

# Access Control for Database Applications: Beyond Policy Enforcement

**Wen Zhang**

UC Berkeley

Aurojit Panda

NYU

Scott Shenker

UC Berkeley & ICSI

**Berkeley** | **EECS**  
ELECTRICAL ENGINEERING & COMPUTER SCIENCES

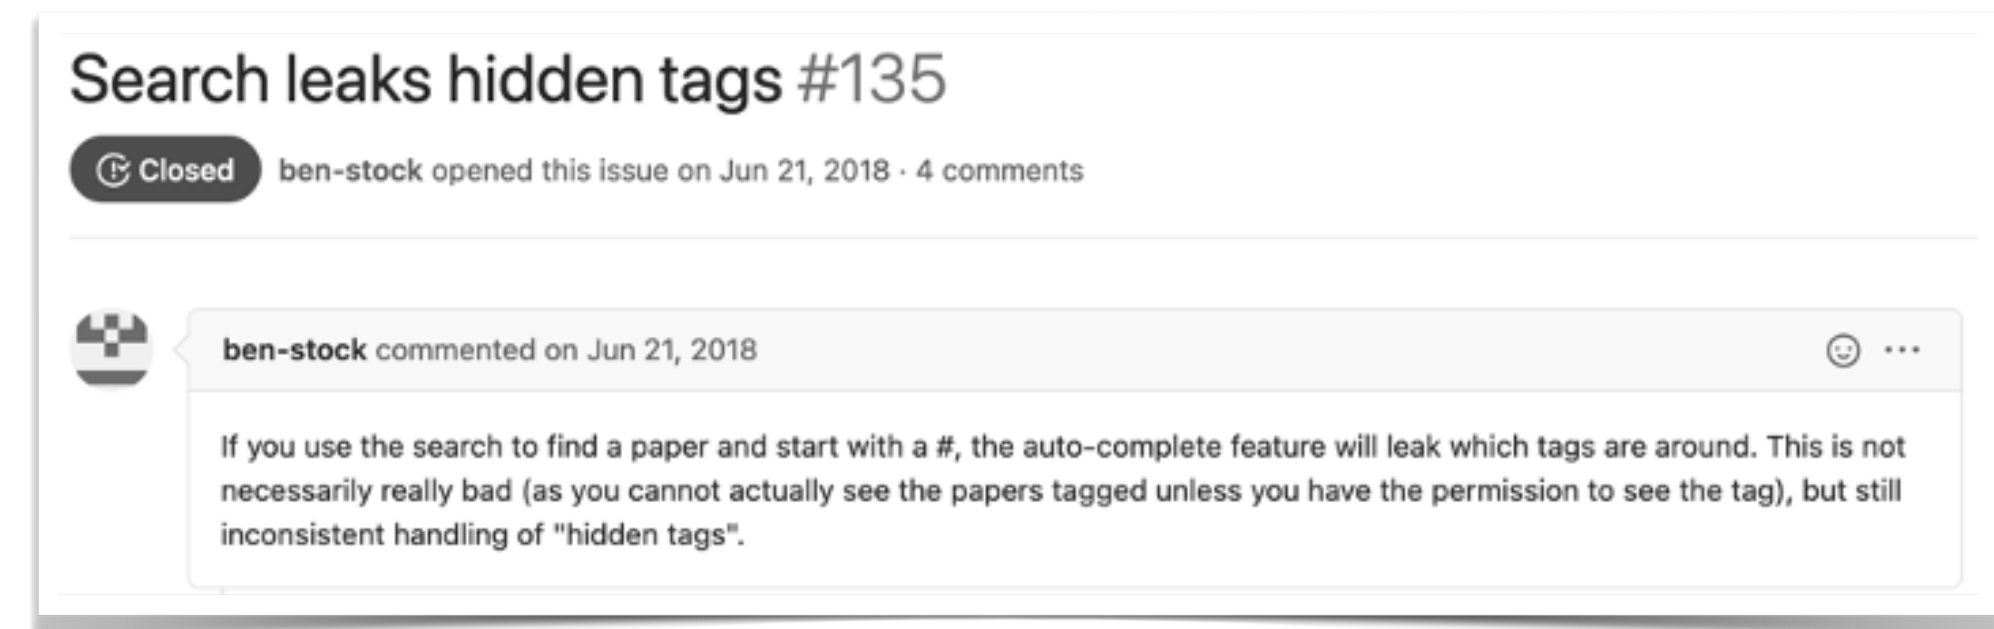
 **NETSYS**

# Apps must take care when revealing sensitive data

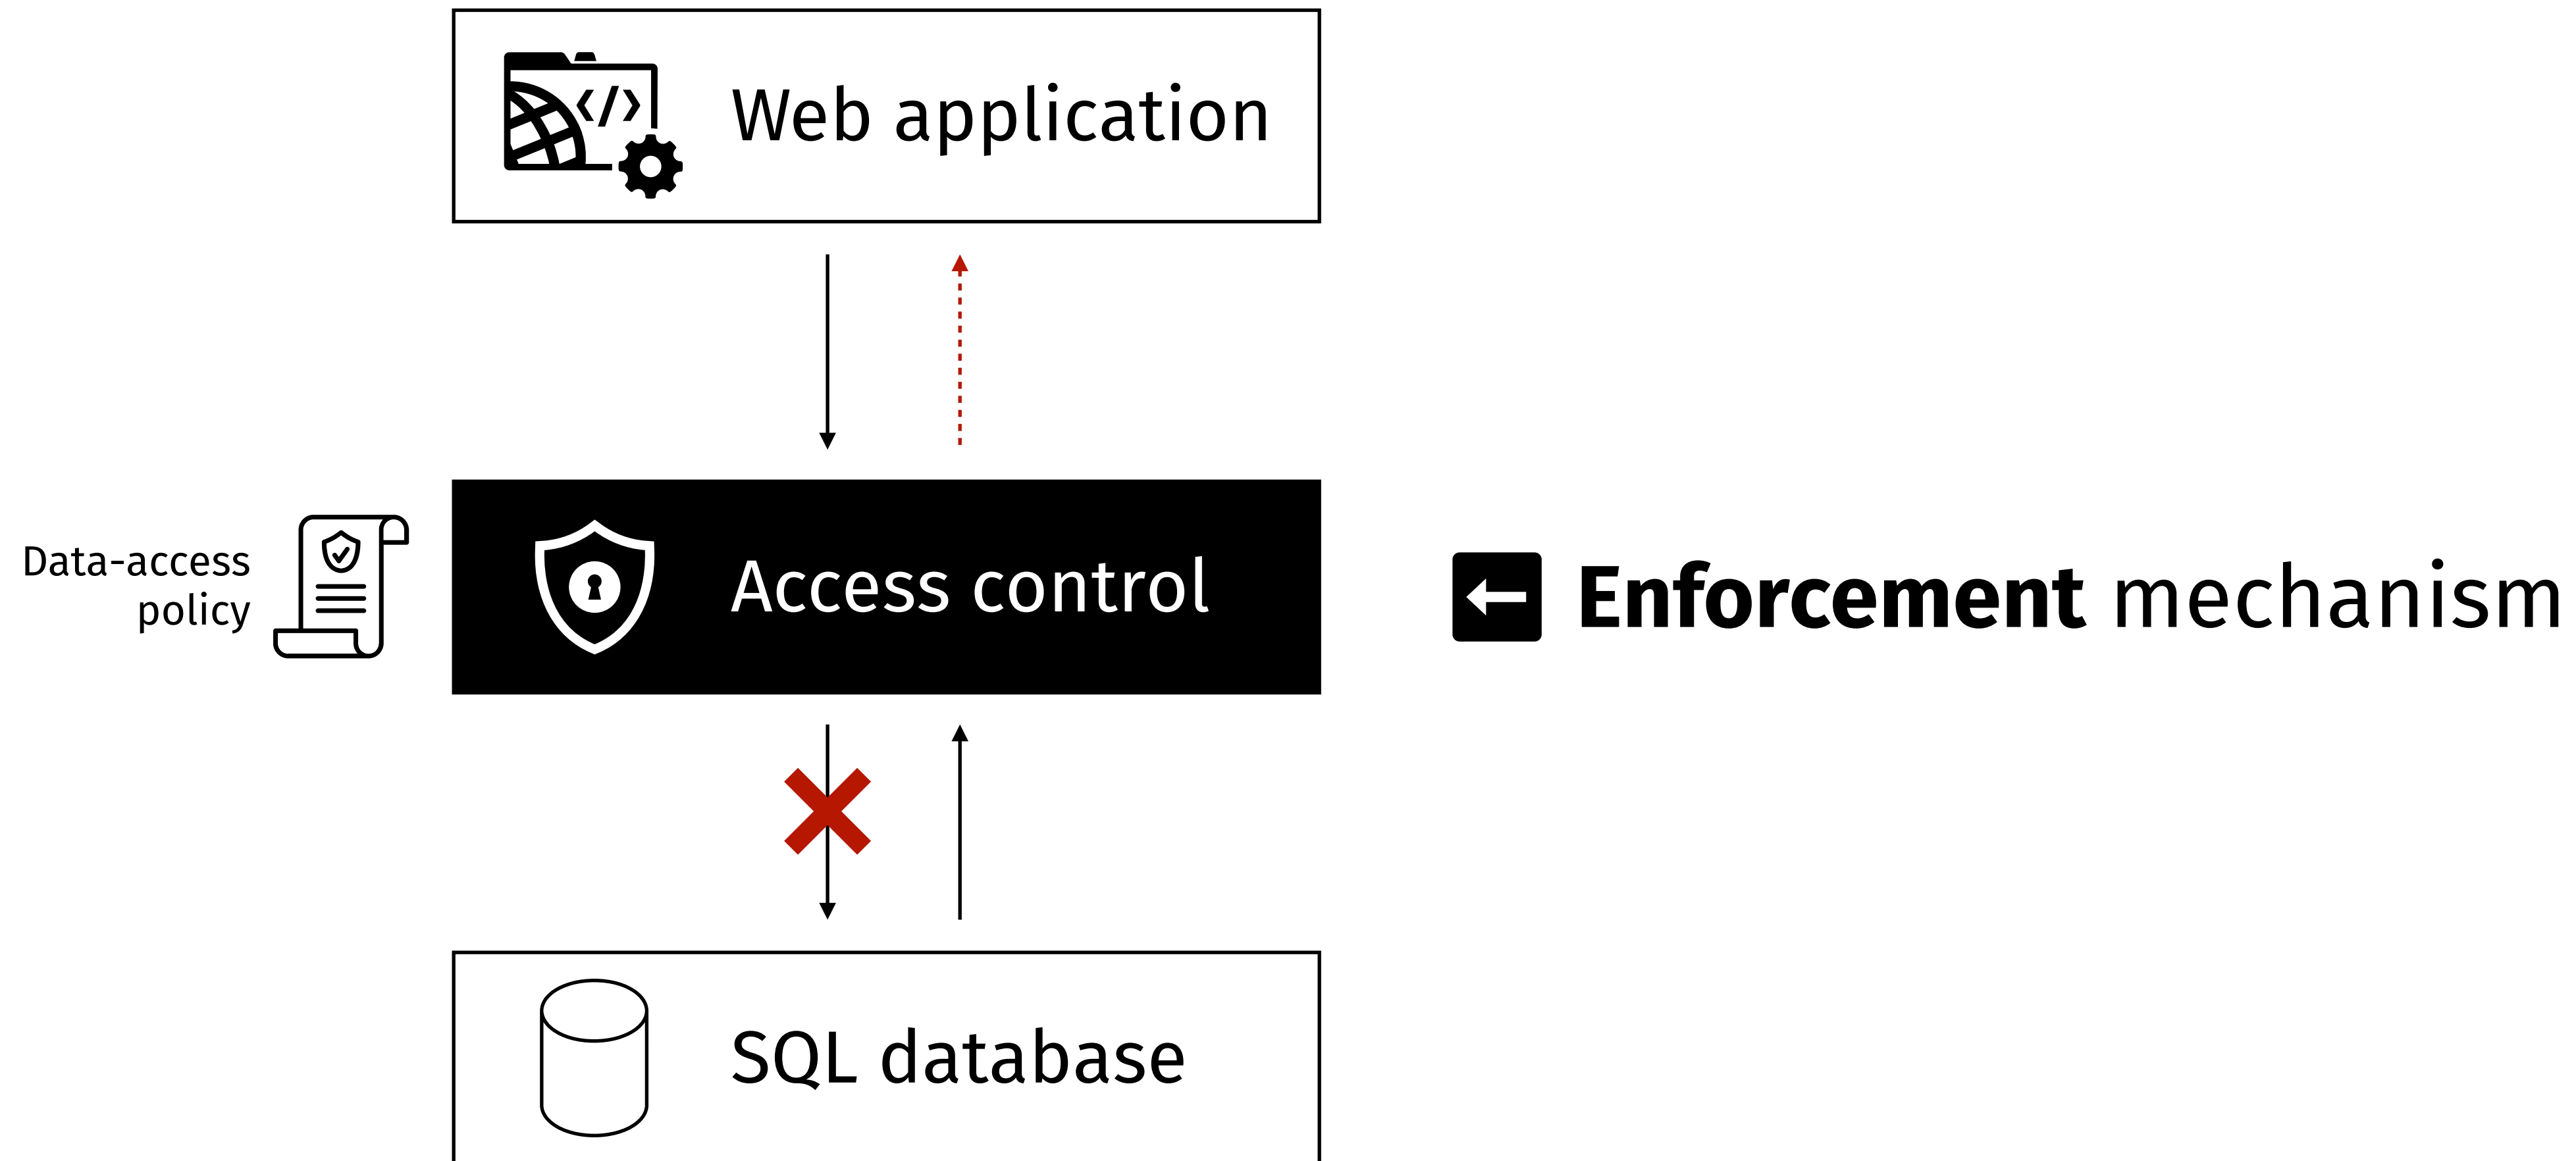
To ensure sensitive data is revealed only to authorized parties, developers insert **permission checks** + **query filters** in their code.

Permission checks and query filters are **easy to miss** or **get wrong**.

When this happens, **sensitive data is leaked** to unauthorized parties.



# Access control to the rescue



# Enforcement: Focus of access-control literature

## Limiting Disclosure in Hippocratic Databases VLDB '04

Kristen LeFevre<sup>†\*</sup> Rakesh Agrawal<sup>†</sup> Vuk Ercegovic<sup>\*</sup> Raghu Ramakrishnan<sup>\*</sup>  
Yirong Xu<sup>†</sup> David DeWitt<sup>\*</sup>

<sup>†</sup>IBM Almaden Research Center, San Jose, CA 95120

<sup>\*</sup>University of Wisconsin, Madison, WI 53706

## Qapla: Policy compliance for database-backed systems USENIX Sec '17

Aastha Mehta<sup>1</sup>, Eslam Elnikety<sup>1</sup>, Katura Harvey<sup>1,2</sup>, Deepak Garg<sup>1</sup>, and Peter Druschel<sup>1</sup>

<sup>1</sup>Max Planck Institute for Software Systems (MPI-SWS), Saarland Informatics Campus

<sup>2</sup>University of Maryland, College Park

## Static Checking of Dynamically-Varying Security Policies in Database-Backed Applications OSDI '10

Adam Chlipala  
*Impredicative LLC*

## Towards Multiverse Databases HotOS '19

Alana Marzoev Lara Timbó Araújo<sup>†</sup> Malte Schwarzkopf Samyukta Yagati  
Eddie Kohler<sup>‡</sup> Robert Morris M. Frans Kaashoek Sam Madden  
*MIT CSAIL* <sup>†</sup>*MIT CSAIL and Airbnb* <sup>‡</sup>*Harvard University*

## STORM: Refinement Types for Secure Web Applications OSDI '21

Nico Lehmann Rose Kunkel Jordan Brown Jean Yang  
*UC San Diego* *UC San Diego* *Independent* *Akita Software*  
Niki Vazou Nadia Polikarpova Deian Stefan Ranjit Jhala  
*IMDEA Software Institute* *UC San Diego* *UC San Diego* *UC San Diego*

## Precise, Dynamic Information Flow for Database-Backed Applications PLDI '16

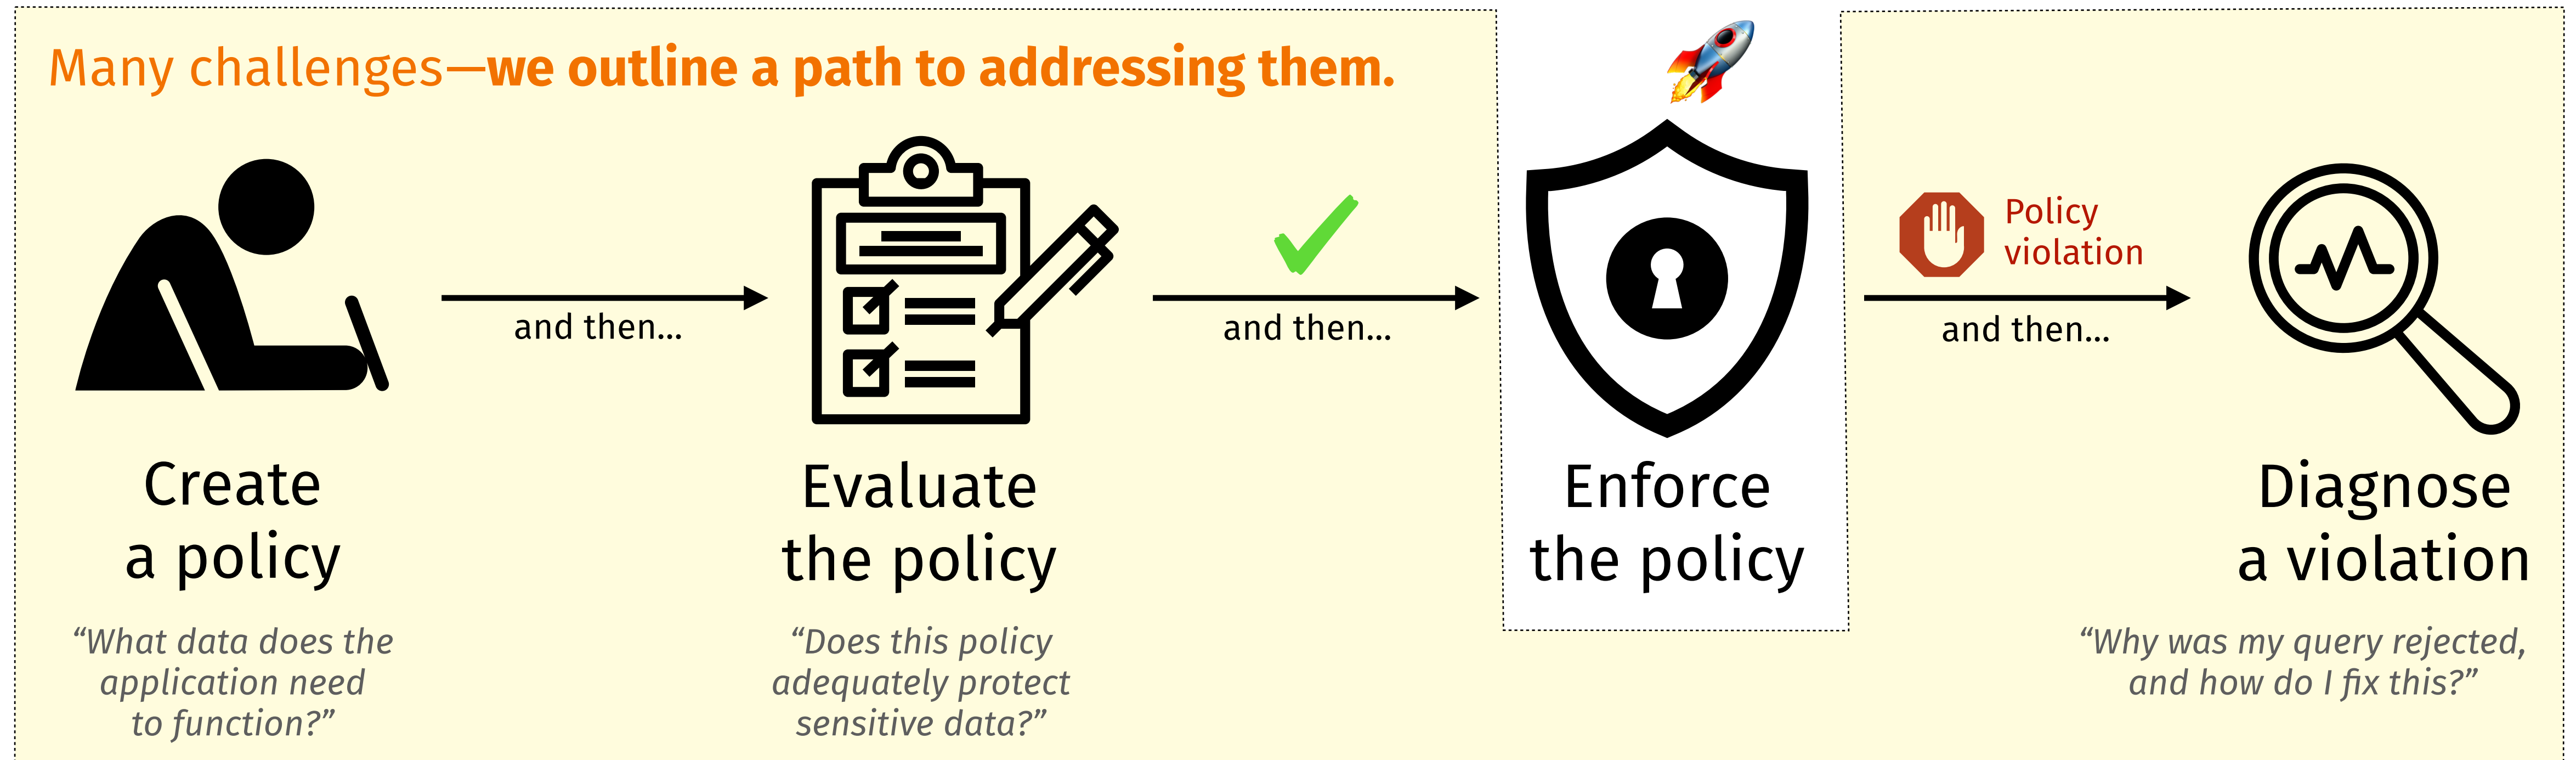
Jean Yang Carnegie Mellon University and Harvard Medical School, USA	Travis Hance Dropbox, USA	Thomas H. Austin San Jose State University, USA
Armando Solar-Lezama Massachusetts Institute of Technology, USA	Cormac Flanagan University of California, Santa Cruz, USA	Stephen Chong Harvard University, USA

## Blockaid: Data Access Policy Enforcement for Web Applications OSDI '22

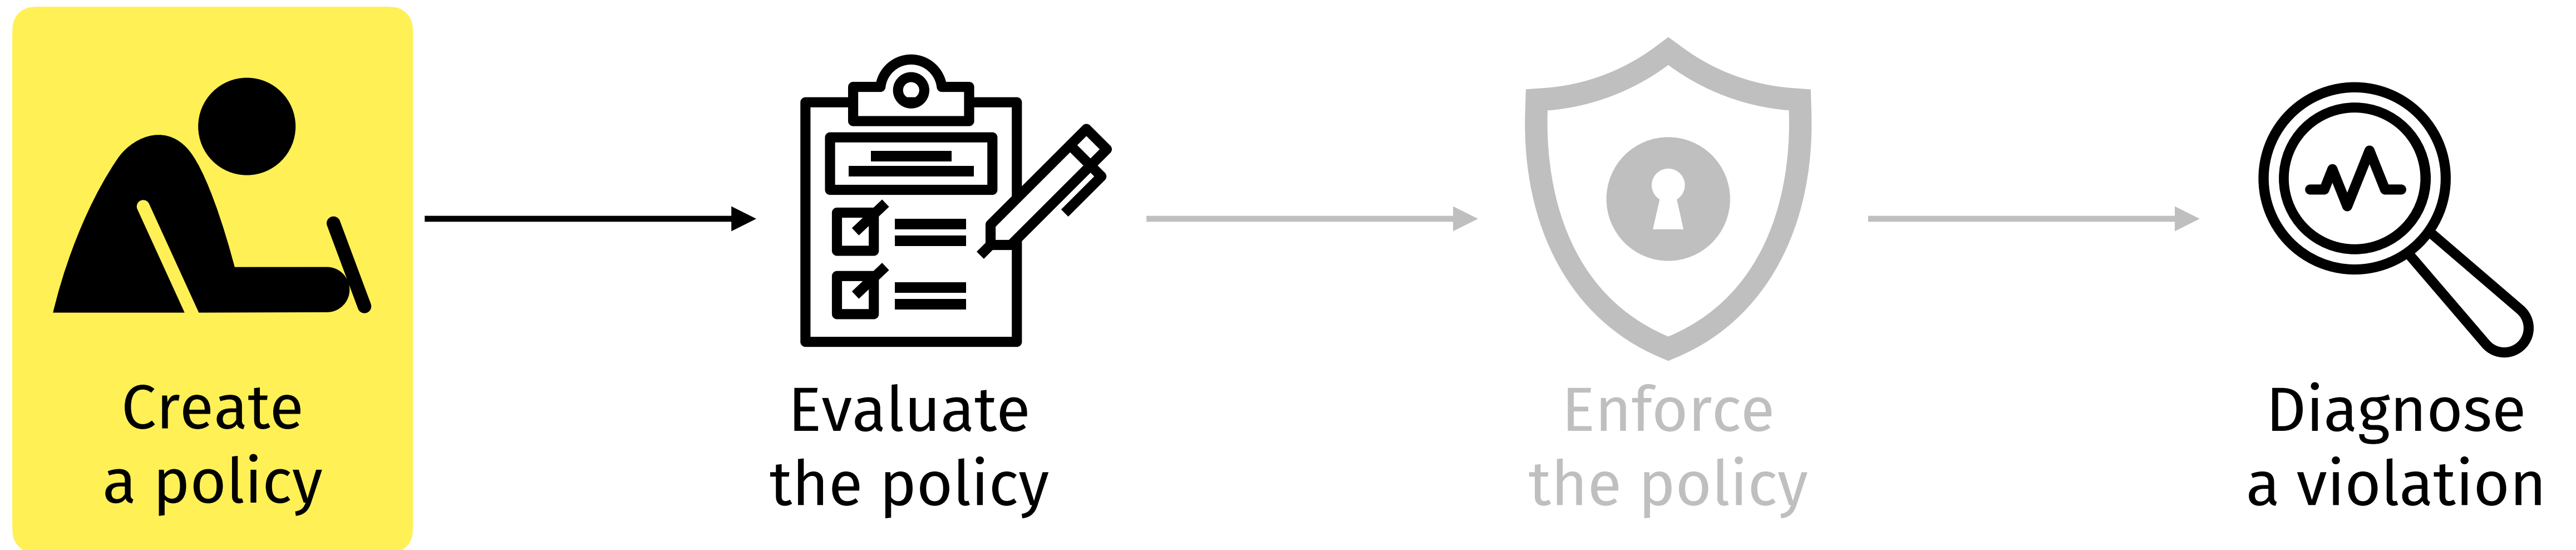
Wen Zhang<sup>1</sup> Eric Sheng<sup>2,\*</sup> Michael Chang<sup>1</sup> Aurojit Panda<sup>3</sup> Mooly Sagiv<sup>4</sup> Scott Shenker<sup>1,5</sup>  
<sup>1</sup>*UC Berkeley* <sup>2</sup>*Yugabyte* <sup>3</sup>*NYU* <sup>4</sup>*Tel Aviv University* <sup>5</sup>*ICSI*

# The full life-cycle of access control

You're a web admin. You want to **deploy access control** on your web app.



# Challenges in the access-control life-cycle



# Writing a policy **from scratch** is challenging

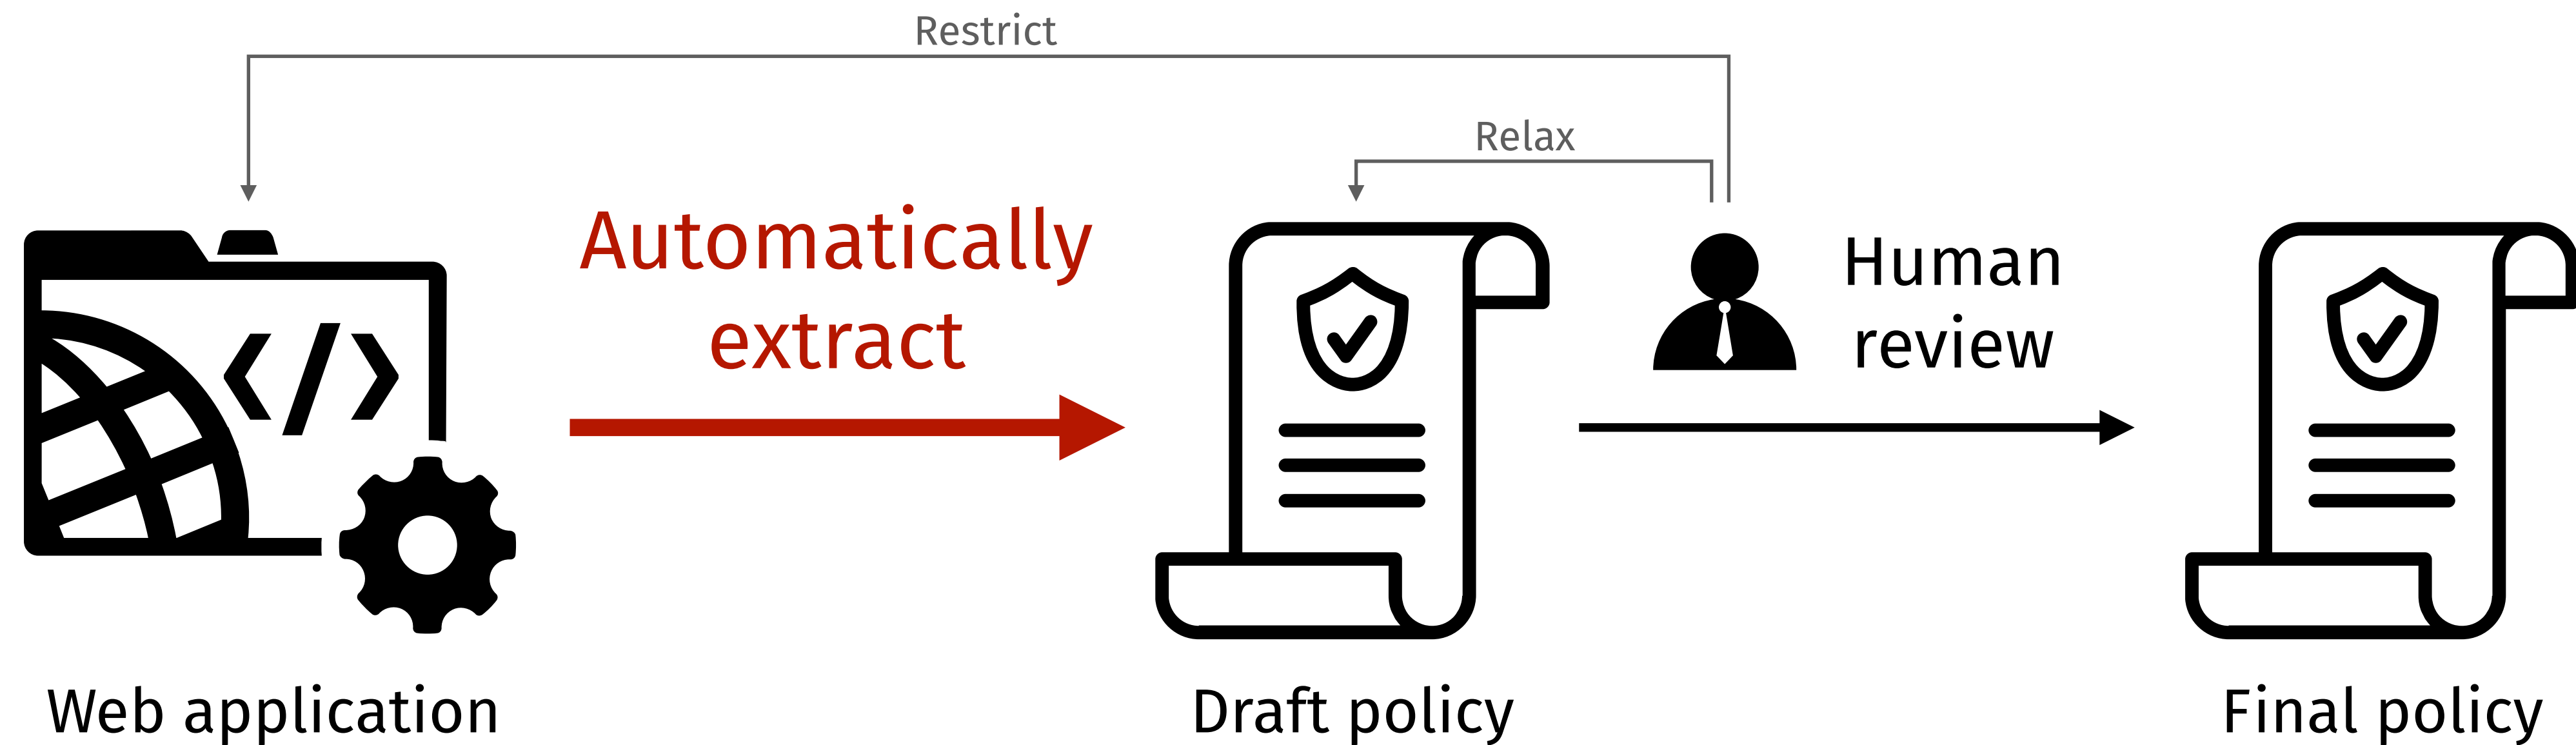
- You're an admin, wanting to **write a data-access policy** for your application.
  1. You map out and understand all the intended data accesses...
  2. And write down a good policy.
- This is **tedious** and **error-prone**.
  - As you wade through thousands upon thousands of line of code...
  - You can easily miss an **edge case**, or make a **typo** in the policy.

# Can we make policy creation **easier**?

- Writing a policy inherently **requires human input**.
  - Balance application needs vs social, regulatory requirements.
- But we should **automate** as much as possible.
- Assumption:
  - Compared to writing a policy from scratch...
  - Reviewing and refining a **draft policy** is much easier.




# Proposal: Policy extraction



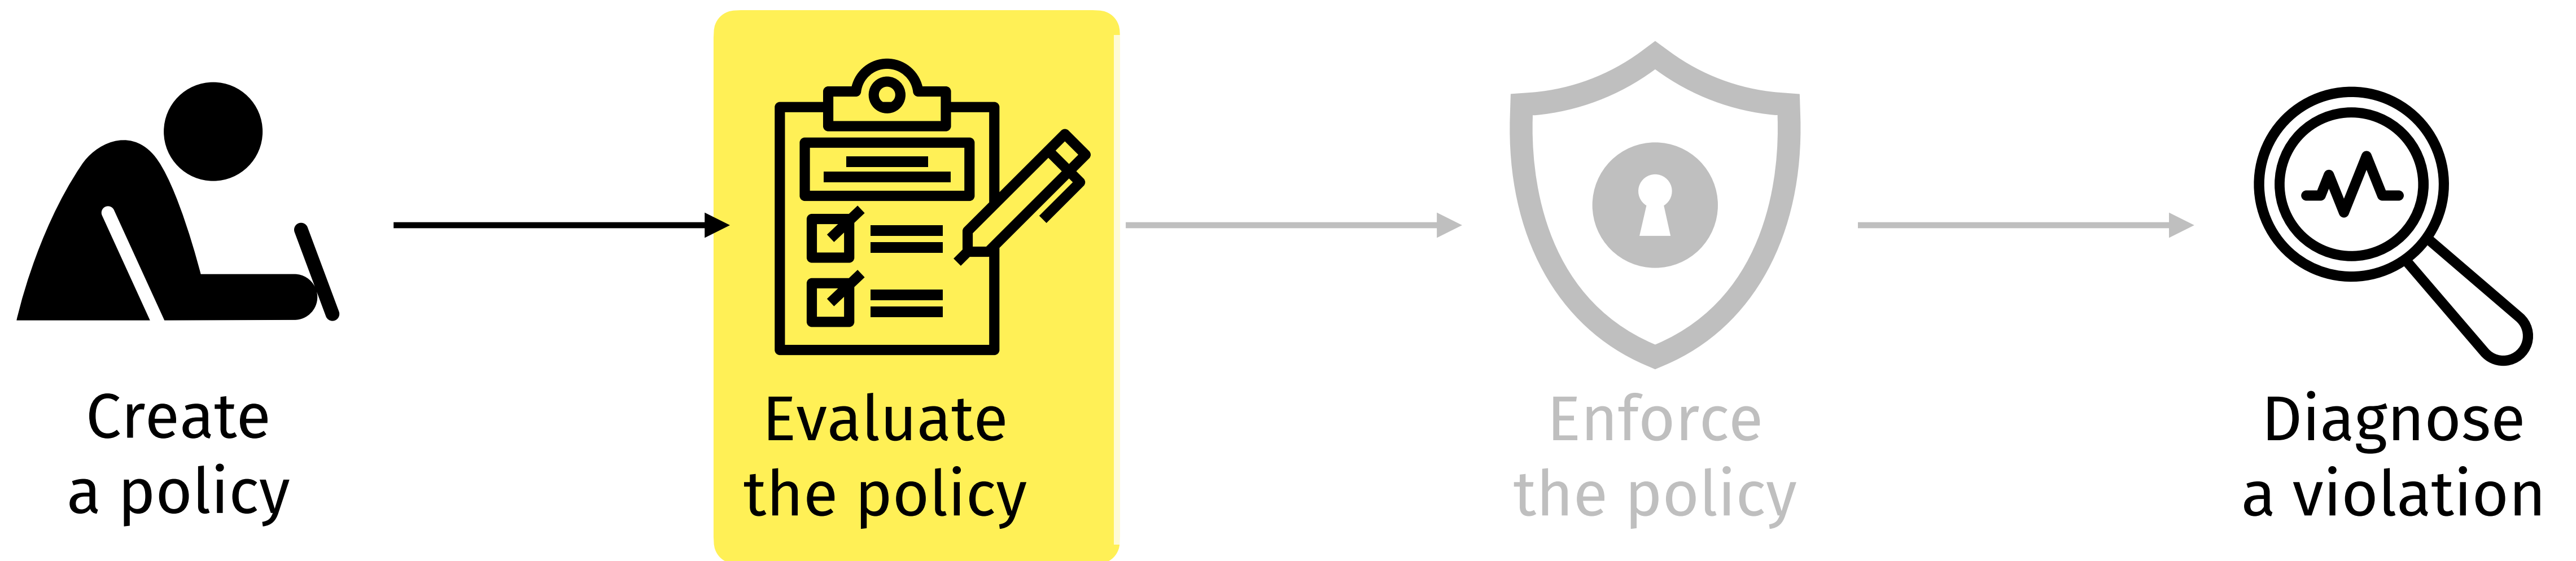
Ideally: Strictest policy that allows all possible access.

# How to extract draft policy?



- Language-based: **Symbolic execution**.
  - Explore program paths ➡ Gather (symbolic) queries ➡ Construct policy.
  - **Challenge:** Web languages & frameworks are highly dynamic.
- Language-agnostic: (Run-time) **Specification mining**.
  - Run the application ➡ Observe (concrete) queries ➡ Construct policy.
  - **Challenge:** Must run application on comprehensive input suite.
  - **Challenge:** Must **generalize** concrete queries into policy.
    - Doctor #1 accesses Patient #10 ➡ Doctors can access patients they treat.

# Challenges in the access-control life-cycle



# Does a policy adequately protect sensitive data?

- A policy must strike a **balance** between:
  - Application's need to access data,
  - Admin's need to protect sensitive information.
- **Where does a given policy fall in the balance?**
  - Given a policy for **allowed queries**, and a **sensitive query**...
  - Does the policy **disclose too much** about sensitive query's output?
- To discuss this, we need a **metric** for disclosure.

Possible metric: Reject the sensitive query.

## Reject sensitive query $\neq$ Prevent disclosure

- In a medical-records management system, the policy allows analyst to view:
  1. The doctor assigned to each patient.
  2. The diseases treated by each doctor.
- **Sensitive query**: What disease is patient John being treated for?
  - Query is rejected under the policy...
  - But **significant disclosure** is possible from answers to allowed queries.
  - (E.g., if John is being treated by a doctor who treats only two diseases.)

# State-of-the-art metric: Bayesian privacy

Given an adversary's **prior belief** over what the sensitive data might be...

Compute a **posterior belief** after seeing the allowed data.

Sensitive data **disclosure** = Difference between the prior and the posterior.

- Can seeing allowed data **change an adversary's mind** about sensitive data?
- **Issue:** Must estimate what the adversary believes *a priori*.
- **Hard to model** realistically and **validate** empirically.

# Proposal: Prior-agnostic privacy criteria


We should use metrics that **do not require modeling priors**.

- Two examples from theoretical literature:
  - PQI: Can allowed data imply John **definitely has** pneumonia?
  - NQI: Can allowed data imply John **definitely does not have** pneumonia?
- Coarser-grained than Bayesian criteria, but meaningful regardless of belief.
- **Challenges:**
  - How to **measure** prior-agnostic privacy?
  - How to **present** the result to the admin?

# Challenges beyond enforcement



Create  
a policy



Evaluate  
the policy



Enforce  
the policy



Diagnose  
a violation




# When a query gets rejected...

- One day, the application issues a query that gets **rejected** under the policy.
  - **Why** was the query rejected?
  - **How** do I fix this?
- You're shown: The policy, the offending query, a stack trace.
  - Diagnosing the violation is still difficult—*too much information!*
- **How to better assist the admin in diagnosing such a violation?**
  - Ideally: Give a **small amount of feedback** that the admin can act upon.

# What would the ideal feedback look like?

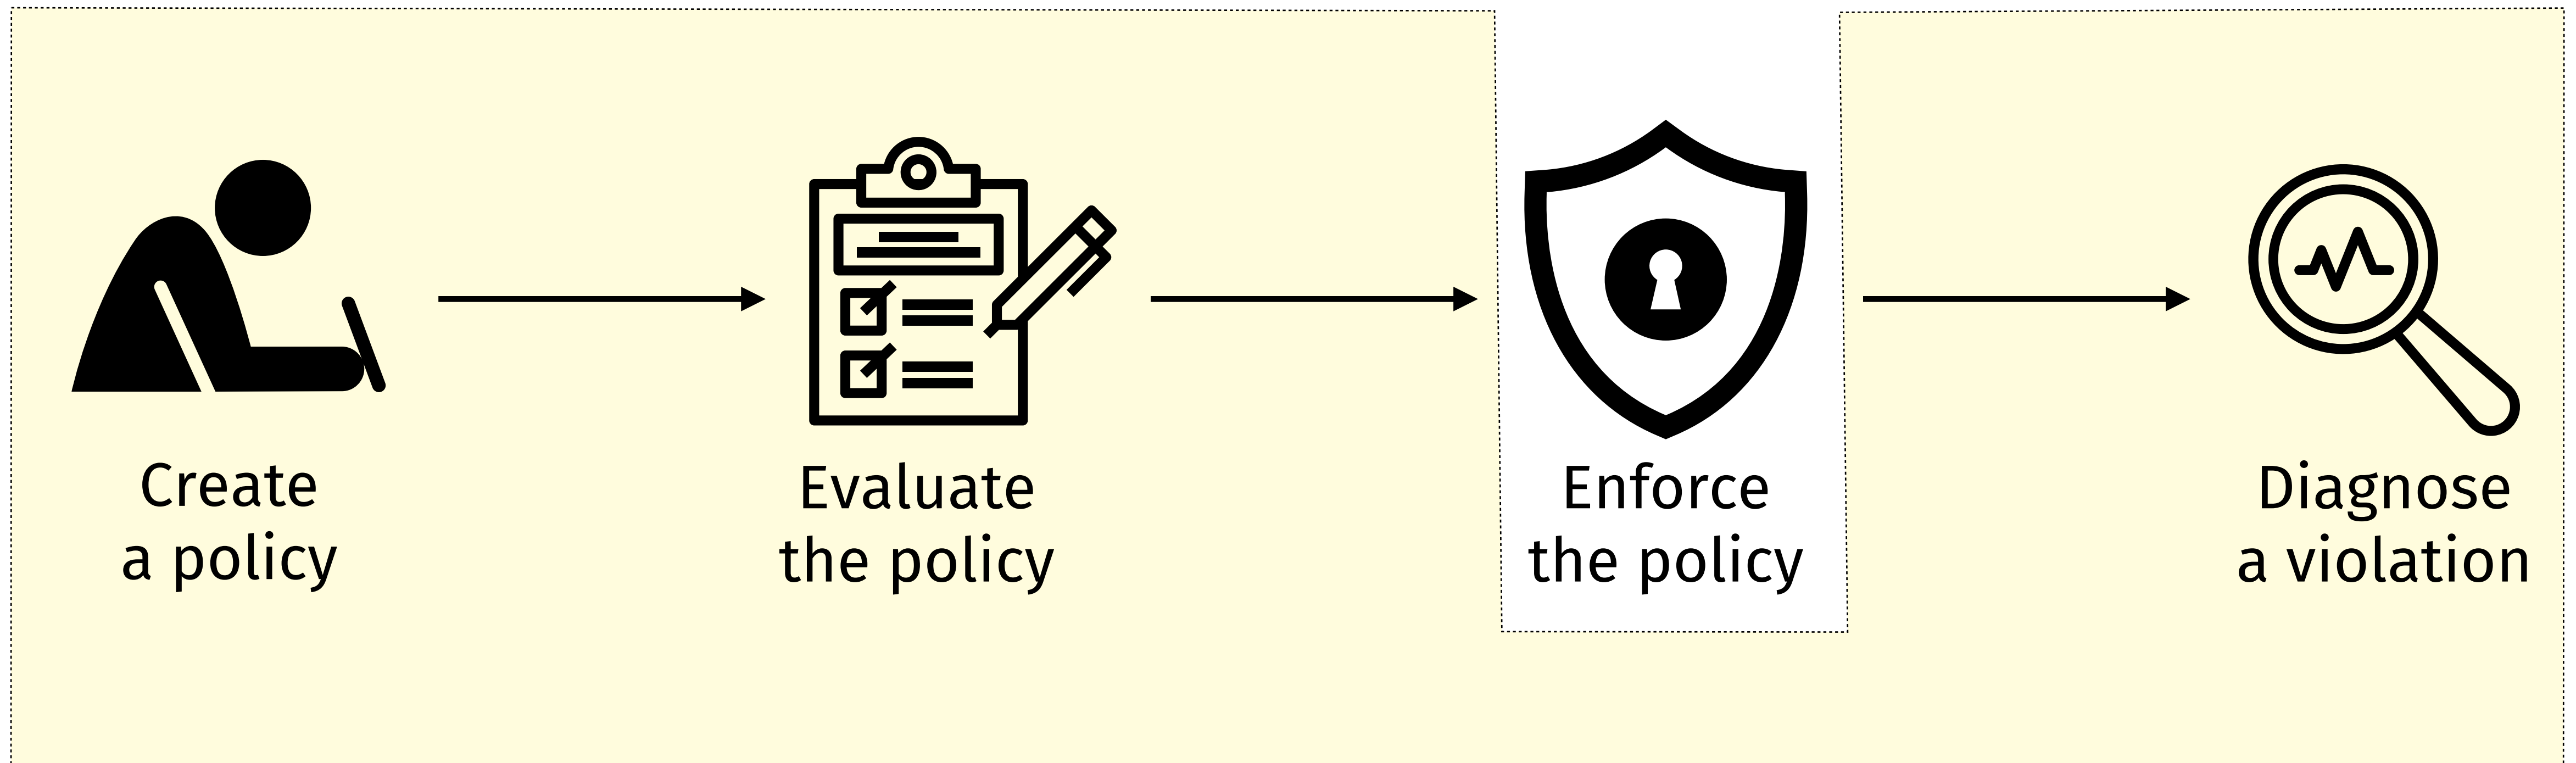
- It is unclear, especially for expressive policies.
- An **allow-list** policy specifies accessible information.
- A query gets **rejected** simply because the information it reveals is **not contained in** “accessible regions”.
- No policy item—or subset of items—is **responsible** for the rejection.
  - **Hard to explain why** the query was rejected.



# Proposal: Generate fixes, show to admin

- Fixing the policy: Grant access to more data.
  - **Approach:** Re-run policy **extraction** on updated source code / input suite.
- Fixing the application: Narrow down the query, or insert access check.
  - **Approach:** View-based query rewriting, abductive reasoning.

# Addressing the **full life-cycle** of access control



Thank you!

Wen Zhang <[zhangwen@cs.berkeley.edu](mailto:zhangwen@cs.berkeley.edu)>