WHITE PAPER

“True SaaS” Machine Data Analytics For DHHS Modern Applications

sumo logic

Sumo Logic is a secure, cloud-native, machine data analytics service, delivering real-time, continuous intelligence across the entire application lifecycle and stack. More than 2,000 customers around the globe rely on Sumo Logic for the analytics and insights to build, run and secure their modern applications and cloud infrastructures.

DHHS has unique certification requirements for cloud analytics. All of the data requires a federally certified solution. Yet due to the nature of the DHHS mission, much of the machine data contains HIPAA protected information. Sumo Logic provides this unique level of protection.

The Federal HIPAA/PII Compliance Challenge

HIPAA compliance demands IT infrastructure and advanced strategies for protection against data privacy risks. Compliant organizations need to be prepared for an investigation of potential security breaches.

And this means maintaining an audit trail that provides key information about any event:

- **What occurred**
- **When it occurred**
- **What caused it**

HIPAA audits require log data retention, routine reviews, and reporting on specific activity within your infrastructure. To comply, you must retain and secure ever-larger activity logs, all while adapting to evolving regulation. Compliant log management proves a challenge due to the sheer size and types of data, and organizations must have log management tools to help them automate audits and demonstrate their compliance.

Protection of HIPAA and PII data, Sumo Logic natively supports the compliance requirements to protect many types of government information. Sumo Logic was built with security at its core, with regular independent third-party audits, SOC 2 Type II attestation, FIPS 140-2 compliance, etc. Our SaaS product will be FedRAMP certified within the coming year.

The Challenge of Rising Complexity and Risk for Traditional Applications

Working with our customers, we continue to see rapid advances in tools and processes used by various enterprise personas to build, run and secure modern applications. IT departments rely on a variety of disparate tools for the development, operational support, and security. This puts a large strain on the shrinking number of IT personnel. A new consolidated approach is required.

Sumo Logic is that new approach. Sumo Logic is a cloud-native, machine data analytics service that unifies logs and time-series metrics to provide full-stack and lifecycle visibility into modern applications—from code to end-user behaviors. Powered by machine-learning algorithms and graphical, visual dashboards, Sumo Logic turns unstructured, semi-structured and structured data (logs and time-series metrics) into continuous intelligence that enables organizations to improve their insights across all machine data, improving application performance and security.
This full-stack visibility, real-time monitoring, and advanced analytics help you ensure the performance of your applications and infrastructure—moving from incidents to insights.

The Rise of Modern Application Architecture

Today’s federal agencies are striving to deliver high performance, highly scalable and always-on digital services. These services are built on custom “modern architectures” built with new technologies like containers, microservices and typically running on cloud platforms like Amazon Web Services (AWS), Microsoft Azure, Google Cloud Platform (GCP), etc.

Improved Posture in Federal Agencies

Sumo Logic helps move from the classic reactive posture to proactive by using the Sumo Logic single pane of glass solution. Federal IT departments can utilize Sumo Logic to visualize the relationship between on-premise and cloud deployments. As government agencies move to the cloud, Sumo Logic can ease the transition with better visibility into cloud applications and underlying infrastructure for AWS, Azure Google, and other CSPs. Sumo Logic Key Capabilities:

- **Full-stack visibility** for gathering event streams from modern architectures at every stage and combining them with time-stamped metrics for comprehensive visibility in real time
- **Machine learning using pattern recognition** and outlier detection to correlate logs and metrics across multiple data sources by reducing hundreds of thousands of results into a handful of meaningful patterns—for identifying errors before and after releases, lowering false positives and comparing behavior across clusters
- **Powerful analytic tools to assist troubleshooting** and root-cause analysis for quickly getting to the source of problems across applications and infrastructure to improve IT operations.

To understand, secure, and speed their operations with Sumo Logic’s next-generation cloud-native machine data analytics require the following capabilities:

Threat Intelligence

Be informed, not overwhelmed, with real-time threat detection and alerts. Integrated threat intelligence allows agencies to prevent breaches and identify adversaries.

DevOps to SecOps

Sumo Logic’s full visibility across the IT pipeline ensures security and insights from development to production. The support for centralized log analytics and metrics can improve audit and compliance and overall security posture for federal agencies. Build, run, and secure applications and infrastructure quickly, at scale, and with built-in compliance and security.

Silo Elimination

DHHS has been tasked to share data across missions. Sumo Logic allows content sharing across large federal communities when appropriate while securing information when privacy is a requirement.

Why SaaS for DhhS

First things first, SaaS is absolutely the way to go, especially when you are considering bringing on a scalable machine data analytics solution to manage your application performance and security. SaaS removes the need for organizations to install and run their management software in their own data centers, or even their cloud, and the time-to-value and low total cost of ownership (not to mention the focus on their core application) makes this the right choice for most applications.

But, the architecture of the SaaS solution also matters. To help put a finer point on things and to help aid in your organization’s decision-making process, this blog will explain why enterprises evaluating a

variety of machine data analytics providers should look for a true multi-tenant solution if they want to get sustainable value from a SaaS service. Unlike most on-premise solutions, Sumo Logic offers unlimited users to every SaaS customer!

Instant and Automated Setup

Multi-tenant software is architected for instant setup and self-serve value. You can sign up for Salesforce and be instantly productive. A hosted model, on the other hand, may require some time to set up, since every new software has to be purpose built (or at least provisioned) for the new tenant. Even vendors who have automated much of the provisioning process require hours, or even days, to set up tenant accounts.

Zero Maintenance Solution

To provide massive operational scale, a multi-tenant model enables customers to configure and customize their own account. Customers can add users, data and even third-party integrations and extensions with no administrative support from the SaaS vendor. Hosted software may provide a few of these capabilities, but typically requires administrative support from the vendor to do major configurations like third-party integrations and extensions.

Scalability and Elasticity

Since multi-tenant software typically supports thousands (or even tens of thousands) of customers, the software is very scalable and elastic to demand. Single tenant models deploy resources to satisfy a single tenant, which might seem more flexible, but it actually more limiting. In the hosted model, a tenant's software is provisioned with fixed resources and it is very challenging and/or time and labor intensive to elastically scale up or down the system.

On the other hand, multi-tenant SaaS models can balance resources across lots of customers. Since customers rarely have exactly the same usage patterns, bursts in resource requirements by one customer can be absorbed by the larger shared resource pool. This is fundamentally the core operating model of the cloud — to provide greater flexibility for the single customer by planning for capacity at the group level. With a single tenant system, if you burst, and didn't provision for your peak needs, you are now starved for resources. This is a particularly vexing issue for machine data analytics where the volume, velocity and variability of data is very unpredictable. Imagine a denial of service (DDoS) attack or a performance outage where your machine data ingest and analytics requirements will spike quite dramatically. A hosted model will cripple the enterprise the most in these situations since the system cannot scale fast enough to support the analytics need. And unfortunately, this is the exact moment when you need the analytics solution the most.

Consistent and Up-to-Date

Since there is "one software" in the multi-tenant model, all tenants are always on the latest (and same) version of the software. The latest version has all the recent capabilities and fixes, thus enabling faster innovation for customers. In the hosted model, the vendor needs to upgrade all instances when making new releases. Some of these vendors may be on different versions of software already, and other times, it could take a fair bit of time for the vendor to upgrade all the hosted software to the latest version.

Security by Design

Because of market and customer requirements, viable multi-tenant software has to be built with security in mind. Since SaaS services have to convince their customers to give up their data to another's control, they have always been held to a higher standard. Whether it's user interaction with the software or the storage of data, these systems must have the right security capabilities (encryption, key management, etc.) and policies (role-based access control, multi-factor authentication/authorization, etc.) in place to protect tenants as well as their end-users and data. Since hosted software is in essence packaged software, it does not have similar exacting standards. In fact, hosted software is meant to run within a customers environment, and was not designed to be secured at scale. After all, the onus of securing the package software is typically passed on to the customer.