Self-Driving IT Operations

Navigating the 5 Levels in the Automation Journey

ca technologies

A Broadcom Company

The Quest for Autonomous IT Operations

It wasn't long ago that the prospect of driver-less cars was solely the stuff of futuristic science fiction, but now these vehicles are part of our very near-term future. It's not a question of if these cars will be unveiled, but when, and their impact will be enormous.

So, if driver-less cars are within reach, where are driver-less IT operations?

When will IT monitoring advance to a nirvana state? A state where AI and machine learning for IT operations, also known as AIOps, transform reactive and backward-looking monitoring into a fully autonomous function that learns and constantly optimizes applications according to the business outcomes they support.

In the automotive arena, cars will advance through multiple levels of automation before they can be fully autonomous. The Society of Automotive Engineers (SAE) International has developed the JS3016 standard, which describes six levels of driving automation, spanning from level 0, or no automation, to level 5, representing full automation. These levels can provide an informative analogy for how organizations can advance to a world of autonomous IT operations using AIOps.

In the following pages, we examine how IT teams will be able to navigate the various levels of automation required.

Levels 0 and 1:

Establishing Algorithmic Noise Reduction and Anomaly Detection

In level 0 cars, all functions are monitored and performed by the driver. This of course is a situation familiar to today's IT operations teams, who are over-reliant upon manual, often historical reports and static analysis.

At level 1, the car is still controlled by the driver, but some driver-assist features may be incorporated into the design, including adaptive cruise control and blind-spot warnings. Similarly, in IT operations, a level 1 state means staff are still behind the wheel so to speak, but analytics can be effectively employed to eliminate false positives and pinpoint anomalies—our own blind-spot equivalents.

While most of today's cars have a plethora of driver-assist wizardry, attaining level 1 in IT operations can be challenging. As IT assets have accumulated over the years, an operational function organized along technology lines becomes pre-disposed to acquiring tools designed for a discrete monitoring purpose. As a result, it's not uncommon for an enterprise to be using upwards of 20 monitoring products. The problem with this approach is that it's no longer scalable.

While level 1 represents the beginning of an autonomous operations journey, with analytics and machine learning controlling a limited set of assistance type functions, this phase is still extremely beneficial for modern organizations. With algorithmic noise reduction and anomaly detection, staff can progress from being passive or reactive passengers with limited control, to applying cumulative machine learnings as their improvement feedback loops. As a result, they become better drivers, ones that are more efficient and less stressed.

Level 2:
Automated Root-Cause

At level 2, self-driving cars deliver partial automation, with the vehicle helping with many functions, such as automated parking and lane departure prevention. Of course, the driver must still be ready to take control and always has full responsibility for safety-critical functions and monitoring.

While we now take partial automation in cars for granted, attaining this level in IT operations has proven more difficult. That's not through lack of monitoring sensors—far from it. Modern containerized applications decomposed into 1000s of microservices can produce a 10-fold increase in metrics. Now it's not uncommon for many IT ops teams to get more than 50,000 alerts per month. That's just too many for staff to realistically process using traditional monitoring, so it's not surprising that more problematic conditions and anomalies are often missed. Even worse, when alarms are relegated to the noise and nuisance bucket, they're easy to ignore completely.

Many operations teams routinely employ a variety of dashboards and agent-based technologies and many of these do a great job at helping teams find the root cause of a problem within a narrow technology domain. Where they fall down, however, is in the area of scalable cross-domain analytics—that is, handling massive increases in data across the entire technology stack and then correlating conditions to find the root cause.

To achieve level 2 AIOps and help teams automate root cause at a true "system level," modern AIOps platforms will require a unified data model. Dynamically built using a time-journaled, directed graph of objects, this model provides the foundation upon which to ingest, correlate and visualize the more complex problems arising across modern distributed applications and microservices.

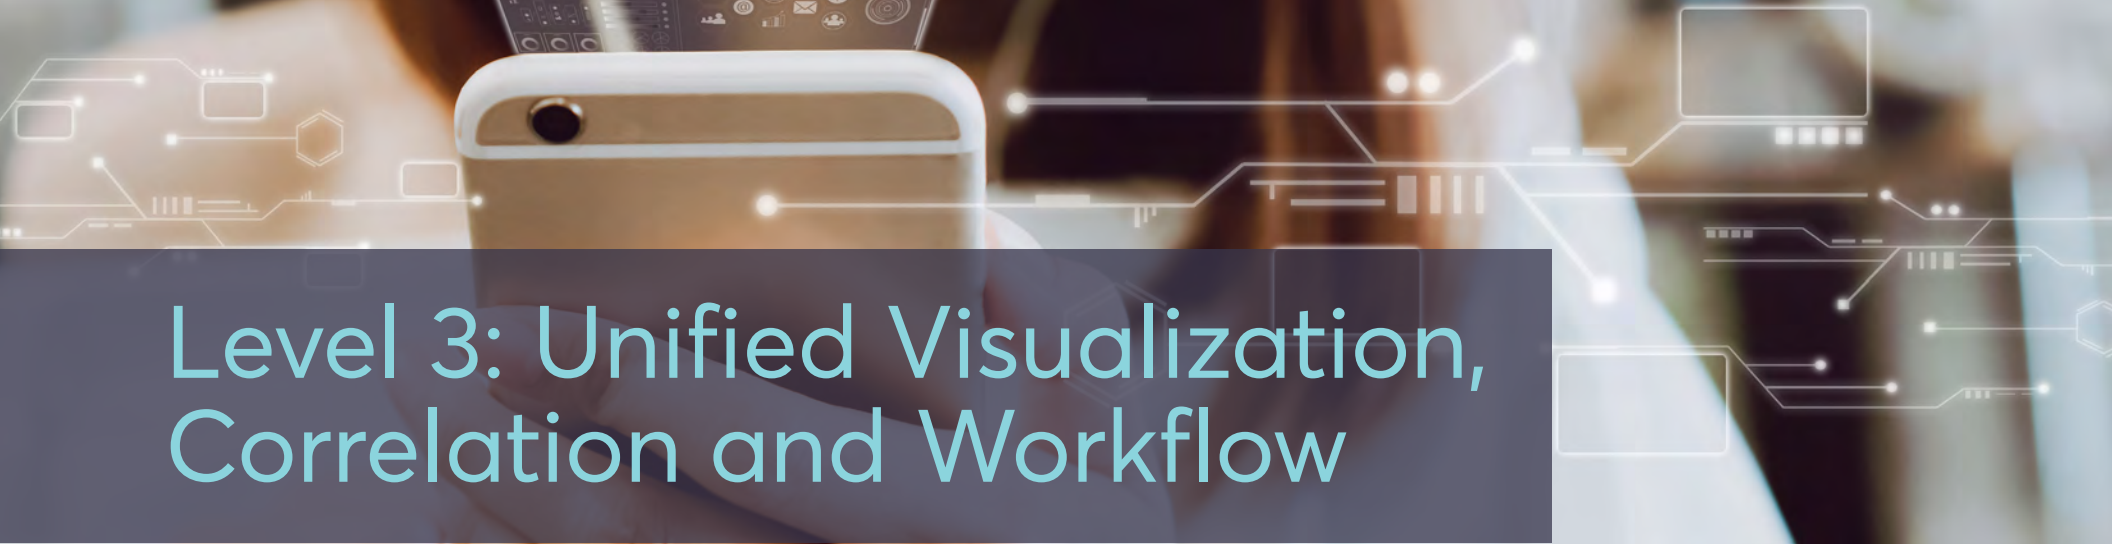

Level 3: Unified Visualization, Correlation and Workflow

In a level 3 car, we have what's commonly referred to as conditional automation. Here, the vehicle is capable of taking full driving control during parts of a journey under certain operating conditions. One example of this automation is found in the Audi A8 sedan, which features the Audi AI traffic jam pilot. This system allows completely hands-free driving when the car is on a freeway with clear lane markings, is traveling under 37 mph and there are no pedestrians or traffic lights.

However, achieving this level of automation requires some serious smarts. For example, the Audi system includes constant traffic map monitoring, 12 ultrasonic parking sensors, four 360 degree cameras, mid and long-range sensors and some innovative laser technology.

Level 3 AIOps requires similar smarts and correlation. With complex conditions, it's not enough to focus on narrow situational slices. Rather, advanced systems should be capable of ingesting a wide variety of data sources into a single data lake, and then correlating and prioritizing based on the potential business impact. For example, just as a level 3 autonomous car would avoid a barrier crash when lane markings have faded (because it also has proximity sensors), level 3 AIOps will be equally dexterous—continually analyzing and correlating across multiple domains to determine how emerging patterns may affect critical business services.

> By enabling a "hands-free" monitoring approach, staff effectiveness and efficiency can be significantly improved.

These types of systems have a profound impact on the IT operations function. By enabling a "hands-free" monitoring approach, staff effectiveness and efficiency can be significantly improved. At this level, the heavy-lifting of alarm and root-cause analysis is conducted by AIOps platforms, while the system is augmenting staff with advanced visualization and automated workflows.

Level 4: Self-Healing Operations

Now let's hop into the seat of a level 4 car. At this level, a car should be able to drive itself safely even if the driver does not respond when asked to intervene. Drivers can safely go to sleep or leave the driver's seat. The car will autonomously accelerate, slow down, pull over or park safely if the driver doesn't take control when requested. Level 4 is considered highly automated; however, full self-driving is only supported in certain conditions or locations.

The machine learning at this level will be highly sophisticated. Deep neural networks will provide advanced situational understanding, detecting and responding to emerging conditions and patterns over time. Sensors inside and outside the car will track driver attentiveness, eye movements and gaze, plus alert on conditions humans can't detect.

In a level 4 AIOps system we need similar intelligence. As with many driving conditions, today's application architectures are highly dynamic. Elastic, serverless architectures mean containers can spin up and down based on demand patterns. Add continuous integration/delivery pipelines and we have systems that can change in minutes. All of this leads to systems producing emergent behaviors that can no longer be predicted; often and aptly called "unknown unknowns."

In such situations, static polling and simplistic correlation as a determinant for automated self-healing will fall short. Nuanced consideration of dynamic system behavior is a critical element of level 4 automation and self-healing applications. Similar to a level 4 car monitoring both driving conditions and the driver, advanced AIOps systems like these provide guardrails for IT operations, remediating problems and automating tasks, while giving control back to staff when needed.

Level 5: Continuous AIOps

Imagine a level 5 car, a fully autonomous vehicle that can drive anywhere. This level represents an advanced stage in which cars continuously perceive, sense and react to a range of conditions. If faults are predicted, the level 5 car will order the part and coordinate repairs. It's also likely that you'll be renting the car rather than owning it, so the vehicle will be learning and adapting to new behaviors, routes and conditions—continuously.

So what will a level 5 continuous AIOps system look like? Well, let's start with that word "continuous." Like a fully autonomous car, advanced AIOps systems will continuously process massive amounts of information at tremendous scale and apply real-time machine learning modules to gain new and deeper insights. No data will be off-limits, with logs, metrics and application performance instrumentation, user-experience data and IoT data all enriching the system with additional context.

These systems eliminate the human cognitive overhead associated with lengthy data gathering, cleansing, correlation and interpretation. Instead, the system uses massive learning sets to increase intelligence and deliver new capabilities over time. These systems will be continuously fixing issues and tuning performance, without operator intervention, and better still, without customers even realizing there's been a problem.

The Payoff

The impact of autonomous AIOps on IT operations will be profound. These systems will dramatically reduce the cost and human capital overhead associated with having valuable staff handling low-value activities. Instead of having staff tied up with mundane, interrupt-driven tasks that increase technical debt and degrade organizational capacity, these organizations will be empowered with optimization knowledge and learnings, so that everything and everyone get stronger.

Following are some of the characteristics that will define organizations with level 5 AIOps systems:

- Rather than fighting fires and fixing repeat problems, teams apply AIOps analytics across the continuous delivery pipeline to determine which applications, code, functions, practices and more correlate to the best performance and business outcomes.
- Instead of waiting for year end to guestimate infrastructure requirements, teams will use AIOps to continuously determine the optimum placement of workloads across elastic infrastructures.
- Rather than constantly having to hunt for the root-cause needle in a massive haystack, teams will use AIOps to provide an intelligent cost-benefit analysis of requests from business leaders, such as outlining the return on investment expected by a requested 100ms performance improvement.

Use Cases

Just as a fully self-driving car will serve many drivers in many conditions, level 5 AIOps will support multiple stakeholders and use cases; extending well beyond what's been considered the normal purview of IT monitoring. Following are a few of the potential examples:

1. Security

Advanced AIOps systems enable teams to correlate application performance data, security information and event management (SIEM) data and external threat intelligence inputs to proactively identify potentially malicious activity and risks—and automatically remediate them.

2. Cloud architecture risk analysis

With these systems, teams can gather all risk metrics associated with a particular public cloud application architecture, projecting failure points and outlining the appropriate exit strategies to executives.

3. Cloud workload optimization

With continuous AIOps, organizations can gain intelligence for continuously tuning and optimizing the management of production cloud workload variables—including instance types, network configurations, storage placement and more.

Accelerate Your AIOps Journey

Traditional methods simply can't continue to scale to support the unpredictability of modern dynamic systems. Consequently, IT teams must replace low-value, labor-intensive monitoring with more advanced AIOps systems. By advancing through the automation levels spelled out in this ebook, organizations will embark on a journey that will deliver massive benefits. By establishing advanced AIOps systems, organizations can benefit from fully automated learning and intelligence that can fuel continuous self-healing and optimized applications.

To get more insights into pursuing your own AIOps journey, check out "The Definitive Guide to AIOps."

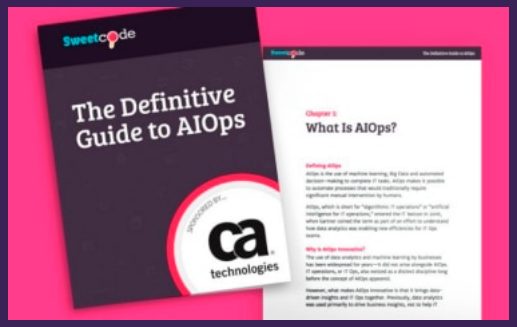

The Definitive Guide to AIOps

CA Technologies (NASDAQ: CA) creates software that fuels transformation for companies and enables them to seize the opportunities of the application economy. Software is at the heart of every business, in every industry. From planning to development to management and security, CA is working with companies worldwide to change the way we live, transact and communicate—across mobile, private and public cloud, distributed and mainframe environments. Learn more at ca.com.

CS_1118

ca technologies
A Broadcom Company