which policy is most suitable as a basis for providing sustainable support to the system administration environment has not been well defined.

Attribute based credential policy for assuring identity trust is symmetric to the well established binary models of trust literature. Credentials suggest binary trust conditions against cloud attributes. The user has to produce several predetermined set of credentials (for e.g. student id's, club membership cards and so on) to gain specific access privileges. The credential provides information about the rights, qualifications, responsibilities and characteristics attributable to its bearer by one or more trusted authorities. In addition, it provides trust information about the cloud providers themselves.

Although credential based access control solves the problem of open distributed systems , there are still serious challenges. Strictly speaking, can one tie credentials to a purported behavior or action of a user? Issues like Identity theft threaten the validity of attributable ownership to a credential. Concerns of session history and behavioral history represent common problems, and the issue becomes even more worrisome for virtual service platforms like a compute cloud. The identity trust relationship for cloud entities should accurately be based on experience, knowledge, recommendation, interactions, and behavior as suggested by the authors [3] particularly when one considers the difficulty in managing information flow within such environments.

These observations have motivated the need to establish systems audit as an enforcement for multiple trust level agreements for compute clouds. The idea is to enforce this functional trust requirement as apart of the VM itself at an attribute level. This is indeed a critical observation. But the observation in of itself seems a poorly understood requirement that has only just started to gain traction. Defining such requirement specification standards within an interoperable hybrid compute cloud is still in it's infancy as noted in the very recent work by the Open Virtualization Format (OVF) working groups. Against this background it is difficult to qualify the dynamic (VM) attribute components unless we become vendor specific. Our own arguments of the audit provisions seek to accommodate the case for accuracy and consistency for the VM identity within a continuously dynamic and public compute cloud.

Arguably if one is able to specify and map clearly the constituents of the dynamic class attributes of the VM, then one will be able infuse into the cloud stack these concerns .Equally success at measuring the multi-attribute values of such VMs will significantly build credence and confidence to the argument that trust as an auditable service for utility compute clouds is a realistic opportunity. Hence an audit framework can only engender VM system trust, if only if, one is able to validate VM attribute interaction history, behavioral history, communication history as suitable concerns within this flexible access control cloud context. Invariably this allows the policy administrator to recognize the type of experience, knowledge and reputation required for using such VM resources. When one thinks of a utility cloud, the concerns of each data center loci being able to provide an autonomous authority of its own decisions and judgments is best measured through attribute based control of such abstract services.

Attribute based access control (ABAC) is both dynamic and scalable. A key to GVMAPA scalability is that the issuers of the credentials can be strangers whose authority is based on their own cloud attributes. The concept suggests inherently a meta data management challenge for the existing VM resources. Changes in the access level of the system administrator changes the attributes that this user has in the system and hence the user privileges. The system can define as many trust levels as it wants and can assign each level to a specific set of access privileges. The system just needs to find a way to audit and monitor the trust levels of the administrator and the regulation is automatically achieved. This GVMAPA framework highlights these concerns. The next sections of the paper discuss the Background and Related literature, Scenario overview, Concepts and Definitions used in the GVMAPA, Scenario Analysis and Discussion, then the Conclusion and Future work.

II. Background and Related Work

Access Control Mechanisms (ACM) has been the hallmark for measuring the confidentiality requirement of most IT security frameworks. However this fact is fast becoming a fallacy, as the dynamic nature of information limits the ability to enforce constraints on flow control. This observation notably has been a driving force behind the National Security Agency (NSA) contributing to the design of the SE-Linux object labelled flexible policy architecture. We also adopt as apart of an attribute based cloud policy the arguments for flexible access control within the virtual compute cloud.

The issue is recognizing any single one of these mechanism as a suitable formal policy standard to enforce or assert levels of trust against any one entity. Typical enforcement mechanisms available are Role Based Access Control, Discretionary Access Control (DAC), Mandatory Access Control (MAC), Team Based Access Control (TBAC), Attribute Based Access control (ABAC), and Coalition Based Access Control. NIST to date represents the only standards agency that has sought to unify any representation of these policy models. NIST qualified RBAC models have fast become the defacto standard for access control policies within the enterprise. Variations for RBAC include LRBAC i.e Location RBAC and TRUSTRBAC are even newer variances of the RBAC standard. RBAC represents to date one of the few attempts at constructing a formal definition to qualify and quantify trust parameters for access control policy users. In of itself the RBAC model is used to mitigate the DAC model.

For Multi-centric systems, the TRUSTRBAC purports the use of interaction and behavioural history to model access privileges. The trusted computing base (TCB) is used to enforce access control in peer to peer environments and this is a well known formalism.

As systems move towards utility clouds, TBAC mechanism presents yet another option, however like the existing RBAC models, dynamism and scalability represent constraints for measuring resource identities. The constant scale of cloud resources makes role based access a poor choice for guarantees on evaluating user identities. Similar to the TBAC model, there is also Coalition Based Access Control (CBAC), which uses a cluster or group based access control model to evaluate system privileges. Yet this model too represents a constraint on the access mechanism required for a cloud environment, given that group based private or public cloud clusters are non dynamic structures. Ferrari and Bertino's work in [7] however purport the use of dynamism for access control environments.

Its premise seems promising for work that can be adaptable for cloud distributed and multi-centric attribute based policy trust. In this scheme the access privileges of code is determined at run time. It is the authors view given his own experience as a system administrator, that the VM attribute access control policy auditing will become critical to the stability of interoperable cloud platforms of the future, as any basis for creating any elements of a realized cloud trust model. Today coalitions wishing to share resources often find themselves with no better alternative than to establish a virtual private network (VPN), and to make such shared resources available via the VPN.

This means that users with access to any of the shared resources within a public cloud by example inappropriately have course grained access control. A Finer grained access control mechanism is achievable through an ABAC extended approach which for this paper is examined as apart of the GVMAPA framework model. Alternative solutions that provide the appropriate granularity are based on identity, local role, or capabilities of the cloud resource requestor. As such, they require foreign requestors to be known to the resource providing organization before access can be authorized, such systems do not scale. To date work on ABAC models are still limited based on the well known ABAC project.

The project however has been able to favourably report advances in trust Negotiation. Some of the achieved goal assumptions of ABAC are-: distributed credential discovery, where various specified algorithms and system credentials types are stored with the issuer ensuring credentials can be found to answer to authorization questions. Additionally the ABAC project has also reported techniques in policy language designs. ABAC language specifications demonstrate that existing trust negotiations meet these requirements. Prominent trust management languages like KeyNote, and Delegation Logic, SPKI 1.0 act as suitable candidate language requirements. ABAC also supports the design of a realistic trust negotiation strategy. The point of view shared by this latter point is predicated on the argument that successful access control to credential content and information is based on the credentials that a negotiator holds.

Critical credentials have high bandwidth covert channels that enable unauthorized access to credential content. ABAC has successfully closed the gap on identification of covert channel concerns in this space. ABAC is foreseen to have a wide bearing impact on military and commercial coalition operations. For example the DARPA/Army Future Combat System requires rapid deployment and joint international interoperability. This requires an authorization infrastructure that can be administered efficiently. The issue becomes even more concerned when the deployment and interoperability constraints are migrated to the cloud.

Other likely military contexts for the deployment of this technology include the US Army Communications – Electronics Command (CECOM) where the joint vision 2010 aims to maximize information Systems integration and interoperability while increasing system/platform effectiveness. Integration of forces will require subjects, including intelligent software agents, from various multiple organizations to establish trust with one another rapidly, automatically, and effectively. ABAC technology is context adaptable, beyond traditional access control that require trust establishment.

ABAC aims to address a fundamental problem confronting dynamic coalitions throughout the military and commercial sectors: how to make authorization decisions without requiring prior local knowledge of each subject in the coalition. More generally, this problem confronts any pair of subjects attempting to establish trust with no prior contact or knowledge of one another. The numbers of situations with respect to handling identities within the cloud computing space are enormous, and currently there are no widely accepted standard solutions.

III. Fundamentals

A. Overview of Case Scenario

In the author's earlier work on modeling Contextual trust concerns for Digital Identities [2], the survey used the University of Technology (UTECH) as the prescribed environment to represent the features and behavior of digital identities. In extending this work for the GVMAPA Model, one canvases the prescribed case scenario in [2] by emphasizing the arguments as a suitable set of concerns that can be reused for utility cloud deployments. The suggestions and arguments herein will both highlight and build on the identity characteristics in [2] for a general virtual machine academic cloud.

The primary idea considers a credential based cloud system where transactional logs are used to both quantify and qualify the client/server attribute behavior for such digital identities. However, parallel work on the synchronization of virtual machine logs within the physical device logs as a measure of attribute consistency is being done, and hence not the focus of the current paper. Interestingly the GVMAPA model is also currently used to determine new research opportunities for distributed identity vulnerability monitoring systems in general. In this paper the measure of what constitutes a digital identity is assumed to be well understood. Hence one approaches the definition, by suggesting that the UTECH Meta verse of a digital identity as a subject or entity, which is a student, an academic staff, a school, or a software application making a request for a UTECH defined resource or system property etc. Evidently such a resource might be a Lecturer's web page, a piece of data within a database, an online credit card transaction and so on. To gain access to the UTECH resource, the

subject lays claim to an identity for which we scale this defined resource as a virtual machine resource as apart of the service architecture for cloud environment.

In this context, VM identities are represented as data objects that contain attributes, preferences and traits. It is these data objects that must show uniqueness of character when used as apart of the virtual compute cloud. Hence for the University these data objects and it's constituent attributes includes student history, past payment behavior, account balances, date of birth and so on.

Preferences represent desires such as a choice of a course elective within your degree major, the semester in which to select a course major, based on the semester course offerings. Preferences could also consider the use of a specific cryptographic signature over another to secure the data. Traits will include those inherent characteristics of the subject. Examples include things like a person's height, weight and so on.

For a subject to use a particular identity in order to access a virtual machine resource, one assigns the subject a credential. As a possible relation to the Josang model of trust, a credential will suffice as the basic requirement for a trust recommendation. In other words one can equally describe attribute credentials as sources of proof to substantiate that a subject has a particular right to a resource within the UTECH metaverse. Essentially credentials are a transferrable form of trust among subjects. When the credentials are presented to the UTECH based security authority or the equivalent policy enforcement agent, the agent authenticates the credentials. Authentication in this context assumes a user name, password and an X.509 certificate. The level of authentication required is proportional to the perceived risk involved in accessing a UTECH resource .

Within the context of the GVMAPA, the attribute credentials are believed to be authentic, the authority retrieves the security policy or sends it to an independent policy decision agent. The policy decision agent uses the security policy and the asserted identity to determine the entitlements and permissions associated with that resource for that particular identity. The GVMAPA also supports that for the prescribed context, an entitlement is an action that becomes triggered, whenever the attribute identity is referenced. Entitlements can be seen as the virtual machine services and resources to which an attribute VM identity is entitled. The detail analysis of the contextual framework for modeling a digital trusted identity in traditional systems is the basis of the earlier work [2][14].

What is evident is that a GVMAPA definition allows for novel opportunities to scope clear open access control policy boundaries for the use of the VM identities. Its articulation can help to demystify the concerns of the very elusive argument of identity representations within say a public versus a private cloud domain within an identity federated network. In other words the need to be clear on what models constitute private versus public domain attributes for a compute cloud has to be calculated as a function of the data center owner's responsibility to be able to demarcate where a VM identity is co-located being able to deterministically state over qualifying the domain context as private versus public this contextual model of identity trust, one should treat security as a parametric set of both identity and authorization being the subject issues to be included. Privacy enforces granular access control of the attributes, preferences and traits associated with an identity from being disseminated beyond the subject's need in any particular transaction. In a cyclic manner, this identity is built on the foundation of good information security and that is dependent on good identity management.

B. Scenario Analysis

At first glance the term digital identity seems foreign to most persons; however it becomes perfectly understood given the chosen context for which it can be used[2]. The presupposition is that trust enforcement is only possible if the context allows one to define the attributes, preferences and the traits required clearly for an identity. In other words a failure to define the specific context leaves one to various assumptions leading to eventual distrust. The recognition of an identity within the physical versus digital world has several parallels of context, however the physical tends to be more appreciated as persons can relate to the types of personal or commercial transactions of such environments. Whereas within the digital world relations are harder to establish for identities as an obvious fact, unless we can parameterize the context.

So one can take the physical world of a University campus for which the GVMAPA is now enforced, and consider the case of a student identification card as a credential that asserts that a person has certain attributes and traits. The identification card contains authorization to perform certain task, for example to attend an examination sitting or a tutorial seminar. So when a person (i.e. a student) wants to write an examination (i.e. perform an action on a resource), the invigilator (i.e. the security authority) examines the student's identification card to determine if it looks real (i.e. determine the validity of the credential) and uses the picture along with the unique identification number and swipe it against the virtual machine reader at the exam entrance (i.e. embedded biometric device and X.509 certificate). Once certain that the card is authentic , the invigilator reads the faculty bar code (i.e. the attribute) from the student card and determines the faculty to which the student belongs so that the exam for which student intends to write matches the subject of the corresponding Faculty from which the student comes (i.e. consults a security policy determined by the University academic administration) and makes policy decisions about permissions associated with the identity.

Now suppose the same student with his ID card was actually sitting a laboratory exam of the same context and was required to purchase print paper for his exam results ?.He supplies a new credential, a student print card which has monetary value assigned to purchase print credits from the lab. The print card is supplied to the invigilator, acting as a policy enforcement point, runs the print card through the local Point of Sale device in the lab which transmits identity attributes from the card (i.e. card

number, expiration date, and card cash surrender value) along with the named printer resource to be accessed to the lab manager , who acts here as the policy decision point and determine whether or not the subject is entitled to print based on the minimum credit value of the card. The invigilator receives the print authorization and completes the transaction.

The Context of Digital Identity within a University Environment

In addition to being able to identify customers in order to provide them with a service, the University as a business has an increasing need to identify employees, systems, resources and services in a systematic way to create business agility and ensure the security of such business assets . This notion of business agility is predicated on the merits of the ***Service Oriented Architecture (SOA)*** approach which is a well established principle. The SOA framework is implicitly assumed for functionality of this GVMAPA model framework.

Past experiences have shown that the traditional University network was considered provably secure if a local firewall and defined perimeters of access control were established, and where this persist as an existing problem, scalable cloud networks represents even more demanding concerns for the virtual perimeter. Hence as the University expands it academic services and widens it's entrepreneurial activities beyond the physical campus limits , the need to integrate internal systems with both local and external business partners and customers has become paramount.

The creation of standards for exchanging data is fast becoming a significant trend to distributed computing embodied within web services. This trend has great implications for how the University functions as a business community: one cannot sit idly by and treat the end points of the network as a secure perimeter. Invariably the suggestion is implicit to qualify the perimeter as an un-trusted perimeter. When the University posits itself as a business entity we cannot any longer only talk about the network and the machines but one needs to talk about the data, documents, people, the actions and partners as they are also extremely important to this framework. This type of security model is definitely more complex in nature as the perimeter to secure is a ***virtual perimeter*** unlike the traditional systems we have been use to. But even if we define a policy , how can we ensure that it is properly implemented across dozens or even hundreds of virtual cloud network systems at the same time control access to resources as granular as fields with a database or even paragraphs within a document? This answer lies within the digital identity management framework that the University would be expected to provide for both its traditional environment and compute clouds. Within the Campus network a digital identity management is built on a set of concepts that provide a framework for addressing the identity in a business context instead of an IT Security context :

- Integrity and non-repudiation
- Confidentiality
- Authentication and authorization
- Identity attribute provisioning

The most significant challenge in each of these areas is interoperability since the goal is to build an enterprise wide trusted University identity cloud management system, even one that federates outside of the University enterprise. In the subsequent section of this paper, the author provides a short discussion on the use of UML 2.0 to simply represent a basic University Person Identity as seen applicable for use within a basic compute cloud. The author's sole intent is to disambiguate any misconceptions as it relates to determining identity attributes. This discussion is then preceded by a short evaluative treatment on how this domain can be used to derive quantifiable constraints that represents a possible cloud identity attribute context graph. A subsequent full paper seeks to use UML modelling to define a multiplicity of other VM instance identities and their attributes as a complete virtual system of network properties. These concerns will provide a suitable base for use within the GVMAPA framework

C. Features

Let's start with taxonomy of all the prescribed identities that we'll need to account for as a basis of our starting definition.

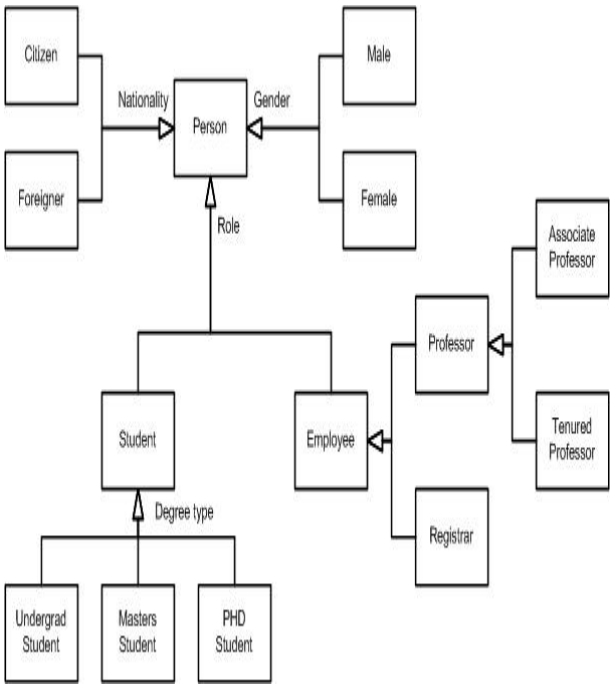

Figure 1.0 Taxonomy for digital identities within the university. (Source [12])

One can further decompose the taxonomy of Figure 1.0 into a basic class diagram:

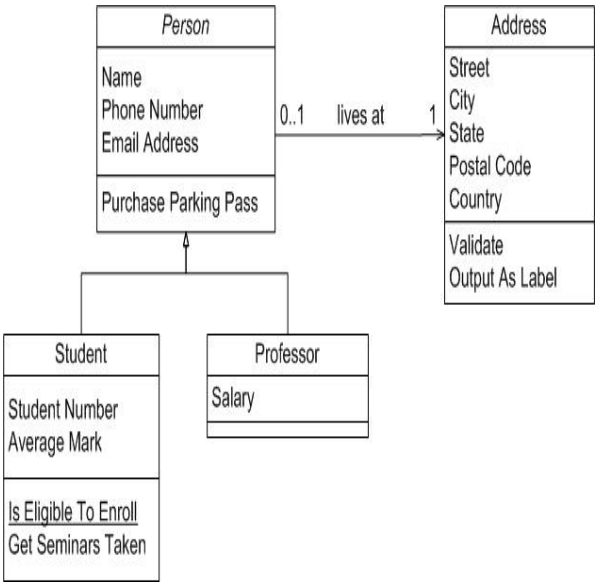

Modelling a Person Identity as apart of an inheritance hierarchy (Source [12])

The idea here is that one can see the constituents of a student identity, and by extension one can determine the components of the identity relationships that are established against this student entity model.

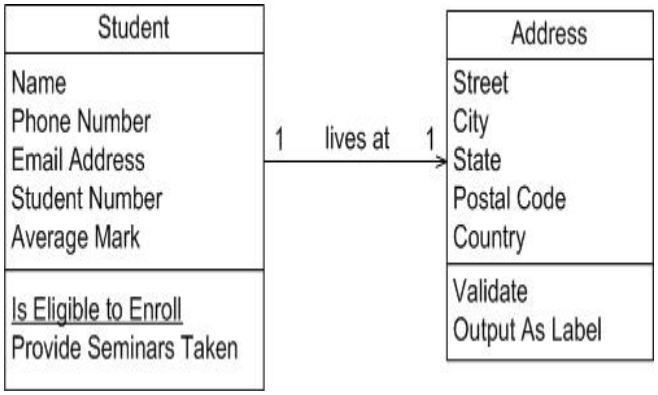

Figure 1.3 Modeling a Student Identity (Source: [12])

In the context of assessing the UTECH scenario the following arguments are adopted. From Fig 1.3, the class level responsibility is a basis for defining the **context action** for this model. Hence a student's ability to enrol becomes a persistent context action for every course in the programme, and hence an assertion of the trust that may be given to that student for every context instance that may arise.

IV. Concepts and Definitions

For the purposes of this work, let's assume that the VM identity attributes are defined as an ordered collection of attributes. These attributes are classified into two (2) types (1) Non Category attributes (2) Category attributes. Non Category attributes are decision making attributes such as a role, subject, location, and time etc. Each non category attribute represents some important features of a particular rule and contains some discrete or continuous value. On the other hand category attribute represent the class of the rule to which it belongs. Typically a category attribute takes the binary values {Accept, Reject}. One has direct inconsistency when two rules present the same cloud attribute policy set that end up with contradictory results.

For example ,suppose that one rule states that a VM user x is allowed access to resource r and another rule states that the same user x is allowed access to the same resource in the same context. Formally one can define attribute inconsistency in the following manner.

Let R be a set of rules (R = {R1, R2,….R_n }, where $R \neq 0$. All rules $R \in \Re$ have a uniform structure, consisting of a number of attribute/value pairs. This can be realistically expected, if one assumes that default values can be used. Each rule comprises of a set of non category attributes A = {A_1, A2,A_n} and one category attribute C. Formally a rule R_i can be written as follow:

$R_i : A_{1} \cdot \wedge A_2 \wedge \ldots\ldots\ldots \wedge A_n \rightarrow C$

For example, consider the following rule.

R: role (Professor) .Λresource (Student Record)Λ action (Write) → Allowed

In this example, the VM policy rule seeks to recognize, Professor, StudentRecord and Write operation as the non category attributes of the rule and allowed is the Category attribute of the rule. Let $v(R_i \cdot A_j)$ denote the value assigned to a VM attribute A_j in the rule R_i

Definition 1 : Rules $R_i . R_j \in \Re$ are mutually inconsistent if

1). $\forall A_k \in A \;\; v(R_i . A_k) = v(R_j . A_k)$
2) $v(R_i . C) \neq v(R_j . C)$

Informally condition 1 in the above definition states that all decision making VM attribute values of R_i are the same as the corresponding attribute values of rule R_j and condition 2 states the category attribute value of rule R_j. Note that one assumes that decision making attributes will be in the same order for all rules, a condition that can easily be satisfied.

One has indirect inconsistency if the two rules are present in different policy sets lead to contradictory conclusions.

Such inconsistencies are difficult to see because they may not be visible at the time of defining policies and can only be triggered only when some specific event occur. For example on the VM server a Professor T is allowed to create a student exam account and Mary a student is allowed to delete accounts. A policy may state that create and delete operation cannot be performed by the same entity or identity. Inconsistency could occur if Professor T delegates his rights to Mary. From the perspective of having a VM policy attribute service within the GVMAPA model that can perform data mining on the attribute identities , one could formally define inconsistency in the following manner.

Definition 2: Rule $R_i . \in \Re$ and $R_j \in \Re^1$ are mutually inconsistent if

1).). $\forall A_k \in A \;\; v(R_i . A_k) = v(R_j . A_k)$

2) $v(R_{i.}\ C) \neq v(R_{j.}C)$

Definition 3
As a basis of determining VM attribute inconsistency , one adopts the use of decision tree classifiers , as expressed in the following algorithm below :

Algorithm 1.0 VM Node Inconsistency Detection Algorithm

Input: Decision tree
Output: Context of VM node inconsistency

1: Let *A*(*bi*) be the set of all attributes present in one branch.
2: Bool *consistent* = *true*;
3: **for** each branch *bi* in Decision tree **do**
4: **if** more than 1 category attribute is assigned to terminal node *bi.tnode* **then**
5: *A*(*bi*) = fetch all attributes of branch(*bi*);
6: **for** each actual rule *Ra* in the policy set **do**
7: **if** *v*(*A*(*Ra*)) = *v*(*A*(*bi*)) **then**
8: Highlight: *Ra* : *A*1 .. . *An* → *C*;
9: **end if**
10: **end for**
11: *consistent* = *false*;
12: **end if**
13: **end for**
14: **if** *consistent* = *true* **then**
15: No inconsistency found;
16: **end if**

In a decision tree, each branch *bi* (from the root to a VM terminal node) represents one rule. In order to detect inconsistency, one will apply the above algorithm. First one checks the terminal node of each branch (Lines: 3- 4). If any terminal node *tnode* contains more than one category (*C*) attribute value (Line: 4), this means that some rules in the policy set are mutually inconsistent. In order to determine which particular rules in the VM policy are mutually inconsistent, first one fetch all the attributes of the particular branch (Line: 5). After that the algorithm will start searching the attribute-values in the actual VM policy set (Lines: 6-10). All the rules in the policy set that contain those attribute-values will be highlighted as inconsistent (Lines: 7- 9). If in a decision tree, no terminal node has more than one category attribute-value then this means that no inconsistency has been found in the policy set (Lines: 14-16).

Definition 4
An identity will always display multi-attribute behavior. The VM policy attribute machine is required to develop, demonstrate and correct inconsistency analysis against these attributes.

Definition 5
An identity's context cloud attributes are quantified by the specified data objects, machine logical addresses, and time allocated to a specified cloud instance based on the formalisms below:

$$e_x^i = \{t_x^i, \{m_1^{i^x}, ..., m_a^{i^x}\}, \{d_1^{i^x}, ..., d_b^{i^x}\}, \{cn_1^{i^x}, .., cn_e^{i^x}\}\}$$

where

$m_a^{i^x} \in M$, $d_b^{i^x} \in D$, $cn_e^{i^x} \in CN$, $1 \le a \le \infty$, $1 \le b \le \infty$, $1 \le e \le \infty$

Let M be the set of all machines, D is the set of all related data and CN is the set of all connected nodes. *M* is assumed to be a collection of physical machines only. *a* is the maximum number of machines, *b* is the maximum number of data transactions and *e* is the maximum number of nodes connected to the UTECH data center. As the number of machines grows across the UTECH data center, the values for *a* and *b* are assumed to have no upper limit as the number of machines and data for a particular context will grow or shrink as needed.

Definition 6
The issuance and renewal of all identity context attribute credentials are handled by an identity Kerberos cloud provider which forms apart of GVMAPA.

Algorithm 2.0 Make Identity Context Graph(*Context* c_a)
1: Create an empty identity graph *G*
2: For each context environ e_x in c_a
3: Create a identity graph node (ign) in *G* for e_x
4: For each machine *m* contained in e_x
5: Build a node *m* in ign e_x for machine *m*
6: For each e_x in c_a
7: For each next element *e* contained in e_x
8: Find the associated node or ign associated to *e*'s source
9: Find the associated node or ign associated to *e*'s target
10: Create a link between source and target
11: Assign a time value to the weight of the edge
12: **return** *G*

V. Scenario Evaluation

In this section an establishment of the Virtual Digital Identity System properties used to describe the concepts and definitions of the GVMAPA highlighted in the previous section is done. The System Digital Identity properties are characteristics and resources of a computer system as in this case the basic model of a University compute cloud network and its environment component resources. This qualification harmonizes the discussion for the UTECH digital identity environment. Each virtualized system property describes one specific attribute at a time of such a system. Equally a representation of all related digital identity properties are depicted in Figure 1.4 below.

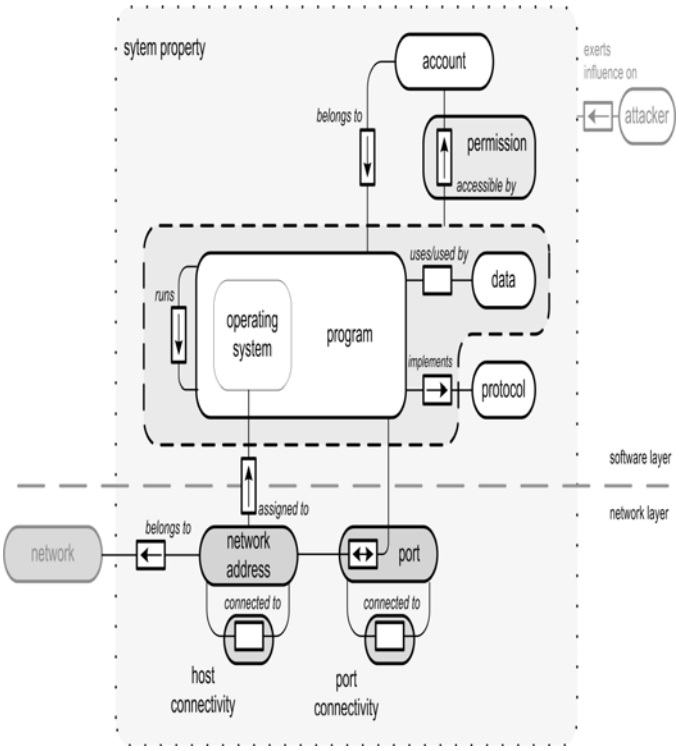

Figure 1.4 (Adopted From [15])

For example, the installed version of an application can be a system property. An application's version is meaningless if it cannot be identified and linked. In this UTECH environment this is critical.

Adopted from Roschke et.al[15],your digital system properties here seeks to describe states a virtual system can be in, e.g., running programs, existing accounts, and existing databases. In the context of the current work an identification of the digital identities within a federated UTECH cloud, helps one to realize the purpose for which these digital properties are significant to the system administrator. Let one contend that these properties are called "Influence properties". Influence properties in the GVMAPA seek to describe the influence an attacker has on system properties by successful exploitation.

If one can establish the influence on such digital properties, it becomes relevant to the model to determine the scope of the impact for any such properties. Here one can define such scope as a "Range properties" by definition. Range properties describe the virtual machine location from which an attacker can perform successful exploitation, e.g., local or remote.

This range on these systems attributes then allows one to determine vulnerability points. Vulnerability requires a precondition and a post-condition, which can be represented by these virtual digital system properties. The vulnerability attack graph scenarios that considers the covert channel clouds that may lead to inconsistent and cloned identities and their impact is currently addressed as independent work within the research group. As a basis for leveraging what is the likely scope of impact on a digital identity cloud resource, will be based on the types of vulnerability constraints to which these VMs are likely to be exposed.

One considers two basic types that could be considered for the descriptions: properties and sets. Properties represent predicates and sets represent a grouping of properties based on binary logic. Both types facilitate a simple evaluation based on matching binary values of True or False. Finally, descriptions link different virtual system states together, one as the requirement and the other as the result of an attack. Based on this properties and sets, one can flexibly describe many different digital identity system attribute states with some of level of accuracy.

Digital system properties are characteristics and resources of a virtual system for which are considered relevant vulnerability information. Each system property describes one specific attribute of such system at a time, whereas properties are related to one another by categories.

For example, the installed version of an application can be a System property. An application's version is meaningless if it cannot be linked to a certain application. Properties and their relations may change over time due to modifications, such that an application may be upgraded to a newer version. System properties can be found in two layers, the network layer and the software layer. The network layer describes properties of interconnected computers, such as virtual mapped network addresses, system properties which are useful to create attack graphs, such as network properties and port numbers. The software layer describes properties of software systems, such as programs, data, and account information. One could define several different virtual digital system properties which are useful to create attack graphs, such as network properties, host connectivity, programs, protocols, data, accounts, and others.

To describe actions performed on systems, influence attribute properties can be used. Influence properties describe the relationship between a potential attacker and the virtual system properties which represent computer resources.

VI. Conclusion and Future Work

This paper has presented the work flow properties that can be used to identify utility compute cloud resources as fine grained attribute constraints for such abstract environments. The arguments are significant in their contribution towards the representation of a Global Virtual Machine Attribute Policy Auditor framework within the system administrator environment. This qualification of these system attributes and their behavior provides the security administrator with a context for monitoring VM identities within hybrid compute cloud environments.

Ongoing work on GVMAPA examines synchronization of the attribute properties for both the physical and virtual disk logs of the VM System environment. Subsequent opportunities of this work is to find and connect more applicable virtual data sources to the system, e.g., historic user and system data can be used for forensics and correlation of IDS alerts over a long period of time. The system will need extensive performance tests and scalability tests especially for enterprise wide distributed VM environments.

References

[1] T. Grandison, E M Maximilien, S.Thorpe, Alba, Towards a Formal Model of Cloud Computing, - July 2010 6th IEEE World Congress.

[2] Sean Thorpe, "Contextual Models of Trust for Digital Identities" IEEE IAS August 2010

[3] Indrakshi Ray, Indrajit Ray and Sudip Chakraborty, An Interoperable Context Sensitive model of trust, Journal of Intelligent Information Systems, 32(1), February 2009. Retrieved from: http://www.cs.colostate.edu/~iray/projects/afosr06 index.html

[4] T. Grandison and M. Sloman. A Survey of Trust in Internet Applications. *IEEE Communications Surveys and Tutorials*, 3(4):2.16, Fourth Quarter 2000.

[5] A.J.I. Jones and B.S. Firozabadi. On the Characterization of a Trusting Agent. Aspects of a Formal Approach. In C.Castelfranchi and Y.Tan, editors, *Trust and Deception in Virtual Societies*, pages 163.174. Kluwer Academic Publishers, 2000.

[6] A. Jøsang. Artificial Reasoning with Subjective Logic. In *Proceedings of the Second*, 2008.

[7] E. Bertino, E. Ferrari and V. Atluri, The specification and enforcement of authorization constraints in workflow management systems, ACM Trans. Inf. Syst. Security. 1999.

[8] A. Jøsang. Logic for Uncertain Probabilities. *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems*, 9(3):279.311, June 2001.

[9] Garth V. Crosby, Lance Hester and Niki Pissinou. "Location-aware, Trust-based Detection and Isolation of Compromised Nodes in Wireless Sensor Networks."December 2006. Retrieved from http://ijns.femto.com.tw/paper_upload/IJNS-2006-12-08-1.pdf

[10] R. Yahalom, B. Klein and T. Beth. Trust Relationship in Secure Systems: A Distributed Authentication Perspective. In *Proceedings of the IEEE Symposium on Security and Privacy*, pages 150.164, Oakland, California, USA, May 1993. IEEE Computer Society.

[11] R. Yahalom and B. Klein. Trust-based Navigation in Distributed Systems. *Computing Systems*, 7(1):45.73, Winter 1994.

[12] "*Introduction to UML 2 Class Diagrams*", located on: http://www.agilemodeling.com/artifacts/classDiagram. html

[13]"*Understanding Digital Identity Management*", *Retrieved from:* http://www.windley.com/docs/DigitalIdentity.pdf

[14] Bo Yang, Min Zhou, Guohan Li,Jiajia Huo,A Secure Reputation System, Journal of Information Assurance and Security ,2007.

[15] Roschke, S., Cheng, F., Schuppenies, R., and Meinel, Ch.: "Towards Unifying Vulnerability Information for attack graph Constructions." In: Proceedings of the 12th Information Security Conference (ISC'09), Springer LNCS,vol. 5735, pp. 218-233, 2009).

Author Biographies

Sean Thorpe holds a MSc in Information Security and a BSc in Computer Science respectively from the University of Westminster London, UK in 2002, and from the University of the West Indies, Mona Campus, Kingston, Jamaica in 1999. Mr. Thorpe joined the University of Technology (UTECH),Jamaica as a Lecturer in 2003 with responsibility for teaching System Security at the undergraduate level. Mr. Thorpe has worked extensively in the IT industry since 1995 as a System Programmer and Oracle DBA before joining academia. He is the 2009 recipient of the Fulbright Visiting faculty Scholarship award to Harvard University, where he explored collaborative research work in the area of Security Metrics. He is also the 2009 winner of the OOPSLA Educational Symposium Award for his innovative computer science teaching methods. His specific research interest includes cloud forensics, and security policies. He is currently a PhD candidate in Computer Science.

Indrajit Ray is an Associate Professor at Colorado State University since 2002. Prior he was an Assistant Professor at the University of Michigan, Melbourne from 1997 to 2001. He earned his PhD from George Mason University, Virginia in summer 1997. He obtained his BSc in Computer Science Degree from Bengal Institute in India in 1984 and then his MSc from Jadvapur University in 1991, also in India. His primary research interests are digital forensics, security policies, access controls, and intrusion detection.

Indrakshi Ray is an Associate Professor at Colorado Sate University since 2002. Her research interest includes access controls, UMLs and workflow security issues for the enterprise.

Tyrone Grandison is Program Manager of the Health Informatics group at IBM Almaden Research Labs. Silicon Valley, California. He holds a PhD from Imperial College, University of London, UK since 2003, and prior obtained his MSc and BSc degrees from the University of the West Indies, Mona Campus, Jamaica. He holds several patents for research in the area of medical database security.

Abbie Barbir holds a PhD from Arizona State University since 1991. He currently heads the OASIS Security Study group 17 responsible for policy formations and standards for open access environments. He has significant industry experience within the telecomm sector including Nortel and the ITU. He is currently a consultant to Bank of America.