The Role of AI and IT in the Digitally Transformed Enterprise

Why the Evolution of Your Business Processes is Making AIOps a Necessity

By Will Cappelli / CTO and VP of Product Strategy, Moogsoft

WHITE PAPER

Moogsoft®

While agility and business fitness has increased by orders of magnitude, the manageability of and the visibility into the underlying IT systems has declined just as dramatically.

Introduction

The last 10 years have witnessed a fundamental transformation in the structure and content of business processes.

In 2007, it was estimated that at most 30% of all identifiable business process events were reflected in or realized by state changes in enterprise IT systems. By the end of 2017, that percentage had grown to close to 50 - 70 percent across the global 2000 and, among technology-intensive businesses like financial services and media and entertainment, the percentage had breached 70 percent.

In other words, over the past 10 years, business processes have become largely digital in the very literal sense that a percentage of revenue-generating business events are, in fact, IT system state changes.

The Fundamentals of the Digital Business Transformation

This fundamental transformation in the nature of business has been enabled by four critical changes in the underlying computing and communications technology base and two deep modifications of IT organizational structure and process.

The technology changes are: 1) the shift to cloud-centered infrastructures; 2) the use of increasingly modular application architectures (culminating in hyper-modularization of container-based systems); 3) growing reliance on big data platforms; and 4) the dominance of mobile, consumer-style application access interfaces.

The organizational modifications are: 1) the emergence of DevOps as a set of principles organizing the relationship between the development function and the IT operations function; and 2) the increasing tendency to treat applications and infrastructure as integrated programmable stack. Taken together these changes and modifications have allowed enterprises to create application and infrastructure stacks to more directly support and extend business processes and, importantly ensured that these digital business processes can themselves change and evolve at breathtaking speed.

Of course, all of this has come with a price. While agility and business fitness has increased by orders of magnitude, the manageability of and the visibility into the underlying IT systems has declined just as dramatically. There are three root causes of the blindness now faced by IT operations professionals.

The Top Three Challenges Facing IT

First, the mapping between the behavior of a business process in its execution and the underlying behaviors of IT system components has become abstract and indirect. In the past, while IT systems supported and enabled relatively few business processes, those relationships between the systems and those few business processes they did support was relatively straight forward. Application or business logic resided logically and physically at one identifiable location in the infrastructure topology while data resided at another well-known location. Access was uni-modal and typically tethered to a desktop of one sort or another. If an incident became manifest, a few data items about each location were usually sufficient to detect the source of the incident and remedies could be applied through local fixes to assets over which the enterprise had direct control.

Now, however, cloud services and dynamic programmable infrastructure in general means that the IT system supporting and enabling a digital business process is spread out across components whose dimensions and locations are ever-shifting. It is just about impossible to map an incident to a logical or physical location and, even if one could, the use of a public cloud compute and storage facility means that the access required to perform a remediation is difficult to achieve without the right contractual terms and incontrovertible proof.

Second, one of the key consequences of the heightened modularity of infrastructure and application stacks is the increased entropy of the data generated by the components that constitute them. In the past, when such stacks could be broken down into somewhere between five and 10 large components, it was relatively easy to infer the state of the overall stack from a relatively small number of data items about each component. Furthermore, the choreography among the different components was often so rigid that a modest amount of data concerning one or two of the components was sufficient to allow the IT operations professional to determine the source of almost any performance problem or incident.

Now, however, with the high levels of modularization first due to object orientation and, then, to containerization, the information extractable from data generated by the various components (which are, in effect, the objects or containers) is highly localized. One cannot infer the state of the entire stack just from observing what is emerging from a small number of strategically selected data sources. Indeed, the only way of seeing, let alone understanding and being in a position to modify, the infrastructure and application stack as a whole, is to be able to gather data from across a very large sub-population of the stack's constituent components. Hence the amount of data required grows by many orders of magnitude and the ability to make sense of that data without some kind of automated assistance becomes virtually impossible .

The ability to make sense of that data without some kind of automated assistance becomes virtually impossible.

Third, in the past, large chunks of most infrastructure and application stacks changed slowly, if at all, and the topologies within which the various components were rigid, their configuration remaining the same from year to year, even decade to decade. This meant that the path from incident or performance problem to cause was well understood in advance. It was often just a question of figuring which path the disturbance had travelled not a question of discovering a new path. Thanks to modularization and cloud centricity, however, the components of the stacks to be managed have become far more ephemeral and their interconnecting topologies correspondingly fluid. Cloud-based VMs can last less than a day while containers can last less than microsecond. Not only have paths from cause to effect become more and more difficult to trace, they have often vanished from the scene long before the source of the incident or performance problem can be ascertained.

The Failure of Monitoring, Event, and Incident Management Technologies

Note that this triple challenge renders most monitoring, event, and incident management solutions, at worst, obsolete or, at best, in need of major supplementation.

Monitoring solutions tend to be yoked to specific domains and much of the information they are able to provide depends on their ability to locate elements taken from a very small subset of the data generated by IT within a topology. Put another way, the ability of traditional monitoring solutions to interpret the data they capture is entirely a function of the a priori knowledge about applications, databases, networks, etc. that are is prebuilt into the solution.

Event management solutions, although less tied to specific domains, nonetheless require rules that determine which events deserve notice to be pre-written and, hence, to be useful require the environment being observed to have a static topology. Also, data sets need to be either small or highly redundant. Otherwise the event management system will be overwhelmed by alerts. Finally classical incident management systems likewise ultimately presuppose a rigid topology and a humanly tractable event stream so that they support well-defined repeatable and traceable deterministic incident and problem resolution processes.

The inability of traditional monitoring technologies, and classical event, and incident management systems to handle modern IT systems and the large, volatile and highly entropic data sets they generate, coupled with the growing significance of IT to business, is precisely what is leading many enterprises to pursue AIOps solutions to the tune of more than $2.5 billion a year. The data sets required merely to observe what is taking place, let alone find the source of issues and anticipate the future, have become so large, so volatile, and so complex that AIOps has become a prerequisite for effective IT operations.

An Overview of the AIOps Category

So what precisely is AIOps? It is a broadly accepted term used to designate the application of AI techniques and algorithms to large, dynamic volumes of data arising in IT operations contexts. The results of this application typically take the form of patterns expressed in symbolic or mathematical form that a) describe regularities in the data; b) call out anomalies or departures from the discovered regularities; and c) reveal the causal structure of the events responsible for generating the underlying data set.

Once the patterns are generated, they may be used to a support of variety of IT operations management processes, including event management, availability and performance management, incident and problem management, change and configuration management, and, ultimately, the overall automation of "IT for IT."

There are two fundamental types of AI techniques and algorithms. The first type of AI is logic-based. It makes fundamental use of rules and mechanized inference. The data sets to which logic-based AI is applied are treated like the conditions of a collection of mostly predetermined rules which are then strung together via the mechanized inference algorithms to derive conclusions about the state of the IT systems from which data was collected. It is important to note that this type of AI has been used in IT operations software and appliances almost since the beginning of commercial enterprise management.

The second type of AI is mathematics-based. Its techniques and algorithms approach the data sets directly and via a number of methods generally called "statistical machine learning" or just ML, construct equations, groupings, or graphs that are capable describing or generating data sets like the one initially worked upon. Mathematics-based AI is not new and, in fact, emerged at exactly the same time as logic-based AI in the late 1950's. It has only been in recent years, however, that it has become economically feasible to work with data sets sufficiently large and complex to make the application of mathematics-based AI worthwhile.

It has been estimated that large enterprises will have spent approximately $2.6 billion on AIOps platforms in 2018, a 25 percent growth over 2017, and all indications are that growth is likely to accelerate in 2019. Why this explosion of interest and investment? Part of this is, of course, due to general business-wide interest in the promise of AI. It is important to note that the uptick in AIOps spending started before the generic AI hype wave began to crest following Google's highly publicized acquisition of the London-based start up, DeepMind. Also, at $2.5 billion, AIOps represents, among the various industrial uses of AI, the largest concentration of actual spend by enterprises on AI-related software.

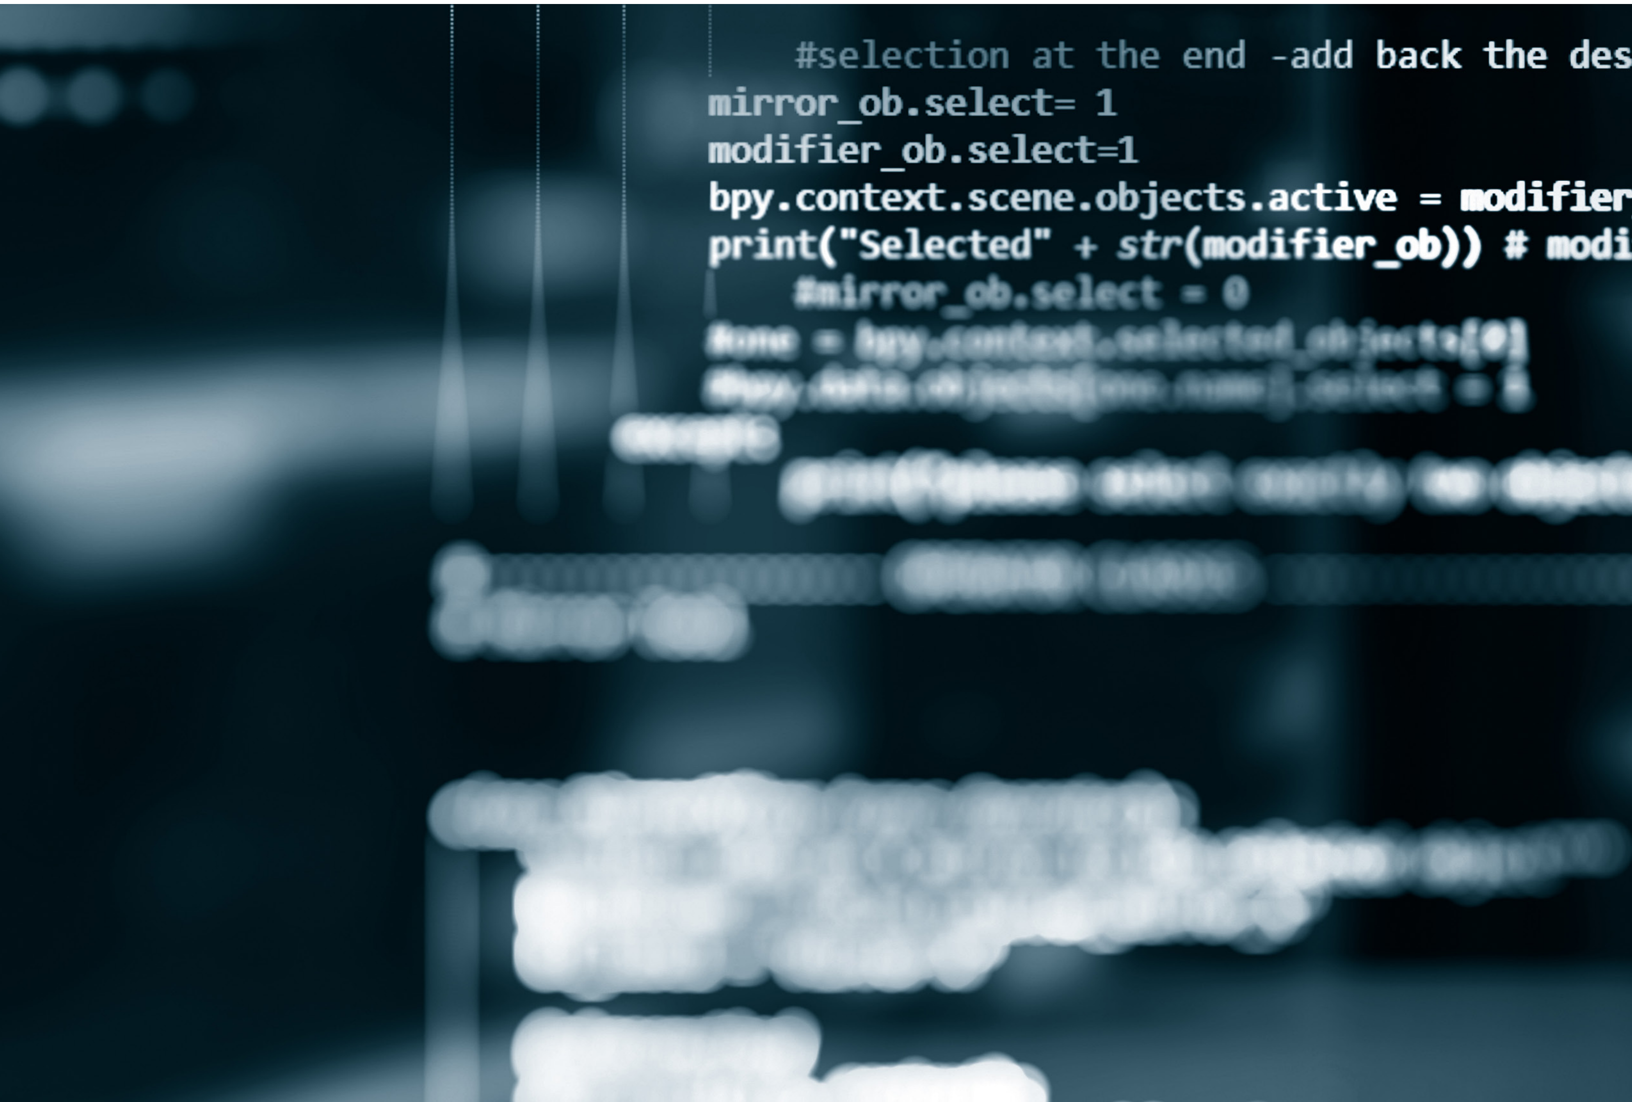

Why the Moogsoft AIOps Platform Is the Solution

Moogsoft's AIOps platform has been widely recognized by the industry as a paradigmatic AIOps technology, particularly of the mathematical sort (although not without some logic-based functional dimensions). Let us see how it copes with the three challenges to traditional monitoring described above.

From 30,000-foot perspective, the AIOps platform combines logic-based and mathematics-based AI, in a manner that reflects the way in which the human brain actually works, ingests streams of complex data being generated either directly from the components of the underlying IT systems themselves or from existing monitoring, event, and incident management capabilities. Rather than relying on a static model to contextualize and interpret the data, it applies a collection of functions directly to the data itself in hopes of discovering

the patterns that govern behavior of the data stream and ultimately to gain an actionable understanding of the environment itself which is originally responsible for generating the data.

In a bit more detail, the process can be sketched as follows: First, data streams into the platform from various sources. As such, it has many shapes and textures and needs to be translated into a form that can be worked upon by the rest of the system (think about how the nervous system converts a variety of sensory inputs into patterns of electrical pulses.). Second, the streams are passed through a filter that attempts to exclude redundant events from consideration, leaving a significantly stripped down stream for further processing. Third, modules internal to the platform try to group event signals into meaningful units or situations that can be eventually tied to what impacts business process execution.

Now, there are three distinct ways of grouping or correlating events: First, they may be grouped with regard to their proximity in Time, Second, they may be grouped with regard to their proximity in space or in the shifting Topology within which the events take place; Third and finally, events may be grouped according to the similarity of attributes described by the Text contained in the fields that make up the structure into which the original raw events have been transformed during the first step. The outcome of these three grouping functions is a collection of situations, essentially packages of events correlated according to Time, Topology, and Text.

Still, the platform is not finished. While correlations are, in and of themselves, immensely useful, giving grounds for prediction of the future and indicators of a problem's root cause. Correlation is not, however, causality. It is only when a causal connection between events can be established, only when one knows that changing a property of a given event will result, without further intervention in a change to a property of some other event, that an IT operations team can truly decide what actions need to be taken on the system to ensure that problems are addressed and business process executions optimized. Note also that it will be impossible to effectively automate IT operations processes unless the robots or run-book execution platforms are fed actual causal information.

In order transform the correlational information present in the situations into causal information, the platform now applies two further algorithms, this time to the situations: Probable Root Cause and Vertex Entropy. Each in its own way attempts to elicit the causal structure that underlies the Time, Topology, and Text based correlations captured earlier. Once the situations have been causally analyzed they are turned over the to Situation Room where they may be collaboratively addressed by the an extended IT operations community.

This layered process of ingesting and comprehending the meaning of events is, of course, not the only way of applying AI to IT operations use cases. Its fundamental architecture was, however, built upon the recognition that monitoring, event, and incident management in a digital business process setting requires the ability to discover and analyze the patterns that emerge from a complex, ever-shifting stream of data that itself reflects the ever-evolving structures that constitute modern IT environments.

Moogsoft

Moogsoft builds AIOps solutions that help IT teams work faster and smarter to provide better customer experiences. With patented algorithms analyzing billions of events daily across the world's most complex IT environments, Moogsoft's unique technology helps enterprise companies such as SAP® SuccessFactors®, Intuit®, GoDaddy™, and HCL Technologies avoid outages and increase their operational agility. To learn more, visit www.moogsoft.com.