splunk> listen to your data™

CASE STUDY

Splunk at Edmunds.com

Driving to New Operational Insight and Efficiency

edmunds.com

"Splunk takes IT data and makes it relevant for non-technical business users. Our custom dashboards built in Splunk provide both the up-to-the-minute and long-term trending data business users need to make the decisions that impact revenue."

John Martin
Senior Director, Application Operations

OVERVIEW

INDUSTRY
- Web Publishing

SPLUNK USE CASES
- Operations Management and Security
- Application Troubleshooting

BUSINESS IMPACT
- Delivered dashboards to C-level executives, providing immediate visibility into key business metrics and analytics
- Achieved 80% reduction in malicious security incidents due to real time reporting and alerts
- Established automated processes for monitoring and preempting infrastructure issues across the IT stack

DATA SOURCES
- Web Application/Load balancer logs: http logs, Apache, WebLogic, F5
- System: syslog-ng, NFS infrastructure
- Firewall/Router logs: Netscreen, Cisco
- Security logs: Sourcefire (IPS), Cisco, Access Control Systems
- Structured data: Oracle RDBMs

The Business

When a company structures the walls and furniture of its award-winning headquarters to mimic the curves of the world's most exciting racetracks, you know you're dealing with a bold organization. Even before any car manufacturer created an online presence, Edmunds.com, which started publishing new and used vehicle guides in 1966, was the first Internet source serving automotive consumers. From that lone site, known as the Electronic Newsstand, Edmunds.com has anticipated and responded to its audiences with four must-see web destinations including a social networking site and the most-read site for auto enthusiasts. Having dropped its print operation in 2006, the entire business relies on a high-performing, reliable and secure IT infrastructure. For this company, the focus of a 2001 Harvard Business School case study, to keep anticipating which car consumers desire, it needs to learn everything it can from its data.

Challenges

Even with way more than its share of logging mechanisms, utilities such as log4j, and homegrown adapters, Edmunds.com couldn't correlate its log data. Not only were logs stored in numerous locations, the many groups dedicated to analyzing these logs each had their own tools, methodologies and agendas. Too much time was spent on disparate analysis by too many people who had more important things to do. And the steady stream of people who needed answers to questions, such as what's the top user agent, had to wait hours.

Enter Splunk

The 50,000 events per minute that occur on the Edmunds.com sites produce 60 to 70 gigabytes of data per day, and enter Splunk Enterprise through syslog, a custom agent for Windows event logs and a custom log4net appender for .Net data.

Availability

Through real-time alerting, daily and weekly reports and historical analysis, Edmunds.com monitors and tracks availability—the good, the bad and the ugly.

The good includes analyzing traffic trends to ensure ad revenue and identify new customer behaviors. The bad covers device failures and security concerns such as port scans and aggressive spidering. The ugly refers to events (mostly errors) that disrupt revenue streams or impact the company image. In addition, Edmunds.com uses summary indexing for statistical analysis on referrers, status, method, URI and User Agent. Combining these across web and application servers lets them understand baseline transaction types to better monitor anomalies.

Visibility

Through a distributed Splunk setup, which segregates sensitive syslog data from non-sensitive data, everyone at Edmunds.com gets the access they need. C-level executives use dashboards showing business analytics. Before Splunk, network operations poured through thousands of lines to figure out why an application didn't work or who was doing what at a particular time. Now there are search forms for

any application that narrow searches by environment, host or time. The team can also perform cross-application mapping, correlating errors between the web and app. tiers.

Security

By normalizing data from Cisco devices, Sourcefire IDSs, Netscreen firewalls and access control systems, all kinds of analysis and reporting become possible. For example, Edmunds.com correlates Sourcefire intrusion events with Cisco denies and creates reports that chart statistics such as intrusion events by device type or by intrusion detection system.

Breakthroughs

After being plagued a dozen times a week with malicious incidents that impacted performance and/or content, the network team set up Splunk alerts to monitor the number of requests coming from a single IP address based on a threshold. Using Splunk, the team immediately uncovers which virtual host and files are being targeted and then takes action. This visibility has decreased weekly attacks by approximately 80%. Splunk has also reduced the mean-time-to-resolution (MTTR) for many other revenue-impacting events.

"Splunk helps us better understand patterns and problems, both known and unknown. Having all our logs in one place and correlating them allows us to determine why something is a problem and what the cause is. With Splunk we get real answers."

John Martin
Senior Director, Application Operations

The combination of a few Splunk search commands and tagging log records with environment, tier and a normalized source name allows Edmunds.com to create weekly "Top X" error reports for the web and application tiers. It also facilitates the build process because they can easily monitor error diffs by build numbers and dates. Both practices have reduced production errors by a factor of ten.

In general, Splunk has improved productivity because people don't require an interpreter to access and make sense of log data. They get the answers they need – quickly. This ability allows the Edmunds.com team to satisfy the "what if" curiosity that keeps a company living up to the boldness, innovation and speed reflected in its surroundings.

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.

250 Brannan St, San Francisco, CA, 94107 info@splunk.com | sales@splunk.com 866-438-7758 | 415-848-8400 www.splunkbase.com