

Artificial Intelligence (AI) for IT Operations



The Cobbler's children have no shoes

It's only occurred to me recently, but historically the IT Operations profession has had some striking similarities to the Cobbler's children and their lack of shoes. With all the architecture, design and development being focused on the customer's requirements, there was little time - or money - left for ensuring that the appropriate tools were in place to effectively and efficiently manage the solution. Of course, today, the tools and technologies are available, so why have our modern-day Cobbler's children been bare footed for so long?

Traditional IT Enterprise Management is based on monitors gathering data and establishing rules that trigger alerts when certain thresholds are breached. This type of approach is somewhat limited though, and in my experience highlights a number of challenges. The biggest of these is the number of false alarms, with only around 1% of alerts actually resulting in an investigable incident.

Not only does that equate to lots of time wasted in dealing with events that are expected, but worst still, it can be prone to errors. On occasions, real incidents were missed because they had previously been considered to be normal. They were normal - until the occasion when combined with other events and alerts that they were no longer the norm!

So what's changed?

Digital transformation is moving the world in which we live at breakneck speed, with the vast majority of this evolution being driven by IT. We have seen increased adoption of cloud technologies and automation. The types and number of devices we have to manage has increased dramatically. No longer do we just manage desktops, servers and printers. Instead, we have laptops, tablets, smart phones and IoT devices. And they are all communicating data about what we are doing and where we are doing it.

Storage systems now deal with Terabytes and Petabytes of data rather than merely the Megabytes of yesteryear. We also have big data solutions that can ingest unstructured data. The growth of AI and machine learning algorithms now allow us to quickly extract insights from data that we previously never knew existed.

The result of all of this is a significant increase in the scale and complexity of our IT systems. A decade ago, we typically managed a small datacentre of servers with an Administrator responsible for around 100 servers. Often those servers were physical, and we named them after our favourite Sci-Fi characters. In our DevOps world, those systems are known as 'Pets' - we name them, we look after them, and if they get sick, we spend lots of time fixing them.

However, in this new cloud-based world we have automated build capabilities and scale that mean we no longer look after our servers if they get sick. Instead, we shoot them! That is, we delete the virtual machine and provision a new one. This approach is known as 'Cattle', and significantly reduces the requirement for system administrators, with a single admin now being able to manage 10,000 virtual machines, if not more.

Identifying what is normal

So, if we piece together this new landscape, we can take all our cloud computing tools and consolidate them. We can gather lots of information and events from all our systems and store it in a vast data lake. Most importantly, we can utilise machine learning techniques in order to identify anomalies, and even start to predict when things may go wrong.

That's where AIOps - or AI for IT Operations comes to the fore. Gartner defines AIOps as:

"An emerging technology that utilises big data, modern machine learning and other advanced analytics technologies to both directly and indirectly enhance the IT operations function with proactive, personal and dynamic insight."

Using the data we have gathered from our monitoring solutions we can now train a machine learning model to understand what is normal. This is great for monitoring and responding to peaks in utilisation. Now our model can learn when the peak usage is and accept that as the norm. Furthermore, should I start seeing an increase in alerts, this can be identified as an anomaly.

In many AIOps tools, this unsupervised learning can be further enhanced by removing any false alarms or associations. So, following a failure event, if a server is incorrectly associated with the failure, the ITOps team can easily update the tool's rules by removing this rogue association. Through a combination of AIOps learning and human review, AIOps platforms can quickly build an accurate picture of your IT infrastructure, without the painstaking work of having to enter all associations, particularly when these change so rapidly.



Automation improves performance

Having learnt how my system performs, and in particular when a failure has occurred, I can now look to leverage any of the automation tools that I have at my disposal. So, if I have discovered an issue with a storage device, I may wish to automatically restart it. In the case of an unhealthy web application server, I may automate my cattle approach and simply kill off the unhealthy server and provision a new one from a trusted image.

In addition, AIOps can also provide additional automation across the IT Operations and IT Service Management (ITSM) functions. We are seeing significant adoption of resolver group prediction within Service Desks, with many of the top ITSM tooling vendors providing tools to support this. For resolver group prediction, machine learning is used to train a model that can provide a prediction of the most likely resolver group, that is the team responsible for fixing an incident reported to an IT Service Desk. Our experience is that a machine learning model is probably better at performing this level of prediction over its human predecessor.

Another benefit for Service Desks - in their quest to increase their first-time fix rates - is that if the model determines that it is the Service Desk that normally resolves a particular issue, we can route users directly back into the Service Desk. Or we can escalate to a manager to ensure we have provided the recommended resolution to the end user, as identified by the machine learning model. This type of resolver group prediction - when used with Virtual Agent technology - can take organisations on the road towards a zero agent Service Desk.

Proactive maintenance before failure occurs

With enough data, it is also possible to train a model to identify the symptoms of failure prior to the failure occurring. In this case, typically the model requires details of real events that can be used to identify when a failure has occurred. However, machine learning can provide additional insight into the typical pattern of activity that occurs prior to the final failure.

A good example of this is where Fujitsu has used data from its systems to build models that can now identify hard disk failures before they actually occur. With this insight, Fujitsu can arrange proactive maintenance activity, providing a replacement before the disk has actually failed, without risking any data loss. In addition, having been given several days' warning, we are also better able to arrange for the maintenance activity when an engineer is either already onsite or nearby, reducing the cost and time to perform the fix. This type of predictive maintenance activity can be adapted to many different scenarios, ranging from monitoring the performance of retail self-service terminals, to IoT devices monitoring water levels.

Managing risk – will I be home in time for tea?

The IT Change Management function has also been impacted by the digital revolution. This has led to an increase in the volume of IT Change Requests across many organisations. But changes to IT always come with risk. The challenge for the Change Manager is managing this risk, working out which changes are risky and those that are not.

Traditionally, the person responsible for identifying the level of risk is the person who knows it best, usually the one performing the change. However, this person may not be the most impartial when it comes to making the decision. They may be so familiar with the task that they simply can't see the risk; or they know that by marking the risk as low means they'll be home in time for tea!

As an IT Operations Manager, in addition to all the process steps there was a huge amount of judgement and gut feel involved when scrutinising a change request. For me, this was often related to how many times I was called about the system at 2am. What level of trust did I have in the engineer implementing the work? Did they have a reputation for fudging together the backout plan.

Finding the needle in the haystack

For many large organisations, these decisions are often made by multiple people, with varying degrees of experience and knowledge. They may be given hundreds of changes to assess each day – they just need to find that needle in the haystack. That's where AIOps comes in. Fortunately, most IT Operations teams keep good records of their IT Service Changes, recording the outcome (success or failure), and often the type and volume of incidents that were caused as a result of the change.

Our machine learning can take all this information and build a model of likely outcome, such that we can predict the level of risk of performing this task. The model would not be biased by whether or not the engineer wants to get home for tea, just what the most likely outcome is based on the information provided.


An augmented approach and shift in skills

Storage and big data platforms now allow us to gather huge amounts of data and store them for several months or even years. Organisations are using or developing DevOps and infrastructure as code-based capabilities, so we are seeing a shift in the skills required in our IT Operations teams, with scripting and coding being a critical skill. AIOps will simply make insights more accessible to our ITOps teams.

No longer will ITOps teams be satisfied with ploughing through hundreds of thousands of the same event and take no action. Instead they will be prepared to actively work to produce better tools and better monitoring of their systems. In addition, IT outages are simply not tolerated by our users, whether these are paying customers or colleagues working within the same organisation.

The IT Operations space has similarities with the challenges faced by our Security Operations Centres (or our DevSecOps team, if combined) who again are looking for the tooling that will return event logs and metrics into real security insight. In addition, by combining our AIOps platform with a ChatOps platform we will allow our IT Operations and Service Management teams to cut through the scale and complexity of the modern IT environment. When issues occur, they can use a modern communication platform that allows easy access to data and insights, whilst collaborating with our colleagues across multiple departments or even service providers.

Now I don't propose it's the end of the IT Operations Manager and Change Manager. What I'm suggesting is that with this additional insight, appropriate focus can be directed to who should review such change proposals and the level of scrutiny required. The good news is this doesn't replace the traditional event and metric gathering tools that we use today. It simply augments it. So, our modern-day Cobbler's children can have their shoes after all.

 To find out more email kevin.yeo@uk.fujitsu.com, or visit www.fujitsu.com/uk/solutions/industry/defence-national-security

About the Author



Kevin Yeo is an experienced IT professional with over 25 years' experience designing, deploying and operating large IT solutions for UK public sector organisations. He is currently working as part of Fujitsu's Defence and National Security Office of the CTO where he has been the Lead Deal Architect for a range of different secure solutions. He has a background in Service Delivery and Service Design and has a keen interest in machine learning, software development, automation, cloud and any innovative, new technology.

He was appointed a Fujitsu Distinguished Engineer in 2013, and in 2017 he lead his team to victory in Fujitsu's first Hackathon "Hello K5" by creating an Alexa app that integrated a bespoke IoT device powered by Fujitsu cloud computing offerings. Kevin is part of the leadership team in Fujitsu's AI Special Interest Group which works to ensure best practise and knowledge within the field of AI and Machine Learning is shared across the company and Fujitsu's customers.

Why Fujitsu?

For over 50 years we have innovated with the MOD, Government Departments and intelligence communities, co-creating new technologies and capabilities. As a result, Fujitsu has around 4,000 security cleared staff and the experience to deliver and manage both generic industry offerings and those tailored to specialist needs at OFFICIAL, SECRET and ABOVE SECRET classifications.

Enabling Your Information Advantage

In today's complex, digital operational environment, never before has information been such a key asset in securing operational advantage. Fujitsu's vision is to provide customers with the means to translate complex data into useful information upon

which to base critical decisions and actions. Transforming this ever-increasing pool of data into meaningful, useful information through analytics, automation and genuine Artificial Intelligence is critical to achieving this goal.

Fujitsu is fully committed to working closely with our customers, and through the use of co-creation will seek to enhance capability both through the acceleration of existing processes, and also through the delivery of truly new capabilities and ways of working. Our approach is based upon maximising both existing investment and best-in-class innovation, delivering the full spectrum of capabilities needed to enable your information advantage.



Contact

Telephone: +44 (0)870 242 7998
Email: askfujitsu@uk.fujitsu.com
Ref: 3972
uk.fujitsu.com

Unclassified. © 2020 FUJITSU. Fujitsu, the Fujitsu logo, are trademarks or registered trademarks of Fujitsu Limited in Japan and other countries. Other company, product and service names may be trademarks or registered trademarks of their respective owners. Technical data subject to modification and delivery subject to availability. Any liability that the data and illustrations are complete, actual or correct is excluded. Designations may be trademarks and/or copyrights of the respective manufacturer, the use of which by third parties for their own purposes may infringe the rights of such owner. ID: 6827-003-03/2020.