Software Architecture for the Cloud

Michael Nygard
www.michaelnygard.com

In This Session

- Taxonomy
- Scenarios
- Examples

Cloud	1) All types of rapidly provisioned, on-demand computing or application services. 2) Any service delivered across the Internet, without requiring hardware or software installation.
Grid	Large scale computing services, typically heterogeneous, distributed, and parallel.
Utility	The “pay for what you use” billing method. May or may not be “cloudlike” in other ways.
Infrastructure as a Service (IaaS)	Virtualized server and network hardware, provisioned by customer control via portal or API.
Platform as a Service (PaaS)	Execution environment with software abstractions over the capabilities and operational management of the underlying infrastructure.
Software as a Service (SaaS)	Application features provided without hardware or software installation. May or may not execute on any of the above environments.

Cloud	1) All types of rapidly provisioned, on-demand computing or application services. 2) Any service delivered across the Internet, without requiring hardware or software installation.
Grid	Large scale computing services, typically heterogeneous, distributed, and parallel.
Utility	The “pay for what you use” billing method. May or may not be “cloudlike” in other ways.
Infrastructure as a Service (IaaS)	Virtualized server and network hardware, provisioned by customer control via portal or API.
Platform as a Service (PaaS)	Execution environment with software abstractions over the capabilities and operational management of the underlying infrastructure.
Software as a Service (SaaS)	Application features provided without hardware or software installation. May or may not execute on any of the above environments.

Grid Computing

- Massively Parallel
- Move Data to Computers
- Move Computation to Data

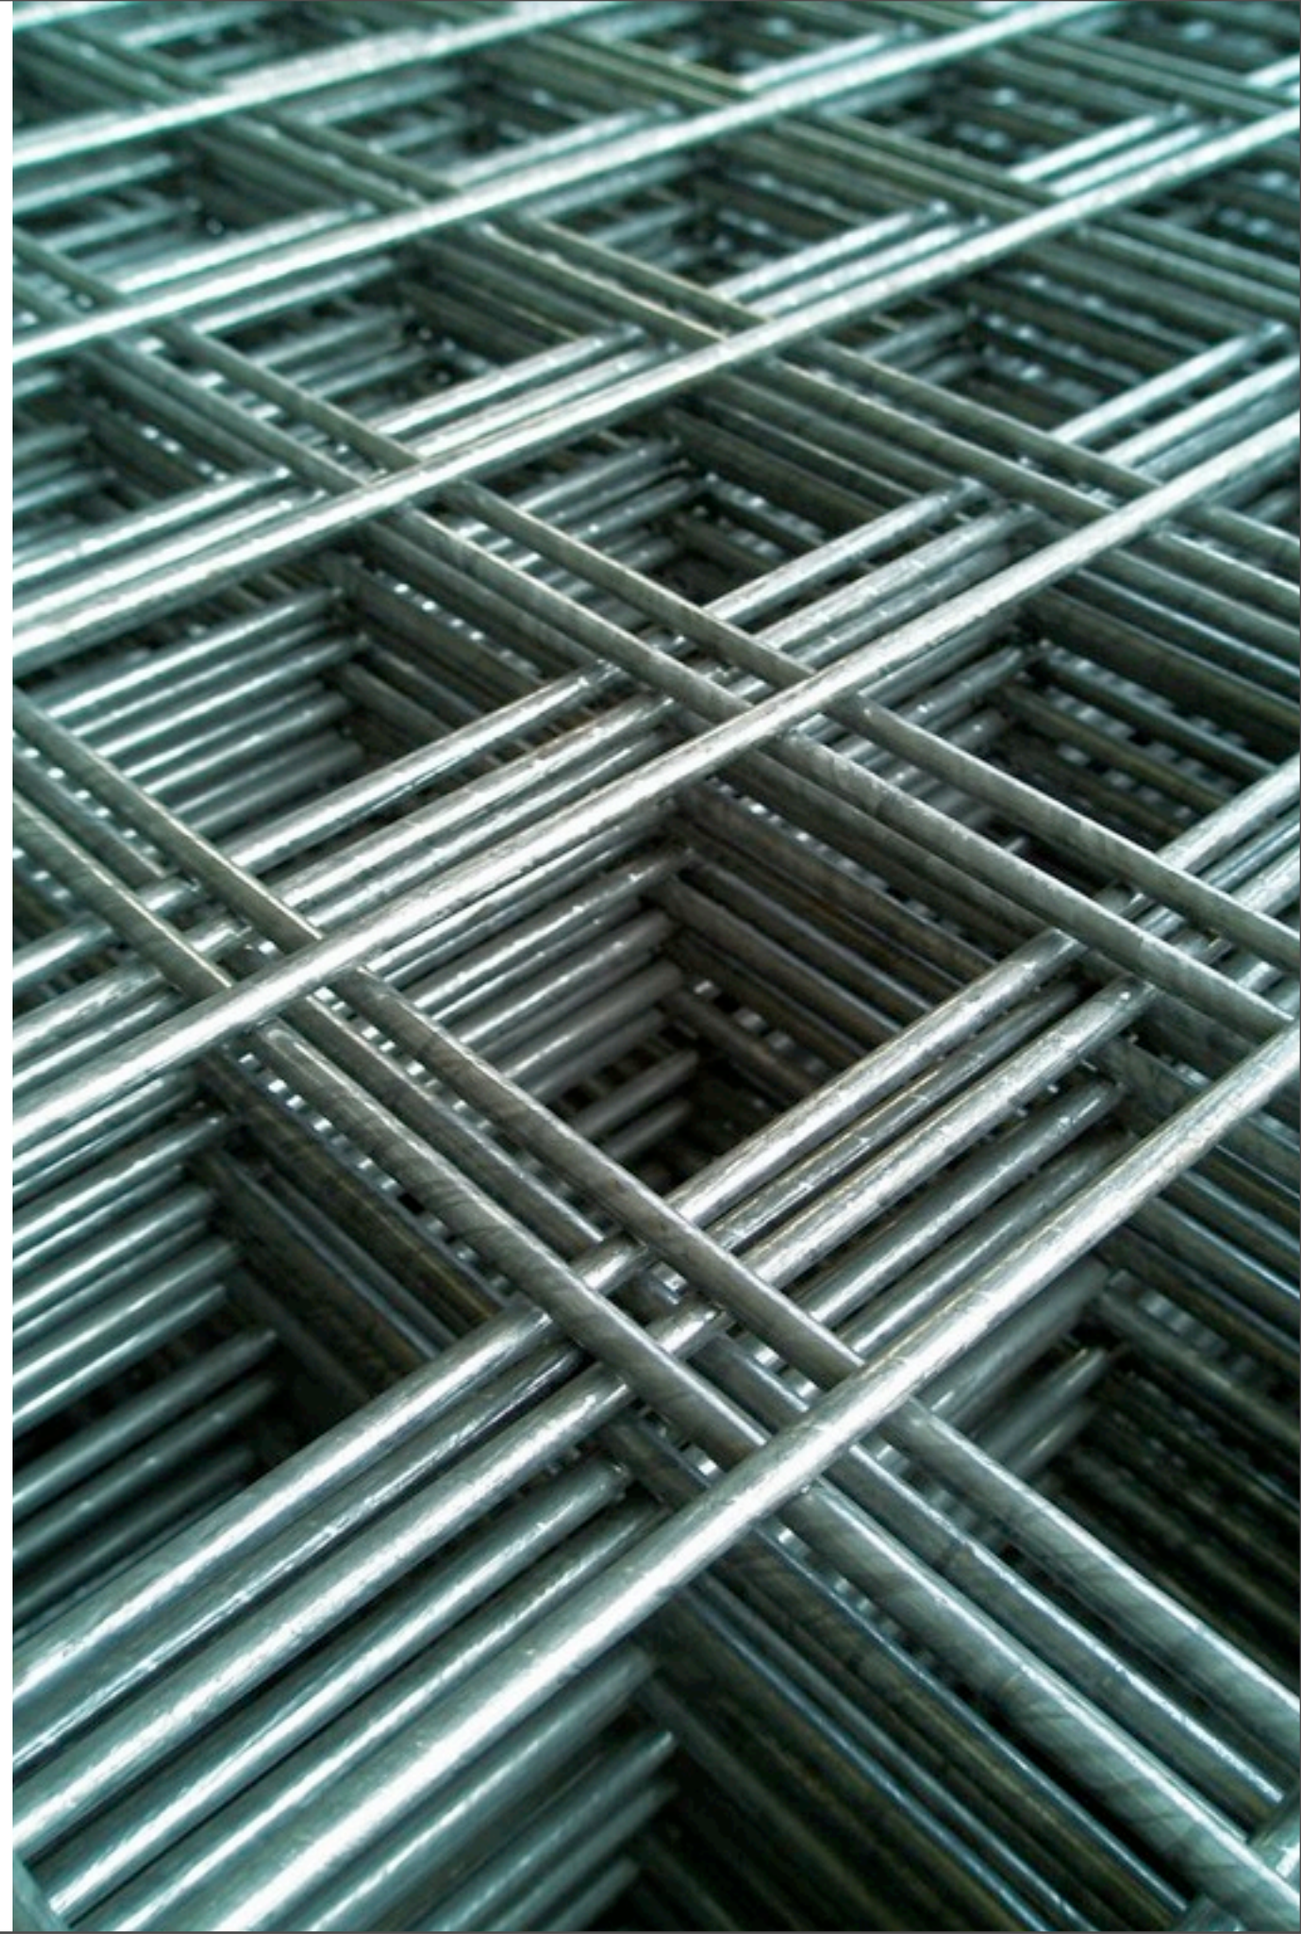

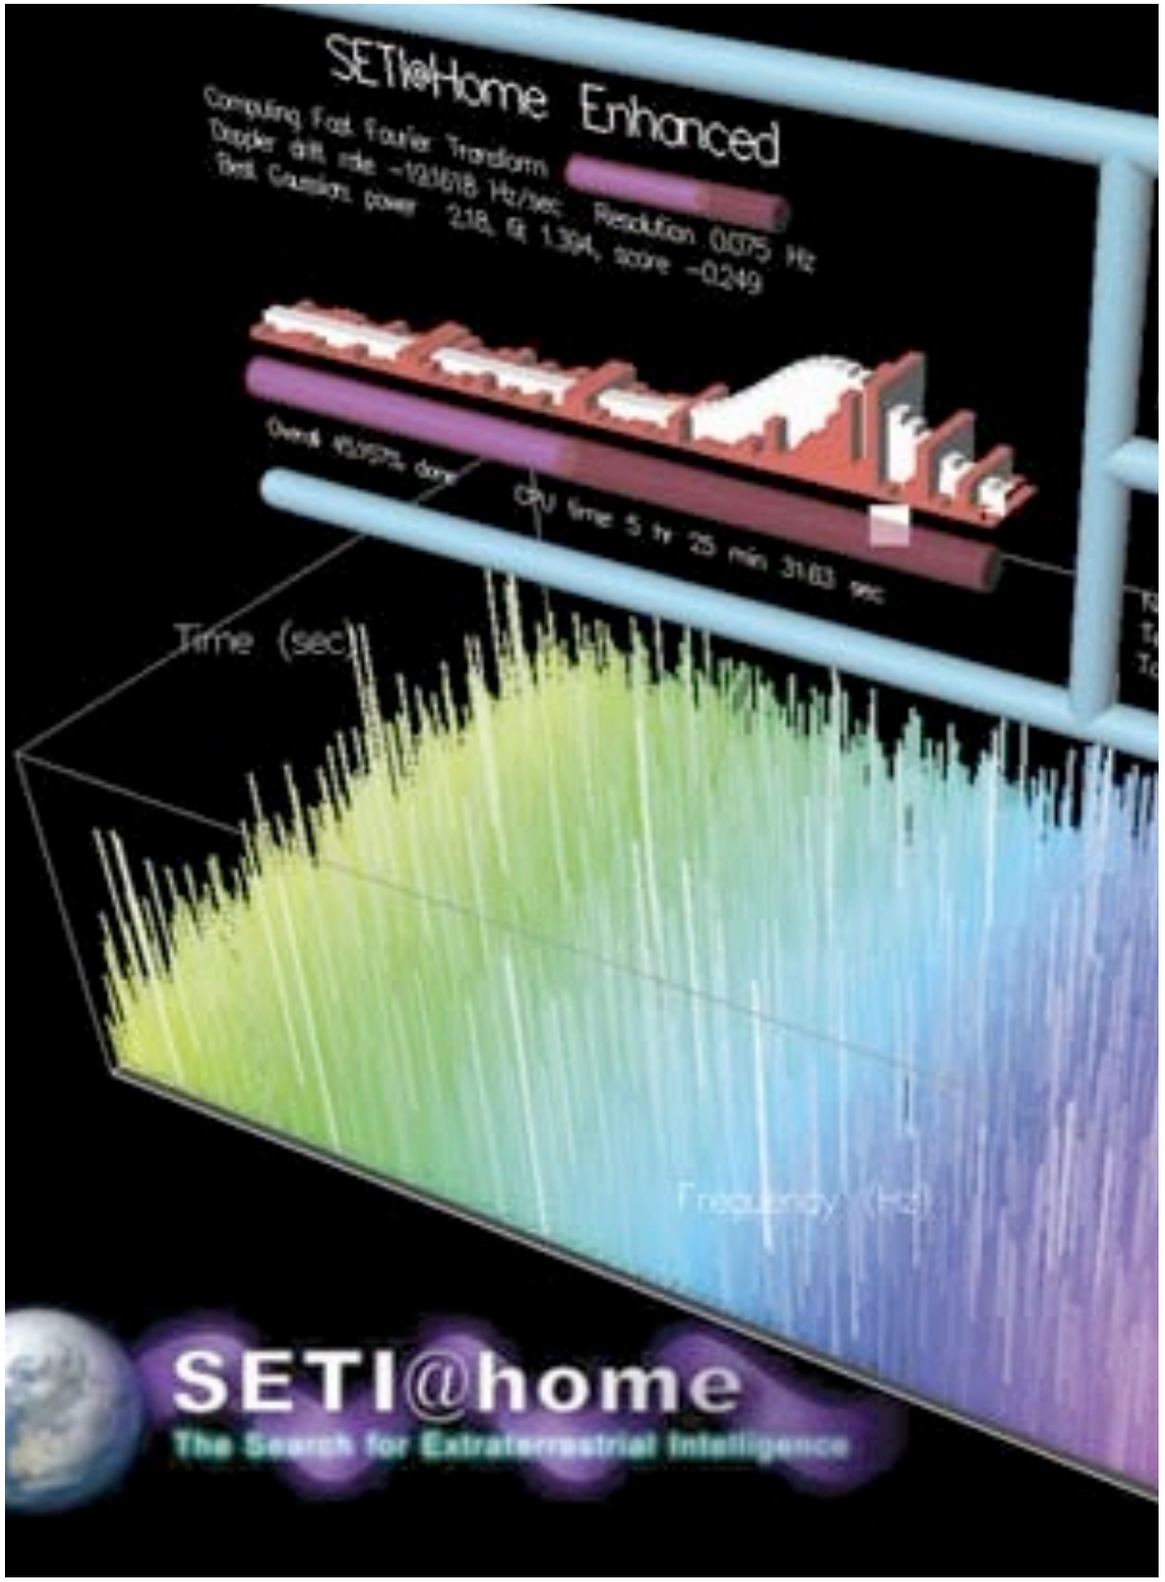

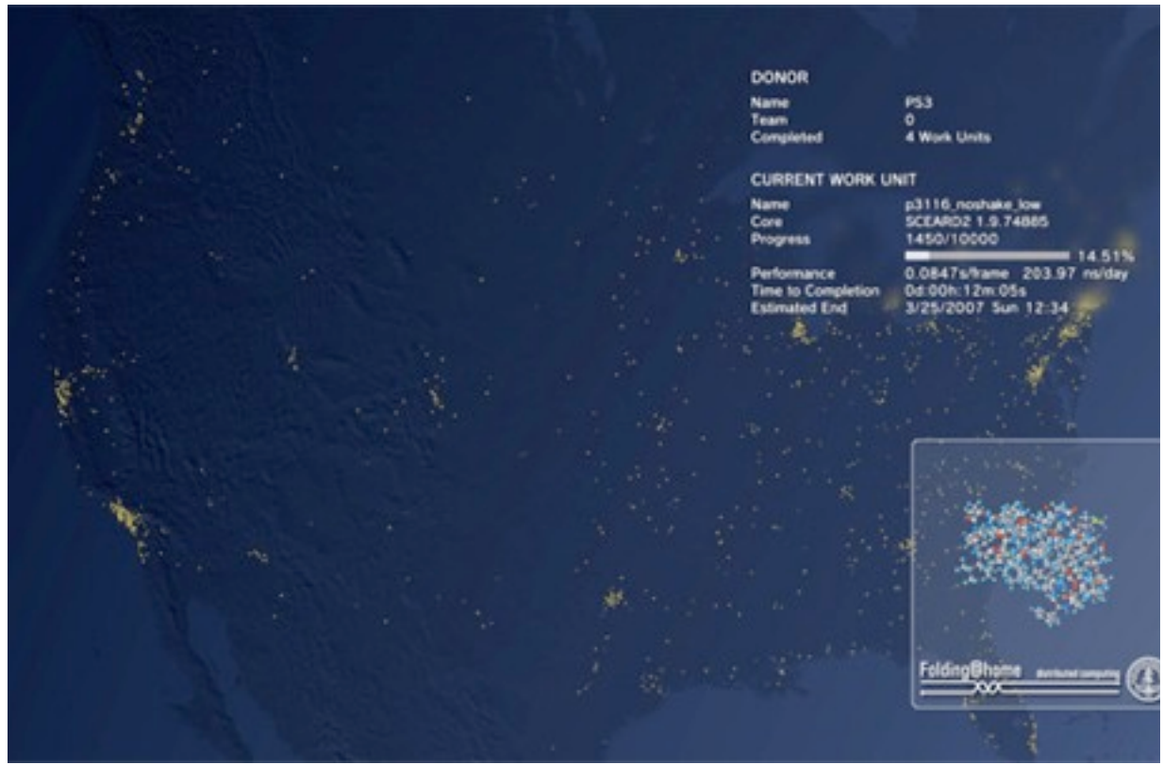

Resource Sharing via Federation

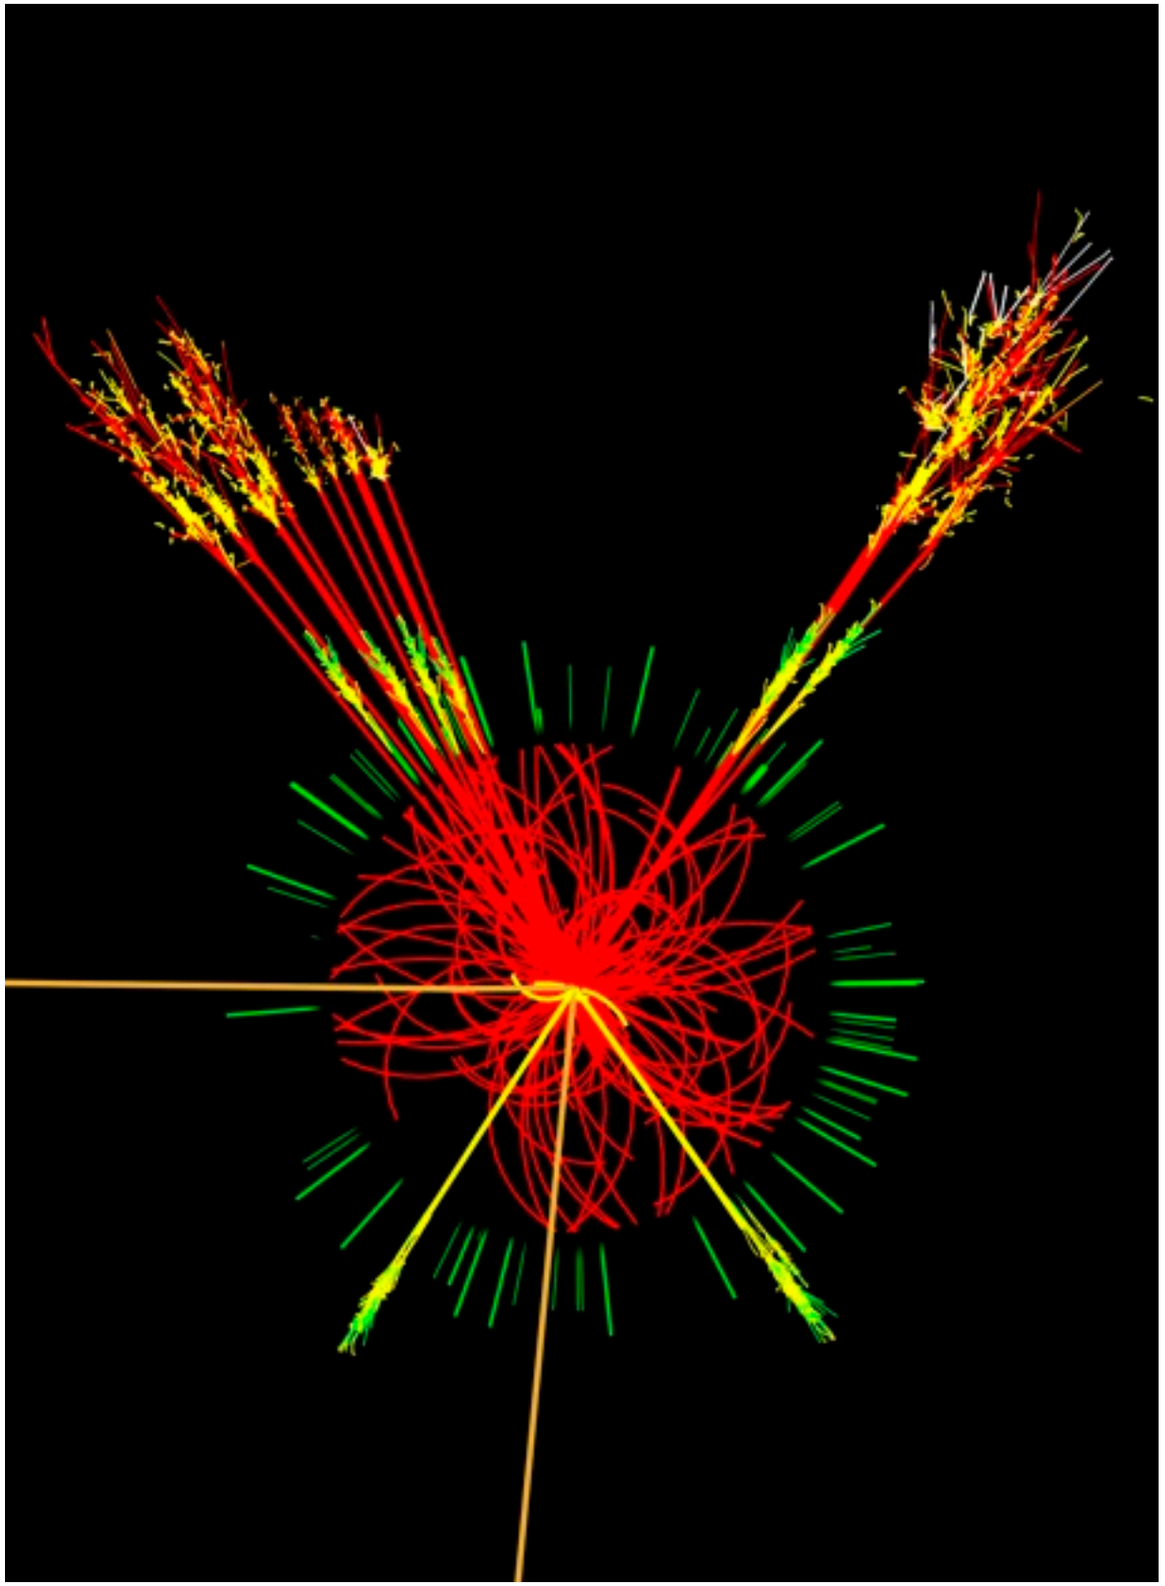

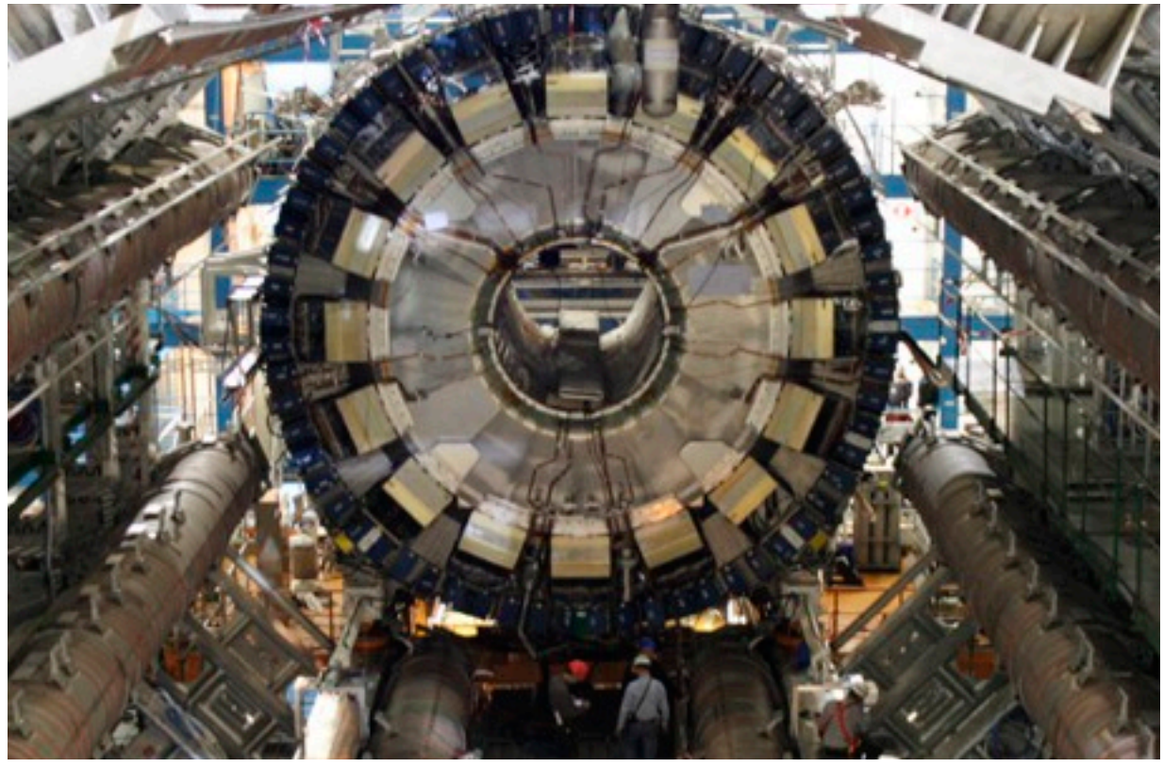

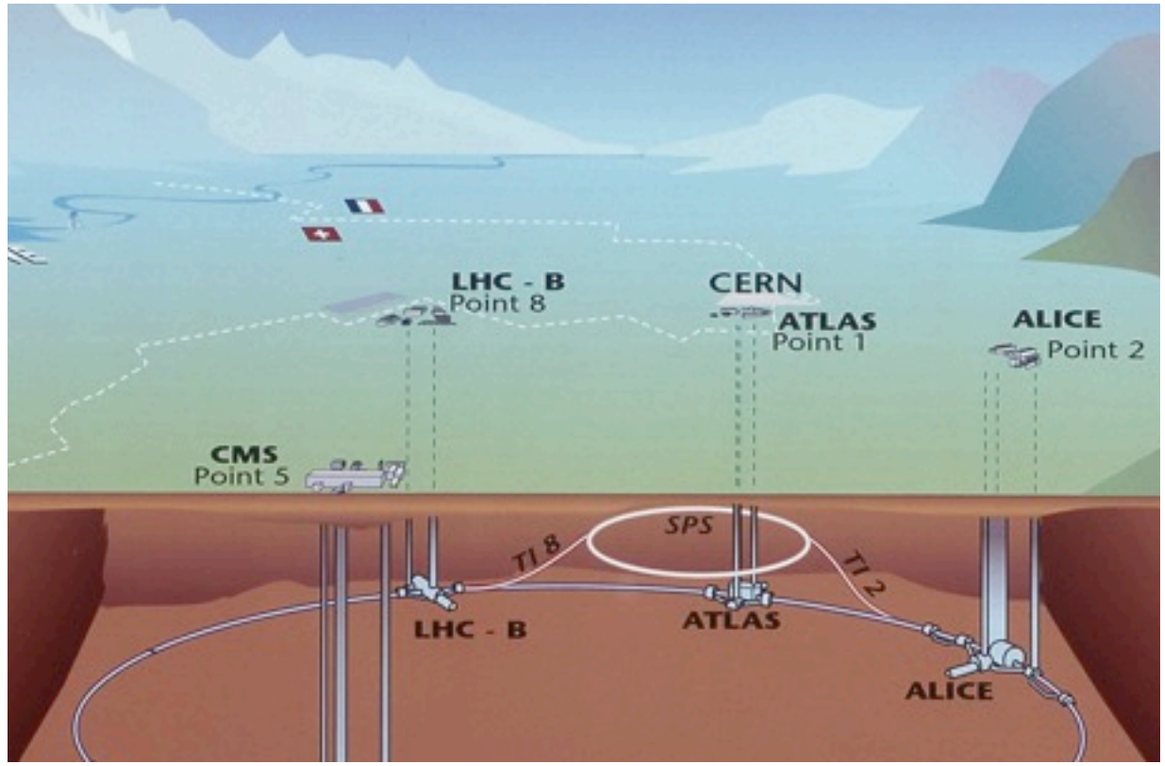

LHC - 15 Terabytes per Collision

Utility Computing

- Just a Billing Method
- Pay for what you use
- Variable demand (usually predictable)
- Financial efficiency
- Ceiling usually capped
- Pricing usually tiered

Holiday Volumes

- Overspending 10 out of 12 months
- Prefer rapid expansion before peak season

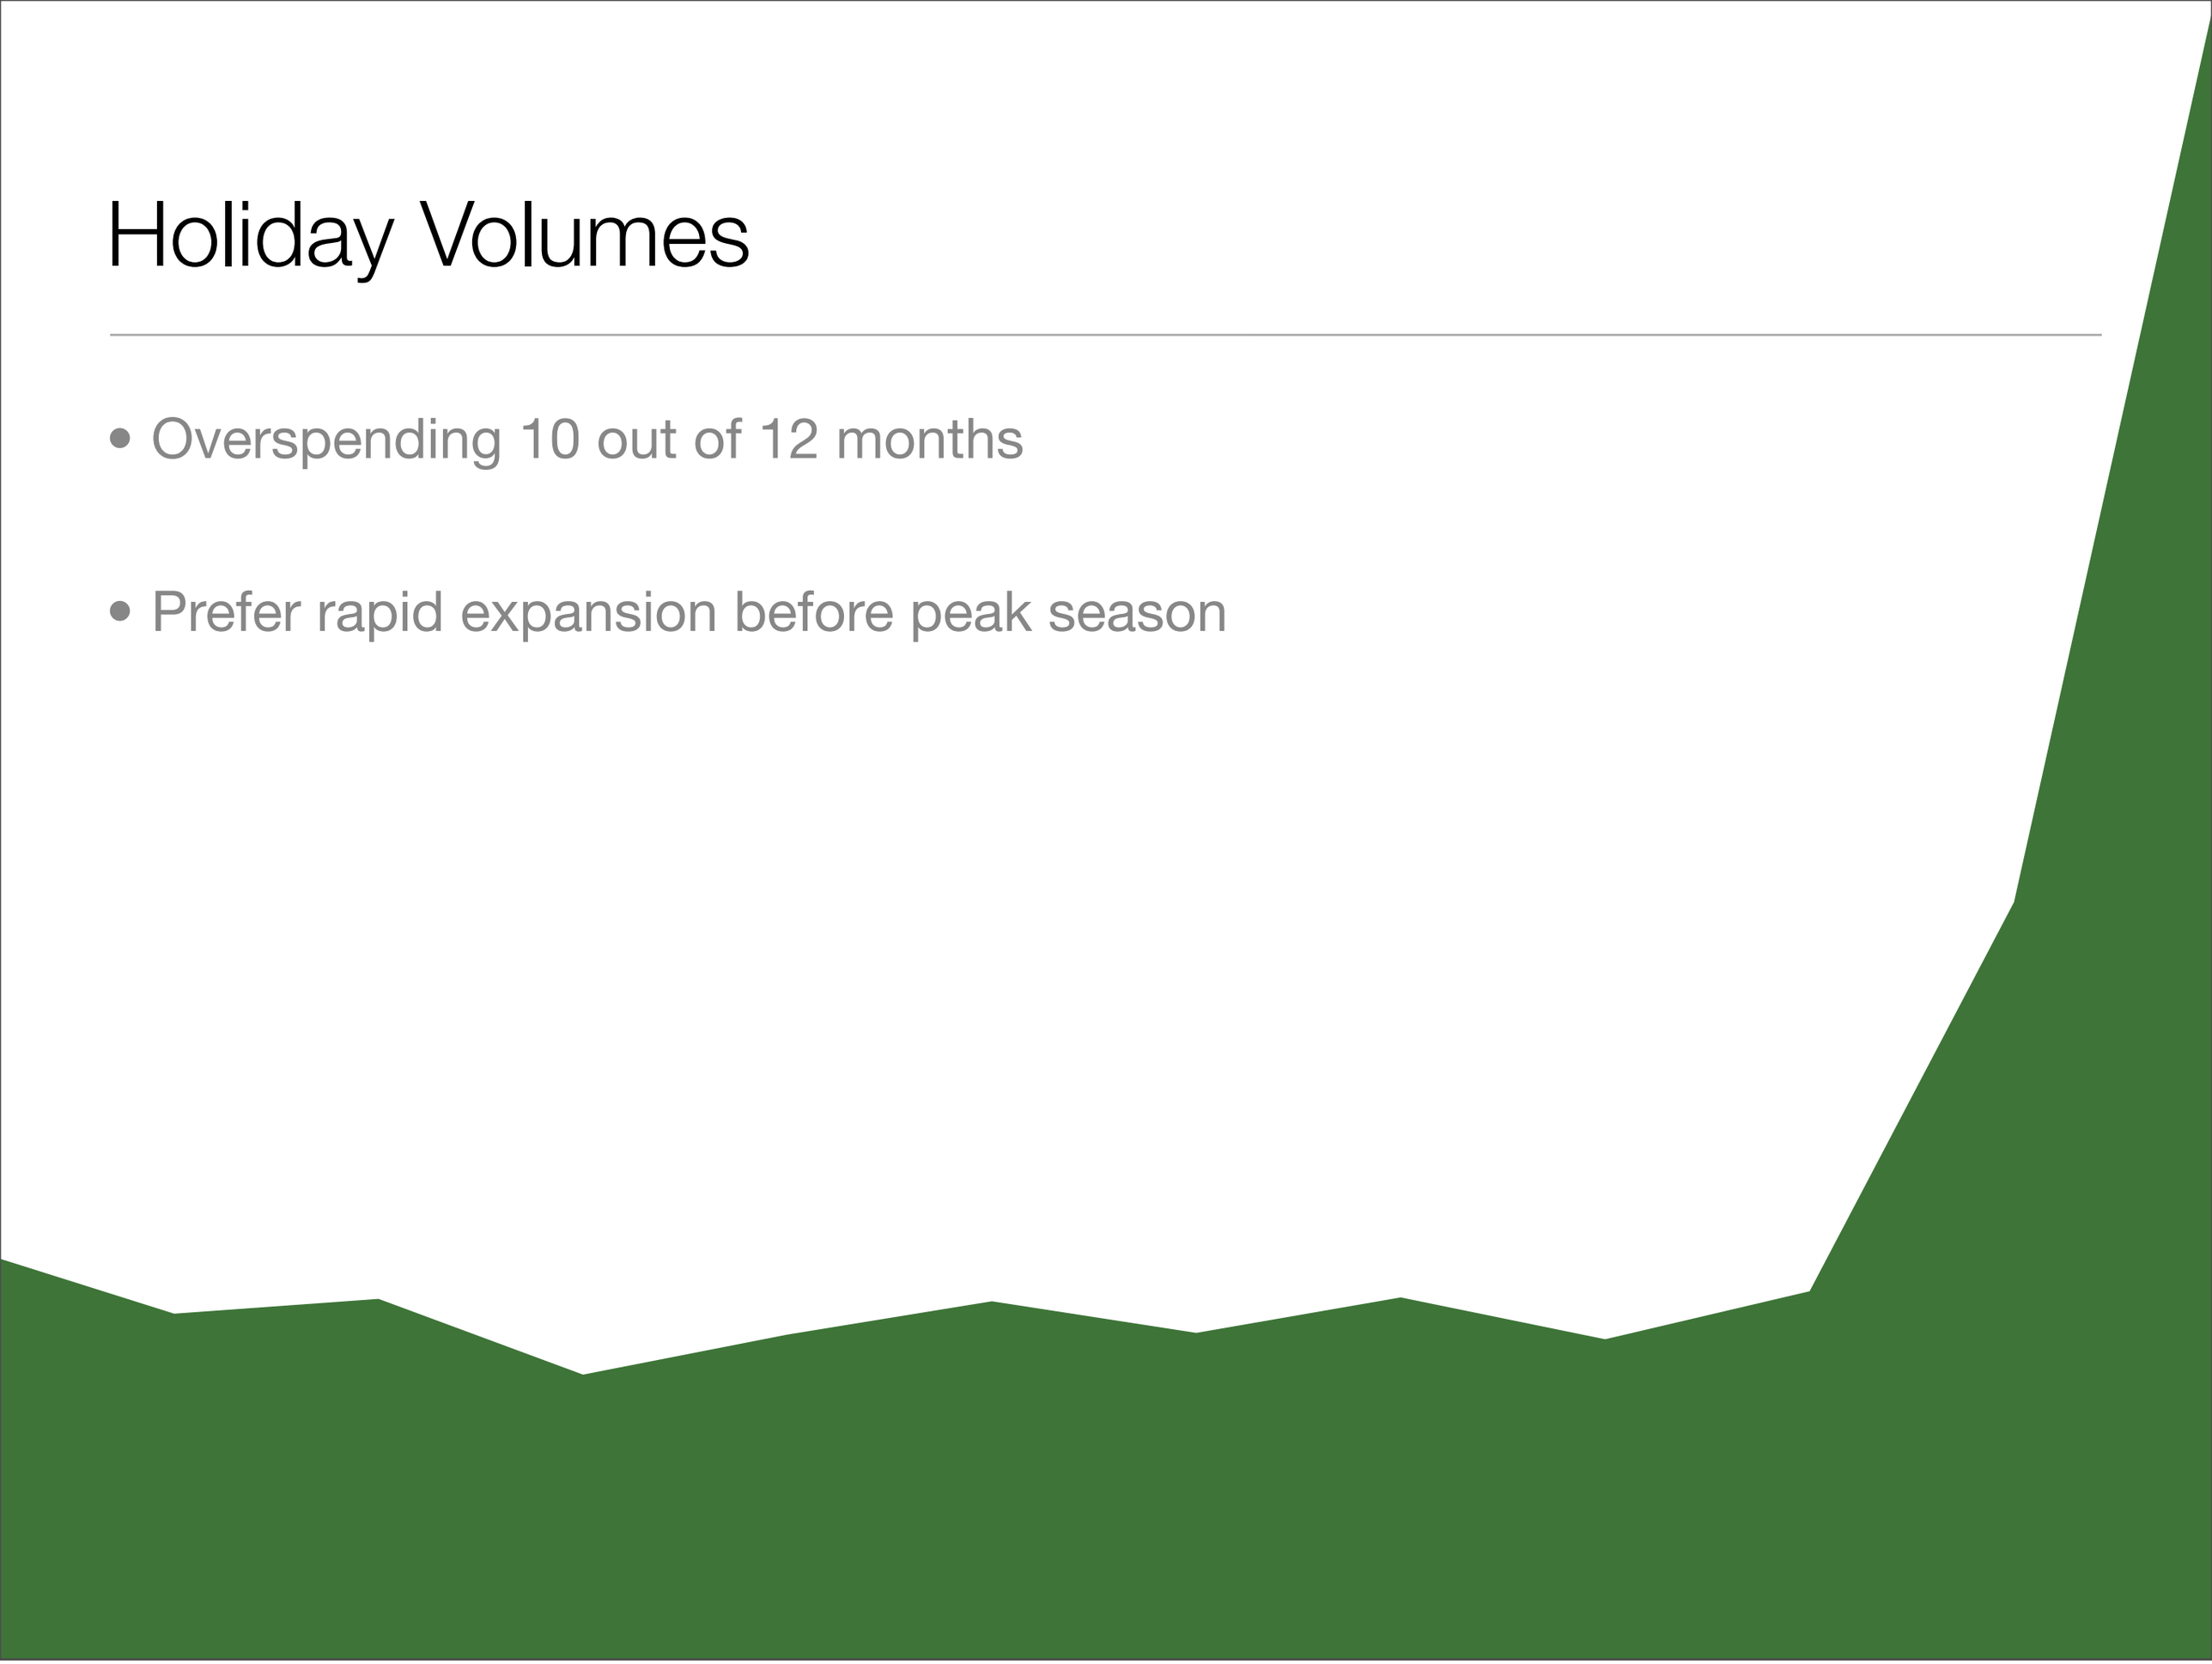

Cloud Computing

Infrastructure as a Service

Four Trends Leading to Clouds

- Virtualization
- Commoditization of hardware
- Horizontally scalable architecture
- Rapid provisioning

Virtualization

- Hardware abstraction
- Allows any operating system to run
- Increased hardware utilization

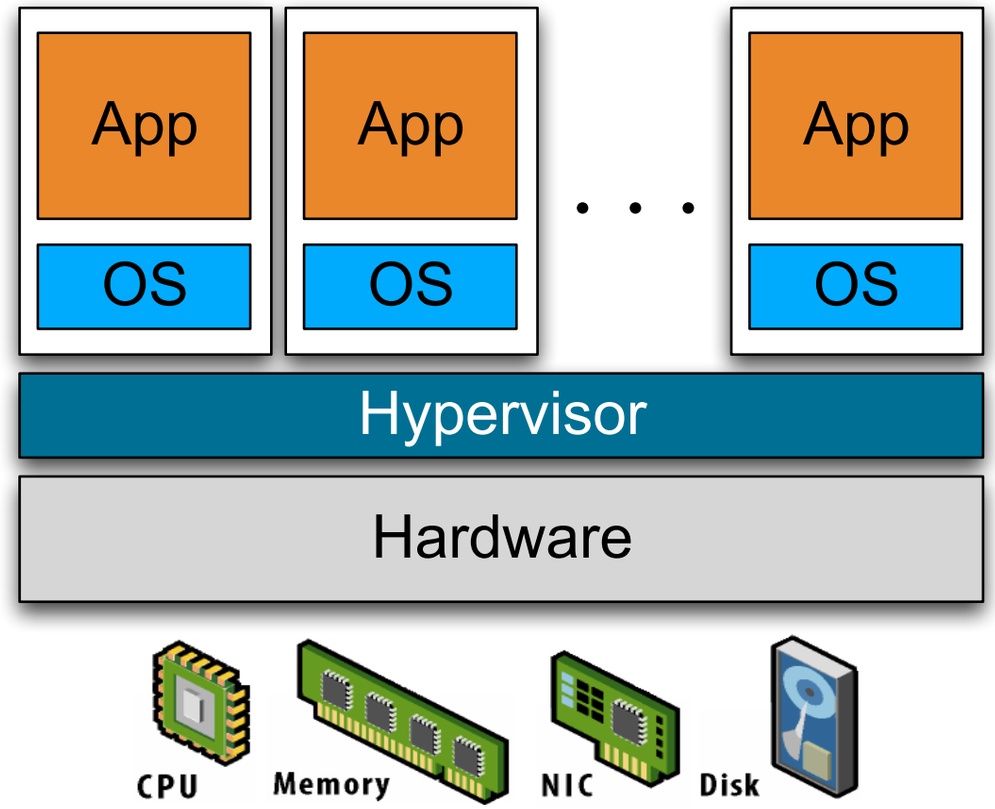

Virtualization

CONSISTENCY!

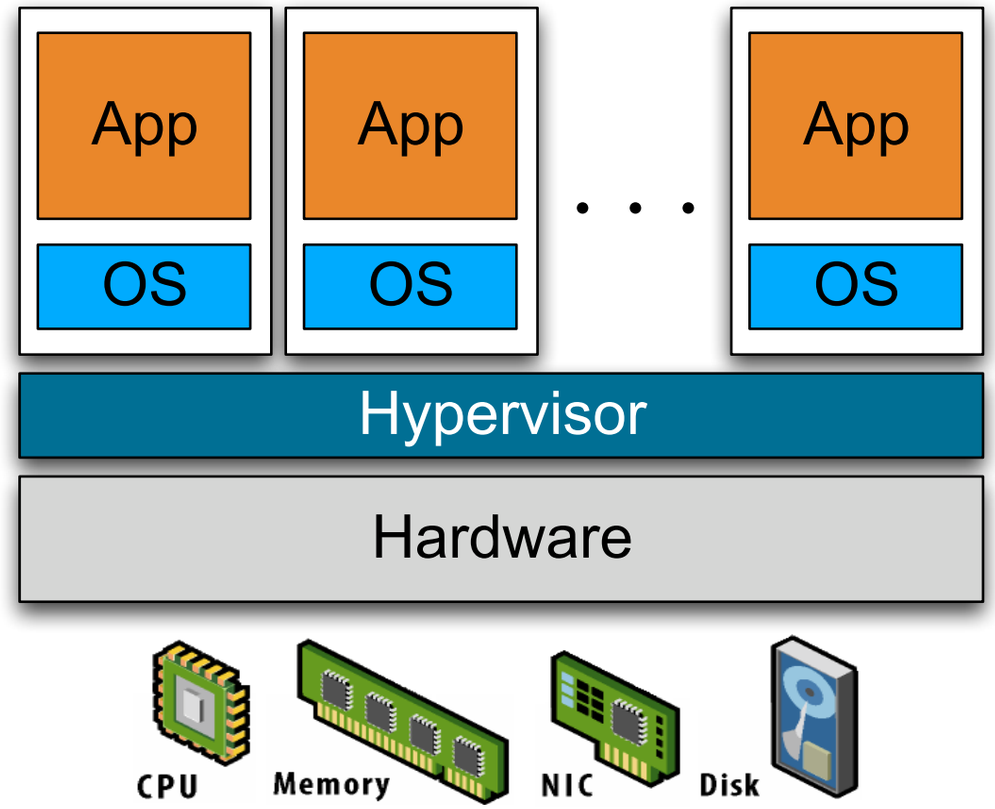

Virtualization

- Factors out variations in devices and drivers
- Applications can be as messy as they like
- Applications can be locked to their own version of the app server, JVM, or OS
- Applications don’t need to share resources or play nicely with others
- Administrators can move virtual machines for hardware maintenance or capacity optimization.

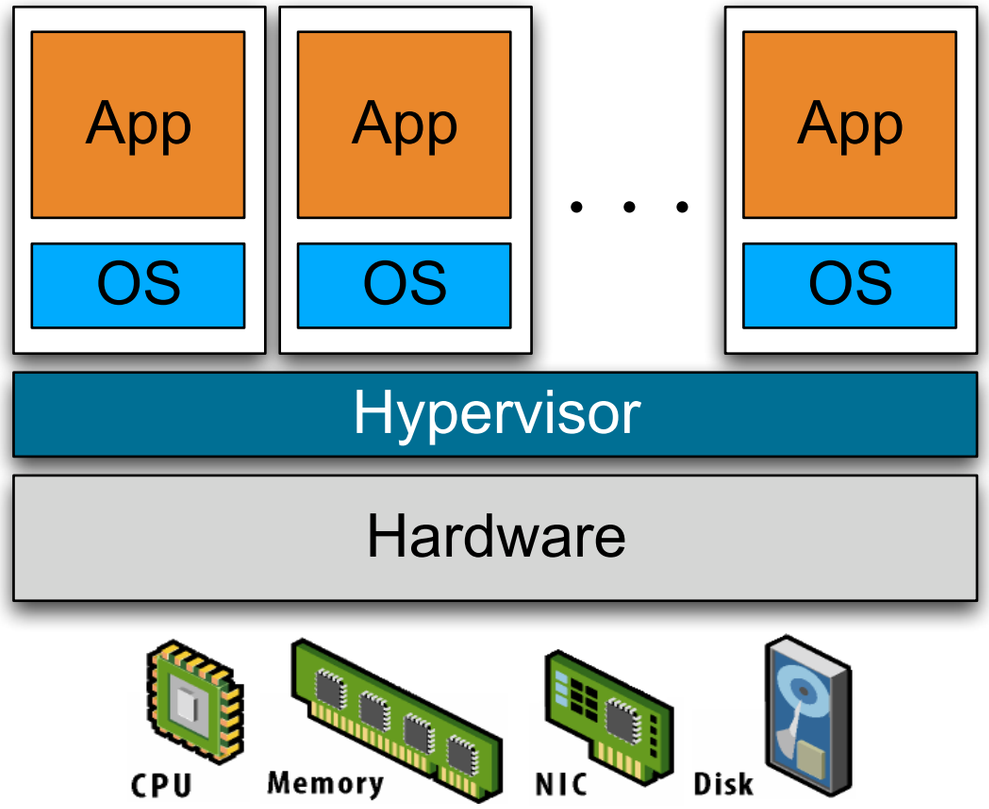

Commoditization

- HP 7410
  16 CPUs
  32 GB RAM
  $250,000

- Dell M600
  8 cores
  32 GB RAM
  $5,800

43 M600s for 1 7410

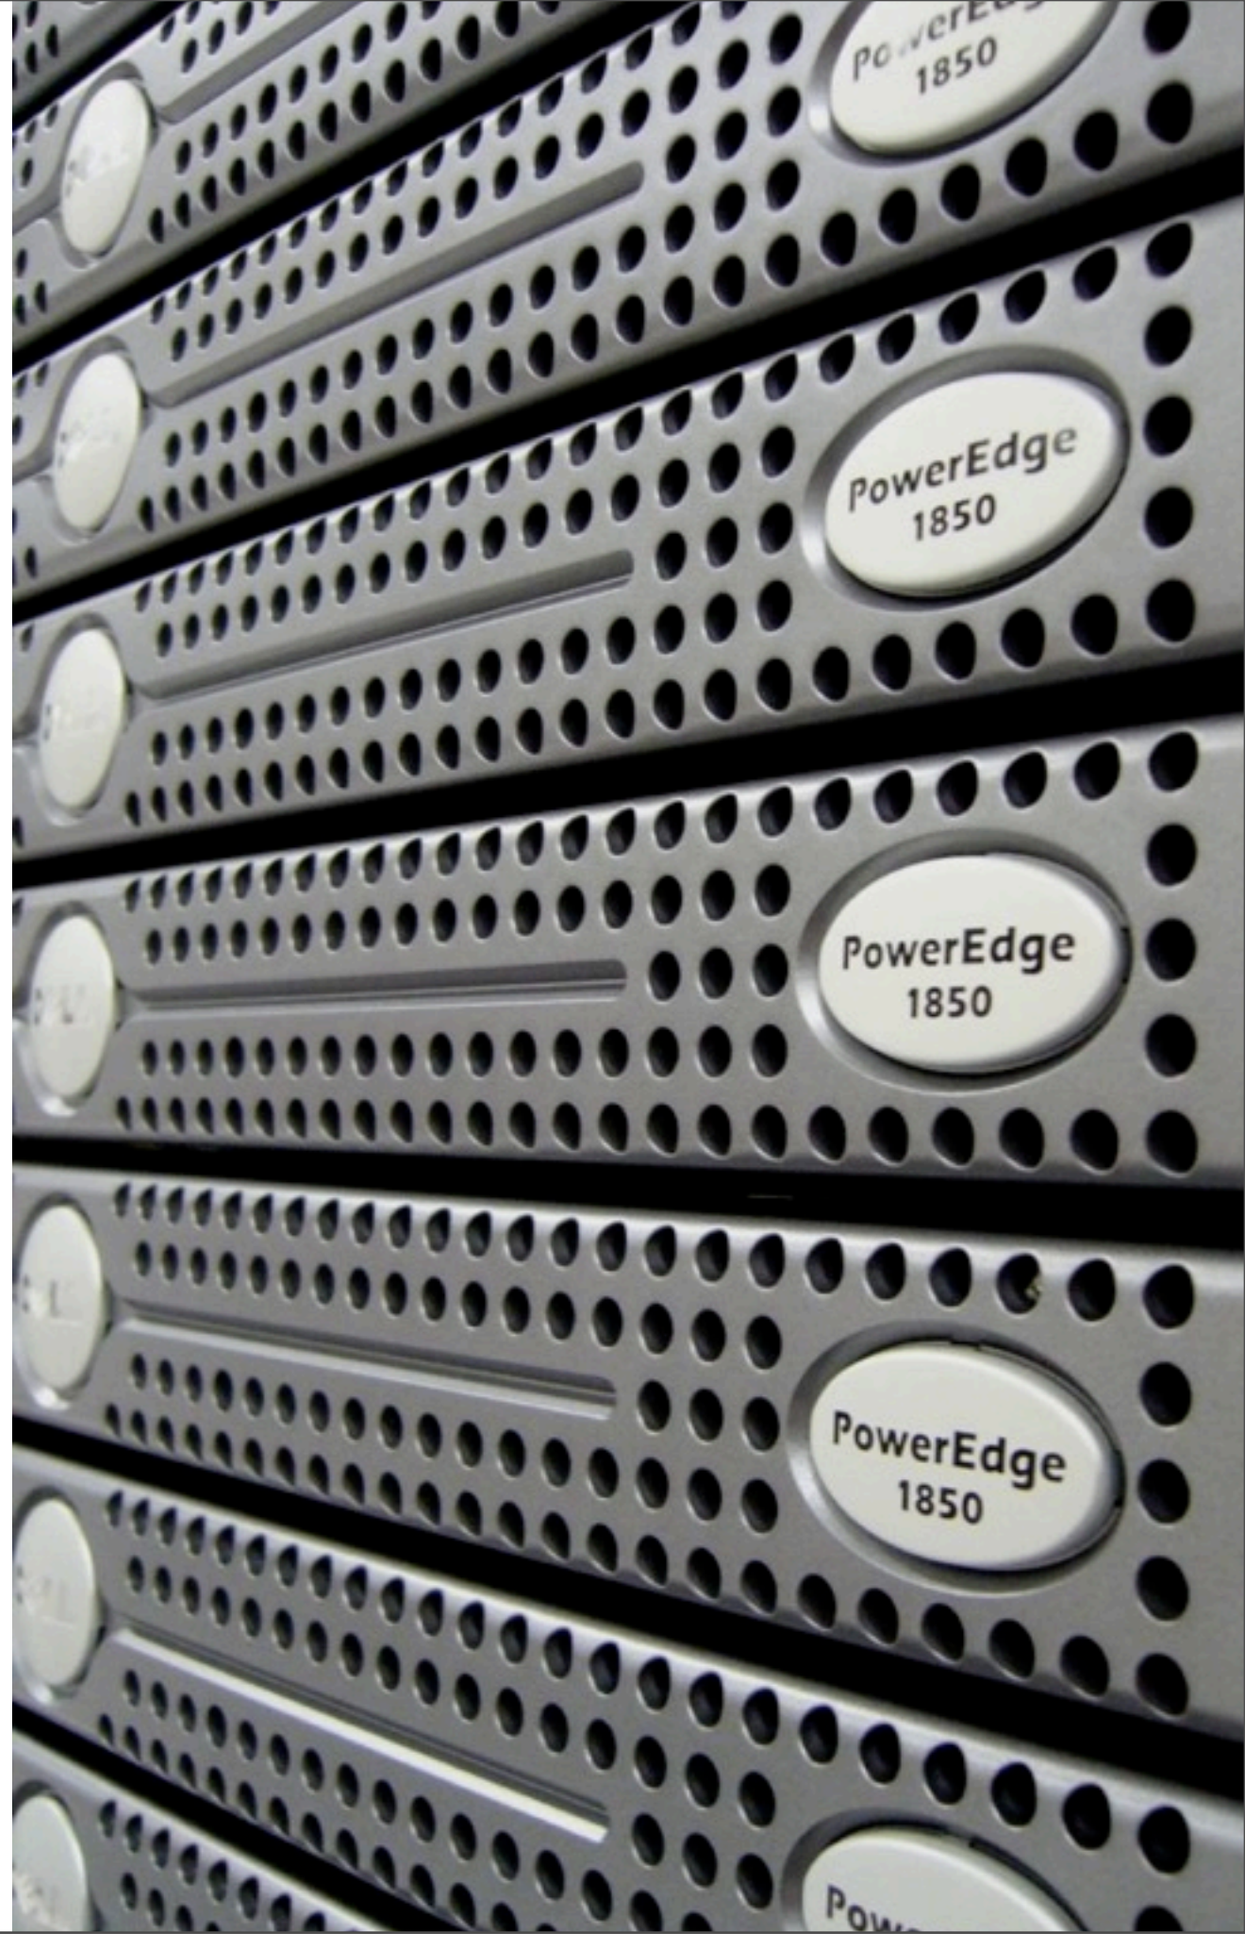

How is a cloud different from plain-old virtualization?

4 Key Questions

Who allocates resources?

Who deploys virtual machines?

How quickly can new resources be allocated?

Is provisioning under human or programmatic control?

4 Key Questions	Virtualization
Who allocates resources?	Administrators
Who deploys virtual machines?	Administrators
How quickly can new resources be allocated?	Depends on the approval process
Is provisioning under human or programmatic control?	Human

4 Key Questions	Virtualization	Cloud
Who allocates resources?	Administrators	Users
Who deploys virtual machines?	Administrators	Users
How quickly can new resources be allocated?	Depends on the approval process	Minutes
Is provisioning under human or programmatic control?	Human	Programmatic

Clouds

- Generic computing platform
- Zero lead time
- Hardware appears homogeneous
- Specialized hardware is abstracted away
  - Load balancers, firewalls, SAN, TBR, etc.

Choosing a cloud provider

- Cloud computing is like corn farming in Iowa...
    - Tiny margin
    - Large fixed costs
    - Scale is essential
- Small providers won’t be around long.

Fog

A cloud that’s near you.

Software Architecture for the Cloud

The Easiest Way to Prepare Your App for the Cloud

The Easiest Way to Prepare Your App for the Cloud

Don’t Do Anything.

Advantages to Gain	Risks to Mitigate
Scalability	Availability
Bundling	Geography
Ephemerality	Ephemerality

Amazon EC2 and S3

Advantages

- Run as many servers as you want
- Pennies per hour
- Virtual server images stored in S3
- Control via web services API

Quirks

- “Clean boot” is really clean
- Local storage not persistent
- IP addresses assigned randomly

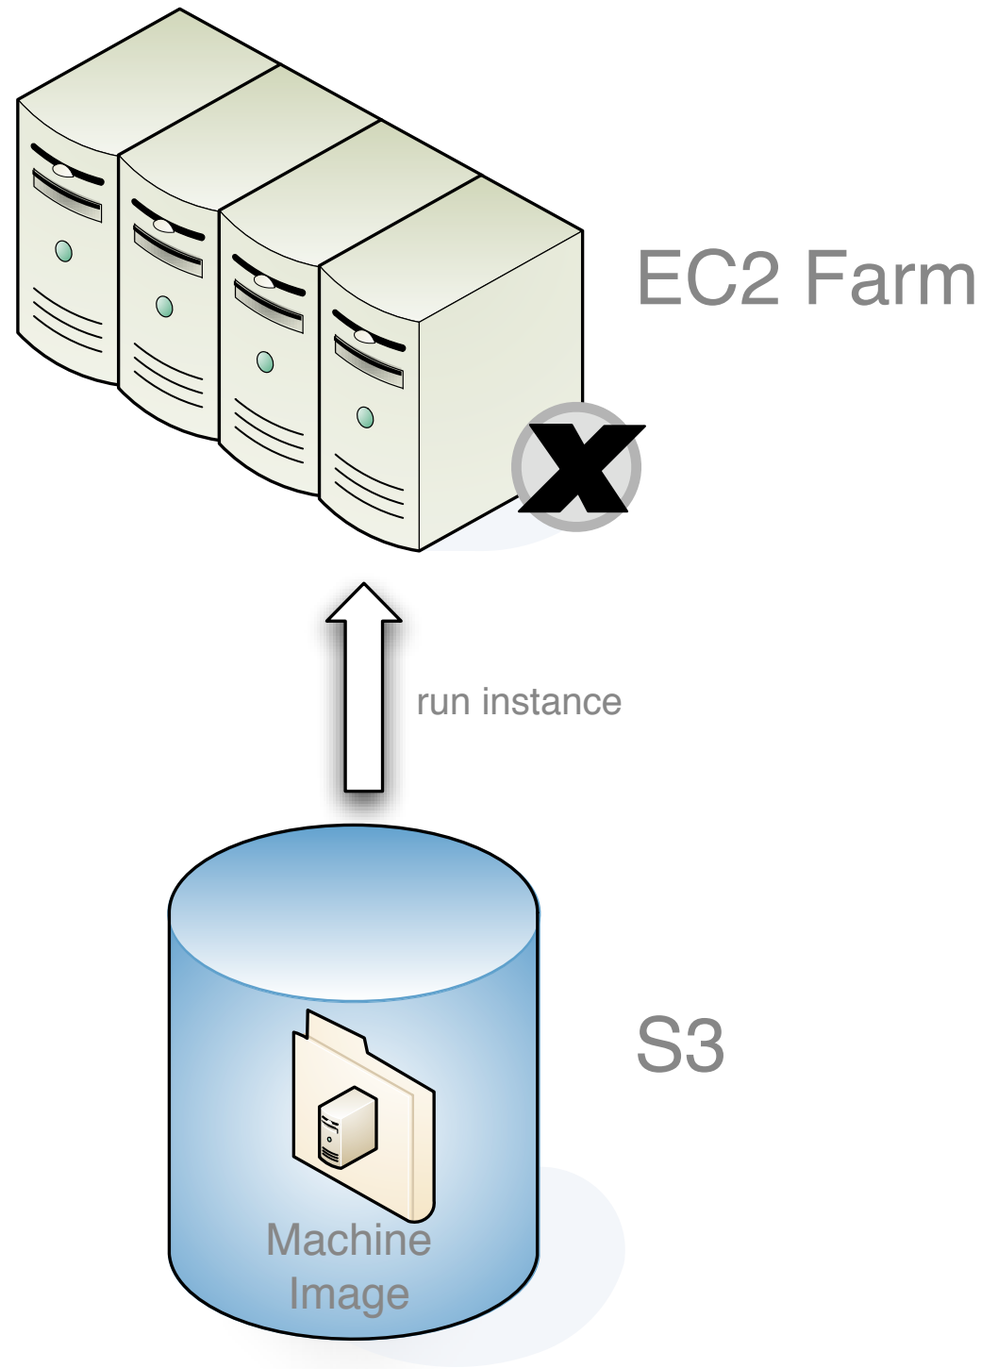

Controlling the Flock

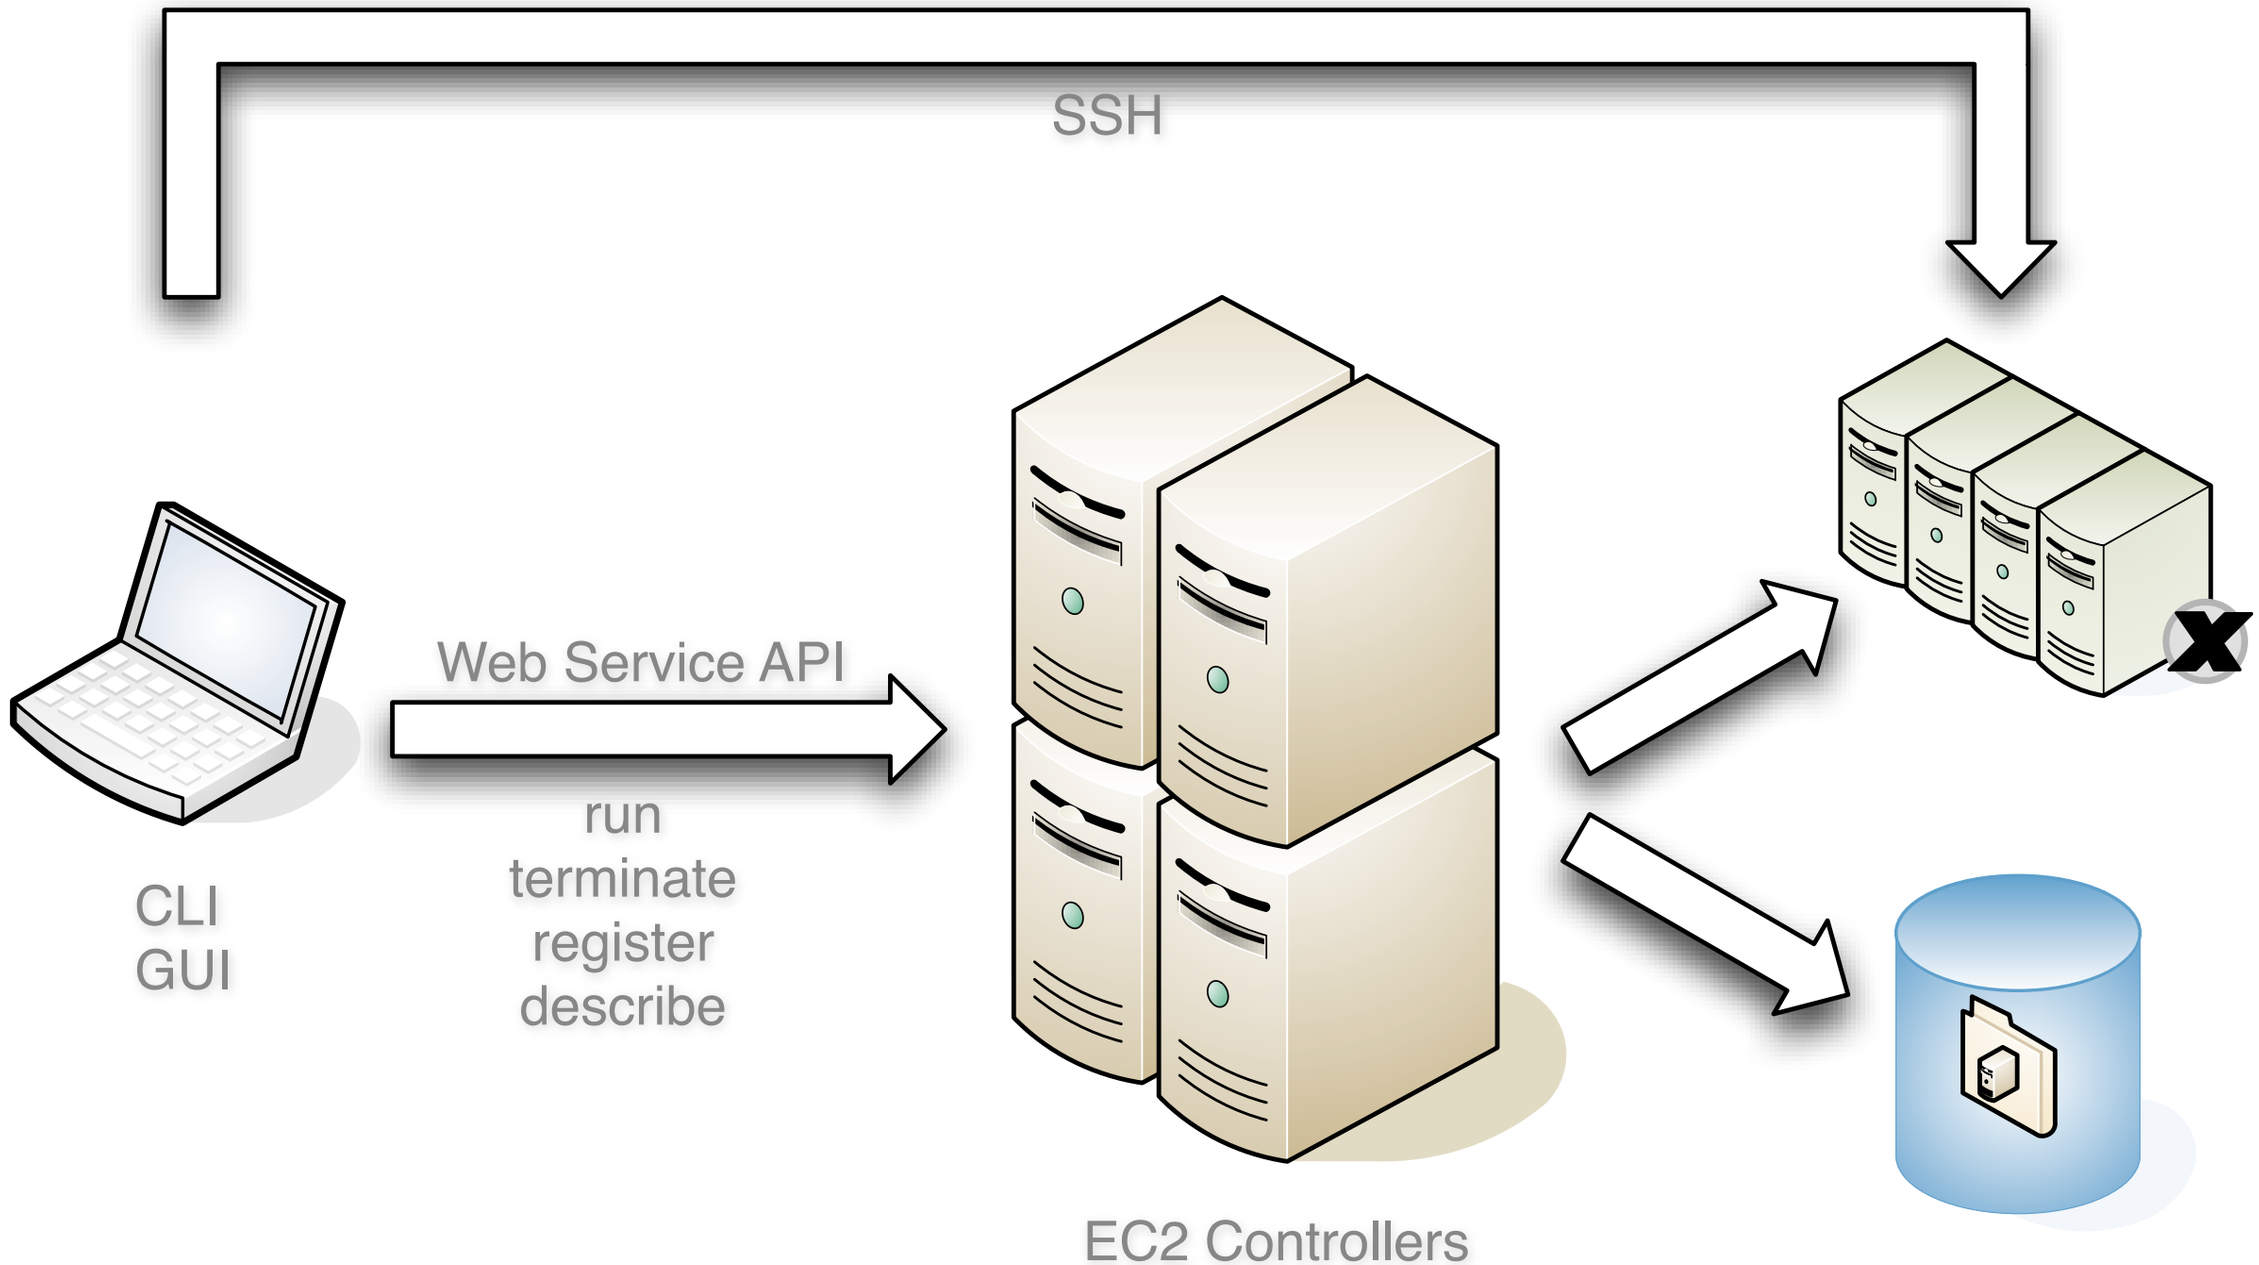

Bundling - Deployment by Virtual Machine

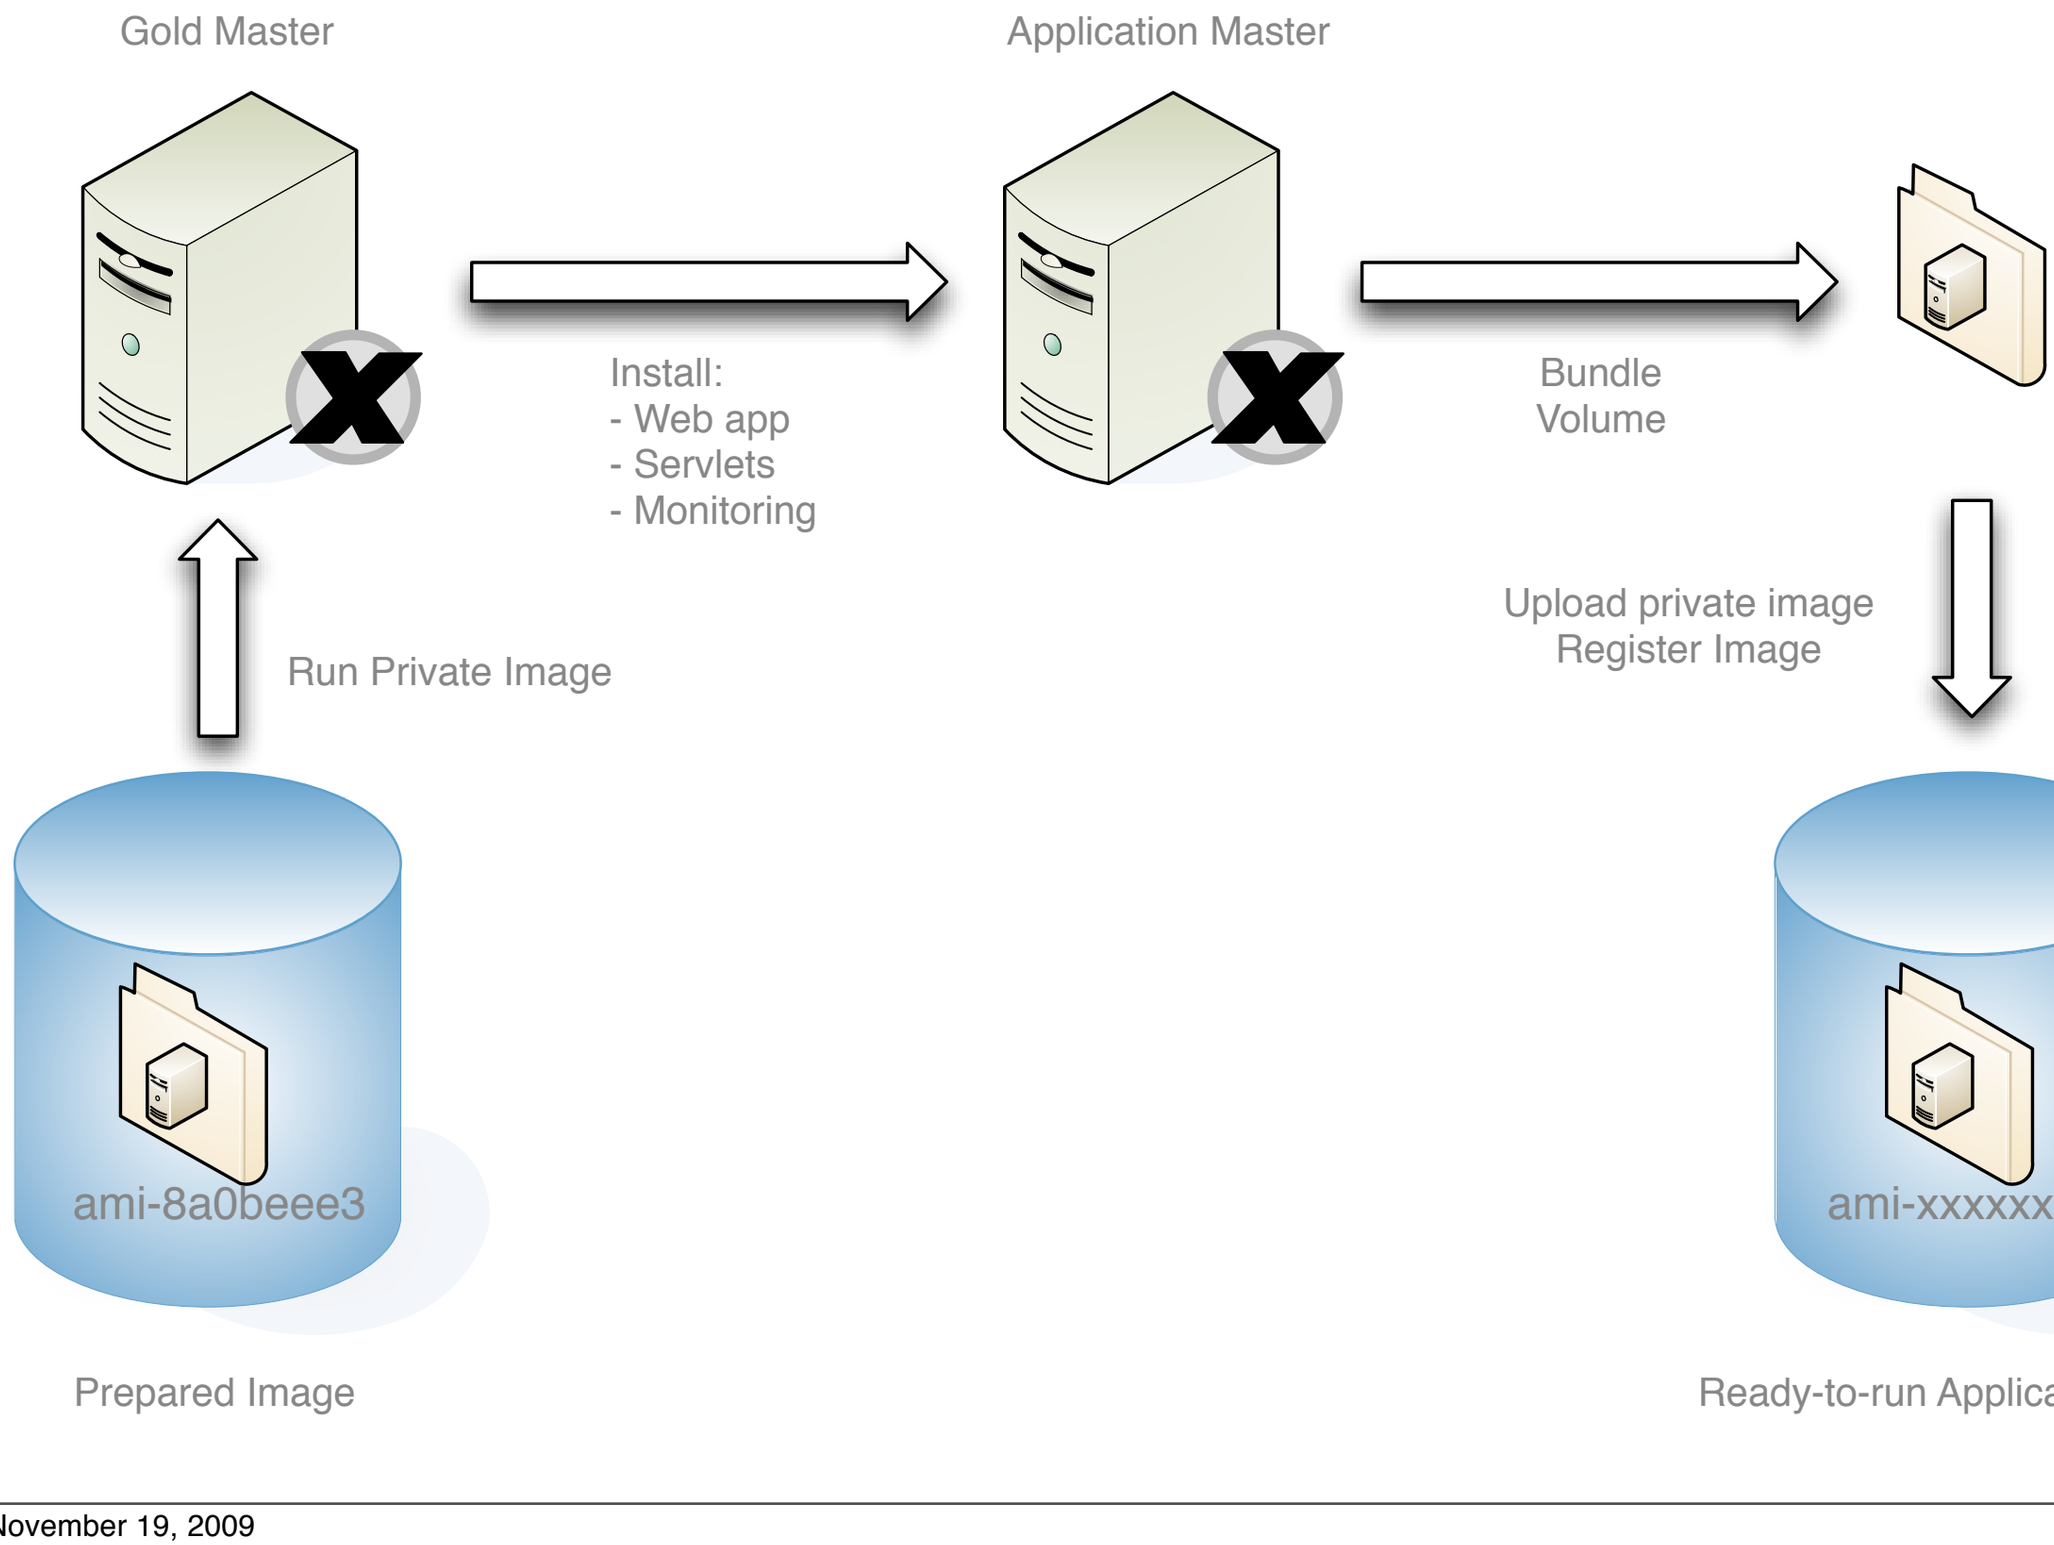

Ready to run

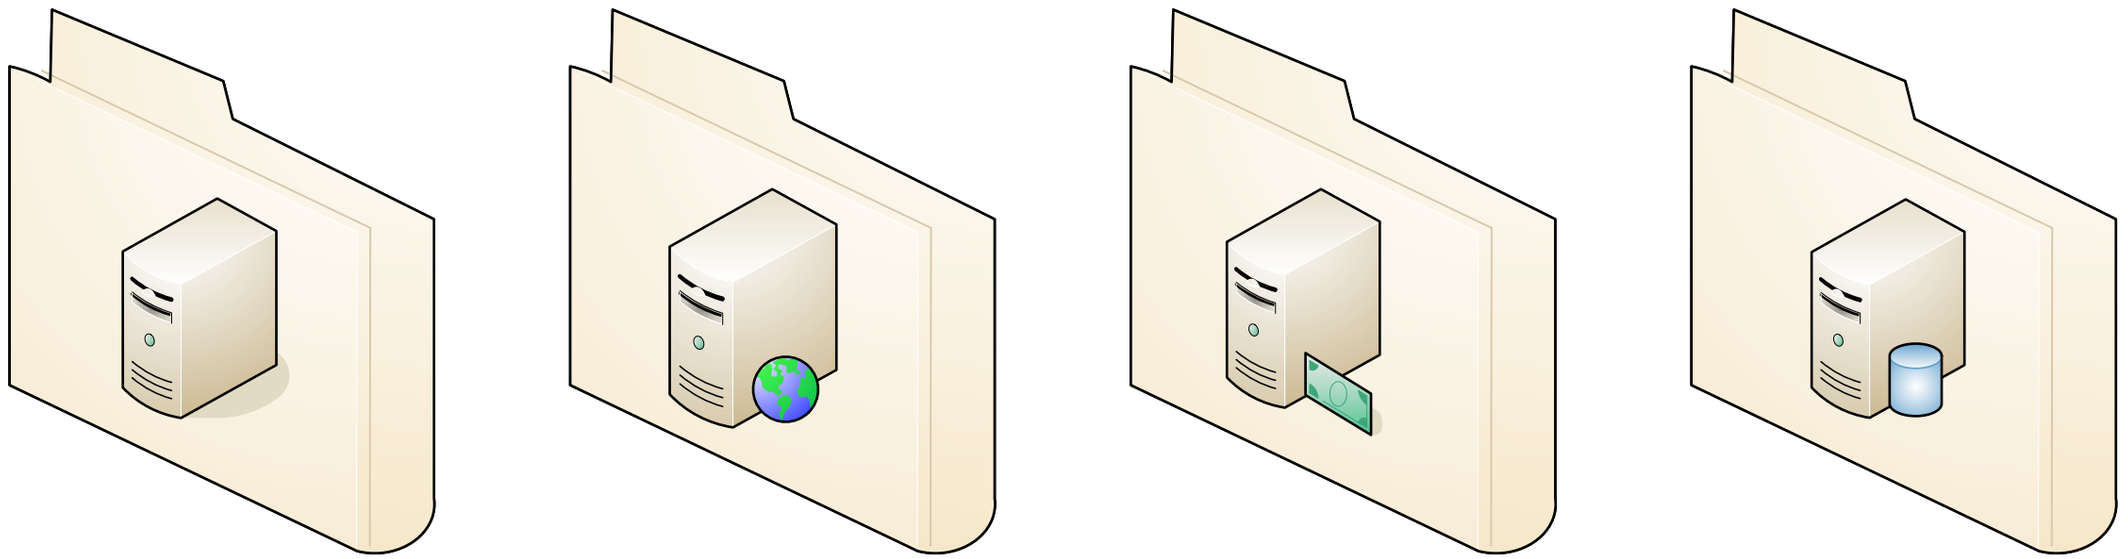

Gold Master Canned Web Server Canned App Server Canned DB Server

Multiple images in your S3 bucket
All “boot ready” and fully configured

The Plus Side of Ephemerality

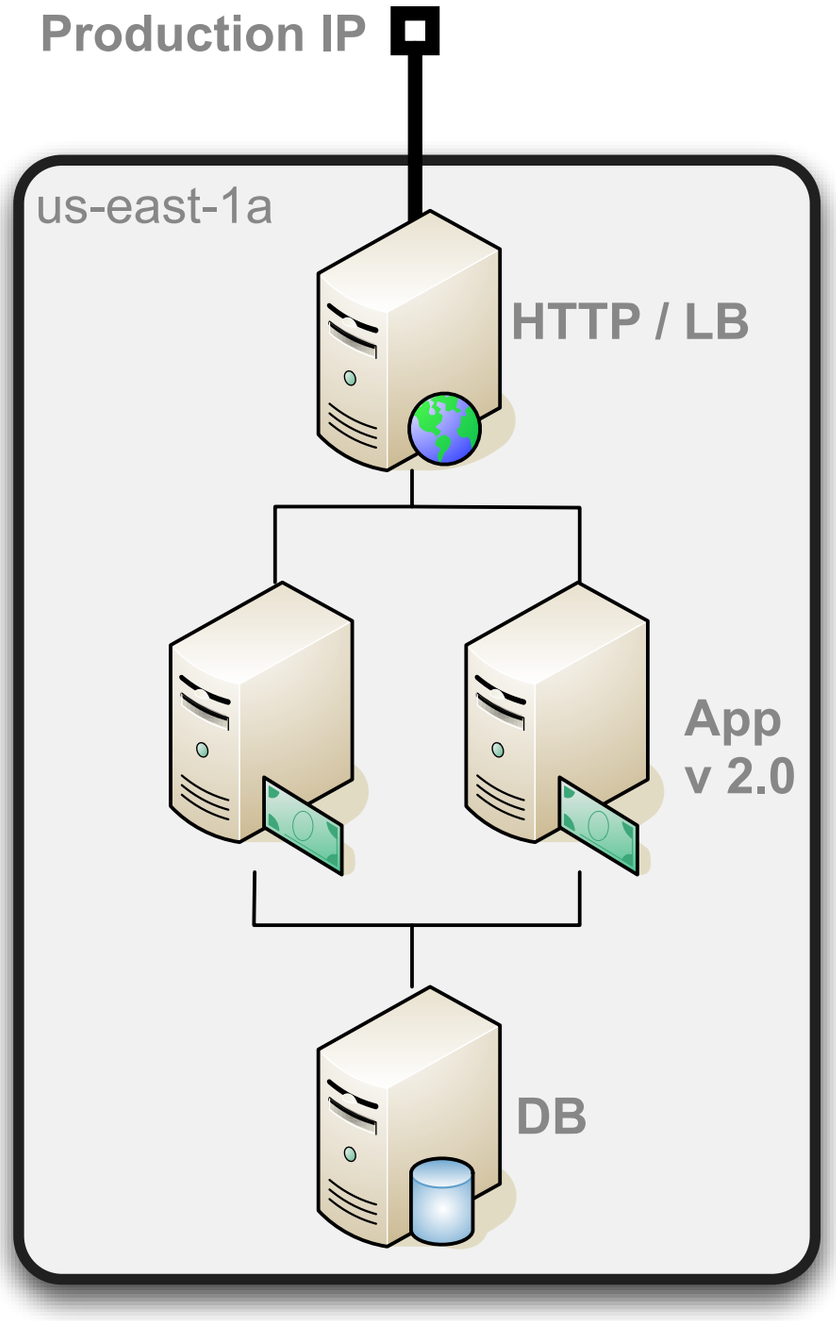

The Plus Side of Ephemerality

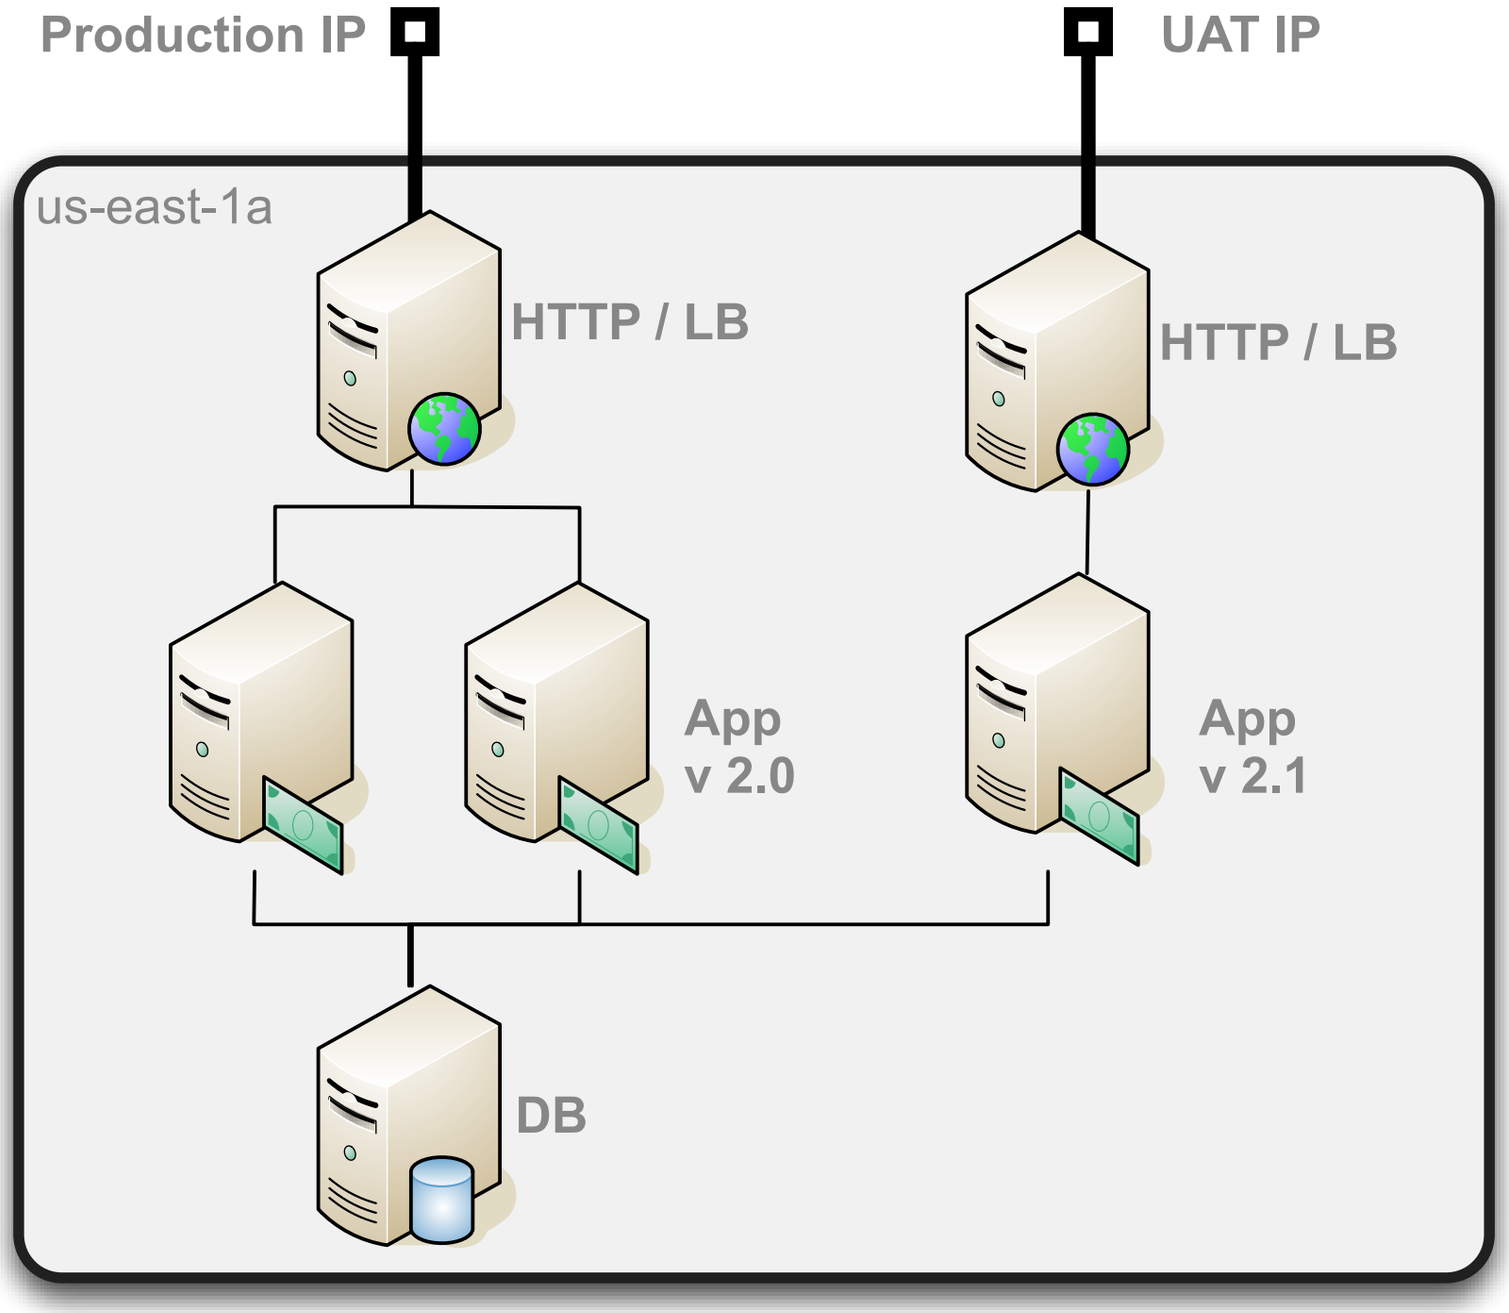

The Plus Side of Ephemerality

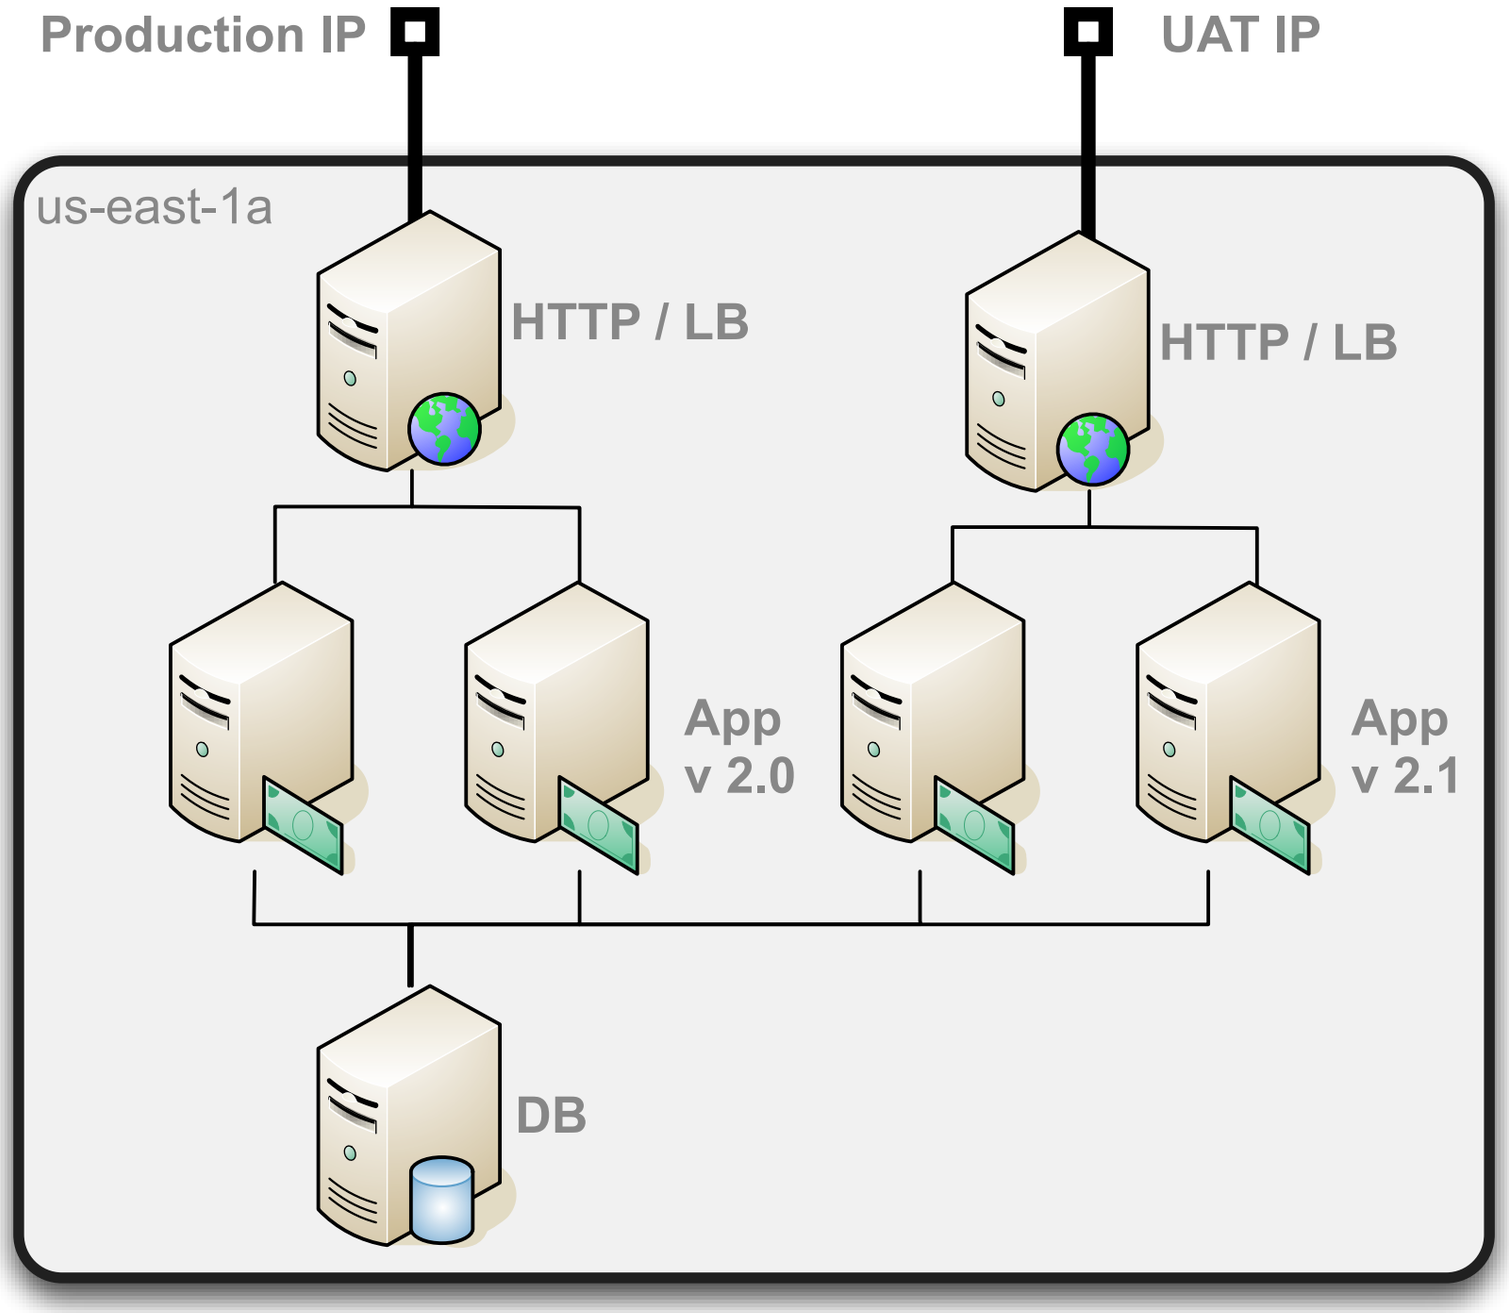

The Plus Side of Ephemerality

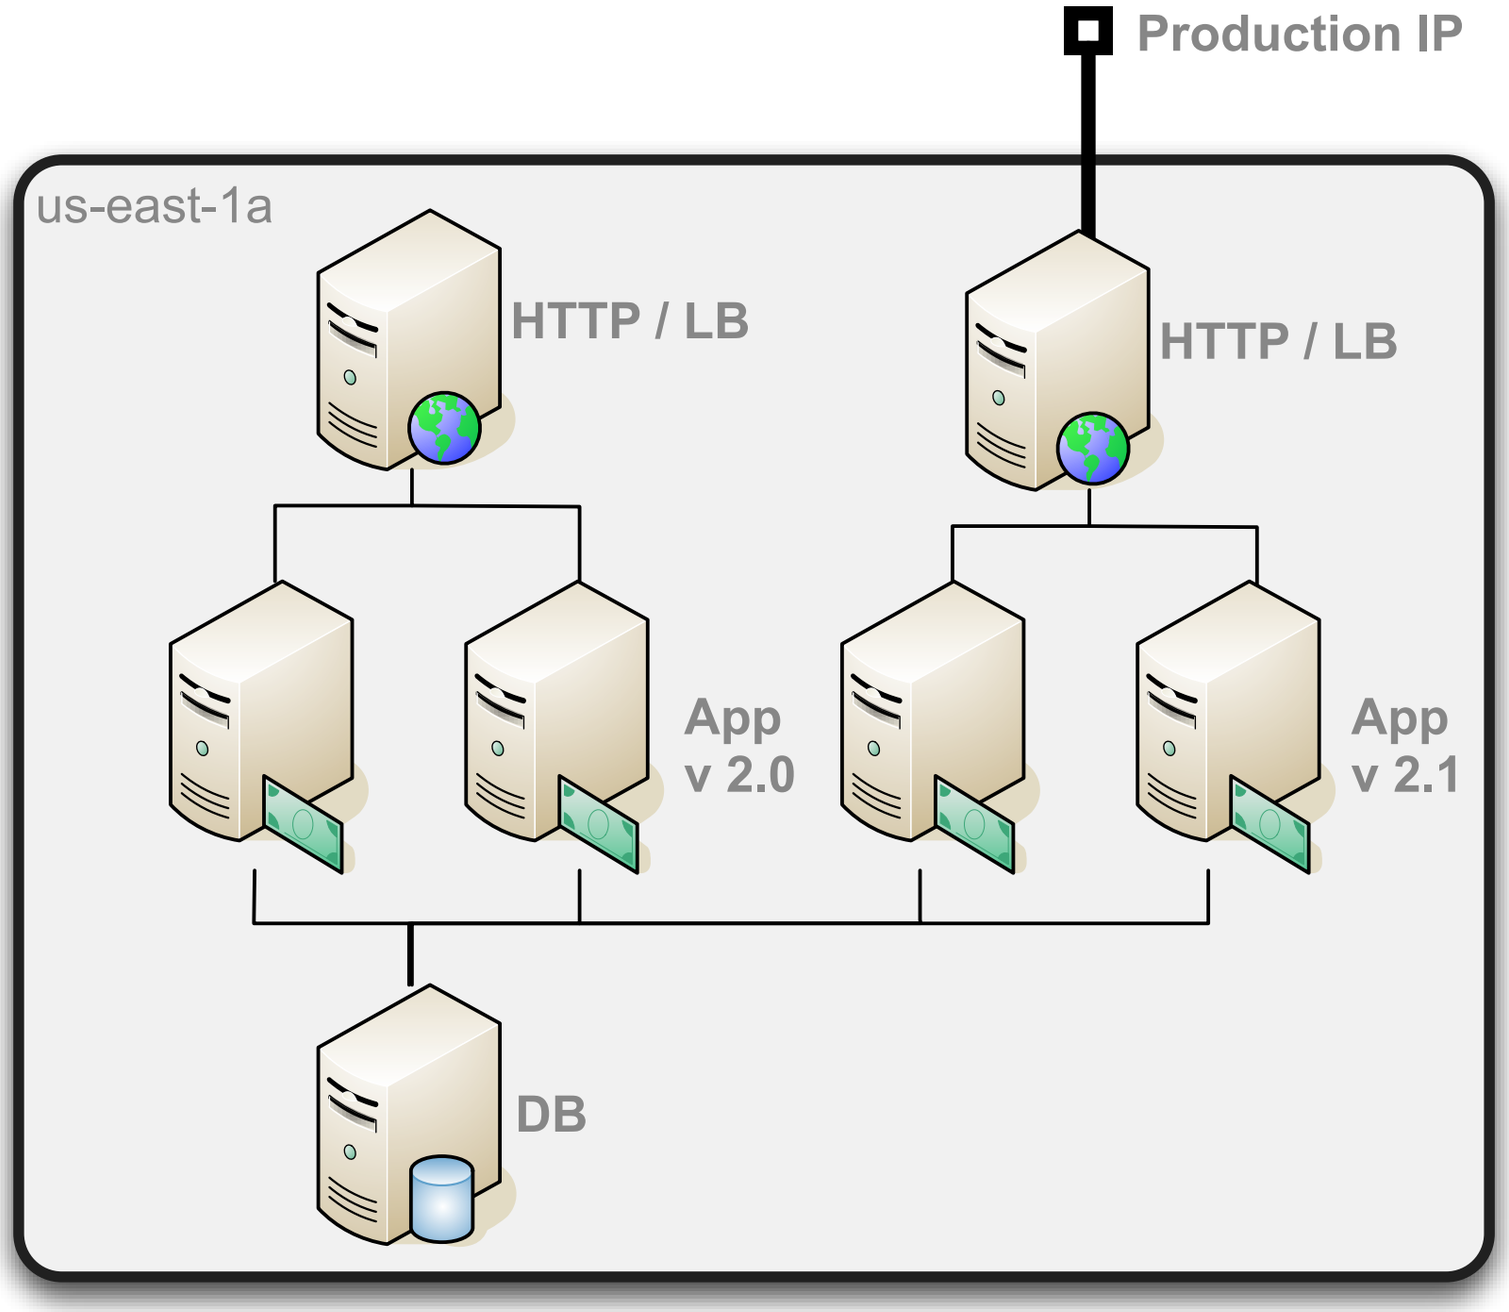

The Plus Side of Ephemerality

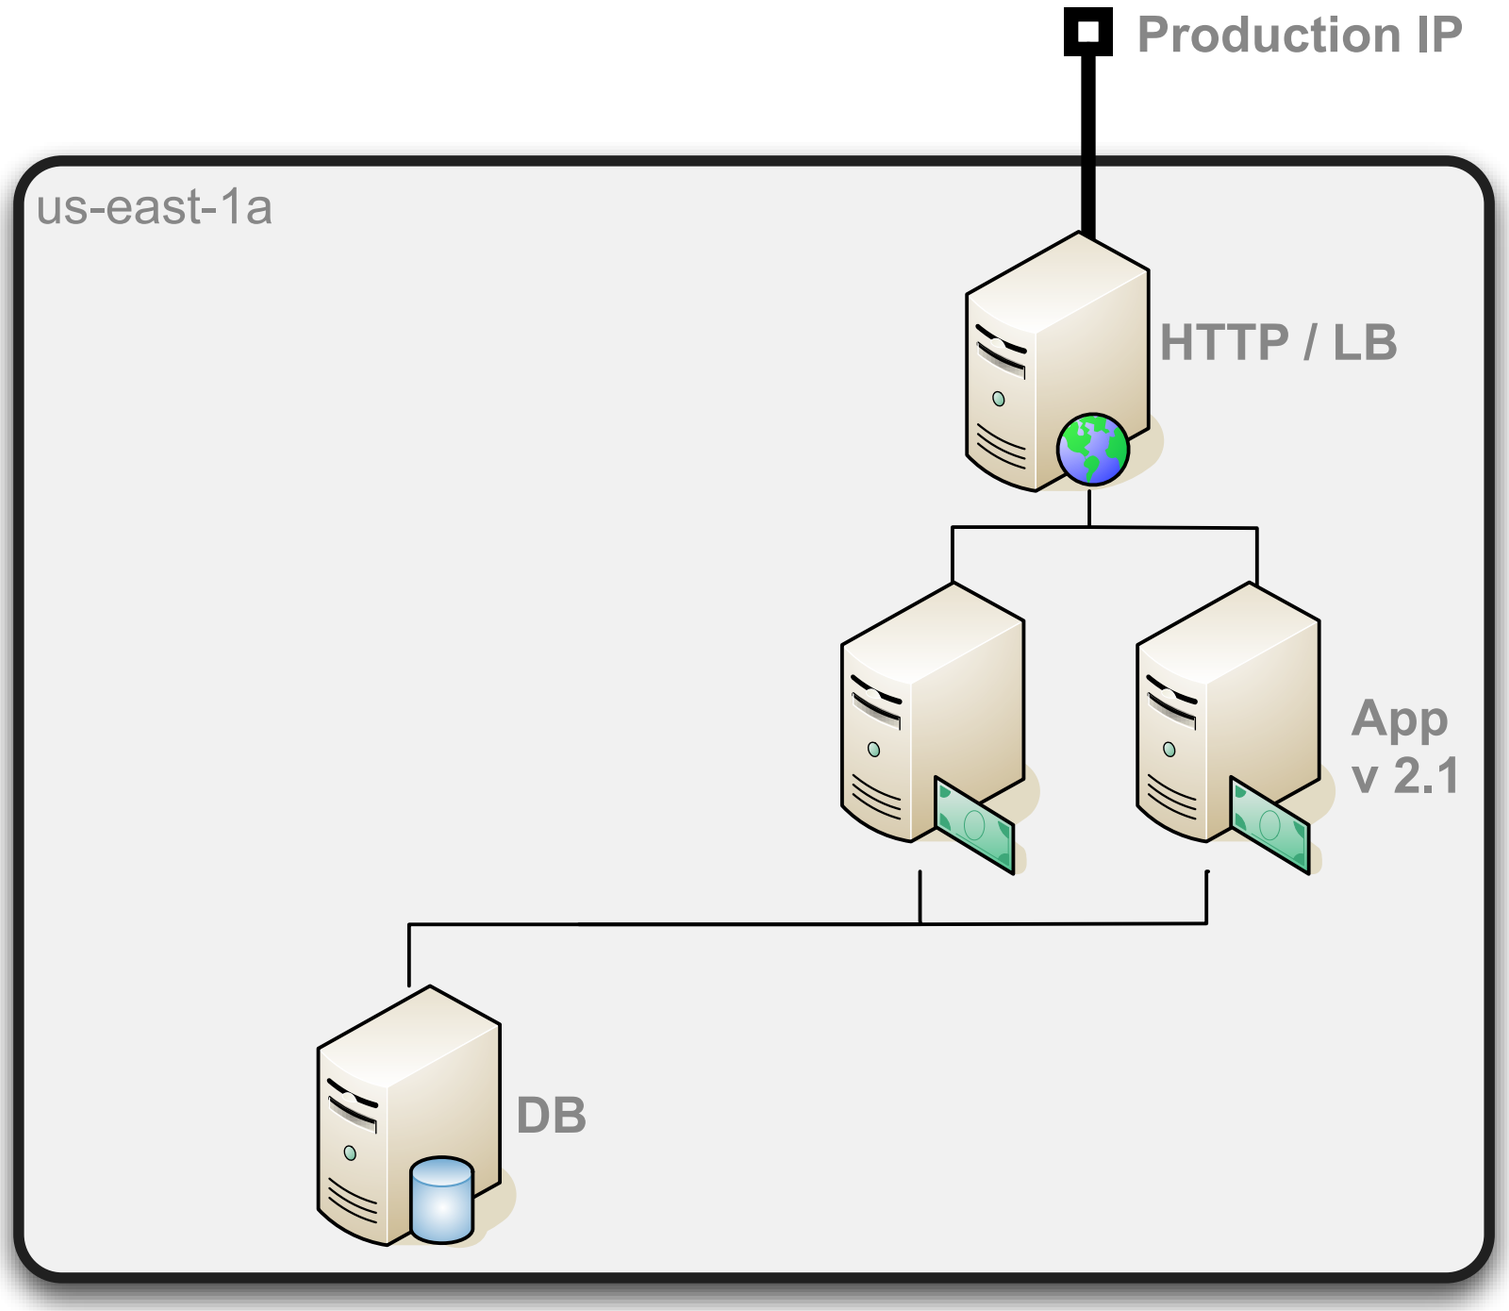

Operational Benefits

- No more deployments to production servers
- Reduced rate of downtime:
    - Operator error
    - Discrepancy between environments

Schema Version Conflict?

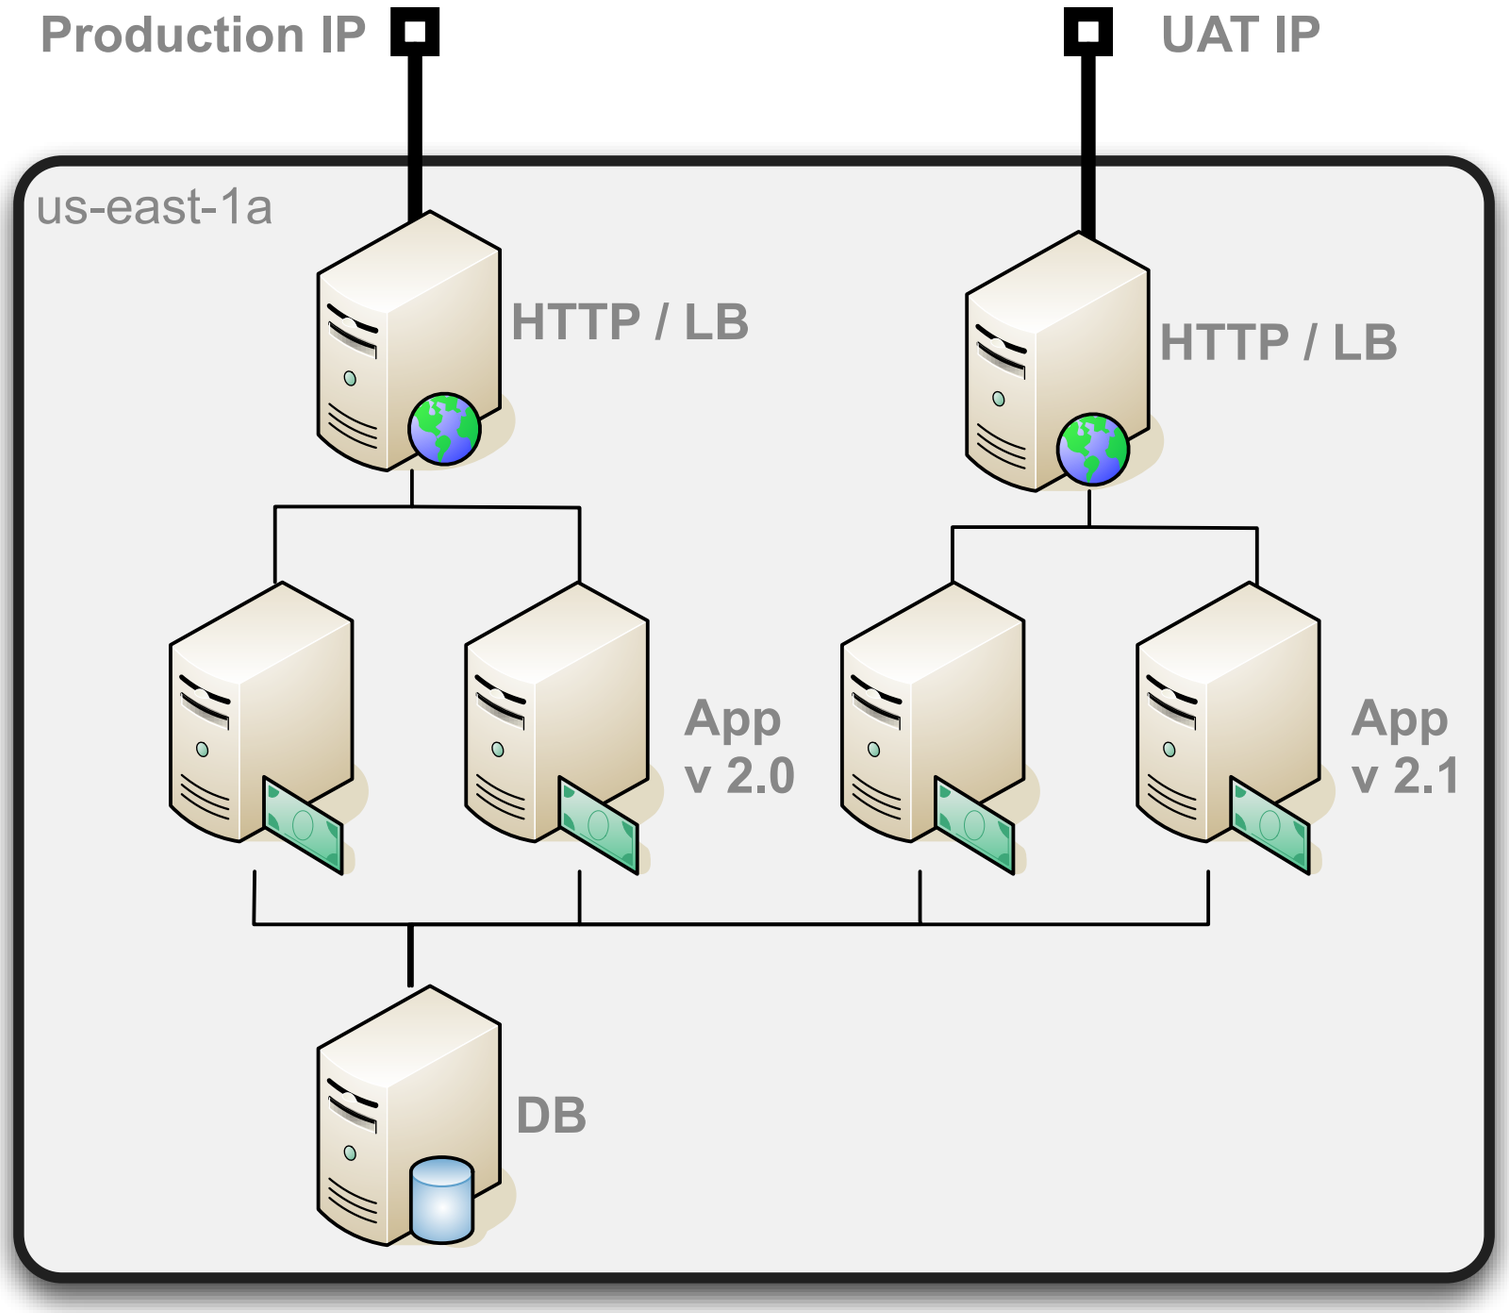

Handling Concurrent Versions

Layer	Strategies
Web Assets	Embed version in URLs. Let old and new assets coexist.
Integration Points	Version all protocols and encodings. Client specifies acceptable versions. Server responds in mutually understood version.
Database Schemas	Apply “Agile Database” techniques. Automate migrations, split into phases: expand first, contract later

Scalability - The Master/Worker Pattern

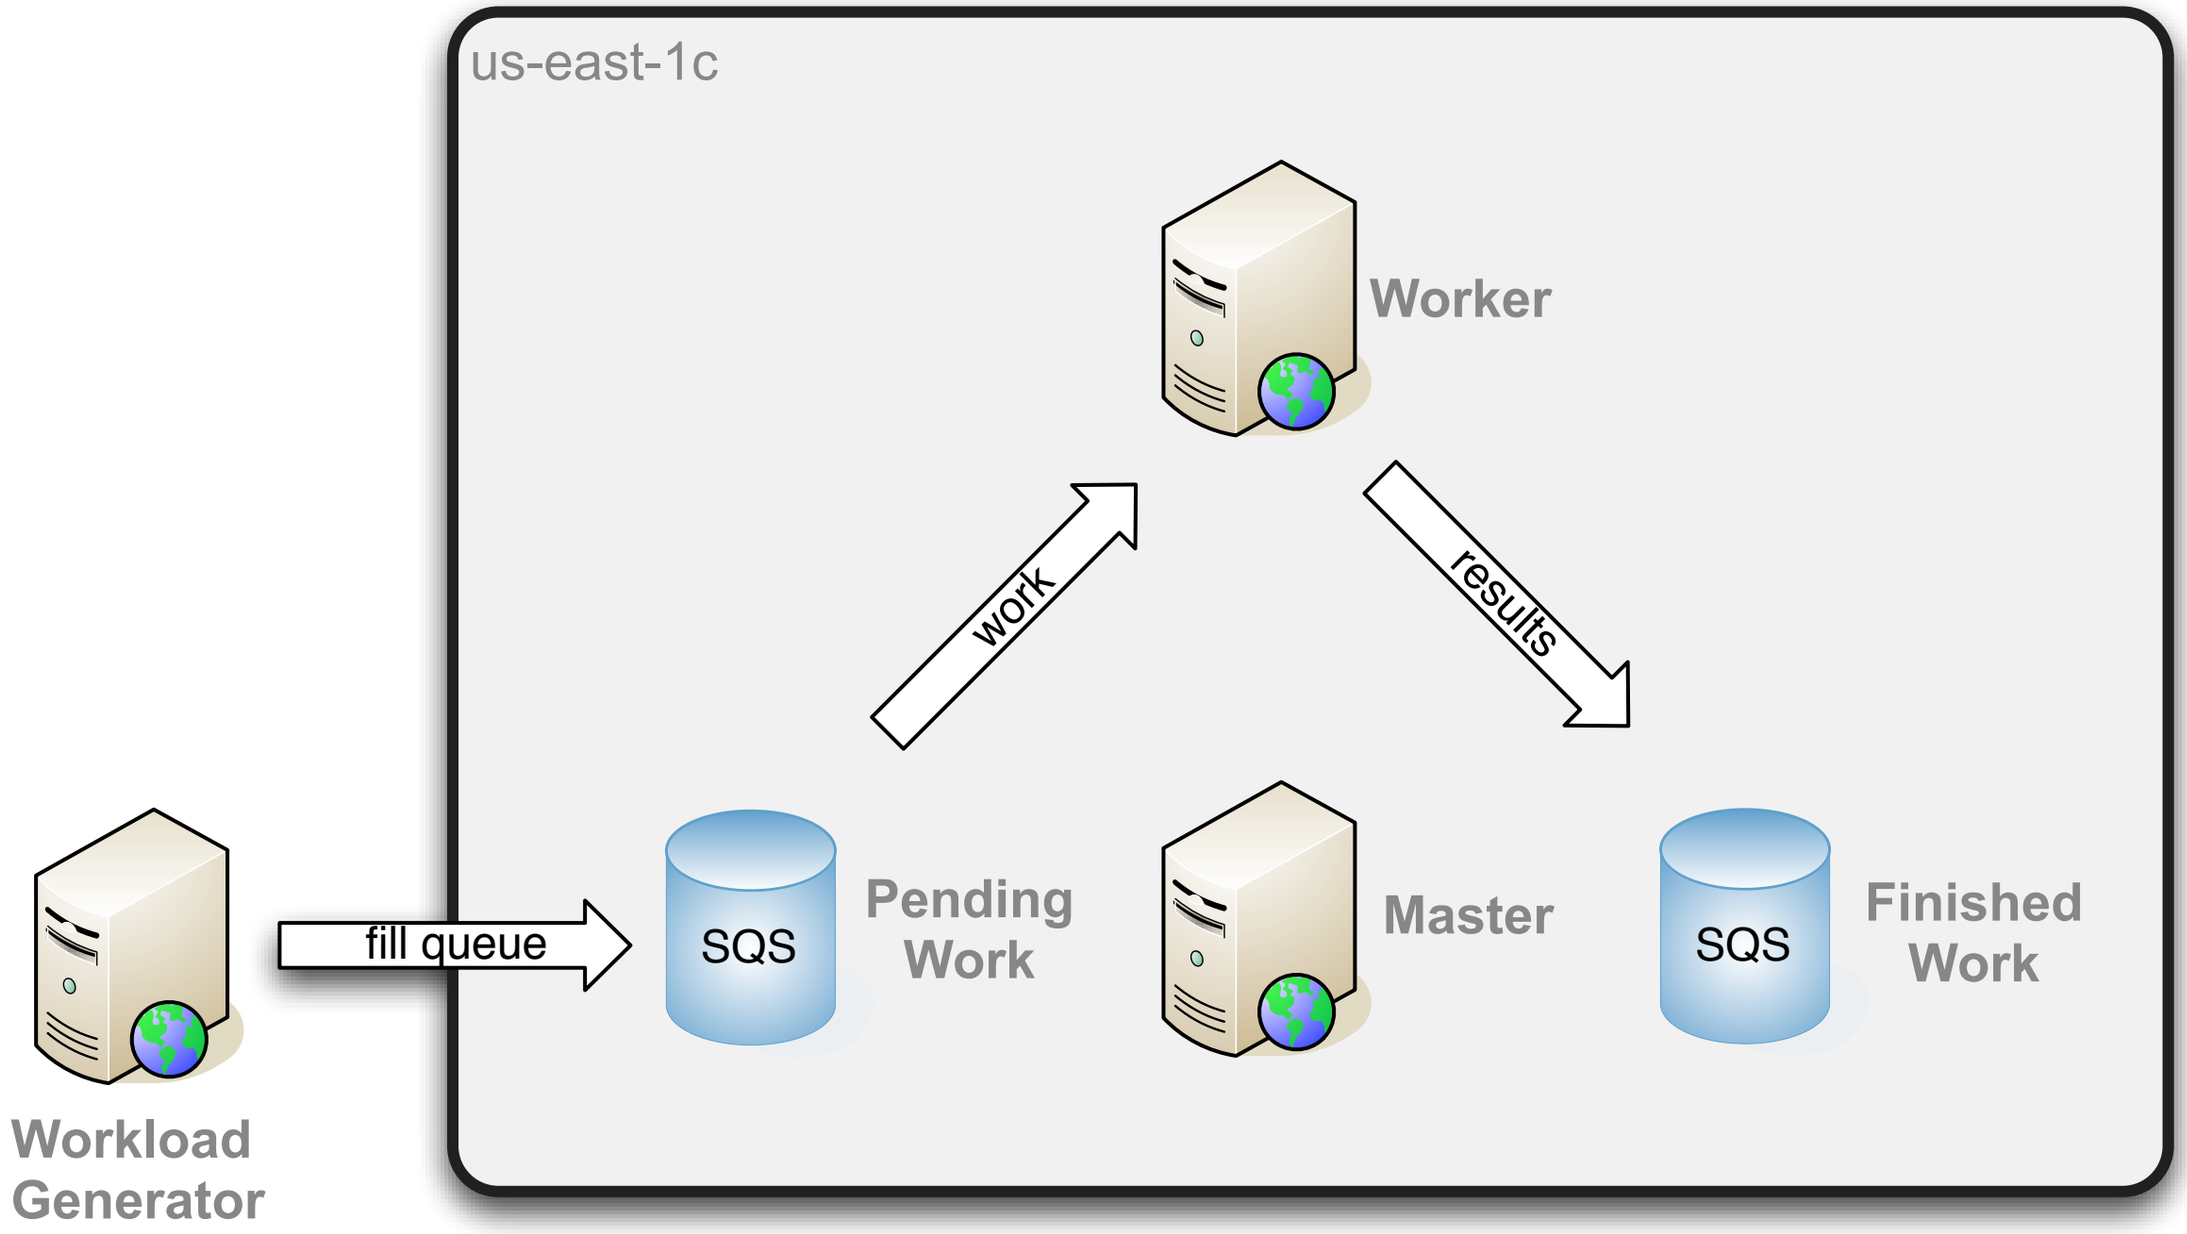

Master/Worker Scenarios

- New York Times - converted 11 million articles from their archive
- Video transcoding/processing - Animoto.com
    - From 50 to 5,000 processing nodes in 1 week
    - Then back down to “a few hundred”

Scalability - Map/Reduce via Hadoop

- Open-source distributed computing
  - Hadoop Core
    - Distributed filesystem (HDFS)
    - Map/Reduce framework
  - HBase
    - High-scalability, distributed database (non-relational)
    - Similar to Google’s BigTable and Amazon’s SimpleDB

Map-Reduce

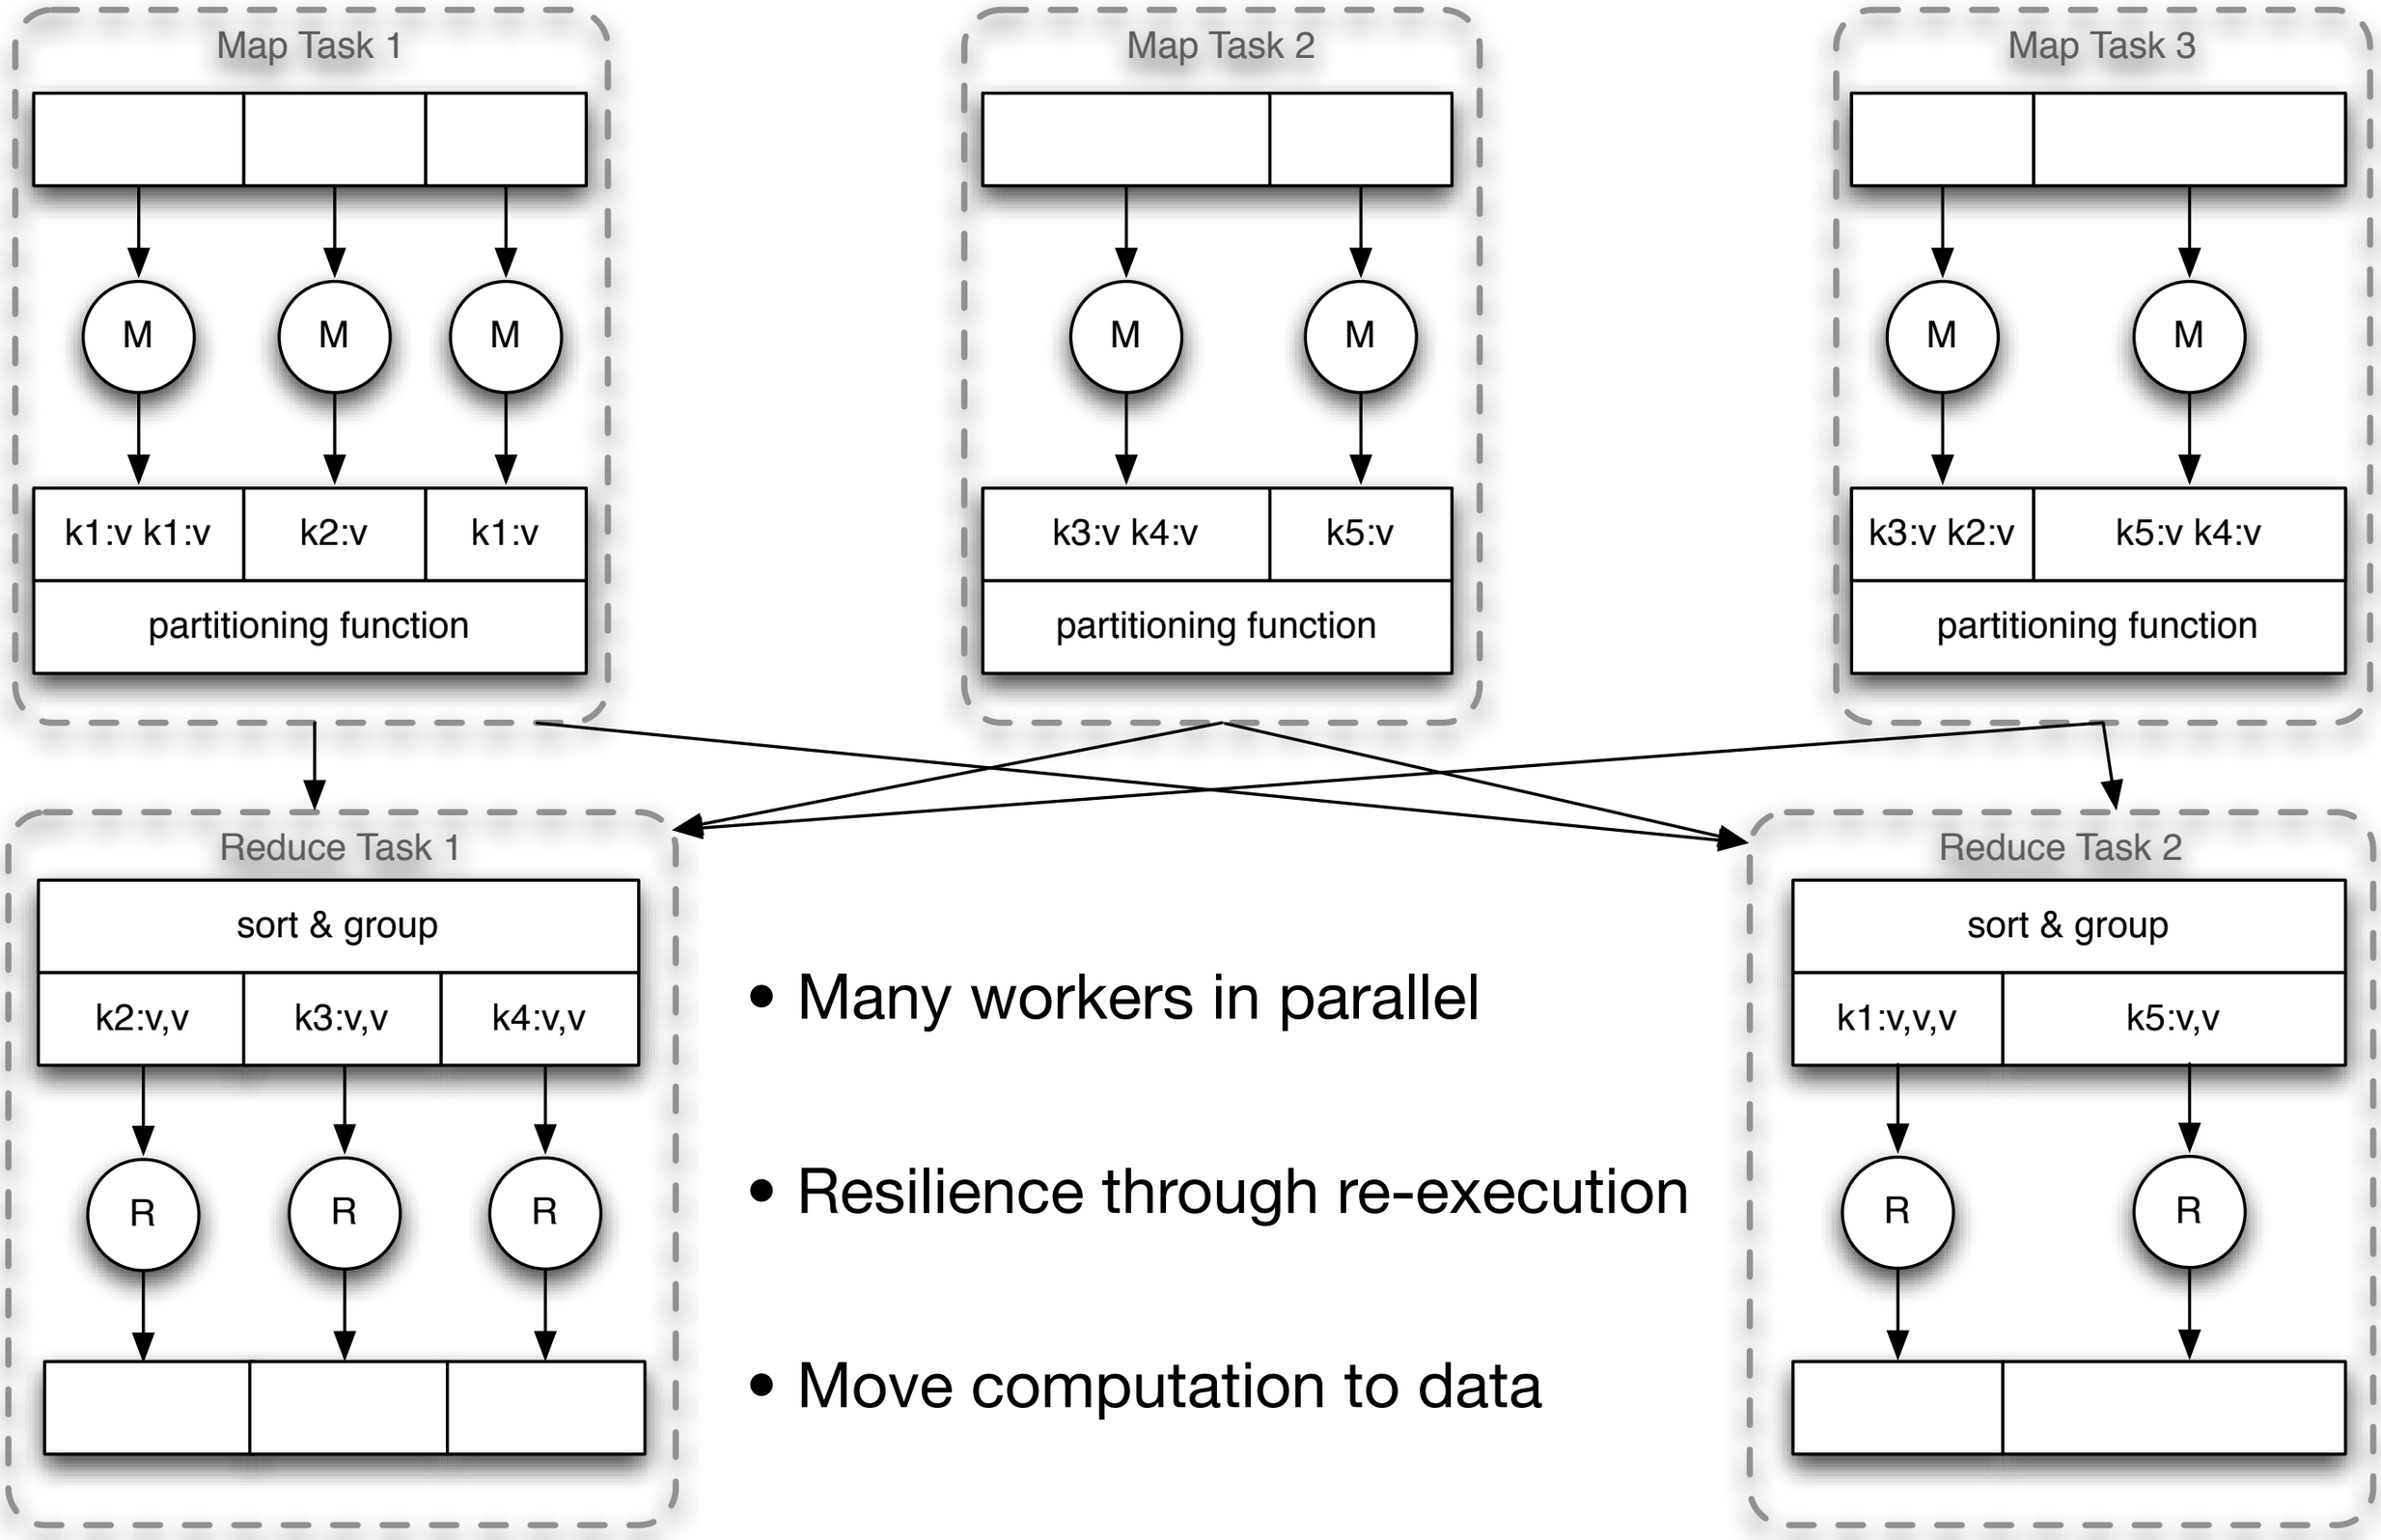

- Many workers in parallel
- Resilience through re-execution
- Move computation to data

Scalability - Horizontal Replication

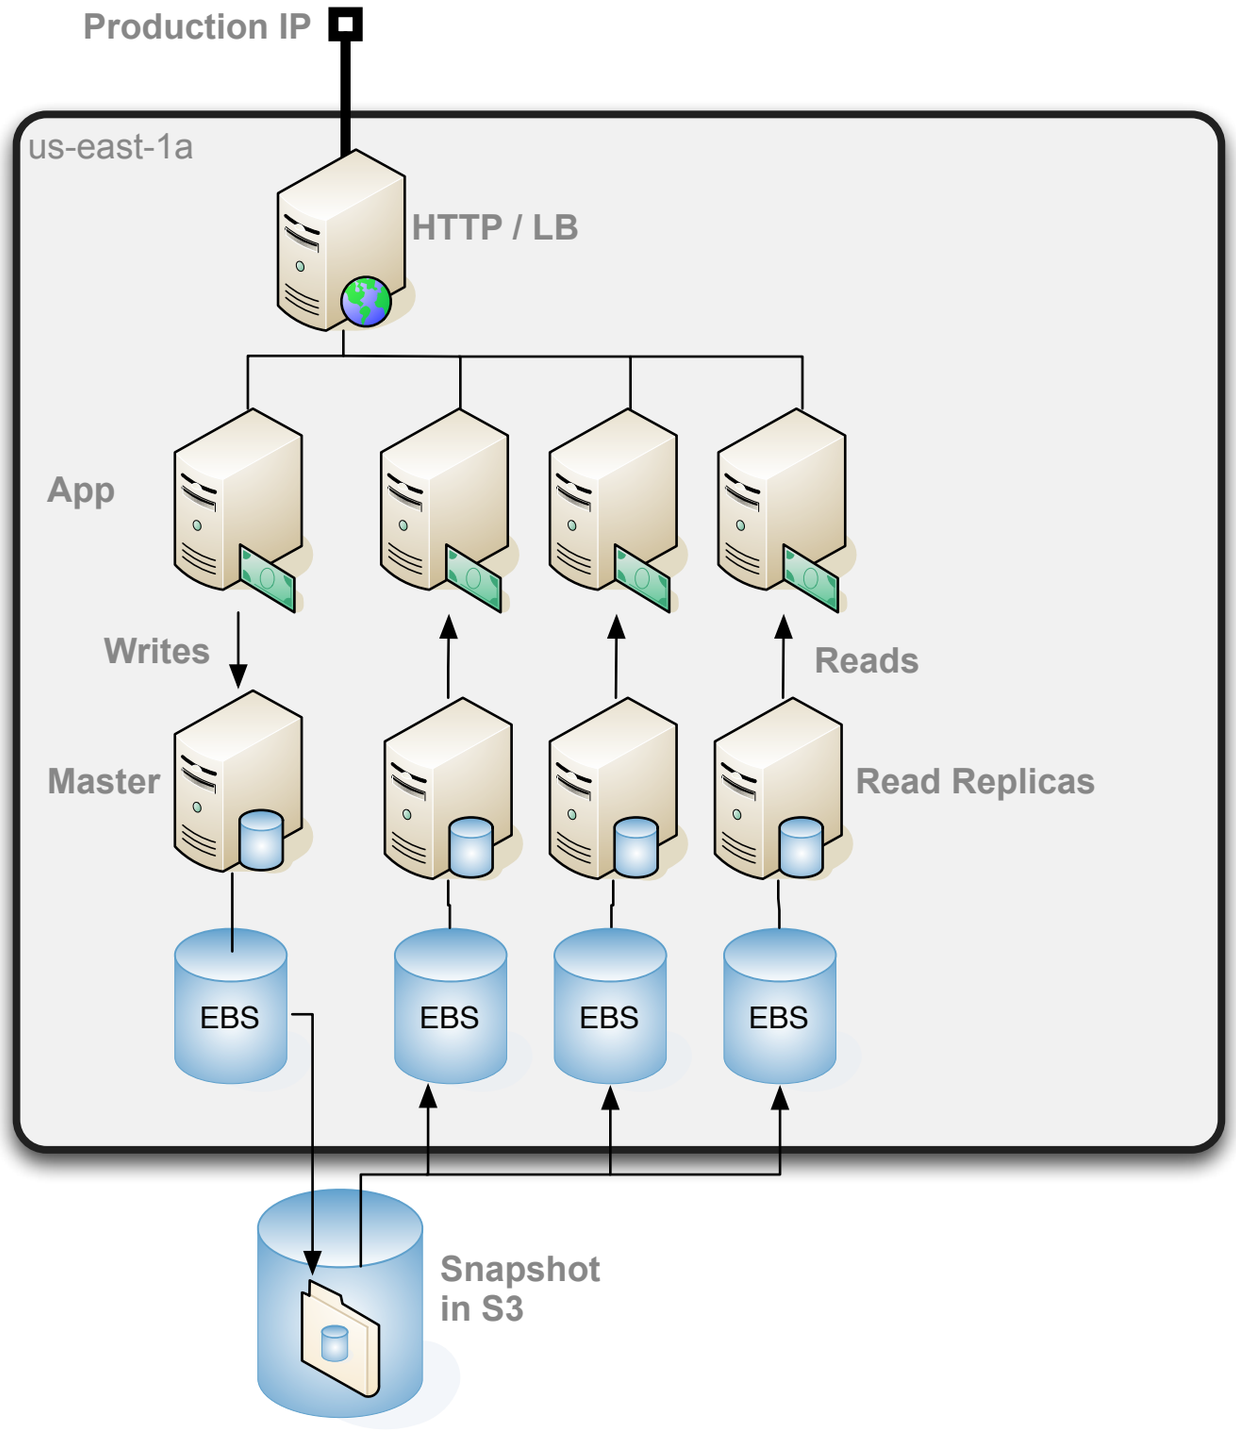

Scalability - Horizontal Replication

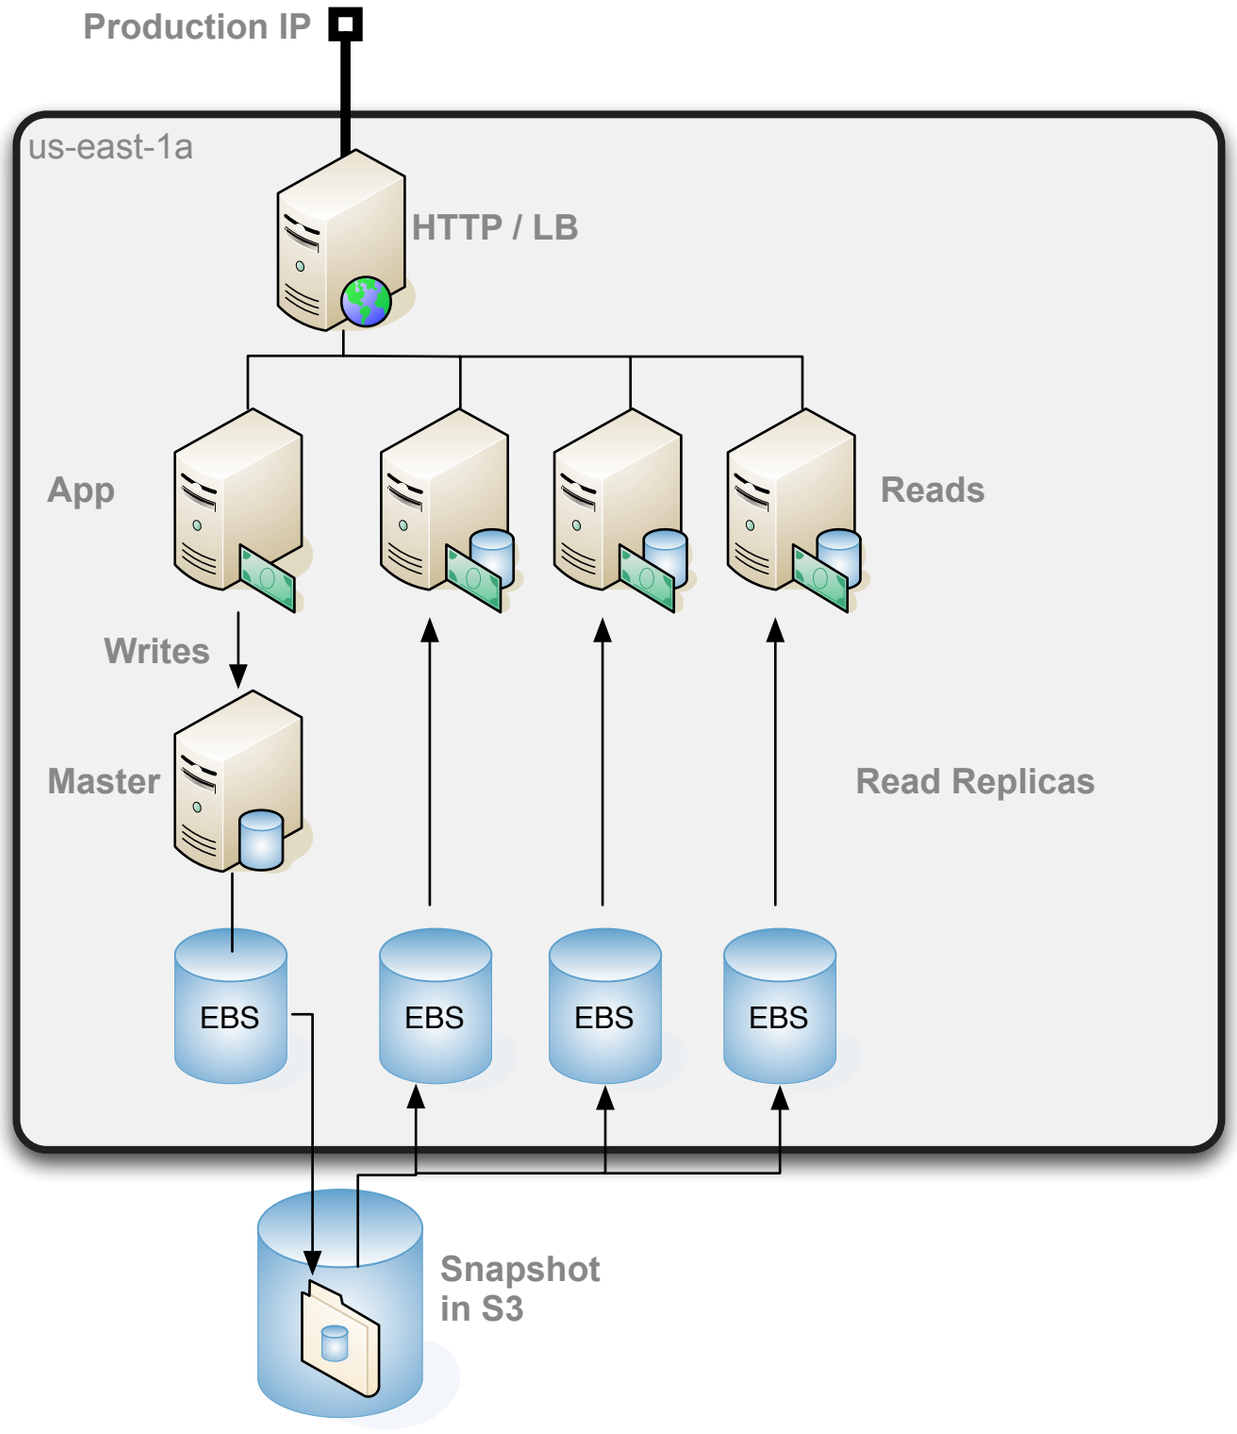

Scalability - Horizontal Replication

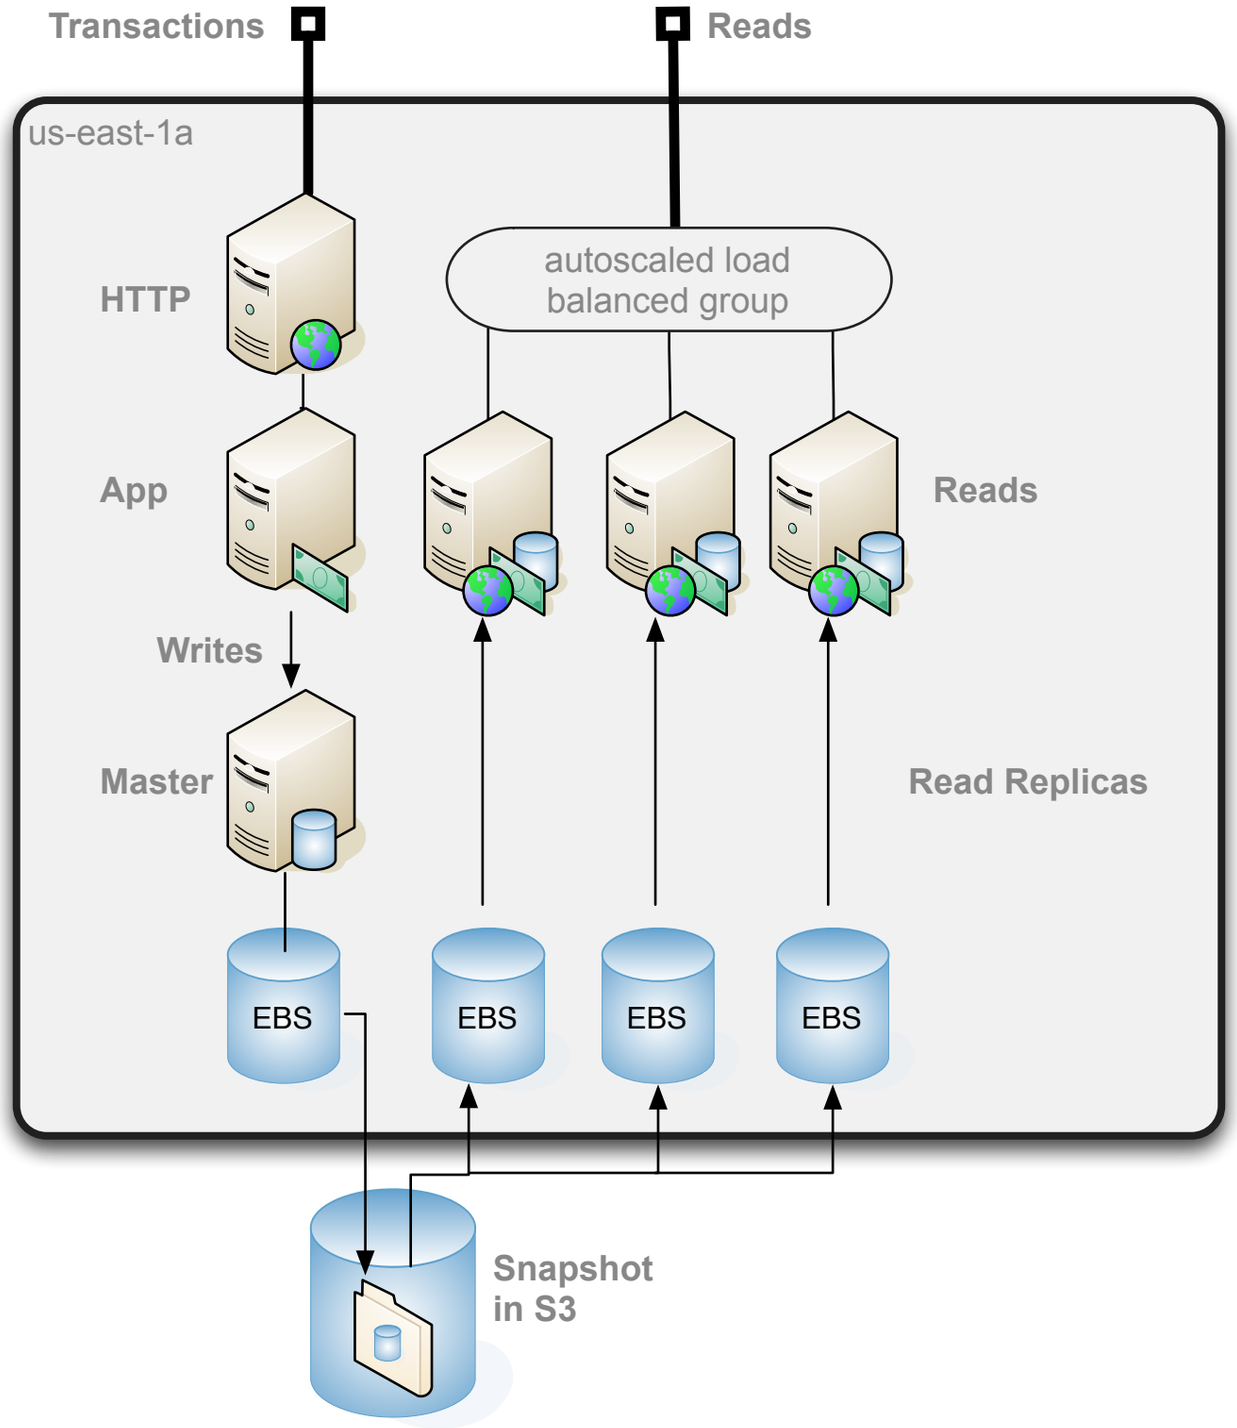

What about “No SQL”?

- Affinity Between Cloud and Scalable Architectures
  - Distributed key-value stores
  - Sharding
  - Eventual consistency

- But not mandatory
  - Not all applications need to be that big
  - Consider both scalability and capacity needs

- Apps that are not cloud-native can still use RDBMS

Scaling with Cache Servers

- Cache servers and In-Memory Data Grids
  - Essential components
  - Serve scalability and capacity needs
  - “Cloud friendly”
- Products to consider
  - memcached
  - GigaSpaces
  - Oracle Coherence
  - Terracotta

Mitigating The Down Side of Ephemerality

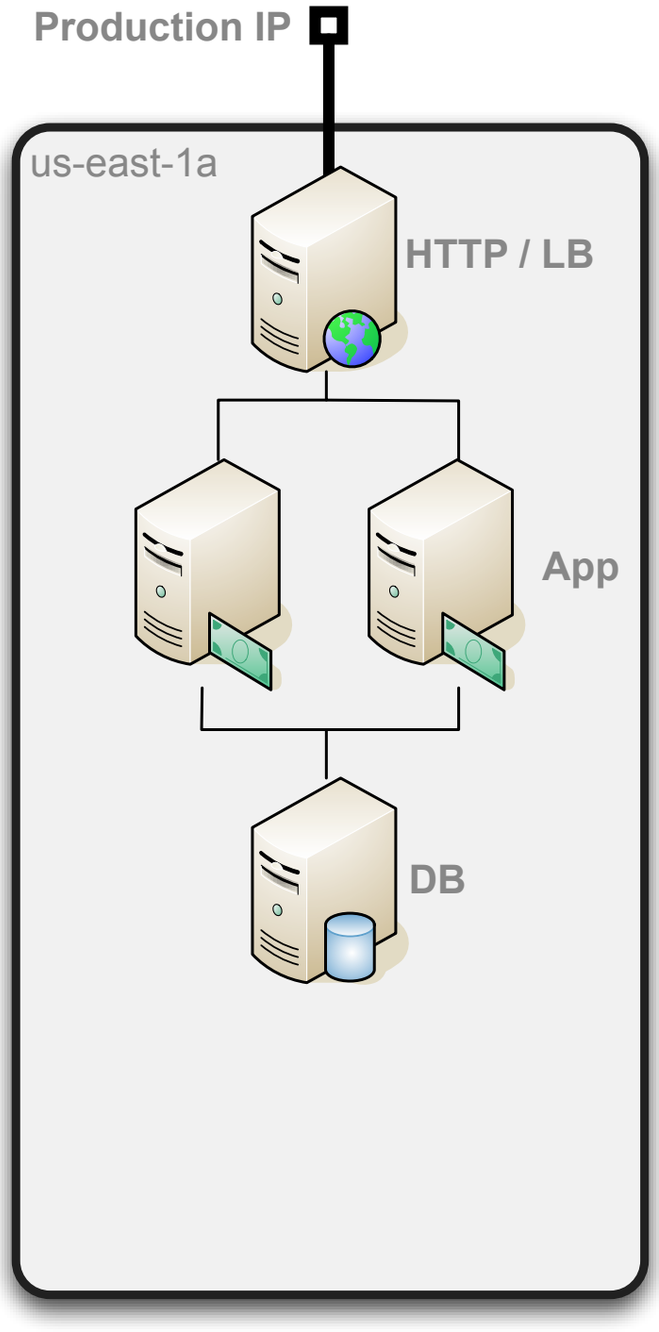

Everything can disappear at any moment.

Mitigating The Down Side of Ephemerality

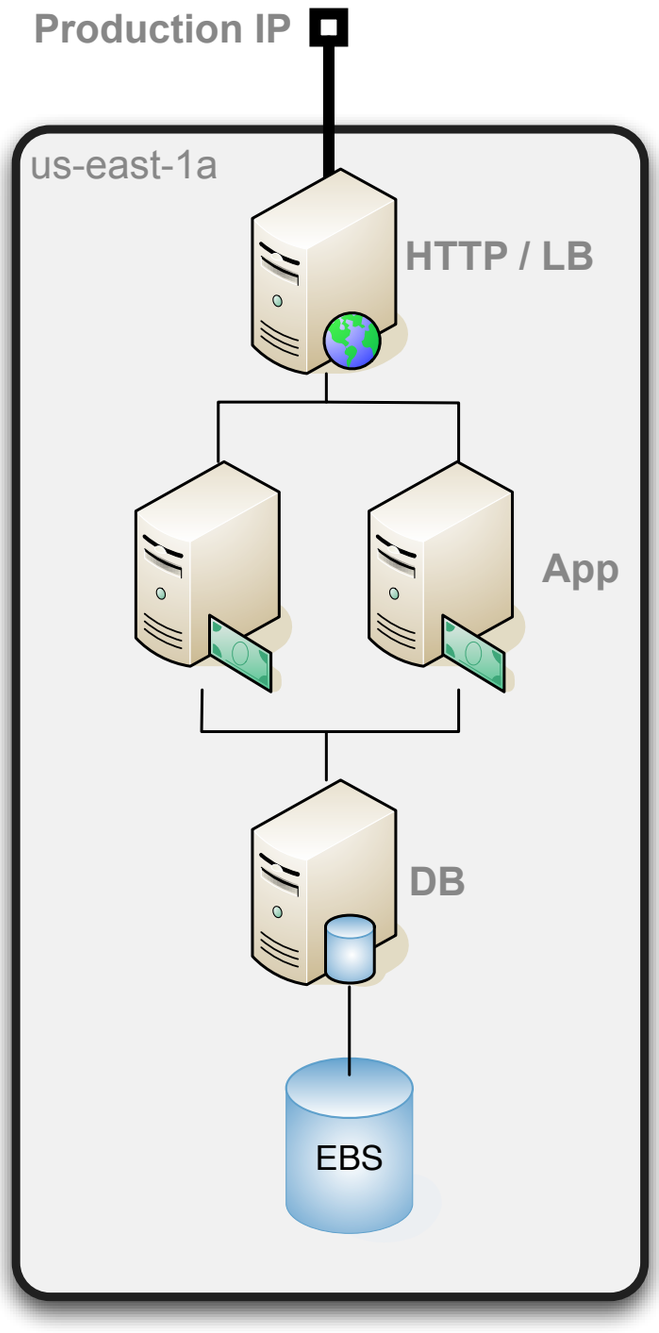

EBS volumes persist between VM restarts

Mitigating The Down Side of Ephemerality

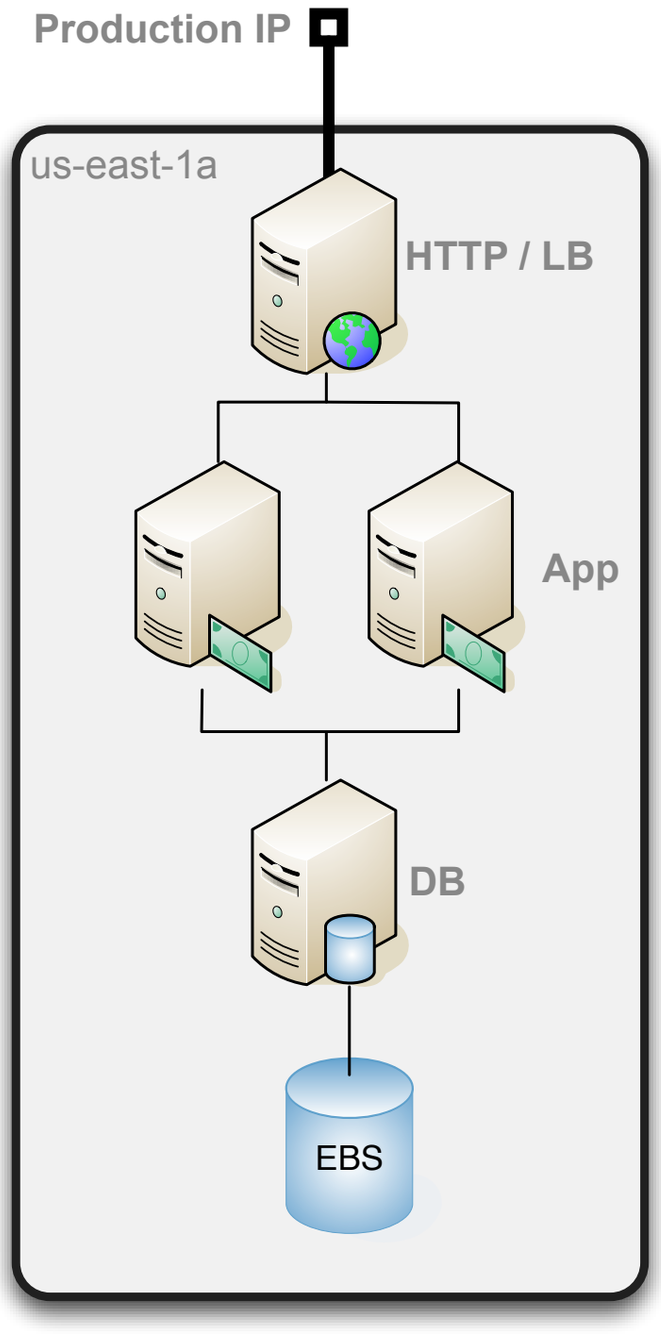

EBS volumes persist between VM restarts but can still be lost.

Mitigating The Down Side of Ephemerality

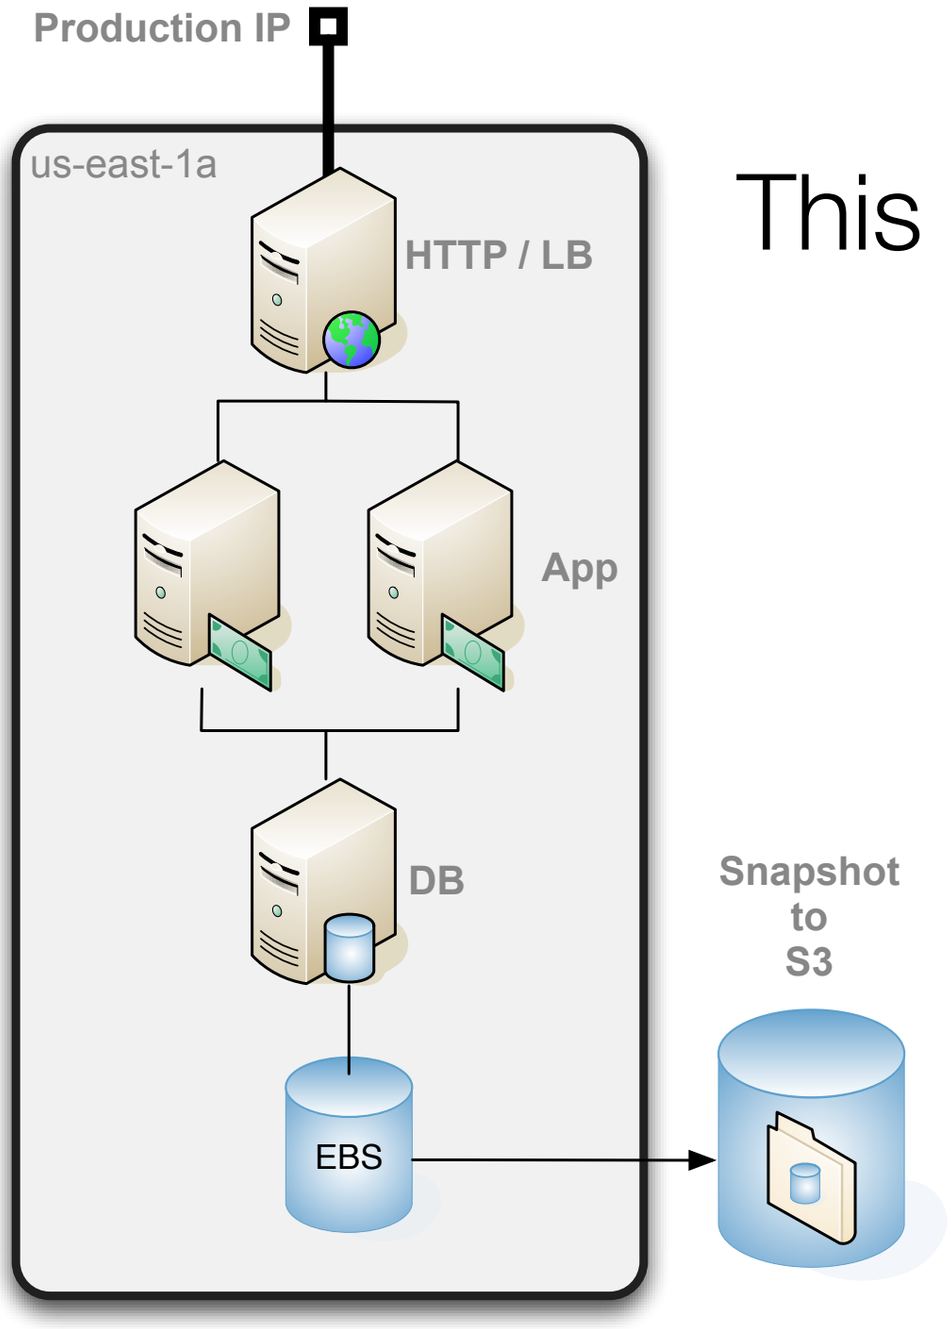

This keeps our data

Mitigating The Down Side of Ephemerality

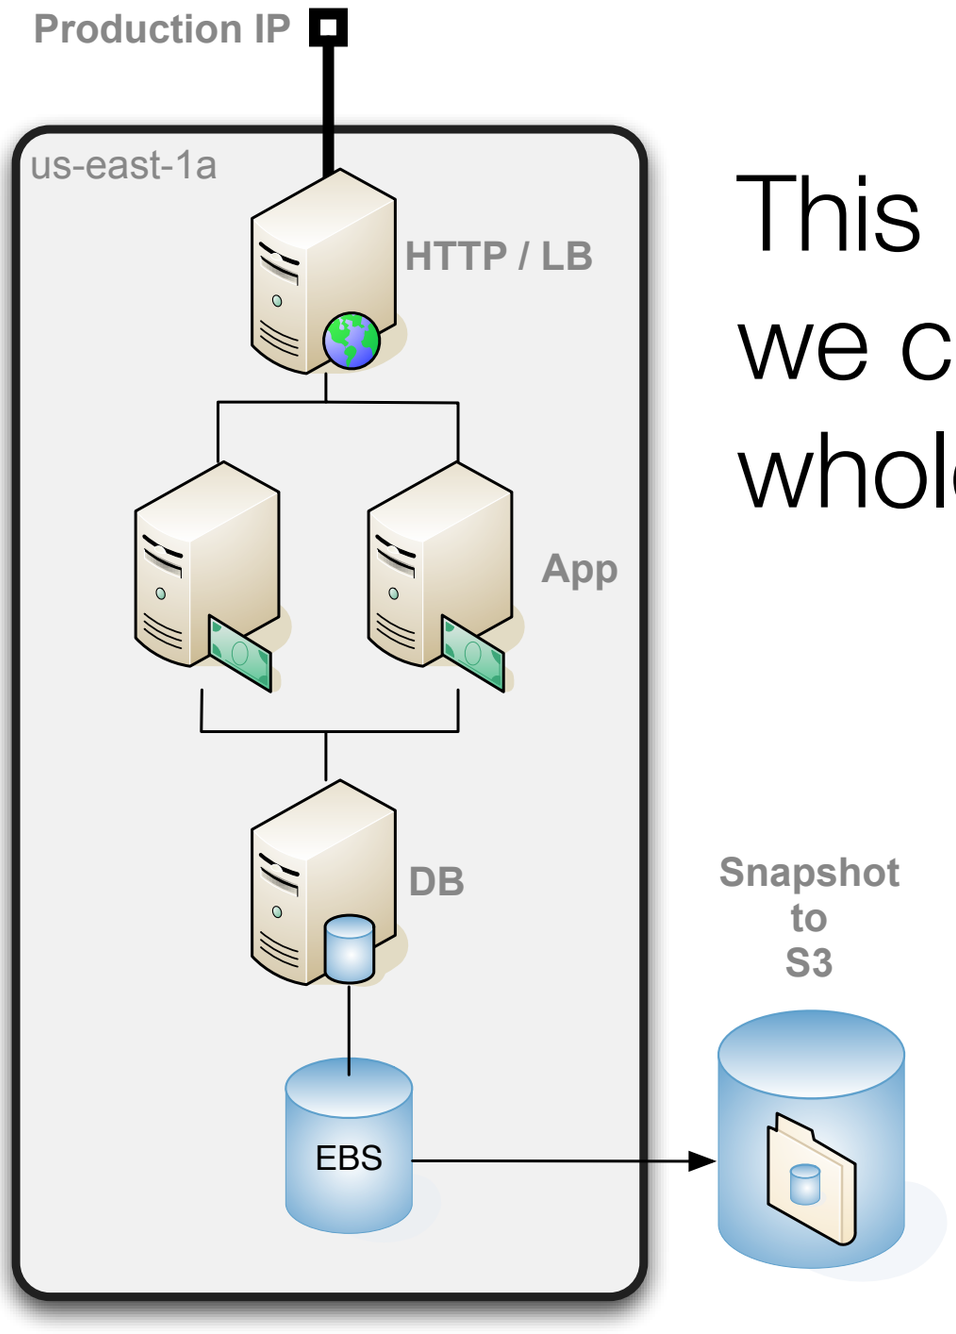

This keeps our data, but we could still lose the whole availability zone.

Mitigating The Down Side of Ephemerality

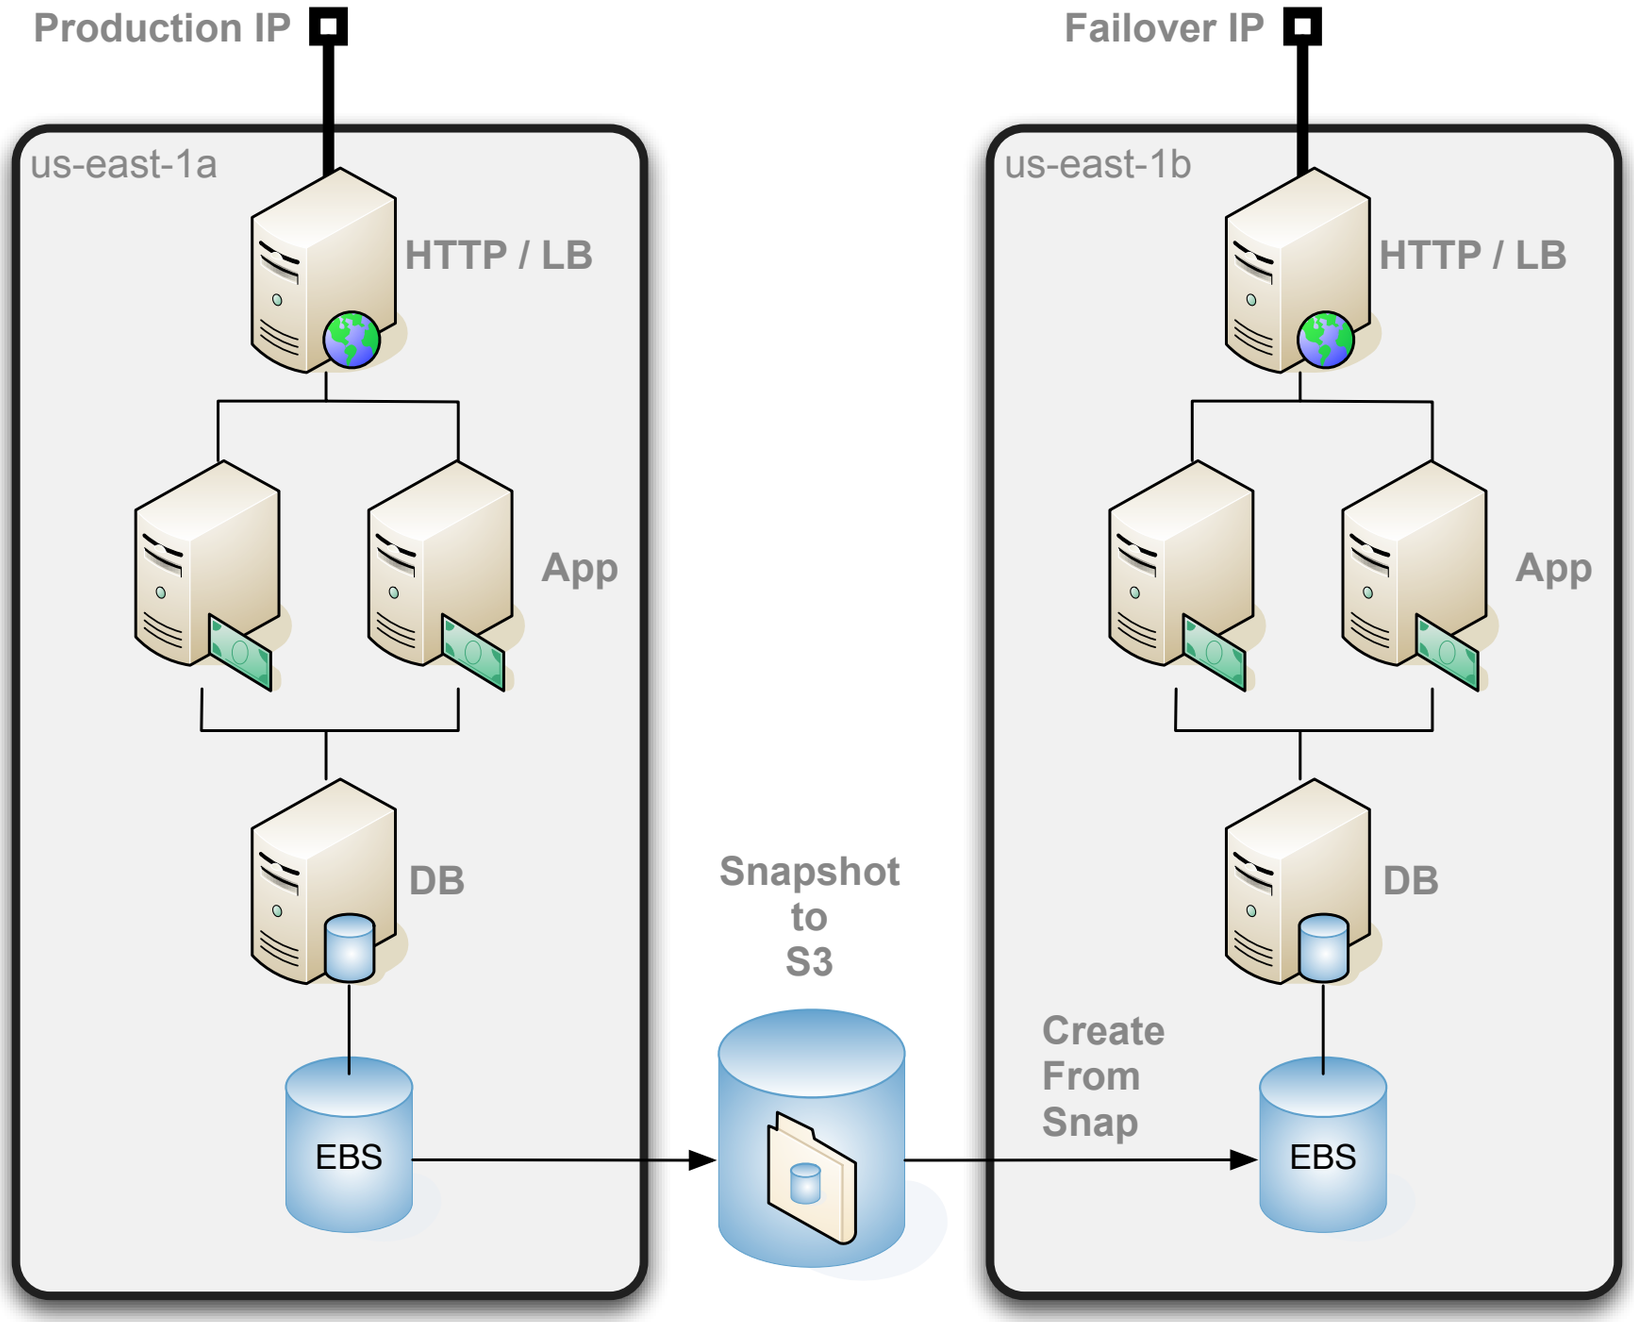

Mitigating The Down Side of Ephemerality

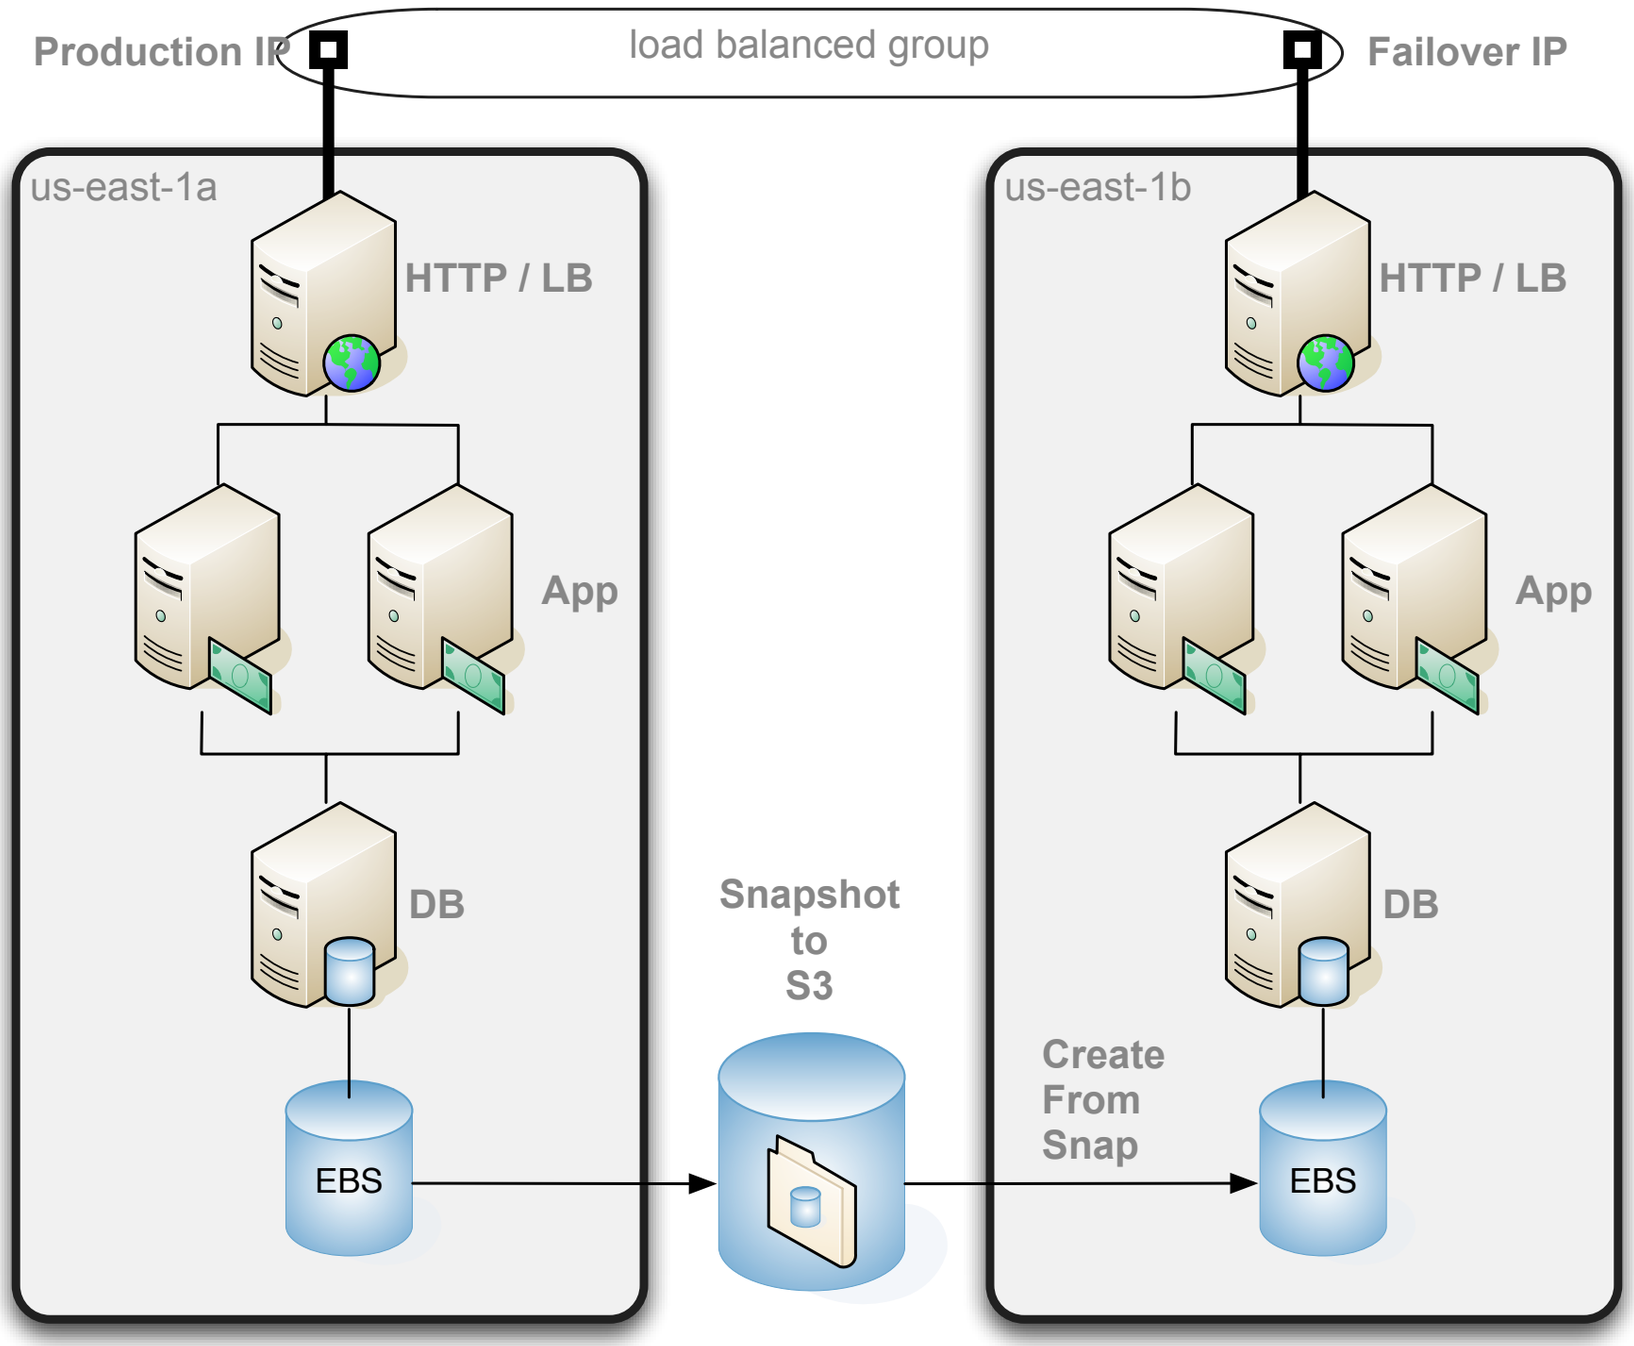

November 19, 2009

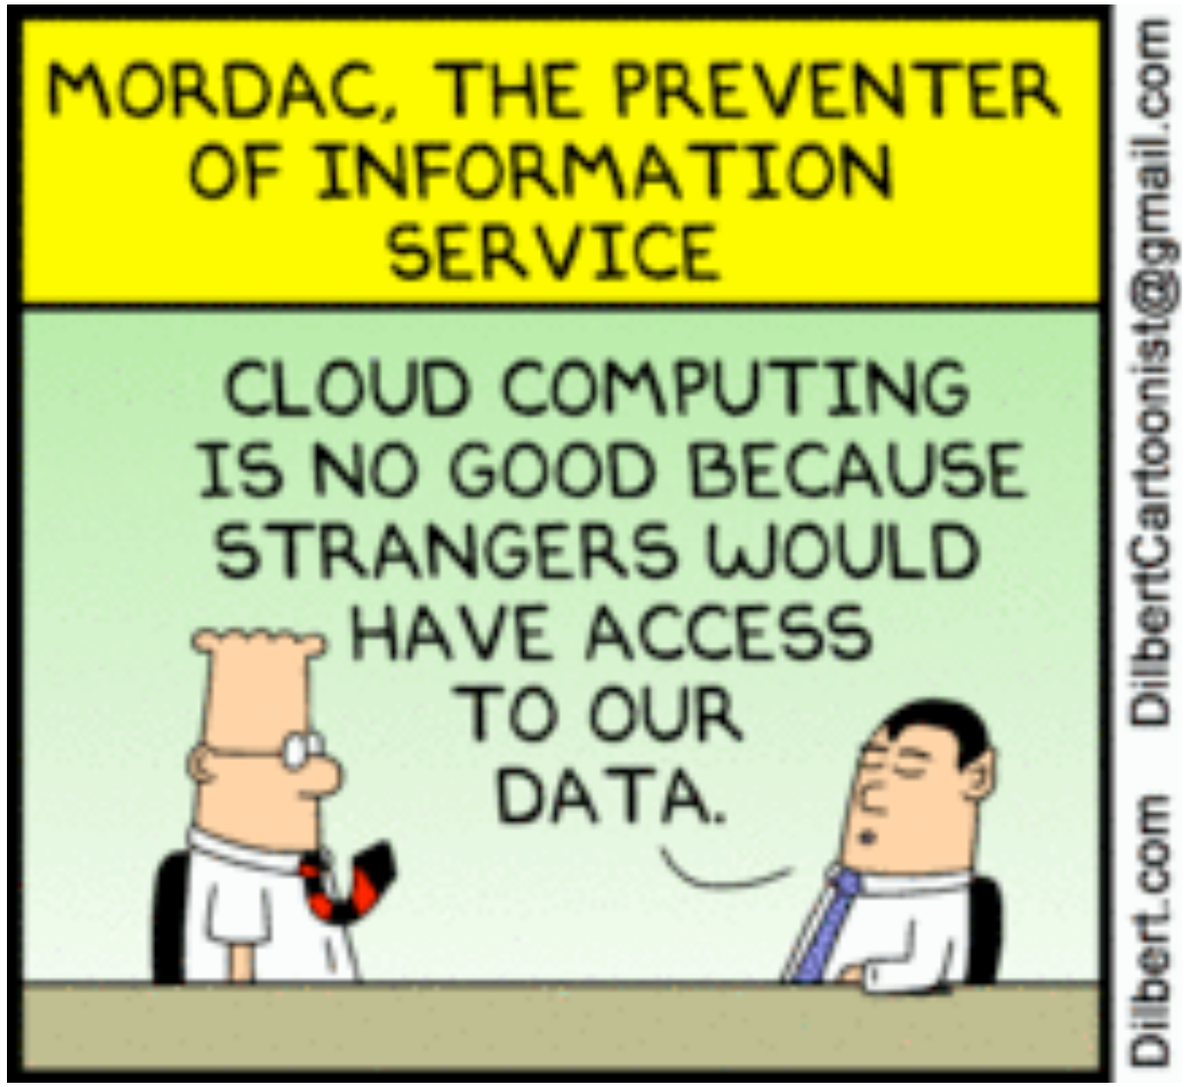

Perceived Risks

- Confidentiality
- Compliance
- Control

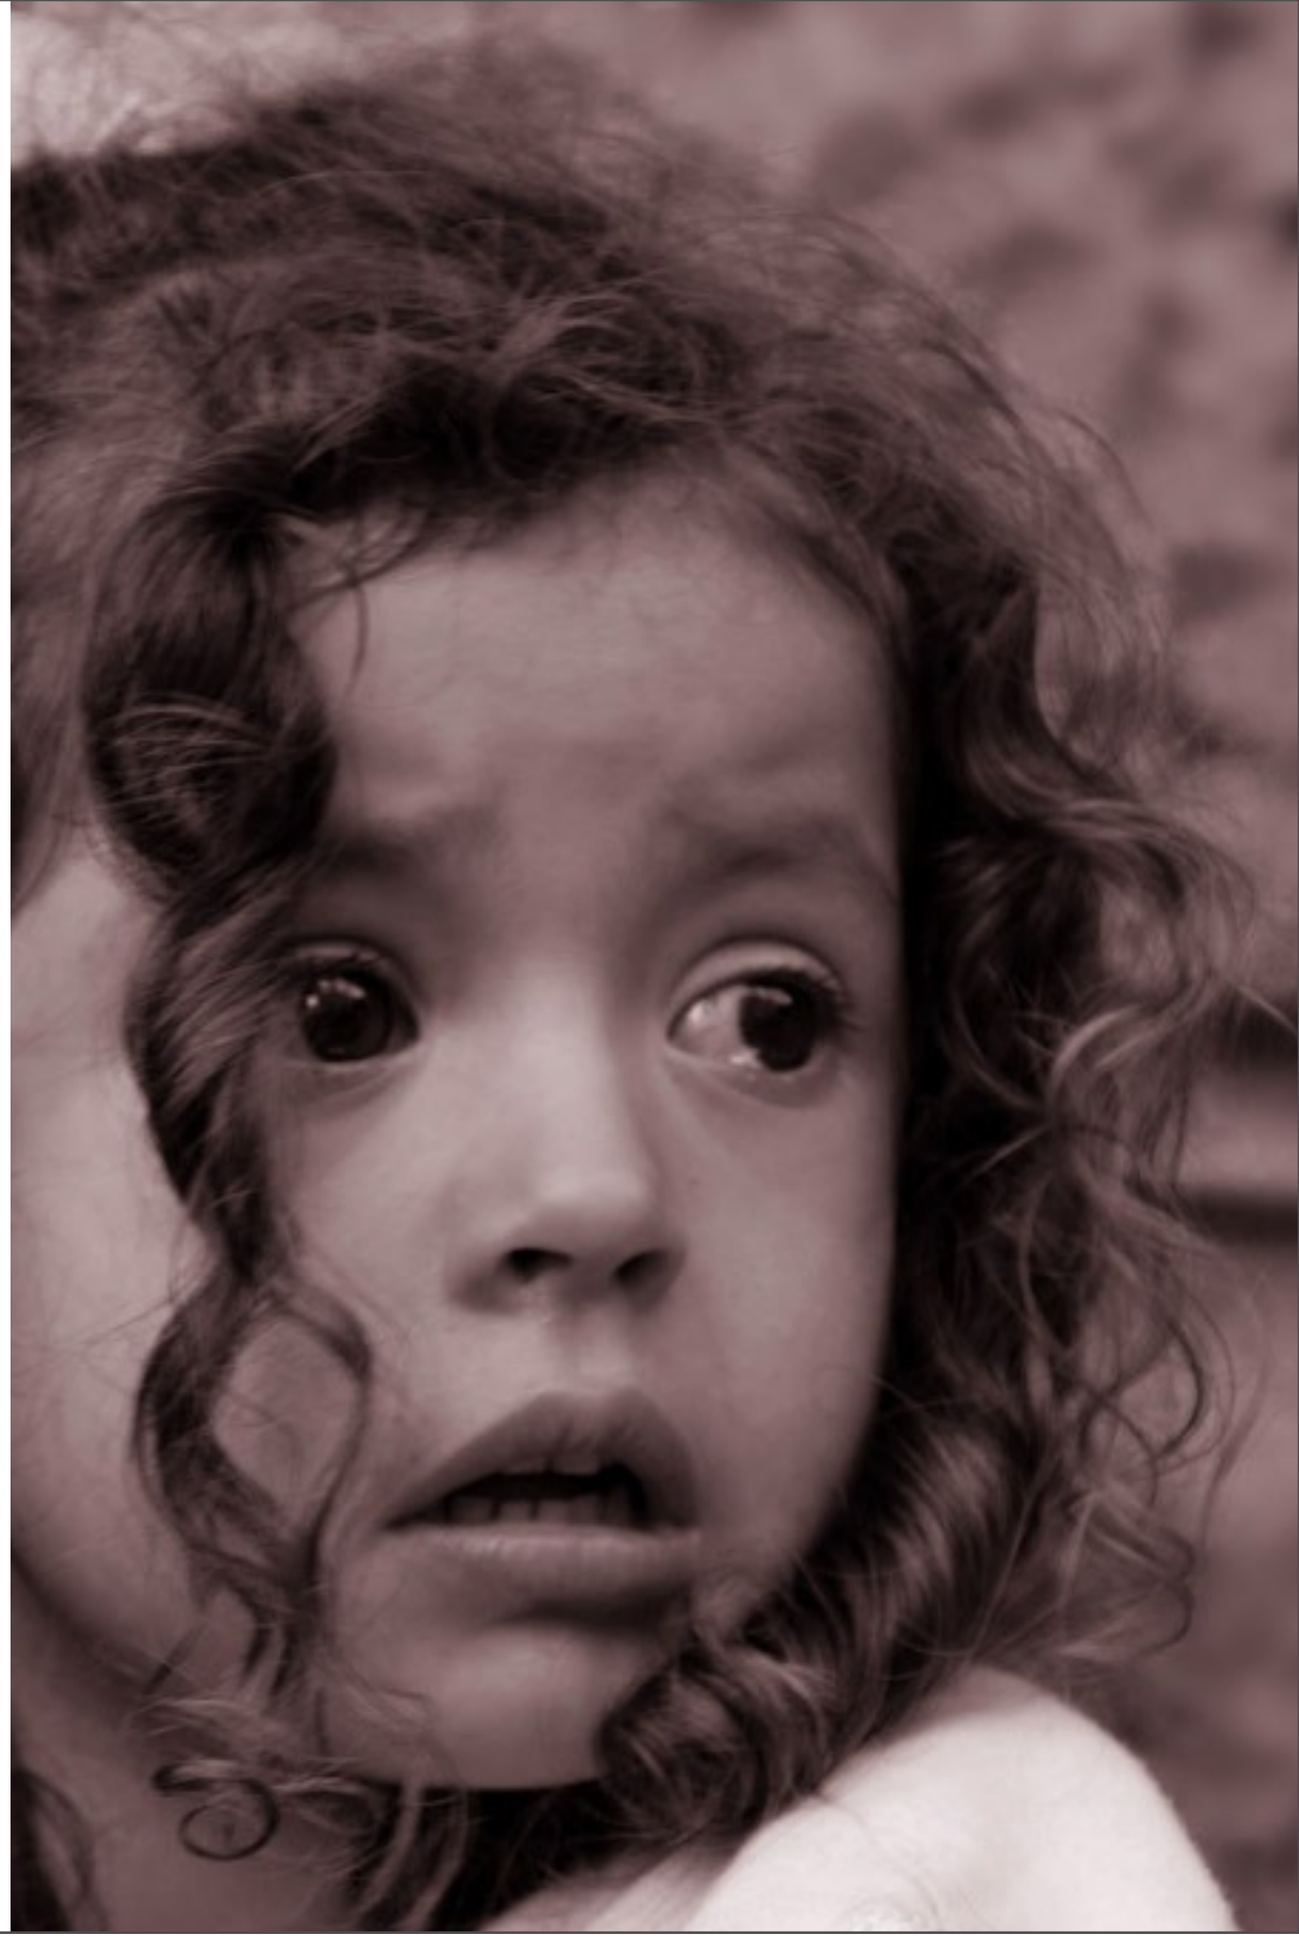

Real Risks

- Control plane threats
- Patent shutdowns
- Lack of risk management info

Security Facilities to Apply

Vulnerability	Countermeasure

Security Facilities to Apply

Vulnerability	Countermeasure
Data in Motion	Transport level encryption *(not new)*

Security Facilities to Apply

Vulnerability	Countermeasure
Data in Motion	Transport level encryption *(not new)*
Data at Rest	Storage level encryption *(not new)*

Security Facilities to Apply

Vulnerability	Countermeasure
Data in Motion	Transport level encryption *(not new)*
Data at Rest	Storage level encryption *(not new)*
Access Control	Security Groups *(comparable to VLANs & firewalls)*

Security Facilities to Apply

Vulnerability	Countermeasure
Data in Motion	Transport level encryption *(not new)*
Data at Rest	Storage level encryption *(not new)*
Access Control	Security Groups *(comparable to VLANs & firewalls)*
Control Plane	Not under your control

Security Facilities to Apply

Vulnerability	Countermeasure
Data in Motion	Transport level encryption *(not new)*
Data at Rest	Storage level encryption *(not new)*
Access Control	Security Groups *(comparable to VLANs & firewalls)*
Control Plane	Not under your control
Multitenancy Concerns	Not under your control

Recently Announced

- On November 12, 2009, Amazon announced that AWS has received it's SAS 70 Type II attestation.
  http://is.gd/4Yclv

- Cloud Security Alliance formed the A6 working group to develop a specification for declarative risk information.
  http://www.scribd.com/Iron%20Fog

Summary

- Cloud Computing brings together four important trends in infrastructure:
  - Virtualization
  - Commoditization of hardware
  - Horizontally scalable architecture
  - Need for rapid provisioning
- It offers significant operational, cost and time-to-market advantages
- It also necessitates changes in architecture

Questions?

Michael Nygard
www.michaelnygard.com