



Scalability Terminology

CS 417 Recitation

Scalability Terminology: Farms, Clones, Partitions, and
Packs: RACS and RAPS - Microsoft Research

Server Systems Growth



- Scale
 - up: replaces servers with larger servers
 - out: add servers
- Farms
 - collection of servers, applications, and data
 - services: mail, database, http, etc
- Farm disaster tolerance
 - Geoplex: collection of farms
 - *active-active*
 - *active-passive*

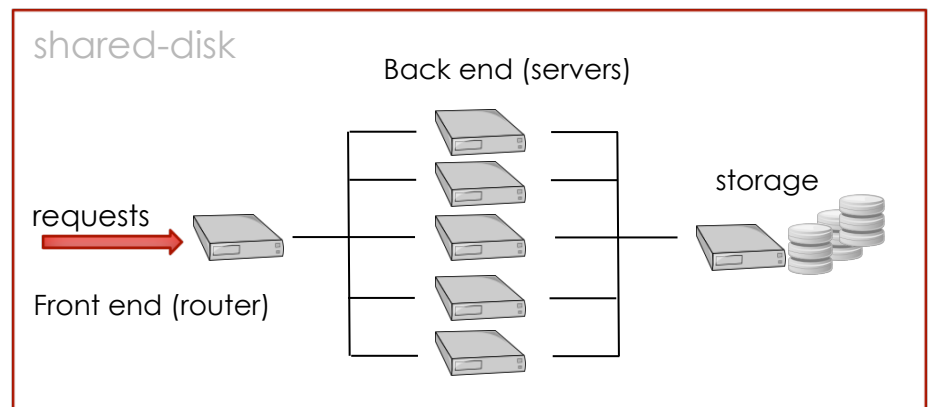
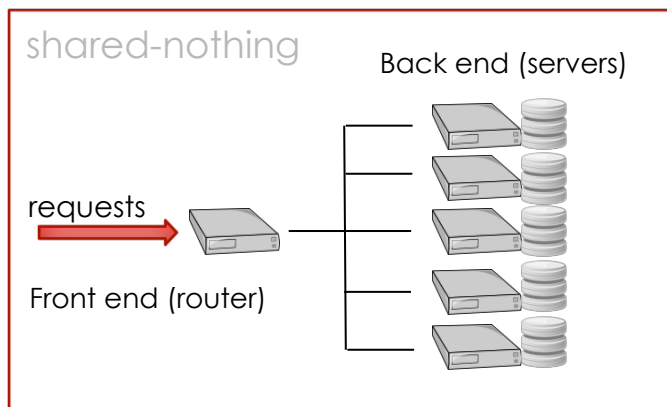
Farm Growth

- Cloning
 - Duplicate hardware, software and data
- Partitioning
 - Duplicate hardware and software
 - Divide the data



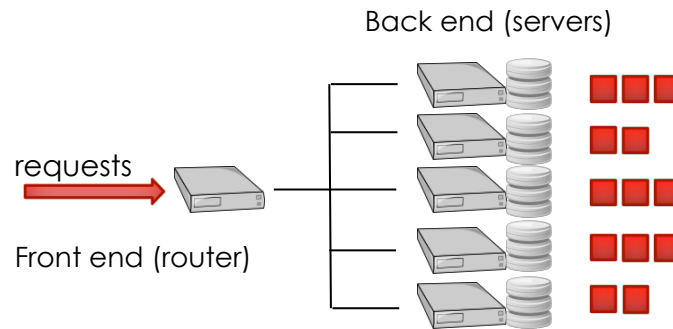
Farm Cloning

- Service replicated on many nodes
 - RACS: Reliable Array of Cloned Services
 - *shared-nothing*
 - *shared-disk*
- Requests are routed to individual nodes



Farm Cloning

- Load balancing
- Failures
 - Masking
 - Availability



Availability	Downtime per year
90%	36.5 days
99%	3.65 days
99.9%	8.76 hours
99.99%	52.56 minutes
99.999%	5.26 minutes
99.9999%	31.5 seconds

Farm Cloning

Shared-nothing RACS

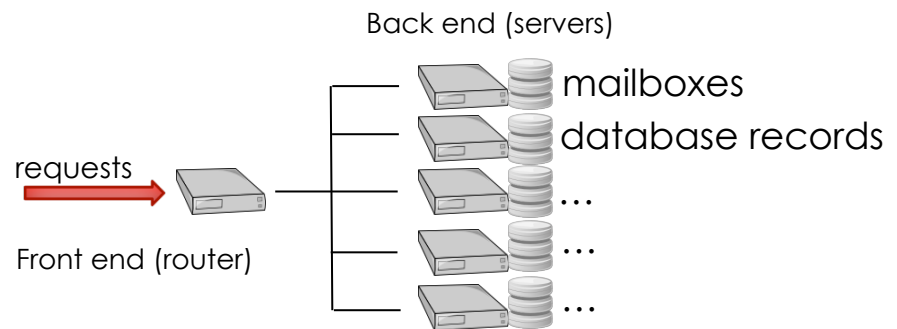
- Pros
 - Availability
 - Scalability
 - Processing power
 - Network and storage bandwidth
- Cons
 - Cannot improve storage capacity
 - Write-intensive services
 - Complex updates

Shared-disk RACS

- Pros
 - all shared-nothing pros
 - reduced storage complexity
- Cons
 - fault-tolerant storage server
 - reduced scalability
 - storage server becomes bottleneck

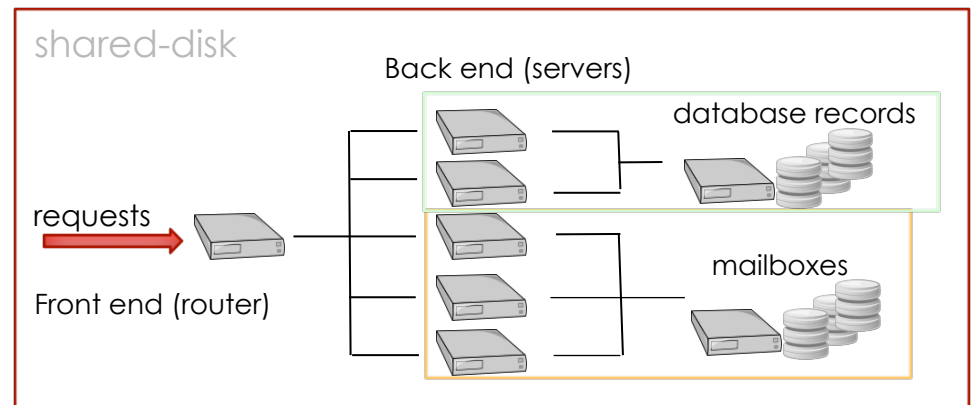
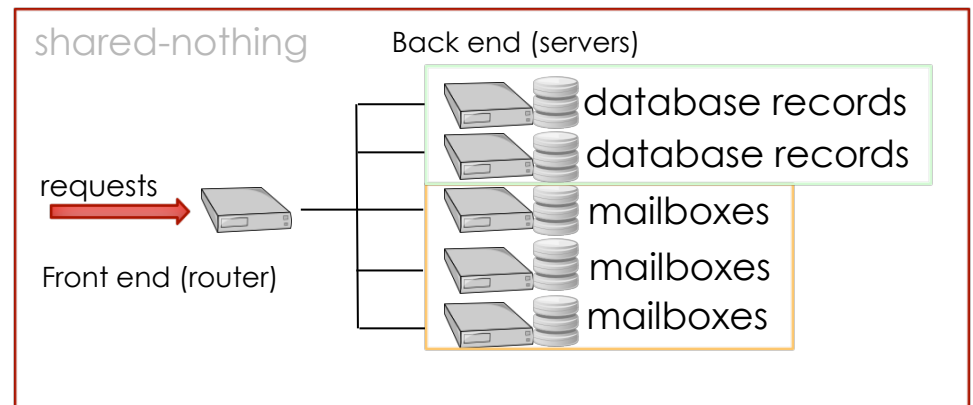
Farm Partitioning

- Duplicates hardware and software but divides data
- Load balancing
- Pros
 - Scalability
 - Processing power
 - Network bandwidth
 - Storage capacity and bandwidth
- Cons
 - Data stored in one place (could add raid)



Farm Partitioning

- Increase availability
 - replicate partitions into *packs*
- Packs
 - *shared-nothing*
 - *active-active*
 - *active-passive*
 - *shared-disk*
- RAPS (shared-nothing)
 - Reliable Array of Partitioned Services



Taxonomy of Scaleability Designs

