

# Access control

Maarten Decat - SecAppDev 2017

[maarten@elimity.com](mailto:maarten@elimity.com)

# What is access control?

Access control is the part of *security* that constrains the *actions* that are performed in a system based on *access control rules*.


- As any security: confidentiality, integrity, availability
- Layer in between (malicious) users and the protected system
- Part of the Trusted Computing Base

# What

1. Not easy to **get right**,  
e.g., what about windows?
2. Difference between access  
**rules** and **mechanism**
3. Different mechanisms have  
different **properties**
4. Different mechanisms support  
different **rules**



# Access control in the physical world



# Access control in software


The image displays several overlapping screenshots illustrating access control in software:

- Facebook Login Page:** A browser window showing the Facebook login interface with fields for "Email or Phone" and "Password", and buttons for "Sign Up" and "Log In".
- Content Unavailable Error:** A message stating "Sorry, this content isn't available right now" with a warning icon, explaining that the link may have expired or the page may not be in.
- Terminal Window 1:** A terminal window titled "maartend@tyrion:~" showing the command `cat /etc/shadow` being executed, resulting in the error: `cat: /etc/shadow: Permission denied`.
- Terminal Window 2:** A terminal window titled "maartend@tyrion:~" showing the command `apachectl start` being executed, resulting in the error: `Failed to start httpd.service: Access denied`.
- 403 Error Page:** A webpage showing a "403 Access to the webpage was denied" error, with the text "You are not authorized to access this web page".
- Banking Interface:** A partial view of a banking interface showing account information for "DECAT MAARTEN" with various account numbers and EUR currency indicators.

# Outline

- Introduction
- Positioning access control
- Access control models
- How to enforce access control
- Some important technologies in practice
- Recap and conclusion

# 10,000m point of view



# But there is more to it

Login Into Your Account

Sign In

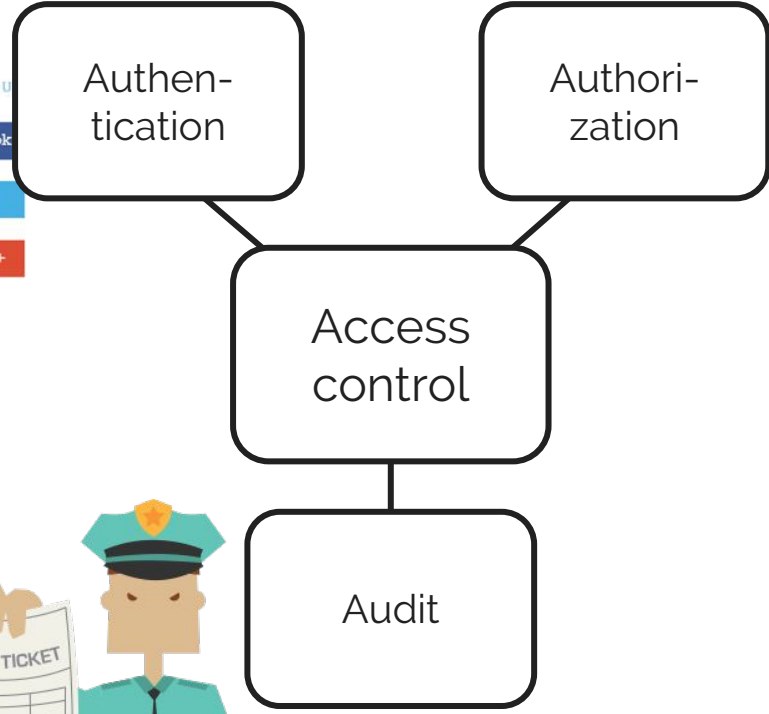
Remember me [forgot password?](#)

Don't have an account? Sign Up

 Sign In with Facebook

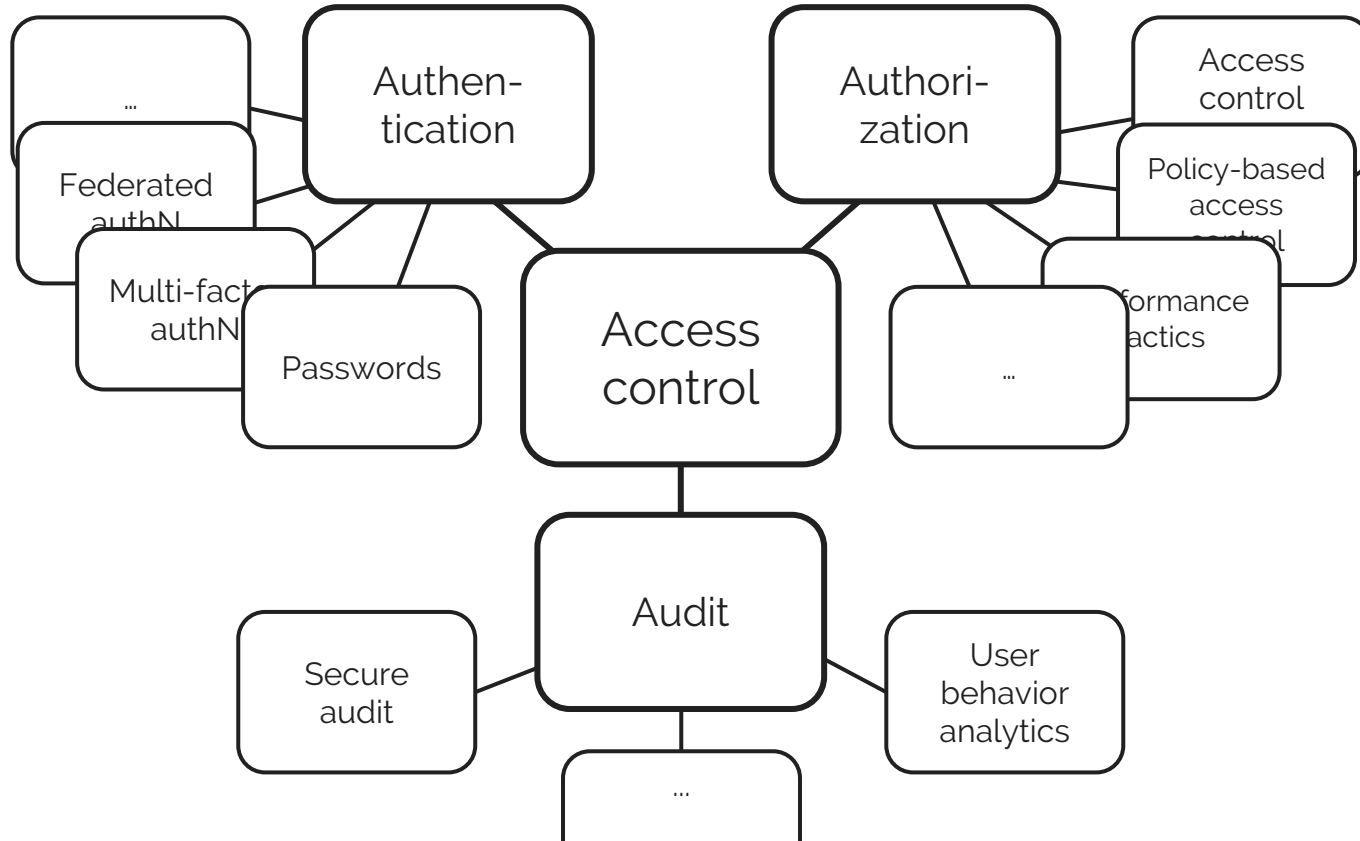
 Sign In with Twitter

 Sign In with Google+



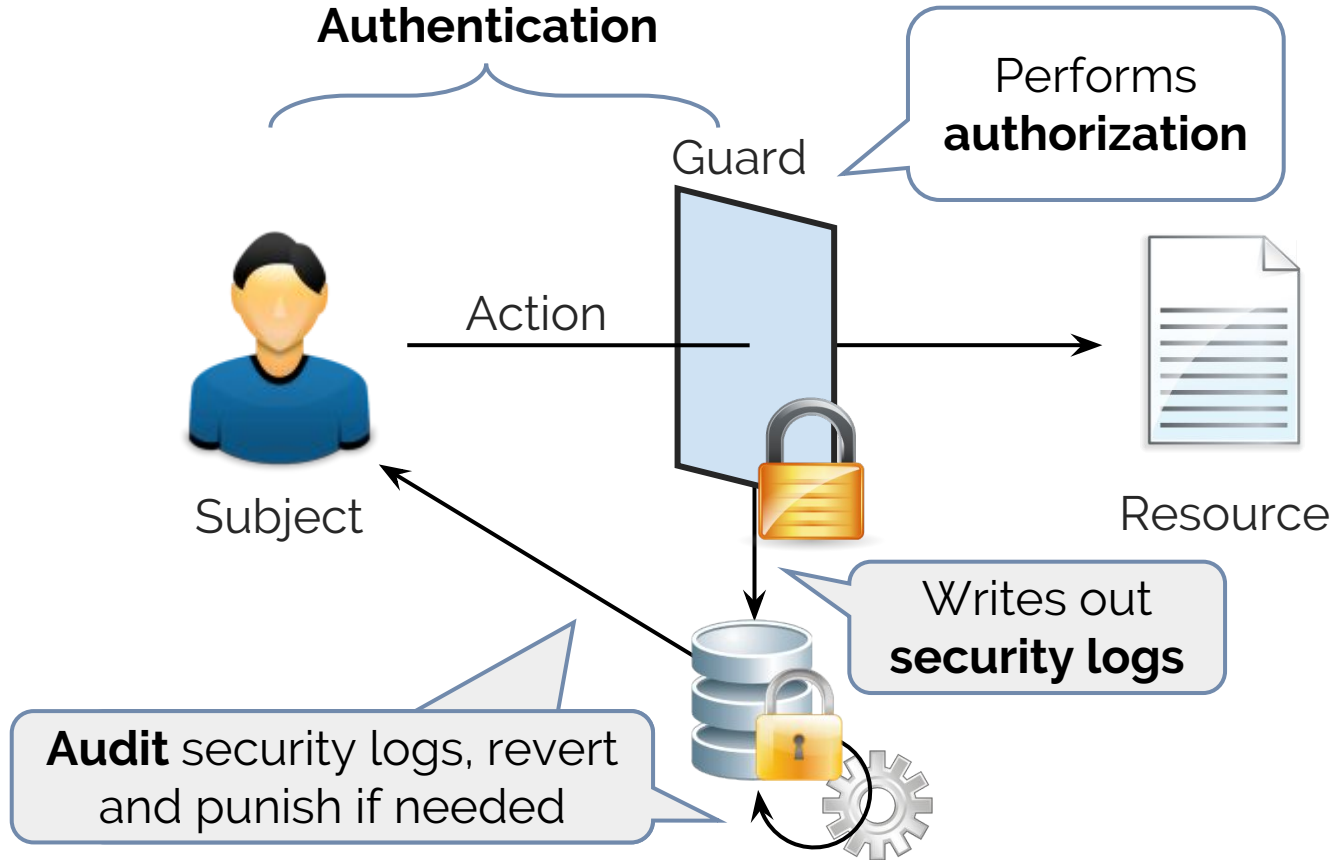


But there is more to it




# 5000m point of view

## Authentication



For the rest of this presentation

**“Access control” = “authorization”**



# Models, policies and mechanisms

- **Guard** is responsible for mediating access
  - Authorize specific actions
  - *Mechanism* that enforces specific *security rules*
- Rules, policies, models and mechanisms
  - **Access rules**: the logical access rules, independent of representation
  - **Mechanism**: low-level implementation of controls
  - **Model**: (formal) representation of how rules can be expressed
- Access control seems straightforward... but is it?

# Access control exists on multiple levels

Level	Subject	Action	Guard	Protected System
Hardware	OS Process	Read memory	CPU	CPU and Memory
Network	Host	Send packets	Firewall	Intranet
Database	Connecting application	SELECT query	DBMS	Data
OS	User	Open file	OS Kernel	Filesystem
Application	User	Read patient file	Application code	Application data


# CWE/SANS Top 25 Software Errors

Rank	Description
5	Missing authentication for critical function
6	Missing authorization
7	Use of hard-coded credentials
8	Missing encryption of sensitive data
10	Reliance on untrusted inputs in a security decision
11	Execution with unnecessary privileges
15	Incorrect authorization
17	Incorrect permission assignment for critical resource
19	Use of a broken or risky cryptographic algorithm
21	Improper restriction of authentication attempts
25	Use of a one-way hash without a salt

# Challenges for access control

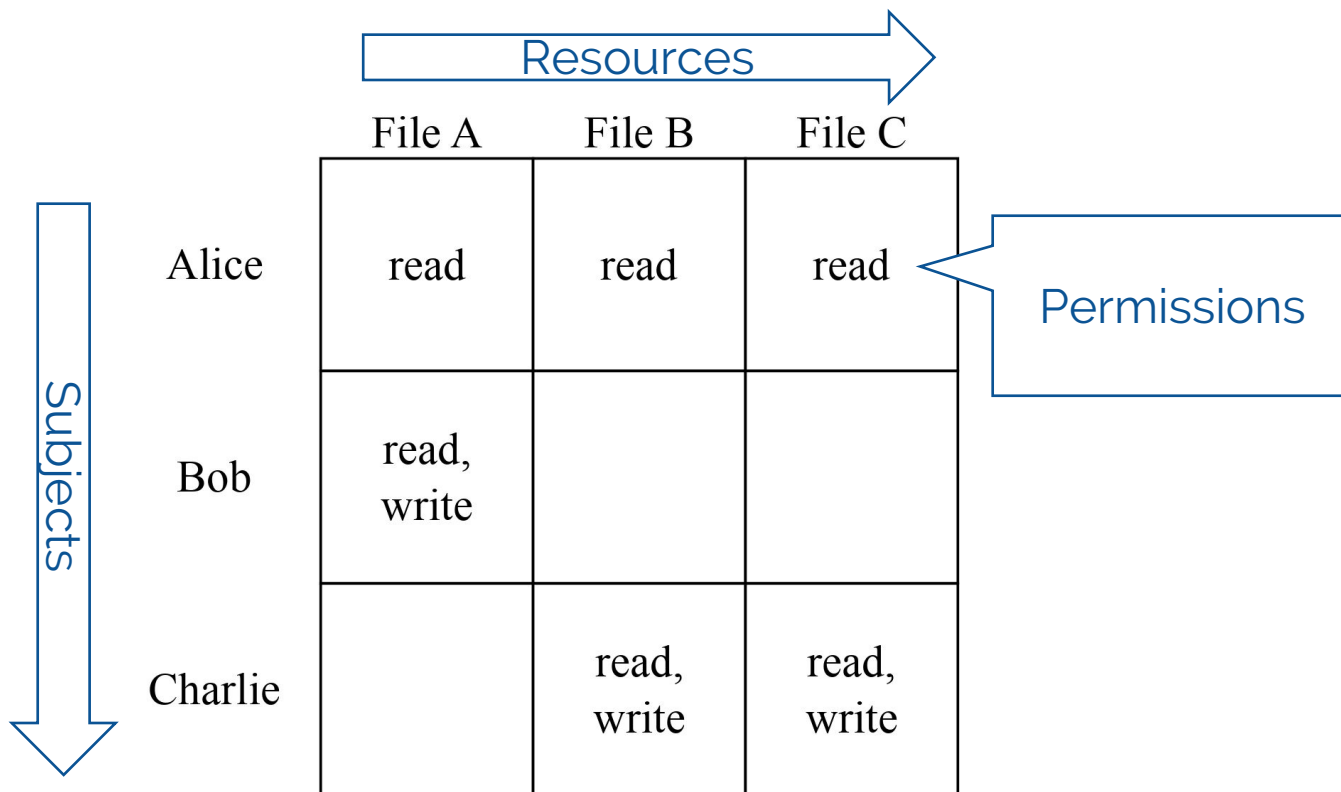
- **Expressiveness**: can the high-level rules be expressed in terms of the access control model of the policy/guard?
- **Performance**: access control decisions are frequent, and must be dealt with quickly
- **Full mediation**: does the guard check *every* action? Does your policy cover every action?
- **Safety**: does the access control mechanism match the policy?

# Outline

- Introduction
  - Positioning access control
  - Access control models
  - How to enforce access control
  - Some important technologies in practice
  - Recap and conclusion
- 
- The basics
  - Who can assign permissions
  - How permissions are assigned
  - Advanced topics



# The basics: the access control matrix



Extensions of the access control matrix:

Who can assign permissions?

# Who can assign permissions?

In general, two approaches:

1. Mandatory access control (MAC)
  - By central authority
2. Discretionary access control (DAC)
  - By subjects themselves

# Mandatory access control (MAC)

- Permissions are assigned by a central authority according to a central policy
  - Good fit within organizations and systems with a strong need for central controls
  - Low flexibility and high management overhead
- Mandatory Access Control in use
  - Often linked to multi-level security systems -> see later on
    - E.g. Government-regulated secrecy systems, military applications
  - Modern operating systems, to separate applications and processes
    - E.g. Windows' *Mandatory Integrity Control*, SELinux, TrustedBSD

# Example: SELinux

- Security-Enhanced Linux
  - “A set of patches to the Linux kernel and some utilities to incorporate a strong, flexible MAC architecture into the major subsystems of the kernel [for] confidentiality and integrity”
  - Activated by default in Fedora, Red Hat Enterprise Linux, etc
- Enforce MAC policy to processes in order to limit access to files and network resources
  - Least privilege
- Policy-based (see later on)
  - Separation of policy from enforcement with well-defined policy interfaces
  - Changing a policy does not require a reboot

# Example: SELinux

user:role:type:level

```
~]$ ls -Z /usr/bin/passwd
```

```
-rwsr-xr-x. root root system_u:object_r:passwd_exec_t:s0 /usr/bin/passwd
```

```
~]$ ls -Z /etc/shadow
```

```
----- . root root system_u:object_r:shadow_t:s0 /etc/shadow
```

## SELinux policies:

- applications running in the `passwd_t` domain can access files labeled with the `shadow_t` type
- the `passwd_t` domain can be entered from the `passwd_exec_t` type

# Discretionary access control (DAC)

- Permissions are set *at the discretion* of the subjects, e.g., the resource owner
  - Highly flexible policy, where permissions can be transferred
  - Lack of central control makes revocation or changes difficult
- Discretionary access control in use
  - Controlling access to files
    - E.g., Windows Access Control Lists (ACL), UNIX file handles
  - Controlling the sharing of personal information
    - E.g., Social networks

# The Graham-Denning Model

- Extends the access control matrix:
  - Subjects are also resources
  - Resources have an **owner**
  - Subjects have a **controller**
  - Permissions can be made **transferrable**
- Matrix can be modified by 7 commands
  - Creating and destroying subjects and resources
  - Granting, transferring and revoking permissions

	Alice	Bob	File 1	File 2	File 3
Alice	control	owner	owner read write	owner	read*
Bob		control		read write	owner read



# The Graham-Denning Model

1. Subject Alice creates object File 1

	<b>Alice</b>	<b>File 1</b>
<b>Alice</b>	control	owner

2. Subject Alice creates subject P1

	<b>Alice</b>	<b>P1</b>
<b>Alice</b>	control	control
<b>P1</b>		

3. Subject Alice destroys object File 1

➔ Alice must own File 1

	<b>Alice</b>	<b>File 1</b>
<b>Alice</b>	control	owner

4. Subject Alice destroys subject P1

➔ Alice must control P1

	<b>Alice</b>	<b>P1</b>
<b>Alice</b>	control	control
<b>P1</b>		

# The Graham-Denning Model

5. Subject Alice grants a right read/read\* on File 1 to P1

➔ Alice must be owner of File 1

	Alice	P1	File 1
Alice	control	control	owner
P1			read

6. Subject Alice transfers a right read/read\* on File 1 to P1

➔ Alice must have a right read\* on File 1

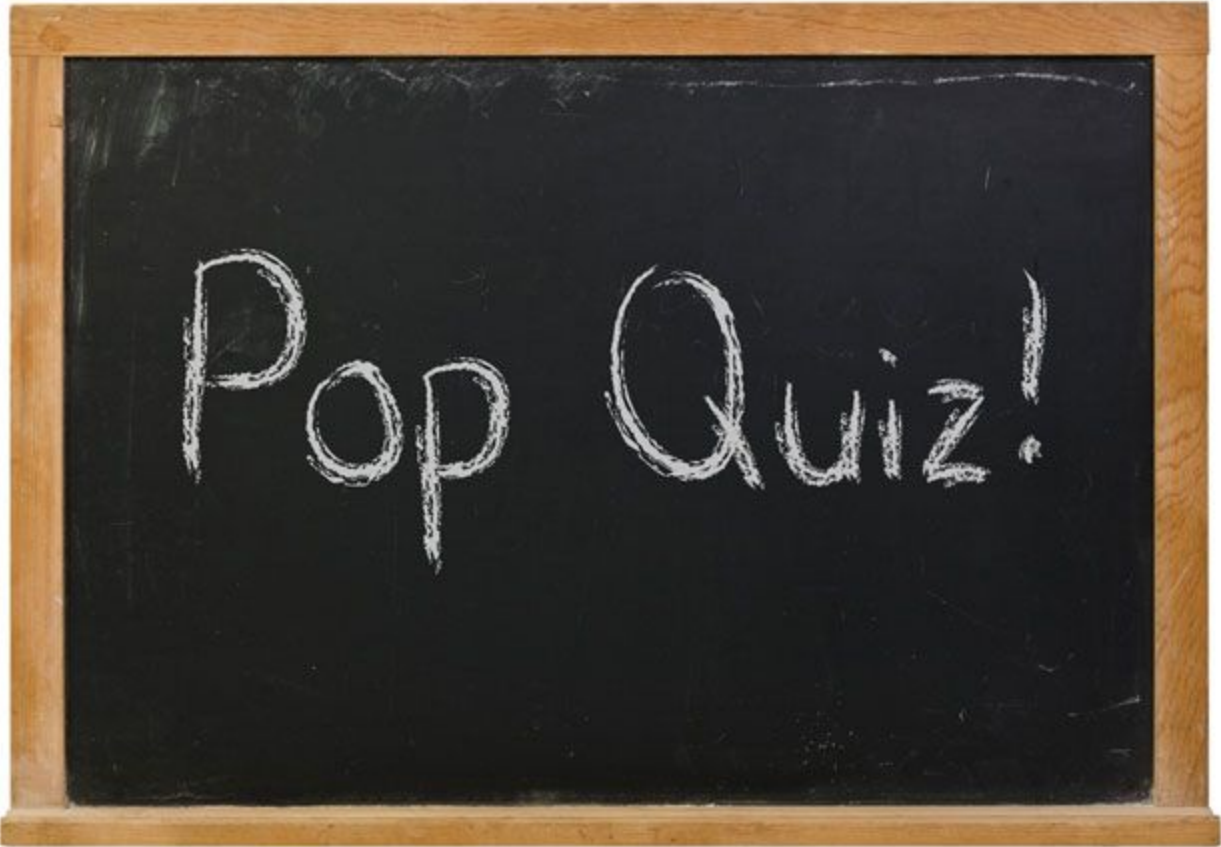
	Alice	P1	File 1
Alice	control	control	read*
P1			read

Only rights with a \* are transferrable

7. Subject Alice deletes a right read/read\* on File 1 from P1

➔ Alice must control P1 or Alice must own File 1

	Alice	P1	File 1
Alice	control	control	read*
P1			read

A rectangular chalkboard with a light-colored wooden frame. The board is black and has the words "Pop Quiz!" written in white chalk. The letters are slightly irregular and have some texture, suggesting they were hand-drawn. The exclamation point is at the end of the word "Quiz!".

Pop Quiz!

# Pop quiz!

How can Alice run a process P1 that can only read File 1?

	Alice	File 1	File 2
Alice	control	owner read write	owner read write

1. Subject Alice creates object File 1

2. Subject Alice creates subject P1

3. Subject Alice destroys object File 1

➡ Alice must own File 1

4. Subject Alice destroys subject P1

➡ Alice must control P1

5. Subject Alice grants a right read/read\* on File 1 to P1 ➡

Alice must be owner of File 1

6. Subject Alice transfers a right r/r\* on File 1 to P1

➡ Alice must have a right read\* on File 1

7. Subject Alice deletes a right r/r\* on File 1 from P1

➡ Alice must control P1 or Alice must own File 1

# Pop quiz!

- Starting state
- Subject Alice creates subject P1
- Subject Alice grants a permission read on resource File 1 to subject P1

	Alice	File 1	File 2
Alice	control	owner read write	owner read write

	Alice	P1	File 1	File 2
Alice	control	owner	owner read write	owner read write
P1		control		

	Alice	P1	File 1	File 2
Alice	control	owner	owner read write	owner read write
P1		control	read	

# More pop quiz!

- Can Alice read File 1?
- Could Alice ever read File 1?
- Could Bob ever read File 1?

	Alice	Bob	File 1	File 2
Alice	control		owner	owner read write
Bob		control		

1. Subject Alice creates object File 1

2. Subject Alice creates subject P1

3. Subject Alice destroys object File 1

➡ Alice must own File 1

4. Subject Alice destroys subject P1

➡ Alice must own P1

5. Subject Alice grants a right read/read\* on File 1 to P1 ➡

Alice must be owner of File 1

6. Subject Alice transfers a right r/r\* on File 1 to P1

➡ Alice must have a right read\* on File 1

7. Subject Alice deletes a right r/r\* on File 1 from P1

➡ Alice must control P1 or Alice must own File 1

# The question of safety

- The access control matrix implements a security policy
  - But DAC allows subjects to specify the access control policy
  - Given a specific starting state of the matrix and a given set of commands, can we prove any properties of all reachable states?
    - E.g. (Bob, Passwords File, Read) will never be granted
- Harrison-Ruzzo-Ullman model
  - Simplified framework, with six commands to manipulate the matrix
  - Impossible to build a security argument for the general case

# Recap: MAC vs DAC

- Two dual approaches
- In practice: combine both
  - Provide some form of discretionary self-management within the constraints of mandatory access rules
    - For example, delegate administration of team resources to an administrator
  - Options:
    - Trust subjects to enforce mandatory policy
    - Audit mandatory policy
    - Enforce mandatory policy



Extensions of the access control matrix:

How are permissions assigned?

# Existing models

- Identity-based access control
- Multi-level access control
- Role-based access control (RBAC)
- Attribute-based access control (ABAC)

# Identity-based access control

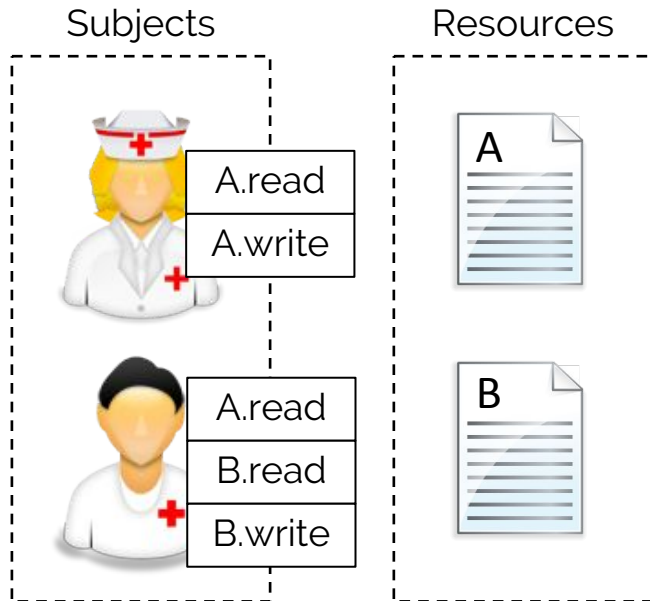
- Assign permissions to individual subjects and resources
  - This is actually again the Access Control Matrix

	File A	File B	File C
Alice	read	read	read
Bob	read, write		
Charlie		read, write	read, write

# Identity-based access control

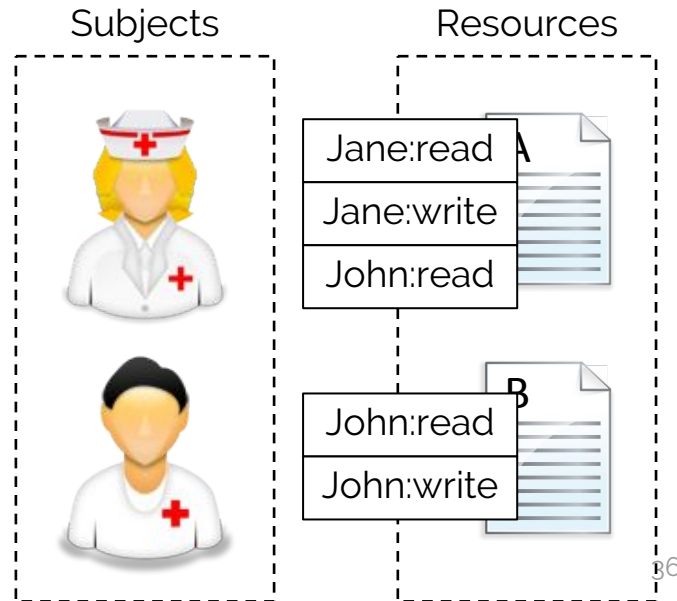
Possible implementations: store **1 big matrix** (not efficient) or:

## Access Control Lists



	File A	File B
Jane	Read Write	
John	Read	Read Write

## Capability Lists





# Identity-based access control

- Disadvantages:
  - Large management effort
    - E.g., “all nurses can read patient files” -> repeat for all nurses
    - E.g., “patients can read their own patient files” -> repeat for all patients
  - Information can be leaked
    - E.g., malicious user
    - E.g., Trojans
    - To address this: control access to information throughout the system
    - Common model for this: multi-level access control

# Multi-level access control

- Sometimes also called Lattice-Based Access Control
- Strict control over information flow
  - Resources are assigned **security classifications**
  - Subjects (and their programs) are assigned **security clearances**
  - These **labels** are organized in a lattice
- Two well-known rule sets:
  - Bell-LaPadula (confidentiality)
  - Biba (integrity)

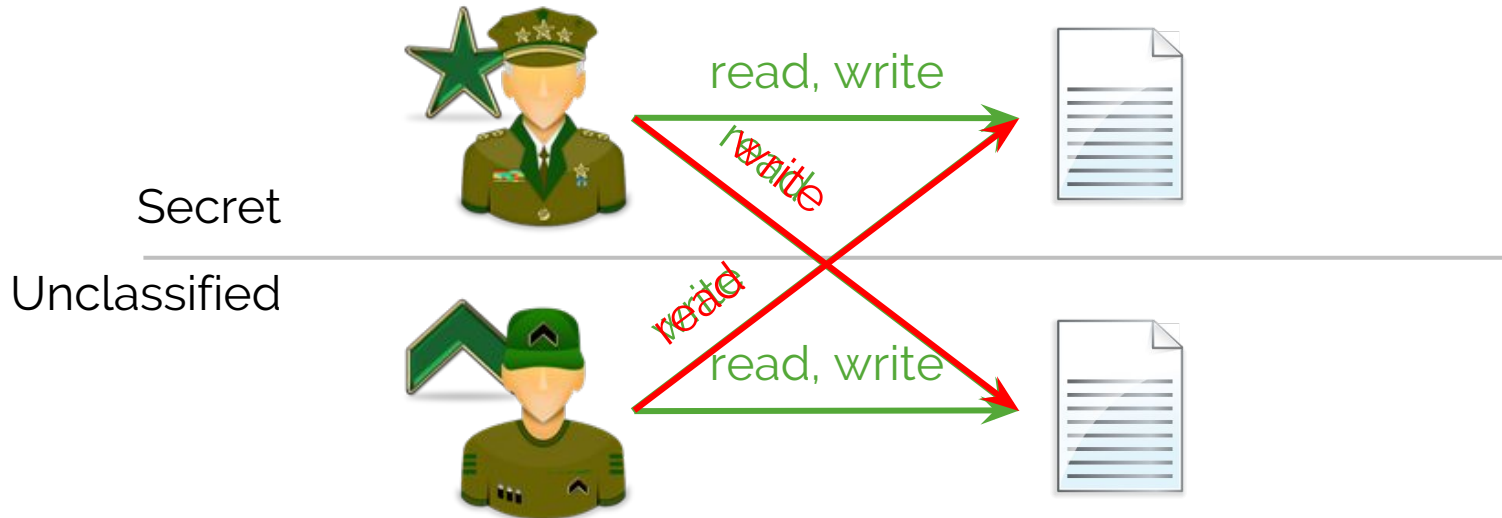


# Multi-level access control

- Model of Bell-LaPadula:

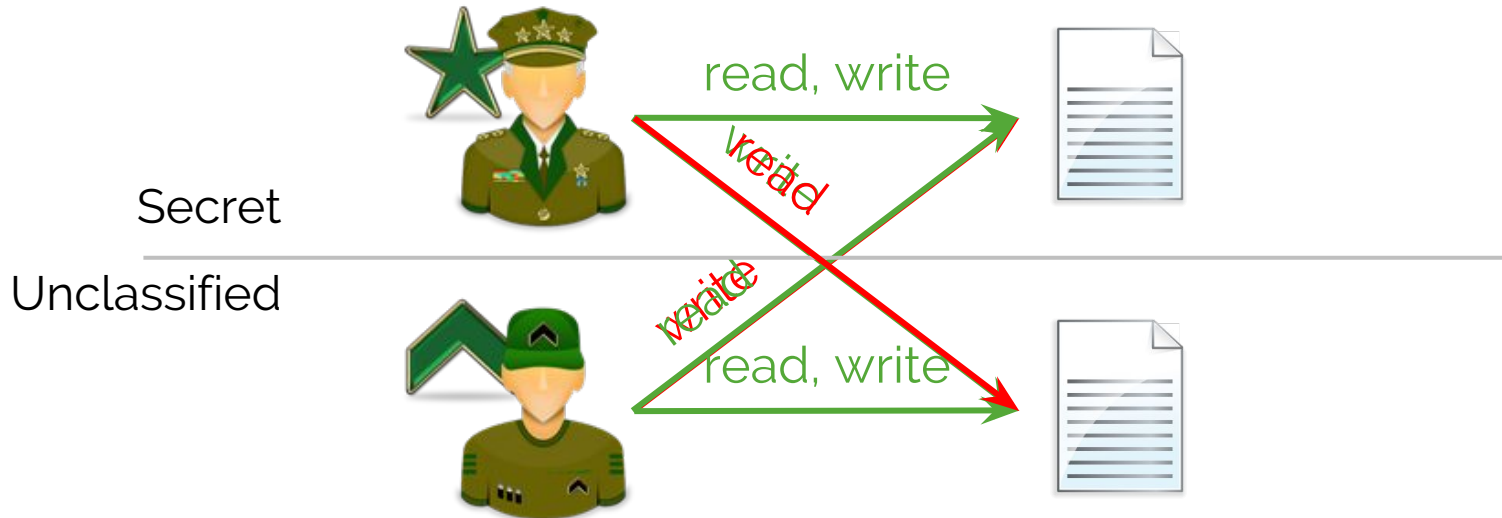
- No read up
- No write down (“☆-property”)

} Confidentiality



# Multi-level access control


- Model of Biba:
    - No write up
    - No read down
- } Integrity





# Multi-level access control

- You want both Bell-LaPadula and Biba
- However, this is not workable in practice
- => Refinement: **Information flow control, taint tracking**



```
var low, high  
if check(high) then  
    low := declassify(high)
```

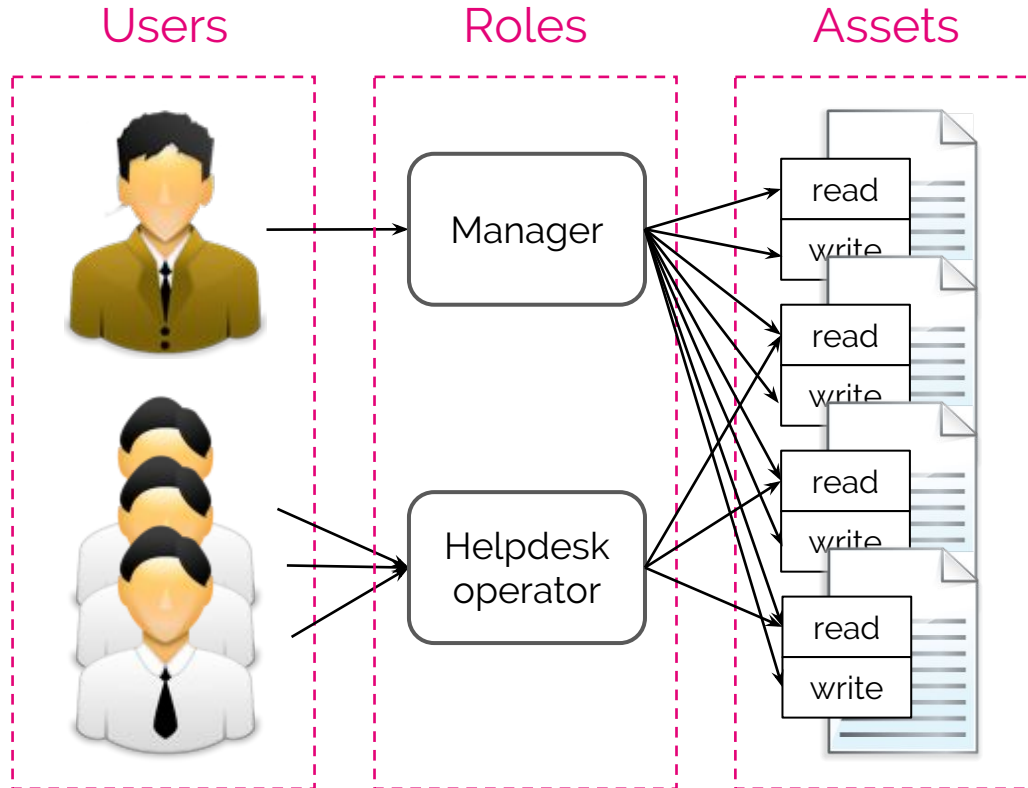
# Multi-level access control in the wild

- Core security feature of Windows Vista and newer
  - Complementary to discretionary access control
  - Control access to securable objects based on integrity level
  - Define the minimum integrity level required to access an object
- Isolate potentially untrustworthy contexts within the OS
  - Used by Google Chrome and Adobe Reader



Name	CPU	Private	Working Set	PID	Name	Integrity	User
svchost.exe		1.872 K	5.940 K	1844	Host Process for Windows S...	System	NT AUTHORITY...
lsass.exe	0.15	4.032 K	11.496 K	484	Local Security Authority Proc...	System	NT AUTHORITY...
lsm.exe	0.06	2.328 K	4.064 K	492	Local Session Manager Serv...	System	NT AUTHORITY...
winlogon.exe	0.01	2.488 K	6.844 K	416	Windows Logon Application	System	NT AUTHORITY...
explorer.exe	0.05	93.444 K	87.964 K	1416	Windows Explorer	Medium	Philippe-PC\Philippe
VBoxTray.exe	0.01	1.640 K	5.488 K	1180	VirtualBox Guest Additions Tr...	Medium	Philippe-PC\Philippe
POWERPNT.EXE	0.01	194.192 K	245.548 K	616	Microsoft PowerPoint	Medium	Philippe-PC\Philippe
WINWORD.EXE		44.144 K	91.400 K	3252	Microsoft Word	Medium	Philippe-PC\Philippe
procexp.exe		2.568 K	7.096 K	2932	Sysintemals Process Explorer	High	Philippe-PC\Philippe
procexp64.exe	0.99	14.356 K	25.040 K	2188	Sysintemals Process Explorer	High	Philippe-PC\Philippe
mspaint.exe		20.520 K	31.064 K	1112	Paint	Medium	Philippe-PC\Philippe
chrome.exe	0.05	44.944 K	72.500 K	236	Google Chrome	Medium	Philippe-PC\Philippe

# Role-based access control (RBAC)

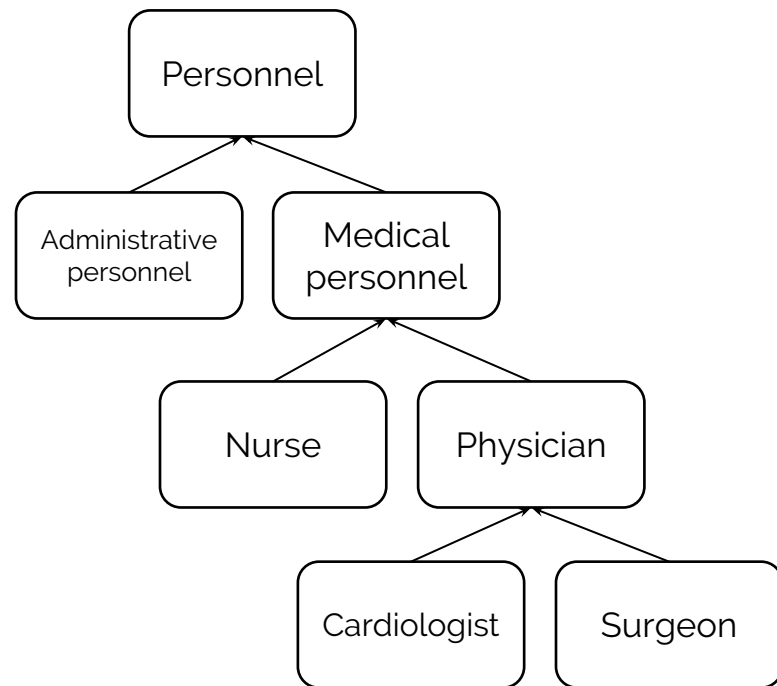


# Role-based access control (RBAC)

- Permissions assigned to roles, roles adopted by users
  - Goal: reduce large number of permissions to limited number of roles
  - Fits well onto the organizational structure of an enterprise
- Originated in research in 1992, NIST standard in 2004
- Immense research field
  - Role hierarchies, role mining, administrative models, delegation, constraints, least privilege, static separation of duty through meta-rules, ...

# Role-based

- Additional features in the NIST standard:
  - Role hierarchies
  - Least privilege through sessions
  - Static separation of duty through meta-rules
  - ...




# RBAC in the wild

- Almost any organization that I know of, employs roles
- Database systems often use and support RBAC
  - E.g., Oracle Enterprise Server
- Application development frameworks
  - Apache Shiro, Spring Security, ...
  - E.g., Java Spring Security:

```
@PreAuthorize("hasRole('manager')")  
public void create(Contact contact);
```

```
@PreAuthorize("hasPermission('delete_contact')")  
public void deleteContact(Contact contact);
```

# The problem with RBAC







# Role-based access control (RBAC)

- Major disadvantage: role explosion
- Reasons:
  - Roles cannot express ownership
    - Requires roles like “owns\_docA”, “owns\_docB”, etc
  - Reality is too fine-grained
    - Often small differences between different persons *in the same job*, leading to yet another role (e.g., “secretary\_with\_colorprint”)
  - Cross-product of multiple hierarchies
    - E.g., “sales\_manager\_for\_belgium\_with\_colorprint\_owns\_docA”
- To address this:
  - In practice: pragmatic choices, e.g., RBAC + ownership, RBAC + permissions, ...
  - Research: large number of extensions proposed

# Role-based access control (RBAC)

- Major disadvantage: role explosion

- Reasons:

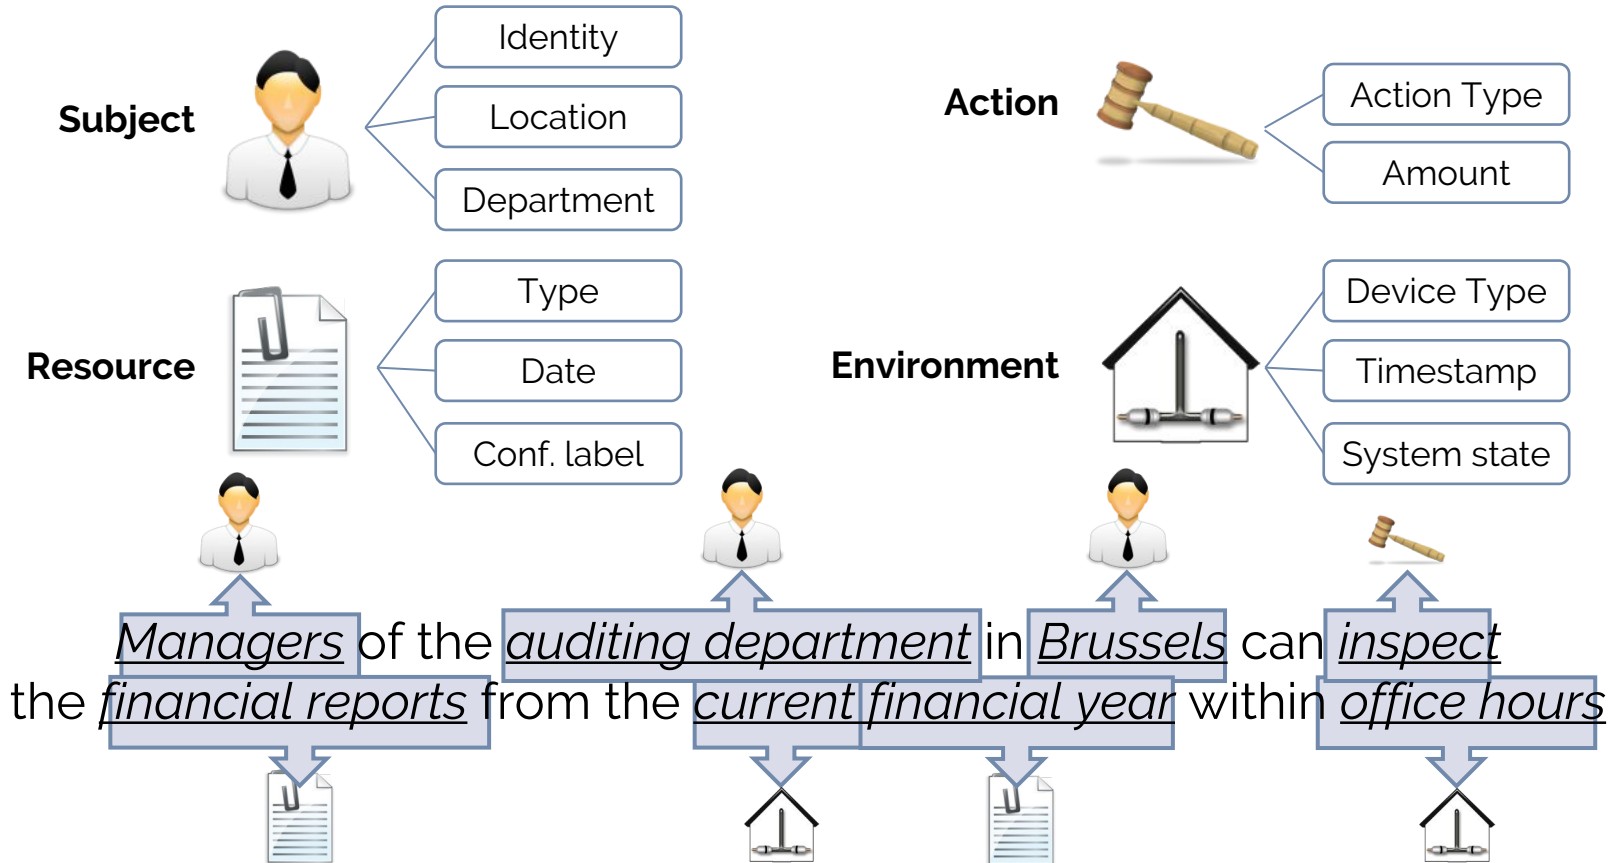
- Roles cannot express owners
  - Requires roles like "owns\_docA"
- Reality is too fine-grained
  - Often small differences between (e.g., "secretary\_with\_colorprint")
- Cross-product of multiple hierarchies
  - E.g., "sales\_manager\_for\_belgium"

PERMISSION	ANONYMOUS USER	AUTHENTICATED USER	ADMINISTRATOR
<b>Comment</b>			
Administer comments and comment settings	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>
View comments	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Post comments	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Skip comment approval	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Edit own comments	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>

- To address this:

- In practice: pragmatic choice for RBAC + ownership, RBAC + permissions, ...
- In research: large number of extensions proposed

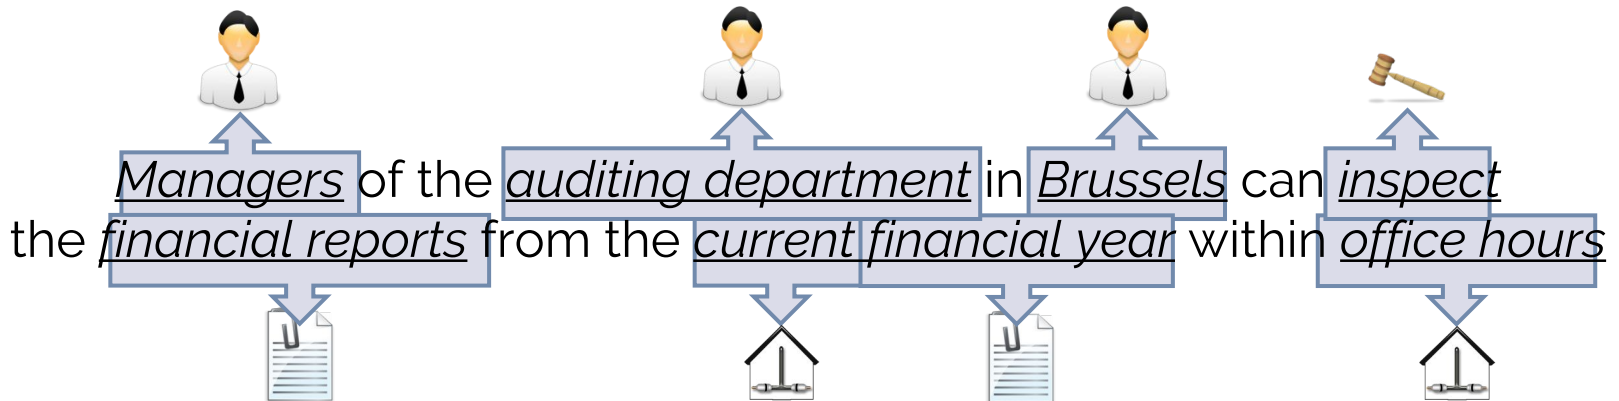
# Attribute-based Access Control (ABAC)



# Attribute-based Access Control (ABAC)

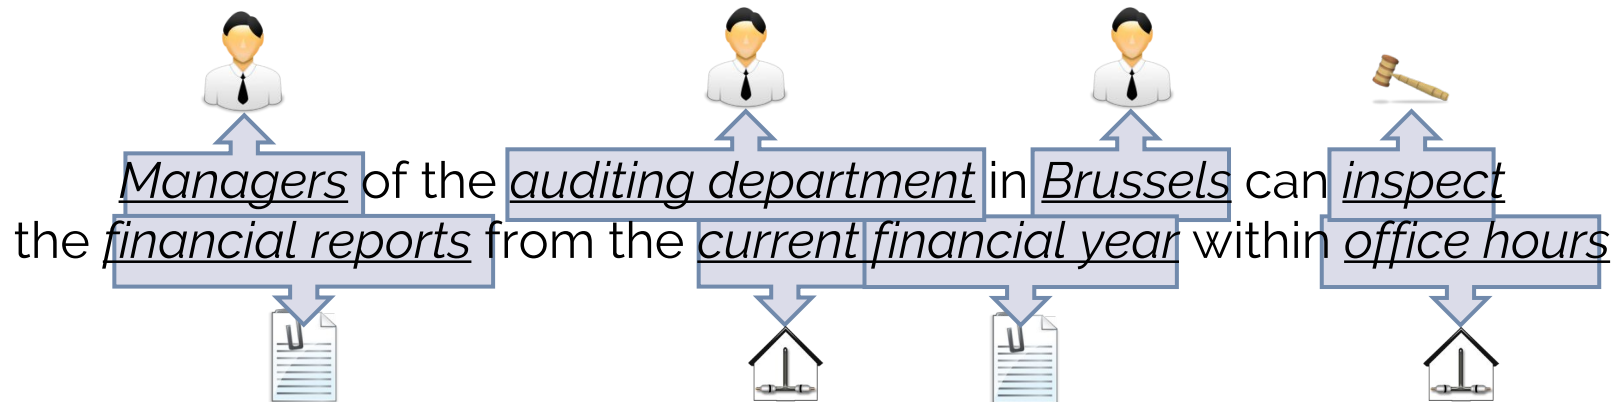
permit if

“manager” in subject.roles and subject.department == “auditing”  
and subject.location == “Brussels” and action == “inspect”  
and resource.type == “financial report”  
and resource.year == environment.current\_year  
and 8h00 < environment.time < 17h00



# Attribute-based Access Control (ABAC)

1. fine-grained access control
2. context-aware access control
3. dynamic access control



# Attribute-based Access Control (ABAC)

- Access decisions are made based on attributes
  - Attributes are key-value properties of the subject, the resource, the action or the environment
  - Results into dynamic and context-aware access control
- Attributes can express many different access control concepts
  - Permissions, roles, groups, departments, time, location, ownership, domain-specific ownership, ...


# Attributes as an enabler for the future

1. Maps better to business policies
2. Provides a new methodology of managing users
3. Attributes can be fetched remotely = good for federated applications
4. You do not need the identity of the subject = good for privacy
5. As a researcher, it looks future-proof
  - a. ABAC supports many advanced policies, e.g., history-based policies, dynamic separation of duty and breaking-the-glass procedures, ...
  - b. Many of the newest access control models can be mapped on attributes, e.g., ReBAC, EBAC [Bogaerts2015], obligations [Park2004], ...
  - c. A lot is still happening in this field, e.g., formal definition of this model and its properties (e.g., [Jin2012a]), languages for expressing attribute-based rules (e.g., [XACML, Crampton2012]), mutable attributes (e.g., [Park2004]), attribute aggregation in federated identity management (e.g., [Chadwick2009]), encryption of attributes (e.g., [Asghar2011]), policy engineering for ABAC (e.g., [Krau2013]), performance (e.g., [Brucker2010]), ...

# Migrating from RBAC to ABAC



Conceptually, three approaches:

## 1. Roles as an attribute

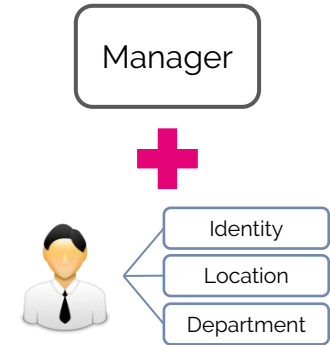


subject.roles

## 2. Dynamic roles



## 3. Constrain roles

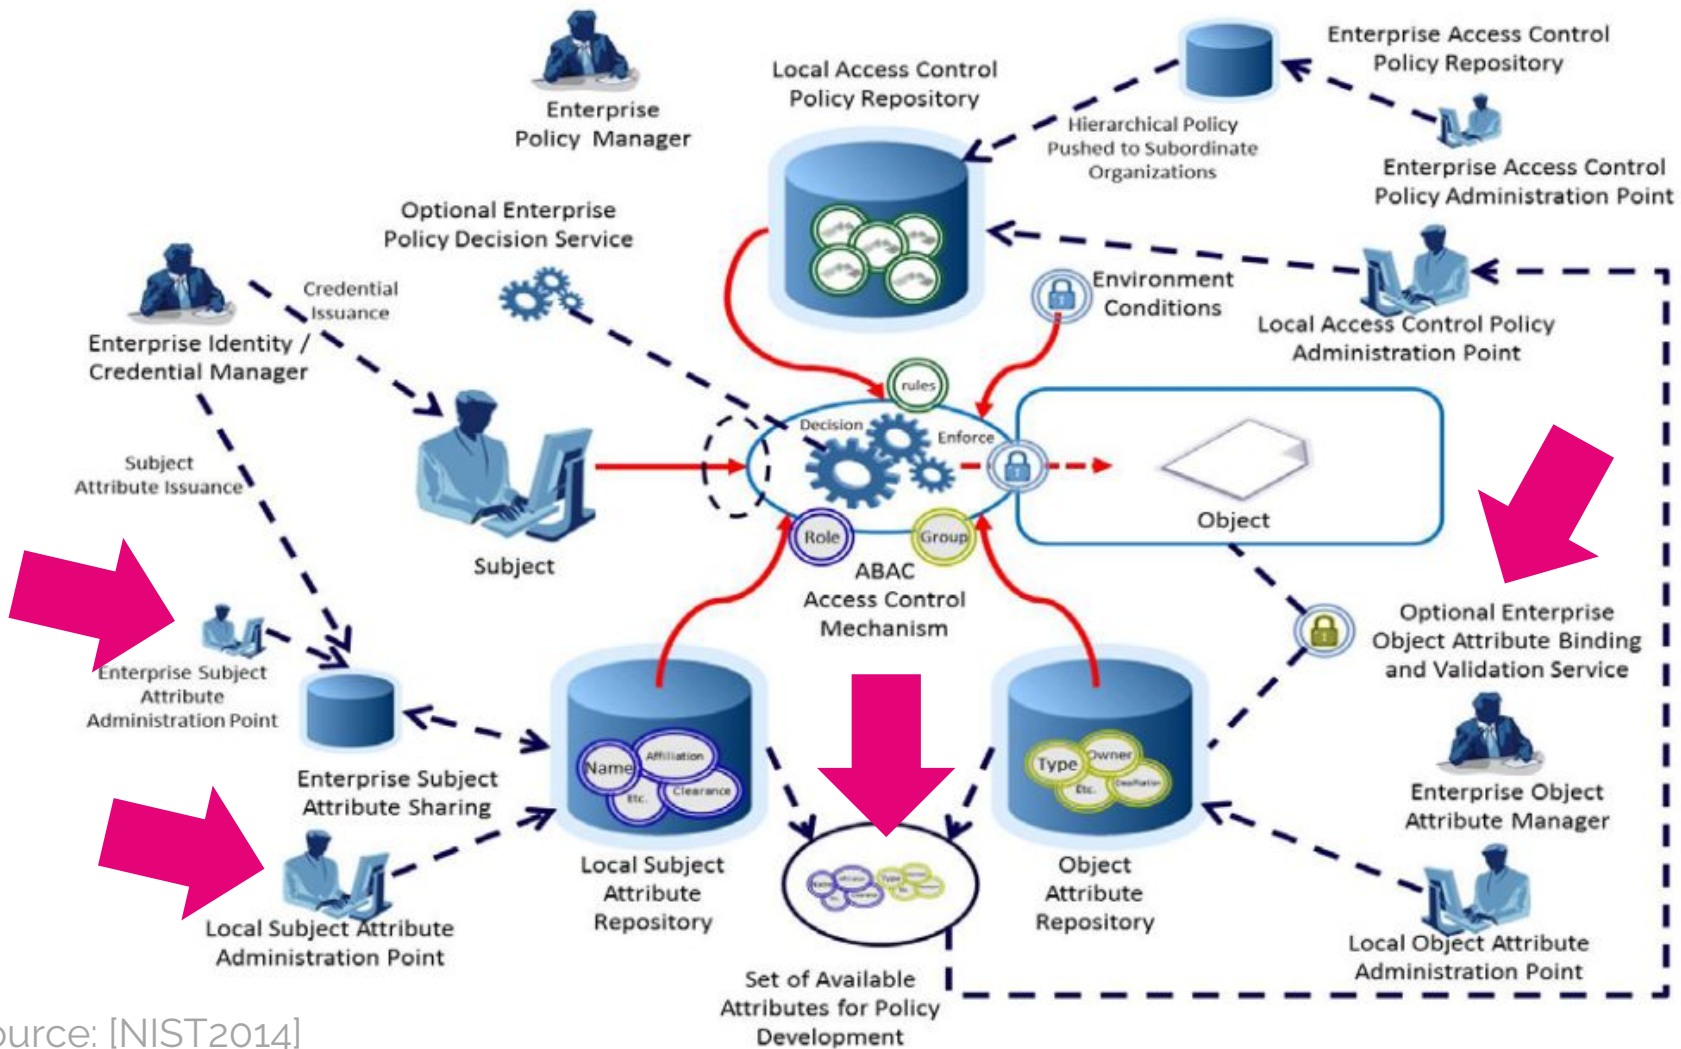


A.read	B.read
B.write	...

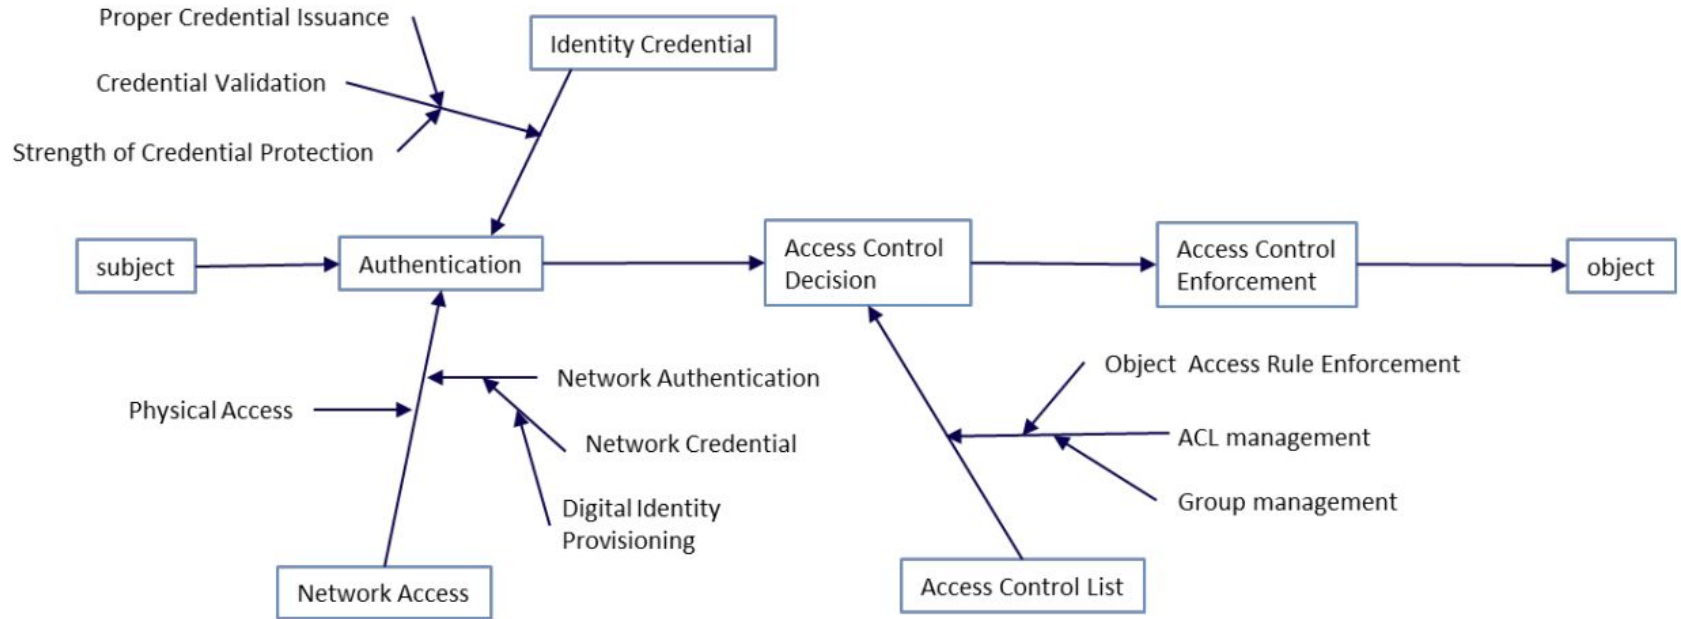


Not all rainbows and unicorns



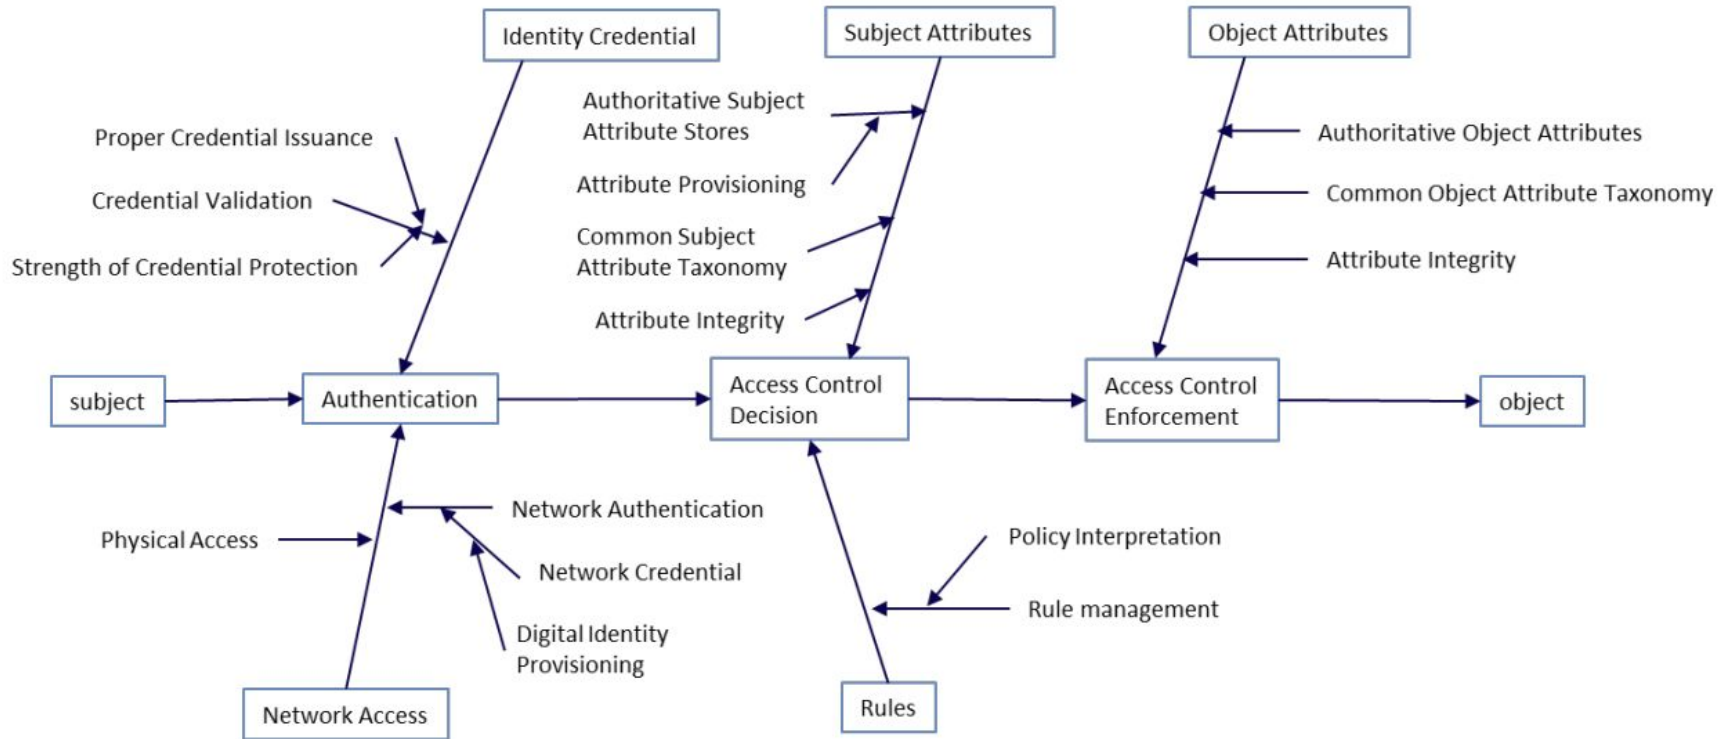


# Not all rainbows and unicorns



Trust chain for Access Control Lists

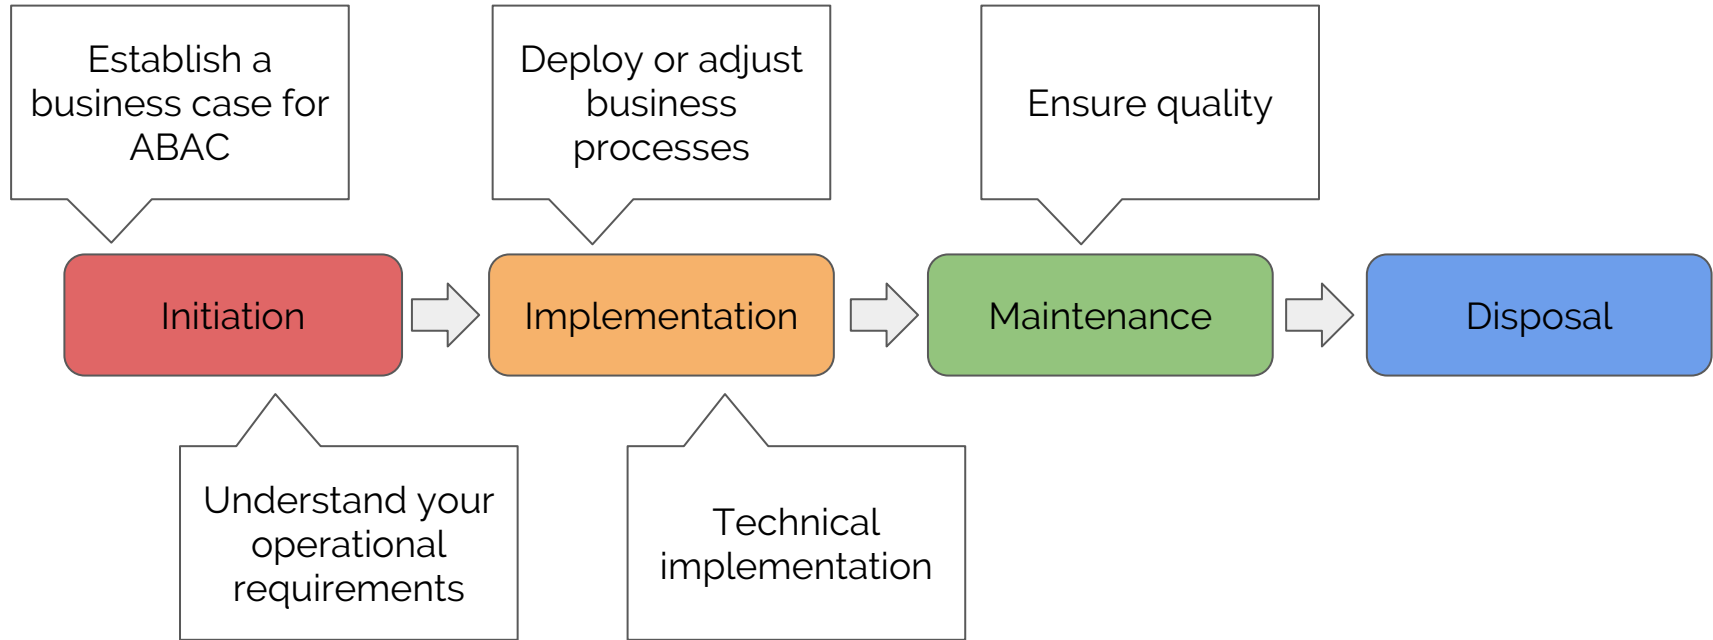
# Not all rainbows and unicorns



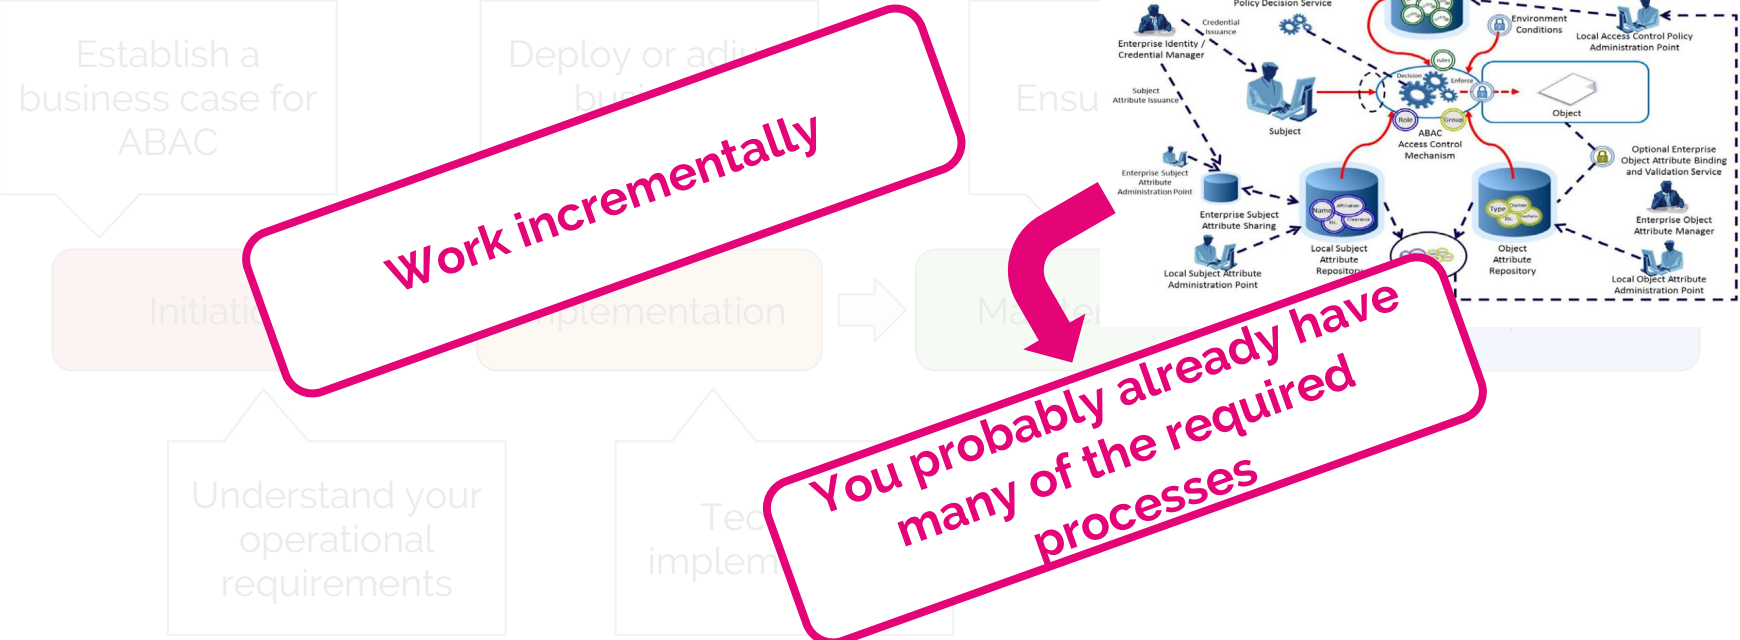
Trust chain for ABAC

“Enterprise ABAC carries with it significant development, implementation, and operations costs as well as a paradigm shift in the way enterprise objects are shared and protected.” -- NIST

# Migrating from RBAC to ABAC, revised



# Migrating from RBAC to ABAC, revised



**Work incrementally**

**You probably already have many of the required processes**

# ABAC: Conclusion

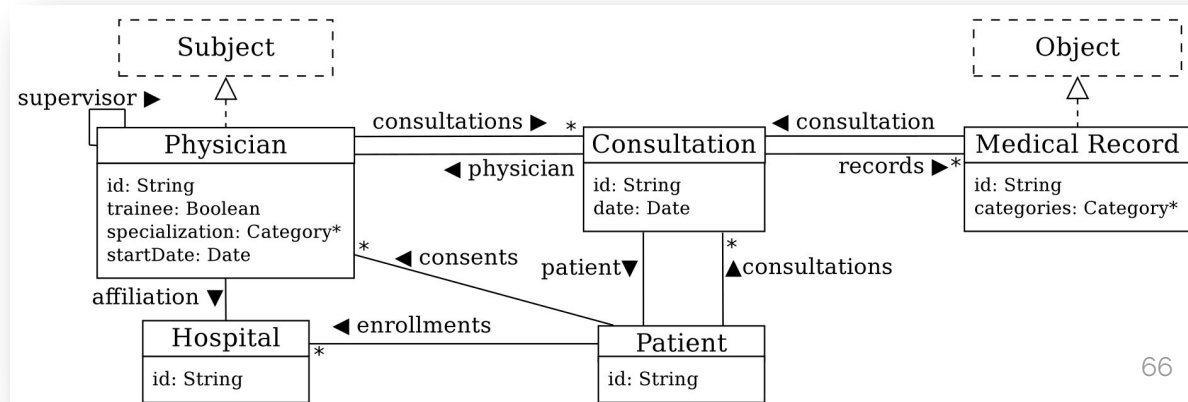
- ABAC brings many interesting improvements compared to previous models
- ABAC is seen by many as the next step in access control
- => Definitely something you should consider, but not a small step to take. **Work incrementally**
- Further reading: [NIST2014]
  - Overview of ABAC, challenges and enterprise considerations



# Advanced topics

# Advanced topics

- Relationship-Based Access Control
  - Originated from social networks
  - Further reading: [Cheng2012, Fong2011]
- Entity-Based Access Control
  - Express access rules in terms of the entities in your application
    - Attributes + relationships
  - Fixes limitations of ABAC
  - I expect a lot of this, but still a long way to go
  - Further reading:
    - [Crampton2014]
    - [Bogaerts2015]



# Advanced topics

- Advanced policy pattern: breaking the glass
  - Enable users to override a deny by “breaking the glass”
  - Common pattern in e-health
    - “A physician should be able to override a deny when a patient is in critical condition”
  - Challenge: *controlled* override
    - Limit who can override a deny (e.g., only physicians of emergency department), limit for which actions a deny can be overridden (e.g., only for reads)
    - Audit these overrides later on, e.g., by writing out logs at override

# Advanced topics

- Advanced policy pattern: separation of duty
  - Separate duties within an organization
  - Statically:
    - E.g., “a manager can never also be a secretary”
    - E.g., “a manager cannot approve his own funding requests”
  - Dynamically:
    - E.g., “if a user has had access to documents of Bank A, he or she is not allowed to access documents of Bank B”
    - Originally described in 1989 as the “Chinese wall policy”, a “commercial security policy” in contrast to “Bell-LaPadula-style policies” [Brewer1989]
  - Very relevant because of Sarbanes-Oxley, but still a hard problem
    - Hard to apply to an organization
    - Hard to implement well (performance issues)

# Advanced topics

- History-based access control
  - E.g., dynamic separation of duty
  - E.g., limit the number of accesses
    - “a user cannot watch more than 10 movies per month”
- Implementation options:
  - Use log files in the policy evaluation
  - Use provenance data in the policy evaluation [Nguyen2012, Nguyen2013]
  - Explicitly update history attributes [Decat2015]

# Advanced topics

- History-based

- E.g., dynamic

- E.g., limit

- “a user  $\in$  “Bank A” in subject.history

**Deny if**

**When resource.owner == “Bank B”,  
apply DenyOverrides to**

**Permit performing**

append(“Bank B”, subject.history)

Obligations

- Implementation options:

- Use log files in the policy evaluation

- Use provenance data in the policy evaluation [Nguyen2012, Nguyen2013]

- Explicitly update history attributes [Decat2015]

# Advanced topics

- Obligations
  - Early definition: “predicates that verify mandatory requirements a subject has to perform before or during a usage exercise” [Park2004]
    - Pre-obligations, ongoing-obligations
    - Examples:
      - User has to agree to terms and conditions (pre)
      - User has to be shown an ad during watching the requested movie (ongoing)
  - More pragmatic definition: action that should be performed with permitting/denying the action
    - Send an e-mail to an administrator on deny to a confidential document
    - Write out log
    - Update attribute


# Outline

- Introduction
- Positioning access control
- Access control models
- How to enforce access control
- Some important technologies in practice
- Recap and conclusion

- Reference monitors
- Access control in application code
  - Policy-based access control

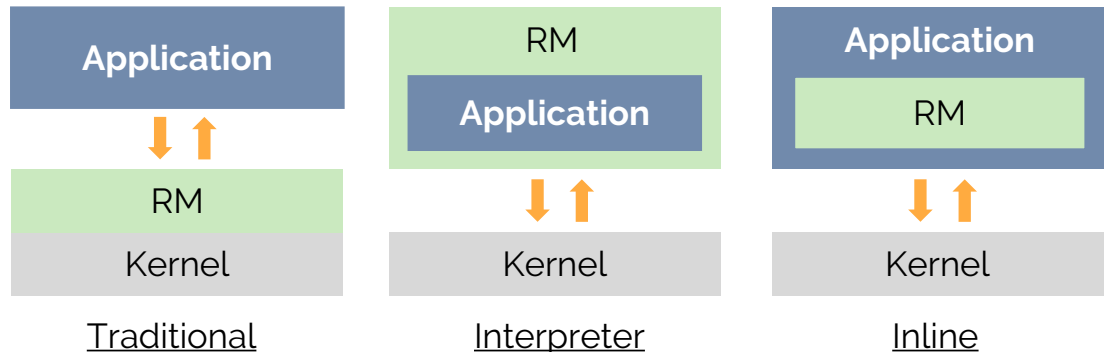


# How to enforce access control




# Reference monitors

- Reference monitors
  - Observe software execution
  - Take remedial action on operations that violate a policy
- Three important security properties
  - Full mediation
  - Tamper proof
  - Verifiable



# Example of a reference monitor

- Antivirus software is implemented as reference monitor
  - Hooks into the OS's system calls to intercept application actions
  - E.g. inspects file contents upon read or write operations
- Good implementation strategy to meet security properties
  - Full mediation: requires coverage of all system calls
  - Tamper proof: requires strong process isolation
  - Verifiable: less straightforward, but possible



# Access control exists on multiple levels

Level	Subject	Action	Guard	Protected System
Hardware	OS Process	Read memory	CPU	CPU and Memory
Network	Host	Send packets	Firewall	Intranet
Database	User	SELECT query	DBMS	User database
OS	User	Open file	OS Kernel	Filesystem
Application	User	Read patient file	Application code	Application data

# Application-level access control

- Rules reason about the concepts in your application
- Add guard to code of your application
- The same holds:
  - Full mediation
  - Tamper proof
  - Verifiable


# Option 1: encode guard and rules in app code

```
public Document getDoc(docId) {
    Doc doc = db.getDoc(docId);
    if (! ("manager" in user.roles
        && doc.owner == user
        && 8h00 < now() < 17h00 )) {
        return null;
    } else {
        return doc;
    }
}
```

- + straightforward
- + you can encode almost anything
- no separation of concerns
- no modularity
  - => hard for reviews
- what if rules change?
  - update application code
  - updates all over the place

## Option 2: modularize

```
public Document getDoc(docId) {
    Doc doc = db.getDoc(docId);
    if (! ("manager" in user.roles
        && doc.owner == user
        && 8h00 < now() < 17h00 )) {
        return null;
    } else {
        return doc;
    }
}
```

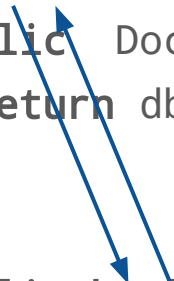


```
@authz(user, "read", result)
public Document getDoc(docId) {
    return db.getDoc(docId);
}
...
public boolean authz(
    user, action, resource) {
    if (!("manager" in user.roles
        && ...)) {
        return true;
    } else {
        return false;
    }
}
```

## Option 2: modularize

- + more modularity: access control logic in 1 place
- no separation of concerns
- ± what if rules change?
  - update application code
  - + updates in one place

```
@authz(user, "read", result)
public Document getDoc(docId) {
    return db.getDoc(docId);
}
...
public boolean authz(
    user, action, resource) {
    if (!(“manager” in user.roles
        && ...)) {
        return true;
    } else {
        return false;
    }
}
```





## Option 2: modularize - Django

settings.py:

```
AUTHENTICATION_BACKENDS = [  
    'mymodule.MyBackend'  
]
```

mymodule/backends.py:

```
class MyBackend(object):  
  
    def has_perm(self, user, perm, obj):  
        if obj.owner == user.id:  
            return True  
        else:  
            return False
```

# Option 2: modularize – Ruby on Rails

In the controller:

```
def show
  @article = Article.find(params[:id])
  authorize! :read, @article
end
```

In the view:

```
<% if can? :update, @article %>
  <%= link_to "Edit",
    edit_article_path(@article) %>
<% end %>
```

The access control code:

```
class Ability
  include CanCan::Ability

  def initialize(user)
    if user.admin?
      can :manage, :all
    else
      can :read, :all
    end
  end
end
```

# Option 2: modularize – Java Spring Security

In the controller:


```
@PreAuthorize("hasPermission(#doc, 'view')")  
public void getDocument(Document doc);
```

In the PermissionEvaluator:


```
boolean hasPermission(Authentication a,  
                       Object resource, String permission) {  
    User user = SecurityUtil.getUserCredential();  
    if(permission == "view" and ...) {  
        return true;  
    } else {  
        return false;  
    }  
}
```

# Option 3: policy-based access control

```
@authz(user, "read", result)
public Document getDoc(docId) {
    return db.getDoc(docId);
}
...
public boolean authz(
    subject, action, resource) {
    if (! ("manager" in user.roles and ...)) {
        return true;
    } else {
        return false;
    }
}}
```



```
@authz(user, "read", result)
public Document getDoc(docId) {
    return db.getDoc(docId);
}
```



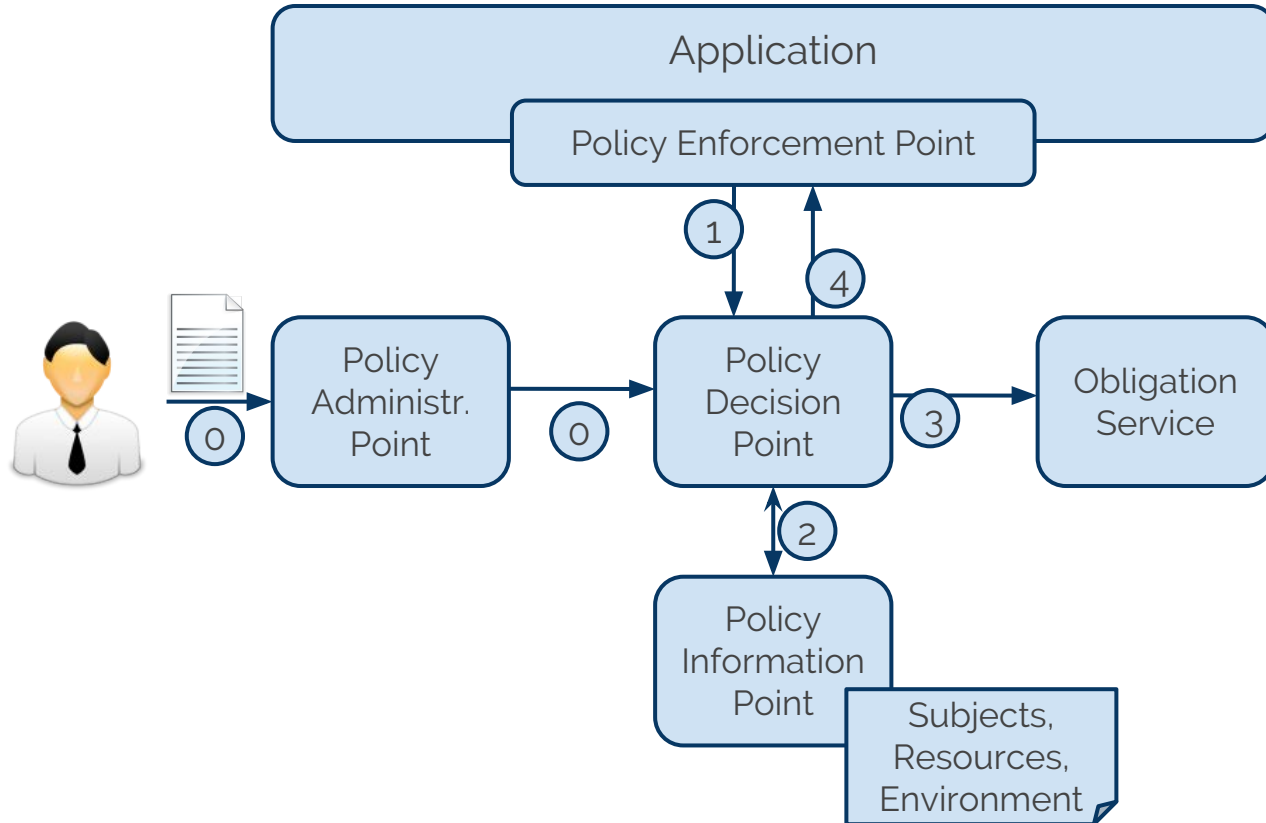
# Option 3: policy-based access control

- Decouple access control rules from application code
  - Express access control rules in a format independent of your programming language
  - In application code: ask the generic question “can this subject perform this action on this resource”?
  - Policy evaluated by specialized component called the Policy Decision Point
  - If policy is stored in a file or a database: change policy at run-time

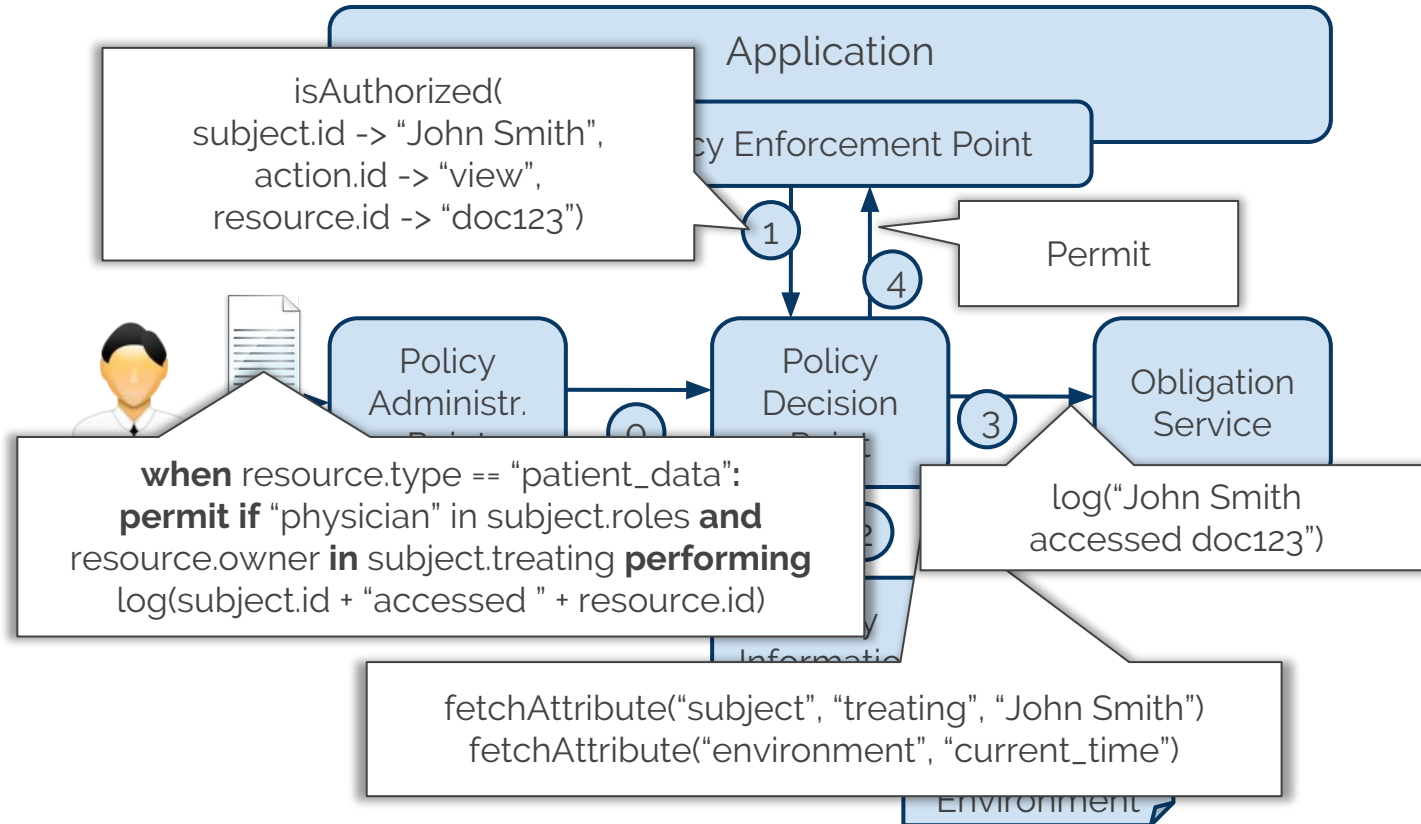
# Advantages of PBAC

- + More modularity: access control logic in 1 place
- + Separation of concerns: policies can be written by non-developer
- + What if rules change?
  - + no updates in application code
  - + updates in a single place
- + Enables your access control policies to easily evolve with your organization
- + Access rules are software artifacts => automated refinement, monitoring, validation, ...

# XACML Reference architecture



# XACML Reference architecture

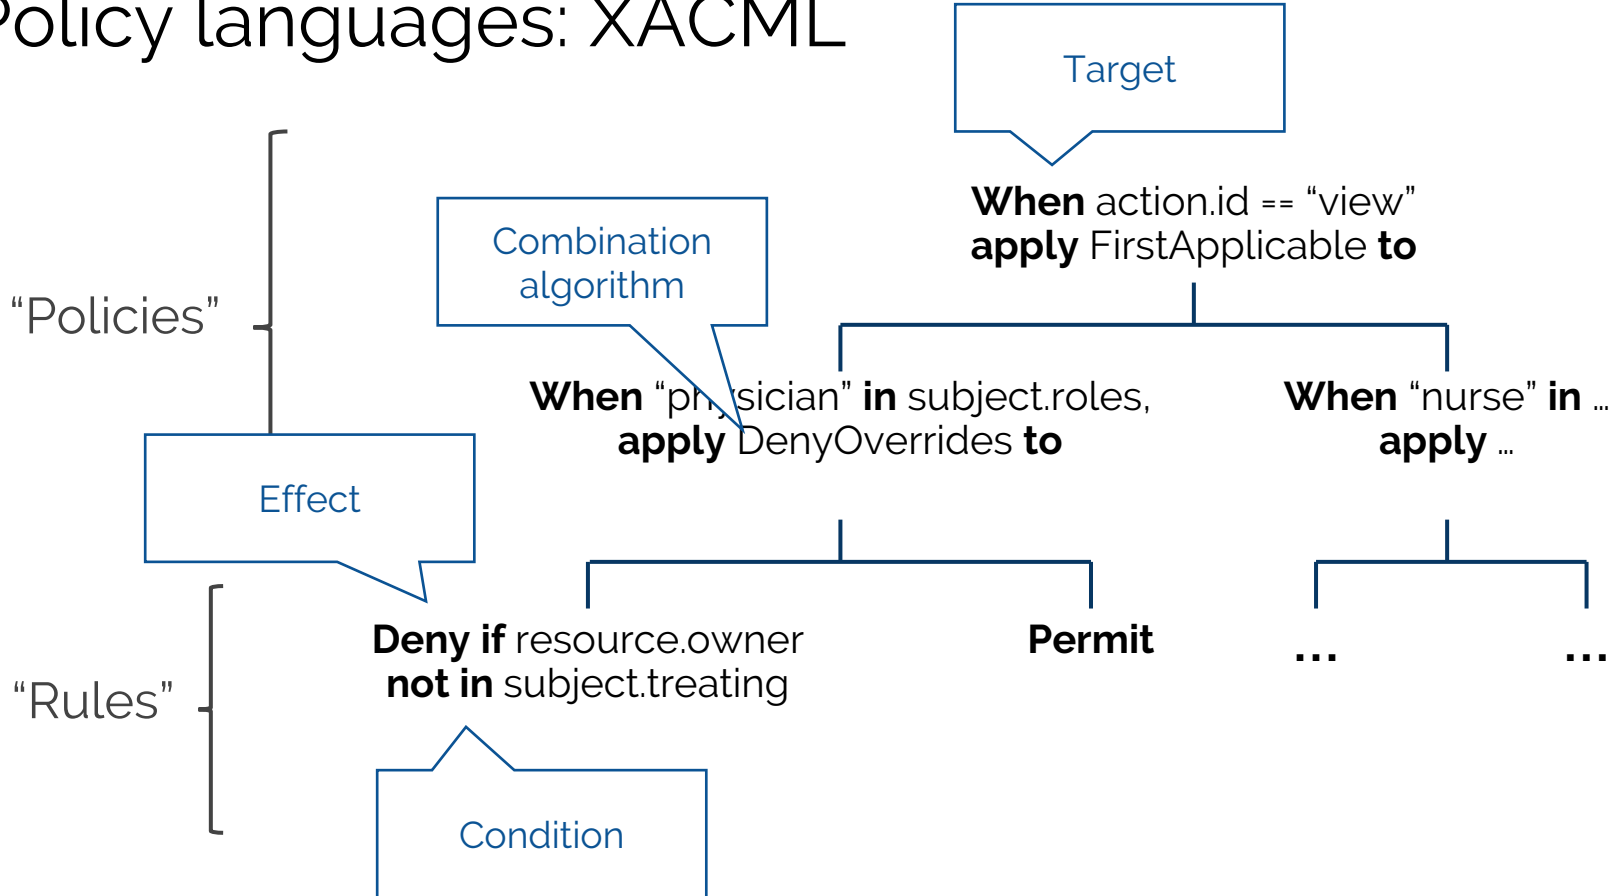




# Policy languages

- A large number of domain-specific policy languages proposed in literature
  - E.g., SPL, Ponder, XACML, Cassandra, SecPAL, ...
- Current major standard: XACML
  - Standardized by OASIS
    - v1.0 ratified in 2003, v3.0 in 2013
  - Attribute-based, tree-structured, obligations
  - XML format

# Policy languages: XACML



# Policy languages: XACML

```

<Rule RuleId="treatings" Effect="Deny">
  <Description>treatings</Description>
  <Condition>
    <Apply FunctionId="string-equal">
      <AttributeValue DataType="string">doc123</AttributeValue>
      <SubjectAttributeDesignator AttributeId="subject:history" DataType="string"/>
    </Apply>
  </Condition>
</Rule>

<Rule RuleId="dynamic-separation-of-duty" Effect="Deny">
  <Description>Dynamic separation of duty</Description>
  <Target>
    <Resources>
      <ResourceMatch MatchId="string-equal">
        <AttributeValue DataType="string">doc123</AttributeValue>
        <ResourceAttributeDesignator AttributeId="resource:id" DataType="string"/>
      </ResourceMatch>
    </Resources>
  </Target>
  <Rule RuleId="deny" Effect="Deny">
    <Description>Deny if viewed other doc</Description>
    <Condition>
      <Apply FunctionId="string-is-in">
        <AttributeValue DataType="string">doc456</AttributeValue>
        <SubjectAttributeDesignator AttributeId="subject:history" DataType="string"/>
      </Apply>
    </Condition>
  </Rule>
  <Rule RuleId="default-permit" Effect="Permit"> </Rule>
  <Obligations>
    <Obligation ObligationId="append-attribute" FulfillOn="Permit">
      <AttributeAssignment AttributeId="value" DataType="string">
        <SubjectAttributeDesignator AttributeId="resource:id" DataType="string"/>
      </AttributeAssignment>
      <AttributeAssignment AttributeId="attribute-id"
        DataType="string">subject:history</AttributeAssignment>
    </Obligation>
  </Obligations>
</Policy>

```

# STAPL

```
Rule("roles") := permit iff ("physician" in subject.roles)
```

```
Rule("ownership") := permit iff (resource.owner in subject.treating)
```

```
Rule("time") := deny iff (env.currentDateTime > (resource.created + 5.days))
```

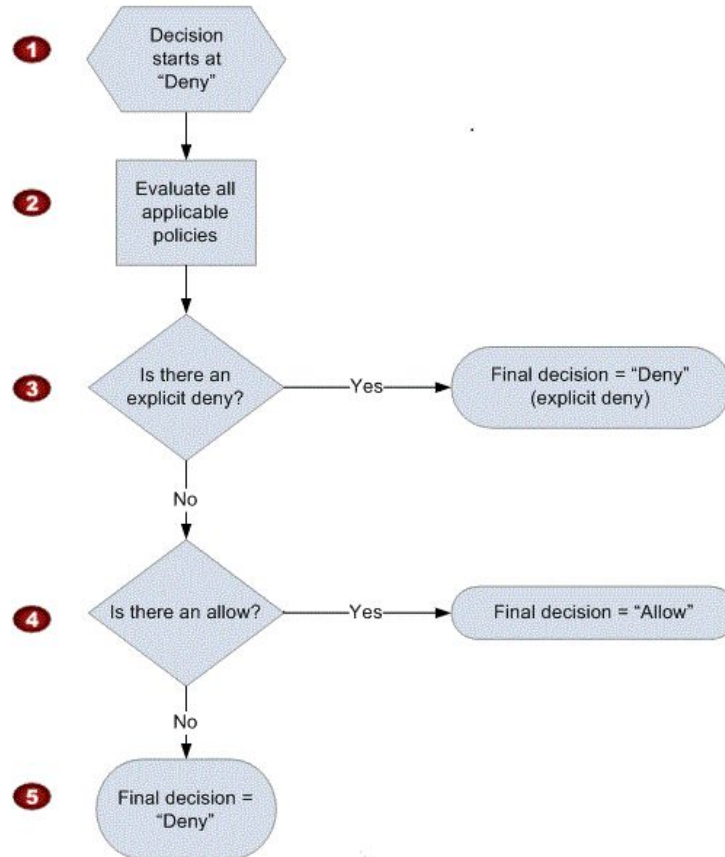
```
Policy("dynamic SoD") := when (resource.id === "doc123") apply Deny0verrides to (  
    Rule("deny") := deny iff ("doc456" in subject.history),  
    defaultPermit  
) performing (append(resource.id, subject.history) on Permit)
```

# PBAC in the wild: Amazon EC2

The screenshot shows the AWS IAM console interface. At the top, there are navigation tabs for 'AWS', 'Services', and 'Edit', along with 'Global' and 'Support' dropdowns. On the left, a sidebar contains navigation links: Dashboard, Search IAM, Details, Groups, Users, Roles, Policies (highlighted), Identity Providers, Account Settings, Credential Report, and Encryption Keys. The main content area is titled 'Description' and contains the text: 'Policy to limit instance creation to specific regions and instance types. See https://forums.aws.amazon.com/thread.jspa?threadID=174503 .'. Below the description are four tabs: 'Policy Document' (selected), 'Attached Entities', 'Policy Versions', and 'Access Advisor'. Under the 'Policy Document' tab, there is an 'Edit' button and a code editor showing a JSON policy document. A blue arrow points to the resource list in the policy document. The policy document is as follows:

```
14     ]
15     },
16     {
17         "Effect": "Allow",
18         "Action": "ec2:*",
19         "Resource": [
20             "arn:aws:ec2:eu-west-1:*:*",
21             "arn:aws:ec2:eu-west-1:*:security-group/*"
22         ],
23         "Condition": {
24             "StringLikeIfExists": {
25                 "ec2:InstanceType": [
26                     "t2.micro",
27                     "t2.small",
28                     "t2.medium"
29                 ]
30             }
31         }
32     },
33     {
34         "Effect": "Allow",
```

# PBAC in the wild: Amazon EC2



# PBAC in the wild: Amazon EC2



## Policy Simulator

Amazon EC2    193 Action(s) se...    **Select All**    **Deselect All**    **Reset Contexts**    **Clear Results**    **Run Simulation**

### ▶ Global Settings ⓘ

Action Settings and Results [193 actions selected. 0 actions not simulated. 63 actions allowed. 130 actions denied.]

	Service	Action	Resource Type	Simulation Resource	Permission
▶	Amazon EC2	AcceptVpcPeeringConne...	vpc-peering-conn...	*	<b>denied</b> Implicitly denied (no matc...
▶	Amazon EC2	ActivateLicense	not required	*	<b>denied</b> Implicitly denied (no matc...
▶	Amazon EC2	AllocateAddress	not required	*	<b>allowed</b> 1 matching statements.
▶	Amazon EC2	AssignPrivatelpAddresses	not required	*	<b>denied</b> Implicitly denied (no matc...
▶	Amazon EC2	AssociateAddress	not required	*	<b>allowed</b> 1 matching statements.
▶	Amazon EC2	AssociateDhcpOptions	not required	*	<b>denied</b> Implicitly denied (no matc...
▶	Amazon EC2	AssociateRouteTable	not required	*	<b>denied</b> Implicitly denied (no matc...
▶	Amazon EC2	AttachClassicLinkVpc	instance,security-...	*	<b>denied</b> Implicitly denied (no matc...
▶	Amazon EC2	AttachInternetGateway	not required	*	<b>denied</b> Implicitly denied (no matc...
▶	Amazon EC2	AttachNetworkInterface	not required	*	<b>denied</b> Implicitly denied (no matc...
▶	Amazon EC2	AttachVolume	instance,volume	*	<b>denied</b> Implicitly denied (no matc...

# Advantages of PBAC

- + More modularity: access control logic in 1 place
- + Separation of concerns: policies can be written by non-developer
- + What if rules change?
  - + no updates in application code
  - + updates in a single place
- + Enables your access control policies to easily evolve with your organization
- + Enables centralizing policies, explicitly managing policies across your organization, refining business policies, ...

**Ideally**



# Not all rainbows and unicorns

- Very interesting technology, great vision to work towards
- But, policy-based access control is (still) very hard in practice:
  - Different way of coding
  - Policy languages are not self-explanatory
  - Requires processes for managing policies within your organization
  - Requires supporting tools such as editors and correctness tests
  - Requires interoperability if you want to centralize authorization for multiple applications
  - Your trusted computing base and trust chains grow significantly
  - ...
  - Plus, from my research experience: inherently hard to decouple authorization logic from an application because these rules should still say something about *this* application

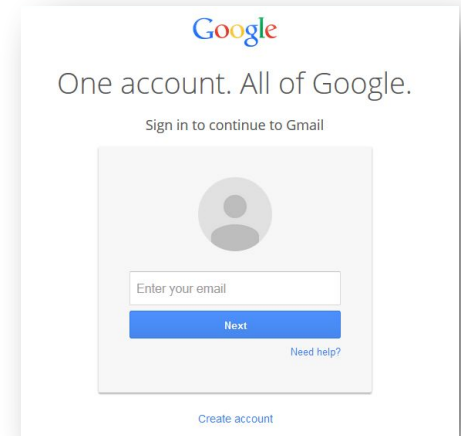
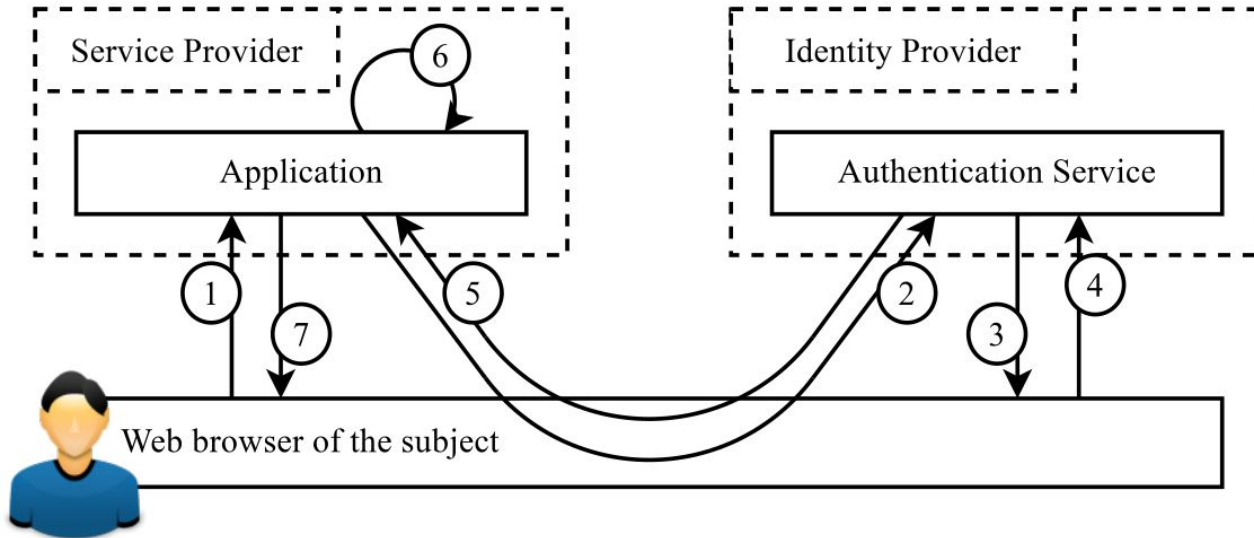
# PBAC: Conclusions

- PBAC:
  - A lot is expected of this technology
  - Enables exciting new stuff
  - But imho currently still too hard to apply in practice
- My recommendation for now:
  - Modularize authorization in your application code (option 2)
    - Provides benefits by itself + future-proof

# Outline

- Introduction
- Positioning access control
- Access control models
- How to enforce access control
- Some important technologies in practice
- Recap and conclusion

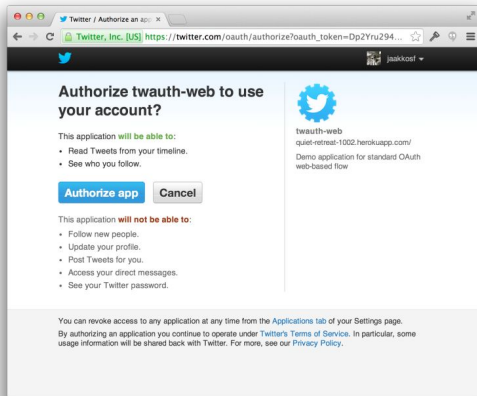
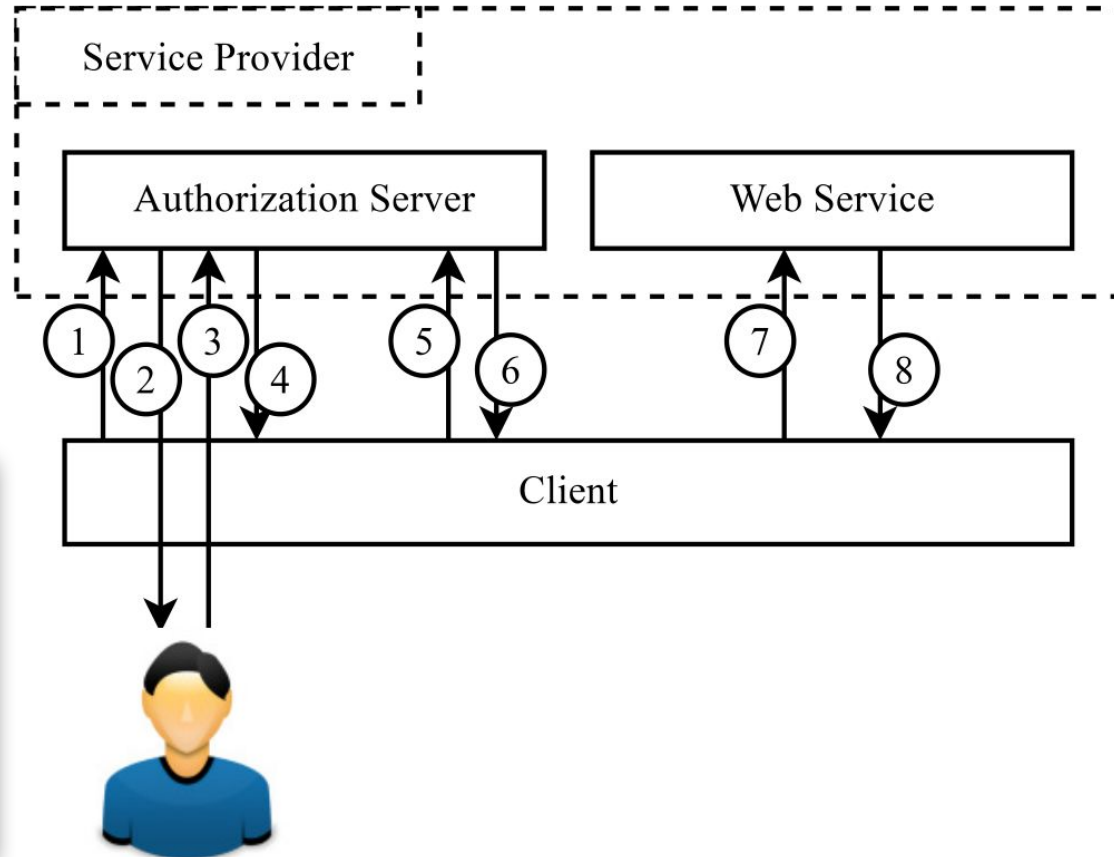
# Federated authentication



# Federated authentication

- Externalizes authorization from a remote application
- Advantages:
  - Lowers the amount of passwords and therefore password reuse
  - Can be used to centralize user mgmt for an organization
  - Removes the need to store passwords in an application
- Standards:
  - OpenID: light-weight, fixed schema, mainly for consumer applications, deprecated
  - SAML: more heavy-weight, extensible, more suitable for enterprise scenarios

# OAuth



# OAuth

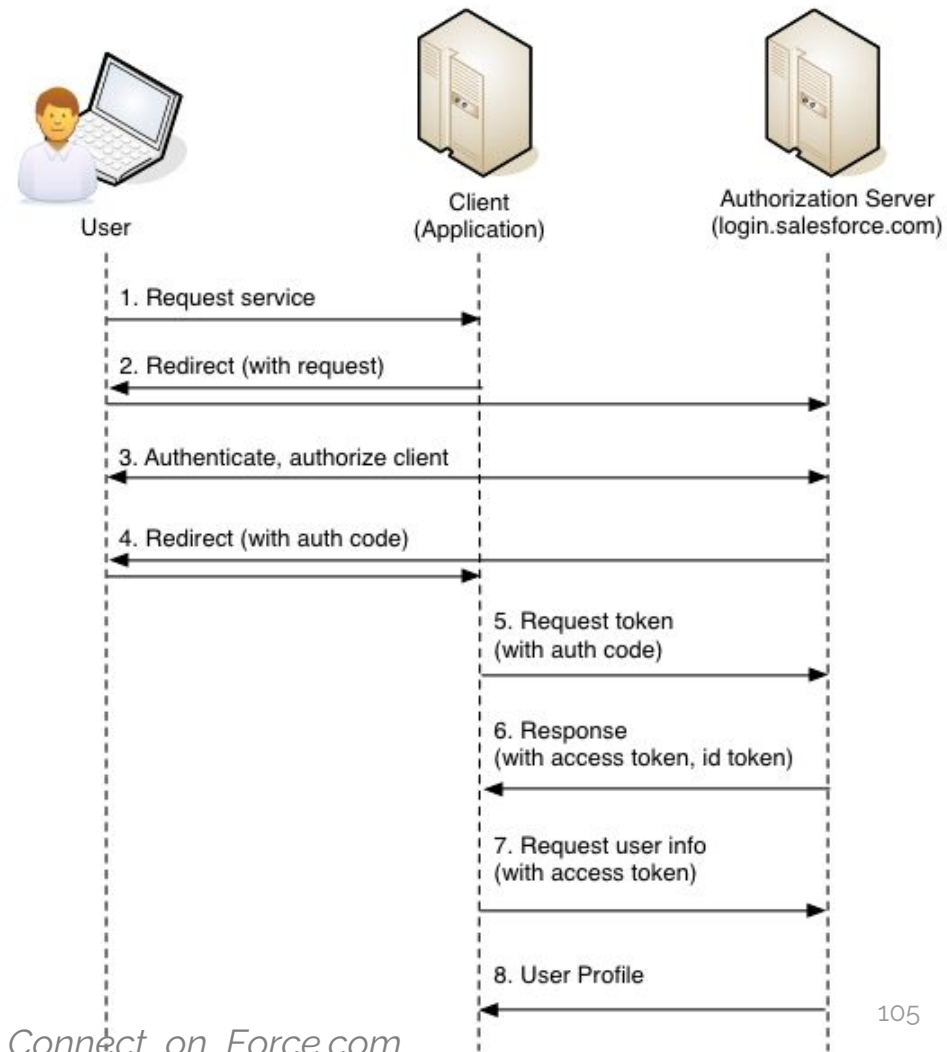
- Constrained delegation of access, mostly to 3<sup>rd</sup> party applications
  - For example, grant a mobile client access to your Twitter stream
  - Also works well with web services and micro-service architectures
- A simplified form of federated authorization
- OAuth 1.0 (2010) was a protocol, OAuth 2.0 (2012) is more a framework
  - Interoperability suffers...

# OpenID Connect

- Identity layer on top of the OAuth 2.0
- Achieves many of the *authentication* features of OpenID, but in a more API-friendly and app-friendly way
  - Get basic user info from AuthZ Server of OAuth, get more details from user mgmt API using the OAuth token
- OpenID is considered deprecated, OpenID Connect (OIDC) is considered the successor



# OpenID Connect



# Outline

- Introduction
- Positioning access control
- Access control models
- How to enforce access control
- Recap and conclusion

# Recap

- Prevent unauthorized access to protected information
  - AAA: authentication, authorization, audit
  - Often domain-specific enforcement and rules
- Different access control models available
  - Who can assign permissions: MAC and/or DAC
  - How permissions are assigned: Identity-based, multi-level, RBAC and ABAC
- How to enforce access control in your application code:
  - Modularize!

# Some final words

- Modern software all depends on access control
- But:
  - Policies are complex to manage in a large organization
    - Choose the minimally complex model for your rules
  - Imperfect because of bugs in the mechanism
    - Make the mechanism as simple as possible
  - Imperfect due to mismatches between policy and mechanism
  - Access control depends on absence of other security bugs
    - Implement least privilege
- After all this, breaches will still occur so prepare and avoid being caught off guard

# References

- [Bogaerts2015] Bogaerts, J., Decat, M., Lagaisse, B., and Joosen, W. Entity-based access control: supporting more expressive access control policies. ACSAC 2015
- [Brewer1989] Brewer, D., and Nash, M. The Chinese Wall security policy. In IEEE Security and Privacy, 1989
- [Cheng2012] Cheng, Y., Park, J., and Sandhu, R. A User-to-User Relationship-Based Access Control Model for Online Social Networks. 2012
- [Crampton2014] Crampton, J., and Sellwood, J. Path Conditions and Principal Matching: A New Approach to Access Control. SACMAT 2014
- [Decat2015] Decat, M., Lagaisse, B., and Joosen, W. Scalable and secure concurrent evaluation of history-based access control policies. ACSAC '15
- [Decat2015b] Decat, M., Bogaerts, J., Lagaisse, B., and Joosen, W. Amusa: Middleware for Efficient Access Control Management of Multi-tenant SaaS Applications. SAC 2015
- [Erlingsson2004] Erlingsson, Úlfar. The inlined reference monitor approach to security policy enforcement. 2004.
- [Fong2011] Fong, P. W. Relationship-based Access Control: Protection Model and Policy Language, CODASPY 2011

# References

- [Graham1972] G. Scott Graham and Peter J. Denning. Protection: principles and practice. AFIPS 1972
- [Harrison1976] Michael A. Harrison, Walter L. Ruzzo, and Jeffrey D. Ullman. Protection in operating systems. 1976
- [Kuhn2010] Kuhn, D. R., Coyne, E. J., and Weil, T. R. Adding attributes to role-based access control. 2010
- [Lampson1971] Lampson, B. W. Protection. ACM SIGOPS Operating Systems Review, 1971
- [Nguyen2012] Park, J., Nguyen, D., and Sandhu, R. A provenance-based access control model. Privacy, Security and Trust (PST) 2012
- [Nguyen2013] Nguyen, D., Park, J., and Sandhu, R. A provenance-based access control model for dynamic separation of duties. Privacy, Security and Trust (PST) 2013
- [NIST2014] Hu, V., Ferraiolo, D., Kuhn, R., Schnitzer, A., Sandlin, K., Miller, R., and Scarfone, K. Guide to Attribute Based Access Control (ABAC) Definition and Considerations. NIST 2014.
- [Park2004] Park, Jaehong, and Ravi Sandhu. The UCON ABC usage control model. ACM Transactions on Information and System Security (TISSEC), 2004

# Accreditation

- Red door: <http://gomighty.com/user/meg/>
- Banking application: <https://kbctouch.kbc.be/>
- Login form:  
<https://w3layouts.com/wp-content/uploads/2014/01/facebook-twitter-google-login.jpg>
- Policy man halt:  
[https://pixabay.com/static/uploads/photo/2012/04/01/18/03/policeman-23796\\_960\\_720.png](https://pixabay.com/static/uploads/photo/2012/04/01/18/03/policeman-23796_960_720.png)
- Policy man traffic fine:  
<http://www.buyautoinsurance.com/wp-content/featured-content/seatbelt/images/traffic-ticket.png>

# Access control

Maarten Decat - SecAppDev 2017

[maarten@elimity.com](mailto:maarten@elimity.com)