

# RED HAT JBOSS DATA GRID 6.1 INTELLIGENT, DISTRIBUTED CACHING

DATASHEET

## RED HAT JBOSS DATA GRID IN ACTION



### FINANCIAL SERVICES TRADING

One of the world's largest financial trading companies required greater scalability, performance, and availability from its trading system.

Red Hat JBoss Data Grid enabled the organization to cost-effectively scale data tiers, without compromising its infrastructure. Additionally, this in-memory data grid offered extremely low latency with minimal to zero downtime.

Even better, Red Hat JBoss Data Grid was implemented with minimal disruption to existing applications, enabling IT and the entire organization to continue to focus on other critical business initiatives.

## THE APPLICATION DESIGN CHALLENGE

Today, businesses demand that their applications deliver higher performance, availability, reliability, flexibility, and scalability than ever before. But the influx of data is creating new obstacles that make it difficult for applications to perform. Data-related obstacles include:

**Unprecedented transaction volumes:** Business is increasingly performed via online applications, which increases web activity and forces applications to handle unprecedented, growing transaction volumes.

**High uptime expectations:** To keep customers loyal and engaged, applications must perform with consistency and deliver seamless customer service, even during peak activity times and unexpected traffic spikes.

**Open hybrid cloud environments:** On-premise or in a cloud? Software, platform, or data-as-a-service? Centralized or distributed? Open source or proprietary? IT infrastructures are more complex than ever before, and organizations need flexible applications that can be used in a variety of open hybrid cloud environments.

**Demand for accurate, real-time information:** Having real-time or near real-time access to accurate information often makes the difference between right and wrong decisions. In today's fast-paced environment, you can't afford to wait for business-critical information.

**Complex and rigid data-tier interaction:** Standing up a new application should be a streamlined and straightforward process. Wrangling with a complex data tier or being restricted by interface limitations is not acceptable for an organization racing to lead the market.

## PRODUCT OVERVIEW

As a leader in the tooling and runtime servers that power numerous high-scale web applications, Red Hat is aware of the challenges application developers face.

Developers need a solution that boosts application performance, meets tough requirements, and offers the flexibility that hybrid cloud environments require. Based upon Infinispan, the popular JBoss® Community project, Red Hat® JBoss Data Grid is appropriate for any type of application that has heavy compute needs. Here's how Red Hat JBoss Data Grid overcomes these challenges.

### FAST, SCALABLE PERFORMANCE

As the amount of data grows, reads and writes to traditional back-end data stores becomes a major performance bottleneck for web applications. By ensuring that the right data can be easily accessed in memory, Red Hat JBoss Data Grid acts as an intermediate layer between relational stores to meet data-retention requirements and promote extremely fast, scalable read-write performance.

### ADAPTABLE RESPONSE TIMES

For many applications, data demands ebb and flow based on different business requirements. Seasonal business, marketing and promotions, and external market changes all push the scaling requirements for a web application. By using Red Hat JBoss Data Grid, applications can elastically spin up and down distributed nodes to guarantee that required response times adapt accordingly.



facebook.com/redhatinc  
@redhatnews

linkedin.com/company/red-hat



## PUBLISHING AND MEDIA

A large, global publisher and distributor of educational books, technology, and media needed to achieve reliable performance and scale during peak seasonal demand and provide a seamless customer experience across their mobile and web applications.

This company used Red Hat JBoss Data Grid to increase application performance, scalability, agility, and resiliency. JBoss Data Grid also acts as a shared, NoSQL data store with data locality and provides long-term flexibility via flexible APIs. Finally, the company used Red Hat JBoss Data Grid to meet their critical business goals, including demanding deadlines, budget limitations, and achieving long-term business agility.

## FLEXIBILITY FOR OPEN HYBRID CLOUD ENVIRONMENTS

Most companies have a diverse IT environment. Applications and infrastructure are on-premise and in the cloud, legacy and contemporary. Unlike proprietary, closed stacks, Red Hat JBoss Data Grid presents a flexible opportunity to work with diverse technologies in any environment, application, platform, or database. Data grids can be used as a data abstraction layer, offering businesses greater independence to unlock themselves from proprietary stacks and achieve the flexibility they need for hybrid solutions.

## STREAMLINED DATABASE INTERACTIONS, FASTER TIME-TO-MARKET

Red Hat JBoss Data Grid decreases time-to-market by simplifying the interaction with traditional databases. As a drop-in, flexible solution with numerous application programming interfaces (APIs), Red Hat JBoss Data Grid boasts compatibility with any application and any platform via multiple wire protocols (i.e. REST, HotRod, memcached). This streamlined approach saves developers time because they don't have to struggle with interface limitations or fight for attention and resources from database teams.

## FEATURES AND BENEFITS

Red Hat JBoss Data Grid has broad industry appeal, offering intelligent, distributed caching for financial services trading, telecommunications, logistics, eCommerce, and call center tracking.

With Red Hat JBoss Data Grid, you get the benefits of scale and high performance without the costs of rewriting or replacing the data tier. Plus, it's available as a subscription, preventing vendor lock-in and saving the licensing costs associated with other data-caching alternatives.

FEATURE	BENEFITS
Cross-datacenter replication	<p>Many applications use geographically distributed datacenters with cross-data-center replication to achieve:</p> <ul style="list-style-type: none"> <li>• Fault tolerance, near caching, zero downtime, and load balancing.</li> <li>• Increased uptime and improved service-level agreements (SLAs).</li> <li>• "Follow the sun" capabilities with geographically dispersed datacenters.</li> <li>• Geographic data partitioning with failover for a better user experience.</li> </ul>
Zero-downtime rolling upgrades	<p>With zero-downtime rolling upgrades, you can upgrade your cluster from Red Hat JBoss Data Grid version 6.1 to the next version without any downtime. This results in increased, uninterrupted performance for remote users and applications.</p>
Deliver into hybrid cloud environments	<p>Red Hat JBoss Data Grid offers an open, flexible opportunity to work with diverse technologies in any environment—on-premise or cloud. The data grid can also act as a data abstraction layer in hybrid cloud solutions.</p>
Management tooling	<p>Red Hat JBoss Data Grid includes a subscription to the Red Hat JBoss Operations Network—a middleware and application management solution that provides a single point of control to deploy, manage, and monitor Red Hat JBoss Middleware applications and services.</p>
Multiple access protocols	<p>It is easy to access the data grid using REST, memcached, HotRod, or through a simple Java™ map API. The variety of access models means any sort of application—legacy or new, Java or native, in-memory or remote—can easily access data in the grid.</p>

FEATURE	BENEFITS
Elastic scaling	Adding and removing storage nodes is achieved simply, is non-disruptive, and scales linearly. Scaling a data grid is as simple as configuring and launching a new process. All the work is completed in the background to ensure data remains distributed and replicated.
Schema-less key value store	Red Hat JBoss Data Grid is an in-memory NoSQL database that offers a simple, flexible way to store different objects without a fixed data model.

## TECHNICAL SPECIFICATIONS

### CORE

- Topology configuration options of local, invalidation, replicated, and distributed modes
- Eviction and expiration processes
- Context Dependency Injection (CDI) integration
- Asynchronous API
- Asynchronous configuration options
- Per-invocation flags
- L1 caching
- Store entries as binary

### DISTRIBUTED

- A distributed system with a single-system image
- Elasticity to add or remove nodes or virtual nodes on-demand
- Load balancing to highly available partitions on multiple nodes via dynamic partitioning
- Non-blocking state transfer to ensure performance and continuous data grid operation when a new node is added or removed
- Data affinity via annotations, node affinity, and key affinity
- Node discovery and failure detection
- High-performance group membership protocol that supports both User Datagram Protocol / Internet Protocol (UDP/IP) and Transmission Control Protocol/Internet Protocol (TCP/IP)
- Multicast and unicast messaging
- Server hinting to ensure owners are not on the same machine, rack, or site

### CONCURRENCY

- Support for Extended Architecture (XA) transactions via Java Transaction API (JTA), transaction recovery, and JTA synchronization
- Locking configuration options: optimistic, pessimistic, and explicit locking as well as deadlock detection. Also, lock striping and single node locking.
- Isolation levels: read committed, repeatable read
- Invocation batching
- Operations with versions

### PROCESSING

- Map and reduce framework and programming model, including mapper, collector, reducer, and collator interfaces
- Distributed tasks to execute a task on some of all nodes in parallel

### INCLUDES DEVELOPER ENVIRONMENT RED HAT JBOSS DEVELOPER STUDIO

Red Hat JBoss Developer Studio is an Eclipse-based integrated developer environment (IDE) bundled with tools for developing, testing, and deploying rich web applications, enterprise applications, and service-oriented architecture (SOA) services. There is no need to buy multiple subscriptions to build and test on Red Hat JBoss products. Or pay production SLAs to get support. It's all included.

**HOW CAN I GET RED HAT JBOSS DATA GRID?**

JBoss Data Grid is generally available from Red Hat and its partners world-wide.

- Parallel and grid processing

**REMOTE**

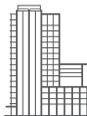
- HTTP/REST, memcached, HotRod or simple map-like API

**STORAGE**

- File system and Java Database Connectivity (JDBC) cache loader/store
- Read-through, write-through (synchronous) and write-behind (asynchronous)
- Activation and passivation via cache store

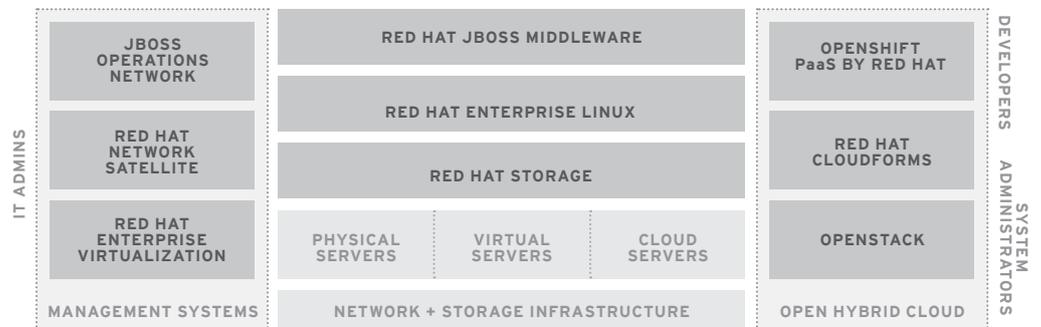
**MANAGEMENT**

- Includes Red Hat JBoss Operations Network
- Management and monitoring options via the Java Management (JMX) API
- Events, notifications, and listeners (synchronous and asynchronous)


**ABOUT RED HAT**

Red Hat is the world's leading provider of open source solutions, using a community-powered approach to provide reliable and high-performing cloud, virtualization, storage, Linux, and middleware technologies. Red Hat also offers award-winning support, training, and consulting services. Red Hat is an S&P company with more than 70 offices spanning the globe, empowering its customers' businesses.

**RED HAT PORTFOLIO** *Learn more at [redhat.com](http://redhat.com).*



facebook.com/redhatinc  
@redhatnews  
linkedin.com/company/red-hat

**NORTH AMERICA**  
1-888-REDHAT1

**EUROPE, MIDDLE EAST AND AFRICA**  
00800 7334 2835  
europe@redhat.com

**ASIA PACIFIC**  
+65 6490 4200  
apac@redhat.com

**LATIN AMERICA**  
+54 11 4329 7300  
latammktg@redhat.com