WHITEPAPER

IN-MEMORY DATA GRIDS

The perfect solution for big data and application performance

IN-MEMORY DATA GRIDS RISE IN POPULARITY

As IT departments struggle to design and implement solutions to manage exponential data growth and meet strict requirements for application scale and performance, many of them are turning to in-memory data grids. In fact, according to 2013 research conducted by Gartner, "although the in-memory data grid market, a key IMC segment, is small, it is likely to grow fast and to reach $1 billion by 2016.".[1]

In-memory data grids are popular because they address two related challenges:

- **Access to big data for real-time use.** As organizations have exponentially increased the volume, velocity, and variability of the type of data they collect, process, and manage, traditional database designs are hard-pressed to keep up. This explosion of data presents challenges related to scalability, infrastructure design, security, timely access for critical decisions, and increased costs due to the increased complexity.
- **Application performance and scale.** Additionally, application development efforts must consider how internal database systems are in use to best determine their ultimate performance. For example, data storage systems and data access may not be able to keep up with the increasing volume of transactions in an organization's mission-critical applications—and lead to an adverse effect on the application's performance.

In-memory data grids present a smart solution posed by these two challenges by:

- Ensuring that the right data is available in memory with simple access. In-memory data grids promote extremely fast, scalable read-write performance.
- Automatically persisting unused data to a file system or maintaining redundant in-memory nodes to ensure high reliability and fault tolerance.
- Elastically spinning up and down distributed nodes.
- Automatically distributing information across the entire cluster so the grid can grow as your scale and performance needs change.

This document will further detail the challenges of big data and the potential effect on application performance. It will also discuss how in-memory data grids can present the perfect solution to these challenges as a high-performing, reliable, cost-effective, and low-impact means to scalability.

---

1 *Source: Gartner, Gartner Says In-Memory Computing Is Racing Towards Mainstream Adoption, http://www.gartner.com/newsroom/id/2405315, April 2013*

Additionally, this paper will introduce Red Hat® JBoss® Data Grid and present a use-case example of how it can help enterprise organizations solve their most critical challenges and achieve better business results.

TWO CHALLENGES POSED BY BIG DATA

Organizations of all sizes are attempting to scale their database systems and infrastructures to store terabyte—and even petabyte—scale data sets. They simply can't afford not to; real-time access to the vast amount of newly available data offers the business an advantage their competitors see as well. As the amount of data continues to grow, there is business value to be had by making better use of the data available to quickly respond to market demands.

Yet, as they attempt to take advantage of big data, previous strategies to manage growth simply won't work. Traditional databases are notoriously difficult and expensive to scale and may not be well-suited to handle the type of data in question. For example, in addition to the structured data that business operations have always relied on—names, addresses, phone numbers, and so on—there is a bulk of new information, such as unstructured data from mobile devices, social media, and log files, that may not easily fit into a relational world. Moreover, the semantics of those types of data require systems to handle massive volumes and real-time velocities in order to truly meet the business requirements.

These three dimensions—volume, velocity, and variability—have changed the game for data storage and access.

- **Volume.** As a result of ongoing business processes and specialized initiatives, organizations are now able to collect much more data than before. To fully understand the impact of increased data volume, consider the proliferation of photos on social sites such as Facebook, which currently hosts over 40 billion user photos.
- **Velocity:** At the same time, real-time business transactions—logistics updates, financial services trades, shipping information—have significantly increased the speed and volume of data that organizations now collect, store, and share via web-based customer service applications.
- **Variability:** Organizations save and manage many different types of structured and unstructured data. Examples of this includes relational data, email, GPS logistics, log files, social media data, metadata, and machine-created data like log files and timestamps. Due to the variability of all this data and due to variability in how an organization collects each type of data, it can be difficult for organizations to create a single view of a customer or a specific order.

As the concept of big data has opened new fields for business, it has redefined the requirements for compute and storage. Traditional proprietary databases and data-warehousing models either can't scale or can't do so in a cost-effective manner due to the high costs of hardware, software licenses, software support, and expertise. As a result, designing and implementing database solutions capable of working with big data, as well as the necessary teams to do so, leads to new set of challenges for IT departments, particularly around reliability, security, and performance.

ADDITIONAL CONSIDERATIONS FOR BIG DATA SOLUTIONS

When it comes to designing a data solution capable of working with big data, there is no single, one-size-fits-all solution. NoSQL, Hadoop, and caching are all viable options for specific problems. Yet when it comes to truly capitalizing on the full potential of big data for powering applications, there are very specialized requirements that can't be overlooked.

RED HAT JBOSS DATA GRID IN ACTION

Financial services trading

One of the world's largest financial trading companies was looking to get greater scalability, performance, and high availability out of its trading system. This goal was driven by the need to better manage and access the volume of data its transactional systems in a reliable manner.

Red Hat JBoss Data Grid gave this financial services company the scalability it required. As a highly scalable, distributed, in-memory database, Red Hat JBoss Data Grid enabled the organization to cost-effectively scale data tiers without making further compromises to its infrastructure.

Additionally, this in-memory data grid gave the company a highly available solution to provide extremely low latency with minimal to zero downtime. Even better, as an open-source solution, Red Hat JBoss Data Grid could be implemented with minimal disruption to existing applications, enabling IT and the entire organization to continue to focus on additional critical business initiatives.

- **Accessing the vast amounts of data at the right time:** As the amount of data grows, traditional data stores will inevitably become a major performance bottleneck for applications. As the cost of memory continues to decrease, the idea of moving much of that data into memory has become a viable, mainstream option and can present the right solution.

  In-memory databases have the core functionality, resilience, and redundancy required to serve as a managed cache, serve as an intermediate layer between relational databases, or stand on their own to meet these requirements. As a result, organizations can fully reap the benefits of access to big data while maintaining operational, real-time performance with in-memory speed.

- **Scalability:** For many applications, data demands may change based on the organization's specific business requirements. For example, seasonal business initiatives, marketing efforts, special promotions, and other variables can all influence the scaling requirements for a particular web application.

  Yet, scaling efforts that consist of adding new database servers are expensive, and complex sharding strategies can be both risky and expensive. And in many cases, simply adding more hardware to a traditional database will not solve the scaling challenge. Organizations need to be able to automatically distribute and partition the data, as well as manage data affinity, in order to scale linearly and seamlessly by simply spinning up new instances.

- **The need to implement non-disruptive technology:** When data becomes a bottleneck, rebuilding a particular application is usually not an attractive idea. Ideally, the solution to data performance and scaling issues should be integrated invisibly into the existing application without disrupting the current IT assets or infrastructure or requiring new expertise.

  For example, if an organization is already using a standard protocol access application data and persist that data to an existing relational database, it should be able to switch over to a new solution without the need for major application rewrites. Ideally, that new solution would continue to work with the existing data solutions for long-term data needs such as archiving, regulatory compliance mandates, serving other back-end business applications, and more.

RED HAT JBOSS DATA GRID

In-memory data grids meet those additional considerations and provide a reliable, transactionally consistent, distributed in-memory data store. In-memory data grids can be used to incrementally extend the performance and scalability of established applications. They can also be used to implement brand-new high-performance and high-scale applications—all based on the idea of using main memory for fast access, distributing data to scale, and working with another master data repository or maintaining duplicate, remote nodes to provide resilience and persistence.

As a leader in the run-time application servers that power numerous high-scale applications, Red Hat is familiar with the challenges posed by real-time access to big data, specifically its effect on application performance. Red Hat JBoss Data Grid address these issues.

JBoss Data Grid provides a way to scale the data tier without an expensive rewrite and using the expertise you have in-house. This gives enterprises a flexible and cost-effective way to improve application performance and get more out of their mission-critical applications.

Based on Infinispan, an open source JBoss Community project, Red Hat JBoss Data Grid provides fast caching along with data grid capabilities such as map/reduce framework, distributed execution and querying. JBoss Data Grid includes the following features, capabilities, and benefits:

- **Schema-less key value store.** Red Hat JBoss Data Grid is an in-memory NoSQL database that provides a simple and flexible means to storing different objects without a fixed data model.
- **Grid-based data storage and distributed processing.** Red Hat JBoss Data Grid is a federated cluster of services to distribute, replicate, and store data, designed to be seamlessly implemented with no additional development work. Processes can be pushed to these data services to distribute workload and provide high performance due to affinity and efficient use of compute power.
- **Elastic scaling.** Adding and removing storage nodes is simple, 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.
- **Multiple access protocols.** It is easy to access the data grid using REST, memcached, Hot Rod, 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.
- **Management tooling.** JBoss Data Grid is available as a subscription from Red Hat and also includes a subscription to Red Hat JBoss Operations Network. JBoss Operations Network is 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.

Red Hat JBoss Data Grid is appropriate for any type of application that has heavy compute needs. And when necessary, the solution integrates with other transactional services such as message queues to ensure reliability and consistency. Typical use-case scenarios for the JBoss Data Grid include financial services trading, logistics, eCommerce, and call center tracking.

Red Hat JBoss Data Grid offers the benefits of scale and high performance without the costs of rewriting or replacing the data tier. As with all Red Hat software, it is fully open source and available as a subscription, preventing vendor lock-in and saving the licensing costs associated with other data-caching alternatives.

CONCLUSION

When it comes to solving the many challenges posed by big data—and the resulting impact on web application performance—Red Hat JBoss Data Grid is a smart solution. As the best way to build high-performing, resilient, and scalable web applications, it can help your organization reap all the potential benefits big data offers for maximum business insight and results.

To learn more about Red Hat JBoss Data Grid, visit **www.redhat.com/jboss**

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.

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