

Case Study

Yammer Powers User Notifications with Riak



Company:
Yammer

Data Issue:
Yammer serves millions of users; all making several status updates throughout the day.

Challenge
Storing and accessing the large volume of activity data in the system to push relevant status updates and other useful event data to users in a timely manner.

Solution
Leveraging Riak as a distributed key-value store enabled Yammer to power an eventually consistent and highly available Notifications module to enhance its users' experience with the enterprise collaboration tool.

A New Feature Prompts a New Data Store

Yammer is an enterprise social network, delivered as a web application via a software-as-a-service (SaaS) model that allows greater horizontal communication between individuals inside a company. Yammer has similar attributes to Facebook, and instant messenger protocols – which means it has a lot of data to manage from its rapidly growing user base. And, because the nature of Yammer is social connections, messages, etc. – the data is constantly being updated.

Recently, Yammer looked to build a new Notifications feature, which allows Yammer users to receive alerts when another user in their network follows them, they are invited into a new group, or if they received a reply to a message they published in the network. “This gives the user a sorted set of notifications, with a typical call to action based on the nature of the notification,” explains Coda Hale, Infrastructure Architect at Yammer. “It is a high value data stream, pushing information users care about into an easily actionable to-do list.”

Evaluating with Dynamo in Mind

Most of Yammer's data is stored in a Postgres data store, but as the Notifications project team was scoping out the project, they discovered that the data powering Notifications was amenable to an eventually consistent key-value data store like the one found in the Riak. platform “Postgres provides consistency of data guarantees at the expense of availability – but we knew we wanted a system that had greater availability by nature,” Hale notes.

So, with availability being an important factor, the project team began looking at projects based on the Amazon Dynamo model. As Hale notes: “The Amazon Dynamo model has probably the best availability story out of all the data stores out there, period.” So, the team looked at various Dynamo-inspired projects and eventually came to Riak and Voldemort.

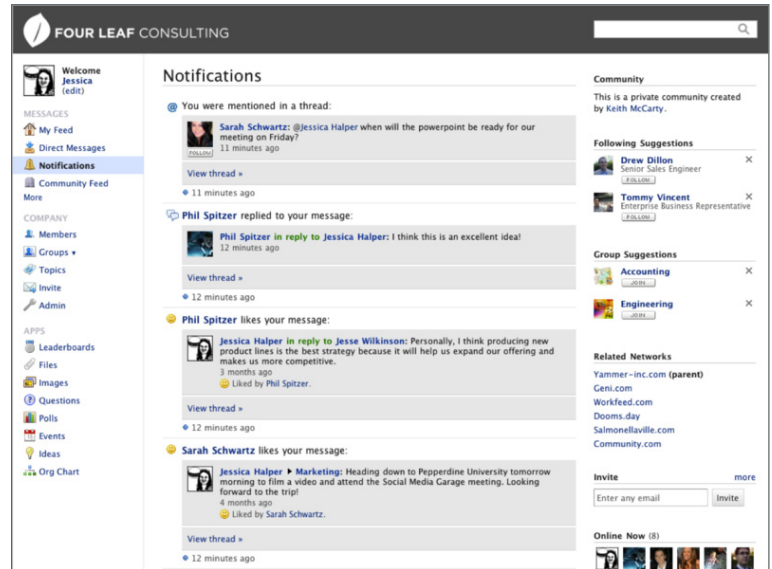
Hale states that his team chose Riak for a number of reasons, the first being the fact that Riak had the most mature footprint. “The operational amenities were really nice,” Hale states. “The architecture and implementation story was well documented, and the tools to do things like open a new node were pretty simple and open.” In addition, Riak won out because of its more mature operational tools and packaging, bigger community, and strong commercial support as compared to the Voldemort project.

Of course, Hale notes that Yammer was especially impressed with the stability and very low latency of Riak in testing. “The latency profile was flat as a pancake,” Hale says. “This made issues like capacity planning and SLAs around the project a lot easier.”

Dealing with Real World Problems

A major benefit to using Riak according to Hale is that the product understands both the needs associated with building distributed systems, as well as the nature of operating distributed systems. “Some traditional database supporters will tell you that non-trivial partitions don’t happen, but they do and having a strong non-relational tool that can provide availability is great,” Hale notes.

But while Hale and his team liked the fact that Riak was stable, and was ideal for managing the distributed nature of Yammer’s data model, he notes that not being familiar with Erlang, the programming language of Riak, was a bit off putting. “I’d heard that Erlang could be difficult to deal with,” Hale says, “but I had our production cluster set up in about fifteen minutes. It has been a complete and total non-issue for us.”



The Yammer Notification module powered by Riak

The Benefit of “Boring”

Perhaps the biggest benefits of Riak in Hale’s mind was the fact that the system simply does what it is designed to do, and its packaging makes it easy to set up a “bulletproof” system. “Not having interoperability issues between the components of the systems was a highly refreshing experience,” Hale notes. “Instead of a solid two days of starting at strange code, we were up and running quickly and have yet to have an alert or any problem.”

As Hale puts it – “It’s a piece of plumbing; it has never been a root cause of any of our problems.” So, with a bulletproof Riak-backed system in place, Yammer now has a robust Notifications module in its social collaboration tool that did not increase its data footprint on its single point of failure, is very low latency and the data powering the notifications highly available thanks to Riak. And given the amount of data moving around one of the most popular social networking tools on the market, not having to worry about a major component of its infrastructure is a huge benefit

About Yammer

Yammer’s founders David Sacks and Adam Pisoni saw an opportunity to apply the social media revolution pioneered by Facebook and Twitter to the workplace. The company launched to the public in September 2008 at the TechCrunch50 Conference and won the grand prize despite strong competition from other great startups. Just two years later, Yammer is used by over 100,000 companies and organizations, including over 80 percent of the Fortune 500.

Yammer is headquartered in San Francisco and is well-funded by top tier investors including Charles River Ventures, Founders Fund,, Emergence Capital, US Venture Partners, SV Angel’s Ron Conway and Goldcrest Investments.

About Basho

Basho Technologies, Inc., founded in January 2008 by a core group of software architects, engineers, and executive leadership from Akamai Technologies, Inc. (Nasdaq:AKAM - News), is headquartered in Cambridge, Massachusetts. Basho produces Riak, a distributed data store that combines extreme fault tolerance, rapid scalability, and ease of use. Designed from the ground up to work with applications that run on the Internet and mobile networks, Riak is particularly well-suited for users of cloud infrastructure such as Amazon’s AWS and Joyent’s Smart platform and is available in both an open source and a paid commercial version. Current customers of Riak include Comcast Corporation, MIG-CAN, and Mochi Media

