Bronto Software Buys Splunk vs. Extending their Hadoop Deployment: An EMA ROI Story

Introduction

Periodically, ENTERPRISE MANAGEMENT ASSOCIATES® (EMA™) analysts author Return On Investment (ROI) studies covering enterprise products that deliver above-average customer value. Splunk Inc. is distinctive in that multiple Splunk customers have provided impressive ROI stories. This EMA ROI study profiles Bronto Software, an award winning marketing solution and services company. It details how Bronto Software utilized Splunk® Enterprise to gain visibility across a Big Data oriented system environment and integrated Splunk with its open source centric infrastructure to save money and enhance IT service.

HIGHLIGHTS

splunk>

Vendor name: Splunk Inc.

Product area: Operational Intelligence

Product name: Splunk® Enterprise™

Product version: 4.3

Bronto

Customer name: Bronto Software

Customer domain: Marketing Solutions for Commerce Focused Companies

Splunk

Background

Splunk was founded in 2004 (http://www.splunk.com NASDAQ: SPLK) and shipped its first software in 2006. Today the company serves over 4,000 customers in 80+ countries. Splunk Enterprise is used by IT professionals to power operational intelligence across application management, application development, IT operations management, security, compliance, Web intelligence and business analytics. At its core, Splunk Enterprise is designed to access machine data and make it useful and valuable. 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 the majority of the Fortune 100. In 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." Its capabilities for data collection, indexing, search and analysis provides a central point of monitoring and visibility into distributed systems, applications and data sources throughout the enterprise.

Product Description

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. Splunk can collect and index data from a vast array of source systems including network traffic, application servers, hypervisors, GPS systems, stock market feeds and structured databases. Extensibility is further supported by Splunk Forwarders, lightweight software agents which broaden the range of data the Splunk platform can collect and securely transmit to Splunk servers.

Splunk was built for big data and scales to collect and index tens of terabytes per day. Many customers keep tens-to-hundreds of terabytes on disk for historical analysis. Splunk's architecture is based on

MapReduce, so as daily volumes and data sources expand, you can scale performance horizontally with the addition of commodity servers. Out-of-the-box reporting and analysis capabilities avoid the need to deploy third-party reporting tools.

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.

Bronto Software

Background

Bronto Software was founded in 2002 and is a privately held provider of cross-channel marketing solutions for commerce-focused companies. Bronto has 130+ employees and is headquartered in Durham, North Carolina. The company serves over 1,000 clients worldwide with its SaaS solutions designed to address campaign management, automated lifecycle messaging, shopping cart abandonment, advanced segmentation, transactional messaging and in-store experience management. Bronto delivers 12 billion emails annually on behalf of its clients and has won multiple industry awards for its products and services.

Technical Team

For this ROI study, EMA interviewed the Manager of Systems Engineering at Bronto Software. This individual helped build the systems engineering group at Bronto and presently manages three full-time systems engineers along with supplemental consulting resources. The group is responsible for monitoring and maintaining multiple systems in support of Bronto's enterprise applications.

Splunk at Bronto Software

Bronto's IT Environment

Like most companies, Bronto Software manages a growing and complex IT environment. In Bronto's case this environment was founded on open source software and its ability to leverage it to support innovation. Along the way, Bronto has benefited greatly from the financial models surrounding open source software. Open source software is foundational at Bronto and has become the company standard within its IT environment. While Bronto's dedication to open source has fostered an atmosphere of do-it-yourself pride, it has also introduced new challenges. When seeking to augment existing technologies, Bronto is sometimes faced with an immature tools market that often fails to keep pace with their unique business problems or provide the analytic capabilities needed to respond to critical processes and services.

One important business function of Bronto Software is the delivery of client emails in support of their marketing campaigns. PostFix, Bronto's open source email platform, sends over 40 million emails a day and produces in excess of 100,000 mail server events per day. This event data contributes to an already complex landscape that requires the aggregation of system level and application level information such as access data, security information and failure/bug data. As this data grew in size and velocity, Bronto determined streaming logs into their MySQL database was incapable of supporting the analytic and performance needs of the group. To accomplish the increased workload Bronto turned to big data solutions. After some experimentation, Bronto settled on two 25-node Hadoop clusters – one designed to leverage HBase for real-time analytics, the other as a compute cluster to execute MapReduce functions.

The Hadoop environment now consists of 35–40 terabytes of data along with a 30–35 terabyte implementation of MySQL storing relational information. As a result of company growth, this environment is expected to reach 100+ terabytes by the end of 2012. With critical data distributed across multiple technology platforms, visibility is nearly impossible across this very complex environment.

The Challenge

Bronto required a solution that would allow engineers to deep dive into system information and events quickly and easily. Insight from this information would allow them to drive better application performance and serve customers more effectively. The solution would need to access data from all supporting platforms at massive scale, provide ad hoc search capabilities, real-time monitoring and alerting and dashboard views for a variety of users. After researching the open source software market, Bronto was unable to locate a tool that fit this profile. Calling on their do-it-yourself heritage, the team investigated the cost of building a big data solution on their own on the Hadoop platform and quickly concluded that the project would require a full-time application developer, extensive resources and ongoing maintenance that would cost in excess of $150,000. Time to market was also a concern; they estimated it would take nearly six months to build this with Hadoop and other open source big data solutions. The cost and time to build were significant enough to have Bronto consider purchasing their first commercial enterprise software solution.

Bronto identified their mission-critical, SaaS email delivery platform as the most important system on which to gain visibility. Prior to Splunk, system engineers were using manual processes to examine log data from the mailing systems to investigate delivery denials. The application is deployed across 19 servers, each hosting 20–30 virtualized instances of the software. These 300+ instances created nearly 25 gigabytes of mail delivery data each day. Each email sent created six new rows of log data, comprising nearly 180 million lines of log data each day. Troubleshooting across all of these systems was a complex challenge, as engineers were forced to manually access these logs, decompress them and use methods to search the text and correlate the results. Bronto estimates that they dedicated a full-time engineering resource to this task each week, effectively costing the company $200,000+ per year in lost manpower.

Bronto implemented the Enterprise trial version of Splunk and within an hour was able to simulate their use case and gain visibility across multiple datasets. As a result, the Manager of Systems Engineering was soon able to make a business case for purchasing the first commercial software solution at Bronto by asking his boss to ask him any question about email delivery denials and then instantly delivering the answer using Splunk. Sales engineers from Splunk helped Bronto develop dashboards that identified the top ten bounce codes by server and top ten bounce domains.

In short course, Bronto made Splunk its first non-open source enterprise solution.

Since its initial production deployment in March of 2011, Splunk has become a critical real-time monitoring and analytics tool for Bronto. They have increased the amount of data for analysis from six weeks of data to an entire year within Splunk and now use their Hadoop infrastructure as an archive mechanism for computing delivery data.

The real-time monitoring and alerting capabilities provided by Splunk have helped Bronto improve their service delivery. They have created over 20 system alerts that allow them to pass performance information on to customers and support the email deliverability team. The engineers can now perform ad hoc searches and correlations across all of their system data to instantly retrieve required information. Over 35 engineers and several IT management personnel access Splunk to view and build dashboards and utilize the ad hoc features of the platform.

True to their open source roots, Bronto leverages the Splunk community to add greater value to their deployment. The team at Bronto has contributed and downloaded several free Apps and Add-ons from SplunkBase, the Splunk Online Community Marketplace (http://splunk-base.splunk.com/apps/).

Future Opportunities

Bronto plans to expand the Splunk alerting capabilities on two fronts. The first is integrating Splunk alerts into the workflow Bronto has built around its open source event-handling platform, Nagios. Getting Splunk deeper into their work processes will deliver even greater value to the team. Secondly, the team will focus on leveraging the alert functionality to enhance their service levels and address system issues in near real time.

Bronto plans to leverage the reporting and dashboard features in Splunk to distill operational intelligence into a format that can benefit employees beyond the IT organization by delivering performance data in formats that can support customer service and marketing.

Return on Investment

Hard ROI	**Before**	**After**	**Savings**
Mail Event Investigation	40 hours per week Calculation: (40(hrs used)*52* $100) = $208,000 per year	10 hours per week Calculation: (10 (hrs Used)* 52 * $100) = $52,000	$156,000+ (Annually)
Buying Splunk versus Extending Hadoop	Estimated development and annual maintenance cost: $200,000	Splunk License: 25,000	$175,000
Total Quantifiable ROI			$156,000 Annually $331,000 First Year
Soft ROI	**Before**	**After**	**Savings**
Mean time to repair	Almost no ability to quickly and accurately address incidents	75% reduction	
Root cause analysis	Almost no ability to quickly and accurately address incidents	Ad hoc capabilities enable team to forensically search all events quickly and easily	
Improved Customer Satisfaction	Unable to locate and communicate issues to clients	Team can now alert clients proactively on failures prior to them identifying the issue	
Improved IT performance	Engineers tasked with time consuming manual tasks	Alerts, Dashboards and ad hoc tools free team up to focus on valuable tasks	

Quotes and Observations

Splunk's go-to-market strategy continues to reap dividends by offering a low barrier to adoption via its trial edition. Prospects can try before they buy and quickly demonstrate the benefits of the solution to fellow stakeholders, creating a quick path to purchase. Bronto Software is an excellent example of quick adoption and immediate ROI.

Quotes:

"This was a much more elegant solution to the problem."

"[Splunk] takes something that's complicated and makes it simple."

"There is nothing else out there like Splunk."

"(When justifying Splunk to his boss) Go ahead ask me anything….boom there it is. Ask me something else."

About EMA