

RED HAT
SUMMIT

LEARN. NETWORK.
EXPERIENCE OPEN SOURCE.

June 11-14, 2013
Boston, MA



RED HAT
SUMMIT


DEFINING THE FUTURE

WITH RED HAT STORAGE GEO-REPLICATION

Dustin L. Black, RHCA
Sr. Technical Account Manager, Red Hat
2013-06-13

Mohit Anchlia
Data Platform Architect, Intuit

Small file workload?
20 MB/s writes?
30-minute geo-replication SLA?
Commodity hardware?



CHALLENGE ACCEPTED

via Rage Comics



Dustin L. Black, RHCA
Sr. Technical Account Manager
Red Hat, Inc.

dustin@redhat.com
@dustinblack




redhat.
CERTIFIED
ARCHITECT


Scale-out Performance, Capacity and Availability

Scale-up Capacity


Red Hat Storage
Server for On-premise



Red Hat Storage
Server for On-premise



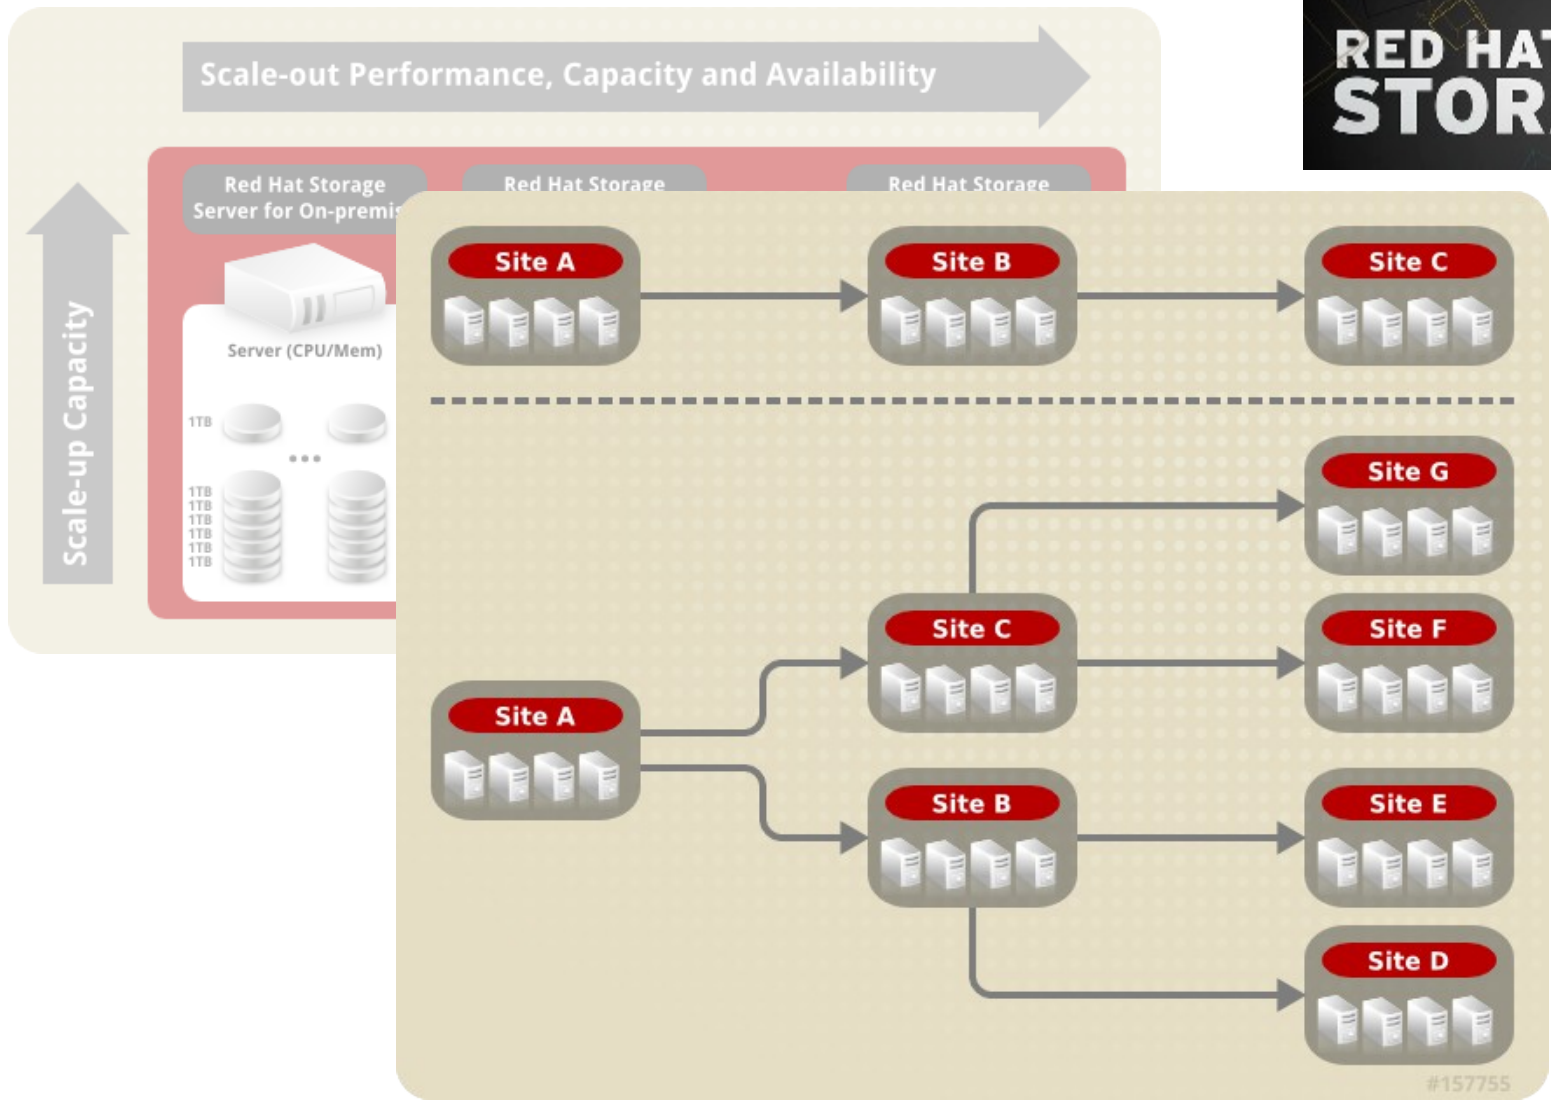
Red Hat Storage
Server for On-premise

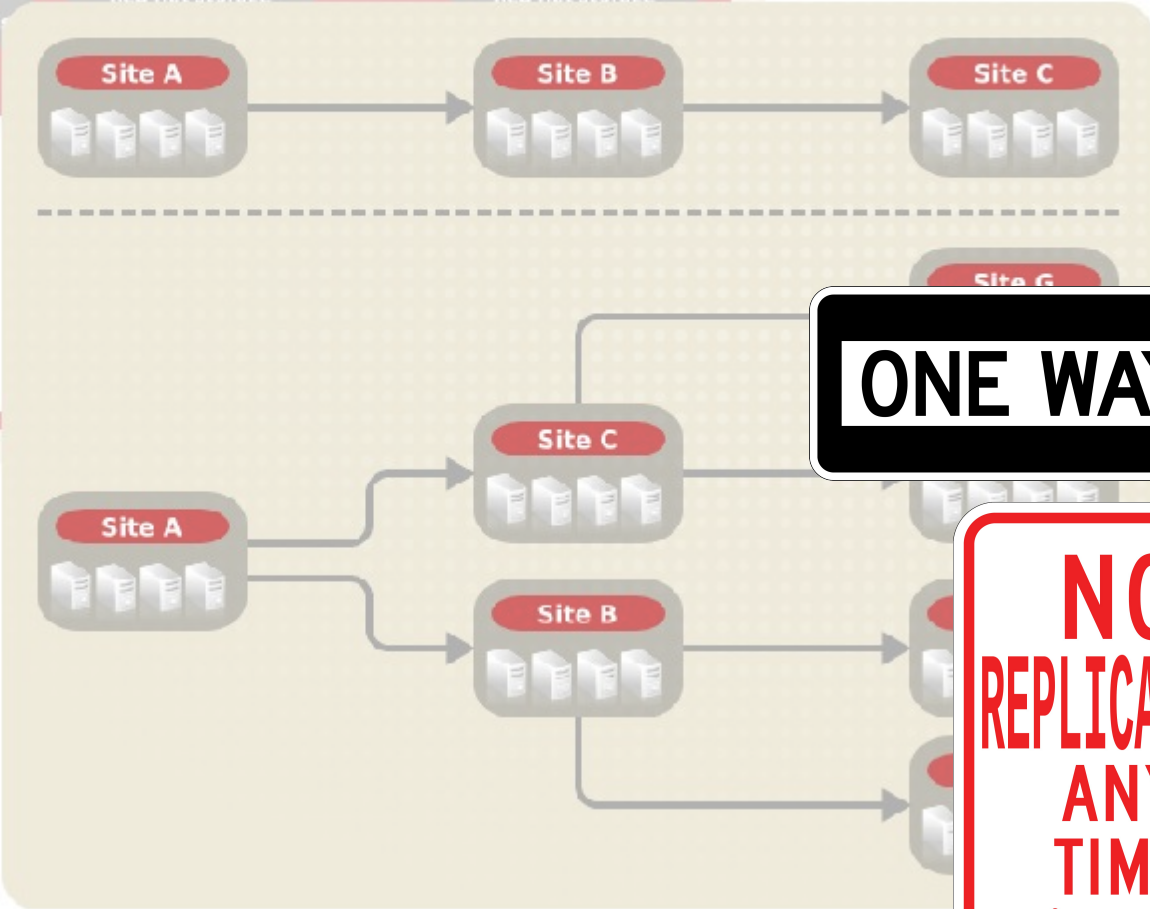
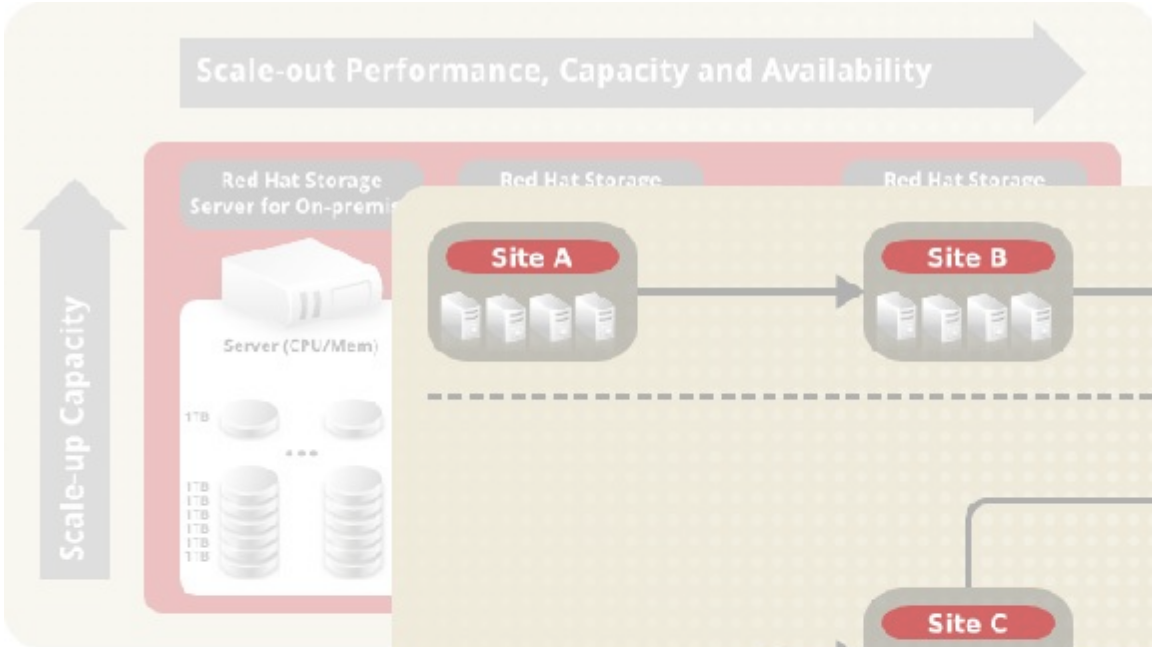


#145075



RED HAT® STORAGE






ONE WAY →

**NO
REPLICATING
ANY
TIME** ←

Scale-out Performance, Capacity and Availability

**RED HAT®
STORAGE**



intuit

TurboTax 





Mohit Anchlia
Data Platform Architect
Intuit, Inc.

mohit_anchlia@intuit.com

intuit®

Background

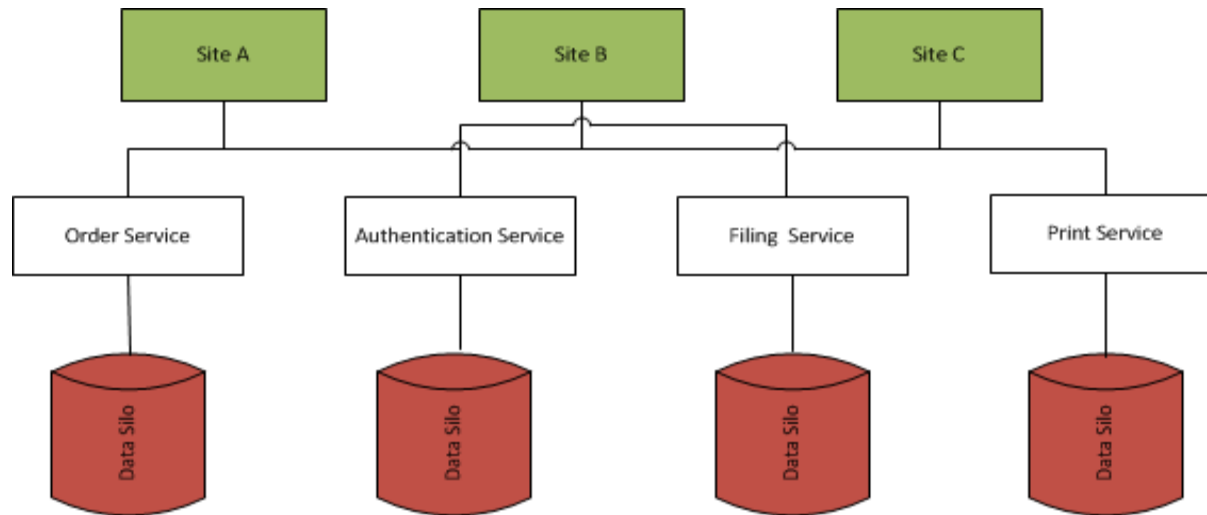
- Intuit is maker of TurboTax, Quicken, Quickbooks and many other products for SBUs.
- Many services work together to deliver awesome product experience

Intuit Proprietary & Confidential

intuit.

Problem Statement (Service explosion)

- Service explosion over the years

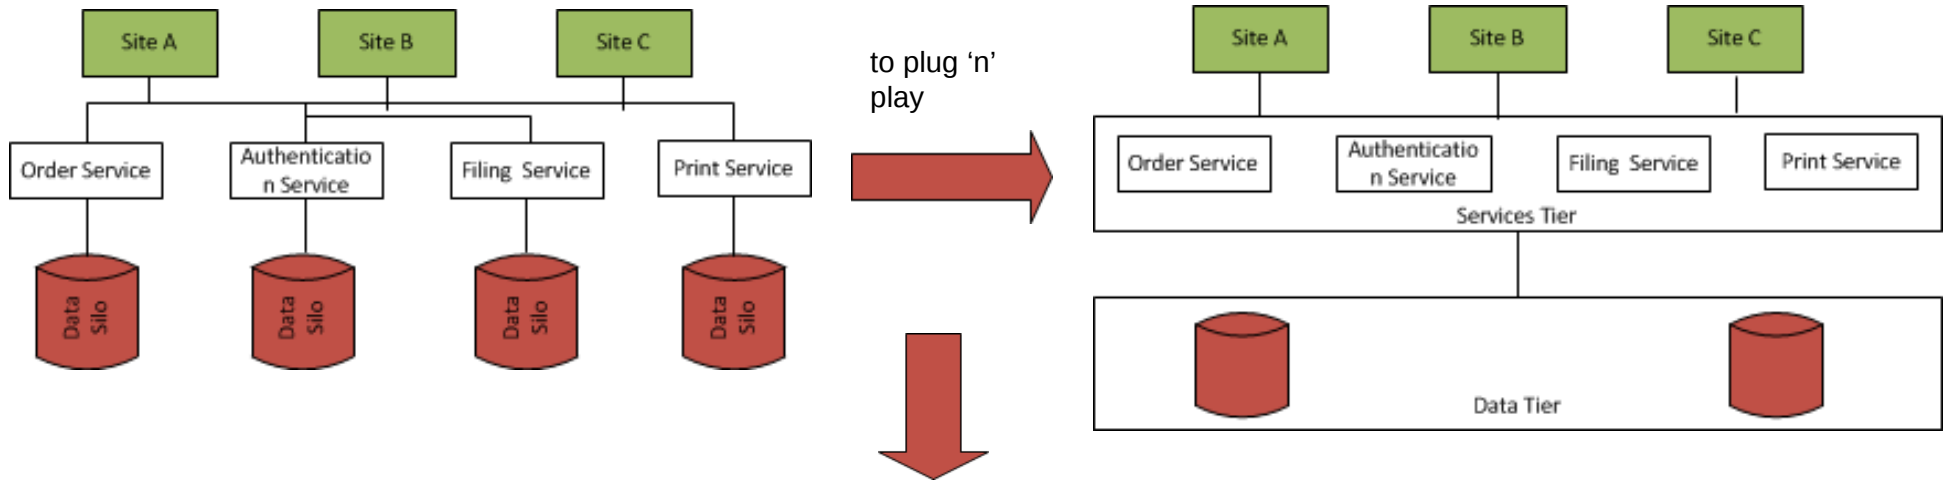


- Code duplication
- Cross cutting concern
- Data silos (information silos)
- Operational challenges - schema design, installs
- Added overhead to test and repeat test in production – slow prototyping



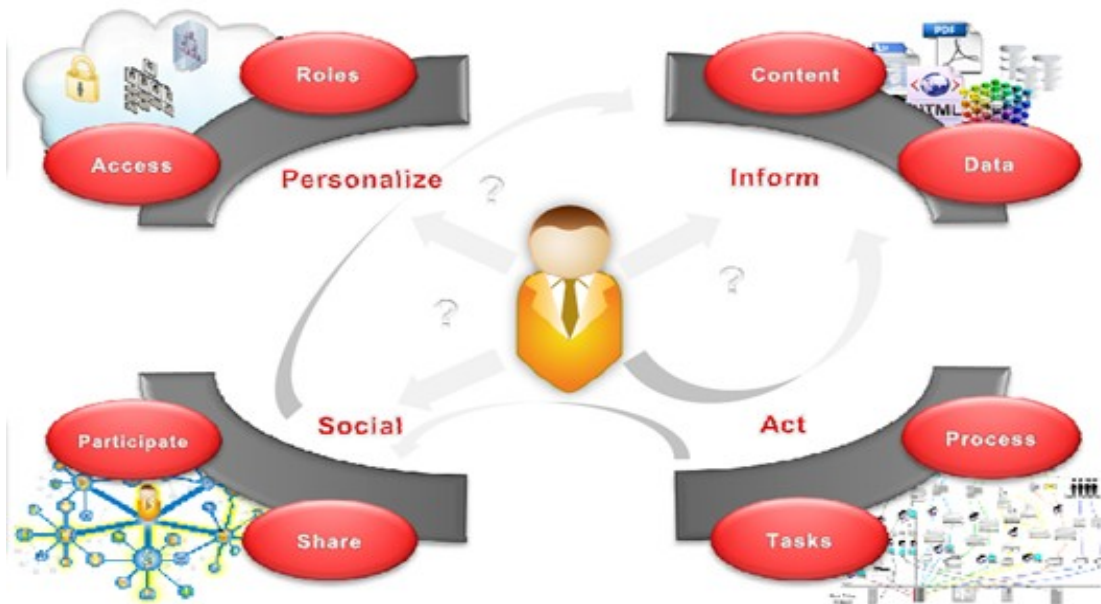
Intuit Proprietary & Confidential

Idea of a Platform



- Highly personalized experience
- Security

- Share data between products, between users




- Brings information together to avoid data silos

- Quick turnaround time
- Plug and play service framework
- Don't need IT and operations

intuit.

Data Platform/Tier

- Principles – Highly Available, Highly Scalable, Fast, Easy to operate software only solution for structured and unstructured data (blobs)
- Projection – Petabyte in 2-3 yrs
- Support – Critical application with 99.99%(5 nines) SLA
- But Wait ...No Stress




Intuit Proprietary & Confidential

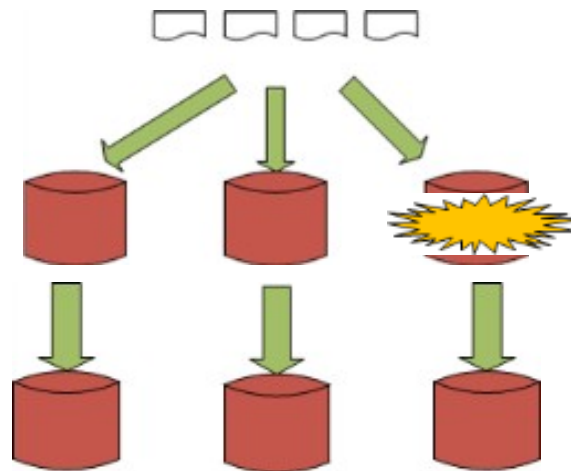
intuit.

Traditional RDBMS?

- Challenges with availability and scalability



- Sharding works well, but introduces new challenges as well



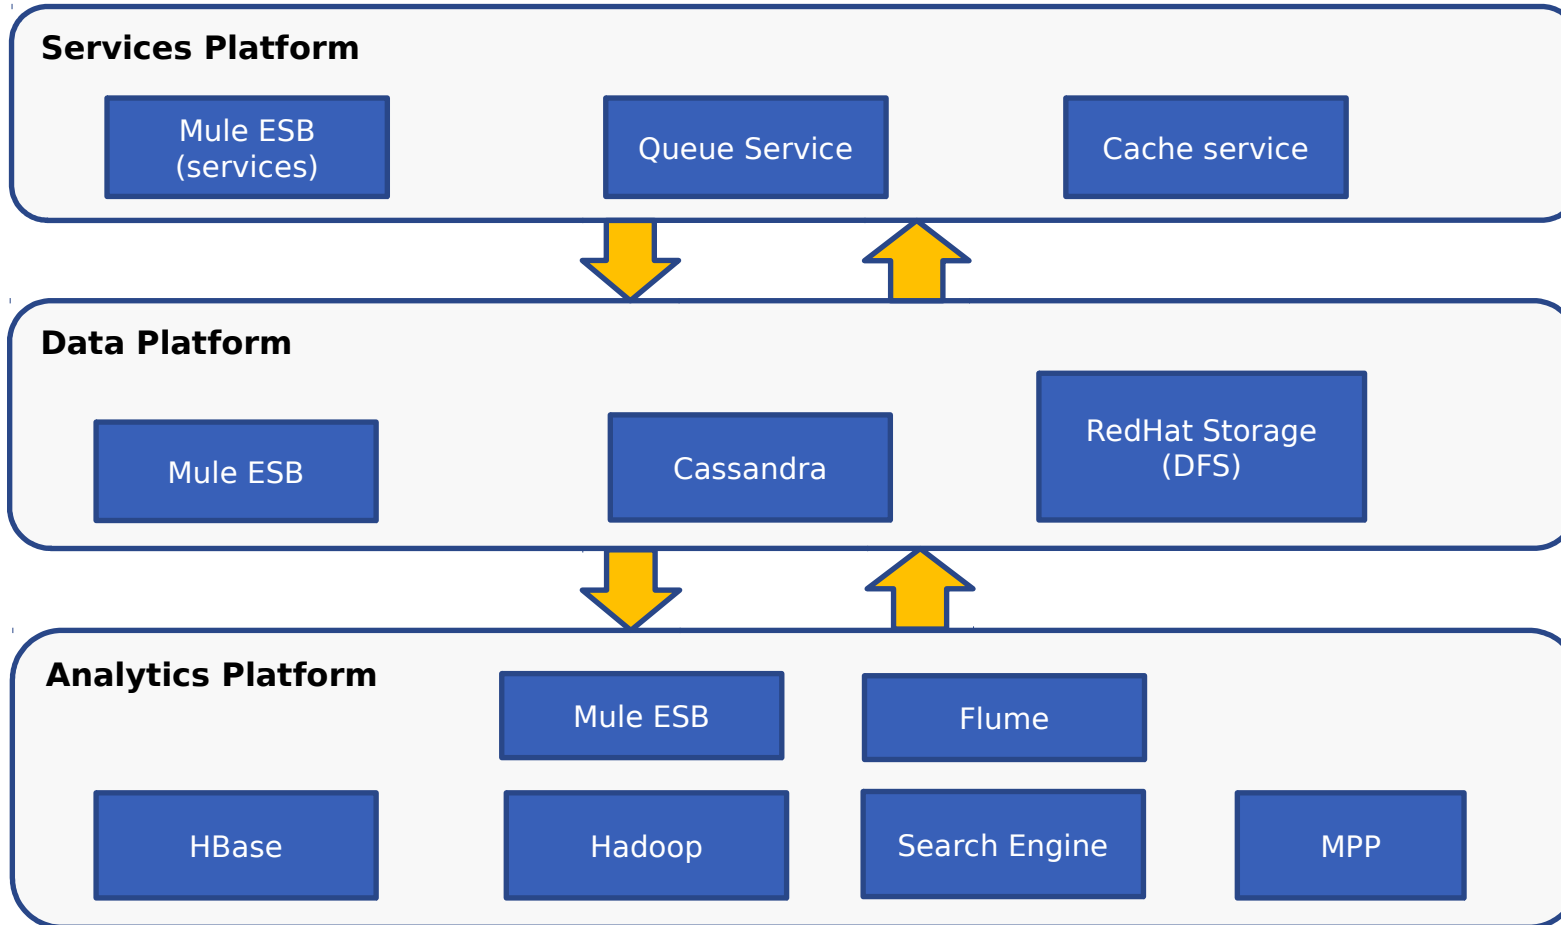
Intuit Proprietary & Confidential

intuit.

Why Red Hat Storage?

- Scalable
 - Easy to scale horizontally to petabytes of data
- Availability
 - Highly Available, can be designed for no SPOF
 - Easy to setup clusters and replication between DC
 - No metadata server required
- Operations
 - Easy to install
- Fast
 - Given the right hardware and RAID cache, Gluster provides low latency response times even for small files.

High Level CFP Stack



- **MuleSoft** ESB for business logic orchestration, with frameworks for additional authoring

Cassandra-powered schemaless database wrapped in entity and relationship logic.
RHS – a distributed file system for blob storage

Hadoop/Hbase/Solr/CEP-to meet batch processing and near real time analytics

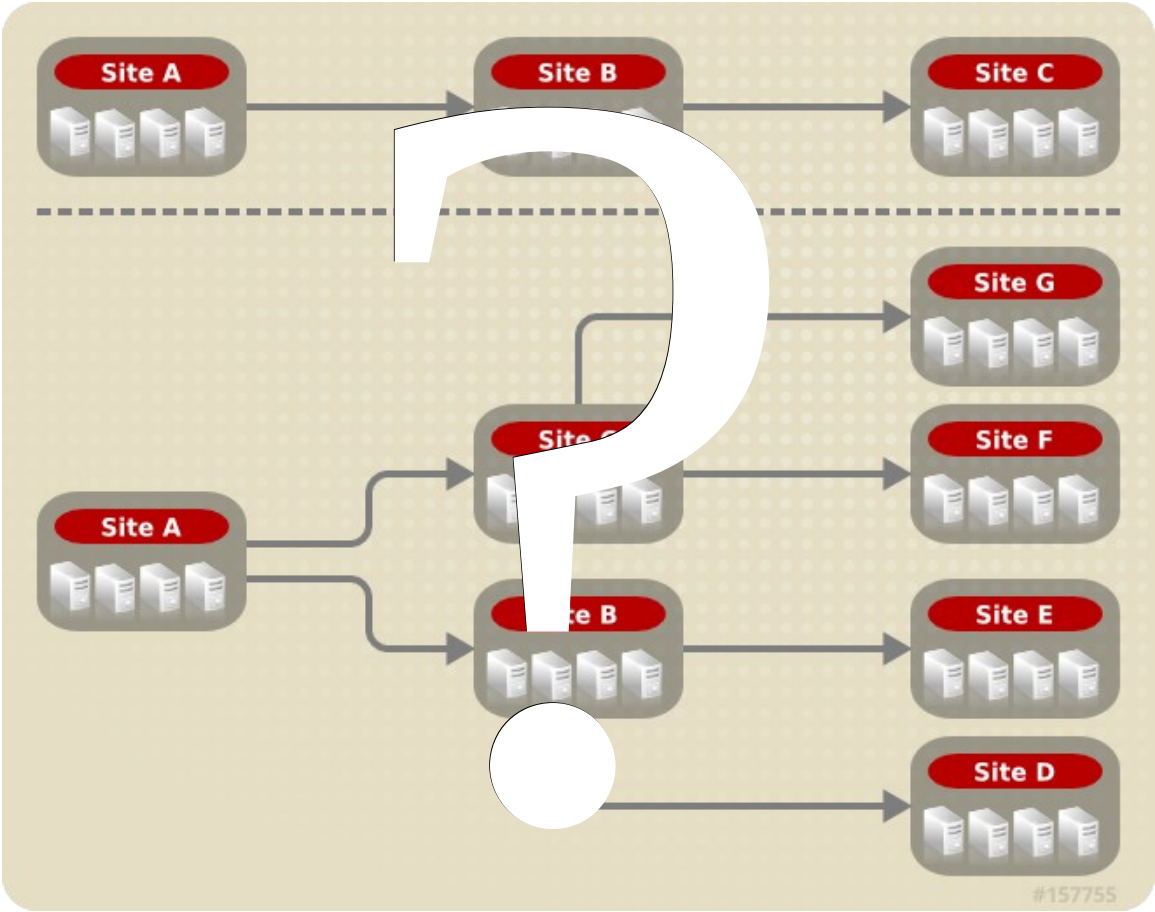
Intuit Proprietary & Confidential

intuit.






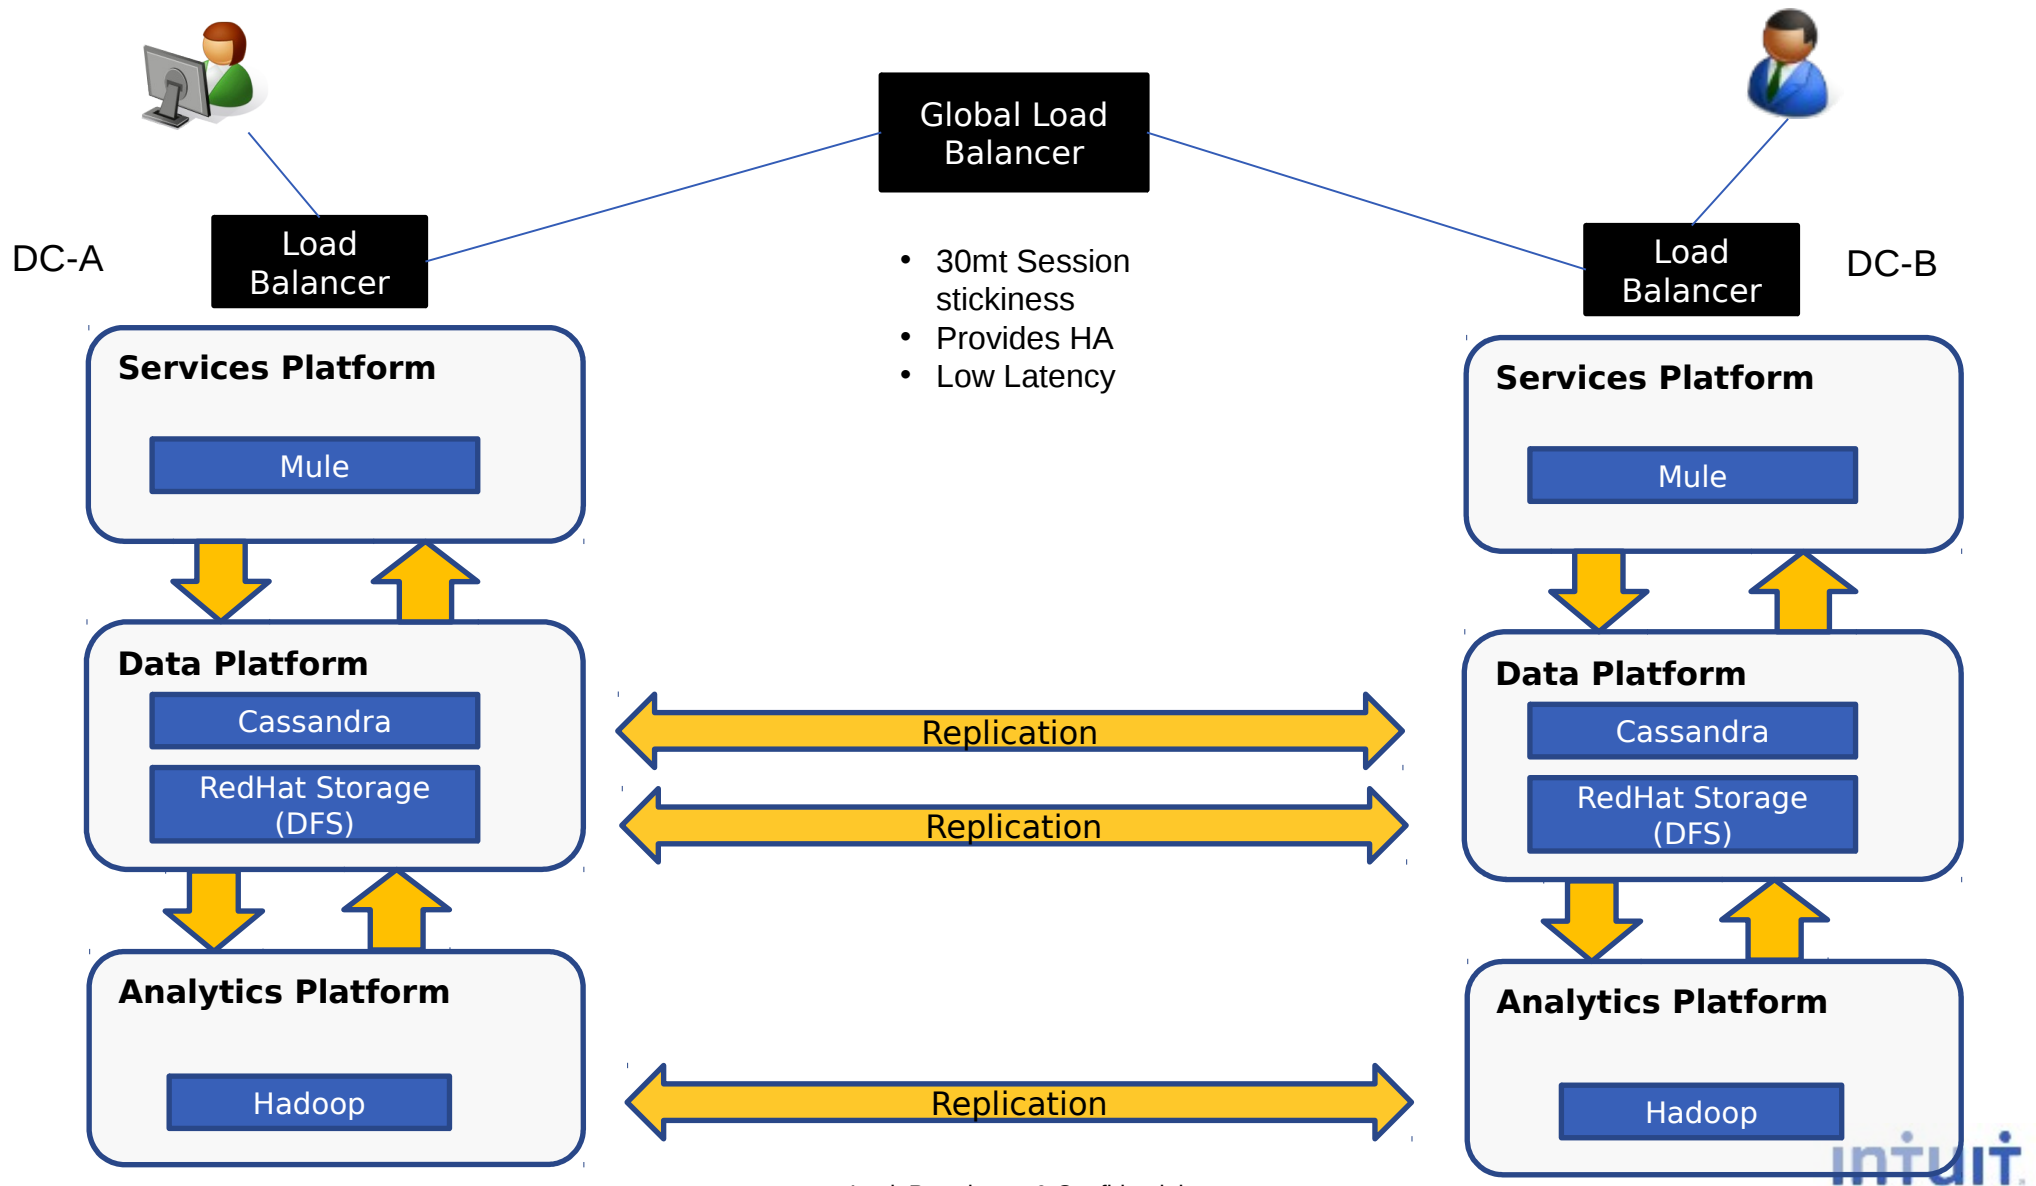
RED HAT® STORAGE




RED HAT® STORAGE



CFP Active/Active Multi-Data Center



Intuit Proprietary & Confidential



redhat.®

intuit.®



via lrargerich on flickr




via flickrfavorites on flickr

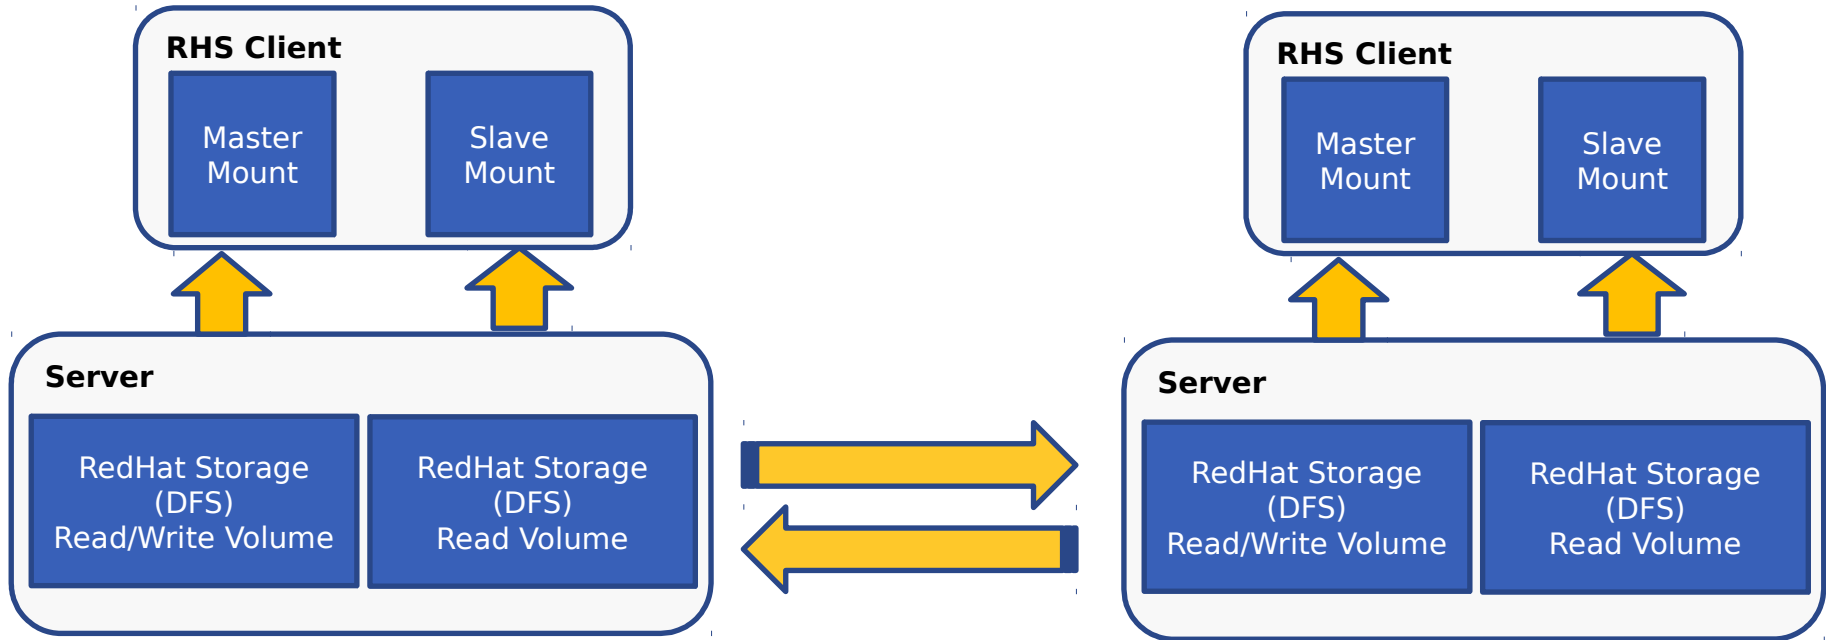
#redhat #rhsummit



RED HAT® STORAGE



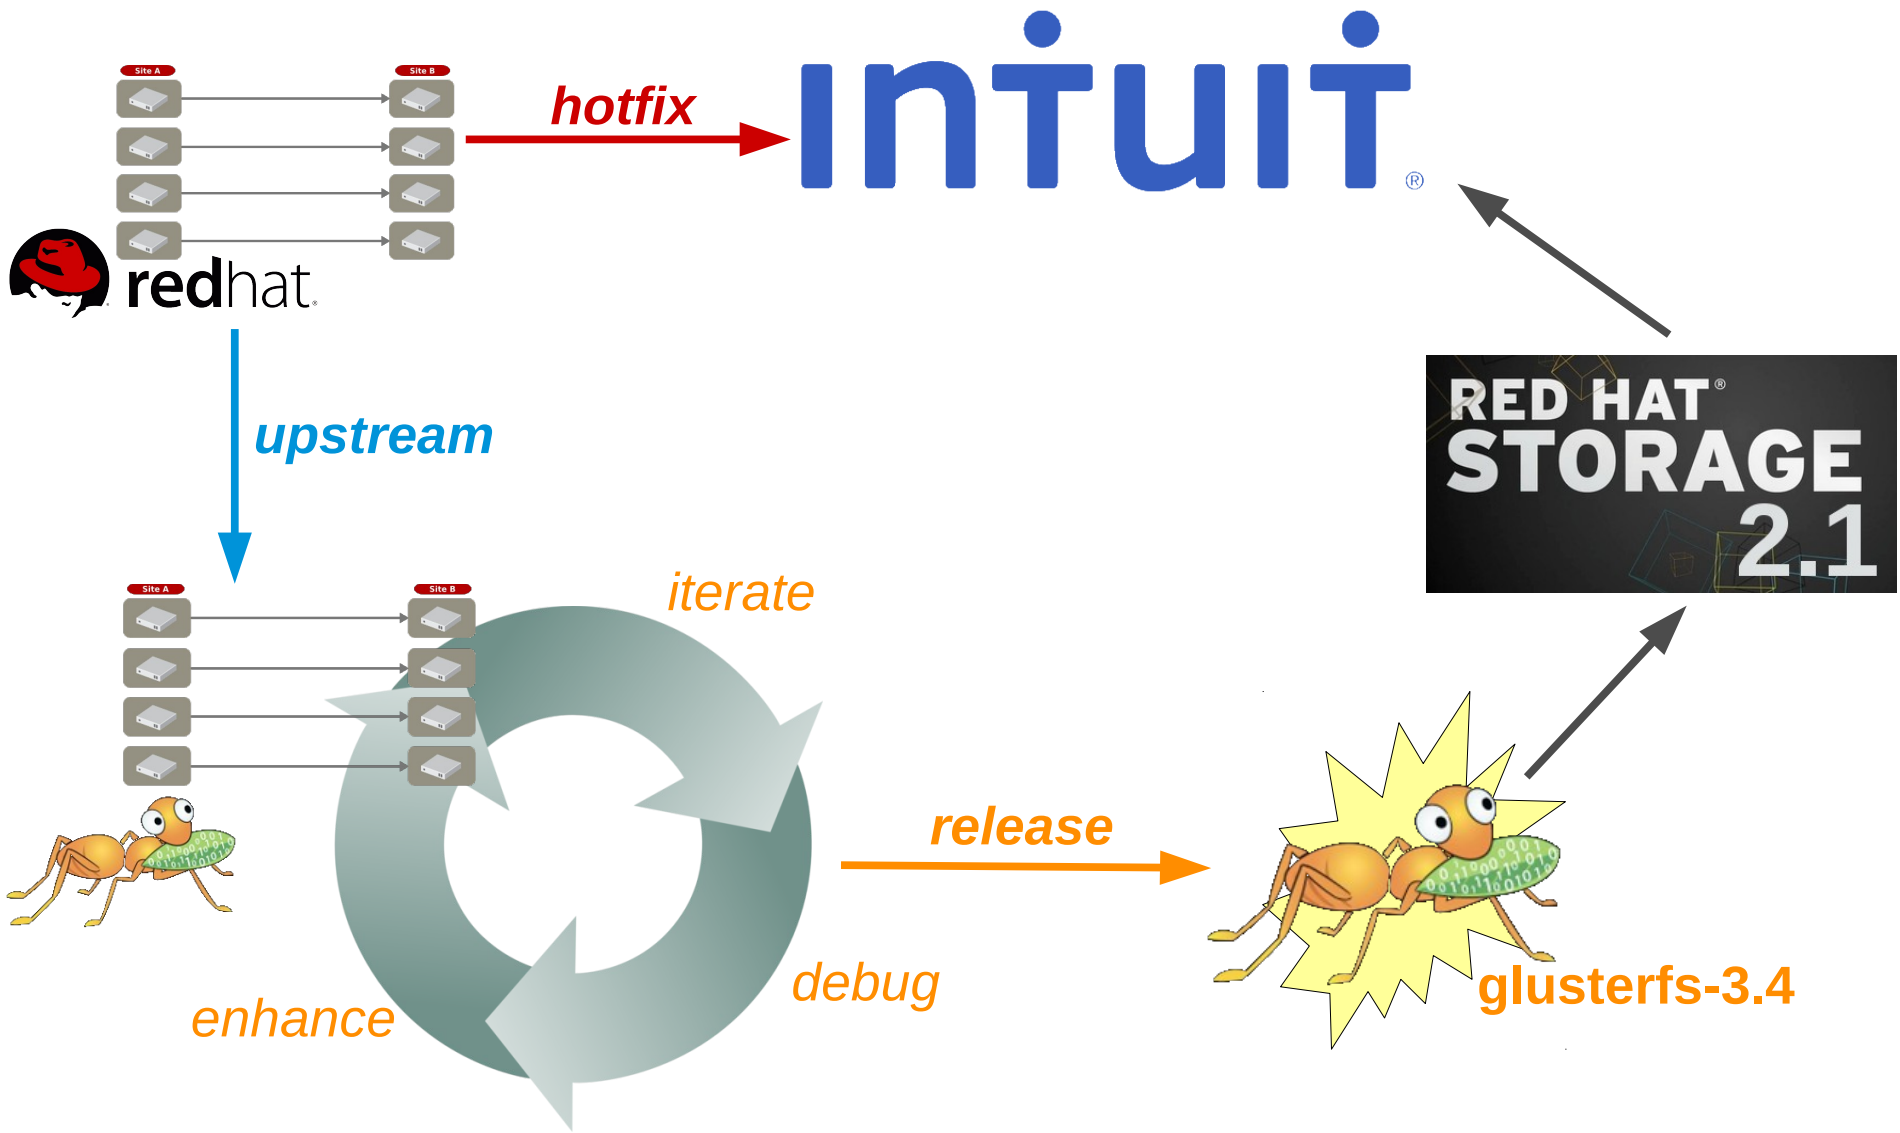
CFP Multi-Data Center Replication



- Use 2 volumes per data center. Each distribute replicate
- Set up 2 geo-replication processes per data center
- Use unique mount point names across entire CFP globally

Intuit Proprietary & Confidential

intuit.





via Laney Griner on flickr

#redhat #rhsummit



RED HAT
SUMMIT

Q&A

Slides at:
<http://people.redhat.com/dblack>

intuit®

Is Hiring!

Contact:
mohit_anchlia@intuit.com

RED HAT
SUMMIT

LEARN. NETWORK.
EXPERIENCE OPEN SOURCE.

June 11-14, 2013
Boston, MA

