



University of Pittsburgh

Insider Threats

Nathalie Baracaldo

Ph.D. Candidate

School of Information Sciences

March 26th, 2015



Insider Attacks



- According to CERT insider attackers are defined as:
 - Currently or previously employed individuals, contractors or business partners that:
 1. are or were legitimately authorized to use some privileges,
 2. decide to exceed or *intentionally* use their privileges to negatively impact an organization

Insider Attacks' Impact




- Insider attacks accounted for **33% of the total incidents** reported
- Monetary losses ranged from \$500 to **\$10 million**
- 75% of the organizations had a negative impact on their operations, and 28% on their reputations

[Computer Crime and Security Survey 2010]

Insider Attacks' Impact

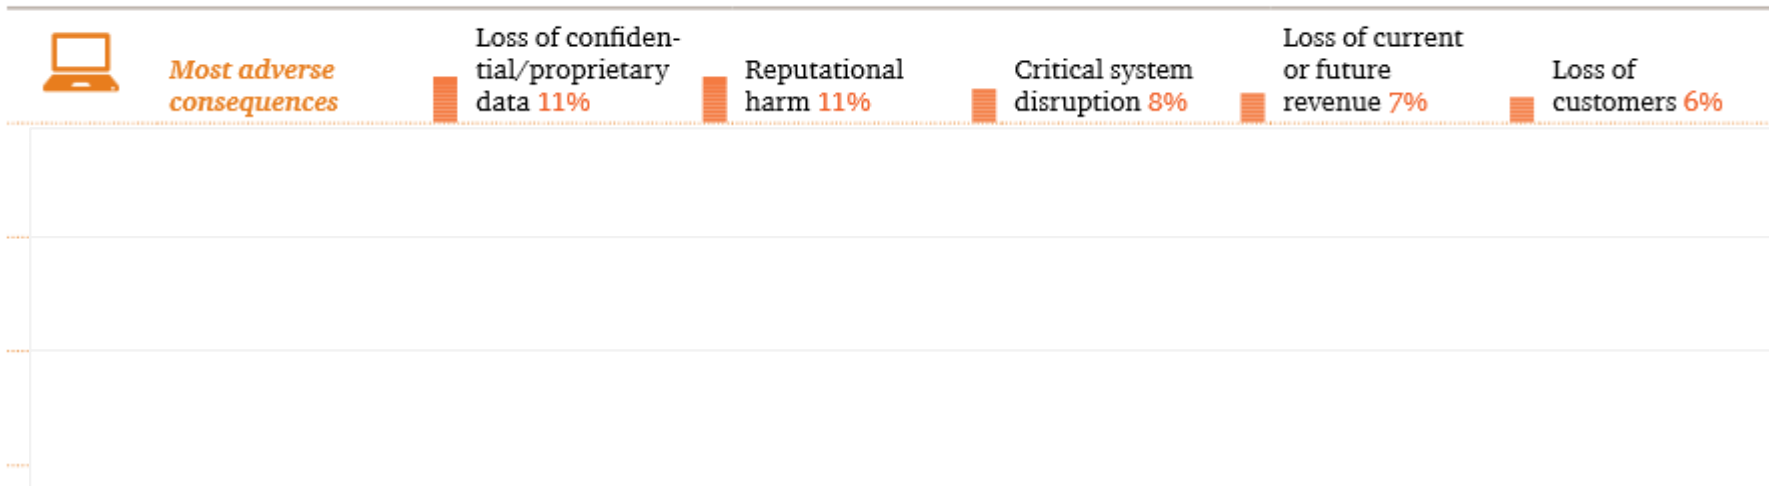
- Caused more monetary damage than attacks performed by outsiders



[Computer Crime and Security Survey 2011]

A closer look

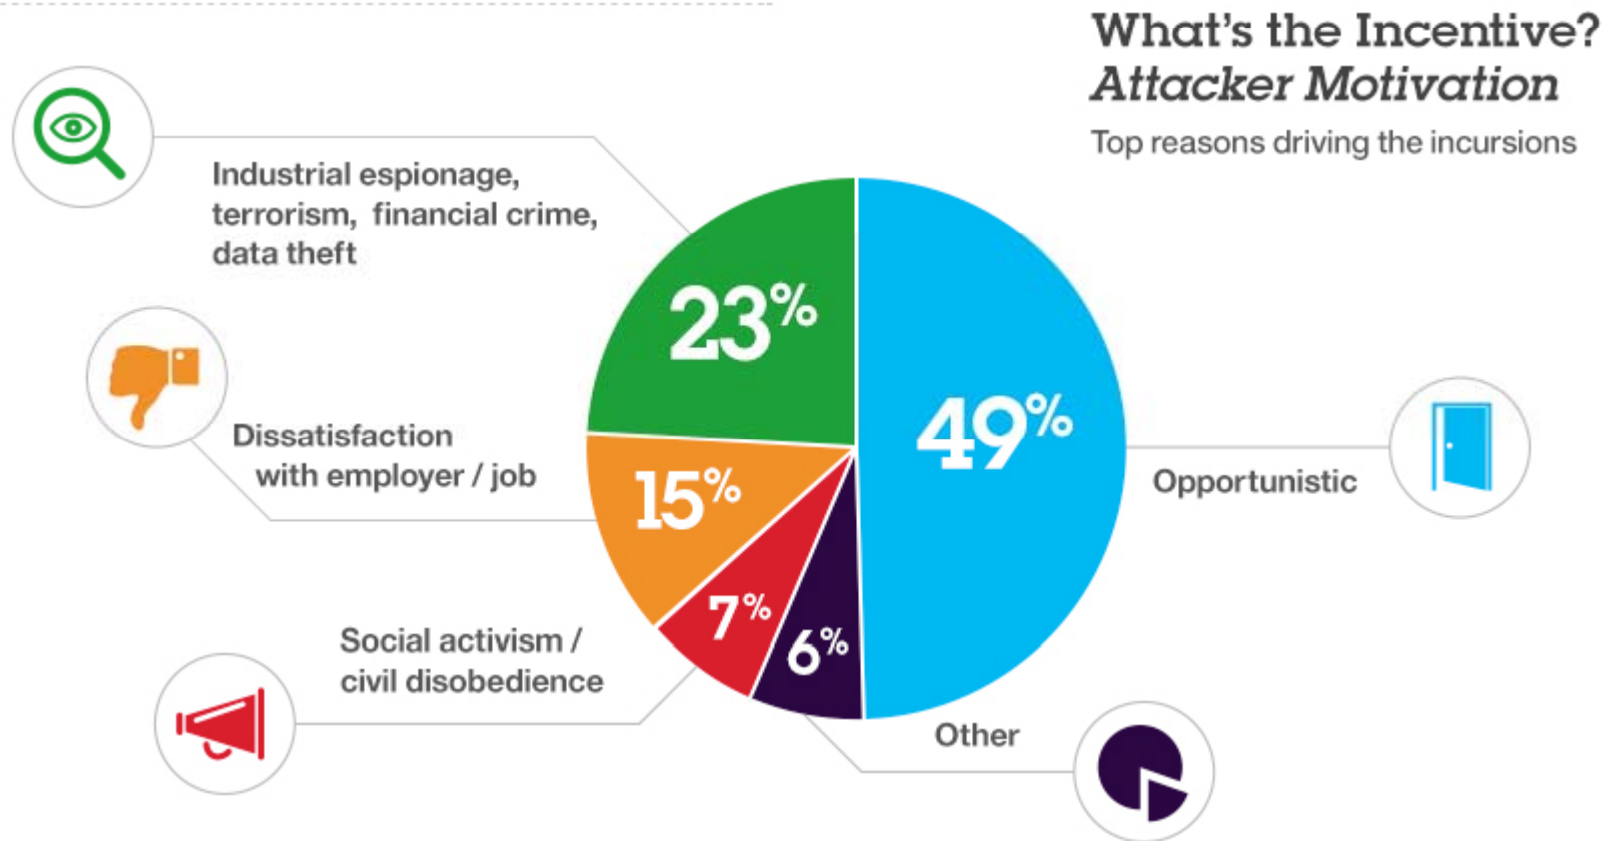
Figure 2: The causes and consequences of cybercrime committed by insiders*



Taken from "US cybercrime: Rising risks, reduced readiness
Key findings from the 2014 US State of Cybercrime Survey"
By PWC

Why do we care about this information?

More numbers...



IBM Cyber Security Index


<http://www-935.ibm.com/services/us/en/security/infographic/cybersecurityindex.html>

Any incidents come to mind?

- Edward Snowden
 - Leakage of confidential information
- Wiki-leaks
- Employee of an electric company
 - 911 service was out of reach for several minutes

Let's classify these attacks according to the attacker's motivation

- Edward Snowden
- Wiki-leaks
- Employee of an electric company
- Insider Trading
- Any examples of an opportunistic attack?



Classification of attacks

- IT sabotage
- Intellectual property theft
- Fraud
- Espionage
- Think about the cloud... Any new types?
 - Curious cloud administrators
 - Stalking
 - Blackmailing or embarrass others
 - Affect political events



According to the CERT

Insider Definition of a Malicious Insider

- “is a current or former employee, contractor, or other business partner who has or had authorized access to an organization’s network, system, or data and **intentionally exceeded** or **misused** that access in a manner that negatively affected the **confidentiality**, **integrity**, or **availability** of the organization’s information or information systems.”

– CERT, Insider Threat Technical Report

Are there unintentional insider
attackers?

A world far from ideal...

- 60 % of respondents reported monetary losses caused by non-malicious insiders *
- Not wise to trust users blindly!

Unintentional Insider: Fishing

- Goals:
 - Obtain user-name and passwords
 - Other confidential information
 - Install virus or spyware



Unintentional Insider: Social Engineering



Some examples of this type of attack




- Any other examples come to mind?

Other types of Unintentional Threats Caused by Insiders



The Human Factor: *How Breaches Occur*

Many elements can contribute to the vulnerability of your organization, however none is more prevalent than the human factor, **which accounts for approximately 80%.**



Taken from **IBM Cyber Security Index**

Unintentional Insider Threat Definition

- *An unintentional insider threat is (1) a current or former employee, contractor, or business partner (2) who has or had authorized access to an organization's network, system, or data and who, (3) through action or inaction without malicious intent (4) **causes harm or substantially increases the probability of future serious harm to the confidentiality, integrity, or availability of the organization's information or information systems.***

"Unintentional Insider Threats: A Foundational Study", CERT 2013

Can we prevent insider threats?


Can we predict these attacks?

- Insider attacks are typically preceded by **technical** and **psychological precursors**




Technical precursors

- Download and use of hacker tools
- Unauthorized access to other users' or customer data
- Setup or use of backdoors
- Transmitting large files
- Etc.



Psychological precursors

- Disgruntlement
- Bad attitude
- Lack of dependability
- Absenteeism
- Etc.



[Greitzer et. al]

Where does this data come from?



Now, let's switch sides!

- As an employee... Do you like to be scrutinized?



An important remark


- It is important to inform employees that they are being monitored!
 - Otherwise it may:
 - increase their disgruntlement levels and their likelihood of becoming an insider
 - reduce the trust employees have in the organization



Some Approaches to Deter Insider Threats


Try to reduce the risk exposure

- Have in place appropriate procedures
- Design adequate access control policies
- Try to predict possible attacks
- Adapt to negative changes in behavior
- We will see some examples!




Example Mitigation 1: Monitor emails, Internet traffic, file transfers, etc.

- Flag suspicious behavior and show an alert to the system administrator



Alerts



Example Mitigation 1 (cont.)

- Any problems?
 - May not be effective against **stealth attackers**
 - This is a problem for all anomaly detection systems
 - May not work if multiple adversaries are **colluding**
 - Much more difficult to flag suspicious behavior if malicious activity is carried out by multiple users
 - **Not automatic**, administrator may miss important alert!



Example Mitigation 2: Use Decoys

- Use file search behavior to identify user's normal search patterns
- Monitor how user is searching his files
- If the current behavior does not match the expected one, **decoy** files are served to the user
- Is this a good solution?



Solution proposed by Salem et. al [Combining a baiting and a user search profiling techniques for masquerade detection](#)

Example Mitigation 2 (cont.)

- Hypothesis of the solution:
 - If an opportunistic malicious colleague is accessing another's user computer, the search behavior would be different
 - In theory, the real user would distinguish fake files vs. real files
- Do you see any problems with this hypothesis?

Example Mitigation 3: Non-technical mitigation strategies

- Educate users
 - Avoid unattended terminals
 - Prevent fishing
 - Prevent social engineering attacks
 - Increase awareness of possible relevant problems e.g., SQL injections in a SW engineering company
- Create a good working environment 😊
 - Disgruntle employees are more likely to become insider attackers
 - Recall that 15% of attacks are committed by unhappy employees



Example Mitigation 4: Implement an Access Control System

- This is a MUST!
- Restrict the access enforcing
 - Separation of duty
 - Least privilege enforcement
- Challenge: Employees need the privileges, but we need to prevent the abuse those permissions



Current Access Control Approaches

- Access control systems are highly static
 - As long as users have the required credentials, they can access the system
 - What about their behavior?
- Require manual verification and input
 - Manual verification of alerts
 - Input of psychological precursors is slow and subjective



Current approaches (*cont.*)

- Do not minimize risk exposure *continuously, automatically* and *optimally*
 - Risk methodologies are performed sporadically (e.g., NIST, Octave, etc.)



Our Proposed Research



- Two concepts:
 - **Trust:** expectation of future behavior based on the history
 - **Risk:** likelihood of a hazardous situation and its consequences if it occurs
- We include **risk** and **trust** in access control systems to adapt to anomalous and suspicious changes in users' behavior

We identify an opportunity to control risk very frequently (for each access request) and

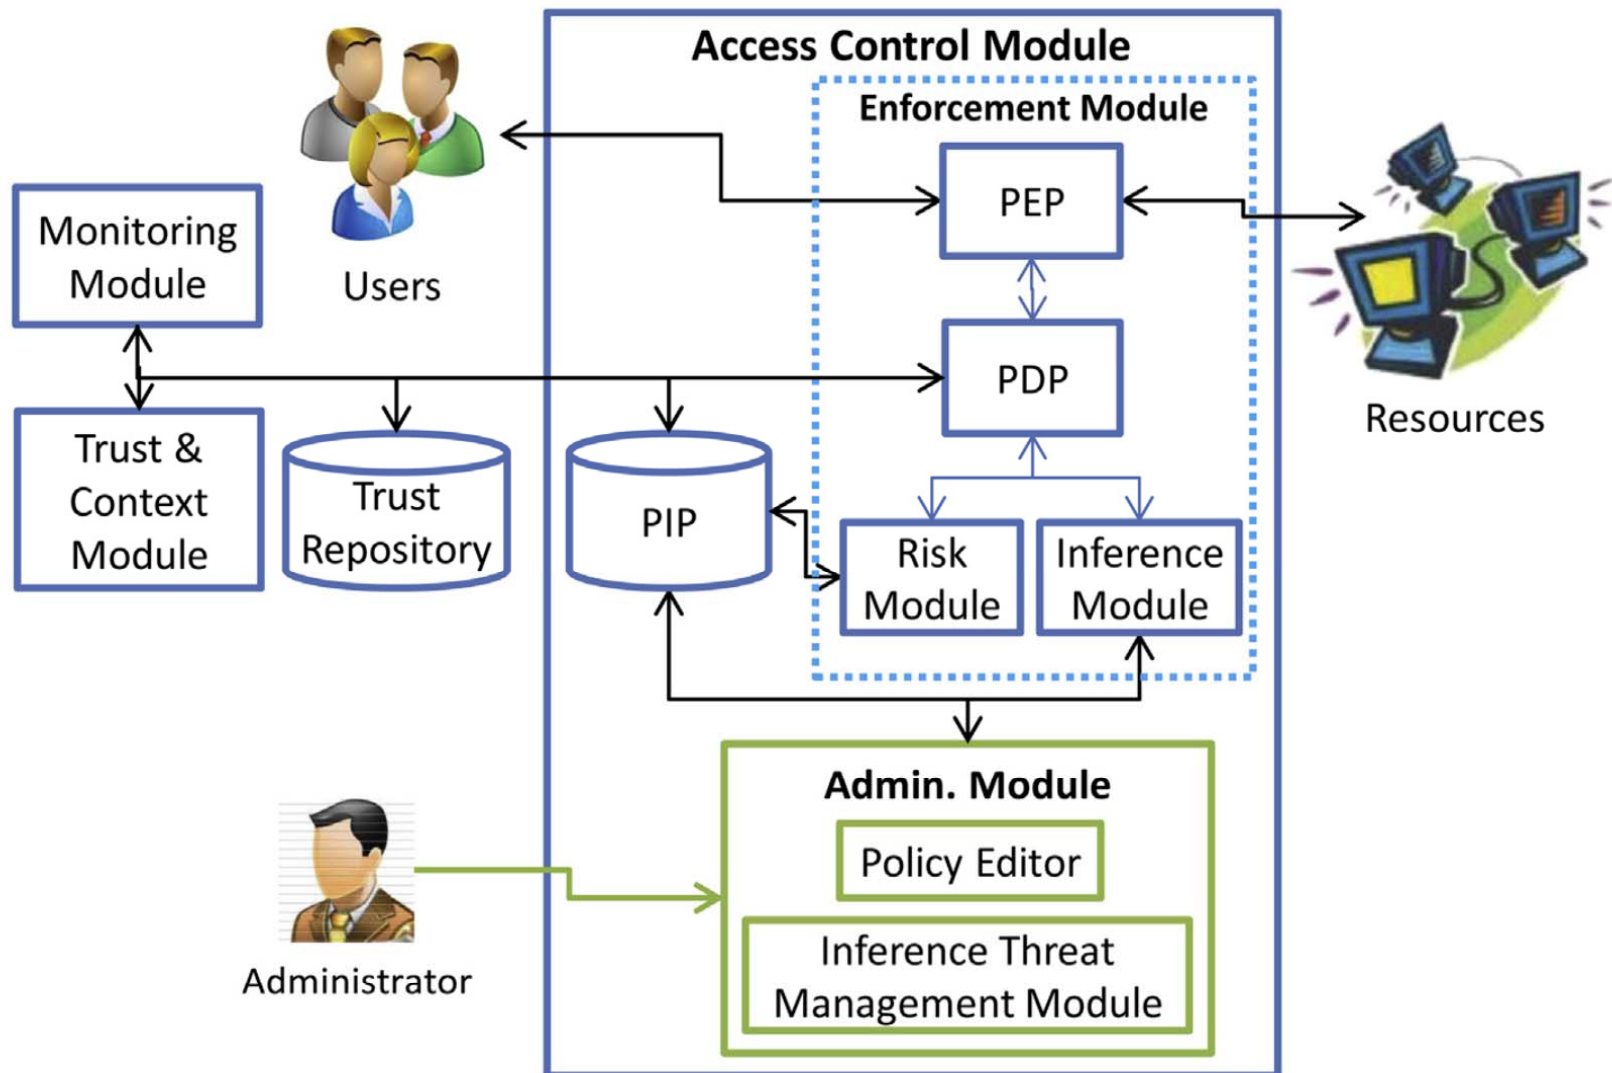
Requirements of Our Adaptive System

1. Enforce separation of duties (SoD) and cardinality constraints
2. Detect suspicious activities, and establish a trust level for each user
 - Different trust values for users depending on the *context*

Requirements (*cont.*)

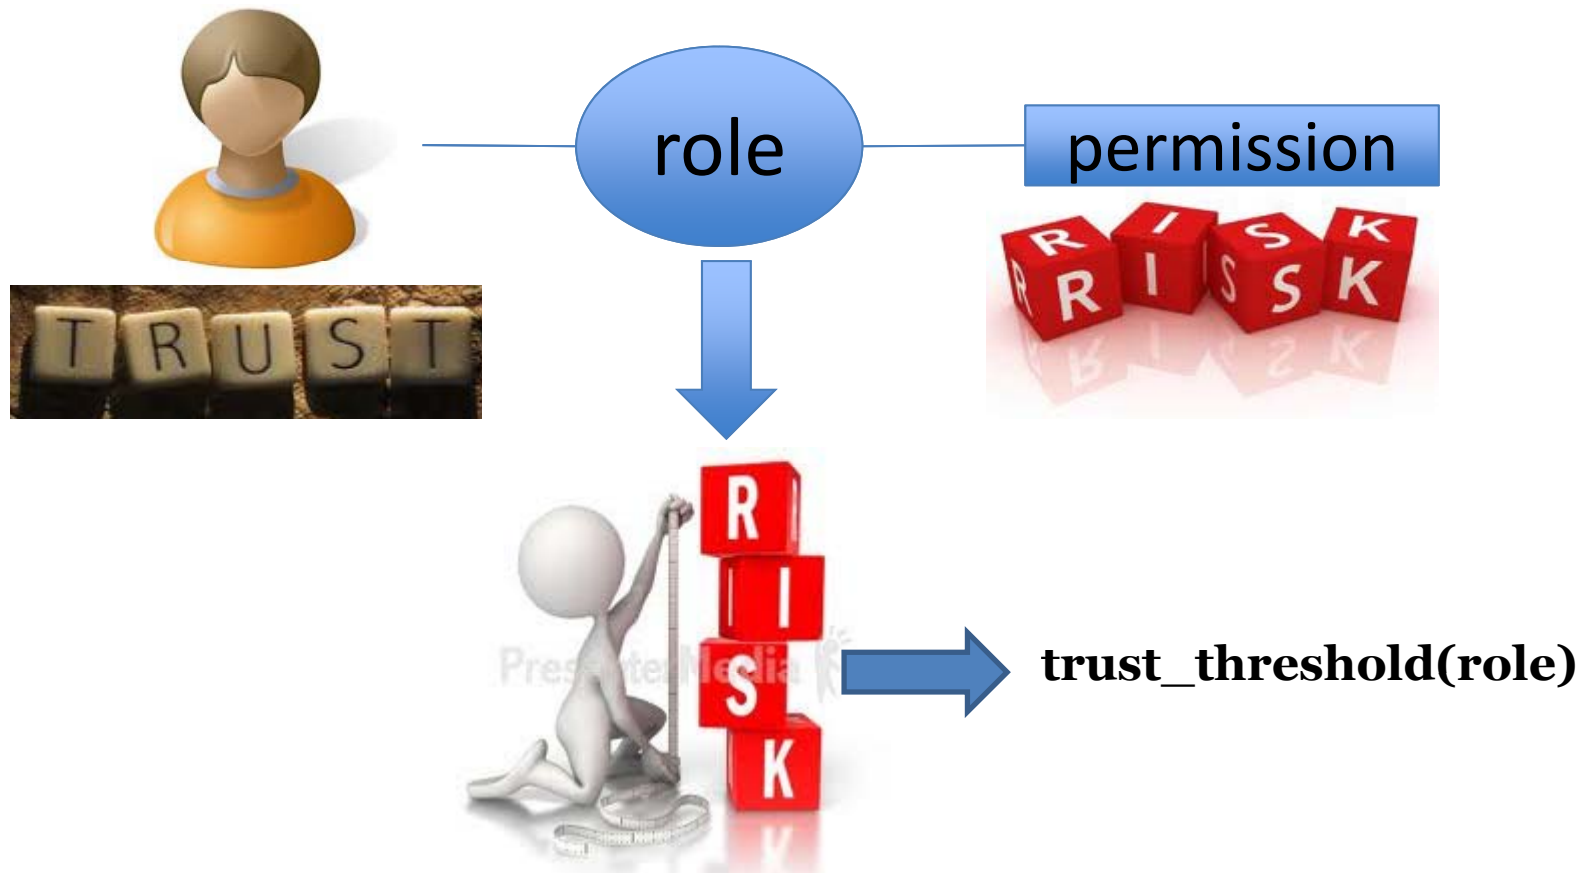
3. Different permissions may have different risks associated with them
 - Adapt to suspicious changes in behavior of users by restricting permissions depending on risk values
4. Risk exposure should be **automatically** reduced, minimizing the impact of possible attacks

Framework Overview



"An adaptive risk management and access control framework to mitigate insider threats", Nathalie Baracaldo and James Joshi, 2013 Paper available: http://www.pitt.edu/~nab62/baracaldo_cose_2013.pdf

In a nutshell...



$\text{authorized}(u, \text{role}) \ \& \ \text{trust}(u, c) \geq \text{trust_threshold}(\text{role})$

Trust value of users

- Each user u is assigned a trust value:
 - $0 \leq \text{trust}(u, c) \leq 1 \rightarrow$ reflects his **behavior**
 - Where c is the context, and u is the user
- Some works exist to calculate this value



Assigning risk to permissions

- Each permission is assigned a risk value according to:
 - The *context*
 - The likelihood of misuse
 - The cost of misuse

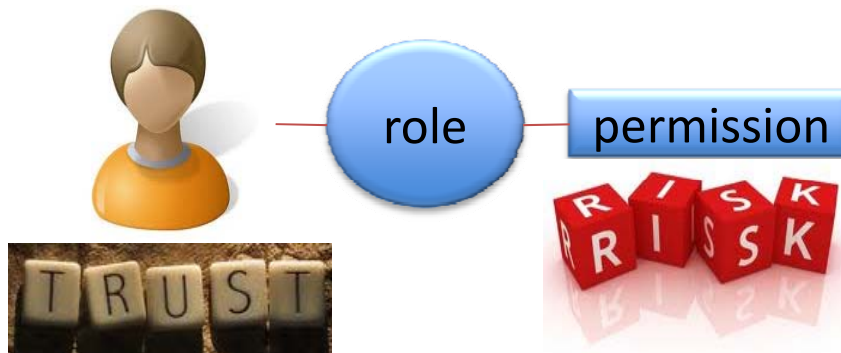


DEFINITION 1. *The risk of permission $p = \langle obj, act \rangle \in P$ in context $c \in C$, written as $rs(p, c)$, is defined as follows:*

$$rs(p, c) = \sum_{x_p \in MaliciousUsage} Pr[x_p | c] * C(x_p)$$

Risk of roles

- The risk of activating a set of roles depends on:
 - Context
 - The user that is going to activate the roles
 - Authorized permissions & their risk
 - Inference risk




Inference risk

- *Inference Threat*: exists when a user is able to infer unauthorized sensitive information through what seems to be innocuous data he is authorized for
- *Inference tuple*:

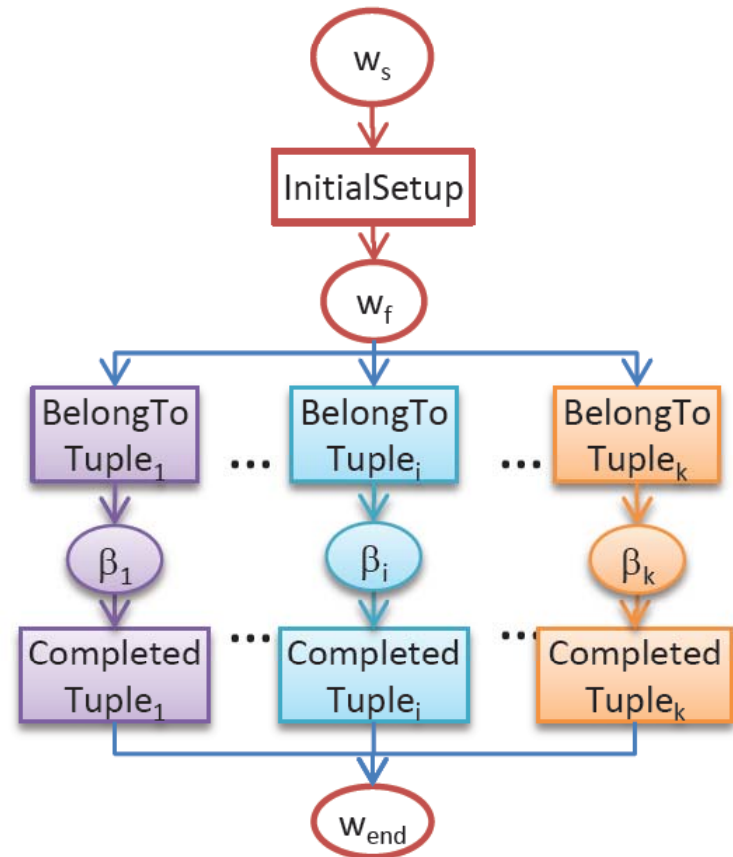
$\langle PS, p_x \rangle$

Shows the minimum information needed (PS) to infer p_x



Inference Risk

- Determine if the user can infer unauthorized information – *colored petri-net*
- Depends on user's access history



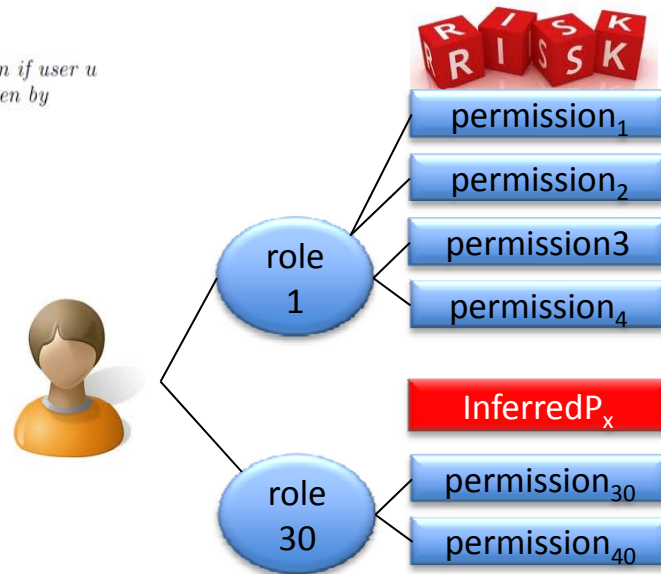
Risk of roles

- Risk exposure of activating a set of roles

DEFINITION 2. The risk exposure of the system if user u activates a set of roles $R' \subseteq R$ in context c is given by

$$rs(R', c, u) = \sum_{p \in \wp} rs(p, c)$$

where $\wp = P_{au}(R') \cup \text{inferred}(\mathcal{H}, u, R')$.



- For a set of roles RS , the **trust threshold** is the normalized version of their risk
- $0 \leq \text{trust_threshold}(RS, c, u) \leq 1$

Automatically reduce the risk exposure

- **Select roles with minimum risk** that also respect the policy constraints & provide the requested permissions

DEFINITION 3. *The Trust-and-Risk Aware Role Activation Optimization Problem for a query $q = \langle u, PS, c \rangle$, consists of finding a solution, R_q , such that:*

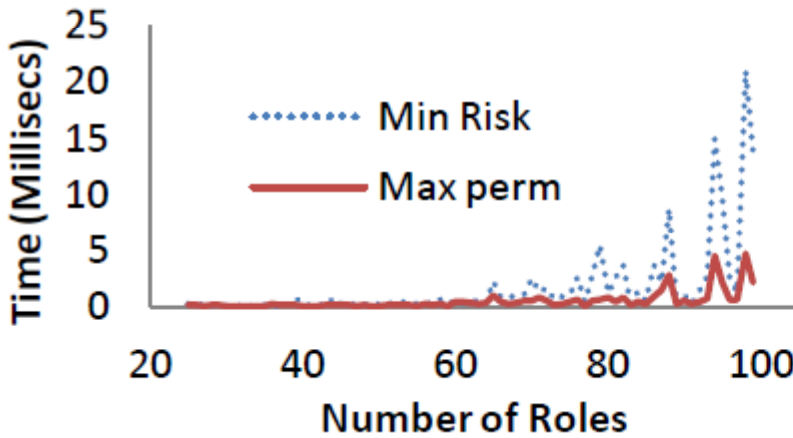
$$\begin{aligned} & \min_{R_q \subseteq \text{authorized}(u)} rs(R_q, c, u) \\ & \text{s.t. } \forall dsod(RS_i, k_i) \in DSoD : |R_q \cap RS_i| < k_i \\ & \forall card(r_c, k) \in CARD \wedge r_c \in R_q : \text{activated}(r_c) + 1 \leq k - 1 \\ & \text{trust}(u, c) \geq \tau(R_q, c, u) \\ & P_{au}(R_q) \supseteq PS \end{aligned}$$



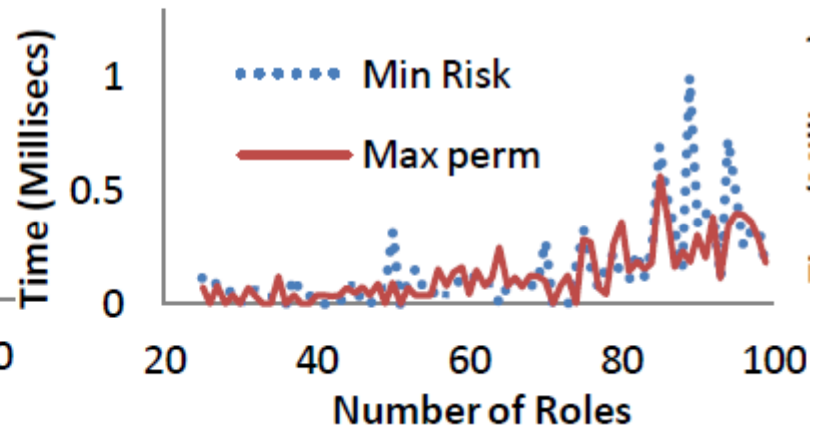
Experimental Setup

- We generated synthetic *well-formed policies*
- Each point represents the average time of running the algorithm for 30 different policies
- We evaluated our algorithm under two different heuristics for several types of policies

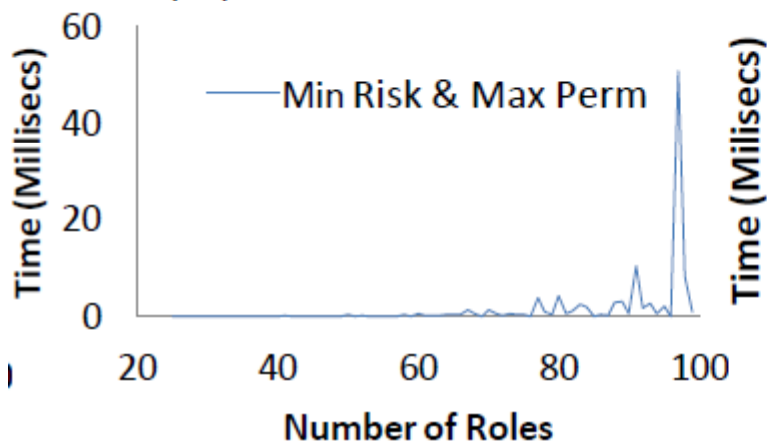
Comparison two selection heuristics
Min risk & *Max perm* for different policy configurations



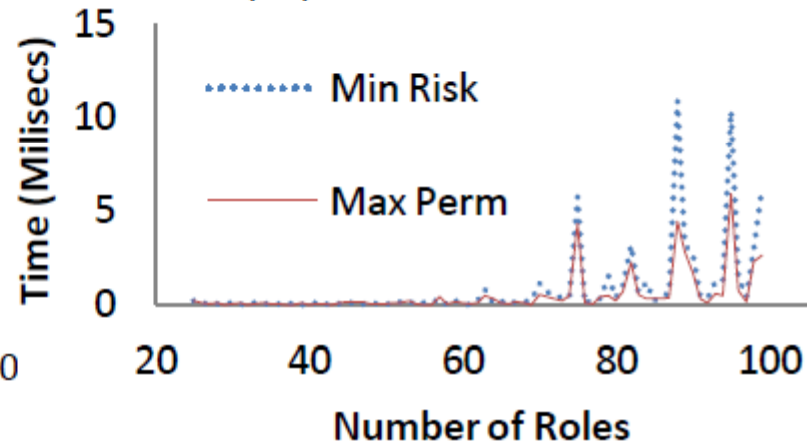
(a) $I:A:IA = 0:0:1$



(b) $I:A:IA = 1:0:0$

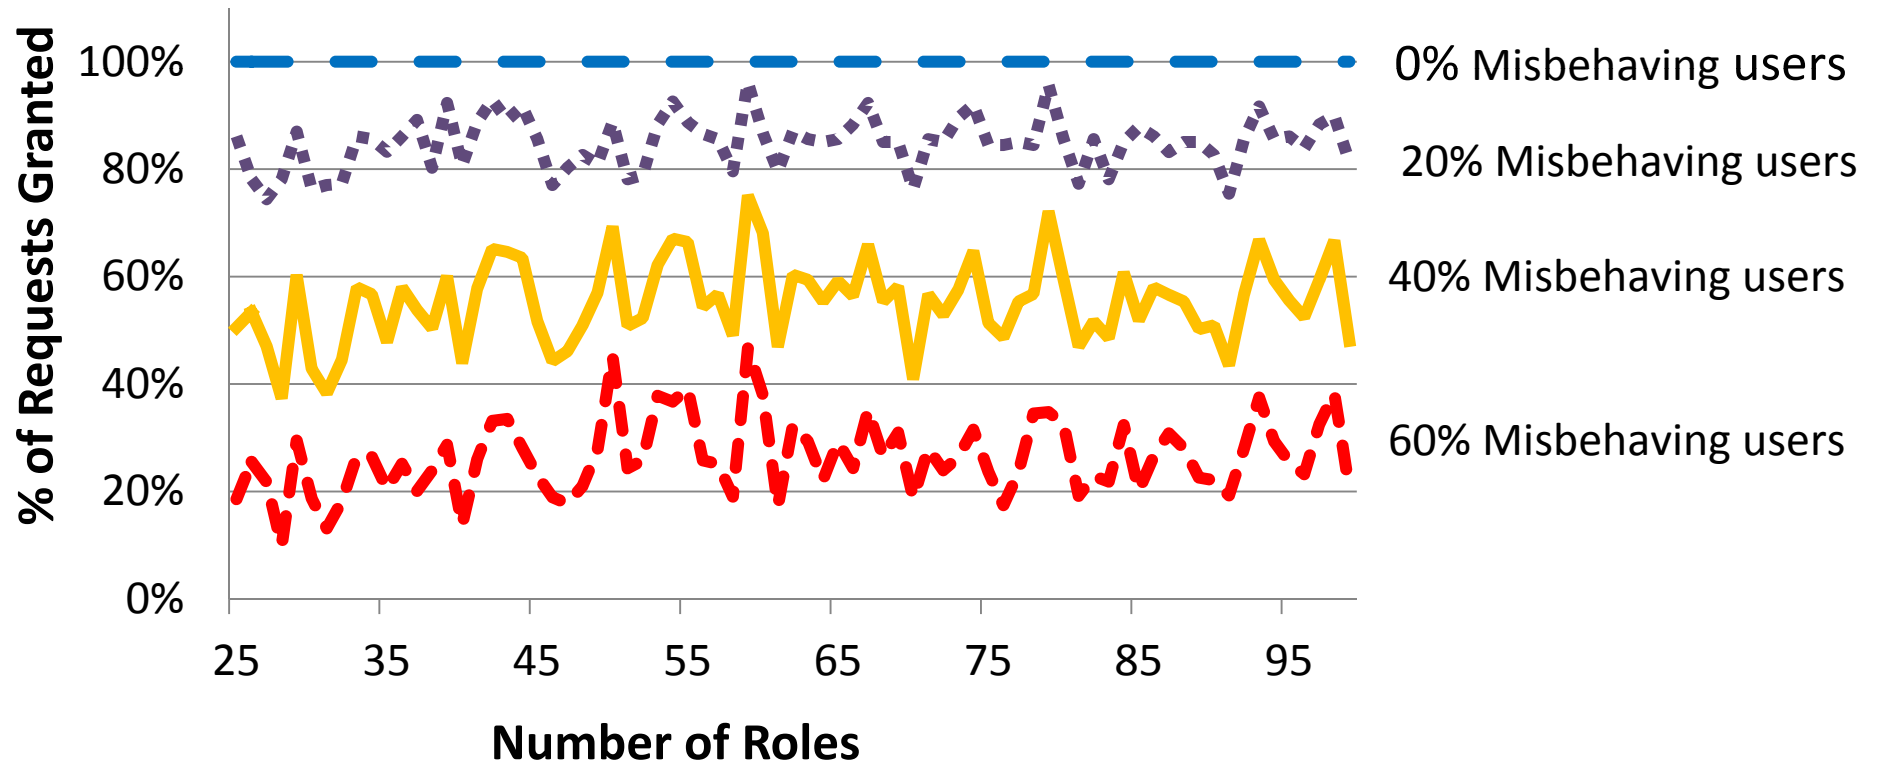


(c) $I:A:IA = 0:1:0$

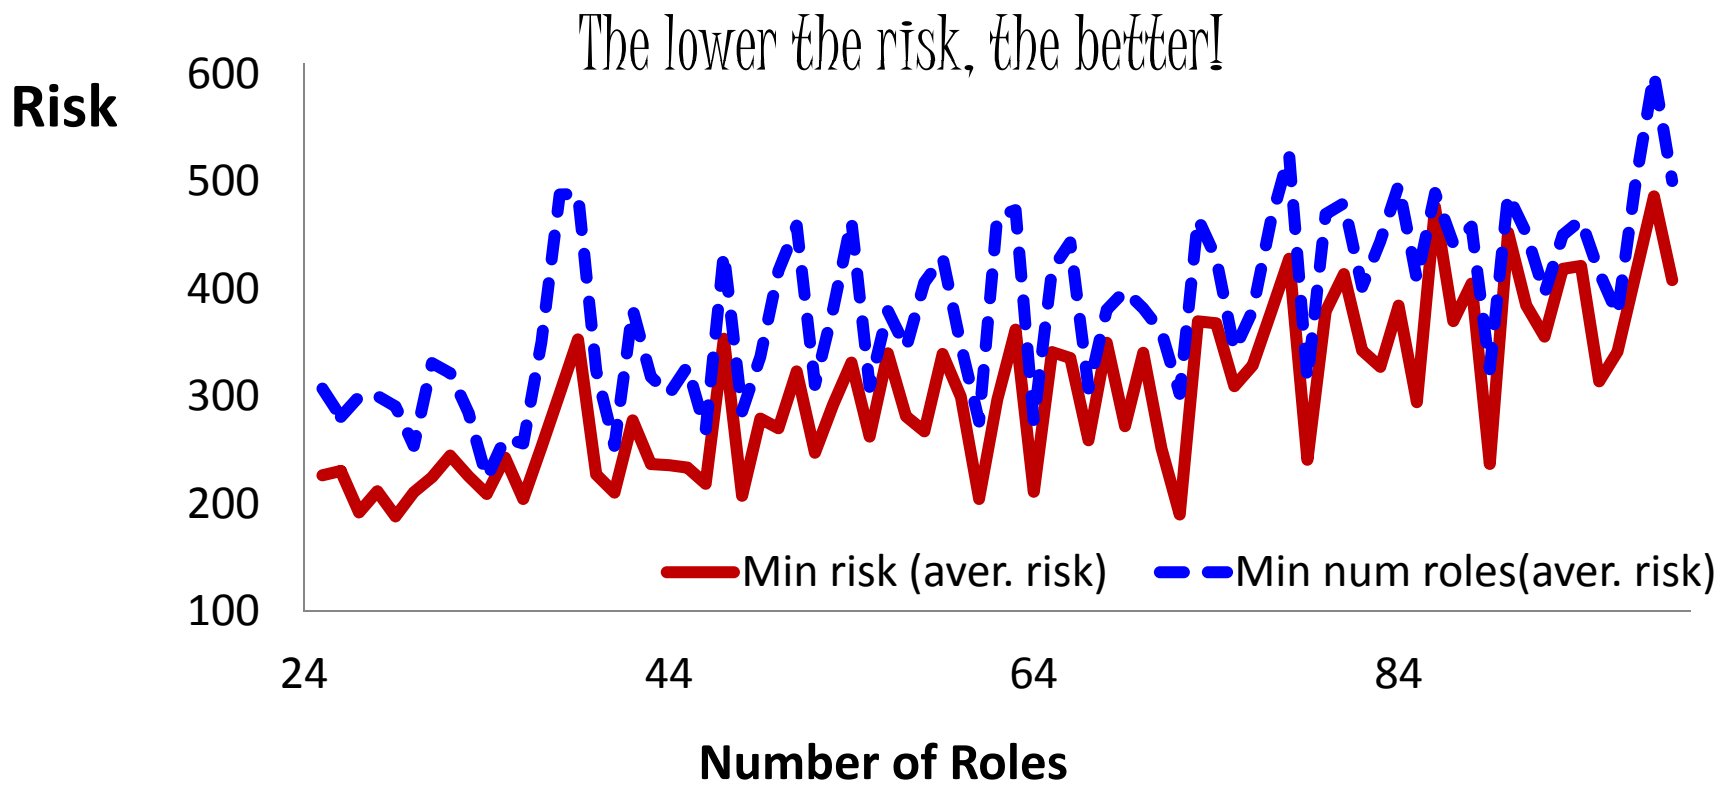


(d) $I:A:IA = 1:1:1$

Granted requests for different percentage of misbehaving users



Comparison of the risk exposure between our system (min. risk) and traditional role activation



How should we mitigate insider threats?

- Mitigation strategies depend on the type of organization
- A **risk assessment analysis** should be performed to define the policies, mechanisms and overall investment
- Remember that multiple technical and **non-technical components** need to be aligned to create a comprehensive solution
- It is also important to have recovery strategies!

Conclusions

So what should we do?

- Be prepared! It is necessary to have a plan to manage insider attacks
 - Decide what mitigation mechanisms are appropriate
 - Have a plan to react in case an insider attack occurs
 - Create the plan before any incident occurs!
- Guidelines: “*Common Sense Guide to Mitigating Insider Threats*”, 4th Edition CERT



Conclusions

- We overviewed **inside threats** and their impacts
- We also explored **unintentional insider threats** and their impact
- We overviewed some solutions to deter insider threats
- This is a challenging threat!

Conclusions (cont.)

- Want to know more?
 - Insider threats
 - The CERT Guide to Insider Threats
<http://resources.sei.cmu.edu/library/asset-view.cfm?assetID=30310>
 - Common Sense Guide to Mitigating Insider Threats, 4th Edition
<http://resources.sei.cmu.edu/library/asset-view.cfm?assetID=34017>
 - General publication list:<http://www.cert.org/insider-threat/publications/index.cfm>
 - Unintentional insider threat:
<http://www.sei.cmu.edu/reports/13tn022.pdf>
 - Use of decoys to deter insider threats:
 - [Baiting inside attackers using **decoy** documents](#)
 - [Combining a baiting and a user **search** profiling techniques for masquerade detection](#)
 - Adaptive access control systems to deter insider threats:
<http://www.pitt.edu/~nab62/research.html> (check papers related to insider threat)