

# How Bluesky Survived

Jaz

(They/Them)

Bsky Infra

@jaz.bsky.social

Who am I?



Who am I?




Who am I?






Who am I?








Who am I?





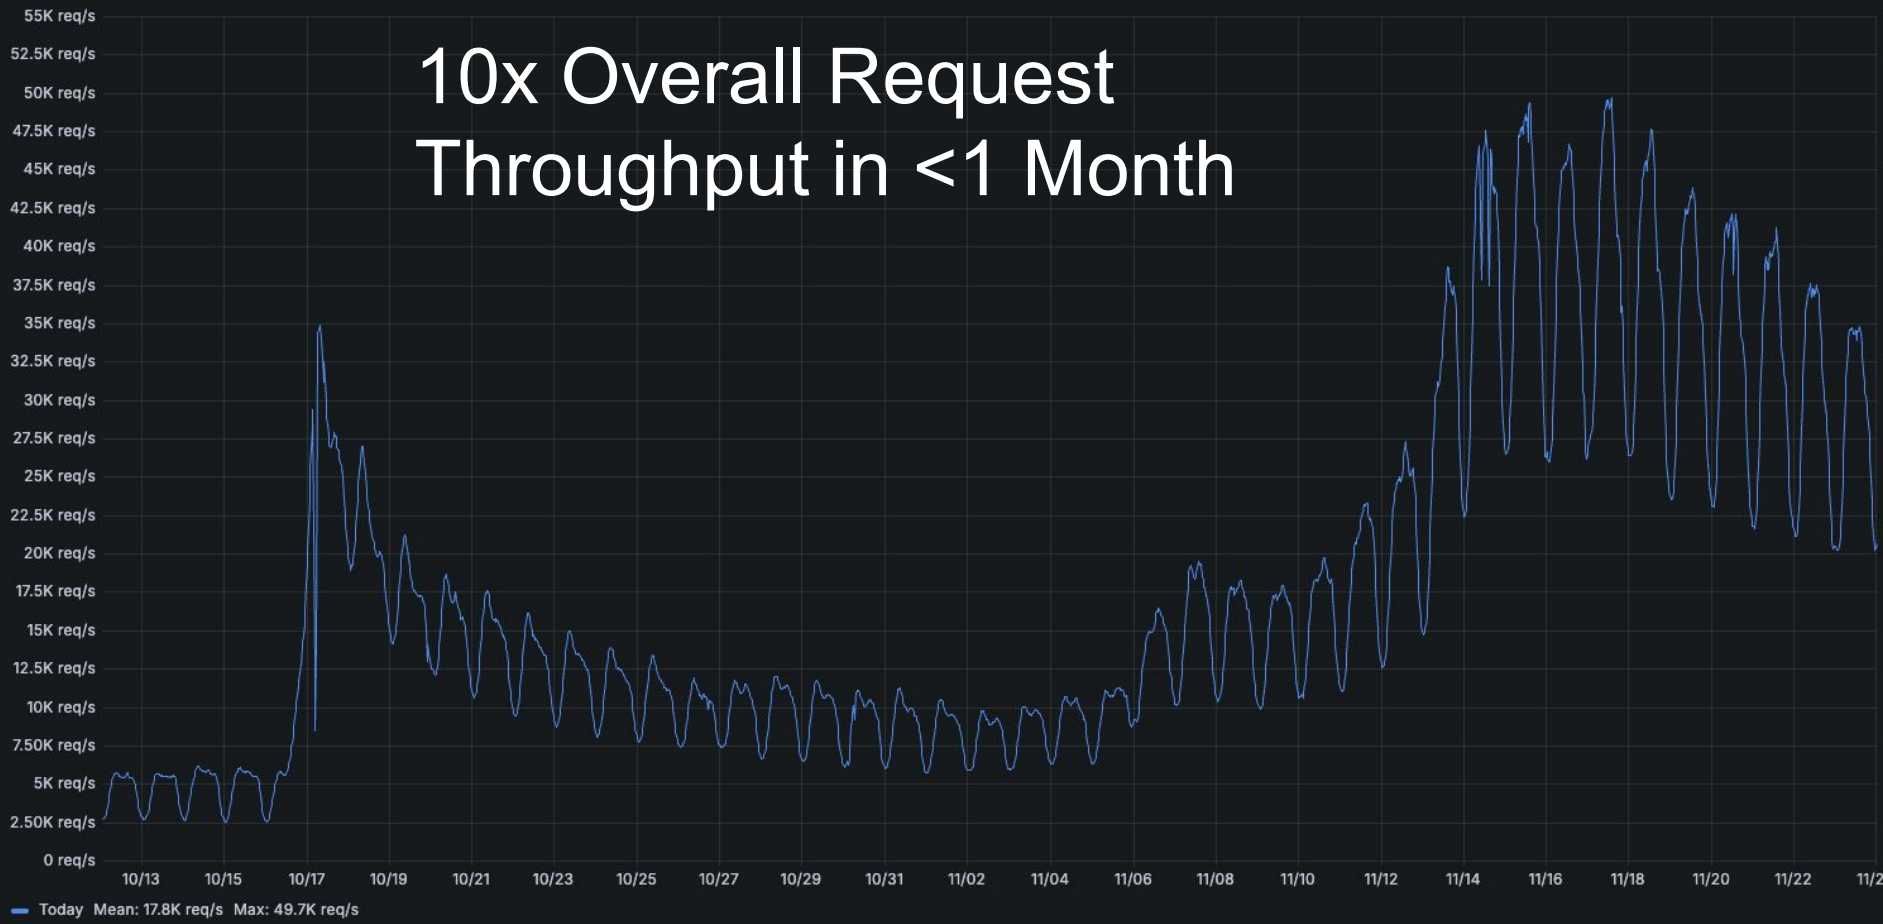


Who am I?



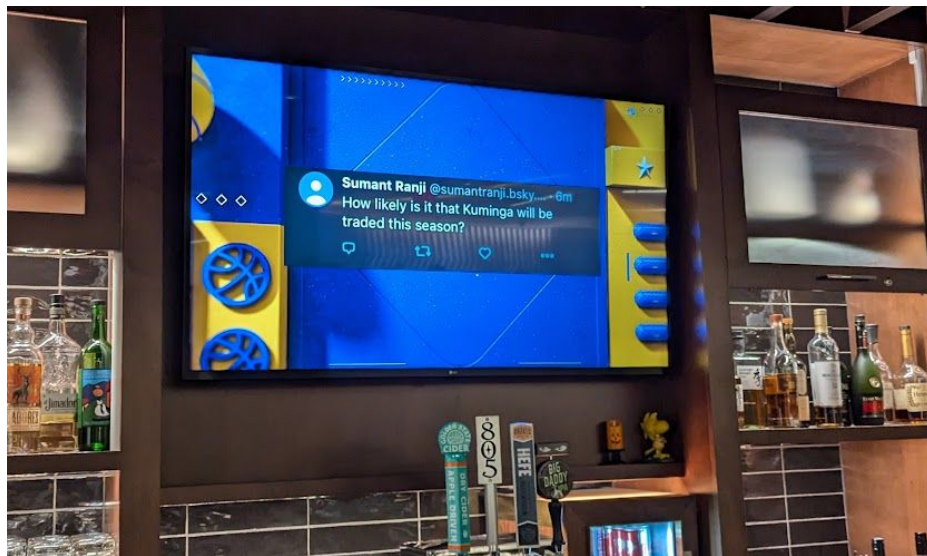
Request Throughput

10x Overall Request  
Throughput in <1 Month

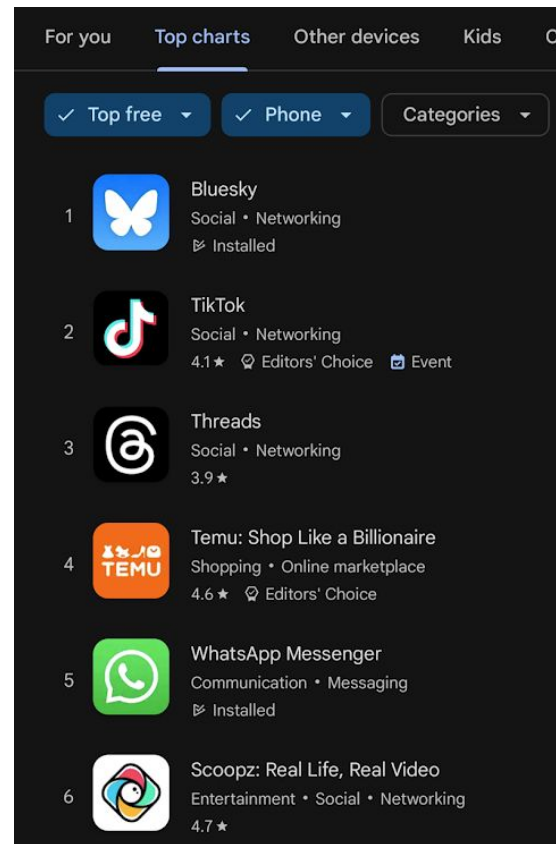




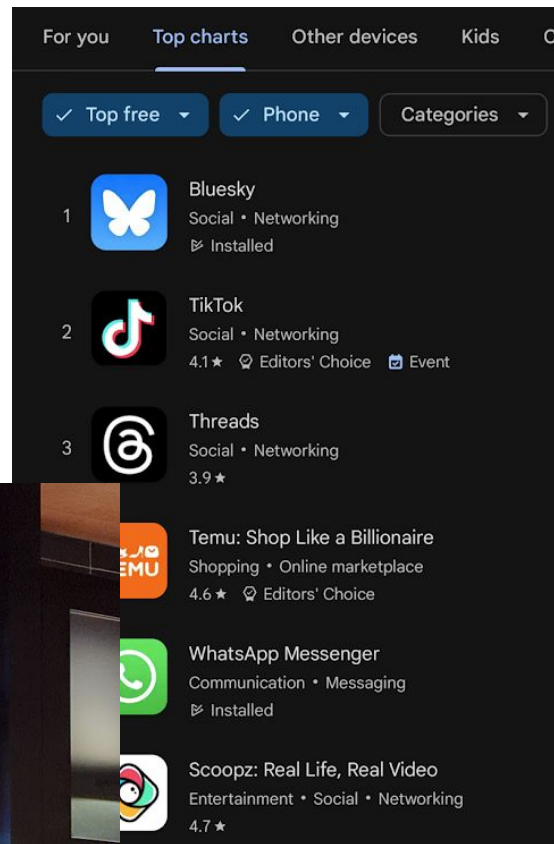
## Bluesky Takes Off



# Bluesky Takes Off



# Bluesky Takes Off



# Mycelial Networks And Bluer Skies

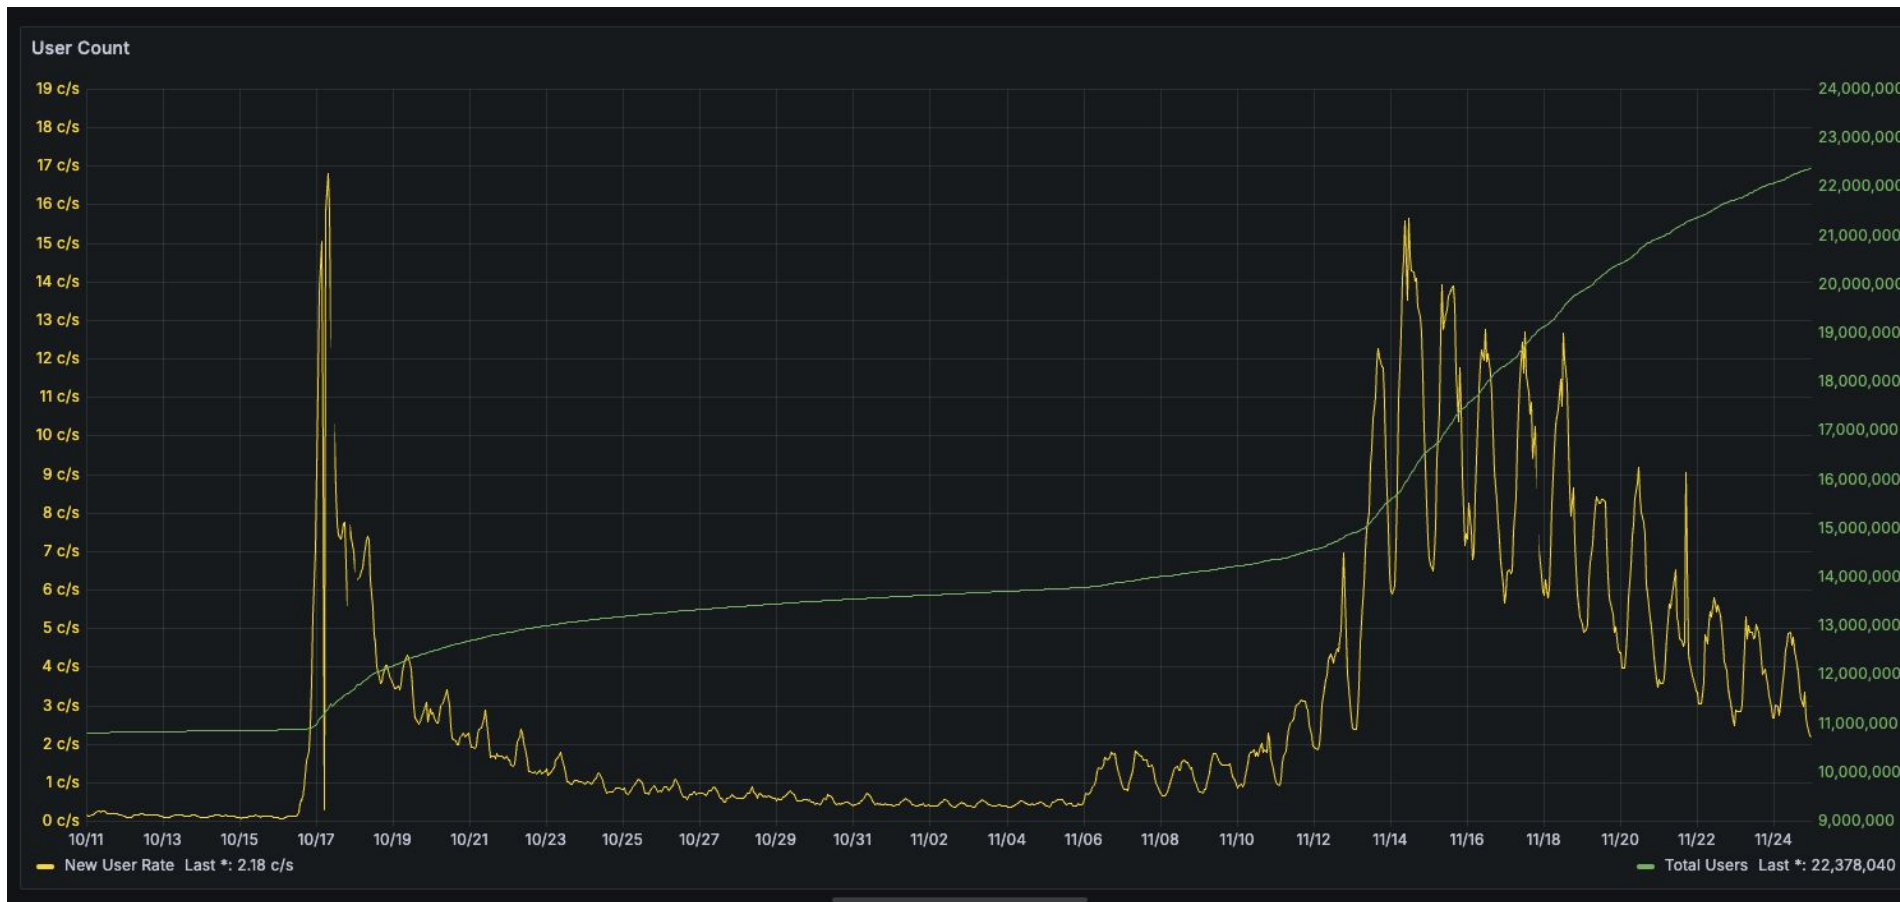




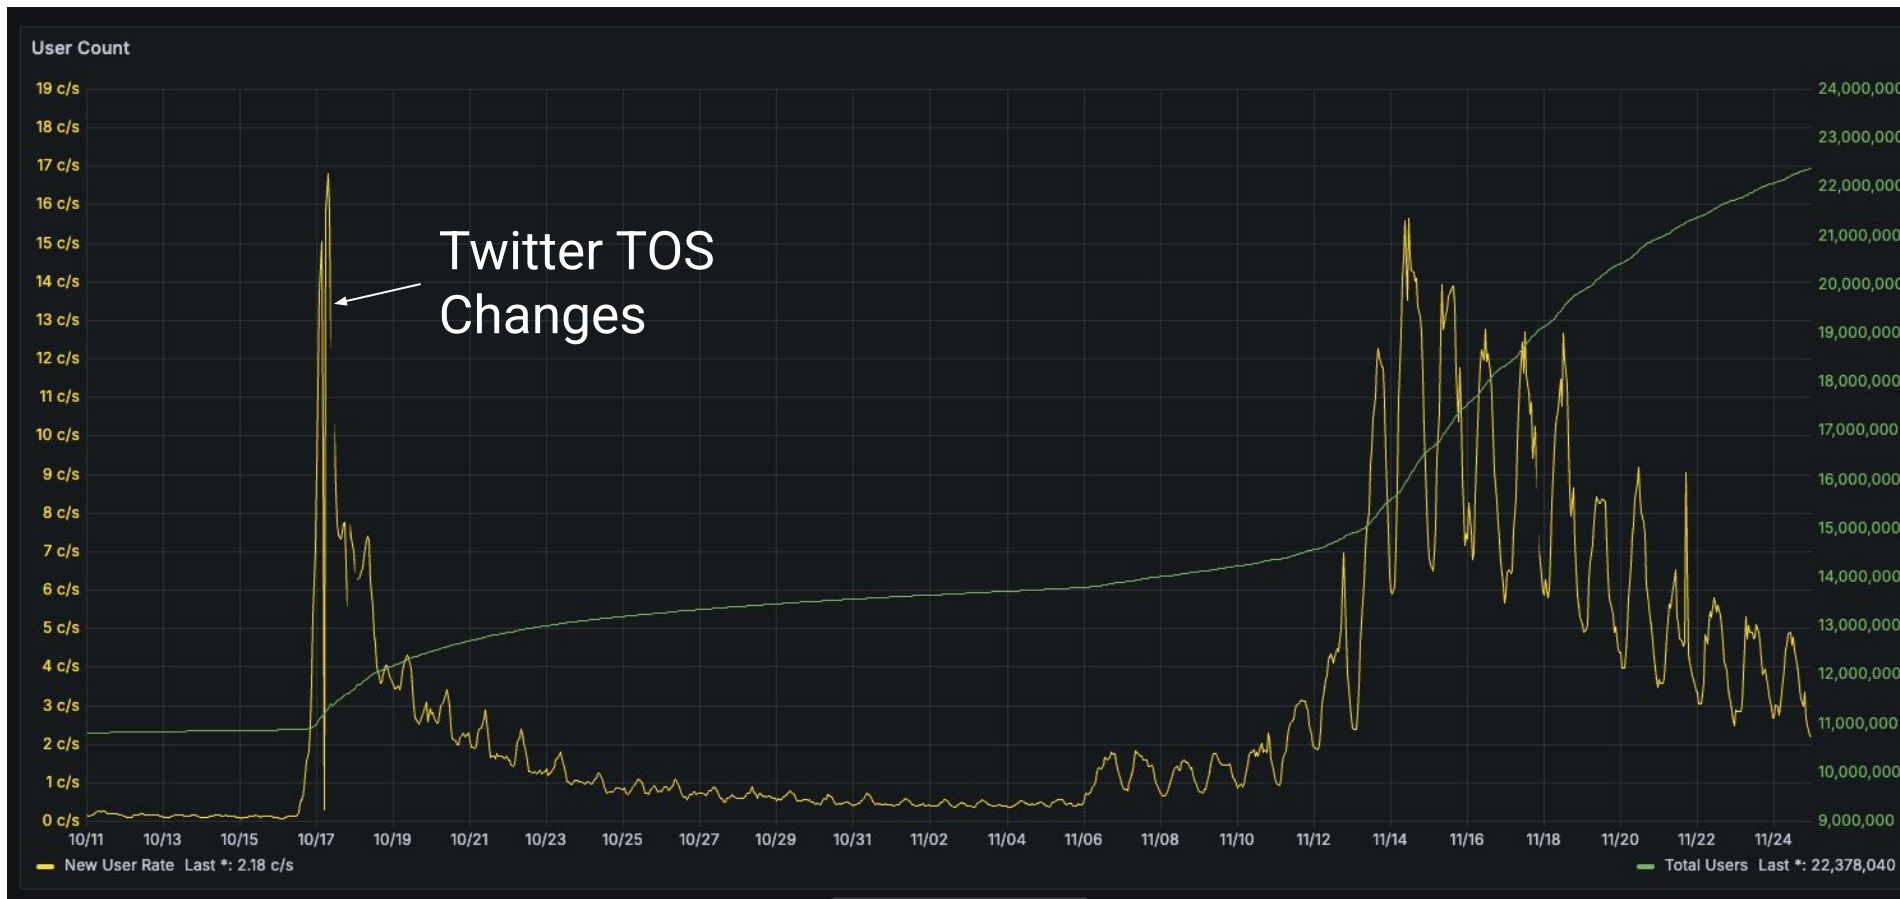
Ready for Action




# User growth



# User growth

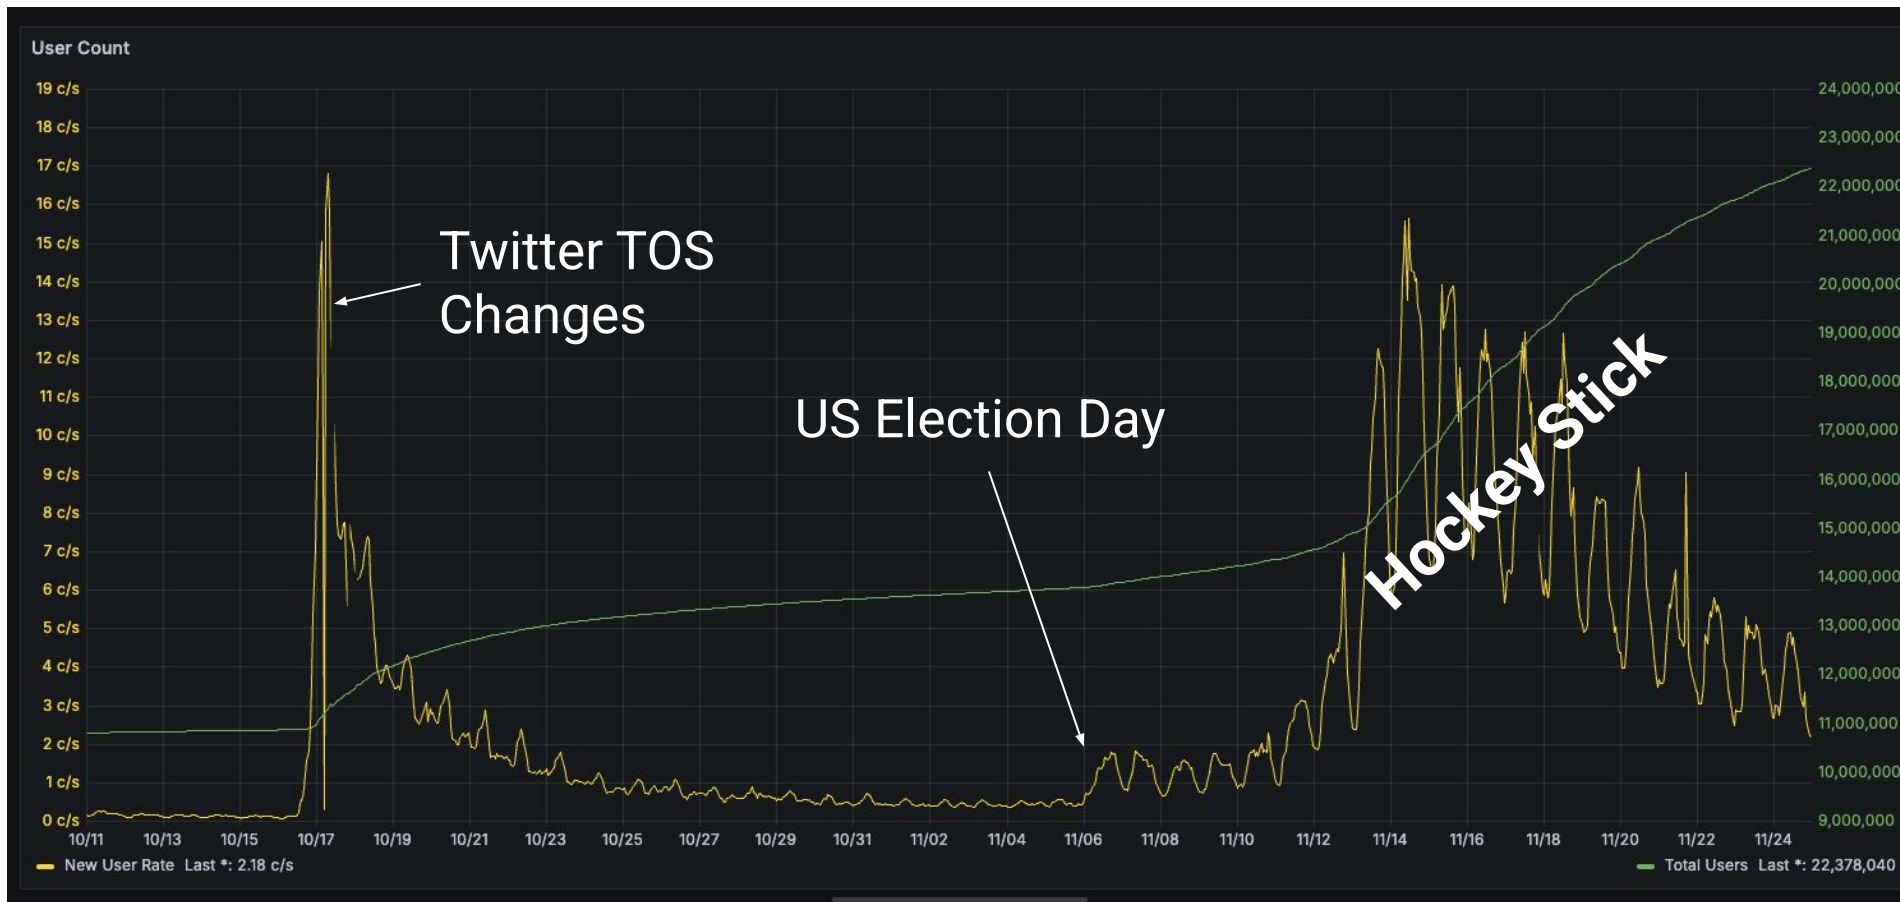


# User growth






# User growth

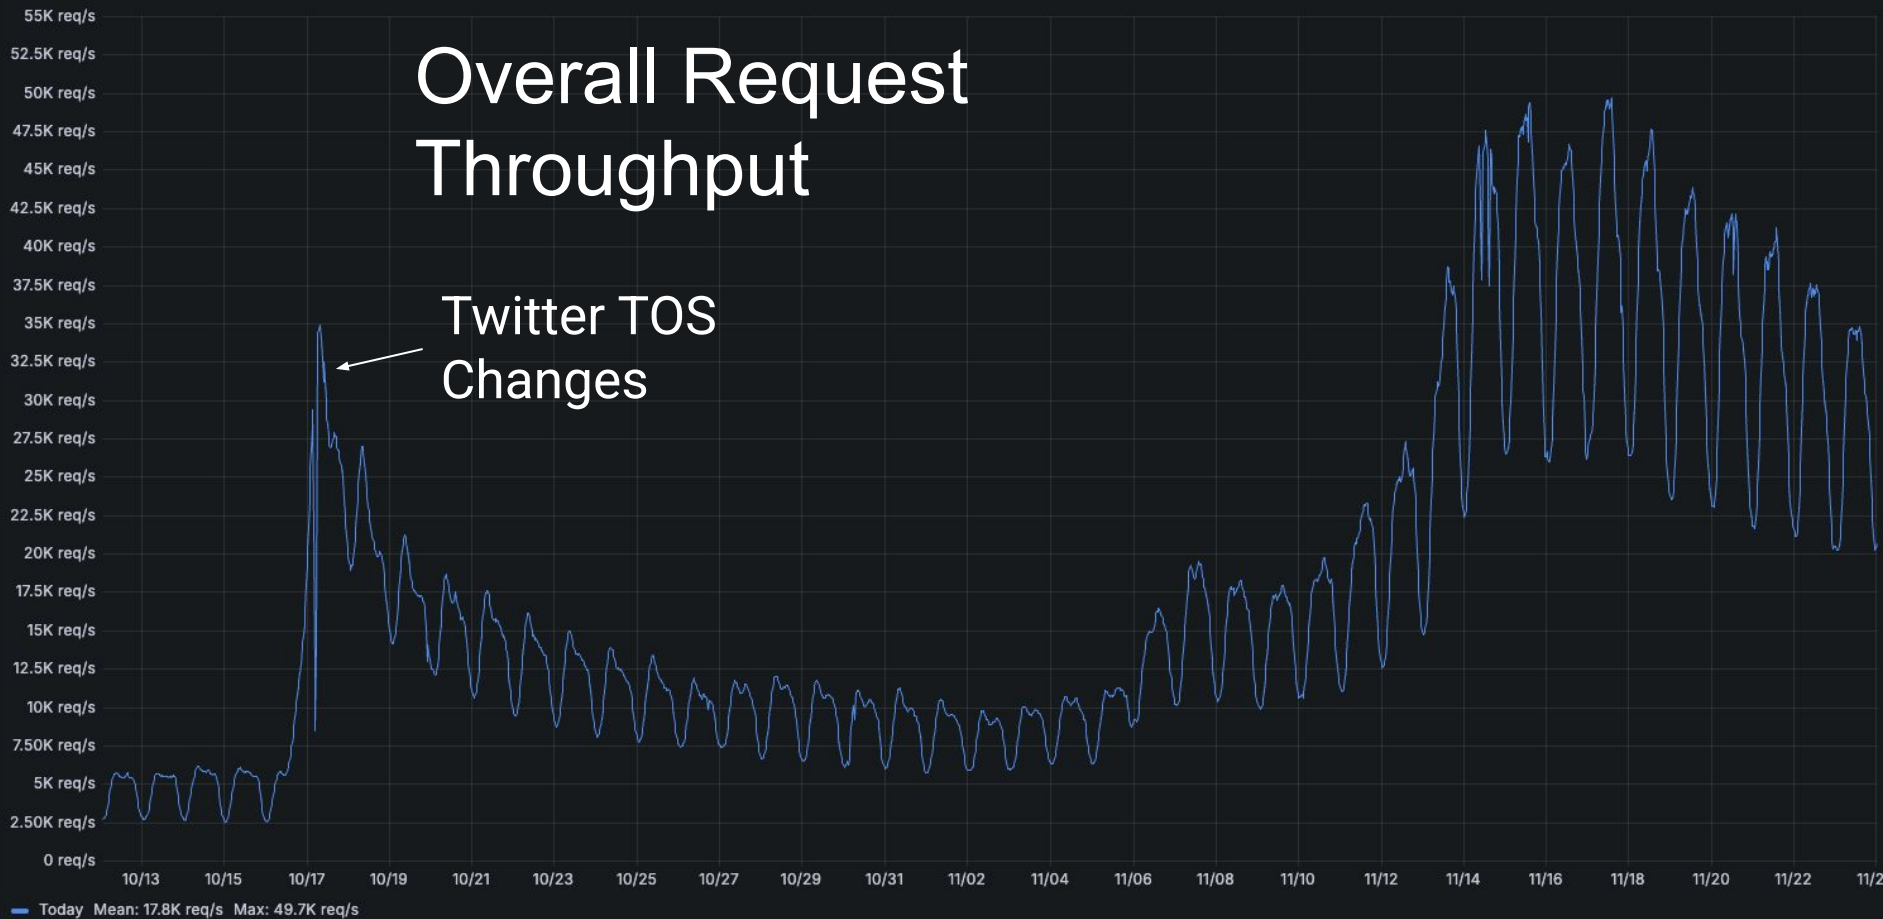


Request Throughput

# Overall Request Throughput



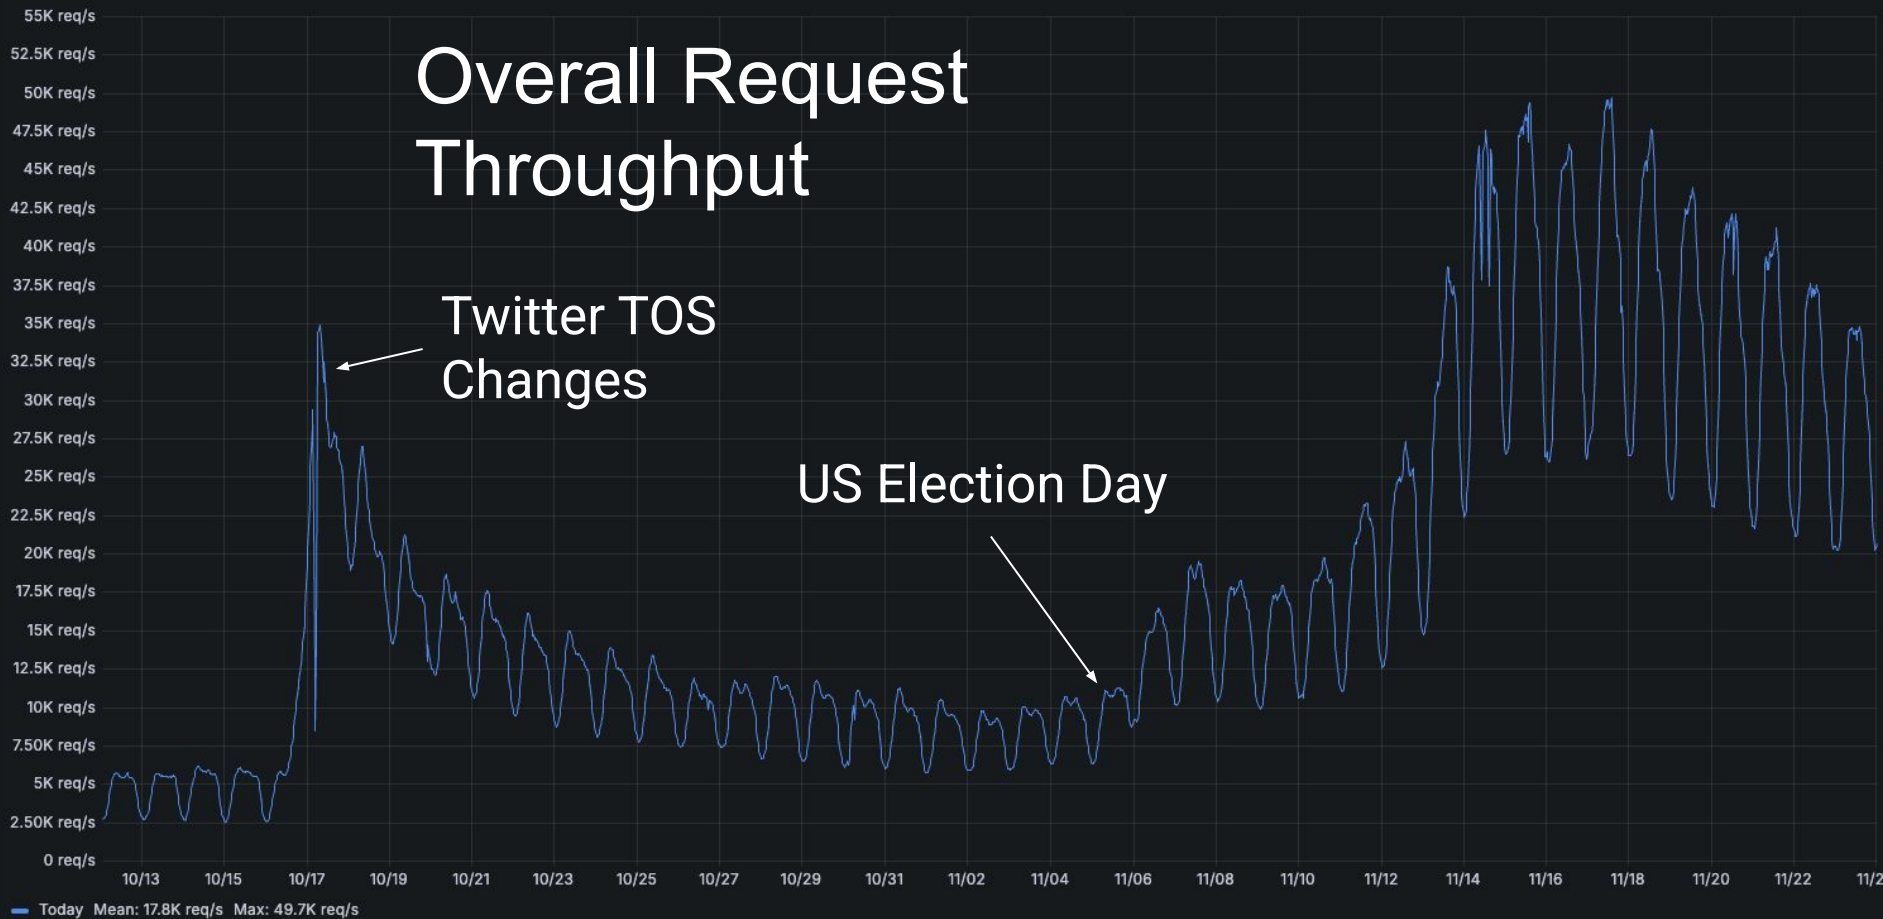
Request Throughput



# Overall Request Throughput

Twitter TOS Changes

Request Throughput




# Overall Request Throughput

Twitter TOS Changes

US Election Day



Request Throughput





# Overall Request Throughput

Twitter TOS Changes

US Election Day

Hockey Stick

# The Fleet



## 4 Modes of Failure

## 4 Modes of Failure

1. Software You Control



## 4 Modes of Failure

1. Software You Control
2. Software You *Don't* Control

## 4 Modes of Failure

1. Software You Control
2. Software You *Don't* Control
3. Hardware You Control

## 4 Modes of Failure

1. Software You Control
2. Software You *Don't* Control
3. Hardware You Control
4. Hardware You *Don't* Control

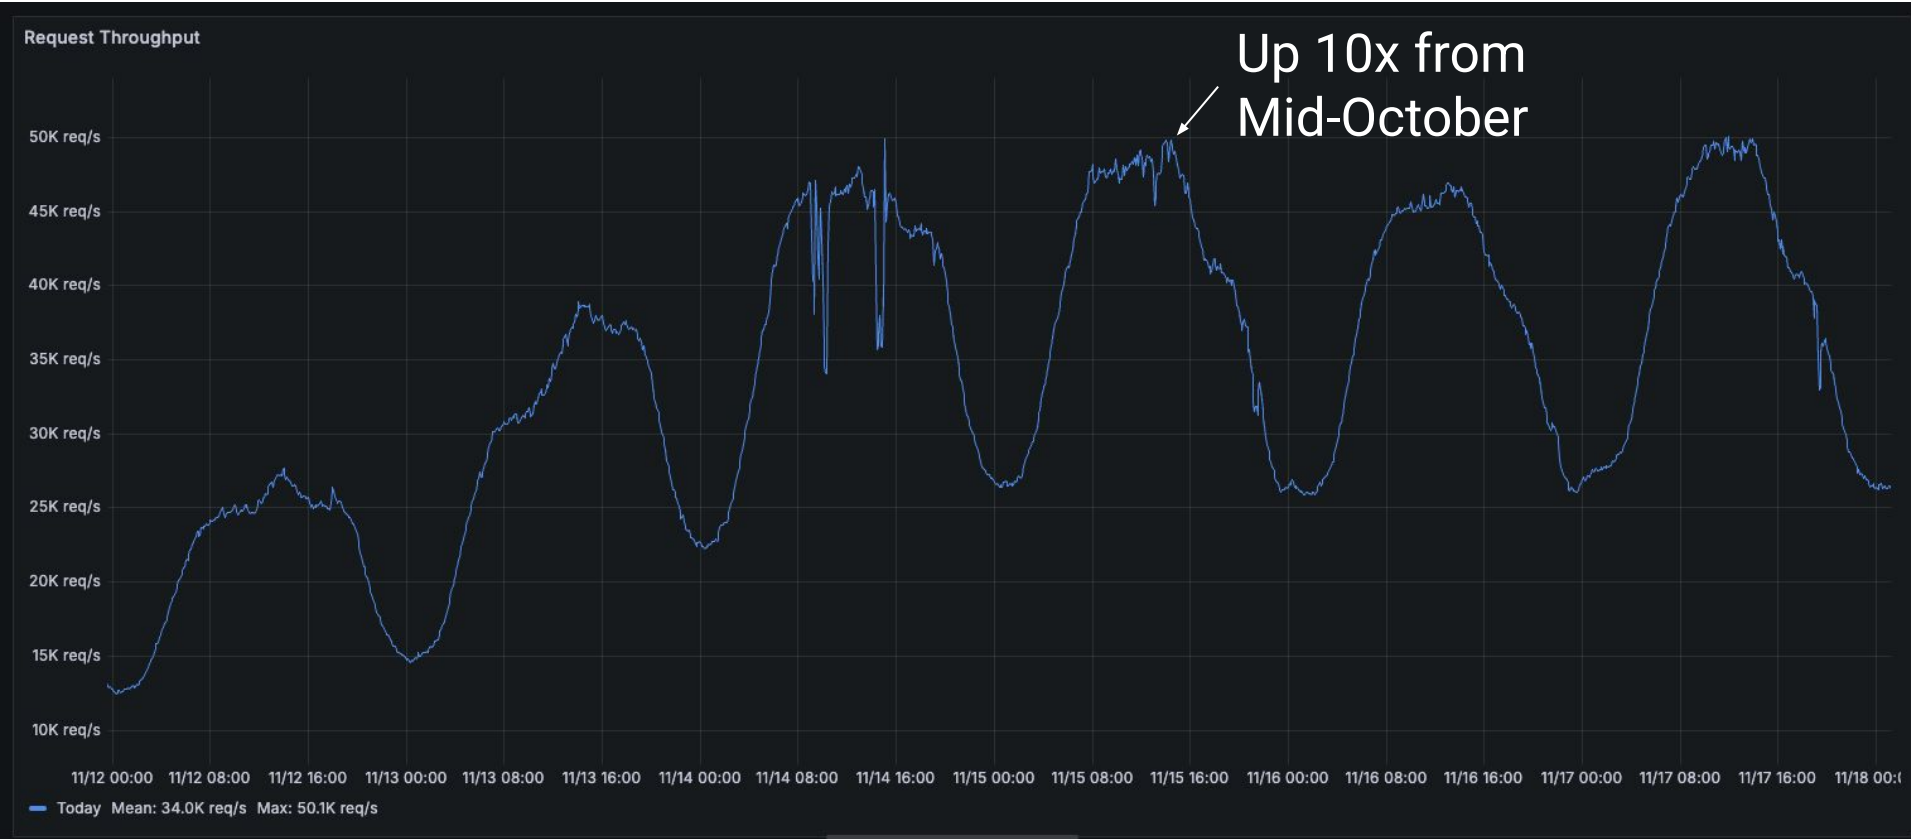
## 4 Modes of Failure

1. Software You Control
2. Software You *Don't* Control
3. Hardware You Control
4. Hardware You *Don't* Control

Each have their own flavors, symptoms, and remedies



# Hell week



# Great American Backhoe Day (Losing PoP2)



# Great American Backhoe Day (Losing PoP2)



# Great American Backhoe Day (Losing PoP2)






# Great American Backhoe Day (Losing PoP2)



# Backhoe Day from PoP2's Perspective


## HA Proxy

### HAProxy Throughput



Name	Mean	Max
2xx	15.4K req/s	21.0K req/s
1 Day Ago	13.0K req/s	17.7K req/s
other	135 req/s	287 req/s
5xx	70.6 req/s	2.57K req/s
4xx	47.2 req/s	235 req/s

### HAProxy Errors



Name	Mean
router2	6.04 req/s
router1	5.32 req/s

## Atlantis

### Atlantis Throughput



Name	Mean
bsky	157K req/s
inferno	4.53K req/s

# External Hardware Failure

External Hardware Failure

Fiber cut



External Hardware Failure

Fiber cut

Single-Homed in PoP2

External Hardware Failure

Fiber cut

Single-Homed in PoP2

**50% of users can't talk to us**

External Hardware Failure

Fiber cut

Single-Homed in PoP2


50% of users can't talk to us

**Remedy: Failover all traffic to PoP1**


# Backhoe Day PoP1 Failover Attempt

## Atlantis

### Atlantis Throughput



### Atlantis Errors



bsky.GetRelationships (internal)	225 req/s
bsky.GetActors (internal)	122 req/s
bsky.GetInteractionCounts (internal)	71.4 req/s
bsky.GetFollowRecords (internal)	29.6 req/s
bsky.GetLikeRecords (internal)	28.3 req/s
bsky.GetPostRecords (internal)	20.3 req/s
bsky.GetBlockExistence (internal)	19.9 req/s

## Backhoe Day PoP1 Failover Attempt

Atlantis

Atlantis Throughput

# Traffic failed over in ~5 minutes

225K req/s  
200K req/s  
175K req/s  
150K req/s  
125K req/s  
100K req/s  
75K req/s  
50K req/s  
25K req/s  
0 req/s

09:00 09:05 09:10 09:15 09:20 09:25 09:30 09:35 09:40

Name Mean

bsky	140K req/s
infernio	4.25K req/s

Atlantis Errors

450 req/s  
400 req/s  
350 req/s  
300 req/s  
250 req/s  
200 req/s  
150 req/s  
100 req/s  
50 req/s  
0 req/s

09:00 09:05 09:10 09:15 09:20 09:25 09:30 09:35 09:40

bsky.GetRelationships (internal)	225 req/s
bsky.GetActors (internal)	122 req/s
bsky.GetInteractionCounts (internal)	71.4 req/s
bsky.GetFollowRecords (internal)	29.6 req/s
bsky.GetLikeRecords (internal)	28.3 req/s
bsky.GetPostRecords (internal)	20.3 req/s
bsky.GetBlockExistence (internal)	19.9 req/s



## Backhoe Day PoP1 Failover Attempt

Traffic failed over in ~5 minutes  
Database cache thrashed **hard**

Atlantis

Atlantis Throughput

225K req/s  
200K req/s  
175K req/s  
150K req/s  
125K req/s  
100K req/s  
75K req/s  
50K req/s  
25K req/s  
0 req/s

09:00 09:05 09:10 09:15 09:20 09:25 09:30 09:35 09:40

Name Mean

bsky	140K req/s
inferno	4.25K req/s

Atlantis Errors

450 req/s  
400 req/s  
350 req/s  
300 req/s  
250 req/s  
200 req/s  
150 req/s  
100 req/s  
50 req/s  
0 req/s

09:00 09:05 09:10 09:15 09:20 09:25 09:30 09:35 09:40

bsky.GetRelationships (internal)	225 req/s
bsky.GetActors (internal)	122 req/s
bsky.GetInteractionCounts (internal)	71.4 req/s
bsky.GetFollowRecords (internal)	29.6 req/s
bsky.GetLikeRecords (internal)	28.3 req/s
bsky.GetPostRecords (internal)	20.3 req/s
bsky.GetBlockExistence (internal)	19.9 req/s

## Backhoe Day PoP1 Failover Attempt

Atlantis

Atlantis Throughput

225K req/s  
200K req/s  
175K req/s  
150K req/s  
125K req/s  
100K req/s  
75K req/s  
50K req/s  
25K req/s  
0 req/s

Traffic failed over in ~5 minutes  
Database cache thrashed **hard**  
Hot shards now twice as hot

09:00 09:05 09:10 09:15 09:20 09:25 09:30 09:35 09:40

Name	Mean
bsky	140K req/s
inferno	4.25K req/s

Atlantis Errors

450 req/s  
400 req/s  
350 req/s  
300 req/s  
250 req/s  
200 req/s  
150 req/s  
100 req/s  
50 req/s  
0 req/s

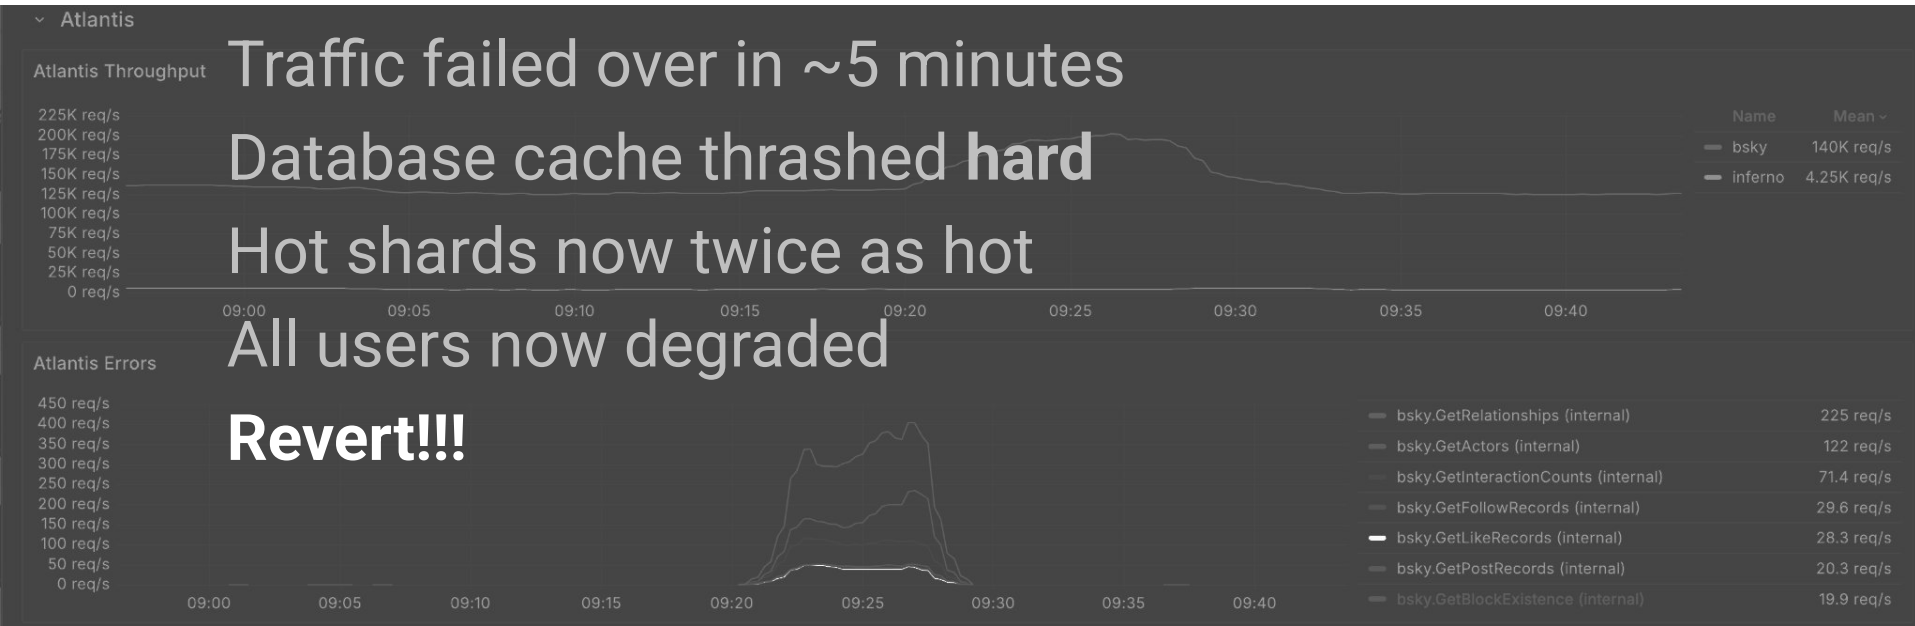
09:00 09:05 09:10 09:15 09:20 09:25 09:30 09:35 09:40

bsky.GetRelationships (internal)	225 req/s
bsky.GetActors (internal)	122 req/s
bsky.GetInteractionCounts (internal)	71.4 req/s
bsky.GetFollowRecords (internal)	29.6 req/s
bsky.GetLikeRecords (internal)	28.3 req/s
bsky.GetPostRecords (internal)	20.3 req/s
bsky.GetBlockExistence (internal)	19.9 req/s

## Backhoe Day PoP1 Failover Attempt



## Backhoe Day PoP1 Failover Attempt



Roll with the punches



Roll with the punches

Call your vendors

Roll with the punches

Call your vendors

**Apologize to your users**

Roll with the punches

Call your vendors

Apologize to your users

**Prioritize dual-home BGP configuration**

Roll with the punches

Call your vendors

Apologize to your users

Prioritize dual-home BGP configuration

**Pray your vendors get things resolved soon**

# Great American Backhoe Day (Losing PoP2)



## Read-path troubles

Request Throughput





# First DB Crash (Nov 17)

## HA Proxy

### HAProxy Throughput



Name	Mean ~	Max
2xx	16.5K req/s	22.1K req/s
1 Day Ago	16.5K req/s	20.8K req/s
5xx	98.8 req/s	4.77K req/s
4xx	60.3 req/s	144 req/s


### HAProxy Errors



Name	Mean ~
router1	98.8 req/s
router2	0 req/s


## Atlantis

### Atlantis Throughput



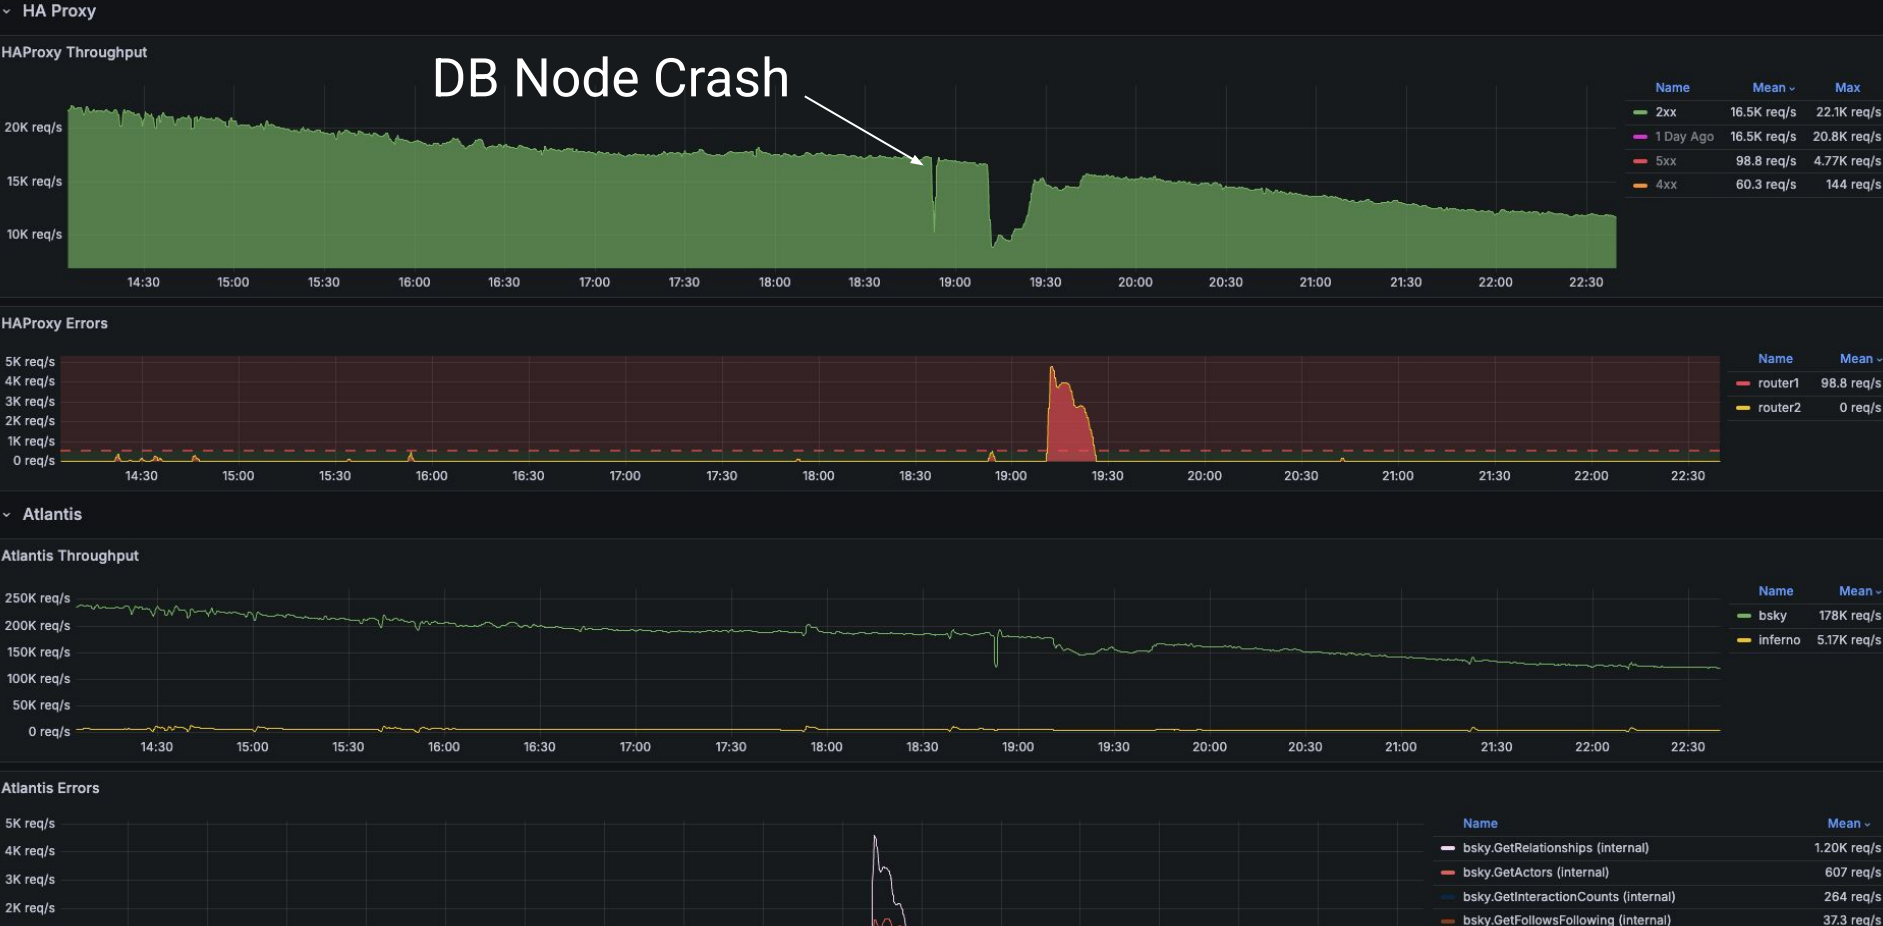
Name	Mean ~
bsky	178K req/s
inferno	5.17K req/s

### Atlantis Errors




Name	Mean ~
bsky.GetRelationships (Internal)	1.20K req/s
bsky.GetActors (Internal)	607 req/s
bsky.GetInteractionCounts (Internal)	264 req/s
bsky.GetFollowsFollowing (Internal)	37.3 req/s


# First DB Crash (Nov 17)



# First DB Crash (Nov 17)




# First DB Crash (Nov 17)



# DB Crash Repeats (Nov 20)




## Atlantis



## And Again (Nov 25)



### Atlantis





# External Software Failure

External Software Failure

Grab logs, coredumps, backtraces

External Software Failure

Grab logs, coredumps, backtraces

**Dig through Github issues**

External Software Failure

Grab logs, coredumps, backtraces

Dig through Github issues

Consult your vendor (if you can)

External Software Failure


Grab logs, coredumps, backtraces

Dig through Github issues

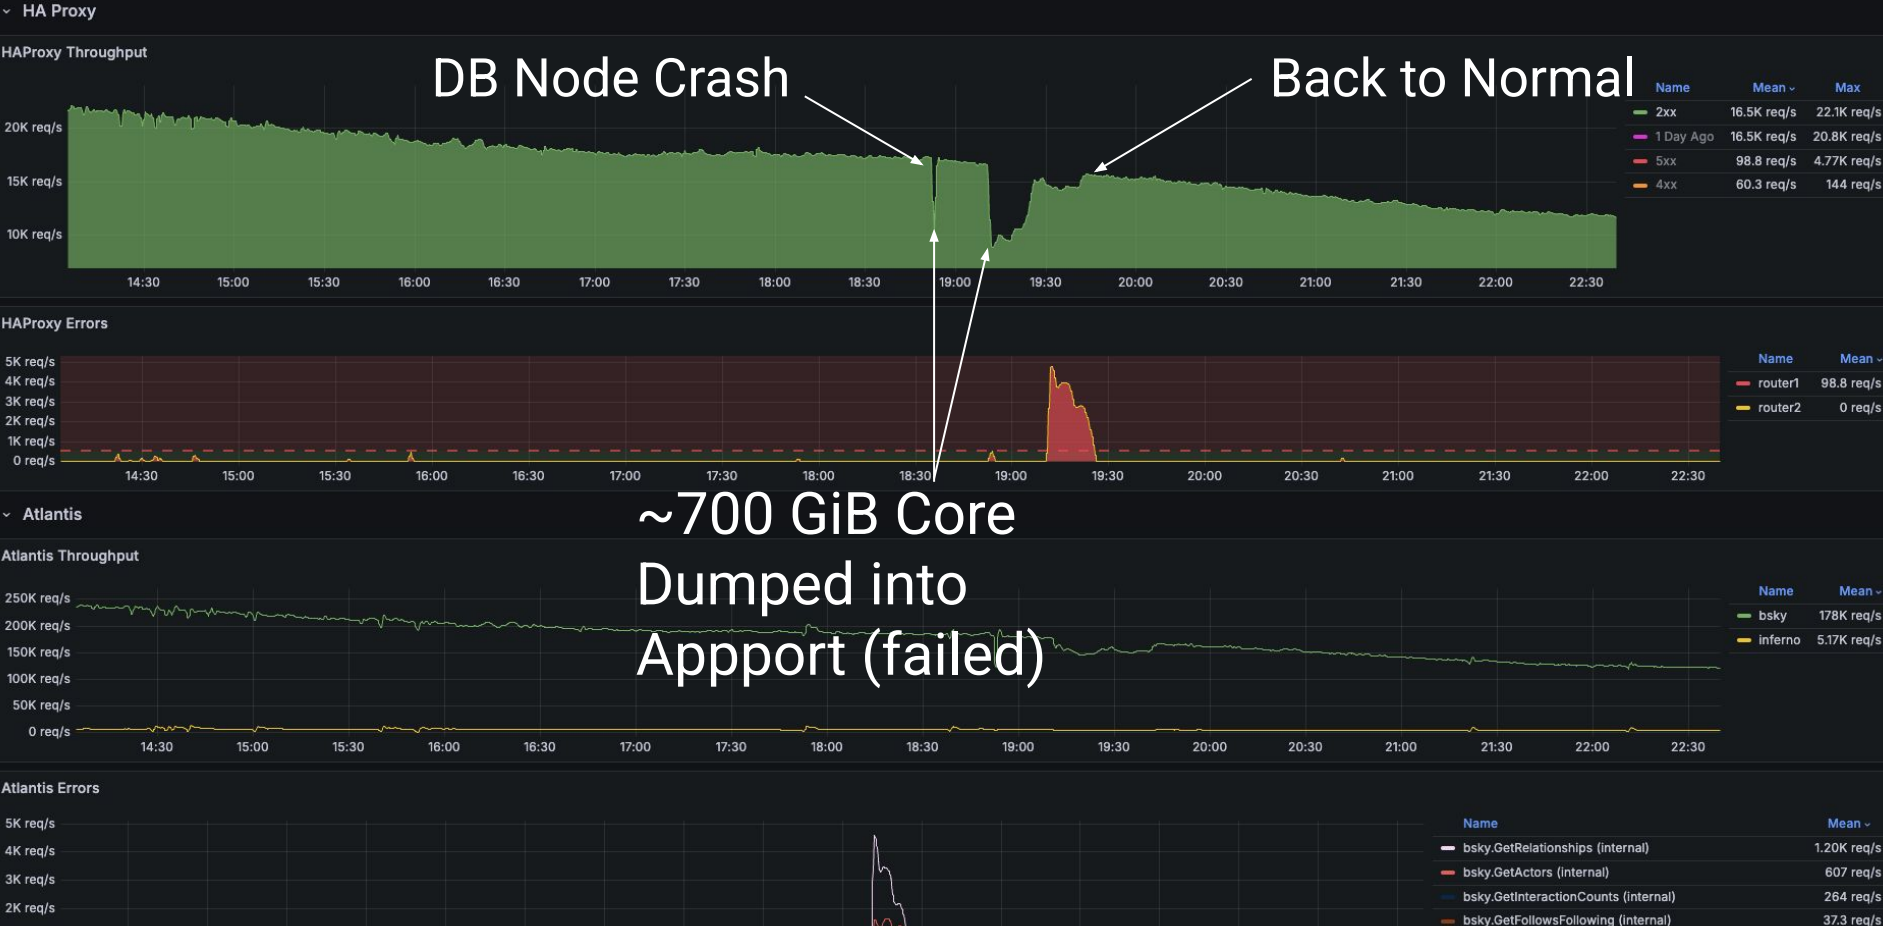
Consult your vendor (if you can)

**Pray and keep digging**

# First DB Crash (Nov 17)




# First DB Crash (Nov 17)





# First DB Crash (Nov 17)



DB installation misconfigured

DB installation misconfigured

Noisy neighbor workloads

DB installation misconfigured

Noisy neighbor workloads

**Aggressive self-throttling**


DB installation misconfigured

Noisy neighbor workloads


Aggressive self-throttling

**No surface visibility**


## Write-path growth



## Firehose throughput growth and troubles




## Relay host dies

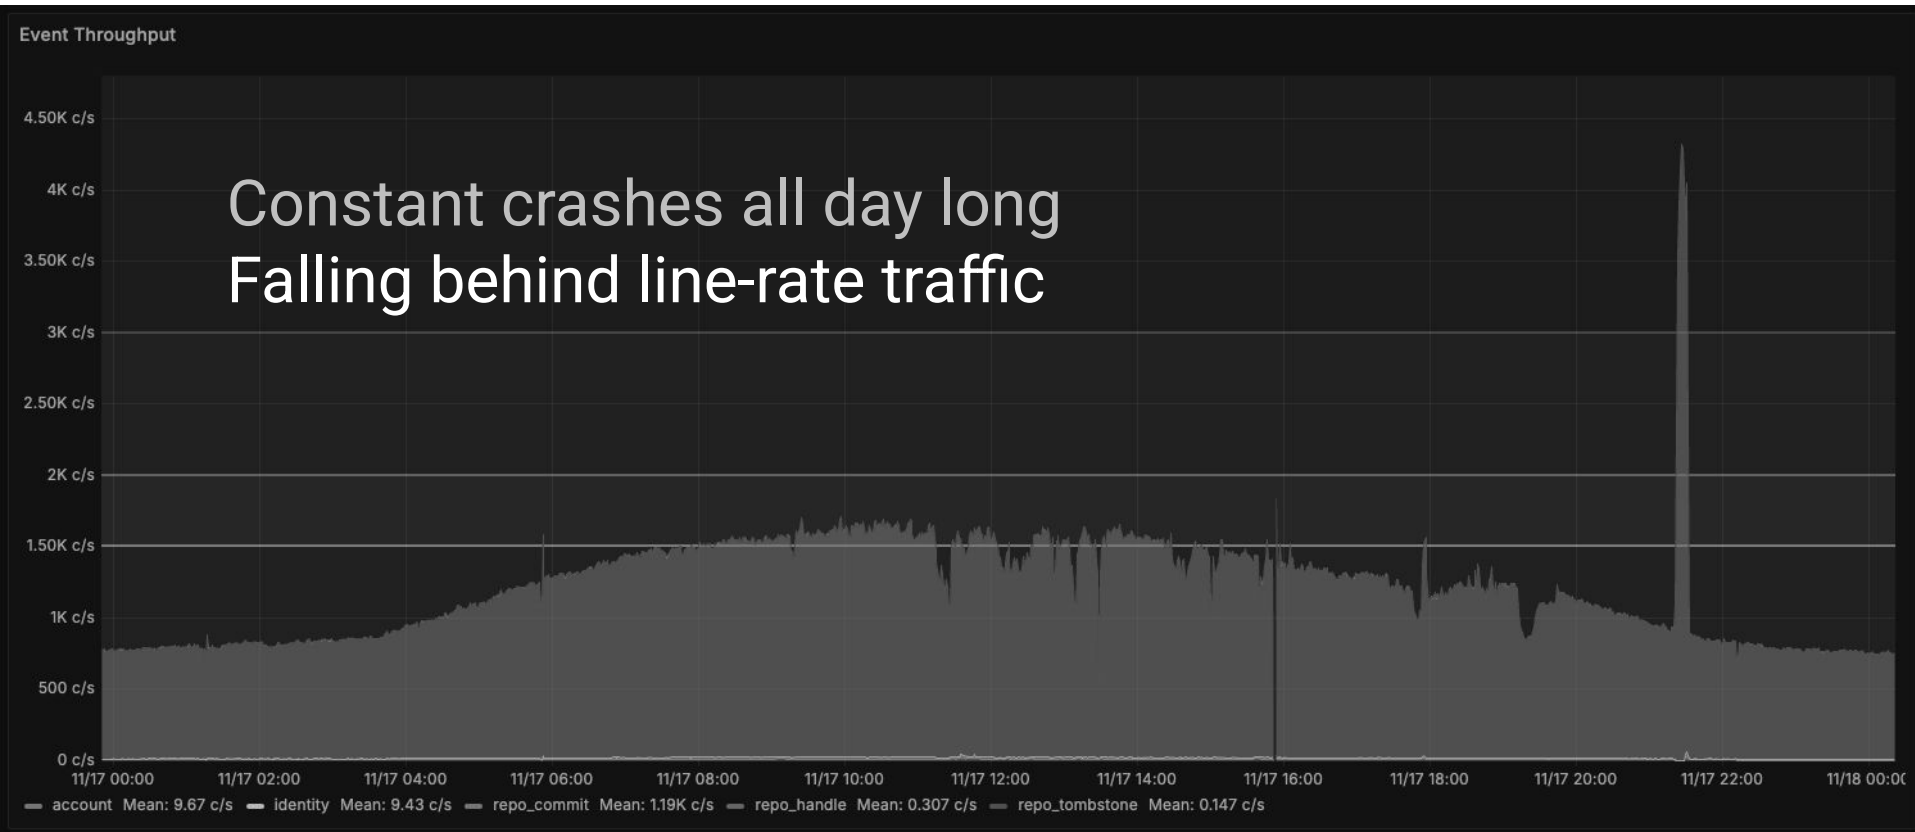




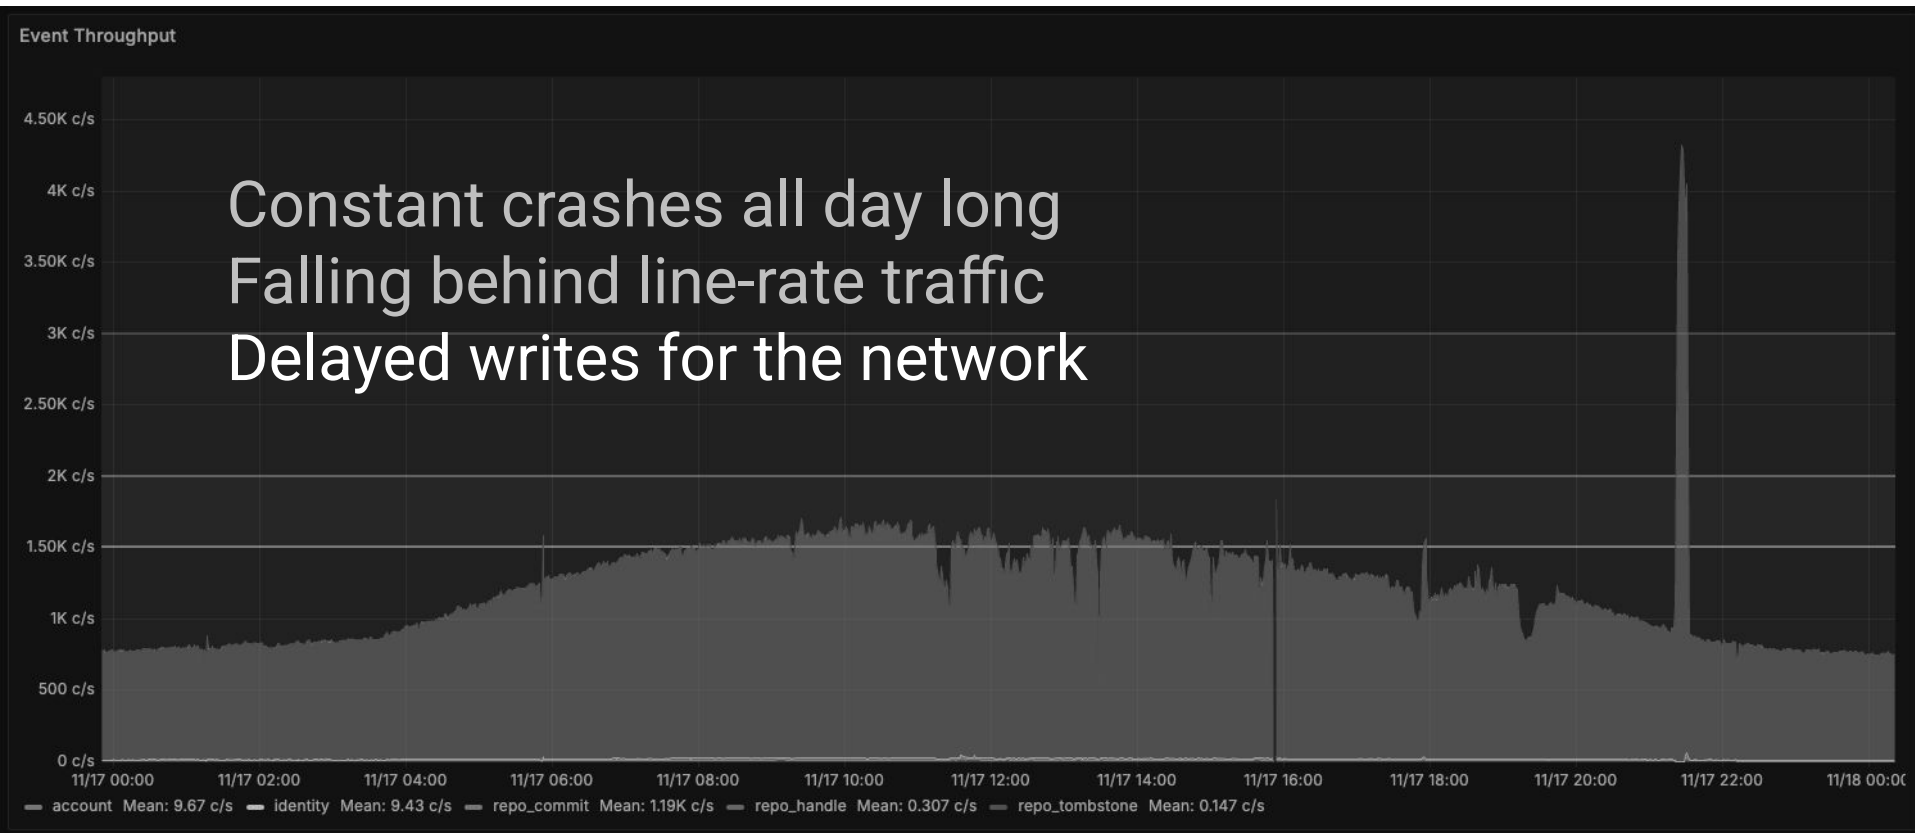
## Relay host dies



## Relay host dies



## Relay host dies




## Relay host dies




# Hardware scaling ticking time bombs

Basic CPU / Mem / Net / Disk

CPU Basic




Memory Basic



Network Traffic Basic




Disk Space Used Basic



## Relay ticking time bombs

### Disk Space Used Basic



# Internal Hardware Failure

Internal Hardware Failure

Disk capacity exhaustion



Internal Hardware Failure

Disk capacity exhaustion

Disk latency spikes (overprovisioning runs out)

Internal Hardware Failure

Disk capacity exhaustion

Disk latency spikes (overprovisioning runs out)

**No free blocks**

Internal Hardware Failure

Disk capacity exhaustion


Disk latency spikes (overprovisioning runs out)

No free blocks

**XFS Corruption**

Overload

System Load



Name	Mean	Last *	Max	Min
Load 1m	305	0.750	3.65 K	0.0600
Load 5m	326	3	3.63 K	1.08
Load 15m	345	5.01	3.65 K	1.22

Shift to lower-load alternate architecture (lucky)

Shift to lower-load alternate architecture (lucky)  
Offload subcomponents of service

Shift to lower-load alternate architecture (lucky)

Offload subcomponents of service

Rotate in hot-spare hardware

Shift to lower-load alternate architecture (lucky)

Offload subcomponents of service


Rotate in hot-spare hardware

**Divide labor among the team**



# Ouroboros

Event Throughput



# Internal Software Failure

Internal Software Failure

Misconfigured proxy deploy

Internal Software Failure

Misconfigured proxy deploy

Triggered by exhaustion

Internal Software Failure

Misconfigured proxy deploy

Triggered by exhaustion

**Miscommunication**

Internal Software Failure

Misconfigured proxy deploy

Triggered by exhaustion


Miscommunication

**Bandaids Stacking**

# Ouroboros

Event Throughput

Public consumers see repeated events



# Ouroboros

Event Throughput

Public consumers see repeated events  
Two hours before we noticed





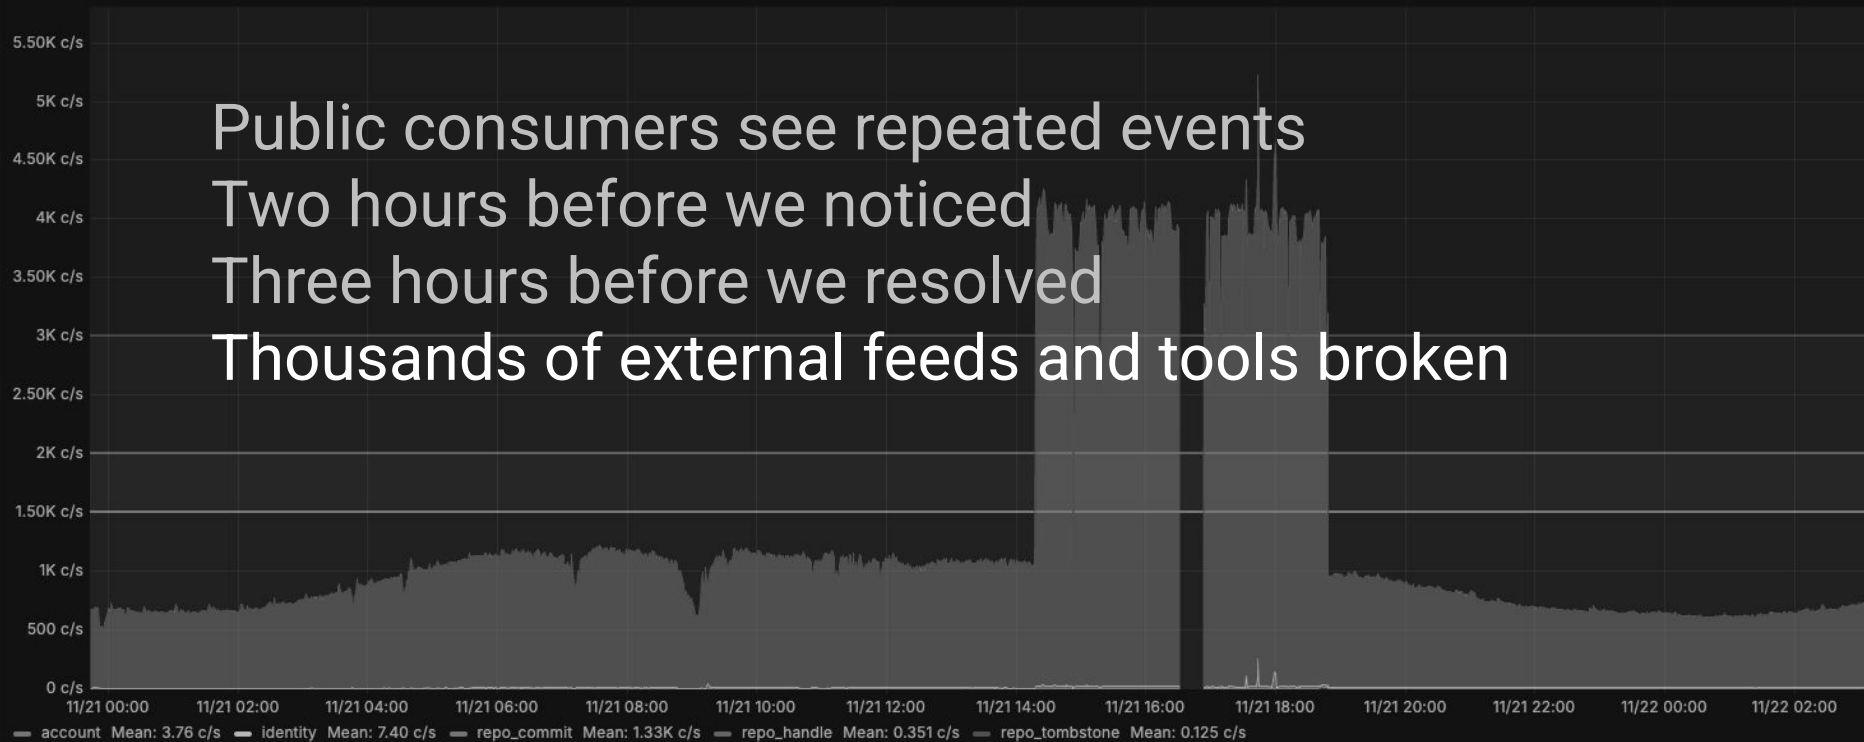
# Ouroboros

Event Throughput



# Ouroboros

Event Throughput



Repair the broken proxy configuration

Repair the broken proxy configuration

Replay missed events for consumers

Repair the broken proxy configuration

Replay missed events for consumers

**Apologize to developers**

Repair the broken proxy configuration

Replay missed events for consumers

Apologize to developers

**Automate future proxy deploys with human approval**

Repair the broken proxy configuration

Replay missed events for consumers

Apologize to developers

Automate future proxy deploys with human approval

**Self-inflicted problems get more frequent when exhausted**

Repair the broken proxy configuration

Replay missed events for consumers

Apologize to developers

Automate future proxy deploys with human approval


Self-inflicted problems get more frequent when exhausted

**When the team is tired, if you can't rest, double-triple check**



# Celeb Caching Problem


Scylla Query P99



Scylla Query Throughput



## Celeb Caching Problem



# Internal Software Failure

Internal Software Failure

Sharded DB means each data partition served by only 3 cores

Internal Software Failure

Sharded DB means each data partition served by only 3 cores

Popular (celebrity) accounts have hot data on read-path

Internal Software Failure

Sharded DB means each data partition served by only 3 cores

Popular (celebrity) accounts have hot data on read-path

**DB Tail latencies for everything spike (random distribution)**

## Internal Software Failure


Sharded DB means each data partition served by only 3 cores

Popular (celebrity) accounts have hot data on read-path

DB Tail latencies for everything spike (random distribution)

Average query fanout means ***all*** API requests feel the pain


## Celeb Caching Solution





## Celeb Caching Solution

Cache Hit Rate




## Celeb Caching Solution

Cache Hit Rate




## Celeb Caching Solution

### Cache Hit Rate



## Celeb Caching Solution

### Cache Hit Rate



# Celeb Caching Result


Scylla Query P99




Scylla Query Throughput



## Celeb Caching Result




## Celeb Caching Result



## Celeb Caching Result

Scylla Query P99




Scylla Query Throughput

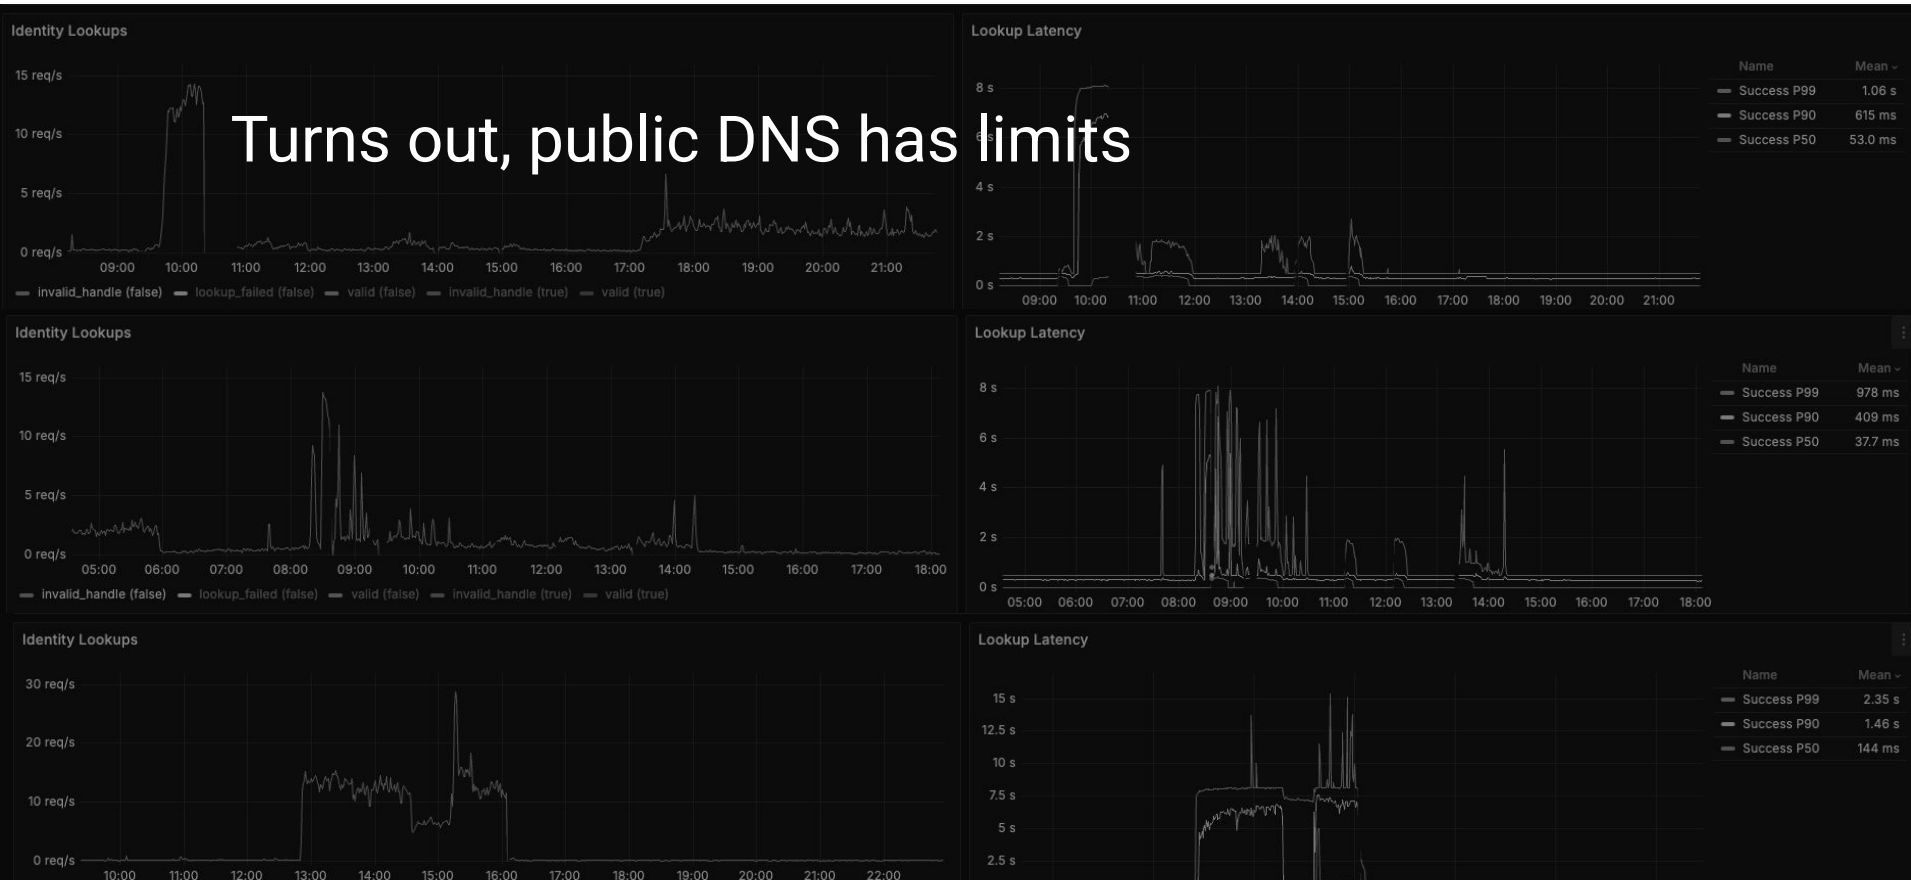




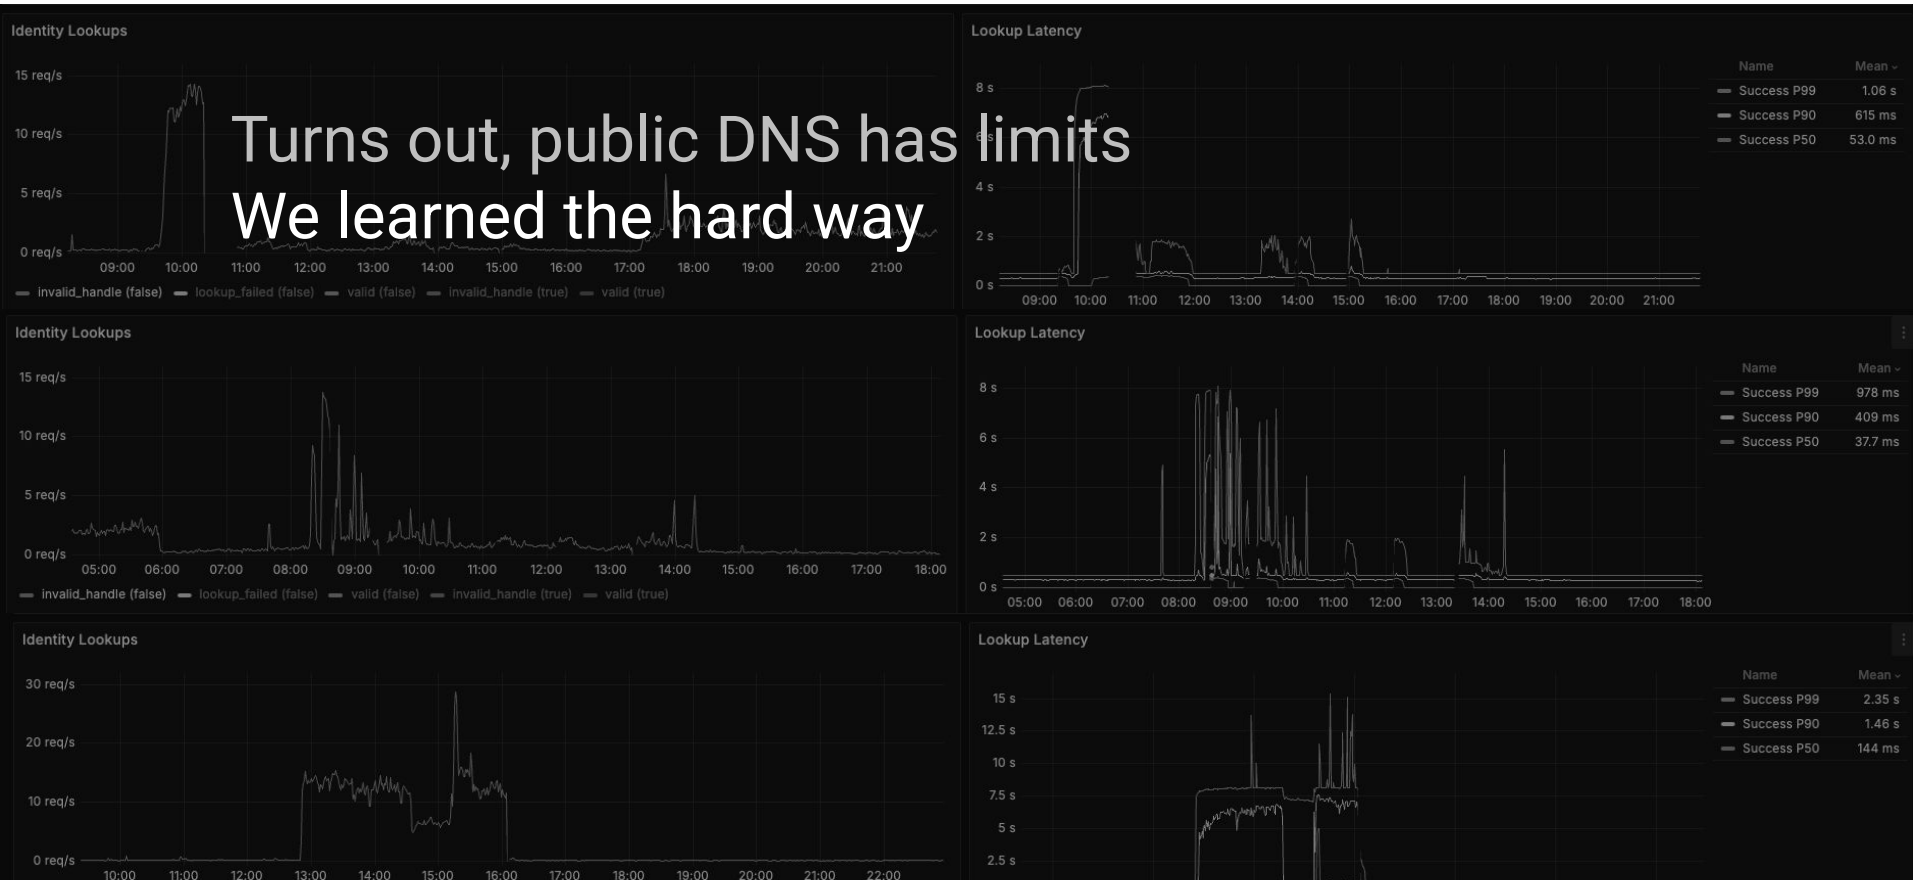
# Invalid Handles



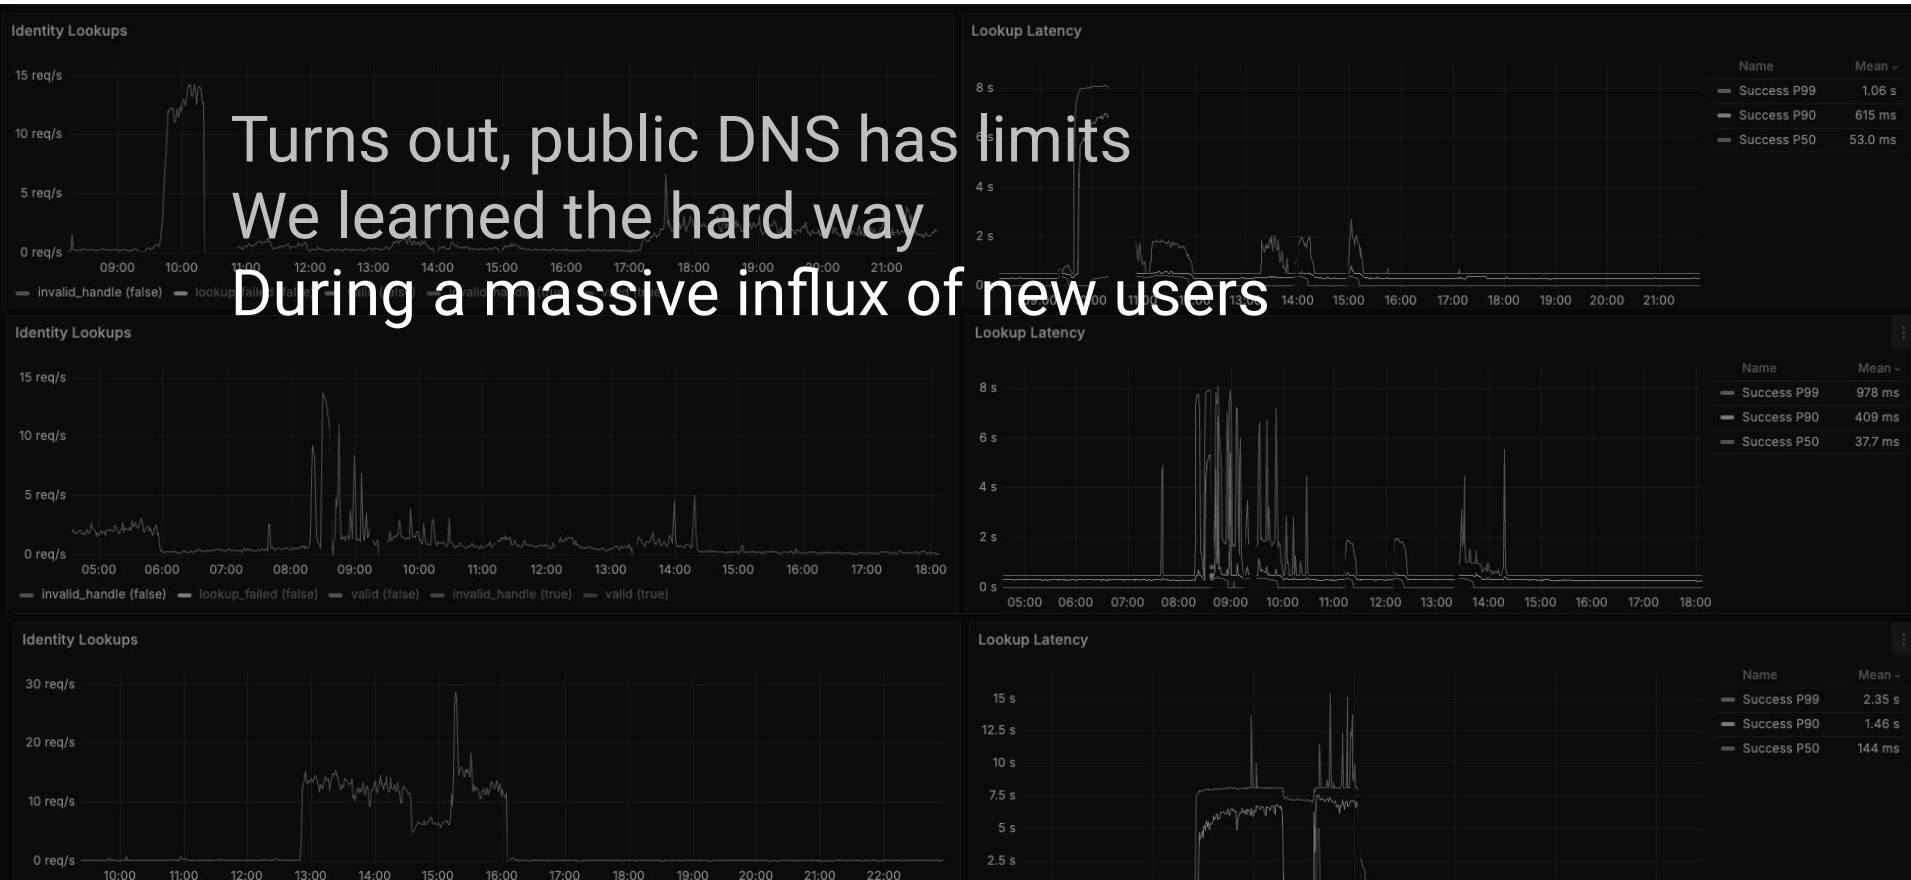
## Invalid Handles



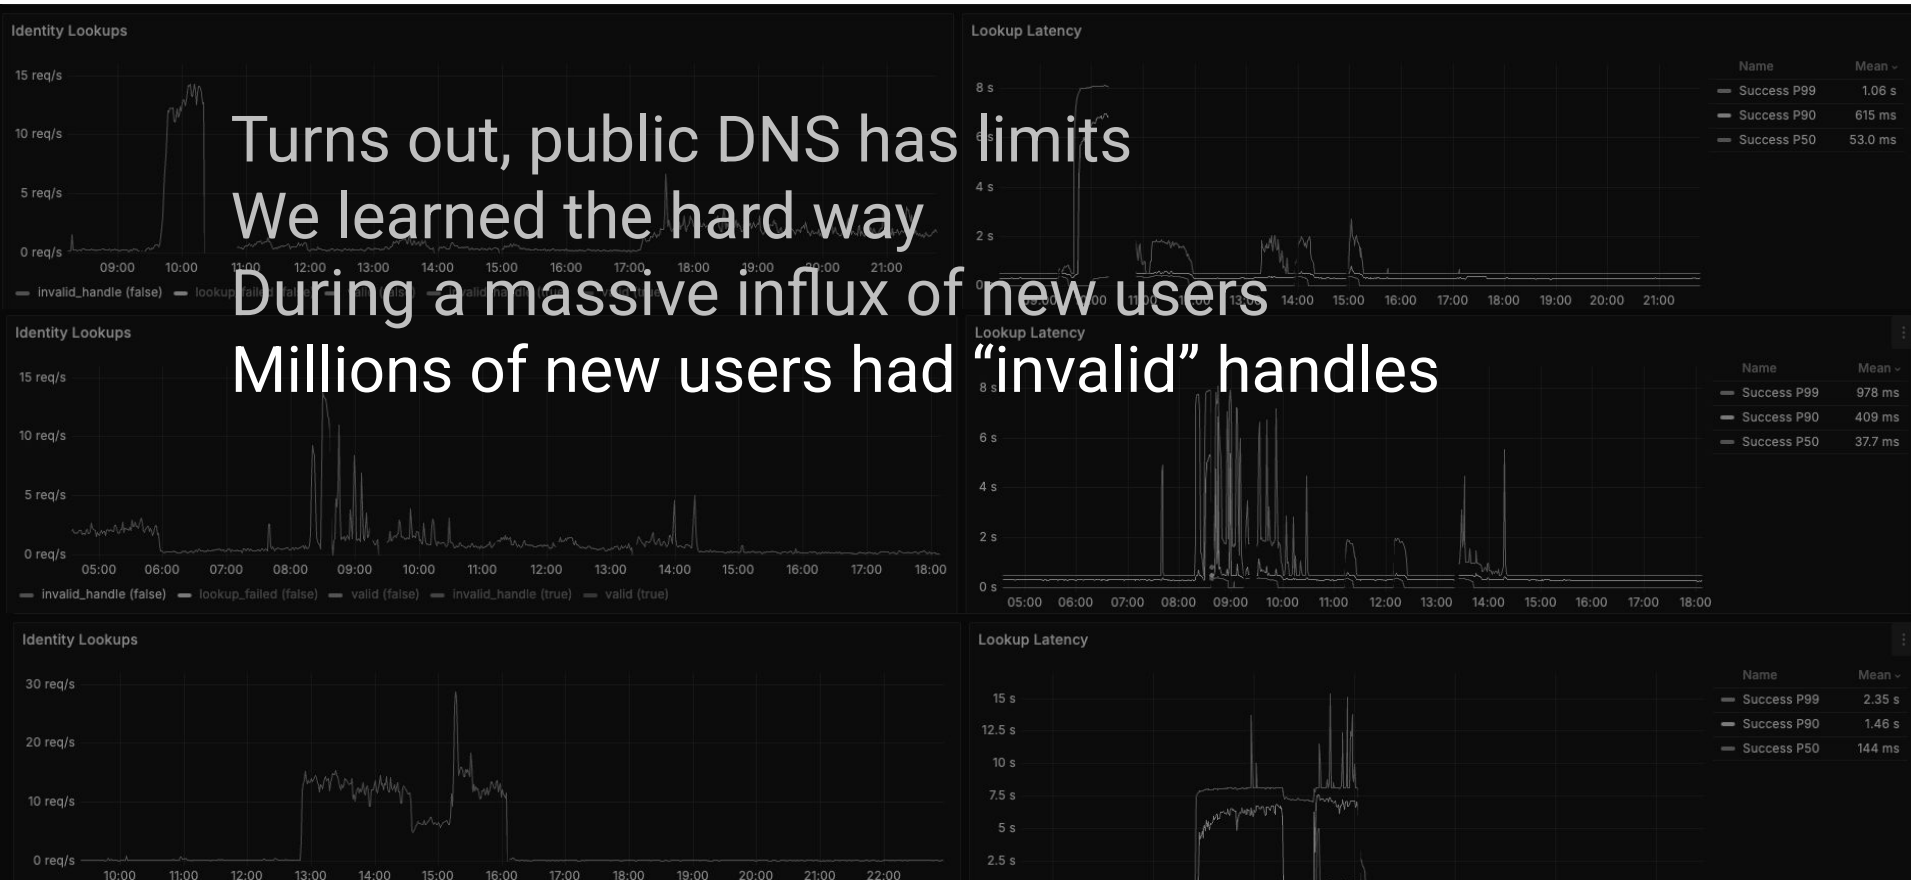
## Invalid Handles



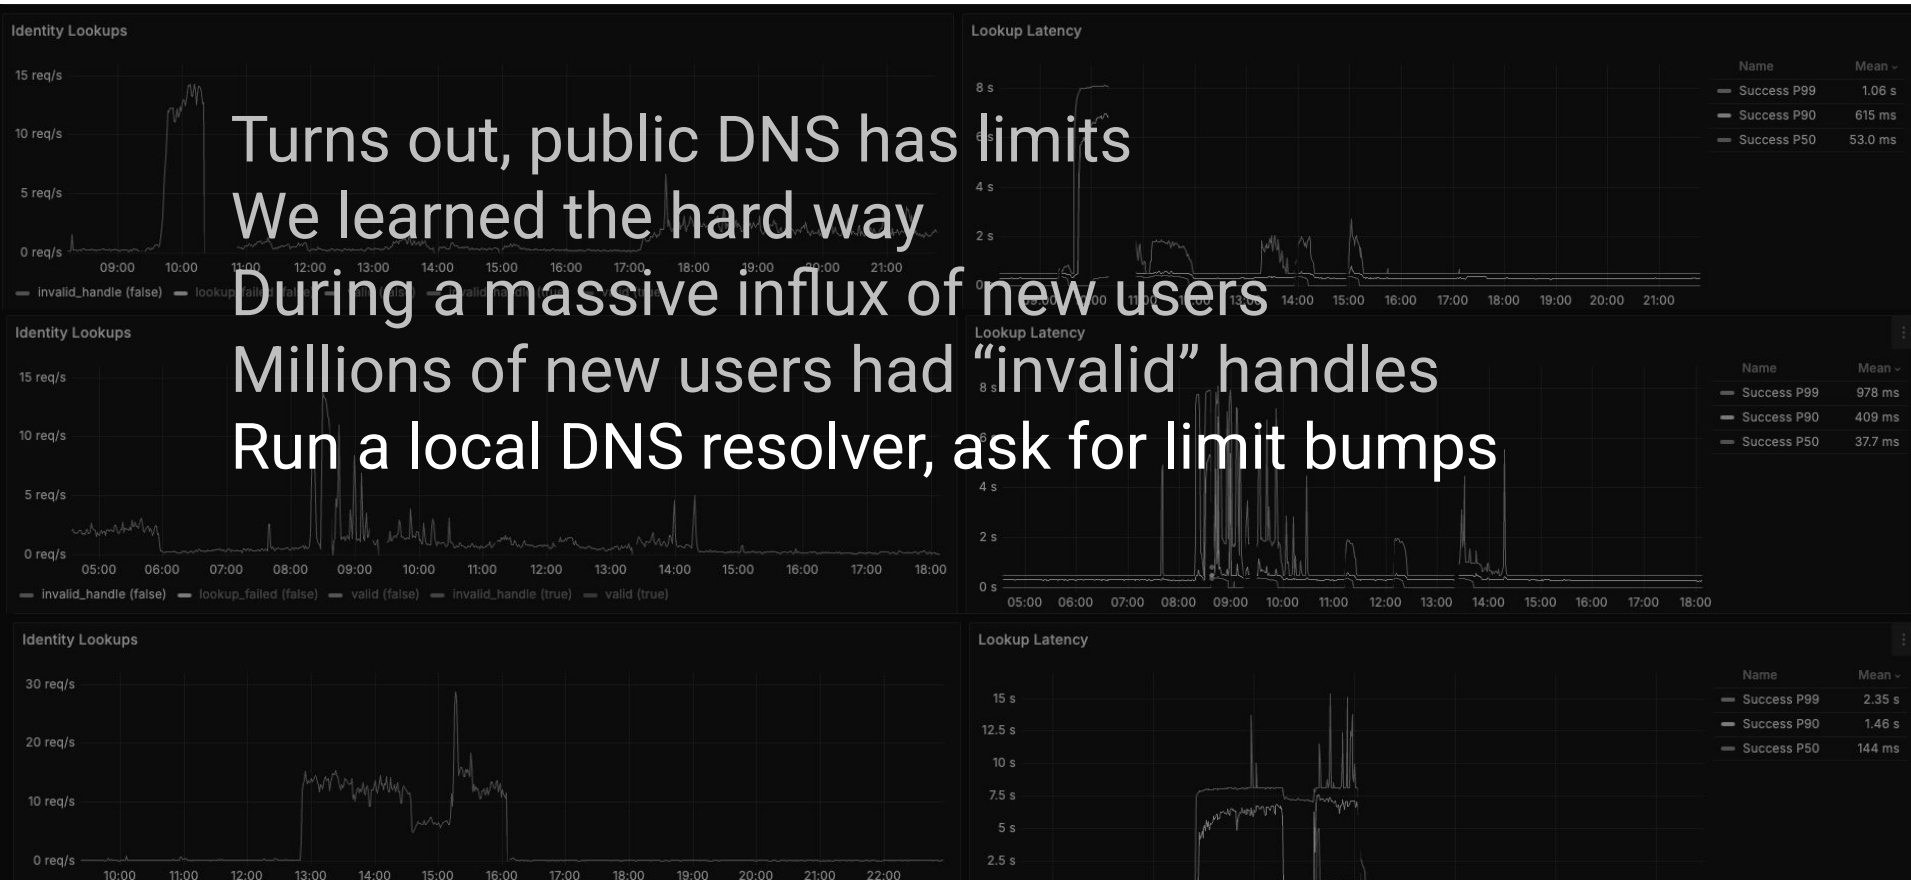
## Invalid Handles




## Invalid Handles



## Invalid Handles




Dropping some o11y ballast



# Search and a Chrome Extension

HTTP Request Throughput



Name	Mean
GET SearchActors 200	631 req/s
GET SearchPosts 200	49.2 req/s
GET SearchPosts top 200	22.4 req/s
GET SearchActors typeahead 200	278 req/s
GET SearchActors 400	0.0219 req/s
GET SearchPosts 400	0.388 req/s
GET SearchPosts top 400	0.00170 req/s
GET SearchActors typeahead 400	0.00185 req/s
GET SearchActors 500	0.480 req/s
GET SearchPosts 500	0.0778 req/s
GET SearchPosts top 500	0.0502 req/s




# Search and a Chrome Extension


HTTP Request Throughput




HTTP Request Throughput




## Discover Growing Pains



## Recovering from Action




## Recovering from Action





## Recovering from Action



# We're Hiring SREs

<https://bsky.social/about/join>



Any Questions?