sumologic

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

A Sumo Logic White Paper

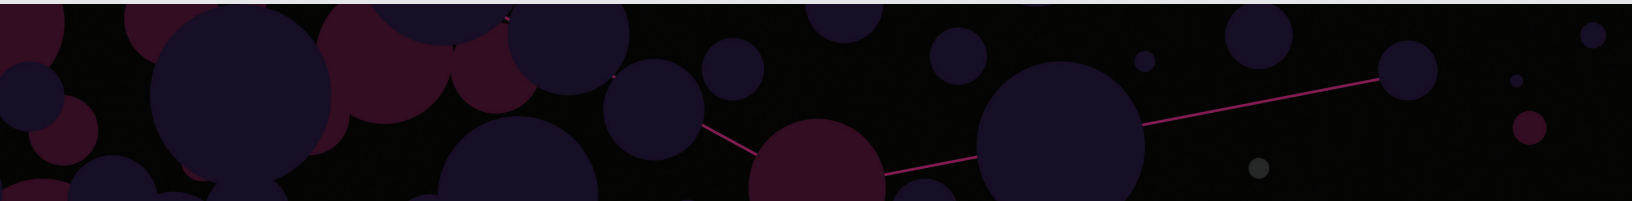

sumologic

305 Main Street
Redwood City, CA 94063

Toll-Free 1 855-LOG-SUMO +
Int'l 1 650 810 8700 +
www.sumologic.com +

Introduction

Managing and analyzing today's huge volume of machine data has never been more challenging or more critical to your business. Within these mountains of log data lies valuable information that can dramatically improve your business performance. Sumo Logic's Next Generation Log Management and Analytics Service provides timely and actionable insights, derived from log data, that not only help solve operational, security and compliance challenges but also provide critical business information.

"... 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 operations, DevOps, and IT security teams face a host of challenges in protecting and managing the critical IT infrastructure that underpins their enterprises. One of those challenges is the sheer volume of log data generated each day, which in some cases can reach into terabytes. IDC estimates compound annual growth in unstructured data at 60 percent. In 2010, a Gartner survey of over 1000 enterprises globally found that 47 percent said data growth was among their top three challenges. A Gartner research director said that 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.

Unlike structured data, such as that contained within ERP/CRM systems and database tables, for which an entire industry of effective business intelligence tools exists, unstructured log data has outgrown the existing generation of tools. Logs are generated by a wide variety of sources, including servers, virtualization infrastructure, network devices, security infrastructure, custom and 3rd-party applications, databases, RFID scanners, and more. Log data (often called machine data) contains invaluable, and most frequently the only, information about the details of the operations of IT systems and applications alike. Due to its unstructured nature and sheer volume and velocity, log data is significantly harder to analyze than structured data. However, it is no less important than structured data, and modern enterprises are realizing the value and the competitive advantage they can gain from harnessing its power.

sumologic

305 Main Street
Redwood City, CA 94063

Toll-Free 1 855-LOG-SUMO +
Int'l 1 650 810 8700 +
www.sumologic.com +

"Use of the cloud overcomes the inherent problems of premise-based solutions, including limits on scalability, manual and configuration heavy analysis, and uncontrolled costs."

Properly analyzed, log data can be the basis for what is often called "operational intelligence," a complement to business intelligence. Because logs are generated by machines and not by humans, it accumulates at a very rapid pace in large enterprises. The immense volume means that IT organizations are often unable to try to derive value from it. Instead, the information is either never collected, or it is collected but never analyzed, or the analysis is poor and not timely.

Current log management technologies are costly to deploy and maintain and need a dedicated staff as well as massive amounts of expensive storage for data archiving. Even the most sophisticated enterprise IT departments are frustrated by how ineffective and inefficient their log management and log analytics capabilities are.

Properly analyzed, log data can:

+ Provide an early warning about problems in production applications or infrastructure

+ Illuminate root causes of network or system failures

+ Enable early discovery of security breaches and compliance issues

+ Provide insights into application metrics and behaviors, and many more

IT staff do the best they can by building home-grown scripts and tools, configuring complex and manual analytics on in-house log management systems, and carefully selecting which data sets they must focus on and which they must let go. The result is limited and manual analysis that only begins to scratch the surface of what the logs can actually do for their enterprises.

Operational problems that can be addressed by properly collecting and analyzing log data are not just IT issues, they are business critical. Quickly determining the root cause of payment time-outs on a retail web application directly impacts the top line. Detecting an issue within a customer facing infrastructure that is threatening to breach service-level agreements helps protect the bottom line. Fast resolution of internal infrastructure failures means your employees can continue to perform business critical tasks.

sumologic

305 Main Street
Redwood City, CA 94063

Toll-Free 1 855-LOG-SUMO +
Int'l 1 650 810 8700 +
www.sumologic.com +

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

In short, enterprises can gain critical business advantages if they collect, manage, and analyze their log data. So what is the best way to accomplish these goals?

Ideally, the solution for collecting, managing, and deriving business insights from logs should be:

+ Scalable: keep pace with the ever-growing log volumes with a predictable and low-cost footprint.
+ Simple: easy and fast to provision with minimal maintenance and configuration
+ Real-time: delivers insights in real-time and reduces time-to-resolution.
+ Secure: encrypts and protects all data in transit and at rest.
+ Analytical: mines data to detect trends and anomalies automatically.
+ Cost-effective: does not require expensive hardware, storage, costly upgrades, or armies of FTEs just to keep operating.

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

Sumo Logic's Log Management and Analytics Service

Sumo Logic's next-generation log management and analytics service delivers actionable insights into application and infrastructure operations and security while dramatically reducing complexity and cost.

As a massively scalable, multi-tenant service, Sumo Logic performs log data collection, processing, storage and analysis within a centralized and highly secure cloud-based platform. It effortlessly handles all of your log data, regardless of volume, type or location, enabling IT teams to uncover operational and security insights buried under terabytes of logs in seconds. As a result, log data is emerging as one of the most strategic business assets within the enterprise.

Sumo Logic is designed from the ground up to handle Big Data-scale environments. Among Sumo Logic's breakthroughs is its near-zero latency Real-Time Forensics™ engine that delivers real-time search results from terabytes of log data. Real-Time Forensics™ makes critical new events occurring

sumologic

305 Main Street
Redwood City, CA 94063

Toll-Free 1 855-LOG-SUMO +
Int'l 1 650 810 8700 +
www.sumologic.com +

within the IT infrastructure instantly available for analysis. Anomalous conditions can be spotted as they occur, enabling operations teams to respond immediately to prevent network outages, eliminate system downtime, resolve application issues and improve SLAs. In short, Sumo Logic reduces mean-time-to-investigation and mean-time-to-resolution dramatically.

Sumo Logic also scales to support orders of magnitude more data than any legacy premise-based log management systems. Its patented Elastic Log Processing™ engine scales each component of the service independently to meet every customer's compute, storage and data processing requirements on demand.

Sumo Logic also takes a unique approach to log data collection. Data is securely and reliably collected through either local collection (via Sumo Logic Collectors) or through hosted collection (via https or directly from Amazon S3). All data is collected in raw, or unstructured format with no need to parse or understand the data upfront; all data processing and parsing is handled in the cloud. By separating collection from processing and parsing, which occur entirely in the Sumo Logic service, there is no need to update complex parsing logic. Consequently performance is significantly improved and management overhead significantly reduced.

To help enterprises manage exploding volumes of log data, Sumo Logic is built around a globally distributed data retention architecture featuring built-in data redundancy. Sumo Logic eliminates the need for costly SAN and NAS infrastructures and removes the complexity of data archiving, backups and restores.

Another major breakthrough is Sumo Logic's patent-pending LogReduce™ technology that eliminates the time-consuming and tedious tasks of manually reading log records, writing scripts and handcrafting queries. LogReduce technology reduces millions of log lines into a handful of human digestible patterns that enable IT teams to get to insights without having to manually writing queries to slice and dice the data. This enables IT teams to quickly find important and emerging system, application and user behavior patterns that would otherwise require days of analysis. Sumo Logic's patent-pending Push Analytics™ technology leverages LogReduce technology to automatically uncover insights and then pushes those insights proactively to IT teams in order to facilitate immediate investigations.

Sumo Logic's next-generation log management and analytics service delivers actionable insights into application and infrastructure operations while dramatically reducing complexity and cost.

sumologic

305 Main Street
Redwood City, CA 94063

Toll-Free 1 855-LOG-SUMO +
Int'l 1 650 810 8700 +
www.sumologic.com +

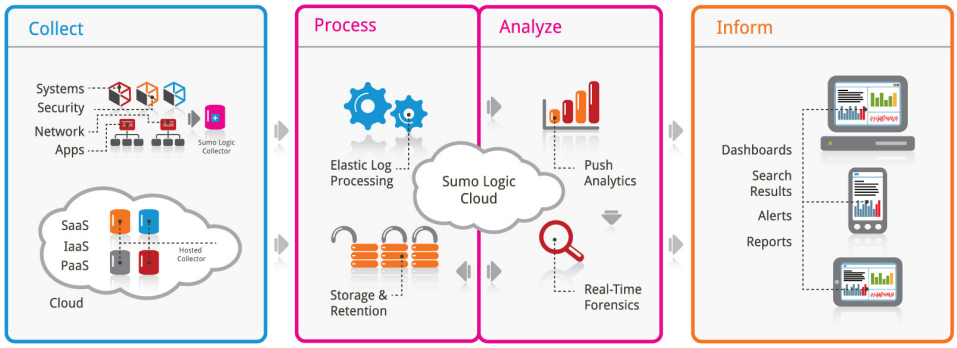

As a cloud-based solution Sumo logic service handles data collection, processing, storage, forensic and analysis through a centralized and highly security platform.

Benefits to the Enterprise

There are many business benefits a modern enterprise can derive from leveraging Sumo Logic Service. Many of those benefits involve improving business operations and reducing the overall cost of managing IT infrastructure and applications.

Better up-time through dramatically faster issue resolution: Sumo Logic Real-time Interactive Analytics enables enterprises to accelerate troubleshooting and root cause analysis. For example, the ability to find that a new exception has occurred 100 times in the last 2 minutes among hundreds of thousands of other uninteresting operational log lines lets you identify and solve the issue as soon as it starts occurring. Faster problem identification means quicker issue resolution and better SLAs, resulting in better business performance and higher customer satisfaction.

Higher productivity through better infrastructure operation: Sumo Logic LogReduce and Push Analytics enable enterprises to uncover deeply hidden problems causing infrastructure failures. As an example, identifying a specific router configuration change that is causing applications to experience failures, in turn causing hundreds of end users to sound an alarm enables you to focus on the actual problem rather than chasing down countless dead-ends. 5 vs. 30 minutes of infrastructure downtime translates into hundreds of hours of productive work protected.

Better products through deeper application analytics: Sumo Logic Real-time Interactive Analytics and Push Analytics enable enterprises to gain a richer perspective of application behavior. The ability to gain instant insights into how a new feature is performing, what the customer adoption of that

sumologic

305 Main Street
Redwood City, CA 94063

Toll-Free 1 855-LOG-SUMO +
Int'l 1 650 810 8700 +
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 impact on business performance."

feature is, and if there are any unforeseen issues with that feature enables you to better understand how your products resonate with your customers. Development, marketing, and sales teams within the enterprises can leverage this insight and build better products, market them more effectively, and have more efficient sales.

Better brand through better security and compliance posture: Sumo Logic Real-time Security Forensics and Push Analytics enables security teams to more quickly and precisely uncover security issues before they cause damage. The ability to detect a deviation in a behavior of a user or a system can shine a light on a compromised user account or a server owned by malware that is about to spread. Detecting and acting on security and compliance issues could mean a difference between pubic image degradation and a stable, trusted brand.

Competitive business through easier adoption of new computing paradigms: Sumo Logic is cloud-based from the ground-up and can collect and analyze logs from physical or virtual infrastructure within on-premise data centers as well as cloud-based infrastructures such as IaaS or PaaS. The ability to have the same level of visibility, no matter where the infrastructure lives, enables companies to leverage new computing paradigms, lower their costs, and gain business agility.

Reduced total cost of ownership through a simple, cloud-based service: 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 of hardware, software, networking and storage, transferring its benefits directly to the bottom line. It also frees highly skilled staff from having to manage and operate a complex 3rd party log management infrastructure. 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.