



Public Library of Science (PLOS)

The Public Library of Science (PLOS) is a nonprofit open access science, technology and medical research publisher and advocacy organization. Founded in 2001 to accelerate progress in science and medicine by transforming research communications, San Francisco-based PLOS is also a technology innovator.

The PLOS website, home to more than 50,000 peer reviewed scientific papers, manages growing visitor traffic - 1K unique sessions per minute and millions of search hits daily – by relying on a largely open-source software stack. To handle the demand of two million article downloads and 12 million article views per month, engineering staffers, led by Director of IT and Operations Clark Hartsock, not only develop applications used by the Library but also contribute to open source projects.

Despite his commitment to open source software, Hartsock is also a realist. “Why reinvent one type of wheel when our company is all about building a different kind of wheel, that hasn’t been created?” he says. This pragmatic attitude led Hartsock to move away from the ELK stack (Elasticsearch, Logstash, and Kibana) to the Devo Data Operations Platform to extract operational and process insights from the growing volumes of data generated by the organization’s web, publishing and scientific communications initiatives.

PLOS had been using the ELK stack to manage data generated by the Library’s applications, web site, VMs; network infrastructure; security devices and software; QA environment, as well as its Nutanix hyper-converged environment.

Hartsock’s team of 12 includes three senior engineers, two DevOps engineers, IT Operations managers, and a Windows administrator. “With our limited resources, we didn’t have the people time to make ELK work the way it was supposed to, the way we needed it to,” he says. “Our DevOps team would have done the work of managing ELK, but they were also rolling out Saltstack to do configuration management. It didn’t make sense to throw an FTE at this when we are busy using emerging technology and new ideas to open up scientific communication to make it faster, more efficient, more connected and more useful.”



INDUSTRY: ACADEMIC/

SCIENTIFIC PUBLISHING

LOCATION: SAN FRANCISCO, CA, USA

CHALLENGE

Replace the ELK/Elastic stack to simplify management of growing volumes of data.

SOLUTION

Replace ELK/Elastic Stack with the Devo Data Operations Platform.

RESULTS

Immediate benefits from real-time insights; reduce resource use by 75%.

“We were looking for operational and process insights from increasing volumes of data. The Library needed to automate getting those insights. Devo fit a need we had with a way cooler solution. Now we have time to invent software and make progress towards our core mission.”

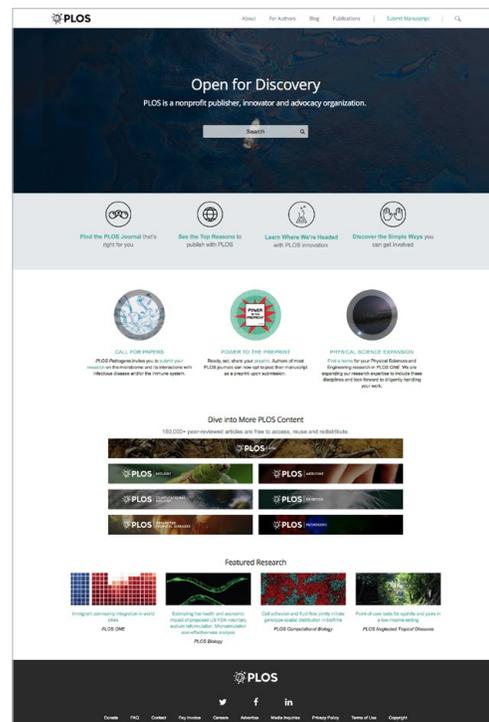
Hartsock and his team had Devo up and running quickly. Because of the volume and disparity of specialized types of data the Library collects and manages, the process took a few weeks, but Hartsock notes the library had standard operations data flowing in to Devo in only an hour.

According to Hartsock, the Library saw an immediate benefit from Devo. “We went from having an opaque mess to having real-time insights with easy and rapid access to data that is otherwise hard to get to. One of our engineers worked with Devo to create a dashboard to aggregate traffic data on the web front end to view visitor activity in real-time – which paper titles people are looking at, where users are coming from, and how site traffic flows. It’s become a key part of our operational workflow.”

Hartsock also has gained operational perspectives that weren’t possible with ELK. “Sometimes people misbehave on the front end. We’re using Devo to detect and block robot.txt exploits and API abuse. Our developers look at the platform all the time to monitor performance changes or disk access changes, and track more under the hood stuff.”

The transition has also enabled the Library to reduce costs. “We recently upgraded our CMS infrastructure,” adds Hartsock. “By looking at performance in Devo, we’ve been able to make changes to our architecture that have reduced resource use by 75%.”

Moving forward, the Public Library of Science plans to integrate tools it uses for internal tracking, including Prometheus and Graphana, with Devo. The Library is also working to be compliant with GDPR. Its BI team, which uses Sisense as a primary tool, is also investigating the use of Devo to improve its access to real-time insights.



“We were looking for operational and process insights from increasing volumes of data. The Library needed to automate getting those insights. Devo fit a need we had with a way cooler solution. Now we have time to invent software and make progress towards our core mission.” —CLARK HARTSOCK, DIRECTOR OF IT AND OPERATIONS, PLOS.

“The AHA! moment for us as we evaluated Devo was the flexibility of the product,” Hartsock concludes. “Most of the time software platform vendors are just trying to sell you stuff. When you look deeper, it’s either a skeletal framework or it’s over the top, pre-done and not flexible. The Library’s web presence is a good portion of our credibility; with real-time insights, my team can add to that value quite a bit. Before we migrated from ELK to Devo, we were walking around not knowing what was going on with the data created by our systems. With the real-time visibility and insights Devo provides, we anticipate building out more use cases and creating more insights with data that will be useful to the Library and its community. Devo hit the ROI sweet spot.”

About Devo

Devo is the leading Data Operations Platform for the digital enterprise. Devo delivers real-time business value from analytics on streaming and historical data to help Fortune 1000 enterprises drive sustained performance and growth. The Devo Data Operations Platform collects, enhances and analyzes machine, business and operational data from across the enterprise. Devo provides real-time analytics and insight for IT operations, security analytics, business analytics, customer insight and log management for the world's leading organizations.

For more information about Devo, visit www.devo.com



ABOUT DEVO

Devo delivers real-time operational and business insights from analytics on streaming and historical data to operations, IT, security and business teams at the world's largest organizations. The Devo Data Operations Platform collects, enhances and analyzes machine, business and operational data, at scale, from across the enterprise.