



WHITEPAPER

The Definitive Guide to Practical AIOps

www.turbonomic.com



The Intro

Today's digital transformation initiatives are aimed at providing better user experiences that drive increased customer engagement and loyalty. Customers are increasingly interacting with enterprises through digital and mobile experiences, and in turn business outcomes are determined by how quickly an enterprise delivers new services to market that surpass those provided by their competitors.

Applications underpin all experiences, and new generations of cloud native applications are architected to enable parallel development, rapid introduction of new services, and personalized customer experiences. Highly distributed applications now employ microservice- and container-based building blocks hosted across complex environments, spanning data centers, public cloud, and the edge. The complexity of managing these environments has given rise to software that can operate beyond human scale and comprehension to tackle the decision making of where and how applications can be supported.

In this eBook, we introduce the current practical applications of AIOps, provide guidance on how and why enterprises should be using AIOps today, and offer insight into what the future holds. The goal is to assist executives and architects in maximizing the business value they can receive from AIOps while preparing for a future where AIOps is an imperative to provide resilient application performance and free employees to focus on driving business innovation.

What is AIOps and Why Does it Matter?

AIOps" stands for "Artificial Intelligence for IT Operations". Simply put, AIOps is the application of Artificial Intelligence to IT Operations. According to Gartner's definition, "Artificial Intelligence (AI) applies advanced analysis and logic-based techniques—including machine learning—to interpret events, support and automate decisions, and to take actions."

It is important to note that AIOps is not associated with mimicking human intelligence. Rather, AIOps seeks to apply algorithms to solve specific problems, often much faster, much more accurately, and at much higher scale than a human could possibly solve the problem. As applications become more distributed and complex, and as the infrastructure those applications run on gets more distributed and complex (often spanning from data centers to public cloud to edge computing) it becomes untenable for applications to perform reliably and efficiently at scale without AIOps. And as enterprises who adopt AIOps are discovering, their employees can be more productive and spend more time on innovation when AIOps is freeing them from troubleshooting and other "keeping the lights on" activities.

A Practical Way to Think About IT Operations

IT organizations have embraced a traditional approach to Operations for several decades. Resources are intentionally overprovisioned to accommodate for future growth and potential unforeseen circumstances. Monitoring tools dedicated to constituent parts of IT (e.g., applications, servers, storage, networks elements, cloud instances) are pervasively used to track access, utilization, latency, adherence to rules, and other KPIs of interest. A combination of monitoring KPIs and generating threshold-based alerts is intended to indicate when actions need to be taken to avoid or mitigate issues or optimize efficiencies.

Most IT organizations spend a significant portion of their time tending to their IT environments. In many enterprises, there are so many alerts being generated by different monitoring tools that a class of monitoring has emerged to filter out all but the most severe performance issues or risks, which are then surfaced for IT staff to investigate and remediate. These tools incorporate advanced analytics and logic-based capabilities to classify alerts that can in all probability be ignored (and so are suppressed from view) so that staff can more quickly identify the root cause of a problem when a significant issue exists or better yet, address a risk before it becomes a big problem.

At its core, Artificial Intelligence for IT Operations, can be thought of as managing two types of challenges:

1. Things that are supposed to not change (e.g., static) or may change in predictable ways (e.g., periodic). AIOps addresses these scenarios through Machine Learning (ML) programs that establish baselines and identify a deviation from normal.
2. Things that may change in unpredictable ways (e.g., dynamic). While including ML capabilities associated with baselines and deviations, in order to appropriately respond to situations that have never been seen before, dynamic optimization is required.

Let's explore these management challenges using common examples in IT.

At its core, Artificial Intelligence for IT Operations, can be thought of as managing two types of challenges:

1. Things that are supposed to not change (e.g., static) or may change in predictable ways (e.g., periodic). AIOps addresses these scenarios through Machine Learning (ML) programs that establish baselines and identify a deviation from normal.
2. Things that may change in unpredictable ways (e.g., dynamic). While including ML capabilities associated with baselines and deviations, in order to appropriately respond to situations that have never been seen before, dynamic optimization is required.

Let's explore these management challenges using common examples in IT.

AIOps Use Case: Deviation From Baseline

Application Performance Management (APM) represents a class of tools that are used by the majority of Fortune 100 companies and large global enterprises. The outcomes sought through APM include facilitating root cause analysis (RCA) and remediation of application performance degradation issues, reduction in the frequency of application performance degradation, and reduction in the number of IT tickets associated with application performance.

Cisco AppDynamics is an example of a leading APM platform that embeds AIOps capabilities in the AppDynamics Cognition Engine. Unlike a human monitoring a wide spectrum of KPIs, the Cognition Engine can ingest, processes, and analyze millions of disparate records every second using streaming analytics technology. Using Machine Learning (ML) algorithms, the Cognition Engine establishes a baseline understanding of what normal looks like, and therefore has an understanding of what a deviation from normal behavior looks like and will elevate this data to greatly speed up root cause analysis in the presence of performance issues. The ML algorithms have the ability to recognize seasonal and periodic patterns, adjusting the baseline accordingly.

Here, AIOps works in concert with IT staff responsible for resolving issues and mitigating risks that may lead to issues. AppDynamics provides visibility into a rich set of KPIs and the analytics that provide actionable insight into what is happening and what to do about it.

AIOps Use Case: Deviation From Baseline

IBM Watson AIOps can incorporate structured and unstructured data from a variety of IT systems, including monitoring tools and workflow platforms. Logs and KPIs, alerts, trouble tickets and topology information can be ingested and analyzed by ML to establish baselines and norms over time. Watson includes AI and natural language processing technologies to correlate unstructured and structured data to provide as much context as possible about an event to speed RCA and remediation. Watson can integrate with collaboration tools to speed resolution and can also compare situations to events that have occurred in the past to learn from past resolution efforts.

IBM Cloud Pak for Watson AIOps is targeted for achieving the outcomes of more comprehensive incident analyses, faster incident resolution and improved incident avoidance. Cross-platform integration serves to bring siloed teams together, aligned around a common understand of the entire environment and circumstances, resulting in better end user experiences and the avoidance of downtime that can inflict both economic and reputational damage.

AIOps Use Case: Dynamic Optimization

Assuring application performance requires insight into the application's fluctuating demand and configuration to dynamically match demand to required resources. Every environment consists of multiple stacks (or layers); at the top is the most important entity, the Application. The purpose of IT operations is to make applications run well. Therefore, everything below the application layer should support the application.

Every resource in the stack can experience contention or issues that will impact the higher layers, all the way up to the application. For example, an instance on AWS EC2 assigned with an EBS volume not suitable for the application IOPS demand may negatively impact application performance despite having ample compute capacity.

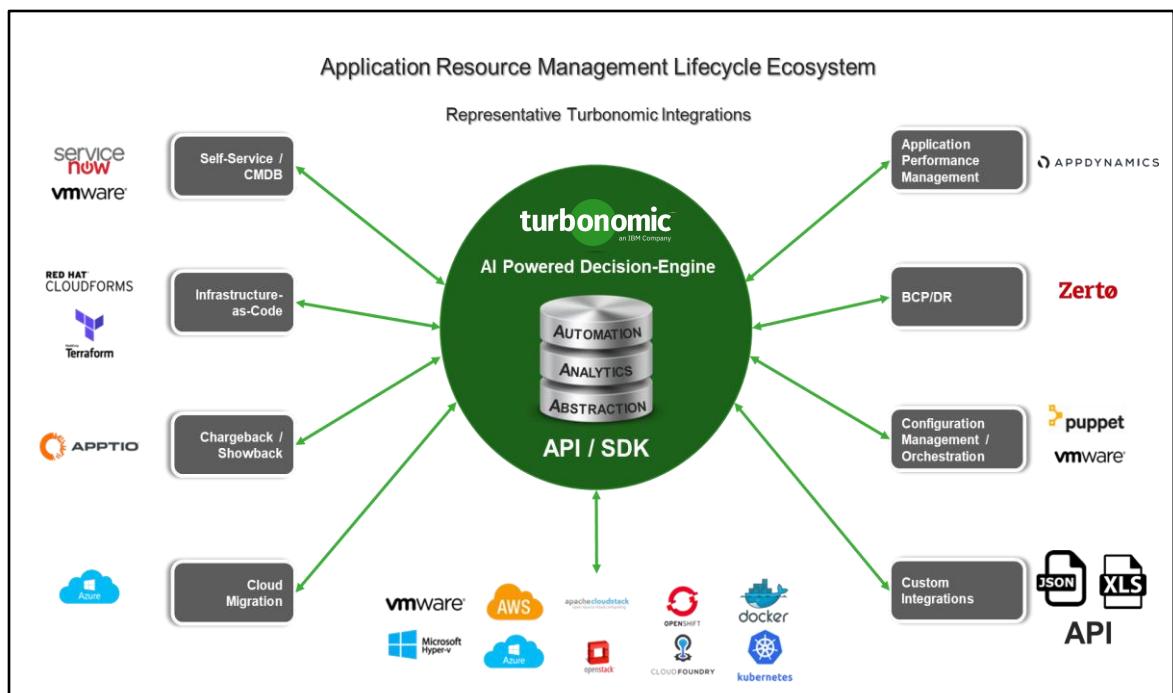
Turbanomic offers full stack visibility and control into every environment it manages. Turbanomic understands the relationships between the layers, the available capacity of resources and fluctuating demand of the application and can execute actions at each layer to assure application performance. Turbanomic has taken an approach of applying economic principles to manage IT resources. Application Resource Management (ARM) is accomplished by abstracting all layers of the IT stack, from discrete physical resources up through every layer to the application components, into a market economy supply chain. At each layer of the stack, analytics compare the available supply in the underlying layer to the demand being requested by the higher layers, to continuously make decisions about how to best meet demand. By continuously optimizing the environment, and making each layer application-aware, IT resources are proactively driven to a desired state of assured performance rather than waiting for something to go wrong and attempting to resolve it.

AIOps Use Case: Service Management Automation

To become agile, companies must embrace automation, from automating the provisioning of infrastructure (for example, Infrastructure as Code) to automating workload and application deployment. Traditional IT resourcing has remained reactive and manual due to siloed resource alerts that arise with limited context in relation to application performance. The burden of resource resolution falls to IT staff. Process automation is useful but insufficient.

The potential of Artificial Intelligence for IT operations (AIOps) is to avoid issues rather than waiting for them to occur and requiring IT staff to resolve them. Preventative measures must be determined and automated in application-aware software. Customers must gain confidence that the decisions made by the software are safe to implement.

Turbonomic is deployed in many of the world's largest and most complex environments, including more than half of the Fortune 100, where multiple teams and stakeholders initially reviewed the actions generated by Turbonomic before automation was turned on. Automation is not an event, but a journey that spans company culture, organizational processes, and IT operations workflow systems, fostering cross-team collaboration through a "single source of truth". Turbonomic integrates with major IT platforms and environments, enabling teams formerly in distinct silos to see their familiar data in the context of the larger environment, and thus seeing and understanding the relevance of initial automation decisions. Through a process of starting with manual actions, followed by actions scheduled during change windows and moving to immediate automated execution, IT organizations can move at their own pace and according to their own policies.

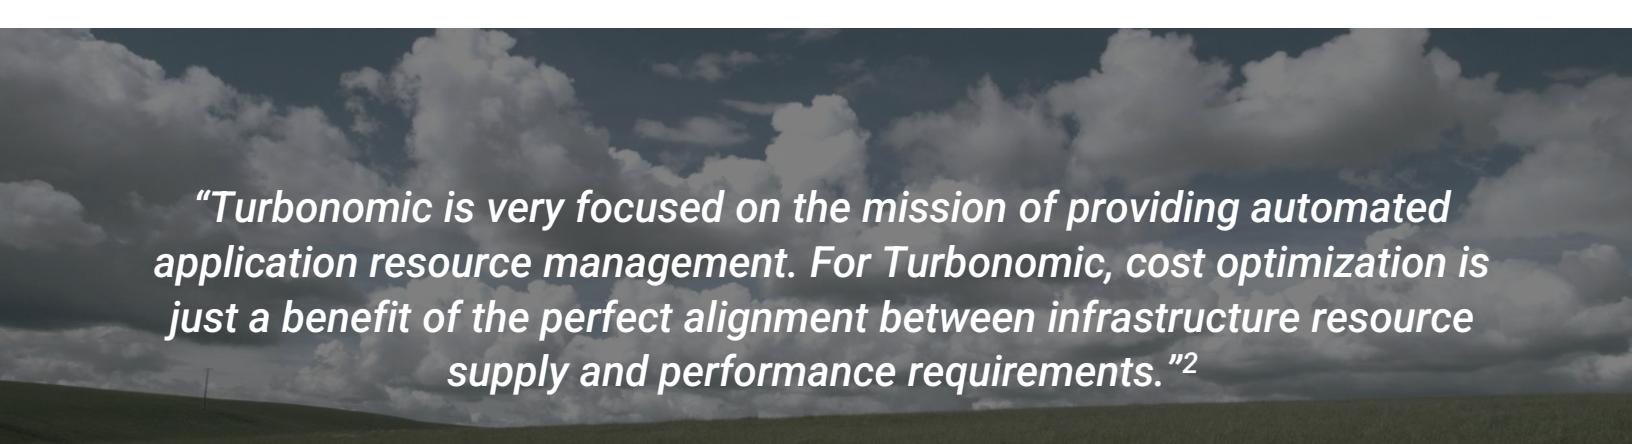


What Problem Are You Trying To Solve?

Most enterprises ignore thousands of alerts each day, manually addressing a subset of those that pose a risk of serious performance issue or threat, often not preventing issues but reacting to them. AIOps tools and platforms that incorporate ML techniques to establish baselines can detect meaningful deviations from baseline. The value they provide manifests in faster resolution of issues and the ability to prevent issues that arise from periodic fluctuations in application demand or IT resource consumption. Unfortunately, they are not able to take customers very far on the path to self-driving IT Operations because they cannot address unpredictable demand in dynamic environments.

As applications get more complex, and the IT resource options introduce new complexities such as the rapid ascent of dynamic and ephemeral container environments, it becomes more difficult to connect the dots between full stack relationships and dependencies. More importantly, the reason that applications are getting more complex and distributed is because digital engagement is continually evolving and becoming more dynamic, leading to unpredictable application demand situations. Applications constructed of hundreds of microservices that can be independently modified to continually update services for thousands or millions of consumers (who are also targets of new types of advertising campaigns) means that unpredictable events will become commonplace. Never-before-seen events may be a regular occurrence.

There is a clear and immediate need for a solution that can address the challenges related to multicloud adoption and containerization of applications that support highly dynamic demand scenarios. Simply put, the problem to be solved is giving applications the resources they need, when they need them. In order to assure application performance, any solution platform must understand application requirements (demand) as well as the relationships between each layer of the IT stack, from the application components to the physical resources that are assembled to build containers and virtual elements in a cloud environment. This understanding of the full stack relationships is what is missing in traditional IT toolsets and approaches.



"Turbonomic is very focused on the mission of providing automated application resource management. For Turbonomic, cost optimization is just a benefit of the perfect alignment between infrastructure resource supply and performance requirements."²

2. Gartner Technical Professional Advice Document G00441626

A Different Approach

The most important objective of IT infrastructure is to provide applications with the resources they need to deliver their services levels. A companion objective is to do this as cost efficiently as possible and adapt to changing environment and application demand scenarios by dynamically adjusting resources over time. Key capabilities include:

1. Application-Aware Optimization
2. Support for On-Premises, Hybrid, and Multicloud Deployments
3. Full Stack Visibility and Control Across the Full Environment
4. Trustworthy & Automatable Actions
5. Enforcement of Business Policy Compliance

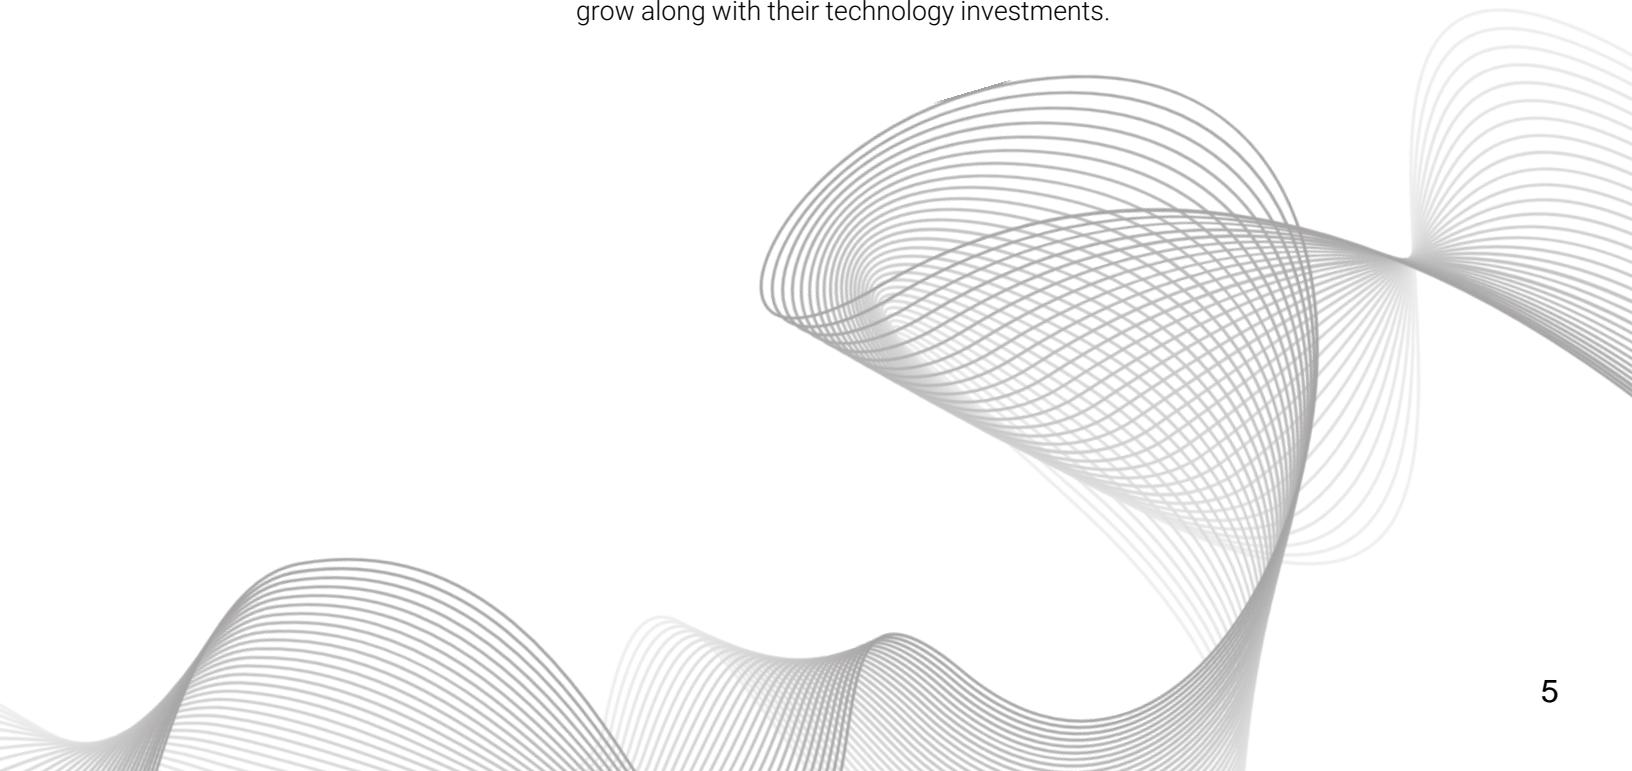
Abstraction, Analytics, Automation

Thus far, we've touched upon a number of use cases for AIOps, including

- Real-time anomaly detection for risk mitigation and problem avoidance
- Faster root cause analysis through event correlation when problems do occur
- Low priority alert suppression, so high priority alerts get better visibility
- Capacity planning and management, based upon predictive analytics
- IT Service Management automation

A primary challenge for large enterprises is one of exploding complexity due to modern applications built on microservices, deployed on modern containerized multicloud infrastructure.

Enterprises are rapidly adopting AIOps with ML embedded in many of their monitoring and management systems, including APM, Service Management, Infrastructure as Code and Configuration Management. Turbonomic serves as a control plane, a modern application hosting platform, to tie those systems together and scaling to millions of managed elements in a single instance. As AIOps technology continues to evolve in independent tools, Turbonomic learns about the changes through its integrations and incorporates that data into the Turbonomic AI engine's decision process, enabling large enterprise customers to grow along with their technology investments.



Conclusion

In today's competitive business climate, applications drive revenue, profits, customer experience and customer retention. Application success is determined by an enterprise's ability to innovate through development teams and IT operations, but application and infrastructure complexity, scale and scope are rising exponentially. AIOps will represent the biggest wave of change to sweep IT organizations in the coming years. Business leaders and IT Operations teams must decouple themselves from reactive tools and embrace technologies that accelerate business innovation, prevent issues and maximize cost efficiencies. Rather than responding to issues when they occur, AIOps offers probabilistic outcomes based upon advanced analytics. In some cases, dependent upon the sophistication of the AIOps platforms and the nature of the action, the action may be scheduled or taken immediately through integration with other IT systems.

Future planning must contemplate the speed at which these technologies evolve and the speed at which organizations will come to understand and trust AIOps generated insights and actions. IT staff will be increasingly called upon to focus on business innovation, and it will be up to intelligent AI software to manage the complexity of IT environments and create a control plane capable of self-driving operations that drive continuous health vs. merely striving to speed mitigation and self-healing.

About Turbonomic, an IBM Company

Turbonomic, an IBM Company, provides Application Resource Management (ARM) software used by customers to assure application performance* and governance by dynamically resourcing applications across hybrid and multicloud environments. Turbonomic Network Performance Management (NPM) provides modern monitoring and analytics solutions to help assure continuous network performance at scale across multivendor networks for enterprises, carriers and managed services providers.

For further information, please visit www.turbonomic.com

*www.turbonomic.com/resources/case-studies

