



AIOps Done Right

Automating the Next Generation
of Enterprise Software



Introduction

AI is driving the next innovation cycle in enterprise software¹, enabling new levels of intelligent automation and vertical integration. As today's enterprise systems increase in size, the benefits of digitization and cloud computing go hand-in-hand with technological complexity and operational risks.

AI-powered software intelligence holds the promise to tackle these challenges and enable a new generation of autonomous cloud enterprise systems.

¹AI Technologies — William Blair Industry Report, June 28, 2018

What's inside



Introduction

The promise of AI



Chapter 1

Anomaly detection and alerting



Chapter 2

Getting the best monitoring data



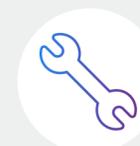
Chapter 3

AI operations and root cause analysis



Chapter 4

Impact analysis and foundational root causes



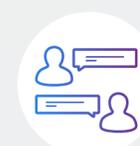
Chapter 5

Auto-remediation



Chapter 6

Automation and system integrations



Chapter 7

Natural language interfaces

Beyond error detection, towards self-healing

Consider this all-too-familiar challenge: An anomaly in a large microservice application triggers a storm of alerts as services around the globe are impacted. As your application contains literally millions of dependencies, how do you find the original error? Conventional monitoring tools are not much help. They collect metrics and raise alerts, but they provide few answers as to what went wrong in the first place.

In contrast, envision an intelligent system that accurately provides the answers—in this case, the technical root cause of the anomaly and how to fix it. Such intelligence, if accurate and reliable, can be trusted to trigger auto-remediation procedures before most users even notice a glitch.

AI and automation are poised to radically change the game in operations. And even more, it's about collecting and applying intelligence along the entire digital value chain, from software development through service delivery all the way to customer interactions. Smart integration and automation will drive the next innovation cycle in enterprise software.

The Promise of AI

Enable autonomous operations, boost innovation, and offer new modes of customer engagement by *automating everything*.



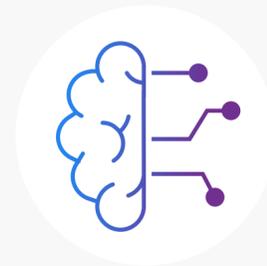
◀ AIOps

Replace a storm of noisy anomaly alerts with accurate and reliable root cause analysis.



◀ Intelligent DevOps

Increase the speed of innovation and software quality through intelligent performance and regression testing.



◀ Auto-remediation

Automate anomaly remediation and performance optimization based on system health and real user demands.



◀ Smart customer engagement

Use business intelligence data to improve customer experience, including automatic remediation of breakdowns and complaints.

A proven record in AIOps

Dynatrace helps the world's most recognized brands to simplify cloud complexity and accelerate digital transformation. Davis—our deterministic, causation-based AI engine—was built into the fabric of the Dynatrace software intelligence platform four years ago, at a time when cloud computing became mainstream and conventional monitoring tools hit a wall. Since then, many leading global brands have relied on Dynatrace to accurately and reliably identify the root cause of performance problems while automating Ops, DevOps, and business processes.

How to avoid closing down 500 supermarkets on a busy Saturday:

Coop, Denmark's largest food retail group, celebrated its 150th anniversary by digitizing its business and moving 80% of their core apps into the cloud.

In 2016, Coop launched their new customer loyalty solution and an updated point-of-sale software. Despite extensive testing, a problem developed soon after the launch in production—checkout registers froze when trying to print out receipts. Suddenly, Coop was facing the prospect of having to close 500 of its stores on a busy Saturday morning because its payment systems were down.

However, two minutes after the first problems occurred in a couple of stores, the Dynatrace monitoring software was able to pinpoint the root cause, a lack of CPU power in the Azure cloud. A major breakdown was avoided by simply spinning up additional resources on the fly.

"Dynatrace is the first I've seen where the AI really shines. Incredible."

Ariel Molina, Sr. Dir., Software Engineering & Enterprise Architecture at Carnival Cruise Line

"Dynatrace, within two minutes came back and said 'you have a problem in your cloud instance', and we spun up extra resources. So we avoided having to close down supermarkets and disappoint customers waiting in line."

Jeppe Hedesgaard Lindberg, Application Performance Manager at Coop Denmark

"We fire up Dynatrace, and immediately the AI goes to work and identifies problems. There's no digging—it's bubbling to the top. It's right there in your face. It just does it for you; it's amazing."

Steve Strout, Director, Platform Engineering Assurant

"The AI paves the way for autonomous operations, enabling us to create auto-remediation workflows that remove the need for human intervention in the resolution of recurring problems."

David Shepherd, Service Delivery Manager, Global IT Service Excellence at Experian

Chapter 1

Anomaly detection and alerting

Insight

The concept of automating operations revolves around better troubleshooting, with the ultimate goal to reduce the Mean Time To Recovery (MTTR). This is accomplished through automatic anomaly detection and alerting, i.e., speedy Mean Time To Discovery (MTTD). However, further reduction of MTTR require automatic root cause analysis.

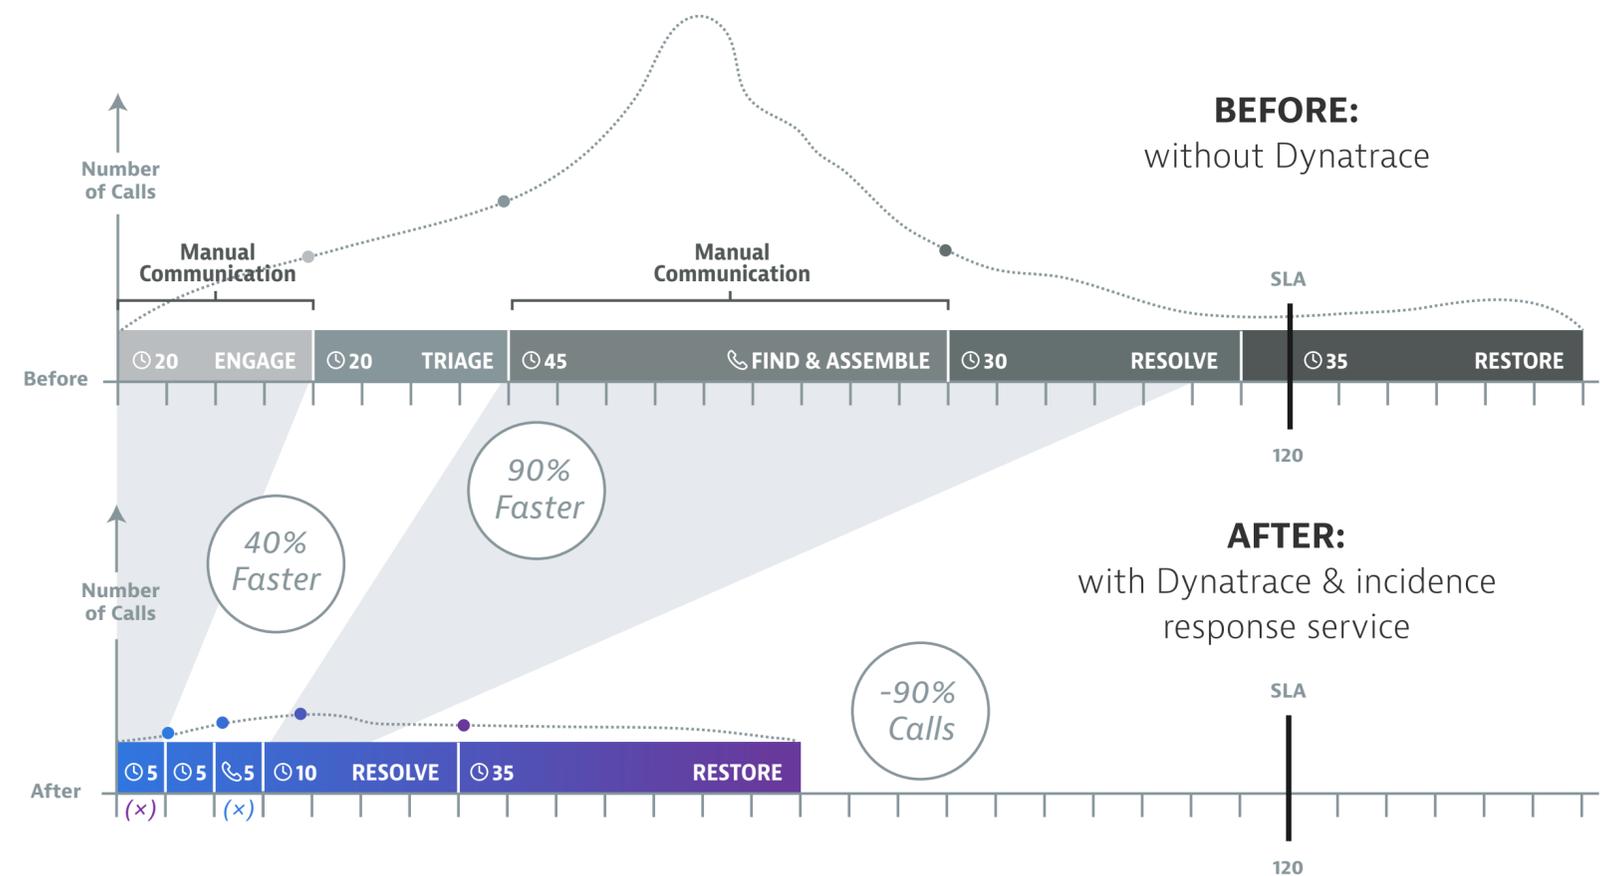
Challenge

Traditional monitoring tools focus on application performance metrics and baselining methods to distinguish normal from faulty behavior. Defining the anomaly thresholds turns out to be a tricky task that requires advanced statistics like machine learning. However, even the best baselining methods prove to be inadequate when it comes to the cloud.

With modern microservice architectures, a single fault impacts a multitude of connected services which subsequently also fail. Therefore, a single problem typically triggers many alerts, which are all justified. This is called an alert storm or noisy alerts.

Conventional monitoring solutions fall short of resolving this issue. It remains up to human operators to make sense of the alerts. Problem triage becomes a time consuming and often frustrating exercise involving war rooms and graveyard shifts.

The only way out is a reliable method for determining the underlying root cause automatically.

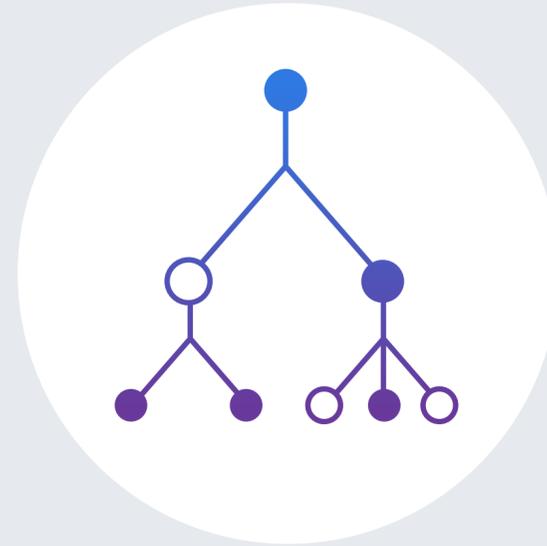


"Incidence response services: xMatters, PagerDuty, VictorOps, Opsgenie"

A large health insurer with 350 hosts captures **900,000 events** per minute and **200,000 measures** per minute.

Fine-tuning individual baselines helps, but it does not fix alert storms. For a real cure, we need to step outside the box and try to find the underlying root cause directly.

There are two very distinct AI-based approaches to reduce alert noise:



Deterministic AI performs a step-by-step fault tree analysis as is common in safety engineering.

Results: Precise identification of the problem root cause

- Works in near real-time
- Explainable results—problem evolution over time can be visualized step-by-step
- Includes technical and foundational root causes as well as impact analysis



Machine learning AI is a statistical approach that correlates metrics, events, and alerts to build a multi-dimensional model of the analyzed system.

Results: A set of correlated alerts; it is still up to human operators to determine the root cause

- Building machine learning models takes time
- Tend to lag behind in dynamic environments
- Some systems suggest likely root causes by accessing historic records created by humans

Chapter 2

Getting the best monitoring data

Insight

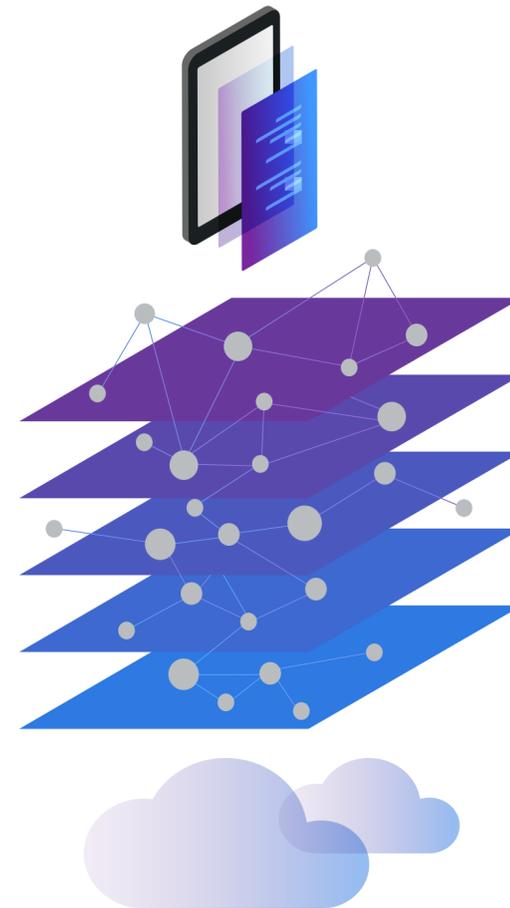
In an environment of disparate monitoring tools, operations personnel are left to make sense of multiple diverse inputs coming from various sources. This increases the likelihood of error in situational awareness and diagnosis.² Currently, only 5% of applications are monitored. The aim is to get full end-to-end visibility.

Challenge

Full system visibility is a necessary precondition for automating operations including solid self-remediation. We need full insight not only into the application—including containers and functions-as-a-service—but also into all layers of the cloud infrastructure, networks, the CI/CD pipeline, and the real user experience. In many cases, data collection itself comes for free, as all major public cloud providers offer monitoring APIs, and open-source tools are abundantly available. However, the following considerations are critical:

- How much manual effort is required for instrumentation and deployment of updates?
- Can the monitoring agents inject themselves into ephemeral components like functions or containers, and do configuration changes require additional manual instrumentation?
- Are the metrics coarsely sampled or high-fidelity?
- Is there enough meta-information and context to build a unifying data model?

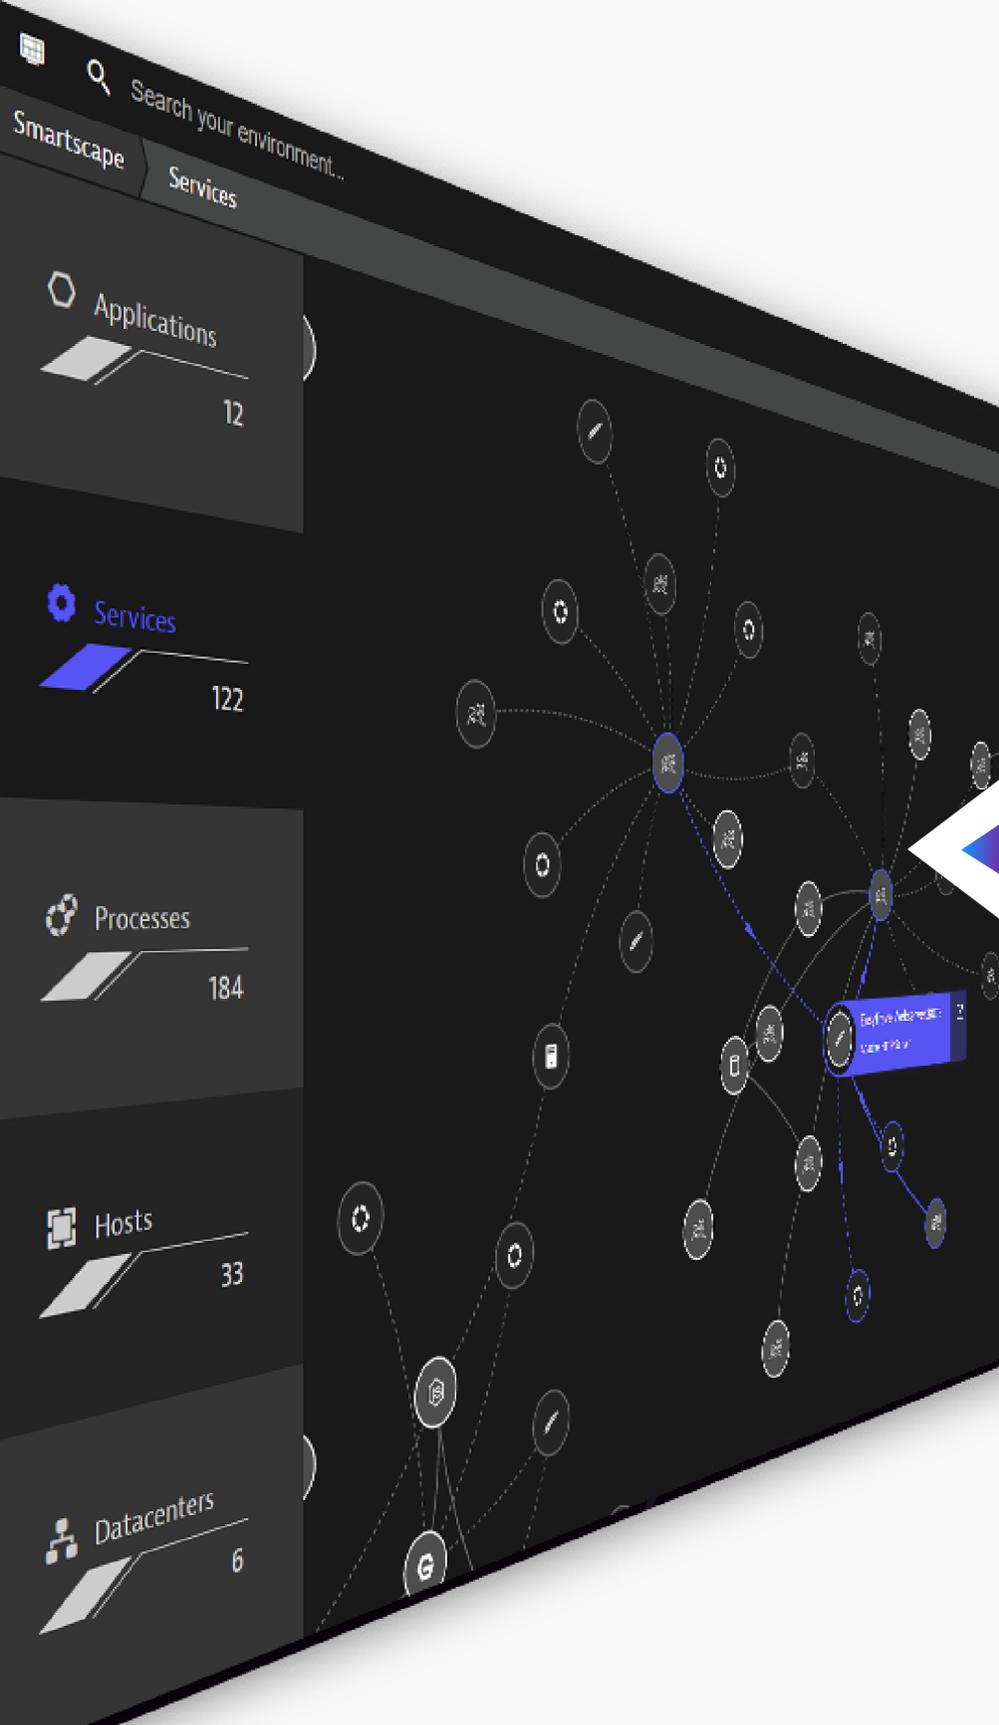
²Use AIOps for a Data-Driven Approach to Improve Insights From IT Operations Monitoring Tools (Gartner Research Note)



A big airline with 2,500 hosts has

432 million

topology updates per day.



Rich data in context

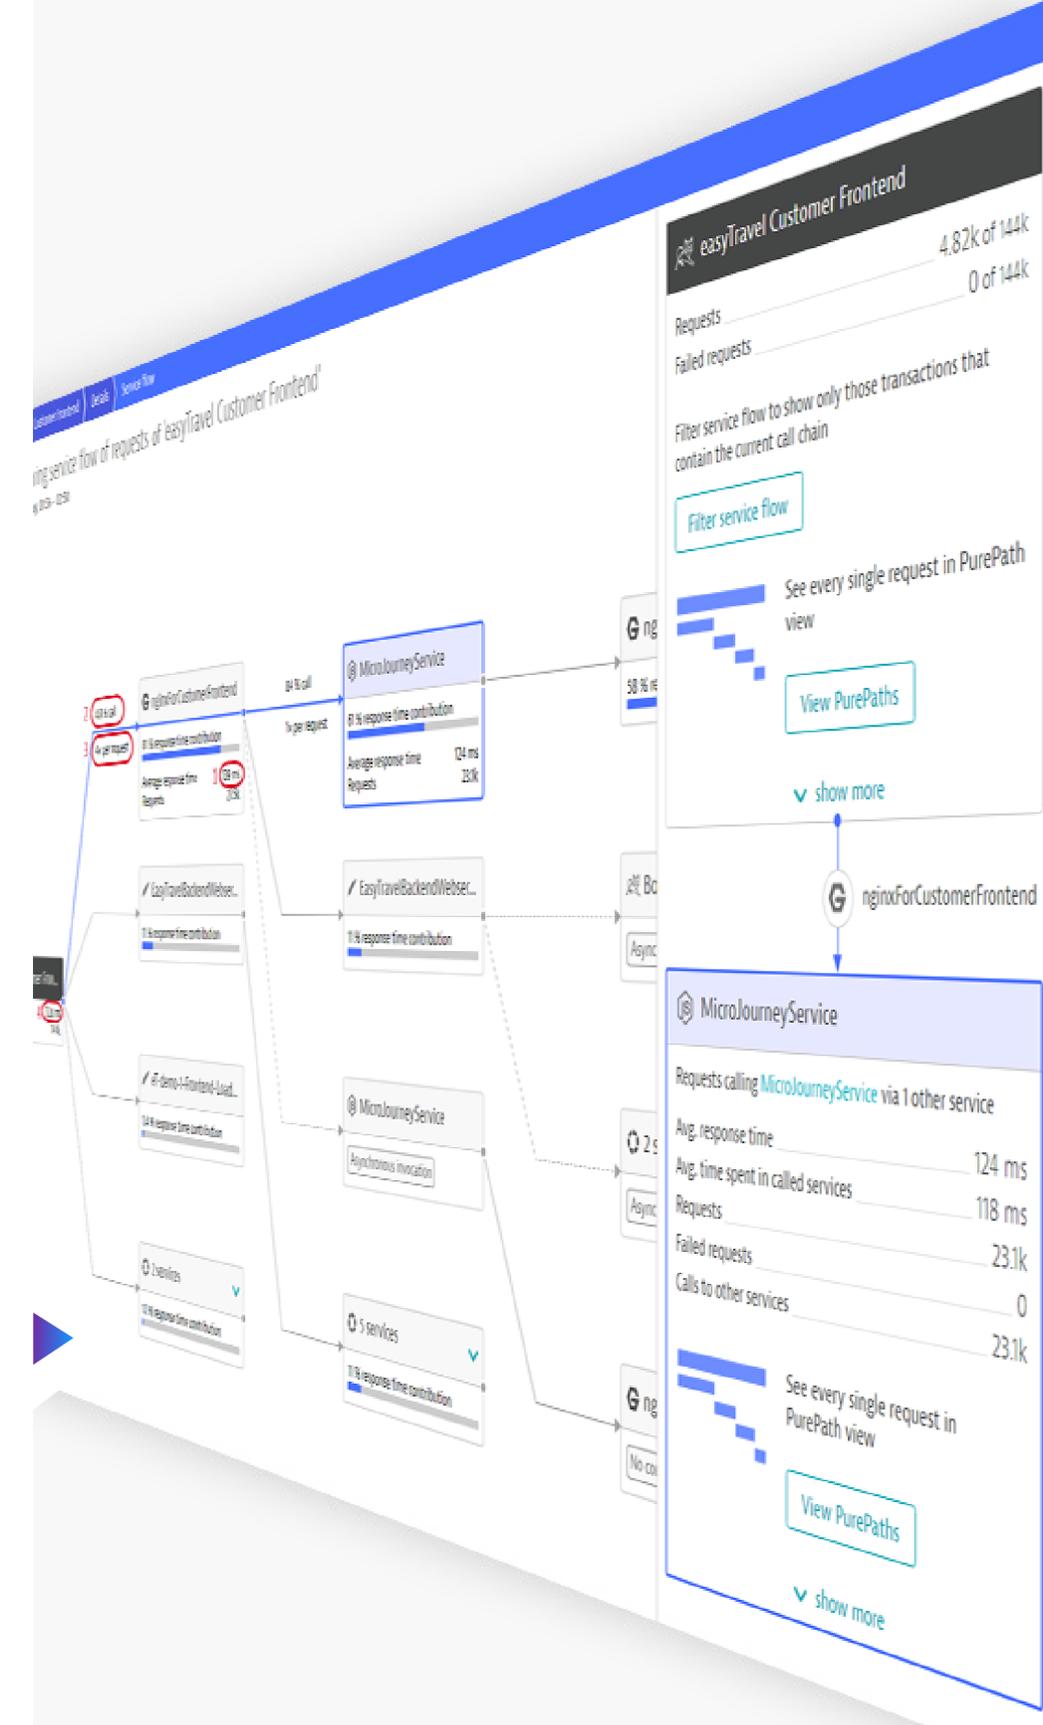
In order to accomplish true root cause analysis, the collected data need to be high-fidelity (minimal or no sampling) and context-rich in order to create real-time topology and service flow maps.

Topology map

A topology map captures and visualizes the entire application environment. This includes the vertical stack (infrastructure, services and processes) and the horizontal dependencies, i.e. all incoming and outgoing call relationships. Leading monitoring solutions provide auto-discovery of new environment components and near real-time updates.

Service flow map

A service flow map offers a transactional view that illustrates the sequence of service calls from the perspective of a single service or request. The difference to topologies is that service flows display a step-by-step sequence of a whole transaction while topologies are higher abstractions and only show general dependences. Service flows require high fidelity data with minimal or no sampling.



Chapter 3

AI Operations and Root Cause Analysis

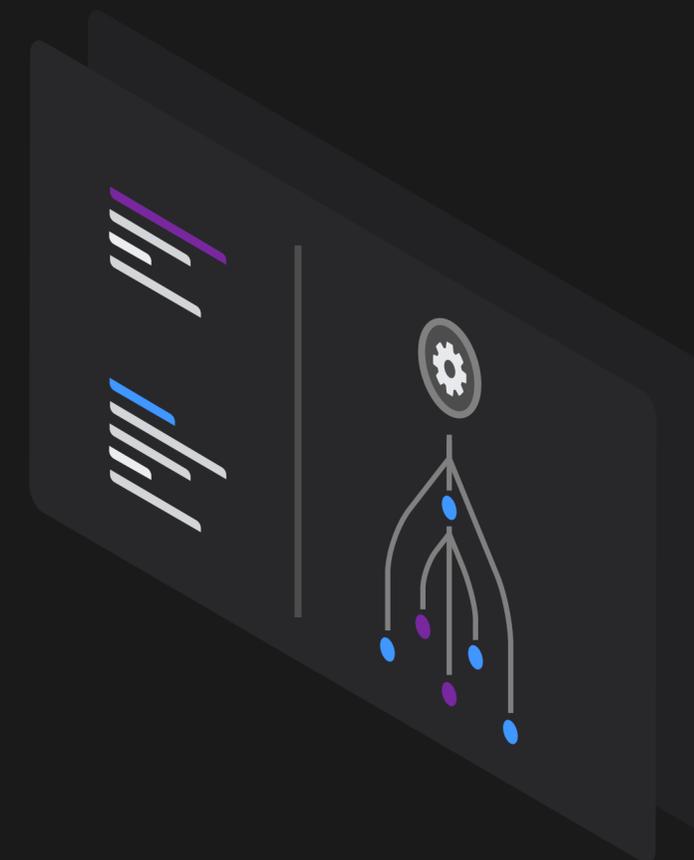
Insight

Enterprises without AI attempt the impossible. Gartner predicts 30% of IT organizations that fail to adopt AI will no longer be operationally viable by 2022.³ As enterprises embrace a hybrid, multi-cloud environment, the sheer volume of data and massive environmental complexity will make it impossible for humans to monitor, comprehend, and take action.

Challenge

We are quickly entering a time when humans will no longer be the main actors to fix IT problems or push code into production. Cloud and AI solutions revolve around automation, so DevOps won't require nearly as much human intervention in the future. For AIOps (truly autonomous cloud operations) to work perfectly, we need a system that can not only identify that something is wrong, but pinpoint the true root cause.

Modern, highly dynamic microservice architectures run in hybrid and multi-cloud environments. Infrastructure and services are spun up and killed within the blink of an eye as loads demand. Determining the root cause of an anomaly requires exponentially more effort than humans can take on.



52 billion

dependencies analyzed per day to find problem root causes for a multinational business systems company with 17,500 hosts.

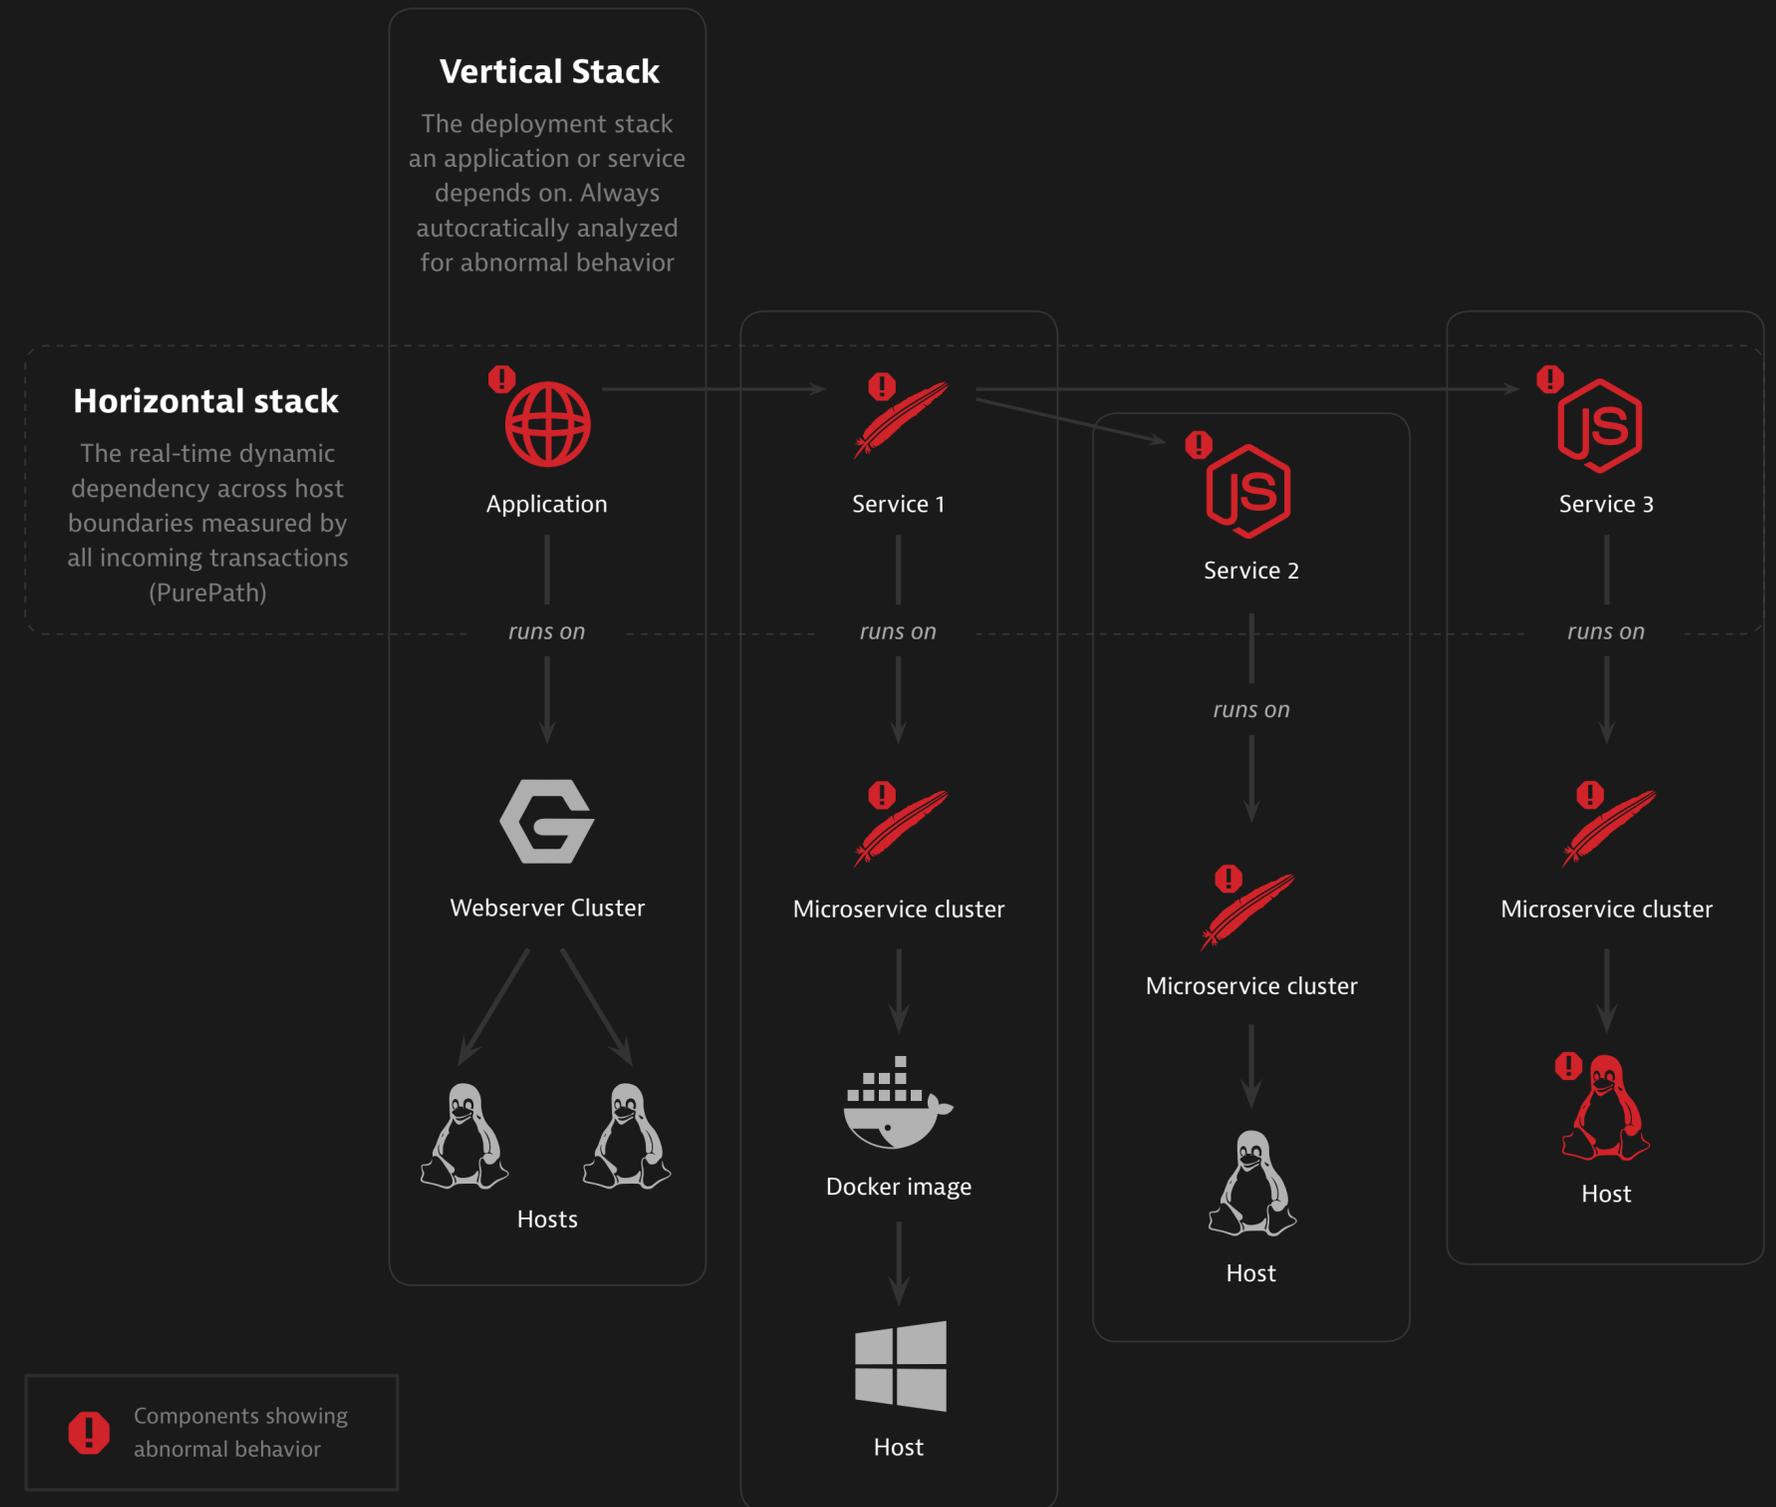
³AI (in a box) for IT Ops—The AIOps 101 you've been looking for.

Root cause analysis with deterministic AI

Davis—the Dynatrace AI engine—uses the application topology and service flow maps together with high-fidelity metrics to perform a fault tree analysis. A fault tree shows all the vertical and horizontal topological dependencies for a given alert. Consider the following example visualized in the chart to the right.

1. A web app exhibits an anomaly, like a reduced response time (see top left in the graphic).
2. Davis first “takes a look” at the vertical stack below and finds that everything performs as expected—no problems there.
3. From here, Davis follows all the transactions and detects a dependency on Service 1 that also shows an anomaly. In addition, all further dependencies (Services 2 and 3) exhibit anomalies as well.
4. The automatic root-cause detection includes all the relevant vertical stacks as shown in the example and ranks the contributors to determine the one with the most negative impact.
5. In this case, the root cause is a CPU saturation in one of the Linux hosts.

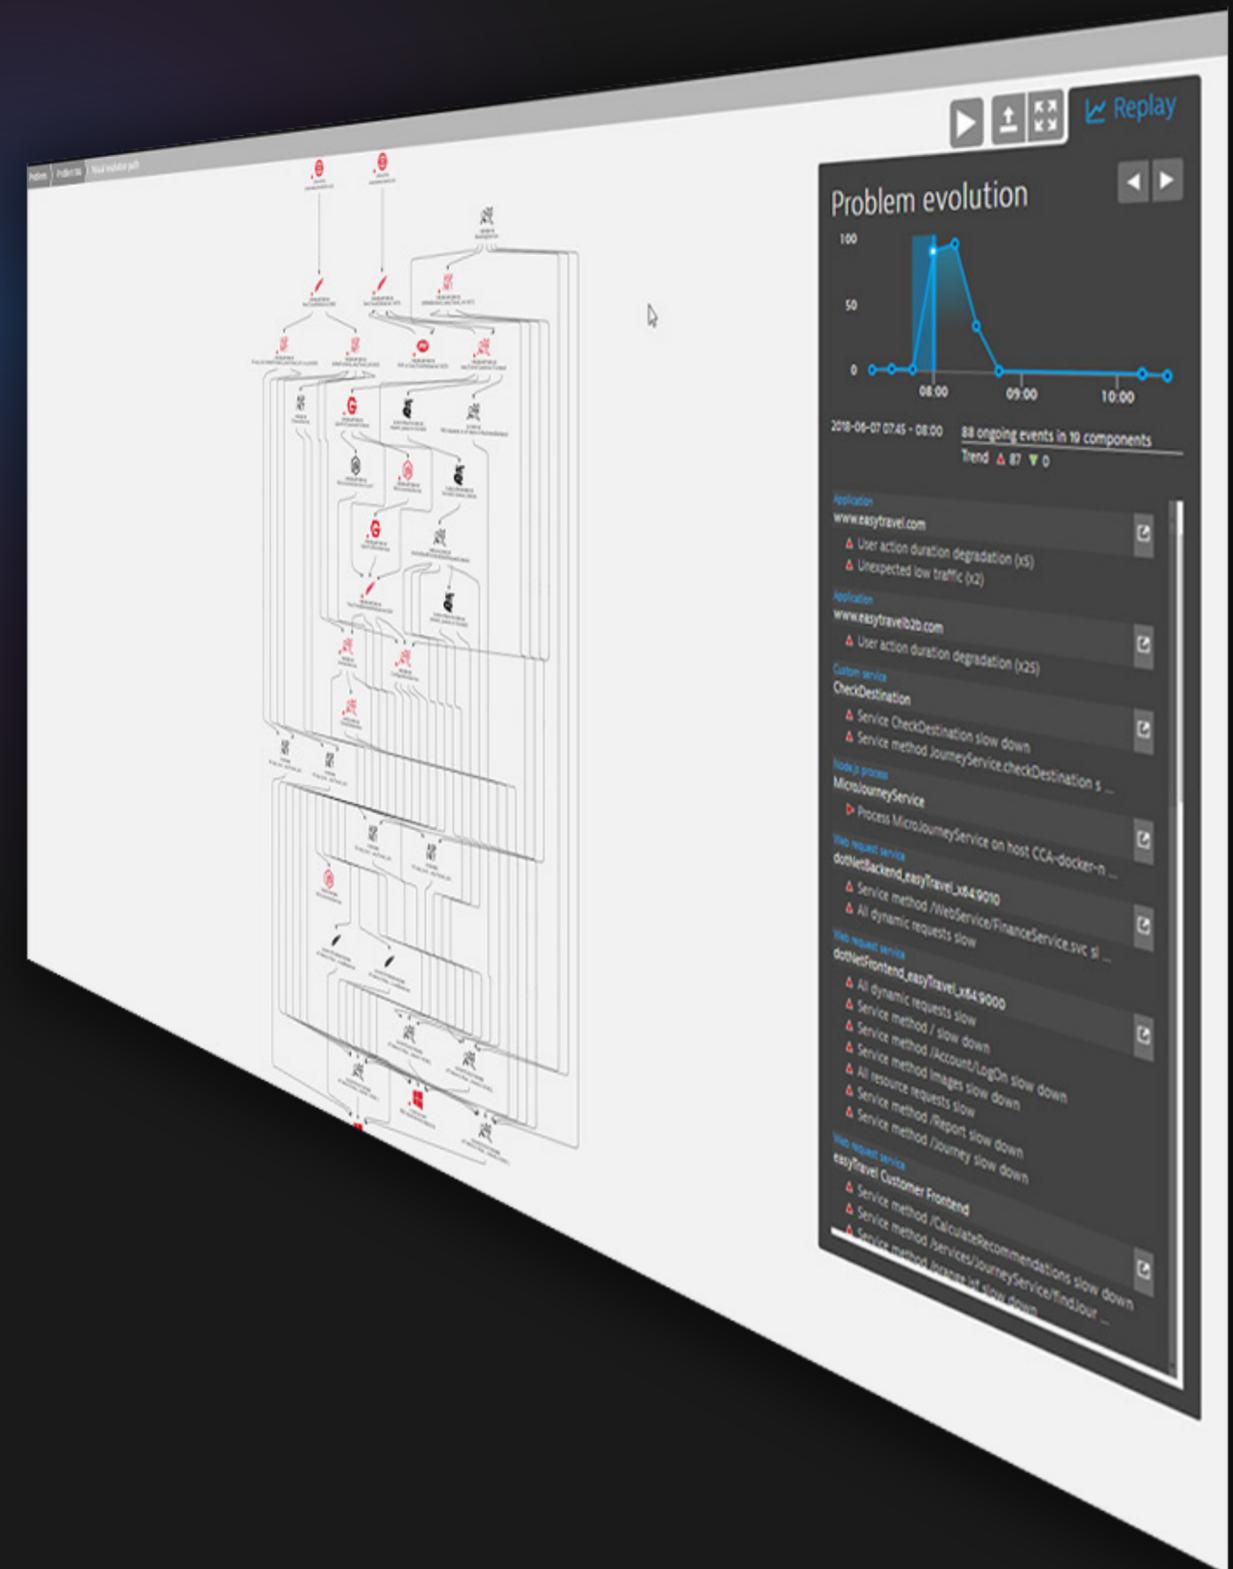
Deterministic AI automatically and accurately determines the technical anomaly root cause. This is a necessary precondition for true AIOps. We’ll go deeper into the requirements auto-remediation in the next sections.



Understanding problem evolution

Deterministic fault tree analysis yields precise, explainable results. This can be used to replay the evolution and resolution of a problem step by step and visualize the affected components in a topology map. This is an extremely powerful feature because it allows the DevOps team to gain a deep understanding of the problem right from the get-go, cutting triage and research time to a minimum.

The problem evolution data is key for auto-remediation. Given that it can be accessed through APIs, remediation sequences can be triggered to resolve a problem with surgical precision and at a speed not achievable by human operators.



Impact analysis and foundational root causes

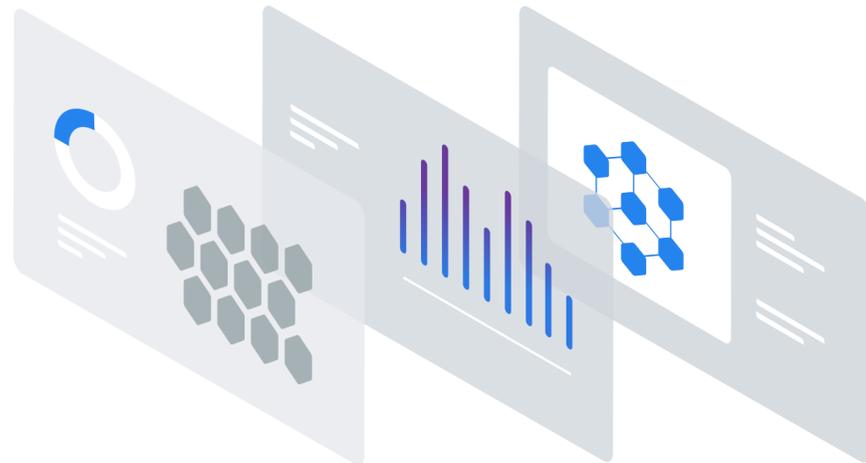
Insight

Infrastructure and services get spun up and killed as needed at a mind-boggling speed in a modern dynamic microservice application. That's the nature of a healthy system.

A disappearing container can be a desired event to optimize resources, or it can be a sign of an unintended disruption that requires immediate mitigation. The AI needs to be able to tell an anomaly from a desired change.

Challenge

A precise and reliable determination of the technical root cause is absolutely essential for auto-remediation, but it is not sufficient. We also need a measure of an anomaly's severity and some indication of what led to the technical root cause in the first place.



Impact severity

Not every disappearing container or host is a problem, and a slow service that nobody uses does not require immediate attention. Therefore, an advanced software intelligence system assesses the severity of a problem:



◀ User impact

How many users have been impacted by a detected problem since it occurred? Ideally, the number should be based on actual real users rather than a statistical extrapolation of historic data.



◀ Service calls impacted

Some parts of the system are not built for human interaction. In this case, the number of impacted service calls is a good estimate of the severity.



◀ Business impact

As software intelligence solutions increasingly cover enterprise systems end-to-end, from user actions all the way to the infrastructure, it is possible to map system performance to business KPIs. A retailer, for example, can measure the dollar value of purchases during a system slowdown and compare it with a reference timeframe in the past.

Foundational root causes

The **technical root** cause determines what is broken.
 The **foundational root** cause specifies why it is broken.

Typical foundational root causes are:

- **Deployments**
 Collecting metrics and events from the CI/CD tool chain makes it possible to link a problem to a specific deployment (and roll it back if needed).
- **Third-party configuration changes**
 These can relate to changes in the underlying cloud infrastructure or a third-party service.
- **Infrastructure availability**
 In many cases, the shutdown or restart of hosts or individual processes causes the problem.

To determine the foundational root causes the AI engine needs to have access to metrics and events from the CI/CD pipeline, ITSM solutions, and other connected tools. Dynatrace provides an API and plug-ins to ingest third-party data into Davis.

 **2 Applications: User action duration degradation**
 Problem 753 detected at Nov 28 06:58–Nov 28 07:54 (was open for 56 minutes). This problem affects real users.

 **Affected applications** 2  **Affected services** 15  **Affected infrastructure** 3

654,998,400
Discrepancies analyzed

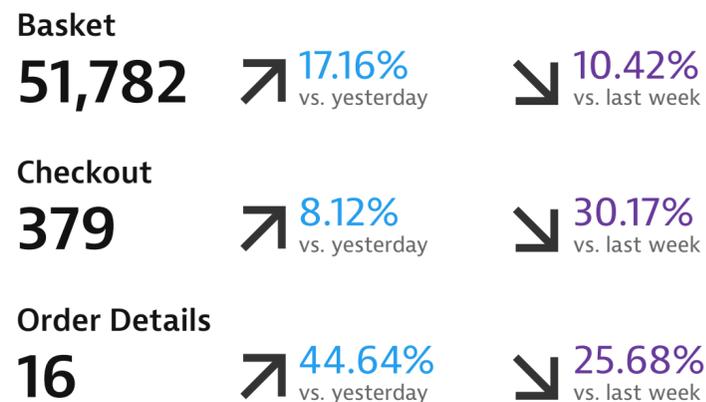
Business Impact Analysis

An analysis of all affected service calls and impacted real users during the first 10 minutes of the problem shows the following potential impact.



Business Metric Analysis

Additional analysis performed on key business metrics such as conversion goals or revenue numbers. Comparisons are done for the Problem timeframe yesterday and a week ago.



Root Cause

Based on our dependency analysis all incidents have the same root cause

 **Check Destination**
Custom service

Response time degradation
 The current response time (19.6 s) exceeds the auto-detected baseline (120 ms) by 16,309%

Affected requests	Service method
551/min	All methods affected

 **BB1-apache-tomcatjms-iis**
Host

CPU saturation
 100% CPU usage

[Analyze logs](#)

Chapter 5

Auto-remediation

Insight

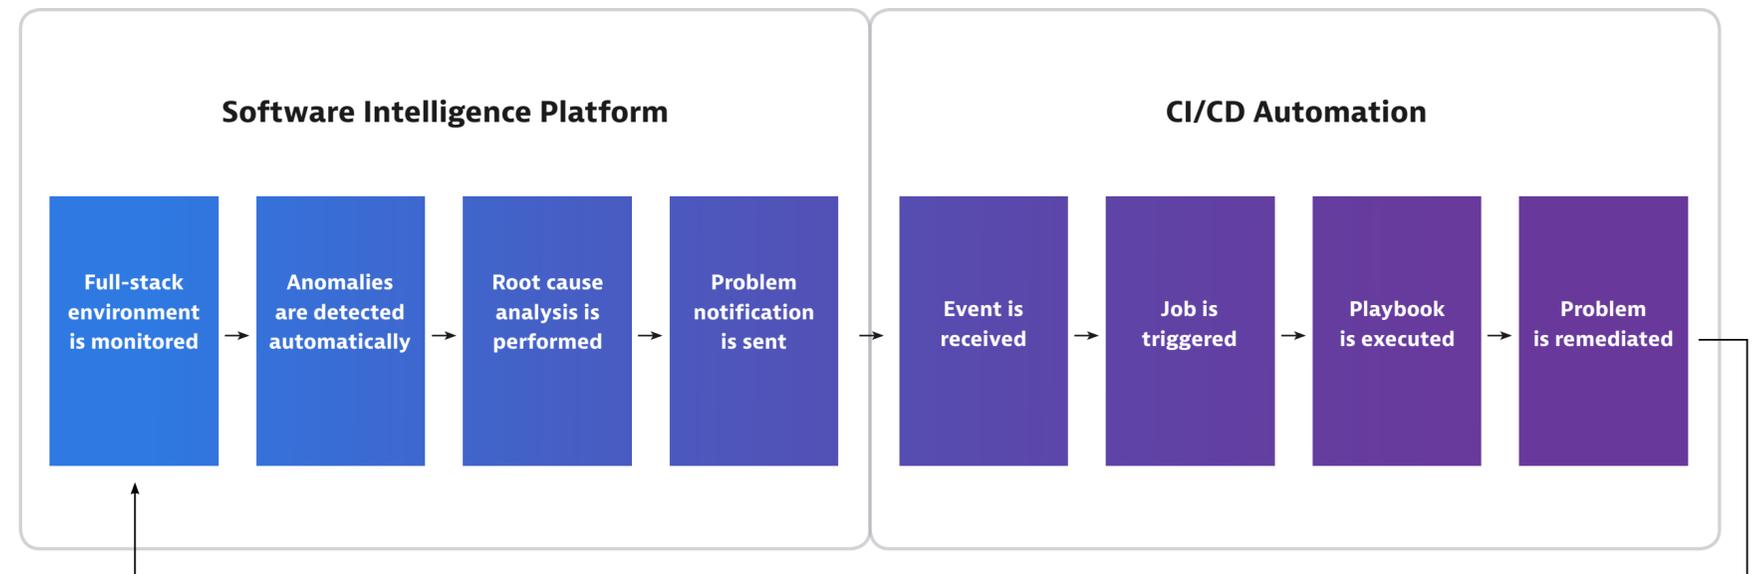
Infrastructure as code and powerful cloud orchestration layers provide the necessary ingredients to automate operations and enable self-remediation. This will not only reduce operational cost but also avoid human error. The key to truly autonomous cloud operations is reliable system health information including deep anomaly root cause and impact analyses.

Challenge

Many cloud platforms offer mechanisms to dynamically adjust resources based on load demand or restart unhealthy hosts and services. Some of these solutions are very advanced—however, they only work within their designed scope. Software intelligence solutions cover the entire enterprise system end-to-end, including hybrid environments where mainframes exist along multiple cloud platforms.

Enabling auto-remediation

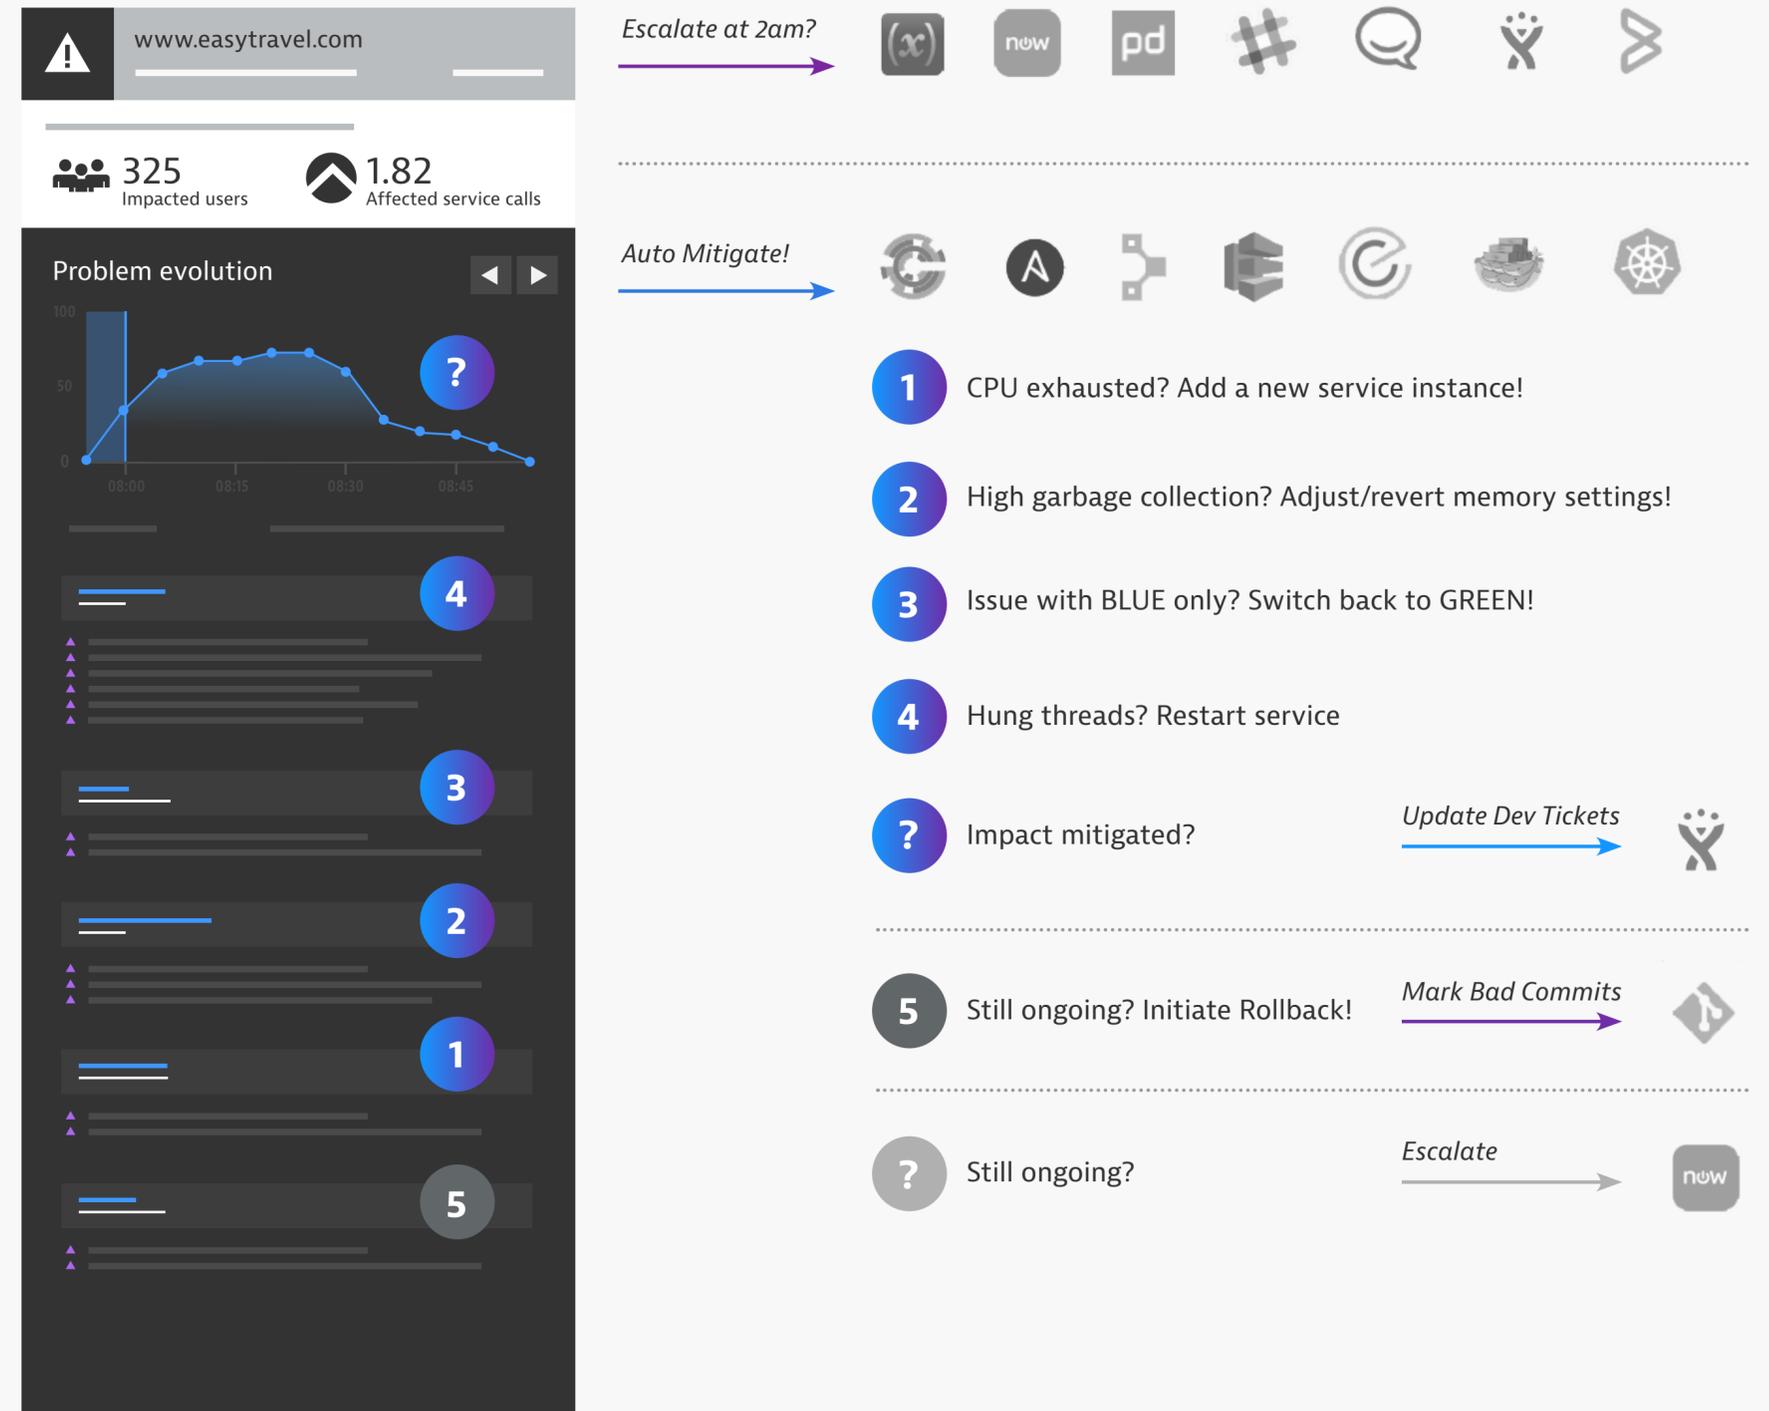
There are many ways of implementing auto-remediation in practice. Typically, the software intelligence platform integrates with CI/CD solutions or with cloud platform configuration layers to execute remediation actions. In any case, the software intelligence solution needs to provide full stack monitoring, automatic anomaly detection, precise root cause analysis and problem notification through APIs.



Complex auto-remediation sequences

This example shows how a precise analysis of the technical root cause, foundational root causes and user/business impact can be used to automate problem resolution through integration with a variety of CI/CD, ITOM, workflow and cloud technologies.

Path to NoOps: Auto-Remediation, Self-Healing...



Automation and system integrations

Insight

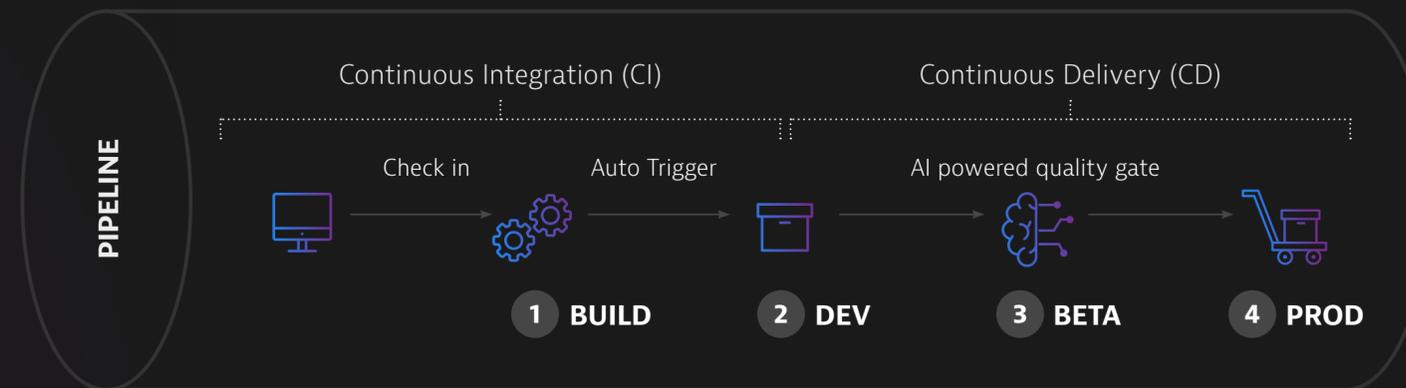
Automation doesn't stop at software operations and auto-remediation in an enterprise grade application environment. Accurate and explainable software intelligence has the capacity to move towards automating the entire digital value chain and to enable novel business processes.

The unbreakable DevOps pipeline

Over the last years, many DevOps teams have come a long way in implementing a CI/CD pipeline that codifies and automates parts of the build, testing, and deployment steps. The goal is to speed up time to market and ensure excellent software quality—to get faster and better. AI-powered software intelligence helps to close existing automation gaps like manual approval steps at decision gates or build validation. It also provides valuable performance signatures to test new builds against production scenarios.⁴

This follows the concept of “shift left”—to use more production data earlier in the development lifecycle to answer the question:

“Is this a good or bad change that we try to push towards production?”



3x
faster build and
test cycles,
50% reduction
in issues.

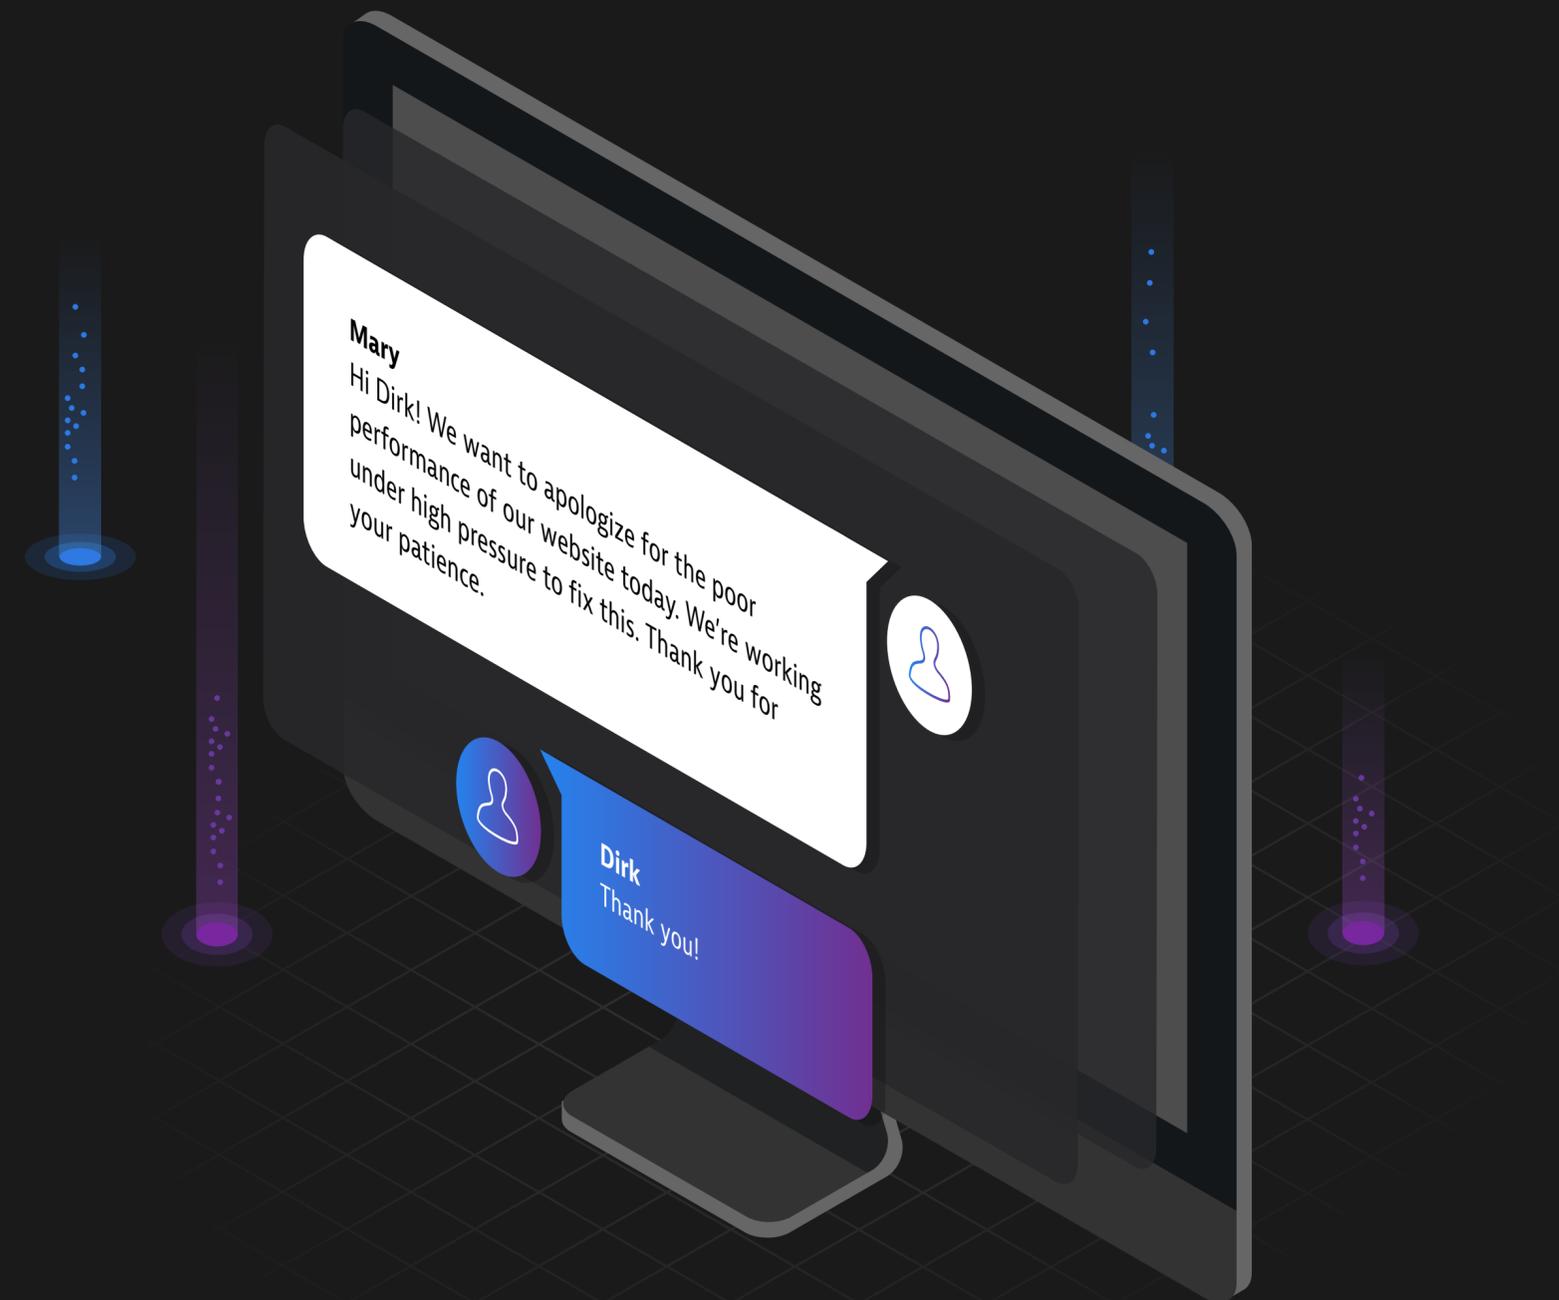
-Verizon Enterprise

⁴<https://www.dynatrace.com/news/blog/shift-left-in-jenkins-how-to-implement-performance-signature-with-dynatrace/>

Automating customer service

Any good software intelligence solution needs to include real user data, and an impact analysis (as described in chapter 4) can be used to ensure customer satisfaction even if something goes wrong.

In case of a breakdown or slowdown, the system can engage autonomously with impacted users. One way is to open a chat window operated by a chatbot behind the scenes and inform the customer about the specific performance issue, then offer to make it up to them by providing discounts, etc.



Chapter 7

Natural language interfaces

Insight

Virtual assistants have emerged as one of the fastest-growing areas of AI and exist in many forms, including voice bots, text bots, and SMS bots.⁵ These technologies have become mainstream with smartphones and home automation systems. When applied to software intelligence platforms, which usually require expert knowledge to navigate, natural language interfaces enable broad user adoption, organizational learning and innovation.

Challenge

Most people are not well-trained performance engineers, but everybody knows how to ask a question. To improve service quality in the eye of the customers, up-to-the-minute information on system health and business KPIs need to be available to customer service reps, PR and marketing specialists, business analysts, software developers and architects, and executive management across the board.

Information and actionable knowledge can be democratized to allow for cross-departmental organizational learning and culture change.

Smart assistants meet software intelligence

A software intelligence solution that holds actionable insights about recent and current problems, their causes, and their impacts on users and the business is well suited for a new and convenient human user interface.

With Dynatrace Davis Assistant users can simply ask a question verbally or use a text-based chat tool to interact with the software intelligence platform, and they'll get a plain-language answer back within the second. Alternatively, for a deep dive, the bot can open a web browser and display the relevant charts. The ultimate level of system intelligence would be smart assistants that identify an error pattern, proactively suggest remediation actions, and only ask for approval to execute them.

⁵AI Technologies—William Blair Industry Report, June 28, 2018



"Give me my morning report."

Good morning. After analyzing 13 applications, here are some interesting findings from the last 24 hours.

- Easytravel.com reported an apdex customer experience rating of 0.85, a 49 percent increase from what was expected; and 7,200 sessions per minute, a 24 percent increase from what was expected.
- Vmware.dynatrace.com had an apdex customer experience rating of 0.97 which is near the expected value of 0.9.



"Are there any problems right now?"

There are currently 11 open problems. The most affected application is easytravel dynatrace-dev, which is being affected by two issues.

Would you like to see a listing of these issues?



(Bot opens a web browser with the pertinent information.)



"What was the customer experience over the weekend?"

Over the weekend, Vmware.easytravel.com was the most heavily used application. The customer experience was excellent, with an average Apdex rating of 0.97. Is there anything else I can help you with?



"What was the total dollar value of bookings?"

\$3.78 million. That's up from the previous weekend by 2 percent.



Software intelligence holds the key to true AIOps, auto-remediation and automation along the entire digital value chain.

We hope this AIOps Done Right ebook has inspired you to take the next step in your digital journey.

Dynatrace is committed to providing enterprises the data and intelligence they need to be successful with their enterprise cloud and digital transformation initiatives, no matter how complex.

Learn more

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



About Dynatrace

Dynatrace provides software intelligence to simplify enterprise cloud complexity and accelerate digital transformation. With AI and complete automation, our all-in-one platform provides answers, not just data, about the performance of applications, the underlying infrastructure and the experience of all users. That's why many of the world's largest enterprises trust Dynatrace to modernize and automate enterprise cloud operations, release better software faster, and deliver unrivaled digital experiences.