JBOSS® ENTERPRISE MIDDLEWARE



DATASHEET

JBOSS DATA GRID 6: BIG DATA = BIG OPPORTUNITIES

Is data changing the way you do business? Is it inventory in your warehouse, or packages on the delivery truck? How about your users' profiles or other information on your website that drives consumers to the right products online. The good news is that data-driven applications enhance online customer experiences, leading to higher customer satisfaction and retention, and increased purchasing. The challenge is that the capture, retention, and use of all of this data can be costly to scale and manage. Solving these Big Data challenges is a big part of IT planning and strategy.

Big Data is a very broad topic and there is no one-size-fits-all solution. NoSQL, Hadoop, and elastic caching are all viable Big Data solutions that depend on specific data needs. However, when it comes to putting Big Data to work powering your web applications, there are some specialized requirements.

PUTTING BIG DATA TO WORK

High performance in-memory database - As the amount of data grows, reads and writes to traditional back-end data stores becomes a major performance bottleneck for web applications. As the cost of memory has continued to decrease, the notion of storing data in memory is becoming mainstream. Additionally, in-memory databases can act as intermediate layers between relational stores to meet data retention requirements. Implementing an in-memory database can significantly improve application performance.

Scales elastically – For many applications, data demands may ebb and flow based on different business requirements. Seasonal business, marketing and promotions, and external market changes can all push the scaling requirements for a web application. Having to add more database servers to solve this issue is neither simple nor inexpensive. In some cases, adding more databases may not even solve the scaling issue. Scaling these types of applications requires a simple way to add more data nodes that are instantly aware of the other nodes in a data grid. It should also be easy to remove nodes as business demands slow down.

Non-intrusive - If data is your bottleneck, rebuilding the application is not an attractive idea. It is better if solving the data issues can be more easily integrated into the existing application. For example, if you are using Hibernate to persist application data to an existing relational database (RDBMS), you should be able to easily switch over to an in-memory data grid without major application rewrites. It should also be simple to enable the data grid to work with other data solutions such as an RDBMS for long-term data needs (archival, regulatory, back-end business applications).

JBOSS° ENTERPRISE MIDDLEWARE



INTRODUCING RED HAT JBOSS DATA GRID 6

As a leader in the tooling and run-time servers that power numerous high-scale web applications, Red Hat is well aware of the Big Data challenge and its impact. Providing a means to more effectively scale the data tier gives enterprises a flexible and cost-effective means to do more with their applications. Based upon Infinispan, the JBoss® Community project, Red Hat® JBoss Data Grid is a distributed in-memory data grid that provides the following capabilities:

- Schemaless key value store Red Hat JBoss Data Grid is a NoSQL database that provides
 the flexibility to store different objects without a fixed data model.
- Grid-based data storage Red Hat JBoss Data Grid is designed to easily replicate data across multiple distributed nodes.
- Elastic scaling Adding and removing storage nodes is achieved simply and is non-disruptive.
- Multiple access protocols It is easy to access to the data grid using REST, memcached, Hot Rod, or simple map-like API.

HOW CAN I USE RED HAT JBOSS DATA GRID?

Red Hat JBoss Data Grid can be used in any type of web application that requires heavy interaction with a data tier. It performs especially well in transactional applications. Typical use cases for data grids include trading, logistics, social, and e commerce applications.

WHAT MAKES RED HAT JBOSS DATA GRID UNIQUE?

Red Hat JBoss Data Grid provides a number of unique capabilities including the ability to significantly improve the performance and scaling for web applications. It is a great way to mitigate the costs of data scaling by reducing the reliance upon relational databases and storage in favor of smaller distributed nodes. Red Hat JBoss Data Grid is also based on open source technologies and is available as a subscription versus the licensing costs associated with other data caching products.

HOW CAN I GET RED HAT JBOSS DATA GRID?

Red Hat JBoss Data Grid is generally available from Red Hat and its partners in North America and EMEA.

SALES AND INQUIRIES

NORTH AMERICA 1-888-REDHAT1 www.redhat.com EUROPE, MIDDLE EAST AND AFRICA 00800 7334 2835 www.europe.redhat.com europe@redhat.com ASIA PACIFIC +65 6490 4200 www.apac.redhat.com apac@redhat.com LATIN AMERICA +54 11 4329 7300 www.latam.redhat.com info-latam@redhat.com