

A faint, grayscale watermark of the Purdue University clock tower is centered in the background. The tower has a pointed roof and a clock face. The words "PURDUE UNIVERSITY" are overlaid on the tower in a serif font, with "PURDUE" in a larger size than "UNIVERSITY".

A Privacy-Preserving
Approach to Policy-Based
Content Dissemination

Ning Shang (Microsoft/Purdue)

Mohamed Nabeel (Purdue)

Federica Paci (Trento/Purdue)

Elisa Bertino (Purdue)

Presented by Mohamed Nabeel

ICDE 2010

5th March 2010

What is policy based content dissemination?

A simplified Health Record (HR)

Contact Info




B



Reports

G

Clinical Record



A doctor can access Clinical Record

Medical History

C

A receptionist can access Contact Info

D

A data analyst can access Lab Reports

H

K

F

I

J

L

M

N

Attribute Based Access Control (ABAC)



How to **access control** such content
in a **push based** system?





Use Encryption!

A simplistic approach

PURDUE
UNIVERSITY

A simplified Health Record (HR)

Contact Info

Clinical Record



Lab Reports

A **doctor** can access **Clinical Record**

Medical History

A **data analyst** can access **Lab Reports**

A **receptionist** can access **Contact Info**

A

B

G

C

D

H

K

F

I

J

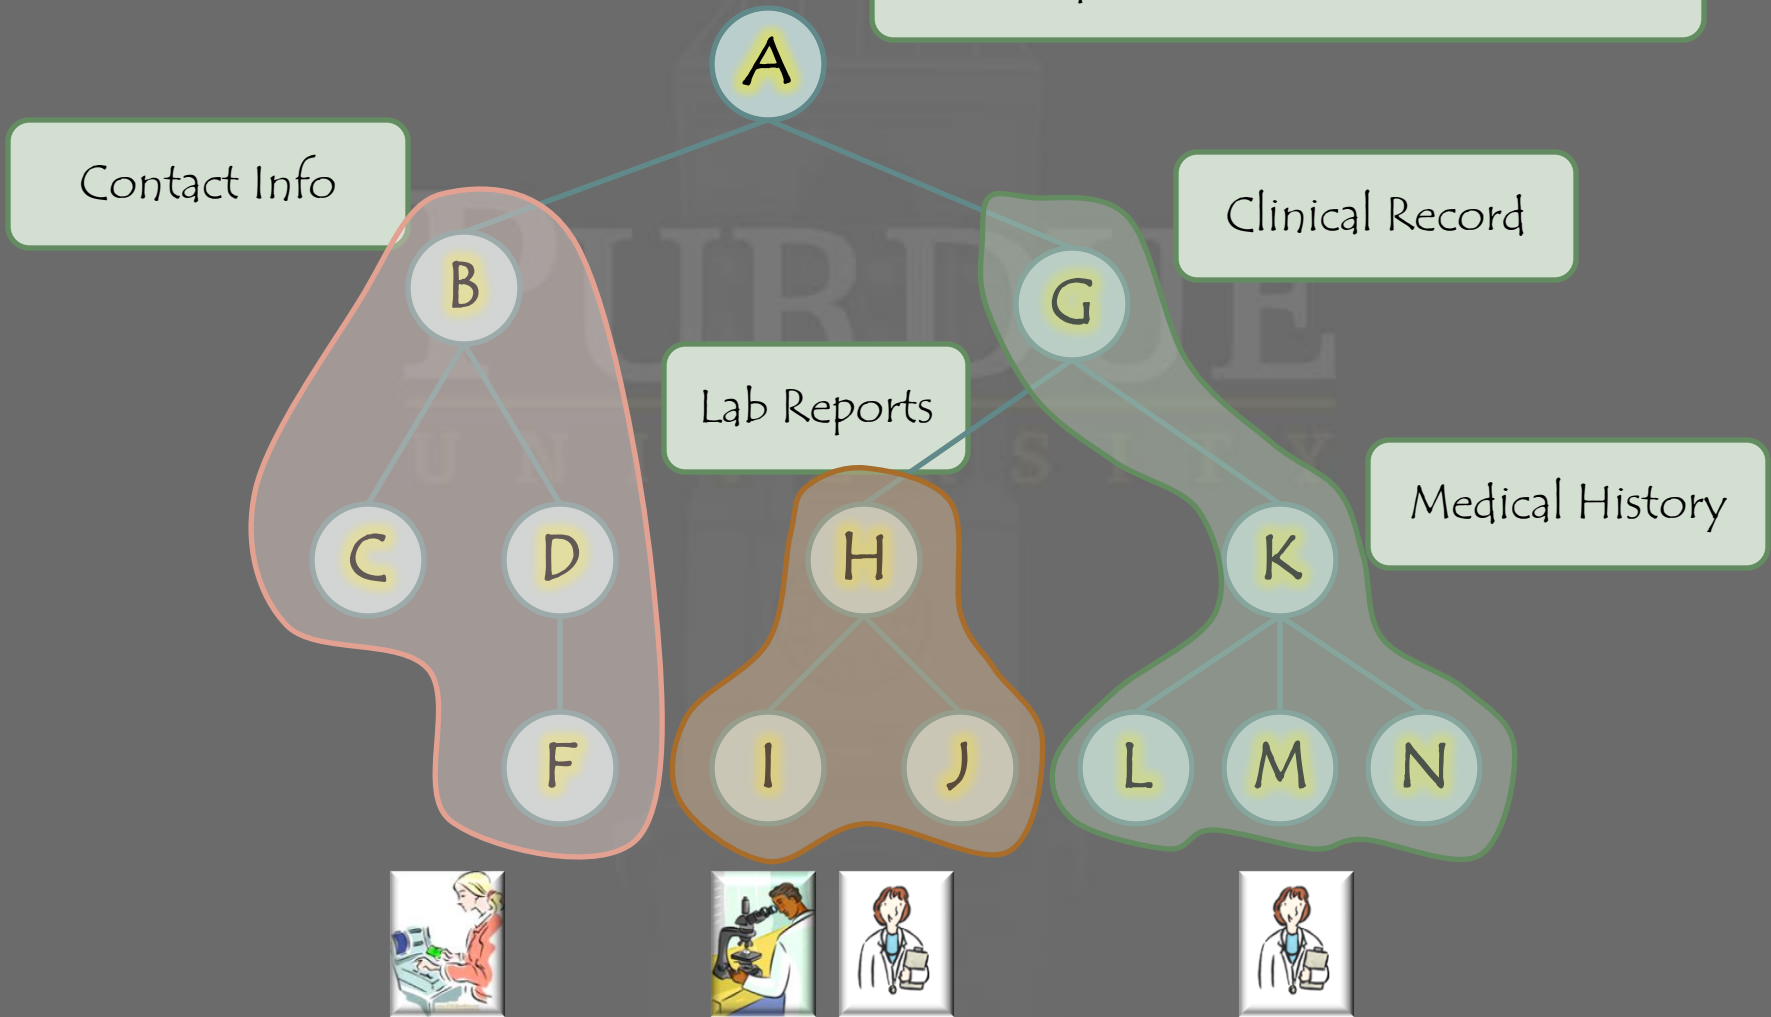
L

M

N

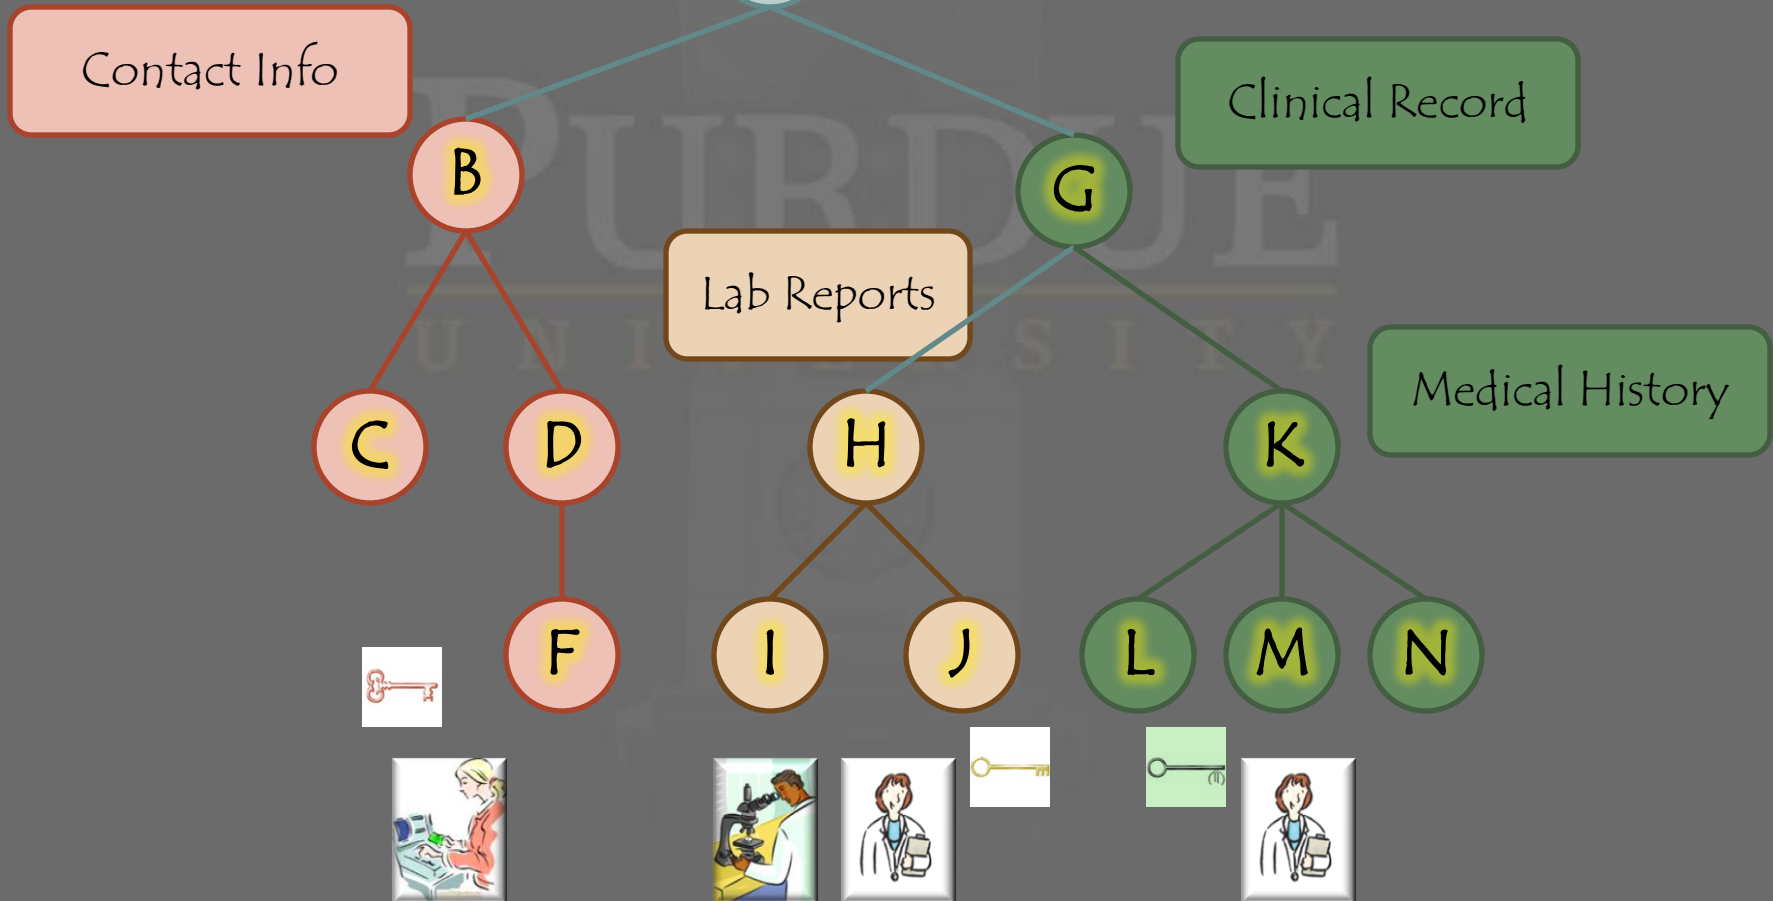


A simplified Health Record (HR)



Each node is encrypted only once!

A simplified Health Record (HR)





It works, but...

Do you see any issues with this approach?



Two issues

Issue #1

User privacy is not preserved



Issue #2

The key management does not scale



We fix these issues




A closer look at the issues



Issue #1: User privacy


PURDUE
UNIVERSITY



Publisher

"I am a **doctor**"

"Here are the **keys** to
decrypt **Medical Records**"



Subscriber

Users have to **reveal** their **attributes**

As you reveal **more credentials**, you
become **identifiable**.

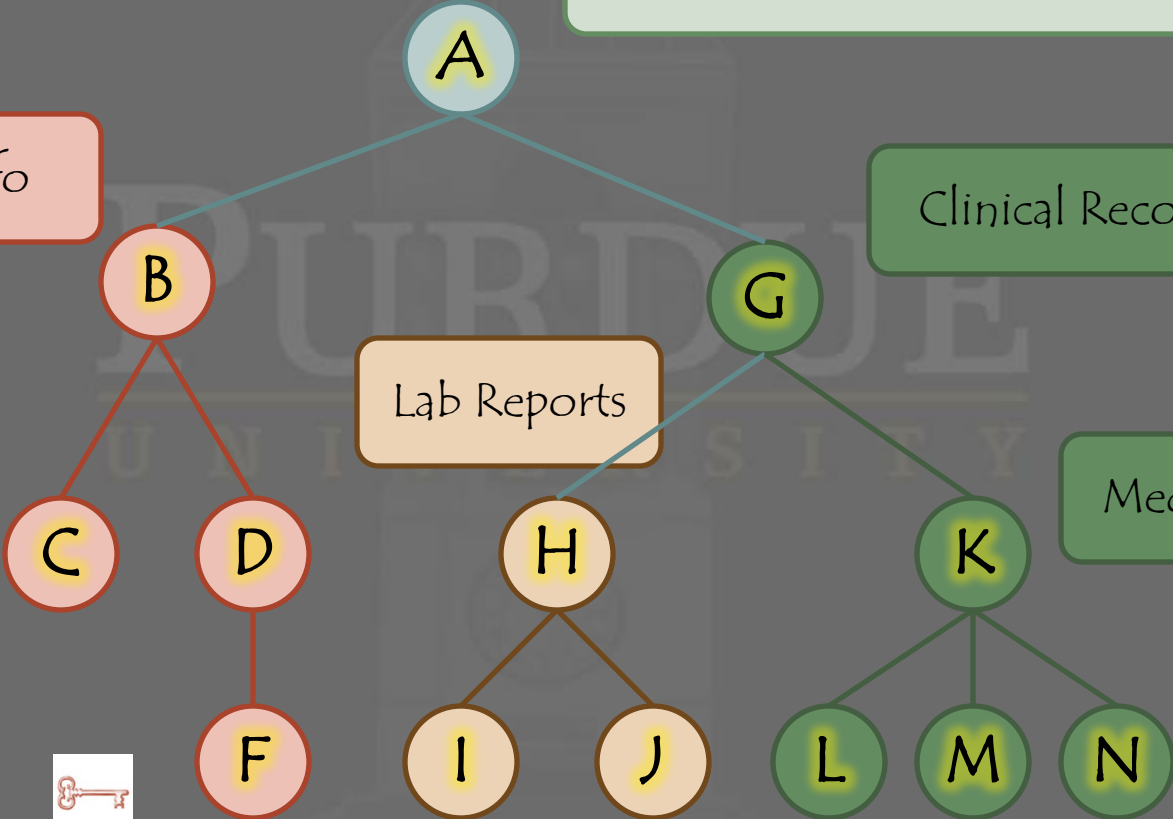
A simplified Health Record (HR)

Contact Info

Clinical Record

Lab Reports

Medical History



Re-keying requires **private communication channels** between the Publisher and each User

Issue #1 is solved if

1. **Credentials** are not revealed to the Publisher.
2. The Publisher does not become aware if a user satisfies a **condition**.
3. A **private communication channel** between the Publisher and each User is not required for re-keying.


Issue #2: key management overhead

Re-keying has $O(n)$ communication
overhead

Issue #2 is solved if multiple users
don't share the same keys

Instead of giving **keys**, give some **secrets** to derive the key using **public information**

We need a **broadcast GKM** scheme



Only subscribers satisfying AC policies are allowed to access

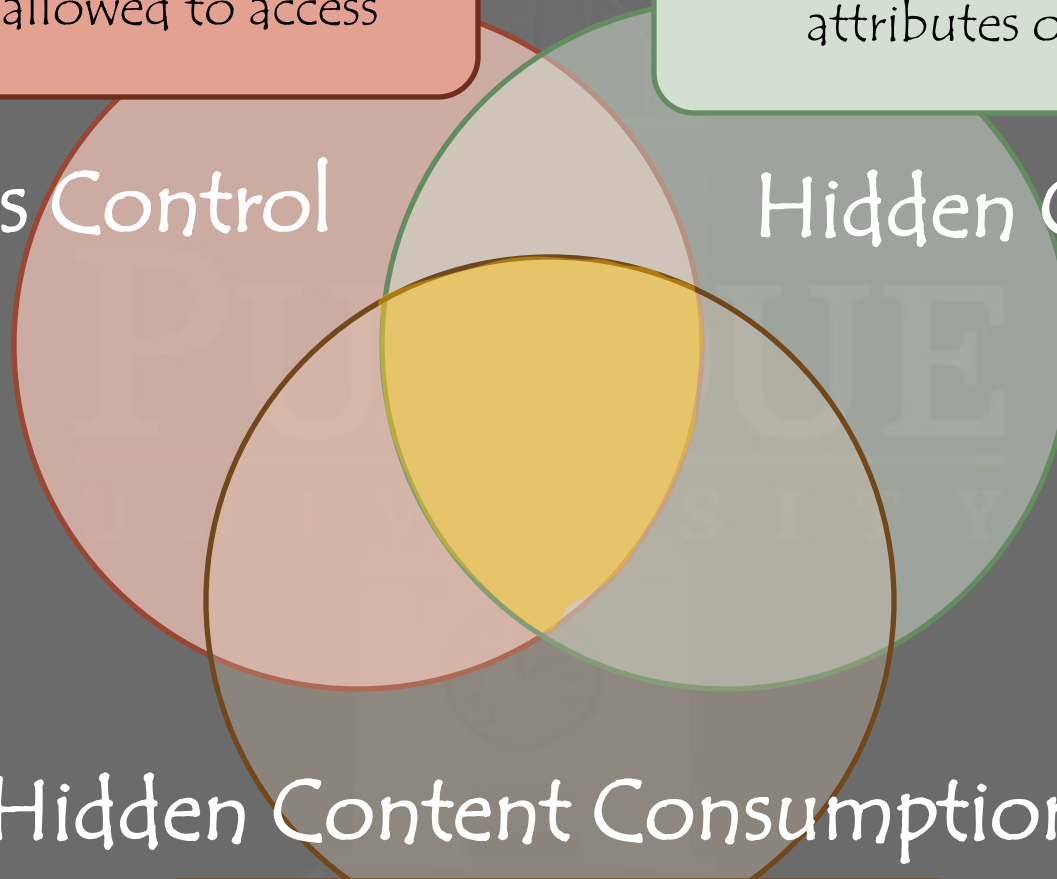
Publisher does not learn the attributes of subscribers

Access Control

Hidden Credentials

Hidden Content Consumption

Publisher does not learn which content subscribers can access



Two building blocks



OCBE (Oblivious Commitment Based Envelope) Protocols*


* J. Li and N. Li, OACerts: Oblivious attribute certificates, 2006

An encrypted message


Unconditionally **hiding** and
computationally **binding**
 $\text{com}(m) = g^m h^r$

Commitment("I am a **doctor**")

Envelope("Here are the **secrets** to
decrypt **Medical Records**")



Publisher



Subscriber

- Publisher does not learn credentials.
- User can open the envelope only if her credential satisfies the condition.

BGKM (Broadcast Group Key Management) Scheme



Public Info


+



S_1



S_2



S_3



$$A = \begin{bmatrix} a_{1,1} & a_{1,2} & a_{1,m} \\ a_{2,1} & a_{2,2} & a_{2,m} \\ \vdots & \vdots & \vdots \\ a_{n,1} & a_{n,2} & a_{n,m} \end{bmatrix}$$

Each row is constructed with secret(s) given to each user



$$B = \begin{bmatrix} b_{1,1} & b_{1,2} & b_{1,m} \\ \vdots & \vdots & \vdots \\ b_{t,1} & b_{t,2} & b_{t,m} \end{bmatrix}$$


The basis of the null space of A

$$m > n \text{ and } t = m - n$$

Publisher constructs an **ACV (Access Control Vector)** hiding the key K and publishes

$$\begin{bmatrix} c_{1,1} & c_{1,2} & c_{1,m} \end{bmatrix} + \begin{bmatrix} K & 0 & 0 \end{bmatrix}$$

A random vector from null space of A



$$\begin{bmatrix} c'_{1,1} & c_{1,2} & c_{1,m} \end{bmatrix}$$

ACV – Access Control Vector

User r can construct at least one row in A (called **KEV – Key Extraction Vector**) using her **secrets** and **public information**

$$\left[\begin{array}{|c|c|} \hline a_{r,1} & a_{r,2} \\ \hline \end{array} \quad \begin{array}{|c|} \hline a_{r,m} \\ \hline \end{array} \right]$$

KEV – Key Extraction Vector

User takes the dot product of ACV
and KEV to obtain the key

ACV – Access Control Vector
KEV – Key Extraction Vector

Putting everything together..



Four Entities

Publisher (Pub)

Subscribers (Sub)

Identity Provider (IdP)

Identity Manager (IdMgr)

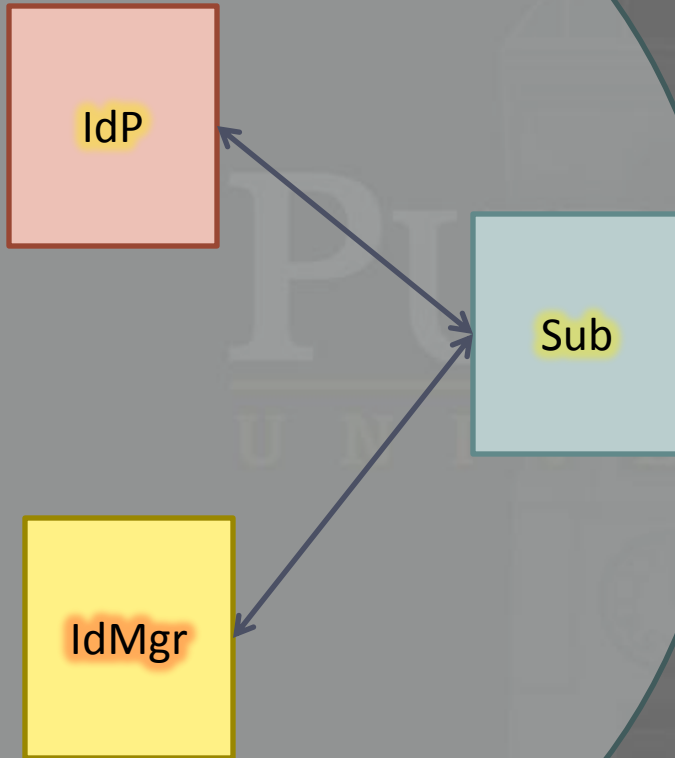
Three Phases

Identity Token Issuance

Identity Token Registration

Document Dissemination

Identity Token Issuance

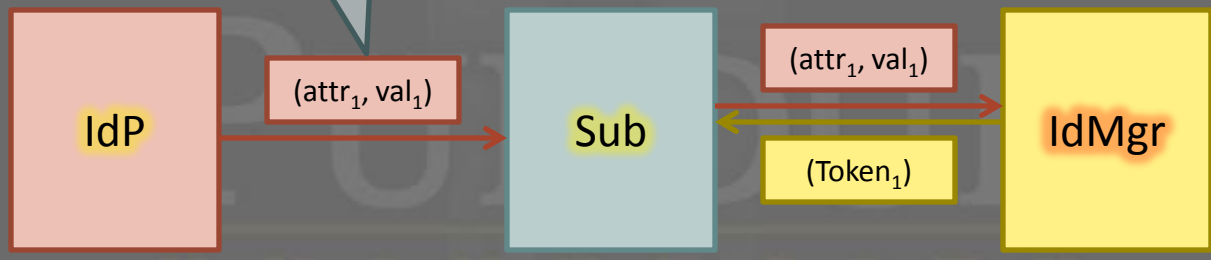


Identity Token Registration
Document Dissemination

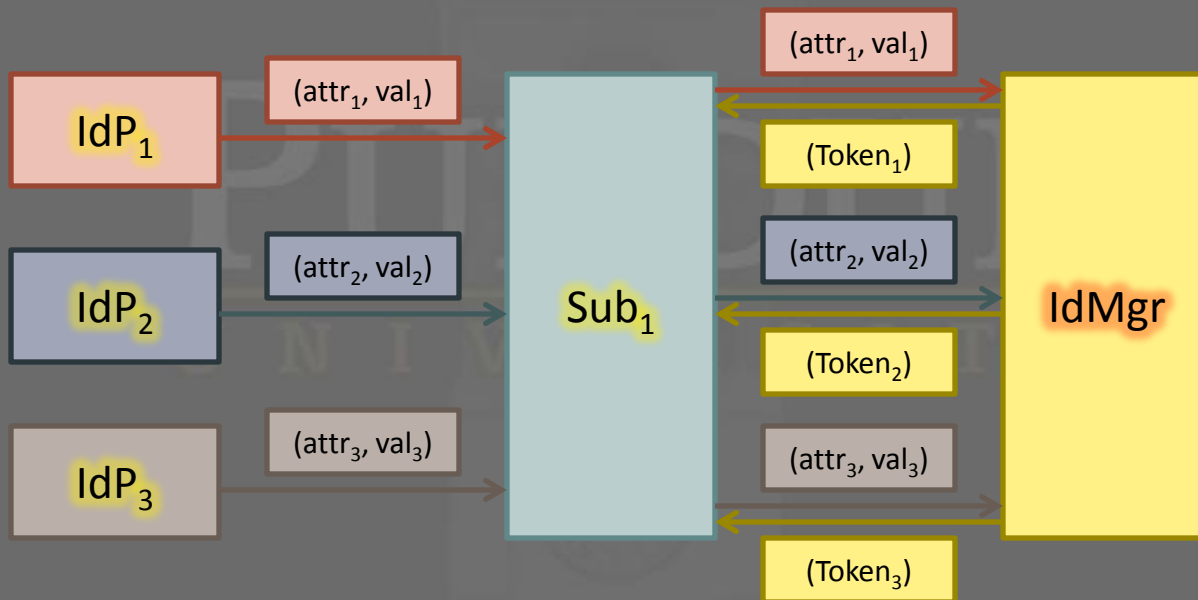
Identity Token Issuance



Certified Identity Attribute
Example: "I am a doctor"



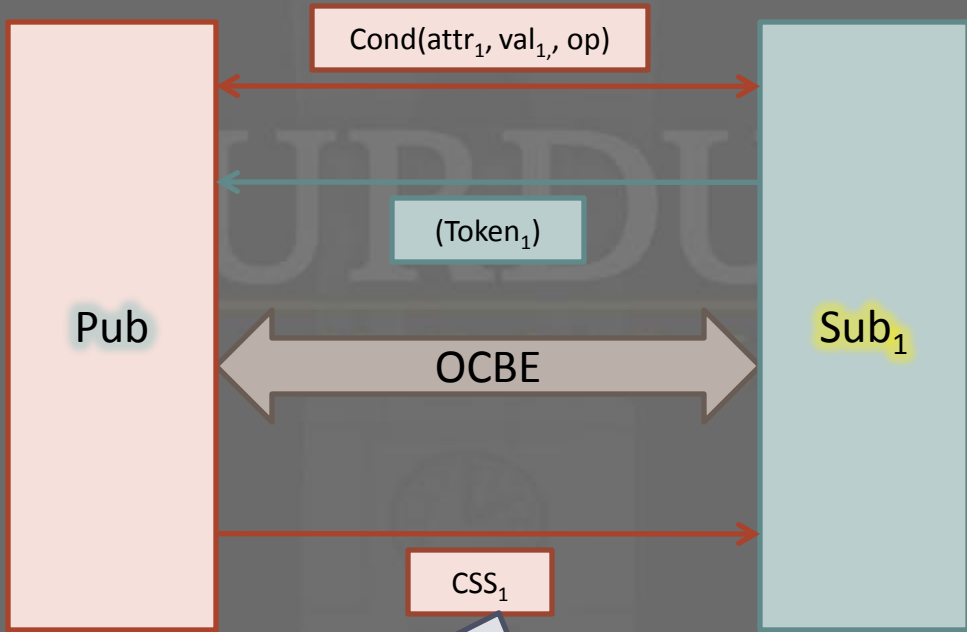
- Identity Token
- Pseudonym
 - Attribute
 - Commitment
 - Signature



Identity Token Registration

PURDUE
UNIVERSITY

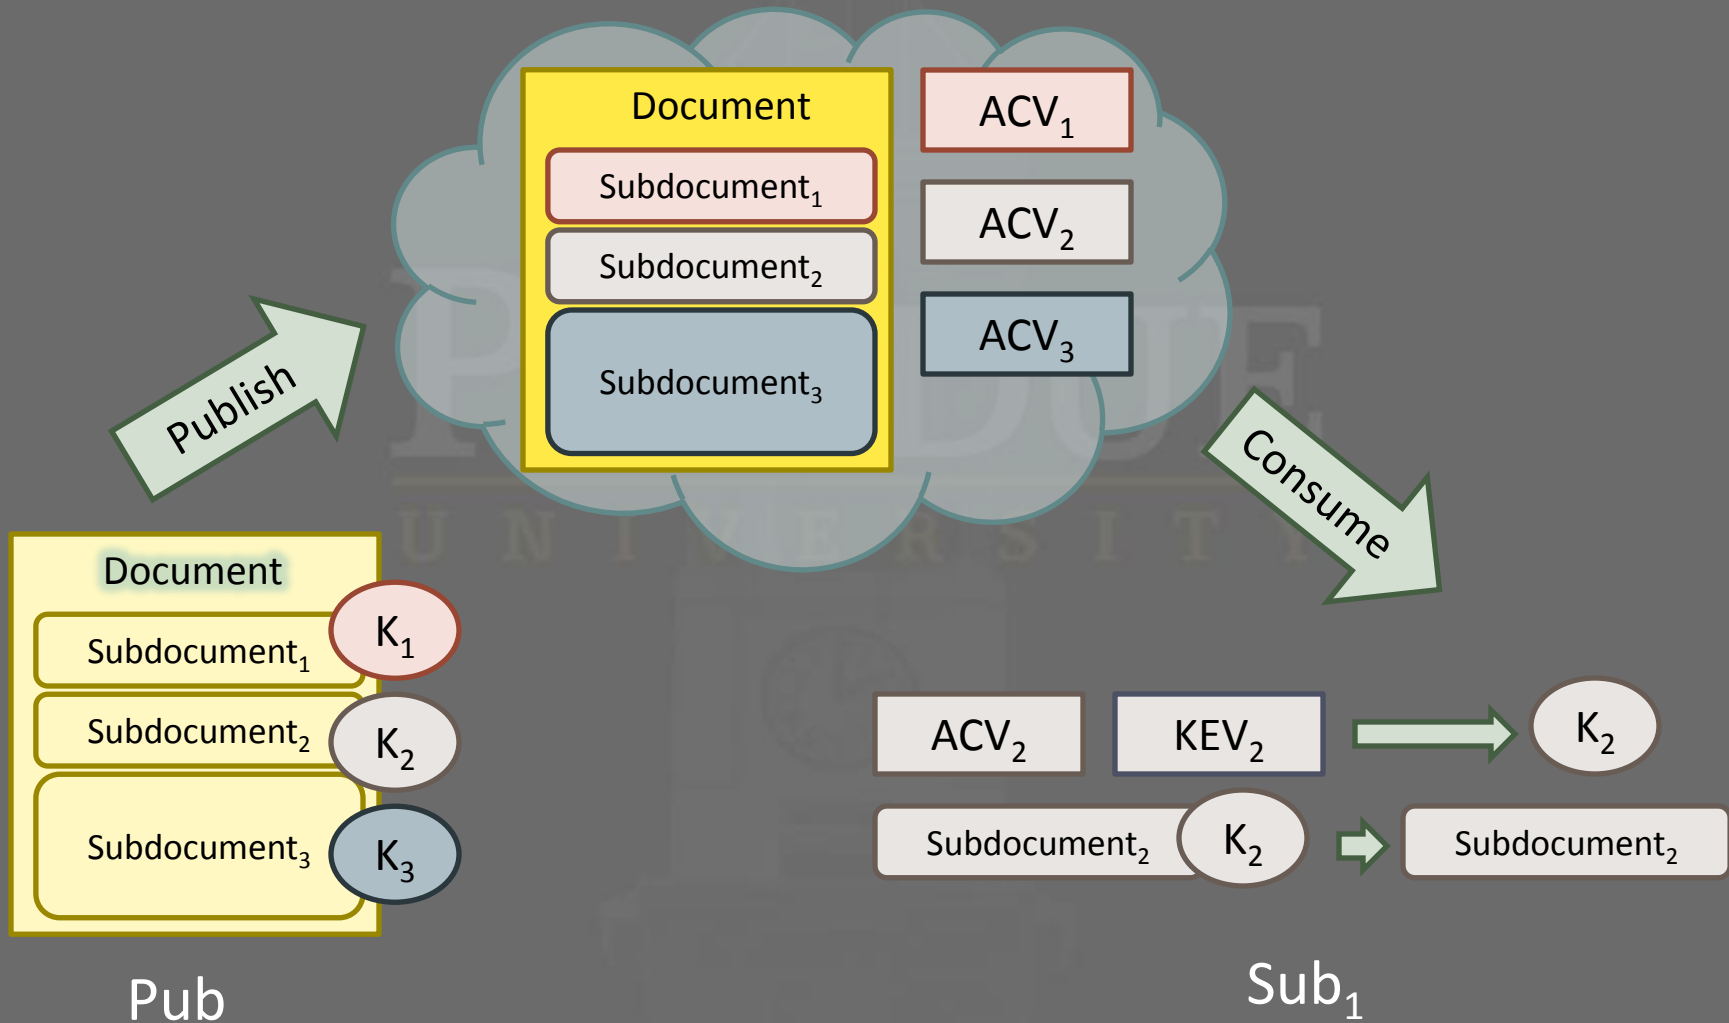
Example: age > 21



Conditional Subscription Secret

Document Dissemination

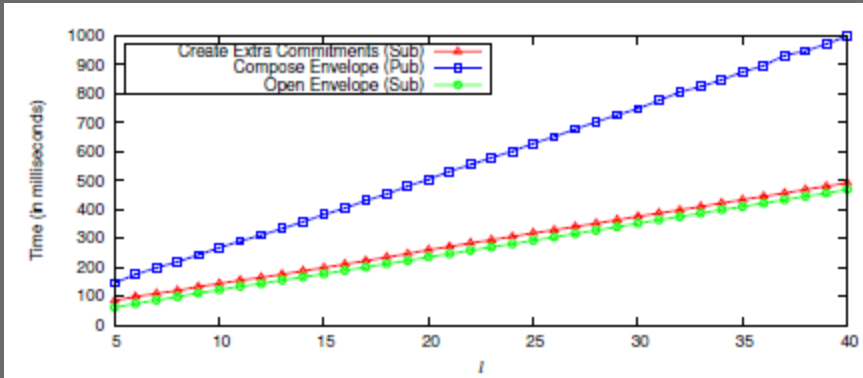





ACV – Access Control Vector
 KEV – Key Extraction Vector

How did we fare?






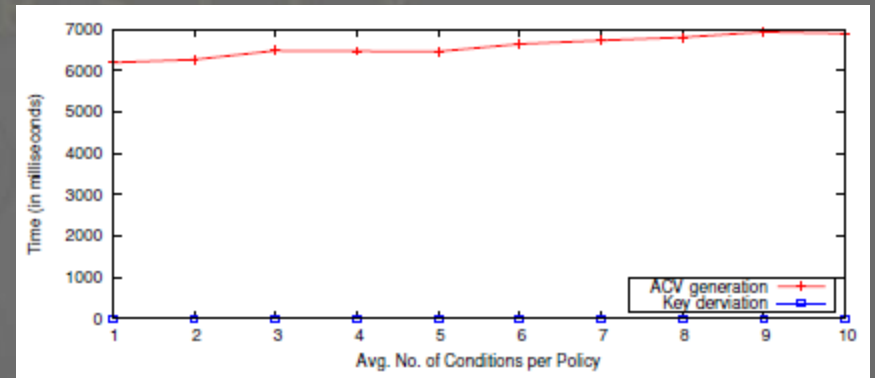
Average computation time for running one round of GE-OCBE protocol



Time to generate an ACV for different user configurations



Key derivation time for different user configurations



ACV generation and key derivation for different number of conditions per policy

Access Control

Hidden Credentials



Hidden Content Consumption

OCBE – Oblivious Commitment Based Envelope
BKGM – Broadcast Group Key Management