

Autodesk Reduces Costs by 28% with Federated Search

Key Challenges

As Autodesk's customer base and digital footprint expanded, its log data volume grew exponentially, costing the company more on observability and infrastructure.

Key Results

Autodesk leverages Splunk's federated search capabilities and intelligent data tiering, sending high-value, frequently searched logs to Splunk for real-time analysis, while less critical ones are stored in Amazon S3 data lakes, reducing costs by 28%.



Industry: [Technology](#)

Products: [Splunk Cloud Platform](#)

Solutions: [Platform](#), [Observability](#)

Capabilities: Federated search and analytics, scalable index, dashboards

Autodesk's mission of "Make Anything" requires everything to be up and running.

What do Disney, Boeing, and BMW have in common? They all rely on Autodesk software. The company provides tools for design, engineering, manufacturing, and entertainment for small businesses and Fortune 500 companies across every major industry, globally. And uptime serves as the foundation that supports the architecture, products, and films Autodesk helps create.

Principal Engineer Jyoti Kumar and his observability team monitor all of Autodesk's applications to ensure 24/7 availability without service degradation or outages that could negatively impact end users. For streamlined troubleshooting and incident analysis, Autodesk turned to [Splunk Cloud Platform](#). With a centralized logging infrastructure built around Splunk, Autodesk engineers have access to a single, unified tool where they can troubleshoot and fix performance issues efficiently.

But as the company's data ingest increased, it became difficult to gain unified insights while still adhering to IT budgets. The solution? Send high-value logs into Splunk, and the rest into Amazon S3 via federated search in Splunk Cloud Platform.

Building higher quality observability data (that costs less)

With [Federated Search for Amazon S3](#) in Splunk Cloud Platform, Autodesk now sends high-value, frequently searched log data to Splunk for real-time analysis, while less critical, infrequently accessed logs — such as ones that are only used once to troubleshoot a single issue — are stored in Amazon S3 data lakes, saving the company roughly 28% in overall ingest costs. "By migrating just one of our existing search tools to Splunk federated search for S3, we saved 78%," says Kumar.

Outcomes

- **Near real-time** metrics for critical applications
- **~28%** cost savings
- **Faster** insights and mean time to resolution (MTTR)

Distinguishing between high- and low-value data (and routing them accordingly) establishes strict logging standards for Autodesk, improving the company's overall data quality. Only logs adhering to these standards are ingested into Splunk; non-compliant logs are redirected to S3 and accessed via federated search. This optimization supports petabyte-scale log management, providing near real-time metrics for critical applications, which has led to healthier Splunk clusters, fewer outages, and improved overall performance — supporting Autodesk's push to reduce MTTR to under 30 minutes.

Maintaining the performance of its Splunk clusters is a priority for Autodesk. Poorly written search queries can lead to high CPU and memory usage, creating a domino effect that slows down data ingestion, degrades query performance, and impacts the reliability of alerts. To combat this, the company is rolling out [Splunk AI Assistant](#) to improve search query quality by allowing users to ask questions in natural language, which it then converts into code. It will also act as a quality check, analyzing existing queries and suggesting optimizations. "By ensuring queries are written efficiently, Splunk AI Assistant will significantly reduce compute cycles, freeing up resources and ensuring our Splunk environment runs smoothly," says Kumar.

"To enhance operational effectiveness, we are focusing on the strategic use of logs. This begins with ensuring all teams adhere to strict standards," continues Kumar. While static alerts are useful, Autodesk leverages machine learning models within [Splunk's AI Toolkit](#) for anomaly detection, which can identify unusual spikes in errors that traditional methods might miss. "This allows us to identify when metrics like error counts suddenly increase beyond their normal patterns, enabling us to catch and address issues much faster," Kumar says.

But running massive data searches can be costly and time-consuming. As Kumar explains, "Every time we run a federated search, it costs a certain number of Data Scan Units (DSUs)." To make things more efficient and cost-effective, his team got proactive. "We looked into our users' history and built an automated scheduled summary search around common queries," says Kumar. The result is a win-win: By running the most popular searches ahead of time, this "summary index approach" helps Autodesk save on compute resources while getting key data to users almost instantly.

Applying a CAD-do attitude to observability

Reducing data silos and blind spots has allowed Autodesk to find and fix issues before they impact customers. "Before consolidating our logging tools into Splunk," explains Kumar, "to troubleshoot one service, you'd log into one tool. And to troubleshoot a second service, you'd log into another tool. We were logging into multiple tools to troubleshoot a single issue." Now, users simply log into Splunk Cloud Platform to see how their downstream and upstream services are performing. This centralized visibility has reduced troubleshooting times and enabled faster resolution, improving the company's overall observability practices.

"Only Splunk can ingest log events, create an infinite number of metrics from them, and utilize it for troubleshooting or alerting," says Kumar.

From Kumar's perspective, Splunk's intuitive Search Processing Language (SPL) and self-service features have driven high internal adoption, mitigating tool sprawl and unifying visibility across teams. "Splunk's SPL language itself is powerful," says Kumar. "It allows us to correlate logs from upstream and downstream services and then produce alerts which will tell us which service is down."

"Other platforms support federated search," continues Kumar, "but only Splunk Cloud Platform provides users with a single interface to query both Splunk and S3 seamlessly." Ultimately, Splunk enables teams across Autodesk to create their own alerts, dashboards, and analytics, supporting faster MTTR and operational efficiency.

Our world is built on Autodesk. Autodesk observability is built on Splunk.



Federated search has been a game changer for us in cost optimization. We route only the critical logs to Splunk and keep everything else in S3 for ad hoc or audit needs.

Jyoti Kumar, Principal Engineer,
Autodesk

Embracing an AI-driven future with Cisco Data Fabric

As Autodesk continues its federated search journey (the company has already signed up for the Federated Analytics 2.0. Alpha Program), there's interest in further integrating Splunk Cloud Platform with Cisco's Data Fabric, particularly when it comes to LLMs and AI assistants. "Storing the vast amount of data required for training AI models requires a preeminent solution," says Kumar. "That's where Cisco Data Fabric comes in."

Cisco's Data Fabric architecture, which enables cost-effective storage and access to massive datasets, aligns with Autodesk's ongoing strategy to optimize infrastructure spend while supporting advanced analytics and AI. "Cisco Data Fabric is a natural evolution toward industry-standard, scalable data lakes," says Kumar.

Autodesk plans to use Cisco Data Fabric to enable ad-hoc analytics on logs and proprietary data, while also supporting Internal MCP agents utilizing Splunk MCP endpoints for improved data access and operational insights.



Cisco's Data Fabric vision resonates with Autodesk's approach — separating operational and analytical data, improving data quality, optimizing for cost, and leveraging large-scale internal data for AI.

Jyoti Kumar, Principal Engineer,
Autodesk

Download Splunk for free or get started with the free cloud trial. Whether cloud, on-premises, or for large or small teams, Splunk has a deployment model that will fit your needs.