sumologic

Sumo Logic Log Management and Analytics Service: Harnessing the Power of Big Data for Real-Time IT

A Sumo Logic White Paper

sumologic

605 Castro Street
Mountain View, CA 94041
P 1 650 625 0183 +

www.sumologic.com +

Introduction

Managing and analyzing today's huge volume of machine data has never been more challenging. Yet within these mountains of log data lies valuable information that can dramatically improve your business performance. Sumo Logic's log management and analytics service provides timely and actionable information, gleaned from any type or source of log data, that not only helps solve operational, security and compliance challenges but provides critical business insights as well.

"... enterprises are trying to organize and analyze log data using the equivalent of a teaspoon, while systems are generating it by the gallon every minute."

The Challenge of Managing Log Data

Today's IT and security departments face a host of challenges in protecting and maintaining the productivity of the IT underpinning of their enterprises. One of those challenges is dealing with the sheer volume of log data generated each day, which in some cases can reach as much as a terabyte.

Evidence of the size of this data-management problem abounds. IDC estimates compound annual growth in all types of data collected at 60 percent. In 2010, Gartner conducted a survey of IT staff at more than 1,000 enterprises globally and found that 47 percent said data growth was among their top three challenges. A Gartner research director said data capacity at large enterprises is growing on average 40 to 60 percent a year, due in major part to an explosion in unstructured data and new regulatory requirements.

Log data comes from a wide variety of sources, including user activity on websites, transaction monitoring, call records, database inquiries, GPS input and network devices. Log data (often called machine data) reflect operational issues related to underlying IT systems, and thus are harder to track and interpret than human-generated data related to functions such as sales, finance, production and personnel. An entire industry, business intelligence, has grown up around tools that analyze these data sources. In contrast, log data is largely unstructured and hard to organize, much less effectively analyze.

605 Castro Street
Mountain View, CA 94041
P 1 650 625 0183 +

www.sumologic.com +

Properly employed, log data can be the basis for what is often called "operational intelligence," a complement to business intelligence. Because the information captured in log data is so specific and detailed, it accumulates at a very rapid pace in large enterprises. The immense volume of data means that IT organizations are often unable or unwilling to try to derive value from it. Instead, the information is either never collected, or is collected but never analyzed. Or the analysis yielded is historical only, not real-time and not a basis for future projections.

Most enterprises, even those with large and sophisticated IT departments, handle log data ineffectively and inefficiently,. Systems to track log data are difficult and costly to deploy and maintain, and massive amounts of storage are needed for archiving. Analyzing all this data is another challenge altogether. Few readily available tools exist to glean meaningful insights, so IT staff do the best they can, burdening themselves with complex tasks of application development, upgrades and system management. The result is limited and unstructured analysis that only begins to scratch the surface of the complete operation.

Some business operate from the erroneous assumption that business intelligence systems alone are capable of providing the IT insights they need to maximize their infrastructure.

"One of the key benefits of Sumo Logic is that their scaling model can easily accommodate our growing volume of log data."

– Roblox

Unlike business intelligence systems, log data analysis can provide insights such as:

+ providing an early warning about problems in production applications;

+ instantly finding root causes of network or system issues; and

+ discovering security problems before they become public breaches.

Furthermore, operational issues recorded in data logs are often not IT issues but business issues. For example, if customers are having valid credit cards rejected when placing an order on a website but the uptick in rejections isn't noticed, that's a business issue. If potential delays or problems in provisioning a new service can be spotted before formal launch, that's a business issue. And if service-level agreements can be maintained through early detection of potential interruptions, that's a business issue.

In short, enterprises can gain critical advantages if they collect, organize and analyze their log data. But, how best can this be accomplished?

sumologic

605 Castro Street
Mountain View, CA 94041
P 1 650 625 0183 +

www.sumologic.com +

"Ooyala provides the underlying technology for large customers like ESPN and Bloomberg. Supporting over 150 million users a month, we must be able to process log data in real time to maintain the health of our IT systems. Sumo Logic allows us to do that."

– Ooyala

Defining an Optimal Solution

An ideal solution to the challenges of handling log data is:

+ Scalable: meeting rising demands and unexpected changes smoothly and at known cost.
+ Timely: delivering information quickly, keeping time-to-resolution as brief as possible.
+ Secure: posing no threat of loss or corruption of vital information.
+ Analytical: mining data to spot significant trends or anomalies.
+ Actionable: presenting solutions that can be implemented immediately.
+ Cost-effective: controlling costs for both staff and equipment, resulting in total cost of ownership below existing or alternative approaches.

Sumo Logic's log management and analytics service meets all these criteria.

The Sumo Logic Solution

Sumo Logic's Log Management and Analytics service offloads enterprises from the complications and growing expense of managing and extracting strategic operational insights from their complex log data infrastructure.

Sumo Logic handles all log data collection, processing, storage, forensics and analysis from a centralized, highly secure, cloud-based platform. Our cloud-based approach overcomes the inherent problems of premise-based solutions, including limits on scalability, inefficient or haphazard analysis, and uncontrolled costs. Sumo Logic's cloud-based solution leads to faster provisioning, simplified deployment and reduced expenses.

Our Real-Time Forensics engine instantly identifies anomalous conditions anywhere in the IT infrastructure, enabling operations teams to respond immediately to prevent network outages, minimize system downtime, resolve application performance issues and ensure that service-level agreements are protected.

To ensure scalability, Sumo Logic employs a patented Elastic Log Processing™ engine that scales each component of the service independently to meet every customer's unique requirements on demand. Sumo Logic also applies an indexing engine capable of scaling to gigabytes of log data per second, enabling real-time analysis.

sumologic

605 Castro Street
Mountain View, CA 94041
P 1 650 625 0183 +

www.sumologic.com +

Sumo Logic also takes a unique approach to collecting log data. All data processing and parsing are handled in the cloud, reducing the processing and parsing requirements of each collector. A library of "out-of-the-box" parsers from Sumo Logic eliminates the need to update complex parsing logic on every collector.

Sumo Logic is built around a globally distributed data retention architecture that keeps all log data available for instant analysis, eliminating the need for an enterprise to manage the cost and complexity of data archiving, backups and restoration.

A major innovation from Sumo Logic is our Push Analytics™, which eliminates the tedious tasks of manually reading log records, writing scripts and handcrafting queries. With Push Analytics, millions of records of log data can be reduced to a shortlist of key insights and reports on changes in applications, and in user and network behavior. Each of these insights is pushed to executive-level officers and technical staff to facilitate immediate investigations or remedial action. As a result, log data is transformed from a confusing, devalued commodity to a true business asset for operational managers and business executives alike.

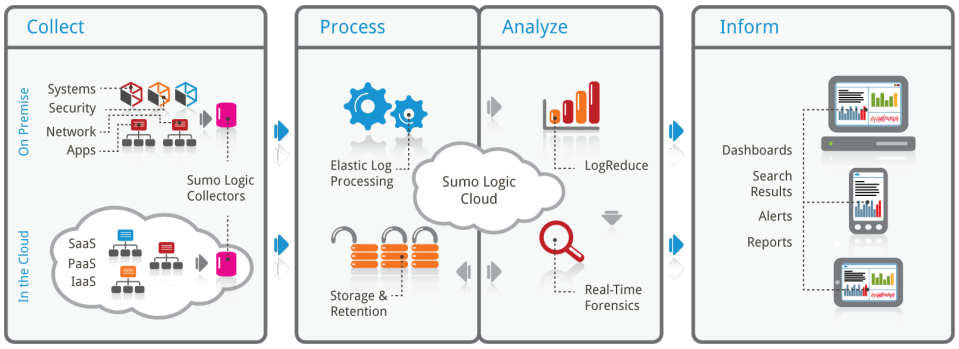

As a cloud-based solution Sumo logic service handles data collection, processing, storage, forensic and analysis through a centralized and highly security platform.

Benefits to the Enterprise

The impact of Sumo Logic's log management and analytics service can be profound. For example, consider an on-line enterprise whose payment system is generating an unusually high number of time-out or failure messages to customers trying to place orders. Every minute this condition persists impacts not only revenue but customer satisfaction. Using Sumo Logic's service, the company would be alerted more rapidly to the problem, reducing the time needed to implement a solution.

sumologic

605 Castro Street
Mountain View, CA 94041
P 1 650 625 0183 +

www.sumologic.com +

"What previously was just a massive collection of raw data can now be transformed into manageable operational insights that can have a direct effect on business performance."

Or consider the example of a telecommunications provider that constantly needs to be alert to potential fraud. With millions of calls made every day over the system, it's easy for evidence of fraudulent activity to get buried and ignored. Sumo Logic's ability to cut through mountains of data and identify key anomalies or trends helps the company to more effectively identify and stop fraud.

Finally, any business that relies increasingly on social media for customer interaction needs to keep abreast of changes in customer behavior and communication patterns. Using Sumo Logic to identify significant trends from operational log data adds a valuable complement to knowledge about customer preferences and habits which can't be gathered in real time.

The benefits of Sumo Logic fall into three categories.

+ Enhanced insight into key operational, security and compliance issues. Terabytes of log data are converted into important insights and then delivered to managers in a meaningful, actionable format. What previously was just a massive collection of raw data is now transformed into manageable operational insights that can have a direct effect on business performance.

+ Greater assurance of business continuity. When problems are spotted early on, or avoided entirely, critical production applications perform as expected, ensuring the business functions optimally. Customer expectations are met or exceeded, service-level agreements are met, and unexpected interruptions are minimized or even eliminated. And if there is a problem, real-time analytics can reduce mean-time-to-investigation and mean-time-to-resolution.

+ Reduced total cost of ownership. Sumo Logic's cloud-based service relieves enterprises of the huge burden of collecting, organizing, storing and analyzing log data. The Sumo Logic solution reduces the ever-rising costs for staff, hardware, software, networking and storage, transferring its benefits directly to the bottom line.

Enterprises trying to cope with the rapid growth in log data volumes now have an alternative to costly and complex premise-based solutions. Sumo Logic's log management and analytics service combines the efficiencies of the cloud with unique analytic tools to deliver a scalable, cost-effective and secure solution that turns mountains of raw information into valuable insights that can directly improve business performance.