# Demystifying the Effects of System Changes on Business Metrics

A guide for SREs

splunk>

It's a dark, stormy evening at the manor, and something — or someone — is stealthily causing disorder. While you attend to the guests, a series of unexpected events leads to chaos. A toppled candlestick sets the drapes ablaze in the ballroom. A lead pipe inexplicably bursts and floods the lounge. Burners spontaneously ignite in the vacant kitchen. As your team files into the dining room and gathers around the table to investigate, questions hang in the air: Were these incidents mere coincidences, or were they plotted by someone in this very room? How will they affect the experience of your highest-paying guests and the manor's operations? And most importantly, *why* did they happen? The game is afoot.

Navigating this situation is like being an SRE, but monitoring every change in your application environment can be trickier than catching the Colonel in the conservatory with a candlestick. That's because planned and unplanned changes happen all the time, and tracking when, where, and why they occur can be complex without the right tools. Whether it's a planned new release of a particular service or an unexpected infrastructure failure, changes like these can impact your business and negatively influence key performance metrics like revenue — which is far from fun and games. With this in mind, it's important to monitor how system changes affect applications and the business so you can immediately deploy resources to fix high-priority issues.

Often, SREs lack the visibility to solve these mysteries, but with the right toolkit, you'll be better equipped to monitor key business metrics, maintain a positive user experience, and respond to issues more quickly.



# Why a concealed view has its consequences

Reliability is crucial for a modern organization. If systems are down, a lot can go wrong. As an SRE, you're at the heart of this: deploying, scaling, and fine-tuning the applications that power businesscritical workflows like checkouts, order fulfillment, account logins, and more. It's especially important that key customer segments, like those who spend the most money with your business, enjoy friction-free experiences. Just like the staff at the manor would rush to help their premier guests first, your team will ideally be able to focus on customer segments that are hit hardest by planned or unplanned changes.

Changes in your application environment may be planned, like software releases and A/B tests; or unplanned, like autoscaling, infrastructure failures, traffic spikes, cyber attacks, and software bugs. Whether or not you see them coming, they can interrupt critical workflows and influence the metrics that matter most to your business. These metrics will vary from company to company — for example, if you work for an online retailer, you'll likely keep availability, page rendering speed, and revenue per minute under close surveillance. However, with so much complexity and countless variables, it can be difficult to quickly uncover the reasons behind shifts in metrics.

Like a detective examining fingerprints with a magnifying glass, the capability to zoom in on specific system changes can help you understand these shifts. But you don't want to comb the entire manor with your magnifying glass, so first, you'll need tips on where to start your investigation. If you're unable to get this guidance that will help you deep-dive into the performance of a specific service, you may have a harder time figuring out the root cause of the issues you've found.

In most cases, organizations don't have granular visibility into how system changes influence the business. With limited visibility, issues are missed, and the consequences can be catastrophic:

- Downtime and unhappy customers can lead to decreased revenue, which has a domino effect on employee retention, bonuses, and promotions.
- Teams waste valuable time investigating problems in war rooms, which results in less time for more interesting work or contributing to strategic initiatives.
- On-call becomes more stressful because it takes longer to find problems and route them to the person responsible for correcting them.
- IT staff become burned out and frustrated due to being the first point of blame every time an issue occurs.
- In the aftermath, postmortems take up a ton of time that could otherwise be devoted to automating processes, working on new things, and making life easier.



### Solving the mystery: key strategies

To improve visibility into how changes impact business metrics, add these strategies to your detective toolkit:

#### **Clearly define key business metrics**

Which services have the greatest impact on your organization's critical metrics? Identify these and establish benchmarks for what are considered "normal" operations. You need visibility into time over time comparisons (like week/week), planned traffic fluctuations (like a surge in traffic on a major holiday), and unplanned traffic fluctuations (like a viral social media mention or media coverage.)

#### Implement strategic tagging

Tag your data with information about your customers and versions — like where the service is deployed, what version it's running, the spend level of the customer making the request, the customer ID, etc. This will help you navigate data more easily to pinpoint and remedy issues arising from planned or unplanned changes.

#### Monitor the performance of key workflows

Like a detective who connects clues on a bulletin board with a red string, group combinations of microservices that perform key backend functions (such as checkouts, logins, etc.) for clarity. Filter these workflows in your observability tool to better comprehend their performance and single out the root cause faster.

#### Isolate performance for specific impacted segments

OpenTelemetry includes a flexible and easy way to tag data. You can use the tags to filter out extraneous traffic from your view. Isolating specific segments (like traffic from an iOS app) eliminates red herrings that can cloud your investigation, accelerating your analysis of who is affected, why, and when. With the right observability tool, you can see performance data separated by these tags in one view, making it effortless to determine the cause of performance problems.

Once you've zeroed in on specific workflows and segments (a smaller set of data), you can quickly see if services are working as intended post-change, if any services are causing mischief, and who owns the repair.



Splunk Tag Spotlight breaks down key metrics by tag, including for key metrics like product version.



#### **Engage developers with shared context**

Every good detective keeps track of their investigations in a logbook to help rule suspects in or out, record evidence, and recognize connections among clues. For the same reason, it's important for your team to have all key telemetry data (such as metrics, logs, and traces) a click away. This makes it easier to filter out any irrelevant data. Once you have a strong grasp of the issue, you can either fix it or share it with the developers who own the service in question. This approach lets developers see all the related logs and metrics in one place, enabling them to analyze incident data and find a resolution quickly.

#### **Collaboration and continuous improvement**

Nobody is an island — no matter how quickly you can isolate problems, you will need to be able to work with peers in ITOps, security, and engineering to rule out all possible causes and to make sure that infrastructure or security issues are fixed. Filtering and analyzing data will help you present a compelling case to your partners and show where the problem is, ultimately making it easier and faster to get customers back up and running.

Using the right observability platform makes it easy to find the root cause of the problem, thanks to tagged data and logs being integrated with the overall troubleshooting workflow. This enables everyone to have the same data about a possible incident and quickly fix the problem. Even if you took all of these steps tomorrow, it would take time to see the effects, so a collaborative, iterative approach is essential. Lean into improvements based on blameless postmortems so you can address bottlenecks and enhance processes over time.







### **Real-world examples**

At Splunk, we often hear teams say they need support to achieve the steps we outlined above. **Splunk Observability Cloud** can do it all and supports segmenting transactions (traces) with basically anything, including customer type or spending level. This helps you better understand how system changes impact the business. Read on to learn about the results real-world companies have seen with these solutions.





6

### Lenovo

With Splunk Observability Cloud, the Lenovo team uses full-fidelity data and predictive analytics to monitor infrastructure performance at cloud scale. "What's most lovely about Splunk is we benefit hugely from having centralized, customizable analytics dashboards that collate and analyze transactions in real time, ensuring that we respond to customers in a timely manner while spotting errors and latency at a glance," says director of operations for the online and eCommerce platform at Lenovo. As a result, the average time it takes to recover from a system failure has now gone from 30 minutes to about five minutes. Faster troubleshooting means increased productivity for developers and lower revenue loss for the organization.

"The fast-performing and lightweight user interface is extremely convenient, too," Lenovo's director of operations adds. "Technical data is more readily understood when presented in a graphical form like Splunk's Service Map. Our developers love the pictorial topology that clearly displays the relationship between different services created for a certain application, as they can now better collaborate with other development teams who are building services for the same application. Apart from the powerful visualization, another bonus is the seamless integration with Microsoft Teams for alerting purposes." Equally impressive are the heatmap visualizations of Lenovo's entire infrastructure and traffic light status reporting, as well as the advanced host monitoring function, which allows the team to capture the real-time status of each server. On Black Friday 2020, Lenovo offered a doorbuster deal on computer products and gave away a limited number of gaming products as incentive gifts. While Lenovo had expected a sudden surge in sales and web traffic, the spike turned out to be a staggering 300% higher than the same period in 2019. Thanks to Splunk Observability Cloud, Lenovo's online shop maintained 100% uptime — with zero outages or digital crises — and delivered a flawless shopping experience despite the massive increase in traffic to its website and mobile app. During these peak times, the Splunk support team also stood by to provide Lenovo with 100% peace of mind.

Lenovo was able to maintain 100% uptime despite a 300% increase in web traffic on Black Friday 2020 with Splunk Observability Cloud.





"The more complex our architecture became, the harder it was for us to detect problems," says Alejandro Comisario, executive vice president of engineering at Rappi. But by switching to Splunk Observability Cloud, Rappi has end-to-end visibility into its distributed microservices-based architecture, which includes Amazon Elastic Container Service and Kubernetes clusters. With a more agile approach and real-time observability from Splunk, the Rappi IT team now efficiently manages more than 1,000 microservices, 6,000 hosts and 15,000 containers — all while slashing mean time to resolution (MTTR) by over 90%.

Increased demand breeds a higher expectation for reliable, resilient services across Rappi's mobile app and infrastructure. Splunk Observability Cloud is a key component of Rappi's success, helping the team see, understand and act on real-time data in one place. "A single dashboard provides data for engineering, DevOps, site reliability engineering, SecOps, peer engineering and microservices, operations and business metrics," Comisario says. "If something happens at Rappi, and we don't see it on our Splunk dashboard, it's actually not happening at all."

With a more agile approach and real-time observability from Splunk, the Rappi IT team now efficiently manages more than 1,000 microservices, 6,000 hosts and 15,000 containers — all while slashing mean time to resolution (MTTR) by over 90%.





## RENT THE RUNWAY

To keep operations running seamlessly, Rent the Runway relies on dozens of complex services across its multi-cloud architecture to keep tabs on everything. With instant visualizations on Splunk dashboards, teams have a one-stop shop for critical metrics across the company's sprawling environment, which enables them to identify and repair problems before impacting customer experience.

"If a customer can't get a dress or is having issues checking out due to a bug in the experience, we need to know so we can address it quickly," says staff software engineer Shane Ryan. "We've leveled up our monitoring game in the last four to five years with Splunk. We no longer have to wait for our user systems to alert us to issues, as we've seen the value of alerting across our front and backend systems. Now we can get ahead of the issues — and when there is an incident, it doesn't have to be all hands on deck." Before, it wasn't uncommon to need two dozen developers on a call when an incident occurred. Aki Yamada, staff engineer with Rent the Runway from the early days, recalls: "Before we started using Splunk, every resolution was bespoke — logging into production machines to analyze logs and run scripts — but Splunk enables us to answer questions about application history with simple queries."

Now, with full visibility across warehousing and consumer apps, teams can monitor what they need to manage and involve fewer people for incident resolution. The result: increased customer satisfaction — and an improved employee experience.

Splunk has helped teams at Rent the Runway boost MTTR by 94%, prevent unplanned downtime and offer exceptional customer experiences.







Next time you need to put on your detective hat while keeping customers — and yourself — happy, come back to the strategies outlined in this guide. When you're able to easily focus on the issues that impact your customers and your business the most, your team and company will reap the benefits.

To learn more about observability strategies, check out How OpenTelemetry Builds a Robust Observability Practice.

Keep the conversation going with Splunk.





Splunk, Splunk> and Turn Data Into Doing are trademarks and registered trademarks of Splunk Inc. in the United States and other countries. All other brand names, product names or trademarks belong to their respective owners. © 2024 Splunk Inc. All rights reserved.

