



5 challenges to achieving observability at scale

Using automation and intelligence to overcome obstacles.



Introduction

Successful digital transformation requires every application and digital service, and the dynamic multi-cloud platforms they run on, to work perfectly. All the time.

But these dynamic, highly distributed cloud-native technologies are fundamentally different than their predecessors. The resulting complexity brought on by microservices, containers, and software-defined cloud infrastructure is overwhelming at web scale. It's all beyond the limits of human teams to manage and scale on their own.

To understand everything going on in these ever-changing environments, all of the time, observability needs to scale.

What teams are up against

6

Challenge One

The complexity of dynamic multi-cloud environments

10

Challenge Two

Monitoring dynamic microservices and containers in real time

14

Challenge Three

The volume, velocity, and variety of data and alerts

18

Challenge Four

Siloed Infra, Dev, Ops, Apps, and Biz teams

22


Challenge Five

Knowing what efforts drive positive business impact

More tools aren't the answer

Some teams mistakenly try to solve this 'observability at scale' problem by adopting more siloed monitoring tools — and spending more time on manual configuration, incurring more technical debt, and struggling to identify issues and prioritize efforts with the greatest impact.

As cloud complexity continues, this approach becomes increasingly unsustainable for even the most experienced teams, who are continuously bogged down in manual-intensive tasks that decrease effectiveness to achieving what matters most.



of applications in enterprise organizations are **not monitored** due to siloed tools and burdensome manual effort.

The shift to intelligent observability

To scale observability, **enterprise organizations must fundamentally transform the way they work** to innovate faster, keep up with constantly changing tech stacks, and reduce risk across teams.

This scale happens when teams shift from simply observing and reacting to issues as they arise, to a culture of proactive understanding and optimization. This unlocks the ability to anticipate, predict and even auto-remediate problems that matter most to the business.

In deciding how to accelerate digital transformation, companies need to understand that every decision is an investment in achieving the original goal of observability: to proactively and efficiently improve user experiences that drive business growth.



Automation and intelligence are essential

Whether selecting a DIY approach, buying another cheap tool, or investing in a strategic platform, everything costs time, money, people, and quality. Prioritizing value and speed of delivery to the business and customers is paramount to finding success in this dynamic multi-cloud world.

Automation and intelligence are essential to transform how teams work to quickly and efficiently achieve observability at enterprise scale.


Requirements		Results
Complete coverage	→	More productivity and time to innovate
Automation everywhere	→	Higher quality releases
Real-time feedback	→	Better customer experiences
Precise answers	→	Reduced risk
Cross-collaboration	→	Accelerated business outcomes

Challenge 1:


The complexity of dynamic multicloud environments

The rate at which new technologies are available and implemented is increasing, exploding the complexity that results from unmanageable volumes and speed of data emitted by dynamic environments.


This makes it near impossible for IT teams to manually understand how everything is related in context, all of the time. **So, teams must find ways to automate the understanding of this data and context to accelerate digital transformation.**




Teams often fail at digital transformation because they're:



Hindered by disconnected data silos
that prevent understanding of entity relationships and interdependencies




Lacking understanding and context of upstream and downstream system impacts from potential changes



Forced to prioritize manual instrumentation and mundane tasks over developing new features

These shortcomings introduce unnecessary risk and burden developers with repetitive toil,
ultimately hurting digital transformation efforts and driving innovation forward.



of an IT teams' time is spent on manual tasks, on average.

— [Dynatrace 2020 Global CIO Report](#)

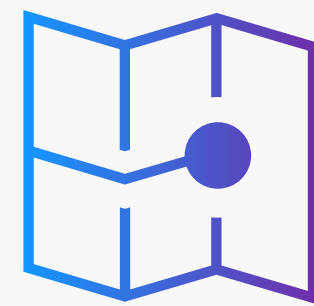
How to overcome it

Automation is an absolute necessity to not only handle the scale of every single component in an enterprise ecosystem, but also understand all the interdependencies.

You can't hire your way to observability at scale. Understanding dynamic multicloud environments requires an automated approach that can multiply productivity of your existing team and shift effort from manual tasks to driving tangible business results.




To scale observability and eliminate blind spots across increasingly complex and expanding environments, teams need automation powered by:



Topology mapping

that continuously maps components, cloud services, and ever-changing relationships between potentially billions of interdependencies



Auto-discovery

of new components to prevent gaps in coverage in real-time



No-code approach

to better leverage skilled developers on proactive optimization efforts and business-driving innovation projects


This continuous automation and always-on context gives teams confidence in keeping up with dynamic technology stacks to digitally transform faster, without the ongoing burden of constant deployments and manual maintenance in attempt to slowly gain more coverage and understanding.

Challenge 2:


Monitoring dynamic microservices and containers in real-time

Short-lived containers and microservices, like those managed in Kubernetes, provide the required speed and agility to successfully modernize. However, the dynamic nature of technologies that can spin up and down within seconds introduces several major issues to scaling observability for these technologies.

This all results in a lack of understanding of internal states of the application, other interdependent components that microservices rely on, and even the impact on users.



IT teams are still blind to what's happening in their dynamic environments and actioning on incomplete data because they:




Don't understand the relationships
between containers
and upstream components
that can impact them



Can't connect end-to-end tracing from real users
accessing these microservices,
to the nodes, the services
and containers they run on



Lack real-time visibility
into exactly what's inside
the workloads running
within containers

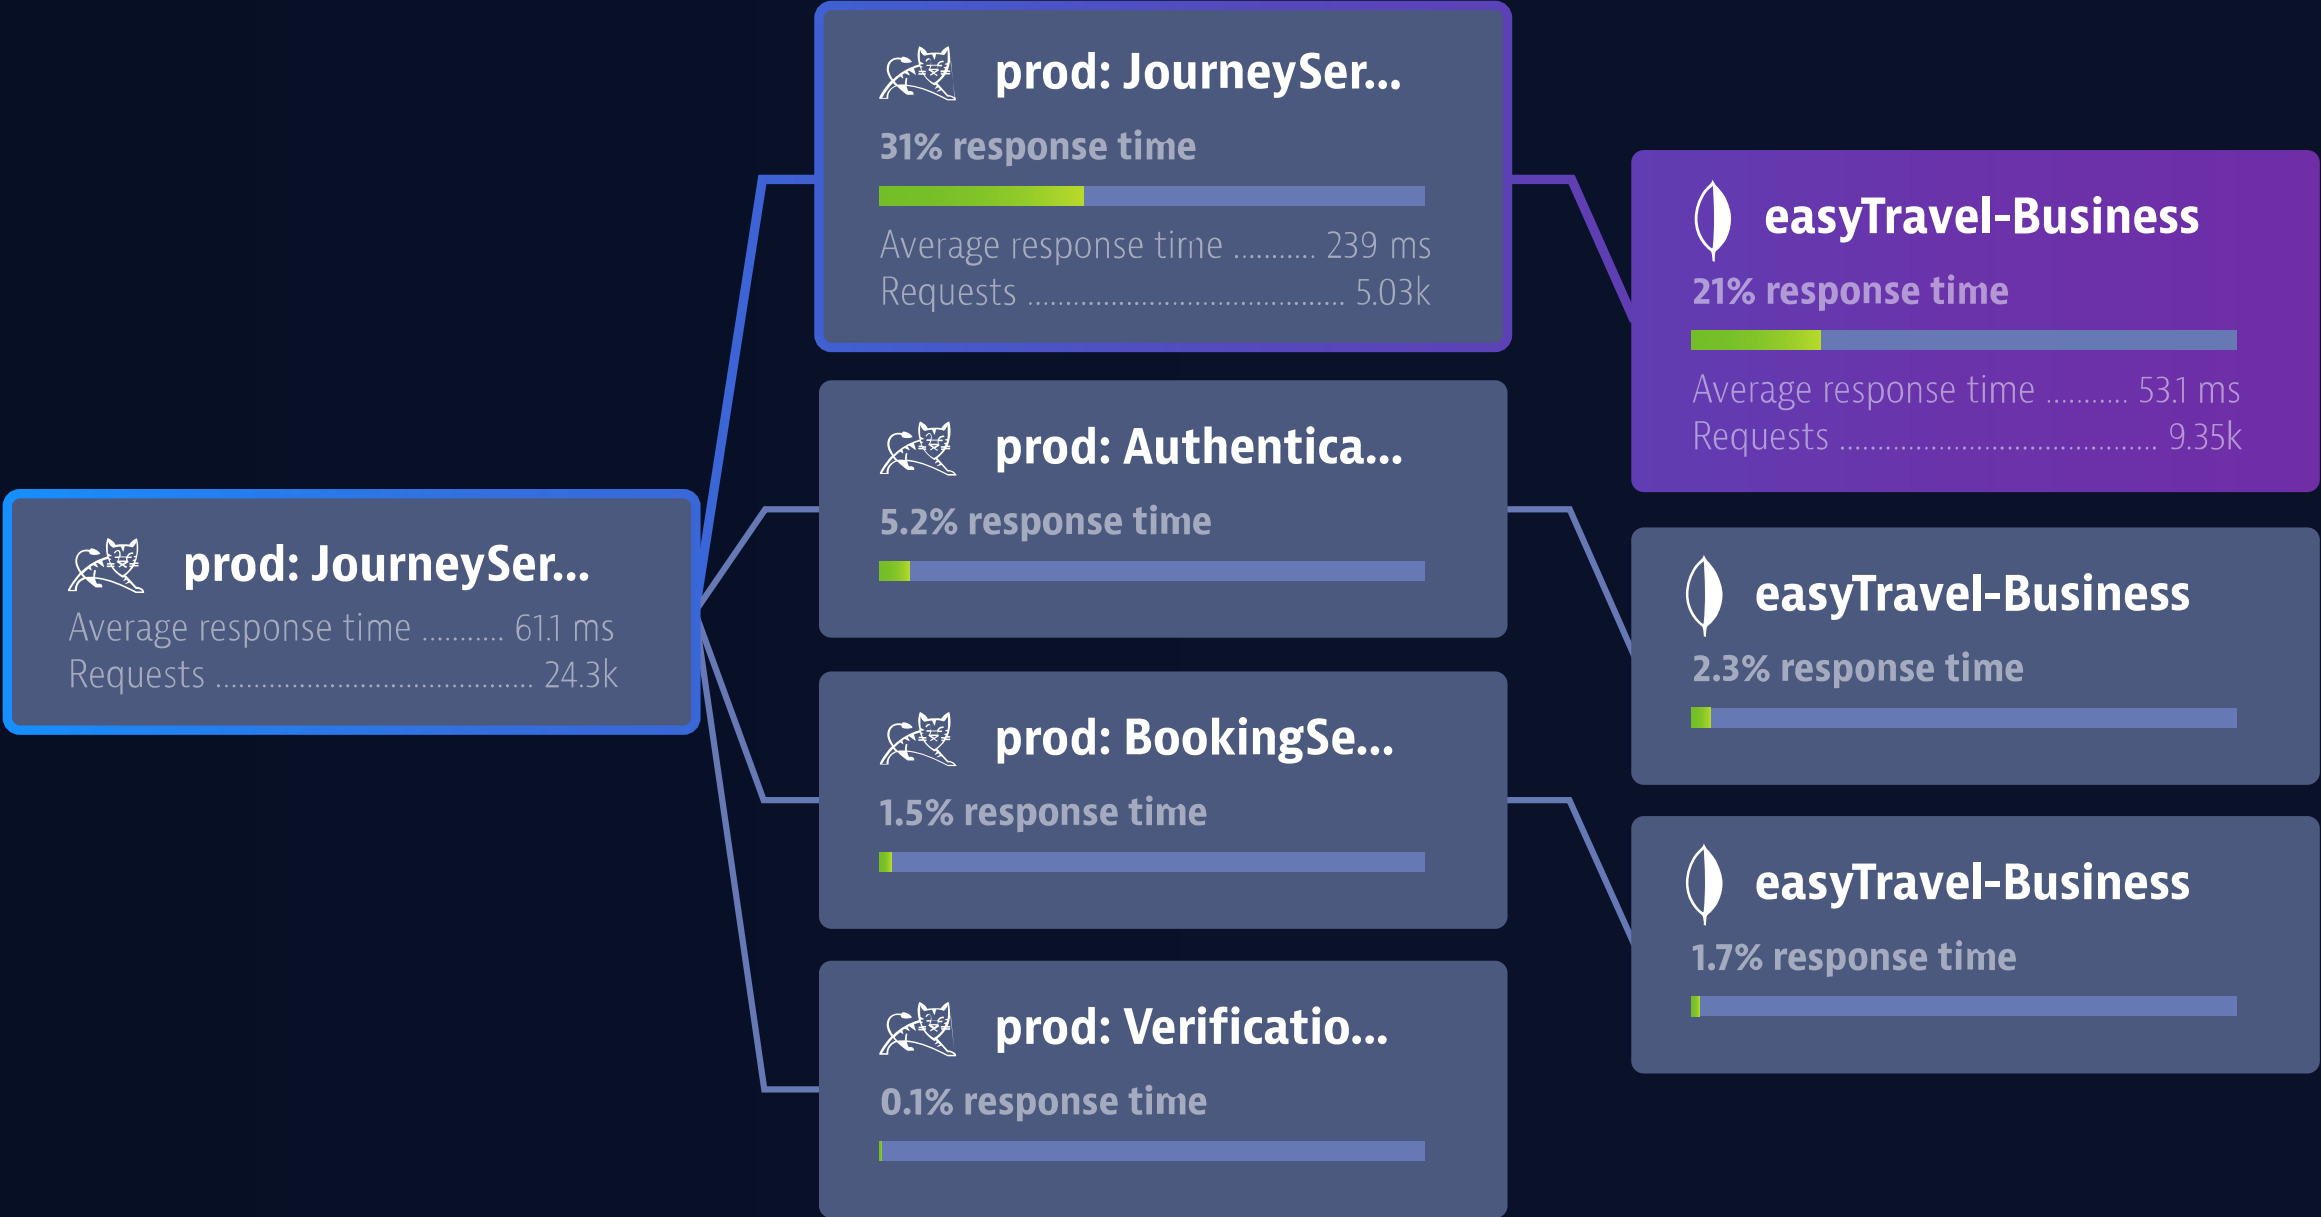


of CIOs say
monitoring containerized
microservices in real-time
is almost impossible

— Dynatrace 2020 Global CIO Report

How to overcome it

Enterprises need observability to scale across their multicloud, including cloud, legacy, and hybrid environments, to handle the dynamic nature of Kubernetes and containers.



To ensure everything's accounted for, no matter how short lived,
teams need real-time intelligence and automation with:




Automatic discovery

of containers at start-up,
along with all things running
inside each workload



Topology context

external to containers,
since anomalies often occur
outside of Kubernetes
nodes, pods, containers,
and clusters



Full-stack visibility

all the way from the pod,
through the cloud provider
and application, to the user
to understand the end-to-end
business impact

With this speed, automation, and context applied to containers and microservices,
**IT teams can continuously understand system behavior and the true origin of anomalies
can be easily isolated and precisely pinpointed at scale.**

Challenge 3:

The volume, velocity, and variety of data and alerts

Dynamic multi-cloud environments are exponentially increasing the amount of telemetry data emitted, and overwhelmed teams are still stuck trying to monitor every data point and make sense of it all.

Already constrained IT resources are stuck reacting to each new problem after users and business goals are already impacted, trying to observe what's happening by manually building, maintaining, and constantly watching potentially thousands of dashboards.



However, **this approach doesn't scale** and persists the same challenges that cannot be solved using the same manual-intensive philosophy:



Defining and redefining "normal" for anomaly thresholds that constantly change with dynamic environments and seasonality



Monitoring "unknown unknowns" — issues you aren't aware of, don't understand, and don't monitor



Siloed data sending mixed signals that multiply alert storms, intensifying team fatigue and unnecessary war rooms



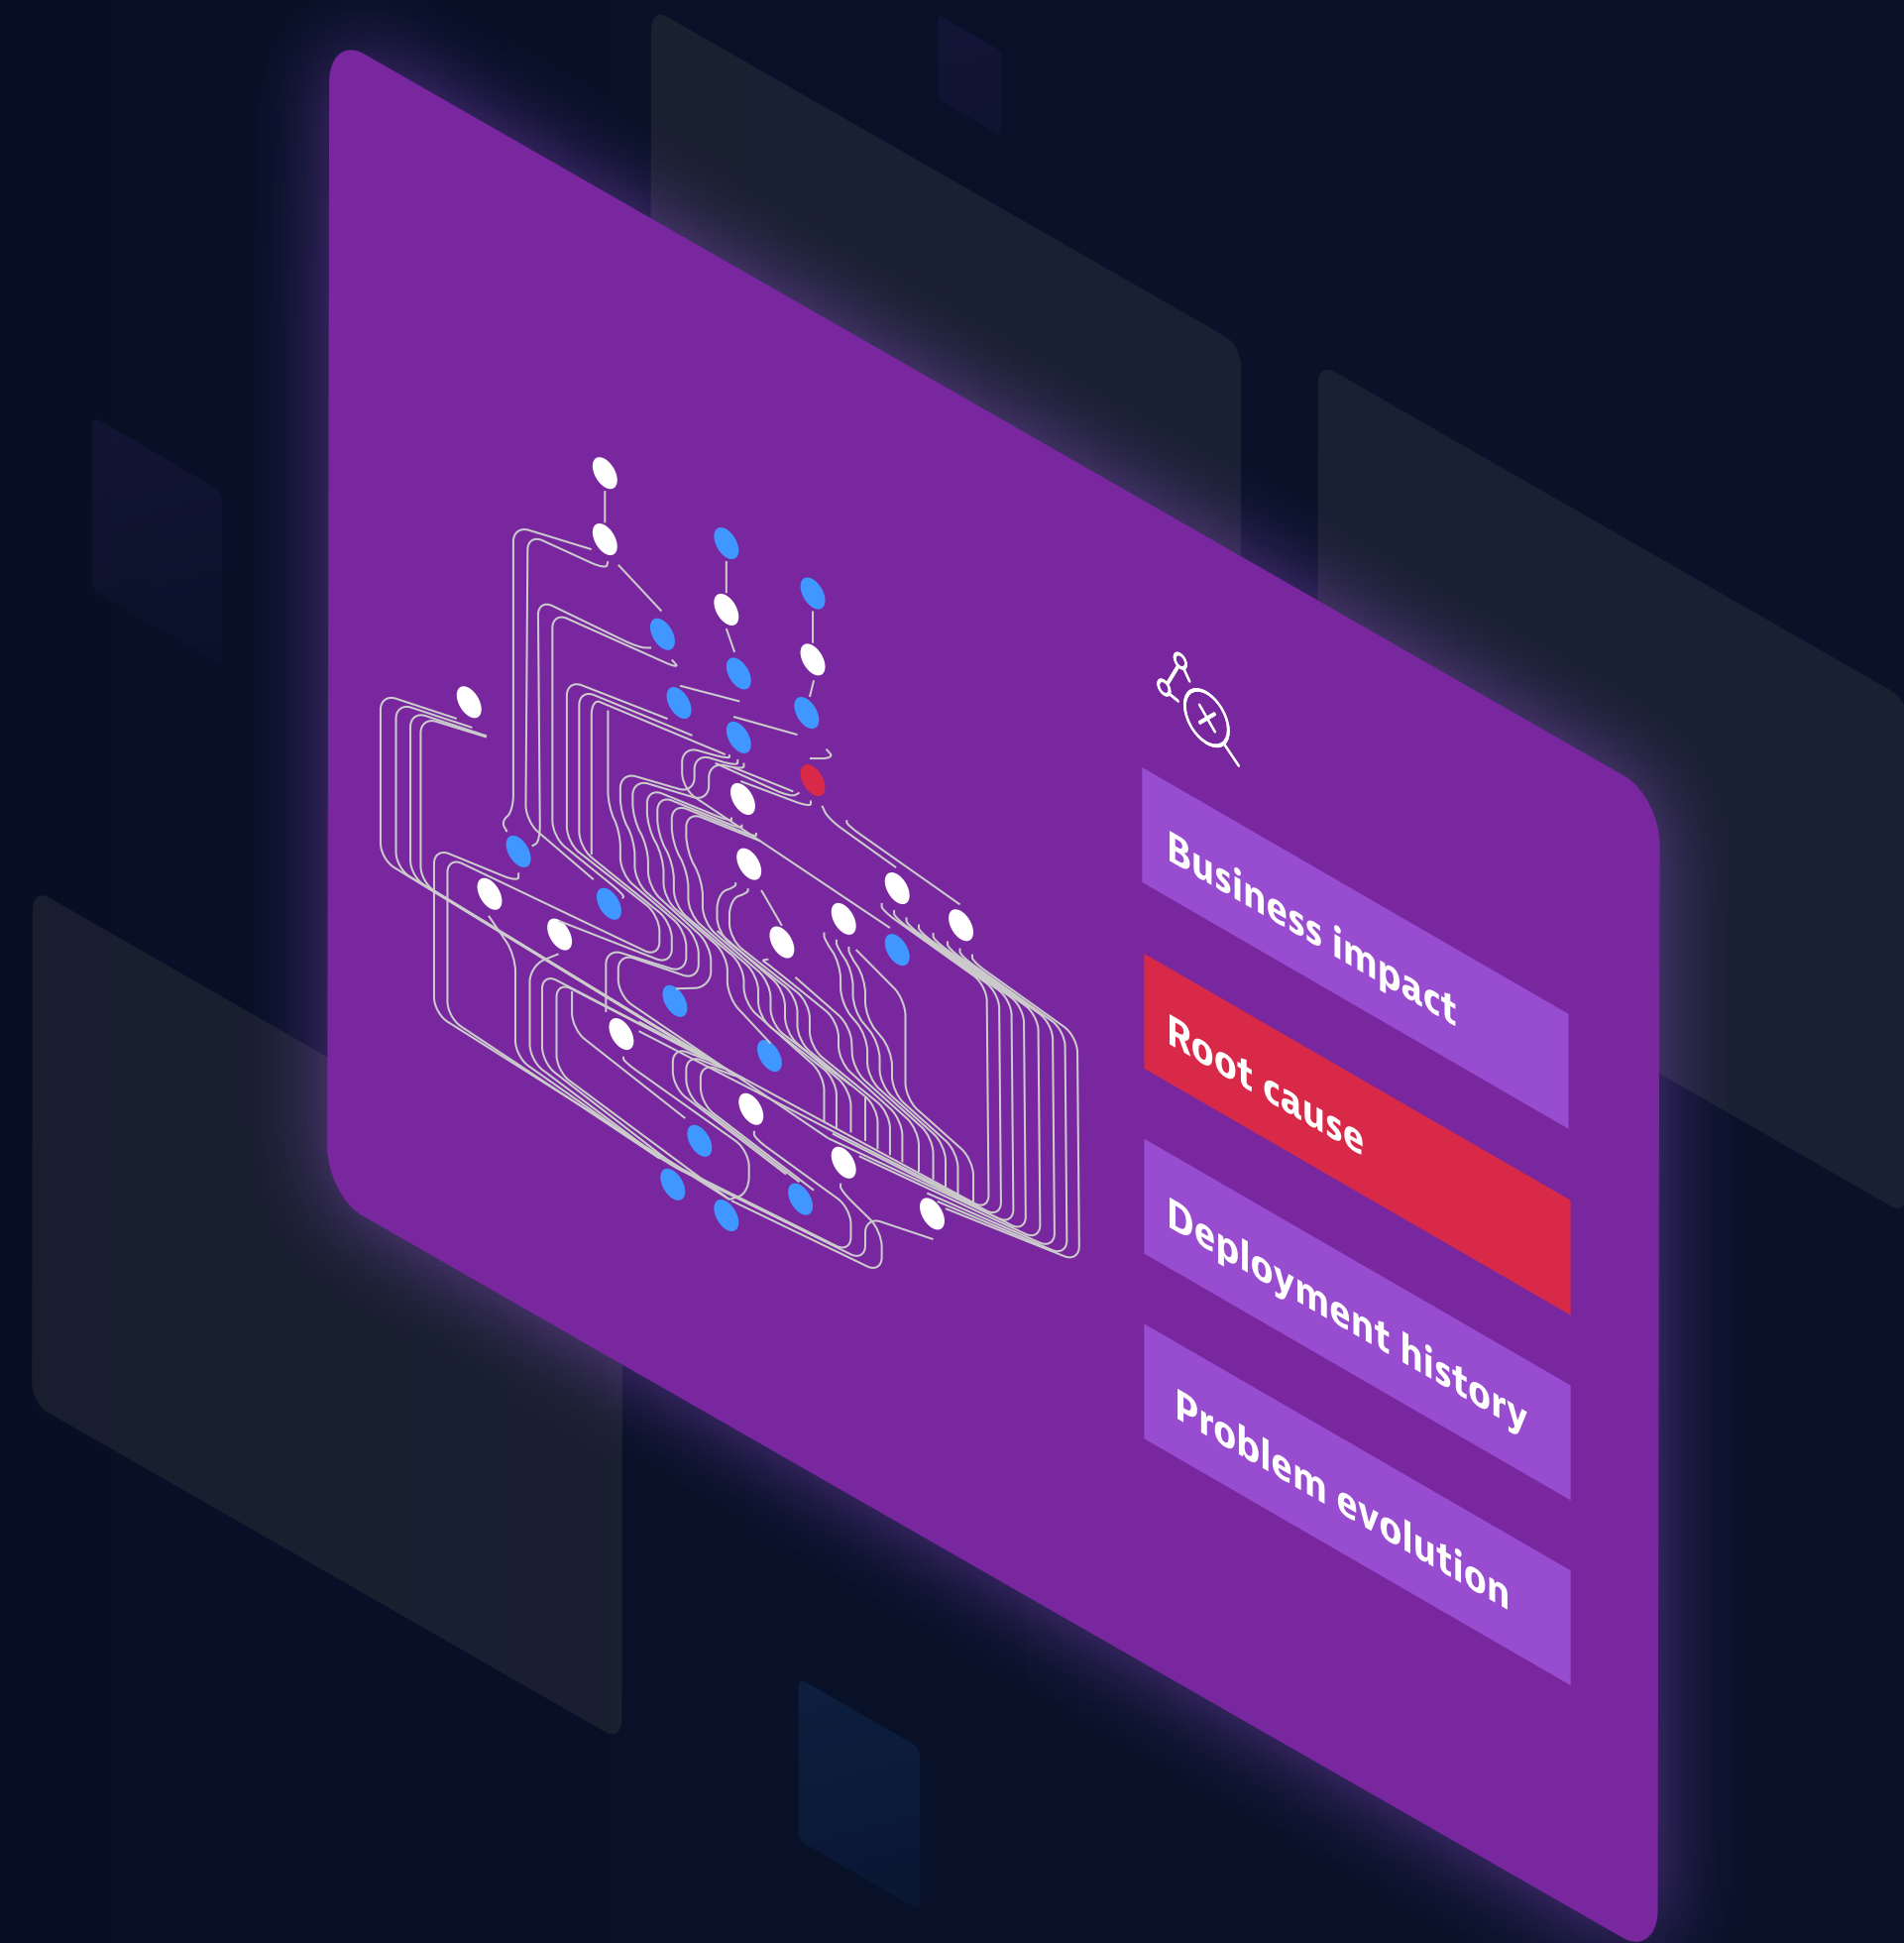
Multiple teams struggling to pinpoint issues across different tools to guess the root cause, causing more finger-pointing and blaming

All of this forces teams to spend even more of their time "keeping the lights on" by guessing about the problem, priority, and diagnosis, rather than continuously optimizing and resolving issues before users are impacted.

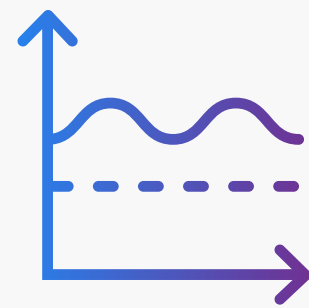
How to overcome it

It's clear that AI is needed to continuously and instantly understand when and why anomalies occur. But the only way to transform from reactive to proactive, is having an AI that doesn't need to learn or be trained.

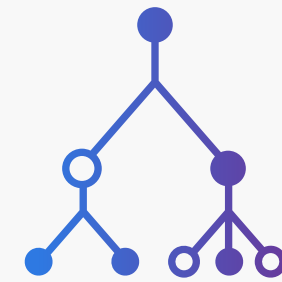
Because dynamic multi-cloud environments can change within seconds, AI needs to know precise answers and be able to anticipate and auto-remediate issues before business impact.



A few critical capabilities of AI
that **enable observability at scale:**




Auto-adaptive threshold baselining for anomaly detection to prioritize what really matters



Intelligent grouping of related anomalies into a single problem to eliminate redundant work across teams



Always-on causation-based AI with code-level analysis that processes billions of dependencies with complete fault tree analysis to instantly deliver answers



Integrating answers with context from external systems (like ServiceNow and other ITSMs) to broaden workflow automation across multiple teams

The goal of causation-based AI is to provide answers to engineering, infrastructure, operations, and application teams and empower them to focus on the things that matter. **Delivering one precise answer for each issue that everyone understands can transform teams away from finger-pointing to efficient cross-team collaboration that drives business outcomes.**

Challenge 4:

Siloed Infra, Dev, Ops, Apps, and Biz teams

New cloud-native technologies require more solutions to instrument and monitor, but teams are already drowning in tool sprawl. This tool sprawl aggravates silos that hurt innovation, decrease software quality, and reduce collaboration.




Each different tool and point solution amplifies these silos,
with the negative effects spreading across each team that continues to struggle
identifying and resolving issues and optimizations with the highest impact.



Data

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 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

How to overcome it

To eliminate these silos, a solution can't simply stitch it all together. It has to bring together teams through a single common language. Bridging these gaps with a single source of truth removes confusion and multiplies productivity across teams.

This cross-team collaboration and more efficient working environment boosts the speed of value-add product features and optimizations that **drive better user experiences.**




Several key requirements **enable teams to collaborate more efficiently** towards the same technical and business SLIs/SLOs:



**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 tying together
the entire software lifecycle**

from feature development,
testing, releases, and ongoing
optimizations to innovate
faster with higher quality

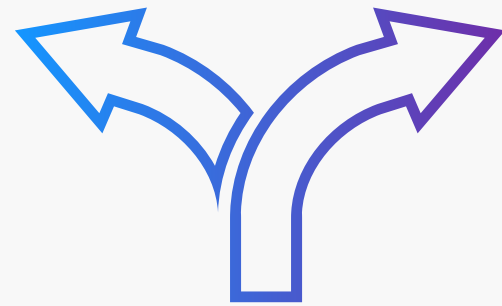
Challenge 5:

Knowing which efforts drive positive business impact

Even with complete visibility to back-end components, a lack of front-end user perspective diminishes much of the tangible value that organizations aim to achieve with observability efforts.



Without front-end application performance,
major risks to the user experience are exposed:



**Disconnected front-end
and back-end perspectives,**

hurting understanding of
technology's impact on users
and business objectives



Critical blind spots

like mobile app crashes,
3rd party services, CDN,
and front-end errors
still exist



Disparate solutions

to attempt observability
for mobile and edge-device
channels, forcing teams
to leave some
applications ignored



**No consideration
of employees working
from home,**

potentially
damaging their ability
to access required resources
they need to deliver frictionless
customer experiences

**Neglecting the end-user experience of applications obstructs the ability
to prioritize optimizations and issues based on greatest business impact.**

When teams only look at technology by itself, IT efforts may not align with business priorities.

How to overcome it

An outside-in user perspective of the application is needed to create a feedback loop from back-end technology teams to product, digital, and business teams, ensuring the entire cloud stack is supporting expected outcomes.



To include user experience into a more intelligent observability approach,
organizations need to connect front and back-end perspective to 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 to holistic user
experience across channels



All-in-one platform

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

**To achieve observability that scales across channels, customers, employees, and all types of applications,
back-end and front-end application performance must be connected.** Only then can teams across IT, product,
and business prioritize and align efforts that drive the bottom line.

Conclusion

To achieve observability at scale for dynamic multi-cloud environments at the speed needed to exceed customer expectations and business goals, a fundamentally different approach is required.

Continuing to waste effort on manual instrumentation and configuration, digging through siloed data, and working on the wrong things prevents teams from making progress, and ultimately from achieving strategic business goals.

Automated and intelligent observability is needed.

Dynatrace helps transform the way you work with:

Intelligent observability — See it all down to code-level, at scale

Continuous automation — Stay ahead of modern, dynamic multi-clouds

Precise Intelligence — Go from guessing to knowing



Our smarter approach to observability helps teams turn AI into ROI, and drive:

99%

Fewer IT tickets

From 700 tickets a week to just 7.



[Learn more](#)

20%

Higher cart value

Order-from-table mobile application drives higher value than order from bar.



[Learn more](#)

75%

Faster innovation delivered

With 75% MTTR and 4x productivity increase.



[Learn more](#)

Software intelligence for the enterprise cloud

Click the link to take the next step in your digital journey and see what we can do for you.

Learn more

If you're ready to learn more, please visit dynatrace.com/platform for assets, resources, and a **free 15-day trial**.

About Dynatrace

[Dynatrace](#) (NYSE: DT) exists to make the world's software work perfectly. Our unified software intelligence platform combines broad and deep observability and continuous runtime application security with the most advanced AIOps to provide answers and intelligent automation from data at enormous scale. This enables innovators to modernize and automate cloud operations, deliver software faster and more securely, and ensure flawless digital experiences. That is why the world's largest organizations trust the Dynatrace® platform to accelerate digital transformation.

 [dynatrace.com blog](#)  [@dynatrace](#)

