RED HAT® JBOSS®
MIDDLEWARE

redhat.

RED HAT JBOSS DATA GRID

TECHNOLOGY OVERVIEW

SCHOLASTIC BUILDS STORIA, A DIGITAL READING MANAGEMENT SOLUTION FOR E-BOOKS

Scholastic is a large, global publisher and distributor of educational books, technology, and media. They wanted to give users a seamless experience between web and mobile applications, which was particularly challenging to accomplish during back-to-school traffic spikes.

With JBoss Data Grid, Scholastic was able to achieve a resilient, reliable scalable environment with zero downtime, and improved response time from over 10 seconds to under two seconds. With Red Hat and JBoss Data Grid, Scholastic both met immediate budget limitations and deadlines and became able to respond to longer-term business and technical demands.

INTRODUCTION

Red Hat® JBoss® Data Grid, an in-memory distributed database, provides fast access to data and elastic scalability for large data volumes and allows developers to meet tough requirements related to data retention, response time, and uptime. JBoss Data Grid offers an alternative or complementary layer to the existing data tier and compatibility with applications written in any language, using any framework and platform via multiple APIs such as memcached, Hot Rod, and REST. With management tooling and other advanced functionality included, JBoss Data Grid is an ideal technology for supercharging today's modern applications while giving organizations a cost-effective, adaptable solution for meeting both immediate and long term IT goals.

THE APPLICATION DESIGN CHALLENGE

Competitive businesses need their applications to provide seamless customer experiences with fast response times. Additionally, businesses rely on the data generated by these applications to make critical business decisions. This 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 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, 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.

Developers need a solution that boosts application performance, meets tough requirements, and offers the flexibility that hybrid cloud environments require to innovate. Red Hat JBoss Data Grid overcomes these challenges by offering:

- High availability to access data-retention.
- Elastic scalability to maintain fast response times.
- Support for open hybrid cloud environments.
- The flexibility to deploy your application data, however best fits your needs.

TRUE CORPORATION ELIMINATES A BILLING SYSTEM BOTTLENECK

A system bottleneck was discovered through engagement in the billing system upgrade, and while not constant, the bottleneck impacted overall performance. JBoss Data Grid was deployed, improving average response time from 10 seconds to 50 milliseconds.

After rounds of evaluation and testing, True bought a JBoss Data Grid subscription and engaged Red Hat Consulting to implement in-memory caching to reduce the interaction between the integration layer and billing. Red Hat JBoss Data Grid will also be used to cache frequently accessed information such as product catalog, promotions, and discount information

PRODUCT COMPONENTS

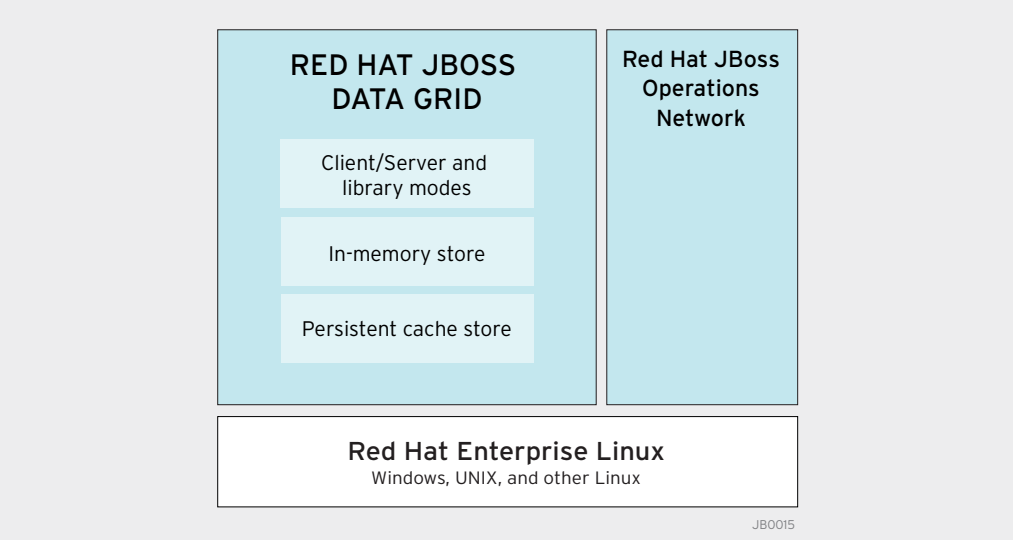

JBoss Data Grid consists of:

- Two deployment modes:
  - Library mode for deployment with the same Java™ Virtual Machine (JVM) as a Java application.
  - Client/Server mode to separate the application from its cache, allowing JBoss Data Grid to scale and be maintained independently. With Client/Server mode, JBoss Data Grid can connect to any type of application (Java, C++, C#, and so on) via the memcached, Hot Rod, or REST protocol.
- The in-memory store, which is the main component for storage and retrieval of the key-value entries.
- A persistent cache store for optional permanent storage of cached entries and restoration after a data grid shutdown.

Additionally, a subscription to JBoss Data Grid includes JBoss Operations Network for a single point of control to deploy, manage, and monitor JBoss Data Grid and your other Red Hat JBoss Middleware applications and services.

POPULAR CONFIGURATIONS

With JBoss Data Grid, you decide how to store and access your data. If real-time access is your priority, you might decide to deploy JBoss Data Grid in a replicated configuration, where every entry is replicated across the grid. Or perhaps superior elastic scalability is more important. For that scenario, a distributed data grid may be your best option.

REPLICATED CONFIGURATION

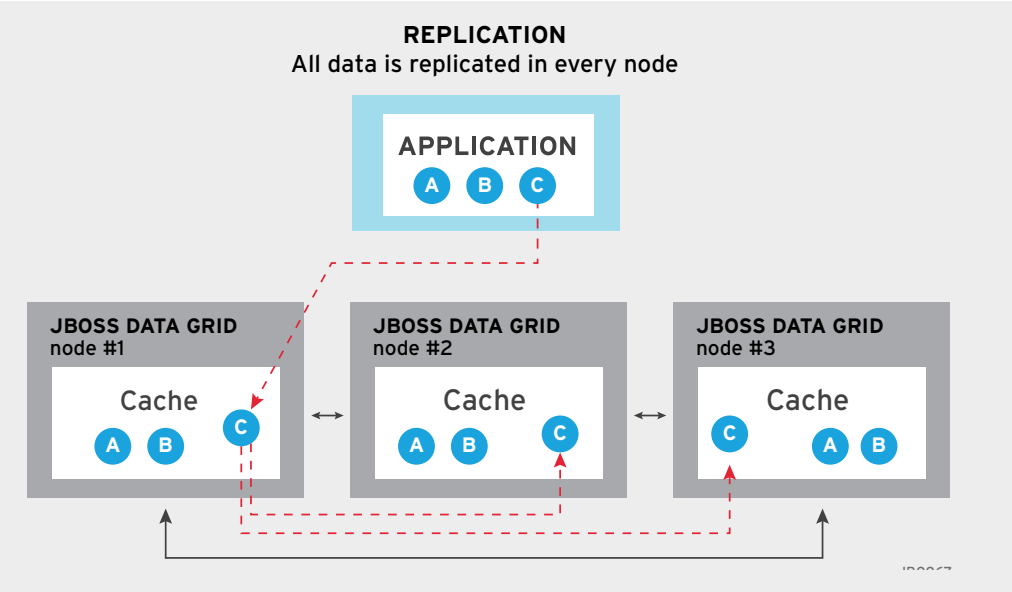

In the replicated configuration, the same exact data exists on all nodes of the clustered cache. This can be ideal for applications that require extremely high fault tolerance, and real-time read access by multiple applications. The amount of data available for in-memory processing is limited by the amount of memory on each server.

DISTRIBUTED CONFIGURATION

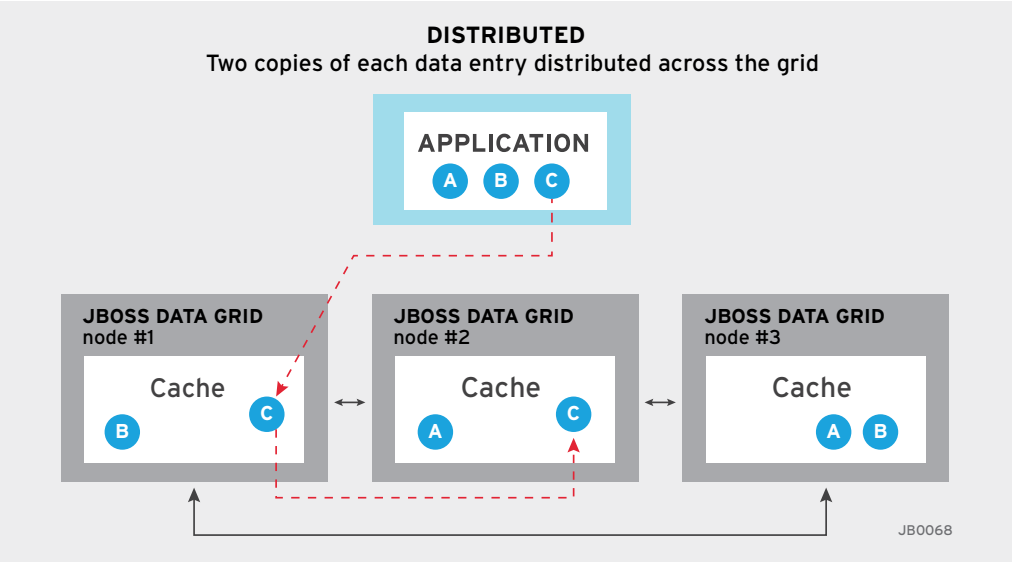

With the distributed configuration, you can take full advantage of JBoss Data Grid's elastic scalability. Choose how many copies of each data entry you want in the grid, and give your applications fast access to shared data. Ideal for elastic data sets that experience large fluctuations, periodicity or unpredictability, this configuration allows you to scale out above the available memory of a single node.

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

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 with 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
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.
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: • High availability to meet SLA requirements for data within and across data centers. • Load balancing and resource efficiency 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.
Security for sensitive data	Ensure only trusted users, service and applications are able to access and manage your sensitive and personal identification data in the grid.

FEATURE	BENEFIT
Multiple access protocols with data interoperability	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.
Zero-downtime rolling upgrades	Upgrade your cluster from one version to the next 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

RUNTIMES CONTAINERS

- Red Hat JBoss Enterprise Application Platform
- Red Hat JBoss Web Server
- Apache Karaf
- Oracle WebLogic

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 of 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)

RED HAT JBOSS DEVELOPER STUDIO

Also included with 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.

TRY JBOSS DATA GRID WITH THESE OTHER JBOSS MIDDLEWARE PRODUCTS:

- Red Hat JBoss Enterprise Application Platform
- Red Hat JBoss Data Virtualization

HOW CAN I GET RED HAT JBOSS DATA GRID?

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

PROCESSING AND DISTRIBUTION

- Map and reduce framework and programming model, including mapper, collector, reducer, and collator interfaces
- Distributed tasks to execute a task on some or 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 cluster communication 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; 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, C++); C# Hot Rod client offered in technology preview

SECURITY

- Secure communication between Client/Server (over Hot Rod) using TLS (SSL)
- Secure communication between server nodes using user-specific cryptography algorithms as supported by the Java Cryptography Architecture (JCA)
- Authentication using SSL Client certificate (SSL) between client-to-server and node-to-node
- Role-based authorization and access control to cache manager and caches

STORAGE

- SingleFileCacheStore for fast read and write performance where the index of keys is stored in-memory
- LevelDB cache store for fast read and writes of large amounts of data persisted locally at each node where the index of keys is not stored in-memory
- JPA cache store to persist data to the database while preserving the data schema
- Java Database Connectivity (JDBC) cache loader/store
- Read-though, 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

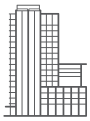

ABOUT RED HAT

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 80 offices spanning the globe, empowering its customers' businesses.

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

redhat.com
#12350827_v2_0714