dynatrace | Red Hat

Containers in multicloud environments: the challenge of scale

Introduction

The future of enterprise computing is being constructed in a proliferation of dynamic, containerized services operating across multicloud environments. For enterprises, this architectural shift has become essential. Microservices are more scalable, more customizable and quicker to update than traditionally constructed monolithic applications. Deploying them at scale is the only sustainable pathway to a competitive advantage in the 21st century, which for most organizations will involve accelerated DevOps, AI-driven cloud-based analytics, and best-in-class digital experience for end users.

Red Hat® OpenShift® has emerged as the market-leading platform for Kubernetes-based containerized applications running across multi-cloud environments. Red Hat® OpenShift® takes the orchestration functionality of Kubernetes and adds what developers and operations need to get containers running and scaling efficiently with minimum risk, including interfaces for networking, storage, and load balancing. Red Hat® OpenShift® allows organizations to:

build and run cloud-native microservices fast, at scale, for new apps and updated apps

scale and manage their infrastructure to support these applications

draw upon a wide ecosystem of partners using open-source cloud-native technologies

Using containers to deploy microservices is the kind of architectural shift that challenges thinking on many fronts. For example, in the environments that Kubernetes and Red Hat® OpenShift® enable, every service has a real-time customer (another service, a customer, a developer or operations). Accordingly, understanding the status of every service becomes important.

The challenge here is scale. We're no longer monitoring the performance of large monolithic applications running on readily identifiable nodes and clusters.

Instead, we're looking at the performance of hundreds or thousands of services that consume compute resources in an abstracted fashion. Human beings and the traditional monitoring systems we use have limits. Simply watching data on glass is no longer enough.

Microservices make available continual real-time feeds of data, which makes monitoring microservices a big-data problem.

New ecosystems, new challenges: from APM to observability

The long-established market for Application Performance & Monitoring (APM) is based on software that measures the health of applications and infrastructure by looking at metrics such as resource consumption, response times and error rates. Over time, it becomes possible to compare new readings with historical parameters for typical or acceptable performance. Typically, these systems are configured to generate alerts when performance deteriorates beyond a pre-set level. The next stage, finding a solution, is the job of developers or operations staff.

What's wrong with this traditional approach in the context of microservices? As it turns out, quite a lot.

First, there's the upfront challenge of knowing what to monitor, or what to make observable. In a traditional environment, it's simply a question of attaching agents (code-based instructions pass data back to the APM system) to a handful of obvious points in the ecosystem (servers, VMs, storage, applications and so on). In an environment with hundreds (or thousands) of microservices making calls on each other, this process — known as instrumentation — becomes far more complex and labor intensive. In each case, an IT professional needs to understand what

is running inside each container, and modify images to include an agent that will monitor performance and adjust configuration properties. In a microservices environment, the result is a vast multiplication of labour, with a good deal of additional scope for human error.

Next, there's the problem of collecting and monitoring the sheer volume, variety and velocity of data generated by a system whose main characteristic is dispersed complexity. Traditional APM relies upon sampling metrics. Microservices make available continual real-time feeds of data, which makes monitoring microservices a big-data problem. This requires something new: a scalable, specialized back end that can handle large quantities of time-series data.

Finally, there's the ultimate goal of all this activity: understanding what is causing a slowdown or failure. Traditional monitoring works by offering IT professionals

opportunities to correlate performance anomalies with the way in which services and applications consume resources. In a microservices environment, expecting human beings to comprehend dependencies, quantify business impact, and get to the root cause is a much taller order. If we think in terms of the traditional analogy of finding a needle in a haystack, it's clear that the haystack is far larger, not least because of the enlarged datasets thrown off by a microservices architecture. The needle itself is much smaller, and there may be a number of them involved in the performance degradation. At this point, the limitations of humans trying to make sense of data on glass become all too apparent — especially in cross-functional teams, where different specialists can very often arrive at different diagnoses.

Microservices architectures bring with them an urgent need for a single source of truth, an automated solution that provides root-cause analysis on which everyone can agree.

Microservices architectures bring with them an urgent need for a single source of truth, an automated solution that provides root-cause analysis on which everyone can agree. With a single source of precise answers about the state of microservices, the environment, and their dependencies, teams can spend less time finger-pointing and more time innovating.

As enterprises switch from traditional architecture to microservices and containers, all of these factors combine to create a new requirement: observability. Organizations need to possess all of the performance data in context that makes a difference to the key metrics of the business. The relevant performance data could include a sequence of calls between microservices, a value written to a database, or a user's mouse click as they access goods, services, and information on the corporate website. The key business metrics could include market share, cashflow, stock levels, or customer satisfaction. In the era of microservices, observability means understanding how the levels of performance made visible by data impact on those business metrics. The goal is to address sources of disruption as quickly as possible, before damage becomes evident to the wider world.

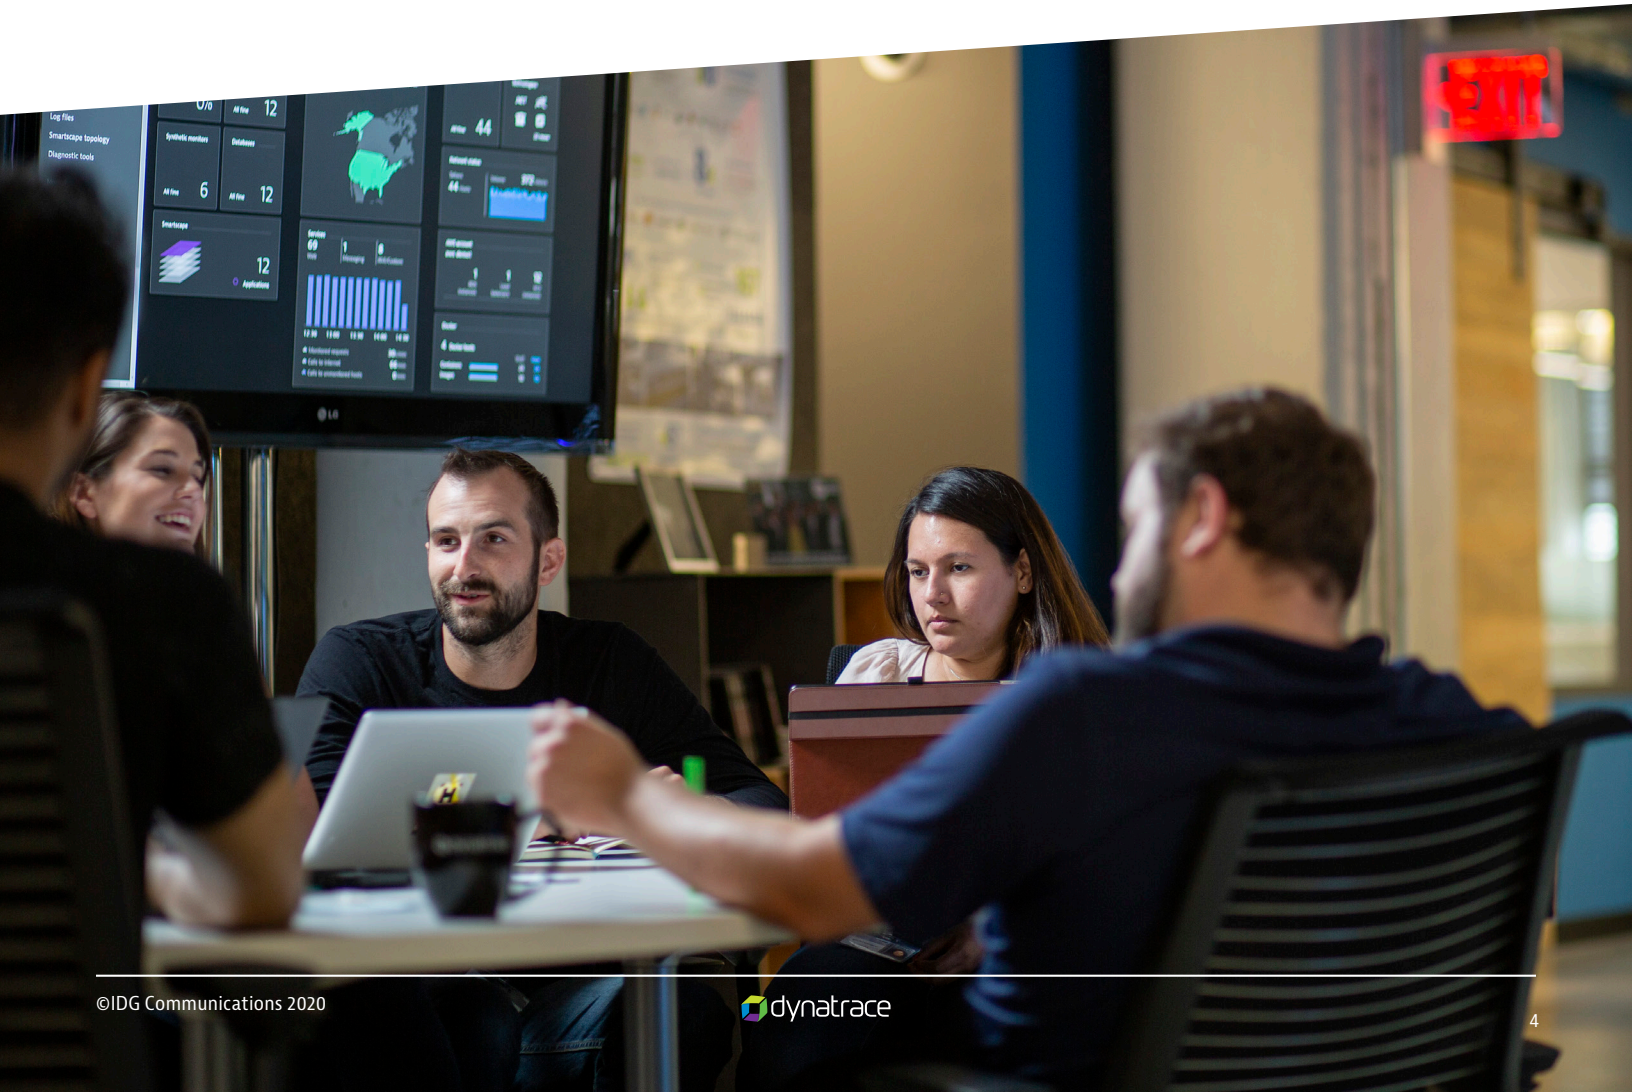

The solution: Dynatrace on Red Hat® OpenShift®

Dynatrace is the only container-aware observability platform with built-in monitoring support for Kubernetes and Red Hat® OpenShift®. It addresses the challenge of observability that accompanies the transformation to a distributed model of microservices, containers and cloud-native applications.

Automated discovery and rapid deployment

Dynatrace automatically discovers and captures data, detecting Docker and CRI-O containers and the microservices running within them. Its continuous automation removes the need to manually deploy agents, tweak code, or change images.

End-to-end granularity

Dynatrace's data capture is comprehensive, ingesting metrics, traces and logs from hosts, networks, services, applications and platforms including Red Hat® OpenShift®. At deployment and in real-time, Dynatrace maps the ecosystem from end-to-end, making sense of calls and dependencies between services. The result is Service Flow, a sophisticated visual UI that shows how a single service triggers other services and how each component contributes to response time.

Built on deterministic AI, delivering root-cause analysis

Dynatrace incorporates deterministic AI for real-time diagnosis of anomalies and deteriorating performance, which puts an end to the alert storms that often accompany a monitoring crisis. Instead of wasting valuable time trying to comprehend the storm, teams can quickly remediate the precise root cause.

Go fast, don't break things

Dynatrace was built with both developers and operations in mind. It integrates with a wide range of CI/CD tools, allowing DevOps teams to start performance checks earlier and speed up feedback loops throughout the development lifecycle.

Cloud-native intelligence at enterprise scale

Like the cloud-native applications it monitors, Dynatrace is built to scale, thanks to its grid-based architecture. Out-of-the-box, it will auto-discover over 100,000 hosts. Fine-grained role-based access controls make Dynatrace securely accessible to large and dispersed teams throughout the organization.

To effectively monitor the vital signs of 21st century multicloud architectures, organizations must harness continuous automation and AI-assistance to accelerate human cognition and productivity.

Automatic and intelligent observability for containers, Kubernetes, and Red Hat® OpenShift®

Modern applications are modular and scalable, built for the cloud and constructed on the proven foundations of containers, Kubernetes and Red Hat® OpenShift®.

The upside is measured in improved agility and faster time to market. The downside is increased complexity and more potential points of failure.

For organizations, the imperative to manage these complex ecosystems is ever increasing. However, we can only manage what we can measure. And we can only measure what we can observe accurately.

To effectively monitor the vital signs of 21st century multicloud architectures, organizations must harness continuous automation and AI-assistance to accelerate human cognition and productivity. In these dynamic ecosystems, Dynatrace tames complexity, eliminates manual routines, and delivers the fast, accurate root-cause analysis on which competitive advantage depends.

CIO FROM IDG

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.

dynatrace.com blog @dynatrace

About Red Hat

Red Hat® OpenShift® is an enterprise-ready Kubernetes container platform with full-stack automated operations to manage hybrid cloud, multicloud, and edge deployments. Red Hat OpenShift is optimized to improve developer productivity and promote innovation.

redhat.com blog @redhat

dynatrace | Red Hat

05.11.21 12751_WP_jw