

# relational → riak

## Introduction

This whitepaper looks at why companies choose Riak over a relational database. We focus specifically on availability, scalability, and the key/value data model. Then we analyze the decision points that should be considered when choosing a non-relational solution and review data modeling, querying, and consistency guarantees. Finally, we end with simple patterns for building common applications in Riak using its key/value design, dealing with data conflicts that emerge in an eventually consistent system, and discuss multi-datacenter replication.

## Why Move From Relational to Riak?

### High Availability

Strongly consistent operations provide applications with guarantees that read operations will reflect the last successful update to your database. They are an important part of relational database systems, enabling operations (like transactions) that are essential for some types of applications. Strong consistency is relatively straightforward when there is one large single relational database server. Once your dataset grows beyond the capacity of a single machine, it becomes necessary to scale the database and operate in a distributed fashion.

Relational databases typically address the challenge of availability with a master/replica architecture, where the topology of a cluster is comprised of a single master and multiple replicas. Under this configuration, the master is responsible for accepting all write operations and coordinating with replicas to apply the updates in a consistent manner. Read requests can either be proxied through the master or sent directly to the replicas.

But what happens when your relational master server fails? The database will favor consistency and reject write

### Riak At-a-Glance

- **Distributed NoSQL database**
- **Helps you store, manage, and secure unstructured data**
- **Fault-tolerant, highly available platform**
- **Lower cost to operationalize than traditional relational databases**
- **Scales both up and out; Easy to manage at scale**
- **Add commodity hardware, as needed, for additional capacity**

### APPLICATIONS, USERS, AND MACHINES GENERATE UNSTRUCTURED DATA



Cloud



Mobile



Social


### DOWNTIME CAN COST YOU MILLIONS

For many of today's applications and platforms, high availability is the most important requirement.

operations until the failure is resolved. This can lead to a window of write unavailability, which is unacceptable in many application designs. Most master/replica architectures recognize that a master is a single point of failure and will perform automatic master failover (a replica is elected as a new master when failure of the master is detected).

In contrast, Riak is a masterless system designed for high availability, even in the event of hardware failures or network partitions. Any server (termed a "node" in Riak) can serve any incoming request and all data is replicated across multiple nodes. If a node experiences an outage, other nodes will continue to service read and write requests.

**RIAK IS MASTERLESS,  
DEPLOYED AS A CLUSTER OF NODES**



Further, if a node becomes unavailable to the rest of the cluster, a neighboring node will take over the

responsibilities of the missing node. The neighboring node will pass new or updated data (termed “objects”) back to the original node once it rejoins the cluster. This process is called “hinted handoff” and it ensures that read and write availability is maintained automatically to minimize your operational burden when nodes fail or comeback online.

For many of today’s applications and platforms, high availability has a direct impact on revenue. A few examples include: cloud services, online retail, shopping carts, gaming and betting, and advertising. Further, lack of availability can damage user trust and result in a poor user experience for many social media and chat applications, websites, and mobile applications. Riak provides the high availability needed for your critical applications. It provides a better solution than relational databases when availability is a key requirement.

## Minimizing the Cost of Scale

Relational databases scale by increasing the server and storage capacity. Special, expensive versions of the database may technically distribute load across multiple machines, but they rely on shared storage that often becomes a bottleneck. These systems are not designed to run on commodity hardware. An Oracle RAC system can cost millions to store a mere 20TB of data, a tiny number in the realm of Big Data where petabytes, exabytes, and zettabytes are the common units of measure. Some Riak customers ingest over 20TB of data per day.

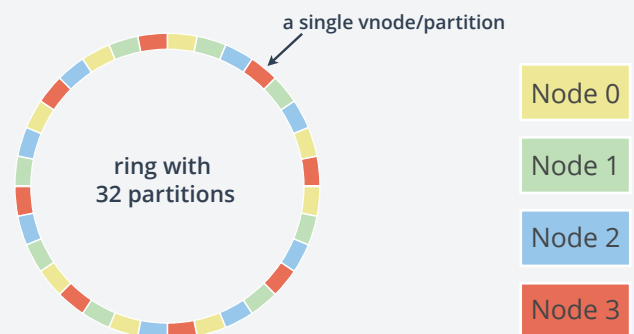
Another way databases scale is to use sharding. Sharding distributes data across several database servers. A common example of this would be putting your user data for differing geographical regions (e.g., US and EU) on different machines, or using an alphabetical or numerical order to split data. While this seems simple, sharding is complex and inherently inflexible.

First, writing and maintaining custom sharding logic increases the overhead of operating and developing an application on the database. Significant growth of data, or traffic, typically means significant, often manual, re-sharding projects. Determining how to intelligently split the dataset without negatively impacting performance, operations, and development presents a substantial challenge — especially when dealing with “big data,” rapid scale, or peak loads. Rapidly growing applications frequently outpace an existing

sharding scheme and, when the data in a shard grows too large, the shard must again be split. While several “auto”-sharding technologies have emerged in recent years, these methods are often imprecise and still require manual intervention. Sharding can often lead to “hot spots” in the database where a few physical machines become responsible for storing and serving a disproportionately high amount of both data and requests. Hot spots can lead to unpredictable latency and degraded performance.

In Riak, data is automatically distributed evenly across nodes using consistent hashing. Consistent hashing ensures data is evenly distributed around the cluster and new nodes can be added with automatic, minimal reshuffling of data. This significantly decreases risky hot spots in the database and lowers the operational burden of scaling.

**THE RIAK “RING” ARCHITECTURE**



Riak stores data using a simple key/value model. Data entries in Riak are referred to as “objects.” Key/value pairs are logically grouped together in namespaces called “buckets.” As you write new keys to Riak, the object’s bucket/key pair is hashed. The resulting value maps onto a 160-bit integer space. This integer space can be conceptualized as a ring that is used to determine where data is placed on physical machines.

**Riak is easier to manage, doesn’t require sharding, and runs on commodity hardware. This provides you with significantly reduced costs of scale over traditional relational databases.**

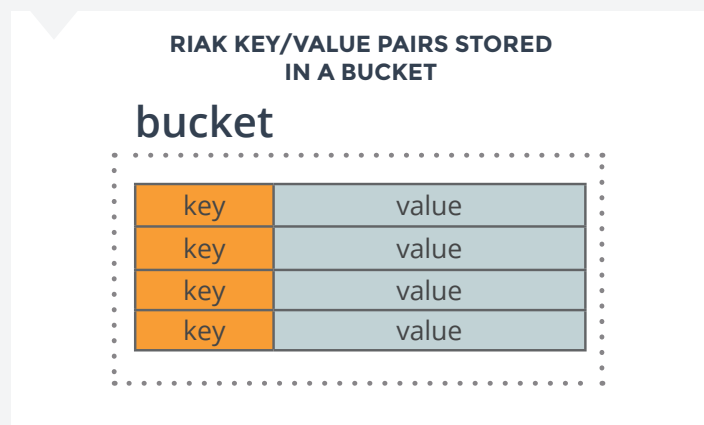
Riak tokenizes the total key space into a fixed number of equally sized partitions (default is 64). Each partition owns the given range of values on the ring and is responsible for

all buckets and keys that, when hashed, fall into that range. A virtual node (a “vnode”) is the process that manages each of these partitions. Physical machines evenly divide responsibility for vnodes.

Hashing and shared responsibility for keys across nodes ensures that data in Riak is evenly distributed. When machines are added, data is rebalanced automatically with no downtime. New machines take responsibility for their share of data by assuming ownership of some of the partitions; existing cluster members hand off the relevant partitions and the associated data. The new node continues claiming partitions until data ownership is equal. The ring state is shared around the cluster by means of a “gossip protocol.” Whenever a node changes its claim on the ring, it announces (i.e., “gossips”) this change to other nodes so that those nodes can respond appropriately. Nodes also periodically re-announce what they know in case any nodes happened to miss previous updates.

## Simple Data Models

With the tremendous growth in data — especially big data and new types of gaming, social, and mobile applications — there is a growing need to store unstructured data (i.e., data that does not fit easily in the rigid data model of relational systems). This unstructured data is more effectively stored in a simple and powerful key/value store.



Riak stores all types of objects. Riak, as a key/value platform, provides more flexibility and simplicity than a traditional relational database. Riak provides multi-model support including key/value, search, and object storage ensuring you

**The relational data model is complex and inflexible for storing massive amounts of unstructured data.**

can easily store and retrieve documents, images, videos, and much more.

With Riak there are no join operations since the key/value model has no concept of columns and rows. Riak is queried via HTTP requests, via the protocol buffers API, or through various client libraries. There are several powerful querying options:

- **Riak Search:** Integration with Apache Solr provides full-text search and support for Solr’s client query APIs.
- **Secondary Indexes:** Secondary Indexes (2i) give developers the ability to tag an object stored in Riak with one or more query values. Indexes can be either integers or strings, and can be queried by either exact matches or ranges of values.
- **MapReduce:** Developers can leverage Riak MapReduce for tasks like filtering documents by tag, counting words in documents, and extracting links to related data. It offers support for JavaScript and Erlang.

# Operational & Development Considerations

## Powerful Data Modeling In Riak

The table below illustrates key/value mappings for common application types. Values in Riak are stored on disk as binaries — JSON or XML documents, images, text, etc.

Application Type	Key	Value
Session	User/Session ID	Session Data
Advertising	Campaign ID	Ad Content
Logs	Date	Log File
Sensor	Date, Date/Time	Sensor Updates
User Data	Login, eMail, UUID	User Attributes
Content	Title, Integer	Text, JSON/XML/HTML Document, Images, etc.

Consider, for example, one of Riak's common use cases, storing user and session data. In a relational database, the "users" table is well known, providing a unique identifier per user, along with a series of identifying information for each user in individual columns, such as:

- First name
- Last name
- Interests
- Site Visits Counter
- Paid Account Identifier

In this example, the data can then be used to correlate or count paid users and/or common interests via a series of SQL queries against the row/column structure of the users table.

Riak, in contrast, provides flexibility in how this data can be modeled based upon the application use case. It may be desirable to create a Users bucket with the UserName (or Unique Identifier) as the key and a JSON object storing all user attributes as the value.

### USERS BUCKET

Key	Value
Unique Identifier (e.g. UserName)	JSON Object (stores all user attributes)

#### YOU CAN ALSO LEVERAGE THE POWER OF RIAK DATA TYPES BY CREATING A MAP TYPE FOR EACH USER, STORING:

- User's first and last name strings in registers
- User Interests as a set
- User visits in a counter
- A paid account identifier in a flag

You can also leverage the power of Riak Data Types by creating a map type for each user, storing:

- User's first and last name strings in registers
- User interests as a set
- User visits in a counter
- A paid account identifier in a flag

One of the best ways to enable application interaction with objects (a key/value pair) in Riak is to provide structured bucket and key names for the objects. This approach often involves wrapping information about the object in the object's location data itself.

For example, appending a timestamp, UUID, or Geographical coordinate to a key's name allows for fine-grained queryability via simple lookup to locate and retrieve a specific set of information. Leveraging the same naming mechanism as created for users (UniqueID as the key) enables, in a separate sessions bucket, storing the UUID append with a timestamp as the key and the session data (in binary format) as the object. In this way, using the same UUID, it is possible to obtain both user and session data stored in different buckets and in different formats.

## Resolving Data Conflicts

In any system that replicates data, conflicts can arise (e.g., if two clients update the same object at the exact same time or if not all updates have yet reached hardware that is experiencing lag). As discussed earlier, Riak is “eventually consistent” — while data is always available, not all replicas may have the most recent update at the same time, causing brief periods (generally on the order of milliseconds) of inconsistency while all state changes are synchronized.

However, Riak provides features to detect and help resolve the statistically small number of incidents when data conflicts occur. When a read request is performed, Riak looks up all replicas for that object. By default, Riak will return the most updated version. You can also resolve conflicts using Riak Data Types to simplify the burden of producing data convergence at the application level by absorbing the complexity into Riak itself.

Further, when an outdated object is discovered as part of a read request, Riak will automatically update the out-of-sync replica to make it consistent. Read Repair, a self-healing property of the database, will even update a replica that returns a “not\_found” in the event that a node loses it due to physical failure.

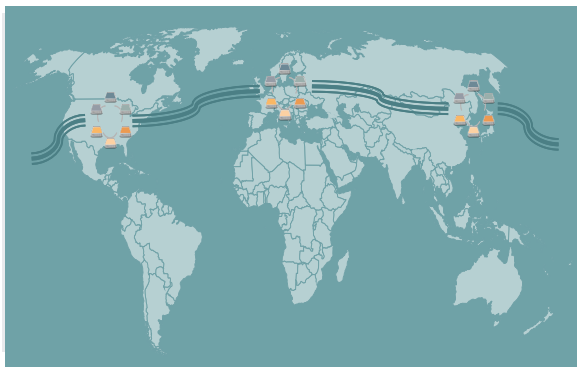
Riak also includes “Active Anti-Entropy,” which is an automatic self-healing property that runs in the background. Rather than waiting for a read request to trigger a replica repair (as with Read Repair), Active Anti-Entropy constantly uses a hash tree exchange to compare replicas of objects and automatically repairs or updates any that are divergent, missing, or corrupt. This can be beneficial for large clusters storing “stale” data.

## Multi-Datacenter Operations

Multi-site replication is critical for many of today’s platforms and applications. Not only does replication across multiple clusters provide geographic data locality (i.e., the ability to serve global traffic at low-latencies), but it can also be an integral part of a disaster recovery or backup strategy. You can also use multi-site replication to maintain secondary data stores, both for failover as well as for performing intensive computation without disrupting production load.

Multi-site replication in Riak works differently than the typical approach of relational databases. In Riak’s multi-datacenter replication, one cluster acts as a “primary cluster.” The primary cluster handles replication requests from one or more “secondary clusters” (generally located in datacenters in other regions or countries). If the datacenter with the primary cluster goes down, a secondary cluster can take over as the primary cluster. In this sense, Riak’s multi-datacenter capabilities are “masterless.”

### MULTI-DATACENTER OPERATIONS



In multi-datacenter replication, there are two primary modes of operation — full-sync and real-time. In full-sync mode, a complete synchronization occurs between primary and secondary cluster(s). In real-time mode, continual, incremental synchronization occurs and replication is triggered by new updates. Full-sync is performed upon initial connection of a secondary cluster, and then periodically (by default, every 6 hours). Full-sync is also triggered if the TCP connection between primary and secondary clusters is severed and then recovered.

Data transfer is unidirectional (primary → secondary). However, bidirectional synchronization can be achieved by configuring a pair of connections between clusters.

## Conclusion

Modeling data in any non-relational solution requires a different way of thinking about the data itself. Rather than an assumption that all data cleanly fits into a structure of rows and columns, the data domain can be overlaid on the core Key/Value store (Riak) in a variety of ways. There are, however, distinct tradeoffs and benefits to understand:

Choosing a database solution affects your infrastructure, budget, and staff. After deciding between a relational and NoSQL database, you need to figure out which specific solution will best meet your organization's needs. Riak is designed to deliver maximum data availability, to scale linearly using commodity hardware, and to provide powerful yet simple data models for storing massive amounts of unstructured data.

““ The new consolidated platform will capture 2.25 billion weather data points 15 times per hour. As with any large-scale, algorithmic-type modeling, the more data you have, the better the weather predictions will be. NoSQL won out over relational databases primarily for its scalability, but Riak was chosen for its simplicity and ease of administration at high scale. It won out over Apache Cassandra, MongoDB and Hadoop. ””

– The Weather Company



### TRADEOFFS AND BENEFITS

#### Relational Databases have:

- Tables
- Foreign keys and constraints
- ACID
- Sophisticated query planners
- Declarative query language (SQL)

#### Riak has:

- A Key/Value model where the value is any unstructured data
- More data redundancy that provides better availability
- Eventual consistency
- Simplified query capabilities
- Riak Search

#### What you will gain:

- More flexible, fluid designs
- More natural data representations
- Scaling without pain
- Faster response times
- Reduced operational complexity

#### Picking the right database for your team requires a careful understanding of:

- The requirements of your application or platform
- What developer productivity means to you
- The operational conditions you need
- How different database solutions support your goals.

## Next Steps

- To read about Riak and Riak Enterprise, including more in-depth technical details for developers and operators, visit our documentation portal at <http://docs.basho.com/riak/latest/>.
- Riak is available open source for download at <http://basho.com/resources/downloads/>.
- If you are interested in Riak Enterprise and would like to discuss your possible use case, please contact us at [techtalk@basho.com](mailto:techtalk@basho.com).
- For more information visit [www.basho.com](http://www.basho.com) or follow us on Twitter at [www.twitter.com/basho](http://www.twitter.com/basho).

