Exposing server performance to network managers through passive network measurements

Jeff Terrell
Dept. of Computer Science
University of North Carolina at Chapel Hill
October 19, 2008

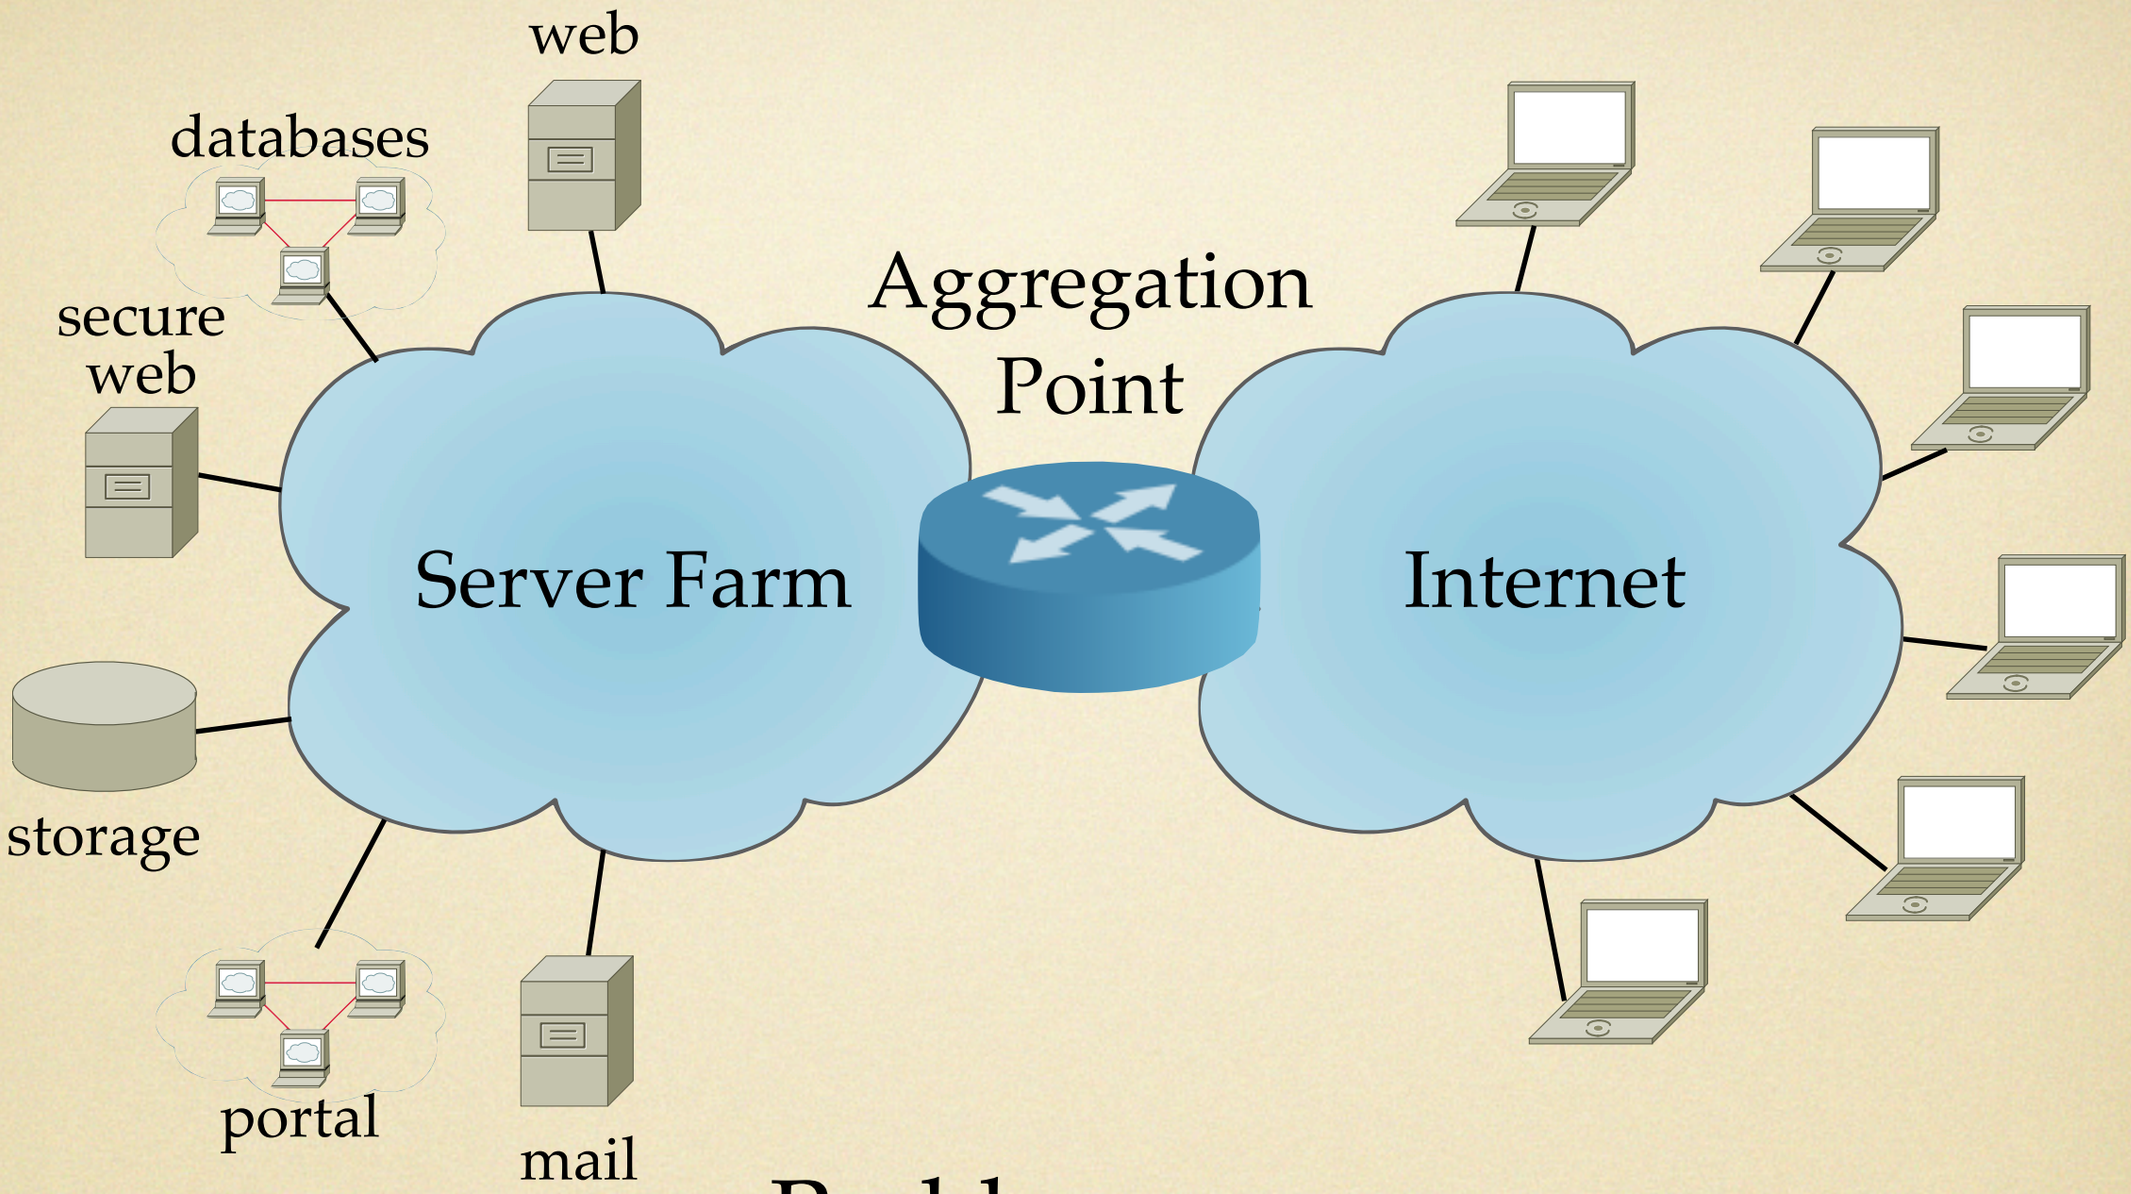

Problem:

Monitor server performance

1. **passively** (no instrumenting or probes)
2. **pervasively** (all servers)
3. in **real-time**

Monitoring Traffic

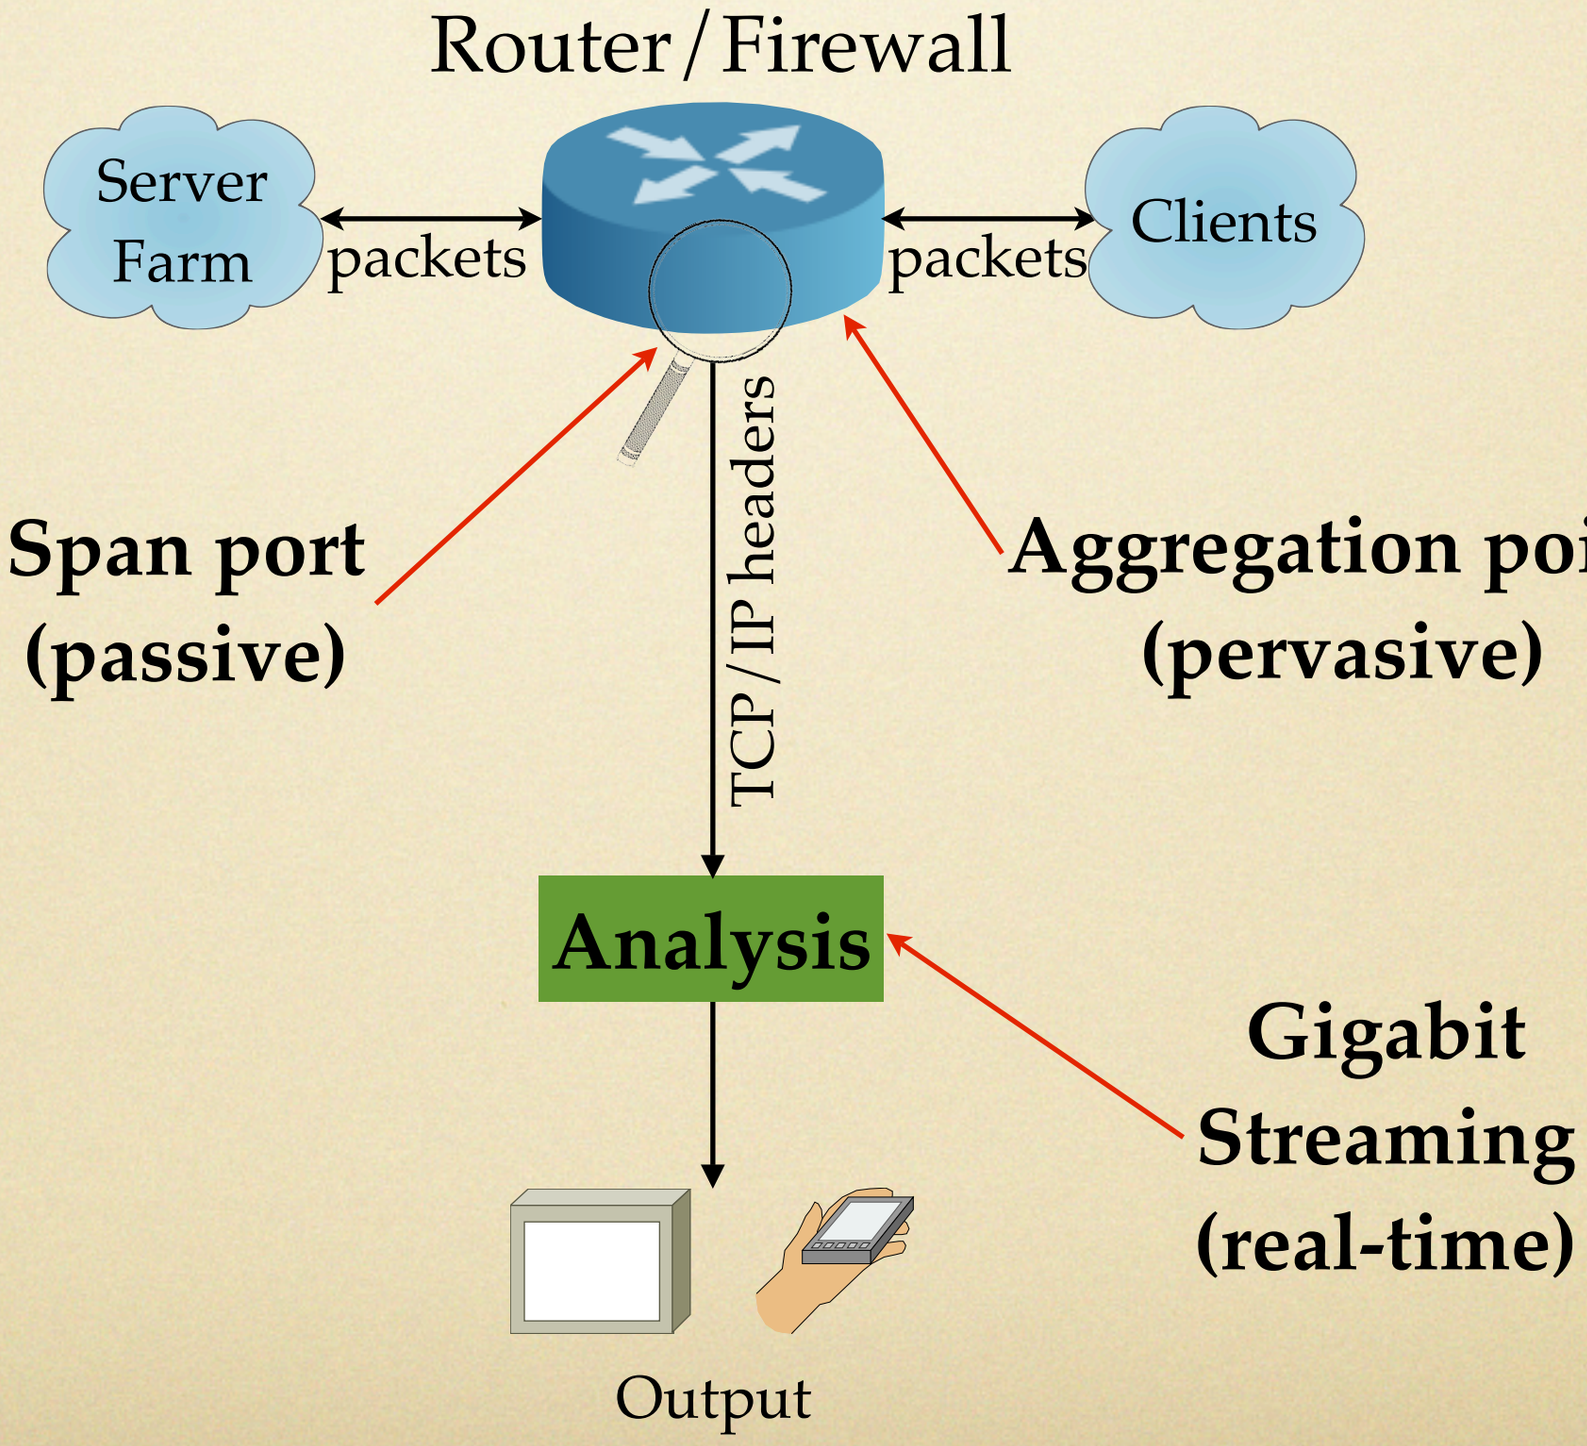

Monitoring Traffic

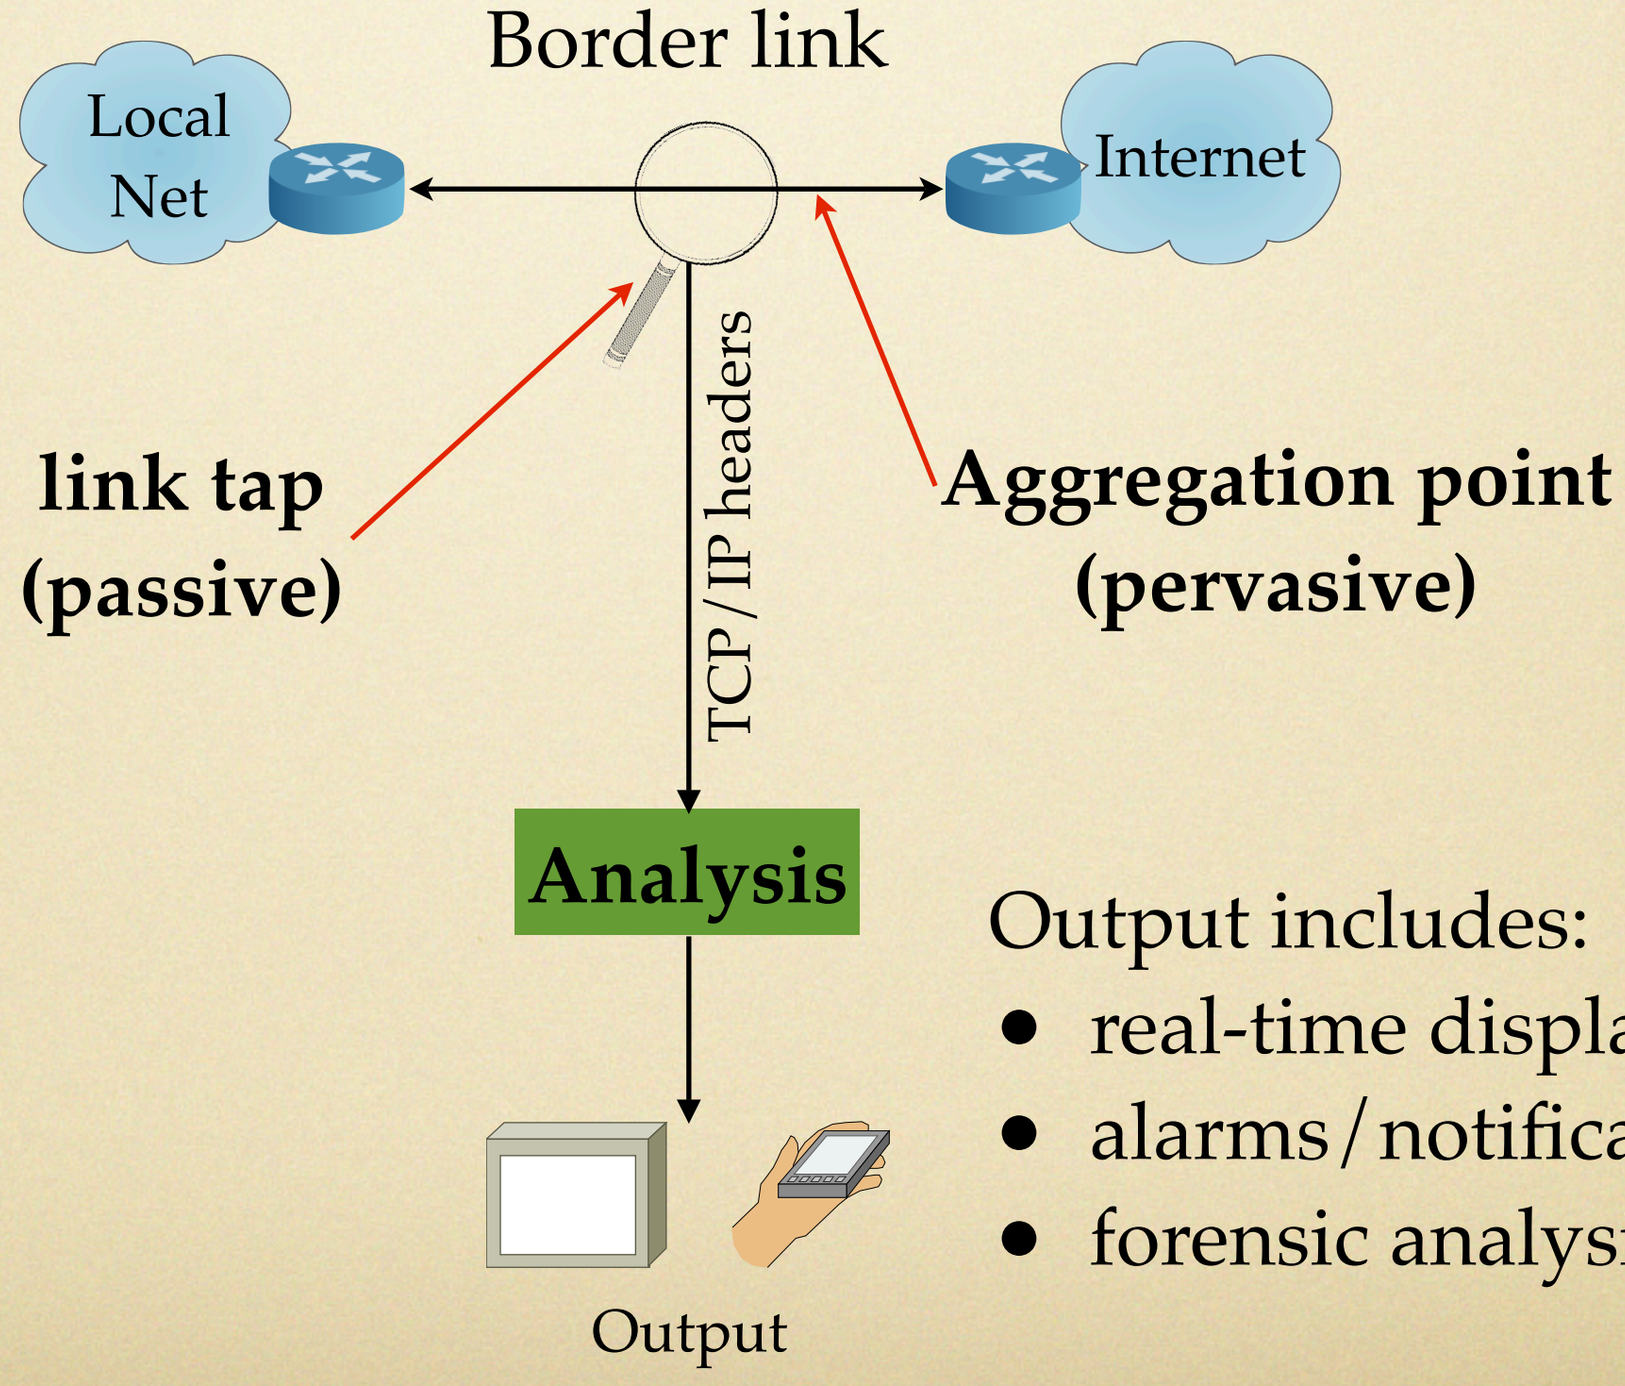

Output includes:

- real-time display
- alarms/notifications
- forensic analysis

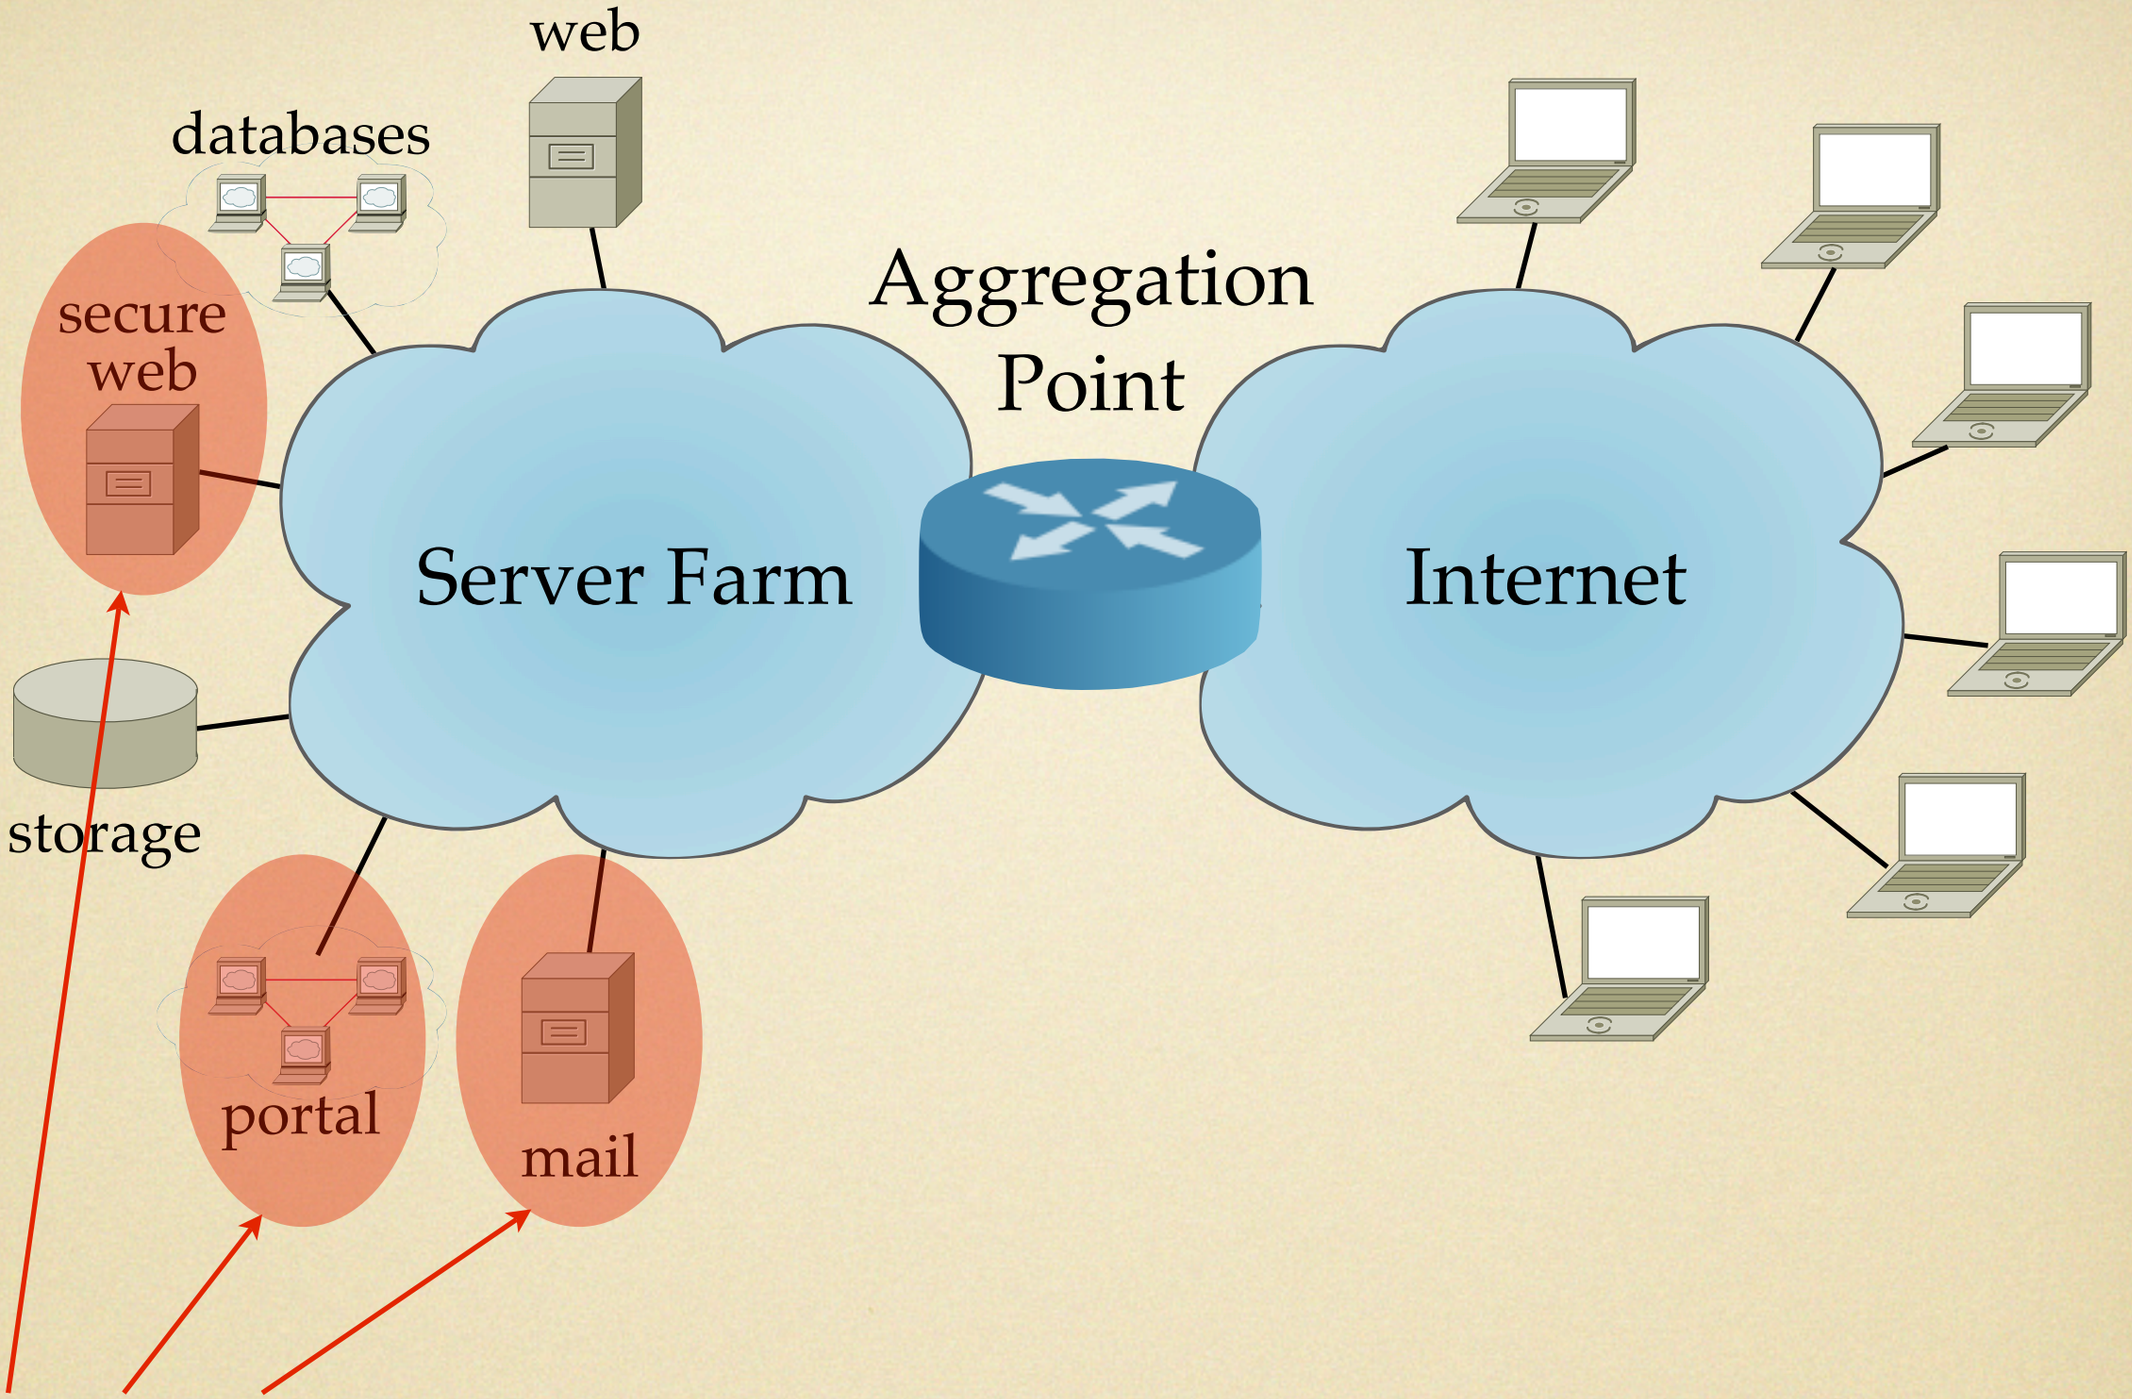

SSL traffic (i.e. **no accessible payload**)

TCP Connection Vectors

- A **connection vector** is a representation of the application-level dialog in a TCP connection.
  - For example:

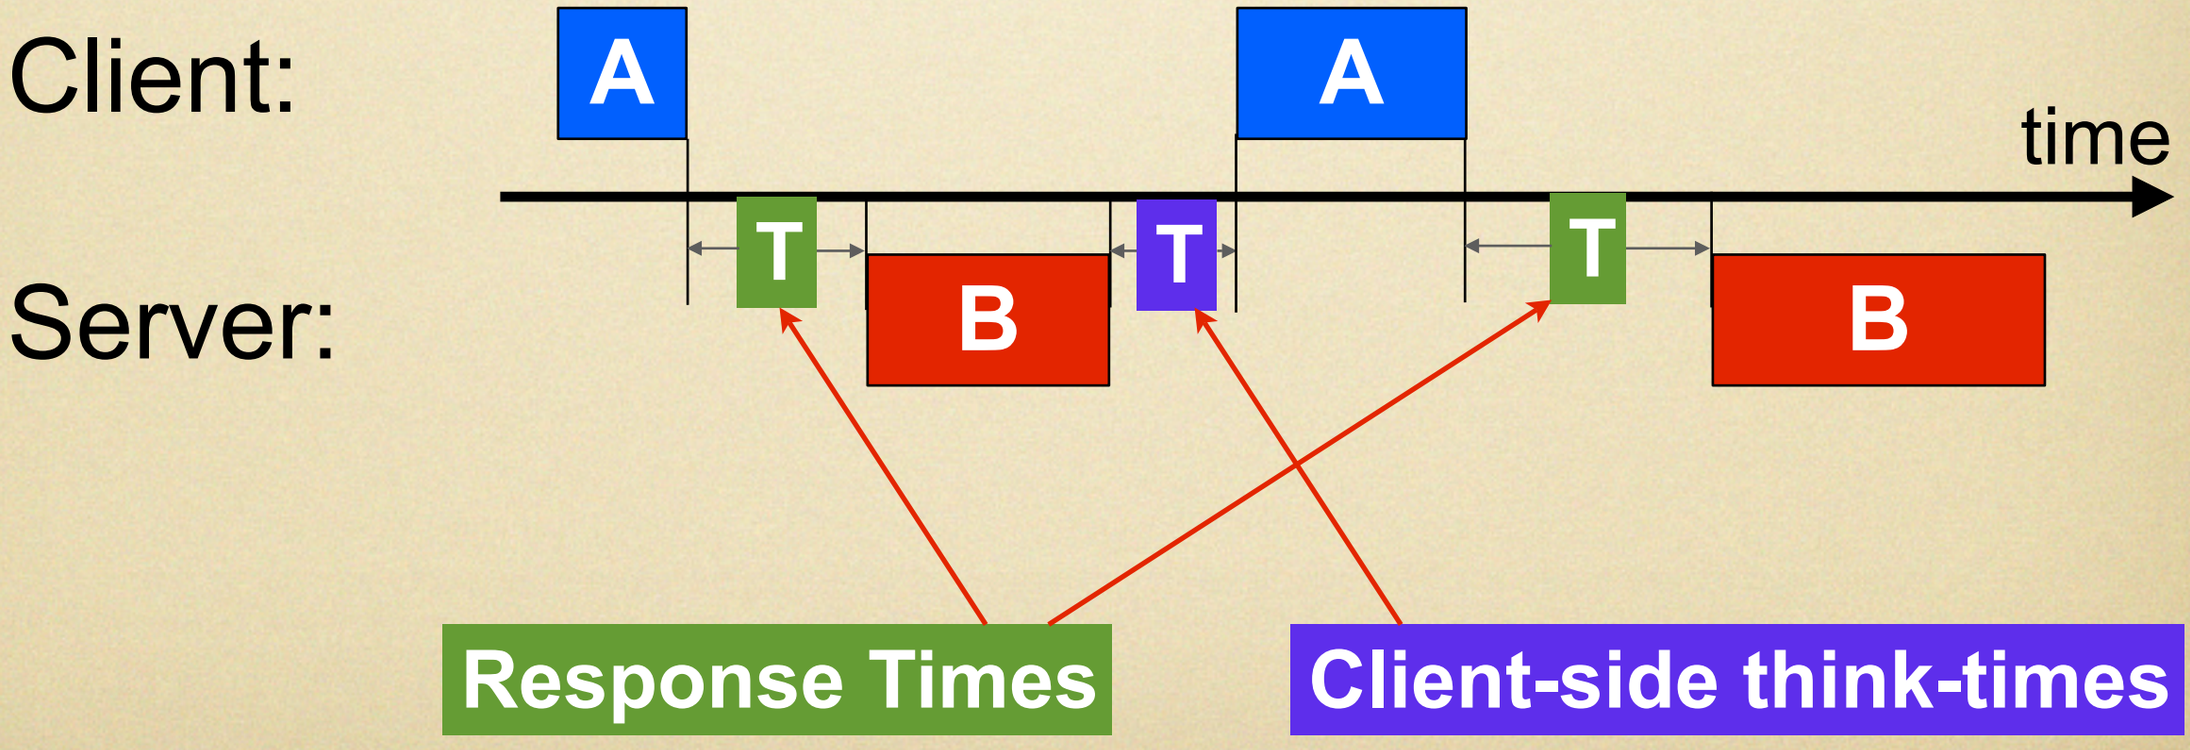

Constructing connection vectors

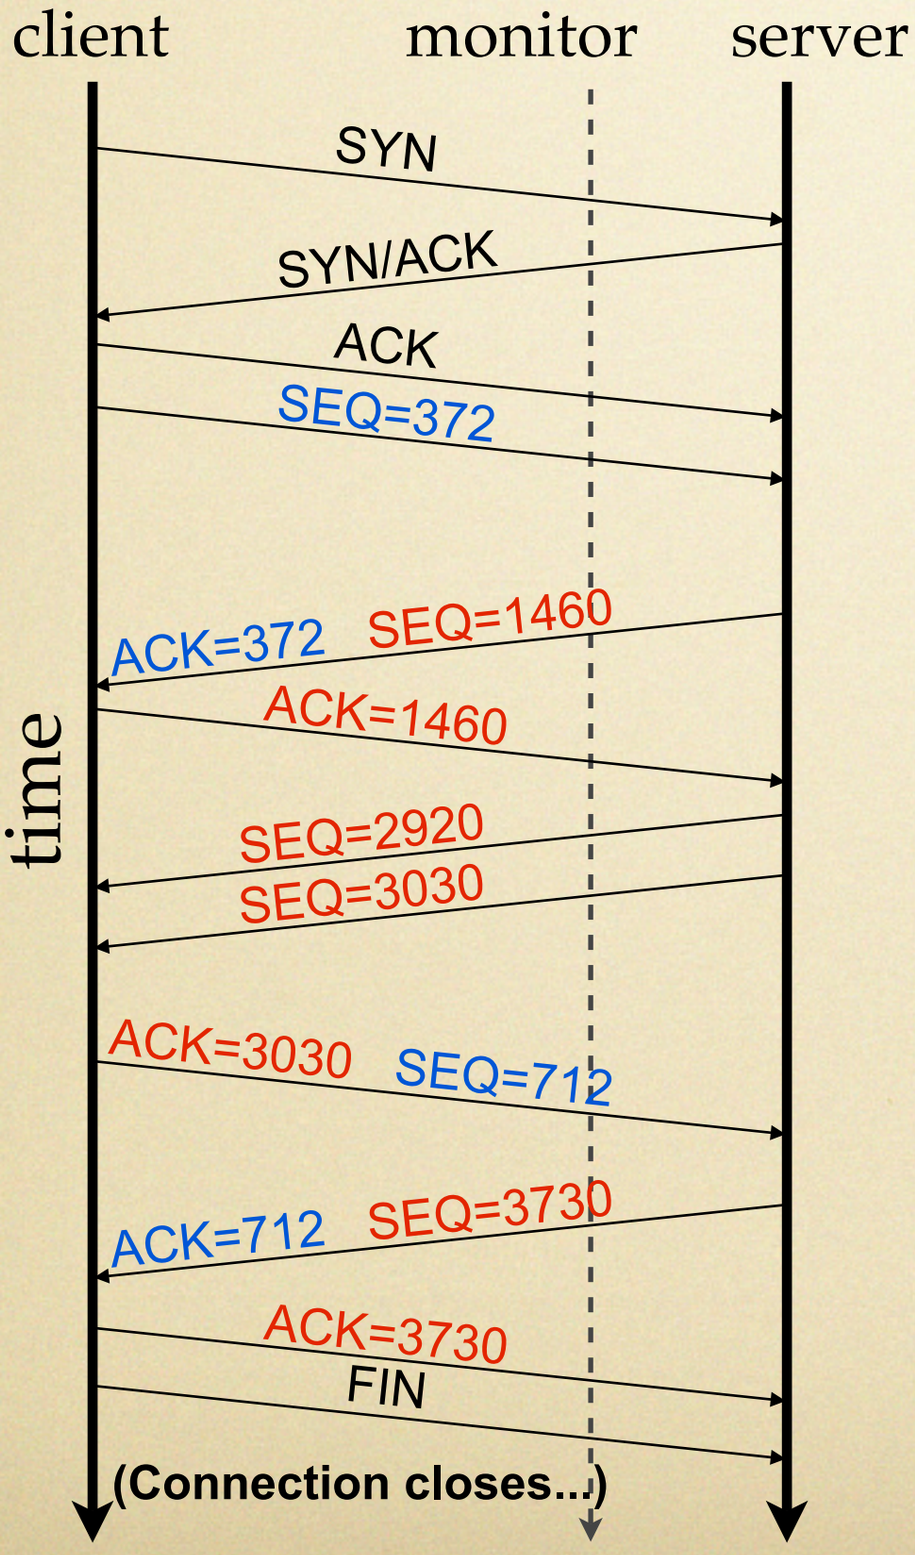

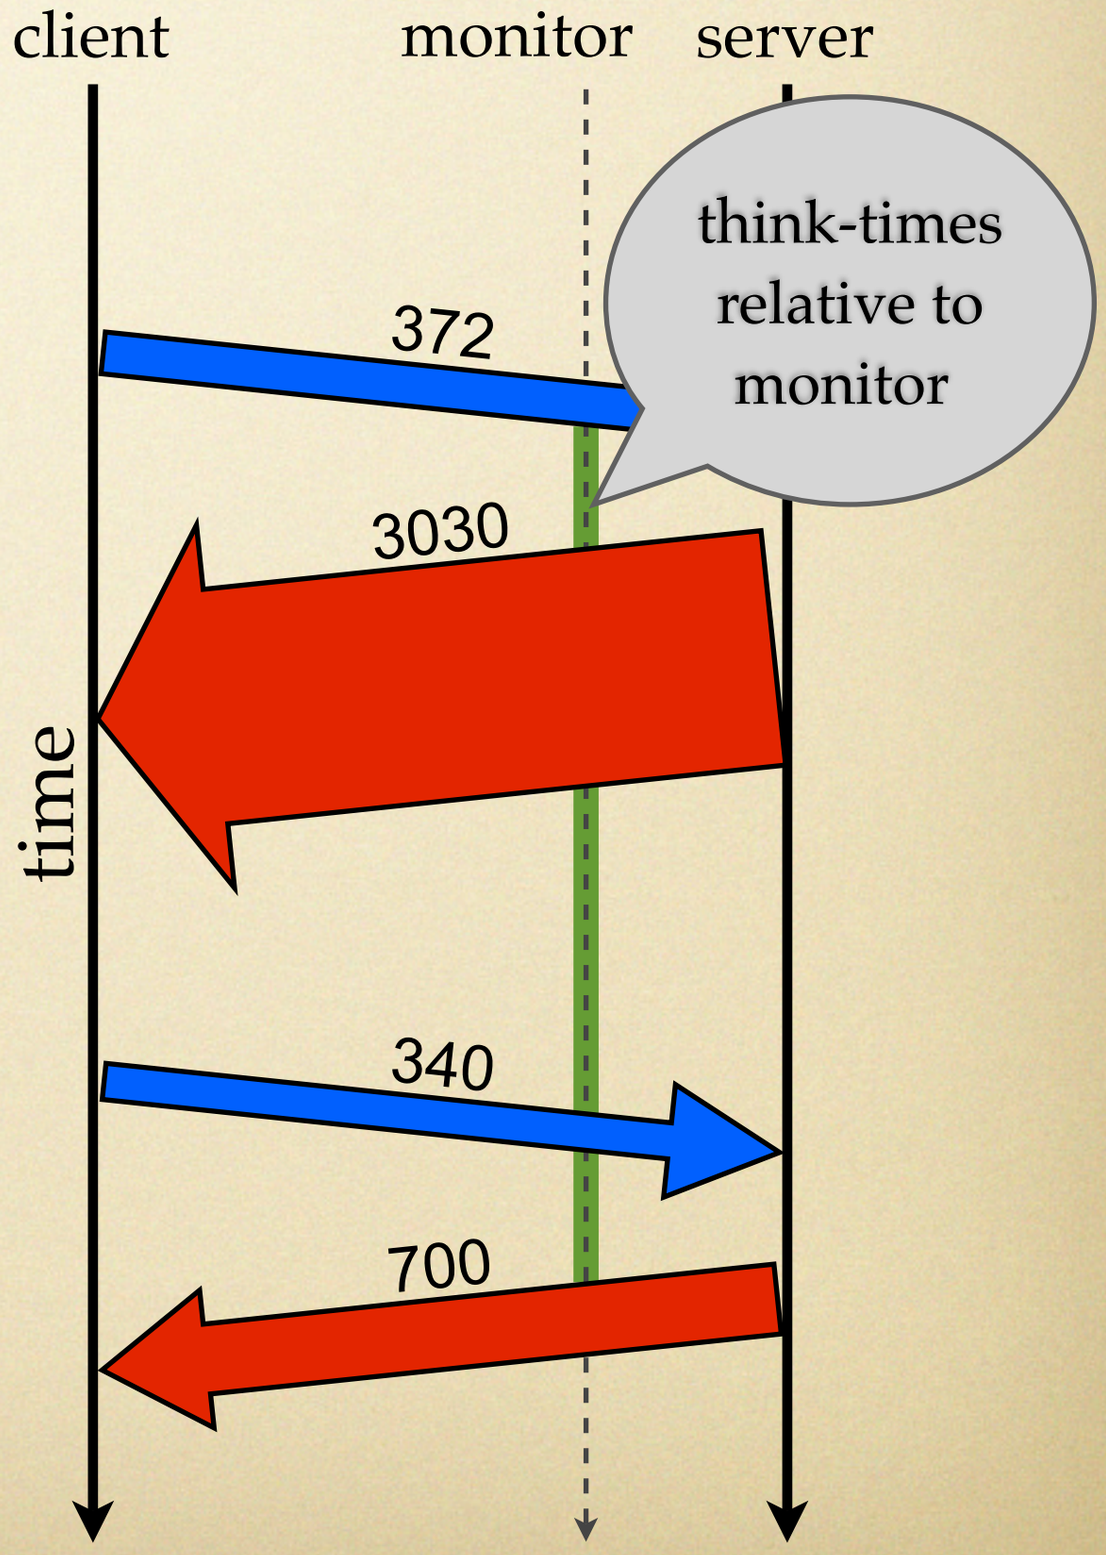

Needed Measurements

- *Application-level* measurements from TCP/IP headers:
  - server **response time**
  - **count** of application-level requests/responses
    - per server (i.e. **server load**)
    - per connection (i.e. **dialog length**)
  - **size** of application-level requests/responses
  - connection **duration**

Viability of Netflow

- What can Netflow provide?
  - server response time **- No**
  - count of *application-level* requests/responses **- No**
    - per server (i.e. server load) **- sort of**
    - per connection (i.e. dialog length) **- No**
  - size of *application-level* requests/responses **- No**
  - connection duration **- sort of**

Previous approach

- Previous work by Felix Hernandez-Campos on building connection vectors.

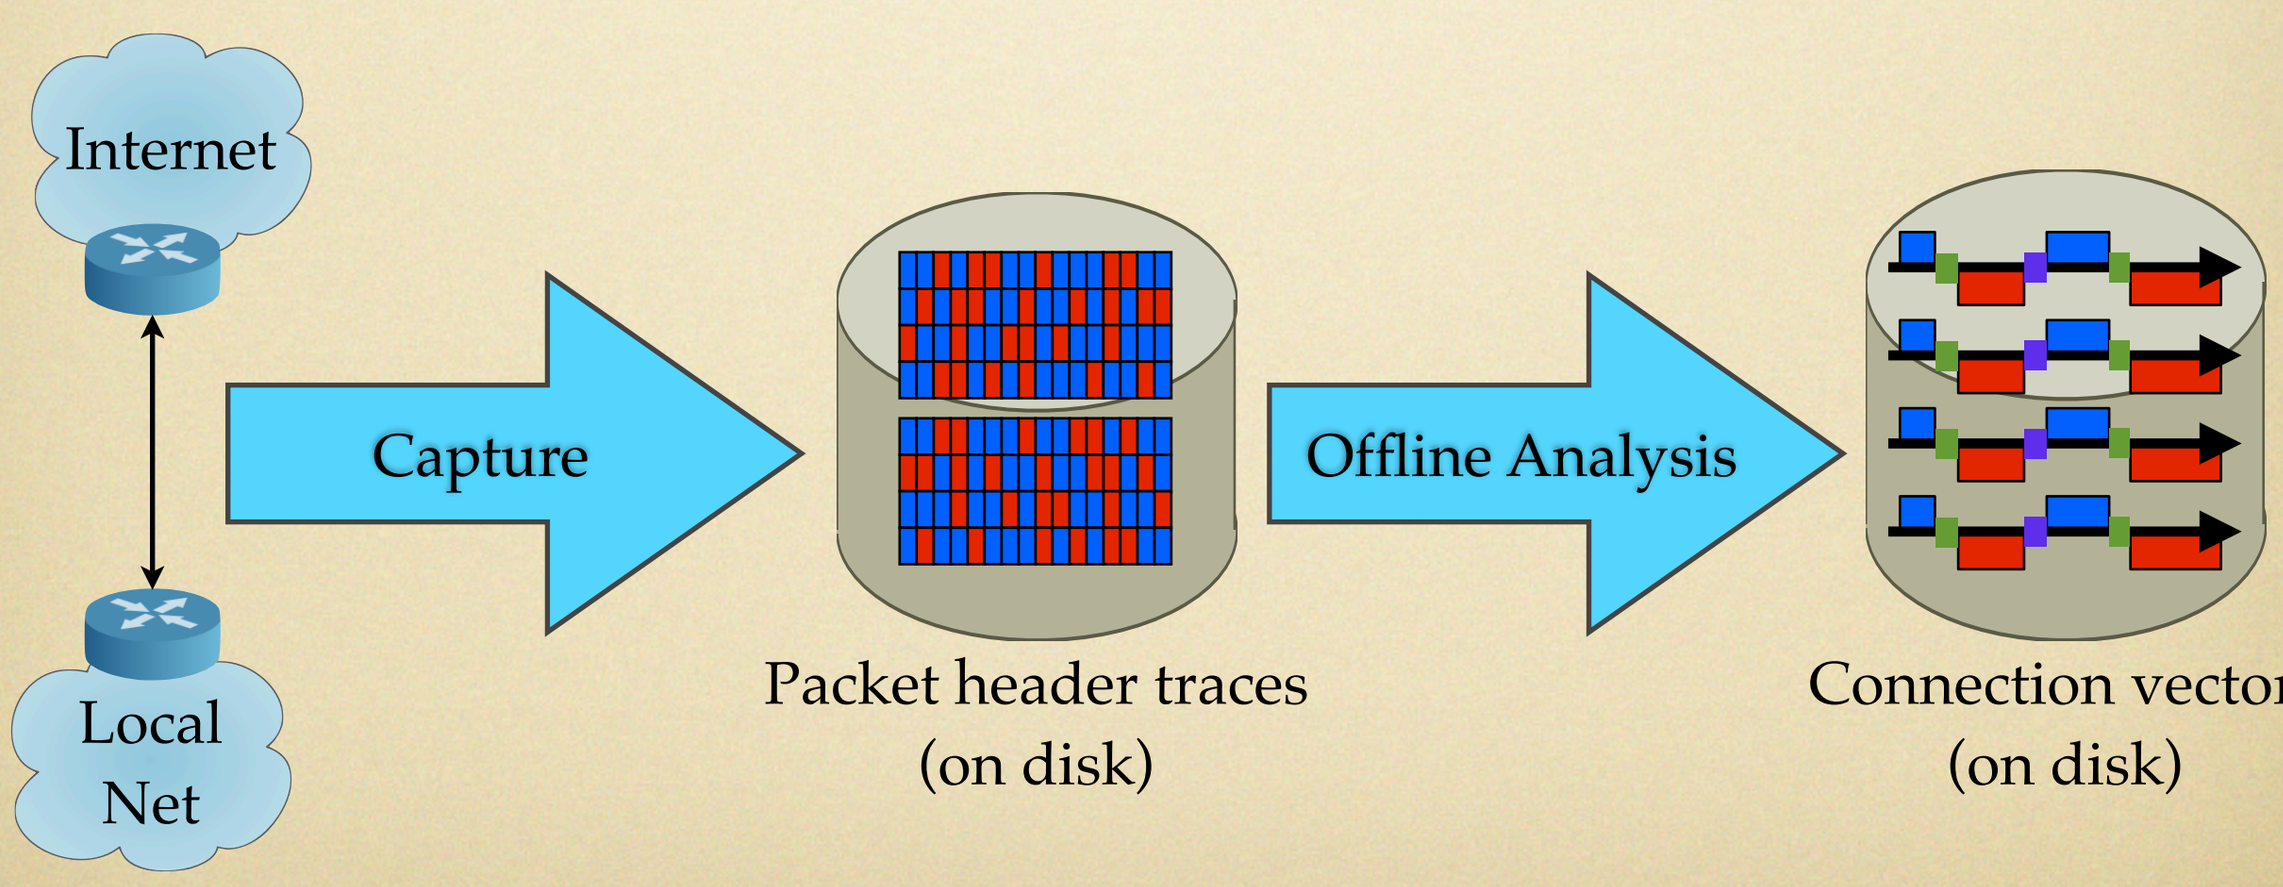

Our approach

- Our innovation: build connection vectors **online**, with a **single pass**.

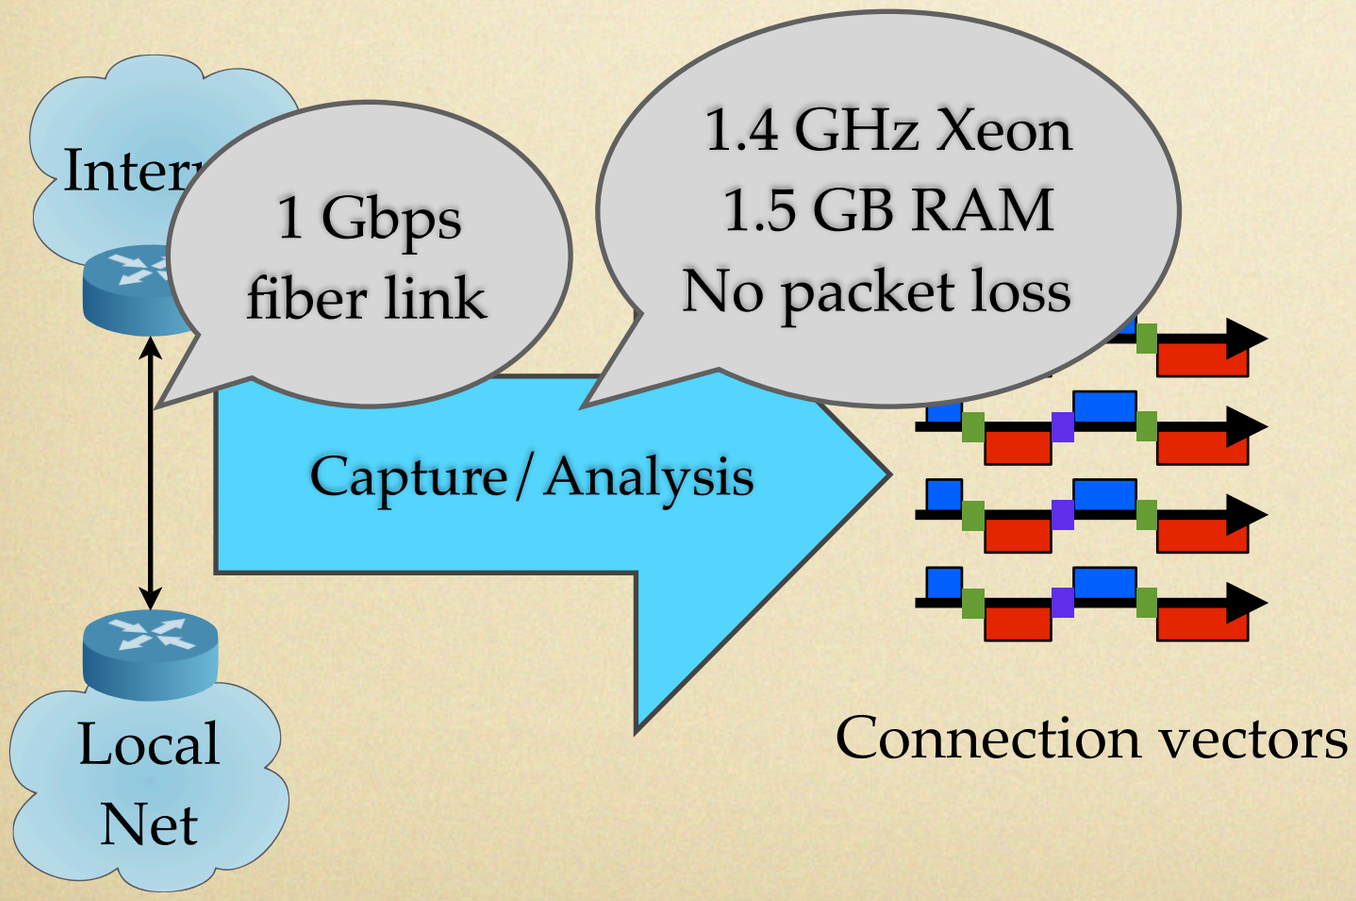

- Now, no intermediate files
- Capability for **continuous** measurement
- Elements of connection vectors available **immediately**
- Capability for online understanding of server performance

adudump

- The tool we wrote to do this is called **adudump**.
- Here's the output of **adudump** for an example connection:

```
TYPE  TIMESTAMP          LOCAL_HOST   DIR  REMOTE_HOST      OTHER_INFO
SYN: 1202706002.650917 190.40.1.180.443 < 221.151.95.184.62015
RTT:  1202706002.651967 190.40.1.180.443 > 221.151.95.184.62015 0.001050
SEQ: 1202706002.681395 190.40.1.180.443 < 221.151.95.184.62015
ADU: 1202706002.688748 190.40.1.180.443 < 221.151.95.184.62015 163 SEQ 0.000542
ADU: 1202706002.733813 190.40.1.180.443 > 221.151.95.184.62015 2886 SEQ 0.045041
ADU: 1202706002.738254 190.40.1.180.443 < 221.151.95.184.62015 198 SEQ 0.004441
ADU: 1202706002.801408 190.40.1.180.443 > 221.151.95.184.62015 59 SEQ
END: 1202706002.821701 190.40.1.180.443 < 221.151.95.184.62015
```

computing all kinds of things in real-time...contextual information as well as ADUs...

Data

<u>For this paper:</u>

- 66 days
- 1.54 TB (uncompressed)
- 16.8 billion requests and responses
- 1.6 billion connections

<u>Overall:</u>

- 180 days
- 3.35 TB (uncompressed)
- 34.8 billion requests and responses
- 4.0 billion connections

Case study: the incident

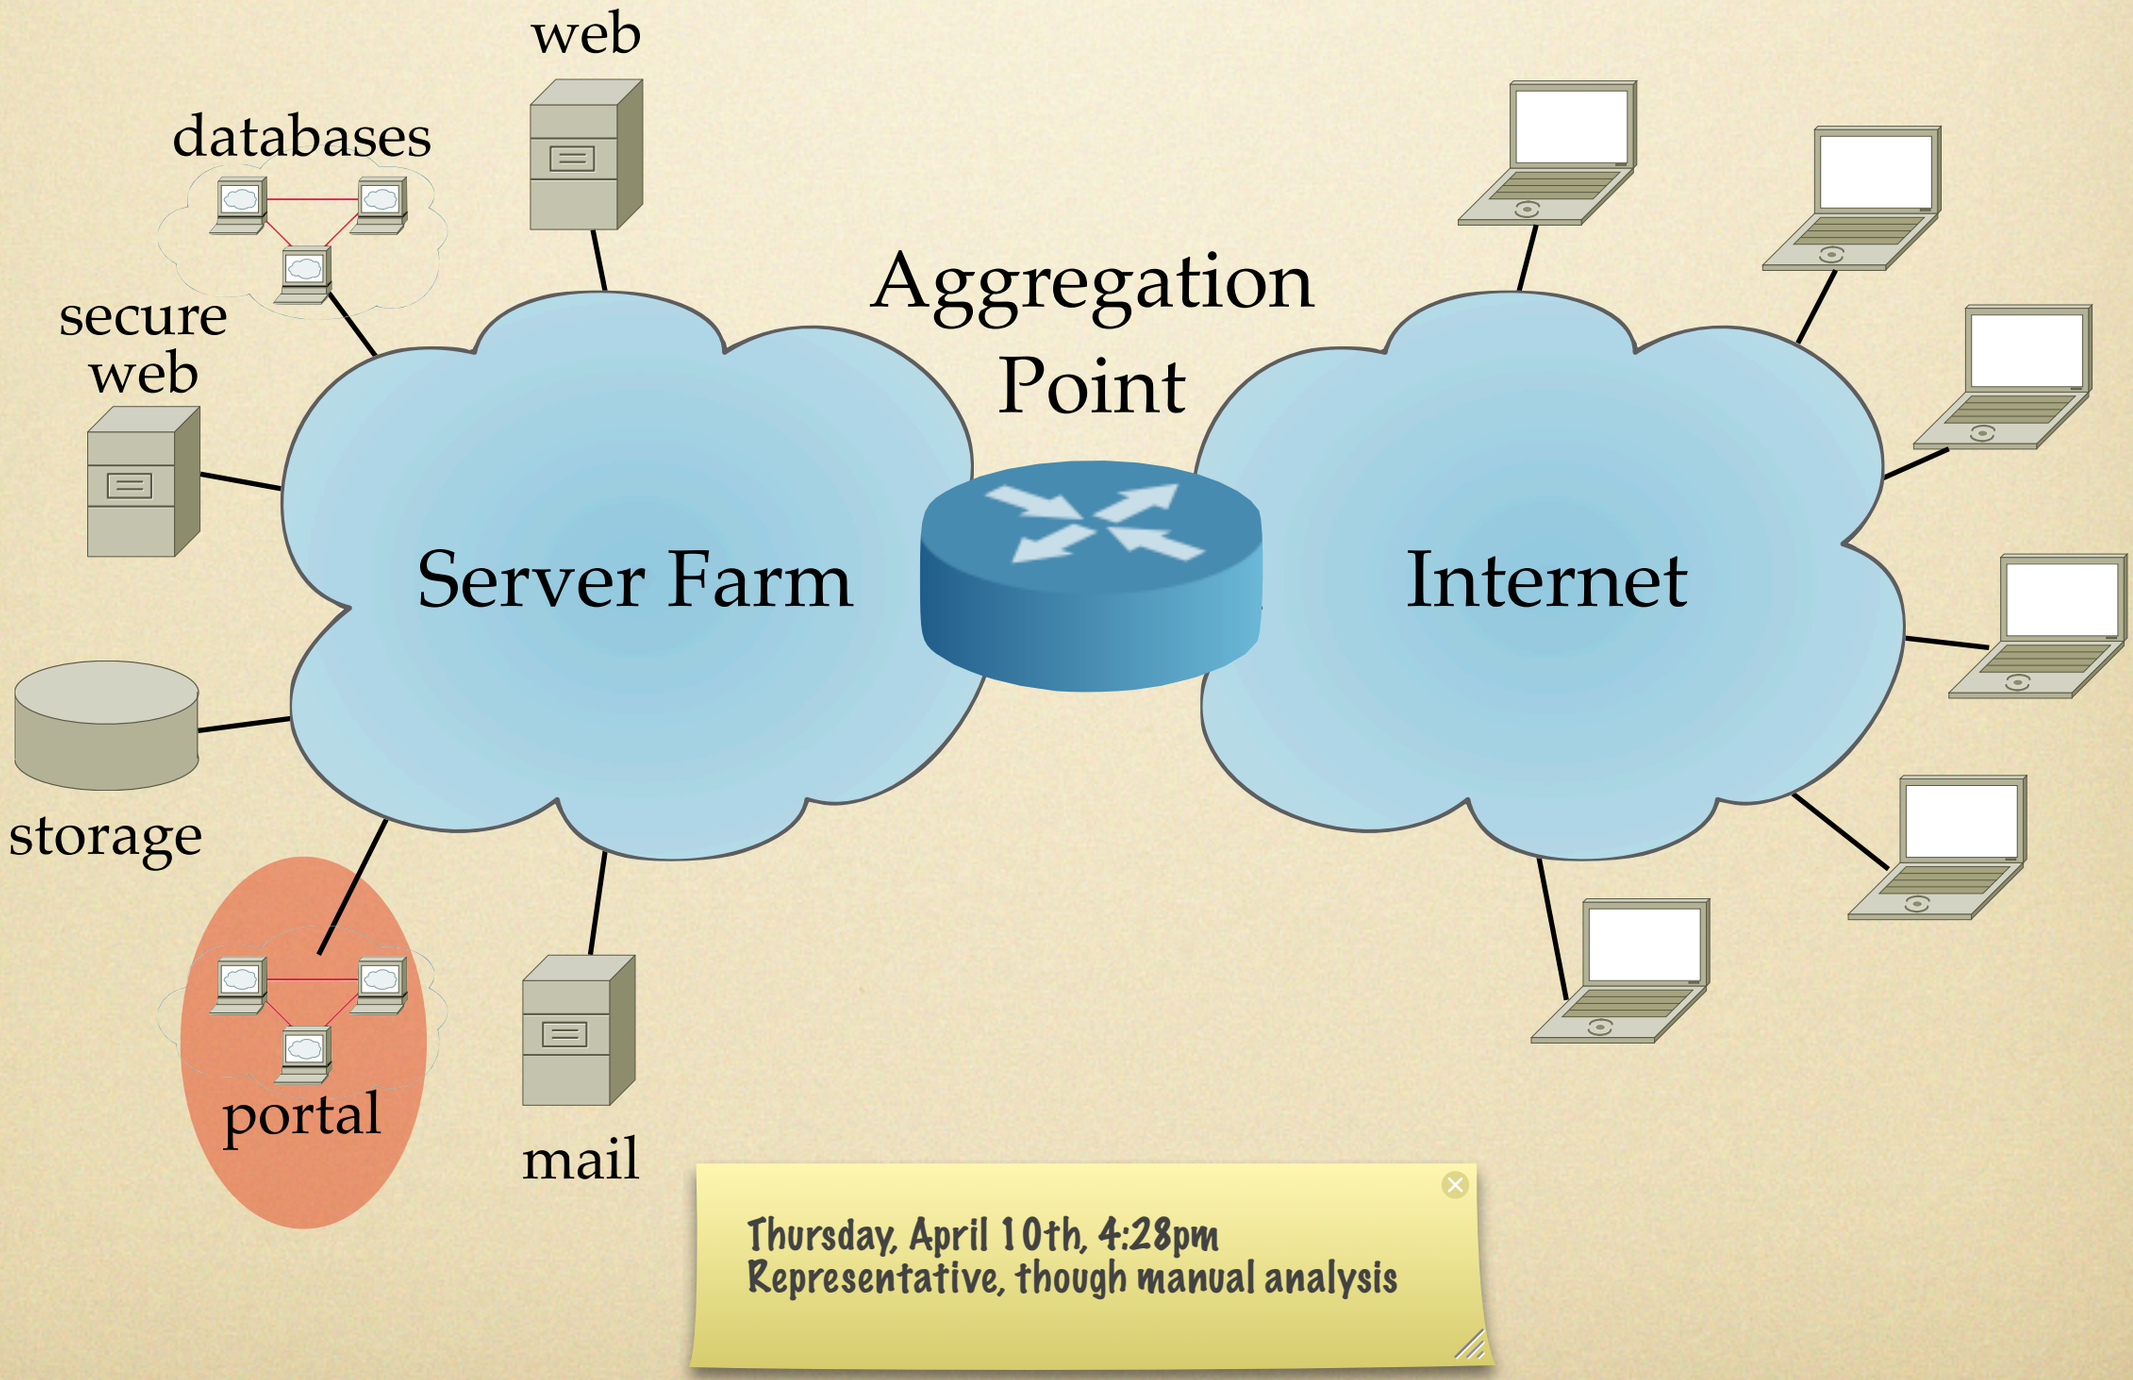

Thursday, April 10th, 4:28pm
Representative, though manual analysis

Case study: the issue

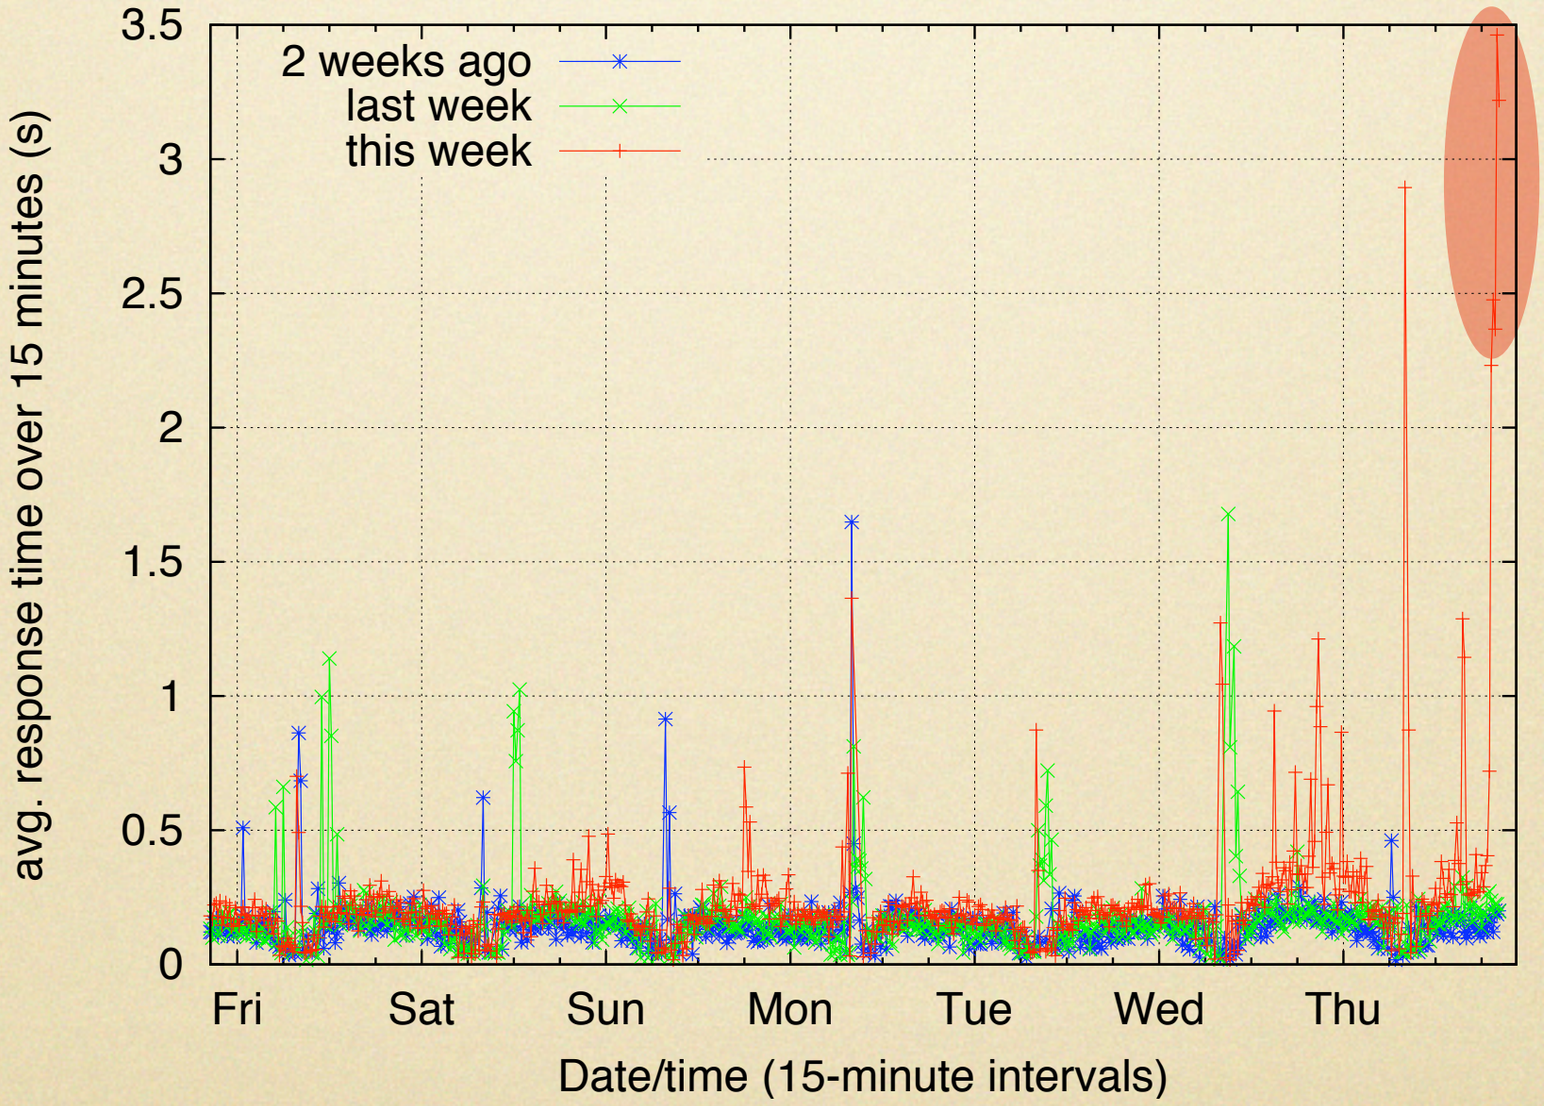

Case study: the issue

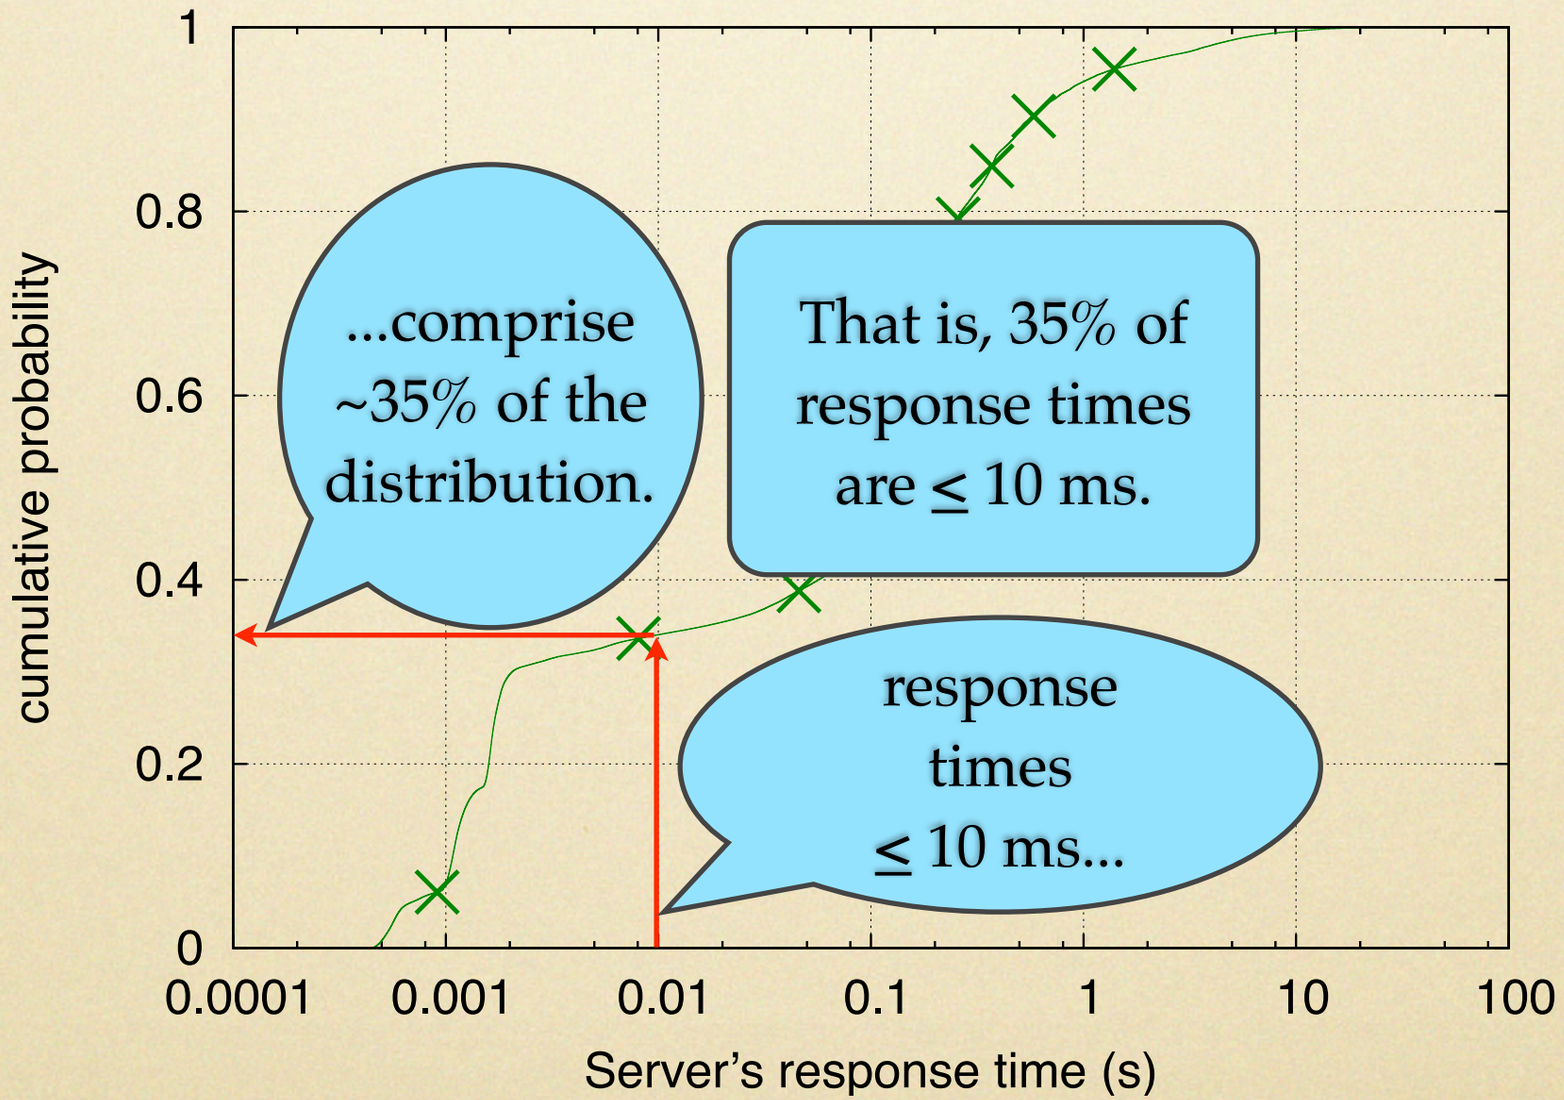

Case study: the issue

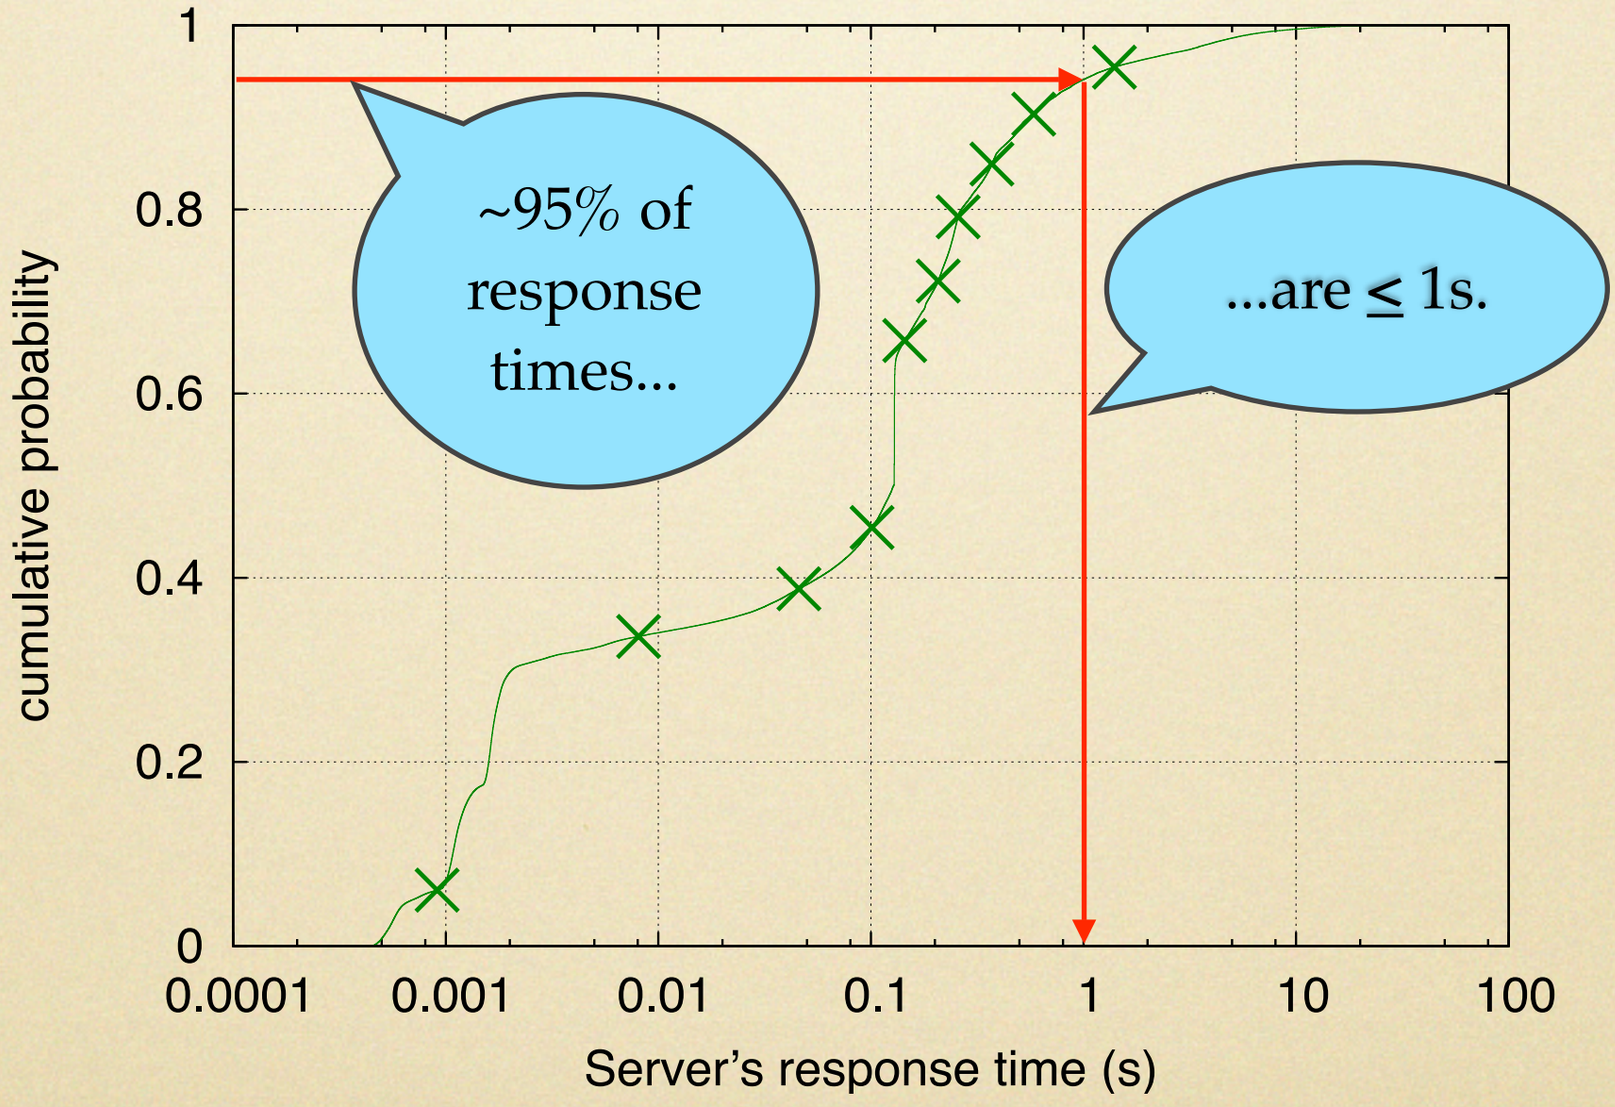

Case study: the issue

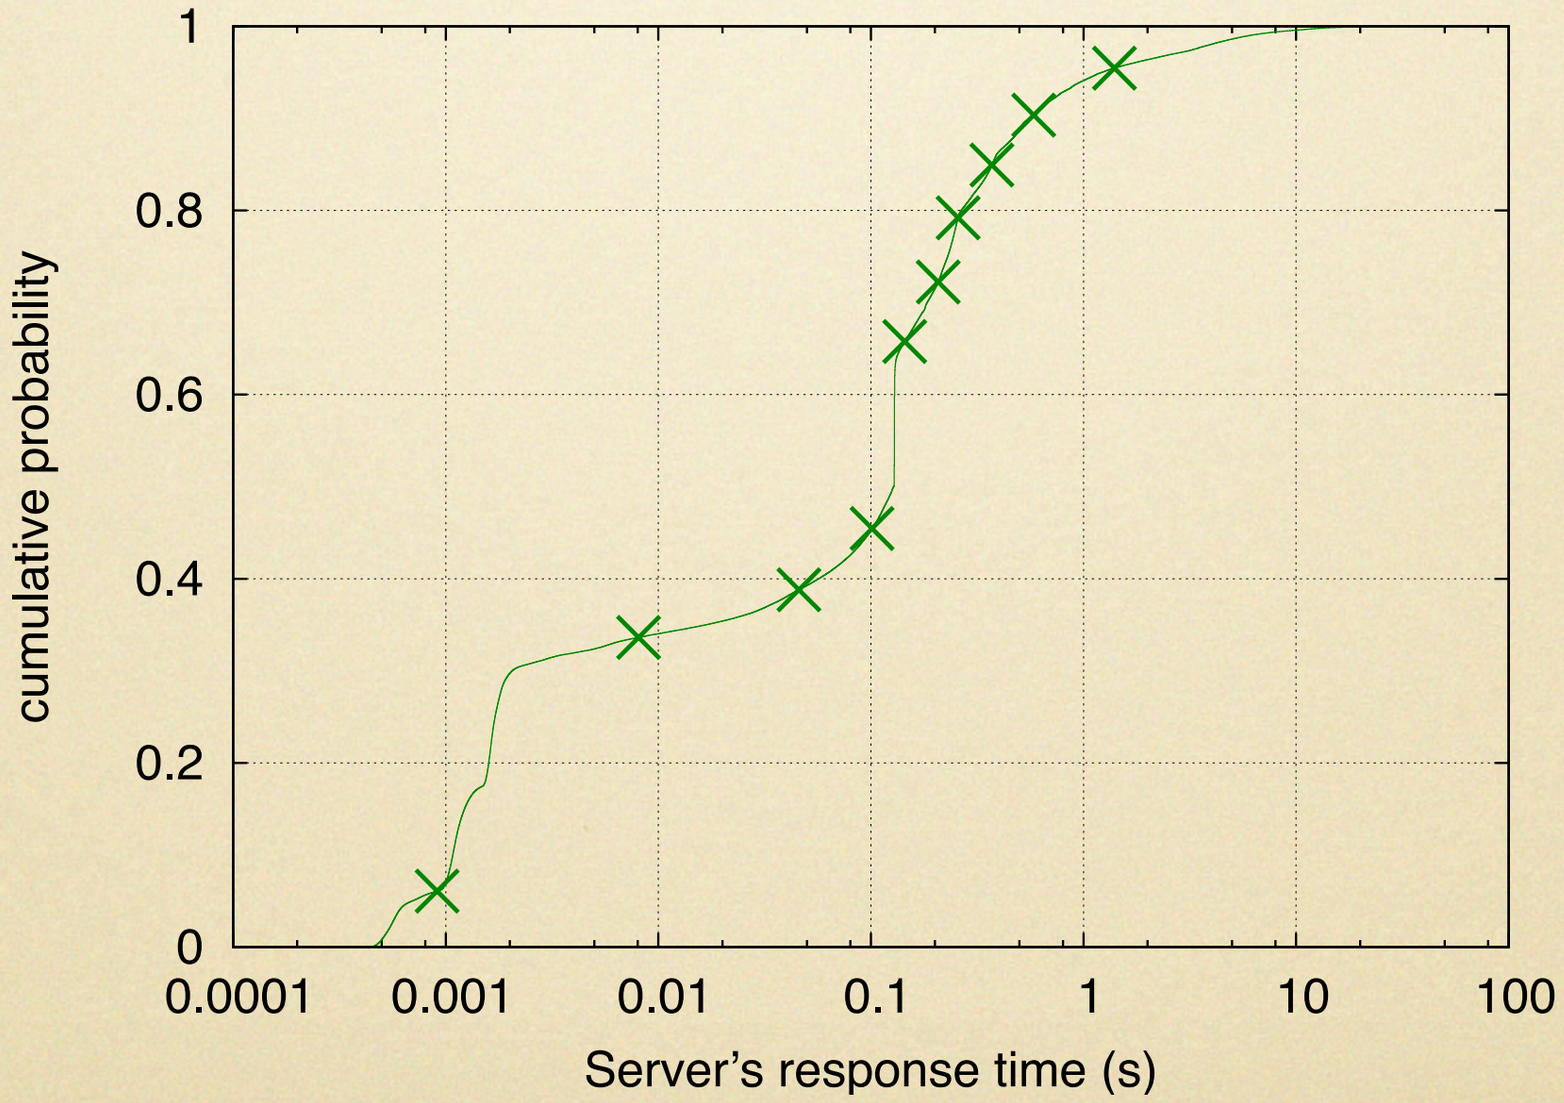

Case study: the issue

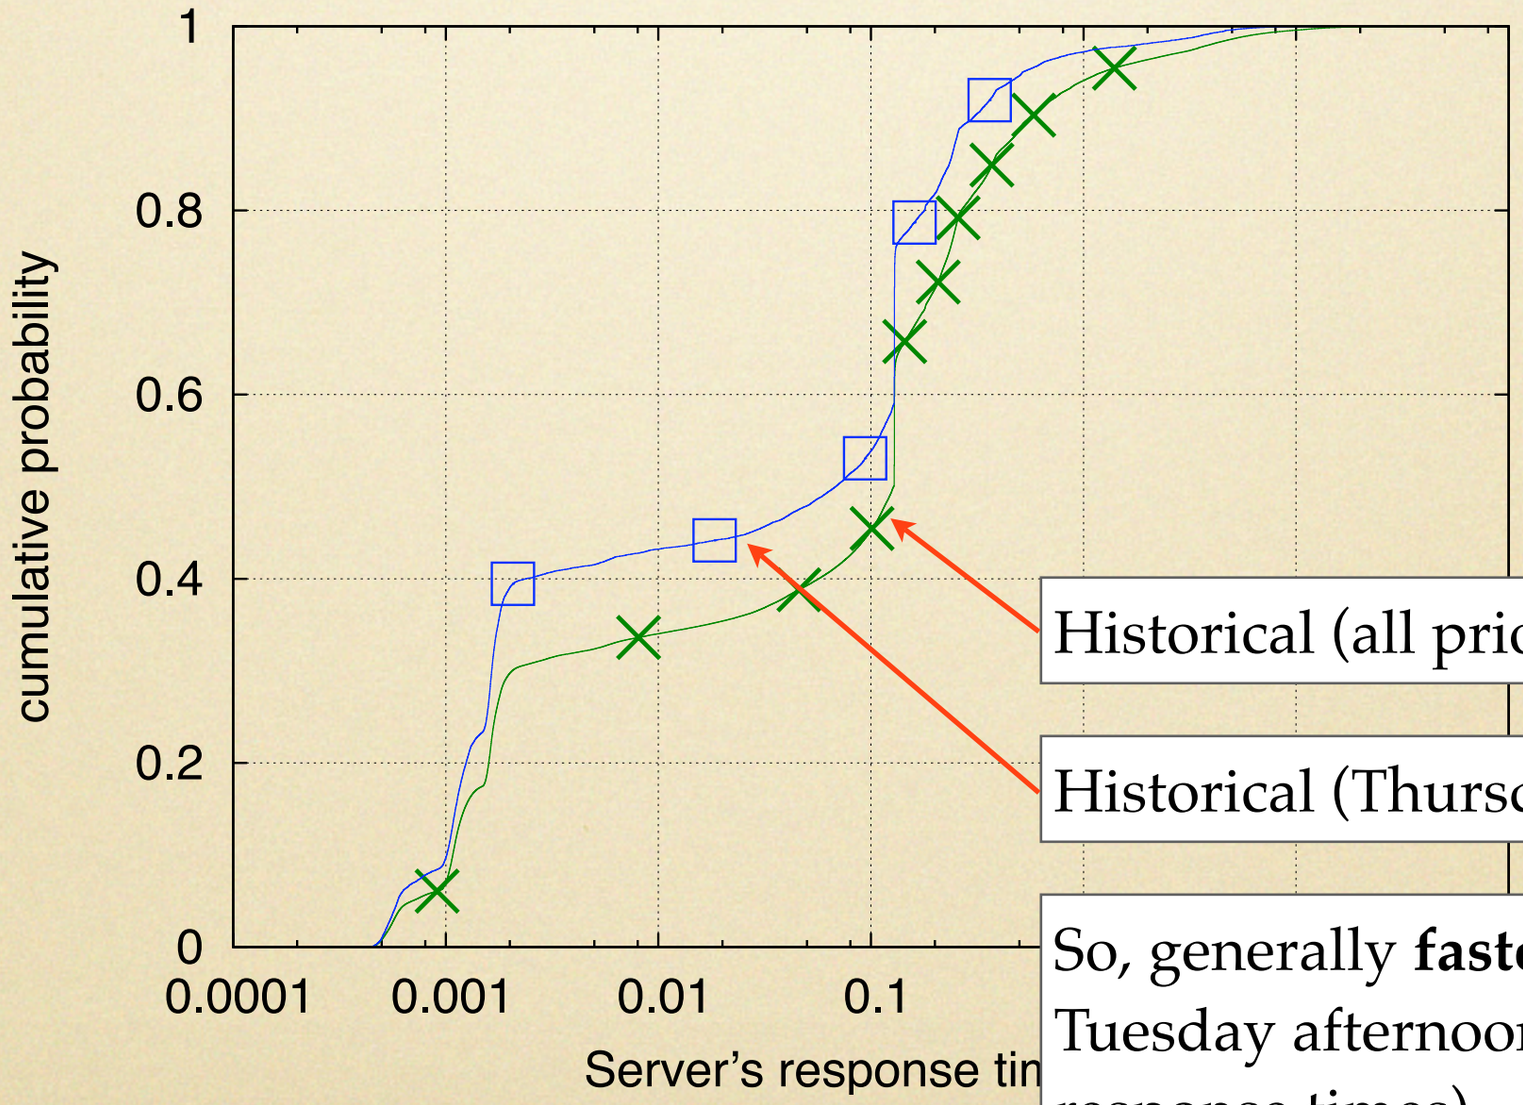

Case study: the issue

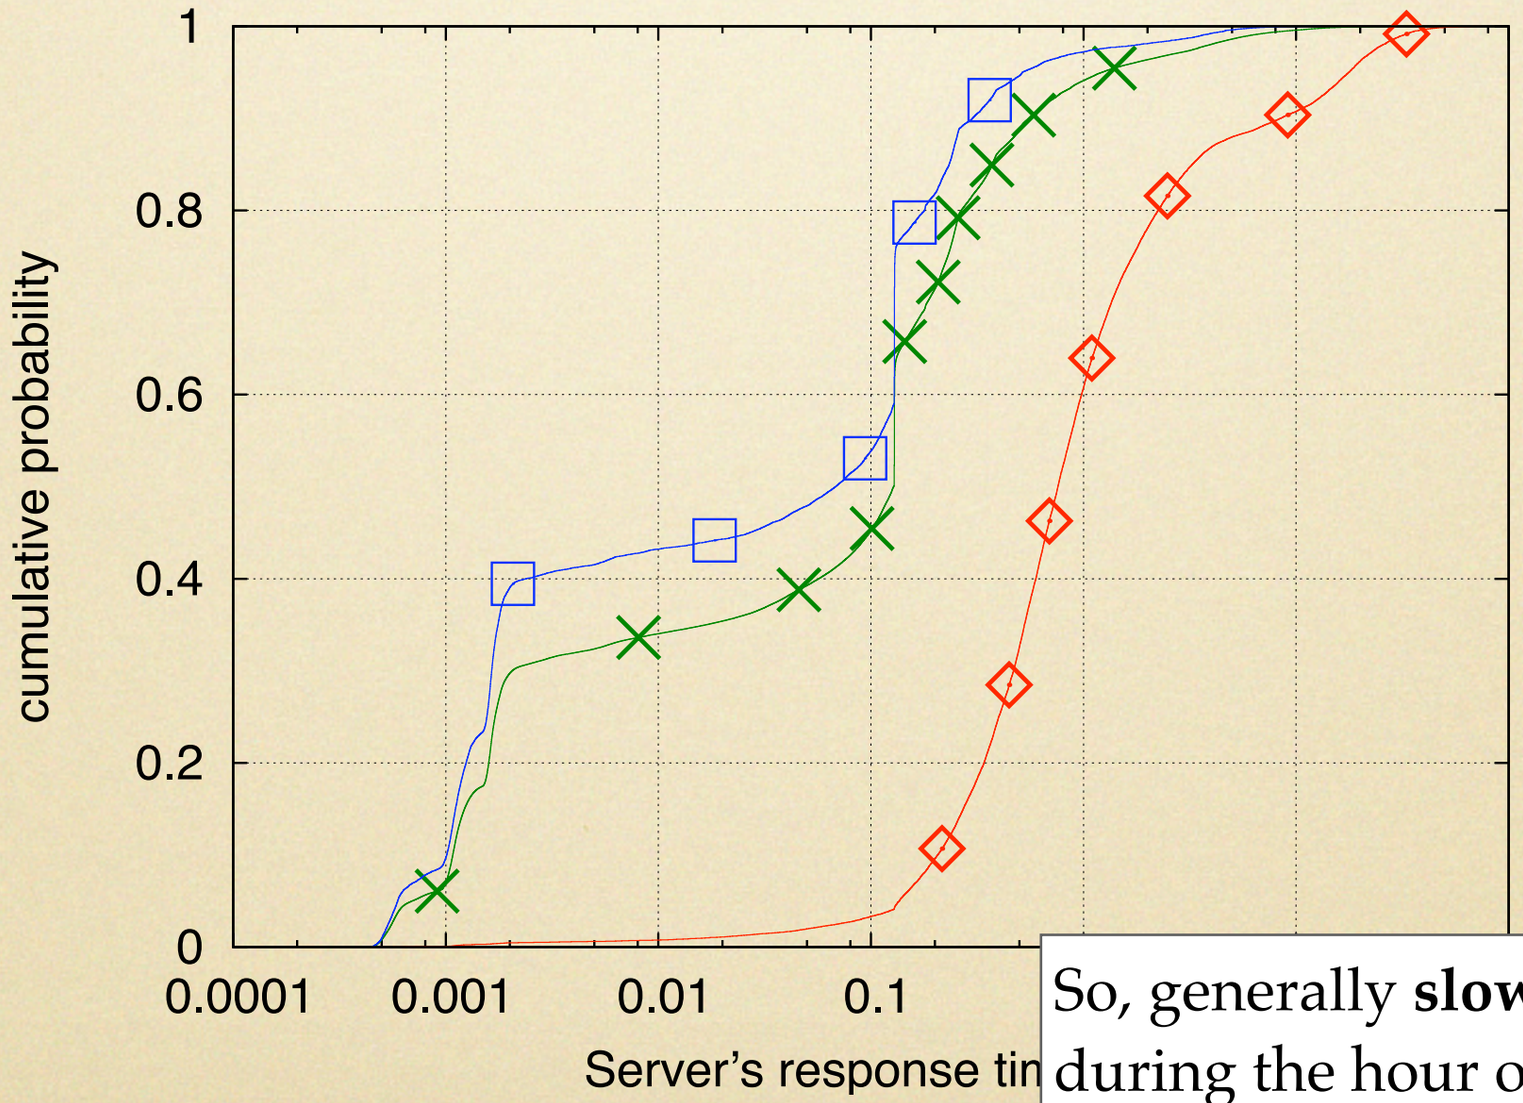

So, generally **slower** responses during the hour of the incident.

Case study: investigation

- What could cause this incident?
  - Larger requests (more processing required)

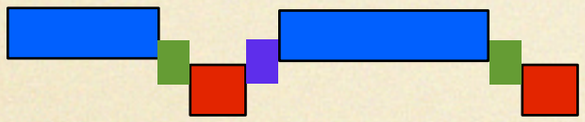

  - Larger responses (implying more processing)

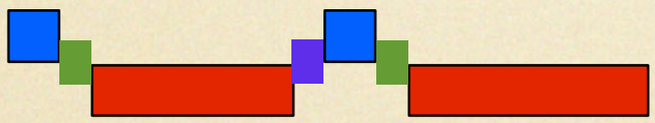

  - More requests per connection (more work)

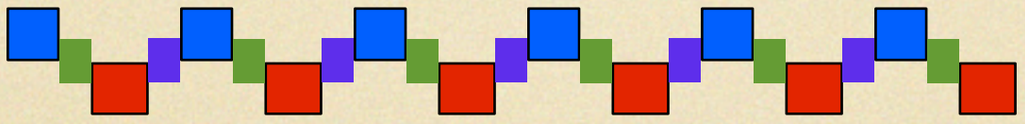

  - More requests per time unit (more work)

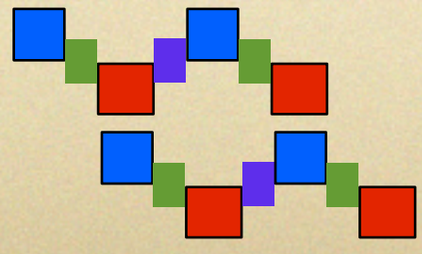

Case study: investigation

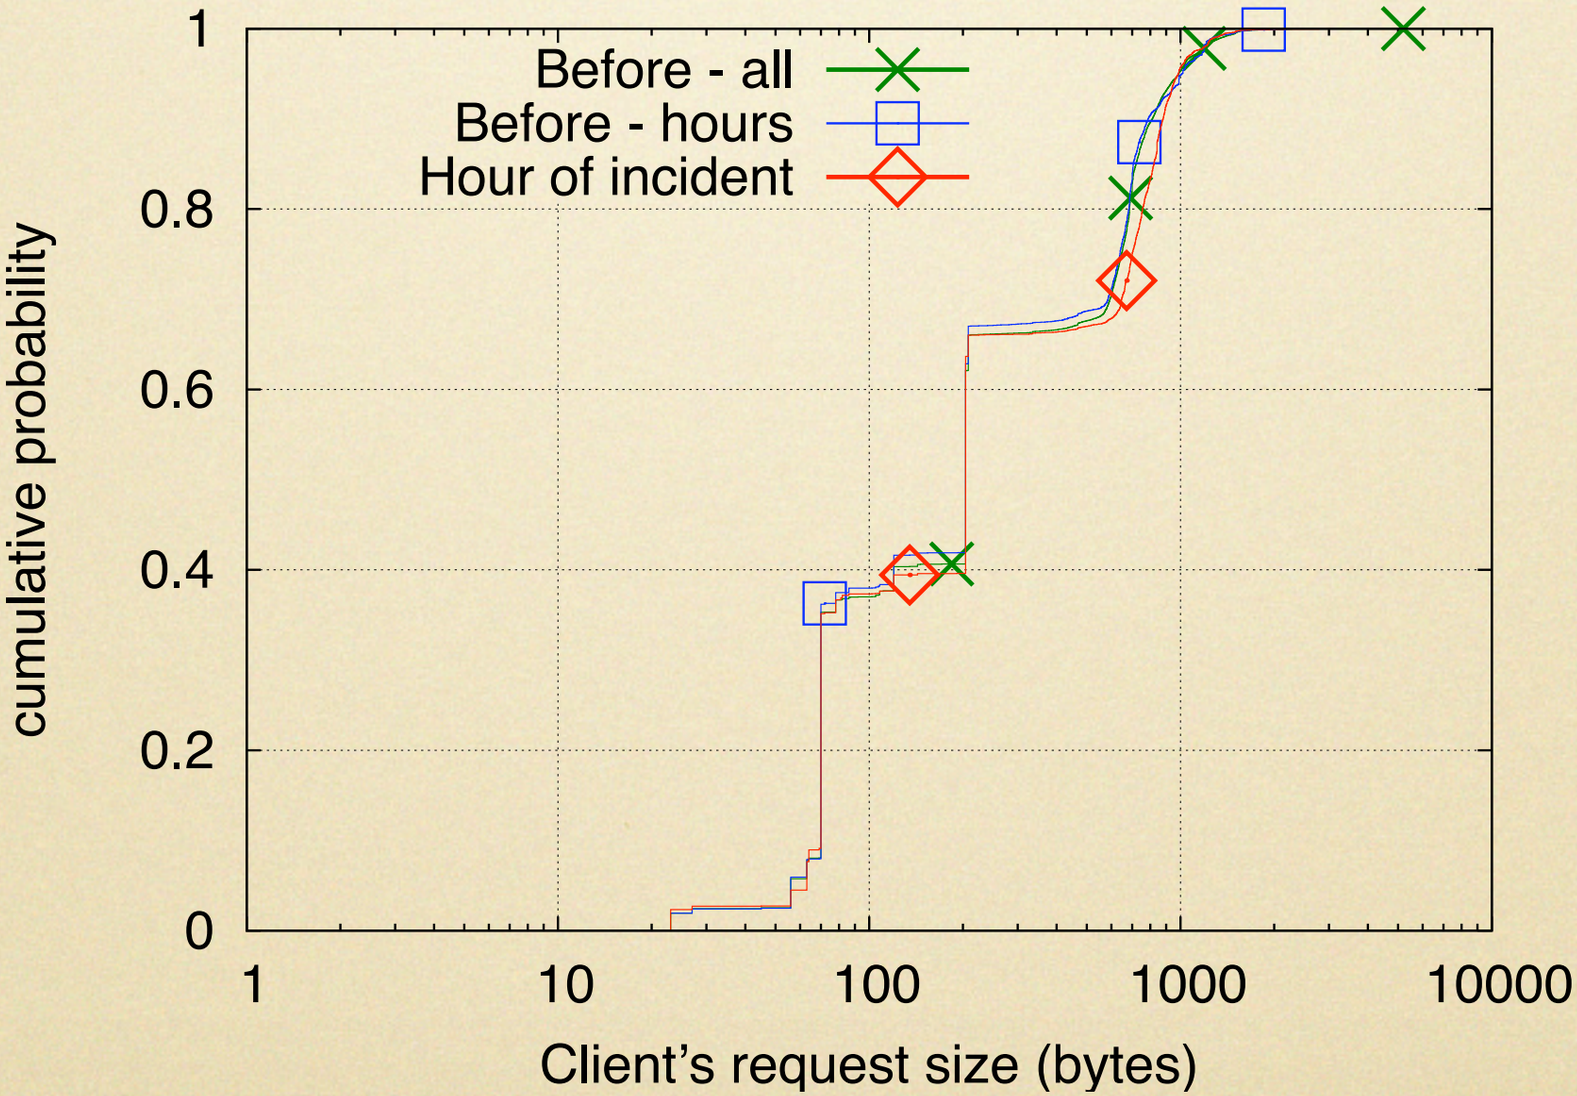

Case study: investigation

- What could cause this incident?
  - Larger requests (more processing required)

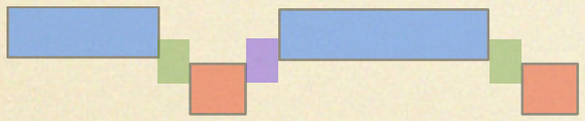

  - Larger responses (implying more processing)

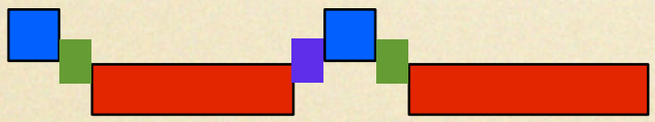

  - More requests per connection (more work)

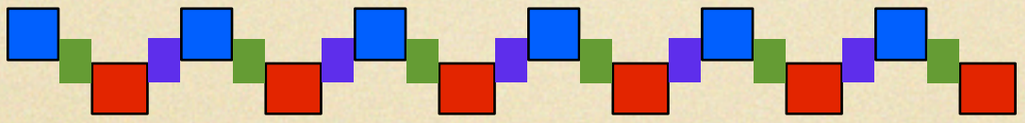

  - More requests per time unit (more work)

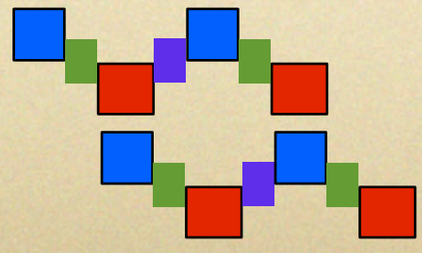

Case study: investigation

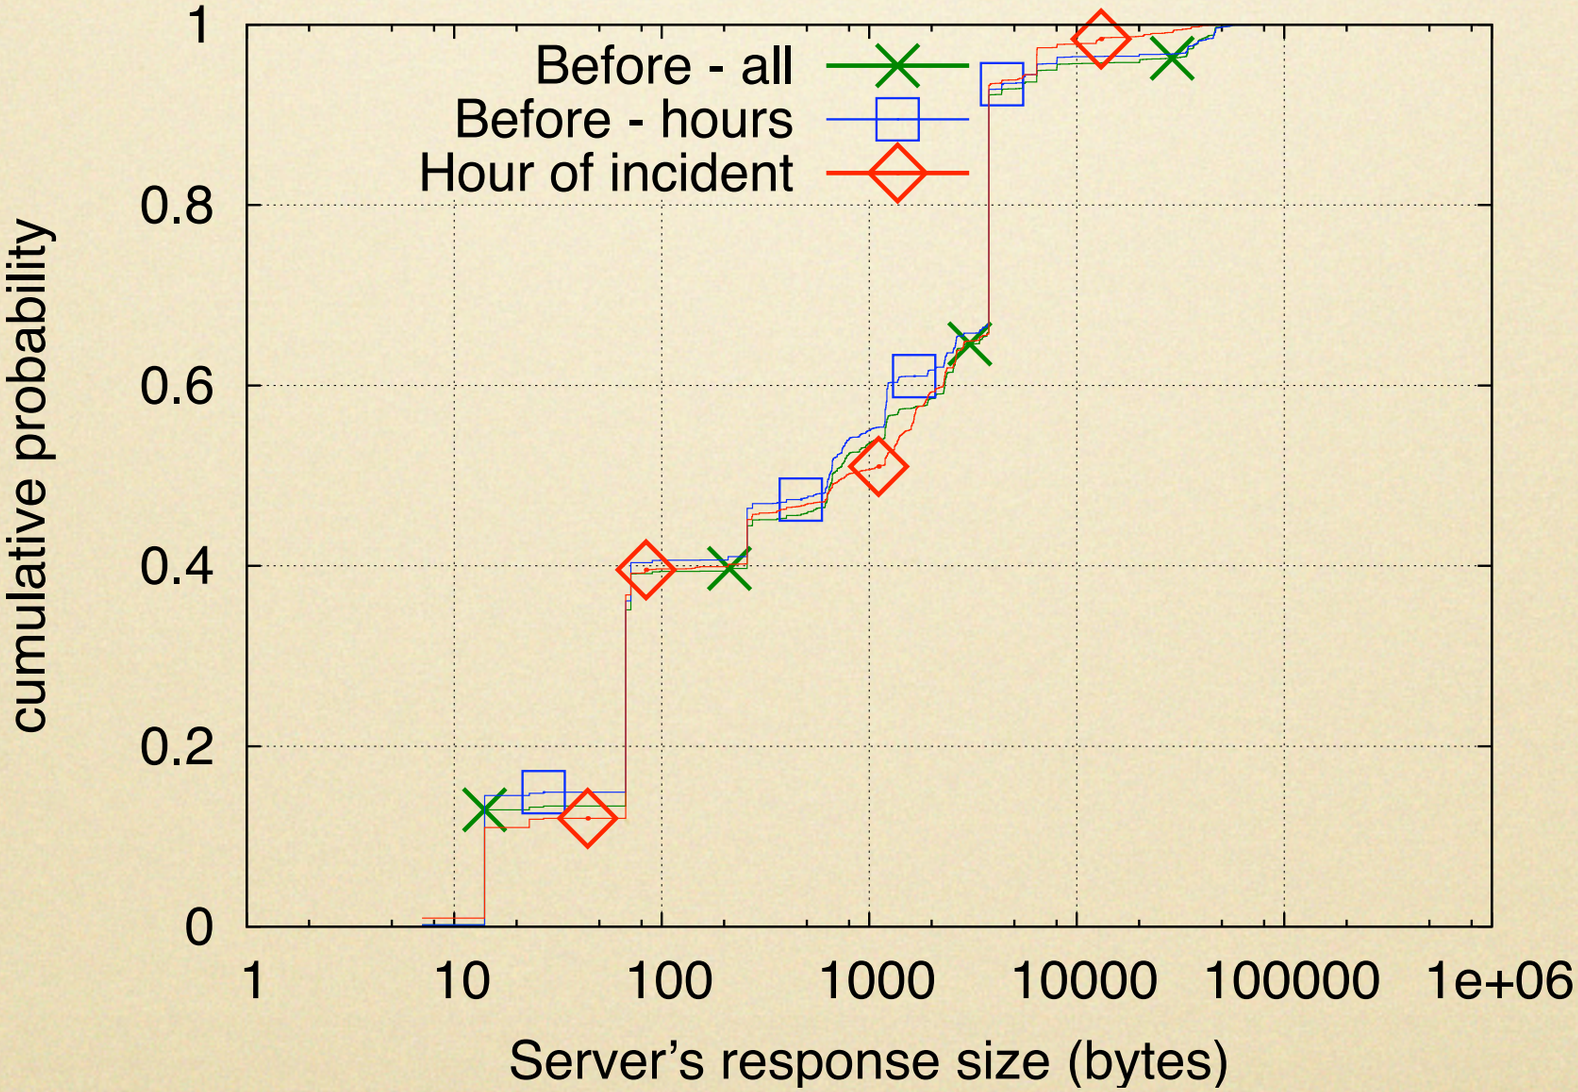

Case study: investigation

- What could cause this incident?
  - Larger requests (more processing required)

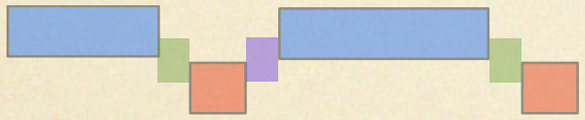

  - Larger responses (implying more processing)

- More requests per connection (more work)

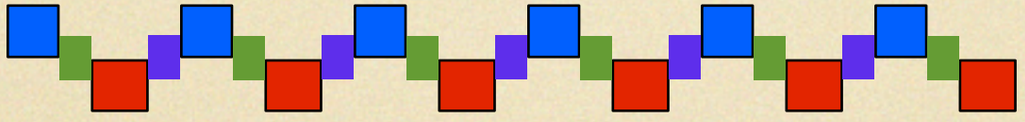

- More requests per time unit (more work)

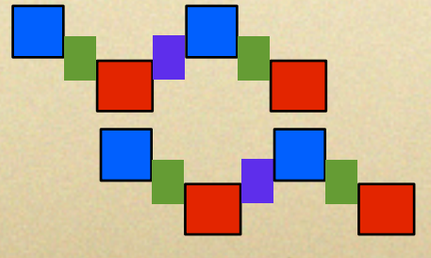

Case study: investigation

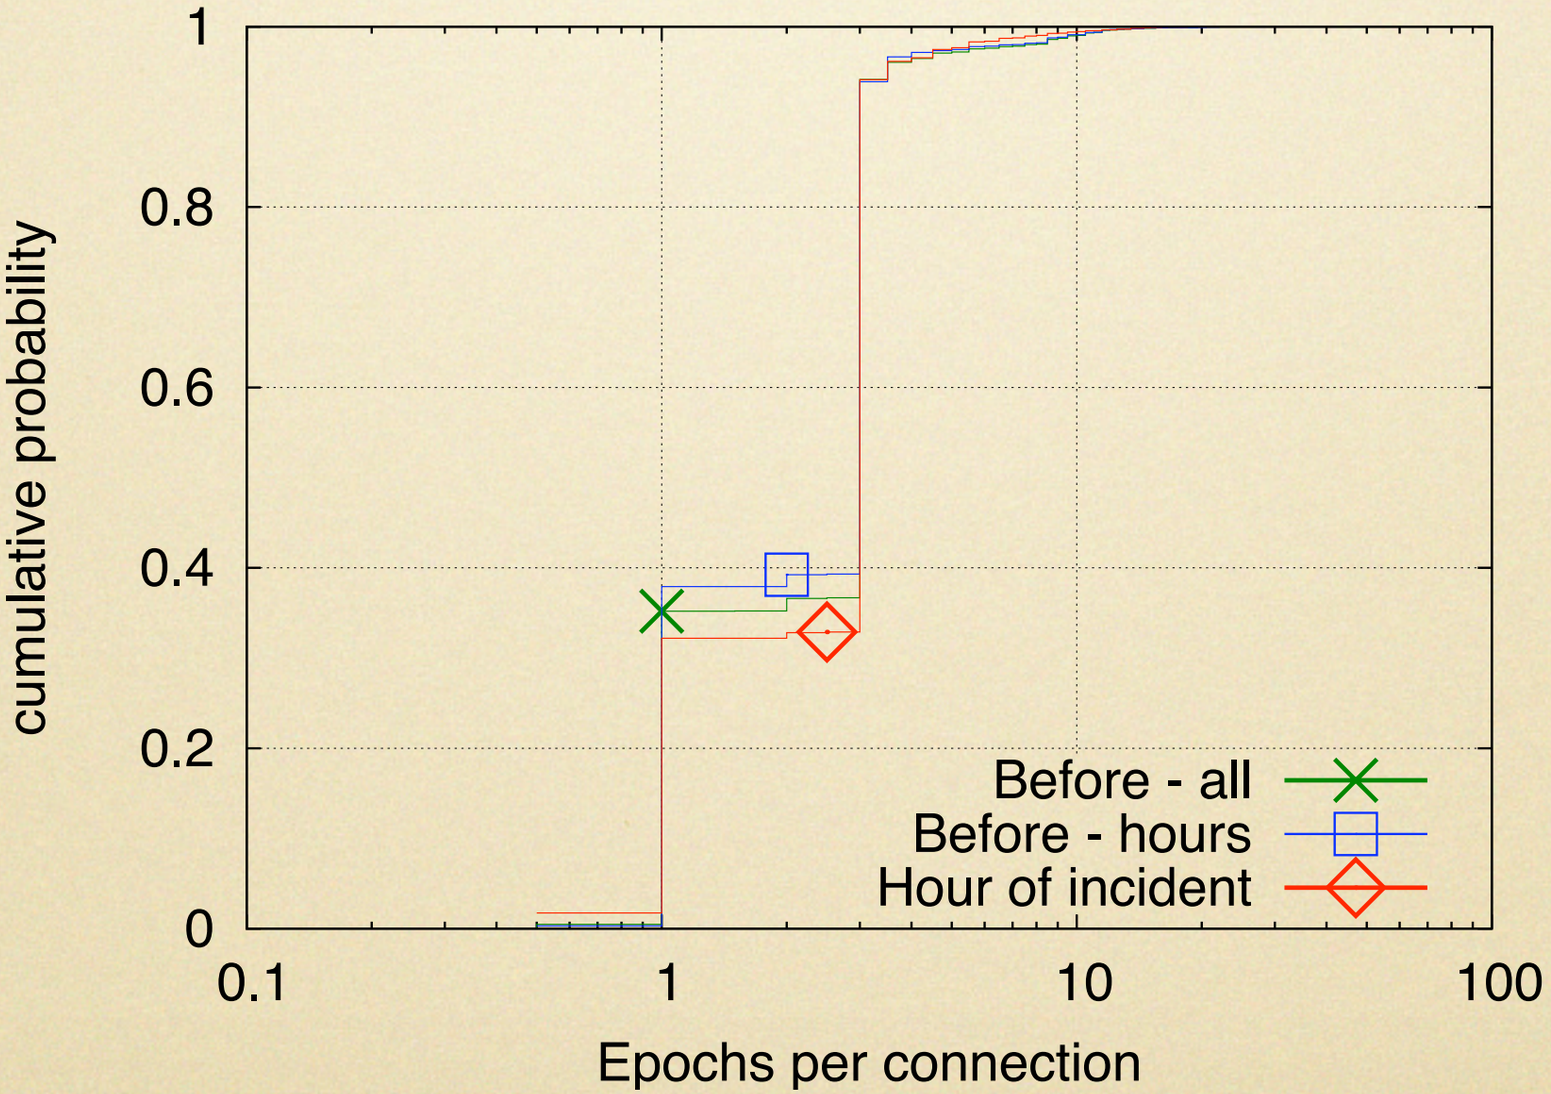

Case study: investigation

- What could cause this incident?
  - Larger requests (more processing required)

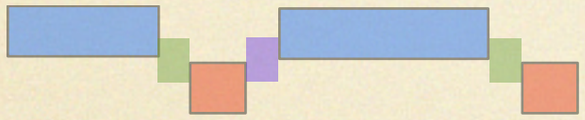

  - Larger responses (implying more processing)

  - More requests per connection (more work)

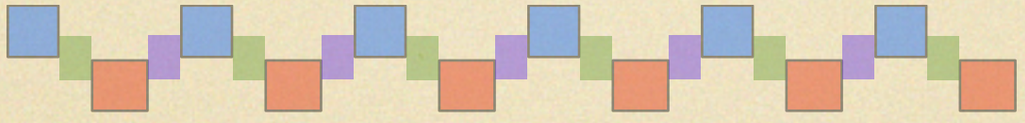

  - More requests per time unit (more work)

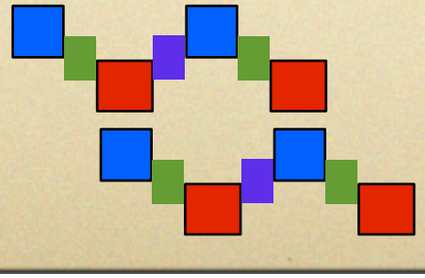

Case study: investigation

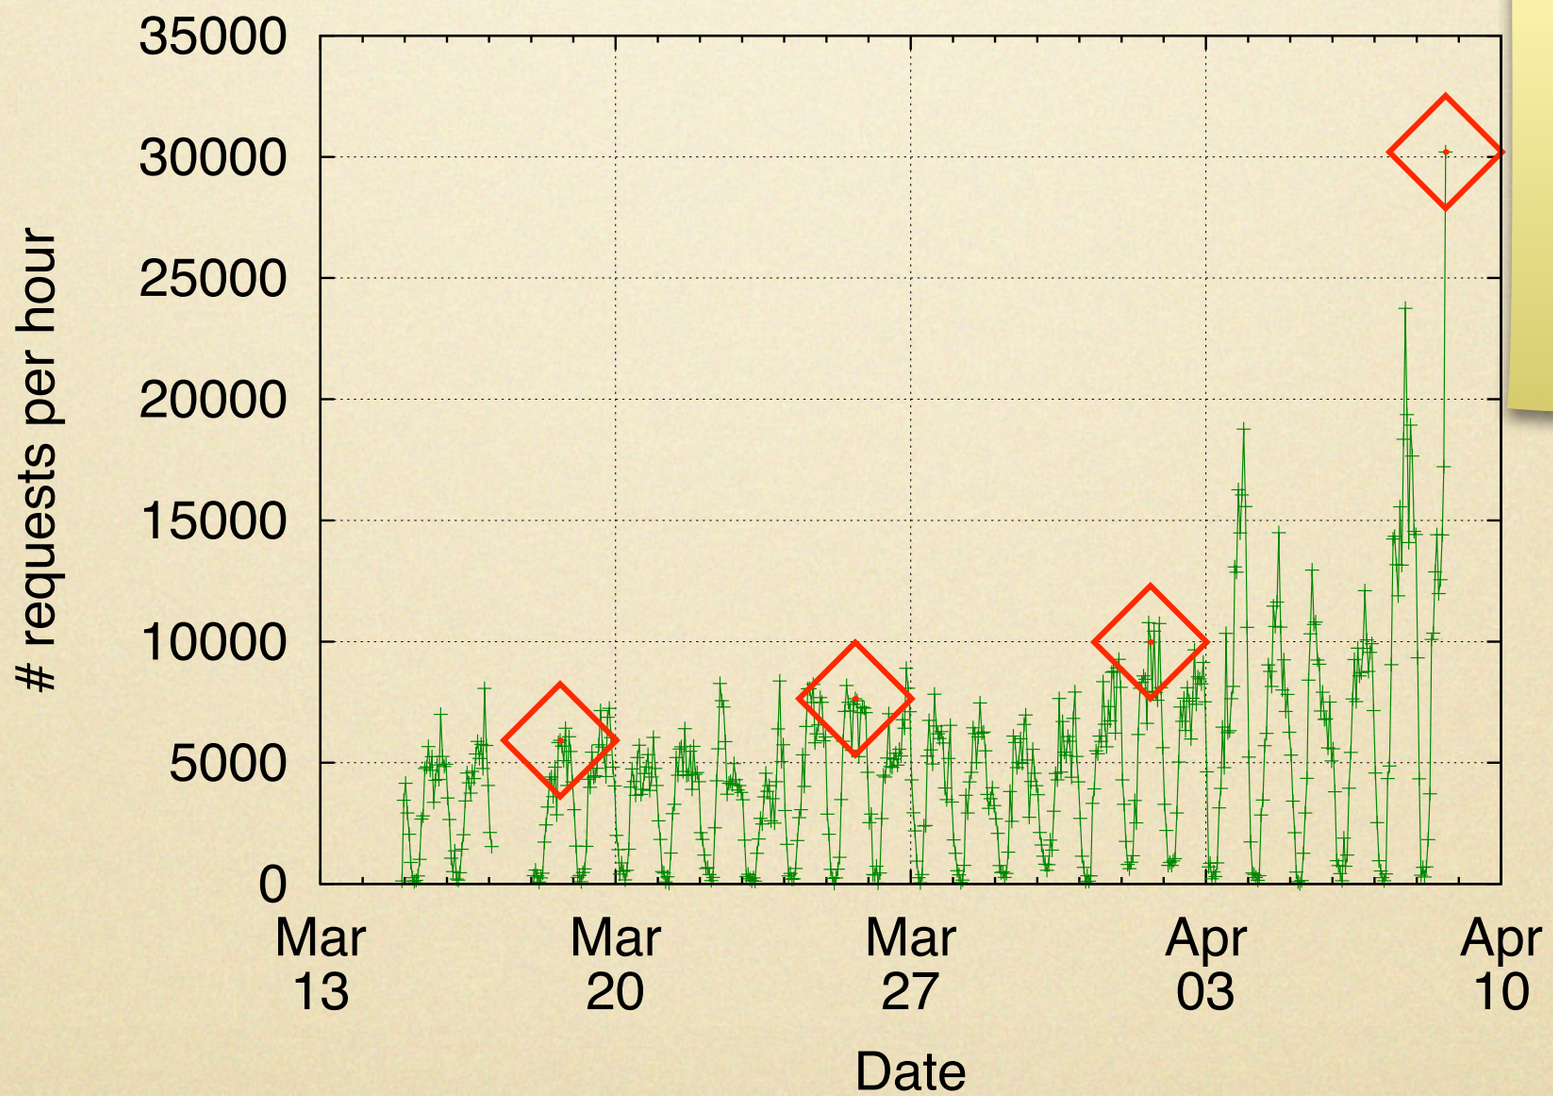

point is that this is an app-level measurement of load...NOT that we couldn't have figured out high load via other mechanisms.

Case study: investigation

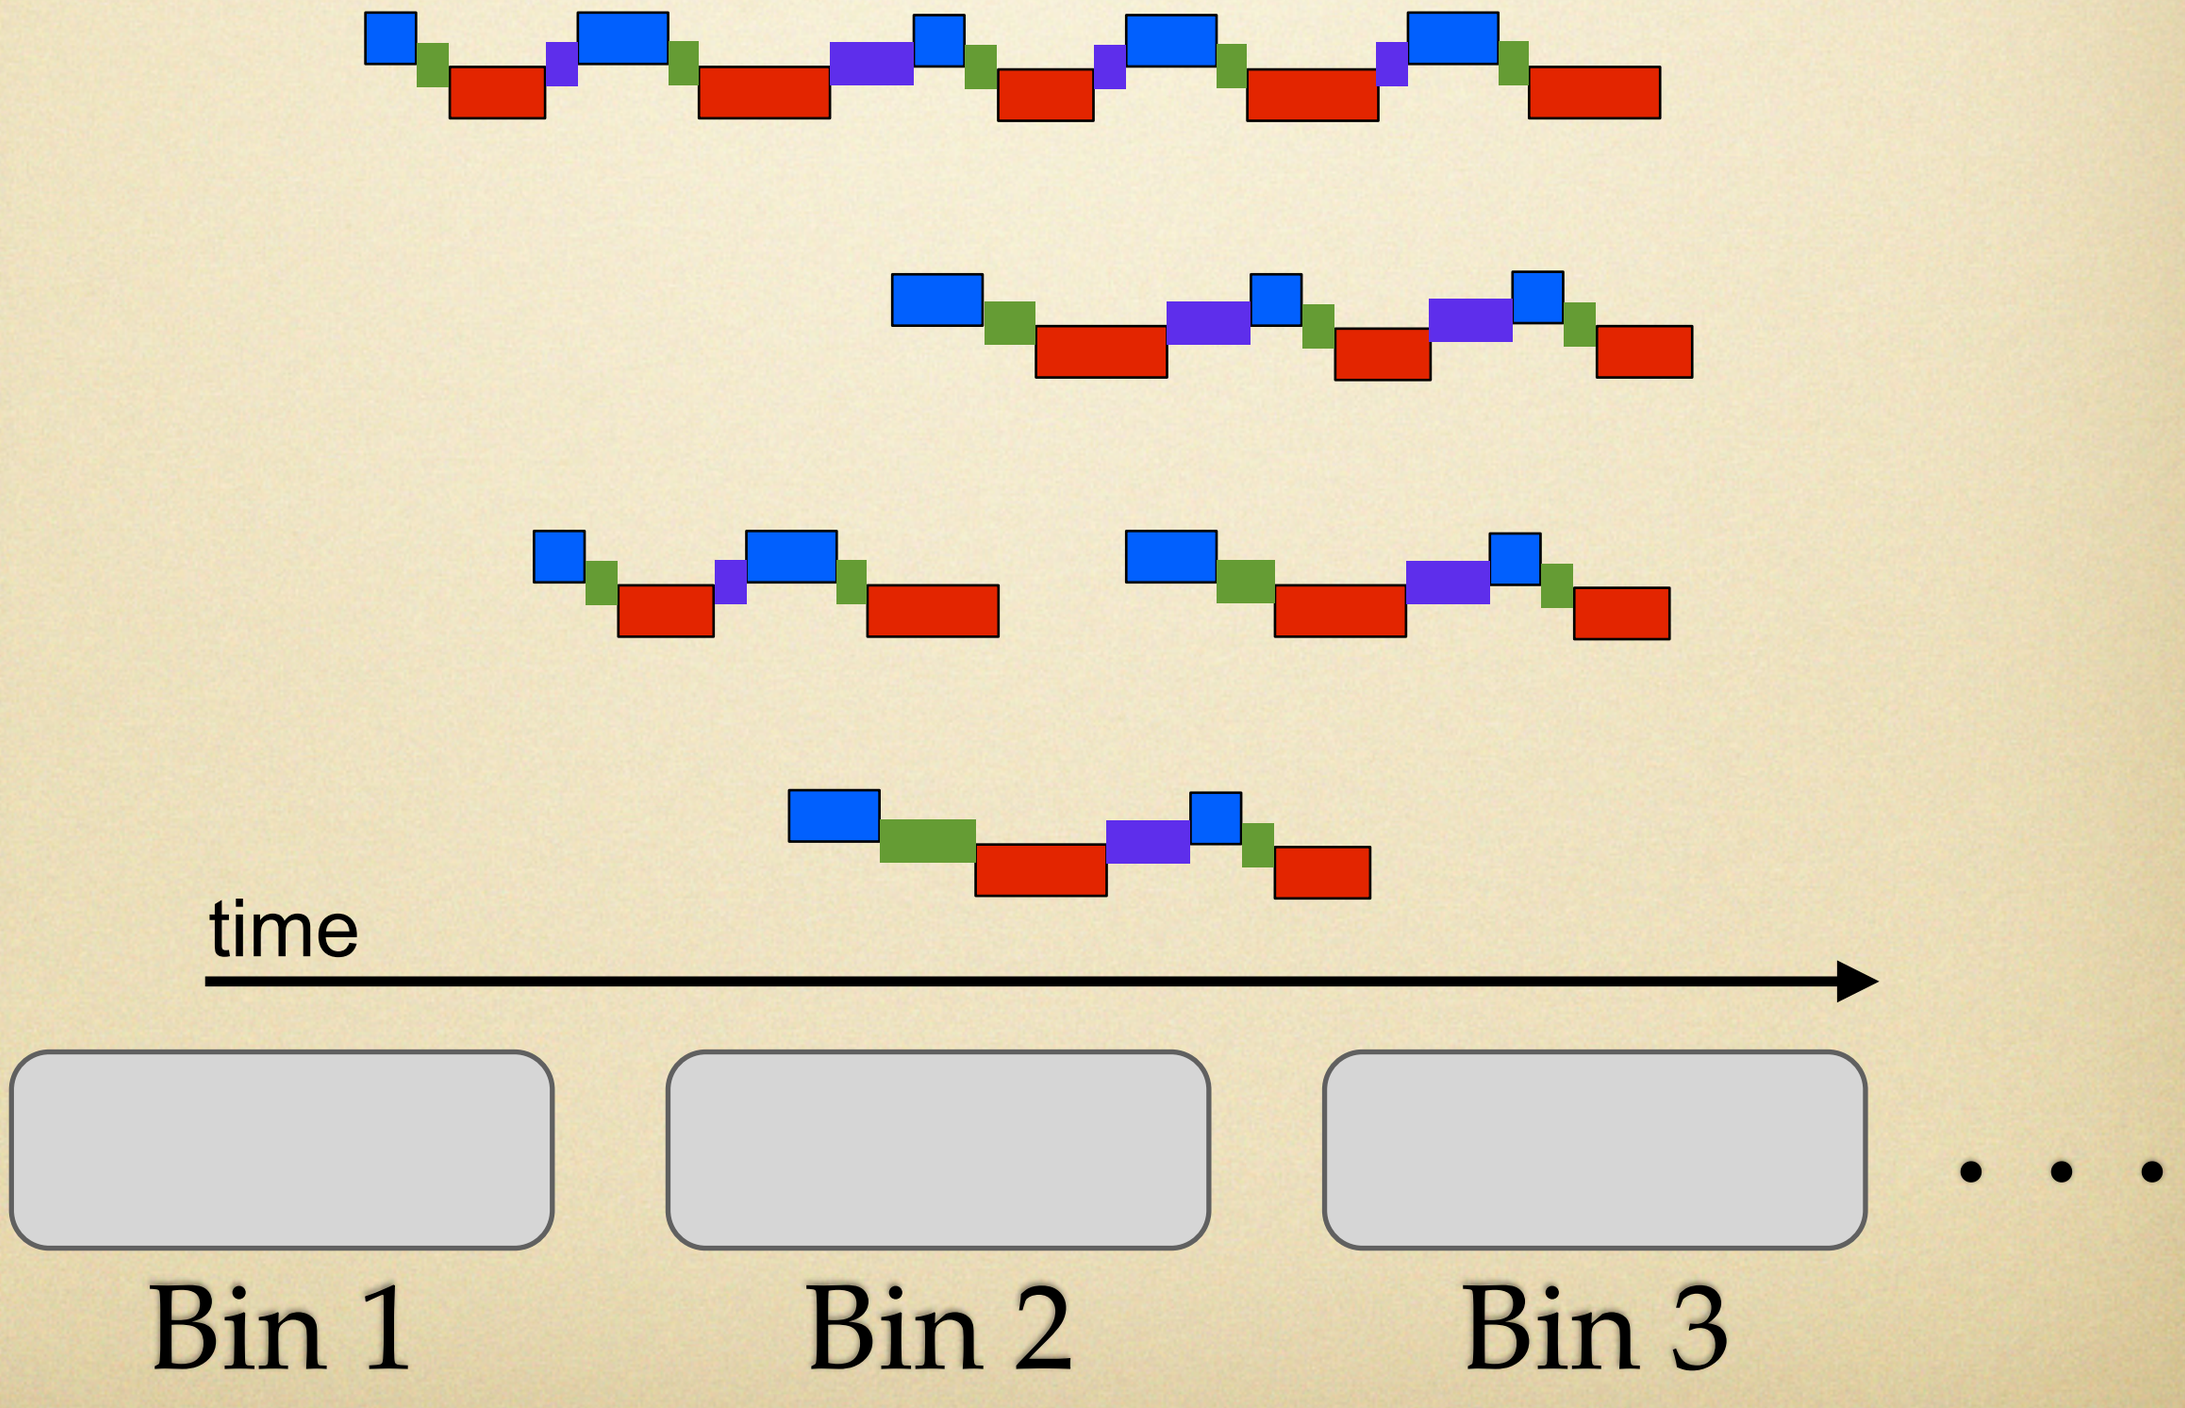

Case study: investigation

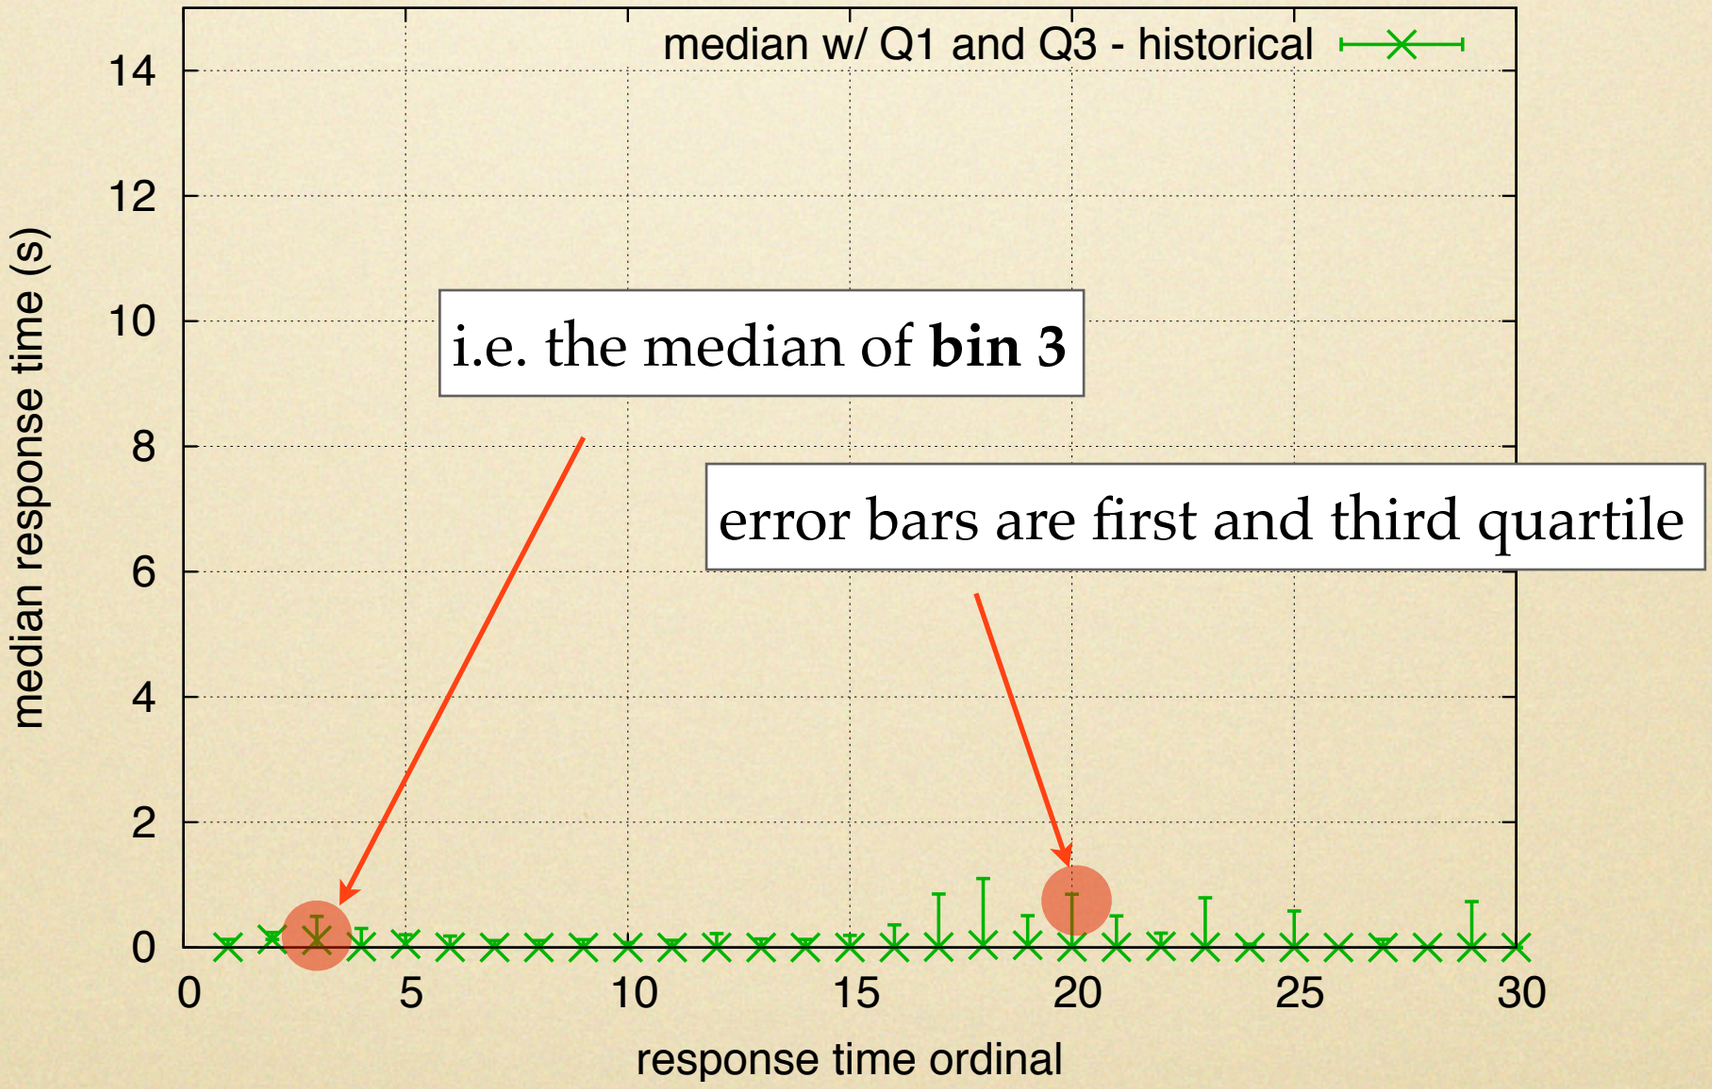

Case study: investigation

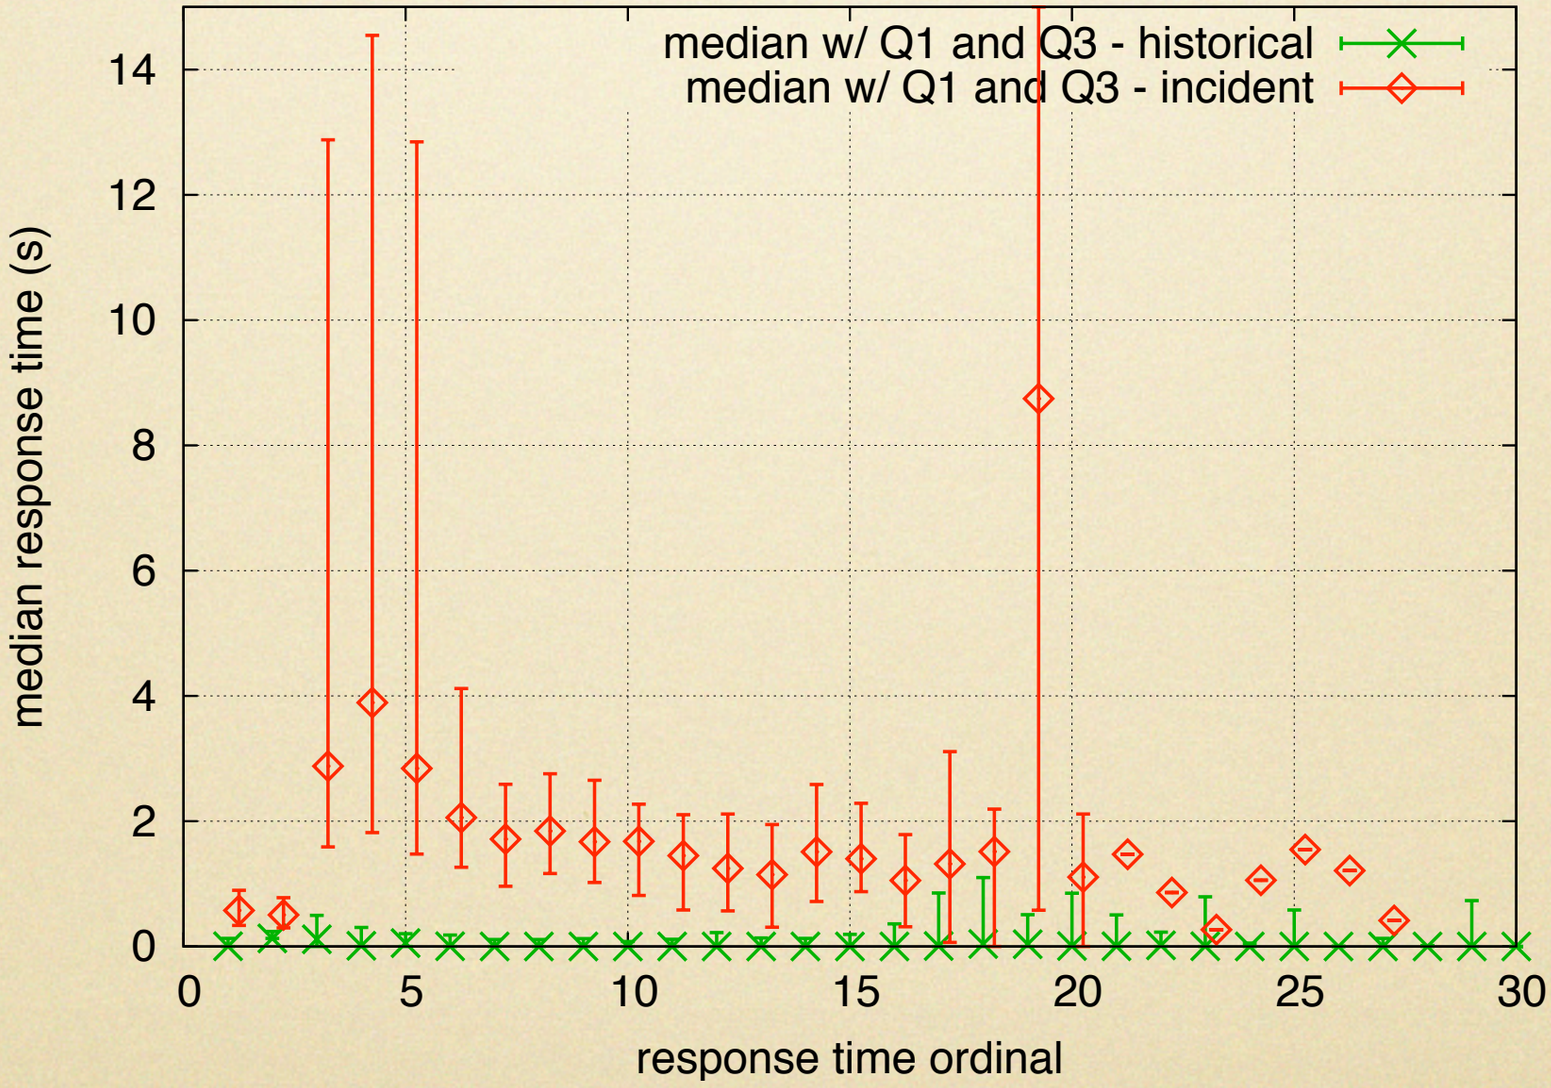

Case study: investigation

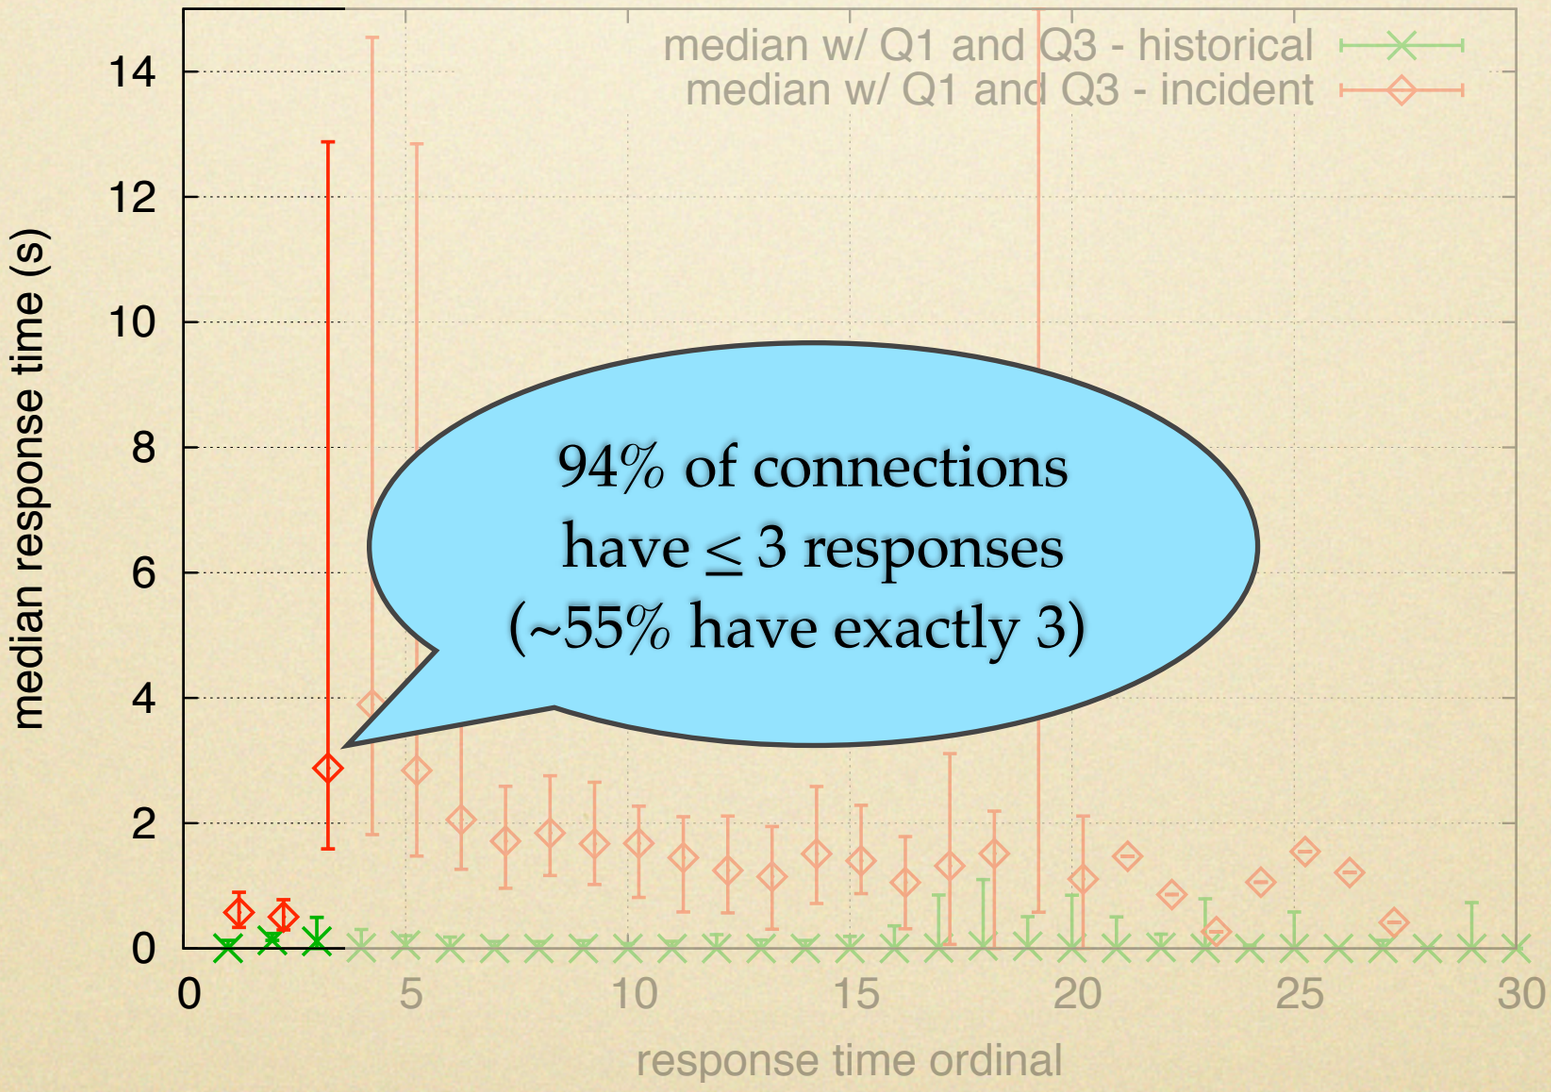

Conclusions

- Achieved monitoring of server performance:
  - for all servers, of any type
  - in real-time, at gigabit speeds,
  - on older hardware,
  - completely passively.
- `adudump` data provides diagnostic insight into performance issues.

Questions?