Recitation 02/12: Consistency, Availability and Fault Tolerance

Sagar Jha

Distributed Systems

- Nodes connected by links
- What is a node?
- Nodes can act independently, fail independently
- What is a failure?
- What is a system?
- Hiding distribution from the external world

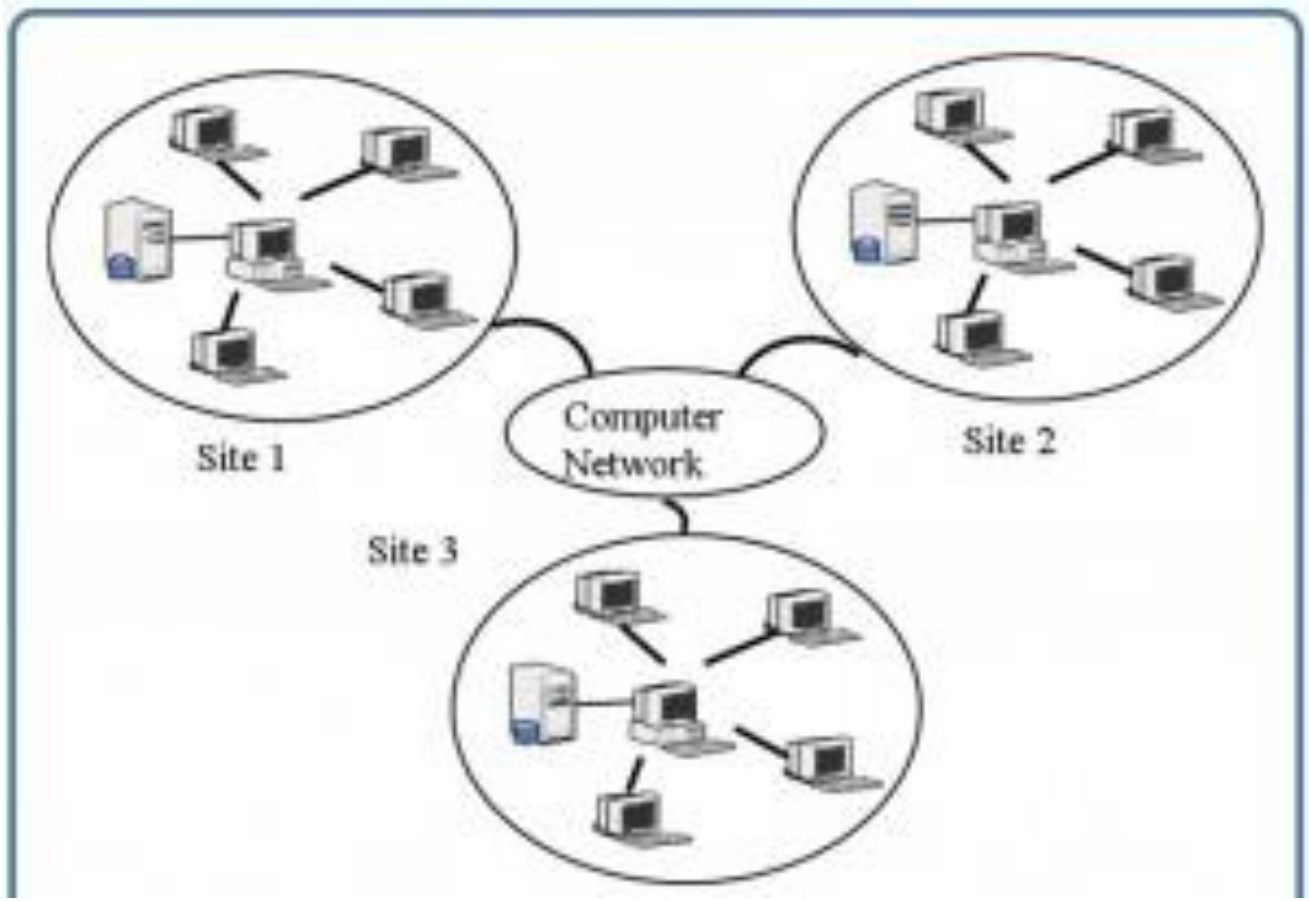

Why do we build systems?

- Availability
  - Single most important goal – responsive to clients, other systems and internal components
  - Measure (in %) – The fraction of time a system is available
  - Availability is with the constraint of being consistent
- Consistency
  - Systems maintain state and respond to queries based on the state
  - State should be consistent, State and data should not be lost
  - Different interpretations
    - Preserve system guarantees, protect against hazards in critical systems

Failures complicate Availability and Consistency

- Recap – We discussed how nodes can fail and then we discussed our goals of availability and consistency
- How common are some of these failures?
  - Failures every 8 hours in a Google service
  - Probability of any one node failing is high in a big datacenter
- Failures can compromise availability
- Failures can compromise consistency

Consequences of losing availability & consistency

- Availability – You built the system to be available, right?
  - Services can lose millions of dollars if they're unavailable
  - Unavailability can drive users away
  - five-nines availability

- Consistency – What's the point if it's not consistent?
  - System can become non-operational
  - Real-world impact – loss of lives, property

Fault tolerance – Central problem in Distributed Systems

- “Availability and Consistency despite failures”
- Some examples of fault tolerance in real-life:
  - Applying to multiple schools
  - Producing multiple kids (in old times)
  - Portfolio diversification
- Strategies for fault tolerance
  - Replication

CAP – Trade-off b/w consistency and availability in cases of limited failures

- By Eric Brewer, C – Consistency, A – Availability, P – Partition tolerance
- Easy to see that failures can impact A & C
- CAP stresses on the trade-off between maintaining one over the other during failures
- It is a rule of thumb
- What is partition tolerance?