Reliability

10:50am – 12:10pm
Tuesday, 10/9

Haonan Lu

University of Southern California
and Princeton University

Web Services Are Important

Services Need to Be Reliable!

Web Services Are Huge

[1][2]

- 2.5 B – content items shared
- 2.7 B – “likes”
- 300 M – photos uploaded
- 105 TB – data scanned
- 500 TB – new data ingested

[1] Facebook data science. https://www.facebook.com/data
[2] “How Big Is Facebook’s Data?” https://goo.gl/bBN2ch

Service Is Distributed

Failure Happens

- Causes
  - Software bugs, misconfiguration, etc.
  - Hardware, power cut, natural disasters, etc.
- Scope
  - Single machine, rack, datacenter, etc.

Failure Is Inevitable

Failure Is Inevitable

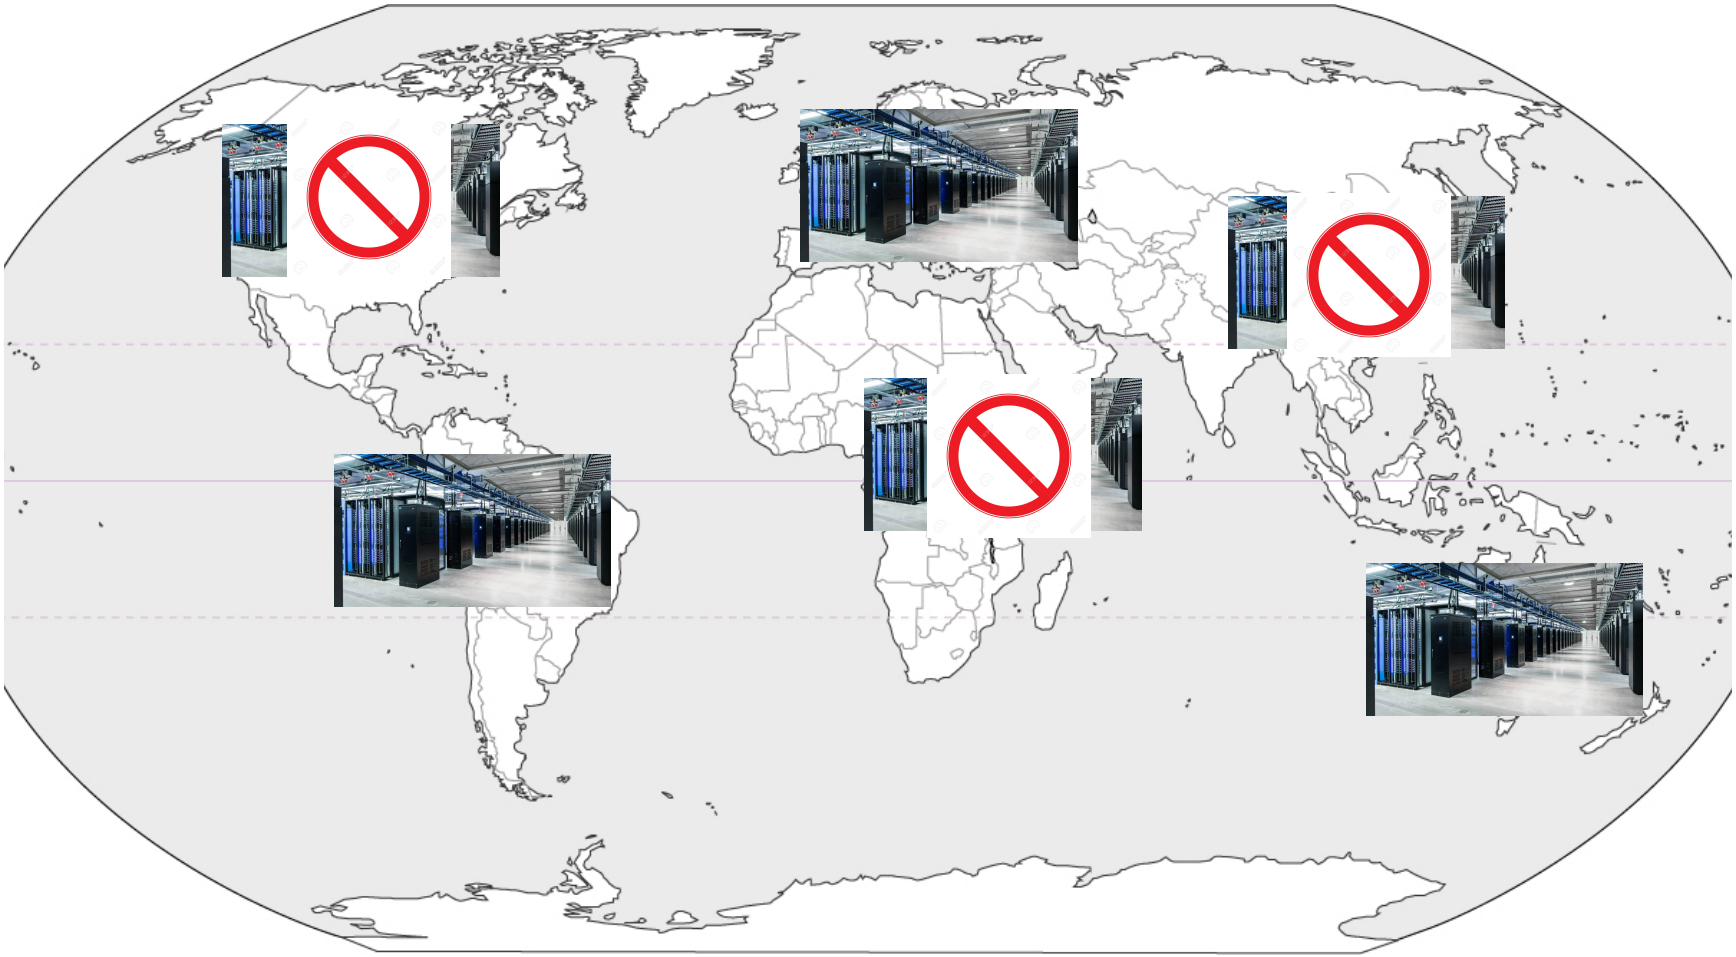

System Reliability

- How well does a system tolerate failures
  - How effective?
  - How efficient?

  } "Fuzzylog", "SAUCR"
  - Minimizing user impact? "Maelstrom"
  - Etc.

State Machine Replication
– replicas like a single copy

- Multiple copies (replicas)
- Same input
- Same order

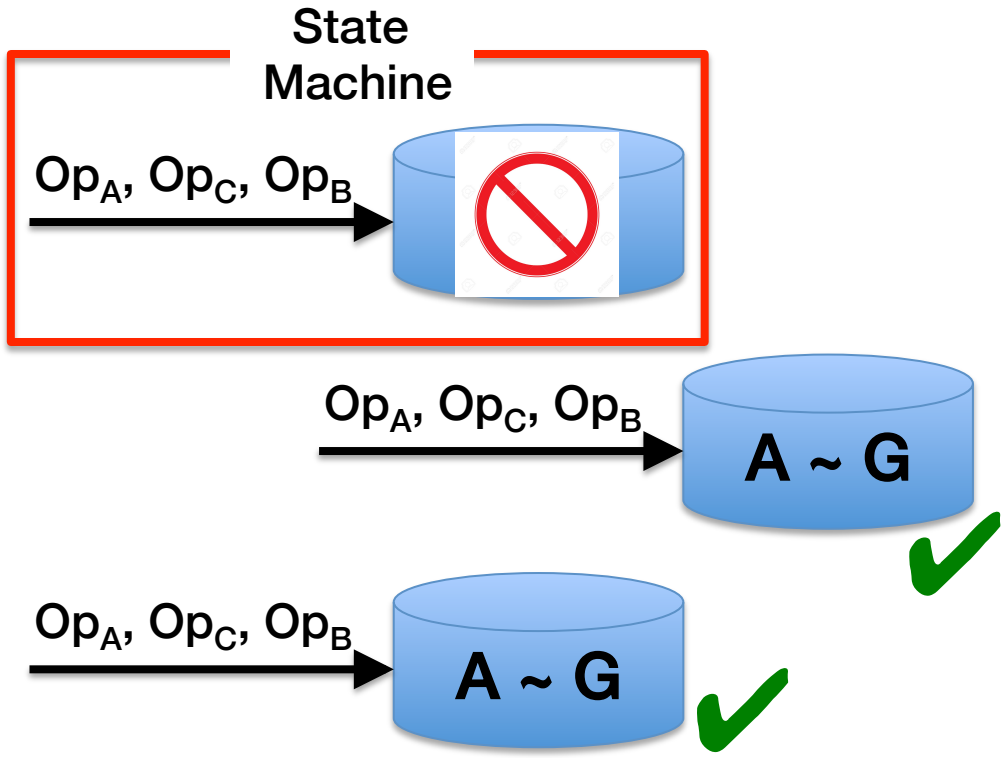

SMR Protocols

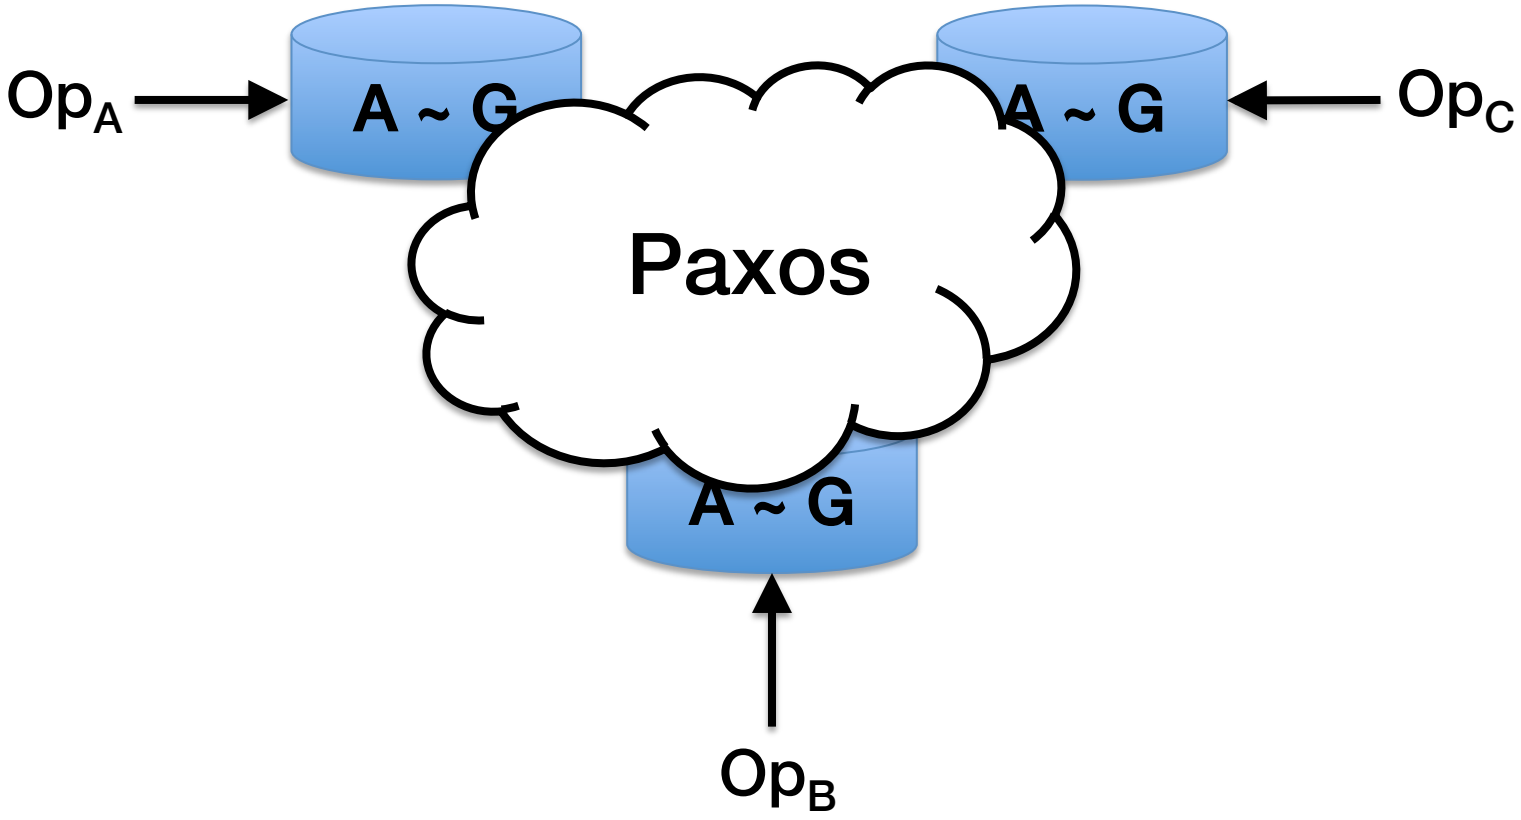

SMR Protocols

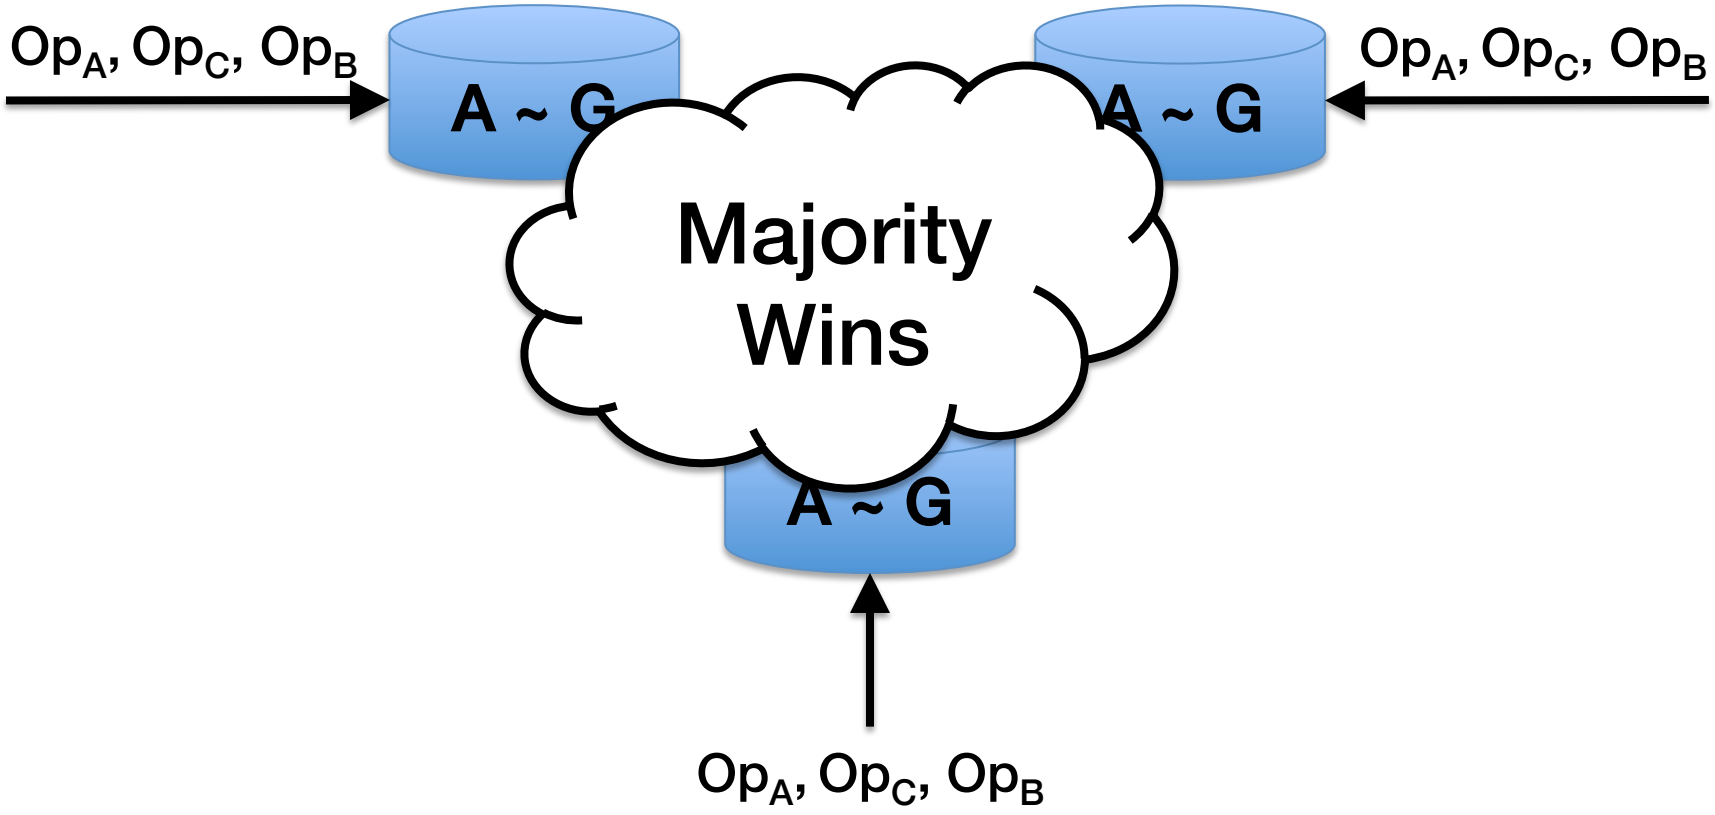

SMR Protocols

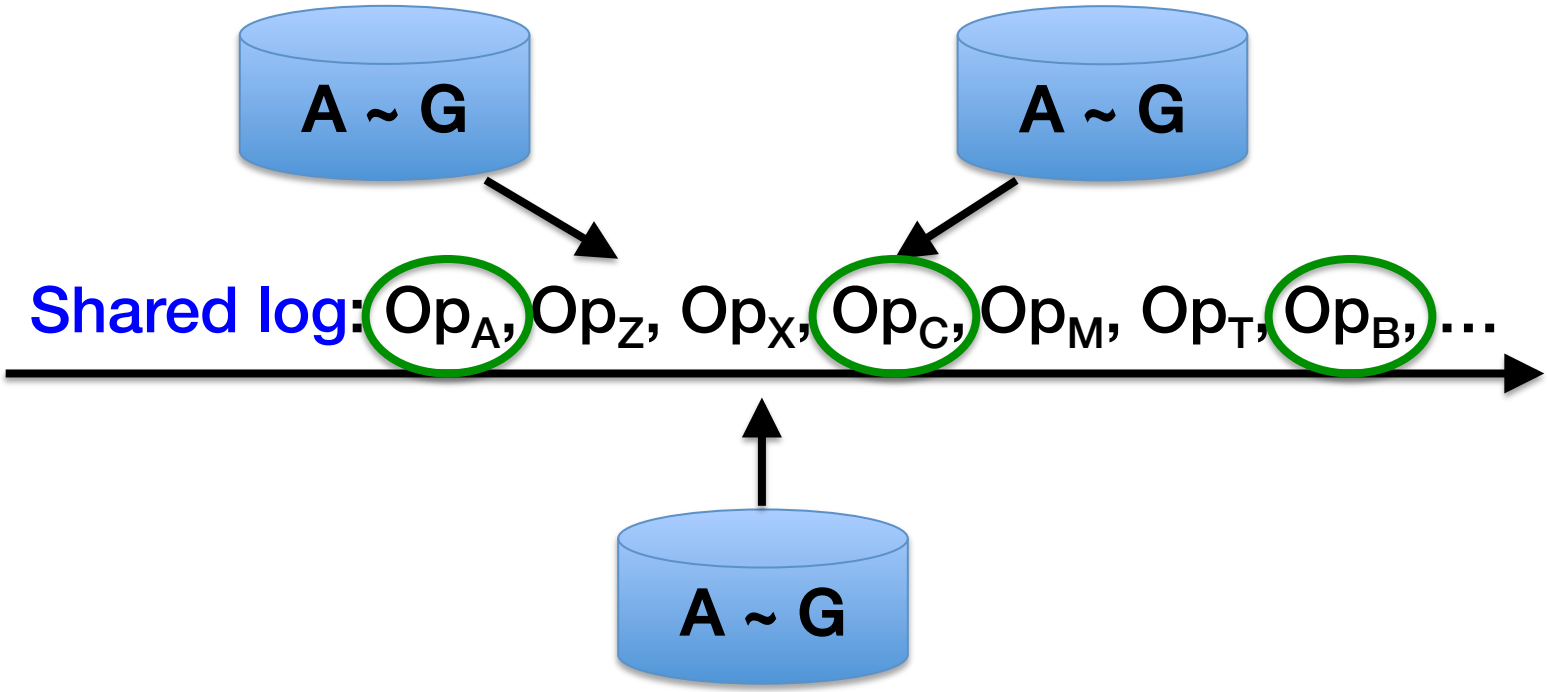

“The FuzzyLog: A Partially Ordered Shared Log”

- How do they solve the problem of expensive total ordering in shared log approach?

“Fault-Tolerance, Fast and Slow: Exploiting Failure Asynchrony in Distributed Systems”

- Memory storing states
  - Fast, but less durable
- Disk storing states
  - Durable, but slow

} Sweet spot?

“Maelstrom: Mitigating Datacenter-level Disasters by Draining Interdependent Traffic Safely and Efficiently”

- How does Facebook serve user requests when an entire datacenter is down?

- How do they test/evolve the failure recovery subsystem?

“Taming Performance Variability”

- Performance reliability of hardware
  - How does service providers control the variability of devices provided for users?
  - How do users cope with hardware variability when running experiments?

Conclusion

- Reliable web services are important!
  - Good user experience & revenue
- Reliable web services are challenging!
  - Large scale & failures everywhere
- Solution: state machine replication
- Talks in reliability session
  - More efficient protocols
  - How Facebook deals with DC disasters

Thank you !

Service Is Distributed

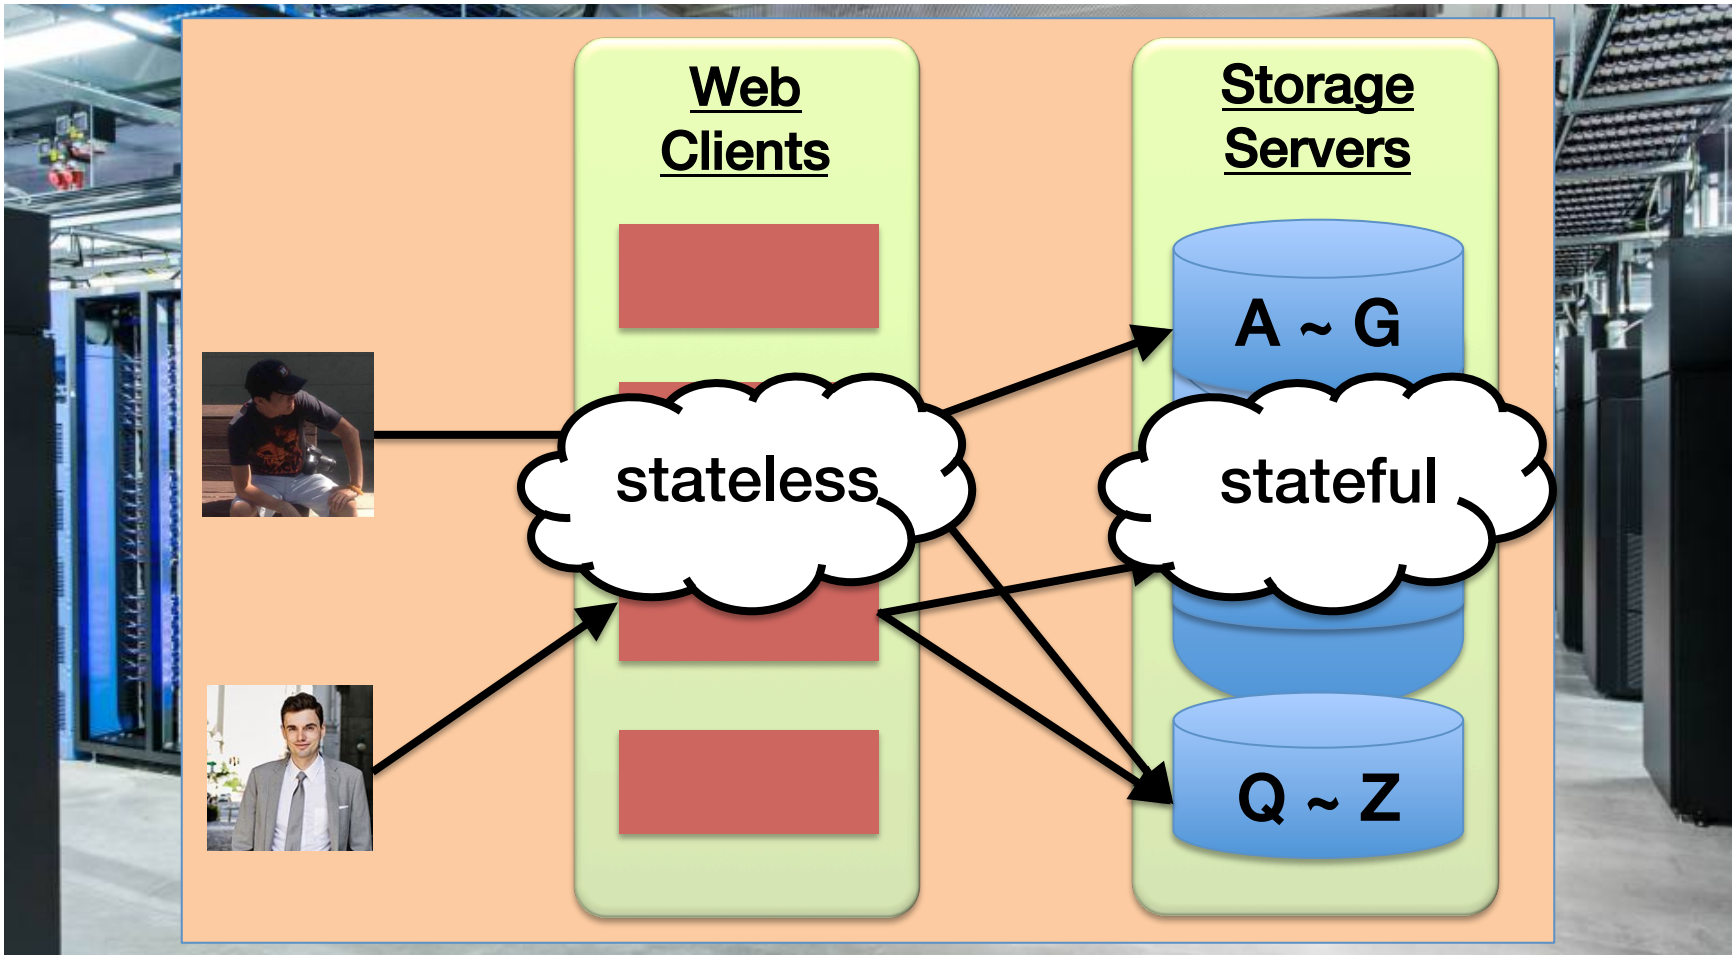

Failure Is Inevitable

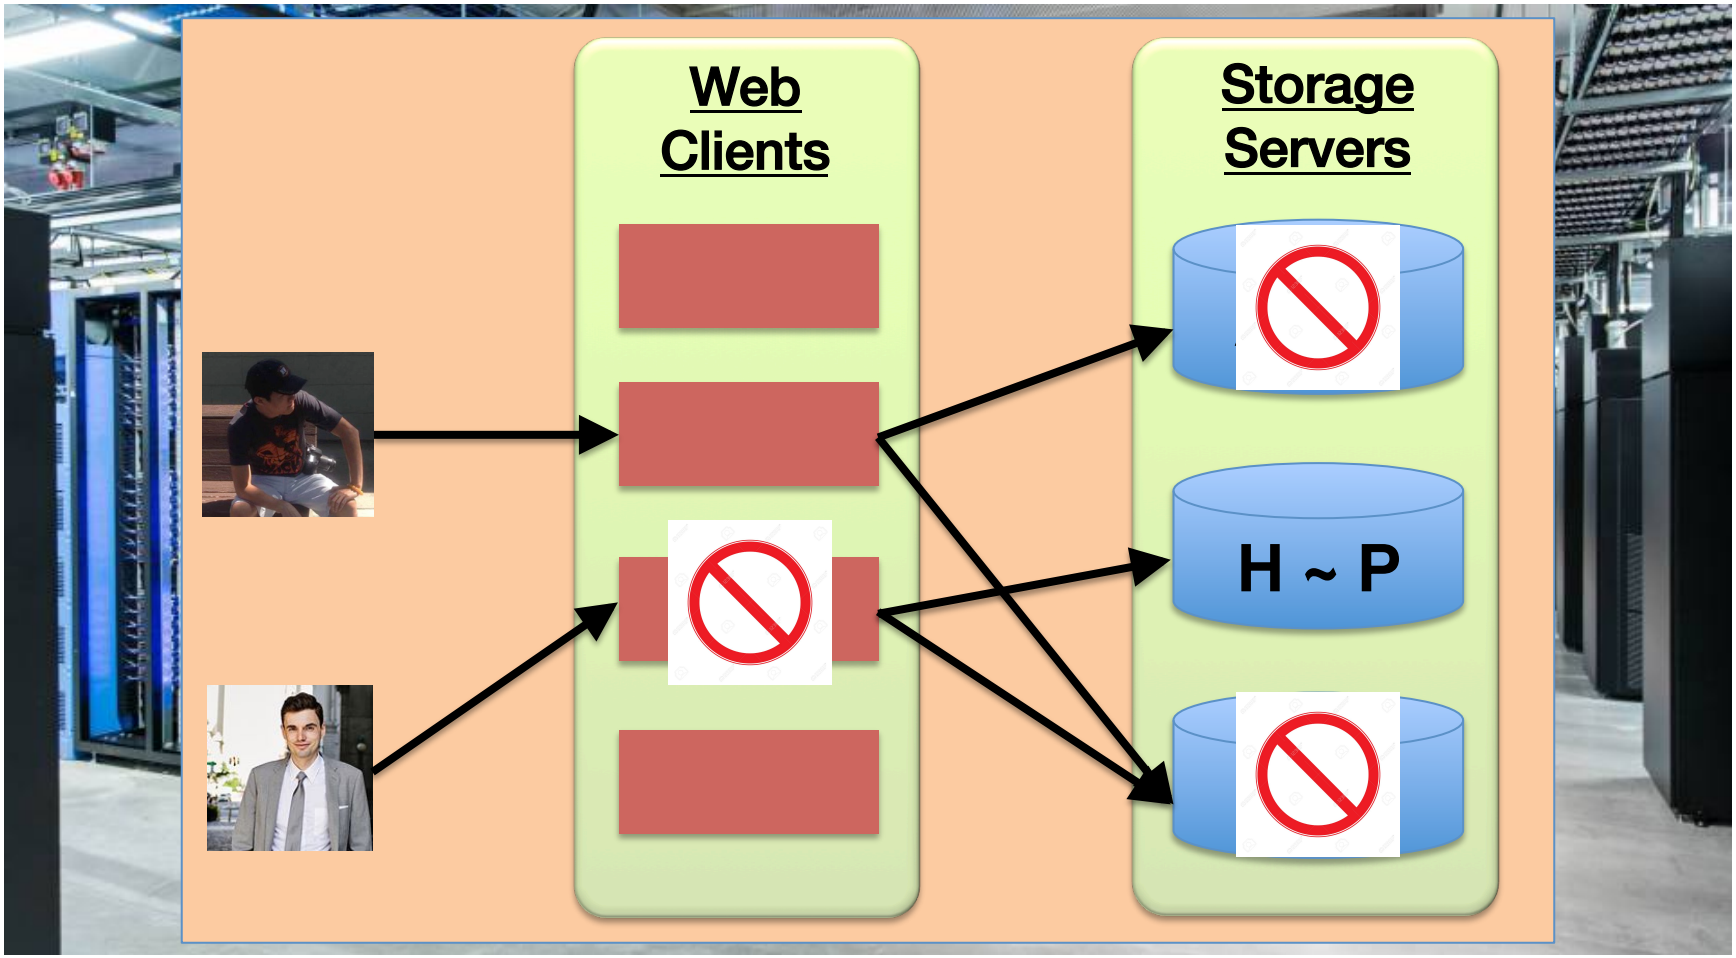

Failure Is Inevitable

- "the probability of seemingly strange behavior can be made very small. However, the distributed nature of the system dictates that this probability can never be zero."
  - P. R. Johnson and R. H. Thomas. *Maintenance of duplicate databases.* RFC 677, Jan. 1975.

System Reliability

- How well does a system tolerate failures
  - How effective/efficient are the mechanisms
  - How fast is the recovery
  - How well do they avoid impact on users
  - Etc.

SMR Protocols

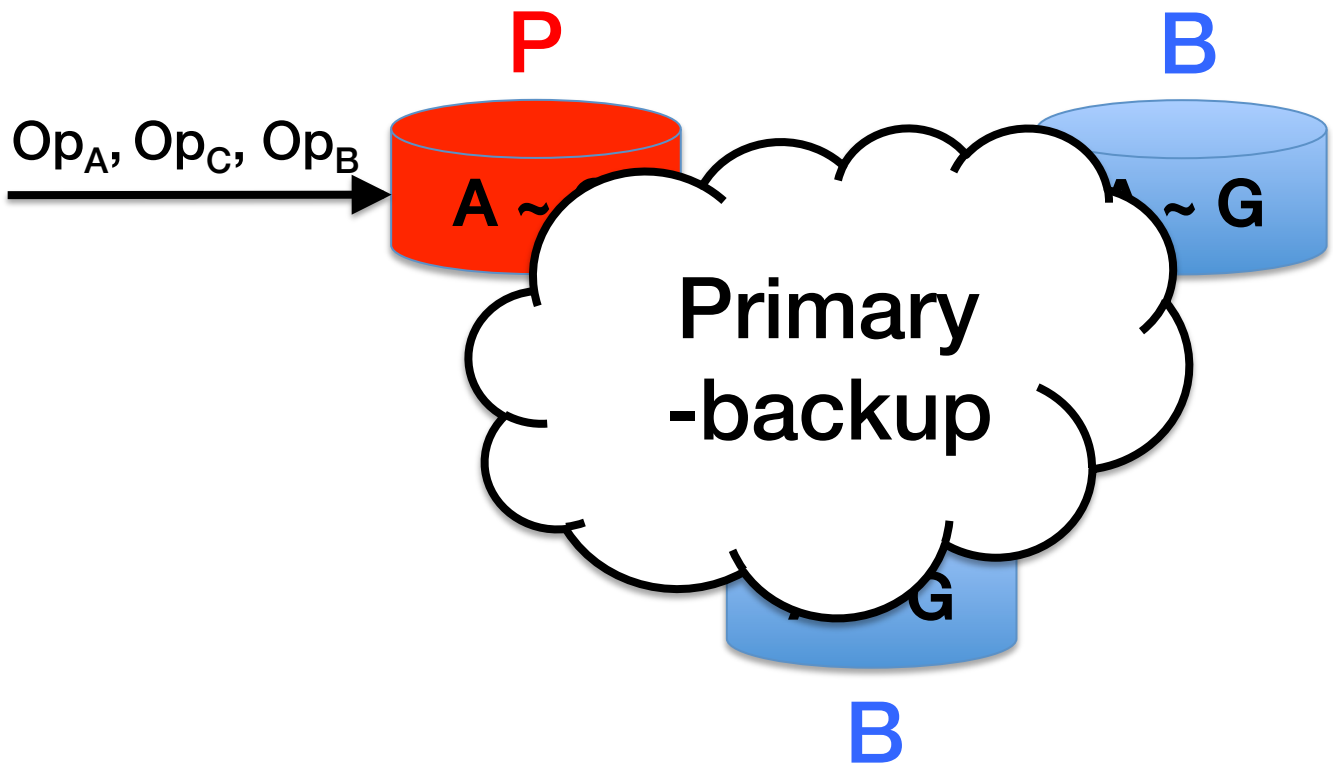

SMR Protocols