Why Your Encrypted Database Is Not Secure

Paul Grubbs Tom Ristenpart
Vitaly Shmatikov

Outsourced Applications Today

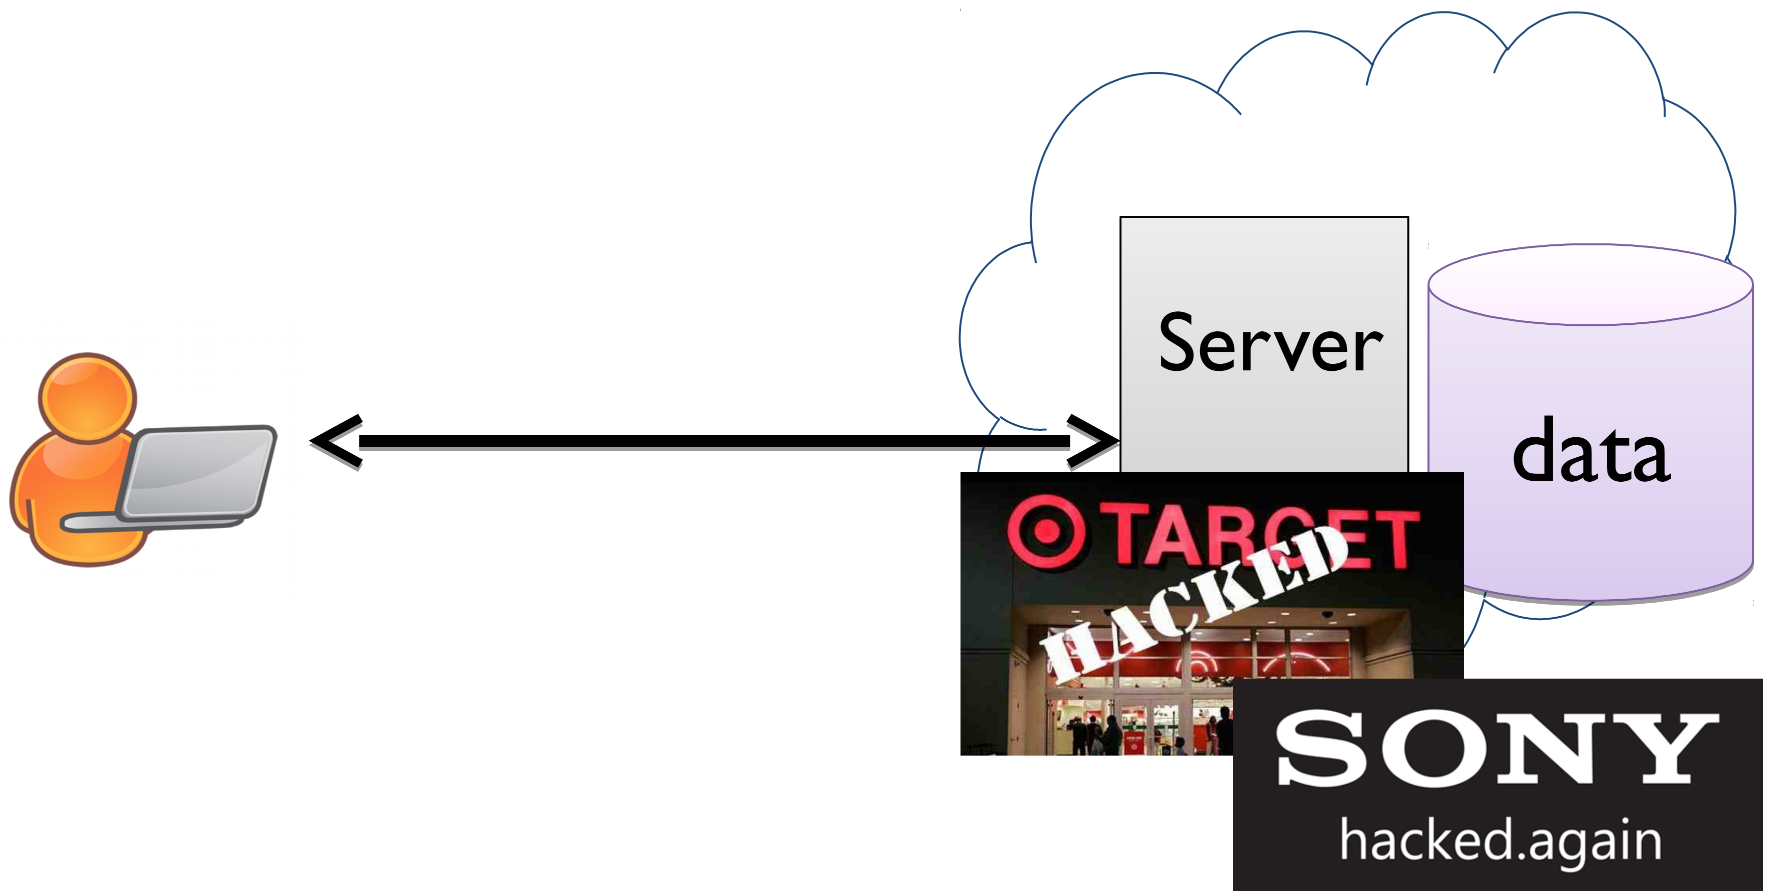

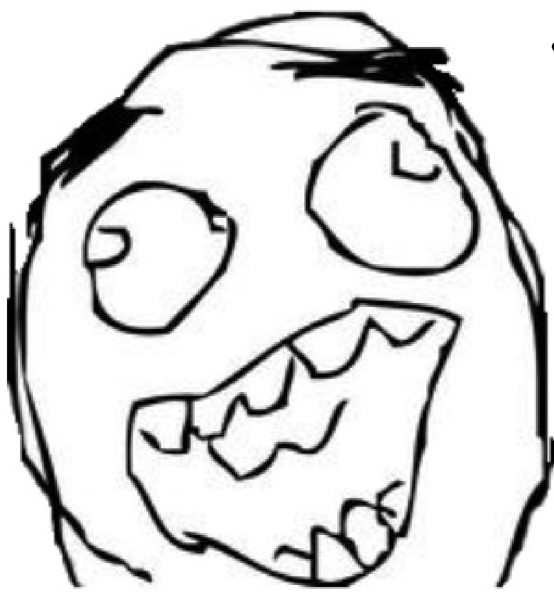

Encrypt
the data!

Encrypt the Data

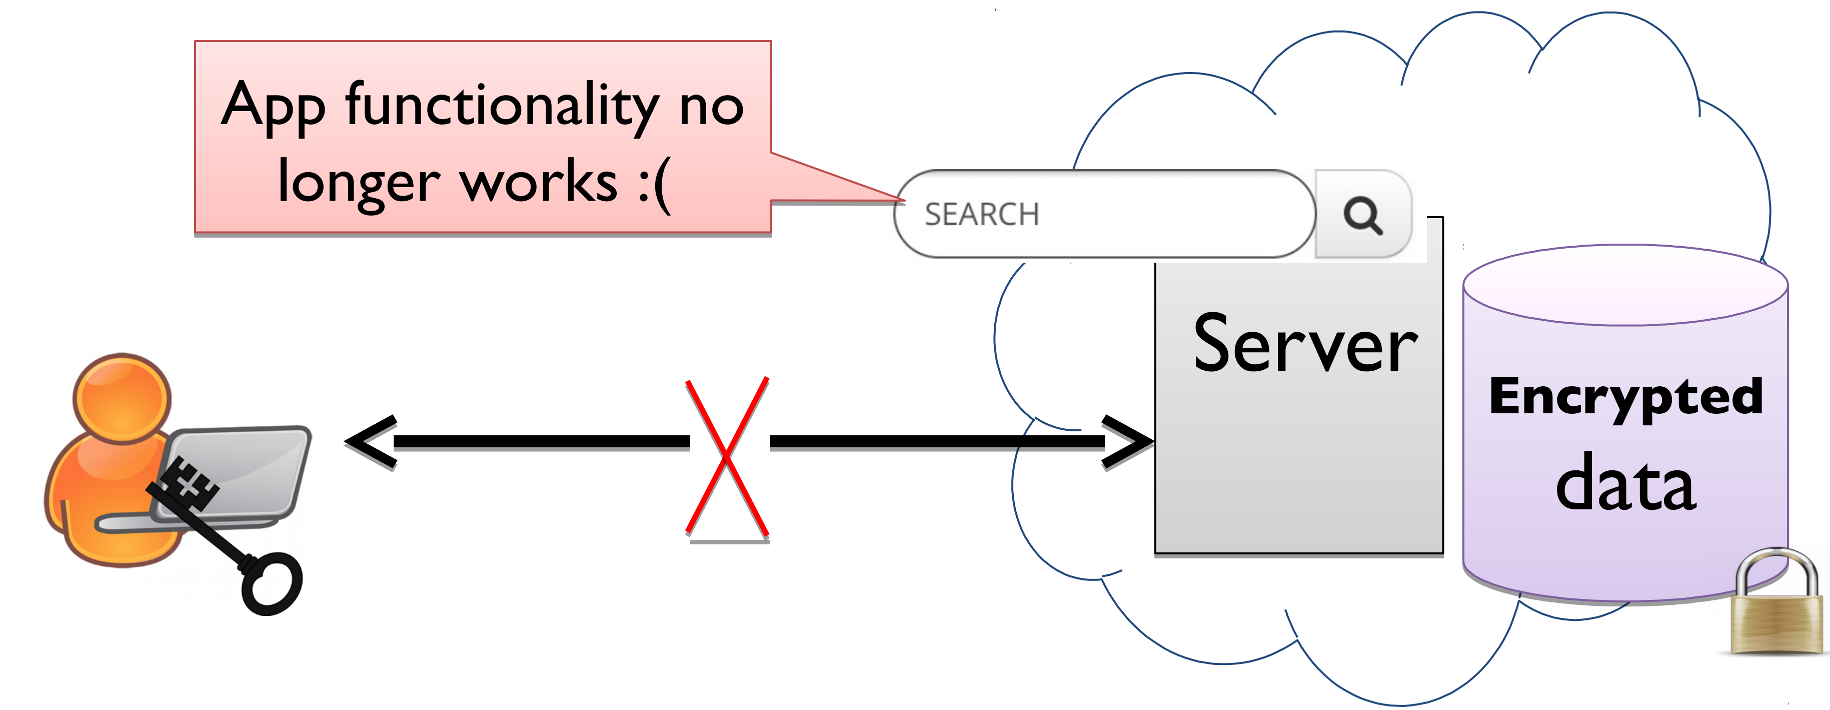

Encrypt the Data

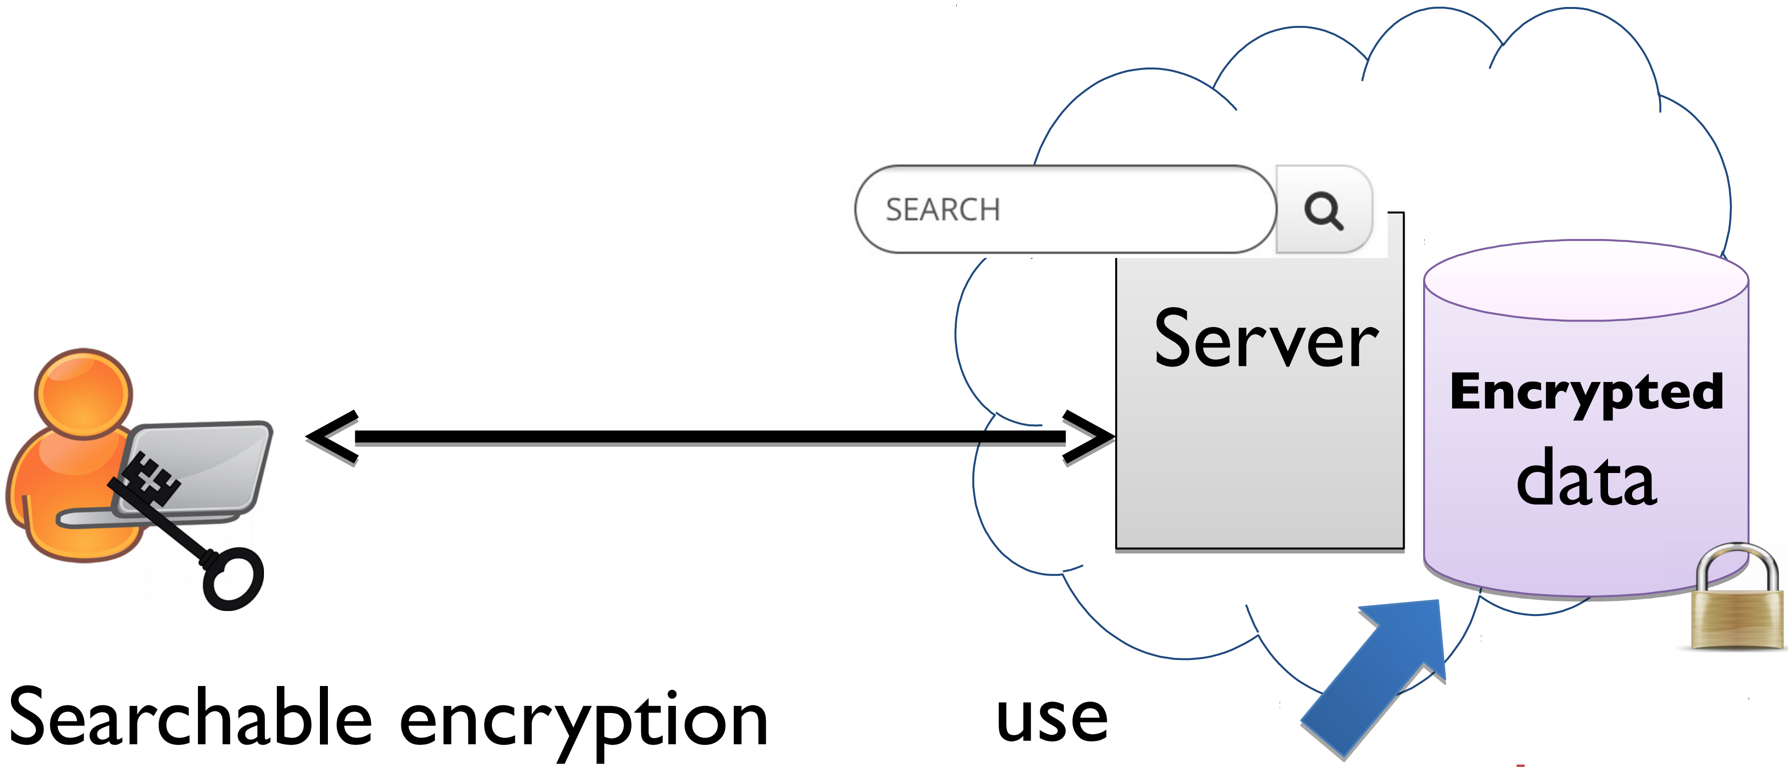

- Searchable encryption
- Deterministic encryption
- Order-revealing encryption

use **property-revealing encryption** (PRE)

Building "Secure" Systems

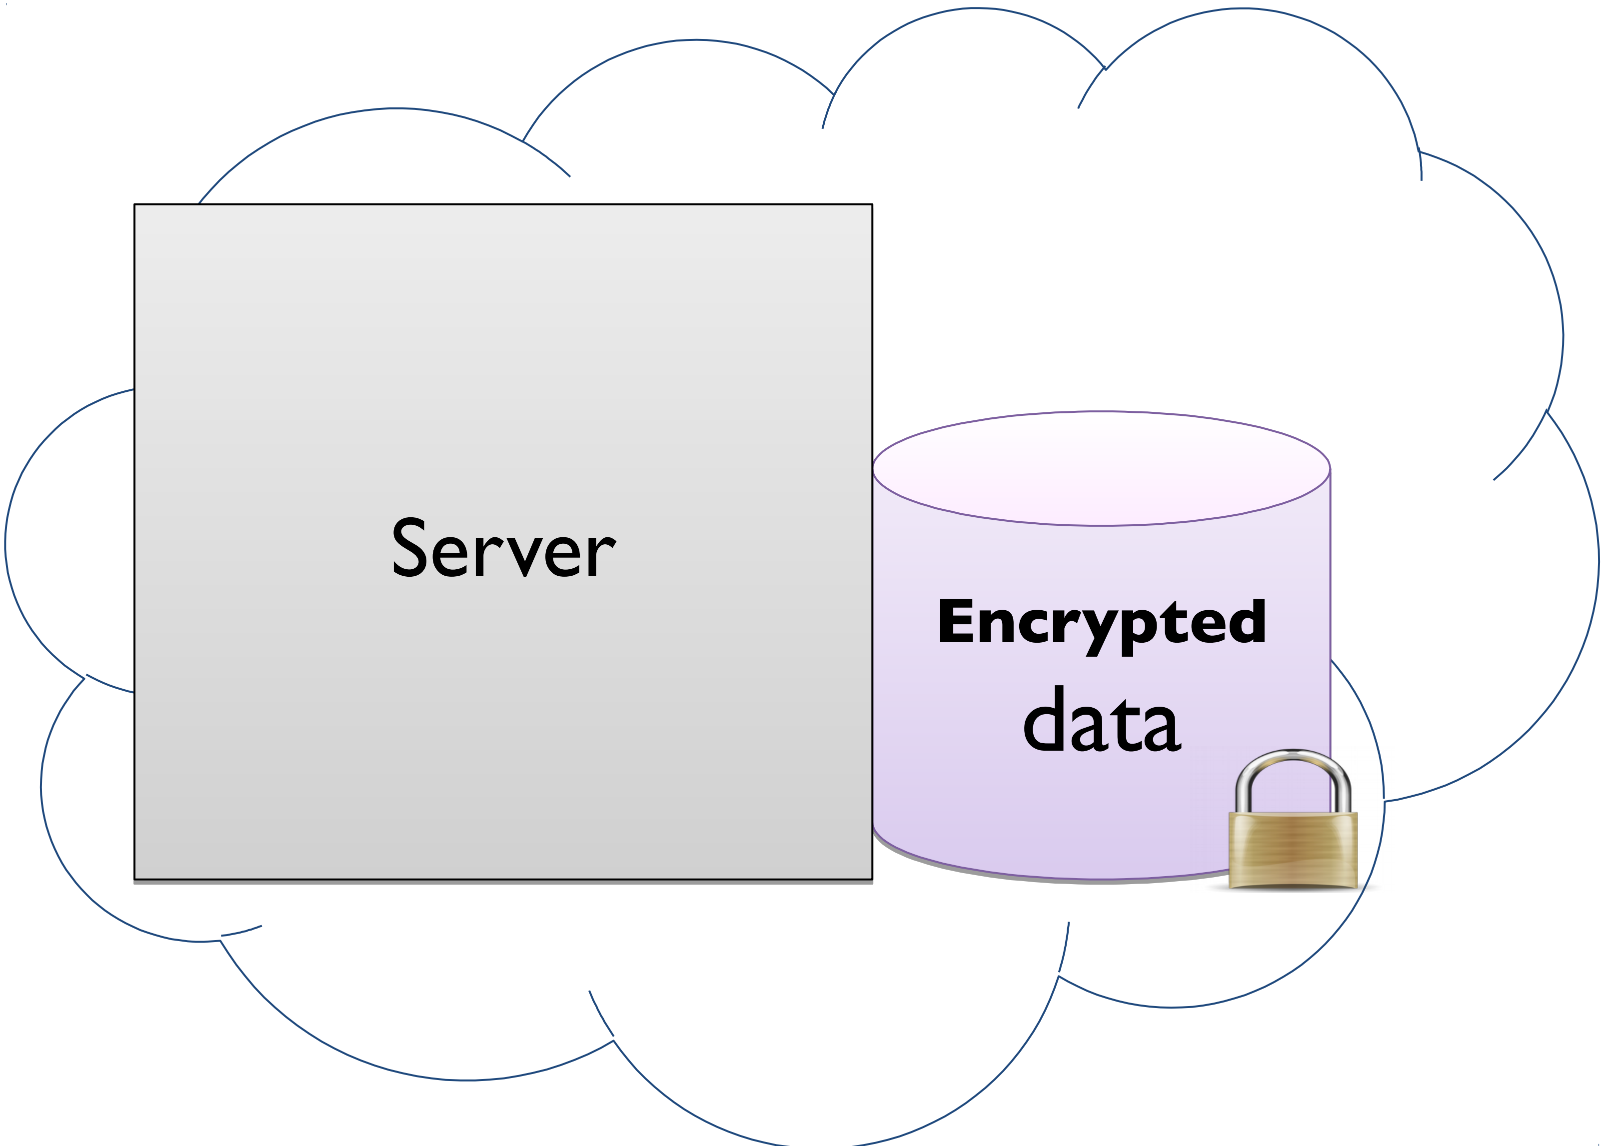

Building “Secure” Systems

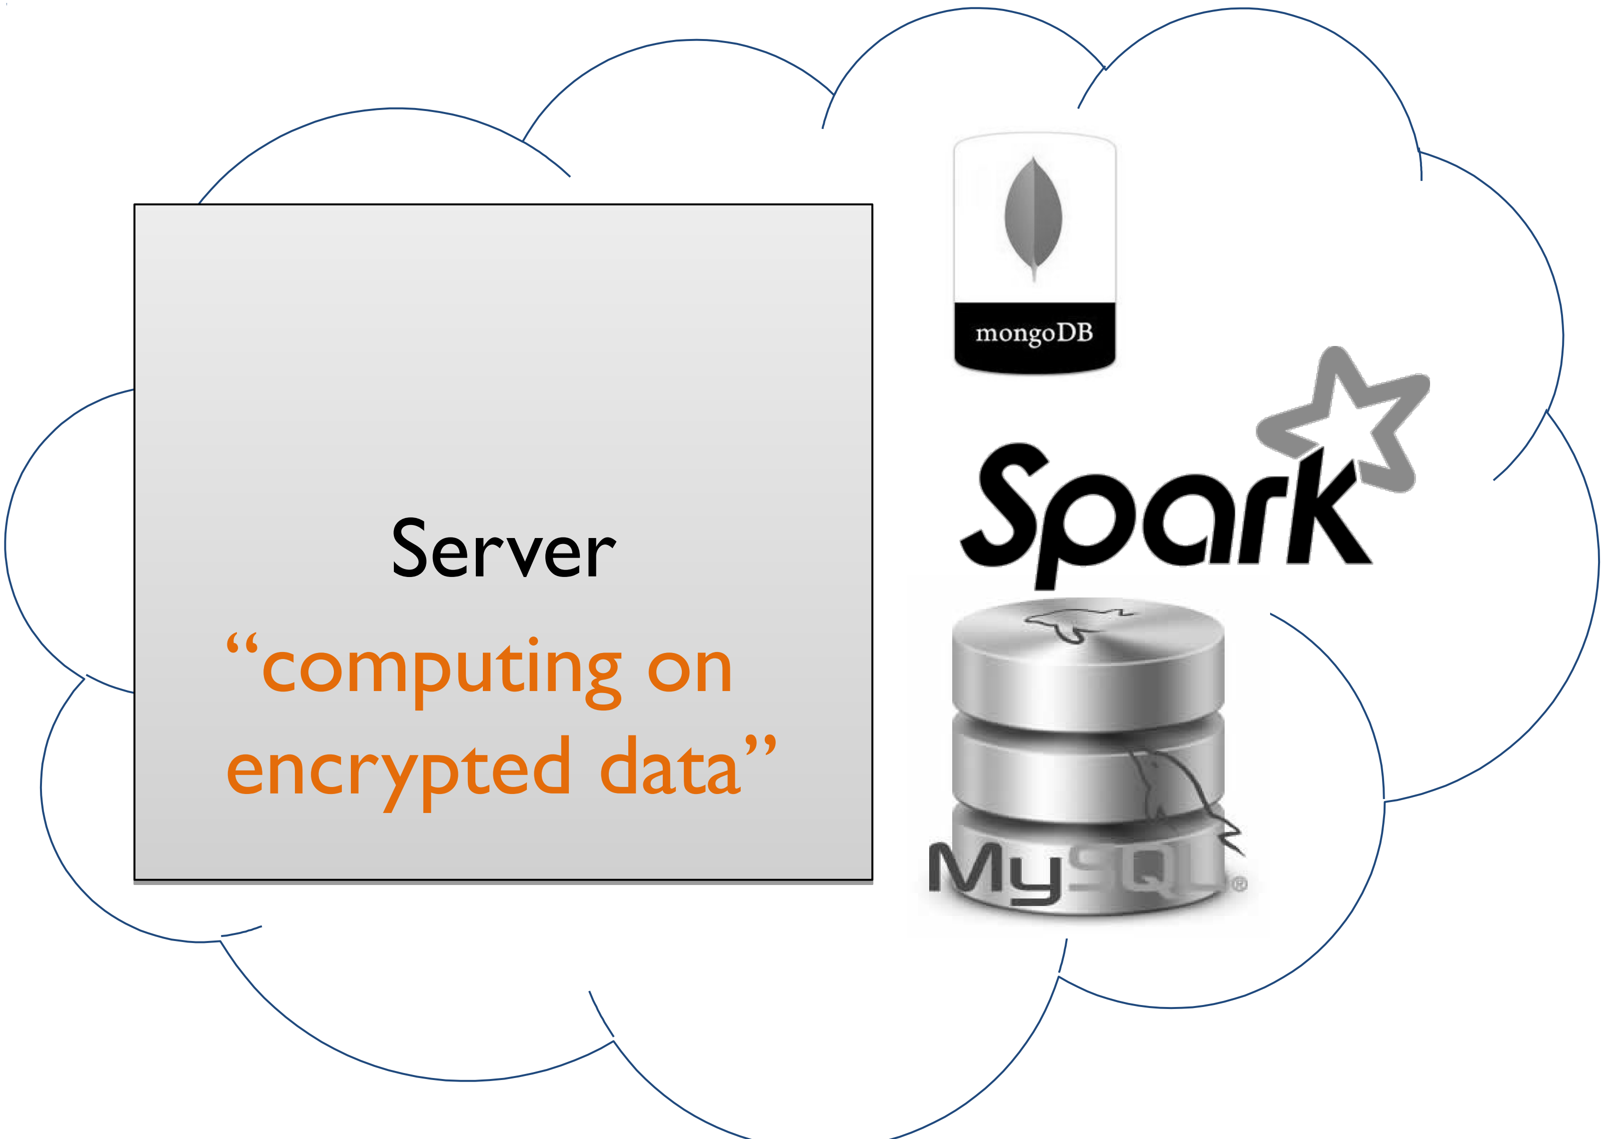

Building "Secure" Systems

- CryptDB (SOSP 2011)
- Mylar (NSDI 2014)
- Seabed (OSDI 2016)
- Arx
- Many others
- Lots of industry and government interest!!

What They Claim

CryptDB is a system that provides practical and provable confidentiality

Using the "sensitive" annotation, CryptDB ensures that, even if an attacker steals an encrypted database, the database does not leak the values of sensitive fields, even if the attacker has side information.

Mylar, a platform for building web applications, which protects data confidentiality against attackers with *full access to servers*.

the server's encrypted database provides semantic security,

strong security guarantees: it provides an IND-CPA-like security to the database, which reveals nothing beyond sizing information.

"Magically Flexible Cryptography"

CryptDB, on the other hand, manages to emulate fully homomorphic encryption for most of the functions of the SQL databases used in many applications, computing only with encrypted data and adding just 15% to 26% to those applications' computing time.

Claims

em ... n

the

of

has side information

Fallacy #1

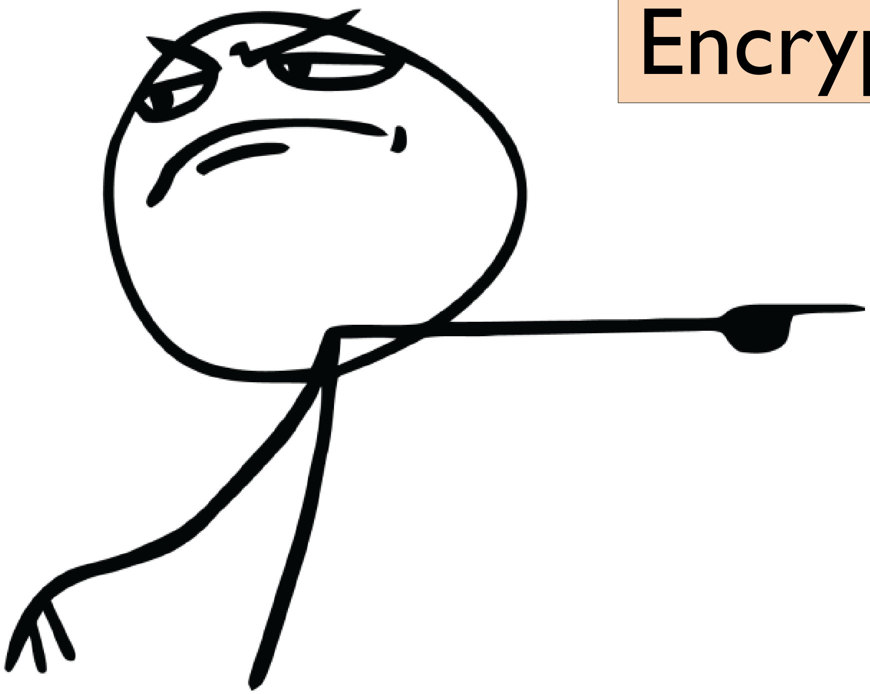

Encryption scheme is "secure"

does not mean

The system is "secure"

What This Talk Is About

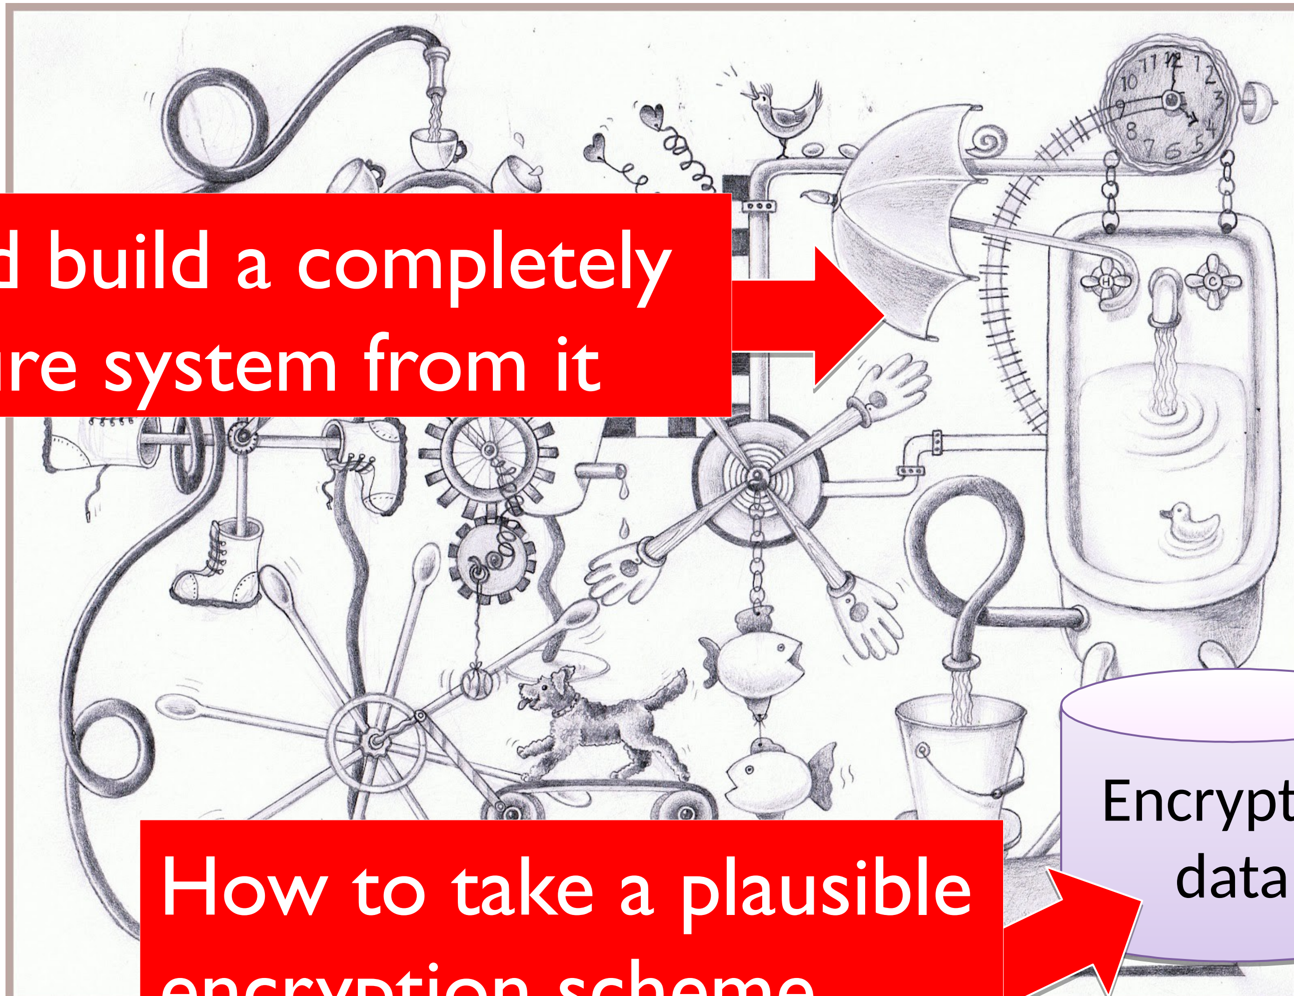

… and build a completely insecure system from it

How to take a plausible encryption scheme

Encrypted data

Unsafe at Any Speed

- CryptDB (SOSP 2011)
- Mylar (NSDI 2014)
- Seabed (OSDI 2016)
- Arx
- Many others
- Lots of industry d
government intere !

If you look at an actual commodity DBMS …

mongoDB

Spark

… insecure under <u>ANY</u> real-world attack

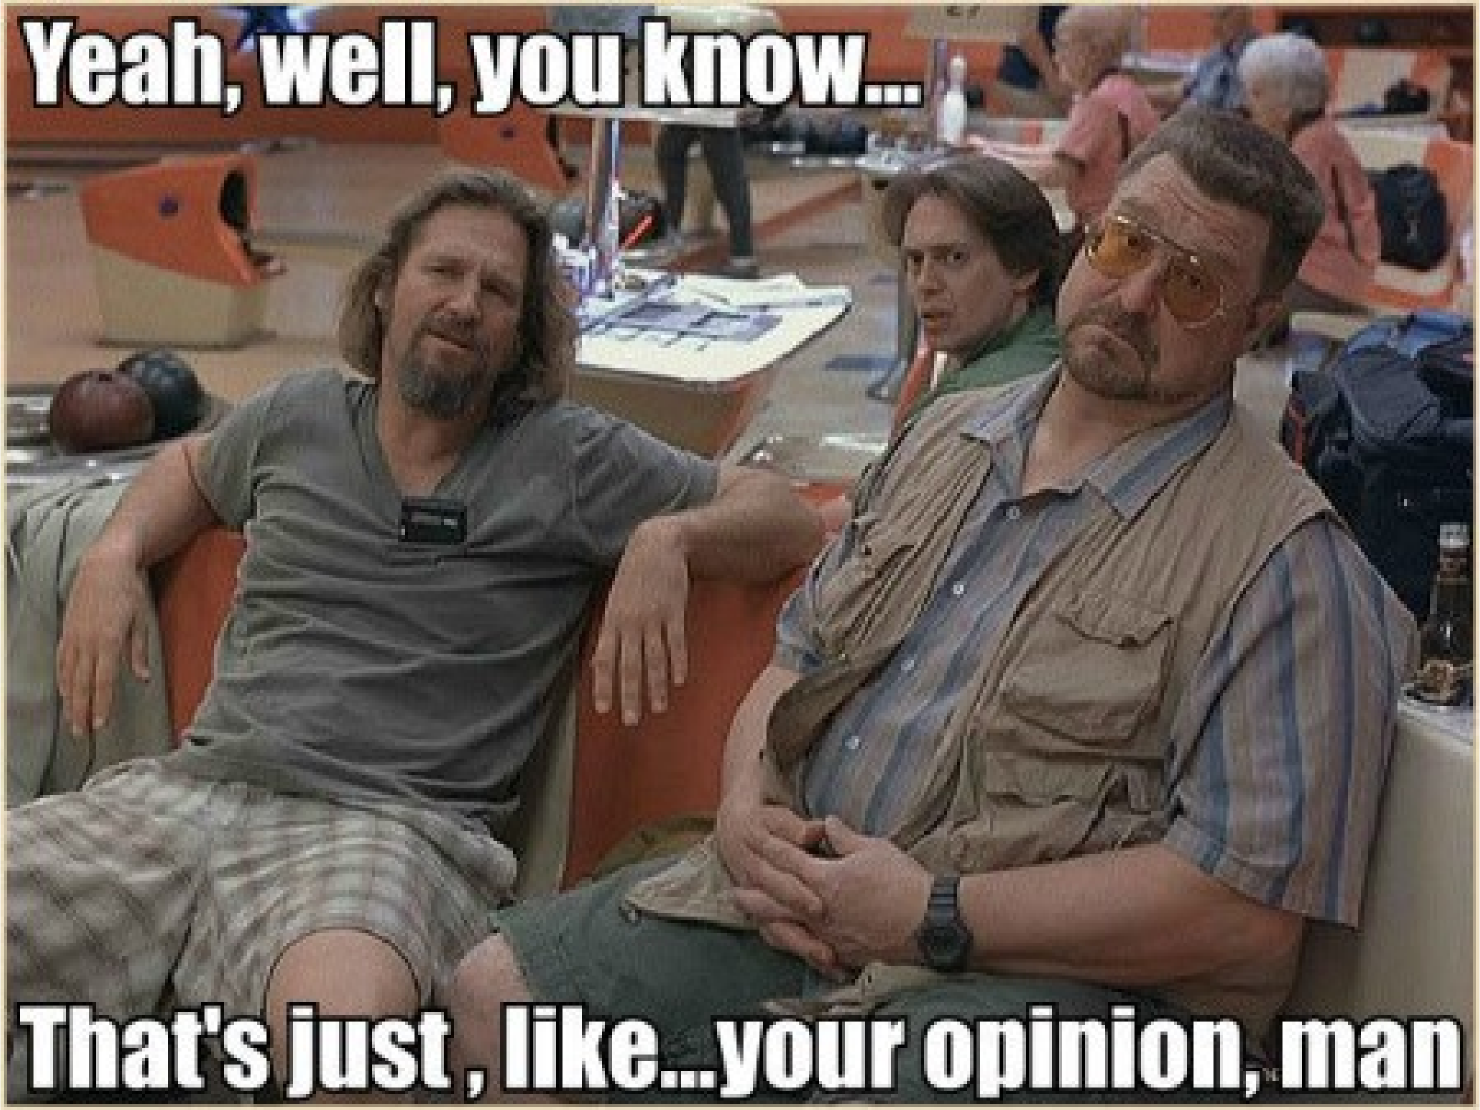
Yeah, well, you know...
That's just , like...your opinion, man

Threat Models

“Snapshot”

Persistent
passive

Active

Claims Meet Reality

- Secure against active attacks: false
  - Grubbs et al. "Breaking web applications built on top of encrypted data" (CCS 2016)
- Secure against "snapshot" attacks: false
  - Grubbs et al. "Why your encrypted database is not secure" (HotOS 2017)
- Sensitivity analysis helps: false
  - Bindschaedler et al. "The tao of inference in privacy-protected databases" (forthcoming)

Security Against Active Attacks

3.4 Threat model

Threats. Both the application and the database servers can be *fully* controlled by an adversary: the adversary may obtain all data from the server, cause the server to send arbitrary responses to web browsers, etc. This model subsumes a wide range of real-world security problems, from bugs in server software to insider attacks.

Mylar also allows some user machines to be controlled by the adversary, and to collude with the server. This may be either because the adversary is a user of the application, or because the adversary broke into a user's machine.

We call this adversary *active*, in contrast to a *passive* adversary that eavesdrops on all information at the server, but does not make any changes, so that the server responds to all client requests as if it were not compromised.

Mylar

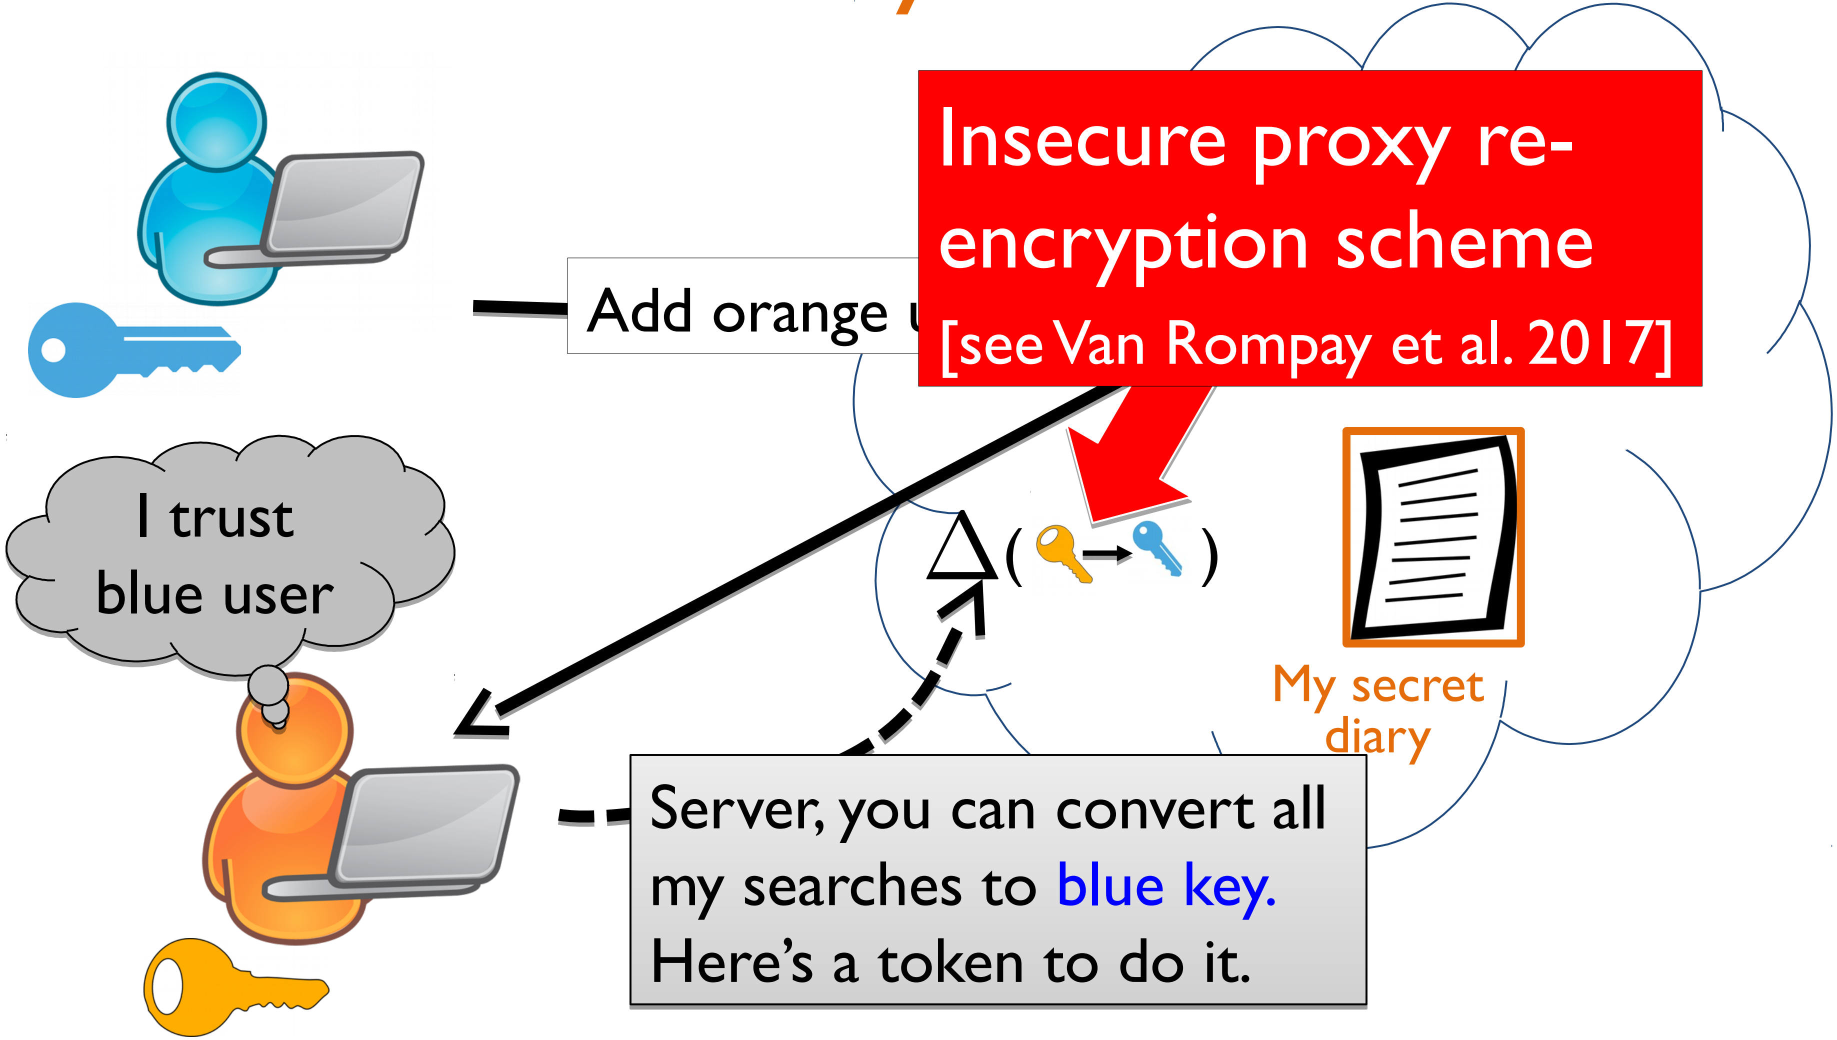

Insecure proxy re-encryption scheme
[see Van Rompay et al. 2017]
Add orange u
I trust blue user
My secret diary
Server, you can convert all my searches to blue key. Here's a token to do it.

Mylar Under Active Attack

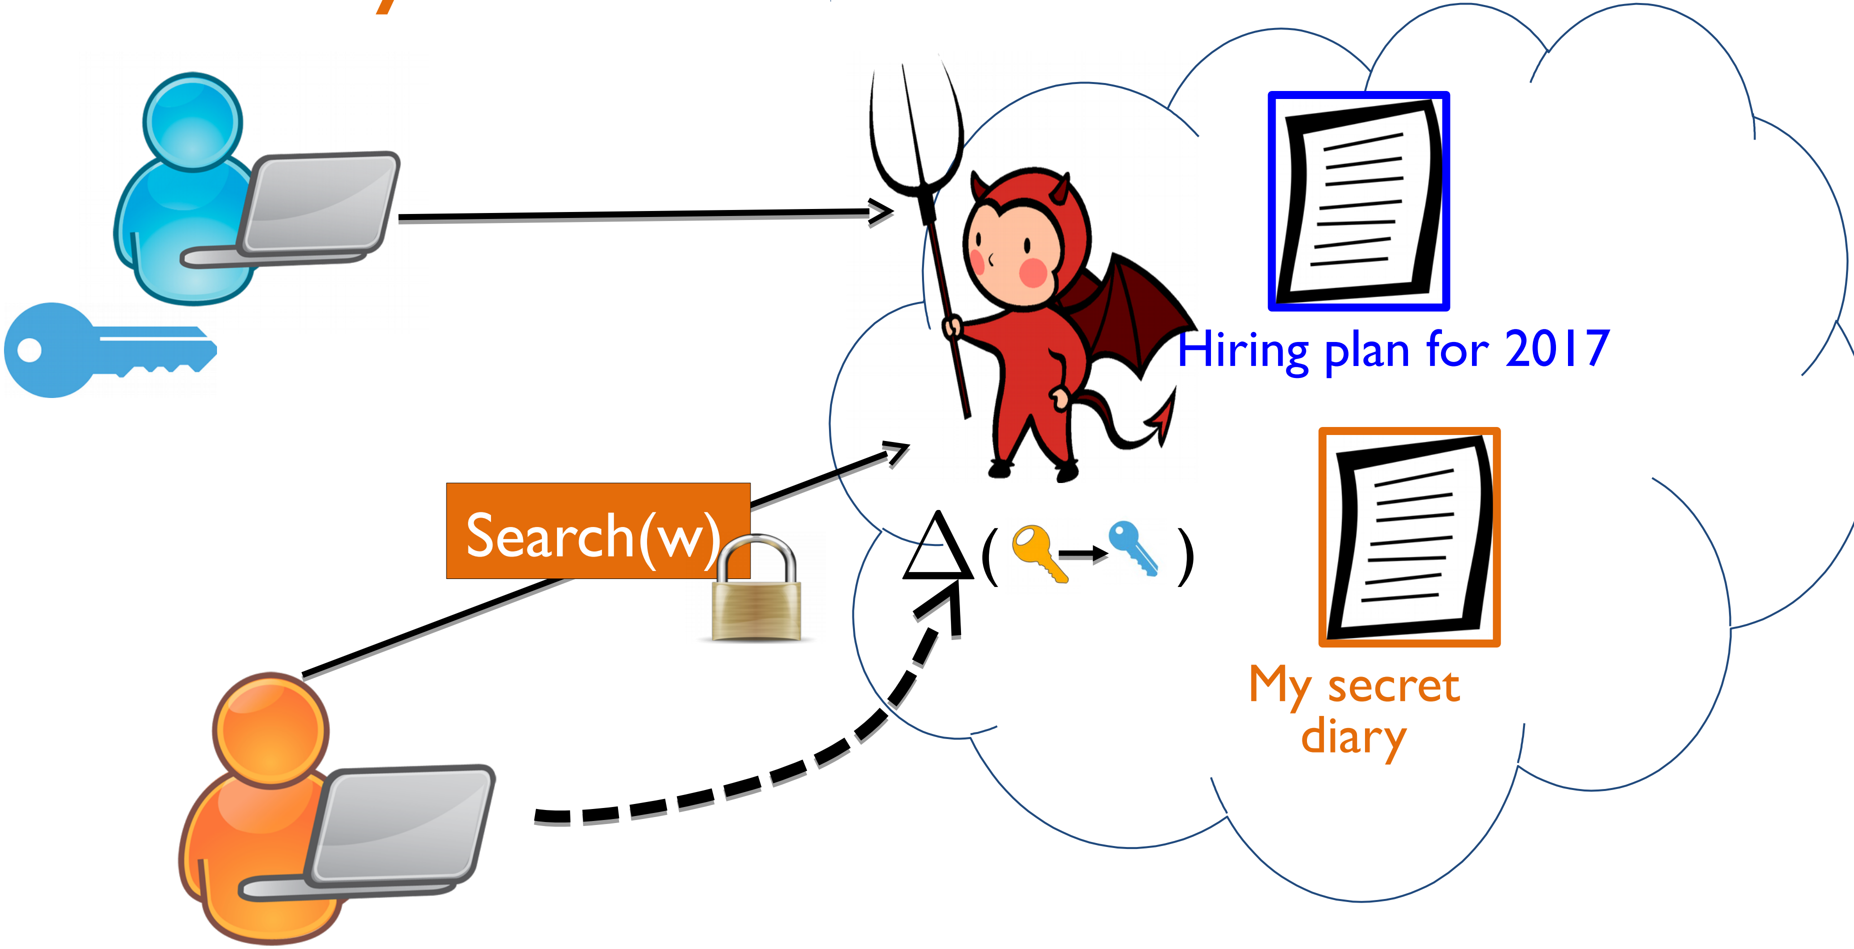

3.4 Threat model

Threats. Both the application and the database servers
can b
may
send
subs
from

some user machines … collude with the server… **because the adversary broke into a user's machine**

Mylar also allows some user machines to be controlled by the adversary, and to collude with the server. This may be either because the adversary is a user of the application, or because the adversary broke into a user's machine.

We call this adversary *active*, in contrast to a *passive* adversary that eavesdrops on all information at the server, but does not make any changes, so that the server responds to all client requests as if it were not compromised.

Mylar Under Active Attack

Guarantees. Mylar protects a data item's confidentiality in the face of arbitrary server compromises, as long as none of the users with access to that data item use a compromised machine.

... as long as none of the users with access to that data item use a compromised machine

Mylar Under Active Attack

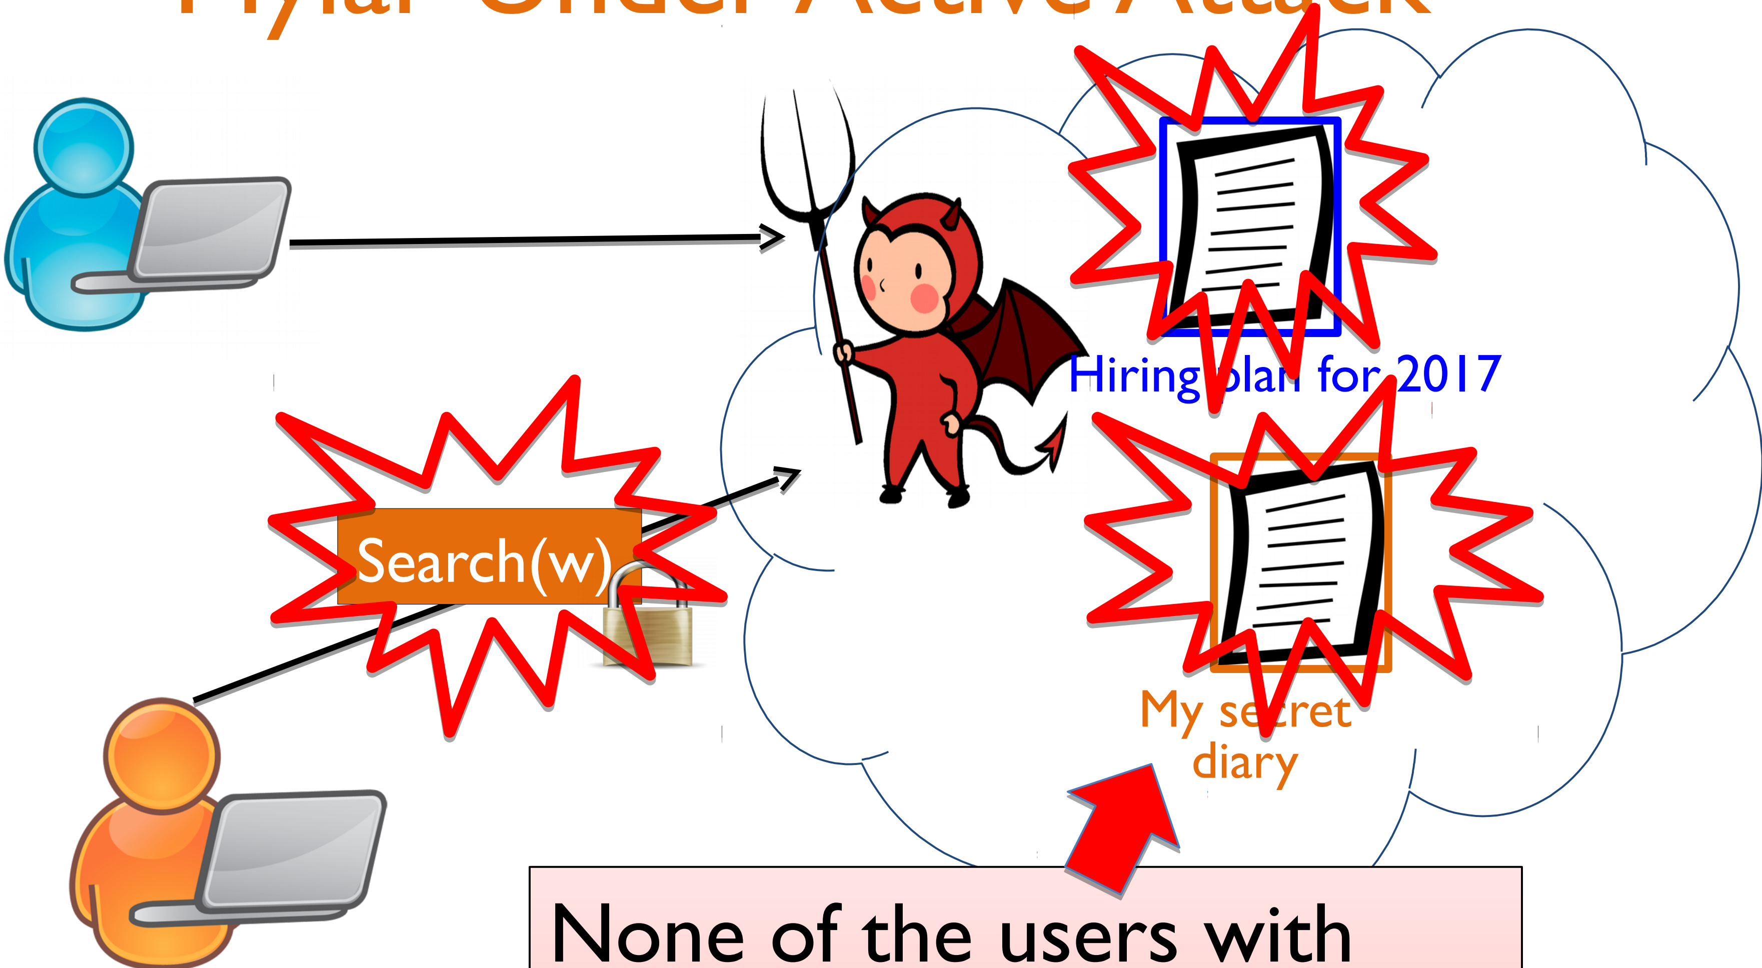

None of the users with access to this data item use a compromised machine

Mylar in a Hospital

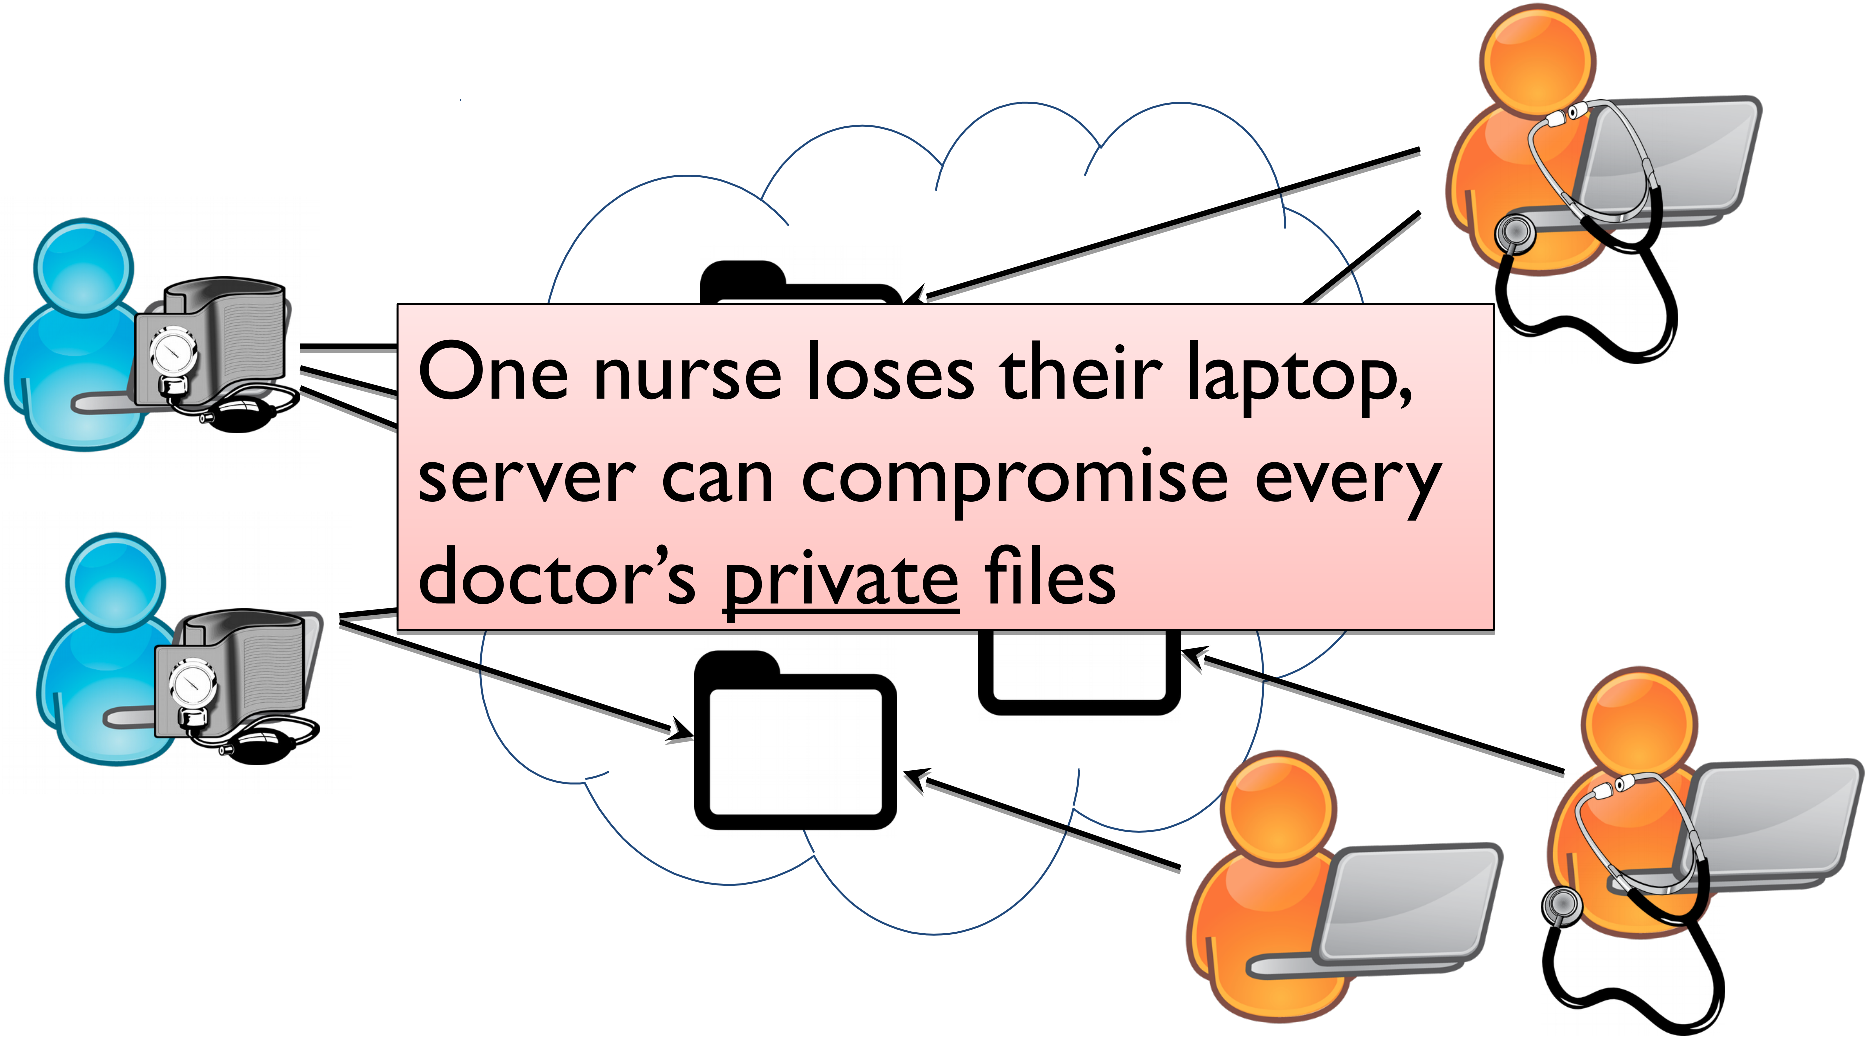

One nurse loses their laptop, server can compromise every doctor's private files

"Snapshot" Threat Model

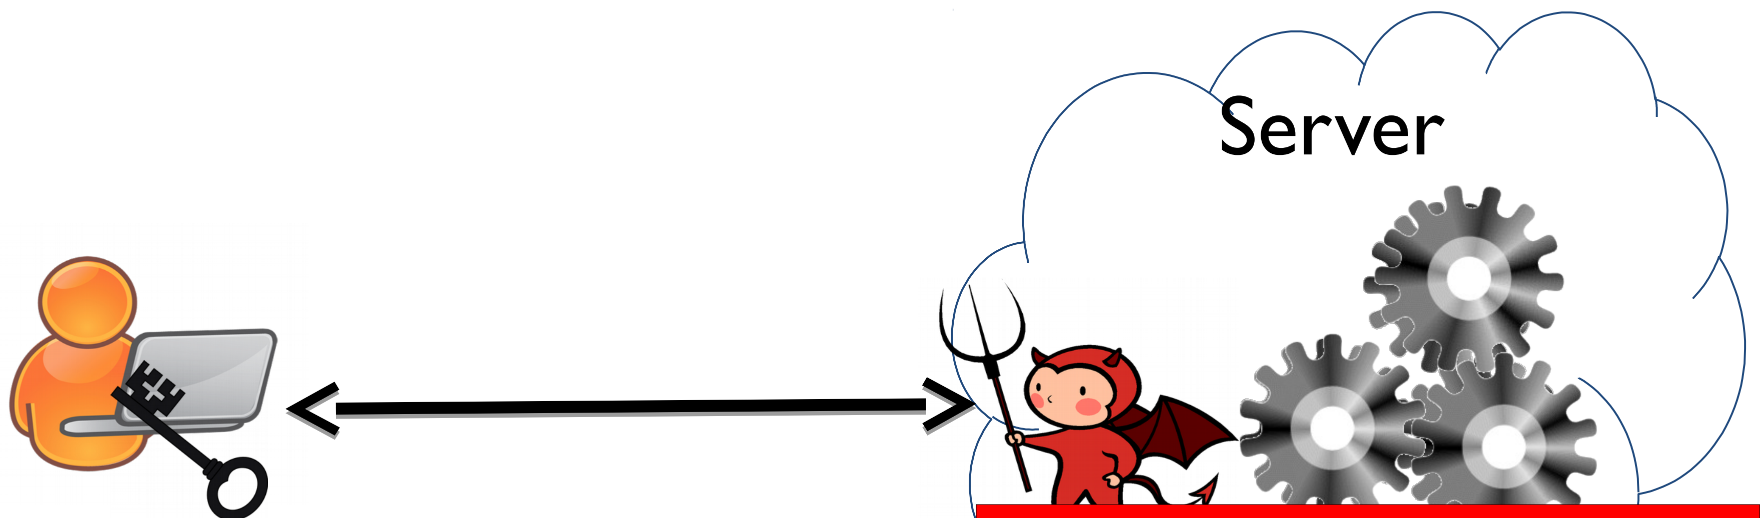

False in any realistic snapshot attack on a commodity DBMS

Existing systems explicitly claim security

… assuming there are no queries in the snapshot

A Simple System Abstraction

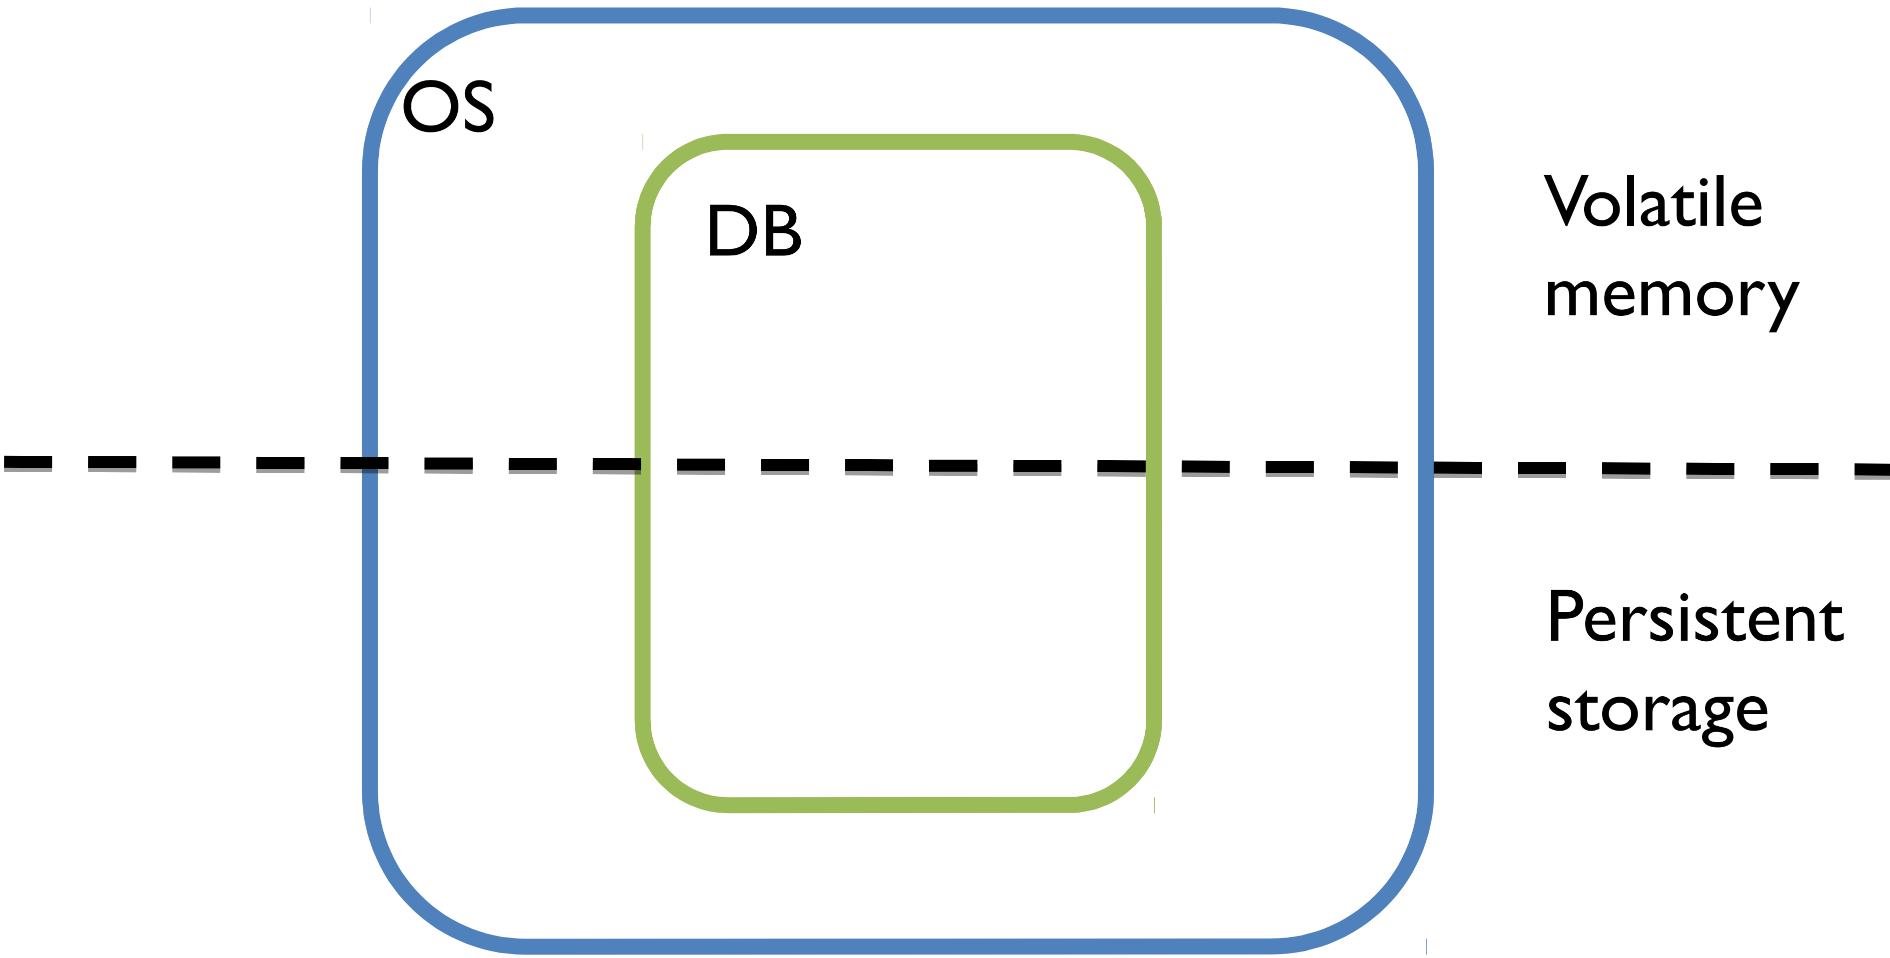

Actual Attacks

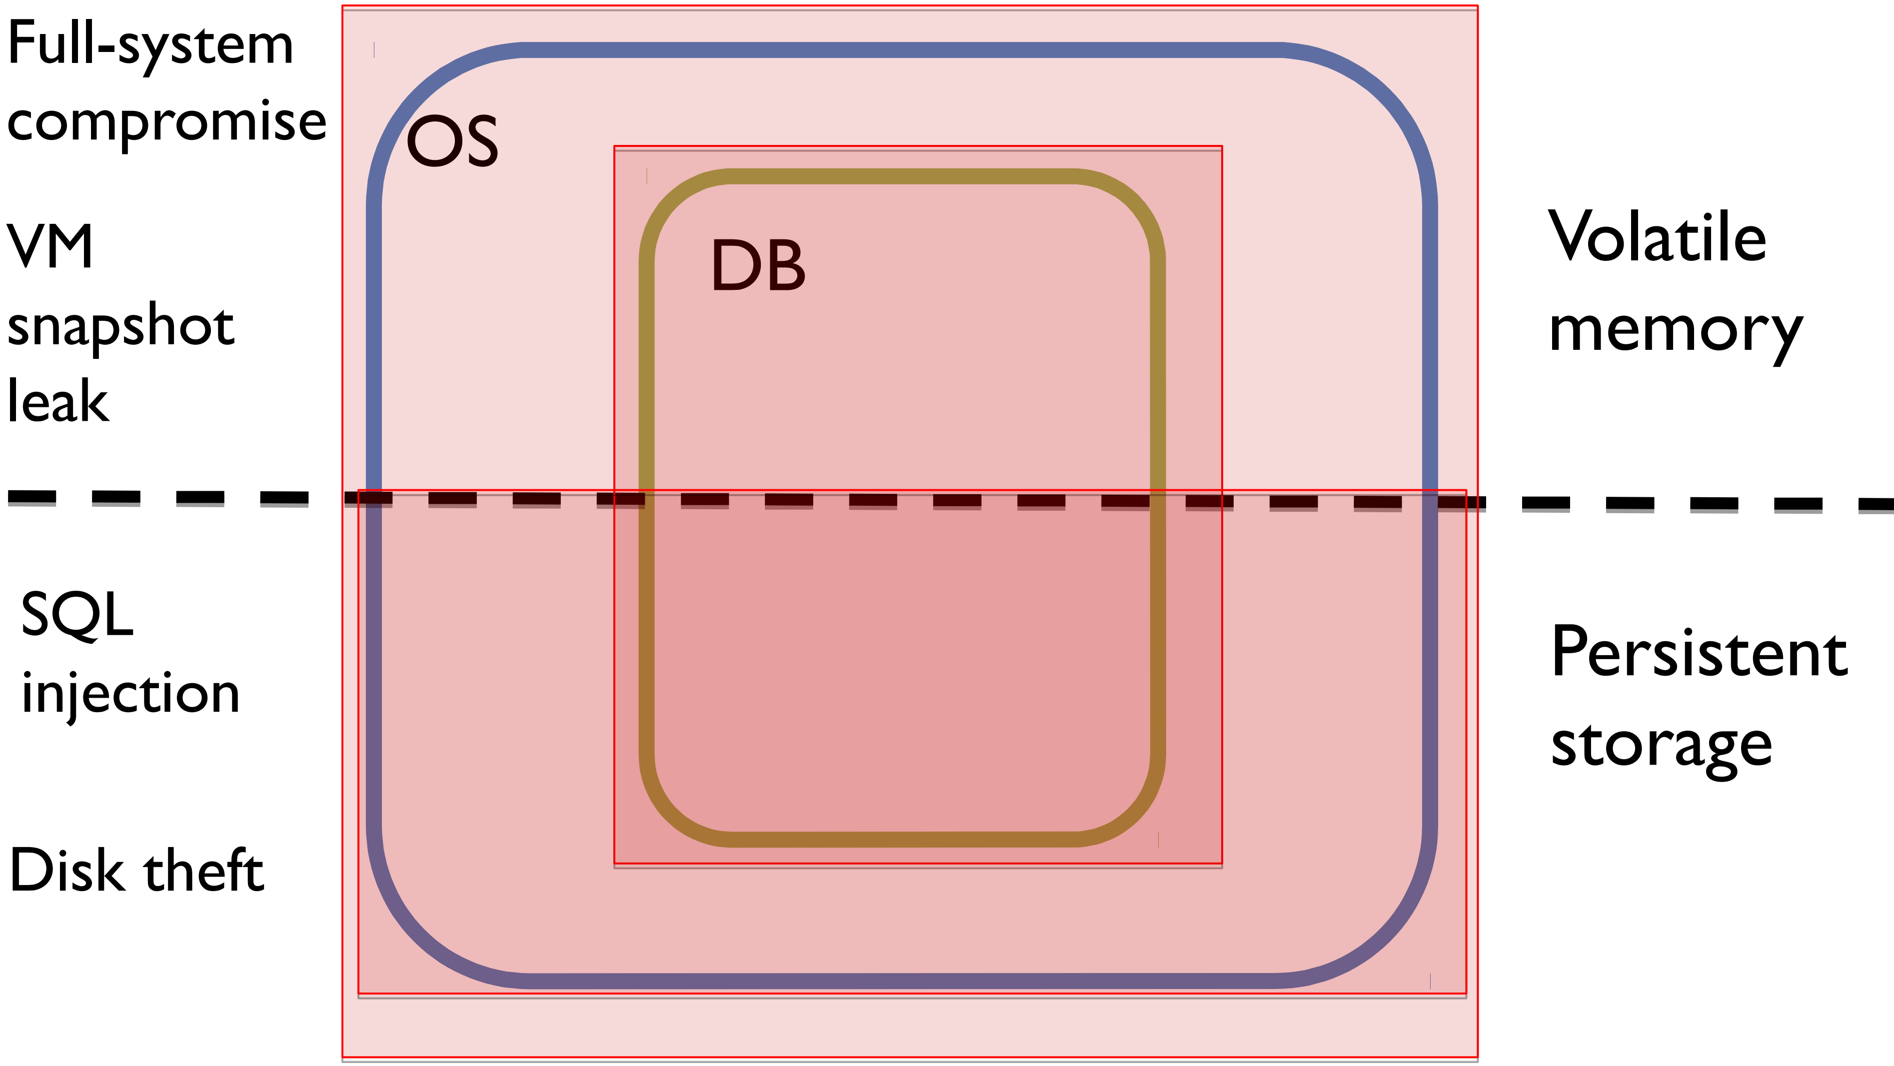

Case Study: MySQL

similar issues in any other commodity DBMS

Attack	What MySQL leaks	Failed encrypted database
Disk theft	MVCC data structures	Arx's range query index
SQL Injection	Past query statistics	Seabed's SPLASHE scheme
Full system compromise or VM snapshot leak	Text of past queries	CryptDB, Lewi/Wu, etc.

Disk Theft

Healthcare IT News

Privacy & Security

Stolen laptop lea
notification for 2(
patients

If this is your threat model, just use full-disk encryption

SC MEDIA

› SC US
SC UK

NEWS CYBERCRIME NETWORK SECURITY PRODUCT REVIEWS IN DEPTH EVE

THE CYBERSECURITY SOU

SC Magazine US > Blogs > The Data Breach Blog > Hard drive stolen from Jackson Memorial Hospital

Hard drive stolen from Jackson Memorial Hospital

Logs on Disk

General query log (not widely used)

Binary log records modifications, used for replication and recovery

Data modification queries can be reconstructed from these logs [FHMW '10, FKSHW '12]

Multi-version concurrency control using log data structures

In all modern SQL databases!

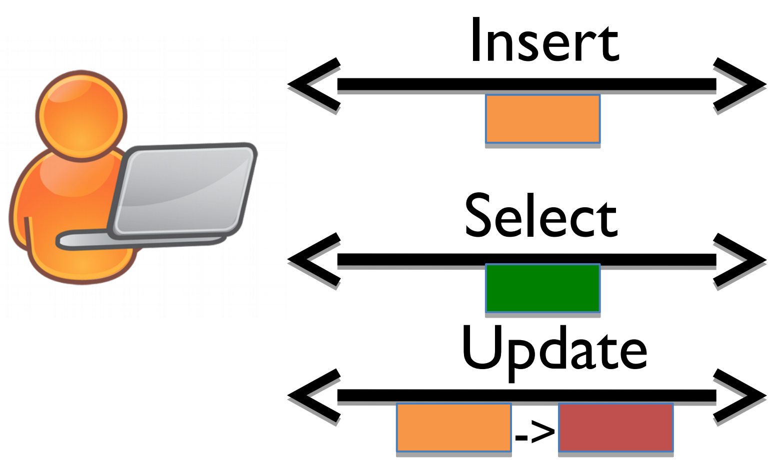

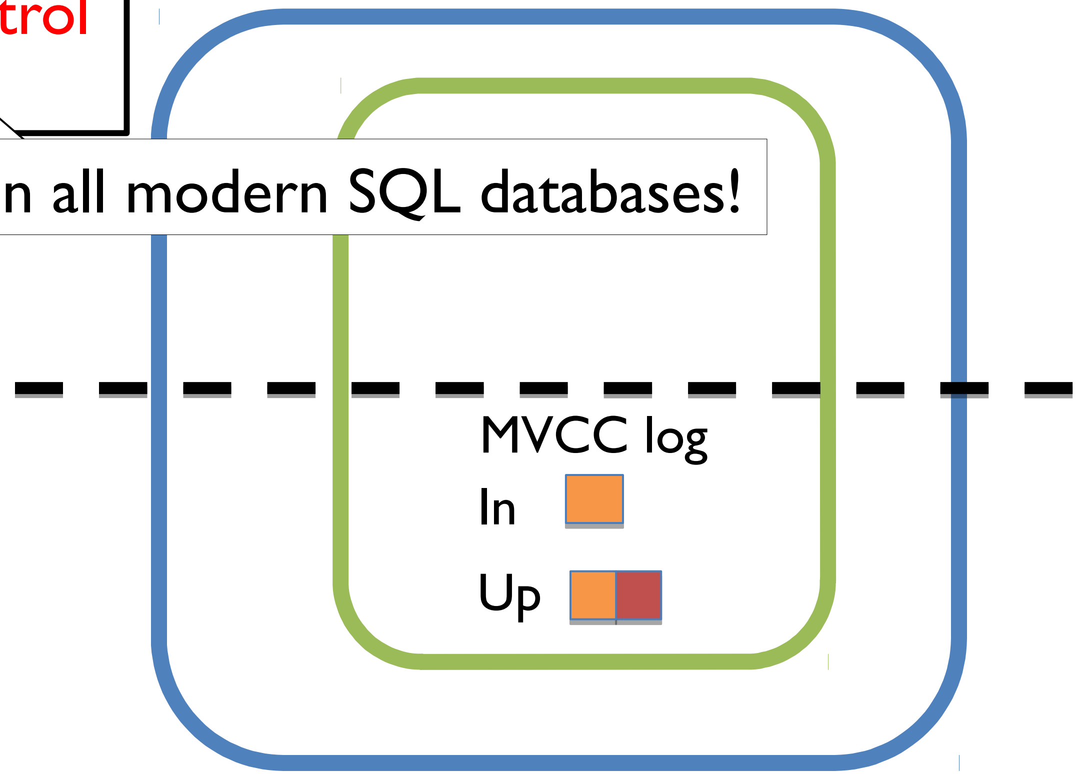

Poddar et al.

Arx

Range queries via chained garbled circuits
Tree nodes become consumed, need replacing

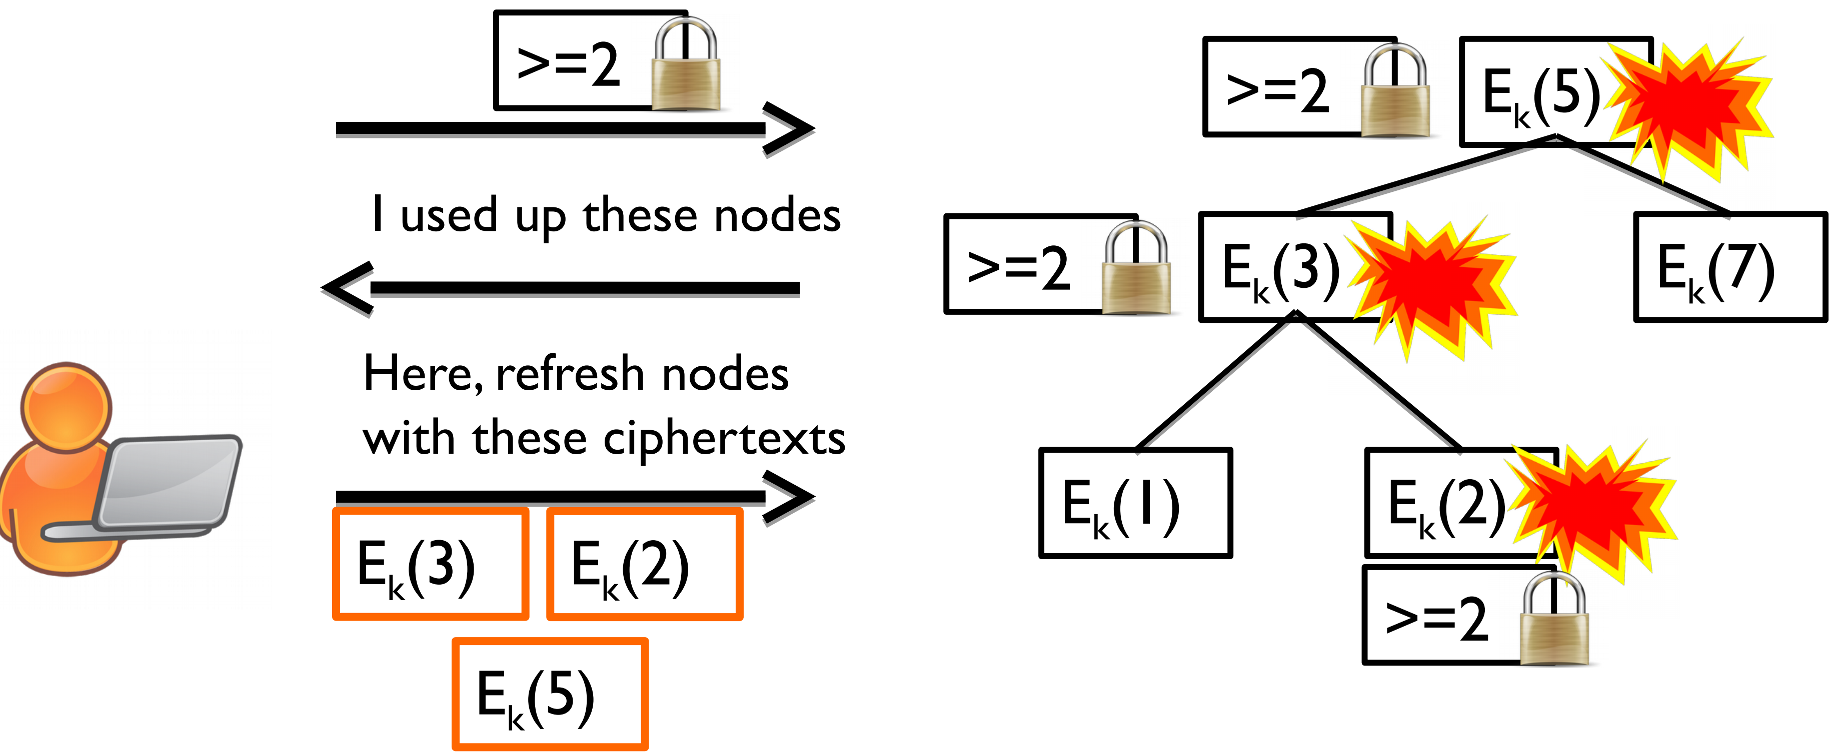

Poddar et al.

Security Claim for Arx

“Arx protects the database with the same level of security as regular AES-based encryption”

Arx Under Snapshot Attack

Range queries via chained garbled circuits

Tree nodes become consumed, need replacing

Consumed nodes immediately replaced, stored in MVCC log

Query access pattern recorded on disk

Snapshot attacker can recover queries and plaintexts using variants of attacks from [GSBNR - S&P '17]

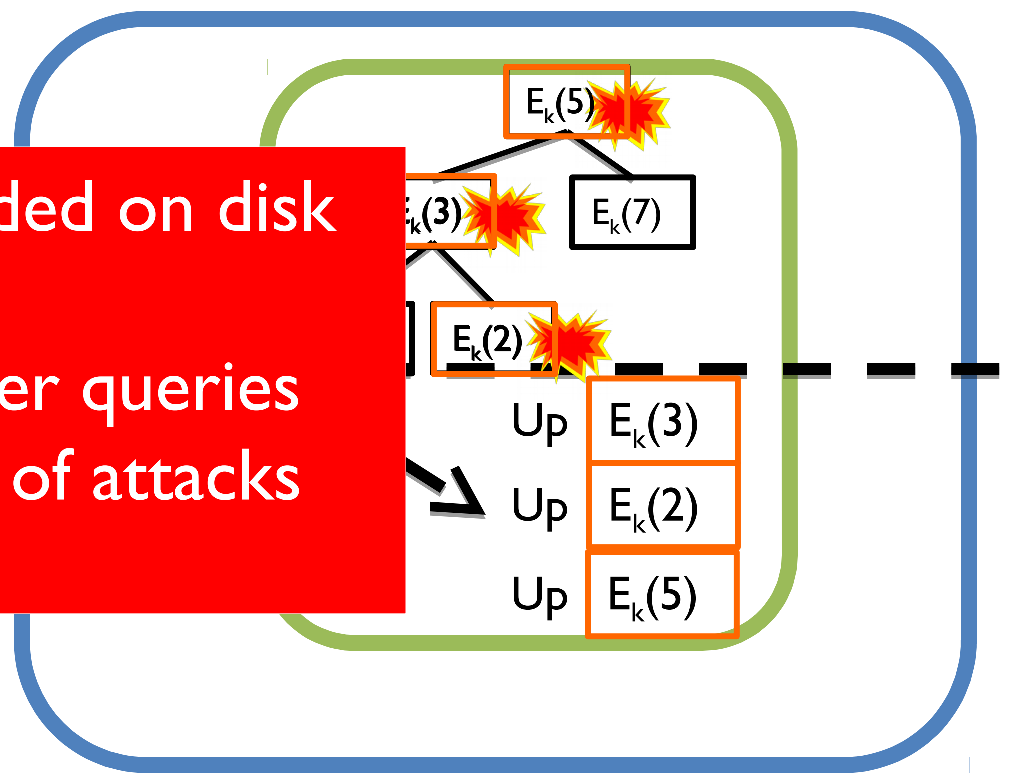

SQL Injection

Attack	What MySQL leaks	Failed encrypted database
Disk theft	MVCC data structures	Arx's range query index
SQL Injection	Past query statistics	Seabed's SPLASHE scheme
Full system compromise or VM snapshot leak	Text of past queries	CryptDB, Lewi/Wu, etc.

SQL Injection

SQL injection accounted for 51% of all Web application attacks in 2016 (source: Akamai)

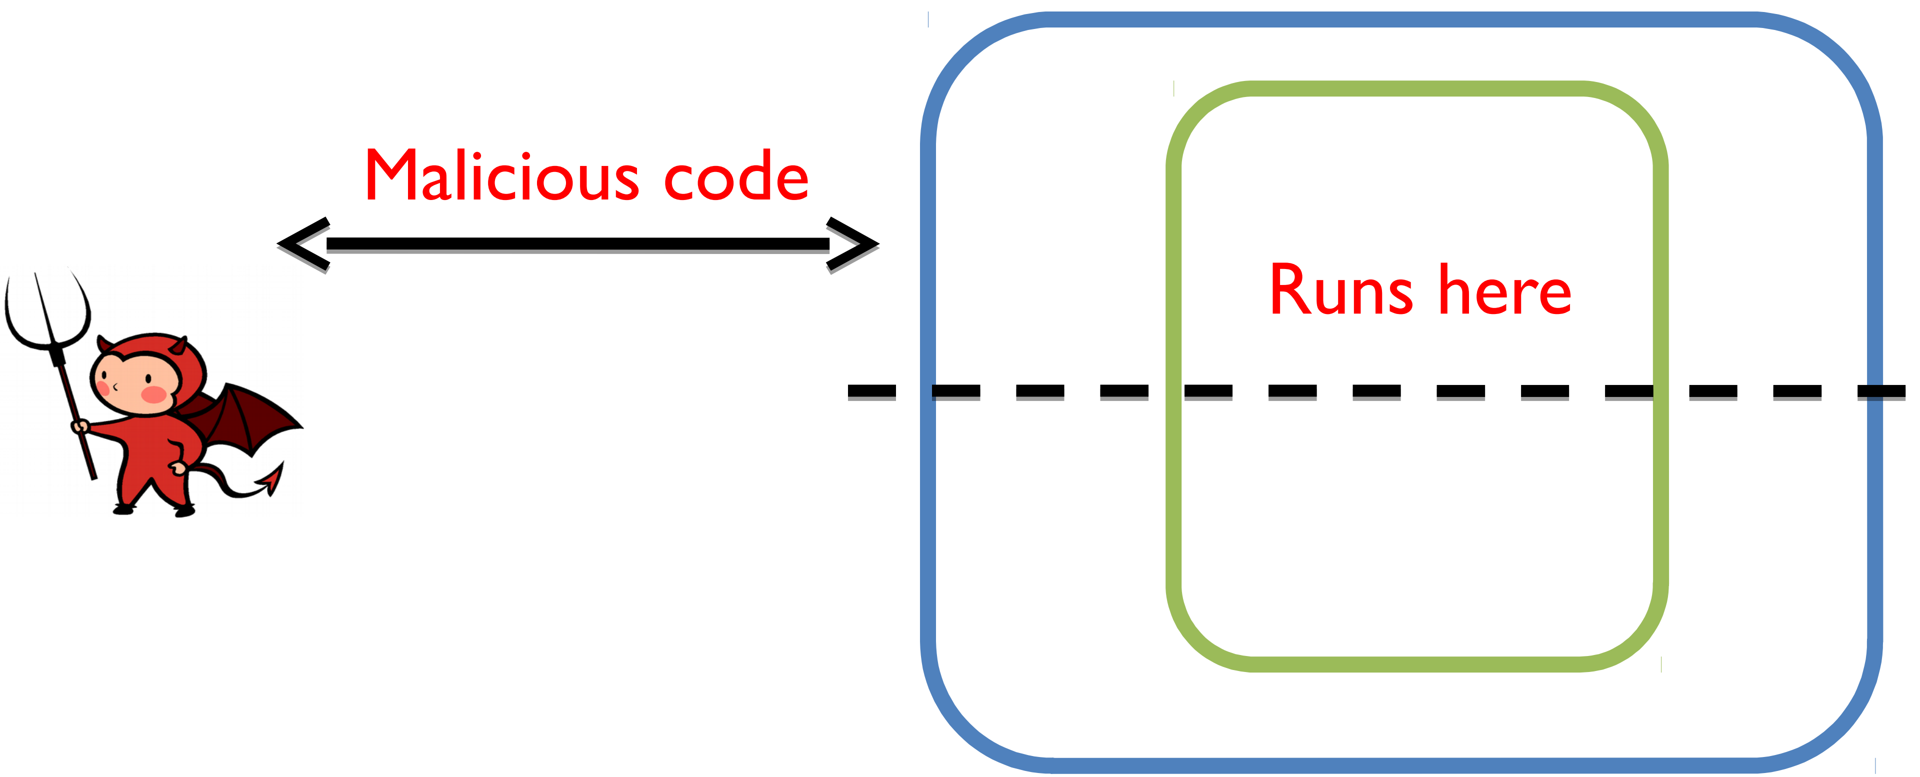

Diagnostic Tables

information_schema stores current query for all users, contents of buffer cache

performance_schema stores current query for all threads, statistics for past queries

Separate counts for queries which involve different columns

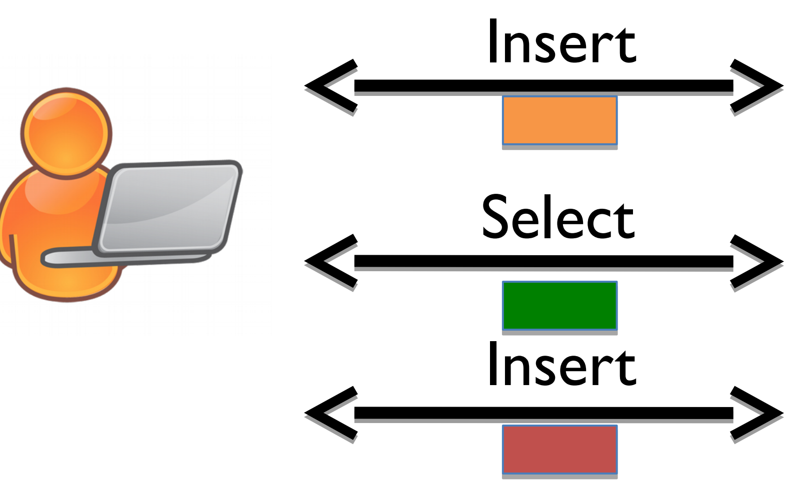

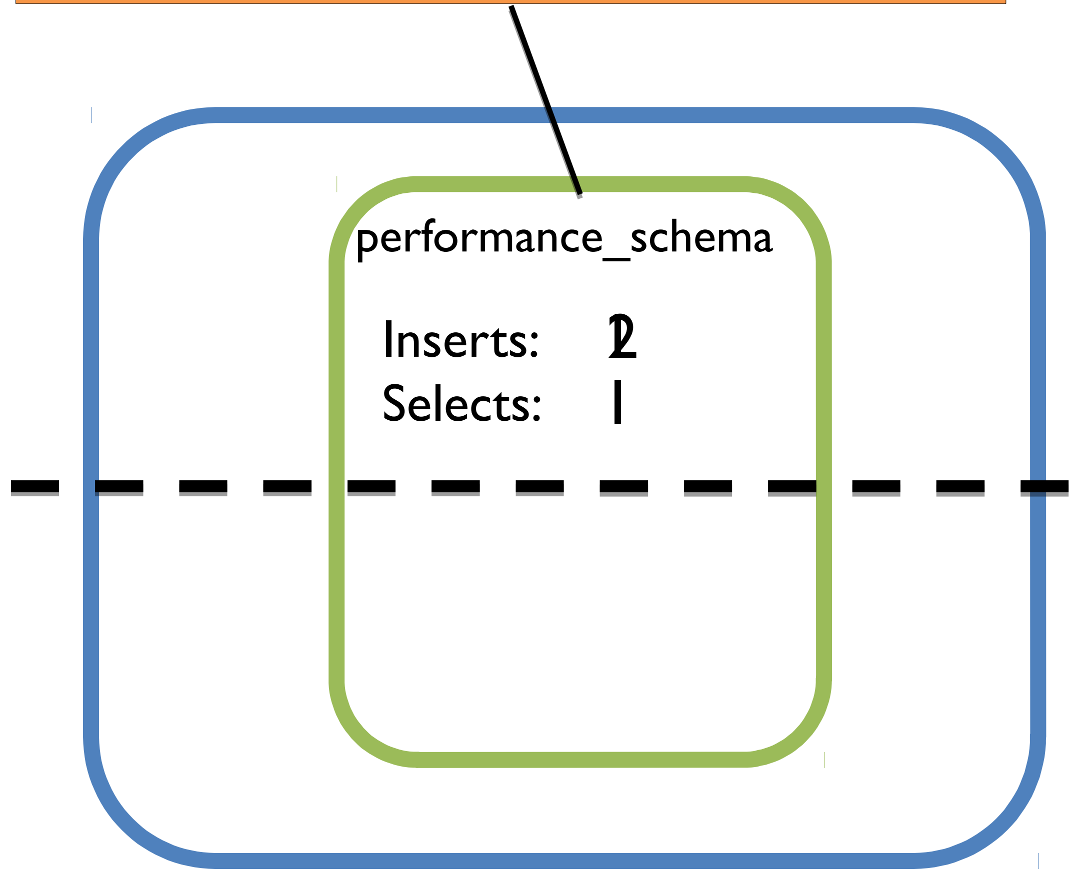

Problem: Frequency Analysis

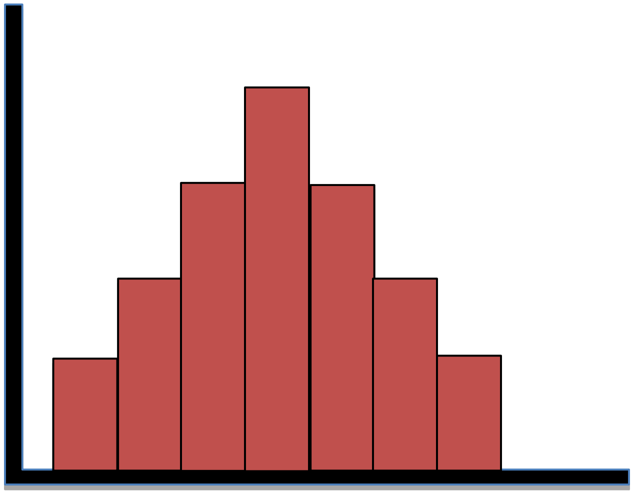

Name	Has given this talk before
Paul Grubbs	1
Thomas Ristenpart	0
Vitaly Shmatikov	0

Order-preserving encryption reveals histogram of plaintext values

This is how Naveed et al. used frequency analysis to break CryptDB: match histogram to auxiliary model of data distribution

Papadimitriou et al. (OSDI 2016)

Seabed

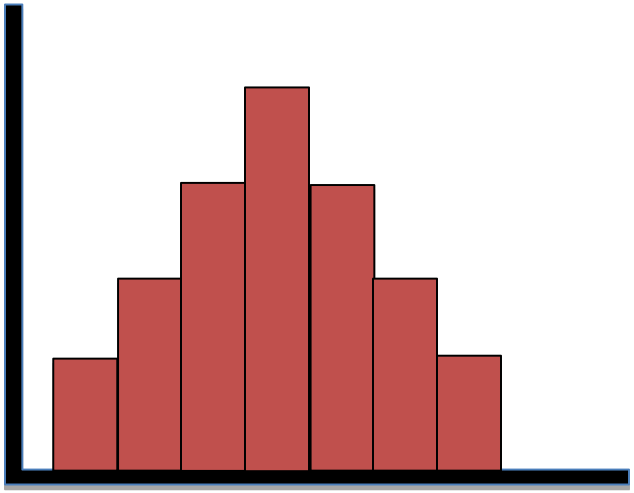

Name	Has given this talk before
Paul Grubbs	1
Thomas Ristenpart	0
Vitaly Shmatikov	0

("Has …"=1) ("Has …"=0)

Name	C2	C3
aspoiwnpoinio	$E_k(1)$	$E_k(0)$
petryoiueytiew	$E_k(0)$	$E_k(1)$
Xncmxncmbcn	$E_k(0)$	$E_k(1)$

Each possible plaintext gets its own column

WHERE clause transformed to correct column

SELECT Count("Has … ") WHERE "Has …"=1 → SELECT Count(C2)

Separate counts for queries which involve different columns

Example

Plaintext Schema

country	salary
USA	100000
USA	100000
Canada	200000
USA	300000
Canada	500000
Canada	800000
India	100000
India	100000
Chile	200000
Iraq	300000
China	500000
Japan	800000
Israel	130000
U.K.	210000

Schema with Enhanced SPLASHE

country	salaryUSA	salaryCanada	salaryOthers
DET(Chile)	ASHE(100000)	ASHE(0)	ASHE(0)
DET(Iraq)	ASHE(100000)	ASHE(0)	ASHE(0)
DET(China)	ASHE(0)	ASHE(200000)	ASHE(0)
DET(Japan)	ASHE(300000)	ASHE(0)	ASHE(0)
DET(Israel)	ASHE(0)	ASHE(500000)	ASHE(0)
DET(U.K.)	ASHE(0)	ASHE(800000)	ASHE(0)
DET(India)	ASHE(0)	ASHE(0)	ASHE(100000)
DET(India)	ASHE(0)	ASHE(0)	ASHE(100000)
DET(Chile)	ASHE(0)	ASHE(0)	ASHE(200000)
DET(Iraq)	ASHE(0)	ASHE(0)	ASHE(300000)
DET(China)	ASHE(0)	ASHE(0)	ASHE(500000)
DET(Japan)	ASHE(0)	ASHE(0)	ASHE(800000)
DET(Israel)	ASHE(0)	ASHE(0)	ASHE(130000)
DET(U.K)	ASHE(0)	ASHE(0)	ASHE(210000)

SQLi Extracts Diagnostic Tables

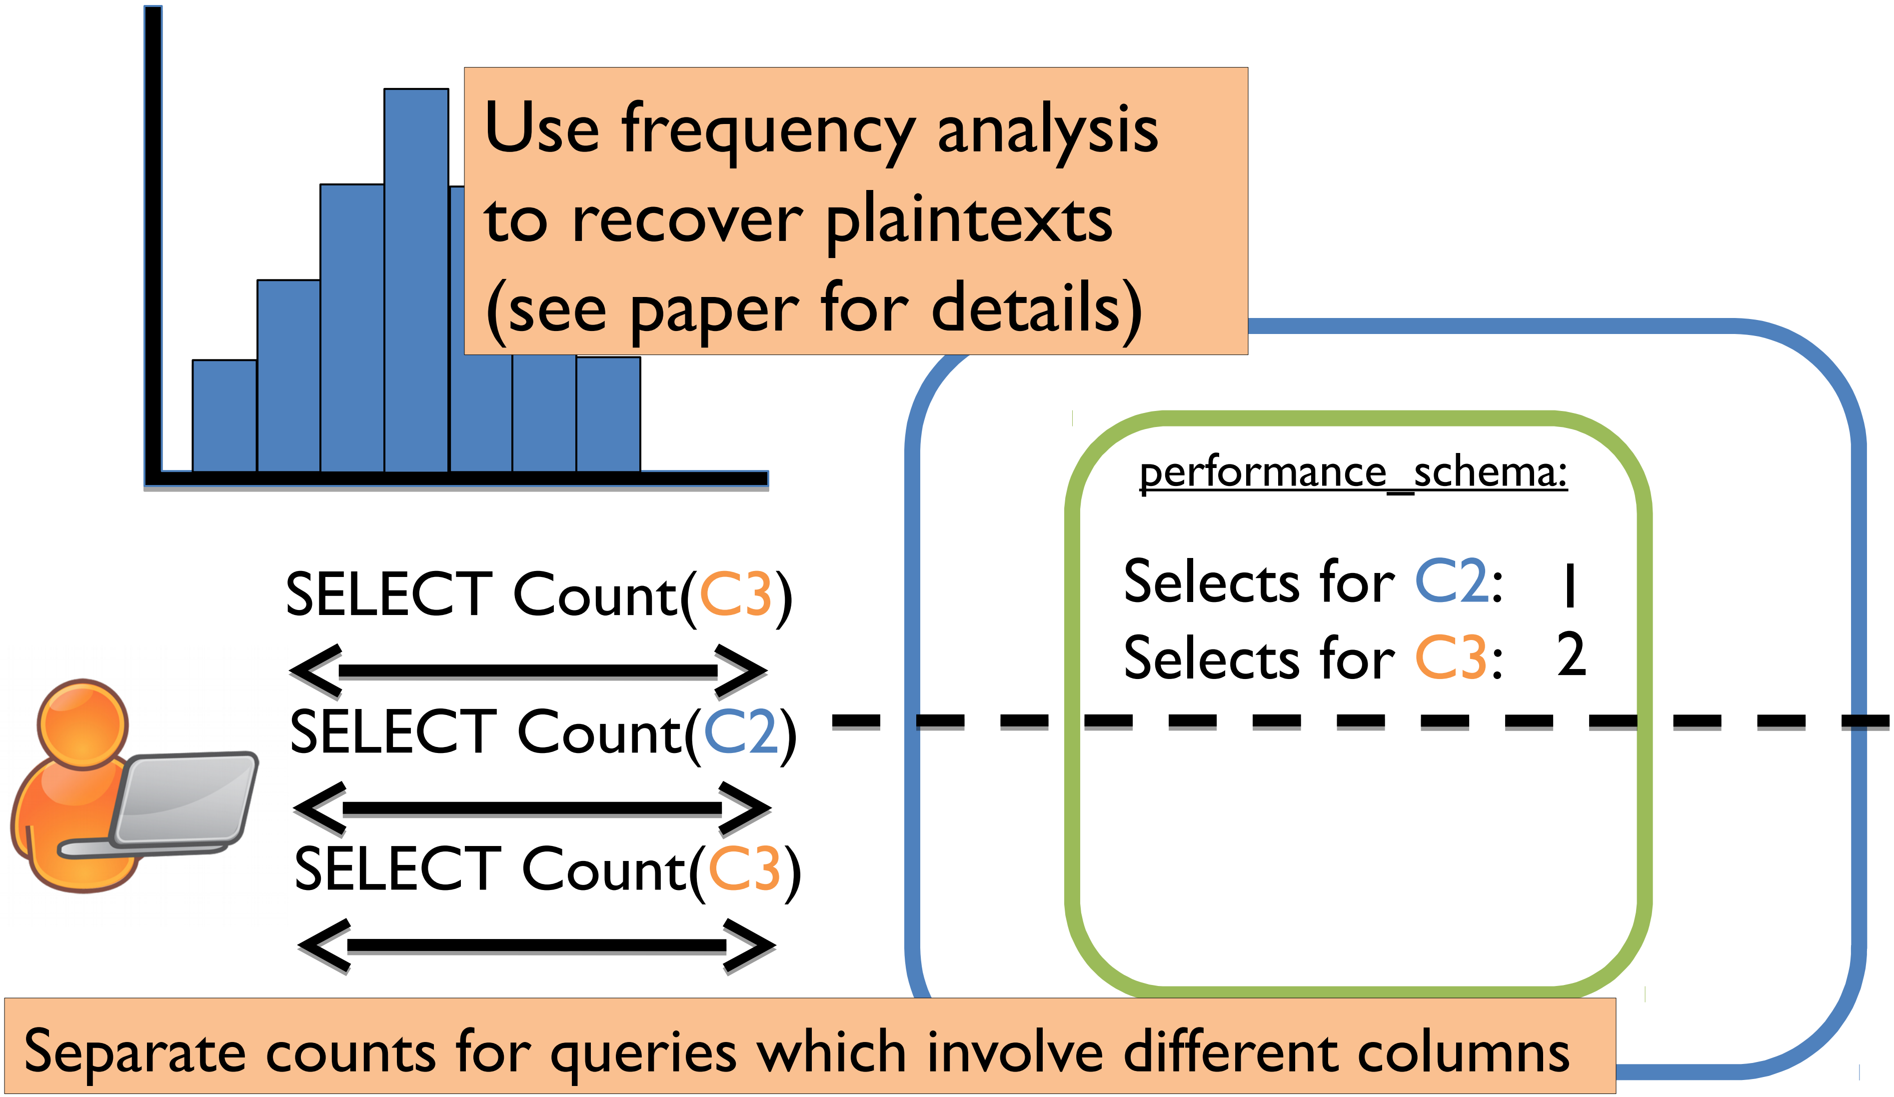

Full-System Snapshot

Attack	What MySQL leaks	Failed encrypted database
Disk theft	MVCC data structures	Arx's range query index
SQL Injection	Past query statistics	Seabed's SPLASHE scheme
Full system compromise or VM snapshot leak	Text of past queries	CryptDB, Lewi/Wu, etc.

Full-System Compromise

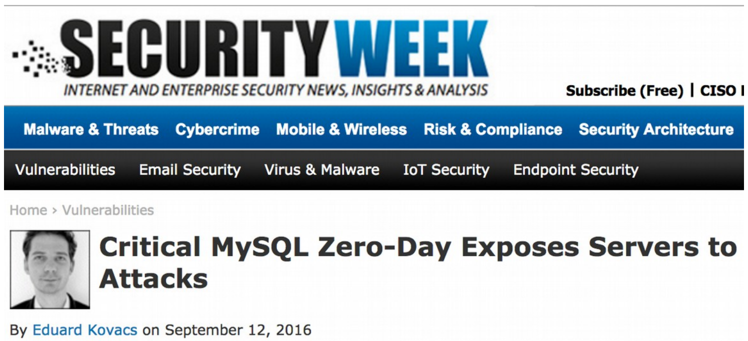
SECURITYWEEK
INTERNET AND ENTERPRISE SECURITY NEWS, INSIGHTS & ANALYSIS
Subscribe (Free) | CISO
Malware & Threats Cybercrime Mobile & Wireless Risk & Compliance Security Architecture
Vulnerabilities Email Security Virus & Malware IoT Security Endpoint Security
Home › Vulnerabilities

Critical MySQL Zero-Day Exposes Servers to Attacks

By Eduard Kovacs on September 12, 2016

Leakage of sensitive data at OS level is well-studied [CPGR, DLJKSXSW]

We focus on DBMS address space, things inaccessible to users

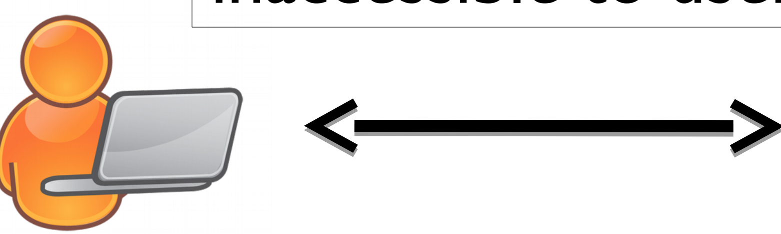

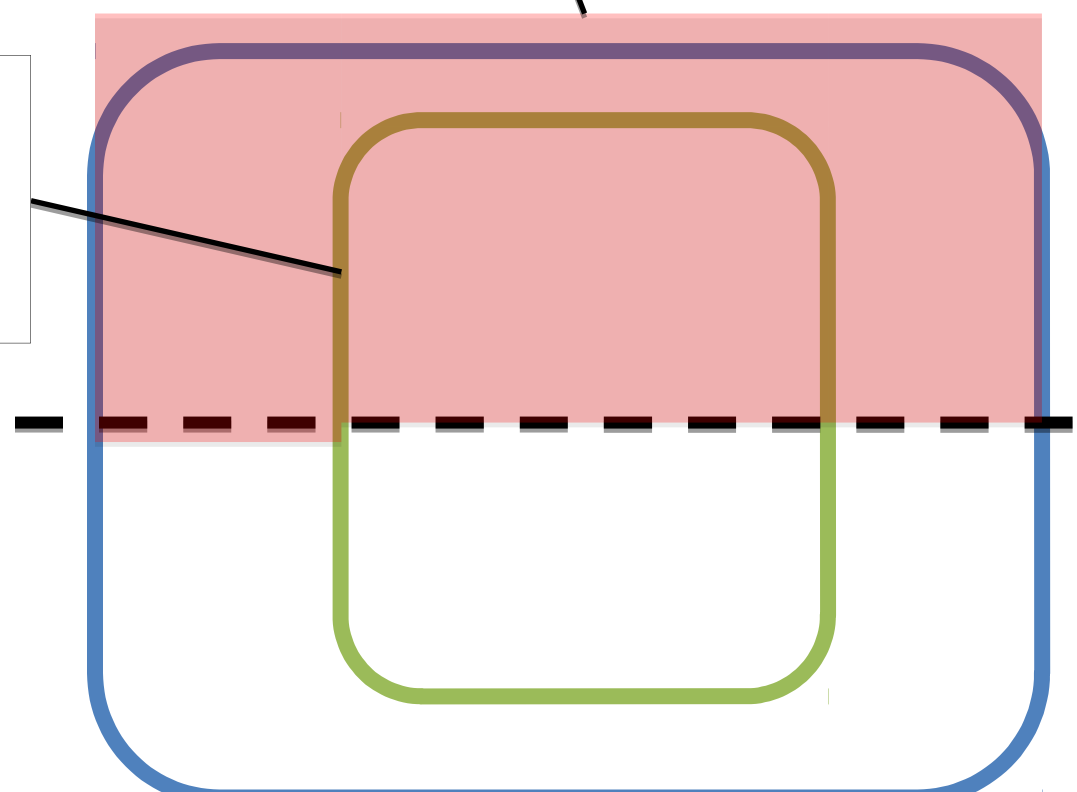

Data Structures and Caches

Adaptive hash index tracks pages accesses, indexes automatically
MySQL query cache stores select queries and results
Other query caches (memcached)

essential for performance!

MySQL manages internal heaps, does not zero freed memory!

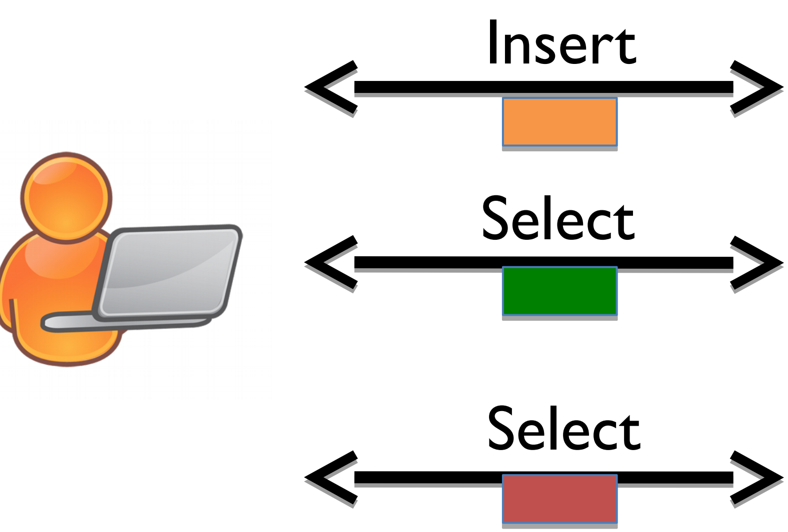

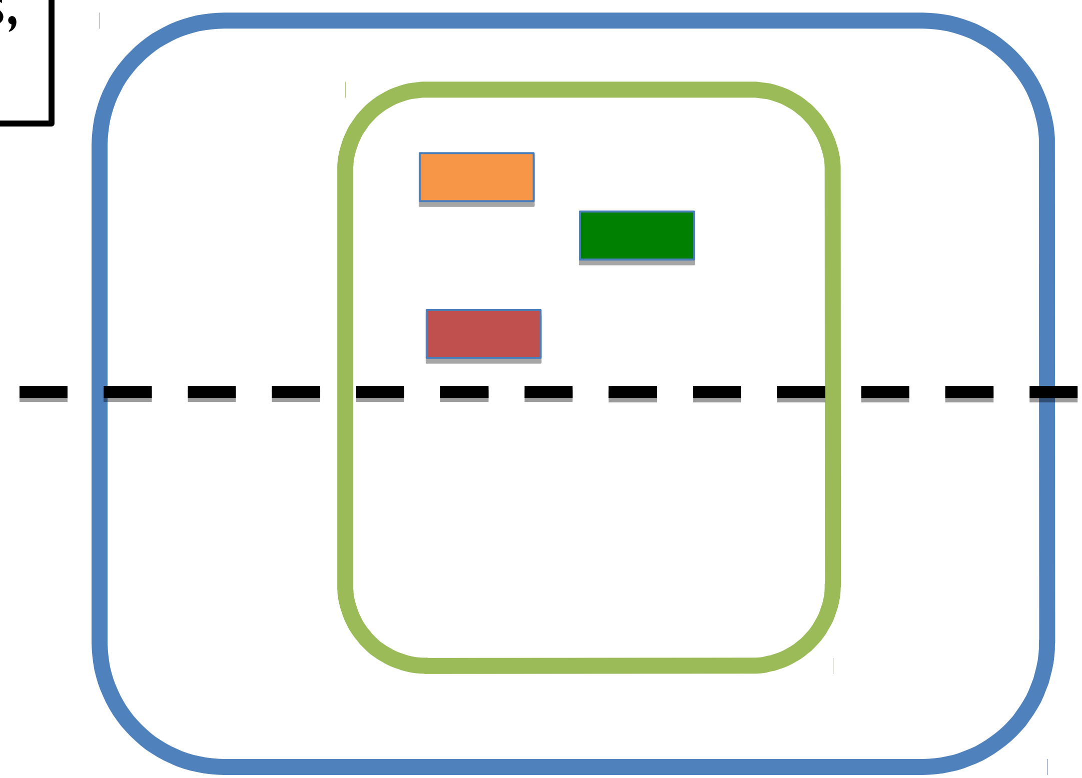

Token-Based Systems

CryptDB, Mylar, Lewi-Wu, other searchable encryption schemes cannot be semantically secure if attacker sees a single search token

1,000 random selects…

Waited a while…

100,000 more random selects…

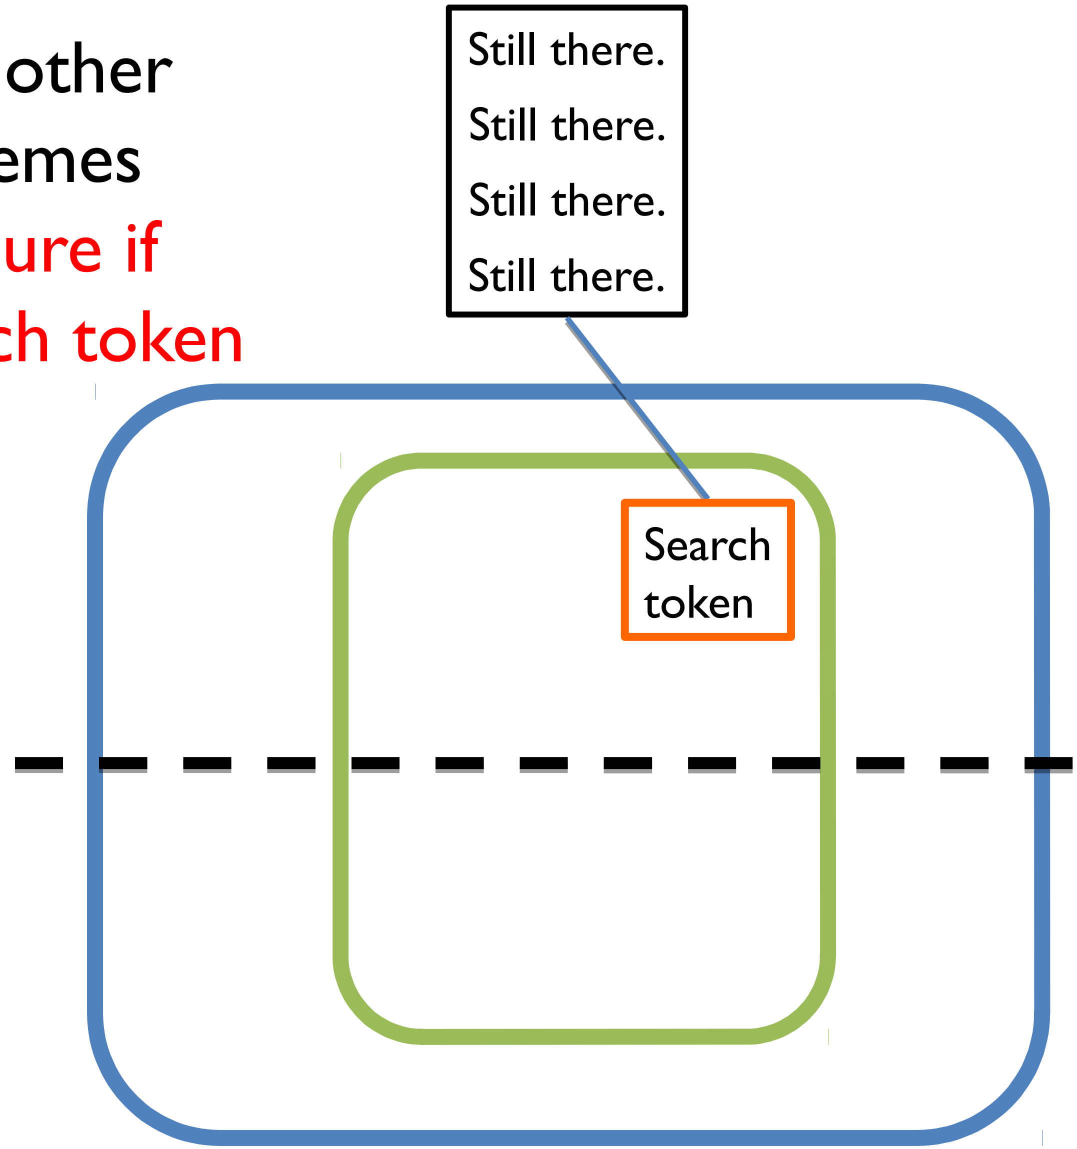

Let Me Make Myself Perfectly Clear

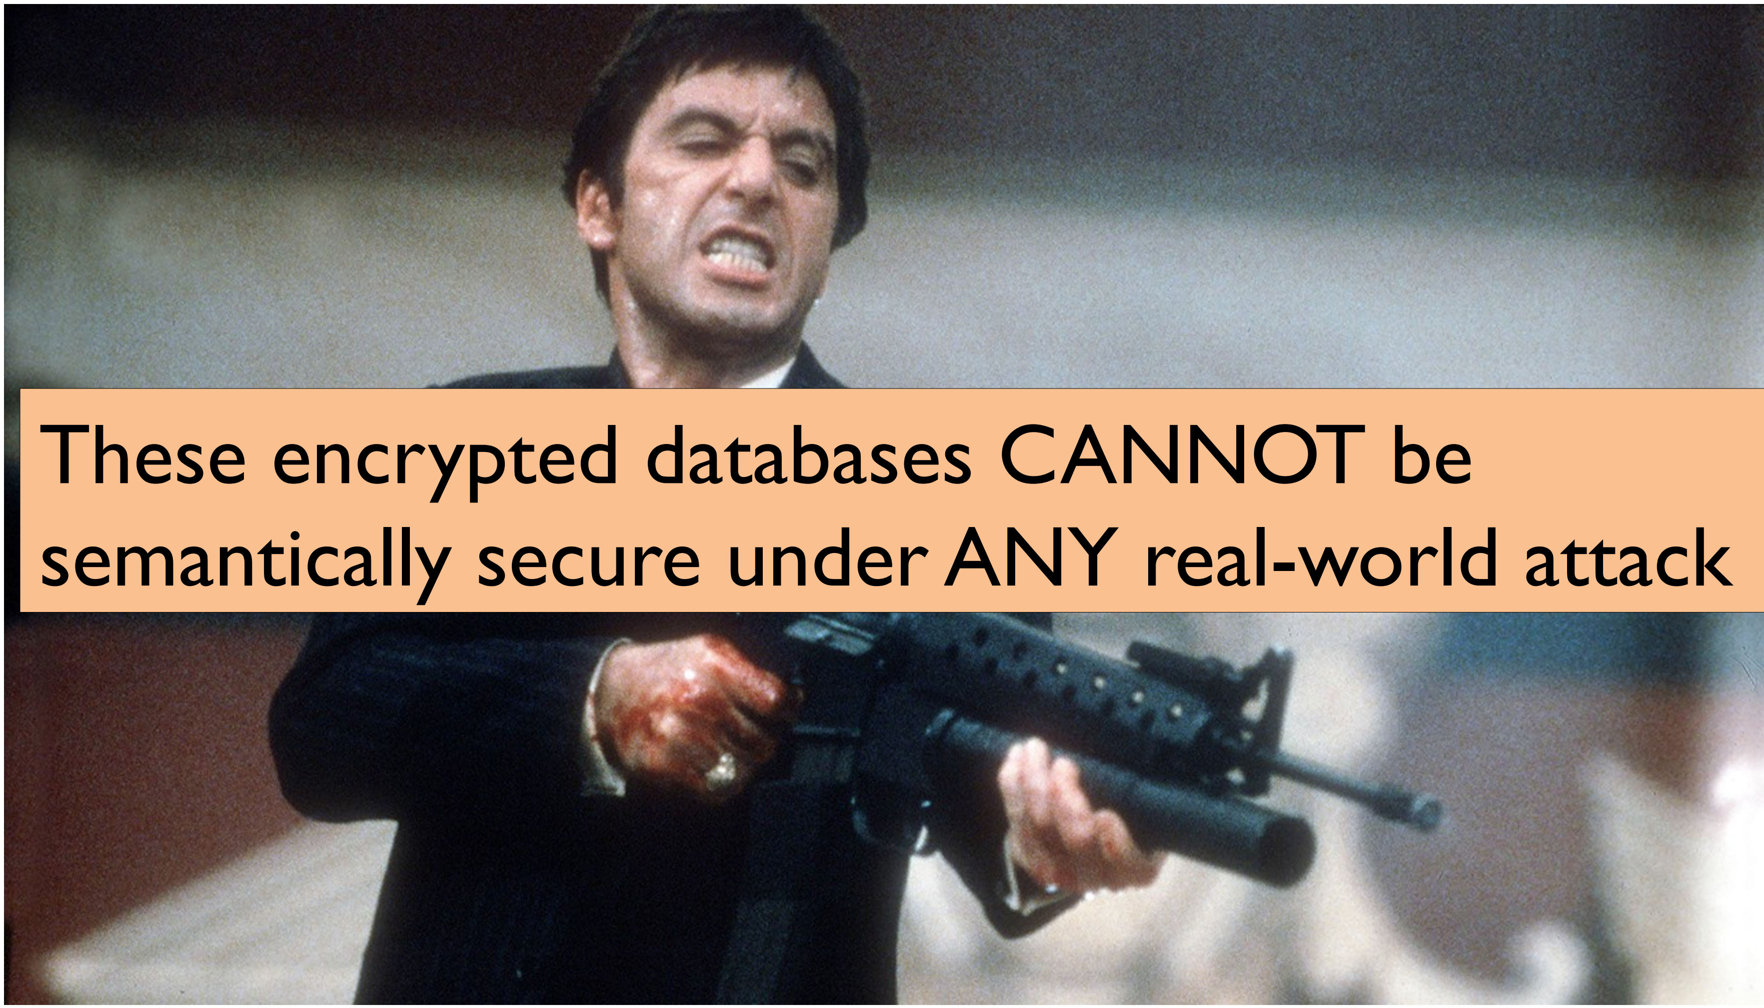

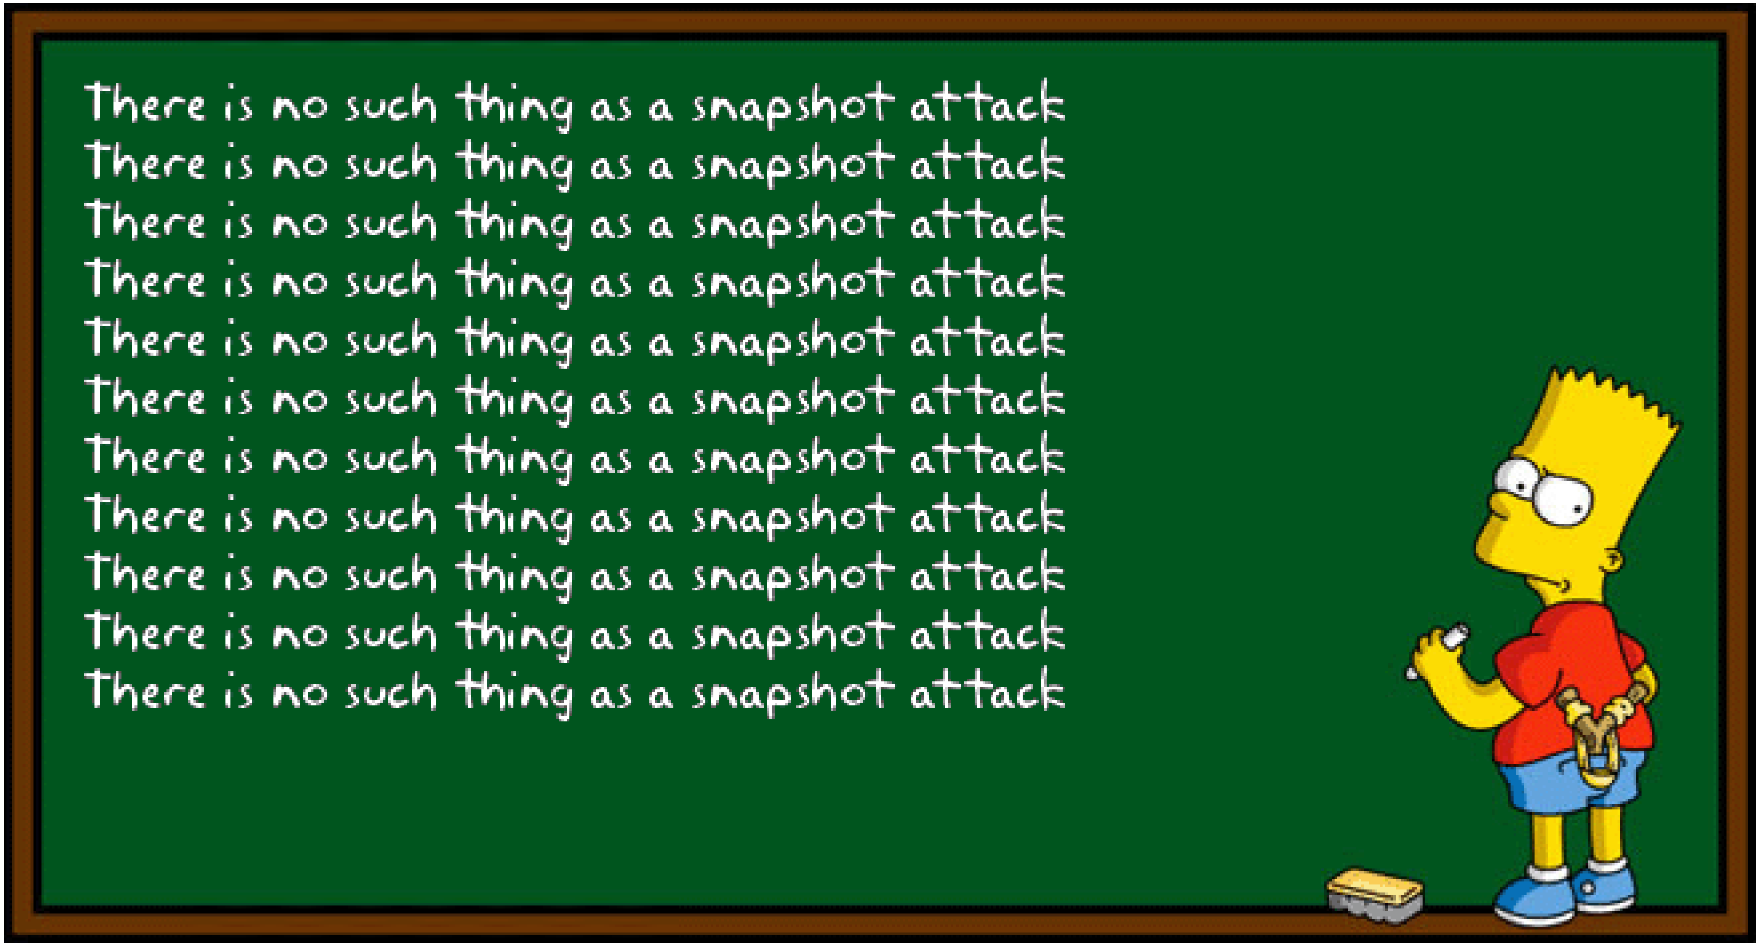
There is no such thing as a snapshot attack
There is no such thing as a snapshot attack
There is no such thing as a snapshot attack
There is no such thing as a snapshot attack
There is no such thing as a snapshot attack
There is no such thing as a snapshot attack
There is no such thing as a snapshot attack
There is no such thing as a snapshot attack
There is no such thing as a snapshot attack
There is no such thing as a snapshot attack
There is no such thing as a snapshot attack

“I Will Build My Own Database”

You can try…

Transaction logs needed to support ACID
Log-structured storage
Caching
Adaptive data structures adjust to workload

… everything in modern databases leaks information about past queries

Sensitivity Analysis

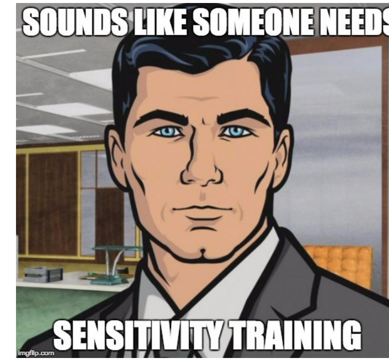

SSN	Name	Ethnicity	Date Of Birth	Sex	ZIP	Marital Status	Problem
		black	09/27/64	male	02139	divorced	obesity
		black	09/30/64	male	02139	divorced	hypertension
		black	04/18/64	male	02139	married	chest pain
		black	04/15/64	male	02139	married	chest pain
		black	09/15/64	male	02138	married	shortness of breath
		caucasian	03/13/63	male	02141	married	hypertension
		caucasian	03/18/63	male	02141	married	shortness of breath
		caucasian	09/13/64	female	02138	married	shortness of breath
		caucasian	09/07/64	female	02138	married	obesity
		caucasian	05/14/61	female	02138	single	chest pain
		caucasian	05/08/61	female	02138	single	obesity

order-preserving encryption

deterministic encryption

“strong” encryption

can sort

can check for equality

Auxiliary Data

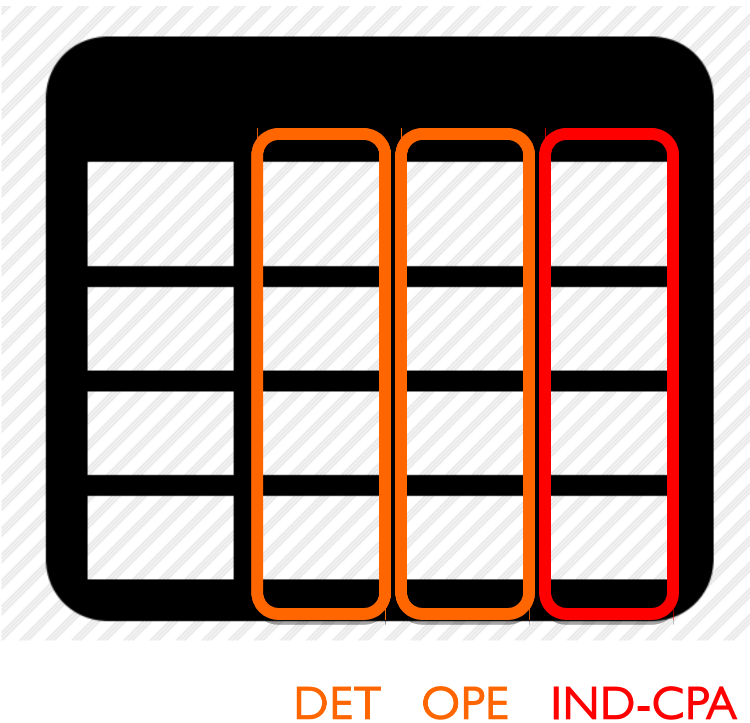

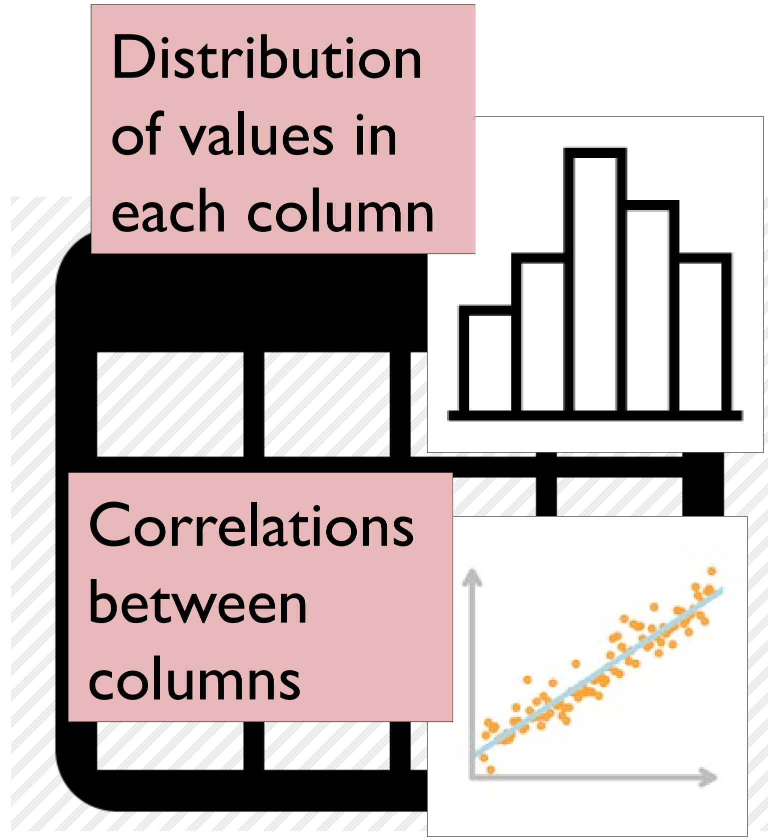

Public auxiliary data
(e.g., previous release of similar datasets)

Bayesian Inference

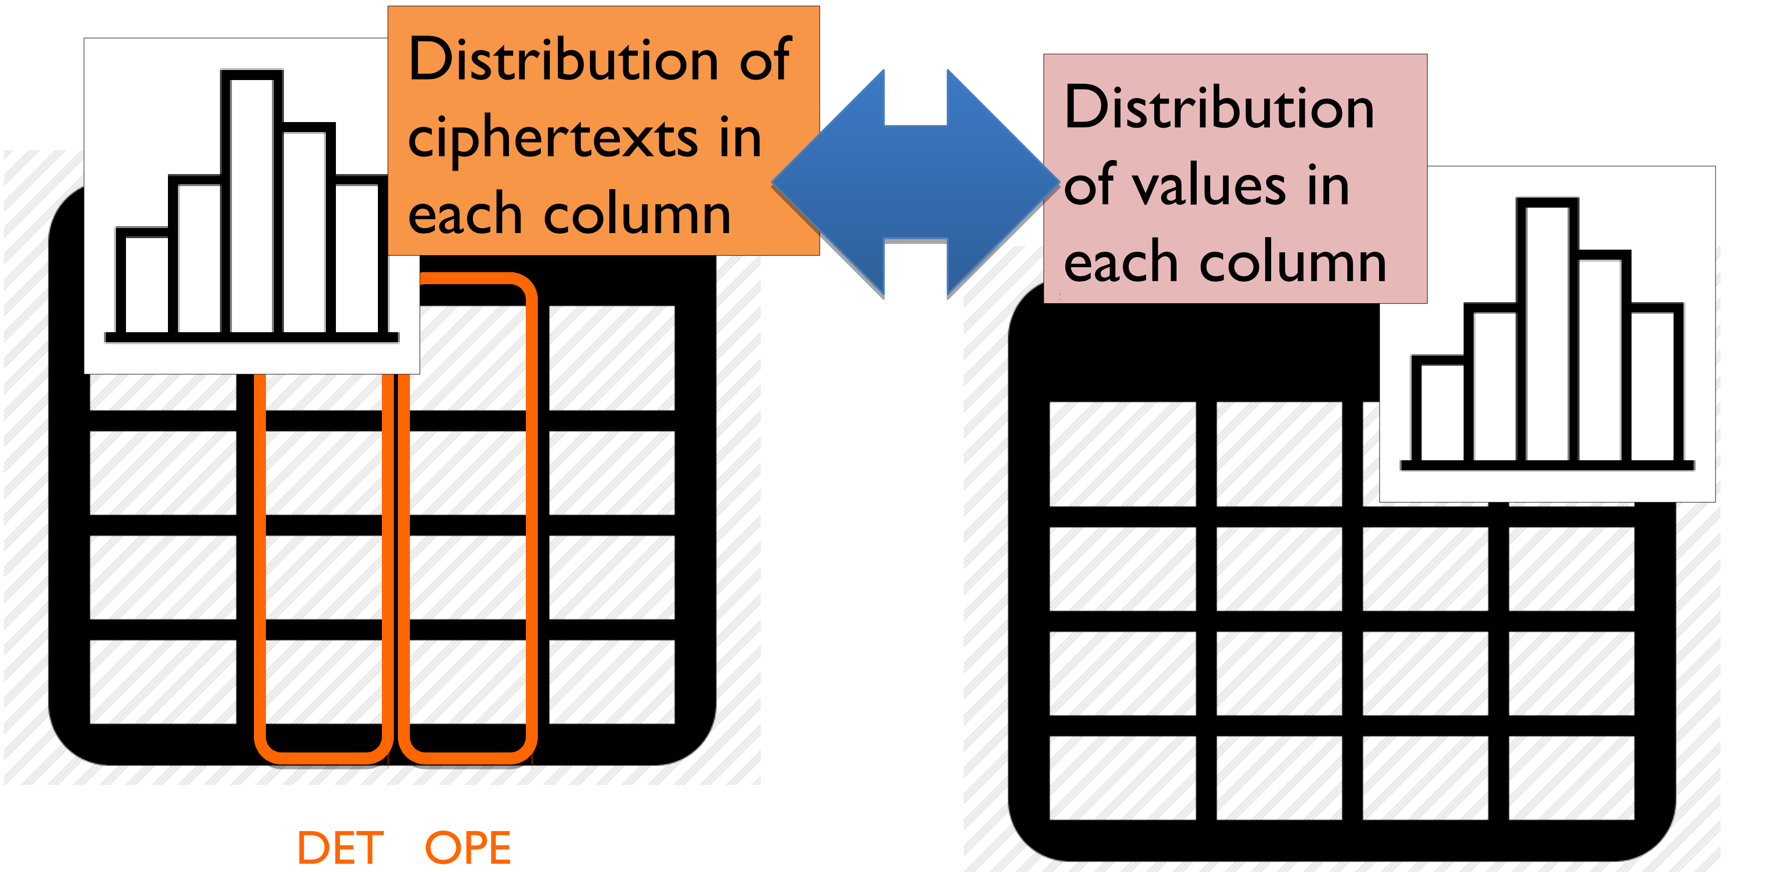

Public auxiliary data
(e.g., previous release of similar dataset)

Bindschaedler et al.

Multinomial Attack

Observed ciphertexts

Plaintext distribution (from auxiliary data)

$$\mathbb{P}\mathrm{r}\{\mathbf{f} = f \mid \vec{c}\,; \rho\} = \frac{\mathbb{P}\mathrm{r}\{\vec{c} \mid \mathbf{f} = f\,; \rho\} \cdot \mathbb{P}\mathrm{r}\{\mathbf{f} = f\,; \rho\}}{\mathbb{P}\mathrm{r}\{\vec{c}\,; \rho\}}$$

$$f_{\max} = \arg\max_{f} \; \mathbb{P}\mathrm{r}\{\mathbf{f} = f \mid \vec{c}\,; \rho\}$$

$$= \arg\max_{f} \; \mathbb{P}\mathrm{r}\{\vec{c} \mid \mathbf{f} = f\,; \rho\} \cdot \mathbb{P}\mathrm{r}\{\mathbf{f} = f\,; \rho\}$$

Most likely mapping of ciphertexts to plaintexts

Density of multinomial distribution

$$\mathbb{P}\mathrm{r}\{\vec{c} \mid \mathbf{f} = f\,; \rho\} = \mathbb{P}\mathrm{r}\{c_1, c_2, \ldots, c_n \mid \mathbf{f} = f\,; \rho\} = K_c \prod_{i=1}^{m} \rho_i^{c_{f(i)}}$$

<u>*Bindschaedler et al.*</u>

Multinomial Attack

- Optimal
  - Maximum likelihood estimator for deterministic ORE
- Outperforms previous heuristics
  - Naveed et al. frequency analysis (CCS 2015)
  - Grubbs et al. non-crossing attacks (Oakland 2017)
- Extends to multiple columns
  - Condition distribution on previously recovered plaintexts for a dependent column

Inferring "Sensitive" Columns

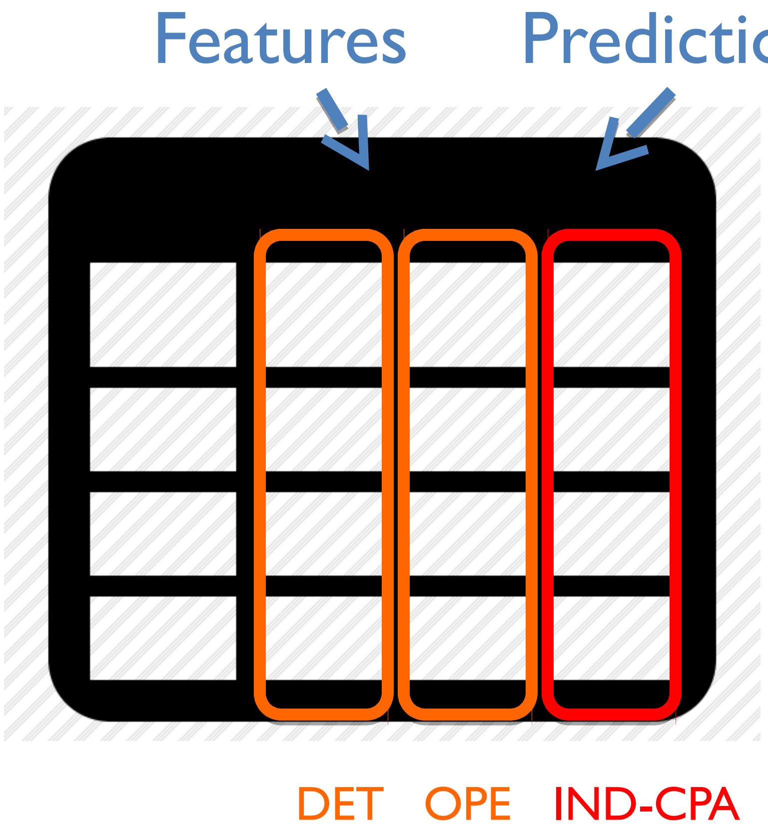

Multinomial attack!

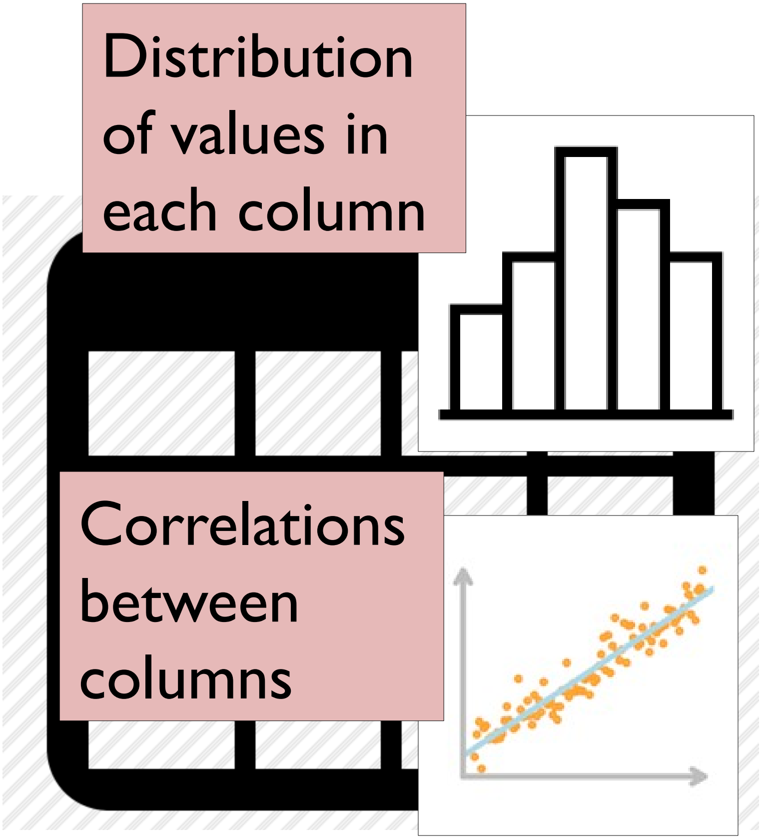

Public auxiliary data
(e.g., previous release of similar dataset)

Let's Try with Real Data

- Over 7 million hospital discharge records each year
- Demographic + medical attributes

- Over 3 million records each year
- Demographic attributes, income

- Data dump from 2015 hack
- Names and addresses of over 600,000 police officers

Empirical Results

- HCUP-NIS hospital discharge records
  - Infer if patient has a mental health or substance abuse condition with 97% accuracy
  - … mood disorder with 96% accuracy
- U.S. Census American Community survey
  - Recover 90% of PRE-encrypted attributes
  - Infer income to within $8.4K
- Fraternal Order of Police (FOP) data dump
  - Exact home addresses of 5,500 police officers in PA

Remember

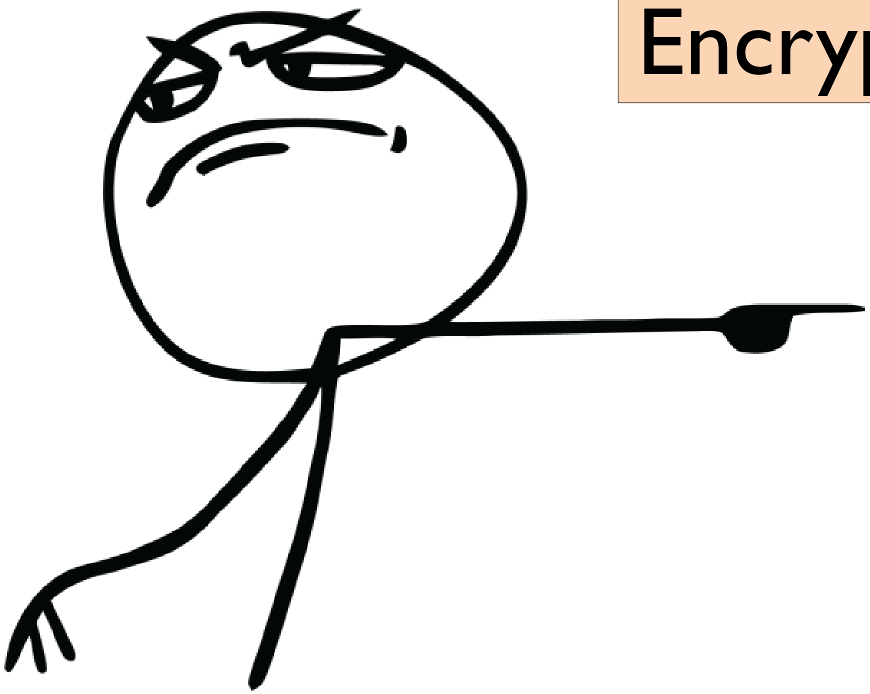

Encryption scheme is “secure”

does <u>not</u> mean

The system is “secure”

Advice to Practitioners