

turbonomic

THE CIO GUIDE

— *to* —

AUTONOMOUS IT

“Without change there is no innovation, creativity, or incentive for improvement. Those who initiate change will have a better opportunity to manage the change that is inevitable.”

- William Pollard

EXECUTIVE SUMMARY

IT Trends are Increasing Complexity

Organizations are adopting cloud computing to accelerate service delivery. Some try to deliver cloud economies of scale in their private data centers with the mantra “automate everything,” a philosophy often simpler in theory than practice. Others have opted to leverage public cloud resources for the added benefit of the pay-as-you-go model, but are finding it difficult to keep costs in check. Regardless of approach, cloud technology poses the same challenge IT has faced for decades: how to assure application performance while minimizing costs.



Within this new territory, however, a number of trends take addressing this problem to heightened levels of complexity and difficulty:

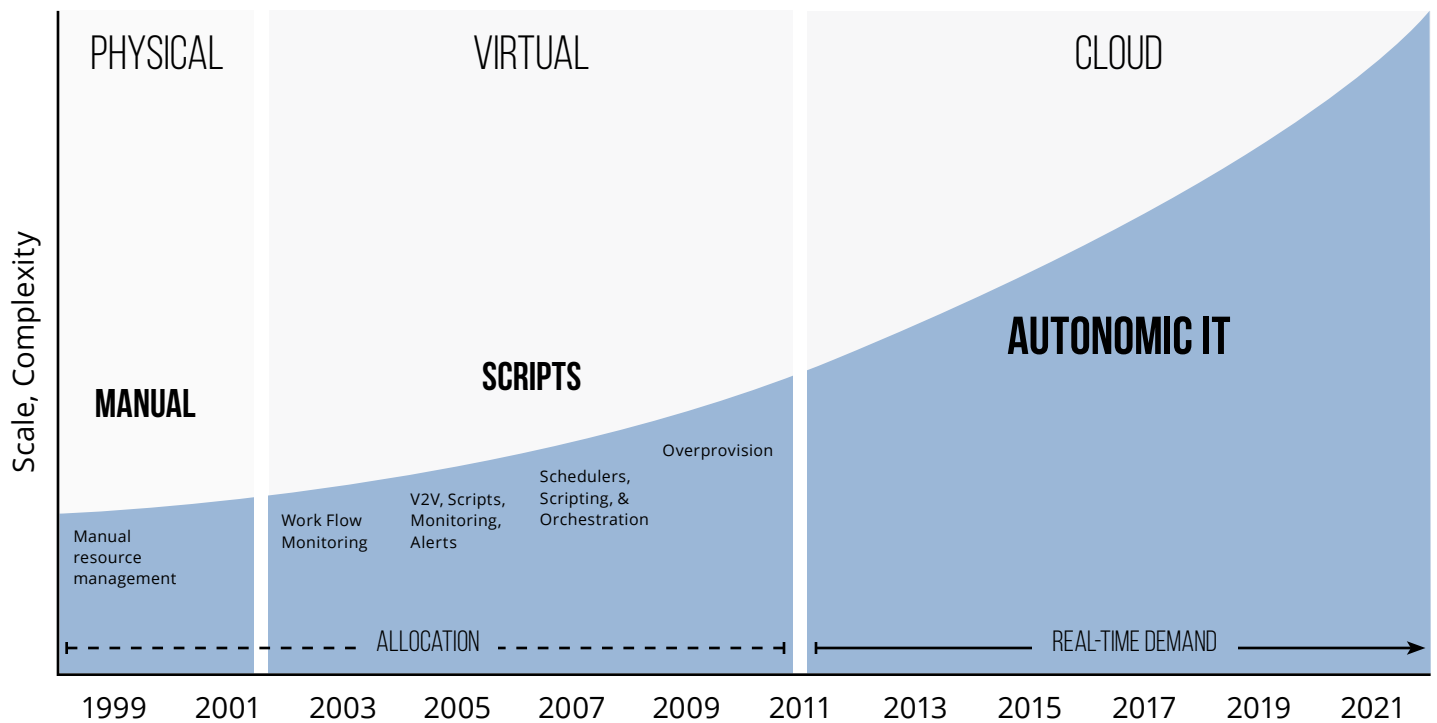
- **Cloud computing and self-service portals** increase the speed, frequency, and volume of changes to a given environment. [A recent report by McKinsey found that cloud-savvy leaders can provision x86 servers in a private cloud within 60 minutes, or in a public cloud within 30 minutes](#) - gone are the days where a lead time of days or weeks was acceptable.
- **Public clouds offer speed and elasticity**, but come with the challenge of controlling the cost overruns. Too often, developers and applications teams will spin up cloud instances faster than what IT can deliver in private cloud. This “Shadow IT” results in unknown, unmanaged, and often abandoned cloud resources, which cause surprising spikes in overall spend. Accordingly, [RightScale’s 2017 State of the Cloud Survey reports that optimizing costs in the cloud is the top initiative for surveyed organizations.](#)
- **Containers and microservice architectures** allow (with extensive automation) development teams to push small, but numerous, changes to parts of their applications rapidly and repeatedly. [Netflix, for example, has enabled its developers to launch hundreds of software changes a day.](#) Enhanced speed is made possible, but with it come exponentially more moving pieces in the environment.
- **Multiple technologies are a given - multiple clouds will soon follow.** “Heterogeneous” is no longer just about multiple hypervisors, databases, or storage flavors in an environment, but multiple environments, whether on-premises, in a cloud, or across multiple clouds. [For example, the RightScale Survey cited above also found that 85 percent of enterprises have a multi-cloud strategy.](#) Can your organization afford to have a different approach for performance-cost management for every workload and every cloud?

Cloud technology poses the same challenge IT has faced for decades: how to assure application performance while minimizing costs.

It quickly becomes apparent that the self-service enabled by cloud computing and the microservice application architectures made possible by containers create highly dynamic environments with more moving parts. At the same time, the industry offers a multitude of technologies, clouds, and payment models for IT to service their business and control costs. Every organization is different, as are their choices about what solutions best suit them. IT is now managing more heterogeneous environments than ever—and naturally, that heterogeneity extends to clouds.

These trends make assuring application performance while maximizing efficiency considerably more challenging. What’s more, this challenge is ushered in by an era in which enterprises are becoming technology companies as a matter of competitive advantage. IT is under greater pressure to both facilitate and support this rapid development, therefore leveraging cloud technologies with rapidly increasing adoption as a response; Morgan Stanley’s recent AlphaWise CIO Survey predicts that by the end of 2020, 49% of application workloads will reside in the public cloud (up from 19% in 2017).

The complexity of infrastructure management has been increasing exponentially since inception. Virtualization pushed it to new limits. Today, cloud and cloud native technologies do the same, but at even greater scale.



Given the undeniable criticality of IT to the business, the CIO that cannot deliver predictable performance while minimizing cost puts their business (and their job) at considerable risk. As they strive to adapt to the trends outlined above, these leaders are realizing the limitations of traditional approaches. Monitoring, alerts, schedulers, scripting, and orchestration cannot assure performance alone. They risk human error and operate in a break-fix loop that only corrects errors once they have already occurred.

It is time for CIOs to rethink their approach to IT.

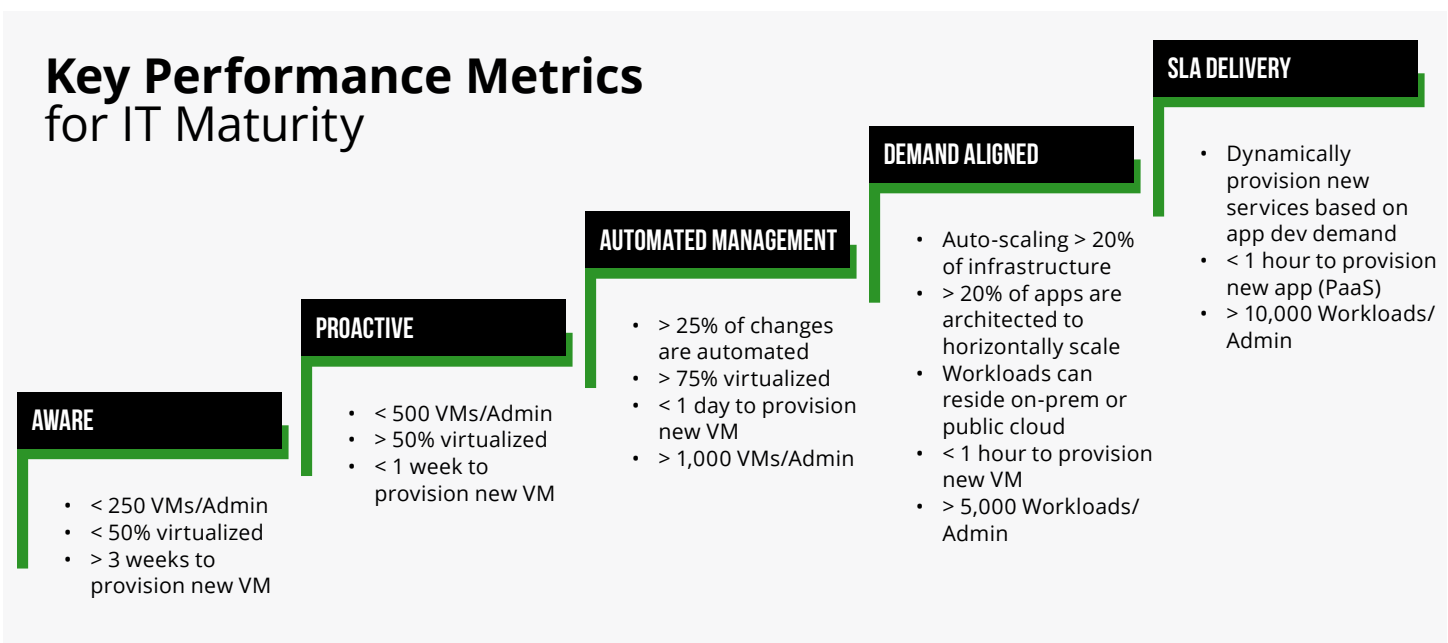
This document is for the CIO that seeks to transform their business through better and faster IT service delivery. It will discuss five stages of IT maturity and how—no matter where an organization exists today—adopting an autonomic approach that allows workloads to self-manage on any infrastructure or cloud can help.

THE 5 STAGES OF IT MATURITY

Organizations today vary in their level of IT maturity. The ability to adopt new technologies and processes, pressure from the business, as well as the beliefs and aspirations of CIOs all play a role in how and why organizations mature differently. The remainder of this document describes five stages of IT maturity to enable CIOs to understand where their organization is today and where they can aspire to reach. The stages are as follows:

- 1. Aware:** Collect and review performance and cost metrics.
- 2. Proactive:** Proactively make changes to assure performance and control cost.
- 3. Automated Management:** Continuously assure performance within cost constraints.
- 4. Demand Aligned:** Auto-scale apps and infrastructure based on real-time demand.
- 5. SLA Delivery:** Deliver differentiated service levels within cost and business constraints.

Key performance metrics make it possible to quickly identify where an organization is in terms of IT maturity. These KPIs include ratio of VMs per admin, percent virtualized, and time to provision to workloads. Determining where your organization resides is the first step towards progress.



No matter the stage in which an organization exists, the performance-cost challenge that the industry has grappled with since its beginning can only be solved by software that enables environments to self-manage. As environments scale and become more complex, the need for self-sustaining and autonomic infrastructure increases.

Adopting an autonomic approach to IT requires a change in mindset across the organization. But whatever their stage of IT maturity, CIOs and their organizations can benefit from the value that autonomic technology provides.

Adopting an autonomic approach to IT requires a change in mindset across the organization.

STAGE 1: AWARE

For organizations residing in the first stage of IT maturity, performance is a matter of uptime and cost is the common language between IT and the business. CIOs want their teams to resolve issues as quickly as possible and they demonstrate value to the business primarily through cost-cutting.

Goals for Aware Organizations

- Reduce trouble shooting
- Improve density

Key Challenge in Stage 1: *Reactive, Allocation-Based Approaches are Inefficient*

For many organizations, monitoring, balancing the environment, and allocation-based resource provisioning have always been the way to manage IT. Unfortunately, these approaches address performance degradation after the fact, typically resulting in an operationally inefficient multi-team blame game. Meanwhile, allocation-based approaches that take no account of real-time demand inevitably result in over-provisioning, given that historical peaks and app-owner demands are driving these decisions.

Key Opportunity in Stage 1: *Understand the Environment and Specific Risks*

Today, IT organizations can not only get full-stack, multi-cloud visibility into their environment, but actions that explain the current performance and/or efficiency risks and how to avoid them. With a holistic understanding of their environment and the software-derived answers for how to optimize it, organizations can mature beyond break-fix, allocation-based approaches that do not scale. However, in order to achieve continued scale, organizations must look towards operating proactively rather than simply understanding where risk exists and how to mitigate it: enlightened CIOs are challenging this status quo, realizing the operational necessity of avoiding performance degradation altogether and the potential efficiency gains to be had from provisioning resources based on real-time demand.

ORGANIZATIONAL SNAPSHOT

1. Service defined on annual basis
2. Changes planned weeks in advance
3. Infrastructure budgeted for peak allocation
4. Monitoring production for uptime
5. <250 VMs/Admin
6. <50% virtualized
7. >3 weeks to provision new VM

STAGE 2: PROACTIVE

ORGANIZATIONAL SNAPSHOT

1. Plans based on known cyclical changes
2. Changes only during planned windows
3. Budget based on historical utilization
4. Show-back/charge-back policy in place
5. Uptime SLAs
6. <500 VMs/Admin
7. >50% virtualized
8. <1 week for new VM

CIOs of proactive IT organizations prioritize maximizing infrastructure ROI and staff productivity. They believe addressing performance issues before they occur is necessary to those ends and that software does a better job of it than people.

Goals for Proactive Organizations

- Maximize infrastructure ROI
- Maximize staff productivity

Key Challenge in Stage 2: *Threshold-Based Policies are Costly*

Traditional “proactive” approaches operate on thresholds or policies set by administrators, which involve a mix of trial-and-error to set up and over-provisioning to mitigate performance risks. Performance improves in this stage, but at great cost to team productivity and infrastructure densities.

Key Opportunity in Stage 2:

Execute Decisions to Address Real-Time Resource Demand

When workloads self-manage, they make placement, sizing, and provisioning decisions based on the resources they need, when they need them. Using an autonomic platform, administrators can diligently execute the decisions provided, freeing themselves from the heavy lift of defining thresholds and policies, monitoring, and balancing the environment. This empowers individuals to manage 10 times more workloads and reduce provisioning time by nearly 80%. Because the platform’s decisions are based on real-time resource demands, administrators increase resource utilization by at least 20% and workload performance improves upwards of 20-30%. While the manual execution of the platform’s decisions increases FTE productivity and avoids performance risks, however, the approach depends on people being continuously available to execute.

STAGE 3: AUTOMATED MANAGEMENT

In the third stage of IT maturity, CIOs look to increase the speed of delivery, while enabling teams to manage more workloads with fewer administrators. They do so by offloading non-value-added tasks to automation, which frees up FTE resources to focus on higher impact work.

Goals for Automated Management Organizations

- Increase speed of delivery
- Manage more workloads with fewer admins

Key Challenge in Stage 3: *More Frequent Deployments Heighten Risk and Complexity*

As CIOs prioritize self-service to speed up delivery and enable development teams, their organizations are under pressure to do so without risking performance or adding FTEs. Be it private cloud or public cloud, more frequent deployments increase changes to the environment, putting performance at risk as operational resources are stretched thin. Ongoing and initial workload placement is critical to safely increasing the speed of delivery, thus the organization invests countless engineering hours to define automation policies or thresholds for every new workload.

Key Opportunity in Stage 3: *Automate Placement Decisions to Accelerate Speed of Deployment*

With an autonomic platform, organizations at this stage can automate initial and continuous placement decisions to lighten the load. Team productivity improves upwards of 20-40%. They can also safely increase the speed of delivery by more than 75%, knowing that the environment will self-adjust and optimize with any new workload and change in resource demand. However, if the goal is to truly keep up with business and application team needs, IT cannot stop at automating placement decisions.

ORGANIZATIONAL SNAPSHOT

1. Management of non-critical workloads is automated
2. Workload TCO is known and managed
3. Applications are managed holistically (multi-tier)
4. Deployment templates are standardized
5. Performance is measured in >20% of apps
6. >25% of changes are automated
7. >75% virtualized
8. <1 day for new VM
9. >1,000 VMs/Admin

STAGE 4: DEMAND ALIGNED

ORGANIZATIONAL SNAPSHOT

1. Seamlessly on-board new services with no scheduling required
2. Failure is assumed
3. No change windows
4. Workloads can reside on-prem or public cloud, with public cloud costs considered in placement/scaling
5. Auto-scaling >20% of infrastructure
6. >20% of apps are architected to horizontally scale
7. <1 hour for new VM
8. >5,000 Workloads/Admin

Organizations in the fourth stage of IT reliably and rapidly deliver services at minimal cost, having completely removed people from day-to-day infrastructure management. They not only automate placement decisions, but also real-time sizing and scaling, and are now focused on effectively managing cost and performance differentials between environments to scale applications and minimize complexity throughout the process.

Goals for Demand Aligned Organizations

- Maximize efficiency across multiple environments
- Minimize complexity while enabling rapid horizontal scaling

Key Challenge in Stage 4: *Standard Auto-Scaling Approaches Rely on Pre-Defined Policies & Thresholds*

Typically in this stage, more than 20% of the apps are architected to horizontally scale and as such, the infrastructure has been engineered to auto-scale as needed, using traditional policy/ threshold approaches. In most cases, public cloud resources are also leveraged for their speed of self-service and/or elasticity. This influx of users, applications, and resource consumption options increases the frequency and magnitude of changes occurring within the environment, while continuing to rely on allocation-based scaling approaches and best-guess determinations when it comes to where workloads should reside. This can result in wasted dollars and FTE hours if the correct decisions aren't made in real-time, which is difficult to do in dynamic environments without the help of software.

Key Opportunity in Stage 4: *Autonomically Execute Workload Placement, Provisioning, and Scaling Decisions in Real-Time*

The benefits of decision automation software have operational, performance, and cost implications. Letting software dynamically optimize the environment based on real-time demand not only assures the performance of applications, but improves it by as much as 60%. At the same time, organizations can fully exhaust the elasticity of public cloud resources, while reducing the infrastructure footprint on-prem by 40-60%. A fully self-managing environment allows the business and application teams to get IT services exactly when they need them; but, the most mature IT organizations understand that the customer is king, and the business benefits from differentiated service levels based on an application's criticality.

STAGE 5: SLA DELIVERY

Organizations in the fifth stage of IT maturity have achieved rapid service delivery and now strive to deliver on sophisticated SLAs. They can dynamically provision new services based on application or developer demand, in some cases provisioning new apps in less than an hour via a PaaS model. Furthermore, these organizations manage costs for public and private cloud infrastructure, with some dynamically bursting to public cloud or on-demand infrastructure.

Goals for SLA Delivery Organizations

- Deliver on sophisticated SLAs through enabling differentiated service levels
- Achieve higher levels of abstraction geared towards greater agility

ORGANIZATIONAL SNAPSHOT

1. Dynamically provision new services based on app dev demand
2. System adjusts in real-time based on QoS SLAs and cost constraints
3. Cost for public and private infrastructure is managed holistically
4. Dynamic bursting to public cloud or on-demand infrastructure is enabled (>40% is cloud/ infrastructure agnostic)
5. Re-architecting for "server-less" computing
6. <1 hour for new app (PaaS)
7. >10,000 workloads/Admin

Key Challenge in Stage 5: *Traditional Approaches to Differentiated QoS Delivery Create Silos*

Latency in revenue-generating applications has expensive consequences compared to, for example, corporate back-office applications. These organizations understand this distinction, and rely on segregating the environment to deliver differentiated quality of service (QoS) across applications and adhere to SLAs. While providing benefits to the customer, this approach inevitably results in wasted resources due to the silos it creates. Traditionally, CIOs will justify this cost given the business criticality of certain applications – customer satisfaction is paramount. At scale, however, these losses can be untenable.

Key Opportunity in Stage 5: *Autonomically Define QoS*

The complexity of achieving differentiated services levels while flattening the infrastructure to maximize efficiency necessitates autonomic IT. Only software can simultaneously ensure rapid service delivery, app-specific QoS levels, reliable performance and maximum efficiency across heterogeneous environments. With an autonomic platform, administrators can simply define QoS levels for an application and the environment adjusts in real-time to meet those SLAs while abiding by cost constraints.

CONCLUSION

The holy grail of IT is to simultaneously achieve rapid service delivery and unparalleled performance—both in terms of speed and reliability—while minimizing costs. Every CIO is leading their organization on a journey in pursuit of this goal, and every route is different. No matter where you and your teams are in terms of IT maturity, however, leveraging software to enable environments to self-manage has benefits from day one *and* as you progress towards IT service excellence, transforming the experience of leaders, administrators, and end-users alike.

Today, over 1700 customers, such as Barclays, Merck, JPMorgan Chase, and Travelport are on this journey, leveraging Turbonomic's autonomic platform to progress their development towards SLA delivery.



CUSTOMER SNAPSHOTS

Autonomic Snapshot

TRIBUNE MEDIA

Industry: Conglomerate (television broadcasting)

Challenge: In 2014, Tribune split its two main business units into two separate companies, Tribune Media and Tribune Publishing. IT needed to divide the virtual environment and also had to plan for a full data center move from Chicago to Raleigh, NC.

Scope: Turbonomic manages 100% of the environment, including storage (NetApp and Pure Storage) and 250 virtual desktops.

Automation: vMotion

Savings: Estimated \$400,000 in savings and consolidated 1,300 VMs to 800 VMs.

“We used Turbonomic throughout the move for a number of projects, including planning and running what-if scenarios to ensure we had the right resources in our new data center, as well as to ensure our applications didn’t suffer from the changes we were making.”

Michael Cannella

Enterprise Cloud Architect

“In our Corporate environment, Turbonomic’s automation has been crucial. It handles workload migrations between our hosts, and corporate datastores. We’ve seen an 80% reduction in the need to manually move and balance hosts; this is time we now use to make actual engineering decisions versus daily clean up.”

Matt McColm

Senior Data Center Engineer

Autonomic Snapshot

EDF RENEWABLE ENERGY

Industry: Renewable Energy

Challenge: Inability to guarantee performance of mission-critical applications; inefficient use of virtual and human resources through manual management.

Scope: Turbonomic operates across 3 data centers and 40 remote sites.

Size: ~1,300 VMs

Automation: Ongoing VM Placement

Results: 80% reduction in manual management

CUSTOMER SNAPSHOTS

Autonomic Snapshot

DIZZION



Industry: Service Provider

Challenge: Dizzion had created their own cloud environment from the ground up to deliver the most robust and rapid virtual desktop solutions available. While in a period of hyper-growth, Dizzion saw more than 100% year-over-year growth. Their great success made it more difficult to ensure high-performing services and accurately scale their infrastructure.

Scope: 8 cloud environments across the globe

Automation: Ongoing VM placement and sizing

Results: Reduced purchasing lifecycle from eight new servers every two months to eight servers every two quarters, allowing Dizzion to efficiently scale at a critical time in its growth.

“Turbonomic knows exactly where to place and how to size our workloads so that our customers are having the best possible experience. This allows us to be hands-off with the management of our infrastructure, and to focus on continuing to give our customers the flexibility and quality experience they’ve come to expect from Dizzion.”

Rob Green

CTO

“The ability of Turbonomic to show such a fast ROI was incredible. It’s the only software company I’ve worked with that can do that.”

Mathew Rumphol

Manager, Infrastructure Engineering

Autonomic Snapshot

BLACKBAUD

blackbaud

Industry: eCommerce Software

Challenge: With the busy season right around the corner, Blackbaud was experiencing significant network latency impacting workload performance that would reduce their ability to process transactions quickly.

Scope: 1,300 workloads across 4 data centers, leveraging VCE Vblock™ infrastructure powered by vSphere®, Cisco UCS® and EMC XtremIO®

Automation: Automated recommendations

Results: With Turbonomic, Blackbaud processed 52% more transactions on Giving Tuesday 2015 than the previous year, totalling over \$39.6 million.

TURBONOMIC MATURITY MODEL

The Turbonomic Maturity Model was developed in partnership with over 1,700 leading organizations that have trusted Turbonomic to enable them to assure performance while minimizing costs.

Financial Services	Technology	Healthcare	Service Provider	Insurance	Telecom/Media	Retail	Manufacturing	Energy	Government

WHAT IS TURBONOMIC

[Turbonomic's Autonomic Platform](#) enables heterogeneous environments to self-manage to assure the performance of any application in any cloud or infrastructure. Turbonomic's patented decision engine dynamically analyzes application demand and allocates shared resources in real time to maintain a continuous state of application health. Matching supply to demand results in concrete actions that the system takes constantly such as: workload sizing and instance templating, placement, and provisioning, as well as horizontal autoscaling of applications. Generally, the larger and more complex a customer's environment and the greater the pressure to cut IT costs, the greater the need for self-managing infrastructure and cloud.

Customers automate decisions progressively as their proficiency with the platform increases. VM placement, for example, is the most common first step in decision automation. Then, customers may automate storage motion or vertical and horizontal scaling. With each level of automation, they realize measurable performance and efficiency benefits.

To learn how Turbonomic can help you and your organization, contact sales@turbonomic.com or visit turbonomic.com.

*Customers
automate decisions
progressively as
their proficiency
with the platform
increases.*

ABOUT THE AUTHORS

Charles Crouchman

Chief Technology Officer, Turbonomic

Prior to joining Turbonomic, Charles held senior executive positions at several technology startups including Cirba, Mformation Technologies, Opalis Software, and Cybermation, spanning product development, operations, and engineering. He brings 20 years of experience in the IT infrastructure and cloud automation industry.

Asena Woodward

Senior Product Marketing Manager, Turbonomic

In her seven years prior to joining Turbonomic in 2015, Asena worked with various early-stage startup and adolescent tech companies. She specializes in understanding new markets and educating customers on transformative technologies.