Delivering High Performance, Resilient Digital Services

An enterprise-grade, full-stack, service-centric, AI-powered approach to Observability

The new digital normal

Business changed forever in 2020. Forced to accelerate their digital transformation plans to meet new requirements, organizations are now addressing the technology challenges that came with it. Whether engaging customers online via e-commerce, offering distance learning, delivering medical care through telemedicine services, or operating a complex supply chain with enterprise resource planning (ERP) systems, organizations now rely on digital services more than ever before to power their most critical business activities and customer experiences.

But antiquated and siloed approaches to monitoring and managing digital services have left organizations unable to meet the new demands of digital transformation. The cost of service interruptions and outages has exploded, leading to expensive, if not fatal outcomes for organizations. Breaches of contractual or regulatory service level agreements can lead to costly fines. Most severely, the failure to meet customer expectations with digital services can lead to churn and poor brand reputation.

Organizations have been stuck with outdated approaches to monitoring digital services because of the hangover of legacy development and operations practices combined with traditional service architectures. As organizations have adopted more agile practices for software development, more functionally integrated approaches to production operations, and more elastic, ephemeral, and containerized architectures for delivering services, outdated approaches to monitoring have created more blind spots, not more visibility.

Now is the time for organizations to adopt a new approach for meeting the needs of the new digital normal. This is the opportunity:

- To continuously and proactively direct both automated and manual responses to incidents using data and analytics
- To enhance the delivery of services across complex, hybrid technology ecosystems with a service-centric, analytics-driven approach that continuously improves service health across the tech stack
- To increase operational efficiency, reduce time-to-market, maximize service uptime, and deliver better user experiences

The evolution of service operations

Historically, the waterfall nature of application development combined with the traditional architectures of services led to development and operations teams with three types of roles: the service delivery manager, the applications performance manager and the technical/infrastructure/network operations manager. By effectively holding responsibility for each of the three layers of the digital service stack

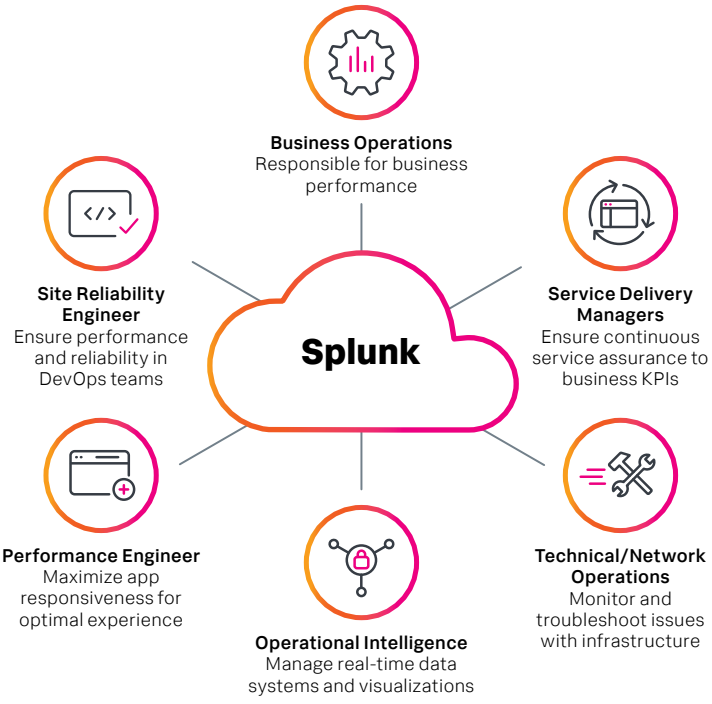

Figure 1: Everyone across development and operations must have access to common data

(business, user/applications, and infrastructure), these three roles were able to monitor their respective layers of the digital service stack without much interaction despite having logical dependencies on one another. Thus siloed monitoring tools were born.

While these three roles still exist in most organizations, there has also been recognition for the need for more functionally integrated roles to deliver digital services in complex and hybrid technology ecosystems. Cloud native environments, in particular, have empowered development teams to become much more agile and continuous than ever before. This acceleration and agility has enabled development processes to "shift left" and empower developers to focus on quality earlier in the development cycle. At the same time, newer environments have brought service architectures that are much more elastic, ephemeral and containerized. This leads to complex service maps of the relationships between different components of the digital service stack. The increased velocity of continuous integration and development combined with the increased architectural complexity is what has motivated the growth of roles like site reliability engineering.

Site Reliability Engineers (SREs) are a hybrid of application performance managers and technical/network operations managers. They ensure the performance and availability of elastic, microservices-based applications and infrastructure. SREs have driven new operational practices both for application release, such as canary testing and chaos engineering, as well as for incident response, such as programmatic auto-scaling and auto-remediation.

This evolution has demanded that everyone in the development and operations team have access to one thing: data. Instrumentation and telemetry is critical to giving Service Delivery Managers, Application Performance Managers, Technical/Network Operations Managers and Site Reliability Engineers complete visibility into the health of the service stack. But as the complexity of digital service architectures increases, data exhaust increases as well. A modern digital service that spans multiple sets of microservice-based applications and elastic infrastructure will generate an extreme volume, velocity and variety of alerts, events, logs, metrics, and traces. Data incorrectly analyzed by domain creates costly visibility gaps that prohibit development and operations teams from delivering highly available, high-performing digital services.

So as development practices become more iterative and service architectures become more dynamic, the approach to finding, fixing, and preventing issues within digital service operations must evolve too. New approaches to observability fundamentally reframe outages and performance degradation as predictable and preventable, rather than as sudden and costly emergencies.

The four modern operations challenges

Modern teams face four critical challenges to delivering continuous response and service improvement.

1. **Managing complex, ephemeral hybrid service architectures.**
   Very few organizations manage a static technology portfolio. Rapidly evolving cloud-native architectures can lead to "service drift," where configuration databases and service meshes are never actually synchronized to the operating environment. As a result, operations teams often have outdated and incomplete understanding of the infrastructure, application and user-experience dependencies behind a service. Without this understanding, it can be difficult to determine which stakeholders are most qualified to investigate and diagnose potential service interruptions.
2. **Struggling with siloed and inflexible monitoring tools.**
   The specialized nature of legacy monitoring tools have inadvertently created a sprawl of isolated dashboards and noisy alerting systems in operations teams. This in turn has made it much harder — ironically — to understand the overall health of a service and efficiently remediate problems. Operations teams suffer fatigue and burnout from dealing with alert storms. Misalignment between teams who see conflicting views of data lengthens the mean time to find and fix incidents.

3. **Triaging service issues with incomplete, sampled data.**
   Operations teams can't identify critical issues if the data is scattered across silos, and must be sampled before it can be analyzed. They also can't proactively detect anomalies in service behavior and prevent interruptions before they become outages. Without the data necessary to fully understand what happened, post-incident reviews don't surface the kind of insights necessary to develop automated preventive measures for the future.
4. **Ensuring service reliability and business performance.**
   The ultimate goal, of course, should be to ensure that business services are performing optimally and meeting business objectives. If an organization wants to tackle the challenges of the new digital normal, it's vitally important to be able to continuously correlate infrastructure, application and user experience metrics to business outcomes.

The limitations of custom, domain-specific approaches

Organizations have reacted to the data explosion by creating new analytics, data science and intelligence teams. Together they're responsible for instrumenting service stacks so that people with different roles across the operations community have the real-time visibility required for monitoring and troubleshooting.

Operations teams are often tempted to build their own custom Observability systems with open-source analytics and visualization frameworks in the hope of getting flexibility and customization. Far more often, though, this approach actually makes it harder for them to get the flexibility they need the most.

In order to keep up with inevitable changes in the technology ecosystem, organizations need the ability to ingest and process any type of alert, event, log, metric and trace data in real-time without relying on proprietary data collection technologies.

Further, organizations need the ability to understand the application and infrastructure service topology as catalogued by a configuration database or service map. This is necessary in order to proactively prevent degradations of service, as defined by business teams, and identify the root cause of issues when they occur. But without a standardized way of correlating insights from operational data to business KPIs, it's nearly impossible to predict the impact of potential service interruptions or provide visibility into the positive impact of successful digital service operations on the business.

This is precisely why the key to sustainable success with Observability is the selection of the right data and analytics platform. Without the right analytics platform to make the right data available to the right monitoring tools at the right time, operations teams are at the mercy of the tech stack and forced to manually collect, index, correlate and build custom visualizations leading to high total cost of ownership.

The right analytics platform, on the other hand, provides both the flexibility and power required to completely extract the most relevant signals from data exhaust, quickly answer questions about the health and performance of the service stack, and accurately triage both automated and manual responses. Using the right analytics-driven approach to Observability, organizations can find, fix, and prevent more issues, more accurately and more quickly.

The advantage of domain-agnostic observability

Addressing the four challenges of operations teams in this complex, fast-changing environment requires a domain-agnostic approach to Observability that addresses all of the objectives of the development and operations teams using:

- a data plane based on real-time, open-source data collection that can collect any type of alert, event, log, metric, and trace data at scale,
- a control plane based on a real-time and streaming-enabled data platform for processing all data, without partial tracing, sampling, or aggregation, into actionable insights within minutes,
- targeted AI-driven applications to understand and monitor different aspects of service health to inform both automated and manual actions.

Splunk's unique domain-agnostic approach helps teams understand the impact of change on service performance and health, while dramatically reducing the mean time to remediation when performance or health degrades.

Greater flexibility with open source data collection and scalable data processing

Splunk is a leading contributor to the OpenTelemetry project and has always built solutions based on collectd-based data collectors. Unlike the proprietary instrumentation of legacy monitoring tools, Splunk's future-proofed approach to data collection empowers operations teams to analyze nearly any data type today and easily incorporate new data types from infrastructure, applications and user experiences yet to be released. This eliminates vendor lock-in and maximizes long term ROI.

Splunk's Data-to-Everything Platform is a non-relational data platform that provides fast streaming, filtering, aggregation and analysis of real-time data. Thousands of organizations have adopted it for a wide variety of use cases. Operations teams use Splunk's platform to power a comprehensive set of activities from infrastructure and application monitoring and troubleshooting to business service monitoring.

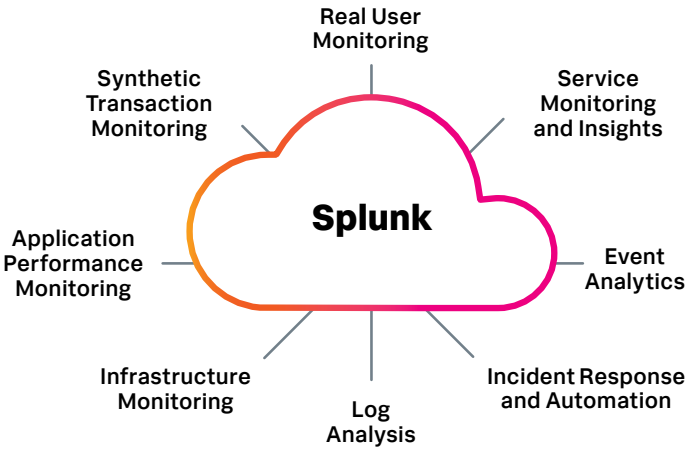

Figure 2: Organizations use Splunk to monitor across the full stack

Splunk's portfolio of observability solutions

Splunk's enterprise-grade, full-stack, service-centric, AI-powered portfolio of observability solutions provides operations teams with targeted tools to find, fix, and prevent more issues, more accurately within seconds, including:

- **Splunk Infrastructure Monitoring:** Real-time, metrics-based infrastructure monitoring with pre-built dashboards and intelligent alerting
- **Splunk Log Observer:** Infinite log analytics for rapid troubleshooting of application faults
- **Splunk Application Performance Management:** Real-time, NoSample™ based full-fidelity application performance monitoring and AI-driven directed troubleshooting
- **Splunk Real User Monitoring:** Real-time, full-fidelity user transaction monitoring
- **Splunk Synthetics Monitoring and Web Optimization:** Simulates end-user requests to measure and optimize availability and performance
- **Splunk IT Service Intelligence:** KPI-driven event analytics and service insights to prevent service interruptions before they happen
- **Splunk On-Call:** Intelligent and automated incident response with collaboration.

The new service-centric operations team

Modern operations teams don't need to suffer the fragmented visibility of legacy tools and long-term inflexibility of customized analytics solutions. Using domain-agnostic Observability solutions, development and operations teams can turn the extreme volume, velocity and variety of operational data into targeted insights to find, fix, and prevent more issues, more accurately within seconds.

Figure 3: Splunk helps find, fix, and prevent more issues, more accurately across the stack

The four challenges of operations teams can now be turned into opportunities for improvement. The flexibility of complex, ever-changing service architectures can be matched with a flexible approach to ensuring their health. The alert noise created by siloed monitoring tools can be filtered out using machine learning, alleviating fatigue and burnout. Data can be collected and correlated across different teams, without loss of fidelity, in a centralized platform to provide everyone an accurate and complete view. Operational insights can also be mapped to business KPIs to create cross-functional alignment and prevent service interruptions before they impact end users.

This modern approach empowers operations teams to pursue many new opportunities:

- First, the increased operational efficiency from improved transparency across the entire technology stack helps operations teams find more issues, more accurately within seconds.
- Second, the reduced time-to-market from greater collaboration can help operations stakeholders fix more issues, more quickly.
- Third, the ability to meet service level objectives more effectively using predictive analytics prevents severe issues from impacting end users all together.

All of this transforms the productivity of operations teams and empowers them to spend more time on innovation and capability development vs. never-ending cycles of fighting problems.

This is why Splunk's innovations in domain-agnostic Observability give everyone across the operations organization the visibility needed to respond to the demands of the new digital normal. Splunk is the platform of choice for 92 of the Fortune 100, and is recognized as a market leader in Observability, IT Operations Management (ITOM) and AIOps by multiple industry analyst firms. Using a service-centric, analytics-driven approach, operations teams can continuously ensure the health, resiliency, and performance of complex digital services across any user experience, application, and infrastructure, whether on-premises or in the cloud.

Visit splunk.com for more information.

splunk>

Learn more: www.splunk.com/asksales

www.splunk.com

21-19853-Splunk-Delivering High Performance Resilient Digital Services-SG-104