

FlowFence: Practical Data Protection for Emerging IoT Application Frameworks

Earlence Fernandes, Justin Paupore, Amir Rahmati, Daniel Simionato,
Mauro Conti, Atul Prakash

University of Michigan, University of Padova

Published at USENIX Security 2016



Wearables/Quantified Self



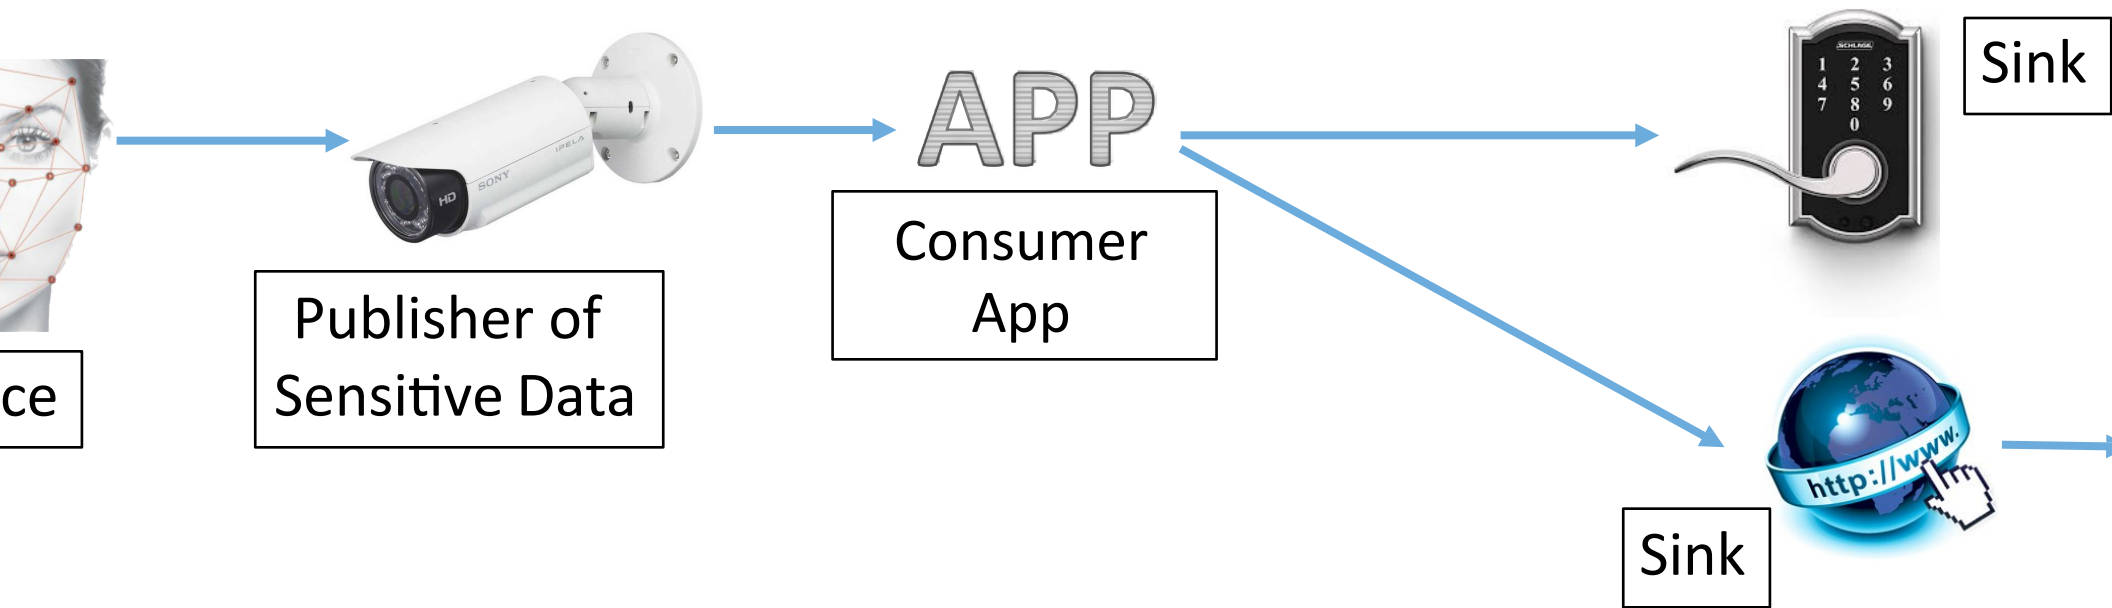
Connected Healthcare



Smart Homes



Emerging IoT App Frameworks



unlock door if face is recognized
home-owner can check activity
from Internet

- App needs to compute sensitive data to provide useful service
- But has the potential to leak data

Existing IoT frameworks have permission based access control



Smart home API

[Smart Homes]



Google Fit API

[Wearables]



Android Sensor API

[Quantified Self]

e.g., capability.lockCodes in SmartThings (pincode)
FITNESS_BODY_READ scope in Google Fit (heart rate)

- Permissions control what data an app can access
- Permissions do not control how apps use data, once they have access

FlowFence

Flow-control is a first-class primitive

OS-based flow control

Component-level information tracking
Flow enforcement through label policies



Language-based flow control

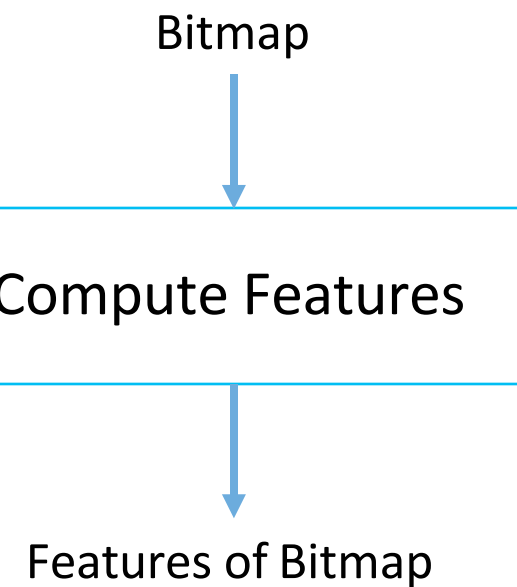
- Restructure apps to obey flow
- Developer declares flows



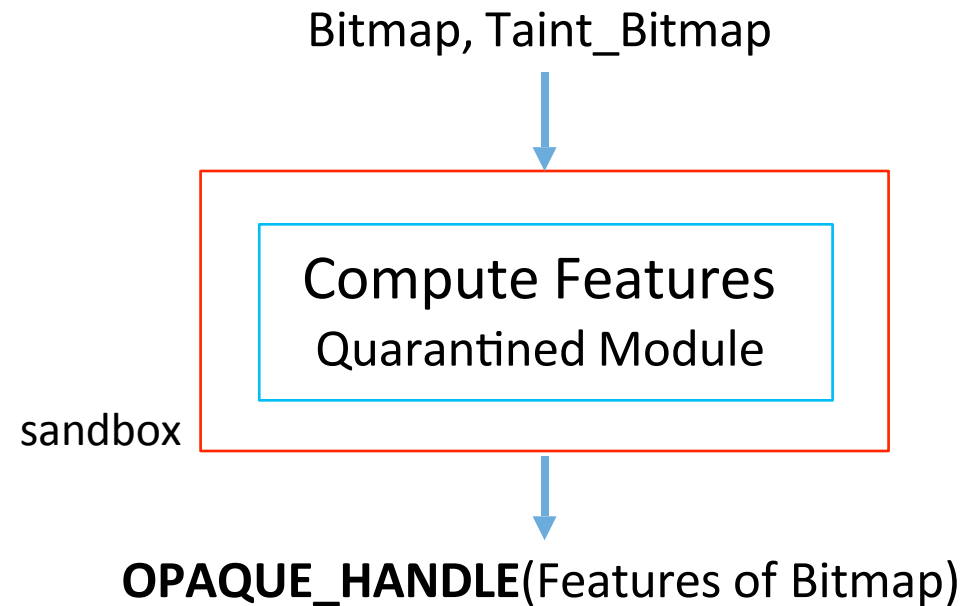
FlowFence

- Support of diverse publishers and consumers of data, with publisher and consumer flow policies
- Allows use of existing languages, tools, and OSes

Quarantined Modules and Opaque Handles

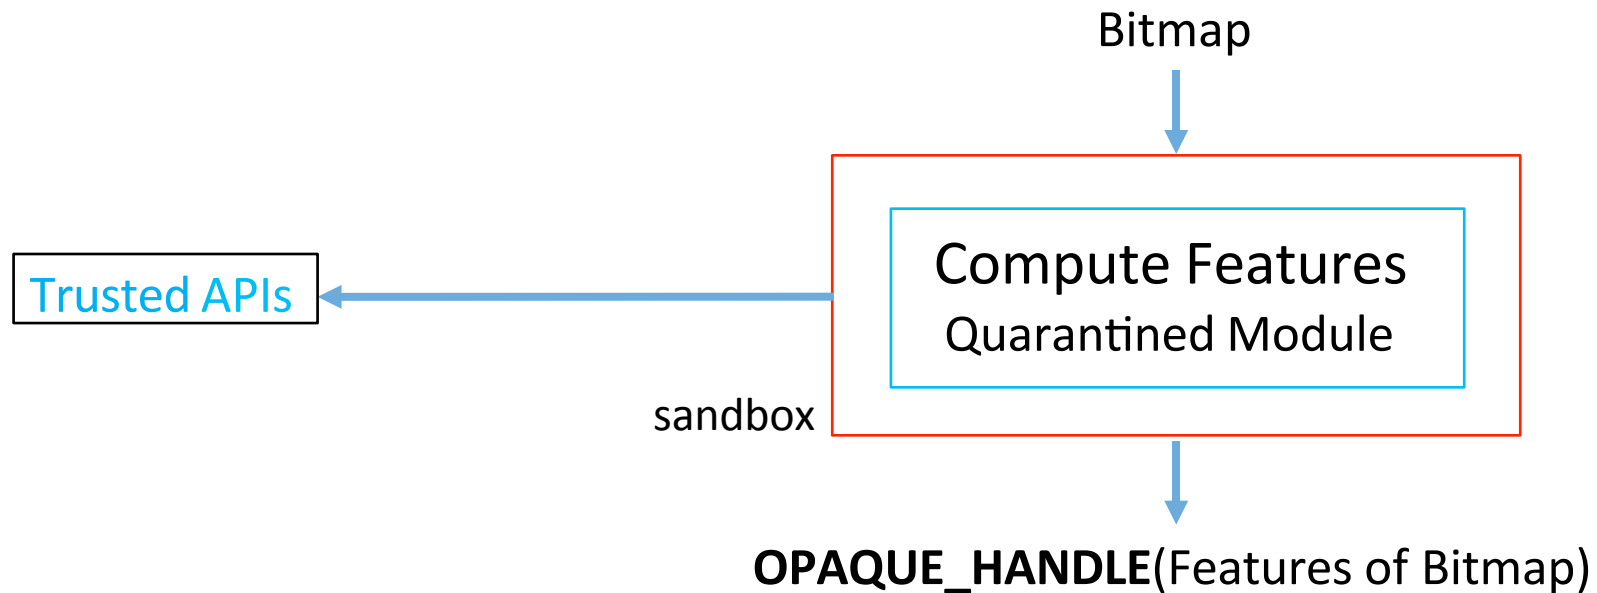


The computation runs with the rights to access sensitive bitmap data



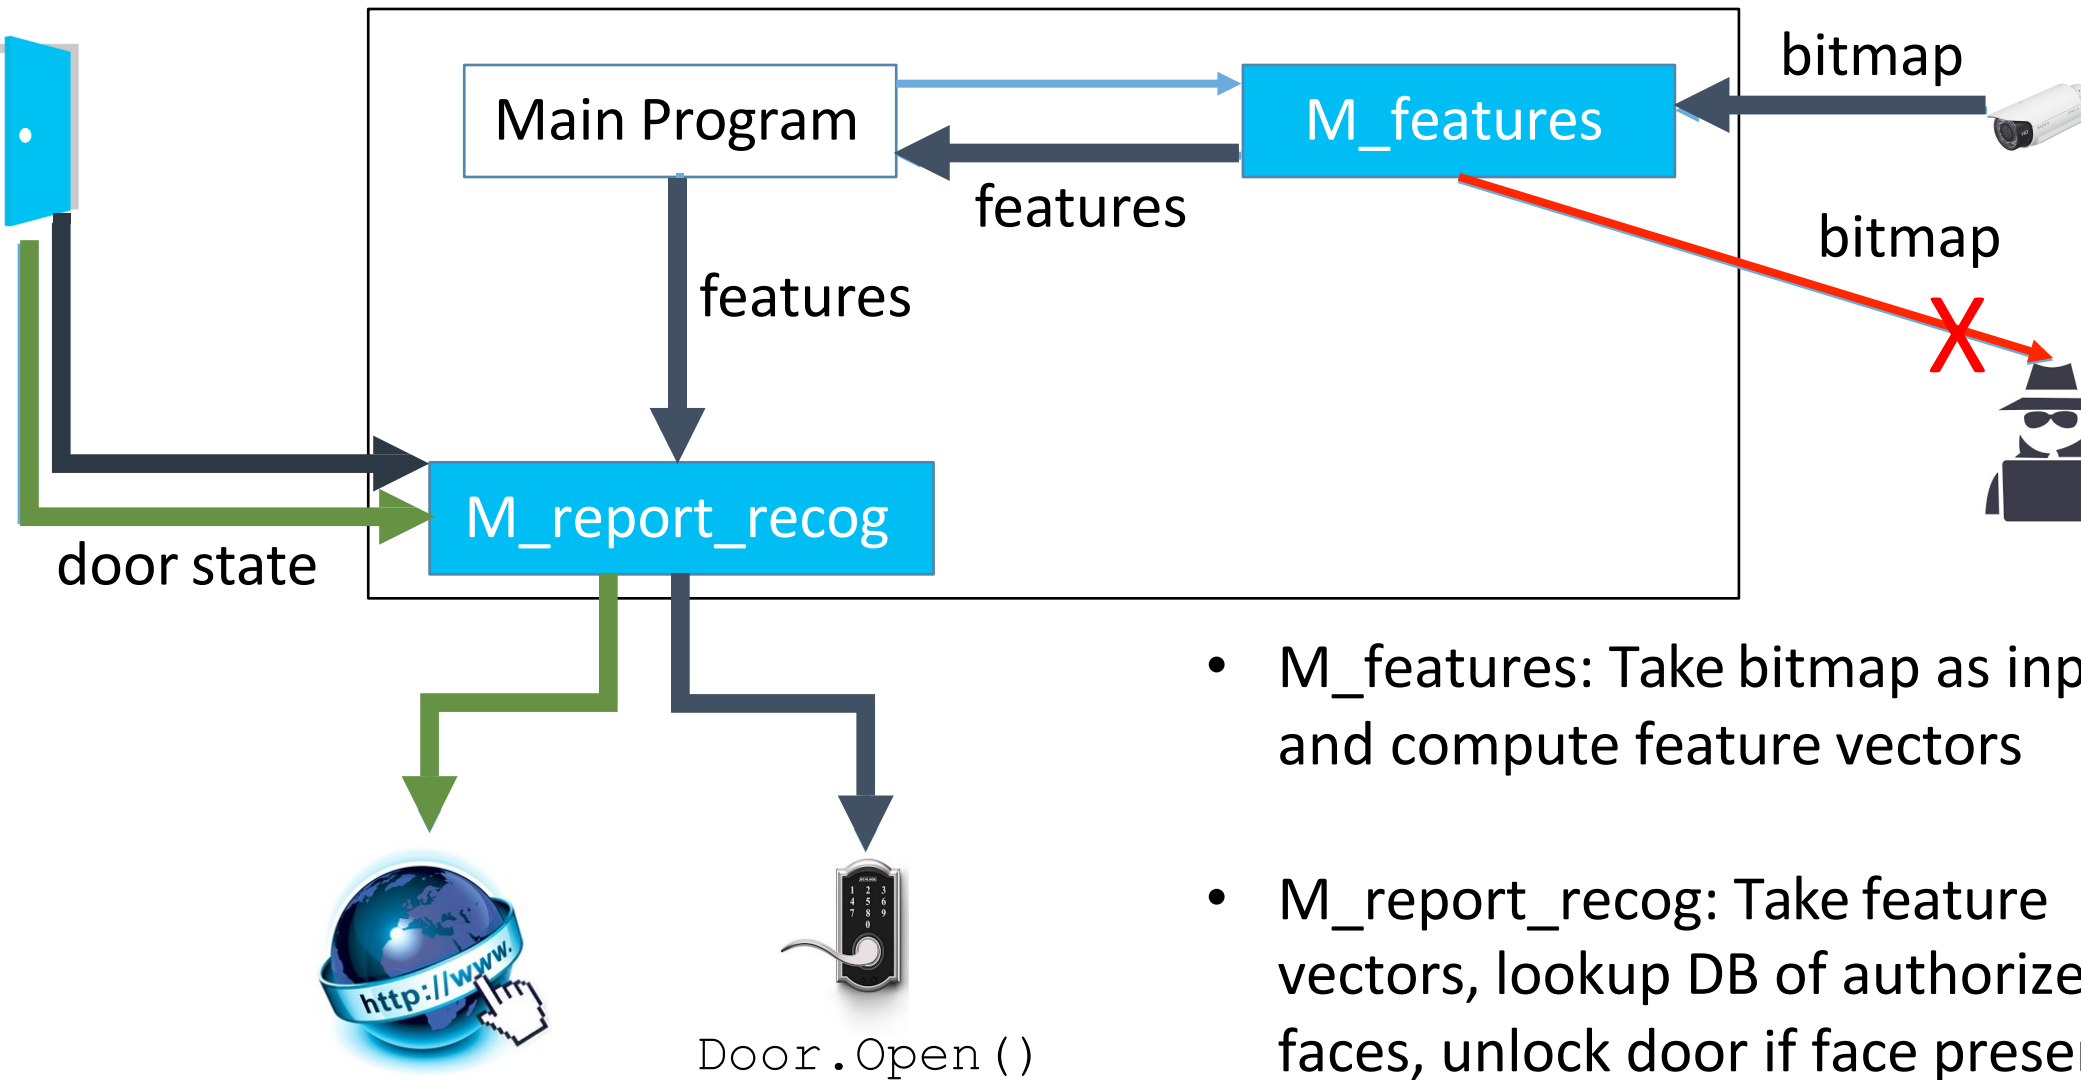
- Submit a computation that runs in a sandbox
- All sensitive data is available only in sandbox

Quarantined Modules and Opaque Handles



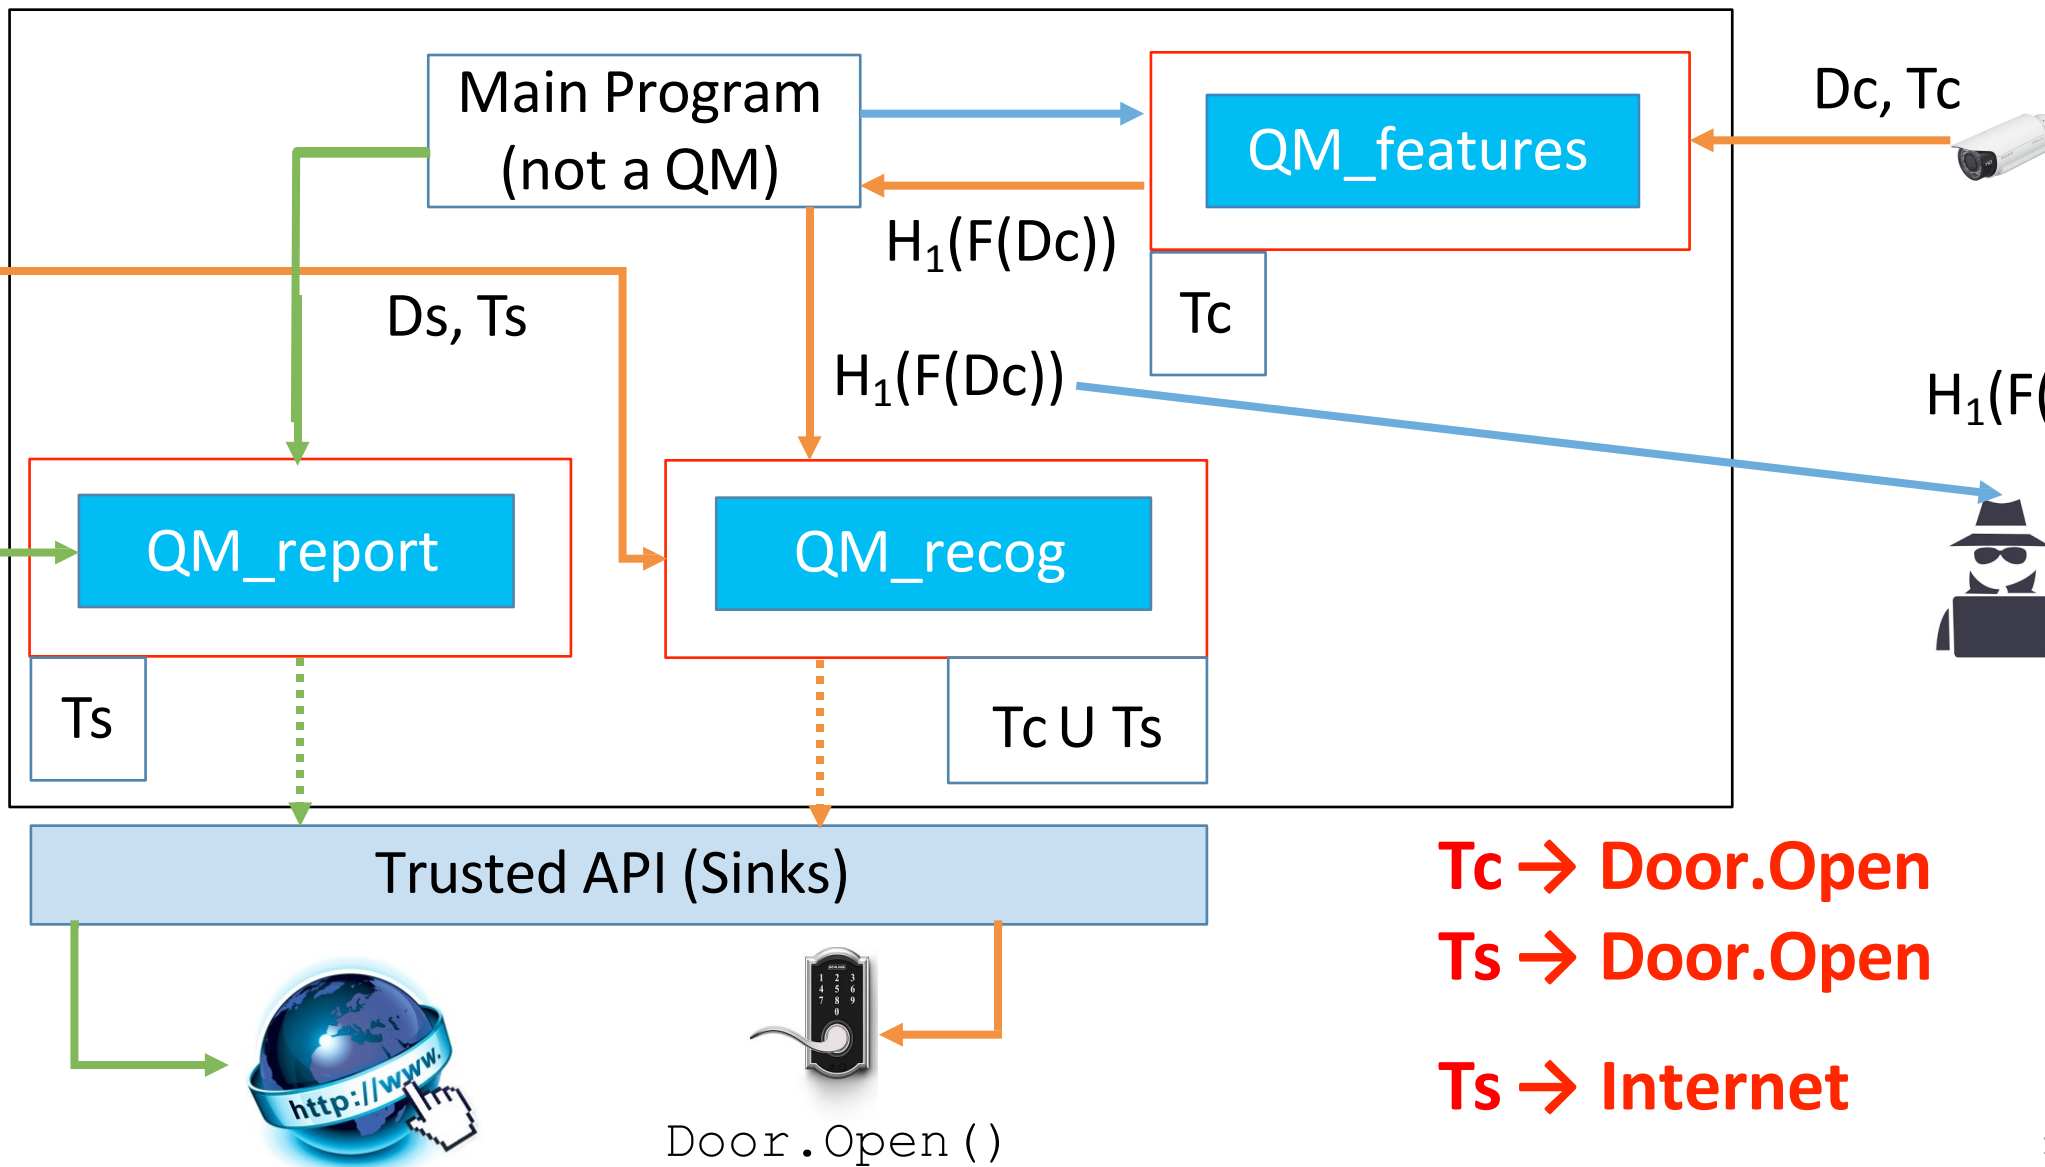
Quarantined Modules can also access FlowFence-provided Trusted APIs
Trusted APIs check the taint labels of the caller against a flow policy

Face Recognition App Example

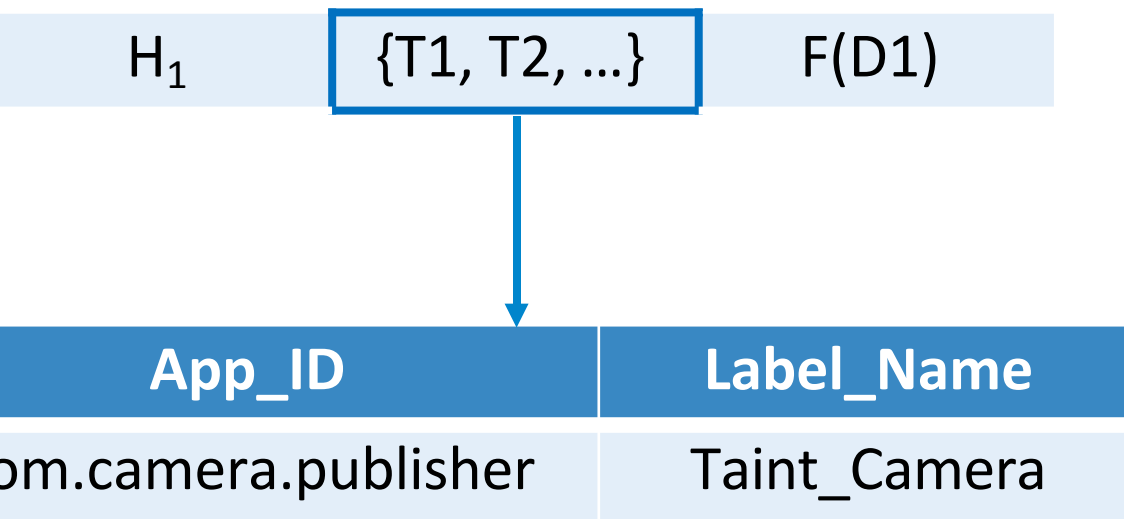


- **M_features**: Take bitmap as input and compute feature vectors
- **M_report_recog**: Take feature vectors, lookup DB of authorized faces, unlock door if face present
Report door state

FlowFence – Refactored App



Taint Labels and Flow Policies



Example Policy

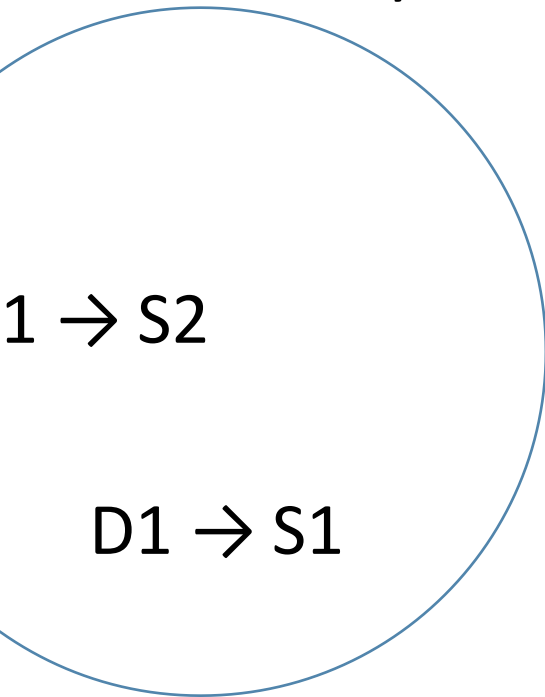
```
{  
  Taint_Camera → UI,  
  Taint_HeartR → Internet  
}
```

App_ID – unique application identifier on the underlying OS

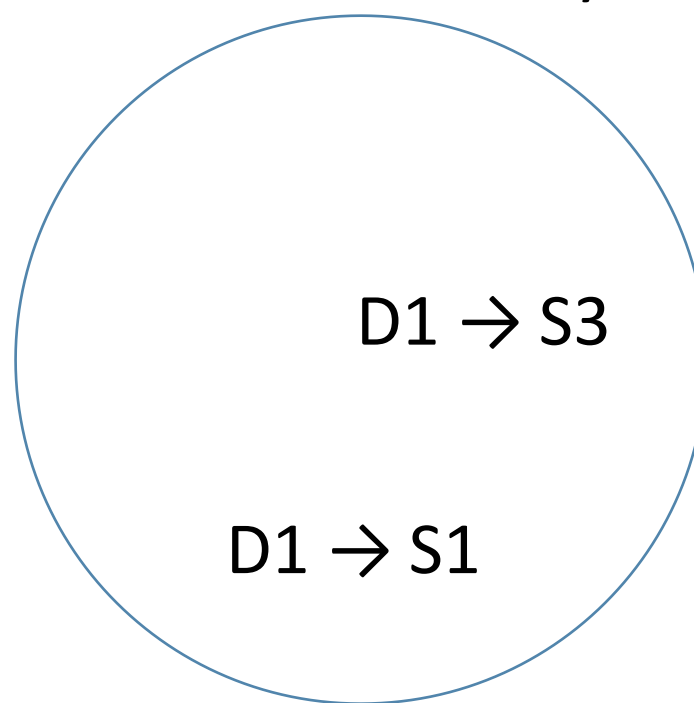
Label_Name – well-known string that identifies the type of data

Publisher and Consumer Flow Policies

Publisher Policy



Consumer Policy




```
{ Publisher;  
  Taint_Camera → UI  
}
```

```
{ Consumer;  
  Taint_Camera → Door.C  
  Taint_DoorState → Doc  
  Taint_DoorState → Inte  
}
```

Publisher and Consumer Flow Policies

Publisher Policy



Consumer Policy

Prompt



Automatically Approved

```
{ Publisher;  
  Taint_Camera → UI  
}
```

```
{ Consumer;  
  Taint_Camera → Door.C  
  Taint_DoorState → Doc  
  Taint_DoorState → Inte  
}
```

Data Sharing Mechanisms in Current IoT Frameworks



Smart home API

[Smart Homes]



Google Fit API

[Wearables]



Android Sensor API

[Quantified Self]

- **Polling Interface**
 - App checks for new data
- **Callback Interface**
 - App is called when new data available
- **Device Independence**
 - E.g., many types of heart rate sensors produce “heart beat” data
 - Consumers should only need to **specify “what”** data they want, **without specifying “how”**

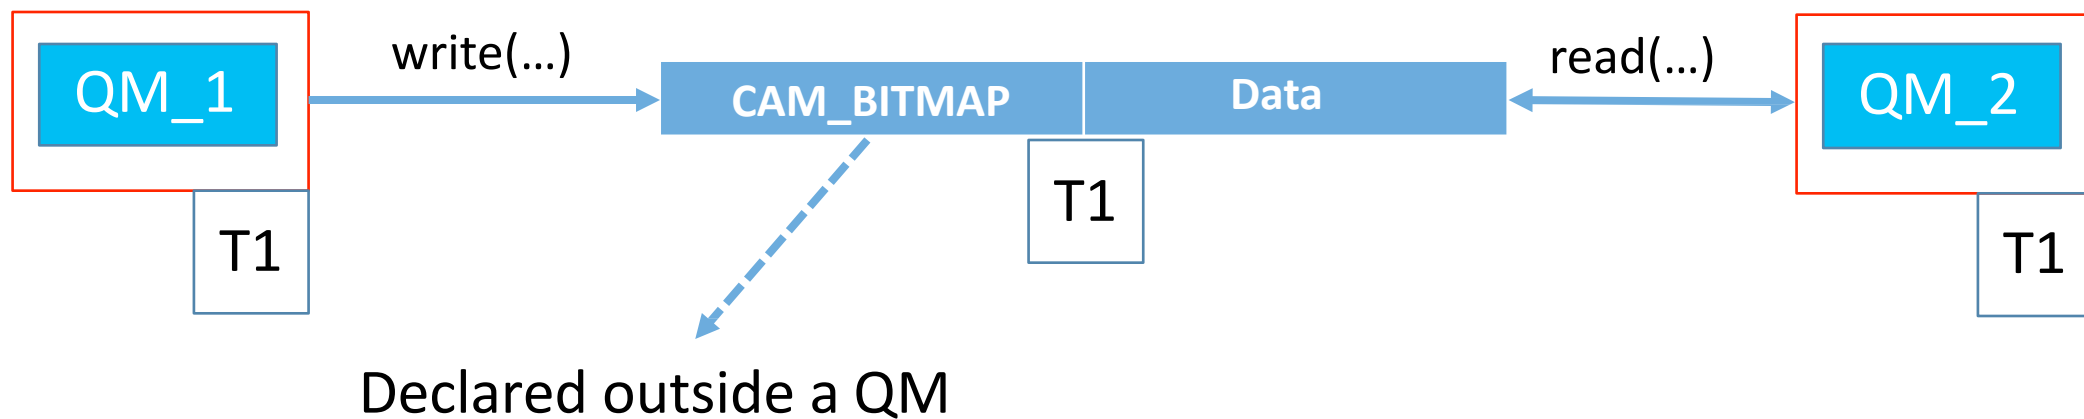
Key-Value Store – Polling Interface/Device Independence

Each app gets a single Key-Value Store

An app can only write to its own Key-Value Store

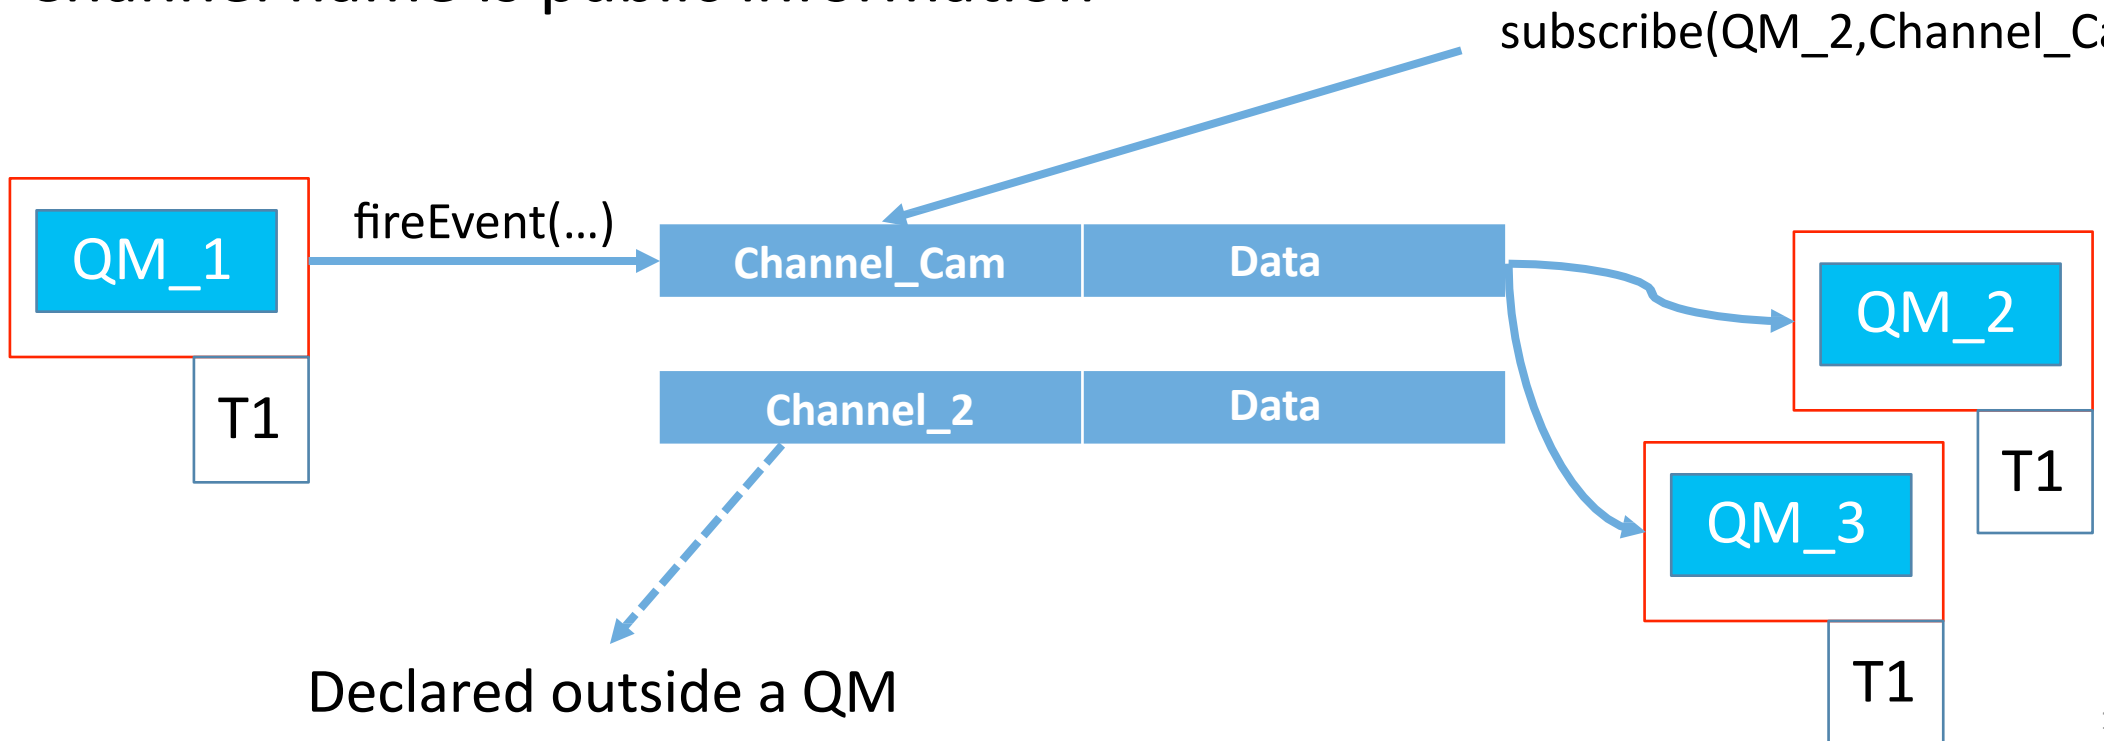
Apps can read from any Key-Value Store

Keys are public information because consumers need to know about the



Event Channels – Callback Interface/Device Independence

Apps can declare statically in code, their intended channels
Only the owner of a channel can fire an event
Channel name is public information



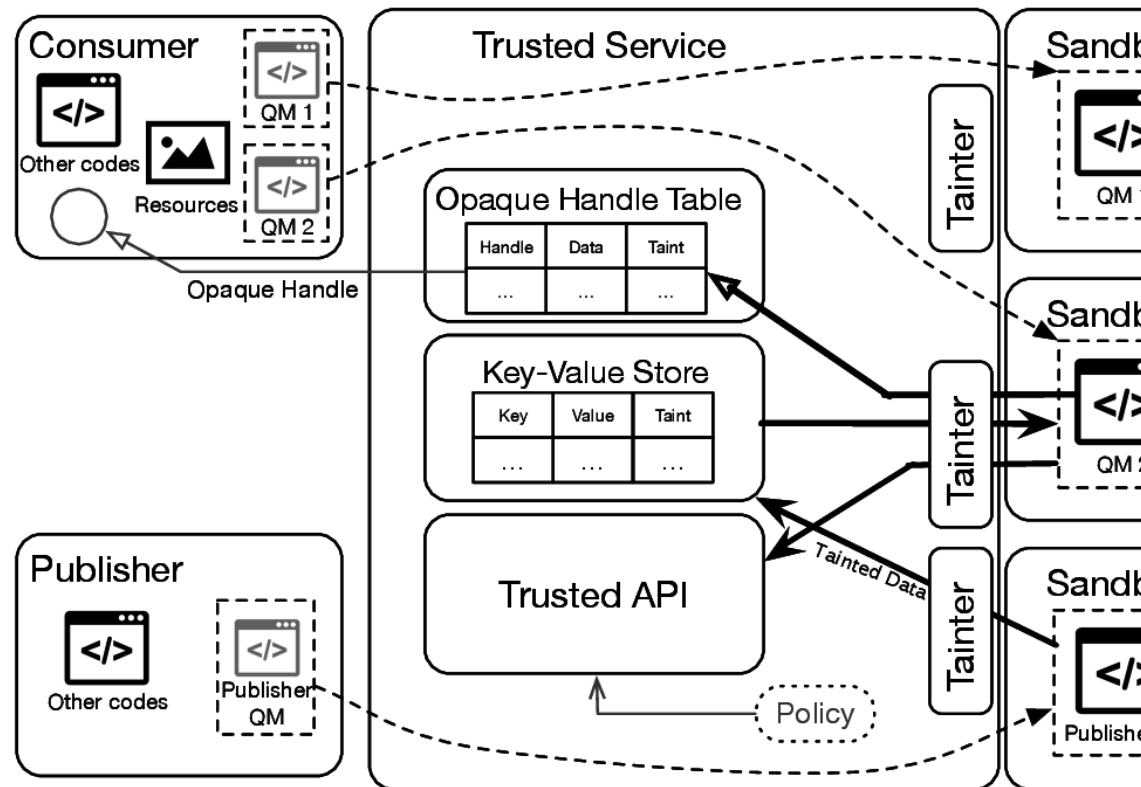
FlowFence Implementation

IoT Architectures

- Cloud
- Hub



- `isolatedProcess = true` for sandboxes
- Supports native code



Evaluation Overview

What is the overhead on a micro-level in terms of computation and memory?

Per-Sandbox Memory Overhead	2.7 MB
QM Call Latency	92 ms
Data Transfer b/w into Sandbox	31.5 MB/s

Comparable to IoT device ops over area-network, e.g., Nest, SmartThings

Nest cam peak bandwidth is 1.2 M

Can FlowFence support real IoT apps securely?

Ported 3 Existing IoT Apps: SmartLights, FaceDoor, HeartRateMonitor

Required adding less than 140 lines per app; FlowFence isolates flows

What is the impact of FlowFence on macro-performance?

FaceDoor Recognition Latency	5% average increase
HeartRateMonitor Throughput	0.2 fps reduction on average
SmartLights end-to-end latency	+110 ms on average

Porting IoT Apps to FlowFence

App	Data Security Risk	Original LoC	FlowFence LoC	Flow Request
SmartLights	Can leak location information	118	193	Loc → Switch
FaceDoor	Can leak images of people	322	456	Cam → Lock, Doorstate → Lock Doorstate → Net
HeartRateMon	Can leak images and heart rate	257	346	Cam → UI

SmartLights, FaceDoor – 2 days of porting effort each, HeartMon – 1 day of porting effort

Macro-performance of Ported Apps

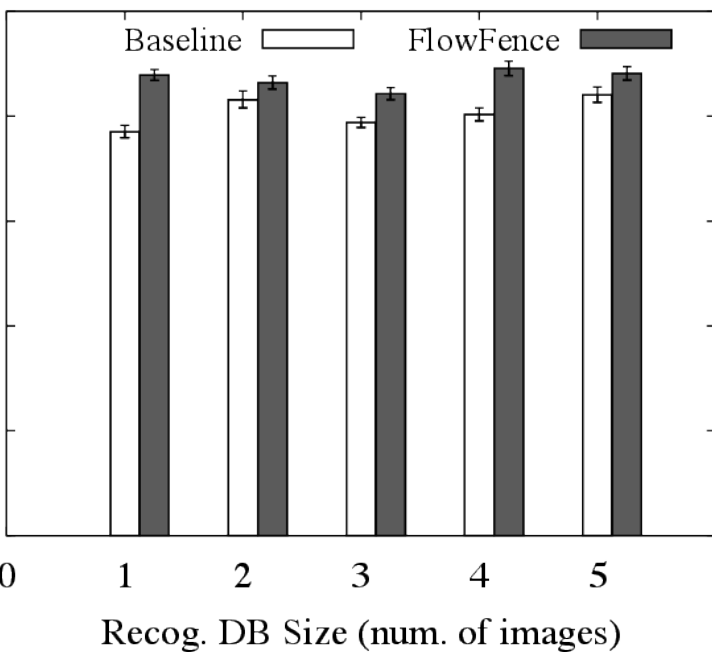
FaceDoor Enroll Latency

Baseline	811 ms (SD = 37.1)
FlowFence	937 ms (SD = 60.4)

SmartLights End-To-End Latency

Baseline	160 ms (SD = 69.9)
FlowFence	270 ms (SD = 96.1)

FaceDoor Recognition Latency (612x816 pixels)



HeartRateMon Throughput

Throughput w/o Image Processing	23.0 (SD=0.7) fps	22.9 (SD=0.7) fps
Throughput w/ Image Processing	22.9 (SD=0.7) fps	22.7 (SD=0.7) fps

Summary

Emerging IoT App Frameworks only support permission-based access control
Malicious apps can steal sensitive data easily

FlowFence explicitly embeds control and data flows within app structure;
Developers must split their apps into:

- Set of communicating Quarantined Modules with the unit of communication being Opaque Handles – taint tracked, opaque refs to data
- Non-sensitive code that orchestrates QM execution

FlowFence supports publisher and consumer flow policies that enable building secure IoT apps

We ported 3 existing IoT apps in 5 days; Each app required adding < 140 LoC

Macro-performance tests on ported apps indicate FlowFence overhead is reasonable: e.g., 4.9% latency overhead to recog. a face & unlock a door

Discussion

What's the limitation of FlowFence?

How is the usability of FlowFence to developers and users?

How to improve FlowFence?

What makes protecting IoT challenging?

Is FlowFence able to mitigate the attacks we discussed in last class?

Instruction-Level Flow Analysis Techniques

Dynamic Taint Tracking

fine granularity
no developer effort
high computational overhead
may need special h/w for acceleration
implicit flows can leak information
limited OS/Language flexibility

IoT devices (and hubs) have constrained hardware

OS and Language Diversity;
[Supports Rapid Development]

Static Taint Tracking

fine granularity
no developer effort
implicit flows can leak information
explicit and async. code can leak information

Fundamental Trigger-Action
Nature of IoT apps = Lots of
async. code