18734: Foundations of Privacy

Bootstrapping Privacy Compliance in Big Data Systems

Anupam Datta

Fall 2014

Privacy Compliance for Bing

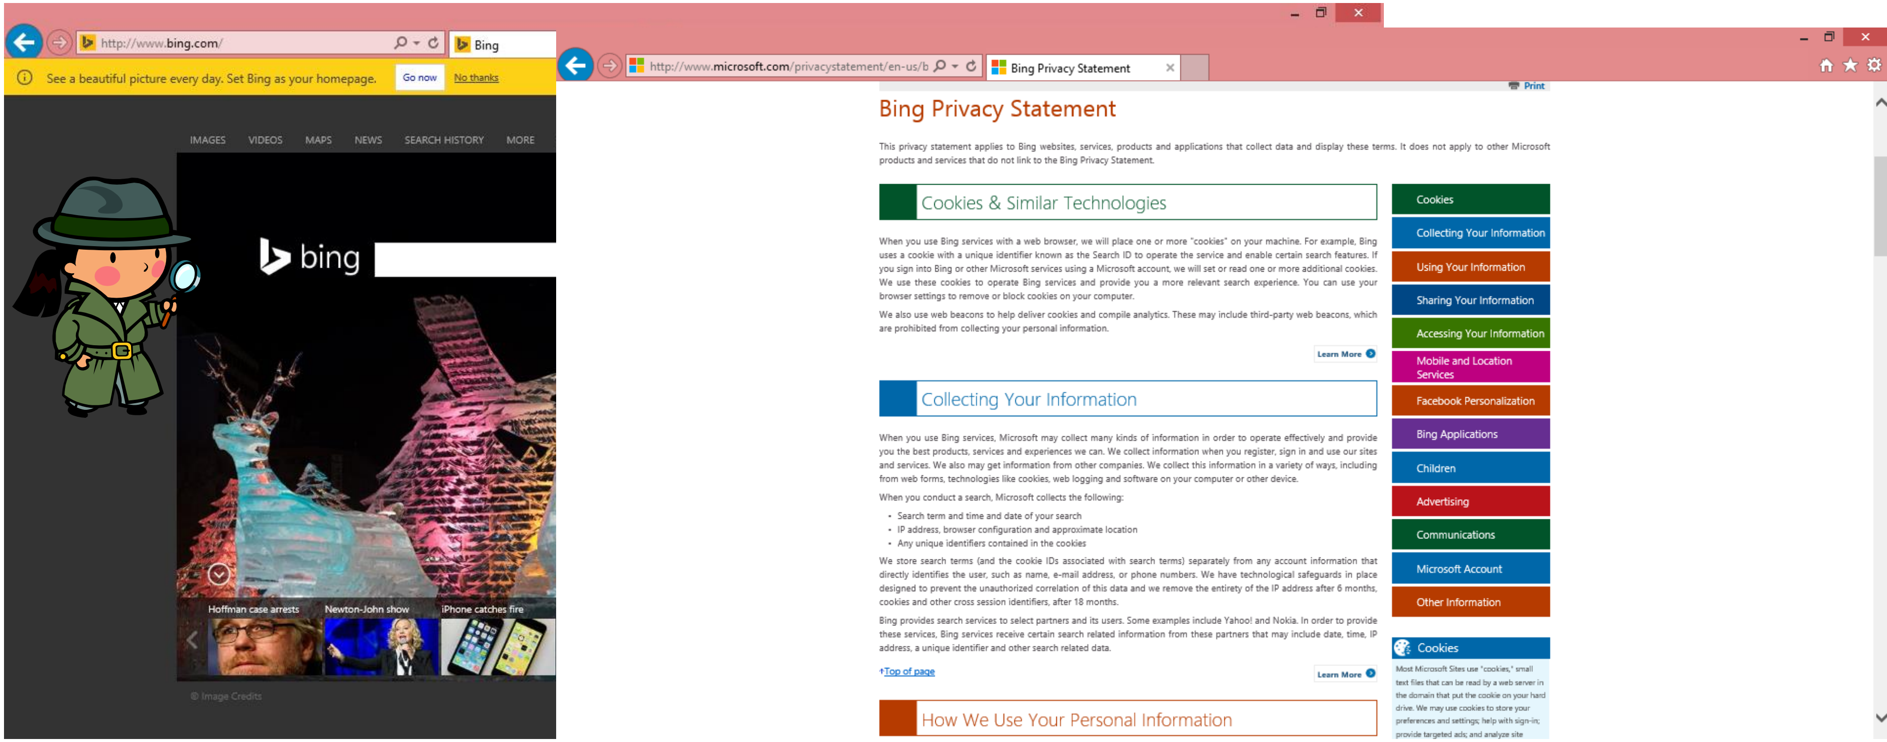

Bing Privacy Statement

This privacy statement applies to Bing websites, services, products and applications that collect data and display these terms. It does not apply to other Microsoft products and services that do not link to the Bing Privacy Statement.

Cookies & Similar Technologies

When you use Bing services with a web browser, we will place one or more "cookies" on your machine. For example, Bing uses a cookie with a unique identifier known as the Search ID to operate the service and enable certain search features. If you sign into Bing or other Microsoft services using a Microsoft account, we will set or read one or more additional cookies. We use these cookies to operate Bing services and provide you a more relevant search experience. You can use your browser settings to remove or block cookies on your computer.

We also use web beacons to help deliver cookies and compile analytics. These may include third-party web beacons, which are prohibited from collecting your personal information.

Learn More

Collecting Your Information

When you use Bing services, Microsoft may collect many kinds of information in order to operate effectively and provide you the best products, services and experiences we can. We collect information when you register, sign in and use our sites and services. We also may get information from other companies. We collect this information in a variety of ways, including from web forms, technologies like cookies, web logging and software on your computer or other device.

When you conduct a search, Microsoft collects the following:

- Search term and time and date of your search
- IP address, browser configuration and approximate location
- Any unique identifiers contained in the cookies

We store search terms (and the cookie IDs associated with search terms) separately from any account information that directly identifies the user, such as name, e-mail address, or phone numbers. We have technological safeguards in place designed to prevent the unauthorized correlation of this data and we remove the entirety of the IP address after 6 months, cookies and other cross session identifiers, after 18 months.

Bing provides search services to select partners and its users. Some examples include Yahoo! and Nokia. In order to provide these services, Bing services receive certain search related information from these partners that may include date, time, IP address, a unique identifier and other search related data.

↑Top of page

Learn More

How We Use Your Personal Information

Cookies
Collecting Your Information
Using Your Information
Sharing Your Information
Accessing Your Information
Mobile and Location Services
Facebook Personalization
Bing Applications
Children
Advertising
Communications
Microsoft Account
Other Information

Cookies

Most Microsoft Sites use "cookies," small text files that can be read by a web server in the domain that put the cookie on your hard drive. We may use cookies to store your preferences and settings; help with sign-in; provide targeted ads; and analyze site

<u>Setting:</u>

- Auditor has access to source code

The Privacy Compliance Challenge

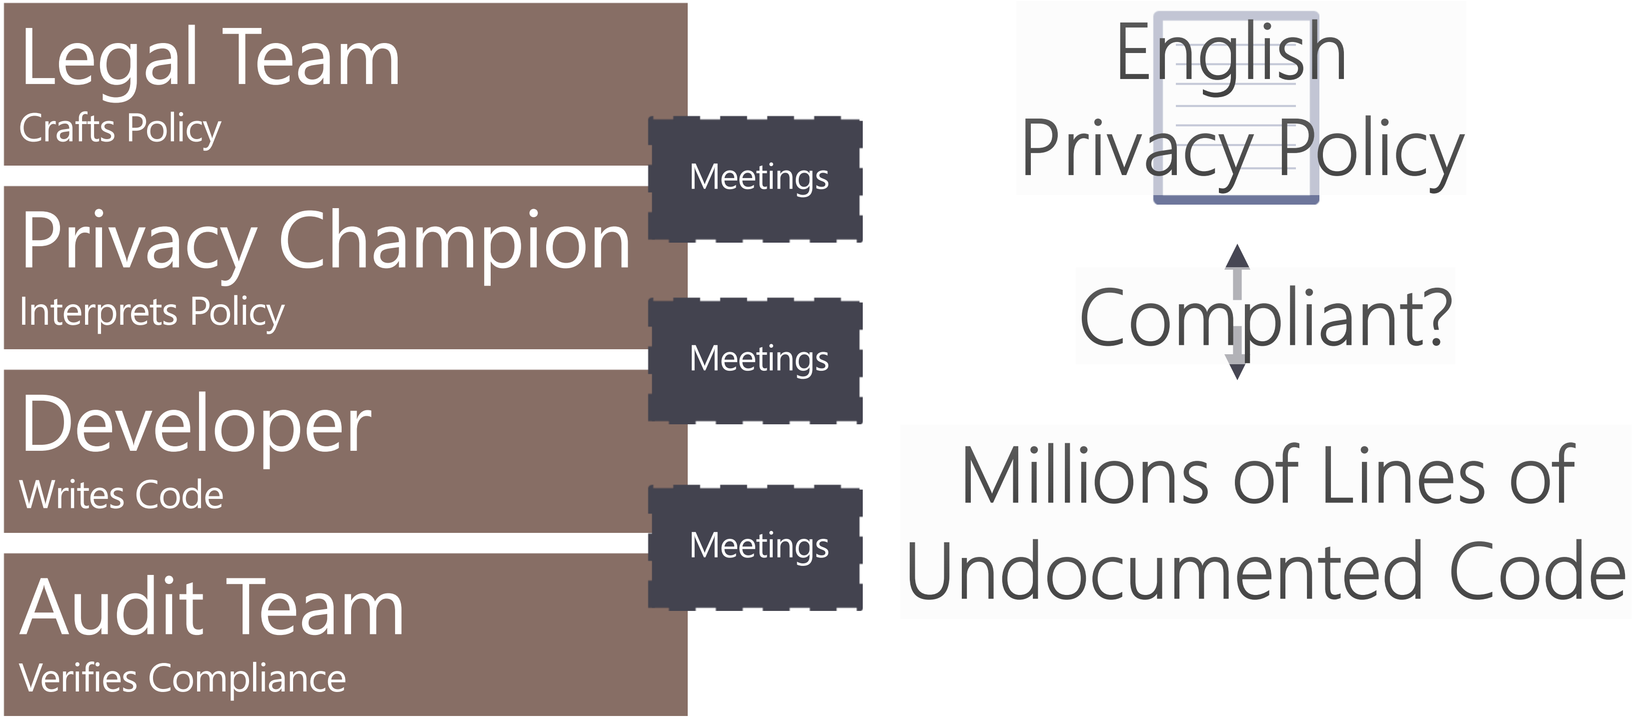

Legal Team

Crafts Policy

Meetings

Privacy Champion

Interprets Policy

Meetings

Developer

Writes Code

Meetings

Audit Team

Verifies Compliance

English Privacy Policy

Compliant?

Millions of Lines of Undocumented Code

A Streamlined Audit Workflow

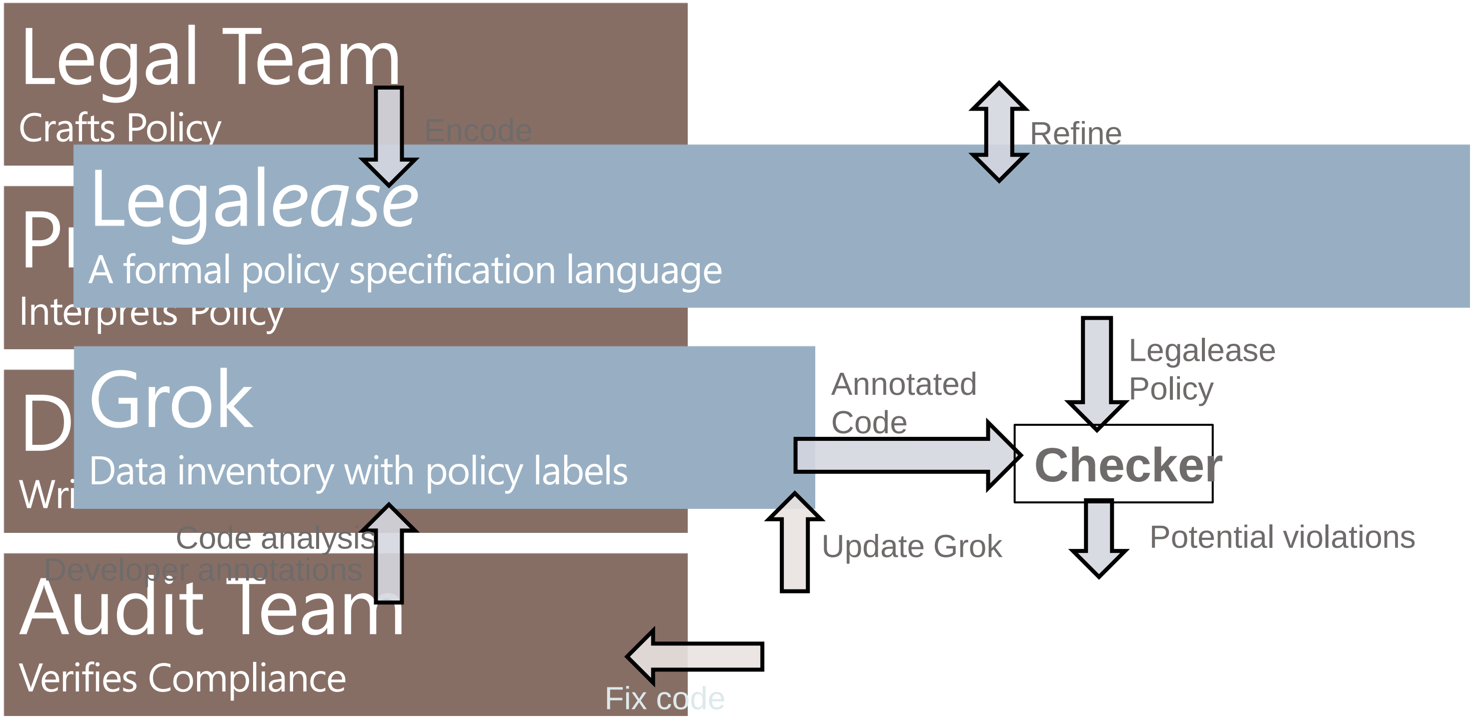

Legal Team

Crafts Policy

Encode

Refine

Legal*ease*

A formal policy specification language

Pr

Interprets Policy

Grok

D

Data inventory with policy labels

Wri

Annotated Code

Legalease Policy

Checker

Update Grok

Potential violations

Code analysis

Developer annotations

Audit Team

Verifies Compliance

Fix code

A Streamlined Audit Workflow

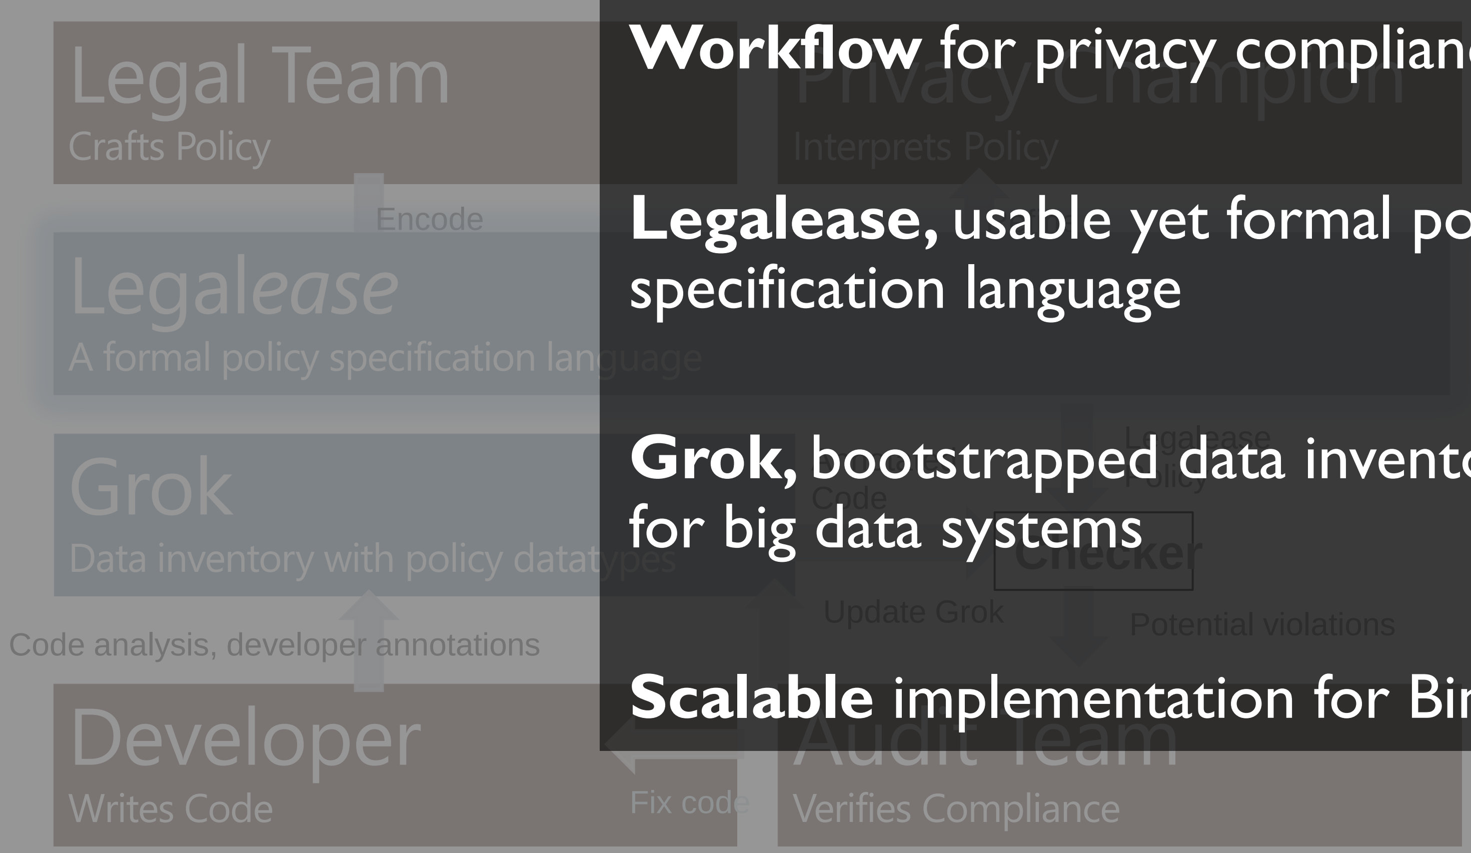

Legal Team

Crafts Policy

Privacy Champion

Interprets Policy

Encode

Legal*ease*

A formal policy specification language

Grok

Data inventory with policy datatypes

Code analysis, developer annotations

Code

Legalease Policy

Checker

Update Grok

Potential violations

Developer

Writes Code

Fix code

Audit Team

Verifies Compliance

Workflow for privacy compliance

Legalease, usable yet formal policy specification language

Grok, bootstrapped data inventory for big data systems

Scalable implementation for Bing

A Streamlined Audit Workflow

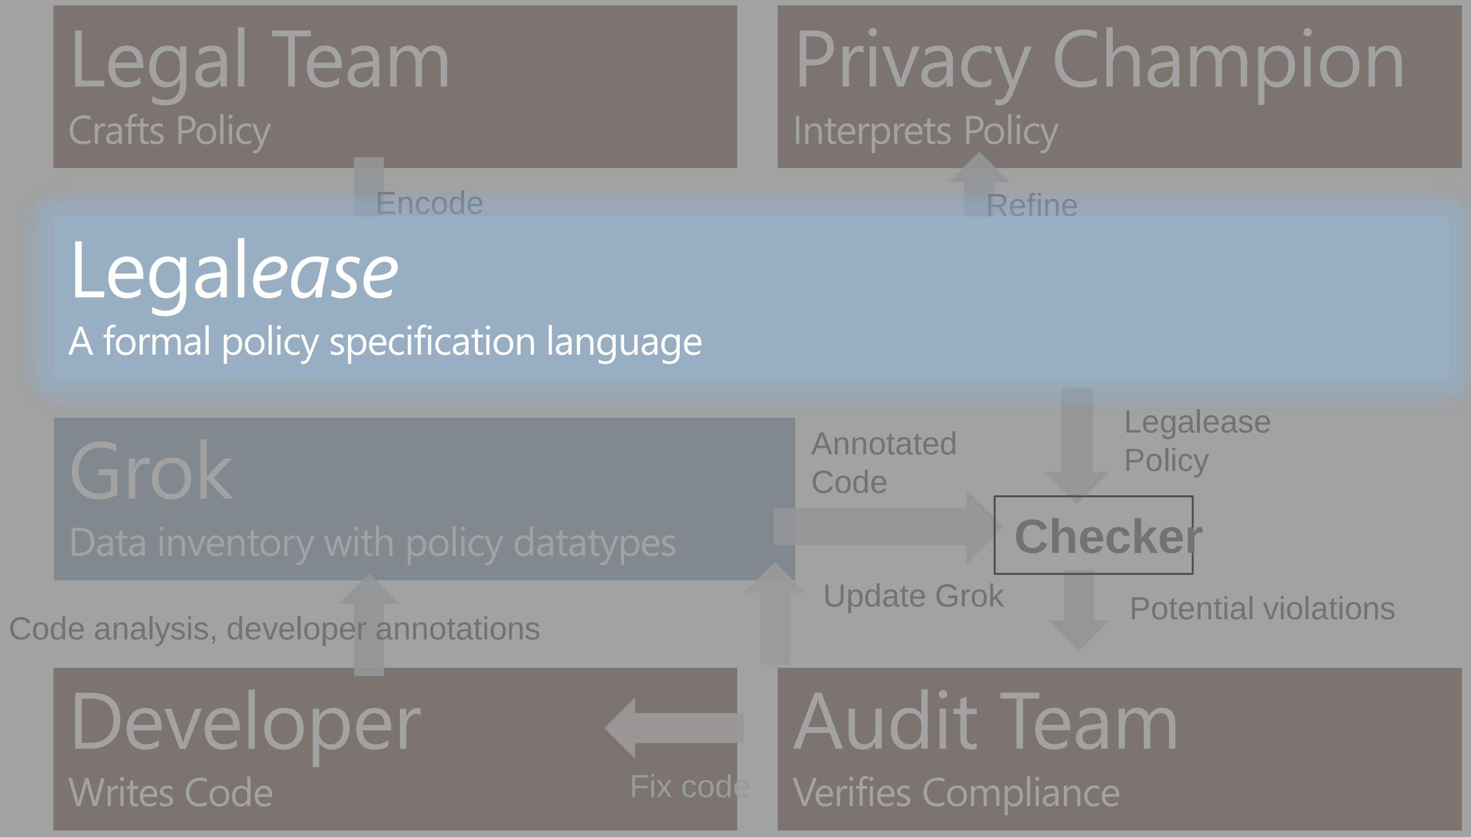

Legalease : Syntax

$$
\begin{array}{rcl}
\text{Policy Clause } C & ::= & D \mid A \\
\text{Deny Clause } D & ::= & \text{DENY } T_1 \cdots T_n \text{ EXCEPT } A_1 \cdots A_m \\
 & & \mid \text{DENY } T_1 \cdots T_n \\
\text{Allow Clause } A & ::= & \text{ALLOW } T_1 \cdots T_n \text{ EXCEPT } D_1 \cdots D_m \\
 & & \mid \text{ALLOW } T_1 \cdots T_n \\
\text{Attribute } T & ::= & \langle\text{attribute-name}\rangle\ v_1 \cdots v_l \\
\text{Value } v & ::= & \langle\text{attribute-value}\rangle
\end{array}
$$

Legalease

DENY *Datatype* IPAddress

UseForPurpose Advertising

We will **not** use **full IP Address** for **Advertising**.

Legalease

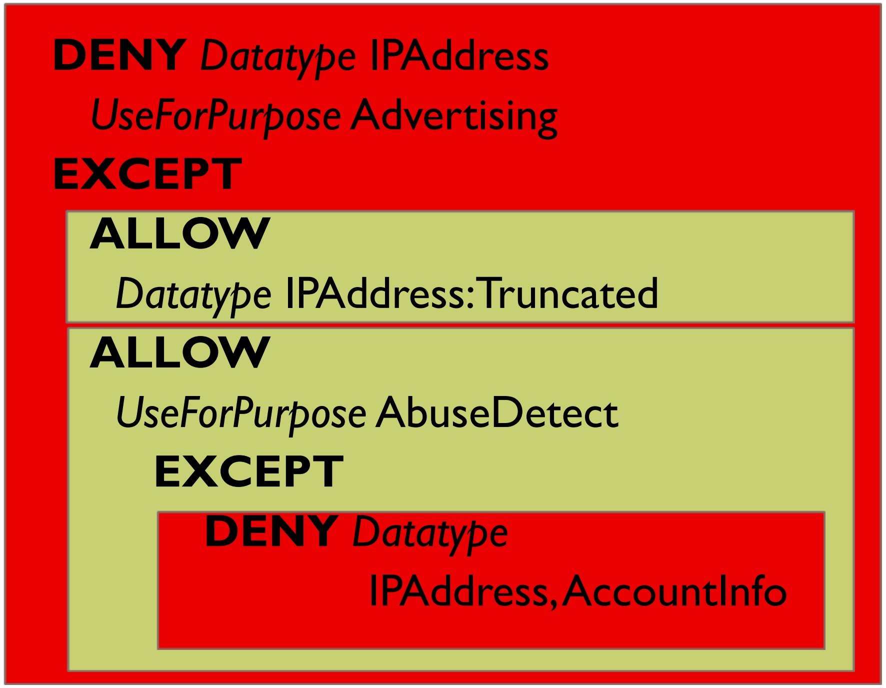

```
DENY Datatype IPAddress
  UseForPurpose Advertising
EXCEPT
  ALLOW
    Datatype IPAddress:Truncated
  ALLOW
    UseForPurpose AbuseDetect
      EXCEPT
        DENY Datatype
                IPAddress, AccountInfo
```

We will **not** use **full IP Address** for **Advertising**. IP Address may be used for **detecting abuse**. In such cases, it will not be combined with **account information.**

Designed for Usability

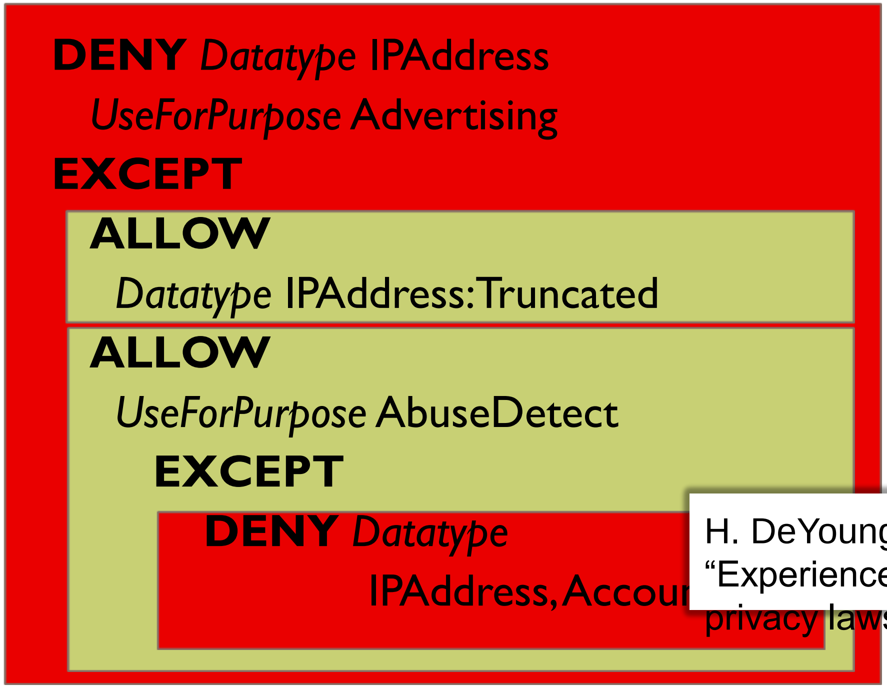

DENY *Datatype* IPAddress
UseForPurpose Advertising
EXCEPT
ALLOW
Datatype IPAddress:Truncated
ALLOW
UseForPurpose AbuseDetect
EXCEPT
DENY *Datatype*
IPAddress,Accour

Exceptions

- How legal texts are structured
- One-to one correspondence

Local Reasoning

- Each exception refines its immediate parent
- Formally proven property

H. DeYoung, D. Garg, L. Jia, D. Kaynar, and A. Datta, “Experiences in the logical specification of the HIPAA and GLBA privacy laws”

Legalease : In Action

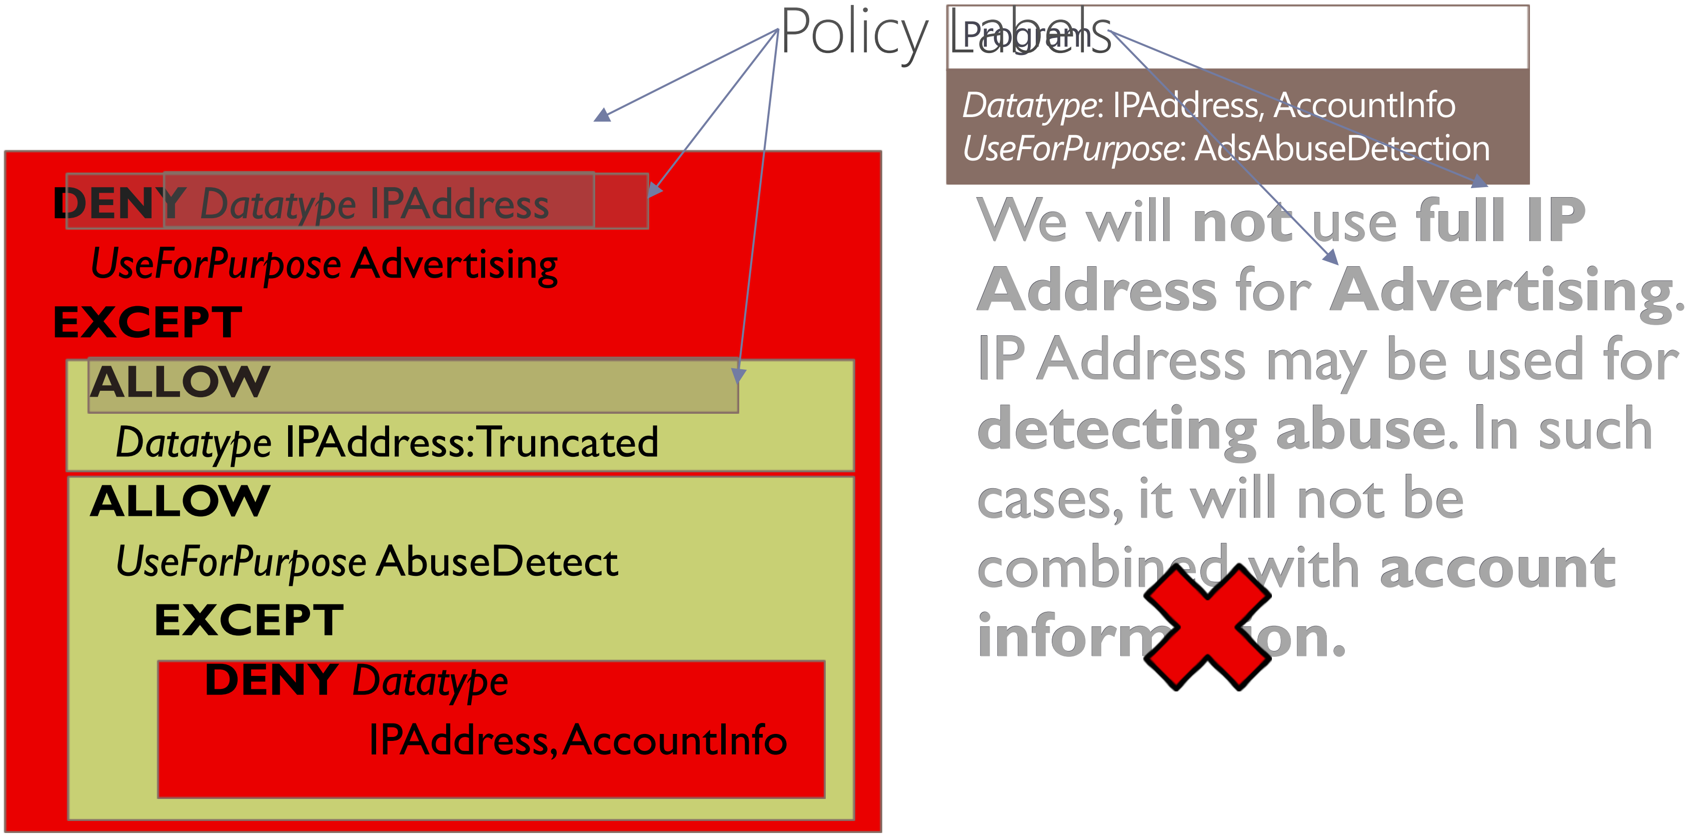

Policy Labels

Programs

Datatype: IPAddress, AccountInfo
UseForPurpose: AdsAbuseDetection

DENY *Datatype* IPAddress
UseForPurpose Advertising
EXCEPT
ALLOW
Datatype IPAddress:Truncated
ALLOW
UseForPurpose AbuseDetect
EXCEPT
DENY *Datatype*
IPAddress,AccountInfo

We will **not** use **full IP Address** for **Advertising**. IP Address may be used for **detecting abuse**. In such cases, it will not be combined with **account information.**

A Lattice of Policy Labels

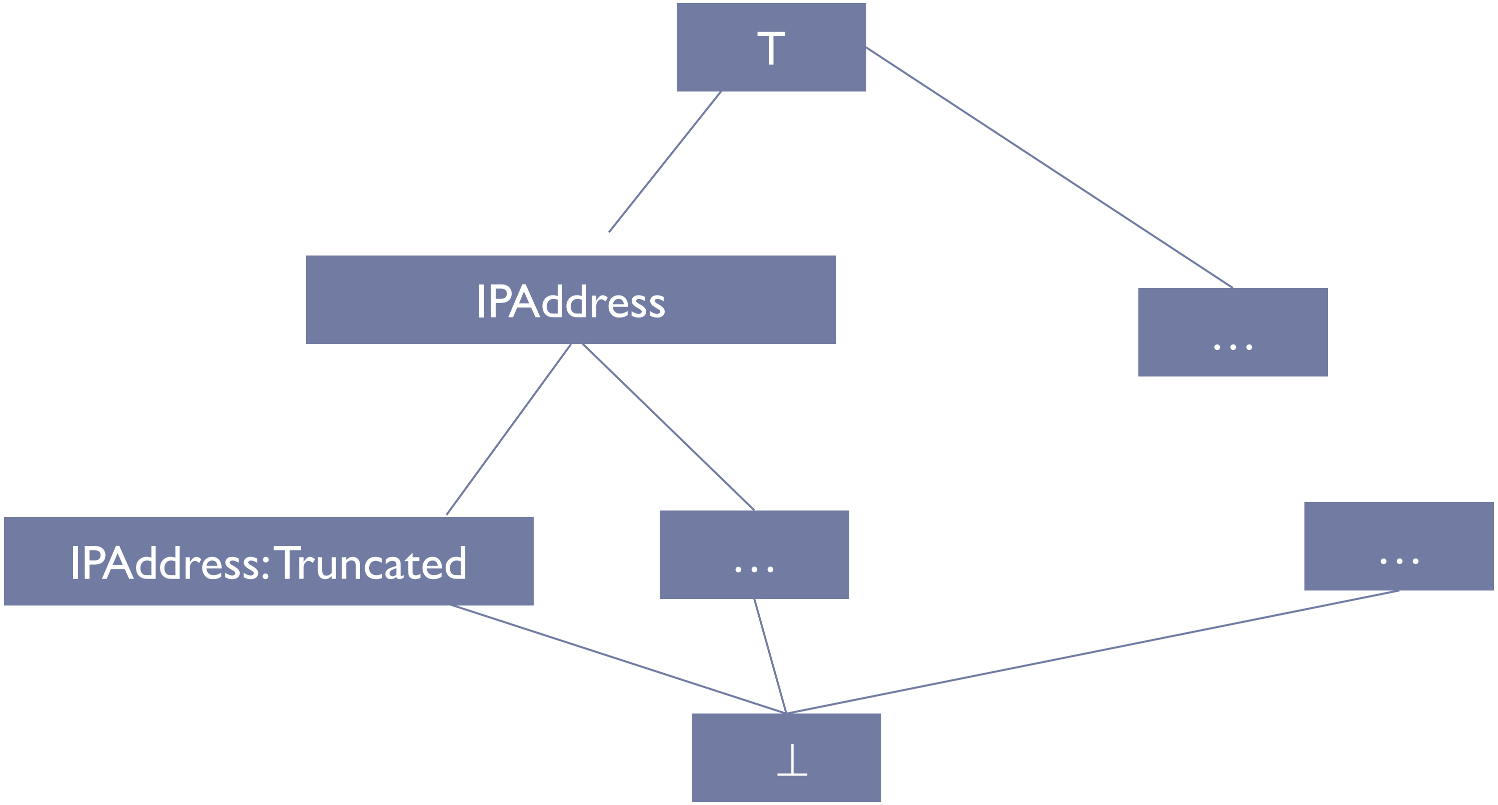

- If “IPAddress” use is allowed then so is everything below it
- If “IPAddress:Truncated” use is denied then so is everything above it

Legalease Usability

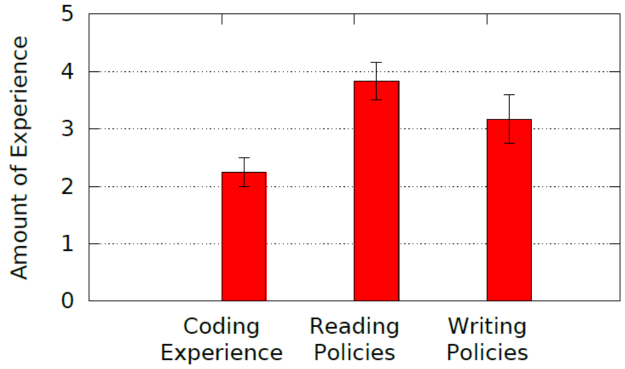

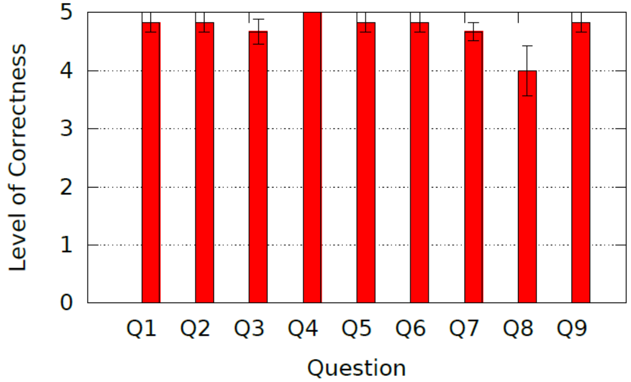

Survey taken by 12 policy authors within Microsoft

- Encode Bing data usage policy after a brief tutorial

Time spent

- 2.4 mins on the tutorial
- 14.3 mins on encoding policy

High overall correctness

A Streamlined Audit Workflow

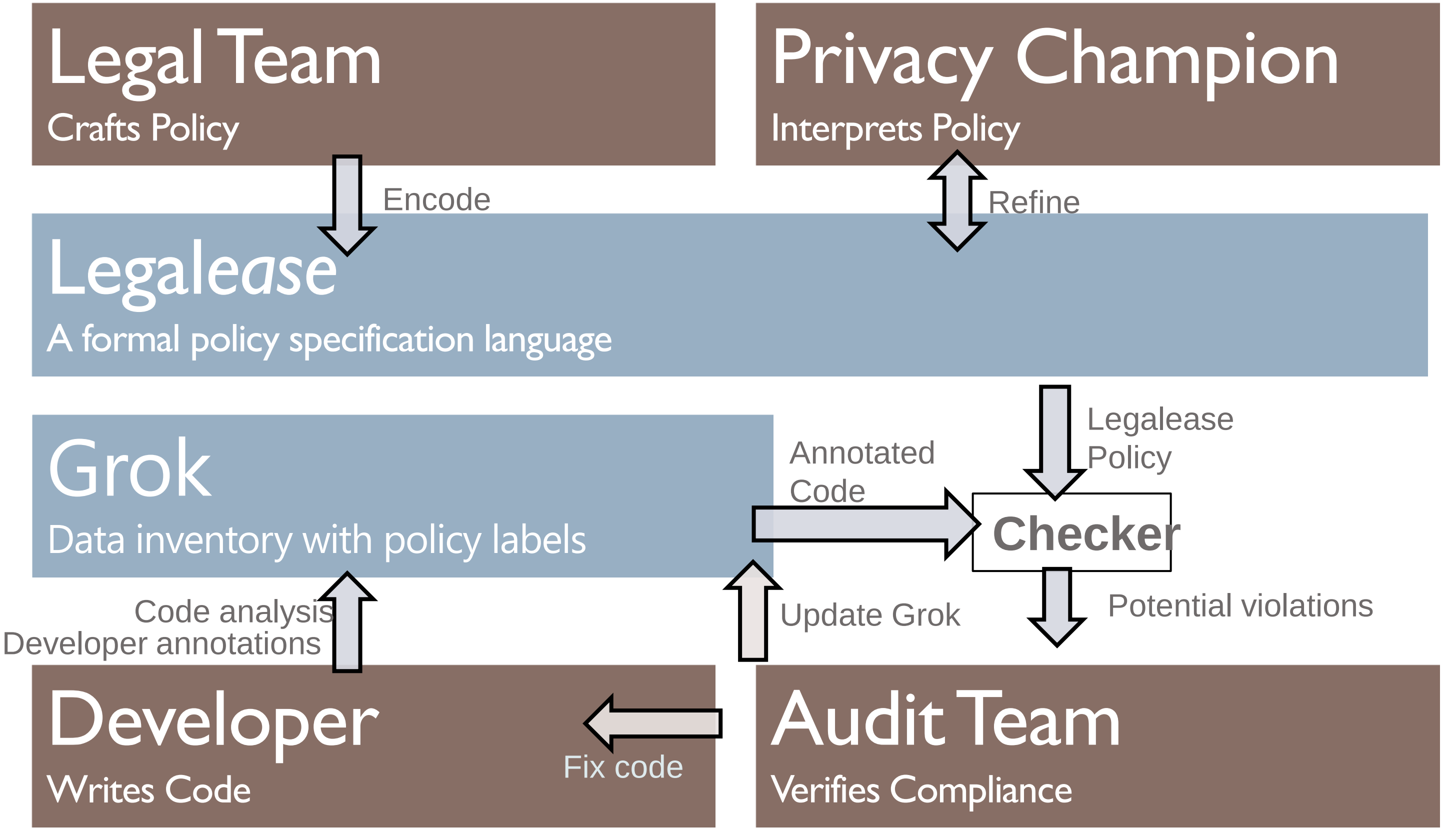

Legal Team

Crafts Policy

Privacy Champion

Interprets Policy

Encode

Refine

Legalease

A formal policy specification language

Grok

Data inventory with policy labels

Annotated Code

Legalease Policy

Checker

Code analysis

Developer annotations

Update Grok

Potential violations

Developer

Writes Code

Fix code

Audit Team

Verifies Compliance

A Streamlined Audit Workflow

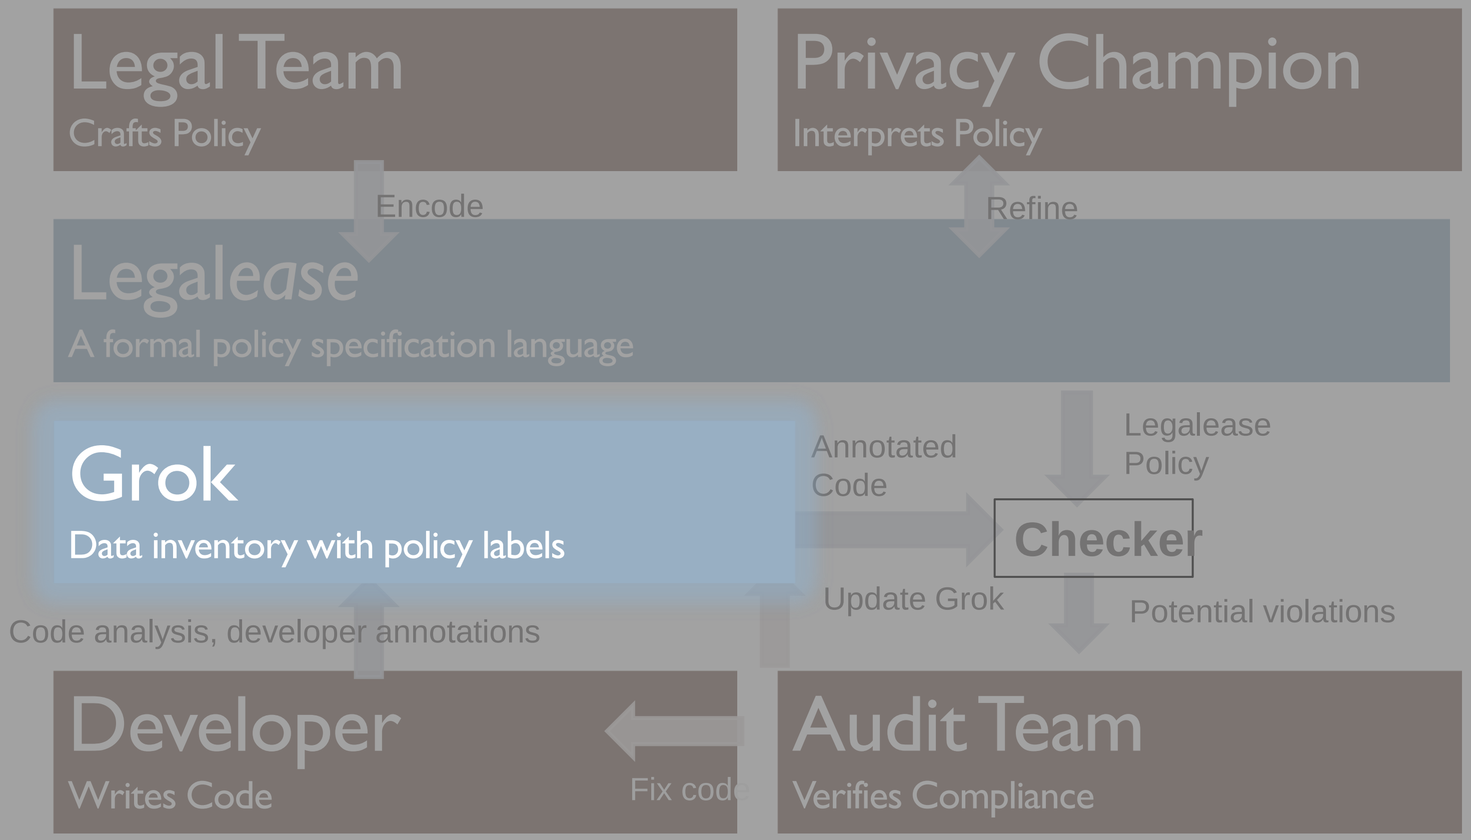

Map-Reduce Programming Systems

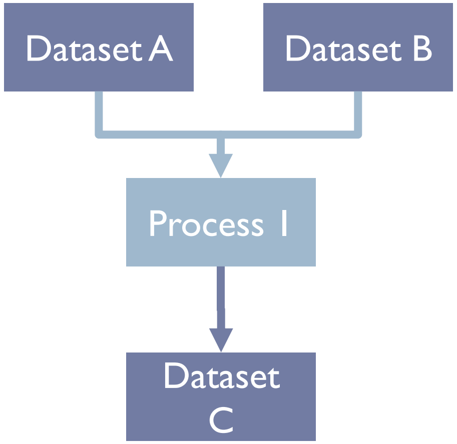

Scope, Hive, Dremel

Data in the form of Tables

Code Transforms Columns to Columns

- No Shared State
- Limited Hidden Flows

```
users =
   SELECT _name, _age FROM datasetAB
user_tag =
   SELECT GenerateTag(_name, _age)
      FROM users
OUTPUT user_tag TO datasetC
```

Grok

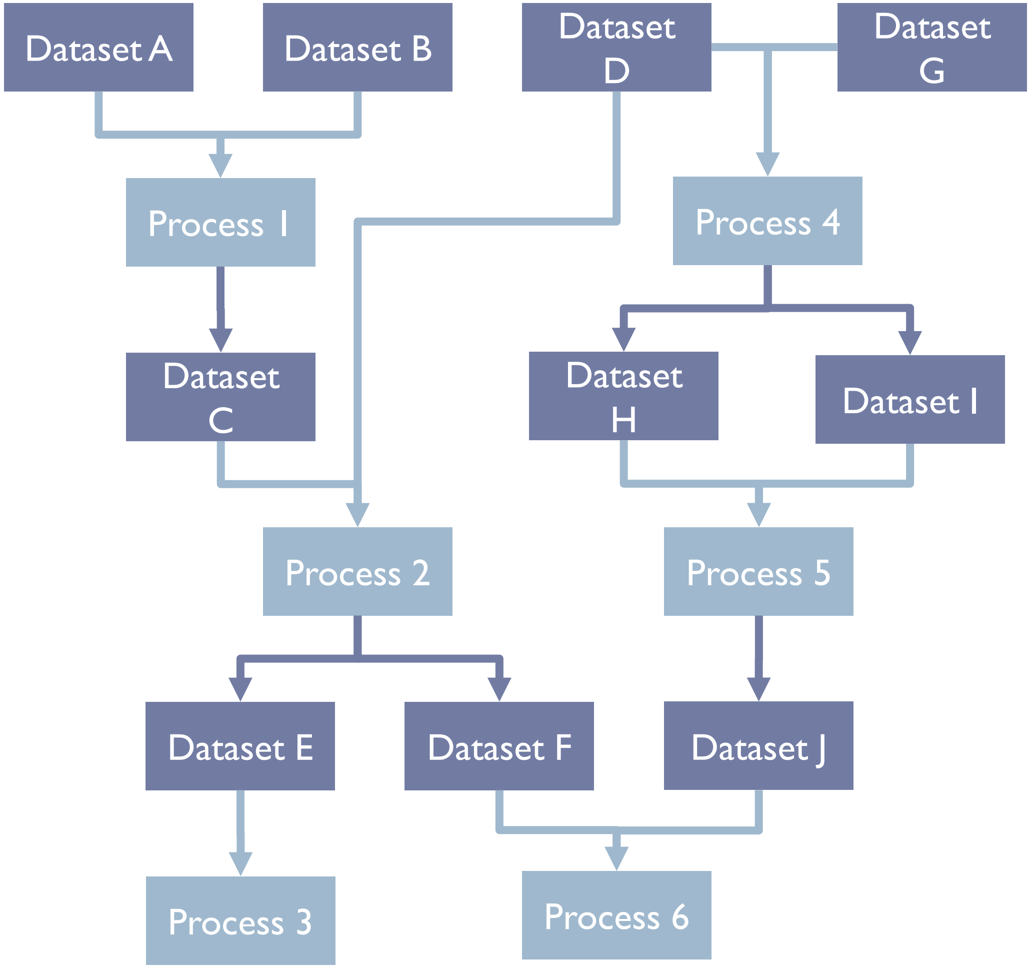
Dataset A
Dataset B
Dataset D
Dataset G
Process 1
Process 4
Dataset C
Dataset H
Dataset I
Process 2
Process 5
Dataset E
Dataset F
Dataset J
Process 3
Process 6

Grok

Purpose Labels

Annotate programs with purpose labels

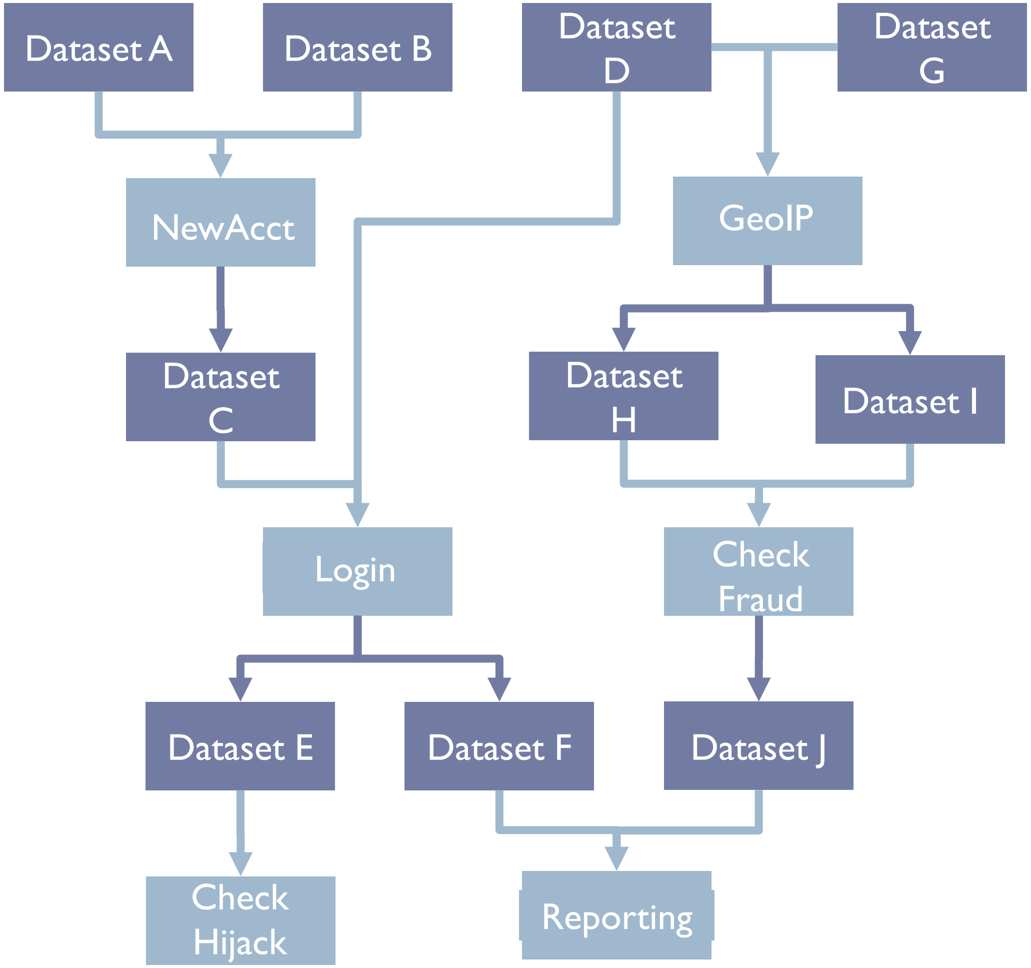

Grok

Purpose Labels

Annotate programs with purpose labels

Initial Data Labels

Heuristics and Annotations

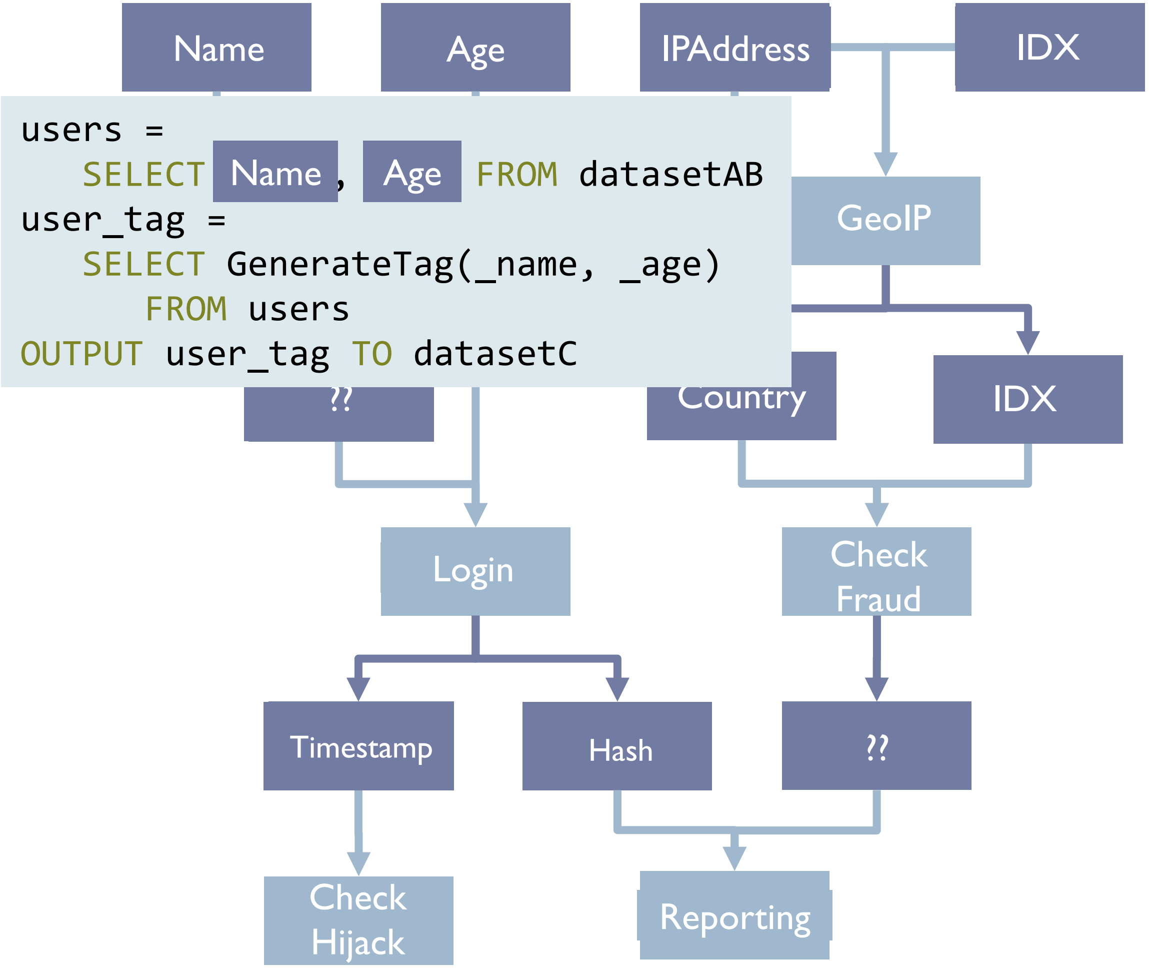

```
users =
   SELECT Name , Age FROM datasetAB
user_tag =
   SELECT GenerateTag(_name, _age)
      FROM users
OUTPUT user_tag TO datasetC
```

Grok

Purpose Labels

Annotate programs with purpose labels

Initial Data Labels

Heuristics and Annotations

Flow Labels

Source labels propagated via data flow graph

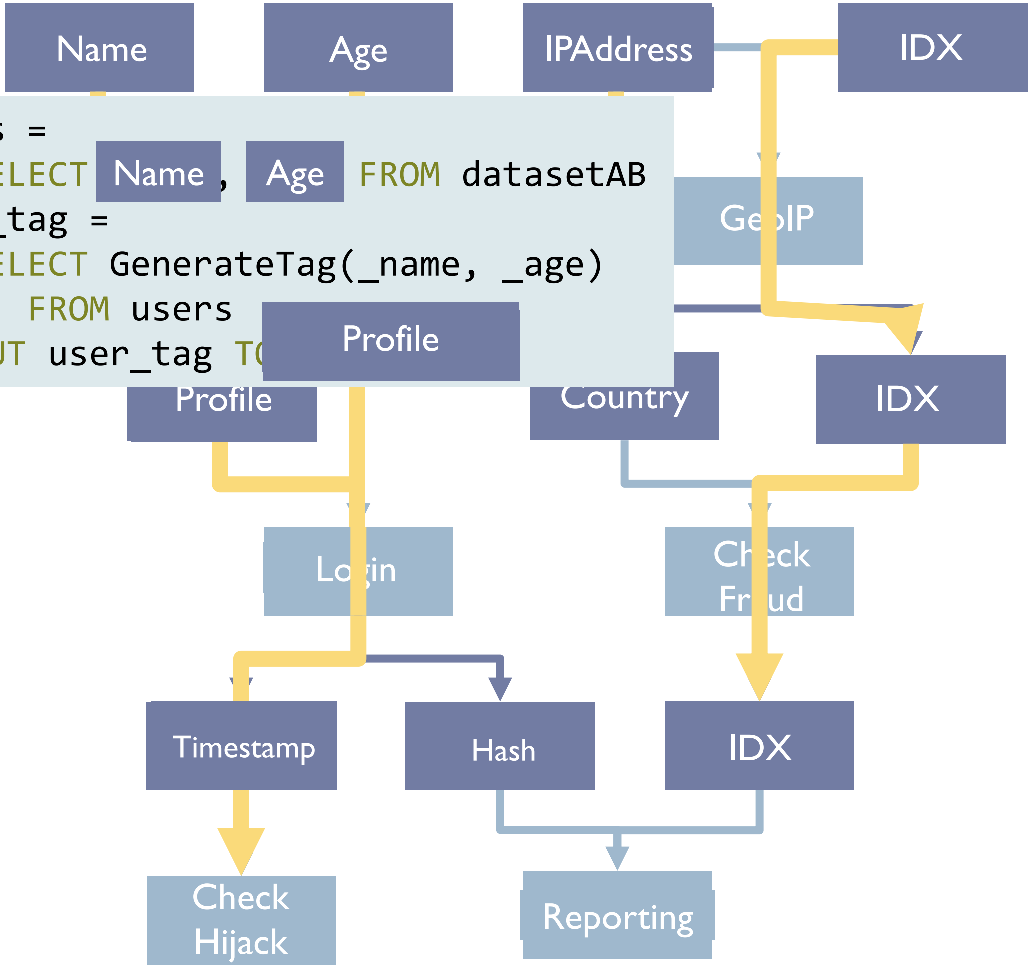

```
users =
   SELECT Name, Age FROM datasetAB
user_tag =
   SELECT GenerateTag(_name, _age)
      FROM users
OUTPUT user_tag TO Profile
```

D. E. Denning. “A lattice model of secure information flow”

A Lattice of Policy Labels

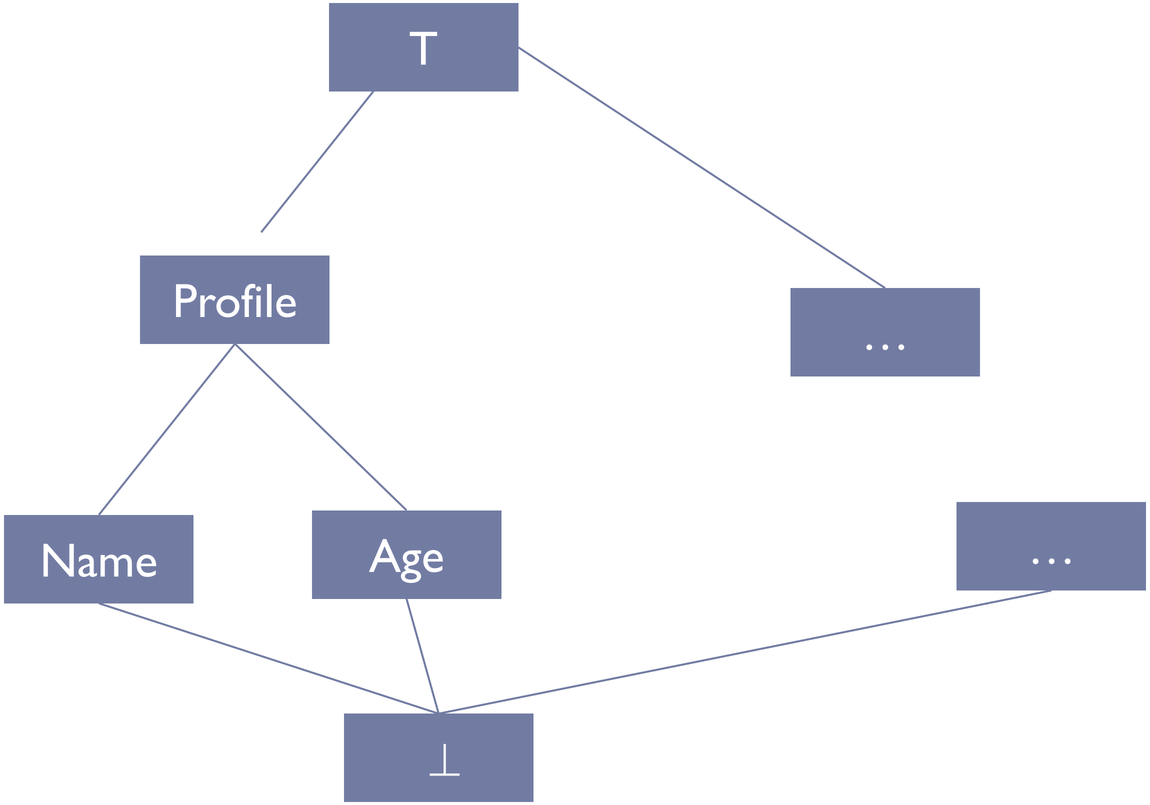

- If “Profile” use is allowed then so is everything below it
- If “Name” use is denied then so is everything above it

Implicit flows

```
users =
   SELECT Name , Age FROM datasetAB

users_35 =
   SELECT _name
      FROM users
      WHERE (_age > 35)

OUTPUT users_35 TO Profile
```

Beyond direct flows discussed in healthcare audit examples

Map-Reduce

Map

Operate on rows

in parallel

eg. filtering

Reduce

Combine groups of rows

eg. aggregation

```
users =
   SELECT Name, Age FROM datasetAB

users_35 =
   SELECT _name, _age
      FROM users
      WHERE (_age > 35)

ages_35 =
   SELECT _age, COUNT(_name) AS Profile
      FROM users_35
      GROUP BY _age

OUTPUT ages_35 TO datasetC
```

Combine Noisy Sources

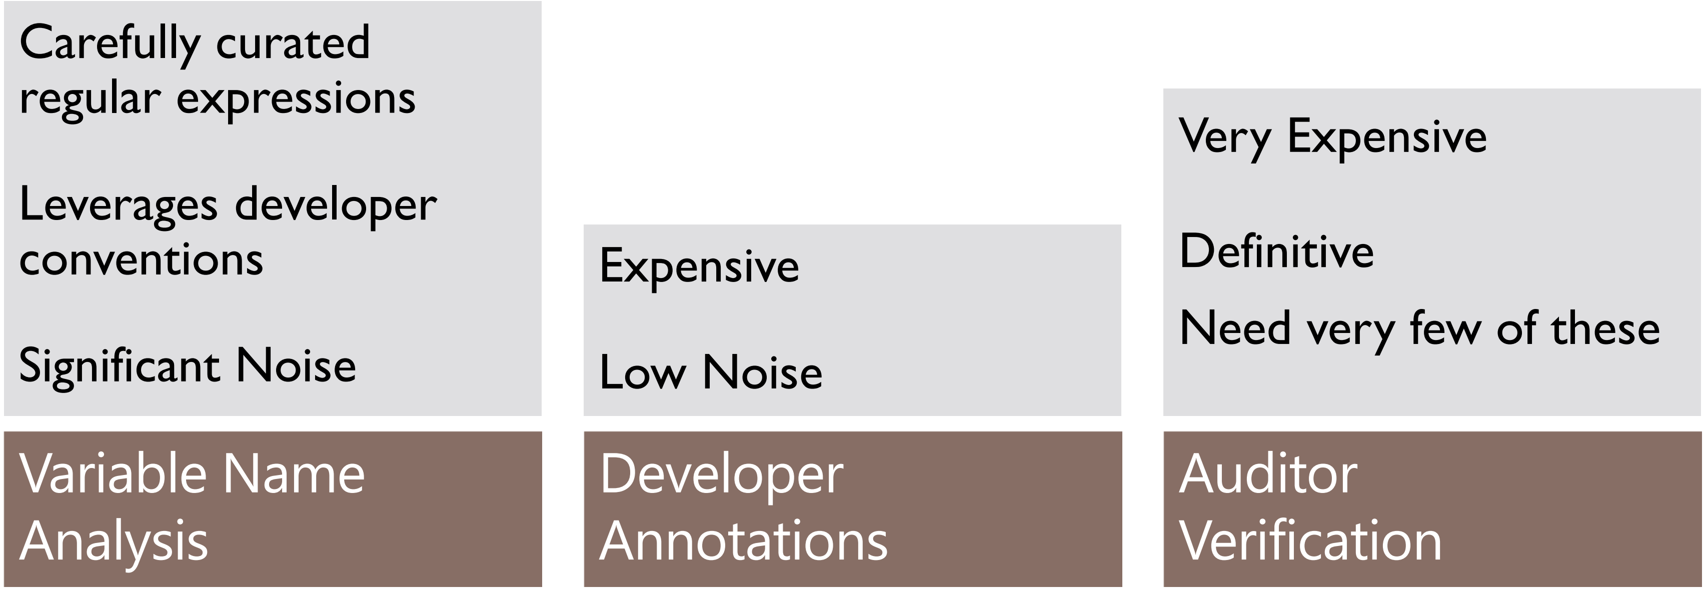

Variable Name Analysis

Carefully curated regular expressions

Leverages developer conventions

Significant Noise

Developer Annotations

Expensive

Low Noise

Auditor Verification

Very Expensive

Definitive

Need very few of these

Why Bootstrapping Grok Works

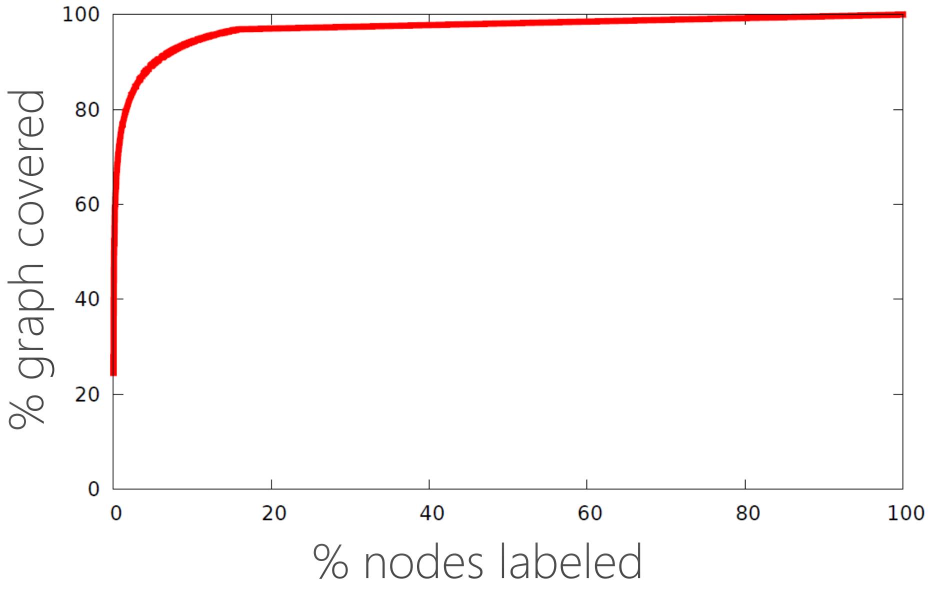

— Pick the nodes which will label the most of the graph

~200 annotations label 60% of nodes

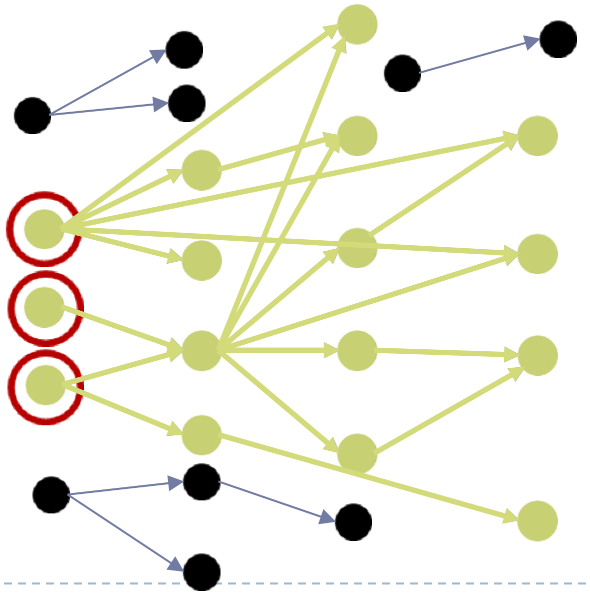

A small number of annotations is enough to get off the ground.

Scale

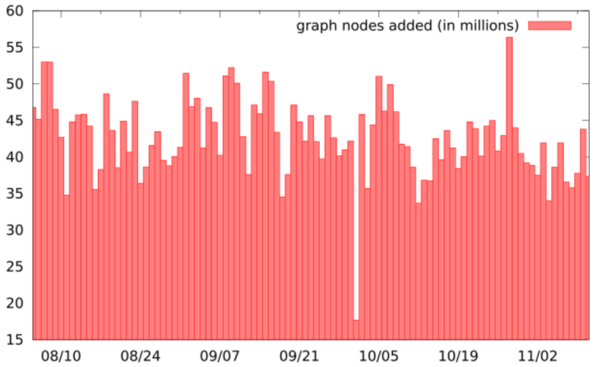

Fig. 9. Number of GROK data flow graph nodes added each day

- 77,000 jobs run each day
  - By 7000 entities
  - 300 functional groups
- 1.1 million unique lines of code
  - 21% changes on avg, daily
  - 46 million table schemas
  - 32 million files
- Manual audit infeasible
- Information flow analysis takes ~30 mins daily

A Streamlined Audit Workflow

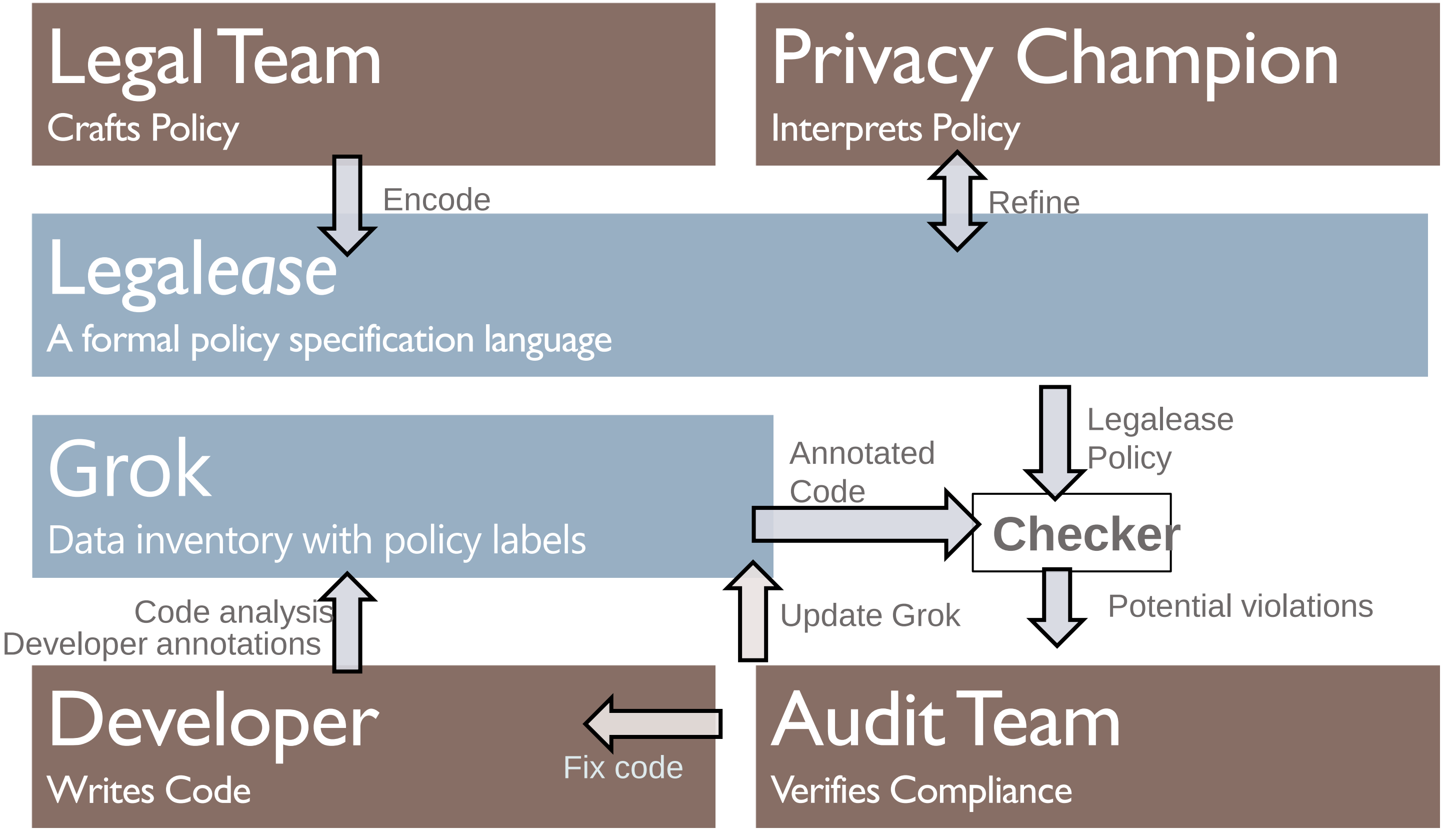

A Streamlined Audit Workflow

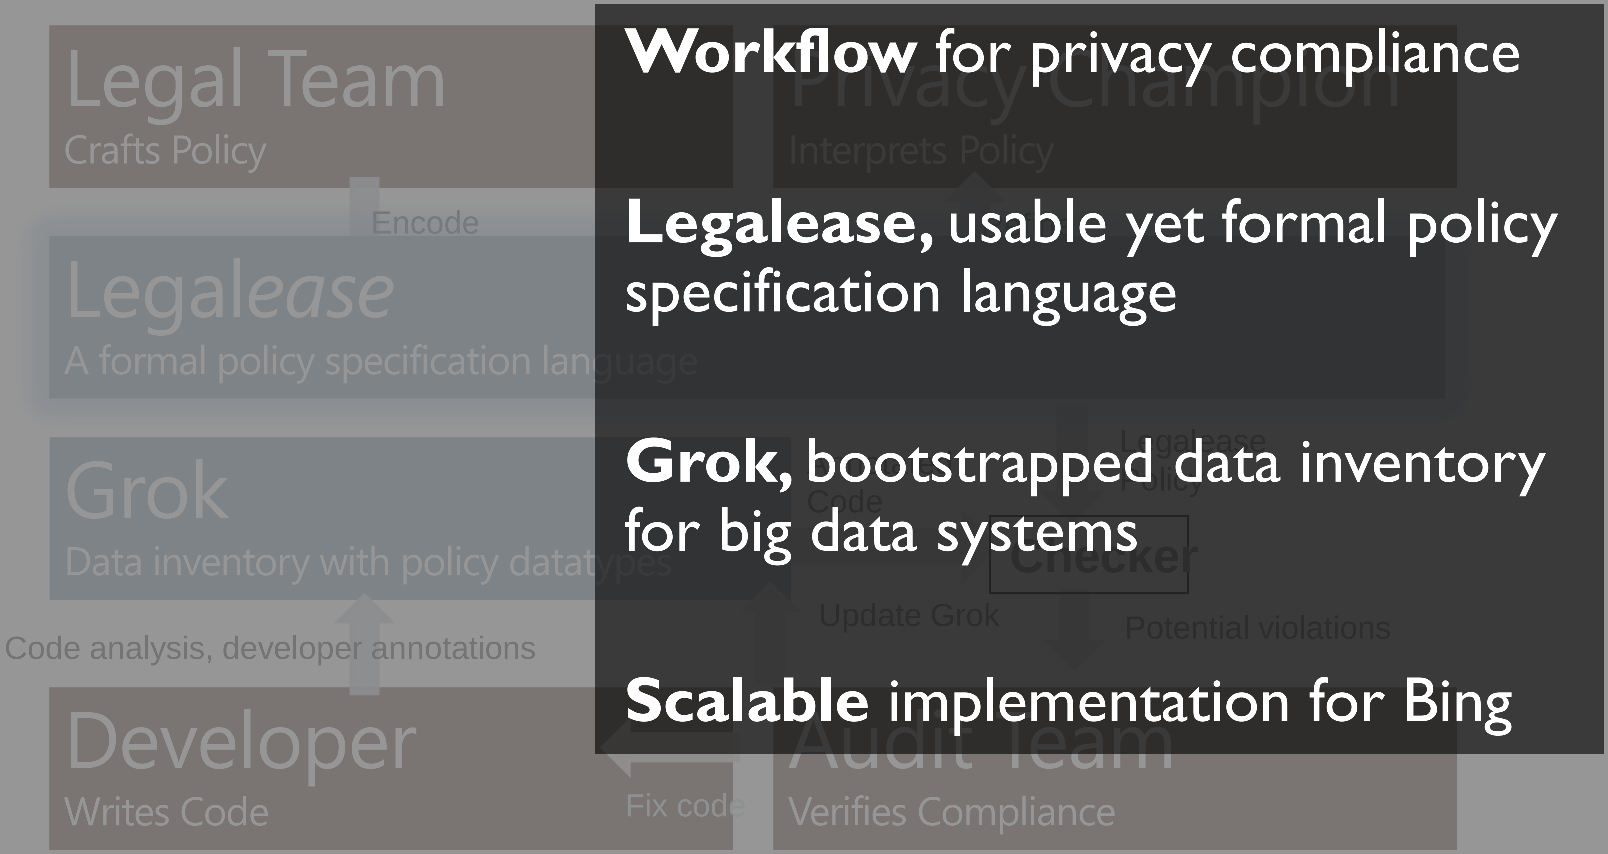

Workflow for privacy compliance

Legalease, usable yet formal policy specification language

Grok, bootstrapped data inventory for big data systems

Scalable implementation for Bing

Reference

- S. Sen, S. Guha, A. Datta, S. Rajamani, J. Tsai, J. M. Wing, Bootstrapping Privacy Compliance in Big Data Systems, in *Proceedings of 35th IEEE Symposium on Security and Privacy*, May 2014.

Policy Labels : Datatypes

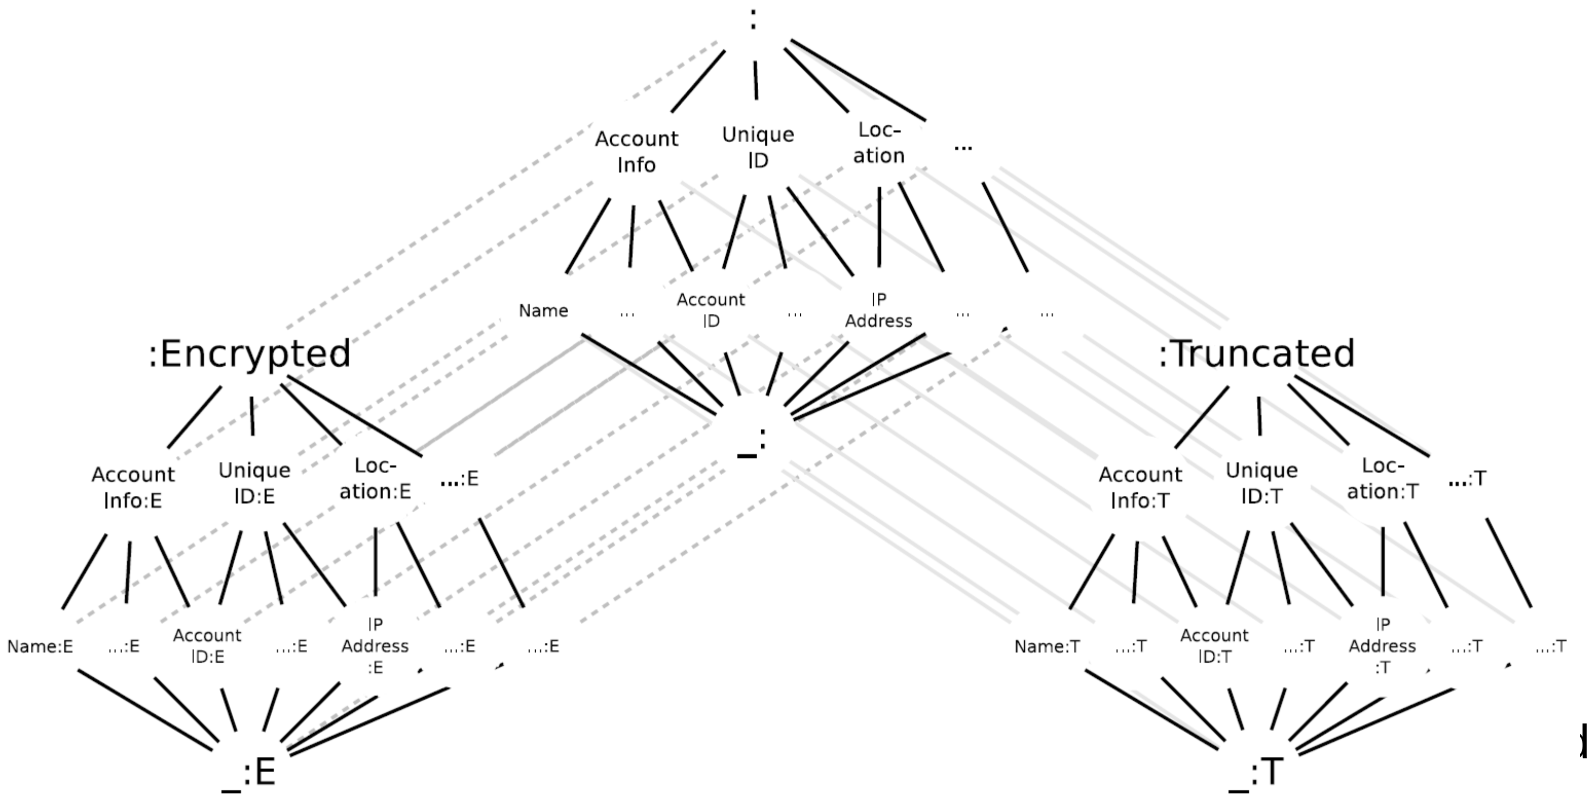

Policy Types : Concept Lattices

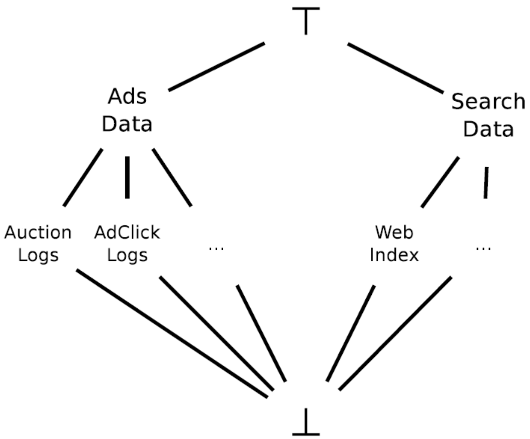

InStore Lattice

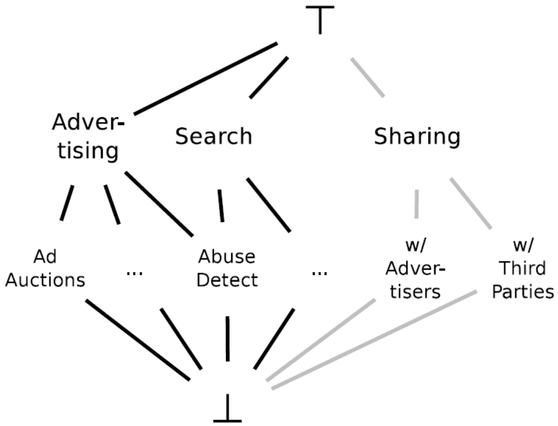

UseForPurpose Lattice

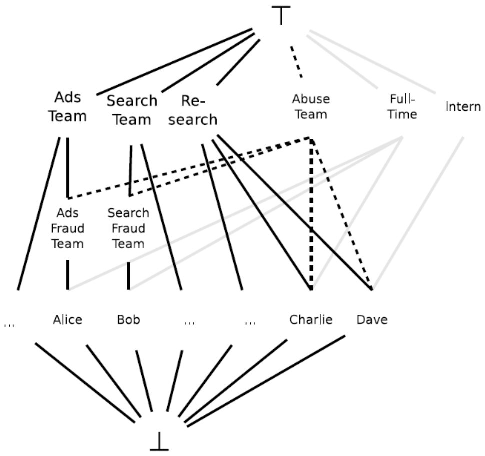

AccessByRole Lattice

Formal Semantics

$$\frac{\boxed{T^G \sqsubseteq T^C} \quad \exists_i D_i \ \ \textsf{denies} \ \ T^G}{\textsf{ALLOW}\ T^C\ \textsf{EXCEPT}\ D_1 \cdots D_m \ \ \textsf{denies} \ \ T^G} \ (\mathrm{A}_2)$$

Based on Lattice Orderings on Policy Types

Formal Semantics

$$\frac{T^G \sqsubseteq T^C \quad \boxed{\exists_i D_i \ \ \text{denies} \ \ T^G}}{\text{ALLOW } T^C \text{ EXCEPT} \boxed{D_1 \cdots D_m} \text{denies} \ \ T^G} \ (\mathrm{A}_2)$$

- Recursively check exceptions
  - ALLOW clauses have DENY clauses as exceptions
  - Top Level clause determines Blacklist/Whitelist

Formal Semantics

$$\frac{T^G \sqsubseteq T^C \quad \exists_i D_i \ \text{denies} \ T^G}{\text{ALLOW } T^C \text{ EXCEPT } D_1 \cdots D_m \ \text{denies} \ T^G} \ (\mathrm{A}_2)$$

Structural properties about semantics

- Always returns an unambiguous answer
- Weakening lattice orderings makes policy more permissive

Encoding

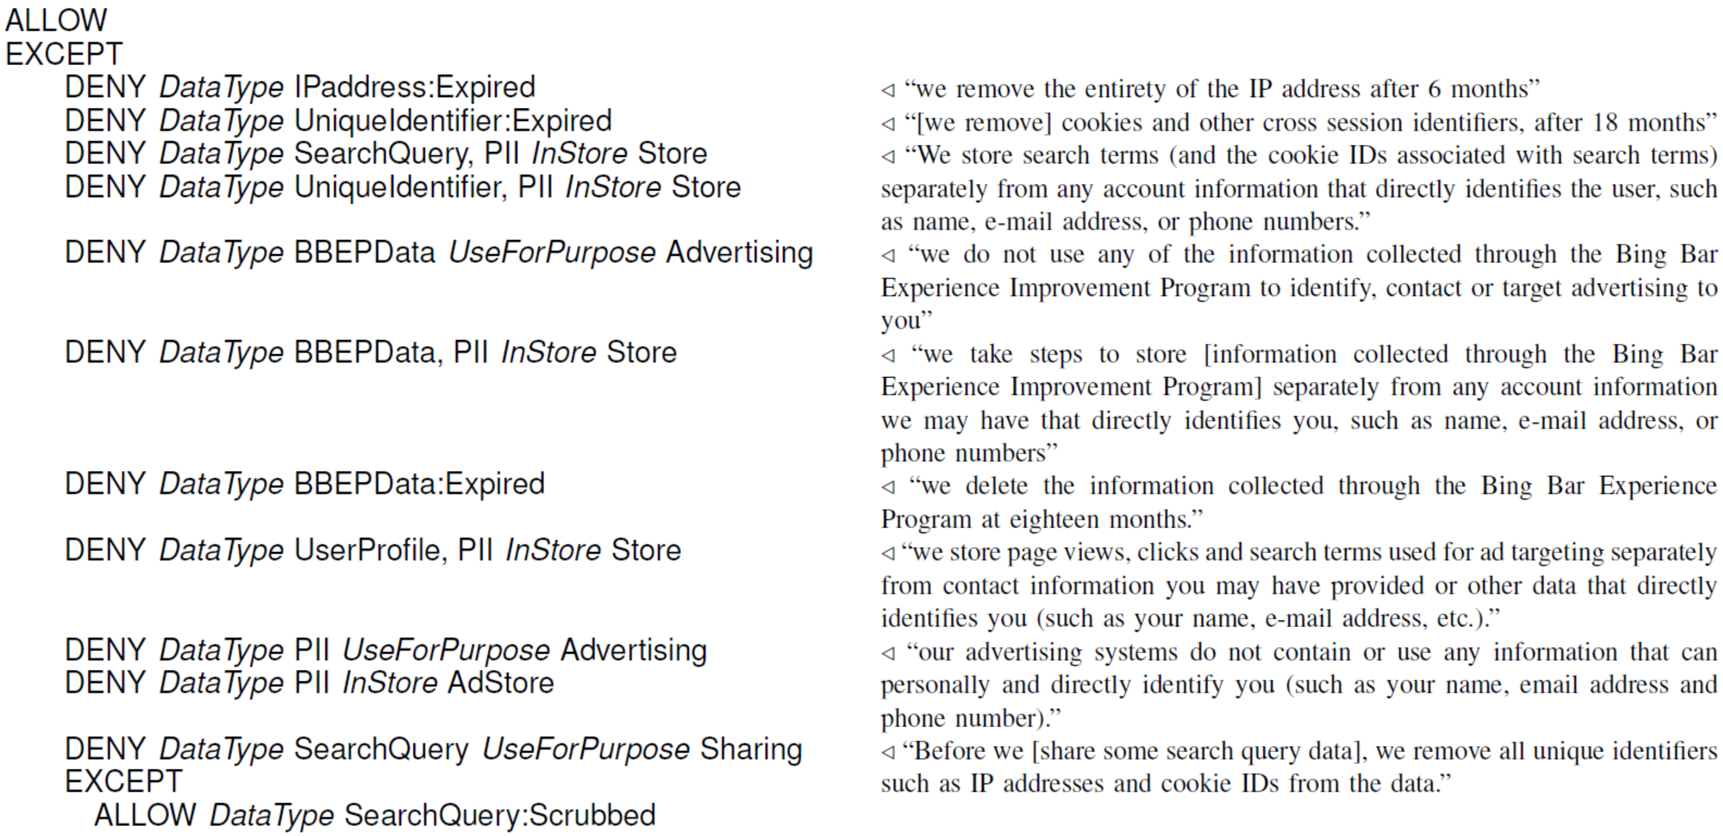

ALLOW	
EXCEPT	
DENY *DataType* IPaddress:Expired	◁ “we remove the entirety of the IP address after 6 months”
DENY *DataType* UniqueIdentifier:Expired	◁ “[we remove] cookies and other cross session identifiers, after 18 months”
DENY *DataType* SearchQuery, PII *InStore* Store DENY *DataType* UniqueIdentifier, PII *InStore* Store	◁ “We store search terms (and the cookie IDs associated with search terms) separately from any account information that directly identifies the user, such as name, e-mail address, or phone numbers.”
DENY *DataType* BBEPData *UseForPurpose* Advertising	◁ “we do not use any of the information collected through the Bing Bar Experience Improvement Program to identify, contact or target advertising to you”
DENY *DataType* BBEPData, PII *InStore* Store	◁ “we take steps to store [information collected through the Bing Bar Experience Improvement Program] separately from any account information we may have that directly identifies you, such as name, e-mail address, or phone numbers”
DENY *DataType* BBEPData:Expired	◁ “we delete the information collected through the Bing Bar Experience Program at eighteen months.”
DENY *DataType* UserProfile, PII *InStore* Store	◁ “we store page views, clicks and search terms used for ad targeting separately from contact information you may have provided or other data that directly identifies you (such as your name, e-mail address, etc.).”
DENY *DataType* PII *UseForPurpose* Advertising DENY *DataType* PII *InStore* AdStore	◁ “our advertising systems do not contain or use any information that can personally and directly identify you (such as your name, email address and phone number).”
DENY *DataType* SearchQuery *UseForPurpose* Sharing EXCEPT ALLOW *DataType* SearchQuery:Scrubbed	◁ “Before we [share some search query data], we remove all unique identifiers such as IP addresses and cookie IDs from the data.”