

YOUR DATA ALWAYS AVAILABLE FOR APPLICATIONS AND USERS

Companies rely on data to power their day-to-day operations. It is imperative that this data be always available. Even minutes of application downtime can mean lost sales, a poor user experience, and a bruised brand. This can add up to millions in lost revenue. Most databases work at small scale, but how do you scale out, up, and down predictably and linearly as your data grows?

You need a different database. Basho Riak® KV Enterprise is a distributed NoSQL database architected to meet your application needs. Riak KV provides high availability and massive scalability. Riak KV can be operationalized at lower costs than traditional relational databases and is easy to manage at scale.

Riak KV integrates with Apache Spark, Redis Caching, Apache Solr, and Apache Mesos to reduce the complexity of integrating and deploying other Big Data technologies.

“ *The Weather Company requires an architecture that is both flexible and reliable, allowing us to deliver higher accuracy through superior data. Basho has been a valuable partner in our transformation, and Riak KV has proven to be a critical component as the NoSQL distributed database powering our new platform.* ”

– Bryson Koehler, EVP and CIO, The Weather Company

RIAK KV BENEFITS

GLOBAL AVAILABILITY

A distributed database with advanced local and multi-cluster replication means your data is always available.

MASSIVE SCALABILITY

Automatic data distribution in the cluster and the ease of adding nodes mean near-linear performance increase as your data grows.

OPERATIONAL SIMPLICITY

Easy to run, easy to add nodes to your cluster. Operations are powerful yet simple. Ensure your operations team sleeps better.

FAULT TOLERANCE

A masterless, multi-node architecture ensures no data loss in the event of network or hardware failures.

FAST DATA ACCESS

Your users expect your application to be fast. Low latency means your data requests are served predictably even during peak times.

FLEXIBLE DATA MODEL

Key/value NoSQL data model provides flexibility with no pre-defined schema.

GET TO THE CLOUD

For business continuity or to meet your peak demand, Riak KV excels in private, public, and hybrid cloud deployments.

OBJECT STORAGE

Multi-model support from a single platform via integration with Riak S2 for large object storage.

SIMPLIFIED DEVELOPMENT

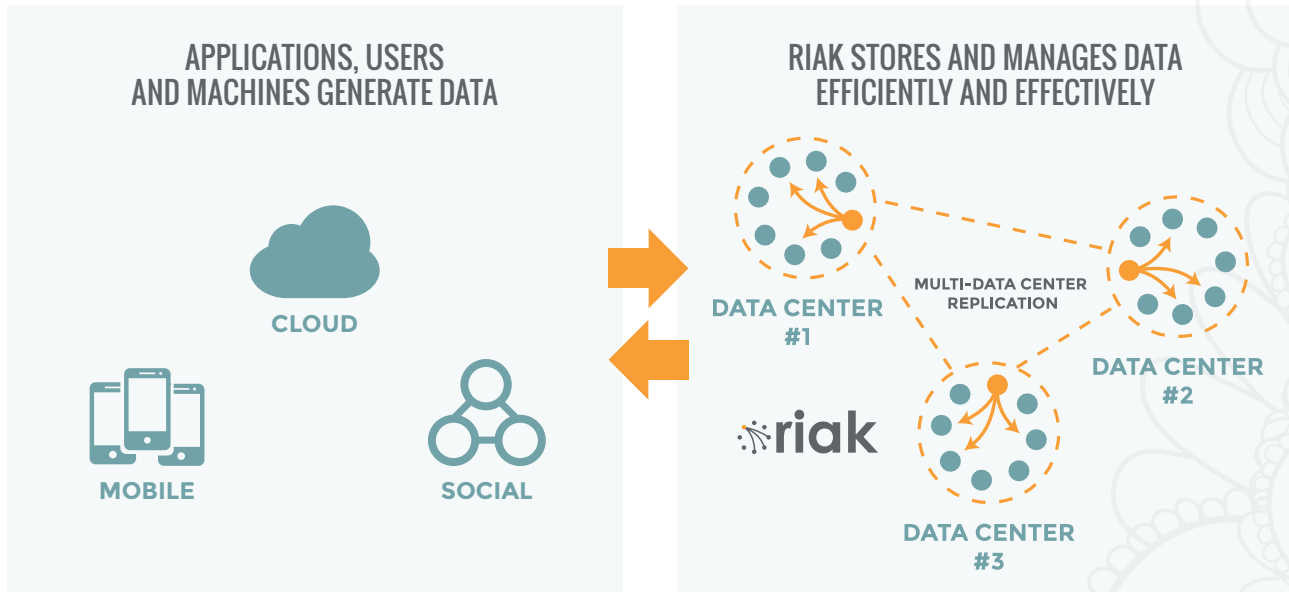
Extensive documentation and packaging tools have you up and running in minutes, and powerful APIs are easy to use.

LOWER TOTAL COST OF OWNERSHIP

Lower cost to operationalize than traditional relational databases.

WEB APPLICATION CHALLENGES

Traditional relational databases were not designed to cope with modern application requirements. Today's IT departments must support global access by millions of users on mobile devices. This not only requires geographic data distribution, but also the ability to handle massive amounts of data. Riak KV is architected to better handle a variety of web application challenges, including tracking user sessions, storing fast-growing unstructured or connected device data, and ensuring globally distributed reads and writes are fast.



USER SESSIONS

Session data is used by an application as it interacts with the end user. Session data is typically passed from the end users' client (browser, phone, etc.) to the server where it is stored, awaiting the return of new session data containing changes from the user. This session data is often critical to ensure user engagement, to download content or software, and to complete transactions or purchases. Riak KV is uniquely architected to handle this type of data. It is designed to never lose a write and to scale horizontally, so that even on peak days all your users' actions are completed seamlessly.

CONNECTED DEVICE DATA

The Internet of Things (IoT) and web applications gather and host vast quantities of data generated frequently, often by thousands or even millions of devices. This data can be time-series data updated at hour, minute, second, or even millisecond intervals. Riak KV is a data repository that scales in an unbounded and cost-effective manner in order to retain this quickly generated — and often unstructured — data. Riak KV enables application processing of this data to generate useful conclusions and actionable information. It is designed to scale horizontally with commodity hardware, making it easy for administrators to grow the data collection repository without creating complex sharding.

GLOBAL APPLICATIONS

Web applications are global. Your users are everywhere. Riak KV has an innovative database architecture that provides fast read and write functionality for globally distributed data. Riak KV is designed for a masterless configuration. This means that administrators can deploy multiple Riak KV clusters and then replicate to keep them all synchronized. For example, if a write is received by Cluster A, then Cluster A will assure that the write is replicated to Clusters B - Z.

RIAK KV FOR YOUR APPLICATIONS

RIAK KV

HIGH AVAILABILITY

Riak KV replicates and retrieves data intelligently, making it available for read and write operations even in failure conditions.

SCALABILITY

Riak KV automatically distributes data around the cluster and yields a near-linear performance increase as capacity is added.

OPERATIONAL SIMPLICITY

Riak KV allows you to add machines to the cluster easily, without a large operational burden.

LOW LATENCY

Riak KV is designed to store data and serve requests predictably and quickly, even during peak times.

FAULT TOLERANCE

Riak KV is fault tolerant so you can lose access to nodes due to network partition or hardware failure and never lose data.

ROBUST APIs and CLIENT LIBRARIES

PBC and HTTP APIs provide developer flexibility to meet your application needs. Supported languages include: Java, Ruby, Python, C#, Erlang, .NET, and Node.js.

RIAK KV DATA TYPES

A powerful, easy-to-use way to store data and handle merge conflicts. Conflicts are resolved automatically with conflict-free replicated data types (CRDTs): flags, registers, counters, sets, and maps. You don't have to write code to deal with data conflicts.

MULTI-MODEL

Riak supports multiple distributed NoSQL database models, including key/value, time series, and large object storage.

RIAK KV ENTERPRISE

Riak KV Enterprise includes ALL the features of Riak KV plus multi-cluster replication, monitoring, and customer support.

MULTI-CLUSTER REPLICATION (MULTI-DATACENTER REPLICATION)

Customers use cluster replication to serve global traffic, maintain active backups, run secondary analytics clusters, or meet disaster recovery and regulatory requirements. You can configure clusters to meet your business needs:

- **Active Cluster Configuration** – Achieve data locality by serving clients at low latency from the nearest datacenter.
- **Availability Zones** – Efficient multi-cluster replication and data redundancy within a geographic region (such as a coast or a country).
- **Secondary Analytics Clusters** – Replicate data from the primary cluster (responsible for serving all production requests) to a secondary cluster on which analytic and other computations can be performed.
- **Private, Public, or Hybrid Cloud** – Clusters can span the globe and support your cloud implementation. Replication ensures scalability with availability across your cloud.

MONITORING

Riak KV Enterprise supports both SNMP, shipping with an SNMP server built in, and JMX monitoring.

24 / 7 CUSTOMER SUPPORT

Riak KV Enterprise includes 24 / 7 access to Basho's Client Services team, including 1-hour response time for emergency production help. Basho's support team has extensive experience with production Riak KV installations and has worked on some of the largest Riak KV clusters in the world. Enterprise licensees have unlimited access to that experience and knowledge. Basho provides SLAs based upon the severity of the issue with 24x7 coverage.

Masterless




Multiple nodes operate to form a single cluster, making Riak KV masterless

Scalable



Scale out, up, and down predictably and linearly by adding nodes to the cluster

Global



Replicate data across multiple datacenters, globally, leveraging Riak KV Enterprise

RIAK KV CUSTOMERS

Businesses building I/O intensive, critical data applications that require a flexible, highly available, powerful yet operationally simple database platform will benefit with Riak KV.

Some of the most data intensive applications and businesses are using Riak KV:

The
Weather
Company

Virgin
america

yammer™

YAHOO!
JAPAN™

BEST
BUY

bet365

NHS

RIK
GAMES

Tapjoy™

Comcast.

View more of our users at basho.com/riak-users/

GET STARTED WITH RIAK KV

If you are interested in Riak KV Enterprise and would like to discuss your possible use case, please contact us at techtalk@basho.com.

To read about Riak KV and Riak KV Enterprise, including more in-depth technical details for developers and operators, visit our documentation portal at <http://basho.com/resources/downloads>.

Riak KV is available open source for download at <http://docs.basho.com/riak/latest>.

For more information visit www.basho.com or follow us on Twitter at www.twitter.com/basho.



ABOUT BASHO TECHNOLOGIES

Basho, the creator of the world's most resilient databases, is dedicated to developing disruptive technology that simplifies enterprises' most critical distributed systems data management challenges. Basho has attracted one of the most talented groups of engineers and technical experts ever assembled devoted exclusively to solving some of the most complex distributed systems challenges presented by Big Data and IoT.

Basho's database, Riak® KV, the industry leading distributed NoSQL database, is used by fast growing Web businesses and by one-third of the Fortune 50 to power their critical Web, mobile and social applications. Built on the same foundation, Basho introduced Riak TS, which is the first enterprise-ready NoSQL database specifically optimized to store, query and analyze time series data. Basho also provides Riak integrations for a variety of Big Data technologies like Apache Spark, Redis, Mesos, and Apache Solr.

For more information visit Basho.com which is full of interesting use cases, customer case studies and product detail, or docs.basho.com for technical documentation.



BASHO TECHNOLOGIES, INC
10900 NE 8TH STREET
SEATTLE, WA 98004

617.714.1700 // WWW.BASHO.COM