

RED HAT JBOSS DATA GRID 6.2

DATASHEET

THE APPLICATION DESIGN CHALLENGE

Today, businesses demand that their applications provide seamless customer experiences with fast response times. But the influx of data is creating new obstacles that make it difficult for applications to meet availability, reliability, flexibility, and scalability requirements. Data-related obstacles include:

- **Unprecedented transaction volumes.** Business is increasingly performed via online applications, which increases web activity and forces applications to handle unprecedented growth in transaction volumes.
- **High up-time 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 good and bad 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. Struggling with a complex data tier or being restricted by interface limitations is not acceptable for an organization racing to lead the market.



facebook.com/redhatinc

[@redhatnews](https://twitter.com/redhatnews)

linkedin.com/company/red-hat

WHAT IS A DATA GRID?

A data grid is an in-memory distributed database designed for fast access to large volumes of data and scalability. Data grids commonly provide a complementary layer to the relational database and the application.

Key characteristics of a data grid:

- In-memory, distributed caching
- Map/reduce
- Elastic scalability
- Advanced querying
- Data replication
- Processing for streaming data
- Transaction capabilities

PRODUCT OVERVIEW

As a leader in the tooling and runtime servers that power numerous high-scale 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 to innovate. Here's how Red Hat® JBoss® Data Grid overcomes these challenges.

HIGH AVAILABILITY TO ACCESS DATA

As the amount of data grows, reads and writes to traditional back-end data stores cause major performance bottlenecks for web applications. By ensuring that the right data is in-memory and highly available within and across datacenters, Red Hat JBoss Data Grid acts as a complementary layer to the application and its relational stores. This allows applications to meet data-retention requirements, promote extremely fast, scalable read-write performance, and meet up-time service-level agreements (SLAs).

MAINTAIN FAST RESPONSE TIMES WITH ELASTIC SCALE

Many applications struggle to adapt to surges in read and write requests during events such as seasonal business spikes, marketing and promotions, and external market changes. These applications need a cost-effective, non-disruptive way to achieve elastic scale and guarantee required response times. JBoss Data Grid makes it simple to add or remove nodes. And, all the work that ensures your data remains distributed and replicated is performed in the background.

FLEXIBILITY FOR OPEN HYBRID CLOUD ENVIRONMENTS

Most companies have diverse IT environments. Applications and infrastructure are on-premise and in the cloud, legacy and contemporary. Applications are often tightly coupled to proprietary, closed databases, preventing businesses from achieving the flexibility they need for hybrid solutions. JBoss Data Grid can act as a data abstraction layer, decoupling the application, cache, and database. As a result, businesses gain control over the life cycle, maintenance, and costs of each component independently.

MEET IMMEDIATE AND LONG TERM I.T. GOALS

Red Hat JBoss Data Grid is an adaptable solution that allows you to meet both immediate and long term IT goals. Avoid vendor lock-in and licensing costs with a straightforward open source subscription. Gain the flexibility to deploy your data, your way. Connect the data grid to any application and any platform via multiple protocols (for example, REST, Hot Rod, and memcached) and a Java™ API. With JBoss Data Grid, you have the flexibility you need to innovate faster, in a smarter way.

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 no 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.

FEATURES AND BENEFITS

Red Hat JBoss Data Grid has broad industry appeal, offering fast in-memory distributed data access and scalability for financial services trading, telecommunications, logistics, e-commerce, and call-center tracking.

Based on Infinispan, the popular JBoss Community project, Red Hat JBoss Data Grid is appropriate for any type of application that has heavy compute needs. 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.

FEATURE	BENEFIT
Querying	Easily search and find an object without needing to know the object’s assigned key. For example, you can search for red cars using a metadata match or for books on a specific topic using a full text search and relevance sorting.
Cross-datacenter replication	With Red Hat JBoss Data Grid, applications can replicate across data centers and achieve: <ul style="list-style-type: none"> • High availability to meet SLA requirements for data within and across data centers. • Load balancing and resource efficiencies via a “follow the sun” approach. • Geographic data partitioning with failover.
Distributed execution and map/reduce API	Quickly process large volumes of data and support long-running compute applications with the map/reduce API. Map/reduce uses the data grid’s parallel processing capabilities to move batch workloads away from the client application and divide it across the distributed system.
Zero-downtime rolling upgrades	Starting with Red Hat JBoss Data Grid 6.1, you can upgrade your cluster from one version to the next version without any downtime. This results in increased, uninterrupted performance for remote users and applications.
Management tooling	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.
Multiple access protocols with data compatibility	Multiple access models with data compatibility means any sort of application—legacy or new, Java or native, in-memory or remote—can easily access and share data in the grid. Applications can access the data grid remotely using REST, memcached, or Hot Rod (Java or C++), or locally through a simple Java™ API. And with compatibility mode, data can be shared (read and written) between applications using different access models.
Key value store	Red Hat JBoss Data Grid is an in-memory NoSQL database that offers a simple, flexible way to store a variety of data without the constraints of a fixed data model.



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. The company is now able to meet critical business goals, including demanding deadlines, budget limitations, and long-term business agility.

MANAGEMENT AND MONITORING TOOLING WITH RED HAT JBOSS OPERATIONS NETWORK

Included with a subscription to Red Hat JBoss Data Grid, Red Hat JBoss Operations Network provides a single point of control to deploy, manage, and monitor Red Hat JBoss Middleware applications. The built-in management and monitoring tooling allows you to effectively administer all your Red Hat JBoss Middleware application environments, helping you improve operational efficiency, reduce costs, and ensure a positive experience for your users.

TECHNICAL HIGHLIGHTS

CORE

- Topology configuration options of local, invalidation, replicated, and distributed modes
- Simple map-like Java API for data access in Library mode
- Eviction and expiration processes
- Context Dependency Injection (CDI) integration
- Asynchronous API
- Asynchronous configuration options
- Per-invocation flags
- L1 caching
- Store entries as binary
- Events, notifications, and listeners (synchronous and asynchronous)
- Data interoperability between protocols, different languages of the same protocol, and deployment modes.

PROCESSING AND DISTRIBUTION

- Map and reduce framework and programming model, including mapper, collector, reducer, and collator interfaces
- Querying to find data without needing to know its assigned key. (Fully supported in library mode and remote querying over Hot Rod is offered in technology preview)
- Distributed tasks to execute a task on some of all nodes in parallel
- Parallel and grid processing
- Consistent hashing and topology aware
- 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

**RED HAT JBOSS
DEVELOPER STUDIO**

Included in JBoss Data Grid is Red Hat JBoss Developer Studio, 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.

CONCURRENCY

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

REMOTE CLIENTS

HTTP/REST, memcached, and HotRod (Java or C++)

SECURITY

- TLS (SSL) for server authentication and encryption over Hot Rod; certificate-based client authentication also available

STORAGE

- SingleFileCacheStore for the fastest read and write performance; an index of keys stored in-memory
- LevelDB cache store for fast read and writes of large amounts of data persisted locally at each node; the index of keys not necessarily stored in-memory
- Java Database Connectivity (JDBC) cache loader/store
- Read-through, write-through (synchronous) and write-behind (asynchronous)
- Activation and passivation via cache store

MANAGEMENT

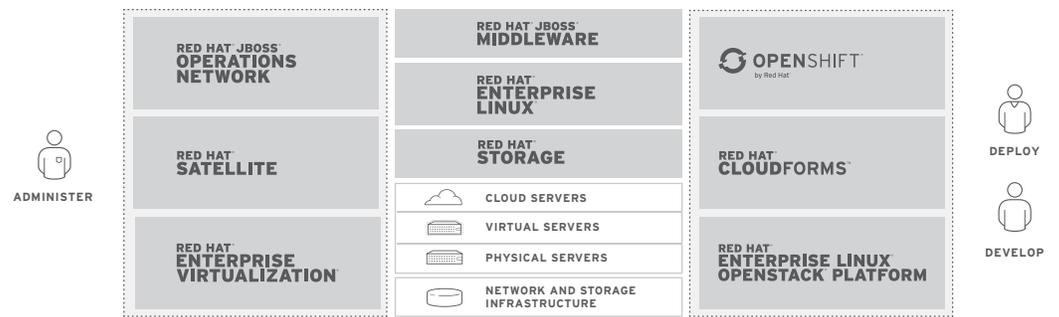
- Includes Red Hat JBoss Operations Network
- Management and monitoring options via the Java Management (JMX) API
- Command line interface



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



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
info-latam@redhat.com