



University at Buffalo

Department of Computer Science  
and Engineering

School of Engineering and Applied Sciences

# DISTRIBUTED STORAGE SYSTEMS


CSE 501

Haonan Lu

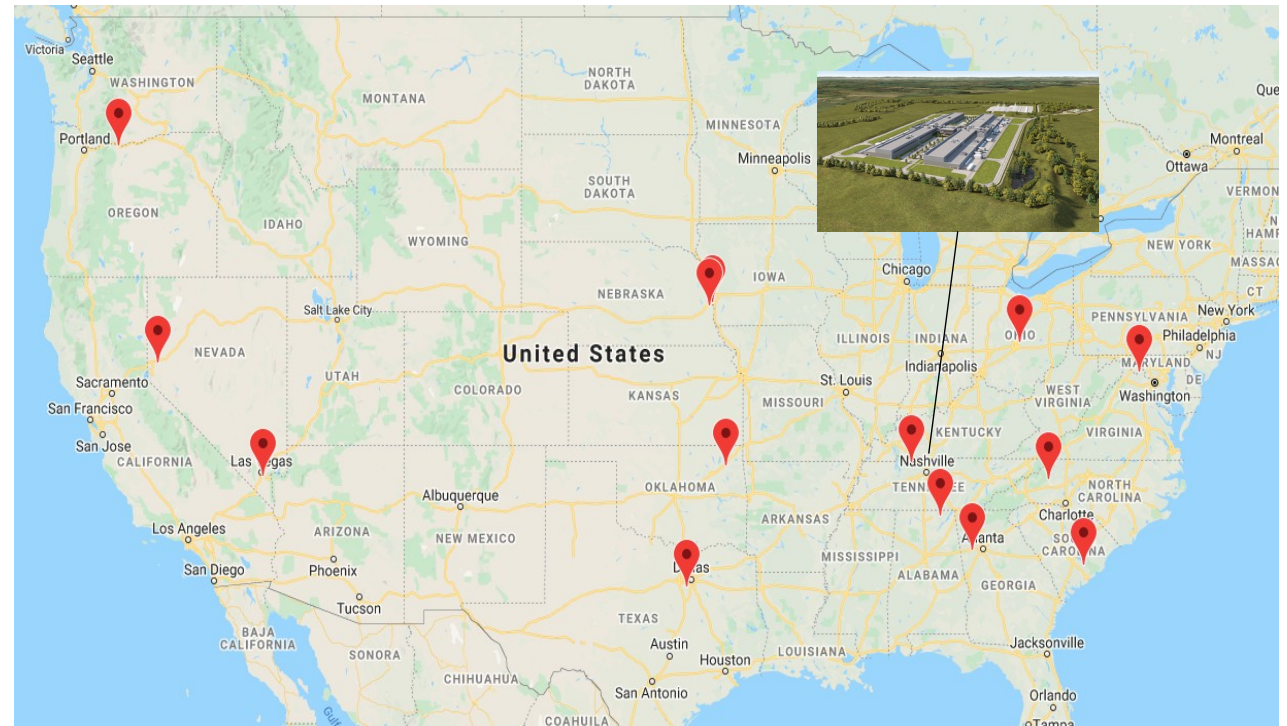
A row of racks of servers inside a typical datacenter



A typical  
datacenter

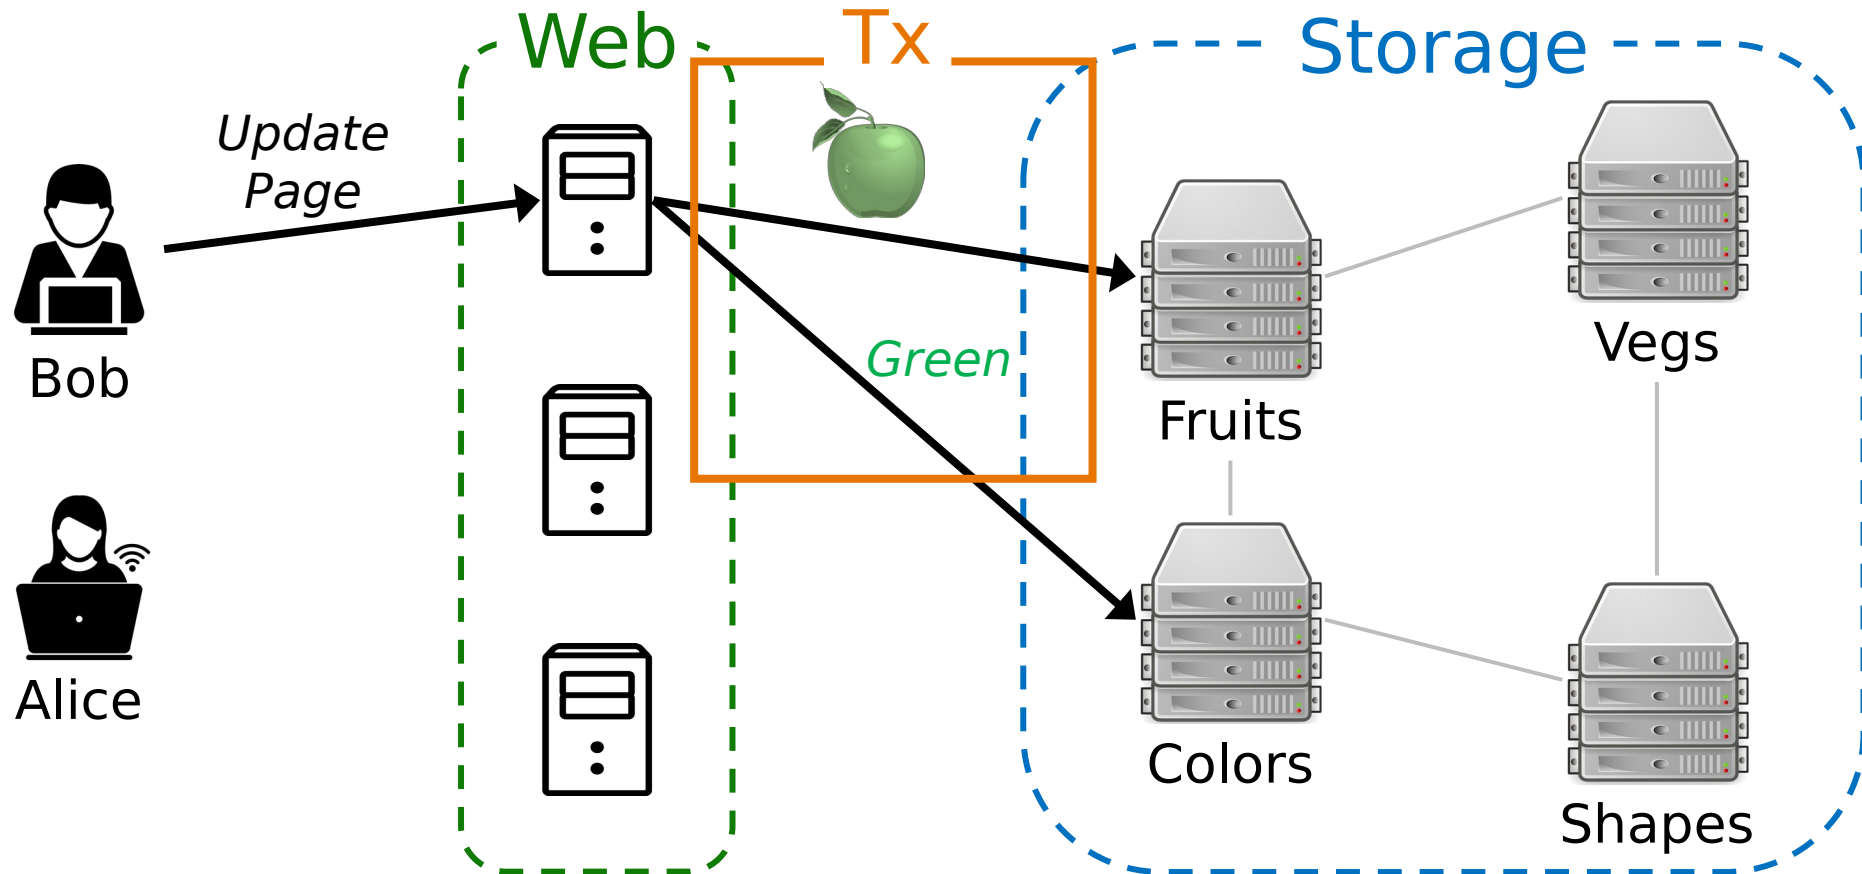


Geographically  
dispersed  
datacenters



# Distributed Storage Systems

## Data Sharding Enables Large Scale



# Research Problems

- Performance
  - System protocols achieving higher throughput and lower latency
  - Our research looks at it from a fundamental perspective
- Correctness
  - Verify the system indeed provides the claimed guarantees
- Easy adoption
  - Deploy in production systems --> practical impacts