QUT isi
Information Security Institute

An Approach To Access Control In Dynamic Environments

Ed Dawson
Farzad Salim
Jason Reid
Uwe Dulleck

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

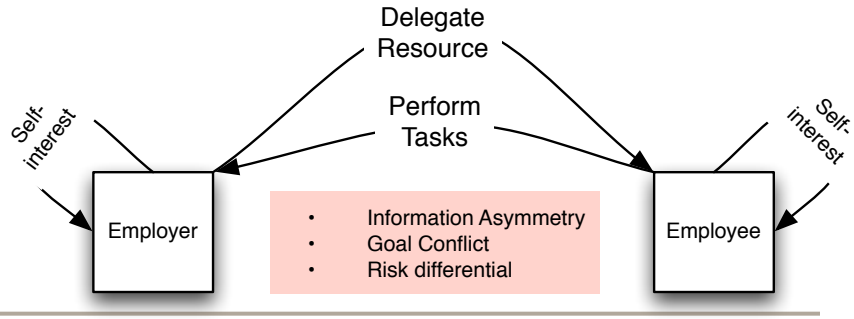

- Employer only has *incomplete information* about the resource the employee needs to complete his task
  - Sometimes the user even has a more realistic view about what resources are required to perform a task
- Our monitoring, detection and audit mechanisms are *imperfect*
  - Sometimes we don't even 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

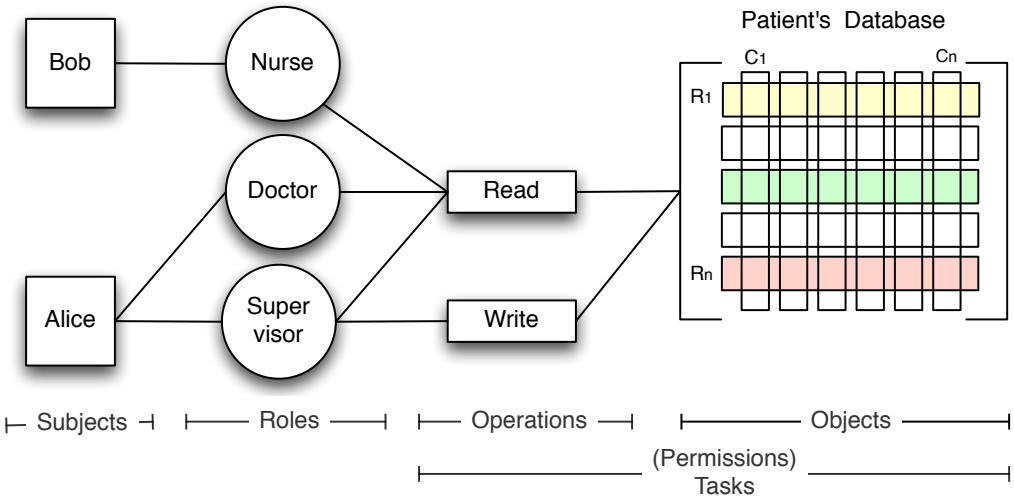

Motivating Scenario (cont.)

- 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)

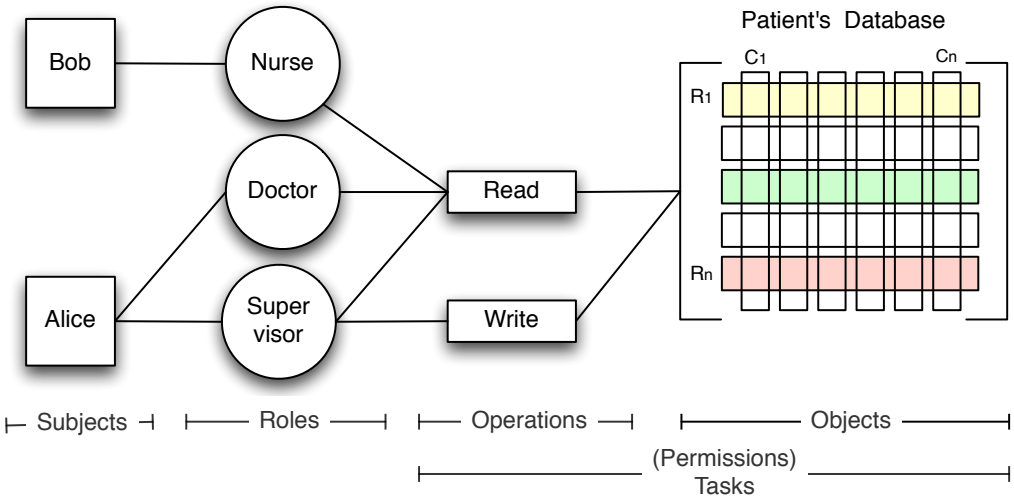

Motivating Scenario (cont.)

- 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

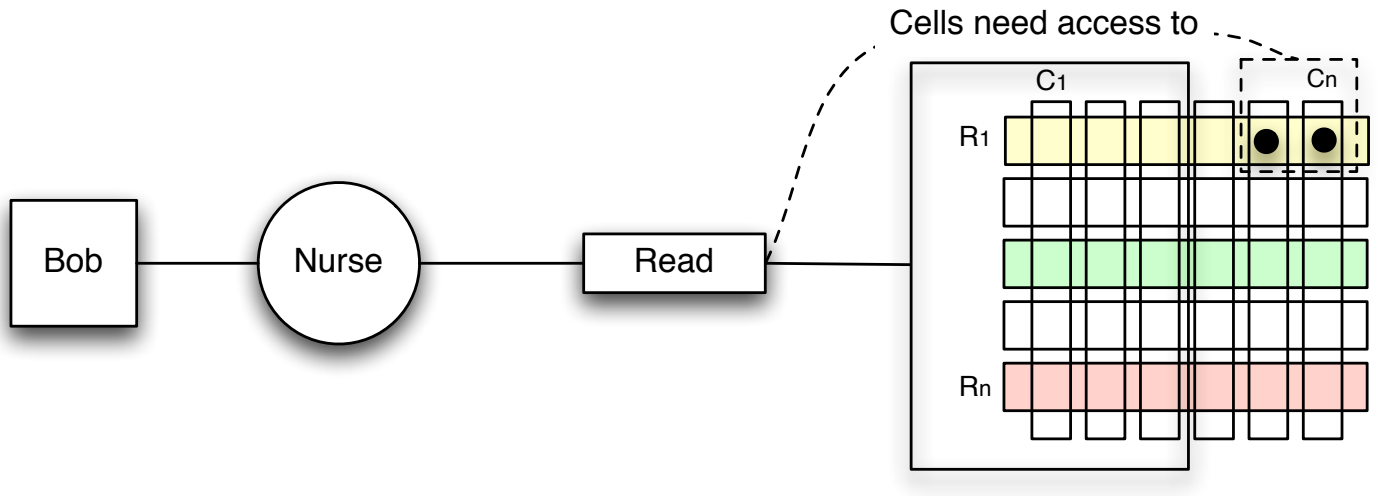

Over-entitlement

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

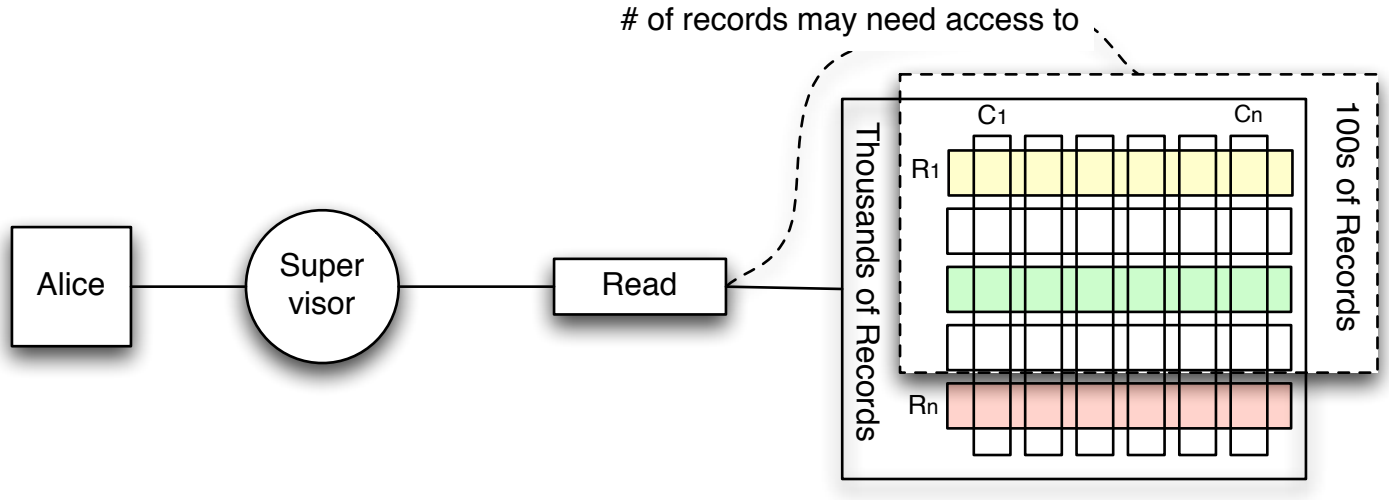

Optimal policy

- 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."

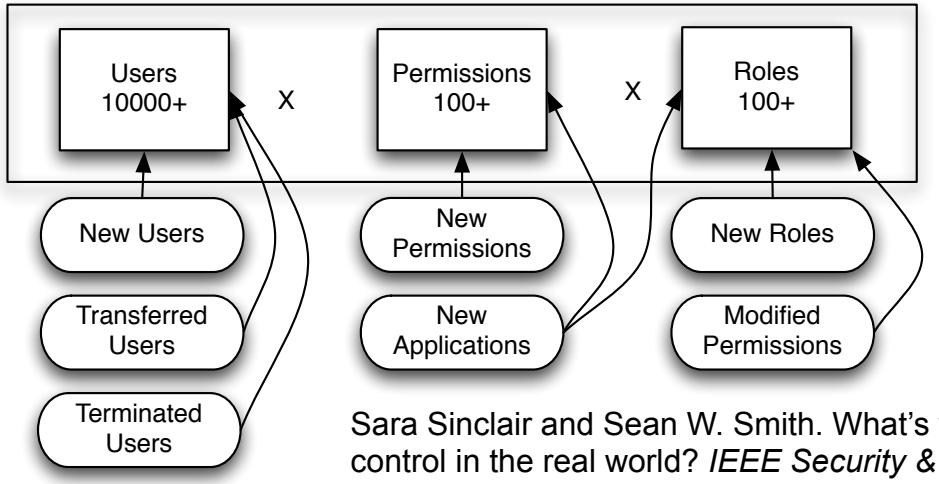

Sara Sinclair and Sean W. Smith. What's wrong with access control in the real world? *IEEE Security & Privacy*, 8(4):74–77, 2010.

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

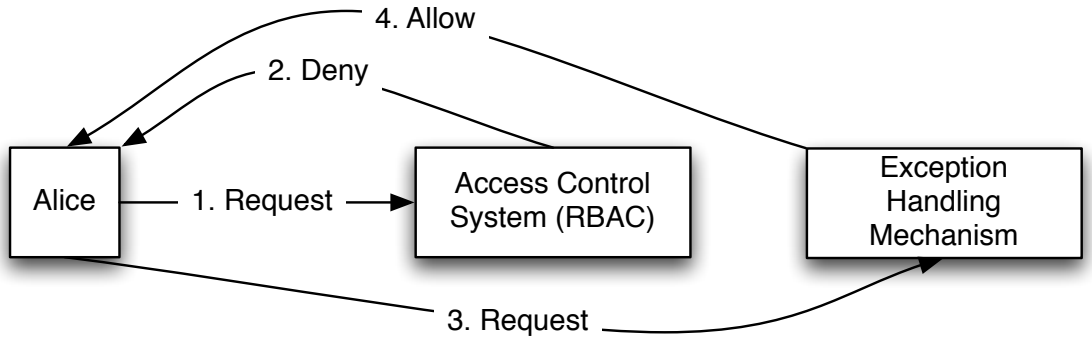

Abundance of Exceptions

- 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

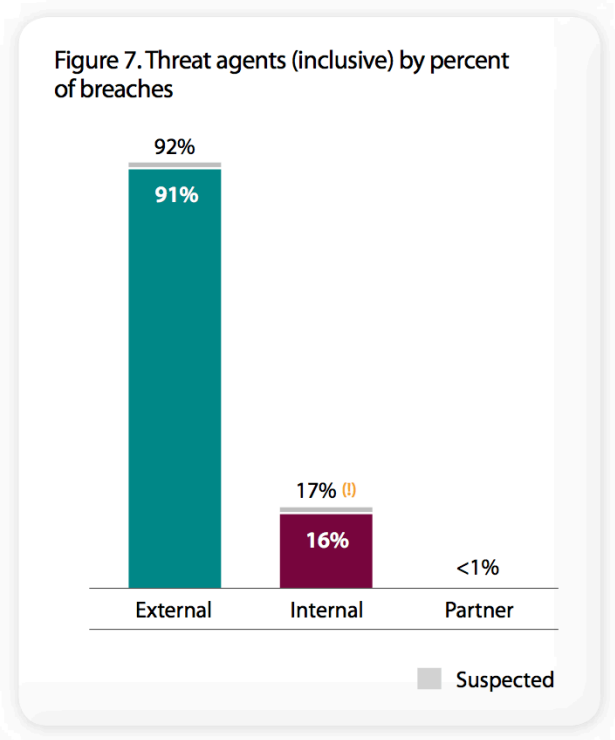

Diagram Source: 2011 Verizon Data Breach Investigations Report

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

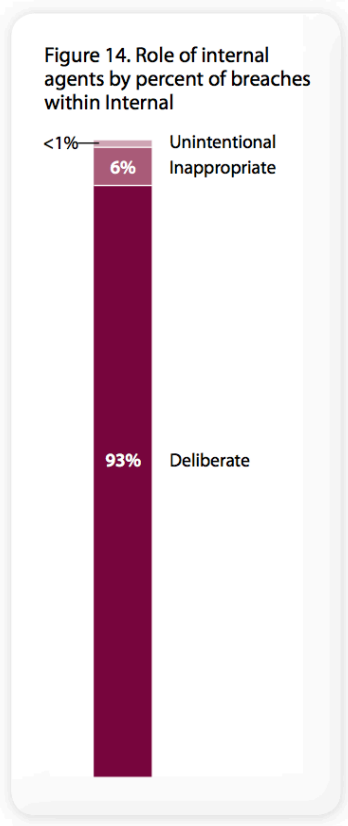

Diagram Source: 2011 Verizon Data Breach Investigations Report

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

Los Angeles Times | California | Local

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

Related Articles

- NORTH HOLLYWOOD - Hospital Plans Day-Care Center Apr 07, 1992
- CAMARILLO - Hospital Use Down, 8 Workers Laid Off May 13, 1992
- Ban on Smoking Imposed at UCLA Medical Center Dec 24, 1985

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 Permanete 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

BBC Low graphics | Help Search

NEWS Watch ONE-MINUTE WORLD NEWS

News Front Page
Africa
Americas
Asia-Pacific
Europe
Middle East
South Asia
UK
Business
Health
Medical notes
Science & Environment
Technology
Entertainment
Also in the news

Page last updated at 10:06 GMT, Friday, 17 July 2009 11:06 UK

E-mail this to a friend Printable version

Octuplets' hospital privacy fine

The hospital where octuplets were born in January has been fined for a second time for failing to protect the family's medical privacy.

Suleman has attracted criticism

Nadya Suleman attracted worldwide attention after giving birth to eight babies at Kaiser Permanente's Bellflower hospital in Los Angeles.

The hospital was fined $250,000 in May over staff looking at Suleman's records inappropriately.

The new $187,500 fine is for similar breaches of the babies' privacy.

QUT isi Information Security Institute

Insider Problem: Some Statistics (cont.)

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

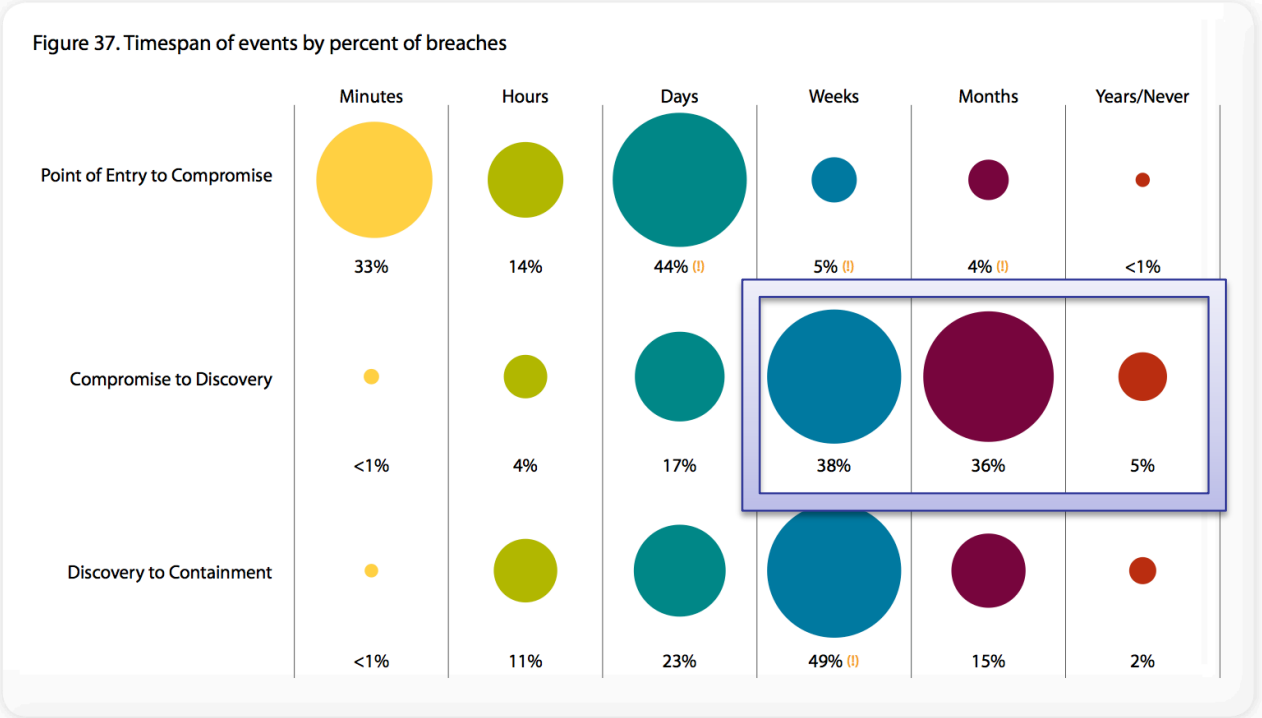

Diagram Source: 2011 Verizon Data Breach Investigations Report

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)

Related Work: 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.

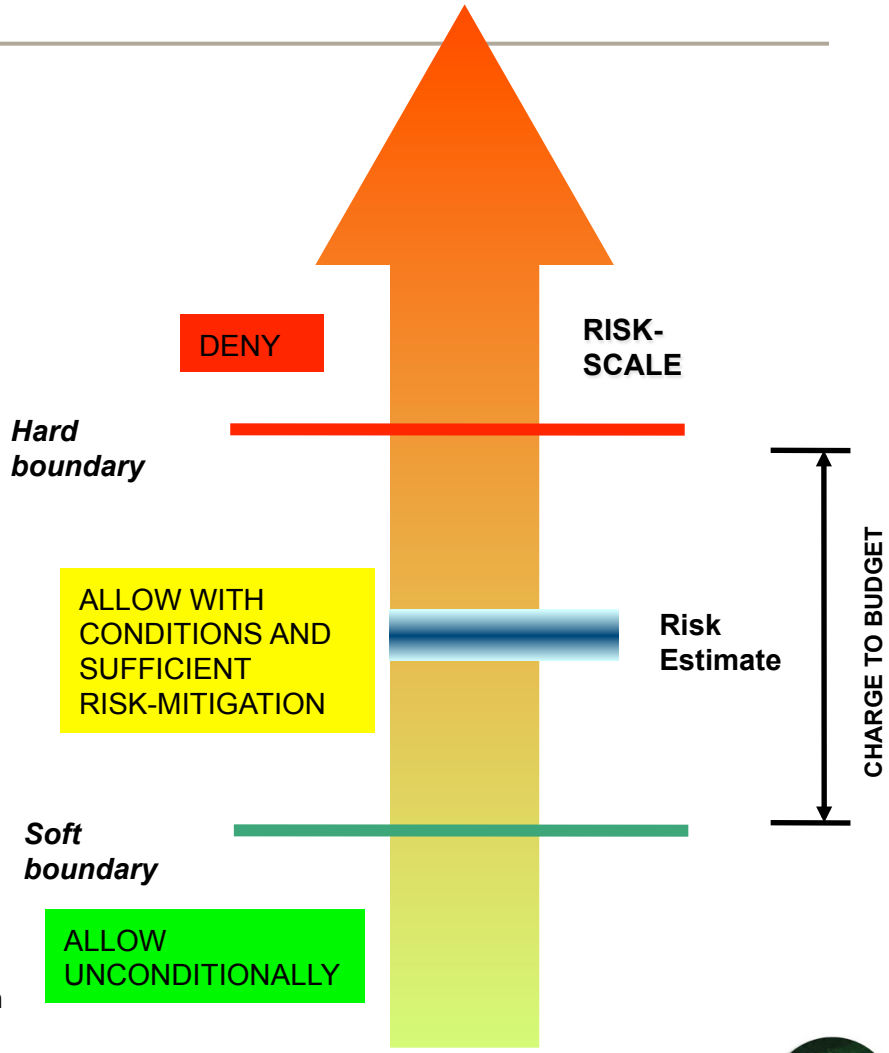

Image Credit: Cheng et.al

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.

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 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

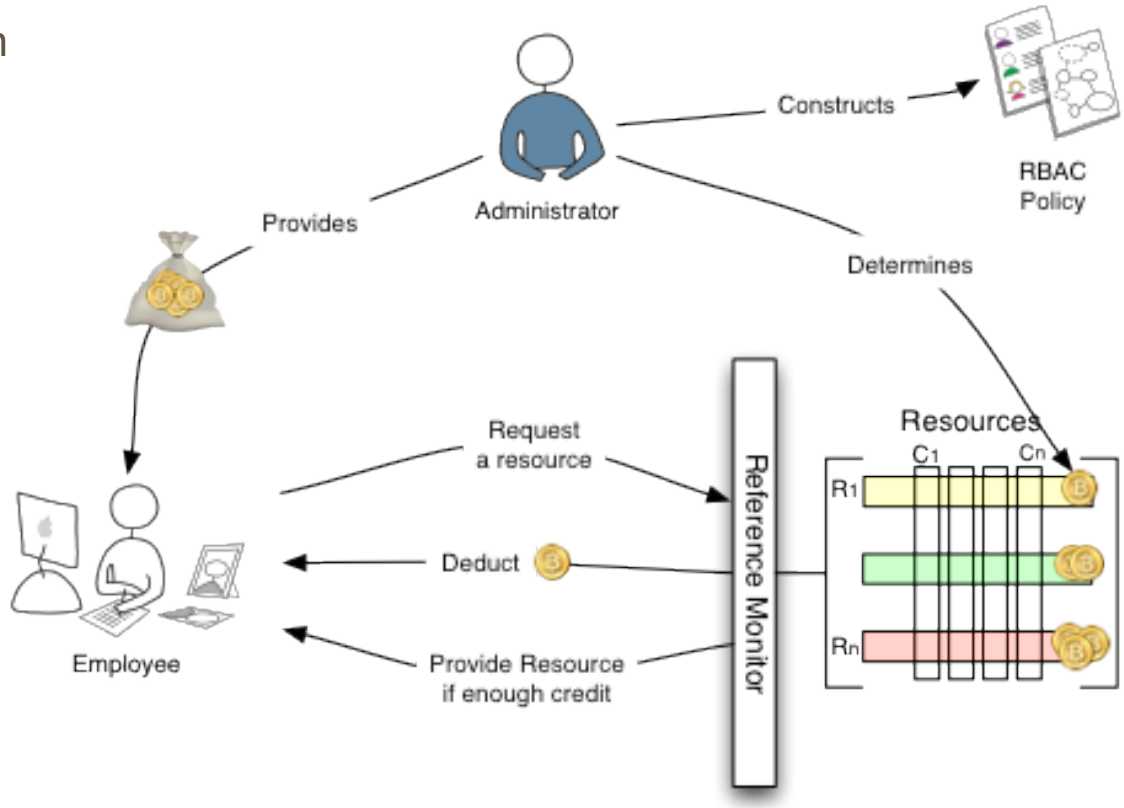

Our approach in a Nutshell

- 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?

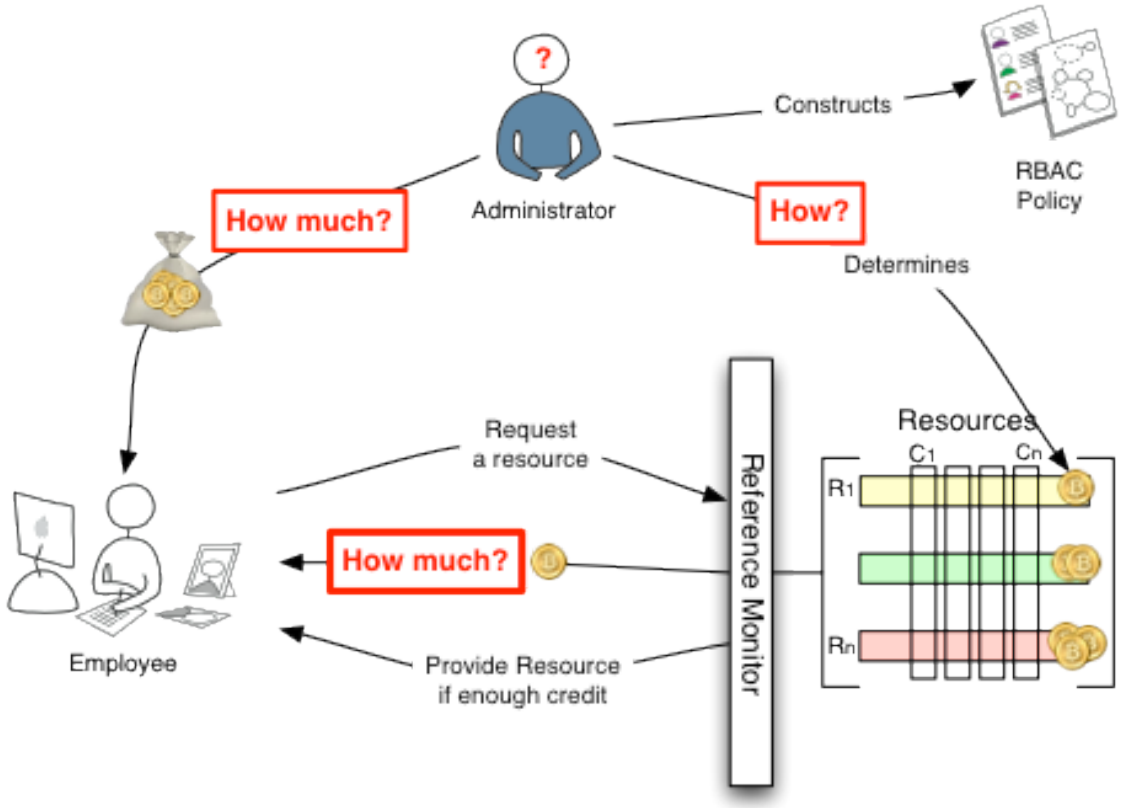

Basic Questions

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

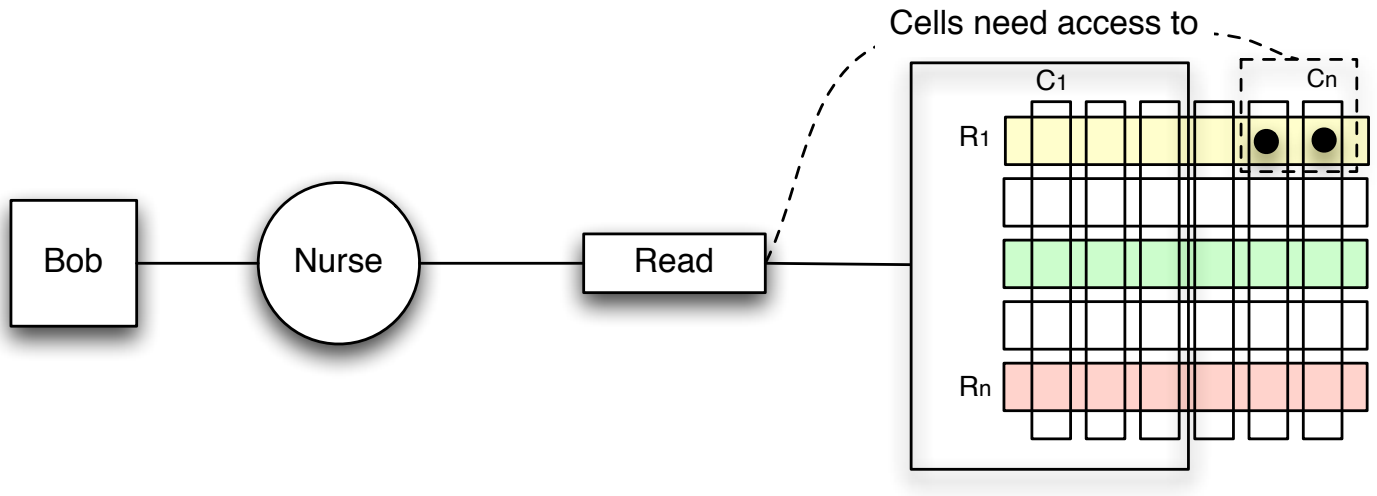

Basic Questions

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

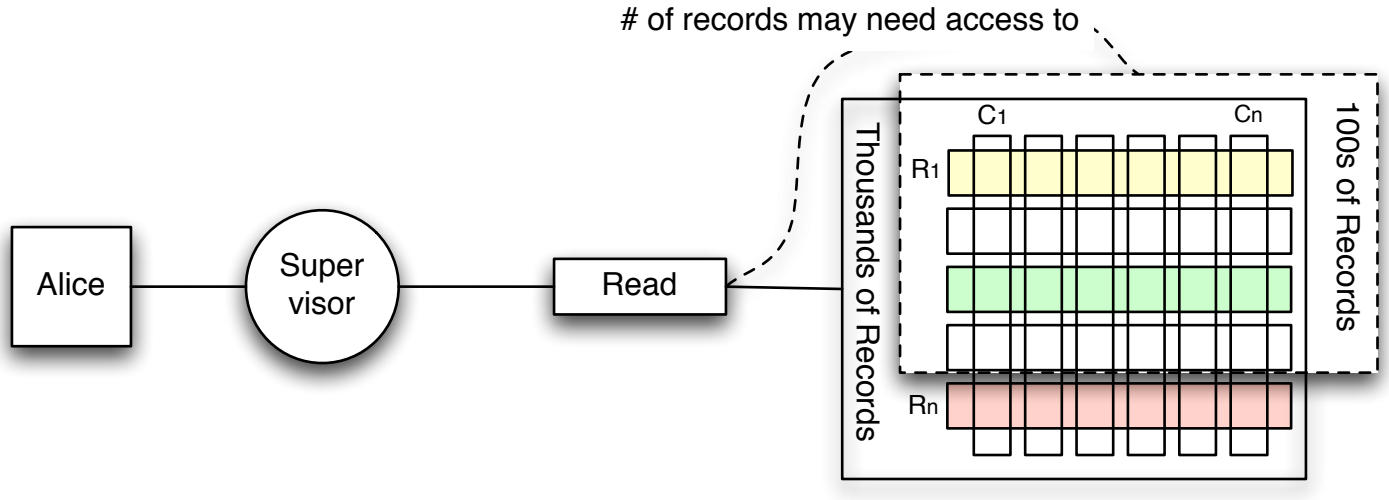

Contemplating an answer to the question

- 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:

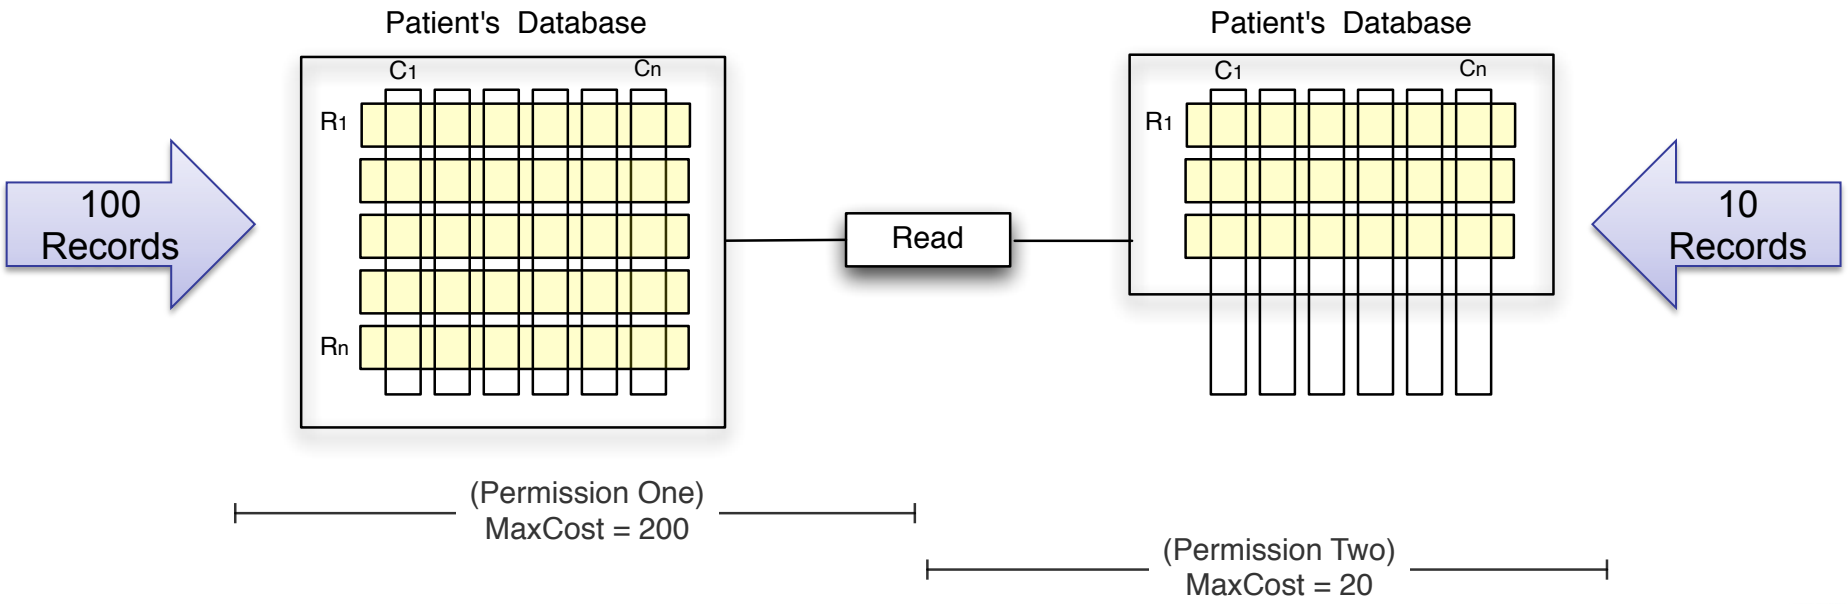

Implications: Permission Comparison

- 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)

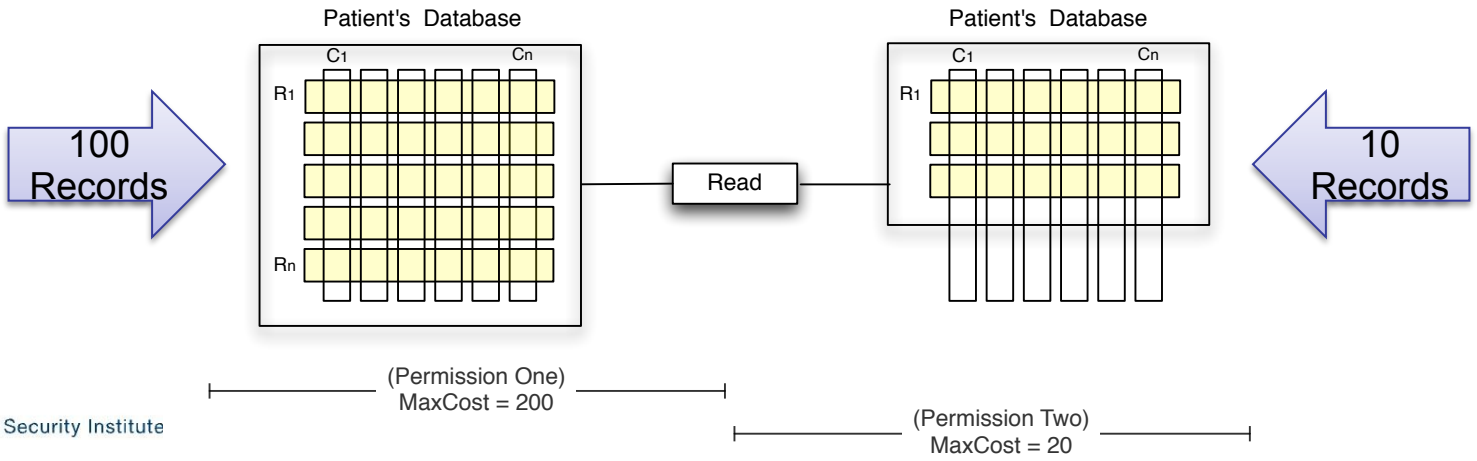

Missing In RBAC

- 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

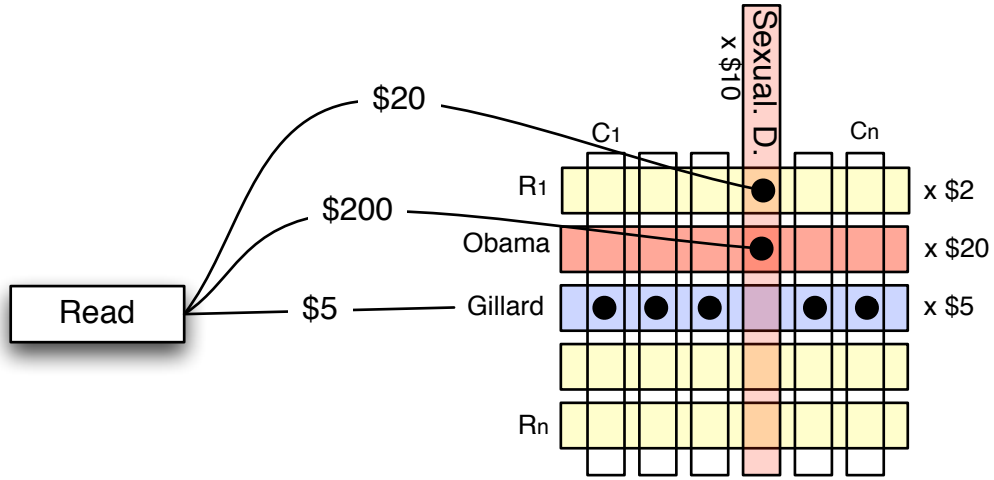

Deriving Role's Weight

- 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.

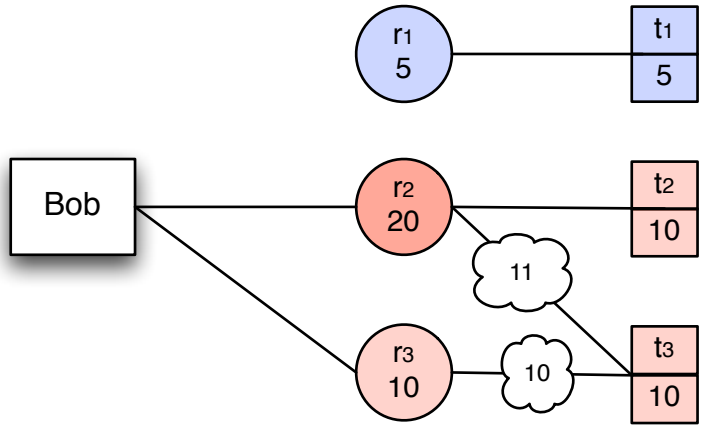

Administration: Side Effect of Role's Weight

- 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. t_4 costs \$10 through r_4 instead of \$6 if t_3 and t_5 are removed from the role

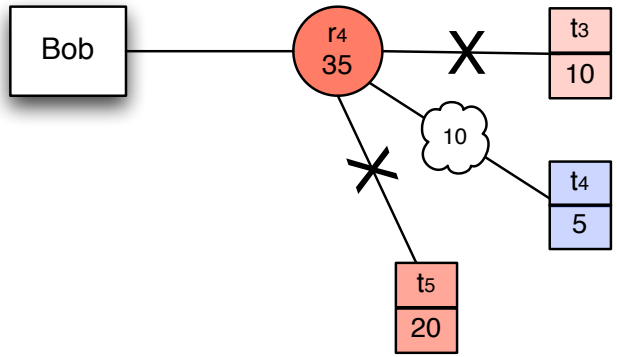

Incorporating Exceptions

- 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)?

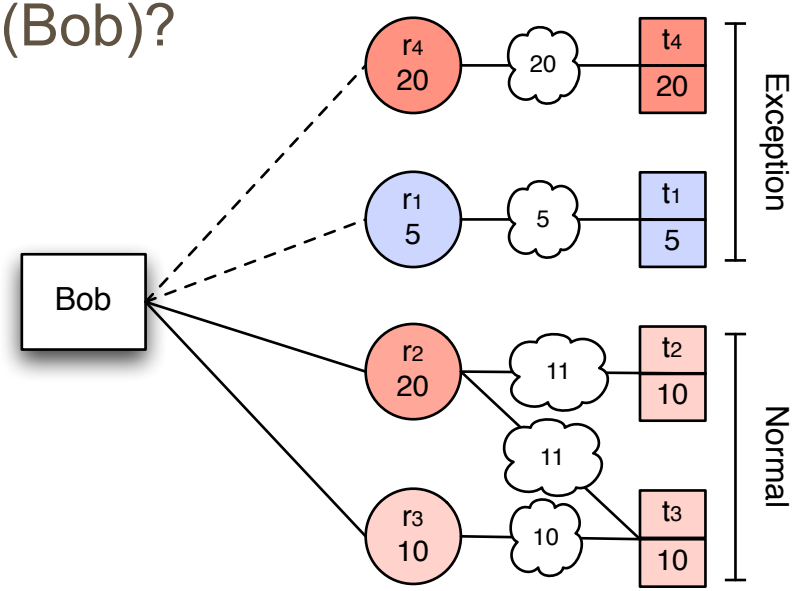

Incorporating Exceptions: Factoring in Probability of Misuse

- 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 $\beta = 1$

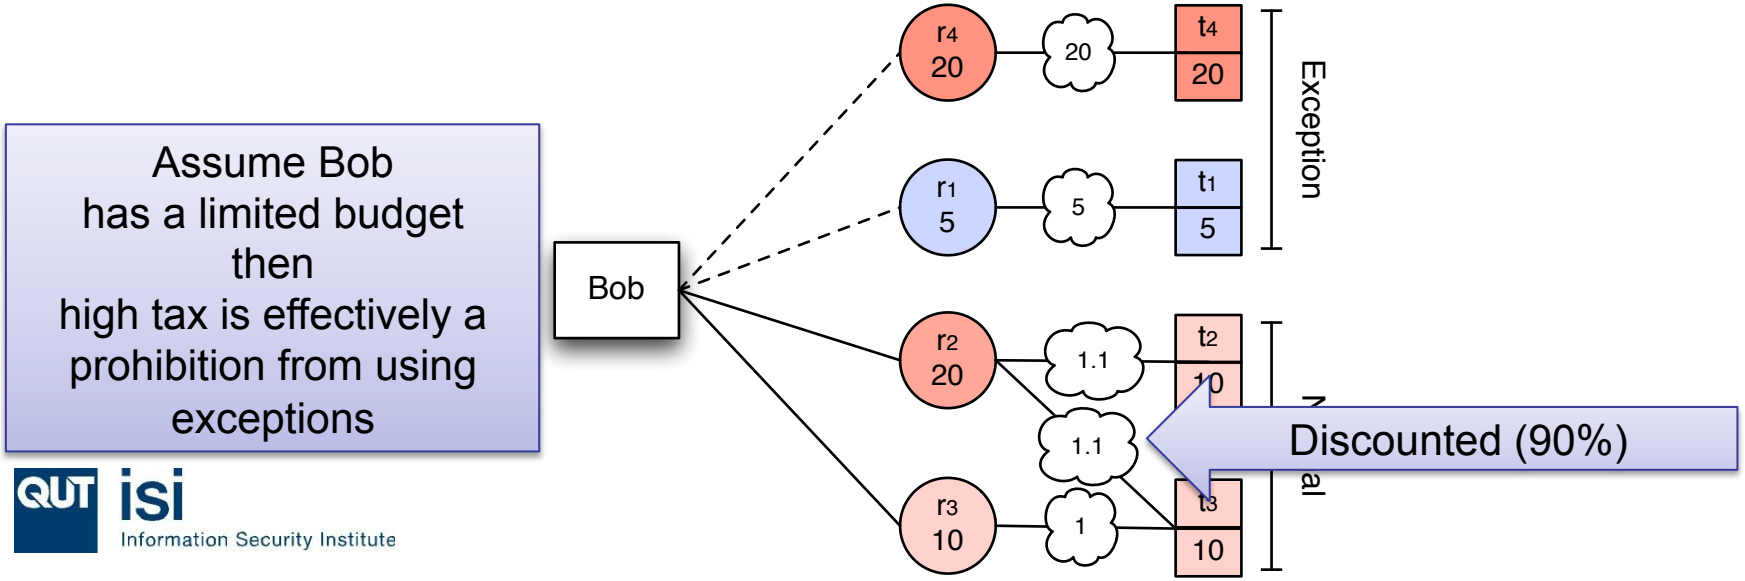

QUT isi Information Security Institute

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

① 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

② 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

② 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

③ 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

④ 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

④ 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.