

Splunk® at Socialize

Social Networking Firm Leverages Splunk for Big Data Analytics to Drive Revenue

socialize

“Splunk gives Socialize powerful analytics with which to leverage our data for business development, all without resorting to complex MapReduce routines. Analyses that once required weeks of coding are now done in hours.”

Chief Technology Officer
Socialize

OVERVIEW

INDUSTRY

- Social networking / Mobile developers

SPLUNK USE CASES

- Big data analytics
- Cloud monitoring/management
- Real-time operational intelligence
- Real-time metrics for customers
- Application management / diagnostics

BUSINESS IMPACT

- Revenue producing business intelligence
- Greater value for developers and publishers
- Data driven, real-time decision making
- Efficient and reliable delivery of services
- Reduced staffing needs

DATA SOURCES

- Mobile social networking applications, databases, servers, and mobile SDKs
- Mobile application advertising campaigns
- Daily API requests numbering in the millions

APPLICATIONS

- Splunk SDK for Python
- Google Maps for Splunk

The Business

Founded in 2008, Socialize (www.getsocialize.com) was created to provide developers with a platform to build applications for mobile devices. When Socialize's leadership team noticed that consumers often stopped using apps after thirty days, they attributed this declining engagement to users' inability to communicate with each other. As the CTO of Socialize noted, "It was as though people with common interests were in the same room but were unable to see or speak with one another."

Consequently, Socialize began offering tools that let developers add social networking capabilities to their mobile apps, allowing consumers to interact and share content. Socialize now also provides developers with metrics about consumer usage that enable them to refine their offerings. Over 68 million people currently use "Socialized" apps.

The San Francisco-based company soon recognized that it could leverage its rapidly growing trove of data about consumers and their interests. It offered access to these users to advertisers via real-time bidding (RTB) exchanges, allowing marketers to target consumers most interested in their products or services with advertising on mobile apps.

Challenges

Socialize confronted a host of challenges because its businesses are data intensive and rapidly growing, supporting daily API requests that number in the millions. On the developer side, the firm ingests and processes over 20 million data points daily, each data point being a user interaction with a "Socialized" app such as a click or a comment, and there are 100,000 new users every day. It must extract intelligence from this data, such as how many new users it gained over the past month. It needs to keep developers informed of the usage of their applications, presenting analytics for each app in real time. It also must help developers debug their apps, remediating, for example, slow response times. Additionally, in order to ensure desired service levels, Socialize must troubleshoot its own infrastructure, which extends into the cloud.

Socialize's RTB business also presented a challenging issue. The RTB exchanges are automated and demand a response time of only 100 milliseconds (ms.). Socialize must offer marketers an advertising campaign for a mobile app and accept the highest bid all within a tenth of a second. To make effective bidding decisions, Socialize therefore needs to analyze hundreds of gigabytes of data, with literally millions of API requests sent per day. For example, it might present 20 million consumers who have shown interest in automobiles to advertisers eager to reach this demographic. Analyzing this volume of data with a MySQL database was time consuming and error prone. Moreover, the firm required sophisticated reporting capabilities in order to provide advertisers with metrics on their campaigns, such as the number of impressions and their cost per thousand.

Enter Splunk

Seeking to expand and automate their business and operational processes, Socialize's team installed Splunk software at the end of 2011. After a fast deployment, they fed

a large volume of logs into the platform and posed questions using natural language queries, seeking to determine whether Splunk would be an effective solution for business analytics. Splunk easily extracted answers to every question from the data.

“Splunk collects and analyzes machine-generated data from IT infrastructures, but it also offers invaluable insight into user-generated data,” noted the CTO of Socialize. “What makes Splunk special is the ease and precision with which we can extract business intelligence from hundreds of gigabytes of data, then graphically display any metrics we want in dashboards.”

In response to staff queries, Socialize started to send data from sources like databases, app servers and mobile SDKs into Splunk for indexing, building charts and graphics. Because of Splunk’s ability to work within public and private clouds, Socialize can store its data in Amazon EC2, allowing the firm to economize on operational costs.

Breakthroughs

Splunk has empowered Socialize with the intelligence to meet the needs of advertisers, application developers and publishers, and to grow its business. The platform indexes 200–300 gigabytes of data every day and this is doubling every month. According to Socialize, “Splunk gives Socialize powerful analytics with which to leverage our data for business development, all without resorting to complex MapReduce routines. Analyses that once required weeks of laborious coding are now done in hours.”

Thanks to Splunk, Socialize enables publishers to capture and display all key metrics for their mobile apps in dashboards. Splunk also provides developers with event logs and machine-generated data for debugging their applications. Relying on the Splunk SDK for Python, the platform alerts developers of any issues in their applications’ performance. Similarly, Socialize’s operational staff depends on Splunk’s reporting and notifications to ensure the firm’s distributed IT infrastructure operates within prescribed thresholds. Additionally, the staff uses Splunk as a diagnostic tool to check that a new hardware or software system is working properly before putting it into production.

Building on the Splunk REST API to integrate at the application level for business intelligence, reporting and alerting, Socialize exposes Splunk data to its customers through highly customized dashboards. Specifically, Socialize presently provides several advertisers with dashboards reporting metrics on their campaigns and it intends to extend this intelligence gathering to all of its advertisers to deliver real-time visibility into every ad campaign.

Splunk is also an essential part of Socialize’s real-time bidding operations. Splunk forwards pre-computed aggregated data to a MySQL database every minute. The data is then put into an in-memory Memcached cluster that enables Socialize to respond to bidding requests within the 100 ms. time constraint of the RTB exchanges. Socialize utilizes Splunk’s analytics to determine the efficacy of campaigns and to make more informed bidding decisions. For example, it uses Google Maps for Splunk to identify the regions where users are most responsive. It then can target these regions for particular campaigns, yielding greater value and returns.

By serving as a diagnostic tool, Splunk helps both developers and Socialize’s engineers to debug applications and systems and ensure their performance. Importantly, because Splunk automates the delivery of data to developers, Socialize’s staff no longer has to manually support them, eliminating hours of tedious work routines. Herein lies Splunk’s additional value. By automating and simplifying intelligence, reporting and alerts, Splunk helps to drive down costs and allows Socialize to run its entire business with only eight employees. “Here at Socialize, it’s crazy to do anything manually because we have Splunk to do all of our data collection, analytics and retrieval for us,” concludes the CTO. “The platform gives us the visibility, analytics and insights into our business that drive our success.”

Free Download

Download [Splunk](#) for free. You’ll get a Splunk Enterprise license for 60 days and you can index up to 500 megabytes of data per day. After 60 days, or anytime before then, you can convert to a perpetual Free license or purchase an Enterprise license by contacting sales@splunk.com.