The ITOps Guide to Data Management

Simmering down the data noise to boost resilience and observability

splunk>
a CISCO company

The heat is on. The tension in the kitchen is tangible as a team of chefs rushes to impress hungry guests on a busy Saturday night. Dozens of ingredients and cooking utensils are strewn across the kitchen, causing confusion and slow preparation. With so much to keep track of, problems are going undetected — as chefs rush to plate their dishes, a pot of soup boils over and leaves a puddle on the floor. The noise is overwhelming as pots and pans sizzle and clash, chefs shout incoming orders, and timers blare in unison. The kitchen is on fire, but each chef must maintain composure to keep customers happy.

As an ITOps professional, you're not laboring over a hot stove in your day-to-day workflow, but this chaotic scene may still sound familiar as you deal with growing data volumes in your evolving digital landscape. Gartner predicts that large enterprises will triple their unstructured data capacity across on-premises, edge, and public cloud locations by 2028, compared with mid-2023. To be successful, your team needs to be able to make sense of all this data — because lack of visibility leads to slower detection, investigation, and response, resulting in costly problems like downtime, silos, and more.

So, how can you prevent these "fires" from happening in the first place? The answer is **data management**.

Poor data management stirs up trouble

Slow MTTx

A chef preoccupied with tracking down all the ingredients they need may be slower to realize that food is burning on the stove, or that customers across the restaurant are losing their patience over delayed service. Similarly, in the disorganized "kitchen" that is your digital landscape, massive (and growing) data volumes can make you slower to anticipate and notice problems. When your ITOps team can't study siloed data holistically, you'll be slower to discover monitoring and troubleshooting issues. If you're not efficiently collecting, managing, and storing growing data volumes, something is bound to slip through the cracks. This may cause you to miss critical insights hidden within large volumes of data, delaying incident detection and resolution and impacting user experience.

High costs

Service errors within your digital landscape can lead to unhappy customers, impacting revenue. Think of it this way: When restaurant patrons are dissatisfied, food and time are wasted, discounts are often expected, and negative reviews push potential customers to take their business elsewhere. In the same way, inefficient data management can impact your business's bottom line. Poor data visibility creates blind spots, hindering effective decision-making and operational efficiency. This results in costly unplanned downtime and lost revenue. Lack of control over data volume can also lead to expensive war rooms, troubleshooting delays, and poor user experiences. And finally, manual dashboard creation and data handling can increase administrative costs.

To save money, your ITOps team needs observability tools that help navigate massive amounts of data and proactively detect incidents. This takes your digital "kitchen" from unorganized and chaotic to streamlined and efficient — enabling you to detect and prevent issues before they damage your business's reputation or customer loyalty.

Capability and data silos

A skilled chef doesn't need a specialized gadget for every separate task in the kitchen — that would be inefficient, clunky, and costly. All they need is just one high-quality knife that can slice, dice, and multitask. The same is true for your ITOps team. Oftentimes, teams have too many data management tools to reconcile. This can lead to more toil, lack of unified visibility, and longer MTTD/MTTR. If your team is investing in a tool for a single use case (like filter and route capability) while it offers no operational value otherwise, you're likely overspending, and your capabilities and teams are siloed.

Silos don't stop there. Enterprises deal with vast and growing amounts of structured and unstructured data from diverse sources like cloud providers, IoT devices, internal systems, and remote users. Without proper visibility, it's challenging to manage cloud usage and data, leading to data and resource sprawl. In addition, data that is decentralized across tools and storage locations can lead to duplicate, untraceable data, making it difficult to gain context and pull insights. As a result, it's unclear whether available data is complete, accurate, and relevant, limiting the ability to make data-driven decisions.

Data overwhelm

A busy restaurant is ultimately a good thing. However, when the kitchen is unorganized and unprepared for a high volume of detailed orders, the staff can quickly get overwhelmed. The same is true for your organization. As your team struggles to generate, capture, store, protect, and analyze more data now than ever, your existing systems can turn into a giant pressure cooker. Vast amounts of data in various formats can make it difficult to maintain consistency and usability.

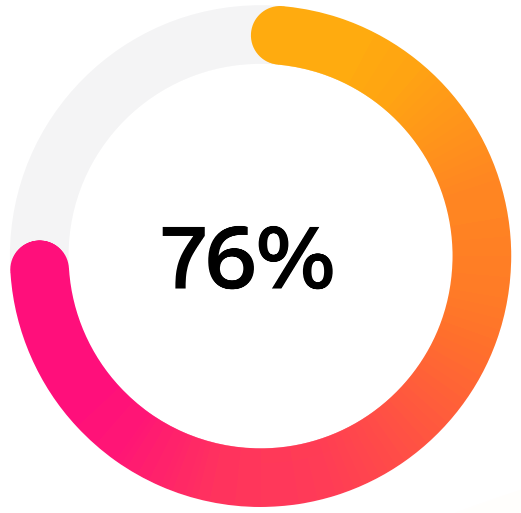

of tech executives in Splunk's report, The Hidden Cost of Downtime, agree that data sprawl was the number one challenge to managing downtime

A recipe for success

So, what is the recipe for observability success? The answer is **data management**. If your team is ready to conquer all of the challenges we mentioned previously, a solid data management strategy will lead to:

Improved MTTx

Effective data management means less noise in your alerts and dashboards because you can filter out the data you don't want, mask sensitive data, and enrich or transform it to the structure that best suits your needs. Splunk data management helps you easily enrich and transform data, converting the logs to metrics to get data in a manageable size and shape so you can spot and resolve problems quickly. Alerts become more timely, accurate, and efficient, improving MTTx and normalizing data to speed investigations. Your "kitchen" goes from loud and chaotic to systematic and orderly.

When you can store data in ways that minimize costs and maximize value, you'll experience more intelligent event management, better alerting, and guided root cause analysis. This reduces MTTR, prevents outages, and streamlines processes.

Maximized value of data

Splunk data management can lower operational costs by allowing you to filter, transform, and route data to the most cost-effective storage location for your use case. For example, it can filter out low-value data and send it to third-party, lower-cost storage like Amazon S3. This helps you optimize your data volume so you don't have to pay for low-value data that is likely to be searched less frequently, taking up space, and costing you money.

An optimized database makes it easier to uncover critical insights about digital systems, including the source of customer and business-impacting issues. In other words, with the right data management tools, teams can accelerate root cause analysis and reduce their mean time to detect and respond to problems. This helps them avoid costly downtime and associated costs.

Capability consolidation

To ensure you're not blowing your budget on single-use tools, you need a unified platform where security and observability practitioners can access the same data sets and apply them to multiple use cases. The unified capabilities within Splunk will reduce the cost of single-use or single-capability tools, while further eliminating data and team silos.

End-to-end visibility

When a kitchen is organized, a head chef has eyes on all processes, and chefs have specific recipes to follow, it's easier to maintain quality control. In the same way, visibility into data helps you adhere to operational resilience mandates, resolve production issues faster, and optimize resources. Visibility is foundational for building a leading observability practice and setting up resilient digital systems. When you combine Splunk's data management capabilities with federated search (and soon federated analytics), you can access data at rest without having to reingest it into Splunk. This allows you to view all your data, regardless of location. You can also still create dashboards and view analytics for that data.

Splunk data management gives you control over your data, its format, storage locations, and what tools can access it to manage the right balance of cost to value.

Key tools

Splunk data management includes key tools that help you troubleshoot faster and gain seamless, enterprise-wide visibility into massive volumes of data cost-effectively. By focusing on the right data, moving quickly, and not wasting time on low-value data, you can improve MTTx, optimize costs, consolidate tools and their capabilities, and gain end-to-end visibility.

Splunk helps you:

- Aggregate data smoothly across complex multi-cloud, hybrid environments
- Normalize and consolidate alerts
- Generate insights into log and metric data with powerful search language and rich visualizations
- Optimize data to accelerate detection, investigation, and response
- Mask, filter, and transform data to normalize results, manage noisy data volumes, and accelerate investigation and response
- Break down data silos and integrate data from diverse sources to provide a unified view that facilitates quicker and more informed decision-making
- Eliminate data duplication and uncover insights with advanced indexing, powerful search capabilities, and comprehensive analytics
- Optimize storage, processing costs, and data value with intelligent data tiering and summary indexing
- Convert logs to metrics for more optimized monitoring

In a busy restaurant, chefs need to have detailed oversight of kitchen processes so they can deliver an exceptional customer experience. To keep your organization running just as smoothly, Splunk unifies data management so your ITOps team can improve visibility and gain insights while optimizing costs. With Splunk, ITOps teams can eliminate data silos, control the flow of data, and expand data access through federation to gain choice and efficiency without sacrifice. Our highly scalable, flexible platform boosts resilience and observability by helping you simmer down the data noise, find and fix issues faster, improve system performance, save on downtime and operating costs, and speed up decision-making.

Get started today with Splunk data management and federated search.

splunk>
a CISCO company