



# Building systems that compute on encrypted data

Raluca Ada Popa  
MIT

Security experts warn of increasing data breaches and privacy risks

**LivingSocial Hacked — More Than 50 Million Customer Names, Emails, Birthdates and Encrypted Passwords**

agents were

Suggested Content

Accessed (Intern Some Victims of Online Hacking Edge Into the Lig

APRIL 26, 2013 AT 1:15 PM PT

LivingSocial, the daily owned in part by A

**Compromise of confidential data is prevalent**

WordPress firm Automattic suffers root-level hack

Don't Trust Facebook with employee's Revelations


**Privacy, security still top cloud concerns**

Asia Cloud Forum editors | November 13, 2013

Asia Cloud Forum

An online survey of Microsoft partners has revealed that traditional concerns about

# Problem setup



no computation

storage



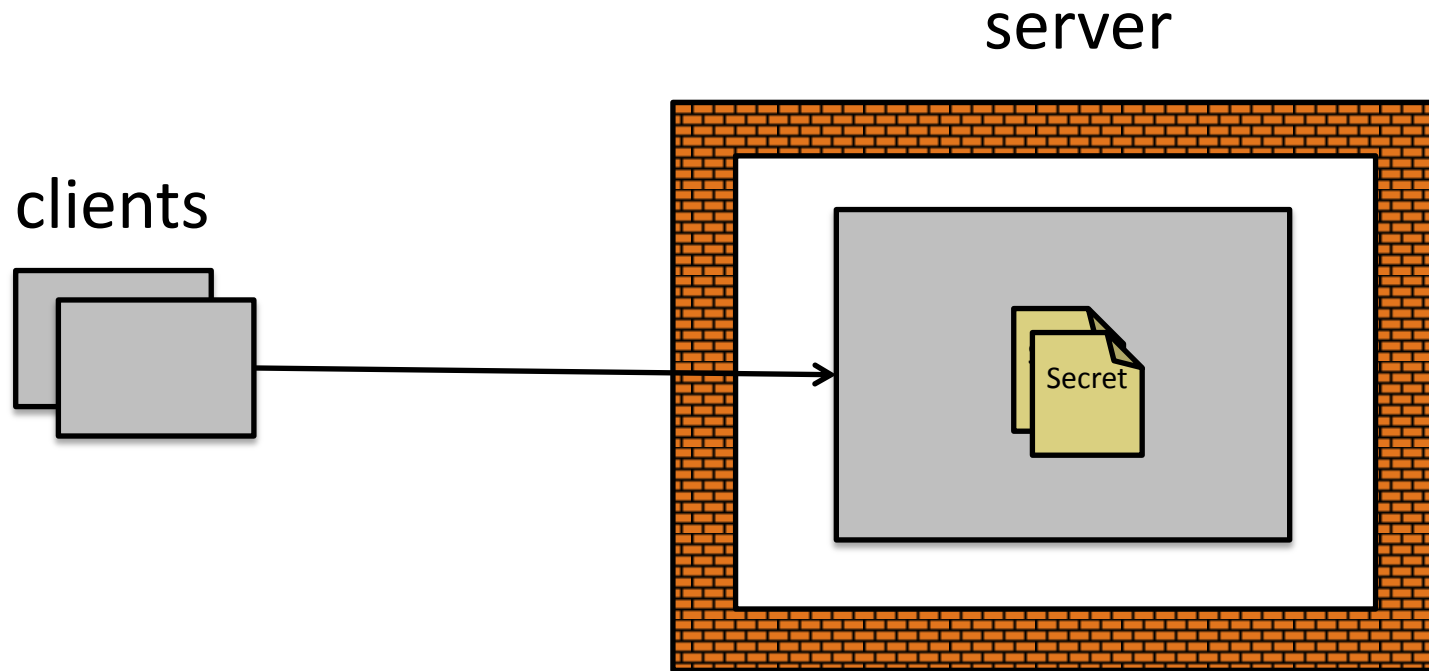
encryption

computation

databases, web applications, mobile applications, machine learning, etc.

??

# Current systems strategy



Prevent attackers from breaking into servers

# Lots of existing work

- Checks at the operating-system level
- Language-based enforcement of a security policy
- Static or dynamic analysis of application code
- Checks at the network level
- Trusted hardware

...


**Data still leaks even with  
these mechanisms**

because

**attackers eventually break in!**

# Attacker examples


## Attacker:



increasingly many companies store data on external clouds



accessed private data according to



## Reason they succeed:

software is complex

**insiders: legitimate server access!**

e.g., physical access


# My work

Systems that protect confidentiality even against  
attackers with access to all server data



# My approach

Servers store, process, and compute on encrypted data *in a practical way*



# Computing on encrypted data in cryptography

[Rivest-Adleman-Dertouzos'78]

Fully homomorphic encryption (**FHE**) [Gentry'09]

**prohibitively slow, e.g., slowdown X 1,000,000,000**

My work: **practical systems**


real-world  
performance

+

large class of  
real  
applications

+

meaningful  
security




# My contributions




System:

Server under attack:

Databases: **CryptDB** [SOSP'11][CACM'12]  
mOPE, adjJOIN  
[Oakland'13]



Web apps: **Mylar** [NSDI'14]  
multi-key search



Mobile apps: **PrivStats** [CCS'11]  
**VPriv** [Usenix Security'09]

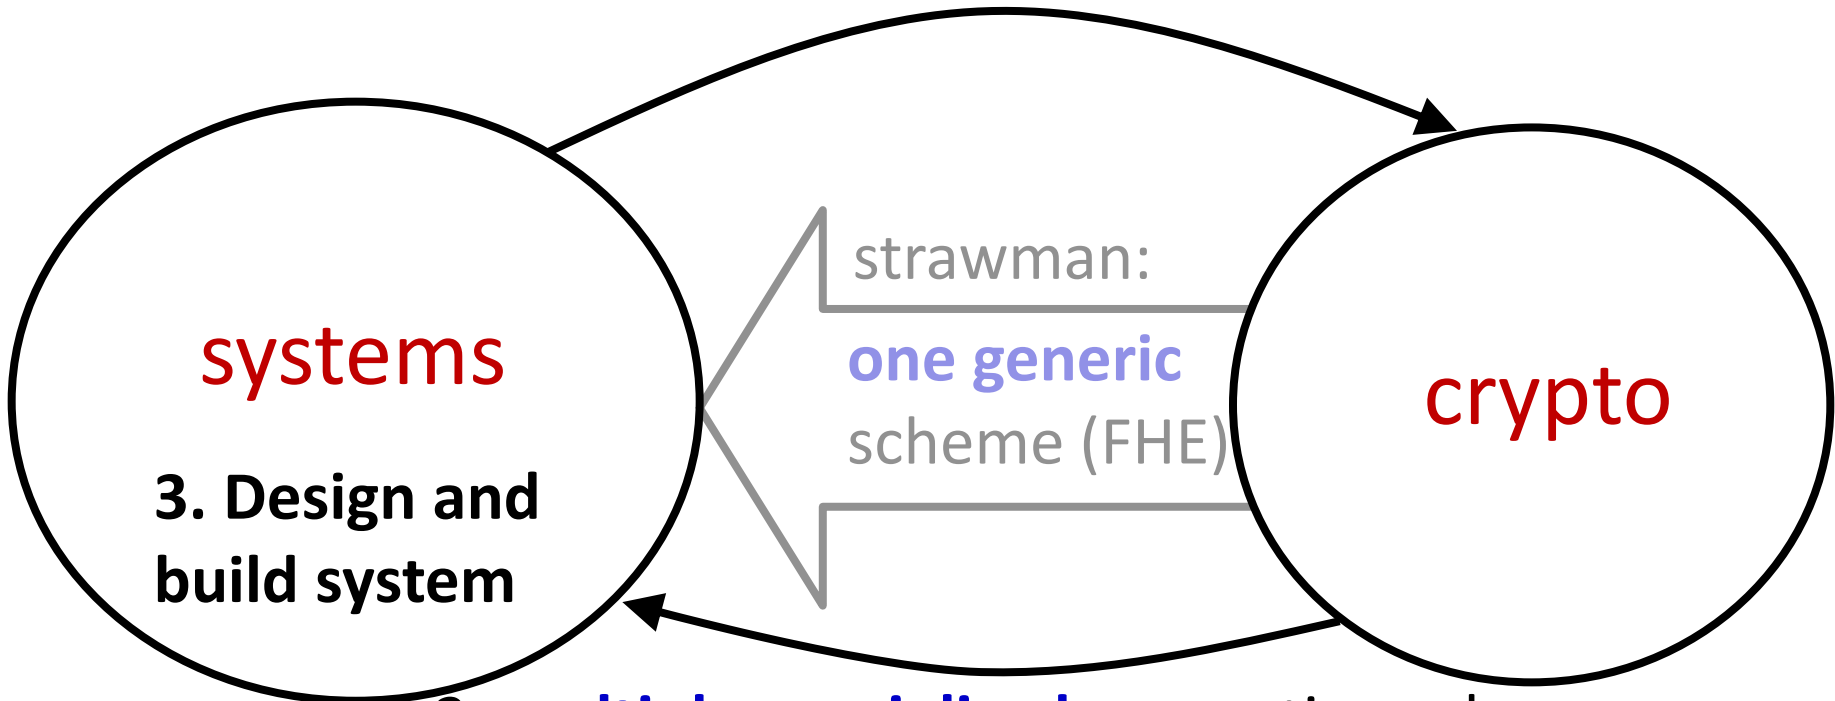


Theory:

In general: **Functional encryption** [STOC'13] [CRYPTO'13]

# Combine systems and cryptography

1. identify core operations needed



2. **multiple specialized** encryption schemes

New schemes:

- mOPE, adjJOIN for CryptDB
- multi-key search for Mylar


# My contributions




System:

Server under attack:


Databases: **CryptDB**



Web apps: **Mylar**



Mobile apps: **PrivStats**  
**VPriv**



Theory:

In general: **Functional encryption**

# CryptDB

[SOSP'11: **Popa**-Redfield-Zeldovich-Balakrishnan]

First **practical** database system (DBMS) to process most SQL queries on encrypted data


# Related work

- **Systems work:** [Hacigumus et al.'02][Damiani et al.'03][Ciriani et al.'09]
  - no formal confidentiality guarantees
  - restricted functionality
  - client-side filtering
- **Theory work:**
  - General computation: FHE [Gentry'09]
    - very strong security: **forces slowdown - many queries must always scan and return the whole DB**
    - prohibitively slow ( $10^9\times$ )
  - Specialized schemes [Amanatidis et al.'07][Song et al.'00][Boldyreva et al.'09]

# Setup

trusted client-side

under passive attack



Use cases:

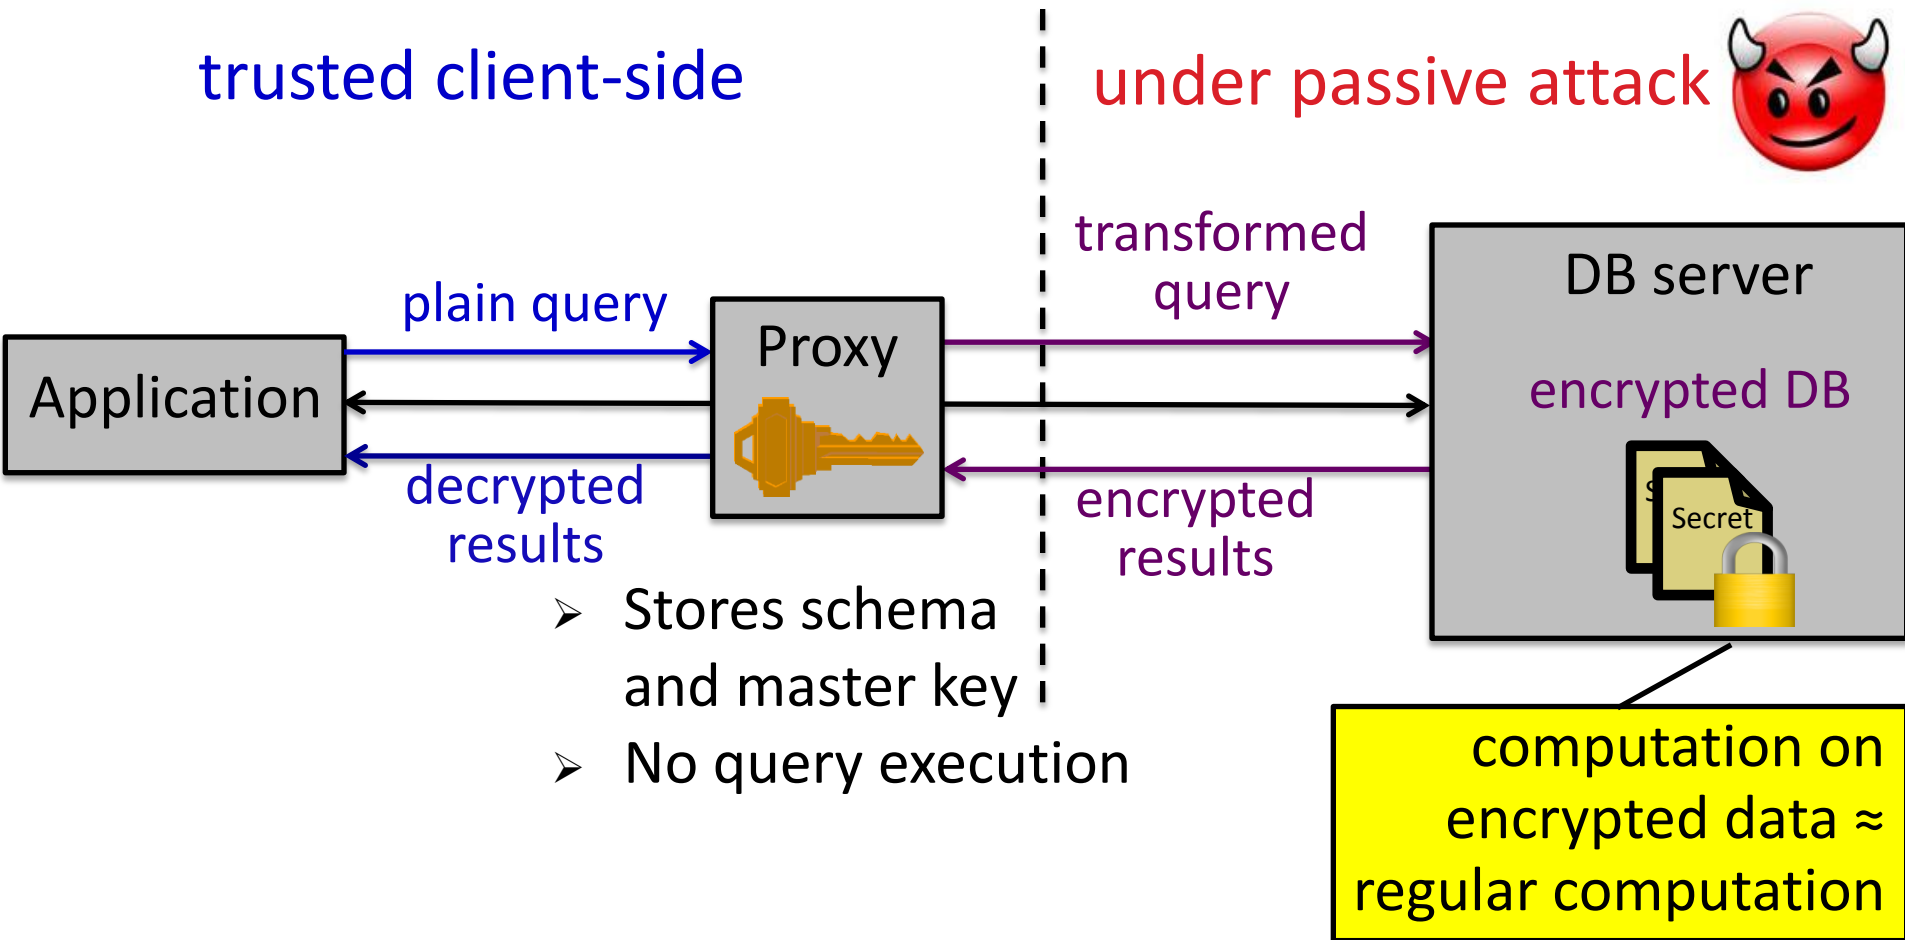
- Outsource DB to the cloud (DBaaS)
  - e.g. Encrypted BigQuery
- Local cluster: hide DB content from sys. admins.



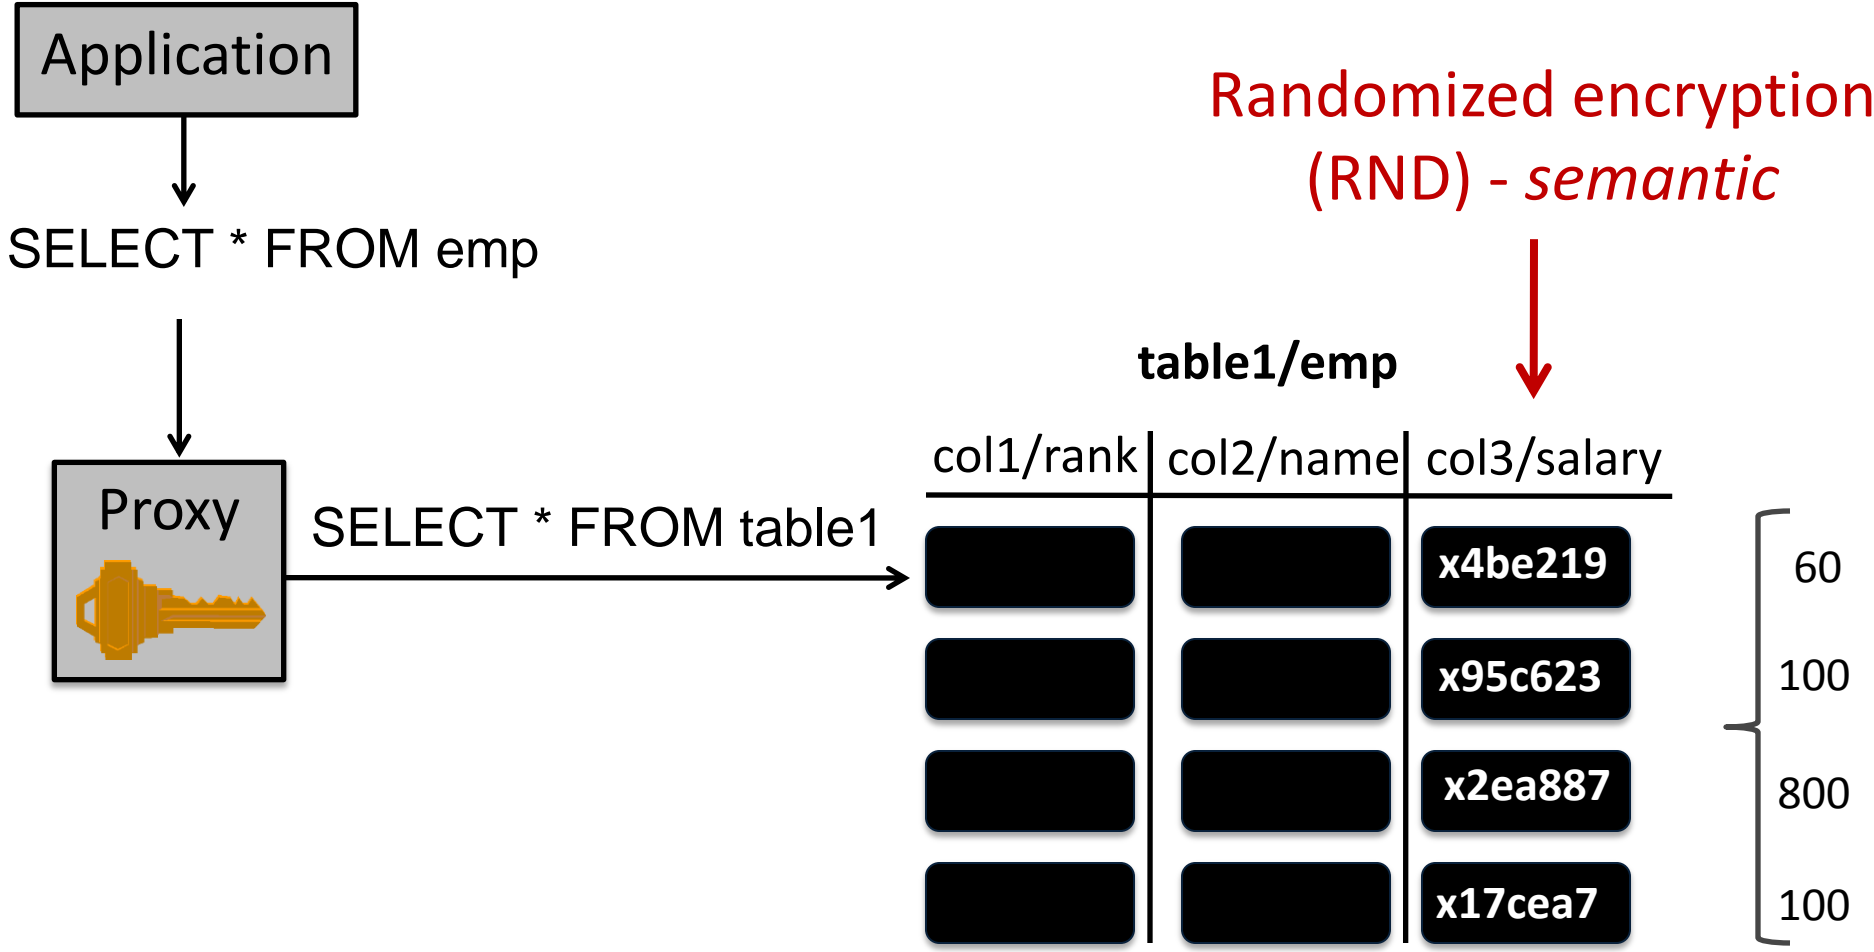
# Setup

trusted client-side

under passive attack



# Example



# Example

Application

SELECT \* FROM emp  
WHERE salary = 100

Proxy  


SELECT \* FROM table1  
WHERE col3 = x5a8c34

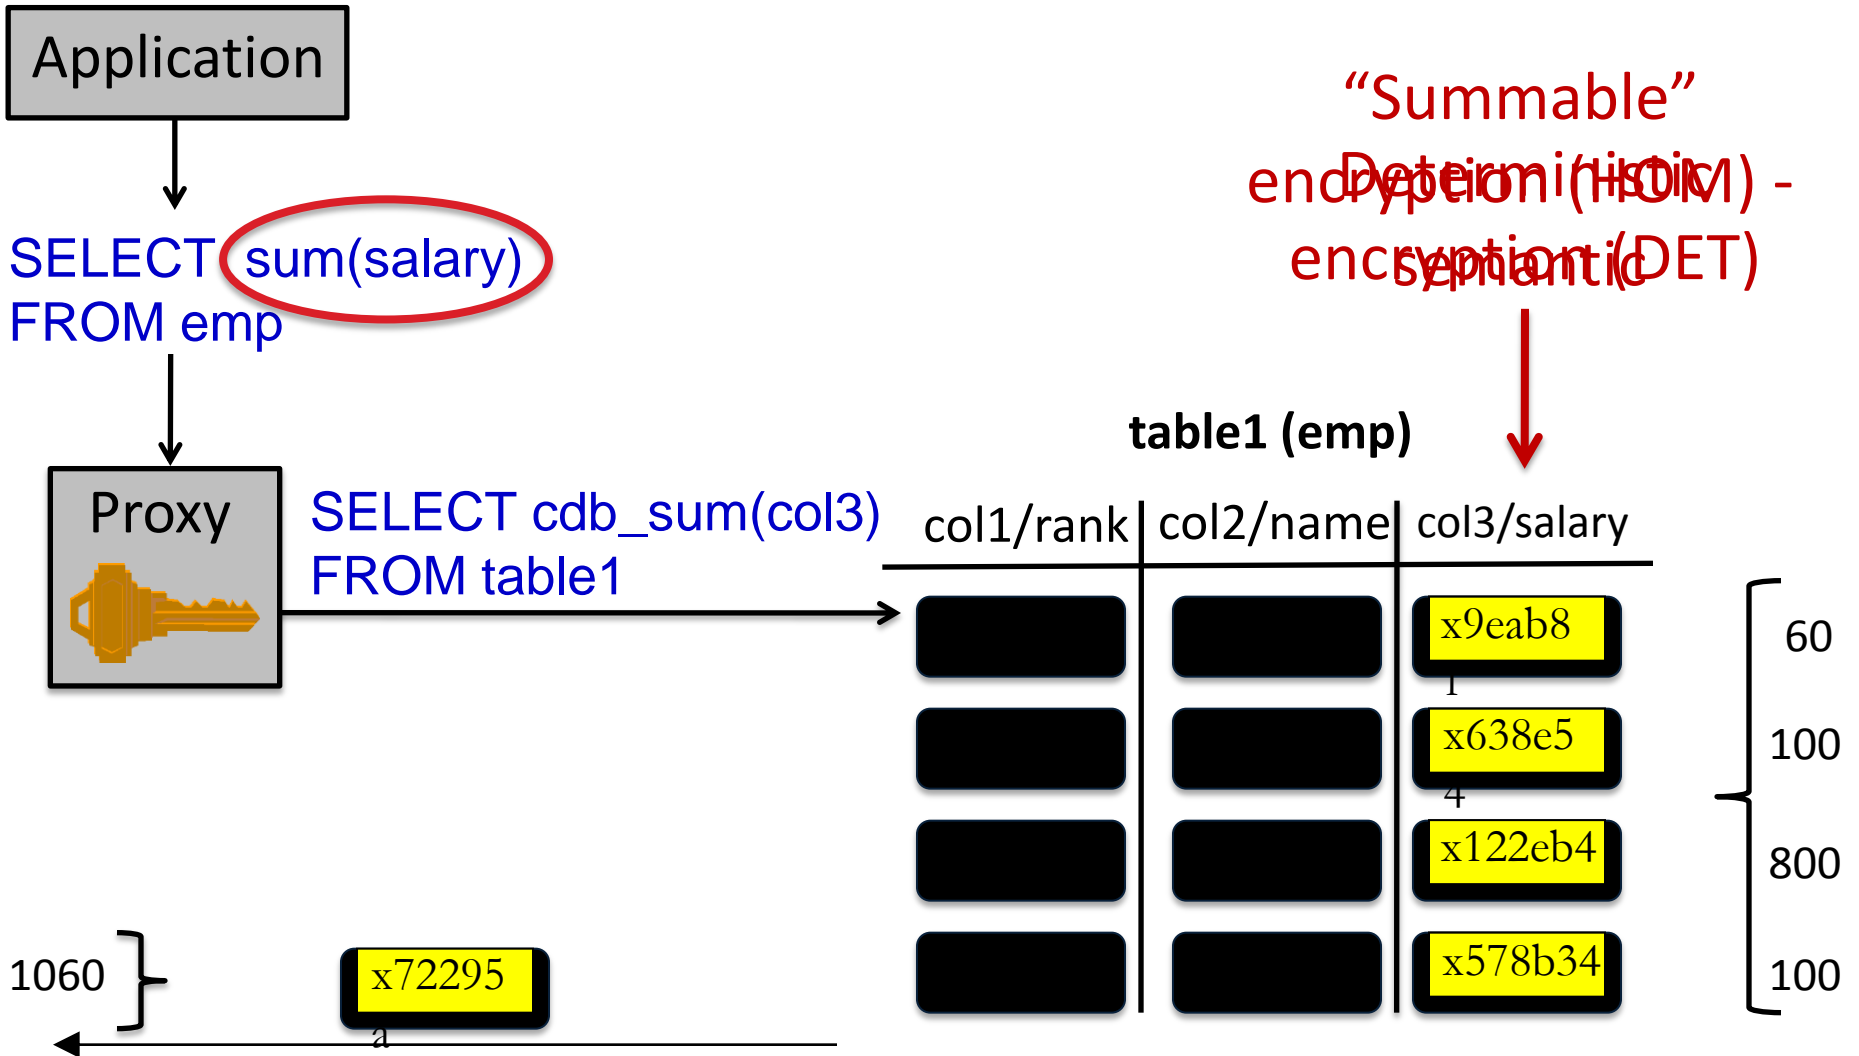
Random encryption  
(DET)

table1/emp

col1/rank	col2/name	col3/salary	
		x934bc1	60
		x5a8c34	100 ✓
		x84a21c	800
		x5a8c34	100 ✓

		x5a8c34
		x5a8c34

# Example



# Techniques

1. Use SQL-aware set of efficient encryption schemes (meta technique!)



Most SQL can be implemented with a few core operations

2. Adjust encryption of data based on queries

3. Query rewriting algorithm

# 1. SQL-aware encryption schemes


Security	Scheme	Constructio	Function	SQL operations:
≈ semantic security	RND	<sup>n</sup> AES in UFE	data moving	e.g., SELECT, UPDATE, DELETE, INSERT, COUNT
	HOM	Paillier	addition	e.g., SUM, +
	SEARCH	Song et al., '00	word search	restricted ILIKE
reveals only repeat pattern	DET	AES in CMC	equality	e.g., =, !=, IN, GROUP BY, DISTINCT
	JOIN	our new scheme	join	
reveals only order	OPE	our new scheme [Oakland'13]	order	e.g., >, <, ORDER BY, ASC, DESC, MAX, MIN, GREATEST, LEAST

$x < y \iff \text{Enc}(x) < \text{Enc}(y)$

# How to encrypt each data item?


- Goals:
1. Support queries
  2. Use most secure encryption schemes

Challenge: may not know queries ahead of time



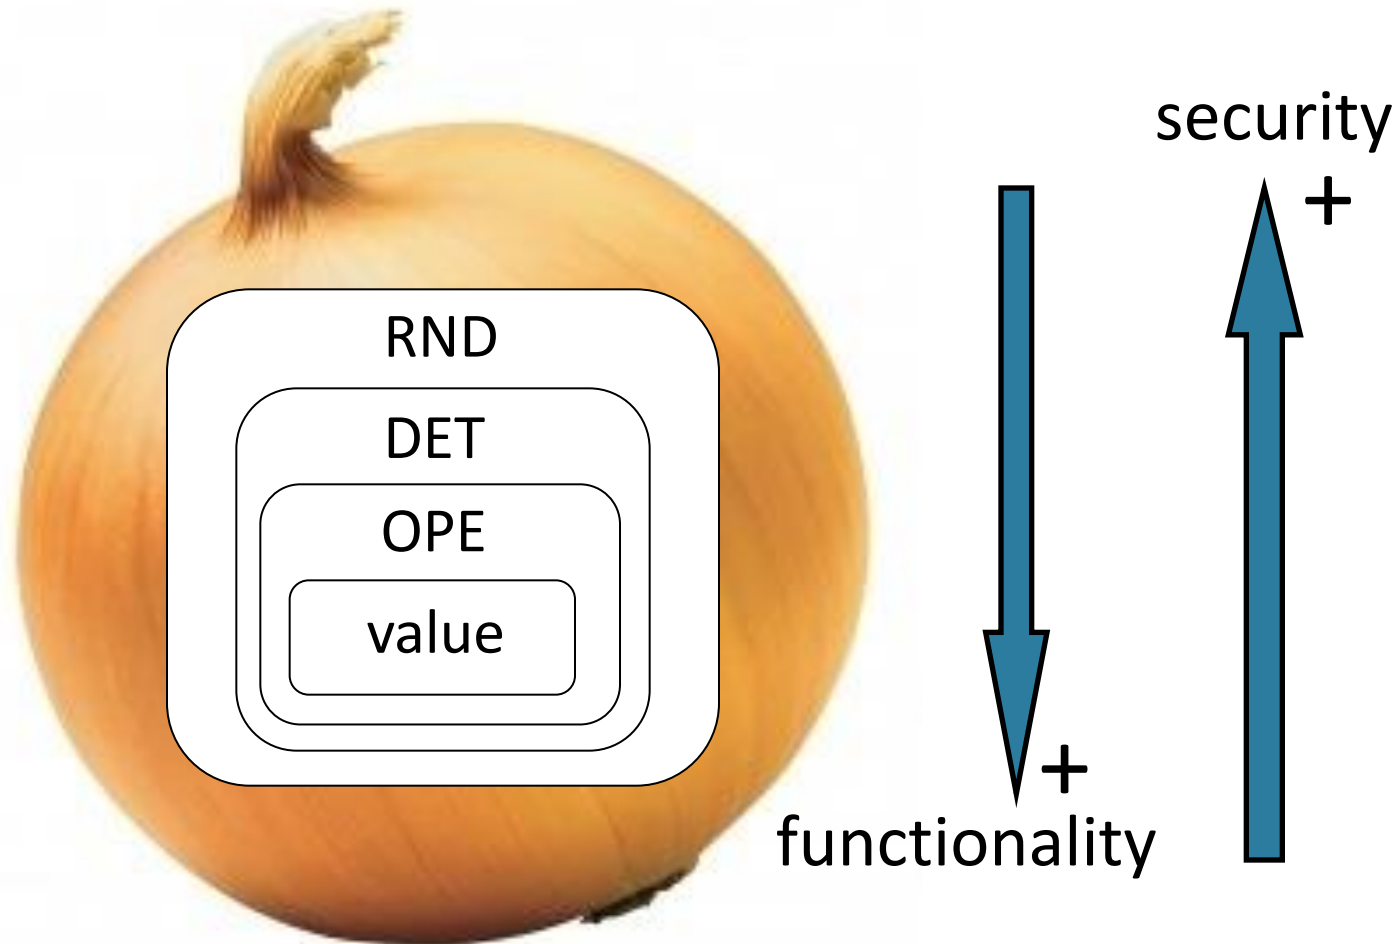
Leaks order!

# Onion





# Onion of encryptions




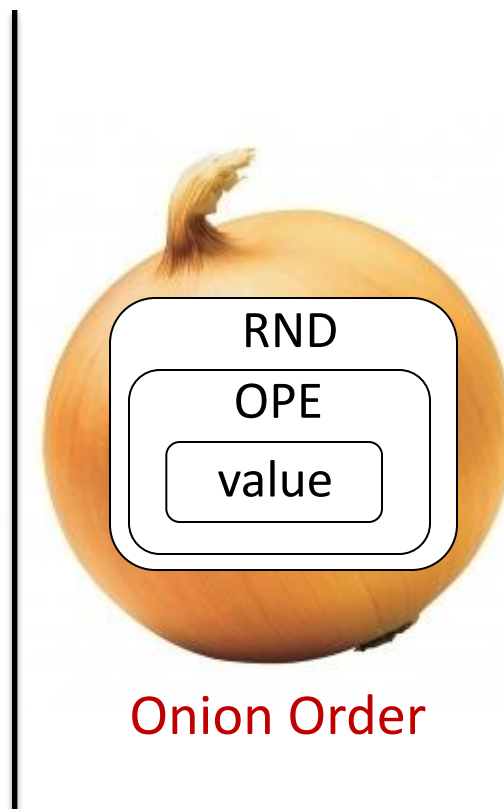

Adjust encryption: strip off layer of the onion

# Onions of encryptions

1 column

3 columns

each value →



Same key for all items in a column for same onion layer

# Onion evolution

- Start out the database with the most secure encryption scheme
- If needed, adjust onion level
  - Proxy gives decryption key to server
  - Proxy remembers onion layer for columns

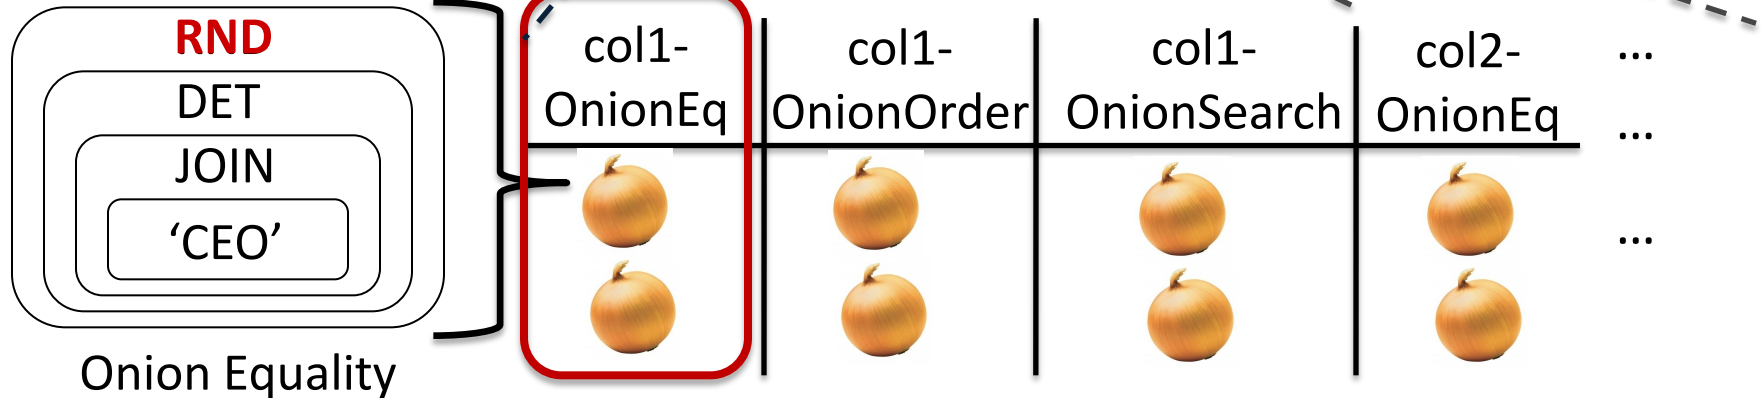
Lowest onion level is never removed

# Example

Logical table:


emp:		
rank	name	salary
'CEO'		
'worker'		

Physical table:



```
SELECT * FROM emp WHERE rank = 'CEO'
```

# Example (cont'd)



```
SELECT * FROM emp WHERE rank = 'CEO'
```



```
UPDATE table1 SET col1-OnionEq =
```

```
Decrypt_RND(key, col1-OnionEq)
```

```
SELECT * FROM table1 WHERE col1-OnionEq = xda5c0407
```

# Security threshold

Data owner can specify minimum level of security

```
CREATE TABLE emp (... , credit_card SENSITIVE integer, ...)
```



RND, HOM, DET for unique fields  
≈ semantic security

# Security guarantee

Columns annotated as sensitive have semantic security (or similar).

Encryption schemes exposed for each column are the most secure enabling queries.

equality → repeats

sum → semantic

no filter → semantic


*common in practice*

# Limitations & Workarounds

Queries not supported:

- More complex operators, e.g., trigonometry
- Certain combinations of encryption schemes:
  - e.g., salary + raise > 100K

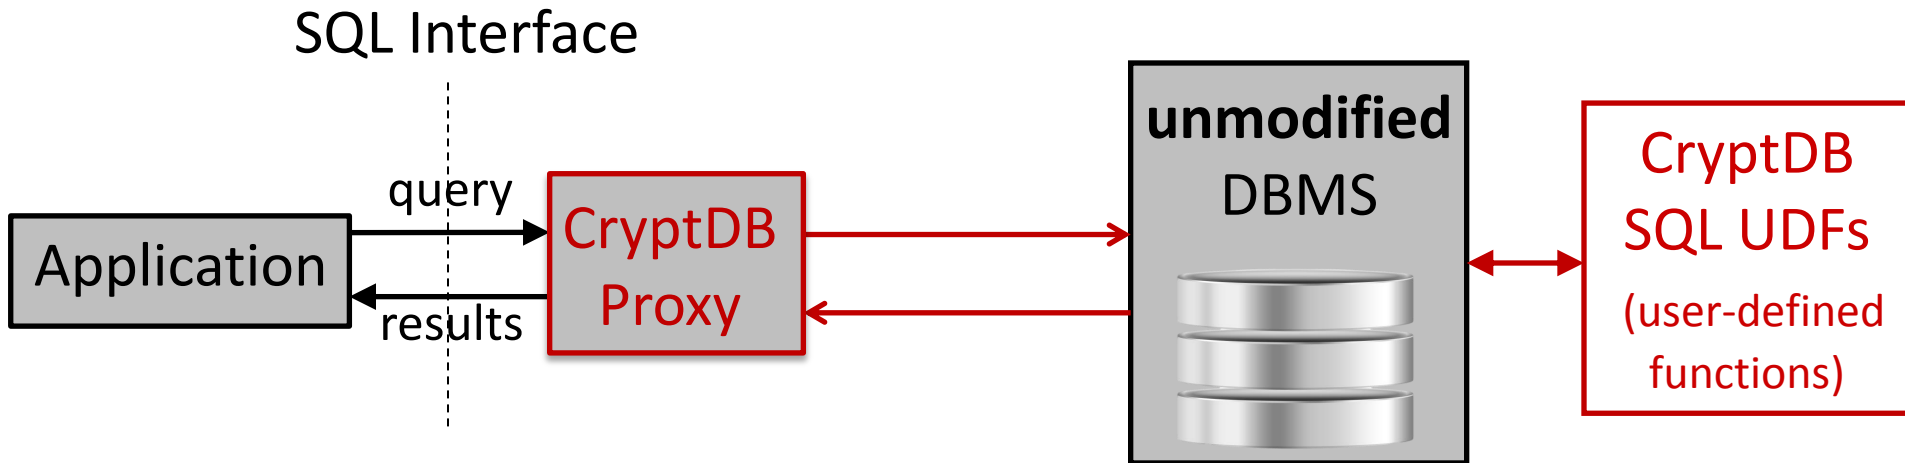
HOM



use query splitting, query rewriting



# Implementation



**No change to the DBMS!**

**Largely no change to apps!**

# Evaluation

1. Does it support real queries/applications?
2. What is the resulting confidentiality level?
3. What is the performance overhead?

# Real queries/applications

Application	Encrypted columns	# cols with queries not supported
phpBB	23	0
HotCRP	22	0
grad-apply	103	0
TPC-C	92	0
sql.mit.edu	128,840	1,094

apps with sensitive columns

tens of thousands of apps

`SELECT 1/log(series_no+1.2) ...  
... WHERE sin(latitude + PI()) ...`

# Confidentiality level

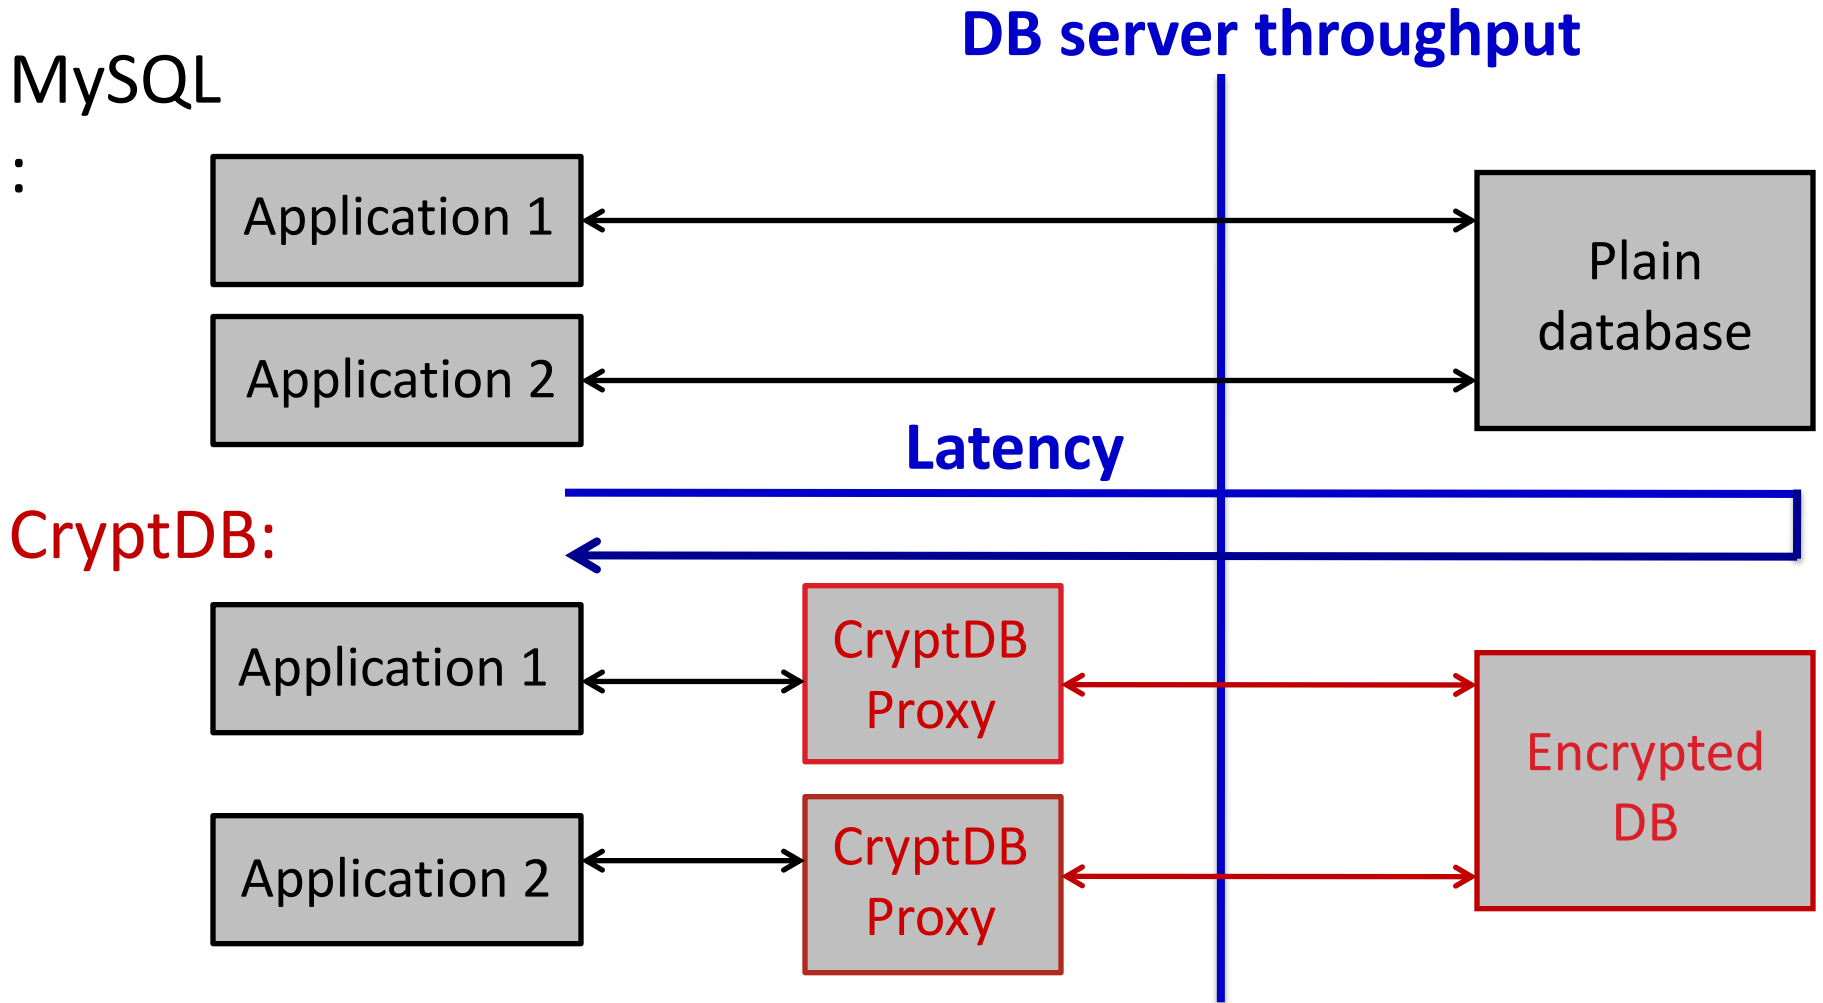
Final onion state

Application	Encrypted columns	Min level: ≈semantic	Min level: DET/JOIN	Min level: OPE
phpBB	23	21	1	1
HotCRP	22	18	1	2
grad-apply	103	95	6	2
TPC-C	92	65	19	8
sql.mit.edu	128,840	80,053	34,212	13,131

Most  
columns at  
semantic

Most columns at  
OPE were less  
sensitive

# Performance




Hardware: 2.4 GHz Intel Xeon E5620 – 8 cores, 12 GB RAM

# TPC-C performance

Latency (per query): 0.10ms MySQL vs. 0.72ms CryptDB

Throughput loss over MySQL: 26%



No cryptography at the DB server in the steady state!

# Adoption

<http://css.csail.mit.edu/cryptdb/>

Google

Encrypted BigQuery [\[http://code.google.com/p/encrypted-bigquery-client/\]](http://code.google.com/p/encrypted-bigquery-client/)



Úlfar Erlingsson,  
head of security  
research, Google

“CryptDB was really eye-opening in establishing the practicality of providing a SQL-like query interface to an encrypted database”

“CryptDB was [...] directly influential on the design and implementation of Encrypted BigQuery.”



SEED implemented on top of the SAP HANA DBMS



Encrypted version of the D4M Accumulo NoSQL engine

[sql.mit.edu](http://sql.mit.edu)


Users opted-in to run Wordpress over our CryptDB source code

Demo




# Attack to all servers?

users






CryptDB  
SQL queries on  
encrypted DB

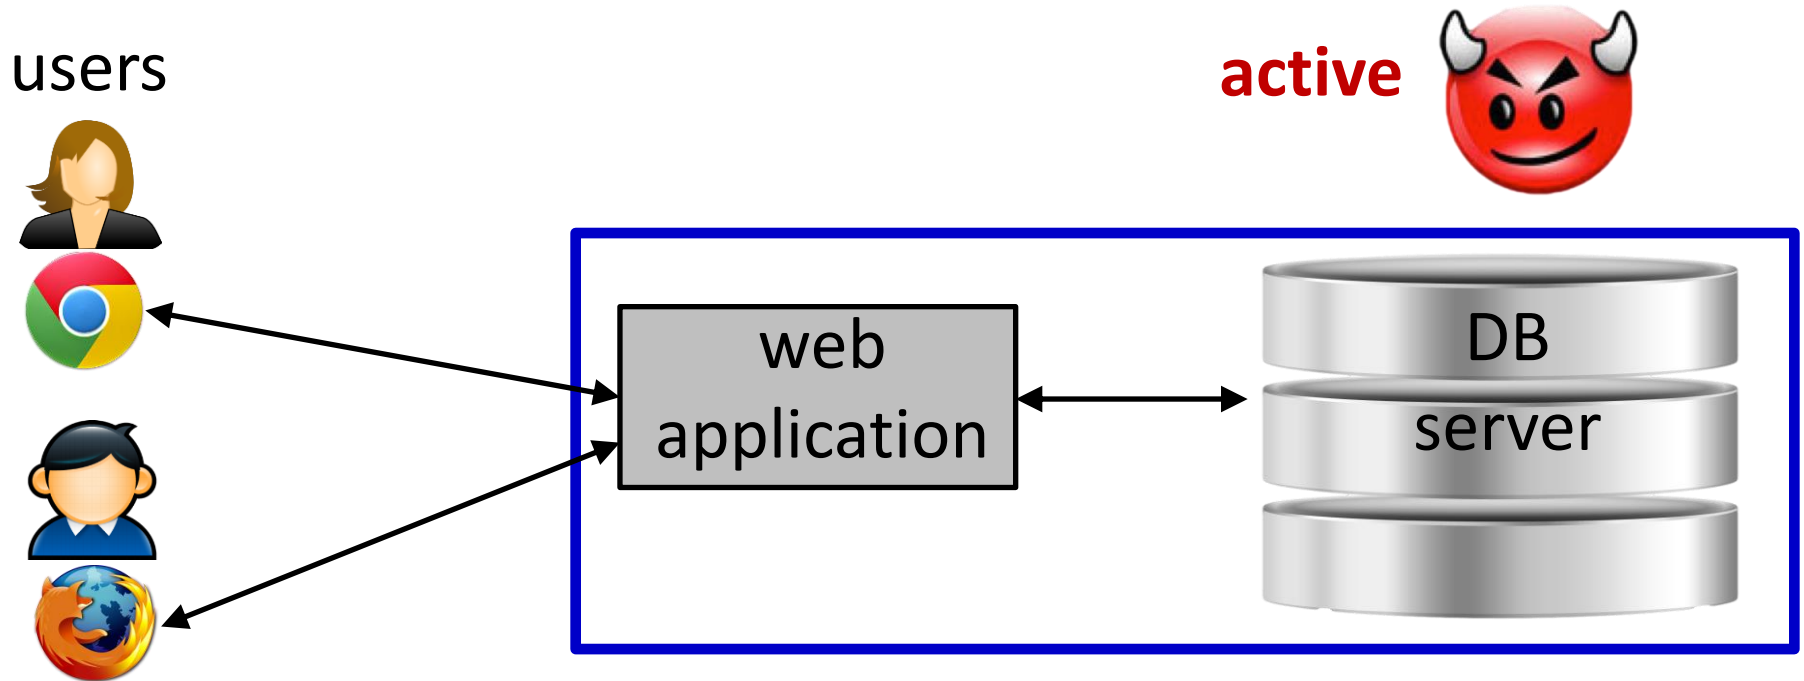
# Attack to all servers?



users

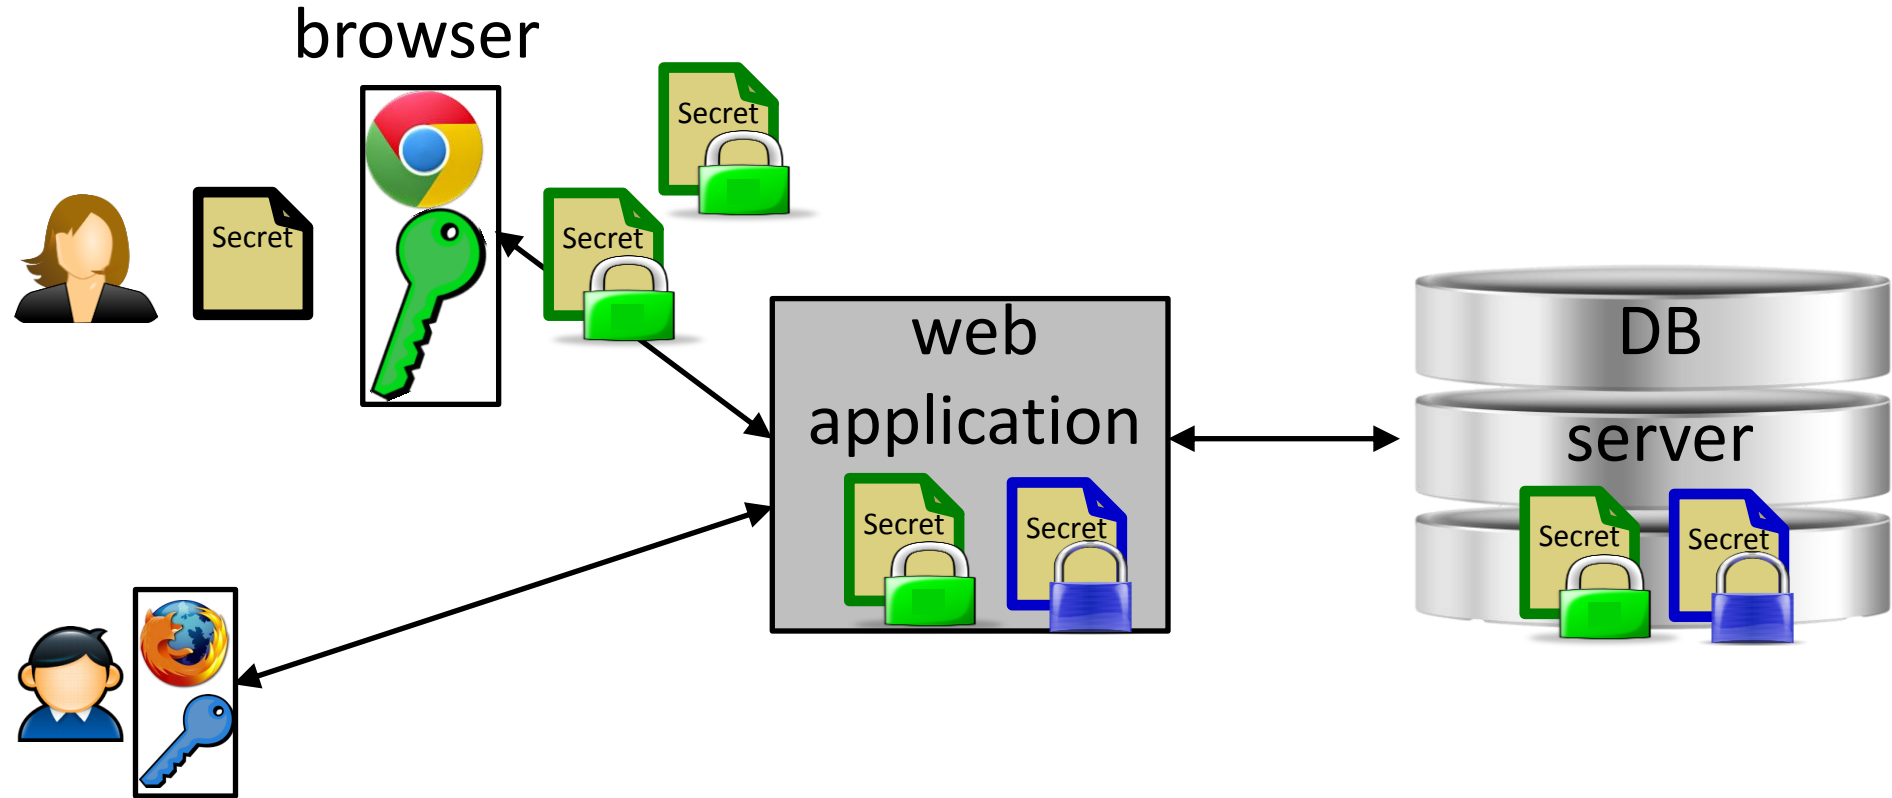


# Mylar



- Framework for building web applications
- Protects confidentiality against attacks to all servers

# Overview



Plaintext data exists only in browsers

# Computation in web applications

1. Mylar is a client-side application framework
2. Non client-side computation: **meta technique!**

- data sharing
- search

## Challenges

- Active attacker
  - key certification
- Multiple keys
  - multi-key search

# Applications

<http://css.csail.mit.edu/mylar/>

chat    medical    class website    forum    calendar    photo sharing

Few developer annotations to secure an application,  
modest overhead



The screenshot shows the login interface for the Endometriosis App. At the top left is the Newton-Wellesley Hospital (NWH) logo, a blue shield with white waves and the letters 'NWH'. To its right, the text 'NEWTON-WELLESLEY HOSPITAL' is displayed in a large, grey, serif font. Below the logo and hospital name, the title 'Endometriosis App' is written in a bold, blue, sans-serif font. Underneath the title, a grey instruction reads 'Please sign in using your email and your password'. There are two input fields: 'Email Address' and 'Password', both with light grey borders. Below the 'Password' field are two buttons: a dark grey 'Sign in' button and an orange 'Forgot Password' button. On the left side of the login form, there is a yellow padlock icon with a black keyhole, indicating a secure login process.


# My contributions




**System:**

**Server under attack:**


Databases: **CryptDB** [SOSP'11][CACM'12]  
mOPE, adjJOIN  
[Oakland'13]



Web apps: **Mylar** [NSDI'14]  
multi-key search




Mobile apps: **PrivStats** [CCS'11]  
**VPriv** [Usenix Security'09]



**Theory:** **Functional encryption** [STOC'13] [CRYPTO'13]

- Proof of concept for general functions
- Solved old open problem: **reusable garbled circuits**

# System design principles



Assume all server data will leak!

Store, process, and compute on encrypted data.

Technique for practicality:

1. identify core operations
2. use an efficient encryption scheme for each



# Future work

Other systems computing on encrypted data:

Genomics analytics and machine learning


# Future work

## Other systems computing on encrypted data:

Genomics analytics and machine learning

### Big data & compression

big data



compressed  
big data



encrypted big data



compressed  
big data

How to compute on it??

# Future work

## Other systems computing on encrypted data:


Genomics analytics and machine learning

Big data & compression

## Security beyond confidentiality:

Correctness of computation

## Client-side security



# Collaborators

**CryptDB:** Catherine Redfield, Nikolai Zeldovich, Hari Balakrishnan, Aaron Burrows

**Mylar:** Steven Valdez, Jonas Helfer, Nikolai Zeldovich,  
Frans M. Kaashoek, Hari Balakrishnan

**PrivStats, VPriv:** Andrew Blumberg, Hari Balakrishnan, Frank H. Li

**Functional encryption:** Shafi Goldwasser, Yael Kalai, Vinod  
Vaikuntanathan, Nikolai Zeldovich

and others for other projects.

# Future work

## Other systems computing on encrypted data:


Genomics analytics and machine learning

Big data & compression

## Security beyond confidentiality:

Correctness of computation

## Client-side security



**THANK  
YOU!**