MySQL in the Cloud

Tricks and Tradeoffs

Thorsten von Eicken

CTO RightScale

MySQL & Amazon EC2 @RightScale

- Operating in Amazon EC2 since fall 2006
- Cloud Computing Management System
- Replicated MySQL product offering mid '07
  - 100's of customer deployments
  - Automation to launch slaves, failover, backup, …
- Animoto.com scaled to 4000 servers in 3.5 days

Cloud Computing & Databases

1. Infinite computing resources
2. Available on-demand
3. Pay per use
4. Delivered across the Internet

- Fully Automatable Infrastructure

As many DB machines as you want!

No cost for short periods

Save work, More reliable

Cloud Computing Vision

- No up-front costs

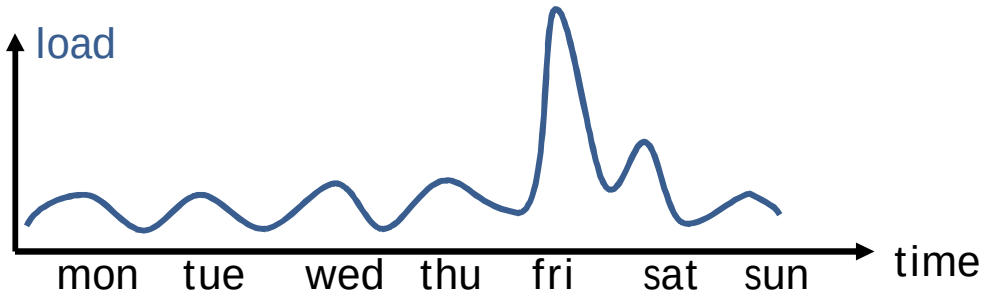

- Pay for average, not peak

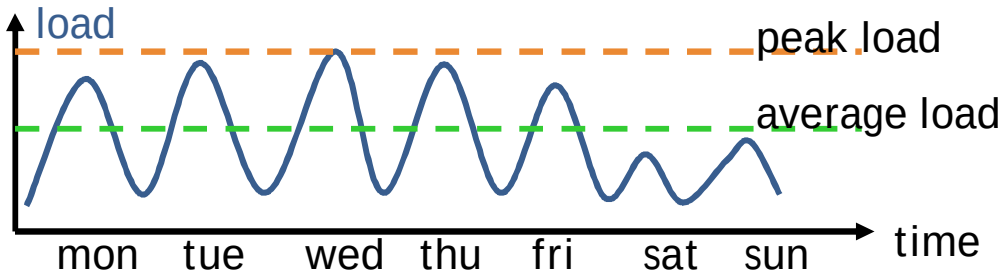

- Easy Dev & Test

- Scale batch apps

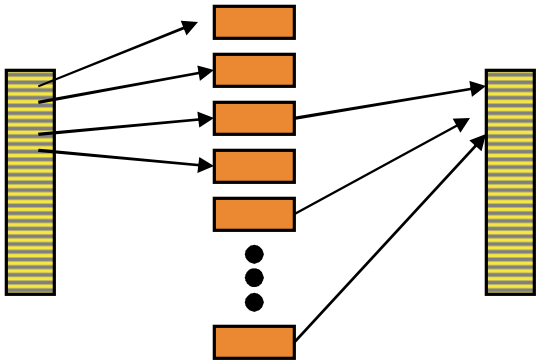

Scale-up/down Automation

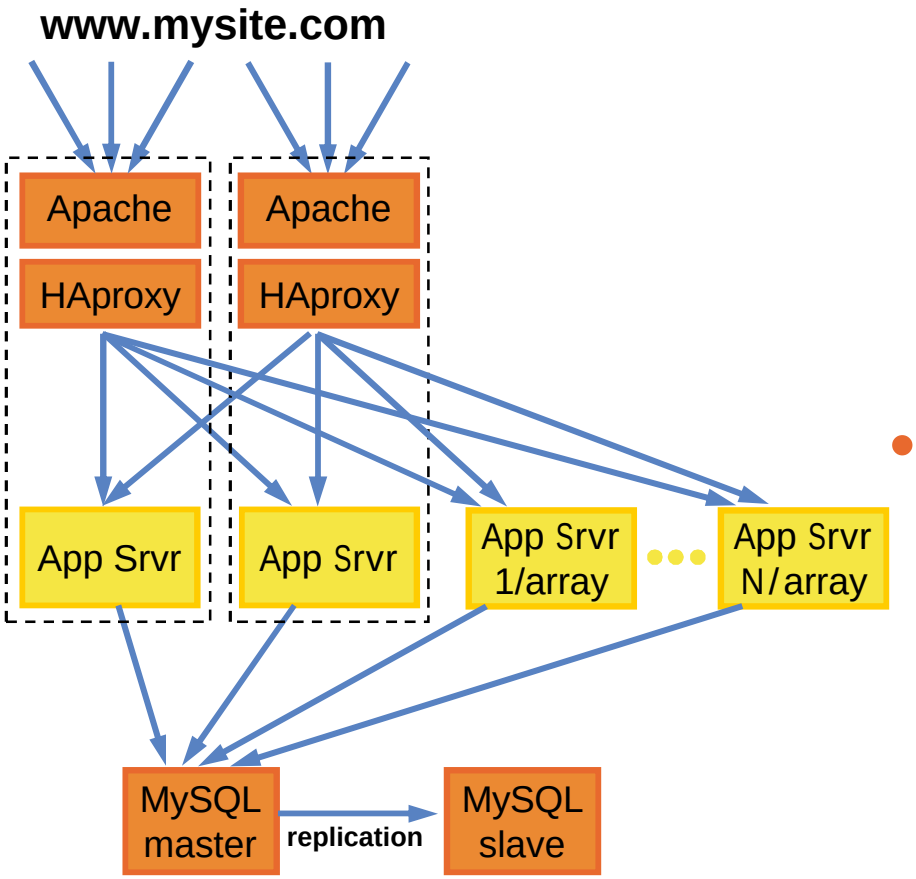

- Scale-up
  - App server load > 50%
  - Majority voting
  - Add app servers
  - Tie into load bal & db
- Scale-down
  - App server load < 20%
  - Majority voting

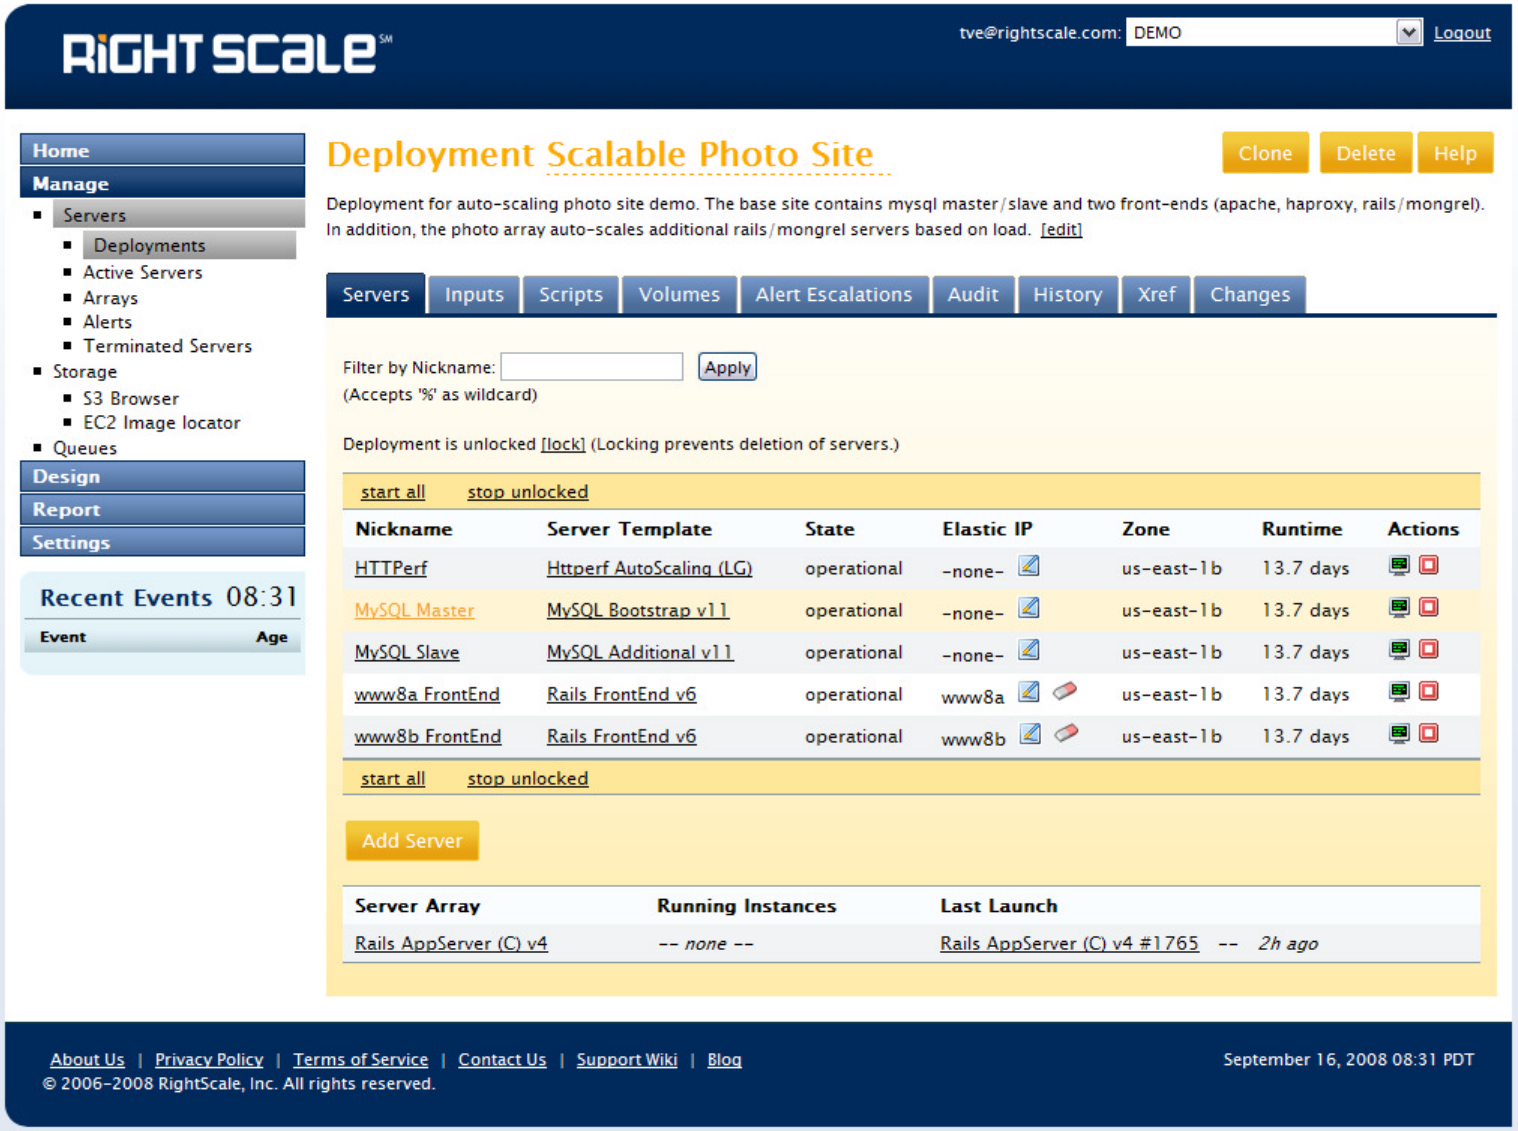

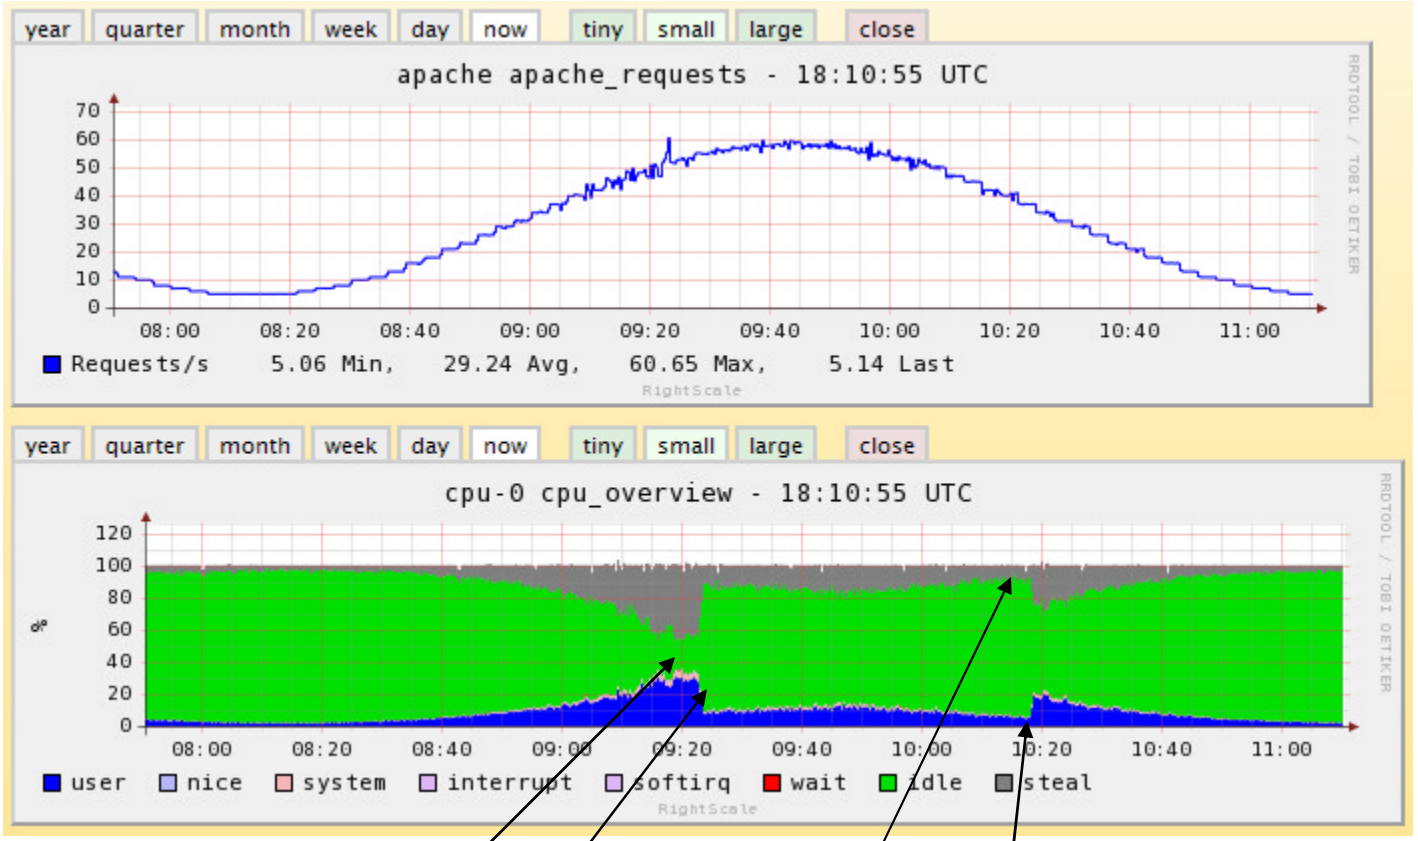

Load exceeds threshold

Additional servers operational

Load drops below threshold

Additional servers terminated

Starcut (Olympics)

- Low cost to entry, scale up as demand grows
- Repeatable, ”clean-room” design promotes stability and process
- Great sandbox for testing and benchmarking

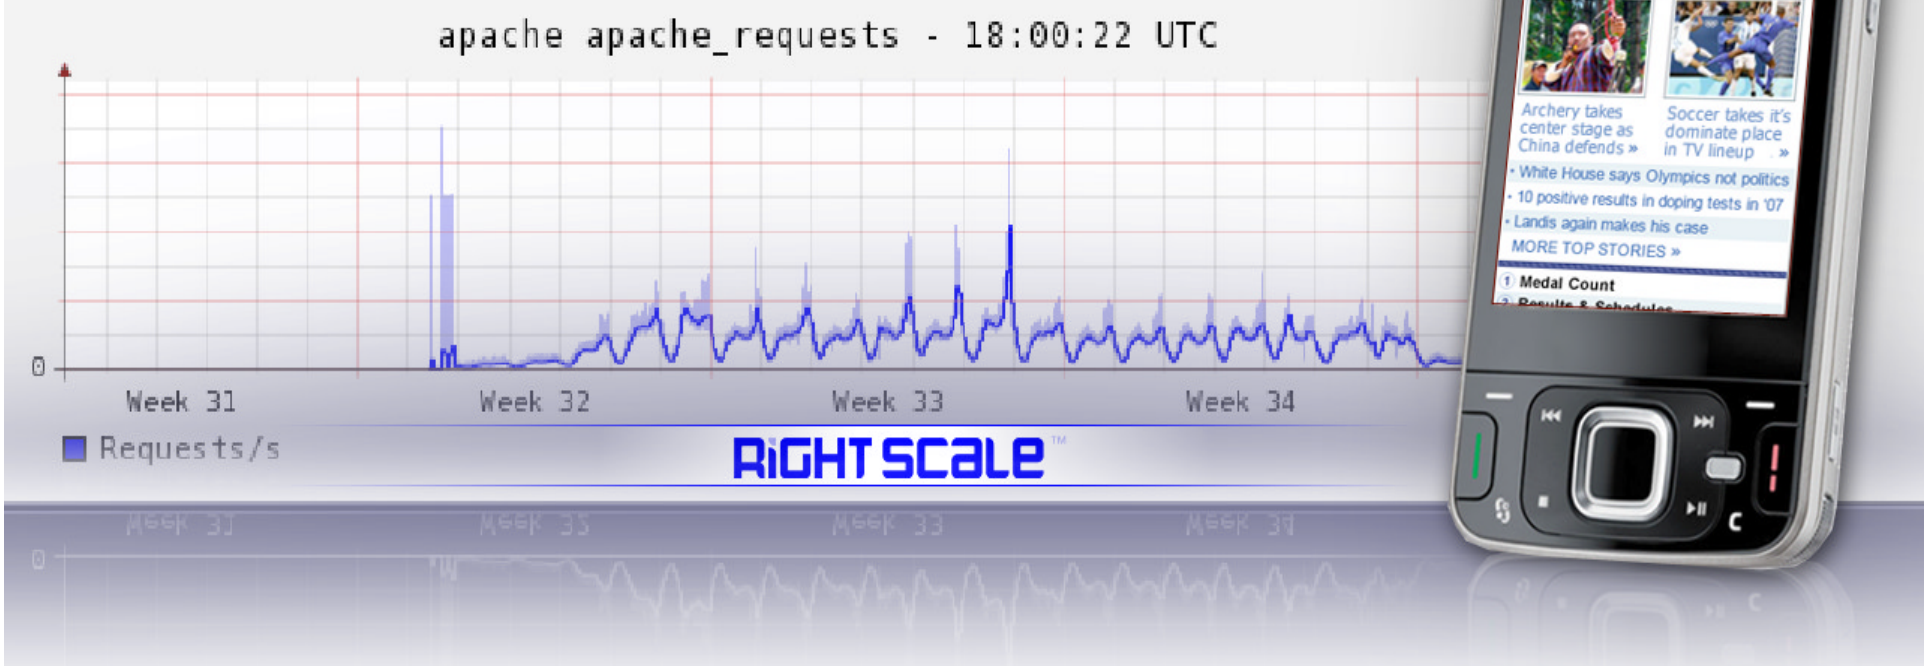

Starcut (Stocks)

- Daily scale up/down to save costs
- Global reach for services is needed (Europe, USA, Asia)
- Big name clients demand carrier-grade reliability
- Full 24/7, 365 support expectations to us and us to our vendors

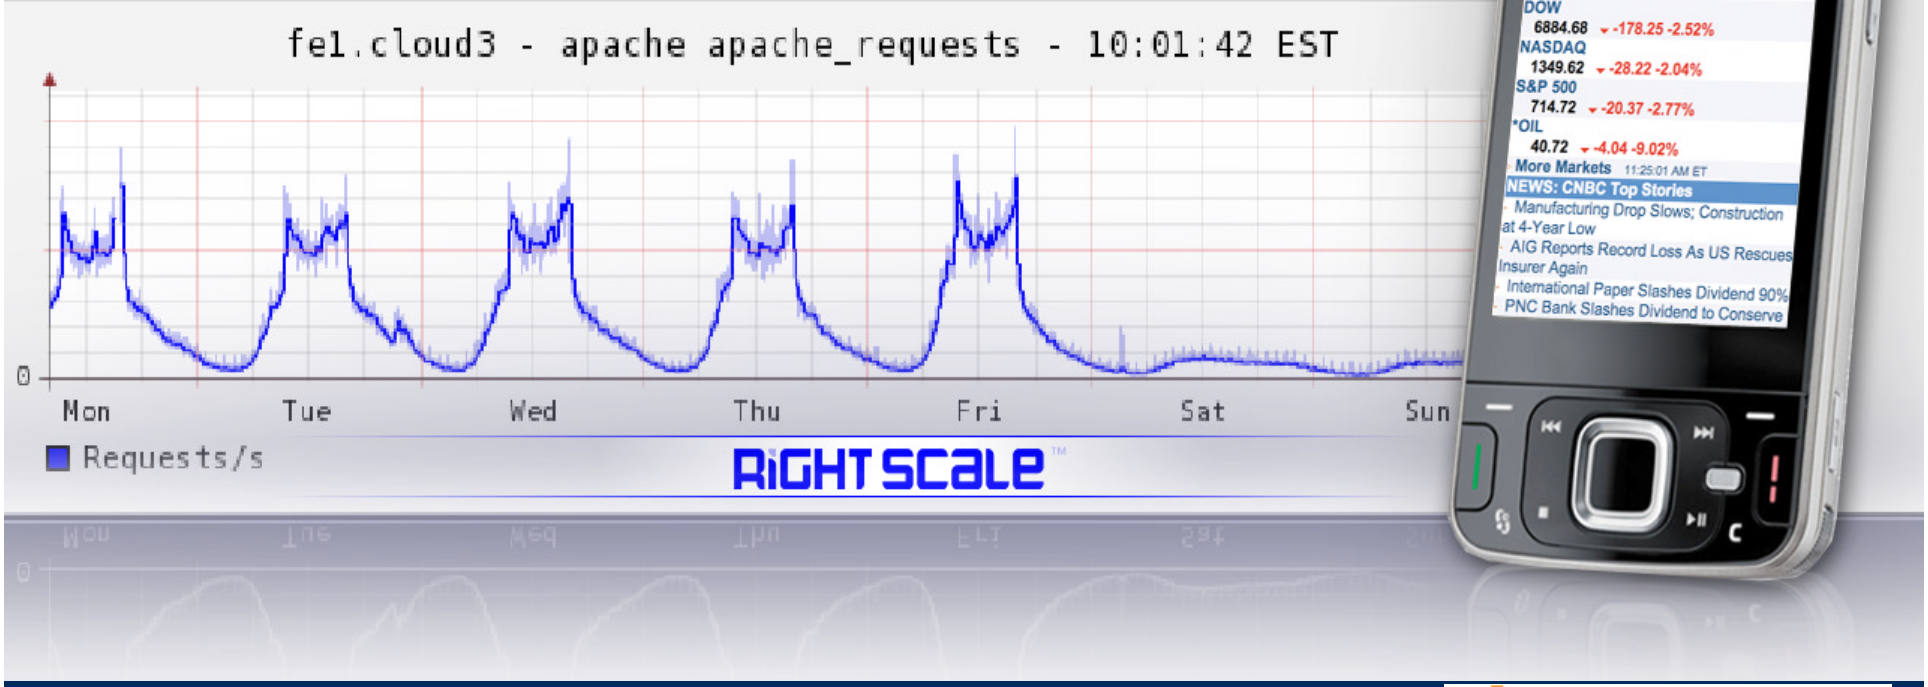

Std MySQL/EBS configuration
&
Storage Array in the Sky

Amazon EC2 Storage Architecture

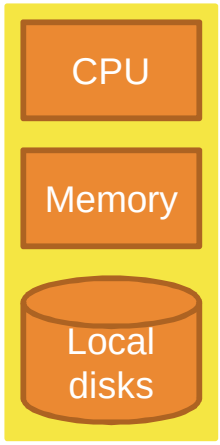

½ core … 8 cores — $0.10-$0.80/hr

1.7GB … 16GB

1 x 160GB … 4 x 420GB

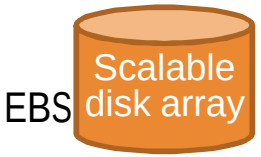

1GB … 1TB volumes — $0.10/GB/mo

1 … volumes / server — $0.10/M-io-ops

network attached, redundant

redundant back-end store — $0.15/GB/mo

S3 put/get semantics — $0.10/K-ops

Elastic Block Store (EBS)

- Create volume: 1GB – 1TB size
- Attach to any one instance
  in same availability zone
- Pay $0.10/GB per month
  + $0.10/million I/O ops
- Snapshot backup a volume to S3
- Create new volume from snapshot

Snapshots

- Point-in-time snapshot
- Written to S3
  - (Not to bucket!)

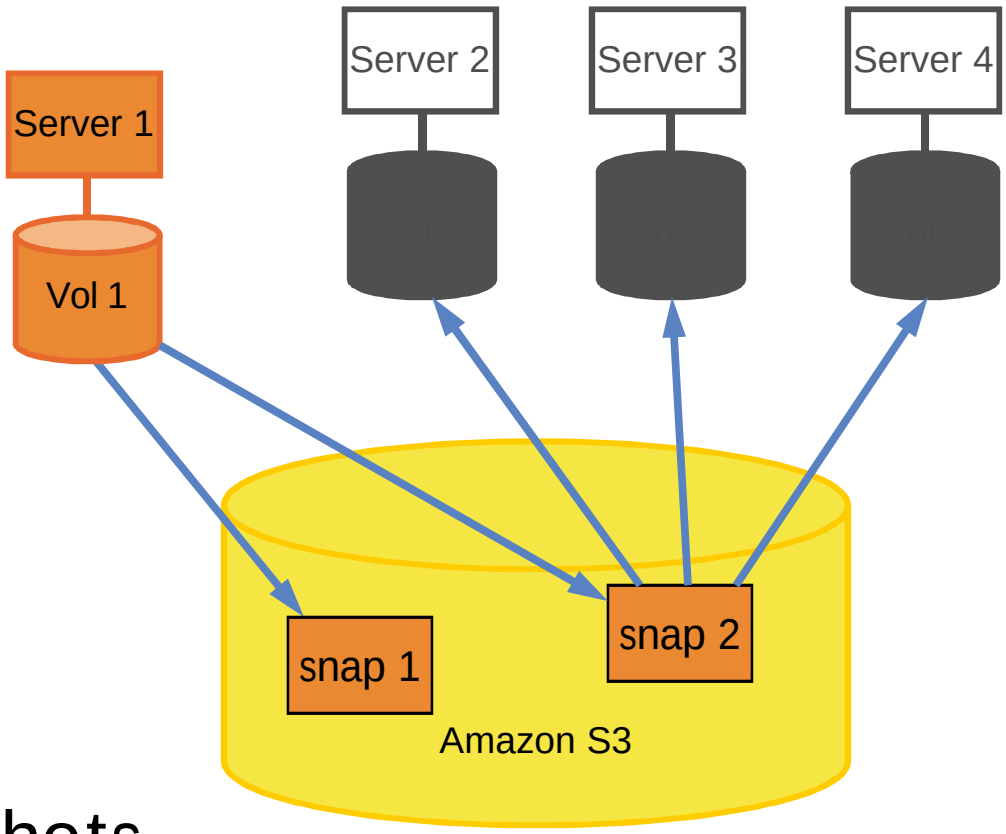

- Incremental snapshots
- Restore to new volume
  - Instantaneous (lazy restore)

Incremental Snapshots

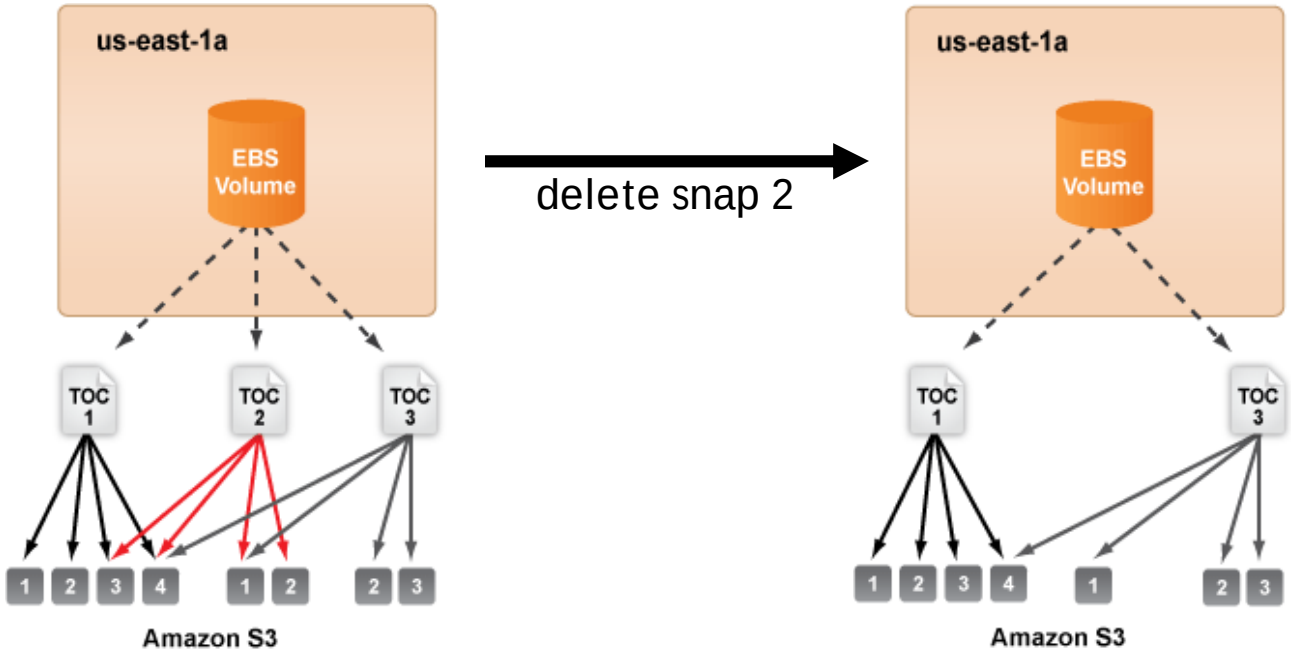

- Table of Contents vs. Data Blocks
- “Space used” difficult to gauge
- Frequent snapshots: minimal cost

Keeping Safe

- Freeze data while snapshotting
  - XFS filesystem has freeze command
  - But: 2.6.18 kernel bug causes hung filesystem (ouch!)
- Volumes are fragile in transit!
  - Unmounting can leave data in flight
  - Mounting mistake will be destructive
- Use snapshots for safety!
  - *Always* snapshot after unmount
  - Create fresh volume from snapshot, mount that, retry in case of problem

MySQL/EBS Replication

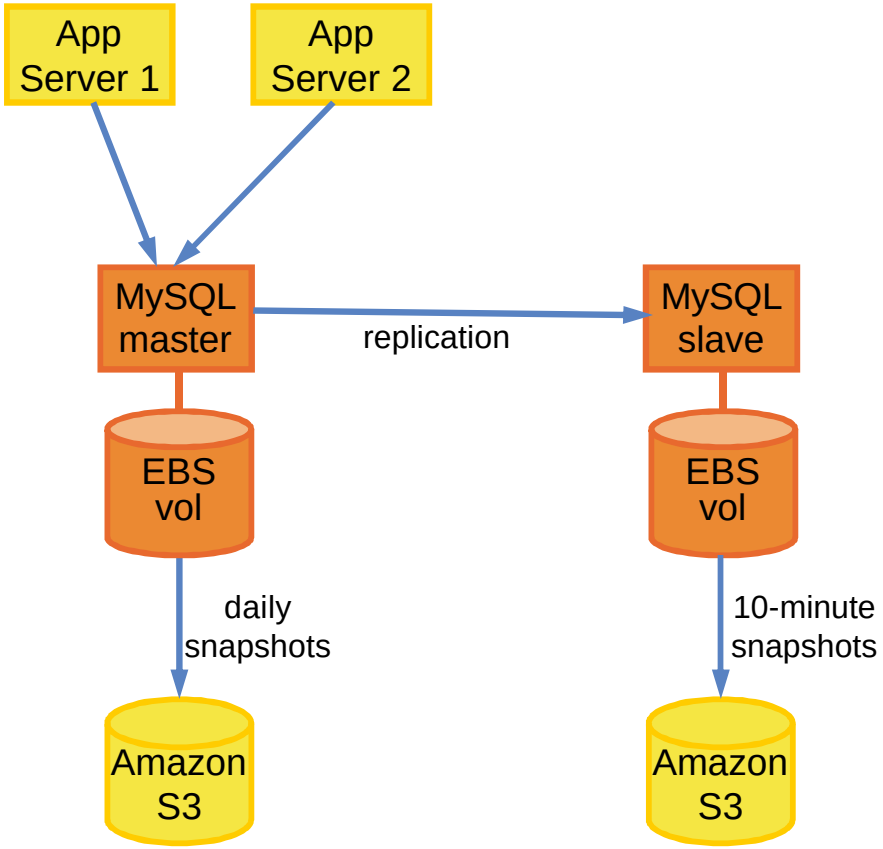

Database operations pain points

• Master hardware failure	no replica ‘til hw fixed
• Master software failure	no replica ‘til root caused
• Testing schema change	extra system: 99% idle
• Migrate to new hardware	2x rack space, SAN ports, …
• New DB requirement	machines, ports, disks, …
• QA, test, stage, DR, demo, …	always to few, always too idle

Failure? Allocate fresh resources!

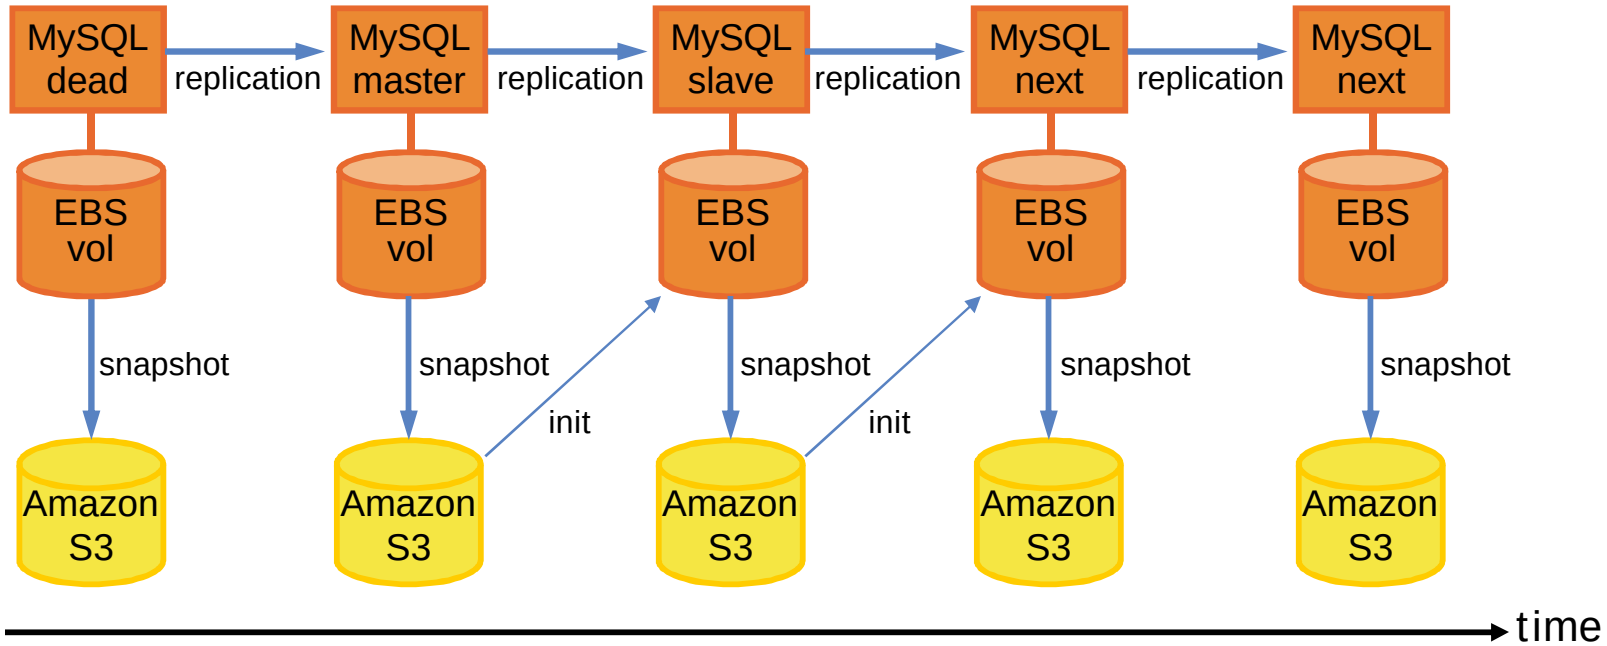

- Extra resources are free for short periods! (e.g. 4-core/16GB + 1TB EBS => <$30/day

Vertical Database Scaling

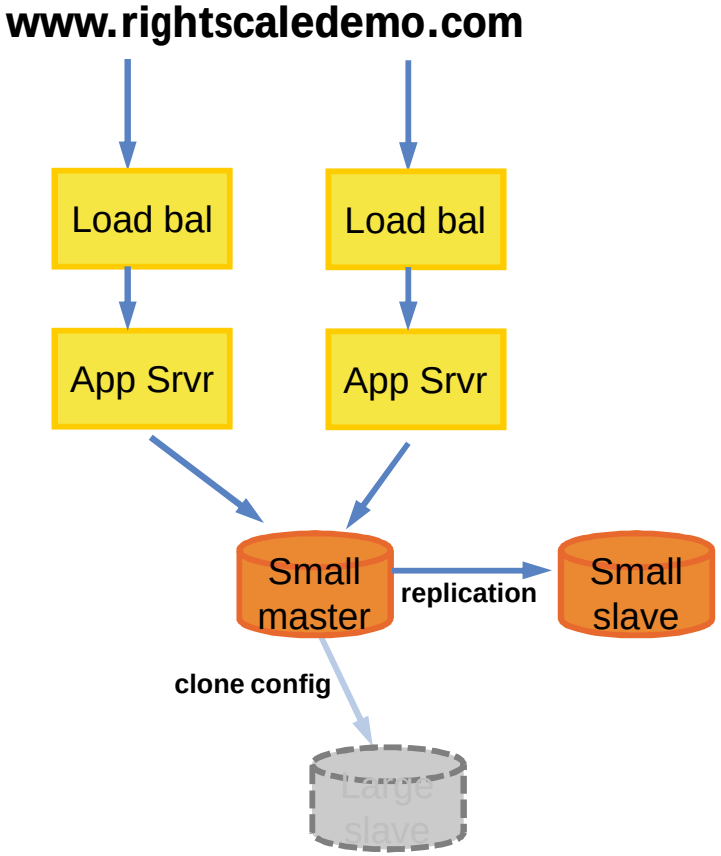

1 Clone Server Config

Vertical Database Scaling

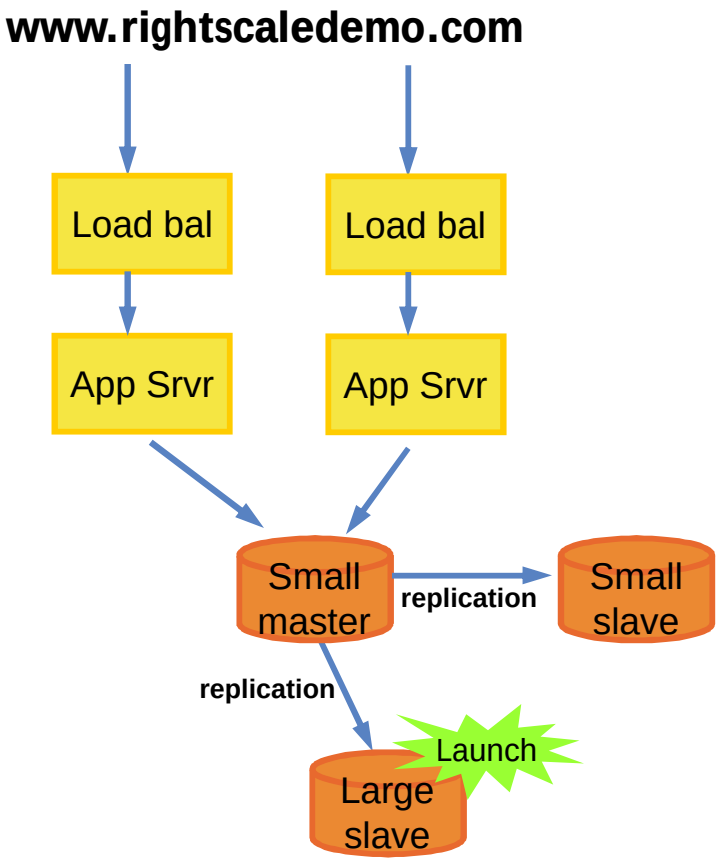

1. Clone Server Config
2. Launch large slave

Vertical Database Scaling

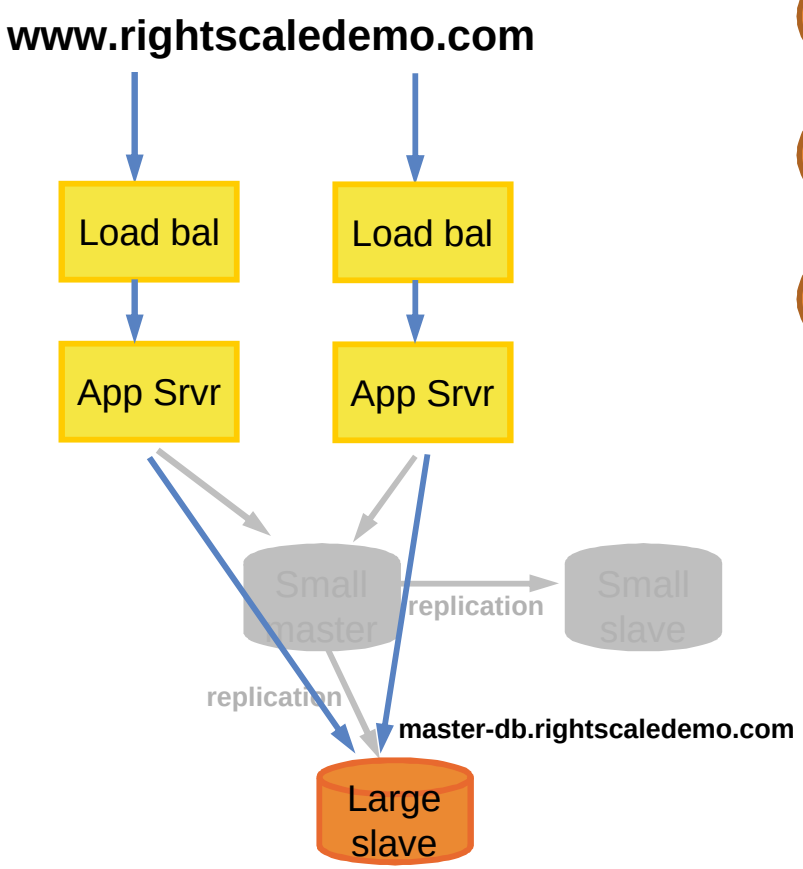

1. Clone Server
2. Launch large slave
3. Promote to master

Vertical Database Scaling

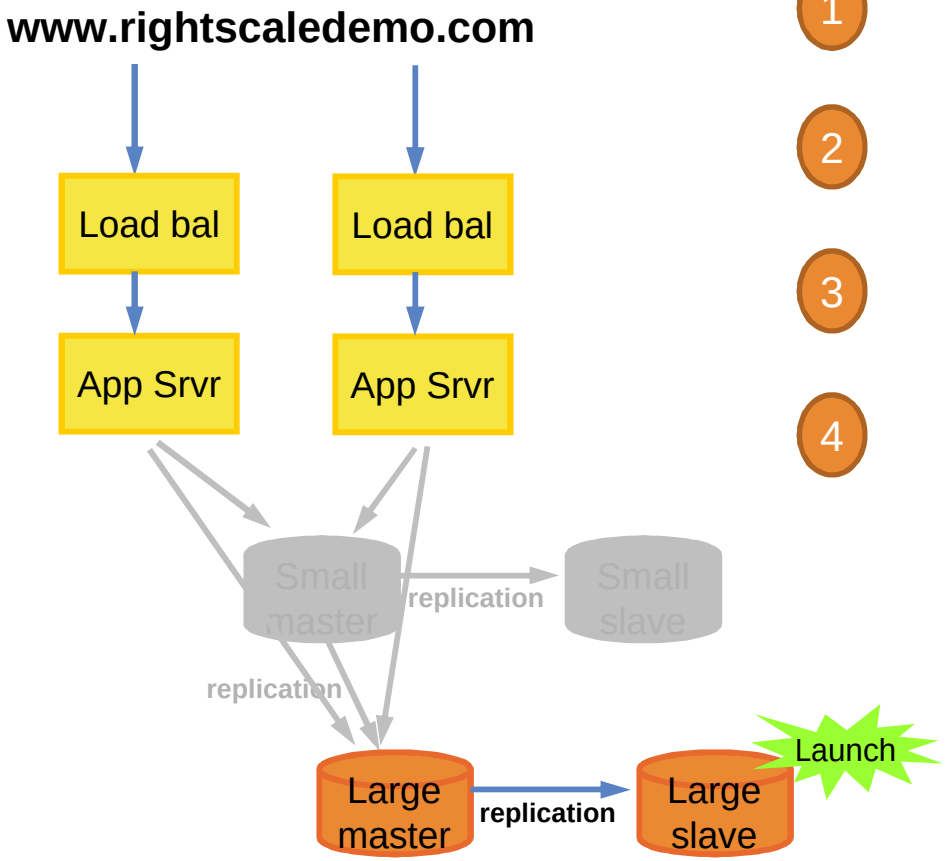

1. Clone Server
2. Launch large slave
3. Promote to master
4. Launch new slave

Watch demo in webinar recording at www.rightscale.com/webinars

Cloud Principle #1: grab the next one!

- Database management
  - Always roll forward, never fail-back
  - Animoto upgraded L -> XL @600 servers
- App server management
  - Always roll forward, never fix in-place
  - www1 bust? Launch www3! Then troubleshoot
- Good sysadmins have the urge to fix… ☹

MySQL/EBS Toolbox server

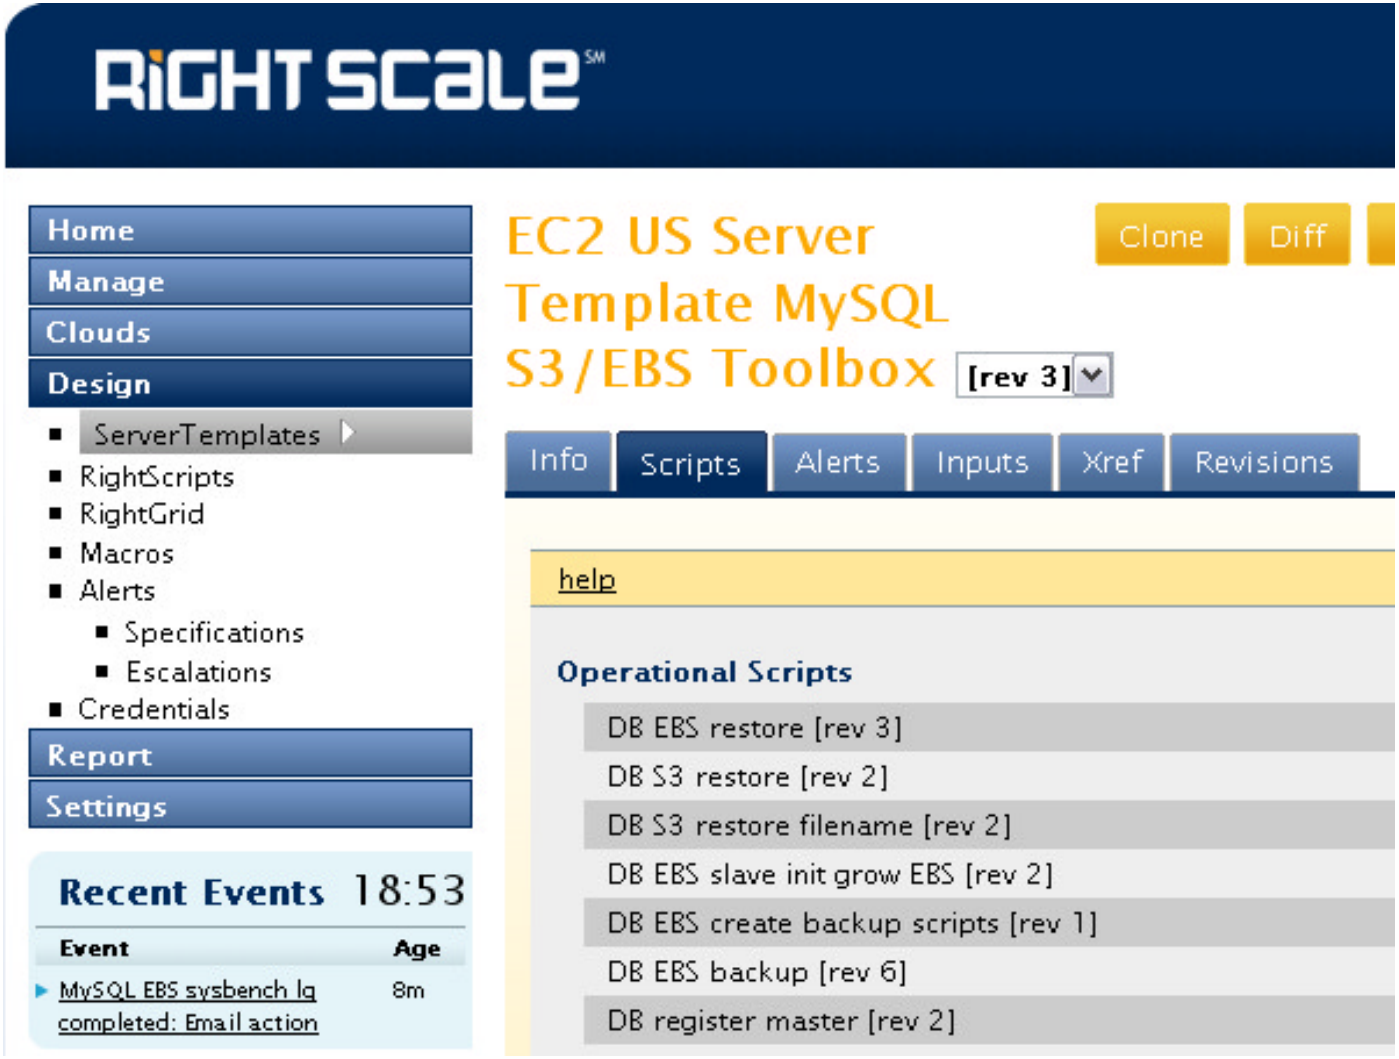

Many small databases

Database pain cycle:

1. Recognize the DB is important
2. Buy expensive DB hardware
3. Consolidate DBs onto expensive hardware
4. Recognize the DB is now even more critical
5. Buy more expensive DB hardware
6. Repeat!

Cloud Computing = Infinite servers and disks
+ Fully automatable

Many Small Databases

- Sharding
  - Split by user
  - Need to manage mappings

- Read slaves
  - Many slave to increase read performance
  - Need to manage failover

EBS Performance

- Functionality (snapshots) trump performance
- Very difficult to benchmark
  - Multiple levels of cache, some hidden
- Sequential I/O limited by network
  - 70-90MB/sec typ, 100MB/sec max
  - 4 local drives striped yield more
- Random I/O (sustained)
  - 400-1000 read ops/sec
  - 100-400 write ops/sec
- Stripe to increase performance

Clouds are Multiplying!

- Amazon US + EU
- GoGrid, Flexiscale
- Eucalyptus open source release
- RackSpace CloudServers
- Sun cloud announcement
- And many more…

- Operating principles are the same
  - Storage architectures differ
  - From better local drives to traditional SANs

Public vs. Private Cloud

- Public cloud:
  - shared cloud computing infrastructure
  - connected to the public internet
- Private cloud:
  - cloud infrastructure dedicated to a single party
  - may not be connected to the public internet
- Hybrid cloud
  - private and public clouds used together
  - leverage the benefits of both

Cloud Definition:

1. Infinite computing resources
2. Available on-demand
3. Pay per use
4. Delivered across the Internet

Hybrid Cloud Use-Cases

- *Develop in public, deploy in private*
  - Outsourced/distributed development
  - flexibility of test & dev resources
- *Develop in private, deploy in public*
  - existing in house dev resources
  - ability to scale in public cloud
  - public-facing app
- *Private core, public expansion*
  - Transaction intensive in-house core
  - Batch processing in public cloud

RightScale Supports Hybrid Clouds

1. Commandeer some machines
2. Install Eucalyptus
3. Register with RightScale
4. The power of RightScale for your cloud
5. Invite friends to your cloud

Multi-Cloud Architecture

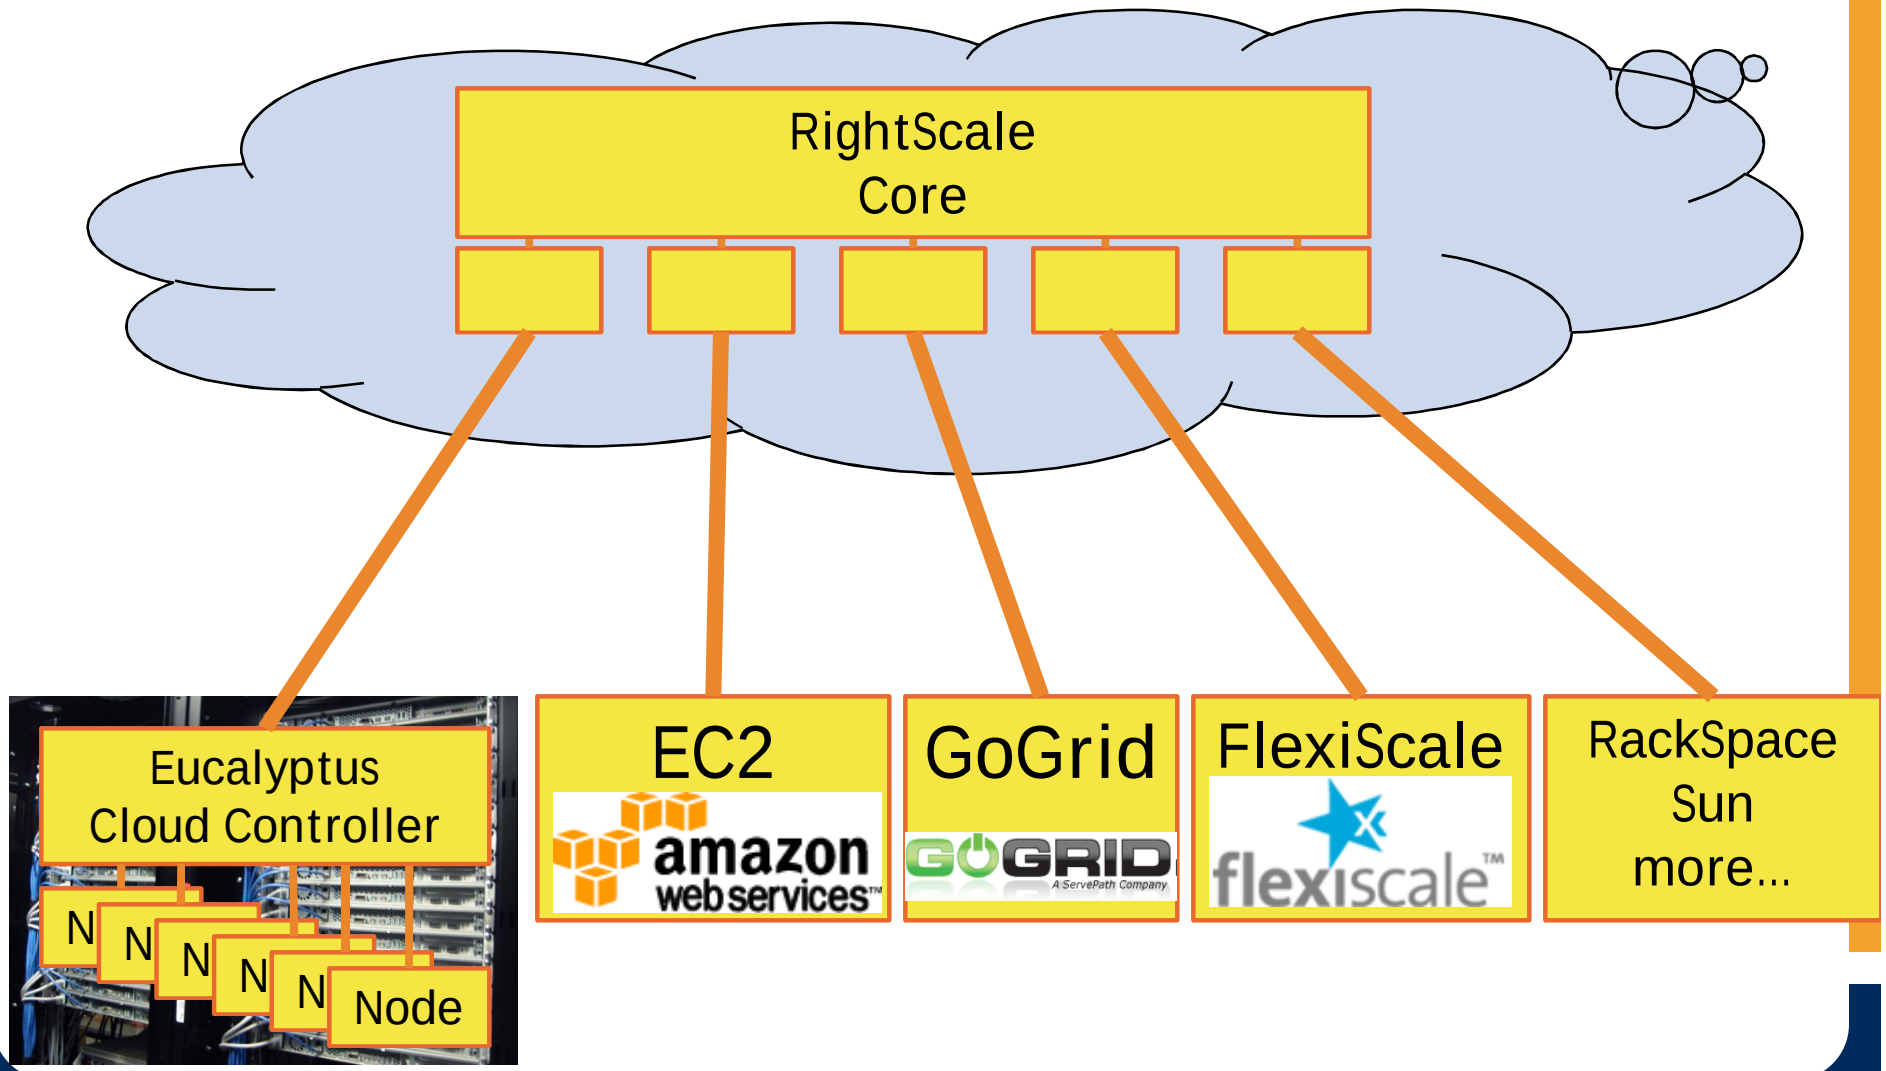

Clouds Surpass Traditional Hosting

- Servers on-demand
- Disk Volumes on-demand
- Multiple datacenters
- Reassign IPs / VIPs across datacenters
- Multiple geographies
- Multiple providers

- **Automation tying it all together**

RightScale
Cloud Management Platform

- Automation
- Multi-Cloud
- Fully transparent
- Fully customizable
- Replicated MySQL/EBS ServerTemplates