# Leading Online Travel Company Leverages Splunk to Consolidate Data and Tools Across Its Ecommerce Group: An EMA ROI Study

## Introduction

Periodically, EMA™ authors Return on Investment (ROI) studies covering enterprise management products that demonstrate above-average customer value. Splunk is distinctive in that multiple Splunk customers have provided impressive ROI stories.

This EMA™ ROI study profiles a leading online travel company. It details the innovative ways this organization has utilized Splunk to consolidate its ecommerce log environment, reduce redundant and disparate tool sets and deliver a more highly optimized user experience for its online clientele. This EMA™ study documents significant hard and soft ROI benefits of the Splunk solution.

## Product Description

With over 4,800 customers since its founding in 2004, Splunk has won a considerable following among IT organizations that are often highly vocal about the distinctive flexibility Splunk provides for operational intelligence.

Splunk provides a combination of capabilities for data collection, indexing, search and analysis that give organizations substantial freedom for better understanding their operational data. The Splunk product is purely software, simply installed and readily extended. Splunk can collect machine-generated data from a wide variety of sources. Its indexing and search capabilities free organizations from much of the overhead of competing approaches that require costly and time consuming normalization and rationalization of data before it can be made useful.

The Splunk platform is essentially very simple, centered on a Splunk server, with user access enabled via a Web console. This architecture can be extended across multiple data centers and Splunk servers, with role-based access controls that facilitate tailoring of reports and analysis to individual users and restricting access to sensitive information when required. Extensibility is further supported by Splunk forwarders, essentially lightweight software agents, which broaden the range of data the Splunk platform can collect and transmit to Splunk servers.

Splunk's flexibility and ease of deployment has resulted in an annual growth rate in the high double-digits according to the company. Splunk is currently in use by over 50 of the Fortune 100. In the second quarter of 2011, Splunk was granted U.S. Patent No. 7,937,344 for organizing and understanding machine data through the use of a "machine data web." The company has recently augmented its offerings with a hosted version of the technology, Splunk Storm. Splunk Storm offers Splunk as an elastic, multi-tenant service, able to monitor both Cloud-based and on-premises environments as well as leveraging the power and extensibility of cloud computing for data analysis.

## Interviewee

Senior Director of Infrastructure, Architecture and Emerging Technologies

## Company

This online travel company helps millions of business and leisure travelers every year to plan their air travel, make hotel bookings, reserve their vacation packages and more. In addition to a number of robust branded sites in North America, the organization also has numerous sites around the world that feature localized offers.

## Problem Scenario

The online travel industry is highly competitive and operates in a real-time web-based business environment. Serving customers in a targeted and on-demand fashion puts severe demands on IT infrastructures stretching their ability to serve the needs of travel consumers. Supporting this business model is expensive and complex, requiring a well-orchestrated chain of applications, processes and systems that enhance the customer experience while optimizing revenue potential.

Through acquisition and organic growth the IT landscape at this particular online travel company has become widely dispersed (silo oriented) and highly complex. As the organization has grown, so has the number of tools used to gain insight into system performance. In the logging space alone, the company was supporting 20 different solutions that ranged from recognized industry products to in-house developed tools and scripts. Most of these solutions were unable to communicate or share data between themselves, making it nearly impossible for the IT team to follow critical chain reactions from one environment to another – which is entirely common in a distributed, scaled out infrastructure.

Solving critical process issues and troubleshooting across 20 different tools was too slow and it ultimately became virtually impossible. In the end, supporting these tools also added extra expense for the company and negatively impacted customer service. The company had two goals with its first Splunk project: eliminate these tools and the costs associated with them; create a common consolidated platform that the entire team could leverage to identify and diagnose system and process failures. The ability to execute a quick Root Cause Analysis (RCA) on any incident was critical to the team.

The company's website properties take an advanced approach to serving the requests of customers. Web pages consist of multiple components; each can be customized for an individual user experience. Algorithms run behind the websites to determine the order in which hotels are shown based on past behavior patterns, page referrals, search keywords and user preferences. Ads within the pages are also optimized for each customer interaction based on similar algorithms.

The content for each of these components comes from multiple applications, and sources in and outside of the company. When a process fails, the site is forced to deliver un-optimized content or, under extreme circumstances, no content at all. Each of these failures is extremely complex and can rely on five or more execution points within the company's technology stack. Reducing these failures requires a clear understanding of the cascading events across all contributing systems; each event requires manpower and time to resolve.

## Acquisition Story

In December 2011, the company began to consolidate systems. At that time they were monitoring 54 separate database instances, all running SQL Server, and each dedicated to logging events from throughout their systems. The IT group drafted a proof of concept document that focused on tool elimination, reduced failure rates, quicker diagnoses and infrastructure savings. The business case projected savings of $7 million dollars over the first two years of the project.

The company downloaded a free copy of Splunk at http://www.splunk.com/download and focused on delivering a pilot project. The Emerging Technologies team approached several business units who

were struggling with older technology that was at its limits. The groups each had a specific pain point around time-consuming RCA discovery and the resulting impact on outages. The IT group did an initial pilot program that covered monitoring 1,000 servers and was able to stand the project up within 30 days.

The team realized immediate success; with all of the system information in Splunk, event discovery became nearly immediate. The team was able to dramatically improve their RCA process and deploy hot fixes at a much quicker rate. During the initial pilot, Splunk helped reduce the company's response time on a significant failure to less than an hour from nearly four hours. The immediate success of the pilot caused word of mouth to spread to other business units, creating a viral response to Splunk and a high demand for the solution. With this, the Emerging Technologies team moved to broaden the deployment across IT and the business.

## Outcomes

The company's initial success with Splunk motivated them to leverage the solution widely across many of their systems. During the first six months of deployment, Splunk had reached just over 4,000 servers and by the end of the year, Splunk was connected to over 11,200 hosts within the 12,000 device total environment. Today Splunk is indexing and monitoring 869 different source types and 227,000 data sources for RCA, performance and web and business analytics. At present time, the Splunk platform has analyzed over 139 billion events.

The adoption of Splunk has enabled the online travel company to report significant ROI from various sources. (See Table 1) Having a centralized tool that serves the nearly 3,000 employees who need access to this information in addition to a consolidated platform has enabled the team to author over 600 alerts and 50+ Splunk applications to better manage and streamline the ecommerce environment.

Splunk allows the team to monitor their infrastructure proactively and spot conditions that lead to failures in real time instead of after the fact. Response times for each RCA has dropped from 4 hours per event to less than an hour for each. Based on a 100K FTE rate, this reduction in time is saving the company $158.63 per incident – with an average of 15 incidents per week the company realizes an annual cost savings of $119,831. The reduced rate of incidence has also allowed the team to focus on innovation instead of system operations. Because of the tremendous focus the company has applied to this solution, incident rates have dropped by 90%.

The consolidation of all of this data using Splunk allowed the company to decommission nearly 200 servers last year; as a result, they needed to purchase 100 fewer this year than originally planned. The elimination of these servers added to the project ROI on many fronts:

1. Elimination of Capital Expense to purchase new equipment
2. Data center personnel, management and energy cost reduction
3. Software license elimination. (Microsoft SQL Server, Operating Systems)
4. Elimination of tool licenses after consolidation
5. Reduction of SAN storage costs

According to the company, the infrastructure savings have totaled 2.75 million dollars per quarter since the launch of the project and could total nearly $22 million over the course of two years.

The company has indicated that the risk aversion value supplied by Splunk alone delivers ROI that far eclipses the financial investment they have made in the Splunk platform.

The team has also identified several areas of soft ROI connected to the project. In support of their marketing unit, they have built a dashboard application in Splunk that details the top-performing ad referral links. The SEO/SEM Analytics team combines the link data with additional system data such

as transaction IDs, click traffic, click stream conversions and outside data to refine their ad campaigns. They use Omniture and TeaLeaf to analyze web traffic but had been unable to combine this string of information and processes together using either of these platforms. The team now uses a Splunk dashboard to optimize their online ad spending.

The company achieved an 82% reduction in Mean Time to Repair (MTTR) and a 90% success rate in root cause identification (versus an industry standard of 70-80%). As the team continues to author alerts and build on the platform, they anticipate the percentage of execution will continue to climb.

Splunk has also enhanced the company's ability to deliver its customers targeted content and offers that promote customer service and help generate additional revenues for the company.

## Hard and Soft ROI Tables

| Hard ROI | Before | After | Savings |
|---|---|---|---|
| Troubleshooting time per transaction | 4 hours, 15 per week= 60 hours per week<br><br>Calculation:<br>(60 (hrs used) * 52 * $48)*<br><br>= $149,760 per year | 40 minutes, 15 per week= 10 hour per week<br><br>Calculation:<br>(10 (hrs used) * 52 * $48)*<br><br>= $24,960 per year | **$124,800+ (annually)** |
| Infrastructure Savings and Tools Consolidation | | Elimination of 300+ servers<br><br>Capital Expense<br><br>Data center personnel,<br><br>Data Center mgt and energy cost reduction<br><br>Software license Microsoft SQL Server, Operating Systems<br><br>Tool licenses after consolidation<br><br>Reduction of SAN storage costs | **$11,000,000 (annually)** |
| Major Shutdown Prevention | $3 Million per incident | Based on Bit attack incident | **$3,000,000** |
| | | **Total Quantifiable Hard ROI** | **$14,124,800 (annually)** |
| **Soft ROI** | **Before** | **After** | |
| Mean Time to Repair (MTTR) | Almost no ability to quickly and accurately address incidents | 82% reduction | |
| Root cause identification | Almost no ability to quickly and accurately address incidents | 90% of incident root causes identified versus industry average of 70-80%** | |
| Marketing SEO/SEM | Unable to integrate data and systems to supply business with required analytics | Re-directing ad spend to top performing referral sources and ads (optimization for their estimated multi-million dollar ad spending) | |
| Improved Customer Satisfaction | Variable outages and untraceable system failures. | 1. Customers receive the best possible offers and opportunities when interacting with the company's services<br><br>2. The real-time performance of the Splunk platform combined with alerts enables the company to act quickly to stop service failures | |

Table 1: Hard and Soft ROI

* Based on industry average
** Based on EMA Research

## Quotes and Observations

The implementation at this online travel company continues to grow beyond just log data. The Emerging Technologies team is now investigating ways to gain better insight into more traditional data-driven systems and is continuing to build Splunk Apps and alerts to support its growing community of internal users.

**Quotes:**

*"The ROI we have achieved with Splunk has exceeded 25X our original investment"*

*"We've saved tens of millions of dollars in outages from a revenue perspective over the last six months just doing avoidance rate"*

*"…the first people saw really quick value and that translated into getting more adoption much, much quicker via word of mouth."*