



Bringing Reactive Applications to the Java Virtual Machine

Gawker leverages Typesafe for extreme scale

Recently, Gawker Media made the move to consolidate their fragmented infrastructure under a JVM based umbrella; the Typesafe Platform was a natural choice for this company with extreme scalability requirements that experiences over 80 million unique page visitors monthly.

About Gawker Media

Gawker Media is the publisher of some of the best-loved titles on the web including the eponymous Gawker and gadget sensation Gizmodo. Founded in 2002 by Nick Denton, the influential media group now produces eight original brands with a collective audience of tens of millions of US readers. Attracting fans and critics alike for their inimitable delivery of news, scandal, and entertainment, the Gawker Media properties are heralded as everything from 'deliciously wicked' to 'the biggest blog in the world.'

The Challenge

After 7 years of PHP development, the development team at Gawker Media started to bump into issues in maintaining a large 300,000 line PHP codebase. Additionally difficulty with managing two separate production environments (LAMP and JVM) was making development inflexible and inefficient.

Gawker maintained a development team split between New York and Budapest, with each team responsible for one technology. Besides the usual problems encountered with distributed teams, Gawker had the additional side effect that because the current systems were built on a mix of Java and PHP, work couldn't be shared easily across those teams due to lack of expertise.

It was clear that moving to a common platform would help in concentrating development efforts on building the application instead of dealing with a multitude of environments, deployment processes, etc. Furthermore, this new platform should give them a solid base with which to build their next generation platform on so that they could accommodate their aggressive growth requirements.

Having defined the need, Gawker set upon the arduous task of evaluating the technology marketplace. Since they were already familiar with deploying and managing JVM-based solutions, they wanted to leverage that experience moving forward, and it was also important that any technology that they adopt

would interoperate well with existing Java components and libraries as the migration to the new platform needed to be a gradual one.

The Team

Tom Plunkett, Gawker's CTO, lead the search for the new platform. He and his team looked at a number of technologies, including Node.js, Python and Django and had even built some smaller projects in them to gauge the level of difficulty in developing and managing them in the real world. What became clear, though was that these technologies, while good, were not complete enough to meet all of Gawker's requirements, and a more flexible and scalable solution was needed.

The Solution

Tom had long been interested in the notion of Functional Programming languages due to their productivity and statically typed nature and he soon found Scala, the programming language created by Martin Odersky. Scala's strong functional programming features, strong multithreading support and the interoperability with Gawker's current Java codebase - as well as any outside Java library - was a lot more compelling than any competing technology could bring.

Another factor influencing Tom's decision was the caliber of companies that had already adopted Scala and other components of the Typesafe Platform, such as Twitter, Foursquare and LinkedIn. These internet-powerhouses had scaling requirements very similar to those faced by Gawker so he felt that he was in good company adopting Scala.

The Deliverable

Gawker had code running *in production* only four months after the first project was started, which includes the time it took for the team of PHP developers to learn the basics of Scala.

Today Gawker has had about 65,000 lines of Scala code in production for a little over a year. The now 20 person distributed development team is working concurrently on two different projects:

- Building out the new system and platform.
- Maintaining the existing Scala-based system (essentially an API with a JavaScript heavy Web Application).

Tom cites one of the greatest benefits enjoyed today is that the new system is highly modularized, where the requirements drive where the code lives, rather than the limitations imposed due to choice of language.

Furthermore, it is easier to move developers between projects, thanks to the common platform. The slightly higher development time compared to a dynamic language like PHP evens out in the long-term, and shows advantages for Scala, thanks to the cleaner code base and the better tools available. New developers can become productive with the new system a lot faster than with the old one.

As their requirements have evolved over time, Gawker has adopted the additional components from the Typesafe Platform, and now leverages Akka - a toolkit and runtime for building highly concurrent, distributed, and fault tolerant event-driven applications, Slick - a library that provides seamless data access

and Play - a framework based on a lightweight, stateless, web-friendly architecture that features predictable and minimal resource consumption for highly-scalable applications.

Akka is being used to populate search and other related indexes while Play is the framework for all of their web projects, as it proved to be flexible enough to use the components they need, and replace those they don't, while providing an easy deployment process. As of June 2013, all of the Gawker services were moved to Play, Akka, Slick and Scala.

Conclusion

The Typesafe Platform aids Gawker in maintaining a very aggressive development and deployment cycle, all in a fail-forward manner and now it is the default choice for all of their new services.

*The **Typesafe Platform** is a modern software platform that makes it easy for developers to build scalable software applications. It combines **Play Framework**, **Akka** runtime, the **Scala** programming language, and robust developer tools in a simple package that integrates seamlessly with existing Java infrastructure. Commercial support and maintenance is available for the Typesafe Platform through the **Typesafe Subscription**.*