

**SUMMIT**

**JBoss  
WORLD**

**PRESENTED BY RED HAT**

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

[www.theredhatsummit.com](http://www.theredhatsummit.com)

# Optimizing MRG Messaging Performance

Andy Goldstein  
Principal Architect  
Amentra, A Red Hat Company

James Kirkland  
Solutions Architect  
Red Hat

May 5, 2011

**SUMMIT**

JBoss  
WORLD

PRESENTED BY RED HAT



# Agenda

- MRG Messaging Overview
- What is Performance?
- General Tuning Guidelines
- Optimizing Throughput
- Optimizing Bandwidth
- Optimizing Latency
- Q&A

**SUMMIT**

JBoss  
WORLD

PRESENTED BY RED HAT

3



# Red Hat Enterprise MRG Messaging

- Open and unencumbered
- Fully featured
- Proven and reliable
- Highly performant & scalable for diverse workloads
- True interoperability with different application platforms



**SUMMIT**

JBoss  
WORLD

PRESENTED BY RED HAT



# Core MRG Messaging Features

- Open specification for wire level protocol - AMQP
- Supports multiple message distribution patterns
- Reliable messaging
- Transactions - local and distributed
- Interoperable clients

**SUMMIT**

JBoss  
WORLD

PRESENTED BY RED HAT

5



# Advanced MRG Messaging Features

- Federation
- Management framework (QMF)
- Active-Active Clustering
- SASL authentication
- SSL encryption
- ACLs

**SUMMIT**

JBoss  
WORLD

PRESENTED BY RED HAT

6



# MRG Fundamentals

- Broker
- Exchanges
  - e.g., Direct, Topic, Fanout, Header, XML
- Routing Key
- Binding Key
- Queue
- Connection
- Session
- Sender / receiver



**SUMMIT**

JBoss  
WORLD

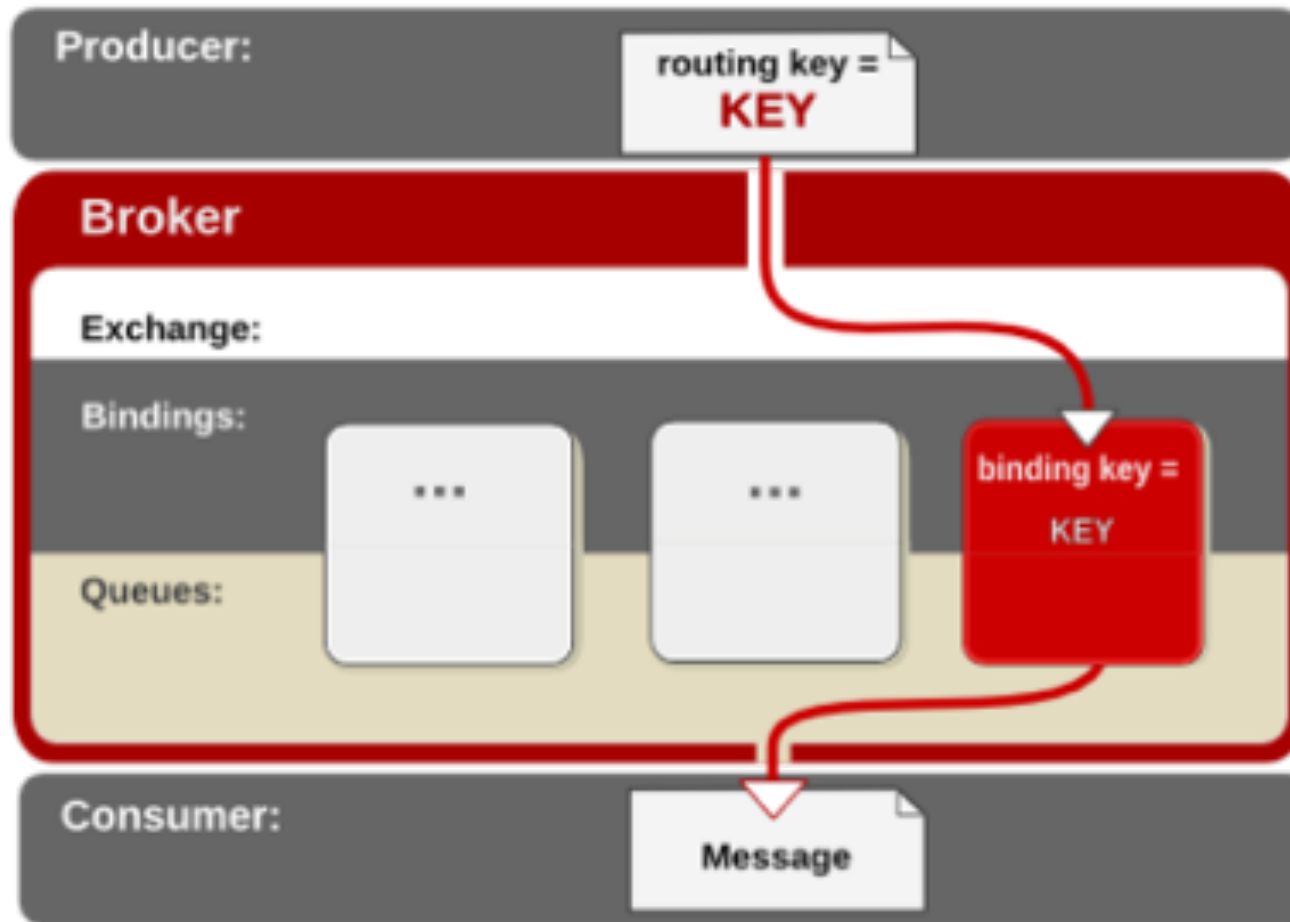
PRESENTED BY RED HAT

7



# Exchange examples

## Direct Exchange



Keys directly match

**SUMMIT**

JBoss  
WORLD

PRESENTED BY RED HAT

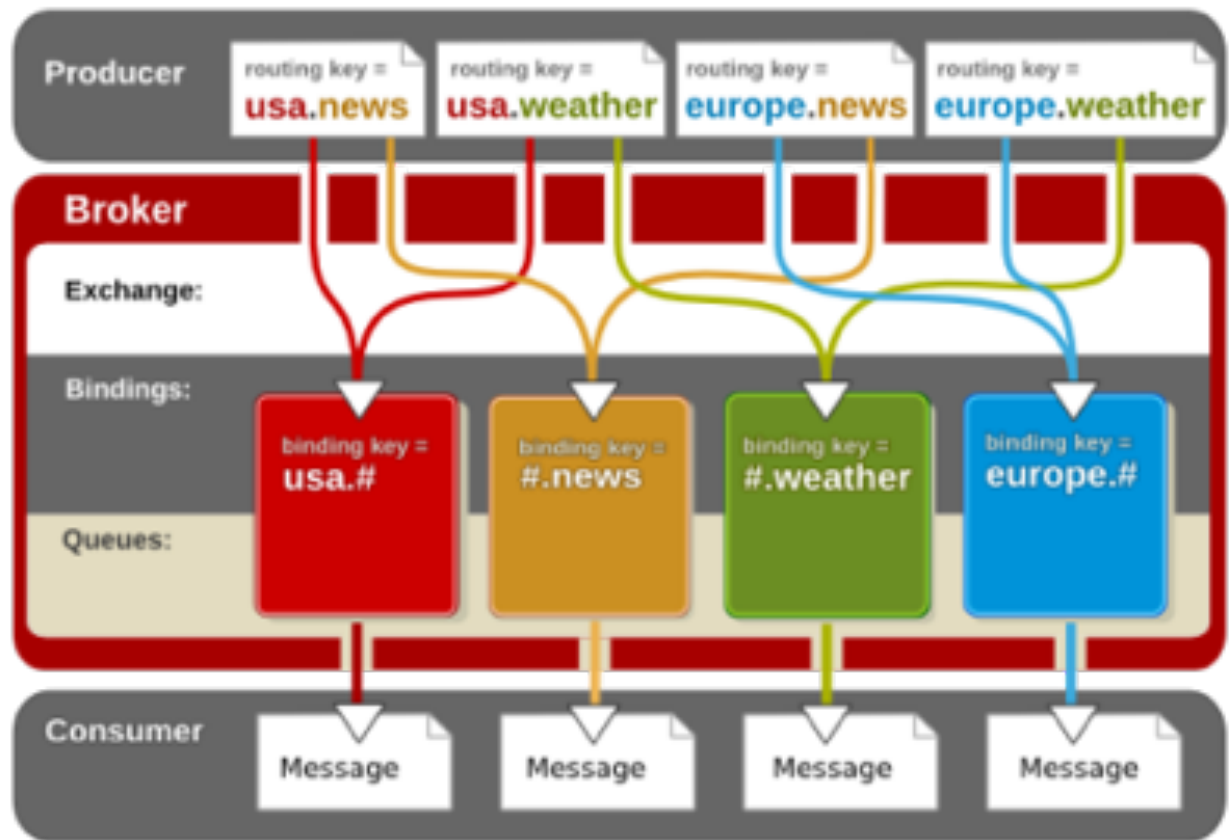




# Exchange examples

Keys matched on wildcards

## Topic Exchange



**SUMMIT**

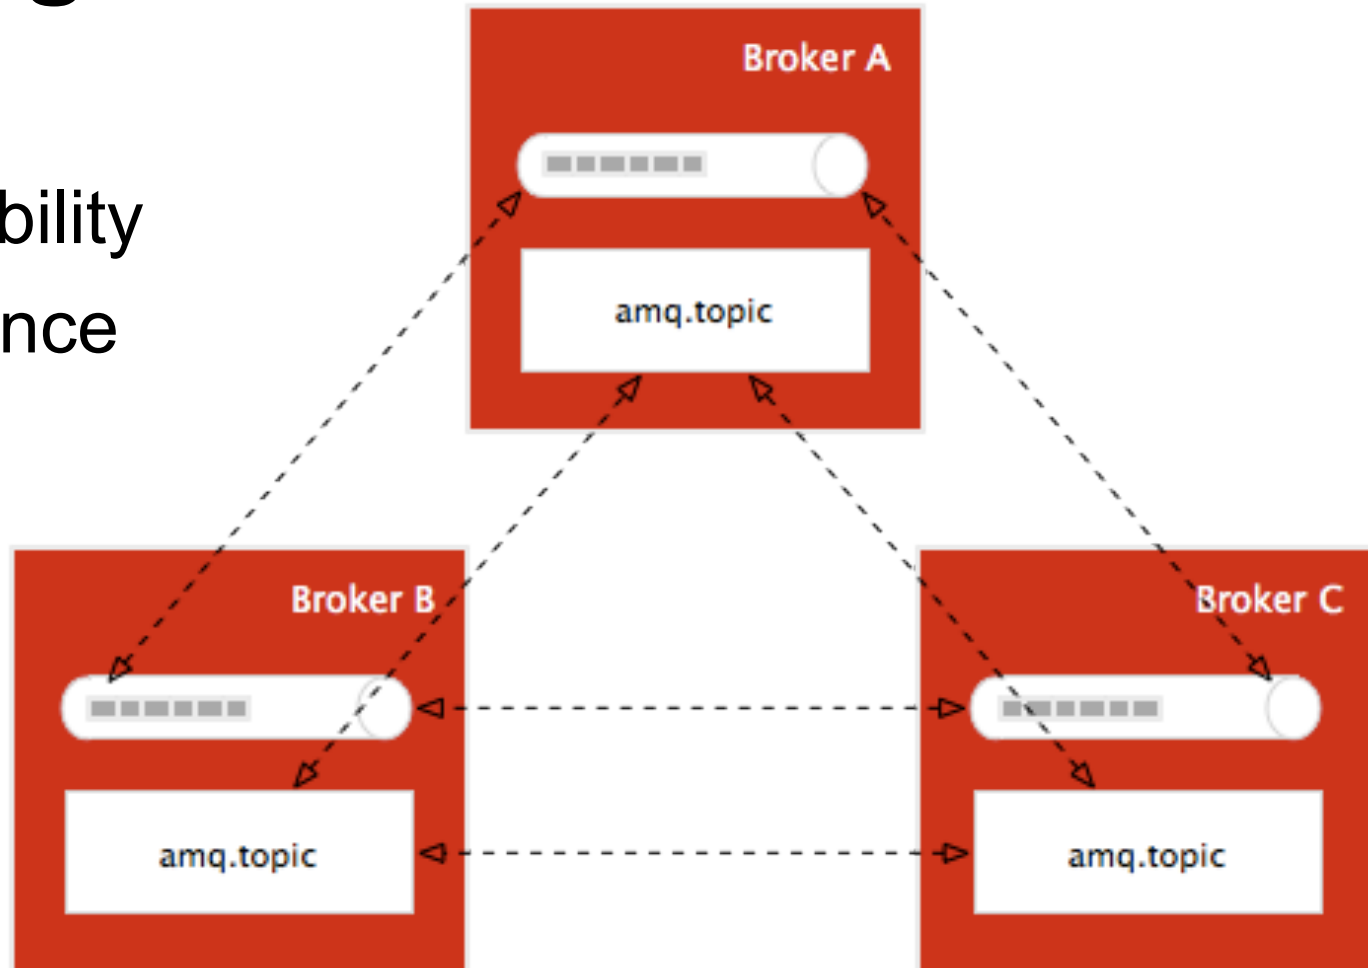
JBoss  
WORLD

PRESENTED BY RED HAT



# Clustering

- Clustering
  - High availability
  - Fault tolerance



**SUMMIT**

**JBoss  
WORLD**


PRESENTED BY RED HAT

10



# Federation

- Federation
  - Forward messages between brokers




**SUMMIT**

JBoss  
WORLD

PRESENTED BY RED HAT

11





# What is performance?

**SUMMIT**

**JBoss  
WORLD**

**PRESENTED BY RED HAT**

12



# ***What are the key measures of messaging performance?***



**Throughput = Messages / second**

**Bandwidth = MB / second**



**Latency = Transit time / message**

**SUMMIT**

**JBoss  
WORLD**

**PRESENTED BY RED HAT**

**13**



# General Tuning Guidelines

- Broker deployment architecture
- Exchange types
- Broker configuration
- Processor affinity
- Persistence and durability
  - Journal size
  - Queue depth



**SUMMIT**

JBoss  
WORLD

PRESENTED BY RED HAT

14



# Clustering and Federation

- Clustering

- Multicast token ring for active-active replication
- Performance implications
- Use RHCS to prevent split-brain
- Client failover timing

- Federation

- Only 1 connection between 2 broker peers
- Acknowledgements

**SUMMIT**

JBoss  
WORLD

PRESENTED BY RED HAT

15



# Optimizing throughput - Msg/sec

- Connection per thread
- Sender capacity
  - 1 - 4,300 msgs/sec
  - 30 - 25,000 msgs/sec
  - 50 - 30,000 msgs/sec
  - 100 - 34,000 msgs/sec
  - 1000 - 41,000 msgs/sec

**SUMMIT**

**JBoss  
WORLD**

PRESENTED BY RED HAT

16





# Optimizing throughput - Msg/sec

- Receiver capacity
  - 1 - 4,300 msgs/sec
  - 30 - 20,000 msgs/sec
  - 50 - 34,000 msgs/sec
  - 100 - 39,000 msgs/sec
  - 1000 - 42,000 msgs/sec

**SUMMIT**

JBoss  
WORLD

PRESENTED BY RED HAT

17



# Optimizing throughput - Msg/sec

- Message acknowledgement
  - Ack every message: 27,000 messages/sec
  - Ack in batches of 2: 36,000 messages/sec
  - Ack in batches of 10: 43,000 messages/sec
  - No acks: 45,000 messages/sec

**SUMMIT**

JBoss  
WORLD

PRESENTED BY RED HAT

18



# Optimizing bandwidth - MB/sec

- Minimum AMQP 0-10 frame size - 12 bytes

Fewer frames

+ More data per frame

Higher bandwidth utilization

- “Container” message option

**SUMMIT**

JBoss  
WORLD

PRESENTED BY RED HAT

19



# Optimizing bandwidth - MB/sec

- Utilize the fastest networking technology available
  - 10 Gigabit Ethernet
  - Infiniband
- Put messaging and cluster traffic on dedicated network

**SUMMIT**

JBoss  
WORLD


PRESENTED BY RED HAT

20




# Optimizing bandwidth - MB/sec

64-Byte Throughput



4096 Bytes Throughput



**SUMMIT**

**JBoss  
WORLD**

**PRESENTED BY RED HAT**

21



# Optimizing latency

- OS induced latency
- Realtime kernel
- Make the right I/O choice
- TCP nodelay
- Look for latency in your application

**SUMMIT**

**JBoss  
WORLD**

**PRESENTED BY RED HAT**

22



# Latency - Dealing with SMI

- Disable SMIs
  - Dynamic Power Savings Mode
  - CPU Utilization monitoring
  - P-state monitoring
  - ECC monitoring
- Benefits both RHEL & MRG operating environments

**SUMMIT**

JBoss  
WORLD


PRESENTED BY RED HAT

23




# Latency - Dealing with SMIs

Latency spikes with standard BIOS settings



Latency when SMIs disabled in BIOS



Source: <ftp://ftp.hp.com/pub/c-products/servers/linux/realtime/HPConWS-panel-Trieloff-Fisher-090914.pdf>

**SUMMIT**

**JBoss  
WORLD**

**PRESENTED BY RED HAT**

24





# I/O Technology Choices

- Network
  - InfiniBand
  - 10gigE
  - RDMA
- Disk
  - FC / FCoE
  - iSCSI
  - Infiniband



**SUMMIT**

**JBoss  
WORLD**

PRESENTED BY RED HAT

25



# What is RDMA?

- Remote Direct Memory Access
  - Originated with InfiniBand
  - Adopted in 10gigE as RoCE
- Direct memory to memory data copy
- Little CPU or OS overhead
- Latest iteration is QDR (40 Gbit/sec bandwidth and sub 3  $\mu$ s latency)



**SUMMIT**


**JBoss  
WORLD**

**PRESENTED BY RED HAT**

26



# Network Type & Latency



**SUMMIT**

**JBoss  
WORLD**

**PRESENTED BY RED HAT**

27



# Operating System Tuning

- NUMA
  - Maximize CPU cache hits
  - Localize memory
- Move IRQ handlers
- Tune networking parameters

**SUMMIT**

JBoss  
WORLD

PRESENTED BY RED HAT

28



# Q&A



**SUMMIT**

**JBoss  
WORLD**

**PRESENTED BY RED HAT**

29



**LIKE US ON FACEBOOK**

[www.facebook.com/redhatinc](http://www.facebook.com/redhatinc)

**FOLLOW US ON TWITTER**

[www.twitter.com/redhatsummit](http://www.twitter.com/redhatsummit)

**TWEET ABOUT IT**

#redhat

**READ THE BLOG**

[summitblog.redhat.com](http://summitblog.redhat.com)

**GIVE US FEEDBACK**

[www.redhat.com/summit/survey](http://www.redhat.com/summit/survey)

**SUMMIT**

**JBoss  
WORLD**

**PRESENTED BY RED HAT**

