

# FAST, SCALABLE, HIGHLY AVAILABLE APPLICATIONS

Improve performance with Red Hat JBoss Data Grid

TECHNOLOGY OVERVIEW

## BENEFITS

- Fast, low latency data processing through the use of memory (RAM) and distributed parallel execution
- Linear scalability through data partitioning and distribution across cluster nodes
- High availability through data replication across cluster nodes
- Disaster recovery and fault tolerance through cross data center geo replication
- Development flexibility and productivity with highly versatile and functionally rich NoSQL data store
- Comprehensive data security with encryption and role-based access
- Offload shared data services, enable real-time data compute and in-memory analytics and event processing

## INTRODUCTION

For a new business initiative, even a few seconds of delay can mean the difference between success and failure, as positive user experiences become increasingly dependent on application performance and quality. Data bottlenecks are more common as organizations quickly process larger volumes and greater varieties of data to meet customer expectations. Red Hat® JBoss® Data Grid is an in-memory data grid and NoSQL data store solution that helps applications access, process, and analyze data at in-memory speed to deliver superior user experiences.

## DATA GROWTH INCREASES I.T. COMPLEXITY

With the emergence of technologies such as cloud, big data, Internet of Things (IoT), and mobile, businesses need their applications to deliver higher performance, availability, reliability, flexibility, and scalability than ever before. But massive data growth is creating new obstacles that make it difficult for applications to meet these demands.

Scaling the data tier creates both technical and economic challenges for organizations. Scaling up requires additional hardware and database software licenses, while scaling out requires complex data partitioning or clustering technologies. With the implementation of cloud, Platform-as-a-Service (PaaS), and container-based infrastructures, these challenges become even more complicated. Whether data is hosted on-premise or in the cloud, in a centralized or distributed architecture, using open source or proprietary solutions, IT infrastructures are more complex than ever before. Organizations need flexible applications that can be used in a variety of open hybrid cloud environments.

## A SCALABLE, FLEXIBLE SOLUTION FOR APPLICATION DATA

To meet the challenges of IT complexity and data growth, data grids provide flexibility and elasticity to help organizations achieve the full benefits of Platform-as-a-Service and microservices architectures, while also helping applications run effectively in the cloud.

In-memory data grids, such as Red Hat JBoss Data Grid, give applications a scalable in-memory repository for rapidly changing application data. This solution eliminates disk bottlenecks and minimizes the use of cloud-based persistent storage. In addition, in-memory data grids enable transparent sharing of application data across a pool of instances to simplify design and reduce development time. This distributed data management system for application data:

- Uses RAM to store information for rapid, low latency response and very high throughput.
- Keeps copies of information synced across multiple servers for continuous availability, information reliability, and linear scalability.

Based on Infinispan, a JBoss community project, Red Hat JBoss Data Grid helps applications with heavy compute needs gain the benefits of scalability and high performance without the costs of rewriting or replacing the data tier.



[facebook.com/redhatinc](https://facebook.com/redhatinc)

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

[linkedin.com/company/red-hat](https://linkedin.com/company/red-hat)

## FEATURES

- On-demand capacity
- Linear scalability
- Fault-tolerant partitioned data caching and processing
- Grid-wide listeners and event-driven computation
- Rich querying, indexing, and Lucene search
- Real-time continuous query
- Distributed execution framework
- Read-through, write-through, and write-behind persistence
- Automatic, dynamic, and transparent management of partitioned and replicated data
- Cross-datacenter replication
- Rolling upgrades
- Configurable ACID transaction management
- Full range of security features
- Familiar JBoss Developer Studio environment

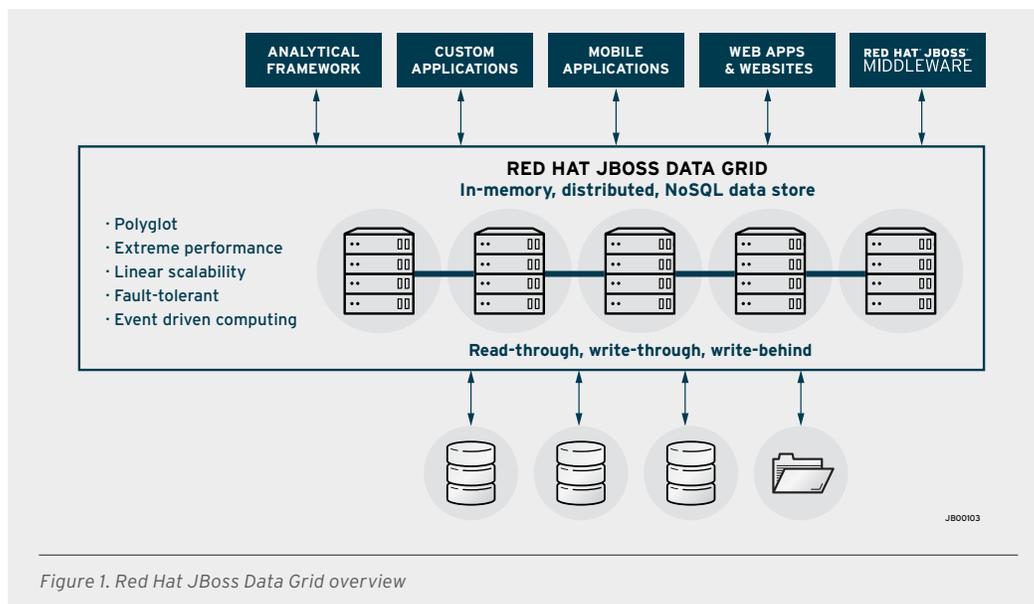


Figure 1. Red Hat JBoss Data Grid overview

With Red Hat JBoss Data Grid, organizations can improve application performance and scalability for faster decision making and greater productivity, resulting in a better customer experience.

## FEATURES AND BENEFITS

To support modern data management requirements with rapid data processing, elastic scalability, and high availability, Red Hat JBoss Data Grid offers:

- **NoSQL data store.** Provides simple, flexible storage for a variety of data without the constraints of a fixed data model. Red Hat JBoss Data Grid can be configured to fully participate in transactions.
- **Apache Spark and Hadoop integration.** Offers full support as an in-memory data store for Apache Spark and Hadoop, with support for Spark resilient distributed datasets (RDDs) and Discretized Streams (Dstreams), as well as Hadoop I/O format.
- **Rich querying.** Provides easy search for objects using values and ranges, without the need for key-based lookups or an object's exact location. Continuous queries provide the latest results in real time, without polling.
- **Polyglot client and access protocol support.** Offers read/write capabilities that let applications written in multiple programming languages easily access and share data. Applications can access the data grid remotely, using REST, Memcached, or Hot Rod – for Java™, C++, and .NET – or locally, through a Java application programming interface (API). Support for Java applications includes JSR107, CDI, and Spring Cache APIs, while all other application languages are supported using popular REST and Memcached protocols. In addition, Node.js client application support is available as a technology preview.

- **Distributed parallel execution.** Quickly process large volumes of data and support long-running compute applications. Simplified Map-Reduce parallel operations, based on the Java 8 Stream API, let developers process data declaratively and take advantage of multicore architecture. Developers can also complete parallel processing for multiple data operations on each Red Hat JBoss Data Grid cluster node and collect the resulting data without writing specific code.
- **Event-driven processing.** Enables real-time response – such as distributed parallel execution for processing large volumes of data – to data change events throughout the data grid. In addition, JBoss Data Grid now supports execution of stored tasks and scripts execution, letting remote clients invoke named tasks or scripts on the server similar to execution of stored procedures or triggers on a database. This capability brings data closer to compute logic – for example, co-located in-memory – for superior performance.
- **Flexible persistence.** Increase the lifespan of information in the memory for improved durability through support for both shared nothing and shared database – RDBMS or NoSQL – architectures. A combination of eviction and passivation ensures that only frequently required information is stored in-memory, while other data is stored in external storage.
- **Comprehensive security.** Meet strict requirements with secured communications between clients and servers and between server nodes within a secure cluster. Authentication, role-based authorization, and access control are integrated with existing security and identity structures to give only trusted users, services, and applications access to the data grid.
- **Cross-datacenter replication.** Replicate applications across datacenters and achieve high availability to meet service-level agreement (SLA) requirements for data within and across datacenters.
- **Rolling upgrades.** Upgrade your cluster without downtime for continuous, uninterrupted operations for remote users and applications.
- **Cloud-ready deployment.** Decouple applications, caches, and databases for independent control of the life cycle, maintenance, and cost of each component using Red Hat JBoss Data Grid as a data abstraction layer. Red Hat JBoss Data Grid can be deployed in on-premise, cloud, or hybrid environments to support both legacy and modern applications, hosted on-premise and in the cloud. Red Hat JBoss Data Grid for xPaaS provides in-memory speed and elastic data management for cloud applications running on OpenShift by Red Hat.

## ENTERPRISE USE CASES

Red Hat JBoss Data Grid provides value as a standard architectural component in application infrastructures for a variety of real-world scenarios and use cases.

### DATA CACHING AND TRANSIENT DATA STORAGE

In the most common data grid use cases, data caching and transient data storage, data grids such as Red Hat JBoss Data Grid are deployed as a rapid in-memory data store for the most frequently accessed data in applications. As a variation of data caching, data grids are often used to store transient data – for example, web sessions and shopping cart data – in e-commerce applications. As a result, these applications gain improved performance and scalability. In addition, these applications access database management systems (DBMS) and transactional backend systems less, resulting in reduced operating costs for these systems.

**CERTIFIED  
INTEGRATION WITH:**

- Red Hat JBoss Enterprise Application Platform
- Red Hat JBoss Fuse
- Red Hat JBoss BRMS
- Red Hat JBoss Data Virtualization
- Red Hat JBoss Web Server
- Spring Framework

Learn more at

[redhat.com/en/technologies/  
jboss-middleware/data-grid](https://redhat.com/en/technologies/jboss-middleware/data-grid)

**PRIMARY DATA STORE**

Red Hat JBoss Data Grid is an in-memory key-value data store, similar to a NoSQL database, and can be used by applications as their primary data store for rapid access to in-memory data, although data may also be persisted for recovery, backup, and archiving. In addition, applications can perform parallel distributed workload execution, run rich queries, manage transactions, scale as needed, and recover from network or system faults. With support for the Java 8 Stream API, Red Hat JBoss Data Grid simplifies the development of high-performance, data-intensive applications. It performs data processing operations in parallel while abstracting low-level multithreading logic to let developers concentrate on the data and related operations.

**LOW LATENCY COMPUTE GRID**

Data grids bring data physically closer to data processing to reduce latency and increase application performance. Red Hat JBoss Data Grid enables scale-out architecture that deploys application logic next to in-memory data in each node, rather than sending large sets of data to the compute nodes via wire. Network traffic is significantly reduced, resulting in dramatically increased application performance. Red Hat JBoss Data Grid also supports event-driven computing by executing application logic as data changes occur in the cluster – a key capability for real-time compute and analytics such as fraud detection and risk management applications.

**BIG DATA AND INTERNET OF THINGS**

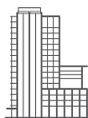
Data grids are well-suited to handle the three Vs of big data: velocity, variability, and volume. To support the velocity needs of big data, data grids support hundreds of thousands of in-memory data updates per second. Data grids support big data's variability in a manner similar to NoSQL data stores. Finally, data grids can be clustered and scaled to support large volumes of data.

IoT devices generate massive volumes of data, often at frequent intervals. JBoss Data Grid provides storage of tens of terabytes of data, with faster response times and near-instant analytics. As a result, IoT data can be processed at nearly the same speed that it is generated.

## STAY COMPETITIVE WITH MODERN DATA MANAGEMENT

Data management is a critical issue for nearly every business. To stay competitive, organizations need to take risks, track success, and correct issues quickly – while also supporting continual growth and evolving to take advantage of mobile computing, big data, the Internet of Things, cloud computing, and other new and emerging technologies. Traditional methods of data persistence and management can increase costs and risk while inhibiting business growth. In-memory data grids use cost-effective technologies to provide data management without disrupting business operations. With Red Hat JBoss Data Grid, organizations can avoid the limitations of legacy technology and focus on developing and applying application logic to achieve success.

- Learn more about Red Hat JBoss Data Grid at [redhat.com/en/technologies/jboss-middleware/data-grid](http://redhat.com/en/technologies/jboss-middleware/data-grid)
- Learn more about JBoss Data Grid at [jboss.org/products/datagrid/overview](http://jboss.org/products/datagrid/overview)
- Access JBoss Data Grid resources at [jboss.org/products/datagrid/resources](http://jboss.org/products/datagrid/resources)
- Get involved in the JBoss Data Grid community at [jboss.org/products/datagrid/community](http://jboss.org/products/datagrid/community)



### ABOUT RED HAT

Red Hat is the world's leading provider of open source software solutions, using a community-powered approach to provide reliable and high-performing cloud, Linux, middleware, storage, and virtualization technologies. Red Hat also offers award-winning support, training, and consulting services. As a connective hub in a global network of enterprises, partners, and open source communities, Red Hat helps create relevant, innovative technologies that liberate resources for growth and prepare customers for the future of IT.



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