

L19: Replicated state machines

Frans Kaashoek
6.033 Spring 2013

Single-copy consistency

Operations appear to execute as if there is only a single copy of the data

Replicated state machine

- A general approach to building consistent replicas of a server:
 - Start with same initial state
 - Provide each replica with the same input and in same order
 - Ensure all inputs are deterministic

RSM is hard in practice [comp.risks]

- Leap day 2012, Microsoft's Azure
 - Configuration error crashed all servers
- April 2011, Amazon's EC2
 - Replication storm exhausted all storage
- September 2010, Facebook clients
 - Clients overwhelmed database
- September 2009, Google's gmail
 - Request routers overloaded

Complete outages

RSM with view server

- Non-backups reject forwarded requests
- Primary in view i must be backup in view $i-1$
- Non-primaries reject client requests
- Primary must wait for each backup to accept request
- New backups copy state from primary

Summary

- Single-copy consistency
- Replicated state machine
- Consensus with view server
- Wednesday: distributed consensus