



# Advance your AWS cloud environment with automatic & intelligent observability

---

5 game-changing ways to enhance hybrid cloud performance



# Outline

- 1 Executive summary
- 2 Modern cloud environments demand more
- 3 5 game-changing ways to enhance AWS hybrid cloud performance
- 4 Let's take action

# Executive Summary


Today's organizations face huge pressure to keep up with the always-on, always-connected digital economy. Consumer demand for instant access to digital services and expectations for constant improvement and greater convenience force organizations into ongoing cycles of innovation.

AWS offers an industry-leading cloud platform that allows companies to move faster and better meet the needs of their customers and their business. However, the transformation to cloud isn't simple—it's a fundamental shift in how applications are built, deployed, and operated. Additional technologies like microservices, containers, and serverless infrastructure increase technology touchpoints across the cloud environment into millions of lines of code and billions of dependencies, making management of these dynamic environments extremely complex.

This is where automatic and intelligent observability leap past traditional three pillar observability. It is quickly becoming a recognized lifeline for organizations looking to not only accelerate their digital transformation, but also help automate operations for more agile and confident business performance.

This ebook will help dev, ops, SRE, and business teams transform faster and massively increase efficiencies by looking at five recognized ways that automatic and intelligent observability can help you:

- Move more confidently and quickly
- Operate and scale more simply
- Drive innovation and optimize business results



<sup>1</sup>In software, observability refers to the extent that the internal status and performance of a system can be inferred from its externally available data.

## Modern cloud environments demand more

The value of observability is increasing as companies strive to understand and accelerate their digital transformation using massive cloud and cloud technologies—environments that are inherently hard to observe and operate due to their dynamic nature and complexity.

Most traditional monitoring or observability tools focus on collecting and aggregating three principal data types—metrics, traces, and logs, the so-called three pillars of observability. But in modern environments, these legacy solutions have limitations for scalability and the ability to “know the unknown”. Today’s environments need more. The business operation as a whole needs to be more observable—and not just on dashboards.


In today’s AWS hybrid-cloud environments, organizations need clear, fast answers to address problems, questions, and decisions for better business operations.

That’s where Dynatrace comes in. Dynatrace has innovated and advanced traditional observability to include automation and intelligence. This distinction is enabled by three unique technology capabilities:



### **Automation for scalability and completeness**

achieved through continuous and automatic discovery and instrumentation with zero manual configuration



### **Real-time topology mapping for context across the full stack,**

self-discovering what’s new in hybrid environments, continuously building, updating, and maintaining a complete entity map of everything in real-time across millions to billions of dependencies



### **Causation-based, code-level AI for precise answers**

all powered by an AI engine, Davis®, that provides contextual, actionable answers to problems and issues through real-time, root-cause analysis

**Automation and intelligence are essential to high caliber, advanced observability.**

Requirements		Results
Precise, actionable answers	→	Speed and reduced risk
Automatic complete coverage	→	Innovative productivity and business performance
Automation everywhere	→	Higher quality releases and operations
Real-time feedback	→	Better business decisions and customer experiences
Tighter cross-collaboration	→	Accelerated business outcomes

## Automatic and intelligent observability for scalability and completeness

Most observability approaches require developers to manually instrument their code. In environments with thousands of hosts and microservices that dynamically scale across global, hybrid cloud infrastructure, this is an exponential, futile effort. With the infusion of automatic and intelligent observability, Dynatrace provides:



### **Auto-discovery**

Upon install, Dynatrace OneAgent automatically detects all applications, containers, services, processes, and infrastructure in real time with no manual configuration



### **Auto-instrumentation**

System components are automatically instrumented with zero configuration or code change and not only collect metrics, logs and traces, but also user experience



### **Auto-baselining**

Dynatrace's smart baselining automatically learns "normal" performance behaviors and adapts dynamically in real time as the environment changes

## **Dynatrace awarded AWS Machine Learning Competency**

Dynatrace has received AWS Machine Learning Competency status in the new Applied Artificial Intelligence (Applied AI) category. This designation reflects AWS's recognition that Dynatrace has demonstrated deep experience and proven customer success building AI-powered solutions on AWS to help some of the world's largest organizations accelerate digital transformation.

## Real-time topology mapping for context across the full stack

Metrics, traces, and log data are frequently stored without meaningful context that tie them together. This makes assessing true holistic system health and understanding the impact of problems inaccessible and unattainable. With the infusion of automation and intelligent observability, Dynatrace adds end-to-end information, including user experience data, and provides:



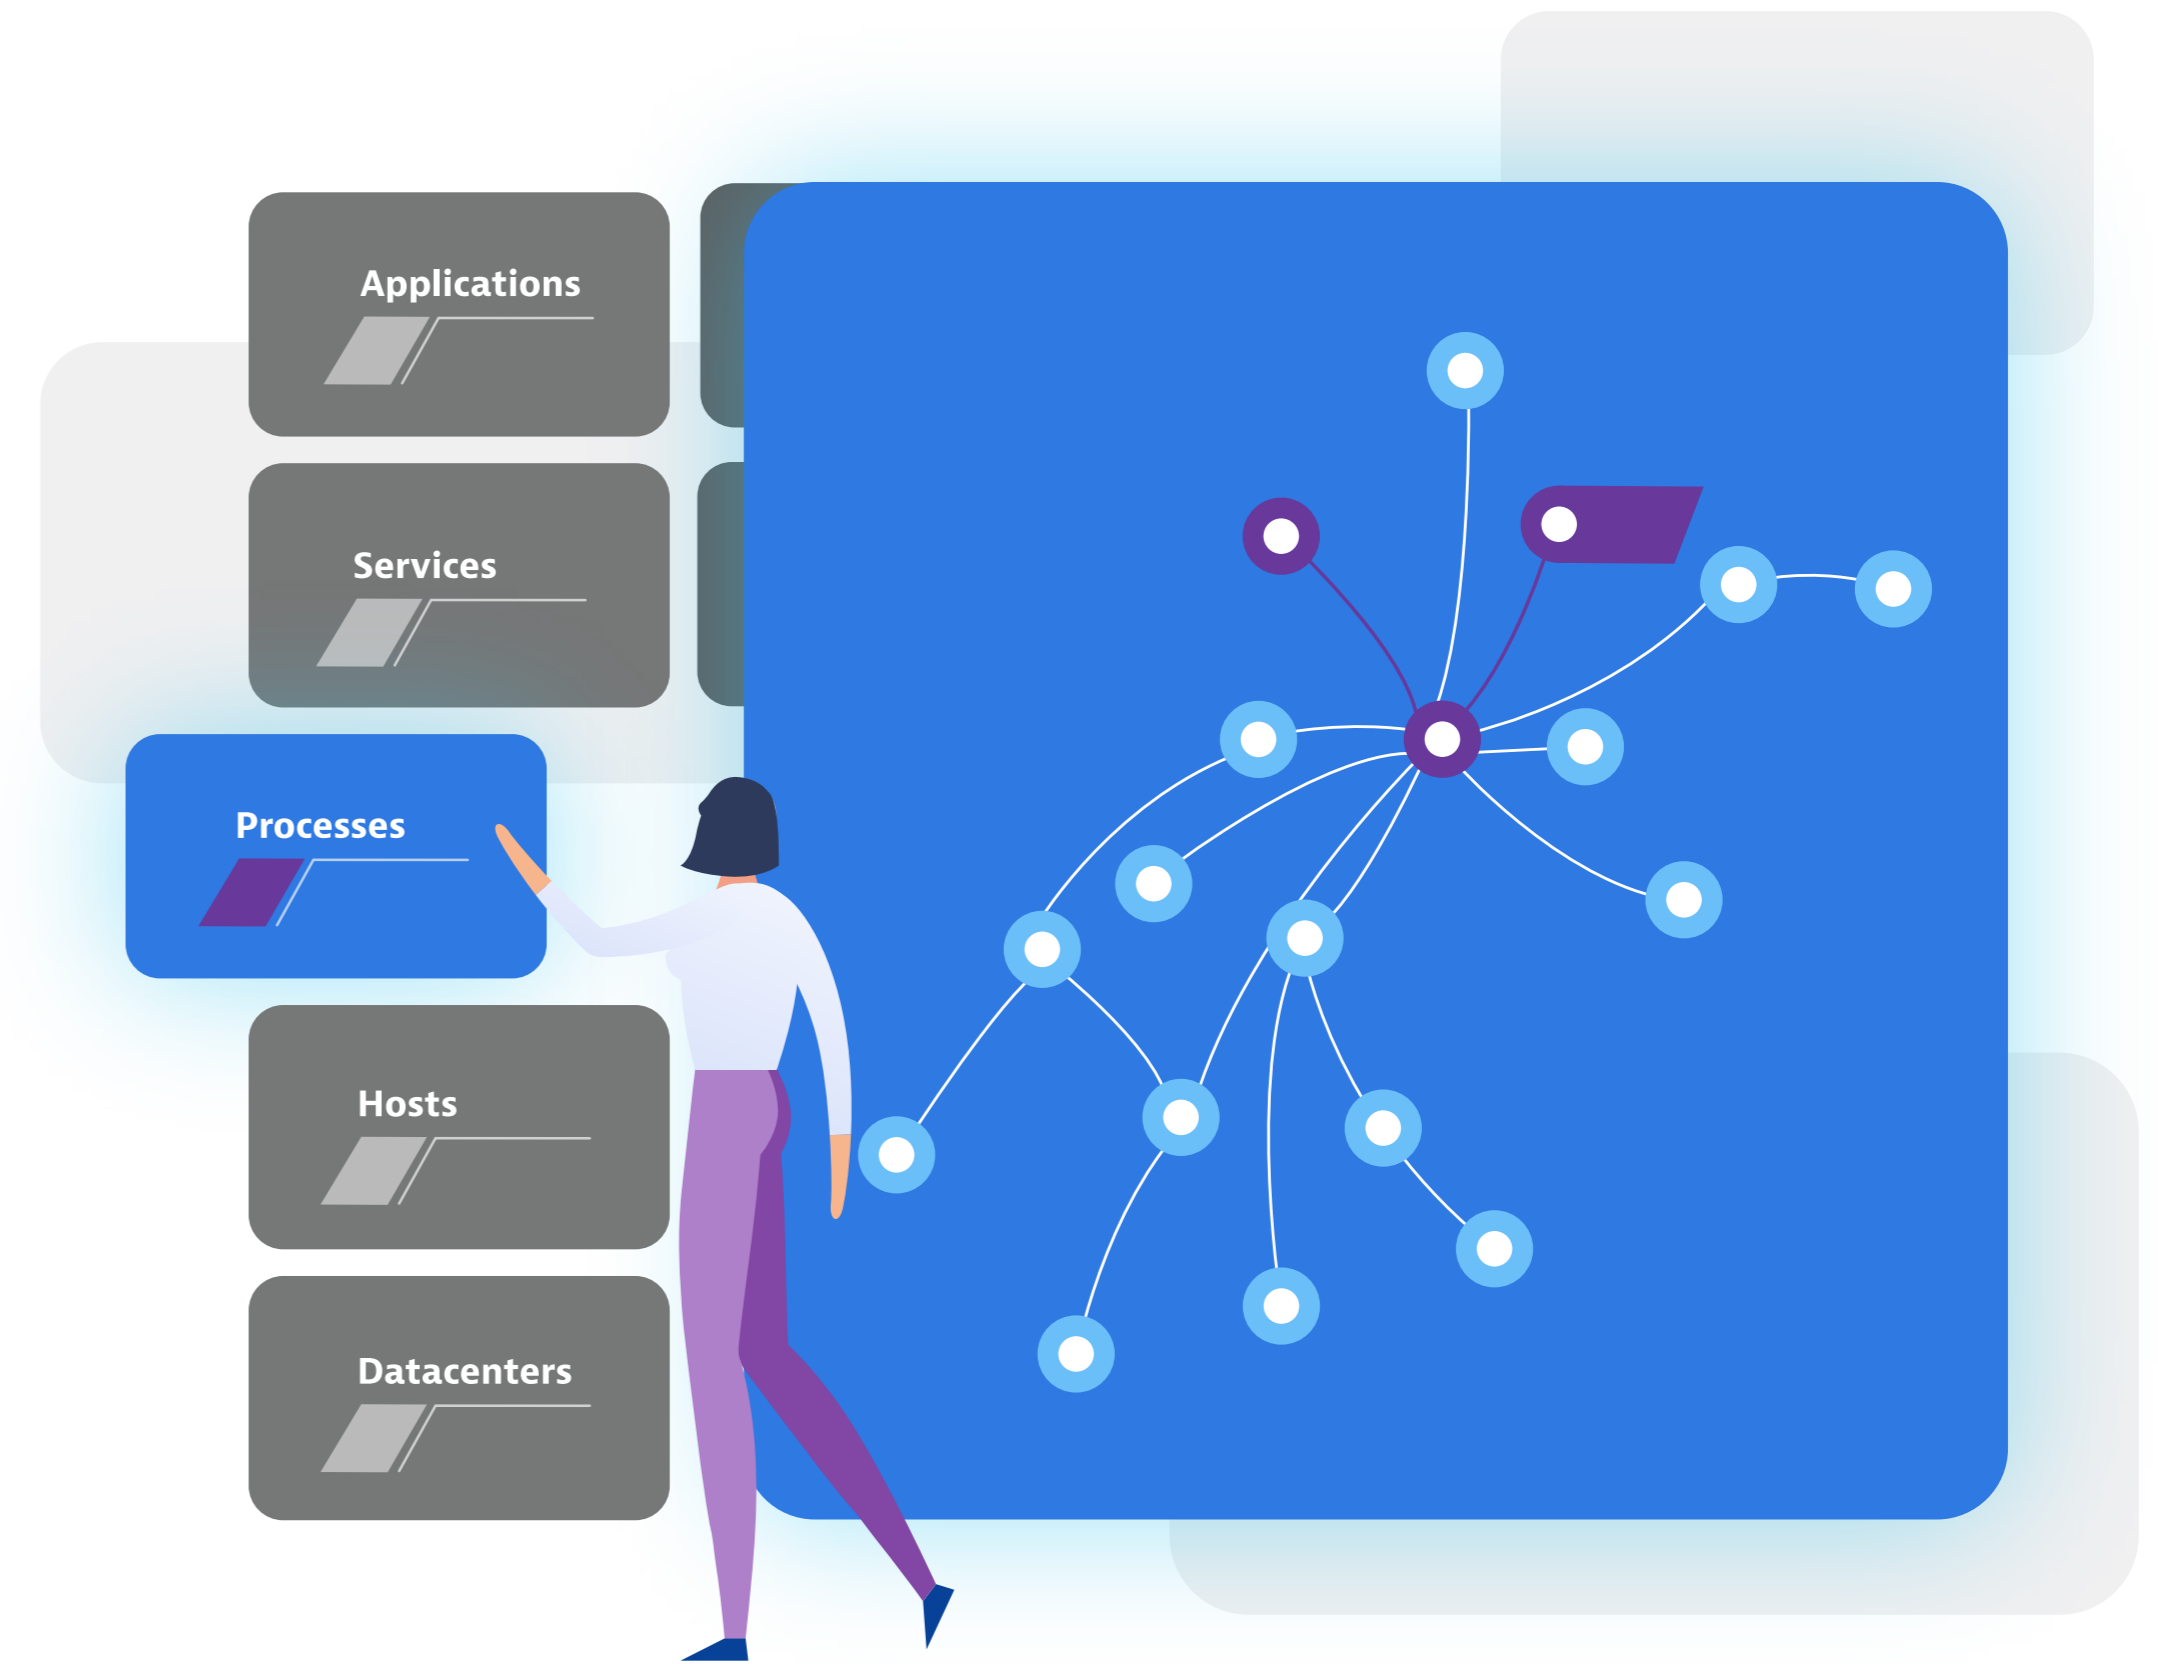
### Real-time topology mapping

A continuous mapping of components, cloud services, and ever-changing relationships between billions of interdependencies, removing silos



### Automatic and always-on

AI reaching across massive cloud environments, constantly tracking and mapping in real-time, compared to less accurate, simple, time-based correlations



## Causation-based, code-level AI for precise answers

Traditional observability solutions offer little information beyond dashboard visualizations. Those dashboards require manual analysis of data and good educated guesses, sometimes in painful war rooms.

Dynatrace's AI engine, Davis, is a causal AI engine that sits at the core of the Dynatrace platform. It automatically processes high-fidelity logs, metrics, traces, and real user data from a range of environments to serve precise answers that are prioritized by business impact.

- Davis is built from the ground up for the cloud. As the core of the Dynatrace platform, it sees all data and dependencies across the full stack, including third-party data, independent of origin.
- Unlike traditional ML approaches, Davis doesn't need time to learn, and it doesn't make statistical guesses. By leveraging the Smartscape®, Davis already knows the true business impact, and alerts only when problems occur.

## Looking ahead: OpenTelemetry

OpenTelemetry is an open-source project spearheaded by the Cloud Native Computing Foundation (CNCF) with the aim of making software more observable and establishing telemetry as a built-in feature of cloud-native software.

This initiative is broadly supported by the open-source community, as well as industry-leading contributors like Dynatrace. We are working in several areas including auto-instrumentation, interoperability, and enterprise-grade solutions.

Dynatrace contributes knowledge from building our own industry-leading technology where logs, metrics, and traces are collected using AI and automation.

### What makes Dynatrace unique in the OpenTelemetry space?

Dynatrace adds the most observability value in the industry with actionable answers on top of data — not just more data on glass.

- Automation ingests real-time data continually into an AI engine
- All information is automatically produced in actionable context, with code-level details, end-user experience, and entity relationships.

As OpenTelemetry is more widely adopted as such, it will serve as an additional data source that will further extend the already advanced breadth of what Dynatrace automatically and intelligently collects, ingests, and observes into its AI engine, Davis.

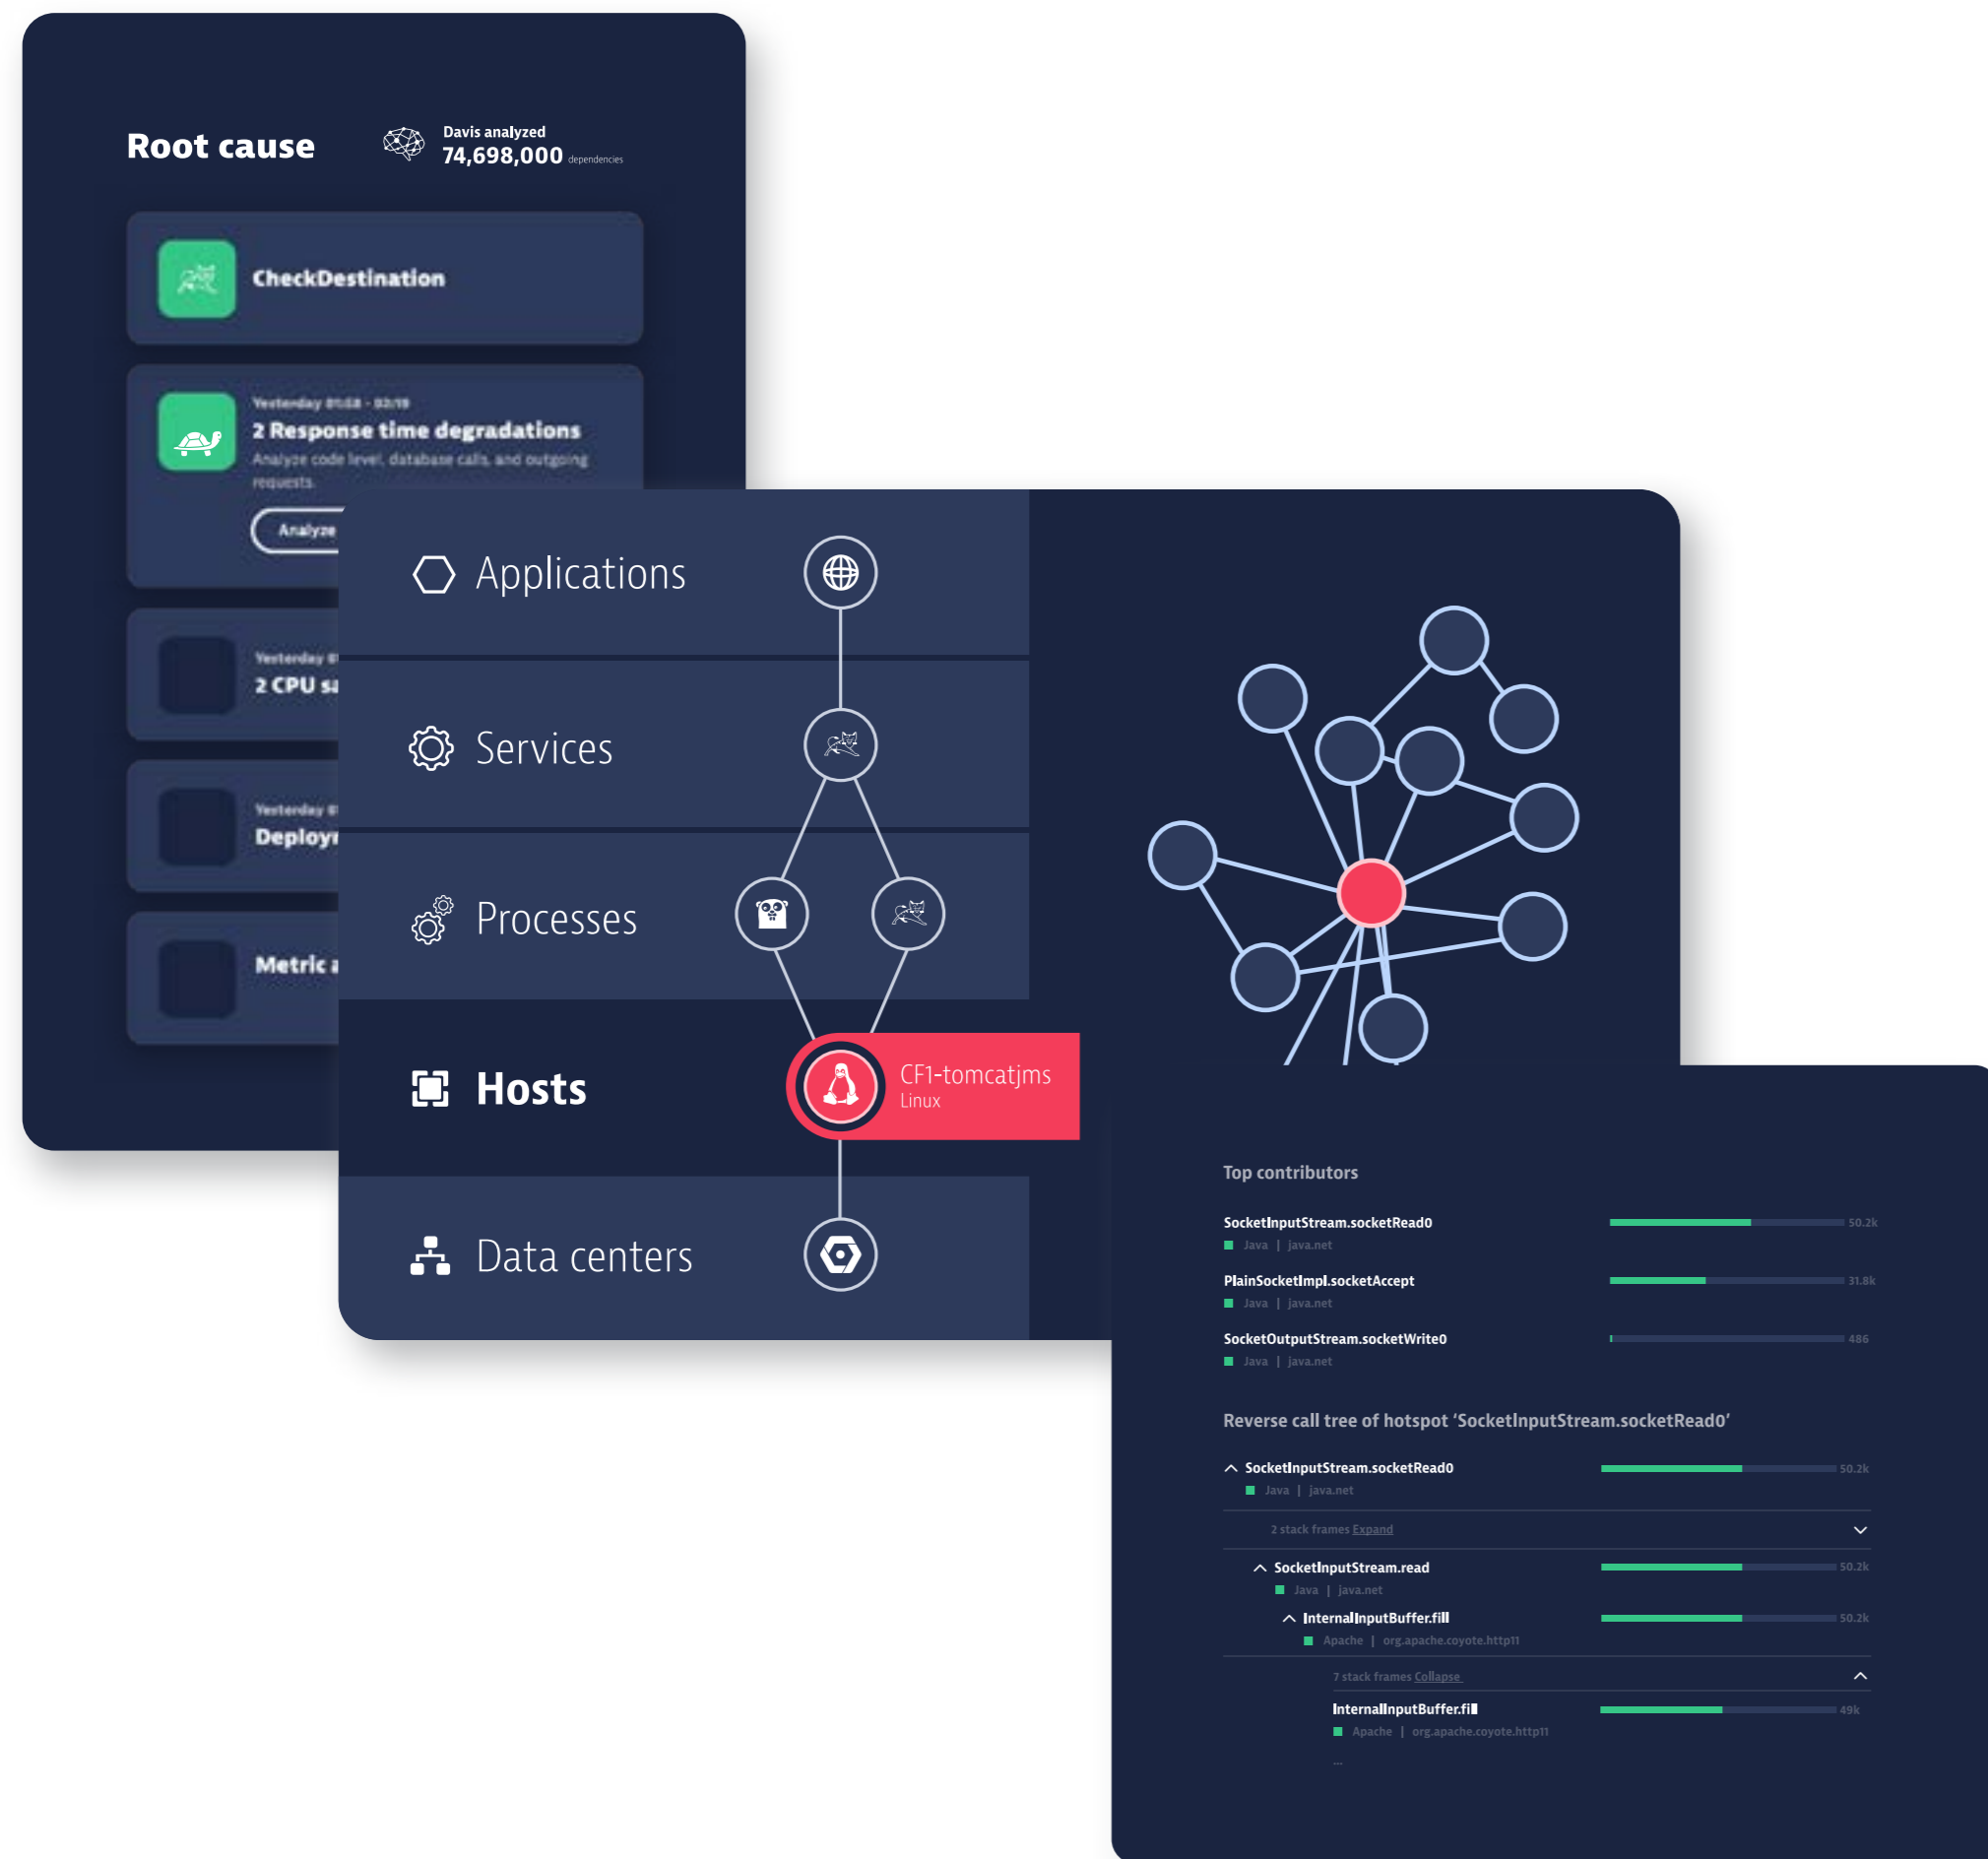




# 5 game-changing ways to enhance AWS hybrid cloud performance

This section shines a light on five ways that companies can advance the performance of their AWS hybrid cloud environment using automatic and intelligent observability to raise the performance for operations, DevOps, and digital experience initiatives, and collaborate more tightly on value-added tasks that improve business outcomes.





# 01 Use automatic and intelligent observability to accelerate modernization

A lack of a clear vision can be the most challenging problem for a cloud modernization, and it can be extremely costly for your business if you don't prepare properly for it.

As you modernize and move apps to the AWS cloud, there are three generally acknowledged areas to focus on around the decision and execution process:




## Plan better

- Evaluate that the app is well suited for cloud with essential early insights
- Assess that the new design and cloud architecture will work well and be effective



## Execute faster

- Avoid problems that impact service delivery and cause delays by identifying issues quickly, in full context
- Resolve problems faster, reducing overall project risk



## Optimize operations

- Feel confident that expectation is met for service delivery through clarity on performance before and after migration to the cloud



## Plan better

Get a clear picture and be an expert on your environment with good data. This is essential for accurate assessments, better decisions, and optimized design. Dynatrace automatically detects all your hosts, processes, services and technologies, where they run, and which services they support. With better planning, you'll understand and be able to answer questions like:

- Which technologies do we have in our organization? Where do they run?
- Which technologies are candidates to be moved and which ones aren't?
- For which technologies do we have alternative offerings?

What if you don't see unmonitored hosts that are a part of the environment you want to move? Not a problem—Dynatrace reveals everything it detects, even if it's not being monitored. This is critical, as you don't want to migrate a service and introduce costs or issues because you were unaware of a dependency.

Now you can move forward and make confident architectural decisions and feel more comfortable in how a new design and cloud architecture might work.



## Execute faster

Stay on top of things and avoid problems that impact service delivery by analyzing details of traffic dependencies and adding predictability. This will help you make "if/then" decisions based on future costs as well as other data, traffic, and resource decisions while ensuring you are meeting your success criteria. You can evaluate and consider questions like:

- Exactly which services do we have?
- What are the current usage and resource consumption patterns of services
- What will it cost us when running in the cloud?
- What services can be migrated in isolation and which ones have a tighter dependency and where?
- What happens if I move this group of servers?

Now you'll be able to understand, answer questions and make decisions based on a clear view of services, dependencies, and resource consumption.



## Optimize operations

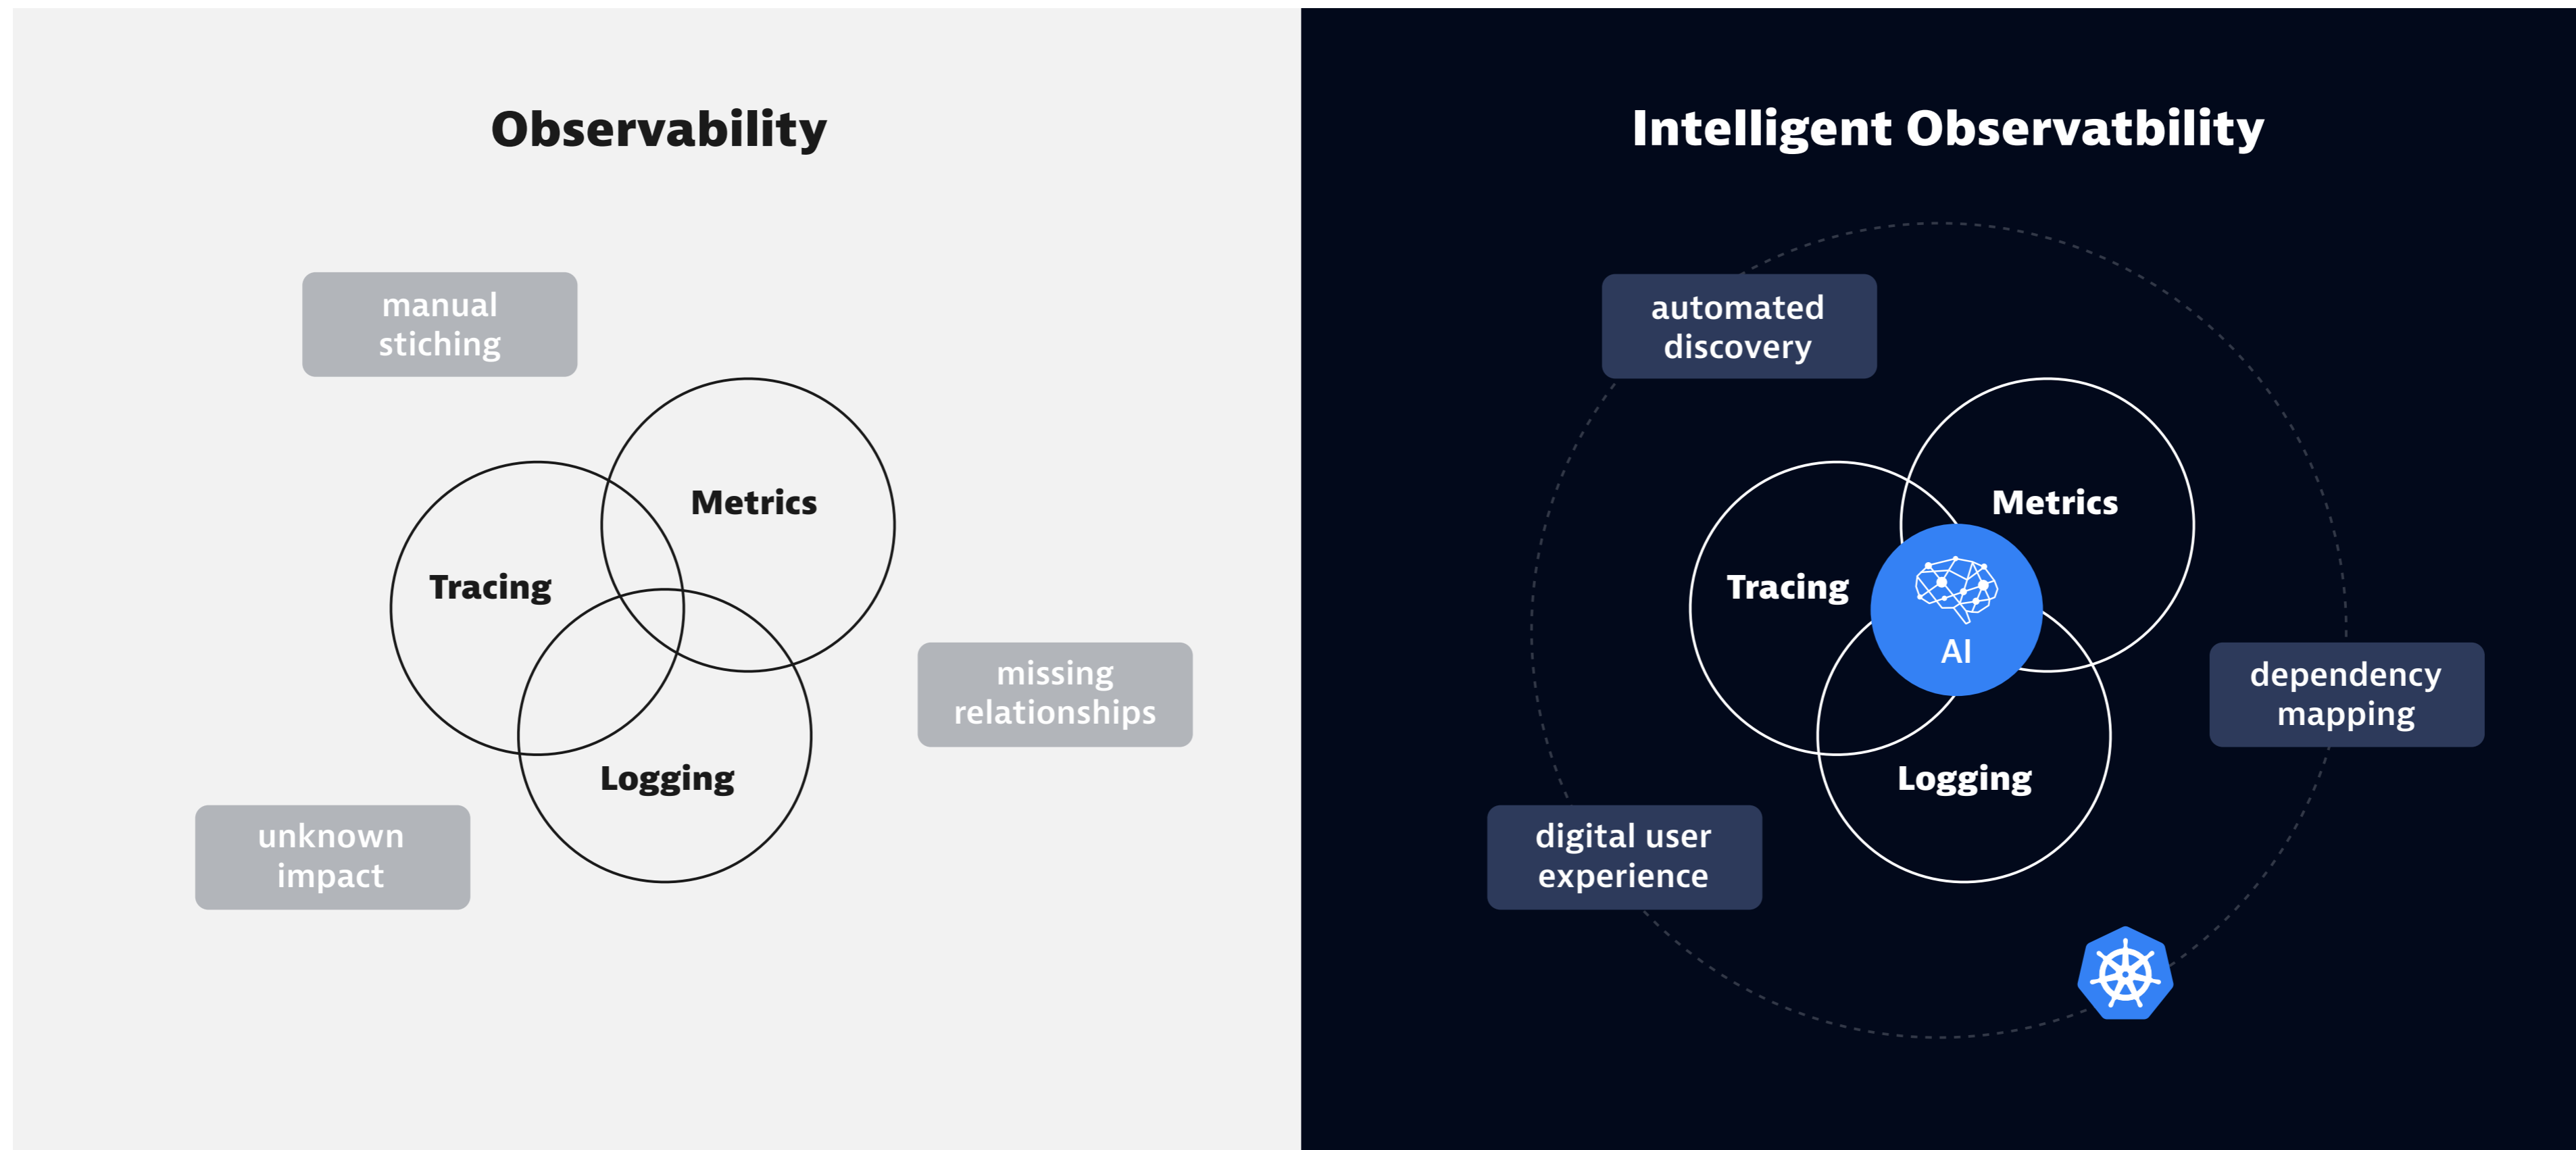
Next let's look at identifying areas for further optimization and places to build in more confident platform and business performance. Get data on components like back-end services, front-end API services, and user facing features, then examine this data through the lens of business benefits. You can start to look for things like:

- How does our performance of the environment compare to pre-modernization?
- What performance touchpoints stand between your top revenue generating features and your data center?
- Would moving a feature closer to the end-user improve the experience?
- Are there features using dedicated hardware better suited for on-demand cloud infrastructure?

Now you can focus automated and intelligent observability on things like customer experience, better features and functions, and more innovative ways to improve business impact. This stage also helps you identify opportunities to build in automation and further modernize operations.

**Dynatrace delivers extremely high-fidelity data for decisions across these three important modernization processes.**

By using automatic and intelligent observability for discovery and instrumentation, topology dependency mapping, full stack context, and actionable answers, you can accelerate modernization and ensure every app is available, functional, fast, and fully optimized across all channels.



## Three essential tips from Dynatrace to extend your observability for your modernization project

### ★ Essential tip 1:

Start with a complete and total view of your entire environment. Deploy the [Dynatrace OneAgent](#) for continuous and automatic discovery and instrumentation for completeness and scalability, all with zero configuration. Enrich data and context using [Dynatrace APIs](#) to third-party data—*independent of origin*—a widely used best practice to infuse customer reports with performance and stability insights.

### ★ Essential tip 2:

Immerse yourself in the Dynatrace Smartscape, a topology mapping model that streamlines the layout of interdependencies across your full stack—all applications, containers, services, processes, and infrastructure in real-time. Again, with no manual configuration or code changes, you now have a broad view of your environment including entity relationships, code-level detail, and user experience—all in context. One Dynatrace customer, a large airline, has 432 million topology updates per day across 2,500 hosts. [See how you might visualize your own Smartscape environment topology here.](#)

### ★ Essential tip 3:

Increase your confidence with [Davis](#), our unique AI engine that sits at the core of the Dynatrace platform and automatically processes high-fidelity logs, metrics, traces, and real user data from hybrid, hybrid environments to serve up precise answers, prioritized by business impact. To learn more of how Dynatrace's AI is industry-leading and unique, see our ebook, [AIOps Done Right](#).

## 02

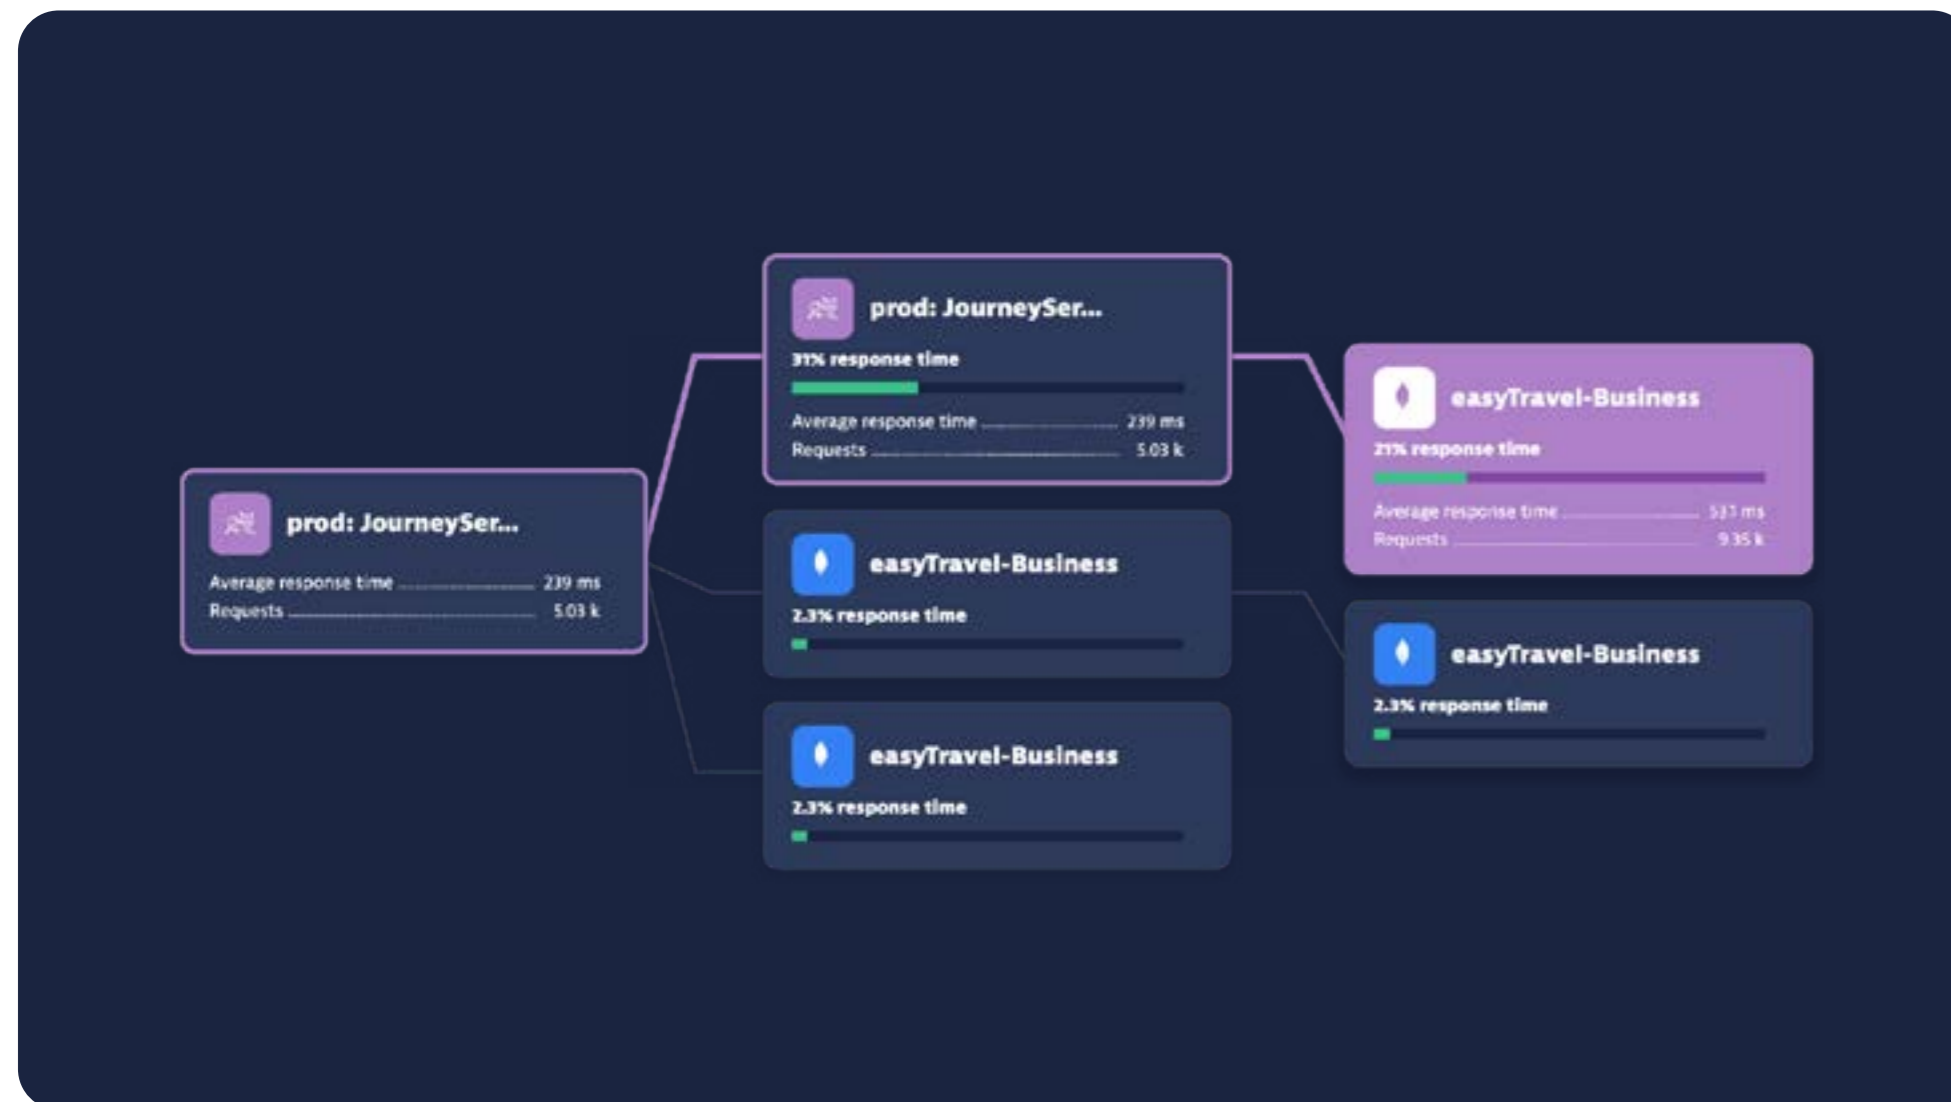
# Leverage automatic and intelligent observability for microservices, containers, and serverless

Microservices, containers, and serverless cloud technology bring speed and agility to the modernization process. They help teams become more flexible and able to bring highly scalable apps to market faster.

However, the dynamic nature of these technologies spinning up and down within seconds introduces several major issues that can impact performance. This extreme complexity bogs down your IT team, because they don't have important and critical insights, for example:

- They can't see the relationships between these cloud technologies and the upstream components that impact them. It's like flying blind when trying to determine the root cause of issues.
- They can't connect end-to-end tracing from real users accessing the microservices, to the nodes, to the services and containers they run. While it's possible to track and tally connections, it requires time and effort and is prone to errors.
- They don't have real-time visibility into exactly what's working inside the workloads running within the containers. This is a big gap in understanding the impact on interactions with cloud and web-based services.

The environment demands more than traditional observability, and it should be an integrated and seamless part of managing your AWS cloud environment.



Dynatrace's automatic and intelligent observability delivers a broad view of AWS hybrid cloud environments that includes metrics, logs, and traces and a full topology with entity relationships, code-level detail, and user experience—all in context. It provides:



Automatic discovery of all cloud technologies at start-up (along with things running inside workloads)



Dependency mapping that enables a multidimensional view of your interdependencies, including all connections between containers inside pods, their services, worker nodes, and data center, along with incoming and outgoing connections along the vertical stack



Impact measuring at scale all the way from the pod, through the cloud and application to the users, which highlights hot spots, identifies problems, breaks down silos, and delivers business impact along with root cause and actionable answers

This leaves no digital stone unturned and enables near-instant insight into any impact on your application. This is the only AI-based intelligent automation that identifies and prioritizes alerts from applications and infrastructure without changing code, container images, or deployments.

With this critical information in one centralized interface, all teams within the software development lifecycle will be able to operate from a single source of truth and resolve issues faster. Another benefit? Zero code changes are required to collect and make sense of this data at scale.

## Staying on top of the latest enhancements of AWS Lambda

AWS Lambda is the fastest growing technology for serverless workloads and helps developers innovate faster. It removes the burden of managing underlying infrastructure and is becoming very popular for cloud-native application environments.

The platform is excitedly evolving with new features and functionalities and Dynatrace—with our close partnership with AWS and the AWS Lambda team—supports these enhancements right from the start. This gives you the early, fast benefits from improved performance and a head-start on your serverless projects and objectives.

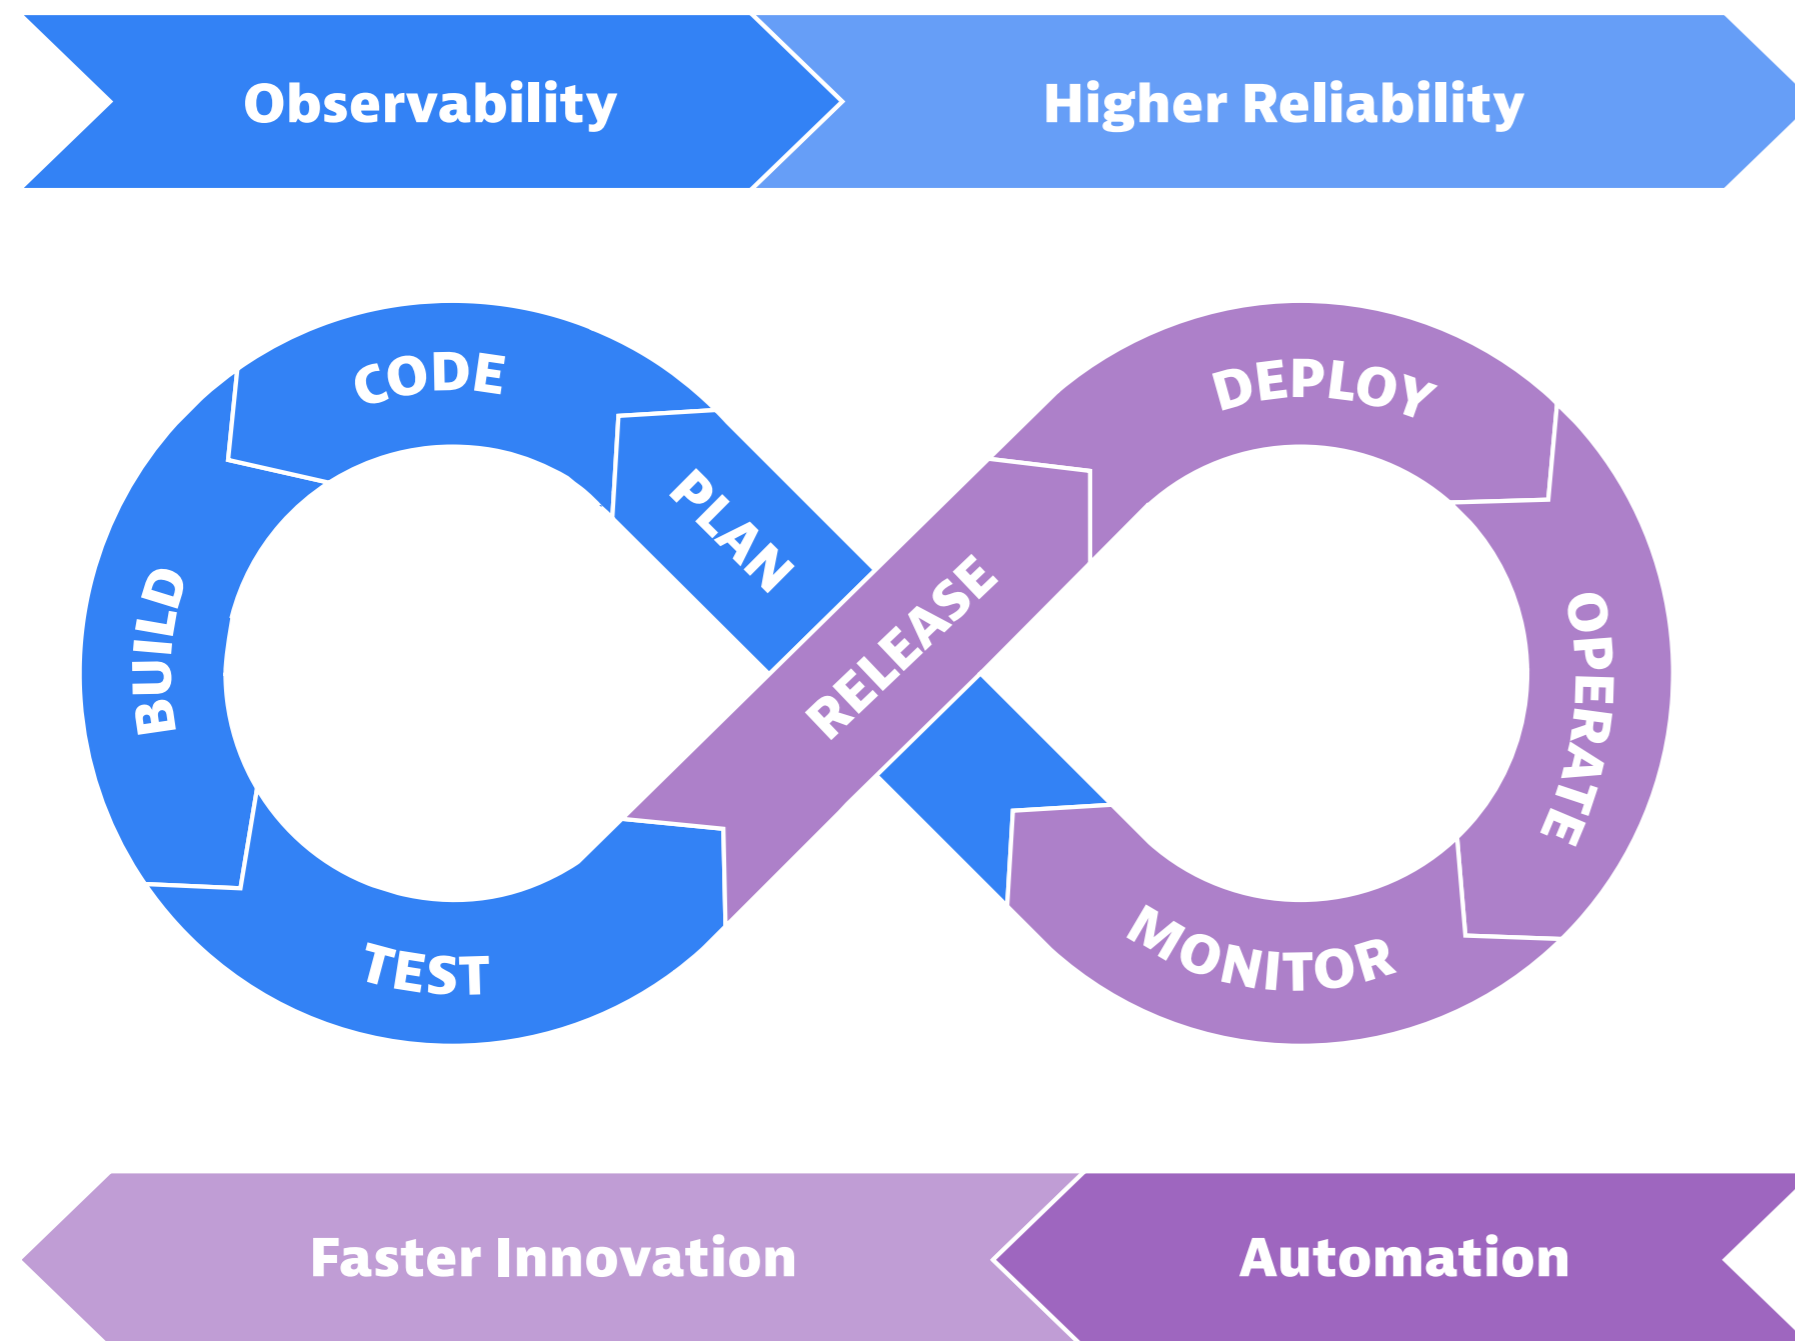
Dynatrace's automatic and intelligent observability boosts end-to-end visibility and helps you understand how the whole environment impacts your customer-facing web applications, mobile apps, or APIs through:

- Automatic and intelligent end-to-end observability of AWS Lambda functions
- Seamless end-to-end distributed tracing/visibility
- Automatic observability and root cause analysis for DevOps, cloud and apps teams
- Insights into how serverless functions are impacting customer facing applications

We help you increase scalability, lower costs and improve the performance of your application by using AI and automation to better understand and optimize the platform.



AWS Lambda Ready



**Automation and Observability**  
are fundamental to DevOps

## 03

### Use automatic and intelligent observability for advancing DevOps and SRE practices

As DevOps teams and SRE practices mature, these teams are realizing that the need for automated and intelligent observability isn't just for operational performance. They are embracing the value of this as a "shift left" strategy to build in better, more automated processes in earlier stages of development and pre-production.

In dynamic, cloud-native architectures, ensuring high-quality, secure, and fast release cycles is incredibly time-consuming, with manual processes and siloed visibility across the application lifecycle. The need for automated operations, tighter collaboration, and continuous delivery has become critical to keep pace with digital business demands to release faster, more efficient software that meets the organizations business at scale.



Dynatrace can transform how development, DevOps, and Site Reliability Engineering (SRE) teams develop and manage cloud-native applications by infusing automatic and intelligent observability to increase software quality with Shift-Left SLO and increase reliability with Shift-Right progressive delivery.

**Automation and orchestration for quality checking** of pre-production applications against SLO-based quality gates to ensure error-free code moves through each stage of the delivery pipeline, without manual checks or intervention.

**Automated, closed-loop remediation** of releases that fail in production, including roll-back of canary or blue/green deployments, management of feature flags, orchestration of remediation runbooks, triggering incident management workflows, and precise context for developers through AI-powered root-cause analysis.

**Automatic release inventory and version comparison** to continuously evaluate the performance of individual release versions and, if an issue is found, automatically restore the most stable version, helping teams consistently deliver the highest quality releases.

This raises the bar and transforms how companies can develop and manage cloud-native applications. You can release better software faster, at the speed the business requires, through the addition of automatic and intelligent observability earlier in the pipeline. The standard is built in with the infusion of AI and automation. It fosters greater collaboration between DevOps teams and SREs to understand customer requirements and define SLIs and SLOs. By leveraging automation, teams spend less time on issues and more time on innovation for the business and customers.

## Gartner® on observability in DevOps and SREs

“Applications that have been developed with observability will allow for significantly faster investigations into outages, particularly in complex environments. This will result in shorter resolution times and improved development efficiencies allowing fast turnaround times for new features.”

— Gartner, Innovation Insight for Observability, Padraig Byrne and Josh Chessman, September 28, 2020


GARTNER is a registered trademark and service mark of Gartner, Inc. and/or its affiliates in the U.S. and internationally and is used herein with permission. All rights reserved.

### Read more:

[Accelerate DevOps and Scale SRE with Service Level Objectives \(SLOs\)](#)

[A guide to event-driven SRE-inspired DevOps](#)

## Drive better business outcomes



## 04 Use automatic and intelligent observability to go beyond operational use cases to strategic performance and business optimization

Almost every organization today strives to differentiate through customer experience. Accelerated by the global pandemic, digital channels have become even more dominant. Rapidly changing preferences have changed how organizations interact with their customers and how they deliver goods and services. Across every industry, competitors are one click away, affecting how we view and protect customer loyalty. As a result, the emphasis on user experience — that subset of customer experience specific to user interaction with your websites and apps — has similarly increased.

Delivery of zero-friction digital experience is essential, and the bar is very high. The digital experience is measured against the best-in-class from Google, Facebook, and other digital leaders.

If your observability tool fails to give you a clear answer as to why there is a sudden drop in your conversion rate, you'll certainly never know that the reason is a failing microservice. For example, if you can't see business outcomes and observability data in the same distributed trace, you'll have to capture transaction path details from the user's click to every single database call, which makes it impossible to identify the actual root cause of a decrease in conversions.

## Dynatrace puts digital experience at the forefront

Dynatrace's automatic and intelligent observability helps organizations understand how customer services are being consumed, which in turn helps them prioritize optimizations and improvements based on greatest business impact. It provides a feedback loop from back-end technology teams to product, digital, and business teams, ensuring the entire cloud stack is supporting expected outcomes.

### Organizations can gain:



Complete insight of technology's impact on user experience and business KPIs like revenue, conversions, and feature adoption



Observability and monitoring across web, mobile, and IoT to gain understanding of the holistic user experience across channels



A simple, all-in-one platform to optimize end-user experience for both customers and employees, no matter where they are in the world

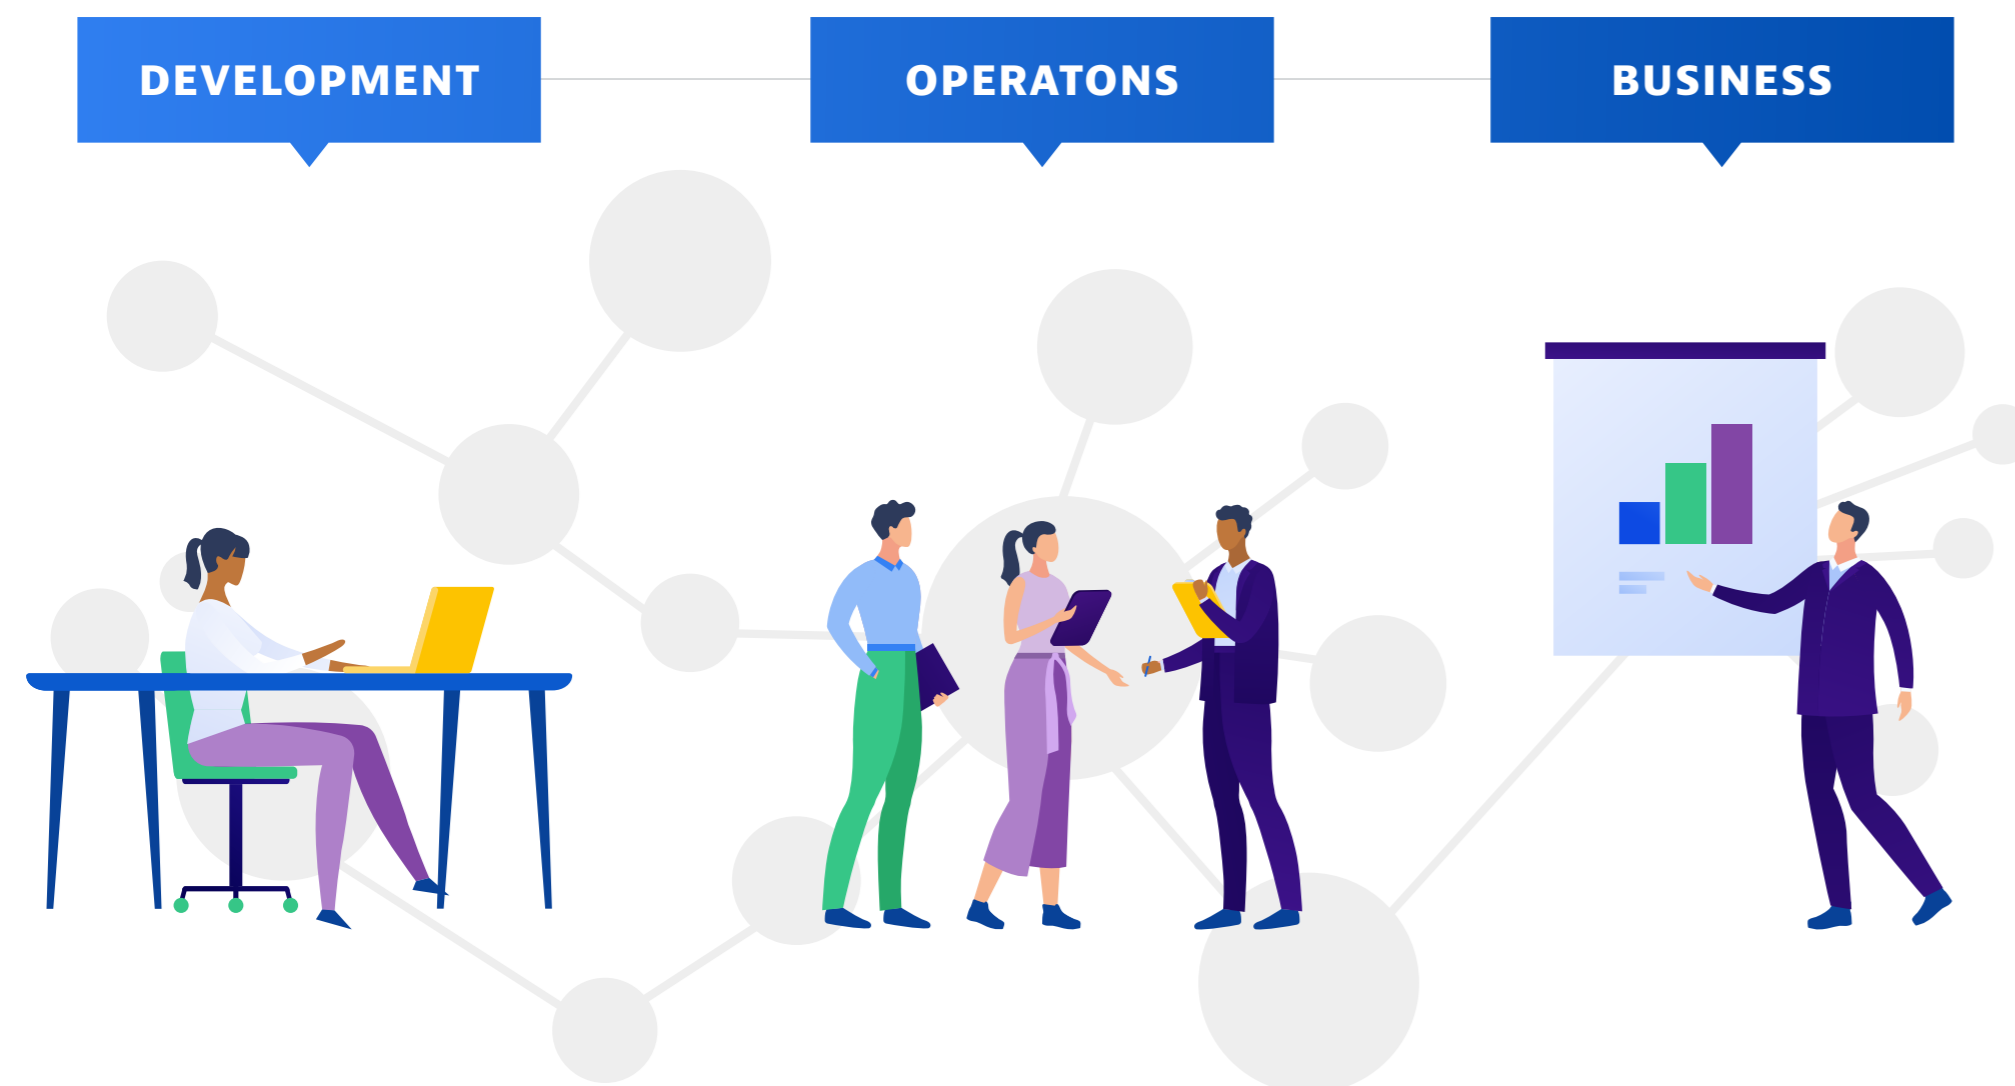
Dynatrace is the only solution to capture the full visibility of customer experience across every digital transaction and provide gapless insight into the complete end-user journey on your app. With these unique capabilities, you can go beyond operational use cases to strategic performance business optimization.

## Vitality focuses Dynatrace on customer experience

"We have successfully built our cloud-native applications on AWS, and Dynatrace's AI and automation ensure they are fast, efficient, and predictable. Dynatrace's deep integrations with AWS, paired with its AI expertise, enables us to find anomalies in our applications and user journeys before they impact business outcomes. The platform's automation has enabled us to improve customer experience through faster responses to customer requests and freeing up time for our teams to innovate."

— David Priestley, Chief Digital Officer, Vitality.

## Digital Transformation is a team sport



## 05

### Use automatic and intelligent observability for tighter cross-team collaboration of biz dev and ops

Tool sprawl aggravates silos, hurts innovation, decreases application quality, and reduces collaboration. Each different tool amplifies the negative effects spreading across each team, leading them to continue struggling to identify and resolve issues and optimizations with the highest impact. Pressure to resolve problems without understanding cause and effect can result in a poor understanding of what to prioritize. After efforts, teams are often at a loss to confirm that they have received the desired business outcome.

What contributes to collaboration challenges?



#### Data

The lack of connective tissue inflicts time-consuming and error-prone joining of disparate data models.




#### Environments

Isolated observability and monitoring across pre-prod and production environments hurt speed and quality of Shift-Left efforts for DevOps and SRE teams.



#### Platforms

Multiple tools for multi-cloud or hybrid cloud platforms create observability blind spots for infrastructure and platform operators.




#### Teams

When each team receives alerts and symptoms in a vacuum, problems and blame are passed "over the wall" to others.

Dynatrace's automatic and intelligent observability platform is a pro collaborator for development, operations, and business teams. It has rich capabilities like [Digital Business Analytics](#), [Session Replay](#), [Real User Monitoring](#), and more that allow you to capture the right metrics and communicate across teams in the same language.

How does this work? Several key requirements enable teams to collaborate more efficiently towards the same technical and business goals:



Single data model to scale observability across all layers and components across the full tech stack



Shared context that facilitates cross-team collaboration, with flexibility to slice and dice across infrastructure, applications, operations, and business data



Seamlessly ties together the entire software lifecycle from feature development, testing, releases, and ongoing optimizations to innovate faster with higher quality

Support your entire digital team — from Development to Operations to Business — with a platform that works seamlessly across your full stack to improve collaboration and drive better business outcomes. Dynatrace eliminates silos and brings teams together through a common language, bridging gaps with a single source of truth. This encourages collaboration and accelerates the speed of value-add product features and optimizations that drive better user experience and greater business impact.

## Effective BizDevOps collaboration starts with good analysis

Dynatrace's Digital Business Analytics and Insights connects the application performance and user experience to business KPIs and action. IT can prioritize resources and investments that can most effectively deliver desired business outcomes.

Precise end user performance and quality data allows you to identify and understand how and what to optimize; you can know in real-time when business KPIs — conversions, quotes, payments, registrations, purchases, etc. — degrade and fall short of expected performance.

You'll know how user experience, app performance, errors, and new features or releases impact these KPIs, and you'll have a common view of business metrics — including page names and audience segments — through a shared business lens.

With one platform that delivers deep visibility and continuous feedback, you'll be alerted to business anomalies courtesy of our AI engine, which packages root cause and actionable answers to improve business KPIs.

# Summary

As companies move to digital services via enterprise cloud, the limits of traditional monitoring tools designed for earlier architectures are not adequate. Organizations need a new way of understanding these dynamic environments, even beyond observability. Today, automatic and intelligent observability is required to manage these vast, complex cloud technology environments.

This means that insights are available to all of the stakeholders and teams that support the delivery of business services.



**For operations,**  
it's optimal performance and reduced time to find and resolve issues.



**For dev teams,**  
it's better data and automation, less manual work, and tighter collaboration with operations.



**For business teams,**  
it's the right services and customer experience.

These all add up to better, more successful competitive business impact.

These five game-changers give you an example of how Dynatrace can uniquely help you master application and business performance in your AWS cloud environment. It integrates and enhances AWS's CloudWatch and enables operations, DevOps, SRE, and business teams to raise the bar in bringing better performance and business results to an organization.

## Dynatrace as an industry leader in observability

We believe, Dynatrace is the only observability platform that places applications and services in real-time context with the environments they run on, providing distributed tracing enriched with code-level analysis and digital experience data. In fact, in the 2022 Gartner® Magic Quadrant™ for Application Performance Monitoring, Dynatrace was named a Leader, positioned best for completeness of vision. Dynatrace also scored highest in 4 of 5 Use Case in 2022 Gartner Critical Capabilities for Application Performance Monitoring<sup>3</sup>. We believe, this recognition demonstrates its usefulness as the solution for drawing key insights from observability-optimized environments.

**Download a complimentary copy of this report.**

[Download](#)

# Take action: Begin accelerating your transformation now

[Explore the platform >](#) [See support technologies >](#)

## Test drive Dynatrace.

[Start free trial](#)



### About Dynatrace

Dynatrace provides software intelligence to simplify cloud complexity and accelerate digital transformation. With automatic and intelligent observability at scale, our all-in-one platform delivers precise answers about the performance and security of applications, the underlying infrastructure, and the experience of all users to enable organizations to innovate faster, collaborate more efficiently, and deliver more value with dramatically less effort. That's why many of the world's largest enterprises trust Dynatrace® to modernize and automate cloud operations, release better software faster, and deliver unrivalled digital experiences.

Gartner, Magic Quadrant for Application Performance Monitoring Federico De Silva, Padraig Byrne, Josh Chessman, 9 April 2022 Gartner, Critical Capabilities for Application Performance Monitoring, Padraig Byrne, Federico De Silva, Josh Chessman, 14 April 2022 Gartner does not endorse any vendor, product or service depicted in its research publications and does not advise technology users to select only those vendors with the highest ratings or other designation. Gartner research publications consist of the opinions of Gartner's Research & Advisory organization and should not be construed as statements of fact. Gartner disclaims all warranties, expressed or implied, with respect to this research, including any warranties of merchantability or fitness for a particular purpose. GARTNER and MAGIC QUADRANT are registered trademark and service mark of Gartner, Inc. and/or its affiliates in the U.S. and internationally and are used herein with permission. All rights reserved.