

RIAK KV[®] ENTERPRISE SOFTWARE DRIVES COMCAST HIGH AVAILABILITY OBJECT STORE

Comcast Cable is one of the nation's leading providers of entertainment, information, and communication products and services. Comcast is principally involved in the operation of cable systems through Comcast Cable and in the development, production, and distribution of entertainment, news, sports, and other content for global audiences through NBCUniversal. Comcast Cable is one of the nation's largest video, high-speed Internet, and phone providers to residential and business customers. Comcast is the majority owner and manager of NBCUniversal, which owns and operates entertainment and news cable networks, the NBC and Telemundo broadcast networks, local television station groups, television production operations, a major motion picture company, and theme parks.

CHALLENGE

Increase developer efficiency while retaining robustness and maintainability.

SOLUTION

Use Riak KV[®] Enterprise Software for highly-available flexible object storage system.



DATA DRIVES DEVELOPMENT SHIFT

In June 2009, Comcast embarked on the HOSS (Highly-Available Object Storage System) project to provide an internally-accessible, scalable key-value store with an HTTP interface for supporting its application development. According to Jon Moore, Technical Fellow at Comcast Interactive Media, "Our problem wasn't performance and scale as much as it was development efficiency. The key-value store is a simple model that allows us to store data on-the-fly, giving our DBA's the freedom to work on other issues." Jon went on to add, "A KV store is flat out simpler to use and manage than comparable, more general RDBMS technologies, if the use case fits. The key-value store may not be as flexible in expressive power as the RDBMS, but there are a lot of use cases where it's sufficient."

The Comcast Team initially deployed their system using an externally-hosted key-value elastic storage service because it was a readily-available, already-deployed system; hence they could start providing an internal key-value service quickly. While their externally-hosted service met many of their requirements, the additional overhead associated with data encryption and latency associated with multi-data center replication motivated the team at Comcast to look for an alternative elastic storage solution that would allow them to bring the system in-house and give them more control.

HOSS REQUIREMENTS

- High Availability
- Developer Flexibility
- Java HTTP interface
- Multi-Datacenter Replication
- Schemaless Key/Value Semantics

“*...we don't have a dedicated Riak KV administrator—because we don't need one.*”

- Jonathan Moore, Technical Fellow, Comcast Interactive Media

COMCAST SELECTS RIAK KV® ENTERPRISE SOFTWARE

In October 2009, while attending the NoSQL East Conference in Atlanta, Georgia, Comcast team members sat in on Basho founder and Chief Technology Officer, Justin Sheehy's presentation of Riak KV Enterprise Software. What attracted the team most to Riak KV was the fact that it was inspired by the principles of Amazon's Dynamo, and its dedication to reliability and robustness. In essence, Riak KV Enterprise Software would allow Comcast Cable to meet and exceed the benefits of an externally-hosted key-value elastic storage service, inside the Enterprise, behind the firewall. In 2009 Basho was in the early stages of developing Riak KV. Comcast realized that committing to Riak KV would require a significant investment in time and money. After a

period of due diligence, Comcast Interactive Media group decided that the benefits, once the product reached a particular point of maturity, would far outweigh the cost of the investment. In April 2010, Comcast decided to move forward with Riak KV.

Initially, while helping to test and harden Riak KV, the Comcast team used Riak KV in parallel with an externally-hosted service, primarily as a write-only mirror. During this period the Comcast development team worked in a high-level pseudo-partnership with Basho, testing and hardening Riak.

THE INVESTMENT PAYS OFF

By November 2010, after successfully meeting all of the necessary diligence requirements, Comcast Cable began using Riak KV as the primary backing datastore for HOSS. Today the data in HOSS is used across xfinity.comcast.net and xfinitytv.comcast.net, as well as mobile applications like XfinityTV IOS and Android apps to power features like streaming premium content and remote DVR control, among others. With Riak KV, Comcast Cable was able to lay the groundwork for future projects similar in nature across the enterprise, enabling fast time-to-market of reliable, maintainable services.

According to Jonathan, "On a day-to-day basis, other than figuring out how much capacity we want to add, we don't have a dedicated Riak KV

administrator— because we don't need one. Riak KV provides a lot of operational simplicity for us. Largely, it 'just works,' and has been doing so since we put production traffic against it in November 2010; with no major outages whatsoever. We're about to go through adding capacity to the cluster, and indeed it seems the long pole in that tent is acquiring hardware, and not getting it added to the cluster, which is awesome." According to Jonathan, by using Riak KV, Comcast developers can get a new dataset started or modify schemas on the fly, without having to coordinate with a separate DBA group. Comcast Cable discovered that it was so easy to add new datasets that they experienced an explosion of data and requests, requiring a renewed focus on capacity management and data governance.

ABOUT COMCAST CORPORATION

Comcast Corporation (Nasdaq: CMCSA, CMCSK) (www.comcast.com) is one of the world's leading media, entertainment, and communications companies. Comcast is principally involved in the operation of cable systems through Comcast Cable and in the development, production, and distribution of entertainment, news, sports, and other content for global audiences through NBCUniversal. Comcast Cable is one of the nation's largest video, high-speed Internet, and phone providers to residential and business customers. Comcast is the majority owner and manager of NBCUniversal, which owns and operates entertainment and news cable networks, the NBC and Telemundo broadcast networks, local television station groups, television production operations, a major motion picture company, and theme parks.

ABOUT BASHO

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.