

Zillow Helps Customers Find Their Way Home With Machine Learning



Executive summary

Ranking among the top 30 websites in the U.S., Zillow is the leading real estate and rental marketplace. Buying a home is one of the most important lifetime decisions a person will make, and Zillow is dedicated to empowering consumers with data to make informed choices. The company uses machine learning to quickly process and analyze the enormous volume of data that keeps its website and other business operations running smoothly. Since deploying Splunk Enterprise and the Splunk Machine Learning Toolkit, Zillow has seen benefits including:

- Improving customer website experience
- Speeding time to market with access to data in real time
- Saving millions of dollars by reducing website outages

Why Splunk

Zillow's website traffic is generated by customers on mobile devices, desktops and tablets who are looking at photos, watching videos and gleaning information about home locations, proximity to schools and more. Previously, when the occasional website outage occurred, Zillow executives were concerned about customer satisfaction, risking the loss of millions of dollars in advertising revenue or failing to meet contractual agreements with companies that license Zillow data. "Before Splunk we were blind," says Jerome Ibanes, data architect at Zillow. "If the website failed at 2 a.m., I would have to painstakingly go through terabytes of logs to figure out what was wrong."

Continuing to use legacy tools was not an option: Zillow needed the right solutions to process and analyze constantly changing unstructured data. Previously, Zillow had to move terabytes of data across multiple tools to leverage machine learning and gain operational insight. Zillow initially deployed Splunk Enterprise to look at machine data for IT troubleshooting. That success led to additional uses for the Splunk platform, from forecasting website traffic and monitoring its application testing environment to learning about customers visiting the website to serve them with the best content.

Industry

- Technology

Splunk Use Cases

- IT operations
- Application delivery
- Business analytics
- Security

Challenges

- Website outages impacted customer experience
- Website outages affected advertising revenue and compromised contractual agreements with companies that license Zillow data
- Identifying root cause of incidents took up to hours
- Product developers waited for one to two days to access data

Business Impact

- Improving customer experience
- Saving millions of dollars by reducing website outages
- Reducing incident root cause from hours to minutes or seconds
- Speeding time to market by providing product developers with real-time data

Data Sources

- Web servers
- Java servers
- Python applications
- Databases
- Firewalls
- Load balancers
- Mobile devices

Splunk Products

- Splunk Enterprise
- Splunk Machine Learning Toolkit

Fast analysis reduces outages

Today, Zillow brings a range of data sources, including logs from web servers, database servers, Java servers, Python applications, firewalls and load balancers, along with mobile device data and customer metrics, into Splunk Enterprise. Using the all-in-one Splunk platform with machine learning, Ibanes and his colleagues can normalize the data and use the Machine Learning Toolkit to take advantage of an open-source library of algorithms to build customized machine learning models.

“With Splunk Enterprise and the Splunk Machine Learning Toolkit, we are now able to reduce the impact of an outage from hours to just a few minutes or seconds by identifying the root cause faster,” Ibanes says. “We can create alerts, identify outliers and escalate issues to the appropriate teams.”

Real-time data speeds time to market

According to Ibanes, Zillow is able to move quickly even though no human being can process the data volume Zillow handles daily. Access to real-time data analytics has enabled the company to ship code to production multiple times daily with minimal customer impact. Previously, developers had to wait a

“We want to know as early as we can if we ship a product that affects a large number of buyers, sellers or renters visiting the website so we can help our customers as much as we can. With Splunk, we finally have good real-time customer visibility. The richness of Splunk is that no matter what you want to do with it, it just delivers.”

Jerome Ibanes
Data Architect, Zillow

“With Splunk Enterprise and the Splunk Machine Learning Toolkit, we are now able to reduce the impact of an outage from hours to just a few minutes or seconds by identifying the root cause faster. We can create alerts, identify outliers and escalate issues to the appropriate teams.”

Jerome Ibanes
Data Architect, Zillow

day or two for the analytics team to process data so they could find out, for example, that a new release affected a subset of customers. With the Splunk Machine Learning Toolkit, developers bake their own metrics into their code, it is ingested in Splunk Enterprise, and they have their own data to analyze in minutes. By processing data in real time, developers are empowered and the company has better overall business visibility. For example, if the number of renters on the website drops, developers know they have to roll back a release.

“We want to know as early as possible if we ship a product that affects a large number of buyers, sellers or renters visiting the website so we can help our customers as much as we can,” Ibanes concludes.

“With Splunk, we finally have good real-time customer visibility. The richness of Splunk is that no matter what you want to do with it, it just delivers.”

Download Splunk for free or get started with the free cloud trial. Whether cloud, on-premises, or for large or small teams, Splunk has a deployment model that will fit your needs.



Learn more: www.splunk.com/asksales

www.splunk.com