Infrastructure Monitoring 101: The Power to Predict and Prevent

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

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.

Our 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 that kind of 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, and that all (if any) issues or outages are resolved as soon as they arise. Each of these technology layers emit volumes of data that contain the information required to monitor, troubleshoot and ultimately improve those experiences.

For many years, the answer was to monitor pieces of infrastructure separately — by switching screens or administrators swiveling from monitor to monitor — but that method 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 complexity to IT infrastructure.

Collectively, the people, practices and processes that keep the infrastructure running are called IT operations. But 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.

But swivel chairs and assorted monitors are not enough to effectively monitor the multitude of separate linked systems — and thinking about IT operations holistically is only part of the answer. You need a solution to provide that centralized view. 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. Let's take a moment to look at the complex web of IT infrastructure in greater detail.

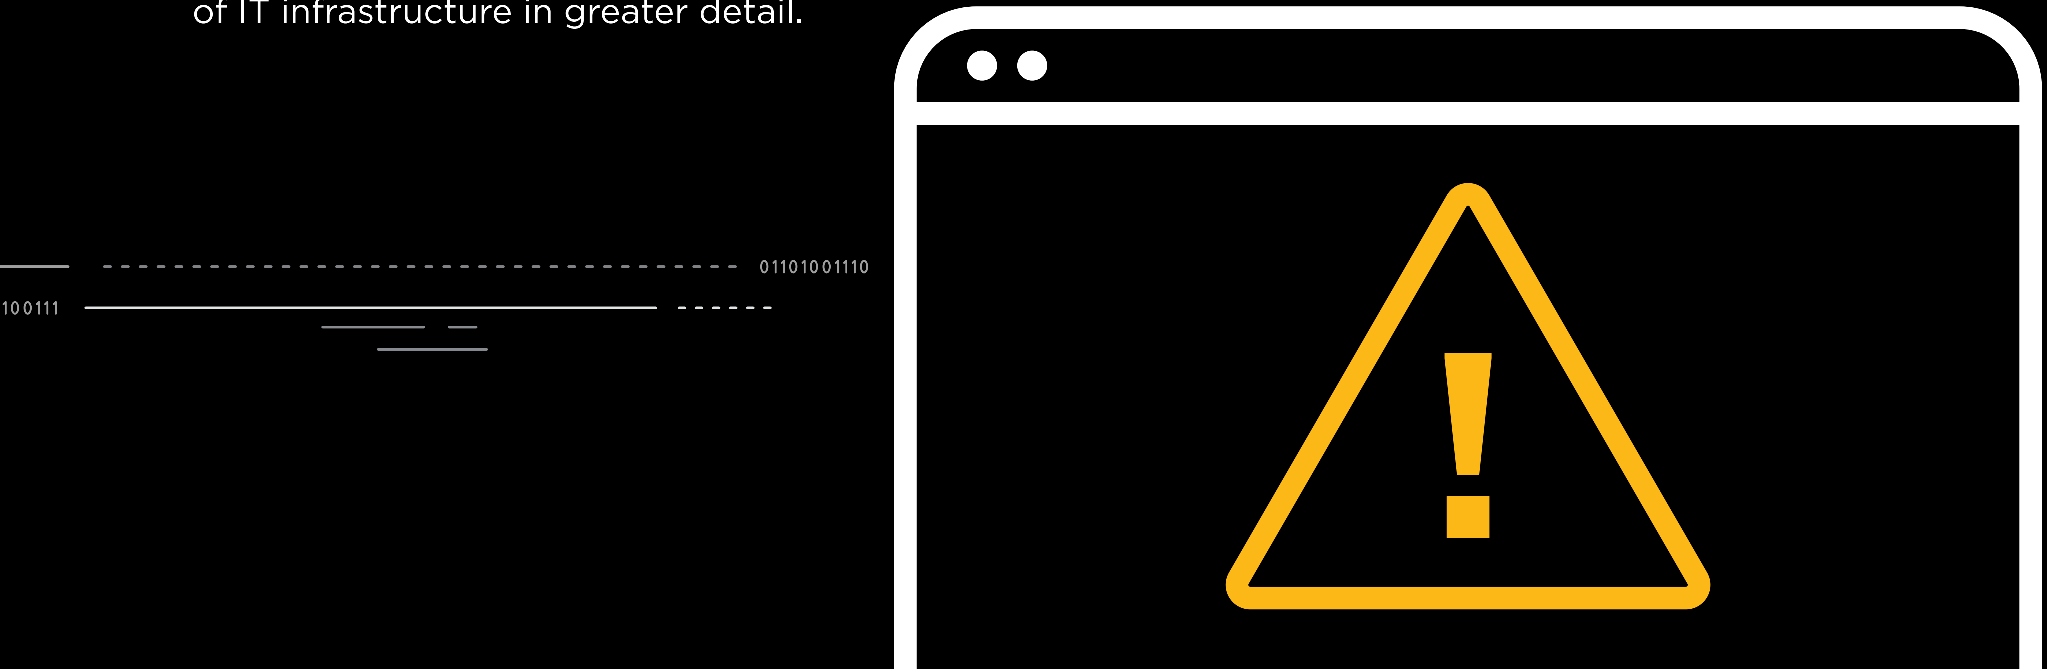

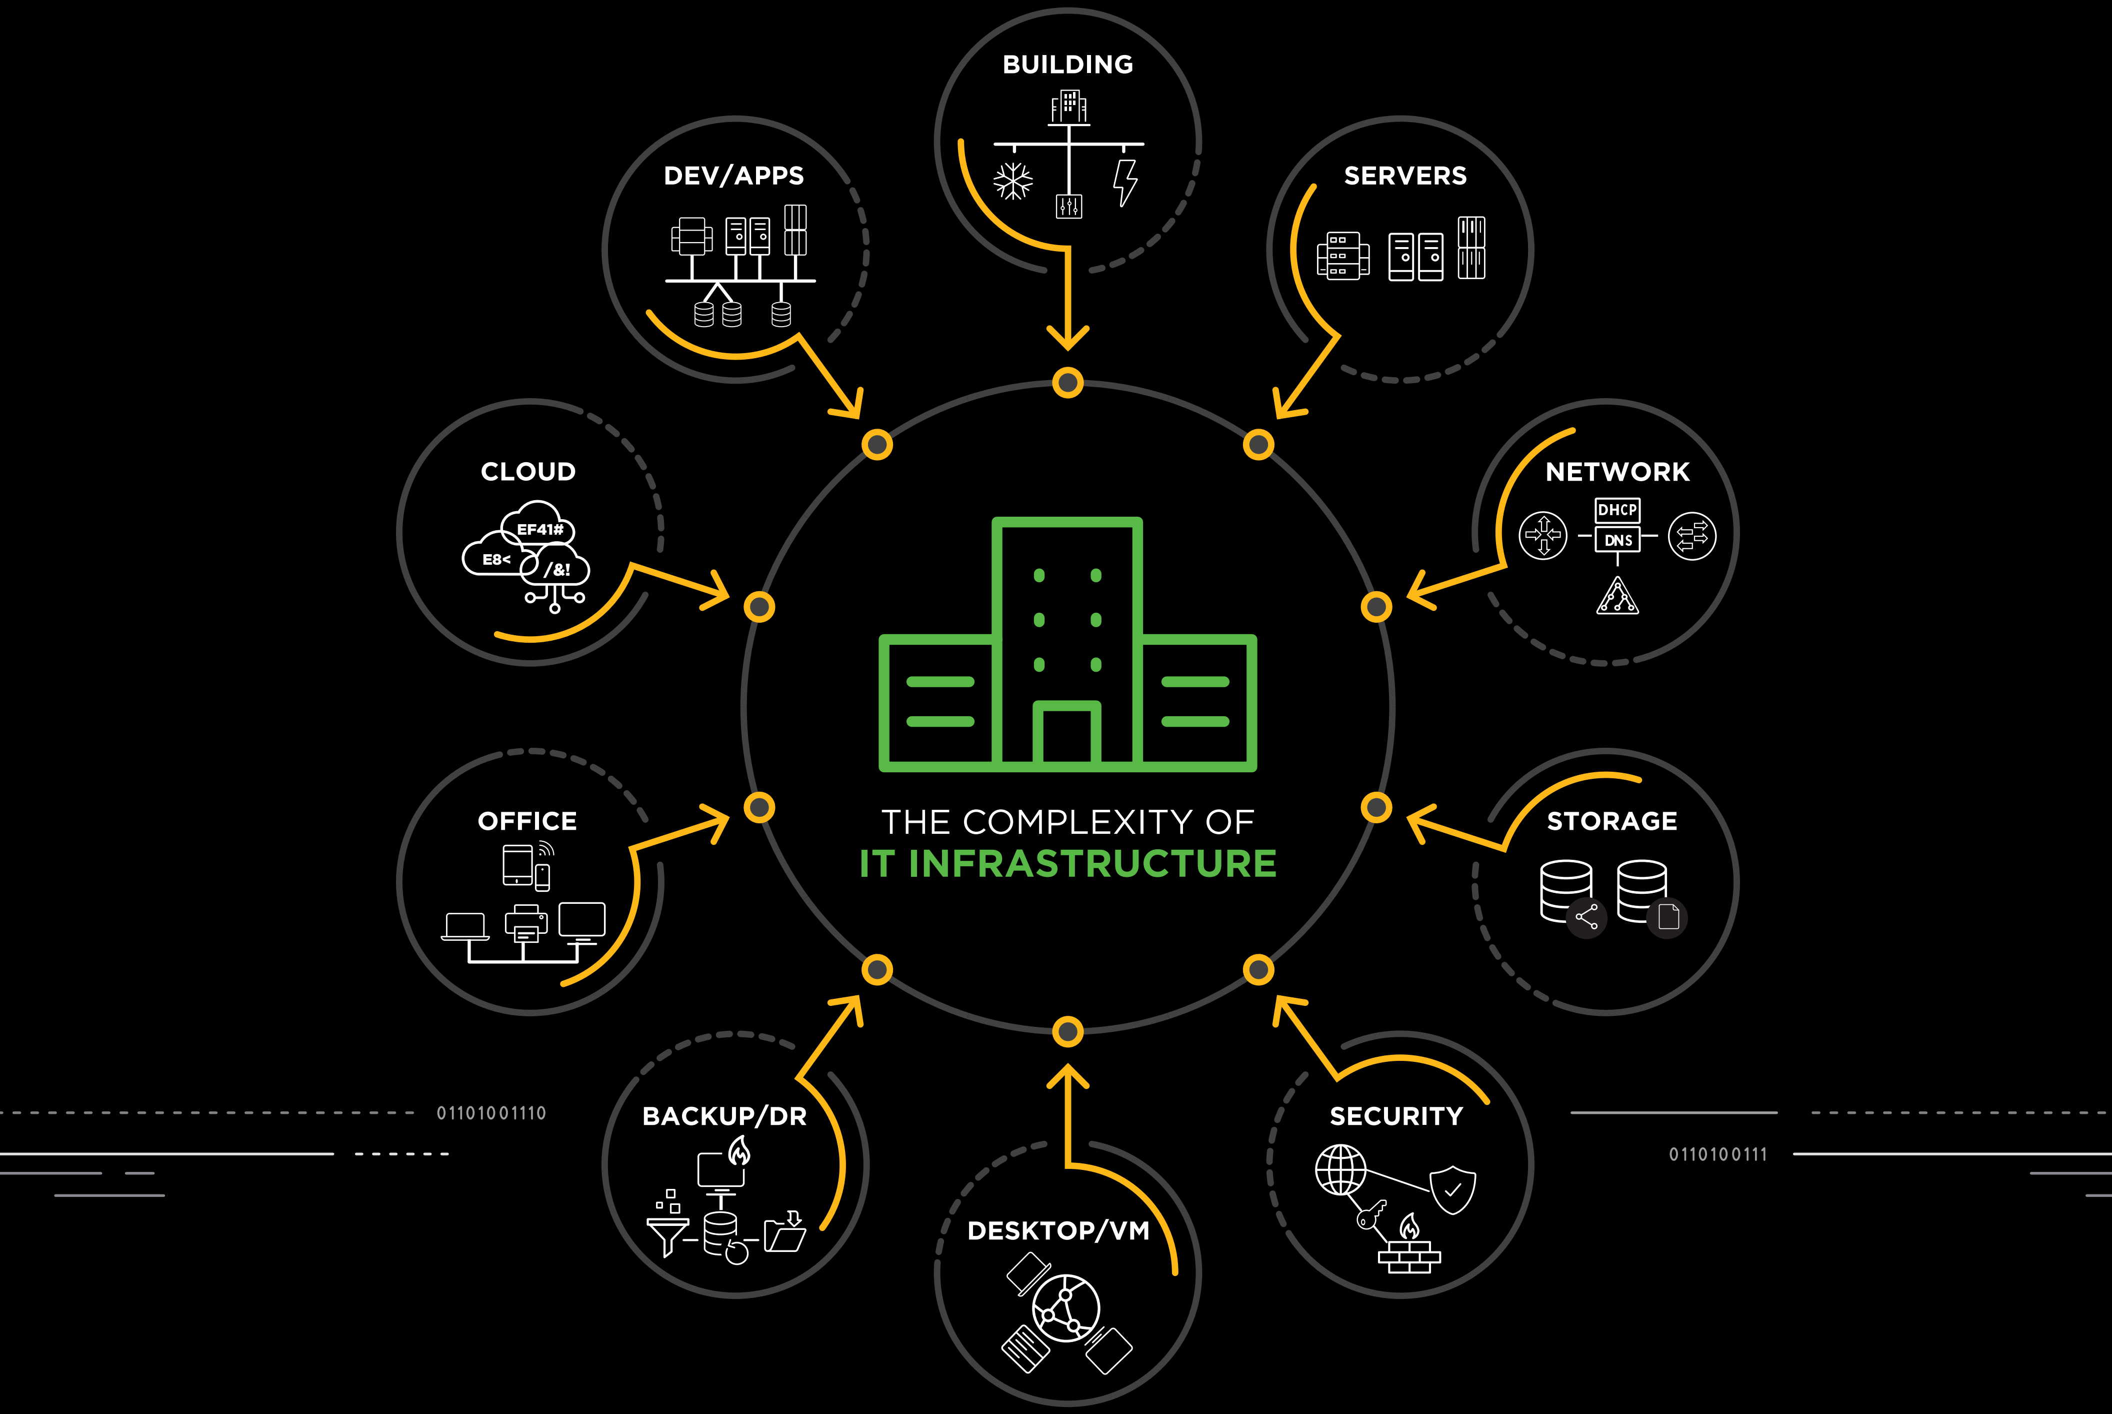

THE COMPLEXITY OF
IT INFRASTRUCTURE

BUILDING

SERVERS

NETWORK

STORAGE

SECURITY

DESKTOP/VM

BACKUP/DR

OFFICE

CLOUD

DEV/APPS

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%.*

What's more alarming about these outages is that 80% could have been prevented; they were caused principally by human error, power outages and network and configuration issues.

> "80% of these outages could have been prevented; they were caused by human error, power outages, network and configuration issues."

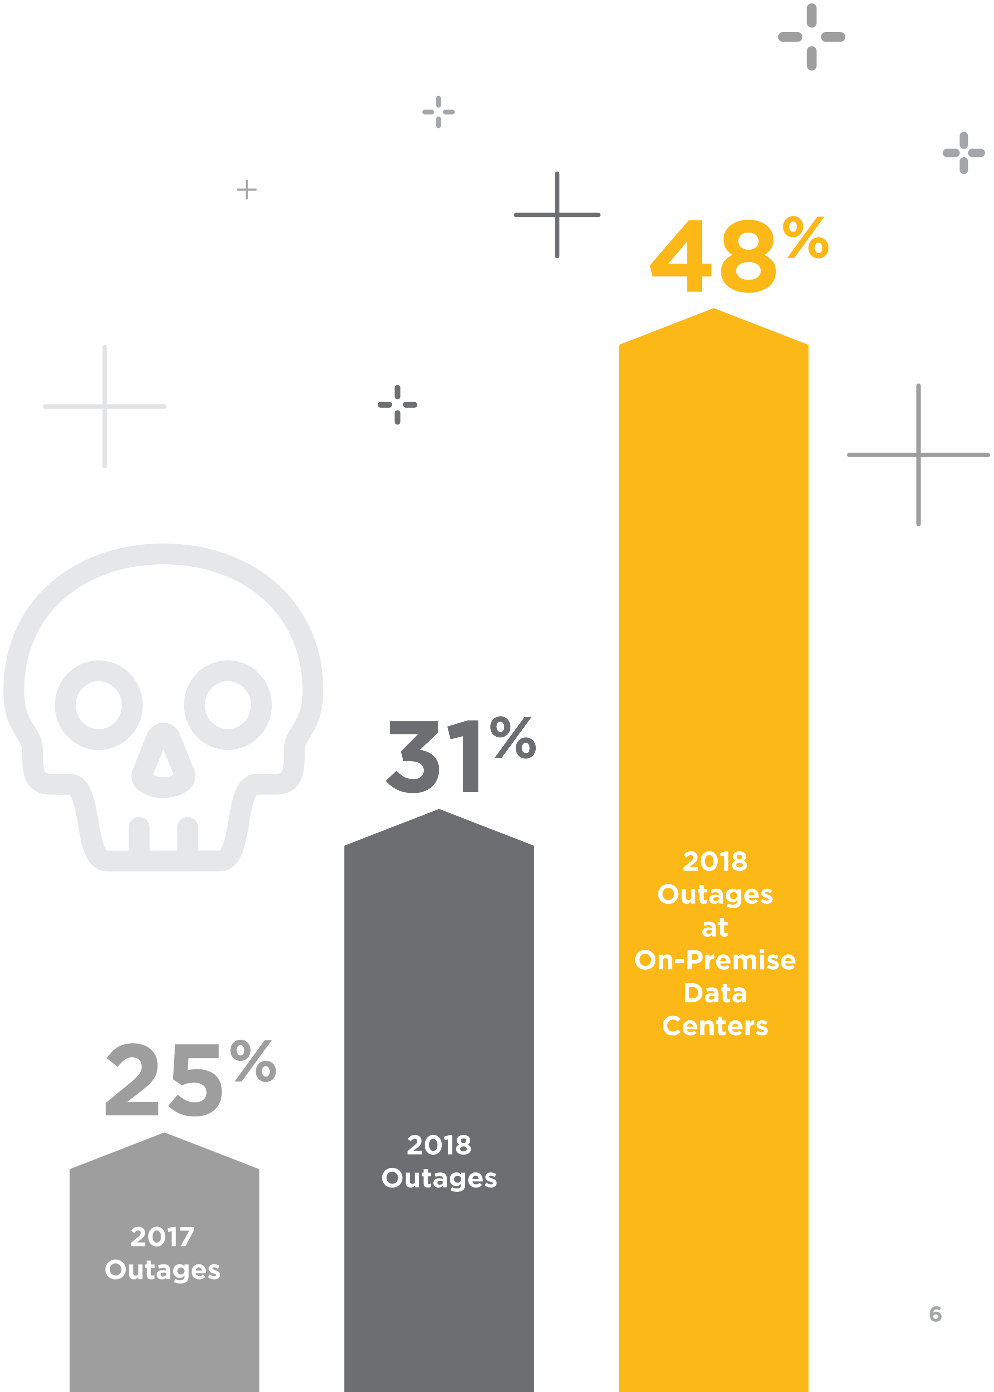

Another angle on that statistic: because of largely preventable errors, almost half of employees and users experienced issues with their apps and services in 2018. That kind of disruption can result in thousands of employee hours wasted, customer dissatisfaction and, ultimately, loss of business.

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 2018:

- On February 15, Google experienced a NoSQL database issue that resulted in high latency, impacting users' ability to log in. In turn, this affected apps and services like Snapchat and even Pokémon GO.
- On March 29, iomart experienced a failure when a farmer cut a fiber optic line, resulting in connectivity losses in Glasgow and Edinburgh, Scotland, and Manchester, England, impacting government operations in the northern United Kingdom.
- And on June 1, Visa experienced a hardware failure that knocked out services in Europe, halting essentially every Visa card transaction or ATM withdrawal.

*Source: Uptime Institute 2018 (8th Annual Data Center Survey)

The best way to ensure that issues are resolved quickly — or prevented altogether — is to monitor and troubleshoot underlying infrastructure. While observing any one element of the infrastructure stack is a straightforward proposition — there are plenty of tools available for monitoring individual pieces of the puzzle — observing each piece individually introduces a host of additional problems to the fore.

How to Cut Through the Fog: Why Infrastructure Monitoring Is Difficult

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. Gaining visibility into the infrastructure as a whole is essential — and fundamentally problematic.

A per-layer monitoring practice leads to siloed teams and incompatible views of the 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 a different language, 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 metrics. The following is a list of observable metrics and events that we have found to be critical when monitoring the infrastructure stack. These sources can be split into three groups:

METRICS

Numbers describing a particular process or activity measured over intervals of time.

- **System metrics** (CPU, memory, disk)
- **Infrastructure metrics** (AWS CloudWatch)
- **Web tracking scripts** (Google Analytics)
- **Application agents** (APM, error tracking)
- **Business metrics** (revenue, customer signups, bounce rate, cart abandonment)

EVENTS

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, Apache, MySQL, AWS)

TRACES

Data that shows which line of code is failing to gain better visibility at the individual user level for events that have occurred.

01101001110

0110100111

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 practice of constructing systems and applications to collect metrics and logs — creating them with the idea that administrators will watch over the system holistically. This is not the same as having all of your monitoring go through a single individual or team, but rather giving all roles across the stack visibility into the system as a whole. When infrastructure, operations and development teams understand the relationship of their roles to the performance of the entire system, channels of communication open up, allowing teams to solve problems more efficiently or prevent them altogether.

Let's look more closely at the layers of the IT stack:

Servers

A high-quality user experience depends on effectively monitoring the systems that support the product. It allows administrators and ITOps personnel to see resource usage patterns as well as 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, revenue.

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, for instance) 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.

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.

Learn more about the ways that network monitoring can impact ITOps in the Essential Guide to Machine Data: Network Machine Data.

Virtualization

Virtualization has revolutionized the modern data center. 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 data center. Visibility, or a lack thereof, is probably the biggest challenge.

Across virtualized machines, data center 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 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, Microsoft Hyper-V 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 complex hybrid cloud infrastructures can be opaque, leading to gaps in ITOps teams' understanding of the system as a whole. One pitfall is often overlooked by organizations eager to get the benefits of cloud; organizations often overspend on cloud services — on deprecated or unused services, unknown redundancies or excessive resource spending. Ingesting all of the cloud infrastructure data into a single environment, and 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.

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, for instance. 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.

For all the benefits that containers bring to IT organizations, they can also make cloud-based application management more complex. Some of the challenges they present include:

- **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 and complexity 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, nodes 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 experiences driving 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 infrastructural issues and reduce mean time to failure (MTTF) detection, 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.

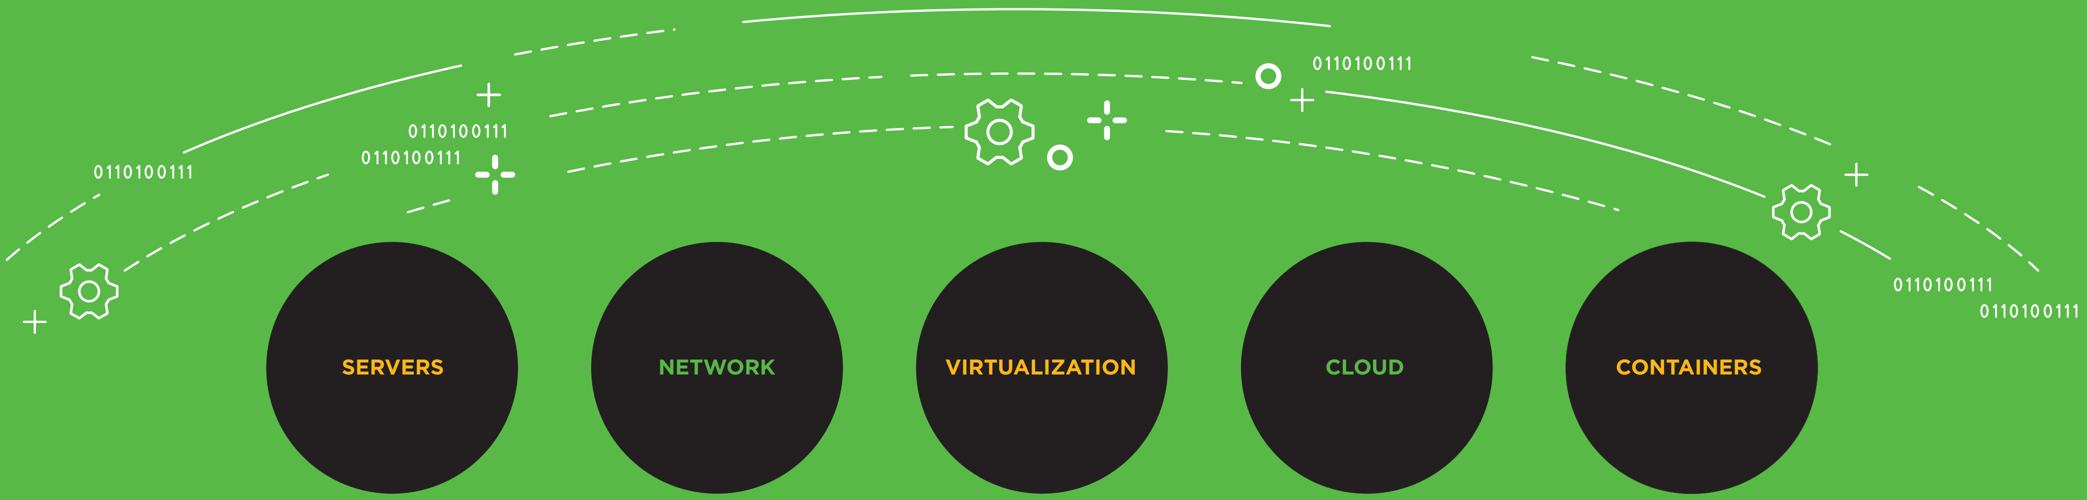

A Strategy for IT Infrastructure Monitoring

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 the data that is being collected is fundamentally unmanageable by humans. Observability allows for asking questions and having systems manage themselves using artificial intelligence (AI) and machine learning (ML) for sophisticated analytics.

Adding AI and ML to an infrastructure monitoring tool unlocks truly powerful opportunities for the ITOps team. ITOps can use artificial intelligence and machine learning 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.

01101001110

0110100111

The Full **Monitoring Stack**

IT infrastructure monitoring enables ITOps teams to get out from underneath the crush of reactive monitoring and crisis management. It also directly benefits other areas of the business by providing critical insight into the systems those groups rely on for their work.

APM

Application performance monitoring has already extended into other elements of the stack. Giving APM teams visibility into the entire infrastructure is a logical expansion of that access. Application speed and uptime — for internal enterprise apps and consumer apps alike — is directly tied to an organization's profitability. Knowing where in the infrastructure 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 have a significant impact on the ability of network administrators to do their jobs. With a more complete view of the infrastructure that supports the network (and that the network, in turn, supports) network performance monitoring and diagnostics (NPMD) can improve their mean times to detection, 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.

Customers Who've Succeeded With Infrastructure Monitoring

Entrust Datacard

When consumers, citizens and employees make purchases, cross borders, access e-government services or log on to secure networks, they expect their transactions to be secure and seamless. Entrust Datacard, with its varied portfolio of innovative security solutions, provides the foundation to enable frictionless, secure transactions, tied to trusted identities. Recently, the company needed unified infrastructure monitoring and metrics to drive operational success during the development of an innovative new cloud service. Since deploying Splunk Enterprise running on Amazon Web Services (AWS), Splunk App for Infrastructure and VictorOps, Entrust Datacard has seen benefits including:

- Modernizing operations, introducing automation and delivering software faster.
- Proactive and collaborative monitoring to ensure a positive customer experience.
- Reducing the number of monitoring tools required while increasing coverage.

Learn more about Entrust Datacard's infrastructure monitoring success.

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 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 freed to 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.

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.

Splunk App for Infrastructure

Splunk App for Infrastructure (SAI) is an IT infrastructure monitoring solution that unifies and correlates metrics and logs for a seamless monitoring and troubleshooting experience, simplifying the monitoring and observability needs of sysadmins and site reliability engineers. Intelligent investigations make spotting trends and finding the root cause of server, OS and Amazon Web Services (AWS) problems easier.

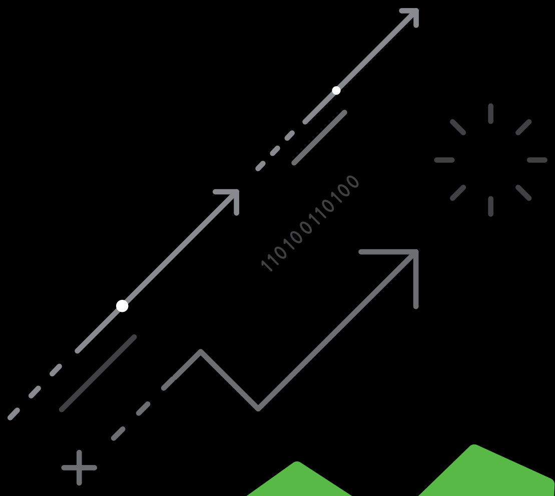

What does Splunk App for Infrastructure monitor?

METRICS

Metrics from your on-premises and cloud-based host that help identify performance trends and issues across the infrastructure.

EVENTS

Logs and events from your hosts with a detailed record of errors, changes and events that can help teams isolate the root cause.

Monitoring metrics that matter

The Splunk App for Infrastructure provides curated, unified metrics and logs focused on infrastructure performance monitoring. Easily analyze metrics by defining and grouping entities. Collect, index, search and build visualizations based on metrics.

Advanced alerting for faster triage

Perform root cause analysis faster with SAI's custom triggered alerting at a group or entity level. Triage alerts more effectively by understanding which conditions triggered the alert, assess the severity of the alert and view all triggered alerts to assess what actions to take.

Visualizations for real-time monitoring

SAI allows for monitoring in real-time with prebuilt visualizations. Monitor performance of hybrid infrastructures by entity, including CPU, network, memory, disk, system load, custom-defined dimensions and more. Monitor single entities or groups of entities. Drill down into an entity or group to review details or troubleshoot an issue.

Correlations pinpoint performance trends

Investigate performance with correlations across metrics and logs from your infrastructure. Analyze performance metrics for a single entity or a group of entities. Determine poorly performing entities by metrics, or determine a point in time when multiple entities began performing in a similar way. View and search for entities in a group, or view all groups an entity belongs to for easy navigation in a chart or list.

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 SAI directly into Splunk IT Service Intelligence (ITSI) to search and analyze across multiple layers of the IT stack, or drill down into the raw infrastructure logs or metrics for advanced troubleshooting. Integrate entities and groups from SAI directly into ITSI as services with just a few clicks using a single integration interface in ITSI.

The Splunk App for Infrastructure is an IT infrastructure monitoring solution that unifies and correlates metrics and logs for a seamless monitoring and troubleshooting experience, simplifying the monitoring and observability needs of sysadmins and site reliability engineers. Intelligent investigations make spotting trends and finding the root cause of server, OS and Amazon Web Services problems faster.

What does Splunk App for Infrastructure monitor?

- Metrics from your on-premises and cloud-based hosts help identify performance trends and issues in your infrastructure.
- Logs and events from your hosts provide you with a detailed record of errors, changes and events that can help you isolate the root cause.

Get up and running in minutes

Get started today with a free trial of Spunk Enterprise and download the Splunk App for Infrastructure from Splunkbase. The Splunk App for Infrastructure is available as a free app with a Splunk Enterprise license. Streamline troubleshooting and monitoring workflows by sending infrastructure data from SAI into Splunk IT Service Intelligence. Drill into SAI data directly from Splunk ITSI for deeper insights and better understand incident responses with SAI alert details in Splunk VictorOps. SAI integrations with Splunk ITSI and Splunk VictorOps encourage sharing insights across teams and allow for better collaboration and increased productivity.

→ Want to learn more?
Visit the Splunk App for Infrastructure product page.

ABOUT SPLUNK.

Splunk Inc. makes data accessible, usable and valuable to everyone.

→ Want to learn more?
Visit the Splunk App for Infrastructure product page.

01101001110

0110100111