



# Spectrum

## High-Bandwidth Anonymous Broadcast

Zachary Newman ([zjn@mit.edu](mailto:zjn@mit.edu)) & Sacha Servan-Schreiber ([3s@mit.edu](mailto:3s@mit.edu))

Joint work with Sridhar Devadas (MIT)

# Elevator Pitch

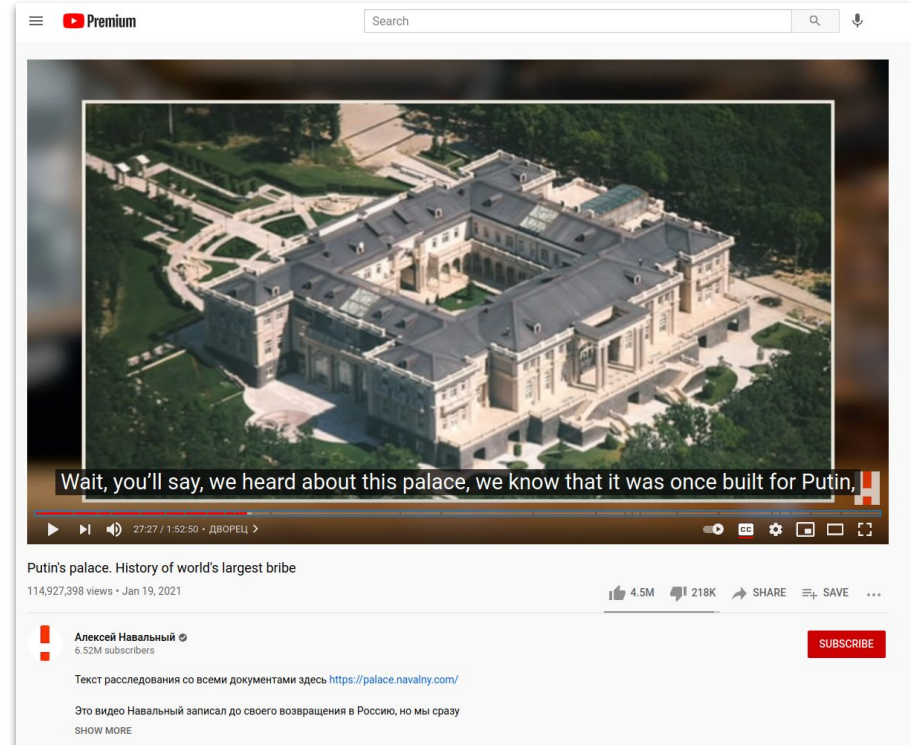
**Previous systems: anonymous Twitter.**

**Us: anonymous Twitch.**

## **Highlights:**

- 3–140x faster; overhead ~200 bytes compared to non-private broadcast
- Security against active attacks
- Share full-length documentary in 1.5 hours (w/ 10K users)

# Whistleblowers expose corruption



# But often at great cost

## Whistleblowing—*Is It Really Worth the Consequences?*

by Kayla L. Delk, JD, BSN, RN

This article explores general principles of state whistleblower laws and alerts nurses to considerations when deciding whether to report an employer violation. Reporting an employer violation can be difficult for any employee but especially for nurses because nurses are torn between their desire and duty to advocate for clients' safety and their desire to maintain employment. The author suggests questions to consider when deciding to "blow the whistle" and alerts nurses to statutes of limitations that may affect when nurses must report violations. The article also illuminates policy and procedural issues for various states that affect how and to whom nurses report violations to protect themselves under whistleblower protection laws. Finally, this article explores personal and professional consequences that nurses should consider before reporting violations.


This article explores when and how to "blow the whistle" and alerts nurses to the consequences of blowing the whistle. However, this is merely a general overview of state whistleblower laws; nurses who intend to blow the whistle should seek legal advice from an attorney knowledgeable about each state's specific statutes regarding whistleblowing.

### BEFORE BLOWING THE WHISTLE

Because whistleblowing can have deleterious effects on nurses' professional and personal lives, they

## Putin critic Navalny jailed in Russia despite protests

© 2 February



Navalny says the latest case against him was fabricated

MEDIA


## Panama Papers journalist killed by car bomb

By Yaron Steinbuch


October 16, 2017 | 3:24pm | Updated




# Can they publish anonymously?




# Can they publish anonymously?



# Can they publish anonymously?




# Can they publish anonymously?







# Can they publish anonymously?



# Can they publish anonymously?



# Can they publish anonymously?



# Dining Cryptographer Nets (DC-nets)

# Simple XOR-based scheme (DC-net)

$$m_i = \mathbf{m} \quad \text{Broadcaster}$$

# Simple XOR-based scheme (DC-net)

$m_i = \mathbf{m}$     Broadcaster

$m_i = 0$     Other clients

# Simple XOR-based scheme (DC-net)

$m_i = \mathbf{m}$     Broadcaster

$m_i = 0$     Other clients

$r_i \oplus m_i$     to Server A

$r_i$     to Server B

# Simple XOR-based scheme (DC-net)

$m_i = \mathbf{m}$     Broadcaster

$m_i = 0$     Other clients

$r_i \oplus m_i$     to Server A

$r_i$     to Server B

$$\begin{aligned}
 \mathbf{m} &= \overbrace{\bigoplus_i^N (r_i \oplus m_i)}^{\text{Server A}} \oplus \overbrace{\bigoplus_i^N r_i}^{\text{Server B}} = \overbrace{\bigoplus_i^N (r_i \oplus r_i)}^{(\text{commutativity of xor})} \oplus \bigoplus_i^N m_i \\
 &= \underbrace{0 \oplus \dots \oplus \mathbf{m} \oplus \dots \oplus 0}_{\text{origin of } \mathbf{m} \text{ is hidden}}.
 \end{aligned}$$



# Simple XOR-based scheme (DC-net)

$m_i = \mathbf{m}$     Broadcaster


$r_i \oplus m_i$     to Server A

$m_i = 0$     Other clients


$r_i$     to Server B

$$\begin{aligned} \mathbf{m} &= \overbrace{\bigoplus_i^N (r_i \oplus m_i)}^{\text{Server A}} \oplus \overbrace{\bigoplus_i^N r_i}^{\text{Server B}} = \overbrace{\bigoplus_i^N (r_i \oplus r_i)}^{(\text{commutativity of xor})} \oplus \bigoplus_i^N m_i \\ &= \underbrace{0 \oplus \dots \oplus \mathbf{m} \oplus \dots \oplus 0}_{\text{origin of } \mathbf{m} \text{ is hidden}}. \end{aligned}$$

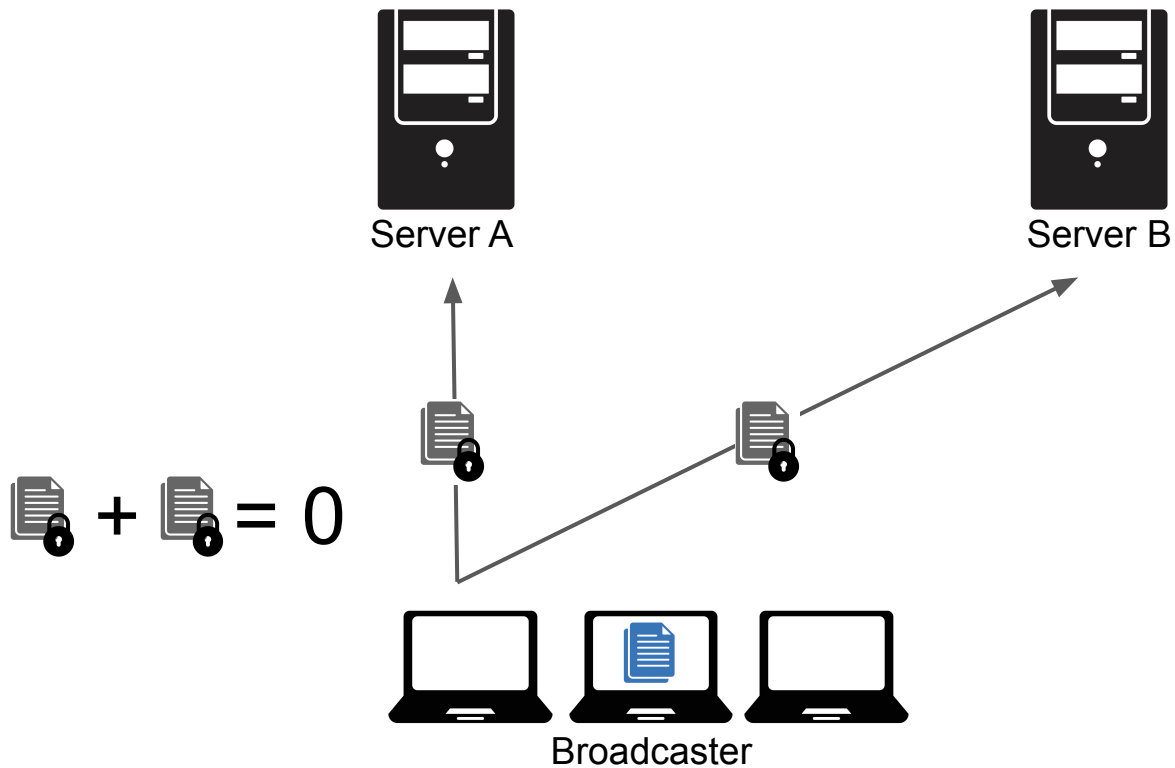
# A toy protocol




# A toy protocol




# A toy protocol




# A toy protocol



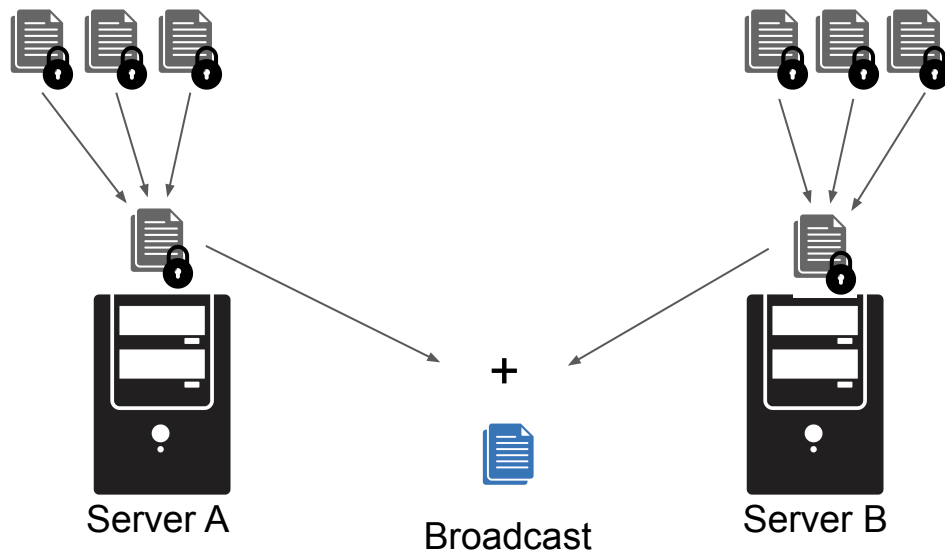
# A toy protocol



# A toy protocol



# A toy protocol



Broadcast origin hidden



# DC-Nets


## Pros

- **Fast:** server XORs requests
- **Metadata private:** fully hides broadcast source
- **High bandwidth:** throughput matches client upload bandwidth



## Cons

- **Not-scalable:** only one broadcaster (unless repeated  $\Rightarrow$  extra bandwidth)
- **Insecure:** client can *undetectably* disrupt the broadcast

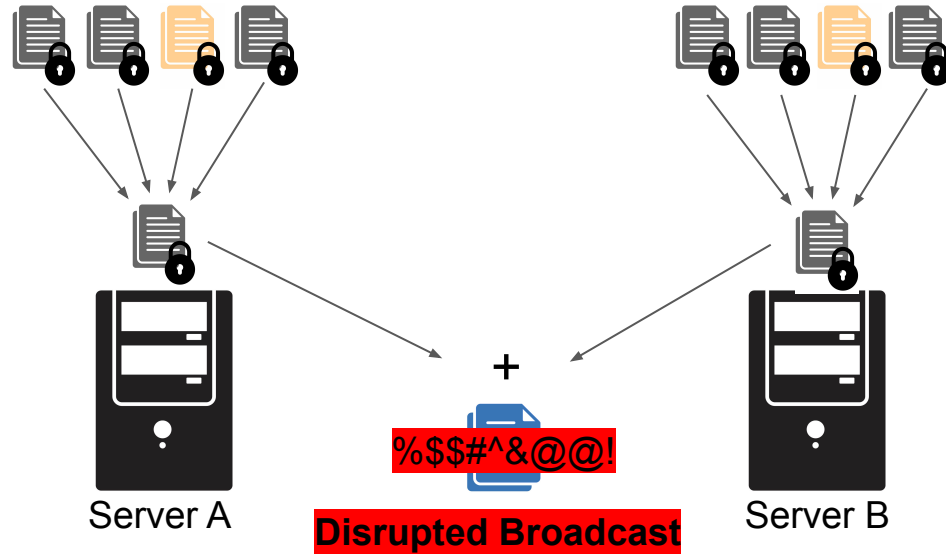
# Broadcast Disruption



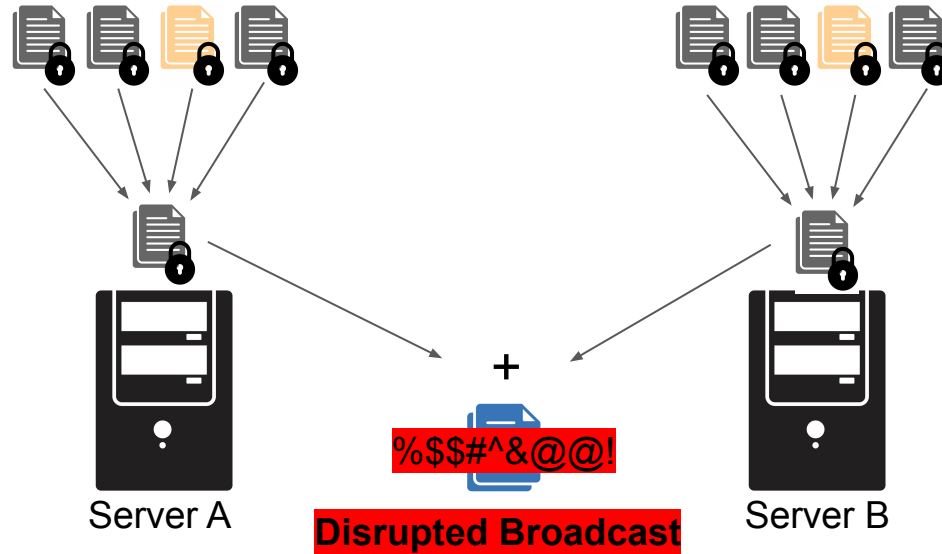
# Broadcast Disruption



# Broadcast Disruption



# Broadcast Disruption




Malicious users can corrupt broadcasts = **censorship**

# Definition: A Broadcast Channel

A **channel** is an allocated slot to which a message is written

- Multiple *simultaneous* broadcasts can be supported with multiple channels
  - A user can write to a channel without clobbering another message on another slot
- Servers can't tell which user is writing to which channel



# Existing Approaches for Preventing Disruption

## Riposte [CGBM15]

- Allocate  $\Omega(\# \text{ users })$  channel slots
- Clients pick **one random channel** to write to
  - Prevents collision with malicious writes w.h.p.
- Additional non-colluding “audit” server required to prevent malicious clients

## Blinder [APY20]

- Allocate  $\Omega(\# \text{ users })$  channel slots
- Clients pick **one random channel** to write to (like Riposte)
- Honest-majority MPC instead of audit server to prevent malicious clients

# Existing Approaches for Preventing Disruption

Common Theme: **# channels  $\gg$  # users**

⇒ more server work than DC-net

⇒ only efficient with small messages

sts our Riposte  
160-byte rows

and up to 1 million clients. Second, we show that Blinder scales to  
large messages of up to 10 kilobytes (required by some applications

Can we do better?



# Existing Approaches for Preventing Disruption

Common Theme: **# channels  $\gg$  # users**

$\Rightarrow$  more server work than DC-net

$\Rightarrow$  only efficient with small messages

sts our Riposte  
160-byte rows

and up to 1 million clients. Second, we show that Blinder scales to  
large messages of up to 10 kilobytes (required by some applications

Can we do better?

1. One channel per broadcaster?

# Existing Approaches for Preventing Disruption

Common Theme: **# channels  $\gg$  # users**

⇒ more server work than DC-net

⇒ only efficient with small messages

sts our Riposte  
160-byte rows

and up to 1 million clients. Second, we show that Blinder scales to  
large messages of up to 10 kilobytes (required by some applications

Can we do better?

1. One channel per broadcaster?
2. Optimize for few broadcasters, many users?

# Our approach

- One channel per broadcaster  $\Rightarrow$  need *authorized* writes
  - Otherwise, malicious clients can disrupt the broadcast

# Our approach

- One channel per broadcaster  $\Rightarrow$  need *authorized* writes
  - Otherwise, malicious clients can disrupt the broadcast
- Solution: **Blind** message authentication for access control
  - To write to a channel, must know its secret key

# Our approach


- One channel per broadcaster  $\Rightarrow$  need *authorized* writes
  - Otherwise, malicious clients can disrupt the broadcast
- Solution: **Blind** message authentication for access control
  - To write to a channel, must know its secret key
  - Servers do not know which client is the broadcaster

# Our approach

- One channel per broadcaster  $\Rightarrow$  need *authorized* writes
  - Otherwise, malicious clients can disrupt the broadcast
- Solution: **Blind** message authentication for access control
  - To write to a channel, must know its secret key
  - Servers do not know which client is the broadcaster
  - Server work:  $O(\# \text{ channels})$ , instead of  $\Omega(\# \text{ users})$


# Our approach

- One channel per broadcaster  $\Rightarrow$  need *authorized* writes
  - Otherwise, malicious clients can disrupt the broadcast
- Solution: **Blind** message authentication for access control
  - To write to a channel, must know its secret key
  - Servers do not know which client is the broadcaster
  - Server work:  $O(\# \text{ channels})$ , instead of  $\Omega(\# \text{ users})$



# Our approach

- One channel per broadcaster  $\Rightarrow$  need *authorized* writes
  - Otherwise, malicious clients can disrupt the broadcast
- Solution: **Blind** message authentication for access control
  - To write to a channel, must know its secret key
  - Servers do not know which client is the broadcaster
  - Server work:  $O(\# \text{ channels})$ , instead of  $\Omega(\# \text{ users})$






# Our approach

- One channel per broadcaster  $\Rightarrow$  need *authorized* writes
  - Otherwise, malicious clients can disrupt the broadcast
- Solution: **Blind** message authentication for access control
  - To write to a channel, must know its secret key
  - Servers do not know which client is the broadcaster
  - Server work:  $O(\# \text{ channels})$ , instead of  $\Omega(\# \text{ users})$

no auth for “dummy” requests





# Our approach

- One channel per broadcaster  $\Rightarrow$  need *authorized* writes
  - Otherwise, malicious clients can disrupt the broadcast
- Solution: **Blind** message authentication for access control
  - To write to a channel, must know its secret key
  - Servers do not know which client is the broadcaster
  - Server work:  $O(\# \text{ channels})$ , instead of  $\Omega(\# \text{ users})$

servers can't tell what key  
was used (if any)



Key	Channel
$a_1$	
$a_2$	

Protocol

# Outline

1. Supporting multiple channels
2. Blind message authentication
3. Preventing de-anonymization attacks

# Multiple channels (almost) for free

- Example scenario:
  - 5 broadcasters, each wants to share a 1MB message
  - If repeated in parallel then *every user* must upload 5 MB
  - Quickly saturates upload bandwidth

# Multiple channels (almost) for free


- Example scenario:
  - 5 broadcasters, each wants to share a 1MB message
  - If repeated in parallel then *every user* must upload 5 MB
  - Quickly saturates upload bandwidth
- Want: each user only uploads ~1MB!
  - Servers must still extract 5 “real” messages
  - From many random-looking messages

# Multiple channels (almost) for free

- Example scenario:
  - 5 broadcasters, each wants to share a 1MB message
  - If repeated in parallel then *every user* must upload 5 MB
  - Quickly saturates upload bandwidth
- Want: each user only uploads ~1MB!
  - Servers must still extract 5 “real” messages
  - From many random-looking messages
- Must write to *every* channel for *every* message
  - Otherwise the servers learn who is broadcasting

# Multiple channels (almost) for free

- Valid writes
  - *Point functions*







## Multiple channels (almost) for free



- Valid writes

- *Point functions*
- Or all-0 messages





# Multiple channels (almost) for free

- Valid writes
  - *Point functions*
  - Or all-0 messages
- Want to send *short secret-shares* of our write

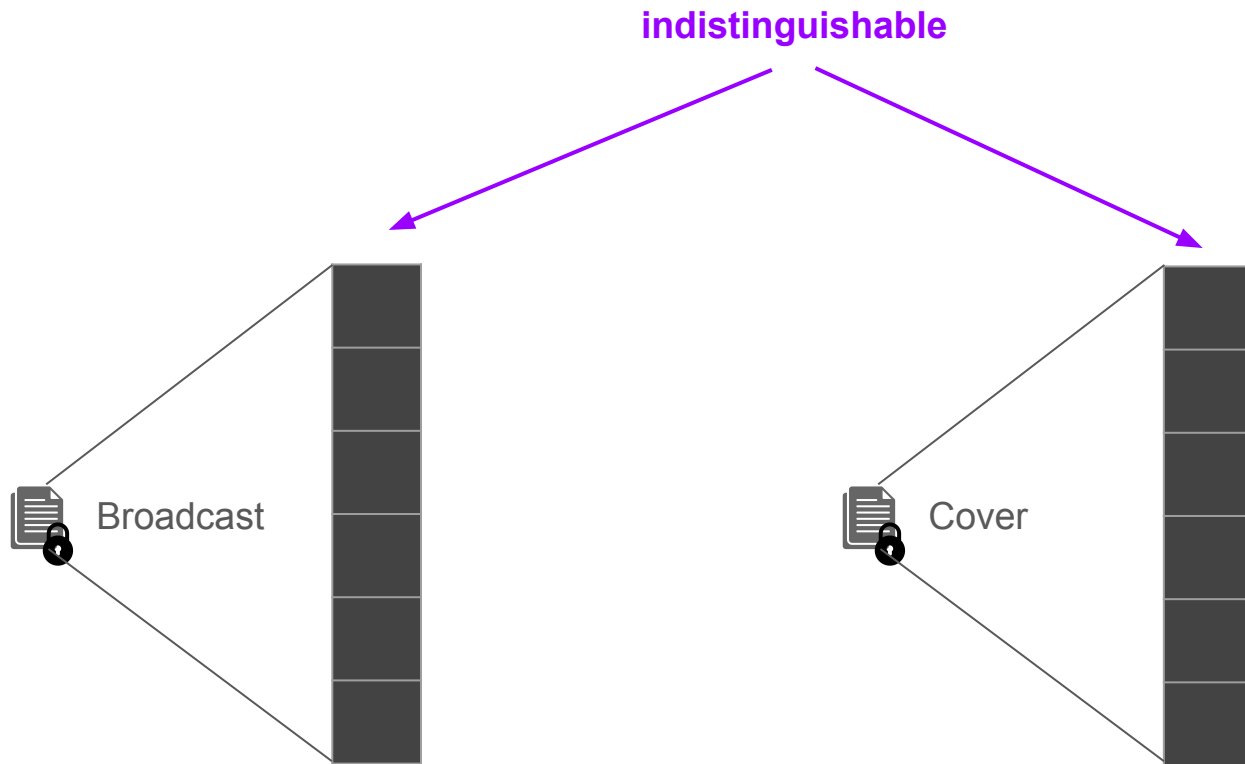


# Distributed point function (DPF) [GI14]

- Want: short representation
- Want: secret shares




# Distributed point function (DPF) [GI14]




# Distributed point function (DPF) [GI14]

Compress and secret-share a point function




# Distributed point function (DPF) [GI14]

Compress and secret-share a point function



# Distributed point function (DPF) [GI14]

Compress and secret-share a point function



# Blind message authentication

- For DPF, want to check (each channel): “is 0 **OR** authorized”
  - must perform check on secret shares
  - minimize interaction
  - must allow 0 messages, but they should look the same
  - malicious server shouldn't be able to impersonate malicious client




# Blind message authentication

- For DPF, want to check (each channel): “is 0 **OR** authorized”
- MAC the message [CW79,WC81]
  - can check over secret shares
  - verify that tagger knows a secret
  - *anybody* can tag empty messages

# Blind message authentication

- For DPF, want to check: “every channel is 0 **OR** client knows the secret”
- MAC the message [CW79,WC81]

MAC tag


$$t = a \cdot m$$

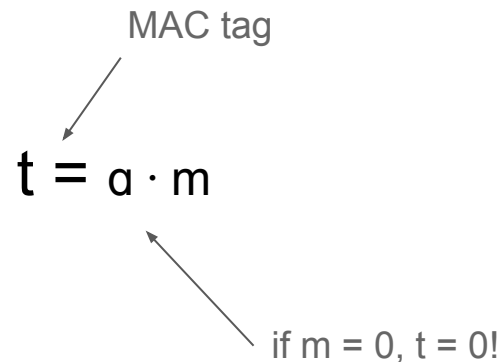
# Blind message authentication

- For DPF, want to check: “every channel is 0 **OR** client knows the secret”
- MAC the message [CW79,WC81]

MAC tag

$$t = a \cdot m$$

if  $m = 0$ ,  $t = 0$ !



# Blind message authentication

- For DPF, want to check: “every channel is 0 **OR** client knows the secret”
- MAC the message [CW79,WC81]
- Can compute *locally* on additive secret-shares

The diagram illustrates the local computation of a MAC tag. On the left, the text "known to servers (for now)" has two arrows pointing to the terms  $a \cdot \llbracket m \rrbracket_1$  and  $a \cdot \llbracket m \rrbracket_2$ . These two terms are separated by a plus sign  $+$ . To the right of the plus sign is the equation  $= t = a \cdot m$ . An arrow points from the text "MAC tag" to the variable  $t$  in this equation.

$$\begin{array}{l} \text{known to servers} \\ \text{(for now)} \end{array} \begin{array}{l} \nearrow a \cdot \llbracket m \rrbracket_1 \\ \searrow a \cdot \llbracket m \rrbracket_2 \end{array} + = t = a \cdot m$$

MAC tag

# Blind message authentication

- For DPF, want to check: “every channel is 0 **OR** client knows the secret”
- MAC the message [CW79,WC81]
- Can compute *locally* on additive secret-shares

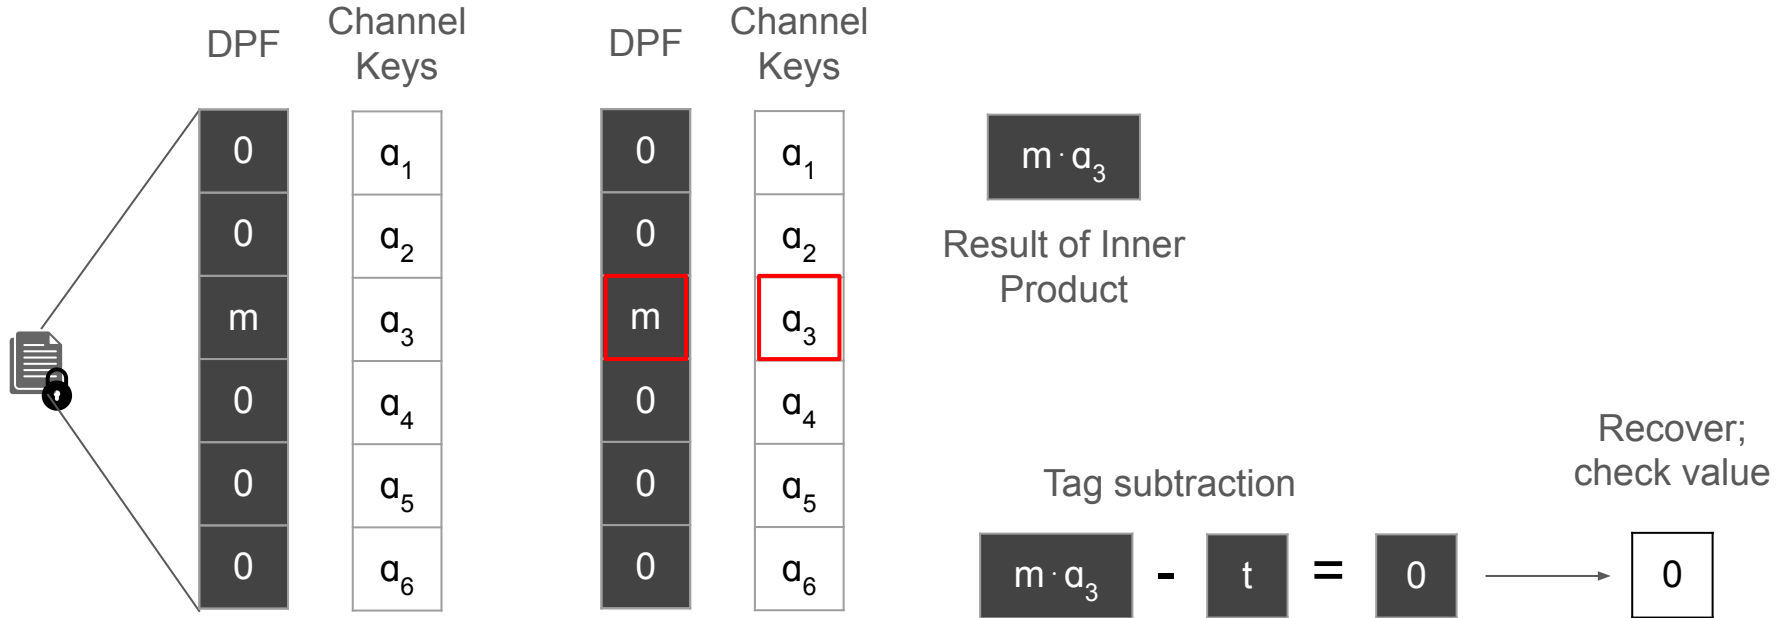
$$\begin{array}{ccccc} & & & & \text{secret shares of MAC tag} \\ & & & & \swarrow \\ & & a \cdot [m]_1 - [t]_1 & & \\ \nearrow & & + & & \searrow \\ \text{known to servers} & & & & \\ \text{(for now)} & & & & \\ \searrow & & a \cdot [m]_2 - [t]_2 & & \\ & & & & \end{array} = 0$$

# Blind message authentication

- For DPF, want to check: “every channel is 0 **OR** client knows the secret”
- MAC the message [CW79,WC81]
- Can compute *locally* on additive secret-shares
- Multiple channels: inner product


$$\begin{aligned} & \langle (a_1, \dots, a_L), ([m_1]_1, \dots, [m_L]_1) \rangle - [t]_1 \\ & \quad + \\ & \langle (a_1, \dots, a_L), ([m_1]_2, \dots, [m_L]_2) \rangle - [t]_2 \end{aligned} = 0$$

# Blind message authentication



# Problem: Client-server collusion

$(a_1, \dots, a_L)$




Server A

$(a_1, \dots, a_L)$



Server B

$a$





# Blind message authentication

- For DPF, want to check: “every channel is 0 **OR** client knows the secret”
- MAC the message [CW79,WC81]
- Can compute *locally* on additive secret-shares
- Multiple channels: inner product

$$\begin{aligned} & \langle (a_1, \dots, a_L), ([m_1]_1, \dots, [m_L]_1) \rangle - [t]_1 \\ & \quad + \\ & \langle (a_1, \dots, a_L), ([m_1]_2, \dots, [m_L]_2) \rangle - [t]_2 \\ & \quad = 0 \end{aligned}$$


# Blind message authentication

- For DPF, want to check: “every channel is 0 **OR** client knows the secret”
- MAC the message [CW79,WC81]
- Can compute *locally* on additive secret-shares
- Multiple channels: inner product
- In the exponent: hide keys from servers!

$$\begin{aligned}
 &g^{\langle (a_1, \dots, a_L), ([m_1]_1, \dots, [m_L]_1) \rangle - [t]_1} \\
 &\quad \cdot \\
 &g^{\langle (a_1, \dots, a_L), ([m_1]_2, \dots, [m_L]_2) \rangle - [t]_2} = g^0
 \end{aligned}$$

# Client-server collusion

$(g^{a_1}, \dots, g^{a_L})$




Server A

$(g^{a_1}, \dots, g^{a_L})$



Server B




# Deanonymization Attack


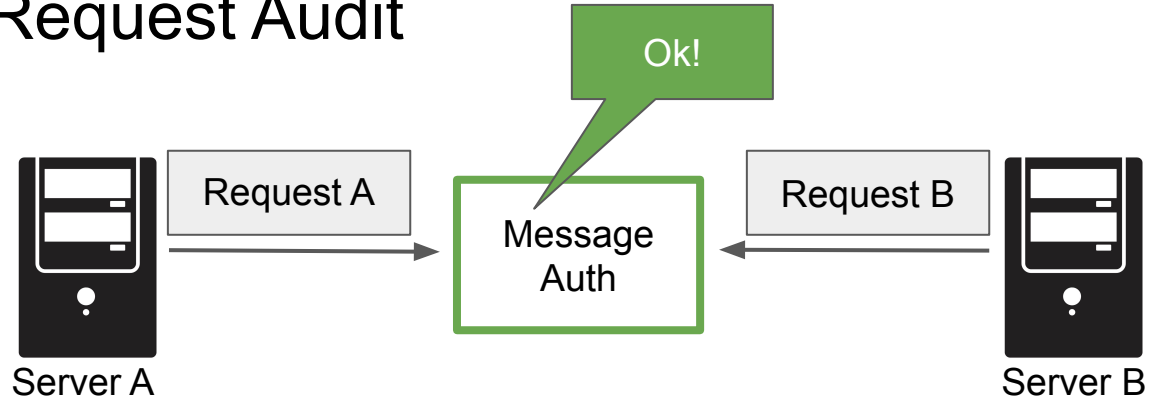
# Availability vs. Privacy

- In anonymous broadcast Availability  $\Rightarrow$  Privacy
  - Denying a client removes them from the anonymity set
- Prior work: disruptive clients ignored
  - Can't have availability and privacy at the same time!
  - Makes Riposte / Spectrum susceptible to an “audit attack”
  - A malicious server can artificially shrink the anonymity set

# Recap: Request Audit



# Recap: Request Audit




# Audit Attack

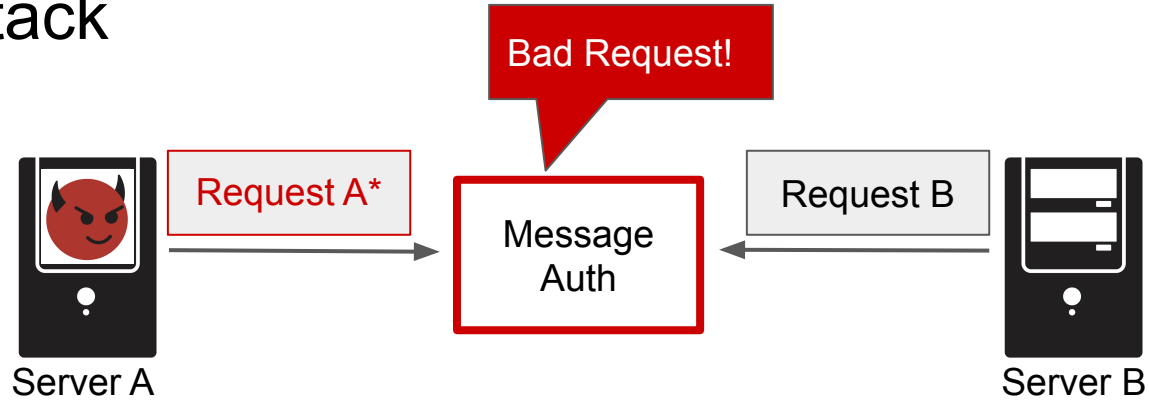




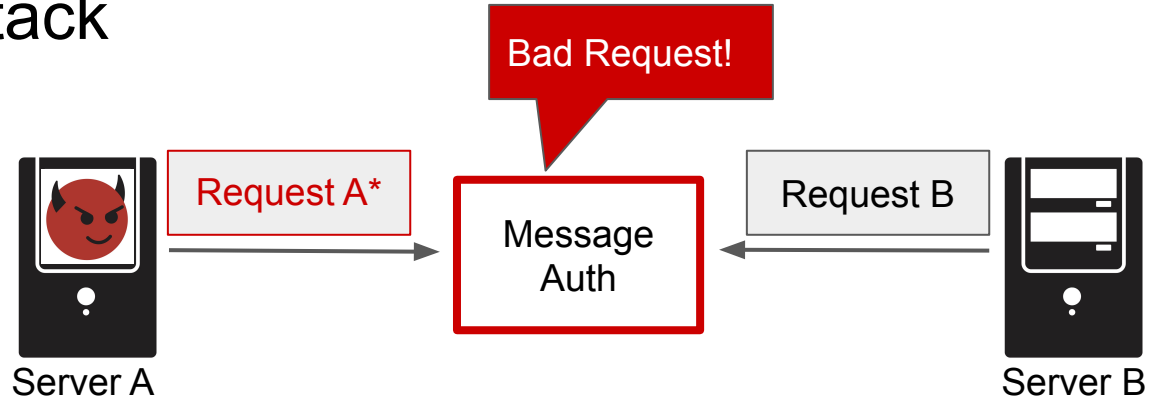
# Audit Attack




# Audit Attack




# Audit Attack



# Audit Attack



Server A



0


Broadcast




Server B




# Audit Attack



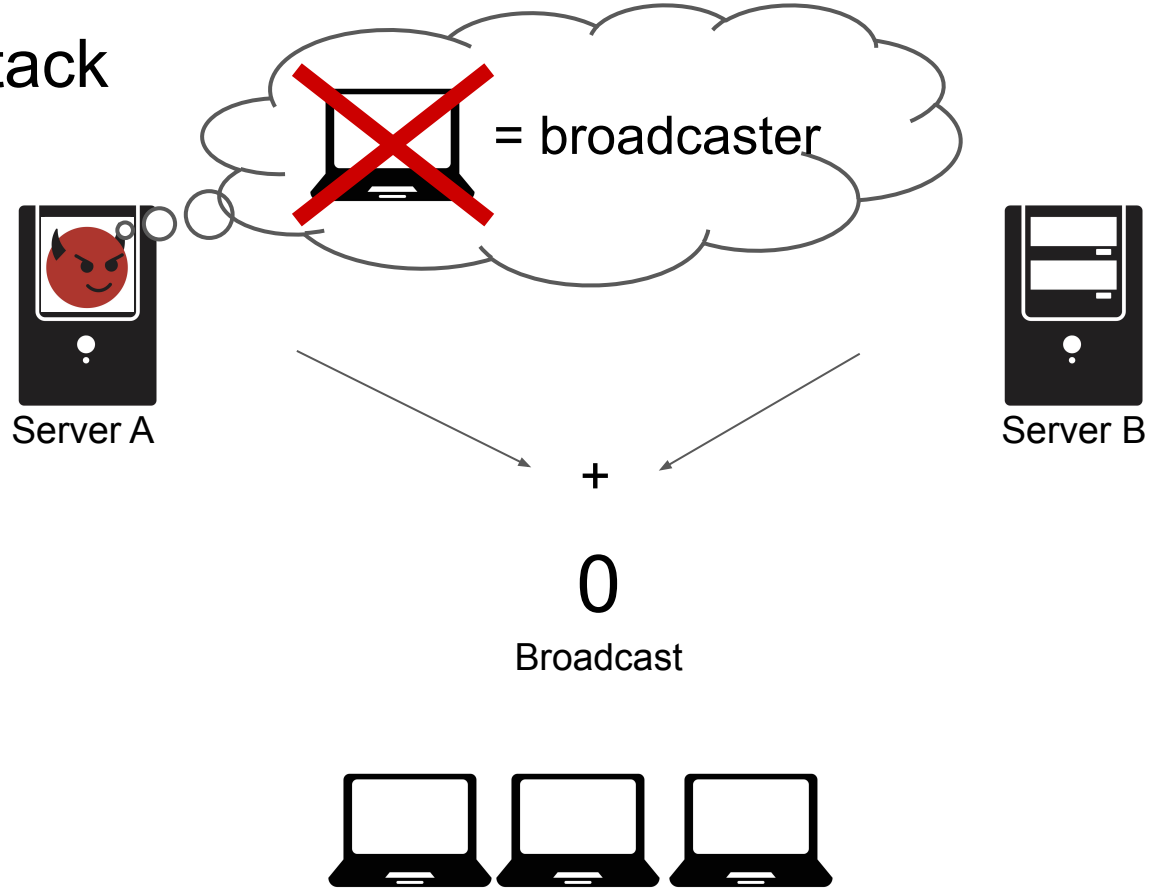
Server A



Server B



# Audit Attack



# Why this is challenging

Can't trust any client

- client could be attempting disruption

Can't trust any server

- “other” server could have caused the failure

# BlameGame: Stopping the Audit Attack

Who's cheating? Is it the client? Is it the server?

**Want: abort Spectrum if a server is malicious.**



# BlameGame: Stopping the Audit Attack

**Key idea:** “commit” each server to its (private) audit input


# BlameGame: Stopping the Audit Attack

**Key idea:** “commit” each server to its (private) audit input

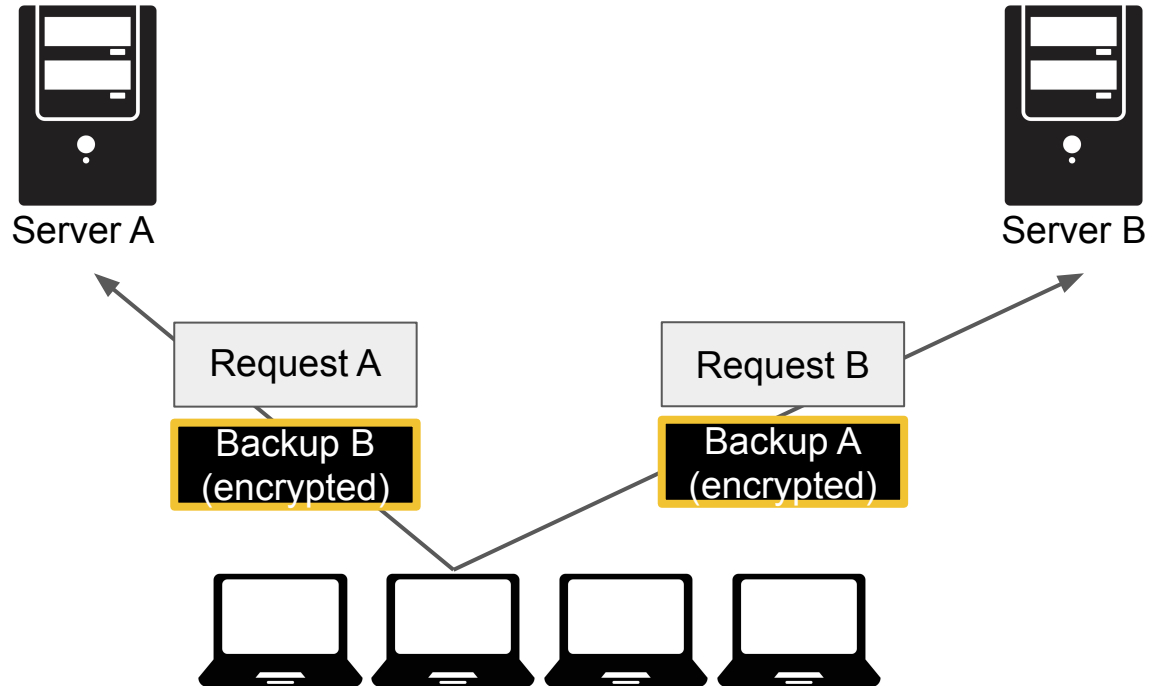
In a nutshell:

1. Clients send “backup” requests encrypted under other server’s public key
2. **If audit fails:** swap backups, decrypt, and try again
3. **If backup fails:** servers prove compliance
4. Blames malicious *server* or *client*


# Step 1: Do regular request audit



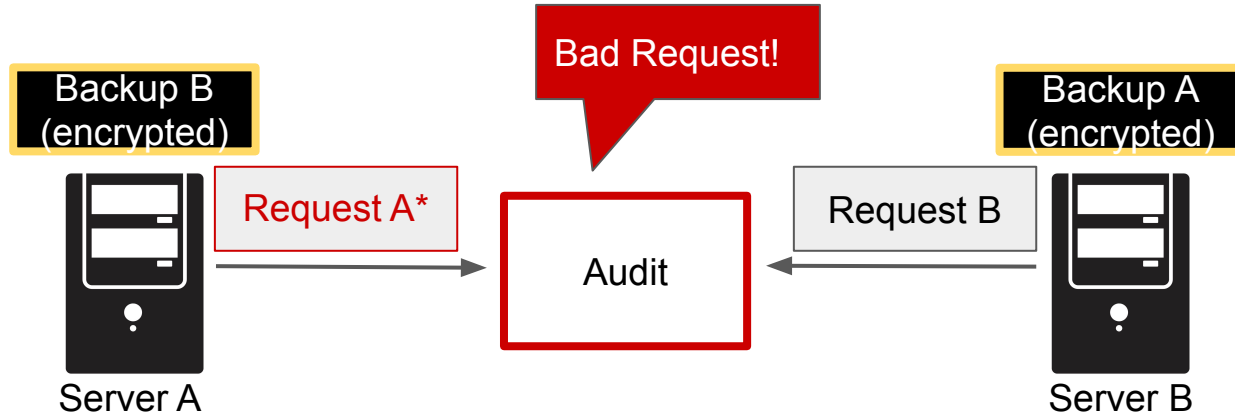
# Step 1: Do regular request audit



# Step 1: Do regular request audit



# Step 1: Do regular request audit



## Step 2: Do backup audit

Backup B  
(encrypted)



Server A


Backup A  
(encrypted)



Server B



## Step 2: Do backup audit





## Step 2: Do backup audit

Backup A  
(encrypted)



Server A

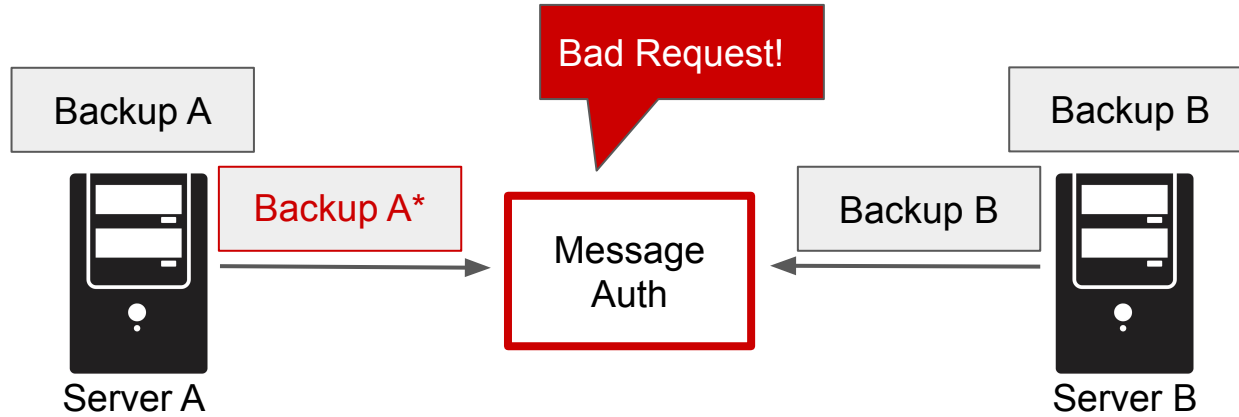
Backup B  
(encrypted)




Server B



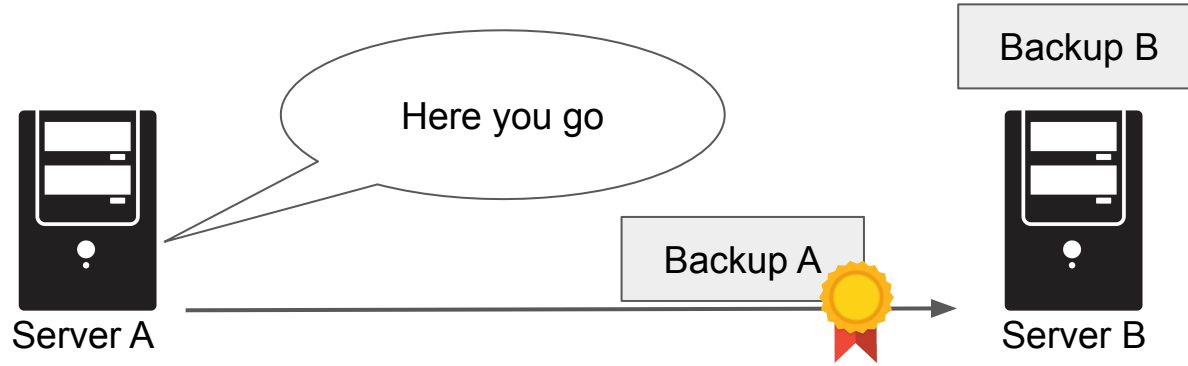
## Step 2: Do backup audit



## Step 3: Prove innocence




## Step 3: Prove innocence




## Step 4: Assign blame



Server A

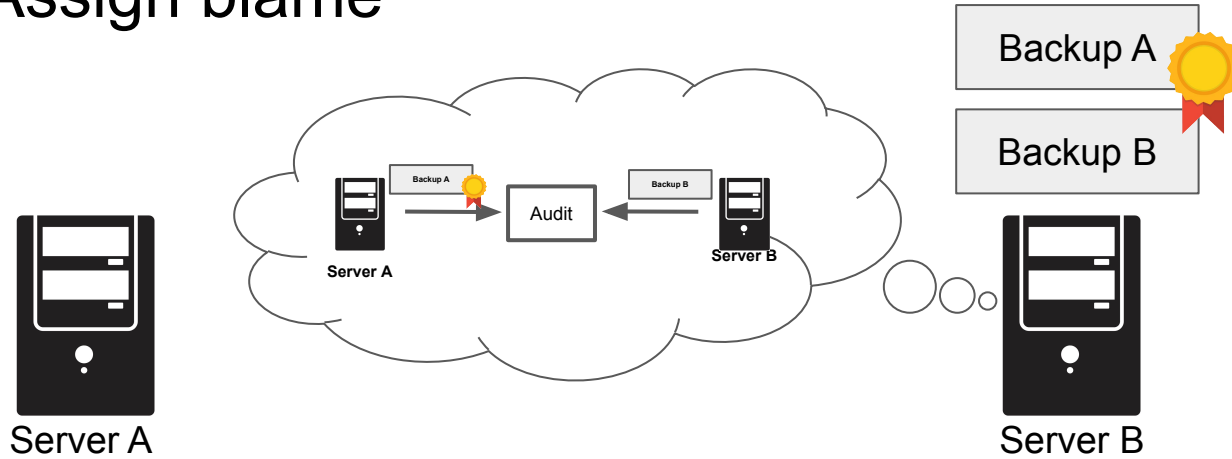


Server B

1. **Check Proof** 
2. Simulate Audit
3. Blame: *Client* or *Server*



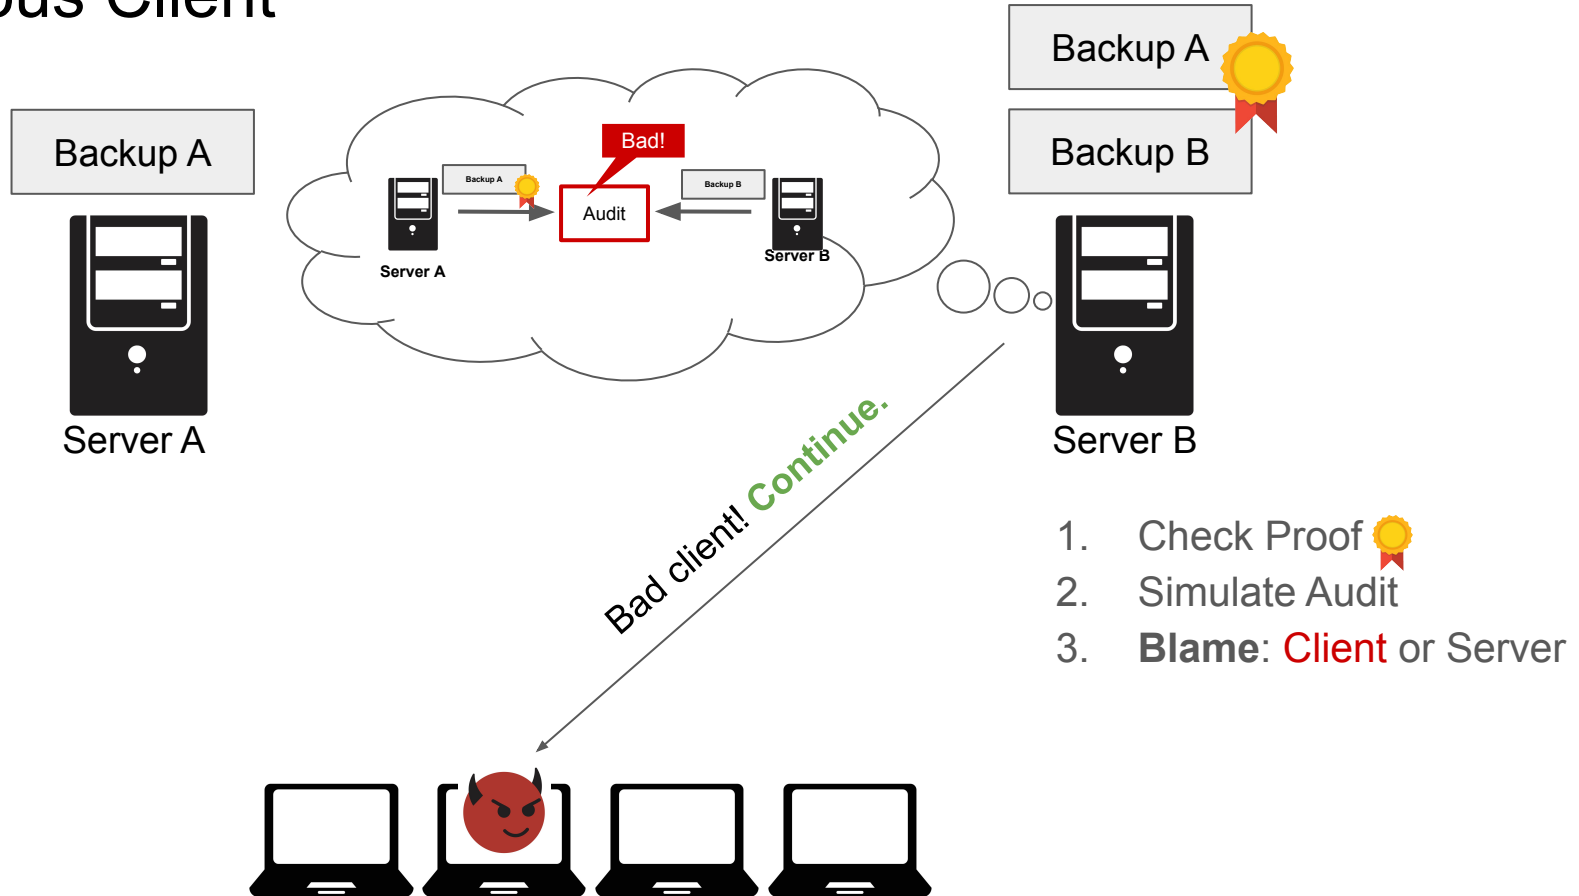
## Step 4: Assign blame




1. Check Proof 🏆
2. **Simulate Audit**
3. Blame: *Client* or *Server*



# Malicious Client



# Malicious Server



1. Check Proof 🏆
2. Simulate Audit
3. **Blame:** Client or **Server**





# BlameGame in practice

- Now only one shot at de-anonymization before Spectrum is aborted
- Only invoked following a failed audit of a request (few in practice)
- Backup request size independent of message length
- Also applies to other anonymous broadcast protocols


# BlameGame in practice

- Now only one shot at de-anonymization before Spectrum is aborted
- Only invoked following a failed audit of a request (few in practice)
- **Backup request size independent of message length**
- Also applies to other anonymous broadcast protocols

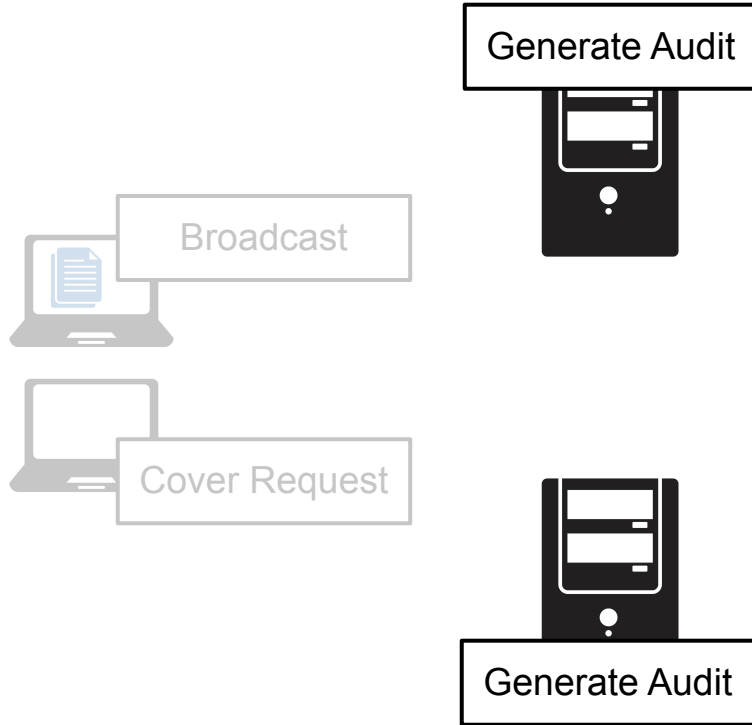
BlameGame (per failed audit)	Backup Request per client	Audit per client	Decryption once per client
	140 bytes	200 bytes	10 $\mu$ s

Spectrum

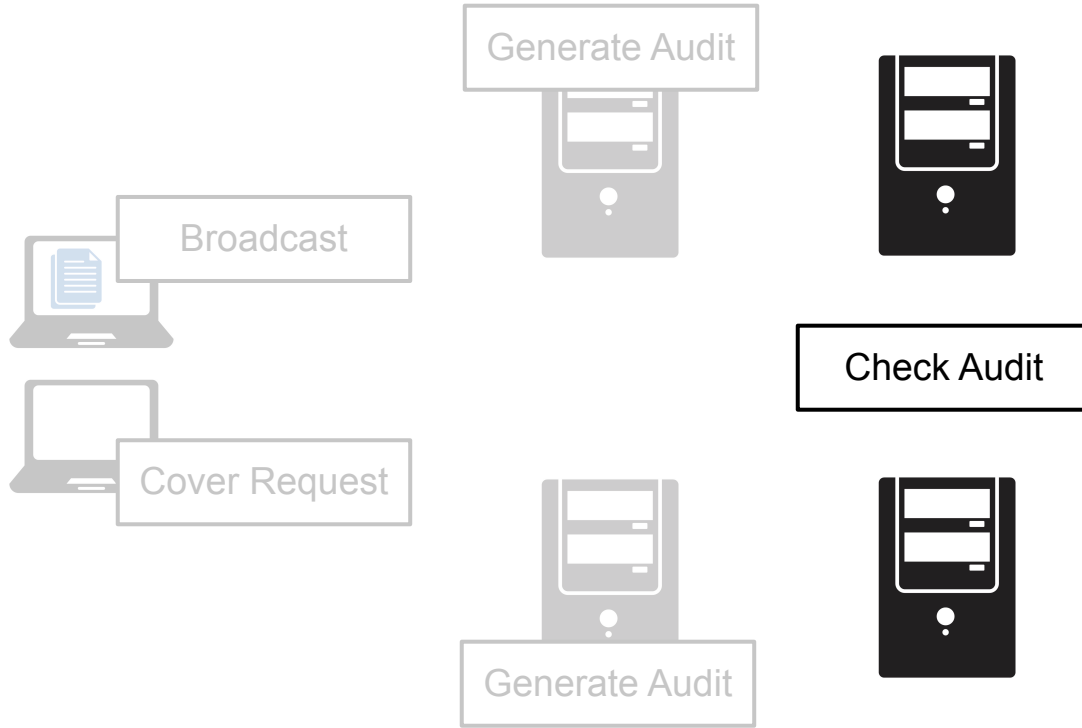
# Putting it all together



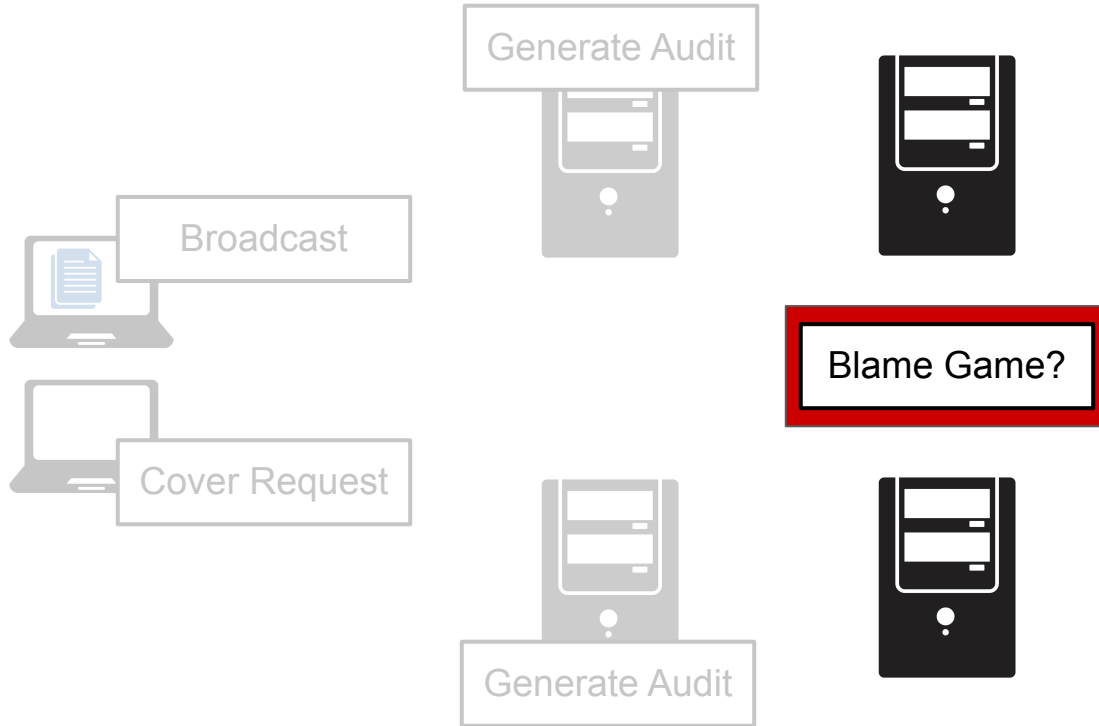
# Putting it all together



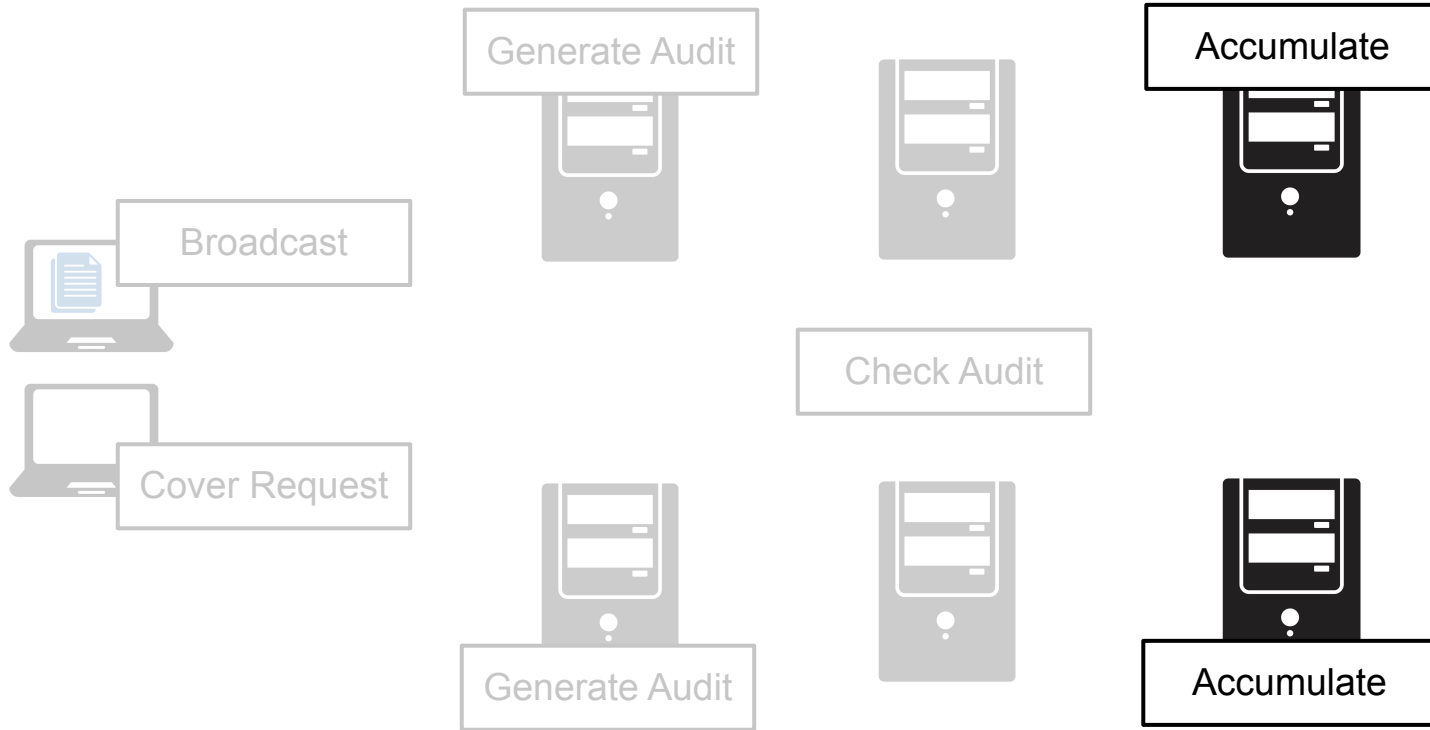
# Putting it all together



# Putting it all together

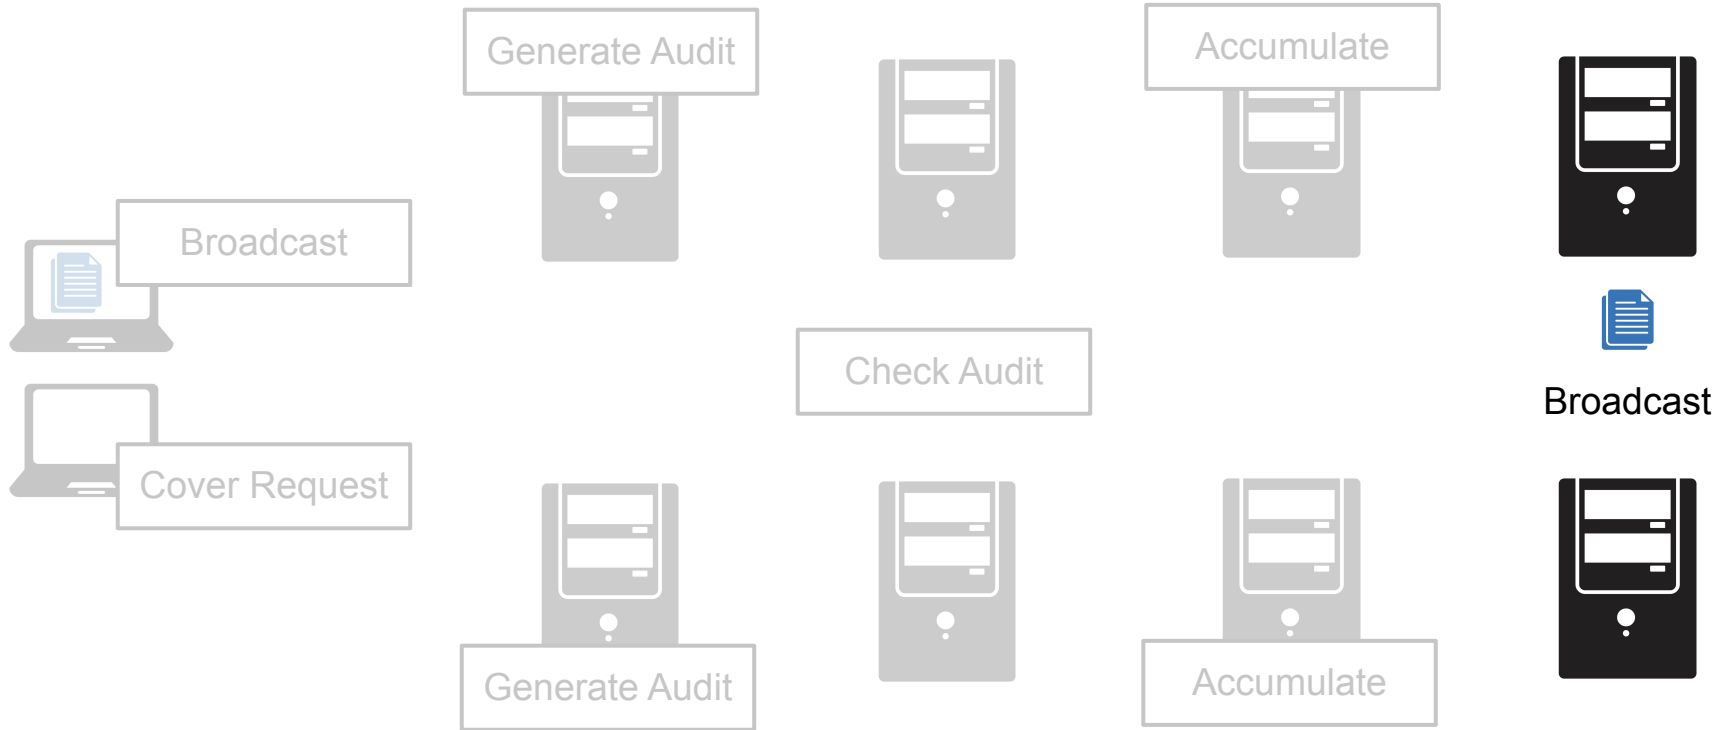


# Putting it all together





# Putting it all together



# Evaluation

# Implementation

- GitHub: [znewman01/spectrum-impl](https://github.com/znewman01/spectrum-impl) (written in Rust)
  - Compare with prior work: Riposte, Express, Blinder
- Terraform templates and scripts for reproducing experiments
  - Including prior work

## README.md

### spectrum-impl

build passing

Implementation and experiments for the [Spectrum paper](#).

**Disclaimer:** research code, not for production use.


```
$ python -m experiments spectrum experiments.json
• [infrastructure] Environment(instance_type='c5.4xlarge', client_machines=1, worker_machines=2)
i [infrastructure] no changes to apply
✓ [infrastructure] connected (SSH)
✓ [infrastructure] etcd healthy

• Experiment(clients=50, channels=3, message_size=10240, instance_type='c5.4xlarge', clients_per_machine=50, workers_per_machine=1, worker_machines_per_group=1, protocol=Symmetric(security=16))
✓ [experiment] 5375 queries in 1232ms => 4362 qps

• [infrastructure] Environment(instance_type='c5.4xlarge', client_machines=2, worker_machines=2)
✓ [infrastructure] found changes to apply:
  • aws_instance.client[1] will be created
✓ [infrastructure] created
✓ [infrastructure] connected (SSH)
✓ [infrastructure] etcd healthy


• Experiment(clients=100, channels=3, message_size=10240, instance_type='c5.4xlarge', clients_per_machine=50, workers_per_machine=1, worker_machines_per_group=1, protocol=Symmetric(security=16))
■ [experiment] running
```

# Spectrum: really fast with one channel & BIG messages

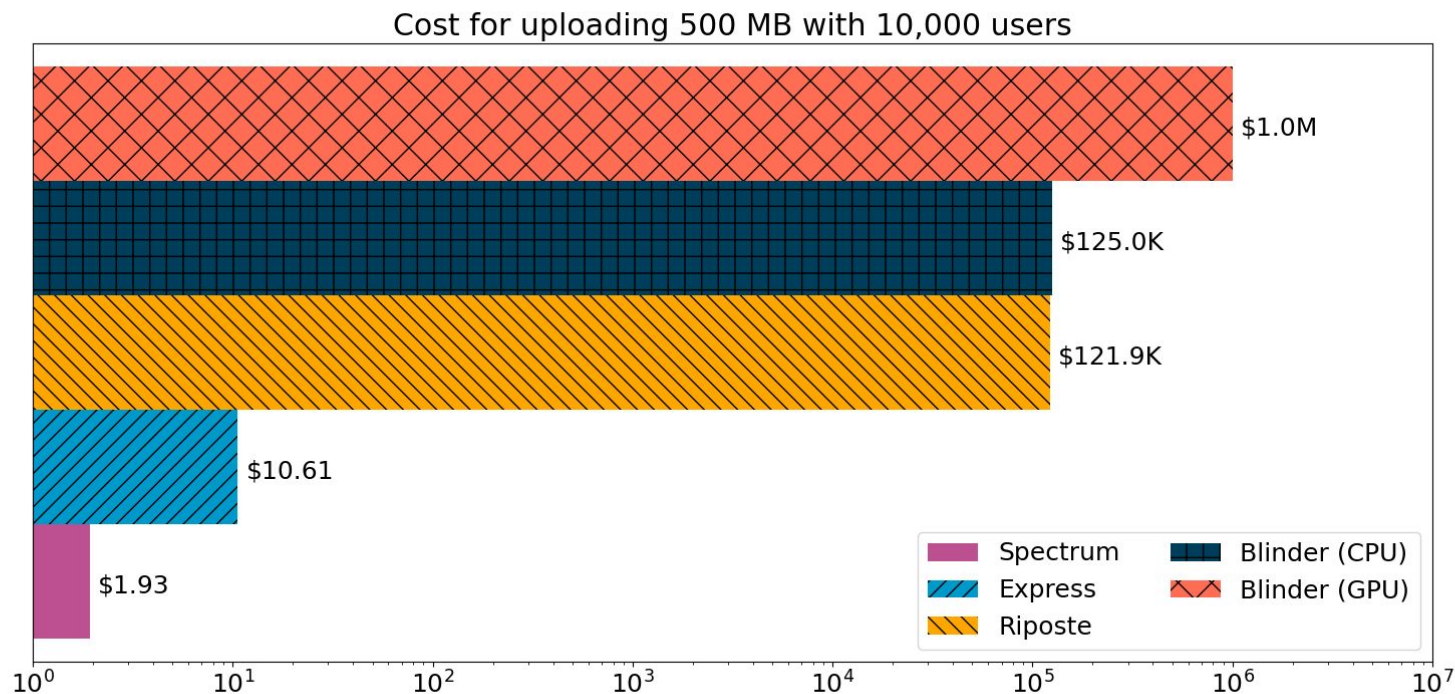


# Spectrum: still fast with many channels (10K users)


(higher is better)



# Cost to upload a full-length documentary



# Can parallelize each server



To share (anonymity set 10K)

Document	Size (MB)	Time	Cost (\$ USD)
PDF (e.g., ePrint for Spectrum)	1	10 s	\$ 0.01
Podcast (1 hr.)	50	8m 30s	\$ 0.19
Documentary (2 hr. @ 720p)	500	1h 24m	\$ 1.93



# Thank You

NSDI'22 (to appear)  
ePrint: [ia.cr/2021/325](https://ia.cr/2021/325)

Zack Newman: [zjn@mit.edu](mailto:zjn@mit.edu)  
Sacha Servan-Schreiber: [3s@mit.edu](mailto:3s@mit.edu)