

# Managing Containers & Multicloud Complexity at Scale

**turbonomic**

The background features a dark green gradient. On the right side, there is a large, solid green circle. Below it, a series of concentric, semi-transparent grey circles are arranged in a grid-like pattern, creating a sense of depth and movement. In the bottom left corner, there are several overlapping, wavy lines in a lighter shade of green, resembling a stylized landscape or data flow.

## EXECUTIVE SUMMARY

# Managing Containers & Multicloud Complexity at Scale

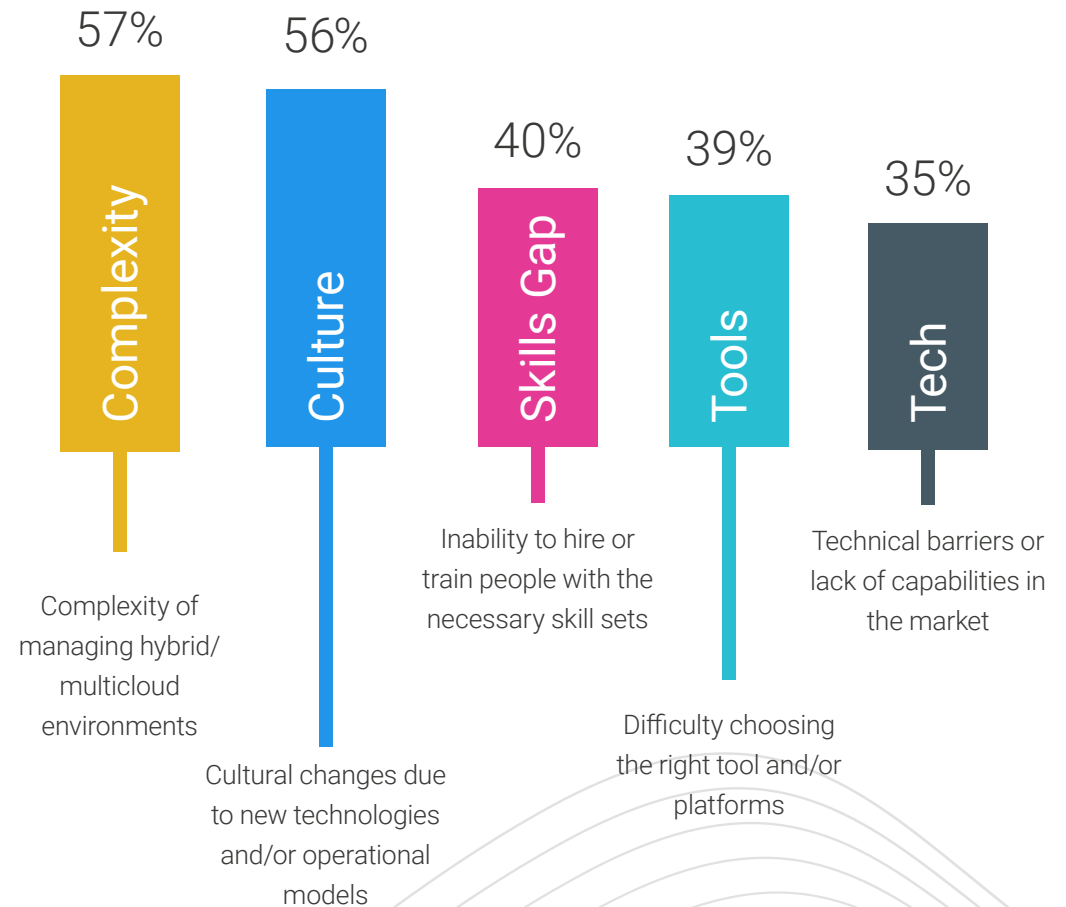
In 2019, the difference between industry leaders and their competitors rests on their applications: how quickly can they innovate, bringing new features and offerings to market, and how well do those applications perform for their customers. To bring ideas to market faster, organizations are adopting container technologies and driving multicloud strategies. These same trends, however, create complexity in how applications get the compute, storage, and network resources they need to run.

**TABLE OF CONTENTS:**

Why Complexity is Your Greatest Challenge.....	3
Containers & Multicloud Complexity Puts Performance at Risk.....	9
Application Resource Management for Kubernetes.....	12

# Why is Complexity Your Greatest Challenge?

In the 2019 State of Multicloud Report, complexity (57%) and culture (56%) were cited as the top challenges to organizations achieving their goals. There are a few reasons why the complexity of managing hybrid and multicloud environments is so challenging. It all comes down to what container technology is enabling. And, in the face of this new complexity, people and teams must adopt new operational models to keep up.

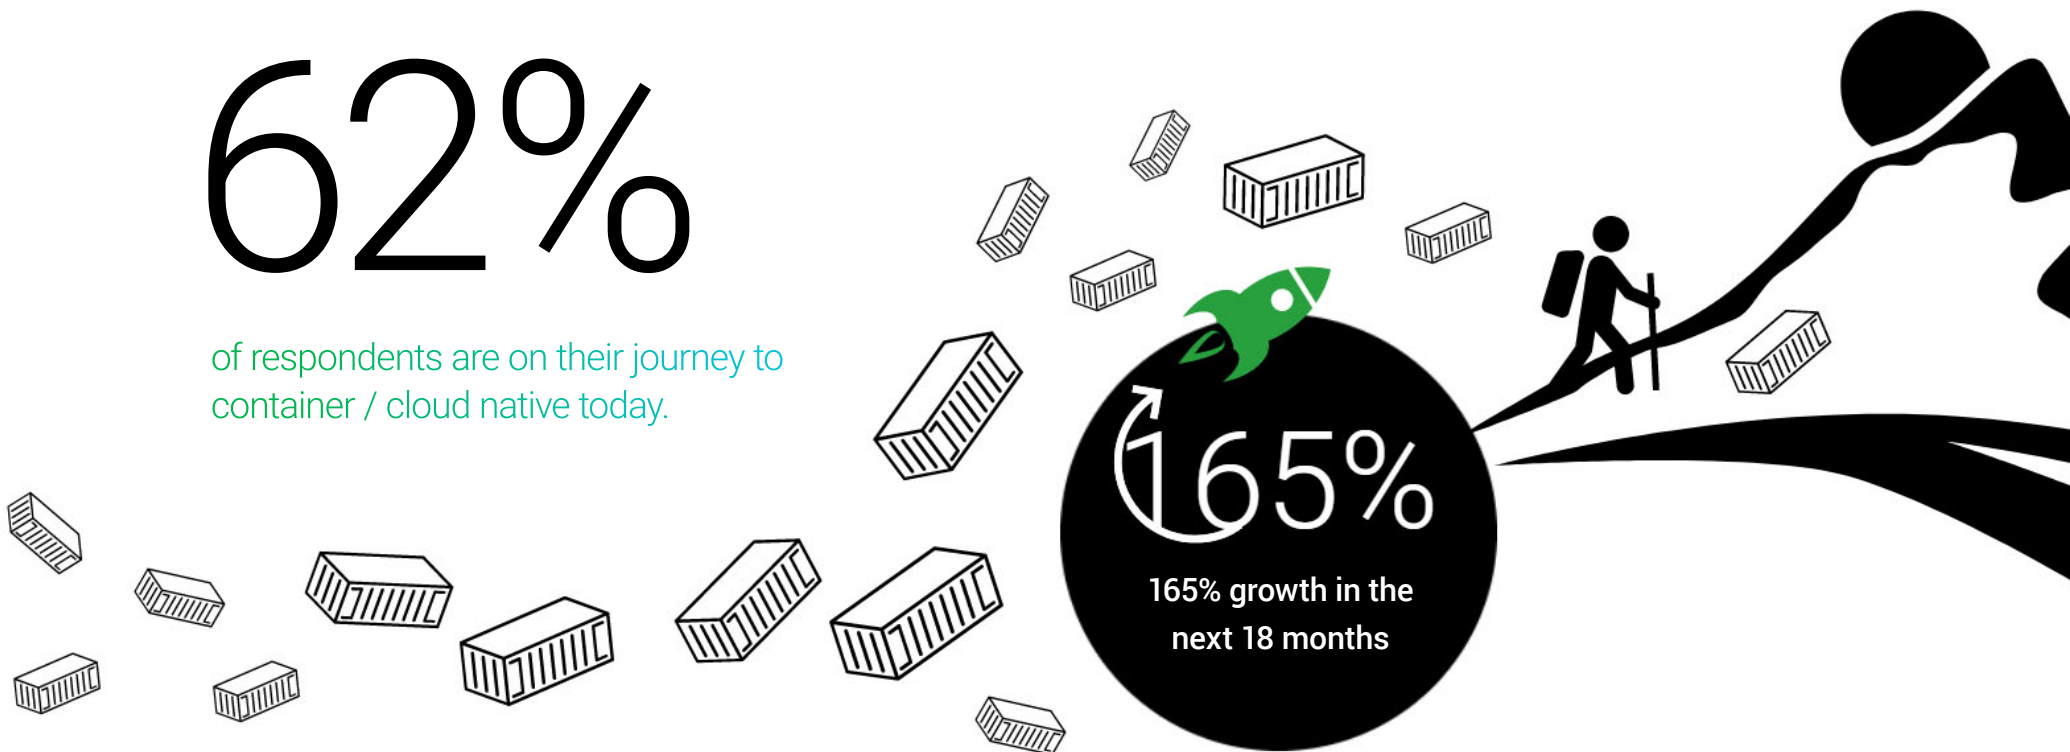
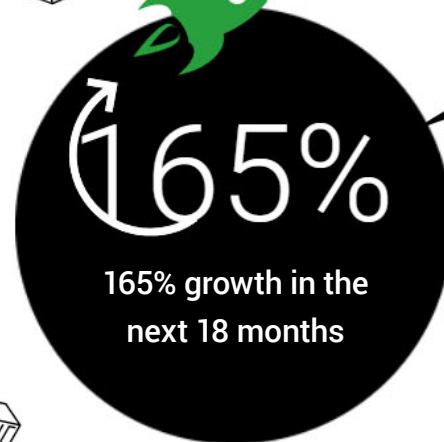


Source: Turbonomic 2019 State of Multicloud

The 2019 State of the Multicloud report found that 62% of respondents are on their containers and cloud native journey. In the next 18 months container adoption is expected to grow 165%.

# 62%

of respondents are on their journey to container / cloud native today.




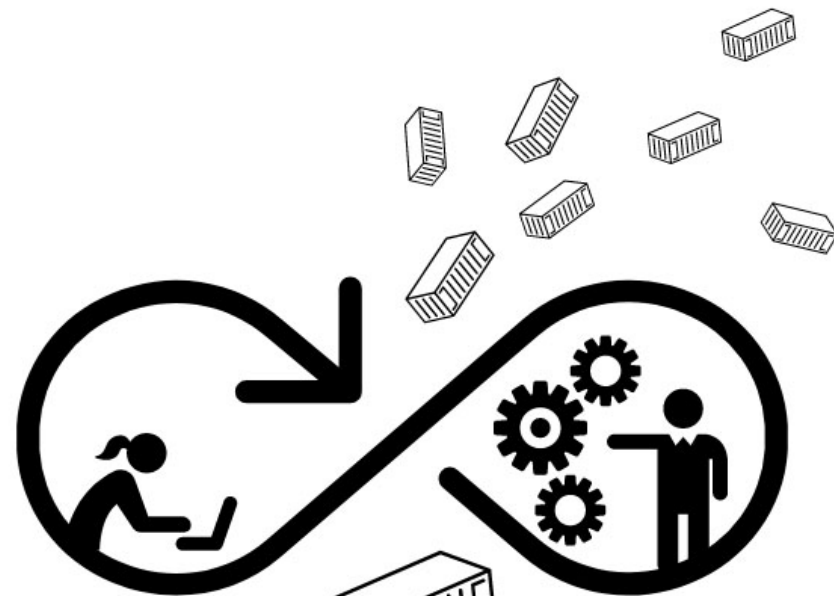
As a result...

# 1 There are more changes to the environment.

Containerization has enabled developers to rapidly and continuously deploy new features and updates to their applications. In the 2018 DORA State of DevOps report, respondents cited significant improvements as a result of implementing CI/CD:




Deployment frequency	Weekly-monthly	→	Hourly-daily
Change lead time	One – six months	→	One – seven days
Change failure rate	46-60%	→	0-15%

DevOps teams' continuous implementation/delivery of smaller and more frequent development releases leads to more changes to environments than ever before. With speed comes the need to have a way to manage constant change in production. 1 of respondents are on their journey to container / cloud native today



## 2 There are more pieces to manage. (Hello, microservices.)

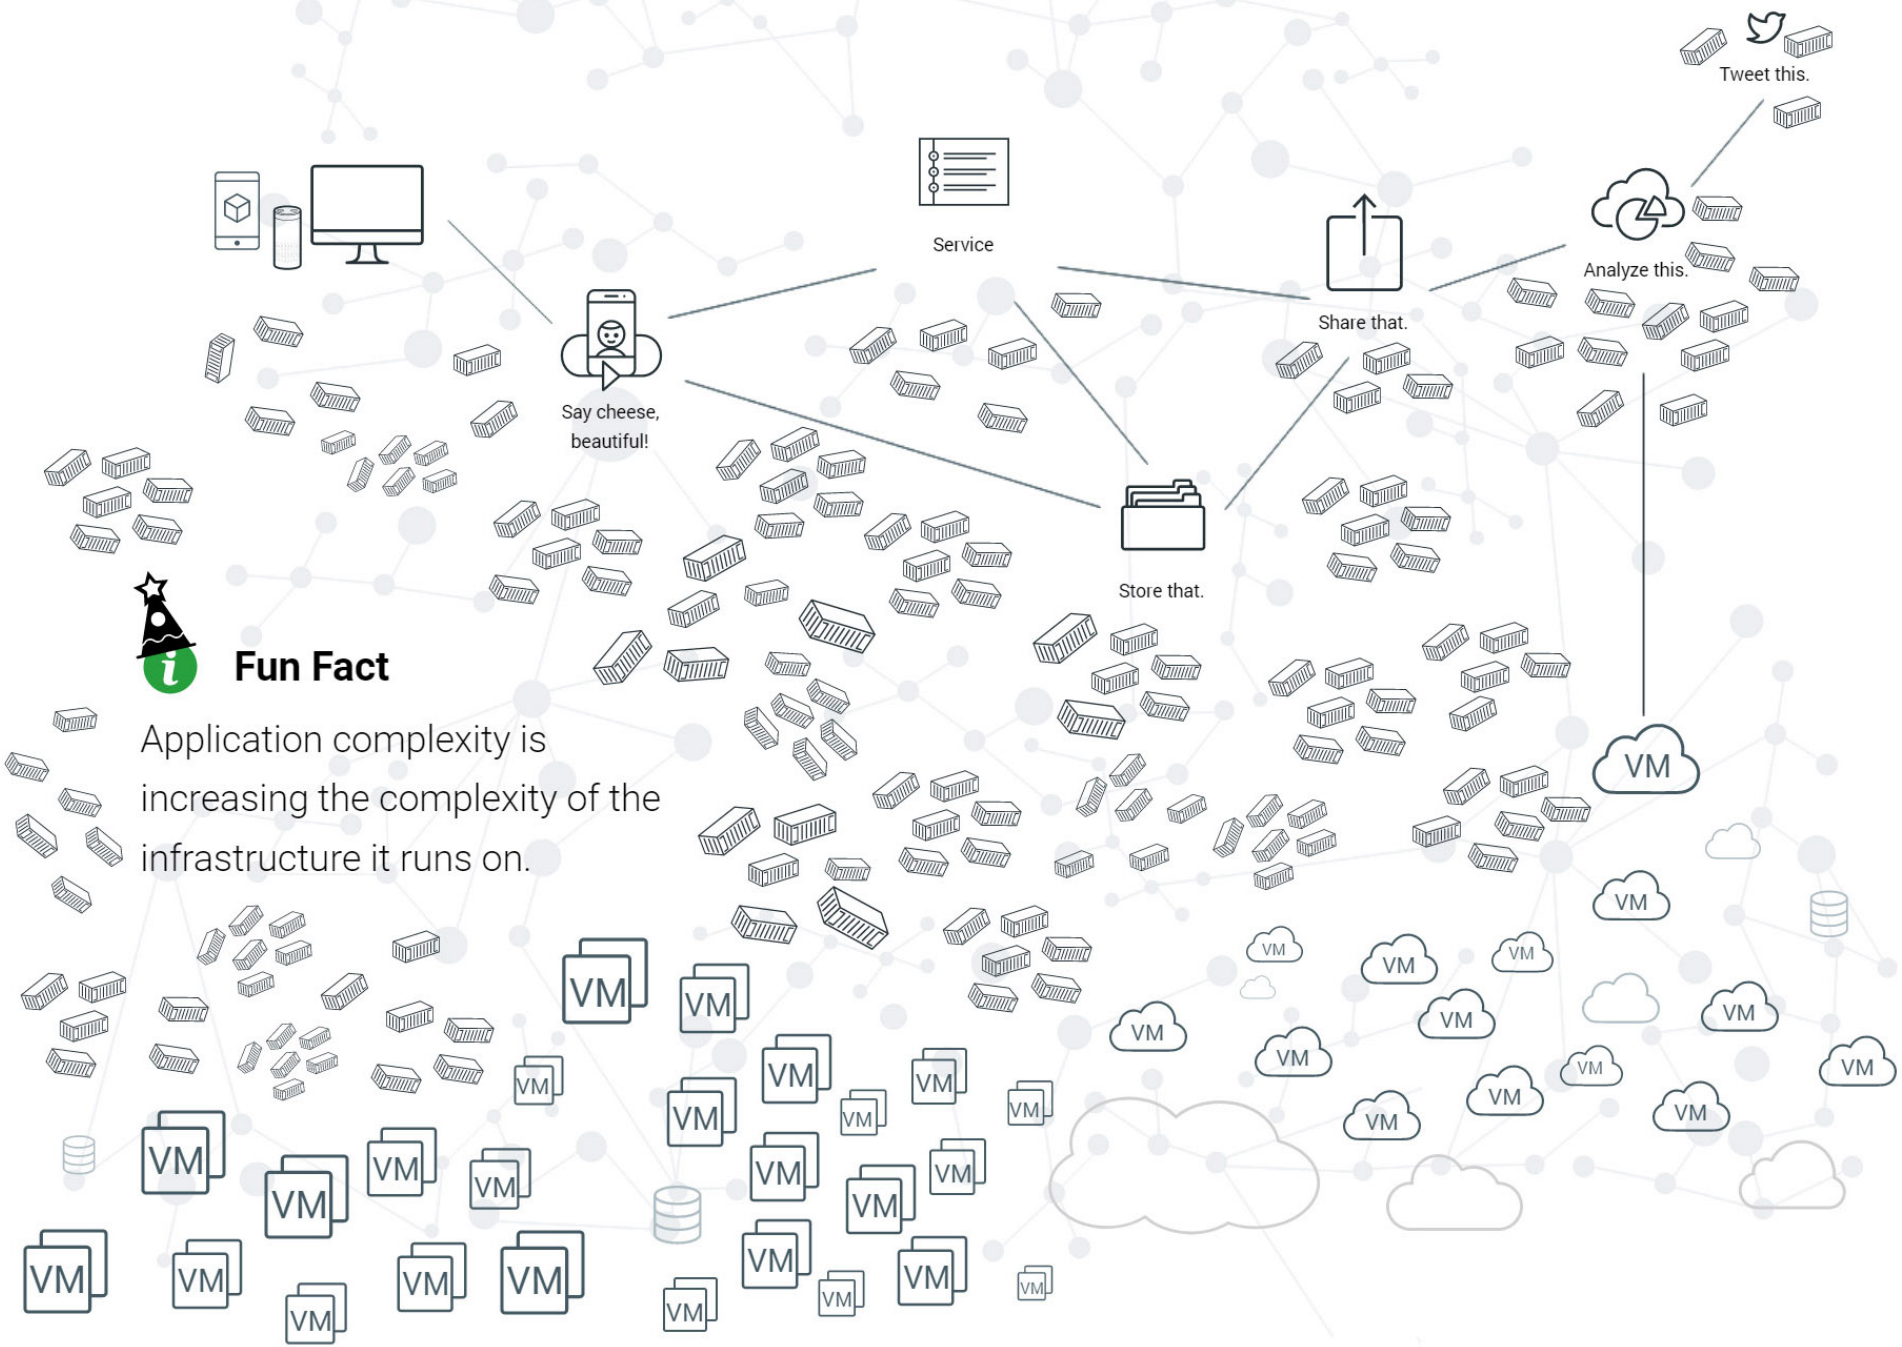
With containers, we can now break down an application into microservices. Loosely coupled services have the benefit of reducing interdependencies within applications, but for operations teams it means there are also more components to manage.



Morgan Stanley forecasts a 23% 5-year CAGR in total workloads and container instances through 2023.




Source: Morgan Stanley Research "A New Software Stack for the Digital Era," May 2019



### Fun Fact

Application complexity is increasing the complexity of the infrastructure it runs on.




### 3 More Types of Infrastructure to Manage, Multicloud & Beyond

The portability of containers allows organizations to build applications once and run anywhere--across multiple clouds, data centers, edge, and IoT. Now we are even seeing single applications distributed across a heterogeneous mix of infrastructure.

More than 50% of large enterprises will deploy at least six edge computing use cases deployed for IoT or immersive experiences by year-end 2023, versus less than 1% in 2019.

1%  
Today

>50%  
2023



Source: Source: Gartner "Exploring the Edge: 12 Frontiers of Edge Computing" May 2019 (ID G00388219)



# Container & Multicloud Complexity Puts Performance at Risk

Increased scale, along with the unpredictable and ever-changing demand of these workloads often leads to resource congestion, directly affecting your customer experience.




# 46%



of mobile web users are unlikely to return to a website that they had trouble accessing in the past.




Source: Morgan Stanley Research "A New Software Stack for the Digital Era," May 2019



Demand is ever-changing.



Causing congestion at any point.



# How are you solving for resource congestion today?



## Break




You know you must adapt how your organization manages application resources to keep pace. But what the industry has offered you is oddly similar to what we've been doing for the last several decades.

## Fix



We design and automate systems around break-fix principles: “self-healing” and “failing gracefully” operate under the assumption that resource congestion is inevitable. The task of monitoring applications and systems is put to software, but still relies on reactive triggers and thresholds that your people define (“guess at”).

## Repeat





And different teams are responsible for different layers of the application stack, each group using their own set of tools to make resource decisions within their specific domain, operating in an isolated way that doesn't reflect the interdependencies that exist between each layer.

# Application Resource Management

What if your applications automatically got the resources they need when they needed them?

Turbonomic uses a top-down, application-driven approach to automatically ensure that your applications always get exactly the resources they need to perform. The software understands every layer of the application stack and its dependencies, as well as the multidimensional resource needs that impact the quality of service (QoS) of your applications. With this application-driven, truly full-stack understanding, Turbonomic is able to continuously make the right resource decisions at the right time, whether it's scaling, placement, or capacity. It automatically executes these actions using the software-defined levers that your environment already exposes, for example, Kubernetes.



Turbonomic is the only vendor to deliver a platform where applications continuously and automatically get the resources they need, when they need them.

# Managing Kubernetes performance at scale with Turbonomic

Only Turbonomic stitches your Kubernetes clusters to the underlying infrastructure and makes resource decisions that understand all the interdependencies in the stack, including:

## 1 How should you size containers

**Container Resizing:** Only Turbonomic continuously analyzes containers' real-time demand for CPU and memory to automatically determine (and adjust) the right CPU and memory limits and requests. These actions can be executed in real-time, or as part of your existing deployment process. No need to set these thresholds for every single service!

## 2 When do you need to reschedule pods? To which node?

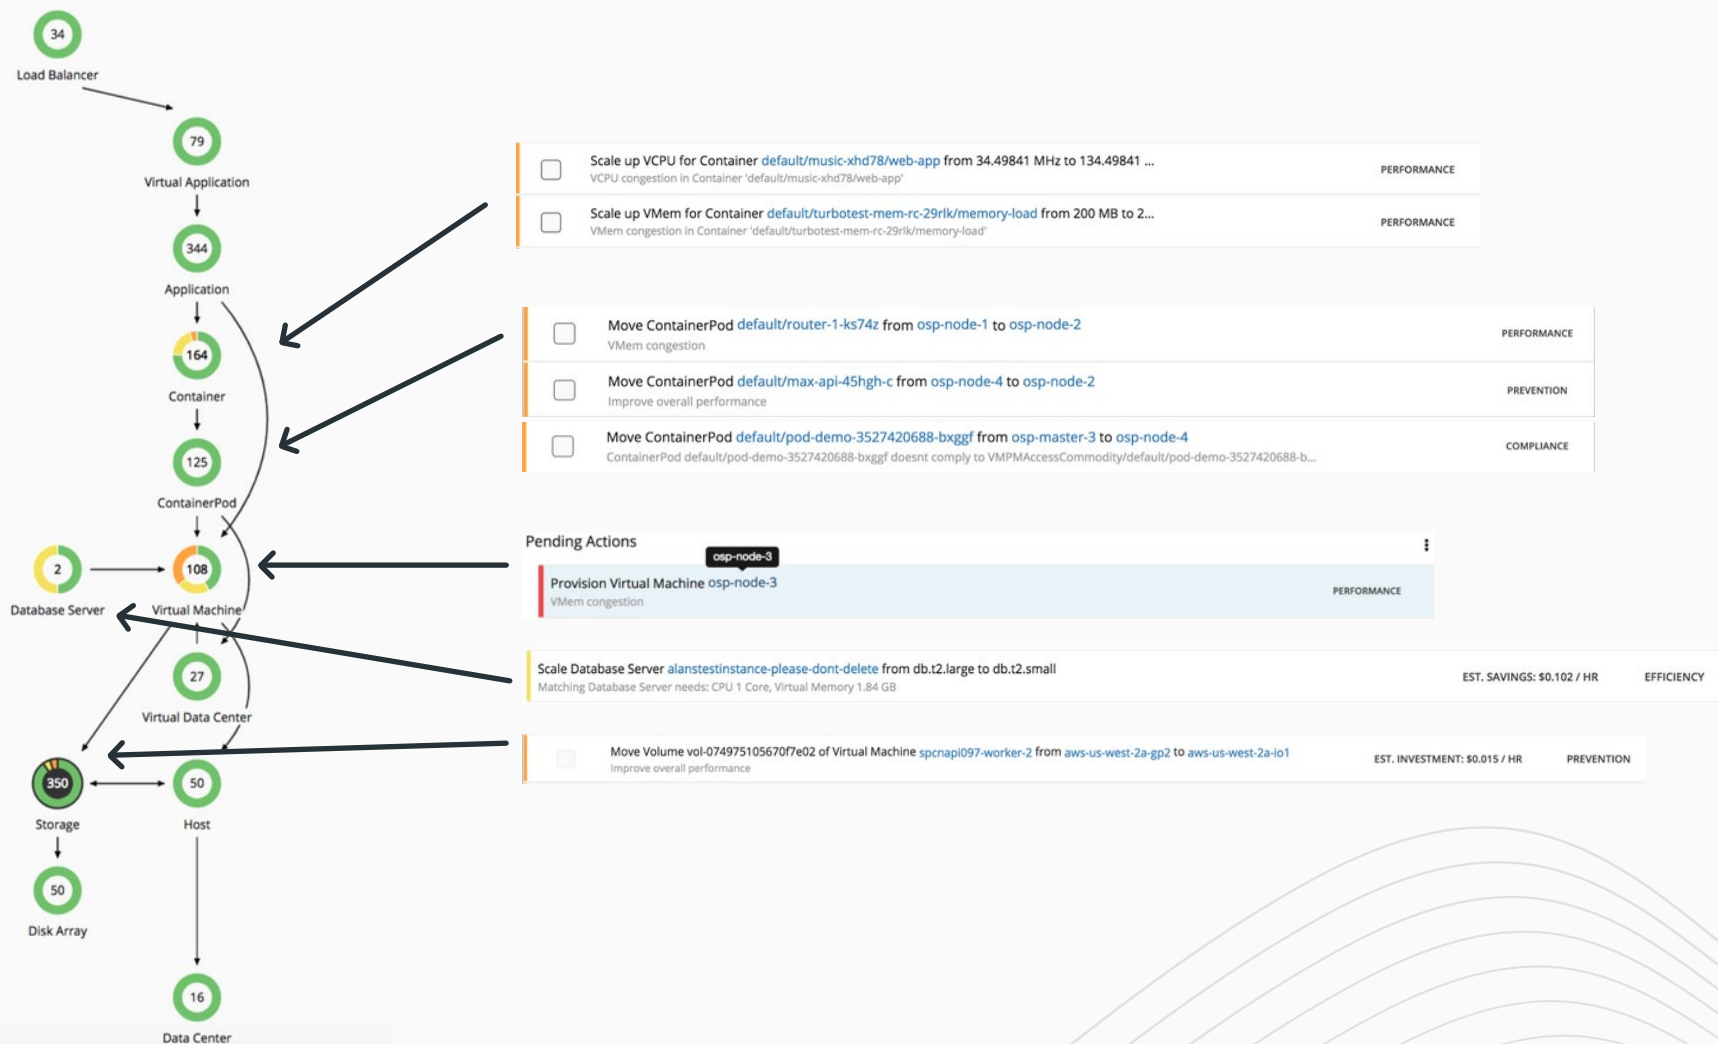
**Continuous placement:** Only Turbonomic can reschedule existing pods while assuring high availability of your service to prevent congestion from noisy neighbors or node resource starvation. Smart continuous placement also can reduce resource fragmentation allowing for more pods to be deployed

## 3 When do you need to scale out (or back) the cluster? By how much?

**Demand-Based, Infrastructure-Aware Cluster Scaling:** Only Turbonomic can simultaneously control the Kubernetes cluster and its underlying infrastructure. Node scaling is based on real-time CPU and memory demand with complete knowledge of the underlying resource availability--whether private cloud, public cloud, or bare metal.



# Only Turbonomic automatically provides the right actions for every layer of the stack.



# Conclusion

Applications and the infrastructure they run on will continue to change. The approach that Turbonomic takes to solve the problem of how to continuously assure the performance of your applications is founded on fundamentals that will not change: applications will always have resource needs that must be met by a supply chain of resource providers and consumers. Investing in a platform that takes a top-down, application-driven approach is to invest in one that will carry you through cloud native, multicloud, and beyond

