Infrastructure Monitoring 101

The Power to Predict and Prevent

The ability to see what's happening across an organization's infrastructure helps teams to predict and prevent outages

splunk> | aws

Infrastructure Is About More Than Keeping the Lights On

Customer experience — often the frontend of a mobile, web or business application — has become one of the most important metrics for success for global organizations. These experiences rely on layers of interconnected technologies that work together to deliver information, transactions and interactions to an end user. As the experiences grow in complexity, so does the technology.

Apps and services are expected to work quickly and seamlessly on any number of devices, from different kinds of networks and in different locations around the globe. Supporting a connected experience — one that is secure and personalized, constantly improving and with little-to-no downtime — requires many different interconnected technologies to function in concert. Each of these technology layers emits volumes of data that contain the information required to monitor, troubleshoot and ultimately improve those experiences — issues or outages are resolved as soon as they arise.

For many years, IT teams were monitoring pieces of infrastructure separately, but that method creates silos and isn't scalable, and it certainly isn't practical. The advent of microservices, serverless architecture and cloud computing has given us improvements in efficiency, but has also introduced new kinds of considerations to IT infrastructure, and new monitoring challenges.

Having knowledgeable teams and sophisticated systems in place are only part of the job; they have to accommodate the business's need for rapid, constant change while keeping the systems in good operating order. That can only be done with a solution that provides a holistic view and that can scale with the business. Something that can help ITOps teams see the big picture and dig into the details when it's required.

Complex IT Infrastructures Are More Likely to Fail

When someone taps their way through an app, they rarely (if ever) consider the various technology stacks that work together to make that experience possible. What does the complex web of IT infrastructure look like?

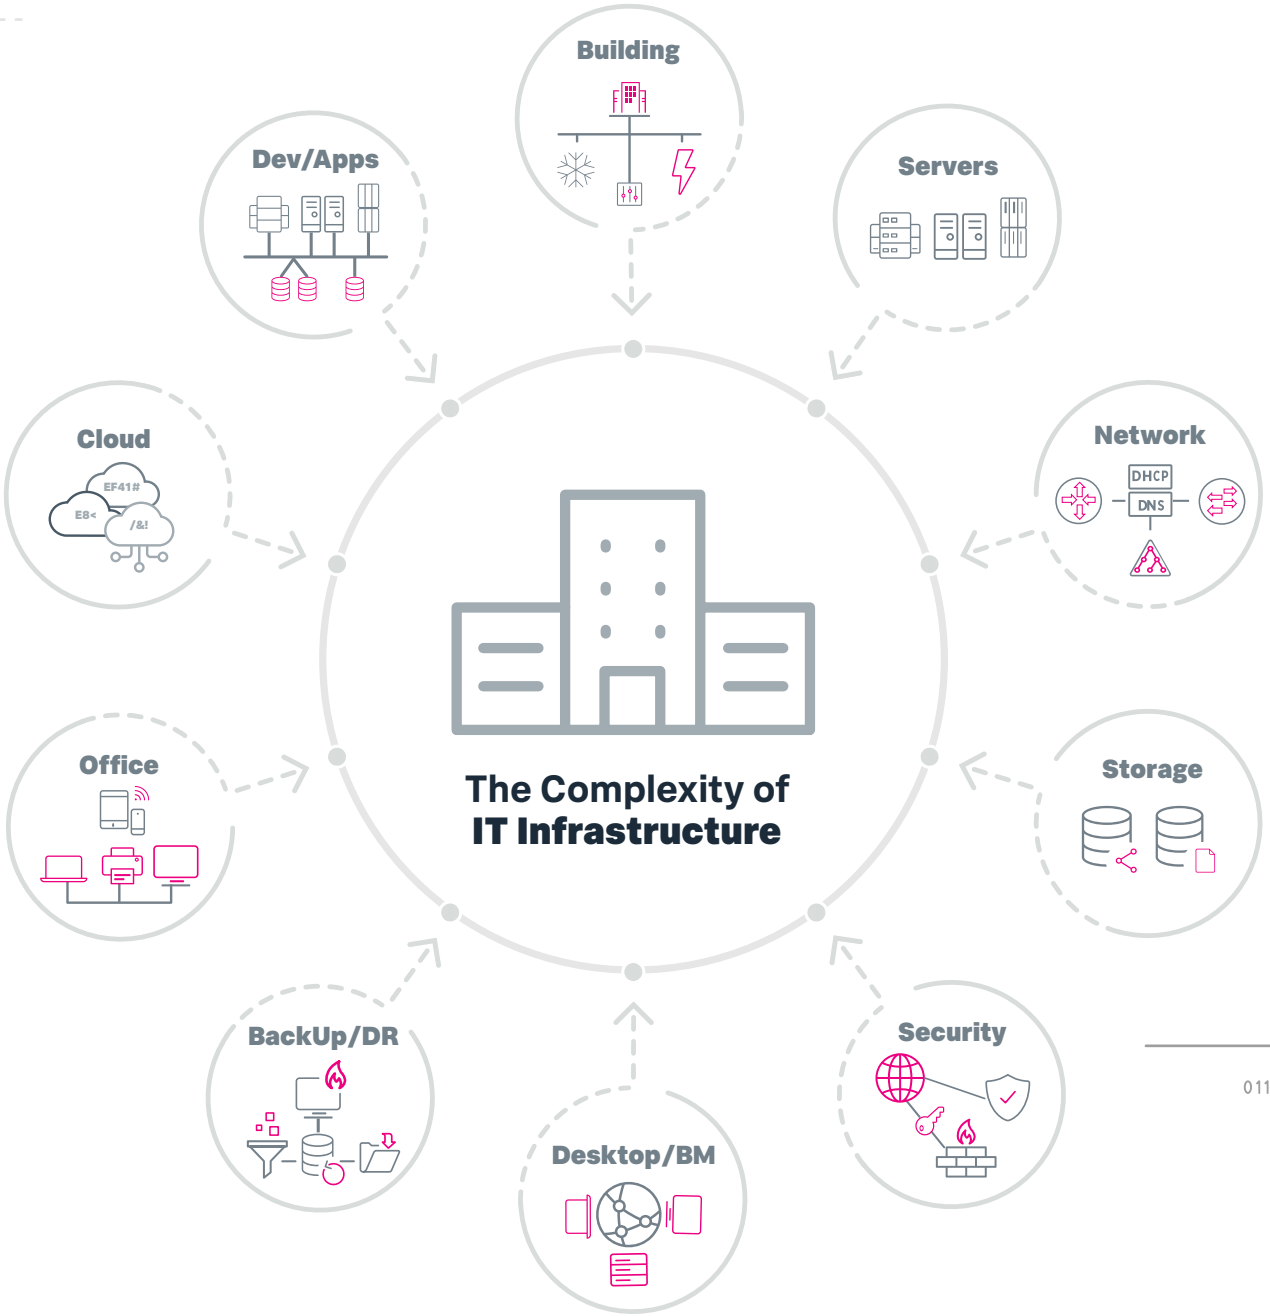
Building
Dev/Apps
Servers
Cloud
Network
DHCP
DNS
The Complexity of
IT Infrastructure
Office
Storage
BackUp/DR
Security
Desktop/BM

More complexity = more room for failure

As we see in the preceding graphic, modern IT infrastructure is an extraordinarily complex system of interconnected technologies, each of which has the potential to run into issues or fail outright. And with more components being added to these stacks as technology evolves, new opportunities for outages arise. In fact, between 2017 and 2018, instances of outages or "server service degradation periods" increased from 25% to 31%, and if we look at on-premises data centers, that number rises to 48%.*

One troubling fact revealed in Uptime Institute's 2020 datacenter survey is that only about half of organizations actually calculate the cost of a downtime incident. This number is trending up, probably as a result of the cost impact and publicity that results from service interruptions.

60% of data center's outages could have been prevented with better management, processes or configuration.

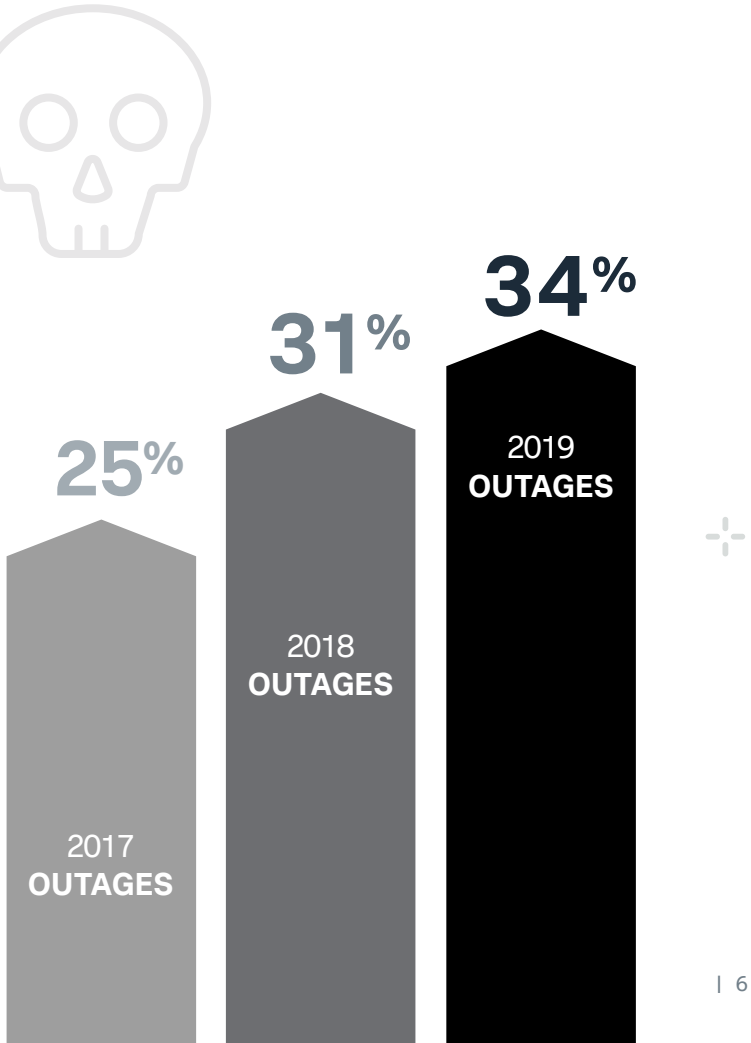

Seventy-eight percent of organizations say they had an IT service outage in the past three years — a higher percentage than in previous years — and 41% classified it as minimal or negligible. Outages in these categories signal bigger problems and are troubling more for their frequency than for their singular impact. When asked about significant, serious or severe outages — which can cause substantial financial and reputational damage — 31% have been affected.**

About 20% of organizations had a serious or severe outage in the past three years — that is, an outage that was costly, caused reputational damage and, in some cases, had major other implications. Nearly a third of all outages cause financial or reputational damage.

Outages are increasingly costly for businesses. In 2020, a greater percentage of outages cost more than $1 million (now nearly one in six rather than one in ten, as in 2019), and a greater percentage cost between $100,000 and $1 million (40% vs. 28%).

Another angle on that statistic: because of largely preventable errors, almost half of employees and users experienced issues with their apps and services. That kind of disruption can result in thousands of employee hours wasted, customer dissatisfaction and, ultimately, loss of business.

Companies of all sizes can be affected by these outages, and because companies frequently rely on each other's infrastructures for their products and services, there is a cascading effect throughout the connected systems when one service goes out. For example, these three major outages occurred in 2020:

- **U.S. mobile operator T-Mobile suffered a major outage on June 15.** Customers reported problems with their mobile phone service, mobile internet connection and their ability to text friends and family. Reports came in over a period of almost ten hours that evening, peaking with 113,980 reports in one 15-minute period.
- **On August 25,** Slack suffered a service outage that affected users in the U.K. and western and southern Europe. Slack users faced troubles with files, messages and connecting to Slack.
- **And on September 23,** Tesla suffered an hour-long global network outage with its internal systems that left several Tesla owners unable to connect to their cars through the mobile app or the website. Tesla's energy products, Tesla solar, and Powerwall home battery systems were inoperative too. The outage was due to an internal break of their application programming interface (API).

The best way to ensure that issues are resolved quickly — or prevented altogether — is to monitor and troubleshoot the underlying infrastructure as well as the mission-critical apps and services that run on them. While observing any one element of the infrastructure stack is a straightforward proposition, observing each piece individually introduces a host of additional problems.

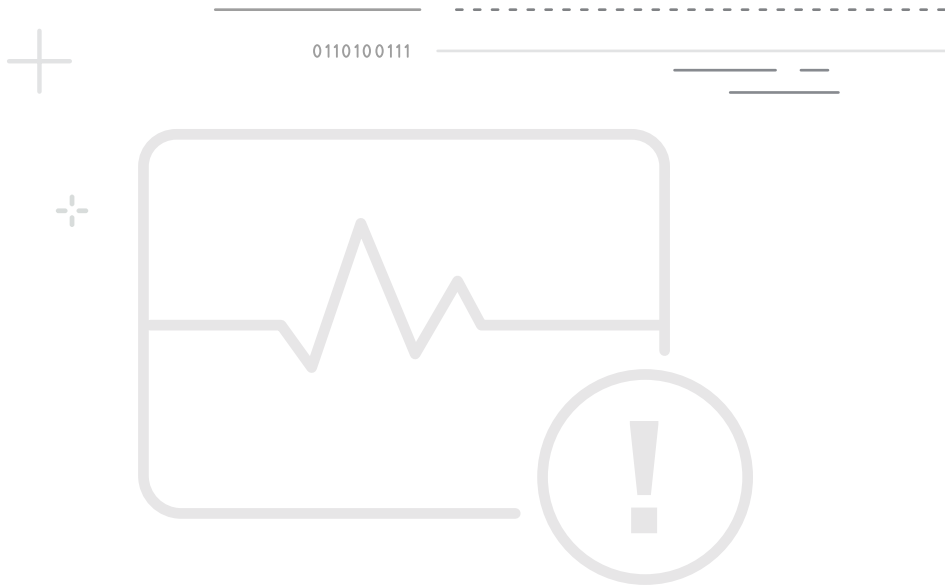

* Source: Uptime Institute 2018 (8th annual Data Center Survey)
** Source: Uptime Institute 2020 (10th annual Data Center Survey)

How to Cut Through the Fog
Better Manage Infrastructure Monitoring

The visibility problem

We can think of ITOps as a stack of physical and logical layers, each with its own technologies, systems and services, and each with a corresponding team or individual responsible for monitoring and maintaining it. This makes gaining visibility into the infrastructure as a whole fundamentally problematic, despite being essential.

A per-layer monitoring practice leads to siloed teams and incompatible views of data. Each layer has different vital metrics, different monitoring tools and dashboards and different personnel behind the keyboard. In practice, per-layer monitoring means people looking at limited information using different languages, leading to difficulties detecting and investigating outages and issues as well as restoring service.

Different types of data created by IT infrastructure

Analysts including Gartner, Forrester, IDC and Computing UK have all developed their own set of essential data. The following is a list of observable telemetry data that we have found to be critical for monitoring the infrastructure stack. These essential data types can be categorized into three groups:

Metrics

Numbers that give us insights about a process or an activity, or the status of an underlying system, network or storage. Generally, metrics are measured over time — often referred to as a time series.

- **System Metrics** (CPU Use, Memory Use, Disk I/O)
- **App Metrics** (Rate, Errors, Duration)
- **Business metrics** (revenue, customer signups, bounce rate, cart abandonment)

Traces

A trace is the record of the progression of a request through an application, including all of its myriad services.

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**

Logs

Immutable records of discrete events that happen over time. Event logs exist in plain or structured text, or binary.

- **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)

Observability is key to a successful IT monitoring solution

One way to avoid the problems of per-layer monitoring is building with observability in mind. Observability is the natural evolution of what we used to call monitoring. Observability recognizes that today's infrastructure and applications are living, breathing organisms that evolve at a much faster rate than ever before. Observability encompasses all of the things we used to do in monitoring, like watching for known failure conditions, and extends it to support the challenges of today's applications, like being prepared for all the unknown failure conditions.

IT stack layers:

Servers

A high-quality user experience depends on effective monitoring of the systems that support the product. It allows administrators and ITOps personnel to see resource usage patterns and optimize the servers keeping websites and applications running smoothly.

Server operating systems routinely record a variety of operational, security, error and debugging data such as system libraries loaded during boot, application processes open, network connections, file systems mounted and system memory usage. The level of detail is configurable by the system administrator; however, there are sufficient options to provide a complete picture of system activity throughout its lifetime. Having visibility into these pieces of server data and monitoring them proactively can help teams find resolutions more quickly or prevent outages altogether.

Imagine a gaming company whose users depend on reliable, high-speed access to a web app — not terribly hard to picture, is it? Having immediate visibility and insight into server performance would be critical to that company's success. The ability to quickly resolve server-based issues (or predict and avoid them altogether) would have a significant impact on the product's uptime and directly impact customer satisfaction and, ultimately, revenues.

Having a single tool from which to monitor the health of servers — one that correlates event data and log data into a seamless experience — enables ITOps teams to quickly isolate what is driving the failure (like memory usage on a single server) and resolve it. It also facilitates proactivity. The ability to create alerts and automations within the monitoring tool saves teams time and allows them to focus their efforts on other tasks.

Serverless

Serverless computing offers a number of advantages over traditional cloud-based or server-centric infrastructure. For many developers, serverless architectures offer greater scalability, more flexibility, and quicker time to release, all at a reduced cost. Serverless apps are deployed in containers that automatically launch on demand when called.

Under a standard Infrastructure-as-a-Service (IaaS) cloud computing model, users purchase units of capacity, meaning you pay a public cloud provider for always-on server components to run your apps. It's the user's responsibility to scale up server capacity during times of high demand and to scale down when that capacity is no longer needed. The cloud infrastructure necessary to run an app is active even when the app isn't being used.

With serverless architecture, by contrast, apps are launched only as needed. When an event triggers app code to run, the public cloud provider dynamically allocates resources for that code. The user stops paying when the code finishes executing. In addition to the cost and efficiency benefits, serverless frees developers from routine and menial tasks associated with app scaling and server provisioning.

With serverless, routine tasks such as managing the operating system and file system, security patches, load balancing, capacity management, scaling, logging and monitoring are all offloaded to a cloud services provider such as Amazon Web Services (AWS).

Network

While each organization's needs and data sources will vary, there are reasons for monitoring network data that are common across companies and institutions:

- Protecting corporate networks from attacks.
- Providing visibility into network traffic.
- Determining the role of the network in the overall availability and performance of critical services.

Monitoring a network means more than having visibility into the state of the hardware that supports that network, like routers, switches, etc. It includes monitoring network event logs, activities across the network infrastructure, traffic bottlenecks or suspicious behavior.

Virtualization

Virtualization has revolutionized the modern datacenter. Whether it be network, server, application or desktop virtualization, each offers numerous benefits such as cost savings, physical server consolidation, dynamic load balancing, ease of migrations and more. While these benefits are compelling, virtualization has also introduced a new level of complexity to managing the datacenter. Visibility, or a lack thereof, is probably the biggest challenge.

Across virtualized machines, datacenter administrators lack the necessary visibility to help them solve problems faced by their application owners. Capturing and storing all the relevant data at full fidelity is vital to truly understanding application performance, especially when mission-critical applications run in virtualized environments. Visualizing this data within the context of data from other technology tiers is essential to understanding exactly which events in which tier are causing problems and impacting performance. Correlating, trending and analyzing virtualization data and data from other technology tiers such as storage, networks and operating systems is a big data problem.

Gaining insight into virtual deployments and making essential correlations with the applications and other parts of the infrastructure — by monitoring the resource usage on virtual environments like VMware and others — is vital to efficiently managing resources and gaining the benefits of virtualization.

Cloud

Running workloads in a cloud environment is not "set it and forget it." ITOps teams still need to monitor the performance, usage, security and availability of the cloud infrastructure continuously. And with the right solutions, it's possible to manage IT systems and derive actionable insights from all of the data in one system, even if the services are running in hybrid environments.

When an organization migrates its services to a cloud platform (or between cloud platforms), for instance, having end-to-end visibility into every stage of the migration can help teams establish baseline performance, monitor services during the transition and ensure that all services are running optimally after the transition is over.

Services running on hybrid and cloud infrastructures can be opaque, leading to gaps in ITOps teams' understanding of the system as a whole. Organizations eager to get the benefits of cloud often overspend on cloud services — on deprecated or unused services, unknown redundancies or excessive resources. Ingesting all of the cloud infrastructure data into a single environment, replacing the multitude of individual monitoring tools with a consolidated solution, can provide an understanding of how resources are performing and being used, allowing for optimization of utilities and billing.

Public Cloud	**Hybrid Cloud**
AWS	On-Prem
	Private/Public Cloud Mix

Infrastructure monitoring should provide out-of-the-box, end-to-end visibility into all stages of cloud migration — before, during and after — and full visibility into public cloud IaaS. The right monitoring solution will simplify the multitude of monitoring tools, and allow you to monitor your entire stack in one place. Teams can collaborate more efficiently, with greater visibility into resources. Built-in dashboards and accurate alerts provide shorter mean time to detect (MTTD), helping resolve issues before they impact operations.

Kubernetes and Containers

Since the introduction of the concept in 2013, adoption of containers has skyrocketed across technology organizations. They share some conceptual features with virtual machines, but they differ in a few essential ways. The easiest way to understand a container is to think of it as exactly that — a container — a receptacle that holds something securely and can be used to transport its contents. A software container performs a similar function. It allows developers to package an application's code, configuration files, libraries, system tools, and everything else needed to execute that app into a self-contained unit, so that they can move the package and run it anywhere with ease.

Containers enable a number of significant benefits to organizations, developers and users — faster deployment, smaller footprints and consistency across environments. But containers, like virtual machines, have their own system metrics that need to be monitored, and with many containers running side-by-side, the task of monitoring, optimizing and troubleshooting them becomes much more complicated.

Cloud-native infrastructure such as containers, Kubernetes and serverless are highly dynamic and ephemeral. When the cloud infrastructure only lives for minutes, the monitoring solution needs to detect and enable automatic remediation within seconds.

For all the benefits that containers bring to IT organizations, they bring new considerations that must be addressed including:

- **Significant blind spots:** Containers are designed to be disposable. Because of this, they introduce several layers of abstraction between the application and the underlying hardware to ensure portability and scalability. This all contributes to a significant blind spot when it comes to conventional monitoring.
- **Increased need to record:** The easy portability of so many interdependent components creates an increased need to maintain telemetry data to ensure observability into the performance and reliability of the application, container and orchestration platform.
- **The importance of visualizations:** The scale introduced by containers and container orchestration requires the ability to both visualize the environment to gain immediate insight into your infrastructure health but also be able to zoom in and view the health and performance of containers, node and pods. The right monitoring solution should provide this workflow.

A good container monitoring solution enables ITOps to stay on top of a dynamic container-based environment by unifying container data with other infrastructure data to provide better contextualization and root cause analysis.

Learn more about container monitoring in **The Essential Guide to Container Monitoring.**

Layers of visibility

Each layer of the IT stack mentioned in the last section presents its own challenges in regards to visibility, which get compounded when organizations work to monitor the stack as a whole. And monitoring the stack as a whole is essential — it's what supports the development and usage of the applications and drives customer and employee experiences.

Having a solution that provides a holistic view of the infrastructure alongside detailed views of individual components is vital if an organization wants to proactively tackle infrastructure issues and reduce mean time to detection (MTTD), investigation and restoration. It's also an essential piece of future planning; knowing how the infrastructure has performed historically, and how it's performing in real time, provides invaluable insights that reduce complexity when integrating new technologies and building new experiences for users and employees.

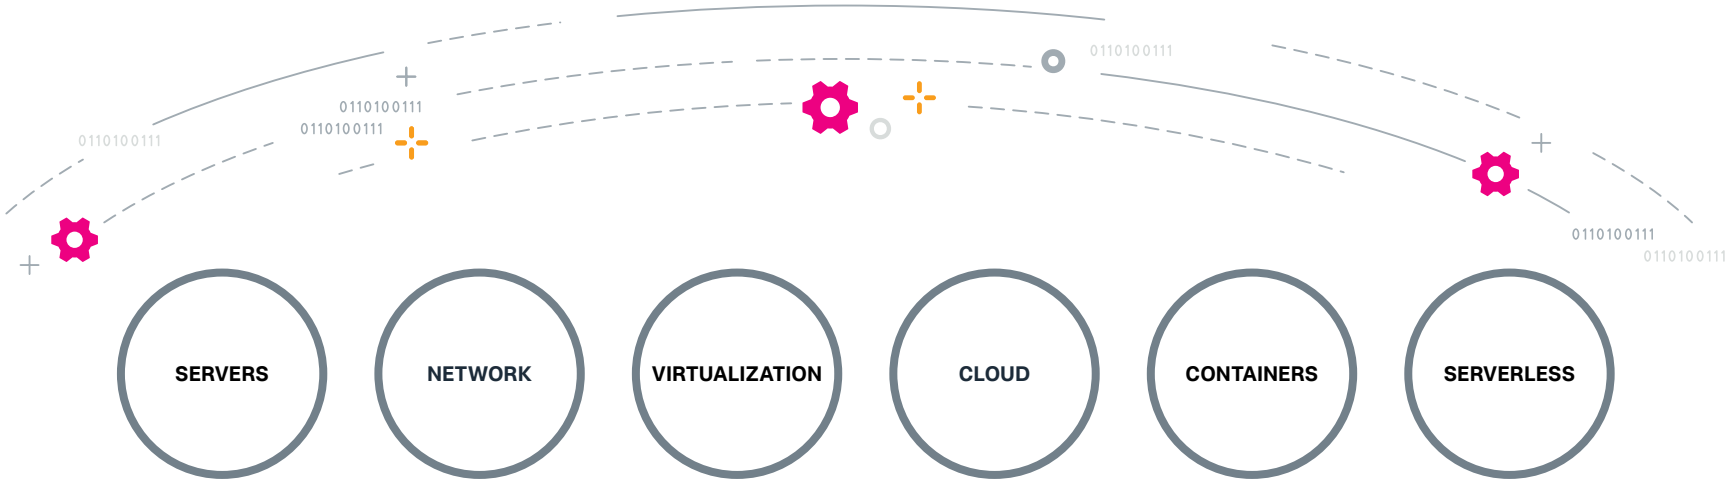

The Importance of IT Infrastructure Monitoring Strategy

Developing an IT infrastructure monitoring strategy will help ITOps teams avoid spending too much time struggling with increasing system complexity and maintaining the tools that were supposed to make monitoring easier and more reliable. To combat these challenges, system administrators and site reliability engineers need a clear view of performance and availability across the infrastructure as a whole.

A strong infrastructure monitoring strategy consists of two key principles:

Centralized and observable data

Separate monitoring tools for each layer of the IT infrastructure are a fundamental issue when it comes to understanding the health of the whole system and solving any problems that arise within it. The answer to the problem is to have a single tool that ingests all of the data and provides onboard correlation and alerting functionality.

A single platform with a unified experience that provides ITOps with access to all the information across domains opens up opportunities for cross-functional investigation and holistic end-to-end infrastructure monitoring. It removes blind spots from the system and, as a result, reduces mean time to resolution (MTTR) because teams can more quickly identify the problem, fix it and move forward.

AI/ML enabled

The volume, velocity and variety of new data must be managed by the right solution. Adding AI and ML to an infrastructure monitoring tool unlocks powerful opportunities for the ITOps team. ITOps can use AI and ML to replace standard monitoring procedures and use predictive algorithms to tackle problems before they arise.

The biggest benefit of an AI/ML-powered monitoring system is the enormous savings in time and effort on the part of ITOps teams. When repetitive tasks and processes are automated, ITOps teams have the bandwidth to do the kinds of things AI and ML are ill-equipped to do — creative problem solving, upgrading existing technologies and planning for the future.

The Full Monitoring Stack

IT infrastructure monitoring frees ITOps teams from the crush of reactive monitoring and crisis management. With critical insights into the systems they rely on, ITOps teams benefit from increased observability into business operations and previously “dark” data.

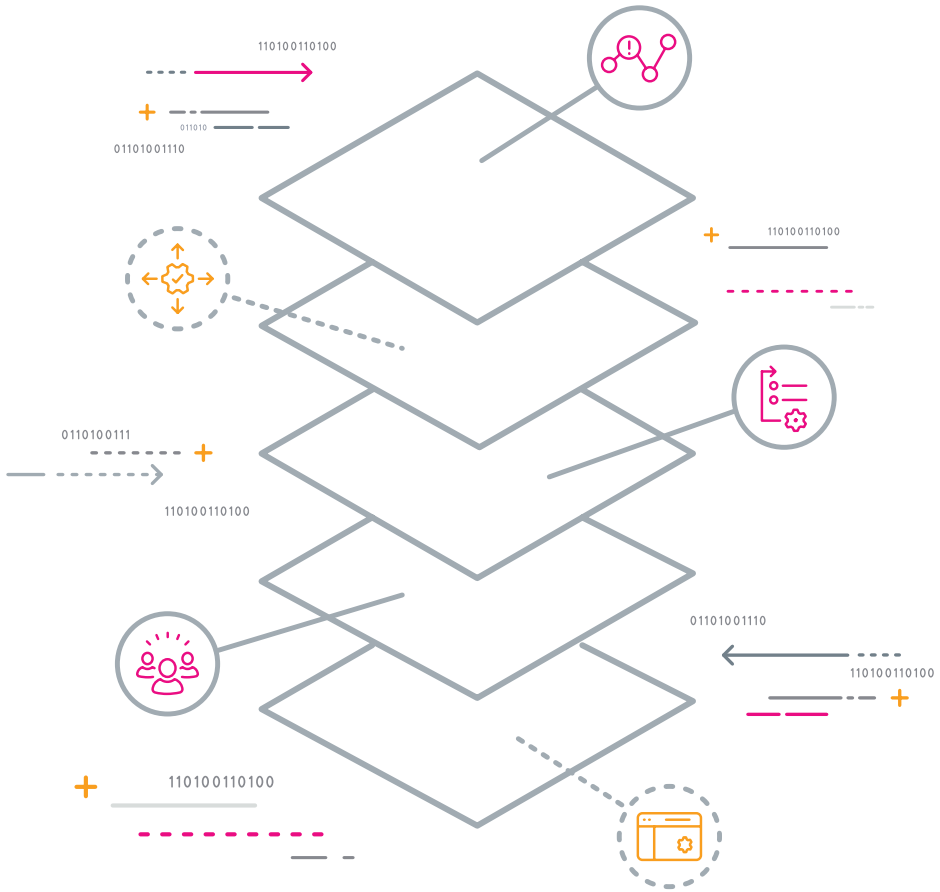

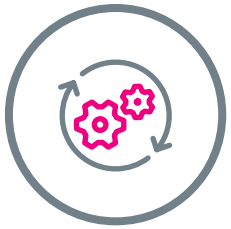

APM

While infrastructure monitoring can help identify if there is a problem, application performance monitoring helps teams locate where the problem is occuring. APM tools are designed to ensure applications provide the right level of service without interruption. Application speed and uptime — for internal enterprise apps and consumer apps — is directly tied to an organization's profitability. Knowing where in the environment an outage originates can result in much faster incident resolution, reducing the consequences of the outage by a significant margin.

NPMD

Improvements to infrastructure monitoring provide new opportunities for network administrators — especially in network performance monitoring and diagnostics (NPMD). With a more complete view of the infrastructure that supports the network, sysadmins can improve their mean times to detection (MTTD), investigation and resolution. And with the implementation of AI and ML, they can use predictive analytics to prevent or minimize outages, altogether.

AIOps

AIOps takes the concept of AI and ML features and expands on it. Rather than having specialized functions for these intelligent systems, AIOps brings AI and ML to every user and every IT use case, so that nearly any function across the business can leverage AI to get ahead. With data and intelligence from the entire infrastructure stack informing employee decisions across the organization, new opportunities and efficiencies become possible.

Observability

Observability goes beyond just monitoring. Customers use point tools to monitor and investigate performance issues across different stages of the cloud journey, leading to fragmented operational data trapped in siloed tools. Manual correlation of vast amounts of operational data from various sources is neither accurate nor in real time, leading to higher MTTD and MTTR. Bringing together infrastructure monitoring, APM, and logs gives teams the true end-to-end visibility they need to ensure system performance.

Customers Who've Succeeded With Infrastructure Monitoring

Acquia

Acquia helps companies build digital customer experiences. As its user base grew, the company needed better insight into its customers' instances and quicker access to data it could trust. Acquia turned to Splunk to monitor its growing AWS environment. The solution enables Acquia's engineering team to release code faster and more reliably, helps technical support teams troubleshoot issues in real time, and even gives Acquia's customers the ability to directly monitor the capacity of their own services. By relying on Splunk, Acquia has reaped a variety of benefits, including:

- Shorter mean time to resolution (MTTR)
- Fewer disruptions
- Slashed support times and overall happier customers
- Less burden on Acquia's technical team

Learn more about **Acquia's cloud monitoring success**.

Acquia

Namely

Namely is an all-in-one HR solution with data-driven analytics that give companies incredible insight into how to best manage their people. Since reliability is a critical requirement for HR solutions, Namely needed a monitoring tool that would help ensure its clients experienced seamless performance while processing essential payroll, benefits, HR and time management transactions.

Splunk provides Namely with real-time monitoring across its advanced microservices architecture. This has allowed Namely to:

- Accelerate product development with confidence
- Develop more advanced features
- Focus the engineering team on enhancements to the Namely platform, providing its clients with a first-class product for building better workplaces

Learn more about **Namely's microservices monitoring success**.

Namely

Imprivata

Imprivata, the healthcare IT security company, provides healthcare organizations globally with a security and identity platform that delivers ubiquitous access, positive identity management and multifactor authentication (MFA). Imprivata enables healthcare securely by establishing trust between people, technology and information to address critical compliance and security challenges while improving productivity and the patient experience. Migrating to Splunk Cloud, Imprivata has seen benefits including:

- DevOps teams can focus on high-priority business needs.
- Streamlined security compliance.
- Avoiding the cost of massive on-premises storage infrastructure.
- Disaster recovery and business continuity of critical Splunk services.

Learn more about **Imprivata's container monitoring success**.

imprivata

CloudShare

CloudShare provides cloud-based solutions that make it easy for application professionals to work in the cloud. Users can efficiently create virtual machine environments, collaborate with others and deploy projects into production, with no background in IT required. The firm needed a way to collect and correlate critical performance and business metrics from thousands of virtual servers. Since deploying Splunk Enterprise and the Splunk App for VMware, CloudShare has seen benefits including:

- Increased customer conversion and retention rates.
- Improved capacity planning based on a better understanding of usage patterns.
- End-to-end visibility and correlation of business and operational data.

Learn more about **CloudShare's virtualization monitoring success**.

cloudshare

Splunk Infrastructure Monitoring

Splunk Infrastructure Monitoring is the most comprehensive, flexible and scalable infrastructure monitoring solution for your entire IT landscape — on-prem, hybrid/cloud — leveraging data from any source, at any scale, in real time. IT teams can deliver on ever-increasing customer expectations by avoiding even seconds of downtime.

What does Splunk monitor?

Any environment at any scale

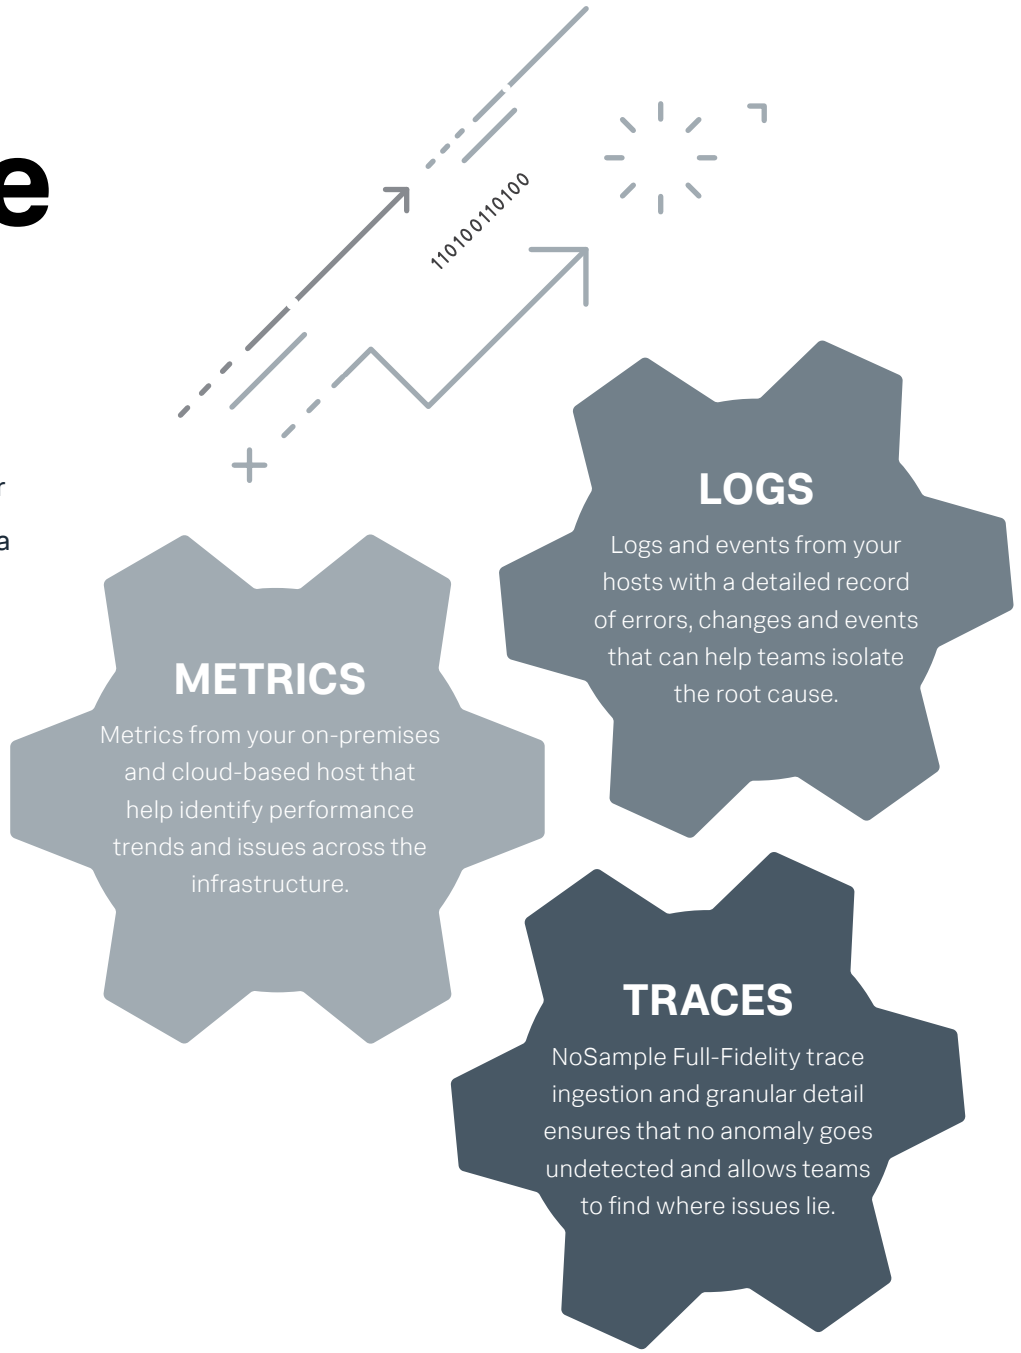

Advanced AI-driven alerting for faster triage

Find and fix performance issues in seconds before they impact end-users. With built-in data science, and a comprehensive library of data science-driven functions, Splunk Infrastructure Monitoring instantly and accurately alerts on dynamic thresholds, multiple conditions, and complex rules to dramatically reduce mean time to detect (MTTD). Alert preview helps prevent alert storms.

Instant visualizations for real-time monitoring

Built on a streaming architecture, Splunk Infrastructure Monitoring lets you interact with your data in real time. Whether built-in or customized, charts and dashboards update in real time with the metrics that matter most to you — instead of waiting minutes, if not hours, with most batch querying monitoring tools. See a live heatmap of your entire infrastructure in one unified view, for full-stack visibility.

Enrich infrastructure data with service context

Combine infrastructure data with data across your entire environment for a holistic view of IT and business performance. Send infrastructure data from Infrastructure Monitoring directly into Splunk IT Service Intelligence (ITSI) to search and analyze across multiple layers of the IT stack. Deep-linking to Splunk Cloud enables in-context monitoring and investigation. Understand interdependencies with built-in correlation with all APM traces for faster troubleshooting.

Splunk Infrastructure Monitoring

Infrastructure teams spend too much time struggling with system complexity and the tools that were supposed to make monitoring easier. To combat these challenges, system administrators and site reliability engineers need a clear view of infrastructure performance and availability.

Splunk Infrastructure Monitoring is the industry's most powerful analytics-driven cloud monitoring and investigation solution for all environments.

Get up and running in minutes

Get started today with a free trial of Splunk Infrastructure Monitoring and eliminate fragmented operational data trapped in siloed tools so you can deliver on ever-increasing customer expectations by avoiding even seconds of downtime.

Want to learn more? Read the Splunk Infrastructure Monitoring Product Brief

About Splunk and AWS

Splunk and AWS are uniquely positioned to help organizations achieve digital transformation success with their highly integrated data-driven cloud adoption and modernization offerings. With industry-leading cloud infrastructure from AWS, combined with the Splunk® Data-to-Everything™ Platform, companies can innovate with confidence, migrate and modernize existing environments, and scale without limits, with data at the center of every business outcome.

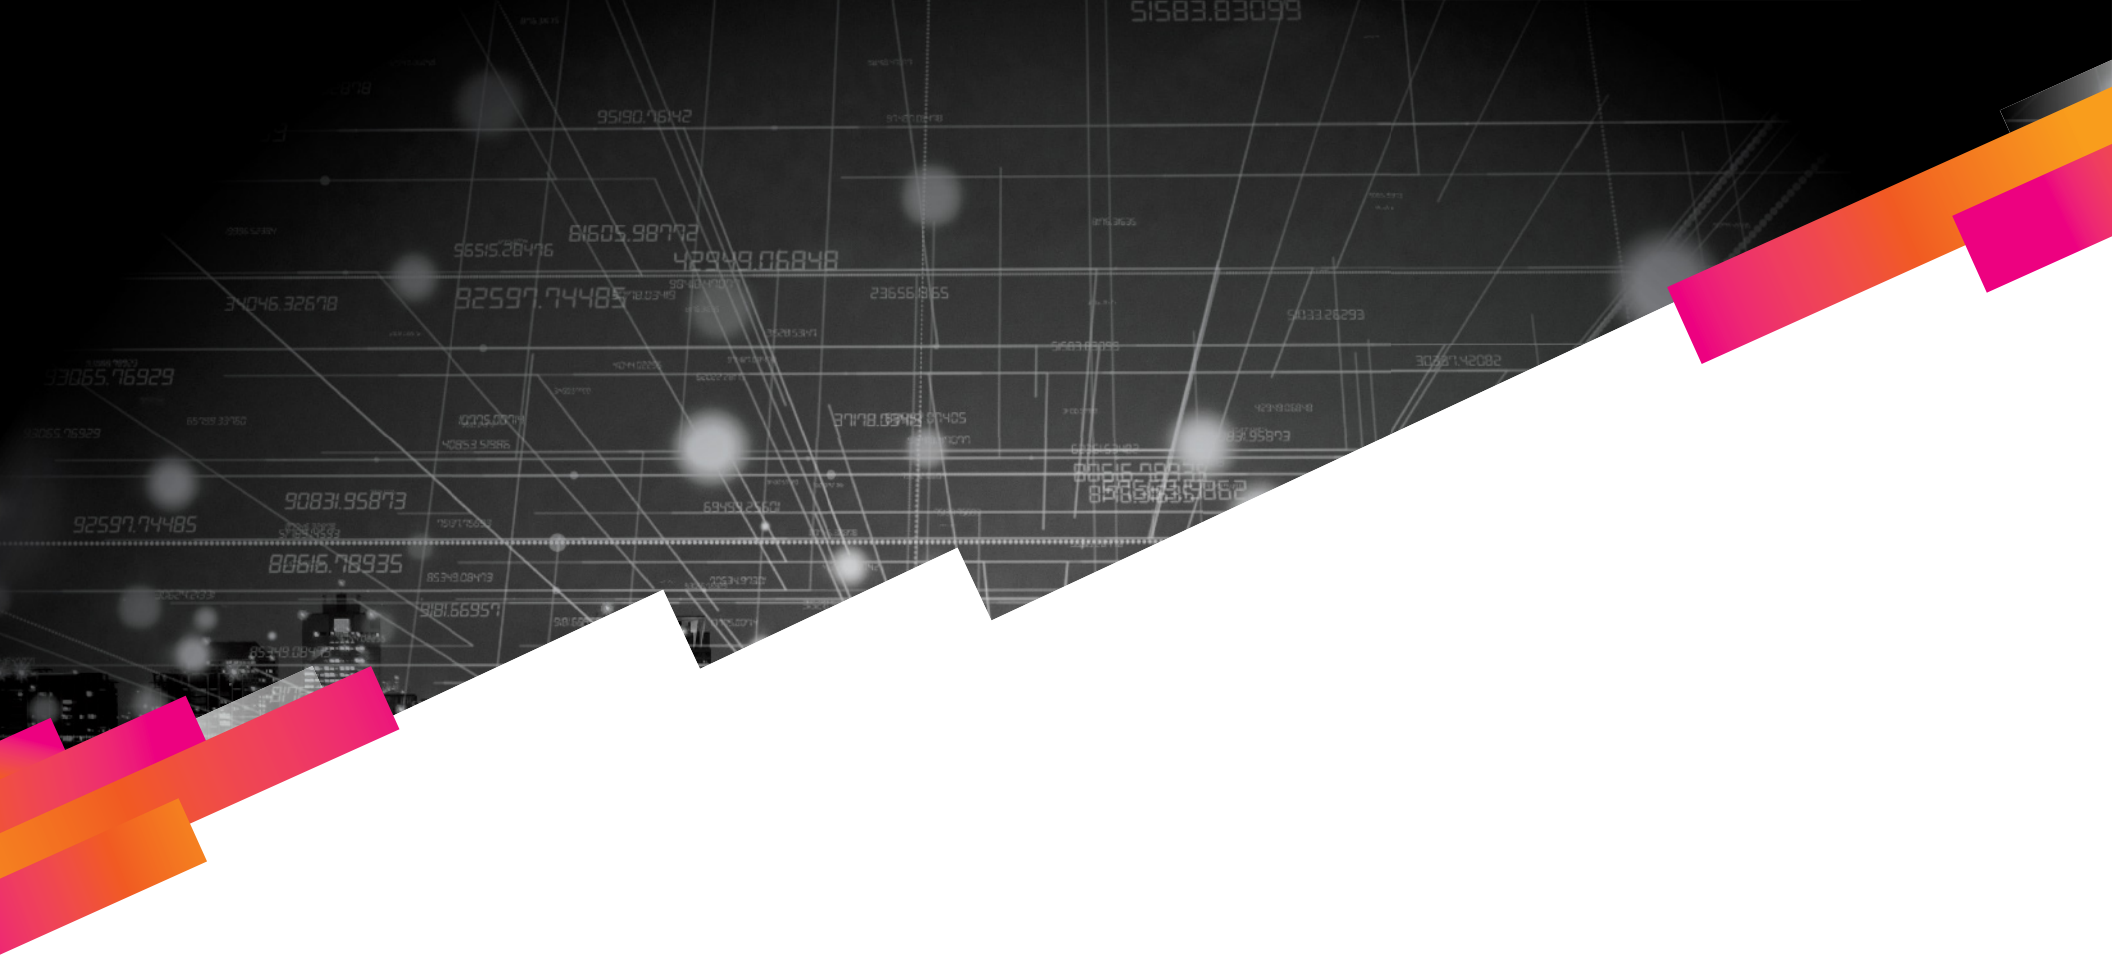

Get Started.

Visit Splunk App for Infrastructure product page.

Learn More

splunk> | aws

21-10581-SPLK-Infrastructure Monitoring 101 The Power to Predict and Prevent-114-EB