



# Technical Overview

Simple, Scalable, Object Storage Software

# Table of Contents

---

<b>Table of Contents</b> .....	<b>1</b>
<b>Introduction &amp; Overview</b> .....	<b>1</b>
<b>Architecture</b> .....	<b>2</b>
<b>How it Works</b> .....	<b>2</b>
<b>APIs and Interfaces</b> .....	<b>3</b>
Administration and Users.....	3
Storage API.....	3
Multipart Upload.....	3
Storage and Usage Statistics .....	3
<b>Operations</b> .....	<b>4</b>
Monitoring and Metrics .....	4
Scaling Out .....	4
<b>Riak S2 Enterprise:</b> .....	<b>4</b>
<b>Multi-Datacenter Replication</b> .....	<b>4</b>
<b>Next Steps</b> .....	<b>5</b>
<b>Getting Started</b> .....	<b>5</b>

# Introduction & Overview

---

Riak S2 (formerly Riak CS) is enterprise object storage software. It is open source and built on top of Riak, offering S3-compatible, multi-tenant large object storage. It can be used for building public or private clouds, or as reliable storage for applications and platforms.

*Use cases for Riak S2 include:*

- Storage for images, text, documents, videos and other large objects
- Internal storage providing on-demand capacity for business units
- A foundational storage layer for public clouds and cloud services
- Infrastructure for moving applications off of Amazon S3, or for hybrid solutions
- Redundant storage for backup and disaster recovery scenarios

*The primary features of Riak S2 are:*

- Highly available, fault-tolerant storage
- Large object support
- S3-compatible API, authentication, and authorization
- Multi-tenancy with per-user reporting on usage and network I/O
- Simple operational model for adding capacity

*Riak S2 Enterprise, a commercial offering, adds:*

- Multi-datacenter replication for active backups, disaster recover, and data locality
- Per-node or capacity-based pricing, inclusive of 24/7 support

This technical brief covers Riak S2 architecture, APIs, and operations with a brief note on Riak S2 Enterprise pricing. Additional technical details on Riak S2 can be found in the [online documentation](#).

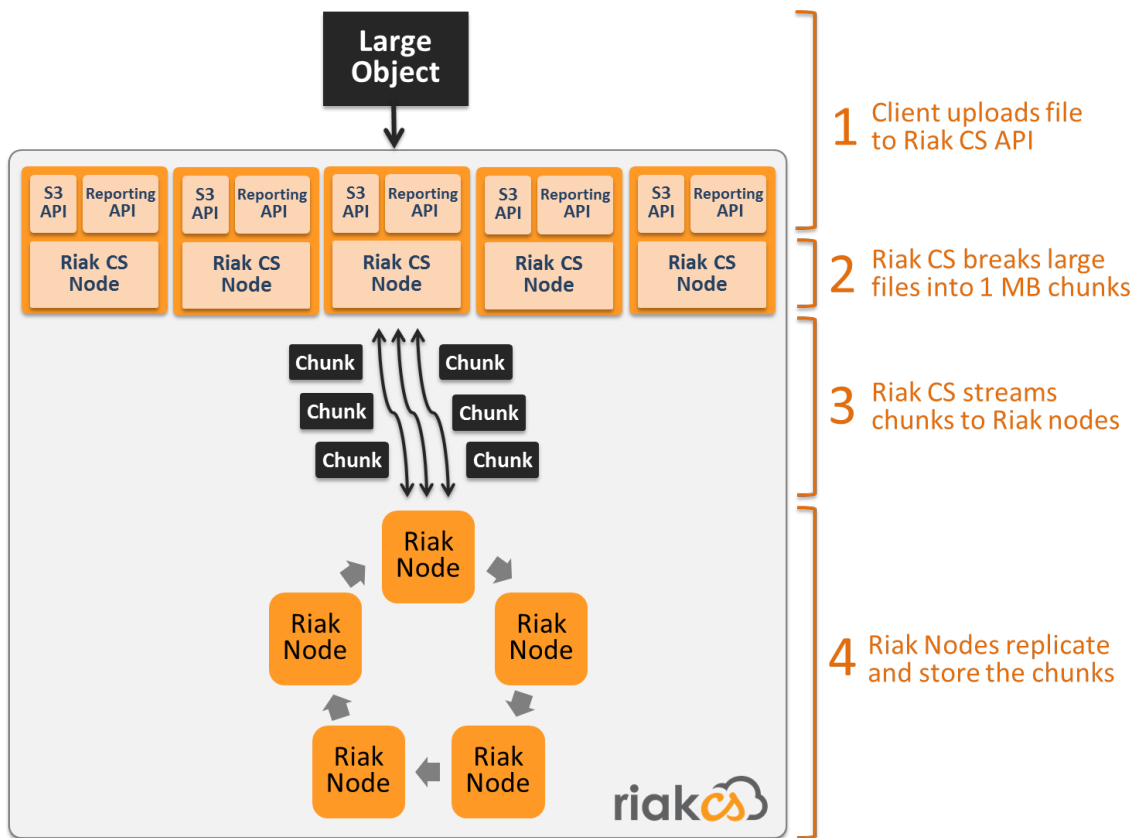
# Architecture

Riak S2 is a cloud storage product built on top of the open source, distributed database Riak. Riak S2 handles large file uploads and exposes an S3-compatible API, user and administrative functions, and usage reporting. The underlying Riak cluster stores objects and metadata for retrieval - providing replication, fault tolerance and high availability for data stored in the system. We recommend a one-to-one mapping of Riak KV (formerly named Riak) to Riak S2 nodes, so there should be one Riak KV node for every Riak S2 node. Ideally, Riak KV and Riak S2 nodes should be run on the same physical machine to minimize unnecessary network traffic.

## How it Works

In a Riak S2 system, any node can respond to client requests - there is no master node and each node has the same responsibilities. Since data is replicated (three replicas per object by default) and other nodes automatically take over the responsibility of failed or non-communicative nodes, data remains available even in the event of node failure or network partition. More information on how Riak handles failure conditions can be found in the [Riak documentation](#).

When an object is uploaded via the storage API, Riak S2 breaks the object into smaller chunks that are streamed, written, and replicated in Riak. Each chunk is associated with metadata for later retrieval. The diagram below provides a visualization.



## APIs and Interfaces

Riak S2 has HTTP interfaces for handling storage operations, authentication, administration, and per-user storage and bandwidth statistics. Below is an overview and additional details can be found in the [Riak S2 documentation](#).

## Administration and Users

Riak S2 exposes [an interface](#) for user creation, disablement, and credential management. Once created and issued credentials, users are able to authenticate, create buckets, upload and download files to their account, retrieve account information, obtain new credentials, or disable their account through the API. Data access can also be defined based on IP address. Riak S2 can be configured so only administrators can create users, or so that anyone can create a user directly. Administrator accounts have additional special privileges: they can retrieve a list of all users in the system and query the user account information of any user.

Riak S2 uses a request serializer to ensure that global entities (Riak S2 users and bucket names) are unique in the system. One instance of this serializer (Stanchion) must be installed in every Riak S2 system.

Riak S2 supports the standard S3 authentication scheme, with support for header and query parameter authorization. In addition, [integration with OpenStack's Keystone authentication](#) is available.

## Storage API

The Riak S2 storage interface is [S3-compatible](#) and can be configured for use with existing S3 tools, including s3cmd, official S3 libraries, and clients such as [Fog](#) and Boto. Like S3, the API has simple, RESTful GET, PUT, Copy, and DELETE operations for objects and buckets in addition to support GET range queries. Riak S2 also supports the addition of arbitrary metadata to objects so users can store more useful information about their objects. Additionally, S3-style ACLs are provided for managing object and bucket permissions. Riak S2 also supports [integration with OpenStack's Object Storage API](#).

## Multipart Upload

As of the Riak S2 1.3 release, Riak S2 supports [multipart uploads](#). At its simplest, multipart upload allows the uploading of a single large object as a series of parts. Once the individual parts have been uploaded, Riak S2 presents the data as a single object. The parts can be between 5MB and 5GB in size.

## Storage and Usage Statistics

Operators can use the Riak S2 [storage](#), [usage](#), and network statistics to support use cases like accounting, subscription, billing, or multi-group utilization for public or private clouds. Riak S2 will report information on how much storage a named user is consuming and the network operations related to accessing it. This data can be queried on the default timespan "now," or as a range from start time through end time. Access statistics are reported as bytes in and bytes out for both object and bucket operations. Reporting of this information can be scheduled for a set interval or manually triggered. Additional details billing and usage reporting are available in the [docs](#).

## Operations

Riak S2 is designed to be easy to operate and scale, with monitoring and metrics solutions and a simple model for adding capacity as needed.

## Monitoring and Metrics

Riak S2 exposes [stats](#) on critical operations that are accessible via HTTP request. Additionally, Riak S2 and Riak nodes both have DTrace support for analysis of running systems. In version 1.5 Syslog support was added to facilitate log aggregation.

## Scaling Out

Adding new capacity to a Riak S2 cluster requires installing both Riak S2 and Riak on a new physical node and joining the node to the cluster using several simple management commands. Riak automatically redistributes data in the system so all nodes have equal responsibility, preventing storage hot spots and decreasing the operational burden of adding new nodes.

As part of the Riak S2 Enterprise license, Basho offers set up, configuration, and operational support for the implementation. Basho also provides [Chef recipes](#) to automatically install Riak S2 on all nodes in the system.



# Riak S2 Enterprise: Multi-Datacenter Replication

Riak S2 Enterprise is a commercial extension of Riak S2 that provides multi-datacenter replication and 24/7 support.

Multi-datacenter replication for Riak S2 Enterprise provides two modes of object replication: full-sync and real-time sync. Data is streamed over a TCP connection, and multi-datacenter replication has support for SSL so data can be securely replicated between sites.

In Riak S2, large objects are broken into blocks and streamed to the underlying Riak cluster on write, where they are replicated for high availability (three replicas by default). A manifest for each object is maintained so that blocks can be retrieved from the cluster and the full object presented to clients. For multi-datacenter replication in Riak S2 Enterprise, global information for users, bucket information, and manifests are streamed in real-time from a primary implementation to a secondary site, so global state is maintained across locations. Objects can then be replicated in either full-sync or real-time sync mode.

In full-sync, objects are replicated from a primary Riak S2 Enterprise implementation to a secondary site on a configurable interval – the default is six hours. During full-sync replication, each cluster computes a hash for each key's block value. Key/block pairs are compared, and the primary site streams any missing blocks or updates needed to the secondary site.

Real-time sync is triggered when an update is sent from a client to a primary Riak S2 Enterprise implementation. Once replicated in the first location, the updates are streamed in real-time to the secondary site. But what if a client requests an object from the secondary cluster and not all of its blocks have been replicated to that cluster? With Riak S2 Enterprise multi-datacenter replication, the secondary cluster will request any missing blocks from the primary cluster so the client request can be served quickly and accurately.

Riak S2 Enterprise has two pricing models. Riak S2 Enterprise can be purchased on a per-node, annual subscription or perpetual license. The per-node price is inclusive of the Riak S2 nodes as well as the underlying Riak nodes required. Riak S2 Enterprise can also be purchased in blocks of raw storage. Under this model, at higher storage tiers, the effective per-gig price of the license is lower. Both models include 24/7 support. [Request a developer trial](#) of Riak S2 Enterprise to try it on your own infrastructure.

## Next Steps

---

### Getting Started

Riak S2 packages are available on our [downloads page](#) and source code is available on [Github](#).

If you are interested in Riak S2 Enterprise, please [contact us](#). We would love to talk to you about your possible use case and needs.

#### ABOUT BASHO

Basho is a distributed systems company dedicated to developing disruptive technology that simplify enterprises' most critical 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 issues presented by scaling distributed systems.

Basho's distributed database, Riak® KV, the industry leading distributed NoSQL database, and Basho's cloud storage software, Riak® S2, are used by fast growing Web businesses and by one third of the Fortune 50 to power their critical Web, mobile and social applications. The Basho Data Platform helps enterprises reduce the complexity of supporting Big Data applications by integrating Riak KV and Riak S2 with Apache Spark™, Redis and Apache Solr™. Basho is the organizer of RICON — a distributed systems conference. Riak is the registered trademark of Basho Technologies, inc.