Supercharge Your IT Monitoring

With the Three Pillars of Observability

Get the most from your data with metrics, traces and logs

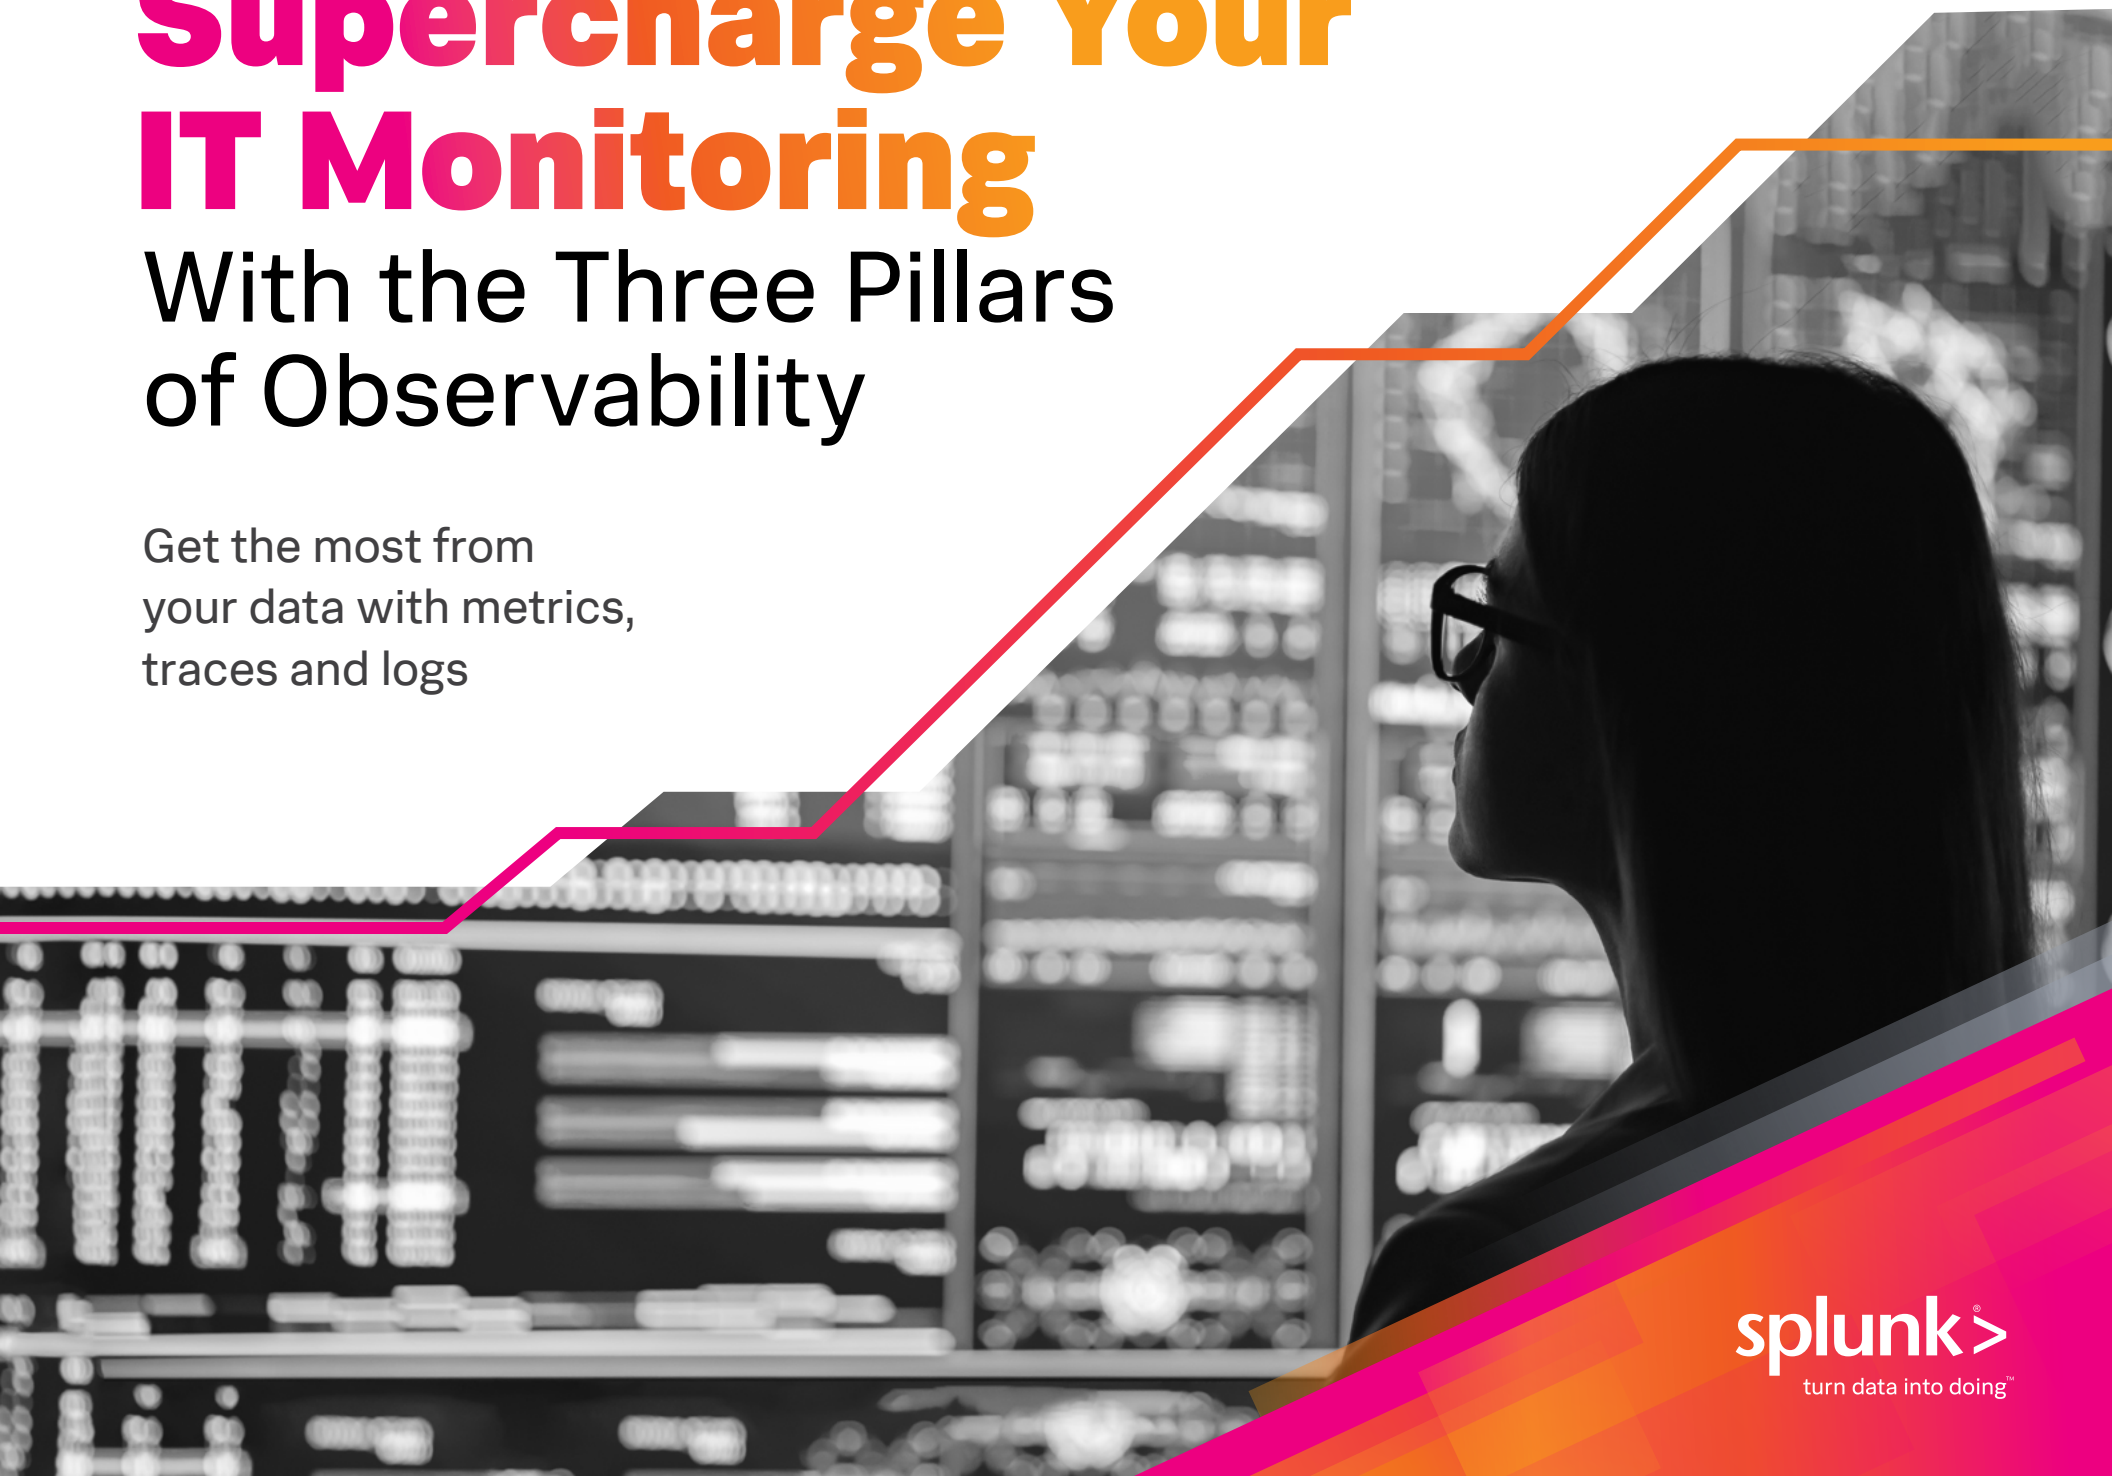

splunk>

turn data into doing™

One of the biggest challenges for any IT organization is keeping ahead of the changes in the monitoring landscape. You need to make sure you're using the best tools and tactics to safeguard your infrastructure and applications. You need to evaluate new technologies and separate the useful from the hype. Analysts and vendors introduce new terms all the time, and keeping track of them can be a big distraction.

At this point, you may be familiar with the concept of observability. You may have heard that it can help you get insights into your applications and infrastructure in real time. You may already be using observability tools, or are considering the best way to adopt them. We will discuss three fundamentals of observability — metrics, traces and logs — how they work together and how they can help you get started on the journey to full observability.

Metrics, traces and logs can help you address three of the most vital issues for data-driven, digital organizations:

- Complexity
- Cost
- Customer experience

Complexity

Cloud adoption, cloud security and growing infrastructures mean massive amounts of data that are impossible for human operators to manage and analyze. Detecting underlying issues and pinpointing root causes is difficult and takes too much time.

Modern IT departments use a wide variety of tools, often purchased from different vendors at different times, that track events and generate data in different formats. It can be an enormous challenge to correlate all of this data and derive insight from it.

Complex tool sets can also force IT to monitor with one tool and troubleshoot with another. Some tools can only do metric-based monitoring, others can only do logging. Different teams are forced to use different tools and oftentimes these tools don't integrate, or may use an entirely different language altogether, creating even more silos and discouraging collaboration.

Cost

Understandably, many organizations turn to the cloud for the flexibility it offers. But that can lead to growing budget requests and more trouble with procurement as cloud spending rises. Worst of all, a lot of that spending is unnecessary. If you don't have full visibility into your entire cloud stack, you may waste capacity on abandoned projects and inefficient utilization, running up bills you shouldn't have to pay.

Customer experience

No matter how organizations evolve, the definition of customer experience is still based on fundamentals. Outages and performance degradations are the enemy, especially when they cause breaks in business continuity. The cause is often outdated, insufficient monitoring solutions. You don't have the time to wait hours to identify, troubleshoot and fix an issue — not if you want to stay in business. In short, IT teams need to work in real time, solve issues quickly and move on.

IT Departments Need a Single Solution

And that solution should offer holistic monitoring across on-prem, hybrid/multicloud environments, leveraging all data from any source and at any scale. It's a simple fix, but one that vexes IT teams struggling with legacy systems.

Every company selling an IT monitoring solution talks about their ability to analyze data, but the differentiators are in the details. It's not just about analyzing data. It's about what data and where it comes from.

It's Time to Get Serious About Observability

Observability has been called everything from a tech buzzword to a "monitoring-on-steroids" must-have. The truth is more involved — especially given the increased complexity of the modern infrastructure and the undisputed need for better monitoring higher in the stack, and deeper in the system.

Teams requiring operational visibility have expanded beyond sysadmins and ITOps analysts — even developers are taking greater ownership of knowing what's going on for a better customer experience. To effectively do this, all roles need visibility inside their entire architecture — from third-party apps and services to their own — to fix and eventually prevent problems. When that capability is built-in — the premise of observability — it makes visibility easier, enables greater insight and leaves more time for more strategic initiatives.

Distributed systems create problems for IT departments tasked with seeing the big picture because each node in a system may have different owners with different requirements and different priorities. Observability solutions allow the IT department to get all the relevant data they need without relying on other teams to supply it.

Different teams inside an organization use application information differently:

Development teams want to see how their applications are performing based on how users interact with them in real time.

DevOps teams need to deploy code quickly, keep it up-to-date and track changes.

By implementing the principles of observability into your IT monitoring solutions, you can tailor the outcomes to match the needs of every user.

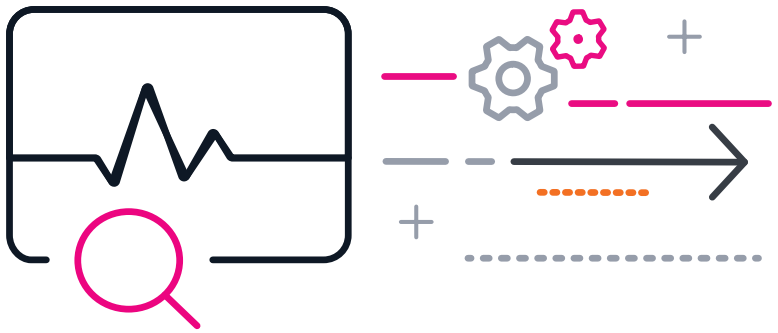

Harnessing the Power of Metrics, Traces and Logs

Observability is based on three types of telemetry data: metrics, traces and logs, often called the “Three Pillars of Observability.” Individually, they can provide you with information to pinpoint issues and root causes, but when taken together their power is significantly magnified.

If you go far enough back in the history of computing, you will find that there were no logs, no traces and no metrics. The ”application” pretty much took over the system and ran to completion — or not.

But even back then, there was a concept for finding problems within systems, called “wolf fencing,” formalized around 1982. The name came from an idea for finding a wolf in Alaska. First, you build a fence across the middle of the state, listen for the wolf to howl and determine which half it’s in. Take that segment, split it in two and repeat until you find the wolf.

A troubleshooter using wolf fencing to identify an issue in an application might add a statement that reads, for example, “I made it to line 148.” When printed, that was essentially a clue that said, “We don’t know where the issue is yet, but we know it’s not before line 148.”

Formalizing this as a standard method of operations led to the concept of logs — which reveal what is causing an issue — and, in turn, to new ways of debugging.

As our environments and applications have increased in complexity, metrics and traces were added to logs to form a more comprehensive picture.

It’s All About the Data

Three Pillars of Observability

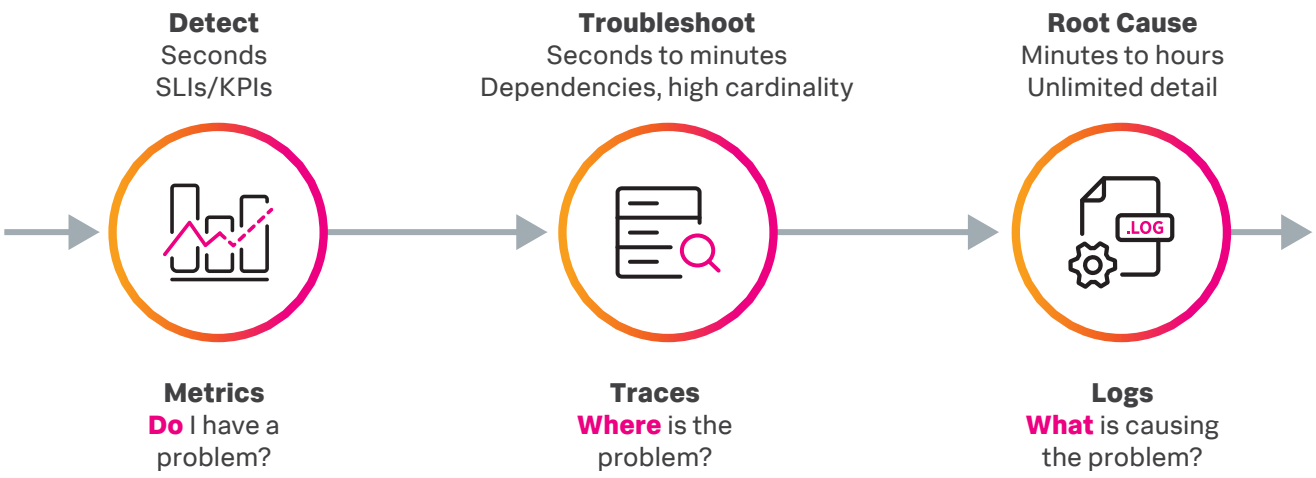

Debugging for complex systems is an iterative process

- First, start with a high-level metric.
- Then, drill down and detangle based on fine-grained data and observations.
- Finally, make the right deductions based on the provided evidence.

Everything is an event, but an event is not everything

First of all, we have to understand that everything that happens can be considered an event. If it's recorded, it's an event. If it wasn't recorded, it didn't happen. Metrics, traces and logs are all events that overlap, but they provide different types of information that, taken together, paint a complete picture.

Modern applications provide such a complex array of information that it can be difficult to even know what to look at. There are too many interconnected components. Only by bringing together metrics, traces and logs can you determine where to look to pinpoint an issue.

Metrics

Metrics help to answer some of the most fundamental questions that the IT department faces. Is there a slowdown in system performance that's affecting customers? Are employees having trouble logging in? Is there an unusually high volume of traffic? Is our rate of customer churn going up?

Metrics are numerical data points captured over time that can be compressed, stored, processed and retrieved far more efficiently than events. You can easily correlate metrics data with other event data to be alerted to what just happened (metrics), and why (logs). Metrics can be considered the most valuable of the three pillars, if taken individually, because they are generated more frequently and by everything, from operating systems to applications. Because metrics come from so many sources, correlating them can provide a more complete view of an issue.

Metrics can come from servers, applications, IoT sensors or just about any machine data-generating object containing numerical, time-series data points. Common examples of metrics that you may be familiar with are system measurements like CPU, memory or disk space; infrastructure measurements from AWS CloudWatch; and measurements from IoT devices like temperature readings or GPS location (e.g. latitude-longitude pairs over time).

Metrics differ from log data in that they can be stored and optimized more efficiently for querying. They don't contain the rich information of a log, but they do present a specific measurement of a system over time.

Common metrics include:

- System metrics (CPU use, memory use, disk I/O)
- App metrics (rate, errors, duration)
- Business metrics (revenue, customer signups, bounce rate, cart abandonment)

Traces

Traces do exactly what the name suggests — they trace the path of an event through the network. A trace can help you identify where an event regularly occurs or a bottleneck is occurring. If customers are having trouble logging in, for example, a trace can find the database preventing them from getting access.

Traces help pull together the data provided by metrics and logs for a more complete picture of a system's performance over time. In a modern distributed IT environment that includes containerized applications and microservices, a request or action can travel through a variety of systems. A trace incorporates all of the information to create a map of the journey and what occurred along the way.

A single trace typically captures data about:

- Spans (service name, operation name, duration and other metadata)
- Errors
- Duration of important operations within each service
- Custom attributes

They All Overlap

Logs can yield metrics

Traces can yield metrics

But you need all three

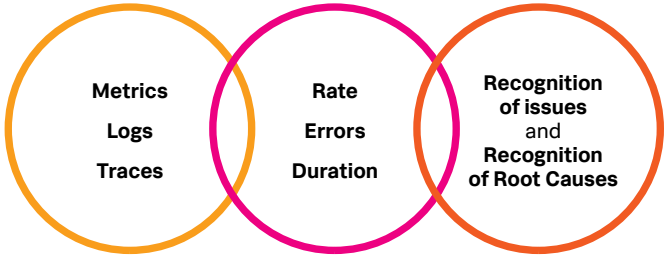

While all three data types are treated separately, they actually all overlap. Logs can yield metrics. Traces can yield metrics. Metrics can point you to the right trace or the right log.

With metrics, traces and logs, you can recognize issues faster to find the underlying root cause.

Logs

Logs are system-generated records of events that happen within an application. Modern IT systems generate reams and reams of log files tracking everything that happens. IT monitoring systems like Splunk were built on the capability to analyze log data and use it to identify and troubleshoot system problems and prevent them before they happen. They provide more information and context on why a problem occurred, rather than just the data that identifies the event.

The challenge of using logs to identify and remediate issues is one of volume; so many systems generate so much log information that finding the most important clues can be difficult. There are multiple log formats used by multiple systems in an organization's network, so the path to an issue resolution can be obscured.

Log data can include:

- System and server logs (syslog, journald)
- Firewall and intrusion detection system logs
- Social media feeds (Twitter, etc.)
- Application, platform and server logs (log4j, log4net, Apache, MySQL, AWS)

Work for different groups	
An alert on one service (Metric)	SRE, Ops
Leads to a timeout error (Tracing)	DevOps Engineering, SWE
Leads to an infrastructure problem (Metric)	DevOps Engineering, SRE
Leads to a configuration issue (Metric)	DevOps Engineering, Ops
Leads to a memory leak in an app (Log)	Developer/SWE

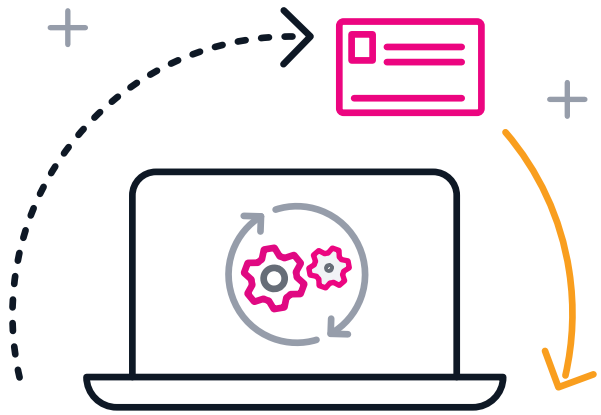

The Power of Three

For some organizations, metrics may provide enough information for them to do the bulk of their troubleshooting. For IT departments with simple infrastructures that generate comparatively small amounts of data, they may get what they need from metrics alone. And not every organization needs to identify problems quickly. Traditional businesses with little or no digital presence other than an informational website don't need to instantaneously identify and troubleshoot customer-facing issues.

But if you're reading this, you're most likely responsible for system performance in a fast-moving, digitally-driven business. For you, combining metrics, traces and logs is another major step forward in the modernization of IT.

Get more effective insights from your data

Just as modern monitoring solutions alerted you to issues before they happened, an observability solution that combines metrics, traces and logs can bring you better and more complete insights faster, making your job even easier.

Scale and grow at the speed you need

Observability is more than just a convenience. Forward-thinking companies are embracing observability as a competitive advantage that helps them scale and grow at the rapid pace of digital transformation. To be prepared, they're making future-proof investments on their monitoring tools.

Correlate data from your entire network — including containers and microservices

If you're using containers and microservices, the benefits to you are even clearer. The additional context provided by tracing is vital for troubleshooting across hybrid environments, as it provides significantly more information about where in a distributed network the issue may have occurred.

Accelerate your move to the cloud

Finally, if you're moving to the cloud, you can expect another quantum leap in the amount of system-generated data you will need to monitor and understand. Observability is tailor-made for the cloud.

There's a reason why observability is one of the most talked about topics in IT. If you haven't considered how to implement observability into your environment, the ideal first step is to consider how your system uses metrics, traces and logs.

Why Splunk?

Splunk is the industry's **only analytics-driven, multicloud monitoring solution** for all environments. Splunk gives you the speed, scale and insights you need to master your IT challenges. Splunk is the only solution that allows you to:

- Pinpoint root cause with speedy real-time troubleshooting, find and fix issues as soon as they arise and get to resolution in just seconds
- Correlate data from multiple data sources in multiple formats from multiple tools and get actionable insights in one solution
- Get insights fast from across your entire environment, whether you're on-prem, hybrid, multicloud or using containers and microservices
- Get started quickly and easily with hundreds of out-of-the-box integrations, pre-built charts and dashboards, and automatic service discovery
- Drive value faster for your organization with easy-to-use, high-performance monitoring and troubleshooting that enables you to quickly detect and resolve issues
- Reduce cost and complexity by consolidating monitoring tools and standardizing on the market-leading data platform
- Future-proof your investment with a comprehensive, scalable and flexible data-driven solution that grows with your organization
- Analyze and correlate data for valuable insights that reduce event noise and predict future degradation

Splunk provides the most comprehensive, robust and flexible troubleshooting and monitoring solution across on-prem, hybrid/multi-cloud environments at any scale.

Splunk is recognized by leading analysts as the best solution in the market for ITIM and ITOM, as well as a leader in cloud observability solutions.

The world's leading organizations, including 90% of the Fortune 100 rely on Splunk.

The Most Comprehensive Set of **Observability Capabilities**

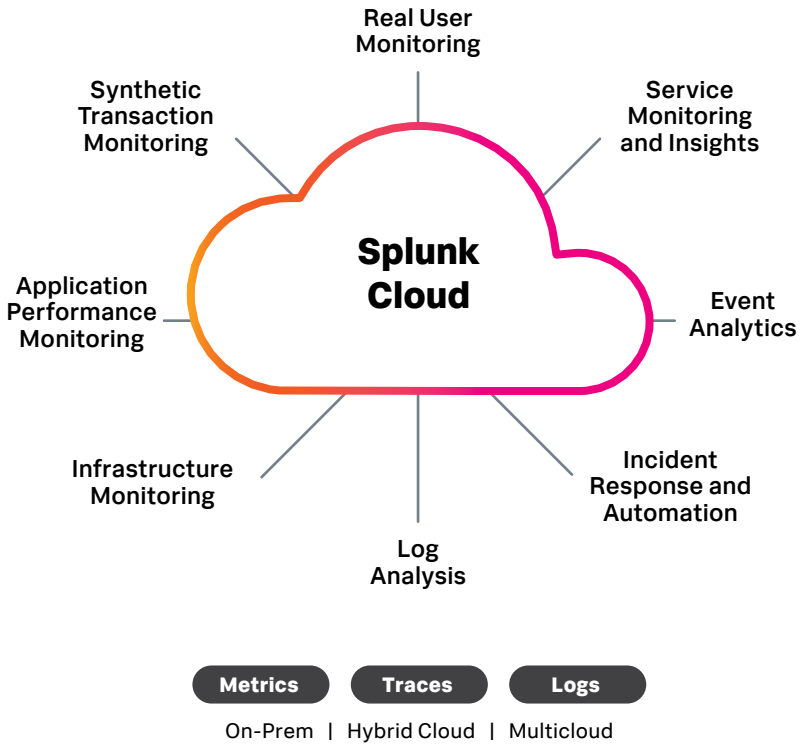

Customer Story: Quantum Metric

As more industries realize the need for accelerated digital transformation, more diverse companies are flocking to Quantum Metric to maximize their potential. For the 2021 unicorn, an influx of customers meant an even bigger influx of data — and an increasingly complex engineering environment that includes everything from Kubernetes clusters to Docker engines.

"We want to do the same thing as our customers: build better products, move quickly, iterate, be able to experiment — and do so safely," says Brent Miller, senior director of cloud operations. "Solving a problem across different use cases requires an observability solution that is extensible and robust enough to meet those use cases without forcing us down one path," adds Eric Irwin, director of engineering.

Quantum Metric needed a flexible observability solution that would help them and their customers build better products, faster. That's why they chose the Splunk Observability Cloud.

Thanks to full-fidelity ingestion of metrics, traces and logs, the team now sees what's happening across their infrastructure and applications — insights that otherwise would be impossible to account for. With full-stack, end-to-end visibility, they can be sure that demo sites are running and understand how their services are working together to deliver value to their customers.

Data-Driven Outcomes

$80k

saved by switching to Splunk, thanks to better downsizing analysis and capacity planning

96%

faster application development, increasing developer productivity

95%

reduction in pending CI jobs due to better assessment of capacity needs

Conclusion

There is a reason you're always being told you need to stay on top of your data. The reason is simple, and you probably know it already. The volume of data you need to monitor and understand as an IT professional will continue to grow and grow. The pace of change will only increase. The key to your value as an IT professional lies in your ability to stay ahead of these changes and help drive business success. Observability provides the best technology to help you make your tasks as easy, efficient and effective as possible.

To recap:

- Metrics, traces and logs provide us data on the operation of our infrastructure, services and applications.
- Each piece fills a unique need for different use cases.
- Each piece works together to give us the complete picture.
- Together, they give us the holistic visibility for monitoring, analysis and response to changes in our environments.

Ready to get the most out of your logs, metrics and traces? Jumpstart your modernization journey today with a free Splunk Observability Cloud trial.

21-19531-Splunk-Supercharge Your IT Monitoring With the Three Pillars of Observability-EB-106

splunk> turn data into doing™