Global Distribution Company Relies on Splunk for High-Performance Web Monitoring and Real-Time IT/Business Insights

Introduction

Periodically, Enterprise Management Associates (EMA) conducts case studies on enterprise management products that demonstrate above-average customer value. Splunk (NASDAQ: SPLK) is distinctive in that multiple customers have provided impressive stories documenting the Splunk value proposition in support of application management, network management, security, and a variety of other IT- and business-related use cases.

This case study details the use of Splunk by a global components distributor. Not only is Splunk being used to manage and optimize the experience of customers placing orders on this company's ecommerce website, it is also being utilized internally for security monitoring and for detailed insights into the content of machine-generated data across the business.

In addition, Splunk is now being used in multiple departments within the Line of Business (LOB). LOB users are utilizing Splunk's indexing, searching, and reporting capabilities to address issues arising within their own operational processes that would otherwise go unsolved.

Vendor and Product Description: Splunk

"*Machine-generated data is one of the fastest growing and [most] complex areas of big data. It's also one of the most valuable, containing a definitive record of all user transactions, customer behavior, machine behavior, security threats, fraudulent activity, and more. Splunk turns machine data into valuable insights, no matter what business you're in. It's what we call 'Operational Intelligence.'*

Operational Intelligence gives you a real-time understanding of what's happening across your IT systems and technology infrastructure so you can make informed decisions."[1]

Founded in 2004, Splunk became a publicly traded company in April 2012.

Splunk Customer Profiled in This Study

This EMA ROI Case Study profiles an industrial and commercial components distributor. Interacting with customers primarily via its corporate website, the company's unique differentiator lies in a combination of a strong customer focus and an ability to supply customers with a wide variety of parts and components that are difficult to find elsewhere.

Customers often purchase in small quantities and need orders delivered "yesterday." And because of the business-critical and often time-sensitive nature of the products it sells, the company focuses heavily on accelerating the ordering process from web order to shipping. The IT organization plays a big part in this process and sees its role as extending customer satisfaction to all aspects of the applications and services it delivers.

Interviewees

Due to internal corporate policy, neither the company name nor the interviewee names can be publicly cited. However, in July 2015, EMA analysts personally interviewed two IT spokespersons with the following titles and roles:

[1] Downloaded from http://www.splunk.com/content/splunkcom/en_us/resources/operational-intelligence.html on August 4, 2015.

IT Solutions Architect – At this company, the Solutions Architect sets direction for IT projects and initiatives. The company relies on his broad company and technological context to help develop a big-picture vision for present and future IT direction. He also oversees and reviews the work of several project teams working within his area of expertise.

IT Solutions Manager – The Solutions Manager assigns concrete objectives to the Solution Architect's vision and authors achievable milestones to be carried out by a project team. She also manages the day-to-day project work and responsibilities of the Splunk implementation team.

Business Scenario

Because fast delivery is so critical to customers, as one interviewee says, "We get those materials to our customers as fast as possible, usually the next day. So we aim to pack orders placed in the morning for pick-up in late afternoon or early evening."[2]

Each step in the delivery chain—from online ordering to shipping—is equally important to ensuring that customers receive orders in a timely manner. As one of the interviewees reported, "We have a close relationship with our warehouse operations. Obviously, if we're trying to deliver items the next day, we've got to get them out the door very quickly. Once we take the order, it must be located, pulled, boxed, and on a truck in less than an hour if it's going to be shipped the same day. And the entire process has to integrate with our financial and business operations."

The interviewee goes on to say, "All of our applications are homegrown, so we have a fairly complex IT environment and a large development staff. We have a lot of custom code, which can be a double-edged sword. Custom code gives us the flexibility to customize systems the way we want, but it's also more difficult because we have 100% of the responsibility for supporting the applications and for finding and fixing any issues on our own. We don't have the scale of, say, an SAP system, which so many companies are using and which is tested and supported by the vendor."

The Splunk purchase was aimed at simplifying support for mission-critical applications running across diverse hardware and software elements, many of which write to log files to communicate back to human administrators.

Acquisition

The support team was finding that troubleshooting execution errors of any kind took far too long. On average, three to four hours per incident were being spent to to comb through log files across many servers to find the "right" error code.

And while the troubleshooting process itself was expensive, time consuming, and often required collaboration across multiple personnel and teams, a bigger challenge was the impact of website issues on customers. During the three or four hours required to troubleshoot and fix a failure or an outage, customers had a poor experience with site performance. These web performance issues inconvenienced customers, and this was the deciding factor in the decision to purchase a monitoring and analytics solution.

Before deciding on Splunk, the team conducted a bake-off between Splunk and a leading systems-focused solution that had some log analysis capabilities. In the end they chose Splunk for its power in consolidating and analyzing machine data from a wide variety of sources—and for the speed at which queries and reports could be produced.

Initial acquisition and deployment costs totaled approximately $144,000.

[2] All quotes cited are from one of the two interviews conducted for this study.

Deployment and Operation

Initial deployment and log collection were reportedly not difficult to set up. Team members did struggle in the beginning with learning the Splunk query language (based on "regular expressions"), a query methodology that is totally different from the more familiar Structured Query Language (SQL) used by most programmers. However, one "Splunk ninja" in the company took ownership, learned the product "forwards and backwards," and wrote a training guide for use by other employees. Today, the company has approximately 50 users and a dozen developers who are "pretty darn good" at creating Splunk Apps and custom queries and reports.

The team also reports that most of the deployment-related challenges were associated with the customization needed to support their use cases. They created custom dashboards, reports, and alerts rather than relying solely on those delivered by Splunk out of the box. They also discovered that Splunk can be applied to a wide variety of use cases and extended to meet the monitoring and analytics-related needs of multiple stakeholders.

One of the spokespersons interviewed for this report states, "There's been good uptake and a lot of people are making use of the product."

Outcomes

Splunk is delivering value in a wide variety of ways, not all of which are specific to IT. LOB employees in multiple departments have been very innovative in extending the Splunk value proposition to a diverse array of use cases:

Alerting and analytics supporting day-to-day IT administration

The IT team has been very successful in developing a set of automated alerts notifying them of server failures and slowdowns. By consolidating and analyzing production logs, the team can rapidly determine the key factors leading to website degradation. The IT team has also developed alerts that automatically fire off messages when significant issues occur.

Examples of these alerts and day-to-day management benefits include the following:

- *Alerts based on custom thresholds for individual servers:* Alerts are triggered by custom thresholds that are set based on past performance. One alert on which they rely heavily, for example, monitors web server performance. They are notified if one web server is considerably slower than the rest, and the server can then be removed from the load balancer configuration until the issue is fixed.
- *Alerts on specific groups of servers:* Another alert tells them if the entire "fleet" of web servers is functioning poorly or has slower response times. "This has been our favorite indicator of bad health. If the entire fleet is having a problem, that's probably indicative of something going wrong with a shared resource, or perhaps a bad migration has caused the problem."
- *Failover notification:* In this environment, SQL Server failover clustering ensures uptime when a server failure occurs. Alerts are fired if a SQL Server instance fails over, giving IT a heads-up so they can remove the server from the cluster and fix the failed server before it impacts operations.
- *Proactive versus reactive website monitoring:* "Before Splunk, we would really have to wait for a user to give us a call or for a customer to say, 'Hey, the website hadn't been working,' or, 'Hey, the website seems pretty slow. What's going on?' With Splunk, we know before our customer does, and we can correct that issue behind the scenes before it impacts our customers at all."
- *Real-time visibility into problems and their sources, and confirmation of "fixes:* "When a problem is happening, we know almost instantly when a corrective action is taken whether it remediates the problem or not. An awesome example of this is an outage where our CDN [Content Delivery Network] provider was caching a bad version of our homepage, which blocked people from getting to our site. Throughout the morning, we had received something like nine phone calls from

> customers. However, when we went back through the logs and counted the number of unique visitor IDs, [we saw that] 1,700 people had seen that error. When we asked the vendor to flash the cache, we could see the problem go away almost instantaneously by watching Splunk. This is an awesome example of how what we would have thought was a tiny little problem was actually a really big deal. And when it was fixed, we knew right away."

Log aggregation for centralized security monitoring

"We're using [Splunk] as a way to aggregate all of the information we get from our third-party security products like our firewalls, our antivirus, etc. We have the ambition to eventually create a unified view of our security events to help us correlate across these different views and detect security issues as they are unfolding."

"It's just very easy to be able to go to one place and use the common search language to search all these logs. A lot of times it's also faster than the tools that are built into the individual products, some of which have the ability to search through logs. You get much better performance and find information much faster by using Splunk."

Improved searching and reporting of data generated by third-party solutions

"As an example, we use <commercial solution> to track user activity on the Internet and block access to sites that people shouldn't be going to. We needed to do some research on a site which several users had accessed and which we suspected of being an Adware site. <Commercial solution> has a search function, but it is painfully slow and the interfaces are poor.

"When I remembered that the data was in Splunk, I got what I needed in minutes and it was great. And I was able to search across multiple users at once and do it in seconds as opposed to hours or maybe never getting an answer if I had used the product's built-in interface. Splunk is made to index the data in all sorts of different ways and to enable you to find the data you're looking for very quickly.

"So in this case—and it's similar with other products as well—we may already have the ability to search the logs within the product, but they don't do it as well or as fast as Splunk."

Cross-team and DevOps collaborations

Within IT, it has historically been difficult to recreate execution problems after the fact. One particular difficulty lies in the communications between Operations and Development. Operations communicates the symptom ("the server was slow and the CPU was pegged at 100%"), but Development lacks enough information to recreate and fix the issue.

This team is utilizing Splunk to take a snapshot of issues when they occur in order to share issues with other teams. One of the interviewees states, "I recently used Splunk to solve a problem and wanted to share what I did with other team members. I was able to just paste two Splunk queries in an email and wrap some text around it. People could then log on to the specific machine and run the two queries to see what I meant. It was really cool."

Warehouse operations, business-focused real-time data analysis, and Microsoft Exchange troubleshooting

As employees across the organization become aware of the power and versatility of Splunk's machine data analysis capabilities, multiple unexpected use cases are cropping up. Some examples include the following:

- Integration between custom VB.NET code and Splunk monitor for conveyor outages in the warehouse can provide better visibility into when conveyor-related outages are occurring and why. The results have included improved warehouse operations and efficiency, improved agility in fulfilling customer orders, and analysis of real-time data for business-facing reporting.

- LOB can access minute-by-minute quantification of the number of users interacting with the corporate website at a given time, quantification of the number of users actually clicking on the site, a rolling total of number of orders placed during the day, visibility to items users are searching for, and the part numbers they are clicking on.
- The IT team also uses Splunk to troubleshoot Exchange problems. In the past, such problems had resulted in undelivered mail. Using Splunk, they are able to diagnose problems on the fly, determine which email messages were not delivered, and resolve Exchange-related issues far faster than would have been possible without Splunk.

Calculations of Hard and Soft ROI

Hard ROI

Source	Before	After	Savings
Salary savings from reducing time spent on day-to-day troubleshooting	2 people, 4 hours on average per incident, twice weekly--manually searching logs for troubleshooting purposes (approximately 16 hours weekly)	10 minutes per incident (twice weekly) (20 minutes total)	**Time savings of approximately 16 hours weekly translates to $58,240 annually** 16 hours * $70 per hour loaded salary for IT tech = $1120 weekly, $58,240 annually[3]
Reduction in MTTD (Mean Time to Diagnose)	Same as above	Less time spent on troubleshooting and diagnosis means issues are fixed faster.	**24X per incident improvement in MTTD means *significantly less* adverse user impact** 20 minutes versus 480 minutes per incident.
Salary savings from reducing time spent troubleshooting “micro-outages” (basically smaller issues with unknown sources)	Average of 30 hours per week spent digging through logs to find source of issues (Average 7.5 people * 4 hours per incident)	Micro-outages identified via Splunk reporting, alerts configured. Staff now alerted by Splunk (versus by user calls), eliminating need to sift through log data to determine the source of a micro-outage.	**Time savings of 30 hours weekly, on average, translates to $109,200 annually** Average 7.5 people * 4 hours per incident = 30 hours = $2100 weekly ($109,200 annually)
Resolution of mail-related issues	Routine mail delivery failures happen approximately 2X annually, more difficult issues 4X annually. Routine issues required (average) 1.5 people * 8 hours = 12 hours annually; non-routine required 3.5 people * 8 hours = 28 hours annually. Total= 40 hours annually ($2800)	Routine issues now require: 1.5 people * 30 minutes = .75 hours annually Non-routine issues: 3.5 people * 1.5 hours = 5.25 hours annually Total = 6 hours annually ($420)	**Savings of $2380 annually** ($2800 - $420)
Hard Savings Totals			**$169,820 annual personnel cost savings reallocated to business-facing projects** **24X improvement in Mean Time To Diagnose** **Recouped initial $144,000 Splunk investment in 1 year plus additional $25,820**

[3] www.payscale.com, http://www.payscale.com/research/US/Job=Systems_Administrator/Salary, downloaded 10/7/15, calculated loaded rate of $70 per hour

Soft and/or Unquantifiable ROI

ROI Source	Before	After	Outcome
DevOps interactions: Splunk enables more precise communication of operational issues back to Development for bug fixes	Code-related app and server problems were detected by Ops, but there was no efficient way to communicate cause back to Dev so code could be fixed.	Ops sends Dev a Splunk query statement. Dev can run the statement to see exact log entry where failure occurred.	Better communication between Dev and Ops means higher-quality code. Issues are eliminated over time, closing the loop on errors.
Identification, quantification, and remediation of "micro outages"	Micro-outages were often ignored as having no significant impact. Such "glitches" slowed server response time, but weren't serious enough to actually take the server down. Unknown to the staff, however, these errors were happening "dozens of times" per day.	Using Splunk reporting, Ops team was able to quantify micro-outages in terms of numbers and performance impact. This quantification made the problems more compelling and encouraged time investments in fixing them.	• Improved code quality • Elimination of micro-outages and resulting performance "slowdowns" • Better customer experience
Ability to quantify web performance	Variable site response time, no good way to quantify	Measurable response times, typically staying under 200 milliseconds per request	Better performance visibility ensures consistently high levels of service quality for customers
Better tracking of security-related issues	Multiple security-related products, no way to consolidate data from all sources into a single view of "how secure are we?"	Using Splunk to aggregate logs from third-party security products, server logs, and other sources to get visibility to events and how they interrelate	Consolidated view of security-related metrics in context with one another makes it simpler to detect security-related anomalies and secure organizational borders
Better visibility into warehouse operations (Team has integrated VB.NET code with Splunk to monitor conveyor outages in the warehouse.)	Little visibility into frequency of conveyor outages *or* into historical sources of outages	Better visibility to *when* conveyor-related outages are occurring and *why*	• Improved warehouse operations and efficiency • Streamlined supply chain • Increased agility in fulfillment of customer orders
Business reporting (Analysis of real time data for business-facing reporting)	No real-time visibility of web metrics	Used Splunk analysis and reporting to deliver real-time insights into web usage: • Minute-by-minute quantification of number of users interacting with corporate website at any given time • Minute-by-minute quantification of number of users actually clicking on the site • Rolling total of number of orders placed during the day • Visibility of items users are searching for and the part numbers they are clicking on	• Elevated the role of IT in supporting the business • Extended value proposition of Splunk to strategic business objectives
MS Exchange troubleshooting	Exchange problems resulted in undelivered mail	Used Splunk to diagnose problem source on the fly and determine exactly which email messages were not delivered	Faster resolution of Exchange-related issues than would have been possible without Splunk

Quotes

“We’re the only team in the department that’s actually focused on day-to-day usage of Splunk. Other teams have made use of it for their own troubleshooting, when they’re doing migrations, etc. They don’t necessarily need to learn Splunk, but they have wanted to because they see how interesting and powerful the product is. I think this is truly a success in and of itself because there was no mandate. People are truly excited about the product.”

“There are maybe 35 [or] 40 people outside of our project team that have access to Splunk normally. I think we’ve received close to 10 unsolicited, ‘Splunk is awesome; this is helping me so much’ emails, just completely out of the blue. We didn’t ask, ‘Hey, what do you guys think of this?’ They just said, ‘Boom, we love this!’”

“Splunk lets us know within five minutes if things are behaving abnormally, and we know right away if it’s one server or all of them—and we’re able to get to the heart of the problem sooner.”

“With the way [our company is] growing our systems, and especially the ambitious goals we have had for our website and apps, we need a tool like Splunk to manage all the multiple underlying components. Trying to manage the entire enterprise by manual investigation is a losing game. Splunk has helped us stay on top of issues and address them before they become problems.”

About EMA

Founded in 1996, Enterprise Management Associates (EMA) is a leading industry analyst firm that provides deep insight across the full spectrum of IT and data management technologies. EMA analysts leverage a unique combination of practical experience, insight into industry best practices, and in-depth knowledge of current and planned vendor solutions to help EMA’s clients achieve their goals. Learn more about EMA research, analysis, and consulting services for enterprise line of business users, IT professionals and IT vendors at www.enterprisemanagement.com or blogs.enterprisemanagement.com. You can also follow EMA on Twitter, Facebook or LinkedIn. 3274.102315