



CASE STUDY

ZULILY

zulily is a publicly traded, Seattle-based online retailer. The company launched in 2010 with a mission of bringing its customers (primarily moms) fresh and interesting product every day. The company has over four and a half million active customers and expects to do over \$1B in sales for 2014.

Unlike search-based e-commerce sites whose users come looking for specific items, the zulily customer visits the company's web site and mobile app on a daily basis to browse and discover products. The company focuses on crafting engaging, unique and compelling experiences for its customer.

The zulily experience promises "something special every day," and creating that experience is no small feat. To do so, the company's 500 merchandising professionals create over 100 daily "events" that launch more than 9,000 unique products each day.

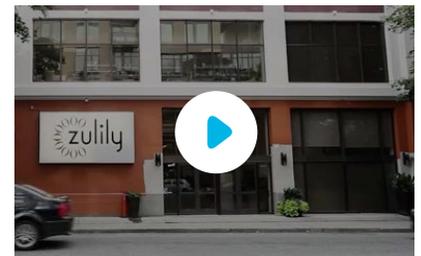
CHALLENGE

To turn this raw content into an engaging experience for customers, zulily has invested heavily in personalization. Showing the right sales event and images to the right member at the right time is critical for zulily. A mother with a 6-year old daughter should have a completely different experience from that of an expecting mother.

Accomplishing this level of personalization required that zulily build systems to understand members coming to its web site and then instantly determine what to show them. To do this, zulily's systems must capture, integrate and analyze many different inputs from a wide variety of sources.

KEY HIGHLIGHTS

- Personalized customer interactions
- From data vision to reality in four months
- Accelerated decision making from weeks to minutes



WATCH: [zulily Does Hadoop](#)

The company was founded with a data platform built on a relational database, but two years after launch the number of events, SKUs, customers and interactions were growing too rapidly for that system to keep up. To continue delivering relevant, customized content to its rapidly growing customer base, zulily knew that it needed to modernize its platform.

For personalization at scale, zulily built a Hadoop-based system for collecting clickstream data from across the company's web, mobile and email engagement channels. This system allowed the company to turn clickstream data into engines that produce personalized, targeted and relevant product recommendations.

The zulily platform helped it achieve a new level of precision and maturity in its ability to personalize its customers' experiences, but one challenge still remained. It still had a silo of structured data in one place (including transactions, customer records and product details), which was separate from its clickstream data in Hadoop.

As Hasbe and his team further defined the company's requirements, they formed a vision for what the company now calls the zulily Data Platform (ZDP). They planned to make ZDP a primary, central repository for all of the business' data. It would:

- Support the company's efforts to enhance the customer experience through better personalization and targeting,
- Give the company's business analysts easy access to all of the company's information,
- Allow the team to make smarter business decisions without needing IT support, and
- Scale to support the company's growth over the long term.

To meet these goals, Hasbe and his team created a modern data architecture that combined the strengths of both Apache™ Hadoop® and cloud computing to deliver a highly scalable unified data platform for structured and unstructured data.

“ This platform is allowing us to take questions that we couldn't answer to the point where we can answer them. ”

Luke Friang, Chief Information Officer, Zulily

“We really struggled to integrate and analyze the data across the two different silos,” said Sudhir Hasbe, director of software engineering for data services at zulily. “We often found ourselves making decisions based exclusively on a single type of data, or needing to get developers involved to produce new reports or analyses.”

The need to constantly involve the company's development team was expensive and time consuming, and it distracted the developers from focusing on their own priorities. Due to the complexity of its siloed data platform, the company found itself limited in its ability to agilely respond to changes in the marketplace or company strategy.

SOLUTION

After zulily thoroughly examined its analytical priorities and the challenges posed by its current infrastructure, the company concluded that it needed to move beyond its legacy relational database.

“We knew we couldn't build what we needed on traditional relational database technology,” said Hasbe. “We would have to innovate and take advantage of the latest technologies available in the industry for managing both structured and unstructured data.”

ZDP is based on:

- Hortonworks Data Platform (HDP®). With YARN as its architectural center, HDP provides a data platform for multi-workload data processing across an array of processing methods – from batch through interactive to real-time, supported by key capabilities required of an enterprise data platform – spanning governance, security and operations.
- Google Cloud Platform (GCP), Google's public cloud infrastructure-as-a-service (IaaS) offering.
- Google BigQuery, a cloud-based tool for super-fast data query across massive data sets.
- Tableau, a visualization and reporting tool suite.

After identifying its path forward, the zulily team was able to move from its vision to an in-production data platform in a mere four months. It will migrate all existing data processing to the new platform by the end of 2014.

“Our new platform enables analytics scenarios that were difficult to achieve with our former technology stack,” said Hasbe. “And we now have the ability to scale both storage and analytics on demand. To finally be ahead of the company's growth curve is exciting for us.”

RESULTS

For Luke Friang, zulily's chief information officer, the ZDP platform creates important new opportunities for the company.

"Data is everything to us, yet we really struggled with how to properly consume and harvest a mass of data to provide our customers with a great experience," said Friang. "Our new platform empowers us to use data all over the business. It drives the content of the email that our customers receive in the morning. It drives how and when we ship customers the products they order. It drives what customers see in the mobile app versus what customers see in a browser on their computer. It's allowing us to make sure that we're tailoring customer experiences appropriately, throughout the entire lifecycle as a zulily customer."

From Friang's perspective, it all comes down to supporting the business' ability to derive new insights and make quick decisions.

"This platform is allowing us to take questions that we couldn't answer to the point where we can answer them," he said. "It is allowing us to accelerate decision making processes from weeks, days and hours to minutes, seconds and milliseconds. From an off-line analytics activity to a real-time decision-making processes embedded within a piece of software. That's the value."

"Hortonworks' depth of knowledge was invaluable to us in this process," added Friang. "The responsiveness of their team, and their ability to get things done and get issues fixed, were key to our ability to get ZDP off the ground."

About Hortonworks

Hortonworks is an industry leading innovator that creates, distributes and supports enterprise-ready open and Connected Data Platforms and Modern Data Applications that deliver actionable intelligence from all data: data-in-motion and data-at-rest. Hortonworks is focused on driving innovation in open source communities such as Apache Hadoop, Apache NiFi and Apache Spark. Along with its 1,800+ partners, Hortonworks provides the expertise, training and services that allow customers to unlock transformational value for their organizations across any line of business.

Contact

For further information visit
www.hortonworks.com

+1 408 675-0983
+1 855 8-HORTON
INTL: +44 (0) 20 3826 1405

