



GumGum

GumGum turns to Sumo Logic to support massive AWS-hosted deep-learning environment for breakthrough in-image processing

Overview

Company

- GumGum

Industry

- Computer Vision

Headquarters

- Santa Monica, California, USA

Size

- 250 Employees

Use Case

- Operations

Business Challenge

GumGum needed to establish a robust, centralized control system to coordinate its sprawling, highly complex distributed computing environment. In addition to streamlining ongoing operations, this initiative would help identify and correct problems and thus reduce revenue-sapping response latency.

Solution

The company rolled out Sumo Logic's cloud-native machine data analytics technologies to enable real-time visibility into its entire computing portfolio. Extensive dashboards supplied instantaneous views of operational status, which made it possible to detect and rectify problems much more rapidly.

Results

- GumGum is now able to freely scale its full technology stack without needing

to be concerned about machine data aggregation and analytics

- Intercepting and repairing application faults has significantly diminished events that had resulted in revenue loss
- The software development and DevOps teams benefit from a newly-optimized application rollout process

Founded in 2008, GumGum is a pioneer in the burgeoning computer vision industry. Their patented software is capable of automatically detecting objects within images and videos, while also applying natural language processing to extract meaning from the text that often accompanies visual content. There are numerous applications for this technology, with advertising and sports serving as two of the most notable use cases. For example, a pet food manufacturer may want to place a contextually relevant, highly targeted advertisement for their brand of cat food. The GumGum system can identify pictures on a website that contain cats, and then directly overlay the pertinent message within the images.



To power such a broad range of advanced capabilities, GumGum fields an exceptionally formidable computing environment. Entirely based in the Amazon Web Services (AWS) cloud, it employs thousands of servers to drive its deep learning frameworks. The company utilizes a considerable collection of AWS products, such as:

- Elastic Compute Cloud (EC2)
- Simple Storage Service (S3)
- Elastic MapReduce (EMR)
- Data Pipeline
- Auto Scaling
- Load Balancer
- Simple Queuing Service (SQS)
- Simple Notification Service (SNS)
- Kinesis

GumGum's technology portfolio goes far beyond AWS products, however, also incorporating open source solutions such as Apache Cassandra and Kafka. To enable fault-tolerance while coordinating administrative efforts, the company groups its far-flung array of servers into clusters; an individual cluster may itself contain hundreds of nodes.

Unsurprisingly, this fast-growing environment had proven very challenging to oversee efficiently. GumGum's operational team quickly discovered that attempting to examine individual servers to uncover problems (and then correlate their potential downstream impact on other servers) was cumbersome and ineffective. Failing to catch and repair these routine problems had a devastating impact on their business, which must carry out real-time bidding for potential advertisers. With mandated response times of less than 60 milliseconds, any abnormal latency means that GumGum can't promptly respond to a bid, thereby depriving the company of significant revenue.

GumGum carried out two separate endeavors to strengthen its ability to manage such a large collection of servers. First, it briefly deployed Splunk's original cloud solution, but soon realized that there were significant scaling issues that couldn't be overcome. The next, more serious effort entailed rolling out an internally managed solution using the Elastic Stack. This package was previously known as the ELK Stack, and is comprised of three open source projects: Elasticsearch, Logstash, and Kibana.

The Elastic Stack approach imposed unacceptable operational and administrative requirements – including maintaining a large, dedicated cluster – and demanded excessive staff resources to oversee it.

Beyond these open-ended obligations, GumGum's technical team also found that its architecture made it difficult to set up alerts, required configuring up Logstash parsers on individual servers, and necessitated time-consuming regular expression (regex) alterations prior to consuming any new custom logs.

Based on the subpar experiences with Splunk and the Elastic Stack, GumGum sought alternative approaches to aggregating and managing their machine data. Since a 100% cloud-based solution was an integral precondition, Sumo Logic's born-in-the-cloud offering was a natural candidate.

“We rely on an intricate, highly advanced collection of hardware and software. By centralizing and simplifying operational management of our sophisticated architecture, Sumo Logic lets us focus on growing our core business.”

Vaibhav Puranik, VP of Engineering, Big Data & Platform, GumGum

To measure its effectiveness, the GumGum DevOps team – in close collaboration with the Big Data group – carried out an extensive proof-of-concept (POC) that was targeted at the company's largest and most complex application: in-image advertising. The evaluation team was confident that if Sumo Logic could gracefully manage that scenario, all other use cases would be feasible as well. During the POC – which took approximately two months to complete – Sumo Logic impressed the evaluation team with its:

- Ease of use
- Alerting
- Aggregation
- User interface

On top of all these benefits, Sumo Logic was also a less expensive alternative. Consequently, the result of the POC was that the DevOps and Big Data teams both signed off on the Sumo Logic selection.

Today, Sumo Logic has been deployed throughout the organization and is actively utilized by more than 60 users. The teams that are most involved with Sumo Logic include the Big Data group – which is responsible for the advertising offering, the data engineering/



data science group which processes millions of pages each day using massive Cassandra clusters, and the back-end server management group.

Sumo Logic has transformed how GumGum applies machine data to help operate its entire technology architecture. Three of the most notable improvements include:

Flexibility

Sumo Logic supplies the GumGum user community with numerous alternatives to interact with their machine data. For example, some groups elected to begin including special logging flags in their application code. These flags, when surfaced through Sumo Logic, could be applied to large numbers of servers - frequently only for brief periods of time. This additional visibility was often all that was necessary to help the developer identify the culprit behind an outage or other problem.

Alternatively, other software developers eschewed modifying their code in favor of defining dashboards to display the health of their most critical services. These visual aids provide an instantaneous indication that a problem is brewing, helping to guide the team to carry out corrective action.

Issue detection and correction

GumGum's software architecture orchestrates multiple services that collaborate to achieve a goal. This means that if a particular service experiences a failure, it's likely driven - at least in part - by a fault somewhere else. Tracking a distributed computing issue to its source is a problem that plagues many organizations. By delivering easy-to-add, fine-grained alerting capabilities - which are far more sophisticated than the CloudWatch alerts offered by AWS - Sumo Logic has made it much easier for GumGum's technical teams to catch errors before they can harm the ad bidding process. In fact, many software teams now use Sumo Logic to monitor LiveTail when rolling out new application code. This helps unmask errors before the software is fully deployed.

Scalability

GumGum underpins its services with many autoscaled clusters. This means that the exact quantity of servers in operation at any point in time continually varies, and is heavily dependent on the overall workload: during peak periods, hundreds of servers are necessary, while quiet periods only require a fraction of this number.

Although autoscaling is an ideal approach for fine-tuning the ideal amount of servers, it also introduces the hazard of inadvertently losing disk-based machine-generated data during an automated server shutdown. Fortunately, by immediately transmitting this information to Sumo Logic's cloud-hosted solution, GumGum can confidently scale up and down without putting any of its critical operational data at risk.

Now that GumGum's has completed the upgrade to Sumo Logic for its most mission-critical services and projects, the next step will be to continue implementing additional use cases. One planned initiative is to leverage Sumo Logic to analyze machine data generated by other AWS services and systems.

About Sumo Logic

Sumo Logic is the leading cloud-native, machine data analytics platform that delivers continuous intelligence across the entire application lifecycle and stack. More than 1,600 customers around the globe rely on Sumo Logic for the analytics and insights to build, run and secure their modern applications and cloud infrastructures. With Sumo Logic, customers gain a service-model advantage to accelerate their shift to continuous innovation, increasing competitive advantage, business value and growth.

Founded in 2010, Sumo Logic is a privately held company based in Redwood City, CA and is backed by Accel Partners, DFJ, Greylock Partners, IVP, Sapphire Ventures, Sequoia Capital and Sutter Hill Ventures. For more information, visit www.sumologic.com.