

SUMMIT

**JBoss
WORLD**

PRESENTED BY RED HAT

**LEARN. NETWORK.
EXPERIENCE OPEN SOURCE.**

www.theredhatsummit.com

Cloud-Scale Applications using Redundant Arrays of Inexpensive Nodes (RAIN)

Rob Hirschfeld

Sr. Cloud Solutions Architect, Dell Inc.

SUMMIT

JBoss
WORLD

PRESENTED BY RED HAT



Agenda

Situation:

Reliable applications on unreliable infrastructure

Virtual Integrated System:

Path to simplification

Service Profiles:

Cloud-based applications

Calls to Action:

Modify architectures to stripe workloads

SUMMIT

JBoss
WORLD

PRESENTED BY RED HAT



Products Highlighted

Cloud Technologies

Dell Cloud Solutions

Joyent Smart DataCenter

Zeus ZXTM



Application Technologies

JBoss Clustering



SUMMIT

JBoss
WORLD

PRESENTED BY RED HAT





There are

TWO ROADS TO THE CLOUDS

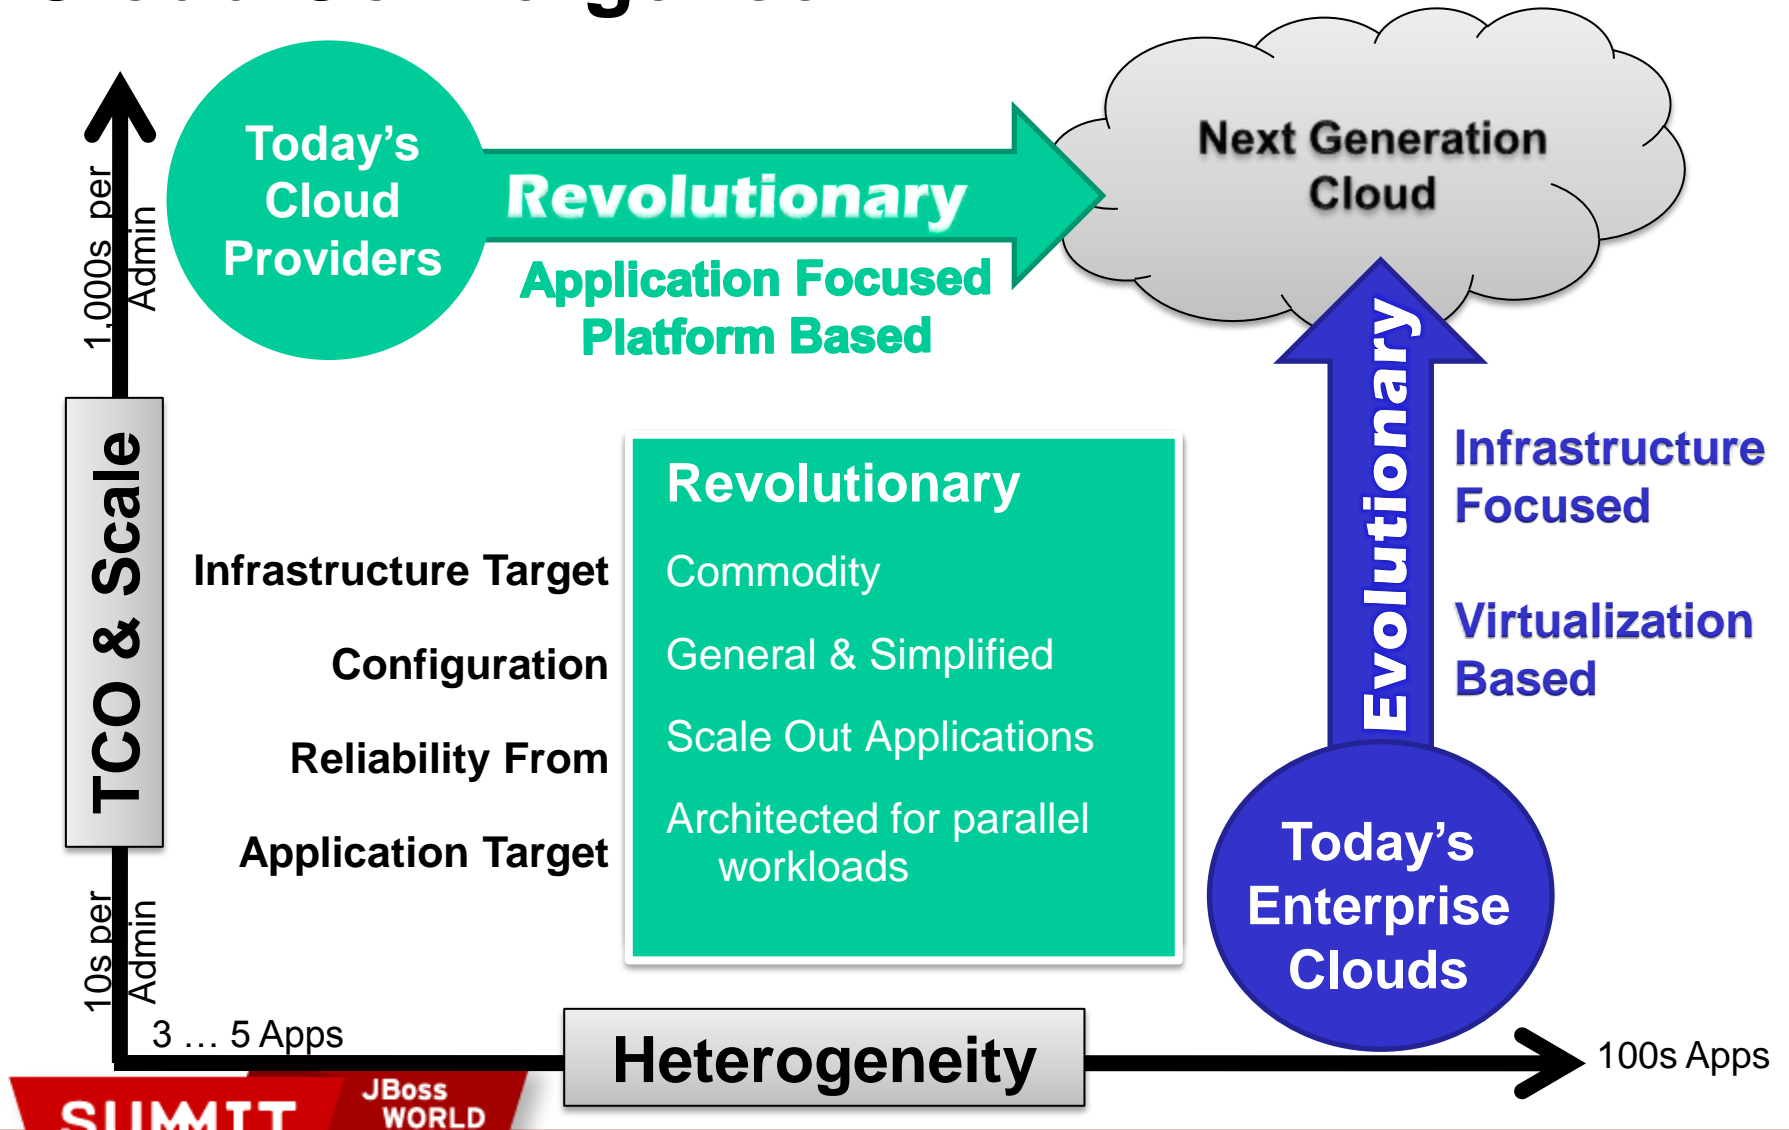
SUMMIT

**JBoss
WORLD**

PRESENTED BY RED HAT



Cloud Convergence



SUMMIT

**JBoss
WORLD**

PRESENTED BY RED HAT





First, we need to think about

REVOLUTIONARY INFRASTRUCTURE

SUMMIT

**JBoss
WORLD**

PRESENTED BY RED HAT



Hardware Redundancy is...

- Expensive
 - Requires premium equipment
 - Often duplicated (or N+1) capacity
- Complex
 - Delicate configurations requiring expert training
 - Sophisticated logic
- Scale Up, not out
 - Adding discrete capacity adds exponential complexity
- Hard to Verify
 - Faults are unusual and may not be verifiable in production

SUMMIT

**JBoss
WORLD**

PRESENTED BY RED HAT



What is Revolutionary Infrastructure?

Infrastructure:

Commodity

General & Simplified

Applications:

Scale Out Applications

Architected for parallel workloads

SUMMIT

JBoss
WORLD

PRESENTED BY RED HAT






Commodity, General & Simplified

- Lower cost
 - Strip away all packaging from CPU, RAM & Disk
 - Scrutinize every infrastructure investment
- Able to scale quickly
 - Easy to incrementally add more resources
 - Burst to other service providers
- Platform agnostic
 - Does not care about hardware or operating system
 - Does not have fragile networking or configuration












Two Approaches to Handling Failures

Evolutionary (*system SLA*)






Revolutionary (*application SLA*)



Avoid Failure:

Loss is significant impact




(larger servers = greater % loss)



Ignore Failure:

Lost is immaterial impact

(commodity servers = more likely loss)



SUMMIT

JBoss
WORLD




PRESENTED BY RED HAT



Two Recovery Modes

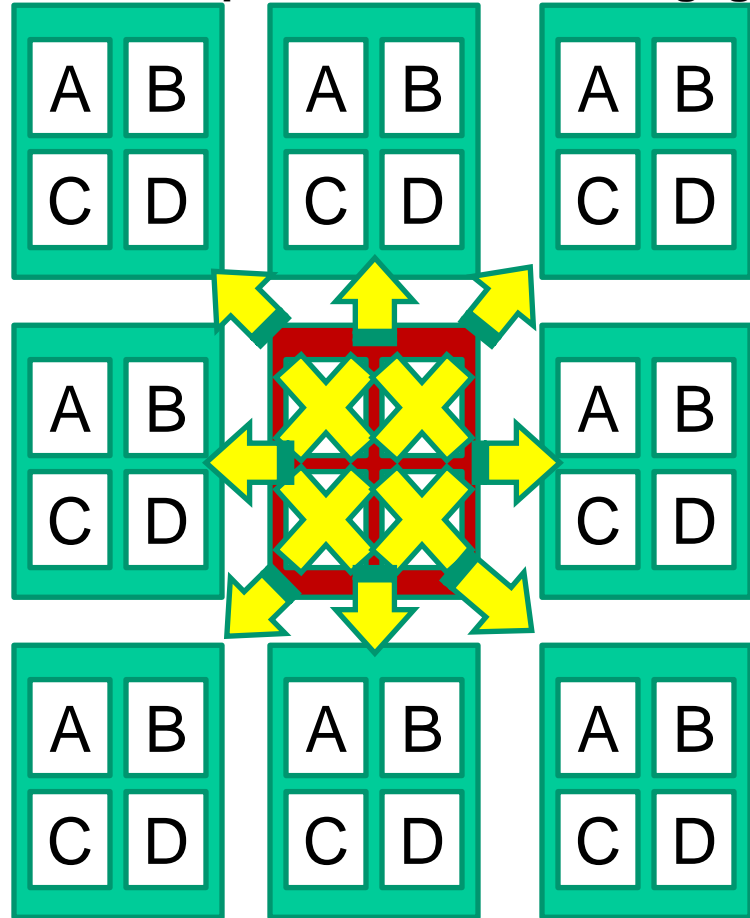
Kill Low Priority Application (D)

Failure causes oversubscription



All operations continue

Failure impact distributed & negligible



SUMMIT

JBoss
WORLD

PRESENTED BY RED HAT



Rules for Revolutionary Infrastructure

1. Complexity does not scale
2. Failure is inevitable



SUMMIT

**JBoss
WORLD**

PRESENTED BY RED HAT



Now, let's talk about

REVOLUTIONARY APPLICATIONS



SUMMIT

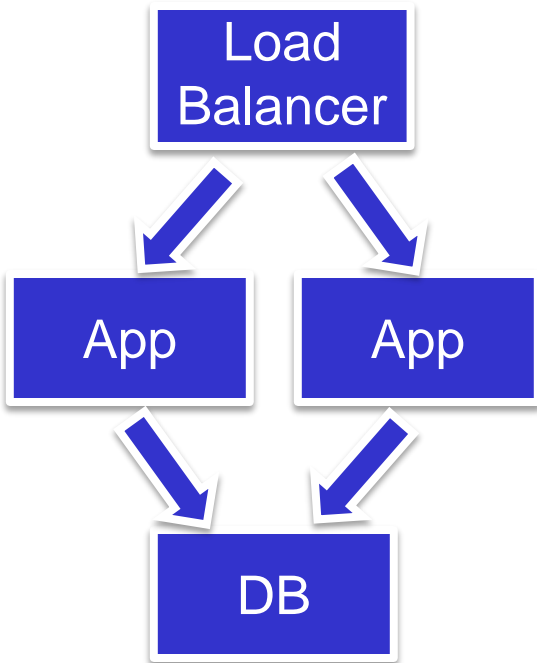
**JBoss
WORLD**

PRESENTED BY RED HAT

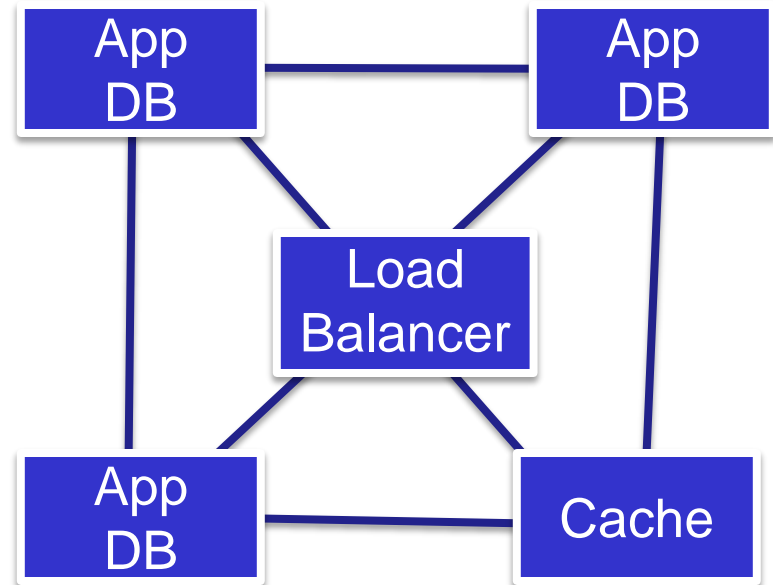


Two Architecture Models

Hardware: Strict Tiers



Software: Flexible Tiers



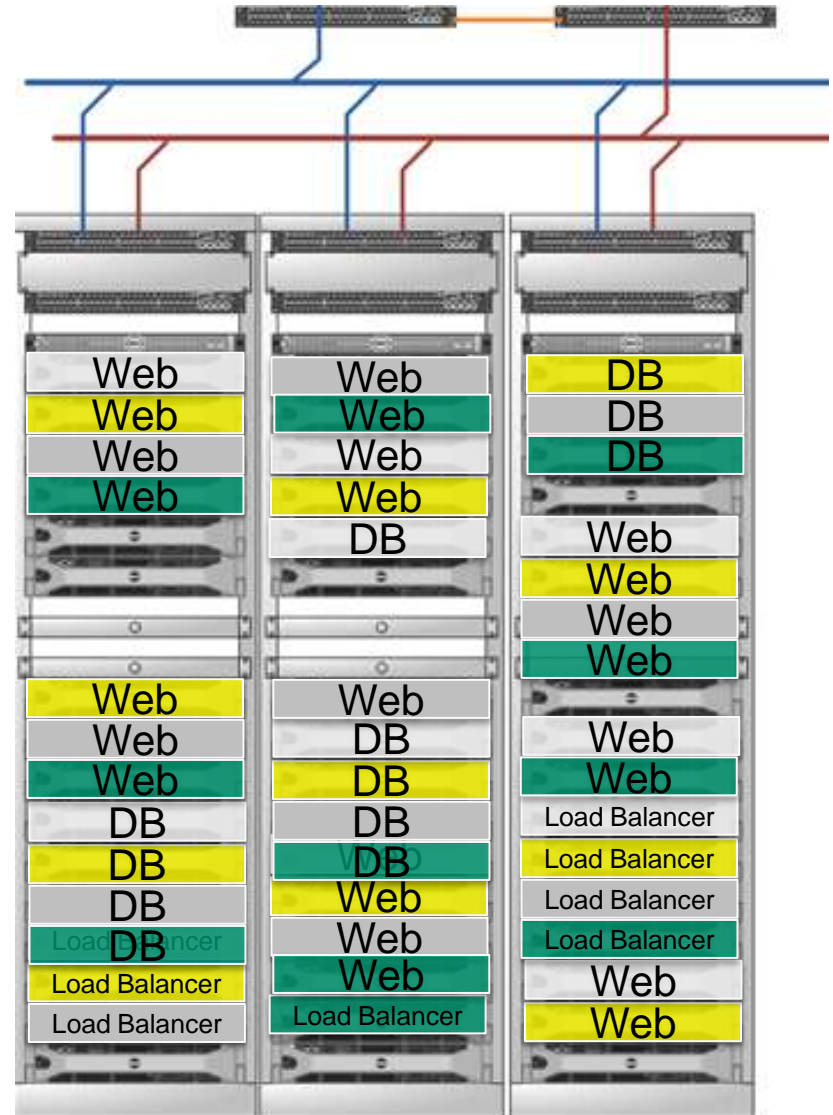
Scale-Out Deployment

To distribute load and improve fault tolerance, application components are distributed throughout the node infrastructure.

Eliminates single points of failure

Maximizes burst capability

Enables failure isolation and fault tolerance



SUMMIT

JBoss
WORLD

PRESENTED BY RED HAT



When hardware is expected to fail, Applications must

Eliminate single points of failure

Store data redundantly

Eliminate locks & races

Self-monitor state

Recover automatically from failure

SUMMIT

JBoss
WORLD

PRESENTED BY RED HAT





Scale-out applications use a

REDUNDANT ARRAY OF INEXPENSIVE NODES (RAIN)

SUMMIT

**JBoss
WORLD**

PRESENTED BY RED HAT



RAIN means Application Striping

RAIN assumes that infrastructure nodes are:

- **Idle** because cloud traffic loads are cyclic
- **Unreliable** because of simplified commodities
- **Divisible** because renting ten 1/10ths of a server costs roughly the same as a whole one
- **Burstable** because 1/10th servers can sometimes consume that extra 9/10th server

RAIN concepts maps to RAID storage concepts


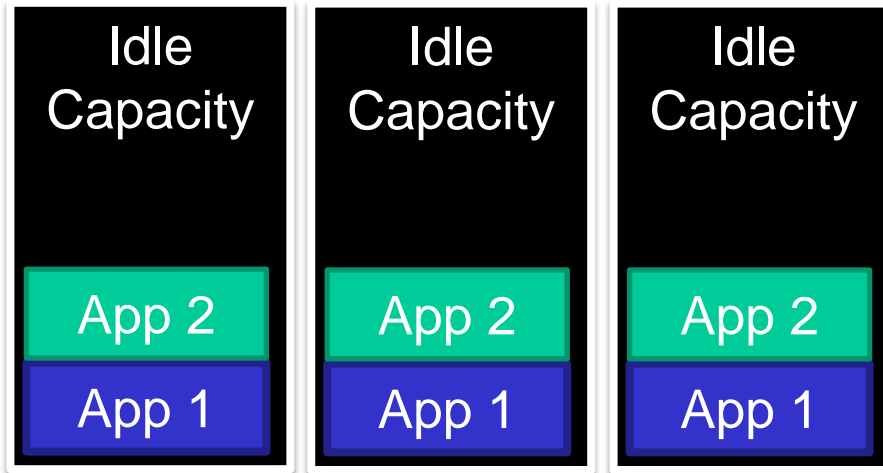
SUMMIT

JBoss
WORLD

PRESENTED BY RED HAT



Divisible + Burstable = Scalable



Normally $3 \times 25\% = 75\%$

Bursts to $3 \times 75\% = 225\%$



SUMMIT

JBoss
WORLD

PRESENTED BY RED HAT



Divisible + Burstable = Reliable



Normally $3 \times 25\% = 75\%$

Recovers as $2 \times 37\% = 75\%$

SUMMIT

JBoss
WORLD

PRESENTED BY RED HAT





RAIN 0: Spanning

Benefit:

Improves workload capacity

Problem:

Does not address reliability



This is the typical application architecture that drives IT departments to create complex, highly redundant & robust infrastructure.

SUMMIT

JBoss
WORLD

PRESENTED BY RED HAT



RAIN 1: Mirroring

Active/Passive Architecture

Benefit

Improves fault tolerance


Problems

More complex to configure

More complex to manage

Passive resources are wasted overhead

Scale is limited to capacity of single node



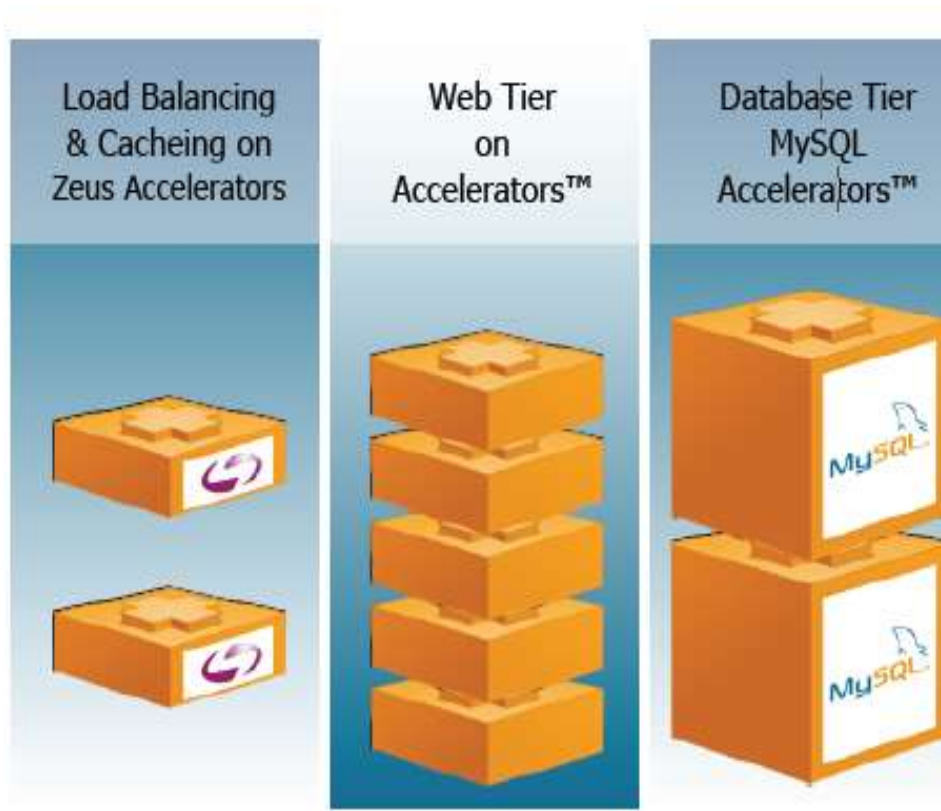
SUMMIT

JBoss
WORLD

PRESENTED BY RED HAT



Classic Scale-Out Architecture



SUMMIT

JBoss
WORLD

PRESENTED BY RED HAT



RAIN 5: Striping

Benefit

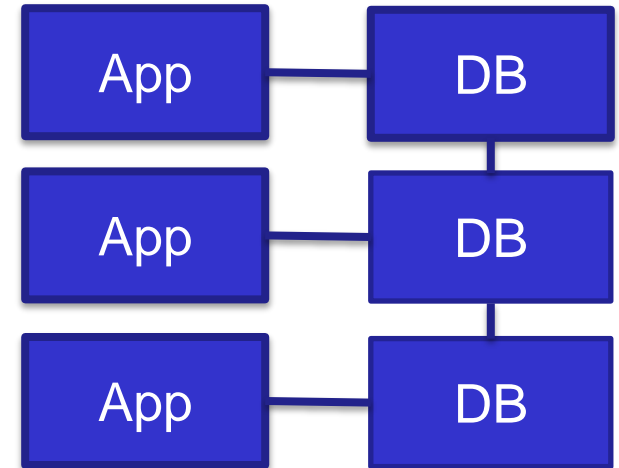
Highly fault tolerant

Scales out gracefully

Problems

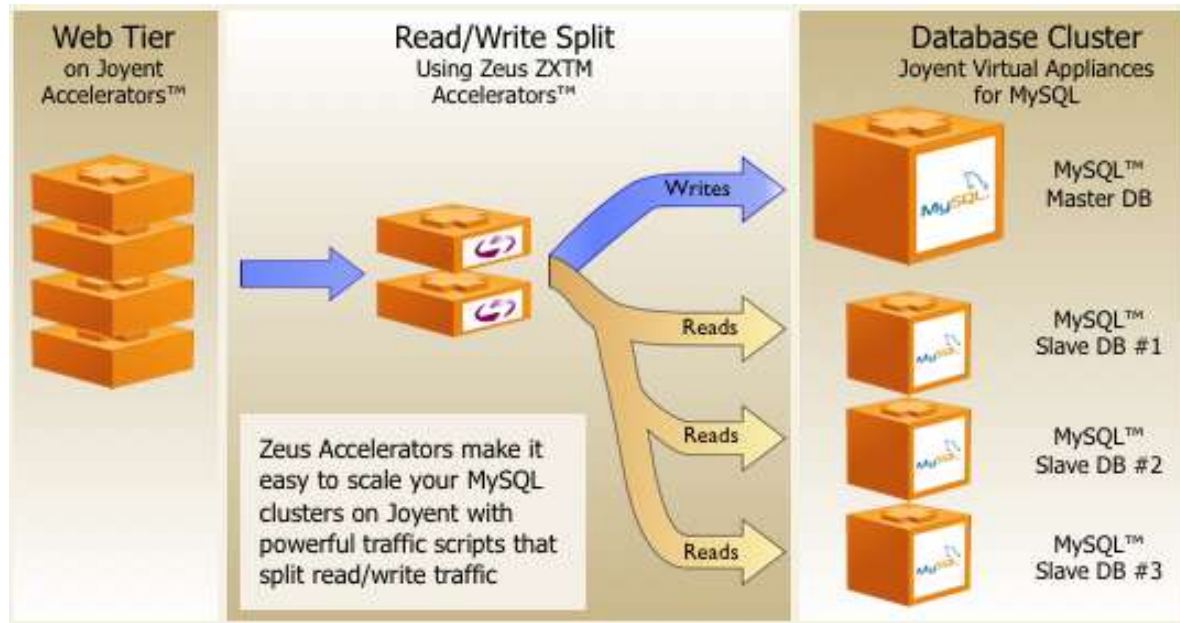
Requires data synchronization

May require application changes



RAIN 5: Example of Scale-Out Database

The Cloud Solution for Web Applications can use Zeus ZXTM to scale-out databases by intelligently splitting SQL reads (SELECT) and writes (UPDATE/INSERT/DELETE)



SUMMIT

**JBoss
WORLD**

PRESENTED BY RED HAT



RAIN 5: Striping + Cache

Benefit

Performance benefits

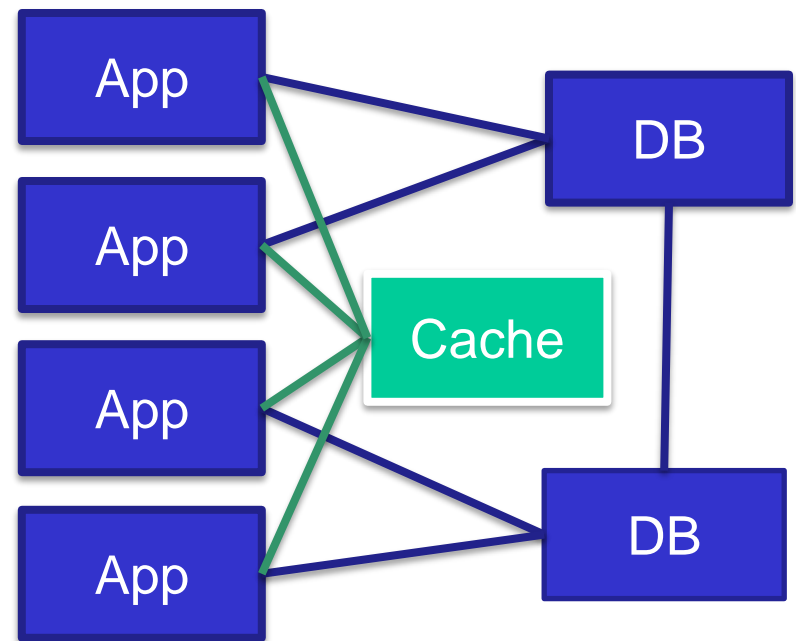
Addresses locking issues

Reduces database loads

Problems

Require application changes

Potential for cache failure



SUMMIT

JBoss
WORLD

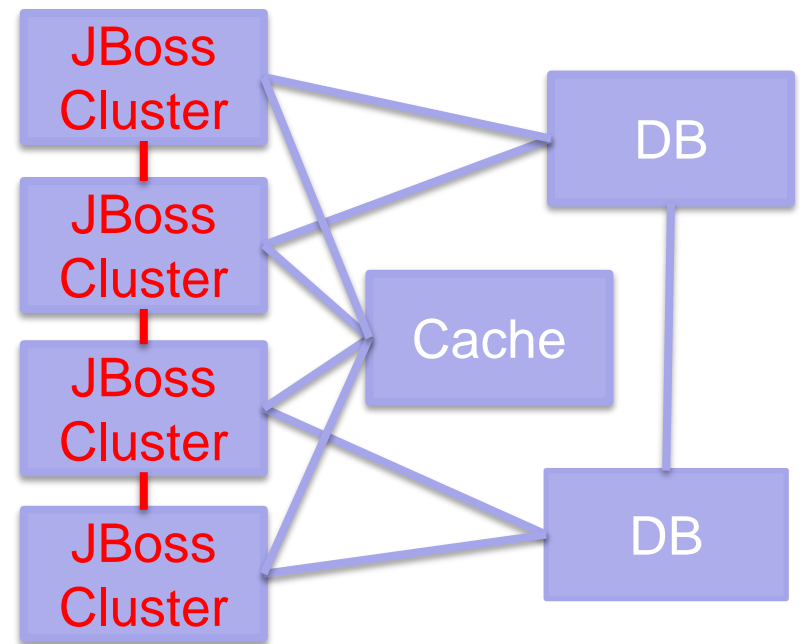
PRESENTED BY RED HAT



JBoss Clustering

Features

- Automatic discovery. Nodes in cluster find each other with no additional configuration.
- Cluster-wide replicated JNDI context
- Failover and load-balancing for JNDI, RMI and all EJB types
- Stateful Session Bean state replication
- HTTP Session replication
- High Availability JMS
- Farming. Distributed deployment of JBoss components. Deploying on one node deploys on all nodes.



RAIN JBON: Just a Bunch Of Nodes

Mixed Function, Inter-changable Units

Benefit

Highly scalable

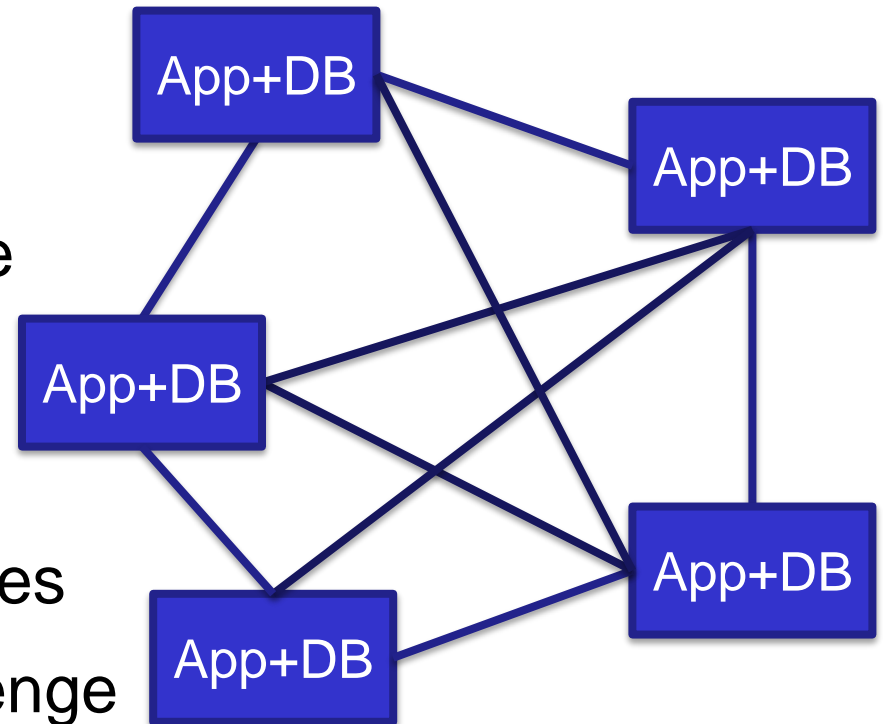
Geographically distributable

Highly fault tolerant

Problems

Requires application changes

Data synchronization challenge



SUMMIT

JBoss
WORLD

PRESENTED BY RED HAT



RAIN JBON + Cache


Benefit

Performance benefits

Addresses synchronization

Problems

Potential for cache failure



SUMMIT

JBoss
WORLD

PRESENTED BY RED HAT



Benefits of Virtualized Load Balancers


Dell/Joyent Cloud Solution features virtualized load balancers.

Dedicated virtual load balancers per application set is key item for isolation and flexibility.

Allows implementing active-active load balancer even small apps.

Avoids:

- Added complexity from shared load physical balancers
- Dependencies between applications
- Cost steps in scaling infrastructure
- Networking bottlenecks



SUMMIT

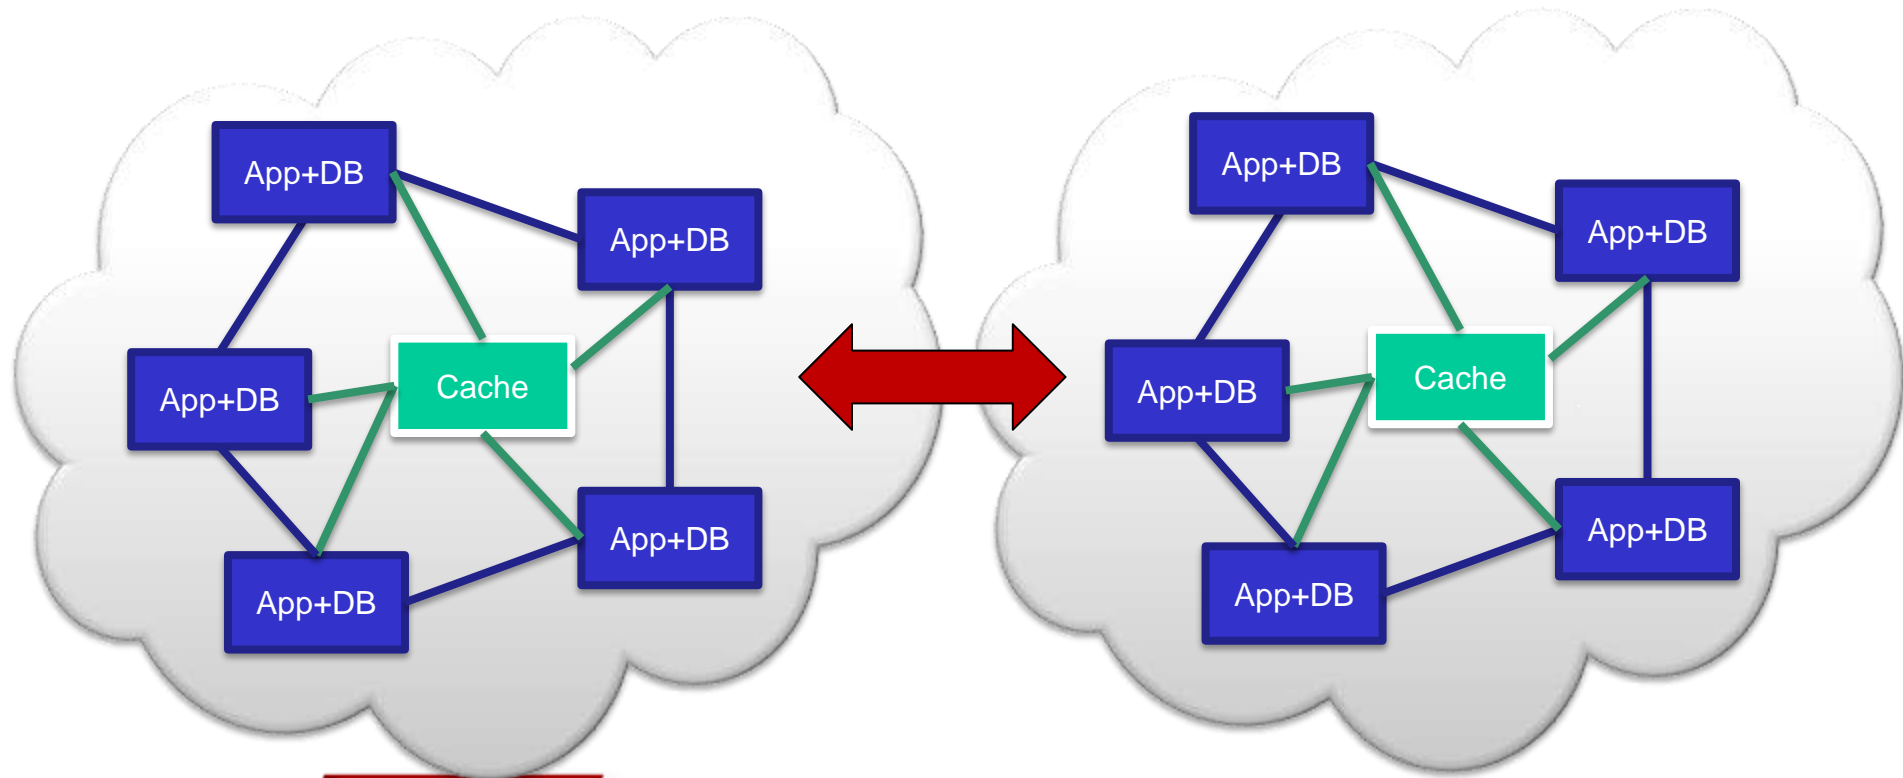
JBoss
WORLD

PRESENTED BY RED HAT



RAIN 200: Multi-Cloud

More modular designs scale to multiple clouds



SUMMIT

JBoss
WORLD

PRESENTED BY RED HAT



Cloud Architecture Best Practices for RAIN 5+

- Distinguish between cache-able and persisted data
- Consider non-SQL, replicating databases
- Use SOA to partition functionality
- Simplify aggressively
- Prepare and test for failure
- Isolate & distribute functions
- Architect with Dev-Ops in mind





Dark Cycle Compute &
SCALE-OUT DEMO

SUMMIT

JBoss
WORLD

PRESENTED BY RED HAT



Presented by

ROB HIRSCHFELD

Sr. Cloud Solution Architect

Rob_Hirschfeld@Dell.com

Blog: <http://RobHirschfeld.com>



SUMMIT

JBoss
WORLD

PRESENTED BY RED HAT



FOLLOW US ON TWITTER

www.twitter.com/redhatsummit

TWEET ABOUT IT

[#summitjbw](https://twitter.com/summitjbw)

READ THE BLOG

<http://summitblog.redhat.com/>

SUMMIT

JBoss
WORLD

PRESENTED BY RED HAT

