QUT isi Information Security Institute

Access Control

Theory and Recent Research

Ed Dawson
Jason Reid
Farzad Salim

a university for the real world

Agenda

- Introduction – access control in context
- Overview of authorisation and access control models
- Information sharing – motivation and challenges
- ISI access control research
  - An approach to access control in dynamic environments
  - Access control for privacy compliance
  - Logical and physical access control convergence using a Building Information Model (BIM)

Information Security

- Access Control is a key element of Information Security but what is information security?

  Definition: "the protection of information from a wide range of threats in order to ensure business continuity, minimize business risks, and maximize return on investments and business opportunities"

  Source: AS/NZS ISO/IEC 27002 Information technology - Security techniques - Code of practice for information security management

Information Security

- The practice of information security has traditionally focused on developing and maintaining systems and processes that protect information and information systems by preserving
  - Confidentiality, Integrity, Availability
- A more recent interpretation of the term places greater emphasis on
  - privacy protection
  - governance and compliance
  - Business continuity management
- Guided by strategic risk management

Fundamental Security Properties of Information

- Confidentiality: The property that information is not made available or disclosed to unauthorized individuals, entities or processes
- Integrity: The property that data has not been altered or destroyed in an unauthorized manner
- Availability: The property of being accessible and usable upon demand by an authorized entity

Source: ISO 7498-2:1989 Information processing systems -- Open Systems Interconnection -- Basic Reference Model -- Part 2: Security Architecture

Information Security and Authorisation

- Formally, the objective of information security is the preservation of the security properties: confidentiality, integrity and availability
- Notice that each property is defined by reference to actions that are *authorised*
- What does authorise mean? Some dictionary definitions:
  - authorize: To give legal or formal warrant to (a person) to do something; to empower, permit authoritatively (Oxford English Dictionary, 2nd Edition)
  - authorization: The conferment of legality; formal warrant, or sanction. (Oxford English Dictionary, 2nd Edition)
- An information security definition:
  - authorization: The granting of rights, which includes the granting of access based on access rights (ISO 7498-2)

ISO 7498-2:1989 Information processing systems -- Open Systems Interconnection -- Basic Reference Model -- Part 2: Security Architecture

Authorisation and Access Control

- Authorization is about 'who can do what' - a policy that defines allowed actions with respect to information – the term clearly encompasses policy definition
- Does authorisation also include decision making and enforcement of the policy when responding to requests (access control)?
  - There is some disagreement on this
  - A meaning that includes both policy definition and decision making is widespread in the literature e.g.
    - In standards: "*Authorization: The granting of rights, which includes the granting of access based on access rights*" (ISO 7498-2)
    - In textbooks: "*Authorization refers to a yes or no decision as to whether a user is granted access to a system resource*." Ferraiolo, Kuhn & Chandramouli, Role-Based Access Control, Artech House, 2003
  - Some authors argue that the term should only apply to policy definition – see Gollman, Computer Security, 3rd Edition, Wiley, 2011, p 387.
  - This presentation uses the term in the wider sense of policy definition and decision making

Authorisation and Authentication

- Entity authentication is a prerequisite for authorisation
- Why? Because the concept of identity is often central to how the access rules are formulated and enforced:
  - e.g., John Smith: employee number 675324 can read and write file *report.doc*
  - To enforce the policy, the system needs to know if it is interacting with the real Bob or an imposter
- Authentication establishes a degree of confidence in a claim.
- Entity authentication establishes confidence that a person can rightfully claim an association with a unique identifier (such as a username or a staff number, or a name and date of birth) that distinguishes them from other persons within a domain (e.g., the domain of current and past employees of Acme Corp).
- An authorisation policy does not have to be based on identity (and thus rely on authenticating identity). It can be based on attributes of the user e.g. a website that is only accessible to people 18 years and older

Entity Authentication

- A person can authenticate their claimed identity in three ways:
  1. by something they know, where that knowledge is a secret that is not known by others e.g., a password or PIN;
  2. by something they have, where the possessed artefact is uniquely identifiable and difficult to duplicate or counterfeit e.g., a smart card that stores a secret cryptographic key;
  3. by something they are, where the individual physical or behavioural characteristic is reliably measureable, difficult to replicate or impersonate and unique among persons in the domain. Physical or behavioural characteristics of this type are known as biometrics.
- A relying party can increase their level of confidence that a person can rightfully claim an association with a unique identifier by requiring the individual to prove their claim using multiple factors.

Access Control

- **Access Control**: The process of granting or denying specific requests by authenticated users to obtain and use information resources and related information processing services (objects)
- **Subject:** a process executing on behalf of an authenticated user and associated with their identifier
- **Entity Authentication**: The process of establishing confidence in the authenticity of a entity's claimed identity. Authentication is a prerequisite for access control.
- **Access Control Policy**: The rules that determine whether a request should be allowed or denied
- **Access Control Mechanism:** Low level hardware and software functions that enforce the policy – together they implement a **reference monitor** to mediate access to objects
- **Access Privilege:** the ability to perform an action on a specific object (an object+action pair)

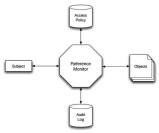

Access Control Policy Types

- Two main types of access control policy – discretionary and mandatory
- Discretionary Access Control (DAC) characteristics:
  - **User control**: Users have the discretion to grant access privileges to other users
  - **Identity based**: Access decisions are made on the basis of a user's identity
  - **Policy is owner-centric**: Resource owners can grant access privileges to other users without restriction
  - **Uncontrolled propagation:** Users can propagate access privileges directly or indirectly without restriction
- Mandatory Access Control (MAC) characteristics:
  - **Central control**: Access privilege determined by an administrative authority – not by users – users do not 'own' information
  - **Controlled propagation:** users cannot propagate or grant their access privileges to other users

See: Messaoud Benantar, *Access Control Systems: Security Identity and Trust Models, Springer, New York, 2006*

Discretionary Access Control (DAC)

- Access Matrix is the basis of DAC
- Matrix is sparse and inefficient to store – mostly empty cells
- Usually stored by column with the object - known as an **Access Control List** (ACL)
  - ACL's make it inefficient to determine all access privileges of a specific user – each object's ACL must be examined
- Called a **Capability** when stored as a row
  - Capability-based: hard to determine all users that have access to a particular object
  - Privilege revocation is difficult
  - Capability forgery must be prevented
  - Useful in distributed systems

Subjects	Objects		
	File 1	File 2	Program 1
Alice	Read		execute
Bob	Read Write		Read Execute
Carol		Read Write	

Weaknesses of DAC

- Corporate and government users do not 'own' information – the corporation does. Regulations and the corporate security policy do not permit users the discretion to determine who should have access
- Vulnerable to Trojan Horse attacks
  - A Trojan horse is an apparently useful program which also performs hidden malicious actions – e.g. a free screen saver that installs a keystroke logger
  - With DAC any process initiated by a user will execute with all the user's privileges – thus DAC does not observe the principle of 'least privilege'.
  - If the user is tricked into executing a malicious program it can do anything the user can do – read, copy, delete files etc.
  - In a DAC system a Trojan horse can leak sensitive information – unlike MAC systems which control information flow
  - Users may be trustworthy but prevalence of malware means the programs they execute may not be
  - MAC MLS developed in part to address this problem of untrusted software

MAC

- Multi Level Security (MLS) as formalized by Bell and La Padula is the most common MAC policy
- Used to enforce military security policy
- Each object and subject is assigned:
  - A **Security Level**: Hierarchically ordered set: Top Secret (TS), Secret (S), Classified (C), Unclassfied (UC) where TS>S>C>UC
  - Set of **Categories:** unordered labels which represent functional or competence areas e.g. Army, Nuclear, NATO
  - Access class c1 **dominates** (≥) access class c2 if the security level of c1 is greater than or equal to that of c2 and the categories of c1 include those of c2
  - Access classes + dominance relationship form a mathematical structure called a **lattice**

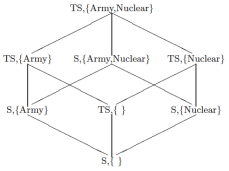

P Samarati and S de Vimercati. Access control: policies, models, and mechanisms. In Foundations of Security Analysis and Design, volume 2171 LNCS, pp 137-196. Springer, 2001.

MLS for Confidentiality

- Security level assigned to users reflects their trustworthiness – known as their **clearance**
- Users connect to a MLS system at a level that is dominated by their clearance e.g. Secret cleared user can connect at Unclassified, Classified or Secret
- Reference monitor enforces two rules:
  - **No read up:** subject cannot read information whose access class dominates their clearance – e.g. subject connected at Secret cannot read a Top Secret object
  - **No write down:** subject can only write information if the access class of the object dominates their access class e.g. a subject connected at Secret cannot write to a (lower) classified object but can write to Secret or Top Secret
- Together these rules prevent information flow from higher access classes to lower access classes

MLS for Confidentiality

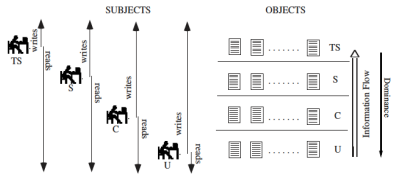

- Note: No write down rule means a user cleared to Top Secret must connect at Unclassified if they wish to write to an Unclassified object – e.g. a memo to a subordinate

Image source: P Samarati and S de Vimercati. Access control: policies, models, and mechanisms. In Foundations of Security Analysis and Design, volume 2171 LNCS, pp 137-196. Springer, 2001.

Weaknesses of MAC

- Rigid classification makes information sharing difficult
- Downgrading information to lower levels is (intentionally) difficult but often operationally necessary – in MLS users are not supposed to be able to alter an object's classification – but part of the object may be less sensitive
- A classification system based on hierarchy of information sensitivity levels is hard to apply in non-military settings
- The system is safe but inflexible – inappropriate for most commercial, government, healthcare applications
- Information leakage can still occur through covert channels

RBAC

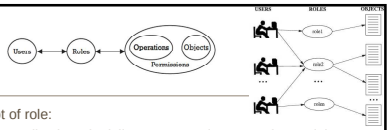

- Based on the concept of role:
  - a role is a coherent collection of privileges appropriate to perform a job function e.g. doctor, nurse, teller
  - Users are assigned to roles based on their competency and responsibility and acquire the privileges necessary to perform their job function indirectly through role membership
  - Important difference to DAC - privileges are not assigned directly to users – this makes security administration more efficient
  - Role membership can be assigned and revoked easily as a person's responsibilities change
  - If business processes change privileges can be added to or removed from a role instead of making the change multiple times (for each affected user) as in a DAC system
  - RBAC is efficient because role/privilege associations change less frequently than user/privilege associations - work flow processes are relatively stable whereas user/task assignments are not

Image source: P Samarati and S de Vimercati. Access control: policies, models, and mechanisms. In Foundations of Security Analysis and Design, volume 2171 LNCS, pp 137-196. Springer, 2001.

RBAC Features

- Permissions can be more transaction or business process oriented – e.g. open account, prescribe medication:
  - this maps well to applications architected according to object-oriented principles - permissions map to methods
- Supports organisational control principles such as separation of duty – e.g. the same officer can't raise a purchase order and approve the payment.
- Supports the security principle known as 'least privilege' – particularly compared to DAC. Only the privileges of assigned and active roles are available to a user (assuming the system supports user controlled role activation/deactivation).

QUT isi Information Security Institute

Hierarchy of RBAC models

- Formalised by NIST and adopted as an ANSI Standard
  - Core RBAC – consists of users, roles, operations, objects and permissions (operation on an object)
  - Hierarchical RBAC – roles can contain other roles and inherit their permissions e.g. the role of Surgeon inherits the permissions of the role of Doctor
  - Static Constrained RBAC – adds constraints which are enforced when users are assigned to roles.
    - Static separation of duty where membership of two roles can be defined as incompatible – e.g. the same user cannot be assigned membership of both Teller and Auditor roles.
    - Cardinality constraints – e.g. only one user can be assigned membership of the role of Branch Manager at any one time
  - Dynamic Constrained RBAC – constraints enforced at the time of role activation e.g. a user cannot simultaneously activate two incompatible roles. Also dynamic cardinality to control role activations

QUT isi Information Security Institute

ANSI/INCITS 359-2004 Information Technology - Role Based Access Control International Committee for Information Technology Standards (formerly NCITS) / 03-Feb-2004

RBAC Weaknesses

- Role engineering is hard – determining the permission set required for a role is not an easy task
- Ensuring fine access control granularity and enforcing the principle of least privilege means the number of roles typically increases rapidly but this makes administration harder. Thus the tension between security and simple, efficient management remains
- RBAC lacks the user-driven flexibility of DAC – this is a strength from the perspective of policy enforcement but the organisation must be able to formulate a workable policy. This can be hard in dynamic environments
- Many business processes are complicated – the question of whether a role should hold a specific privilege is often determined by complex rules which need to be dynamically evaluated at runtime
- RBAC support in many commercial products is still very basic – mostly limited to core RBAC – thus many of the advantages of hierarchies and constraints have failed to materialise

QUT isi Information Security Institute

QUT isi Information Security Institute

Access Control for Information Sharing

a university for the real world

Information Sharing

- The term *information sharing* typically refers to policy-controlled information flow between security domains
  - Thus, at least two distinct organisational entities are involved, each of which has its own security policies, procedures, and systems
  - Differences and potential incompatibility between participant's policies, procedures and systems makes controlled information sharing a difficult problem
  - Control over information sharing is important because an organisation that releases sensitive information to another still needs to know that the information will be protected, even though it is no longer under its direct control

QUT isi Information Security Institute

Why is Information Sharing Important

- Heightened concerns about national security since 9/11
  - Counter terrorism efforts are based in part on sifting large volumes of information from many different sources to identify suspicious patterns that indicate terrorist activity
  - Critical infrastructure operators (electricity, gas, water, transport, communications etc.) need to share information about attacks, vulnerabilities and emergency response plans to address cyber and physical threats
    - Critical infrastructure is highly interdependent – no sector can continue operating if another fails
    - Therefore, response to threats needs to be coordinated
- Rapid adoption of new business models that rely heavily on outsourcing and partnering

QUT isi Information Security Institute

Authorisation for Information Sharing

- **Information Sharing Challenge** - to make sure that the right people and processes have quick access to the right information without exposing a risk that information might leak to unauthorised persons
- Info sharing networks typically host a large number of users from 100's of different organisations
- Users have access based on 'need to know'
  - 'Need to know' is based on their legitimate objective - what they are trying to do – linked to their role in the organisation
  - Access rights need to change as roles/objectives change
  - Rights need to be revoked when employment changes
  - Rights need to change as situation changes – eg in response to a crisis event
- Major Challenge: keeping access rights for individuals up-to-date
- Current authorisation models are inadequate for dynamic environment of information sharing

MITRE JASON Report on Information Sharing

- Report considered information sharing to protect national security
- Key finding: Traditional authorisation models/policies are too rigid to allow for recent emergence of information sharing.
- Organisations have resorted to various ad-hoc means to share information:
  - users have been granted near-blanket access rights or "temporary" authorisations that are never revoked;
  - Report concludes that new access control models are needed that better support dynamic, collaborative environments

MITRE Corporation Jason Program Office. Horizontal integration: Broader access models for realizing information dominance. Technical Report JSR-04-132, MITRE Corporation, 2004.

Authorisation Models for Information Sharing

- First, a snapshot of what we are investigating - Objectives-Based Access Control
  - Traditional approaches – trade-off between competing objectives occurs outside the model
  - Result is a static policy
  - We are investigating how to include objectives and late trade-off within the model
  - Aim is to support dynamic policy that is sensitive to changes in opportunities and threats in the environment
  - Based on decision theory and related techniques from field of economics

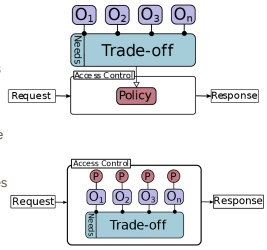

Now in more detail ….

- How to determine who must be authorised?
  - Two questions must be answered:
    1. What are the system's objectives?
    2. Who must be given authorisation based on the objectives?

Traditional Perspective

- Information Systems
  - Used to be isolated.
  - Users were known.
    - They were considered within the system.
- So one could determine:
  - Who must do what (i.e., job functions), and
  - This was the basis for the given access.

Requires a Perspective Change

- Information Systems as a collection of interconnected entities:
  - Users (other entities) are outside the system.
  - They change independently (e.g., their reputation, motivation, objectives …).
- We are shifting from closed isolated systems toward interconnected collaborating entities.
  - There is a shift from the defensive posture of need-to-know which emphasizes confidentiality TO:
  - need-to-share which aims to capture greater benefits through wider information availability (whilst controlling risk)

Dynamic & Uncertain Environments

- With such interconnectedness comes:
  - Dynamism: changes may happen frequently.
  - Uncertainty: changes are unpredictable.
- But what does this mean for an authorisation system?

Dynamic Posture – Private Sector Actor

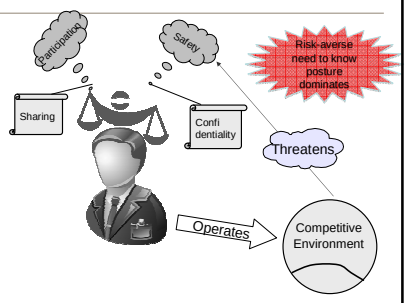

Dynamic Posture – Private Sector Actor

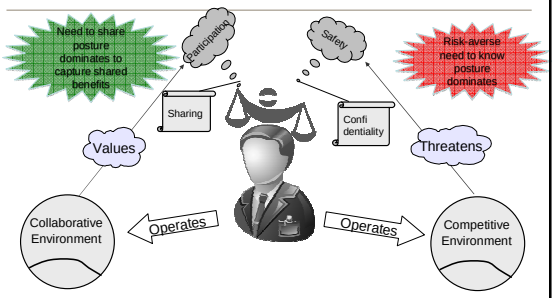

Missing Link!

- The link between objectives, policy, and access decisions is outside the scope of access control systems.
- Basis of current access control systems: an already-made trade-off between a range of system objectives.

Weakness of Existing Access Control Models

- Existing models assume the objectives of the system are known.

Weakness of Existing Access Control Models

- They all assumed the objectives of the system are known.
- A trade-off analysis is made before hand to address *needs*.

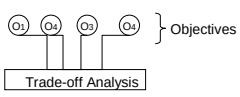

Weakness of Existing Access Control Models

- They all assumed the objectives of the system is known.
- A trade-off analysis is made before hand to address needs.
- A set of rules (policy) can be made based on the trade-off.

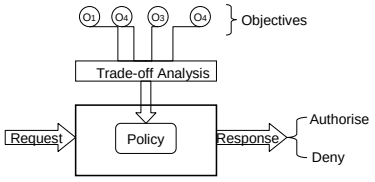

So what is wrong with this?

- Probably nothing:
  - For deterministic, static environment!
  - Or dynamic and uncertain environments, where the system simply does not care about changes.

So what is wrong with this?

- However with a couple of realistic assumptions:
  - Access control is actually part of a system with more than one objective.
  - Environment may change and that will effect what becomes important:
  - (changing needs → changing authorisation posture).
- Then: regardless of how comprehensive the policy, it cannot take account for all unforeseen circumstances. There will be unexpected needs that demand an exception.
- Ad-hoc exceptions result in unaccounted risk.

Assumption of new approaches

- Recent proposals aim to address these problems (e.g. Cheng et al.)
- They assume:
  - Environment is dynamic and uncertain.
  - System must respond to the unpredictable changes of the environment, forces/incentives to:
    - Complete a mission,
    - Address national security,
    - Respond to an emergency,
    - Be commercially competitive.

P Cheng, P Rohatgi, C Keser, P. Karger, G Wagner, and A Schuett Reninger. Fuzzy multi-level security: An experiment on quantified risk-adaptive access control. In IEEE Symposium on Security and Privacy, pages 222–230, 2007.

Risk based approaches

- Rough estimate of risk for each access request
- Non-binary access control decisions
- Allow, Deny, Allow with conditions and risk-mitigation.
  - Boundaries depend on system-wide risk tolerance.
  - Limit individual's risk-taking by ***risk budget*** and post-access consequences such as auditing.
- Risk is charged against the risk budget.

P Cheng, P Rohatgi, C Keser, P. Karger, G Wagner, and A Schuett Reninger. Fuzzy multi-level security: An experiment on quantified risk-adaptive access control. In IEEE Symposium on Security and Privacy, pages 222–230, 2007.

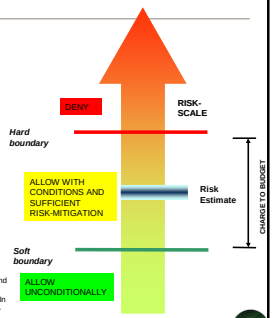

Image Credit: Cheng et.al

Risk Market

- An organisation releases to its internal market the total amount of risk units available to employees:
  - This is based on how much risk an organisation is willing to take.
- Risk Budget:
  - Every employee is given a certain amount of a common currency to purchase risk units from the risk market (personal risk budget)
- Opportunity:
  - An employee may see an opportunity (possibility for maximising what the organisation values).
  - However, the current conservative access level does not allow the employee to execute the necessary actions.
  - Traditional access control models would stop the employee right here.

Risk Market (contd.)

- Given their risk budget an employee can purchase the required risk units that allow access to the required information.
- Employees can pursue perceived opportunities and if it turns out to be beneficial for the organisation, a reward can be given (e.g., more risk budget)
- Note that:
  - There is only a limited amount of risk in the market.
  - The employee risk taking behaviour is constrained by their risk budget.

Shortcomings

- Designed for authorisation based on Multilevel Security (MLS) used in military and intelligence circles.
  - models assume information recipients have a security clearance – used to estimate unauthorised disclosure risk
  - not all users (in commercial sectors) have a clearance – e.g. emergency services and other civilian personnel – other estimates of risk required
- Designed to address under-entitlement problem rather than over-entitlement (we will discuss this later)
- The notion of magnitude of risk is static – derived primarily from the gap between user's clearance and object's classification
- The methodology to determine budget is not explicitly discussed
- The notion of external punishment and reward is assumed – outside the model

Our Research Directions - Authorisation Models for Information Sharing

- Apply techniques and theory from economics to evaluate costs and assist in making trade-offs
- Apply game theoretic techniques to formally analyse user's incentives and payoffs
- Apply budget based approaches to non-military settings
- Analyse complexity and applicability of this approach

QUT isi Information Security Institute

An Approach To Access Control In Dynamic Environments

Ed Dawson
Farzad Salim
Jason Reid
Uwe Dulleck

a university for the real world

Agenda

- Access control: authorisation perspective
  - difficulties we face in constructing a *correct* security policy
  - the implications of making a closed world assumption
- Motivating scenario
  - some recent statistics about Insider's problem
- Related work
- Main objectives
- Budget-based approach to access control (RBAC model)
- Security implications
- Future work
- Conclusion

Access Control: Authorisation

- Resource allocation problem: Given a set of resources and a set of users
  - who should access which resources and sometimes how much of these resources?
- The objective of the exercise varies:
  - to ensure confidentiality or integrity of *information resource* is preserved
    - Why? Fearing *undesirable consequences*
  - to maximize profit (reduce cost) if the resources are of economic value
    - Internet Bandwidth, Storage, Printer, etc.

When is it easy?

- When the resource provider (employer - she) is *well* informed about the operational needs of the user (employee – he)
  - She can predict exactly what resources the employee needs to perform the job
  - Request for extra resources can be safely ignored
- Or When there is a *perfect* usage monitoring mechanism in place, hence users can be held *liable*
  - So misuses of the resources can be detected and punished
- Or when there is no conflict of *goals (incentives)* between the resource provider and the user
  - Hence, it is as if the resource provider is performing the job

Real World Complexity

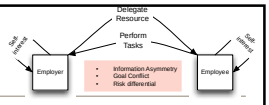

- Employer only has *incomplete information* about the resource the employee needs to complete his task
  - Sometimes the user even has a more realistic view about what resources are required to perform a task
- Our monitoring, detection and audit mechanisms are *imperfect*
  - Sometimes we don't even find-out a resource has been misused (stolen)
  - Sometimes we find out and its too late, or proofs are flimsy and users cannot be held liable
- Users are human beings – they are *self-interested* and act *strategically* to increase their payoffs
  - Most of the time misusing resources is attractive, even if it leads with some probability to being fired or prosecuted

Motivating Scenario

- Assume we have a hospital with a collection of sensitive information to protect (for simplicity assume privacy reasons)
- Our hospital has a set of individuals working as
  - doctors, nurses, interns, paramedics, etc.
  - Some full-time employees some part-time, working only one day a week, etc.
- These individuals need access to segments of patient's record and some times other relatives (not the patients) records in order to provide diagnosis
  - What segment and who else's information is involved can only be decided at the time of giving care
  - The best judge of this is the care provider who is examining the patient

Motivating Scenario (cont.)

- assume we have already adopted a database with support for RBAC (Role-based Access Control)
  - so we introduced a collection of roles (job functions)
  - associated them with a set of operations that can be performed on the records (permissions)
  - and finally we associated our employees with the roles

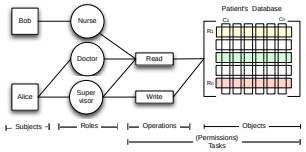

Motivating Scenario (cont.)

- so SQL queries (e.g., SELECT (read), UPDATE (write), etc) to tables (e.g., finance, patients, etc.) in our databases are can be intercepted
- and only if there is a user-role-permission link predefined in the RBAC policy, is the employee authorised to access (operation) the information (object)

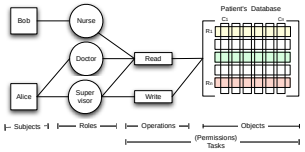

Motivating Scenario (cont.)

- The problem is, as naïve as it sounds, we can't say for certain whether an employee's request has to be **denied**:
  - This is in contrary to other areas (military) where a closed-world assumption is more realistic to make (deny unless authorised)
  - Recall that the care giver has information advantage – information asymmetry, they know what records they need
  - there is also an aspect of time criticality, patients diagnosis must not be delayed because our ill constructed policy denies the access!
- The irony is, the information that we collect can be very sensitive and can't be left unprotected either!
  - So in summary, we don't now *exactly* who has to be denied access, but we know that NOT every request has to be allowed either

Motivating Scenario (cont.)

- Despite our uncertainties about who exactly needs what we can say the following about appropriate access:
  1. We know that the misuses of some records (objects) have more severe *undesirable consequences*
     - for example, we have health records of some well-known individuals, and the leak of these records has more severe reputational damage than normal individuals (leading to monetary loss)
  2. We can predefine *only approximately:*
     - the roles and their permissions
     - the employees to assign to these roles based on: their *job function* NOT trustworthiness (roles in RBAC implicitly encapsulate both)
       - for example, an employee, an administrator or a nurse may change their attitude (become destructive) if informed about being laid off, even though her role (job position) is still the same.
     - the trustworthiness is more volatile

Motivating Scenario (cont.)

- Despite our uncertainties about who exactly needs what we can say the following about appropriate access:
  3. We also have some sense of *average usage frequency*
     - for example, we know by experience that a full-time nurse, usually provide care for about 50-80 patients per week for GP's this number is between 70-100, while surgeons attend can attend to between 20-30 patients, etc.
  4. Finally we have some knowledge about the *toxic combination* of permissions
     - for example, we know that having access to patients Full Name, Address and Sexual History may have more *undesirable consequences* than accessing only one of these information

In a Nutshell

- We know now that regardless of how much time we spend and how much analytical effort we put in constructing a policy, it will:
  - *under-entitle* some users
  - *over-entitle* some others

Under-entitlement

- **Under-entitlement:** users legitimate access to resources may be denied:
  - Leads to the loss of productivity
  - In a hospital emergency case may interrupt providing care and put patients life in danger

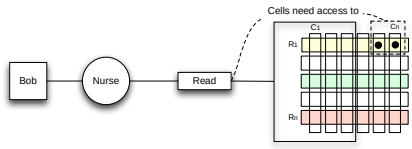

Over-entitlement

- **Over-entitlement:** some users usually acquire excess of permissions that can be misused

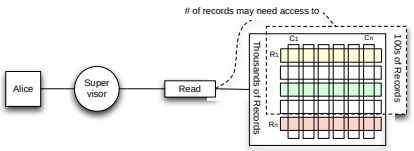

Optimal policy

- Why is it hard to *construct* and *maintain* an optimal policy?
- By optimal we mean:
  - to allocate each user the level of access they need to do their job - **no more** and **no less**
- The straight answer is the *information asymmetry* between the policy writer and the users,
- Why there is such an information gap?
  1. Environmental Unpredictability, and
  2. Human (behavioral) unpredictability

Sources of Uncertainty: Environmental Factors

- Although the sheer number of employees, roles and objects to secure adds to the complexity of specifying a optimal policy (management issue – not interested)
  - The real problem as Sinclair et. al., puts it is the *dynamicity* of the environment.

"During a few months of the review, one business group of 3,000 people witnessed 1,000 changes to organizational structure; in the space of a few weeks, 158 users in another group had changed job positions."

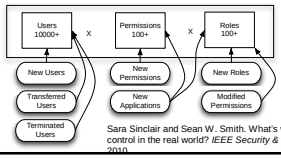

Sara Sinclair and Sean W. Smith. What's wrong with access control in the real world? *IEEE Security & Privacy*, 8(4):74–77, 2010.

Sources of Uncertainty: Environmental Factors

- Even when we assume the policy can be updated to reflect the changes, there are always access needs that are *unpredictable*:
  - In a hospital emergency access to patients health record may be needed by Interns (who normally don't get to access such information)
  - A fire in a building may require uncleared individuals to acquire access to sensitive information
- And in a real setting you can't ignore unpredictable needs either. So your policy may theoretically ensure security but it is not necessarily optimal!
  - Since security is not the only objective

Sources of Uncertainty: Human Factor

- Permissions are eventually executed on behalf of human users
- Humans are *self-interested* individuals
  - Does not necessarily mean they are malicious
- Self-interested users may change their behaviour with respect to their preference (or payoff function) that is ***private***
  - Authorised users may misuse their permissions for personal benefit, e.g., steal customer records (**Insiders Problem*)***
  - regardless of the accuracy of the vetting techniques we can only "guess" how a user will behave
- Interesting thing is, the optimality of any policy is directly dependent on how the authorised users choose to act!

How Current Models Address Under-Entitlement

- Since traditional authorisation models/policies are too rigid to allow for emergency conditions, organisations have resorted to various ad-hoc means:
  - users have been granted near-blanket access rights or "temporary" authorisations that are never revoked (MITRE JASON)
  - There is a common conception that there is correlation between the length of an individual's employment and the number of permissions they hold

MITRE Corporation Jason Program Office. Horizontal integration: Broader access models for realizing information dominance. Technical Report JSR-04-132, MITRE Corporation, 2004.

How Current Models Address Under-Entitlement (Cont.)

- Or an Exception Mechanism is adopted:
  - So when a user's access request is denied by the access control model (RBAC), user can flag it as an exception and proceed with access.
- Existing flexible models such as *break-the-glass* use this approach
  - Assumption is through appropriate audit and recovery mechanisms employees misuses of exceptions can be detected and rolled back
- This is very common in healthcare systems
- But it has lead to abundance of exceptions

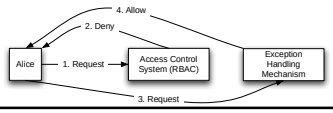

Abundance of Exceptions

- A field study of 8 Norwegian hospitals that had implemented RBAC system found:
  - 74% of the staff were assigned the permission to override denied access requests
  - 54% of active health records (i.e. those accessed in a one month period) had been accessed through this exception mechanism
  - 17% of all record accesses occurred through the exception mechanism
- It seems now that normal access has become an exception now!

Lillian Røstad and Ole Edsberg. A study of access control requirements for healthcare systems based on audit trails from access logs. In *Computer Security Applications Conference, 2006. ACSAC '06. 22nd Annual*, pages 175 –186, 2006

Misuse Detection Problem

- Exceptions increase the risk of misuse
  - more access, larger misuse probability
- Use of exceptions has to be analyzed to ensure they have been used appropriately
  - but in reality, administrator's resources (time) are scarce!
  - plethora of exceptions makes it very hard to investigate (e.g., through access log analysis) and identify the misuse cases (reduces *verifiability*)
- **Less verifiability** leads to **reduced user's liability**
  - this acts as a positive feedback for opportunity seeking employees who wouldn't have used exceptions to misuse resources, if they didn't think they could get away with it!

Insider Problem: Some statistics

- When there is little or no liability:
  - we will start having ***insider problem!***
    - Insiders are those *authorised* users who misuse their permissions for personal benefit – monetary, revenge, etc.
- According to CSI/FBI Computer Crime (2005) report that survey various industries (health, banking, etc.):
  - 56% of respondent reported some sort of insider misuse – the remaining 44% simply did not know where there has been!
  - insider's misuses accounted for 54% of the total losses due to attacks – about US$70 billion
  - The survey suggested that the actual loss could be even larger as many institutions do not report such attacks due to bad publicity… (more active government interventions needed?)

Insider Problem: Some Statistics (cont.)

- The recent 2011 Verizon Data Breach Investigations Report suggests a insiders' share of overall attack is reduced to **17%**
- However the report mentions:
  - Not that the number of insider attack is reduced!
  - The number of external attacks has increased substantially

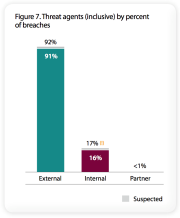

Diagram Source: 2011 Verizon Data Breach Investigations Report

Insider Problem: Some Statistics (cont.)

- Another interesting point:
  - nearly all internal breaches (**93%**) were the result of deliberate malicious activity rather than an unintentional and accidental misuse

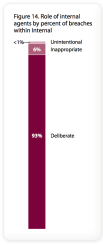

Diagram Source: 2011 Verizon Data Breach Investigations Report

Insider Problem: Some Statistics (cont.)

- Safe from no one:
  - Also, anyone in the origination despite their role, clearance or presumed trustworthiness may misuse their permissions!

Table 10. Targets of social tactics by percent of breaches within Social

Regular employee/end-user	80%
Finance/accounting staff	33%
Human resources staff	30%
Customer (B2C)	8%
Executive/upper management	5%
Helpdesk staff	3%
System/network administrator	1%
Unknown	1%

Diagram Source: 2011 Verizon Data Breach Investigations Report

Insider Problem: Some Statistics (cont.)

- Finally, **79%** of attacks take **more than weeks** to be detected!

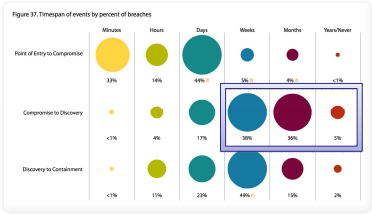

Diagram Source: 2011 Verizon Data Breach Investigations Report

Research Objectives

1. To be able to specify an *upper-bound* on the damage any user may inflict, regardless of their role and assumed trustworthiness
   - currently, a role with 'SELECT' permission to a database can virtually make a 'record dump' of all the records
   - quality is important – recall some records are more important for us (can hurt us more)
   - quantity is also important as well, e.g., thousands of lost record makes its way to news

Research Objectives

2. To align users incentives to observe the least 'privilege principle'
   - currently, an authorised users perspective it does not make any difference to copy all the records from the database or a single record
     - A: "SELECT * FROM database" and B: "SELECT * FROM WHERE patient-id = #" are equal from the users perspective
     - but from security perspective B can be very risky in terms of the amount of information provided to the user
   - there must be some sort of burden put on users to communicate the potential cost their actions expose the organisation to.

Research Objectives (cont.)

3. To allow users to gain permissions that have not been pre- assigned to them (i.e., due to the incomplete knowledge of the administrator)
   - a systematic approach is needed rather than an ad-hoc mechanisms to allow the unaccountable exceptions to be allowed.
4. To facilitate misuse monitoring and detection,
   - misuse detection and monitoring is currently a separate machinery unrelated to the access control models (RBAC)

Our approach in a Nutshell

- We use a budget based approach in RBAC setting
- Budget is a proxy for administrators uncertainty about which permissions (operation) on objects the employee exactly needs
- Observe that there is NO direct link between the 'reference monitor' and the 'RBAC policy'
- The reference monitor is only concerned about the availability of user budget
- No external punishment/reward machinery outside the model is assumed

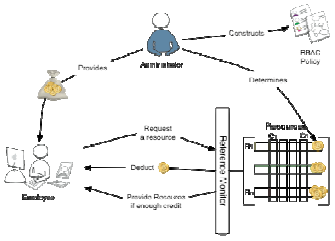

Our approach in a Nutshell

- Questions:
  - What does cost mean?
  - how to determine this for permissions?
  - how much budget to provide to employees?
  - how much should they be charged for access?
  - what problems does it solve?

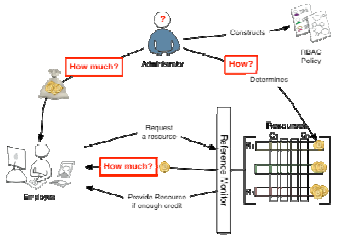

Basic Questions

- Back to our hospital example
- Should we allow Bob to access these extra cells for which he has not been given permission?

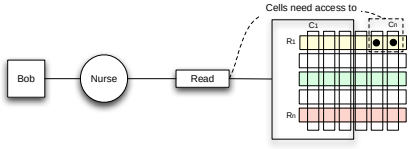

Basic Questions

- How about Alice, should she be simply allowed to copy 1000s of records, for which she already has permission for?

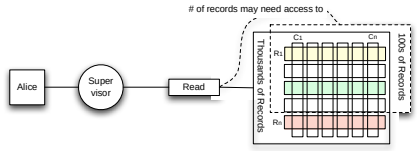

Contemplating an answer to the question

- We need to ask:
  - What is the **expected undesirable consequence** of providing the access!
  - In other words, realistically,
    - it is not so much *who* is requesting to access the resource,
    - but *how* can they misuse the resource! – what are the possible ways it can be misused and the *consequences* of each usage for us
    - its not enough however to have the list of undesirable consequences from misuse,
    - the *probability* of occurrence of these undesirable states need to be known to determine the expected value (risk exposure).

Can we disentangle the question?

- Obviously, determining the probability of possible consequences is not easy – very contextual, also subjective
- Let us put aside the question of what is the subjective probability of each undesirable consequence
- Can we then say anything useful about consequence alone?
  - commonsense – only those we consider to be probable at all: $0 < c <= 1$

Worst Possible Consequence is Useful

- So we are left with
  - list of consequences of a permission (operation on an object),
  - with no probability for occurrence of each consequence
    - so no regard to *who* is actually using the permission
- This is actually the part we are relatively good at
  - we usually estimate the *worst possible* outcome of a decision
- However we are not so good with probabilities

Permission's Maximum Cost

- So what is the *maximum cost* of a permission
  - when resources have an *intrinsic* value it is intuitive:
    - In controlling access to resources such as printer, the cost of a permission "print a document" can be:
      - the unit cost of a print per page X the number of pages printed.
    - In controlling access to limited resources such as bandwidth where quality of service is important the cost of usage may be driven from
      - The marginal cost of the facility, and
      - Extra premium for cost imposed on others, e.g., excessive crowding

Permission's Maximum Cost

- Maximum cost of an operation on *information resources* that do not have an *intrinsic* value can also be determined by the same logic, even though it may be less intuitive.
  - The value of these resources depends on the cost of most undesirable misuse!
    - the cost of undesirable misuses are is application dependent
    - they can be the cost to:
      - reconstruct lost data,
      - restore the integrity of the fabricated or intercepted data or
      - pay the functional liabilities for public disclosure of confidential or private data

Example: Permission's Maximum Cost

- Consider our hospital with one table of patient records, that is only concerned about complying with Government privacy Act.
  - Consider a legislation is in place that for each customer record stolen the hospital incurs $2 fine

Example: Permission's Maximum Cost

- So that without any knowledge about the identity (or intention) of who is accessing a record we can estimate the maximum cost of a permission:

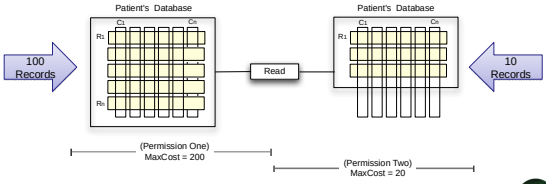

Implications: Permission Comparison

- The explicit assignment of maximum cost to permissions, even though an approximate measure, has two important advantages:
  - It quantifies (even if approximate) the potential upper-bound cost that any operation may incur.
  - more importantly, it provides a basis for a **relative comparison** of permissions

MaxCost(P1) > MaxCost(P2)

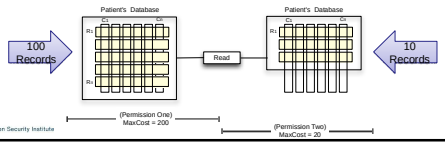

Missing In RBAC

- Note that in RBAC:
  - **no cost** associated to permissions
  - the model cannot distinguish between the level of harm two permissions may cause
  - e.g., from RBAC perspective there is no way to compare 'dumping a whole table of records' to 'reading a single record'

Granularity of Cost Function

- The cost of permissions can be very fine grained:
  - depending on the completeness of administrator's information
  - consider our previous example; we know
    - the patients record has a **column** of 'History of sexual diseases'
    - and there are celebrities as our patients
    - these *factors* are essentially *multipliers* for the cost of permissions accessing these cells

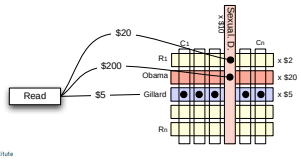

Deriving Role's Weight

- In RBAC the unit of decision making is role
- Roles are simply **grouping of permissions** that can be assigned to users
- They may differ in terms of:
  - **quantity** of permissions (i.e., in traditional RBAC), and **now**
  - **quality** of their permissions as well
    - the extent of undesirable consequences from the misuses of its permissions
- The question is: how to nudge users to use less costly roles when possible:
  - our goal is similar to the well known concept of keeping '*root*' or '*administrator*' accounts (in operating systems) for 'administrative' permissions only.
  - so far this has been just a recommendation – No way for the model to enforce it.

Nudging Users: Least Privilege Principle

- we use arithmetic summation over the cost of role's permissions to derive role's weights
- then we use the weight of a role as the *multiplier* for the permissions that is being accessed through the role (details in our paper)
- as the result, using a permission from a cheaper role turns out to be cheaper
  - For example: It is cheaper for Bob to execute t3 through r3 than r2.

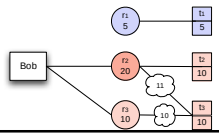

Administration: Side Effect of Role's Weight

- Since the cost of a role is proportional to the number of its permissions, for users who are assigned to roles, the unnecessary permissions are no longer considered as "free permissions".
  - this is in contrast to current practice, where it is beneficial to users to overestimate the permissions they need to perform their job and demand that administrators assign as many permissions to the roles as possible (i.e., permission hoarding).
    - E.g. t4 costs \$10 through r4 instead of \$6 if t3 and t5 are removed from the role

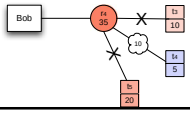

Incorporating Exceptions

- How about if Bob is attempting to access t_1, which has not already been assigned to him by the administrator (through user-role assignment)?
  - With current flat pricing there is no difference between the normal access or an exception
  - How to ensure the burden of acquiring exception is carried by the user (Bob)?

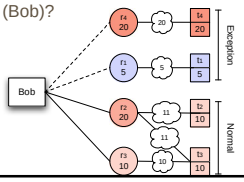

Incorporating Exceptions: Factoring in Probability of Misuse

- We *price discriminate* between users based on RBAC policy – crafted by the administrator – approximation of operational needs
  - For normal access, Bob receives a *discount* on the price of the permissions ($\alpha \in [0, 1]$)
  - To provide a means to make some of exceptions impossible or very costly, we also introduce a *tax* on exceptions ($\beta \geq 1$)
- For example, assume $\alpha = 0.1$ and $\beta = 1$

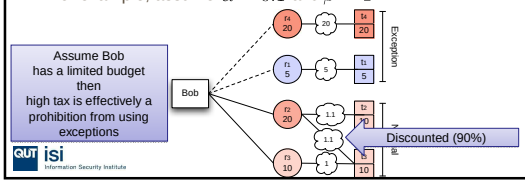

Rate of Tax and Discounts

- Determining the tax/discount rate is application specific and can be very elaborate by taking into account several factors including:
  - the roles a user already possesses
  - the relevance of these roles and the role that is being used to make the exception possible.
- e.g. a doctor's exception to read the record of a patient's parents (search for potential genetic causes) may be considered as "relevant", hence taxed substantially less than a finance manager who is requesting the same exception
- role mining techniques that provide the quantitative notion of *distance* between roles may be adopted for deriving tax rates.

User's Budget: How Much To Allocate

- Budget a virtual currency allocated to users by the administrator
- Allocated to users periodically to pay for the permissions, at the time of access
  - In general budget is estimated with respect to the *frequency* and *cost of permissions* users (or more abstractly roles) may need for a given period (e.g. week)
    - E.g. a fulltime nurse on average can attend 50 patients, so
      - Nurse's budget = cost of reading a record X 50
- We assume
  - user's budget is *limited* and *non forgeable*,
  - administrator is trusted!

User's Budget: How Much To Allocate

- To determine a user's budget for a period:
  - for each role that has been associated to the user (in RBAC policy)
    1. we know the normal frequency of the permissions and the cost of permissions of the role
    2. so this gives us the base budget that members of the role need to perform their job
    3. then we multiply that budget by the discount rate
       - this ensures high power users (e.g., administrators) are not allocated a very large budget that can be used for exceptions
  - But wait!
    - this only provide users with the budget we think they require based on their operational needs!
    - remember users behaviour may change towards misusing resources (e.g., disgruntled employees) – even though their job position remains the same

User's Budget: How Much To Allocate

- So we adjust the users budget by indicators of users behaviour change (benevolence to malicious)
  - these indicators are driven from imperfect estimators
    - we have not specified any specific machinery
    - but monitoring users action may provide some insight
      - (e.g., AZALIA is a tool developed by Bishop et al., for reasoning about (e.g., disgruntled) users misuse probability through analyzing blogs, browsing patterns, etc.)
      - output from intrusion detection systems can be another imperfect indicator of potential misuses
  - so in a sense user's budget may be reduced when we are allocating budget for the period
    - as the users propensity to misuse approaches to 1, users budget approaches to 0

Security Implications

- Effective Monitoring
  - Monitoring and analysis of users' budgets provides a uniform mechanism for:
    1. better understanding of access needs, as well as
    2. detecting misuses

Security Implications

- There may be two reasons for budgets to be exhausted before the period ends:
  1. *Erroneous Budget Allocation:*
     - The administrator may have incorrectly predicted a user's access requirements for the given period.
     - Precisely the error in the proposed model can be due to an incorrect user-role or role-permission assignments, or under-estimation of the frequency of permission usage
     - Regardless of the source of error however, the abrupt exhaustion of a user's budget and their inability to perform their job demands that the administrator maintain an accurate picture about users budget needs.
     - We envisage that in a long-run through analysing budget spending patterns, users' budget can be quantified with a sufficient proximity to the actual need.

Security Implications

- There may be two reasons for budgets to be exhausted before the period ends:
  2. *Permission Misuse:*
     - When the allocated budget is adequate, a user's budget exhaustion flags the potential misuse of permissions.
     - This feature can improve the efficiency of monitoring and audit as the administrator can focus on those users whose budget has been exhausted, rather than needing to audit and verify all accesses or exceptions.
     - Note that, not only misuses can be detected when budget is exhausted, but also when the users 'remaining budget' to 'remaining duration' ratio falls below a threshold (e.g., 0.2).
     - The administrator can also focus on monitoring those exceptions which have a tax rate above a defined threshold.

Security Implications

- Addressing Impersonation Attacks
  - An outsider may acquire the credentials of an employee and access the system,
    - by guessing or key logging a password, or through social engineering means, etc.
    - The consequences of a successful impersonation attack in a traditional access control model can be devastating as such attacks are difficult to detect or prevent.
      - The adversary can access any and all resources for which the legitimate user held privileges without affecting the actual user's access capabilities

Security Implications

- Addressing Impersonation Attacks
  - The implications may not be as devastating in our model
    - Even though the attack is still possible, any access by the attacker is counted against the user's budget.
    - Hence, the users can detect the reduction in their budgets
    - Even if such detection doesn't happen, the consequences of such attacks are strictly limited by the available users' budget for the period

Security Implications

- Addressing Denial of Service (Query Flooding) Attacks
  - the malicious user sends a large number of select or update queries to a targeted database
  - Current techniques to detect/prevent such attacks require comprehensive analysis of query log files and assumptions about *normal* patterns of access that so far suffer from high incidence of false-positives
  - In the proposed model such attacks will have a little impact and will be easy to detect,
    - a user's ability to send a query is bounded by their limited budget, the queries from users with **no budget** can be intercepted by a proxy server that sits between the client and the database
    - also the exhaustion of which will lead to termination of the attack and potentially, misuse detection

Security Implications

- Addressing Escalation Attack
  - One criticism of the proposed model may be that it potentially allows malicious users to acquire unwarranted permissions.
    - Although this criticism is not only applied to our model, as escalations already happening in reality through exception mechanisms

Security Implications

- Addressing Escalation Attack
  - In our model
    - the aggregate amount of damage that may be incurred is restricted by the budget allocated to users.
    - Further, the budget allocation function is parameterized by the outcome of online monitoring mechanisms such as intrusion detection systems to adjust the users' disposable budget based on their estimated propensity to misuse permissions
    - Also the administrator has additional control over the escalations through personalising the evaluations of escalation tax, which could take into account the users' application specific factors such as trustworthiness, need, and access history into account.

Future Work

- We would like to explore what techniques can be used to estimate tax rates:
  - as we mentioned before, there are some work is being done by role engineering community where the distance between the roles are measured based on the relevance and weight of permissions.
- We would be interested in implementing and deploying a budget-based module to interact with the current RBAC security modules in database applications
- We would like to examine if our approach can address some of the problems in cross-organizational information sharing
  - it's hard for the information provider to determine what information the receiving organization needs
  - since organizations are independent entities they can change more frequently – and they may be less liable

Concluding Remarks

- Security is not 'the objective' in most real world commercial organizations:
  - maximizing profit, getting the job done on-time are far more tangible objectives
  - if the security policy conflicts with these objectives it will be by passed one way or another
- It is very difficult to know who exactly needs what resources
  - Only some approximation can be made
- Budget can be a useful proxy to deal with administrator's incomplete knowledge about users access needs
  - instead of treating an RBAC policy as the bible for decision making
  - use it as a reference to discriminate the price of permissions for users
  - roughly estimate and allocate budget to users
    - at the end of each period observe the remaining or exhaustion and refine the budget

More Information

- Farzad Salim, Jason Reid, and Ed Dawson. **Towards authorisation models for secure information sharing: A survey and research agenda.** *The ISC International Journal of Information Security (ISeCure)*, 2:67– 85, 2010.
- Farzad Salim, Jason Reid, Uwe Dulleck, and Ed Dawson. **Towards a game theoretic approach to authorisation**. In *Decision and Game Theory for Security (GameSec)*, volume 6442 of *LNCS*, pages 208–219, Springer/Heidelberg, 2010.
- Farzad Salim, Jason Reid, Uwe Dulleck, and Ed Dawson. **An approach to access control under uncertainty.** To appear in *Availability, Reliability and Security (ARES)*, IEEE Computer Society, 2011.

QUT isi Information Security Institute

Access Control & Privacy

a university for the real world

What is Privacy?

- Clark defines privacy as:
  the interest that individuals have in sustaining a 'personal space', free from interference by other people and organisations.
- Clark defines information privacy as:
  the interest an individual has in controlling, or at least significantly influencing, the handling of data about themselves.
- Information privacy and confidentiality are related but DIFFERENT CONCEPTS

QUT isi Information Security Institute

Roger Clark, Introduction to Dataveillance and Information Privacy, and Definitions of Terms. Available at http://www.rogerclarke.com/DV/Intro.html#Priv

Privacy Violation – Sources of Harm

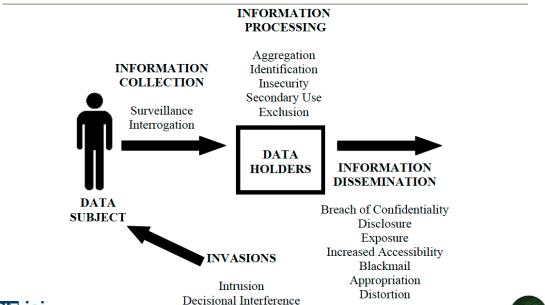

QUT isi Information Security Institute

Image Source: Daniel J. Solove, "A Taxonomy of Privacy", University of Pennsylvania Law Review, Vol. 154, No. 3, p. 477, January 2006

Privacy Legislation

- Privacy compliance is a major challenge for organisations that collect personal information
- Personal information needs to be handled according to relevant privacy regulation:
  - Generally based on Information Privacy Principles (IPPs) contained in the *OECD Guidelines on the Protection of Privacy and Transborder Flows of Personal Data*
  - IPPs regulate the way organisations collect, store, use and disclose personal information about individuals
  - IPPs impose positive obligations and constraints

QUT isi Information Security Institute

Purpose of Collection & Disclosure

- **Purpose Specification Principle** – organisations must inform individuals who provide personal information of:
  - *the **purpose** for which the information is being collected*
  - *whether the information will be **disclosed** to another agency or organisation*
- **Security Safeguards Principle** – organisations must protect records against loss, **unauthorised access**, use, modification or disclosure

QUT isi Information Security Institute

Use Must be Consistent with Disclosed Purpose

- **Use Limitation Principle** - information that was obtained for a particular purpose shall not be used for any other purpose unless:
  - *the individual concerned has consented to use for the other purpose*
  - *Another relevant exception applies (e.g. imminent threat to life, law enforcement etc.)*

Privacy Compliance

- Compliance currently achieved through:
  - The actions of people who understand the rules
  - Administrative processes
- Rules are not explicitly embedded in information systems
  - Violations possible when people are not aware of their responsibilities
  - Problem is arguably manageable when systems and processes are relatively static
  - Service Oriented Architecture/Web Services makes compliance much harder

SOA & Web Services

- **Old way**: Information "silos" and inter-agency system incompatibility **indirectly enforced limitations** over use and disclosure of personal information
- **New way**: systems are being progressively re-engineered according to principles of service orientation
- SOA supports **dynamic** service composition
  - Reduced friction in information exchange
  - More people have potential access to more information
  - Makes privacy policy enforcement more difficult

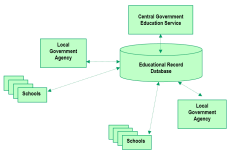

Image source: UN E-Government Survey 2008, United Nations, New York 2008, available at http://unpan1.un.org/intradoc/groups/public/documents/UN/UNPAN028607.pdf

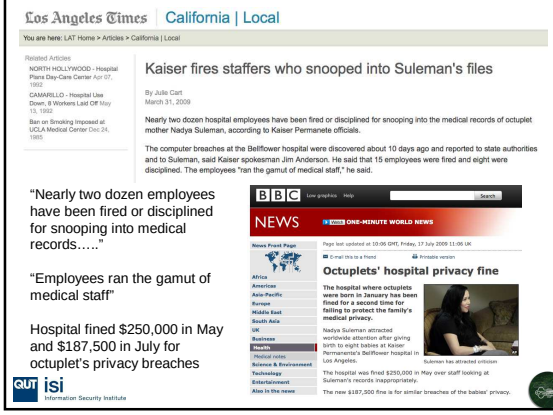
Los Angeles Times | California | Local

You are here: LAT Home > Articles > California | Local

Kaiser fires staffers who snooped into Suleman's files

By Julie Cart
March 31, 2009

Nearly two dozen hospital employees have been fired or disciplined for snooping into the medical records of octuplet mother Nadya Suleman, according to Kaiser Permanente officials.

The computer breaches at the Bellflower hospital were discovered about 10 days ago and reported to state authorities and to Suleman, said Kaiser spokesman Jim Anderson. He said that 15 employees were fired and eight were disciplined. The employees "ran the gamut of medical staff," he said.

"Nearly two dozen employees have been fired or disciplined for snooping into medical records….."

"Employees ran the gamut of medical staff"

Hospital fined $250,000 in May and $187,500 in July for octuplet's privacy breaches

How can ICT Support Privacy Compliance?

- The main opportunity: system-enforced purpose-based limitations on use & disclosure:
  - Requires a new access control model that comprehends:
    - Collection purpose – an attribute of data
    - Access purpose – an attribute of a query/request
  - Authorisation is dependent on user's intention as well as their identity
  - **How?** an enforceable privacy policy expression language

Policy Language Requirements

- Policy language must express these elements:
  1. **Users** (distinguish use from disclosure)
  2. **Actions** (read, write, update, insert, application level functions)
  3. **Personal data objects** (unique, fine-grained referencing capability)
  4. **Purposes** (collection and access)
  5. **Obligations** (e.g. delete 1 year after transaction)
  6. **Conditions** (Boolean predicates over previous 5 elements, data content and environment)
- Policy statements must be composable (requires support for precedence and conflict resolution)
- Tight integration with authorisation framework

Karjoth, G. and Schunter, M. (2002). A privacy policy model for enterprises. In CSFW '02: Proceedings of the 15th IEEE Workshop on Computer Security Foundations, page 271, Washington, DC, USA.

Other Opportunities for Supporting Privacy Compliance with ICT

- Use ICT to systematically document:
  - purpose of collection notification details
  - Provenance metadata (collection - who, when, from whom)
  - Audit trails
    - Access details via disclosure exceptions (law enforcement etc.)
    - Basis of determinations on "reasonable grounds" for disclosure exceptions e.g. threats to safety, criminal investigation
    - Detect unauthorised access
- Support access and amendment obligations
  - Catalogue personal information holdings – standard reports that retrieve all PI for an individual
- Check accuracy of personal information before use

QUT isi Information Security Institute

Privacy Policy Languages and Enforcement

a university for the real world

Necessary Capabilities for Privacy Compliance in Information Systems

1. Author privacy policies in a formal language, which enables rigorous analysis and verification.
2. Establish the correspondence between a privacy policy expressed in a formal language and the natural language version on which it is based.
3. Derive access control and audit polices to enforce formal language privacy policies.
4. Translate derived access control policies into enforceable, system specific authorisation rules and configurations.
5. Evaluate fine-grained, system-specific authorisation rules, policies and configurations to ensure they implement and enforce the required formal language privacy policy, which itself, accurately captures the natural language policy.

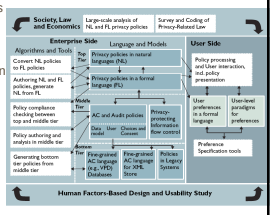

Antón, A. I., Bertino, E., Li, N., and Yu, T. (2007). A roadmap for comprehensive online privacy policy management. Commun. ACM, 50(7):109–116.

P3P - Privacy Policy Specification Language

- (Only) widely deployed privacy policy language
- XML P3P statements: data group (purpose, recipient, retention)
- E-Commerce focus evident in predefined set of purposes
- Not enforceable:
  - ambiguous semantics
  - no support for general conditions and obligations
  - not composable
  - no interface to an authorisation system
  - a language for making privacy promises – not enforcing them

The Platform for Privacy Preferences 1.0 (P3P1.0) Specification, W3C Recommendation, 16 April 2002. Available at http://www.w3.org/TR/P3P/

Enforceable Privacy Policy Languages – EPAL and XACML

- EPAL is an IBM proprietary language
  - Submitted to W3C for standardisation in 2003 but this has not progressed
  - Specialised privacy policy language
  - Rules evaluated in order - first match terminates evaluation
  - Can't compose rules from different authors without careful analysis and integration
- XACML is the leading general-purpose authorisation language
  - Ratified as a standard by OASIS in 2003
  - Privacy Profile gives it a notion of *purpose*
  - Nothing in the literature describing how to implement privacy policies with XACML and Privacy Profile
- EPAL and XACML have a lot in common

EPAL and XACML - Common Features

- Architecture based on PDP and PEP
- Policy & rule structure
  - Target: an applicability filter (subject, resource, action, environment) for policies and rules
  - Condition: Boolean predicate
  - Rule effect: permit or deny
- Abstractness
  - They don't specify how the PEP should enforce the PDP's decisions

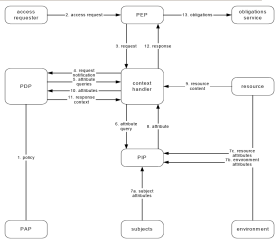

Moses, T. (2005). Extensible Access Control Markup Language (XACML) version 2.0 & Privacy Policy Profile of XACML Version 2.0. OASIS Standard, February 2005.

XACML Advantages Compared to EPAL

- Supports nested policies – inserted by reference
- More powerful rule and policy combining algorithms
- Supports multiple requesting subjects (e.g. user and software component)
- Supports policy directed error handling
- Policies can reference protected information content (XML docs and XPath)
- Multiple responses to a single query (hierarchical resources)

"EPAL contains no privacy specific features that are not already supported by XACML. EPAL lacks significant features included in XACML and that are important in many enterprise privacy situations"

Anderson, A.H. (2005). Comparing Two Privacy Policy Languages: EPAL and XACML, Sun Microsystems Laboratories Technical Report 2005-147

Research Question

- Can XACML be used to record and enforce privacy policies over data stored in relational databases?
  - Focus on application software based on web services
  - Focus on relational databases because they are still widely used and store lots of personal information
  - No examples in the literature of XACML + RDBMS + Privacy
    - XML formatted records have better coverage

Policy Enforcement for Relational Databases

- Four different levels of enforcement granularity required to enforce privacy policies
  - Table Level: allows a policy to be implemented uniformly to all data in the table.
    - How? SQL GRANT statement (SELECT, UPDATE, INSERT, DELETE)
  - Column Level: allows a policy to be implemented uniformly to all data in the column.
    - How? SQL VIEW
  - Row Level: allows a policy to be implemented uniformly to a row.
    - How? Query rewriting – add a predicate (WHERE clause)
  - Cell Level: allows a policy to be implemented individual cells in the same column.
    - How? Query rewriting - add a predicate (WHERE clause)

Hippocratic Database

- Privacy enforcing authorisation for RDBMS introduced by Agrawal et al 2002
- Hippocratic database should be responsible for protecting its own data.
- This would require modifications to SQL since it does not support a notion of purpose – can use query re-writing instead
- Model introduces:
  - Privacy Policies Table: (purpose, table name, attribute name, external recipients, retention period).
  - Privacy Authorisations Table: (purpose, table name, attribute name, authorised groups).
- Doesn't do cell level enforcement (no opt-in/opt-out)
- LeFevre et al 2004 does cell level
- Policies enforced by re-writing queries to include additional predicates (e.g. a WHERE clauses that restricts the records returned by a SELECT statement).

Agrawal, R., Kiernan, J., Srikant, R., and Xu, Y. (2002). Hippocratic databases. In VLDB '02: Proceedings of the 28th international conference on Very Large Data Bases, pages 143–154. VLDB Endowment.

LeFevre, K., Agrawal, R., Ercegovac, V., Ramakrishnan, R., Xu, Y., and DeWitt, D. (2004). Limiting disclosure in Hippocratic databases. In VLDB '04: Proceedings of the Thirtieth international conference on Very large data bases, pages 108–119.

Case Study and Demonstration System

a work in progress overview

Development credits: Sunil Ghelawat and James Mackie

Case Study Scenario

- Hypothetical State Government-run rebate scheme for solar hot water heaters
  - Opt-in options for marketing
  - Opt-in option for disclosure

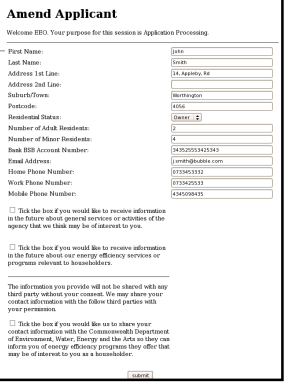

Amend Applicant

Welcome EEO. Your purpose for this session is Application Processing.

First Name:	John
Last Name:	Smith
Address 1st Line:	14, Appleby, Rd
Address 2nd Line:	
Suburb/Town:	Worthington
Postcode:	4056
Residential Status:	Owner
Number of Adult Residents:	2
Number of Minor Residents:	4
Bank BSB Account Number:	343525553425343
Email Address:	jsmith@dubbie.com
Home Phone Number:	0733453332
Work Phone Number:	0733425533
Mobile Phone Number:	4345098435

☐ Tick the box if you would like to receive information in the future about general services or activities of the agency that we think may be of interest to you.

☐ Tick the box if you would like to receive information in the future about our energy efficiency services or programs relevant to householders.

The information you provide will not be shared with any third party without your consent. We may share your contact information with the follow third parties with your permission.

☐ Tick the box if you would like us to share your contact information with the Commonwealth Department of Environment, Water, Energy and the Arts so they can inform you of energy efficiency programs they offer that may be of interest to you as a householder.

submit

Demonstration System Components

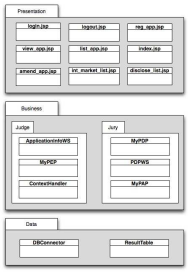

- User interface – JSP
- Business logic – Web Services JDK 6
- JBoss XACML with modifications

Database

- Backward compatible schema design – add privacy awareness to an existing set of tables
- Adopts principles of Hippocratic Database
  - Policy data stored in the database
- Policy metadata tables
  - Purposes
  - Disclosure
  - Obligations
- Supports mandatory policy (column) and opt-in policy (cell)
- Must access DB through PEP which is privacy policy aware

Table Name	Description
applicant_details	Information provided by rebate scheme applicants.
mandatory_policy_permitted_use	Lists mandatory use purposes for each applicant details attribute. These policies are not subject to modification via an individual's opt-in or opt-out choices.
mandatory_policy_permitted_disclosure	Lists mandatory disclosure policies linked to the receiving entity for each applicant details attribute. These policies are not subject to modification via an individual's opt-in or opt-out choices.
policy_permitted_uses	Consented uses for each data cell in applicant _details to which an optional-opt-in policy applies. Multiple policies can apply to a single field.
policy_permitted_disclosure	Consented disclosures for each data cell in applicant _details to which an optional-opt-in disclosure policy applies. Multiple policies can apply to a single field.
policy_obligations	Notified information handling obligations (retention, notification, audit etc.) Multiple policies can apply to a single field.
metadata_use_purpose	Usage or access purpose categories.
metadata_disclosure	Disclosure organisations/entities.
metadata_obligations	Obligation types.
applicant_details_columns	Column names for the applicant details table – used to reference fine grained policy entries in permitted uses and permitted disclosures tables.
residential_status_type	Possible values for residential status.
roles	Application roles and their identifiers.
role_privileges	Table granularity access privileges for each role describing the allowed modes of access (i.e. Select, Update, Delete, Insert).
role_allowed_purposes	Authorised purposes for each role

Purpose-Based Authorisation

Personal Information Attribute	1 Application Processing	2 Scheme Audit	3 General Marketing	4 Energy Eff Marketing	5 Energy Eff Analysis
First Name	M	M	O	O	P
Last Name	M	M	O	O	P
Address 1	M	M	O	O	P
Address 2	M	M	O	O	P
Suburb/Town	M	M	O	O	P
Postcode	M	M	O	O	P
Residential Status	M	M	P	P	M
No. Adult Residents	M	P	P	P	M
No. Minor Residents	M	P	P	P	M
Bank BSB and Account Number	M	M	P	P	P
Email address	M	P	O	O	P
Phone Number: Home	M	P	O	O	P
Phone Number: Work	M	P	O	O	P
Phone Number: Mobile	M	P	O	O	P

M=Mandatory, P=Prohibited, O=Optional

Role-Based Authorisation

- Role support based on RBAC Profile
- Fine-grained privileges are held by purposes not roles
- Not safe to infer purpose from role
- Roles are authorised to exercise purposes
- A user's session is associated with a role and a purpose

Please login

Please Select your role

Audit Officer

Please Select your intended purpose

Application Processing

Login

Collection Purpose	Description	Authorised Roles
Application processing	Information access for any purpose associated with the processing of the application and administration of the SWHR scheme	SWHR Applicant Energy Efficiency Officer
Scheme audit	Information access for the purpose of auditing the SWHR scheme to ensure it complies with the scheme guidelines and relevant organisational policies and legislative requirements	Audit Officer
General Marketing	Information access for the purpose of general marketing activities carried out by the department	Marketing Officer
Energy Efficiency Marketing	Information access for the purpose of marketing related to the department's energy efficiency programs	Marketing Officer Energy Efficiency Officer
Energy Efficiency Analysis	Information access for the purpose of analysis and planning related to energy efficiency programs	Energy Efficiency Officer

Cell Level Policies

- Database schema supports cell level policies:
  - necessary for opt-in purpose and disclosure
- XACML inefficient for cell-level policy evaluation:
  - Separate PDP request for each non-mandatory attribute for each record
  - Not practical
- We use a query re-writing approach in context handler
  - the database selects the records that match the criteria (access purpose, recipient)
  - considerably more efficient
  - XACML still plays an important role with this approach

Conclusion

- XACML can be used to enforce privacy policies
- Higher assurance of privacy compliance
- Rule and policy combining algorithms are very powerful
- Integration of privacy and role based authorisation in the same framework delivers useful features
- Fine grained access purpose enforcement presents a challenge for XACML – better handled by the database

QUT isi Information Security Institute

An Authorisation System for an Airport Information Model

Ed Dawson
Nimalaprakasan Skandhakumar
Jason Reid

a university for the real world

Project Vision

"To develop a complex systems approach for the integrated design, engineering, management and operation of airport systems"…

"To develop tools to manage airport effectiveness and balance conflicting security, economic and passenger-driven pressures" …

QUT isi Information Security Institute

Grant Structure

- 15 named investigators.
- 33 Partners.
- ~$8M Value.

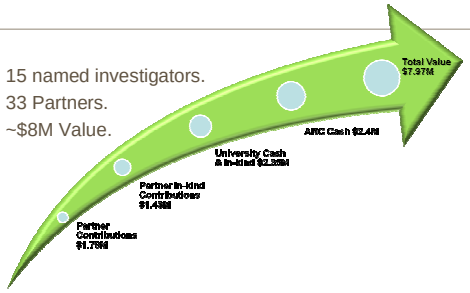

QUT isi Information Security Institute

Airports of the Future Partners

Airports
- Brisbane
- Cairns
- Canberra
- Hobart
- Kingaroy
- Mackay
- Melbourne
- Newcastle
- NT Airports (Darwin, Alice Springs)
- Perth
- QLD Airports (Gold Coast, Townsville, Mt Isa)
- Rockhampton
- Sunshine Coast
- Schiphol (NL)

Agencies
- Australian Airports Association
- Australian Customs Service
- Australian Federal Police
- Australian Quarantine & Insp. Service
- Airports Coordination Australia
- Department of Infrastructure, Transport, Regional Development & Local Govt
- IATA
- Tourism & Transport Forum
- Department of Immigration & Citizenship
- Australian Crime Commission

Service Providers
- ISS

Airlines
- Emirates
- Virgin Blue

Research
- QUT
- UTS
- Uni Melbourne
- ECU
- MIT (USA)
- TU Delft (NL)

QUT isi Information Security Institute

The Airport as a Complex System

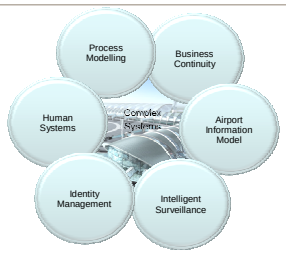

QUT isi Information Security Institute

Motivation

- Reliable knowledge of identity of passengers and staff is critical for secure and efficient airport operation
- Airport personnel identity management drives entitlement provisioning for:
  - Physical access controls
  - Airport information systems (Travel information management, baggage handling, facilities management etc.)
- Staff require access across a range of systems and physical spaces to do their jobs – assets/systems managed by different entities
- Without integration:
  - Management costs are high
  - Access entitlements are hard to keep up-to-date
  - Unnecessary privilege allocation presents a security risk
  - Users need multiple cards/credentials

QUT isi Information Security Institute

Airport Information Model

- An intelligent building information model
- BIM – a tool to aid building design, construction and management
- AIM – a tool for operational decision support, scenario analysis and facility management
- Leverage spatial cues for application and data integration/analysis
- Provide a more intuitive interface for user
- Improve decision-making support for security operators
- Support incident response

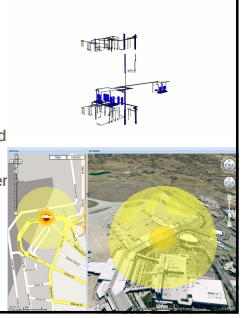

Research Objectives

- Conceptualise ideas related to spatiotemporal access control, building information modeling and converged physical and logical access control systems
- Develop an authorisation framework that uses the concept of an Airport Information Model
- Support authorisation rules based on spatiotemporal constraints
- Unify access control for physical spaces and information systems
- Enable automated provisioning and de-provisioning of access privileges
- Support required level of control over personal data to comply with privacy laws

Authorisation Framework

- Authorisation Framework for an Airport Information Model
- Explore the concept of spatial zones and logical zones
- Spatiotemporal authorisation rules
- Spatiotemporal constraints for users and resources
- Spatiotemporal reasoning in access control decision making
  - Example: can't access sensitive functions of Human Resource Management application unless PACS has admitted user to building (policy is no remote access)

Physical and Logical Access Control

- Access control to physical locations and information systems
- Merge physical security and information security operations
- Enhance security management and access policy definition
- Enable two way communications between systems in decision making

Conclusion

- Develop a proof of concept authorisation framework based on existing technologies and the specific requirements of the project.
- Conceptualise ideas related to spatiotemporal access control, building information modeling and converged physical and logical access control systems.

Research Objectives

- Develop an authorisation system:
  - To support new architecture and approach of Airport Information Model (AIM)
  - To converge Physical Access Control (PACS) and Logical Access Control into a single framework
    - Investigating PACS integration standards: OSIPS and BACNET
    - Investigating Building Information Model standards and application architecture
  - To enable authorisation rules based on spatiotemporal constraints
    - Location of the user
    - Location of the service/resource
    - Example: can't log into workstation unless PACS has admitted user to building

Some References

- Salim, Farzad, Reid, Jason, Dulleck, Uwe, & Dawson, Edward (2011) "An Approach to Access Control Under Uncertainty" In Proceedings of The Sixth International Conference on Availability, Reliability and Security (ARES 2011), 22-26 August 2011, Vienna, Austria.
- Salim, Farzad, Reid, Jason, & Dawson, Edward (2010) Towards authorisation models for secure information sharing: a survey and research agenda. ISeCure, The ISC International Journal of Information Security, 2(2), pp. 69-87.
- Salim, Farzad, Reid, Jason, Dulleck, Uwe, & Dawson, Edward (2010) Towards a game theoretic authorisation model. In: Conference on Decision and Game Theory for Security (GameSec 2010), 22-23 November 2010, Berlin, Germany.

QUT isi Information Security Institute