



WHITE PAPER



A Broadcom Point of View

How Automation.ai Delivers Scalable, Powerful, and Agile Machine Learning

Table of Contents

Executive Summary	3
Evolving IT Landscape and Operations	3
Introducing Digital BizOps from Broadcom, Powered by Automation.ai	4
How We Do It: The Automation.ai Architecture	5
Building a knowledge graph to optimize algorithm usage	5
Harnessing domain expertise	5
Constraining problem scope and employing small, reusable components	6
Employing ontological abstractions to establish a product-agnostic, software intelligence platform	6
Developing an open, flexible architecture	7
Employing an intelligent, flexible, scalable data model	7
The Advantages of Automation.ai	8
Contextualized, correlated intelligence	8
Current, complete, and accurate model of the run-time environment	9
Conclusion	10

Executive Summary

Today's business, IT, and development teams face unprecedented demands. In response, these teams have to gain unprecedented levels of speed, intelligence, and sophistication. Today, Broadcom offers the innovative solutions that help teams meet these new mandates. Digital BizOps from Broadcom integrates business, development, and operations data to generate actionable insights—enabling teams to continuously improve decision making and the execution of digital initiatives, so they yield enhanced business outcomes. Digital BizOps is powered by Automation.ai, the industry's first AI-driven software intelligence platform. This paper offers a detailed look at the architectural approaches Broadcom has taken in developing Automation.ai, and the unique advantages the platform provides.

Evolving IT Landscape and Operations

In recent years, the IT landscape has seen massive transformation, and the rate of change only continues to accelerate. In the not-too-distant past, monolithic applications and static infrastructures were the norm. Today, applications, and the way they're developed, delivered, monitored, and supported looks very different. Following are just a few of the aspects of this change:

- **Proliferating technologies.** Environments are now composed of a diverse mix of virtualization technologies, clouds, containers, microservices, orchestration systems, and more. Environments continue to grow increasingly dynamic and ephemeral. Increasing layers and types of security technologies continue to be implemented as well.
- **Expanding toolsets.** It isn't just that the volume and complexity of technologies has been increasing. These evolving environments accelerate the velocity of change and require continuous analysis. As a result, the number of tools adopted to manage these complex environments has also continued to expand. Now, it's not uncommon for IT organizations to be employing hundreds of tools for managing IT operations, security, and development.
- **Changing approaches.** Approaches like DevOps, BizOps, site reliability engineering (SRE), and continuous integration/continuous delivery (CI/CD) are emerging as norms, which all serve to accelerate the rate of change. These emerging paradigms place fundamentally different demands on team members, requiring staff to move from being specialists with a focus on specific tools and technologies to generalists who are dedicated to managing services.
- **Escalating demands.** For today's enterprises, the quality and performance of digital services is increasingly intertwined with business performance. It is paramount to bring innovative, compelling digital services to market faster. Teams need to maximize service levels and security at all times. At the same time, wringing maximum productivity out of staff and technology investments continues to grow increasingly critical.

In recent years, the IT landscape has seen massive transformation, and the rate of change only continues to accelerate.

All these factors are necessitating a fundamental change in the way business, IT, and development organizations work. These organizations simply can't continue to rely on legacy models and approaches to meet their charters. Moving forward, rigid, siloed tools and data sets, labor-intensive workflows, and limited, disjointed visibility are an increasingly imposing impediment to internal teams' ability to meet their charters and improve business outcomes. To fully leverage the business value, organizations need to bring AI and machine learning into their workflows. By leveraging both AI and humans, organizations can make better decisions than using either one alone.

Introducing Digital BizOps from Broadcom, Powered by Automation.ai

To keep in front of their organization's rapidly evolving technological environments, business requirements, and security threats, teams need new AI and machine learning technologies that augment and automate decision making. Broadcom is focused on delivering advanced solutions that address these emerging requirements.

Digital BizOps from Broadcom integrates business, development, and operations data to generate actionable insights. Powered by Automation.ai, this solution enables teams to establish continuous improvement in the business outcomes of digital initiatives.

Digital BizOps works across the digital business lifecycle, creating a continuous feedback loop that enables enterprises to enhance decision making, transform customer experience, increase employee productivity, improve operational efficiency, and speed innovation.

Automation.ai is the industry's first AI-driven software intelligence platform. Automation.ai harnesses the power of advanced AI, machine learning, and Internet-scale, open-source frameworks to transform massive volumes of enterprise data into actionable insights. Automation.ai offers this unparalleled combination of features:

- **AI-driven services.** The platform provides a predefined, prepackaged set of automated, AI-driven services, including for analysis, correlation, recommendation, and remediation.
- **Openness.** The platform ingests data from across the software development lifecycle, and from a comprehensive range of systems, including Broadcom products, third-party tools, and open-source platforms. The solution relies on ontological abstraction to integrate diverse sets of IT and business data that comes from different domains. Through this approach, the platform can build a unifying model that describes all aspects of the enterprise environment, which means that semantically rich context can be transferred from one domain to the other. Through this unified visibility, teams can eliminate the gap that's persisted between IT, development, and business organizations.
- **Continuous learning.** Automation.ai continuously learns the activities and architectures of the enterprise, and it validates and improves decision making methodologies based on real-world outcomes.

To keep in front of their organization's rapidly evolving technological environments, business requirements, and security threats, teams need new AI and machine learning technologies that augment and automate decision making.

- **Extensibility.** The platform can be run entirely independently or incorporated into a customer's existing AI and machine-learning frameworks.
- **Multi-cloud support.** Automation.ai employs Kubernetes-based orchestration and all of its components are fully containerized, enabling efficient deployment across multiple ecosystems, including public and private cloud environments.

How We Do It: The Automation.ai Architecture

In the realm of astronomy, scientists have long been able predict solar eclipses, and over the years, they have been able to do so with ever-increasing specificity and certainty. They are able to do so because models of the solar system have been getting better with each technological advancement. Today, we have comprehensive models and calculations that enable us to determine exactly when an eclipse will appear, where it will be most visible, and so on. Through our solutions, we sought to deliver this kind of comprehensive model for business, IT operations, and development.

In developing Automation.ai, we set out to create a modern platform for organizations. Instead of trying to rearchitect our existing tools, we chose to establish an intelligence layer that could sit on top of a number of specific solutions. We also set out to make it possible to access data from across a number of silos, while preserving the context of the source content and enabling this context to be shared efficiently. This contextual, comprehensive visibility is vital in delivering the actionable insights today's decision makers, developers, and operations teams need.

In the following sections, we offer a detailed look at some of the key approaches we've taken in developing Automation.ai.

Building a knowledge graph to optimize algorithm usage

Often, when it comes to deriving value from machine learning, it isn't the algorithms themselves that matter. Typically, what matters is the way these algorithms are orchestrated and scoped. To be able to use algorithms effectively, we set out to maintain a multi-domain knowledge graph that describes the enterprise in great detail. Once this detailed information is available, our platform decides what machine learning technique to apply, and which specific data set to apply it to. In addition, the platform evolves dynamically as the organization and environment change.

In effect, we sought to match the situational awareness of a real human analyst so that our platform can pick and apply the right analysis technique to the right data set, dynamically, based on what it discovers about the environment. This approach means that any specific analysis we apply can change and adapt to the enterprise as it evolves. This is very different than the rule-based expert systems of the past, which were simply too brittle for today's dynamic enterprises.

Harnessing domain expertise

Domain experts can leverage data effectively because they know what questions to ask. To harness this domain expertise, we are pursuing a couple key efforts:

- **Studying users.** We have been doing extensive research into how individuals work, including in terms of how they interact with our products and other vendors' tools, to see how they go about leveraging data to complete tasks and solve problems.
- **Working with domain experts.** We collaborate extensively with a number of domain experts, both within Broadcom as well as in customer and partner organizations, to identify how they test their theories against data.

So far, we've interviewed hundreds of expert analysts about the type of analysis they do under recurring situations. We asked them what signals matter for them while testing their hypothesis about a situation, including what patterns they scan, what correlations they seek, and so on. We then captured these heuristics with machine learning robots so that our system can leverage this expertise and apply it correctly to the right scenario, technologies, and problems.

For example, a root-cause-analysis robot running on software components can determine that a Java Virtual Machine (JVM) encountered in the knowledge graph should be evaluated for memory leaks. The domain expertise would recommend the checking of certain patterns across a multitude of signals, including memory utilization, garbage collection (GC) policy, GC frequency, CPU utilization, and so on. To capture this particular know-how, our robots could apply the right classifier and model to evaluate the status of the JVM to scan the same pattern that human analysts rely on to make a determination.

Constraining problem scope and employing small, reusable components

Powerful machine learning techniques are often costly to run, and their efficacy typically increases with input curation. Consequently, to be most effective, machine learning has to be employed within the right guardrails. It is important to constrain the scope of the problem you're trying to solve.

When automobile manufacturers started to employ machine learning to deliver natural language processing-based in-car assistants, they struggled for some time. However, these days these systems are very effective, which is primarily enabled by restricting the possible words to those related to making commands to a car.¹ Automation.ai employs a similar philosophy, focusing on narrowly defined tasks and using small, reusable components to build robots. Through this approach, we can make robots fast and efficient to run, and extend their utility.

We've created one robot that has the sole responsibility of detecting incidents, for example, while another will handle incident response. Each of these robots can be managed independently, for example, they don't need to be run on the same computer or built by the same team. Further, these robots can be enhanced and optimized on independent schedules, according to evolving priorities.

Employing ontological abstractions to establish a product-agnostic, software intelligence platform

In developing the machine learning architecture for Automation.ai, we've focused on establishing capabilities around ontologies, rather than being tied to specific products. For example, in the area of application performance management (APM), we didn't focus on developing around any specific product, which can have distinct collection methods, terminology, and so on. Instead, we focused on the common, industry-accepted ontology that all APM solutions share. Consequently, our architecture can work for all APM solutions, including those from Broadcom as well as third parties.

At the same time, it's important to recognize the fact that ontologies vary across domains. For example, while an infrastructure monitoring ontology will be concerned with elements like routers and switches, a DevOps ontology will be focused on testing and production rules. That's why we've built our architecture to accommodate different ontologies, including APM, infrastructure, networks, DevOps, security, and more. Most importantly, the platform can incorporate and integrate the intelligence from these different domains.

Developing an open, flexible architecture

In the market today, many topology approaches are closed in nature, bound by specific technological approaches and linear models. By contrast, Automation.ai employs an open, source-agnostic approach.

The reality is that there's not one specific way to approach how technology is architected, developed, and managed. In fact, it's very common to see that two similar-sized organizations in the same industry will typically take very different approaches. Rather than requiring teams to adopt a specific way to think about their operations, our platform enables teams to enjoy maximum flexibility, both in the near term and the long term, without requiring them to change their approaches. This type of flexibility is critical: Today, teams can't be certain about what technology they may need to monitor in six months or a year; they do know they'll need to adapt quickly to any new technology, while continuing to support their legacy implementations.

The platform's architecture is flexible in several key ways:

- **Data source extensibility.** The platform's architecture is not bound by any specific product, but features an open data lake, algorithms, and more. Customers can easily accommodate new data sources, including those from multiple Broadcom solutions as well as solutions from a wide range of third-party vendors. As a result, the platform can easily address new use cases as implementations and business needs evolve. For example, teams can run queries against a log management platform, and import results of those queries into the Automation.ai time-series database to establish a single source of truth that algorithms can run against.
- **Architecture extensibility.** With Automation.ai, Broadcom, partners, and customers can introduce entirely new ontologies, without having to make any architectural changes.
- **Ontology extensibility.** Teams can add different properties onto existing ontologies, and so easily accommodate organization-specific information, including tribal knowledge, naming or classification approaches, and so on.
- **Robot extensibility.** The Automation.ai architecture can efficiently accommodate new robots as needed, while at the same time, enabling each robot to be employed against a unified, consistent data set. For example, a customer may have an existing robot that detects incidents or anomalies. Their partners or internal teams can subsequently develop a new robot that takes some action against an incident, such as handling intelligent routing. Teams can also enrich the incident with additional insights, or add their own data science based on proprietary internal intelligence.

In addition, by employing Automation.ai's documented, public APIs, customers can use external machine learning tools to access all the platform's correlated data, knowledge graphs, and more. As mentioned above, we are committed to providing maximum flexibility to customers, both over the near and long term, and that is especially true for customers who want to pursue their own data science initiatives. Toward that end, we're focused on delivering optimized frameworks and APIs that make it fast and efficient for customers to leverage their data science code and algorithms on our platform.

Employing an intelligent, flexible, scalable data model

In development, vendors can choose whether to create data models that are very general or very specific. While very specific models may be pragmatic in the near term, they typically require manual modifications to accommodate new technologies, data types, and so on. Over time, these specific models tend to become brittle, leaving teams unable to react quickly to changes.

In developing the data model for Automation.ai, we've employed a patented, property-graph based approach that has history awareness. This approach is structured based on entities, relationships, and their properties, and is journaled over time. With property graphs, complex relational lookups can be done instantaneously. As a result, they provide an excellent structure for doing ontological inference. By comparison, using a traditional relational database management system (RDBMS) for this model would require an impractical amount of join queries between schemas and tables, introducing an unacceptable level of performance-degrading latency.

The Automation.ai model uses properties to describe the qualitative specifics of a given entity. A customer may have monitoring agents that discover some properties of an element, such as the manufacturer, firmware version, software type, and so on. In addition, teams can leverage other data sources to further decorate the element with additional facts.

Automation.ai builds on this approach by employing a time-journaling approach within a time-aware graph database. This combination is a key design principle, and represents a unique, patented approach. Within the platform, all events are captured in a time journal. These time-stamped records represent an immutable data point, and a valuable way to establish incremental observation of an environment.

Through this approach, users and algorithmic analysts can observe structures and cross-domain interdependencies. For example, this enables a team responsible for networks to also see how changes in the broader environment may affect specific devices, such as when an application update causes traffic to start being sent to a different router, causing an unexpected workload spike. Line of business leaders can assess how a backend infrastructure upgrade affected a business metric like sales per hour. In development scenarios, this approach enables teams to track build deployments and compare them with changes in application performance, so they can more quickly and intuitively spot issues and optimization opportunities.

The Advantages of Automation.ai

By employing the differentiated approaches above, the Automation.ai platform provides a wide range of advantages to customers:

- **Maximizing data utility.** With Automation.ai, teams can combine data from all products into one model, so algorithms can be used to evaluate everything—maximizing the value of data captured and aggregated.
- **Facilitating effective analysis.** By describing the realities of an environment with a unified data model, the platform enables teams to have robots take on a range of analysis use cases. With these capabilities, Automation.ai enables teams to conduct analysis that can be very broad in scope and also very precise.
- **Powering intelligent automation.** The platform's combination of a unified data model and applied expertise enable teams to establish intelligent automation that offloads a range of tasks from human analysts.

The following sections offer more details on the unique advantages Automation.ai provides.

Contextualized, correlated intelligence

If operations teams are only looking at infrastructure monitoring, they may see there's a corrupted RAID drive. At that point, they'll have to look in other tools and work with other teams to assess the impact, determine whether it is important, and so on. This effort can demand a lot of time, effort, and expertise.

Automation.ai streamlines these efforts by providing the power of multi-layer visibility and correlation. With the platform, teams can unify intelligence from different domains. The platform can incorporate different ontologies to create a uniform playing field for data that algorithms can run on. As different data sources are incorporated, they are integrated into the property graph, building intelligence based on what is known about an entity.

The platform provides cross-domain context transfer that enables silos to be broken down. Teams can leverage what they know from one product and what they just learned from another product to react more swiftly and intelligently. For example, if a problem arises in hardware, operators can now evaluate its impact within the context of the application or end user experience. They can therefore see immediately that the hardware issue is the cause of a delay in a critical service, such as shopping cart transactions, and escalate the matter accordingly.

Current, complete, and accurate model of the run-time environment

Among disparate teams, establishing and maintaining a current picture of an IT environment has been a significant challenge. An architect can draw up an intended deployment diagram. As soon as operations deploys it, the details of that architecture can start to change. For any number of reasons, the components may not be deployed exactly as designed. Systems may immediately need to be scaled up or down. Particularly as environments continue to grow more complex and dynamic, gaining a clear, accurate picture of the operating environment only gets more difficult and time consuming.

One of the significant advantages of Automation.ai is that it enables teams to establish—and maintain—a current, complete, and accurate picture of the run-time environment. Instead of being constrained by visibility into a single architectural layer or technology, teams can view a model that's based on the multi-dimensional reality of the environment. Teams don't have to rely on old diagrams, or keep manually creating models that quickly become out of date.

Teams can leverage visibility that affords a much higher degree of accuracy. In addition, through Automation.ai, robots can be employed to apply machine learning approaches to manage clustering, correlating, and analyzing these comprehensive data sets.

They can work from a unified data model that provides comprehensive descriptions of environments and how they operate. Rather than just a static, point-in-time snapshot, the platform tracks the history of all events and relations. This enables teams to see everything happening in the data center, and to track changes and trends within the environment through a user interface and algorithmic analysts.

This multi-layered visibility means an analyst can immediately identify that a power supply failed, understand which software is running on the affected machine, and see how application performance, end users, and even business metrics are being affected.

The platform also enables teams to get invaluable insights into how the environment evolves. For example, an organization may be running applications with 20 distributed components, all interacting with each other. With Automation.ai, teams can immediately see that, after one of the components was upgraded, a bug was introduced. Further, they can then see how that bug affected other components.

Conclusion

Digital BizOps represents a modern solution that uniquely empowers business, IT, and development teams to meet their ever-intensifying demands. Digital BizOps is powered by Automation.ai, a software intelligence platform that provides an unparalleled array of capabilities. Through its differentiated architecture, Automation.ai enables teams to gain a unified view of their dynamic business and IT landscape, and leverage the automation and insights needed to realize optimized decision making, operational execution, and business outcomes.

To learn more, be sure to visit the **Automation.ai** web site.

About Broadcom

Broadcom Inc. (NASDAQ: AVGO) is a global technology leader that designs, develops and supplies a broad range of semiconductor and infrastructure software solutions. Broadcom's category-leading product portfolio serves critical markets including data center, networking, enterprise software, broadband, wireless, storage and industrial. Our solutions include data center networking and storage, enterprise and mainframe software focused on automation, monitoring and security, smartphone components, telecoms and factory automation. For more information, go to www.broadcom.com.

¹ IEEE, "Speech recognition interactive system for vehicle," <https://ieeexplore.ieee.org/document/8064929>

Broadcom, the pulse logo, Connecting everything, CA Technologies, the CA technologies logo, and Automic are among the trademarks of Broadcom and/or its affiliates in the United States, certain other countries, and/or the EU.

Copyright © 2019 Broadcom. All Rights Reserved.



The term "Broadcom" refers to Broadcom Inc. and/or its subsidiaries. For more information, please visit www.broadcom.com. Broadcom reserves the right to make changes without further notice to any products or data herein to improve reliability, function, or design. Information furnished by Broadcom is believed to be accurate and reliable. However, Broadcom does not assume any liability arising out of the application or use of this information, nor the application or use of any product or circuit described herein, neither does it convey any license under its patent rights nor the rights of others.